

---

---

**Information technology —  
Programming languages — Ada**

*Technologies de l'information — Langages de programmation — Ada*

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 8652:2023



STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 8652:2023



**COPYRIGHT PROTECTED DOCUMENT**

© ISO/IEC 2023

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office  
CP 401 • Ch. de Blandonnet 8  
CH-1214 Vernier, Geneva  
Phone: +41 22 749 01 11  
Email: [copyright@iso.org](mailto:copyright@iso.org)  
Website: [www.iso.org](http://www.iso.org)

Published in Switzerland

# Contents

|                                                        |      |
|--------------------------------------------------------|------|
| Foreword .....                                         | xiii |
| Introduction .....                                     | xv   |
| 1 Scope .....                                          | 1    |
| 2 Normative references .....                           | 1    |
| 3 Terms and definitions .....                          | 2    |
| 3.1 Types, objects, and their properties .....         | 2    |
| 3.2 Subprograms and their properties .....             | 7    |
| 3.3 Other syntactic constructs .....                   | 7    |
| 3.4 Runtime actions .....                              | 10   |
| 3.5 Exceptional situations .....                       | 10   |
| 4 General .....                                        | 11   |
| 4.1 Structure .....                                    | 11   |
| 4.2 Conformity of an Implementation .....              | 12   |
| 4.3 Method of Description and Syntax Notation .....    | 14   |
| 4.4 Classification of Errors .....                     | 15   |
| 5 Lexical Elements .....                               | 16   |
| 5.1 Character Set .....                                | 16   |
| 5.2 Lexical Elements, Separators, and Delimiters ..... | 18   |
| 5.3 Identifiers .....                                  | 19   |
| 5.4 Numeric Literals .....                             | 20   |
| 5.4.1 Decimal Literals .....                           | 20   |
| 5.4.2 Based Literals .....                             | 21   |
| 5.5 Character Literals .....                           | 21   |
| 5.6 String Literals .....                              | 21   |
| 5.7 Comments .....                                     | 22   |
| 5.8 Pragmas .....                                      | 22   |
| 5.9 Reserved Words .....                               | 25   |
| 6 Declarations and Types .....                         | 25   |
| 6.1 Declarations .....                                 | 25   |
| 6.2 Types and Subtypes .....                           | 27   |
| 6.2.1 Type Declarations .....                          | 28   |
| 6.2.2 Subtype Declarations .....                       | 29   |
| 6.2.3 Classification of Operations .....               | 30   |
| 6.2.4 Subtype Predicates .....                         | 31   |
| 6.3 Objects and Named Numbers .....                    | 34   |
| 6.3.1 Object Declarations .....                        | 36   |
| 6.3.2 Number Declarations .....                        | 39   |
| 6.4 Derived Types and Classes .....                    | 39   |
| 6.4.1 Derivation Classes .....                         | 42   |
| 6.5 Scalar Types .....                                 | 44   |
| 6.5.1 Enumeration Types .....                          | 47   |
| 6.5.2 Character Types .....                            | 48   |
| 6.5.3 Boolean Types .....                              | 49   |
| 6.5.4 Integer Types .....                              | 49   |
| 6.5.5 Operations of Discrete Types .....               | 52   |
| 6.5.6 Real Types .....                                 | 53   |
| 6.5.7 Floating Point Types .....                       | 54   |

|        |                                                         |     |
|--------|---------------------------------------------------------|-----|
| 6.5.8  | Operations of Floating Point Types .....                | 55  |
| 6.5.9  | Fixed Point Types.....                                  | 56  |
| 6.5.10 | Operations of Fixed Point Types.....                    | 58  |
| 6.6    | Array Types.....                                        | 59  |
| 6.6.1  | Index Constraints and Discrete Ranges .....             | 61  |
| 6.6.2  | Operations of Array Types.....                          | 62  |
| 6.6.3  | String Types .....                                      | 63  |
| 6.7    | Discriminants.....                                      | 63  |
| 6.7.1  | Discriminant Constraints.....                           | 66  |
| 6.7.2  | Operations of Discriminated Types.....                  | 67  |
| 6.8    | Record Types.....                                       | 67  |
| 6.8.1  | Variant Parts and Discrete Choices .....                | 70  |
| 6.9    | Tagged Types and Type Extensions .....                  | 72  |
| 6.9.1  | Type Extensions.....                                    | 75  |
| 6.9.2  | Dispatching Operations of Tagged Types.....             | 76  |
| 6.9.3  | Abstract Types and Subprograms.....                     | 79  |
| 6.9.4  | Interface Types .....                                   | 80  |
| 6.10   | Access Types .....                                      | 82  |
| 6.10.1 | Incomplete Type Declarations .....                      | 85  |
| 6.10.2 | Operations of Access Types .....                        | 87  |
| 6.11   | Declarative Parts.....                                  | 93  |
| 6.11.1 | Completions of Declarations .....                       | 93  |
| 7      | Names and Expressions.....                              | 94  |
| 7.1    | Names.....                                              | 94  |
| 7.1.1  | Indexed Components .....                                | 95  |
| 7.1.2  | Slices .....                                            | 96  |
| 7.1.3  | Selected Components.....                                | 97  |
| 7.1.4  | Attributes .....                                        | 99  |
| 7.1.5  | User-Defined References.....                            | 100 |
| 7.1.6  | User-Defined Indexing .....                             | 101 |
| 7.2    | Literals .....                                          | 103 |
| 7.2.1  | User-Defined Literals.....                              | 104 |
| 7.3    | Aggregates .....                                        | 106 |
| 7.3.1  | Record Aggregates.....                                  | 106 |
| 7.3.2  | Extension Aggregates .....                              | 108 |
| 7.3.3  | Array Aggregates.....                                   | 110 |
| 7.3.4  | Delta Aggregates.....                                   | 114 |
| 7.3.5  | Container Aggregates .....                              | 115 |
| 7.4    | Expressions .....                                       | 121 |
| 7.5    | Operators and Expression Evaluation.....                | 123 |
| 7.5.1  | Logical Operators and Short-circuit Control Forms ..... | 124 |
| 7.5.2  | Relational Operators and Membership Tests.....          | 125 |
| 7.5.3  | Binary Adding Operators .....                           | 128 |
| 7.5.4  | Unary Adding Operators.....                             | 129 |
| 7.5.5  | Multiplying Operators .....                             | 129 |
| 7.5.6  | Highest Precedence Operators.....                       | 132 |
| 7.5.7  | Conditional Expressions.....                            | 132 |
| 7.5.8  | Quantified Expressions .....                            | 134 |
| 7.5.9  | Declare Expressions.....                                | 135 |
| 7.5.10 | Reduction Expressions.....                              | 136 |
| 7.6    | Type Conversions.....                                   | 139 |
| 7.7    | Qualified Expressions.....                              | 144 |
| 7.8    | Allocators .....                                        | 144 |
| 7.9    | Static Expressions and Static Subtypes.....             | 146 |
| 7.9.1  | Statically Matching Constraints and Subtypes .....      | 150 |



|        |                                                                |     |
|--------|----------------------------------------------------------------|-----|
| 7.10   | Image Attributes .....                                         | 151 |
| 8      | Statements .....                                               | 154 |
| 8.1    | Simple and Compound Statements - Sequences of Statements ..... | 155 |
| 8.2    | Assignment Statements .....                                    | 156 |
| 8.2.1  | Target Name Symbols .....                                      | 158 |
| 8.3    | If Statements .....                                            | 158 |
| 8.4    | Case Statements .....                                          | 159 |
| 8.5    | Loop Statements .....                                          | 160 |
| 8.5.1  | User-Defined Iterator Types .....                              | 164 |
| 8.5.2  | Generalized Loop Iteration .....                               | 166 |
| 8.5.3  | Procedural Iterators .....                                     | 169 |
| 8.6    | Block Statements .....                                         | 172 |
| 8.6.1  | Parallel Block Statements .....                                | 172 |
| 8.7    | Exit Statements .....                                          | 174 |
| 8.8    | Goto Statements .....                                          | 175 |
| 9      | Subprograms .....                                              | 175 |
| 9.1    | Subprogram Declarations .....                                  | 176 |
| 9.1.1  | Preconditions and Postconditions .....                         | 178 |
| 9.1.2  | The Global and Global'Class Aspects .....                      | 183 |
| 9.2    | Formal Parameter Modes .....                                   | 186 |
| 9.3    | Subprogram Bodies .....                                        | 187 |
| 9.3.1  | Conformance Rules .....                                        | 188 |
| 9.3.2  | Inline Expansion of Subprograms .....                          | 190 |
| 9.4    | Subprogram Calls .....                                         | 190 |
| 9.4.1  | Parameter Associations .....                                   | 192 |
| 9.5    | Return Statements .....                                        | 194 |
| 9.5.1  | Nonreturning Subprograms .....                                 | 197 |
| 9.6    | Overloading of Operators .....                                 | 198 |
| 9.7    | Null Procedures .....                                          | 199 |
| 9.8    | Expression Functions .....                                     | 199 |
| 10     | Packages .....                                                 | 201 |
| 10.1   | Package Specifications and Declarations .....                  | 201 |
| 10.2   | Package Bodies .....                                           | 202 |
| 10.3   | Private Types and Private Extensions .....                     | 203 |
| 10.3.1 | Private Operations .....                                       | 205 |
| 10.3.2 | Type Invariants .....                                          | 208 |
| 10.3.3 | Default Initial Conditions .....                               | 211 |
| 10.3.4 | Stable Properties of a Type .....                              | 211 |
| 10.4   | Deferred Constants .....                                       | 213 |
| 10.5   | Limited Types .....                                            | 214 |
| 10.6   | Assignment and Finalization .....                              | 216 |
| 10.6.1 | Completion and Finalization .....                              | 219 |
| 11     | Visibility Rules .....                                         | 221 |
| 11.1   | Declarative Region .....                                       | 221 |
| 11.2   | Scope of Declarations .....                                    | 222 |
| 11.3   | Visibility .....                                               | 223 |
| 11.3.1 | Overriding Indicators .....                                    | 226 |
| 11.4   | Use Clauses .....                                              | 227 |
| 11.5   | Renaming Declarations .....                                    | 228 |
| 11.5.1 | Object Renaming Declarations .....                             | 228 |
| 11.5.2 | Exception Renaming Declarations .....                          | 230 |
| 11.5.3 | Package Renaming Declarations .....                            | 230 |
| 11.5.4 | Subprogram Renaming Declarations .....                         | 230 |

|         |                                                             |     |
|---------|-------------------------------------------------------------|-----|
| 11.5.5  | Generic Renaming Declarations .....                         | 232 |
| 11.6    | The Context of Overload Resolution .....                    | 233 |
| 12      | Tasks and Synchronization .....                             | 235 |
| 12.1    | Task Units and Task Objects .....                           | 236 |
| 12.2    | Task Execution - Task Activation .....                      | 238 |
| 12.3    | Task Dependence - Termination of Tasks .....                | 239 |
| 12.4    | Protected Units and Protected Objects .....                 | 241 |
| 12.5    | Intertask Communication .....                               | 244 |
| 12.5.1  | Protected Subprograms and Protected Actions .....           | 248 |
| 12.5.2  | Entries and Accept Statements .....                         | 250 |
| 12.5.3  | Entry Calls .....                                           | 253 |
| 12.5.4  | Requeue Statements .....                                    | 255 |
| 12.6    | Delay Statements, Duration, and Time .....                  | 257 |
| 12.6.1  | Formatting, Time Zones, and other operations for Time ..... | 260 |
| 12.7    | Select Statements .....                                     | 266 |
| 12.7.1  | Selective Accept .....                                      | 266 |
| 12.7.2  | Timed Entry Calls .....                                     | 268 |
| 12.7.3  | Conditional Entry Calls .....                               | 269 |
| 12.7.4  | Asynchronous Transfer of Control .....                      | 269 |
| 12.8    | Abort of a Task - Abort of a Sequence of Statements .....   | 270 |
| 12.9    | Task and Entry Attributes .....                             | 272 |
| 12.10   | Shared Variables .....                                      | 272 |
| 12.10.1 | Conflict Check Policies .....                               | 274 |
| 12.11   | Example of Tasking and Synchronization .....                | 276 |
| 13      | Program Structure and Compilation Issues .....              | 277 |
| 13.1    | Separate Compilation .....                                  | 277 |
| 13.1.1  | Compilation Units - Library Units .....                     | 278 |
| 13.1.2  | Context Clauses - With Clauses .....                        | 281 |
| 13.1.3  | Subunits of Compilation Units .....                         | 283 |
| 13.1.4  | The Compilation Process .....                               | 284 |
| 13.1.5  | Pragmas and Program Units .....                             | 285 |
| 13.1.6  | Environment-Level Visibility Rules .....                    | 286 |
| 13.2    | Program Execution .....                                     | 286 |
| 13.2.1  | Elaboration Control .....                                   | 288 |
| 14      | Exceptions .....                                            | 291 |
| 14.1    | Exception Declarations .....                                | 292 |
| 14.2    | Exception Handlers .....                                    | 292 |
| 14.3    | Raise Statements and Raise Expressions .....                | 293 |
| 14.4    | Exception Handling .....                                    | 294 |
| 14.4.1  | The Package Exceptions .....                                | 295 |
| 14.4.2  | Pragmas Assert and Assertion_Policy .....                   | 297 |
| 14.4.3  | Example of Exception Handling .....                         | 299 |
| 14.5    | Suppressing Checks .....                                    | 300 |
| 14.6    | Exceptions and Optimization .....                           | 303 |
| 15      | Generic Units .....                                         | 304 |
| 15.1    | Generic Declarations .....                                  | 304 |
| 15.2    | Generic Bodies .....                                        | 306 |
| 15.3    | Generic Instantiation .....                                 | 307 |
| 15.4    | Formal Objects .....                                        | 309 |
| 15.5    | Formal Types .....                                          | 310 |
| 15.5.1  | Formal Private and Derived Types .....                      | 312 |
| 15.5.2  | Formal Scalar Types .....                                   | 314 |
| 15.5.3  | Formal Array Types .....                                    | 314 |

|                     |                                                                      |     |
|---------------------|----------------------------------------------------------------------|-----|
| 15.5.4              | Formal Access Types .....                                            | 315 |
| 15.5.5              | Formal Interface Types .....                                         | 316 |
| 15.6                | Formal Subprograms .....                                             | 316 |
| 15.7                | Formal Packages .....                                                | 319 |
| 15.8                | Example of a Generic Package .....                                   | 321 |
| 16                  | Representation Issues.....                                           | 322 |
| 16.1                | Operational and Representation Aspects.....                          | 322 |
| 16.1.1              | Aspect Specifications.....                                           | 326 |
| 16.2                | Packed Types .....                                                   | 329 |
| 16.3                | Operational and Representation Attributes .....                      | 330 |
| 16.4                | Enumeration Representation Clauses .....                             | 337 |
| 16.5                | Record Layout.....                                                   | 338 |
| 16.5.1              | Record Representation Clauses .....                                  | 338 |
| 16.5.2              | Storage Place Attributes.....                                        | 341 |
| 16.5.3              | Bit Ordering.....                                                    | 341 |
| 16.6                | Change of Representation .....                                       | 342 |
| 16.7                | The Package System .....                                             | 342 |
| 16.7.1              | The Package System.Storage_Elements.....                             | 344 |
| 16.7.2              | The Package System.Address_To_Access_Conversions.....                | 345 |
| 16.8                | Machine Code Insertions .....                                        | 346 |
| 16.9                | Unchecked Type Conversions.....                                      | 347 |
| 16.9.1              | Data Validity .....                                                  | 348 |
| 16.9.2              | The Valid Attribute .....                                            | 349 |
| 16.10               | Unchecked Access Value Creation.....                                 | 349 |
| 16.11               | Storage Management .....                                             | 350 |
| 16.11.1             | Storage Allocation Attributes.....                                   | 353 |
| 16.11.2             | Unchecked Storage Deallocation.....                                  | 354 |
| 16.11.3             | Default Storage Pools.....                                           | 355 |
| 16.11.4             | Storage Subpools .....                                               | 356 |
| 16.11.5             | Subpool Reclamation .....                                            | 359 |
| 16.11.6             | Storage Subpool Example.....                                         | 359 |
| 16.12               | Pragma Restrictions and Pragma Profile .....                         | 362 |
| 16.12.1             | Language-Defined Restrictions and Profiles.....                      | 363 |
| 16.13               | Streams .....                                                        | 365 |
| 16.13.1             | The Streams Subsystem .....                                          | 365 |
| 16.13.2             | Stream-Oriented Attributes .....                                     | 368 |
| 16.14               | Freezing Rules .....                                                 | 373 |
| Annex A (normative) | Predefined Language Environment.....                                 | 376 |
| A.1                 | The Package Standard .....                                           | 379 |
| A.2                 | The Package Ada.....                                                 | 384 |
| A.3                 | Character Handling.....                                              | 384 |
| A.3.1               | The Packages Characters, Wide_Characters, and Wide_Wide_Characters.. | 384 |
| A.3.2               | The Package Characters.Handling.....                                 | 385 |
| A.3.3               | The Package Characters.Latin_1.....                                  | 387 |
| A.3.4               | The Package Characters.Conversions .....                             | 395 |
| A.3.5               | The Package Wide_Characters.Handling.....                            | 397 |
| A.3.6               | The Package Wide_Wide_Characters.Handling .....                      | 400 |
| A.4                 | String Handling .....                                                | 400 |
| A.4.1               | The Package Strings.....                                             | 400 |
| A.4.2               | The Package Strings.Maps .....                                       | 401 |
| A.4.3               | Fixed-Length String Handling .....                                   | 404 |
| A.4.4               | Bounded-Length String Handling .....                                 | 412 |
| A.4.5               | Unbounded-Length String Handling.....                                | 419 |
| A.4.6               | String-Handling Sets and Mappings.....                               | 425 |

|         |                                                                                   |     |
|---------|-----------------------------------------------------------------------------------|-----|
| A.4.7   | Wide_String Handling .....                                                        | 425 |
| A.4.8   | Wide_Wide_String Handling .....                                                   | 428 |
| A.4.9   | String Hashing .....                                                              | 430 |
| A.4.10  | String Comparison .....                                                           | 431 |
| A.4.11  | String Encoding .....                                                             | 432 |
| A.4.12  | Universal Text Buffers .....                                                      | 437 |
| A.5     | The Numerics Packages .....                                                       | 439 |
| A.5.1   | Elementary Functions .....                                                        | 440 |
| A.5.2   | Random Number Generation .....                                                    | 443 |
| A.5.3   | Attributes of Floating Point Types .....                                          | 448 |
| A.5.4   | Attributes of Fixed Point Types .....                                             | 452 |
| A.5.5   | Big Numbers .....                                                                 | 453 |
| A.5.6   | Big Integers .....                                                                | 453 |
| A.5.7   | Big Reals .....                                                                   | 455 |
| A.6     | Input-Output .....                                                                | 457 |
| A.7     | External Files and File Objects .....                                             | 457 |
| A.8     | Sequential and Direct Files .....                                                 | 458 |
| A.8.1   | The Generic Package Sequential_IO .....                                           | 459 |
| A.8.2   | File Management .....                                                             | 460 |
| A.8.3   | Sequential Input-Output Operations .....                                          | 462 |
| A.8.4   | The Generic Package Direct_IO .....                                               | 463 |
| A.8.5   | Direct Input-Output Operations .....                                              | 464 |
| A.9     | The Generic Package Storage_IO .....                                              | 465 |
| A.10    | Text Input-Output .....                                                           | 466 |
| A.10.1  | The Package Text_IO .....                                                         | 467 |
| A.10.2  | Text File Management .....                                                        | 475 |
| A.10.3  | Default Input, Output, and Error Files .....                                      | 475 |
| A.10.4  | Specification of Line and Page Lengths .....                                      | 476 |
| A.10.5  | Operations on Columns, Lines, and Pages .....                                     | 477 |
| A.10.6  | Get and Put Procedures .....                                                      | 480 |
| A.10.7  | Input-Output of Characters and Strings .....                                      | 481 |
| A.10.8  | Input-Output for Integer Types .....                                              | 483 |
| A.10.9  | Input-Output for Real Types .....                                                 | 485 |
| A.10.10 | Input-Output for Enumeration Types .....                                          | 487 |
| A.10.11 | Input-Output for Bounded Strings .....                                            | 488 |
| A.10.12 | Input-Output for Unbounded Strings .....                                          | 489 |
| A.11    | Wide Text Input-Output and Wide Wide Text Input-Output .....                      | 491 |
| A.12    | Stream Input-Output .....                                                         | 491 |
| A.12.1  | The Package Streams.Stream_IO .....                                               | 491 |
| A.12.2  | The Package Text_IO.Text_Streams .....                                            | 494 |
| A.12.3  | The Package Wide_Text_IO.Text_Streams .....                                       | 495 |
| A.12.4  | The Package Wide_Wide_Text_IO.Text_Streams .....                                  | 495 |
| A.13    | Exceptions in Input-Output .....                                                  | 495 |
| A.14    | File Sharing .....                                                                | 497 |
| A.15    | The Package Command_Line .....                                                    | 497 |
| A.15.1  | The Packages Wide_Command_Line and Wide_Wide_Command_Line .....                   | 498 |
| A.16    | The Package Directories .....                                                     | 498 |
| A.16.1  | The Package Directories.Hierarchical_File_Names .....                             | 506 |
| A.16.2  | The Packages Wide_Directories and Wide_Wide_Directories .....                     | 507 |
| A.17    | The Package Environment_Variables .....                                           | 507 |
| A.17.1  | The Packages Wide_Environment_Variables and Wide_Wide_Environment_Variables ..... | 509 |
| A.18    | Containers .....                                                                  | 509 |
| A.18.1  | The Package Containers .....                                                      | 511 |
| A.18.2  | The Generic Package Containers.Vectors .....                                      | 511 |

|                              |                                                                     |            |
|------------------------------|---------------------------------------------------------------------|------------|
| A.18.3                       | The Generic Package Containers.Doubly_Linked_Lists .....            | 545        |
| A.18.4                       | Maps.....                                                           | 567        |
| A.18.5                       | The Generic Package Containers.Hashed_Maps .....                    | 577        |
| A.18.6                       | The Generic Package Containers.Ordered_Maps .....                   | 586        |
| A.18.7                       | Sets.....                                                           | 597        |
| A.18.8                       | The Generic Package Containers.Hashed_Sets.....                     | 608        |
| A.18.9                       | The Generic Package Containers.Ordered_Sets.....                    | 618        |
| A.18.10                      | The Generic Package Containers.Multiway_Trees .....                 | 631        |
| A.18.11                      | The Generic Package Containers.Indefinite_Vectors .....             | 667        |
| A.18.12                      | The Generic Package Containers.Indefinite_Doubly_Linked_Lists ..... | 668        |
| A.18.13                      | The Generic Package Containers.Indefinite_Hashed_Maps .....         | 669        |
| A.18.14                      | The Generic Package Containers.Indefinite_Ordered_Maps .....        | 669        |
| A.18.15                      | The Generic Package Containers.Indefinite_Hashed_Sets.....          | 670        |
| A.18.16                      | The Generic Package Containers.Indefinite_Ordered_Sets.....         | 670        |
| A.18.17                      | The Generic Package Containers.Indefinite_Multiway_Trees.....       | 670        |
| A.18.18                      | The Generic Package Containers.Indefinite_Holders.....              | 671        |
| A.18.19                      | The Generic Package Containers.Bounded_Vectors.....                 | 676        |
| A.18.20                      | The Generic Package Containers.Bounded_Doubly_Linked_Lists.....     | 677        |
| A.18.21                      | The Generic Package Containers.Bounded_Hashed_Maps.....             | 679        |
| A.18.22                      | The Generic Package Containers.Bounded_Ordered_Maps.....            | 680        |
| A.18.23                      | The Generic Package Containers.Bounded_Hashed_Sets .....            | 682        |
| A.18.24                      | The Generic Package Containers.Bounded_Ordered_Sets .....           | 684        |
| A.18.25                      | The Generic Package Containers.Bounded_Multiway_Trees .....         | 685        |
| A.18.26                      | Array Sorting.....                                                  | 687        |
| A.18.27                      | The Generic Package Containers.Synchronized_Queue_Interfaces .....  | 689        |
| A.18.28                      | The Generic Package Containers.Unbounded_Synchronized_Queues.....   | 690        |
| A.18.29                      | The Generic Package Containers.Bounded_Synchronized_Queues .....    | 690        |
| A.18.30                      | The Generic Package Containers.Unbounded_Priority_Queues .....      | 691        |
| A.18.31                      | The Generic Package Containers.Bounded_Priority_Queues.....         | 693        |
| A.18.32                      | The Generic Package Containers.Bounded_Indefinite_Holders.....      | 694        |
| A.18.33                      | Example of Container Use .....                                      | 694        |
| A.19                         | The Package Locales .....                                           | 697        |
| <b>Annex B (normative)</b>   | <b>Interface to Other Languages .....</b>                           | <b>698</b> |
| B.1                          | Interfacing Aspects .....                                           | 698        |
| B.2                          | The Package Interfaces .....                                        | 701        |
| B.3                          | Interfacing with C and C++ .....                                    | 702        |
| B.3.1                        | The Package Interfaces.C.Strings.....                               | 709        |
| B.3.2                        | The Generic Package Interfaces.C.Pointers .....                     | 712        |
| B.3.3                        | Unchecked Union Types.....                                          | 714        |
| B.4                          | Interfacing with COBOL.....                                         | 716        |
| B.5                          | Interfacing with Fortran .....                                      | 722        |
| <b>Annex C (informative)</b> | <b>Systems Programming.....</b>                                     | <b>725</b> |
| C.1                          | Access to Machine Operations .....                                  | 725        |
| C.2                          | Required Representation Support.....                                | 726        |
| C.3                          | Interrupt Support.....                                              | 726        |
| C.3.1                        | Protected Procedure Handlers .....                                  | 728        |
| C.3.2                        | The Package Interrupts .....                                        | 730        |
| C.4                          | Preelaboration Requirements.....                                    | 732        |
| C.5                          | Aspect Discard_Names.....                                           | 732        |
| C.6                          | Shared Variable Control.....                                        | 733        |
| C.6.1                        | The Package System.Atomic_Operations .....                          | 736        |
| C.6.2                        | The Package System.Atomic_Operations.Exchange .....                 | 736        |
| C.6.3                        | The Package System.Atomic_Operations.Test_and_Set .....             | 737        |
| C.6.4                        | The Package System.Atomic_Operations.Integer_Arithmetic.....        | 738        |



|                       |                                                               |     |
|-----------------------|---------------------------------------------------------------|-----|
| C.6.5                 | The Package System.Atomic_Operations.Modular_Arithmetic ..... | 739 |
| C.7                   | Task Information .....                                        | 740 |
| C.7.1                 | The Package Task_Identification.....                          | 740 |
| C.7.2                 | The Package Task_Attributes.....                              | 742 |
| C.7.3                 | The Package Task_Termination.....                             | 744 |
| Annex D (informative) | Real-Time Systems.....                                        | 746 |
| D.1                   | Task Priorities.....                                          | 746 |
| D.2                   | Priority Scheduling .....                                     | 748 |
| D.2.1                 | The Task Dispatching Model.....                               | 748 |
| D.2.2                 | Task Dispatching Pragmas.....                                 | 750 |
| D.2.3                 | Preemptive Dispatching .....                                  | 751 |
| D.2.4                 | Non-Preemptive Dispatching.....                               | 752 |
| D.2.5                 | Round Robin Dispatching .....                                 | 753 |
| D.2.6                 | Earliest Deadline First Dispatching.....                      | 754 |
| D.3                   | Priority Ceiling Locking.....                                 | 757 |
| D.4                   | Entry Queuing Policies.....                                   | 759 |
| D.4.1                 | Admission Policies.....                                       | 761 |
| D.5                   | Dynamic Priorities .....                                      | 761 |
| D.5.1                 | Dynamic Priorities for Tasks .....                            | 761 |
| D.5.2                 | Dynamic Priorities for Protected Objects .....                | 762 |
| D.6                   | Preemptive Abort .....                                        | 763 |
| D.7                   | Tasking Restrictions .....                                    | 764 |
| D.8                   | Monotonic Time.....                                           | 767 |
| D.9                   | Delay Accuracy.....                                           | 770 |
| D.10                  | Synchronous Task Control .....                                | 771 |
| D.10.1                | Synchronous Barriers.....                                     | 772 |
| D.11                  | Asynchronous Task Control.....                                | 773 |
| D.12                  | Other Optimizations and Determinism Rules.....                | 774 |
| D.13                  | The Ravenscar and Jorvik Profiles.....                        | 775 |
| D.14                  | Execution Time.....                                           | 777 |
| D.14.1                | Execution Time Timers.....                                    | 779 |
| D.14.2                | Group Execution Time Budgets.....                             | 780 |
| D.14.3                | Execution Time of Interrupt Handlers .....                    | 783 |
| D.15                  | Timing Events .....                                           | 783 |
| D.16                  | Multiprocessor Implementation.....                            | 785 |
| D.16.1                | Multiprocessor Dispatching Domains .....                      | 786 |
| Annex E (informative) | Distributed Systems.....                                      | 789 |
| E.1                   | Partitions .....                                              | 789 |
| E.2                   | Categorization of Library Units .....                         | 790 |
| E.2.1                 | Shared Passive Library Units.....                             | 791 |
| E.2.2                 | Remote Types Library Units.....                               | 792 |
| E.2.3                 | Remote Call Interface Library Units.....                      | 793 |
| E.3                   | Consistency of a Distributed System.....                      | 794 |
| E.4                   | Remote Subprogram Calls .....                                 | 795 |
| E.4.1                 | Asynchronous Remote Calls.....                                | 796 |
| E.4.2                 | Example of Use of a Remote Access-to-Class-Wide Type .....    | 797 |
| E.5                   | Partition Communication Subsystem.....                        | 798 |
| Annex F (informative) | Information Systems.....                                      | 801 |
| F.1                   | Machine_Radix Attribute Definition Clause .....               | 801 |
| F.2                   | The Package Decimal.....                                      | 801 |
| F.3                   | Edited Output for Decimal Types .....                         | 802 |
| F.3.1                 | Picture String Formation.....                                 | 804 |
| F.3.2                 | Edited Output Generation .....                                | 807 |
| F.3.3                 | The Package Text_IO.Editing.....                              | 811 |

|                              |                                                             |            |
|------------------------------|-------------------------------------------------------------|------------|
| F.3.4                        | The Package Wide_Text_IO.Editing .....                      | 814        |
| F.3.5                        | The Package Wide_Wide_Text_IO.Editing .....                 | 814        |
| <b>Annex G (informative)</b> | <b>Numerics .....</b>                                       | <b>815</b> |
| G.1                          | Complex Arithmetic .....                                    | 815        |
| G.1.1                        | Complex Types .....                                         | 815        |
| G.1.2                        | Complex Elementary Functions .....                          | 819        |
| G.1.3                        | Complex Input-Output .....                                  | 823        |
| G.1.4                        | The Package Wide_Text_IO.Complex_IO .....                   | 825        |
| G.1.5                        | The Package Wide_Wide_Text_IO.Complex_IO .....              | 825        |
| G.2                          | Numeric Performance Requirements .....                      | 825        |
| G.2.1                        | Model of Floating Point Arithmetic .....                    | 826        |
| G.2.2                        | Model-Oriented Attributes of Floating Point Types .....     | 827        |
| G.2.3                        | Model of Fixed Point Arithmetic .....                       | 828        |
| G.2.4                        | Accuracy Requirements for the Elementary Functions .....    | 829        |
| G.2.5                        | Performance Requirements for Random Number Generation ..... | 831        |
| G.2.6                        | Accuracy Requirements for Complex Arithmetic .....          | 831        |
| G.3                          | Vector and Matrix Manipulation .....                        | 833        |
| G.3.1                        | Real Vectors and Matrices .....                             | 834        |
| G.3.2                        | Complex Vectors and Matrices .....                          | 838        |
| <b>Annex H (informative)</b> | <b>High Integrity Systems .....</b>                         | <b>849</b> |
| H.1                          | Pragma Normalize_Scalars .....                              | 849        |
| H.2                          | Documentation of Implementation Decisions .....             | 850        |
| H.3                          | Reviewable Object Code .....                                | 850        |
| H.3.1                        | Pragma Reviewable .....                                     | 850        |
| H.3.2                        | Pragma Inspection_Point .....                               | 851        |
| H.4                          | High Integrity Restrictions .....                           | 852        |
| H.4.1                        | Aspect No_Controlled_Parts .....                            | 855        |
| H.5                          | Pragma Detect_Blocking .....                                | 855        |
| H.6                          | Pragma Partition_Elaboration_Policy .....                   | 855        |
| H.7                          | Extensions to Global and Global'Class Aspects .....         | 856        |
| H.7.1                        | The Use_Formal and Dispatching Aspects .....                | 857        |
| <b>Annex I (normative)</b>   | <b>Obsolescent Features .....</b>                           | <b>860</b> |
| I.1                          | Renamings of Library Units .....                            | 860        |
| I.2                          | Allowed Replacements of Characters .....                    | 860        |
| I.3                          | Reduced Accuracy Subtypes .....                             | 861        |
| I.4                          | The Constrained Attribute .....                             | 861        |
| I.5                          | ASCII .....                                                 | 862        |
| I.6                          | Numeric_Error .....                                         | 862        |
| I.7                          | At Clauses .....                                            | 862        |
| I.7.1                        | Interrupt Entries .....                                     | 863        |
| I.8                          | Mod Clauses .....                                           | 864        |
| I.9                          | The Storage_Size Attribute .....                            | 864        |
| I.10                         | Specific Suppression of Checks .....                        | 865        |
| I.11                         | The Class Attribute of Untagged Incomplete Types .....      | 865        |
| I.12                         | Pragma Interface .....                                      | 865        |
| I.13                         | Dependence Restriction Identifiers .....                    | 865        |
| I.14                         | Character and Wide_Character Conversion Functions .....     | 866        |
| I.15                         | Aspect-related Pragas .....                                 | 866        |
| I.15.1                       | Pragma Inline .....                                         | 867        |
| I.15.2                       | Pragma No_Return .....                                      | 867        |
| I.15.3                       | Pragma Pack .....                                           | 867        |
| I.15.4                       | Pragma Storage_Size .....                                   | 868        |
| I.15.5                       | Interfacing Pragas .....                                    | 868        |
| I.15.6                       | Pragma Unchecked_Union .....                                | 869        |

|                       |                                                    |     |
|-----------------------|----------------------------------------------------|-----|
| I.15.7                | Pragmas Interrupt_Handler and Attach_Handler ..... | 869 |
| I.15.8                | Shared Variable Pragmas .....                      | 870 |
| I.15.9                | Pragma CPU .....                                   | 870 |
| I.15.10               | Pragma Dispatching_Domain .....                    | 871 |
| I.15.11               | Pragmas Priority and Interrupt_Priority .....      | 871 |
| I.15.12               | Pragma Relative_Deadline .....                     | 872 |
| I.15.13               | Pragma Asynchronous .....                          | 872 |
| I.15.14               | Elaboration Control Pragmas .....                  | 872 |
| I.15.15               | Distribution Pragmas .....                         | 873 |
| Annex J (informative) | Language-Defined Aspects and Attributes .....      | 875 |
| J.1                   | Language-Defined Aspects .....                     | 875 |
| J.2                   | Language-Defined Attributes .....                  | 879 |
| Annex K (informative) | Language-Defined Pragmas .....                     | 897 |
| Annex L (informative) | Summary of Documentation Requirements .....        | 899 |
| L.1                   | Specific Documentation Requirements .....          | 899 |
| L.2                   | Implementation-Defined Characteristics .....       | 901 |
| L.3                   | Implementation Advice .....                        | 907 |
| Annex M (informative) | Syntax Summary .....                               | 916 |
| M.1                   | Syntax Rules .....                                 | 916 |
| M.2                   | Syntax Cross Reference .....                       | 941 |
| Annex N (informative) | Language-Defined Entities .....                    | 967 |
| N.1                   | Language-Defined Packages .....                    | 967 |
| N.2                   | Language-Defined Types and Subtypes .....          | 970 |
| N.3                   | Language-Defined Subprograms .....                 | 974 |
| N.4                   | Language-Defined Exceptions .....                  | 987 |
| N.5                   | Language-Defined Objects .....                     | 988 |
| Bibliography          | .....                                              | 993 |
| Index                 | .....                                              | 994 |



## Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see [www.iso.org/directives](http://www.iso.org/directives) or [www.iec.ch/members\\_experts/refdocs](http://www.iec.ch/members_experts/refdocs)).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see [www.iso.org/patents](http://www.iso.org/patents)) or the IEC list of patent declarations received (see <https://patents.iec.ch>).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see [www.iso.org/iso/foreword.html](http://www.iso.org/iso/foreword.html). In the IEC, see [www.iec.ch/understanding-standards](http://www.iec.ch/understanding-standards).

This document was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 22, *Programming languages, their environments and system software interfaces*.

This fourth edition cancels and replaces the third edition (ISO/IEC 8652:2012), which has been technically revised. It also incorporates the Technical Corrigendum ISO/IEC 8652:2012/Cor.1:2016.

The main changes are as follows:

- improved support for parallel execution is provided via the introduction of parallel loops, parallel blocks, parallel container iteration, and parallel reduction;
- more precise specification of subprogram interfaces is supported via the new aspects Global, GlobalClass, and Nonblocking. The Global aspects, in particular, help to determine whether two constructs can safely execute in parallel;
- Pre and Post aspects can now be specified for access-to-subprogram types and for generic formal subprograms; a postcondition for the default initialization of a type can be specified using the new Default\_Initial\_Condition aspect;
- the behavior of many predefined container operations is now more precisely specified by using pre- and postcondition specifications instead of English descriptions; a restricted (“stable”) view for most containers is introduced to support more efficient iteration;
- more flexible uses of static expressions are supported via the introduction of static expression functions along with fewer restrictions on static strings;
- the Image attribute is supported for nonscalar types, and a user-specifiable attribute Put\_Image is provided, which determines the value of the Image attribute for a user-defined type;

- the use of numeric and string literals is generalized to allow their use with other categories of types, via the new aspects `Integer_Literal`, `Real_Literal`, and `String_Literal`;
- array and record aggregates are made more flexible: index parameters are allowed in an array aggregate to define the components as a function of their array index; discriminants can be defined more flexibly within an aggregate for a variant record type;
- new types of aggregates are provided: delta aggregates to allow the construction of a new object by incremental updates to an existing object; container aggregates to allow construction of an object of a container type by directly specifying its elements;
- a shorthand is provided, using the token '@', to refer to the target of an assignment statement in the expression defining its new value;
- declare expressions are provided that permit the definition and use of local constants or renamings, to allow a large expression to be simplified by defining common parts as named entities;
- support for lightweight iteration is added via the introduction of procedural iterators;
- support for the map-reduce programming strategy is added via the introduction of reduction expressions;
- for constructs that use iterators of any sort, a filter can be specified that restricts the elements produced by the iteration to those that satisfy the condition of the filter;
- predefined packages supporting arbitrary-precision integer and real arithmetic are provided;
- the Jorvik profile is introduced to support hard real-time applications that want to go beyond the restrictions of the Ravenscar profile.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at [www.iso.org/members.html](http://www.iso.org/members.html) and [www.iec.ch/national-committees](http://www.iec.ch/national-committees).

# Introduction

## Design Goals

Ada was originally designed with three overriding concerns: program reliability and maintenance, programming as a human activity, and efficiency. The 1995 revision to the language was designed to provide greater flexibility and extensibility, additional control over storage management and synchronization, and standardized packages oriented toward supporting important application areas, while at the same time retaining the original emphasis on reliability, maintainability, and efficiency. Subsequent editions, including this fourth edition, have provided further flexibility and added more standardized packages within the framework provided by the 1995 revision.

The need for languages that promote reliability and simplify maintenance is well established. Hence emphasis was placed on program readability over ease of writing. For example, the rules of the language require that program variables be explicitly declared and that their type be specified. Since the type of a variable is invariant, compilers can ensure that operations on variables are compatible with the properties intended for objects of the type. Furthermore, error-prone notations have been avoided, and the syntax of the language avoids the use of encoded forms in favor of more English-like constructs. Finally, the language offers support for separate compilation of program units in a way that facilitates program development and maintenance, and which provides the same degree of checking between units as within a unit.

Concern for the human programmer was also stressed during the design. Above all, an attempt was made to keep to a relatively small number of underlying concepts integrated in a consistent and systematic way while continuing to avoid the pitfalls of excessive involution. The design especially aims to provide language constructs that correspond intuitively to the normal expectations of users.

Like many other human activities, the development of programs is becoming ever more decentralized and distributed. Consequently, the ability to assemble a program from independently produced software components continues to be a central idea in the design. The concepts of packages, of private types, and of generic units are directly related to this idea, which has ramifications in many other aspects of the language. An allied concern is the maintenance of programs to match changing requirements; type extension and the hierarchical library enable a program to be modified while minimizing disturbance to existing tested and trusted components.

No language can avoid the problem of efficiency. Languages that require over-elaborate compilers, or that lead to the inefficient use of storage or execution time, force these inefficiencies on all machines and on all programs. Every construct of the language was examined in the light of present implementation techniques. Any proposed construct whose implementation was unclear or that required excessive machine resources was rejected. Parallel constructs were introduced to simplify making safe and efficient use of modern multicore architectures.

## Language Summary

An Ada program is composed of one or more program units. Program units can be subprograms (which define executable algorithms), packages (which define collections of entities), task units (which define concurrent computations), protected units (which define operations for the coordinated sharing of data between tasks), or generic units (which define parameterized forms of packages and subprograms). Each program unit normally consists of two parts: a specification, containing the information that is visible to other units, and a body, containing the implementation details, which are not visible to other units. Most program units can be compiled separately.

This distinction of the specification and body, and the ability to compile units separately, allows a program to be designed, written, and tested as a set of largely independent software components.

An Ada program will normally make use of a library of program units of general utility. The language provides means whereby individual organizations can construct their own libraries. All libraries are structured in a hierarchical manner; this enables the logical decomposition of a subsystem into

individual components. The text of a separately compiled program unit names the library units it requires.

### *Program Units*

A subprogram is the basic unit for expressing an algorithm. There are two kinds of subprograms: procedures and functions. A procedure is the means of invoking a series of actions. For example, it can read data, update variables, or produce some output. It can have parameters, to provide a controlled means of passing information between the procedure and the point of call. A function is the means of invoking the computation of a value. It is similar to a procedure, but in addition will return a result.

A package is the basic unit for defining a collection of logically related entities. For example, a package can be used to define a set of type declarations and associated operations. Portions of a package can be hidden from the user, thus allowing access only to the logical properties expressed by the package specification.

Subprogram and package units can be compiled separately and arranged in hierarchies of parent and child units giving fine control over visibility of the logical properties and their detailed implementation.

A task unit is the basic unit for defining a task whose sequence of actions can be executed concurrently with those of other tasks. Such tasks can be implemented on multicomputers, multiprocessors, or with interleaved execution on a single processor. A task unit can define either a single executing task or a task type permitting the creation of any number of similar tasks.

A protected unit is the basic unit for defining protected operations for the coordinated use of data shared between tasks. Simple mutual exclusion is provided automatically, and more elaborate sharing protocols can be defined. A protected operation can either be a subprogram or an entry. A protected entry specifies a Boolean expression (an entry barrier) that blocks the execution of the body until it evaluates to True. A protected unit can define a single protected object or a protected type permitting the creation of several similar objects.

### *Declarations and Statements*

The body of a program unit generally contains two parts: a declarative part, which defines the logical entities to be used in the program unit, and a sequence of statements, which defines the execution of the program unit.

The declarative part associates names with declared entities. For example, a name can denote a type, a constant, a variable, or an exception. A declarative part also introduces the names and parameters of other nested subprograms, packages, task units, protected units, and generic units to be used in the program unit.

The sequence of statements describes a sequence of actions to be performed. The statements are executed in succession (unless a transfer of control causes execution to continue from another place).

An assignment statement changes the value of a variable. A procedure call invokes execution of a procedure after associating any actual parameters provided at the call with the corresponding formal parameters.

Case statements and if statements allow the selection of an enclosed sequence of statements based on the value of an expression or on the value of a condition.

The loop statement provides the basic iterative mechanism in the language. A loop statement specifies a sequence of statements that are executed repeatedly as directed by an iteration scheme, or until an exit statement is encountered.

A block statement comprises a sequence of statements preceded by the declaration of local entities used by the statements.

Certain statements are associated with concurrent execution. A delay statement delays the execution of a task for a specified duration or until a specified time. An entry call statement is written as a procedure call statement; it requests an operation on a task or on a protected object, blocking the caller until the operation can be performed. A called task can accept an entry call by executing a corresponding accept

statement, which specifies the actions then to be performed as part of the rendezvous with the calling task. An entry call on a protected object is processed when the corresponding entry barrier evaluates to true, whereupon the body of the entry is executed. The requeue statement permits the provision of a service as a number of related activities with preference control. One form of the select statement allows a selective wait for one of several alternative rendezvous. Other forms of the select statement allow conditional or timed entry calls and the asynchronous transfer of control in response to some triggering event. Various parallel constructs, including parallel loops and parallel blocks, support the initiation of multiple logical threads of control designed to execute in parallel when multiple processors are available.

Execution of a program unit can encounter error situations in which normal program execution cannot continue. For example, an arithmetic computation can exceed the maximum allowed value of a number, or an attempt can be made to access an array component by using an incorrect index value. To deal with such error situations, the statements of a program unit can be textually followed by exception handlers that specify the actions to be taken when the error situation arises. Exceptions can be raised explicitly by a raise statement.

### *Data Types*

Every object in the language has a type, which characterizes a set of values and a set of applicable operations. The main categories of types are elementary types (comprising enumeration, numeric, and access types) and composite types (including array and record types).

An enumeration type defines an ordered set of distinct enumeration literals, for example a list of states or an alphabet of characters. The enumeration types Boolean, Character, Wide\_Character, and Wide\_Wide\_Character are predefined.

Numeric types provide a means of performing exact or approximate numerical computations. Exact computations use integer types, which denote sets of consecutive integers. Approximate computations use either fixed point types, with absolute bounds on the error, or floating point types, with relative bounds on the error. The numeric types Integer, Float, and Duration are predefined.

Composite types allow definitions of structured objects with related components. The composite types in the language include arrays and records. An array is an object with indexed components of the same type. A record is an object with named components of possibly different types. Task and protected types are also forms of composite types. The array types String, Wide\_String, and Wide\_Wide\_String are predefined.

Record, task, and protected types can have special components called discriminants which parameterize the type. Variant record structures that depend on the values of discriminants can be defined within a record type.

Access types allow the construction of linked data structures. A value of an access type represents a reference to an object declared as aliased or to an object created by the evaluation of an allocator. Several variables of an access type can designate the same object, and components of one object can designate the same or other objects. Both the elements in such linked data structures and their relation to other elements can be altered during program execution. Access types also permit references to subprograms to be stored, passed as parameters, and ultimately dereferenced as part of an indirect call.

Private types permit restricted views of a type. A private type can be defined in a package so that only the logically necessary properties are made visible to the users of the type. The full structural details that are externally irrelevant are then only available within the package and any child units.

From any type a new type can be defined by derivation. A type, together with its derivatives (both direct and indirect) form a derivation class. Class-wide operations can be defined that accept as a parameter an operand of any type in a derivation class. For record and private types, the derivatives can be extensions of the parent type. Types that support these object-oriented capabilities of class-wide operations and type extension are tagged, so that the specific type of an operand within a derivation class can be identified at run time. When an operation of a tagged type is applied to an operand whose specific type is not known until run time, implicit dispatching is performed based on the tag of the operand.

Interface types provide abstract models from which other interfaces and types can be composed and derived. This provides a reliable form of multiple inheritance. Interface types can also be implemented by task types and protected types thereby enabling concurrent programming and inheritance to be merged.

The concept of a type is further refined by the concept of a subtype, whereby a user can constrain the set of allowed values of a type. Subtypes can be used to define subranges of scalar types, arrays with a limited set of index values, and records and private types with particular discriminant values.

#### *Other Facilities*

Aspect clauses can be used to specify the mapping between types and features of an underlying machine. For example, the user can specify that objects of a given type must be represented with a given number of bits, or that the components of a record must be represented using a given storage layout. Other features allow the controlled use of low level, nonportable, or implementation-dependent aspects, including the direct insertion of machine code.

Aspect clauses can also be used to specify more abstract properties of program entities, such as the pre- and postconditions of a subprogram, or the invariant for a private type. Additional aspects are specifiable to allow user-defined types to use constructs of the language, such as literals, aggregates, or indexing, normally reserved for particular language-defined categories of types, such as numeric types, record types, or array types.

The predefined environment of the language provides for input-output and other capabilities by means of standard library packages. Input-output is supported for values of user-defined as well as of predefined types. Standard means of representing values in display form are also provided.

The predefined standard library packages provide facilities such as string manipulation, containers of various kinds (vectors, lists, maps, etc.), mathematical functions, random number generation, and access to the execution environment.

The specialized annexes define further predefined library packages and facilities with emphasis on areas such as real-time scheduling, interrupt handling, distributed systems, numerical computation, and high-integrity systems.

Finally, the language provides a powerful means of parameterization of program units, called generic program units. The generic parameters can be types and subprograms (as well as objects and packages) and so allow general algorithms and data structures to be defined that are applicable to all types of a given class.



# Information technology — Programming Languages — Ada

## 1 Scope

This document specifies the form and meaning of programs written in Ada. Its purpose is to promote the portability of Ada programs to a variety of computing systems.

This document specifies:

- The form of a program written in Ada;
- The effect of translating and executing such a program;
- The manner in which program units can be combined to form Ada programs;
- The language-defined library units that a conforming implementation is required to supply;
- The permissible variations in conformance to the rules of this document, and the manner in which they are to be documented;
- Those violations of the requirements of this document that a conforming implementation is required to detect, and the effect of attempting to translate or execute a program containing such violations;
- Those violations of the requirements of this document that a conforming implementation is not required to detect.

This document does not specify:

- The means whereby a program written in Ada is transformed into object code executable by a processor;
- The means whereby translation or execution of programs is invoked and the executing units are controlled;
- The size or speed of the object code, or the relative execution speed of different language constructs;
- The form or contents of any listings produced by implementations; in particular, the form or contents of error or warning messages;
- The effect of unspecified execution;
- The size of a program or program unit that will exceed the capacity of a particular conforming implementation.

## 2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO 639-3:2007, *Codes for the representation of names of languages — Part 3: Alpha-3 code for comprehensive coverage of languages*

ISO 3166-1:2020, *Codes for the representation of names of countries and their subdivisions — Part 1: Country codes*

ISO/IEC 10646:2020, *Information technology — Universal coded character set (UCS)*



### 3 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

ISO and IEC maintain terminology databases for use in standardization at the following addresses:

- ISO Online browsing platform: available at <https://www.iso.org/obp>
- IEC Electropedia: available at <https://www.electropedia.org/>

#### 3.1 Types, objects, and their properties

##### 3.1.1

##### **abstract type**

tagged type intended for use as an ancestor of other types, but which is not allowed to have objects of its own

##### 3.1.2

##### **access type**

type that has values that designate aliased objects

Note 1 to entry: Access types correspond to “pointer types” or “reference types” in some other languages.

##### 3.1.3

##### **accessibility level**

representation of the lifetime of an entity in terms of the level of dynamic nesting within which the entity is known to exist

##### 3.1.4

##### **aliased view**

view of an object that can be designated by an access value

Note 1 to entry: Objects allocated by allocators are aliased. Objects can also be explicitly declared as aliased with the reserved word aliased. The Access attribute can be used to create an access value designating an aliased object.

##### 3.1.5

##### **ancestor of a type**

type itself or, in the case of a type derived from other types, its parent type or one of its progenitor types or one of their ancestors

Note 1 to entry: Ancestor and descendant are inverse relationships.

##### 3.1.6

##### **array type**

composite type whose components are all of the same type

##### 3.1.7

##### **aspect**

specifiable property of an entity

Note 1 to entry: An aspect can be specified by an `aspect_specification` on the declaration of the entity. Some aspects can be queried via attributes.

##### 3.1.8

##### **attribute**

characteristic or property of an entity that can be queried, and in some cases specified

##### 3.1.9

##### **category of types**

set of types with one or more common properties, such as primitive operations

Note 1 to entry: A category of types that is closed under derivation is also known as a class.

**3.1.10****character type**

enumeration type whose values include characters

**3.1.11****class of types**

set of types that is closed under derivation, which means that if a given type is in the class, then all types derived from that type are also in the class

Note 1 to entry: The set of types of a class share common properties, such as their primitive operations.

**3.1.12****composite type**

type with components, such as an array or record

**3.1.13****controlled type**

type that supports user-defined assignment and finalization

Note 1 to entry: Objects are always finalized before being destroyed.

**3.1.14****default initial condition**

property that holds for every default-initialized object of a given type

**3.1.15****derived type**

type defined in terms of a parent type and zero or more progenitor types given in a derived type definition

Note 1 to entry: A derived type inherits properties such as components and primitive operations from its parent and progenitors.

Note 2 to entry: A type together with the types derived from it (directly or indirectly) form a derivation class.

**3.1.16****descendant of a type**

type itself or a type derived (directly or indirectly) from it

Note 1 to entry: Descendant and ancestor are inverse relationships.

**3.1.17****discrete type**

type that is either an integer type or an enumeration type

**3.1.18****discriminant**

parameter for a composite type, which can control, for example, the bounds of a component that is an array

Note 1 to entry: A discriminant for a task type can be used to pass data to a task of the type upon its creation.

**3.1.19****elementary type**

type that does not have components

**3.1.20****enumeration type**

type defined by an enumeration of its values, which can be denoted by identifiers or character literals

**3.1.21****full type**

type that defines a full view

### 3.1.22

#### **full view**

view of a type that reveals all of its properties

Note 1 to entry: There can be other views of the type that reveal fewer properties.

### 3.1.23

#### **incomplete type**

type that defines an incomplete view

Note 1 to entry: Incomplete types can be used for defining recursive data structures.

### 3.1.24

#### **incomplete view**

view of a type that reveals minimal properties

Note 1 to entry: The remaining properties are defined by the full view given elsewhere.

### 3.1.25

#### **indexable container type**

type that has user-defined behavior for indexing, via the `Constant_Indexing` or `Variable_Indexing` aspects

### 3.1.26

#### **integer type**

type that represents signed or modular integers

Note 1 to entry: A signed integer type has a base range that includes both positive and negative numbers, and has operations that can raise an exception when the result is outside the base range. A modular type has a base range whose lower bound is zero, and has operations with “wraparound” semantics. Modular types subsume what are called “unsigned types” in some other languages.

### 3.1.27

#### **interface type**

abstract tagged type that has no components or concrete operations except possibly null procedures

Note 1 to entry: Interface types are used for composing other interfaces and tagged types and thereby provide multiple inheritance. Only an interface type can be used as a progenitor of another type.

### 3.1.28

#### **invariant**

assertion that is expected to be `True` for all objects of a given private type when viewed from outside the defining package

### 3.1.29

#### **iterable container type**

type that has user-defined behavior for iteration, via the `Default_Iterator` and `Iterator_Element` aspects

### 3.1.30

#### **limited type**

type for which copying (such as in an `assignment_statement`) is not allowed

Note 1 to entry: All types are either limited types or nonlimited types.

### 3.1.31

#### **needed component**

component of a record type or record extension that is required to have its value specified within a given aggregate

### 3.1.32

#### **nominal subtype**

subtype specified when a view of an object is defined

**3.1.33****nonlimited type**

type for which copying is allowed

**3.1.34****object**

entity that contains a value, and is either a constant or a variable

Note 1 to entry: An object is created by an `object_declaration` or by an `allocator`. A formal parameter is (a view of) an object. A subcomponent of an object is an object.

**3.1.35****operational aspect**

aspect that indicates a logical property of an entity, such as the precondition of a subprogram, or the procedure used to write a given type of object to a stream

**3.1.36****parent of a derived type**

first ancestor type given in the definition of the derived type

Note 1 to entry: The parent can be almost any kind of type, including an interface type.

**3.1.37****partial view**

view of a type that reveals only some of its properties

Note 1 to entry: The remaining properties are defined by the full view given elsewhere.

**3.1.38****primitive operations of a type**

operations (such as subprograms) declared together with the type declarations

Note 1 to entry: Primitive operations are inherited by other types in the same derivation class of types.

**3.1.39****private extension**

type that extends another type, with the additional properties hidden from its clients

**3.1.40****private type**

type that defines a partial view

Note 1 to entry: Private types can be used for defining abstractions that hide unnecessary details from their clients.

**3.1.41****progenitor**

type given in the interface list, if any, of an interface, task, protected, or derived type definition

Note 1 to entry: A progenitor is always an interface type.

**3.1.42****protected type**

composite type whose components are accessible only through one of its protected operations, which synchronize concurrent access by multiple tasks

**3.1.43****real type**

type that has values that are approximations of the real numbers

Note 1 to entry: Floating point and fixed point types are real types.

**3.1.44****record extension**

type that extends another type optionally with additional components

**3.1.45**

**record type**

composite type consisting of zero or more named components, possibly of different types

**3.1.46**

**reference type**

type that has user-defined behavior for “.all”, defined by the Implicit\_Dereference aspect

**3.1.47**

**representation aspect**

aspect that indicates how an entity is mapped onto the underlying hardware, for example the size or alignment of an object

**3.1.48**

**scalar type**

either a discrete type or a real type

**3.1.49**

**stable property**

characteristic associated with objects of a given type that is preserved by many of the primitive operations of the type

**3.1.50**

**storage pool object**

object associated with one or more access types from which the storage for objects created by allocators of the access type(s) is obtained

Note 1 to entry: Some storage pools can be partitioned into subpools in order to support finer-grained storage management.

**3.1.51**

**stream**

sequence of elements that can be used, along with the stream-oriented attributes, to support marshalling and unmarshalling of values of most types

**3.1.52**

**subtype**

type together with optional constraints, null exclusions, and predicates, which constrain the values of the type to the subset that satisfies the implied conditions

**3.1.53**

**synchronized**

safely operated on by multiple tasks concurrently

Note 1 to entry: Synchronized is used to qualify entities, as in a synchronized interface.

**3.1.54**

**tagged type**

type whose objects each have a run-time type tag, which indicates the specific type for which the object was originally created

Note 1 to entry: Tagged types can be extended with additional components.

**3.1.55**

**task type**

composite type used to represent active entities which execute concurrently and that can communicate via queued task entries

Note 1 to entry: The top-level task of a partition is called the environment task.

**3.1.56****type**

defining characteristic of each object and expression of the language, with an associated set of values, and a set of primitive operations that implement the fundamental aspects of its semantics

Note 1 to entry: Types are grouped into categories. Most language-defined categories of types are also classes of types.

**3.1.57****view of an entity**

representation of an entity that reveals some or all of the properties of the entity

Note 1 to entry: A single entity can have multiple views.

**3.2 Subprograms and their properties****3.2.1****function**

form of subprogram that returns a result and can be called as part of an expression

**3.2.2****overriding operation**

operation that replaces an inherited primitive operation

Note 1 to entry: Operations can be marked explicitly as overriding or not overriding.

**3.2.3****postcondition**

assertion that is expected to be True when a given subprogram returns normally

**3.2.4****precondition**

assertion that is expected to be True when a given subprogram is called

**3.2.5****procedure**

form of subprogram that does not return a result and can only be invoked by a statement

**3.2.6****subprogram**

unit of a program that can be brought into execution in various contexts, with the invocation being a subprogram call that can parameterize the effect of the subprogram through the passing of operands

Note 1 to entry: There are two forms of subprograms: functions, which return values, and procedures, which do not.

**3.3 Other syntactic constructs****3.3.1****aggregate**

construct used to define a value of a composite type by specifying the values of the components of the type

**3.3.2****compilation unit**

program unit that is separately compiled

Note 1 to entry: A `compilation_unit` contains either the declaration, the body, or a renaming of a program unit.

**3.3.3****construct**

piece of text (explicit or implicit) that is an instance of a syntactic category defined under Syntax



### 3.3.4

#### **container**

structured object that represents a collection of elements all of the same (potentially class-wide) type, such as a vector or a tree

Note 1 to entry: Several predefined container types are provided by the children of package Ada.Containers (see A.18.1).

### 3.3.5

#### **container aggregate**

construct used to define a value of a type that represents a collection of elements, by explicitly specifying the elements in the collection

### 3.3.6

#### **core language**

clause or annex in which are defined language constructs or capabilities that are provided by all conforming implementations

Note 1 to entry: A construct is said to be part of the core language if it is defined in a core language clause or annex.

### 3.3.7

#### **declaration**

language construct that associates a name with (a view of) an entity

Note 1 to entry: A declaration can appear explicitly in the program text (an explicit declaration), or can be supposed to occur at a given place in the text as a consequence of the semantics of another construct (an implicit declaration).

### 3.3.8

#### **generic instance**

nongeneric unit created by the instantiation of a generic unit

### 3.3.9

#### **generic unit**

template for a (nongeneric) program unit

Note 1 to entry: The template can be parameterized by objects, types, subprograms, and packages.

Note 2 to entry: Generic units can be used to perform the role that macros sometimes play in other languages.

### 3.3.10

#### **iterator**

construct that is used to loop over the elements of an array or container

Note 1 to entry: Iterators can be user defined, and can perform arbitrary computations to access elements from a container.

### 3.3.11

#### **iterator filter**

construct that is used to restrict the elements produced by an iteration to those for which a boolean condition evaluates to True

### 3.3.12

#### **library unit**

separately compiled program unit, which is a package, a subprogram, or a generic unit

Note 1 to entry: Library units can have other (logically nested) library units as children, and can have other program units physically nested within them. A root library unit, together with its children and grandchildren and so on, form a subsystem.

**3.3.13****master construct**

one of certain executable constructs for which there can be objects or tasks whose lifetime ends when the construct completes

Note 1 to entry: Execution of a master construct is a master, with which objects and tasks are associated for the purposes of waiting and finalization.

**3.3.14****needed compilation unit**

compilation unit that is necessary to produce an executable partition, because some entity declared or defined within the unit is used elsewhere in the partition

**3.3.15****package**

program unit that defines the interface to a group of logically related entities, along with their implementation

Note 1 to entry: Typically, a package contains the declaration of a type (often a private type or private extension) along with the declarations of primitive subprograms of the type, which can be called from outside the package, while their inner workings remain hidden from outside users.

**3.3.16****parallel construct**

executable construct that defines multiple activities of a single task that can proceed in parallel, via the execution of multiple logical threads of control

**3.3.17****partition**

part of a program, which consists of a set of interdependent library units

Note 1 to entry: Each partition can run in a separate address space, possibly on a separate computer. A program can contain just one partition, or it can be distributed across multiple partitions, which can execute concurrently.

**3.3.18****pragma**

compiler directive to provide control over and above that provided by the other syntactic constructs of the language

Note 1 to entry: There are language-defined pragmas that give instructions for optimization, listing control, etc. An implementation can support additional (implementation-defined) pragmas.

**3.3.19****program**

set of partitions, each of which can execute in a separate address space, possibly on a separate computer

**3.3.20****program unit**

language construct that is a package, a task unit, a protected unit, a protected entry, a generic unit, or an explicitly declared subprogram other than an enumeration literal

Note 1 to entry: Certain kinds of program units can be separately compiled. Alternatively, they can appear physically nested within other program units.

**3.3.21****reduction expression**

expression that defines how to map or transform a collection of values into a new set of values, and then summarize the values by applying an operation to reduce the set to a single value

**3.3.22****renaming**

declaration that does not define a new entity, but instead defines a new view of an existing entity

### 3.3.23

#### **specialized needs annex**

annex in which are defined language constructs or capabilities that are not necessarily provided by all conforming implementations

### 3.3.24

#### **subunit**

body of a program unit that can be compiled separately from its enclosing program unit

## 3.4 Runtime actions

### 3.4.1

#### **assertion**

boolean expression that is expected to be True at run time at certain specified places

Note 1 to entry: Certain pragmas and aspects define various kinds of assertions.

### 3.4.2

#### **elaboration**

process by which a declaration achieves its run-time effect

Note 1 to entry: Elaboration is one of the forms of execution.

### 3.4.3

#### **evaluation**

process by which an expression achieves its run-time effect

Note 1 to entry: Evaluation is one of the forms of execution.

### 3.4.4

#### **execution**

process by which a construct achieves its run-time effect

Note 1 to entry: Execution of a declaration is also called elaboration. Execution of an expression is also called evaluation.

### 3.4.5

#### **logical thread of control**

activity within the execution of a program that can proceed in parallel with other activities of the same task, or of separate tasks

### 3.4.6

#### **master**

execution of a master construct

Note 1 to entry: Each object and task is associated with a master. When a master is left, associated tasks are awaited and associated objects are finalized.

## 3.5 Exceptional situations

### 3.5.1

#### **check**

test made during execution to determine whether a language rule has been violated

### 3.5.2

#### **exception**

kind of exceptional situation

### 3.5.3

#### **exception occurrence**

run-time occurrence of an exceptional situation

**3.5.4****handle an exception**

perform some actions in response to the arising of an exception

**3.5.5****raise an exception**

abandon normal program execution so as to draw attention to the fact that the corresponding situation has arisen

**3.5.6****suppress a check**

assert that the check cannot fail, and request that the compiler optimize by disabling the check

Note 1 to entry: The compiler is not required to honor this request. Suppressing checks that can fail can cause a program to behave in arbitrary ways.

**4 General****4.1 Structure**

This document contains sixteen clauses, fourteen annexes, a bibliography, and an index.

All implementations shall conform to the core of the Ada language, which consists of:

- Clauses 1 to 16
- Annex A, “Predefined Language Environment”
- Annex B, “Interface to Other Languages”
- Annex I, “Obsolescent Features”

In addition, an implementation may conform separately to one or more Specialized Needs Annexes, which define features that are needed by certain application areas:

- Annex C, “Systems Programming”
- Annex D, “Real-Time Systems”
- Annex E, “Distributed Systems”
- Annex F, “Information Systems”
- Annex G, “Numerics”
- Annex H, “High Integrity Systems”

The following Annexes provide additional information:

- Annex J, “Language-Defined Aspects and Attributes”
- Annex K, “Language-Defined Pragmas”
- Annex L, “Summary of Documentation Requirements”
- Annex M, “Syntax Summary”
- Annex N, “Language-Defined Entities”

Each clause is divided into subclauses that have a common structure. Each clause and subclause first introduces its subject. After the introductory text, text is labeled with the following headings:

*Syntax*

Syntax rules (indented).

*Name Resolution Rules*

Compile-time rules that are used in name resolution, including overload resolution.

*Legality Rules*

Rules that are enforced at compile time. A construct is *legal* if it obeys all of the Legality Rules.

*Static Semantics*

A definition of the compile-time effect of each construct.

*Post-Compilation Rules*

Rules that are enforced before running a partition. A partition is legal if its compilation units are legal and it obeys all of the Post-Compilation Rules.

*Dynamic Semantics*

A definition of the run-time effect of each construct.

*Bounded (Run-Time) Errors*

Situations that result in bounded (run-time) errors (see 4.4).

*Erroneous Execution*

Situations that result in erroneous execution (see 4.4).

*Implementation Requirements*

Additional requirements for conforming implementations.

*Documentation Requirements*

Documentation requirements for conforming implementations.

*Metrics*

Metrics that are specified for the time/space properties of the execution of certain language constructs.

*Implementation Permissions*

Additional permissions given to the implementer.

*Implementation Advice*

Optional advice given to the implementer. It is implementation defined whether or not a given recommendation is obeyed.

NOTE Notes emphasize consequences of the rules described in the (sub)clause or elsewhere.

EXAMPLE Examples illustrate the possible forms of the constructs described.

## 4.2 Conformity of an Implementation

*Implementation Requirements*

A conforming implementation shall:

- Translate and correctly execute legal programs written in Ada, provided that they are not so large as to exceed the capacity of the implementation;
- Identify all programs or program units that are so large as to exceed the capacity of the implementation (or raise an appropriate exception at run time);
- Identify all programs or program units that contain errors whose detection is required by this document;
- Supply all language-defined library units required by this document;

- Contain no variations except those explicitly permitted by this document, or those that are impossible or impractical to avoid given the implementation's execution environment;
- Specify all such variations in the manner prescribed by this document.

The *external effect* of the execution of an Ada program is defined in terms of its interactions with its external environment. The following are defined as *external interactions*:

- Any interaction with an external file (see A.7);
- The execution of certain `code_statements` (see 16.8); which `code_statements` cause external interactions is implementation defined.
- Any call on an imported subprogram (see Annex B), including any parameters passed to it;
- Any result returned or exception propagated from a main subprogram (see 13.2) or an exported subprogram (see Annex B) to an external caller;
- Any read or update of an atomic or volatile object (see C.6);
- The values of imported and exported objects (see Annex B) at the time of any other interaction with the external environment.

A conforming implementation of this document shall produce for the execution of a given Ada program a set of interactions with the external environment whose order and timing are consistent with the definitions and requirements of this document for the semantics of the given program.

An implementation that conforms to this document shall support each capability required by the core language as specified. In addition, an implementation that conforms to this document may conform to one or more Specialized Needs Annexes (or to none). Conformance to a Specialized Needs Annex (Annex C, Annex D, Annex E, Annex F, Annex G, and Annex H) means that each capability required by the Annex shall be provided as specified.

An implementation conforming to this document may provide additional aspects, attributes, library units, and pragmas. However, it shall not provide any aspect, attribute, library unit, or pragma having the same name as an aspect, attribute, library unit, or pragma (respectively) specified in a Specialized Needs Annex unless the provided construct is either as specified in the Specialized Needs Annex or is more limited in capability than that required by the Annex. A program that attempts to use an unsupported capability of an Annex shall either be identified by the implementation before run time or shall raise an exception at run time.

For an implementation that conforms to this document, the implementation of a language-defined unit shall abide by all postconditions, type invariants, and default initial conditions specified for the unit by this document (see 14.4.2).

#### Documentation Requirements

Certain aspects of the semantics are defined to be either *implementation defined* or *unspecified*. In such cases, the set of possible effects is specified, and the implementation may choose any effect in the set. Implementations shall document their behavior in implementation-defined situations, but documentation is not required for unspecified situations. The implementation-defined characteristics are summarized in L.2.

The implementation may choose to document implementation-defined behavior either by documenting what happens in general, or by providing some mechanism for the user to determine what happens in a particular case.

#### Implementation Advice

If an implementation detects the use of an unsupported Specialized Needs Annex feature at run time, it should raise `Program_Error` if feasible.

If an implementation wishes to provide implementation-defined extensions to the functionality of a language-defined library unit, it should normally do so by adding children to the library unit.

NOTE The above requirements imply that an implementation conforming to this document can support some of the capabilities required by a Specialized Needs Annex without supporting all required capabilities.

### 4.3 Method of Description and Syntax Notation

The form of an Ada program is described by means of a context-free syntax together with context-dependent requirements expressed by narrative rules.

The meaning of Ada programs is described by means of narrative rules defining both the effects of each construct and the composition rules for constructs.

The context-free syntax of the language is described using a simple variant of Backus-Naur Form. In particular:

- Lower case words in a fixed-width font, some containing embedded underlines, are used to denote syntactic categories, for example:

```
case_statement
```

- Boldface words are used to denote reserved words, for example:

```
array
```

- Square brackets enclose optional items. Thus the two following rules are equivalent.

```
simple_return_statement ::= return [expression];  
simple_return_statement ::= return; | return expression;
```

- Curly brackets enclose a repeated item. The item may appear zero or more times; the repetitions occur from left to right as with an equivalent left-recursive rule. Thus the two following rules are equivalent.

```
term ::= factor {multiplying_operator factor}  
term ::= factor | term multiplying_operator factor
```

- A vertical line separates alternative items, for example:

```
constraint ::= scalar_constraint | composite_constraint
```

- For symbols used in this notation (square brackets, curly brackets, and the vertical line), the symbols when surrounded by ' represent themselves, for example:

```
discrete_choice_list ::= discrete_choice {' discrete_choice}  
named_container_aggregate ::=  
    '[' container_element_association_list ']'
```

- If the name of any syntactic category starts with an italicized part, it is equivalent to the category name without the italicized part. The italicized part is intended to convey some semantic information. For example *subtype\_name* and *task\_name* are both equivalent to *name* alone.

The delimiters, compound delimiters, reserved words, and `numeric_literals` are exclusively made of the characters whose code point is between 16#20# and 16#7E#, inclusively. The special characters for which names are defined in this document (see 5.1) belong to the same range. For example, the character E in the definition of `exponent` is the character whose name is “LATIN CAPITAL LETTER E”, not “GREEK CAPITAL LETTER EPSILON”.

When this document mentions the conversion of some character or sequence of characters to upper case, it means the character or sequence of characters obtained by using simple upper case mapping, as defined by documents referenced in Clause 2 of ISO/IEC 10646:2020.

A *syntactic category* is a nonterminal in the grammar defined in BNF under “Syntax”. Names of syntactic categories are set in a different font, like `this`.

A *construct* is a piece of text (explicit or implicit) that is an instance of a syntactic category defined under “Syntax”.



A *constituent* of a construct is the construct itself, or any construct appearing within it.

Whenever the run-time semantics defines certain actions to happen in an *arbitrary order*, this means that the implementation shall arrange for these actions to occur in a way that is equivalent to some sequential order, following the rules that result from that sequential order. When evaluations are defined to happen in an arbitrary order, with conversion of the results to some subtypes, or with some runtime checks, the evaluations, conversions, and checks may be arbitrarily interspersed, so long as each expression is evaluated before converting or checking its value. Note that the effect of a program can depend on the order chosen by the implementation. This can happen, for example, if two actual parameters of a given call have side effects.

NOTE 1 The syntax rules describing structured constructs are presented in a form that corresponds to the recommended paragraphing. For example, an `if_statement` is defined as:

```
if_statement ::=
    if condition then
        sequence_of_statements
    {elsif condition then
        sequence_of_statements}
    [else
        sequence_of_statements]
    end if;
```

NOTE 2 The line breaks and indentation in the syntax rules indicate the recommended line breaks and indentation in the corresponding constructs. The preferred places for other line breaks are after semicolons.

## 4.4 Classification of Errors

### Implementation Requirements

The language definition classifies errors into several different categories:

- Errors that are required to be detected prior to run time by every Ada implementation;

These errors correspond to any violation of a rule given in this document, other than those listed below. In particular, violation of any rule that uses the terms *shall*, *allowed*, *permitted*, *legal*, or *illegal* belongs to this category. Any program that contains such an error is not a legal Ada program; on the other hand, the fact that a program is legal does not mean, *per se*, that the program is free from other forms of error.

The rules are further classified as either compile time rules, or post compilation rules, depending on whether a violation has to be detected at the time a compilation unit is submitted to the compiler, or may be postponed until the time a compilation unit is incorporated into a partition of a program.

- Errors that are required to be detected at run time by the execution of an Ada program;

The corresponding error situations are associated with the names of the predefined exceptions. Every Ada compiler is required to generate code that raises the corresponding exception if such an error situation arises during program execution. If such an error situation is certain to arise in every execution of a construct, then an implementation is allowed (although not required) to report this fact at compilation time.

- Bounded errors;

The language rules define certain kinds of errors that are not expected to be detected either prior to or during run time, but if not detected, the range of possible effects shall be bounded. The errors of this category are called *bounded errors*. The possible effects of a given bounded error are specified for each such error, but in any case one possible effect of a bounded error is the raising of the exception `Program_Error`.

- Erroneous execution.

In addition to bounded errors, the language rules define certain kinds of errors as leading to *erroneous execution*. Like bounded errors, the implementation is not expected to detect such errors

either prior to or during run time. Unlike bounded errors, there is no language-specified bound on the possible effect of erroneous execution; the effect is in general not predictable.

#### Implementation Permissions

An implementation may provide *nonstandard modes* of operation. Typically these modes would be selected by a `pragma` or by a command line switch when the compiler is invoked. When operating in a nonstandard mode, the implementation may reject `compilation_units` that do not conform to additional requirements associated with the mode, such as an excessive number of warnings or violation of coding style guidelines. Similarly, in a nonstandard mode, the implementation may apply special optimizations or alternative algorithms that are only meaningful for programs that satisfy certain criteria specified by the implementation. In any case, an implementation shall support a *standard* mode that conforms to the requirements of this document; in particular, in the standard mode, all legal `compilation_units` shall be accepted.

#### Implementation Advice

If an implementation detects a bounded error or erroneous execution, it should raise `Program_Error`.

## 5 Lexical Elements

The text of a program consists of the texts of one or more `compilations`. The text of a `compilation` is a sequence of lexical elements, each composed of characters; the rules of composition are given in this clause. `Pragmas`, which provide certain information for the compiler, are also described in this clause.

### 5.1 Character Set

The character repertoire for the text of an Ada program consists of the entire coding space described by the ISO/IEC 10646:2020 Universal Coded Character Set. This coding space is organized in *planes*, each plane comprising 65536 characters.

#### Syntax

A `character` is defined by this document for each cell in the coding space described by ISO/IEC 10646:2020, regardless of whether or not ISO/IEC 10646:2020 allocates a character to that cell.

#### Static Semantics

The coded representation for characters is implementation defined (it can be a representation that is not defined within ISO/IEC 10646:2020). A character whose relative code point in its plane is 16#FFFE# or 16#FFFF# is not allowed anywhere in the text of a program. The only characters allowed outside of comments are those in categories `other_format`, `format_effector`, and `graphic_character`.

The semantics of an Ada program whose text is not in Normalization Form C (as defined by Clause 22 of ISO/IEC 10646:2020) is implementation defined.

The description of the language definition in this document uses the character properties General Category, Simple Uppercase Mapping, Uppercase Mapping, and Special Case Condition of the documents referenced by Clause 2 of ISO/IEC 10646:2020. The actual set of graphic symbols used by an implementation for the visual representation of the text of an Ada program is not specified.

Characters are categorized as follows:

`letter_uppercase`

Any character whose General Category is defined to be "Letter, Uppercase".

`letter_lowercase`

Any character whose General Category is defined to be "Letter, Lowercase".

|                                     |                                                                                                                                                                                                                                                                                         |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>letter_titlecase</code>       | Any character whose General Category is defined to be “Letter, Titlecase”.                                                                                                                                                                                                              |
| <code>letter_modifier</code>        | Any character whose General Category is defined to be “Letter, Modifier”.                                                                                                                                                                                                               |
| <code>letter_other</code>           | Any character whose General Category is defined to be “Letter, Other”.                                                                                                                                                                                                                  |
| <code>mark_non_spacing</code>       | Any character whose General Category is defined to be “Mark, Non-Spacing”.                                                                                                                                                                                                              |
| <code>mark_spacing_combining</code> | Any character whose General Category is defined to be “Mark, Spacing Combining”.                                                                                                                                                                                                        |
| <code>number_decimal</code>         | Any character whose General Category is defined to be “Number, Decimal”.                                                                                                                                                                                                                |
| <code>number_letter</code>          | Any character whose General Category is defined to be “Number, Letter”.                                                                                                                                                                                                                 |
| <code>punctuation_connector</code>  | Any character whose General Category is defined to be “Punctuation, Connector”.                                                                                                                                                                                                         |
| <code>other_format</code>           | Any character whose General Category is defined to be “Other, Format”.                                                                                                                                                                                                                  |
| <code>separator_space</code>        | Any character whose General Category is defined to be “Separator, Space”.                                                                                                                                                                                                               |
| <code>separator_line</code>         | Any character whose General Category is defined to be “Separator, Line”.                                                                                                                                                                                                                |
| <code>separator_paragraph</code>    | Any character whose General Category is defined to be “Separator, Paragraph”.                                                                                                                                                                                                           |
| <code>format_effector</code>        | The characters whose code points are 16#09# (CHARACTER TABULATION), 16#0A# (LINE FEED), 16#0B# (LINE TABULATION), 16#0C# (FORM FEED), 16#0D# (CARRIAGE RETURN), 16#85# (NEXT LINE), and the characters in categories <code>separator_line</code> and <code>separator_paragraph</code> . |
| <code>other_control</code>          | Any character whose General Category is defined to be “Other, Control”, and which is not defined to be a <code>format_effector</code> .                                                                                                                                                 |
| <code>other_private_use</code>      | Any character whose General Category is defined to be “Other, Private Use”.                                                                                                                                                                                                             |
| <code>other_surrogate</code>        | Any character whose General Category is defined to be “Other, Surrogate”.                                                                                                                                                                                                               |
| <code>graphic_character</code>      | Any character that is not in the categories <code>other_control</code> , <code>other_private_use</code> , <code>other_surrogate</code> , <code>format_effector</code> , and whose relative code point in its plane is neither 16#FFFE# nor 16#FFFF#.                                    |

The following names are used when referring to certain characters (the first name is that given in ISO/IEC 10646:2020):

| graphic symbol | name             | graphic symbol | name           |
|----------------|------------------|----------------|----------------|
| "              | quotation mark   | :              | colon          |
| #              | number sign      | ;              | semicolon      |
| &              | ampersand        | <              | less-than sign |
| '              | apostrophe, tick | =              | equals sign    |

|   |                        |   |                      |
|---|------------------------|---|----------------------|
| ( | left parenthesis       | > | greater-than sign    |
| ) | right parenthesis      | - | low line, underline  |
| * | asterisk, multiply     |   | vertical line        |
| + | plus sign              | / | solidus, divide      |
| , | comma                  | ! | exclamation point    |
| - | hyphen-minus, minus    | % | percent sign         |
| . | full stop, dot, point  | [ | left square bracket  |
| @ | commercial at, at sign | ] | right square bracket |

#### Implementation Requirements

An Ada implementation shall accept Ada source code in UTF-8 encoding, with or without a BOM (see A.4.11), where every character is represented by its code point. The character pair CARRIAGE RETURN/LINE FEED (code points 16#0D# 16#0A#) signifies a single end of line (see 5.2); every other occurrence of a `format_effector` other than the character whose code point position is 16#09# (CHARACTER TABULATION) also signifies a single end of line.

#### Implementation Permissions

The categories defined above, as well as case mapping and folding, may be based on an implementation-defined version of ISO/IEC 10646 (2003 edition or later).

NOTE The characters in categories `other_control`, `other_private_use`, and `other_surrogate` are only allowed in comments.

## 5.2 Lexical Elements, Separators, and Delimiters

#### Static Semantics

The text of a program consists of the texts of one or more compilations. The text of each compilation is a sequence of separate *lexical elements*. Each lexical element is formed from a sequence of characters, and is either a delimiter, an identifier, a reserved word, a `numeric_literal`, a `character_literal`, a `string_literal`, or a comment. The meaning of a program depends only on the particular sequences of lexical elements that form its compilations, excluding comments.

The text of a compilation is divided into *lines*. In general, the representation for an end of line is implementation defined. However, a sequence of one or more `format_effectors` other than the character whose code point is 16#09# (CHARACTER TABULATION) signifies at least one end of line.

In some cases an explicit *separator* is required to separate adjacent lexical elements. A separator is any of a `separator_space`, a `format_effector`, or the end of a line, as follows:

- A separator space is a separator except within a comment, a `string_literal`, or a `character_literal`.
- The character whose code point is 16#09# (CHARACTER TABULATION) is a separator except within a comment.
- The end of a line is always a separator.

One or more separators are allowed between any two adjacent lexical elements, before the first of each compilation, or after the last. At least one separator is required between an identifier, a reserved word, or a `numeric_literal` and an adjacent identifier, reserved word, or `numeric_literal`.

One or more `other_format` characters are allowed anywhere that a separator is; any such characters have no effect on the meaning of an Ada program.

A *delimiter* is either one of the following characters:

& ' ( ) \* + , - . / : ; < = > @ [ ] |

or one of the following *compound delimiters* each composed of two adjacent special characters

=> .. \*\* := /= >= <= << >> <>

Each of the special characters listed for single character delimiters is a single delimiter except if this character is used as a character of a compound delimiter, or as a character of a comment, `string_literal`, `character_literal`, or `numeric_literal`.

The following names are used when referring to compound delimiters:

| delimiter | name                                 |
|-----------|--------------------------------------|
| =>        | arrow                                |
| ..        | double dot                           |
| **        | double star, exponentiate            |
| :=        | assignment (pronounced: “becomes”)   |
| /=        | inequality (pronounced: “not equal”) |
| >=        | greater than or equal                |
| <=        | less than or equal                   |
| <<        | left label bracket                   |
| >>        | right label bracket                  |
| <>        | box                                  |

#### Implementation Requirements

An implementation shall support lines of at least 200 characters in length, not counting any characters used to signify the end of a line. An implementation shall support lexical elements of at least 200 characters in length. The maximum supported line length and lexical element length are implementation defined.

### 5.3 Identifiers

Identifiers are used as names.

#### Syntax

```

identifier ::=
  identifier_start { identifier_start | identifier_extend }
identifier_start ::=
  letter_uppercase
| letter_lowercase
| letter_titlecase
| letter_modifier
| letter_other
| number_letter
identifier_extend ::=
  mark_non_spacing
| mark_spacing_combining
| number_decimal
| punctuation_connector

```

An identifier shall not contain two consecutive characters in category `punctuation_connector`, or end with a character in that category.

*Legality Rules*

An identifier shall only contain characters that may be present in Normalization Form KC as defined by Clause 22 of ISO/IEC 10646:2020.

*Static Semantics*

Two identifiers are considered the same if they consist of the same sequence of characters after applying locale-independent simple case folding, as defined by documents referenced in Clause 2 of ISO/IEC 10646:2020.

After applying simple case folding, an identifier shall not be identical to a reserved word.

*Implementation Permissions*

In a nonstandard mode, an implementation may support other upper/lower case equivalence rules for identifiers, to accommodate local conventions.

NOTE Identifiers differing only in the use of corresponding upper and lower case letters are considered the same.

EXAMPLE *Examples of identifiers:*

|            |    |             |                 |        |
|------------|----|-------------|-----------------|--------|
| Count      | X  | Get_Symbol  | Ethelyn         | Marion |
| Snobol_4   | X1 | Page_Count  | Store_Next_Item |        |
| Πλάτων     | -- | Plato       |                 |        |
| Чайковский | -- | Tchaikovsky |                 |        |
| θ φ        | -- | Angles      |                 |        |

## 5.4 Numeric Literals

There are two kinds of numeric\_literals, *real literals* and *integer literals*. A real literal is a numeric\_literal that includes a point; an integer literal is a numeric\_literal without a point.

*Syntax*

```
numeric_literal ::= decimal_literal | based_literal
```

NOTE The type of an integer literal is *universal\_integer*. The type of a real literal is *universal\_real*.

### 5.4.1 Decimal Literals

A decimal\_literal is a numeric\_literal in the conventional decimal notation (that is, the base is ten).

*Syntax*

```
decimal_literal ::= numeral [numeral] [exponent]
numeral ::= digit {[underline] digit}
exponent ::= E [+] numeral | E - numeral
digit ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
```

An exponent for an integer literal shall not have a minus sign.

*Static Semantics*

An underline character in a numeric\_literal does not affect its meaning. The letter E of an exponent can be written either in lower case or in upper case, with the same meaning.

An exponent indicates the power of ten by which the value of the decimal\_literal without the exponent is to be multiplied to obtain the value of the decimal\_literal with the exponent.

EXAMPLE *Examples of decimal literals:*

```
12          0          1E6      123_456      -- integer literals
12.0        0.0        0.456    3.14159_26 -- real literals
```

## 5.4.2 Based Literals

A `based_literal` is a `numeric_literal` expressed in a form that specifies the base explicitly.

### Syntax

```
based_literal ::=
    base # based_numeral [based_numeral] # [exponent]
base ::= numeral
based_numeral ::=
    extended_digit {[underline] extended_digit}
extended_digit ::= digit | A | B | C | D | E | F
```

### Legality Rules

The *base* (the numeric value of the decimal numeral preceding the first #) shall be at least two and at most sixteen. The extended\_digits A through F represent the digits ten through fifteen, respectively. The value of each extended\_digit of a based\_literal shall be less than the base.

### Static Semantics

The conventional meaning of based notation is assumed. An exponent indicates the power of the base by which the value of the `based_literal` without the exponent is to be multiplied to obtain the value of the `based_literal` with the exponent. The base and the exponent, if any, are in decimal notation.

The extended\_digits A through F can be written either in lower case or in upper case, with the same meaning.

EXAMPLE *Examples of based literals:*

```
2#1111_1111#  16#FF#      016#0ff#  -- integer literals of value 255
16#E#E1       2#1110_0000#  -- integer literals of value 224
16#F.FF#E+2   2#1.1111_1111_1110#E11 -- real literals of value 4095.0
```

## 5.5 Character Literals

A `character_literal` is formed by enclosing a graphic character between two apostrophe characters.

### Syntax

```
character_literal ::= 'graphic_character'
```

NOTE A `character_literal` is an enumeration literal of a character type. See 6.5.2.

EXAMPLE *Examples of character literals:*

```
'A'      '*'      ' '      ' '
'L'      'л'      'Λ'      -- Various els.
'∞'      'ℵ'      -- Big numbers - infinity and aleph.
```

## 5.6 String Literals

A `string_literal` is formed by a sequence of graphic characters (possibly none) enclosed between two quotation marks used as string brackets. They are used to represent `operator_symbols` (see 9.1), values of a string type (see 7.2), and array subaggregates (see 7.3.3).



*Syntax*

```
string_literal ::= "{string_element}"
string_element ::= "" | non_quotation_mark_graphic_character
```

A `string_element` is either a pair of quotation marks (`""`), or a single `graphic_character` other than a quotation mark.

*Static Semantics*

The *sequence of characters* of a `string_literal` is formed from the sequence of `string_elements` between the bracketing quotation marks, in the given order, with a `string_element` that is `""` becoming a single quotation mark in the sequence of characters, and any other `string_element` being reproduced in the sequence.

A *null string literal* is a `string_literal` with no `string_elements` between the quotation marks.

NOTE 1 An end of line cannot appear in a `string_literal`.

NOTE 2 No transformation is performed on the sequence of characters of a `string_literal`.

EXAMPLE *Examples of string literals:*

```
"Message of the day:"

""                -- a null string literal
" "  "A"  """"    -- three string literals of length 1

"Characters such as $, %, and } are allowed in string literals"
"Archimedes said ""Εύρηκα""
"Volume of cylinder ( $\pi r^2 h$ ) = "
```

## 5.7 Comments

A comment starts with two adjacent hyphens and extends up to the end of the line.

*Syntax*

```
comment ::= --{non_end_of_line_character}
```

A comment may appear on any line of a program.

*Static Semantics*

The presence or absence of comments has no influence on whether a program is legal or illegal. Furthermore, comments do not influence the meaning of a program; their sole purpose is the enlightenment of the human reader.

EXAMPLE *Examples of comments:*

```
-- the last sentence above echoes the Algol 68 report

end;  -- processing of Line is complete

-- a long comment can be split onto
-- two or more consecutive lines

----- the first two hyphens start the comment
```

## 5.8 Pragmas

A pragma is a compiler directive. There are language-defined pragmas that give instructions for optimization, listing control, etc. An implementation may support additional (implementation-defined) pragmas.

*Syntax*

```

pragma ::=
  pragma identifier
  [(pragma_argument_association {, pragma_argument_association})];
pragma_argument_association ::=
  [pragma_argument_identifier =>] name
  | [pragma_argument_identifier =>] expression
  | pragma_argument_aspect_mark => name
  | pragma_argument_aspect_mark => expression

```

In a pragma, any pragma\_argument\_associations without a *pragma\_argument\_identifier* or *pragma\_argument\_aspect\_mark* shall precede any associations with a *pragma\_argument\_identifier* or *pragma\_argument\_aspect\_mark*.

Pragmas are only allowed at the following places in a program:

- After a semicolon delimiter, but not within a *formal\_part*, *discriminant\_part*, or *declare\_expression*.
- At any place where the syntax rules allow a construct defined by a syntactic category whose name ends with “declaration”, “item”, “statement”, “clause”, or “alternative”, or one of the syntactic categories *variant* or *exception\_handler*; but not in place of such a construct if the construct is required, or is part of a list that is required to have at least one such construct.
- In place of a statement in a *sequence\_of\_statements*.
- At any place where a *compilation\_unit* is allowed.

Additional syntax rules and placement restrictions exist for specific pragmas.

The *name* of a pragma is the identifier following the reserved word **pragma**. The name or expression of a *pragma\_argument\_association* is a *pragma\_argument*.

An *identifier specific to a pragma* is an identifier or reserved word that is used in a pragma argument with special meaning for that pragma.

*Static Semantics*

If an implementation does not recognize the name of a pragma, then it has no effect on the semantics of the program. Inside such a pragma, the only rules that apply are the Syntax Rules.

*Dynamic Semantics*

Any pragma that appears at the place of an executable construct is executed. Unless otherwise specified for a particular pragma, this execution consists of the evaluation of each evaluable pragma argument in an arbitrary order.

*Implementation Requirements*

The implementation shall give a warning message for an unrecognized pragma name.

*Implementation Permissions*

An implementation may provide implementation-defined pragmas; the name of an implementation-defined pragma shall differ from those of the language-defined pragmas.

An implementation may ignore an unrecognized pragma even if it violates some of the Syntax Rules, if detecting the syntax error is too complex.

*Implementation Advice*

Normally, implementation-defined pragmas should have no semantic effect for error-free programs; that is, if the implementation-defined pragmas in a working program are replaced with unrecognized pragmas, the program should still be legal, and should still have the same semantics.

Normally, an implementation should not define pragmas that can make an illegal program legal, except as follows:

- A pragma used to complete a declaration;
- A pragma used to configure the environment by adding, removing, or replacing `library_items`.

*Syntax*

The forms of List, Page, and Optimize pragmas are as follows:

**pragma** List(identifier);

**pragma** Page;

**pragma** Optimize(identifier);

Other pragmas are defined throughout this document, and are summarized in Annex K.

*Static Semantics*

A `pragma List` takes one of the identifiers `On` or `Off` as the single argument. This pragma is allowed anywhere a pragma is allowed. It specifies that listing of the compilation is to be continued or suspended until a List pragma with the opposite argument is given within the same compilation. The pragma itself is always listed if the compiler is producing a listing.

A `pragma Page` is allowed anywhere a pragma is allowed. It specifies that the program text which follows the pragma should start on a new page (if the compiler is currently producing a listing).

A `pragma Optimize` takes one of the identifiers `Time`, `Space`, or `Off` as the single argument. This pragma is allowed anywhere a pragma is allowed, and it applies until the end of the immediately enclosing declarative region, or for a pragma at the place of a `compilation_unit`, to the end of the compilation. It gives advice to the implementation as to whether time or space is the primary optimization criterion, or that optional optimizations should be turned off. It is implementation defined how this advice is followed.

EXAMPLE *Examples of pragmas:*

```
pragma List(Off); -- turn off listing generation
pragma Optimize(Off); -- turn off optional optimizations
pragma Assertion_Policy(Check); -- check assertions
pragma Assert(Exists(File_Name),
    Message => "Nonexistent file"); -- assert file exists
```

## 5.9 Reserved Words

### Syntax

The following are the *reserved words*. Within a program, some or all of the letters of a reserved word may be in upper case.

|                 |                  |                   |                     |
|-----------------|------------------|-------------------|---------------------|
| <b>abort</b>    | <b>else</b>      | <b>new</b>        | <b>return</b>       |
| <b>abs</b>      | <b>elsif</b>     | <b>not</b>        | <b>reverse</b>      |
| <b>abstract</b> | <b>end</b>       | <b>null</b>       | <b>select</b>       |
| <b>accept</b>   | <b>entry</b>     | <b>of</b>         | <b>separate</b>     |
| <b>access</b>   | <b>exception</b> | <b>or</b>         | <b>some</b>         |
| <b>aliased</b>  | <b>exit</b>      | <b>others</b>     | <b>subtype</b>      |
| <b>all</b>      | <b>for</b>       | <b>out</b>        | <b>synchronized</b> |
| <b>and</b>      | <b>function</b>  | <b>overriding</b> | <b>tagged</b>       |
| <b>array</b>    | <b>generic</b>   | <b>package</b>    | <b>task</b>         |
| <b>at</b>       | <b>goto</b>      | <b>parallel</b>   | <b>terminate</b>    |
| <b>begin</b>    | <b>if</b>        | <b>pragma</b>     | <b>then</b>         |
| <b>body</b>     | <b>in</b>        | <b>private</b>    | <b>type</b>         |
| <b>case</b>     | <b>interface</b> | <b>procedure</b>  | <b>until</b>        |
| <b>constant</b> | <b>is</b>        | <b>protected</b>  | <b>use</b>          |
| <b>declare</b>  | <b>limited</b>   | <b>raise</b>      | <b>when</b>         |
| <b>delay</b>    | <b>loop</b>      | <b>range</b>      | <b>while</b>        |
| <b>delta</b>    | <b>mod</b>       | <b>record</b>     | <b>with</b>         |
| <b>digits</b>   |                  | <b>rem</b>        | <b>xor</b>          |
| <b>do</b>       |                  | <b>renames</b>    |                     |
|                 |                  | <b>requeue</b>    |                     |

NOTE The reserved words appear in **lower case boldface** in this document, except when used in the designator of an attribute (see 7.1.4). Lower case boldface is also used for a reserved word in a `string_literal` used as an `operator_symbol`. This is merely a convention — programs can be written in whatever typeface is desired and available.

## 6 Declarations and Types

This clause describes the types in the language and the rules for declaring constants, variables, and named numbers.

### 6.1 Declarations

The language defines several kinds of named *entities* that are declared by declarations. The entity's *name* is defined by the declaration, usually by a `defining_identifier`, but sometimes by a `defining_character_literal` or `defining_operator_symbol`. There are also entities that are not directly declared; some of these are elements of other entities, or are allocated dynamically. Such entities can be denoted using `indexed_component`, `selected_component`, or dereference names (see 7.1).

There are several forms of declaration. A `basic_declaration` is a form of declaration defined as follows.

### Syntax

```
basic_declaration ::=
  type_declaration           | subtype_declaration
| object_declaration         | number_declaration
| subprogram_declaration     | abstract_subprogram_declaration
| null_procedure_declaration | expression_function_declaration
```

|                       |                      |
|-----------------------|----------------------|
| package_declaration   | renaming_declaration |
| exception_declaration | generic_declaration  |
| generic_instantiation |                      |

defining\_identifier ::= identifier

#### Static Semantics

A *declaration* is a language construct that associates a name with (a view of) an entity. A declaration may appear explicitly in the program text (an *explicit* declaration), or may be supposed to occur at a given place in the text as a consequence of the semantics of another construct (an *implicit* declaration).

Each of the following is defined to be a declaration: any basic\_declaration; an enumeration\_literal\_specification; a discriminant\_specification; a component\_declaration; a defining\_identifier of an iterated\_component\_association; a loop\_parameter\_specification; a defining\_identifier of a chunk\_specification; an iterator\_specification; a defining\_identifier of an iterator\_parameter\_specification; a parameter\_specification; a subprogram\_body; an extended\_return\_object\_declaration; an entry\_declaration; an entry\_index\_specification; a choice\_parameter\_specification; a generic\_formal\_parameter\_declaration.

All declarations contain a *definition* for a *view* of an entity. A view consists of an identification of the entity (the entity *of* the view), plus view-specific characteristics that affect the use of the entity through that view (such as mode of access to an object, formal parameter names and defaults for a subprogram, or visibility to components of a type). In most cases, a declaration also contains the definition for the entity itself (a renaming\_declaration is an example of a declaration that does not define a new entity, but instead defines a view of an existing entity (see 11.5)).

When it is clear from context, the term *object* is used in place of *view of an object*. Similarly, the terms *type* and *subtype* are used in place of *view of a type* and *view of a subtype*, respectively.

For each declaration, the language rules define a certain region of text called the *scope* of the declaration (see 11.2). Most declarations associate an identifier with a declared entity. Within its scope, and only there, there are places where it is possible to use the identifier to refer to the declaration, the view it defines, and the associated entity; these places are defined by the visibility rules (see 11.3). At such places the identifier is said to be a *name* of the entity (the direct\_name or selector\_name); the name is said to *denote* the declaration, the view, and the associated entity (see 11.6). The declaration is said to *declare* the name, the view, and in most cases, the entity itself.

As an alternative to an identifier, an enumeration literal can be declared with a character\_literal as its name (see 6.5.1), and a function can be declared with an operator\_symbol as its name (see 9.1).

The syntax rules use the terms defining\_identifier, defining\_character\_literal, and defining\_operator\_symbol for the defining occurrence of a name; these are collectively called *defining names*. The terms direct\_name and selector\_name are used for usage occurrences of identifiers, character\_literals, and operator\_symbols. These are collectively called *usage names*.

#### Dynamic Semantics

The process by which a construct achieves its run-time effect is called *execution*. This process is also called *elaboration* for declarations and *evaluation* for expressions. One of the terms execution, elaboration, or evaluation is defined by this document for each construct that has a run-time effect.

NOTE At compile time, the declaration of an entity *declares* the entity. At run time, the elaboration of the declaration *creates* the entity.

## 6.2 Types and Subtypes

### Static Semantics

A *type* is characterized by a set of values, and a set of *primitive operations* which implement the fundamental aspects of its semantics. An *object* of a given type is a run-time entity that contains (has) a value of the type.

Types are grouped into *categories* of types. There exist several *language-defined categories* of types reflecting the similarity of their values and primitive operations. Most categories of types form *classes* of types. *Elementary* types are those whose values are logically indivisible; *composite* types are those whose values are composed of *component* values.

The elementary types are the *scalar* types (*discrete* and *real*) and the *access* types (whose values provide access to objects or subprograms). Discrete types are either *integer* types or are defined by enumeration of their values (*enumeration* types). Real types are either *floating point* types or *fixed point* types.

The composite types are the *record* types, *record extensions*, *array* types, *interface* types, *task* types, and *protected* types.

There can be multiple views of a type with varying sets of operations. An *incomplete* type represents an incomplete view (see 6.10.1) of a type with a very restricted usage, providing support for recursive data structures. A *private* type or *private extension* represents a partial view (see 10.3) of a type, providing support for data abstraction. The full view (see 6.2.1) of a type represents its complete definition. An incomplete or partial view is considered a composite type, even if the full view is not.

Certain composite types (and views thereof) have special components called *discriminants* whose values affect the presence, constraints, or initialization of other components. Discriminants can be thought of as parameters of the type.

The term *subcomponent* is used in this document in place of the term *component* to indicate either a component, or a component of another subcomponent. Where other subcomponents are excluded, the term *component* is used instead. Similarly, a *part* of an object or value is used to mean the whole object or value, or any set of its subcomponents. The terms *component*, *subcomponent*, and *part* are also applied to a type meaning the component, subcomponent, or part of objects and values of the type.

The set of possible values for an object of a given type can be subjected to a condition that is called a *constraint* (the case of a *null constraint* that specifies no restriction is also included); the rules for which values satisfy a given kind of constraint are given in 6.5 for *range\_constraints*, 6.6.1 for *index\_constraints*, and 6.7.1 for *discriminant\_constraints*. The set of possible values for an object of an access type can also be subjected to a condition that excludes the null value (see 6.10).

A *subtype* of a given type is a combination of the type, a constraint on values of the type, and certain attributes specific to the subtype. The given type is called the *type of the subtype*. Similarly, the associated constraint is called the *constraint of the subtype*. The set of values of a subtype consists of the values of its type that satisfy its constraint and any exclusion of the null value. Such values *belong* to the subtype. The other values of the type are *outside* the subtype.

A subtype is called an *unconstrained* subtype if its type has unknown discriminants, or if its type allows range, index, or discriminant constraints, but the subtype does not impose such a constraint; otherwise, the subtype is called a *constrained* subtype (since it has no unconstrained characteristics).

NOTE Any set of types can be called a “category” of types, and any set of types that is closed under derivation (see 6.4) can be called a “class” of types. However, only certain categories and classes are used in the description of the rules of the language — generally those that have their own particular set of primitive operations (see 6.2.3), or that correspond to a set of types that are matched by a given kind of generic formal type (see 15.5). The following are examples of “interesting” *language-defined classes*: elementary, scalar, discrete, enumeration, character, boolean, integer, signed integer, modular, real, floating point, fixed point, ordinary fixed point, decimal fixed point, numeric, access, access-to-object, access-to-subprogram, composite, array, string, (untagged) record, tagged, task, protected, nonlimited. Special syntax is provided to define types in each of these classes. In addition to these

classes, the following are examples of “interesting” *language-defined categories*: abstract, incomplete, interface, limited, private, record.

These language-defined categories are organized like this:

- all types
  - elementary
    - scalar
      - discrete
        - enumeration
          - character
          - boolean
          - other enumeration
        - integer
          - signed integer
          - modular integer
      - real
        - floating point
        - fixed point
          - ordinary fixed point
          - decimal fixed point
    - access
      - access-to-object
      - access-to-subprogram
    - composite
      - untagged
        - array
          - string
          - other array
        - record
        - task
        - protected
        - tagged (including interfaces)
          - nonlimited tagged record
          - limited tagged
            - limited tagged record
            - synchronized tagged
              - tagged task
              - tagged protected

There are other categories, such as “numeric” and “discriminated”, which represent other categorization dimensions, but do not fit into the above strictly hierarchical picture.

### 6.2.1 Type Declarations

A `type_declaration` declares a type and its first subtype.

*Syntax*

```

type_declaration ::= full_type_declaration
| incomplete_type_declaration
| private_type_declaration
| private_extension_declaration

full_type_declaration ::=
  type defining_identifier [known_discriminant_part] is type_definition
  [aspect_specification];
| task_type_declaration
| protected_type_declaration

type_definition ::=
  enumeration_type_definition | integer_type_definition
| real_type_definition | array_type_definition
| record_type_definition | access_type_definition
  
```



| derived\_type\_definition | interface\_type\_definition

#### Legality Rules

A given type shall not have a subcomponent whose type is the given type itself.

#### Static Semantics

The `defining_identifier` of a `type_declaration` denotes the *first subtype* of the type. The `known_discriminant_part`, if any, defines the discriminants of the type (see 6.7). The remainder of the `type_declaration` defines the remaining characteristics of (the view of) the type.

A type defined by a `type_declaration` is a *named type*; such a type has one or more nameable subtypes. Certain other forms of declaration also include type definitions as part of the declaration for an object. The type defined by such a declaration is *anonymous* — it has no nameable subtypes. For explanatory purposes, this document sometimes refers to an anonymous type by a pseudo-name, written in italics, and uses such pseudo-names at places where the syntax normally requires an identifier. For a named type whose first subtype is T, this document sometimes refers to the type of T as simply “the type T”.

A named type that is declared by a `full_type_declaration`, or an anonymous type that is defined by an `access_definition` or as part of declaring an object of the type, is called a *full type*. The declaration of a full type also declares the *full view* of the type. The `type_definition`, `task_definition`, `protected_definition`, or `access_definition` that defines a full type is called a *full type definition*. Types declared by other forms of type declaration are not separate types; they are partial or incomplete views of some full type.

The definition of a type implicitly declares certain *predefined operators* that operate on the type, according to what classes the type belongs, as specified in 7.5.

The *predefined types* (for example the types Boolean, Wide\_Character, Integer, *root\_integer*, and *universal\_integer*) are the types that are defined in a predefined library package called Standard; this package also includes the (implicit) declarations of their predefined operators. The package Standard is described in A.1.

#### Dynamic Semantics

The elaboration of a `full_type_declaration` consists of the elaboration of the full type definition. Each elaboration of a full type definition creates a distinct type and its first subtype.

EXAMPLE 1 *Examples of type definitions:*

```
(White, Red, Yellow, Green, Blue, Brown, Black)
range 1 .. 72
array(1 .. 10) of Integer
```

EXAMPLE 2 *Examples of type declarations:*

```
type Color is (White, Red, Yellow, Green, Blue, Brown, Black);
type Column is range 1 .. 72;
type Table is array(1 .. 10) of Integer;
```

NOTE Each of the above examples declares a named type. The identifier given denotes the first subtype of the type. Other named subtypes of the type can be declared with `subtype_declarations` (see 6.2.2). Although names do not directly denote types, a phrase like “the type Column” is sometimes used in this document to refer to the type of Column, where Column denotes the first subtype of the type. For an example of the definition of an anonymous type, see the declaration of the array Color\_Table in 6.3.1; its type is anonymous — it has no nameable subtypes.

## 6.2.2 Subtype Declarations

A `subtype_declaration` declares a subtype of some previously declared type, as defined by a `subtype_indication`.

*Syntax*

```

subtype_declaration ::=
  subtype defining_identifier is subtype_indication
    [aspect_specification];
subtype_indication ::= [null_exclusion] subtype_mark [constraint]
subtype_mark ::= subtype_name
constraint ::= scalar_constraint | composite_constraint
scalar_constraint ::=
  range_constraint | digits_constraint | delta_constraint
composite_constraint ::=
  index_constraint | discriminant_constraint

```

*Name Resolution Rules*

A `subtype_mark` shall resolve to denote a subtype. The type *determined* by a `subtype_mark` is the type of the subtype denoted by the `subtype_mark`.

*Dynamic Semantics*

The elaboration of a `subtype_declaration` consists of the elaboration of the `subtype_indication`. The elaboration of a `subtype_indication` creates a new subtype. If the `subtype_indication` does not include a constraint, the new subtype has the same (possibly null) constraint as that denoted by the `subtype_mark`. The elaboration of a `subtype_indication` that includes a constraint proceeds as follows:

- The constraint is first elaborated.
- A check is then made that the constraint is *compatible* with the subtype denoted by the `subtype_mark`.

The condition imposed by a constraint is the condition obtained after elaboration of the constraint. The rules defining compatibility are given for each form of constraint in the appropriate subclause. These rules are such that if a constraint is *compatible* with a subtype, then the condition imposed by the constraint cannot contradict any condition already imposed by the subtype on its values. The exception `Constraint_Error` is raised if any check of compatibility fails.

NOTE A `scalar_constraint` can be applied to a subtype of an appropriate scalar type (see 6.5, 6.5.9, and I.3), even if the subtype is already constrained. On the other hand, a `composite_constraint` can be applied to a composite subtype (or an access-to-composite subtype) only if the composite subtype is unconstrained (see 6.6.1 and 6.7.1).

EXAMPLE *Examples of subtype declarations:*

```

subtype Rainbow    is Color range Red .. Blue;           -- see 6.2.1
subtype Red_Blue   is Rainbow;
subtype Int         is Integer;
subtype Small_Int   is Integer range -10 .. 10;
subtype Up_To_K     is Column range 1 .. K;               -- see 6.2.1
subtype Square      is Matrix(1 .. 10, 1 .. 10);          -- see 6.6
subtype Male        is Person(Sex => M);                  -- see 6.10.1
subtype Binop_Ref   is not null Binop_Ptr;                 -- see 6.10

```

### 6.2.3 Classification of Operations

*Static Semantics*

An operation *operates on a type T* if it yields a value of type *T*, if it has an operand whose expected type (see 11.6) is *T*, or if it has an access parameter or access result type (see 9.1) designating *T*. A predefined operator, or other language-defined operation such as assignment or a membership test,

that operates on a type, is called a *predefined operation* of the type. The *primitive operations* of a type are the predefined operations of the type, plus any user-defined primitive subprograms.

The *primitive subprograms* of a specific type are defined as follows:

- The predefined operators of the type (see 7.5);
- For a derived type, the inherited (see 6.4) user-defined subprograms;
- For an enumeration type, the enumeration literals (which are considered parameterless functions — see 6.5.1);
- For a specific type declared immediately within a `package_specification`, any subprograms (in addition to the enumeration literals) that are explicitly declared immediately within the same `package_specification` and that operate on the type;
- For a specific type with an explicitly declared primitive "=" operator whose result type is Boolean, the corresponding "/=" operator (see 9.6);
- For a nonformal type, any subprograms not covered above that are explicitly declared immediately within the same declarative region as the type and that override (see 11.3) other implicitly declared primitive subprograms of the type.

A primitive subprogram whose designator is an `operator_symbol` is called a *primitive operator*.

#### 6.2.4 Subtype Predicates

The language-defined *predicate aspects* `Static_Predicate` and `Dynamic_Predicate` may be used to define properties of subtypes. A *predicate specification* is an `aspect_specification` for one of the two predicate aspects. General rules for aspects and `aspect_specifications` are found in Clause 16 (16.1 and 16.1.1 respectively). The predicate aspects are assertion aspects (see 14.4.2). The predicate aspects are not inherited, but their effects are additive, as defined below.

##### Name Resolution Rules

The expected type for a predicate `aspect_expression` is any boolean type.

##### Static Semantics

A predicate specification may be given on a `type_declaration` or a `subtype_declaration`, and applies to the declared subtype. In addition, predicate specifications apply to certain other subtypes:

- For a (first) subtype defined by a type declaration, any predicates of parent or progenitor subtypes apply.
- For a subtype created by a `subtype_indication`, the predicate of the subtype denoted by the `subtype_mark` applies.

Predicate checks are defined to be *enabled* or *disabled* for a given subtype as follows:

- If a subtype is declared by a `type_declaration` or `subtype_declaration` that includes a predicate specification, then:
  - if performing checks is required by the `Static_Predicate` assertion policy (see 14.4.2) and the declaration includes a `Static_Predicate` specification, then predicate checks are enabled for the subtype;
  - if performing checks is required by the `Dynamic_Predicate` assertion policy (see 14.4.2) and the declaration includes a `Dynamic_Predicate` specification, then predicate checks are enabled for the subtype;
  - otherwise, predicate checks are disabled for the subtype, regardless of whether predicate checking is enabled for any other subtypes mentioned in the declaration;

- If a subtype is defined by a type declaration that does not include a predicate specification, then predicate checks are enabled for the subtype if and only if any predicate checks are enabled for parent or progenitor subtypes;
- If a subtype is created by a `subtype_indication` other than in one of the previous cases, then predicate checks are enabled for the subtype if and only if predicate checks are enabled for the subtype denoted by the `subtype_mark`;
- Otherwise, predicate checks are disabled for the given subtype.

For a subtype with a directly-specified predicate aspect, the following additional language-defined aspect may be specified with an `aspect_specification` (see 16.1.1):

#### `Predicate_Failure`

This aspect shall be specified by an `expression`, which determines the action to be performed when a predicate check fails because a directly-specified predicate aspect of the subtype evaluates to False, as explained below.

#### *Name Resolution Rules*

The expected type for the `Predicate_Failure` expression is String.

#### *Legality Rules*

The `expression` of a `Static_Predicate` specification shall be *predicate-static*; that is, one of the following:

- a static expression;
- a membership test whose `tested_simple_expression` is the current instance, and whose `membership_choice_list` meets the requirements for a static membership test (see 7.9);
- a `case_expression` whose `selecting_expression` is the current instance, and whose `dependent_expressions` are static expressions;
- a call to a predefined equality or ordering operator, where one operand is the current instance, and the other is a static expression;
- a call to a predefined boolean operator **and**, **or**, **xor**, or **not**, where each operand is predicate-static;
- a short-circuit control form where both operands are predicate-static; or
- a parenthesized predicate-static expression.

A predicate shall not be specified for an incomplete subtype.

If a predicate applies to a subtype, then that predicate shall not mention any other subtype to which the same predicate applies.

An `index_subtype`, `discrete_range` of an `index_constraint` or `slice`, or a `discrete_subtype_definition` of a `constrained_array_definition`, `entry_declaration`, or `entry_index_specification` shall not denote a subtype to which predicate specifications apply.

The prefix of an `attribute_reference` whose `attribute_designator` is First, Last, or Range shall not denote a scalar subtype to which predicate specifications apply.

The `discrete_subtype_definition` of a `loop_parameter_specification` shall not denote a nonstatic subtype to which predicate specifications apply or any subtype to which `Dynamic_Predicate` specifications apply.

The `discrete_choice` of a `named_array_aggregate` shall not denote a nonstatic subtype to which predicate specifications apply.

In addition to the places where Legality Rules normally apply (see 15.3), these rules apply also in the private part of an instance of a generic unit.

If any of the above Legality Rules is violated in an instance of a generic unit, `Program_Error` is raised at the point of the violation.

To determine whether a value *satisfies the predicates* of a subtype *S*, the following tests are performed in the following order, until one of the tests fails, in which case the predicates are not satisfied and no further tests are performed, or all of the tests succeed, in which case the predicates are satisfied:

- the value is first tested to determine whether it satisfies any constraints or any null exclusion of *S*;
- then:
  - if *S* is a first subtype, the value is tested to determine whether it satisfies the predicates of the parent and progenitor subtypes (if any) of *S* (in an arbitrary order), after a (view) conversion of the value to the corresponding parent or progenitor type;
  - if *S* is defined by a `subtype_indication`, the value is tested to determine whether it satisfies the predicates of the subtype denoted by the `subtype_mark` of the `subtype_indication`;
- finally, if *S* is defined by a declaration to which one or more predicate specifications apply, the predicates are evaluated (in an arbitrary order) to test that all of them yield True for the given value.

If predicate checks are enabled for a given subtype, then:

On a subtype conversion, a check is performed that the operand satisfies the predicates of the target subtype, except for certain view conversions (see 7.6). In addition, after normal completion and leaving of a subprogram, for each **in out** or **out** parameter that is passed by reference, a check is performed that the value of the parameter satisfies the predicates of the subtype of the actual. For an object created by an `object_declaration` with no explicit initialization expression, or by an uninitialized allocator, if the types of any parts have specified `Default_Value` or `Default_Component_Value` aspects, or any subcomponents have `default_expressions`, a check is performed that the value of the created object satisfies the predicates of the nominal subtype.

If any of the predicate checks fail, `Assertion_Error` is raised, unless the subtype whose directly-specified predicate aspect evaluated to False also has a directly-specified `Predicate_Failure` aspect. In that case, the specified `Predicate_Failure` expression is evaluated; if the evaluation of the `Predicate_Failure` expression propagates an exception occurrence, then this occurrence is propagated for the failure of the predicate check; otherwise, `Assertion_Error` is raised, with an associated message string defined by the value of the `Predicate_Failure` expression. In the absence of such a `Predicate_Failure` aspect, an implementation-defined message string is associated with the `Assertion_Error` exception.

NOTE 1 A predicate specification does not cause a subtype to be considered constrained.

NOTE 2 A `Static_Predicate`, like a constraint, always remains True for all objects of the subtype, except in the case of uninitialized variables and other invalid values. A `Dynamic_Predicate`, on the other hand, is checked as specified above, but can become False at other times. For example, the predicate of a record subtype is not checked when a subcomponent is modified.

NOTE 3 No predicates apply to the base subtype of a scalar type; every value of a scalar type *T* is considered to satisfy the predicates of *T*'Base.

NOTE 4 `Predicate_Failure` expressions are never evaluated during the evaluation of a membership test (see 7.5.2) or Valid attribute (see 16.9.2).

NOTE 5 A `Predicate_Failure` expression can be a `raise_expression` (see 14.3).

EXAMPLE 1 Examples of predicates applied to scalar types:

```
subtype Basic_Letter is Character -- See A.3.2 for "basic letter".
with Static_Predicate => Basic_Letter in 'A'..'Z' | 'a'..'z' | 'Æ' |
                                     'æ' | 'Ð' | 'ð' | 'Þ' | 'þ' | 'ß';
```

```

subtype Even_Integer is Integer
  with Dynamic_Predicate => Even_Integer mod 2 = 0,
        Predicate_Failure => "Even_Integer must be a multiple of 2";

```

EXAMPLE 2 *Text\_IO* (see A.10.1) could have used predicates to describe some common exceptional conditions as follows:

```

with Ada.IO_Exceptions;
package Ada.Text_IO is

  type File_Type is limited private;

  subtype Open_File_Type is File_Type
    with Dynamic_Predicate => Is_Open (Open_File_Type),
          Predicate_Failure => raise Status_Error with "File not open";

  subtype Input_File_Type is Open_File_Type
    with Dynamic_Predicate => Mode (Input_File_Type) = In_File,
          Predicate_Failure => raise Mode_Error
            with "Cannot read file: " & Name (Input_File_Type);

  subtype Output_File_Type is Open_File_Type
    with Dynamic_Predicate => Mode (Output_File_Type) /= In_File,
          Predicate_Failure => raise Mode_Error
            with "Cannot write file: " & Name (Output_File_Type);

  ...

  function Mode (File : in Open_File_Type) return File_Mode;
  function Name (File : in Open_File_Type) return String;
  function Form (File : in Open_File_Type) return String;

  ...

  procedure Get (File : in Input_File_Type; Item : out Character);
  procedure Put (File : in Output_File_Type; Item : in Character);

  ...

  -- Similarly for all of the other input and output subprograms.

```

### 6.3 Objects and Named Numbers

Objects are created at run time and contain a value of a given type. An object can be created and initialized as part of elaborating a declaration, evaluating an allocator, aggregate, or function\_call, or passing a parameter by copy. Prior to reclaiming the storage for an object, it is finalized if necessary (see 10.6.1).

#### Static Semantics

All of the following are objects:

- the entity declared by an object\_declaration;
- a formal parameter of a subprogram, entry, or generic subprogram;
- a generic formal object;
- a loop parameter;
- the index parameter of an iterated\_component\_association;
- the chunk parameter of a chunk\_specification;
- a choice parameter of an exception\_handler;
- an entry index of an entry\_body;
- the result of dereferencing an access-to-object value (see 7.1);
- the return object of a function;
- the result of evaluating an aggregate;



- a value conversion or `qualified_expression` whose operand denotes an object;
- a component, slice, or view conversion of another object.

An object is either a *constant* object or a *variable* object. Similarly, a view of an object is either a *constant* or a *variable*. All views of a constant elementary object are constant. All views of a constant composite object are constant, except for parts that are of controlled or immutably limited types; variable views of those parts and their subcomponents may exist. In this sense, objects of controlled and immutably limited types are *inherently mutable*. A constant view of an object cannot be used to modify its value. The terms constant and variable by themselves refer to constant and variable views of objects.

A constant object is known to have no variable views if it does not have a part that is immutably limited, or of a controlled type, private type, or private extension.

The value of an object is *read* when the value of any part of the object is evaluated, or when the value of an enclosing object is evaluated. The value of a variable is *updated* when an assignment is performed to any part of the variable, or when an assignment is performed to an enclosing object.

Whether a view of an object is constant or variable is determined by the definition of the view. The following (and no others) represent variables:

- an object declared by an `object_declaration` without the reserved word **constant**;
- a formal parameter of mode **in out** or **out**;
- a generic formal object of mode **in out**;
- a non-discriminant component of a variable;
- a `slice` of a variable;
- a loop parameter that is specified to be a variable for a generalized loop (see 8.5.2);
- a view conversion of a variable;
- a dereference of an access-to-variable value;
- the return object declared by an `extended_return_statement` without the reserved word **constant**;
- the current instance of a type other than a protected type, if the current instance is an object and not a value (see 11.6);
- the current instance of a protected unit except within the body of a protected function of that protected unit, or within a function declared immediately within the body of the protected unit;
- an `attribute_reference` where the attribute is defined to denote a variable (for example, the `Storage_Pool` attribute – see 16.11).

At the place where a view of an object is defined, a *nominal subtype* is associated with the view. The *nominal type* of a view is the type of the nominal subtype of the view. The object's *actual subtype* (that is, its subtype) can be more restrictive than the nominal subtype of the view; it always is more restrictive if the nominal subtype is an *indefinite subtype*. A subtype is an indefinite subtype if it is an unconstrained array subtype, or if it has unknown discriminants or unconstrained discriminants without defaults (see 6.7); otherwise, the subtype is a *definite subtype* (all elementary subtypes are definite subtypes). A class-wide subtype is defined to have unknown discriminants, and is therefore an indefinite subtype. An indefinite subtype does not by itself provide enough information to create an object; an additional `constraint` or explicit initialization expression is necessary (see 6.3.1). A component cannot have an indefinite nominal subtype.

A view of a composite object is *known to be constrained* if:

- its nominal subtype is constrained and not an untagged partial view, and it is neither a value conversion nor a `qualified_expression`; or
- its nominal subtype is indefinite; or



- its type is immutably limited (see 10.5); or
- it is part of a stand-alone constant (including a generic formal object of mode **in**); or
- it is part of a formal parameter of mode **in**; or
- it is part of the object denoted by a `function_call` or aggregate; or
- it is a value conversion or `qualified_expression` where the operand denotes a view of a composite object that is known to be constrained; or
- it is part of a constant return object of an `extended_return_statement`; or
- it is a dereference of a pool-specific access type, and there is no ancestor of its type that has a constrained partial view.

For the purposes of determining within a generic body whether an object is known to be constrained:

- if a subtype is a descendant of an untagged generic formal private or derived type, and the subtype is not an unconstrained array subtype, it is not considered indefinite and is considered to have a constrained partial view;
- if a subtype is a descendant of a formal access type, it is not considered pool-specific.

A *named number* provides a name for a numeric value known at compile time. It is declared by a `number_declaration`.

NOTE 1 A constant cannot be the target of an assignment operation, nor be passed as an **in out** or **out** parameter, between its initialization and finalization, if any.

NOTE 2 The value of a constant object cannot be changed after its initialization, except in some cases where the object has a controlled or immutably limited part (see 10.5, 10.6, and 16.9.1).

NOTE 3 The nominal and actual subtypes of an elementary object are always the same. For a discriminated or array object, if the nominal subtype is constrained, then so is the actual subtype.

### 6.3.1 Object Declarations

An `object_declaration` declares a *stand-alone* object with a given nominal subtype and, optionally, an explicit initial value given by an initialization expression. For an array, access, task, or protected object, the `object_declaration` may include the definition of the (anonymous) type of the object.

#### Syntax

```
object_declaration ::=
  defining_identifier_list:
    [aliased] [constant] subtype_indication [:= expression]
    [aspect_specification];
| defining_identifier_list:
    [aliased] [constant] access_definition [:= expression]
    [aspect_specification];
| defining_identifier_list:
    [aliased] [constant] array_type_definition [:= expression]
    [aspect_specification];
| single_task_declaration
| single_protected_declaration

defining_identifier_list ::=
  defining_identifier {, defining_identifier}
```

#### Name Resolution Rules

For an `object_declaration` with an expression following the compound delimiter `:=`, the type expected for the expression is that of the object. This expression is called the *initialization expression*.

*Legality Rules*

An `object_declaration` without the reserved word **constant** declares a variable object. If it has a `subtype_indication` or an `array_type_definition` that defines an indefinite subtype, then there shall be an initialization expression.

*Static Semantics*

An `object_declaration` with the reserved word **constant** declares a constant object. If it has an initialization expression, then it is called a *full constant declaration*. Otherwise, it is called a *deferred constant declaration*. The rules for deferred constant declarations are given in 10.4. The rules for full constant declarations are given in this subclause.

Any declaration that includes a `defining_identifier_list` with more than one `defining_identifier` is equivalent to a series of declarations each containing one `defining_identifier` from the list, with the rest of the text of the declaration copied for each declaration in the series, in the same order as the list. The remainder of this document relies on this equivalence; explanations are given for declarations with a single `defining_identifier`.

The `subtype_indication`, `access_definition`, or full type definition of an `object_declaration` defines the nominal subtype of the object. The `object_declaration` declares an object of the type of the nominal subtype.

A component of an object is said to *require late initialization* if:

- it has an access discriminant value constrained by a per-object expression; or
- it has an initialization expression that includes a name denoting an access discriminant; or
- it has an initialization expression that includes a reference to the current instance of the type either by name or implicitly as the target object of a call.

*Dynamic Semantics*

If a composite object declared by an `object_declaration` has an unconstrained nominal subtype, then if this subtype is indefinite or the object is constant the actual subtype of this object is constrained. The constraint is determined by the bounds or discriminants (if any) of its initial value; the object is said to be *constrained by its initial value*. When not constrained by its initial value, the actual and nominal subtypes of the object are the same. If its actual subtype is constrained, the object is called a *constrained object*.

For an `object_declaration` without an initialization expression, any initial values for the object or its subcomponents are determined by the *implicit initial values* defined for its nominal subtype, as follows:

- The implicit initial value for an access subtype is the null value of the access type.
- The implicit initial value for a scalar subtype that has the `Default_Value` aspect specified is the value of that aspect converted to the nominal subtype (which can raise `Constraint_Error` — see 7.6);
- The implicit initial (and only) value for each discriminant of a constrained discriminated subtype is defined by the subtype.
- For a (definite) composite subtype, the implicit initial value of each component with a `default_expression` is obtained by evaluation of this expression and conversion to the component's nominal subtype (which can raise `Constraint_Error`), unless the component is a discriminant of a constrained subtype (the previous case), or is in an excluded variant (see 6.8.1). For each component that does not have a `default_expression`, if the composite subtype has the `Default_Component_Value` aspect specified, the implicit initial value is the value of that aspect converted to the component's nominal subtype; otherwise, any implicit initial values are those determined by the component's nominal subtype.
- For a protected or task subtype, there is an implicit component (an entry queue) corresponding to each entry, with its implicit initial value being an empty queue.

The elaboration of an `object_declaration` proceeds in the following sequence of steps:

1. The `subtype_indication`, `access_definition`, `array_type_definition`, `single_task_declaration`, or `single_protected_declaration` is first elaborated. This creates the nominal subtype (and the anonymous type in the last four cases).
2. If the `object_declaration` includes an initialization expression, the (explicit) initial value is obtained by evaluating the expression and converting it to the nominal subtype (which can raise `Constraint_Error` — see 7.6).
3. The object is created, and, if there is not an initialization expression, the object is *initialized by default*. When an object is initialized by default, any per-object constraints (see 6.8) are elaborated and any implicit initial values for the object or for its subcomponents are obtained as determined by the nominal subtype. Any initial values (whether explicit or implicit) are assigned to the object or to the corresponding subcomponents. As described in 8.2 and 10.6, `Initialize` and `Adjust` procedures can be called.

For the third step above, evaluations and assignments are performed in an arbitrary order subject to the following restrictions:

- Assignment to any part of the object is preceded by the evaluation of the value that is to be assigned.
- The evaluation of a `default_expression` that includes the name of a discriminant is preceded by the assignment to that discriminant.
- The evaluation of the `default_expression` for any component that depends on a discriminant is preceded by the assignment to that discriminant.
- The assignments to any components, including implicit components, not requiring late initialization precede the initial value evaluations for any components requiring late initialization; if two components both require late initialization, then assignments to parts of the component occurring earlier in the order of the component declarations precede the initial value evaluations of the component occurring later.

There is no implicit initial value defined for a scalar subtype unless the `Default_Value` aspect has been specified for the type. In the absence of an explicit initialization or the specification of the `Default_Value` aspect, a newly created scalar object can have a value that does not belong to its subtype (see 16.9.1 and H.1).

NOTE 1 Implicit initial values are not defined for an indefinite subtype, because if an object's nominal subtype is indefinite, an explicit initial value is required.

NOTE 2 As indicated above, a stand-alone object is an object declared by an `object_declaration`. Similar definitions apply to “stand-alone constant” and “stand-alone variable”. A subcomponent of an object is not a stand-alone object, nor is an object that is created by an allocator. An object declared by a `loop_parameter_specification`, `iterator_specification`, `iterated_component_association`, `chunk_specification`, `parameter_specification`, `entry_index_specification`, `choice_parameter_specification`, `extended_return_statement`, or a formal `object_declaration` of mode **in out** is not considered a stand-alone object.

NOTE 3 The type of a stand-alone object cannot be abstract (see 6.9.3).

EXAMPLE 1 *Example of a multiple object declaration:*

```
-- the multiple object declaration
John, Paul : not null Person_Name := new Person(Sex => M); -- see 6.10.1
-- is equivalent to the two single object declarations in the order given
John : not null Person_Name := new Person(Sex => M);
Paul : not null Person_Name := new Person(Sex => M);
```

**EXAMPLE 2** *Examples of variable declarations:*

```

Count, Sum : Integer;
Size       : Integer range 0 .. 10_000 := 0;
Sorted     : Boolean := False;
Color_Table : array(1 .. Max) of Color;
Option     : Bit_Vector(1 .. 10) := (others => True); -- see 6.6
Hello      : aliased String := "Hi, world.";
 $\theta$ ,  $\varphi$  : Float range  $-\pi$  ..  $+\pi$ ;

```

**EXAMPLE 3** *Examples of constant declarations:*

```

Limit      : constant Integer := 10_000;
Low_Limit  : constant Integer := Limit/10;
Tolerance  : constant Real := Dispersion(1.15);
A_String   : constant String := "A";
Hello_Msg  : constant access String := Hello'Access; -- see 6.10.2

```

**6.3.2 Number Declarations**

A `number_declaration` declares a named number.

*Syntax*

```

number_declaration ::=
    defining_identifier_list : constant := static_expression;

```

*Name Resolution Rules*

The `static_expression` given for a `number_declaration` is expected to be of any numeric type.

A name that denotes a `number_declaration` is interpreted as a value of a universal type, unless the expected type for the name is a non-numeric type with an `Integer_Literal` or `Real_Literal` aspect, in which case it is interpreted to be of its expected type.

*Legality Rules*

The `static_expression` given for a `number_declaration` shall be a static expression, as defined by 7.9.

*Static Semantics*

The named number denotes a value of type *universal\_integer* if the type of the `static_expression` is an integer type. The named number denotes a value of type *universal\_real* if the type of the `static_expression` is a real type.

The value denoted by the named number is the value of the `static_expression`, converted to the corresponding universal type.

*Dynamic Semantics*

The elaboration of a `number_declaration` has no effect.

**EXAMPLE** *Examples of number declarations:*

```

Two_Pi      : constant := 2.0*Ada.Numerics.Pi;    -- a real number (see A.5)
Max         : constant := 500;                    -- an integer number
Max_Line_Size : constant := Max/6;                -- the integer 83
Power_16    : constant := 2**16;                  -- the integer 65_536
One, Un, Eins : constant := 1;                    -- three different names
for 1

```

**6.4 Derived Types and Classes**

A `derived_type_definition` defines a *derived type* (and its first subtype) whose characteristics are *derived* from those of a parent type, and possibly from progenitor types.

A *class of types* is a set of types that is closed under derivation; that is, if the parent or a progenitor type of a derived type belongs to a class, then so does the derived type. By saying that a particular group of types forms a class, we are saying that all derivatives of a type in the set inherit the characteristics that define that set. The more general term *category of types* is used for a set of types whose defining characteristics are not necessarily inherited by derivatives; for example, limited, abstract, and interface are all categories of types, but not classes of types.

#### Syntax

```
derived_type_definition ::=
  [abstract] limited new parent_subtype_indication
  [[and interface_list] record_extension_part]
```

#### Legality Rules

The *parent\_subtype\_indication* defines the *parent subtype*; its type is the *parent type*. The *interface\_list* defines the progenitor types (see 6.9.4). A derived type has one parent type and zero or more progenitor types.

A type shall be completely defined (see 6.11.1) prior to being specified as the parent type in a *derived\_type\_definition* — the full\_type\_declarations for the parent type and any of its subcomponents have to precede the *derived\_type\_definition*.

If there is a *record\_extension\_part*, the derived type is called a *record extension* of the parent type. A *record\_extension\_part* shall be provided if and only if the parent type is a tagged type. An *interface\_list* shall be provided only if the parent type is a tagged type.

If the reserved word **limited** appears in a *derived\_type\_definition*, the parent type shall be a limited type. If the parent type is a tagged formal type, then in addition to the places where Legality Rules normally apply (see 15.3), this rule applies also in the private part of an instance of a generic unit.

#### Static Semantics

The first subtype of the derived type is unconstrained if a *known\_discriminant\_part* is provided in the declaration of the derived type, or if the parent subtype is unconstrained. Otherwise, the constraint of the first subtype *corresponds* to that of the parent subtype in the following sense: it is the same as that of the parent subtype except that for a range constraint (implicit or explicit), the value of each bound of its range is replaced by the corresponding value of the derived type.

The first subtype of the derived type excludes null (see 6.10) if and only if the parent subtype excludes null.

The *characteristics* and implicitly declared primitive subprograms of the derived type are defined as follows:

- If the parent type or a progenitor type belongs to a class of types, then the derived type also belongs to that class. The following sets of types, as well as any higher-level sets composed from them, are classes in this sense, and hence the characteristics defining these classes are inherited by derived types from their parent or progenitor types: signed integer, modular integer, ordinary fixed, decimal fixed, floating point, enumeration, boolean, character, access-to-constant, general access-to-variable, pool-specific access-to-variable, access-to-subprogram, array, string, non-array composite, unlimited, untagged record, tagged, task, protected, and synchronized tagged.
- If the parent type is an elementary type or an array type, then the set of possible values of the derived type is a copy of the set of possible values of the parent type. For a scalar type, the base range of the derived type is the same as that of the parent type.
- If the parent type is a composite type other than an array type, then the components, protected subprograms, and entries that are declared for the derived type are as follows:
  - The discriminants specified by a new *known\_discriminant\_part*, if there is one; otherwise, each discriminant of the parent type (implicitly declared in the same order with the

same specifications) — in the latter case, the discriminants are said to be *inherited*, or if unknown in the parent, are also unknown in the derived type;

- Each nondiscriminant component, entry, and protected subprogram of the parent type, implicitly declared in the same order with the same declarations; these components, entries, and protected subprograms are said to be *inherited*;
- Each component declared in a `record_extension_part`, if any.

Declarations of components, protected subprograms, and entries, whether implicit or explicit, occur immediately within the declarative region of the type, in the order indicated above, following the parent `subtype_indication`.

- For each predefined operator of the parent type, there is a corresponding predefined operator of the derived type.
- For each user-defined primitive subprogram (other than a user-defined equality operator — see below) of the parent type or of a progenitor type that already exists at the place of the `derived_type_definition`, there exists a corresponding *inherited* primitive subprogram of the derived type with the same defining name. Primitive user-defined equality operators of the parent type and any progenitor types are also inherited by the derived type, except when the derived type is a nonlimited record extension, and the inherited operator would have a profile that is type conformant with the profile of the corresponding predefined equality operator; in this case, the user-defined equality operator is not inherited, but is rather incorporated into the implementation of the predefined equality operator of the record extension (see 7.5.2).

The profile of an inherited subprogram (including an inherited enumeration literal) is obtained from the profile of the corresponding (user-defined) primitive subprogram of the parent or progenitor type, after systematic replacement of each subtype of its profile (see 9.1) that is of the parent or progenitor type, other than those subtypes found in the designated profile of an `access_definition`, with a *corresponding subtype* of the derived type. For a given subtype of the parent or progenitor type, the corresponding subtype of the derived type is defined as follows:

- If the declaration of the derived type has neither a `known_discriminant_part` nor a `record_extension_part`, then the corresponding subtype has a constraint that corresponds (as defined above for the first subtype of the derived type) to that of the given subtype.
- If the derived type is a record extension, then the corresponding subtype is the first subtype of the derived type.
- If the derived type has a new `known_discriminant_part` but is not a record extension, then the corresponding subtype is constrained to those values that when converted to the parent type belong to the given subtype (see 7.6).

The same formal parameters have `default_expressions` in the profile of the inherited subprogram. Any type mismatch due to the systematic replacement of the parent or progenitor type by the derived type is handled as part of the normal type conversion associated with parameter passing — see 9.4.1.

If a primitive subprogram of the parent or progenitor type is visible at the place of the `derived_type_definition`, then the corresponding inherited subprogram is implicitly declared immediately after the `derived_type_definition`. Otherwise, the inherited subprogram is implicitly declared later or not at all, as explained in 10.3.1.

A derived type can also be defined by a `private_extension_declaration` (see 10.3) or a `formal_derived_type_definition` (see 15.5.1). Such a derived type is a partial view of the corresponding full or actual type.

All numeric types are derived types, in that they are implicitly derived from a corresponding root numeric type (see 6.5.4 and 6.5.6).



The elaboration of a `derived_type_definition` creates the derived type and its first subtype, and consists of the elaboration of the `subtype_indication` and the `record_extension_part`, if any. If the `subtype_indication` depends on a discriminant, then only those expressions that do not depend on a discriminant are evaluated.

For the execution of a call on an inherited subprogram, a call on the corresponding primitive subprogram of the parent or progenitor type is performed; the normal conversion of each actual parameter to the subtype of the corresponding formal parameter (see 9.4.1) performs any necessary type conversion as well. If the result type of the inherited subprogram is the derived type, the result of calling the subprogram of the parent or progenitor is converted to the derived type, or in the case of a null extension, extended to the derived type using the equivalent of an `extension_aggregate` with the original result as the `ancestor_part` and **null record** as the `record_component_association_list`.

NOTE 1 Classes are closed under derivation — any class that contains a type also contains its derivatives. Operations available for a given class of types are available for the derived types in that class.

NOTE 2 Evaluating an inherited enumeration literal is equivalent to evaluating the corresponding enumeration literal of the parent type, and then converting the result to the derived type. This follows from their equivalence to parameterless functions.

NOTE 3 A generic subprogram is not a subprogram, and hence cannot be a primitive subprogram and cannot be inherited by a derived type. On the other hand, an instance of a generic subprogram can be a primitive subprogram, and hence can be inherited.

NOTE 4 If the parent type is an access type, then the parent and the derived type share the same storage pool; there is a **null** access value for the derived type and it is the implicit initial value for the type. See 6.10.

NOTE 5 If the parent type is a boolean type, the predefined relational operators of the derived type deliver a result of the predefined type Boolean (see 7.5.2). If the parent type is an integer type, the right operand of the predefined exponentiation operator is of the predefined type Integer (see 7.5.6).

NOTE 6 Any discriminants of the parent type are either all inherited, or completely replaced with a new set of discriminants.

NOTE 7 For an inherited subprogram, the subtype of a formal parameter of the derived type can be such that it has no value in common with the first subtype of the derived type.

NOTE 8 If the reserved word **abstract** is given in the declaration of a type, the type is abstract (see 6.9.3).

NOTE 9 An interface type that has a progenitor type “is derived from” that type. A `derived_type_definition`, however, never defines an interface type.

NOTE 10 It is illegal for the parent type of a `derived_type_definition` to be a synchronized tagged type.

EXAMPLE *Examples of derived type declarations:*

```
type Local_Coordinate is new Coordinate;    -- two different types
type Midweek is new Day range Tue .. Thu;  -- see 6.5.1
type Counter is new Positive;              -- same range as Positive

type Special_Key is new Key_Manager.Key;    -- see 10.3.1
-- the inherited subprograms have the following specifications:
--     procedure Get_Key(K : out Special_Key);
--     function "<"(X,Y : Special_Key) return Boolean;
```

### 6.4.1 Derivation Classes

In addition to the various language-defined classes of types, types can be grouped into *derivation classes*.

A derived type is *derived from* its parent type *directly*; it is derived *indirectly* from any type from which its parent type is derived. A derived type, interface type, type extension, task type, protected type, or formal derived type is also derived from every ancestor of each of its progenitor types, if any. The



derivation class of types for a type *T* (also called the class *rooted at T*) is the set consisting of *T* (the *root type* of the class) and all types derived from *T* (directly or indirectly) plus any associated universal or class-wide types (defined below).

Every type is one of a *specific* type, a *class-wide* type, or a *universal* type. A specific type is one defined by a `type_declaration`, a `formal_type_declaration`, or a full type definition embedded in another construct. Class-wide and universal types are implicitly defined, to act as representatives for an entire class of types, as follows:

#### Class-wide types

Class-wide types are defined for (and belong to) each derivation class rooted at a tagged type (see 6.9). Given a subtype *S* of a tagged type *T*, *S'Class* is the `subtype_mark` for a corresponding subtype of the tagged class-wide type *T'Class*. Such types are called “class-wide” because when a formal parameter is defined to be of a class-wide type *T'Class*, an actual parameter of any type in the derivation class rooted at *T* is acceptable (see 11.6).

The set of values for a class-wide type *T'Class* is the discriminated union of the set of values of each specific type in the derivation class rooted at *T* (the tag acts as the implicit discriminant — see 6.9). Class-wide types have no primitive subprograms of their own. However, as explained in 6.9.2, operands of a class-wide type *T'Class* can be used as part of a dispatching call on a primitive subprogram of the type *T*. The only components (including discriminants) of *T'Class* that are visible are those of *T*. If *S* is a first subtype, then *S'Class* is a first subtype.

#### Universal types

Universal types are defined for (and belong to) the integer, real, fixed point, and access classes, and are referred to in this document as respectively, *universal\_integer*, *universal\_real*, *universal\_fixed*, and *universal\_access*. These are analogous to class-wide types for these language-defined elementary classes. As with class-wide types, if a formal parameter is of a universal type, then an actual parameter of any type in the corresponding class is acceptable. In addition, a value of a universal type (including an integer or real `numeric_literal`, or the literal **null**) is “universal” in that it is acceptable where some particular type in the class is expected (see 11.6).

The set of values of a universal type is the undiscriminated union of the set of values possible for any definable type in the associated class. Like class-wide types, universal types have no primitive subprograms of their own. However, their “universality” allows them to be used as operands with the primitive subprograms of any type in the corresponding class.

The integer and real numeric classes each have a specific root type in addition to their universal type, named respectively *root\_integer* and *root\_real*.

A class-wide or universal type is said to *cover* all of the types in its class. In addition, *universal\_integer* covers a type that has a specified `Integer_Literal` aspect, while *universal\_real* covers a type that has a specified `Real_Literal` aspect (see 7.2.1). A specific type covers only itself.

A specific type *T2* is defined to be a *descendant* of a type *T1* if *T2* is the same as *T1*, or if *T2* is derived (directly or indirectly) from *T1*. A class-wide type *T2'Class* is defined to be a descendant of type *T1* if *T2* is a descendant of *T1*. Similarly, the numeric universal types are defined to be descendants of the root types of their classes. If a type *T2* is a descendant of a type *T1*, then *T1* is called an *ancestor* of *T2*. An *ultimate ancestor* of a type is an ancestor of that type that is not itself a descendant of any other type. Every untagged type has a unique ultimate ancestor.

An inherited component (including an inherited discriminant) of a derived type is inherited *from* a given ancestor of the type if the corresponding component was inherited by each derived type in the chain of derivations going back to the given ancestor.

NOTE Because operands of a universal type are acceptable to the predefined operators of any type in their class, ambiguity can result. For *universal\_integer* and *universal\_real*, this potential ambiguity is resolved by giving a preference (see 11.6) to the predefined operators of the corresponding root types (*root\_integer* and *root\_real*, respectively). Hence, in an apparently ambiguous expression like

$$1 + 4 < 7$$

where each of the literals is of type *universal\_integer*, the predefined operators of *root\_integer* will be preferred over those of other specific integer types, thereby resolving the ambiguity.

## 6.5 Scalar Types

*Scalar* types comprise enumeration types, integer types, and real types. Enumeration types and integer types are called *discrete* types; each value of a discrete type has a *position number* which is an integer value. Integer types and real types are called *numeric* types. All scalar types are ordered, that is, all relational operators are predefined for their values.

### Syntax

```
range_constraint ::= range range
range ::= range_attribute_reference
| simple_expression..simple_expression
```

A *range* has a *lower bound* and an *upper bound* and specifies a subset of the values of some scalar type (the *type of the range*). A range with lower bound L and upper bound R is described by “L .. R”. If R is less than L, then the range is a *null range*, and specifies an empty set of values. Otherwise, the range specifies the values of the type from the lower bound to the upper bound, inclusive. A value *belongs* to a range if it is of the type of the range, and is in the subset of values specified by the range. A value *satisfies* a range constraint if it belongs to the associated range. One range is *included* in another if all values that belong to the first range also belong to the second.

### Name Resolution Rules

For a *subtype\_indication* containing a *range\_constraint*, either directly or as part of some other *scalar\_constraint*, the type of the range shall resolve to that of the type determined by the *subtype\_mark* of the *subtype\_indication*. For a range of a given type, the *simple\_expressions* of the range (likewise, the *simple\_expressions* of the equivalent range for a *range\_attribute\_reference*) are expected to be of the type of the range.

### Static Semantics

The *base range* of a scalar type is the range of finite values of the type that can be represented in every unconstrained object of the type; it is also the range supported at a minimum for intermediate values during the evaluation of expressions involving predefined operators of the type.

A constrained scalar subtype is one to which a range constraint applies. The *range* of a constrained scalar subtype is the range associated with the range constraint of the subtype. The *range* of an unconstrained scalar subtype is the base range of its type.

### Dynamic Semantics

A range is *compatible* with a scalar subtype if and only if it is either a null range or each bound of the range belongs to the range of the subtype. A *range\_constraint* is *compatible* with a scalar subtype if and only if its range is compatible with the subtype.

The elaboration of a *range\_constraint* consists of the evaluation of the *range*. The evaluation of a *range* determines a lower bound and an upper bound. If *simple\_expressions* are given to specify bounds, the evaluation of the *range* evaluates these *simple\_expressions* in an arbitrary order, and converts them to the type of the range. If a *range\_attribute\_reference* is given, the evaluation of the *range* consists of the evaluation of the *range\_attribute\_reference*.

### Attributes

For every scalar subtype S, the following attributes are defined:

**S'First**      S'First denotes the lower bound of the range of S. The value of this attribute is of the type of S.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| S'Last            | S'Last denotes the upper bound of the range of S. The value of this attribute is of the type of S.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| S'Range           | S'Range is equivalent to the <code>range S'First .. S'Last</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| S'Base            | S'Base denotes an unconstrained subtype of the type of S. This unconstrained subtype is called the <i>base subtype</i> of the type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| S'Min             | <p>S'Min denotes a function with the following specification:</p> <pre>function S'Min(Left, Right : S'Base) return S'Base</pre> <p>The function returns the lesser of the values of the two parameters.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| S'Max             | <p>S'Max denotes a function with the following specification:</p> <pre>function S'Max(Left, Right : S'Base) return S'Base</pre> <p>The function returns the greater of the values of the two parameters.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| S'Succ            | <p>S'Succ denotes a function with the following specification:</p> <pre>function S'Succ(Arg : S'Base) return S'Base</pre> <p>For an enumeration type, the function returns the value whose position number is one more than that of the value of <i>Arg</i>; <i>Constraint_Error</i> is raised if there is no such value of the type. For an integer type, the function returns the result of adding one to the value of <i>Arg</i>. For a fixed point type, the function returns the result of adding <i>small</i> to the value of <i>Arg</i>. For a floating point type, the function returns the machine number (as defined in 6.5.7) immediately above the value of <i>Arg</i>; <i>Constraint_Error</i> is raised if there is no such machine number.</p>               |
| S'Pred            | <p>S'Pred denotes a function with the following specification:</p> <pre>function S'Pred(Arg : S'Base) return S'Base</pre> <p>For an enumeration type, the function returns the value whose position number is one less than that of the value of <i>Arg</i>; <i>Constraint_Error</i> is raised if there is no such value of the type. For an integer type, the function returns the result of subtracting one from the value of <i>Arg</i>. For a fixed point type, the function returns the result of subtracting <i>small</i> from the value of <i>Arg</i>. For a floating point type, the function returns the machine number (as defined in 6.5.7) immediately below the value of <i>Arg</i>; <i>Constraint_Error</i> is raised if there is no such machine number.</p> |
| S'Wide_Wide_Width | S'Wide_Wide_Width denotes the maximum length of a <i>Wide_Wide_String</i> returned by <i>S'Wide_Wide_Image</i> over all values of the subtype S, assuming a default implementation of <i>S'Put_Image</i> . It denotes zero for a subtype that has a null range. Its type is <i>universal_integer</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| S'Wide_Width      | S'Wide_Width denotes the maximum length of a <i>Wide_String</i> returned by <i>S'Wide_Image</i> over all values of the subtype S, assuming a default implementation of <i>S'Put_Image</i> . It denotes zero for a subtype that has a null range. Its type is <i>universal_integer</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| S'Width           | S'Width denotes the maximum length of a <i>String</i> returned by <i>S'Image</i> over all values of the subtype S, assuming a default implementation of <i>S'Put_Image</i> . It denotes zero for a subtype that has a null range. Its type is <i>universal_integer</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| S'Wide_Wide_Value | <p>S'Wide_Wide_Value denotes a function with the following specification:</p> <pre>function S'Wide_Wide_Value(Arg : Wide_Wide_String) return S'Base</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

This function returns a value given an image of the value as a `Wide_Wide_String`, ignoring any leading or trailing spaces.

For the evaluation of a call on `S'Wide_Wide_Value` for an enumeration subtype `S`, if the sequence of characters of the parameter (ignoring leading and trailing spaces) has the syntax of an enumeration literal and if it corresponds to a literal of the type of `S` (or corresponds to the result of `S'Wide_Wide_Image` for a nongraphic character of the type), the result is the corresponding enumeration value; otherwise, `Constraint_Error` is raised.

For the evaluation of a call on `S'Wide_Wide_Value` for an integer subtype `S`, if the sequence of characters of the parameter (ignoring leading and trailing spaces) has the syntax of an integer literal, with an optional leading sign character (plus or minus for a signed type; only plus for a modular type), and the corresponding numeric value belongs to the base range of the type of `S`, then that value is the result; otherwise, `Constraint_Error` is raised.

For the evaluation of a call on `S'Wide_Wide_Value` for a real subtype `S`, if the sequence of characters of the parameter (ignoring leading and trailing spaces) has the syntax of one of the following:

- `numeric_literal`
- `numeral.[exponent]`
- `.numeral[exponent]`
- `base#based_numeral#[exponent]`
- `base#.based_numeral#[exponent]`

with an optional leading sign character (plus or minus), and if the corresponding numeric value belongs to the base range of the type of `S`, then that value is the result; otherwise, `Constraint_Error` is raised. The sign of a zero value is preserved (positive if none has been specified) if `S'Signed_Zeros` is `True`.

#### `S'Wide_Value`

`S'Wide_Value` denotes a function with the following specification:

```
function S'Wide_Value (Arg : Wide_String)
return S'Base
```

This function returns a value given an image of the value as a `Wide_String`, ignoring any leading or trailing spaces.

For the evaluation of a call on `S'Wide_Value` for an enumeration subtype `S`, if the sequence of characters of the parameter (ignoring leading and trailing spaces) has the syntax of an enumeration literal and if it corresponds to a literal of the type of `S` (or corresponds to the result of `S'Wide_Image` for a value of the type, assuming a default implementation of `S'Put_Image`), the result is the corresponding enumeration value; otherwise, `Constraint_Error` is raised. For a numeric subtype `S`, the evaluation of a call on `S'Wide_Value` with `Arg` of type `Wide_String` is equivalent to a call on `S'Wide_Wide_Value` for a corresponding `Arg` of type `Wide_Wide_String`.

#### `S'Value`

`S'Value` denotes a function with the following specification:

```
function S'Value (Arg : String)
return S'Base
```

This function returns a value given an image of the value as a `String`, ignoring any leading or trailing spaces.

For the evaluation of a call on `S'Value` for an enumeration subtype `S`, if the sequence of characters of the parameter (ignoring leading and trailing spaces) has the syntax of an enumeration literal and if it corresponds to a literal of the type of `S` (or corresponds to the result of `S'Image` for a value of the type, assuming a default implementation of `S'Put_Image`), the result is the corresponding enumeration value; otherwise, `Constraint_Error` is raised. For a numeric subtype `S`, the evaluation of a call on `S'Value` with `Arg` of type `String` is

equivalent to a call on `S'Wide_Wide_Value` for a corresponding *Arg* of type `Wide_Wide_String`.

#### Implementation Permissions

An implementation may extend the `Wide_Wide_Value`, `Wide_Value`, `Value`, `Wide_Wide_Image`, `Wide_Image`, and `Image` attributes of a floating point type to support special values such as infinities and NaNs.

An implementation may extend the `Wide_Wide_Value`, `Wide_Value`, and `Value` attributes of a character type to accept strings of the form “Hex\_ hhhhhhhh” (ignoring case) for any character (not just the ones for which `Wide_Wide_Image` would produce that form — see 6.5.2), as well as three-character strings of the form “X”, where *X* is any character, including nongraphic characters.

#### Static Semantics

For a scalar type, the following language-defined representation aspect may be specified with an `aspect_specification` (see 16.1.1):

##### Default\_Value

This aspect shall be specified by a static expression, and that expression shall be explicit, even if the aspect has a boolean type. `Default_Value` shall be specified only on a `full_type_declaration`.

If a derived type inherits a boolean `Default_Value` aspect, the aspect may be specified to have any value for the derived type. If a derived type *T* does not inherit a `Default_Value` aspect, it shall not specify such an aspect if it inherits a primitive subprogram that has a parameter of type *T* of mode **out**.

#### Name Resolution Rules

The expected type for the `expression` specified for the `Default_Value` aspect is the type defined by the `full_type_declaration` on which it appears.

NOTE 1 The evaluation of `S'First` or `S'Last` never raises an exception. If a scalar subtype *S* has a nonnull range, `S'First` and `S'Last` belong to this range. These values can, for example, always be assigned to a variable of subtype *S*.

NOTE 2 For a subtype of a scalar type, the result delivered by the attributes `Succ`, `Pred`, and `Value` can be outside to the subtype; similarly, the actual parameters of the attributes `Succ`, `Pred`, and `Image` can also be outside the subtype.

NOTE 3 For any value *V* (including any nongraphic character) of an enumeration subtype *S* without a specified `Put_Image` (see 7.10), `S'Value(S'Image(V))` equals *V*, as do `S'Wide_Value(S'Wide_Image(V))` and `S'Wide_Wide_Value(S'Wide_Wide_Image(V))`. None of these expressions ever raise `Constraint_Error`.

##### EXAMPLE 1 Examples of ranges:

```
-10 .. 10
X .. X+1
0.0 .. 2.0*Pi
Red .. Green      -- see 6.5.1
1 .. 0             -- a null range
Table'Range       -- a range attribute reference (see 6.6)
```

##### EXAMPLE 2 Examples of range constraints:

```
range -999.0 .. +999.0
range S'First+1 .. S'Last-1
```

## 6.5.1 Enumeration Types

An `enumeration_type_definition` defines an enumeration type.

#### Syntax

```
enumeration_type_definition ::=
  (enumeration_literal_specification { enumeration_literal_specification })
enumeration_literal_specification ::=
```

```

defining_identifier | defining_character_literal
defining_character_literal ::= character_literal

```

#### Legality Rules

The `defining_identifiers` in upper case and the `defining_character_literals` listed in an `enumeration_type_definition` shall be distinct.

#### Static Semantics

Each `enumeration_literal_specification` is the explicit declaration of the corresponding *enumeration literal*: it declares a parameterless function, whose defining name is the `defining_identifier` or `defining_character_literal`, and whose result subtype is the base subtype of the enumeration type.

Each enumeration literal corresponds to a distinct value of the enumeration type, and to a distinct position number. The position number of the value of the first listed enumeration literal is zero; the position number of the value of each subsequent enumeration literal is one more than that of its predecessor in the list.

The predefined order relations between values of the enumeration type follow the order of corresponding position numbers.

If the same `defining_identifier` or `defining_character_literal` is specified in more than one `enumeration_type_definition`, the corresponding enumeration literals are said to be *overloaded*. At any place where an overloaded enumeration literal occurs in the text of a program, the type of the enumeration literal has to be determinable from the context (see 11.6).

#### Dynamic Semantics

The elaboration of an `enumeration_type_definition` creates the enumeration type and its first subtype, which is constrained to the base range of the type.

When called, the parameterless function associated with an enumeration literal returns the corresponding value of the enumeration type.

NOTE If an enumeration literal occurs in a context that does not otherwise suffice to determine the type of the literal, then qualification by the name of the enumeration type is one way to resolve the ambiguity (see 7.7).

EXAMPLE *Examples of enumeration types and subtypes:*

```

type Day is (Mon, Tue, Wed, Thu, Fri, Sat, Sun);
type Month_Name is (January, February, March, April, May, June, July,
August, September, October, November, December);
type Suit is (Clubs, Diamonds, Hearts, Spades);
type Gender is (M, F);
type Level is (Low, Medium, Urgent);
type Color is (White, Red, Yellow, Green, Blue, Brown, Black);
type Light is (Red, Amber, Green); -- Red and Green are overloaded
type Hexa is ('A', 'B', 'C', 'D', 'E', 'F');
type Mixed is ('A', 'B', '*', B, None, '?', '%');

subtype Weekday is Day range Mon .. Fri;
subtype Major is Suit range Hearts .. Spades;
subtype Rainbow is Color range Red .. Blue;
-- the Color Red, not the Light

```

## 6.5.2 Character Types

#### Static Semantics

An enumeration type is said to be a *character type* if at least one of its enumeration literals is a `character_literal`.



The predefined type `Character` is a character type whose values correspond to the 256 code points of Row 00 (also known as Latin-1) of the ISO/IEC 10646:2020 Basic Multilingual Plane (BMP). Each of the graphic characters of Row 00 of the BMP has a corresponding `character_literal` in `Character`. Each of the nongraphic characters of Row 00 has a corresponding language-defined name, which is not usable as an enumeration literal, but which is usable with the attributes `Image`, `Wide_Image`, `Wide_Wide_Image`, `Value`, `Wide_Value`, and `Wide_Wide_Value`; these names are given in the definition of type `Character` in A.1, but are set in *italics*.

The predefined type `Wide_Character` is a character type whose values correspond to the 65536 code points of the ISO/IEC 10646:2020 Basic Multilingual Plane (BMP). Each of the graphic characters of the BMP has a corresponding `character_literal` in `Wide_Character`. The first 256 values of `Wide_Character` have the same `character_literal` or language-defined name as defined for `Character`. Each of the `graphic_characters` has a corresponding `character_literal`.

The predefined type `Wide_Wide_Character` is a character type whose values correspond to the 2147483648 code points of the ISO/IEC 10646:2020 character set. Each of the `graphic_characters` has a corresponding `character_literal` in `Wide_Wide_Character`. The first 65536 values of `Wide_Wide_Character` have the same `character_literal` or language-defined name as defined for `Wide_Character`.

The characters whose code point is larger than 16#FF# and which are not `graphic_characters` have language-defined names which are formed by appending to the string "Hex\_" the representation of their code point in hexadecimal as eight extended digits. As with other language-defined names, these names are usable only with the attributes `(Wide_)Wide_Image` and `(Wide_)Wide_Value`; they are not usable as enumeration literals.

NOTE 1 The language-defined library package `Characters.Latin_1` (see A.3.3) includes the declaration of constants denoting control characters, lower case characters, and special characters of the predefined type `Character`.

NOTE 2 A conventional character set such as *EBCDIC* can be declared as a character type; the internal codes of the characters can be specified by an `enumeration_representation_clause` as explained in 16.4.

EXAMPLE *Example of a character type:*

```
type Roman_Digit is ('I', 'V', 'X', 'L', 'C', 'D', 'M');
```

### 6.5.3 Boolean Types

#### Static Semantics

There is a predefined enumeration type named `Boolean`, declared in the visible part of package `Standard`. It has the two enumeration literals `False` and `True` ordered with the relation `False < True`. Any descendant of the predefined type `Boolean` is called a *boolean* type.

### 6.5.4 Integer Types

An `integer_type_definition` defines an integer type; it defines either a *signed* integer type, or a *modular* integer type. The base range of a signed integer type includes at least the values of the specified range. A modular type is an integer type with all arithmetic modulo a specified positive *modulus*; such a type corresponds to an unsigned type with wrap-around semantics.

#### Syntax

```
integer_type_definition ::=
    signed_integer_type_definition | modular_type_definition
signed_integer_type_definition ::=
    range static_simple_expression .. static_simple_expression
modular_type_definition ::= mod static_expression
```



## Name Resolution Rules

Each `simple_expression` in a `signed_integer_type_definition` is expected to be of any integer type; they can be of different integer types. The expression in a `modular_type_definition` is likewise expected to be of any integer type.

## Legality Rules

The `simple_expressions` of a `signed_integer_type_definition` shall be static, and their values shall be in the range `System.Min_Int .. System.Max_Int`.

The expression of a `modular_type_definition` shall be static, and its value (the *modulus*) shall be positive, and shall be no greater than `System.Max_Binary_Modulus` if a power of 2, or no greater than `System.Max_Nonbinary_Modulus` if not.

## Static Semantics

The set of values for a signed integer type is the (infinite) set of mathematical integers, though only values of the base range of the type are fully supported for run-time operations. The set of values for a modular integer type are the values from 0 to one less than the modulus, inclusive.

A `signed_integer_type_definition` defines an integer type whose base range includes at least the values of the `simple_expressions` and is symmetric about zero, excepting possibly an extra negative value. A `signed_integer_type_definition` also defines a constrained first subtype of the type, with a range whose bounds are given by the values of the `simple_expressions`, converted to the type being defined.

A `modular_type_definition` defines a modular type whose base range is from zero to one less than the given modulus. A `modular_type_definition` also defines a constrained first subtype of the type with a range that is the same as the base range of the type.

There is a predefined signed integer subtype named `Integer`, declared in the visible part of package `Standard`. It is constrained to the base range of its type.

`Integer` has two predefined subtypes, declared in the visible part of package `Standard`:

```
subtype Natural is Integer range 0 .. Integer'Last;  
subtype Positive is Integer range 1 .. Integer'Last;
```

A type defined by an `integer_type_definition` is implicitly derived from *root\_integer*, an anonymous predefined (specific) integer type, whose base range is `System.Min_Int .. System.Max_Int`. However, the base range of the new type is not inherited from *root\_integer*, but is instead determined by the range or modulus specified by the `integer_type_definition`. Integer literals are all of the type *universal\_integer*, the universal type (see 6.4.1) for the class rooted at *root\_integer*, allowing their use with the operations of any integer type.

The *position number* of an integer value is equal to the value.

For every modular subtype *S*, the following attributes are defined:

*S*'Mod      *S*'Mod denotes a function with the following specification:

```
function S'Mod (Arg : universal_integer)  
  return S'Base
```

This function returns *Arg* **mod** *S*'Modulus, as a value of the type of *S*.

*S*'Modulus   *S*'Modulus yields the modulus of the type of *S*, as a value of the type *universal\_integer*.

## Dynamic Semantics

The elaboration of an `integer_type_definition` creates the integer type and its first subtype.

For a modular type, if the result of the execution of a predefined operator (see 7.5) is outside the base range of the type, the result is reduced modulo the modulus of the type to a value that is within the base range of the type.

For a signed integer type, the exception `Constraint_Error` is raised by the execution of an operation that cannot deliver the correct result because it is outside the base range of the type. For any integer type, `Constraint_Error` is raised by the operators `"/`, `"rem`", and `"mod`" if the right operand is zero.

#### Implementation Requirements

In an implementation, the range of `Integer` shall include the range  $-2^{15}+1 .. +2^{15}-1$ .

If `Long_Integer` is predefined for an implementation, then its range shall include the range  $-2^{31}+1 .. +2^{31}-1$ .

`System.Max_Binary_Modulus` shall be at least  $2^{16}$ .

#### Implementation Permissions

For the execution of a predefined operation of a signed integer type, it is optional to raise `Constraint_Error` if the result is outside the base range of the type, so long as the correct result is produced.

An implementation may provide additional predefined signed integer types, declared in the visible part of Standard, whose first subtypes have names of the form `Short_Integer`, `Long_Integer`, `Short_Short_Integer`, `Long_Long_Integer`, etc. Different predefined integer types are allowed to have the same base range. However, the range of `Integer` should be no wider than that of `Long_Integer`. Similarly, the range of `Short_Integer` (if provided) should be no wider than `Integer`. Corresponding recommendations apply to any other predefined integer types. An implementation may support base ranges for which there is no corresponding named integer type. The range of each first subtype should be the base range of its type.

An implementation may provide *nonstandard integer types*, descendants of *root\_integer* that are declared outside of the specification of package Standard, which may have different characteristics than a type defined by an *integer\_type\_definition*. For example, a nonstandard integer type can have an asymmetric base range or it can be disallowed as an array or loop index (a very long integer). Any type descended from a nonstandard integer type is also nonstandard. An implementation may place arbitrary restrictions on the use of such types; it is implementation defined whether operators that are predefined for "any integer type" are defined for a particular nonstandard integer type. In any case, such types are not permitted as *explicit\_generic\_actual\_parameters* for formal scalar types — see 15.5.2.

For a one's complement machine, the high bound of the base range of a modular type whose modulus is one less than a power of 2 may be equal to the modulus, rather than one less than the modulus. It is implementation defined for which powers of 2, if any, this permission is exercised.

For a one's complement machine, implementations may support nonbinary modulus values greater than `System.Max_Nonbinary_Modulus`. It is implementation defined which specific values greater than `System.Max_Nonbinary_Modulus`, if any, are supported.

#### Implementation Advice

An implementation should support `Long_Integer` in addition to `Integer` if the target machine supports 32-bit (or longer) arithmetic. No other named integer subtypes are recommended for package Standard. Instead, appropriate named integer subtypes should be provided in the library package Interfaces (see B.2).

An implementation for a two's complement machine should support modular types with a binary modulus up to `System.Max_Int*2+2`. An implementation should support a nonbinary modulus up to `Integer'Last`.

NOTE 1 Integer literals are of the anonymous predefined integer type *universal\_integer*. Other integer types have no literals. However, the overload resolution rules (see 11.6) allow expressions of the type *universal\_integer* whenever an integer type is expected.

NOTE 2 The same arithmetic operators are predefined for all signed integer types defined by a *signed\_integer\_type\_definition* (see 7.5). For modular types, these same operators are predefined, plus

bit-wise logical operators (**and**, **or**, **xor**, and **not**). In addition, for the unsigned types declared in the language-defined package Interfaces (see B.2), functions are defined that provide bit-wise shifting and rotating.

NOTE 3 Modular types match a `generic_formal_parameter_declaration` of the form "**type T is mod** <>"; signed integer types match "**type T is range** <>"; (see 15.5.2).

EXAMPLE *Examples of integer types and subtypes:*

```
type Page_Num is range 1 .. 2_000;
type Line_Size is range 1 .. Max_Line_Size;

subtype Small_Int is Integer range -10 .. 10;
subtype Column_Ptr is Line_Size range 1 .. 10;
subtype Buffer_Size is Integer range 0 .. Max;

type Byte is mod 256; -- an unsigned byte
type Hash_Index is mod 97; -- modulus is prime
```

### 6.5.5 Operations of Discrete Types

#### *Static Semantics*

For every discrete subtype S, the following attributes are defined:

S'Pos S'Pos denotes a function with the following specification:

```
function S'Pos (Arg : S'Base)
return universal_integer
```

This function returns the position number of the value of Arg, as a value of type *universal\_integer*.

S'Val S'Val denotes a function with the following specification:

```
function S'Val (Arg : universal_integer)
return S'Base
```

This function returns a value of the type of S whose position number equals the value of Arg. For the evaluation of a call on S'Val, if there is no value in the base range of its type with the given position number, Constraint\_Error is raised.

For every static discrete subtype S for which there exists at least one value belonging to S that satisfies the predicates of S, the following attributes are defined:

S'First\_Valid

S'First\_Valid denotes the smallest value that belongs to S and satisfies the predicates of S. The value of this attribute is of the type of S.

S'Last\_Valid

S'Last\_Valid denotes the largest value that belongs to S and satisfies the predicates of S. The value of this attribute is of the type of S.

First\_Valid and Last\_Valid attribute references are always static expressions. Any explicit predicate of S can only have been specified by a Static\_Predicate aspect.

#### *Implementation Advice*

For the evaluation of a call on S'Pos for an enumeration subtype, if the value of the operand does not correspond to the internal code for any enumeration literal of its type (perhaps due to an uninitialized variable), then the implementation should raise Program\_Error. This is particularly important for enumeration types with noncontiguous internal codes specified by an enumeration\_representation\_clause.

NOTE 1 Indexing and loop iteration use values of discrete types.

NOTE 2 The predefined operations of a discrete type include the assignment operation, qualification, the membership tests, and the relational operators; for a boolean type they include the short-circuit control forms and the logical operators; for an integer type they include type conversion to and from other numeric types, as well as the binary and unary adding operators – and +, the multiplying operators, the unary operator **abs**, and the

exponentiation operator. The assignment operation is described in 8.2. The other predefined operations are described in Clause 7.

NOTE 3 As for all types, objects of a discrete type have Size and Address attributes (see 16.3).

NOTE 4 For a subtype of a discrete type, the result delivered by the attribute Val can be outside the subtype; similarly, the actual parameter of the attribute Pos can also be outside the subtype. The following relations are satisfied (in the absence of an exception) by these attributes:

$$\begin{aligned} S'Val(S'Pos(X)) &= X \\ S'Pos(S'Val(N)) &= N \end{aligned}$$

EXAMPLE *Examples of attributes of discrete subtypes:*

```
-- For the types and subtypes declared in 6.5.1 the following hold:
-- Color'First   = White,   Color'Last   = Black
-- Rainbow'First = Red,     Rainbow'Last = Blue
-- Color'Succ(Blue) = Rainbow'Succ(Blue) = Brown
-- Color'Pos(Blue)  = Rainbow'Pos(Blue)  = 4
-- Color'Val(0)     = Rainbow'Val(0)     = White
```

### 6.5.6 Real Types

Real types provide approximations to the real numbers, with relative bounds on errors for floating point types, and with absolute bounds for fixed point types.

#### Syntax

```
real_type_definition ::=
  floating_point_definition | fixed_point_definition
```

#### Static Semantics

A type defined by a *real\_type\_definition* is implicitly derived from *root\_real*, an anonymous predefined (specific) real type. Hence, all real types, whether floating point or fixed point, are in the derivation class rooted at *root\_real*.

Real literals are all of the type *universal\_real*, the universal type (see 6.4.1) for the class rooted at *root\_real*, allowing their use with the operations of any real type. Certain multiplying operators have a result type of *universal\_fixed* (see 7.5.5), the universal type for the class of fixed point types, allowing the result of the multiplication or division to be used where any specific fixed point type is expected.

#### Dynamic Semantics

The elaboration of a *real\_type\_definition* consists of the elaboration of the *floating\_point\_definition* or the *fixed\_point\_definition*.

#### Implementation Requirements

An implementation shall perform the run-time evaluation of a use of a predefined operator of *root\_real* with an accuracy at least as great as that of any floating point type definable by a *floating\_point\_definition*.

#### Implementation Permissions

For the execution of a predefined operation of a real type, it is optional to raise *Constraint\_Error* if the result is outside the base range of the type, so long as the correct result is produced, or the *Machine\_Overflows* attribute of the type is *False* (see G.2.1).

An implementation may provide *nonstandard real types*, descendants of *root\_real* that are declared outside of the specification of package *Standard*, which may have different characteristics than a type defined by a *real\_type\_definition*. For example, a nonstandard real type can have an asymmetric or unsigned base range, or its predefined operations can wrap around or “saturate” rather than overflow (modular or saturating arithmetic), or it can have a different accuracy model than is standard (see G.2.1). Any type descended from a nonstandard real type is also nonstandard. An implementation

may place arbitrary restrictions on the use of such types; it is implementation defined whether operators that are predefined for “any real type” are defined for a particular nonstandard real type. In any case, such types are not permitted as `explicit_generic_actual_parameters` for formal scalar types — see 15.5.2.

NOTE As stated, real literals are of the anonymous predefined real type *universal\_real*. Other real types have no literals. However, the overload resolution rules (see 11.6) allow expressions of the type *universal\_real* whenever a real type is expected.

### 6.5.7 Floating Point Types

For floating point types, the error bound is specified as a relative precision by giving the required minimum number of significant decimal digits.

#### Syntax

```
floating_point_definition ::=
  digits static_expression [real_range_specification]
real_range_specification ::=
  range static_simple_expression .. static_simple_expression
```

#### Name Resolution Rules

The *requested decimal precision*, which is the minimum number of significant decimal digits required for the floating point type, is specified by the value of the `expression` given after the reserved word **digits**. This expression is expected to be of any integer type.

Each `simple_expression` of a `real_range_specification` is expected to be of any real type; the types can be different.

#### Legality Rules

The requested decimal precision shall be specified by a static expression whose value is positive and no greater than `System.Max_Base_Digits`. Each `simple_expression` of a `real_range_specification` shall also be static. If the `real_range_specification` is omitted, the requested decimal precision shall be no greater than `System.Max_Digits`.

A `floating_point_definition` is illegal if the implementation does not support a floating point type that satisfies the requested decimal precision and range.

#### Static Semantics

The set of values for a floating point type is the (infinite) set of rational numbers. The *machine numbers* of a floating point type are the values of the type that can be represented exactly in every unconstrained variable of the type. The base range (see 6.5) of a floating point type is symmetric around zero, except that it can include some extra negative values in some implementations.

The *base decimal precision* of a floating point type is the number of decimal digits of precision representable in objects of the type. The *safe range* of a floating point type is that part of its base range for which the accuracy corresponding to the base decimal precision is preserved by all predefined operations.

A `floating_point_definition` defines a floating point type whose base decimal precision is no less than the requested decimal precision. If a `real_range_specification` is given, the safe range of the floating point type (and hence, also its base range) includes at least the values of the simple expressions given in the `real_range_specification`. If a `real_range_specification` is not given, the safe (and base) range of the type includes at least the values of the range  $-10.0^{*(4*D)} .. +10.0^{*(4*D)}$  where  $D$  is the requested decimal precision. The safe range can include other values as well. The attributes `Safe_First` and `Safe_Last` give the actual bounds of the safe range.

A `floating_point_definition` also defines a first subtype of the type. If a `real_range_specification` is given, then the subtype is constrained to a range whose bounds are



given by a conversion of the values of the `simple_expressions` of the `real_range_specification` to the type being defined. Otherwise, the subtype is unconstrained.

There is a predefined, unconstrained, floating point subtype named `Float`, declared in the visible part of package `Standard`.

#### Dynamic Semantics

The elaboration of a `floating_point_definition` creates the floating point type and its first subtype.

#### Implementation Requirements

In an implementation that supports floating point types with 6 or more digits of precision, the requested decimal precision for `Float` shall be at least 6.

If `Long_Float` is predefined for an implementation, then its requested decimal precision shall be at least 11.

#### Implementation Permissions

An implementation is allowed to provide additional predefined floating point types, declared in the visible part of `Standard`, whose (unconstrained) first subtypes have names of the form `Short_Float`, `Long_Float`, `Short_Short_Float`, `Long_Long_Float`, etc. Different predefined floating point types are allowed to have the same base decimal precision. However, the precision of `Float` should be no greater than that of `Long_Float`. Similarly, the precision of `Short_Float` (if provided) should be no greater than `Float`. Corresponding recommendations apply to any other predefined floating point types. An implementation may support base decimal precisions for which there is no corresponding named floating point type.

#### Implementation Advice

An implementation should support `Long_Float` in addition to `Float` if the target machine supports 11 or more digits of precision. No other named floating point subtypes are recommended for package `Standard`. Instead, appropriate named floating point subtypes should be provided in the library package `Interfaces` (see B.2).

NOTE If a floating point subtype is unconstrained, then assignments to variables of the subtype involve only `Overflow_Checks`, never `Range_Checks`.

EXAMPLE *Examples of floating point types and subtypes:*

```
type Coefficient is digits 10 range -1.0 .. 1.0;
type Real is digits 8;
type Mass is digits 7 range 0.0 .. 1.0E35;
subtype Probability is Real range 0.0 .. 1.0;
-- a subtype with a smaller range
```

## 6.5.8 Operations of Floating Point Types

#### Static Semantics

The following attribute is defined for every floating point subtype `S`:

**S'Digits**     `S'Digits` denotes the requested decimal precision for the subtype `S`. The value of this attribute is of the type *universal\_integer*. The requested decimal precision of the base subtype of a floating point type `T` is defined to be the largest value of  $d$  for which  $\text{ceiling}(d * \log(10) / \log(T'Machine\_Radix)) + g \leq T'Model\_Mantissa$  where  $g$  is 0 if `Machine_Radix` is a positive power of 10 and 1 otherwise.

NOTE 1 The predefined operations of a floating point type include the assignment operation, qualification, the membership tests, and explicit conversion to and from other numeric types. They also include the relational operators and the following predefined arithmetic operators: the binary and unary adding operators `-` and `+`, certain multiplying operators, the unary operator `abs`, and the exponentiation operator.



NOTE 2 As for all types, objects of a floating point type have Size and Address attributes (see 16.3). Other attributes of floating point types are defined in A.5.3.

### 6.5.9 Fixed Point Types

A fixed point type is either an ordinary fixed point type, or a decimal fixed point type. The error bound of a fixed point type is specified as an absolute value, called the *delta* of the fixed point type.

#### Syntax

```
fixed_point_definition ::=
  ordinary_fixed_point_definition | decimal_fixed_point_definition
ordinary_fixed_point_definition ::=
  delta static_expression real_range_specification
decimal_fixed_point_definition ::=
  delta static_expression digits static_expression [real_range_specification]
digits_constraint ::=
  digits static_simple_expression [range_constraint]
```

#### Name Resolution Rules

For a type defined by a `fixed_point_definition`, the *delta* of the type is specified by the value of the expression given after the reserved word **delta**; this expression is expected to be of any real type. For a type defined by a `decimal_fixed_point_definition` (a *decimal* fixed point type), the number of significant decimal digits for its first subtype (the *digits* of the first subtype) is specified by the expression given after the reserved word **digits**; this expression is expected to be of any integer type.

The `simple_expression` of a `digits_constraint` is expected to be of any integer type.

#### Legality Rules

In a `fixed_point_definition` or `digits_constraint`, the expressions given after the reserved words **delta** and **digits** shall be static; their values shall be positive.

The set of values of a fixed point type comprise the integral multiples of a number called the *small* of the type. The *machine numbers* of a fixed point type are the values of the type that can be represented exactly in every unconstrained variable of the type. For a type defined by an `ordinary_fixed_point_definition` (an *ordinary* fixed point type), the *small* may be specified by an `attribute_definition_clause` (see 16.3); if so specified, it shall be no greater than the *delta* of the type. If not specified, the *small* of an ordinary fixed point type is an implementation-defined power of two less than or equal to the *delta*.

For a decimal fixed point type, the *small* equals the *delta*; the *delta* shall be a power of 10. If a `real_range_specification` is given, both bounds of the range shall be in the range  $-(10^{**digits}-1)*delta .. +(10^{**digits}-1)*delta$ .

A `fixed_point_definition` is illegal if the implementation does not support a fixed point type with the given *small* and specified range or *digits*.

For a `subtype_indication` with a `digits_constraint`, the `subtype_mark` shall denote a decimal fixed point subtype.

#### Static Semantics

The base range (see 6.5) of a fixed point type is symmetric around zero, except possibly for an extra negative value in some implementations.

An `ordinary_fixed_point_definition` defines an ordinary fixed point type whose base range includes at least all multiples of *small* that are between the bounds specified in the `real_range_specification`. The base range of the type does not necessarily include the specified

bounds themselves. An `ordinary_fixed_point_definition` also defines a constrained first subtype of the type, with each bound of its range given by the closer to zero of:

- the value of the conversion to the fixed point type of the corresponding expression of the `real_range_specification`;
- the corresponding bound of the base range.

A `decimal_fixed_point_definition` defines a decimal fixed point type whose base range includes at least the range  $-(10^{**digits-1}) * delta .. +(10^{**digits-1}) * delta$ . A `decimal_fixed_point_definition` also defines a constrained first subtype of the type. If a `real_range_specification` is given, the bounds of the first subtype are given by a conversion of the values of the expressions of the `real_range_specification`. Otherwise, the range of the first subtype is  $-(10^{**digits-1}) * delta .. +(10^{**digits-1}) * delta$ .

#### Dynamic Semantics

The elaboration of a `fixed_point_definition` creates the fixed point type and its first subtype.

For a `digits_constraint` on a decimal fixed point subtype with a given `delta`, if it does not have a `range_constraint`, then it specifies an implicit range  $-(10^{**D-1}) * delta .. +(10^{**D-1}) * delta$ , where `D` is the value of the `simple_expression`. A `digits_constraint` is *compatible* with a decimal fixed point subtype if the value of the `simple_expression` is no greater than the `digits` of the subtype, and if it specifies (explicitly or implicitly) a range that is compatible with the subtype.

The elaboration of a `digits_constraint` consists of the elaboration of the `range_constraint`, if any. If a `range_constraint` is given, a check is made that the bounds of the range are both in the range  $-(10^{**D-1}) * delta .. +(10^{**D-1}) * delta$ , where `D` is the value of the (static) `simple_expression` given after the reserved word **digits**. If this check fails, `Constraint_Error` is raised.

#### Implementation Requirements

The implementation shall support at least 24 bits of precision (including the sign bit) for fixed point types.

#### Implementation Permissions

Implementations are permitted to support only *smalls* that are a power of two. In particular, all decimal fixed point type declarations can be disallowed. Note however that conformance with the Information Systems Annex requires support for decimal *smalls*, and decimal fixed point type declarations with `digits` up to at least 18.

NOTE The specified bounds themselves can be outside the base range of an ordinary fixed point type so that the range specification can be given in a natural way, such as:

```
type Fraction is delta 2.0**(-15) range -1.0 .. 1.0;
```

With 2's complement hardware, such a type would typically have a signed 16-bit representation, using 1 bit for the sign and 15 bits for fraction, resulting in a base range of  $-1.0 .. 1.0 - 2.0^{**(-15)}$ .

EXAMPLE *Examples of fixed point types and subtypes:*

```
type Volt is delta 0.125 range 0.0 .. 255.0;
-- A pure fraction which requires all the available
-- space in a word can be declared as the type Fraction:
type Fraction is delta System.Fine_Delta range -1.0 .. 1.0;
-- Fraction'Last = 1.0 - System.Fine_Delta

type Money is delta 0.01 digits 15; -- decimal fixed point
subtype Salary is Money digits 10;
-- Money'Last = 10.0**13 - 0.01, Salary'Last = 10.0**8 - 0.01
```

### 6.5.10 Operations of Fixed Point Types

#### Static Semantics

The following attributes are defined for every fixed point subtype S:

- S'Small** S'Small denotes the *small* of the type of S. The value of this attribute is of the type *universal\_real*. Small may be specified for nonderived ordinary fixed point types via an *attribute\_definition\_clause* (see 16.3); the expression of such a clause shall be static and positive.
- S'Delta** S'Delta denotes the *delta* of the fixed point subtype S. The value of this attribute is of the type *universal\_real*.
- S'Fore** S'Fore yields the minimum number of characters needed before the decimal point for the decimal representation of any value of the subtype S, assuming that the representation does not include an exponent, but includes a one-character prefix that is either a minus sign or a space. (This minimum number does not include superfluous zeros or underlines, and is at least 2.) The value of this attribute is of the type *universal\_integer*.
- S'Aft** S'Aft yields the number of decimal digits needed after the decimal point to accommodate the *delta* of the subtype S, unless the *delta* of the subtype S is greater than 0.1, in which case the attribute yields the value one. (S'Aft is the smallest positive integer N for which  $(10^{**}N) * S'Delta$  is greater than or equal to one.) The value of this attribute is of the type *universal\_integer*.

The following additional attributes are defined for every decimal fixed point subtype S:

- S'Digits** S'Digits denotes the *digits* of the decimal fixed point subtype S, which corresponds to the number of decimal digits that are representable in objects of the subtype. The value of this attribute is of the type *universal\_integer*. Its value is determined as follows:
- For a first subtype or a subtype defined by a *subtype\_indication* with a *digits\_constraint*, the digits is the value of the expression given after the reserved word **digits**;
  - For a subtype defined by a *subtype\_indication* without a *digits\_constraint*, the digits of the subtype is the same as that of the subtype denoted by the *subtype\_mark* in the *subtype\_indication*;
  - The digits of a base subtype is the largest integer *D* such that the range  $-(10^{**}D-1)*delta .. +(10^{**}D-1)*delta$  is included in the base range of the type.
- S'Scale** S'Scale denotes the *scale* of the subtype S, defined as the value N such that  $S'Delta = 10.0^{**}(-N)$ . The scale indicates the position of the point relative to the rightmost significant digits of values of subtype S. The value of this attribute is of the type *universal\_integer*.
- S'Round** S'Round denotes a function with the following specification:
- ```

function S'Round(X : universal_real)
return S'Base
  
```

The function returns the value obtained by rounding *X* (away from 0, if *X* is midway between two values of the type of S).

**NOTE 1** All subtypes of a fixed point type will have the same value for the Delta attribute, in the absence of *delta\_constraints* (see I.3).

**NOTE 2** S'Scale is not always the same as S'Aft for a decimal subtype; for example, if S'Delta = 1.0 then S'Aft is 1 while S'Scale is 0.

**NOTE 3** The predefined operations of a fixed point type include the assignment operation, qualification, the membership tests, and explicit conversion to and from other numeric types. They also include the relational operators and the following predefined arithmetic operators: the binary and unary adding operators – and +, multiplying operators, and the unary operator **abs**.

NOTE 4 As for all types, objects of a fixed point type have Size and Address attributes (see 16.3). Other attributes of fixed point types are defined in A.5.4.

## 6.6 Array Types

An *array* object is a composite object consisting of components which all have the same subtype. The name for a component of an array uses one or more index values belonging to specified discrete types. The value of an array object is a composite value consisting of the values of the components.

### Syntax

```
array_type_definition ::=
  unconstrained_array_definition | constrained_array_definition
unconstrained_array_definition ::=
  array(index_subtype_definition { index_subtype_definition })
  of component_definition
index_subtype_definition ::= subtype_mark range <>
constrained_array_definition ::=
  array (discrete_subtype_definition { discrete_subtype_definition })
  of component_definition
discrete_subtype_definition ::= discrete_subtype_indication | range
component_definition ::=
  [aliased] subtype_indication
  | [aliased] access_definition
```

### Name Resolution Rules

For a *discrete\_subtype\_definition* that is a range, the range shall resolve to be of some specific discrete type; which discrete type shall be determined without using any context other than the bounds of the range itself (plus the preference for *root\_integer* — see 11.6).

### Legality Rules

Each *index\_subtype\_definition* or *discrete\_subtype\_definition* in an *array\_type\_definition* defines an *index subtype*; its type (the *index type*) shall be discrete.

The subtype defined by the *subtype\_indication* of a *component\_definition* (the *component subtype*) shall be a definite subtype.

### Static Semantics

An array is characterized by the number of indices (the *dimensionality* of the array), the type and position of each index, the lower and upper bounds for each index, and the subtype of the components. The order of the indices is significant.

A one-dimensional array has a distinct component for each possible index value. A multidimensional array has a distinct component for each possible sequence of index values that can be formed by selecting one value for each index position (in the given order). The possible values for a given index are all the values between the lower and upper bounds, inclusive; this range of values is called the *index range*. The *bounds* of an array are the bounds of its index ranges. The *length* of a dimension of an array is the number of values of the index range of the dimension (zero for a null range). The *length* of a one-dimensional array is the length of its only dimension.

An *array\_type\_definition* defines an array type and its first subtype. For each object of this array type, the number of indices, the type and position of each index, and the subtype of the components are as in the type definition; the values of the lower and upper bounds for each index belong to the corresponding index subtype of its type, except for null arrays (see 6.6.1).

An *unconstrained\_array\_definition* defines an array type with an unconstrained first subtype. Each *index\_subtype\_definition* defines the corresponding index subtype to be the

subtype denoted by the `subtype_mark`. The compound delimiter `<>` (called a *box*) of an `index_subtype_definition` stands for an undefined range (different objects of the type can have different bounds).

A `constrained_array_definition` defines an array type with a constrained first subtype. Each `discrete_subtype_definition` defines the corresponding index subtype, as well as the corresponding index range for the constrained first subtype. The *constraint* of the first subtype consists of the bounds of the index ranges.

The discrete subtype defined by a `discrete_subtype_definition` is either that defined by the `subtype_indication`, or a subtype determined by the range as follows:

- If the type of the range resolves to *root\_integer*, then the `discrete_subtype_definition` defines a subtype of the predefined type Integer with bounds given by a conversion to Integer of the bounds of the range;
- Otherwise, the `discrete_subtype_definition` defines a subtype of the type of the range, with the bounds given by the range.

The component definition of an `array_type_definition` defines the nominal subtype of the components. If the reserved word **aliased** appears in the component definition, then each component of the array is aliased (see 6.10).

#### Dynamic Semantics

The elaboration of an `array_type_definition` creates the array type and its first subtype, and consists of the elaboration of any `discrete_subtype_definitions` and the `component_definition`.

The elaboration of a `discrete_subtype_definition` that does not contain any per-object expressions creates the discrete subtype, and consists of the elaboration of the `subtype_indication` or the evaluation of the range. The elaboration of a `discrete_subtype_definition` that contains one or more per-object expressions is defined in 6.8. The elaboration of a `component_definition` in an `array_type_definition` consists of the elaboration of the `subtype_indication` or `access_definition`. The elaboration of any `discrete_subtype_definitions` and the elaboration of the `component_definition` are performed in an arbitrary order.

#### Static Semantics

For an array type with a scalar component type, the following language-defined representation aspect may be specified with an `aspect_specification` (see 16.1.1):

##### Default\_Component\_Value

This aspect shall be specified by a static expression, and that expression shall be explicit, even if the aspect has a boolean type. `Default_Component_Value` shall be specified only on a `full_type_declaration`.

If a derived type inherits a boolean `Default_Component_Value` aspect, the aspect may be specified to have any value for the derived type.

#### Name Resolution Rules

The expected type for the expression specified for the `Default_Component_Value` aspect is the component type of the array type defined by the `full_type_declaration` on which it appears.

NOTE 1 All components of an array have the same subtype. In particular, for an array of components that are one-dimensional arrays, this means that all components have the same bounds and hence the same length.

NOTE 2 Each elaboration of an `array_type_definition` creates a distinct array type. A consequence of this is that each object whose `object_declaration` contains an `array_type_definition` is of its own unique type.



EXAMPLE 1 *Examples of type declarations with unconstrained array definitions:*

```
type Vector      is array(Integer range <>) of Real;
type Matrix      is array(Integer range <>, Integer range <>) of Real;
type Bit_Vector  is array(Integer range <>) of Boolean;
type Roman       is array(Positive range <>) of Roman_Digit; -- see 6.5.2
```

EXAMPLE 2 *Examples of type declarations with constrained array definitions:*

```
type Table       is array(1 .. 10) of Integer;
type Schedule    is array(Day) of Boolean;
type Line        is array(1 .. Max_Line_Size) of Character;
```

EXAMPLE 3 *Examples of object declarations with array type definitions:*

```
Grid      : array(1 .. 80, 1 .. 100) of Boolean;
Mix       : array(Color range Red .. Green) of Boolean;
Msg_Table : constant array(Error_Code) of access constant String :=
    (Too_Big => new String("Result too big"), Too_Small => ...);
Page      : array(Positive range <>) of Line := -- an array of arrays
    (1 | 50 => Line'(1 | Line'Last => '+', others => '-'), -- see 7.3.3
     2 .. 49 => Line'(1 | Line'Last => '|', others => ' '));
-- Page is constrained by its initial value to (1..50)
```

### 6.6.1 Index Constraints and Discrete Ranges

An `index_constraint` determines the range of possible values for every index of an array subtype, and thereby the corresponding array bounds.

*Syntax*

```
index_constraint ::= (discrete_range{, discrete_range})
discrete_range ::= discrete_subtype_indication | range
```

*Name Resolution Rules*

The type of a `discrete_range` is the type of the subtype defined by the `subtype_indication`, or the type of the range. For an `index_constraint`, each `discrete_range` shall resolve to be of the type of the corresponding index.

*Legality Rules*

An `index_constraint` shall appear only in a `subtype_indication` whose `subtype_mark` denotes either an unconstrained array subtype, or an unconstrained access subtype whose designated subtype is an unconstrained array subtype; in either case, the `index_constraint` shall provide a `discrete_range` for each index of the array type.

*Static Semantics*

A `discrete_range` defines a range whose bounds are given by the range, or by the range of the subtype defined by the `subtype_indication`.

*Dynamic Semantics*

An `index_constraint` is *compatible* with an unconstrained array subtype if and only if the index range defined by each `discrete_range` is compatible (see 6.5) with the corresponding index subtype. If any of the `discrete_ranges` defines a null range, any array thus constrained is a *null array*, having no components. An array value *satisfies* an `index_constraint` if at each index position the array value and the `index_constraint` have the same index bounds.

The elaboration of an `index_constraint` consists of the evaluation of the `discrete_range(s)`, in an arbitrary order. The evaluation of a `discrete_range` consists of the elaboration of the `subtype_indication` or the evaluation of the range.



NOTE 1 The elaboration of a `subtype_indication` consisting of a `subtype_mark` followed by an `index_constraint` checks the compatibility of the `index_constraint` with the `subtype_mark` (see 6.2.2).

NOTE 2 Even if an array value does not satisfy the index constraint of an array subtype, `Constraint_Error` is not raised on conversion to the array subtype, so long as the length of each dimension of the array value and the array subtype match. See 7.6.

EXAMPLE 1 *Examples of array declarations including an index constraint:*

```
Board      : Matrix(1 .. 8, 1 .. 8);  -- see 6.6
Rectangle  : Matrix(1 .. 20, 1 .. 30);
Inverse    : Matrix(1 .. N, 1 .. N);  -- N can be nonstatic
Filter     : Bit_Vector(0 .. 31);     -- see 6.6
```

EXAMPLE 2 *Example of array declaration with a constrained array subtype:*

```
My_Schedule : Schedule;  -- all arrays of type Schedule have the same bounds
```

EXAMPLE 3 *Example of record type with a component that is an array:*

```
type Var_Line(Length : Natural) is
  record
    Image : String(1 .. Length);
  end record;

Null_Line : Var_Line(0);  -- Null_Line.Image is a null array
```

## 6.6.2 Operations of Array Types

### Legality Rules

The argument `N` used in the `attribute_designators` for the `N`-th dimension of an array shall be a static expression of some integer type. The value of `N` shall be positive (nonzero) and no greater than the dimensionality of the array.

### Static Semantics

The following attributes are defined for a prefix `A` that is of an array type (after any implicit dereference), or denotes a constrained array subtype:

**A'First**     `A'First` denotes the lower bound of the first index range; its type is the corresponding index type.

**A'First(N)**     `A'First(N)` denotes the lower bound of the `N`-th index range; its type is the corresponding index type.

**A'Last**     `A'Last` denotes the upper bound of the first index range; its type is the corresponding index type.

**A'Last(N)**     `A'Last(N)` denotes the upper bound of the `N`-th index range; its type is the corresponding index type.

**A'Range**     `A'Range` is equivalent to the range `A'First .. A'Last`, except that the prefix `A` is only evaluated once.

**A'Range(N)**     `A'Range(N)` is equivalent to the range `A'First(N) .. A'Last(N)`, except that the prefix `A` is only evaluated once.

**A'Length**     `A'Length` denotes the number of values of the first index range (zero for a null range); its type is *universal\_integer*.

**A'Length(N)**     `A'Length(N)` denotes the number of values of the `N`-th index range (zero for a null range); its type is *universal\_integer*.

## Implementation Advice

An implementation should normally represent multidimensional arrays in row-major order, consistent with the notation used for multidimensional array aggregates (see 7.3.3). However, if convention Fortran is specified for a multidimensional array type, then column-major order should be used instead (see B.5).

NOTE 1 The attribute\_references A'First and A'First(1) denote the same value. A similar relation exists for the attribute\_references A'Last, A'Range, and A'Length. The following relation is satisfied (except for a null array) by the above attributes if the index type is an integer type:

$$A'Length(N) = A'Last(N) - A'First(N) + 1$$

NOTE 2 An array type is limited if its component type is limited (see 10.5).

NOTE 3 The predefined operations of an array type include the membership tests, qualification, and explicit conversion. If the array type is not limited, they also include assignment and the predefined equality operators. For a one-dimensional array type, they include the predefined concatenation operators (if nonlimited) and, if the component type is discrete, the predefined relational operators; if the component type is boolean, the predefined logical operators are also included.

NOTE 4 A component of an array can be named with an indexed\_component. A value of an array type can be specified with an array\_aggregate. For a one-dimensional array type, a slice of the array can be named; also, string literals are defined if the component type is a character type.

EXAMPLE Examples (using arrays declared in the examples of 6.6.1):

```
-- Filter'First      = 0      Filter'Last      = 31      Filter'Length = 32
-- Rectangle'Last(1) = 20     Rectangle'Last(2) = 30
```

### 6.6.3 String Types

## Static Semantics

A one-dimensional array type whose component type is a character type is called a *string type*.

There are three predefined string types, String, Wide\_String, and Wide\_Wide\_String, each indexed by values of the predefined subtype Positive; these are declared in the visible part of package Standard:

```
subtype Positive is Integer range 1 .. Integer'Last;
type String is array(Positive range <>) of Character;
type Wide_String is array(Positive range <>) of Wide_Character;
type Wide_Wide_String is array(Positive range <>) of Wide_Wide_Character;
```

NOTE String literals (see 5.6 and 7.2) are defined for all string types. The concatenation operator & is predefined for string types, as for all nonlimited one-dimensional array types. The ordering operators <, <=, >, and >= are predefined for string types, as for all one-dimensional discrete array types; these ordering operators correspond to lexicographic order (see 7.5.2).

EXAMPLE Examples of string objects:

```
Stars      : String(1 .. 120) := (1 .. 120 => '*' );
Question   : constant String  := "How many characters?";
-- Question'First = 1, Question'Last
= 20
-- Question'Length = 20 (the number
of characters)
Ask_Twice  : String := Question & Question; -- constrained to (1..40)
Ninety_Six : constant Roman  := "XCVI"; -- see 6.5.2 and 6.6
```

### 6.7 Discriminants

A composite type (other than an array or interface type) can have discriminants, which parameterize the type. A known\_discriminant\_part specifies the discriminants of a composite type. A discriminant of an object is a component of the object, and is either of a discrete type or an access type. An unknown\_discriminant\_part in the declaration of a view of a type specifies that the

discriminants of the type are unknown for the given view; all subtypes of such a view are indefinite subtypes.

#### Syntax

```
discriminant_part ::= unknown_discriminant_part | known_discriminant_part
unknown_discriminant_part ::= (<>)
known_discriminant_part ::=
  (discriminant_specification { discriminant_specification })
discriminant_specification ::=
  defining_identifier_list: [null_exclusion] subtype_mark
  [:= default_expression]
  [aspect_specification]
| defining_identifier_list: access_definition
  [:= default_expression]
  [aspect_specification]
default_expression ::= expression
```

#### Name Resolution Rules

The expected type for the `default_expression` of a `discriminant_specification` is that of the corresponding discriminant.

#### Legality Rules

A `discriminant_part` is only permitted in a declaration for a composite type that is not an array or interface type (this includes generic formal types). A type declared with a `known_discriminant_part` is called a *discriminated* type, as is a type that inherits (known) discriminants.

The subtype of a discriminant may be defined by an optional `null_exclusion` and a `subtype_mark`, in which case the `subtype_mark` shall denote a discrete or access subtype, or it may be defined by an `access_definition`. A discriminant that is defined by an `access_definition` is called an *access discriminant* and is of an anonymous access type.

Default expressions shall be provided either for all or for none of the discriminants of a `known_discriminant_part`. No default expressions are permitted in a `known_discriminant_part` in a declaration of a nonlimited tagged type or a generic formal type.

A `discriminant_specification` for an access discriminant may have a `default_expression` only in the declaration for an immutably limited type (see 10.5). In addition to the places where Legality Rules normally apply (see 15.3), this rule applies also in the private part of an instance of a generic unit.

For a type defined by a `derived_type_definition`, if a `known_discriminant_part` is provided in its declaration, then:

- The parent subtype shall be constrained;
- If the parent type is not a tagged type, then each discriminant of the derived type shall be used in the constraint defining the parent subtype;
- If a discriminant is used in the constraint defining the parent subtype, the subtype of the discriminant shall be statically compatible (see 7.9.1) with the subtype of the corresponding parent discriminant.

#### Static Semantics

A `discriminant_specification` declares a discriminant; the `subtype_mark` denotes its subtype unless it is an access discriminant, in which case the discriminant's subtype is the anonymous access-to-variable subtype defined by the `access_definition`.

For a type defined by a `derived_type_definition`, each discriminant of the parent type is either inherited, constrained to equal some new discriminant of the derived type, or constrained to the value of an expression. When inherited or constrained to equal some new discriminant, the parent discriminant and the discriminant of the derived type are said to *correspond*. Two discriminants also correspond if there is some common discriminant to which they both correspond. A discriminant corresponds to itself as well. If a discriminant of a parent type is constrained to a specific value by a `derived_type_definition`, then that discriminant is said to be *specified* by that `derived_type_definition`.

A constraint that appears within the definition of a discriminated type *depends on a discriminant* of the type if it names the discriminant as a bound or discriminant value. A `component_definition` depends on a discriminant if its `constraint` depends on the discriminant, or on a discriminant that corresponds to it.

A component *depends on a discriminant* if:

- Its `component_definition` depends on the discriminant; or
- It is declared in a `variant_part` that is governed by the discriminant; or
- It is a component inherited as part of a `derived_type_definition`, and the constraint of the `parent_subtype_indication` depends on the discriminant; or
- It is a subcomponent of a component that depends on the discriminant.

Each value of a discriminated type includes a value for each component of the type that does not depend on a discriminant; this includes the discriminants themselves. The values of discriminants determine which other component values are present in the value of the discriminated type.

A type declared with a `known_discriminant_part` is said to have *known discriminants*; its first subtype is unconstrained. A type declared with an `unknown_discriminant_part` is said to have *unknown discriminants*. A type declared without a `discriminant_part` has no discriminants, unless it is a derived type; if derived, such a type has the same sort of discriminants (known, unknown, or none) as its parent (or ancestor) type. A tagged class-wide type also has unknown discriminants. Any subtype of a type with unknown discriminants is an unconstrained and indefinite subtype (see 6.2 and 6.3).

#### Dynamic Semantics

For an access discriminant, its `access_definition` is elaborated when the value of the access discriminant is defined: by evaluation of its `default_expression`, by elaboration of a `discriminant_constraint`, or by an assignment that initializes the enclosing object.

NOTE 1 If a discriminated type has `default_expressions` for its discriminants, then unconstrained variables of the type are permitted, and the values of the discriminants can be changed by an assignment to such a variable. If defaults are not provided for the discriminants, then all variables of the type are constrained, either by explicit constraint or by their initial value; the values of the discriminants of such a variable cannot be changed after initialization.

NOTE 2 The `default_expression` for a discriminant of a type is evaluated when an object of an unconstrained subtype of the type is created.

NOTE 3 Assignment to a discriminant of an object (after its initialization) is not allowed, since the name of a discriminant is a constant; neither `assignment_statements` nor assignments inherent in passing as an **in out** or **out** parameter are allowed. Note however that the value of a discriminant can be changed by assigning to the enclosing object, presuming it is an unconstrained variable.

NOTE 4 A discriminant that is of a named access type is not called an access discriminant; that term is used only for discriminants defined by an `access_definition`.

EXAMPLE *Examples of discriminated types:*

```

type Buffer(Size : Buffer_Size := 100) is           -- see 6.5.4
  record
    Pos   : Buffer_Size := 0;
    Value : String(1 .. Size);
  end record;

type Matrix_Rec(Rows, Columns : Integer) is
  record
    Mat : Matrix(1 .. Rows, 1 .. Columns);           -- see 6.6
  end record;

type Square(Side : Integer) is new
  Matrix_Rec(Rows => Side, Columns => Side);

type Double_Square(Number : Integer) is
  record
    Left  : Square(Number);
    Right : Square(Number);
  end record;

task type Worker(Prio : System.Priority; Buf : access Buffer
  with Priority => Prio is -- see D.1
  -- discriminants used to parameterize the task type (see 12.1)
  entry Fill;
  entry Drain;
end Worker;

```

### 6.7.1 Discriminant Constraints

A `discriminant_constraint` specifies the values of the discriminants for a given discriminated type.

*Syntax*

```

discriminant_constraint ::=
  (discriminant_association { discriminant_association })
discriminant_association ::=
  [discriminant_selector_name {'|' discriminant_selector_name} =>] expression

```

A `discriminant_association` is said to be *named* if it has one or more *discriminant\_selector\_names*; it is otherwise said to be *positional*. In a `discriminant_constraint`, any positional associations shall precede any named associations.

*Name Resolution Rules*

Each selector name of a named `discriminant_association` shall resolve to denote a discriminant of the subtype being constrained; the discriminants so named are the *associated discriminants* of the named association. For a positional association, the *associated discriminant* is the one whose `discriminant_specification` occurred in the corresponding position in the `known_discriminant_part` that defined the discriminants of the subtype being constrained.

The expected type for the expression in a `discriminant_association` is that of the associated discriminant(s).

*Legality Rules*

A `discriminant_constraint` is only allowed in a `subtype_indication` whose `subtype_mark` denotes either an unconstrained discriminated subtype, or an unconstrained access subtype whose designated subtype is an unconstrained discriminated subtype. However, in the case of an access subtype, a `discriminant_constraint` is legal only if any dereference of a value of the access type is known to be constrained (see 6.3). In addition to the places where Legality Rules normally apply (see 15.3), these rules apply also in the private part of an instance of a generic unit.

A named `discriminant_association` with more than one `selector_name` is allowed only if the named discriminants are all of the same type. A `discriminant_constraint` shall provide exactly one value for each discriminant of the subtype being constrained.

#### Dynamic Semantics

A `discriminant_constraint` is *compatible* with an unconstrained discriminated subtype if each discriminant value belongs to the subtype of the corresponding discriminant.

A composite value *satisfies* a discriminant constraint if and only if each discriminant of the composite value has the value imposed by the discriminant constraint.

For the elaboration of a `discriminant_constraint`, the expressions in the `discriminant_associations` are evaluated in an arbitrary order and converted to the type of the associated discriminant (which can raise `Constraint_Error` — see 7.6); the expression of a named association is evaluated (and converted) once for each associated discriminant. The result of each evaluation and conversion is the value imposed by the constraint for the associated discriminant.

NOTE The rules of the language ensure that a discriminant of an object always has a value, either from explicit or implicit initialization.

EXAMPLE Examples (using types declared above in 6.7):

```
Large    : Buffer(200);  -- constrained, always 200 characters
                        -- (explicit discriminant value)
Message  : Buffer;      -- unconstrained, initially 100 characters
                        -- (default discriminant value)
Basis    : Square(5);   -- constrained, always 5 by 5
Illegal  : Square;      -- illegal, a Square has to be constrained
```

## 6.7.2 Operations of Discriminated Types

If a discriminated type has default expressions for its discriminants, then unconstrained variables of the type are permitted, and the discriminants of such a variable can be changed by assignment to the variable. For a formal parameter of such a type, an attribute is provided to determine whether the corresponding actual parameter is constrained or unconstrained.

#### Static Semantics

For a prefix `A` that is of a discriminated type (after any implicit dereference), the following attribute is defined:

`A'Constrained`

Yields the value `True` if `A` denotes a constant, a value, a tagged object, or a constrained variable, and `False` otherwise. The value of this attribute is of the predefined type `Boolean`.

#### Erroneous Execution

The execution of a construct is erroneous if the construct has a constituent that is a `name` denoting a subcomponent that depends on discriminants, and the value of any of these discriminants is changed by this execution between evaluating the `name` and the last use (within this execution) of the subcomponent denoted by the `name`.

## 6.8 Record Types

A record object is a composite object consisting of named components. The value of a record object is a composite value consisting of the values of the components.

#### Syntax

```
record_type_definition ::= [[abstract] tagged] [limited] record_definition
record_definition ::=
record
```



```

    component_list
  end record [record_identifier]
| null record
component_list ::=
  component_item{component_item}
| {component_item} variant_part
| null;
component_item ::= component_declaration | aspect_clause
component_declaration ::=
  defining_identifier_list: component_definition [= default_expression]
  [aspect_specification];

```

If a *record\_identifier* appears at the end of the *record\_definition*, it shall repeat the *defining\_identifier* of the enclosing *full\_type\_declaration*.

#### Name Resolution Rules

The expected type for the *default\_expression*, if any, in a *component\_declaration* is the type of the component.

#### Legality Rules

Each *component\_declaration* declares a component of the record type. Besides components declared by *component\_declarations*, the components of a record type include any components declared by *discriminant\_specifications* of the record type declaration. The identifiers of all components of a record type shall be distinct.

Within a *type\_declaration*, a name that denotes a component, protected subprogram, or entry of the type is allowed only in the following cases:

- A name that denotes any component, protected subprogram, or entry is allowed within an *aspect\_specification*, an operational item, or a representation item that occurs within the declaration of the composite type.
- A name that denotes a noninherited discriminant is allowed within the declaration of the type, but not within the *discriminant\_part*. If the discriminant is used to define the constraint of a component, the bounds of an entry family, or the constraint of the parent subtype in a *derived\_type\_definition*, then its name shall appear alone as a *direct\_name* (not as part of a larger expression or expanded name). A discriminant shall not be used to define the constraint of a scalar component.

If the name of the current instance of a type (see 11.6) is used to define the constraint of a component, then it shall appear as a *direct\_name* that is the prefix of an *attribute\_reference* whose result is of an access type, and the *attribute\_reference* shall appear alone.

#### Static Semantics

If a *record\_type\_definition* includes the reserved word **limited**, the type is called an *explicitly limited record type*.

The *component\_definition* of a *component\_declaration* defines the (nominal) subtype of the component. If the reserved word **aliased** appears in the *component\_definition*, then the component is aliased (see 6.10).

If the *component\_list* of a record type is defined by the reserved word **null** and there are no discriminants, then the record type has no components and all records of the type are *null records*. A *record\_definition* of **null record** is equivalent to **record null; end record**.

## Dynamic Semantics

The elaboration of a `record_type_definition` creates the record type and its first subtype, and consists of the elaboration of the `record_definition`. The elaboration of a `record_definition` consists of the elaboration of its `component_list`, if any.

The elaboration of a `component_list` consists of the elaboration of the `component_items` and `variant_part`, if any, in the order in which they appear. The elaboration of a `component_declaration` consists of the elaboration of the `component_definition`.

Within the definition of a composite type, if a `component_definition` or `discrete_subtype_definition` (see 12.5.2) includes a name that denotes a discriminant of the type, or that is an `attribute_reference` whose prefix denotes the current instance of the type, the expression containing the name is called a *per-object expression*, and the constraint or range being defined is called a *per-object constraint*. For the elaboration of a `component_definition` of a `component_declaration` or the `discrete_subtype_definition` of an `entry_declaration` for an entry family (see 12.5.2), if the component subtype is defined by an `access_definition` or if the constraint or range of the `subtype_indication` or `discrete_subtype_definition` is not a per-object constraint, then the `access_definition`, `subtype_indication`, or `discrete_subtype_definition` is elaborated. On the other hand, if the constraint or range is a per-object constraint, then the elaboration consists of the evaluation of any included expression that is not part of a per-object expression. Each such expression is evaluated once unless it is part of a named association in a discriminant constraint, in which case it is evaluated once for each associated discriminant.

When a per-object constraint is elaborated (as part of creating an object), each per-object expression of the constraint is evaluated. For other expressions, the values determined during the elaboration of the `component_definition` or `entry_declaration` are used. Any checks associated with the enclosing `subtype_indication` or `discrete_subtype_definition` are performed, including the subtype compatibility check (see 6.2.2), and the associated subtype is created.

NOTE 1 A `component_declaration` with several identifiers is equivalent to a sequence of single `component_declarations`, as explained in 6.3.1.

NOTE 2 The `default_expression` of a record component is only evaluated upon the creation of a default-initialized object of the record type (presuming the object has the component, if it is in a `variant_part` — see 6.3.1).

NOTE 3 The subtype defined by a `component_definition` (see 6.6) has to be a definite subtype.

NOTE 4 If a record type does not have a `variant_part`, then the same components are present in all values of the type.

NOTE 5 A record type is limited if it has the reserved word **limited** in its definition, or if any of its components are limited (see 10.5).

NOTE 6 The predefined operations of a record type include membership tests, qualification, and explicit conversion. If the record type is nonlimited, they also include assignment and the predefined equality operators.

NOTE 7 A component of a record can be named with a `selected_component`. A value of a record can be specified with a `record_aggregate`.

EXAMPLE 1 *Examples of record type declarations:*

```
type Date is
  record
    Day   : Integer range 1 .. 31;
    Month : Month_Name;           -- see 6.5.1
    Year  : Integer range 0 .. 4000;
  end record;
```

```

type Complex is
  record
    Re : Real := 0.0;
    Im : Real := 0.0;
  end record Complex;

```

#### EXAMPLE 2 Examples of record variables:

```

Tomorrow, Yesterday : Date;
A, B, C : Complex;

-- both components of A, B, and C are implicitly initialized to zero

```

### 6.8.1 Variant Parts and Discrete Choices

A record type with a `variant_part` specifies alternative lists of components. Each variant defines the components for the value or values of the discriminant covered by its `discrete_choice_list`.

#### Syntax

```

variant_part ::=
  case discriminant_direct_name is
    variant
    {variant}
  end case;

variant ::=
  when discrete_choice_list =>
    component_list

discrete_choice_list ::= discrete_choice {'| discrete_choice}

discrete_choice ::=
  choice_expression | discrete_subtype_indication | range | others

```

#### Name Resolution Rules

The *discriminant\_direct\_name* shall resolve to denote a discriminant (called the *discriminant of the variant\_part*) specified in the known *discriminant\_part* of the full\_type\_declaration that contains the *variant\_part*. The expected type for each *discrete\_choice* in a variant is the type of the discriminant of the *variant\_part*.

#### Legality Rules

The discriminant of the *variant\_part* shall be of a discrete type.

The choice\_expressions, subtype\_indications, and ranges given as *discrete\_choices* in a *variant\_part* shall be static. The *discrete\_choice* **others** shall appear alone in a *discrete\_choice\_list*, and such a *discrete\_choice\_list*, if it appears, shall be the last one in the enclosing construct.

A *discrete\_choice* is defined to *cover a value* in the following cases:

- A *discrete\_choice* that is a *choice\_expression* covers a value if the value equals the value of the *choice\_expression* converted to the expected type.
- A *discrete\_choice* that is a *subtype\_indication* covers all values (possibly none) that belong to the subtype and that satisfy the static predicates of the subtype (see 6.2.4).
- A *discrete\_choice* that is a range covers all values (possibly none) that belong to the range.
- The *discrete\_choice* **others** covers all values of its expected type that are not covered by previous *discrete\_choice\_lists* of the same construct.

A *discrete\_choice\_list* covers a value if one of its *discrete\_choices* covers the value.

The possible values of the discriminant of a `variant_part` shall be covered as follows:

- If the discriminant is of a static constrained scalar subtype then, except within an instance of a generic unit, each **non-others** `discrete_choice` shall cover only values in that subtype that satisfy its predicates, and each value of that subtype that satisfies its predicates shall be covered by some `discrete_choice` (either explicitly or by **others**);
- If the type of the discriminant is a descendant of a generic formal scalar type, then the `variant_part` shall have an **others** `discrete_choice`;
- Otherwise, each value of the base range of the type of the discriminant shall be covered (either explicitly or by **others**).

Two distinct `discrete_choices` of a `variant_part` shall not cover the same value.

#### Static Semantics

If the `component_list` of a variant is specified by **null**, the variant has no components.

The discriminant of a `variant_part` is said to *govern* the `variant_part` and its variants. In addition, the discriminant of a derived type governs a `variant_part` and its variants if it corresponds (see 6.7) to the discriminant of the `variant_part`.

#### Dynamic Semantics

A record value contains the values of the components of a particular variant only if the value of the discriminant governing the variant is covered by the `discrete_choice_list` of the variant. This rule applies in turn to any further variant that is, itself, included in the `component_list` of the given variant.

When an object of a discriminated type *T* is initialized by default, `Constraint_Error` is raised if no `discrete_choice_list` of any variant of a `variant_part` of *T* covers the value of the discriminant that governs the `variant_part`. When a `variant_part` appears in the `component_list` of another variant *V*, this test is only applied if the value of the discriminant governing *V* is covered by the `discrete_choice_list` of *V*.

The elaboration of a `variant_part` consists of the elaboration of the `component_list` of each variant in the order in which they appear.

EXAMPLE 1 *Example of record type with a variant part:*

```

type Device is (Printer, Disk, Drum);
type State is (Open, Closed);

type Peripheral (Unit : Device := Disk) is
  record
    Status : State;
    case Unit is
      when Printer =>
        Line_Count : Integer range 1 .. Page_Size;
      when others =>
        Cylinder   : Cylinder_Index;
        Track      : Track_Number;
    end case;
  end record;

```

EXAMPLE 2 *Examples of record subtypes:*

```

subtype Drum_Unit is Peripheral (Drum);
subtype Disk_Unit is Peripheral (Disk);

```

EXAMPLE 3 *Examples of constrained record variables:*

```

Writer   : Peripheral (Unit => Printer);
Archive  : Disk_Unit;

```

## 6.9 Tagged Types and Type Extensions

Tagged types and type extensions support object-oriented programming, based on inheritance with extension and run-time polymorphism via *dispatching operations*.

### Static Semantics

A record type or private type that has the reserved word **tagged** in its declaration is called a *tagged* type. In addition, an interface type is a tagged type, as is a task or protected type derived from an interface (see 6.9.4). When deriving from a tagged type, as for any derived type, additional primitive subprograms may be defined, and inherited primitive subprograms may be overridden. The derived type is called an *extension* of its ancestor types, or simply a *type extension*.

Every type extension is also a tagged type, and is a *record extension* or a *private extension* of some other tagged type, or a noninterface synchronized tagged type (see 6.9.4). A record extension is defined by a `derived_type_definition` with a `record_extension_part` (see 6.9.1), which may include the definition of additional components. A private extension, which is a partial view of a record extension or of a synchronized tagged type, can be declared in the visible part of a package (see 10.3) or in a generic formal part (see 15.5.1).

An object of a tagged type has an associated (run-time) *tag* that identifies the specific tagged type used to create the object originally. The tag of an operand of a class-wide tagged type *T* controls which subprogram body is to be executed when a primitive subprogram of type *T* is applied to the operand (see 6.9.2); using a tag to control which body to execute is called *dispatching*.

The tag of a specific tagged type identifies the `full_type_declaration` of the type, and for a type extension, is sufficient to uniquely identify the type among all descendants of the same ancestor. If a declaration for a tagged type occurs within a `generic_package_declaration`, then the corresponding type declarations in distinct instances of the generic package are associated with distinct tags. For a tagged type that is local to a generic package body and with all of its ancestors (if any) also local to the generic body, the language does not specify whether repeated instantiations of the generic body result in distinct tags.

The following language-defined library package exists:

```
package Ada.Tags
  with Preelaborate, Nonblocking, Global => in out synchronized is
  type Tag is private
    with Preelaborable_Initialization;

  No_Tag : constant Tag;

  function Expanded_Name(T : Tag) return String;
  function Wide_Expanded_Name(T : Tag) return Wide_String;
  function Wide_Wide_Expanded_Name(T : Tag) return Wide_Wide_String;
  function External_Tag(T : Tag) return String;
  function Internal_Tag(External : String) return Tag;

  function Descendant_Tag(External : String; Ancestor : Tag) return Tag;
  function Is_Descendant_At_Same_Level(Descendant, Ancestor : Tag)
    return Boolean;

  function Parent_Tag (T : Tag) return Tag;

  type Tag_Array is array (Positive range <>) of Tag;

  function Interface_Ancestor_Tags (T : Tag) return Tag_Array;

  function Is_Abstract (T : Tag) return Boolean;

  Tag_Error : exception;

private
  ... -- not specified by the language
end Ada.Tags;
```

No\_Tag is the default initial value of type Tag.

The function `Wide_Wide_Expanded_Name` returns the full expanded name of the first subtype of the specific type identified by the tag, in upper case, starting with a root library unit. The result is implementation defined if the type is declared within an unnamed `block_statement`.

The function `Expanded_Name` (respectively, `Wide_Expanded_Name`) returns the same sequence of graphic characters as that defined for `Wide_Wide_Expanded_Name`, if all the graphic characters are defined in `Character` (respectively, `Wide_Character`); otherwise, the sequence of characters is implementation defined, but no shorter than that returned by `Wide_Wide_Expanded_Name` for the same value of the argument.

The function `External_Tag` returns a string to be used in an external representation for the given tag. The call `External_Tag(S'Tag)` is equivalent to the `attribute_reference` `S'External_Tag` (see 16.3).

The string returned by the functions `Expanded_Name`, `Wide_Expanded_Name`, `Wide_Wide_Expanded_Name`, and `External_Tag` has lower bound 1.

The function `Internal_Tag` returns a tag that corresponds to the given external tag, or raises `Tag_Error` if the given string is not the external tag for any specific type of the partition. `Tag_Error` is also raised if the specific type identified is a library-level type whose tag has not yet been created (see 16.14).

The function `Descendant_Tag` returns the (internal) tag for the type that corresponds to the given external tag and is both a descendant of the type identified by the Ancestor tag and has the same accessibility level as the identified ancestor. `Tag_Error` is raised if External is not the external tag for such a type. `Tag_Error` is also raised if the specific type identified is a library-level type whose tag has not yet been created, or if the given external tag identifies more than one type that has the appropriate Ancestor and accessibility level.

The function `Is_Descendant_At_Same_Level` returns `True` if the Descendant tag identifies a type that is both a descendant of the type identified by Ancestor and at the same accessibility level. If not, it returns `False`.

For the purposes of the dynamic semantics of functions `Descendant_Tag` and `Is_Descendant_At_Same_Level`, a tagged type `T2` is a *descendant* of a type `T1` if it is the same as `T1`, or if its parent type or one of its progenitor types is a descendant of type `T1` by this rule, even if at the point of the declaration of `T2`, one of the derivations in the chain is not visible.

The function `Parent_Tag` returns the tag of the parent type of the type whose tag is `T`. If the type does not have a parent type (that is, it was not defined by a `derived_type_definition`), then `No_Tag` is returned.

The function `Interface_Ancessor_Tags` returns an array containing the tag of each interface ancestor type of the type whose tag is `T`, other than `T` itself. The lower bound of the returned array is 1, and the order of the returned tags is unspecified. Each tag appears in the result exactly once. If the type whose tag is `T` has no interface ancestors, a null array is returned.

The function `Is_Abstract` returns `True` if the type whose tag is `T` is abstract, and `False` otherwise.

For every subtype `S` of a tagged type `T` (specific or class-wide), the following attributes are defined:

- |                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>S'Class</code> | <code>S'Class</code> denotes a subtype of the class-wide type (called <i>T'Class</i> in this document) for the class rooted at <i>T</i> (or if <i>S</i> already denotes a class-wide subtype, then <code>S'Class</code> is the same as <i>S</i> ).<br><br><code>S'Class</code> is unconstrained. However, if <i>S</i> is constrained, then the values of <code>S'Class</code> are only those that when converted to the type <i>T</i> belong to <i>S</i> . |
| <code>S'Tag</code>   | <code>S'Tag</code> denotes the tag of the type <i>T</i> (or if <i>T</i> is class-wide, the tag of the root type of the corresponding class). The value of this attribute is of type <code>Tag</code> .                                                                                                                                                                                                                                                     |

Given a `prefix` `X` that is of a class-wide tagged type (after any implicit dereference), the following attribute is defined:

- |                    |                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------|
| <code>X'Tag</code> | <code>X'Tag</code> denotes the tag of <i>X</i> . The value of this attribute is of type <code>Tag</code> . |
|--------------------|------------------------------------------------------------------------------------------------------------|

The following language-defined generic function exists:



```

generic
  type T (<>) is abstract tagged limited private;
  type Parameters (<>) is limited private;
  with function Constructor (Params : not null access Parameters)
    return T is abstract;
function Ada.Tags.Generic_Dispatching_Constructor
  (The_Tag : Tag;
   Params : not null access Parameters) return T'Class
  with Preelaborate, Convention => Intrinsic,
       Nonblocking, Global => in out synchronized;

```

Tags.Generic\_Dispatching\_Constructor provides a mechanism to create an object of an appropriate type from just a tag value. The function Constructor is expected to create the object given a reference to an object of type Parameters.

#### Dynamic Semantics

The tag associated with an object of a tagged type is determined as follows:

- The tag of a stand-alone object, a component, or an aggregate of a specific tagged type *T* identifies *T*.
- The tag of an object created by an allocator for an access type with a specific designated tagged type *T* identifies *T*.
- The tag of an object of a class-wide tagged type is that of its initialization expression.
- The tag of the result returned by a function whose result type is a specific tagged type *T* identifies *T*.
- The tag of the result returned by a function with a class-wide result type is that of the return object.

The tag is preserved by type conversion and by parameter passing. The tag of a value is the tag of the associated object (see 9.2).

Tag\_Error is raised by a call of Descendant\_Tag, Expanded\_Name, External\_Tag, Interface\_Ancessor\_Tags, Is\_Abstract, Is\_Descendant\_At\_Same\_Level, Parent\_Tag, Wide\_Expanded\_Name, or Wide\_Wide\_Expanded\_Name if any tag passed is No\_Tag.

An instance of Tags.Generic\_Dispatching\_Constructor raises Tag\_Error if The\_Tag does not represent a concrete descendant of T or if the innermost master (see 10.6.1) of this descendant is not also a master of the instance. Otherwise, it dispatches to the primitive function denoted by the formal Constructor for the type identified by The\_Tag, passing Params, and returns the result. Any exception raised by the function is propagated.

#### Erroneous Execution

If an internal tag provided to an instance of Tags.Generic\_Dispatching\_Constructor or to any subprogram declared in package Tags identifies either a type that is not library-level and whose tag has not been created (see 16.14), or a type that does not exist in the partition at the time of the call, then execution is erroneous.

#### Implementation Permissions

The implementation of Internal\_Tag and Descendant\_Tag may raise Tag\_Error if no specific type corresponding to the string External passed as a parameter exists in the partition at the time the function is called, or if there is no such type whose innermost master is a master of the point of the function call.

#### Implementation Advice

Internal\_Tag should return the tag of a type, if one exists, whose innermost master is a master of the point of the function call.

NOTE 1 A type declared with the reserved word **tagged** is normally declared in a package\_specification, so that new primitive subprograms can be declared for it.

NOTE 2 Once an object has been created, its tag never changes.

NOTE 3 Class-wide types are defined to have unknown discriminants (see 6.7). This means that, by the rules in 6.7 for objects with unknown discriminants, objects of a class-wide type are illegal unless they are explicitly initialized (whether created by an `object_declaration` or an `allocator`), and that aggregates are illegal unless they are explicitly qualified with a specific type when their expected type is class-wide.

NOTE 4 The capability provided by `Tags.Generic_Dispatching_Constructor` is sometimes known as a *factory*.

EXAMPLE *Examples of tagged record types:*

```

type Point is tagged
  record
    X, Y : Real := 0.0;
  end record;

type Expression is tagged null record;
  -- Components will be added by each extension

```

## 6.9.1 Type Extensions

Every type extension is a tagged type, and is a *record extension* or a *private extension* of some other tagged type, or a noninterface synchronized tagged type.

### Syntax

```
record_extension_part ::= with record_definition
```

### Legality Rules

The parent type of a record extension shall not be a class-wide type nor shall it be a synchronized tagged type (see 6.9.4). If the parent type or any progenitor is nonlimited, then each of the components of the `record_extension_part` shall be nonlimited. In addition to the places where Legality Rules normally apply (see 15.3), these rules apply also in the private part of an instance of a generic unit.

Within the body of a generic unit, or the body of any of its descendant library units, a tagged type shall not be declared as a descendant of a formal type declared within the formal part of the generic unit.

### Static Semantics

A record extension is a *null extension* if its declaration has no `known_discriminant_part` and its `record_extension_part` includes no component declarations.

In the case where the (compile-time) view of an object *X* is of a tagged type *T1* or *T1'Class* and the (run-time) tag of *X* is *T2'Tag*, only the components (if any) of *X* that are components of *T1* (or that are discriminants which correspond to a discriminant of *T1*) are said to be *components of the nominal type* of *X*. Similarly, only parts (respectively, subcomponents) of *T1* are parts (respectively, subcomponents) of the nominal type of *X*.

### Dynamic Semantics

The elaboration of a `record_extension_part` consists of the elaboration of the `record_definition`.

NOTE 1 The term “type extension” refers to a type as a whole. The term “extension part” refers to the piece of text that defines the additional components (if any) the type extension has relative to its specified ancestor type.

NOTE 2 When an extension is declared immediately within a body, primitive subprograms are inherited and are overridable, but new primitive subprograms cannot be added.

NOTE 3 By the rules given in 6.8, a name that denotes a component (including a discriminant) of the parent type is illegal within the `record_extension_part`. Similarly, a name that denotes a component defined within the `record_extension_part` is illegal within the `record_extension_part`. A name that denotes a discriminant of the record extension is legal, providing that it refers to a discriminant defined in a new `known_discriminant_part` in the enclosing type declaration.

NOTE 4 By the rules given in 11.3, each visible component of a record extension will have a unique name, whether the component is (visibly) inherited from the parent type or declared in the `record_extension_part`.

EXAMPLE Examples of record extensions (of types defined above in 6.9):

```

type Painted_Point is new Point with
  record
    Paint : Color := White;
  end record;
  -- Components X and Y are inherited

Origin : constant Painted_Point := (X | Y => 0.0, Paint => Black);

type Literal is new Expression with
  record                                -- a leaf in an Expression tree
    Value : Real;
  end record;

type Expr_Ptr is access all Expression'Class;
  -- see 6.9

type Binary_Operation is new Expression with
  record                                -- an internal node in an Expression tree
    Left, Right : Expr_Ptr;
  end record;

type Addition is new Binary_Operation with null record;
type Subtraction is new Binary_Operation with null record;
  -- No additional components needed for these extensions

Tree : Expr_Ptr :=                -- A tree representation of "5.0 + (13.0-7.0)"
  new Addition'(
    Left => new Literal'(Value => 5.0),
    Right => new Subtraction'(
      Left => new Literal'(Value => 13.0),
      Right => new Literal'(Value => 7.0));

```

## 6.9.2 Dispatching Operations of Tagged Types

The primitive subprograms of a tagged type, the subprograms declared by `formal_abstract_subprogram_declarations`, the `Put_Image` attribute (see 7.10) of a specific tagged type, and the stream attributes of a specific tagged type that are available (see 16.13.2) at the end of the declaration list where the type is declared are called *dispatching operations*. A dispatching operation can be called using a statically determined *controlling* tag, in which case the body to be executed is determined at compile time. Alternatively, the controlling tag can be dynamically determined, in which case the call *dispatches* to a body that is determined at run time; such a call is termed a *dispatching call*. As explained below, the properties of the operands and the context of a particular call on a dispatching operation determine how the controlling tag is determined, and hence whether or not the call is a dispatching call. Run-time polymorphism is achieved when a dispatching operation is called by a dispatching call.

### Static Semantics

A call on a dispatching operation is a call whose name or prefix denotes the declaration of a dispatching operation. A *controlling operand* in a call on a dispatching operation of a tagged type *T* is one whose corresponding formal parameter is of type *T* or is of an anonymous access type with designated type *T*; the corresponding formal parameter is called a *controlling formal parameter*. If the controlling formal parameter is an access parameter, the controlling operand is the object designated by the actual parameter, rather than the actual parameter itself. If the call is to a (primitive) function with result type *T* (a *function with a controlling result*), then the call has a *controlling result* — the context of the call can control the dispatching. Similarly, if the call is to a function with an access result type designating *T* (a *function with a controlling access result*), then the call has a *controlling access result*, and the context can similarly control dispatching.

A name or expression of a tagged type is either *statically* tagged, *dynamically* tagged, or *tag indeterminate*, according to whether, when used as a controlling operand, the tag that controls dispatching is determined statically by the operand's (specific) type, dynamically by its tag at run time, or from context. A `qualified_expression` or `parenthesized expression` is statically, dynamically, or indeterminately tagged according to its operand. A `conditional_expression` is statically,

dynamically, or indeterminately tagged according to rules given in 7.5.7. A `declare_expression` is statically, dynamically, or indeterminately tagged according to its `body_expression`. For other kinds of names and expressions, this is determined as follows:

- The name or expression is *statically tagged* if it is of a specific tagged type and, if it is a call with a controlling result or controlling access result, it has at least one statically tagged controlling operand;
- The name or expression is *dynamically tagged* if it is of a class-wide type, or it is a call with a controlling result or controlling access result and at least one dynamically tagged controlling operand;
- The name or expression is *tag indeterminate* if it is a call with a controlling result or controlling access result, all of whose controlling operands (if any) are tag indeterminate.

A `type_conversion` is statically or dynamically tagged according to whether the type determined by the `subtype_mark` is specific or class-wide, respectively. For an object that is designated by an expression whose expected type is an anonymous access-to-specific tagged type, the object is dynamically tagged if the expression, ignoring enclosing parentheses, is of the form `X'Access`, where `X` is of a class-wide type, or is of the form `new T(...)`, where `T` denotes a class-wide subtype. Otherwise, the object is statically or dynamically tagged according to whether the designated type of the type of the expression is specific or class-wide, respectively.

#### Legality Rules

A call on a dispatching operation shall not have both dynamically tagged and statically tagged controlling operands.

If the expected type for an expression or name is some specific tagged type, then the expression or name shall not be dynamically tagged unless it is a controlling operand in a call on a dispatching operation. Similarly, if the expected type for an expression is an anonymous access-to-specific tagged type, then the object designated by the expression shall not be dynamically tagged unless it is a controlling operand in a call on a dispatching operation.

In the declaration of a dispatching operation of a tagged type, everywhere a subtype of the tagged type appears as a subtype of the profile (see 9.1), it shall statically match the first subtype of the tagged type. If the dispatching operation overrides an inherited subprogram, it shall be subtype conformant with the inherited subprogram. The convention of an inherited dispatching operation is the convention of the corresponding primitive operation of the parent or progenitor type. The default convention of a dispatching operation that overrides an inherited primitive operation is the convention of the inherited operation; if the operation overrides multiple inherited operations, then they shall all have the same convention. An explicitly declared dispatching operation shall not be of convention *Intrinsic*.

The `default_expression` for a controlling formal parameter of a dispatching operation shall be tag indeterminate.

If a dispatching operation is defined by a `subprogram_renaming_declaration` or the instantiation of a generic subprogram, any access parameter of the renamed subprogram or the generic subprogram that corresponds to a controlling access parameter of the dispatching operation, shall have a subtype that excludes null.

A given subprogram shall not be a dispatching operation of two or more distinct tagged types.

The explicit declaration of a primitive subprogram of a tagged type shall occur before the type is frozen (see 16.14). For example, new dispatching operations cannot be added after objects or values of the type exist, nor after deriving a record extension from it, nor after a body.

#### Dynamic Semantics

For the execution of a call on a dispatching operation of a type *T*, the *controlling tag value* determines which subprogram body is executed. The controlling tag value is defined as follows:

- If one or more controlling operands are statically tagged, then the controlling tag value is *statically determined* to be the tag of *T*.
- If one or more controlling operands are dynamically tagged, then the controlling tag value is not statically determined, but is rather determined by the tags of the controlling operands. If there is more than one dynamically tagged controlling operand, a check is made that they all have the same tag. If this check fails, `Constraint_Error` is raised unless the call is a `function_call` whose name denotes the declaration of an equality operator (predefined or user defined) that returns Boolean, in which case the result of the call is defined to indicate inequality, and no `subprogram_body` is executed. This check is performed prior to evaluating any tag-indeterminate controlling operands.
- If all of the controlling operands (if any) are tag-indeterminate, then:
  - If the call has a controlling result or controlling access result and is itself, or designates, a (possibly parenthesized or qualified) controlling operand of an enclosing call on a dispatching operation of a descendant of type *T*, then its controlling tag value is determined by the controlling tag value of this enclosing call;
  - If the call has a controlling result or controlling access result and (possibly parenthesized, qualified, or dereferenced) is the expression of an `assignment_statement` whose target is of a class-wide type, then its controlling tag value is determined by the target;
  - Otherwise, the controlling tag value is statically determined to be the tag of type *T*.

For the execution of a call on a dispatching operation, the action performed is determined by the properties of the corresponding dispatching operation of the specific type identified by the controlling tag value:

- if the corresponding operation is explicitly declared for this type, even if the declaration occurs in a private part, then the action comprises an invocation of the explicit body for the operation;
- if the corresponding operation is implicitly declared for this type and is implemented by an entry or protected subprogram (see 12.1 and 12.4), then the action comprises a call on this entry or protected subprogram, with the target object being given by the first actual parameter of the call, and the actual parameters of the entry or protected subprogram being given by the remaining actual parameters of the call, if any;
- if the corresponding operation is a predefined operator then the action comprises an invocation of that operator;
- otherwise, the action is the same as the action for the corresponding operation of the parent type or progenitor type from which the operation was inherited except that additional invariant checks (see 10.3.2) and class-wide postcondition checks (see 9.1.1) may apply. If there is more than one such corresponding operation, the action is that for the operation that is not a null procedure, if any; otherwise, the action is that of an arbitrary one of the operations.

NOTE 1 The body to be executed for a call on a dispatching operation is determined by the tag; it does not matter whether that tag is determined statically or dynamically, and it does not matter whether the subprogram's declaration is visible at the place of the call.

NOTE 2 This subclause covers calls on dispatching subprograms of a tagged type. Rules for tagged type membership tests are described in 7.5.2. Controlling tag determination for an `assignment_statement` is described in 8.2.

NOTE 3 A dispatching call can dispatch to a body whose declaration is not visible at the place of the call.

NOTE 4 A call through an access-to-subprogram value is never a dispatching call, even if the access value designates a dispatching operation. Similarly a call whose prefix denotes a `subprogram_renaming_declaration` cannot be a dispatching call unless the renaming itself is the declaration of a primitive subprogram.



### 6.9.3 Abstract Types and Subprograms

An *abstract type* is a tagged type intended for use as an ancestor of other types, but which is not allowed to have objects of its own. An *abstract subprogram* is a subprogram that has no body, but is intended to be overridden at some point when inherited. Because objects of an abstract type cannot be created, a dispatching call to an abstract subprogram always dispatches to some overriding body.

#### Syntax

```
abstract_subprogram_declaration ::=
  [overriding_indicator]
  subprogram_specification is abstract
  [aspect_specification];
```

#### Static Semantics

Interface types (see 6.9.4) are abstract types. In addition, a tagged type that has the reserved word **abstract** in its declaration is an abstract type. The class-wide type (see 6.4.1) rooted at an abstract type is not itself an abstract type.

#### Legality Rules

Only a tagged type shall have the reserved word **abstract** in its declaration.

A subprogram declared by an `abstract_subprogram_declaration` or a `formal_abstract_subprogram_declaration` (see 15.6) is an *abstract subprogram*. If it is a primitive subprogram of a tagged type, then the tagged type shall be abstract.

If a type has an implicitly declared primitive subprogram that is inherited or is a predefined operator, and the corresponding primitive subprogram of the parent or ancestor type is abstract or is a function with a controlling access result, or if a type other than a nonabstract null extension inherits a function with a controlling result, then:

- If the type is abstract or untagged, the implicitly declared subprogram is *abstract*.
- Otherwise, the subprogram shall be overridden with a nonabstract subprogram or, in the case of a private extension inheriting a nonabstract function with a controlling result, have a full type that is a null extension; for a type declared in the visible part of a package, the overriding may be either in the visible or the private part. Such a subprogram is said to *require overriding*. However, if the type is a generic formal type, the subprogram is allowed to be inherited as is, without being overridden for the formal type itself; a nonabstract version will necessarily be provided by the actual type.

A call on an abstract subprogram shall be a dispatching call; nondispatching calls to an abstract subprogram are not allowed. In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

If the name or prefix given in an `iterator_procedure_call` (see 8.5.3) denotes an abstract subprogram, the subprogram shall be a dispatching subprogram.

The type of an aggregate, of an object created by an `object_declaration` or an `allocator`, or of a generic formal object of mode **in**, shall not be abstract. The type of the target of an assignment operation (see 8.2) shall not be abstract. The type of a component shall not be abstract. If the result type of a function is abstract, then the function shall be abstract. If a function has an access result type designating an abstract type, then the function shall be abstract. The type denoted by a `return_subtype_indication` (see 9.5) shall not be abstract. A generic function shall not have an abstract result type or an access result type designating an abstract type.

If a partial view is not abstract, the corresponding full view shall not be abstract. If a generic formal type is abstract, then for each primitive subprogram of the formal that is not abstract, the corresponding primitive subprogram of the actual shall not be abstract.

For an abstract type declared in a visible part, an abstract primitive subprogram shall not be declared in the private part, unless it is overriding an abstract subprogram implicitly declared in the visible part. For a tagged type declared in a visible part, a primitive function with a controlling result or a



controlling access result shall not be declared in the private part, unless it is overriding a function implicitly declared in the visible part.

A generic actual subprogram shall not be an abstract subprogram unless the generic formal subprogram is declared by a `formal_abstract_subprogram_declaration`. The prefix of an `attribute_reference` for the `Access`, `Unchecked_Access`, or `Address` attributes shall not denote an abstract subprogram.

#### Dynamic Semantics

The elaboration of an `abstract_subprogram_declaration` has no effect.

NOTE 1 Abstractness is not inherited; a type is abstract only if the reserved word **abstract** is used in the declaration of the type extension.

NOTE 2 A class-wide type is never abstract. Even if a class is rooted at an abstract type, the class-wide type for the class is not abstract, and an object of the class-wide type can be created; the tag of such an object will identify some nonabstract type in the class.

EXAMPLE *Example of an abstract type representing a set of natural numbers:*

```
package Sets is
  subtype Element_Type is Natural;
  type Set is abstract tagged null record;
  function Empty return Set is abstract;
  function Union(Left, Right : Set) return Set is abstract;
  function Intersection(Left, Right : Set) return Set is abstract;
  function Unit_Set(Element : Element_Type) return Set is abstract;
  procedure Take(Element : out Element_Type;
    From : in out Set) is abstract;
end Sets;
```

Given the above abstract type, one can derive various (nonabstract) extensions of the type, representing alternative implementations of a set. One possibility is to use a bit vector, but impose an upper bound on the largest element representable, while another possible implementation is a hash table, trading off space for flexibility.

## 6.9.4 Interface Types

An interface type is an abstract tagged type that provides a restricted form of multiple inheritance. A tagged type, task type, or protected type may have one or more interface types as ancestors.

#### Syntax

```
interface_type_definition ::=
  [limited | task | protected | synchronized] interface [and interface_list]
interface_list ::= interface_subtype_mark {and interface_subtype_mark}
```

#### Static Semantics

An interface type (also called an *interface*) is a specific abstract tagged type that is defined by an `interface_type_definition`.

An interface with the reserved word **limited**, **task**, **protected**, or **synchronized** in its definition is termed, respectively, a *limited interface*, a *task interface*, a *protected interface*, or a *synchronized interface*. In addition, all task and protected interfaces are synchronized interfaces, and all synchronized interfaces are limited interfaces.

A task or protected type derived from an interface is a tagged type. Such a tagged type is called a *synchronized tagged type*, as are synchronized interfaces and private extensions whose declaration includes the reserved word **synchronized**.

A task interface is an abstract task type. A protected interface is an abstract protected type.

An interface type has no components.

An *interface\_subtype\_mark* in an *interface\_list* names a *progenitor subtype*; its type is the *progenitor type*. An interface type inherits user-defined primitive subprograms from each progenitor type in the same way that a derived type inherits user-defined primitive subprograms from its progenitor types (see 6.4).

#### Legality Rules

All user-defined primitive subprograms of an interface type shall be abstract subprograms or null procedures.

The type of a subtype named in an *interface\_list* shall be an interface type.

A type derived from a nonlimited interface shall be nonlimited.

An interface derived from a task interface shall include the reserved word **task** in its definition; any other type derived from a task interface shall be a private extension or a task type declared by a task declaration (see 12.1).

An interface derived from a protected interface shall include the reserved word **protected** in its definition; any other type derived from a protected interface shall be a private extension or a protected type declared by a protected declaration (see 12.4).

An interface derived from a synchronized interface shall include one of the reserved words **task**, **protected**, or **synchronized** in its definition; any other type derived from a synchronized interface shall be a private extension, a task type declared by a task declaration, or a protected type declared by a protected declaration.

No type shall be derived from both a task interface and a protected interface.

In addition to the places where Legality Rules normally apply (see 15.3), these rules apply also in the private part of an instance of a generic unit.

#### Dynamic Semantics

The elaboration of an *interface\_type\_definition* creates the interface type and its first subtype.

NOTE Nonlimited interface types have predefined nonabstract equality operators. These can be overridden with user-defined abstract equality operators. Such operators will then require an explicit overriding for any nonabstract descendant of the interface.

EXAMPLE 1 *Example of a limited interface and a synchronized interface extending it:*

```

type Queue is limited interface;
procedure Append(Q : in out Queue; Person : in Person_Name) is abstract;
procedure Remove_First(Q : in out Queue;
                        Person : out Person_Name) is abstract;
function Cur_Count(Q : in Queue) return Natural is abstract;
function Max_Count(Q : in Queue) return Natural is abstract;
-- See 6.10.1 for Person_Name.
Queue_Error : exception;
-- Append raises Queue_Error if Cur_Count(Q) = Max_Count(Q)
-- Remove_First raises Queue_Error if Cur_Count(Q) = 0

type Synchronized_Queue is synchronized interface and Queue; -- see 12.11
procedure Append_Wait(Q : in out Synchronized_Queue;
                      Person : in Person_Name) is abstract;
procedure Remove_First_Wait(Q : in out Synchronized_Queue;
                             Person : out Person_Name) is abstract;
...

```

```

procedure Transfer(From   : in out Queue'Class;
                   To     : in out Queue'Class;
                   Number : in      Natural := 1) is
    Person : Person_Name;
begin
    for I in 1..Number loop
        Remove_First(From, Person);
        Append(To, Person);
    end loop;
end Transfer;

```

This defines a Queue interface defining a queue of people. (A similar design is possible to define any kind of queue simply by replacing Person\_Name by an appropriate type.) The Queue interface has four dispatching operations, Append, Remove\_First, Cur\_Count, and Max\_Count. The body of a class-wide operation, Transfer is also shown. Every nonabstract extension of Queue will provide implementations for at least its four dispatching operations, as they are abstract. Any object of a type derived from Queue can be passed to Transfer as either the From or the To operand. The two operands can be of different types in a given call.

The Synchronized\_Queue interface inherits the four dispatching operations from Queue and adds two additional dispatching operations, which wait if necessary rather than raising the Queue\_Error exception. This synchronized interface can only be implemented by a task or protected type, and as such ensures safe concurrent access.

*Example use of the interface:*

```

type Fast_Food_Queue is new Queue with record ...;
procedure Append(Q : in out Fast_Food_Queue; Person : in Person_Name);
procedure Remove_First(Q : in out Fast_Food_Queue;
                      Person : out Person_Name);
function Cur_Count(Q : in Fast_Food_Queue) return Natural;
function Max_Count(Q : in Fast_Food_Queue) return Natural;
...
Cashier, Counter : Fast_Food_Queue;
...
-- Add Casey (see 6.10.1) to the cashier's queue:
Append (Cashier, Casey);
-- After payment, move Casey to the sandwich counter queue:
Transfer (Cashier, Counter);
...

```

An interface such as Queue can be used directly as the parent of a new type (as shown here), or can be used as a progenitor when a type is derived. In either case, the primitive operations of the interface are inherited. For Queue, the implementation of the four inherited routines will necessarily be provided. Inside the call of Transfer, calls will dispatch to the implementations of Append and Remove\_First for type Fast\_Food\_Queue.

**EXAMPLE 2** *Example of a task interface:*

```

type Serial_Device is task interface; -- see 12.1
procedure Read (Dev : in Serial_Device; C : out Character) is abstract;
procedure Write(Dev : in Serial_Device; C : in Character) is abstract;

```

The Serial\_Device interface has two dispatching operations which are intended to be implemented by task entries (see 9.1).

## 6.10 Access Types

A value of an access type (an *access value*) provides indirect access to the object or subprogram it *designates*. Depending on its type, an access value can designate either subprograms, objects created by allocators (see 7.8), or more generally *aliased* objects of an appropriate type.

*Syntax*

```

access_type_definition ::=
    [null_exclusion] access_to_object_definition
  | [null_exclusion] access_to_subprogram_definition
access_to_object_definition ::=
    access [general_access_modifier] subtype_indication

```

```

general_access_modifier ::= all | constant
access_to_subprogram_definition ::=
    access [protected] procedure parameter_profile
    | access [protected] function parameter_and_result_profile
null_exclusion ::= not null
access_definition ::=
    [null_exclusion] access [constant] subtype_mark
    | [null_exclusion] access [protected] procedure parameter_profile
    | [null_exclusion] access [protected] function parameter_and_result_profile

```

#### Static Semantics

There are two kinds of access types, *access-to-object* types, whose values designate objects, and *access-to-subprogram* types, whose values designate subprograms. Associated with an access-to-object type is a *storage pool*; several access types may share the same storage pool. All descendants of an access type share the same storage pool. A storage pool is an area of storage used to hold dynamically allocated objects (called *pool elements*) created by allocators; storage pools are described further in 16.11.

Access-to-object types are further subdivided into *pool-specific* access types, whose values can designate only the elements of their associated storage pool, and *general* access types, whose values can designate the elements of any storage pool, as well as aliased objects created by declarations rather than allocators, and aliased subcomponents of other objects.

A view of an object is defined to be *aliased* if it is defined by an object\_declaration, component\_definition, parameter\_specification, or extended\_return\_object\_declaration with the reserved word **aliased**, or by a renaming of an aliased view. In addition, the dereference of an access-to-object value denotes an aliased view, as does a view conversion (see 7.6) of an aliased view. A qualified\_expression denotes an aliased view when the operand denotes an aliased view. The current instance of an immutably limited type (see 10.5) is defined to be aliased. Finally, a formal parameter or generic formal object of a tagged type is defined to be aliased. Aliased views are the ones that can be designated by an access value.

An access\_to\_object\_definition defines an access-to-object type and its first subtype; the subtype\_indication defines the *designated subtype* of the access type. If a general\_access\_modifier appears, then the access type is a general access type. If the modifier is the reserved word **constant**, then the type is an *access-to-constant type*; a designated object cannot be updated through a value of such a type. If the modifier is the reserved word **all**, then the type is an *access-to-variable type*; a designated object can be both read and updated through a value of such a type. If no general\_access\_modifier appears in the access\_to\_object\_definition, the access type is a pool-specific access-to-variable type.

An access\_to\_subprogram\_definition defines an access-to-subprogram type and its first subtype; the parameter\_profile or parameter\_and\_result\_profile defines the *designated profile* of the access type. There is a *calling convention* associated with the designated profile; only subprograms with this calling convention can be designated by values of the access type. By default, the calling convention is “*protected*” if the reserved word **protected** appears, and “*Ada*” otherwise. See Annex B for how to override this default.

An access\_definition defines an anonymous general access type or an anonymous access-to-subprogram type. For a general access type, the subtype\_mark denotes its *designated subtype*; if the general\_access\_modifier **constant** appears, the type is an access-to-constant type; otherwise, it is an access-to-variable type. For an access-to-subprogram type, the parameter\_profile or parameter\_and\_result\_profile denotes its *designated profile*.

For each access type, there is a null access value designating no entity at all, which can be obtained by (implicitly) converting the literal **null** to the access type. The null value of an access type is the default initial value of the type. Nonnull values of an access-to-object type are obtained by evaluating an allocator, which returns an access value designating a newly created object (see 6.10.2), or in the

case of a general access-to-object type, evaluating an `attribute_reference` for the `Access` or `Unchecked_Access` attribute of an aliased view of an object. Nonnull values of an access-to-subprogram type are obtained by evaluating an `attribute_reference` for the `Access` attribute of a nonintrinsic subprogram.

A `null_exclusion` in a construct specifies that the null value does not belong to the access subtype defined by the construct, that is, the access subtype *excludes null*. In addition, the anonymous access subtype defined by the `access_definition` for a controlling access parameter (see 6.9.2) excludes null. Finally, for a `subtype_indication` without a `null_exclusion`, the subtype denoted by the `subtype_indication` excludes null if and only if the subtype denoted by the `subtype_mark` in the `subtype_indication` excludes null.

All subtypes of an access-to-subprogram type are constrained. The first subtype of a type defined by an `access_definition` or an `access_to_object_definition` is unconstrained if the designated subtype is an unconstrained array or discriminated subtype; otherwise, it is constrained.

#### Legality Rules

If a `subtype_indication`, `discriminant_specification`, `parameter_specification`, `parameter_and_result_profile`, `object_renaming_declaration`, or `formal_object_declaration` has a `null_exclusion`, the `subtype_mark` in that construct shall denote an access subtype that does not exclude null.

#### Dynamic Semantics

A `composite_constraint` is *compatible* with an unconstrained access subtype if it is compatible with the designated subtype. A `null_exclusion` is compatible with any access subtype that does not exclude null. An access value *satisfies* a `composite_constraint` of an access subtype if it equals the null value of its type or if it designates an object whose value satisfies the constraint. An access value satisfies an exclusion of the null value if it does not equal the null value of its type.

The elaboration of an `access_type_definition` creates the access type and its first subtype. For an access-to-object type, this elaboration includes the elaboration of the `subtype_indication`, which creates the designated subtype.

The elaboration of an `access_definition` creates an anonymous access type.

NOTE 1 Access values are called “pointers” or “references” in some other languages.

NOTE 2 Each access-to-object type has an associated storage pool; several access types can share the same pool. An object can be created in the storage pool of an access type by an `allocator` (see 7.8) for the access type. A storage pool (roughly) corresponds to what some other languages call a “heap”. See 16.11 for a discussion of pools.

NOTE 3 Only index constraints and `discriminant_constraints` can be applied to access types (see 6.6.1 and 6.7.1).

EXAMPLE 1 *Examples of access-to-object types:*

```
type Frame is access Matrix;  -- see 6.6
type Peripheral_Ref is not null access Peripheral;  -- see 6.8.1
type Binop_Ptr is access all Binary_Operation'Class;
                                     -- general access-to-class-wide,
see 6.9.1
```

EXAMPLE 2 *Example of an access subtype:*

```
subtype Drum_Ref is Peripheral_Ref(Drum);  -- see 6.8.1
```

EXAMPLE 3 *Example of an access-to-subprogram type:*

```

type Message_Procedure is access procedure (M : in String := "Error!");
procedure Default_Message_Procedure(M : in String);
Give_Message : Message_Procedure := Default_Message_Procedure'Access;
...
procedure Other_Procedure(M : in String);
...
Give_Message := Other_Procedure'Access;
...
Give_Message("File not found."); -- call with parameter (.all is optional)
Give_Message.all;                -- call with no parameters

```

### 6.10.1 Incomplete Type Declarations

There are no particular limitations on the designated type of an access type. In particular, the type of a component of the designated type can be another access type, or even the same access type. This permits mutually dependent and recursive access types. An `incomplete_type_declaration` can be used to introduce a type to be used as a designated type, while deferring its full definition to a subsequent `full_type_declaration`.

#### Syntax

```

incomplete_type_declaration ::=
  type defining_identifier [discriminant_part] [is tagged];

```

#### Static Semantics

An `incomplete_type_declaration` declares an *incomplete view* of a type and its first subtype; the first subtype is unconstrained if a `discriminant_part` appears. If the `incomplete_type_declaration` includes the reserved word **tagged**, it declares a *tagged incomplete view*. If  $T$  denotes a tagged incomplete view, then  $T^{\text{Class}}$  denotes a tagged incomplete view. An incomplete view of a type is a limited view of the type (see 10.5).

Given an access type  $A$  whose designated type  $T$  is an incomplete view, a dereference of a value of type  $A$  also has this incomplete view except when:

- it occurs within the immediate scope of the completion of  $T$ , or
- it occurs within the scope of a `nonlimited_with_clause` that mentions a library package in whose visible part the completion of  $T$  is declared, or
- it occurs within the scope of the completion of  $T$  and  $T$  is an incomplete view declared by an `incomplete_type_declaration`.

In these cases, the dereference has the view of  $T$  visible at the point of the dereference.

Similarly, if a `subtype_mark` denotes a `subtype_declaration` defining a subtype of an incomplete view  $T$ , the `subtype_mark` denotes an incomplete view except under the same three circumstances given above, in which case it denotes the view of  $T$  visible at the point of the `subtype_mark`.

#### Legality Rules

An `incomplete_type_declaration` requires a completion, which shall be a `type_declaration` other than an `incomplete_type_declaration`. If the `incomplete_type_declaration` occurs immediately within either the visible part of a `package_specification` or a `declarative_part`, then the `type_declaration` shall occur later and immediately within this visible part or `declarative_part`. If the `incomplete_type_declaration` occurs immediately within the private part of a given `package_specification`, then the `type_declaration` shall occur later and immediately within either the private part itself, or the `declarative_part` of the corresponding `package_body`.



If an `incomplete_type_declaration` includes the reserved word **tagged**, then a `type_declaration` that completes it shall declare a tagged type. If an `incomplete_type_declaration` has a `known_discriminant_part`, then a `type_declaration` that completes it shall have a fully conforming (explicit) `known_discriminant_part` (see 9.3.1). If an `incomplete_type_declaration` has no `discriminant_part` (or an `unknown_discriminant_part`), then a corresponding `type_declaration` is nevertheless allowed to have discriminants, either explicitly, or inherited via derivation.

A name that denotes an incomplete view of a type may be used as follows:

- as the `subtype_mark` in the `subtype_indication` of an `access_to_object_definition`; the only form of constraint allowed in this `subtype_indication` is a `discriminant_constraint` (a `null_exclusion` is not allowed);
- as the `subtype_mark` in the `subtype_indication` of a `subtype_declaration`; the `subtype_indication` shall not have a `null_exclusion` or a constraint;
- as the `subtype_mark` in an `access_definition` for an access-to-object type;
- as the `subtype_mark` defining the subtype of a parameter or result in a profile occurring within a `basic_declaration`;
- as a generic actual parameter whose corresponding generic formal parameter is a formal incomplete type (see 15.5.1).

If such a name denotes a tagged incomplete view, it may also be used:

- as the `subtype_mark` defining the subtype of a parameter in the profile for a `subprogram_body`, `entry_body`, or `accept_statement`;
- as the prefix of an `attribute_reference` whose `attribute_designator` is `Class`; such an `attribute_reference` is restricted to the uses allowed here; it denotes a tagged incomplete view.

If any of the above uses occurs as part of the declaration of a primitive subprogram of the incomplete view, and the declaration occurs immediately within the private part of a package, then the completion of the incomplete view shall also occur immediately within the private part; it shall not be deferred to the package body.

No other uses of a name that denotes an incomplete view of a type are allowed.

A prefix that denotes an object shall not be of an incomplete view. An actual parameter in a call shall not be of an untagged incomplete view. The result object of a function call shall not be of an incomplete view. A prefix shall not denote a subprogram having a formal parameter of an untagged incomplete view, nor a return type that is an incomplete view.

The controlling operand or controlling result of a dispatching call shall not be of an incomplete view if the operand or result is dynamically tagged.

#### Dynamic Semantics

The elaboration of an `incomplete_type_declaration` has no effect.

NOTE 1 Within a `declarative_part`, an `incomplete_type_declaration` and a corresponding `full_type_declaration` cannot be separated by an intervening body. This is because, by the rules given in 16.14, a type is illegal if it is not completely defined before it is frozen, and a body freezes all types declared prior to it in the same `declarative_part`.

NOTE 2 A name that denotes an object of an incomplete view is defined to be of a limited type. Hence, the target of an assignment statement cannot be of an incomplete view.

EXAMPLE 1 *Example of a recursive type:*

```
type Cell; -- incomplete type declaration
type Link is access Cell;
```

```

type Cell is
  record
    Value  : Integer;
    Succ   : Link;
    Pred   : Link;
  end record;

Head  : Link := new Cell'(0, null, null);
Next  : Link := Head.Succ;

```

EXAMPLE 2 *Examples of mutually dependent access types:*

```

type Person(<>);      -- incomplete type declaration
type Car is tagged;  -- incomplete type declaration

type Person_Name is access Person;
type Car_Name    is access all Car'Class;

type Car is tagged
  record
    Number  : Integer;
    Owner   : Person_Name;
  end record;

type Person(Sex : Gender) is
  record
    Name     : String(1 .. 20);
    Birth    : Date;
    Age      : Integer range 0 .. 130;
    Vehicle  : Car_Name;
    case Sex is
      when M => Wife           : Person_Name(Sex => F);
      when F => Husband       : Person_Name(Sex => M);
    end case;
  end record;

My_Car, Your_Car, Next_Car : Car_Name := new Car;  -- see 7.8
Casey : Person_Name := new Person(M);
...
Casey.Vehicle := Your_Car;

```

## 6.10.2 Operations of Access Types

The attribute `Access` is used to create access values designating aliased objects and nonintrinsic subprograms. The “accessibility” rules prevent dangling references (in the absence of uses of certain unchecked features — see Clause 16).

### *Name Resolution Rules*

For an `attribute_reference` with `attribute_designator` `Access` (or `Unchecked_Access` — see 16.10), the expected type shall be a single access type *A* such that:

- *A* is an access-to-object type with designated type *D* and the type of the `prefix` is *D*'Class or is covered by *D*, or
- *A* is an access-to-subprogram type whose designated profile is type conformant with that of the `prefix`.

The `prefix` of such an `attribute_reference` is never interpreted as an `implicit_dereference` or a `parameterless_function_call` (see 7.1.4). The designated type or profile of the expected type of the `attribute_reference` is the expected type or profile for the `prefix`.

### *Static Semantics*

The accessibility rules, which prevent dangling references, are written in terms of *accessibility levels*, which reflect the run-time nesting of *masters*. As explained in 10.6.1, a master is the execution of a certain construct (called a *master construct*), such as a `subprogram_body`. An accessibility level is

*deeper than* another if it is more deeply nested at run time. For example, an object declared local to a called subprogram has a deeper accessibility level than an object declared local to the calling subprogram. The accessibility rules for access types require that the accessibility level of an object designated by an access value be no deeper than that of the access type. This ensures that the object will live at least as long as the access type, which in turn ensures that the access value cannot later designate an object that no longer exists. The `Unchecked_Access` attribute may be used to circumvent the accessibility rules.

A given accessibility level is said to be *statically deeper* than another if the given level is known at compile time (as defined below) to be deeper than the other for all possible executions. In most cases, accessibility is enforced at compile time by Legality Rules. Run-time accessibility checks are also used, since the Legality Rules do not cover certain cases involving access parameters and generic packages.

Each master, and each entity and view created by it, has an accessibility level; when two levels are defined to be the same, the accessibility levels of the two associated entities are said to be *tied* to each other. Accessibility levels are defined as follows:

- The accessibility level of a given master is deeper than that of each dynamically enclosing master, and deeper than that of each master upon which the task executing the given master directly depends (see 12.3).
- An entity or view defined by a declaration and created as part of its elaboration has the same accessibility level as the innermost master of the declaration except in the cases of renaming and derived access types described below. A formal parameter of a callable entity has the same accessibility level as the master representing the invocation of the entity.
- The accessibility level of a view of an object or subprogram defined by a `renaming_declaration` is the same as that of the renamed view, unless the renaming is of a formal subprogram, in which case the accessibility level is that of the instance.
- The accessibility level of a view conversion, `qualified_expression`, or parenthesized expression, is the same as that of the operand.
- The accessibility level of a conditional expression (see 7.5.7) is the accessibility level of the evaluated *dependent\_expression*.
- The accessibility level of a `declare_expression` (see 7.5.9) is the accessibility level of the *body\_expression*.
- The accessibility level of an aggregate that is used (in its entirety) to directly initialize part of an object is that of the object being initialized. In other contexts, the accessibility level of an aggregate is that of the innermost master that evaluates the aggregate. Corresponding rules apply to a value conversion (see 7.6).
- The accessibility level of the result of a function call is that of the *master of the function call*, which is determined by the point of call as follows:
  - If the result type at the point of the function (or access-to-function type) declaration is a composite type, and the result is used (in its entirety) to directly initialize part of an object, the master is that of the object being initialized. In the case where the initialized object is a coextension (see below) that becomes a coextension of another object, the master is that of the eventual object to which the coextension will be transferred.
  - If the result is of an anonymous access type and is converted to a (named or anonymous) access type, the master is determined following the rules given below for determining the master of an object created by an allocator (even if the access result is of an access-to-subprogram type);
  - If the call itself defines the result of a function *F*, or has an accessibility level that is tied to the result of such a function *F*, then the master of the call is that of the master of the call invoking *F*;
  - In other cases, the master of the call is that of the innermost master that evaluates the function call.

In the case of a call to a function whose result type is an anonymous access type, the accessibility level of the type of the result of the function call is also determined by the point of call as described above.

- Within a return statement, the accessibility level of the return object is that of the execution of the return statement. If the return statement completes normally by returning from the function, then prior to leaving the function, the accessibility level of the return object changes to be a level determined by the point of call, as does the level of any coextensions (see below) of the return object.
- The accessibility level of a derived access type is the same as that of its ultimate ancestor.
- The accessibility level of the anonymous access type defined by an `access_definition` of an `object_renaming_declaration` is the same as that of the renamed view.
- The accessibility level of the anonymous access type defined by an `access_definition` of a `loop_parameter_subtype_indication` is that of the loop parameter.
- The accessibility level of the anonymous access type of an access discriminant in the `subtype_indication` or `qualified_expression` of an allocator, or in the `expression` or `return_subtype_indication` of a return statement is determined as follows:
  - If the value of the access discriminant is determined by a `discriminant_association` in a `subtype_indication`, the accessibility level of the object or subprogram designated by the associated value (or library level if the value is null);
  - If the value of the access discriminant is determined by a `default_expression` in the declaration of the discriminant, the level of the object or subprogram designated by the associated value (or library level if null);
  - If the value of the access discriminant is determined by a `record_component_association` in an aggregate, the accessibility level of the object or subprogram designated by the associated value (or library level if the value is null);
  - In other cases, where the value of the access discriminant is determined by an object with an unconstrained nominal subtype, the accessibility level of the object.
- The accessibility level of the anonymous access type of an access discriminant in any other context is that of the enclosing object.
- The accessibility level of the anonymous access type of an access parameter specifying an access-to-object type is the same as that of the view designated by the actual (or library-level if the actual is null).
- The accessibility level of the anonymous access type of an access parameter specifying an access-to-subprogram type is deeper than that of any master; all such anonymous access types have this same level.
- The accessibility level of the anonymous access subtype defined by a `return_subtype_indication` that is an `access_definition` (see 9.5) is that of the result subtype of the enclosing function.
- The accessibility level of the type of a stand-alone object of an anonymous access-to-object type is the same as the accessibility level of the type of the access value most recently assigned to the object; accessibility checks ensure that this is never deeper than that of the declaration of the stand-alone object.
- The accessibility level of an object created by an `allocator` is the same as that of the access type, except for an `allocator` of an anonymous access type (an *anonymous allocator*) in certain contexts, as follows: For an anonymous allocator that defines the result of a function with an access result, the accessibility level is determined as though the `allocator` were in place of the call of the function; in the special case of a call that is the operand of a type conversion, the level is that of the target access type of the conversion. For an anonymous allocator defining the value of an access

parameter, the accessibility level is that of the innermost master of the call. For an anonymous allocator whose type is that of a stand-alone object of an anonymous access-to-object type, the accessibility level is that of the declaration of the stand-alone object. For one defining an access discriminant, the accessibility level is determined as follows:

- for an `allocator` used to define the discriminant of an object, the level of the object;
- for an `allocator` used to define the constraint in a `subtype_indication` in any other context, the level of the master that elaborates the `subtype_indication`.

In the first case, the allocated object is said to be a *coextension* of the object whose discriminant designates it, as well as of any object of which the discriminated object is itself a coextension or subcomponent. If the allocated object is a coextension of an anonymous object representing the result of an aggregate or function call that is used (in its entirety) to directly initialize a part of an object, after the result is assigned, the coextension becomes a coextension of the object being initialized and is no longer considered a coextension of the anonymous object. All coextensions of an object (which have not thus been transferred by such an initialization) are finalized when the object is finalized (see 10.6.1).

- Within a return statement, the accessibility level of the anonymous access type of an access result is that of the master of the call.
- The accessibility level of a view of an object or subprogram designated by an access value is the same as that of the access type.
- The accessibility level of a component, protected subprogram, or entry of (a view of) a composite object is the same as that of (the view of) the composite object.

In the above rules, the operative constituents of a `name` or `expression` (see 7.4) are considered to be used in a given context if the enclosing `name` or `expression` is used in that context.

One accessibility level is defined to be *statically deeper* than another in the following cases:

- For a master construct that is statically nested within another master construct, the accessibility level of the inner master construct is statically deeper than that of the outer master construct.
- The accessibility level of the anonymous access type of an access parameter specifying an access-to-subprogram type is statically deeper than that of any master; all such anonymous access types have this same level.
- The statically deeper relationship does not apply to the accessibility level of the following:
  - the anonymous type of an access parameter specifying an access-to-object type;
  - the type of a stand-alone object of an anonymous access-to-object type;
  - a `raise_expression`;
  - a descendant of a generic formal type;
  - a descendant of a type declared in a generic formal package.
- When the statically deeper relationship does not apply, the accessibility level is not considered to be statically deeper, nor statically shallower, than any other.
- When within a function body or the return expression of an expression function, the accessibility level of the master representing an execution of the function is statically deeper than that of the master of the function call invoking that execution, independent of how the master of the function call is determined (see above).
- For determining whether one level is statically deeper than another when within a generic package body, the generic package is presumed to be instantiated at the same level as where it was declared; runtime checks are required in the case of more deeply nested instantiations.



- For determining whether one level is statically deeper than another when within the declarative region of a `type_declaration`, the current instance of the type is presumed to be an object created at a deeper level than that of the type.

Notwithstanding other rules given above, the accessibility level of an entity that is tied to that of an explicitly aliased formal parameter of an enclosing function is considered (both statically and dynamically) to be the same as that of an entity whose accessibility level is tied to that of the return object of that function.

The accessibility level of all library units is called the *library level*; a library-level declaration or entity is one whose accessibility level is the library level.

The following attribute is defined for a `prefix` `X` that denotes an aliased view of an object:

**X'Access** X'Access yields an access value that designates the object denoted by `X`. The type of X'Access is an access-to-object type, as determined by the expected type. The expected type shall be a general access type. `X` shall denote an aliased view of an object, including possibly the current instance (see 11.6) of a limited type within its definition, or a formal parameter or generic formal object of a tagged type. The view denoted by the `prefix` `X` shall satisfy the following additional requirements, presuming the expected type for X'Access is the general access type `A` with designated type `D`:

- If `A` is an access-to-variable type, then the view shall be a variable; on the other hand, if `A` is an access-to-constant type, the view may be either a constant or a variable.
- The view shall not be a subcomponent that depends on discriminants of an object unless the object is known to be constrained.
- If `A` is a named access type and `D` is a tagged type, then the type of the view shall be covered by `D`; if `A` is anonymous and `D` is tagged, then the type of the view shall be either `D'Class` or a type covered by `D`; if `D` is untagged, then the type of the view shall be `D`, and either:
  - the designated subtype of `A` shall statically match the nominal subtype of the view; or
  - `D` shall be discriminated in its full view and unconstrained in any partial view, and the designated subtype of `A` shall be unconstrained.
- The accessibility level of the view shall not be statically deeper than that of the access type `A`.

In addition to the places where Legality Rules normally apply (see 15.3), these requirements apply also in the private part of an instance of a generic unit.

A check is made that the accessibility level of `X` is not deeper than that of the access type `A`. If this check fails, `Program_Error` is raised.

If the nominal subtype of `X` does not statically match the designated subtype of `A`, a view conversion of `X` to the designated subtype is evaluated (which can raise `Constraint_Error` — see 7.6) and the value of X'Access designates that view.

The following attribute is defined for a `prefix` `P` that denotes a subprogram:

**P'Access** P'Access yields an access value that designates the subprogram denoted by `P`. The type of P'Access is an access-to-subprogram type (`S`), as determined by the expected type. The accessibility level of `P` shall not be statically deeper than that of `S`. If `S` is nonblocking, `P` shall be nonblocking. In addition to the places where Legality Rules normally apply (see 15.3), these rules apply also in the private part of an instance of a generic unit. The profile of `P` shall be subtype conformant with the designated profile of `S`, and shall not be *Intrinsic*. If the subprogram denoted by `P` is declared within a generic unit, and the expression P'Access occurs within the body of that generic unit or within the body of a generic unit declared



within the declarative region of the generic unit, then the ultimate ancestor of *S* shall be either a nonformal type declared within the generic unit or an anonymous access type of an access parameter.

#### Legality Rules

An expression is said to have *distributed accessibility* if it is

- a `conditional_expression` (see 7.5.7); or
- a `declare_expression` (see 7.5.9) whose *body\_expression* has distributed accessibility; or
- a view conversion, `qualified_expression`, or parenthesized expression whose operand has distributed accessibility.

The statically deeper relationship does not apply to the accessibility level of an expression having distributed accessibility; that is, such an accessibility level is not considered to be statically deeper, nor statically shallower, than any other.

Any static accessibility requirement that is imposed on an expression that has distributed accessibility (or on its type) is instead imposed on the *dependent\_expressions* of the underlying `conditional_expression`. This rule is applied recursively if a *dependent\_expression* also has distributed accessibility.

NOTE 1 The `Unchecked_Access` attribute yields the same result as the `Access` attribute for objects, but has fewer restrictions (see 16.10). There are other predefined operations that yield access values: an allocator can be used to create an object, and return an access value that designates it (see 7.8); evaluating the literal **null** yields a null access value that designates no entity at all (see 7.2).

NOTE 2 The predefined operations of an access type also include the assignment operation, qualification, and membership tests. Explicit conversion is allowed between general access types with matching designated subtypes; explicit conversion is allowed between access-to-subprogram types with subtype conformant profiles (see 7.6). Named access types have predefined equality operators; anonymous access types do not, but they can use the predefined equality operators for *universal\_access* (see 7.5.2).

NOTE 3 The object or subprogram designated by an access value can be named with a dereference, either an `explicit_dereference` or an `implicit_dereference`. See 7.1.

NOTE 4 A call through the dereference of an access-to-subprogram value is never a dispatching call.

NOTE 5 The `Access` attribute for subprograms and parameters of an anonymous access-to-subprogram type can be used together to implement “downward closures” — that is, to pass a more-nested subprogram as a parameter to a less-nested subprogram, as can be appropriate for an iterator abstraction or numerical integration. Downward closures can also be implemented using generic formal subprograms (see 15.6). Unlike for objects, there is no `Unchecked_Access` attribute for subprograms.

NOTE 6 Using an access-to-class-wide tagged type with a dispatching operation is a potentially more structured alternative to using an access-to-subprogram type.

NOTE 7 An implementation can consider two access-to-subprogram values to be unequal, even though they designate the same subprogram. For instance, this can happen because one points directly to the subprogram, while the other points to a special prologue that performs an `Elaboration_Check` and then jumps to the subprogram. See 7.5.2.

EXAMPLE *Example of use of the Access attribute:*

```
Becky : Person_Name := new Person(F);           -- see 6.10.1
Cars  : array (1..2) of aliased Car;
...
Becky.Vehicle := Cars(1)'Access;
Casey.Vehicle := Cars(2)'Access;
```

## 6.11 Declarative Parts

A `declarative_part` contains `declarative_items` (possibly none).

### Syntax

```

declarative_part ::= {declarative_item}
declarative_item ::=
  basic_declarative_item | body
basic_declarative_item ::=
  basic_declaration | aspect_clause | use_clause
body ::= proper_body | body_stub
proper_body ::=
  subprogram_body | package_body | task_body | protected_body

```

### Static Semantics

The list of `declarative_items` of a `declarative_part` is called the *declaration list* of the `declarative_part`.

### Dynamic Semantics

The elaboration of a `declarative_part` consists of the elaboration of the `declarative_items`, if any, in the order in which they are given in the `declarative_part`.

An elaborable construct is in the *elaborated* state after the normal completion of its elaboration. Prior to that, it is *not yet elaborated*.

For a construct that attempts to use a body, a check (Elaboration\_Check) is performed, as follows:

- For a call to a (non-protected) subprogram that has an explicit body, a check is made that the body is already elaborated. This check and the evaluations of any actual parameters of the call are done in an arbitrary order.
- For a call to a protected operation of a protected type (that has a body — no check is performed if the protected type is imported — see B.1), a check is made that the `protected_body` is already elaborated. This check and the evaluations of any actual parameters of the call are done in an arbitrary order.
- For the activation of a task, a check is made by the activator that the `task_body` is already elaborated. If two or more tasks are being activated together (see 12.2), as the result of the elaboration of a `declarative_part` or the initialization for the object created by an allocator, this check is done for all of them before activating any of them.
- For the instantiation of a generic unit that has a body, a check is made that this body is already elaborated. This check and the evaluation of any `explicit_generic_actual_parameters` of the instantiation are done in an arbitrary order.

The exception `Program_Error` is raised if any of these checks fails.

### 6.11.1 Completions of Declarations

Declarations sometimes come in two parts. A declaration that requires a second part is said to *require completion*. The second part is called the *completion* of the declaration (and of the entity declared), and is either another declaration, a body, or a pragma. A *body* is a body, an `entry_body`, a `null_procedure_declaration` or an `expression_function_declaration` that completes another declaration, or a renaming-as-body (see 11.5.4).

*Name Resolution Rules*

A construct that can be a completion is interpreted as the completion of a prior declaration only if:

- The declaration and the completion occur immediately within the same declarative region;
- The defining name or `defining_program_unit_name` in the completion is the same as in the declaration, or in the case of a `pragma`, the `pragma` applies to the declaration;
- If the declaration is overloadable, then the completion either has a type-conformant profile, or is a `pragma`.

*Legality Rules*

An implicit declaration shall not have a completion. For any explicit declaration that is specified to *require completion*, there shall be a corresponding explicit completion, unless the declared entity is imported (see B.1).

At most one completion is allowed for a given declaration. Additional requirements on completions appear where each kind of completion is defined.

A type is *completely defined* at a place that is after its full type definition (if it has one) and after all of its subcomponent types are completely defined. A type shall be completely defined before it is frozen (see 16.14 and 10.3).

NOTE 1 Completions are in principle allowed for any kind of explicit declaration. However, for some kinds of declaration, the only allowed completion is an implementation-defined `pragma`, and implementations are not required to have any such `pragmas`.

NOTE 2 There are rules that prevent premature uses of declarations that have a corresponding completion. The Elaboration Checks of 6.11 prevent such uses at run time for subprograms, protected operations, tasks, and generic units. The freezing rules (see 16.14) prevent, at compile time, premature uses of other entities such as private types and deferred constants.

## 7 Names and Expressions

The rules applicable to the different forms of name and expression, and to their evaluation, are given in this clause.

### 7.1 Names

Names can denote declared entities, whether declared explicitly or implicitly (see 6.1). Names can also denote objects or subprograms designated by access values; the results of `type_conversions` or `function_calls`; subcomponents and slices of objects and values; protected subprograms, single entries, entry families, and entries in families of entries. Finally, names can denote attributes of any of the foregoing.

*Syntax*

```

name ::=
    direct_name           | explicit_dereference
  | indexed_component     | slice
  | selected_component    | attribute_reference
  | type_conversion       | function_call
  | character_literal     | qualified_expression
  | generalized_reference | generalized_indexing
  | target_name

direct_name ::= identifier | operator_symbol

prefix ::= name | implicit_dereference

explicit_dereference ::= name.all
implicit_dereference ::= name

```

Certain forms of name (`indexed_components`, `selected_components`, `slices`, and `attribute_references`) include a prefix that is either itself a name that denotes some related entity, or an `implicit_dereference` of an access value that designates some related entity.

#### Name Resolution Rules

The name in a *dereference* (either an `implicit_dereference` or an `explicit_dereference`) is expected to be of any access type.

#### Static Semantics

If the type of the name in a dereference is some access-to-object type *T*, then the dereference denotes a view of an object, the *nominal subtype* of the view being the designated subtype of *T*. If the designated subtype has unconstrained discriminants, the (actual) subtype of the view is constrained by the values of the discriminants of the designated object, except when there is a partial view of the type of the designated subtype that does not have discriminants, in which case the dereference is not constrained by its discriminant values.

If the type of the name in a dereference is some access-to-subprogram type *S*, then the dereference denotes a view of a subprogram, the *profile* of the view being the designated profile of *S*.

#### Dynamic Semantics

The evaluation of a name determines the entity denoted by the name. This evaluation has no other effect for a name that is a `direct_name` or a `character_literal`.

The evaluation of a name that has a prefix includes the evaluation of the prefix. The evaluation of a prefix consists of the evaluation of the name or the `implicit_dereference`. The prefix denotes the entity denoted by the name or the `implicit_dereference`.

The evaluation of a dereference consists of the evaluation of the name and the determination of the object or subprogram that is designated by the value of the name. A check is made that the value of the name is not the null access value. `Constraint_Error` is raised if this check fails. The dereference denotes the object or subprogram designated by the value of the name.

#### EXAMPLE 1 Examples of direct names:

```
Pi      -- the direct name of a number (see 6.3.2)
Limit   -- the direct name of a constant (see 6.3.1)
Count   -- the direct name of a scalar variable (see 6.3.1)
Board   -- the direct name of an array variable (see 6.6.1)
Matrix  -- the direct name of a type (see 6.6)
Random  -- the direct name of a function (see 9.1)
Error   -- the direct name of an exception (see 14.1)
```

#### EXAMPLE 2 Examples of dereferences:

```
Next_Car.all    -- explicit dereference denoting the object designated
                -- by the access variable Next_Car (see 6.10.1)
Next_Car.Owner  -- selected component with implicit dereference;
                -- same as Next_Car.all.Owner
```

### 7.1.1 Indexed Components

An `indexed_component` denotes either a component of an array or an entry in a family of entries.

#### Syntax

```
indexed_component ::= prefix(expression {expression})
```

#### Name Resolution Rules

The prefix of an `indexed_component` with a given number of expressions shall resolve to denote an array (after any implicit dereference) with the corresponding number of index positions, or

shall resolve to denote an entry family of a task or protected object (in which case there shall be only one expression).

The expected type for each expression is the corresponding index type.

#### Static Semantics

When the `prefix` denotes an array, the `indexed_component` denotes the component of the array with the specified index value(s). The nominal subtype of the `indexed_component` is the component subtype of the array type.

When the `prefix` denotes an entry family, the `indexed_component` denotes the individual entry of the entry family with the specified index value.

#### Dynamic Semantics

For the evaluation of an `indexed_component`, the `prefix` and the expressions are evaluated in an arbitrary order. The value of each expression is converted to the corresponding index type. A check is made that each index value belongs to the corresponding index range of the array or entry family denoted by the `prefix`. `Constraint_Error` is raised if this check fails.

EXAMPLE *Examples of indexed components:*

```
My_Schedule(Sat)  -- a component of a one-dimensional array (see 6.6.1)
Page(10)          -- a component of a one-dimensional array (see 6.6)
Board(M, J + 1)   -- a component of a two-dimensional array (see 6.6.1)
Page(10)(20)      -- a component of a component (see 6.6)
Request(Medium)   -- an entry in a family of entries (see 12.1)
Next_Frame(L)(M, N) -- a component of a function call (see 9.1)
```

Distinct notations are used for components of multidimensional arrays (such as `Board`) and arrays of arrays (such as `Page`). The components of an array of arrays are arrays and can therefore be indexed. Thus `Page(10)(20)` denotes the 20th component of `Page(10)`. In the last example `Next_Frame(L)` is a function call returning an access value that designates a two-dimensional array.

### 7.1.2 Slices

A `slice` denotes a one-dimensional array formed by a sequence of consecutive components of a one-dimensional array. A `slice` of a variable is a variable; a `slice` of a constant is a constant; a `slice` of a value is a value.

#### Syntax

```
slice ::= prefix(discrete_range)
```

#### Name Resolution Rules

The `prefix` of a `slice` shall resolve to denote a one-dimensional array (after any implicit dereference).

The expected type for the `discrete_range` of a `slice` is the index type of the array type.

#### Static Semantics

A `slice` denotes a one-dimensional array formed by the sequence of consecutive components of the array denoted by the `prefix`, corresponding to the range of values of the index given by the `discrete_range`.

The type of the `slice` is that of the `prefix`. Its bounds are those defined by the `discrete_range`.

#### Dynamic Semantics

For the evaluation of a `slice`, the `prefix` and the `discrete_range` are evaluated in an arbitrary order. If the `slice` is not a *null slice* (a `slice` where the `discrete_range` is a null range), then a

check is made that the bounds of the `discrete_range` belong to the index range of the array denoted by the `prefix`. `Constraint_Error` is raised if this check fails.

NOTE 1 By the rules given in 6.10.2, a `slice` is illegal as the `prefix` of an `Access attribute_reference`, even if the components or the array as a whole are aliased.

NOTE 2 For a one-dimensional array `A`, the `slice` `A(N .. N)` denotes an array that has only one component; its type is the type of `A`. On the other hand, `A(N)` denotes a component of the array `A` and has the corresponding component type.

EXAMPLE *Examples of slices:*

```
Stars(1 .. 15)           -- a slice of 15 characters           (see 6.6.3)
Page(10 .. 10 + Size)   -- a slice of 1 + Size components   (see 6.6)
Page(L) (A .. B)        -- a slice of the array Page(L)      (see 6.6)
Stars(1 .. 0)           -- a null slice                     (see 6.6.3)
My_Schedule(Weekday)    -- bounds given by subtype (see 6.6.1 and 6.5.1)
Stars(5 .. 15) (K)      -- same as Stars(K)                 (see 6.6.3)
                        -- provided that K is in 5 .. 15
```

### 7.1.3 Selected Components

`Selected_components` are used to denote components (including discriminants), entries, entry families, and protected subprograms; they are also used as expanded names as described below.

#### Syntax

```
selected_component ::= prefix.selector_name
selector_name ::= identifier | character_literal | operator_symbol
```

#### Name Resolution Rules

A `selected_component` is called an *expanded name* if, according to the visibility rules, at least one possible interpretation of its `prefix` denotes a package or an enclosing named construct (directly, not through a `subprogram_renaming_declaration` or `generic_renaming_declaration`).

A `selected_component` that is not an expanded name shall resolve to denote one of the following:

- A component (including a discriminant):

The `prefix` shall resolve to denote an object or value of some non-array composite type (after any implicit dereference). The `selector_name` shall resolve to denote a discriminant specification of the type, or, unless the type is a protected type, a component declaration of the type. The `selected_component` denotes the corresponding component of the object or value.

- A single entry, an entry family, or a protected subprogram:

The `prefix` shall resolve to denote an object or value of some task or protected type (after any implicit dereference). The `selector_name` shall resolve to denote an `entry_declaration` or `subprogram_declaration` occurring (implicitly or explicitly) within the visible part of that type. The `selected_component` denotes the corresponding entry, entry family, or protected subprogram.

- A view of a subprogram whose first formal parameter is of a tagged type or is an access parameter whose designated type is tagged:

The `prefix` (after any implicit dereference) shall resolve to denote an object or value of a specific tagged type `T` or class-wide type `T'Class`. The `selector_name` shall resolve to denote a view of a subprogram declared immediately within the declarative region in which an ancestor of the type `T` is declared. The first formal parameter of the subprogram shall be of type `T`, or a class-wide type that covers `T`, or an access parameter designating one of these types. The designator of the subprogram shall not be the same as that of a component of the tagged type visible at the point of the `selected_component`. The subprogram shall not be an implicitly declared primitive



operation of type *T* that overrides an inherited subprogram implemented by an entry or protected subprogram visible at the point of the `selected_component`. The `selected_component` denotes a view of this subprogram that omits the first formal parameter. This view is called a *prefixed view* of the subprogram, and the `prefix` of the `selected_component` (after any implicit dereference) is called the *prefix* of the prefixed view.

An expanded name shall resolve to denote a declaration that occurs immediately within a named declarative region, as follows:

- The `prefix` shall resolve to denote either a package (including the current instance of a generic package, or a rename of a package), or an enclosing named construct.
- The `selector_name` shall resolve to denote a declaration that occurs immediately within the declarative region of the package or enclosing construct (the declaration shall be visible at the place of the expanded name — see 11.3). The expanded name denotes that declaration.
- If the `prefix` does not denote a package, then it shall be a `direct_name` or an expanded name, and it shall resolve to denote a program unit (other than a package), the current instance of a type, a block\_statement, a loop\_statement, or an accept\_statement (in the case of an accept\_statement or entry\_body, no family index is allowed); the expanded name shall occur within the declarative region of this construct. Further, if this construct is a callable construct and the `prefix` denotes more than one such enclosing callable construct, then the expanded name is ambiguous, independently of the `selector_name`.

#### Legality Rules

For a prefixed view of a subprogram whose first formal parameter is an access parameter, the `prefix` shall be legal as the `prefix` of an attribute\_reference with attribute\_designator `Access` appearing as the first actual parameter in a call on the unprefix view of the subprogram.

For a subprogram whose first parameter is of mode **in out** or **out**, or of an anonymous access-to-variable type, the `prefix` of any prefixed view shall denote a variable.

#### Dynamic Semantics

The evaluation of a `selected_component` includes the evaluation of the `prefix`.

For a `selected_component` that denotes a component of a variant, a check is made that the values of the discriminants are such that the value or object denoted by the `prefix` has this component. The exception `Constraint_Error` is raised if this check fails.

#### EXAMPLE 1 Examples of selected components:

```

Tomorrow.Month      -- a record component           (see 6.8)
Next_Car.Owner      -- a record component           (see 6.10.1)
Next_Car.Owner.Age  -- a record component           (see 6.10.1)
-- the previous two lines involve implicit dereferences
Writer.Unit         -- a record component (a discriminant) (see 6.8.1)
Min_Cell(H).Value   -- a record component of the result (see 9.1)
-- of the function call Min_Cell(H)
Cashier.Append      -- a prefixed view of a procedure (see 6.9.4)
Control.Seize       -- an entry of a protected object (see 12.4)
Pool(K).Write       -- an entry of the task Pool(K)   (see 12.1)

```

#### EXAMPLE 2 Examples of expanded names:

```

Key_Manager."<"      -- an operator of the visible part of a
-- package           (see 10.3.1)
Dot_Product.Sum      -- a variable declared in a function body (see 9.1)
Buffer.Pool          -- a variable declared in a protected unit (see 12.11)
Buffer.Read          -- an entry of a protected unit (see 12.11)
Swap.Temp            -- a variable declared in a block statement (see 8.6)
Standard.Boolean     -- the name of a predefined type (see A.1)

```

### 7.1.4 Attributes

An *attribute* is a characteristic of an entity that can be queried via an *attribute\_reference* or a *range\_attribute\_reference*.

#### Syntax

```
attribute_reference ::=
    prefix'attribute_designator
    | reduction_attribute_reference
attribute_designator ::=
    identifier[(static_expression)]
    | Access | Delta | Digits | Mod
range_attribute_reference ::= prefix'range_attribute_designator
range_attribute_designator ::= Range[(static_expression)]
```

#### Name Resolution Rules

In an *attribute\_reference* that is not a *reduction\_attribute\_reference*, if the *attribute\_designator* is for an attribute defined for (at least some) objects of an access type, then the prefix is never interpreted as an *implicit\_dereference*; otherwise (and for all *range\_attribute\_references* and *reduction\_attribute\_references*), if there is a prefix and the type of the name within the prefix is of an access type, the prefix is interpreted as an *implicit\_dereference*. Similarly, if the *attribute\_designator* is for an attribute defined for (at least some) functions, then the prefix is never interpreted as a *parameterless\_function\_call*; otherwise (and for all *range\_attribute\_references* and *reduction\_attribute\_references*), if there is a prefix and the prefix consists of a name that denotes a function, it is interpreted as a *parameterless\_function\_call*.

The expression, if any, in an *attribute\_designator* or *range\_attribute\_designator* is expected to be of any integer type.

#### Legality Rules

The expression, if any, in an *attribute\_designator* or *range\_attribute\_designator* shall be static.

#### Static Semantics

An *attribute\_reference* denotes a value, an object, a subprogram, or some other kind of program entity. Unless explicitly specified otherwise, for an *attribute\_reference* that denotes a value or an object, if its type is scalar, then its nominal subtype is the base subtype of the type; if its type is tagged, its nominal subtype is the first subtype of the type; otherwise, its nominal subtype is a subtype of the type without any constraint, *null\_exclusion*, or predicate. Similarly, unless explicitly specified otherwise, for an *attribute\_reference* that denotes a function, when its result type is scalar, its result subtype is the base subtype of the type, when its result type is tagged, the result subtype is the first subtype of the type, and when the result type is some other type, the result subtype is a subtype of the type without any constraint, *null\_exclusion*, or predicate.

A *range\_attribute\_reference* *X'Range(N)* is equivalent to the range *X'First(N) .. X'Last(N)*, except that the prefix is only evaluated once. Similarly, *X'Range* is equivalent to *X'First .. X'Last*, except that the prefix is only evaluated once.

#### Dynamic Semantics

The evaluation of a *range\_attribute\_reference* or an *attribute\_reference* that is not a *reduction\_attribute\_reference* consists of the evaluation of the prefix. The evaluation of a *reduction\_attribute\_reference* is defined in 7.5.10.

*Implementation Permissions*

An implementation may provide implementation-defined attributes; the identifier for such an implementation-defined attribute shall differ from those of the language-defined attributes.

An implementation may extend the definition of a language-defined attribute by accepting uses of that attribute that would otherwise be illegal in the following cases:

- in order to support compatibility with a previous edition of of this document; or
- in the case of a language-defined attribute whose `prefix` is required by this document to be a floating point subtype, an implementation may accept an `attribute_reference` whose `prefix` is a fixed point subtype; the semantics of such an `attribute_reference` are implementation defined.

NOTE 1 Attributes are defined throughout this document, and are summarized in J.2.

NOTE 2 By the general rules given above, there is no expected type or profile for the name in a `prefix` of an `attribute_reference` (or a `range_attribute_reference`), which means that no context can be used to resolve the name. However, by the rules given in 6.10.2 for the case of the Access attribute, the expected type for the `attribute_reference` will be a single access type, and the resolution of the name can make use of the designated type or profile of this access type.

EXAMPLE *Examples of attributes:*

```
Color'First      -- minimum value of the enumeration type Color      (see
6.5.1)
Rainbow'Base'First -- same as Color'First                          (see 6.5.1)
Real'Digits      -- precision of the type Real                      (see 6.5.7)
Board'Last(2)    -- upper bound of the second dimension of Board    (see
6.6.1)
Board'Range(1)   -- index range of the first dimension of Board     (see
6.6.1)
Pool(K)'Terminated -- True if task Pool(K) is terminated (see 12.1)
Date'Size        -- number of bits for records of type Date        (see 6.8)
Message'Address  -- address of the record variable Message
--                                                         (see 6.7.1)
```

### 7.1.5 User-Defined References

*Static Semantics*

Given a discriminated type *T*, the following type-related operational aspect may be specified:

**Implicit\_Dereference**

This aspect is specified by a name that denotes an access discriminant declared for the type *T*.

A (view of a) type with a specified `Implicit_Dereference` aspect is a *reference type*. A *reference object* is an object of a reference type. The discriminant named by the `Implicit_Dereference` aspect is the *reference discriminant* of the reference type or reference object. A *generalized\_reference* is a name that identifies a reference object, and denotes the object or subprogram designated by the reference discriminant of the reference object.

*Syntax*

```
generalized_reference ::= reference_object_name
```

*Name Resolution Rules*

The expected type for the *reference\_object\_name* in a *generalized\_reference* is any reference type.

*Static Semantics*

The `Implicit_Dereference` aspect is nonoverridable (see 16.1.1).

A `generalized_reference` denotes a view equivalent to that of a dereference of the reference discriminant of the reference object.

Given a reference type  $T$ , the `Implicit_Dereference` aspect is inherited by descendants of type  $T$  if not overridden (which is only permitted if confirming). If a descendant type constrains the value of the reference discriminant of  $T$  by a new discriminant, that new discriminant is the reference discriminant of the descendant. If the descendant type constrains the value of the reference discriminant of  $T$  by an expression other than the name of a new discriminant, a `generalized_reference` that identifies an object of the descendant type denotes the object or subprogram designated by the value of this constraining expression.

#### Dynamic Semantics

The evaluation of a `generalized_reference` consists of the evaluation of the `reference_object_name` and a determination of the object or subprogram designated by the reference discriminant of the named reference object. A check is made that the value of the reference discriminant is not the null access value. `Constraint_Error` is raised if this check fails. The `generalized_reference` denotes the object or subprogram designated by the value of the reference discriminant of the named reference object.

EXAMPLE *Examples of the specification and use of generalized references:*

```

type Barrel is tagged ... -- holds objects of type Element
type Ref_Element(Data : access Element) is limited private
  with Implicit_Dereference => Data;
  -- This Ref_Element type is a "reference" type.
  -- "Data" is its reference discriminant.

function Find (B : aliased in out Barrel; Key : String) return Ref_Element;
  -- Returns a reference to an element of a barrel.

B: aliased Barrel;

...

Find (B, "grape") := Element'(...); -- Assign through a reference.
-- This is equivalent to:
Find (B, "grape").Data.all := Element'(...);

```

### 7.1.6 User-Defined Indexing

#### Static Semantics

Given a tagged type  $T$ , the following type-related, operational aspects may be specified:

##### Constant\_Indexing

This aspect shall be specified by a `name` that denotes one or more functions declared immediately within the same declaration list in which  $T$ , or the declaration completed by  $T$ , is declared. All such functions shall have at least two parameters, the first of which is of type  $T$  or  $T^{\text{Class}}$ , or is an access-to-constant parameter with designated type  $T$  or  $T^{\text{Class}}$ .

##### Variable\_Indexing

This aspect shall be specified by a `name` that denotes one or more functions declared immediately within the same declaration list in which  $T$ , or the declaration completed by  $T$ , is declared. All such functions shall have at least two parameters, the first of which is of type  $T$  or  $T^{\text{Class}}$ , or is an access parameter with designated type  $T$  or  $T^{\text{Class}}$ . All such functions shall have a return type that is a reference type (see 7.1.5), whose reference discriminant is of an access-to-variable type.

These aspects are inherited by descendants of  $T$  (including the class-wide type  $T^{\text{Class}}$ ).

An *indexable container type* is (a view of) a tagged type with at least one of the aspects `Constant_Indexing` or `Variable_Indexing` specified. An *indexable container object* is an object of an

indexable container type. A `generalized_indexing` is a name that denotes the result of calling a function named by a `Constant_Indexing` or `Variable_Indexing` aspect.

The `Constant_Indexing` and `Variable_Indexing` aspects are nonoverridable (see 16.1.1).

#### Legality Rules

If an ancestor of a type *T* is an indexable container type, then any explicit specification of the `Constant_Indexing` or `Variable_Indexing` aspects shall be confirming; that is, the specified name shall match the inherited aspect (see 16.1.1).

In addition to the places where Legality Rules normally apply (see 15.3), this rule applies also in the private part of an instance of a generic unit.

A `generalized_indexing` is illegal if the equivalent prefixed view (see below) is illegal.

#### Syntax

```
generalized_indexing ::= indexable_container_object_prefix actual_parameter_part
```

#### Name Resolution Rules

The expected type for the `indexable_container_object_prefix` of a `generalized_indexing` is any indexable container type.

If the `Constant_Indexing` aspect is specified for the type of the `indexable_container_object_prefix` of a `generalized_indexing`, then the `generalized_indexing` is interpreted as a *constant indexing* under the following circumstances:

- when the `Variable_Indexing` aspect is not specified for the type of the `indexable_container_object_prefix`;
- when the `indexable_container_object_prefix` denotes a constant;
- when the `generalized_indexing` is used within a primary where a name denoting a constant is permitted.

Otherwise, the `generalized_indexing` is interpreted as a *variable indexing*.

When a `generalized_indexing` is interpreted as a constant (or variable) indexing, it is equivalent to a call on a prefixed view of one of the functions named by the `Constant_Indexing` (or `Variable_Indexing`) aspect of the type of the `indexable_container_object_prefix` with the given `actual_parameter_part`, and with the `indexable_container_object_prefix` as the prefix of the prefixed view.

NOTE The `Constant_Indexing` and `Variable_Indexing` aspects cannot be redefined when inherited for a derived type, but the functions that they denote can be modified by overriding or overloading.

EXAMPLE Examples of the specification and use of generalized indexing:

```
type Indexed_Barrel is tagged ...
  with Variable_Indexing => Find;
  -- Indexed_Barrel is an indexable container type,
  -- Find is the generalized indexing operation.

function Find (B : aliased in out Indexed_Barrel; Key : String) return
  Ref_Element;
  -- Return a reference to an element of a barrel (see 7.1.5).

IB: aliased Indexed_Barrel;
```

```

-- All of the following calls are then equivalent:
Find (IB,"pear").Data.all := Element'(...); -- Traditional call
IB.Find ("pear").Data.all := Element'(...); -- Call of prefixed view
IB.Find ("pear")          := Element'(...); -- Implicit dereference (see
7.1.5)
IB      ("pear")          := Element'(...); -- Implicit indexing and
dereference
IB      ("pear").Data.all := Element'(...); -- Implicit indexing only

```

## 7.2 Literals

A *literal* represents a value literally, that is, by means of notation suited to its kind. A literal is either a `numeric_literal`, a `character_literal`, the literal **null**, or a `string_literal`.

### Name Resolution Rules

For a name that consists of a `character_literal`, either its expected type shall be a single character type, in which case it is interpreted as a parameterless `function_call` that yields the corresponding value of the character type, or its expected profile shall correspond to a parameterless function with a character result type, in which case it is interpreted as the name of the corresponding parameterless function declared as part of the character type's definition (see 6.5.1). In either case, the `character_literal` denotes the `enumeration_literal` specification.

The expected type for a primary that is a `string_literal` shall be a single string type or a type with a specified `String_Literal` aspect (see 7.2.1). In either case, the `string_literal` is interpreted to be of its expected type. If the expected type of an integer literal is a type with a specified `Integer_Literal` aspect (see 7.2.1), the literal is interpreted to be of its expected type; otherwise it is interpreted to be of type *universal\_integer*. If the expected type of a real literal is a type with a specified `Real_Literal` aspect (see 7.2.1), it is interpreted to be of its expected type; otherwise, it is interpreted to be of type *universal\_real*.

### Legality Rules

A `character_literal` that is a name shall correspond to a `defining_character_literal` of the expected type, or of the result type of the expected profile.

If the expected type for a `string_literal` is a string type, then for each character of the `string_literal` there shall be a corresponding `defining_character_literal` of the component type of the expected string type.

### Static Semantics

The literal **null** is of type *universal\_access*.

### Dynamic Semantics

If its expected type is a numeric type, the evaluation of a numeric literal yields the represented value. In other cases, the effect of evaluating a numeric literal is determined by the `Integer_Literal` or `Real_Literal` aspect that applies (see 7.2.1).

The evaluation of the literal **null** yields the null value of the expected type.

The evaluation of a `string_literal` that is a primary and has an expected type that is a string type, yields an array value containing the value of each character of the sequence of characters of the `string_literal`, as defined in 5.6. The bounds of this array value are determined according to the rules for `positional_array_aggregates` (see 7.3.3), except that for a null string literal, the upper bound is the predecessor of the lower bound. In other cases, the effect of evaluating a `string_literal` is determined by the `String_Literal` aspect that applies (see 7.2.1).

For the evaluation of a `string_literal` of a string type *T*, a check is made that the value of each character of the `string_literal` belongs to the component subtype of *T*. For the evaluation of a null



string literal of a string type, a check is made that its lower bound is greater than the lower bound of the base range of the index type. The exception `Constraint_Error` is raised if either of these checks fails.

NOTE Enumeration literals that are identifiers rather than `character_literals` follow the normal rules for identifiers when used in a name (see 7.1 and 7.1.3). `Character_literals` used as `selector_names` follow the normal rules for expanded names (see 7.1.3).

EXAMPLE *Examples of literals:*

```
3.14159_26536  -- a real literal
1_345          -- an integer literal
'A'            -- a character literal
"Some Text"    -- a string literal
```

### 7.2.1 User-Defined Literals

Using one or more of the aspects defined below, a type may be specified to allow the use of one or more kinds of literals as values of the type.

#### Static Semantics

The following type-related operational aspects (collectively known as *user-defined literal aspects*) may be specified for a type *T*:

#### Integer\_Literal

This aspect is specified by a *function\_name* that statically denotes a function with a result type of *T* and one **in** parameter that is of type `String` and is not explicitly aliased.

#### Real\_Literal

This aspect is specified by a *function\_name* that statically denotes a function with a result type of *T* and one **in** parameter that is of type `String` and is not explicitly aliased, and optionally a second function (overloading the first) with a result type of *T* and two **in** parameters of type `String` that are not explicitly aliased.

#### String\_Literal

This aspect is specified by a *function\_name* that statically denotes a function with a result type of *T* and one **in** parameter that is of type `Wide_Wide_String` and is not explicitly aliased.

User-defined literal aspects are nonoverridable (see 16.1.1).

When a numeric literal is interpreted as a value of a non-numeric type *T* or a `string_literal` is interpreted as a value of a type *T* that is not a string type (see 7.2), it is equivalent to a call to the subprogram denoted by the corresponding aspect of *T*: the `Integer_Literal` aspect for an integer literal, the `Real_Literal` aspect for a real literal, and the `String_Literal` aspect for a `string_literal`. The actual parameter of this notional call is a `string_literal` representing a sequence of characters that is the same as the sequence of characters in the original numeric literal, or the sequence represented by the original string literal.

Such a literal is said to be a *user-defined literal*.

When a named number that denotes a value of type *universal\_integer* is interpreted as a value of a non-numeric type *T*, it is equivalent to a call to the function denoted by the `Integer_Literal` aspect of *T*. The actual parameter of this notional call is a `String` having a textual representation of a decimal integer literal optionally preceded by a minus sign, representing the same value as the named number.

When a named number that denotes a value of type *universal\_real* is interpreted as a value of a non-numeric type *T*, it is equivalent to a call to the two-parameter function denoted by the `Real_Literal` aspect of *T*, if any. The actual parameters of this notional call are each a `String` with the textual representation of a decimal integer literal, with the first optionally preceded by a minus sign, where the first `String` represents the same value as the numerator, and the second the same value as the denominator, of the named number when represented as a rational number in lowest terms, with a positive denominator.

### Legality Rules

The `Integer_Literal` or `Real_Literal` aspect shall not be specified for a type *T* if the full view of *T* is a numeric type. The `String_Literal` aspect shall not be specified for a type *T* if the full view of *T* is a string type.

For a nonabstract type, the function directly specified for a user-defined literal aspect shall not be abstract.

For a tagged type with a partial view, a user-defined literal aspect shall not be directly specified on the full type.

If a nonabstract tagged type inherits any user-defined literal aspect, then each inherited aspect shall be directly specified as a nonabstract function for the type unless the inherited aspect denotes a nonabstract function, or functions, and the type is a null extension.

If a named number that denotes a value of type *universal\_integer* is interpreted as a value of a non-numeric type *T*, *T* shall have an `Integer_Literal` aspect. If a named number that denotes a value of type *universal\_real* is interpreted as a value of a non-numeric type *T*, *T* shall have a `Real_Literal` aspect, and the aspect shall denote a function that has two **in** parameters, both of type `String`, with result of type *T*.

In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

### Bounded (Run-Time) Errors

It is a bounded error if the evaluation of a literal or named number that has an expected type with a specified user-defined literal aspect propagates an exception. Either `Program_Error` or the exception propagated by the evaluation is raised at the point of use of the value of the literal or named number. If it is recognized prior to run time that evaluation of such a literal or named number will inevitably (if executed) result in such a bounded error, then this may be reported as an error prior to run time.

EXAMPLE *Examples of the specification and use of user-defined literals:*

```

subtype Roman_Character is Wide_Wide_Character
  with Static_Predicate =>
    Roman_Character in 'I' | 'V' | 'X' | 'L' | 'C' | 'D' | 'M';

Max_Roman_Number : constant := 3_999; -- MMMCMXCIX

type Roman_Number is range 1 .. Max_Roman_Number
  with String_Literal => To_Roman_Number;

function To_Roman_Number (S : Wide_Wide_String) return Roman_Number
  with Pre => S'Length > 0 and then
    (for all Char of S => Char in Roman_Character);

function To_Roman_Number (S : Wide_Wide_String) return Roman_Number is
  (declare
    R : constant array (Integer range <>) of Roman_Number :=
      (for D in S'Range => Roman_Digit'Enum_Rep
        (Roman_Digit'Wide_Wide_Value ('' & S(D) & '')));
    -- See 6.5.2 and 16.4
  begin
    [for I in R'Range =>
      (if I < R'Last and then R(I) < R(I + 1) then -1 else 1) * R(I)]
      'Reduce("+", 0)
    );

X : Roman_Number := "III" * "IV" * "XII"; -- 144 (that is, CXLIV)

```

### 7.3 Aggregates

An *aggregate* combines component values into a composite value of an array type, record type, or record extension.

#### Syntax

```
aggregate ::=
  record_aggregate | extension_aggregate | array_aggregate
  | delta_aggregate | container_aggregate
```

#### Name Resolution Rules

The expected type for an aggregate shall be a single array type, a single type with the Aggregate aspect specified, or a single descendant of a record type or of a record extension.

#### Legality Rules

A record\_aggregate or extension\_aggregate shall not be of a class-wide type.

#### Dynamic Semantics

For the evaluation of an aggregate, an anonymous object is created and values for the components or ancestor part are obtained (as described in the subsequent subclause for each kind of the aggregate) and assigned into the corresponding components or ancestor part of the anonymous object. Obtaining the values and the assignments occur in an arbitrary order. The value of the aggregate is the value of this object.

If an aggregate is of a tagged type, a check is made that its value belongs to the first subtype of the type. Constraint\_Error is raised if this check fails.

#### 7.3.1 Record Aggregates

In a record\_aggregate, a value is specified for each component of the record or record extension value, using either a named or a positional association.

#### Syntax

```
record_aggregate ::= (record_component_association_list)
record_component_association_list ::=
  record_component_association {, record_component_association}
  | null record
record_component_association ::=
  [component_choice_list =>] expression
  | component_choice_list => <>
component_choice_list ::=
  component_selector_name {'|' component_selector_name}
  | others
```

A record\_component\_association is a *named component association* if it has a component\_choice\_list; otherwise, it is a *positional component association*. Any positional component associations shall precede any named component associations. If there is a named association with a component\_choice\_list of **others**, it shall come last.

In the record\_component\_association\_list for a record\_aggregate, if there is only one association, it shall be a named association.

#### Name Resolution Rules

The expected type for a record\_aggregate shall be a single record type or record extension.

For the record\_component\_association\_list of a record\_aggregate, all components of the composite value defined by the aggregate are *needed*; for the association list of an

`extension_aggregate`, only those components not determined by the ancestor expression or subtype are needed (see 7.3.2). Each `component_selector_name` in a `record_component_association` of a `record_aggregate` or `extension_aggregate` shall denote a needed component (including possibly a discriminant). Each `component_selector_name` in a `record_component_association` of a `record_delta_aggregate` (see 7.3.4) shall denote a nondiscriminant component of the type of the aggregate.

The expected type for the expression of a `record_component_association` is the type of the *associated* component(s); the associated component(s) are as follows:

- For a positional association, the component (including possibly a discriminant) in the corresponding relative position (in the declarative region of the type), counting only the needed components;
- For a named association with one or more `component_selector_names`, the named component(s);
- For a named association with the reserved word **others**, all needed components that are not associated with some previous association.

#### Legality Rules

If the type of a `record_aggregate` is a record extension, then it shall be a descendant of a record type, through one or more record extensions (and no private extensions).

A `record_component_association_list` shall be **null record** only if the list occurs in a `record_aggregate` or `extension_aggregate`, and there are no components needed for that list.

For a `record_aggregate` or `extension_aggregate`, each `record_component_association` other than an **others** choice with a `<>` shall have at least one associated component, and each needed component shall be associated with exactly one `record_component_association`. For a `record_delta_aggregate`, each `component_selector_name` of each `component_choice_list` shall denote a distinct nondiscriminant component of the type of the aggregate.

If a `record_component_association` with an expression has two or more associated components, all of them shall be of the same type, or all of them shall be of anonymous access types whose subtypes statically match. In addition, Legality Rules are enforced separately for each associated component.

For a `record_aggregate` or `extension_aggregate`, if a `variant_part` *P* is nested within a `variant` *V* that is not selected by the discriminant value governing the `variant_part` enclosing *V*, then there is no restriction on the discriminant governing *P*. Otherwise, the value of the discriminant that governs *P* shall be given by a static expression, or by a nonstatic expression having a constrained static nominal subtype. In this latter case of a nonstatic expression, there shall be exactly one `discrete_choice_list` of *P* that covers each value that belongs to the nominal subtype and satisfies the predicates of the subtype, and there shall be at least one such value.

A `record_component_association` for a discriminant without a `default_expression` shall have an expression rather than `<>`.

A `record_component_association` of the `record_component_association_list` of a `record_delta_aggregate` shall not:

- use the box compound delimiter `<>` rather than an expression;
- have an expression of a limited type;
- omit the `component_choice_list`; or
- have a `component_choice_list` that is an **others** choice.

For a `record_delta_aggregate`, no two *component\_selector\_names* shall denote components declared within different *variants* of the same *variant\_part*.

#### Dynamic Semantics

The evaluation of a `record_aggregate` consists of the evaluation of the `record_component_association_list`.

For the evaluation of a `record_component_association_list`, any per-object constraints (see 6.8) for components specified in the association list are elaborated and any expressions are evaluated and converted to the subtype of the associated component. Any constraint elaborations and expression evaluations (and conversions) occur in an arbitrary order, except that the expression for a discriminant is evaluated (and converted) prior to the elaboration of any per-object constraint that depends on it, which in turn occurs prior to the evaluation and conversion of the expression for the component with the per-object constraint. If the value of a discriminant that governs a selected *variant\_part* is given by a nonstatic expression, and the evaluation of that expression yields a value that does not belong to the nominal subtype of the expression, then `Constraint_Error` is raised.

For a `record_component_association` with an expression, the expression defines the value for the associated component(s). For a `record_component_association` with `<>`, if the *component\_declaration* has a *default\_expression*, that *default\_expression* defines the value for the associated component(s); otherwise, the associated component(s) are initialized by default as for a stand-alone object of the component subtype (see 6.3.1).

The expression of a `record_component_association` is evaluated (and converted) once for each associated component.

NOTE By the rules given above, for a `record_aggregate` with positional associations, expressions specifying discriminant values appear first and in the same order as their corresponding *discriminant\_specifications*, since the known *discriminant\_part* occurs first in the declaration of the type.

EXAMPLE 1 *Example of a record aggregate with positional associations:*

```
(4, July, 1776)                                -- see 6.8
```

EXAMPLE 2 *Examples of record aggregates with named associations:*

```
(Day => 4, Month => July, Year => 1776)
(Month => July, Day => 4, Year => 1776)

(Disk, Closed, Track => 5, Cylinder => 12)        -- see 6.8.1
(Unit => Disk, Status => Closed, Cylinder => 9, Track => 1)
```

EXAMPLE 3 *Examples of component associations with several choices:*

```
(Value => 0, Succ|Pred => new Cell'(0, null, null)) -- see 6.10.1
-- The allocator is evaluated twice:
-- Succ and Pred designate different cells

(Value => 0, Succ|Pred => <>)                      -- see 6.10.1
-- Succ and Pred will be set to null
```

EXAMPLE 4 *Examples of record aggregates for tagged types (see 6.9 and 6.9.1):*

```
Expression'(null record)
Literal'(Value => 0.0)
Painted_Point'(0.0, Pi/2.0, Paint => Red)
```

### 7.3.2 Extension Aggregates

An `extension_aggregate` specifies a value for a type that is a record extension by specifying a value or subtype for an ancestor of the type, followed by associations for any components not determined by the *ancestor\_part*.

*Syntax*

```

extension_aggregate ::=
  (ancestor_part with record_component_association_list)
ancestor_part ::= expression | subtype_mark

```

*Name Resolution Rules*

The expected type for an `extension_aggregate` shall be a single type that is a record extension. If the `ancestor_part` is an expression, it is expected to be of any tagged type.

*Legality Rules*

If the `ancestor_part` is a `subtype_mark`, it shall denote a specific tagged subtype. If the `ancestor_part` is an expression, it shall not be dynamically tagged. The type of the `extension_aggregate` shall be a descendant of the type of the `ancestor_part` (the *ancestor* type), through one or more record extensions (and no private extensions). If the `ancestor_part` is a `subtype_mark`, the view of the ancestor type from which the type is descended (see 10.3.1) shall not have unknown discriminants.

If the type of the `ancestor_part` is limited and at least one component is needed in the `record_component_association_list`, then the `ancestor_part` shall not have an operative constituent expression (see 7.4) that is a call to a function with an unconstrained result subtype.

*Static Semantics*

For the `record_component_association_list` of an `extension_aggregate`, the only components *needed* are those of the composite value defined by the aggregate that are not inherited from the type of the `ancestor_part`, plus any inherited discriminants if the `ancestor_part` is a `subtype_mark` that denotes an unconstrained subtype.

*Dynamic Semantics*

For the evaluation of an `extension_aggregate`, the `record_component_association_list` is evaluated. If the `ancestor_part` is an expression, it is also evaluated; if the `ancestor_part` is a `subtype_mark`, the components of the value of the aggregate not given by the `record_component_association_list` are initialized by default as for an object of the ancestor type. Any implicit initializations or evaluations are performed in an arbitrary order, except that the expression for a discriminant is evaluated prior to any other evaluation or initialization that depends on it.

If the type of the `ancestor_part` has discriminants and the `ancestor_part` is not a `subtype_mark` that denotes an unconstrained subtype, then a check is made that each discriminant determined by the `ancestor_part` has the value specified for a corresponding discriminant, if any, either in the `record_component_association_list`, or in the `derived_type_definition` for some ancestor of the type of the `extension_aggregate`. `Constraint_Error` is raised if this check fails.

NOTE 1 By the rules given in 7.3.1, if all components of the value of the `extension_aggregate` are determined by the `ancestor_part`, then the `record_component_association_list` will be simply **null record**.

NOTE 2 If the `ancestor_part` is a `subtype_mark`, then its type can be abstract. If its type is controlled, then as the last step of evaluating the aggregate, the `Initialize` procedure of the ancestor type is called, unless the `Initialize` procedure is abstract (see 10.6).

EXAMPLE *Examples of extension aggregates (for types defined in 6.9.1):*

```

Painted_Point'(Point with Red)
(Point'(P) with Paint => Black)

(Expression with Left => new Literal'(Value => 1.2),
              Right => new Literal'(Value => 3.4))
Addition'(Binop with null record)
-- presuming Binop is of type Binary_Operation

```



### 7.3.3 Array Aggregates

In an `array_aggregate`, a value is specified for each component of an array, either positionally or by its index. For a `positional_array_aggregate`, the components are given in increasing-index order, with a final **others**, if any, representing any remaining components. For a `named_array_aggregate`, the components are identified by the values covered by the `discrete_choices`.

#### Syntax

```
array_aggregate ::=
  positional_array_aggregate | null_array_aggregate | named_array_aggregate
positional_array_aggregate ::=
  (expression, expression { expression })
  | (expression { expression }, others => expression)
  | (expression { expression }, others => <>)
  | '[' expression { expression } [, others => expression] '['
  | '[' expression { expression }, others => <> '['
null_array_aggregate ::= '[' ']'
named_array_aggregate ::=
  (array_component_association_list)
  | '[' array_component_association_list ']'
array_component_association_list ::=
  array_component_association { array_component_association }
array_component_association ::=
  discrete_choice_list => expression
  | discrete_choice_list => <>
  | iterated_component_association
iterated_component_association ::=
  for defining_identifier in discrete_choice_list => expression
  | for iterator_specification => expression
```

An *n-dimensional* `array_aggregate` is one that is written as *n* levels of nested `array_aggregates` (or at the bottom level, equivalent `string_literals`). For the multidimensional case ( $n \geq 2$ ) the `array_aggregates` (or equivalent `string_literals`) at the *n-1* lower levels are called *subaggregates* of the enclosing *n-dimensional* `array_aggregate`. The expressions of the bottom level subaggregates (or of the `array_aggregate` itself if one-dimensional) are called the *array component expressions* of the enclosing *n-dimensional* `array_aggregate`.

The `defining_identifier` of an `iterated_component_association` declares an *index parameter*, an object of the corresponding index type.

#### Name Resolution Rules

The expected type for an `array_aggregate` (that is not a subaggregate) shall be a single array type. The component type of this array type is the expected type for each array component expression of the `array_aggregate`.

The expected type for each `discrete_choice` in any `discrete_choice_list` of a `named_array_aggregate` is the type of the *corresponding index*; the corresponding index for an `array_aggregate` that is not a subaggregate is the first index of its type; for an (*n-m*)-dimensional subaggregate within an `array_aggregate` of an *n-dimensional* type, the corresponding index is the index in position *m+1*.

#### Legality Rules

An `array_aggregate` of an *n-dimensional* array type shall be written as an *n-dimensional* `array_aggregate`, or as a `null_array_aggregate`.

An **others** choice is allowed for an `array_aggregate` only if an *applicable index constraint* applies to the `array_aggregate`. An applicable index constraint is a constraint provided by certain contexts that can be used to determine the bounds of the array value specified by an `array_aggregate`. Each of the following contexts (and none other) defines an applicable index constraint:

- For an `explicit_actual_parameter`, an `explicit_generic_actual_parameter`, the expression of a return statement, the return expression of an expression function, the initialization expression in an `object_declaration`, or a `default_expression` (for a parameter or a component), when the nominal subtype of the corresponding formal parameter, generic formal parameter, function return object, expression function return object, object, or component is a constrained array subtype, the applicable index constraint is the constraint of the subtype;
- For the expression of an `assignment_statement` where the name denotes an array variable, the applicable index constraint is the constraint of the array variable;
- For the operand of a `qualified_expression` whose `subtype_mark` denotes a constrained array subtype, the applicable index constraint is the constraint of the subtype;
- For a component expression in an aggregate, if the component's nominal subtype is a constrained array subtype, the applicable index constraint is the constraint of the subtype;
- For the *base\_expression* of a `delta_aggregate`, if the nominal subtype of the `delta_aggregate` is a constrained array subtype, the applicable index constraint is the constraint of the subtype;
- For a parenthesized expression, the applicable index constraint is that, if any, defined for the expression;
- For a `conditional_expression` (see 7.5.7), the applicable index constraint for each *dependent\_expression* is that, if any, defined for the `conditional_expression`;
- For a `declare_expression` (see 7.5.9), the applicable index constraint for the *body\_expression* is that, if any, defined for the `declare_expression`.

The applicable index constraint *applies* to an `array_aggregate` that appears in such a context, as well as to any subaggregates thereof. In the case of an `explicit_actual_parameter` (or `default_expression`) for a call on a generic formal subprogram, no applicable index constraint is defined.

The `discrete_choice_list` of an `array_component_association` (including an `iterated_component_association`) is allowed to have a `discrete_choice` that is a `nonstatic_choice_expression` or that is a `subtype_indication` or range that defines a nonstatic or null range, only if it is the single `discrete_choice` of its `discrete_choice_list`, and either there is only one `array_component_association` in the enclosing `array_component_association_list` or the enclosing aggregate is an `array_delta_aggregate`, not an `array_aggregate`.

Either all or none of the `array_component_associations` of an `array_component_association_list` shall be `iterated_component_associations` with an `iterator_specification`.

In a `named_array_aggregate` where all `discrete_choices` are static, no two `discrete_choices` are allowed to cover the same value (see 6.8.1); if there is no **others** choice, the `discrete_choices` taken together shall exactly cover a contiguous sequence of values of the corresponding index type.

A bottom level subaggregate of a multidimensional `array_aggregate` of a given array type is allowed to be a `string_literal` only if the component type of the array type is a character type; each character of such a `string_literal` shall correspond to a `defining_character_literal` of the component type.

*Static Semantics*

A subaggregate that is a `string_literal` is equivalent to one that is a `positional_array_aggregate` of the same length, with each expression being the `character_literal` for the corresponding character of the `string_literal`.

The subtype (and nominal subtype) of an index parameter is the corresponding index subtype.

*Dynamic Semantics*

For an `array_aggregate` that contains only `array_component_associations` that are `iterated_component_associations` with `iterator_specifications`, evaluation proceeds in two steps:

1. Each `iterator_specification` is elaborated (in an arbitrary order) and an iteration is performed solely to determine a maximum count for the number of values produced by the iteration; all of these counts are combined to determine the overall length of the array, and ultimately the limits on the bounds of the array (defined below);
2. A second iteration is performed for each of the `iterator_specifications`, in the order given in the aggregate, and for each value conditionally produced by the iteration (see 8.5 and 8.5.2), the associated expression is evaluated, its value is converted to the component subtype of the array type, and used to define the value of the next component of the array starting at the low bound and proceeding sequentially toward the high bound. A check is made that the second iteration results in an array length no greater than the maximum determined by the first iteration; `Constraint_Error` is raised if this check fails.

The evaluation of any other `array_aggregate` of a given array type proceeds in two steps:

1. Any `discrete_choices` of this aggregate and of its subaggregates are evaluated in an arbitrary order, and converted to the corresponding index type;
2. The array component expressions of the aggregate are evaluated in an arbitrary order and their values are converted to the component subtype of the array type; an array component expression is evaluated once for each associated component.

Each expression in an `array_component_association` defines the value for the associated component(s). For an `array_component_association` with `<>`, the associated component(s) are initialized to the `Default_Component_Value` of the array type if this aspect has been specified for the array type; otherwise, they are initialized by default as for a stand-alone object of the component subtype (see 6.3.1).

During an evaluation of the expression of an `iterated_component_association` with a `discrete_choice_list`, the value of the corresponding index parameter is that of the corresponding index of the corresponding array component. During an evaluation of the expression of an `iterated_component_association` with an `iterator_specification`, the value of the loop parameter of the `iterator_specification` is the value produced by the iteration (as described in 8.5.2).

The bounds of the index range of an `array_aggregate` (including a subaggregate) are determined as follows:

- For an `array_aggregate` with an **others** choice, the bounds are those of the corresponding index range from the applicable index constraint;
- For a `positional_array_aggregate` (or equivalent `string_literal`) without an **others** choice, the lower bound is that of the corresponding index range in the applicable index constraint, if defined, or that of the corresponding index subtype, if not; in either case, the upper bound is determined from the lower bound and the number of expressions (or the length of the `string_literal`);

- For a `null_array_aggregate`, bounds for each dimension are determined as for a `positional_array_aggregate` without an **others** choice that has no expressions for each dimension;
- For a `named_array_aggregate` containing only `iterated_component_associations` with an `iterator_specification`, the lower bound is determined as for a `positional_array_aggregate` without an **others** choice, and the upper bound is determined from the lower bound and the total number of values produced by the second set of iterations;
- For any other `named_array_aggregate` without an **others** choice, the bounds are determined by the smallest and largest index values covered by any `discrete_choice_list`.

For an `array_aggregate`, a check is made that the index range defined by its bounds is compatible with the corresponding index subtype.

For an `array_aggregate` with an **others** choice, a check is made that no expression or `<>` is specified for an index value outside the bounds determined by the applicable index constraint.

For a multidimensional `array_aggregate`, a check is made that all subaggregates that correspond to the same index have the same bounds.

The exception `Constraint_Error` is raised if any of the above checks fail.

#### Implementation Permissions

When evaluating `iterated_component_associations` for an `array_aggregate` that contains only `iterated_component_associations` with `iterator_specifications`, the first step of evaluating an `iterated_component_association` can be omitted if the implementation can determine the maximum number of values by some other means.

NOTE 1 In an `array_aggregate` delimited by parentheses, positional notation can only be used with two or more expressions; a single expression in parentheses is interpreted as a parenthesized expression. An `array_aggregate` delimited by square brackets can be used to specify an array with a single component.

NOTE 2 An index parameter is a constant object (see 6.3).

EXAMPLE 1 *Examples of array aggregates with positional associations:*

```
(7, 9, 5, 1, 3, 2, 4, 8, 6, 0)
Table'(5, 8, 4, 1, others => 0) -- see 6.6
```

EXAMPLE 2 *Examples of array aggregates with named associations:*

```
(1 .. 5 => (1 .. 8 => 0.0)) -- two-dimensional
[1 .. N => new Cell] -- N new cells, in particular for N = 0

Table'(2 | 4 | 10 => 1, others => 0)
Schedule'(Mon .. Fri => True, others => False) -- see 6.6
Schedule'[Wed | Sun => False, others => True]
Vector'(1 => 2.5) -- single-component vector
```

EXAMPLE 3 *Examples of two-dimensional array aggregates:*

```
-- Three aggregates for the same value of
-- subtype Matrix(1..2,1..3) (see 6.6):

((1.1, 1.2, 1.3), (2.1, 2.2, 2.3))
(1 => [1.1, 1.2, 1.3], 2 => [2.1, 2.2, 2.3])
[1 => (1 => 1.1, 2 => 1.2, 3 => 1.3), 2 => (1 => 2.1, 2 => 2.2, 3 => 2.3)]
```

EXAMPLE 4 *Examples of aggregates as initial values:*

```
A : Table := (7, 9, 5, 1, 3, 2, 4, 8, 6, 0); -- A(1)=7, A(10)=0
B : Table := (2 | 4 | 10 => 1, others => 0); -- B(1)=0, B(10)=1
C : constant Matrix := (1 .. 5 => (1 .. 8 => 0.0)); -- C'Last(1)=5,
C'Last(2)=8

D : Bit_Vector(M .. N) := (M .. N => True); -- see 6.6
E : Bit_Vector(M .. N) := (others => True);
F : String(1 .. 1) := (1 => 'F'); -- a one component aggregate: same as "F"
```

```

G : constant Matrix :=
  (for I in 1 .. 4 =>
    (for J in 1 .. 4 =>
      (if I=J then 1.0 else 0.0))); -- Identity matrix

Empty_Matrix : constant Matrix := []; -- A matrix without elements

```

EXAMPLE 5 Example of an array aggregate with defaulted others choice and with an applicable index constraint provided by an enclosing record aggregate:

```

Buffer'(Size => 50, Pos => 1, Value => ('x', others => <>)) -- see 6.7

```

### 7.3.4 Delta Aggregates

Evaluating a (record or array) delta aggregate yields a composite value that starts with a copy of another value of the same type and then assigns to some (but typically not all) components of the copy.

#### Syntax

```

delta_aggregate ::= record_delta_aggregate | array_delta_aggregate
record_delta_aggregate ::=
  (base_expression with delta record_component_association_list)
array_delta_aggregate ::=
  (base_expression with delta array_component_association_list)
  | '[' base_expression with delta array_component_association_list ']'

```

#### Name Resolution Rules

The expected type for a `record_delta_aggregate` shall be a single descendant of a record type or record extension.

The expected type for an `array_delta_aggregate` shall be a single array type.

The expected type for the *base\_expression* of any `delta_aggregate` is the type of the enclosing `delta_aggregate`.

The Name Resolution Rules and Legality Rules for each `record_component_association` of a `record_delta_aggregate` are as defined in 7.3.1.

For an `array_delta_aggregate`, the expected type for each `discrete_choice` in an `array_component_association` is the index type of the type of the `delta_aggregate`.

The expected type of the expression in an `array_component_association` is defined as for an `array_component_association` occurring within an `array_aggregate` of the type of the `delta_aggregate`.

#### Legality Rules

For an `array_delta_aggregate`, the `array_component_association` shall not use the box symbol <>, and the `discrete_choice` shall not be **others**.

For an `array_delta_aggregate`, the dimensionality of the type of the `delta_aggregate` shall be 1.

For an `array_delta_aggregate`, the *base\_expression* and each expression in every `array_component_association` shall be of a nonlimited type.

#### Dynamic Semantics

The evaluation of a `delta_aggregate` begins with the evaluation of the *base\_expression* of the `delta_aggregate`; then that value is used to create and initialize the anonymous object of the aggregate. The bounds of the anonymous object of an `array_delta_aggregate` and the discriminants (if any) of the anonymous object of a `record_delta_aggregate` are those of the

*base\_expression*. If a *record\_delta\_aggregate* is of a specific tagged type, its tag is that of the specific type; if it is of a class-wide type, its tag is that of the *base\_expression*.

For a *record\_delta\_aggregate*, for each component associated with each *record\_component\_association* (in an unspecified order):

- if the associated component belongs to a variant, a check is made that the values of the discriminants are such that the anonymous object has this component. The exception *Constraint\_Error* is raised if this check fails.
- the expression of the *record\_component\_association* is evaluated, converted to the nominal subtype of the associated component, and assigned to the component of the anonymous object.

For an *array\_delta\_aggregate*, for each *discrete\_choice* of each *array\_component\_association* (in the order given in the enclosing *discrete\_choice\_list* and *array\_component\_association\_list*, respectively) the *discrete\_choice* is evaluated; for each represented index value (in ascending order, if the *discrete\_choice* represents a range):

- the index value is converted to the index type of the array type.
- a check is made that the index value belongs to the index range of the anonymous object of the aggregate; *Constraint\_Error* is raised if this check fails.
- the component expression is evaluated, converted to the array component subtype, and assigned to the component of the anonymous object identified by the index value.

EXAMPLE 1 *Examples of use of delta aggregates in a postcondition:*

```
procedure Twelfth (D : in out Date) -- see 6.8 for type Date
  with Post => D = (D'Old with delta Day => 12);

procedure The_Answer (V : in out Vector;
  A, B : in Integer) -- see 6.6 for type Vector
  with Post => V = (V'Old with delta A .. B => 42.0, V'First => 0.0);
```

EXAMPLE 2 *Examples where the base expression is nontrivial:*

```
New_Cell : Cell := (Min_Cell (Head) with delta Value => 42);
-- see 6.10.1 for Cell and Head; 9.1 for Min_Cell

A1 : Vector := ((0 => 1.0, 1 => 2.0, 2 => 3.0)
  with delta Integer(Random * 2.0) => 14.2);
-- see 6.6 for declaration of type Vector
-- see 9.1 for declaration of Random

Tomorrow := ((Yesterday with delta Day => 12)
  with delta Month => April); -- see 6.8
```

EXAMPLE 3 *Example where the base expression is class-wide:*

```
function Translate (P : Point'Class; X, Y : Real) return Point'Class is
  (P with delta X => P.X + X,
   Y => P.Y + Y); -- see 6.9 for declaration of type Point
```

### 7.3.5 Container Aggregates

In a *container\_aggregate*, values are specified for elements of a container; for a *positional\_container\_aggregate*, the elements are given sequentially; for a *named\_container\_aggregate*, the elements are specified by a sequence of key/value pairs, or using an iterator. The Aggregate aspect of the type of the aggregate determines how the elements are combined to form the container.

For a type other than an array type, the following type-related operational aspect may be specified:

Aggregate This aspect is an aggregate of the form:

```
(Empty => name[,
  Add_Named => procedure_name][,
```



```
Add_Unnamed => procedure_name][,
New_Indexed => function_name,
Assign_Indexed => procedure_name])
```

The type for which this aspect is specified is known as the *container type* of the Aggregate aspect. A *procedure\_name* shall be specified for at least one of Add\_Named, Add\_Unnamed, or Assign\_Indexed. If Add\_Named is specified, neither Add\_Unnamed nor Assign\_Indexed shall be specified. Either both or neither of New\_Indexed and Assign\_Indexed shall be specified.

#### Name Resolution Rules

The *name* specified for Empty for an Aggregate aspect shall denote a constant of the container type, or denote exactly one function with a result type of the container type that has no parameters, or that has one **in** parameter of a signed integer type.

The *procedure\_name* specified for Add\_Unnamed for an Aggregate aspect shall denote exactly one procedure that has two parameters, the first an **in out** parameter of the container type, and the second an **in** parameter of some nonlimited type, called the *element type* of the container type.

The *function\_name* specified for New\_Indexed for an Aggregate aspect shall denote exactly one function with a result type of the container type, and two parameters of the same discrete type, with that type being the *key type* of the container type.

The *procedure\_name* specified for Add\_Named or Assign\_Indexed for an Aggregate aspect shall denote exactly one procedure that has three parameters, the first an **in out** parameter of the container type, the second an **in** parameter of a nonlimited type (the *key type* of the container type), and the third, an **in** parameter of a nonlimited type that is called the *element type* of the container type.

#### Legality Rules

If the container type of an Aggregate aspect is a private type, the full type of the container type shall not be an array type. If the container type is limited, the name specified for Empty shall denote a function rather than a constant object.

For an Aggregate aspect, the key type of Assign\_Indexed shall be the same type as that of the parameters of New\_Indexed. Additionally, if both Add\_Unnamed and Assign\_Indexed are specified, the final parameters shall be of the same type — the element type of the container type.

#### Static Semantics

The Aggregate aspect is nonoverridable (see 16.1.1).

#### Syntax

```
container_aggregate ::=
  null_container_aggregate
| positional_container_aggregate
| named_container_aggregate

null_container_aggregate ::= '[' ']'

positional_container_aggregate ::= '[' expression{, expression} ']'

named_container_aggregate ::= '[' container_element_association_list ']'

container_element_association_list ::=
  container_element_association{, container_element_association}

container_element_association ::=
  key_choice_list => expression
| key_choice_list => <>
| iterated_element_association

key_choice_list ::= key_choice {'|' key_choice}
```

```

key_choice ::= key_expression | discrete_range
iterated_element_association ::=
  for loop_parameter_specification[ use key_expression ] => expression
  | for iterator_specification[ use key_expression ] => expression

```

#### Name Resolution Rules

The expected type for a `container_aggregate` shall be a single type for which the Aggregate aspect has been specified. The expected type for each expression of a `container_aggregate` is the element type of the expected type.

The expected type for a `key_expression`, or a `discrete_range` of a `key_choice`, is the key type of the expected type of the aggregate.

#### Legality Rules

The expected type for a `positional_container_aggregate` shall have an Aggregate aspect that includes a specification for an `Add_Unnamed` procedure or an `Assign_Indexed` procedure. The expected type for a `named_container_aggregate` that contains one or more `iterated_element_associations` with a `key_expression` shall have an Aggregate aspect that includes a specification for the `Add_Named` procedure. The expected type for a `named_container_aggregate` that contains one or more `key_choice_lists` shall have an Aggregate aspect that includes a specification for the `Add_Named` or `Assign_Indexed` procedure. A `null_container_aggregate` can be of any type with an Aggregate aspect.

A non-null container aggregate is called an *indexed aggregate* if the expected type *T* of the aggregate specifies an `Assign_Indexed` procedure in its Aggregate aspect, and either there is no `Add_Unnamed` procedure specified for the type, or the aggregate is a `named_container_aggregate` with a `container_element_association` that contains a `key_choice_list` or a `loop_parameter_specification`. The key type of an indexed aggregate is also called the *index type* of the aggregate.

A `container_element_association` with a `<>` rather than an expression, or with a `key_choice` that is a `discrete_range`, is permitted only in an indexed aggregate.

For an `iterated_element_association` without a `key_expression`, if the aggregate is an indexed aggregate or the expected type of the aggregate specifies an `Add_Named` procedure in its Aggregate aspect, then the type of the loop parameter of the `iterated_element_association` shall be the same as the key type of the aggregate.

For a `named_container_aggregate` that is an indexed aggregate, all `container_element_associations` shall contain either a `key_choice_list`, or a `loop_parameter_specification` without a `key_expression` or `iterator_filter`. Furthermore, for such an aggregate, either:

- all `key_choices` shall be static expressions or static ranges, and every `loop_parameter_specification` shall have a `discrete_subtype_definition` that defines a non-null static range, and the set of values of the index type covered by the `key_choices` and the `discrete_subtype_definitions` shall form a contiguous range of values with no duplications; or
- there shall be exactly one `container_element_association`, and if it has a `key_choice_list`, the list shall have exactly one `key_choice`.

#### Dynamic Semantics

The evaluation of a `container_aggregate` starts by creating an anonymous object *A* of the expected type *T*, initialized as follows:

- if the aggregate is an indexed aggregate, from the result of a call on the `New_Indexed` function; the actual parameters in this call represent the lower and upper bound of the aggregate, and are determined as follows:
  - if the aggregate is a `positional_container_aggregate`, the lower bound is the low bound of the subtype of the key parameter of the `Add_Indexed` procedure, and the upper bound has a position number that is the sum of the position number of the lower bound and one less than the number of expressions in the aggregate;
  - if the aggregate is a `named_container_aggregate`, the lower bound is the lowest value covered by a `key_choice_list` or is the low bound of a range defined by a `discrete_subtype_definition` of a `loop_parameter_specification`; the upper bound is the highest value covered by a `key_choice_list` or is the high bound of a range defined by a `discrete_subtype_definition` of a `loop_parameter_specification`.
- if the aggregate is not an indexed aggregate, by assignment from the `Empty` constant, or from a call on the `Empty` function specified in the `Aggregate` aspect. In the case of an `Empty` function with a formal parameter, the actual parameter has the following value:
  - for a `null_container_aggregate`, the value zero;
  - for a `positional_container_aggregate`, the number of expressions;
  - for a `named_container_aggregate` without an `iterated_element_association`, the number of `key_expressions`;
  - for a `named_container_aggregate` where every `iterated_element_association` contains a `loop_parameter_specification`, the total number of elements specified by all of the `container_element_associations`;
  - otherwise, to an implementation-defined value.

The evaluation then proceeds as follows:

- for a `null_container_aggregate`, the anonymous object *A* is the result;
- for a `positional_container_aggregate` of a type with a specified `Add_Unnamed` procedure, each expression is evaluated in an arbitrary order, and the `Add_Unnamed` procedure is invoked in sequence with the anonymous object *A* as the first parameter and the result of evaluating each expression as the second parameter, in the order of the expressions;
- for a `positional_container_aggregate` that is an indexed aggregate, each expression is evaluated in an arbitrary order, and the `Assign_Indexed` procedure is invoked in sequence with the anonymous object *A* as the first parameter, the key value as the second parameter, computed by starting with the low bound of the subtype of the key formal parameter of the `Assign_Indexed` procedure and taking the successor of this value for each successive expression, and the result of evaluating each expression as the third parameter;
- for a `named_container_aggregate` for a type with an `Add_Named` procedure in its `Aggregate` aspect, the `container_element_associations` are evaluated in an arbitrary order:
  - for a `container_element_association` with a `key_choice_list`, for each `key_choice` of the list in an arbitrary order, the `key_choice` is evaluated as is the expression of the `container_element_association` (in an arbitrary order), and the `Add_Named` procedure is invoked once for each value covered by the `key_choice`, with the anonymous object *A* as the first parameter, the value from the `key_choice` as the second parameter, and the result of evaluating the expression as the third parameter;
  - for a `container_element_association` with an `iterated_element_association`, first the `iterated_element_association` is elaborated, then an iteration is performed, and for each value conditionally produced by the iteration (see 8.5 and 8.5.2) the `Add_Named` procedure is invoked with the anonymous object *A* as the first parameter, the result of evaluating the expression as the third parameter, and:

- if there is a *key\_expression*, the result of evaluating the *key\_expression* as the second parameter;
- otherwise, with the loop parameter as the second parameter;
- for a *named\_container\_aggregate* that is an indexed aggregate, the evaluation proceeds as above for the case of *Add\_Named*, but with the *Assign\_Indexed* procedure being invoked instead of *Add\_Named*; in the case of a *container\_element\_association* with a *<>* rather than an expression, the corresponding call on *Assign\_Indexed* is not performed, leaving the component as it was upon return from the *New\_Indexed* function;
- for any other *named\_container\_aggregate*, the *container\_element\_associations* (which are necessarily *iterated\_element\_associations*) are evaluated in the order given; each such evaluation comprises two steps:
  1. the *iterated\_element\_association* is elaborated;
  2. an iteration is performed, and for each value conditionally produced by the iteration (see 8.5 and 8.5.2) the *Add\_Unnamed* procedure is invoked, with the anonymous object *A* as the first parameter and the result of evaluating the expression as the second parameter.

EXAMPLE 1 *Examples of specifying the Aggregate aspect for a Set\_Type, a Map\_Type, and a Vector\_Type:*

```
-- Set_Type is a set-like container type.
type Set_Type is private
  with Aggregate => (Empty      => Empty_Set,
                    Add_Unnamed => Include);
function Empty_Set return Set_Type;

subtype Small_Int is Integer range -1000..1000;

procedure Include (S : in out Set_Type; N : in Small_Int);

-- Map_Type is a map-like container type.
type Map_Type is private
  with Aggregate => (Empty      => Empty_Map,
                    Add_Named  => Add_To_Map);

procedure Add_To_Map (M      : in out Map_Type;
                     Key     : in Integer;
                     Value   : in String);

Empty_Map : constant Map_Type;

-- Vector_Type is an extensible array-like container type.
type Vector_Type is private
  with Aggregate => (Empty      => Empty_Vector,
                    Add_Unnamed => Append_One,
                    New_Indexed  => New_Vector,
                    Assign_Indexed => Assign_Element);

function Empty_Vector (Capacity : Integer := 0) return Vector_Type;
procedure Append_One (V : in out Vector_Type; New_Item : in String);
procedure Assign_Element (V      : in out Vector_Type;
                         Index   : in Positive;
                         Item    : in String);

function New_Vector (First, Last : Positive) return Vector_Type
  with Pre => First = Positive'First;
  -- Vectors are always indexed starting at the
  -- lower bound of their index subtype.

-- Private part not shown.
```

EXAMPLE 2 *Examples of container aggregates for Set\_Type, Map\_Type, and Vector\_Type:*

```
-- Example aggregates using Set_Type.
S : Set_Type;

-- Assign the empty set to S:
S := [];
```

```

-- Is equivalent to:
S := Empty_Set;

-- A positional set aggregate:
S := [1, 2];

-- Is equivalent to:
S := Empty_Set;
Include (S, 1);
Include (S, 2);

-- A set aggregate with an iterated_element_association:
S := [for Item in 1 .. 5 => Item * 2];

-- Is equivalent to:
S := Empty_Set;
for Item in 1 .. 5 loop
    Include (S, Item * 2);
end loop;

-- A set aggregate consisting of two iterated_element_associations:
S := [for Item in 1 .. 5 => Item,
      for Item in 1 .. 5 => -Item];

-- Is equivalent (assuming set semantics) to:
S := Empty_Set;
for Item in 1 .. 5 loop
    Include (S, Item);
end loop;
for Item in -5 .. -1 loop
    Include (S, Item);
end loop;

-- Example aggregates using Map_Type.
M : Map_Type;

-- A simple named map aggregate:
M := [12 => "house", 14 => "beige"];

-- Is equivalent to:
M := Empty_Map;
Add_To_Map (M, 12, "house");
Add_To_Map (M, 14, "beige");

-- Define a table of pairs:
type Pair is record
    Key : Integer;
    Value : access constant String;
end record;

Table : constant array(Positive range <>) of Pair :=
    [(Key => 33, Value => new String'("a nice string")),
     (Key => 44, Value => new String'("an even better string"))];

-- A map aggregate using an iterated_element_association
-- and a key_expression, built from from a table of key/value pairs:
M := [for P of Table use P.Key => P.Value.all];

-- Is equivalent to:
M := Empty_Map;
for P of Table loop
    Add_To_Map (M, P.Key, P.Value.all);
end loop;

-- Create an image table for an array of integers:
Keys : constant array(Positive range <>) of Integer := [2, 3, 5, 7, 11];

-- A map aggregate where the values produced by the
-- iterated_element_association are of the same type as the key
-- (hence a separate key_expression is unnecessary):
M := [for Key of Keys => Integer'Image (Key)];

```

```

-- Is equivalent to:
M := Empty_Map;
for Key of Keys loop
  Add_To_Map (M, Key, Integer'Image (Key));
end loop;

-- Example aggregates using Vector_Type.
V : Vector_Type;

-- A positional vector aggregate:
V := ["abc", "def"];

-- Is equivalent to:
V := Empty_Vector (2);
Append_One (V, "abc");
Append_One (V, "def");

-- An indexed vector aggregate:
V := [1 => "this", 2 => "is", 3 => "a", 4 => "test"];

-- Is equivalent to:
V := New_Vector (1, 4);
Assign_Element (V, 1, "this");
Assign_Element (V, 2, "is");
Assign_Element (V, 3, "a");
Assign_Element (V, 4, "test");

```

## 7.4 Expressions

An *expression* is a formula that defines the computation or retrieval of a value. In this document, the term “expression” refers to a construct of the syntactic category expression or of any of the following categories: choice\_expression, choice\_relation, relation, simple\_expression, term, factor, primary, conditional\_expression, quantified\_expression.

### Syntax

```

expression ::=
  relation {and relation} | relation {and then relation}
  | relation {or relation} | relation {or else relation}
  | relation {xor relation}
choice_expression ::=
  choice_relation {and choice_relation}
  | choice_relation {or choice_relation}
  | choice_relation {xor choice_relation}
  | choice_relation {and then choice_relation}
  | choice_relation {or else choice_relation}
choice_relation ::=
  simple_expression [relational_operator simple_expression]
relation ::=
  simple_expression [relational_operator simple_expression]
  | tested simple_expression [not] in membership_choice_list
  | raise_expression
membership_choice_list ::= membership_choice {'membership_choice}
membership_choice ::= choice simple_expression | range | subtype_mark
simple_expression ::=
  [unary_adding_operator] term {binary_adding_operator term}
term ::= factor {multiplying_operator factor}
factor ::= primary [** primary] | abs primary | not primary
primary ::=

```



```

numeric_literal | null | string_literal | aggregate
| name | allocator | (expression)
| (conditional_expression) | (quantified_expression)
| (declare_expression)

```

#### Name Resolution Rules

A name used as a primary shall resolve to denote an object or a value.

#### Static Semantics

Each expression has a type; it specifies the computation or retrieval of a value of that type.

A primary that is an expression surrounded by ( and ) is known as a *parenthesized expression*.

Every name or expression consists of one or more *operative constituent* names or expressions, only one of which is evaluated as part of evaluating the name or expression (the *evaluated operative constituent*). The operative constituents are determined as follows, according to the form of the expression (or name):

- if the expression is a conditional\_expression, the operative constituents of its *dependent\_expressions*;
- if the expression (or name) is a parenthesized expression, a qualified\_expression, or a view conversion, the operative constituent(s) of its operand;
- if the expression is a declare\_expression, the operative constituent(s) of its *body\_expression*;
- otherwise, the expression (or name) itself.

In certain contexts, we specify that an operative constituent shall (or shall not) be *newly constructed*. This means the operative constituent shall (or shall not) be an aggregate or a function\_call; in either case, a raise\_expression is permitted.

#### Dynamic Semantics

The value of a primary that is a name denoting an object is the value of the object.

An expression of a numeric universal type is evaluated as if it has type *root\_integer* (for *universal\_integer*) or *root\_real* (otherwise) unless the context identifies a specific type (in which case that type is used).

#### Implementation Permissions

For the evaluation of a primary that is a name denoting an object of an unconstrained numeric subtype, if the value of the object is outside the base range of its type, the implementation may either raise Constraint\_Error or return the value of the object.

#### EXAMPLE 1 Examples of primaries:

```

4.0          -- real literal
Pi           -- named number
(1 .. 10 => 0) -- array aggregate
Sum          -- variable
Integer'Last -- attribute
Sine(X)      -- function call
Color'(Blue) -- qualified expression
Real(M*N)    -- conversion
(Line_Count + 10) -- parenthesized expression

```

**EXAMPLE 2** *Examples of expressions:*

|                                        |                                                 |
|----------------------------------------|-------------------------------------------------|
| Volume                                 | -- <i>primary</i>                               |
| <b>not</b> Destroyed                   | -- <i>factor</i>                                |
| 2*Line_Count                           | -- <i>term</i>                                  |
| -4.0                                   | -- <i>simple expression</i>                     |
| -4.0 + A                               | -- <i>simple expression</i>                     |
| B**2 - 4.0*A*C                         | -- <i>simple expression</i>                     |
| R*Sin(θ)*Cos(φ)                        | -- <i>simple expression</i>                     |
| Password(1 .. 3) = "Bwv"               | -- <i>relation</i>                              |
| Count <b>in</b> Small_Int              | -- <i>relation</i>                              |
| Count <b>not in</b> Small_Int          | -- <i>relation</i>                              |
| Index = 0 <b>or</b> Item_Hit           | -- <i>expression</i>                            |
| (Cold <b>and</b> Sunny) <b>or</b> Warm | -- <i>expression (parentheses are required)</i> |
| A**(B**C)                              | -- <i>expression (parentheses are required)</i> |

**7.5 Operators and Expression Evaluation**

The language defines the following six categories of operators (given in order of increasing precedence). The corresponding `operator_symbols`, and only those, can be used as designators in declarations of functions for user-defined operators. See 9.6.

*Syntax*

|                                 |                                        |
|---------------------------------|----------------------------------------|
| logical_operator ::=            | <b>and   or   xor</b>                  |
| relational_operator ::=         | <b>=   &lt;=   &lt;   &gt;   &gt;=</b> |
| binary_adding_operator ::=      | <b>+   -   &amp;</b>                   |
| unary_adding_operator ::=       | <b>+   -</b>                           |
| multiplying_operator ::=        | <b>*   /   mod   rem</b>               |
| highest_precedence_operator ::= | <b>**   abs   not</b>                  |

*Static Semantics*

For a sequence of operators of the same precedence level, the operators are associated with their operands in textual order from left to right. Parentheses can be used to impose specific associations.

For each form of type definition, certain of the above operators are *predefined*; that is, they are implicitly declared immediately after the type definition. For each such implicit operator declaration, the parameters are called Left and Right for *binary* operators; the single parameter is called Right for *unary* operators. An expression of the form X op Y, where op is a binary operator, is equivalent to a `function_call` of the form "op"(X, Y). An expression of the form op Y, where op is a unary operator, is equivalent to a `function_call` of the form "op"(Y). The predefined operators and their effects are described in 7.5.1 through 7.5.6.

*Dynamic Semantics*

The predefined operations on integer types either yield the mathematically correct result or raise the exception `Constraint_Error`. For implementations that support the Numerics Annex, the predefined operations on real types yield results whose accuracy is defined in Annex G, or raise the exception `Constraint_Error`.

*Implementation Requirements*

The implementation of a predefined operator that delivers a result of an integer or fixed point type may raise `Constraint_Error` only if the result is outside the base range of the result type.

The implementation of a predefined operator that delivers a result of a floating point type may raise `Constraint_Error` only if the result is outside the safe range of the result type.

## Implementation Permissions

For a sequence of predefined operators of the same precedence level (and in the absence of parentheses imposing a specific association), an implementation may impose any association of the operators with operands so long as the result produced is an allowed result for the left-to-right association, but ignoring the potential for failure of language-defined checks in either the left-to-right or chosen order of association.

NOTE The two operands of an expression of the form  $X \text{ op } Y$ , where  $\text{op}$  is a binary operator, are evaluated in an arbitrary order, as for any `function_call` (see 9.4).

EXAMPLE *Examples of precedence:*

```

not Sunny or Warm      -- same as (not Sunny) or Warm
X > 4.0 and Y > 0.0      -- same as (X > 4.0) and (Y > 0.0)

-4.0*A**2                -- same as -(4.0 * (A**2))
abs(1 + A) + B           -- same as (abs (1 + A)) + B
Y**(-3)                  -- parentheses are necessary
A / B * C                 -- same as (A/B)*C
A + (B + C)               -- evaluate B + C before adding it to A

```

## 7.5.1 Logical Operators and Short-circuit Control Forms

## Name Resolution Rules

An expression consisting of two relations connected by **and then** or **or else** (a *short-circuit control form*) shall resolve to be of some boolean type; the expected type for both relations is that same boolean type.

## Static Semantics

The following logical operators are predefined for every boolean type  $T$ , for every modular type  $T$ , and for every one-dimensional array type  $T$  whose component type is a boolean type:

```

function "and" (Left, Right : T) return T
function "or"  (Left, Right : T) return T
function "xor" (Left, Right : T) return T

```

For boolean types, the predefined logical operators **and**, **or**, and **xor** perform the conventional operations of conjunction, inclusive disjunction, and exclusive disjunction, respectively.

For modular types, the predefined logical operators are defined on a bit-by-bit basis, using the binary representation of the value of the operands to yield a binary representation for the result, where zero represents False and one represents True. If this result is outside the base range of the type, a final subtraction by the modulus is performed to bring the result into the base range of the type.

The logical operators on arrays are performed on a component-by-component basis on matching components (as for equality — see 7.5.2), using the predefined logical operator for the component type. The bounds of the resulting array are those of the left operand.

## Dynamic Semantics

The short-circuit control forms **and then** and **or else** deliver the same result as the corresponding predefined **and** and **or** operators for boolean types, except that the left operand is always evaluated first, and the right operand is not evaluated if the value of the left operand determines the result.

For the logical operators on arrays, a check is made that for each component of the left operand there is a matching component of the right operand, and vice versa. Also, a check is made that each component of the result belongs to the component subtype. The exception `Constraint_Error` is raised if either of the above checks fails.

NOTE The conventional meaning of the logical operators is given by the following truth table:

| A    | B    | (A <b>and</b> B) | (A <b>or</b> B) | (A <b>xor</b> B) |
|------|------|------------------|-----------------|------------------|
| True | True | True             | True            | False            |

|       |       |       |       |       |
|-------|-------|-------|-------|-------|
| True  | False | False | True  | True  |
| False | True  | False | True  | True  |
| False | False | False | False | False |

EXAMPLE 1 *Examples of logical operators:*

```
Sunny or Warm
Filter(1 .. 10) and Filter(15 .. 24)  -- see 6.6.1
```

EXAMPLE 2 *Examples of short-circuit control forms:*

```
Next_Car.Owner /= null and then Next_Car.Owner.Age > 25  -- see 6.10.1
N = 0 or else A(N) = Hit_Value
```

## 7.5.2 Relational Operators and Membership Tests

The *equality operators* = (equals) and /= (not equals) are predefined for nonlimited types. The other *relational operators* are the *ordering operators* < (less than), <= (less than or equal), > (greater than), and >= (greater than or equal). The ordering operators are predefined for scalar types, and for *discrete array types*, that is, one-dimensional array types whose components are of a discrete type.

A *membership test*, using **in** or **not in**, determines whether or not a value belongs to any given subtype or range, is equal to any given value, has a tag that identifies a type that is covered by a given type, or is convertible to and has an accessibility level appropriate for a given access type. Membership tests are allowed for all types.

### *Name Resolution Rules*

The *tested type* of a membership test is determined by the *membership\_choices* of the *membership\_choice\_list*. Either all *membership\_choices* of the *membership\_choice\_list* shall resolve to the same type, which is the tested type; or each *membership\_choice* shall be of an elementary type, and the tested type shall be covered by each of these elementary types.

If the tested type is tagged, then the *tested\_simple\_expression* shall resolve to be of a type that is convertible (see 7.6) to the tested type; if untagged, the expected type of the *tested\_simple\_expression* is the tested type. The expected type of a *choice\_simple\_expression* in a *membership\_choice*, and of a *simple\_expression* of a range in a *membership\_choice*, is the tested type of the membership operation.

### *Legality Rules*

For a membership test, if the *tested\_simple\_expression* is of a tagged class-wide type, then the tested type shall be (visibly) tagged.

If a membership test includes one or more *choice\_simple\_expressions* and the tested type of the membership test is limited, then the tested type of the membership test shall have a visible primitive equality operator; if the tested type of the membership test is nonlimited with a user-defined primitive equality operator that is defined at a point where the type is limited, the tested type shall be a record type or record extension.

### *Static Semantics*

The result type of a membership test is the predefined type Boolean.

The equality operators are predefined for every specific type *T* that is not limited, and not an anonymous access type, with the following specifications:

```
function "=" (Left, Right : T) return Boolean
function "/"= (Left, Right : T) return Boolean
```

The following additional equality operators for the *universal\_access* type are declared in package Standard for use with anonymous access types:

```
function "=" (Left, Right : universal_access) return Boolean
function "/"= (Left, Right : universal_access) return Boolean
```

The ordering operators are predefined for every specific scalar type *T*, and for every discrete array type *T*, with the following specifications:

```
function "<" (Left, Right : T) return Boolean
function "<=" (Left, Right : T) return Boolean
function ">" (Left, Right : T) return Boolean
function ">=" (Left, Right : T) return Boolean
```

#### *Name Resolution Rules*

At least one of the operands of an equality operator for *universal\_access* shall be of a specific anonymous access type. Unless the predefined equality operator is identified using an expanded name with `prefix` denoting the package `Standard`, neither operand shall be of an access-to-object type whose designated type is *D* or *D'*Class, where *D* has a user-defined primitive equality operator such that:

- its result type is Boolean;
- it is declared immediately within the same declaration list as *D* or any partial or incomplete view of *D*; and
- at least one of its operands is an access parameter with designated type *D*.

#### *Legality Rules*

At least one of the operands of the equality operators for *universal\_access* shall be of type *universal\_access*, or both shall be of access-to-object types, or both shall be of access-to-subprogram types. Further:

- When both are of access-to-object types, the designated types shall be the same or one shall cover the other, and if the designated types are elementary or array types, then the designated subtypes shall statically match;
- When both are of access-to-subprogram types, the designated profiles shall be subtype conformant.

If the profile of an explicitly declared primitive equality operator of an untagged record type is type conformant with that of the corresponding predefined equality operator, the declaration shall occur before the type is frozen. In addition, no type shall have been derived from the untagged record type before the declaration of the primitive equality operator. In addition to the places where Legality Rules normally apply (see 15.3), this rule applies also in the private part of an instance of a generic unit.

#### *Dynamic Semantics*

For discrete types, the predefined relational operators are defined in terms of corresponding mathematical operations on the position numbers of the values of the operands.

For real types, the predefined relational operators are defined in terms of the corresponding mathematical operations on the values of the operands, subject to the accuracy of the type.

Two access-to-object values are equal if they designate the same object, or if both are equal to the null value of the access type.

Two access-to-subprogram values are equal if they are the result of the same evaluation of an `Access attribute_reference`, or if both are equal to the null value of the access type. Two access-to-subprogram values are unequal if they designate different subprograms. It is unspecified whether two access values that designate the same subprogram but are the result of distinct evaluations of `Access attribute_references` are equal or unequal.

For a type extension, predefined equality is defined in terms of the primitive (possibly user-defined) equals operator for the parent type and for any components that have a record type in the extension part, and predefined equality for any other components not inherited from the parent type.

For a private type, if its full type is a record type or a record extension, predefined equality is defined in terms of the primitive equals operator of the full type; otherwise, predefined equality for the private type is that of its full type.

For other composite types, the predefined equality operators (and certain other predefined operations on composite types — see 7.5.1 and 7.6) are defined in terms of the corresponding operation on *matching components*, defined as follows:

- For two composite objects or values of the same non-array type, matching components are those that correspond to the same `component_declaration` or `discriminant_specification`;
- For two one-dimensional arrays of the same type, matching components are those (if any) whose index values match in the following sense: the lower bounds of the index ranges are defined to match, and the successors of matching indices are defined to match;
- For two multidimensional arrays of the same type, matching components are those whose index values match in successive index positions.

The analogous definitions apply if the types of the two objects or values are convertible, rather than being the same.

Given the above definition of matching components, the result of the predefined equals operator for composite types (other than for those composite types covered earlier) is defined as follows:

- If there are no components, the result is defined to be True;
- If there are unmatched components, the result is defined to be False;
- Otherwise, the result is defined in terms of the primitive equals operator for any matching components that are records, and the predefined equals for any other matching components.

If the primitive equals operator for an untagged record type is abstract, then `Program_Error` is raised at the point of any call to that abstract subprogram, implicitly as part of an equality operation on an enclosing composite object, or in an instance of a generic with a formal private type where the actual type is a record type with an abstract `"=`".

For any composite type, the order in which `"=`" is called for components is unspecified. Furthermore, if the result can be determined before calling `"=`" on some components, it is unspecified whether `"=`" is called on those components.

The predefined `"/=`" operator gives the complementary result to the predefined `"=`" operator.

For a discrete array type, the predefined ordering operators correspond to *lexicographic order* using the predefined order relation of the component type: A null array is lexicographically less than any array having at least one component. In the case of nonnull arrays, the left operand is lexicographically less than the right operand if the first component of the left operand is less than that of the right; otherwise, the left operand is lexicographically less than the right operand only if their first components are equal and the tail of the left operand is lexicographically less than that of the right (the *tail* consists of the remaining components beyond the first and can be null).

An *individual membership test* is the membership test of a single `membership_choice`.

For the evaluation of a membership test using **in** whose `membership_choice_list` has a single `membership_choice`, the *tested\_simple\_expression* and the `membership_choice` are evaluated in an arbitrary order; the result is the result of the individual membership test for the `membership_choice`.

For the evaluation of a membership test using **in** whose `membership_choice_list` has more than one `membership_choice`, the *tested\_simple\_expression* of the membership test is evaluated first and the result of the operation is equivalent to that of a sequence consisting of an individual membership test on each `membership_choice` combined with the short-circuit control form **or else**.

An individual membership test yields the result True if:

- The `membership_choice` is a *choice\_simple\_expression*, and the *tested\_simple\_expression* is equal to the value of the `membership_choice`. If the tested type is a record type or a record extension, or is limited at the point where the membership test occurs, the test uses the primitive equality for the type; otherwise, the test uses predefined equality.



- The `membership_choice` is a range and the value of the *tested\_simple\_expression* belongs to the given range.
- The `membership_choice` is a `subtype_mark`, the tested type is scalar, the value of the *tested\_simple\_expression* belongs to the range of the named subtype, and the value satisfies the predicates of the named subtype.
- The `membership_choice` is a `subtype_mark`, the tested type is not scalar, the value of the *tested\_simple\_expression* satisfies any constraints of the named subtype, the value satisfies the predicates of the named subtype, and:
  - if the type of the *tested\_simple\_expression* is class-wide, the value has a tag that identifies a type covered by the tested type;
  - if the tested type is an access type and the named subtype excludes null, the value of the *tested\_simple\_expression* is not null;
  - if the tested type is a general access-to-object type, the type of the *tested\_simple\_expression* is convertible to the tested type and its accessibility level is no deeper than that of the tested type; further, if the designated type of the tested type is tagged and the *tested\_simple\_expression* is nonnull, the tag of the object designated by the value of the *tested\_simple\_expression* is covered by the designated type of the tested type.

Otherwise, the test yields the result `False`.

A membership test using **not in** gives the complementary result to the corresponding membership test using **in**.

#### Implementation Requirements

For all nonlimited types declared in language-defined packages, the `"=`" and `"!="` operators of the type shall behave as if they were the predefined equality operators for the purposes of the equality of composite types and generic formal types.

NOTE If a composite type has components that depend on discriminants, two values of this type have matching components if and only if their discriminants are equal. Two nonnull arrays have matching components if and only if the length of each dimension is the same for both.

EXAMPLE *Examples of expressions involving relational operators and membership tests:*

```
X /= Y
A_String = "A" -- True (see 6.3.1)
"" < A_String and A_String < "Aa" -- True
A_String < "Bb" and A_String < "A " -- True

My_Car = null -- True if My_Car has been set to null (see
6.10.1)
My_Car = Your_Car -- True if we both share the same car
My_Car.all = Your_Car.all -- True if the two cars are identical

N not in 1 .. 10 -- range membership test
Today in Mon .. Fri -- range membership test
Today in Weekday -- subtype membership test (see 6.5.1)
Card in Clubs | Spades -- list membership test (see 6.5.1)
Archive in Disk_Unit -- subtype membership test (see 6.8.1)
Tree.all in Addition'Class -- class membership test (see 6.9.1)
```

### 7.5.3 Binary Adding Operators

#### Static Semantics

The binary adding operators `+` (addition) and `-` (subtraction) are predefined for every specific numeric type *T* with their conventional meaning. They have the following specifications:

```
function "+"(Left, Right : T) return T
function "-"(Left, Right : T) return T
```

The concatenation operators & are predefined for every nonlimited, one-dimensional array type  $T$  with component type  $C$ . They have the following specifications:

```
function "&" (Left : T; Right : T) return T
function "&" (Left : T; Right : C) return T
function "&" (Left : C; Right : T) return T
function "&" (Left : C; Right : C) return T
```

#### Dynamic Semantics

For the evaluation of a concatenation with result type  $T$ , if both operands are of type  $T$ , the result of the concatenation is a one-dimensional array whose length is the sum of the lengths of its operands, and whose components comprise the components of the left operand followed by the components of the right operand. If the left operand is a null array, the result of the concatenation is the right operand. Otherwise, the lower bound of the result is determined as follows:

- If the ultimate ancestor of the array type was defined by a `constrained_array_definition`, then the lower bound of the result is that of the index subtype;
- If the ultimate ancestor of the array type was defined by an `unconstrained_array_definition`, then the lower bound of the result is that of the left operand.

The upper bound is determined by the lower bound and the length. A check is made that the upper bound of the result of the concatenation belongs to the range of the index subtype, unless the result is a null array. `Constraint_Error` is raised if this check fails.

If either operand is of the component type  $C$ , the result of the concatenation is given by the above rules, using in place of such an operand an array having this operand as its only component (converted to the component subtype) and having the lower bound of the index subtype of the array type as its lower bound.

The result of a concatenation is defined in terms of an assignment to an anonymous object, as for any function call (see 9.5).

NOTE As for all predefined operators on modular types, the binary adding operators  $+$  and  $-$  on modular types include a final reduction modulo the modulus if the result is outside the base range of the type.

EXAMPLE *Examples of expressions involving binary adding operators:*

```
Z + 0.1      -- Z has to be of a real type
"A" & "BCD"  -- concatenation of two string literals
'A' & "BCD"  -- concatenation of a character literal and a string literal
'A' & 'A'    -- concatenation of two character literals
```

## 7.5.4 Unary Adding Operators

#### Static Semantics

The unary adding operators  $+$  (identity) and  $-$  (negation) are predefined for every specific numeric type  $T$  with their conventional meaning. They have the following specifications:

```
function "+" (Right : T) return T
function "-" (Right : T) return T
```

NOTE For modular integer types, the unary adding operator  $-$ , when given a nonzero operand, returns the result of subtracting the value of the operand from the modulus; for a zero operand, the result is zero.

## 7.5.5 Multiplying Operators

#### Static Semantics

The multiplying operators  $*$  (multiplication),  $/$  (division), **mod** (modulus), and **rem** (remainder) are predefined for every specific integer type  $T$ :

```

function "*" (Left, Right : T) return T
function "/" (Left, Right : T) return T
function "mod" (Left, Right : T) return T
function "rem" (Left, Right : T) return T

```

Signed integer multiplication has its conventional meaning.

Signed integer division and remainder are defined by the relation:

$$A = (A/B) * B + (A \text{ rem } B)$$

where  $(A \text{ rem } B)$  has the sign of  $A$  and an absolute value less than the absolute value of  $B$ . Signed integer division satisfies the identity:

$$(-A) / B = -(A / B) = A / (-B)$$

The signed integer modulus operator is defined such that the result of  $A \text{ mod } B$  is either zero, or has the sign of  $B$  and an absolute value less than the absolute value of  $B$ ; in addition, for some signed integer value  $N$ , this result satisfies the relation:

$$A = B * N + (A \text{ mod } B)$$

The multiplying operators on modular types are defined in terms of the corresponding signed integer operators, followed by a reduction modulo the modulus if the result is outside the base range of the type (which is only possible for the "\*" operator).

Multiplication and division operators are predefined for every specific floating point type  $T$ :

```

function "*" (Left, Right : T) return T
function "/" (Left, Right : T) return T

```

The following multiplication and division operators, with an operand of the predefined type Integer, are predefined for every specific fixed point type  $T$ :

```

function "*" (Left : T; Right : Integer) return T
function "*" (Left : Integer; Right : T) return T
function "/" (Left : T; Right : Integer) return T

```

All of the above multiplying operators are usable with an operand of an appropriate universal numeric type. The following additional multiplying operators for *root\_real* are predefined, and are usable when both operands are of an appropriate universal or root numeric type, and the result is allowed to be of type *root\_real*, as in a number\_declaration:

```

function "*" (Left, Right : root_real) return root_real
function "/" (Left, Right : root_real) return root_real

function "*" (Left : root_real; Right : root_integer) return root_real
function "*" (Left : root_integer; Right : root_real) return root_real
function "/" (Left : root_real; Right : root_integer) return root_real

```

Multiplication and division between any two fixed point types are provided by the following two predefined operators:

```

function "*" (Left, Right : universal_fixed) return universal_fixed
function "/" (Left, Right : universal_fixed) return universal_fixed

```

#### Name Resolution Rules

The above two fixed-fixed multiplying operators shall not be used in a context where the expected type for the result is itself *universal\_fixed* — the context has to identify some other numeric type to which the result is to be converted, either explicitly or implicitly. Unless the predefined universal operator is identified using an expanded name with *prefix* denoting the package Standard, an explicit conversion is required on the result when using the above fixed-fixed multiplication operator if either operand is of a type having a user-defined primitive multiplication operator such that:

- it is declared immediately within the same declaration list as the type or any partial or incomplete view thereof; and
- both of its formal parameters are of a fixed-point type.

A corresponding requirement applies to the universal fixed-fixed division operator.

#### Dynamic Semantics

The multiplication and division operators for real types have their conventional meaning. For floating point types, the accuracy of the result is determined by the precision of the result type. For decimal fixed point types, the result is truncated toward zero if the mathematical result is between two multiples of the *small* of the specific result type (possibly determined by context); for ordinary fixed point types, if the mathematical result is between two multiples of the *small*, it is unspecified which of the two is the result.

The exception `Constraint_Error` is raised by integer division, **rem**, and **mod** if the right operand is zero. Similarly, for a real type *T* with *T*.Machine\_Overflows True, division by zero raises `Constraint_Error`.

NOTE 1 For positive *A* and *B*, *A*/*B* is the quotient and *A rem B* is the remainder when *A* is divided by *B*. The following relations are satisfied by the **rem** operator:

$$\begin{aligned} A \text{ rem } (-B) &= A \text{ rem } B \\ (-A) \text{ rem } B &= -(A \text{ rem } B) \end{aligned}$$

NOTE 2 For any signed integer *K*, the following identity holds:

$$A \text{ mod } B = (A + K*B) \text{ mod } B$$

The relations between signed integer division, remainder, and modulus are illustrated by the following table:

| <i>A</i> | <i>B</i> | <i>A</i> / <i>B</i> | <i>A rem B</i> | <i>A mod B</i> | <i>A</i> | <i>B</i> | <i>A</i> / <i>B</i> | <i>A rem B</i> | <i>A mod B</i> |
|----------|----------|---------------------|----------------|----------------|----------|----------|---------------------|----------------|----------------|
| 10       | 5        | 2                   | 0              | 0              | -10      | 5        | -2                  | 0              | 0              |
| 11       | 5        | 2                   | 1              | 1              | -11      | 5        | -2                  | -1             | 4              |
| 12       | 5        | 2                   | 2              | 2              | -12      | 5        | -2                  | -2             | 3              |
| 13       | 5        | 2                   | 3              | 3              | -13      | 5        | -2                  | -3             | 2              |
| 14       | 5        | 2                   | 4              | 4              | -14      | 5        | -2                  | -4             | 1              |
| <i>A</i> | <i>B</i> | <i>A</i> / <i>B</i> | <i>A rem B</i> | <i>A mod B</i> | <i>A</i> | <i>B</i> | <i>A</i> / <i>B</i> | <i>A rem B</i> | <i>A mod B</i> |
| 10       | -5       | -2                  | 0              | 0              | -10      | -5       | 2                   | 0              | 0              |
| 11       | -5       | -2                  | 1              | -4             | -11      | -5       | 2                   | -1             | -1             |
| 12       | -5       | -2                  | 2              | -3             | -12      | -5       | 2                   | -2             | -2             |
| 13       | -5       | -2                  | 3              | -2             | -13      | -5       | 2                   | -3             | -3             |
| 14       | -5       | -2                  | 4              | -1             | -14      | -5       | 2                   | -4             | -4             |

EXAMPLE Examples of expressions involving multiplying operators:

```

I : Integer := 1;
J : Integer := 2;
K : Integer := 3;

X : Real := 1.0;      -- see 6.5.7
Y : Real := 2.0;

F : Fraction := 0.25; -- see 6.5.9
G : Fraction := 0.5;

```

| Expression                      | Value | Result Type                                                     |
|---------------------------------|-------|-----------------------------------------------------------------|
| <i>I</i> * <i>J</i>             | 2     | same as <i>I</i> and <i>J</i> , that is, Integer                |
| <i>K</i> / <i>J</i>             | 1     | same as <i>K</i> and <i>J</i> , that is, Integer                |
| <i>K mod J</i>                  | 1     | same as <i>K</i> and <i>J</i> , that is, Integer                |
| <i>X</i> / <i>Y</i>             | 0.5   | same as <i>X</i> and <i>Y</i> , that is, Real                   |
| <i>F</i> / <i>2</i>             | 0.125 | same as <i>F</i> , that is, Fraction                            |
| 3* <i>F</i>                     | 0.75  | same as <i>F</i> , that is, Fraction                            |
| 0.75* <i>G</i>                  | 0.375 | universal_fixed, implicitly convertible to any fixed point type |
| Fraction( <i>F</i> * <i>G</i> ) | 0.125 | Fraction, as stated by the conversion                           |
| Real( <i>J</i> )* <i>Y</i>      | 4.0   | Real, the type of both operands after conversion of <i>J</i>    |

## 7.5.6 Highest Precedence Operators

### Static Semantics

The highest precedence unary operator **abs** (absolute value) is predefined for every specific numeric type  $T$ , with the following specification:

```
function "abs" (Right :  $T$ ) return  $T$ 
```

The highest precedence unary operator **not** (logical negation) is predefined for every boolean type  $T$ , every modular type  $T$ , and for every one-dimensional array type  $T$  whose components are of a boolean type, with the following specification:

```
function "not" (Right :  $T$ ) return  $T$ 
```

The result of the operator **not** for a modular type is defined as the difference between the high bound of the base range of the type and the value of the operand. For a binary modulus, this corresponds to a bit-wise complement of the binary representation of the value of the operand.

The operator **not** that applies to a one-dimensional array of boolean components yields a one-dimensional boolean array with the same bounds; each component of the result is obtained by logical negation of the corresponding component of the operand (that is, the component that has the same index value). A check is made that each component of the result belongs to the component subtype; the exception `Constraint_Error` is raised if this check fails.

The highest precedence *exponentiation* operator **\*\*** is predefined for every specific integer type  $T$  with the following specification:

```
function "***" (Left :  $T$ ; Right : Natural) return  $T$ 
```

Exponentiation is also predefined for every specific floating point type as well as *root\_real*, with the following specification (where  $T$  is *root\_real* or the floating point type):

```
function "***" (Left :  $T$ ; Right : Integer | Base) return  $T$ 
```

The right operand of an exponentiation is the *exponent*. The value of  $X^{**}N$  with the value of the exponent  $N$  positive is the same as the value of  $X * X * \dots * X$  (with  $N-1$  multiplications) except that the multiplications are associated in an arbitrary order. With  $N$  equal to zero, the result is one. With the value of  $N$  negative (only defined for a floating point operand), the result is the reciprocal of the result using the absolute value of  $N$  as the exponent.

### Implementation Permissions

The implementation of exponentiation for the case of a negative exponent is allowed to raise `Constraint_Error` if the intermediate result of the repeated multiplications is outside the safe range of the type, even though the final result (after taking the reciprocal) would not be. (The best machine approximation to the final result in this case would generally be 0.0.)

NOTE As implied by the specification given above for exponentiation of an integer type, a check is made that the exponent is not negative. `Constraint_Error` is raised if this check fails.

## 7.5.7 Conditional Expressions

A `conditional_expression` selects for evaluation at most one of the enclosed `dependent_expressions`, depending on a decision among the alternatives. One kind of `conditional_expression` is the `if_expression`, which selects for evaluation a `dependent_expression` depending on the value of one or more corresponding conditions. The other kind of `conditional_expression` is the `case_expression`, which selects for evaluation one of a number of alternative `dependent_expressions`; the chosen alternative is determined by the value of a `selecting_expression`.

### Syntax

```
conditional_expression ::= if_expression | case_expression
if_expression ::=
```

```

if condition then dependent_expression
{elsif condition then dependent_expression}
[else dependent_expression]

condition ::= boolean_expression

case_expression ::=
  case selecting_expression is
  case_expression_alternative {
    case_expression_alternative}

case_expression_alternative ::=
  when discrete_choice_list =>
    dependent_expression

```

Wherever the Syntax Rules allow an expression, a *conditional\_expression* may be used in place of the expression, so long as it is immediately surrounded by parentheses.

#### Name Resolution Rules

If a *conditional\_expression* is expected to be of a type *T*, then each *dependent\_expression* of the *conditional\_expression* is expected to be of type *T*. Similarly, if a *conditional\_expression* is expected to be of some class of types, then each *dependent\_expression* of the *conditional\_expression* is subject to the same expectation. If a *conditional\_expression* shall resolve to be of a type *T*, then each *dependent\_expression* shall resolve to be of type *T*.

The possible types of a *conditional\_expression* are further determined as follows:

- If the *conditional\_expression* is the operand of a type conversion, the type of the *conditional\_expression* is the target type of the conversion; otherwise,
- If all of the *dependent\_expressions* are of the same type, the type of the *conditional\_expression* is that type; otherwise,
- If a *dependent\_expression* is of an elementary type, the type of the *conditional\_expression* shall be covered by that type; otherwise,
- If the *conditional\_expression* is expected to be of type *T* or shall resolve to type *T*, then the *conditional\_expression* is of type *T*.

A condition is expected to be of any boolean type.

The expected type for the *selecting\_expression* and the *discrete\_choices* are as for case statements (see 8.4).

#### Legality Rules

All of the *dependent\_expressions* shall be convertible (see 7.6) to the type of the *conditional\_expression*.

If the expected type of a *conditional\_expression* is a specific tagged type, all of the *dependent\_expressions* of the *conditional\_expression* shall be dynamically tagged, or none shall be dynamically tagged. In this case, the *conditional\_expression* is dynamically tagged if all of the *dependent\_expressions* are dynamically tagged, is tag-indeterminate if all of the *dependent\_expressions* are tag-indeterminate, and is statically tagged otherwise.

If there is no **else** *dependent\_expression*, the *if\_expression* shall be of a boolean type.

All Legality Rules that apply to the *discrete\_choices* of a *case\_statement* (see 8.4) also apply to the *discrete\_choices* of a *case\_expression* except within an instance of a generic unit.



For the evaluation of an `if_expression`, the condition specified after **if**, and any conditions specified after **elsif**, are evaluated in succession (treating a final **else** as **elsif True then**), until one evaluates to True or all conditions are evaluated and yield False. If a condition evaluates to True, the associated *dependent\_expression* is evaluated, converted to the type of the `if_expression`, and the resulting value is the value of the `if_expression`. Otherwise (when there is no **else** clause), the value of the `if_expression` is True.

For the evaluation of a `case_expression`, the *selecting\_expression* is first evaluated. If the value of the *selecting\_expression* is covered by the *discrete\_choice\_list* of some `case_expression_alternative`, then the *dependent\_expression* of the `case_expression_alternative` is evaluated, converted to the type of the `case_expression`, and the resulting value is the value of the `case_expression`. Otherwise (the value is not covered by any *discrete\_choice\_list*, perhaps due to being outside the base range), `Constraint_Error` is raised.

EXAMPLE 1 *Example of use of an if\_expression:*

```
Put_Line ("Casey is " &
  (if Casey.Sex = M then "Male" else "Female")); -- see 6.10.1
```

EXAMPLE 2 *Example of use of a case\_expression:*

```
function Card_Color (Card : Suit) return Color is -- see 6.5.1
  (case Card is
    when Clubs | Spades => Black,
    when Hearts | Diamonds => Red);
```

## 7.5.8 Quantified Expressions

Quantified expressions provide a way to write universally and existentially quantified predicates over containers and arrays.

### Syntax

```
quantified_expression ::=
  for quantifier loop_parameter_specification => predicate
  | for quantifier iterator_specification => predicate

quantifier ::= all | some

predicate ::= boolean_expression
```

Wherever the Syntax Rules allow an expression, a `quantified_expression` may be used in place of the expression, so long as it is immediately surrounded by parentheses.

### Name Resolution Rules

The expected type of a `quantified_expression` is any Boolean type. The predicate in a `quantified_expression` is expected to be of the same type.

### Dynamic Semantics

For the evaluation of a `quantified_expression`, the *loop\_parameter\_specification* or *iterator\_specification* is first elaborated. The evaluation of a `quantified_expression` then performs an iteration, and evaluates the predicate for each value conditionally produced by the iteration (see 8.5 and 8.5.2).

The value of the `quantified_expression` is determined as follows:

- If the quantifier is **all**, the expression is False if the evaluation of any predicate yields False; evaluation of the `quantified_expression` stops at that point. Otherwise (every predicate has been evaluated and yielded True), the expression is True. Any exception raised by evaluation of the predicate is propagated.

- If the quantifier is **some**, the expression is True if the evaluation of any predicate yields True; evaluation of the `quantified_expression` stops at that point. Otherwise (every predicate has been evaluated and yielded False), the expression is False. Any exception raised by evaluation of the predicate is propagated.

EXAMPLE 1 Example of a quantified expression as a postcondition for a sorting routine on an array *A* with an index subtype *T*:

```
Post => (A'Length < 2 or else
  (for all I in A'First .. T'Pred(A'Last) => A (I) <= A (T'Succ (I))))
```

EXAMPLE 2 Example of use of a quantified expression as an assertion that a positive number *N* is composite (as opposed to prime):

```
pragma Assert (for some X in 2 .. N when X * X <= N => N mod X = 0);
-- see iterator_filter in 8.5
```

### 7.5.9 Declare Expressions

Declare expressions provide a way to declare local constants and object renamings in an expression context.

#### Syntax

```
declare_expression ::=
  declare {declare_item}
  begin body_expression

declare_item ::= object_declaration | object_renaming_declaration
```

Wherever the Syntax Rules allow an expression, a `declare_expression` may be used in place of the expression, so long as it is immediately surrounded by parentheses.

#### Legality Rules

A `declare_item` that is an `object_declaration` shall declare a constant of a nonlimited type.

A `declare_item` that is an `object_renaming_declaration` (see 11.5.1) shall not rename an object of a limited type if any operative constituent of the *object\_name* is a value conversion or is newly constructed (see 7.4).

The following are not allowed within a `declare_expression`: a declaration containing the reserved word **aliased**; the attribute\_designator `Access` or `Unchecked_Access`; or an anonymous access type.

#### Name Resolution Rules

If a `declare_expression` is expected to be of a type *T*, then the *body\_expression* is expected to be of type *T*. Similarly, if a `declare_expression` is expected to be of some class of types, then the *body\_expression* is subject to the same expectation. If a `declare_expression` shall resolve to be of a type *T*, then the *body\_expression* shall resolve to be of type *T*.

The type of a `declare_expression` is the type of the *body\_expression*.

#### Dynamic Semantics

For the evaluation of a `declare_expression`, the `declare_items` are elaborated in order, and then the *body\_expression* is evaluated. The value of the `declare_expression` is that of the *body\_expression*.

EXAMPLE Example of use of a declare expression as a replacement postcondition for `Ada.Containers.Vectors."&"` (see A.18.2):

```

with Post =>
  (declare
    Result renames Vectors."&"Result;
    Length : constant Count_Type := Left.Length + Right.Length;
  begin
    Result.Length = Length and then
    not Tampering_With_Elements_Prohibited (Result) and then
    not Tampering_With_Cursors_Prohibited (Result) and then
    Result.Capacity >= Length)

```

### 7.5.10 Reduction Expressions

Reduction expressions provide a way to map or transform a collection of values into a new set of values, and then summarize the values produced by applying an operation to reduce the set to a single value result. A reduction expression is represented as an `attribute_reference` of the reduction attributes `Reduce` or `Parallel_Reduce`.

#### Syntax

```

reduction_attribute_reference ::=
  value_sequence reduction_attribute_designator
| prefix reduction_attribute_designator

value_sequence ::=
  '[' [parallel[(chunk_specification)] [aspect_specification]]
  iterated_element_association ']'

reduction_attribute_designator ::=
  reduction_identifier(reduction_specification)

reduction_specification ::= reducer_name, initial_value_expression

```

The `iterated_element_association` of a `value_sequence` shall not have a `key_expression`, nor shall it have a `loop_parameter_specification` that has the reserved word **reverse**.

The `chunk_specification`, if any, of a `value_sequence` shall be an `integer_simple_expression`.

#### Name Resolution Rules

The expected type for a `reduction_attribute_reference` shall be a single nonlimited type.

In the remainder of this subclause, we will refer to nonlimited subtypes *Value\_Type* and *Accum\_Type* of a `reduction_attribute_reference`. These subtypes and interpretations of the names and expressions of a `reduction_attribute_reference` are determined by the following rules:

- *Accum\_Type* is a subtype of the expected type of the `reduction_attribute_reference`.
- A *reducer subprogram* is subtype conformant with one of the following specifications:
 

```

function Reducer(Accumulator : Accum_Type;
                  Value : Value_Type) return Accum_Type;

procedure Reducer(Accumulator : in out Accum_Type;
                  Value : in Value_Type);
      
```
- The *reducer\_name* of a `reduction_specification` denotes a reducer subprogram.
- The expected type of an *initial\_value\_expression* of a `reduction_specification` is that of subtype *Accum\_Type*.
- The expected type of the expression of the `iterated_element_association` of a `value_sequence` is that of subtype *Value\_Type*.

*Legality Rules*

If the `reduction_attribute_reference` has a `value_sequence` with the reserved word **parallel**, the subtypes *Accum\_Type* and *Value\_Type* shall statically match.

If the identifier of a `reduction_attribute_designator` is `Parallel_Reduce`, the subtypes *Accum\_Type* and *Value\_Type* shall statically match.

*Static Semantics*

A `reduction_attribute_reference` denotes a value, with its nominal subtype being the subtype of the first parameter of the subprogram denoted by the *reducer\_name*.

*Dynamic Semantics*

For the evaluation of a `value_sequence`, the `iterated_element_association`, the `chunk_specification`, and the `aspect_specification`, if any, are elaborated in an arbitrary order. Next an iteration is performed, and for each value conditionally produced by the iteration (see 8.5 and 8.5.2), the associated expression is evaluated with the loop parameter having this value, which produces a result that is converted to *Value\_Type* and is used to define the next value in the sequence.

If the `value_sequence` does not have the reserved word **parallel**, it is produced as a single sequence of values by a single logical thread of control. If the reserved word **parallel** is present in the `value_sequence`, the enclosing `reduction_attribute_reference` is a parallel construct, and the sequence of values is generated by a parallel iteration (as defined in 8.5, 8.5.1, and 8.5.2), as a set of non-empty, non-overlapping contiguous chunks (*subsequences*) with one logical thread of control (see Clause 12) associated with each subsequence. If there is a `chunk_specification`, it determines the maximum number of chunks, as defined in 8.5; otherwise the maximum number of chunks is implementation defined.

For a `value_sequence` *V*, the following attribute is defined:

*V*'Reduce(*Reducer*, *Initial\_Value*)

This attribute represents a *reduction expression*, and is in the form of a `reduction_attribute_reference`.

The evaluation of a use of this attribute begins by evaluating the parts of the `reduction_attribute_designator` (the *reducer\_name* *Reducer* and the *initial\_value\_expression* *Initial\_Value*), in an arbitrary order. It then initializes the *accumulator* of the reduction expression to the value of the *initial\_value\_expression* (the *initial value*). The `value_sequence` *V* is then evaluated.

If the `value_sequence` does not have the reserved word **parallel**, each value of the `value_sequence` is passed, in order, as the second (*Value*) parameter to a call on *Reducer*, with the first (*Accumulator*) parameter being the prior value of the accumulator, saving the result as the new value of the accumulator. The reduction expression yields the final value of the accumulator.

If the reserved word **parallel** is present in a `value_sequence`, then the (parallel) reduction expression is a parallel construct and the sequence has been partitioned into one or more subsequences (see above) each with its own separate logical thread of control.

Each logical thread of control creates a local accumulator for processing its subsequence. The accumulator for a subsequence is initialized to the first value of the subsequence, and calls on *Reducer* start with the second value of the subsequence (if any). The result for the subsequence is the final value of its local accumulator.

After all logical threads of control of a parallel reduction expression have completed, *Reducer* is called for each subsequence, in the original sequence order, passing the local accumulator for that subsequence as the second (*Value*) parameter, and the overall accumulator (initialized above to the initial value) as the first (*Accumulator*) parameter,

with the result saved back in the overall accumulator. The parallel reduction expression yields the final value of the overall accumulator.

If the evaluation of the `value_sequence` yields an empty sequence of values, the reduction expression yields the initial value.

If an exception is propagated by one of the calls on Reducer, that exception is propagated from the reduction expression. If different exceptions are propagated in different logical threads of control, one is chosen arbitrarily to be propagated from the reduction expression as a whole.

For a prefix `X` of an array type (after any implicit dereference), or that denotes an iterable container object (see 8.5.1), the following attributes are defined:

**X'Reduce**(Reducer, Initial\_Value)

X'Reduce is a reduction expression that yields a result equivalent to replacing the prefix of the attribute with the `value_sequence`:

```
[for Item of X => Item]
```

**X'Parallel\_Reduce**(Reducer, Initial\_Value)

X'Parallel\_Reduce is a reduction expression that yields a result equivalent to replacing the attribute identifier with Reduce and the prefix of the attribute with the `value_sequence`:

```
[parallel for Item of X => Item]
```

*Bounded (Run-Time) Errors*

For a parallel reduction expression, it is a bounded error if the reducer subprogram is not associative. That is, for any arbitrary values of subtype *Value\_Type* `A`, `B`, `C` and a reducer function `R`, the result of `R (A, R (B, C))` should produce a result equal to `R (R (A, B), C)`; it is a bounded error if `R` does not. The possible consequences are Program\_Error, or a result that does not match the equivalent sequential reduction expression due to the order of calls on the reducer subprogram being unspecified in the overall reduction. Analogous rules apply in the case of a reduction procedure.

**EXAMPLE 1** Example of an expression function that returns its result as a reduction expression:

```
function Factorial(N : Natural) return Natural is
  ([for J in 1..N => J] 'Reduce("*", 1));
```

**EXAMPLE 2** Example of a reduction expression that computes the Sine of `X` using a Taylor expansion:

```
function Sine (X : Float; Num_Terms : Positive := 5) return Float is
  ([for I in 1..Num_Terms =>
    (-1.0)**(I-1) * X**(2*I-1)/Float(Factorial(2*I-1))] 'Reduce("+", 0.0));
```

**EXAMPLE 3** Example of a reduction expression that outputs the sum of squares:

```
Put_Line ("Sum of Squares is" &
  Integer'Image([for I in 1 .. 10 => I**2] 'Reduce("+", 0)));
```

**EXAMPLE 4** Example of a reduction expression used to compute the value of `Pi`:

```
-- See 6.5.7.
function Pi (Number_Of_Steps : Natural := 10_000) return Real is
  (1.0 / Real (Number_Of_Steps) *
    [for I in 1 .. Number_Of_Steps =>
      (4.0 / (1.0 + ((Real (I) - 0.5) *
        (1.0 / Real (Number_Of_Steps)))**2))]
    'Reduce("+", 0.0));
```

**EXAMPLE 5** Example of a reduction expression used to calculate the sum of elements of an array of integers:

```
A'Reduce("+", 0) -- See 7.3.3.
```

**EXAMPLE 6** Example of a reduction expression used to determine if all elements in a two-dimensional array of booleans are set to true:

```
Grid'Reduce("and", True) -- See 6.6.
```

EXAMPLE 7 Example of a reduction expression used to calculate the minimum value of an array of integers in parallel:

```
A'Parallel_Reduce(Integer'Min, Integer'Last)
```

EXAMPLE 8 Example of a parallel reduction expression used to calculate the mean of the elements of a two-dimensional array of subtype Matrix (see 6.6) that are greater than 100.0:

```
type Accumulator is record
  Sum    : Real; -- See 6.5.7.
  Count  : Integer;
end record;

function Accumulate (L, R : Accumulator) return Accumulator is
  (Sum    => L.Sum    + R.Sum,
   Count => L.Count + R.Count);

function Average_of_Values_Greater_Than_100 (M : Matrix) return Real is
  (declare
    Acc : constant Accumulator :=
      [parallel for Val of M when Val > 100.0 => (Val, 1)]
      'Reduce(Accumulate, (Sum => 0, Count => 0));
  begin
    Acc.Sum / Real(Acc.Count));
```

## 7.6 Type Conversions

Explicit type conversions, both value conversions and view conversions, are allowed between closely related types as defined below. This subclause also defines rules for value and view conversions to a particular subtype of a type, both explicit ones and those implicit in other constructs.

*Syntax*

```
type_conversion ::=
  subtype_mark(expression)
| subtype_mark(name)
```

The *target subtype* of a type\_conversion is the subtype denoted by the subtype\_mark. The *operand* of a type\_conversion is the expression or name within the parentheses; its type is the *operand type*.

One type is *convertible* to a second type if a type\_conversion with the first type as operand type and the second type as target type is legal according to the rules of this subclause. Two types are convertible if each is convertible to the other.

A type\_conversion is called a *view conversion* if both its target type and operand type are tagged, or if it appears in a call as an actual parameter of mode **out** or **in out**; other type\_conversions are called *value conversions*.

*Name Resolution Rules*

The operand of a type\_conversion is expected to be of any type.

The operand of a view conversion is interpreted only as a name; the operand of a value conversion is interpreted as an expression.

*Legality Rules*

In a view conversion for an untagged type, the target type shall be convertible (back) to the operand type.

If there is a type (other than a root numeric type) that is an ancestor of both the target type and the operand type, or both types are class-wide types, then at least one of the following rules shall apply:

- The target type shall be untagged; or
- The operand type shall be covered by or descended from the target type; or



- The operand type shall be a class-wide type that covers the target type; or
- The operand and target types shall both be class-wide types and the specific type associated with at least one of them shall be an interface type.

If there is no type (other than a root numeric type) that is the ancestor of both the target type and the operand type, and they are not both class-wide types, one of the following rules shall apply:

- If the target type is a numeric type, then the operand type shall be a numeric type.
- If the target type is an array type, then the operand type shall be an array type. Further:
  - The types shall have the same dimensionality;
  - Corresponding index types shall be convertible;
  - The component subtypes shall statically match;
  - If the component types are anonymous access types, then the accessibility level of the operand type shall not be statically deeper than that of the target type;
  - Neither the target type nor the operand type shall be limited;
  - If the target type of a view conversion has aliased components, then so shall the operand type; and
  - The operand type of a view conversion shall not have a tagged, private, or volatile subcomponent.
- If the target type is *universal\_access*, then the operand type shall be an access type.
- If the target type is a general access-to-object type, then the operand type shall be *universal\_access* or an access-to-object type. Further, if the operand type is not *universal\_access*:
  - If the target type is an access-to-variable type, then the operand type shall be an access-to-variable type;
  - If the target designated type is tagged, then the operand designated type shall be convertible to the target designated type;
  - If the target designated type is not tagged, then the designated types shall be the same, and either:
    - the designated subtypes shall statically match; or
    - the designated type shall be discriminated in its full view and unconstrained in any partial view, and one of the designated subtypes shall be unconstrained;
  - The accessibility level of the operand type shall not be statically deeper than that of the target type, unless the target type is an anonymous access type of a stand-alone object. If the target type is that of such a stand-alone object, the accessibility level of the operand type shall not be statically deeper than that of the declaration of the stand-alone object.
- If the target type is a pool-specific access-to-object type, then the operand type shall be *universal\_access*.
- If the target type is an access-to-subprogram type, then the operand type shall be *universal\_access* or an access-to-subprogram type. Further, if the operand type is not *universal\_access*:
  - The designated profiles shall be subtype conformant.
  - The accessibility level of the operand type shall not be statically deeper than that of the target type. If the operand type is declared within a generic body, the target type shall be declared within the generic body.
  - If the target type has a Global aspect other than **in out all** or Unspecified, then each mode of the Global aspect of the operand type shall identify a subset of the variables identified by the corresponding mode of the target type Global aspect, or by the **in out** mode of the target type Global aspect.

- If the target type is nonblocking, the operand type shall be nonblocking.

In addition to the places where Legality Rules normally apply (see 15.3), these rules apply also in the private part of an instance of a generic unit.

#### Static Semantics

A `type_conversion` that is a value conversion denotes the value that is the result of converting the value of the operand to the target subtype.

A `type_conversion` that is a view conversion denotes a view of the object denoted by the operand. This view is a variable of the target type if the operand denotes a variable; otherwise, it is a constant of the target type.

The nominal subtype of a `type_conversion` is its target subtype.

#### Dynamic Semantics

For the evaluation of a `type_conversion` that is a value conversion, the operand is evaluated, and then the value of the operand is *converted* to a *corresponding* value of the target type, if any. If there is no value of the target type that corresponds to the operand value, `Constraint_Error` is raised; this can only happen on conversion to a modular type, and only when the operand value is outside the base range of the modular type. Additional rules follow:

- **Numeric Type Conversion**
  - If the target and the operand types are both integer types, then the result is the value of the target type that corresponds to the same mathematical integer as the operand.
  - If the target type is a decimal fixed point type, then the result is truncated (toward 0) if the value of the operand is not a multiple of the *small* of the target type.
  - If the target type is some other real type, then the result is within the accuracy of the target type (see G.2, for implementations that support the Numerics Annex).
  - If the target type is an integer type and the operand type is real, the result is rounded to the nearest integer (away from zero if exactly halfway between two integers).
- **Enumeration Type Conversion**
  - The result is the value of the target type with the same position number as that of the operand value.
- **Array Type Conversion**
  - If the target subtype is a constrained array subtype, then a check is made that the length of each dimension of the value of the operand equals the length of the corresponding dimension of the target subtype. The bounds of the result are those of the target subtype.
  - If the target subtype is an unconstrained array subtype, then the bounds of the result are obtained by converting each bound of the value of the operand to the corresponding index type of the target type. For each nonnull index range, a check is made that the bounds of the range belong to the corresponding index subtype.
  - In either array case, the value of each component of the result is that of the matching component of the operand value (see 7.5.2).
  - If the component types of the array types are anonymous access types, then a check is made that the accessibility level of the operand type is not deeper than that of the target type.
- **Composite (Non-Array) Type Conversion**
  - The value of each nondiscriminant component of the result is that of the matching component of the operand value.

- The tag of the result is that of the operand. If the operand type is class-wide, a check is made that the tag of the operand identifies a (specific) type that is covered by or descended from the target type.
- For each discriminant of the target type that corresponds to a discriminant of the operand type, its value is that of the corresponding discriminant of the operand value; if it corresponds to more than one discriminant of the operand type, a check is made that all these discriminants are equal in the operand value.
- For each discriminant of the target type that corresponds to a discriminant that is specified by the `derived_type_definition` for some ancestor of the operand type (or if class-wide, some ancestor of the specific type identified by the tag of the operand), its value in the result is that specified by the `derived_type_definition`.
- For each discriminant of the operand type that corresponds to a discriminant that is specified by the `derived_type_definition` for some ancestor of the target type, a check is made that in the operand value it equals the value specified for it.
- For each discriminant of the result, a check is made that its value belongs to its subtype.
- Access Type Conversion
  - For an access-to-object type, a check is made that the accessibility level of the operand type is not deeper than that of the target type, unless the target type is an anonymous access type of a stand-alone object. If the target type is that of such a stand-alone object, a check is made that the accessibility level of the operand type is not deeper than that of the declaration of the stand-alone object; then if the check succeeds, the accessibility level of the target type becomes that of the operand type.
  - If the operand value is null, the result of the conversion is the null value of the target type.
  - If the operand value is not null, then the result designates the same object (or subprogram) as is designated by the operand value, but viewed as being of the target designated subtype (or profile); any checks associated with evaluating a conversion to the target designated subtype are performed.

After conversion of the value to the target type, if the target subtype is constrained, a check is performed that the value satisfies this constraint. If the target subtype excludes null, then a check is made that the value is not null. If predicate checks are enabled for the target subtype (see 6.2.4), a check is performed that the value satisfies the predicates of the target subtype, unless the conversion is:

- a view conversion that is the target of an assignment statement and is not referenced with a `target_name`, or an actual parameter of mode **out**; or
- an implicit subtype conversion of an actual parameter of mode **out** to the nominal subtype of its formal parameter.

For the evaluation of a view conversion, the operand `name` is evaluated, and a new view of the object denoted by the operand is created, whose type is the target type; if the target type is composite, checks are performed as above for a value conversion.

The properties of this new view are as follows:

- If the target type is composite, the bounds or discriminants (if any) of the view are as defined above for a value conversion; each nondiscriminant component of the view denotes the matching component of the operand object; the subtype of the view is constrained if either the target subtype or the operand object is constrained, or if the target subtype is indefinite, or if the operand type is a descendant of the target type and has discriminants that were not inherited from the target type;
- If the target type is tagged, then an assignment to the view assigns to the corresponding part of the object denoted by the operand; otherwise, an assignment to the view assigns to the object, after converting the assigned value to the subtype of the object (which can raise `Constraint_Error`);

- Reading the value of the view yields the result of converting the value of the operand object to the target subtype (which can raise `Constraint_Error`), except if the object is of an elementary type and the view conversion is passed as an **out** parameter; in this latter case, the value of the operand object may be used to initialize the formal parameter without checking against any constraint of the target subtype (as described more precisely in 9.4.1).

If an `Accessibility_Check` fails, `Program_Error` is raised. If a predicate check fails, the effect is as defined in 6.2.4. Any other check associated with a conversion raises `Constraint_Error` if it fails.

Conversion to a type is the same as conversion to an unconstrained subtype of the type.

Evaluation of a value conversion of an object either creates a new anonymous object (similar to the object created by the evaluation of an aggregate or a function call) or yields a new view of the operand object without creating a new object:

- If the target type is a by-reference type and there is a type that is an ancestor of both the target type and the operand type then no new object is created;
- If the target type is an array type having aliased components and the operand type is an array type having unaliased components, then a new object is created;
- If the target type is an elementary type, then a new object is created;
- Otherwise, it is unspecified whether a new object is created.

If a new object is created, then the initialization of that object is an assignment operation.

NOTE 1 In addition to explicit `type_conversions`, type conversions are performed implicitly in situations where the expected type and the actual type of a construct differ, as is permitted by the type resolution rules (see 11.6). For example, an integer literal is of the type *universal integer*, and is implicitly converted when assigned to a target of some specific integer type. Similarly, an actual parameter of a specific tagged type is implicitly converted when the corresponding formal parameter is of a class-wide type.

Even when the expected and actual types are the same, implicit subtype conversions are performed to adjust the array bounds (if any) of an operand to match the desired target subtype, or to raise `Constraint_Error` if the (possibly adjusted) value does not satisfy the constraints of the target subtype.

NOTE 2 A ramification of the overload resolution rules is that the operand of an (explicit) `type_conversion` cannot be an allocator, an aggregate, a `string_literal`, a `character_literal`, or an `attribute_reference` for an `Access` or `Unchecked_Access` attribute. Similarly, such an expression enclosed by parentheses is not allowed. A `qualified_expression` (see 7.7) can be used instead of such a `type_conversion`.

EXAMPLE 1 *Examples of numeric type conversion:*

```
Real(2*J)      -- value is converted to floating point
Integer(1.6)   -- value is 2
Integer(-0.4)  -- value is 0
```

EXAMPLE 2 *Example of conversion between derived types:*

```
type A_Form is new B_Form;
X : A_Form;
Y : B_Form;
X := A_Form(Y);
Y := B_Form(X);  -- the reverse conversion
```

EXAMPLE 3 *Examples of conversions between array types:*

```
type Sequence is array (Integer range <>) of Integer;
subtype Dozen is Sequence(1 .. 12);
Ledger : array(1 .. 100) of Integer;

Sequence(Ledger)           -- bounds are those of Ledger
Sequence(Ledger(31 .. 42)) -- bounds are 31 and 42
Dozen(Ledger(31 .. 42))    -- bounds are those of Dozen
```

## 7.7 Qualified Expressions

A `qualified_expression` is used to state explicitly the type, and to verify the subtype, of an operand that is either an expression or an aggregate.

### Syntax

```
qualified_expression ::=
  subtype_mark'(expression) | subtype_mark'aggregate
```

### Name Resolution Rules

The expected type for the *operand* (the expression or aggregate) is determined by the `subtype_mark`. Furthermore, the operand shall resolve to be either the specified expected type or a universal type that covers it.

### Static Semantics

If the operand of a `qualified_expression` denotes an object, the `qualified_expression` denotes a constant view of that object. The nominal subtype of a `qualified_expression` is the subtype denoted by the `subtype_mark`.

### Dynamic Semantics

The evaluation of a `qualified_expression` evaluates the operand (and if of a universal type, converts it to the type determined by the `subtype_mark`) and checks that its value belongs to the subtype denoted by the `subtype_mark`. The exception `Constraint_Error` is raised if this check fails. Furthermore, if predicate checks are enabled for the subtype denoted by the `subtype_mark`, a check is performed as defined in 6.2.4 that the value satisfies the predicates of the subtype.

NOTE When a given context does not uniquely identify an expected type, a `qualified_expression` can be used to do so. In particular, if an overloaded name or aggregate is passed to an overloaded subprogram, it can be necessary to qualify the operand to resolve its type.

EXAMPLE *Examples of disambiguating expressions using qualification:*

```
type Mask is (Fix, Dec, Exp, Signif);
type Code is (Fix, Cla, Dec, Tnz, Sub);

Print (Mask'(Dec)); -- Dec is of type Mask
Print (Code'(Dec)); -- Dec is of type Code

for J in Code'(Fix) .. Code'(Dec) loop ... -- qualification is necessary for
either Fix or Dec
for J in Code range Fix .. Dec loop ... -- qualification unnecessary
for J in Code'(Fix) .. Dec loop ... -- qualification unnecessary for
Dec

Dozen'(1 | 3 | 5 | 7 => 2, others => 0) -- see 7.6
```

## 7.8 Allocators

The evaluation of an allocator creates an object and yields an access value that designates the object.

### Syntax

```
allocator ::=
  new[subpool_specification] subtype_indication
  | new[subpool_specification] qualified_expression

subpool_specification ::= (subpool_handle_name)
```

For an allocator with a `subtype_indication`, the `subtype_indication` shall not specify a `null_exclusion`.

#### Name Resolution Rules

The expected type for an allocator shall be a single access-to-object type with designated type *D* such that either *D* covers the type determined by the `subtype_mark` of the `subtype_indication` or `qualified_expression`, or the expected type is anonymous and the determined type is *D*'Class. A `subpool_handle_name` is expected to be of any type descended from `Subpool_Handle`, which is the type used to identify a subpool, declared in package `System.Storage_Pools.Subpools` (see 16.11.4).

#### Legality Rules

An *initialized* allocator is an allocator with a `qualified_expression`. An *uninitialized* allocator is one with a `subtype_indication`. In the `subtype_indication` of an uninitialized allocator, a constraint is permitted only if the `subtype_mark` denotes an unconstrained composite subtype; if there is no constraint, then the `subtype_mark` shall denote a definite subtype.

If the type of the allocator is an access-to-constant type, the allocator shall be an initialized allocator.

If a `subpool_specification` is given, the type of the storage pool of the access type shall be a descendant of `Root_Storage_Pool_With_Subpools`.

If the designated type of the type of the allocator is class-wide, the accessibility level of the type determined by the `subtype_indication` or `qualified_expression` shall not be statically deeper than that of the type of the allocator.

If the subtype determined by the `subtype_indication` or `qualified_expression` of the allocator has one or more access discriminants, then the accessibility level of the anonymous access type of each access discriminant shall not be statically deeper than that of the type of the allocator (see 6.10.2).

An allocator shall not be of an access type for which the `Storage_Size` has been specified by a static expression with value zero or is defined by the language to be zero.

If the designated type of the type of the allocator is limited, then the allocator shall not be used to define the value of an access discriminant, unless the discriminated type is immutably limited (see 10.5).

In addition to the places where Legality Rules normally apply (see 15.3), these rules apply also in the private part of an instance of a generic unit.

#### Static Semantics

If the designated type of the type of the allocator is elementary, then the subtype of the created object is the designated subtype. If the designated type is composite, then the subtype of the created object is the designated subtype when the designated subtype is constrained or there is an ancestor of the designated type that has a constrained partial view; otherwise, the created object is constrained by its initial value (even if the designated subtype is unconstrained with defaults).

#### Dynamic Semantics

For the evaluation of an initialized allocator, the evaluation of the `qualified_expression` is performed first. An object of the designated type is created and the value of the `qualified_expression` is converted to the designated subtype and assigned to the object.

For the evaluation of an uninitialized allocator, the elaboration of the `subtype_indication` is performed first. Then:

- If the designated type is elementary, an object of the designated subtype is created and any implicit initial value is assigned;
- If the designated type is composite, an object of the designated type is created with tag, if any, determined by the `subtype_mark` of the `subtype_indication`. This object is then initialized by default (see 6.3.1) using the `subtype_indication` to determine its nominal subtype. A check



is made that the value of the object belongs to the designated subtype. `Constraint_Error` is raised if this check fails. This check and the initialization of the object are performed in an arbitrary order.

For any `allocator`, if the designated type of the type of the `allocator` is class-wide, then a check is made that the master of the type determined by the `subtype_indication`, or by the tag of the value of the `qualified_expression`, includes the elaboration of the type of the `allocator`. If any part of the subtype determined by the `subtype_indication` or `qualified_expression` of the `allocator` (or by the tag of the value if the type of the `qualified_expression` is class-wide) has one or more access discriminants, then a check is made that the accessibility level of the anonymous access type of each access discriminant is not deeper than that of the type of the `allocator`. `Program_Error` is raised if either such check fails.

If the object to be created by an `allocator` has a controlled or protected part, and the finalization of the collection of the type of the `allocator` (see 10.6.1) has started, `Program_Error` is raised.

If the object to be created by an `allocator` contains any tasks, and the master of the type of the `allocator` is completed, and all of the dependent tasks of the master are terminated (see 12.3), then `Program_Error` is raised.

If the `allocator` includes a `subpool_handle_name`, `Constraint_Error` is raised if the subpool handle is **null**. `Program_Error` is raised if the subpool does not *belong* (see 16.11.4) to the storage pool of the access type of the `allocator`.

If the created object contains any tasks, they are activated (see 12.2). Finally, an access value that designates the created object is returned.

#### Bounded (Run-Time) Errors

It is a bounded error if the finalization of the collection of the type (see 10.6.1) of the `allocator` has started. If the error is detected, `Program_Error` is raised. Otherwise, the allocation proceeds normally.

NOTE 1 Allocators cannot create objects of an abstract type. See 6.9.3.

NOTE 2 If any part of the created object is controlled, the initialization includes calls on corresponding `Initialize` or `Adjust` procedures. See 10.6.

NOTE 3 As explained in 16.11, the storage for an object allocated by an `allocator` comes from a storage pool (possibly user defined). The exception `Storage_Error` is raised by an `allocator` if there is not enough storage. Instances of `Unchecked_Deallocation` can be used to explicitly reclaim storage.

NOTE 4 Implementations can, if desired, provide garbage collection.

EXAMPLE *Examples of allocators:*

```
new Cell'(0, null, null)                -- initialized explicitly,
see 6.10.1
new Cell'(Value => 0, Succ => null, Pred => null) -- initialized explicitly
new Cell                                -- not initialized

new Matrix(1 .. 10, 1 .. 20)            -- the bounds only are given
new Matrix'(1 .. 10 => (1 .. 20 => 0.0)) -- initialized explicitly

new Buffer(100)                         -- the discriminant only is
given
new Buffer'(Size => 80, Pos => 0, Value => (1 .. 80 => 'A')) -- initialized
explicitly

Expr_Ptr'(new Literal)                  -- allocator for access-to-class-wide
type, see 6.9.1
Expr_Ptr'(new Literal'(Expression with 3.5)) -- initialized explicitly
```

## 7.9 Static Expressions and Static Subtypes

Certain expressions of a scalar or string type are defined to be static. Similarly, certain discrete ranges are defined to be static, and certain scalar and string subtypes are defined to be static subtypes. *Static* means determinable at compile time, using the declared properties or values of the program entities.

## Static Semantics

A static expression is a scalar or string expression that is one of the following:

- a `numeric_literal` of a numeric type;
- a `string_literal` of a static string subtype;
- a name that denotes the declaration of a static constant;
- a name that denotes a named number, and that is interpreted as a value of a numeric type;
- a `function_call` whose *function\_name* or *function\_prefix* statically denotes a static function, and whose actual parameters, if any (whether given explicitly or by default), are all static expressions;
- an `attribute_reference` that denotes a scalar value, and whose *prefix* denotes a static scalar subtype;
- an `attribute_reference` whose *prefix* statically names a statically constrained array object or array subtype, and whose *attribute\_designator* is First, Last, or Length, with an optional dimension;
- an `attribute_reference` whose *prefix* denotes a non-generic entity that is not declared in a generic unit, and whose *attribute\_designator* is Nonblocking;
- a `type_conversion` whose *subtype\_mark* denotes a static (scalar or string) subtype, and whose operand is a static expression;
- a `qualified_expression` whose *subtype\_mark* denotes a static (scalar or string) subtype, and whose operand is a static expression;
- a membership test whose *tested\_simple\_expression* is a static expression, and whose *membership\_choice\_list* consists only of membership choices that are either static *choice\_simple\_expressions*, static ranges, or *subtype\_marks* that denote a static (scalar or string) subtype;
- a short-circuit control form both of whose relations are static expressions;
- a `conditional_expression` all of whose conditions, *selecting\_expressions*, and *dependent\_expressions* are static expressions;
- a `declare_expression` whose *body\_expression* is static and each of whose declarations, if any, is either the declaration of a static constant or is an *object\_renaming\_declaration* with an *object\_name* that statically names the renamed object;
- a static expression enclosed in parentheses.

A name *statically denotes* an entity if it denotes the entity and:

- It is a `direct_name`, expanded name, or `character_literal`, and it denotes a declaration other than a *renaming\_declaration*; or
- It is an `attribute_reference` whose *prefix* statically denotes some entity; or
- It is a *target\_name* (see 8.2.1) in an *assignment\_statement* whose *variable\_name* statically denotes some entity; or
- It denotes a *renaming\_declaration* with a name that statically denotes the renamed entity.

A name *statically names* an object if it:

- statically denotes the declaration of an object (possibly through one or more renames);
- is a `selected_component` whose *prefix* statically names an object, there is no implicit dereference of the *prefix*, and the *selector\_name* does not denote a *component\_declaration* occurring within a *variant\_part*; or

- is an `indexed_component` whose `prefix` statically names an object, there is no implicit dereference of the `prefix`, the object is statically constrained, and the index expressions of the object are static and have values that are within the range of the index constraint.

For an entity other than an object, a `name` statically names an entity if the `name` statically denotes the entity.

A *static function* is one of the following:

- a predefined operator whose parameter and result types are all scalar types none of which are descendants of formal scalar types;
- a predefined relational operator whose parameters are of a string type that is not a descendant of a formal array type;
- a predefined concatenation operator whose result type is a string type that is not a descendant of a formal array type;
- a shifting or rotating function associated with a modular type declared in package `Interfaces` (see B.2);
- an enumeration literal;
- a static expression function (see 9.8);
- a language-defined attribute that is a function, if the `prefix` denotes a static scalar subtype, and if the parameter and result types are scalar.

In any case, a generic formal subprogram is not a static function.

A *static constant* is a constant view declared by a full constant declaration or an `object_renaming_declaration` with a static nominal subtype, having a value defined by a static scalar expression or by a static string expression, and which satisfies any constraint or predicate that applies to the nominal subtype.

A *static range* is a range whose bounds are static expressions, or a `range_attribute_reference` that is equivalent to such a range. A *static discrete range* is one that is a static range or is a `subtype_indication` that defines a static scalar subtype. The base range of a scalar type is a static range, unless the type is a descendant of a formal scalar type.

A *static subtype* is either a *static scalar subtype* or a *static string subtype*. A static scalar subtype is an unconstrained scalar subtype whose type is not a descendant of a formal type, or a constrained scalar subtype formed by imposing a compatible static constraint on a static scalar subtype. A static string subtype is an unconstrained string subtype whose index subtype and component subtype are static, or a constrained string subtype formed by imposing a compatible static constraint on a static string subtype. In any case, the subtype of a generic formal object of mode **in out**, and the result subtype of a generic formal function, are not static. Also, a subtype is not static if any `Dynamic_Predicate` specifications apply to it.

The different kinds of *static constraint* are defined as follows:

- A null constraint is always static;
- A scalar constraint is static if it has no `range_constraint`, or one with a static range;
- An index constraint is static if each `discrete_range` is static, and each index subtype of the corresponding array type is static;
- A discriminant constraint is static if each `expression` of the constraint is static, and the subtype of each discriminant is static.

In any case, the constraint of the first subtype of a scalar formal type is neither static nor null.

A subtype is *statically constrained* if it is constrained, and its constraint is static. An object is *statically constrained* if its nominal subtype is statically constrained, or if it is a static string constant.

*Legality Rules*

An expression is *statically unevaluated* if it is part of:

- the right operand of a static short-circuit control form whose value is determined by its left operand; or
- a *dependent\_expression* of an *if\_expression* whose associated condition is static and equals False; or
- a condition or *dependent\_expression* of an *if\_expression* where the condition corresponding to at least one preceding *dependent\_expression* of the *if\_expression* is static and equals True; or
- a *dependent\_expression* of a *case\_expression* whose *selecting\_expression* is static and whose value is not covered by the corresponding *discrete\_choice\_list*; or
- a *choice\_simple\_expression* (or a *simple\_expression* of a range that occurs as a *membership\_choice* of a *membership\_choice\_list*) of a static membership test that is preceded in the enclosing *membership\_choice\_list* by another item whose individual membership test (see 7.5.2) statically yields True.

A static expression is evaluated at compile time except when it is statically unevaluated. The compile-time evaluation of a static expression is performed exactly, without performing *Overflow\_Checks*. For a static expression that is evaluated:

- The expression is illegal if its evaluation fails a language-defined check other than *Overflow\_Check*. For the purposes of this evaluation, the assertion policy is assumed to be *Check*.
- If the expression is not part of a larger static expression and the expression is expected to be of a single specific type, then its value shall be within the base range of its expected type. Otherwise, the value may be arbitrarily large or small.
- If the expression is of type *universal\_real* and its expected type is a decimal fixed point type, then its value shall be a multiple of the *small* of the decimal type. This restriction does not apply if the expected type is a descendant of a formal scalar type (or a corresponding actual type in an instance).

In addition to the places where Legality Rules normally apply (see 15.3), the above restrictions also apply in the private part of an instance of a generic unit.

*Implementation Requirements*

For a real static expression that is not part of a larger static expression, and whose expected type is not a descendant of a formal type, the implementation shall round or truncate the value (according to the *Machine\_Rounds* attribute of the expected type) to the nearest machine number of the expected type; if the value is exactly half-way between two machine numbers, the rounding performed is implementation-defined. If the expected type is a descendant of a formal type, or if the static expression appears in the body of an instance of a generic unit and the corresponding expression is nonstatic in the corresponding generic body, then no special rounding or truncating is required — normal accuracy rules apply (see Annex G).

*Implementation Advice*

For a real static expression that is not part of a larger static expression, and whose expected type is not a descendant of a formal type, the rounding should be the same as the default rounding for the target system.

NOTE 1 An expression can be static even if it occurs in a context where staticness is not required.

NOTE 2 A static (or run-time) *type\_conversion* from a real type to an integer type performs rounding. If the operand value is exactly half-way between two integers, the rounding is performed away from zero.

EXAMPLE *Examples of static expressions:*

```
1 + 1      -- 2
abs(-10)*3 -- 30

Kilo : constant := 1000;
Mega : constant := Kilo*Kilo;    -- 1_000_000
Long  : constant := Float'Digits*2;

Half_Pi   : constant := Pi/2;      -- see 6.3.2
Deg_To_Rad : constant := Half_Pi/90;
Rad_To_Deg : constant := 1.0/Deg_To_Rad; -- equivalent to
1.0/((3.14159_26536/2)/90)
```

## 7.9.1 Statically Matching Constraints and Subtypes

### *Static Semantics*

A constraint *statically matches* another constraint if:

- both are null constraints;
- both are static and have equal corresponding bounds or discriminant values;
- both are nonstatic and result from the same elaboration of a constraint of a subtype\_indication or the same evaluation of a range of a discrete\_subtype\_definition; or
- both are nonstatic and come from the same formal\_type\_declaration.

The Global or Global'Class aspects (see 9.1.2) of two entities *statically match* if both consist of a single global\_aspect\_definition where each is the reserved word **null**, or each is of the form “global\_mode global\_designator” with each global\_mode being the same sequence of reserved words and each global\_designator being the same reserved word, or each being a global\_name that statically names the same entity.

A subtype *statically matches* another subtype of the same type if they have statically matching constraints, all predicate specifications that apply to them come from the same declarations, Nonblocking aspects have the same value, global aspects statically match, Object\_Size (see 16.3) has been specified to have a nonconfirming value for either both or neither, and the nonconfirming values, if any, are the same, and, for access subtypes, either both or neither exclude null. Two anonymous access-to-object subtypes statically match if their designated subtypes statically match, and either both or neither exclude null, and either both or neither are access-to-constant. Two anonymous access-to-subprogram subtypes statically match if their designated profiles are subtype conformant, and either both or neither exclude null.

Two ranges of the same type *statically match* if both result from the same evaluation of a range, or if both are static and have equal corresponding bounds.

A constraint is *statically compatible* with a scalar subtype if it statically matches the constraint of the subtype, or if both are static and the constraint is compatible with the subtype. A constraint is *statically compatible* with an access or composite subtype if it statically matches the constraint of the subtype, or if the subtype is unconstrained.

Two statically matching subtypes are statically compatible with each other. In addition, a subtype *S1* is statically compatible with a subtype *S2* if:

- the constraint of *S1* is statically compatible with *S2*, and
- if *S2* excludes null, so does *S1*, and
- either:
  - all predicate specifications that apply to *S2* apply also to *S1*, or
  - both subtypes are static, every value that satisfies the predicates of *S1* also satisfies the predicates of *S2*, and it is not the case that both types each have at least one applicable

predicate specification, predicate checks are enabled (see 14.4.2) for *S2*, and predicate checks are not enabled for *S1*.

## 7.10 Image Attributes

An *image* of a value is a string representing the value in display form. The attributes *Image*, *Wide\_Image*, and *Wide\_Wide\_Image* are available to produce the image of a value as a *String*, *Wide\_String*, or *Wide\_Wide\_String* (respectively). User-defined images for a given type can be implemented by overriding the default implementation of the attribute *Put\_Image*.

### Static Semantics

For every subtype *S* of a type *T* other than *universal\_real* or *universal\_fixed*, the following type-related operational attribute is defined:

*S'Put\_Image*

*S'Put\_Image* denotes a procedure with the following specification:

```
procedure S'Put_Image
  (Buffer : in out
    Ada.Strings.Text_Buffers.Root_Buffer_Type'Class;
   Arg    : in T);
```

The default implementation of *S'Put\_Image* writes (using *Wide\_Wide\_Put*) an *image* of the value of *Arg*.

The *Put\_Image* attribute may be specified for any specific type *T* either via an *attribute\_definition\_clause* or via an *aspect\_specification* specifying the *Put\_Image* aspect of the type. The *Put\_Image* aspect is not inherited, but rather is implicitly composed for derived types, as defined below.

For an *aspect\_specification* or *attribute\_definition\_clause* specifying *Put\_Image*, the subtype of the *Arg* parameter shall be the first subtype or the base subtype if scalar, and the first subtype if not scalar.

The behavior of the default implementation of *S'Put\_Image* depends on the class of *T*.

For an untagged derived type, or a null extension, the default implementation of *T'Put\_Image* invokes the *Put\_Image* for its parent type on a conversion of the parameter of type *T* to the parent type.

For a nonderived elementary type, the implementation is equivalent to:

```
procedure Scalar_Type'Put_Image
  (Buffer : in out Ada.Strings.Text_Buffers.Root_Buffer_Type'Class;
   Arg    : in Scalar_Type) is
begin
  Buffer.Wide_Wide_Put (<described below>);
end Scalar_Type'Put_Image;
```

where the *Wide\_Wide\_String* value written out to the text buffer is defined as follows:

- For an integer type, the image written out is the corresponding decimal literal, without underlines, leading zeros, exponent, or trailing spaces, but with a single leading character that is either a minus sign or a space.
- For an enumeration type, the image written out is either the corresponding identifier in upper case or the corresponding character literal (including the two apostrophes); neither leading nor trailing spaces are included. For a *nongraphic character* (a value of a character type that has no enumeration literal associated with it), the value is a corresponding language-defined name in upper case (for example, the image of the nongraphic character identified as *nul* is "NUL" — the quotes are not part of the image).
- For a floating point type, the image written out is a decimal real literal best approximating the value (rounded away from zero if halfway between) with a single leading character that is either a minus sign or a space, a single digit (that is nonzero unless the value is zero), a decimal point, *S'Digits-1*



(see 6.5.8) digits after the decimal point (but one if S'Digits is one), an upper case E, the sign of the exponent (either + or -), and two or more digits (with leading zeros if necessary) representing the exponent. If S'Signed\_Zeros is True, then the leading character is a minus sign for a negatively signed zero.

- For a fixed point type, the image written out is a decimal real literal best approximating the value (rounded away from zero if halfway between) with a single leading character that is either a minus sign or a space, one or more digits before the decimal point (with no redundant leading zeros), a decimal point, and S'Aft (see 6.5.10) digits after the decimal point.
- For an access type (named or anonymous), the image written out depends on whether the value is **null**. If it is **null**, then the image is "NULL". Otherwise the image is a left parenthesis followed by "ACCESS", a space, and a sequence of graphic characters, other than space or right parenthesis, representing the location of the designated object, followed by a right parenthesis, as in "(ACCESS FF0012AC)".

For a nonnull type extension, the default implementation of T'Put\_Image depends on whether there exists a noninterface ancestor of T (other than T itself) for which the Put\_Image aspect has been directly specified. If so, then T'Put\_Image will generate an image based on extension aggregate syntax where the ancestor type of the extension aggregate is the nearest ancestor type whose Put\_Image aspect has been specified. If no such ancestor exists, then the default implementation of T'Put\_Image is the same as described below for a nonderived record type.

For a specific, nonderived composite type:

- If the default implementation of Put\_Image writes components, the order in which components are written is the same canonical order in which components of a composite type T are written out by the default implementation of T'Write. This is also the order that is used in determining the meaning of a positional aggregate of type T.
- For an array type T, the default implementation of T'Put\_Image generates an image based on named (not positional) array aggregate syntax (with '[' and ']' as the delimiters) using calls to the Put\_Image procedures of the index type(s) and the element type to generate images for values of those types.

The case of a null array is handled specially, using ranges for index bounds and "<>" as a syntactic component-value placeholder.

- For a record type (or, as indicated above, a type extension with no noninterface ancestor specifying Put\_Image), or a protected type, the default implementation of T'Put\_Image generates an image based on named (not positional) record aggregate syntax (except that for a protected type, the initial left parenthesis is followed by "PROTECTED with "). Component names are displayed in upper case, following the rules for the image of an enumeration value. Component values are displayed via calls to the component type's Put\_Image procedure.

The image written out for a record having no components (including any interface type) is "(NULL-RECORD)". The image written out for a componentless protected type is "(PROTECTED NULL RECORD)". In the case of a protected type T, a call to the default implementation of T'Put\_Image begins only one protected (read-only) action.

- For an undiscriminated task type, the default implementation of T'Put\_Image generates an image of the form "(TASK <task\_id\_image>)" where <task\_id\_image> is the result obtained by calling Task\_Identification.Image with the id of the given task and then passing that String to Characters.Conversions.To\_Wide\_Wide\_String.
- For a discriminated task type, the default implementation of T'Put\_Image also includes discriminant values, as in:

```
"(TASK <task_id_image> with D1 => 123, D2 => 456)"
```

For a class-wide type, the default implementation of T'Put\_Image generates an image based on qualified expression syntax. Wide\_Wide\_Put is called with Wide\_Wide\_Expanded\_Name of Arg'Tag. Then S'Put\_Image is called, where S is the specific type identified by Arg'Tag.

T'Put\_Image is the same for both the partial view and full view of T, if T has a partial view.

In the `parameter_and_result_profile` for the default implementation of Put\_Image, the subtype of the *Arg* parameter is the base subtype of *T* if *T* is a scalar type, and the first subtype otherwise. For an `aspect_specification` or `attribute_definition_clause` specifying Put\_Image, the subprogram name shall denote a nonabstract procedure whose second parameter is either of the first subtype of *T*, or as an option when *T* is scalar, the base subtype of *T*.

For every subtype *S* of a type *T*, the following attributes are defined:

S'Wide\_Wide\_Image

S'Wide\_Wide\_Image denotes a function with the following specification:

```
function S'Wide_Wide_Image(Arg : S'Base)
return Wide_Wide_String
```

S'Wide\_Wide\_Image calls S'Put\_Image passing *Arg* (which will typically store a sequence of character values in a text buffer) and then returns the result of retrieving the contents of that buffer with function Wide\_Wide\_Get. The lower bound of the result is one. Any exception propagated by the call of S'Put\_Image is propagated.

S'Wide\_Image

S'Wide\_Image denotes a function with the following specification:

```
function S'Wide_Image(Arg : S'Base)
return Wide_String
```

S'Wide\_Image calls S'Put\_Image passing *Arg* (which will typically store a sequence of character values in a text buffer) and then returns the result of retrieving the contents of that buffer with function Wide\_Get. The lower bound of the result is one. Any exception propagated by the call of S'Put\_Image is propagated.

S'Image S'Image denotes a function with the following specification:

```
function S'Image(Arg : S'Base)
return String
```

S'Image calls S'Put\_Image passing *Arg* (which will typically store a sequence of character values in a text buffer) and then returns the result of retrieving the contents of that buffer with function Get. The lower bound of the result is one. Any exception propagated by the call of S'Put\_Image is propagated.

For a prefix *X* of a type *T* other than *universal\_real* or *universal\_fixed*, the following attributes are defined:

X'Wide\_Wide\_Image

X'Wide\_Wide\_Image denotes the result of calling function S'Wide\_Wide\_Image with *Arg* being *X*, where *S* is the nominal subtype of *X*.

X'Wide\_Image

X'Wide\_Image denotes the result of calling function S'Wide\_Image with *Arg* being *X*, where *S* is the nominal subtype of *X*.

X'Image X'Image denotes the result of calling function S'Image with *Arg* being *X*, where *S* is the nominal subtype of *X*.

#### Implementation Permissions

An implementation may transform the image generated by the default implementation of S'Put\_Image for a composite subtype *S* in the following ways:

- If *S* is a composite subtype, the leading character of the image *M* of a component value or index value is a space, and the immediately preceding character (if any) is an open parenthesis, open bracket, or space, then the leading space of the image *M* may be omitted.

- If *S* is an array subtype, the low bound of the array in each dimension equals the low bound of the corresponding index subtype, and the array value is not a null array value, then positional array aggregate syntax may be used.
- If *S* is an array subtype and the given value can be displayed using `named_array_aggregate` syntax where some `discrete_choice_list` identifies more than one index value by identifying a sequence of one or more ranges and values separated by vertical bars, then this image may be generated instead; this may involve the reordering of component values.
- Similarly, if *S* is a record subtype (or a discriminated type) and the given value can be displayed using named component association syntax where the length of some `component_choice_list` is greater than one, then this image may be generated instead; this may involve the reordering of component values.
- Additional spaces (`Wide_Wide_Chars` with position 32), and calls to the `New_Line` operation of a text buffer, may be inserted to improve readability of the generated image, with the spaces inserted directly or via use of the `Increase_Indent` and `Decrease_Indent` procedures.
- For a string type, implementations may produce an image corresponding to a string literal.
- For an unchecked union type, implementations may raise `Program_Error` or produce some recognizable image (such as `"(UNCHECKED UNION)"`) that does not require reading the discriminants.

For each language-defined nonscalar type *T*, `T'Put_Image` may be specified.

#### Implementation Requirements

For each language-defined container type *T* (that is, each of the `Vector`, `List`, `Map`, `Set`, `Tree`, and `Holder` types defined in the various children of `Ada.Containers`), `T'Put_Image` shall be specified so that `T'Image` produces a result consistent with array aggregate syntax (using '[' and ']' as delimiters) as follows:

- `Vector` images shall be consistent with the default image of an array type with the same index and component types.
- `Map` images shall be consistent with named array aggregate syntax, using key value images in place of discrete choice names. For example, `[Key1 => Value1, Key2 => Value2]`.
- `Set`, `List`, and `Holder` images shall be consistent with positional array aggregate syntax. List elements shall occur in order within an image of a list. The image of an empty holder shall be `[]`.
- `Tree` images (and images of subtrees of trees) shall be consistent with positional array aggregate syntax. For example, `[[1, 2], [111, 222, 333]]`.

For each language-defined nonscalar type *T* that has a primitive language-defined `Image` function whose profile is type conformant with that of `T'Image` (for example, `Ada.Numerics.Float_Random.State` has such an `Image` function), `T'Put_Image` shall be specified so that `T'Image` yields the same result as that `Image` function.

#### Implementation Advice

For each language-defined private type *T*, `T'Image` should generate an image that would be meaningful based only on the relevant public interfaces, as opposed to requiring knowledge of the implementation of the private type.

## 8 Statements

A statement defines an action to be performed upon its execution.

This clause describes the general rules applicable to all statements. Some statements are discussed in later clauses: `Procedure_call_statements` and `return statements` are described in Clause 9. `Entry_call_statements`, `requeue_statements`, `delay_statements`, `accept_statements`, `select_statements`, and `abort_statements` are described in Clause 12. `Raise_`

statements are described in Clause 14, and `code_statements` in Clause 16. The remaining forms of statements are presented in this clause.

## 8.1 Simple and Compound Statements - Sequences of Statements

A statement is either simple or compound. A `simple_statement` encloses no other statement. A `compound_statement` can enclose `simple_statements` and other `compound_statements`. A *parallel construct* is a construct that introduces additional logical threads of control (see Clause 12) without creating a new task. Parallel loops (see 8.5) and `parallel_block_statements` (see 8.6.1) are parallel constructs.

### Syntax

```
sequence_of_statements ::= statement {statement} {label}
statement ::=
  {label} simple_statement | {label} compound_statement
simple_statement ::= null_statement
  | assignment_statement          | exit_statement
  | goto_statement                | procedure_call_statement
  | simple_return_statement        | entry_call_statement
  | requeue_statement             | delay_statement
  | abort_statement               | raise_statement
  | code_statement
compound_statement ::=
  if_statement                    | case_statement
  | loop_statement                | block_statement
  | extended_return_statement
  | parallel_block_statement
  | accept_statement              | select_statement
null_statement ::= null;
label ::= <<label_statement_identifier>>
statement_identifier ::= direct_name
```

The `direct_name` of a `statement_identifier` shall be an identifier (not an `operator_symbol`).

### Name Resolution Rules

The `direct_name` of a `statement_identifier` shall resolve to denote its corresponding implicit declaration (see below).

### Legality Rules

Distinct identifiers shall be used for all `statement_identifiers` that appear in the same body, including `inner_block_statements` but excluding `inner_program_units`.

### Static Semantics

For each `statement_identifier`, there is an implicit declaration (with the specified identifier) at the end of the `declarative_part` of the innermost `block_statement` or body that encloses the `statement_identifier`. The implicit declarations occur in the same order as the `statement_identifiers` occur in the source text. If a usage name denotes such an implicit declaration, the entity it denotes is the `label`, `loop_statement`, or `block_statement` with the given `statement_identifier`.

If one or more labels end a `sequence_of_statements`, an implicit `null_statement` follows the labels before any following constructs.

The execution of a `null_statement` has no effect.

A *transfer of control* is the run-time action of an `exit_statement`, `return statement`, `goto_statement`, or `requeue_statement`, selection of a `terminate_alternative`, raising of an exception, or an abort, which causes the next action performed to be one other than what would normally be expected from the other rules of the language. As explained in 10.6.1, a transfer of control can cause the execution of constructs to be completed and then left, which may trigger finalization.

The execution of a `sequence_of_statements` consists of the execution of the individual statements in succession until the `sequence` is completed.

Within a parallel construct, if a transfer of control out of the construct is initiated by one of the logical threads of control, an attempt is made to *cancel* all other logical threads of control initiated by the parallel construct. Once all other logical threads of control of the construct either complete or are canceled, the transfer of control occurs. If two or more logical threads of control of the same construct initiate such a transfer of control concurrently, one of them is chosen arbitrarily and the others are canceled.

When a logical thread of control is canceled, the cancellation causes it to complete as though it had performed a transfer of control to the point where it would have finished its execution. Such a cancellation is deferred while the logical thread of control is executing within an abort-deferred operation (see 12.8), and may be deferred further, but not past a point where the logical thread initiates a new nested parallel construct or reaches an exception handler that is outside such an abort-deferred operation.

#### Bounded (Run-Time) Errors

During the execution of a parallel construct, it is a bounded error to invoke an operation that is potentially blocking (see 12.5). `Program_Error` is raised if the error is detected by the implementation; otherwise, the execution of the potentially blocking operation can either proceed normally, or it can result in the indefinite blocking of some or all of the logical threads of control making up the current task.

NOTE A `statement_identifier` that appears immediately within the declarative region of a named `loop_statement` or an `accept_statement` is nevertheless implicitly declared immediately within the declarative region of the innermost enclosing body or `block_statement`; in other words, the expanded name for a named statement is not affected by whether the statement occurs inside or outside a named loop or an `accept_statement` — only nesting within `block_statements` is relevant to the form of its expanded name.

EXAMPLE *Examples of labeled statements:*

```
<<Here>> <<Ici>> <<Aqui>> <<Hier>> null;
<<After>> X := 1;
```

## 8.2 Assignment Statements

An `assignment_statement` replaces the current value of a variable with the result of evaluating an expression.

#### Syntax

```
assignment_statement ::=
  variable_name := expression;
```

The execution of an `assignment_statement` includes the evaluation of the expression and the *assignment* of the value of the expression into the *target*. An assignment operation (as opposed to an `assignment_statement`) is performed in other contexts as well, including object initialization and by-copy parameter passing. The *target* of an assignment operation is the view of the object to which a value is being assigned; the target of an `assignment_statement` is the variable denoted by the *variable\_name*.

## Name Resolution Rules

The *variable\_name* of an *assignment\_statement* is expected to be of any type. The expected type for the expression is the type of the target.

## Legality Rules

The target denoted by the *variable\_name* shall be a variable of a nonlimited type.

If the target is of a tagged class-wide type *T*'Class, then the expression shall either be dynamically tagged, or of type *T* and tag-indeterminate (see 6.9.2).

## Dynamic Semantics

For the execution of an *assignment\_statement*, the *variable\_name* and the expression are first evaluated in an arbitrary order.

When the type of the target is class-wide:

- If the expression is tag-indeterminate (see 6.9.2), then the controlling tag value for the expression is the tag of the target;
- Otherwise (the expression is dynamically tagged), a check is made that the tag of the value of the expression is the same as that of the target; if this check fails, *Constraint\_Error* is raised.

The value of the expression is converted to the subtype of the target. The conversion can raise an exception (see 7.6).

In cases involving controlled types, the target is finalized, and an anonymous object can be used as an intermediate in the assignment, as described in 10.6.1. In any case, the converted value of the expression is then *assigned* to the target, which consists of the following two steps:

- The value of the target becomes the converted value.
- If any part of the target is controlled, its value is adjusted as explained in 10.6.

NOTE The tag of an object never changes; in particular, an *assignment\_statement* does not change the tag of the target.

EXAMPLE 1 *Examples of assignment statements:*

```
Value := Max_Value - 1;
Shade := Blue;

Next_Frame(F)(M, N) := 2.5;      -- see 7.1.1
U := Dot_Product(V, W);         -- see 9.3

Writer := (Status => Open, Unit => Printer, Line_Count => 60); -- see 6.8.1
Next.all := (72074, null, Head); -- see 6.10.1
```

EXAMPLE 2 *Examples involving scalar subtype conversions:*

```
I, J : Integer range 1 .. 10 := 5;
K    : Integer range 1 .. 20 := 15;
...

I := J; -- identical ranges
K := J; -- compatible ranges
J := K; -- will raise Constraint_Error if K > 10
```

EXAMPLE 3 *Examples involving array subtype conversions:*

```
A : String(1 .. 31);
B : String(3 .. 33);
...

A := B; -- same number of components

A(1 .. 9) := "tar sauce";
A(4 .. 12) := A(1 .. 9); -- A(1 .. 12) = "tartar sauce"
```



Assignment\_statements are well-defined even in the case of overlapping slices of the same array, because the *variable\_name* and *expression* are both evaluated before copying the value into the variable. In the above example, an implementation yielding  $A(1..12) = \text{"tartartartar"}$  would be incorrect.

### 8.2.1 Target Name Symbols

@, known as the *target name* of an assignment statement, provides an abbreviation to avoid repetition of potentially long names in assignment statements.

#### Syntax

```
target_name ::= @
```

#### Name Resolution Rules

If a *target\_name* occurs in an assignment\_statement *A*, the *variable\_name* *V* of *A* is a complete context. The target name is a constant view of *V*, having the nominal subtype of *V*.

#### Legality Rules

A *target\_name* shall appear only in the expression of an assignment\_statement.

#### Dynamic Semantics

For the execution of an assignment\_statement with one or more *target\_names* appearing in its expression, the *variable\_name* *V* of the assignment\_statement is evaluated first to determine the object denoted by *V*, and then the expression of the assignment\_statement is evaluated with the evaluation of each *target\_name* yielding a constant view of the target whose properties are otherwise identical to those of the view provided by *V*. The remainder of the execution of the assignment\_statement is as given in 8.2.

EXAMPLE Examples of the use of target name symbols:

```
Board(1, 1) := @ + 1.0; -- An abbreviation for Board(1, 1) := Board(1, 1) +
1.0;
-- (Board is declared in 6.6.1).

My_Complex_Array : array (1 .. Max) of Complex; -- See 6.3.2, 6.8.
...
-- Square the element in the Count (see 6.3.1) position:
My_Complex_Array (Count) := (Re => @.Re**2 - @.Im**2,
Im => 2.0 * @.Re * @.Im);
-- A target_name can be used multiple times and
-- as a prefix if desired.
```

### 8.3 If Statements

An if\_statement selects for execution at most one of the enclosed sequences\_of\_statements, depending on the (truth) value of one or more corresponding conditions.

#### Syntax

```
if_statement ::=
  if condition then
    sequence_of_statements
  {elsif condition then
    sequence_of_statements}
  [else
    sequence_of_statements]
  end if;
```

#### Dynamic Semantics

For the execution of an if\_statement, the condition specified after **if**, and any conditions specified after **elsif**, are evaluated in succession (treating a final **else** as **elsif True then**), until one

evaluates to True or all conditions are evaluated and yield False. If a condition evaluates to True, then the corresponding `sequence_of_statements` is executed; otherwise, none of them is executed.

EXAMPLE *Examples of if statements:*

```

if Month = December and Day = 31 then
    Month := January;
    Day   := 1;
    Year  := Year + 1;
end if;

if Line_Too_Short then
    raise Layout_Error;
elsif Line_Full then
    New_Line;
    Put(Item);
else
    Put(Item);
end if;

if My_Car.Owner.Vehicle /= My_Car then           -- see 6.10.1
    Report ("Incorrect data");
end if;

```

## 8.4 Case Statements

A `case_statement` selects for execution one of a number of alternative `sequences_of_statements`; the chosen alternative is defined by the value of an expression.

*Syntax*

```

case_statement ::=
    case selecting_expression is
        case_statement_alternative
        {case_statement_alternative}
    end case;

case_statement_alternative ::=
    when discrete_choice_list =>
        sequence_of_statements

```

*Name Resolution Rules*

The *selecting\_expression* is expected to be of any discrete type. The expected type for each *discrete\_choice* is the type of the *selecting\_expression*.

*Legality Rules*

The *choice\_expressions*, *subtype\_indications*, and ranges given as *discrete\_choices* of a `case_statement` shall be static. A *discrete\_choice* **others**, if present, shall appear alone and in the last *discrete\_choice\_list*.

The possible values of the *selecting\_expression* shall be covered (see 6.8.1) as follows:

- If the *selecting\_expression* is a name (including a *type\_conversion*, *qualified\_expression*, or *function\_call*) having a static and constrained nominal subtype, then each non-**others** *discrete\_choice* shall cover only values in that subtype that satisfy its predicates (see 6.2.4), and each value of that subtype that satisfies its predicates shall be covered by some *discrete\_choice* (either explicitly or by **others**).
- If the type of the *selecting\_expression* is *root\_integer*, *universal\_integer*, or a descendant of a formal scalar type, then the `case_statement` shall have an **others** *discrete\_choice*.
- Otherwise, each value of the base range of the type of the *selecting\_expression* shall be covered (either explicitly or by **others**).

Two distinct `discrete_choices` of a `case_statement` shall not cover the same value.

*Dynamic Semantics*

For the execution of a `case_statement`, the *selecting\_expression* is first evaluated.

If the value of the *selecting\_expression* is covered by the `discrete_choice_list` of some `case_statement_alternative`, then the `sequence_of_statements` of the `_alternative` is executed.

Otherwise (the value is not covered by any `discrete_choice_list`, perhaps due to being outside the base range), `Constraint_Error` is raised.

NOTE The execution of a `case_statement` chooses one and only one alternative. Qualification of the expression of a `case_statement` by a static subtype can often be used to limit the number of choices that can be given explicitly.

EXAMPLE *Examples of case statements:*

```

case Sensor is
  when Elevation => Record_Elevation(Sensor_Value);
  when Azimuth   => Record_Azimuth  (Sensor_Value);
  when Distance  => Record_Distance (Sensor_Value);
  when others    => null;
end case;

case Today is
  when Mon       => Compute_Initial_Balance;
  when Fri       => Compute_Closing_Balance;
  when Tue .. Thu => Generate_Report(Today);
  when Sat .. Sun => null;
end case;

case Bin_Number(Count) is
  when 1         => Update_Bin(1);
  when 2         => Update_Bin(2);
  when 3 | 4 =>
    Empty_Bin(1);
    Empty_Bin(2);
  when others => raise Error;
end case;

```

## 8.5 Loop Statements

A `loop_statement` includes a `sequence_of_statements` that is to be executed repeatedly, zero or more times with the iterations running sequentially or concurrently with one another.

*Syntax*

```

loop_statement ::=
  [loop_statement_identifier:]
  [iteration_scheme] loop
    sequence_of_statements
  end loop [loop_identifier];

iteration_scheme ::= while condition
| for loop_parameter_specification
| for iterator_specification
| [parallel [aspect_specification]]
  for procedural_iterator
| parallel [(chunk_specification)] [aspect_specification]
  for loop_parameter_specification
| parallel [(chunk_specification)] [aspect_specification]
  for iterator_specification

chunk_specification ::=

```

```

integer_simple_expression
| defining_identifier in discrete_subtype_definition
loop_parameter_specification ::=
    defining_identifier in [reverse] discrete_subtype_definition
    [iterator_filter]
iterator_filter ::= when condition

```

If a loop\_statement has a *loop\_statement\_identifier*, then the identifier shall be repeated after the **end loop**; otherwise, there shall not be an identifier after the **end loop**.

An iteration\_scheme that begins with the reserved word **parallel** shall not have the reserved word **reverse** in its loop\_parameter\_specification.

#### Name Resolution Rules

In a chunk\_specification that is an *integer\_simple\_expression*, the *integer\_simple\_expression* is expected to be of any integer type.

#### Static Semantics

A loop\_parameter\_specification declares a *loop parameter*, which is an object whose subtype (and nominal subtype) is that defined by the discrete\_subtype\_definition.

In a chunk\_specification that has a discrete\_subtype\_definition, the chunk\_specification declares a *chunk parameter* object whose subtype (and nominal subtype) is that defined by the discrete\_subtype\_definition.

#### Dynamic Semantics

The *filter* of an *iterator construct* (a loop\_parameter\_specification, iterator\_specification, or procedural\_iterator) is defined to be *satisfied* when there is no iterator\_filter for the iterator construct, or when the condition of the iterator\_filter evaluates to True for a given iteration of the iterator construct.

If a sequence\_of\_statements of a loop\_statement with an iterator construct is said to be *conditionally executed*, then the statements are executed only when the filter of the iterator construct is satisfied.

The loop iterators loop\_parameter\_specification and iterator\_specification can also be used in contexts other than loop\_statements (for example, see 7.3.5 and 7.5.8). In such a context, the iterator *conditionally produces* values in the order specified for the associated construct below or in 8.5.2. The values produced are the values given to the loop parameter when the filter of the iterator construct is satisfied for that value. No value is produced when the condition of an iterator filter evaluates to False.

For the execution of a loop\_statement, the sequence\_of\_statements is executed zero or more times, until the loop\_statement is complete. The loop\_statement is complete when a transfer of control occurs that transfers control out of the loop, or, in the case of an iteration\_scheme, as specified below.

For the execution of a loop\_statement with a **while** iteration\_scheme, the condition is evaluated before each execution of the sequence\_of\_statements; if the value of the condition is True, the sequence\_of\_statements is executed; if False, the execution of the loop\_statement is complete.

If the reserved word **parallel** is present in the iteration\_scheme of a loop\_statement (a *parallel loop*), the iterations are partitioned into one or more *chunks*, each with its own separate logical thread of control (see Clause 12). If a chunk\_specification is present in a parallel loop, it is elaborated first, and the result of the elaboration determines the maximum number of chunks used for the parallel loop. If the chunk\_specification is an *integer\_simple\_expression*, the elaboration

evaluates the expression, and the value of the expression determines the maximum number of chunks. If a `discrete_subtype_definition` is present, the elaboration elaborates the `discrete_subtype_definition`, which defines the subtype of the chunk parameter, and the number of values in this subtype determines the maximum number of chunks. After elaborating the `chunk_specification`, a check is made that the determined maximum number of chunks is greater than zero. If this check fails, `Program_Error` is raised.

For the execution of a `loop_statement` that has an `iteration_scheme` including a `loop_parameter_specification`, after elaborating the `chunk_specification` and `aspect_specification`, if any, the `loop_parameter_specification` is elaborated. This elaborates the `discrete_subtype_definition`, which defines the subtype of the loop parameter. If the `discrete_subtype_definition` defines a subtype with a null range, the execution of the `loop_statement` is complete. Otherwise, the `sequence_of_statements` is conditionally executed once for each value of the discrete subtype defined by the `discrete_subtype_definition` that satisfies the predicates of the subtype (or until the loop is left as a consequence of a transfer of control). Prior to each such iteration, the corresponding value of the discrete subtype is assigned to the loop parameter associated with the given iteration. If the loop is a parallel loop, each chunk has its own logical thread of control with its own copy of the loop parameter; otherwise (a *sequential loop*), a single logical thread of control performs the loop, and there is a single copy of the loop parameter. Each logical thread of control handles a distinct subrange of the values of the subtype of the loop parameter such that all values are covered with no overlaps. Within each logical thread of control, the values are assigned to the loop parameter in increasing order unless the reserved word **reverse** is present, in which case the values are assigned in decreasing order. In the absence of a transfer of control, the associated parallel construct of a `loop_parameter_specification` is complete when all of its logical threads of control are complete.

If a `chunk_specification` with a `discrete_subtype_definition` is present, then the logical thread of control associated with a given chunk has its own copy of the chunk parameter initialized with a distinct value from the discrete subtype defined by the `discrete_subtype_definition`. The values of the chunk parameters are assigned such that they increase with increasing values of the ranges covered by the corresponding loop parameters.

Whether or not a `chunk_specification` is present in a parallel loop, the total number of iterations of the loop represents an upper bound on the number of logical threads of control devoted to the loop.

For details about the execution of a `loop_statement` with the `iteration_scheme` including an `iterator_specification`, see 8.5.2. For details relating to a `procedural_iterator`, see 8.5.3.

NOTE 1 A loop parameter declared by a `loop_parameter_specification` is a constant; it cannot be updated within the `sequence_of_statements` of the loop (see 6.3).

NOTE 2 No separate `object_declaration` is expected for a loop parameter, since the loop parameter is automatically declared by the `loop_parameter_specification`. The scope of a loop parameter extends from the `loop_parameter_specification` to the end of the `loop_statement`, and the visibility rules are such that a loop parameter is only visible within the `sequence_of_statements` of the loop.

NOTE 3 The `discrete_subtype_definition` of a `for` loop is elaborated just once. Use of the reserved word **reverse** does not alter the discrete subtype defined, so that the following `iteration_schemes` are not equivalent; the first has a null range.

```
for J in reverse 1 .. 0
for J in 0 .. 1
```

EXAMPLE 1 *Example of a loop statement without an iteration scheme:*

```
loop
  Get(Current_Character);
  exit when Current_Character = '*';
end loop;
```

EXAMPLE 2 *Example of a loop statement with a **while** iteration scheme:*

```
while Bid(N).Price < Cut_Off.Price loop
  Record_Bid(Bid(N).Price);
  N := N + 1;
end loop;
```

EXAMPLE 3 *Example of a loop statement with a **for** iteration scheme:*

```
for J in Buffer'Range loop      -- works even with a null range
  if Buffer(J) /= Space then
    Put(Buffer(J));
  end if;
end loop;
```

EXAMPLE 4 *Example of a loop statement with a name:*

```
Summation:
  while Next /= Head loop      -- see 6.10.1
    Sum := Sum + Next.Value;
    Next := Next.Succ;
  end loop Summation;
```

EXAMPLE 5 *Example of a simple parallel loop:*

```
-- see 6.6
parallel
for I in Grid'Range(1) loop
  Grid(I, 1) := (for all J in Grid'Range(2) => Grid(I,J) = True);
end loop;
```

EXAMPLE 6 *Example of a parallel loop with a chunk specification:*

```
declare
  subtype Chunk_Number is Natural range 1 .. 8;

  Partial_Sum,
  Partial_Max : array (Chunk_Number) of Natural := (others => 0);
  Partial_Min : array (Chunk_Number) of Natural :=
    (others => Natural'Last);

begin
  parallel (Chunk in Chunk_Number)
  for I in Grid'Range(1) loop
    declare
      True_Count : constant Natural :=
        [for J in Grid'Range(2) =>
          (if Grid(I, J) then 1 else 0)]'Reduce("+", 0);

      begin
        Partial_Sum (Chunk) := @ + True_Count;
        Partial_Min (Chunk) := Natural'Min(@, True_Count);
        Partial_Max (Chunk) := Natural'Max(@, True_Count);
      end;
    end loop;

  Put_Line
    ("Total=" & Partial_Sum'Reduce("+", 0)'Image &
     ", Min=" & Partial_Min'Reduce(Natural'Min, Natural'Last)'Image &
     ", Max=" & Partial_Max'Reduce(Natural'Max, 0)'Image);

end;
```

EXAMPLE 7 *For an example of an iterator\_filter, see 7.5.8.*



### 8.5.1 User-Defined Iterator Types

#### Static Semantics

The following language-defined generic library package exists:

```

generic
  type Cursor;
  with function Has_Element (Position : Cursor) return Boolean;
package Ada.Iterator_Interfaces
  with Pure, Nonblocking => False is

    type Forward_Iterator is limited interface;
    function First (Object : Forward_Iterator) return Cursor is abstract;
    function Next (Object : Forward_Iterator; Position : Cursor)
      return Cursor is abstract;

    type Reversible_Iterator is limited interface and Forward_Iterator;
    function Last (Object : Reversible_Iterator) return Cursor is abstract;
    function Previous (Object : Reversible_Iterator; Position : Cursor)
      return Cursor is abstract;

    type Parallel_Iterator is limited interface and Forward_Iterator;
    subtype Chunk_Index is Positive;
    function Is_Split (Object : Parallel_Iterator)
      return Boolean is abstract;

    procedure Split_Into_Chunks (Object      : in out Parallel_Iterator;
                               Max_Chunks : in      Chunk_Index) is abstract
      with Pre'Class => not Object.Is_Split or else raise Program_Error,
          Post'Class => Object.Is_Split and then
            Object.Chunk_Count <= Max_Chunks;

    function Chunk_Count (Object : Parallel_Iterator)
      return Chunk_Index is abstract
      with Pre'Class => Object.Is_Split or else raise Program_Error;

    function First (Object : Parallel_Iterator;
                   Chunk   : Chunk_Index) return Cursor is abstract
      with Pre'Class => (Object.Is_Split and then
                        Chunk <= Object.Chunk_Count)
                        or else raise Program_Error;

    function Next (Object   : Parallel_Iterator;
                  Position  : Cursor;
                  Chunk     : Chunk_Index) return Cursor is abstract
      with Pre'Class => (Object.Is_Split and then
                        Chunk <= Object.Chunk_Count)
                        or else raise Program_Error;

    type Parallel_Reversible_Iterator is limited interface
      and Parallel_Iterator and Reversible_Iterator;
  end Ada.Iterator_Interfaces;

```

An *iterator type* is a type descended from the Forward\_Iterator interface from some instance of Ada.Iterator\_Interfaces. A *reversible iterator type* is a type descended from the Reversible\_Iterator interface from some instance of Ada.Iterator\_Interfaces. A *parallel iterator type* is a type descended from the Parallel\_Iterator interface from some instance of Ada.Iterator\_Interfaces. A type descended from the Parallel\_Reversible\_Iterator interface from some instance of Ada.Iterator\_Interfaces is both a parallel iterator type and a reversible iterator type. An *iterator object* is an object of an iterator type. A *reversible iterator object* is an object of a reversible iterator type. A *parallel iterator object* is an object of a parallel iterator type. The formal subtype Cursor from the associated instance of Ada.Iterator\_Interfaces is the *iteration cursor subtype* for the iterator type.

The following type-related operational aspects may be specified for an indexable container type *T* (see 7.1.6):

**Default\_Iterator**

This aspect is specified by a *name* that denotes exactly one function declared immediately within the same declaration list in which *T*, or the declaration completed by *T*, is declared, whose first parameter is of type *T* or *T*'Class or an access parameter whose designated type is type *T* or *T*'Class, whose other parameters, if any, have default expressions, and whose result type is an iterator type. This function is the *default iterator function* for *T*. Its result subtype is the *default iterator subtype* for *T*. The iteration cursor subtype for the default iterator subtype is the *default cursor subtype* for *T*. This aspect is inherited by descendants of type *T* (including *T*'Class).

**Iterator\_Element**

This aspect is specified by a *name* that denotes a subtype. This is the *default element subtype* for *T*. This aspect is inherited by descendants of type *T* (including *T*'Class).

**Iterator\_View**

This aspect is specified by a *name* that denotes a type *T2* with the following properties:

- *T2* is declared in the same compilation unit as *T*;
- *T2* is an iterable container type;
- *T2* has a single discriminant which is an access discriminant designating *T*; and
- The default iterator subtypes for *T* and *T2* statically match.

This aspect is never inherited, even by *T*'Class.

An *iterable container type* is an indexable container type with specified *Default\_Iterator* and *Iterator\_Element* aspects. A *reversible iterable container type* is an iterable container type with the default iterator type being a reversible iterator type. A *parallel iterable container type* is an iterable container type with the default iterator type being a parallel iterator type. An *iterable container object* is an object of an iterable container type. A *reversible iterable container object* is an object of a reversible iterable container type. A *parallel iterable container object* is an object of a parallel iterable container type.

The *Default\_Iterator* and *Iterator\_Element* aspects are nonoverridable (see 16.1.1).

*Legality Rules*

The *Constant\_Indexing* aspect (if any) of an iterable container type *T* shall denote exactly one function with the following properties:

- the result type of the function is covered by the default element type of *T* or is a reference type (see 7.1.5) with an access discriminant designating a type covered by the default element type of *T*;
- the type of the second parameter of the function covers the default cursor type for *T*;
- if there are more than two parameters, the additional parameters all have default expressions.

This function (if any) is the *default constant indexing function* for *T*.

The *Variable\_Indexing* aspect (if any) of an iterable container type *T* shall denote exactly one function with the following properties:

- the result type of the function is a reference type (see 7.1.5) with an access discriminant designating a type covered by the default element type of *T*;
- the type of the second parameter of the function covers the default cursor type for *T*;
- if there are more than two parameters, the additional parameters all have default expressions.

This function (if any) is the *default variable indexing function* for *T*.

*Erroneous Execution*

A call on the *First* or *Next* operation on a given *Parallel\_Iterator* object with a given *Chunk* value, which does not propagate an exception, should return a *Cursor* value that either yields *False* when passed to

Has\_Element, or that identifies an element distinct from any Cursor value returned by a call on a First or Next operation on the same Parallel\_Iterator object with a different Chunk value. If the First or Next operations with a Chunk parameter behave in any other manner, execution is erroneous.

### 8.5.2 Generalized Loop Iteration

Generalized forms of loop iteration are provided by an `iterator_specification`.

#### Syntax

```
iterator_specification ::=
  defining_identifier[:loop_parameter_subtype_indication] in [reverse]
    iterator_name[iterator_filter]
  | defining_identifier[:loop_parameter_subtype_indication] of
    [reverse] iterable_name[iterator_filter]

loop_parameter_subtype_indication ::= subtype_indication|access_definition
```

If an `iterator_specification` is for a parallel construct, the reserved word **reverse** shall not appear in the `iterator_specification`.

#### Name Resolution Rules

For the first form of `iterator_specification`, called a *generalized iterator*, the expected type for the `iterator_name` is any iterator type. For the second form of `iterator_specification`, the expected type for the `iterable_name` is any array or iterable container type. If the `iterable_name` denotes an array object, the `iterator_specification` is called an *array component iterator*; otherwise it is called a *container element iterator*.

#### Legality Rules

If the reserved word **reverse** appears, the `iterator_specification` is a *reverse iterator*. If the `iterator_specification` is for a parallel construct, the `iterator_specification` is a *parallel iterator*. Otherwise, it is a *forward iterator*. Forward and reverse iterators are collectively called *sequential iterators*. In a reverse generalized iterator, the `iterator_name` shall be of a reversible iterator type. In a parallel generalized iterator, the `iterator_name` shall be of a parallel iterator type. In a reverse container element iterator, the default iterator type for the type of the `iterable_name` shall be a reversible iterator type. In a parallel container element iterator, the default iterator type for the type of the `iterable_name` shall be of a parallel iterator type.

The subtype defined by the `loop_parameter_subtype_indication`, if any, of a generalized iterator shall statically match the iteration cursor subtype. The subtype defined by the `loop_parameter_subtype_indication`, if any, of an array component iterator shall statically match the component subtype of the type of the `iterable_name`. The subtype defined by the `loop_parameter_subtype_indication`, if any, of a container element iterator shall statically match the default element subtype for the type of the `iterable_name`.

In a container element iterator whose `iterable_name` has type *T*, if the `iterable_name` denotes a constant or the Variable\_Indexing aspect is not specified for *T*, then the Constant\_Indexing aspect shall be specified for *T*.

The `iterator_name` or `iterable_name` of an `iterator_specification` shall not denote a subcomponent that depends on discriminants of an object whose nominal subtype is unconstrained, unless the object is known to be constrained.

A container element iterator is illegal if the call of the default iterator function that creates the loop iterator (see below) is illegal.

A generalized iterator is illegal if the iteration cursor subtype of the `iterator_name` is a limited type at the point of the generalized iterator. A container element iterator is illegal if the default cursor subtype of the type of the `iterable_name` is a limited type at the point of the container element iterator.

## Static Semantics

An *iterator\_specification* declares a *loop parameter*. In a generalized iterator, an array component iterator, or a container element iterator, if a *loop\_parameter\_subtype\_indication* is present, it determines the nominal subtype of the loop parameter. In a generalized iterator, if a *loop\_parameter\_subtype\_indication* is not present, the nominal subtype of the loop parameter is the iteration cursor subtype. In an array component iterator, if a *loop\_parameter\_subtype\_indication* is not present, the nominal subtype of the loop parameter is the component subtype of the type of the *iterable\_name*. In a container element iterator, if a *loop\_parameter\_subtype\_indication* is not present, the nominal subtype of the loop parameter is the default element subtype for the type of the *iterable\_name*.

In a generalized iterator, the loop parameter is a constant. In an array component iterator, the loop parameter is a constant if the *iterable\_name* denotes a constant; otherwise it denotes a variable. In a container element iterator, the loop parameter is a constant if the *iterable\_name* denotes a constant, or if the *Variable\_Indexing* aspect is not specified for the type of the *iterable\_name*; otherwise it is a variable.

## Dynamic Semantics

For the execution of a *loop\_statement* with an *iterator\_specification*, the *iterator\_specification* is first elaborated. This elaboration elaborates the *subtype\_indication*, if any.

For a sequential generalized iterator, the loop parameter is created, the *iterator\_name* is evaluated, and the denoted iterator object becomes the *loop iterator*. In a forward generalized iterator, the operation *First* of the iterator type is called on the loop iterator, to produce the initial value for the loop parameter. If the result of calling *Has\_Element* on the initial value is *False*, then the execution of the *loop\_statement* is complete. Otherwise, the *sequence\_of\_statements* is conditionally executed and then the *Next* operation of the iterator type is called with the loop iterator and the current value of the loop parameter to produce the next value to be assigned to the loop parameter. This repeats until the result of calling *Has\_Element* on the loop parameter is *False*, or the loop is left as a consequence of a transfer of control. For a reverse generalized iterator, the operations *Last* and *Previous* are called rather than *First* and *Next*.

For a parallel generalized iterator, the *chunk\_specification*, if any, of the associated parallel construct, is first elaborated, to determine the maximum number of chunks (see 8.5), and then the operation *Split\_Into\_Chunks* of the iterator type is called, with the determined maximum passed as the *Max\_Chunks* parameter, specifying the upper bound for the number of loop parameter objects (and the number of logical threads of control) to be associated with the iterator. In the absence of a *chunk\_specification*, the maximum number of chunks is determined in an implementation-defined manner.

Upon return from *Split\_Into\_Chunks*, the actual number of chunks for the loop is determined by calling the *Chunk\_Count* operation of the iterator, at which point one logical thread of control is initiated for each chunk, with an associated chunk index in the range from one to the actual number of chunks.

Within each logical thread of control, a loop parameter is created. If a *chunk\_specification* with a *discrete\_subtype\_definition* is present in the associated parallel construct, then a chunk parameter is created and initialized with a value from the discrete subtype defined by the *discrete\_subtype\_definition*, so that the order of the chosen chunk parameter values correspond to the order of the chunk indices associated with the logical threads of control. The operation *First* of the iterator type that has a *Chunk* parameter is called on the loop iterator, with *Chunk* initialized from the corresponding chunk index, to produce the initial value for the loop parameter. If the result of calling *Has\_Element* on this initial value is *False*, then the execution of the logical thread of control is complete. Otherwise, the *sequence\_of\_statements* is conditionally executed, and then the *Next* operation of the iterator type that has a *Chunk* parameter is called with the loop iterator, the current value of the loop parameter, and the corresponding chunk index, to produce the next value to be assigned to the loop parameter. This repeats until the result of calling *Has\_Element*

on the loop parameter is False, or the associated parallel construct is left as a consequence of a transfer of control.

In the absence of a transfer of control, the associated parallel construct of a parallel generalized iterator is complete when all of its logical threads of control are complete.

For an array component iterator, the `chunk_specification` of the associated parallel construct, if any, is first elaborated to determine the maximum number of chunks (see 8.5), and then the `iterable_name` is evaluated and the denoted array object becomes the *array for the loop*. If the array for the loop is a null array, then the execution of the `loop_statement` is complete. Otherwise, the `sequence_of_statements` is conditionally executed with the loop parameter denoting each component of the array for the loop, using a *canonical* order of components, which is last dimension varying fastest (unless the array has convention Fortran, in which case it is first dimension varying fastest). For a forward array component iterator, the iteration starts with the component whose index values are each the first in their index range, and continues in the canonical order. For a reverse array component iterator, the iteration starts with the component whose index values are each the last in their index range, and continues in the reverse of the canonical order. For a parallel array component iterator, the iteration is broken up into contiguous chunks of the canonical order, such that all components are covered with no overlaps; each chunk has its own logical thread of control with its own loop parameter and iteration within each chunk is in the canonical order. The number of chunks is implementation defined, but is limited in the presence of a `chunk_specification` to the determined maximum. The loop iteration proceeds until the `sequence_of_statements` has been conditionally executed for each component of the array for the loop, or until the loop is left as a consequence of a transfer of control.

If a `chunk_specification` with a `discrete_subtype_definition` is present in the associated parallel construct, then the logical thread of control associated with a given chunk has a chunk parameter initialized with a distinct value from the discrete subtype defined by the `discrete_subtype_definition`. The values of the chunk parameters are assigned such that they increase in the canonical order of the starting array components for the chunks.

For a container element iterator, the `chunk_specification` of the associated parallel construct, if any, is first elaborated to determine the maximum number of chunks (see 8.5), and then the `iterable_name` is evaluated. If the container type has `Iterator_View` specified, an object of the `Iterator_View` type is created with the discriminant referencing the iterable container object denoted by the `iterable_name`. This is the *iterable container object for the loop*. Otherwise, the iterable container object denoted by the `iterable_name` becomes the iterable container object for the loop. The default iterator function for the type of the iterable container object for the loop is called on the iterable container object and the result is the *loop iterator*. For a sequential container element iterator, an object of the default cursor subtype is created (the *loop cursor*). For a parallel container element iterator, each chunk of iterations will have its own loop cursor, again of the default cursor subtype.

A container element iterator then proceeds as described above for a generalized iterator, except that each reference to a loop parameter is replaced by a reference to the corresponding loop cursor. For a container element iterator, the loop parameter for each iteration instead denotes an indexing (see 7.1.6) into the iterable container object for the loop, with the only parameter to the indexing being the value of the loop cursor for the given iteration. If the loop parameter is a constant (see above), then the indexing uses the default constant indexing function for the type of the iterable container object for the loop; otherwise it uses the default variable indexing function.

Any exception propagated by the execution of a generalized iterator or container element iterator is propagated by the immediately enclosing loop statement.

EXAMPLE 1 *Example of a parallel generalized loop over an array:*

```
parallel
for Element of Board loop -- See 6.6.1.
  Element := Element * 2.0; -- Double each element of Board, a two-
  dimensional array.
end loop;
```



EXAMPLE 2 For examples of use of generalized iterators, see A.18.33 and the corresponding container packages in A.18.2 and A.18.3.

### 8.5.3 Procedural Iterators

A `procedural_iterator` invokes a user-defined procedure, passing in the body of the enclosing `loop_statement` as a parameter of an anonymous access-to-procedure type, to allow the loop body to be executed repeatedly as part of the invocation of the user-defined procedure.

#### Syntax

```
procedural_iterator ::=
  iterator_parameter_specification of iterator_procedure_call
  [iterator_filter]
iterator_parameter_specification ::=
  formal_part
  | (defining_identifier{, defining_identifier})
iterator_procedure_call ::=
  procedure_name
  | procedure_prefix iterator_actual_parameter_part
iterator_actual_parameter_part ::=
  (iterator_parameter_association{, iterator_parameter_association})
iterator_parameter_association ::=
  parameter_association
  | parameter_association_with_box
parameter_association_with_box ::=
  [formal_parameter_selector_name => ] <>
```

At most one `iterator_parameter_association` within an `iterator_actual_parameter_part` shall be a `parameter_association_with_box`.

#### Name Resolution Rules

The name or prefix given in an `iterator_procedure_call` shall resolve to denote a callable entity *C* (the *iterating procedure*) that is a procedure, or an entry renamed as (viewed as) a procedure. When there is an `iterator_actual_parameter_part`, the prefix can be an implicit\_dereference of an access-to-subprogram value.

An `iterator_procedure_call` without a `parameter_association_with_box` is equivalent to one with an `iterator_actual_parameter_part` with an additional `parameter_association_with_box` at the end, with the *formal\_parameter\_selector\_name* identifying the last formal parameter of the callable entity denoted by the name or prefix.

An `iterator_procedure_call` shall contain at most one `iterator_parameter_association` for each formal parameter of the callable entity *C*. Each formal parameter without an `iterator_parameter_association` shall have a `default_expression` (in the profile of the view of *C* denoted by the name or prefix).

The formal parameter of the callable entity *C* associated with the `parameter_association_with_box` shall be of an anonymous access-to-procedure type *A*.

#### Legality Rules

The anonymous access-to-procedure type *A* shall have at least one formal parameter in its parameter profile. If the `iterator_parameter_specification` is a `formal_part`, then this `formal_part` shall be mode conformant with that of *A*. If the `iterator_parameter_specification` is a list of `defining_identifiers`, the number of formal parameters of *A* shall be the same as the length of this list.



If the name or prefix given in an `iterator_procedure_call` denotes an abstract subprogram, the subprogram shall be a dispatching subprogram.

#### Static Semantics

A `loop_statement` with an `iteration_scheme` that has a `procedural_iterator` is equivalent to a local declaration of a procedure `P` followed by a `procedure_call_statement` that is formed from the `iterator_procedure_call` by replacing the `<>` of the `parameter_association_with_box` with `P'Access`. The `formal_part` of the locally declared procedure `P` is formed from the `formal_part` of the anonymous access-to-procedure type `A`, by replacing the identifier of each formal parameter of this `formal_part` with the identifier of the corresponding formal parameter or element of the list of `defining_identifiers` given in the `iterator_parameter_specification`. The body of `P` consists of the conditionally executed `sequence_of_statements`. The procedure `P` is called the *loop body procedure*.

In a procedural iterator, the `Parallel_Calls` aspect (see 12.10.1) of the loop body procedure is True if the reserved word **parallel** occurs in the corresponding loop statement, and False otherwise.

The following aspects may be specified for a callable entity `S` that has exactly one formal parameter of an anonymous access-to-subprogram type:

#### Allows\_Exit

The `Allows_Exit` aspect is of type Boolean. The specified value shall be static. The `Allows_Exit` aspect of an inherited primitive subprogram is True if `Allows_Exit` is True either for the corresponding subprogram of the progenitor type or for any other inherited subprogram that it overrides. If not specified or inherited as True, the `Allows_Exit` aspect of a callable entity is False. For an entry, only a confirming specification of False is permitted for the `Allows_Exit` aspect.

Specifying the `Allows_Exit` aspect to be True for a subprogram indicates that the subprogram *allows exit*, meaning that it is prepared to be completed by arbitrary transfers of control from the loop body procedure, including propagation of exceptions. A subprogram for which `Allows_Exit` is True should use finalization as appropriate rather than exception handling to recover resources and make any necessary final updates to data structures.

#### Parallel\_Iterator

The `Parallel_Iterator` aspect is of type Boolean. The specified value shall be static. The `Parallel_Iterator` aspect of an inherited primitive subprogram is True if `Parallel_Iterator` is True either for the corresponding subprogram of the progenitor type or for any other inherited subprogram that it overrides. If not specified or inherited as True, the `Parallel_Iterator` aspect of a callable entity is False.

Specifying the `Parallel_Iterator` aspect to be True for a callable entity indicates that the entity is allowed to invoke the loop body procedure from multiple distinct logical threads of control. The `Parallel_Iterator` aspect for a subprogram shall be statically False if the subprogram allows exit.

#### Legality Rules

If a callable entity overrides an inherited dispatching subprogram that allows exit, the overriding callable entity also shall allow exit. If a callable entity overrides an inherited dispatching subprogram that has a True `Parallel_Iterator` aspect, the overriding callable entity also shall have a True `Parallel_Iterator` aspect.

A `loop_statement` with a `procedural_iterator` as its `iteration_scheme` shall begin with the reserved word **parallel** if and only if the callable entity identified in the `iterator_procedure_call` has a `Parallel_iterator` aspect of True.

If the actual parameter of an anonymous access-to-subprogram type, passed in an explicit call of a subprogram for which the `Parallel_Iterator` aspect is `True`, is of the form *P*'Access, the designated subprogram *P* shall have a `Parallel_Calls` aspect `True` (see 12.10.1).

The `sequence_of_statements` of a `loop_statement` with a `procedural_iterator` as its `iteration_scheme` shall contain an `exit_statement`, `return_statement`, `goto_statement`, or `requeue_statement` that leaves the loop only if the callable entity associated with the `procedural_iterator` allows exit.

The `sequence_of_statements` of a `loop_statement` with a `procedural_iterator` as its `iteration_scheme` shall not contain an `accept_statement` whose `entry_declaration` occurs outside the `loop_statement`.

#### Dynamic Semantics

For the execution of a `loop_statement` with an `iteration_scheme` that has a `procedural_iterator`, the procedure denoted by the name or prefix of the `iterator_procedure_call` (the *iterating procedure*) is invoked, passing an access value designating the loop body procedure as a parameter. The iterating procedure then calls the loop body procedure zero or more times and returns, whereupon the `loop_statement` is complete. If the **parallel** reserved word is present, the iterating procedure is allowed to invoke the loop body procedure from multiple distinct logical threads of control. The `aspect_specification`, if any, is elaborated prior to the invocation of the iterating procedure.

#### Bounded (Run-Time) Errors

If the callable entity identified in the `iterator_procedure_call` allows exit, then it is a bounded error for a call of the loop body procedure to be performed from within an abort-deferred operation (see 12.8), unless the entire `loop_statement` was within the same abort-deferred operation. If detected, `Program_Error` is raised at the point of the call; otherwise, a transfer of control from the `sequence_of_statements` of the `loop_statement` will not necessarily terminate the `loop_statement`, and the loop body procedure can be called again.

If a `loop_statement` with the `procedural_iterator` as its `iteration_scheme` (see 8.5) does not begin with the reserved word **parallel**, it is a bounded error if the loop body procedure is invoked from a different logical thread of control than the one that initiates the `loop_statement`. If detected, `Program_Error` is raised; otherwise, conflicts associated with concurrent executions of the loop body procedure can occur without being detected by the applicable conflict check policy (see 12.10.1). Furthermore, propagating an exception or making an attempt to exit in the presence of multiple threads of control will not necessarily terminate the `loop_statement`, deadlock can occur, or the loop body procedure can be called again.

EXAMPLE 1 Example of iterating over a map from *My\_Key\_Type* to *My\_Element\_Type* (see A.18.4):

```

for (C : Cursor) of My_Map.Iterate loop
  Put_Line (My_Key_Type'Image (Key (C)) & " => " &
    My_Element_Type'Image (Element (C)));
end loop;

-- The above is equivalent to:

declare
  procedure P (C : Cursor) is
  begin
    Put_Line (My_Key_Type'Image (Key (c)) & " => " &
      My_Element_Type'Image (Element (C)));
  end P;
begin
  My_Map.Iterate (P'Access);
end;

```

EXAMPLE 2 Example of iterating over the environment variables (see A.17):

```

for (Name, Val) of Ada.Environment_Variables.Iterate(<>) loop
  -- "<>" is optional because it is the last parameter
  Put_Line (Name & " => " & Val);
end loop;

-- The above is equivalent to:

declare
  procedure P (Name : String; Val : String) is
  begin
    Put_Line (Name & " => " & Val);
  end P;
begin
  Ada.Environment_Variables.Iterate (P'Access);
end;

```

## 8.6 Block Statements

A `block_statement` encloses a `handled_sequence_of_statements` optionally preceded by a `declarative_part`.

### Syntax

```

block_statement ::=
  [block_statement_identifier:]
  [declare
    declarative_part]
  begin
    handled_sequence_of_statements
  end [block_identifier];

```

If a `block_statement` has a *block\_statement\_identifier*, then the identifier shall be repeated after the **end**; otherwise, there shall not be an identifier after the **end**.

### Static Semantics

A `block_statement` that has no explicit `declarative_part` has an implicit empty `declarative_part`.

### Dynamic Semantics

The execution of a `block_statement` consists of the elaboration of its `declarative_part` followed by the execution of its `handled_sequence_of_statements`.

EXAMPLE *Example of a block statement with a local variable:*

```

Swap:
  declare
    Temp : Integer;
  begin
    Temp := V; V := U; U := Temp;
  end Swap;

```

### 8.6.1 Parallel Block Statements

A `parallel_block_statement` comprises two or more `sequence_of_statements` separated by **and** where each represents an independent activity that is intended to proceed concurrently with the others.

### Syntax

```

parallel_block_statement ::=
  parallel [(chunk_specification)] [aspect_specification] do
    sequence_of_statements
  and
    sequence_of_statements

```

```

{and
  sequence_of_statements}
end do;

```

The `chunk_specification`, if any, of a `parallel_block_statement` shall be an *integer\_simple\_expression*.

#### Dynamic Semantics

For the execution of a `parallel_block_statement`, the `chunk_specification` and the `aspect_specification`, if any, are elaborated in an arbitrary order. After elaborating the `chunk_specification`, if any, a check is made that the determined maximum number of chunks is greater than zero. If this check fails, `Program_Error` is raised.

Then, the various `sequence_of_statements` are grouped into one or more *chunks*, each with its own logical thread of control (see Clause 12), up to the maximum number of chunks specified by the `chunk_specification`, if any. Within each chunk every `sequence_of_statements` of the chunk is executed in turn, in an arbitrary order. The `parallel_block_statement` is complete once every one of the `sequence_of_statements` has completed, either by reaching the end of its execution, or due to a transfer of control out of the construct by one of the `sequence_of_statements` (see 8.1).

EXAMPLE 1 *Example of a parallel block used to walk a binary tree in parallel:*

```

procedure Traverse (T : Expr_Ptr) is -- see 6.9.1
begin
  if T /= null and then
    T.all in Binary_Operation'Class -- see 6.9.1
  then -- recurse down the binary tree
    parallel do
      Traverse (T.Left);
    and
      Traverse (T.Right);
    and
      Ada.Text_IO.Put_Line
        ("Processing " & Ada.Tags.Expanded_Name (T.Tag));
    end do;
  end if;
end Traverse;

```

EXAMPLE 2 *Example of a parallel block used to search two halves of a string in parallel:*

```

function Search (S : String; Char : Character) return Boolean is
begin
  if S'Length <= 1000 then
    -- Sequential scan
    return (for some C of S => C = Char);
  else
    -- Parallel divide and conquer
    declare
      Mid : constant Positive := S'First + S'Length/2 - 1;
    begin
      parallel do
        for C of S(S'First .. Mid) loop
          if C = Char then
            return True; -- Terminates enclosing do
          end if;
        end loop;
      and
        for C of S(Mid + 1 .. S'Last) loop
          if C = Char then
            return True; -- Terminates enclosing do
          end if;
        end loop;
      end do;
      -- Not found
      return False;
    end;
  end if;
end Search;

```

## 8.7 Exit Statements

An `exit_statement` is used to complete the execution of an enclosing `loop_statement`; the completion is conditional if the `exit_statement` includes a condition.

*Syntax*

```

exit_statement ::=
  exit [loop_name] [when condition];

```

*Name Resolution Rules*

The `loop_name`, if any, in an `exit_statement` shall resolve to denote a `loop_statement`.

*Legality Rules*

Each `exit_statement` *applies to* a `loop_statement`; this is the `loop_statement` being exited. An `exit_statement` with a name is only allowed within the `loop_statement` denoted by the name, and applies to that `loop_statement`. An `exit_statement` without a name is only allowed within a `loop_statement`, and applies to the innermost enclosing one. An `exit_statement` that applies to a given `loop_statement` shall not appear within a body or `accept_statement`, if this construct is itself enclosed by the given `loop_statement`.

*Dynamic Semantics*

For the execution of an `exit_statement`, the condition, if present, is first evaluated. If the value of the condition is True, or if there is no condition, a transfer of control is done to complete the `loop_statement`. If the value of the condition is False, no transfer of control takes place.

NOTE Several nested loops can be exited by an `exit_statement` that names the outer loop.

EXAMPLE *Examples of loops with exit statements:*

```

for N in 1 .. Max_Num_Items loop
  Get_New_Item(New_Item);
  Merge_Item(New_Item, Storage_File);
  exit when New_Item = Terminal_Item;
end loop;

Main_Cycle:
loop
  -- initial statements
  exit Main_Cycle when Found;
  -- final statements
end loop Main_Cycle;

```

## 8.8 Goto Statements

A `goto_statement` specifies an explicit transfer of control from this statement to a target statement with a given label.

### Syntax

```
goto_statement ::= goto label_name;
```

### Name Resolution Rules

The *label\_name* shall resolve to denote a label; the statement with that label is the *target statement*.

### Legality Rules

The innermost *sequence\_of\_statements* that encloses the target statement shall also enclose the `goto_statement`. Furthermore, if a `goto_statement` is enclosed by an `accept_statement` or a body, then the target statement shall not be outside this enclosing construct.

### Dynamic Semantics

The execution of a `goto_statement` transfers control to the target statement, completing the execution of any compound statement that encloses the `goto_statement` but does not enclose the target.

NOTE The above rules allow transfer of control to a statement of an enclosing *sequence\_of\_statements* but not the reverse. Similarly, they prohibit transfers of control such as between alternatives of a *case\_statement*, *if\_statement*, or *select\_statement*; between exception handlers; or from an exception handler of a handled *sequence\_of\_statements* back to its *sequence\_of\_statements*.

EXAMPLE *Example of a loop containing a goto statement:*

```

<<Sort>>
for I in 1 .. N-1 loop
  if A(I) > A(I+1) then
    Exchange(A(I), A(I+1));
    goto Sort;
  end if;
end loop;

```

## 9 Subprograms

A subprogram is a program unit or intrinsic operation whose execution is invoked by a subprogram call. There are two forms of subprogram: procedures and functions. A procedure call is a statement; a function call is an expression and returns a value. The definition of a subprogram can be given in two parts: a subprogram declaration defining its interface, and a *subprogram\_body* defining its execution. Operators and enumeration literals are functions.



A *callable entity* is a subprogram or entry (see Section 9). A callable entity is invoked by a *call*; that is, a subprogram call or entry call. A *callable construct* is a construct that defines the action of a call upon a callable entity: a `subprogram_body`, `entry_body`, or `accept_statement`.

## 9.1 Subprogram Declarations

A `subprogram_declaration` declares a procedure or function.

### Syntax

```
subprogram_declaration ::=
    [overriding_indicator]
    subprogram_specification
    [aspect_specification];
subprogram_specification ::=
    procedure_specification
    | function_specification
procedure_specification ::=
    procedure defining_program_unit_name parameter_profile
function_specification ::=
    function defining_designator parameter_and_result_profile
designator ::= [parent_unit_name.] identifier | operator_symbol
defining_designator ::=
    defining_program_unit_name | defining_operator_symbol
defining_program_unit_name ::= [parent_unit_name.] defining_identifier
```

The optional `parent_unit_name` is only allowed for library units (see 13.1.1).

```
operator_symbol ::= string_literal
```

The sequence of characters in an `operator_symbol` shall form a reserved word, a delimiter, or compound delimiter that corresponds to an operator belonging to one of the six categories of operators defined in 7.5.

```
defining_operator_symbol ::= operator_symbol
parameter_profile ::= [formal_part]
parameter_and_result_profile ::=
    [formal_part] return [null_exclusion] subtype_mark
    | [formal_part] return access_definition
formal_part ::=
    (parameter_specification { parameter_specification })
parameter_specification ::=
    defining_identifier_list : [aliased] mode [null_exclusion] subtype_mark
    | := default_expression
    [aspect_specification]
    | defining_identifier_list : access_definition
    | := default_expression
    [aspect_specification]
mode ::= [in] | in out | out
```

### Name Resolution Rules

A *formal parameter* is an object directly visible within a `subprogram_body` that represents the actual parameter passed to the subprogram in a call; it is declared by a `parameter_specification`. For a formal parameter, the expected type for its `default_expression`, if any, is that of the formal parameter.

*Legality Rules*

The *parameter mode* of a formal parameter conveys the direction of information transfer with the actual parameter: **in**, **in out**, or **out**. Mode **in** is the default, and is the mode of a parameter defined by an `access_definition`.

A `default_expression` is only allowed in a `parameter_specification` for a formal parameter of mode **in**.

A `subprogram_declaration` or a `generic_subprogram_declaration` requires a completion unless the `Import` aspect (see B.1) is `True` for the declaration; the completion shall be a body or a `renaming_declaration` (see 11.5). A completion is not allowed for an `abstract_subprogram_declaration` (see 6.9.3), a `null_procedure_declaration` (see 9.7), or an `expression_function_declaration` (see 9.8).

A name that denotes a formal parameter is not allowed within the `formal_part` in which it is declared, nor within the `formal_part` of a corresponding body or `accept_statement`.

*Static Semantics*

The *profile* of (a view of) a callable entity is either a `parameter_profile` or `parameter_and_result_profile`; it embodies information about the interface to that entity — for example, the profile includes information about parameters passed to the callable entity. All callable entities have a profile — enumeration literals, other subprograms, and entries. An access-to-subprogram type has a designated profile. Associated with a profile is a calling convention. A `subprogram_declaration` declares a procedure or a function, as indicated by the initial reserved word, with name and profile as given by its specification.

The nominal subtype of a formal parameter is the subtype determined by the optional `null_exclusion` and the `subtype_mark`, or defined by the `access_definition`, in the `parameter_specification`. The nominal subtype of a function result is the subtype determined by the optional `null_exclusion` and the `subtype_mark`, or defined by the `access_definition`, in the `parameter_and_result_profile`.

An *explicitly aliased parameter* is a formal parameter whose `parameter_specification` includes the reserved word **aliased**.

An *access parameter* is a formal **in** parameter specified by an `access_definition`. An *access result type* is a function result type specified by an `access_definition`. An access parameter or result type is of an anonymous access type (see 6.10). Access parameters of an access-to-object type allow dispatching calls to be controlled by access values. Access parameters of an access-to-subprogram type permit calls to subprograms passed as parameters irrespective of their accessibility level.

The *subtypes of a profile* are:

- For any non-access parameters, the nominal subtype of the parameter.
- For any access parameters of an access-to-object type, the designated subtype of the parameter type.
- For any access parameters of an access-to-subprogram type, the subtypes of the designated profile of the parameter type.
- For any non-access result, the nominal subtype of the function result.
- For any access result type of an access-to-object type, the designated subtype of the result type.
- For any access result type of an access-to-subprogram type, the subtypes of the designated profile of the result type.

The *types of a profile* are the types of those subtypes.

A subprogram declared by an `abstract_subprogram_declaration` is abstract; a subprogram declared by a `subprogram_declaration` is not. See 6.9.3. Similarly, a procedure declared by a

`null_procedure_declaration` is a null procedure; a procedure declared by a `subprogram_declaration` is not. See 9.7. Finally, a function declared by an `expression_function_declaration` is an expression function; a function declared by a `subprogram_declaration` is not. See 9.8.

An `overriding_indicator` is used to indicate whether overriding is intended. See 11.3.1.

#### Dynamic Semantics

The elaboration of a `subprogram_declaration` has no effect.

NOTE 1 A `parameter_specification` with several identifiers is equivalent to a sequence of single `parameter_specifications`, as explained in 6.3.

NOTE 2 Abstract subprograms do not have bodies, and cannot be used in a nondispatching call (see 6.9.3).

NOTE 3 The evaluation of `default_expressions` is caused by certain calls, as described in 9.4.1. They are not evaluated during the elaboration of the subprogram declaration.

NOTE 4 Subprograms can be called recursively and can be called concurrently from multiple tasks.

EXAMPLE 1 *Examples of subprogram declarations:*

```

procedure Traverse_Tree;
procedure Increment(X : in out Integer);
procedure Right_Indent(Margin : out Line_Size); -- see 6.5.4
procedure Switch(From, To : in out Link); -- see 6.10.1

function Random return Probability; -- see 6.5.7

function Min_Cell(X : Link) return Cell; -- see 6.10.1
function Next_Frame(K : Positive) return Frame; -- see 6.10
function Dot_Product(Left, Right : Vector) return Real; -- see 6.6
function Find(B : aliased in out Barrel; Key : String) return Real; -- see 7.1.5

function "*" (Left, Right : Matrix) return Matrix; -- see 6.6

```

EXAMPLE 2 *Examples of `in` parameters with default expressions:*

```

procedure Print_Header(Pages : in Natural;
    Header : in Line := (1 .. Line'Last => ' '); -- see 6.6
    Center : in Boolean := True);

```

### 9.1.1 Preconditions and Postconditions

For a noninstance subprogram (including a generic formal subprogram), a generic subprogram, an entry, or an access-to-subprogram type, the following language-defined assertion aspects may be specified with an `aspect_specification` (see 16.1.1):

- |            |                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Pre        | This aspect specifies a specific precondition for a callable entity or an access-to-subprogram type; it shall be specified by an <code>expression</code> , called a <i>specific precondition expression</i> . If not specified for an entity, the specific precondition expression for the entity is the enumeration literal <code>True</code> .                                                                                  |
| Pre'Class  | This aspect specifies a class-wide precondition for a dispatching operation of a tagged type and its descendants; it shall be specified by an <code>expression</code> , called a <i>class-wide precondition expression</i> . If not specified for an entity, then if no other class-wide precondition applies to the entity, the class-wide precondition expression for the entity is the enumeration literal <code>True</code> . |
| Post       | This aspect specifies a specific postcondition for a callable entity or an access-to-subprogram type; it shall be specified by an <code>expression</code> , called a <i>specific postcondition expression</i> . If not specified for an entity, the specific postcondition expression for the entity is the enumeration literal <code>True</code> .                                                                               |
| Post'Class | This aspect specifies a class-wide postcondition for a dispatching operation of a tagged type and its descendants; it shall be specified by an <code>expression</code> , called a <i>class-wide</i>                                                                                                                                                                                                                               |

*postcondition expression*. If not specified for an entity, the class-wide postcondition expression for the entity is the enumeration literal `True`.

#### *Name Resolution Rules*

The expected type for a precondition or postcondition expression is any boolean type.

Within the expression for a Pre'Class or Post'Class aspect for a primitive subprogram *S* of a tagged type *T*, a `name` that denotes a formal parameter (or *S*'Result) of type *T* is interpreted as though it had a (notional) nonabstract type *NT* that is a formal derived type whose ancestor type is *T*, with directly visible primitive operations. Similarly, a `name` that denotes a formal access parameter (or *S*'Result for an access result) of type *access-to-T* is interpreted as having type *access-to-NT*. The result of this interpretation is that the only operations that can be applied to such `names` are those defined for such a formal derived type.

For an `attribute_reference` with `attribute_designator` *Old*, if the attribute reference has an expected type (or class of types) or shall resolve to a given type, the same applies to the `prefix`; otherwise, the `prefix` shall be resolved independently of context.

#### *Legality Rules*

The Pre or Post aspect shall not be specified for an abstract subprogram or a null procedure. Only the Pre'Class and Post'Class aspects may be specified for such a subprogram.

If a type *T* has an implicitly declared subprogram *P* inherited from a parent type *T1* and a homograph (see 11.3) of *P* from a progenitor type *T2*, and

- the corresponding primitive subprogram *P1* of type *T1* is neither null nor abstract; and
- the class-wide precondition expression `True` does not apply to *P1* (implicitly or explicitly); and
- there is a class-wide precondition expression that applies to the corresponding primitive subprogram *P2* of *T2* that does not fully conform to any class-wide precondition expression that applies to *P1*,

then:

- If the type *T* is abstract, the implicitly declared subprogram *P* is *abstract*.
- Otherwise, the subprogram *P* *requires overriding* and shall be overridden with a nonabstract subprogram.

If a renaming of a subprogram or entry *S1* overrides an inherited subprogram *S2*, then the overriding is illegal unless each class-wide precondition expression that applies to *S1* fully conforms to some class-wide precondition expression that applies to *S2* and each class-wide precondition expression that applies to *S2* fully conforms to some class-wide precondition expression that applies to *S1*.

Pre'Class shall not be specified for an overriding primitive subprogram of a tagged type *T* unless the Pre'Class aspect is specified for the corresponding primitive subprogram of some ancestor of *T*.

In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

#### *Static Semantics*

If a Pre'Class or Post'Class aspect is specified for a primitive subprogram *S* of a tagged type *T*, or such an aspect defaults to `True`, then a corresponding expression also applies to the corresponding primitive subprogram *S* of each descendant of *T* (including *T* itself). The *corresponding expression* is constructed from the associated expression as follows:

- References to formal parameters of *S* (or to *S* itself) are replaced with references to the corresponding formal parameters of the corresponding inherited or overriding subprogram *S* (or to the corresponding subprogram *S* itself).

If the primitive subprogram *S* is not abstract, but the given descendant of *T* is abstract, then a nondispatching call on *S* is illegal if any Pre'Class or Post'Class aspect that applies to *S* is other than a

static boolean expression. Similarly, a primitive subprogram of an abstract type *T*, to which a non-static Pre'Class or Post'Class aspect applies, shall neither be the prefix of an Access attribute\_reference, nor shall it be a generic actual subprogram for a formal subprogram declared by a formal\_concrete\_subprogram\_declaration.

If performing checks is required by the Pre, Pre'Class, Post, or Post'Class assertion policies (see 14.4.2) in effect at the point of a corresponding aspect specification applicable to a given subprogram, entry, or access-to-subprogram type, then the respective precondition or postcondition expressions are considered *enabled*.

A subexpression of a postcondition expression is *known on entry* if it is any of:

- a static subexpression (see 7.9);
- a literal whose type does not have any Integer\_Literal, Real\_Literal, or String\_Literal aspect specified, or the function specified by such an attribute has aspect Global specified to be **null**;
- a name statically denoting a full constant declaration which is known to have no variable views (see 6.3);
- a name statically denoting a nonaliased **in** parameter of an elementary type;
- an Old attribute\_reference;
- an invocation of a predefined operator where all of the operands are known on entry;
- a function call where the function has aspect Global => **null** where all of the actual parameters are known on entry;
- a selected\_component of a known-on-entry prefix;
- an indexed\_component of a known-on-entry prefix where all index expressions are known on entry;
- a parenthesized known-on-entry expression;
- a qualified\_expression or type\_conversion whose operand is a known-on-entry expression;
- a conditional\_expression where all of the conditions, *selecting*\_expressions, and *dependent*\_expressions are known on entry.

A subexpression of a postcondition expression is *unconditionally evaluated*, *conditionally evaluated*, or *repeatedly evaluated*. A subexpression is considered unconditionally evaluated unless it is conditionally evaluated or repeatedly evaluated.

The following subexpressions are repeatedly evaluated:

- A subexpression of a predicate of a quantified\_expression;
- A subexpression of the expression of an array\_component\_association;
- A subexpression of the expression of a container\_element\_association.

For a subexpression that is conditionally evaluated, there is a set of *determining expressions* that determine whether the subexpression is actually evaluated at run time. Subexpressions that are conditionally evaluated and their determining expressions are as follows:

- For an if\_expression that is not repeatedly evaluated, a subexpression of any part other than the first condition is conditionally evaluated, and its determining expressions include all conditions of the if\_expression that precede the subexpression textually;
- For a case\_expression that is not repeatedly evaluated, a subexpression of any *dependent\_expression* is conditionally evaluated, and its determining expressions include the *selecting\_expression* of the case\_expression;

- For a short-circuit control form that is not repeatedly evaluated, a subexpression of the right-hand operand is conditionally evaluated, and its determining expressions include the left-hand operand of the short-circuit control form;
- For a membership test that is not repeatedly evaluated, a subexpression of a `membership_choice` other than the first is conditionally evaluated, and its determining expressions include the *tested\_simple\_expression* and the preceding *membership\_choices* of the membership test.

A conditionally evaluated subexpression is *determined to be unevaluated* at run time if its set of determining expressions are all known on entry, and when evaluated on entry their values are such that the given subexpression is not evaluated.

For a `prefix` *X* that denotes an object of a nonlimited type, the following attribute is defined:

**X'Old** Each X'Old in a postcondition expression that is enabled, other than those that occur in subexpressions that are determined to be unevaluated, denotes a constant that is implicitly declared at the beginning of the subprogram body, entry body, or accept statement.

The implicitly declared entity denoted by each occurrence of X'Old is declared as follows:

- If *X* is of an anonymous access type defined by an `access_definition` *A* then

*X'Old* : **constant** *A* := *X*;

- If *X* is of a specific tagged type *T* then

*anonymous* : **constant** *T*'Class := *T*'Class(*X*);  
*X'Old* : *T* **renames** *T*(*anonymous*);

where the name *X'Old* denotes the object renaming.

- Otherwise

*X'Old* : **constant** *S* := *X*;

where *S* is the nominal subtype of *X*. This includes the case where the type of *S* is an anonymous array type or a universal type.

The type and nominal subtype of X'Old are as implied by the above definitions.

Reference to this attribute is only allowed within a postcondition expression. The `prefix` of an `Old` `attribute_reference` shall not contain a `Result` `attribute_reference`, nor an `Old` `attribute_reference`, nor a use of an entity declared within the postcondition expression but not within `prefix` itself (for example, the loop parameter of an enclosing `quantified_expression`). The `prefix` of an `Old` `attribute_reference` shall statically name (see 7.9) an entity, unless the `attribute_reference` is unconditionally evaluated, or is conditionally evaluated where all of the determining expressions are known on entry.

For a `prefix` *F* that denotes a function declaration or an access-to-function type, the following attribute is defined:

**F'Result** Within a postcondition expression for *F*, denotes the return object of the function call for which the postcondition expression is evaluated. The type of this attribute is that of the result subtype of the function or access-to-function type except within a `Post'Class` postcondition expression for a function with a controlling result or with a controlling access result; in those cases the type of the attribute is described above as part of the Name Resolution Rules for `Post'Class`.

Use of this attribute is allowed only within a postcondition expression for *F*.

For a `prefix` *E* that denotes an entry declaration of an entry family (see 12.5.2), the following attribute is defined:



**E'Index** Within a precondition or postcondition expression for entry family E, denotes the value of the entry index for the call of E. The nominal subtype of this attribute is the entry index subtype.

Use of this attribute is allowed only within a precondition or postcondition expression for E.

*Dynamic Semantics*

Upon a call of the subprogram or entry, after evaluating any actual parameters, precondition checks are performed as follows:

- The specific precondition check begins with the evaluation of the specific precondition expression that applies to the subprogram or entry, if it is enabled; if the expression evaluates to **False**, **Assertions.Assertion\_Error** is raised; if the expression is not enabled, the check succeeds.
- The class-wide precondition check begins with the evaluation of any enabled class-wide precondition expressions that apply to the subprogram or entry. If and only if all the class-wide precondition expressions evaluate to **False**, **Assertions.Assertion\_Error** is raised.

The precondition checks are performed in an arbitrary order, and if any of the class-wide precondition expressions evaluate to **True**, it is not specified whether the other class-wide precondition expressions are evaluated. The precondition checks and any check for elaboration of the subprogram body are performed in an arbitrary order. In a call on a protected operation, the checks are performed before starting the protected action. For an entry call, the checks are performed prior to checking whether the entry is open.

Upon successful return from a call of the subprogram or entry, prior to copying back any by-copy **in out** or **out** parameters, the postcondition check is performed. This consists of the evaluation of any enabled specific and class-wide postcondition expressions that apply to the subprogram or entry. If any of the postcondition expressions evaluate to **False**, then **Assertions.Assertion\_Error** is raised. The postcondition expressions are evaluated in an arbitrary order, and if any postcondition expression evaluates to **False**, it is not specified whether any other postcondition expressions are evaluated. The postcondition check, and any constraint or predicate checks associated with **in out** or **out** parameters are performed in an arbitrary order.

For a call to a task entry, the postcondition check is performed before the end of the rendezvous; for a call to a protected operation, the postcondition check is performed before the end of the protected action of the call. The postcondition check for any call is performed before the finalization of any implicitly-declared constants associated (as described above) with **Old attribute\_references** but after the finalization of any other entities whose accessibility level is that of the execution of the callable construct.

If a precondition or postcondition check fails, the exception is raised at the point of the call; the exception cannot be handled inside the called subprogram or entry. Similarly, any exception raised by the evaluation of a precondition or postcondition expression is raised at the point of call.

For any call to a subprogram or entry *S* (including dispatching calls), the checks that are performed to verify specific precondition expressions and specific and class-wide postcondition expressions are determined by those for the subprogram or entry actually invoked. Note that the class-wide postcondition expressions verified by the postcondition check that is part of a call on a primitive subprogram of type *T* includes all class-wide postcondition expressions originating in any progenitor of *T*, even if the primitive subprogram called is inherited from a type *T1* and some of the postcondition expressions do not apply to the corresponding primitive subprogram of *T1*. Any operations within a class-wide postcondition expression that were resolved as primitive operations of the (notional) formal derived type *NT*, are in the evaluation of the postcondition bound to the corresponding operations of the type identified by the controlling tag of the call on *S*. This applies to both dispatching and non-dispatching calls on *S*.

The class-wide precondition check for a call to a subprogram or entry *S* consists solely of checking the class-wide precondition expressions that apply to the denoted callable entity (not necessarily to the one that is invoked). Any operations within such an expression that were resolved as primitive

operations of the (notional) formal derived type *NT* are in the evaluation of the precondition bound to the corresponding operations of the type identified by the controlling tag of the call on *S*. This applies to both dispatching and non-dispatching calls on *S*.

For the purposes of the above rules, a call on an inherited subprogram is considered to involve a call on a subprogram *S'* whose body consists only of a call (with appropriate conversions) on the non-inherited subprogram *S* from which the inherited subprogram was derived. It is not specified whether class-wide precondition or postcondition expressions that are equivalent (with respect to which non-inherited function bodies are executed) for *S* and *S'* are evaluated once or twice. If evaluated only once, the value returned is used for both associated checks.

For a call via an access-to-subprogram value, precondition and postcondition checks performed are as determined by the subprogram or entry denoted by the prefix of the Access attribute reference that produced the value. In addition, a precondition check of any precondition expression associated with the access-to-subprogram type is performed. Similarly, a postcondition check of any postcondition expression associated with the access-to-subprogram type is performed.

For a call on a generic formal subprogram, precondition and postcondition checks performed are as determined by the subprogram or entry denoted by the actual subprogram, along with any specific precondition and specific postcondition of the formal subprogram itself.

#### Implementation Permissions

An implementation may evaluate a known-on-entry subexpression of a postcondition expression of an entity at the place where X'Old constants are created for the entity, with the normal evaluation of the postcondition expression, or both.

NOTE 1 A precondition is checked just before the call. If another task can change any value that the precondition expression depends on, the precondition can evaluate to False within the subprogram or entry body.

NOTE 2 For an example of the use of these aspects and attributes, see the Streams Subsystem definitions in 16.13.1.

### 9.1.2 The Global and Global'Class Aspects

The Global and Global'Class aspects of a program unit are used to identify the objects global to the unit that can be read or written during its execution.

#### Syntax

```
global_aspect_definition ::=
    null
  | Unspecified
  | global_mode global_designator
  | (global_aspect_element { global_aspect_element })

global_aspect_element ::=
    global_mode global_set
  | global_mode all
  | global_mode synchronized

global_mode ::=
    basic_global_mode
  | extended_global_mode

basic_global_mode ::= in | in out | out

global_set ::= global_name { global_name }

global_designator ::= all | synchronized | global_name

global_name ::= object_name | package_name
```

*Name Resolution Rules*

A `global_name` shall resolve to statically name an object or a package (including a limited view of a package).

*Static Semantics*

For a subprogram, an entry, an access-to-subprogram type, a task unit, a protected unit, or a library package or generic library package, the following language-defined aspect may be specified with an `aspect_specification` (see 16.1.1):

**Global** The Global aspect shall be specified with a `global_aspect_definition`.

The Global aspect identifies the set of variables (which, for the purposes of this clause, includes all constants except those which are known to have no variable views (see 6.3)) that are global to a callable entity or task body, and that are read or updated as part of the execution of the callable entity or task body. If specified for a protected unit, it refers to all of the protected operations of the protected unit. Constants of any type may also be mentioned in a Global aspect.

If not specified or otherwise defined below, the aspect defaults to the Global aspect for the enclosing library unit if the entity is declared at library level, and to Unspecified otherwise. If not specified for a library unit, the aspect defaults to `Global => null` for a library unit that is declared Pure, and to `Global => Unspecified` otherwise.

For a dispatching subprogram, the following language-defined aspect may be specified with an `aspect_specification` (see 16.1.1):

**Global'Class**

The Global'Class aspect shall be specified with a `global_aspect_definition`. This aspect identifies an upper bound on the set of variables global to a dispatching operation that can be read or updated as a result of a dispatching call on the operation. If not specified, the aspect defaults to the Global aspect for the dispatching subprogram.

Together, we refer to the Global and Global'Class aspects as *global* aspects.

A `global_aspect_definition` defines the Global or Global'Class aspect of some entity. The Global aspect identifies the sets of global variables that can be read, written, or modified as a side effect of executing the operation(s) associated with the entity. The Global'Class aspect associated with a dispatching operation of type *T* represents a restriction on the Global aspect on a corresponding operation of any descendant of type *T*.

The Global aspect for a callable entity defines the global variables that can be referenced as part of a call on the entity, including any assertion expressions that apply to the call (even if not enabled), such as preconditions, postconditions, predicates, and type invariants.

The Global aspect for an access-to-subprogram object (or subtype) identifies the global variables that can be referenced when calling via the object (or any object of that subtype) including assertion expressions that apply.

For a predefined operator of an elementary type, the function representing an enumeration literal, or any other static function (see 7.9), the Global aspect is **null**. For a predefined operator of a composite type, the Global aspect of the operator defaults to that of the enclosing library unit (unless a Global aspect is specified for the type — see H.7).

The following is defined in terms of operations that are performed by or on behalf of an entity. The rules on operations apply to the entity(s) associated with those operations.

The global variables associated with any `global_mode` can be read as a side effect of an operation. The **in out** and **out** `global_modes` together identify the set of global variables that can be updated as a side effect of an operation. Creating an access-to-variable value that designates an object is considered an update of the designated object, and creating an access-to-constant value that designates an object is considered a read of the designated object.

The overall set of objects associated with each `global_mode` includes all objects identified for the mode in the `global_aspect_definition`.

A `global_set` identifies a *global variable set* as follows:

- **all** identifies the set of all global variables;
- **synchronized** identifies the set of all synchronized variables (see 12.10), as well as variables of a composite type all of whose non-discriminant subcomponents are synchronized;
- `global_name{, global_name}` identifies the union of the sets of variables identified by the `global_names` in the list, for the following forms of `global_name`:
  - *object\_name* identifies the specified global variable (or constant);
  - *package\_name* identifies the set of all variables declared in the private part or body of the package, or anywhere within a private descendant of the package.

#### Legality Rules

Within a `global_aspect_definition`, a given `global_mode` shall be specified at most once. Similarly, within a `global_aspect_definition`, a given entity shall be named at most once by a `global_name`.

If an entity (other than a library package or generic library package) has a Global aspect other than Unspecified or **in out all**, then the associated operation(s) shall read only those variables global to the entity that are within the global variable set associated with the **in**, **in out**, or **out** `global_modes`, and the operation(s) shall update only those variables global to the entity that are within the global variable set associated with either the **in out** or **out** `global_modes`. In the absence of the No\_Hidden\_Indirect\_Globals restriction (see H.4), this ignores objects reached via a dereference of an access value. The above rule includes any possible Global effects of calls occurring during the execution of the operation, except for the following excluded calls:

- calls to formal subprograms;
- calls associated with operations on formal subtypes;
- calls through formal objects of an access-to-subprogram type;
- calls through access-to-subprogram parameters;
- calls on operations with Global aspect Unspecified.

The possible Global effects of these excluded calls (other than those that are Unspecified) are taken into account by the caller of the original operation, by presuming they occur at least once during its execution. For calls that are not excluded, the possible Global effects of the call are those permitted by the Global aspect of the associated entity, or by its Global'Class aspect if a dispatching call.

If a Global aspect other than Unspecified or **in out all** applies to an access-to-subprogram type, then the prefix of an Access attribute\_reference producing a value of such a type shall denote a subprogram whose Global aspect is not Unspecified and is *covered* by that of the result type, where a global aspect *G1* is *covered* by a global aspect *G2* if the set of variables that *G1* identifies as readable or updatable is a subset of the corresponding set for *G2*. Similarly on a conversion to such a type, the operand shall be of a named access-to-subprogram type whose Global aspect is covered by that of the target type.

If an implementation-defined `global_mode` applies to a given set of variables, an implementation-defined rule determines what sort of references to them are permitted.

For a subprogram that is a dispatching operation of a tagged type *T*, each mode of its Global aspect shall identify a subset of the variables identified by the corresponding mode, or by the **in out** mode, of the Global'Class aspect of a corresponding dispatching subprogram of any ancestor of *T*, unless the aspect of that ancestor is Unspecified.

*Implementation Permissions*

An implementation can allow some references to a constant object which are not accounted for by the Global or Global'Class aspect when it is considered a variable in the above rules, if the implementation can determine that the object is in fact immutable.

Implementations may perform additional checks on calls to operations with an Unspecified Global aspect to ensure that they do not violate any limitations associated with the point of call.

Implementations may extend the syntax or semantics of the Global aspect in an implementation-defined manner; for example, supporting additional `global_modes`.

NOTE For an example of the use of these aspects, see the Vector container definition in A.18.2.

## 9.2 Formal Parameter Modes

A `parameter_specification` declares a formal parameter of mode **in**, **in out**, or **out**.

*Static Semantics*

A parameter is passed either *by copy* or *by reference*. When a parameter is passed by copy, the formal parameter denotes a separate object from the actual parameter, and any information transfer between the two occurs only before and after executing the subprogram. When a parameter is passed by reference, the formal parameter denotes (a view of) the object denoted by the actual parameter; reads and updates of the formal parameter directly reference the actual parameter object.

A type is a *by-copy type* if it is an elementary type, or if it is a descendant of a private type whose full type is a by-copy type. A parameter of a by-copy type is passed by copy, unless the formal parameter is explicitly aliased.

A type is a *by-reference type* if it is a descendant of one of the following:

- a tagged type;
- a task or protected type;
- an explicitly limited record type;
- a composite type with a subcomponent of a by-reference type;
- a private type whose full type is a by-reference type.

A parameter of a by-reference type is passed by reference, as is an explicitly aliased parameter of any type. Each value of a by-reference type has an associated object. For a value conversion, the associated object is the anonymous result object if such an object is created (see 7.6); otherwise it is the associated object of the operand. In other cases, the object associated with the evaluated operative constituent of the `name` or `expression` (see 7.4) determines its associated object.

For other parameters, it is unspecified whether the parameter is passed by copy or by reference.

*Bounded (Run-Time) Errors*

If one `name` denotes a part of a formal parameter, and a second `name` denotes a part of a distinct formal parameter or an object that is not part of a formal parameter, then the two `names` are considered *distinct access paths*. If an object is of a type for which the parameter passing mechanism is not specified and is not an explicitly aliased parameter, then it is a bounded error to assign to the object via one access path, and then read the value of the object via a distinct access path, unless the first access path denotes a part of a formal parameter that no longer exists at the point of the second access (due to leaving the corresponding callable construct). The possible consequences are that `Program_Error` is raised, or the newly assigned value is read, or some old value of the object is read.

NOTE 1 The mode of a formal parameter describes the direction of information transfer to or from the `subprogram_body` (see 9.1).

NOTE 2 A formal parameter of mode **in** is a constant view (see 6.3); it cannot be updated within the `subprogram_body`.

NOTE 3 A formal parameter of mode **out** can be uninitialized at the start of the `subprogram_body` (see 9.4.1).

### 9.3 Subprogram Bodies

A `subprogram_body` specifies the execution of a subprogram.

#### Syntax

```
subprogram_body ::=
  [overriding_indicator]
  subprogram_specification
  [aspect_specification] is
  declarative_part
  begin
    handled_sequence_of_statements
  end [designator];
```

If a designator appears at the end of a `subprogram_body`, it shall repeat the `defining_designator` of the `subprogram_specification`.

#### Legality Rules

In contrast to other bodies, a `subprogram_body` is allowed to be defined without it being the completion of a previous declaration, in which case the body declares the subprogram. If the body is a completion, it shall be the completion of a `subprogram_declaration` or `generic_subprogram_declaration`. The profile of a `subprogram_body` that completes a declaration shall conform fully to that of the declaration.

#### Static Semantics

A `subprogram_body` is considered a declaration. It can either complete a previous declaration, or itself be the initial declaration of the subprogram.

#### Dynamic Semantics

The elaboration of a nongeneric `subprogram_body` has no other effect than to establish that the subprogram can from then on be called without failing the `Elaboration_Check`.

The execution of a `subprogram_body` is invoked by a subprogram call. For this execution the `declarative_part` is elaborated, and the `handled_sequence_of_statements` is then executed.

EXAMPLE 1 *Example of procedure body:*

```
procedure Push(E : in Element_Type; S : in out Stack) is
begin
  if S.Index = S.Size then
    raise Stack_Overflow;
  else
    S.Index := S.Index + 1;
    S.Space(S.Index) := E;
  end if;
end Push;
```

EXAMPLE 2 *Example of a function body:*

```
function Dot_Product(Left, Right : Vector) return Real is
  Sum : Real := 0.0;
begin
  Check(Left'First = Right'First and Left'Last = Right'Last);
  for J in Left'Range loop
    Sum := Sum + Left(J)*Right(J);
  end loop;
  return Sum;
end Dot_Product;
```



### 9.3.1 Conformance Rules

When subprogram profiles are given in more than one place, they are required to conform in one of four ways: type conformance, mode conformance, subtype conformance, or full conformance.

#### Static Semantics

As explained in B.1, a *convention* can be specified for an entity. Unless this document states otherwise, the default convention of an entity is Ada. For a callable entity or access-to-subprogram type, the convention is called the *calling convention*. The following conventions are defined by the language:

- The default calling convention for any subprogram not listed below is *Ada*. The Convention aspect may be specified to override the default calling convention (see B.1).
- The *Intrinsic* calling convention represents subprograms that are “built in” to the compiler. The default calling convention is *Intrinsic* for the following:
  - an enumeration literal;
  - a `"/=` operator declared implicitly due to the declaration of `=` (see 9.6);
  - any other implicitly declared subprogram unless it is a dispatching operation of a tagged type;
  - an inherited subprogram of a generic formal tagged type with unknown discriminants;
  - an attribute that is a subprogram;
  - a subprogram declared immediately within a `protected body`;
  - any prefixed view of a subprogram (see 7.1.3) without synchronization kind (see 12.5) `By_Entry` or `By_Protected_Procedure`.

The Access attribute is not allowed for *Intrinsic* subprograms.

- The default calling convention is *protected* for a `protected` subprogram, for a prefixed view of a subprogram with a synchronization kind of `By_Protected_Procedure`, and for an access-to-subprogram type with the reserved word **protected** in its definition.
- The default calling convention is *entry* for an entry and for a prefixed view of a subprogram with a synchronization kind of `By_Entry`.
- The calling convention for an anonymous access-to-subprogram parameter or anonymous access-to-subprogram result is *protected* if the reserved word **protected** appears in its definition; otherwise, it is the convention of the entity that has the parameter or result, unless that entity has convention *protected*, *entry*, or *Intrinsic*, in which case the convention is Ada.
- If not specified above as *Intrinsic*, the calling convention for any inherited or overriding dispatching operation of a tagged type is that of the corresponding subprogram of the parent type. The default calling convention for a new dispatching operation of a tagged type is the convention of the type.

Of these four conventions, only Ada and *Intrinsic* are allowed as a *convention\_identifier* in the specification of a Convention aspect.

Two profiles are *type conformant* if they have the same number of parameters, and both have a result if either does, and corresponding parameter and result types are the same, or, for access parameters or access results, corresponding designated types are the same, or corresponding designated profiles are type conformant.

Two profiles are *mode conformant* if:

- they are type conformant; and
- corresponding parameters have identical modes and both or neither are explicitly aliased parameters; and
- for corresponding access parameters and any access result type, the designated subtypes statically match and either both or neither are access-to-constant, or the designated profiles are subtype conformant.

Two profiles are *subtype conformant* if they are mode conformant, corresponding subtypes of the profile statically match, and the associated calling conventions are the same. The profile of a generic formal subprogram is not subtype conformant with any other profile.

Two profiles are *fully conformant* if they are subtype conformant, if they have access-to-subprogram results whose designated profiles are fully conformant, and for corresponding parameters:

- they have the same names; and
- both or neither have `null_exclusions`; and
- neither have `default_expressions`, or they both have `default_expressions` that are fully conformant with one another; and
- for access-to-subprogram parameters, the designated profiles are fully conformant.

Two expressions are *fully conformant* if, after replacing each use of an operator with the equivalent `function_call`:

- each constituent construct of one corresponds to an instance of the same syntactic category in the other, except that an expanded name may correspond to a `direct_name` (or `character_literal`) or to a different expanded name in the other; and
- corresponding `defining_identifiers` occurring within the two expressions are the same; and
- each `direct_name`, `character_literal`, and `selector_name` that is not part of the prefix of an expanded name in one denotes the same declaration as the corresponding `direct_name`, `character_literal`, or `selector_name` in the other, or they denote corresponding declarations occurring within the two expressions; and
- each `attribute_designator` in one is the same as the corresponding `attribute_designator` in the other; and
- each `primary` that is a literal in one is a user-defined literal if and only if the corresponding literal in the other is also a user-defined literal. Furthermore, if neither are user-defined literals then they shall have the same values, but they may have differing textual representations; if both are user-defined literals then they shall have the same textual representation.

Two `known_discriminant_parts` are *fully conformant* if they have the same number of discriminants, and discriminants in the same positions have the same names, statically matching subtypes, and `default_expressions` that are fully conformant with one another.

Two `discrete_subtype_definitions` are *fully conformant* if they are both `subtype_indications` or are both ranges, the `subtype_marks` (if any) denote the same subtype, and the corresponding `simple_expressions` of the ranges (if any) fully conform.

The *prefixed view profile* of a subprogram is the profile obtained by omitting the first parameter of that subprogram. There is no prefixed view profile for a parameterless subprogram. For the purposes of defining subtype and mode conformance, the convention of a prefixed view profile is considered to match that of either an entry or a protected operation.

#### Implementation Permissions

An implementation may declare an operator declared in a language-defined library unit to be intrinsic.

NOTE Any conformance requirements between `aspect_specifications` that are part of a profile or `known_discriminant_part` are defined by the semantics of each particular aspect. In particular, there is no general requirement for `aspect_specifications` to match in conforming profiles or discriminant parts.

### 9.3.2 Inline Expansion of Subprograms

Subprograms may be expanded in line at the call site.

#### Static Semantics

For a callable entity or a generic subprogram, the following language-defined representation aspect may be specified:

**Inline** The type of aspect `Inline` is Boolean. When aspect `Inline` is True for a callable entity, inline expansion is desired for all calls to that entity. When aspect `Inline` is True for a generic subprogram, inline expansion is desired for all calls to all instances of that generic subprogram.

If directly specified, the `aspect_definition` shall be a static expression. This aspect is never inherited; if not directly specified, the aspect is False.

#### Implementation Permissions

For each call, an implementation is free to follow or to ignore the recommendation determined by the `Inline` aspect.

## 9.4 Subprogram Calls

A *subprogram call* is either a `procedure_call_statement` or a `function_call`; it invokes the execution of the `subprogram_body`. The call specifies the association of the actual parameters, if any, with formal parameters of the subprogram.

#### Syntax

```

procedure_call_statement ::=
    procedure_name;
    | procedure_prefix actual_parameter_part;
function_call ::=
    function_name
    | function_prefix actual_parameter_part
actual_parameter_part ::=
    (parameter_association {, parameter_association})
parameter_association ::=
    [formal_parameter_selector_name =>] explicit_actual_parameter
explicit_actual_parameter ::= expression | variable_name

```

A `parameter_association` is *named* or *positional* according to whether or not the `formal_parameter_selector_name` is specified. For the `parameter_associations` of a single `actual_parameter_part` or `iterator_actual_parameter_part`, any positional associations shall precede any named associations. Named associations are not allowed if the prefix in a subprogram call is an `attribute_reference`.

#### Name Resolution Rules

The name or prefix given in a `procedure_call_statement` shall resolve to denote a callable entity that is a procedure, or an entry renamed as (viewed as) a procedure. The name or prefix given in a `function_call` shall resolve to denote a callable entity that is a function. The name or prefix shall not resolve to denote an abstract subprogram unless it is also a dispatching subprogram. When there is an `actual_parameter_part`, the prefix can be an `implicit_dereference` of an access-to-subprogram value.

A subprogram call shall contain at most one association for each formal parameter. Each formal parameter without an association shall have a `default_expression` (in the profile of the view denoted by the name or prefix). This rule is an overloading rule (see 11.6).

*Static Semantics*

If the name or prefix of a subprogram call denotes a prefixed view (see 7.1.3), the subprogram call is equivalent to a call on the underlying subprogram, with the first actual parameter being provided by the prefix of the prefixed view (or the Access attribute of this prefix if the first formal parameter is an access parameter), and the remaining actual parameters given by the `actual_parameter_part`, if any.

*Dynamic Semantics*

For the execution of a subprogram call, the name or prefix of the call is evaluated, and each `parameter_association` is evaluated (see 9.4.1). If a `default_expression` is used, an implicit `parameter_association` is assumed for this rule. These evaluations are done in an arbitrary order. The `subprogram_body` is then executed, or a call on an entry or protected subprogram is performed (see 6.9.2). Finally, if the subprogram completes normally, then after it is left, any necessary assigning back of formal to actual parameters occurs (see 9.4.1).

The exception `Program_Error` is raised at the point of a `function_call` if the function completes normally without executing a return statement.

A `function_call` denotes a constant, as defined in 9.5; the nominal subtype of the constant is given by the nominal subtype of the function result.

**EXAMPLE 1** *Examples of procedure calls:*

```
Traverse_Tree; -- see 9.1
Print_Header(128, Title, True); -- see 9.1
Switch(From => X, To => Next); -- see 9.1
Print_Header(128, Header => Title, Center => True); -- see 9.1
Print_Header(Header => Title, Center => True, Pages => 128); -- see 9.1
```

**EXAMPLE 2** *Examples of function calls:*

```
Dot_Product(U, V) -- see 9.1 and 9.3
Clock -- see 12.6
F.all -- presuming F is of an access-to-subprogram type — see 6.10
```

**EXAMPLE 3** *Examples of procedures with default expressions:*

```
procedure Activate(Process : in Process_Name;
  After : in Process_Name := No_Process;
  Wait : in Duration := 0.0;
  Prior : in Boolean := False);

procedure Pair(Left, Right : in Person_Name := new Person(M)); -- see 6.10.1
```

**EXAMPLE 4** *Examples of their calls:*

```
Activate(X);
Activate(X, After => Y);
Activate(X, Wait => 60.0, Prior => True);
Activate(X, Y, 10.0, False);

Pair;
Pair(Left => new Person(F), Right => new Person(M));
```

**NOTE** If a `default_expression` is used for two or more parameters in a multiple `parameter_specification`, the `default_expression` is evaluated once for each omitted parameter. Hence in the above examples, the two calls of `Pair` are equivalent.

**EXAMPLE 5** *Examples of overloaded subprograms:*

```
procedure Put(X : in Integer);
procedure Put(X : in String);

procedure Set(Tint : in Color);
procedure Set(Signal : in Light);
```

EXAMPLE 6 *Examples of their calls:*

```
Put(28);
Put("no possible ambiguity here");

Set(Tint => Red);
Set(Signal => Red);
Set(Color'(Red));

-- Set(Red) would be ambiguous since Red can
-- denote a value either of type Color or of type Light
```

#### 9.4.1 Parameter Associations

A parameter association defines the association between an actual parameter and a formal parameter.

##### *Name Resolution Rules*

The *formal\_parameter\_selector\_name* of a named *parameter\_association* shall resolve to denote a *parameter\_specification* of the view being called; this is the formal parameter of the association. The formal parameter for a positional *parameter\_association* is the parameter with the corresponding position in the formal part of the view being called.

The *actual parameter* is either the *explicit\_actual\_parameter* given in a *parameter\_association* for a given formal parameter, or the corresponding *default\_expression* if no *parameter\_association* is given for the formal parameter. The expected type for an actual parameter is the type of the corresponding formal parameter.

If the mode is **in**, the actual is interpreted as an expression; otherwise, the actual is interpreted only as a name, if possible.

##### *Legality Rules*

If the mode is **in out** or **out**, the actual shall be a name that denotes a variable.

If the mode is **out**, the actual parameter is a view conversion, and the type of the formal parameter is a scalar type, then

- neither the target type nor the operand type has the *Default\_Value* aspect specified; or
- both the target type and the operand type shall have the *Default\_Value* aspect specified, and there shall exist a type (other than a root numeric type) that is an ancestor of both the target type and the operand type.

In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

If the formal parameter is an explicitly aliased parameter, the type of the actual parameter shall be tagged or the actual parameter shall be an aliased view of an object. Further, if the formal parameter subtype *F* is untagged:

- the subtype *F* shall statically match the nominal subtype of the actual object; or
- the subtype *F* shall be unconstrained, discriminated in its full view, and unconstrained in any partial view.

In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

In a function call, the accessibility level of the actual object for each explicitly aliased parameter shall not be statically deeper than the accessibility level of the master of the call (see 6.10.2).

Two names are *known to denote the same object* if:

- both names statically denote the same stand-alone object or parameter; or
- both names are *selected\_components*, their prefixes are known to denote the same object, and their *selector\_names* denote the same component; or

- both `names` are dereferences (implicit or explicit) and the dereferenced `names` are known to denote the same object; or
- both `names` are `indexed_components`, their `prefixes` are known to denote the same object, and each of the pairs of corresponding index values are either both static expressions with the same static value or both `names` that are known to denote the same object; or
- both `names` are `slices`, their `prefixes` are known to denote the same object, and the two `slices` have statically matching index constraints; or
- one of the two `names` statically denotes a renaming declaration whose renamed `object_name` is known to denote the same object as the other, the `prefix` of any dereference within the renamed `object_name` is not a variable, and any expression within the renamed `object_name` contains no references to variables nor calls on nonstatic functions.

Two names are *known to refer to the same object* if

- The two `names` are known to denote the same object; or
- One of the `names` is a `selected_component`, `indexed_component`, or `slice` and its `prefix` is known to refer to the same object as the other `name`; or
- One of the two `names` statically denotes a renaming declaration whose renamed `object_name` is known to refer to the same object as the other `name`.

If a call *C* has two or more parameters of mode **in out** or **out** that are of an elementary type, then the call is legal only if:

- For each `name N` denoting an object of an elementary type that is passed as a parameter of mode **in out** or **out** to the call *C*, there is no other `name` among the other parameters of mode **in out** or **out** to *C* that is known to denote the same object.

If a construct *C* has two or more direct constituents that are `names` or `expressions` whose evaluation may occur in an arbitrary order, at least one of which contains a function call with an **in out** or **out** parameter, then the construct is legal only if:

- For each `name N` that is passed as a parameter of mode **in out** or **out** to some inner function call *C2* (not including the construct *C* itself), there is no other `name` anywhere within a direct constituent of the construct *C* other than the one containing *C2*, that is known to refer to the same object.

For the purposes of checking this rule:

- For an array aggregate, an expression associated with a `discrete_choice_list` that has two or more discrete choices, or that has a nonstatic range, is considered as two or more separate occurrences of the expression;
- For a record aggregate:
  - The expression of a `record_component_association` is considered to occur once for each associated component; and
  - The `default_expression` for each `record_component_association` with `<>` for which the associated component has a `default_expression` is considered part of the aggregate;
- For a call, any `default_expression` evaluated as part of the call is considered part of the call.

#### Dynamic Semantics

For the evaluation of a `parameter_association`:

- The actual parameter is first evaluated.
- For an access parameter, the `access_definition` is elaborated, which creates the anonymous access type.



- For a parameter (of any mode) that is passed by reference (see 9.2), a view conversion of the actual parameter to the nominal subtype of the formal parameter is evaluated, and the formal parameter denotes that conversion.
- For an **in** or **in out** parameter that is passed by copy (see 9.2), the formal parameter object is created, and the value of the actual parameter is converted to the nominal subtype of the formal parameter and assigned to the formal.
- For an **out** parameter that is passed by copy, the formal parameter object is created, and:
  - For an access type, the formal parameter is initialized from the value of the actual, without checking whether the value satisfies any constraints, predicates, or null exclusions, but including any dynamic accessibility checks associated with a conversion to the type of the formal parameter.
  - For a scalar type that has the `Default_Value` aspect specified, the formal parameter is initialized from the value of the actual, without checking that the value satisfies any constraint or any predicate.
  - For a composite type with discriminants or that has implicit initial values for any subcomponents (see 6.3.1), the behavior is as for an **in out** parameter passed by copy, except that no predicate check is performed.
  - For any other type, the formal parameter is uninitialized. If composite, a view conversion of the actual parameter to the nominal subtype of the formal is evaluated (which can raise `Constraint_Error`), and the actual subtype of the formal is that of the view conversion. If elementary, the actual subtype of the formal is given by its nominal subtype.
  - Furthermore, if the type is a scalar type, and the actual parameter is a view conversion, then `Program_Error` is raised if either the target or the operand type has the `Default_Value` aspect specified, unless they both have the `Default_Value` aspect specified, and there is a type (other than a root numeric type) that is an ancestor of both the target type and the operand type.
- In a function call, for each explicitly aliased parameter, a check is made that the accessibility level of the master of the actual object is not deeper than that of the master of the call (see 6.10.2).

A formal parameter of mode **in out** or **out** with discriminants is constrained if either its nominal subtype or the actual parameter is constrained.

After normal completion and leaving of a subprogram, for each **in out** or **out** parameter that is passed by copy, the value of the formal parameter is converted to the subtype of the variable given as the actual parameter and assigned to it. These conversions and assignments occur in an arbitrary order.

#### *Erroneous Execution*

If the nominal subtype of a formal parameter with discriminants is constrained or indefinite, and the parameter is passed by reference, then the execution of the call is erroneous if the value of any discriminant of the actual is changed while the formal parameter exists (that is, before leaving the corresponding callable construct).

#### *Implementation Permissions*

If the actual parameter in a `parameter_association` with mode **out** is a view conversion between two access types that do not share a common ancestor type, the implementation may pass in the null value of the type of the formal parameter instead of the value of the actual parameter. It is implementation-defined under what circumstances the implementation passes in the null value.

## 9.5 Return Statements

A `simple_return_statement` or `extended_return_statement` (collectively called a *return statement*) is used to complete the execution of the innermost enclosing `subprogram_body`, `entry_body`, or `accept_statement`.

*Syntax*

```

simple_return_statement ::= return [expression];
extended_return_object_declaration ::=
  defining_identifier: [aliased] [constant] return_subtype_indication
    [= expression]
    [aspect_specification]
extended_return_statement ::=
  return extended_return_object_declaration [do
    handled_sequence_of_statements
  end return];
return_subtype_indication ::= subtype_indication | access_definition

```

*Name Resolution Rules*

The *result subtype* of a function is the subtype denoted by the `subtype_mark`, or defined by the `access_definition`, after the reserved word **return** in the profile of the function. The expected type for the expression, if any, of a `simple_return_statement` is the result type of the corresponding function. The expected type for the expression of an `extended_return_object_declaration` is that of the `return_subtype_indication`.

*Legality Rules*

A return statement shall be within a callable construct, and it *applies to* the innermost callable construct or `extended_return_statement` that contains it. A return statement shall not be within a body that is within the construct to which the return statement applies.

A function body shall contain at least one return statement that applies to the function body, unless the function contains `code_statements`. A `simple_return_statement` shall include an expression if and only if it applies to a function body. An `extended_return_statement` shall apply to a function body. An `extended_return_object_declaration` with the reserved word **constant** shall include an expression.

The *expression of an extended\_return\_statement* is the expression (if any) of the `extended_return_object_declaration` of the `extended_return_statement`.

For an `extended_return_statement` that applies to a function body:

- If the result subtype of the function is defined by a `subtype_mark`, the `return_subtype_indication` shall be a `subtype_indication`. The type of the `subtype_indication` shall be covered by the result type of the function. The subtype defined by the `subtype_indication` shall be statically compatible with the result subtype of the function; if the result type of the function is elementary, the two subtypes shall statically match. If the result subtype of the function is indefinite, then the subtype defined by the `subtype_indication` shall be a definite subtype, or there shall be an expression.
- If the result subtype of the function is defined by an `access_definition`, the `return_subtype_indication` shall be an `access_definition`. The subtype defined by the `access_definition` shall statically match the result subtype of the function. The accessibility level of this anonymous access subtype is that of the result subtype.
- If the result subtype of the function is class-wide, the accessibility level of the type of the subtype defined by the `return_subtype_indication` shall not be statically deeper than that of the master that elaborated the function body.

For any return statement that applies to a function body:

- If the result subtype of the function is limited, then the expression of the return statement (if any) shall meet the restrictions described in 10.5.

- If the result subtype of the function is class-wide, the accessibility level of the type of the expression (if any) of the return statement shall not be statically deeper than that of the master that elaborated the function body.
- If the subtype determined by the expression of the `simple_return_statement` or by the `return_subtype_indication` has one or more access discriminants, the accessibility level of the anonymous access type of each access discriminant shall not be statically deeper than that of the master that elaborated the function body.

If the reserved word **aliased** is present in an `extended_return_object_declaration`, the type of the extended return object shall be immutably limited.

#### Static Semantics

Within an `extended_return_statement`, the *return object* is declared with the given `defining_identifier`, with the nominal subtype defined by the `return_subtype_indication`. An `extended_return_statement` with the reserved word **constant** is a full constant declaration that declares the return object to be a constant object.

#### Dynamic Semantics

For the execution of an `extended_return_statement`, the `subtype_indication` or `access_definition` is elaborated. This creates the nominal subtype of the return object. If there is an expression, it is evaluated and converted to the nominal subtype (which can raise `Constraint_Error` — see 7.6); the return object is created and the converted value is assigned to the return object. Otherwise, the return object is created and initialized by default as for a stand-alone object of its nominal subtype (see 6.3.1). If the nominal subtype is indefinite, the return object is constrained by its initial value. A check is made that the value of the return object belongs to the function result subtype. `Constraint_Error` is raised if this check fails.

For the execution of a `simple_return_statement`, the expression (if any) is first evaluated, converted to the result subtype, and then is assigned to the anonymous *return object*.

If the return object has any parts that are tasks, the activation of those tasks does not occur until after the function returns (see 12.2).

If the result type of a function is a specific tagged type, the tag of the return object is that of the result type. If the result type is class-wide, the tag of the return object is that of the value of the expression of the return statement, unless the return object is defined by an `extended_return_object_declaration` with a `subtype_indication` that is specific, in which case it is that of the type of the `subtype_indication`. A check is made that the master of the type identified by the tag of the result includes the elaboration of the master that elaborated the function body. If this check fails, `Program_Error` is raised.

For the execution of an `extended_return_statement`, the `handled_sequence_of_statements` is executed. Within this `handled_sequence_of_statements`, the execution of a `simple_return_statement` that applies to the `extended_return_statement` causes a transfer of control that completes the `extended_return_statement`. Upon completion of a return statement that applies to a callable construct by the normal completion of a `simple_return_statement` or by reaching the **end return** of an `extended_return_statement`, a transfer of control is performed which completes the execution of the callable construct, and returns to the caller.

If the result subtype of the function is defined by an `access_definition` designating a specific tagged type *T*, a check is made that the result value is null or the tag of the object designated by the result value identifies *T*. `Constraint_Error` is raised if this check fails.

If any part of the specific type of the return object of a function (or coextension thereof) has one or more access discriminants whose value is not constrained by the result subtype of the function, a check is made that the accessibility level of the anonymous access type of each access discriminant, as determined by the expression or the `return_subtype_indication` of the return statement, is

not deeper than the level of the master of the call (see 6.10.2). If this check fails, `Program_Error` is raised.

A check is performed that the return value satisfies the predicates of the return subtype. If this check fails, the effect is as defined in 6.2.4.

In the case of a function, the `function_call` denotes a constant view of the return object.

#### Implementation Permissions

For a function call used to initialize a composite object with a constrained nominal subtype or used to initialize a return object that is built in place into such an object:

- If the result subtype of the function is constrained, and conversion of an object of this subtype to the subtype of the object being initialized would raise `Constraint_Error`, then `Constraint_Error` may be raised before calling the function.
- If the result subtype of the function is unconstrained, and a return statement is executed such that the return object is known to be constrained, and conversion of the return object to the subtype of the object being initialized would raise `Constraint_Error`, then `Constraint_Error` may be raised at the point of the call (after abandoning the execution of the function body).

EXAMPLE Examples of return statements:

```

return;                                -- in a procedure body, entry_body,
                                         -- accept_statement, or
extended_return_statement
return Key_Value (Last_Index);          -- in a function body
return Node : Cell do                  -- in a function body, see 6.10.1 for Cell
    Node.Value := Result;
    Node.Succ := Next_Node;
end return;

```

### 9.5.1 Nonreturning Subprograms

Specifying aspect `No_Return` to have the value `True` indicates that a subprogram cannot return normally; it may, for example, propagate an exception or loop forever.

#### Static Semantics

For a subprogram or generic subprogram, the following language-defined representation aspect may be specified:

**No\_Return** The type of aspect `No_Return` is Boolean. When aspect `No_Return` is `True` for an entity, the entity is said to be *nonreturning*.

If directly specified, the `aspect_definition` shall be a static expression. When not directly specified, if the subprogram is a primitive subprogram inherited by a derived type, then the aspect is `True` if any corresponding subprogram of the parent or progenitor types is nonreturning. Otherwise, the aspect is `False`.

If a generic subprogram is nonreturning, then so are its instances. If a subprogram declared within a generic unit is nonreturning, then so are the corresponding copies of that subprogram in instances.

#### Legality Rules

Aspect `No_Return` shall not be specified for a null procedure nor an instance of a generic unit.

A return statement shall not apply to a nonreturning procedure or generic procedure.

Any return statement that applies to a nonreturning function or generic function shall be a `simple_return_statement` with an expression that is a `raise_expression`, a call on a nonreturning function, or a parenthesized expression of one of these.

A subprogram shall be nonreturning if it overrides a dispatching nonreturning subprogram. In addition to the places where Legality Rules normally apply (see 15.3), this rule applies also in the private part of an instance of a generic unit.

If a renaming-as-body completes a nonreturning subprogram declaration, then the renamed subprogram shall be nonreturning.

#### Dynamic Semantics

If the body of a nonreturning procedure completes normally, Program\_Error is raised at the point of the call.

EXAMPLE *Example of a specification of a No\_Return aspect:*

```
procedure Fail(Msg : String)  -- raises Fatal_Error exception
with No_Return;
-- Inform compiler and reader that procedure never returns normally
```

## 9.6 Overloading of Operators

An *operator* is a function whose designator is an operator\_symbol. Operators, like other functions, may be overloaded.

#### Name Resolution Rules

Each use of a unary or binary operator is equivalent to a function\_call with function\_prefix being the corresponding operator\_symbol, and with (respectively) one or two positional actual parameters being the operand(s) of the operator (in order).

#### Legality Rules

The subprogram\_specification of a unary or binary operator shall have one or two parameters, respectively. The parameters shall be of mode **in**. A generic function instantiation whose designator is an operator\_symbol is only allowed if the specification of the generic function has the corresponding number of parameters, and they are all of mode **in**.

Default\_expressions are not allowed for the parameters of an operator (whether the operator is declared with an explicit subprogram\_specification or by a generic\_instantiation).

An explicit declaration of **"/=** shall not have a result type of the predefined type Boolean.

#### Static Semantics

An explicit declaration of **"/=** whose result type is Boolean implicitly declares an operator **"/=** that gives the complementary result.

NOTE The operators **++** and **--** are both unary and binary operators, and hence can be overloaded with both one- and two-parameter functions.

EXAMPLE *Examples of user-defined operators:*

```
function "+" (Left, Right : Matrix) return Matrix;
function "+" (Left, Right : Vector) return Vector;

-- assuming that A, B, and C are of the type Vector
-- the following two statements are equivalent:

A := B + C;
A := "+" (B, C);
```

## 9.7 Null Procedures

A `null_procedure_declaration` provides a shorthand to declare a procedure with an empty body.

### Syntax

```
null_procedure_declaration ::=
  [overriding_indicator]
  procedure_specification is null
  [aspect_specification];
```

### Legality Rules

If a `null_procedure_declaration` is a completion, it shall be the completion of a `subprogram_declaration` or `generic_subprogram_declaration`. The profile of a `null_procedure_declaration` that completes a declaration shall conform fully to that of the declaration.

### Static Semantics

A `null_procedure_declaration` that is not a completion declares a *null procedure*. A completion is not allowed for a `null_procedure_declaration`; however, a `null_procedure_declaration` can complete a previous declaration.

### Dynamic Semantics

The execution of a null procedure is invoked by a subprogram call. For the execution of a subprogram call on a null procedure, or on a procedure completed with a `null_procedure_declaration`, the execution of the `subprogram_body` has no effect.

The elaboration of a `null_procedure_declaration` has no other effect than to establish that the null procedure can be called without failing the `Elaboration_Check`.

EXAMPLE *Example of the declaration of a null procedure:*

```
procedure Simplify(Expr : in out Expression) is null; -- see 6.9
-- By default, Simplify does nothing, but it can be overridden in extensions
of Expression
```

## 9.8 Expression Functions

An `expression_function_declaration` provides a shorthand to declare a function whose body consists of a single return statement.

### Syntax

```
expression_function_declaration ::=
  [overriding_indicator]
  function_specification is
    (expression)
    [aspect_specification];
| [overriding_indicator]
  function_specification is
    aggregate
    [aspect_specification];
```

### Name Resolution Rules

The expected type for the expression or aggregate of an `expression_function_declaration` is the result type (see 9.5) of the function.



## Static Semantics

An `expression_function_declaration` that is not a completion declares an *expression function*. The *return expression of an expression function* is the expression or aggregate of the `expression_function_declaration`. A completion is not allowed for an `expression_function_declaration`; however, an `expression_function_declaration` can complete a previous declaration.

A *potentially static expression* is defined in the same way as a static expression except that

- a name denoting a formal parameter of an expression function is a potentially static expression; and
- each use of “static expression” in the definition of “static expression” is replaced with a corresponding use of “potentially static expression” in the definition of “potentially static expression”.

The following language-defined representation aspect may be specified for an expression function:

**Static** The type of aspect Static is Boolean. When aspect Static is True for an expression function, the function is a *static expression function*. If directly specified, the `aspect_definition` shall be a static expression.

The Static value for an inherited function is True if some corresponding primitive function of the parent or progenitor type is a static expression function; otherwise, if not directly specified, the aspect is False.

A static expression function is a static function; see 7.9.

## Legality Rules

If an `expression_function_declaration` is a completion, it shall be the completion of a `subprogram_declaration` or `generic_subprogram_declaration`. The profile of an `expression_function_declaration` that completes a declaration shall conform fully to that of the declaration.

If the result subtype has one or more unconstrained access discriminants, the accessibility level of the anonymous access type of each access discriminant, as determined by the expression or aggregate of the `expression_function_declaration`, shall not be statically deeper than that of the master that elaborated the `expression_function_declaration`.

Aspect Static shall be specified to have the value True only if the associated `expression_function_declaration`:

- is not a completion;
- has an expression that is a potentially static expression;
- contains no calls to itself;
- each parameter (if any) is of mode **in** and is of a static subtype;
- has a result subtype that is a static subtype;
- has no applicable precondition or postcondition expression; and
- for result type *R*, if the function is a boundary entity for type *R* (see 10.3.2), no type invariant applies to type *R*; if *R* has a component type *C*, a similar rule applies to *C*.

## Dynamic Semantics

The execution of an expression function is invoked by a subprogram call. For the execution of a subprogram call on an expression function, or on a function completed with a `expression_function_declaration`, the execution of the `subprogram_body` executes an implicit function body containing only a `simple_return_statement` whose expression is the return expression of the expression function.

The elaboration of an `expression_function_declaration` has no other effect than to establish that the expression function can be called without failing the `Elaboration_Check`.

EXAMPLE *Example of an expression function:*

```
function Is-Origin (P : in Point) return Boolean is -- see 6.9
(P.X = 0.0 and P.Y = 0.0);
```

## 10 Packages

Packages are program units that allow the specification of groups of logically related entities. Typically, a package contains the declaration of a type (often a private type or private extension) along with the declarations of primitive subprograms of the type, which can be called from outside the package, while their inner workings remain hidden from outside users.

### 10.1 Package Specifications and Declarations

A package is generally provided in two parts: a `package_specification` and a `package_body`. Every package has a `package_specification`, but not all packages have a `package_body`.

#### Syntax

```
package_declaration ::= package_specification;
package_specification ::=
  package defining_program_unit_name
  [aspect_specification] is
  {basic_declarative_item}
  [private
   {basic_declarative_item}]
  end [[parent_unit_name.]identifier]
```

If an identifier or `parent_unit_name.identifier` appears at the end of a `package_specification`, then this sequence of lexical elements shall repeat the `defining_program_unit_name`.

#### Legality Rules

A `package_declaration` or `generic_package_declaration` requires a completion (a body) if it contains any `basic_declarative_item` that requires a completion, but whose completion is not in its `package_specification`.

#### Static Semantics

The first list of `basic_declarative_items` of a `package_specification` of a package other than a generic formal package is called the *visible part* of the package. The optional list of `basic_declarative_items` after the reserved word **private** (of any `package_specification`) is called the *private part* of the package. If the reserved word **private** does not appear, the package has an implicit empty private part. Each list of `basic_declarative_items` of a `package_specification` forms a *declaration list* of the package.

An entity declared in the private part of a package is visible only within the declarative region of the package itself (including any child units — see 13.1.1). In contrast, expanded names denoting entities declared in the visible part can be used even outside the package; furthermore, direct visibility of such entities can be achieved by means of `use_clauses` (see 7.1.3 and 11.4).

#### Dynamic Semantics

The elaboration of a `package_declaration` consists of the elaboration of its `basic_declarative_items` in the given order.

NOTE 1 The visible part of a package contains all the information that another program unit is able to know about the package.

NOTE 2 If a declaration occurs immediately within the specification of a package, and the declaration has a corresponding completion that is a body, then that body has to occur immediately within the body of the package.

EXAMPLE *Example of a package declaration:*

```

package Rational_Numbers is
  type Rational is
    record
      Numerator    : Integer;
      Denominator  : Positive;
    end record;
  function "=" (X,Y : Rational) return Boolean;
  function "/"  (X,Y : Integer) return Rational;  -- to construct a
rational number
  function "+"  (X,Y : Rational) return Rational;
  function "-"  (X,Y : Rational) return Rational;
  function "*"  (X,Y : Rational) return Rational;
  function "/"  (X,Y : Rational) return Rational;
end Rational_Numbers;

```

There are also many examples of package declarations in the predefined language environment (see Annex A).

## 10.2 Package Bodies

In contrast to the entities declared in the visible part of a package, the entities declared in the `package_body` are visible only within the `package_body` itself. As a consequence, a package with a `package_body` can be used for the construction of a group of related subprograms in which the logical operations available to clients are clearly isolated from the internal entities.

*Syntax*

```

package_body ::=
  package body defining_program_unit_name
  [aspect_specification] is
  declarative_part
  [begin
    handled_sequence_of_statements]
  end [[parent_unit_name.]identifier];

```

If an identifier or `parent_unit_name.identifier` appears at the end of a `package_body`, then this sequence of lexical elements shall repeat the `defining_program_unit_name`.

*Legality Rules*

A `package_body` shall be the completion of a previous `package_declaration` or `generic_package_declaration`. A library `package_declaration` or library `generic_package_declaration` shall not have a body unless it requires a body; the `Elaborate_Body` aspect can be used to require a library `unit_declaration` to have a body (see 13.2.1) if it would not otherwise require one.

*Static Semantics*

In any `package_body` without statements there is an implicit `null_statement`. For any `package_declaration` without an explicit completion, there is an implicit `package_body` containing a single `null_statement`. For a noninstance, nonlibrary package, this body occurs at the end of the `declarative_part` of the innermost enclosing program unit or `block_statement`; if there are several such packages, the order of the implicit `package_bodies` is unspecified. (For an instance, the implicit `package_body` occurs at the place of the instantiation (see 15.3). For a library package, the place is partially determined by the elaboration dependences (see Clause 13).)

## Dynamic Semantics

For the elaboration of a nongeneric `package_body`, its `declarative_part` is first elaborated, and its `handled_sequence_of_statements` is then executed.

NOTE 1 A variable declared in the body of a package is only visible within this body and, consequently, its value can only be changed within the `package_body`. In the absence of local tasks, the value of such a variable remains unchanged between calls issued from outside the package to subprograms declared in the visible part. The properties of such a variable are similar to those of a “static” variable of C.

NOTE 2 The elaboration of the body of a subprogram explicitly declared in the visible part of a package is caused by the elaboration of the body of the package. Hence a call of such a subprogram by an outside program unit raises the exception `Program_Error` if the call takes place before the elaboration of the `package_body` (see 6.11).

EXAMPLE *Example of a package body (see 10.1):*

```
package body Rational_Numbers is
  procedure Same_Denominator (X,Y : in out Rational) is
  begin
    -- reduces X and Y to the same denominator:
    ...
  end Same_Denominator;

  function "=" (X,Y : Rational) return Boolean is
    U : Rational := X;
    V : Rational := Y;
  begin
    Same_Denominator (U,V);
    return U.Numerator = V.Numerator;
  end "=";

  function "/" (X,Y : Integer) return Rational is
  begin
    if Y > 0 then
      return (Numerator => X, Denominator => Y);
    else
      return (Numerator => -X, Denominator => -Y);
    end if;
  end "/";

  function "+" (X,Y : Rational) return Rational is ... end "+";
  function "-" (X,Y : Rational) return Rational is ... end "-";
  function "*" (X,Y : Rational) return Rational is ... end "*";
  function "/" (X,Y : Rational) return Rational is ... end "/";

end Rational_Numbers;
```

### 10.3 Private Types and Private Extensions

The declaration (in the visible part of a package) of a type as a private type or private extension serves to separate the characteristics that can be used directly by outside program units (that is, the logical properties) from other characteristics whose direct use is confined to the package (the details of the definition of the type itself). See 6.9.1 for an overview of type extensions.

## Syntax

```
private_type_declaration ::=
  type defining_identifier [discriminant_part] is [[abstract] tagged] [limited] private
  [aspect_specification];

private_extension_declaration ::=
  type defining_identifier [discriminant_part] is
  [abstract] [limited | synchronized] new ancestor_subtype_indication
  [and interface_list] with private
  [aspect_specification];
```

A `private_type_declaration` or `private_extension_declaration` declares a *partial view* of the type; such a declaration is allowed only as a `declarative_item` of the visible part of a package, and it requires a completion, which shall be a `full_type_declaration` that occurs as a `declarative_item` of the private part of the package. The view of the type declared by the `full_type_declaration` is called the *full view*. A generic formal private type or a generic formal private extension is also a partial view.

A type shall be completely defined before it is frozen (see 6.11.1 and 16.14). Thus, neither the declaration of a variable of a partial view of a type, nor the creation by an `allocator` of an object of the partial view are allowed before the full declaration of the type. Similarly, before the full declaration, the name of the partial view cannot be used in a `generic_instantiation` or in a representation item.

A private type is limited if its declaration includes the reserved word **limited**; a private extension is limited if its ancestor type is a limited type that is not an interface type, or if the reserved word **limited** or **synchronized** appears in its definition. If the partial view is nonlimited, then the full view shall be nonlimited. If a tagged partial view is limited, then the full view shall be limited. On the other hand, if an untagged partial view is limited, the full view may be limited or nonlimited.

If the partial view is tagged, then the full view shall be tagged. On the other hand, if the partial view is untagged, then the full view may be tagged or untagged. In the case where the partial view is untagged and the full view is tagged, no derivatives of the partial view are allowed within the immediate scope of the partial view; derivatives of the full view are allowed.

If a full type has a partial view that is tagged, then:

- the partial view shall be a synchronized tagged type (see 6.9.4) if and only if the full type is a synchronized tagged type;
- the partial view shall be a descendant of an interface type (see 3.9.4) if and only if the full type is a descendant of the interface type.

The *ancestor subtype* of a `private_extension_declaration` is the subtype defined by the `ancestor_subtype_indication`; the ancestor type shall be a specific tagged type. The full view of a private extension shall be derived (directly or indirectly) from the ancestor type. In addition to the places where Legality Rules normally apply (see 15.3), the requirement that the ancestor be specific applies also in the private part of an instance of a generic unit.

If the reserved word **limited** appears in a `private_extension_declaration`, the ancestor type shall be a limited type. If the reserved word **synchronized** appears in a `private_extension_declaration`, the ancestor type shall be a limited interface.

If the declaration of a partial view includes a `known_discriminant_part`, then the `full_type_declaration` shall have a fully conforming (explicit) `known_discriminant_part` (see 9.3.1). The ancestor subtype may be unconstrained; the parent subtype of the full view is required to be constrained (see 6.7).

If a private extension inherits known discriminants from the ancestor subtype, then the full view shall also inherit its discriminants from the ancestor subtype, and the parent subtype of the full view shall be constrained if and only if the ancestor subtype is constrained.

If the `full_type_declaration` for a private extension includes a `derived_type_definition`, then the reserved word **limited** shall appear in the `full_type_declaration` if and only if it also appears in the `private_extension_declaration`.

If a partial view has unknown discriminants, then the `full_type_declaration` may define a definite or an indefinite subtype, with or without discriminants.

If a partial view has neither known nor unknown discriminants, then the `full_type_declaration` shall define a definite subtype.

If the ancestor subtype of a private extension has constrained discriminants, then the parent subtype of the full view shall impose a statically matching constraint on those discriminants.

#### Static Semantics

A `private_type_declaration` declares a private type and its first subtype. Similarly, a `private_extension_declaration` declares a private extension and its first subtype.

A declaration of a partial view and the corresponding `full_type_declaration` define two views of a single type. The declaration of a partial view together with the visible part define the operations that are available to outside program units; the declaration of the full view together with the private part define other operations whose direct use is possible only within the declarative region of the package itself. Moreover, within the scope of the declaration of the full view, the characteristics (see 6.4) of the type are determined by the full view; in particular, within its scope, the full view determines the classes that include the type, which components, entries, and protected subprograms are visible, what attributes and other predefined operations are allowed, and whether the first subtype is static. See 10.3.1.

For a private extension, the characteristics (including components, but excluding discriminants if there is a new `discriminant_part` specified), predefined operators, and inherited user-defined primitive subprograms are determined by its ancestor type and its progenitor types (if any), in the same way that those of a record extension are determined by those of its parent type and its progenitor types (see 6.4 and 10.3.1).

#### Dynamic Semantics

The elaboration of a `private_type_declaration` creates a partial view of a type. The elaboration of a `private_extension_declaration` elaborates the *ancestor\_subtype\_indication*, and creates a partial view of a type.

NOTE 1 The partial view of a type as declared by a `private_type_declaration` is defined to be a composite view (in 6.2). The full view of the type can be elementary or composite. A private extension is also composite, as is its full view.

NOTE 2 Declaring a private type with an `unknown_discriminant_part` is a way of preventing clients from creating uninitialized objects of the type; they are then forced to initialize each object by calling some operation declared in the visible part of the package.

NOTE 3 The ancestor type specified in a `private_extension_declaration` and the parent type specified in the corresponding declaration of a record extension given in the private part can be different. If the ancestor type is not an interface type, the parent type of the full view can be any descendant of the ancestor type. In this case, for a primitive subprogram that is inherited from the ancestor type and not overridden, the formal parameter names and default expressions (if any) come from the corresponding primitive subprogram of the specified ancestor type, while the body comes from the corresponding primitive subprogram of the parent type of the full view. See 6.9.2.

NOTE 4 If the ancestor type specified in a `private_extension_declaration` is an interface type, the parent type can be any type so long as the full view is a descendant of the ancestor type. The progenitor types specified in a `private_extension_declaration` and the progenitor types specified in the corresponding declaration of a record extension given in the private part are not necessarily the same — it is only necessary that the private extension and the record extension be descended from the same set of interfaces.

EXAMPLE 1 *Examples of private type declarations:*

```
type Key is private;  
type File_Name is limited private;
```

EXAMPLE 2 *Example of a private extension declaration:*

```
type List is new Ada.Finalization.Controlled with private;
```

### 10.3.1 Private Operations

For a type declared in the visible part of a package or generic package, certain operations on the type do not become visible until later in the package — either in the private part or the body. Such *private operations* are available only inside the declarative region of the package or generic package.



The predefined operators that exist for a given type are determined by the classes to which the type belongs. For example, an integer type has a predefined "+" operator. In most cases, the predefined operators of a type are declared immediately after the definition of the type; the exceptions are explained below. Inherited subprograms are also implicitly declared immediately after the definition of the type, except as stated below.

For a composite type, the characteristics (see 10.3) of the type are determined in part by the characteristics of its component types. At the place where the composite type is declared, the only characteristics of component types used are those characteristics visible at that place. If later immediately within the declarative region in which the composite type is declared additional characteristics become visible for a component type, then any corresponding characteristics become visible for the composite type. Any additional predefined operators are implicitly declared at that place. If there is no such place, then additional predefined operators are not declared at all, but they still exist.

The corresponding rule applies to a type defined by a `derived_type_definition`, if there is a place immediately within the declarative region in which the type is declared where additional characteristics of its parent type become visible.

For example, an array type whose component type is limited private becomes nonlimited if the full view of the component type is nonlimited and visible at some later place immediately within the declarative region in which the array type is declared. In such a case, the predefined "=" operator is implicitly declared at that place, and assignment is allowed after that place.

The characteristics and constraints of the designated subtype of an access type follow a somewhat different rule. The view of the designated subtype of (a view of) an access type at a given place is determined by the view of the designated subtype that is visible at that place, rather than the view at the place where the access type is declared.

A type is a *descendant* of the full view of some ancestor of its parent type only if the current view it has of its parent is a descendant of the full view of that ancestor. More generally, at any given place, a type is descended from the same view of an ancestor as that from which the current view of its parent is descended. This view determines what characteristics are inherited from the ancestor, and, for example, whether the type is considered to be a descendant of a record type, or a descendant only through record extensions of a more distant ancestor.

Furthermore, it is possible for there to be places where a derived type is known to be derived indirectly from an ancestor type, but is not a descendant of even a partial view of the ancestor type, because the parent of this derived type is not visibly a descendant of the ancestor. In this case, the derived type inherits no characteristics from that ancestor, but nevertheless is within the derivation class of the ancestor for the purposes of type conversion, the "covers" relationship, and matching against a formal derived type. In this case the derived type is effectively a *descendant* of an incomplete view of the ancestor.

Inherited primitive subprograms follow a different rule. For a `derived_type_definition`, each inherited primitive subprogram is implicitly declared at the earliest place, if any, immediately within the declarative region in which the `type_declaration` occurs, but after the `type_declaration`, where the corresponding declaration from the parent is visible. If there is no such place, then the inherited subprogram is not declared at all, but it still exists. For a tagged type, it is possible to dispatch to an inherited subprogram that is not declared at all.

For a `private_extension_declaration`, each inherited subprogram is declared immediately after the `private_extension_declaration` if the corresponding declaration from the ancestor is visible at that place. Otherwise, the inherited subprogram is not declared for the private extension, though it can be for the full type.

The Class attribute is defined for tagged subtypes in 6.9. In addition, for every subtype S of an untagged private type whose full view is tagged, the following attribute is defined:

**S'Class** Denotes the class-wide subtype corresponding to the full view of S. This attribute is allowed only from the beginning of the private part in which the full view is declared, until the declaration of the full view. After the full view, the Class attribute of the full view can be used.

**NOTE 1** Because a partial view and a full view are two different views of one and the same type, outside of the defining package the characteristics of the type are those defined by the visible part. Within these outside program units the type is just a private type or private extension, and any language rule that applies only to another class of types does not apply. The fact that the full declaration can implement a private type with a type of a particular class (for example, as an array type) is relevant only within the declarative region of the package itself including any child units.

The consequences of this actual implementation are, however, valid everywhere. For example: any default initialization of components takes place; the attribute Size provides the size of the full view; finalization is still done for controlled components of the full view; task dependence rules still apply to components that are task objects.

**NOTE 2** Partial views provide initialization, membership tests, selected components for the selection of discriminants and inherited components, qualification, and explicit conversion. Nonlimited partial views also allow use of assignment statements.

**NOTE 3** For a subtype S of a partial view, S'Size is defined (see 16.3). For an object A of a partial view, the attributes A'Size and A'Address are defined (see 16.3). The Position, First\_Bit, and Last\_Bit attributes are also defined for discriminants and inherited components.

**EXAMPLE** *Example of a type with private operations:*

```

package Key_Manager is
  type Key is private;
  Null_Key : constant Key; -- a deferred constant declaration (see 10.4)
  procedure Get_Key(K : out Key);
  function "<" (X, Y : Key) return Boolean;
private
  type Key is new Natural;
  Null_Key : constant Key := Key'First;
end Key_Manager;

package body Key_Manager is
  Last_Key : Key := Null_Key;
  procedure Get_Key(K : out Key) is
  begin
    Last_Key := Last_Key + 1;
    K := Last_Key;
  end Get_Key;

  function "<" (X, Y : Key) return Boolean is
  begin
    return Natural(X) < Natural(Y);
  end "<";
end Key_Manager;

```

Outside of the package Key\_Manager, the operations available for objects of type Key include assignment, the comparison for equality or inequality, the procedure Get\_Key and the operator "<"; they do not include other relational operators such as ">=", or arithmetic operators.

The explicitly declared operator "<" hides the predefined operator "<" implicitly declared by the full\_type\_declaration. Within the body of the function, an explicit conversion of X and Y to the subtype Natural is necessary to invoke the "<" operator of the parent type. Alternatively, the result of the function can be written as **not** (X >= Y), since the operator ">=" is not redefined.

The value of the variable Last\_Key, declared in the package body, remains unchanged between calls of the procedure Get\_Key. (See also the NOTES of 10.2.)

### 10.3.2 Type Invariants

For a private type, private extension, or interface, the following language-defined assertion aspects may be specified with an `aspect_specification` (see 16.1.1):

#### Type\_Invariant

This aspect shall be specified by an expression, called an *invariant expression*. Type\_Invariant may be specified on a `private_type_declaration`, on a `private_extension_declaration`, or on a `full_type_declaration` that declares the completion of a private type or private extension.

#### Type\_Invariant'Class

This aspect shall be specified by an expression, called an *invariant expression*. Type\_Invariant'Class may be specified on a `private_type_declaration`, a `private_extension_declaration`, or a `full_type_declaration` for an interface type. Type\_Invariant'Class determines a *class-wide type invariant* for a tagged type. The Type\_Invariant'Class aspect is not inherited, but its effects are additive, as defined below.

#### Name Resolution Rules

The expected type for an invariant expression is any boolean type.

Within an invariant expression, the identifier of the first subtype of the associated type denotes the current instance of the type. Within an invariant expression for the Type\_Invariant aspect of a type *T*, the type of this current instance is *T*. Within an invariant expression for the Type\_Invariant'Class aspect of a type *T*, the type of this current instance is interpreted as though it had a (notional) nonabstract type *NT* that is a visible formal derived type whose ancestor type is *T*. The effect of this interpretation is that the only operations that can be applied to this current instance are those defined for such a formal derived type.

#### Legality Rules

The Type\_Invariant'Class aspect shall not be specified for an untagged type. The Type\_Invariant aspect shall not be specified for an abstract type.

If a type extension occurs immediately within the visible part of a package specification, at a point where a private operation of some ancestor is visible and inherited, and a Type\_Invariant'Class expression applies to that ancestor, then the inherited operation shall be abstract or shall be overridden.

#### Static Semantics

If the Type\_Invariant aspect is specified for a type *T*, then the invariant expression applies to *T*.

If the Type\_Invariant'Class aspect is specified for a tagged type *T*, then a *corresponding expression* also applies to each nonabstract descendant *T1* of *T* (including *T* itself if it is nonabstract). The corresponding expression is constructed from the associated expression as follows:

- References to nondiscriminant components of *T* (or to *T* itself) are replaced with references to the corresponding components of *T1* (or to *T1* as a whole).
- References to discriminants of *T* are replaced with references to the corresponding discriminant of *T1*, or to the specified value for the discriminant, if the discriminant is specified by the `derived_type_definition` for some type that is an ancestor of *T1* and a descendant of *T* (see 6.7).

For a nonabstract type *T*, a callable entity is said to be a *boundary entity* for *T* if it is declared within the immediate scope of *T* (or by an instance of a generic unit, and the generic is declared within the immediate scope of type *T*), or is the Read or Input stream-oriented attribute of type *T*, and either:

- *T* is a private type or a private extension and the callable entity is visible outside the immediate scope of type *T* or overrides an inherited operation that is visible outside the immediate scope of *T*;
- or

- $T$  is a record extension, and the callable entity is a primitive operation visible outside the immediate scope of type  $T$  or overrides an inherited operation that is visible outside the immediate scope of  $T$ .

*Dynamic Semantics*

If one or more invariant expressions apply to a nonabstract type  $T$ , then an invariant check is performed at the following places, on the specified object(s):

- After successful initialization of an object of type  $T$  by default (see 6.3.1), the check is performed on the new object unless the partial view of  $T$  has unknown discriminants;
- After successful explicit initialization of the completion of a deferred constant whose nominal type has a part of type  $T$ , if the completion is inside the immediate scope of the full view of  $T$ , and the deferred constant is visible outside the immediate scope of  $T$ , the check is performed on the part(s) of type  $T$ ;
- After successful conversion to type  $T$ , the check is performed on the result of the conversion;
- For a view conversion, outside the immediate scope of  $T$ , that converts from a descendant of  $T$  (including  $T$  itself) to an ancestor of type  $T$  (other than  $T$  itself), a check is performed on the part of the object that is of type  $T$ :
  - after assigning to the view conversion; and
  - after successful return from a call that passes the view conversion as an **in out** or **out** parameter.
- Upon successful return from a call on any callable entity which is a boundary entity for  $T$ , an invariant check is performed on each object which is subject to an invariant check for  $T$ . In the case of a call to a protected operation, the check is performed before the end of the protected action. In the case of a call to a task entry, the check is performed before the end of the rendezvous. The following objects of a callable entity are subject to an invariant check for  $T$ :
  - a result with a nominal type that has a part of type  $T$ ;
  - an **out** or **in out** parameter whose nominal type has a part of type  $T$ ;
  - an access-to-object parameter or result whose designated nominal type has a part of type  $T$ ; or
  - for a procedure or entry, an **in** parameter whose nominal type has a part of type  $T$ .

If the nominal type of a formal parameter (or the designated nominal type of an access-to-object parameter or result) is incomplete at the point of the declaration of the callable entity, and if the completion of that incomplete type does not occur in the same declaration list as the incomplete declaration, then for purposes of the preceding rules the nominal type is considered to have no parts of type  $T$ .

- For a view conversion to a class-wide type occurring within the immediate scope of  $T$ , from a specific type that is a descendant of  $T$  (including  $T$  itself), a check is performed on the part of the object that is of type  $T$ .

If performing checks is required by the Type\_Invariant or Type\_Invariant'Class assertion policies (see 14.4.2) in effect at the point of the corresponding aspect specification applicable to a given type, then the respective invariant expression is considered *enabled*.

The invariant check consists of the evaluation of each enabled invariant expression that applies to  $T$ , on each of the objects specified above. If any of these evaluate to False, Assertions.Assertion\_Error is raised at the point of the object initialization, conversion, or call. If a given call requires more than one evaluation of an invariant expression, either for multiple objects of a single type or for multiple types with invariants, the evaluations are performed in an arbitrary order, and if one of them evaluates to False, it is not specified whether the others are evaluated. Any invariant check is performed prior to copying back any by-copy **in out** or **out** parameters. Invariant checks, any postcondition check, and any constraint or predicate checks associated with **in out** or **out** parameters are performed in an arbitrary order.

For an invariant check on a value of type *T1* based on a class-wide invariant expression inherited from an ancestor type *T*, any operations within the invariant expression that were resolved as primitive operations of the (notional) formal derived type *NT* are bound to the corresponding operations of type *T1* in the evaluation of the invariant expression for the check on *T1*.

The invariant checks performed on a call are determined by the subprogram or entry actually invoked, whether directly, as part of a dispatching call, or as part of a call through an access-to-subprogram value.

NOTE For a call of a primitive subprogram of type *NT* that is inherited from type *T*, the specified checks of the specific invariants of both the types *NT* and *T* are performed. For a call of a primitive subprogram of type *NT* that is overridden for type *NT*, the specified checks of the specific invariants of only type *NT* are performed.

EXAMPLE *Example of a work scheduler where only urgent work can be scheduled for weekends:*

```

package Work_Orders is
  -- See 3.5.1 for type declarations of Level, Day, and Weekday
  type Work_Order is private with
    Type_Invariant => Day_Scheduled (Work_Order) in Weekday
                    or else Priority (Work_Order) = Urgent;
  function Schedule_Work (Urgency : in Level;
                        To_Occur : in Day) return Work_Order
    with Pre => Urgency = Urgent or else To_Occur in Weekday;
  function Day_Scheduled (Order : in Work_Order) return Day;
  function Priority (Order : in Work_Order) return Level;
  procedure Change_Priority (Order : in out Work_Order;
                          New_Priority in Level;
                          Changed : out Boolean)
    with Post => Changed = (Day_Scheduled(Order) in Weekday
                          or else Priority(Order) = Urgent);

private
  type Work_Order is record
    Scheduled : Day;
    Urgency : Level;
  end record;
end Work_Orders;

package body Work_Orders is
  function Schedule_Work (Urgency : in Level;
                        To_Occur : in Day) return Work_Order is
    (Scheduled => To_Occur, Urgency => Urgency);
  function Day_Scheduled (Order : in Work_Order) return Day is
    (Order.Scheduled);
  function Priority (Order : in Work_Order) return Level is
    (Order.Urgency);
  procedure Change_Priority (Order : in out Work_Order;
                          New_Priority : in Level;
                          Changed : out Boolean) is
    begin
      -- Ensure type invariant is not violated
      if Order.Urgency = Urgent or else (Order.Scheduled in Weekday) then
        Changed := True;
        Order.Urgency := New_Priority;
      else
        Changed := False;
      end if;
    end Change_Priority;
end Work_Orders;

```

### 10.3.3 Default Initial Conditions

For a private type or private extension (including a generic formal type), the following language-defined assertion aspect may be specified with an `aspect_specification` (see 16.1.1):

#### Default\_Initial\_Condition

This aspect shall be specified by an expression, called a *default initial condition expression*. `Default_Initial_Condition` may be specified on a `private_type_declaration`, a `private_extension_declaration`, a `formal_private_type_definition`, or a `formal_derived_type_definition`. The `Default_Initial_Condition` aspect is not inherited, but its effects are additive, as defined below.

#### Name Resolution Rules

The expected type for a default initial condition expression is any boolean type.

Within a default initial condition expression associated with a declaration for a type T, a name that denotes the declaration is interpreted as a current instance of a notional (nonabstract) formal derived type NT with ancestor T, that has directly visible primitive operations.

#### Legality Rules

The `Default_Initial_Condition` aspect shall not be specified for a type whose partial view has unknown discriminants, whether explicitly declared or inherited.

#### Static Semantics

If the `Default_Initial_Condition` aspect is specified for a type T, then the default initial condition expression applies to T and to all descendants of T.

#### Dynamic Semantics

If one or more default initial condition expressions apply to a (nonabstract) type T, then a default initial condition check is performed after successful initialization of an object of type T by default (see 6.3.1). In the case of a controlled type, the check is performed after the call to the type's Initialize procedure (see 10.6).

If performing checks is required by the `Default_Initial_Condition` assertion policy (see 14.4.2) in effect at the point of the corresponding `aspect_specification` applicable to a given type, then the respective default initial condition expression is considered enabled.

The default initial condition check consists of the evaluation of each enabled default initial condition expression that applies to T. Any operations within such an expression that were resolved as primitive operations of the (notional) formal derived type NT, are in the evaluation of the expression resolved as for a formal derived type in an instance with T as the actual type for NT (see 15.5.1). These evaluations, if there are more than one, are performed in an arbitrary order. If any of these evaluate to False, `Assertions.Assertion_Error` is raised at the point of the object initialization.

For a generic formal type T, default initial condition checks performed are as determined by the actual type, along with any default initial condition of the formal type itself.

#### Implementation Permissions

Implementations may extend the syntax or semantics of the `Default_Initial_Condition` aspect in an implementation-defined manner.

NOTE For an example of the use of this aspect, see the Vector container definition in A.18.2.

### 10.3.4 Stable Properties of a Type

Certain characteristics of an object of a given type are unchanged by most of the primitive operations of the type. Such characteristics are called *stable properties* of the type.



A *property function* of type *T* is a function with a single parameter of type *T* or of a class-wide type that covers *T*.

A *type property aspect definition* is a list of names written in the syntax of a *positional\_array\_aggregate*. A *subprogram property aspect definition* is a list of names, each optionally preceded by reserved word **not**, also written in the syntax of a *positional\_array\_aggregate*.

For a nonformal private type, nonformal private extension, or full type that does not have a partial view, the following language-defined aspects may be specified with an *aspect\_specification* (see 16.1.1):

#### Stable\_Properties

This aspect shall be specified by a type property aspect definition; each *name* shall statically denote a single property function of the type. This aspect defines the *specific stable property functions* of the associated type.

#### Stable\_Properties'Class

This aspect shall be specified by a type property aspect definition; each *name* shall statically denote a single property function of the type. This aspect defines the *class-wide stable property functions* of the associated type. Unlike most class-wide aspects, *Stable\_Properties'Class* is not inherited by descendant types and subprograms, but the enhanced class-wide postconditions are inherited in the normal manner.

The specific and class-wide stable properties of a type together comprise the stable properties of the type.

For a primitive subprogram, the following language-defined aspects may be specified with an *aspect\_specification* (see 16.1.1):

#### Stable\_Properties

This aspect shall be specified by a subprogram property aspect definition; each *name* shall statically denote a single property function of a type for which the associated subprogram is primitive.

#### Stable\_Properties'Class

This aspect shall be specified by a subprogram property aspect definition; each *name* shall statically denote a single property function of a tagged type for which the associated subprogram is primitive. Unlike most class-wide aspects, *Stable\_Properties'Class* is not inherited by descendant subprograms, but the enhanced class-wide postconditions are inherited in the normal manner.

#### Legality Rules

A stable property function of a type *T* shall have a nonlimited return type and shall be:

- a primitive function with a single parameter of mode **in** of type *T*; or
- a function that is declared immediately within the declarative region in which an ancestor type of *T* is declared and has a single parameter of mode **in** of a class-wide type that covers *T*.

In a subprogram property aspect definition for a subprogram *S*:

- all or none of the items shall be preceded by **not**;
- any property functions mentioned after **not** shall be a stable property function of a type for which *S* is primitive.

If a *subprogram\_renaming\_declaration* declares a primitive subprogram of a type *T*, then the renamed callable entity shall also be a primitive subprogram of type *T* and the two primitive subprograms shall have the same specific stable property functions and the same class-wide stable property functions (see below).

## Static Semantics

For a primitive subprogram  $S$  of a type  $T$ , the specific stable property functions of  $S$  for type  $T$  are:

- if  $S$  has an aspect `Stable_Properties` specified that does not include **not**, those functions denoted in the aspect `Stable_Properties` for  $S$  that have a parameter of  $T$  or  $T$ 'Class;
- if  $S$  has an aspect `Stable_Properties` specified that includes **not**, those functions denoted in the aspect `Stable_Properties` for  $T$ , excluding those denoted in the aspect `Stable_Properties` for  $S$ ;
- if  $S$  does not have an aspect `Stable_Properties`, those functions denoted in the aspect `Stable_Properties` for  $T$ , if any.

A similar definition applies for class-wide stable property functions by replacing aspect `Stable_Properties` with aspect `Stable_Properties'Class` in the above definition.

The *explicit* specific postcondition expression for a subprogram  $S$  is the expression directly specified for  $S$  with the `Post` aspect. Similarly, the *explicit* class-wide postcondition expression for a subprogram  $S$  is the expression directly specified for  $S$  with the `Post'Class` aspect.

For a primitive subprogram  $S$  of a type  $T$  that has a parameter  $P$  of type  $T$ , the parameter is *excluded from stable property checks* if:

- $S$  is a stable property function of  $T$ ;
- $P$  has mode **out**;
- the Global aspect of  $S$  is **null**, and  $P$  has mode **in** and the mode is not overridden by a global aspect.

For every primitive subprogram  $S$  of a type  $T$  that is not an abstract subprogram or null procedure, the specific postcondition expression of  $S$  is modified to include expressions of the form  $F(P) = F(P) \text{ 'Old}$ , all **anded** with each other and any explicit specific postcondition expression, with one such equality included for each specific stable property function  $F$  of  $S$  for type  $T$  that does not occur in the explicit specific postcondition expression of  $S$ , and  $P$  is each parameter of  $S$  that has type  $T$  and is not excluded from stable property checks. The resulting specific postcondition expression of  $S$  is used in place of the explicit specific postcondition expression of  $S$  when interpreting the meaning of the postcondition as defined in 9.1.1.

For every primitive subprogram  $S$  of a type  $T$ , the class-wide postcondition expression of  $S$  is modified to include expressions of the form  $F(P) = F(P) \text{ 'Old}$ , all **anded** with each other and any explicit class-wide postcondition expression, with one such equality included for each class-wide stable property function  $F$  of  $S$  for type  $T$  that does not occur in any class-wide postcondition expression that applies to  $S$ , and  $P$  is each parameter of  $S$  that has type  $T$  and is not excluded from stable property checks. The resulting class-wide postcondition expression of  $S$  is used in place of the explicit class-wide postcondition expression of  $S$  when interpreting the meaning of the postcondition as defined in 9.1.1.

The equality operation that is used in the aforementioned equality expressions is as described in the case of an individual membership test whose `membership_choice` is a `choice_simple_expression` (see 7.5.2).

The `Post` expression additions described above are enabled or disabled depending on the `Post` assertion policy that is in effect at the point of declaration of the subprogram  $S$ . A similar rule applies to the `Post'Class` expression additions.

NOTE For an example of the use of these aspects, see the Vector container definition in A.18.2.

## 10.4 Deferred Constants

Deferred constant declarations may be used to declare constants in the visible part of a package, but with the value of the constant given in the private part. They may also be used to declare constants imported from other languages (see Annex B).

*Legality Rules*

A *deferred constant declaration* is an `object_declaration` with the reserved word **constant** but no initialization expression. The constant declared by a deferred constant declaration is called a *deferred constant*. Unless the Import aspect (see B.1) is True for a deferred constant declaration, the deferred constant declaration requires a completion, which shall be a full constant declaration (called the *full declaration* of the deferred constant).

A deferred constant declaration that is completed by a full constant declaration shall occur immediately within the visible part of a `package_specification`. For this case, the following additional rules apply to the corresponding full declaration:

- The full declaration shall occur immediately within the private part of the same package;
- The deferred and full constants shall have the same type, or shall have statically matching anonymous access subtypes;
- If the deferred constant declaration includes a `subtype_indication` *S* that defines a constrained subtype, then the constraint defined by the `subtype_indication` in the full declaration shall match the constraint defined by *S* statically. On the other hand, if the subtype of the deferred constant is unconstrained, then the full declaration is still allowed to impose a constraint. The constant itself will be constrained, like all constants;
- If the deferred constant declaration includes the reserved word **aliased**, then the full declaration shall also;
- If the subtype of the deferred constant declaration excludes null, the subtype of the full declaration shall also exclude null.

A deferred constant declaration for which the Import aspect is True can appear anywhere that an `object_declaration` is allowed, and has no full constant declaration.

The completion of a deferred constant declaration shall occur before the constant is frozen (see 16.14).

*Dynamic Semantics*

The elaboration of a deferred constant declaration elaborates the `subtype_indication`, `access_definition`, or (only allowed in the case of an imported constant) the `array_type_definition`.

NOTE The full constant declaration for a deferred constant that is of a given private type or private extension is not allowed before the corresponding `full_type_declaration`. This is a consequence of the freezing rules for types (see 16.14).

EXAMPLE *Examples of deferred constant declarations:*

```
Null_Key : constant Key;           -- see 10.3.1
CPU_Identifier : constant String(1..8)
  with Import => True, Convention => Assembler, Link_Name => "CPU_ID";
  -- see B.1
```

## 10.5 Limited Types

A limited type is (a view of) a type for which copying (such as for an `assignment_statement`) is not allowed. A nonlimited type is a (view of a) type for which copying is allowed.

*Legality Rules*

If a tagged record type has any limited components, then the reserved word **limited** shall appear in its `record_type_definition`. If the reserved word **limited** appears in the definition of a `derived_type_definition`, its parent type and any progenitor interfaces shall be limited.

In the following contexts, an `expression` of a limited type is permitted only if each of its operative constituents is newly constructed (see 7.4):

- the initialization expression of an object\_declaration (see 6.3.1)
- the default\_expression of a component\_declaration (see 6.8)
- the expression of a record\_component\_association (see 7.3.1)
- the expression for an ancestor\_part of an extension\_aggregate (see 7.3.2)
- an expression of a positional\_array\_aggregate or the expression of an array\_component\_association (see 7.3.3)
- the base\_expression of a record\_delta\_aggregate (see 7.3.4)
- the qualified\_expression of an initialized allocator (see 7.8)
- the expression of a return statement (see 9.5)
- the return expression of an expression function (see 9.8)
- the default\_expression or actual parameter for a formal object of mode **in** (see 15.4)

#### Static Semantics

A view of a type is *limited* if it is one of the following:

- a type with the reserved word **limited**, **synchronized**, **task**, or **protected** in its definition;
- a class-wide type whose specific type is limited;
- a composite type with a limited component;
- an incomplete view;
- a derived type whose parent is limited and is not an interface.

Otherwise, the type is nonlimited.

There are no predefined equality operators for a limited type.

A type is *immutably limited* if it is one of the following:

- An explicitly limited record type;
- A record extension with the reserved word **limited**;
- A nonformal limited private type that is tagged or has at least one access discriminant with a default\_expression;
- A task type, a protected type, or a synchronized interface;
- A type derived from an immutably limited type.

A descendant of a generic formal limited private type is presumed to be immutably limited except within the body of a generic unit or a body declared within the declarative region of a generic unit, if the formal type is declared within the formal part of the generic unit.

NOTE 1 While it is allowed to write initializations of limited objects, such initializations never copy a limited object. The source of such an assignment operation will be an aggregate or function\_call, and such aggregates and function\_calls will be built directly in the target object (see 10.6).

NOTE 2 As illustrated in 10.3.1, an untagged limited type can become nonlimited under certain circumstances.

EXAMPLE *Example of a package with a limited type:*

```
package IO_Package is
  type File_Name is limited private;
```

```

procedure Open (F : in out File_Name);
procedure Close(F : in out File_Name);
procedure Read (F : in File_Name; Item : out Integer);
procedure Write(F : in File_Name; Item : in Integer);
private
  type File_Name is
    limited record
      Internal_Name : Integer := 0;
    end record;
end IO_Package;

package body IO_Package is
  Limit : constant := 200;
  type File_Descriptor is record ... end record;
  Directory : array (1 .. Limit) of File_Descriptor;
  ...
  procedure Open (F : in out File_Name) is ... end;
  procedure Close(F : in out File_Name) is ... end;
  procedure Read (F : in File_Name; Item : out Integer) is ... end;
  procedure Write(F : in File_Name; Item : in Integer) is ... end;
begin
  ...
end IO_Package;

```

In the example, an outside subprogram making use of IO\_Package can obtain a file name by calling Open and later use it in calls to Read and Write. Thus, outside the package, a file name obtained from Open acts as a kind of password; its internal properties (such as containing a numeric value) are not known and no other operations (such as addition or comparison of internal names) can be performed on a file name. Most importantly, clients of the package cannot make copies of objects of type File\_Name.

This example is characteristic of any case where complete control over the operations of a type is desired. Such packages serve a dual purpose. They prevent a user from making use of the internal structure of the type. They also implement the notion of an encapsulated data type where the only operations on the type are those given in the package specification.

The fact that the full view of File\_Name is explicitly declared **limited** means that parameter passing will always be by reference and function results will always be built directly in the result object (see 9.2 and 9.5).

## 10.6 Assignment and Finalization

Three kinds of actions are fundamental to the manipulation of objects: initialization, finalization, and assignment. Every object is initialized, either explicitly or by default, after being created (for example, by an object\_declaration or allocator). Every object is finalized before being destroyed (for example, by leaving a subprogram\_body containing an object\_declaration, or by a call to an instance of Unchecked\_Deallocation). An assignment operation is used as part of assignment\_statements, explicit initialization, parameter passing, and other operations.

Default definitions for these three fundamental operations are provided by the language, but a *controlled* type gives the user additional control over parts of these operations. In particular, the user can define, for a controlled type, an Initialize procedure which is invoked immediately after the normal default initialization of a controlled object, a Finalize procedure which is invoked immediately before finalization of any of the components of a controlled object, and an Adjust procedure which is invoked as the last step of an assignment to a (nonlimited) controlled object.

### Static Semantics

The following language-defined library package exists:

```

package Ada.Finalization
  with Pure, Nonblocking => False is
    type Controlled is abstract tagged private
      with Preelaborable_Initialization;

```

```

procedure Initialize (Object : in out Controlled) is null;
procedure Adjust      (Object : in out Controlled) is null;
procedure Finalize    (Object : in out Controlled) is null;

type Limited_Controlled is abstract tagged limited private
    with FreeLaborable_Initialization;

procedure Initialize (Object : in out Limited_Controlled) is null;
procedure Finalize   (Object : in out Limited_Controlled) is null;
private
    ... -- not specified by the language
end Ada.Finalization;

```

A controlled type is a descendant of Controlled or Limited\_Controlled. The predefined "=" operator of type Controlled always returns True, since this operator is incorporated into the implementation of the predefined equality operator of types derived from Controlled, as explained in 7.5.2. The type Limited\_Controlled is like Controlled, except that it is limited and it lacks the primitive subprogram Adjust.

A type is said to *need finalization* if:

- it is a controlled type, a task type or a protected type; or
- it has a component whose type needs finalization; or
- it is a class-wide type; or
- it is a partial view whose full view needs finalization; or
- it is one of a number of language-defined types that are explicitly defined to need finalization.

#### Dynamic Semantics

During the elaboration or evaluation of a construct that causes an object to be initialized by default, for every controlled subcomponent of the object that is not assigned an initial value (as defined in 6.3.1), Initialize is called on that subcomponent. Similarly, if the object that is initialized by default as a whole is controlled, Initialize is called on the object.

For an `extension_aggregate` whose `ancestor_part` is a `subtype_mark` denoting a controlled subtype, the Initialize procedure of the ancestor type is called, unless that Initialize procedure is abstract.

Initialize and other initialization operations are done in an arbitrary order, except as follows. Initialize is applied to an object after initialization of its subcomponents, if any (including both implicit initialization and Initialize calls). If an object has a component with an access discriminant constrained by a per-object expression, Initialize is applied to this component after any components that do not have such discriminants. For an object with several components with such a discriminant, Initialize is applied to them in order of their `component_declarations`. For an allocator, any task activations follow all calls on Initialize.

When a target object with any controlled parts is assigned a value, either when created or in a subsequent `assignment_statement`, the *assignment operation* proceeds as follows:

- The value of the target becomes the assigned value.
- The value of the target is *adjusted*.

To adjust the value of a composite object, the values of the components of the object are first adjusted in an arbitrary order, and then, if the object is nonlimited controlled, Adjust is called. Adjusting the value of an elementary object has no effect, nor does adjusting the value of a composite object with no controlled parts.

For an `assignment_statement`, after the name and expression have been evaluated, and any conversion (including constraint checking) has been done, an anonymous object is created, and the value is assigned into it; that is, the assignment operation is applied. (Assignment includes value adjustment.) The target of the `assignment_statement` is then finalized. The value of the



anonymous object is then assigned into the target of the `assignment_statement`. Finally, the anonymous object is finalized. As explained below, the implementation may eliminate the intermediate anonymous object, so this description subsumes the one given in 8.2.

When a function call or `aggregate` is used to initialize an object, the result of the function call or `aggregate` is an anonymous object, which is assigned into the newly-created object. For such an assignment, the anonymous object may be *built in place*, in which case the assignment does not involve any copying. Under certain circumstances, the anonymous object is required to be built in place. In particular:

- If the full type of any part of the object is immutably limited, the anonymous object is built in place.
- In the case of an `aggregate`, if the full type of any part of the newly-created object is controlled, the anonymous object is built in place.
- In other cases, it is unspecified whether the anonymous object is built in place.

Notwithstanding what this document says elsewhere, if an object is built in place:

- Upon successful completion of the return statement or `aggregate`, the anonymous object *mutates into* the newly-created object; that is, the anonymous object ceases to exist, and the newly-created object appears in its place.
- Finalization is not performed on the anonymous object.
- Adjustment is not performed on the newly-created object.
- All access values that designate parts of the anonymous object now designate the corresponding parts of the newly-created object.
- All renamings of parts of the anonymous object now denote views of the corresponding parts of the newly-created object.
- Coextensions of the anonymous object become coextensions of the newly-created object.

#### Implementation Permissions

An implementation is allowed to relax the above rules for `assignment_statements` in the following ways:

- If an object is assigned the value of that same object, the implementation may omit the entire assignment.
- For assignment of a noncontrolled type, the implementation may finalize and assign each component of the variable separately (rather than finalizing the entire variable and assigning the entire new value) unless a discriminant of the variable is changed by the assignment.
- The implementation may avoid creating an anonymous object if the value being assigned is the result of evaluating a `name` denoting an object (the source object) whose storage cannot overlap with the target. If the source object can overlap with the target object, then the implementation can avoid the need for an intermediary anonymous object by exercising one of the above permissions and perform the assignment one component at a time (for an overlapping array assignment), or not at all (for an assignment where the target and the source of the assignment are the same object).

Furthermore, an implementation is permitted to omit implicit Initialize, Adjust, and Finalize calls and associated assignment operations on an object of a nonlimited controlled type provided that:

- any omitted Initialize call is not a call on a user-defined Initialize procedure, and
- any usage of the value of the object after the implicit Initialize or Adjust call and before any subsequent Finalize call on the object does not change the external effect of the program, and
- after the omission of such calls and operations, any execution of the program that executes an Initialize or Adjust call on an object or initializes an object by an `aggregate` will also later execute a Finalize call on the object and will always do so prior to assigning a new value to the object, and
- the assignment operations associated with omitted Adjust calls are also omitted.

This permission applies to Adjust and Finalize calls even if the implicit calls have additional external effects.

### 10.6.1 Completion and Finalization

This subclause defines *completion* and *leaving* of the execution of constructs and entities. A *master* is the execution of a construct that includes finalization of local objects after it is complete (and after waiting for any local tasks — see 12.3), but before leaving. Other constructs and entities are left immediately upon completion.

#### Dynamic Semantics

The execution of a construct or entity is *complete* when the end of that execution has been reached, or when a transfer of control (see 8.1) causes it to be abandoned. Completion due to reaching the end of execution, or due to the transfer of control of an `exit_statement`, `return_statement`, `goto_statement`, or `requeue_statement` or of the selection of a `terminate_alternative` is *normal completion*. Completion is *abnormal* otherwise — when control is transferred out of a construct due to abort or the raising of an exception.

After execution of a construct or entity is complete, it is *left*, meaning that execution continues with the next action, as defined for the execution that is taking place. Leaving an execution happens immediately after its completion, except in the case of the execution of a *master construct*: a body other than a `package_body`; a `statement`; or an `expression`, `function_call`, or `range` that is not part of an `enclosing expression`, `function_call`, `range`, or `simple_statement` other than a `simple_return_statement`. The term *master* by itself refers to the execution of a master construct. A master is finalized after it is complete, and before it is left.

For the *finalization* of a master, dependent tasks are first awaited, as explained in 12.3. Then each object whose accessibility level is the same as that of the master is finalized if the object was successfully initialized and still exists. These actions are performed whether the master is left by reaching the last statement or via a transfer of control. When a transfer of control causes completion of an execution, each included master is finalized in order, from innermost outward.

For the *finalization* of an object:

- If the full type of the object is an elementary type, finalization has no effect;
- If the full type of the object is a tagged type, and the tag of the object identifies a controlled type, the Finalize procedure of that controlled type is called;
- If the full type of the object is a protected type, or if the full type of the object is a tagged type and the tag of the object identifies a protected type, the actions defined in 12.4 are performed;
- If the full type of the object is a composite type, then after performing the above actions, if any, every component of the object is finalized in an arbitrary order, except as follows: if the object has a component with an access discriminant constrained by a per-object expression, this component is finalized before any components that do not have such discriminants; for an object with several components with such a discriminant, they are finalized in the reverse of the order of their `component_declarations`;
- If the object has coextensions (see 6.10.2), each coextension is finalized after the object whose access discriminant designates it.

Immediately before an instance of `Unchecked_Deallocation` reclaims the storage of an object, the object is finalized. If an instance of `Unchecked_Deallocation` is never applied to an object created by an `allocator`, the object will still exist when the corresponding master completes, and it will be finalized then.

The finalization of a master performs finalization of objects created by declarations in the master in the reverse order of their creation. After the finalization of a master is complete, the objects finalized as part of its finalization cease to *exist*, as do any types and subtypes defined and created within the master.

Each nonderived access type *T* has an associated *collection*, which is the set of objects created by allocators of *T*, or of types derived from *T*. `Unchecked_Deallocation` removes an object from its collection. Finalization of a collection consists of finalization of each object in the collection, in an arbitrary order. The collection of an access type is an object implicitly declared at the following place:

- For a named access type, the first freezing point (see 16.14) of the type.
- For the type of an access parameter, the call that contains the `allocator`.
- For the type of an access result, within the master of the call (see 6.10.2).
- For any other anonymous access type, the first freezing point of the innermost enclosing declaration.

The target of an `assignment_statement` is finalized before copying in the new value, as explained in 10.6.

The master of an object is the master enclosing its creation whose accessibility level (see 6.10.2) is equal to that of the object, except in the case of an anonymous object representing the result of an aggregate or function call. If such an anonymous object is part of the result of evaluating the actual parameter expression for an explicitly aliased parameter of a function call, the master of the object is the innermost master enclosing the evaluation of the aggregate or function call, excluding the aggregate or function call itself. Otherwise, the master of such an anonymous object is the innermost master enclosing the evaluation of the aggregate or function call, which may be the aggregate or function call itself.

In the case of an `expression` that is a master, finalization of any (anonymous) objects occurs after completing evaluation of the `expression` and all use of the objects, prior to starting the execution of any subsequent construct.

#### *Bounded (Run-Time) Errors*

It is a bounded error for a call on `Finalize` or `Adjust` that occurs as part of object finalization or assignment to propagate an exception. The possible consequences depend on what action invoked the `Finalize` or `Adjust` operation:

- For a `Finalize` invoked as part of an `assignment_statement`, `Program_Error` is raised at that point.
- For an `Adjust` invoked as part of assignment operations other than those invoked as part of an `assignment_statement`, some of the adjustments due to be performed can be performed, and then `Program_Error` is raised. During its propagation, finalization may be applied to objects whose `Adjust` failed. For an `Adjust` invoked as part of an `assignment_statement`, any other adjustments due to be performed are performed, and then `Program_Error` is raised.
- For a `Finalize` invoked as part of a call on an instance of `Unchecked_Deallocation`, any other finalizations due to be performed are performed, and then `Program_Error` is raised.
- For a `Finalize` invoked due to reaching the end of the execution of a master, any other finalizations associated with the master are performed, and `Program_Error` is raised immediately after leaving the master.
- For a `Finalize` invoked by the transfer of control of an `exit_statement`, `return_statement`, `goto_statement`, or `requeue_statement`, `Program_Error` is raised no earlier than after the finalization of the master being finalized when the exception occurred, and no later than the point where normal execution would have continued. Any other finalizations due to be performed up to that point are performed before raising `Program_Error`.
- For a `Finalize` invoked by a transfer of control that is due to raising an exception, any other finalizations due to be performed for the same master are performed; `Program_Error` is raised immediately after leaving the master.
- For a `Finalize` invoked by a transfer of control due to an abort or selection of a terminate alternative, the exception is ignored; any other finalizations due to be performed are performed.

*Implementation Permissions*

If the execution of an `allocator` propagates an exception, any parts of the allocated object that were successfully initialized may be finalized as part of the finalization of the innermost master enclosing the `allocator`.

The implementation may finalize objects created by `allocators` for an access type whose storage pool supports subpools (see 16.11.4) as if the objects were created (in an arbitrary order) at the point where the storage pool was elaborated instead of at the first freezing point of the access type.

NOTE 1 The rules of Clause 13 imply that immediately prior to partition termination, Finalize operations are applied to library-level controlled objects (including those created by `allocators` of library-level access types, except those already finalized). This occurs after waiting for library-level tasks to terminate.

NOTE 2 A constant is only constant between its initialization and finalization. Both initialization and finalization are allowed to change the value of a constant.

NOTE 3 Abort is deferred during certain operations related to controlled types, as explained in 12.8. Those rules prevent an abort from causing a controlled object to be left in an ill-defined state.

NOTE 4 The Finalize procedure is called upon finalization of a controlled object, even if Finalize was called earlier, either explicitly or as part of an assignment; hence, if a controlled type is visibly controlled (implying that its Finalize primitive is directly callable), or is nonlimited (implying that assignment is allowed), its Finalize procedure is ideally designed to have no ill effect if it is applied a second time to the same object.

## 11 Visibility Rules

The rules defining the scope of declarations and the rules defining which identifiers, `character_literals`, and `operator_symbols` are visible at (or from) various places in the text of the program are described in this clause. The formulation of these rules uses the notion of a declarative region.

As explained in Clause 6, a declaration declares a view of an entity and associates a defining name with that view. The view comprises an identification of the viewed entity, and possibly additional properties. A usage name denotes a declaration. It also denotes the view declared by that declaration, and denotes the entity of that view. Thus, two different usage names can denote two different views of the same entity; in this case they denote the same entity.

### 11.1 Declarative Region

*Static Semantics*

For each of the following constructs, there is a portion of the program text called its *declarative region*, within which nested declarations can occur:

- any declaration, other than that of an enumeration type, that is not a completion of a previous declaration;
- an `access_definition`;
- an `iterated_component_association`;
- an `iterated_element_association`;
- a `quantified_expression`;
- a `declare_expression`;
- a `block_statement`;
- a `loop_statement`;
- an `extended_return_statement`;
- an `accept_statement`;

— an `exception_handler`.

The declarative region includes the text of the construct together with additional text determined (recursively), as follows:

- If a declaration is included, so is its completion, if any.
- If the declaration of a library unit (including Standard — see 13.1.1) is included, so are the declarations of any child units (and their completions, by the previous rule). The child declarations occur after the declaration.
- If a `body_stub` is included, so is the corresponding subunit.
- If a `type_declaration` is included, then so is a corresponding `record_representation_clause`, if any.

The declarative region of a declaration is also called the *declarative region* of any view or entity declared by the declaration.

A declaration occurs *immediately within* a declarative region if this region is the innermost declarative region that encloses the declaration (the *immediately enclosing* declarative region), not counting the declarative region (if any) associated with the declaration itself.

A declaration is *local* to a declarative region if the declaration occurs immediately within the declarative region. An entity is *local* to a declarative region if the entity is declared by a declaration that is local to the declarative region.

A declaration is *global* to a declarative region if the declaration occurs immediately within another declarative region that encloses the declarative region. An entity is *global* to a declarative region if the entity is declared by a declaration that is global to the declarative region.

NOTE 1 The children of a parent library unit are inside the parent's declarative region, even though they do not occur inside the parent's declaration or body. This implies that one can use (for example) "P.Q" to refer to a child of P whose defining name is Q, and that after "use P;" Q can refer (directly) to that child.

NOTE 2 As explained above and in 13.1.1, all library units are descendants of Standard, and so are contained in the declarative region of Standard. They are *not* inside the declaration or body of Standard, but they *are* inside its declarative region.

NOTE 3 For a declarative region that comes in multiple parts, the text of the declarative region does not include any of the text that appears between the parts. Thus, when a portion of a declarative region is said to extend from one place to another in the declarative region, the portion does not contain any of the text that appears between the parts of the declarative region.

## 11.2 Scope of Declarations

For each declaration, the language rules define a certain portion of the program text called the *scope* of the declaration. The scope of a declaration is also called the scope of any view or entity declared by the declaration. Within the scope of an entity, and only there, there are places where it is legal to refer to the declared entity. These places are defined by the rules of visibility and overloading.

### Static Semantics

The *immediate scope* of a declaration is a portion of the declarative region immediately enclosing the declaration. The immediate scope starts at the beginning of the declaration, except in the case of an overloadable declaration, in which case the immediate scope starts just after the place where the profile of the callable entity is determined (which is at the end of the `_specification` for the callable entity, or at the end of the `generic_instantiation` if an instance). The immediate scope extends to the end of the declarative region, with the following exceptions:

- The immediate scope of a `library_item` includes only its semantic dependents.
- The immediate scope of a declaration in the private part of a library unit does not include the visible part of any public descendant of that library unit.



The *visible part* of (a view of) an entity is a portion of the text of its declaration containing declarations that are visible from outside. The *private part* of (a view of) an entity that has a visible part contains all declarations within the declaration of (the view of) the entity, except those in the visible part; these are not visible from outside. Visible and private parts are defined only for these kinds of entities: callable entities, other program units, and composite types.

- The visible part of a view of a callable entity is its profile.
- The visible part of a composite type other than a task or protected type consists of the declarations of all components declared (explicitly or implicitly) within the `type_declaration`.
- The visible part of a generic unit includes the `generic_formal_part`. For a generic package, it also includes the first list of `basic_declarative_items` of the `package_specification`. For a generic subprogram, it also includes the profile.
- The visible part of a package, task unit, or protected unit consists of declarations in the program unit's declaration other than those following the reserved word **private**, if any; see 10.1 and 15.7 for packages, 12.1 for task units, and 12.4 for protected units.

The scope of a declaration always contains the immediate scope of the declaration. In addition, for a given declaration that occurs immediately within the visible part of an outer declaration, or is a public child of an outer declaration, the scope of the given declaration extends to the end of the scope of the outer declaration, except that the scope of a `library_item` includes only its semantic dependents.

The scope of an `attribute_definition_clause` is identical to the scope of a declaration that would occur at the point of the `attribute_definition_clause`. The scope of an `aspect_specification` is identical to the scope of the associated declaration.

The immediate scope of a declaration is also the immediate scope of the entity or view declared by the declaration. Similarly, the scope of a declaration is also the scope of the entity or view declared by the declaration.

The immediate scope of a pragma that is not used as a configuration pragma is defined to be the region extending from immediately after the pragma to the end of the declarative region immediately enclosing the pragma.

NOTE There are notations for denoting visible declarations that are not directly visible. For example, `parameter_specifications` are in the visible part of a `subprogram_declaration` so that they can be used in named-notation calls appearing outside the called subprogram. For another example, declarations of the visible part of a package can be denoted by expanded names appearing outside the package, and can be made directly visible by a `use_clause`.

### 11.3 Visibility

The *visibility rules*, given below, determine which declarations are visible and directly visible at each place within a program. The visibility rules apply to both explicit and implicit declarations.

#### Static Semantics

A declaration is defined to be *directly visible* at places where a name consisting of only an `identifier` or `operator_symbol` is sufficient to denote the declaration; that is, no `selected_component` notation or special context (such as preceding `=>` in a named association) is necessary to denote the declaration. A declaration is defined to be *visible* wherever it is directly visible, as well as at other places where some name (such as a `selected_component`) can denote the declaration.

The syntactic category `direct_name` is used to indicate contexts where direct visibility is required. The syntactic category `selector_name` is used to indicate contexts where visibility, but not direct visibility, is required.

There are two kinds of direct visibility: *immediate visibility* and *use-visibility*. A declaration is immediately visible at a place if it is directly visible because the place is within its immediate scope. A



declaration is use-visible if it is directly visible because of a `use_clause` (see 11.4). Both conditions can apply.

A declaration can be *hidden*, either from direct visibility, or from all visibility, within certain parts of its scope. Where *hidden from all visibility*, it is not visible at all (neither using a `direct_name` nor a `selector_name`). Where *hidden from direct visibility*, only direct visibility is lost; visibility using a `selector_name` is still possible.

Two or more declarations are *overloaded* if they all have the same defining name and there is a place where they are all directly visible.

The declarations of callable entities (including enumeration literals) are *overloadable*, meaning that overloading is allowed for them.

Two declarations are *homographs* if they have the same defining name, and, if both are overloadable, their profiles are type conformant. An inner declaration hides any outer homograph from direct visibility.

Two homographs are not generally allowed immediately within the same declarative region unless one *overrides* the other (see Legality Rules below). The only declarations that are *overridable* are the implicit declarations for predefined operators and inherited primitive subprograms. A declaration overrides another homograph that occurs immediately within the same declarative region in the following cases:

- A declaration that is not overridable overrides one that is overridable, regardless of which declaration occurs first;
- The implicit declaration of an inherited operator overrides that of a predefined operator;
- An implicit declaration of an inherited subprogram overrides a previous implicit declaration of an inherited subprogram.
- If two or more homographs are implicitly declared at the same place:
  - If at least one is a subprogram that is neither a null procedure nor an abstract subprogram, and does not require overriding (see 6.9.3), then they override those that are null procedures, abstract subprograms, or require overriding. If more than one such homograph remains that is not thus overridden, then they are all hidden from all visibility.
  - Otherwise (all are null procedures, abstract subprograms, or require overriding), then any null procedure overrides all abstract subprograms and all subprograms that require overriding; if more than one such homograph remains that is not thus overridden, then if the profiles of the remaining homographs are all fully conformant with one another, one is chosen arbitrarily; if not, they are all hidden from all visibility.
- For an implicit declaration of a primitive subprogram in a generic unit, there is a copy of this declaration in an instance. However, a whole new set of primitive subprograms is implicitly declared for each type declared within the visible part of the instance. These new declarations occur immediately after the type declaration, and override the copied ones. The copied ones can be called only from within the instance; the new ones can be called only from outside the instance, although for tagged types, the body of a new one can be executed by a call to an old one.

A declaration is visible within its scope, except where hidden from all visibility, as follows:

- An overridden declaration is hidden from all visibility within the scope of the overriding declaration.
- A declaration is hidden from all visibility until the end of the declaration, except:
  - For a record type or record extension, the declaration is hidden from all visibility only until the reserved word **record**;
  - For a `package_declaration`, `generic_package_declaration`, `subprogram_body`, or `expression_function_declaration`, the declaration is hidden from all visibility only until the reserved word **is** of the declaration;

- For a task declaration or protected declaration, the declaration is hidden from all visibility only until the reserved word **with** of the declaration if there is one, or the reserved word **is** of the declaration if there is no **with**.
- If the completion of a declaration is a declaration, then within the scope of the completion, the first declaration is hidden from all visibility. Similarly, a `discriminant_specification` or `parameter_specification` is hidden within the scope of a corresponding `discriminant_specification` or `parameter_specification` of a corresponding completion, or of a corresponding `accept_statement`.
- The declaration of a library unit (including a `library_unit_renaming_declaration`) is hidden from all visibility at places outside its declarative region that are not within the scope of a `nonlimited_with_clause` that mentions it. The limited view of a library package is hidden from all visibility at places that are not within the scope of a `limited_with_clause` that mentions it; in addition, the limited view is hidden from all visibility within the declarative region of the package, as well as within the scope of any `nonlimited_with_clause` that mentions the package. Where the declaration of the limited view of a package is visible, any name that denotes the package denotes the limited view, including those provided by a package renaming.
- For each declaration or renaming of a generic unit as a child of some parent generic package, there is a corresponding declaration nested immediately within each instance of the parent. Such a nested declaration is hidden from all visibility except at places that are within the scope of a `with_clause` that mentions the child.

A declaration with a `defining_identifier` or `defining_operator_symbol` is immediately visible (and hence directly visible) within its immediate scope except where hidden from direct visibility, as follows:

- A declaration is hidden from direct visibility within the immediate scope of a homograph of the declaration, if the homograph occurs within an inner declarative region;
- A declaration is also hidden from direct visibility where hidden from all visibility.

An `attribute_definition_clause` or an `aspect_specification` is *visible* everywhere within its scope.

#### Name Resolution Rules

A `direct_name` shall resolve to denote a directly visible declaration whose defining name is the same as the `direct_name`. A `selector_name` shall resolve to denote a visible declaration whose defining name is the same as the `selector_name`.

These rules on visibility and direct visibility do not apply in a `context_clause`, a `parent_unit_name`, or a pragma that appears at the place of a `compilation_unit`. For those contexts, see the rules in 13.1.6.

#### Legality Rules

A nonoverridable declaration is illegal if there is a homograph occurring immediately within the same declarative region that is visible at the place of the declaration, and is not hidden from all visibility by the nonoverridable declaration. In addition, a type extension is illegal if somewhere within its immediate scope it has two visible components with the same name. Similarly, the `context_clause` for a compilation unit is illegal if it mentions (in a `with_clause`) some library unit, and there is a homograph of the library unit that is visible at the place of the compilation unit, and the homograph and the mentioned library unit are both declared immediately within the same declarative region. These rules also apply to dispatching operations declared in the visible part of an instance of a generic unit. However, they do not apply to other overloadable declarations in an instance; such declarations may have type conformant profiles in the instance, so long as the corresponding declarations in the generic were not type conformant.

NOTE 1 Visibility for compilation units follows from the definition of the environment in 13.1.4, except that it is necessary to apply a `with_clause` to obtain visibility to a `library_unit_declaration` or `library_unit_renaming_declaration`.

NOTE 2 In addition to the visibility rules given above, the meaning of the occurrence of a `direct_name` or `selector_name` at a given place in the text can depend on the overloading rules (see 11.6).

NOTE 3 Not all contexts where an identifier, `character_literal`, or `operator_symbol` are allowed require visibility of a corresponding declaration. Contexts where visibility is not required are identified by using one of these three syntactic categories directly in a syntax rule, rather than using `direct_name` or `selector_name`.

### 11.3.1 Overriding Indicators

An `overriding_indicator` is used to declare that an operation is intended to **override** (or not **override**) an inherited operation.

#### Syntax

`overriding_indicator ::= [not] overriding`

#### Legality Rules

If an `abstract_subprogram_declaration`, `null_procedure_declaration`, `expression_function_declaration`, `subprogram_body`, `subprogram_body_stub`, `subprogram_renaming_declaration`, `generic_instantiation` of a `subprogram`, or `subprogram_declaration` other than a `protected_subprogram` has an `overriding_indicator`, then:

- the operation shall be a primitive operation for some type;
- if the `overriding_indicator` is **overriding**, then the operation shall override a homograph at the place of the declaration or body;
- if the `overriding_indicator` is **not overriding**, then the operation shall not override any homograph (at any place).

In addition to the places where Legality Rules normally apply, these rules also apply in the private part of an instance of a generic unit.

NOTE Rules for `overriding_indicators` of task and protected entries and of protected subprograms are found in 12.5.2 and 12.4, respectively.

EXAMPLE *Example of use of an overriding indicator when declaring a security queue derived from the Queue interface of 6.9.4:*

```
type Security_Queue is new Queue with record ...;

overriding
procedure Append(Q : in out Security_Queue; Person : in Person_Name);

overriding
procedure Remove_First(Q : in out Security_Queue; Person : out Person_Name);

overriding
function Cur_Count(Q : in Security_Queue) return Natural;

overriding
function Max_Count(Q : in Security_Queue) return Natural;

not overriding
procedure Arrest(Q : in out Security_Queue; Person : in Person_Name);
```

*The first four subprogram declarations guarantee that these subprograms will override the four subprograms inherited from the Queue interface. A misspelling in one of these subprograms will be detected at compile time by the implementation. Conversely, the declaration of Arrest guarantees that this is a new operation.*

## 11.4 Use Clauses

A `use_package_clause` achieves direct visibility of declarations that appear in the visible part of a package; a `use_type_clause` achieves direct visibility of the primitive operators of a type.

### Syntax

```
use_clause ::= use_package_clause | use_type_clause
use_package_clause ::= use package_name {, package_name};
use_type_clause ::= use [all] type subtype_mark {, subtype_mark};
```

### Legality Rules

A `package_name` of a `use_package_clause` shall denote a nonlimited view of a package.

### Static Semantics

For each `use_clause`, there is a certain region of text called the *scope* of the `use_clause`. For a `use_clause` within a `context_clause` of a `library_unit_declaration` or `library_unit_renaming_declaration`, the scope is the entire declarative region of the declaration. For a `use_clause` within a `context_clause` of a body, the scope is the entire body and any subunits (including multiply nested subunits). The scope does not include `context_clauses` themselves.

For a `use_clause` immediately within a declarative region, the scope is the portion of the declarative region starting just after the `use_clause` and extending to the end of the declarative region. However, the scope of a `use_clause` in the private part of a library unit does not include the visible part of any public descendant of that library unit.

A package is *named* in a `use_package_clause` if it is denoted by a `package_name` of that clause. A type is *named* in a `use_type_clause` if it is determined by a `subtype_mark` of that clause.

For each package named in a `use_package_clause` whose scope encloses a place, each declaration that occurs immediately within the declarative region of the package is *potentially use-visible* at this place if the declaration is visible at this place. For each type  $T$  or  $T^{\text{Class}}$  named in a `use_type_clause` whose scope encloses a place, the declaration of each primitive operator of type  $T$  is *potentially use-visible* at this place if its declaration is visible at this place. If a `use_type_clause` whose scope encloses a place includes the reserved word **all**, then the following entities are also *potentially use-visible* at this place if the declaration of the entity is visible at this place:

- Each primitive subprogram of  $T$  including each enumeration literal (if any);
- Each subprogram that is declared immediately within the declarative region in which an ancestor type of  $T$  is declared and that operates on a class-wide type that covers  $T$ .

Certain implicit declarations may become *potentially use-visible* in certain contexts as described in 15.6.

A declaration is *use-visible* if it is *potentially use-visible*, except in these naming-conflict cases:

- A *potentially use-visible* declaration is not *use-visible* if the place considered is within the immediate scope of a homograph of the declaration.
- *Potentially use-visible* declarations that have the same identifier are not *use-visible* unless each of them is an overloadable declaration.

### Dynamic Semantics

The elaboration of a `use_clause` has no effect.

EXAMPLE 1 Example of a use clause in a context clause:

```
with Ada.Calendar; use Ada;
```

EXAMPLE 2 *Example of a use type clause:*

```
use type Rational_Numbers.Rational; -- see 10.1
Two_Thirds: Rational_Numbers.Rational := 2/3;
```

## 11.5 Renaming Declarations

A `renaming_declaration` declares another name for an entity, such as an object, exception, package, subprogram, entry, or generic unit. Alternatively, a `subprogram_renaming_declaration` can be the completion of a previous `subprogram_declaration`.

### Syntax

```
renaming_declaration ::=
  object_renaming_declaration
| exception_renaming_declaration
| package_renaming_declaration
| subprogram_renaming_declaration
| generic_renaming_declaration
```

### Dynamic Semantics

The elaboration of a `renaming_declaration` evaluates the name that follows the reserved word **renames** and thereby determines the view and entity denoted by this name (the *renamed view* and *renamed entity*). A name that denotes the `renaming_declaration` denotes (a new view of) the renamed entity.

NOTE 1 Renaming can be used to resolve name conflicts and to act as a shorthand. Renaming with a different identifier or operator\_symbol does not hide the old name; the new name and the old name can be visible at different places.

NOTE 2 A subtype defined without any additional constraint can be used to achieve the effect of renaming another subtype (including a task or protected subtype) as in

```
subtype Mode is Ada.Text_IO.File_Mode;
```

### 11.5.1 Object Renaming Declarations

An `object_renaming_declaration` is used to rename an object or value.

### Syntax

```
object_renaming_declaration ::=
  defining_identifier [: [null_exclusion] subtype_mark] renames object_name
  [aspect_specification];
| defining_identifier : access_definition renames object_name
  [aspect_specification];
```

### Name Resolution Rules

The type of the *object\_name* shall resolve to the type determined by the `subtype_mark`, if present. If no `subtype_mark` or `access_definition` is present, the expected type of the *object\_name* is any type.

In the case where the type is defined by an `access_definition`, the type of the *object\_name* shall resolve to an anonymous access type. If the anonymous access type is an access-to-object type, the type of the *object\_name* shall have the same designated type as that of the `access_definition`. If the anonymous access type is an access-to-subprogram type, the type of the *object\_name* shall have a designated profile that is type conformant with that of the `access_definition`.

### Legality Rules

The renamed entity shall be an object or value.

In the case where the type is defined by an `access_definition`, the type of the renamed entity and the type defined by the `access_definition`:

- shall both be access-to-object types with statically matching designated subtypes and with both or neither being access-to-constant types; or
- shall both be access-to-subprogram types with subtype conformant designated profiles.

For an `object_renaming_declaration` with a `null_exclusion` or an `access_definition` that has a `null_exclusion`, the subtype of the `object_name` shall exclude null. In addition, if the `object_renaming_declaration` occurs within the body of a generic unit *G* or within the body of a generic unit declared within the declarative region of generic unit *G*, then:

- if the `object_name` statically denotes a generic formal object of mode **in out** of *G*, then the declaration of that object shall have a `null_exclusion`;
- if the `object_name` statically denotes a call of a generic formal function of *G*, then the declaration of the result of that function shall have a `null_exclusion`.

In the case where the `object_name` is a `qualified_expression` with a nominal subtype *S* and whose expression is a name that denotes an object *Q*:

- if *S* is an elementary subtype, then:
  - *Q* shall be a constant other than a dereference of an access type; or
  - the nominal subtype of *Q* shall be statically compatible with *S*; or
  - *S* shall statically match the base subtype of its type if scalar, or the first subtype of its type if an access type.
- if *S* is a composite subtype, then *Q* shall be known to be constrained or *S* shall statically match the first subtype of its type.

The renamed entity shall not be a subcomponent that depends on discriminants of an object whose nominal subtype is unconstrained unless the object is known to be constrained. A `slice` of an array shall not be renamed if this restriction disallows renaming of the array.

In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

#### Static Semantics

An `object_renaming_declaration` declares a new view of the renamed entity whose properties are identical to those of the renamed view. Thus, the properties of the renamed entity are not affected by the `renaming_declaration`. In particular, its nominal subtype, whether it is a value or an object, its value if it is an object, and whether or not it is a constant, are unaffected; similarly, the constraints and other properties of its nominal subtype are not affected by renaming (any constraint implied by the `subtype_mark` or `access_definition` of the `object_renaming_declaration` is ignored).

EXAMPLE 1 *Example of renaming an object:*

```
declare
  L : Person renames Leftmost_Person; -- see 6.10.1
begin
  L.Age := L.Age + 1;
end;
```

EXAMPLE 2 *Example of renaming a value:*

```
Uno renames One; -- see 6.3.2
```



### 11.5.2 Exception Renaming Declarations

An `exception_renaming_declaration` is used to rename an exception.

#### Syntax

```
exception_renaming_declaration ::=
  defining_identifier : exception renames exception_name
  [aspect_specification];
```

#### Legality Rules

The renamed entity shall be an exception.

#### Static Semantics

An `exception_renaming_declaration` declares a new view of the renamed exception.

EXAMPLE *Example of renaming an exception:*

```
EOF : exception renames Ada.IO_Exceptions.End_Error; -- see A.13
```

### 11.5.3 Package Renaming Declarations

A `package_renaming_declaration` is used to rename a package.

#### Syntax

```
package_renaming_declaration ::=
  package defining_program_unit_name renames package_name
  [aspect_specification];
```

#### Legality Rules

The renamed entity shall be a package.

If the `package_name` of a `package_renaming_declaration` denotes a limited view of a package *P*, then a name that denotes the `package_renaming_declaration` shall occur only within the immediate scope of the renaming or the scope of a `with_clause` that mentions the package *P* or, if *P* is a nested package, the innermost library package enclosing *P*.

#### Static Semantics

A `package_renaming_declaration` declares a new view of the renamed package.

At places where the declaration of the limited view of the renamed package is visible, a name that denotes the `package_renaming_declaration` denotes a limited view of the package (see 13.1.1).

EXAMPLE *Example of renaming a package:*

```
package TM renames Table_Manager;
```

### 11.5.4 Subprogram Renaming Declarations

A `subprogram_renaming_declaration` can serve as the completion of a `subprogram_declaration`; such a renaming\_declaration is called a *renaming-as-body*. A `subprogram_renaming_declaration` that is not a completion is called a *renaming-as-declaration*, and is used to rename a subprogram (possibly an enumeration literal) or an entry.

#### Syntax

```
subprogram_renaming_declaration ::=
  [overriding_indicator]
  subprogram_specification renames callable_entity_name
  [aspect_specification];
```

*Name Resolution Rules*

The expected profile for the *callable\_entity\_name* is the profile given in the *subprogram\_specification*.

*Legality Rules*

The profile of a renaming-as-declaration shall be mode conformant with that of the renamed callable entity.

For a parameter or result subtype of the *subprogram\_specification* that has an explicit *null\_exclusion*:

- if the *callable\_entity\_name* statically denotes a generic formal subprogram of a generic unit *G*, and the *subprogram\_renaming\_declaration* occurs within the body of a generic unit *G* or within the body of a generic unit declared within the declarative region of the generic unit *G*, then the corresponding parameter or result subtype of the formal subprogram of *G* shall have a *null\_exclusion*;
- otherwise, the subtype of the corresponding parameter or result type of the renamed callable entity shall exclude null. In addition to the places where Legality Rules normally apply (see 15.3), this rule applies also in the private part of an instance of a generic unit.

The profile of a renaming-as-body shall conform fully to that of the declaration it completes. If the renaming-as-body completes that declaration before the subprogram it declares is frozen, the profile shall be mode conformant with that of the renamed callable entity and the subprogram it declares takes its convention from the renamed subprogram; otherwise, the profile shall be subtype conformant with that of the renamed callable entity and the convention of the renamed subprogram shall not be Intrinsic. A renaming-as-body is illegal if the declaration occurs before the subprogram whose declaration it completes is frozen, and the renaming renames the subprogram itself, through one or more subprogram renaming declarations, none of whose subprograms has been frozen.

The *callable\_entity\_name* of a renaming shall not denote a subprogram that requires overriding (see 6.9.3).

The *callable\_entity\_name* of a renaming-as-body shall not denote an abstract subprogram.

If the *callable\_entity\_name* of a renaming is a prefixed view, the prefix of that view shall denote an object for which renaming is allowed.

A name that denotes a formal parameter of the *subprogram\_specification* is not allowed within the *callable\_entity\_name*.

*Static Semantics*

A renaming-as-declaration declares a new view of the renamed entity. The profile of this new view takes its subtypes, parameter modes, and calling convention from the original profile of the callable entity while taking the formal parameter names and *default\_expressions* from the profile given in the *subprogram\_renaming\_declaration*. The new view is a function or procedure, never an entry.

*Dynamic Semantics*

For a call to a subprogram whose body is given as a renaming-as-body, the execution of the renaming-as-body is equivalent to the execution of a *subprogram\_body* that simply calls the renamed subprogram with its formal parameters as the actual parameters and, if it is a function, returns the value of the call.

For a call on a renaming of a dispatching subprogram that is overridden, if the overriding occurred before the renaming, then the body executed is that of the overriding declaration, even if the overriding declaration is not visible at the place of the renaming; otherwise, the inherited or predefined subprogram is called. A corresponding rule applies to a call on a renaming of a predefined equality operator for an untagged record type.

*Bounded (Run-Time) Errors*

If a subprogram directly or indirectly renames itself, then it is a bounded error to call that subprogram. Possible consequences are that `Program_Error` or `Storage_Error` is raised, or that the call results in infinite recursion.

NOTE 1 A procedure can only be renamed as a procedure. A function whose `defining_designator` is either an identifier or an `operator_symbol` can be renamed with either an identifier or an `operator_symbol`; for renaming as an operator, the subprogram specification given in the `renaming_declaration` is subject to the rules given in 9.6 for operator declarations. Enumeration literals can be renamed as functions; similarly, `attribute_references` that denote functions (such as references to `Succ` and `Pred`) can be renamed as functions. An entry can only be renamed as a procedure; the new name is only allowed to appear in contexts that allow a procedure name. An entry of a family can be renamed, but an entry family cannot be renamed as a whole.

NOTE 2 The operators of the root numeric types cannot be renamed because the types in the profile are anonymous, so the corresponding specifications cannot be written; the same holds for certain attributes, such as `Pos`.

NOTE 3 The primitiveness of a renaming-as-declaration is determined by its profile, and by where it occurs, as for any declaration of (a view of) a subprogram; primitiveness is not determined by the renamed view. In order to perform a dispatching call, the subprogram name has to denote a primitive subprogram, not a nonprimitive renaming of a primitive subprogram.

EXAMPLE 1 *Examples of subprogram renaming declarations:*

```
procedure My_Write(C : in Character) renames Pool(K).Write; -- see 7.1.3
function Real_Plus(Left, Right : Real ) return Real renames "+";
function Int_Plus (Left, Right : Integer) return Integer renames "+";

function Rouge return Color renames Red; -- see 6.5.1
function Rot return Color renames Red;
function Rosso return Color renames Rouge;

function Next(X : Color) return Color renames Color'Succ; -- see 6.5.1
```

EXAMPLE 2 *Example of a subprogram renaming declaration with new parameter names:*

```
function "*" (X,Y : Vector) return Real renames Dot_Product; -- see 9.1
```

EXAMPLE 3 *Example of a subprogram renaming declaration with a new default expression:*

```
function Minimum(L : Link := Head) return Cell renames Min_Cell; -- see 9.1
```

### 11.5.5 Generic Renaming Declarations

A `generic_renaming_declaration` is used to rename a generic unit.

*Syntax*

```
generic_renaming_declaration ::=
  generic_package      defining_program_unit_name renames generic_package_name
    [aspect_specification];
| generic_procedure    defining_program_unit_name renames generic_procedure_name
    [aspect_specification];
| generic_function     defining_program_unit_name renames generic_function_name
    [aspect_specification];
```

*Legality Rules*

The renamed entity shall be a generic unit of the corresponding kind.

*Static Semantics*

A `generic_renaming_declaration` declares a new view of the renamed generic unit.

NOTE Although the properties of the new view are the same as those of the renamed view, the place where the `generic_renaming_declaration` occurs can affect the legality of subsequent renamings and instantiations

that denote the `generic_renaming_declaration`, in particular if the renamed generic unit is a library unit (see 13.1.1).

EXAMPLE *Example of renaming a generic unit:*

```
generic package Enum_IO renames Ada.Text_IO Enumeration_IO;  -- see A.10.10
```

## 11.6 The Context of Overload Resolution

Because declarations can be overloaded, it is possible for an occurrence of a usage name to have more than one possible interpretation; in most cases, ambiguity is disallowed. This subclause describes how the possible interpretations resolve to the actual interpretation.

Certain rules of the language (the Name Resolution Rules) are considered “overloading rules”. If a possible interpretation violates an overloading rule, it is assumed not to be the intended interpretation; some other possible interpretation is assumed to be the actual interpretation. On the other hand, violations of nonoverloading rules do not affect which interpretation is chosen; instead, they cause the construct to be illegal. To be legal, there usually has to be exactly one acceptable interpretation of a construct that is a “complete context”, not counting any nested complete contexts.

The syntax rules of the language and the visibility rules given in 11.3 determine the possible interpretations. Most type checking rules (rules that require a particular type, or a particular class of types, for example) are overloading rules. Various rules for the matching of formal and actual parameters are overloading rules.

### *Name Resolution Rules*

Overload resolution is applied separately to each *complete context*, not counting inner complete contexts. Each of the following constructs is a *complete context*:

- A `context_item`.
- A `declarative_item` or declaration.
- A `statement`.
- A `pragma_argument_association`.
- The *selecting\_expression* of a `case_statement` or `case_expression`.
- The *variable\_name* of an `assignment_statement` *A*, if the expression of *A* contains one or more `target_names`.

An (overall) *interpretation* of a complete context embodies its meaning, and includes the following information about the constituents of the complete context, not including constituents of inner complete contexts:

- for each constituent of the complete context, to which syntactic categories it belongs, and by which syntax rules; and
- for each usage name, which declaration it denotes (and, therefore, which view and which entity it denotes); and
- for a complete context that is a `declarative_item`, whether or not it is a completion of a declaration, and (if so) which declaration it completes.

A *possible interpretation* is one that obeys the syntax rules and the visibility rules. An *acceptable interpretation* is a possible interpretation that obeys the *overloading rules*, that is, those rules that specify an expected type or expected profile, or specify how a construct shall *resolve* or be *interpreted*.

The *interpretation* of a constituent of a complete context is determined from the overall interpretation of the complete context as a whole. Thus, for example, “interpreted as a `function_call`”, means that the construct's interpretation says that it belongs to the syntactic category `function_call`.

Each occurrence of a usage name *denotes* the declaration determined by its interpretation. It also denotes the view declared by its denoted declaration, except in the following cases:

- If a usage name appears within the declarative region of a `type_declaration` and denotes that same `type_declaration`, then it denotes the *current instance* of the type (rather than the type itself); the current instance of a type is the object or value of the type that is associated with the execution that evaluates the usage name. Similarly, if a usage name appears within the declarative region of a `subtype_declaration` and denotes that same `subtype_declaration`, then it denotes the current instance of the subtype. These rules do not apply if the usage name appears within the `subtype_mark` of an `access_definition` for an access-to-object type, or within the subtype of a parameter or result of an access-to-subprogram type.

Within an `aspect_specification` for a type or subtype, the current instance represents a value of the type; it is not an object. Unless otherwise specified, the nominal subtype of this value is given by the subtype itself (the first subtype in the case of a `type_declaration`), prior to applying any predicate specified directly on the type or subtype. If the type or subtype is by-reference, the associated object of the value is the object associated (see 9.2) with the evaluation of the usage name.

- If a usage name appears within the declarative region of a `generic_declaration` (but not within its `generic_formal_part`) and it denotes that same `generic_declaration`, then it denotes the *current instance* of the generic unit (rather than the generic unit itself). See also 15.3.

A usage name that denotes a view also denotes the entity of that view.

The *expected type* for a given `expression`, `name`, or other construct determines, according to the *type resolution rules* given below, the types considered for the construct during overload resolution. The type resolution rules provide support for class-wide programming, universal literals, dispatching operations, and anonymous access types:

- If a construct is expected to be of any type in a class of types, or of the universal or class-wide type for a class, then the type of the construct shall resolve to a type in that class or to a universal type that covers the class.
- If the expected type for a construct is a specific type *T*, then the type of the construct shall resolve either to *T*, or:
  - to *T*'Class; or
  - to a universal type that covers *T*; or
  - when *T* is a specific anonymous access-to-object type (see 6.10) with designated type *D*, to an access-to-object type whose designated type is *D*'Class or is covered by *D*; or
  - when *T* is a named general access-to-object type (see 6.10) with designated type *D*, to an anonymous access-to-object type whose designated type covers or is covered by *D*; or
  - when *T* is an anonymous access-to-subprogram type (see 6.10), to an access-to-subprogram type whose designated profile is type conformant with that of *T*.

In certain contexts, such as in a `subprogram_renaming_declaration`, the Name Resolution Rules define an *expected profile* for a given name; in such cases, the name shall resolve to the name of a callable entity whose profile is type conformant with the expected profile.

#### Legality Rules

When a construct is one that requires that its expected type be a *single* type in a given class, the type of the construct shall be determinable solely from the context in which the construct appears, excluding the construct itself, but using the requirement that it be in the given class. Furthermore, the context shall not be one that expects any type in some class that contains types of the given class; in particular, the construct shall not be the operand of a `type_conversion`.

Other than for the *tested\_simple\_expression* of a membership test, if the expected type for a `name` or `expression` is not the same as the actual type of the `name` or `expression`, the actual type shall be convertible to the expected type (see 7.6); further, if the expected type is a named access-to-object type with designated type *D1* and the actual type is an anonymous access-to-object type with



designated type *D2*, then *D1* shall cover *D2*, and the `name` or `expression` shall denote a view with an accessibility level for which the statically deeper relationship applies; in particular it shall not denote an access parameter nor a stand-alone access object.

A complete context shall have at least one acceptable interpretation; if there is exactly one, then that one is chosen.

There is a *preference* for the primitive operators (and *ranges*) of the root numeric types *root\_integer* and *root\_real*. In particular, if two acceptable interpretations of a constituent of a complete context differ only in that one is for a primitive operator (or *range*) of the type *root\_integer* or *root\_real*, and the other is not, the interpretation using the primitive operator (or *range*) of the root numeric type is *preferred*.

Similarly, there is a preference for the equality operators of the *universal\_access* type (see 7.5.2). If two acceptable interpretations of a constituent of a complete context differ only in that one is for an equality operator of the *universal\_access* type, and the other is not, the interpretation using the equality operator of the *universal\_access* type is preferred.

For a complete context, if there is exactly one overall acceptable interpretation where each constituent's interpretation is the same as or preferred (in the above sense) over those in all other overall acceptable interpretations, then that one overall acceptable interpretation is chosen. Otherwise, the complete context is *ambiguous*.

A complete context other than a `pragma_argument_association` shall not be ambiguous.

A complete context that is a `pragma_argument_association` is allowed to be ambiguous (unless otherwise specified for the particular `pragma`), but only if every acceptable interpretation of the `pragma` argument is as a `name` that statically denotes a callable entity. Such a `name` denotes all of the declarations determined by its interpretations, and all of the views declared by these declarations.

NOTE If a usage name has only one acceptable interpretation, then it denotes the corresponding entity. However, this does not mean that the usage name is necessarily legal since other requirements exist which are not considered for overload resolution; for example, the fact that an expression is static, whether an object is constant, mode and subtype conformance rules, freezing rules, order of elaboration, and so on.

Similarly, subtypes are not considered for overload resolution (the violation of a constraint does not make a program illegal but raises an exception during program execution).

## 12 Tasks and Synchronization

The execution of an Ada program consists of the execution of one or more *tasks*. Each task represents a separable activity that proceeds independently and concurrently between the points where it *interacts* with other tasks. A single task, when within the context of a parallel construct, can represent multiple *logical threads of control* which can proceed in parallel; in other contexts, each task represents one logical thread of control.

The various forms of task interaction are described in this clause, and include:

- the activation and termination of a task;
- a call on a protected subprogram of a *protected object*, providing exclusive read-write access, or concurrent read-only access to shared data;
- a call on an entry, either of another task, allowing for synchronous communication with that task, or of a protected object, allowing for asynchronous communication with one or more other tasks using that same protected object;
- a timed operation, including a simple delay statement, a timed entry call or accept, or a timed asynchronous select statement (see next item);
- an asynchronous transfer of control as part of an asynchronous select statement, where a task stops what it is doing and begins execution at a different point in response to the completion of an entry call or the expiration of a delay;



— an abort statement, allowing one task to cause the termination of another task.

In addition, tasks can communicate indirectly by reading and updating (unprotected) shared variables, presuming the access is properly synchronized through some other kind of task interaction.

#### Static Semantics

The properties of a task are defined by a corresponding task declaration and `task_body`, which together define a program unit called a *task unit*.

#### Dynamic Semantics

Over time, tasks proceed through various *states*. A task is initially *inactive*; upon activation, and prior to its *termination* it is either *blocked* (as part of some task interaction) or *ready* to run. While ready, a task competes for the available *execution resources* that it requires to run. In the context of a parallel construct, a single task can utilize multiple processing resources simultaneously.

NOTE Concurrent task execution can be implemented on multicomputers, multiprocessors, or with interleaved execution on a single physical processor. On the other hand, whenever an implementation can determine that the required semantic effects can be achieved when parts of the execution of a single logical thread of control are performed by different physical processors acting in parallel, it can choose to perform them in this way.

## 12.1 Task Units and Task Objects

A task unit is declared by a *task declaration*, which has a corresponding `task_body`. A task declaration may be a *task\_type\_declaration*, in which case it declares a named task type; alternatively, it may be a *single\_task\_declaration*, in which case it defines an anonymous task type, as well as declaring a named task object of that type.

#### Syntax

```
task_type_declaration ::=
  task type defining_identifier [known_discriminant_part]
    [aspect_specification] [is
      [new interface_list with]
      task_definition];

single_task_declaration ::=
  task defining_identifier
    [aspect_specification] [is
      [new interface_list with]
      task_definition];

task_definition ::=
  {task_item}
  [private
    {task_item}]
  end [task_identifier]

task_item ::= entry_declaration | aspect_clause

task_body ::=
  task body defining_identifier
    [aspect_specification] is
    declarative_part
  begin
    handled_sequence_of_statements
  end [task_identifier];
```

If a *task\_identifier* appears at the end of a *task\_definition* or *task\_body*, it shall repeat the *defining\_identifier*.

*Static Semantics*

A `task_definition` defines a task type and its first subtype. The first list of `task_items` of a `task_definition`, together with the `known_discriminant_part`, if any, is called the visible part of the task unit. The optional list of `task_items` after the reserved word **private** is called the private part of the task unit.

For a task declaration without a `task_definition`, a `task_definition` without `task_items` is assumed.

For a task declaration with an `interface_list`, the task type inherits user-defined primitive subprograms from each progenitor type (see 6.9.4), in the same way that a derived type inherits user-defined primitive subprograms from its progenitor types (see 6.4). If the first parameter of a primitive inherited subprogram is of the task type or an access parameter designating the task type, and there is an `entry_declaration` for a single entry with the same identifier within the task declaration, whose profile is type conformant with the prefixed view profile of the inherited subprogram, the inherited subprogram is said to be *implemented* by the conforming task entry using an implicitly declared nonabstract subprogram which has the same profile as the inherited subprogram and which overrides it.

*Legality Rules*

A task declaration requires a completion, which shall be a `task_body`, and every `task_body` shall be the completion of some task declaration.

Each `interface_subtype_mark` of an `interface_list` appearing within a task declaration shall denote a limited interface type that is not a protected interface.

The prefixed view profile of an explicitly declared primitive subprogram of a tagged task type shall not be type conformant with any entry of the task type, if the subprogram has the same defining name as the entry and the first parameter of the subprogram is of the task type or is an access parameter designating the task type.

For each primitive subprogram inherited by the type declared by a task declaration, at most one of the following shall apply:

- the inherited subprogram is overridden with a primitive subprogram of the task type, in which case the overriding subprogram shall be subtype conformant with the inherited subprogram and not abstract; or
- the inherited subprogram is implemented by a single entry of the task type; in which case its prefixed view profile shall be subtype conformant with that of the task entry.

If neither applies, the inherited subprogram shall be a null procedure. In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

*Dynamic Semantics*

The elaboration of a task declaration elaborates the `task_definition`. The elaboration of a `single_task_declaration` also creates an object of an (anonymous) task type.

The elaboration of a `task_definition` creates the task type and its first subtype; it also includes the elaboration of the `entry_declarations` in the given order.

As part of the initialization of a task object, any `aspect_clauses` and any per-object constraints associated with `entry_declarations` of the corresponding `task_definition` are elaborated in the given order.

The elaboration of a `task_body` has no effect other than to establish that tasks of the type can from then on be activated without failing the `Elaboration_Check`.

The execution of a `task_body` is invoked by the activation of a task of the corresponding type (see 12.2).

The content of a task object of a given task type includes:

- The values of the discriminants of the task object, if any;
- An entry queue for each entry of the task object;
- A representation of the state of the associated task.

NOTE 1 Other than in an `access_definition`, the name of a task unit within the declaration or body of the task unit denotes the current instance of the unit (see 11.6), rather than the first subtype of the corresponding task type (and thus the name cannot be used as a `subtype_mark`).

NOTE 2 The notation of a `selected_component` can be used to denote a discriminant of a task (see 7.1.3). Within a task unit, the name of a discriminant of the task type denotes the corresponding discriminant of the current instance of the unit.

NOTE 3 A task type is a limited type (see 10.5), and hence precludes use of `assignment_statements` and predefined equality operators. If a programmer wants to write an application that stores and exchanges task identities, they can do so by defining an access type designating the corresponding task objects and by using access values for identification purposes. Assignment is available for such an access type as for any access type. Alternatively, if the implementation supports the Systems Programming Annex, the Identity attribute can be used for task identification (see C.7.1).

EXAMPLE 1 *Examples of declarations of task types:*

```
task type Server is
  entry Next_Work_Item(WI : in Work_Item);
  entry Shut_Down;
end Server;

task type Keyboard_Driver(ID : Keyboard_ID := New_ID) is
  new Serial_Device with -- see 6.9.4
  entry Read (C : out Character);
  entry Write(C : in Character);
end Keyboard_Driver;
```

EXAMPLE 2 *Examples of declarations of single tasks:*

```
task Controller is
  entry Request(Level)(D : Item); -- a family of entries
end Controller;

task Parser is
  entry Next_Lexeme(L : in Lexical_Element);
  entry Next_Action(A : out Parser_Action);
end;

task User; -- has no entries
```

EXAMPLE 3 *Examples of task objects:*

```
Agent      : Server;
Teletype   : Keyboard_Driver(TTY_ID);
Pool       : array(1 .. 10) of Keyboard_Driver;
```

EXAMPLE 4 *Example of access type designating task objects:*

```
type Keyboard is access Keyboard_Driver;
Terminal : Keyboard := new Keyboard_Driver(Term_ID);
```

## 12.2 Task Execution - Task Activation

### Dynamic Semantics

The execution of a task of a given task type consists of the execution of the corresponding `task_body`. The initial part of this execution is called the *activation* of the task; it consists of the elaboration of the `declarative_part` of the `task_body`. Should an exception be propagated by the elaboration of its `declarative_part`, the activation of the task is defined to have *failed*, and it becomes a completed task.

A task object (which represents one task) can be a part of a stand-alone object, of an object created by an `allocator`, or of an anonymous object of a limited type, or a coextension of one of these. All tasks that are part or coextensions of any of the stand-alone objects created by the elaboration of `object_declarations` (or `generic_associations` of formal objects of mode **in**) of a single declarative region are activated together. All tasks that are part or coextensions of a single object that is not a stand-alone object are activated together.

For the tasks of a given declarative region, the activations are initiated within the context of the `handled_sequence_of_statements` (and its associated `exception_handlers` if any — see 14.2), just prior to executing the statements of the `handled_sequence_of_statements`. For a package without an explicit body or an explicit `handled_sequence_of_statements`, an implicit body or an implicit `null_statement` is assumed, as defined in 10.2.

For tasks that are part or coextensions of a single object that is not a stand-alone object, activations are initiated after completing any initialization of the outermost object enclosing these tasks, prior to performing any other operation on the outermost object. In particular, for tasks that are part or coextensions of the object created by the evaluation of an `allocator`, the activations are initiated as the last step of evaluating the `allocator`, prior to returning the new access value. For tasks that are part or coextensions of an object that is the result of a function call, the activations are not initiated until after the function returns.

The task that created the new tasks and initiated their activations (the *activator*) is blocked until all of these activations complete (successfully or not). Once all of these activations are complete, if the activation of any of the tasks has failed (due to the propagation of an exception), `Tasking_Error` is raised in the activator, at the place at which it initiated the activations. Otherwise, the activator proceeds with its execution normally. Any tasks that are aborted prior to completing their activation are ignored when determining whether to raise `Tasking_Error`.

If the master that directly encloses the point where the activation of a task *T* would be initiated, completes before the activation of *T* is initiated, *T* becomes terminated and is never activated. Furthermore, if a return statement is left such that the return object is not returned to the caller, any task that was created as a part of the return object or one of its coextensions immediately becomes terminated and is never activated.

NOTE 1 An entry of a task can be called before the task has been activated.

NOTE 2 If several tasks are activated together, the execution of any of these tasks can proceed without waiting until the end of the activation of the other tasks.

NOTE 3 A task can become completed during its activation either because of an exception or because it is aborted (see 12.8).

EXAMPLE *Example of task activation:*

```

procedure P is
  A, B : Server;      -- elaborate the task objects A, B
  C    : Server;      -- elaborate the task object C
begin
  -- the tasks A, B, C are activated together before the first statement
  ...
end;
```

## 12.3 Task Dependence - Termination of Tasks

### *Dynamic Semantics*

Each task (other than an environment task — see 13.2) *depends* on one or more masters (see 10.6.1), as follows:

- If the task is created by the evaluation of an `allocator` for a given named access type, it depends on each master that includes the elaboration of the declaration of the ultimate ancestor of the given access type.

- If the task is created by the elaboration of an `object_declaration`, it depends on each master that includes this elaboration.
- Otherwise, the task depends on the master of the outermost object of which it is a part (as determined by the accessibility level of that object — see 6.10.2 and 10.6.1), as well as on any master whose execution includes that of the master of the outermost object.

Furthermore, if a task depends on a given master, it is defined to depend on the task that executes the master, and (recursively) on any master of that task.

A task is said to be *completed* when the execution of its corresponding `task_body` is completed. A task is said to be *terminated* when any finalization of the `task_body` has been performed (see 10.6.1). The first step of finalizing a master (including a `task_body`) is to wait for the termination of any tasks dependent on the master. The task executing the master is blocked until all the dependents have terminated. Any remaining finalization is then performed and the master is left.

Completion of a task (and the corresponding `task_body`) can occur when the task is blocked at a `select_statement` with an open `terminate_alternative` (see 12.7.1); the open `terminate_alternative` is selected if and only if the following conditions are satisfied:

- The task depends on some completed master; and
- Each task that depends on the master considered is either already terminated or similarly blocked at a `select_statement` with an open `terminate_alternative`.

When both conditions are satisfied, the task considered becomes completed, together with all tasks that depend on the master considered that are not yet completed.

NOTE 1 The full view of a limited private type can be a task type, or can have subcomponents of a task type. Creation of an object of such a type creates dependences according to the full type.

NOTE 2 An `object_renaming_declaration` defines a new view of an existing entity and hence creates no further dependence.

NOTE 3 The rules given for the collective completion of a group of tasks all blocked on `select_statements` with open `terminate_alternatives` ensure that the collective completion can occur only when there are no remaining active tasks that can call one of the tasks being collectively completed.

NOTE 4 If two or more tasks are blocked on `select_statements` with open `terminate_alternatives`, and become completed collectively, their finalization actions proceed concurrently.

NOTE 5 The completion of a task can occur due to any of the following:

- the raising of an exception during the elaboration of the `declarative_part` of the corresponding `task_body`;
- the completion of the `handled_sequence_of_statements` of the corresponding `task_body`;
- the selection of an open `terminate_alternative` of a `select_statement` in the corresponding `task_body`;
- the abort of the task.

EXAMPLE *Example of task dependence:*

```

declare
  type Global is access Server;           -- see 12.1
  A, B : Server;
  G    : Global;
begin
  -- activation of A and B
  declare
    type Local is access Server;
    X : Global := new Server; -- activation of X.all
    L : Local  := new Server; -- activation of L.all
    C : Server;
  begin
    -- activation of C
    G := X; -- both G and X designate the same task object
    ...
  end; -- await termination of C and L.all (but not X.all)
  ...
end; -- await termination of A, B, and G.all

```

## 12.4 Protected Units and Protected Objects

A *protected object* provides coordinated access to shared data, through calls on its visible *protected operations*, which can be *protected subprograms* or *protected entries*. A *protected unit* is declared by a *protected declaration*, which has a corresponding *protected\_body*. A protected declaration may be a *protected\_type\_declaration*, in which case it declares a named protected type; alternatively, it may be a *single\_protected\_declaration*, in which case it defines an anonymous protected type, as well as declaring a named protected object of that type.

*Syntax*

```

protected_type_declaration ::=
protected type defining_identifier [known_discriminant_part]
  [aspect_specification] is
  [new interface_list with]
  protected_definition;

single_protected_declaration ::=
protected defining_identifier
  [aspect_specification] is
  [new interface_list with]
  protected_definition;

protected_definition ::=
  {protected_operation_declaration}
  [private
    {protected_element_declaration}]
  end [protected_identifier]

protected_operation_declaration ::= subprogram_declaration
  | entry_declaration
  | aspect_clause

protected_element_declaration ::= protected_operation_declaration
  | component_declaration

protected_body ::=
protected body defining_identifier
  [aspect_specification] is
  {protected_operation_item}
  end [protected_identifier];

protected_operation_item ::= subprogram_declaration
  | subprogram_body

```



```

| null_procedure_declaration
| expression_function_declaration
| entry_body
| aspect_clause

```

If a *protected\_identifier* appears at the end of a *protected\_definition* or *protected\_body*, it shall repeat the *defining\_identifier*.

#### Static Semantics

A *protected\_definition* defines a *protected type* and its first subtype. The list of *protected\_operation\_declarations* of a *protected\_definition*, together with the *known\_discriminant\_part*, if any, is called the *visible part* of the *protected unit*. The optional list of *protected\_element\_declarations* after the reserved word **private** is called the *private part* of the *protected unit*.

For a *protected declaration* with an *interface\_list*, the *protected type* inherits user-defined primitive subprograms from each progenitor type (see 6.9.4), in the same way that a derived type inherits user-defined primitive subprograms from its progenitor types (see 6.4). If the first parameter of a primitive inherited subprogram is of the *protected type* or an access parameter designating the *protected type*, and there is a *protected\_operation\_declaration* for a *protected subprogram* or single entry with the same identifier within the *protected declaration*, whose profile is type conformant with the prefixed view profile of the inherited subprogram, the inherited subprogram is said to be *implemented* by the conforming *protected subprogram* or entry using an implicitly declared nonabstract subprogram which has the same profile as the inherited subprogram and which overrides it.

#### Legality Rules

A *protected declaration* requires a completion, which shall be a *protected\_body*, and every *protected\_body* shall be the completion of some *protected declaration*.

Each *interface\_subtype\_mark* of an *interface\_list* appearing within a *protected declaration* shall denote a limited interface type that is not a task interface.

The prefixed view profile of an explicitly declared primitive subprogram of a tagged *protected type* shall not be type conformant with any *protected operation* of the *protected type*, if the subprogram has the same defining name as the *protected operation* and the first parameter of the subprogram is of the *protected type* or is an access parameter designating the *protected type*.

For each primitive subprogram inherited by the type declared by a *protected declaration*, at most one of the following shall apply:

- the inherited subprogram is overridden with a primitive subprogram of the *protected type*, in which case the overriding subprogram shall be subtype conformant with the inherited subprogram and not abstract; or
- the inherited subprogram is implemented by a *protected subprogram* or single entry of the *protected type*, in which case its prefixed view profile shall be subtype conformant with that of the *protected subprogram* or entry.

If neither applies, the inherited subprogram shall be a null procedure. In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

If an inherited subprogram is implemented by a *protected procedure* or an entry, then the first parameter of the inherited subprogram shall be of mode **out** or **in out**, or an access-to-variable parameter. If an inherited subprogram is implemented by a *protected function*, then the first parameter of the inherited subprogram shall be of mode **in**, but not an access-to-variable parameter.

If a *protected subprogram declaration* has an *overriding\_indicator*, then at the point of the declaration:

- if the `overriding_indicator` is **overriding**, then the subprogram shall implement an inherited subprogram;
- if the `overriding_indicator` is **not overriding**, then the subprogram shall not implement any inherited subprogram.

In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

#### Dynamic Semantics

The elaboration of a `protected_declaration` elaborates the `protected_definition`. The elaboration of a `single_protected_declaration` also creates an object of an (anonymous) `protected` type.

The elaboration of a `protected_definition` creates the `protected` type and its first subtype; it also includes the elaboration of the `component_declarations` and `protected_operation_declarations` in the given order.

As part of the initialization of a `protected` object, any per-object constraints (see 6.8) are elaborated.

The elaboration of a `protected_body` has no other effect than to establish that `protected` operations of the type can from then on be called without failing the `Elaboration_Check`.

The content of an object of a given `protected` type includes:

- The values of the components of the `protected` object, including (implicitly) an entry queue for each entry declared for the `protected` object;
- A representation of the state of the execution resource *associated* with the `protected` object (one such resource is associated with each `protected` object).

The execution resource associated with a `protected` object has to be acquired to read or update any components of the `protected` object; it can be acquired (as part of a `protected` action — see 12.5.1) either for concurrent read-only access, or for exclusive read-write access.

As the first step of the *finalization* of a `protected` object, each call remaining on any entry queue of the object is removed from its queue and `Program_Error` is raised at the place of the corresponding `entry_call_statement`.

#### Bounded (Run-Time) Errors

It is a bounded error to call an entry or subprogram of a `protected` object after that object is finalized. If the error is detected, `Program_Error` is raised. Otherwise, the call proceeds normally, which may leave a task queued forever.

NOTE 1 Within the declaration or body of a `protected` unit other than in an `access_definition`, the name of the `protected` unit denotes the current instance of the unit (see 11.6), rather than the first subtype of the corresponding `protected` type (and thus the name cannot be used as a `subtype_mark`).

NOTE 2 A `selected_component` can be used to denote a discriminant of a `protected` object (see 7.1.3). Within a `protected` unit, the name of a discriminant of the `protected` type denotes the corresponding discriminant of the current instance of the unit.

NOTE 3 A `protected` type is a limited type (see 10.5), and hence precludes use of `assignment_statements` and predefined equality operators.

NOTE 4 The bodies of the `protected` operations given in the `protected_body` define the actions that take place upon calls to the `protected` operations.

NOTE 5 The declarations in the private part are only visible within the private part and the body of the `protected` unit.

EXAMPLE 1 *Example of declaration of protected type and corresponding body:*

```
protected type Resource is
  entry Seize;
  procedure Release;
private
  Busy : Boolean := False;
end Resource;

protected body Resource is
  entry Seize when not Busy is
  begin
    Busy := True;
  end Seize;

  procedure Release is
  begin
    Busy := False;
  end Release;
end Resource;
```

EXAMPLE 2 *Example of a single protected declaration and corresponding body:*

```
protected Shared_Array is
  -- Index, Item, and Item_Array are global types
  function Component (N : in Index) return Item;
  procedure Set_Component (N : in Index; E : in Item);
private
  Table : Item_Array(Index) := (others => Null_Item);
end Shared_Array;

protected body Shared_Array is
  function Component (N : in Index) return Item is
  begin
    return Table(N);
  end Component;

  procedure Set_Component (N : in Index; E : in Item) is
  begin
    Table(N) := E;
  end Set_Component;
end Shared_Array;
```

EXAMPLE 3 *Examples of protected objects:*

```
Control : Resource;
Flags : array(1..100) of Resource;
```

## 12.5 Intertask Communication

The primary means for intertask communication is provided by calls on entries and protected subprograms. Calls on protected subprograms allow coordinated access to shared data objects. Entry calls allow for blocking the caller until a given condition is satisfied (namely, that the corresponding entry is open — see 12.5.3), and then communicating data or control information directly with another task or indirectly via a shared protected object.

### Static Semantics

When a name or prefix denotes an entry, protected subprogram, or a prefixed view of a primitive subprogram of a limited interface whose first parameter is a controlling parameter, the name or prefix determines a *target object*, as follows:

- If it is a `direct_name` or `expanded name` that denotes the declaration (or body) of the operation, then the target object is implicitly specified to be the current instance of the task or protected unit immediately enclosing the operation; a call using such a name is defined to be an *internal call*;

- If it is a `selected_component` that is not an expanded name, then the target object is explicitly specified to be the object denoted by the `prefix` of the name; a call using such a name is defined to be an *external call*;
- If the name or `prefix` is a dereference (implicit or explicit) of an access-to-protected-subprogram value, then the target object is determined by the `prefix` of the Access attribute `reference` that produced the access value originally; a call using such a name is defined to be an *external call*;
- If the name or `prefix` denotes a `subprogram_renaming_declaration`, then the target object is as determined by the name of the renamed entity.

A call on an entry or a protected subprogram either uses a name or `prefix` that determines a target object implicitly, as above, or is a call on (a non-prefixed view of) a primitive subprogram of a limited interface whose first parameter is a controlling parameter, in which case the target object is identified explicitly by the first parameter. This latter case is an *external call*.

A corresponding definition of target object applies to a `requeue_statement` (see 12.5.4), with a corresponding distinction between an *internal requeue* and an *external requeue*.

#### Legality Rules

If a name or `prefix` determines a target object, and the name denotes a protected entry or procedure, then the target object shall be a variable, unless the `prefix` is for an attribute `reference` to the Count attribute (see 12.9).

An internal call on a protected function shall not occur within a precondition expression (see 9.1.1) of a protected operation nor within a `default_expression` of a `parameter_specification` of a protected operation.

#### Dynamic Semantics

Within the body of a protected operation, the current instance (see 11.6) of the immediately enclosing protected unit is determined by the target object specified (implicitly or explicitly) in the call (or `requeue`) on the protected operation.

Any call on a protected procedure or entry of a target protected object is defined to be an update to the object, as is a `requeue` on such an entry.

#### Syntax

`synchronization_kind ::= By_Entry | By_Protected_Procedure | Optional`

#### Static Semantics

For the declaration of a primitive procedure of a synchronized tagged type the following language-defined representation aspect may be specified with an `aspect_specification` (see 16.1.1):

#### Synchronization

If specified, the aspect definition shall be a `synchronization_kind`.

Inherited subprograms inherit the Synchronization aspect, if any, from the corresponding subprogram of the parent or progenitor type. If an overriding operation does not have a directly specified Synchronization aspect then the Synchronization aspect of the inherited operation is inherited by the overriding operation.

#### Legality Rules

The `synchronization_kind By_Protected_Procedure` shall not be applied to a primitive procedure of a task interface.

A procedure for which the specified `synchronization_kind` is `By_Entry` shall be implemented by an entry. A procedure for which the specified `synchronization_kind` is `By_Protected_Procedure` shall be implemented by a protected procedure. A procedure for which the specified

`synchronization_kind` is Optional may be implemented by an entry or by a procedure (including a protected procedure).

If a primitive procedure overrides an inherited operation for which the Synchronization aspect has been specified to be `By_Entry` or `By_Protected_Procedure`, then any specification of the aspect Synchronization applied to the overriding operation shall have the same `synchronization_kind`.

In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

#### Static Semantics

For a program unit, task entry, formal package, formal subprogram, formal object of an anonymous access-to-subprogram type, enumeration literal, and for a subtype (including a formal subtype), the following language-defined operational aspect is defined:

#### Nonblocking

This aspect specifies the blocking restriction for the entity; it shall be specified by a static Boolean expression. The `aspect_definition` can be omitted from the specification of this aspect; in that case, the aspect for the entity is `True`.

The Nonblocking aspect may be specified for all entities for which it is defined, except for protected operations and task entries. In particular, Nonblocking may be specified for generic formal parameters.

When aspect Nonblocking is `False` for an entity, the entity can contain a potentially blocking operation; such an entity *allows blocking*. If the aspect is `True` for an entity, the entity is said to be *nonblocking*.

For a generic instantiation and entities declared within such an instance, the aspect is determined by the Nonblocking aspect for the corresponding entity of the generic unit, **and** with the Nonblocking aspects of the actual generic parameters *used* by the entity. If the aspect is directly specified for an instance, the specified expression shall have the same value as the Nonblocking aspect of the instance (after **and**ing with the aspects of the used actual parameters). In the absence of a `Use_Formal` aspect, all actual generic parameters are presumed to be *used* by an entity (see H.7.1).

For a (protected or task) entry, the Nonblocking aspect is `False`.

For an enumeration literal, the Nonblocking aspect is `True`.

For a predefined operator of an elementary type, the Nonblocking aspect is `True`. For a predefined operator of a composite type, the Nonblocking aspect of the operator is the same as the Nonblocking aspect for the type.

For a dereference of an access-to-subprogram type, the Nonblocking aspect of the designated subprogram is that of the access-to-subprogram type.

For the base subtype of a scalar (sub)type, the Nonblocking aspect is `True`.

For an inherited primitive dispatching subprogram that is null or abstract, the subprogram is nonblocking if and only if a corresponding subprogram of at least one ancestor is nonblocking. For any other inherited subprogram, it is nonblocking if and only if the corresponding subprogram of the parent is nonblocking.

Unless directly specified, overriding of dispatching operations inherit this aspect.

Unless directly specified, for a formal subtype, formal package, or formal subprogram, the Nonblocking aspect is that of the actual subtype, package, or subprogram.

Unless directly specified, for a non-first subtype *S*, the Nonblocking aspect is that of the subtype identified in the `subtype_indication` defining *S*; unless directly specified for the first subtype of a derived type, the Nonblocking aspect is that of the ancestor subtype.

Unless directly specified, for any other program unit, first subtype, or formal object, the Nonblocking aspect of the entity is determined by the Nonblocking aspect for the innermost program unit enclosing the entity.

If not specified for a library unit, the Nonblocking aspect is True if the library unit is declared pure, or False otherwise.

The following are defined to be *potentially blocking* operations:

- a `select_statement`;
- an `accept_statement`;
- an `entry_call_statement`, or a call on a procedure that renames or is implemented by an entry;
- a `delay_statement`;
- an `abort_statement`;
- task creation or activation;
- during a protected action, an external call on a protected subprogram (or an external requeue) with the same target object as that of the protected action.

If a language-defined subprogram allows blocking, then a call on the subprogram is a potentially blocking operation.

#### Legality Rules

A portion of program text is called a *nonblocking region* if it is anywhere within a parallel construct, or if the innermost enclosing program unit is nonblocking. A nonblocking region shall not contain any of the following:

- a `select_statement`;
- an `accept_statement`;
- a `delay_statement`;
- an `abort_statement`;
- task creation or activation.

Furthermore, a parallel construct shall neither contain a call on a callable entity for which the Nonblocking aspect is False, nor shall it contain a call on a callable entity declared within a generic unit that uses a generic formal parameter with Nonblocking aspect False (see Use\_Formal aspect in H.7.1).

Finally, a nonblocking region that is outside of a parallel construct shall not contain a call on a callable entity for which the Nonblocking aspect is False, unless the region is within a generic unit and the callable entity is associated with a generic formal parameter of the generic unit, or the call is within the `aspect_definition` of an assertion aspect for an entity that allows blocking.

For the purposes of the above rules, an `entry_body` is considered nonblocking if the immediately enclosing protected unit is nonblocking.

For a subtype for which aspect Nonblocking is True, any predicate expression that applies to the subtype shall only contain constructs that are allowed immediately within a nonblocking program unit.

A subprogram shall be nonblocking if it overrides a nonblocking dispatching operation. An entry shall not implement a nonblocking procedure. If an inherited dispatching subprogram allows blocking, then the corresponding subprogram of each ancestor shall allow blocking.

It is illegal to directly specify aspect Nonblocking for the first subtype of the full view of a type that has a partial view. If the Nonblocking aspect of the full view is inherited, it shall have the same value as that of the partial view, or have the value True.



Aspect Nonblocking shall be directly specified for the first subtype of a derived type only if it has the same value as the Nonblocking aspect of the ancestor subtype or if it is specified True. Aspect Nonblocking shall be directly specified for a nonfirst subtype *S* only if it has the same value as the Nonblocking aspect of the subtype identified in the `subtype_indication` defining *S* or if it is specified True.

For an access-to-object type that is nonblocking, the Allocate, Deallocate, and Storage\_Size operations on its storage pool shall be nonblocking.

For a composite type that is nonblocking:

- All component subtypes shall be nonblocking;
- For a record type or extension, every call in the `default_expression` of a component (including discriminants) shall call an operation that is nonblocking;
- For a controlled type, the Initialize, Finalize, and Adjust (if any) subprograms shall be nonblocking.

The predefined equality operator for a composite type, unless it is for a record type or record extension and the operator is overridden by a primitive equality operator, is illegal if it is nonblocking and:

- for a record type or record extension, the parent primitive "=" allows blocking; or
- some component is of a type *T*, and:
  - *T* is a record type or record extension that has a primitive "=" that allows blocking; or
  - *T* is neither a record type nor a record extension, and *T* has a predefined "=" that allows blocking.

In a generic instantiation:

- the actual subprogram corresponding to a nonblocking formal subprogram shall be nonblocking (an actual that is an entry is not permitted in this case);
- the actual subtype corresponding to a nonblocking formal subtype shall be nonblocking;
- the actual object corresponding to a formal object of a nonblocking access-to-subprogram type shall be of a nonblocking access-to-subprogram type;
- the actual instance corresponding to a nonblocking formal package shall be nonblocking.

In addition to the places where Legality Rules normally apply (see 15.3), the above rules also apply in the private part of an instance of a generic unit.

NOTE The synchronization kind `By_Protected_Procedure` implies that the operation will not block.

### 12.5.1 Protected Subprograms and Protected Actions

A *protected subprogram* is a subprogram declared immediately within a `protected_definition`. Protected procedures provide exclusive read-write access to the data of a protected object; protected functions provide concurrent read-only access to the data.

#### Static Semantics

Within the body of a protected function (or a function declared immediately within a `protected_body`), the current instance of the enclosing protected unit is defined to be a constant (that is, its subcomponents may be read but not updated). Within the body of a protected procedure (or a procedure declared immediately within a `protected_body`), and within an `entry_body`, the current instance is defined to be a variable (updating is permitted).

For a type declared by a `protected_type_declaration` or for the anonymous type of an object declared by a `single_protected_declaration`, the following language-defined type-related representation aspect may be specified:

### Exclusive\_Functions

The type of aspect `Exclusive_Functions` is Boolean. If not specified (including by inheritance), the aspect is False.

A value of True for this aspect indicates that protected functions behave in the same way as protected procedures with respect to mutual exclusion and queue servicing (see below).

A protected procedure or entry is an *exclusive* protected operation. A protected function of a protected type *P* is an exclusive protected operation if the `Exclusive_Functions` aspect of *P* is True.

#### Dynamic Semantics

For the execution of a call on a protected subprogram, the evaluation of the `name` or `prefix` and of the parameter associations, and any assigning back of **in out** or **out** parameters, proceeds as for a normal subprogram call (see 9.4). If the call is an internal call (see 12.5), the body of the subprogram is executed as for a normal subprogram call. If the call is an external call, then the body of the subprogram is executed as part of a new *protected action* on the target protected object; the protected action completes after the body of the subprogram is executed. A protected action can also be started by an entry call (see 12.5.3).

A new protected action is not started on a protected object while another protected action on the same protected object is underway, unless both actions are the result of a call on a nonexclusive protected function. This rule is expressible in terms of the execution resource associated with the protected object:

- *Starting* a protected action on a protected object corresponds to *acquiring* the execution resource associated with the protected object, either for exclusive read-write access if the protected action is for a call on an exclusive protected operation, or for concurrent read-only access otherwise;
- *Completing* the protected action corresponds to *releasing* the associated execution resource.

After performing an exclusive protected operation on a protected object, but prior to completing the associated protected action, the entry queues (if any) of the protected object are serviced (see 12.5.3).

If a parallel construct occurs within a protected action, no new logical threads of control are created. Instead, each element of the parallel construct that would have become a separate logical thread of control executes on the logical thread of control that is performing the protected action. If there are multiple such elements initiated at the same point, they execute in an arbitrary order.

#### Bounded (Run-Time) Errors

During a protected action, it is a bounded error to invoke an operation that is potentially blocking (see 12.5).

If the bounded error is detected, `Program_Error` is raised. If not detected, the bounded error can result in deadlock or a (nested) protected action on the same target object.

During a protected action, a call on a subprogram whose body contains a potentially blocking operation is a bounded error. If the bounded error is detected, `Program_Error` is raised; otherwise, the call proceeds normally.

NOTE 1 If two tasks both try to start a protected action on a protected object, and at most one is calling a protected nonexclusive function, then only one of the tasks can proceed. Although the other task cannot proceed, it is not considered blocked, and it can be consuming processing resources while it awaits its turn. Unless there is an admission policy (see D.4.1) in effect, there is no language-defined ordering or queuing presumed for tasks competing to start a protected action — on a multiprocessor such tasks can use busy-waiting; for further monoprocessor and multiprocessor considerations, see D.3.

NOTE 2 The body of a protected unit can contain declarations and bodies for local subprograms. These are not visible outside the protected unit.

NOTE 3 The body of a protected function can contain internal calls on other protected functions, but not protected procedures, because the current instance is a constant. On the other hand, the body of a protected procedure can contain internal calls on both protected functions and procedures.

NOTE 4 From within a protected action, an internal call on a protected subprogram, or an external call on a protected subprogram with a different target object is not considered a potentially blocking operation.

NOTE 5 The aspect Nonblocking can be specified True on the definition of a protected unit in order to reject most attempts to use potentially blocking operations within the protected unit (see 12.5). The pragma Detect\_Blocking can be used to ensure that any remaining executions of potentially blocking operations during a protected action raise Program\_Error. See H.5.

EXAMPLE *Examples of protected subprogram calls (see 12.4):*

```
Shared_Array.Set_Component(N, E);
E := Shared_Array.Component(M);
Control.Release;
```

## 12.5.2 Entries and Accept Statements

Entry\_declarations, with the corresponding entry\_bodies or accept\_statements, are used to define potentially queued operations on tasks and protected objects.

*Syntax*

```
entry_declaration ::=
  [overriding_indicator]
  entry defining_identifier [(discrete_subtype_definition)] parameter_profile
  [aspect_specification];

accept_statement ::=
  accept entry_direct_name [(entry_index)] parameter_profile [do
    handled_sequence_of_statements
  end [entry_identifier]];

entry_index ::= expression

entry_body ::=
  entry defining_identifier entry_body_formal_part
  [aspect_specification]
  entry_barrier is
  declarative_part
  begin
  handled_sequence_of_statements
  end [entry_identifier];

entry_body_formal_part ::= [(entry_index_specification)] parameter_profile

entry_barrier ::= when condition

entry_index_specification ::=
  for defining_identifier in discrete_subtype_definition
  [aspect_specification]
```

If an *entry* identifier appears at the end of an *accept\_statement*, it shall repeat the *entry*-*direct\_name*. If an *entry\_identifier* appears at the end of an *entry\_body*, it shall repeat the *defining\_identifier*.

An *entry\_declaration* is allowed only in a protected or task declaration.

An *overriding\_indicator* is not allowed in an *entry\_declaration* that includes a *discrete\_subtype\_definition*.

*Name Resolution Rules*

In an *accept\_statement*, the expected profile for the *entry\_direct\_name* is that of the *entry\_declaration*; the expected type for an *entry\_index* is that of the subtype defined by the *discrete\_subtype\_definition* of the corresponding *entry\_declaration*.

Within the *handled\_sequence\_of\_statements* of an *accept\_statement*, if a selected\_component has a prefix that denotes the corresponding *entry\_declaration*, then the entity

denoted by the prefix is the `accept_statement`, and the `selected_component` is interpreted as an expanded name (see 7.1.3); the `selector_name` of the `selected_component` has to be the identifier for some formal parameter of the `accept_statement`.

#### Legality Rules

An `entry_declaration` in a task declaration shall not contain a specification for an access parameter (see 6.10).

If an `entry_declaration` has an `overriding_indicator`, then at the point of the declaration:

- if the `overriding_indicator` is **overriding**, then the entry shall implement an inherited subprogram;
- if the `overriding_indicator` is **not overriding**, then the entry shall not implement any inherited subprogram.

In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

For an `accept_statement`, the innermost enclosing body shall be a `task_body`, and the `entry_direct_name` shall denote an `entry_declaration` in the corresponding task declaration; the profile of the `accept_statement` shall conform fully to that of the corresponding `entry_declaration`. An `accept_statement` shall have a parenthesized `entry_index` if and only if the corresponding `entry_declaration` has a `discrete_subtype_definition`.

An `accept_statement` shall not be within another `accept_statement` that corresponds to the same `entry_declaration`, nor within an `asynchronous_select` inner to the enclosing `task_body`.

An `entry_declaration` of a protected unit requires a completion, which shall be an `entry_body`, and every `entry_body` shall be the completion of an `entry_declaration` of a protected unit. The profile of the `entry_body` shall conform fully to that of the corresponding declaration.

An `entry_body_formal_part` shall have an `entry_index_specification` if and only if the corresponding `entry_declaration` has a `discrete_subtype_definition`. In this case, the `discrete_subtype_definitions` of the `entry_declaration` and the `entry_index_specification` shall fully conform to one another (see 9.3.1).

A name that denotes a formal parameter of an `entry_body` is not allowed within the `entry_barrier` of the `entry_body`.

#### Static Semantics

The parameter modes defined for parameters in the `parameter_profile` of an `entry_declaration` are the same as for a `subprogram_declaration` and have the same meaning (see 9.2).

An `entry_declaration` with a `discrete_subtype_definition` (see 6.6) declares a *family* of distinct entries having the same profile, with one such entry for each value of the *entry index subtype* defined by the `discrete_subtype_definition`. A name for an entry of a family takes the form of an `indexed_component`, where the prefix denotes the `entry_declaration` for the family, and the index value identifies the entry within the family. The term *single entry* is used to refer to any entry other than an entry of an entry family.

In the `entry_body` for an entry family, the `entry_index_specification` declares a named constant whose subtype is the `entry index subtype` defined by the corresponding `entry_declaration`; the value of the *named entry index* identifies which entry of the family was called.

The elaboration of an `entry_declaration` for an entry family consists of the elaboration of the `discrete_subtype_definition`, as described in 6.8. The elaboration of an `entry_declaration` for a single entry has no effect.

The actions to be performed when an entry is called are specified by the corresponding `accept_statements` (if any) for an entry of a task unit, and by the corresponding `entry_body` for an entry of a protected unit.

The interaction between a task that calls an entry and an accepting task is called a *rendezvous*.

For the execution of an `accept_statement`, the `entry_index`, if any, is first evaluated and converted to the entry index subtype; this index value identifies which entry of the family is to be accepted. Further execution of the `accept_statement` is then blocked until a caller of the corresponding entry is selected (see 12.5.3), whereupon the `handled_sequence_of_statements`, if any, of the `accept_statement` is executed, with the formal parameters associated with the corresponding actual parameters of the selected entry call. Execution of the rendezvous consists of the execution of the `handled_sequence_of_statements`, performance of any postcondition or type invariant checks associated with the entry, and any initialization or finalization associated with these checks, as described in 9.1.1 and 10.3.2. After execution of the rendezvous, the `accept_statement` completes and is left. The two tasks then proceed independently. When an exception is propagated from the `handled_sequence_of_statements` of an `accept_statement`, the same exception is also raised by the execution of the corresponding `entry_call_statement`.

An `entry_body` is executed when the condition of the `entry_barrier` evaluates to True and a caller of the corresponding single entry, or entry of the corresponding entry family, has been selected (see 12.5.3). For the execution of the `entry_body`, the `declarative_part` of the `entry_body` is elaborated, and the `handled_sequence_of_statements` of the body is executed, as for the execution of a `subprogram_body`. The value of the named entry index, if any, is determined by the value of the entry index specified in the `entry_name` of the selected entry call (or intermediate `requeue_statement` — see 12.5.4).

NOTE 1 A task entry has corresponding `accept_statements` (zero or more), whereas a protected entry has a corresponding `entry_body` (exactly one).

NOTE 2 A consequence of the rule regarding the allowed placements of `accept_statements` is that a task can execute `accept_statements` only for its own entries.

NOTE 3 A `return_statement` (see 9.5) or a `requeue_statement` (see 12.5.4) can be used to complete the execution of an `accept_statement` or an `entry_body`.

NOTE 4 The condition in the `entry_barrier` can reference anything visible except the formal parameters of the entry. This includes the entry index (if any), the components (including discriminants) of the protected object, the Count attribute of an entry of that protected object, and data global to the protected unit.

The restriction against referencing the formal parameters within an `entry_barrier` ensures that all calls of the same entry see the same barrier value. If it is necessary to look at the parameters of an entry call before deciding whether to handle it, the `entry_barrier` can be “**when** True” and the caller can be requeued (on some private entry) when its parameters indicate that it cannot be handled immediately.

EXAMPLE 1 *Examples of entry declarations:*

```
entry Read(V : out Item);
entry Seize;
entry Request(Level) (D : Item); -- a family of entries
```

EXAMPLE 2 *Examples of accept statements:*

```
accept Shut_Down;

accept Read(V : out Item) do
  V := Local_Item;
end Read;
```

```

accept Request(Low) (D : Item) do
    ...
end Request;

```

### 12.5.3 Entry Calls

An `entry_call_statement` (an *entry call*) can appear in various contexts. A *simple* entry call is a stand-alone statement that represents an unconditional call on an entry of a target task or a protected object. Entry calls can also appear as part of `select_statements` (see 12.7).

#### Syntax

```
entry_call_statement ::= entry_name[actual_parameter_part];
```

#### Name Resolution Rules

The `entry_name` given in an `entry_call_statement` shall resolve to denote an entry. The rules for parameter associations are the same as for subprogram calls (see 9.4 and 9.4.1).

#### Static Semantics

The `entry_name` of an `entry_call_statement` specifies (explicitly or implicitly) the target object of the call, the entry or entry family, and the entry index, if any (see 12.5).

#### Dynamic Semantics

Under certain circumstances (detailed below), an entry of a task or protected object is checked to see whether it is *open* or *closed*:

- An entry of a task is open if the task is blocked on an `accept_statement` that corresponds to the entry (see 12.5.2), or on a `selective_accept` (see 12.7.1) with an open `accept_alternative` that corresponds to the entry; otherwise, it is closed.
- An entry of a protected object is open if the `condition` of the `entry_barrier` of the corresponding `entry_body` evaluates to True; otherwise, it is closed. If the evaluation of the `condition` propagates an exception, the exception `Program_Error` is propagated to all current callers of all entries of the protected object.

For the execution of an `entry_call_statement`, evaluation of the name and of the parameter associations is as for a subprogram call (see 9.4). The entry call is then *issued*: For a call on an entry of a protected object, a new protected action is started on the object (see 12.5.1). The named entry is checked to see if it is open; if open, the entry call is said to be *selected immediately*, and the execution of the call proceeds as follows:

- For a call on an open entry of a task, the accepting task becomes ready and continues the execution of the corresponding `accept_statement` (see 12.5.2).
- For a call on an open entry of a protected object, the corresponding `entry_body` is executed (see 12.5.2) as part of the protected action.

If the `accept_statement` or `entry_body` completes other than by a `requeue` (see 12.5.4), return is made to the caller (after servicing the entry queues — see below); any necessary assigning back of formal to actual parameters occurs, as for a subprogram call (see 9.4.1); such assignments take place outside of any protected action.

If the named entry is closed, the entry call is added to an *entry queue* (as part of the protected action, for a call on a protected entry), and the call remains queued until it is selected or cancelled; there is a separate (logical) entry queue for each entry of a given task or protected object (including each entry of an entry family).

When a queued call is *selected*, it is removed from its entry queue. Selecting a queued call from a particular entry queue is called *servicing* the entry queue. An entry with queued calls can be serviced under the following circumstances:



- When the associated task reaches a corresponding `accept_statement`, or a `selective_accept` with a corresponding open `accept_alternative`;
- If after performing, as part of a protected action on the associated protected object, an exclusive protected operation on the object, the entry is checked and found to be open.

If there is at least one call on a queue corresponding to an open entry, then one such call is selected according to the *entry queuing policy* in effect (see below), and the corresponding `accept_statement` or `entry_body` is executed as above for an entry call that is selected immediately.

The entry queuing policy controls selection among queued calls both for task and protected entry queues. The default entry queuing policy is to select calls on a given entry queue in order of arrival. If calls from two or more queues are simultaneously eligible for selection, the default entry queuing policy does not specify which queue is serviced first. Other entry queuing policies can be specified by pragmas (see D.4).

For a protected object, the above servicing of entry queues continues until there are no open entries with queued calls, at which point the protected action completes.

For an entry call that is added to a queue, and that is not the `triggering_statement` of an `asynchronous_select` (see 12.7.4), the calling task is blocked until the call is cancelled, or the call is selected and a corresponding `accept_statement` or `entry_body` completes without requeuing. In addition, the calling task is blocked during a rendezvous.

An attempt can be made to cancel an entry call upon an abort (see 12.8) and as part of certain forms of `select_statement` (see 12.7.2, 12.7.3, and 12.7.4). The cancellation does not take place until a point (if any) when the call is on some entry queue, and not protected from cancellation as part of a requeue (see 12.5.4); at such a point, the call is removed from the entry queue and the call completes due to the cancellation. The cancellation of a call on an entry of a protected object is a protected action, and as such cannot take place while any other protected action is occurring on the protected object. Like any protected action, it includes servicing of the entry queues (in case some entry barrier depends on a Count attribute).

A call on an entry of a task that has already completed its execution raises the exception `Tasking_Error` at the point of the call; similarly, this exception is raised at the point of the call if the called task completes its execution or becomes abnormal before accepting the call or completing the rendezvous (see 12.8). This applies equally to a simple entry call and to an entry call as part of a `select_statement`.

#### Implementation Permissions

An implementation may perform the sequence of steps of a protected action using any thread of control; it can be a thread other than that of the task that started the protected action. If an `entry_body` completes without requeuing, then the corresponding calling task may be made ready without waiting for the entire protected action to complete.

When the entry of a protected object is checked to see whether it is open, the implementation can bypass reevaluating the condition of the corresponding `entry_barrier` if no variable or attribute referenced by the condition (directly or indirectly) has been altered by the execution (or cancellation) of a call to an exclusive protected operation of the object since the condition was last evaluated.

An implementation may evaluate the conditions of all `entry_barriers` of a given protected object any time any entry of the object is checked to see if it is open.

When an attempt is made to cancel an entry call, the implementation can use a thread of control other than that of the task (or interrupt) that initiated the cancellation; in particular, it may use the thread of control of the caller itself to attempt the cancellation, even if this can allow the entry call to be selected in the interim.

NOTE 1 If an exception is raised during the execution of an `entry_body`, it is propagated to the corresponding caller (see 14.4).

NOTE 2 For a call on a protected entry, the entry is checked to see if it is open prior to queuing the call, and again thereafter if its `Count` attribute (see 12.9) is referenced in some entry barrier.

NOTE 3 In addition to simple entry calls, the language permits timed, conditional, and asynchronous entry calls (see 12.7.2, 12.7.3, and see 12.7.4).

NOTE 4 The `condition` of an `entry_barrier` is allowed to be evaluated by an implementation more often than strictly necessary, even if the evaluation can have side effects. On the other hand, an implementation can avoid reevaluating the `condition` if nothing it references was updated by an intervening protected action on the protected object, even if the `condition` references some global variable that is updated by an action performed from outside of a protected action.

EXAMPLE *Examples of entry calls:*

|                                                  |             |
|--------------------------------------------------|-------------|
| <code>Agent.Shut_Down;</code>                    | -- see 12.1 |
| <code>Parser.Next_Lexeme(E);</code>              | -- see 12.1 |
| <code>Pool(5).Read(Next_Char);</code>            | -- see 12.1 |
| <code>Controller.Request(Low)(Some_Item);</code> | -- see 12.1 |
| <code>Flags(3).Seize;</code>                     | -- see 12.4 |

### 12.5.4 Requeue Statements

A `requeue_statement` can be used to complete an `accept_statement` or `entry_body`, while redirecting the corresponding entry call to a new (or the same) entry queue. Such a *requeue* can be performed with or without allowing an intermediate cancellation of the call, due to an abort or the expiration of a delay.

*Syntax*

`requeue_statement ::= requeue procedure_or_entry_name [with abort];`

*Name Resolution Rules*

The *procedure\_or\_entry\_name* of a `requeue_statement` shall resolve to denote a procedure or an entry (the *requeue target*). The profile of the entry, or the profile or prefixed profile of the procedure, shall either have no parameters, or be type conformant (see 9.3.1) with the profile of the innermost enclosing `entry_body` or `accept_statement`.

*Legality Rules*

A `requeue_statement` shall be within a callable construct that is either an `entry_body` or an `accept_statement`, and this construct shall be the innermost enclosing body or callable construct.

If the *requeue target* has parameters, then its (prefixed) profile shall be subtype conformant with the profile of the innermost enclosing callable construct.

Given a `requeue_statement` where the innermost enclosing callable construct is for an entry *E1*, for every specific or class-wide postcondition expression *P1* that applies to *E1*, there shall exist a postcondition expression *P2* that applies to the *requeue target* *E2* such that

- *P1* is fully conformant with the expression produced by replacing each reference in *P2* to a formal parameter of *E2* with a reference to the corresponding formal parameter of *E1*; and
- if *P1* is enabled, then *P2* is also enabled.

The *requeue target* shall not have an applicable specific or class-wide postcondition that includes an `Old` or `Index attribute_reference`.

If the *requeue target* is declared immediately within the `task_definition` of a named task type or the `protected_definition` of a named protected type, and if the *requeue statement* occurs within the body of that type, and if the *requeue* is an external *requeue*, then the *requeue target* shall not have a specific or class-wide postcondition which includes a name denoting either the current instance of that type or any entity declared within the declaration of that type.

If the target is a procedure, the name shall denote a renaming of an entry, or shall denote a view or a prefixed view of a primitive subprogram of a synchronized interface, where the first parameter of the unprefix view of the primitive subprogram shall be a controlling parameter, and the Synchronization aspect shall be specified with `synchronization_kind` `By_Entry` for the primitive subprogram.

In a `requeue_statement` of an `accept_statement` of some task unit, either the target object shall be a part of a formal parameter of the `accept_statement`, or the accessibility level of the target object shall not be equal to or statically deeper than any enclosing `accept_statement` of the task unit. In a `requeue_statement` of an `entry_body` of some protected unit, either the target object shall be a part of a formal parameter of the `entry_body`, or the accessibility level of the target object shall not be statically deeper than that of the `entry_declaration` for the `entry_body`.

#### Dynamic Semantics

The execution of a `requeue_statement` begins with the following sequence of steps:

1. The *procedure\_or\_entry\_name* is evaluated. This includes evaluation of the *prefix* (if any) identifying the target task or protected object and of the *expression* (if any) identifying the entry within an entry family.
2. If the target object is not a part of a formal parameter of the innermost enclosing callable construct, a check is made that the accessibility level of the target object is not equal to or deeper than the level of the innermost enclosing callable construct. If this check fails, `Program_Error` is raised.
3. Precondition checks are performed as for a call to the requeue target.
4. The `entry_body` or `accept_statement` enclosing the `requeue_statement` is then completed, finalized, and left (see 10.6.1).

For the execution of a requeue on an entry of a target task, after leaving the enclosing callable construct, the named entry is checked to see if it is open and the requeued call is either selected immediately or queued, as for a normal entry call (see 12.5.3).

For the execution of a requeue on an entry of a target protected object, after leaving the enclosing callable construct:

- if the requeue is an internal requeue (that is, the requeue is back on an entry of the same protected object — see 12.5), the call is added to the queue of the named entry and the ongoing protected action continues (see 12.5.1);
- if the requeue is an external requeue (that is, the target protected object is not implicitly the same as the current object — see 12.5), a protected action is started on the target object and proceeds as for a normal entry call (see 12.5.3).

If the requeue target named in the `requeue_statement` has formal parameters, then during the execution of the `accept_statement` or `entry_body` corresponding to the new entry and during the checking of any preconditions of the new entry, the formal parameters denote the same objects as did the corresponding formal parameters of the callable construct completed by the requeue. In any case, no parameters are specified in a `requeue_statement`; any parameter passing is implicit.

If the `requeue_statement` includes the reserved words **with abort** (it is a *requeue-with-abort*), then:

- if the original entry call has been aborted (see 12.8), then the requeue acts as an abort completion point for the call, and the call is cancelled and no requeue is performed;
- if the original entry call was timed (or conditional), then the original expiration time is the expiration time for the requeued call.

If the reserved words **with abort** do not appear, then the call remains protected against cancellation while queued as the result of the `requeue_statement`.

NOTE A requeue is permitted from a single entry to an entry of an entry family, or vice versa. The entry index, if any, plays no part in the subtype conformance check between the profiles of the two entries; an entry index is part of the *entry\_name* for an entry of a family.

EXAMPLE *Examples of requeue statements:*

```
requeue Request(Medium) with abort;
-- requeue on a member of an entry family of the current
task, see 12.1
requeue Flags(I).Seize;
-- requeue on an entry of an array component, see 12.4
```

## 12.6 Delay Statements, Duration, and Time

A *delay\_statement* is used to block further execution until a specified *expiration time* is reached. The expiration time can be specified either as a particular point in time (in a *delay\_until\_statement*), or in seconds from the current time (in a *delay\_relative\_statement*). The language-defined package *Calendar* provides definitions for a type *Time* and associated operations, including a function *Clock* that returns the current time.

### Syntax

```
delay_statement ::= delay_until_statement | delay_relative_statement
delay_until_statement ::= delay until delay_expression;
delay_relative_statement ::= delay delay_expression;
```

### Name Resolution Rules

The expected type for the *delay\_expression* in a *delay\_relative\_statement* is the predefined type *Duration*. The *delay\_expression* in a *delay\_until\_statement* is expected to be of any unlimited type.

### Legality Rules

There can be multiple time bases, each with a corresponding clock, and a corresponding *time type*. The type of the *delay\_expression* in a *delay\_until\_statement* shall be a time type — either the type *Time* defined in the language-defined package *Calendar* (see below), the type *Time* in the package *Real\_Time* (see D.8), or some other implementation-defined time type.

### Static Semantics

There is a predefined fixed point type named *Duration*, declared in the visible part of package *Standard*; a value of type *Duration* is used to represent the length of an interval of time, expressed in seconds. The type *Duration* is not specific to a particular time base, but can be used with any time base.

A value of the type *Time* in package *Calendar*, or of some other time type, represents a time as reported by a corresponding clock.

The following language-defined library package exists:

```
package Ada.Calendar
with Nonblocking, Global => in out synchronized is
type Time is private;

subtype Year_Number is Integer range 1901 .. 2399;
subtype Month_Number is Integer range 1 .. 12;
subtype Day_Number is Integer range 1 .. 31;
subtype Day_Duration is Duration range 0.0 .. 86_400.0;

function Clock return Time;
```

```

function Year    (Date : Time) return Year_Number;
function Month   (Date : Time) return Month_Number;
function Day     (Date : Time) return Day_Number;
function Seconds (Date : Time) return Day_Duration;

procedure Split (Date : in Time;
                 Year  : out Year_Number;
                 Month  : out Month_Number;
                 Day    : out Day_Number;
                 Seconds : out Day_Duration);

function Time_Of (Year  : Year_Number;
                 Month  : Month_Number;
                 Day    : Day_Number;
                 Seconds : Day_Duration := 0.0)
    return Time;

function "+" (Left : Time;   Right : Duration) return Time;
function "+" (Left : Duration; Right : Time) return Time;
function "-" (Left : Time;   Right : Duration) return Time;
function "-" (Left : Time;   Right : Time) return Duration;

function "<" (Left, Right : Time) return Boolean;
function "<=" (Left, Right : Time) return Boolean;
function ">" (Left, Right : Time) return Boolean;
function ">=" (Left, Right : Time) return Boolean;

Time_Error : exception;

private
    ... -- not specified by the language
end Ada.Calendar;

```

#### Dynamic Semantics

For the execution of a `delay_statement`, the `delay_expression` is first evaluated. For a `delay_until_statement`, the expiration time for the delay is the value of the `delay_expression`, in the time base associated with the type of the expression. For a `delay_relative_statement`, the expiration time is defined as the current time, in the time base associated with relative delays, plus the value of the `delay_expression` converted to the type `Duration`, and then rounded up to the next clock tick. The time base associated with relative delays is as defined in D.9 or is implementation defined.

The task executing a `delay_statement` is blocked until the expiration time is reached, at which point it becomes ready again. If the expiration time has already passed, the task is not blocked.

If an attempt is made to *cancel* the `delay_statement` (as part of an `asynchronous_select` or `abort` — see 12.7.4 and 12.8), the statement is cancelled if the expiration time has not yet passed, thereby completing the `delay_statement`.

The time base associated with the type `Time` of package `Calendar` is implementation defined. The function `Clock` of package `Calendar` returns a value representing the current time for this time base. The implementation-defined value of the named number `System.Tick` (see 16.7) is an approximation of the length of the real-time interval during which the value of `Calendar.Clock` remains constant.

The functions `Year`, `Month`, `Day`, and `Seconds` return the corresponding values for a given value of the type `Time`, as appropriate to an implementation-defined time zone; the procedure `Split` returns all four corresponding values. Conversely, the function `Time_Of` combines a year number, a month number, a day number, and a duration, into a value of type `Time`. The operators `+` and `-` for addition and subtraction of times and durations, and the relational operators for times, have the conventional meaning.

If `Time_Of` is called with a seconds value of 86\_400.0, the value returned is equal to the value of `Time_Of` for the next day with a seconds value of 0.0. The value returned by the function `Seconds` or through the `Seconds` parameter of the procedure `Split` is always less than 86\_400.0.

The exception `Time_Error` is raised by the function `Time_Of` if the actual parameters do not form a proper date. This exception is also raised by the operators `+` and `-` if the result is not representable in the type `Time` or `Duration`, as appropriate. This exception is also raised by the functions `Year`, `Month`, `Day`, and `Seconds` and the procedure `Split` if the year number of the given date is outside of the range of the subtype `Year_Number`.

#### Implementation Requirements

The implementation of the type `Duration` shall allow representation of time intervals (both positive and negative) up to at least 86400 seconds (one day); `Duration'Small` shall not be greater than twenty milliseconds. The implementation of the type `Time` shall allow representation of all dates with year numbers in the range of `Year_Number`; it may allow representation of other dates as well (both earlier and later).

#### Implementation Permissions

An implementation may define additional time types.

An implementation may raise `Time_Error` if the value of a *delay\_expression* in a *delay\_until\_statement* of a *select\_statement* represents a time more than 90 days past the current time. The actual limit, if any, is implementation-defined.

#### Implementation Advice

Whenever possible in an implementation, the value of `Duration'Small` should be no greater than 100 microseconds.

The time base for *delay\_relative\_statements* should be monotonic; it can be different than the time base as used for `Calendar.Clock`.

NOTE 1 A *delay\_relative\_statement* with a negative value of the *delay\_expression* is equivalent to one with a zero value.

NOTE 2 A *delay\_statement* can be executed by the environment task; consequently *delay\_statements* can be executed as part of the elaboration of a *library\_item* or the execution of the main subprogram. Such statements delay the environment task (see 13.2).

NOTE 3 A *delay\_statement* is an abort completion point and a potentially blocking operation, even if the task is not actually blocked.

NOTE 4 There is no necessary relationship between `System.Tick` (the resolution of the clock of package `Calendar`) and `Duration'Small` (the *small* of type `Duration`).

NOTE 5 Additional requirements associated with *delay\_statements* are given in D.9.

EXAMPLE 1 *Example of a relative delay statement:*

```
delay 3.0;  -- delay 3.0 seconds
```

EXAMPLE 2 *Example of a periodic task:*

```
declare
  use Ada.Calendar;
  Next_Time : Time := Clock + Period;
                                -- Period is a global constant of type Duration
begin
  loop                          -- repeated every Period seconds
    delay until Next_Time;
    ... -- perform some actions
    Next_Time := Next_Time + Period;
  end loop;
end;
```



## 12.6.1 Formatting, Time Zones, and other operations for Time

## Static Semantics

The following language-defined library packages exist:

```

package Ada.Calendar.Time_Zones
  with Nonblocking, Global => in out synchronized is
    -- Time zone manipulation:
    type Time_Offset is range -28*60 .. 28*60;
    Unknown_Zone_Error : exception;
    function Local_Time_Offset (Date : Time := Clock) return Time_Offset;
    function UTC_Time_Offset (Date : Time := Clock) return Time_Offset
      renames Local_Time_Offset;
end Ada.Calendar.Time_Zones;

package Ada.Calendar.Arithmetic
  with Nonblocking, Global => in out synchronized is
    -- Arithmetic on days:
    type Day_Count is range
      -366*(1+Year_Number'Last - Year_Number'First)
      ..
      366*(1+Year_Number'Last - Year_Number'First);
    subtype Leap_Seconds_Count is Integer range -2047 .. 2047;
    procedure Difference (Left, Right : in Time;
      Days : out Day_Count;
      Seconds : out Duration;
      Leap_Seconds : out Leap_Seconds_Count);

    function "+" (Left : Time; Right : Day_Count) return Time;
    function "+" (Left : Day_Count; Right : Time) return Time;
    function "-" (Left : Time; Right : Day_Count) return Time;
    function "-" (Left, Right : Time) return Day_Count;
end Ada.Calendar.Arithmetic;

with Ada.Calendar.Time_Zones;
package Ada.Calendar.Formatting
  with Nonblocking, Global => in out synchronized is
    -- Day of the week:
    type Day_Name is (Monday, Tuesday, Wednesday, Thursday,
      Friday, Saturday, Sunday);
    function Day_of_Week (Date : Time) return Day_Name;
    -- Hours:Minutes:Seconds access:
    subtype Hour_Number is Natural range 0 .. 23;
    subtype Minute_Number is Natural range 0 .. 59;
    subtype Second_Number is Natural range 0 .. 59;
    subtype Second_Duration is Day_Duration range 0.0 .. 1.0;
    function Year (Date : Time;
      Time_Zone : Time_Zones.Time_Offset := 0)
      return Year_Number;
    function Month (Date : Time;
      Time_Zone : Time_Zones.Time_Offset := 0)
      return Month_Number;

```

```

function Day      (Date : Time;
                    Time_Zone : Time_Zones.Time_Offset := 0)
    return Day_Number;

function Hour      (Date : Time;
                    Time_Zone : Time_Zones.Time_Offset := 0)
    return Hour_Number;

function Minute     (Date : Time;
                    Time_Zone : Time_Zones.Time_Offset := 0)
    return Minute_Number;

function Second     (Date : Time)
    return Second_Number;

function Sub_Second (Date : Time)
    return Second_Duration;

function Seconds_Of (Hour   : Hour_Number;
                    Minute  : Minute_Number;
                    Second   : Second_Number := 0;
                    Sub_Second : Second_Duration := 0.0)
    return Day_Duration;

procedure Split (Seconds : in Day_Duration;
                Hour      : out Hour_Number;
                Minute    : out Minute_Number;
                Second     : out Second_Number;
                Sub_Second : out Second_Duration);

function Time_Of (Year      : Year_Number;
                Month      : Month_Number;
                Day         : Day_Number;
                Hour        : Hour_Number;
                Minute      : Minute_Number;
                Second       : Second_Number;
                Sub_Second  : Second_Duration := 0.0;
                Leap_Second : Boolean := False;
                Time_Zone   : Time_Zones.Time_Offset := 0)
    return Time;

function Time_Of (Year      : Year_Number;
                Month      : Month_Number;
                Day         : Day_Number;
                Seconds     : Day_Duration := 0.0;
                Leap_Second : Boolean := False;
                Time_Zone   : Time_Zones.Time_Offset := 0)
    return Time;

procedure Split (Date      : in Time;
                Year        : out Year_Number;
                Month       : out Month_Number;
                Day         : out Day_Number;
                Hour        : out Hour_Number;
                Minute      : out Minute_Number;
                Second       : out Second_Number;
                Sub_Second  : out Second_Duration;
                Time_Zone   : in Time_Zones.Time_Offset := 0);

procedure Split (Date      : in Time;
                Year        : out Year_Number;
                Month       : out Month_Number;
                Day         : out Day_Number;
                Hour        : out Hour_Number;
                Minute      : out Minute_Number;
                Second       : out Second_Number;
                Sub_Second  : out Second_Duration;
                Leap_Second : out Boolean;
                Time_Zone   : in Time_Zones.Time_Offset := 0);

```

```

procedure Split (Date      : in Time;
                  Year      : out Year_Number;
                  Month     : out Month_Number;
                  Day       : out Day_Number;
                  Seconds   : out Day_Duration;
                  Leap_Second : out Boolean;
                  Time_Zone  : in Time_Zones.Time_Offset := 0);

-- Simple image and value:
function Image (Date : Time;
                Include_Time_Fraction : Boolean := False;
                Time_Zone : Time_Zones.Time_Offset := 0) return String;

function Local_Image (Date : Time;
                    Include_Time_Fraction : Boolean := False)
    return String is
    (Image (Date, Include_Time_Fraction,
            Time_Zones.Local_Time_Offset (Date)));

function Value (Date : String;
                Time_Zone : Time_Zones.Time_Offset := 0) return Time;

function Image (Elapsed_Time : Duration;
                Include_Time_Fraction : Boolean := False) return String;

function Value (Elapsed_Time : String) return Duration;

end Ada.Calendar.Formatting;

```

Type `Time_Offset` represents for a given locality at a given moment the number of minutes the local time is, at that moment, ahead (+) or behind (-) Coordinated Universal Time (abbreviated UTC). The `Time_Offset` for UTC is zero.

```

function Local_Time_Offset (Date : Time := Clock) return Time_Offset;

```

Returns, as a number of minutes, the `Time_Offset` of the implementation-defined time zone of Calendar, at the time Date. If the time zone of the Calendar implementation is unknown, then `Unknown_Zone_Error` is raised.

```

procedure Difference (Left, Right : in Time;
                    Days : out Day_Count;
                    Seconds : out Duration;
                    Leap_Seconds : out Leap_Seconds_Count);

```

Returns the difference between Left and Right. Days is the number of days of difference, Seconds is the remainder seconds of difference excluding leap seconds, and Leap\_Seconds is the number of leap seconds. If `Left < Right`, then `Seconds ≤ 0.0`, `Days ≤ 0`, and `Leap_Seconds ≤ 0`. Otherwise, all values are nonnegative. The absolute value of Seconds is always less than 86\_400.0. For the returned values, if `Days = 0`, then `Seconds + Duration(Leap_Seconds) = Calendar "-" (Left, Right)`.

```

function "+" (Left : Time; Right : Day_Count) return Time;
function "+" (Left : Day_Count; Right : Time) return Time;

```

Adds a number of days to a time value. `Time_Error` is raised if the result is not representable as a value of type Time.

```

function "-" (Left : Time; Right : Day_Count) return Time;

```

Subtracts a number of days from a time value. `Time_Error` is raised if the result is not representable as a value of type Time.

```

function "-" (Left, Right : Time) return Day_Count;

```

Subtracts two time values, and returns the number of days between them. This is the same value that `Difference` would return in Days.

**function** Day\_of\_Week (Date : Time) **return** Day\_Name;

Returns the day of the week for Time. This is based on the Year, Month, and Day values of Time.

**function** Year (Date : Time;  
Time\_Zone : Time\_Zones.Time\_Offset := 0)  
**return** Year\_Number;

Returns the year for Date, as appropriate for the specified time zone offset.

**function** Month (Date : Time;  
Time\_Zone : Time\_Zones.Time\_Offset := 0)  
**return** Month\_Number;

Returns the month for Date, as appropriate for the specified time zone offset.

**function** Day (Date : Time;  
Time\_Zone : Time\_Zones.Time\_Offset := 0)  
**return** Day\_Number;

Returns the day number for Date, as appropriate for the specified time zone offset.

**function** Hour (Date : Time;  
Time\_Zone : Time\_Zones.Time\_Offset := 0)  
**return** Hour\_Number;

Returns the hour for Date, as appropriate for the specified time zone offset.

**function** Minute (Date : Time;  
Time\_Zone : Time\_Zones.Time\_Offset := 0)  
**return** Minute\_Number;

Returns the minute within the hour for Date, as appropriate for the specified time zone offset.

**function** Second (Date : Time)  
**return** Second\_Number;

Returns the second within the hour and minute for Date.

**function** Sub\_Second (Date : Time)  
**return** Second\_Duration;

Returns the fraction of second for Date (this has the same accuracy as Day\_Duration). The value returned is always less than 1.0.

**function** Seconds\_Of (Hour : Hour\_Number;  
Minute : Minute\_Number;  
Second : Second\_Number := 0;  
Sub\_Second : Second\_Duration := 0.0)  
**return** Day\_Duration;

Returns a Day\_Duration value for the combination of the given Hour, Minute, Second, and Sub\_Second. This value can be used in Calendar.Time\_Of as well as the argument to Calendar."+" and Calendar."-". If Seconds\_Of is called with a Sub\_Second value of 1.0, the value returned is equal to the value of Seconds\_Of for the next second with a Sub\_Second value of 0.0.

**procedure** Split (Seconds : **in** Day\_Duration;  
Hour : **out** Hour\_Number;  
Minute : **out** Minute\_Number;  
Second : **out** Second\_Number;  
Sub\_Second : **out** Second\_Duration);

Splits Seconds into Hour, Minute, Second and Sub\_Second in such a way that the resulting values all belong to their respective subtypes. The value returned in the Sub\_Second parameter is always less than 1.0. If Seconds = 86400.0, Split propagates Time\_Error.

```

function Time_Of (Year      : Year_Number;
                  Month     : Month_Number;
                  Day       : Day_Number;
                  Hour      : Hour_Number;
                  Minute    : Minute_Number;
                  Second     : Second_Number;
                  Sub_Second : Second_Duration := 0.0;
                  Leap_Second: Boolean := False;
                  Time_Zone  : Time_Zones.Time_Offset := 0)
return Time;

```

If Leap\_Second is False, returns a Time built from the date and time values, relative to the specified time zone offset. If Leap\_Second is True, returns the Time that represents the time within the leap second that is one second later than the time specified by the other parameters. Time\_Error is raised if the parameters do not form a proper date or time. If Time\_Of is called with a Sub\_Second value of 1.0, the value returned is equal to the value of Time\_Of for the next second with a Sub\_Second value of 0.0.

```

function Time_Of (Year      : Year_Number;
                  Month     : Month_Number;
                  Day       : Day_Number;
                  Seconds   : Day_Duration := 0.0;
                  Leap_Second: Boolean := False;
                  Time_Zone  : Time_Zones.Time_Offset := 0)
return Time;

```

If Leap\_Second is False, returns a Time built from the date and time values, relative to the specified time zone offset. If Leap\_Second is True, returns the Time that represents the time within the leap second that is one second later than the time specified by the other parameters. Time\_Error is raised if the parameters do not form a proper date or time. If Time\_Of is called with a Seconds value of 86\_400.0, the value returned is equal to the value of Time\_Of for the next day with a Seconds value of 0.0.

```

procedure Split (Date      : in Time;
                 Year      : out Year_Number;
                 Month     : out Month_Number;
                 Day       : out Day_Number;
                 Hour      : out Hour_Number;
                 Minute    : out Minute_Number;
                 Second     : out Second_Number;
                 Sub_Second : out Second_Duration;
                 Leap_Second: out Boolean;
                 Time_Zone  : in Time_Zones.Time_Offset := 0);

```

If Date does not represent a time within a leap second, splits Date into its constituent parts (Year, Month, Day, Hour, Minute, Second, Sub\_Second), relative to the specified time zone offset, and sets Leap\_Second to False. If Date represents a time within a leap second, set the constituent parts to values corresponding to a time one second earlier than that given by Date, relative to the specified time zone offset, and sets Leap\_Seconds to True. The value returned in the Sub\_Second parameter is always less than 1.0.

```

procedure Split (Date      : in Time;
                 Year      : out Year_Number;
                 Month     : out Month_Number;
                 Day       : out Day_Number;
                 Hour      : out Hour_Number;
                 Minute    : out Minute_Number;
                 Second     : out Second_Number;
                 Sub_Second : out Second_Duration;
                 Time_Zone  : in Time_Zones.Time_Offset := 0);

```

Splits Date into its constituent parts (Year, Month, Day, Hour, Minute, Second, Sub\_Second), relative to the specified time zone offset. The value returned in the Sub\_Second parameter is always less than 1.0.

```

procedure Split (Date      : in Time;
                  Year      : out Year_Number;
                  Month     : out Month_Number;
                  Day       : out Day_Number;
                  Seconds   : out Day_Duration;
                  Leap_Second : out Boolean;
                  Time_Zone  : in Time_Zones.Time_Offset := 0);

```

If Date does not represent a time within a leap second, splits Date into its constituent parts (Year, Month, Day, Seconds), relative to the specified time zone offset, and sets Leap\_Second to False. If Date represents a time within a leap second, set the constituent parts to values corresponding to a time one second earlier than that given by Date, relative to the specified time zone offset, and sets Leap\_Seconds to True. The value returned in the Seconds parameter is always less than 86\_400.0.

```

function Image (Date : Time;
                 Include_Time_Fraction : Boolean := False;
                 Time_Zone : Time_Zones.Time_Offset := 0) return String;

```

Returns a string form of the Date relative to the given Time\_Zone. The format is "Year-Month-Day Hour:Minute:Second", where the Year is a 4-digit value, and all others are 2-digit values, of the functions defined in Calendar and Calendar.Formatting, including a leading zero, if needed. The separators between the values are a minus, another minus, a colon, and a single space between the Day and Hour. If Include\_Time\_Fraction is True, the integer part of Sub\_Seconds\*100 is suffixed to the string as a point followed by a 2-digit value.

```

function Value (Date : String;
                 Time_Zone : Time_Zones.Time_Offset := 0) return Time;

```

Returns a Time value for the image given as Date, relative to the given time zone. Constraint\_Error is raised if the string is not formatted as described for Image, or the function cannot interpret the given string as a Time value.

```

function Image (Elapsed_Time : Duration;
                 Include_Time_Fraction : Boolean := False) return String;

```

Returns a string form of the Elapsed\_Time. The format is "Hour:Minute:Second", where all values are 2-digit values, including a leading zero, if necessary. The separators between the values are colons. If Include\_Time\_Fraction is True, the integer part of Sub\_Seconds\*100 is suffixed to the string as a point followed by a 2-digit value. If Elapsed\_Time < 0.0, the result is Image (**abs** Elapsed\_Time, Include\_Time\_Fraction) prefixed with a minus sign. If **abs** Elapsed\_Time represents 100 hours or more, the result is implementation-defined.

```

function Value (Elapsed_Time : String) return Duration;

```

Returns a Duration value for the image given as Elapsed\_Time. Constraint\_Error is raised if the string is not formatted as described for Image, or the function cannot interpret the given string as a Duration value.

#### Implementation Advice

An implementation should support leap seconds if the target system supports them. If leap seconds are not supported, Difference should return zero for Leap\_Seconds, Split should return False for Leap\_Second, and Time\_Of should raise Time\_Error if Leap\_Second is True.

NOTE 1 The implementation-defined time zone of package Calendar can be the local time zone. Local\_Time\_Offset always returns the difference relative to the implementation-defined time zone of package Calendar. If Local\_Time\_Offset does not raise Unknown\_Zone\_Error, UTC time can be safely calculated (within the accuracy of the underlying time-base).

NOTE 2 Calling Split on the results of subtracting Duration(Local\_Time\_Offset\*60) from Clock provides the components (hours, minutes, and so on) of the UTC time. In the United States, for example, Local\_Time\_Offset will generally be negative.



## 12.7 Select Statements

There are four forms of the `select_statement`. One form provides a selective wait for one or more `select_alternatives`. Two provide timed and conditional entry calls. The fourth provides asynchronous transfer of control.

*Syntax*

```
select_statement ::=
  selective_accept
| timed_entry_call
| conditional_entry_call
| asynchronous_select
```

EXAMPLE *Example of a select statement:*

```
select
  accept Driver_Awake_Signal;
or
  delay 30.0*Seconds;
  Stop_The_Train;
end select;
```

### 12.7.1 Selective Accept

This form of the `select_statement` allows a combination of waiting for, and selecting from, one or more alternatives. The selection may depend on conditions associated with each alternative of the `selective_accept`.

*Syntax*

```
selective_accept ::=
select
  [guard]
  select_alternative
{ or
  [guard]
  select_alternative }
[ else
  sequence_of_statements ]
end select;

guard ::= when condition =>

select_alternative ::=
  accept_alternative
| delay_alternative
| terminate_alternative

accept_alternative ::=
  accept_statement [sequence_of_statements]

delay_alternative ::=
  delay_statement [sequence_of_statements]

terminate_alternative ::= terminate;
```

A `selective_accept` shall contain at least one `accept_alternative`. In addition, it can contain:

- a `terminate_alternative` (only one); or
- one or more `delay_alternatives`; or
- an *else part* (the reserved word **else** followed by a `sequence_of_statements`).

These three possibilities are mutually exclusive.

*Legality Rules*

If a `selective_accept` contains more than one `delay_alternative`, then all shall be `delay_relative_statements`, or all shall be `delay_until_statements` for the same time type.

*Dynamic Semantics*

A `select_alternative` is said to be *open* if it is not immediately preceded by a guard, or if the condition of its guard evaluates to True. It is said to be *closed* otherwise.

For the execution of a `selective_accept`, any guard conditions are evaluated; open alternatives are thus determined. For an open `delay_alternative`, the *delay\_expression* is also evaluated. Similarly, for an open `accept_alternative` for an entry of a family, the *entry\_index* is also evaluated. These evaluations are performed in an arbitrary order, except that a *delay\_expression* or *entry\_index* is not evaluated until after evaluating the corresponding condition, if any. Selection and execution of one open alternative, or of the else part, then completes the execution of the `selective_accept`; the rules for this selection are described below.

Open `accept_alternatives` are first considered. Selection of one such alternative takes place immediately if the corresponding entry already has queued calls. If several alternatives can thus be selected, one of them is selected according to the entry queuing policy in effect (see 12.5.3 and D.4). When such an alternative is selected, the selected call is removed from its entry queue and the *handled\_sequence\_of\_statements* (if any) of the corresponding `accept_statement` is executed; after the rendezvous completes any subsequent *sequence\_of\_statements* of the alternative is executed. If no selection is immediately possible (in the above sense) and there is no else part, the task blocks until an open alternative can be selected.

Selection of the other forms of alternative or of an else part is performed as follows:

- An open `delay_alternative` is selected when its expiration time is reached if no `accept_alternative` or other `delay_alternative` can be selected prior to the expiration time. If several `delay_alternatives` have this same expiration time, one of them is selected according to the queuing policy in effect (see D.4); the default queuing policy chooses arbitrarily among the `delay_alternatives` whose expiration time has passed.
- The else part is selected and its *sequence\_of\_statements* is executed if no `accept_alternative` can immediately be selected; in particular, if all alternatives are closed.
- An open `terminate_alternative` is selected if the conditions stated at the end of 12.3 are satisfied.

The exception `Program_Error` is raised if all alternatives are closed and there is no else part.

NOTE A `selective_accept` can have several open `delay_alternatives`. A `selective_accept` can have several open `accept_alternatives` for the same entry.

EXAMPLE *Example of a task body with a selective accept:*

```

task body Server is
  Current_Work_Item : Work_Item;
begin
  loop
    select
      accept Next_Work_Item(WI : in Work_Item) do
        Current_Work_Item := WI;
      end;
      Process_Work_Item(Current_Work_Item);
    or
      accept Shut_Down;
      exit;      -- Premature shut down requested
    or
      terminate; -- Normal shutdown at end of scope
    end select;
  end loop;
end Server;

```

### 12.7.2 Timed Entry Calls

A `timed_entry_call` issues an entry call that is cancelled if the call (or a requeue-with-abort of the call) is not selected before the expiration time is reached. A procedure call may appear rather than an entry call for cases where the procedure can be implemented by an entry.

#### *Syntax*

```

timed_entry_call ::=
  select
    entry_call_alternative
  or
    delay_alternative
  end select;

entry_call_alternative ::=
  procedure_or_entry_call [sequence_of_statements]

procedure_or_entry_call ::=
  procedure_call_statement | entry_call_statement

```

#### *Legality Rules*

If a `procedure_call_statement` is used for a `procedure_or_entry_call`, the *procedure\_name* or *procedure\_prefix* of the `procedure_call_statement` shall statically denote an entry renamed as a procedure or (a view of) a primitive subprogram of a limited interface whose first parameter is a controlling parameter (see 6.9.2).

#### *Dynamic Semantics*

For the execution of a `timed_entry_call`, the *entry\_name*, *procedure\_name*, or *procedure\_prefix*, and any actual parameters are evaluated, as for a simple entry call (see 12.5.3) or procedure call (see 9.4). The expiration time (see 12.6) for the call is determined by evaluating the *delay\_expression* of the `delay_alternative`. If the call is an entry call or a call on a procedure implemented by an entry, the entry call is then issued. Otherwise, the call proceeds as described in 9.4 for a procedure call, followed by the *sequence\_of\_statements* of the `entry_call_alternative`; the *sequence\_of\_statements* of the `delay_alternative` is ignored.

If the call is queued (including due to a requeue-with-abort), and not selected before the expiration time is reached, an attempt to cancel the call is made. If the call completes due to the cancellation, the optional *sequence\_of\_statements* of the `delay_alternative` is executed; if the entry call completes normally, the optional *sequence\_of\_statements* of the `entry_call_alternative` is executed.

EXAMPLE *Example of a timed entry call:*

```
select
  Controller.Request(Medium) (Some_Item);
or
  delay 45.0;
  -- controller too busy, try something else
end select;
```

### 12.7.3 Conditional Entry Calls

A `conditional_entry_call` issues an entry call that is then cancelled if it is not selected immediately (or if a `requeue-with-abort` of the call is not selected immediately). A procedure call may appear rather than an entry call for cases where the procedure can be implemented by an entry.

*Syntax*

```
conditional_entry_call ::=
  select
    entry_call_alternative
  else
    sequence_of_statements
  end select;
```

*Dynamic Semantics*

The execution of a `conditional_entry_call` is defined to be equivalent to the execution of a `timed_entry_call` with a `delay_alternative` specifying an immediate expiration time and the same `sequence_of_statements` as given after the reserved word **else**.

NOTE A `conditional_entry_call` can briefly increase the Count attribute of the entry, even if the conditional call is not selected.

EXAMPLE *Example of a conditional entry call:*

```
procedure Spin(R : in out Resource) is -- see 12.4
begin
  loop
    select
      R.Seize;
      return;
    else
      null; -- busy waiting
    end select;
  end loop;
end;
```

### 12.7.4 Asynchronous Transfer of Control

An asynchronous `select_statement` provides asynchronous transfer of control upon completion of an entry call or the expiration of a delay.

*Syntax*

```
asynchronous_select ::=
  select
    triggering_alternative
  then abort
    abortable_part
  end select;

triggering_alternative ::= triggering_statement [sequence_of_statements]
triggering_statement ::= procedure_or_entry_call | delay_statement
abortable_part ::= sequence_of_statements
```

For the execution of an `asynchronous_select` whose `triggering_statement` is a `procedure_or_entry_call`, the `entry_name`, `procedure_name`, or `procedure_prefix`, and actual parameters are evaluated as for a simple entry call (see 12.5.3) or procedure call (see 9.4). If the call is an entry call or a call on a procedure implemented by an entry, the entry call is issued. If the entry call is queued (or requeued-with-abort), then the `abortable_part` is executed. If the entry call is selected immediately, and never requeued-with-abort, then the `abortable_part` is never started. If the call is on a procedure that is not implemented by an entry, the call proceeds as described in 9.4, followed by the `sequence_of_statements` of the `triggering_alternative`; the `abortable_part` is never started.

For the execution of an `asynchronous_select` whose `triggering_statement` is a `delay_statement`, the `delay_expression` is evaluated and the expiration time is determined, as for a normal `delay_statement`. If the expiration time has not already passed, the `abortable_part` is executed.

If the `abortable_part` completes and is left prior to completion of the `triggering_statement`, an attempt to cancel the `triggering_statement` is made. If the attempt to cancel succeeds (see 12.5.3 and 12.6), the `asynchronous_select` is complete.

If the `triggering_statement` completes other than due to cancellation, the `abortable_part` is aborted (if started but not yet completed — see 12.8). If the `triggering_statement` completes normally, the optional `sequence_of_statements` of the `triggering_alternative` is executed after the `abortable_part` is left.

EXAMPLE 1 *Example of a main command loop for a command interpreter:*

```
loop
  select
    Terminal.Wait_For_Interrupt;
    Put_Line("Interrupted");
  then abort
    -- This will be abandoned upon terminal interrupt
    Put_Line("-> ");
    Get_Line(Command, Last);
    Process_Command(Command(1..Last));
  end select;
end loop;
```

EXAMPLE 2 *Example of a time-limited calculation:*

```
select
  delay 5.0;
  Put_Line("Calculation does not converge");
then abort
  -- This calculation is expected to finish in 5.0 seconds;
  -- if not, it is assumed to diverge.
  Horribly_Complicated_Recursive_Function(X, Y);
end select;
```

Note that these examples presume that there are abort completion points (see 12.8) within the execution of the `abortable_part`.

## 12.8 Abort of a Task - Abort of a Sequence of Statements

An `abort_statement` causes one or more tasks to become abnormal, thus preventing any further interaction with such tasks. The completion of the `triggering_statement` of an `asynchronous_select` causes a `sequence_of_statements` to be aborted.

### Syntax

```
abort_statement ::= abort task_name { task_name};
```

## Name Resolution Rules

Each *task\_name* is expected to be of any task type; each can be of a different task type.

## Dynamic Semantics

For the execution of an *abort\_statement*, the given *task\_names* are evaluated in an arbitrary order. Each named task is then *aborted*, which consists of making the task *abnormal* and aborting the execution of the corresponding *task\_body*, unless it is already completed.

When the execution of a construct is *aborted* (including that of a *task\_body* or of a *sequence\_of\_statements*), the execution of every construct included within the aborted execution is also aborted, except for executions included within the execution of an *abort-deferred* operation; the execution of an abort-deferred operation continues to completion without being affected by the abort; the following are the abort-deferred operations:

- a protected action;
- waiting for an entry call to complete (after having initiated the attempt to cancel it — see below);
- waiting for the termination of dependent tasks;
- the execution of an Initialize procedure as the last step of the default initialization of a controlled object;
- the execution of a Finalize procedure as part of the finalization of a controlled object;
- an assignment operation to an object with a controlled part.

The last three of these are discussed further in 10.6.

When a master is aborted, all tasks that depend on that master are aborted.

The order in which tasks become abnormal as the result of an *abort\_statement* or the abort of a *sequence\_of\_statements* is not specified by the language.

If the execution of an entry call is aborted, an immediate attempt is made to cancel the entry call (see 12.5.3). If the execution of a construct is aborted at a time when the execution is blocked, other than for an entry call, at a point that is outside the execution of an abort-deferred operation, then the execution of the construct completes immediately. For an abort due to an *abort\_statement*, these immediate effects occur before the execution of the *abort\_statement* completes. Other than for these immediate cases, the execution of a construct that is aborted does not necessarily complete before the *abort\_statement* completes. However, the execution of the aborted construct completes no later than its next *abort completion point* (if any) that occurs outside of an abort-deferred operation; the following are abort completion points for an execution:

- the point where the execution initiates the activation of another task;
- the end of the activation of a task;
- a point within a parallel construct where a new logical thread of control is created;
- the end of a parallel construct;
- the start or end of the execution of an entry call, *accept\_statement*, *delay\_statement*, or *abort\_statement*;
- the start of the execution of a *select\_statement*, or of the *sequence\_of\_statements* of an *exception\_handler*.

## Bounded (Run-Time) Errors

An attempt to execute an *asynchronous\_select* as part of the execution of an abort-deferred operation is a bounded error. Similarly, an attempt to create a task that depends on a master that is included entirely within the execution of an abort-deferred operation is a bounded error. In both cases, *Program\_Error* is raised if the error is detected by the implementation; otherwise, the operations



proceed as they would outside an abort-deferred operation, except that an abort of the `abortable_part` or the created task does not necessarily have an effect.

*Erroneous Execution*

If an assignment operation completes prematurely due to an abort, the assignment is said to be *disrupted*; the target of the assignment or its parts can become abnormal, and certain subsequent uses of the object can be erroneous, as explained in 16.9.1.

NOTE 1 An `abort_statement` is best used only in situations requiring unconditional termination.

NOTE 2 A task is allowed to abort any task it can name, including itself.

NOTE 3 Additional requirements associated with abort are given in D.6.

## 12.9 Task and Entry Attributes

*Dynamic Semantics*

For a `prefix T` that is of a task type (after any implicit dereference), the following attributes are defined:

**T'Callable** Yields the value `True` when the task denoted by `T` is *callable*, and `False` otherwise; a task is callable unless it is completed or abnormal. The value of this attribute is of the predefined type `Boolean`.

**T'Terminated**

Yields the value `True` if the task denoted by `T` is terminated, and `False` otherwise. The value of this attribute is of the predefined type `Boolean`.

For a `prefix E` that denotes an entry of a task or protected unit, the following attribute is defined. This attribute is only allowed within the body of the task or protected unit, but excluding, in the case of an entry of a task unit, within any program unit that is, itself, inner to the body of the task unit.

**E'Count** Yields the number of calls presently queued on the entry `E` of the current instance of the unit. The value of this attribute is of the type *universal\_integer*.

NOTE 1 For the `Count` attribute, the entry can be either a single entry or an entry of a family. The name of the entry or entry family can be either a `direct_name` or an expanded name.

NOTE 2 Within task units, by interrogating the attribute `E'Count` an algorithm can allow for the increase of the value of this attribute for incoming entry calls, and its decrease, for example with `timed_entry_calls`. A `conditional_entry_call` can also briefly increase this value, even if the conditional call is not accepted.

NOTE 3 Within protected units, by interrogating the attribute `E'Count` in the `entry_barrier` for the entry `E` an algorithm can allow for the evaluation of the `condition` of the barrier both before and after queuing a given caller.

### 12.10 Shared Variables

*Static Semantics*

If two different objects, including nonoverlapping parts of the same object, are *independently addressable*, they can be manipulated concurrently by two different logical threads of control without synchronization, unless both are subcomponents of the same full access object, and either is nonatomic (see C.6). Any two nonoverlapping objects are independently addressable if either object is specified as independently addressable (see C.6). Otherwise, two nonoverlapping objects are independently addressable except when they are both parts of a composite object for which a nonconfirming value is specified for any of the following representation aspects: (record) `Layout`, `Component_Size`, `Pack`, `Atomic`, or `Convention`; in this case it is unspecified whether the parts are independently addressable.

Separate logical threads of control normally proceed independently and concurrently with one another. However, task interactions can be used to synchronize the actions of two or more logical threads of control to allow, for example, meaningful communication by the direct updating and reading of variables shared between them. The actions of two different logical threads of control are synchronized in this sense when an action of one *signals* an action of the other; an action A1 is defined to signal an action A2 under the following circumstances:

- If A1 and A2 are part of the execution of the same task, and the language rules require A1 to be performed before A2;
- If A1 is the action of an activator that initiates the activation of a task, and A2 is part of the execution of the task that is activated;
- If A1 is part of the activation of a task, and A2 is the action of waiting for completion of the activation;
- If A1 is part of the execution of a task, and A2 is the action of waiting for the termination of the task;
- If A1 is the termination of a task T, and A2 is either an evaluation of the expression T.Terminated that results in True, or a call to Ada.Task\_Identification.Is\_Terminated with an actual parameter that identifies T and a result of True (see C.7.1);
- If A1 is the action of issuing an entry call, and A2 is part of the corresponding execution of the appropriate `entry_body` or `accept_statement`;
- If A1 is part of the execution of an `accept_statement` or `entry_body`, and A2 is the action of returning from the corresponding entry call;
- If A1 is part of the execution of a protected procedure body or `entry_body` for a given protected object, and A2 is part of a later execution of an `entry_body` for the same protected object;
- If A1 signals some action that in turn signals A2.

Action A1 is defined to *potentially signal* action A2 if A1 signals A2, if action A1 and A2 occur as part of the execution of the same logical thread of control, and the language rules permit action A1 to precede action A2, or if action A1 potentially signals some action that in turn potentially signals A2.

Two actions are defined to be *sequential* if one of the following is true:

- One action signals the other;
- Both actions occur as part of the execution of the same logical thread of control;
- Both actions occur as part of protected actions on the same protected object, and at least one of the actions is part of a call on an exclusive protected operation of the protected object.

Aspect Atomic or aspect Atomic\_Components may also be specified to ensure that certain reads and updates are sequential — see C.6.

Two actions that are not sequential are defined to be *concurrent* actions.

Two actions are defined to *conflict* if one action assigns to an object, and the other action reads or assigns to a part of the same object (or of a neighboring object if the two are not independently addressable). The action comprising a call on a subprogram or an entry is defined to *potentially conflict* with another action if the Global aspect (or Global'Class aspect in the case of a dispatching call) of the called subprogram or entry is such that a conflicting action would be possible during the execution of the call. Similarly, two calls are considered to potentially conflict if they each have Global (or Global'Class in the case of a dispatching call) aspects such that conflicting actions would be possible during the execution of the calls. Finally, two actions that conflict are also considered to potentially conflict.

A *synchronized* object is an object of a task or protected type, an atomic object (see C.6), a suspension object (see D.10), or a synchronous barrier (see D.10.1). Operations on such objects are necessarily sequential with respect to one another, and hence are never considered to conflict.

*Erroneous Execution*

The execution of two concurrent actions is erroneous if the actions make conflicting uses of a shared variable (or neighboring variables that are not independently addressable).

**12.10.1 Conflict Check Policies**

This subclause determines what checks are performed relating to possible concurrent conflicting actions (see 12.10).

*Syntax*

The form of a `pragma Conflict_Check_Policy` is as follows:

**pragma** `Conflict_Check_Policy` (*policy\_identifier* [, *policy\_identifier*]);

A `pragma Conflict_Check_Policy` is allowed only immediately within a `declarative_part`, a `package_specification`, or as a configuration `pragma`.

*Legality Rules*

Each *policy\_identifier* shall be one of `No_Parallel_Conflict_Checks`, `Known_Parallel_Conflict_Checks`, `All_Parallel_Conflict_Checks`, `No_Tasking_Conflict_Checks`, `Known_Tasking_Conflict_Checks`, `All_Tasking_Conflict_Checks`, `No_Conflict_Checks`, `Known_Conflict_Checks`, `All_Conflict_Checks`, or an implementation-defined conflict check policy. If two *policy\_identifiers* are given, one shall include the word `Parallel` and one shall include the word `Tasking`. If only one *policy\_identifier* is given, it shall not include the word `Parallel` or `Tasking`.

A `pragma Conflict_Check_Policy` given in a `declarative_part` or immediately within a `package_specification` applies from the place of the `pragma` to the end of the innermost enclosing `declarative_region`. The region for a `pragma Conflict_Check_Policy` given as a configuration `pragma` is the `declarative_region` for the entire compilation unit (or units) to which it applies.

If a `pragma Conflict_Check_Policy` applies to a `generic_instantiation`, then the `pragma Conflict_Check_Policy` applies to the entire instance.

If multiple `Conflict_Check_Policy` `pragmas` apply to a given construct, the conflict check policy is determined by the one in the innermost enclosing region. If no `Conflict_Check_Policy` `pragma` applies to a construct, the policy is (`All_Parallel_Conflict_Checks`, `No_Tasking_Conflict_Checks`) (see below).

Certain potentially conflicting actions are disallowed according to which conflict check policies apply at the place where the action or actions occur, as follows:

**No\_Parallel\_Conflict\_Checks**

This policy imposes no restrictions on concurrent actions arising from parallel constructs.

**No\_Tasking\_Conflict\_Checks**

This policy imposes no restrictions on concurrent actions arising from tasking constructs.

**Known\_Parallel\_Conflict\_Checks**

If this policy applies to two concurrent actions appearing within parallel constructs, they are disallowed if they are known to denote the same object (see 9.4.1) with uses that conflict. For the purposes of this check, any parallel loop may be presumed to involve multiple concurrent iterations. Also, for the purposes of deciding whether two actions are concurrent, it is enough for the logical threads of control in which they occur to be concurrent at any point in their execution, unless all of the following are true:

- the shared object is volatile;
- the two logical threads of control are both known to also refer to a shared synchronized object; and
- each thread whose potentially conflicting action updates the shared volatile object, also updates this shared synchronized object.

**Known\_Tasking\_Conflict\_Checks**

If this policy applies to two concurrent actions appearing within the same compilation unit, at least one of which appears within a task body but not within a parallel construct, they are disallowed if they are known to denote the same object (see 9.4.1) with uses that conflict, and neither potentially signals the other (see 12.10). For the purposes of this check, any named task type may be presumed to have multiple instances. Also, for the purposes of deciding whether two actions are concurrent, it is enough for the tasks in which they occur to be concurrent at any point in their execution, unless all of the following are true:

- the shared object is volatile;
- the two tasks are both known to also refer to a shared synchronized object; and
- each task whose potentially conflicting action updates the shared volatile object, also updates this shared synchronized object.

**All\_Parallel\_Conflict\_Checks**

This policy includes the restrictions imposed by the Known\_Parallel\_Conflict\_Checks policy, and in addition disallows a parallel construct from reading or updating a variable that is global to the construct, unless it is a synchronized object, or unless the construct is a parallel loop, and the global variable is a part of a component of an array denoted by an indexed component with at least one index expression that statically denotes the loop parameter of the `loop_parameter_specification` or the chunk parameter of the parallel loop.

**All\_Tasking\_Conflict\_Checks**

This policy includes the restrictions imposed by the Known\_Tasking\_Conflict\_Checks policy, and in addition disallows a task body from reading or updating a variable that is global to the task body, unless it is a synchronized object.

**No\_Conflict\_Checks, Known\_Conflict\_Checks, All\_Conflict\_Checks**

These are shorthands for (No\_Parallel\_Conflict\_Checks, No\_Tasking\_Conflict\_Checks), (Known\_Parallel\_Conflict\_Checks, Known\_Tasking\_Conflict\_Checks), and (All\_Parallel\_Conflict\_Checks, All\_Tasking\_Conflict\_Checks), respectively.

*Static Semantics*

For a subprogram, the following language-defined representation aspect may be specified:

**Parallel\_Calls**

The Parallel\_Calls aspect is of type Boolean. The specified value shall be static. The Parallel\_Calls aspect of an inherited primitive subprogram is True if Parallel\_Calls is True either for the corresponding subprogram of the progenitor type or for any other inherited subprogram that it overrides. If not specified or inherited as True, the Parallel\_Calls aspect of a subprogram is False.

Specifying the Parallel\_Calls aspect to be True for a subprogram indicates that the subprogram can be safely called in parallel. Conflict checks (if required by the Conflict\_Check\_Policy in effect) are made on the subprogram assuming that multiple concurrent calls exist. Such checks can then be omitted on a call of the subprogram in a parallel iteration context.

*Implementation Permissions*

When the conflict check policy Known\_Parallel\_Conflict\_Checks or All\_Parallel\_Conflict\_Checks applies, the implementation may disallow two concurrent actions appearing within parallel constructs if the implementation can prove they will at run-time denote the same object with uses that conflict. Similarly, when the conflict check policy Known\_Tasking\_Conflict\_Checks or All\_Tasking\_Conflict\_Checks applies, the implementation may disallow two concurrent actions, at least one of which appears within a task body but not within a parallel construct, if the implementation can prove they will at run-time denote the same object with uses that conflict.

## 12.11 Example of Tasking and Synchronization

EXAMPLE The following example defines a buffer protected object to smooth variations between the speed of output of a producing task and the speed of input of some consuming task. For instance, the producing task can have the following structure:

```
task Producer;

task body Producer is
  Person : Person_Name; -- see 6.10.1
begin
  loop
    ... -- simulate arrival of the next customer
    Buffer.Append_Wait(Person);
    exit when Person = null;
  end loop;
end Producer;
```

and the consuming task can have the following structure:

```
task Consumer;

task body Consumer is
  Person : Person_Name;
begin
  loop
    Buffer.Remove_First_Wait(Person);
    exit when Person = null;
    ... -- simulate serving a customer
  end loop;
end Consumer;
```

The buffer object contains an internal array of person names managed in a round-robin fashion. The array has two indices, an In\_Index denoting the index for the next input person name and an Out\_Index denoting the index for the next output person name.

The Buffer is defined as an extension of the Synchronized\_Queue interface (see 6.9.4), and as such promises to implement the abstraction defined by that interface. By doing so, the Buffer can be passed to the Transfer class-wide operation defined for objects of a type covered by Queue'Class.

```
type Person_Name_Array is array (Positive range <>)
  of Person_Name; -- see 6.10.1

protected Buffer is new Synchronized_Queue with -- see 6.9.4
  entry Append_Wait(Person : in Person_Name);
  entry Remove_First_Wait(Person : out Person_Name);
  function Cur_Count return Natural;
  function Max_Count return Natural;
  procedure Append(Person : in Person_Name);
  procedure Remove_First(Person : out Person_Name);
private
  Pool : Person_Name_Array(1 .. 100);
  Count : Natural := 0;
  In_Index, Out_Index : Positive := 1;
end Buffer;

protected body Buffer is
  entry Append_Wait(Person : in Person_Name)
    when Count < Pool.Length is
  begin
    Append(Person);
  end Append_Wait;
```

```

procedure Append(Person : in Person_Name) is
begin
    if Count = Pool'Length then
        raise Queue_Error with "Buffer Full"; -- see 14.3
    end if;
    Pool(In_Index) := Person;
    In_Index      := (In_Index mod Pool'Length) + 1;
    Count        := Count + 1;
end Append;

entry Remove_First_Wait(Person : out Person_Name)
    when Count > 0 is
begin
    Remove_First(Person);
end Remove_First_Wait;

procedure Remove_First(Person : out Person_Name) is
begin
    if Count = 0 then
        raise Queue_Error with "Buffer Empty"; -- see 14.3
    end if;
    Person      := Pool(Out_Index);
    Out_Index   := (Out_Index mod Pool'Length) + 1;
    Count       := Count - 1;
end Remove_First;

function Cur_Count return Natural is
begin
    return Buffer.Count;
end Cur_Count;

function Max_Count return Natural is
begin
    return Pool'Length;
end Max_Count;
end Buffer;

```

## 13 Program Structure and Compilation Issues

The overall structure of programs and the facilities for separate compilation are described in this clause. A *program* is a set of *partitions*, each of which may execute in a separate address space, possibly on a separate computer.

As explained below, a partition is constructed from *library units*. Syntactically, the declaration of a library unit is a *library\_item*, as is the body of a library unit. An implementation may support a concept of a *program library* (or simply, a “library”), which contains *library\_items* and their subunits. Library units may be organized into a hierarchy of children, grandchildren, and so on.

This clause has two subclauses: 13.1 discusses compile-time issues related to separate compilation. 13.2 discusses issues related to what is traditionally known as “link time” and “run time” — building and executing partitions.

### 13.1 Separate Compilation

A *program unit* is either a package, a task unit, a protected unit, a protected entry, a generic unit, or an explicitly declared subprogram other than an enumeration literal. Certain kinds of program units can be separately compiled. Alternatively, they can appear physically nested within other program units.

The text of a program can be submitted to the compiler in one or more compilations. Each compilation is a succession of *compilation\_units*. A *compilation\_unit* contains either the declaration, the body, or a renaming of a program unit. The representation for a compilation is implementation-defined.



A library unit is a separately compiled program unit, and is a package, subprogram, or generic unit. Library units may have other (logically nested) library units as children, and may have other program units physically nested within them. A root library unit, together with its children and grandchildren and so on, form a *subsystem*.

#### Implementation Permissions

An implementation may impose implementation-defined restrictions on compilations that contain multiple compilation\_units.

### 13.1.1 Compilation Units - Library Units

A `library_item` is a compilation unit that is the declaration, body, or renaming of a library unit. Each library unit (except Standard) has a *parent unit*, which is a library package or generic library package. A library unit is a *child* of its parent unit. The *root* library units are the children of the predefined library package Standard.

#### Syntax

```

compilation ::= {compilation_unit}
compilation_unit ::=
  context_clause library_item
| context_clause subunit
library_item ::= [private] library_unit_declaration
| library_unit_body
| [private] library_unit_renaming_declaration
library_unit_declaration ::=
  subprogram_declaration | package_declaration
| generic_declaration | generic_instantiation
library_unit_renaming_declaration ::=
  package_renaming_declaration
| generic_renaming_declaration
| subprogram_renaming_declaration
library_unit_body ::= subprogram_body | package_body
parent_unit_name ::= name

```

An overriding\_indicator is not allowed in a subprogram\_declaration, generic\_instantiation, or subprogram\_renaming\_declaration that declares a library unit.

A *library unit* is a program unit that is declared by a `library_item`. When a program unit is a library unit, the prefix “library” is used to refer to it (or “generic library” if generic), as well as to its declaration and body, as in “library procedure”, “library package\_body”, or “generic library package”. The term *compilation unit* is used to refer to a `compilation_unit`. When the meaning is clear from context, the term is also used to refer to the `library_item` of a `compilation_unit` or to the proper\_body of a subunit (that is, the `compilation_unit` without the `context_clause` and the **separate** (`parent_unit_name`)).

The *parent declaration* of a `library_item` (and of the library unit) is the declaration denoted by the `parent_unit_name`, if any, of the `defining_program_unit_name` of the `library_item`. If there is no `parent_unit_name`, the parent declaration is the declaration of Standard, the `library_item` is a *root library\_item*, and the library unit (renaming) is a *root library unit* (renaming). The declaration and body of Standard itself have no parent declaration. The *parent unit* of a `library_item` or library unit is the library unit declared by its parent declaration.

The children of a library unit occur immediately within the declarative region of the declaration of the library unit. The *ancestors* of a library unit are itself, its parent, its parent's parent, and so on. (Standard is an ancestor of every library unit.) The *descendant* relation is the inverse of the ancestor relation.

A `library_unit_declaration` or a `library_unit_renaming_declaration` is *private* if the declaration is immediately preceded by the reserved word **private**; it is otherwise *public*. A library unit is private or public according to its declaration. The *public descendants* of a library unit are the library unit itself, and the public descendants of its public children. Its other descendants are *private descendants*.

For each `library_package_declaration` in the environment, there is an implicit declaration of a *limited view* of that library package. The limited view of a package contains:

- For each `package_declaration` occurring immediately within the visible part, a declaration of the limited view of that package, with the same `defining_program_unit_name`.
- For each `type_declaration` occurring immediately within the visible part that is not an `incomplete_type_declaration`, an incomplete view of the type with no `discriminant_part`; if the `type_declaration` is tagged, then the view is a tagged incomplete view.

The limited view of a `library_package_declaration` is private if that `library_package_declaration` is immediately preceded by the reserved word **private**.

There is no syntax for declaring limited views of packages, because they are always implicit. The implicit declaration of a limited view of a library package is not the declaration of a library unit (the `library_package_declaration` is); nonetheless, it is a `library_item`. The implicit declaration of the limited view of a library package forms an (implicit) compilation unit whose `context_clause` is empty.

A `library_package_declaration` is the completion of the declaration of its limited view.

#### Legality Rules

The parent unit of a `library_item` shall be a library package or generic library package.

If a `defining_program_unit_name` of a given declaration or body has a `parent_unit_name`, then the given declaration or body shall be a `library_item`. The body of a program unit shall be a `library_item` if and only if the declaration of the program unit is a `library_item`. In a `library_unit_renaming_declaration`, the (old) name shall denote a `library_item`.

A `parent_unit_name` (which can be used within a `defining_program_unit_name` of a `library_item` and in the **separate** clause of a subunit), and each of its prefixes, shall not denote a renaming declaration. On the other hand, a name that denotes a `library_unit_renaming_declaration` is allowed in a `nonlimited_with_clause` and other places where the name of a library unit is allowed.

If a library package is an instance of a generic package, then every child of the library package shall either be itself an instance or be a renaming of a library unit.

A child of a generic library package shall either be itself a generic unit or be a renaming of some other child of the same generic unit.

A child of a parent generic package shall be instantiated or renamed only within the declarative region of the parent generic.

For each child *C* of some parent generic package *P*, there is a corresponding declaration *C* nested immediately within each instance of *P*. For the purposes of this rule, if a child *C* itself has a child *D*, each corresponding declaration for *C* has a corresponding child *D*. The corresponding declaration for a child within an instance is visible only within the scope of a `with_clause` that mentions the (original) child generic unit.

A library subprogram shall not override a primitive subprogram.

The defining name of a function that is a compilation unit shall not be an `operator_symbol`.

#### Static Semantics

A `subprogram_renaming_declaration` that is a `library_unit_renaming_declaration` is a `renaming-as-declaration`, not a `renaming-as-body`.

There are two kinds of dependences among compilation units:

- The *semantic dependences* (see below) are the ones necessary to check the compile-time rules across compilation unit boundaries; a compilation unit depends semantically on the other compilation units necessary to determine its legality. The visibility rules are based on the semantic dependences.
- The *elaboration dependences* (see 13.2) determine the order of elaboration of `library_items`.

A `library_item` depends semantically upon its parent declaration. A subunit depends semantically upon its parent body. A `library_unit_body` depends semantically upon the corresponding `library_unit_declaration`, if any. The declaration of the limited view of a library package depends semantically upon the declaration of the limited view of its parent. The declaration of a library package depends semantically upon the declaration of its limited view. A compilation unit depends semantically upon each `library_item` mentioned in a `with_clause` of the compilation unit. In addition, if a given compilation unit contains an `attribute_reference` of a type defined in another compilation unit, then the given compilation unit depends semantically upon the other compilation unit. The semantic dependence relationship is transitive.

#### Dynamic Semantics

The elaboration of the declaration of the limited view of a package has no effect.

NOTE 1 A simple program can consist of a single compilation unit. A compilation can have no compilation units; for example, its text can consist of pragmas.

NOTE 2 The designator of a library function cannot be an `operator_symbol`, but a nonlibrary `renaming_declaration` is allowed to rename a library function as an operator. Within a partition, two library subprograms are required to have distinct names and hence cannot overload each other. However, `renaming_declarations` are allowed to define overloaded names for such subprograms, and a locally declared subprogram is allowed to overload a library subprogram. The expanded name `Standard.L` can be used to denote a root library unit `L` (unless the declaration of `Standard` is hidden) since root library unit declarations occur immediately within the declarative region of package `Standard`.

EXAMPLE *Examples of library units:*

```
package Rational_Numbers.IO is -- public child of Rational_Numbers, see 10.1
  procedure Put(R : in Rational);
  procedure Get(R : out Rational);
end Rational_Numbers.IO;

private procedure Rational_Numbers.Reduce(R : in out Rational);
                                -- private child of Rational_Numbers

with Rational_Numbers.Reduce; -- refer to a private child
package body Rational_Numbers is
  ...
end Rational_Numbers;

with Rational_Numbers.IO; use Rational_Numbers;
with Ada.Text_io;          -- see A.10
procedure Main is          -- a root library procedure
  R : Rational;
begin
  R := 5/3;                -- construct a rational number, see 10.1
  Ada.Text_IO.Put("The answer is: ");
  IO.Put(R);
  Ada.Text_IO.New_Line;
end Main;
```

```

with Rational_Numbers.IO;
package Rational_IO renames Rational_Numbers.IO;
                                -- a library unit renaming declaration

```

Each of the above `library_items` can be submitted to the compiler separately.

### 13.1.2 Context Clauses - With Clauses

A `context_clause` is used to specify the `library_items` whose names are needed within a compilation unit.

#### Syntax

```

context_clause ::= {context_item}
context_item  ::= with_clause | use_clause
with_clause   ::= limited_with_clause | nonlimited_with_clause
limited_with_clause ::= limited [private] with library_unit_name {, library_unit_name};
nonlimited_with_clause ::= [private] with library_unit_name {, library_unit_name};

```

#### Name Resolution Rules

The *scope* of a `with_clause` that appears on a `library_unit_declaration` or `library_unit_renaming_declaration` consists of the entire declarative region of the declaration, which includes all children and subunits. The scope of a `with_clause` that appears on a body consists of the body, which includes all subunits.

A `library_item` (and the corresponding library unit) is *named* in a `with_clause` if it is denoted by a `library_unit_name` in the `with_clause`. A `library_item` (and the corresponding library unit) is *mentioned* in a `with_clause` if it is named in the `with_clause` or if it is denoted by a prefix in the `with_clause`.

Outside its own declarative region, the declaration or renaming of a library unit can be visible only within the scope of a `with_clause` that mentions it. The visibility of the declaration or renaming of a library unit otherwise follows from its placement in the environment.

#### Legality Rules

If a `with_clause` of a given `compilation_unit` mentions a private child of some library unit, then the given `compilation_unit` shall be one of:

- the declaration, body, or subunit of a private descendant of that library unit;
- the body or subunit of a public descendant of that library unit, but not a subprogram body acting as a subprogram declaration (see 13.1.4); or
- the declaration of a public descendant of that library unit, in which case the `with_clause` shall include the reserved word **private**.

A name denoting a `library_item` (or the corresponding declaration for a child of a generic within an instance — see 13.1.1), if it is visible only due to being mentioned in one or more `with_clauses` of a unit *U* that include the reserved word **private**, shall appear only within:

- a private part;
- a body of a public descendant of *U*, but not within the `subprogram_specification` of a body of a subprogram that is a public descendant of *U*;
- a private descendant of *U* or its body; or
- a pragma within a context clause.

A `library_item` mentioned in a `limited_with_clause` shall be the implicit declaration of the limited view of a library package, not the declaration of a subprogram, generic unit, generic instance, or a renaming.

A `limited_with_clause` shall not appear on a `library_unit_body`, `subunit`, or `library_unit_renaming_declaration`.

A `limited_with_clause` that names a library package shall not appear:

- in the `context_clause` for the explicit declaration of the named library package or any of its descendants;
- within a `context_clause` for a `library_item` that is within the scope of a `nonlimited_with_clause` that mentions the same library package; or
- within a `context_clause` for a `library_item` that is within the scope of a `use_clause` that names an entity declared within the declarative region of the library package.

NOTE A `library_item` mentioned in a `nonlimited_with_clause` of a compilation unit is visible within the compilation unit and hence acts just like an ordinary declaration. Thus, within a compilation unit that mentions its declaration, the name of a library package can be given in `use_clauses` and can be used to form expanded names, a library subprogram can be called, and instances of a generic library unit can be declared. If a child of a parent generic package is mentioned in a `nonlimited_with_clause`, then the corresponding declaration nested within each visible instance is visible within the compilation unit. Similarly, a `library_item` mentioned in a `limited_with_clause` of a compilation unit is visible within the compilation unit and thus can be used to form expanded names.

EXAMPLE *Examples of use of with clauses, limited with clauses, and private with clauses:*

```
package Office is
end Office;

with Ada.Strings.Unbounded;
package Office.Locations is
  type Location is new Ada.Strings.Unbounded.Unbounded_String;
end Office.Locations;

limited with Office.Departments; -- types are incomplete
private with Office.Locations;  -- only visible in private part
package Office.Employees is
  type Employee is private;

  function Dept_Of(Emp : Employee) return access Departments.Department;
  procedure Assign_Dept(Emp : in out Employee;
                        Dept : access Departments.Department);

  ...
private
  type Employee is
    record
      Dept : access Departments.Department;
      Loc : Locations.Location;
      ...
    end record;
end Office.Employees;

limited with Office.Employees;
package Office.Departments is
  type Department is ...;

  function Manager_Of(Dept : Department) return access Employees.Employee;
  procedure Assign_Manager(Dept : in out Department;
                           Mgr : access Employees.Employee);

  ...
end Office.Departments;
```

The `limited_with_clause` can be used to support mutually dependent abstractions that are split across multiple packages. In this case, an employee is assigned to a department, and a department has a manager who is an employee. If a `with_clause` with the reserved word **private** appears on one library unit and mentions a second library unit, it

*provides visibility to the second library unit, but restricts that visibility to the private part and body of the first unit. The compiler checks that no use is made of the second unit in the visible part of the first unit.*

### 13.1.3 Subunits of Compilation Units

Subunits are like child units, with these (important) differences: subunits support the separate compilation of bodies only (not declarations); the parent contains a `body_stub` to indicate the existence and place of each of its subunits; declarations appearing in the parent's body can be visible within the subunits.

#### Syntax

```
body_stub ::=
  subprogram_body_stub | package_body_stub
  | task_body_stub | protected_body_stub

subprogram_body_stub ::=
  [overriding_indicator]
  subprogram_specification is separate
  [aspect_specification];

package_body_stub ::=
  package body defining_identifier is separate
  [aspect_specification];

task_body_stub ::=
  task body defining_identifier is separate
  [aspect_specification];

protected_body_stub ::=
  protected body defining_identifier is separate
  [aspect_specification];

subunit ::= separate (parent_unit_name) proper_body
```

#### Legality Rules

The *parent body* of a subunit is the body of the program unit denoted by its `parent_unit_name`. The term *subunit* is used to refer to a subunit and also to the `proper_body` of a subunit. The *subunits of a program unit* include any subunit that names that program unit as its parent, as well as any subunit that names such a subunit as its parent (recursively).

The parent body of a subunit shall be present in the current environment, and shall contain a corresponding `body_stub` with the same `defining_identifier` as the subunit.

A `package_body_stub` shall be the completion of a `package_declaration` or `generic_package_declaration`; a `task_body_stub` shall be the completion of a `task_declaration`; a `protected_body_stub` shall be the completion of a `protected_declaration`.

In contrast, a `subprogram_body_stub` can be defined without it being the completion of a previous declaration, in which case the `_stub` declares the subprogram. If the `_stub` is a completion, it shall be the completion of a `subprogram_declaration` or `generic_subprogram_declaration`. The profile of a `subprogram_body_stub` that completes a declaration shall conform fully to that of the declaration.

A subunit that corresponds to a `body_stub` shall be of the same kind (`package_`, `subprogram_`, `task_`, or `protected_`) as the `body_stub`. The profile of a `subprogram_body` subunit shall be fully conformant to that of the corresponding `body_stub`.

A `body_stub` shall appear immediately within the `declarative_part` of a compilation unit body. This rule does not apply within an instance of a generic unit.

The `defining_identifiers` of all `body_stubs` that appear immediately within a particular `declarative_part` shall be distinct.



## Post-Compilation Rules

For each `body_stub`, there shall be a subunit containing the corresponding `proper_body`.

NOTE The rules in 13.1.4 say that a `body_stub` is equivalent to the corresponding `proper_body`. This implies:

- Visibility within a subunit is the visibility that would be obtained at the place of the corresponding `body_stub` (within the parent body) if the `context_clause` of the subunit were appended to that of the parent body.
- The effect of the elaboration of a `body_stub` is to elaborate the subunit.

EXAMPLE 1 *Example that defines package Parent without subunits:*

```
package Parent is
  procedure Inner;
end Parent;

with Ada.Text_IO;
package body Parent is
  Variable : String := "Hello, there.";
  procedure Inner is
  begin
    Ada.Text_IO.Put_Line(Variable);
  end Inner;
end Parent;
```

EXAMPLE 2 *Example showing how the body of procedure Inner can be turned into a subunit by rewriting the package body as follows (with the declaration of Parent remaining the same):*

```
package body Parent is
  Variable : String := "Hello, there.";
  procedure Inner is separate;
end Parent;

with Ada.Text_IO;
separate(Parent)
procedure Inner is
begin
  Ada.Text_IO.Put_Line(Variable);
end Inner;
```

### 13.1.4 The Compilation Process

Each compilation unit submitted to the compiler is compiled in the context of an *environment* `declarative_part` (or simply, an *environment*), which is a conceptual `declarative_part` that forms the outermost declarative region of the context of any compilation. At run time, an environment forms the `declarative_part` of the body of the environment task of a partition (see 13.2).

The `declarative_items` of the environment are `library_items` appearing in an order such that there are no forward semantic dependences. Each included subunit occurs in place of the corresponding stub. The visibility rules apply as if the environment were the outermost declarative region, except that `with_clauses` are necessary to make declarations of library units visible (see 13.1.2).

The mechanisms for creating an environment and for adding and replacing compilation units within an environment are implementation defined. The mechanisms for adding a compilation unit mentioned in a `limited_with_clause` to an environment are implementation defined.

## Name Resolution Rules

If a `library_unit_body` that is a `subprogram_body` is submitted to the compiler, it is interpreted only as a completion if a `library_unit_declaration` with the same `defining_program_unit_name` already exists in the environment for a subprogram other than an instance of a generic subprogram or for a generic subprogram (even if the profile of the body is not type

conformant with that of the declaration); otherwise, the `subprogram_body` is interpreted as both the declaration and body of a library subprogram.

#### *Legality Rules*

When a compilation unit is compiled, all compilation units upon which it depends semantically shall already exist in the environment; the set of these compilation units shall be *consistent* in the sense that the new compilation unit shall not semantically depend (directly or indirectly) on two different versions of the same compilation unit, nor on an earlier version of itself.

#### *Implementation Permissions*

The implementation may require that a compilation unit be legal before it can be mentioned in a `limited_with_clause` or it can be inserted into the environment.

When a compilation unit that declares or renames a library unit is added to the environment, the implementation may remove from the environment any preexisting `library_item` or subunit with the same full expanded name. When a compilation unit that is a subunit or the body of a library unit is added to the environment, the implementation may remove from the environment any preexisting version of the same compilation unit. When a compilation unit that contains a `body_stub` is added to the environment, the implementation may remove any preexisting `library_item` or subunit with the same full expanded name as the `body_stub`. When a given compilation unit is removed from the environment, the implementation may also remove any compilation unit that depends semantically upon the given one. If the given compilation unit contains the body of a subprogram for which `aspect Inline` is `True`, the implementation may also remove any compilation unit containing a call to that subprogram.

NOTE 1 The rules of the language are enforced across compilation and compilation unit boundaries, just as they are enforced within a single compilation unit.

NOTE 2 An implementation can support a concept of a *library*, which contains `library_items`. If multiple libraries are supported, the implementation can document how a single environment is constructed when a compilation unit is submitted to the compiler. Naming conflicts between different libraries can, for example, be resolved by treating each library as the root of a hierarchy of child library units.

NOTE 3 A compilation unit containing an instantiation of a separately compiled generic unit does not semantically depend on the body of the generic unit. Therefore, replacing the generic body in the environment does not result in the removal of the compilation unit containing the instantiation.

### 13.1.5 Pragmas and Program Units

This subclause discusses pragmas related to compilations.

#### *Post-Compilation Rules*

Certain pragmas are defined to be *configuration pragmas*; they shall appear before the first `compilation_unit` of a compilation. They are generally used to select a partition-wide or system-wide option. The pragma applies to all `compilation_units` appearing in the compilation, unless there are none, in which case it applies to all future `compilation_units` compiled into the same environment.

#### *Implementation Permissions*

An implementation may require that configuration pragmas that select partition-wide or system-wide options be compiled when the environment contains no `library_items` other than those of the predefined environment. In this case, the implementation shall still accept configuration pragmas in individual compilations that confirm the initially selected partition-wide or system-wide options.

### 13.1.6 Environment-Level Visibility Rules

The normal visibility rules do not apply within a `parent_unit_name` or a `context_clause`, nor within a `pragma` that appears at the place of a compilation unit. The special visibility rules for those contexts are given here.

#### Static Semantics

Within the `parent_unit_name` at the beginning of an explicit `library_item`, and within a `nonlimited_with_clause`, the only declarations that are visible are those that are explicit `library_items` of the environment, and the only declarations that are directly visible are those that are explicit root `library_items` of the environment. Within a `limited_with_clause`, the only declarations that are visible are those that are the implicit declaration of the limited view of a library package of the environment, and the only declarations that are directly visible are those that are the implicit declaration of the limited view of a root library package.

Within a `use_clause` or `pragma` that is within a `context_clause`, each `library_item` mentioned in a previous `with_clause` of the same `context_clause` is visible, and each root `library_item` so mentioned is directly visible. In addition, within such a `use_clause`, if a given declaration is visible or directly visible, each declaration that occurs immediately within the given declaration's visible part is also visible. No other declarations are visible or directly visible.

Within the `parent_unit_name` of a subunit, `library_items` are visible as they are in the `parent_unit_name` of a `library_item`; in addition, the declaration corresponding to each `body_stub` in the environment is also visible.

Within a `pragma` that appears at the place of a compilation unit, the immediately preceding `library_item` and each of its ancestors is visible. The ancestor root `library_item` is directly visible.

Notwithstanding the rules of 7.1.3, an expanded name in a `with_clause`, a `pragma` in a `context_clause`, or a `pragma` that appears at the place of a compilation unit may consist of a prefix that denotes a generic package and a `selector_name` that denotes a child of that generic package. (The child is necessarily a generic unit; see 13.1.1.)

## 13.2 Program Execution

An Ada *program* consists of a set of *partitions*, which can execute in parallel with one another, possibly in a separate address space, and possibly on a separate computer.

#### Post-Compilation Rules

A partition is a program or part of a program that can be invoked from outside the Ada implementation. For example, on many systems, a partition can be an executable file generated by the system linker. The user can *explicitly assign* library units to a partition. The assignment is done in an implementation-defined manner. The compilation units included in a partition are those of the explicitly assigned library units, as well as other compilation units *needed by* those library units. The compilation units needed by a given compilation unit (the *needed compilation units*) are determined as follows (unless specified otherwise via an implementation-defined `pragma`, or by some other implementation-defined means):

- A compilation unit is a needed compilation unit of itself;
- If a compilation unit is among the needed compilation units, then so are any compilation units upon which it depends semantically;
- If a `library_unit_declaration` is among the needed compilation units, then so is any corresponding `library_unit_body`;
- If a compilation unit with stubs is among the needed compilation units, then so are any corresponding subunits;

- If the (implicit) declaration of the limited view of a library package is among the needed compilation units, then so is the explicit declaration of the library package.

The user can optionally designate (in an implementation-defined manner) one subprogram as the *main subprogram* for the partition. A main subprogram, if specified, shall be a subprogram.

Each partition has an anonymous *environment task*, which is an implicit outermost task whose execution elaborates the `library_items` of the `environment declarative_part`, and then calls the main subprogram, if there is one. A partition's execution is that of its tasks.

The order of elaboration of library units is determined primarily by the *elaboration dependences*. There is an elaboration dependence of a given `library_item` upon another if the given `library_item` or any of its subunits depends semantically on the other `library_item`. In addition, if a given `library_item` or any of its subunits has a pragma `Elaborate` or `Elaborate_All` that names another library unit, then there is an elaboration dependence of the given `library_item` upon the body of the other library unit, and, for `Elaborate_All` only, upon each `library_item` that is a needed compilation unit of the declaration of the other library unit.

The environment task for a partition has the following structure:

```
task Environment_Task;
task body Environment_Task is
... (1) -- The environment declarative_part
        -- (that is, the sequence of library_items) goes here.
begin
... (2) -- Call the main subprogram, if there is one.
end Environment_Task;
```

The environment `declarative_part` at (1) is a sequence of `declarative_items` consisting of copies of the `library_items` included in the partition. The order of elaboration of `library_items` is the order in which they appear in the `environment declarative_part`:

- The order of all included `library_items` is such that there are no forward elaboration dependences.
- Any included `library_unit` declaration for which `aspect Elaborate_Body` is `True` is immediately followed by its `library_unit_body`, if included.
- All `library_items` declared `pure` occur before any that are not declared `pure`.
- All preelaborated `library_items` occur before any that are not preelaborated.

There shall be a total order of the `library_items` that obeys the above rules. The order is otherwise implementation defined.

The full expanded names of the library units and subunits included in a given partition shall be distinct.

The `sequence_of_statements` of the environment task (see (2) above) consists of either:

- A call to the main subprogram, if the partition has one. If the main subprogram has parameters, they are passed; where the actuals come from is implementation defined. What happens to the result of a main function is also implementation defined.

or:

- A `null_statement`, if there is no main subprogram.

The mechanisms for building and running partitions are implementation defined. These can be combined into one operation, as, for example, in dynamic linking, or “load-and-go” systems.

#### Dynamic Semantics

The execution of a program consists of the execution of a set of partitions. Further details are implementation defined. The execution of a partition starts with the execution of its environment task, ends when the environment task terminates, and includes the executions of all tasks of the partition. The execution of the (implicit) `task_body` of the environment task acts as a master for all other tasks

created as part of the execution of the partition. When the environment task completes (normally or abnormally), it waits for the termination of all such tasks, and then finalizes any remaining objects of the partition.

#### *Bounded (Run-Time) Errors*

Once the environment task has awaited the termination of all other tasks of the partition, any further attempt to create a task (during finalization) is a bounded error, and may result in the raising of `Program_Error` either upon creation or activation of the task. If such a task is activated, it is not specified whether the task is awaited prior to termination of the environment task.

#### *Implementation Requirements*

The implementation shall ensure that all compilation units included in a partition are consistent with one another, and are legal according to the rules of the language.

#### *Implementation Permissions*

The kind of partition described in this subclause is known as an *active* partition. An implementation is allowed to support other kinds of partitions, with implementation-defined semantics.

An implementation may restrict the kinds of subprograms it supports as main subprograms. However, an implementation is required to support all main subprograms that are public parameterless library procedures.

If the environment task completes abnormally, the implementation may abort any dependent tasks.

NOTE 1 An implementation can provide inter-partition communication mechanism(s) via special packages and pragmas. Standard pragmas for distribution and methods for specifying inter-partition communication are defined in Annex E. If no such mechanisms are provided, then each partition is isolated from all others, and behaves as a program in and of itself.

NOTE 2 Partitions are not required to run in separate address spaces. For example, an implementation can support dynamic linking via the partition concept.

NOTE 3 An order of elaboration of `library_items` that is consistent with the partial ordering defined above does not always ensure that each `library_unit_body` is elaborated before any other compilation unit whose elaboration necessitates that the `library_unit_body` be already elaborated. (In particular, there is no requirement that the body of a library unit be elaborated as soon as possible after the `library_unit_declaration` is elaborated, unless the pragmas or aspects in 13.2.1 are used.)

NOTE 4 A partition (active or otherwise) does not necessarily have a main subprogram. In such a case, all the work done by the partition would be done by elaboration of various `library_items`, and by tasks created by that elaboration. Passive partitions, which cannot have main subprograms, are defined in Annex E.

### 13.2.1 Elaboration Control

This subclause defines aspects and pragmas that help control the elaboration order of `library_items`.

#### *Legality Rules*

An elaborable construct is preelaborable unless its elaboration performs any of the following actions:

- The execution of a `statement` other than a `null_statement`.
- A call to a subprogram other than:
  - a static function;
  - an instance of `Unchecked_Conversion` (see 16.9);
  - a function declared in `System.Storage_Elements` (see 16.7.1); or
  - the functions `To_Pointer` and `To_Address` declared in an instance of `System.Address_to_Access_Conversions` (see 16.7.2).



- The evaluation of a `primary` that is a `name` of an object, unless the `name` is a static expression, or statically denotes a discriminant of an enclosing type.
- The creation of an object (including a component) that is initialized by default, if its type does not have preelaborable initialization. Similarly, the evaluation of an `extension_aggregate` with an ancestor `subtype_mark` denoting a subtype of such a type.
- The elaboration of any elaborable construct that is not preelaborable.

A generic declaration is preelaborable unless every instance would perform one of the above actions.

A generic body is preelaborable only if elaboration of a corresponding instance body would not perform any such actions, presuming that:

- the actual for each discriminated formal derived type, formal private type, or formal private extension declared within the formal part of the generic unit is a type that does not have preelaborable initialization, unless the `Preelaborable_Initialization` aspect was specified for the formal type;
- the actual for each formal type is nonstatic;
- the actual for each formal object is nonstatic; and
- the actual for each formal subprogram is a user-defined subprogram.

When the library unit aspect (see 16.1.1) `Preelaborate` of a program unit is `True`, the unit is said to be *preelaborated*. When the `Preelaborate` aspect is specified `True` for a library unit, all compilation units of the library unit are preelaborated. The declaration and body of a preelaborated library unit, and all subunits that are elaborated as part of elaborating the library unit, shall be preelaborable. All compilation units of a preelaborated library unit shall depend semantically only on declared pure or preelaborated `library_items`. In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit. If a library unit is preelaborated, then its declaration, if any, and body, if any, are elaborated prior to all nonpreelaborated `library_items` of the partition.

The following rules specify which entities have *preelaborable initialization*, namely that the `Preelaborable_Initialization` aspect of the entity is `True`:

- The partial view of a private type or private extension, a protected type without `entry_declarations`, a generic formal private type, or a generic formal derived type, has preelaborable initialization if and only if the `Preelaborable_Initialization` aspect has been specified `True` for them. A protected type with `entry_declarations` or a task type never has preelaborable initialization. The `Preelaborable_Initialization` aspect of a partial view of a type may be specified as `False`, even if the full view of the type has preelaborable initialization. Similarly, a generic formal type may be specified with `Preelaborable_Initialization` `False`, even if the actual type in an instance has preelaborable initialization.
- A component (including a discriminant) of a record or protected type has preelaborable initialization if its declaration includes a `default_expression` whose execution does not perform any actions prohibited in preelaborable constructs as described above, or if its declaration does not include a default expression and its type has preelaborable initialization.
- A derived type has preelaborable initialization if its parent type has preelaborable initialization and if the noninherited components all have preelaborable initialization. However, a controlled type with an `Initialize` procedure that is not a null procedure does not have preelaborable initialization.
- A view of a type has preelaborable initialization if it is an elementary type, an array type whose component type has preelaborable initialization, a record type whose components all have preelaborable initialization, or an interface type.

The following attribute is defined for a nonformal composite subtype `S` declared within the visible part of a package or a generic package, or a generic formal private subtype or formal derived subtype:



**S'Preelaborable\_Initialization**

This attribute is of Boolean type, and its value reflects whether the type of S has preelaborable initialization. The value of this attribute, the type-related `Preelaborable_Initialization` aspect, may be specified for any type for which the attribute is defined. The value shall be specified by a static expression, unless the type is not a formal type but is nevertheless declared within a generic package. In this latter case, the value may also be specified by references to the `Preelaborable_Initialization` attribute of one or more formal types visible at the point of the declaration of the composite type, conjoined with **and**.

If the `Preelaborable_Initialization` aspect is specified True for a private type or a private extension, the full view of the type shall have preelaborable initialization. If the aspect is specified True for a protected type, the protected type shall not have entries, and each component of the protected type shall have preelaborable initialization. If the aspect is specified True for a generic formal type, then in a `generic_instantiation` the corresponding actual type shall have preelaborable initialization. If the aspect definition includes one or more `Preelaborable_Initialization` attribute references, then the full view of the type shall have preelaborable initialization presuming the types mentioned in the prefixes of the attribute references all have preelaborable initialization. For any other composite type, the aspect shall be specified statically True or False only if it is confirming. In addition to the places where Legality Rules normally apply (see 15.3), these rules apply also in the private part of an instance of a generic unit.

*Implementation Advice*

In an implementation, a type declared in a preelaborated package should have the same representation in every elaboration of a given version of the package, whether the elaborations occur in distinct executions of the same program, or in executions of distinct programs or partitions that include the given version.

*Static Semantics*

A *pure* program unit is a preelaborable program unit whose elaboration does not perform any of the following actions:

- the elaboration of a variable declaration;
- the evaluation of an `allocator` of an access-to-variable type; for the purposes of this rule, the partial view of a type is presumed to have nonvisible components whose default initialization evaluates such an `allocator`;
- the elaboration of the declaration of a nonderived named access-to-variable type unless the `Storage_Size` of the type has been specified by a static expression with value zero or is defined by the language to be zero;
- the elaboration of the declaration of a nonderived named access-to-constant type for which the `Storage_Size` has been specified by an expression other than a static expression with value zero;
- the elaboration of any program unit that is not pure.

A generic declaration is pure unless every instance would perform one of the above actions.

A generic body is pure only if elaboration of a corresponding instance body would not perform any such actions presuming any composite formal types have nonvisible components whose default initialization evaluates an `allocator` of an access-to-variable type.

The `Storage_Size` for an anonymous access-to-variable type declared at library level in a library unit that is declared pure is defined to be zero.

*Legality Rules*

When the library unit aspect `Pure` of a program unit is True, the unit is said to be *declared pure*. When the `Pure` aspect is specified True for a library unit, all compilation units of the library unit are declared pure. In addition, the limited view of any library package is declared pure. The declaration and body of

a declared pure library unit, and all subunits that are elaborated as part of elaborating the library unit, shall be pure. All compilation units of a declared pure library unit shall depend semantically only on declared pure `library_items`. In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit. Furthermore, the full view of any partial view declared in the visible part of a declared pure library unit that has any available stream attributes shall support external streaming (see 16.13.2).

#### *Erroneous Execution*

Execution is erroneous if some operation (other than the initialization or finalization of the object) modifies the value of a constant object declared at library-level in a pure package.

#### *Implementation Permissions*

If a library unit is declared pure, then the implementation is permitted to omit a call on a library-level subprogram of the library unit if the results are not needed after the call. In addition, the implementation may omit a call on such a subprogram and simply reuse the results produced by an earlier call on the same subprogram, provided that none of the parameters nor any object accessible via access values from the parameters have any part that is of a type whose full type is an immutably limited type, and the addresses and values of all by-reference actual parameters, the values of all by-copy-in actual parameters, and the values of all objects accessible via access values from the parameters, are the same as they were at the earlier call. This permission applies even if the subprogram produces other side effects when called.

#### *Syntax*

The following `pragmas` are defined with the given forms:

**pragma** Elaborate(*library\_unit\_name*{, *library\_unit\_name*});

**pragma** Elaborate\_All(*library\_unit\_name*{, *library\_unit\_name*});

A `pragma` Elaborate or Elaborate\_All is only allowed within a `context_clause`.

#### *Legality Rules*

If the aspect `Elaborate_Body` is `True` for a declaration, then the declaration requires a completion (a body).

The *library\_unit\_name* of a `pragma` Elaborate or Elaborate\_All shall denote a nonlimited view of a library unit.

#### *Static Semantics*

A `pragma` Elaborate specifies that the body of the named library unit is elaborated before the current `library_item`. A `pragma` Elaborate\_All specifies that each `library_item` that is needed by the named library unit declaration is elaborated before the current `library_item`.

If the `Elaborate_Body` aspect of a library unit is `True`, the body of the library unit is elaborated immediately after its declaration.

NOTE 1 A preelaborated library unit can have nonpreelaborable children.

NOTE 2 A library unit that is declared pure can have impure children.

## 14 Exceptions

This clause defines the facilities for dealing with errors or other exceptional situations that arise during program execution. An *exception* represents a kind of exceptional situation; an occurrence of such a situation (at run time) is called an *exception occurrence*. To *raise* an exception is to abandon normal program execution so as to draw attention to the fact that the corresponding situation has arisen. Performing some actions in response to the arising of an exception is called *handling* the exception.

An `exception_declaration` declares a name for an exception. An exception can be raised explicitly (for example, by a `raise_statement`) or implicitly (for example, by the failure of a language-defined

check). When an exception arises, control can be transferred to a user-provided `exception_handler` at the end of a `handled_sequence_of_statements`, or it can be propagated to a dynamically enclosing execution.

## 14.1 Exception Declarations

An `exception_declaration` declares a name for an exception.

### Syntax

```
exception_declaration ::= defining_identifier_list: exception
[aspect_specification];
```

### Static Semantics

Each single `exception_declaration` declares a name for a different exception. If a generic unit includes an `exception_declaration`, the `exception_declarations` implicitly generated by different instantiations of the generic unit refer to distinct exceptions (but all have the same `defining_identifier`). The particular exception denoted by an exception name is determined at compilation time and is the same regardless of how many times the `exception_declaration` is elaborated.

The *predefined* exceptions are the ones declared in the declaration of package `Standard`: `Constraint_Error`, `Program_Error`, `Storage_Error`, and `Tasking_Error`; one of them is raised when a language-defined check fails.

### Dynamic Semantics

The elaboration of an `exception_declaration` has no effect.

The execution of any construct raises `Storage_Error` if there is insufficient storage for that execution. The amount of storage necessary for the execution of constructs is unspecified.

EXAMPLE *Examples of user-defined exception declarations:*

```
Singular : exception;
Error    : exception;
Overflow, Underflow : exception;
```

## 14.2 Exception Handlers

The response to one or more exceptions is specified by an `exception_handler`.

### Syntax

```
handled_sequence_of_statements ::=
sequence_of_statements
[exception
exception_handler
{exception_handler}]
exception_handler ::=
when [choice_parameter_specification:]
exception_choice {'|' exception_choice} =>
sequence_of_statements
choice_parameter_specification ::= defining_identifier
exception_choice ::= exception_name | others
```

### Legality Rules

An *exception\_name* of an `exception_choice` shall denote an exception.

A choice with an *exception\_name* covers the named exception. A choice with **others** covers all exceptions not named by previous choices of the same `handled_sequence_of_statements`. Two

choices in different `exception_handlers` of the same `handled_sequence_of_statements` shall not cover the same exception.

A choice with **others** is allowed only for the last handler of a `handled_sequence_of_statements` and as the only choice of that handler.

An *exception\_name* of a choice shall not denote an exception declared in a generic formal package.

#### Static Semantics

A *choice\_parameter\_specification* declares a *choice parameter*, which is a constant object of type `Exception_Occurrence` (see 14.4.1). During the handling of an exception occurrence, the choice parameter, if any, of the handler represents the exception occurrence that is being handled.

#### Dynamic Semantics

The execution of a `handled_sequence_of_statements` consists of the execution of the `sequence_of_statements`. The optional handlers are used to handle any exceptions that are propagated by the `sequence_of_statements`.

EXAMPLE *Example of an exception handler:*

```
begin
  Open(File, In_File, "input.txt");    -- see A.8.2
exception
  when E : Name_Error =>
    Put("Cannot open input file : ");
    Put_Line(Ada.Exceptions.Exception_Message(E)); -- see 14.4.1
  raise;
end;
```

### 14.3 Raise Statements and Raise Expressions

A *raise\_statement* raises an exception.

#### Syntax

```
raise_statement ::= raise;
                  | raise exception_name [with string_expression];
raise_expression ::= raise exception_name [with string_simple_expression]
```

If a *raise\_expression* appears within the expression of one of the following contexts, the *raise\_expression* shall appear within a pair of parentheses within the expression:

- *object\_declaration*;
- *modular\_type\_definition*;
- *floating\_point\_definition*;
- *ordinary\_fixed\_point\_definition*;
- *decimal\_fixed\_point\_definition*;
- *default\_expression*;
- *ancestor\_part*.

#### Legality Rules

The *exception\_name*, if any, of a *raise\_statement* or *raise\_expression* shall denote an exception. A *raise\_statement* with no *exception\_name* (that is, a *re-raise statement*) shall be within a handler, but not within a body enclosed by that handler.

## Name Resolution Rules

The *string\_expression* or *string\_simple\_expression*, if any, of a *raise\_statement* or *raise\_expression* is expected to be of type *String*.

The expected type for a *raise\_expression* shall be any single type.

## Dynamic Semantics

To *raise an exception* is to raise a new occurrence of that exception, as explained in 14.4. For the execution of a *raise\_statement* with an *exception\_name*, the named exception is raised. Similarly, for the evaluation of a *raise\_expression*, the named exception is raised. In both of these cases, if a *string\_expression* or *string\_simple\_expression* is present, the expression is evaluated and its value is associated with the exception occurrence. For the execution of a re-raise statement, the exception occurrence that caused transfer of control to the innermost enclosing handler is raised again.

NOTE If the evaluation of a *string\_expression* or *string\_simple\_expression* raises an exception, that exception is propagated instead of the one denoted by the *exception\_name* of the *raise\_statement* or *raise\_expression*.

EXAMPLE *Examples of raise statements:*

```
raise Ada.IO_Exceptions.Name_Error;    -- see A.13
raise Queue_Error with "Buffer Full";  -- see 12.11

raise;                                -- re-raise the current exception

-- For an example of a raise expression, see the Streams Subsystem definitions
in 16.13.1.
```

## 14.4 Exception Handling

When an exception occurrence is raised, normal program execution is abandoned and control is transferred to an applicable *exception\_handler*, if any. To *handle* an exception occurrence is to respond to the exceptional event. To *propagate* an exception occurrence is to raise it again in another context; that is, to fail to respond to the exceptional event in the present context.

## Dynamic Semantics

Within a given task, if the execution of construct *a* is defined by this document to consist (in part) of the execution of construct *b*, then while *b* is executing, the execution of *a* is said to *dynamically enclose* the execution of *b*. The *innermost dynamically enclosing* execution of a given execution is the dynamically enclosing execution that started most recently.

When an exception occurrence is raised by the execution of a given construct, the rest of the execution of that construct is *abandoned*; that is, any portions of the execution that have not yet taken place are not performed. The construct is first completed, and then left, as explained in 10.6.1. Then:

- If the construct is a *task\_body*, the exception does not propagate further;
- If the construct is the *sequence\_of\_statements* of a *handled\_sequence\_of\_statements* that has a handler with a choice covering the exception, the occurrence is handled by that handler;
- Otherwise, the occurrence is *propagated* to the innermost dynamically enclosing execution, which means that the occurrence is raised again in that context.

When an occurrence is *handled* by a given handler, the *choice\_parameter\_specification*, if any, is first elaborated, which creates the choice parameter and initializes it to the occurrence. Then, the *sequence\_of\_statements* of the handler is executed; this execution replaces the abandoned portion of the execution of the *sequence\_of\_statements*.

NOTE Exceptions raised in a *declarative\_part* of a body are not handled by the handlers of the *handled\_sequence\_of\_statements* of that body.

### 14.4.1 The Package Exceptions

#### Static Semantics

The following language-defined library package exists:

```

with Ada.Streams;
package Ada.Exceptions
  with Preelaborate, Nonblocking, Global => in out synchronized is
  type Exception_Id is private
    with Preelaborable_Initialization;
  Null_Id : constant Exception_Id;
  function Exception_Name(Id : Exception_Id) return String;
  function Wide_Exception_Name(Id : Exception_Id) return Wide_String;
  function Wide_Wide_Exception_Name(Id : Exception_Id)
    return Wide_Wide_String;

  type Exception_Occurrence is limited private
    with Preelaborable_Initialization;
  type Exception_Occurrence_Access is access all Exception_Occurrence;
  Null_Occurrence : constant Exception_Occurrence;

  procedure Raise_Exception(E : in Exception_Id;
    Message : in String := "");

    with No_Return;
  function Exception_Message(X : Exception_Occurrence) return String;
  procedure Reraise_Occurrence(X : in Exception_Occurrence);

  function Exception_Identity(X : Exception_Occurrence)
    return Exception_Id;
  function Exception_Name(X : Exception_Occurrence) return String;
    -- Same as Exception_Name(Exception_Identity(X)).
  function Wide_Exception_Name(X : Exception_Occurrence)
    return Wide_String;
    -- Same as Wide_Exception_Name(Exception_Identity(X)).
  function Wide_Wide_Exception_Name(X : Exception_Occurrence)
    return Wide_Wide_String;
    -- Same as Wide_Wide_Exception_Name(Exception_Identity(X)).
  function Exception_Information(X : Exception_Occurrence) return String;

  procedure Save_Occurrence(Target : out Exception_Occurrence;
    Source : in Exception_Occurrence);
  function Save_Occurrence(Source : Exception_Occurrence)
    return Exception_Occurrence_Access;

  procedure Read_Exception_Occurrence
    (Stream : not null access Ada.Streams.Root_Stream_Type'Class;
    Item : out Exception_Occurrence);
  procedure Write_Exception_Occurrence
    (Stream : not null access Ada.Streams.Root_Stream_Type'Class;
    Item : in Exception_Occurrence);

  for Exception_Occurrence'Read use Read_Exception_Occurrence;
  for Exception_Occurrence'Write use Write_Exception_Occurrence;

private
  ... -- not specified by the language
end Ada.Exceptions;

```

Each distinct exception is represented by a distinct value of type `Exception_Id`. `Null_Id` does not represent any exception, and is the default initial value of type `Exception_Id`. Each occurrence of an exception is represented by a value of type `Exception_Occurrence`. `Null_Occurrence` does not represent any exception occurrence, and is the default initial value of type `Exception_Occurrence`.

For a prefix `E` that denotes an exception, the following attribute is defined:

`E'Identity` `E'Identity` returns the unique identity of the exception. The type of this attribute is `Exception_Id`.



`Raise_Exception` raises a new occurrence of the identified exception.

`Exception_Message` returns the message associated with the given `Exception_Occurrence`. For an occurrence raised by a call to `Raise_Exception`, the message is the `Message` parameter passed to `Raise_Exception`. For the occurrence raised by a `raise_statement` or `raise_expression` with an `exception_name` and a `string_expression` or `string_simple_expression`, the message is the `string_expression` or `string_simple_expression`. For the occurrence raised by a `raise_statement` or `raise_expression` with an `exception_name` but without a `string_expression` or `string_simple_expression`, the message is a string giving implementation-defined information about the exception occurrence. For an occurrence originally raised in some other manner (including by the failure of a language-defined check), the message is an unspecified string. In all cases, `Exception_Message` returns a string with lower bound 1.

`Reraise_Occurrence` reraises the specified exception occurrence.

`Exception_Identity` returns the identity of the exception of the occurrence.

The `Wide_Wide_Exception_Name` functions return the full expanded name of the exception, in upper case, starting with a root library unit. For an exception declared immediately within package `Standard`, the `defining_identifier` is returned. The result is implementation defined if the exception is declared within an unnamed `block_statement`.

The `Exception_Name` functions (respectively, `Wide_Exception_Name`) return the same sequence of graphic characters as that defined for `Wide_Wide_Exception_Name`, if all the graphic characters are defined in `Character` (respectively, `Wide_Character`); otherwise, the sequence of characters is implementation defined, but no shorter than that returned by `Wide_Wide_Exception_Name` for the same value of the argument.

The string returned by the `Exception_Name`, `Wide_Exception_Name`, and `Wide_Wide_Exception_Name` functions has lower bound 1.

`Exception_Information` returns implementation-defined information about the exception occurrence. The returned string has lower bound 1.

`Reraise_Occurrence` has no effect in the case of `Null_Occurrence`. `Raise_Exception` and `Exception_Name` raise `Constraint_Error` for a `Null_Id`. `Exception_Message`, `Exception_Name`, and `Exception_Information` raise `Constraint_Error` for a `Null_Occurrence`. `Exception_Identity` applied to `Null_Occurrence` returns `Null_Id`.

The `Save_Occurrence` procedure copies the `Source` to the `Target`. The `Save_Occurrence` function uses an `allocator` of type `Exception_Occurrence_Access` to create a new object, copies the `Source` to this new object, and returns an access value designating this new object; the result may be deallocated using an instance of `Unchecked_Deallocation`.

`Write_Exception_Occurrence` writes a representation of an exception occurrence to a stream; `Read_Exception_Occurrence` reconstructs an exception occurrence from a stream (including one written in a different partition).

#### *Implementation Permissions*

An implementation of `Exception_Name` in a space-constrained environment may return the `defining_identifier` instead of the full expanded name.

The string returned by `Exception_Message` may be truncated (to no less than 200 characters) by the `Save_Occurrence` procedure (not the function), the `Reraise_Occurrence` procedure, and the re-raise statement.

#### *Implementation Advice*

`Exception_Message` (by default) and `Exception_Information` should produce information useful for debugging. `Exception_Message` should be short (about one line), whereas `Exception_Information` can be

long. `Exception_Message` should not include the `Exception_Name`. `Exception_Information` should include both the `Exception_Name` and the `Exception_Message`.

NOTE UTF-8 encoding (see A.4.11) can be used to represent non-ASCII characters in exception messages.

#### 14.4.2 Pragma Assert and Assertion\_Policy

Pragma Assert is used to assert the truth of a boolean expression at a point within a sequence of declarations or statements.

Assert pragmas, subtype predicates (see 6.2.4), preconditions and postconditions (see 9.1.1), type invariants (see 10.3.2), and default initial conditions (see 10.3.3) are collectively referred to as *assertions*; their boolean expressions are referred to as *assertion expressions*.

Pragma `Assertion_Policy` is used to control whether assertions are to be ignored by the implementation, checked at run time, or handled in some implementation-defined manner.

##### Syntax

The form of a pragma Assert is as follows:

```
pragma Assert([Check=>] boolean_expression [, [Message=>] string_expression]);
```

A pragma Assert is allowed at the place where a `declarative_item` or a statement is allowed.

The form of a pragma `Assertion_Policy` is as follows:

```
pragma Assertion_Policy(policy_identifier);  
pragma Assertion_Policy(  
    assertion_aspect_mark => policy_identifier  
    {, assertion_aspect_mark => policy_identifier});
```

A pragma `Assertion_Policy` is allowed only immediately within a `declarative_part`, immediately within a `package_specification`, or as a configuration pragma.

##### Name Resolution Rules

The expected type for the *boolean\_expression* of a pragma Assert is any boolean type. The expected type for the *string\_expression* of a pragma Assert is type String.

##### Legality Rules

The *assertion\_aspect\_mark* of a pragma `Assertion_Policy` shall identify an *assertion aspect*, namely one of `Assert`, `Static_Predicate`, `Dynamic_Predicate`, `Pre`, `Pre'Class`, `Post`, `Post'Class`, `Type_Invariant`, `Type_Invariant'Class`, `Default_Initial_Condition`, or some implementation-defined (assertion) *aspect\_mark*. The *policy\_identifier* shall be either `Check`, `Ignore`, or some implementation-defined identifier.

##### Static Semantics

A pragma `Assertion_Policy` determines for each assertion aspect named in the `pragma_argument_associations` whether assertions of the given aspect are to be enforced by a runtime check. The *policy\_identifier* `Check` requires that assertion expressions of the given aspect be checked that they evaluate to True at the points specified for the given aspect; the *policy\_identifier* `Ignore` requires that the assertion expression not be evaluated at these points, and the runtime checks not be performed. Note that for subtype predicate aspects (see 6.2.4), even when the applicable `Assertion_Policy` is `Ignore`, the predicate will still be evaluated as part of membership tests and `Valid attribute_references`, and if static, will still have an effect on loop iteration over the subtype, and the selection of `case_statement_alternatives` and variants.

If no *assertion\_aspect\_marks* are specified in the pragma, the specified policy applies to all assertion aspects.

A `pragma Assertion_Policy` applies to the named assertion aspects in a specific region, and applies to all assertion expressions associated with those aspects specified in that region. A `pragma Assertion_Policy` given in a `declarative_part` or immediately within a `package_specification` applies from the place of the `pragma` to the end of the innermost enclosing declarative region. The region for a `pragma Assertion_Policy` given as a configuration `pragma` is the declarative region for the entire compilation unit (or units) to which it applies.

If a `pragma Assertion_Policy` applies to a `generic_instantiation`, then the `pragma Assertion_Policy` applies to the entire instance.

If multiple `Assertion_Policy` `pragmas` apply to a given construct for a given assertion aspect, the assertion policy is determined by the one in the innermost enclosing region of a `pragma Assertion_Policy` specifying a policy for the assertion aspect. If no such `Assertion_Policy` `pragma` exists, the policy is implementation defined.

The following language-defined library package exists:

```
package Ada.Assertions
  with Pure is

  Assertion_Error : exception;

  procedure Assert(Check : in Boolean);
  procedure Assert(Check : in Boolean; Message : in String);

end Ada.Assertions;
```

A compilation unit containing a check for an assertion (including a `pragma Assert`) has a semantic dependence on the Assertions library unit.

#### *Dynamic Semantics*

If performing checks is required by the `Assert` assertion policy in effect at the place of a `pragma Assert`, the elaboration of the `pragma` consists of evaluating the boolean expression, and if the result is `False`, evaluating the `Message` argument, if any, and raising the exception `Assertions.Assertion_Error`, with a message if the `Message` argument is provided.

Calling the procedure `Assertions.Assert` without a `Message` parameter is equivalent to:

```
if Check = False then
  raise Ada.Assertions.Assertion_Error;
end if;
```

Calling the procedure `Assertions.Assert` with a `Message` parameter is equivalent to:

```
if Check = False then
  raise Ada.Assertions.Assertion_Error with Message;
end if;
```

The procedures `Assertions.Assert` have these effects independently of the assertion policy in effect.

#### *Bounded (Run-Time) Errors*

It is a bounded error to invoke a potentially blocking operation (see 12.5.1) during the evaluation of an assertion expression associated with a call on, or return from, a protected operation. If the bounded error is detected, `Program_Error` is raised. If not detected, execution proceeds normally, but if it is invoked within a protected action, it can result in deadlock or a (nested) protected action.

#### *Implementation Requirements*

Any postcondition expression, type invariant expression, or default initial condition expression occurring in the specification of a language-defined unit is enabled (see 9.1.1, 10.3.2, and 10.3.3).

The evaluation of any such postcondition, type invariant, or default initial condition expression shall either yield `True` or propagate an exception from a `raise_expression` that appears within the assertion expression.

Any precondition expression occurring in the specification of a language-defined unit is enabled (see 9.1.1) unless suppressed (see 14.5). Similarly, any predicate checks for a subtype occurring in the specification of a language-defined unit are enabled (see 6.2.4) unless suppressed.

#### Implementation Permissions

Assertion\_Error may be declared by renaming an implementation-defined exception from another package.

Implementations may define their own assertion policies.

If the result of a function call in an assertion is not used to determine the value of the assertion expression, an implementation is permitted to omit the function call. This permission applies even if the function has side effects.

An implementation may disallow the specification of an assertion expression if the evaluation of the expression has a side effect such that an immediate reevaluation of the expression can produce a different value. Similarly, an implementation may disallow the specification of an assertion expression that is checked as part of a call on or return from a callable entity *C*, if the evaluation of the expression has a side effect such that the evaluation of some other assertion expression associated with the same call of (or return from) *C* can produce a different value than in the case when the first expression had not been evaluated.

### 14.4.3 Example of Exception Handling

EXAMPLE Exception handling can be used to separate the detection of an error from the response to that error:

```
package File_System is
  type Data_Type is ...;
  type File_Handle is limited private;

  File_Not_Found : exception;
  procedure Open(F : in out File_Handle; Name : String);
    -- raises File_Not_Found if named file does not exist

  End_Of_File : exception;
  procedure Read(F : in out File_Handle; Data : out Data_Type);
    -- raises End_Of_File if the file is not open

  ...
private
  ...
end File_System;
package body File_System is
  ...

  procedure Open(F : in out File_Handle; Name : String) is
  begin
    if File_Exists(Name) then
      ...
    else
      raise File_Not_Found with "File not found: " & Name & ".";
    end if;
  end Open;

  procedure Read(F : in out File_Handle; Data : out Data_Type) is
  begin
    if F.Current_Position <= F.Last_Position then
      ...
    else
      raise End_Of_File;
    end if;
  end Read;

  ...
end File_System;
```

```

with Ada.Text_IO;
with Ada.Exceptions;
with File_System; use File_System;
use Ada;
procedure Main is
  Verbosity_Desired : Boolean := ...;
begin
  ... -- call operations in File_System
exception
  when End_Of_File =>
    Close(Some_File);
  when Not_Found_Error : File_Not_Found =>
    Text_IO.Put_Line(Exceptions.Exception_Message(Not_Found_Error));
  when The_Error : others =>
    Text_IO.Put_Line("Unknown error:");
    if Verbosity_Desired then
      Text_IO.Put_Line(Exceptions.Exception_Information(The_Error));
    else
      Text_IO.Put_Line(Exceptions.Exception_Name(The_Error));
      Text_IO.Put_Line(Exceptions.Exception_Message(The_Error));
    end if;
  raise;
end Main;

```

In the above example, the `File_System` package contains information about detecting certain exceptional situations, but it does not specify how to handle those situations. Procedure `Main` specifies how to handle them; other clients of `File_System` can have different handlers, even though the exceptional situations arise from the same basic causes.

## 14.5 Suppressing Checks

*Checking pragmas* give instructions to an implementation on handling language-defined checks. A `pragma Suppress` gives permission to an implementation to omit certain language-defined checks, while a `pragma Unsuppress` revokes the permission to omit checks.

A *language-defined check* (or simply, a “check”) is one of the situations defined by this document that requires a check to be made at run time to determine whether some condition is true. A check *fails* when the condition being checked is `False`, causing an exception to be raised.

### Syntax

The forms of checking pragmas are as follows:

```
pragma Suppress(identifier);
```

```
pragma Unsuppress(identifier);
```

A checking pragma is allowed only immediately within a `declarative_part`, immediately within a `package_specification`, or as a configuration pragma.

### Legality Rules

The `identifier` shall be the name of a check.

### Static Semantics

A checking pragma applies to the named check in a specific region, and applies to all entities in that region. A checking pragma given in a `declarative_part` or immediately within a `package_specification` applies from the place of the pragma to the end of the innermost enclosing declarative region. The region for a checking pragma given as a configuration pragma is the declarative region for the entire compilation unit (or units) to which it applies.

If a checking pragma applies to a `generic_instantiation`, then the checking pragma also applies to the entire instance.

A `pragma Suppress` gives permission to an implementation to omit the named check (or every check in the case of `All_Checks`) for any entities to which it applies. If permission has been given to suppress a given check, the check is said to be *suppressed*.

A `pragma Unsuppress` revokes the permission to omit the named check (or every check in the case of `All_Checks`) given by any `pragma Suppress` that applies at the point of the `pragma Unsuppress`. The permission is revoked for the region to which the `pragma Unsuppress` applies. If there is no such permission at the point of a `pragma Unsuppress`, then the `pragma` has no effect. A later `pragma Suppress` can renew the permission.

The following are the language-defined checks:

- The following checks correspond to situations in which the exception `Constraint_Error` is raised upon failure of a language-defined check.

#### Access\_Check

When evaluating a dereference (explicit or implicit), check that the value of the `name` is not **null**. When converting to a subtype that excludes null, check that the converted value is not **null**.

#### Discriminant\_Check

Check that the discriminants of a composite value have the values imposed by a discriminant constraint. Also, when accessing a record component, check that it exists for the current discriminant values.

#### Division\_Check

Check that the second operand is not zero for the operations `/`, **rem** and **mod**.

#### Index\_Check

Check that the bounds of an array value are equal to the corresponding bounds of an index constraint. Also, when accessing a component of an array object, check for each dimension that the given index value belongs to the range defined by the bounds of the array object. Also, when accessing a slice of an array object, check that the given discrete range is compatible with the range defined by the bounds of the array object.

#### Length\_Check

Check that two arrays have matching components, in the case of array subtype conversions, and logical operators for arrays of boolean components.

#### Overflow\_Check

Check that a scalar value is within the base range of its type, in cases where the implementation chooses to raise an exception instead of returning the correct mathematical result.

#### Range\_Check

Check that a scalar value satisfies a range constraint. Also, for the elaboration of a `subtype_indication`, check that the `constraint` (if present) is compatible with the subtype denoted by the `subtype_mark`. Also, for an aggregate, check that an index or discriminant value belongs to the corresponding subtype. Also, check that when the result of an operation yields an array, the value of each component belongs to the component subtype. Also, for the attributes `Value`, `Wide_Value`, and `Wide_Wide_Value`, check that the given string has the appropriate syntax and value for the base subtype of the `prefix` of the `attribute_reference`.

#### Tag\_Check

Check that operand tags in a dispatching call are all equal. Check for the correct tag on tagged type conversions, for an `assignment_statement`, and when returning a tagged limited object from a function.

- The following checks correspond to situations in which the exception `Program_Error` is raised upon failure of a language-defined check.



#### Accessibility\_Check

Check the accessibility level of an entity or view.

#### Allocation\_Check

For an `allocator`, check that the master of any tasks to be created by the `allocator` is not yet completed or some dependents have not yet terminated, and that the finalization of the collection has not started.

#### Elaboration\_Check

When a subprogram or protected entry is called, a task activation is accomplished, or a generic instantiation is elaborated, check that the body of the corresponding unit has already been elaborated.

#### Program\_Error\_Check

Other language-defined checks that raise `Program_Error`: that subtypes with predicates are not used to index an array in a generic unit; that the maximum number of chunks is greater than zero; that the default value of an out parameter is convertible; that there is no misuse of functions in a generic with a class-wide actual type; that there are not colliding `External_Tag` values; that there is no misuse of operations of unchecked union types.

- The following check corresponds to situations in which the exception `Storage_Error` is raised upon failure of a language-defined check.

#### Storage\_Check

Check that evaluation of an `allocator` does not require more space than is available for a storage pool. Check that the space available for a task or subprogram has not been exceeded.

- The following check corresponds to situations in which the exception `Tasking_Error` is raised upon failure of a language-defined check.

#### Tasking\_Check

Check that all tasks activated successfully. Check that a called task has not yet terminated.

- The following checks correspond to situations in which the exception `Assertion_Error` is raised upon failure of a language-defined check. For a language-defined unit *U* associated with one of these checks in the list below, the check refers to performance of checks associated with the `Pre`, `Static_Predicate`, and `Dynamic_Predicate` aspects associated with any entity declared in a descendant of *U*, or in an instance of a generic unit which is, or is declared in, a descendant of *U*. Each check is associated with one or more units:

#### Calendar\_Assertion\_Check

`Calendar`.

#### Characters\_Assertion\_Check

`Characters`, `Wide_Characters`, and `Wide_Wide_Characters`.

#### Containers\_Assertion\_Check

`Containers`.

#### Interfaces\_Assertion\_Check

`Interfaces`.

#### IO\_Assertion\_Check

`Sequential_IO`, `Direct_IO`, `Text_IO`, `Wide_Text_IO`, `Wide_Wide_Text_IO`, `Storage_IO`, `Streams.Stream_IO`, and `Directories`.

#### Numerics\_Assertion\_Check

`Numerics`.

#### Strings\_Assertion\_Check

`Strings`.

#### System\_Assertion\_Check

`System`.

- The following check corresponds to all situations in which any predefined exception is raised upon failure of a check.

#### All\_Checks

Represents the union of all checks; suppressing All\_Checks suppresses all checks other than those associated with assertions. In addition, an implementation is allowed (but not required) to behave as if a pragma Assertion\_Policy(Ignore) applies to any region to which pragma Suppress(All\_Checks) applies.

#### Erroneous Execution

If a given check has been suppressed, and the corresponding error situation occurs, the execution of the program is erroneous. Similarly, if a precondition check has been suppressed and the evaluation of the precondition would have raised an exception, execution is erroneous.

#### Implementation Permissions

An implementation is allowed to place restrictions on checking pragmas, subject only to the requirement that pragma Unsuppress shall allow any check names supported by pragma Suppress. An implementation is allowed to add additional check names, with implementation-defined semantics. When Overflow\_Check has been suppressed, an implementation may also suppress an unspecified subset of the Range\_Checks.

An implementation may support an additional parameter on pragma Unsuppress similar to the one allowed for pragma Suppress (see I.10). The meaning of such a parameter is implementation-defined.

#### Implementation Advice

The implementation should minimize the code executed for checks that have been suppressed.

NOTE 1 There is no guarantee that a suppressed check is actually removed; hence a pragma Suppress is useful only to improve efficiency.

NOTE 2 It is possible to give both a pragma Suppress and Unsuppress for the same check immediately within the same declarative\_part. In that case, the last pragma given determines whether or not the check is suppressed. Similarly, it is possible to resuppress a check which has been unsuppressed by giving a pragma Suppress in an inner declarative region.

EXAMPLE Examples of suppressing and unsuppressing checks:

```
pragma Suppress(Index_Check);
pragma Unsuppress(Overflow_Check);
```

## 14.6 Exceptions and Optimization

This subclause gives permission to the implementation to perform certain “optimizations” that do not necessarily preserve the canonical semantics.

#### Dynamic Semantics

The rest of this document (outside this subclause) defines the *canonical semantics* of the language. The canonical semantics of a given (legal) program determines a set of possible external effects that can result from the execution of the program with given inputs.

As explained in 4.2, the external effect of a program is defined in terms of its interactions with its external environment. Hence, the implementation can perform any internal actions whatsoever, in any order or in parallel, so long as the external effect of the execution of the program is one that is allowed by the canonical semantics, or by the rules of this subclause.

#### Implementation Permissions

The following additional permissions are granted to the implementation:

- An implementation can omit raising an exception when a language-defined check fails. Instead, the operation that failed the check can simply yield an *undefined result*. The exception is required to be

raised by the implementation only if, in the absence of raising it, the value of this undefined result would have some effect on the external interactions of the program. In determining this, the implementation shall not presume that an undefined result has a value that belongs to its subtype, nor even to the base range of its type, if scalar. Having removed the raise of the exception, the canonical semantics will in general allow the implementation to omit the code for the check, and some or all of the operation itself.

- If an exception is raised due to the failure of a language-defined check, then upon reaching the corresponding `exception_handler` (or the termination of the task, if none), the external interactions that have occurred have to reflect only that the exception was raised somewhere within the execution of the `sequence_of_statements` with the handler (or the `task_body`), possibly earlier (or later if the interactions are independent of the result of the checked operation) than that defined by the canonical semantics, but not within the execution of some abort-deferred operation or *independent* subprogram that does not dynamically enclose the execution of the construct whose check failed. An independent subprogram is one that is defined outside the library unit containing the construct whose check failed, and for which the `Inline` aspect is `False`. Any assignment that occurred outside of such abort-deferred operations or independent subprograms can be disrupted by the raising of the exception, causing the object or its parts to become abnormal, and certain subsequent uses of the object to be erroneous, as explained in 16.9.1.

NOTE The permissions granted by this subclause can have an effect on the semantics of a program only if the program fails a language-defined check.

## 15 Generic Units

A *generic unit* is a program unit that is either a generic subprogram or a generic package. A generic unit is a *template*, which can be parameterized, and from which corresponding (nongeneric) subprograms or packages can be obtained. The resulting program units are said to be *instances* of the original generic unit.

A generic unit is declared by a `generic_declaration`. This form of declaration has a `generic_formal_part` declaring any generic formal parameters. An instance of a generic unit is obtained as the result of a `generic_instantiation` with appropriate generic actual parameters for the generic formal parameters. An instance of a generic subprogram is a subprogram. An instance of a generic package is a package.

Generic units are templates. As templates they do not have the properties that are specific to their nongeneric counterparts. For example, a generic subprogram can be instantiated but it cannot be called. In contrast, an instance of a generic subprogram is a (nongeneric) subprogram; hence, this instance can be called but it cannot be used to produce further instances.

### 15.1 Generic Declarations

A `generic_declaration` declares a generic unit, which is either a generic subprogram or a generic package. A `generic_declaration` includes a `generic_formal_part` declaring any generic formal parameters. A generic formal parameter can be an object; alternatively (unlike a parameter of a subprogram), it can be a type, a subprogram, or a package.

*Syntax*

```
generic_declaration ::=
  generic_subprogram_declaration | generic_package_declaration
generic_subprogram_declaration ::=
  generic_formal_part subprogram_specification
  [aspect_specification];
generic_package_declaration ::=
  generic_formal_part package_specification;
generic_formal_part ::=
```

```

generic {generic_formal_parameter_declaration | use_clause}
generic_formal_parameter_declaration ::=
    formal_object_declaration
  | formal_type_declaration
  | formal_subprogram_declaration
  | formal_package_declaration

```

The only form of subtype\_indication allowed within a generic\_formal\_part is a subtype\_mark (that is, the subtype\_indication shall not include an explicit constraint). The defining name of a generic subprogram shall be an identifier (not an operator\_symbol).

#### Static Semantics

A generic\_declaration declares a generic unit — a generic package, generic procedure, or generic function, as appropriate.

An entity is a *generic formal* entity if it is declared by a generic\_formal\_parameter\_declaration. “Generic formal”, or simply “formal”, is used as a prefix in referring to objects, subtypes (and types), functions, procedures and packages, that are generic formal entities, as well as to their respective declarations. Examples: “generic formal procedure” or a “formal integer type declaration”.

The list of generic\_formal\_parameter\_declarations of a generic\_formal\_part form a *declaration list* of the generic unit.

#### Dynamic Semantics

The elaboration of a generic\_declaration has no effect.

NOTE 1 Outside a generic unit a name that denotes the generic\_declaration denotes the generic unit. In contrast, within the declarative region of the generic unit, a name that denotes the generic\_declaration denotes the current instance.

NOTE 2 Within a generic subprogram body, the name of this program unit acts as the name of a subprogram. Hence this name can be overloaded, and it can appear in a recursive call of the current instance. For the same reason, this name cannot appear after the reserved word **new** in a (recursive) generic\_instantiation.

NOTE 3 A default\_expression or default\_name appearing in a generic\_formal\_part is not evaluated during elaboration of the generic\_formal\_part; instead, it is evaluated when used. However, the usual visibility rules apply to any name used in a default, with name resolution performed based on the location of the name within the generic\_formal\_part.

EXAMPLE 1 Examples of generic formal parts:

```

generic      -- parameterless
generic
    Size : Natural;  -- formal object
generic
    Length : Integer := 200;          -- formal object with a default
expression
    Area   : Integer := Length*Length; -- formal object with a default
expression
generic
    type Item is private;          -- formal type
    type Index is (<>);             -- formal type
    type Row is array (Index range <>) of Item; -- formal type
    with function "<" (X, Y : Item) return Boolean; -- formal subprogram

```

EXAMPLE 2 *Examples of generic declarations declaring generic subprograms Exchange and Squaring:*

```
generic
  type Elem is private;
procedure Exchange(U, V : in out Elem);

generic
  type Item (<>) is private;
  with function "*" (U, V : Item) return Item is <>;
function Squaring(X : Item) return Item;
```

EXAMPLE 3 *Example of a generic declaration declaring a generic package:*

```
generic
  type Item is private;
  type Vector is array (Positive range <>) of Item;
  with function Sum(X, Y : Item) return Item;
package On_Vectors is
  function Sum (A, B : Vector) return Vector;
  function Sigma(A : Vector) return Item;
  Length_Error : exception;
end On_Vectors;
```

## 15.2 Generic Bodies

The body of a generic unit (a *generic body*) is a template for the instance bodies. The syntax of a generic body is identical to that of a nongeneric body.

### Dynamic Semantics

The elaboration of a generic body has no other effect than to establish that the generic unit can from then on be instantiated without failing the Elaboration\_Check. If the generic body is a child of a generic package, then its elaboration establishes that each corresponding declaration nested in an instance of the parent (see 13.1.1) can from then on be instantiated without failing the Elaboration\_Check.

NOTE The syntax of generic subprograms implies that a generic subprogram body is always the completion of a declaration.

EXAMPLE 1 *Example of a generic procedure body:*

```
procedure Exchange(U, V : in out Elem) is -- see 15.1
  T : Elem; -- the generic formal type
begin
  T := U;
  U := V;
  V := T;
end Exchange;
```

EXAMPLE 2 *Example of a generic function body:*

```
function Squaring(X : Item) return Item is -- see 15.1
begin
  return X*X; -- the formal operator "*"
end Squaring;
```

EXAMPLE 3 *Example of a generic package body:*

```
package body On_Vectors is -- see 15.1
  function Sum(A, B : Vector) return Vector is
    Result : Vector(A'Range); -- the formal type Vector
    Bias : constant Integer := B'First - A'First;
  begin
    if A'Length /= B'Length then
      raise Length_Error;
    end if;
```

```

    for N in A'Range loop
        Result(N) := Sum(A(N), B(N + Bias)); -- the formal function Sum
    end loop;
    return Result;
end Sum;

function Sigma(A : Vector) return Item is
    Total : Item := A(A'First); -- the formal type Item
begin
    for N in A'First + 1 .. A'Last loop
        Total := Sum(Total, A(N)); -- the formal function Sum
    end loop;
    return Total;
end Sigma;
end On_Vectors;

```

### 15.3 Generic Instantiation

An instance of a generic unit is declared by a `generic_instantiation`.

#### Syntax

```

generic_instantiation ::=
    package defining_program_unit_name is
        new generic_package_name [generic_actual_part]
            [aspect_specification];
    | [overriding_indicator]
        procedure defining_program_unit_name is
            new generic_procedure_name [generic_actual_part]
                [aspect_specification];
    | [overriding_indicator]
        function defining_designator is
            new generic_function_name [generic_actual_part]
                [aspect_specification];

generic_actual_part ::=
    (generic_association { generic_association })

generic_association ::=
    [generic_formal_parameter_selector_name =>] explicit_generic_actual_parameter

explicit_generic_actual_parameter ::= expression | variable_name
    | subprogram_name | entry_name | subtype_mark
    | package_instance_name

```

A `generic_association` is *named* or *positional* according to whether or not the *generic\_formal\_parameter\_selector\_name* is specified. Any positional associations shall precede any named associations.

The *generic actual parameter* is either the `explicit_generic_actual_parameter` given in a `generic_association` for each formal, or the corresponding `default_expression`, `default_subtype_mark`, or `default_name` if no `generic_association` is given for the formal. When the meaning is clear from context, the term “generic actual”, or simply “actual”, is used as a synonym for “generic actual parameter” and also for the view denoted by one, or the value of one.

#### Legality Rules

In a `generic_instantiation` for a particular kind of program unit (package, procedure, or function), the name shall denote a generic unit of the corresponding kind (generic package, generic procedure, or generic function, respectively).

The *generic\_formal\_parameter\_selector\_name* of a named `generic_association` shall denote a `generic_formal_parameter_declaration` of the generic unit being instantiated. If two or more



formal subprograms have the same defining name, then named associations are not allowed for the corresponding actuals.

The `generic_formal_parameter_declaration` for a positional `generic_association` is the parameter with the corresponding position in the `generic_formal_part` of the generic unit being instantiated.

A `generic_instantiation` shall contain at most one `generic_association` for each formal. Each formal without an association shall have a `default_expression`, `default_subtype_mark`, or `subprogram_default`.

In a generic unit, Legality Rules are enforced at compile time of the `generic_declaration` and generic body, given the properties of the formals. In the visible part and formal part of an instance, Legality Rules are enforced at compile time of the `generic_instantiation`, given the properties of the actuals. In other parts of an instance, Legality Rules are not enforced; this rule does not apply when a given rule explicitly specifies otherwise.

#### Static Semantics

A `generic_instantiation` declares an instance; it is equivalent to the instance declaration (a `package_declaration` or `subprogram_declaration`) immediately followed by the instance body, both at the place of the instantiation.

The instance is a copy of the text of the template. Each use of a formal parameter becomes (in the copy) a use of the actual, as explained below. An instance of a generic package is a package, that of a generic procedure is a procedure, and that of a generic function is a function.

The interpretation of each construct within a generic declaration or body is determined using the overloading rules when that generic declaration or body is compiled. In an instance, the interpretation of each (copied) construct is the same, except in the case of a name that denotes the `generic_declaration` or some declaration within the generic unit; the corresponding name in the instance then denotes the corresponding copy of the denoted declaration. The overloading rules do not apply in the instance.

In an instance, a `generic_formal_parameter_declaration` declares a view whose properties are identical to those of the actual, except when specified otherwise (in particular, see 9.1.1, 15.4 and 15.6). Similarly, for a declaration within a `generic_formal_parameter_declaration`, the corresponding declaration in an instance declares a view whose properties are identical to the corresponding declaration within the declaration of the actual.

Implicit declarations are also copied, and a name that denotes an implicit declaration in the generic denotes the corresponding copy in the instance. However, for a type declared within the visible part of the generic, a whole new set of primitive subprograms is implicitly declared for use outside the instance, and may differ from the copied set if the properties of the type in some way depend on the properties of some actual type specified in the instantiation. For example, if the type in the generic is derived from a formal private type, then in the instance the type will inherit subprograms from the corresponding actual type.

These new implicit declarations occur immediately after the type declaration in the instance, and override the copied ones. The copied ones can be called only from within the instance; the new ones can be called only from outside the instance, although for tagged types, the body of a new one can be executed by a call to an old one.

In the visible part of an instance, an explicit declaration overrides an implicit declaration if they are homographs, as described in 11.3. On the other hand, an explicit declaration in the private part of an instance overrides an implicit declaration in the instance, only if the corresponding explicit declaration in the generic overrides a corresponding implicit declaration in the generic. Corresponding rules apply to the other kinds of overriding described in 11.3.

## Post-Compilation Rules

Recursive generic instantiation is not allowed in the following sense: if a given generic unit includes an instantiation of a second generic unit, then the instance generated by this instantiation shall not include an instance of the first generic unit (whether this instance is generated directly, or indirectly by intermediate instantiations).

## Dynamic Semantics

For the elaboration of a `generic_instantiation`, each `generic_association` is first evaluated. If a default is used, an implicit `generic_association` is assumed for this rule. These evaluations are done in an arbitrary order, except that the evaluation for a default actual takes place after the evaluation for another actual if the default includes a `name` that denotes the other one. Finally, the instance declaration and body are elaborated.

For the evaluation of a `generic_association` the `generic actual` parameter is evaluated. Additional actions are performed in the case of a formal object of mode **in** (see 15.4).

NOTE If a formal type is not tagged, then the type is treated as an untagged type within the generic body. Deriving from such a type in a generic body is permitted; the new type does not get a `new` tag value, even if the actual is tagged. Overriding operations for such a derived type cannot be dispatched to from outside the instance.

EXAMPLE 1 *Examples of generic instantiations (see 15.1):*

```

procedure Swap is new Exchange(Elem => Integer);
procedure Swap is new Exchange(Character);      -- Swap is overloaded
function Square is new Squaring(Integer);      -- "*" of Integer used by
default
function Square1 is new Squaring(Item => Matrix, "*" => Matrix_Product);
function Square2 is new Squaring(Matrix, Matrix_Product); -- same as previous
package Int_Vectors is new On_Vectors(Integer, Table, "+");

```

EXAMPLE 2 *Examples of uses of instantiated units:*

```

Swap(A, B);
A := Square(A);

T : Table(1 .. 5) := (10, 20, 30, 40, 50);
N : Integer := Int_Vectors.Sigma(T); -- 150
                                     -- (see 15.2 for the body of Sigma)

use Int_Vectors;
M : Integer := Sigma(T); -- 150

```

## 15.4 Formal Objects

A generic formal object can be used to pass a value or variable to a generic unit.

## Syntax

```

formal_object_declaration ::=
  defining_identifier_list:mode [null_exclusion] subtype_mark
    [:= default_expression]
    [aspect_specification];
| defining_identifier_list:mode access_definition
    [:= default_expression]
    [aspect_specification];

```

## Name Resolution Rules

The expected type for the `default_expression`, if any, of a formal object is the type of the formal object.

For a generic formal object of mode **in**, the expected type for the actual is the type of the formal.

For a generic formal object of mode **in out**, the type of the actual shall resolve to the type determined by the `subtype_mark`, or for a `formal_object_declaration` with an `access_definition`, to

a specific anonymous access type. If the anonymous access type is an access-to-object type, the type of the actual shall have the same designated type as that of the `access_definition`. If the anonymous access type is an access-to-subprogram type, the type of the actual shall have a designated profile which is type conformant with that of the `access_definition`.

#### Legality Rules

If a generic formal object has a `default_expression`, then the mode shall be **in** (either explicitly or by default); otherwise, its mode shall be either **in** or **in out**.

For a generic formal object of mode **in**, the actual shall be an `expression`. For a generic formal object of mode **in out**, the actual shall be a name that denotes a variable for which renaming is allowed (see 11.5.1).

In the case where the type of the formal is defined by an `access_definition`, the type of the actual and the type of the formal:

- shall both be access-to-object types with statically matching designated subtypes and with both or neither being access-to-constant types; or
- shall both be access-to-subprogram types with subtype conformant designated profiles.

For a `formal_object_declaration` of mode **in out** with a `null_exclusion` or an `access_definition` that has a `null_exclusion`, the subtype of the actual matching the `formal_object_declaration` shall exclude null. In addition, if the actual matching the `formal_object_declaration` statically denotes the generic formal object of mode **in out** of another generic unit *G*, and the instantiation containing the actual occurs within the body of *G* or within the body of a generic unit declared within the declarative region of *G*, then the declaration of the formal object of *G* shall have a `null_exclusion`. In addition to the places where Legality Rules normally apply (see 15.3), this rule applies also in the private part of an instance of a generic unit.

#### Static Semantics

A `formal_object_declaration` declares a generic formal object. The default mode is **in**. For a formal object of mode **in**, the nominal subtype is the one denoted by the `subtype_mark` or `access_definition` in the declaration of the formal. For a formal object of mode **in out**, its type is determined by the `subtype_mark` or `access_definition` in the declaration; its nominal subtype is nonstatic, even if the `subtype_mark` denotes a static subtype; for a composite type, its nominal subtype is unconstrained if the first subtype of the type is unconstrained, even if the `subtype_mark` denotes a constrained subtype.

In an instance, a `formal_object_declaration` of mode **in** is a *full constant declaration* and declares a new stand-alone constant object whose initialization expression is the actual, whereas a `formal_object_declaration` of mode **in out** declares a view whose properties are identical to those of the actual.

#### Dynamic Semantics

For the evaluation of a `generic_association` for a formal object of mode **in**, a constant object is created, the value of the actual parameter is converted to the nominal subtype of the formal object, and assigned to the object, including any value adjustment — see 10.6.

## 15.5 Formal Types

A generic formal subtype can be used to pass to a generic unit a subtype whose type is in a certain category of types.

#### Syntax

```
formal_type_declaration ::=
    formal_complete_type_declaration
    | formal_incomplete_type_declaration
```

```

formal_complete_type_declaration ::=
    type defining_identifier[discriminant_part] is formal_type_definition
    [or use default_subtype_mark] [aspect_specification];
formal_incomplete_type_declaration ::=
    type defining_identifier[discriminant_part] [is tagged]
    [or use default_subtype_mark];
formal_type_definition ::=
    formal_private_type_definition
  | formal_derived_type_definition
  | formal_discrete_type_definition
  | formal_signed_integer_type_definition
  | formal_modular_type_definition
  | formal_floating_point_definition
  | formal_ordinary_fixed_point_definition
  | formal_decimal_fixed_point_definition
  | formal_array_type_definition
  | formal_access_type_definition
  | formal_interface_type_definition

```

#### Legality Rules

For a generic formal subtype, the actual shall be a `subtype_mark`; it denotes the *(generic) actual subtype*.

#### Static Semantics

A `formal_type_declaration` declares a *(generic) formal type*, and its first subtype, the *(generic) formal subtype*.

The form of a `formal_type_definition` *determines a category (of types)* to which the formal type belongs. For a `formal_private_type_definition` the reserved words **tagged** and **limited** indicate the category of types (see 15.5.1). The reserved word **tagged** also plays this role in the case of a `formal_incomplete_type_declaration`. For a `formal_derived_type_definition` the category of types is the derivation class rooted at the ancestor type. For other formal types, the name of the syntactic category indicates the category of types; a `formal_discrete_type_definition` defines a discrete type, and so on.

#### Legality Rules

The actual type shall be in the category determined for the formal.

The `default_subtype_mark`, if any, shall denote a subtype which is allowed as an actual subtype for the formal type.

#### Static Semantics

The formal type also belongs to each category that contains the determined category. The primitive subprograms of the type are as for any type in the determined category. For a formal type other than a formal derived type, these are the predefined operators of the type. For an elementary formal type, the predefined operators are implicitly declared immediately after the declaration of the formal type. For a composite formal type, the predefined operators are implicitly declared either immediately after the declaration of the formal type, or later immediately within the declarative region in which the type is declared according to the rules of 10.3.1. In an instance, the copy of such an implicit declaration declares a view of the predefined operator of the actual type, even if this operator has been overridden for the actual type and even if it is never declared for the actual type, unless the actual type is an untagged record type, in which case it declares a view of the primitive (equality) operator. The rules specific to formal derived types are given in 15.5.1.

NOTE 1 Generic formal types, like all types, are not named. Instead, a name can denote a generic formal subtype. Within a generic unit, a generic formal type is considered as being distinct from all other (formal or nonformal) types.

NOTE 2 A `discriminant_part` is allowed only for certain kinds of types, and therefore only for certain kinds of generic formal types. See 6.7.

EXAMPLE 1 *Examples of generic formal types:*

```
type Item is private;
type Buffer (Length : Natural) is limited private;

type Enum is (<>);
type Int is range <>;
type Angle is delta <>;
type Mass is digits <>;

type Table is array (Enum) of Item;
```

EXAMPLE 2 *Example of a generic formal part declaring a formal integer type:*

```
generic
  type Rank is range <>;
  First : Rank := Rank'First;
  Second : Rank := First + 1; -- the operator "+" of the type Rank
```

### 15.5.1 Formal Private and Derived Types

In its most general form, the category determined for a formal private type is all types, but the category can be restricted to only nonlimited types or to only tagged types. Similarly, the category for a formal incomplete type is all types but the category can be restricted to only tagged types; unlike other formal types, the actual type can be incompletely defined, and not ready to be frozen (see 16.14). The category determined for a formal derived type is the derivation class rooted at the ancestor type.

#### Syntax

```
formal_private_type_definition ::= [[abstract] tagged] [limited] private
formal_derived_type_definition ::=
  [abstract] [limited | synchronized] new
  subtype_mark [[and interface_list] with private]
```

#### Legality Rules

If a generic formal type declaration has a known `discriminant_part`, then it shall not include a `default_expression` for a discriminant.

The *ancestor subtype* of a formal derived type is the subtype denoted by the `subtype_mark` of the `formal_derived_type_definition`. For a formal derived type declaration, the reserved words **with private** shall appear if and only if the ancestor type is a tagged type; in this case the formal derived type is a private extension of the ancestor type and the ancestor shall not be a class-wide type. Similarly, an `interface_list` or the optional reserved words **abstract** or **synchronized** shall appear only if the ancestor type is a tagged type. The reserved word **limited** or **synchronized** shall appear only if the ancestor type and any progenitor types are limited types. The reserved word **synchronized** shall appear (rather than **limited**) if the ancestor type or any of the progenitor types are synchronized interfaces. The ancestor type shall be a limited interface if the reserved word **synchronized** appears.

The actual type for a formal derived type shall be a descendant of the ancestor type and every progenitor of the formal type. The actual type for a formal derived type shall be tagged if and only if the formal derived type is a private extension. If the reserved word **synchronized** appears in the declaration of the formal derived type, the actual type shall be a synchronized tagged type.

If a formal private or derived subtype is definite, then the actual subtype shall also be definite. If the formal type is nonlimited, the actual type shall be nonlimited.

A `formal_incomplete_type_declaration` declares a formal incomplete type. The only view of a formal incomplete type is an incomplete view. Thus, a formal incomplete type is subject to the same usage restrictions as any other incomplete type — see 6.10.1.

For a generic formal derived type with no `discriminant_part`, the actual subtype shall be statically compatible with the ancestor subtype. Furthermore:

- If the ancestor subtype is constrained, the actual subtype shall be constrained;
- If the ancestor subtype is an unconstrained access or composite subtype, the actual subtype shall be unconstrained.
- If the ancestor subtype is an unconstrained discriminated subtype, then the actual shall have the same number of discriminants, and each discriminant of the actual shall correspond to a discriminant of the ancestor, in the sense of 6.7.
- If the ancestor subtype is an access subtype, the actual subtype shall exclude null if and only if the ancestor subtype excludes null.

The declaration of a formal derived type shall not have a `known_discriminant_part`. For a generic formal private or incomplete type with a `known_discriminant_part`:

- The actual type shall be a type with the same number of discriminants.
- The actual subtype shall be unconstrained.
- The subtype of each discriminant of the actual type shall statically match the subtype of the corresponding discriminant of the formal type.

For a generic formal type with an `unknown_discriminant_part`, the actual may have discriminants, though that is not required, and may be definite or indefinite.

When enforcing Legality Rules, for the purposes of determining within a generic body whether a type is unconstrained in any partial view, a discriminated subtype is considered to have a constrained partial view if it is a descendant of an untagged generic formal private or derived type.

#### *Static Semantics*

The category determined for a formal private type is as follows:

##### *Type Definition*

##### *Determined Category*

**limited private**

the category of all types

**private**

the category of all nonlimited types

**tagged limited private**

the category of all tagged types

**tagged private**

the category of all nonlimited tagged types

The presence of the reserved word **abstract** determines whether the actual type may be abstract.

The category determined for a formal incomplete type is the category of all types, unless the `formal_type_declaration` includes the reserved word **tagged**; in this case, it is the category of all tagged types.

A formal private or derived type is a private or derived type, respectively. A formal derived tagged type is a private extension. A formal private or derived type is abstract if the reserved word **abstract** appears in its declaration.

For a formal derived type, the characteristics (including components, but excluding discriminants if there is a new `discriminant_part`), predefined operators, and inherited user-defined primitive subprograms are determined by its ancestor type and its progenitor types (if any), in the same way that those of a derived type are determined by those of its parent type and its progenitor types (see 6.4 and 10.3.1).

In an instance, the copy of an implicit declaration of a primitive subprogram of a formal derived type declares a view of the corresponding primitive subprogram of the ancestor or progenitor of the formal derived type, even if this primitive has been overridden for the actual type and even if it is never declared for the actual type. When the ancestor or progenitor of the formal derived type is itself a formal type, the copy of the implicit declaration declares a view of the corresponding copied operation of the ancestor or progenitor. In the case of a formal private extension, however, the tag of the formal



type is that of the actual type, so if the tag in a call is statically determined to be that of the formal type, the body executed will be that corresponding to the actual type.

In an instance, the implicitly composed and additive aspects (see 16.1.1) of a formal type are those of the actual; for a nonoverridable aspect, a formal derived type inherits the aspect if the ancestor or any progenitor has the aspect, according to the rules given in 16.1.

For a `prefix S` that denotes a formal indefinite subtype, the following attribute is defined:

**S'Definite** S'Definite yields True if the actual subtype corresponding to `S` is definite; otherwise, it yields False. The value of this attribute is of the predefined type Boolean.

#### Dynamic Semantics

In the case where a formal type has unknown discriminants, and the actual type is a class-wide type `T`Class:

- For the purposes of defining the primitive operations of the formal type, each of the primitive operations of the actual type is considered to be a subprogram (with an intrinsic calling convention — see 9.3.1) whose body consists of a dispatching call upon the corresponding operation of `T`, with its formal parameters as the actual parameters. If it is a function, the result of the dispatching call is returned.
- If the corresponding operation of `T` has no controlling formal parameters, then the controlling tag value is determined by the context of the call, according to the rules for tag-indeterminate calls (see 6.9.2 and 8.2). In the case where the tag would be statically determined to be that of the formal type, the call raises `Program_Error`. If such a function is renamed, any call on the renaming raises `Program_Error`.

NOTE 1 The actual type can be abstract only if the formal type is abstract (see 6.9.3).

NOTE 2 If the formal has a `discriminant_part`, the actual can be either definite or indefinite. Otherwise, the actual can only be definite.

### 15.5.2 Formal Scalar Types

A *formal scalar type* is one defined by any of the `formal_type_definitions` in this subclause. The category determined for a formal scalar type is the category of all discrete, signed integer, modular, floating point, ordinary fixed point, or decimal types.

#### Syntax

```
formal_discrete_type_definition ::= (<>)
formal_signed_integer_type_definition ::= range <>
formal_modular_type_definition ::= mod <>
formal_floating_point_definition ::= digits <>
formal_ordinary_fixed_point_definition ::= delta <>
formal_decimal_fixed_point_definition ::= delta <> digits <>
```

#### Legality Rules

The actual type for a formal scalar type shall not be a nonstandard numeric type.

### 15.5.3 Formal Array Types

The category determined for a formal array type is the category of all array types.

#### Syntax

```
formal_array_type_definition ::= array_type_definition
```

*Legality Rules*

The only form of `discrete_subtype_definition` that is allowed within the declaration of a generic formal (constrained) array subtype is a `subtype_mark`.

For a formal array subtype, the actual subtype shall satisfy the following conditions:

- The formal array type and the actual array type shall have the same dimensionality; the formal subtype and the actual subtype shall be either both constrained or both unconstrained.
- For each index position, the index types shall be the same, and the index subtypes (if unconstrained), or the index ranges (if constrained), shall statically match (see 7.9.1).
- The component subtypes of the formal and actual array types shall statically match.
- If the formal type has aliased components, then so shall the actual.

EXAMPLE *Example of formal array types:*

```
-- given the generic package
generic
  type Item    is private;
  type Index   is (<>);
  type Vector  is array (Index range <>) of Item;
  type Table   is array (Index) of Item;
package P is
  ...
end P;

-- and the types
type Mix      is array (Color range <>) of Boolean;
type Option   is array (Color) of Boolean;

-- then Mix can match Vector and Option can match Table
package R is new P(Item => Boolean, Index => Color,
                  Vector => Mix,      Table => Option);

-- Note that Mix cannot match Table and Option cannot match Vector
```

#### 15.5.4 Formal Access Types

The category determined for a formal access type is the category of all access types.

*Syntax*

```
formal_access_type_definition ::= access_type_definition
```

*Legality Rules*

For a formal access-to-object type, the designated subtypes of the formal and actual types shall statically match.

If and only if the `general_access_modifier` **constant** applies to the formal, the actual shall be an access-to-constant type. If the `general_access_modifier` **all** applies to the formal, then the actual shall be a general access-to-variable type (see 6.10). If and only if the formal subtype excludes null, the actual subtype shall exclude null.

For a formal access-to-subprogram subtype, the designated profiles of the formal and the actual shall be subtype conformant.

EXAMPLE *Example of formal access types:*

```
-- the formal types of the generic package
```

```

generic
  type Node is private;
  type Link is access Node;
package P is
  ...
end P;

-- can be matched by the actual types

type Car;
type Car_Name is access Car;

type Car is
  record
    Pred, Succ : Car_Name;
    Number      : License_Number;
    Owner       : Person;
  end record;

-- in the following generic instantiation

package R is new P(Node => Car, Link => Car_Name);

```

### 15.5.5 Formal Interface Types

The category determined for a formal interface type is the category of all interface types.

#### Syntax

```
formal_interface_type_definition ::= interface_type_definition
```

#### Legality Rules

The actual type shall be a descendant of every progenitor of the formal type.

The actual type shall be a limited, task, protected, or synchronized interface if and only if the formal type is also, respectively, a limited, task, protected, or synchronized interface.

**EXAMPLE** Example of the use of a generic with a formal interface type, to establish a standard interface that all tasks will implement so they can be managed appropriately by an application-specific scheduler:

```

type Root_Work_Item is tagged private;

generic
  type Managed_Task is task interface;
  type Work_Item(<>) is new Root_Work_Item with private;
package Server_Manager is
  task type Server is new Managed_Task with
    entry Start(Data : in out Work_Item);
  end Server;
end Server_Manager;

```

## 15.6 Formal Subprograms

Formal subprograms can be used to pass callable entities to a generic unit.

#### Syntax

```

formal_subprogram_declaration ::= formal_concrete_subprogram_declaration
| formal_abstract_subprogram_declaration

formal_concrete_subprogram_declaration ::=
  with subprogram_specification [is subprogram_default]
  [aspect_specification];

formal_abstract_subprogram_declaration ::=
  with subprogram_specification is abstract [subprogram_default]
  [aspect_specification];

subprogram_default ::= default_name | <> | null

```

`default_name ::= name`

A `subprogram_default` of **null** shall not be specified for a formal function or for a `formal_abstract_subprogram_declaration`.

#### *Name Resolution Rules*

The expected profile for the `default_name`, if any, is that of the formal subprogram.

For a generic formal subprogram, the expected profile for the actual is that of the formal subprogram.

#### *Legality Rules*

The profiles of the formal and any named default shall be mode conformant.

The profiles of the formal and actual shall be mode conformant.

For a parameter or result subtype of a `formal_subprogram_declaration` that has an explicit `null_exclusion`:

- if the actual matching the `formal_subprogram_declaration` statically denotes a generic formal subprogram of another generic unit *G*, and the instantiation containing the actual occurs within the body of a generic unit *G* or within the body of a generic unit declared within the declarative region of the generic unit *G*, then the corresponding parameter or result type of the formal subprogram of *G* shall have a `null_exclusion`;
- otherwise, the subtype of the corresponding parameter or result type of the actual matching the `formal_subprogram_declaration` shall exclude null. In addition to the places where Legality Rules normally apply (see 15.3), this rule applies also in the private part of an instance of a generic unit.

If the named default, if any, is a prefixed view, the prefix of that view shall denote an object for which renaming is allowed (see 11.5.1). Similarly, if the actual subprogram in an instantiation is a prefixed view, the prefix of that view shall denote an object for which renaming is allowed.

If a formal parameter of a `formal_abstract_subprogram_declaration` is of a specific tagged type *T* or of an anonymous access type designating a specific tagged type *T*, *T* is called a *controlling type* of the `formal_abstract_subprogram_declaration`. Similarly, if the result of a `formal_abstract_subprogram_declaration` for a function is of a specific tagged type *T* or of an anonymous access type designating a specific tagged type *T*, *T* is called a controlling type of the `formal_abstract_subprogram_declaration`. A `formal_abstract_subprogram_declaration` shall have exactly one controlling type, and that type shall not be incomplete.

The actual subprogram for a `formal_abstract_subprogram_declaration` shall be:

- a dispatching operation of the controlling type; or
- if the controlling type is a formal type, and the actual type corresponding to that formal type is a specific type *T*, a dispatching operation of type *T*; or
- if the controlling type is a formal type, and the actual type is a class-wide type *T'*Class, an implicitly declared subprogram corresponding to a primitive operation of type *T* (see Static Semantics below).

#### *Static Semantics*

A `formal_subprogram_declaration` declares a generic formal subprogram. The types of the formal parameters and result, if any, of the formal subprogram are those determined by the `subtype_marks` given in the `formal_subprogram_declaration`; however, independent of the particular subtypes that are denoted by the `subtype_marks`, the nominal subtypes of the formal parameters and result, if any, are defined to be nonstatic, and unconstrained if of an array type (no applicable index constraint is provided in a call on a formal subprogram). In an instance, a `formal_subprogram_declaration` declares a view of the actual. The profile of this view takes its subtypes and calling convention from the original profile of the actual entity, while taking the formal

parameter names and `default_expressions` from the profile given in the `formal_subprogram_declaration`. The view is a function or procedure, never an entry.

If a `subtype_mark` in the profile of the `formal_subprogram_declaration` denotes a formal private or formal derived type and the actual type for this formal type is a class-wide type `T'Class`, then for the purposes of resolving the corresponding actual subprogram at the point of the instantiation, certain implicit declarations may be available as possible resolutions as follows:

For each primitive subprogram of `T` that is directly visible at the point of the instantiation, and that has at least one controlling formal parameter, a corresponding implicitly declared subprogram with the same defining name, and having the same profile as the primitive subprogram except that `T` is systematically replaced by `T'Class` in the types of its profile, is potentially use-visible. The body of such a subprogram is as defined in 15.5.1 for primitive subprograms of a formal type when the actual type is class-wide.

If a generic unit has a `subprogram_default` specified by a box, and the corresponding actual parameter is omitted, then it is equivalent to an explicit actual parameter that is a usage name identical to the defining name of the formal.

If a generic unit has a `subprogram_default` specified by the reserved word **null**, and the corresponding actual parameter is omitted, then it is equivalent to an explicit actual parameter that is a null procedure having the profile given in the `formal_subprogram_declaration`.

The subprogram declared by a `formal_abstract_subprogram_declaration` with a controlling type `T` is a dispatching operation of type `T`.

NOTE 1 The matching rules for formal subprograms state requirements that are similar to those applying to `subprogram_renaming_declarations` (see 11.5.4). In particular, the name of a parameter of the formal subprogram can be different from that of the corresponding parameter of the actual subprogram; similarly, for these parameters, `default_expressions` can be different.

NOTE 2 The subtype specified for a formal parameter of a generic formal subprogram can be any visible subtype, including a generic formal subtype of the same generic `formal_part`.

NOTE 3 A formal subprogram is matched by an attribute of a type if the attribute is a function with a matching specification. An enumeration literal of a given type matches a parameterless formal function whose result type is the given type.

NOTE 4 A `default_name` denotes an entity that is visible or directly visible at the place of the `generic_declaration`; a box used as a default is equivalent to a name that denotes an entity that is directly visible at the place of the `generic_instantiation`.

NOTE 5 The actual subprogram cannot be abstract unless the formal subprogram is a `formal_abstract_subprogram_declaration` (see 6.9.3).

NOTE 6 The subprogram declared by a `formal_abstract_subprogram_declaration` is an abstract subprogram. All calls on a subprogram declared by a `formal_abstract_subprogram_declaration` are limited to dispatching calls. See 6.9.3.

NOTE 7 A null procedure as a subprogram default has convention *Intrinsic* (see 9.3.1).

EXAMPLE *Examples of generic formal subprograms:*

```
with function "+"(X, Y : Item) return Item is <>;
with function Image(X : Enum) return String is Enum'Image;
with procedure Update is Default_Update;
with procedure Pre_Action(X : in Item) is null; -- defaults to no action
with procedure Write(S : not null access Root_Stream_Type'Class;
                    Desc : Descriptor)
is abstract Descriptor'Write; -- see 16.13.2
-- Dispatching operation on Descriptor with default
-- given the generic procedure declaration

generic
  with procedure Action (X : in Item);
procedure Iterate(Seq : in Item_Sequence);
```

```
-- and the procedure
procedure Put_Item(X : in Item);
-- the following instantiation is possible
procedure Put_List is new Iterate(Action => Put_Item);
```

## 15.7 Formal Packages

Formal packages can be used to pass packages to a generic unit. The `formal_package_declaration` declares that the formal package is an instance of a given generic package. Upon instantiation, the actual package has to be an instance of that generic package.

### Syntax

```
formal_package_declaration ::=
  with package defining_identifier is new
    generic_package_name formal_package_actual_part
    [aspect_specification];
formal_package_actual_part ::=
  ([others =>] <>)
  | [generic_actual_part]
  | (formal_package_association { formal_package_association } [, others => <>])
formal_package_association ::=
  generic_association
  | generic_formal_parameter_selector_name => <>
```

Any positional `formal_package_associations` shall precede any named `formal_package_associations`.

### Legality Rules

The `generic_package_name` shall denote a generic package (the *template* for the formal package); the formal package is an instance of the template.

The `generic_formal_parameter_selector_name` of a `formal_package_association` shall denote a `generic_formal_parameter_declaration` of the template. If two or more formal subprograms of the template have the same defining name, then named associations are not allowed for the corresponding actuals.

A `formal_package_actual_part` shall contain at most one `formal_package_association` for each formal parameter. If the `formal_package_actual_part` does not include “**others** => <>”, each formal parameter without an association shall have a `default_expression` or `subprogram_default`.

The rules for matching between `formal_package_associations` and the generic formal parameters of the template are as follows:

- If all of the `formal_package_associations` are given by generic associations, the `explicit_generic_actual_parameters` of the `formal_package_associations` shall be legal for an instantiation of the template.
- If a `formal_package_association` for a formal type *T* of the template is given by <>, then the `formal_package_association` for any other `generic_formal_parameter_declaration` of the template that mentions *T* directly or indirectly shall also be given by <>.

The actual shall be an instance of the template. If the `formal_package_actual_part` is (<>) or (**others** => <>), then the actual may be any instance of the template; otherwise, certain of the actual parameters of the actual instance shall match the corresponding actual parameters of the formal package, determined as follows:



- If the `formal_package_actual_part` includes `generic_associations` as well as associations with `<>`, then only the actual parameters specified explicitly with `generic_associations` are required to match;
- Otherwise, all actual parameters shall match, whether any actual parameter is given explicitly or by default.

The rules for matching of actual parameters between the actual instance and the formal package are as follows:

- For a formal object of mode **in**, the actuals match if they are static expressions with the same value, or if they statically denote the same constant, or if they are both the literal **null**.
- For a formal subtype, the actuals match if they denote statically matching subtypes.
- For other kinds of formals, the actuals match if they statically denote the same entity.

For the purposes of matching, any actual parameter that is the name of a formal object of mode **in** is replaced by the formal object's actual expression (recursively).

#### Static Semantics

A `formal_package_declaration` declares a generic formal package.

The visible part of a formal package includes the first list of `basic_declarative_items` of the `package_specification`. In addition, for each actual parameter that is not required to match, a copy of the declaration of the corresponding formal parameter of the template is included in the visible part of the formal package. If the copied declaration is for a formal type, copies of the implicit declarations of the primitive subprograms of the formal type are also included in the visible part of the formal package.

For the purposes of matching, if the actual instance *A* is itself a formal package, then the actual parameters of *A* are those specified explicitly or implicitly in the `formal_package_actual_part` for *A*, plus, for those not specified, the copies of the formal parameters of the template included in the visible part of *A*.

EXAMPLE 1 *Example of a generic package with formal package parameters:*

```
with Ada.Containers.Ordered_Maps; -- see A.18.6
generic
  with package Mapping_1 is new Ada.Containers.Ordered_Maps(<>);
  with package Mapping_2 is new Ada.Containers.Ordered_Maps
    (Key_Type => Mapping_1.Element_Type,
     others => <>);
package Ordered_Join is
  -- Provide a "join" between two mappings
  subtype Key_Type is Mapping_1.Key_Type;
  subtype Element_Type is Mapping_2.Element_Type;
  function Lookup(Key : Key_Type) return Element_Type;
  ...
end Ordered_Join;
```

EXAMPLE 2 *Example of an instantiation of a package with formal packages:*

```
with Ada.Containers.Ordered_Maps;
package Symbol_Package is
  subtype Key_String is String(1..5);
  type String_Id is ...
  type Symbol_Info is ...
  package String_Table is new Ada.Containers.Ordered_Maps
    (Key_Type => Key_String,
     Element_Type => String_Id);
```

```

package Symbol_Table is new Ada.Containers.Ordered_Maps
    (Key_Type => String_Id,
     Element_Type => Symbol_Info);

package String_Info is new Ordered_Join(Mapping_1 => String_Table,
                                         Mapping_2 => Symbol_Table);

Apple_Info : constant Symbol_Info := String_Info.Lookup("Apple");
end Symbol_Package;

```

## 15.8 Example of a Generic Package

EXAMPLE The following example provides a possible formulation of stacks by means of a generic package. The size of each stack and the type of the stack elements are provided as generic formal parameters.

```

generic
    Size : Positive;
    type Item is private;
package Stack is
    procedure Push(E : in Item);
    procedure Pop (E : out Item);
    Overflow, Underflow : exception;
end Stack;

package body Stack is
    type Table is array (Positive range <>) of Item;
    Space : Table(1 .. Size);
    Index : Natural := 0;

    procedure Push(E : in Item) is
    begin
        if Index >= Size then
            raise Overflow;
        end if;
        Index := Index + 1;
        Space(Index) := E;
    end Push;

    procedure Pop(E : out Item) is
    begin
        if Index = 0 then
            raise Underflow;
        end if;
        E := Space(Index);
        Index := Index - 1;
    end Pop;
end Stack;

```

Instances of this generic package can be obtained as follows:

```

package Stack_Int is new Stack(Size => 200, Item => Integer);
package Stack_Bool is new Stack(100, Boolean);

```

Thereafter, the procedures of the instantiated packages can be called as follows:

```

Stack_Int.Push(N);
Stack_Bool.Push(True);

```

Alternatively, a generic formulation of the type Stack can be given as follows (package body omitted):

```
generic
  type Item is private;
package On_Stacks is
  type Stack(Size : Positive) is limited private;
  procedure Push(S : in out Stack; E : in Item);
  procedure Pop (S : in out Stack; E : out Item);
  Overflow, Underflow : exception;
private
  type Table is array (Positive range <>) of Item;
  type Stack(Size : Positive) is
    record
      Space : Table(1 .. Size);
      Index : Natural := 0;
    end record;
end On_Stacks;
```

In order to use such a package, an instance has to be created and thereafter stacks of the corresponding type can be declared:

```
declare
  package Stack_Real is new On_Stacks(Real); use Stack_Real;
  S : Stack(100);
begin
  ...
  Push(S, 2.54);
  ...
end;
```

## 16 Representation Issues

This clause describes features for querying and controlling certain aspects of entities and for interfacing to hardware.

### 16.1 Operational and Representation Aspects

Two kinds of aspects of entities can be specified: representation aspects and operational aspects. Representation aspects affect how the types and other entities of the language are to be mapped onto the underlying machine. Operational aspects determine other properties of entities.

Either kind of aspect of an entity may be specified by means of an *aspect\_specification* (see 16.1.1), which is an optional element of most kinds of declarations and applies to the entity or entities being declared. Aspects may also be specified by certain other constructs occurring subsequent to the declaration of the affected entity: a representation aspect value may be specified by means of a representation item and an operational aspect value may be specified by means of an operational item.

There are six kinds of *representation items*: *attribute\_definition\_clauses* for representation attributes, *enumeration\_representation\_clauses*, *record\_representation\_clauses*, *at\_clauses*, *component\_clauses*, and *representation pragmas*. They can be provided to give more efficient representation or to interface with features that are outside the domain of the language (for example, peripheral hardware).

An *operational item* is an *attribute\_definition\_clause* for an operational attribute.

An operational item or a representation item applies to an entity identified by a *local\_name*, which denotes an entity declared local to the current declarative region, or a library unit declared immediately preceding a representation pragma in a compilation.

*Syntax*

```
aspect_clause ::= attribute_definition_clause
| enumeration_representation_clause
```

```

| record_representation_clause
| at_clause
local_name ::= direct_name
| direct_name'attribute_designator
| library_unit_name

```

A representation pragma is allowed only at places where an `aspect_clause` or `compilation_unit` is allowed.

#### *Name Resolution Rules*

In an operational item or representation item, if the `local_name` is a `direct_name`, then it shall resolve to denote a declaration (or, in the case of a pragma, one or more declarations) that occurs immediately within the same declarative region as the item. If the `local_name` has an `attribute_designator`, then it shall resolve to denote an implementation-defined component (see 16.5.1) or a class-wide type implicitly declared immediately within the same declarative region as the item. A `local_name` that is a `library_unit_name` (only permitted in a representation pragma) shall resolve to denote the `library_item` that immediately precedes (except for other pragmas) the representation pragma.

#### *Legality Rules*

The `local_name` of an `aspect_clause` or representation pragma shall statically denote an entity (or, in the case of a pragma, one or more entities) declared immediately preceding it in a compilation, or within the same declarative part, package\_specification, task\_definition, protected\_definition, or record\_definition as the representation or operational item. If a `local_name` denotes a local callable entity, it may do so through a local subprogram\_renaming\_declaration (as a way to resolve ambiguity in the presence of overloading); otherwise, the `local_name` shall not denote a renaming\_declaration.

The *representation* of an object consists of a certain number of bits (the *size* of the object). For an object of an elementary type, these are the bits that are normally read or updated by the machine code when loading, storing, or operating-on the value of the object. For an object of a composite type, these are the bits reserved for this object, and include bits occupied by subcomponents of the object. If the size of an object is greater than that of its subtype, the additional bits are padding bits. For an elementary object, these padding bits are normally read and updated along with the others. For a composite object, it is unspecified whether padding bits are read or updated in any given composite operation.

A representation item *directly specifies* a *representation aspect* of the entity denoted by the `local_name`, except in the case of a type-related representation item, whose `local_name` shall denote a first subtype, and which directly specifies an aspect of the subtype's type. A representation item that names a subtype is either *subtype-specific* (Size, Object\_Size, and Alignment clauses) or *type-related* (all others).

An operational item *directly specifies* an *operational aspect* of the entity denoted by the `local_name`, except in the case of a type-related operational item, whose `local_name` shall denote a first subtype, and which directly specifies an aspect of the type of the subtype.

Aspects that can be specified for types and subtypes are also classified into type-related or subtype-specific aspects. Representation aspects that can be specified for types and subtypes are considered type-related unless specified otherwise. In contrast, the classification of operational aspects are given with the definition of the aspect. Type-related aspects have the same value for all subtypes of (a view of) a type, while subtype-specific aspects may differ for different subtypes of the same type.

A representation item or operational item that directly specifies an aspect of an entity shall appear before the entity is frozen (see 16.14).

A representation aspect of a subtype or type shall not be specified (whether by a representation item or an `aspect_specification`) before the type is completely defined (see 6.11.1).

If a `representation item`, `operational item`, `library unit pragma` (see I.15), or `aspect_specification` is given that directly specifies an aspect of an entity, then it is illegal to give another `representation item`, `operational item`, `library unit pragma`, or `aspect_specification` that directly specifies the same aspect of the entity.

Unless otherwise specified, it is illegal to specify an operational or representation aspect of a generic formal parameter.

A *by-reference primitive* is a user-defined primitive subprogram for a type *T* that has an access result designating type *T*, or that has a formal parameter that is an access parameter designating type *T* or is aliased and of type *T*. It is illegal to specify a nonconfirming type-related representation aspect for an untagged type *T* if it is derived from a by-reference type or inherits one or more by-reference primitives, or if one or more types have been derived from *T* prior to the specification of the aspect and type *T* is a by-reference type or defines one or more by-reference primitives that are inherited by these descendants.

If a type-related aspect is defined for the partial view of a type, then it has the same definition for the full view of the type, except for certain Boolean-valued operational aspects where the language specifies that the partial view can have the value `False` even when the full view has the value `True`. Type-related aspects cannot be specified, and are not defined for an incomplete view of a type. Representation aspects of a generic formal parameter are the same as those of the actual. Specification of a type-related representation aspect is not allowed for a descendant of a generic formal untagged type.

The specification of the `Size` aspect for a given subtype, or the size or storage place for an object (including a component) of a given subtype, shall allow for enough storage space to accommodate any value of the subtype.

The specification of certain language-defined aspects is not required to be supported by all implementations; in such an implementation, the specification for such an aspect is illegal or raises an exception at run time.

A `type_declaration` is illegal if it has one or more progenitors, and a nonconfirming value was specified for a representation aspect of an ancestor, and this conflicts with the representation of some other ancestor. The cases that cause conflicts are implementation defined.

When specifying an aspect that denotes a subprogram, the profile of the subprogram shall be mode conformant with the one required for the aspect, and the convention shall be `Ada`. Additional requirements are defined for particular aspects.

#### Static Semantics

If two subtypes statically match, then their subtype-specific aspects (for example, `Size` and `Alignment`) are the same.

A derived type inherits each type-related representation aspect of its parent type that was directly specified before the declaration of the derived type, or (in the case where the parent is derived) that was inherited by the parent type from the grandparent type. A derived subtype inherits each subtype-specific representation aspect of its parent subtype that was directly specified before the declaration of the derived type, or (in the case where the parent is derived) that was inherited by the parent subtype from the grandparent subtype, but only if the parent subtype statically matches the first subtype of the parent type. An inherited representation aspect is overridden by a subsequent `aspect_specification` or `representation item` that specifies a different value for the same aspect of the type or subtype.

In contrast, whether type-related operational aspects are inherited by a derived type depends on each specific aspect; unless specified, an operational aspect is not inherited. When type-related operational aspects are inherited by a derived type, aspects that were directly specified by `aspect_specifications` or `operational items` that are visible at any point within the immediate scope of the derived type declaration, or (in the case where the parent is derived) that were inherited by the parent type from the grandparent type, are inherited. An inherited operational aspect is

overridden by an `aspect_specification` or operational item that specifies the same aspect of the type.

When a type-related operational aspect is inherited, the rules for inheritance depend on the nature of the aspect (see 16.1.1). Unless otherwise specified for a given aspect, these rules are as follows:

- For an operational aspect that is a value, the inherited aspect has the same value;
- For an operational aspect that is a name:
  - if the name denotes one or more primitive subprograms of the type, the inherited aspect is a name that denotes the corresponding primitive subprogram(s) of the derived type;
  - otherwise, the inherited aspect is a name that denotes the same entity or entities as the original aspect;
- For an operational aspect that is an identifier specific to the aspect, the inherited aspect is the same identifier;
- For an operational aspect that is an expression or an aggregate, the inherited aspect is a corresponding expression or aggregate where each name, value, and identifier follows these same rules for inheritance.

Each aspect of representation of an entity is as follows:

- If the aspect is *specified* for the entity, meaning that it is either directly specified or inherited, then that aspect of the entity is as specified, except in the case of `Storage_Size`, which specifies a minimum.
- If an aspect of representation of an entity is not specified, it is chosen by default in an unspecified manner.

If an operational aspect is *specified* for an entity (meaning that it is either directly specified or, if type-related or subtype-specific, inherited), then that aspect of the entity is as specified. Otherwise, the aspect of the entity has the default value for that aspect. For aspects that are neither type-related nor subtype-specific, the terms “specified” and “directly specified” are equivalent.

An `aspect_specification` or representation item that specifies a representation aspect that would have been chosen in the absence of the `aspect_specification` or representation item is said to be *confirming*. The aspect value specified in this case is said to be a *confirming* representation aspect value. Other values of the aspect are said to be *nonconfirming*, as are the `aspect_specifications` and representation items that specified them. Similarly, an `aspect_specification` or operational item that specifies an operational aspect to be the same as the definition it would have by default is said to be *confirming*; otherwise it is *nonconfirming*.

#### Dynamic Semantics

For the elaboration of an `aspect_clause`, any evaluable constructs within it are evaluated.

#### Implementation Permissions

An implementation may interpret representation aspects in an implementation-defined manner. An implementation may place implementation-defined restrictions on the specification of representation aspects. A *recommended level of support* is defined for the specification of representation aspects and related features in each subclause. These recommendations are changed to requirements for implementations that support the Systems Programming Annex (see C.2).

#### Implementation Advice

The recommended level of support for the specification of all representation aspects is qualified as follows:

- A confirming specification for a representation aspect should be supported.



- An implementation is not required to support the specification for a representation aspect that contains nonstatic expressions, unless each nonstatic expression is a *name* that statically denotes a constant declared before the entity.
- An implementation is not required to support a specification for the *Object\_Size* or *Size* for a given composite subtype, nor the size or storage place for an object (including a component) of a given composite subtype, unless the constraints on the subtype and its composite subcomponents (if any) are all static constraints.
- An implementation is not required to support specifying a nonconfirming representation aspect value if it can cause an aliased object or an object of a by-reference type to be allocated at a nonaddressable location or, when the alignment attribute of the subtype of such an object is nonzero, at an address that is not an integral multiple of that alignment.
- An implementation is not required to support specifying a nonconfirming representation aspect value if it can cause an aliased object of an elementary type to have a size other than that which would have been chosen by default.
- An implementation is not required to support specifying a nonconfirming representation aspect value if it can cause an aliased object of a composite type, or an object whose type is by-reference, to have a size smaller than that which would have been chosen by default.
- An implementation is not required to support specifying a nonconfirming subtype-specific representation aspect value for an indefinite or abstract subtype.

For purposes of these rules, the determination of whether specifying a representation aspect value for a type *can cause* an object to have some property is based solely on the properties of the type itself, not on any available information about how the type is used. In particular, it presumes that minimally aligned objects of this type can be declared at some point.

NOTE Aspects that can be specified are defined throughout this document, and are summarized in J.1.

### 16.1.1 Aspect Specifications

Certain representation or operational aspects of an entity may be specified as part of its declaration using an *aspect\_specification*, rather than using a separate representation or operational item. The declaration with the *aspect\_specification* is termed the *associated declaration*.

#### Syntax

```

aspect_specification ::=
  with aspect_mark [=> aspect_definition] {
    aspect_mark [=> aspect_definition] }

aspect_mark ::= aspect_identifier ['Class']

aspect_definition ::=
  name | expression | identifier
  | aggregate | global_aspect_definition

```

#### Name Resolution Rules

An *aspect\_mark* identifies an aspect of the entity defined by the associated declaration (the *associated entity*); the aspect denotes an object, a value, an expression, an aggregate, a subprogram, or some other kind of entity. If the *aspect\_mark* identifies:

- an aspect that denotes an object, the *aspect\_definition* shall be a *name*. The expected type for the *name* is the type of the identified aspect of the associated entity;
- an aspect that is a value or an expression, the *aspect\_definition* shall be an *expression*. The expected type for the *expression* is the type of the identified aspect of the associated entity;
- an aspect that is an aggregate, the aspect definition shall be an *expression* that is an aggregate, with the form of the aggregate determined by the identified aspect;

- an aspect that denotes a subprogram, the `aspect_definition` shall be a name; the expected profile for the name is the profile required for the aspect of the associated entity;
- an aspect that denotes some other kind of entity, the `aspect_definition` shall be a name, and the name shall resolve to denote an entity of the appropriate kind;
- an aspect that is given by an identifier specific to the aspect, the `aspect_definition` shall be an identifier, and the identifier shall be one of the identifiers specific to the identified aspect.

The usage names in an `aspect_definition` associated with a declaration are not resolved at the point of the associated declaration, but rather are resolved at the end of the immediately enclosing declaration list, or in the case of the declaration of a library unit, at the end of the visible part of the entity.

If the associated declaration is for a subprogram, entry, or access-to-subprogram type, the names of the formal parameters are directly visible within the `aspect_definition`, as are certain attributes, as specified elsewhere in this document for the identified aspect. If the associated declaration is a `type_declaration`, within the `aspect_definition` the names of any visible components, protected subprograms, and entries are directly visible, and the name of the first subtype denotes the current instance of the type (see 11.6). If the associated declaration is a `subtype_declaration`, within the `aspect_definition` the name of the new subtype denotes the current instance of the subtype.

#### Legality Rules

If the first freezing point of the associated entity comes before the end of the immediately enclosing declaration list, then each usage name in the `aspect_definition` shall resolve to the same entity at the first freezing point as it does at the end of the immediately enclosing declaration list.

An expression or name that causes freezing of an entity shall not occur within an `aspect_specification` that specifies a representation or operational aspect of that entity.

At most one occurrence of each `aspect_mark` is allowed within a single `aspect_specification`. The aspect identified by the `aspect_mark` shall be an aspect that can be specified for the associated entity (or view of the entity defined by the associated declaration).

The `aspect_definition` associated with a given `aspect_mark` may be omitted only when the `aspect_mark` identifies an aspect of a boolean type, in which case it is equivalent to the `aspect_definition` being specified as `True`.

If the `aspect_mark` includes 'Class, then the associated entity shall be a tagged type or a primitive subprogram of a tagged type.

Unless otherwise specified for a specific aspect, a language-defined aspect cannot be specified on a `renaming_declaration` or a `generic_formal_parameter_declaration`.

Unless specified otherwise, a language-defined aspect shall not be specified in an `aspect_specification` given on a completion of a program unit.

If an aspect of a derived type is inherited from an ancestor type and has the boolean value `True`, the inherited value shall not be overridden to have the value `False` for the derived type, unless otherwise specified in this document.

If a given aspect is type-related and inherited, then within an `aspect_definition` for the aspect, if a name resolves to denote multiple visible subprograms, all or none of the denoted subprograms shall be primitives of the associated type.

Certain type-related aspects are defined to be *nonoverridable*; all such aspects are inherited by derived types according to the rules given in 16.1. Any legality rule associated with a nonoverridable aspect is re-checked for the derived type, if the derived type is not abstract. Certain type-related and subtype-specific aspects are defined to be *additive*; such aspects are not inherited, but they can *apply* to the types derived from, or the subtypes based on, the original type or subtype, as defined for each such

aspect. Finally, certain type-related aspects are *implicitly composed*; such aspects are not inherited, but rather a default implementation for a derived type is provided, as defined for each such aspect, based on that of its parent type, presuming the aspect for the parent type is available where the derived type is declared, plus those of any new components added as part of a type extension.

If a nonoverridable aspect is directly specified for a type *T*, then any explicit specification of that aspect for any descendant of *T* (other than *T* itself) shall be *confirming*. In the case of an aspect that is a *name*, this means that the specified *name* shall *match* the inherited aspect in the sense that it shall denote the same declarations as would the inherited *name*. Similarly, for an aspect that is an *expression* or an *aggregate*, confirming means the defining *expression* is fully conformant (see 9.3.1) with the defining *expression* for the inherited aspect, with the added rule that an identifier that is specific to the aspect is the same as the corresponding identifier in the inherited aspect.

If a full type has a partial view, and a given nonoverridable aspect is allowed for both the full view and the partial view, then the given aspect for the partial view and the full view shall be the same: the aspect shall be directly specified only on the partial view; if the full type inherits the aspect, then a matching definition shall be specified (directly or by inheritance) for the partial view.

If a type inherits a nonoverridable aspect from multiple ancestors, the value of the aspect inherited from any given ancestor shall be confirming of the values inherited from all other ancestors.

In addition to the places where Legality Rules normally apply (see 15.3), these rules about nonoverridable aspects also apply in the private part of an instance of a generic unit.

#### Static Semantics

Depending on which aspect is identified by the *aspect\_mark*, an *aspect\_definition* specifies:

- a *name* that denotes a subprogram, object, or other kind of entity;
- an *expression* (other than an *aggregate*), which is either evaluated to produce a single value, or which (as in a precondition) is to be evaluated at particular points during later execution;
- an *identifier* specific to the aspect; or
- an *aggregate*, which is positional or named, and is composed of elements of any of these four kinds of constructs.

The identified aspect of the associated entity, or in some cases, the view of the entity defined by the declaration, is as specified by the *aspect\_definition* (or by the default of True when boolean). Whether an *aspect\_specification* *applies* to an entity or only to the particular view of the entity defined by the declaration is determined by the *aspect\_mark* and the kind of entity. The following aspects are view specific:

- An aspect specified on an *object\_declaration*;
- An aspect specified on a *subprogram\_declaration*;
- An aspect specified on a *renaming\_declaration*.

All other *aspect\_specifications* are associated with the entity, and *apply* to all views of the entity, unless otherwise specified in this document.

If the *aspect\_mark* includes 'Class (a *class-wide aspect*), then, unless specified otherwise for a particular class-wide aspect:

- if the associated entity is a tagged type, the specification *applies* to all descendants of the type;
- if the associated entity is a primitive subprogram of a tagged type *T*, the specification *applies* to the corresponding primitive subprogram of all descendants of *T*.

All specifiable operational and representation attributes may be specified with an *aspect\_specification* instead of an *attribute\_definition\_clause* (see 16.3).

Some aspects are defined to be *library unit aspects*. Library unit aspects are of type Boolean. The expression specifying a library unit aspect shall be static. Library unit aspects are defined for all

program units, but shall be specified only for library units. Notwithstanding what this document says elsewhere, the expression of a library unit aspect is resolved and evaluated at the point where it occurs in the `aspect_specification`, rather than the first freezing point of the associated unit.

In addition, other operational and representation aspects not associated with specifiable attributes or representation pragmas may be specified, as specified elsewhere in this document.

If a Legality Rule or Static Semantics rule only applies when a particular aspect has been specified, the aspect is considered to have been specified only when the `aspect_specification` or `attribute_definition_clause` is visible (see 11.3) at the point of the application of the rule.

Alternative legality and semantics rules may apply for particular aspects, as specified elsewhere in this document.

#### *Dynamic Semantics*

At the freezing point of the associated entity, the `aspect_specification` is elaborated. When appearing in a construct other than a declaration, an `aspect_specification` is elaborated as part of the execution of the construct. The elaboration of the `aspect_specification` consists of the elaboration of each `aspect_definition` in an arbitrary order. The elaboration of an `aspect_definition` includes the evaluation of any name or expression that is part of the `aspect_definition` unless the part is itself an expression. If the corresponding aspect (or part thereof) represents an expression (as in a precondition), the elaboration of that part has no effect; the expression is evaluated later at points within the execution as specified elsewhere in this document for the particular aspect.

#### *Implementation Permissions*

Implementations may support implementation-defined aspects. The `aspect_specification` for an implementation-defined aspect may use an implementation-defined syntax for the `aspect_definition`, and may follow implementation-defined legality and semantics rules.

An implementation may ignore the specification of an unrecognized aspect; if an implementation chooses to ignore such an aspect specification (as opposed to rejecting it), then it has no effect on the semantics of the program except for possibly (and this is not required) the rejection of syntax errors within the `aspect_definition`.

## 16.2 Packed Types

The Pack aspect having the value True specifies that storage minimization should be the main criterion when selecting the representation of a composite type.

#### *Static Semantics*

For a full type declaration of a composite type, the following language-defined representation aspect may be specified:

**Pack** The type of aspect Pack is Boolean. When aspect Pack is True for a type, the type (or the extension part) is said to be *packed*. For a type extension, the parent part is packed as for the parent type, and specifying Pack causes packing only of the extension part.

If directly specified, the `aspect_definition` shall be a static expression. If not specified (including by inheritance), the aspect is False.

#### *Implementation Advice*

If a type is packed, then the implementation should try to minimize storage allocated to objects of the type, possibly at the expense of speed of accessing components, subject to reasonable complexity in addressing calculations.

The recommended level of support for the Pack aspect is:

- Any component of a packed type that is of a by-reference type, that is specified as independently addressable, or that contains an aliased part, shall be aligned according to the alignment of its subtype.
- For a packed record type, the components should be packed as tightly as possible subject to the above alignment requirements, the Sizes of the component subtypes, and any `record_representation_clause` that applies to the type; the implementation is allowed to reorder components or cross aligned word boundaries to improve the packing. A component whose Size is greater than the word size may be allocated an integral number of words.
- For a packed array type, if the Size of the component subtype is less than or equal to the word size, `Component_Size` should be less than or equal to the Size of the component subtype, rounded up to the nearest factor of the word size, unless this would violate the above alignment requirements.

### 16.3 Operational and Representation Attributes

The values of certain implementation-dependent characteristics can be obtained by interrogating appropriate operational or representation attributes. Some of these attributes are specifiable via an `attribute_definition_clause`.

#### Syntax

```
attribute_definition_clause ::=
  for local_name'attribute_designator use expression;
| for local_name'attribute_designator use name;
```

#### Name Resolution Rules

For an `attribute_definition_clause` that specifies an attribute that denotes a value, the form with an expression shall be used. Otherwise, the form with a name shall be used.

For an `attribute_definition_clause` that specifies an attribute that denotes a value or an object, the expected type for the expression or name is that of the attribute. For an `attribute_definition_clause` that specifies an attribute that denotes a subprogram, the expected profile for the name is the profile required for the attribute. For an `attribute_definition_clause` that specifies an attribute that denotes some other kind of entity, the name shall resolve to denote an entity of the appropriate kind.

#### Legality Rules

An `attribute_designator` is allowed in an `attribute_definition_clause` only if this document explicitly allows it, or for an implementation-defined attribute if the implementation allows it. Each specifiable attribute constitutes an operational aspect or aspect of representation; the name of the aspect is that of the attribute.

#### Static Semantics

A *Size clause* is an `attribute_definition_clause` whose `attribute_designator` is `Size`. Similar definitions apply to the other specifiable attributes.

A *storage element* is an addressable element of storage in the machine. A *word* is the largest amount of storage that can be conveniently and efficiently manipulated by the hardware, given the implementation's run-time model. A word consists of an integral number of storage elements.

A *machine scalar* is an amount of storage that can be conveniently and efficiently loaded, stored, or operated upon by the hardware. Machine scalars consist of an integral number of storage elements. The set of machine scalars is implementation defined, but includes at least the storage element and the word. Machine scalars are used to interpret `component_clauses` when the nondefault bit ordering applies.

The following representation attributes are defined: `Address`, `Alignment`, `Size`, `Object_Size`, `Storage_Size`, `Component_Size`, `Has_Same_Storage`, and `Overlaps_Storage`.



For a *prefix* X that denotes an object, program unit, or label:

**X'Address** Denotes the address of the first of the storage elements allocated to X. For a program unit or label, this value refers to the machine code associated with the corresponding body or statement. The value of this attribute is of type *System.Address*.

The prefix of X'Address shall not statically denote a subprogram that has convention *Intrinsic*. X'Address raises *Program\_Error* if X denotes a subprogram that has convention *Intrinsic*.

Address may be specified for stand-alone objects and for program units via an *attribute\_definition\_clause*.

#### *Erroneous Execution*

If an Address is specified, it is the programmer's responsibility to ensure that the address is valid and appropriate for the entity and its use; otherwise, program execution is erroneous.

#### *Implementation Advice*

For an array X, X'Address should point at the first component of the array, and not at the array bounds.

The recommended level of support for the Address attribute is:

- X'Address should produce a useful result if X is an object that is aliased or of a by-reference type, or is an entity whose Address has been specified.
- An implementation should support Address clauses for imported subprograms.
- If the Address of an object is specified, or it is imported or exported, then the implementation should not perform optimizations based on assumptions of no aliases.

**NOTE 1** The specification of a link name with the *Link\_Name* aspect (see B.1) for a subprogram or object is an alternative to explicit specification of its link-time address, allowing a link-time directive to place the subprogram or object within memory.

**NOTE 2** The rules for the Size attribute imply, for an aliased object X, that if X'Size = *Storage\_Unit*, then X'Address points at a storage element containing all of the bits of X, and only the bits of X.

#### *Static Semantics*

For a *prefix* X that denotes an object:

**X'Alignment**

The value of this attribute is of type *universal\_integer*, and nonnegative; zero means that the object is not necessarily aligned on a storage element boundary. If X'Alignment is not zero, then X is aligned on a storage unit boundary and X'Address is an integral multiple of X'Alignment (that is, the Address modulo the Alignment is zero).

Alignment may be specified for stand-alone objects via an *attribute\_definition\_clause*; the expression of such a clause shall be static, and its value nonnegative.

For every subtype S:

**S'Alignment**

The value of this attribute is of type *universal\_integer*, and nonnegative.

For an object X of subtype S, if S'Alignment is not zero, then X'Alignment is a nonzero integral multiple of S'Alignment unless specified otherwise by a representation item.

Alignment may be specified for first subtypes via an *attribute\_definition\_clause*; the expression of such a clause shall be static, and its value nonnegative.

#### *Erroneous Execution*

Program execution is erroneous if an Address clause is given that conflicts with the Alignment.



For an object that is not allocated under control of the implementation, execution is erroneous if the object is not aligned according to its Alignment.

*Implementation Advice*

For any tagged specific subtype *S*, *S*'Class'Alignment should equal *S*'Alignment.

The recommended level of support for the Alignment attribute for subtypes is:

- An implementation should support an Alignment clause for a discrete type, fixed point type, record type, or array type, specifying an Alignment value that is zero or a power of two, subject to the following:
- An implementation is not required to support an Alignment clause for a signed integer type specifying an Alignment greater than the largest Alignment value that is ever chosen by default by the implementation for any signed integer type. A corresponding limitation may be imposed for modular integer types, fixed point types, enumeration types, record types, and array types.
- An implementation is not required to support a nonconfirming Alignment clause that can cause the creation of an object of an elementary type that cannot be easily loaded and stored by available machine instructions.
- An implementation is not required to support an Alignment specified for a derived tagged type that is not a multiple of the Alignment of the parent type. An implementation is not required to support a nonconfirming Alignment specified for a derived untagged by-reference type.

The recommended level of support for the Alignment attribute for objects is:

- For stand-alone library-level objects of statically constrained subtypes, the implementation should support all Alignments supported by the target linker. For example, page alignment is likely to be supported for such objects, but not for subtypes.
- For other objects, an implementation should at least support the alignments supported for their subtype, subject to the following:
- An implementation is not required to support Alignments specified for objects of a by-reference type or for objects of types containing aliased subcomponents if the specified Alignment is not a multiple of the Alignment of the subtype of the object.

NOTE 3 Alignment is a subtype-specific attribute.

NOTE 4 A `component_clause`, `Component_Size` clause, or specifying the `Pack` aspect as `True` can override a specified Alignment.

*Static Semantics*

For a prefix *X* that denotes an object:

**X'Size** Denotes the size in bits of the representation of the object. The value of this attribute is of the type *universal\_integer*.

Size may be specified for stand-alone objects via an `attribute_definition_clause`; the expression of such a clause shall be static and its value nonnegative.

*Implementation Advice*

The size of an array object should not include its bounds.

The recommended level of support for the Size attribute of objects is the same as for subtypes (see below), except that only a confirming Size clause is required to be supported for an aliased elementary object.

*Static Semantics*

For every subtype *S*:

**S'Size** If *S* is definite, denotes the size (in bits) that the implementation would choose for the following objects of subtype *S*:

- A record component of subtype *S* when the record type is packed.
- The formal parameter of an instance of *Unchecked\_Conversion* that converts from subtype *S* to some other subtype.

If *S* is indefinite, the meaning is implementation defined. The value of this attribute is of the type *universal\_integer*. The Size of an object is at least as large as that of its subtype, unless the object's Size is determined by a Size clause, a component\_clause, or a Component\_Size clause. Size may be specified for first subtypes via an *attribute\_definition\_clause*; the expression of such a clause shall be static and its value nonnegative.

#### Implementation Requirements

In an implementation, Boolean'Size shall be 1.

#### Implementation Advice

If the Size of a subtype is nonconfirming and allows for efficient independent addressability (see 12.10) on the target architecture, then the Object\_Size of the subtype should have the same value in the absence of an explicit specification of a different value.

A Size clause on a composite subtype should not affect the internal layout of components.

The recommended level of support for the Size attribute of subtypes is:

- The Size (if not specified) of a static discrete or fixed point subtype should be the number of bits necessary to represent each value belonging to the subtype using an unbiased representation, leaving space for a sign bit only if the subtype contains negative values. If such a subtype is a first subtype, then an implementation should support a specified Size for it that reflects this representation.
- For a subtype implemented with levels of indirection, the Size should include the size of the pointers, but not the size of what they point at.
- An implementation should support a Size clause for a discrete type, fixed point type, record type, or array type, subject to the following:
  - An implementation is not required to support a Size clause for a signed integer type specifying a Size greater than that of the largest signed integer type supported by the implementation in the absence of a size clause (that is, when the size is chosen by default). A corresponding limitation may be imposed for modular integer types, fixed point types, enumeration types, record types, and array types.
  - A nonconfirming size clause for the first subtype of a derived untagged by-reference type is not required to be supported.

NOTE 5 Size is a subtype-specific attribute.

NOTE 6 A component\_clause or Component\_Size clause can override a specified Size. Aspect Pack cannot.

#### Static Semantics

For every subtype *S*:

*S*'Object\_Size

If *S* is definite, denotes the size (in bits) of a stand-alone aliased object, or a component of subtype *S* in the absence of an *aspect\_specification* or representation item that specifies the size of the object or component. If *S* is indefinite, the meaning is implementation-defined. The value of this attribute is of the type *universal\_integer*. If not specified otherwise, the Object\_Size of a subtype is at least as large as the Size of the subtype. Object\_Size may be specified via an *attribute\_definition\_clause*; the expression of such a clause shall be static and its value nonnegative. All aliased objects with nominal subtype *S* have the size *S*'Object\_Size. In the absence of an explicit specification, the Object\_Size of a subtype *S* defined by a *subtype\_indication* without a constraint, is

that of the value of the `Object_Size` of the subtype denoted by the `subtype_mark` of the `subtype_indication`, at the point of this definition.

*Implementation Advice*

If `S` is a definite first subtype and `S'Object_Size` is not specified, `S'Object_Size` should be the smallest multiple of the storage element size larger than or equal to `S'Size` that is consistent with the alignment of `S`.

If `X` denotes an object (including a component) of subtype `S`, `X'Size` should equal `S'Object_Size`, unless:

- `X'Size` is specified; or
- `X` is a nonaliased stand-alone object; or
- The size of `X` is determined by a `component_clause` or `Component_Size` clause; or
- The type containing component `X` is packed.

An `Object_Size` clause on a composite type should not affect the internal layout of components.

The recommended level of support for the `Object_Size` attribute of subtypes is:

- If `S` is a static signed integer subtype, the implementation should support the specification of `S'Object_Size` to match the size of any signed integer base subtype provided by the implementation that is no smaller than `S'Size`. Corresponding support is expected for modular integer subtypes, fixed point subtypes, and enumeration subtypes.
- If `S` is an array or record subtype with static constraints and `S` is not a first subtype of a derived untagged by-reference type, the implementation should support the specification of `S'Object_Size` to be any multiple of the storage element size that is consistent with the alignment of `S`, that is no smaller than `S'Size`, and that is no larger than that of the largest composite subtype supported by the implementation.
- If `S` is some other subtype, only confirming specifications of `Object_Size` are required to be supported.

*Static Semantics*

For a prefix `T` that denotes a task object (after any implicit dereference):

`T'Storage_Size`

Denotes the number of storage elements reserved for the task. The value of this attribute is of the type *universal\_integer*. The `Storage_Size` includes the size of the task's stack, if any. The language does not specify whether or not it includes other storage associated with the task (such as the “task control block” used by some implementations.) If the aspect `Storage_Size` is specified for the type of the object, the value of the `Storage_Size` attribute is at least the value determined by the aspect.

Aspect `Storage_Size` specifies the amount of storage to be reserved for the execution of a task.

*Static Semantics*

For a task type (including the anonymous type of a `single_task_declaration`), the following language-defined representation aspect may be specified:

`Storage_Size`

The `Storage_Size` aspect is an *expression*, which shall be of any integer type.

*Legality Rules*

The `Storage_Size` aspect shall not be specified for a task interface type.

*Dynamic Semantics*

When a task object is created, the `expression` (if any) associated with the `Storage_Size` aspect of its type is evaluated; the `Storage_Size` attribute of the newly created task object is at least the value of the `expression`.

At the point of task object creation, or upon task activation, `Storage_Error` is raised if there is insufficient free storage to accommodate the requested `Storage_Size`.

*Static Semantics*

For a `prefix X` that denotes an array subtype or array object (after any implicit dereference):

`X'Component_Size`

Denotes the size in bits of components of the type of `X`. The value of this attribute is of type *universal\_integer*.

`Component_Size` may be specified for array types via an `attribute_definition_clause`; the expression of such a clause shall be static, and its value nonnegative.

*Implementation Advice*

The recommended level of support for the `Component_Size` attribute is:

- An implementation is not required to support specified `Component_Sizes` that are less than the `Size` of the component subtype.
- An implementation should support specified `Component_Sizes` that are factors and multiples of the word size. For such `Component_Sizes`, the array should contain no gaps between components. For other `Component_Sizes` (if supported), the array should contain no gaps between components when `Pack` is also specified; the implementation should forbid this combination in cases where it cannot support a no-gaps representation.

*Static Semantics*

For a `prefix X` that denotes an object:

`X'Has_Same_Storage`

`X'Has_Same_Storage` denotes a function with the following specification:

```
function X'Has_Same_Storage (Arg : any_type)
return Boolean
```

The actual parameter shall be a name that denotes an object. The object denoted by the actual parameter can be of any type. This function evaluates the names of the objects involved. It returns `True` if the representation of the object denoted by the actual parameter occupies exactly the same bits as the representation of the object denoted by `X` and the objects occupy at least one bit; otherwise, it returns `False`.

For a `prefix X` that denotes an object:

`X'Overlaps_Storage`

`X'Overlaps_Storage` denotes a function with the following specification:

```
function X'Overlaps_Storage (Arg : any_type)
return Boolean
```

The actual parameter shall be a name that denotes an object. The object denoted by the actual parameter can be of any type. This function evaluates the names of the objects involved and returns `True` if the representation of the object denoted by the actual parameter shares at least one bit with the representation of the object denoted by `X`; otherwise, it returns `False`.

NOTE 7 `X'Has_Same_Storage(Y)` implies `X'Overlaps_Storage(Y)`.

NOTE 8 `X'Has_Same_Storage(Y)` and `X'Overlaps_Storage(Y)` are not considered to be reads of `X` and `Y`.

*Static Semantics*

The following type-related operational attribute is defined: `External_Tag`.

For every subtype `S` of a tagged type `T` (specific or class-wide):

`S'External_Tag`

`S'External_Tag` denotes an external string representation for `S'Tag`; it is of the predefined type `String`. `External_Tag` may be specified for a specific tagged type via an `attribute_definition_clause`; the expression of such a clause shall be static. The default external tag representation is implementation defined. See 16.13.2. The value of `External_Tag` is never inherited; the default value is always used unless a new value is directly specified for a type.

*Dynamic Semantics*

If a user-specified external tag `S'External_Tag` is the same as `T'External_Tag` for some other tagged type declared by a different declaration in the partition, `Program_Error` is raised by the elaboration of the `attribute_definition_clause`.

*Implementation Requirements*

In an implementation, the default external tag for each specific tagged type declared in a partition shall be distinct, so long as the type is declared outside an instance of a generic body. If the compilation unit in which a given tagged type is declared, and all compilation units on which it semantically depends, are the same in two different partitions, then the external tag for the type shall be the same in the two partitions. What it means for a compilation unit to be the same in two different partitions is implementation defined. At a minimum, if the compilation unit is not recompiled between building the two different partitions that include it, the compilation unit is considered the same in the two partitions.

*Implementation Permissions*

If a user-specified external tag `S'External_Tag` is the same as `T'External_Tag` for some other tagged type declared by a different declaration in the partition, the partition may be rejected.

NOTE 9 The following language-defined attributes are specifiable, at least for some of the kinds of entities to which they apply: `Address`, `Alignment`, `Bit_Order`, `Component_Size`, `External_Tag`, `Input`, `Machine_Radix`, `Output`, `Read`, `Size`, `Small`, `Storage_Pool`, `Storage_Size`, `Stream_Size`, and `Write`.

NOTE 10 It follows from the general rules in 16.1 that if one writes “`for X'Size use Y;`” then the `X'Size attribute_reference` will return `Y` (assuming the implementation allows the `Size` clause). The same is true for all of the specifiable attributes except `Storage_Size`.

EXAMPLE *Examples of attribute definition clauses:*

```
Byte : constant := 8;
Page : constant := 2**12;

type Medium is range 0 .. 65_000;
for Medium'Size use 2*Byte;
for Medium'Alignment use 2;
Device_Register : Medium;
for Device_Register'Size use Medium'Size;
for Device_Register'Address use
  System.Storage_Elements.To_Address(16#FFFF_0020#);

type Short is delta 0.01 range -100.0 .. 100.0;
for Short'Size use 15;

for Car_Name'Storage_Size use -- specify access type's storage pool size
  2000*((Car'Size/System.Storage_Unit) +1); -- approximately 2000 cars

function My_Input(Stream : not null access Ada.Streams.Root_Stream_Type'Class)
  return T;
for T'Input use My_Input; -- see 16.13.2
```

In the Size clause for Short, fifteen bits is the minimum necessary, since the type definition requires Short's Small  $\leq 2^{**}(-7)$ .

## 16.4 Enumeration Representation Clauses

An `enumeration_representation_clause` specifies the internal codes for enumeration literals.

### Syntax

```
enumeration_representation_clause ::=
  for first_subtype_local_name use enumeration_aggregate;
enumeration_aggregate ::= array_aggregate
```

### Name Resolution Rules

The `enumeration_aggregate` shall be written as a one-dimensional `array_aggregate`, for which the index subtype is the unconstrained subtype of the enumeration type, and each component expression is expected to be of any integer type.

### Legality Rules

The `first_subtype_local_name` of an `enumeration_representation_clause` shall denote an enumeration subtype.

Each component of the `array_aggregate` shall be given by an expression rather than a `<>`. The expressions given in the `array_aggregate` shall be static, and shall specify distinct integer codes for each value of the enumeration type; the associated integer codes shall satisfy the predefined ordering relation of the type.

### Static Semantics

An `enumeration_representation_clause` specifies the *coding* aspect of representation. The coding consists of the *internal code* for each enumeration literal, that is, the integral value used internally to represent each literal.

### Implementation Requirements

For nonboolean enumeration types, if the coding is not specified for the type, then for each value of the type, the internal code shall be equal to its position number.

### Implementation Advice

The recommended level of support for `enumeration_representation_clauses` is:

- An implementation should support at least the internal codes in the range `System.Min_Int .. System.Max_Int`. An implementation is not required to support `enumeration_representation_clauses` for boolean types.

### Static Semantics

For every discrete subtype `S`, the following attributes are defined:

`S'Enum_Rep`

`S'Enum_Rep` denotes a function with the following specification:

```
function S'Enum_Rep (Arg : S'Base) return universal_integer
```

This function returns the representation value of the value of `Arg`, as a value of type *universal\_integer*. The *representation value* is the internal code specified in an enumeration representation clause, if any, for the type corresponding to the value of `Arg`, and otherwise is the position number of the value.

`S'Enum_Val`

`S'Enum_Val` denotes a function with the following specification:

```
function S'Enum_Val (Arg : universal_integer) return S'Base
```



This function returns a value of the type of *S* whose representation value equals the value of *Arg*. For the evaluation of a call on *S'Enum\_Val*, if there is no value in the base range of its type with the given representation value, *Constraint\_Error* is raised.

NOTE Attribute *Enum\_Rep* can be used to query the internal codes used for an enumeration type; attribute *Enum\_Val* can be used to convert from an internal code to an enumeration value. The other attributes of the type, such as *Succ*, *Pred*, and *Pos*, are unaffected by an *enumeration\_representation\_clause*. For example, *Pos* always returns the position number, *not* an internal integer code that was specified in an *enumeration\_representation\_clause*.

EXAMPLE *Examples of enumeration representation clauses:*

```
type Mix_Code is (ADD, SUB, MUL, LDA, STA, STZ);

for Mix_Code use
  (ADD => 1, SUB => 2, MUL => 3, LDA => 8, STA => 24, STZ =>33);

-- See 6.5.2.
for Roman_Digit use ('I' => 1,
                      'V' => 5,
                      'X' => 10,
                      'L' => 50,
                      'C' => 100,
                      'D' => 500,
                      'M' => 1000);

-- For an example of the use of attribute Enum_Rep, see 7.2.1.
```

## 16.5 Record Layout

The (*record*) *layout* aspect of representation consists of the *storage places* for some or all components, that is, storage place attributes of the components. The layout can be specified with a *record\_representation\_clause*.

### 16.5.1 Record Representation Clauses

A *record\_representation\_clause* specifies the storage representation of records and record extensions, that is, the order, position, and size of components (including discriminants, if any).

*Syntax*

```
record_representation_clause ::=
  for first_subtype_local_name use
    record [mod_clause]
      {component_clause}
    end record [local_name];

component_clause ::=
  component_local_name at position range first_bit..last_bit;

position ::= static_expression

first_bit ::= static_simple_expression

last_bit ::= static_simple_expression
```

If a *local\_name* appears at the end of the *record\_representation\_clause*, it shall repeat the *first\_subtype\_local\_name*.

*Name Resolution Rules*

Each *position*, *first\_bit*, and *last\_bit* is expected to be of any integer type.

*Legality Rules*

The *first\_subtype\_local\_name* of a *record\_representation\_clause* shall denote a specific record or record extension subtype.

If the *component\_local\_name* is a *direct\_name*, the *local\_name* shall denote a component of the type. For a record extension, the component shall not be inherited, and shall not be a discriminant that corresponds to a discriminant of the parent type. If the *component\_local\_name* has an *attribute\_designator*, the *direct\_name* of the *local\_name* shall denote either the declaration of the type or a component of the type, and the *attribute\_designator* shall denote an implementation-defined implicit component of the type.

The *position*, *first\_bit*, and *last\_bit* shall be static expressions. The value of *position* and *first\_bit* shall be nonnegative. The value of *last\_bit* shall be no less than *first\_bit* - 1.

If the nondefault bit ordering applies to the type, then either:

- the value of *last\_bit* shall be less than the size of the largest machine scalar; or
- the value of *first\_bit* shall be zero and the value of *last\_bit* + 1 shall be a multiple of *System.Storage\_Unit*.

At most one *component\_clause* is allowed for each component of the type, including for each discriminant (*component\_clauses* may be given for some, all, or none of the components). Storage places within a *component\_list* shall not overlap, unless they are for components in distinct variants of the same *variant\_part*.

A name that denotes a component of a type is not allowed within a *record\_representation\_clause* for the type, except as the *component\_local\_name* of a *component\_clause*.

#### Static Semantics

A *record\_representation\_clause* (without the *mod\_clause*) specifies the layout.

If the default bit ordering applies to the type, the *position*, *first\_bit*, and *last\_bit* of each *component\_clause* directly specify the position and size of the corresponding component.

If the nondefault bit ordering applies to the type, then the layout is determined as follows:

- the *component\_clauses* for which the value of *last\_bit* is greater than or equal to the size of the largest machine scalar directly specify the position and size of the corresponding component;
- for other *component\_clauses*, all of the components having the same value of *position* are considered to be part of a single machine scalar, located at that *position*; this machine scalar has a size which is the smallest machine scalar size larger than the largest *last\_bit* for all *component\_clauses* at that *position*; the *first\_bit* and *last\_bit* of each *component\_clause* are then interpreted as bit offsets in this machine scalar.

A *record\_representation\_clause* for a record extension does not override the layout of the parent part; if the layout was specified for the parent type, it is inherited by the record extension.

#### Implementation Permissions

An implementation may generate implementation-defined components (for example, one containing the offset of another component). An implementation may generate names that denote such implementation-defined components; such names shall be implementation-defined *attribute\_references*. An implementation may allow such implementation-defined names to be used in *record\_representation\_clauses*. An implementation can restrict such *component\_clauses* in any manner it sees fit.

If a *record\_representation\_clause* is given for an untagged derived type, the storage place attributes for all of the components of the derived type may differ from those of the corresponding components of the parent type, even for components whose storage place is not specified explicitly in the *record\_representation\_clause*.

The recommended level of support for `record_representation_clauses` is:

- An implementation should support machine scalars that correspond to all of the integer, floating point, and address formats supported by the machine.
- An implementation should support storage places that can be extracted with a load, mask, shift sequence of machine code, and set with a load, shift, mask, store sequence, given the available machine instructions and run-time model.
- A storage place should be supported if its size is equal to the Size of the component subtype, and it starts and ends on a boundary that obeys the Alignment of the component subtype.
- For a component with a subtype whose Size is less than the word size, any storage place that does not cross an aligned word boundary should be supported.
- An implementation may reserve a storage place for the tag field of a tagged type, and disallow other components from overlapping that place.
- An implementation is not required to support a `component_clause` for a component of an extension part if the storage place is not after the storage places of all components of the parent type, whether or not those storage places had been specified.

NOTE If no `component_clause` is given for a component, then the choice of the storage place for the component is left to the implementation. If `component_clauses` are given for all components, the `record_representation_clause` completely specifies the representation of the type and will be obeyed exactly by the implementation.

EXAMPLE *Example of specifying the layout of a record type:*

```
Word : constant := 4; -- storage element is byte, 4 bytes per word

type State is (A,M,W,P);
type Mode is (Fix, Dec, Exp, Signif);

type Byte_Mask is array (0..7) of Boolean with Component_Size => 1;
type State_Mask is array (State) of Boolean with Component_Size => 1;
type Mode_Mask is array (Mode) of Boolean with Component_Size => 1;

type Program_Status_Word is
  record
    System_Mask : Byte_Mask;
    Protection_Key : Integer range 0 .. 3;
    Machine_State : State_Mask;
    Interrupt_Cause : Interruption_Code;
    Ilc : Integer range 0 .. 3;
    Cc : Integer range 0 .. 3;
    Program_Mask : Mode_Mask;
    Inst_Address : Address;
  end record;

for Program_Status_Word use
  record
    System_Mask at 0*Word range 0 .. 7;
    Protection_Key at 0*Word range 10 .. 11; -- bits 8,9 unused
    Machine_State at 0*Word range 12 .. 15;
    Interrupt_Cause at 0*Word range 16 .. 31;
    Ilc at 1*Word range 0 .. 1; -- second word
    Cc at 1*Word range 2 .. 3;
    Program_Mask at 1*Word range 4 .. 7;
    Inst_Address at 1*Word range 8 .. 31;
  end record;

for Program_Status_Word'Size use 8*System.Storage_Unit;
for Program_Status_Word'Alignment use 8;
```

The `record_representation_clause` defines the record layout. The Size clause guarantees that (at least) eight storage elements are used for objects of the type. The Alignment clause guarantees that aliased, imported, or exported objects of the type will have addresses divisible by eight.

## 16.5.2 Storage Place Attributes

### Static Semantics

For a component C of a composite, non-array object R, the *storage place attributes* are defined:

#### R.C'Position

If the nondefault bit ordering applies to the composite type, and if a `component_clause` specifies the placement of C, denotes the value given for the `position` of the `component_clause`; otherwise, denotes the same value as `R.C'Address - R'Address`. The value of this attribute is of the type *universal\_integer*.

#### R.C'First\_Bit

If the nondefault bit ordering applies to the composite type, and if a `component_clause` specifies the placement of C, denotes the value given for the `first_bit` of the `component_clause`; otherwise, denotes the offset, from the start of the first of the storage elements occupied by C, of the first bit occupied by C. This offset is measured in bits. The first bit of a storage element is numbered zero. The value of this attribute is of the type *universal\_integer*.

#### R.C'Last\_Bit

If the nondefault bit ordering applies to the composite type, and if a `component_clause` specifies the placement of C, denotes the value given for the `last_bit` of the `component_clause`; otherwise, denotes the offset, from the start of the first of the storage elements occupied by C, of the last bit occupied by C. This offset is measured in bits. The value of this attribute is of the type *universal\_integer*.

### Implementation Advice

If a component is represented using some form of pointer (such as an offset) to the actual data of the component, and this data is contiguous with the rest of the object, then the storage place attributes should reflect the place of the actual data, not the pointer. If a component is allocated discontinuously from the rest of the object, then a warning should be generated upon reference to one of its storage place attributes.

## 16.5.3 Bit Ordering

The `Bit_Order` attribute specifies the interpretation of the storage place attributes.

### Static Semantics

A bit ordering is a method of interpreting the meaning of the storage place attributes. `High_Order_First` (known in the vernacular as “big endian”) means that the first bit of a storage element (bit 0) is the most significant bit (interpreting the sequence of bits that represent a component as an unsigned integer value). `Low_Order_First` (known in the vernacular as “little endian”) means the opposite: the first bit is the least significant.

For every specific record subtype S, the following representation attribute is defined:

#### S'Bit\_Order

Denotes the bit ordering for the type of S. The value of this attribute is of type `System.Bit_Order`. `Bit_Order` may be specified for specific record types via an `attribute_definition_clause`; the expression of such a clause shall be static.

If `Word_Size = Storage_Unit`, the default bit ordering is implementation defined. If `Word_Size > Storage_Unit`, the default bit ordering is the same as the ordering of storage elements in a word, when interpreted as an integer.

The storage place attributes of a component of a type are interpreted according to the bit ordering of the type.

The recommended level of support for the nondefault bit ordering is:

- The implementation should support the nondefault bit ordering in addition to the default bit ordering.

NOTE Bit\_Order clauses make it possible to write `record_representation_clauses` that can be ported between machines having different bit ordering. They do not guarantee transparent exchange of data between such machines.

## 16.6 Change of Representation

A `type_conversion` (see 7.6) can be used to convert between two different representations of the same array or record. To convert an array from one representation to another, two array types with matching component subtypes and convertible index types are required. If one type has `Pack` specified and the other does not, then explicit conversion can be used to pack or unpack an array.

To convert an untagged record from one representation to another, two record types with a common ancestor type are required. Distinct representations can then be specified for the record types, and explicit conversion between the types can be used to effect a change in representation.

EXAMPLE *Example of change of representation:*

```
-- Packed_Descriptor and Descriptor are two different types
-- with identical characteristics, apart from their
-- representation

type Descriptor is
  record
    -- components of a descriptor
  end record;

type Packed_Descriptor is new Descriptor;

for Packed_Descriptor use
  record
    -- component clauses for some or for all components
  end record;

-- Change of representation can now be accomplished by explicit type
-- conversions:

D : Descriptor;
P : Packed_Descriptor;

P := Packed_Descriptor(D);  -- pack D
D := Descriptor(P);        -- unpack P
```

## 16.7 The Package System

For each implementation there is a library package called `System` which includes the definitions of certain configuration-dependent characteristics.

The following language-defined library package exists:

```
package System
  with Pure is

  type Name is implementation-defined-enumeration-type;
  System_Name : constant Name := implementation-defined;

  -- System-Dependent Named Numbers:

  Min_Int      : constant := root_integer'First;
  Max_Int      : constant := root_integer'Last;

  Max_Binary_Modulus : constant := implementation-defined;
  Max_Nonbinary_Modulus : constant := implementation-defined;
```

```

Max_Base_Digits      : constant := root_real'Digits;
Max_Digits           : constant := implementation-defined;
Max_Mantissa         : constant := implementation-defined;
Fine_Delta           : constant := implementation-defined;
Tick                 : constant := implementation-defined;

-- Storage-related Declarations:

type Address is implementation-defined;
Null_Address : constant Address;

Storage_Unit : constant := implementation-defined;
Word_Size    : constant := implementation-defined * Storage_Unit;
Memory_Size  : constant := implementation-defined;

-- Address Comparison:
function "<" (Left, Right : Address) return Boolean
    with Convention => Intrinsic;
function "<=" (Left, Right : Address) return Boolean
    with Convention => Intrinsic;
function ">" (Left, Right : Address) return Boolean
    with Convention => Intrinsic;
function ">=" (Left, Right : Address) return Boolean
    with Convention => Intrinsic;
function "=" (Left, Right : Address) return Boolean
    with Convention => Intrinsic;
-- function "/=" (Left, Right : Address) return Boolean;
-- "/=" is implicitly defined

-- Other System-Dependent Declarations:
type Bit_Order is (High_Order_First, Low_Order_First);
Default_Bit_Order : constant Bit_Order := implementation-defined;

-- Priority-related declarations (see D.1):
subtype Any_Priority is Integer range implementation-defined;
subtype Priority is Any_Priority range Any_Priority'First ..
    implementation-defined;
subtype Interrupt_Priority is Any_Priority range Priority'Last+1 ..
    Any_Priority'Last;

Default_Priority : constant Priority :=
    (Priority'First + Priority'Last)/2;

private
    ... -- not specified by the language
end System;

```

Name is an enumeration subtype. Values of type Name are the names of alternative machine configurations handled by the implementation. System\_Name represents the current machine configuration.

The named numbers Fine\_Delta and Tick are of the type *universal\_real*; the others are of the type *universal\_integer*.

The meanings of the named numbers are:

**Min\_Int**    The smallest (most negative) value allowed for the expressions of a *signed\_integer\_type\_definition*.

**Max\_Int**    The largest (most positive) value allowed for the expressions of a *signed\_integer\_type\_definition*.

**Max\_Binary\_Modulus**

A power of two such that it, and all lesser positive powers of two, are allowed as the modulus of a *modular\_type\_definition*.

**Max\_Nonbinary\_Modulus**

A value such that it, and all lesser positive integers, are allowed as the modulus of a *modular\_type\_definition*.



Max\_Base\_Digits

The largest value allowed for the requested decimal precision in a `floating_point_`-definition.

Max\_Digits

The largest value allowed for the requested decimal precision in a `floating_point_`-definition that has no `real_range_specification`. `Max_Digits` is less than or equal to `Max_Base_Digits`.

Max\_Mantissa

The largest possible number of binary digits in the mantissa of machine numbers of a user-defined ordinary fixed point type. (The mantissa is defined in Annex G.)

Fine\_Delta

The smallest delta allowed in an `ordinary_fixed_point_definition` that has the `real_range_specification` **range** `-1.0 .. 1.0`.

Tick

A period in seconds approximating the real time interval during which the value of `Calendar.Clock` remains constant.

Storage\_Unit

The number of bits per storage element.

Word\_Size

The number of bits per word.

Memory\_Size

An implementation-defined value that is intended to reflect the memory size of the configuration in storage elements.

Address is a definite, nonlimited type with preelaborable initialization (see 13.2.1). Address represents machine addresses capable of addressing individual storage elements. `Null_Address` is an address that is distinct from the address of any object or program unit.

`Default_Bit_Order` shall be a static constant. See 16.5.3 for an explanation of `Bit_Order` and `Default_Bit_Order`.

*Implementation Permissions*

An implementation may add additional implementation-defined declarations to package `System` and its children. However, it is usually better for the implementation to provide additional functionality via implementation-defined children of `System`.

*Implementation Advice*

Address should be a private type.

NOTE There are also some language-defined child packages of `System` defined elsewhere.

## 16.7.1 The Package `System.Storage_Elements`

*Static Semantics*

The following language-defined library package exists:

```
package System.Storage_Elements
with Pure is

  type Storage_Offset is range implementation-defined;
  subtype Storage_Count is Storage_Offset range 0..Storage_Offset'Last;
  type Storage_Element is mod implementation-defined;
  for Storage_Element'Size use Storage_Unit;
  type Storage_Array is array
    (Storage_Offset range <>) of aliased Storage_Element;
  for Storage_Array'Component_Size use Storage_Unit;
```

```

-- Address Arithmetic:
function "+"(Left : Address; Right : Storage_Offset) return Address
  with Convention => Intrinsic;
function "+"(Left : Storage_Offset; Right : Address) return Address
  with Convention => Intrinsic;
function "-"(Left : Address; Right : Storage_Offset) return Address
  with Convention => Intrinsic;
function "-"(Left, Right : Address) return Storage_Offset
  with Convention => Intrinsic;

function "mod"(Left : Address; Right : Storage_Offset)
  return Storage_Offset
  with Convention => Intrinsic;

-- Conversion to/from integers:

type Integer_Address is implementation-defined;
function To_Address(Value : Integer_Address) return Address
  with Convention => Intrinsic;
function To_Integer(Value : Address) return Integer_Address
  with Convention => Intrinsic;

end System.Storage_Elements;

```

Storage\_Element represents a storage element. Storage\_Offset represents an offset in storage elements. Storage\_Count represents a number of storage elements. Storage\_Array represents a contiguous sequence of storage elements.

Integer\_Address is a (signed or modular) integer subtype. To\_Address and To\_Integer convert back and forth between this type and Address.

#### Implementation Requirements

Storage\_Offset'Last shall be greater than or equal to Integer'Last or the largest possible storage offset, whichever is smaller. Storage\_Offset'First shall be  $\leq (-\text{Storage\_Offset'Last})$ .

#### Implementation Advice

Operations in System and its children should reflect the target environment semantics as closely as is reasonable. For example, on most machines, it makes sense for address arithmetic to “wrap around”. Operations that do not make sense should raise Program\_Error.

### 16.7.2 The Package System.Address\_To\_Access\_Conversions

#### Static Semantics

The following language-defined generic library package exists:

```

generic
  type Object(<>) is limited private;
package System.Address_To_Access_Conversions
  with Preelaborate, Nonblocking, Global => in out synchronized is
  type Object_Pointer is access all Object;
  function To_Pointer(Value : Address) return Object_Pointer
    with Convention => Intrinsic;
  function To_Address(Value : Object_Pointer) return Address
    with Convention => Intrinsic;

end System.Address_To_Access_Conversions;

```

The To\_Pointer and To\_Address subprograms convert back and forth between values of types Object\_Pointer and Address. To\_Pointer(X'Address) is equal to X'Unchecked\_Access for any X that allows Unchecked\_Access. To\_Pointer(Null\_Address) returns null. For other addresses, the behavior is unspecified. To\_Address(null) returns Null\_Address. To\_Address(Y), where Y  $\neq$  null, returns Y.all'Address.

*Implementation Permissions*

An implementation may place restrictions on instantiations of `Address_To_Access_Conversions`.

## 16.8 Machine Code Insertions

A machine code insertion can be achieved by a call to a subprogram whose `sequence_of_statements` contains `code_statements`.

*Syntax*

```
code_statement ::= qualified_expression;
```

A `code_statement` is only allowed in the `handled_sequence_of_statements` of a `subprogram_body`. If a `subprogram_body` contains any `code_statements`, then within this `subprogram_body` the only allowed form of statement is a `code_statement` (labeled or not), the only allowed declarative items are `use_clauses`, and no exception handler is allowed (comments and pragmas are allowed as usual).

*Name Resolution Rules*

The `qualified_expression` is expected to be of any type.

*Legality Rules*

The `qualified_expression` shall be of a type declared in package `System.Machine_Code`.

A `code_statement` shall appear only within the scope of a `with_clause` that mentions package `System.Machine_Code`.

*Static Semantics*

The contents of the library package `System.Machine_Code` (if provided) are implementation defined. The meaning of `code_statements` is implementation defined. Typically, each `qualified_expression` represents a machine instruction or assembly directive.

*Implementation Permissions*

An implementation may place restrictions on `code_statements`. An implementation is not required to provide package `System.Machine_Code`.

NOTE 1 An implementation can provide implementation-defined pragmas specifying register conventions and calling conventions.

NOTE 2 Machine code functions are exempt from the rule that a return statement is required. In fact, return statements are forbidden, since only `code_statements` are allowed.

NOTE 3 Intrinsic subprograms (see 9.3.1) can also be used to achieve machine code insertions. Interface to assembly language can be achieved using the features in Annex B.

EXAMPLE *Example of a code statement:*

```
M : Mask;
procedure Set_Mask
  with Inline;

procedure Set_Mask is
  use System.Machine_Code; -- assume "with System.Machine_Code;" appears
  somewhere above
begin
  SI_Format'(Code => SSM, B => M'Base_Reg, D => M'Disp);
  -- Base_Reg and Disp are implementation-defined attributes
end Set_Mask;
```

## 16.9 Unchecked Type Conversions

An unchecked type conversion can be achieved by a call to an instance of the generic function `Unchecked_Conversion`.

### *Static Semantics*

The following language-defined generic library function exists:

```
generic
  type Source(<>) is limited private;
  type Target(<>) is limited private;
function Ada.Unchecked_Conversion(S : Source) return Target
  with Pure, Nonblocking, Convention => Intrinsic;
```

### *Dynamic Semantics*

The size of the formal parameter `S` in an instance of `Unchecked_Conversion` is that of its subtype. This is the actual subtype passed to `Source`, except when the actual is an unconstrained composite subtype, in which case the subtype is constrained by the bounds or discriminants of the value of the actual expression passed to `S`.

If all of the following are true, the effect of an unchecked conversion is to return the value of an object of the target subtype whose representation is the same as that of the source object `S`:

- `S`'Size = `Target`'Size.
- `S`'Alignment is a multiple of `Target`'Alignment or `Target`'Alignment is zero.
- The target subtype is not an unconstrained composite subtype.
- `S` and the target subtype both have a contiguous representation.
- The representation of `S` is a representation of an object of the target subtype.

Otherwise, if the result type is scalar, the result of the function is implementation defined, and can have an invalid representation (see 16.9.1). If the result type is nonscalar, the effect is implementation defined; in particular, the result can be abnormal (see 16.9.1).

### *Implementation Permissions*

An implementation may return the result of an unchecked conversion by reference, if the `Source` type is not a by-copy type. In this case, the result of the unchecked conversion represents simply a different (read-only) view of the operand of the conversion.

An implementation may place restrictions on `Unchecked_Conversion`.

### *Implementation Advice*

Since the Size of an array object generally does not include its bounds, the bounds should not be part of the converted data.

The implementation should not generate unnecessary runtime checks to ensure that the representation of `S` is a representation of the target type. It should take advantage of the permission to return by reference when possible. Restrictions on unchecked conversions should be avoided unless required by the target environment.

The recommended level of support for unchecked conversions is:

- Unchecked conversions should be supported and should be reversible in the cases where this subclause defines the result. To enable meaningful use of unchecked conversion, a contiguous representation should be used for elementary subtypes, for statically constrained array subtypes whose component subtype is one of the subtypes described in this paragraph, and for record subtypes without discriminants whose component subtypes are described in this paragraph.

### 16.9.1 Data Validity

Certain actions that can potentially lead to erroneous execution are not directly erroneous, but instead can cause objects to become *abnormal*. Subsequent uses of abnormal objects can be erroneous.

A scalar object can have an *invalid representation*, which means that the object's representation does not represent any value of the object's subtype. The primary cause of invalid representations is uninitialized variables.

Abnormal objects and invalid representations are explained in this subclause.

#### Dynamic Semantics

When an object is first created, and any explicit or default initializations have been performed, the object and all of its parts are in the *normal* state. Subsequent operations generally leave them normal. However, an object or part of an object can become *abnormal* in the following ways:

- An assignment to the object is disrupted due to an abort (see 12.8) or due to the failure of a language-defined check (see 14.6).
- The object is not scalar, and is passed to an **in out** or **out** parameter of an imported procedure, the Read procedure of an instance of Sequential\_IO, Direct\_IO, or Storage\_IO, or the stream attribute T'Read, if after return from the procedure the representation of the parameter does not represent a value of the parameter's subtype.
- The object is the return object of a function call of a nonscalar type, and the function is an imported function, an instance of Unchecked\_Conversion, or the stream attribute T'Input, if after return from the function the representation of the return object does not represent a value of the function's subtype.

For an imported object, it is the programmer's responsibility to ensure that the object remains in a normal state.

Whether or not an object actually becomes abnormal in these cases is not specified. An abnormal object becomes normal again upon successful completion of an assignment to the object as a whole.

#### Erroneous Execution

It is erroneous to evaluate a `primary` that is a name denoting an abnormal object, or to evaluate a `prefix` that denotes an abnormal object.

#### Bounded (Run-Time) Errors

If the representation of a scalar object does not represent a value of the object's subtype (perhaps because the object was not initialized), the object is said to have an *invalid representation*. It is a bounded error to evaluate the value of such an object. If the error is detected, either Constraint\_Error or Program\_Error is raised. Otherwise, execution continues using the invalid representation. The rules of the language outside this subclause assume that all objects have valid representations. The semantics of operations on invalid representations are as follows:

- If the representation of the object represents a value of the object's type, the value of the type is used.
- If the representation of the object does not represent a value of the object's type, the semantics of operations on such representations is implementation-defined, but does not by itself lead to erroneous or unpredictable execution, or to other objects becoming abnormal.

#### Erroneous Execution

A call to an imported function or an instance of Unchecked\_Conversion is erroneous if the result is scalar, the result object has an invalid representation, and the result is used other than as the expression of an `assignment_statement` or an `object_declaration`, as the *object\_name* of an `object_renaming_declaration`, or as the `prefix` of a Valid attribute. If such a result object is used as the source of an assignment, and the assigned value is an invalid representation for the target

of the assignment, then any use of the target object prior to a further assignment to the target object, other than as the `prefix` of a Valid attribute reference, is erroneous.

The dereference of an access value is erroneous if it does not designate an object of an appropriate type or a subprogram with an appropriate profile, if it designates a nonexistent object, or if it is an access-to-variable value that designates a constant object and it did not originate from an `attribute_reference` applied to an aliased variable view of a controlled or immutably limited object. An access value whose dereference is erroneous can exist, for example, because of `Unchecked_Deallocation`, `Unchecked_Access`, or `Unchecked_Conversion`.

NOTE Objects can become abnormal due to other kinds of actions that directly update the object's representation; such actions are generally considered directly erroneous, however.

### 16.9.2 The Valid Attribute

The Valid attribute can be used to check the validity of data produced by unchecked conversion, input, interface to foreign languages, and the like.

#### *Static Semantics*

For a `prefix` X that denotes a scalar object (after any implicit dereference), the following attribute is defined:

**X'Valid** Yields True if and only if the object denoted by X is normal, has a valid representation, and then, if the preceding conditions hold, the value of X also satisfies the predicates of the nominal subtype of X. The value of this attribute is of the predefined type Boolean.

NOTE 1 Invalid data can be created in the following cases (not counting erroneous or unpredictable execution):

- an uninitialized scalar object,
- the result of an unchecked conversion,
- input,
- interface to another language (including machine code),
- aborting an assignment,
- disrupting an assignment due to the failure of a language-defined check (see 14.6), and
- use of an object whose Address has been specified.

NOTE 2 Determining whether X is normal and has a valid representation as part of the evaluation of X'Valid is not considered to include an evaluation of X; hence, it is not an error to check the validity of an object that is invalid or abnormal. Determining whether X satisfies the predicates of its nominal subtype can include an evaluation of X, but only after it has been determined that X has a valid representation.

If X is volatile, the evaluation of X'Valid is considered a read of X.

NOTE 3 The Valid attribute can be used to check the result of calling an instance of `Unchecked_Conversion` (or any other operation that can return invalid values). However, an exception handler is still useful because implementations are permitted to raise `Constraint_Error` or `Program_Error` if they detect the use of an invalid representation (see 16.9.1).

### 16.10 Unchecked Access Value Creation

The attribute `Unchecked_Access` is used to create access values in an unsafe manner — the programmer is responsible for preventing “dangling references”.

#### *Static Semantics*

The following attribute is defined for a `prefix` X that denotes an aliased view of an object:

**X'Unchecked\_Access**

All rules and semantics that apply to X'Access (see 6.10.2) apply also to X'Unchecked\_Access, except that, for the purposes of accessibility rules and checks, it is as if X were declared immediately within a library package.



NOTE 1 This attribute is provided to support the situation where a local object is to be inserted into a global linked data structure, when the programmer knows that it will always be removed from the data structure prior to exiting the object's scope. The Access attribute would be illegal in this case (see 6.10.2).

NOTE 2 There is no Unchecked\_Access attribute for subprograms.

## 16.11 Storage Management

Each access-to-object type has an associated storage pool. The storage allocated by an allocator comes from the pool; instances of Unchecked\_Deallocation return storage to the pool. Several access types can share the same pool.

A storage pool is a variable of a type in the class rooted at Root\_Storage\_Pool, which is an abstract limited controlled type. By default, the implementation chooses a *standard storage pool* for each access-to-object type. The user may define new pool types, and may override the choice of pool for an access-to-object type by specifying Storage\_Pool for the type.

### Legality Rules

If Storage\_Pool is specified for a given access type, Storage\_Size shall not be specified for it.

### Static Semantics

The following language-defined library package exists:

```
with Ada.Finalization;
with System.Storage_Elements;
package System.Storage_Pools
  with Pure, Nonblocking => False is
  type Root_Storage_Pool is
    abstract new Ada.Finalization.Limited_Controlled with private
    with Preelaborable_Initialization;

  procedure Allocate(
    Pool : in out Root_Storage_Pool;
    Storage_Address : out Address;
    Size_In_Storage_Elements : in Storage_Elements.Storage_Count;
    Alignment : in Storage_Elements.Storage_Count) is abstract;

  procedure Deallocate(
    Pool : in out Root_Storage_Pool;
    Storage_Address : in Address;
    Size_In_Storage_Elements : in Storage_Elements.Storage_Count;
    Alignment : in Storage_Elements.Storage_Count) is abstract;

  function Storage_Size(Pool : Root_Storage_Pool)
    return Storage_Elements.Storage_Count is abstract;

private
  ... -- not specified by the language
end System.Storage_Pools;
```

A *storage pool type* (or *pool type*) is a descendant of Root\_Storage\_Pool. The *elements* of a storage pool are the objects allocated in the pool by allocators.

For every access-to-object subtype S, the following representation attributes are defined:

S'Storage\_Pool

Denotes the storage pool of the type of S. The type of this attribute is Root\_Storage\_Pool'Class.

S'Storage\_Size

Yields the result of calling Storage\_Size(S'Storage\_Pool), which is intended to be a measure of the number of storage elements reserved for the pool. The type of this attribute is *universal\_integer*.

`Storage_Size` or `Storage_Pool` may be specified for a nonderived access-to-object type via an `attribute_definition_clause`; the name in a `Storage_Pool` clause shall denote a variable. If the nominal subtype of the name specified for `Storage_Pool` is nonblocking (see 12.5), then the primitive `Allocate`, `Deallocate`, and `Storage_Size` subprograms of that type shall be nonblocking. Additionally, if the pool is one that supports subpools (see 16.11.4), the primitive `Default_Subpool_for_Pool`, `Allocate_From_Subpool`, and `Deallocate_Subpool` subprograms shall be nonblocking.

An `allocator` of a type  $T$  that does not support subpools allocates storage from  $T$ 's storage pool. If the storage pool is a user-defined object, then the storage is allocated by calling `Allocate` as described below. `Allocators` for types that support subpools are described in 16.11.4.

If `Storage_Pool` is not specified for a type defined by an `access_to_object_definition`, then the implementation chooses a standard storage pool for it in an implementation-defined manner. In this case, the exception `Storage_Error` is raised by an `allocator` if there is not enough storage. It is implementation defined whether or not the implementation provides user-accessible names for the standard pool type(s).

The type(s) of the standard pool(s), and the primitive `Allocate`, `Deallocate`, and `Storage_Size` subprograms for the standard pool(s) are nonblocking. Concurrent invocations of these subprograms do not conflict with one another (see 12.10) when applied to standard storage pools.

If `Storage_Size` is specified for an access type  $T$ , an implementation defined pool  $P$  is used for the type. The `Storage_Size` of  $P$  is at least that requested, and the storage for  $P$  is reclaimed when the master containing the declaration of the access type is left. If the implementation cannot satisfy the request, `Storage_Error` is raised at the freezing point of type  $T$ . The storage pool  $P$  is used only for `allocators` returning type  $T$  or other access types specified to use  $T$ 's `Storage_Pool`. `Storage_Error` is raised by an `allocator` returning such a type if the storage space of  $P$  is exhausted (additional memory is not allocated). The type of  $P$ , and the primitive `Allocate`, `Deallocate`, and `Storage_Size` subprograms of  $P$  are nonblocking.

If neither `Storage_Pool` nor `Storage_Size` are specified, then the meaning of `Storage_Size` is implementation defined.

If `Storage_Pool` is specified for an access type, then the specified pool is used.

The effect of calling `Allocate` and `Deallocate` for a standard storage pool directly (rather than implicitly via an `allocator` or an instance of `Unchecked_Deallocation`) is unspecified.

#### *Erroneous Execution*

If `Storage_Pool` is specified for an access type, then if `Allocate` can satisfy the request, it should allocate a contiguous block of memory, and return the address of the first storage element in `Storage_Address`. The block should contain `Size_In_Storage_Elements` storage elements, and should be aligned according to `Alignment`. The allocated storage should not be used for any other purpose while the pool element remains in existence. If the request cannot be satisfied, then `Allocate` should propagate an exception (such as `Storage_Error`). If `Allocate` behaves in any other manner, then the program execution is erroneous.

#### *Implementation Requirements*

The `Allocate` procedure of a user-defined storage pool object  $P$  may be called by the implementation only to allocate storage for a type  $T$  whose pool is  $P$ , only at the following points:

- During the execution of an `allocator` of type  $T$ ;
- During the execution of a return statement for a function whose result is built-in-place in the result of an `allocator` of type  $T$ ;
- During the execution of an assignment operation with a target of an allocated object of type  $T$  with a part that has an unconstrained discriminated subtype with defaults.

For each of the calls of `Allocate` described above,  $P$  (equivalent to  $T$ 's `Storage_Pool`) is passed as the `Pool` parameter. The `Size_In_Storage_Elements` parameter indicates the number of storage elements to be

allocated, and is no more than  $D'_{\text{Max\_Size\_In\_Storage\_Elements}}$ , where  $D$  is the designated subtype of  $T$ . The Alignment parameter is a nonzero integral multiple of  $D'_{\text{Alignment}}$  if  $D$  is a specific type, and otherwise is a nonzero integral multiple of the alignment of the specific type identified by the tag of the object being created; it is unspecified if there is no such value. The Alignment parameter is no more than  $D'_{\text{Max\_Alignment\_For\_Allocation}}$ . The result returned in the Storage\_Address parameter is used as the address of the allocated storage, which is a contiguous block of memory of Size\_In\_Storage\_Elements storage elements. Any exception propagated by Allocate is propagated by the construct that contained the call.

The number of calls to Allocate that will be used to implement an allocator for any particular type is unspecified. The number of calls to Deallocate that will be used to implement an instance of Unchecked\_Deallocation (see 16.11.2) for any particular object is the same as the number of Allocate calls for that object.

The Deallocate procedure of a user-defined storage pool object  $P$  may be called by the implementation to deallocate storage for a type  $T$  whose pool is  $P$  only at the places when an Allocate call is allowed for  $P$ , during the execution of an instance of Unchecked\_Deallocation for  $T$ , or as part of the finalization of the collection of  $T$ . For such a call of Deallocate,  $P$  (equivalent to  $T'_{\text{Storage\_Pool}}$ ) is passed as the Pool parameter. The value of the Storage\_Address parameter for a call to Deallocate is the value returned in the Storage\_Address parameter of the corresponding successful call to Allocate. The values of the Size\_In\_Storage\_Elements and Alignment parameters are the same values passed to the corresponding Allocate call. Any exception propagated by Deallocate is propagated by the construct that contained the call.

#### Documentation Requirements

An implementation shall document the set of values that a user-defined Allocate procedure has to accept for the Alignment parameter. An implementation shall document how the standard storage pool is chosen, and how storage is allocated by standard storage pools.

#### Implementation Advice

An implementation should document any cases in which it dynamically allocates heap storage for a purpose other than the evaluation of an allocator.

A default (implementation-provided) storage pool for an access-to-constant type should not have overhead to support deallocation of individual objects.

The storage pool used for an allocator of an anonymous access type should be determined as follows:

- If the allocator is defining a coextension (see 6.10.2) of an object being created by an outer allocator, then the storage pool used for the outer allocator should also be used for the coextension;
- For other access discriminants and access parameters, the storage pool should be created at the point of the allocator, and be reclaimed when the allocated object becomes inaccessible;
- If the allocator defines the result of a function with an access result, the storage pool is determined as though the allocator were in place of the call of the function. If the call is the operand of a type conversion, the storage pool is that of the target access type of the conversion. If the call is itself defining the result of a function with an access result, this rule is applied recursively;
- Otherwise, a default storage pool should be created at the point where the anonymous access type is elaborated; such a storage pool may have no mechanism for the deallocation of individual objects.

NOTE 1 A user-defined storage pool type can be obtained by extending the Root\_Storage\_Pool type, and overriding the primitive subprograms Allocate, Deallocate, and Storage\_Size. A user-defined storage pool can then be obtained by declaring an object of the type extension. The user can override Initialize and Finalize if there is any

desire for nontrivial initialization and finalization for a user-defined pool type. For example, Finalize can reclaim blocks of storage that are allocated separately from the pool object itself.

NOTE 2 The writer of the user-defined allocation and deallocation procedures, and users of `allocators` for the associated access type, are responsible for dealing with any interactions with tasking. In particular:

- If the `allocators` are used in different tasks, they require mutual exclusion.
- If they are used inside protected objects, they cannot block.
- If they are used by interrupt handlers (see C.3), the mutual exclusion mechanism has to work properly in that context.

NOTE 3 The primitives `Allocate`, `Deallocate`, and `Storage_Size` are declared as abstract (see 6.9.3), and therefore they have to be overridden when a new (nonabstract) storage pool type is declared.

EXAMPLE To associate an access type with a storage pool object, the user first declares a pool object of some type derived from `Root_Storage_Pool`. Then, the user defines its `Storage_Pool` attribute, as follows:

```
Pool_Object : Some_Storage_Pool_Type;

type T is access Designated;
for T'Storage_Pool use Pool_Object;
```

Another access type can be added to an existing storage pool, via:

```
for T2'Storage_Pool use T'Storage_Pool;
```

The semantics of this is implementation defined for a standard storage pool.

As usual, a derivative of `Root_Storage_Pool` can define additional operations. For example, consider the `Mark_Release_Pool_Type` defined in 16.11.6, that has two additional operations, `Mark` and `Release`, the following is a possible use:

```
type Mark_Release_Pool_Type
  (Pool_Size : Storage_Elements.Storage_Count)
  is new Subpools.Root_Storage_Pool_With_Subpools with private;
  -- As defined in package MR_Pool, see 16.11.6

...

Our_Pool : Mark_Release_Pool_Type (Pool_Size => 2000);
My_Mark : Subpool_Handle; -- As declared in 16.11.6

type Acc is access ...;
for Acc'Storage_Pool use Our_Pool;
...

My_Mark := Mark(Our_Pool);
... -- Allocate objects using "new (My_Mark) Designated(...)".
Release(My_Mark); -- Finalize objects and reclaim storage.
```

### 16.11.1 Storage Allocation Attributes

The `Max_Size_In_Storage_Elements` and `Max_Alignment_For_Allocation` attributes may be useful in writing user-defined pool types.

#### Static Semantics

For every subtype `S`, the following attributes are defined:

`S'Max_Size_In_Storage_Elements`

Denotes the maximum value for `Size_In_Storage_Elements` that can be requested by the implementation via `Allocate` for an access type whose designated subtype is `S`. The value of this attribute is of type *universal\_integer*.

`S'Max_Alignment_For_Allocation`

Denotes the maximum value for `Alignment` that can be requested by the implementation via `Allocate` for an access type whose designated subtype is `S`. The value of this attribute is of type *universal\_integer*.

For a type with access discriminants, if the implementation allocates space for a coextension in the same pool as that of the object having the access discriminant, then these attributes account for any calls on `Allocate` that can be performed to provide space for such coextensions.

### 16.11.2 Unchecked Storage Deallocation

Unchecked storage deallocation of an object designated by a value of an access type is achieved by a call to an instance of the generic procedure `Unchecked_Deallocation`.

#### Static Semantics

The following language-defined generic library procedure exists:

```
generic
  type Object(<>) is limited private;
  type Name is access Object;
  procedure Ada.Unchecked_Deallocation(X : in out Name)
    with Preelaborate, Nonblocking,
    Global => in out Name'Storage_Pool,
    Convention => Intrinsic;
```

#### Legality Rules

A call on an instance of `Unchecked_Deallocation` is illegal if the actual access type of the instance is a type for which the `Storage_Size` has been specified by a static expression with value zero or is defined by the language to be zero. In addition to the places where Legality Rules normally apply (see 15.3), this rule applies also in the private part of an instance of a generic unit.

#### Dynamic Semantics

Given an instance of `Unchecked_Deallocation` declared as follows:

```
procedure Free is
  new Ada.Unchecked_Deallocation(
    object_subtype_name, access_to_variable_subtype_name);
```

Procedure `Free` has the following effect:

1. After executing `Free(X)`, the value of `X` is **null**.
2. `Free(X)`, when `X` is already equal to **null**, has no effect.
3. `Free(X)`, when `X` is not equal to **null** first performs finalization of the object designated by `X` (and any coextensions of the object — see 6.10.2), as described in 10.6.1. It then deallocates the storage occupied by the object designated by `X` (and any coextensions). If the storage pool is a user-defined object, then the storage is deallocated by calling `Deallocate` as described in 16.11. There is one exception: if the object being freed contains tasks, it is unspecified whether the object is deallocated.

After the finalization step of `Free(X)`, the object designated by `X`, and any subcomponents (and coextensions) thereof, no longer exist; their storage can be reused for other purposes.

#### Bounded (Run-Time) Errors

It is a bounded error to free a discriminated, unterminated task object. The possible consequences are:

- No exception is raised.
- `Program_Error` or `Tasking_Error` is raised at the point of the deallocation.
- `Program_Error` or `Tasking_Error` is raised in the task the next time it references any of the discriminants.

In the first two cases, the storage for the discriminants (and for any enclosing object if it is designated by an access discriminant of the task) is not reclaimed prior to task termination.

An access value that designates a nonexistent object is called a *dangling reference*.

If a dangling reference is dereferenced (implicitly or explicitly), execution is erroneous (see below). If there is no explicit or implicit dereference, then it is a bounded error to evaluate an expression whose result is a dangling reference. If the error is detected, either `Constraint_Error` or `Program_Error` is raised. Otherwise, execution proceeds normally, but with the possibility that the access value designates some other existing object.

#### *Erroneous Execution*

Evaluating a name that denotes a nonexistent object, or a protected subprogram or subprogram renaming whose associated object (if any) is nonexistent, is erroneous. The execution of a call to an instance of `Unchecked_Deallocation` is erroneous if the object was created other than by an `allocator` for an access type whose pool is `Name'Storage_Pool`.

#### *Implementation Advice*

For a standard storage pool, `Free` should actually reclaim the storage.

A call on an instance of `Unchecked_Deallocation` with a nonnull access value should raise `Program_Error` if the actual access type of the instance is a type for which the `Storage_Size` has been specified to be zero or is defined by the language to be zero.

NOTE 1 The rules here that refer to `Free` apply to any instance of `Unchecked_Deallocation`.

NOTE 2 `Unchecked_Deallocation` cannot be instantiated for an access-to-constant type. This is implied by the rules of 15.5.4.

### 16.11.3 Default Storage Pools

`Pragma` and aspect `Default_Storage_Pool` specify the storage pool that will be used in the absence of an explicit specification of a storage pool or storage size for an access type.

#### *Syntax*

The form of a `pragma Default_Storage_Pool` is as follows:

```
pragma Default_Storage_Pool (storage_pool_indicator);
```

```
storage_pool_indicator ::= storage_pool_name | null | Standard
```

A `pragma Default_Storage_Pool` is allowed immediately within the visible part of a `package_specification`, immediately within a `declarative_part`, or as a configuration `pragma`.

#### *Name Resolution Rules*

The `storage_pool_name` is expected to be of type `Root_Storage_Pool'Class`.

#### *Legality Rules*

The `storage_pool_name` shall denote a variable.

The Standard `storage_pool_indicator` is an identifier specific to a `pragma` (see 5.8) and does not denote any declaration. If the `storage_pool_indicator` is `Standard`, then there shall not be a declaration with defining identifier `Standard` that is immediately visible at the point of the `pragma`, other than package `Standard` itself.

If the `pragma` is used as a configuration `pragma`, the `storage_pool_indicator` shall be either `null` or `Standard`, and it defines the *default pool* to be the given `storage_pool_indicator` within all applicable compilation units (see 13.1.5), except within the immediate scope of another `pragma Default_Storage_Pool`. Otherwise, the `pragma` occurs immediately within a sequence of declarations, and it defines the default pool within the immediate scope of the `pragma` to be the given `storage_pool_indicator`, except within the immediate scope of a later `pragma Default_Storage_Pool`. Thus, an inner `pragma` overrides an outer one.



A `pragma Default_Storage_Pool` shall not be used as a configuration pragma that applies to a compilation unit that is within the immediate scope of another `pragma Default_Storage_Pool`.

*Static Semantics*

The language-defined aspect `Default_Storage_Pool` may be specified for a generic instance; it defines the default pool for access types within an instance. .

The `Default_Storage_Pool` aspect may be specified as `Standard`, which is an identifier specific to an aspect (see 16.1.1) and defines the default pool to be `Standard`. In this case, there shall not be a declaration with `defining_identifier Standard` that is immediately visible at the point of the aspect specification, other than package `Standard` itself.

Otherwise, the expected type for the `Default_Storage_Pool` aspect is `Root_Storage_Pool'Class` and the `aspect_definition` shall be a name that denotes a variable. This aspect overrides any `Default_Storage_Pool` pragma that applies to the generic unit; if the aspect is not specified, the default pool of the instance is that defined for the generic unit.

The effect of specifying the aspect `Default_Storage_Pool` on an instance of a language-defined generic unit is implementation-defined.

For nonderived access types declared in places where the default pool is defined by the pragma or aspect, their `Storage_Pool` or `Storage_Size` attribute is determined as follows, unless `Storage_Pool` or `Storage_Size` is specified for the type:

- If the default pool is **null**, the `Storage_Size` attribute is defined by the language to be zero. Therefore, an `allocator` for such a type is illegal.
- If the default pool is neither **null** nor `Standard`, the `Storage_Pool` attribute is that pool.

Otherwise (including when the default pool is specified as `Standard`), the standard storage pool is used for the type as described in 16.11.

*Implementation Permissions*

An object created by an `allocator` that is passed as the actual parameter to an access parameter may be allocated on the stack, and automatically reclaimed, regardless of the default pool.

NOTE `Default_Storage_Pool` can be used with restrictions `No_Coextensions` and `No_Access_Parameter_Allocators` (see H.4) to ensure that all `allocators` use the default pool.

## 16.11.4 Storage Subpools

This subclause defines a package to support the partitioning of a storage pool into subpools. A subpool may be specified as the default to be used for allocation from the associated storage pool, or a particular subpool may be specified as part of an `allocator` (see 7.8).

*Static Semantics*

The following language-defined library package exists:

```
package System.Storage_Pools.Subpools
  with Preelaborate, Global => in out synchronized is

  type Root_Storage_Pool_With_Subpools is
    abstract new Root_Storage_Pool with private
    with Preelaborable_Initialization;

  type Root_Subpool is abstract tagged limited private
    with Preelaborable_Initialization;

  type Subpool_Handle is access all Root_Subpool'Class;
  for Subpool_Handle'Storage_Size use 0;

  function Create_Subpool (Pool : in out Root_Storage_Pool_With_Subpools)
    return not null Subpool_Handle is abstract;

  -- The following operations are intended for pool implementers:
```

```

function Pool_of_Subpool (Subpool : not null Subpool_Handle)
    return access Root_Storage_Pool_With_Subpools'Class;

procedure Set_Pool_of_Subpool (
    Subpool : in not null Subpool_Handle;
    To : in out Root_Storage_Pool_With_Subpools'Class)
    with Global => overriding in out Subpool;

procedure Allocate_From_Subpool (
    Pool : in out Root_Storage_Pool_With_Subpools;
    Storage_Address : out Address;
    Size_In_Storage_Elements : in Storage_Elements.Storage_Count;
    Alignment : in Storage_Elements.Storage_Count;
    Subpool : in not null Subpool_Handle) is abstract
    with Pre'Class => Pool_of_Subpool(Subpool) = Pool'Access,
        Global => overriding in out Subpool;

procedure Deallocate_Subpool (
    Pool : in out Root_Storage_Pool_With_Subpools;
    Subpool : in out Subpool_Handle) is abstract
    with Pre'Class => Pool_of_Subpool(Subpool) = Pool'Access;

function Default_Subpool_for_Pool (
    Pool : in out Root_Storage_Pool_With_Subpools)
    return not null Subpool_Handle;

overriding
procedure Allocate (
    Pool : in out Root_Storage_Pool_With_Subpools;
    Storage_Address : out Address;
    Size_In_Storage_Elements : in Storage_Elements.Storage_Count;
    Alignment : in Storage_Elements.Storage_Count);

overriding
procedure Deallocate (
    Pool : in out Root_Storage_Pool_With_Subpools;
    Storage_Address : in Address;
    Size_In_Storage_Elements : in Storage_Elements.Storage_Count;
    Alignment : in Storage_Elements.Storage_Count) is null;

overriding
function Storage_Size (Pool : Root_Storage_Pool_With_Subpools)
    return Storage_Elements.Storage_Count
    is (Storage_Elements.Storage_Count'Last);

private
    ... -- not specified by the language
end System.Storage_Pools.Subpools;

```

A *subpool* is a separately reclaimable portion of a storage pool, identified by an object of type `Subpool_Handle` (a *subpool handle*). A subpool handle also identifies the enclosing storage pool, a *storage pool that supports subpools*, which is a storage pool whose type is descended from `Root_Storage_Pool_With_Subpools`. A subpool is created by calling `Create_Subpool` or a similar constructor; the constructor returns the subpool handle.

A *subpool object* is an object of a type descended from `Root_Subpool`. Typically, subpool objects are managed by the containing storage pool; only the handles have to be exposed to clients of the storage pool. Subpool objects are designated by subpool handles, and are the run-time representation of a subpool.

Each subpool *belongs* to a single storage pool (which will always be a pool that supports subpools). An access to the pool that a subpool belongs to can be obtained by calling `Pool_of_Subpool` with the subpool handle. `Set_Pool_of_Subpool` causes the subpool of the subpool handle to belong to the given pool; this is intended to be called from subpool constructors like `Create_Subpool`. `Set_Pool_of_Subpool` propagates `Program_Error` if the subpool already belongs to a pool. If `Set_Pool_of_Subpool` has not yet been called for a subpool, `Pool_of_Subpool` returns **null**.

When an `allocator` for a type whose storage pool supports subpools is evaluated, a call is made on `Allocate_From_Subpool` passing in a `Subpool_Handle`, in addition to the parameters as defined for calls on `Allocate` (see 16.11). The subpool designated by the `subpool_handle_name` is used, if specified in an `allocator`. Otherwise, `Default_Subpool_for_Pool` of the `Pool` is used to provide a subpool handle. All requirements on the `Allocate` procedure also apply to `Allocate_from_Subpool`.

#### Legality Rules

If a storage pool that supports subpools is specified as the `Storage_Pool` for an access type, the access type is called a *subpool access type*. A subpool access type shall be a pool-specific access type.

The accessibility level of a subpool access type shall not be statically deeper than that of the storage pool object. If the specified storage pool object is a storage pool that supports subpools, then the `name` that denotes the object shall not denote part of a formal parameter, nor shall it denote part of a dereference of a value of a non-library-level general access type. In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

#### Dynamic Semantics

When an access type with a specified storage pool is frozen (see 16.14), if the tag of the storage pool object identifies a storage pool that supports subpools, the following checks are made:

- the `name` used to specify the storage pool object does not denote part of a formal parameter nor part of a dereference of a value of a non-library-level general access type; and
- the accessibility level of the access type is not deeper than that of the storage pool object.

`Program_Error` is raised if either of these checks fail.

A call to `Subpools.Allocate(P, Addr, Size, Align)` does the following:

```
Allocate_From_Subpool
  (Root_Storage_Pool_With_Subpools'Class(P),
   Addr, Size, Align,
   Subpool => Default_Subpool_for_Pool
              (Root_Storage_Pool_With_Subpools'Class(P)));
```

An `allocator` that allocates in a subpool raises `Program_Error` if the allocated object has task parts.

Unless overridden, `Default_Subpool_for_Pool` propagates `Program_Error`.

#### Erroneous Execution

If `Allocate_From_Subpool` does not meet one or more of the requirements on the `Allocate` procedure as given in the Erroneous Execution rules of 16.11, then the program execution is erroneous.

#### Implementation Permissions

When an `allocator` for a type whose storage pool is of type `Root_Storage_Pool'Class` is evaluated, but supports subpools, the implementation may call `Allocate` rather than `Allocate_From_Subpool`. This will have the same effect, so long as `Allocate` has not been overridden.

NOTE 1 A user-defined storage pool type that supports subpools can be implemented by extending the `Root_Storage_Pool_With_Subpools` type, and overriding the primitive subprograms `Create_Subpool`, `Allocate_From_Subpool`, and `Deallocate_Subpool`. `Create_Subpool` is expected to call `Set_Pool_Of_Subpool` before returning the subpool handle. To make use of such a pool, a user can declare an object of the type extension, can use it to define the `Storage_Pool` attribute of one or more access types, and then can invoke `Create_Subpool` to obtain subpool handles associated with the pool.

NOTE 2 A user-defined storage pool type that supports subpools can define additional subpool constructors similar to `Create_Subpool` (these typically will have additional parameters).

NOTE 3 The pool implementor can override `Default_Subpool_For_Pool` if they want the pool to support a default subpool for the pool. The implementor can override `Deallocate` if individual object reclamation is to be supported, and can override `Storage_Size` if there is some limit on the total size of the storage pool. The implementor can override `Initialize` and `Finalize` if there is any desire for nontrivial initialization and finalization for the pool as a

whole. For example, Finalize can reclaim blocks of storage that are allocated over and above the space occupied by the pool object itself. The pool implementor can extend the Root\_Subpool type as necessary to carry additional information with each subpool provided by Create\_Subpool.

### 16.11.5 Subpool Reclamation

A subpool may be explicitly deallocated using Unchecked\_Deallocate\_Subpool.

#### Static Semantics

The following language-defined library procedure exists:

```
with System.Storage_Pools.Subpools;
procedure Ada.Unchecked_Deallocate_Subpool
  (Subpool : in out System.Storage_Pools.Subpools.Subpool_Handle)
with Global => in out all;
```

If Subpool is **null**, a call on Unchecked\_Deallocate\_Subpool has no effect. Otherwise, the subpool is finalized, and Subpool is set to **null**.

Finalization of a subpool has the following effects in the given order:

1. Any of the objects allocated from the subpool that still exist are finalized in an arbitrary order;
2. All of the objects allocated from the subpool cease to exist;
3. The following dispatching call is then made:

```
Deallocate_Subpool(Pool_of_Subpool(Subpool).all, Subpool);
```

4. The subpool ceases to belong to any pool.

Finalization of a Root\_Storage\_Pool\_With\_Subpools object finalizes all subpools that belong to that pool that have not yet been finalized.

### 16.11.6 Storage Subpool Example

EXAMPLE The following example is a simple but complete implementation of the classic Mark/Release pool using subpools:

```
with System.Storage_Pools.Subpools;
with System.Storage_Elements;
with Ada.Unchecked_Deallocate_Subpool;
package MR_Pool is
  use System.Storage_Pools;
  -- For uses of Subpools.
  use System.Storage_Elements;
  -- For uses of Storage_Count and Storage_Array.
  -- Mark and Release work in a stack fashion, and allocations are not
  allowed
  -- from a subpool other than the one at the top of the stack. This is also
  -- the default pool.
  subtype Subpool_Handle is Subpools.Subpool_Handle;
  type Mark_Release_Pool_Type (Pool_Size : Storage_Count) is new
    Subpools.Root_Storage_Pool_With_Subpools with private;
  function Mark (Pool : in out Mark_Release_Pool_Type)
    return not null Subpool_Handle;
  procedure Release (Subpool : in out Subpool_Handle) renames
    Ada.Unchecked_Deallocate_Subpool;
```

**private**

```

type MR_Subpool is new Subpools.Root_Subpool with record
    Start : Storage_Count;
end record;
subtype Subpool_Indexes is Positive range 1 .. 10;
type Subpool_Array is array (Subpool_Indexes) of aliased MR_Subpool;

type Mark_Release_Pool_Type (Pool_Size : Storage_Count) is new
    Subpools.Root_Storage_Pool_With_Subpools with record
        Storage      : Storage_Array (0 .. Pool_Size);
        Next_Allocation : Storage_Count := 0;
        Markers       : Subpool_Array;
        Current_Pool   : Subpool_Indexes := 1;
end record;

overriding
function Create_Subpool (Pool : in out Mark_Release_Pool_Type)
    return not null Subpool_Handle;

function Mark (Pool : in out Mark_Release_Pool_Type)
    return not null Subpool_Handle renames Create_Subpool;

overriding
procedure Allocate_From_Subpool (
    Pool : in out Mark_Release_Pool_Type;
    Storage_Address : out System.Address;
    Size_In_Storage_Elements : in Storage_Count;
    Alignment : in Storage_Count;
    Subpool : not null Subpool_Handle);

overriding
procedure Deallocate_Subpool (
    Pool : in out Mark_Release_Pool_Type;
    Subpool : in out Subpool_Handle);

overriding
function Default_Subpool_for_Pool (Pool : in out Mark_Release_Pool_Type)
    return not null Subpool_Handle;

overriding
procedure Initialize (Pool : in out Mark_Release_Pool_Type);
    -- We don't need Finalize.

```

**end** MR\_Pool;

**package body** MR\_Pool **is**

```

    use type Subpool_Handle;

    procedure Initialize (Pool : in out Mark_Release_Pool_Type) is
        -- Initialize the first default subpool.
    begin
        Pool.Markers(1).Start := 1;
        Subpools.Set_Pool_of_Subpool
            (Pool.Markers(1)'Unchecked_Access, Pool);
    end Initialize;

    function Create_Subpool (Pool : in out Mark_Release_Pool_Type)
        return not null Subpool_Handle is
        -- Mark the current allocation location.
    begin
        if Pool.Current_Pool = Subpool_Indexes'Last then
            raise Storage_Error; -- No more subpools.
        end if;
        Pool.Current_Pool := Pool.Current_Pool + 1; -- Move to the next subpool

```

```

    return Result : constant not null Subpool_Handle :=
        Pool.Markers(Pool.Current_Pool)'Unchecked_Access
do
    Pool.Markers(Pool.Current_Pool).Start := Pool.Next_Allocation;
    Subpools.Set_Pool_of_Subpool (Result, Pool);
end return;
end Create_Subpool;

procedure Deallocate_Subpool (
    Pool : in out Mark_Release_Pool_Type;
    Subpool : in out Subpool_Handle) is
begin
    if Subpool /= Pool.Markers(Pool.Current_Pool)'Unchecked_Access then
        raise Program_Error; -- Only the last marked subpool can be released.
    end if;
    if Pool.Current_Pool /= 1 then
        Pool.Next_Allocation := Pool.Markers(Pool.Current_Pool).Start;
        Pool.Current_Pool := Pool.Current_Pool - 1; -- Move to the previous
subpool
    else -- Reinitialize the default subpool:
        Pool.Next_Allocation := 1;
        Subpools.Set_Pool_of_Subpool
            (Pool.Markers(1)'Unchecked_Access, Pool);
    end if;
end Deallocate_Subpool;

function Default_Subpool_for_Pool (Pool : in out Mark_Release_Pool_Type)
    return not null Subpool_Handle is
begin
    return Pool.Markers(Pool.Current_Pool)'Unchecked_Access;
end Default_Subpool_for_Pool;

procedure Allocate_From_Subpool (
    Pool : in out Mark_Release_Pool_Type;
    Storage_Address : out System.Address;
    Size_In_Storage_Elements : in Storage_Count;
    Alignment : in Storage_Count;
    Subpool : not null Subpool_Handle) is
begin
    if Subpool /= Pool.Markers(Pool.Current_Pool)'Unchecked_Access then
        raise Program_Error; -- Only the last marked subpool can be used for
allocations.
    end if;

    -- Check for the maximum supported alignment, which is the alignment of
the storage area:
    if Alignment > Pool.Storage'Alignment then
        raise Program_Error;
    end if;

    -- Correct the alignment if necessary:
    Pool.Next_Allocation := Pool.Next_Allocation +
        ((-Pool.Next_Allocation) mod Alignment);
    if Pool.Next_Allocation + Size_In_Storage_Elements >
        Pool.Pool_Size then
        raise Storage_Error; -- Out of space.
    end if;
    Storage_Address := Pool.Storage (Pool.Next_Allocation)'Address;
    Pool.Next_Allocation :=
        Pool.Next_Allocation + Size_In_Storage_Elements;
end Allocate_From_Subpool;
end MR_Pool;

```



## 16.12 Pragma Restrictions and Pragma Profile

A `pragma Restrictions` expresses the user's intent to abide by certain restrictions. A `pragma Profile` expresses the user's intent to abide by a set of Restrictions or other specified run-time policies. These may facilitate the construction of simpler run-time environments.

### *Syntax*

The form of a `pragma Restrictions` is as follows:

```
pragma Restrictions(restriction{, restriction});
restriction ::= restriction_identifier
| restriction_parameter_identifier => restriction_parameter_argument
restriction_parameter_argument ::= name | expression
```

### *Name Resolution Rules*

Unless otherwise specified for a particular restriction, the `expression` is expected to be of any integer type.

### *Legality Rules*

Unless otherwise specified for a particular restriction, the `expression` shall be static, and its value shall be nonnegative.

### *Post-Compilation Rules*

A `pragma Restrictions` is a configuration `pragma`. If a `pragma Restrictions` applies to any compilation unit included in the partition, this may impose either (or both) of two kinds of requirements, as specified for the particular restriction:

- A restriction may impose requirements on some or all of the units comprising the partition. Unless otherwise specified for a particular restriction, such a requirement applies to all of the units comprising the partition and is enforced via a post-compilation check.
- A restriction may impose requirements on the run-time behavior of the program, as indicated by the specification of run-time behavior associated with a violation of the requirement.

For the purpose of checking whether a partition contains constructs that violate any restriction (unless specified otherwise for a particular restriction):

- Generic instances are logically expanded at the point of instantiation;
- If an object of a type is declared or allocated and not explicitly initialized, then all expressions appearing in the definition for the type and any of its ancestors are presumed to be used;
- A default `expression` for a formal parameter or a generic formal object is considered to be used if and only if the corresponding actual parameter is not provided in a given call or instantiation.

### *Implementation Permissions*

An implementation may provide implementation-defined restrictions; the identifier for an implementation-defined restriction shall differ from those of the language-defined restrictions.

An implementation may place limitations on the values of the `expression` that are supported, and limitations on the supported combinations of restrictions. The consequences of violating such limitations are implementation defined.

An implementation is permitted to omit restriction checks for code that is recognized at compile time to be unreachable and for which no code is generated.

Whenever enforcement of a restriction is not required prior to execution, an implementation may nevertheless enforce the restriction prior to execution of a partition to which the restriction applies, provided that every execution of the partition would violate the restriction.

*Syntax*

The form of a pragma Profile is as follows:

**pragma** Profile (*profile\_identifier* {, *profile\_pragma\_argument\_association*});

*Legality Rules*

The *profile\_identifier* shall be the name of a usage profile. The semantics of any *profile\_pragma\_argument\_associations* are defined by the usage profile specified by the *profile\_identifier*.

*Static Semantics*

A profile is equivalent to the set of configuration pragmas that is defined for each usage profile.

*Post-Compilation Rules*

A pragma Profile is a configuration pragma. There may be more than one pragma Profile for a partition.

*Implementation Permissions*

An implementation may provide implementation-defined usage profiles; the identifier for an implementation-defined usage profile shall differ from those of the language-defined usage profiles.

NOTE 1 Restrictions intended to facilitate the construction of efficient tasking run-time systems are defined in D.7. Restrictions intended for use when constructing high integrity systems are defined in H.4.

NOTE 2 An implementation has to enforce the restrictions in cases where enforcement is required, even if it chooses not to take advantage of the restrictions in terms of efficiency.

## 16.12.1 Language-Defined Restrictions and Profiles

*Static Semantics*

The following *restriction\_identifiers* are language defined (additional restrictions are defined in the Specialized Needs Annexes):

### No\_Implementation\_Aspect\_Specifications

There are no implementation-defined aspects specified by an *aspect\_specification*. This restriction applies only to the current compilation or environment, not the entire partition.

### No\_Implementation\_Attributes

There are no implementation-defined attributes. This restriction applies only to the current compilation or environment, not the entire partition.

### No\_Implementation\_Identifiers

There are no usage names that denote declarations with implementation-defined identifiers that occur within language-defined packages or instances of language-defined generic packages. Such identifiers can arise as follows:

- The following language-defined packages and generic packages allow implementation-defined identifiers:
  - package System (see 16.7);
  - package Standard (see A.1);
  - package Ada.Command\_Line (see A.15);
  - package Interfaces.C (see B.3);
  - package Interfaces.C.Strings (see B.3.1);
  - package Interfaces.C.Pointers (see B.3.2);
  - package Interfaces.COBOLE (see B.4);

- package Interfaces.Fortran (see B.5);
- The following language-defined packages contain only implementation-defined identifiers:
  - package System.Machine\_Code (see 16.8);
  - package Ada.Directories.Information (see A.16);
  - nested Implementation packages of the Queue containers (see A.18.28-31);
  - package Interfaces (see B.2);
  - package Ada.Interrupts.Names (see C.3.2).

For package Standard, Standard.Long\_Integer and Standard.Long\_Float are considered language-defined identifiers, but identifiers such as Standard.Short\_Short\_Integer are considered implementation-defined.

This restriction applies only to the current compilation or environment, not the entire partition.

#### No\_Implementation\_Pragmas

There are no implementation-defined pragmas or pragma arguments. This restriction applies only to the current compilation or environment, not the entire partition.

#### No\_Implementation\_Units

There is no mention in the `context_clause` of any implementation-defined descendants of packages Ada, Interfaces, or System. This restriction applies only to the current compilation or environment, not the entire partition.

#### No\_Obsolescent\_Features

There is no use of language features defined in Annex J. It is implementation defined whether uses of the renamings of I.1 and of the pragmas of I.15 are detected by this restriction. This restriction applies only to the current compilation or environment, not the entire partition.

The following *restriction\_parameter\_identifiers* are language defined:

#### No\_Dependence

Specifies a library unit on which there are no semantic dependences.

#### No\_Specification\_of\_Aspect

Identifies an aspect for which no `aspect_specification`, `attribute_definition_clause`, or pragma is given.

#### No\_Use\_Of\_Attribute

Identifies an attribute for which no `attribute_reference` or `attribute_definition_clause` is given.

#### No\_Use\_Of\_Pragma

Identifies a pragma which is not to be used.

#### No\_Unrecognized\_Aspects

There are no `aspect_specifications` having an unrecognized *aspect\_identifier*. This restriction applies only to the current compilation or environment, not the entire partition.

#### No\_Unrecognized\_Pragmas

There are no pragmas having an unrecognized *pragma identifier*. This restriction applies only to the current compilation or environment, not the entire partition.

*Legality Rules*

The `restriction_parameter_argument` of a `No_Dependence` restriction shall be a name; the name shall have the form of a full expanded name of a library unit, but can be a name that has no corresponding unit currently present in the environment.

The `restriction_parameter_argument` of a `No_Specification_of_Aspect` restriction shall be an identifier; this is an identifier specific to a pragma (see 5.8) and does not denote any declaration.

The `restriction_parameter_argument` of a `No_Use_Of_Attribute` restriction shall be an identifier or one of the reserved words `Access`, `Delta`, `Digits`, `Mod`, or `Range`; this is an identifier specific to a pragma.

The `restriction_parameter_argument` of a `No_Use_Of_Pragma` restriction shall be an identifier or the reserved word `Interface`; this is an identifier specific to a pragma.

*Post-Compilation Rules*

No compilation unit included in the partition shall depend semantically on the library unit identified by the name of a `No_Dependence` restriction.

*Static Semantics*

The following `profile_identifier` is language defined:

`No_Implementation_Extensions`

For usage profile `No_Implementation_Extensions`, there shall be no `profile_pragma_argument_` associations.

The `No_Implementation_Extensions` usage profile is equivalent to the following restrictions:

```
No_Implementation_Aspect_Specifications,
No_Implementation_Attributes,
No_Implementation_Identifiers,
No_Implementation_Pragmas,
No_Implementation_Units.
```

## 16.13 Streams

A *stream* is a sequence of elements comprising values from possibly different types and allowing sequential access to these values. A *stream type* is a type in the class whose root type is `Streams.Root_Stream_Type`. A stream type may be implemented in various ways, such as an external sequential file, an internal buffer, or a network channel.

### 16.13.1 The Streams Subsystem

*Static Semantics*

The abstract type `Root_Stream_Type` is the root type of the class of stream types. The types in this class represent different kinds of streams. A new stream type is defined by extending the root type (or some other stream type), overriding the `Read` and `Write` operations, and optionally defining additional primitive subprograms, according to the requirements of the particular kind of stream. The predefined stream-oriented attributes like `T'Read` and `T'Write` make dispatching calls on the `Read` and `Write` procedures of the `Root_Stream_Type`. (User-defined `T'Read` and `T'Write` attributes can also make such calls, or can call the `Read` and `Write` attributes of other types.)

The library package `Ada.Streams` has the following declaration:

```
package Ada.Streams
with Pure, Nonblocking => False is
    type Root_Stream_Type is abstract tagged limited private
        with Preelaborable_Initialization;
```

```

type Stream_Element is mod implementation-defined;
type Stream_Element_Offset is range implementation-defined;
subtype Stream_Element_Count is
    Stream_Element_Offset range 0..Stream_Element_Offset'Last;
type Stream_Element_Array is
    array (Stream_Element_Offset range <>) of aliased Stream_Element;

procedure Read(
    Stream : in out Root_Stream_Type;
    Item   : out Stream_Element_Array;
    Last   : out Stream_Element_Offset) is abstract;

procedure Write(
    Stream : in out Root_Stream_Type;
    Item   : in Stream_Element_Array) is abstract;

private
    ... -- not specified by the language
end Ada.Streams;

```

The Read operation transfers stream elements from the specified stream to fill the array Item. Elements are transferred until Item'Length elements have been transferred or until the end of the stream is reached. If any elements are transferred, the index of the last stream element transferred is returned in Last. Otherwise, Item'First - 1 is returned in Last. Last is less than Item'Last only if the end of the stream is reached.

The Write operation appends Item to the specified stream.

Three additional packages provide stream implementations that do not make use of any file operations. These packages provide the same operations, with Streams.Storage providing an abstract interface, and two child packages providing implementations of that interface. The difference is that for Streams.Storage.Bounded, the maximum storage is bounded.

The library package Ada.Streams.Storage has the following declaration:

```

package Ada.Streams.Storage
with Pure, Nonblocking is

    type Storage_Stream_Type is abstract new Root_Stream_Type with private;

    function Element_Count (Stream : Storage_Stream_Type)
        return Stream_Element_Count is abstract;

    procedure Clear (Stream : in out Storage_Stream_Type) is abstract;

private
    ... -- not specified by the language
end Ada.Streams.Storage;

```

The library package Ada.Streams.Storage.Unbounded has the following declaration:

```

package Ada.Streams.Storage.Unbounded
with Prelaborated, Nonblocking, Global => in out synchronized is

    type Stream_Type is new Storage_Stream_Type with private
        with Default_Initial_Condition =>
            Element_Count (Stream_Type) = 0;

```

```

overriding
procedure Read (
  Stream : in out Stream_Type;
  Item    : out Stream_Element_Array;
  Last    : out Stream_Element_Offset)
with Post =>
  (declare
    Num_Read : constant Stream_Element_Count :=
      Stream_Element_Count'Min
      (Element_Count(Stream)'Old, Item'Length);
  begin
    Last = Num_Read + Item'First - 1 and
    Element_Count (Stream) =
      Element_Count (Stream)'Old - Num_Read);

overriding
procedure Write (
  Stream : in out Stream_Type;
  Item    : in Stream_Element_Array)
with Post =>
  Element_Count (Stream) =
  Element_Count (Stream)'Old + Item'Length;

overriding
function Element_Count (Stream : Stream_Type)
return Stream_Element_Count;

overriding
procedure Clear (Stream : in out Stream_Type)
with Post => Element_Count (Stream) = 0;

private
  ... -- not specified by the language
end Ada.Streams.Storage.Unbounded;

```

The library package Ada.Streams.Storage.Bounded has the following declaration:

```

package Ada.Streams.Storage.Bounded
with Pure, Nonblocking is

  type Stream_Type (Max_Elements : Stream_Element_Count)
  is new Storage_Stream_Type with private
  with Default_Initial_Condition =>
    Element_Count (Stream_Type) = 0;

overriding
procedure Read (
  Stream : in out Stream_Type;
  Item    : out Stream_Element_Array;
  Last    : out Stream_Element_Offset)
with Post =>
  (declare
    Num_Read : constant Stream_Element_Count :=
      Stream_Element_Count'Min
      (Element_Count(Stream)'Old, Item'Length);
  begin
    Last = Num_Read + Item'First - 1 and
    Element_Count (Stream) =
      Element_Count (Stream)'Old - Num_Read);

overriding
procedure Write (
  Stream : in out Stream_Type;
  Item    : in Stream_Element_Array)
with Pre =>
  Element_Count (Stream) + Item'Length <= Stream.Max_Elements
  or else (raise Constraint_Error),
  Post =>
  Element_Count (Stream) =
  Element_Count (Stream)'Old + Item'Length;

```



```

overriding
function Element_Count (Stream : Stream_Type)
  return Stream_Element_Count
  with Post => Element_Count'Result <= Stream.Max_Elements;

overriding
procedure Clear (Stream : in out Stream_Type)
  with Post => Element_Count (Stream) = 0;

private
  ... -- not specified by the language
end Ada.Streams.Storage.Bounded;

```

The Element\_Count functions return the number of stream elements that are available for reading from the given stream.

The Read and Write procedures behave as described for package Ada.Streams above. Stream elements are read in FIFO (first-in, first-out) order; stream elements are available for reading immediately after they are written.

The Clear procedures remove any available stream elements from the given stream.

#### Implementation Permissions

If Stream\_Element'Size is not a multiple of System.Storage\_Unit, then the components of Stream\_Element\_Array will not be aliased.

#### Implementation Advice

Streams.Storage.Bounded.Stream\_Type objects should be implemented without implicit pointers or dynamic allocation.

NOTE 1 See A.12.1 for an example of extending type Root\_Stream\_Type.

NOTE 2 If the end of stream has been reached, and Item\_First is Stream\_Element\_Offset'First, Read will raise Constraint\_Error.

### 16.13.2 Stream-Oriented Attributes

The type-related operational attributes Write, Read, Output, and Input convert values to a stream of elements and reconstruct values from a stream.

#### Static Semantics

For every subtype *S* of an elementary type *T*, the following representation attribute is defined:

S'Stream\_Size

Denotes the number of bits read from or written to a stream by the default implementations of S'Read and S'Write. Hence, the number of stream elements required per item of elementary type *T* is:

$$T'Stream\_Size / Ada.Streams.Stream\_Element'Size$$

The value of this attribute is of type *universal\_integer* and is a multiple of Stream\_Element'Size.

Stream\_Size may be specified for first subtypes via an attribute\_definition\_clause; the expression of such a clause shall be static, nonnegative, and a multiple of Stream\_Element'Size.

#### Implementation Advice

If not specified, the value of Stream\_Size for an elementary type should be the number of bits that corresponds to the minimum number of stream elements required by the first subtype of the type, rounded up to the nearest factor or multiple of the word size that is also a multiple of the stream element size.

The recommended level of support for the `Stream_Size` attribute is:

- A `Stream_Size` clause should be supported for a discrete or fixed point type  $T$  if the specified `Stream_Size` is a multiple of `Stream_Element'Size` and is no less than the size of the first subtype of  $T$ , and no greater than the size of the largest type of the same elementary class (signed integer, modular integer, enumeration, ordinary fixed point, or decimal fixed point).

#### Static Semantics

For every subtype  $S$  of a specific type  $T$ , the following attributes are defined.

**S'Write** S'Write denotes a procedure with the following specification:

```
procedure S'Write(  
    Stream : not null access Ada.Streams.Root_Stream_Type'Class;  
    Item : in  $T$ )
```

S'Write writes the value of *Item* to *Stream*.

**S'Read** S'Read denotes a procedure with the following specification:

```
procedure S'Read(  
    Stream : not null access Ada.Streams.Root_Stream_Type'Class;  
    Item : out  $T$ )
```

S'Read reads the value of *Item* from *Stream*.

The default implementations of the `Write` and `Read` attributes, where available, execute as follows:

For nonderived elementary types, `Read` reads (and `Write` writes) the number of stream elements implied by the `Stream_Size` for the type  $T$ ; the representation of those stream elements is implementation defined. For nonderived composite types, the `Write` or `Read` attribute for each component (excluding those, if any, that are not components of the nominal type of the object) is called in canonical order, which is last dimension varying fastest for an array (unless the convention of the array is Fortran, in which case it is first dimension varying fastest), and positional aggregate order for a record. Bounds are not included in the stream if  $T$  is an array type. If  $T$  is a discriminated type, discriminants are included only if they have defaults. If  $T$  is a tagged type, the tag is not included.

For type extensions, the `Write` or `Read` attribute for the parent type is called, followed by the `Write` or `Read` attribute of each component of the extension part, in canonical order. For a limited type extension, if the attribute of the parent type or any progenitor type of  $T$  is available anywhere within the immediate scope of  $T$ , and the attribute of the parent type or the type of any of the extension components is not available at the freezing point of  $T$ , then the attribute of  $T$  shall be directly specified. For untagged derived types, the `Write` (resp. `Read`) attribute invokes the corresponding attribute of the parent type, if the attribute is available for the parent type.

If  $T$  is a discriminated type and its discriminants have defaults, then `S'Read` first reads the discriminants from the stream without modifying *Item*. `S'Read` then creates an object of type  $T$  constrained by these discriminants. The value of this object is then converted to the subtype of *Item* and is assigned to *Item*. Finally, the `Read` attribute for each nondiscriminant component of *Item* is called in canonical order as described above. Normal default initialization and finalization take place for the created object.

`Constraint_Error` is raised by the predefined `Write` attribute if the value of the elementary item is outside the range of values representable using `Stream_Size` bits. For a signed integer type, an enumeration type, or a fixed point type, the range is unsigned only if the integer code for the lower bound of the first subtype is nonnegative, and a (symmetric) signed range that covers all values of the first subtype would require more than `Stream_Size` bits; otherwise, the range is signed.

For every subtype `S'Class` of a class-wide type `T'Class`:

**S'Class'Write**

S'Class'Write denotes a procedure with the following specification:

```

procedure S'Class'Write(
    Stream : not null access Ada.Streams.Root_Stream_Type'Class;
    Item   : in T'Class)

```

Dispatches to the subprogram denoted by the Write attribute of the specific type identified by the tag of Item.

#### S'Class'Read

S'Class'Read denotes a procedure with the following specification:

```

procedure S'Class'Read(
    Stream : not null access Ada.Streams.Root_Stream_Type'Class;
    Item   : out T'Class)

```

Dispatches to the subprogram denoted by the Read attribute of the specific type identified by the tag of Item.

#### Static Semantics

For every subtype *S* of a specific type *T*, the following attributes are defined.

**S'Output** S'Output denotes a procedure with the following specification:

```

procedure S'Output(
    Stream : not null access Ada.Streams.Root_Stream_Type'Class;
    Item   : in T)

```

S'Output writes the value of *Item* to *Stream*, including any bounds or discriminants.

**S'Input** S'Input denotes a function with the following specification:

```

function S'Input(
    Stream : not null access Ada.Streams.Root_Stream_Type'Class)
return T

```

S'Input reads and returns one value from *Stream*, using any bounds or discriminants written by a corresponding S'Output to determine how much to read.

For an untagged derived type, the default implementation of the Output (resp. Input) attribute invokes the corresponding attribute of the parent type, if the attribute is available for the parent type. For any other type, the default implementations of the Output and Input attributes, where available, execute as follows:

- If *T* is an array type, S'Output first writes the bounds, and S'Input first reads the bounds. If *T* has discriminants without defaults, S'Output first writes the discriminants (using the Write attribute of the discriminant type for each), and S'Input first reads the discriminants (using the Read attribute of the discriminant type for each).
- S'Output then calls S'Write to write the value of *Item* to the stream. S'Input then creates an object of type *T*, with the bounds or (when without defaults) the discriminants, if any, taken from the stream, passes it to S'Read, and returns the value of the object. If *T* has discriminants, then this object is unconstrained if and only the discriminants have defaults. Normal default initialization and finalization take place for this object (see 6.3.1, 10.6, and 10.6.1).

If *T* is an abstract type, then S'Input is an abstract function.

For every subtype S'Class of a class-wide type *T'Class*:

#### S'Class'Output

S'Class'Output denotes a procedure with the following specification:

```

procedure S'Class'Output(
    Stream : not null access Ada.Streams.Root_Stream_Type'Class;
    Item   : in T'Class)

```

First writes the external tag of *Item* to *Stream* (by calling String'Output(*Stream*, Tags.External\_Tag(*Item*'Tag)) — see 6.9) and then dispatches to the subprogram denoted by the Output attribute of the specific type identified by the tag. Tag\_Error is raised if the tag of Item identifies a type declared at an accessibility level deeper than that of S.

**S'Class'Input**

S'Class'Input denotes a function with the following specification:

```
function S'Class'Input (
    Stream : not null access Ada.Streams.Root_Stream_Type'Class)
    return T'Class
```

First reads the external tag from *Stream* and determines the corresponding internal tag (by calling Tags.Descendant\_Tag(String'Input(*Stream*), S'Tag) which can raise Tag\_Error — see 6.9) and then dispatches to the subprogram denoted by the Input attribute of the specific type identified by the internal tag; returns that result. If the specific type identified by the internal tag is abstract, Constraint\_Error is raised.

In the default implementation of Read and Input for a composite type, for each scalar component that is a discriminant or that has an implicit initial value, a check is made that the value returned by Read for the component belongs to its subtype. Constraint\_Error is raised if this check fails. For other scalar components, no check is made. For each component that is of an access type, if the implementation can detect that the value returned by Read for the component is not a value of its subtype, Constraint\_Error is raised. If the value is not a value of its subtype and this error is not detected, the component has an abnormal value, and erroneous execution can result (see 16.9.1). In the default implementation of Read for a composite type with defaulted discriminants, if the actual parameter of Read is constrained, a check is made that the discriminants read from the stream are equal to those of the actual parameter. Constraint\_Error is raised if this check fails.

It is unspecified at which point and in which order these checks are performed. In particular, if Constraint\_Error is raised due to the failure of one of these checks, it is unspecified how many stream elements have been read from the stream.

In the default implementation of Read and Input for a type, End\_Error is raised if the end of the stream is reached before the reading of a value of the type is completed.

The Nonblocking aspect is statically True and the Global aspect is **null** for the default implementations of stream-oriented attributes for elementary types. For the default implementations of stream-oriented attributes for composite types, the value of the Nonblocking aspect is that of the first subtype, and the Global aspect defaults to that of the first subtype. A default implementation of a stream-oriented attribute that has the Nonblocking aspect statically True is considered a nonblocking region. The aspect Dispatching (see H.7.1) is Read(Stream) for the default implementations of the stream-oriented attributes Read, Read'Class, Input, and Input'Class; the aspect Dispatching is Write(Stream) for the default implementations of the stream-oriented attributes Write, Write'Class, Output, and Output'Class.

The stream-oriented attributes may be specified for any type via an attribute\_definition\_clause. Alternatively, each of the specific stream-oriented attributes may be specified using an aspect\_specification on any type\_declaration, with the aspect name being the corresponding attribute name. Each of the class-wide stream-oriented attributes may be specified using an aspect\_specification for a tagged type *T* using the name of the stream-oriented attribute followed by 'Class; such class-wide aspects do not apply to other descendants of *T*. If not directly specified, a default implementation of a stream-oriented attribute is implicitly composed for a nonlimited type, and for certain limited types, as defined above.

The subprogram name given in such an attribute\_definition\_clause or aspect\_specification shall statically denote a subprogram that is not an abstract subprogram. Furthermore, if a specific stream-oriented attribute is specified for an interface type, the subprogram name given in the attribute\_definition\_clause or aspect\_specification shall statically denote a null procedure.

A stream-oriented attribute for a subtype of a specific type *T* is *available* at places where one of the following conditions is true:

- *T* is nonlimited.

- The `attribute_designator` is Read (resp. Write) and  $T$  is a limited record extension, and the attribute Read (resp. Write) is available for the parent type of  $T$  and for the types of all of the extension components.
- $T$  is a limited untagged derived type, and the attribute is available for the parent type.
- The `attribute_designator` is Input (resp. Output), and  $T$  is a limited type, and the attribute Read (resp. Write) is available for  $T$ .
- The attribute has been specified via an `attribute_definition_clause` or `aspect_specification`, and the `attribute_definition_clause` or `aspect_specification` is visible.

A stream-oriented attribute for a subtype of a class-wide type  $T$ 'Class is available at places where one of the following conditions is true:

- $T$  is nonlimited;
- the attribute has been specified via an `attribute_definition_clause` or `aspect_specification`, and the `attribute_definition_clause` or `aspect_specification` is visible; or
- the corresponding attribute of  $T$  is available, provided that if  $T$  has a partial view, the corresponding attribute is available at the end of the visible part where  $T$  is declared.

An `attribute_reference` for one of the stream-oriented attributes is illegal unless the attribute is available at the place of the `attribute_reference`. Furthermore, an `attribute_reference` for  $T$ 'Input is illegal if  $T$  is an abstract type. In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

Unless available for a parent type, if any, for an untagged type having a task, protected, or explicitly limited record part, the default implementation of each of the Read, Write, Input, and Output attributes raises `Program_Error` and performs no other action.

In the `parameter_and_result_profiles` for the default implementations of the stream-oriented attributes, the subtype of the *Item* parameter is the base subtype of  $T$  if  $T$  is a scalar type, and the first subtype otherwise. The same rule applies to the result of the Input attribute.

For an `attribute_definition_clause` or `aspect_specification` specifying one of these attributes, the subtype of the *Item* parameter shall be the first subtype or the base subtype if scalar, and the first subtype if not scalar. The same rule applies to the result of the Input function.

A type is said to *support external streaming* if Read and Write attributes are provided for sending values of such a type between active partitions, with Write marshalling the representation, and Read unmarshalling the representation. A limited type supports external streaming only if it has available Read and Write attributes. A type with a part that is of a nonremote access type supports external streaming only if that access type or the type of some part that includes the access type component, has Read and Write attributes that have been specified via an `attribute_definition_clause`, and that `attribute_definition_clause` is visible. An anonymous access type does not support external streaming. All other types (including remote access types, see E.2.2) support external streaming.

#### Erroneous Execution

If the internal tag returned by `Descendant_Tag` to  $T$ 'Class'Input identifies a type that is not library-level and whose tag has not been created, or does not exist in the partition at the time of the call, execution is erroneous.

#### Implementation Requirements

For every subtype  $S$  of a language-defined nonlimited specific type  $T$ , the output generated by  $S$ 'Output or  $S$ 'Write shall be readable by  $S$ 'Input or  $S$ 'Read, respectively. This rule applies across partitions if the implementation conforms to the Distributed Systems Annex.



If `Constraint_Error` is raised during a call to `Read` because of failure of one the above checks, the implementation shall ensure that the discriminants of the actual parameter of `Read` are not modified.

#### Implementation Permissions

The number of calls performed by the predefined implementation of the stream-oriented attributes on the `Read` and `Write` operations of the stream type is unspecified. An implementation may take advantage of this permission to perform internal buffering. However, all the calls on the `Read` and `Write` operations of the stream type used to implement an explicit invocation of a stream-oriented attribute shall take place before this invocation returns. An explicit invocation is one appearing explicitly in the program text, possibly through a generic instantiation (see 15.3).

If  $T$  is a discriminated type and its discriminants have defaults, then in two cases an execution of the default implementation of `S'Read` is not required to create an anonymous object of type  $T$ : If the discriminant values that are read in are equal to the corresponding discriminant values of  $Item$ , then creation of a new object of type  $T$  may be bypassed and  $Item$  may be used instead. If they are not equal and  $Item$  is a constrained variable, then `Constraint_Error` may be raised at that point, before any further values are read from the stream and before the object of type  $T$  is created.

A default implementation of `S'Input` that calls the default implementation of `S'Read` may create a constrained anonymous object with discriminants that match those in the stream.

NOTE 1 For a definite subtype  $S$  of a type  $T$ , only  $T'$ Write and  $T'$ Read are necessary to pass an arbitrary value of the subtype through a stream. For an indefinite subtype  $S$  of a type  $T$ ,  $T'$ Output and  $T'$ Input will normally be necessary, since  $T'$ Write and  $T'$ Read do not pass bounds, discriminants, or tags.

NOTE 2 User-specified attributes of `S'Class` are not inherited by other class-wide types descended from  $S$ .

EXAMPLE Example of user-defined `Write` attribute:

```
procedure My_Write(
  Stream : not null access Ada.Streams.Root_Stream_Type'Class;
  Item   : My_Integer'Base);
for My_Integer'Write use My_Write;
```

## 16.14 Freezing Rules

This subclause defines a place in the program text where each declared entity becomes “frozen”. A use of an entity, such as a reference to it by name, or (for a type) an expression of the type, causes freezing of the entity in some contexts, as described below. The Legality Rules forbid certain kinds of uses of an entity in the region of text where it is frozen.

The *freezing* of an entity occurs at one or more places (*freezing points*) in the program text where the representation for the entity has to be fully determined. Each entity is frozen from its first freezing point to the end of the program text (given the ordering of compilation units defined in 13.1.4).

This subclause also defines a place in the program text where the profile of each declared callable entity becomes *frozen*. A use of a callable entity causes freezing of its profile in some contexts, as described below. At the place where the profile of a callable entity becomes frozen, the entity itself becomes frozen.

The end of a `declarative_part`, `protected_body`, or a declaration of a library package or generic library package, causes *freezing* of each entity and profile declared within it, as well as the entity itself in the case of the declaration of a library unit. A `noninstance proper_body`, `body_stub`, or `entry_body` causes freezing of each entity and profile declared before it within the same `declarative_part`.

A construct that (explicitly or implicitly) references an entity can cause the *freezing* of the entity, as defined by subsequent paragraphs. At the place where a construct causes freezing, each `name`, `expression`, `implicit_dereference`, or `range` within the construct causes freezing:

- The occurrence of a `generic_instantiation` causes freezing, except that a `name` which is a generic actual parameter whose corresponding generic formal parameter is a formal incomplete



type (see 15.5.1) does not cause freezing. In addition, if a parameter of the instantiation is defaulted, the `default_expression` or `default_name` for that parameter causes freezing.

- At the occurrence of an `expression_function_declaration` that is a completion, the return expression of the expression function causes freezing.
- At the occurrence of a `renames-as-body` whose `callable_entity_name` denotes an expression function, the return expression of the expression function causes freezing.
- The occurrence of an `object_declaration` that has no corresponding completion causes freezing.
- The declaration of a record extension causes freezing of the parent subtype.
- The declaration of a record extension, interface type, task unit, or protected unit causes freezing of any progenitor types specified in the declaration.
- At the freezing point of the entity associated with an `aspect_specification`, any static expressions within the `aspect_specification` cause freezing, as do expressions or names in `aspect_definitions` for representation aspects, or operational aspects that have a corresponding operational attribute. Similarly, if an `aspect_definition` for an operational aspect, other than an assertion aspect, can affect the Name Resolution, Static Semantics, or Legality Rules of a subsequent construct, then any expressions or names within the `aspect_definition` cause freezing at the freezing point of the associated entity. Any static expressions within an `aspect_specification` also cause freezing at the end of the immediately enclosing declaration list. For the purposes of this rule, if there is no declared entity associated with an `aspect_specification`, the freezing point is considered to occur immediately following the `aspect_specification`.

A static expression (other than within an `aspect_specification`) causes freezing where it occurs. An object name or nonstatic expression causes freezing where it occurs, unless the name or expression is part of a `default_expression`, a `default_name`, the return expression of an expression function, an `aspect_specification`, or a per-object expression of a component's constraint, in which case, the freezing occurs later as part of another construct or at the freezing point of an associated entity.

An implicit call freezes the same entities and profiles that would be frozen by an explicit call. This is true even if the implicit call is removed via implementation permissions.

If an expression is implicitly converted to a type or subtype *T*, then at the place where the expression causes freezing, *T* is frozen.

The following rules define which entities are frozen at the place where a construct causes freezing:

- At the place where an expression causes freezing, the type of the expression is frozen, unless the expression is an enumeration literal used as a `discrete_choice` of the `array_aggregate` of an `enumeration_representation_clause` or as the `aspect_definition` of a specification for `aspect Default_Value`.
- At the place where a function call causes freezing, the profile of the function is frozen. Furthermore, if a parameter of the call is defaulted, the `default_expression` for that parameter causes freezing. If the function call is to an expression function, the return expression of the expression function causes freezing.
- At the place where a `generic_instantiation` causes freezing of a callable entity, the profile of that entity is frozen unless the formal subprogram corresponding to the callable entity has a parameter or result of a formal untagged incomplete type; if the callable entity is an expression function, the return expression of the expression function causes freezing.
- At the place where a use of the `Access` or `Unchecked_Access` attribute whose `prefix` denotes an expression function causes freezing, the return expression of the expression function causes freezing.

- At the place where a `name` causes freezing, the entity denoted by the `name` is frozen, unless the `name` is a `prefix` of an expanded `name`; at the place where an object `name` causes freezing, the nominal subtype associated with the `name` is frozen.
- At the place where an `implicit_dereference` causes freezing, the nominal subtype associated with the `implicit_dereference` is frozen.
- At the place where a `range` causes freezing, the type of the `range` is frozen.
- At the place where an `allocator` causes freezing, the designated subtype of its type is frozen. If the type of the `allocator` is a derived type, then all ancestor types are also frozen.
- At the place where a profile is frozen, each subtype of the profile is frozen. If the corresponding callable entity is a member of an entry family, the index subtype of the family is frozen.
- At the place where a subtype is frozen, its type is frozen. At the place where a type is frozen, any expressions or names within the full type definition cause freezing, other than those that occur within an `access_type_definition` or an `access_definition`; the first subtype, and any component subtypes, index subtypes, and parent subtype of the type are frozen as well. For a specific tagged type, the corresponding class-wide type is frozen as well. For a class-wide type, the corresponding specific type is frozen as well.
- At the place where a specific tagged type is frozen, the primitive subprograms of the type are frozen. At the place where a type is frozen, any subprogram named in an `attribute_definition_clause` for the type is frozen.
- At the place where a construct causes freezing, if the construct includes a check associated with some assertion aspect (independent of whether the check is enabled), or depends on the definition of some operational aspect as part of its Dynamic Semantics, any names or expressions in the `aspect_definition` for the aspect cause freezing.

Notwithstanding the rest of this subclause, freezing an incomplete view has no effect.

#### *Legality Rules*

The explicit declaration of a primitive subprogram of a tagged type shall occur before the type is frozen (see 6.9.2).

A type shall be completely defined before it is frozen (see 6.11.1 and 10.3).

The completion of a deferred constant declaration shall occur before the constant is frozen (see 10.4).

An operational or representation item that directly specifies an aspect of an entity shall appear before the entity is frozen (see 16.1).

#### *Dynamic Semantics*

The tag (see 6.9) of a tagged type `T` is created at the point where `T` is frozen.

## Annex A (normative)

### Predefined Language Environment

This Annex contains the specifications of library units that shall be provided by every implementation. There are three root library units: Ada, Interfaces, and System; other library units are children of these:

#### Standard — A.1

##### Ada — A.2

- Assertions — 14.4.2
- Asynchronous\_Task\_Control — D.11
- Calendar — 12.6
  - Arithmetic — 12.6.1
  - Formatting — 12.6.1
  - Time\_Zones — 12.6.1
- Characters — A.3.1
  - Conversions — A.3.4
  - Handling — A.3.2
  - Latin\_1 — A.3.3
- Command\_Line — A.15
- Complex\_Text\_IO — G.1.3
- Containers — A.18.1
  - Bounded\_Doubly\_Linked\_Lists — A.18.20
  - Bounded\_Hashed\_Maps — A.18.21
  - Bounded\_Hashed\_Sets — A.18.23
  - Bounded\_Indefinite\_Holders — A.18.32
  - Bounded\_Multiway\_Trees — A.18.25
  - Bounded\_Ordered\_Maps — A.18.22
  - Bounded\_Ordered\_Sets — A.18.24
  - Bounded\_Priority\_Queues — A.18.31
  - Bounded\_Synchronized\_Queues — A.18.29
  - Bounded\_Vectors — A.18.19
  - Doubly\_Linked\_Lists — A.18.3
  - Generic\_Array\_Sort — A.18.26
  - Generic\_Constrained\_Array\_Sort — A.18.26
  - Generic\_Sort — A.18.26
  - Hashed\_Maps — A.18.5
  - Hashed\_Sets — A.18.8
  - Indefinite\_Doubly\_Linked\_Lists — A.18.12
  - Indefinite\_Hashed\_Maps — A.18.13
  - Indefinite\_Hashed\_Sets — A.18.15

#### Standard (...continued)

##### Ada (...continued)

##### Containers (...continued)

- Indefinite\_Holders — A.18.18
- Indefinite\_Multiway\_Trees — A.18.17
- Indefinite\_Ordered\_Maps — A.18.14
- Indefinite\_Ordered\_Sets — A.18.16
- Indefinite\_Vectors — A.18.11
- Multiway\_Trees — A.18.10
- Ordered\_Maps — A.18.6
- Ordered\_Sets — A.18.9
- Synchronized\_Queue\_Interfaces — A.18.27
- Unbounded\_Priority\_Queues — A.18.30
- Unbounded\_Synchronized\_Queues — A.18.28
- Vectors — A.18.2
- Decimal — F.2
- Direct\_IO — A.8.4
- Directories — A.16
  - Hierarchical\_File\_Names — A.16.1
- Information — A.16
- Dispatching — D.2.1
  - EDF — D.2.6
  - Non\_Preemptive — D.2.4
  - Round\_Robin — D.2.5
- Dynamic\_Priorities — D.5.1
- Environment\_Variables — A.17
- Exceptions — 14.4.1
- Execution\_Time — D.14
  - Group\_Budgets — D.14.2
  - Interrupts — D.14.3
  - Timers — D.14.1
- Finalization — 10.6
- Float\_Text\_IO — A.10.9
- Float\_Wide\_Text\_IO — A.11
- Float\_Wide\_Wide\_Text\_IO — A.11

## Standard (...continued)

## Ada (...continued)

Integer\_Text\_IO — A.10.8  
 Integer\_Wide\_Text\_IO — A.11  
 Integer\_Wide\_Wide\_Text\_IO — A.11  
 Interrupts — C.3.2  
   Names — C.3.2  
 IO\_Exceptions — A.13  
 Iterator\_Interfaces — 8.5.1  
 Locales — A.19      Numerics — A.5  
   Big\_Numbers — A.5.5  
   Big\_Integers — A.5.6  
   Big\_Reals — A.5.7  
   Complex\_Arrays — G.3.2  
   Complex\_Elementary\_Functions — G.1.2  
   Complex\_Types — G.1.1  
   Discrete\_Random — A.5.2  
   Elementary\_Functions — A.5.1  
   Float\_Random — A.5.2  
   Generic\_Complex\_Arrays — G.3.2  
   Generic\_Complex\_Elementary\_Functions  
     — G.1.2  
   Generic\_Complex\_Types — G.1.1  
   Generic\_Elementary\_Functions — A.5.1  
   Generic\_Real\_Arrays — G.3.1  
   Real\_Arrays — G.3.1  
 Real\_Time — D.8  
   Timing\_Events — D.15  
 Sequential\_IO — A.8.1  
 Storage\_IO — A.9  
 Streams — 16.13.1  
   Storage\_Streams — 16.13.1  
     Bounded\_FIFO\_Streams — 16.13.1  
     FIFO\_Streams — 16.13.1  
   Stream\_IO — A.12.1  
 Strings — A.4.1  
   Bounded — A.4.4  
     Equal\_Case\_Insensitive — A.4.10  
     Hash — A.4.9  
     Hash\_Case\_Insensitive — A.4.9  
     Less\_Case\_Insensitive — A.4.10  
   Equal\_Case\_Insensitive — A.4.10  
   Fixed — A.4.3  
     Equal\_Case\_Insensitive — A.4.10  
     Hash — A.4.9  
     Hash\_Case\_Insensitive — A.4.9  
     Less\_Case\_Insensitive — A.4.10

## Standard (...continued)

## Ada (...continued)

## Strings (...continued)

Hash — A.4.9  
 Hash\_Case\_Insensitive — A.4.9  
 Less\_Case\_Insensitive — A.4.10  
 Maps — A.4.2  
   Constants — A.4.6  
 Text\_Buffers — A.4.12  
   Bounded — A.4.12  
   Unbounded — A.4.12  
 Unbounded — A.4.5  
   Equal\_Case\_Insensitive — A.4.10  
   Hash — A.4.9  
   Hash\_Case\_Insensitive — A.4.9  
   Less\_Case\_Insensitive — A.4.10  
 UTF\_Encoding — A.4.11  
   Conversions — A.4.11  
   Strings — A.4.11  
   Wide\_Strings — A.4.11  
   Wide\_Wide\_Strings — A.4.11  
 Wide\_Bounded — A.4.7  
   Wide\_Equal\_Case\_Insensitive  
     — A.4.7  
   Wide\_Hash — A.4.7  
   Wide\_Hash\_Case\_Insensitive — A.4.7  
 Wide\_Equal\_Case\_Insensitive — A.4.7  
 Wide\_Fixed — A.4.7  
   Wide\_Equal\_Case\_Insensitive  
     — A.4.7  
   Wide\_Hash — A.4.7  
   Wide\_Hash\_Case\_Insensitive — A.4.7  
 Wide\_Hash — A.4.7  
 Wide\_Hash\_Case\_Insensitive — A.4.7  
 Wide\_Maps — A.4.7  
   Wide\_Constants — A.4.7  
 Wide\_Unbounded — A.4.7  
   Wide\_Equal\_Case\_Insensitive  
     — A.4.7  
   Wide\_Hash — A.4.7  
   Wide\_Hash\_Case\_Insensitive — A.4.7  
 Wide\_Wide\_Bounded — A.4.8  
   Wide\_Wide\_Equal\_Case\_Insensitive  
     — A.4.8  
   Wide\_Wide\_Hash — A.4.8  
   Wide\_Wide\_Hash\_Case\_Insensitive  
     — A.4.8

Standard (...continued)

Ada (...continued)

Strings (...continued)

Wide\_Wide\_Equal\_Case\_Insensitive  
— A.4.8

Wide\_Wide\_Fixed — A.4.8

Wide\_Wide\_Equal\_Case\_Insensitive  
— A.4.8

Wide\_Wide\_Hash — A.4.8

Wide\_Wide\_Hash\_Case\_Insensitive  
— A.4.8

Wide\_Wide\_Hash — A.4.8

Wide\_Wide\_Hash\_Case\_Insensitive  
— A.4.8

Wide\_Wide\_Maps — A.4.8

Wide\_Wide\_Constants — A.4.8

Wide\_Wide\_Unbounded — A.4.8

Wide\_Wide\_Equal\_Case\_Insensitive  
— A.4.8

Wide\_Wide\_Hash — A.4.8

Wide\_Wide\_Hash\_Case\_Insensitive  
— A.4.8

Synchronous\_Barriers — D.10.1

Synchronous\_Task\_Control — D.10

EDF — D.10

Tags — 6.9

Generic\_Dispatching\_Constructor — 6.9

Task\_Attributes — C.7.2

Task\_Identification — C.7.1

Task\_Termination — C.7.3

Text\_IO — A.10.1

Bounded\_IO — A.10.11

Complex\_IO — G.1.3

Editing — F.3.3

Text\_Streams — A.12.2

Unbounded\_IO — A.10.12

Unchecked\_Conversion — 16.9

Unchecked\_Deallocate\_Subpool — 16.11.5

Unchecked\_Deallocation — 16.11.2

Wide\_Characters — A.3.1

Handling — A.3.5

Wide\_Command\_Line — A.15.1

Wide\_Directories — A.16.2

Wide\_Environment\_Variables — A.17.1

Standard (...continued)

Ada (...continued)

Wide\_Text\_IO — A.11

Complex\_IO — G.1.4

Editing — F.3.4

Text\_Streams — A.12.3

Wide\_Bounded\_IO — A.11

Wide\_Unbounded\_IO — A.11

Wide\_Wide\_Characters — A.3.1

Handling — A.3.6

Wide\_Wide\_Command\_Line — A.15.1

Wide\_Wide\_Directories — A.16.2

Wide\_Wide\_Environment\_Variables —  
A.17.1

Wide\_Wide\_Text\_IO — A.11

Complex\_IO — G.1.5

Editing — F.3.5

Text\_Streams — A.12.4

Wide\_Wide\_Bounded\_IO — A.11

Wide\_Wide\_Unbounded\_IO — A.11

Interfaces — B.2

C — B.3

Pointers — B.3.2

Strings — B.3.1

COBOL — B.4

Fortran — B.5

System — 16.7

Address\_To\_Access\_Conversions — 16.7.2

Atomic\_Operations — C.6.1

Exchange — C.6.2

Integer\_Arithmetic — C.6.4

Modular\_Arithmetic — C.6.5

Test\_And\_Set — C.6.3

Machine\_Code — 16.8

Multiprocessors — D.16

Dispatching\_Domains — D.16.1

RPC — E.5

Storage\_Elements — 16.7.1

Storage\_Pools — 16.11

Subpools — 16.11.4

*Implementation Requirements*

The implementation shall ensure that concurrent calls on any two (possibly the same) language-defined subprograms perform as specified, so long as all pairs of objects (one from each call) that are either denoted by parameters that can be passed by reference, or are designated by parameters of an access type, are nonoverlapping.

For the purpose of determining whether concurrent calls on text input-output subprograms are required to perform as specified above, when calling a subprogram within Text\_IO or its children that implicitly operates on one of the default input-output files, the subprogram is considered to have a parameter of Current\_Input or Current\_Output (as appropriate).

If a descendant of a language-defined tagged type is declared, the implementation shall ensure that each inherited language-defined subprogram behaves as described in this document. In particular, overriding a language-defined subprogram shall not alter the effect of any inherited language-defined subprogram.

*Implementation Permissions*

The implementation may restrict the replacement of language-defined compilation units. The implementation may restrict children of language-defined library units (other than Standard).

**A.1 The Package Standard**

This subclause outlines the specification of the package Standard containing all predefined identifiers in the language. The corresponding package body is not specified by the language.

The operators that are predefined for the types declared in the package Standard are given in comments since they are implicitly declared. *Italics* are used for pseudo-names of anonymous types (such as *root\_real*) and for undefined information (such as *implementation-defined*).

*Static Semantics*

The library package Standard has the following declaration:

```

package Standard
  with Pure is

  type Boolean is (False, True);

  -- The predefined relational operators for this type are as follows:
  -- function "="      (Left, Right : Boolean'Base) return Boolean;
  -- function "/="     (Left, Right : Boolean'Base) return Boolean;
  -- function "<"      (Left, Right : Boolean'Base) return Boolean;
  -- function "<="     (Left, Right : Boolean'Base) return Boolean;
  -- function ">"      (Left, Right : Boolean'Base) return Boolean;
  -- function ">="     (Left, Right : Boolean'Base) return Boolean;

  -- The predefined logical operators and the predefined logical
  -- negation operator are as follows:
  -- function "and"    (Left, Right : Boolean'Base) return Boolean'Base;
  -- function "or"     (Left, Right : Boolean'Base) return Boolean'Base;
  -- function "xor"    (Left, Right : Boolean'Base) return Boolean'Base;
  -- function "not"   (Right : Boolean'Base) return Boolean'Base;

  -- The integer type root_integer and the
  -- corresponding universal type universal_integer are predefined.
  type Integer is range implementation-defined;
  subtype Natural is Integer range 0 .. Integer'Last;
  subtype Positive is Integer range 1 .. Integer'Last;

```



```

-- The predefined operators for type Integer are as follows:
-- function "=" (Left, Right : Integer'Base) return Boolean;
-- function "/=" (Left, Right : Integer'Base) return Boolean;
-- function "<" (Left, Right : Integer'Base) return Boolean;
-- function "<=" (Left, Right : Integer'Base) return Boolean;
-- function ">" (Left, Right : Integer'Base) return Boolean;
-- function ">=" (Left, Right : Integer'Base) return Boolean;
--
-- function "+" (Right : Integer'Base) return Integer'Base;
-- function "-" (Right : Integer'Base) return Integer'Base;
-- function "abs" (Right : Integer'Base) return Integer'Base;
--
-- function "+" (Left, Right : Integer'Base) return Integer'Base;
-- function "-" (Left, Right : Integer'Base) return Integer'Base;
-- function "*" (Left, Right : Integer'Base) return Integer'Base;
-- function "/" (Left, Right : Integer'Base) return Integer'Base;
-- function "rem" (Left, Right : Integer'Base) return Integer'Base;
-- function "mod" (Left, Right : Integer'Base) return Integer'Base;
--
-- function "**" (Left : Integer'Base; Right : Natural)
--               return Integer'Base;
--
-- The specification of each operator for the type
-- root_integer, or for any additional predefined integer
-- type, is obtained by replacing Integer by the name of the type
-- in the specification of the corresponding operator of the type
-- Integer. The right operand of the exponentiation operator
-- remains as subtype Natural.
--
-- The floating point type root_real and the
-- corresponding universal type universal_real are predefined.
type Float is digits implementation-defined;
-- The predefined operators for this type are as follows:
--
-- function "=" (Left, Right : Float) return Boolean;
-- function "/=" (Left, Right : Float) return Boolean;
-- function "<" (Left, Right : Float) return Boolean;
-- function "<=" (Left, Right : Float) return Boolean;
-- function ">" (Left, Right : Float) return Boolean;
-- function ">=" (Left, Right : Float) return Boolean;
--
-- function "+" (Right : Float) return Float;
-- function "-" (Right : Float) return Float;
-- function "abs" (Right : Float) return Float;
--
-- function "+" (Left, Right : Float) return Float;
-- function "-" (Left, Right : Float) return Float;
-- function "*" (Left, Right : Float) return Float;
-- function "/" (Left, Right : Float) return Float;
--
-- function "**" (Left : Float; Right : Integer'Base) return Float;
--
-- The specification of each operator for the type root_real, or for
-- any additional predefined floating point type, is obtained by
-- replacing Float by the name of the type in the specification of the
-- corresponding operator of the type Float.
--
-- In addition, the following operators are predefined for the root
-- numeric types:
function "*" (Left : root_integer; Right : root_real)
  return root_real;
function "*" (Left : root_real; Right : root_integer)
  return root_real;
function "/" (Left : root_real; Right : root_integer)
  return root_real;

```

```

-- The type universal_fixed is predefined.
-- The only multiplying operators defined between
-- fixed point types are

function "*" (Left : universal_fixed; Right : universal_fixed)
  return universal_fixed;

function "/" (Left : universal_fixed; Right : universal_fixed)
  return universal_fixed;

-- The type universal_access is predefined.
-- The following equality operators are predefined:

function "=" (Left, Right: universal_access) return Boolean;
function "/=" (Left, Right: universal_access) return Boolean;

-- The declaration of type Character is based on the standard ISO 8859-1
character set.

-- There are no character literals corresponding to the positions for
control characters.
-- They are indicated in italics in this definition. See 6.5.2.

type Character is
  (nul,      soh,  stx,  etx,  eot,  enq,  ack,  bel,  --0 (16#00#) .. 7
  (16#07#)
  bs,       ht,   lf,   vt,   ff,   cr,   so,   si,   --8 (16#08#) .. 15
  (16#0F#)

  dle,      dc1,  dc2,  dc3,  dc4,  nak,  syn,  etb,  --16 (16#10#) .. 23
  (16#17#)
  can,      em,   sub,  esc,  fs,   gs,   rs,   us,   --24 (16#18#) .. 31
  (16#1F#)

  ' ',      '!',  '"',  '#',  '$',  '%',  '&',  "'",  --32 (16#20#) .. 39
  (16#27#)
  '(',      ')',  '*',  '+',  ',',  '-',  '.',  '/',  --40 (16#28#) .. 47
  (16#2F#)

  '0',      '1',  '2',  '3',  '4',  '5',  '6',  '7',  --48 (16#30#) .. 55
  (16#37#)
  '8',      '9',  ':',  ';',  '<',  '=',  '>',  '?',  --56 (16#38#) .. 63
  (16#3F#)

  '@',      'A',  'B',  'C',  'D',  'E',  'F',  'G',  --64 (16#40#) .. 71
  (16#47#)
  'H',      'I',  'J',  'K',  'L',  'M',  'N',  'O',  --72 (16#48#) .. 79
  (16#4F#)

  'P',      'Q',  'R',  'S',  'T',  'U',  'V',  'W',  --80 (16#50#) .. 87
  (16#57#)
  'X',      'Y',  'Z',  '[',  '\',  ']',  '^',  '_',  --88 (16#58#) .. 95
  (16#5F#)

  '`',      'a',  'b',  'c',  'd',  'e',  'f',  'g',  --96 (16#60#) .. 103
  (16#67#)
  'h',      'i',  'j',  'k',  'l',  'm',  'n',  'o',  --104 (16#68#) .. 111
  (16#6F#)

  'p',      'q',  'r',  's',  't',  'u',  'v',  'w',  --112 (16#70#) .. 119
  (16#77#)
  'x',      'y',  'z',  '{',  '|',  '}',  '~',  del,  --120 (16#78#) .. 127
  (16#7F#)

  reserved_128, reserved_129, bph, nbh,          --128 (16#80#)
  .. 131 (16#83#)

```

```

        reserved_132, nel, ssa, esa, --132 (16#84#) .. 135
(16#87#)
        hts, htj, vts, pld, plu, ri, ss2, ss3, --136 (16#88#) .. 143
(16#8F#)

        dcs, pul, pu2, sts, cch, mw, spa, epa, --144 (16#90#) .. 151
(16#97#)
        sos, reserved_153, sci, csi, --152 (16#98#)
.. 155 (16#9B#)
        st, osc, pm, apc, --156 (16#9C#) .. 159
(16#9F#)

        ' ', '!', '¢', '£', '¤', '¥', '¦', '§', --160 (16#A0#) .. 167
(16#A7#)
        '¨', '©', 'ª', «', --168 (16#A8#) .. 171
(16#AB#)
        '¬', soft_hyphen, '®', '¯', --172 (16#AC#)
.. 175 (16#AF#)

        '°', '±', '²', '³', '´', 'µ', '¶', '·', --176 (16#B0#) .. 183
(16#B7#)
        '¸', '¹', 'º', '»', '¼', '½', '¾', '¿', --184 (16#B8#) .. 191
(16#BF#)

        'À', 'Á', 'Â', 'Ã', 'Ä', 'Å', 'Æ', 'Ç', --192 (16#C0#) .. 199
(16#C7#)
        'È', 'É', 'Ê', 'Ë', 'Ì', 'Í', 'Î', 'Ï', --200 (16#C8#) .. 207
(16#CF#)

        'Ð', 'Ñ', 'Ò', 'Ó', 'Ô', 'Õ', 'Ö', '×', --208 (16#D0#) .. 215
(16#D7#)
        'Ø', 'Ù', 'Ú', 'Û', 'Ü', 'Ý', 'Þ', 'ß', --216 (16#D8#) .. 223
(16#DF#)

        'à', 'á', 'â', 'ã', 'ä', 'å', 'æ', 'ç', --224 (16#E0#) .. 231
(16#E7#)
        'è', 'é', 'ê', 'ë', 'ì', 'í', 'î', 'ï', --232 (16#E8#) .. 239
(16#EF#)

        'ð', 'ñ', 'ò', 'ó', 'ô', 'õ', 'ö', '÷', --240 (16#F0#) .. 247
(16#F7#)
        'ø', 'ù', 'ú', 'û', 'ü', 'ý', 'þ', 'ÿ'); --248 (16#F8#) .. 255
(16#FF#)

```

-- The predefined operators for the type Character are the same as for  
-- any enumeration type.

-- The declaration of type Wide\_Character is based on the standard ISO/IEC  
10646:2020 BMP character

-- Set. The first 256 positions have the same contents as type Character.  
See 6.5.2.

**type** Wide\_Character **is** (nul, soh ... Hex\_0000FFFE, Hex\_0000FFFF);

-- The declaration of type Wide\_Wide\_Character is based on the full  
-- ISO/IEC 10646:2020 character set. The first 65536 positions have the  
-- same contents as type Wide\_Character. See 6.5.2.

**type** Wide\_Wide\_Character **is** (nul, soh ... Hex\_7FFFFFFE, Hex\_7FFFFFFF);

**for** Wide\_Wide\_Character'Size **use** 32;

**package** ASCII **is** ... **end** ASCII; --Obsolescent; see I.5

```

-- Predefined string types:

type String is array(Positive range <>) of Character
  with Pack;

-- The predefined operators for this type are as follows:

--   function "=" (Left, Right: String) return Boolean;
--   function "/=" (Left, Right: String) return Boolean;
--   function "<" (Left, Right: String) return Boolean;
--   function "<=" (Left, Right: String) return Boolean;
--   function ">" (Left, Right: String) return Boolean;
--   function ">=" (Left, Right: String) return Boolean;

--   function "&" (Left: String; Right: String) return String;
--   function "&" (Left: Character; Right: String) return String;
--   function "&" (Left: String; Right: Character) return String;
--   function "&" (Left: Character; Right: Character) return String;

type Wide_String is array(Positive range <>) of Wide_Character
  with Pack;

-- The predefined operators for this type correspond to those for String.

type Wide_Wide_String is array (Positive range <>)
  of Wide_Wide_Character
  with Pack;

-- The predefined operators for this type correspond to those for String.

type Duration is delta implementation-defined range implementation-defined;
  -- The predefined operators for the type Duration are the same as for
  -- any fixed point type.

-- The predefined exceptions:
Constraint_Error: exception;
Program_Error   : exception;
Storage_Error   : exception;
Tasking_Error   : exception;

end Standard;

```

Standard has no private part.

In each of the types Character, Wide\_Character, and Wide\_Wide\_Character, the character literals for the space character (position 32) and the non-breaking space character (position 160) correspond to different values. Unless indicated otherwise, each occurrence of the character literal ' ' in this document refers to the space character. Similarly, the character literals for hyphen (position 45) and soft hyphen (position 173) correspond to different values. Unless indicated otherwise, each occurrence of the character literal '-' in this document refers to the hyphen character.

#### Dynamic Semantics

Elaboration of the body of Standard has no effect.

#### Implementation Permissions

An implementation may provide additional predefined integer types and additional predefined floating point types. Some or all of these types may be anonymous.

#### Implementation Advice

If an implementation provides additional named predefined integer types, then the names should end with "Integer" as in "Long\_Integer". If an implementation provides additional named predefined floating point types, then the names should end with "Float" as in "Long\_Float".

NOTE 1 Certain aspects of the predefined entities cannot be completely described in the language itself. For example, although the enumeration type Boolean can be written showing the two enumeration literals False and True, the short-circuit control forms cannot be expressed in the language.

NOTE 2 As explained in 11.1 and 13.1.4, the declarative region of the package Standard encloses every library unit and consequently the main subprogram; the declaration of every library unit is assumed to occur within this declarative region. `Library_items` are assumed to be ordered in such a way that there are no forward semantic dependences. However, as explained in 11.3, the only library units that are visible within a given compilation unit are the library units named by all `with_clauses` that apply to the given unit, and moreover, within the declarative region of a given library unit, that library unit itself.

NOTE 3 If all `block_statements` of a program are named, then the name of each program unit can always be written as an expanded name starting with Standard (unless Standard is itself hidden). The name of a library unit cannot be a homograph of a name (such as Integer) that is already declared in Standard.

NOTE 4 The exception `Standard.Numeric_Error` is defined in I.6.

## A.2 The Package Ada

### *Static Semantics*

The following language-defined library package exists:

```
package Ada
  with Pure is
end Ada;
```

Ada serves as the parent of most of the other language-defined library units; its declaration is empty.

### *Legality Rules*

In the standard mode, it is illegal to compile a child of package Ada.

## A.3 Character Handling

This subclause presents the packages related to character processing: an empty declared pure package Characters and child packages Characters.Handling and Characters.Latin\_1. The package Characters.Handling provides classification and conversion functions for Character data, and some simple functions for dealing with Wide\_Character and Wide\_Wide\_Character data. The child package Characters.Latin\_1 declares a set of constants initialized to values of type Character.

### A.3.1 The Packages Characters, Wide\_Characters, and Wide\_Wide\_Characters

#### *Static Semantics*

The library package Characters has the following declaration:

```
package Ada.Characters
  with Pure is
end Ada.Characters;
```

The library package Wide\_Characters has the following declaration:

```
package Ada.Wide_Characters
  with Pure is
end Ada.Wide_Characters;
```

The library package Wide\_Wide\_Characters has the following declaration:

```
package Ada.Wide_Wide_Characters
  with Pure is
end Ada.Wide_Wide_Characters;
```

#### *Implementation Advice*

If an implementation chooses to provide implementation-defined operations on Wide\_Character or Wide\_String (such as collating and sorting, etc.) it should do so by providing child units of Wide\_Characters. Similarly if it chooses to provide implementation-defined operations on Wide\_Wide\_Character or Wide\_Wide\_String it should do so by providing child units of Wide\_Wide\_Characters.

### A.3.2 The Package Characters.Handling

#### Static Semantics

The library package Characters.Handling has the following declaration:

```

with Ada.Characters.Conversions;
package Ada.Characters.Handling
  with Pure is

  --Character classification functions

  function Is_Control          (Item : in Character) return Boolean;
  function Is_Graphic         (Item : in Character) return Boolean;
  function Is_Letter          (Item : in Character) return Boolean;
  function Is_Lower           (Item : in Character) return Boolean;
  function Is_Upper           (Item : in Character) return Boolean;
  function Is_Basic           (Item : in Character) return Boolean;
  function Is_Digit           (Item : in Character) return Boolean;
  function Is_Decimal_Digit   (Item : in Character) return Boolean
  renames Is_Digit;
  function Is_Hexadecimal_Digit (Item : in Character) return Boolean;
  function Is_Alphanumeric     (Item : in Character) return Boolean;
  function Is_Special          (Item : in Character) return Boolean;
  function Is_Line_Terminator  (Item : in Character) return Boolean;
  function Is_Mark             (Item : in Character) return Boolean;
  function Is_Other_Format     (Item : in Character) return Boolean;
  function Is_Punctuation_Connector (Item : in Character) return Boolean;
  function Is_Space            (Item : in Character) return Boolean;
  function Is_NFKC             (Item : in Character) return Boolean;

  --Conversion functions for Character and String

  function To_Lower (Item : in Character) return Character;
  function To_Upper (Item : in Character) return Character;
  function To_Basic (Item : in Character) return Character;

  function To_Lower (Item : in String) return String;
  function To_Upper (Item : in String) return String;
  function To_Basic (Item : in String) return String;

  --Classifications of and conversions between Character and ISO 646

  subtype ISO_646 is
    Character range Character'Val(0) .. Character'Val(127);

  function Is_ISO_646 (Item : in Character) return Boolean;
  function Is_ISO_646 (Item : in String) return Boolean;

  function To_ISO_646 (Item          : in Character;
    Substitute       : in ISO_646 := ' ')
  return ISO_646;

  function To_ISO_646 (Item          : in String;
    Substitute       : in ISO_646 := ' ')
  return String;

  -- The functions Is_Character, Is_String, To_Character, To_String,
  To_Wide_Character,
  -- and To_Wide_String are obsolescent; see I.14.

end Ada.Characters.Handling;

```

In the description below for each function that returns a Boolean result, the effect is described in terms of the conditions under which the value True is returned. If these conditions are not met, then the function returns False.

Each of the following classification functions has a formal Character parameter, Item, and returns a Boolean result.



Is\_Control

True if Item is a control character. A *control character* is a character whose position is in one of the ranges 0..31 or 127..159.

Is\_Graphic

True if Item is a graphic character. A *graphic character* is a character whose position is in one of the ranges 32..126 or 160..255.

Is\_Letter

True if Item is a letter. A *letter* is a character that is in one of the ranges 'A'..'Z' or 'a'..'z', or whose position is in one of the ranges 192..214, 216..246, or 248..255.

Is\_Lower

True if Item is a lower-case letter. A *lower-case letter* is a character that is in the range 'a'..'z', or whose position is in one of the ranges 223..246 or 248..255.

Is\_Upper

True if Item is an upper-case letter. An *upper-case letter* is a character that is in the range 'A'..'Z' or whose position is in one of the ranges 192..214 or 216.. 222.

Is\_Basic

True if Item is a basic letter. A *basic letter* is a character that is in one of the ranges 'A'..'Z' and 'a'..'z', or that is one of the following: 'Æ', 'æ', 'Ð', 'ð', 'Þ', 'þ', or 'ß'.

Is\_Digit

True if Item is a decimal digit. A *decimal digit* is a character in the range '0'..'9'.

Is\_Decimal\_Digit

A renaming of Is\_Digit.

Is\_Hexadecimal\_Digit

True if Item is a hexadecimal digit. A *hexadecimal digit* is a character that is either a decimal digit or that is in one of the ranges 'A' .. 'F' or 'a' .. 'f'.

Is\_Alphanumeric

True if Item is an alphanumeric character. An *alphanumeric character* is a character that is either a letter or a decimal digit.

Is\_Special

True if Item is a special graphic character. A *special graphic character* is a graphic character that is not alphanumeric.

Is\_Line\_Terminator

True if Item is a character with position 10 .. 13 (Line\_Feed, Line\_Tabulation, Form\_Feed, Carriage\_Return) or 133 (Next\_Line).

Is\_Mark

Never True (no value of type Character has categories Mark, Non-Spacing or Mark, Spacing Combining).

Is\_Other\_Format

True if Item is a character with position 173 (Soft\_Hyphen).

Is\_Punctuation\_Connector

True if Item is a character with position 95 ('\_', known as Low\_Line or Underscore).

Is\_Space

True if Item is a character with position 32 (' ') or 160 (No\_Break\_Space).

Is\_NFKC

True if Item can be present in a string normalized to Normalization Form KC (as defined by Clause 22 of ISO/IEC 10646:2020); this includes all characters except those with positions 160, 168, 170, 175, 178, 179, 180, 181, 184, 185, 186, 188, 189, and 190.

Each of the names To\_Lower, To\_Upper, and To\_Basic refers to two functions: one that converts from Character to Character, and the other that converts from String to String. The result of each Character-to-Character function is described below, in terms of the conversion applied to Item, its formal Character parameter. The result of each String-to-String conversion is obtained by applying to each element of the function's String parameter the corresponding Character-to-Character conversion; the result is the null String if the value of the formal parameter is the null String. The lower bound of the result String is 1.

To\_Lower

Returns the corresponding lower-case value for Item if Is\_Upper(Item), and returns Item otherwise.

**To\_Upper** Returns the corresponding upper-case value for Item if Is\_Lower(Item) and Item has an upper-case form, and returns Item otherwise. The lower case letters 'ß' and 'ÿ' do not have upper case forms.

**To\_Basic** Returns the letter corresponding to Item but with no diacritical mark, if Item is a letter but not a basic letter; returns Item otherwise.

The following set of functions test for membership in the ISO 646 character range, or convert between ISO 646 and Character.

**Is\_ISO\_646**

The function whose formal parameter, Item, is of type Character returns True if Item is in the subtype ISO\_646.

**Is\_ISO\_646**

The function whose formal parameter, Item, is of type String returns True if Is\_ISO\_646(Item(I)) is True for each I in Item'Range.

**To\_ISO\_646**

The function whose first formal parameter, Item, is of type Character returns Item if Is\_ISO\_646(Item), and returns the Substitute ISO\_646 character otherwise.

**To\_ISO\_646**

The function whose first formal parameter, Item, is of type String returns the String whose Range is 1..Item'Length and each of whose elements is given by To\_ISO\_646 of the corresponding element in Item.

NOTE 1 A basic letter is a letter without a diacritical mark.

NOTE 2 Except for the hexadecimal digits, basic letters, and ISO\_646 characters, the categories identified in the classification functions form a strict hierarchy:

- Control characters
- Graphic characters
  - Alphanumeric characters
    - Letters
      - Upper-case letters
      - Lower-case letters
    - Decimal digits
  - Special graphic characters

NOTE 3 There are certain characters which are defined to be lower case letters by ISO 10646 and are therefore allowed in identifiers, but are not considered lower case letters by Ada.Characters.Handling.

### A.3.3 The Package Characters.Latin\_1

The package Characters.Latin\_1 declares constants for characters in ISO 8859-1.

*Static Semantics*

The library package Characters.Latin\_1 has the following declaration:

```
package Ada.Characters.Latin_1
with Pure is
```

*-- Control characters:*

```

NUL          : constant Character := Character'Val(0);
SOH          : constant Character := Character'Val(1);
STX          : constant Character := Character'Val(2);
ETX          : constant Character := Character'Val(3);
EOT          : constant Character := Character'Val(4);
ENQ          : constant Character := Character'Val(5);
ACK          : constant Character := Character'Val(6);
BEL          : constant Character := Character'Val(7);
BS           : constant Character := Character'Val(8);
HT           : constant Character := Character'Val(9);
LF           : constant Character := Character'Val(10);
VT           : constant Character := Character'Val(11);
FF           : constant Character := Character'Val(12);
CR           : constant Character := Character'Val(13);
SO           : constant Character := Character'Val(14);
SI           : constant Character := Character'Val(15);

DLE          : constant Character := Character'Val(16);
DC1          : constant Character := Character'Val(17);
DC2          : constant Character := Character'Val(18);
DC3          : constant Character := Character'Val(19);
DC4          : constant Character := Character'Val(20);
NAK          : constant Character := Character'Val(21);
SYN          : constant Character := Character'Val(22);
ETB          : constant Character := Character'Val(23);
CAN          : constant Character := Character'Val(24);
EM           : constant Character := Character'Val(25);
SUB          : constant Character := Character'Val(26);
ESC          : constant Character := Character'Val(27);
FS           : constant Character := Character'Val(28);
GS           : constant Character := Character'Val(29);
RS           : constant Character := Character'Val(30);
US           : constant Character := Character'Val(31);

```

*-- ISO 646 graphic characters:*

```

Space        : constant Character := ' '; -- Character'Val(32)
Exclamation  : constant Character := '!'; -- Character'Val(33)
Quotation    : constant Character := '"'; -- Character'Val(34)
Number_Sign  : constant Character := '#'; -- Character'Val(35)
Dollar_Sign  : constant Character := '$'; -- Character'Val(36)
Percent_Sign : constant Character := '%'; -- Character'Val(37)
Ampersand    : constant Character := '&'; -- Character'Val(38)
Apostrophe   : constant Character := '\''; -- Character'Val(39)
Left_Parenthesis : constant Character := '('; -- Character'Val(40)
Right_Parenthesis : constant Character := ')'; -- Character'Val(41)
Asterisk     : constant Character := '*'; -- Character'Val(42)
Plus_Sign    : constant Character := '+'; -- Character'Val(43)
Comma        : constant Character := ','; -- Character'Val(44)
Hyphen       : constant Character := '-'; -- Character'Val(45)
Minus_Sign   : constant Character renames Hyphen;
Full_Stop    : constant Character := '.'; -- Character'Val(46)
Solidus      : constant Character := '/'; -- Character'Val(47)

```

*-- Decimal digits '0' through '9' are at positions 48 through 57*

```

Colon        : constant Character := ':'; -- Character'Val(58)
Semicolon    : constant Character := ';'; -- Character'Val(59)
Less_Than_Sign : constant Character := '<'; -- Character'Val(60)
Equals_Sign   : constant Character := '='; -- Character'Val(61)
Greater_Than_Sign : constant Character := '>'; -- Character'Val(62)
Question     : constant Character := '?'; -- Character'Val(63)
Commercial_At : constant Character := '@'; -- Character'Val(64)

```

```

-- Letters 'A' through 'Z' are at positions 65 through 90

Left_Square_Bracket  : constant Character := '['; -- Character'Val (91)
Reverse_Solidus      : constant Character := '\'; -- Character'Val (92)
Right_Square_Bracket : constant Character := ']'; -- Character'Val (93)
Circumflex           : constant Character := '^'; -- Character'Val (94)
Low_Line              : constant Character := '_'; -- Character'Val (95)

Grave                 : constant Character := '`'; -- Character'Val (96)
LC_A                  : constant Character := 'a'; -- Character'Val (97)
LC_B                  : constant Character := 'b'; -- Character'Val (98)
LC_C                  : constant Character := 'c'; -- Character'Val (99)
LC_D                  : constant Character := 'd'; -- Character'Val (100)
LC_E                  : constant Character := 'e'; -- Character'Val (101)
LC_F                  : constant Character := 'f'; -- Character'Val (102)
LC_G                  : constant Character := 'g'; -- Character'Val (103)
LC_H                  : constant Character := 'h'; -- Character'Val (104)
LC_I                  : constant Character := 'i'; -- Character'Val (105)
LC_J                  : constant Character := 'j'; -- Character'Val (106)
LC_K                  : constant Character := 'k'; -- Character'Val (107)
LC_L                  : constant Character := 'l'; -- Character'Val (108)
LC_M                  : constant Character := 'm'; -- Character'Val (109)
LC_N                  : constant Character := 'n'; -- Character'Val (110)
LC_O                  : constant Character := 'o'; -- Character'Val (111)

LC_P                  : constant Character := 'p'; -- Character'Val (112)
LC_Q                  : constant Character := 'q'; -- Character'Val (113)
LC_R                  : constant Character := 'r'; -- Character'Val (114)
LC_S                  : constant Character := 's'; -- Character'Val (115)
LC_T                  : constant Character := 't'; -- Character'Val (116)
LC_U                  : constant Character := 'u'; -- Character'Val (117)
LC_V                  : constant Character := 'v'; -- Character'Val (118)
LC_W                  : constant Character := 'w'; -- Character'Val (119)
LC_X                  : constant Character := 'x'; -- Character'Val (120)
LC_Y                  : constant Character := 'y'; -- Character'Val (121)
LC_Z                  : constant Character := 'z'; -- Character'Val (122)
Left_Curly_Bracket    : constant Character := '{'; -- Character'Val (123)
Vertical_Line          : constant Character := '|'; -- Character'Val (124)
Right_Curly_Bracket   : constant Character := '}'; -- Character'Val (125)
Tilde                 : constant Character := '~'; -- Character'Val (126)
DEL                   : constant Character := Character'Val (127);

-- ISO 6429 control characters:

IS4                   : Character renames FS;
IS3                   : Character renames GS;
IS2                   : Character renames RS;
IS1                   : Character renames US;

Reserved_128          : constant Character := Character'Val (128);
Reserved_129          : constant Character := Character'Val (129);
BPH                   : constant Character := Character'Val (130);
NBH                   : constant Character := Character'Val (131);
Reserved_132          : constant Character := Character'Val (132);
NEL                   : constant Character := Character'Val (133);
SSA                   : constant Character := Character'Val (134);
ESA                   : constant Character := Character'Val (135);
HTS                   : constant Character := Character'Val (136);
HTJ                   : constant Character := Character'Val (137);
VTS                   : constant Character := Character'Val (138);
PLD                   : constant Character := Character'Val (139);
PLU                   : constant Character := Character'Val (140);
RI                    : constant Character := Character'Val (141);
SS2                   : constant Character := Character'Val (142);
SS3                   : constant Character := Character'Val (143);

```

```

DCS                : constant Character := Character'Val(144);
PU1                : constant Character := Character'Val(145);
PU2                : constant Character := Character'Val(146);
STS                : constant Character := Character'Val(147);
CCH                : constant Character := Character'Val(148);
MW                 : constant Character := Character'Val(149);
SPA                : constant Character := Character'Val(150);
EPA                : constant Character := Character'Val(151);

SOS                : constant Character := Character'Val(152);
Reserved_153       : constant Character := Character'Val(153);
SCI                : constant Character := Character'Val(154);
CSI                : constant Character := Character'Val(155);
ST                 : constant Character := Character'Val(156);
OSC                : constant Character := Character'Val(157);
PM                 : constant Character := Character'Val(158);
APC                : constant Character := Character'Val(159);

-- Other graphic characters:
-- Character positions 160 (16#A0#) .. 175 (16#AF#):
No_Break_Space      : constant Character := ' '; --
Character'Val(160)
NBSP                : Character renames No_Break_Space;
Inverted_Exclamation : constant Character := '¡'; --
Character'Val(161)
Cent_Sign           : constant Character := '¢'; --
Character'Val(162)
Pound_Sign          : constant Character := '£'; --
Character'Val(163)
Currency_Sign       : constant Character := '¤'; --
Character'Val(164)
Yen_Sign            : constant Character := '¥'; --
Character'Val(165)
Broken_Bar          : constant Character := '¦'; --
Character'Val(166)
Section_Sign        : constant Character := '§'; --
Character'Val(167)
Diaeresis           : constant Character := '¨'; --
Character'Val(168)
Copyright_Sign      : constant Character := '©'; --
Character'Val(169)
Feminine_Ordinal_Indicator : constant Character := 'ª'; --
Character'Val(170)
Left_Angle_Quotation : constant Character := '«'; --
Character'Val(171)
Not_Sign            : constant Character := '¬'; --
Character'Val(172)
Soft_Hyphen         : constant Character := Character'Val(173);
Registered_Trade_Mark_Sign : constant Character := '®'; --
Character'Val(174)
Macron              : constant Character := '¯'; --
Character'Val(175)

```

```

-- Character positions 176 (16#B0#) .. 191 (16#BF#):
Degree_Sign : constant Character := '°'; --
Character'Val(176)
Ring_Above : Character renames Degree_Sign;
Plus_Minus_Sign : constant Character := '±'; --
Character'Val(177)
Superscript_Two : constant Character := '²'; --
Character'Val(178)
Superscript_Three : constant Character := '³'; --
Character'Val(179)
Acute : constant Character := '´'; --
Character'Val(180)
Micro_Sign : constant Character := 'µ'; --
Character'Val(181)
Pilcrow_Sign : constant Character := '¶'; --
Character'Val(182)
Paragraph_Sign : Character renames Pilcrow_Sign;
Middle_Dot : constant Character := '·'; --
Character'Val(183)
Cedilla : constant Character := '¸'; --
Character'Val(184)
Superscript_One : constant Character := '¹'; --
Character'Val(185)
Masculine_Ordinal_Indicator : constant Character := 'º'; --
Character'Val(186)
Right_Angle_Quotation : constant Character := '»'; --
Character'Val(187)
Fraction_One_Quarter : constant Character := '¼'; --
Character'Val(188)
Fraction_One_Half : constant Character := '½'; --
Character'Val(189)
Fraction_Three_Quarters : constant Character := '¾'; --
Character'Val(190)
Inverted_Question : constant Character := '¿'; --
Character'Val(191)

```



```

-- Character positions 192 (16#C0#) .. 207 (16#CF#):
    UC_A_Grave                : constant Character := 'À'; --
Character'Val(192)
    UC_A_Acute                : constant Character := 'Á'; --
Character'Val(193)
    UC_A_Circumflex           : constant Character := 'Â'; --
Character'Val(194)
    UC_A_Tilde                : constant Character := 'Ã'; --
Character'Val(195)
    UC_A_Diaeresis            : constant Character := 'Ä'; --
Character'Val(196)
    UC_A_Ring                 : constant Character := 'Å'; --
Character'Val(197)
    UC_AE_Diphthong           : constant Character := 'Æ'; --
Character'Val(198)
    UC_C_Cedilla              : constant Character := 'Ç'; --
Character'Val(199)
    UC_E_Grave                : constant Character := 'È'; --
Character'Val(200)
    UC_E_Acute                : constant Character := 'É'; --
Character'Val(201)
    UC_E_Circumflex           : constant Character := 'Ê'; --
Character'Val(202)
    UC_E_Diaeresis            : constant Character := 'Ë'; --
Character'Val(203)
    UC_I_Grave                : constant Character := 'Ì'; --
Character'Val(204)
    UC_I_Acute                : constant Character := 'Í'; --
Character'Val(205)
    UC_I_Circumflex           : constant Character := 'Î'; --
Character'Val(206)
    UC_I_Diaeresis            : constant Character := 'Ï'; --
Character'Val(207)

```

```

-- Character positions 208 (16#D0#) .. 223 (16#DF#):
    UC_Icelandic_Eth          : constant Character := 'Ð'; --
Character'Val(208)
    UC_N_Tilde                : constant Character := 'Ñ'; --
Character'Val(209)
    UC_O_Grave                : constant Character := 'Ò'; --
Character'Val(210)
    UC_O_Acute                : constant Character := 'Ó'; --
Character'Val(211)
    UC_O_Circumflex           : constant Character := 'Ô'; --
Character'Val(212)
    UC_O_Tilde                : constant Character := 'Õ'; --
Character'Val(213)
    UC_O_Diaeresis            : constant Character := 'Ö'; --
Character'Val(214)
    Multiplication_Sign        : constant Character := '×'; --
Character'Val(215)
    UC_O_Oblique_Stroke       : constant Character := 'Ø'; --
Character'Val(216)
    UC_U_Grave                : constant Character := 'Ù'; --
Character'Val(217)
    UC_U_Acute                : constant Character := 'Ú'; --
Character'Val(218)
    UC_U_Circumflex           : constant Character := 'Û'; --
Character'Val(219)
    UC_U_Diaeresis            : constant Character := 'Ü'; --
Character'Val(220)
    UC_Y_Acute                : constant Character := 'Ý'; --
Character'Val(221)
    UC_Icelandic_Thorn        : constant Character := 'Þ'; --
Character'Val(222)
    LC_German_Sharp_S         : constant Character := 'ß'; --
Character'Val(223)

```

```

-- Character positions 224 (16#E0#) .. 239 (16#EF#):
  LC_A_Grave                : constant Character := 'à'; --
Character'Val(224)
  LC_A_Acute                : constant Character := 'á'; --
Character'Val(225)
  LC_A_Circumflex          : constant Character := 'â'; --
Character'Val(226)
  LC_A_Tilde               : constant Character := 'ã'; --
Character'Val(227)
  LC_A_Diaeresis           : constant Character := 'ä'; --
Character'Val(228)
  LC_A_Ring                : constant Character := 'å'; --
Character'Val(229)
  LC_AE_Diphthong          : constant Character := 'æ'; --
Character'Val(230)
  LC_C_Cedilla             : constant Character := 'ç'; --
Character'Val(231)
  LC_E_Grave               : constant Character := 'è'; --
Character'Val(232)
  LC_E_Acute               : constant Character := 'é'; --
Character'Val(233)
  LC_E_Circumflex          : constant Character := 'ê'; --
Character'Val(234)
  LC_E_Diaeresis           : constant Character := 'ë'; --
Character'Val(235)
  LC_I_Grave               : constant Character := 'ì'; --
Character'Val(236)
  LC_I_Acute               : constant Character := 'í'; --
Character'Val(237)
  LC_I_Circumflex          : constant Character := 'î'; --
Character'Val(238)
  LC_I_Diaeresis           : constant Character := 'ï'; --
Character'Val(239)

```

```

-- Character positions 240 (16#F0#) .. 255 (16#FF#):
  LC_Icelandic_Eth      : constant Character := 'ð'; --
Character'Val(240)
  LC_N_Tilde            : constant Character := 'ñ'; --
Character'Val(241)
  LC_O_Grave           : constant Character := 'ò'; --
Character'Val(242)
  LC_O_Acute           : constant Character := 'ó'; --
Character'Val(243)
  LC_O_Circumflex      : constant Character := 'ô'; --
Character'Val(244)
  LC_O_Tilde           : constant Character := 'õ'; --
Character'Val(245)
  LC_O_Diaeresis       : constant Character := 'ö'; --
Character'Val(246)
  Division_Sign        : constant Character := '÷'; --
Character'Val(247)
  LC_O_Oblique_Stroke  : constant Character := 'ø'; --
Character'Val(248)
  LC_U_Grave           : constant Character := 'ù'; --
Character'Val(249)
  LC_U_Acute           : constant Character := 'ú'; --
Character'Val(250)
  LC_U_Circumflex      : constant Character := 'û'; --
Character'Val(251)
  LC_U_Diaeresis       : constant Character := 'ü'; --
Character'Val(252)
  LC_Y_Acute           : constant Character := 'ý'; --
Character'Val(253)
  LC_Icelandic_Thorn   : constant Character := 'þ'; --
Character'Val(254)
  LC_Y_Diaeresis       : constant Character := 'ÿ'; --
Character'Val(255)
end Ada.Characters.Latin_1;

```

#### Implementation Permissions

An implementation may provide additional packages as children of Ada.Characters, to declare names for the symbols of the local character set or other character sets.

### A.3.4 The Package Characters.Conversions

#### Static Semantics

The library package Characters.Conversions has the following declaration:

```

package Ada.Characters.Conversions
with Pure is
  function Is_Character (Item : in Wide_Character) return Boolean;
  function Is_String    (Item : in Wide_String)    return Boolean;
  function Is_Character (Item : in Wide_Wide_Character) return Boolean;
  function Is_String    (Item : in Wide_Wide_String) return Boolean;
  function Is_Wide_Character (Item : in Wide_Wide_Character)
    return Boolean;
  function Is_Wide_String   (Item : in Wide_Wide_String)
    return Boolean;

```

```

function To_Wide_Character (Item : in Character) return Wide_Character;
function To_Wide_String    (Item : in String)    return Wide_String;
function To_Wide_Wide_Character (Item : in Character)
    return Wide_Wide_Character;
function To_Wide_Wide_String    (Item : in String)
    return Wide_Wide_String;
function To_Wide_Wide_Character (Item : in Wide_Character)
    return Wide_Wide_Character;
function To_Wide_Wide_String    (Item : in Wide_String)
    return Wide_Wide_String;

function To_Character (Item      : in Wide_Character;
                       Substitute : in Character := ' ')
    return Character;
function To_String    (Item      : in Wide_String;
                       Substitute : in Character := ' ')
    return String;
function To_Character (Item :      in Wide_Wide_Character;
                       Substitute : in Character := ' ')
    return Character;
function To_String    (Item :      in Wide_Wide_String;
                       Substitute : in Character := ' ')
    return String;
function To_Wide_Character (Item :      in Wide_Wide_Character;
                           Substitute : in Wide_Character := ' ')
    return Wide_Character;
function To_Wide_String    (Item :      in Wide_Wide_String;
                           Substitute : in Wide_Character := ' ')
    return Wide_String;
end Ada.Characters.Conversions;

```

The functions in package Characters.Conversions test Wide\_Wide\_Character or Wide\_Character values for membership in Wide\_Character or Character, or convert between corresponding characters of Wide\_Wide\_Character, Wide\_Character, and Character.

```

function Is_Character (Item : in Wide_Character) return Boolean;
    Returns True if Wide_Character'Pos(Item) <= Character'Pos(Character'Last).

function Is_Character (Item : in Wide_Wide_Character) return Boolean;
    Returns True if Wide_Wide_Character'Pos(Item) <= Character'Pos(Character'Last).

function Is_Wide_Character (Item : in Wide_Wide_Character) return Boolean;
    Returns True if Wide_Wide_Character'Pos(Item) <= Wide_Character'Pos(Wide_Character'Last).

function Is_String (Item : in Wide_String)      return Boolean;
function Is_String (Item : in Wide_Wide_String) return Boolean;
    Returns True if Is_Character(Item(I)) is True for each I in Item'Range.

function Is_Wide_String (Item : in Wide_Wide_String) return Boolean;
    Returns True if Is_Wide_Character(Item(I)) is True for each I in Item'Range.

function To_Character (Item :      in Wide_Character;
                       Substitute : in Character := ' ') return Character;
function To_Character (Item :      in Wide_Wide_Character;
                       Substitute : in Character := ' ') return Character;
    Returns the Character corresponding to Item if Is_Character(Item), and returns the Substitute Character otherwise.

function To_Wide_Character (Item : in Character) return Wide_Character;
    Returns the Wide_Character X such that Character'Pos(Item) = Wide_Character'Pos (X).

```

```

function To_Wide_Character (Item : in Wide_Wide_Character;
                           Substitute : in Wide_Character := ' ')
  return Wide_Character;

```

Returns the Wide\_Character corresponding to Item if Is\_Wide\_Character(Item), and returns the Substitute Wide\_Character otherwise.

```

function To_Wide_Wide_Character (Item : in Character)
  return Wide_Wide_Character;

```

Returns the Wide\_Wide\_Character X such that Character'Pos(Item) = Wide\_Wide\_Character'Pos(X).

```

function To_Wide_Wide_Character (Item : in Wide_Character)
  return Wide_Wide_Character;

```

Returns the Wide\_Wide\_Character X such that Wide\_Character'Pos(Item) = Wide\_Wide\_Character'Pos(X).

```

function To_String (Item : in Wide_String;
                   Substitute : in Character := ' ') return String;

```

```

function To_String (Item : in Wide_Wide_String;
                   Substitute : in Character := ' ') return String;

```

Returns the String whose range is 1..Item'Length and each of whose elements is given by To\_Character of the corresponding element in Item.

```

function To_Wide_String (Item : in String) return Wide_String;

```

Returns the Wide\_String whose range is 1..Item'Length and each of whose elements is given by To\_Wide\_Character of the corresponding element in Item.

```

function To_Wide_String (Item : in Wide_Wide_String;
                       Substitute : in Wide_Character := ' ')
  return Wide_String;

```

Returns the Wide\_String whose range is 1..Item'Length and each of whose elements is given by To\_Wide\_Character of the corresponding element in Item with the given Substitute Wide\_Character.

```

function To_Wide_Wide_String (Item : in String) return Wide_Wide_String;

```

```

function To_Wide_Wide_String (Item : in Wide_String)
  return Wide_Wide_String;

```

Returns the Wide\_Wide\_String whose range is 1..Item'Length and each of whose elements is given by To\_Wide\_Wide\_Character of the corresponding element in Item.

### A.3.5 The Package Wide\_Characters.Handling

The package Wide\_Characters.Handling provides operations for classifying Wide\_Characters and case folding for Wide\_Characters.

#### Static Semantics

The library package Wide\_Characters.Handling has the following declaration:

```

package Ada.Wide_Characters.Handling
  with Pure is

  function Character_Set_Version return String;

  function Is_Control (Item : Wide_Character) return Boolean;

  function Is_Letter (Item : Wide_Character) return Boolean;

  function Is_Lower (Item : Wide_Character) return Boolean;

  function Is_Upper (Item : Wide_Character) return Boolean;

  function Is_Basic (Item : Wide_Character) return Boolean;

```



```

function Is_Digit (Item : Wide_Character) return Boolean;
function Is_Decimal_Digit (Item : Wide_Character) return Boolean
    renames Is_Digit;
function Is_Hexadecimal_Digit (Item : Wide_Character) return Boolean;
function Is_Alphanumeric (Item : Wide_Character) return Boolean;
function Is_Special (Item : Wide_Character) return Boolean;
function Is_Line_Terminator (Item : Wide_Character) return Boolean;
function Is_Mark (Item : Wide_Character) return Boolean;
function Is_Other_Format (Item : Wide_Character) return Boolean;
function Is_Punctuation_Connector (Item : Wide_Character) return Boolean;
function Is_Space (Item : Wide_Character) return Boolean;
function Is_NFKC (Item : Wide_Character) return Boolean;
function Is_Graphic (Item : Wide_Character) return Boolean;
function To_Lower (Item : Wide_Character) return Wide_Character;
function To_Upper (Item : Wide_Character) return Wide_Character;
function To_Basic (Item : Wide_Character) return Wide_Character;
function To_Lower (Item : Wide_String) return Wide_String;
function To_Upper (Item : Wide_String) return Wide_String;
function To_Basic (Item : Wide_String) return Wide_String;
end Ada.Wide_Characters.Handling;

```

The subprograms defined in `Wide_Characters.Handling` are locale independent.

```

function Character_Set_Version return String;

```

Returns an implementation-defined identifier that identifies the version of the character set standard that is used for categorizing characters by the implementation.

```

function Is_Control (Item : Wide_Character) return Boolean;

```

Returns True if the `Wide_Character` designated by `Item` is categorized as `other_control`; otherwise returns False.

```

function Is_Letter (Item : Wide_Character) return Boolean;

```

Returns True if the `Wide_Character` designated by `Item` is categorized as `letter_uppercase`, `letter_lowercase`, `letter_titlecase`, `letter_modifier`, `letter_other`, or `number_letter`; otherwise returns False.

```

function Is_Lower (Item : Wide_Character) return Boolean;

```

Returns True if the `Wide_Character` designated by `Item` is categorized as `letter_lowercase`; otherwise returns False.

```

function Is_Upper (Item : Wide_Character) return Boolean;

```

Returns True if the `Wide_Character` designated by `Item` is categorized as `letter_uppercase`; otherwise returns False.

```

function Is_Basic (Item : Wide_Character) return Boolean;

```

Returns True if the `Wide_Character` designated by `Item` has no Decomposition Mapping in the code charts of ISO/IEC 10646:2020; otherwise returns False.

```

function Is_Digit (Item : Wide_Character) return Boolean;

```

Returns True if the `Wide_Character` designated by `Item` is categorized as `number_decimal`; otherwise returns False.

**function** Is\_Hexadecimal\_Digit (Item : Wide\_Character) **return** Boolean;

Returns True if the Wide\_Character designated by Item is categorized as number\_decimal, or is in the range 'A' .. 'F' or 'a' .. 'f'; otherwise returns False.

**function** Is\_Alphanumeric (Item : Wide\_Character) **return** Boolean;

Returns True if the Wide\_Character designated by Item is categorized as letter\_uppercase, letter\_lowercase, letter\_titlecase, letter\_modifier, letter\_other, number\_letter, or number\_decimal; otherwise returns False.

**function** Is\_Special (Item : Wide\_Character) **return** Boolean;

Returns True if the Wide\_Character designated by Item is categorized as graphic\_character, but not categorized as letter\_uppercase, letter\_lowercase, letter\_titlecase, letter\_modifier, letter\_other, number\_letter, or number\_decimal; otherwise returns False.

**function** Is\_Line\_Terminator (Item : Wide\_Character) **return** Boolean;

Returns True if the Wide\_Character designated by Item is categorized as separator\_line or separator\_paragraph, or if Item is a conventional line terminator character (Line\_Feed, Line\_Tabulation, Form\_Feed, Carriage\_Return, Next\_Line); otherwise returns False.

**function** Is\_Mark (Item : Wide\_Character) **return** Boolean;

Returns True if the Wide\_Character designated by Item is categorized as mark\_non\_spacing or mark\_spacing\_combining; otherwise returns False.

**function** Is\_Other\_Format (Item : Wide\_Character) **return** Boolean;

Returns True if the Wide\_Character designated by Item is categorized as other\_format; otherwise returns False.

**function** Is\_Punctuation\_Connector (Item : Wide\_Character) **return** Boolean;

Returns True if the Wide\_Character designated by Item is categorized as punctuation\_connector; otherwise returns False.

**function** Is\_Space (Item : Wide\_Character) **return** Boolean;

Returns True if the Wide\_Character designated by Item is categorized as separator\_space; otherwise returns False.

**function** Is\_NFKC (Item : Wide\_Character) **return** Boolean;

Returns True if the Wide\_Character designated by Item can be present in a string normalized to Normalization Form KC (as defined by Clause 22 of ISO/IEC 10646:2020), otherwise returns False.

**function** Is\_Graphic (Item : Wide\_Character) **return** Boolean;

Returns True if the Wide\_Character designated by Item is categorized as graphic\_character; otherwise returns False.

**function** To\_Lower (Item : Wide\_Character) **return** Wide\_Character;

Returns the Simple Lowercase Mapping as defined by documents referenced in Clause 2 of ISO/IEC 10646:2020 of the Wide\_Character designated by Item. If the Simple Lowercase Mapping does not exist for the Wide\_Character designated by Item, then the value of Item is returned.

**function** To\_Lower (Item : Wide\_String) **return** Wide\_String;

Returns the result of applying the To\_Lower conversion to each Wide\_Character element of the Wide\_String designated by Item. The result is the null Wide\_String if the value of the formal parameter is the null Wide\_String. The lower bound of the result Wide\_String is 1.

**function** To\_Upper (Item : Wide\_Character) **return** Wide\_Character;

Returns the Simple Uppercase Mapping as defined by documents referenced in Clause 2 of ISO/IEC 10646:2020 of the Wide\_Character designated by Item. If the Simple Uppercase Mapping does not exist for the Wide\_Character designated by Item, then the value of Item is returned.

**function** To\_Upper (Item : Wide\_String) **return** Wide\_String;

Returns the result of applying the To\_Upper conversion to each Wide\_Character element of the Wide\_String designated by Item. The result is the null Wide\_String if the value of the formal parameter is the null Wide\_String. The lower bound of the result Wide\_String is 1.

**function** To\_Basic (Item : Wide\_Character) **return** Wide\_Character;

Returns the Wide\_Character whose code point is given by the first value of its Decomposition Mapping in the code charts of ISO/IEC 10646:2020 if any; returns Item otherwise.

**function** To\_Basic (Item : Wide\_String) **return** Wide\_String;

Returns the result of applying the To\_Basic conversion to each Wide\_Character element of the Wide\_String designated by Item. The result is the null Wide\_String if the value of the formal parameter is the null Wide\_String. The lower bound of the result Wide\_String is 1.

#### *Implementation Advice*

The string returned by Character\_Set\_Version should include either “10646:” or “Unicode”.

NOTE 1 The results returned by these functions can depend on which particular version of ISO/IEC 10646 is supported by the implementation (see 5.1).

NOTE 2 The case insensitive equality comparison routines provided in A.4.10 are also available for wide strings (see A.4.7).

### **A.3.6 The Package Wide\_Wide\_Characters.Handling**

The package Wide\_Wide\_Characters.Handling has the same contents as Wide\_Characters.Handling except that each occurrence of Wide\_Character is replaced by Wide\_Wide\_Character, and each occurrence of Wide\_String is replaced by Wide\_Wide\_String.

## **A.4 String Handling**

This subclause presents the specifications of the package Strings and several child packages, which provide facilities for dealing with string data. Fixed-length, bounded-length, and unbounded-length strings are supported, for String, Wide\_String, and Wide\_Wide\_String. The string-handling subprograms include searches for pattern strings and for characters in program-specified sets, translation (via a character-to-character mapping), and transformation (replacing, inserting, overwriting, and deleting of substrings).

### **A.4.1 The Package Strings**

The package Strings provides declarations common to the string handling packages.

#### *Static Semantics*

The library package Strings has the following declaration:

```
package Ada.Strings
with Pure is
```

```

Space      : constant Character      := ' ';
Wide_Space : constant Wide_Character := ' ';
Wide_Wide_Space : constant Wide_Wide_Character := ' ';

Length_Error, Pattern_Error, Index_Error, Translation_Error : exception;

type Alignment is (Left, Right, Center);
type Truncation is (Left, Right, Error);
type Membership is (Inside, Outside);
type Direction is (Forward, Backward);
type Trim_End is (Left, Right, Both);
end Ada.Strings;

```

## A.4.2 The Package Strings.Maps

The package Strings.Maps defines the types, operations, and other entities necessary for character sets and character-to-character mappings.

### Static Semantics

The library package Strings.Maps has the following declaration:

```

package Ada.Strings.Maps
with Pure is
    -- Representation for a set of character values:
    type Character_Set is private
        with Preelaborable_Initialization;

    Null_Set : constant Character_Set;

    type Character_Range is
        record
            Low   : Character;
            High  : Character;
        end record;
    -- Represents Character range Low..High

    type Character_Ranges is array (Positive range <>) of Character_Range;

    function To_Set (Ranges : in Character_Ranges) return Character_Set;
    function To_Set (Span   : in Character_Range) return Character_Set;
    function To_Ranges (Set   : in Character_Set) return Character_Ranges;
    function "=" (Left, Right : in Character_Set) return Boolean;
    function "not" (Right : in Character_Set) return Character_Set;
    function "and" (Left, Right : in Character_Set) return Character_Set;
    function "or" (Left, Right : in Character_Set) return Character_Set;
    function "xor" (Left, Right : in Character_Set) return Character_Set;
    function "-" (Left, Right : in Character_Set) return Character_Set;

    function Is_In (Element : in Character;
                  Set      : in Character_Set)
        return Boolean;

    function Is_Subset (Elements : in Character_Set;
                     Set        : in Character_Set)
        return Boolean;

    function "<=" (Left   : in Character_Set;
                Right  : in Character_Set)
        return Boolean renames Is_Subset;

    -- Alternative representation for a set of character values:
    subtype Character_Sequence is String;

    function To_Set (Sequence : in Character_Sequence) return Character_Set;
    function To_Set (Singleton : in Character) return Character_Set;
    function To_Sequence (Set : in Character_Set) return Character_Sequence;

```

```

-- Representation for a character to character mapping:
type Character_Mapping is private
  with Preelaborable_Initialization;

function Value (Map      : in Character_Mapping;
               Element : in Character)
  return Character;

Identity : constant Character_Mapping;

function To_Mapping (From, To : in Character_Sequence)
  return Character_Mapping;

function To_Domain (Map : in Character_Mapping)
  return Character_Sequence;

function To_Range (Map : in Character_Mapping)
  return Character_Sequence;

type Character_Mapping_Function is
  access function (From : in Character) return Character;

private
  ... -- not specified by the language
end Ada.Strings.Maps;

```

An object of type `Character_Set` represents a set of characters.

`Null_Set` represents the set containing no characters.

An object `Obj` of type `Character_Range` represents the set of characters in the range `Obj.Low .. Obj.High`.

An object `Obj` of type `Character_Ranges` represents the union of the sets corresponding to `Obj(I)` for `I` in `Obj.Range`.

```

function To_Set (Ranges : in Character_Ranges) return Character_Set;

```

If `Ranges.Length=0` then `Null_Set` is returned; otherwise, the returned value represents the set corresponding to `Ranges`.

```

function To_Set (Span : in Character_Range) return Character_Set;

```

The returned value represents the set containing each character in `Span`.

```

function To_Ranges (Set : in Character_Set) return Character_Ranges;

```

If `Set = Null_Set`, then an empty `Character_Ranges` array is returned; otherwise, the shortest array of contiguous ranges of `Character` values in `Set`, in increasing order of `Low`, is returned.

```

function "=" (Left, Right : in Character_Set) return Boolean;

```

The function `"="` returns `True` if `Left` and `Right` represent identical sets, and `False` otherwise.

Each of the logical operators `"not"`, `"and"`, `"or"`, and `"xor"` returns a `Character_Set` value that represents the set obtained by applying the corresponding operation to the set(s) represented by the parameter(s) of the operator. `"-"`(`Left`, `Right`) is equivalent to `"and"`(`Left`, `"not"`(`Right`)).

```

function Is_In (Element : in Character;
               Set      : in Character_Set);
  return Boolean;

```

`Is_In` returns `True` if `Element` is in `Set`, and `False` otherwise.

```

function Is_Subset (Elements : in Character_Set;
                  Set       : in Character_Set)
  return Boolean;

```

`Is_Subset` returns `True` if `Elements` is a subset of `Set`, and `False` otherwise.

```

subtype Character_Sequence is String;

```

The `Character_Sequence` subtype is used to portray a set of character values and also to identify the domain and range of a character mapping.

```
function To_Set (Sequence : in Character_Sequence) return Character_Set;
```

```
function To_Set (Singleton : in Character) return Character_Set;
```

Sequence portrays the set of character values that it explicitly contains (ignoring duplicates). Singleton portrays the set comprising a single Character. Each of the To\_Set functions returns a Character\_Set value that represents the set portrayed by Sequence or Singleton.

```
function To_Sequence (Set : in Character_Set) return Character_Sequence;
```

The function To\_Sequence returns a Character\_Sequence value containing each of the characters in the set represented by Set, in ascending order with no duplicates.

```
type Character_Mapping is private;
```

An object of type Character\_Mapping represents a Character-to-Character mapping.

```
function Value (Map : in Character_Mapping;  
               Element : in Character)  
return Character;
```

The function Value returns the Character value to which Element maps with respect to the mapping represented by Map.

A character *C* *matches* a pattern character *P* with respect to a given Character\_Mapping value Map if Value(Map, C) = P. A string *S* *matches* a pattern string *P* with respect to a given Character\_Mapping if their lengths are the same and if each character in *S* matches its corresponding character in the pattern string *P*.

String handling subprograms that deal with character mappings have parameters whose type is Character\_Mapping.

```
Identity : constant Character_Mapping;
```

Identity maps each Character to itself.

```
function To_Mapping (From, To : in Character_Sequence)  
return Character_Mapping;
```

To\_Mapping produces a Character\_Mapping such that each element of From maps to the corresponding element of To, and each other character maps to itself. If From.Length /= To.Length, or if some character is repeated in From, then Translation\_Error is propagated.

```
function To_Domain (Map : in Character_Mapping) return Character_Sequence;
```

To\_Domain returns the shortest Character\_Sequence value D such that each character not in D maps to itself, and such that the characters in D are in ascending order. The lower bound of D is 1.

```
function To_Range (Map : in Character_Mapping) return Character_Sequence;
```

To\_Range returns the Character\_Sequence value R, such that if D = To\_Domain(Map), then R has the same bounds as D, and D(I) maps to R(I) for each I in D'Range.

An object F of type Character\_Mapping\_Function maps a Character value C to the Character value F.all(C), which is said to *match* C with respect to mapping function F.

NOTE 1 Character\_Mapping and Character\_Mapping\_Function are used both for character equivalence mappings in the search subprograms (such as for case insensitivity) and as transformational mappings in the Translate subprograms.

NOTE 2 To\_Domain(Identity) and To\_Range(Identity) each returns the null string.

EXAMPLE *Example of use of Strings.Maps.To\_Mapping:*

To\_Mapping("ABCD", "ZZAB") returns a Character\_Mapping that maps 'A' and 'B' to 'Z', 'C' to 'A', 'D' to 'B', and each other Character to itself.



### A.4.3 Fixed-Length String Handling

The language-defined package `Strings.Fixed` provides string-handling subprograms for fixed-length strings; that is, for values of type `Standard.String`. Several of these subprograms are procedures that modify the contents of a `String` that is passed as an **out** or an **in out** parameter; each has additional parameters to control the effect when the logical length of the result differs from the parameter's length.

For each function that returns a `String`, the lower bound of the returned value is 1.

The basic model embodied in the package is that a fixed-length string comprises significant characters and possibly padding (with space characters) on either or both ends. When a shorter string is copied to a longer string, padding is inserted, and when a longer string is copied to a shorter one, padding is stripped. The `Move` procedure in `Strings.Fixed`, which takes a `String` as an **out** parameter, allows the programmer to control these effects. Similar control is provided by the string transformation procedures.

#### Static Semantics

The library package `Strings.Fixed` has the following declaration:

```
with Ada.Strings.Maps;
package Ada.Strings.Fixed
  with Preelaborate, Nonblocking, Global => in out synchronized is
-- "Copy" procedure for strings of possibly different lengths
  procedure Move (Source : in String;
                  Target : out String;
                  Drop   : in Truncation := Error;
                  Justify : in Alignment := Left;
                  Pad    : in Character := Space);
-- Search subprograms
  function Index (Source : in String;
                  Pattern : in String;
                  From    : in Positive;
                  Going    : in Direction := Forward;
                  Mapping  : in Maps.Character_Mapping := Maps.Identity)
    return Natural;
  function Index (Source : in String;
                  Pattern : in String;
                  From    : in Positive;
                  Going    : in Direction := Forward;
                  Mapping  : in Maps.Character_Mapping_Function)
    return Natural;
  function Index (Source : in String;
                  Pattern : in String;
                  Going    : in Direction := Forward;
                  Mapping  : in Maps.Character_Mapping
                           := Maps.Identity)
    return Natural;
  function Index (Source : in String;
                  Pattern : in String;
                  Going    : in Direction := Forward;
                  Mapping  : in Maps.Character_Mapping_Function)
    return Natural;
  function Index (Source : in String;
                  Set     : in Maps.Character_Set;
                  From    : in Positive;
                  Test     : in Membership := Inside;
                  Going    : in Direction := Forward)
    return Natural;
```

```

function Index (Source : in String;
                Set      : in Maps.Character_Set;
                Test     : in Membership := Inside;
                Going    : in Direction  := Forward)
    return Natural;

function Index_Non_Blank (Source : in String;
                          From    : in Positive;
                          Going   : in Direction := Forward)
    return Natural;

function Index_Non_Blank (Source : in String;
                          Going   : in Direction := Forward)
    return Natural;

function Count (Source   : in String;
                 Pattern  : in String;
                 Mapping  : in Maps.Character_Mapping
                        := Maps.Identity)
    return Natural;

function Count (Source   : in String;
                 Pattern  : in String;
                 Mapping  : in Maps.Character_Mapping_Function)
    return Natural;

function Count (Source   : in String;
                 Set      : in Maps.Character_Set)
    return Natural;

procedure Find-Token (Source : in String;
                      Set      : in Maps.Character_Set;
                      From     : in Positive;
                      Test     : in Membership;
                      First    : out Positive;
                      Last     : out Natural);

procedure Find-Token (Source : in String;
                      Set      : in Maps.Character_Set;
                      Test     : in Membership;
                      First    : out Positive;
                      Last     : out Natural);

-- String translation subprograms
function Translate (Source : in String;
                   Mapping : in Maps.Character_Mapping)
    return String;

procedure Translate (Source : in out String;
                    Mapping : in Maps.Character_Mapping);

function Translate (Source : in String;
                   Mapping : in Maps.Character_Mapping_Function)
    return String;

procedure Translate (Source : in out String;
                    Mapping : in Maps.Character_Mapping_Function);

-- String transformation subprograms
function Replace_Slice (Source : in String;
                       Low      : in Positive;
                       High     : in Natural;
                       By       : in String)
    return String;

```

```

procedure Replace_Slice (Source   : in out String;
                        Low       : in Positive;
                        High      : in Natural;
                        By        : in String;
                        Drop      : in Truncation := Error;
                        Justify   : in Alignment  := Left;
                        Pad       : in Character  := Space);

function Insert (Source   : in String;
                Before    : in Positive;
                New_Item  : in String)
    return String;

procedure Insert (Source   : in out String;
                Before    : in Positive;
                New_Item  : in String;
                Drop      : in Truncation := Error);

function Overwrite (Source   : in String;
                   Position  : in Positive;
                   New_Item  : in String)
    return String;

procedure Overwrite (Source   : in out String;
                   Position  : in Positive;
                   New_Item  : in String;
                   Drop      : in Truncation := Right);

function Delete (Source   : in String;
                From      : in Positive;
                Through   : in Natural)
    return String;

procedure Delete (Source   : in out String;
                From      : in Positive;
                Through   : in Natural;
                Justify   : in Alignment := Left;
                Pad       : in Character := Space);

--String selector subprograms
function Trim (Source : in String;
              Side   : in Trim_End)
    return String;

procedure Trim (Source : in out String;
              Side     : in Trim_End;
              Justify  : in Alignment := Left;
              Pad      : in Character := Space);

function Trim (Source : in String;
              Left     : in Maps.Character_Set;
              Right    : in Maps.Character_Set)
    return String;

procedure Trim (Source : in out String;
              Left     : in Maps.Character_Set;
              Right    : in Maps.Character_Set;
              Justify  : in Alignment := Strings.Left;
              Pad      : in Character := Space);

function Head (Source : in String;
              Count   : in Natural;
              Pad     : in Character := Space)
    return String;

procedure Head (Source : in out String;
              Count   : in Natural;
              Justify : in Alignment := Left;
              Pad     : in Character := Space);

```

```

function Tail (Source : in String;
               Count  : in Natural;
               Pad    : in Character := Space)
    return String;

procedure Tail (Source  : in out String;
               Count    : in Natural;
               Justify   : in Alignment := Left;
               Pad       : in Character := Space);

--String constructor functions

function "*" (Left   : in Natural;
             Right  : in Character) return String;

function "*" (Left   : in Natural;
             Right  : in String) return String;

end Ada.Strings.Fixed;

```

The effects of the above subprograms are as follows.

```

procedure Move (Source  : in String;
               Target   : out String;
               Drop     : in Truncation := Error;
               Justify  : in Alignment  := Left;
               Pad      : in Character  := Space);

```

The Move procedure copies characters from Source to Target. If Source has the same length as Target, then the effect is to assign Source to Target. If Source is shorter than Target, then:

- If Justify=Left, then Source is copied into the first Source'Length characters of Target.
- If Justify=Right, then Source is copied into the last Source'Length characters of Target.
- If Justify=Center, then Source is copied into the middle Source'Length characters of Target. In this case, if the difference in length between Target and Source is odd, then the extra Pad character is on the right.
- Pad is copied to each Target character not otherwise assigned.

If Source is longer than Target, then the effect is based on Drop.

- If Drop=Left, then the rightmost Target'Length characters of Source are copied into Target.
- If Drop=Right, then the leftmost Target'Length characters of Source are copied into Target.
- If Drop=Error, then the effect depends on the value of the Justify parameter and also on whether any characters in Source other than Pad would fail to be copied:
  - If Justify=Left, and if each of the rightmost Source'Length-Target'Length characters in Source is Pad, then the leftmost Target'Length characters of Source are copied to Target.
  - If Justify=Right, and if each of the leftmost Source'Length-Target'Length characters in Source is Pad, then the rightmost Target'Length characters of Source are copied to Target.
  - Otherwise, Length\_Error is propagated.

```

function Index (Source : in String;
                Pattern : in String;
                From    : in Positive;
                Going   : in Direction := Forward;
                Mapping  : in Maps.Character_Mapping := Maps.Identity)
return Natural;

```

```

function Index (Source : in String;
                Pattern : in String;
                From    : in Positive;
                Going   : in Direction := Forward;
                Mapping  : in Maps.Character_Mapping_Function)
return Natural;

```

Each Index function searches, starting from From, for a slice of Source, with length Pattern'Length, that matches Pattern with respect to Mapping; the parameter Going indicates the direction of the lookup. If Source is the null string, Index returns 0; otherwise, if From is not in Source'Range, then Index\_Error is propagated. If Going = Forward, then Index returns the smallest index I which is greater than or equal to From such that the slice of Source starting at I matches Pattern. If Going = Backward, then Index returns the largest index I such that the slice of Source starting at I matches Pattern and has an upper bound less than or equal to From. If there is no such slice, then 0 is returned. If Pattern is the null string, then Pattern\_Error is propagated.

```

function Index (Source : in String;
                Pattern : in String;
                Going   : in Direction := Forward;
                Mapping  : in Maps.Character_Mapping
                        := Maps.Identity)
return Natural;

```

```

function Index (Source : in String;
                Pattern : in String;
                Going   : in Direction := Forward;
                Mapping  : in Maps.Character_Mapping_Function)
return Natural;

```

If Going = Forward, returns

```
Index (Source, Pattern, Source'First, Forward, Mapping);
```

otherwise, returns

```
Index (Source, Pattern, Source'Last, Backward, Mapping);
```

```

function Index (Source : in String;
                Set      : in Maps.Character_Set;
                From     : in Positive;
                Test      : in Membership := Inside;
                Going     : in Direction := Forward)
return Natural;

```

Index searches for the first or last occurrence of any of a set of characters (when Test=Inside), or any of the complement of a set of characters (when Test=Outside). If Source is the null string, Index returns 0; otherwise, if From is not in Source'Range, then Index\_Error is propagated. Otherwise, it returns the smallest index I >= From (if Going=Forward) or the largest index I <= From (if Going=Backward) such that Source(I) satisfies the Test condition with respect to Set; it returns 0 if there is no such Character in Source.

```

function Index (Source : in String;
                Set      : in Maps.Character_Set;
                Test      : in Membership := Inside;
                Going     : in Direction := Forward)
return Natural;

```

If Going = Forward, returns

```

    Index (Source, Set, Source'First, Test, Forward);
otherwise, returns
    Index (Source, Set, Source'Last, Test, Backward);
function Index_Non_Blank (Source : in String;
                          From   : in Positive;
                          Going  : in Direction := Forward)
    return Natural;
    Returns Index (Source, Maps.To_Set(Space), From, Outside, Going);
function Index_Non_Blank (Source : in String;
                          Going   : in Direction := Forward)
    return Natural;
    Returns Index(Source, Maps.To_Set(Space), Outside, Going)

function Count (Source   : in String;
                 Pattern  : in String;
                 Mapping  : in Maps.Character_Mapping
                        := Maps.Identity)
    return Natural;

function Count (Source   : in String;
                 Pattern  : in String;
                 Mapping  : in Maps.Character_Mapping_Function)
    return Natural;
    Returns the maximum number of nonoverlapping slices of Source that match Pattern with
    respect to Mapping. If Pattern is the null string then Pattern_Error is propagated.

function Count (Source   : in String;
                 Set       : in Maps.Character_Set)
    return Natural;
    Returns the number of occurrences in Source of characters that are in Set.

procedure Find-Token (Source : in String;
                      Set     : in Maps.Character_Set;
                      From    : in Positive;
                      Test    : in Membership;
                      First   : out Positive;
                      Last    : out Natural);

    If Source is not the null string and From is not in Source'Range, then Index_Error is raised.
    Otherwise, First is set to the index of the first character in Source(From .. Source'Last) that
    satisfies the Test condition. Last is set to the largest index such that all characters in
    Source(First .. Last) satisfy the Test condition. If no characters in Source(From .. Source'Last)
    satisfy the Test condition, First is set to From, and Last is set to 0.

procedure Find-Token (Source : in String;
                      Set     : in Maps.Character_Set;
                      Test    : in Membership;
                      First   : out Positive;
                      Last    : out Natural);

    Equivalent to Find-Token (Source, Set, Source'First, Test, First, Last).

```



```
function Translate (Source : in String;
                   Mapping : in Maps.Character_Mapping)
return String;
```

```
function Translate (Source : in String;
                   Mapping : in Maps.Character_Mapping_Function)
return String;
```

Returns the string S whose length is Source'Length and such that S(I) is the character to which Mapping maps the corresponding element of Source, for I in 1..Source'Length.

```
procedure Translate (Source : in out String;
                   Mapping : in Maps.Character_Mapping);
```

```
procedure Translate (Source : in out String;
                   Mapping : in Maps.Character_Mapping_Function);
```

Equivalent to Source := Translate(Source, Mapping).

```
function Replace_Slice (Source : in String;
                       Low      : in Positive;
                       High     : in Natural;
                       By       : in String)
return String;
```

If Low > Source'Last+1, or High < Source'First-1, then Index\_Error is propagated. Otherwise:

- If High >= Low, then the returned string comprises Source(Source'First..Low-1) & By & Source(High+1..Source'Last), but with lower bound 1.
- If High < Low, then the returned string is Insert(Source, Before=>Low, New\_Item=>By).

```
procedure Replace_Slice (Source : in out String;
                       Low      : in Positive;
                       High     : in Natural;
                       By       : in String;
                       Drop     : in Truncation := Error;
                       Justify  : in Alignment  := Left;
                       Pad     : in Character  := Space);
```

Equivalent to Move(Replace\_Slice(Source, Low, High, By), Source, Drop, Justify, Pad).

```
function Insert (Source : in String;
                Before  : in Positive;
                New_Item : in String)
return String;
```

Propagates Index\_Error if Before is not in Source'First .. Source'Last+1; otherwise, returns Source(Source'First..Before-1) & New\_Item & Source(Before..Source'Last), but with lower bound 1.

```
procedure Insert (Source : in out String;
                 Before  : in Positive;
                 New_Item : in String;
                 Drop    : in Truncation := Error);
```

Equivalent to Move(Insert(Source, Before, New\_Item), Source, Drop).

```
function Overwrite (Source : in String;
                   Position : in Positive;
                   New_Item : in String)
return String;
```

Propagates Index\_Error if Position is not in Source'First .. Source'Last+1; otherwise, returns the string obtained from Source by consecutively replacing characters starting at Position with corresponding characters from New\_Item. If the end of Source is reached before the characters in New\_Item are exhausted, the remaining characters from New\_Item are appended to the string.

```

procedure Overwrite (Source   : in out String;
                     Position : in Positive;
                     New_Item : in String;
                     Drop     : in Truncation := Right);

```

Equivalent to Move(Overwrite(Source, Position, New\_Item), Source, Drop).

```

function Delete (Source : in String;
                 From    : in Positive;
                 Through : in Natural)
return String;

```

If From <= Through, the returned string is Replace\_Slice(Source, From, Through, ""); otherwise, it is Source with lower bound 1.

```

procedure Delete (Source : in out String;
                 From    : in Positive;
                 Through : in Natural;
                 Justify : in Alignment := Left;
                 Pad     : in Character := Space);

```

Equivalent to Move(Delete(Source, From, Through), Source, Justify=>Justify, Pad=>Pad).

```

function Trim (Source : in String;
               Side    : in Trim_End)
return String;

```

Returns the string obtained by removing from Source all leading Space characters (if Side = Left), all trailing Space characters (if Side = Right), or all leading and trailing Space characters (if Side = Both).

```

procedure Trim (Source : in out String;
               Side    : in Trim_End;
               Justify : in Alignment := Left;
               Pad     : in Character := Space);

```

Equivalent to Move(Trim(Source, Side), Source, Justify=>Justify, Pad=>Pad).

```

function Trim (Source : in String;
               Left    : in Maps.Character_Set;
               Right   : in Maps.Character_Set)
return String;

```

Returns the string obtained by removing from Source all leading characters in Left and all trailing characters in Right.

```

procedure Trim (Source : in out String;
               Left    : in Maps.Character_Set;
               Right   : in Maps.Character_Set;
               Justify : in Alignment := Strings.Left;
               Pad     : in Character := Space);

```

Equivalent to Move(Trim(Source, Left, Right), Source, Justify=>Justify, Pad=>Pad).

```

function Head (Source : in String;
               Count   : in Natural;
               Pad     : in Character := Space)
return String;

```

Returns a string of length Count. If Count <= Source'Length, the string comprises the first Count characters of Source. Otherwise, its contents are Source concatenated with Count-Source'Length Pad characters.

```

procedure Head (Source : in out String;
                 Count  : in Natural;
                 Justify : in Alignment := Left;
                 Pad     : in Character := Space);

```

Equivalent to Move(Head(Source, Count, Pad), Source, Drop=>Error, Justify=>Justify, Pad=>Pad).

```

function Tail (Source : in String;
                Count  : in Natural;
                Pad     : in Character := Space)
return String;

```

Returns a string of length Count. If Count <= Source'Length, the string comprises the last Count characters of Source. Otherwise, its contents are Count-Source'Length Pad characters concatenated with Source.

```

procedure Tail (Source : in out String;
                 Count  : in Natural;
                 Justify : in Alignment := Left;
                 Pad     : in Character := Space);

```

Equivalent to Move(Tail(Source, Count, Pad), Source, Drop=>Error, Justify=>Justify, Pad=>Pad).

```

function "*" (Left  : in Natural;
              Right : in Character) return String;

```

```

function "*" (Left  : in Natural;
              Right : in String) return String;

```

These functions replicate a character or string a specified number of times. The first function returns a string whose length is Left and each of whose elements is Right. The second function returns a string whose length is Left\*Right'Length and whose value is the null string if Left = 0 and otherwise is (Left-1)\*Right & Right with lower bound 1.

NOTE 1 In the Index and Count functions taking Pattern and Mapping parameters, for there to be a match, the actual String parameter passed to Pattern can contain only characters occurring as target characters of the mapping.

NOTE 2 In the Insert subprograms, inserting at the end of a string is obtained by passing Source'Last+1 as the Before parameter.

NOTE 3 If a null Character\_Mapping\_Function is passed to any of the string handling subprograms, Constraint\_Error is propagated.

#### A.4.4 Bounded-Length String Handling

The language-defined package Strings.Bounded provides a generic package each of whose instances yields a private type Bounded\_String and a set of operations. An object of a particular Bounded\_String type represents a String whose low bound is 1 and whose length can vary conceptually between 0 and a maximum size established at the generic instantiation. The subprograms for fixed-length string handling are either overloaded directly for Bounded\_String, or are modified as necessary to reflect the variability in length. Additionally, since the Bounded\_String type is private, appropriate constructor and selector operations are provided.

##### Static Semantics

The library package Strings.Bounded has the following declaration:

```

with Ada.Strings.Maps;
package Ada.Strings.Bounded
  with Preelaborate, Nonblocking, Global => in out synchronized is

  generic
    Max : Positive;    -- Maximum length of a Bounded_String
  package Generic_Bounded_Length is

```

```

Max_Length : constant Positive := Max;
type Bounded_String is private;
Null_Bounded_String : constant Bounded_String;
subtype Length_Range is Natural range 0 .. Max_Length;
function Length (Source : in Bounded_String) return Length_Range;
-- Conversion, Concatenation, and Selection functions
function To_Bounded_String (Source : in String;
                           Drop   : in Truncation := Error)
    return Bounded_String;
function To_String (Source : in Bounded_String) return String;
procedure Set_Bounded_String
    (Target : out Bounded_String;
     Source : in String;
     Drop   : in Truncation := Error);
function Append (Left, Right : in Bounded_String;
                 Drop         : in Truncation := Error)
    return Bounded_String;
function Append (Left  : in Bounded_String;
                 Right  : in String;
                 Drop   : in Truncation := Error)
    return Bounded_String;
function Append (Left  : in String;
                 Right  : in Bounded_String;
                 Drop   : in Truncation := Error)
    return Bounded_String;
function Append (Left  : in Bounded_String;
                 Right  : in Character;
                 Drop   : in Truncation := Error)
    return Bounded_String;
function Append (Left  : in Character;
                 Right  : in Bounded_String;
                 Drop   : in Truncation := Error)
    return Bounded_String;
procedure Append (Source : in out Bounded_String;
                  New_Item : in Bounded_String;
                  Drop     : in Truncation := Error);
procedure Append (Source : in out Bounded_String;
                  New_Item : in String;
                  Drop     : in Truncation := Error);
procedure Append (Source : in out Bounded_String;
                  New_Item : in Character;
                  Drop     : in Truncation := Error);
function "&" (Left, Right : in Bounded_String)
    return Bounded_String;
function "&" (Left : in Bounded_String; Right : in String)
    return Bounded_String;
function "&" (Left : in String; Right : in Bounded_String)
    return Bounded_String;
function "&" (Left : in Bounded_String; Right : in Character)
    return Bounded_String;
function "&" (Left : in Character; Right : in Bounded_String)
    return Bounded_String;

```

```

function Element (Source : in Bounded_String;
                  Index  : in Positive)
    return Character;

procedure Replace_Element (Source : in out Bounded_String;
                          Index  : in Positive;
                          By      : in Character);

function Slice (Source : in Bounded_String;
               Low      : in Positive;
               High      : in Natural)
    return String;

function Bounded_Slice
    (Source : in Bounded_String;
     Low     : in Positive;
     High     : in Natural)
    return Bounded_String;

procedure Bounded_Slice
    (Source : in Bounded_String;
     Target : out Bounded_String;
     Low     : in Positive;
     High     : in Natural);

function "=" (Left, Right : in Bounded_String) return Boolean;
function "=" (Left : in Bounded_String; Right : in String)
    return Boolean;

function "=" (Left : in String; Right : in Bounded_String)
    return Boolean;

function "<" (Left, Right : in Bounded_String) return Boolean;
function "<" (Left : in Bounded_String; Right : in String)
    return Boolean;

function "<" (Left : in String; Right : in Bounded_String)
    return Boolean;

function "<=" (Left, Right : in Bounded_String) return Boolean;
function "<=" (Left : in Bounded_String; Right : in String)
    return Boolean;

function "<=" (Left : in String; Right : in Bounded_String)
    return Boolean;

function ">" (Left, Right : in Bounded_String) return Boolean;
function ">" (Left : in Bounded_String; Right : in String)
    return Boolean;

function ">" (Left : in String; Right : in Bounded_String)
    return Boolean;

function ">=" (Left, Right : in Bounded_String) return Boolean;
function ">=" (Left : in Bounded_String; Right : in String)
    return Boolean;

function ">=" (Left : in String; Right : in Bounded_String)
    return Boolean;

-- Search subprograms

function Index (Source : in Bounded_String;
               Pattern : in String;
               From     : in Positive;
               Going     : in Direction := Forward;
               Mapping    : in Maps.Character_Mapping := Maps.Identity)
    return Natural;

```

```

function Index (Source : in Bounded_String;
                Pattern : in String;
                From    : in Positive;
                Going   : in Direction := Forward;
                Mapping  : in Maps.Character_Mapping_Function)
    return Natural;

function Index (Source : in Bounded_String;
                Pattern : in String;
                Going   : in Direction := Forward;
                Mapping  : in Maps.Character_Mapping
                        := Maps.Identity)
    return Natural;

function Index (Source : in Bounded_String;
                Pattern : in String;
                Going   : in Direction := Forward;
                Mapping  : in Maps.Character_Mapping_Function)
    return Natural;

function Index (Source : in Bounded_String;
                Set      : in Maps.Character_Set;
                From     : in Positive;
                Test      : in Membership := Inside;
                Going    : in Direction := Forward)
    return Natural;

function Index (Source : in Bounded_String;
                Set      : in Maps.Character_Set;
                Test      : in Membership := Inside;
                Going    : in Direction := Forward)
    return Natural;

function Index_Non_Blank (Source : in Bounded_String;
                        From      : in Positive;
                        Going     : in Direction := Forward)
    return Natural;

function Index_Non_Blank (Source : in Bounded_String;
                        Going     : in Direction := Forward)
    return Natural;

function Count (Source : in Bounded_String;
                Pattern : in String;
                Mapping  : in Maps.Character_Mapping
                        := Maps.Identity)
    return Natural;

function Count (Source : in Bounded_String;
                Pattern : in String;
                Mapping  : in Maps.Character_Mapping_Function)
    return Natural;

function Count (Source : in Bounded_String;
                Set      : in Maps.Character_Set)
    return Natural;

procedure Find-Token (Source : in Bounded_String;
                    Set      : in Maps.Character_Set;
                    From     : in Positive;
                    Test      : in Membership;
                    First     : out Positive;
                    Last      : out Natural);

procedure Find-Token (Source : in Bounded_String;
                    Set      : in Maps.Character_Set;
                    Test      : in Membership;
                    First     : out Positive;
                    Last      : out Natural);

```



```

-- String translation subprograms
function Translate (Source : in Bounded_String;
                   Mapping : in Maps.Character_Mapping)
    return Bounded_String;
procedure Translate (Source : in out Bounded_String;
                   Mapping : in Maps.Character_Mapping);

function Translate (Source : in Bounded_String;
                   Mapping : in Maps.Character_Mapping_Function)
    return Bounded_String;
procedure Translate (Source : in out Bounded_String;
                   Mapping : in Maps.Character_Mapping_Function);

-- String transformation subprograms
function Replace_Slice (Source : in Bounded_String;
                      Low : in Positive;
                      High : in Natural;
                      By : in String;
                      Drop : in Truncation := Error)
    return Bounded_String;
procedure Replace_Slice (Source : in out Bounded_String;
                      Low : in Positive;
                      High : in Natural;
                      By : in String;
                      Drop : in Truncation := Error);

function Insert (Source : in Bounded_String;
                Before : in Positive;
                New_Item : in String;
                Drop : in Truncation := Error)
    return Bounded_String;
procedure Insert (Source : in out Bounded_String;
                Before : in Positive;
                New_Item : in String;
                Drop : in Truncation := Error);

function Overwrite (Source : in Bounded_String;
                   Position : in Positive;
                   New_Item : in String;
                   Drop : in Truncation := Error)
    return Bounded_String;
procedure Overwrite (Source : in out Bounded_String;
                   Position : in Positive;
                   New_Item : in String;
                   Drop : in Truncation := Error);

function Delete (Source : in Bounded_String;
                From : in Positive;
                Through : in Natural)
    return Bounded_String;
procedure Delete (Source : in out Bounded_String;
                From : in Positive;
                Through : in Natural);

--String selector subprograms
function Trim (Source : in Bounded_String;
              Side : in Trim_End)
    return Bounded_String;
procedure Trim (Source : in out Bounded_String;
              Side : in Trim_End);

```

```

function Trim (Source : in Bounded_String;
               Left   : in Maps.Character_Set;
               Right  : in Maps.Character_Set)
    return Bounded_String;

procedure Trim (Source : in out Bounded_String;
               Left   : in Maps.Character_Set;
               Right  : in Maps.Character_Set);

function Head (Source : in Bounded_String;
               Count  : in Natural;
               Pad    : in Character := Space;
               Drop   : in Truncation := Error)
    return Bounded_String;

procedure Head (Source : in out Bounded_String;
               Count  : in Natural;
               Pad    : in Character := Space;
               Drop   : in Truncation := Error);

function Tail (Source : in Bounded_String;
               Count  : in Natural;
               Pad    : in Character := Space;
               Drop   : in Truncation := Error)
    return Bounded_String;

procedure Tail (Source : in out Bounded_String;
               Count  : in Natural;
               Pad    : in Character := Space;
               Drop   : in Truncation := Error);

--String constructor subprograms

function "*" (Left  : in Natural;
              Right : in Character)
    return Bounded_String;

function "*" (Left  : in Natural;
              Right : in String)
    return Bounded_String;

function "*" (Left  : in Natural;
              Right : in Bounded_String)
    return Bounded_String;

function Replicate (Count : in Natural;
                   Item  : in Character;
                   Drop   : in Truncation := Error)
    return Bounded_String;

function Replicate (Count : in Natural;
                   Item  : in String;
                   Drop   : in Truncation := Error)
    return Bounded_String;

function Replicate (Count : in Natural;
                   Item  : in Bounded_String;
                   Drop   : in Truncation := Error)
    return Bounded_String;

private
    ... -- not specified by the language
end Generic_Bounded_Length;

end Ada.Strings.Bounded;

```

Null\_Bounded\_String represents the null string. If an object of type Bounded\_String is not otherwise initialized, it will be initialized to the same value as Null\_Bounded\_String.

```
function Length (Source : in Bounded_String) return Length_Range;
```

The Length function returns the length of the string represented by Source.

```

function To_Bounded_String (Source : in String;
                             Drop    : in Truncation := Error)
    return Bounded_String;

```

If Source.Length <= Max\_Length, then this function returns a Bounded\_String that represents Source. Otherwise, the effect depends on the value of Drop:

- If Drop=Left, then the result is a Bounded\_String that represents the string comprising the rightmost Max\_Length characters of Source.
- If Drop=Right, then the result is a Bounded\_String that represents the string comprising the leftmost Max\_Length characters of Source.
- If Drop=Error, then Strings.Length\_Error is propagated.

```

function To_String (Source : in Bounded_String) return String;

```

To\_String returns the String value with lower bound 1 represented by Source. If B is a Bounded\_String, then B = To\_Bounded\_String(To\_String(B)).

```

procedure Set_Bounded_String
    (Target : out Bounded_String;
     Source : in      String;
     Drop   : in      Truncation := Error);

```

Equivalent to Target := To\_Bounded\_String (Source, Drop);

Each of the Append functions returns a Bounded\_String obtained by concatenating the string or character given or represented by one of the parameters, with the string or character given or represented by the other parameter, and applying To\_Bounded\_String to the concatenation result string, with Drop as provided to the Append function.

Each of the procedures Append(Source, New\_Item, Drop) has the same effect as the corresponding assignment Source := Append(Source, New\_Item, Drop).

Each of the "&" functions has the same effect as the corresponding Append function, with Error as the Drop parameter.

```

function Element (Source : in Bounded_String;
                  Index   : in Positive)
    return Character;

```

Returns the character at position Index in the string represented by Source; propagates Index\_Error if Index > Length(Source).

```

procedure Replace_Element (Source : in out Bounded_String;
                           Index   : in Positive;
                           By      : in Character);

```

Updates Source such that the character at position Index in the string represented by Source is By; propagates Index\_Error if Index > Length(Source).

```

function Slice (Source : in Bounded_String;
                Low     : in Positive;
                High    : in Natural)
    return String;

```

Returns the slice at positions Low through High in the string represented by Source; propagates Index\_Error if Low > Length(Source)+1 or High > Length(Source). The bounds of the returned string are Low and High.

```

function Bounded_Slice
  (Source : in Bounded_String;
   Low    : in Positive;
   High   : in Natural)
  return Bounded_String;

```

Returns the slice at positions Low through High in the string represented by Source as a bounded string; propagates Index\_Error if Low > Length(Source)+1 or High > Length(Source).

```

procedure Bounded_Slice
  (Source : in      Bounded_String;
   Target : out    Bounded_String;
   Low    : in      Positive;
   High   : in      Natural);

```

Equivalent to Target := Bounded\_Slice (Source, Low, High);

Each of the functions "=", "<", ">", "<=", and ">=" returns the same result as the corresponding String operation applied to the String values given or represented by the two parameters.

Each of the search subprograms (Index, Index\_Non\_Blank, Count, Find Token) has the same effect as the corresponding subprogram in Strings.Fixed applied to the string represented by the Bounded\_String parameter.

Each of the Translate subprograms, when applied to a Bounded\_String, has an analogous effect to the corresponding subprogram in Strings.Fixed. For the Translate function, the translation is applied to the string represented by the Bounded\_String parameter, and the result is converted (via To\_Bounded\_String) to a Bounded\_String. For the Translate procedure, the string represented by the Bounded\_String parameter after the translation is given by the Translate function for fixed-length strings applied to the string represented by the original value of the parameter.

Each of the transformation subprograms (Replace\_Slice, Insert, Overwrite, Delete), selector subprograms (Trim, Head, Tail), and constructor functions ("\*") has an effect based on its corresponding subprogram in Strings.Fixed, and Replicate is based on Fixed.\*". In the case of a function, the corresponding fixed-length string subprogram is applied to the string represented by the Bounded\_String parameter. To\_Bounded\_String is applied the result string, with Drop (or Error in the case of Generic\_Bounded\_Length.\*") determining the effect when the string length exceeds Max\_Length. In the case of a procedure, the corresponding function in Strings.Bounded.Generic\_Bounded\_Length is applied, with the result assigned into the Source parameter.

#### *Implementation Advice*

Bounded string objects should not be implemented by implicit pointers and dynamic allocation.

### **A.4.5 Unbounded-Length String Handling**

The language-defined package Strings.Unbounded provides a private type Unbounded\_String and a set of operations. An object of type Unbounded\_String represents a String whose low bound is 1 and whose length can vary conceptually between 0 and Natural'Last. The subprograms for fixed-length string handling are either overloaded directly for Unbounded\_String, or are modified as necessary to reflect the flexibility in length. Since the Unbounded\_String type is private, relevant constructor and selector operations are provided.

#### *Static Semantics*

The library package Strings.Unbounded has the following declaration:

```

with Ada.Strings.Maps;
package Ada.Strings.Unbounded
  with Preelaborate, Nonblocking, Global => in out synchronized is
    type Unbounded_String is private
      with Preelaborable_Initialization;
    Null_Unbounded_String : constant Unbounded_String;

```

```

function Length (Source : in Unbounded_String) return Natural;
type String_Access is access all String;
procedure Free (X : in out String_Access);
-- Conversion, Concatenation, and Selection functions
function To_Unbounded_String (Source : in String)
    return Unbounded_String;
function To_Unbounded_String (Length : in Natural)
    return Unbounded_String;
function To_String (Source : in Unbounded_String) return String;
procedure Set_Unbounded_String
    (Target : out Unbounded_String;
     Source : in String);
procedure Append (Source : in out Unbounded_String;
                  New_Item : in Unbounded_String);
procedure Append (Source : in out Unbounded_String;
                  New_Item : in String);
procedure Append (Source : in out Unbounded_String;
                  New_Item : in Character);
function "&" (Left, Right : in Unbounded_String)
    return Unbounded_String;
function "&" (Left : in Unbounded_String; Right : in String)
    return Unbounded_String;
function "&" (Left : in String; Right : in Unbounded_String)
    return Unbounded_String;
function "&" (Left : in Unbounded_String; Right : in Character)
    return Unbounded_String;
function "&" (Left : in Character; Right : in Unbounded_String)
    return Unbounded_String;
function Element (Source : in Unbounded_String;
                  Index : in Positive)
    return Character;
procedure Replace_Element (Source : in out Unbounded_String;
                           Index : in Positive;
                           By : in Character);
function Slice (Source : in Unbounded_String;
                 Low : in Positive;
                 High : in Natural)
    return String;
function Unbounded_Slice
    (Source : in Unbounded_String;
     Low : in Positive;
     High : in Natural)
    return Unbounded_String;
procedure Unbounded_Slice
    (Source : in Unbounded_String;
     Target : out Unbounded_String;
     Low : in Positive;
     High : in Natural);
function "=" (Left, Right : in Unbounded_String) return Boolean;
function "=" (Left : in Unbounded_String; Right : in String)
    return Boolean;
function "=" (Left : in String; Right : in Unbounded_String)
    return Boolean;

```

```

function "<" (Left, Right : in Unbounded_String) return Boolean;
function "<" (Left : in Unbounded_String; Right : in String)
  return Boolean;
function "<" (Left : in String; Right : in Unbounded_String)
  return Boolean;
function "<=" (Left, Right : in Unbounded_String) return Boolean;
function "<=" (Left : in Unbounded_String; Right : in String)
  return Boolean;
function "<=" (Left : in String; Right : in Unbounded_String)
  return Boolean;
function ">" (Left, Right : in Unbounded_String) return Boolean;
function ">" (Left : in Unbounded_String; Right : in String)
  return Boolean;
function ">" (Left : in String; Right : in Unbounded_String)
  return Boolean;
function ">=" (Left, Right : in Unbounded_String) return Boolean;
function ">=" (Left : in Unbounded_String; Right : in String)
  return Boolean;
function ">=" (Left : in String; Right : in Unbounded_String)
  return Boolean;
-- Search subprograms
function Index (Source : in Unbounded_String;
  Pattern : in String;
  From : in Positive;
  Going : in Direction := Forward;
  Mapping : in Maps.Character_Mapping := Maps.Identity)
  return Natural;
function Index (Source : in Unbounded_String;
  Pattern : in String;
  From : in Positive;
  Going : in Direction := Forward;
  Mapping : in Maps.Character_Mapping_Function)
  return Natural;
function Index (Source : in Unbounded_String;
  Pattern : in String;
  Going : in Direction := Forward;
  Mapping : in Maps.Character_Mapping
    := Maps.Identity)
  return Natural;
function Index (Source : in Unbounded_String;
  Pattern : in String;
  Going : in Direction := Forward;
  Mapping : in Maps.Character_Mapping_Function)
  return Natural;
function Index (Source : in Unbounded_String;
  Set : in Maps.Character_Set;
  From : in Positive;
  Test : in Membership := Inside;
  Going : in Direction := Forward)
  return Natural;
function Index (Source : in Unbounded_String;
  Set : in Maps.Character_Set;
  Test : in Membership := Inside;
  Going : in Direction := Forward) return Natural;

```



```

function Index_Non_Blank (Source : in Unbounded_String;
                          From   : in Positive;
                          Going  : in Direction := Forward)
    return Natural;

function Index_Non_Blank (Source : in Unbounded_String;
                          Going   : in Direction := Forward)
    return Natural;

function Count (Source   : in Unbounded_String;
                Pattern   : in String;
                Mapping    : in Maps.Character_Mapping
                          := Maps.Identity)
    return Natural;

function Count (Source   : in Unbounded_String;
                Pattern   : in String;
                Mapping    : in Maps.Character_Mapping_Function)
    return Natural;

function Count (Source   : in Unbounded_String;
                Set        : in Maps.Character_Set)
    return Natural;

procedure Find-Token (Source : in Unbounded_String;
                     Set      : in Maps.Character_Set;
                     From     : in Positive;
                     Test     : in Membership;
                     First    : out Positive;
                     Last     : out Natural);

procedure Find-Token (Source : in Unbounded_String;
                     Set      : in Maps.Character_Set;
                     Test     : in Membership;
                     First    : out Positive;
                     Last     : out Natural);

-- String translation subprograms
function Translate (Source : in Unbounded_String;
                   Mapping  : in Maps.Character_Mapping)
    return Unbounded_String;

procedure Translate (Source : in out Unbounded_String;
                   Mapping  : in Maps.Character_Mapping);

function Translate (Source : in Unbounded_String;
                   Mapping  : in Maps.Character_Mapping_Function)
    return Unbounded_String;

procedure Translate (Source : in out Unbounded_String;
                   Mapping  : in Maps.Character_Mapping_Function);

-- String transformation subprograms
function Replace_Slice (Source   : in Unbounded_String;
                      Low        : in Positive;
                      High       : in Natural;
                      By         : in String)
    return Unbounded_String;

procedure Replace_Slice (Source   : in out Unbounded_String;
                      Low        : in Positive;
                      High       : in Natural;
                      By         : in String);

function Insert (Source   : in Unbounded_String;
                Before    : in Positive;
                New_Item   : in String)
    return Unbounded_String;

```

```

procedure Insert (Source   : in out Unbounded_String;
                  Before   : in Positive;
                  New_Item  : in String);

function Overwrite (Source   : in Unbounded_String;
                    Position  : in Positive;
                    New_Item  : in String)
    return Unbounded_String;

procedure Overwrite (Source   : in out Unbounded_String;
                    Position  : in Positive;
                    New_Item  : in String);

function Delete (Source   : in Unbounded_String;
                 From      : in Positive;
                 Through   : in Natural)
    return Unbounded_String;

procedure Delete (Source   : in out Unbounded_String;
                 From      : in Positive;
                 Through   : in Natural);

function Trim (Source : in Unbounded_String;
               Side   : in Trim_End)
    return Unbounded_String;

procedure Trim (Source : in out Unbounded_String;
               Side   : in Trim_End);

function Trim (Source : in Unbounded_String;
               Left    : in Maps.Character_Set;
               Right   : in Maps.Character_Set)
    return Unbounded_String;

procedure Trim (Source : in out Unbounded_String;
               Left    : in Maps.Character_Set;
               Right   : in Maps.Character_Set);

function Head (Source : in Unbounded_String;
               Count   : in Natural;
               Pad     : in Character := Space)
    return Unbounded_String;

procedure Head (Source : in out Unbounded_String;
               Count   : in Natural;
               Pad     : in Character := Space);

function Tail (Source : in Unbounded_String;
               Count   : in Natural;
               Pad     : in Character := Space)
    return Unbounded_String;

procedure Tail (Source : in out Unbounded_String;
               Count   : in Natural;
               Pad     : in Character := Space);

function "*" (Left  : in Natural;
              Right : in Character)
    return Unbounded_String;

function "*" (Left  : in Natural;
              Right : in String)
    return Unbounded_String;

function "*" (Left  : in Natural;
              Right : in Unbounded_String)
    return Unbounded_String;

private
    ... -- not specified by the language
end Ada.Strings.Unbounded;

```

The type Unbounded\_String needs finalization (see 10.6).

Null\_Unbounded\_String represents the null String. If an object of type Unbounded\_String is not otherwise initialized, it will be initialized to the same value as Null\_Unbounded\_String.

The function Length returns the length of the String represented by Source.

The type String\_Access provides a (nonprivate) access type for explicit processing of unbounded-length strings. The procedure Free performs an unchecked deallocation of an object of type String\_Access.

The function To\_Unbounded\_String(Source : in String) returns an Unbounded\_String that represents Source. The function To\_Unbounded\_String(Length : in Natural) returns an Unbounded\_String that represents an uninitialized String whose length is Length.

The function To\_String returns the String with lower bound 1 represented by Source. To\_String and To\_Unbounded\_String are related as follows:

- If S is a String, then To\_String(To\_Unbounded\_String(S)) = S.
- If U is an Unbounded\_String, then To\_Unbounded\_String(To\_String(U)) = U.

The procedure Set\_Unbounded\_String sets Target to an Unbounded\_String that represents Source.

For each of the Append procedures, the resulting string represented by the Source parameter is given by the concatenation of the original value of Source and the value of New\_Item.

Each of the "&" functions returns an Unbounded\_String obtained by concatenating the string or character given or represented by one of the parameters, with the string or character given or represented by the other parameter, and applying To\_Unbounded\_String to the concatenation result string.

The Element, Replace\_Element, and Slice subprograms have the same effect as the corresponding bounded-length string subprograms.

The function Unbounded\_Slice returns the slice at positions Low through High in the string represented by Source as an Unbounded\_String. The procedure Unbounded\_Slice sets Target to the Unbounded\_String representing the slice at positions Low through High in the string represented by Source. Both subprograms propagate Index\_Error if Low > Length(Source)+1 or High > Length(Source).

Each of the functions "=", "<", ">", "<=", and ">=" returns the same result as the corresponding String operation applied to the String values given or represented by Left and Right.

Each of the search subprograms (Index, Index\_Non\_Blank, Count, Find-Token) has the same effect as the corresponding subprogram in Strings.Fixed applied to the string represented by the Unbounded\_String parameter.

The Translate function has an analogous effect to the corresponding subprogram in Strings.Fixed. The translation is applied to the string represented by the Unbounded\_String parameter, and the result is converted (via To\_Unbounded\_String) to an Unbounded\_String.

Each of the transformation functions (Replace\_Slice, Insert, Overwrite, Delete), selector functions (Trim, Head, Tail), and constructor functions ("\*") is likewise analogous to its corresponding subprogram in Strings.Fixed. For each of the subprograms, the corresponding fixed-length string subprogram is applied to the string represented by the Unbounded\_String parameter, and To\_Unbounded\_String is applied the result string.

For each of the procedures Translate, Replace\_Slice, Insert, Overwrite, Delete, Trim, Head, and Tail, the resulting string represented by the Source parameter is given by the corresponding function for fixed-length strings applied to the string represented by Source's original value.

#### Implementation Requirements

No storage associated with an Unbounded\_String object shall be lost upon assignment or scope exit.

#### A.4.6 String-Handling Sets and Mappings

The language-defined package `Strings.Maps.Constants` declares `Character_Set` and `Character_Mapping` constants corresponding to classification and conversion functions in package `Characters.Handling`.

*Static Semantics*

The library package `Strings.Maps.Constants` has the following declaration:

```
package Ada.Strings.Maps.Constants
with Pure is

  Control_Set           : constant Character_Set;
  Graphic_Set           : constant Character_Set;
  Letter_Set            : constant Character_Set;
  Lower_Set             : constant Character_Set;
  Upper_Set             : constant Character_Set;
  Basic_Set             : constant Character_Set;
  Decimal_Digit_Set     : constant Character_Set;
  Hexadecimal_Digit_Set : constant Character_Set;
  Alphanumeric_Set      : constant Character_Set;
  Special_Set           : constant Character_Set;
  ISO_646_Set           : constant Character_Set;

  Lower_Case_Map        : constant Character_Mapping;
  --Maps to lower case for letters, else identity
  Upper_Case_Map        : constant Character_Mapping;
  --Maps to upper case for letters, else identity
  Basic_Map             : constant Character_Mapping;
  --Maps to basic letter for letters, else identity

private
  ... -- not specified by the language
end Ada.Strings.Maps.Constants;
```

Each of these constants represents a correspondingly named set of characters or character mapping in `Characters.Handling` (see A.3.2).

NOTE There are certain characters which are defined to be lower case letters by ISO/IEC 10646 and are therefore allowed in identifiers, but are not considered lower case letters by `Ada.Strings.Maps.Constants`.

#### A.4.7 Wide\_String Handling

Facilities for handling strings of `Wide_Character` elements are found in the packages `Strings.Wide_Maps`, `Strings.Wide_Fixed`, `Strings.Wide_Bounded`, `Strings.Wide_Unbounded`, and `Strings.Wide_Maps.Constants`, and in the library functions `Strings.Wide_Hash`, `Strings.Wide_Fixed.Wide_Hash`, `Strings.Wide_Bounded.Wide_Hash`, `Strings.Wide_Unbounded.Wide_Hash`, `Strings.Wide_Hash_Case_Insensitive`, `Strings.Wide_Fixed.Wide_Hash_Case_Insensitive`, `Strings.Wide_Bounded.Wide_Hash_Case_Insensitive`, `Strings.Wide_Unbounded.Wide_Hash_Case_Insensitive`, `Strings.Wide_Equal_Case_Insensitive`, `Strings.Wide_Fixed.Wide_Equal_Case_Insensitive`, `Strings.Wide_Bounded.Wide_Equal_Case_Insensitive`, and `Strings.Wide_Unbounded.Wide_Equal_Case_Insensitive`. They provide the same string-handling operations as the corresponding packages and functions for strings of `Character` elements.

*Static Semantics*

The package `Strings.Wide_Maps` has the following declaration.

```
package Ada.Strings.Wide_Maps
with Preelaborate, Nonblocking, Global => in out synchronized is

  -- Representation for a set of Wide_Character values:
  type Wide_Character_Set is private
    with Preelaborable_Initialization;

  Null_Set : constant Wide_Character_Set;
```

```

type Wide_Character_Range is
  record
    Low   : Wide_Character;
    High  : Wide_Character;
  end record;
-- Represents Wide_Character range Low..High

type Wide_Character_Ranges is array (Positive range <>)
  of Wide_Character_Range;

function To_Set      (Ranges : in Wide_Character_Ranges)
  return Wide_Character_Set;

function To_Set      (Span   : in Wide_Character_Range)
  return Wide_Character_Set;

function To_Ranges (Set      : in Wide_Character_Set)
  return Wide_Character_Ranges;

function "="      (Left, Right : in Wide_Character_Set) return Boolean;

function "not"     (Right : in Wide_Character_Set)
  return Wide_Character_Set;

function "and"     (Left, Right : in Wide_Character_Set)
  return Wide_Character_Set;

function "or"      (Left, Right : in Wide_Character_Set)
  return Wide_Character_Set;

function "xor"     (Left, Right : in Wide_Character_Set)
  return Wide_Character_Set;

function "-"       (Left, Right : in Wide_Character_Set)
  return Wide_Character_Set;

function Is_In (Element : in Wide_Character;
                Set      : in Wide_Character_Set)
  return Boolean;

function Is_Subset (Elements : in Wide_Character_Set;
                   Set       : in Wide_Character_Set)
  return Boolean;

function "<=" (Left   : in Wide_Character_Set;
              Right  : in Wide_Character_Set)
  return Boolean renames Is_Subset;

-- Alternative representation for a set of Wide_Character values:
subtype Wide_Character_Sequence is Wide_String;

function To_Set_Sequence : in Wide_Character_Sequence)
  return Wide_Character_Set;

function To_Set (Singleton : in Wide_Character)
  return Wide_Character_Set;

function To_Sequence (Set : in Wide_Character_Set)
  return Wide_Character_Sequence;

-- Representation for a Wide_Character to Wide_Character mapping:
type Wide_Character_Mapping is private
  with Freelaborable_Initialization;

function Value (Map      : in Wide_Character_Mapping;
               Element : in Wide_Character)
  return Wide_Character;

Identity : constant Wide_Character_Mapping;

function To_Mapping (From, To : in Wide_Character_Sequence)
  return Wide_Character_Mapping;

function To_Domain (Map : in Wide_Character_Mapping)
  return Wide_Character_Sequence;

function To_Range (Map : in Wide_Character_Mapping)
  return Wide_Character_Sequence;

```

```

type Wide_Character_Mapping_Function is
  access function (From : in Wide_Character) return Wide_Character;

private
  ... -- not specified by the language
end Ada.Strings.Wide_Maps;

```

The context clause for each of the packages Strings.Wide\_Fixed, Strings.Wide\_Bounded, and Strings.Wide\_Unbounded identifies Strings.Wide\_Maps instead of Strings.Maps.

Types Wide\_Character\_Set and Wide\_Character\_Mapping need finalization.

For each of the packages Strings.Fixed, Strings.Bounded, Strings.Unbounded, and Strings.Maps.Constants, and for library functions Strings.Hash, Strings.Fixed.Hash, Strings.Bounded.Hash, Strings.Unbounded.Hash, Strings.Hash\_Case\_Insensitive, Strings.Fixed.Hash\_Case\_Insensitive, Strings.Bounded.Hash\_Case\_Insensitive, Strings.Unbounded.Hash\_Case\_Insensitive, Strings.Equal\_Case\_Insensitive, Strings.Fixed.Equal\_Case\_Insensitive, Strings.Bounded.Equal\_Case\_Insensitive, and Strings.Unbounded.Equal\_Case\_Insensitive, the corresponding wide string package or function has the same contents except that

- Wide\_Space replaces Space
- Wide\_Character replaces Character
- Wide\_String replaces String
- Wide\_Character\_Set replaces Character\_Set
- Wide\_Character\_Mapping replaces Character\_Mapping
- Wide\_Character\_Mapping\_Function replaces Character\_Mapping\_Function
- Wide\_Maps replaces Maps
- Bounded\_Wide\_String replaces Bounded\_String
- Null\_Bounded\_Wide\_String replaces Null\_Bounded\_String
- To\_Bounded\_Wide\_String replaces To\_Bounded\_String
- To\_Wide\_String replaces To\_String
- Set\_Bounded\_Wide\_String replaces Set\_Bounded\_String
- Unbounded\_Wide\_String replaces Unbounded\_String
- Null\_Unbounded\_Wide\_String replaces Null\_Unbounded\_String
- Wide\_String\_Access replaces String\_Access
- To\_Unbounded\_Wide\_String replaces To\_Unbounded\_String
- Set\_Unbounded\_Wide\_String replaces Set\_Unbounded\_String

The following additional declaration is present in Strings.Wide\_Maps.Wide\_Constants:

```

Character_Set : constant Wide_Maps.Wide_Character_Set;
--Contains each Wide_Character value WC such that
--Characters.Conversions.Is_Character(WC) is True

```

Each Wide\_Character\_Set constant in the package Strings.Wide\_Maps.Wide\_Constants contains no values outside the Character portion of Wide\_Character. Similarly, each Wide\_Character\_Mapping constant in this package is the identity mapping when applied to any element outside the Character portion of Wide\_Character.

Aspect Pure is replaced by aspects Preelaborate, Nonblocking, Global => **in out synchronized** in Strings.Wide\_Maps.Wide\_Constants.

NOTE If a null Wide\_Character\_Mapping\_Function is passed to any of the Wide\_String handling subprograms, Constraint\_Error is propagated.



#### A.4.8 Wide\_Wide\_String Handling

Facilities for handling strings of Wide\_Wide\_Character elements are found in the packages Strings.Wide\_Wide\_Maps, Strings.Wide\_Wide\_Fixed, Strings.Wide\_Wide\_Bounded, Strings.Wide\_Wide\_Unbounded, and Strings.Wide\_Wide\_Maps.Wide\_Wide\_Constants, and in the library functions Strings.Wide\_Wide\_Hash, Strings.Wide\_Wide\_Fixed.Wide\_Wide\_Hash, Strings.Wide\_Wide\_Bounded.Wide\_Wide\_Hash, Strings.Wide\_Wide\_Unbounded.Wide\_Wide\_Hash, Strings.Wide\_Wide\_Hash\_Case\_Insensitive, Strings.Wide\_Wide\_Fixed.Wide\_Wide\_Hash\_Case\_Insensitive, Strings.Wide\_Wide\_Bounded.Wide\_Wide\_Hash\_Case\_Insensitive, Strings.Wide\_Wide\_Unbounded.Wide\_Wide\_Hash\_Case\_Insensitive, Strings.Wide\_Wide\_Equal\_Case\_Insensitive, Strings.Wide\_Wide\_Fixed.Wide\_Wide\_Equal\_Case\_Insensitive, Strings.Wide\_Wide\_Bounded.Wide\_Wide\_Equal\_Case\_Insensitive, and Strings.Wide\_Wide\_Unbounded.Wide\_Wide\_Equal\_Case\_Insensitive. They provide the same string-handling operations as the corresponding packages and functions for strings of Character elements.

##### Static Semantics

The library package Strings.Wide\_Wide\_Maps has the following declaration.

```
package Ada.Strings.Wide_Wide_Maps
with Preelaborate, Nonblocking, Global => in out synchronized is

-- Representation for a set of Wide_Wide_Character values:
type Wide_Wide_Character_Set is private
  with Preelaborable_Initialization;

Null_Set : constant Wide_Wide_Character_Set;

type Wide_Wide_Character_Range is
  record
    Low   : Wide_Wide_Character;
    High  : Wide_Wide_Character;
  end record;
-- Represents Wide_Wide_Character range Low..High

type Wide_Wide_Character_Ranges is array (Positive range <>)
  of Wide_Wide_Character_Range;

function To_Set (Ranges : in Wide_Wide_Character_Ranges)
  return Wide_Wide_Character_Set;

function To_Set (Span : in Wide_Wide_Character_Range)
  return Wide_Wide_Character_Set;

function To_Ranges (Set : in Wide_Wide_Character_Set)
  return Wide_Wide_Character_Ranges;

function "=" (Left, Right : in Wide_Wide_Character_Set) return Boolean;

function "not" (Right : in Wide_Wide_Character_Set)
  return Wide_Wide_Character_Set;

function "and" (Left, Right : in Wide_Wide_Character_Set)
  return Wide_Wide_Character_Set;

function "or" (Left, Right : in Wide_Wide_Character_Set)
  return Wide_Wide_Character_Set;

function "xor" (Left, Right : in Wide_Wide_Character_Set)
  return Wide_Wide_Character_Set;

function "-" (Left, Right : in Wide_Wide_Character_Set)
  return Wide_Wide_Character_Set;

function Is_In (Element : in Wide_Wide_Character;
               Set      : in Wide_Wide_Character_Set)
  return Boolean;

function Is_Subset (Elements : in Wide_Wide_Character_Set;
                   Set       : in Wide_Wide_Character_Set)
  return Boolean;

function "<=" (Left   : in Wide_Wide_Character_Set;
              Right  : in Wide_Wide_Character_Set)
  return Boolean renames Is_Subset;
```

```

-- Alternative representation for a set of Wide_Wide_Character values:
subtype Wide_Wide_Character_Sequence is Wide_Wide_String;

function To_Set (Sequence : in Wide_Wide_Character_Sequence)
    return Wide_Wide_Character_Set;

function To_Set (Singleton : in Wide_Wide_Character)
    return Wide_Wide_Character_Set;

function To_Sequence (Set : in Wide_Wide_Character_Set)
    return Wide_Wide_Character_Sequence;

-- Representation for a Wide_Wide_Character to Wide_Wide_Character
-- mapping:
type Wide_Wide_Character_Mapping is private
    with Preelaborable_Initialization;

function Value (Map      : in Wide_Wide_Character_Mapping;
                Element : in Wide_Wide_Character)
    return Wide_Wide_Character;

Identity : constant Wide_Wide_Character_Mapping;

function To_Mapping (From, To : in Wide_Wide_Character_Sequence)
    return Wide_Wide_Character_Mapping;

function To_Domain (Map : in Wide_Wide_Character_Mapping)
    return Wide_Wide_Character_Sequence;

function To_Range (Map : in Wide_Wide_Character_Mapping)
    return Wide_Wide_Character_Sequence;

type Wide_Wide_Character_Mapping_Function is
    access function (From : in Wide_Wide_Character)
    return Wide_Wide_Character;

private
    ... -- not specified by the language
end Ada.Strings.Wide_Wide_Maps;

```

The context clause for each of the packages Strings.Wide\_Wide\_Fixed, Strings.Wide\_Wide\_Bounded, and Strings.Wide\_Wide\_Unbounded identifies Strings.Wide\_Wide\_Maps instead of Strings.Maps.

Types Wide\_Wide\_Character\_Set and Wide\_Wide\_Character\_Mapping need finalization.

For each of the packages Strings.Fixed, Strings.Bounded, Strings.Unbounded, and Strings.-Maps.Constants, and for library functions Strings.Hash, Strings.Fixed.Hash, Strings.Bounded.Hash, Strings.Unbounded.Hash, Strings.Hash\_Case\_Insensitive, Strings.Fixed.Hash\_Case\_Insensitive, Strings.-Bounded.Hash\_Case\_Insensitive, Strings.Unbounded.Hash\_Case\_Insensitive, Strings.Equal\_Case\_Insensitive, Strings.Fixed.Equal\_Case\_Insensitive, Strings.Bounded.Equal\_Case\_Insensitive, and Strings.-Unbounded.Equal\_Case\_Insensitive, the corresponding wide wide string package or function has the same contents except that

- Wide\_Space replaces Space
- Wide\_Wide\_Character replaces Character
- Wide\_Wide\_String replaces String
- Wide\_Wide\_Character\_Set replaces Character\_Set
- Wide\_Wide\_Character\_Mapping replaces Character\_Mapping
- Wide\_Wide\_Character\_Mapping\_Function replaces Character\_Mapping\_Function
- Wide\_Wide\_Maps replaces Maps
- Bounded\_Wide\_Wide\_String replaces Bounded\_String
- Null\_Bounded\_Wide\_Wide\_String replaces Null\_Bounded\_String
- To\_Bounded\_Wide\_Wide\_String replaces To\_Bounded\_String
- To\_Wide\_Wide\_String replaces To\_String

- Set\_Bounded\_Wide\_Wide\_String replaces Set\_Bounded\_String
- Unbounded\_Wide\_Wide\_String replaces Unbounded\_String
- Null\_Unbounded\_Wide\_Wide\_String replaces Null\_Unbounded\_String
- Wide\_Wide\_String\_Access replaces String\_Access
- To\_Unbounded\_Wide\_Wide\_String replaces To\_Unbounded\_String
- Set\_Unbounded\_Wide\_Wide\_String replaces Set\_Unbounded\_String

The following additional declarations are present in Strings.Wide\_Wide\_Maps.Wide\_Wide\_Constants:

```
Character_Set : constant Wide_Wide_Maps.Wide_Wide_Character_Set;
-- Contains each Wide_Wide_Character value WWC such that
-- Characters.Conversions.Is_Character(WWC) is True
Wide_Character_Set : constant Wide_Wide_Maps.Wide_Wide_Character_Set;
-- Contains each Wide_Wide_Character value WWC such that
-- Characters.Conversions.Is_Wide_Character(WWC) is True
```

Each Wide\_Wide\_Character\_Set constant in the package Strings.Wide\_Wide\_Maps.Wide\_Wide\_Constants contains no values outside the Character portion of Wide\_Wide\_Character. Similarly, each Wide\_Wide\_Character\_Mapping constant in this package is the identity mapping when applied to any element outside the Character portion of Wide\_Wide\_Character.

Aspect Pure is replaced by aspects Preelaborate, Nonblocking, Global => **in out synchronized** in Strings.Wide\_Wide\_Maps.Wide\_Wide\_Constants.

NOTE If a null Wide\_Wide\_Character\_Mapping\_Function is passed to any of the Wide\_Wide\_String handling subprograms, Constraint\_Error is propagated.

#### A.4.9 String Hashing

##### Static Semantics

The library function Strings.Hash has the following declaration:

```
with Ada.Containers;
function Ada.Strings.Hash (Key : String) return Containers.Hash_Type
with Pure;
```

Returns an implementation-defined value which is a function of the value of Key. If *A* and *B* are strings such that *A* equals *B*, Hash(*A*) equals Hash(*B*).

The library function Strings.Fixed.Hash has the following declaration:

```
with Ada.Containers, Ada.Strings.Hash;
function Ada.Strings.Fixed.Hash (Key : String) return Containers.Hash_Type
renames Ada.Strings.Hash;
```

The generic library function Strings.Bounded.Hash has the following declaration:

```
with Ada.Containers;
generic
  with package Bounded is
    new Ada.Strings.Bounded.Generic_Bounded_Length (<>);
function Ada.Strings.Bounded.Hash (Key : Bounded.Bounded_String)
return Containers.Hash_Type
with Preelaborate, Nonblocking, Global => in out synchronized;
```

Equivalent to Strings.Hash (Bounded.To\_String (Key));

The library function Strings.Unbounded.Hash has the following declaration:

```
with Ada.Containers;
function Ada.Strings.Unbounded.Hash (Key : Unbounded_String)
return Containers.Hash_Type
with Preelaborate, Nonblocking, Global => in out synchronized;
```

Equivalent to Strings.Hash (To\_String (Key));

The library function `Strings.Hash_Case_Insensitive` has the following declaration:

```
with Ada.Containers;  
function Ada.Strings.Hash_Case_Insensitive (Key : String)  
  return Containers.Hash_Type  
with Pure;
```

Returns an implementation-defined value which is a function of the value of `Key`, converted to lower case. If `A` and `B` are strings such that `Strings.Equal_Case_Insensitive (A, B)` (see A.4.10) is `True`, then `Hash_Case_Insensitive(A)` equals `Hash_Case_Insensitive(B)`.

The library function `Strings.Fixed.Hash_Case_Insensitive` has the following declaration:

```
with Ada.Containers, Ada.Strings.Hash_Case_Insensitive;  
function Ada.Strings.Fixed.Hash_Case_Insensitive (Key : String)  
  return Containers.Hash_Type renames Ada.Strings.Hash_Case_Insensitive;
```

The generic library function `Strings.Bounded.Hash_Case_Insensitive` has the following declaration:

```
with Ada.Containers;  
generic  
  with package Bounded is  
    new Ada.Strings.Bounded.Generic_Bounded_Length (<>);  
function Ada.Strings.Bounded.Hash_Case_Insensitive  
  (Key : Bounded.Bounded_String) return Containers.Hash_Type  
with Preelaborate, Nonblocking, Global => in out synchronized;  
  
  Equivalent to Strings.Hash_Case_Insensitive (Bounded.To_String (Key));
```

The library function `Strings.Unbounded.Hash_Case_Insensitive` has the following declaration:

```
with Ada.Containers;  
function Ada.Strings.Unbounded.Hash_Case_Insensitive  
  (Key : Unbounded_String) return Containers.Hash_Type  
with Preelaborate, Nonblocking, Global => in out synchronized;  
  
  Equivalent to Strings.Hash_Case_Insensitive (To_String (Key));
```

#### *Implementation Advice*

The Hash functions should be good hash functions, returning a wide spread of values for different string values. It should be unlikely for similar strings to return the same value.

### A.4.10 String Comparison

#### *Static Semantics*

The library function `Strings.Equal_Case_Insensitive` has the following declaration:

```
function Ada.Strings.Equal_Case_Insensitive (Left, Right : String)  
  return Boolean with Pure;
```

Returns `True` if the strings consist of the same sequence of characters after applying locale-independent simple case folding, as defined by documents referenced in Clause 2 of ISO/IEC 10646:2020. Otherwise, returns `False`. This function uses the same method as is used to determine whether two identifiers are the same.

The library function `Strings.Fixed.Equal_Case_Insensitive` has the following declaration:

```
with Ada.Strings.Equal_Case_Insensitive;  
function Ada.Strings.Fixed.Equal_Case_Insensitive  
  (Left, Right : String) return Boolean  
  renames Ada.Strings.Equal_Case_Insensitive;
```

The generic library function `Strings.Bounded.Equal_Case_Insensitive` has the following declaration:

```
generic
  with package Bounded is
    new Ada.Strings.Bounded.Generic_Bounded_Length (<>);
  function Ada.Strings.Bounded.Equal_Case_Insensitive
    (Left, Right : Bounded.Bounded_String) return Boolean
  with Preelaborate, Nonblocking, Global => in out synchronized;

  Equivalent to Strings.Equal_Case_Insensitive (Bounded.To_String (Left), Bounded.To_String
    (Right));
```

The library function `Strings.Unbounded.Equal_Case_Insensitive` has the following declaration:

```
function Ada.Strings.Unbounded.Equal_Case_Insensitive
  (Left, Right : Unbounded_String) return Boolean
  with Preelaborate, Nonblocking, Global => in out synchronized;

  Equivalent to Strings.Equal_Case_Insensitive (To_String (Left), To_String (Right));
```

The library function `Strings.Less_Case_Insensitive` has the following declaration:

```
function Ada.Strings.Less_Case_Insensitive (Left, Right : String)
  return Boolean with Pure;
```

Performs a lexicographic comparison of strings `Left` and `Right`, converted to lower case.

The library function `Strings.Fixed.Less_Case_Insensitive` has the following declaration:

```
with Ada.Strings.Less_Case_Insensitive;
function Ada.Strings.Fixed.Less_Case_Insensitive
  (Left, Right : String) return Boolean
  renames Ada.Strings.Less_Case_Insensitive;
```

The generic library function `Strings.Bounded.Less_Case_Insensitive` has the following declaration:

```
generic
  with package Bounded is
    new Ada.Strings.Bounded.Generic_Bounded_Length (<>);
  function Ada.Strings.Bounded.Less_Case_Insensitive
    (Left, Right : Bounded.Bounded_String) return Boolean
  with Preelaborate, Nonblocking, Global => in out synchronized;

  Equivalent to Strings.Less_Case_Insensitive (Bounded.To_String (Left), Bounded.To_String
    (Right));
```

The library function `Strings.Unbounded.Less_Case_Insensitive` has the following declaration:

```
function Ada.Strings.Unbounded.Less_Case_Insensitive
  (Left, Right : Unbounded_String) return Boolean
  with Preelaborate, Nonblocking, Global => in out synchronized;

  Equivalent to Strings.Less_Case_Insensitive (To_String (Left), To_String (Right));
```

#### A.4.11 String Encoding

Facilities for encoding, decoding, and converting strings in various character encoding schemes are provided by packages `Strings.UTF_Encoding`, `Strings.UTF_Encoding.Conversions`, `Strings.UTF_Encoding.Strings`, `Strings.UTF_Encoding.Wide_Strings`, and `Strings.UTF_Encoding.Wide_Wide_Strings`.

##### Static Semantics

The encoding library packages have the following declarations:

```
package Ada.Strings.UTF_Encoding
  with Pure is
    -- Declarations common to the string encoding packages
    type Encoding_Scheme is (UTF_8, UTF_16BE, UTF_16LE);
```

```

subtype UTF_String is String;
subtype UTF_8_String is String;
subtype UTF_16_Wide_String is Wide_String;
Encoding_Error : exception;
BOM_8      : constant UTF_8_String :=
    Character'Val(16#EF#) &
    Character'Val(16#BB#) &
    Character'Val(16#BF#);
BOM_16BE : constant UTF_String :=
    Character'Val(16#FE#) &
    Character'Val(16#FF#);
BOM_16LE : constant UTF_String :=
    Character'Val(16#FF#) &
    Character'Val(16#FE#);
BOM_16    : constant UTF_16_Wide_String :=
    (1 => Wide_Character'Val(16#FEFF#));
function Encoding (Item      : UTF_String;
                   Default : Encoding_Scheme := UTF_8)
    return Encoding_Scheme;
end Ada.Strings.UTF_Encoding;
package Ada.Strings.UTF_Encoding.Conversions
with Pure is
    -- Conversions between various encoding schemes
    function Convert (Item      : UTF_String;
                   Input_Scheme : Encoding_Scheme;
                   Output_Scheme : Encoding_Scheme;
                   Output_BOM   : Boolean := False) return UTF_String;

    function Convert (Item      : UTF_String;
                   Input_Scheme : Encoding_Scheme;
                   Output_BOM   : Boolean := False)
    return UTF_16_Wide_String;

    function Convert (Item      : UTF_8_String;
                   Output_BOM   : Boolean := False)
    return UTF_16_Wide_String;

    function Convert (Item      : UTF_16_Wide_String;
                   Output_Scheme : Encoding_Scheme;
                   Output_BOM   : Boolean := False) return UTF_String;

    function Convert (Item      : UTF_16_Wide_String;
                   Output_BOM   : Boolean := False) return UTF_8_String;
end Ada.Strings.UTF_Encoding.Conversions;
package Ada.Strings.UTF_Encoding.Strings
with Pure is
    -- Encoding / decoding between String and various encoding schemes
    function Encode (Item      : String;
                  Output_Scheme : Encoding_Scheme;
                  Output_BOM   : Boolean := False) return UTF_String;

    function Encode (Item      : String;
                  Output_BOM : Boolean := False) return UTF_8_String;

    function Encode (Item      : String;
                  Output_BOM : Boolean := False)
    return UTF_16_Wide_String;

    function Decode (Item      : UTF_String;
                  Input_Scheme : Encoding_Scheme) return String;

    function Decode (Item : UTF_8_String) return String;

```



```

function Decode (Item : UTF_16_Wide_String) return String;
end Ada.Strings.UTF_Encoding.Strings;
package Ada.Strings.UTF_Encoding.Wide_Strings
with Pure is
    -- Encoding / decoding between Wide_String and various encoding schemes
    function Encode (Item      : Wide_String;
                    Output_Scheme : Encoding_Scheme;
                    Output_BOM   : Boolean := False) return UTF_String;

    function Encode (Item      : Wide_String;
                    Output_BOM : Boolean := False) return UTF_8_String;

    function Encode (Item      : Wide_String;
                    Output_BOM : Boolean := False)
        return UTF_16_Wide_String;

    function Decode (Item      : UTF_String;
                    Input_Scheme : Encoding_Scheme) return Wide_String;

    function Decode (Item : UTF_8_String) return Wide_String;

    function Decode (Item : UTF_16_Wide_String) return Wide_String;
end Ada.Strings.UTF_Encoding.Wide_Strings;
package Ada.Strings.UTF_Encoding.Wide_Wide_Strings
with Pure is
    -- Encoding / decoding between Wide_Wide_String and various encoding
    schemes
    function Encode (Item      : Wide_Wide_String;
                    Output_Scheme : Encoding_Scheme;
                    Output_BOM   : Boolean := False) return UTF_String;

    function Encode (Item      : Wide_Wide_String;
                    Output_BOM : Boolean := False) return UTF_8_String;

    function Encode (Item      : Wide_Wide_String;
                    Output_BOM : Boolean := False)
        return UTF_16_Wide_String;

    function Decode (Item      : UTF_String;
                    Input_Scheme : Encoding_Scheme) return Wide_Wide_String;

    function Decode (Item : UTF_8_String) return Wide_Wide_String;

    function Decode (Item : UTF_16_Wide_String) return Wide_Wide_String;
end Ada.Strings.UTF_Encoding.Wide_Wide_Strings;

```

The type Encoding\_Scheme defines encoding schemes. UTF\_8 corresponds to the UTF-8 encoding scheme defined by Annex D of ISO/IEC 10646. UTF\_16BE corresponds to the UTF-16 encoding scheme defined by Annex C of ISO/IEC 10646 in 8 bit, big-endian order; and UTF\_16LE corresponds to the UTF-16 encoding scheme in 8 bit, little-endian order.

The subtype UTF\_String is used to represent a String of 8-bit values containing a sequence of values encoded in one of three ways (UTF-8, UTF-16BE, or UTF-16LE). The subtype UTF\_8\_String is used to represent a String of 8-bit values containing a sequence of values encoded in UTF-8. The subtype UTF\_16\_Wide\_String is used to represent a Wide\_String of 16-bit values containing a sequence of values encoded in UTF-16.

The BOM\_8, BOM\_16BE, BOM\_16LE, and BOM\_16 constants correspond to values used at the start of a string to indicate the encoding.

Each of the Encode functions takes a String, Wide\_String, or Wide\_Wide\_String Item parameter that is assumed to be an array of unencoded characters. Each of the Convert functions takes a UTF\_String, UTF\_8\_String, or UTF\_16\_String Item parameter that is assumed to contain characters whose position values correspond to a valid encoding sequence according to the encoding scheme required by the function or specified by its Input\_Scheme parameter.

Each of the Convert and Encode functions returns a UTF\_String, UTF\_8\_String, or UTF\_16\_String value whose characters have position values that correspond to the encoding of the Item parameter according to the encoding scheme required by the function or specified by its Output\_Scheme parameter. For UTF\_8, no overlong encoding is returned. A BOM is included at the start of the returned string if the Output\_BOM parameter is set to True. The lower bound of the returned string is 1.

Each of the Decode functions takes a UTF\_String, UTF\_8\_String, or UTF\_16\_String Item parameter which is assumed to contain characters whose position values correspond to a valid encoding sequence according to the encoding scheme required by the function or specified by its Input\_Scheme parameter, and returns the corresponding String, Wide\_String, or Wide\_Wide\_String value. The lower bound of the returned string is 1.

For each of the Convert and Decode functions, an initial BOM in the input that matches the expected encoding scheme is ignored, and a different initial BOM causes Encoding\_Error to be propagated.

The exception Encoding\_Error is also propagated in the following situations:

- By a Convert or Decode function when a UTF encoded string contains an invalid encoding sequence.
- By a Convert or Decode function when the expected encoding is UTF-16BE or UTF-16LE and the input string has an odd length.
- By a Decode function yielding a String when the decoding of a sequence results in a code point whose value exceeds 16#FF#.
- By a Decode function yielding a Wide\_String when the decoding of a sequence results in a code point whose value exceeds 16#FFFF#.
- By an Encode function taking a Wide\_String as input when an invalid character appears in the input. In particular, the characters whose position is in the range 16#D800# .. 16#DFFF# are invalid because they conflict with UTF-16 surrogate encodings, and the characters whose position is 16#FFFE# or 16#FFFF# are also invalid because they conflict with BOM codes.

```
function Encoding (Item      : UTF_String;
                  Default   : Encoding_Scheme := UTF_8)
return Encoding_Scheme;
```

Inspects a UTF\_String value to determine whether it starts with a BOM for UTF-8, UTF-16BE, or UTF-16LE. If so, returns the scheme corresponding to the BOM; otherwise, returns the value of Default.

```
function Convert (Item          : UTF_String;
                 Input_Scheme   : Encoding_Scheme;
                 Output_Scheme  : Encoding_Scheme;
                 Output_BOM     : Boolean := False) return UTF_String;
```

Returns the value of Item (originally encoded in UTF-8, UTF-16LE, or UTF-16BE as specified by Input\_Scheme) encoded in one of these three schemes as specified by Output\_Scheme.

```
function Convert (Item          : UTF_String;
                 Input_Scheme   : Encoding_Scheme;
                 Output_BOM     : Boolean := False)
return UTF_16_Wide_String;
```

Returns the value of Item (originally encoded in UTF-8, UTF-16LE, or UTF-16BE as specified by Input\_Scheme) encoded in UTF-16.

```
function Convert (Item          : UTF_8_String;
                 Output_BOM     : Boolean := False)
return UTF_16_Wide_String;
```

Returns the value of Item (originally encoded in UTF-8) encoded in UTF-16.

```
function Convert (Item           : UTF_16_Wide_String;
                  Output_Scheme : Encoding_Scheme;
                  Output_BOM    : Boolean := False) return UTF_String;
```

Returns the value of Item (originally encoded in UTF-16) encoded in UTF-8, UTF-16LE, or UTF-16BE as specified by Output\_Scheme.

```
function Convert (Item           : UTF_16_Wide_String;
                  Output_BOM    : Boolean := False) return UTF_8_String;
```

Returns the value of Item (originally encoded in UTF-16) encoded in UTF-8.

```
function Encode (Item           : String;
                 Output_Scheme : Encoding_Scheme;
                 Output_BOM    : Boolean := False) return UTF_String;
```

Returns the value of Item encoded in UTF-8, UTF-16LE, or UTF-16BE as specified by Output\_Scheme.

```
function Encode (Item           : String;
                 Output_BOM    : Boolean := False) return UTF_8_String;
```

Returns the value of Item encoded in UTF-8.

```
function Encode (Item           : String;
                 Output_BOM    : Boolean := False) return UTF_16_Wide_String;
```

Returns the value of Item encoded in UTF-16.

```
function Decode (Item           : UTF_String;
                 Input_Scheme  : Encoding_Scheme) return String;
```

Returns the result of decoding Item, which is encoded in UTF-8, UTF-16LE, or UTF-16BE as specified by Input\_Scheme.

```
function Decode (Item : UTF_8_String) return String;
```

Returns the result of decoding Item, which is encoded in UTF-8.

```
function Decode (Item : UTF_16_Wide_String) return String;
```

Returns the result of decoding Item, which is encoded in UTF-16.

```
function Encode (Item           : Wide_String;
                 Output_Scheme : Encoding_Scheme;
                 Output_BOM    : Boolean := False) return UTF_String;
```

Returns the value of Item encoded in UTF-8, UTF-16LE, or UTF-16BE as specified by Output\_Scheme.

```
function Encode (Item           : Wide_String;
                 Output_BOM    : Boolean := False) return UTF_8_String;
```

Returns the value of Item encoded in UTF-8.

```
function Encode (Item           : Wide_String;
                 Output_BOM    : Boolean := False) return UTF_16_Wide_String;
```

Returns the value of Item encoded in UTF-16.

```
function Decode (Item           : UTF_String;
                 Input_Scheme  : Encoding_Scheme) return Wide_String;
```

Returns the result of decoding Item, which is encoded in UTF-8, UTF-16LE, or UTF-16BE as specified by Input\_Scheme.

```
function Decode (Item : UTF_8_String) return Wide_String;
```

Returns the result of decoding Item, which is encoded in UTF-8.

```
function Decode (Item : UTF_16_Wide_String) return Wide_String;
```

Returns the result of decoding Item, which is encoded in UTF-16.

```
function Encode (Item           : Wide_Wide_String;
                  Output_Scheme : Encoding_Scheme;
                  Output_BOM    : Boolean := False) return UTF_String;
```

Returns the value of Item encoded in UTF-8, UTF-16LE, or UTF-16BE as specified by Output\_Scheme.

```
function Encode (Item           : Wide_Wide_String;
                  Output_BOM    : Boolean := False) return UTF_8_String;
```

Returns the value of Item encoded in UTF-8.

```
function Encode (Item           : Wide_Wide_String;
                  Output_BOM    : Boolean := False) return UTF_16_Wide_String;
```

Returns the value of Item encoded in UTF-16.

```
function Decode (Item           : UTF_String;
                  Input_Scheme  : Encoding_Scheme) return Wide_Wide_String;
```

Returns the result of decoding Item, which is encoded in UTF-8, UTF-16LE, or UTF-16BE as specified by Input\_Scheme.

```
function Decode (Item : UTF_8_String) return Wide_Wide_String;
```

Returns the result of decoding Item, which is encoded in UTF-8.

```
function Decode (Item : UTF_16_Wide_String) return Wide_Wide_String;
```

Returns the result of decoding Item, which is encoded in UTF-16.

#### *Implementation Advice*

If an implementation supports other encoding schemes, another similar child of Ada.Strings should be defined.

NOTE A BOM (Byte-Order Mark, code position 16#FEFF#) can be included in a file or other entity to indicate the encoding; it is skipped when decoding. Typically, only the first line of a file or other entity contains a BOM. When decoding, the Encoding function can be called on the first line to determine the encoding; this encoding will then be used in subsequent calls to Decode to convert all of the lines to an internal format.

### A.4.12 Universal Text Buffers

A universal text buffer can be used to save and retrieve text of any language-defined string type. The types used to save and retrieve the text can be different.

#### *Static Semantics*

The text buffer library packages have the following declarations:

```
with Ada.Strings.UTF_Encoding.Wide_Wide_Strings;
package Ada.Strings.Text_Buffers
with Pure is

  type Text_Buffer_Count is range 0 .. implementation-defined;
  New_Line_Count : constant Text_Buffer_Count := implementation-defined;
  type Root_Buffer_Type is abstract tagged private
    with Default_Initial_Condition =>
      Current_Indent (Root_Buffer_Type) = 0;
  procedure Put (
    Buffer : in out Root_Buffer_Type;
    Item  : in      String) is abstract;
```

```

procedure Wide_Put (
  Buffer : in out Root_Buffer_Type;
  Item   : in      Wide_String) is abstract;

procedure Wide_Wide_Put (
  Buffer : in out Root_Buffer_Type;
  Item   : in      Wide_Wide_String) is abstract;

procedure Put_UTF_8 (
  Buffer : in out Root_Buffer_Type;
  Item   : in      UTF_Encoding.UTF_8_String) is abstract;

procedure Wide_Put_UTF_16 (
  Buffer : in out Root_Buffer_Type;
  Item   : in      UTF_Encoding.UTF_16_Wide_String) is abstract;

procedure New_Line (Buffer : in out Root_Buffer_Type) is abstract;

Standard_Indent : constant Text_Buffer_Count := 3;

function Current_Indent (
  Buffer : Root_Buffer_Type) return Text_Buffer_Count;

procedure Increase_Indent (
  Buffer : in out Root_Buffer_Type;
  Amount : in      Text_Buffer_Count := Standard_Indent)
with Post'Class =>
  Current_Indent (Buffer) = Current_Indent (Buffer)'Old + Amount;

procedure Decrease_Indent (
  Buffer : in out Root_Buffer_Type;
  Amount : in      Text_Buffer_Count := Standard_Indent)
with Pre'Class =>
  Current_Indent (Buffer) >= Amount
  or else raise Constraint_Error,
  Post'Class =>
  Current_Indent (Buffer) =
    Current_Indent (Buffer)'Old - Amount;

private
  ... -- not specified by the language
end Ada.Strings.Text_Buffers;

package Ada.Strings.Text_Buffers.Unbounded
with Preelaborate, Nonblocking, Global => null is

type Buffer_Type is new Root_Buffer_Type with private;

function Get (
  Buffer : in out Buffer_Type)
return String
with Post'Class =>
  Get'Result'First = 1 and then Current_Indent (Buffer) = 0;

function Wide_Get (
  Buffer : in out Buffer_Type)
return Wide_String
with Post'Class =>
  Wide_Get'Result'First = 1 and then Current_Indent (Buffer) = 0;

function Wide_Wide_Get (
  Buffer : in out Buffer_Type)
return Wide_Wide_String
with Post'Class =>
  Wide_Wide_Get'Result'First = 1
  and then Current_Indent (Buffer) = 0;

function Get_UTF_8 (
  Buffer : in out Buffer_Type)
return UTF_Encoding.UTF_8_String
with Post'Class =>
  Get_UTF_8'Result'First = 1 and then Current_Indent (Buffer) = 0;

```

```

function Wide_Get_UTF_16 (
  Buffer : in out Buffer_Type)
return UTF_Encoding.UTF_16_Wide_String
with Post'Class =>
  Wide_Get_UTF_16'Result'First = 1
  and then Current_Indent (Buffer) = 0;

private
  ... -- not specified by the language, but will include nonabstract
  -- overridings of all inherited subprograms that require overriding.
end Ada.Strings.Text_Buffers.Unbounded;

package Ada.Strings.Text_Buffers.Bounded
with Pure, Nonblocking, Global => null is

  type Buffer_Type (Max_Characters : Text_Buffer_Count)
  is new Root_Buffer_Type with private
  with Default_Initial_Condition => not Text_Truncated (Buffer_Type);

  function Text_Truncated (Buffer : in Buffer_Type) return Boolean;

  -- Get, Wide_Get, Wide_Wide_Get, Get_UTF_8, and Wide_Get_UTF_16
  -- are declared here just as in the Unbounded child.

private
  ... -- not specified by the language, but will include nonabstract
  -- overridings of all inherited subprograms that require overriding.
end Ada.Strings.Text_Buffers.Bounded;

```

Character\_Count returns the number of characters currently stored in a text buffer.

New\_Line stores New\_Line\_Count characters that represent a new line into a text buffer. Current\_Indent returns the current indentation associated with the buffer, with zero meaning there is no indentation in effect; Increase\_Indent and Decrease\_Indent increase or decrease the indentation associated with the buffer.

A call to Put, Wide\_Put, Wide\_Wide\_Put, Put\_UTF\_8, or Wide\_Put\_UTF\_16 stores a sequence of characters into the text buffer, preceded by Current\_Indent(Buffer) spaces (Wide\_Wide\_Characters with position 32) if there is at least one character in Item and it would have been the first character on the current line.

A call to function Get, Wide\_Get, Wide\_Wide\_Get, Get\_UTF\_8, or Wide\_Get\_UTF\_16 returns the same sequence of characters as was present in the calls that stored the characters into the buffer, if representable. For a call to Get, if any character in the sequence is not defined in Character, the result is implementation defined. Similarly, for a call to Wide\_Get, if any character in the sequence is not defined in Wide\_Character, the result is implementation defined. As part of a call on any of the Get functions, the buffer is reset to an empty state, with no stored characters.

In the case of a Buf of type Text\_Buffers.Bounded.Buffer\_Type, Text\_Truncated (Buf) returns True if the various Put procedures together have attempted to store more than Buf.Max\_Characters into Buf. If this function returns True, then the various Get functions return a representation of only the first Buf.Max\_Characters characters that were stored in Buf.

#### Implementation Advice

Bounded buffer objects should be implemented without dynamic allocation.

## A.5 The Numerics Packages

The library package Numerics is the parent of several child units that provide facilities for mathematical computation. One child, the generic package Generic\_Elementary\_Functions, is defined in A.5.1, together with nongeneric equivalents; two others, the package Float\_Random and the generic package Discrete\_Random, are defined in A.5.2. Additional (optional) children are defined in Annex G.



*Static Semantics*

```

package Ada.Numerics
  with Pure is
    Argument_Error : exception;
    Pi : constant :=
      3.14159_26535_89793_23846_26433_83279_50288_41971_69399_37511;
     $\pi$  : constant := Pi;
    e : constant :=
      2.71828_18284_59045_23536_02874_71352_66249_77572_47093_69996;
end Ada.Numerics;

```

The Argument\_Error exception is raised by a subprogram in a child unit of Numerics to signal that one or more of the actual subprogram parameters are outside the domain of the corresponding mathematical function.

*Implementation Permissions*

The implementation may specify the values of Pi and e to a larger number of significant digits.

**A.5.1 Elementary Functions**

Implementation-defined approximations to the mathematical functions known as the “elementary functions” are provided by the subprograms in Numerics.Generic\_Elementary\_Functions. Nongeneric equivalents of this generic package for each of the predefined floating point types are also provided as children of Numerics.

*Static Semantics*

The generic library package Numerics.Generic\_Elementary\_Functions has the following declaration:

```

generic
  type Float_Type is digits <>;

package Ada.Numerics.Generic_Elementary_Functions
  with Pure, Nonblocking is

    function Sqrt      (X : Float_Type'Base) return Float_Type'Base;
    function Log       (X : Float_Type'Base) return Float_Type'Base;
    function Log       (X, Base : Float_Type'Base) return Float_Type'Base;
    function Exp       (X : Float_Type'Base) return Float_Type'Base;
    function "**"      (Left, Right : Float_Type'Base) return Float_Type'Base;

    function Sin       (X : Float_Type'Base) return Float_Type'Base;
    function Sin       (X, Cycle : Float_Type'Base) return Float_Type'Base;
    function Cos       (X : Float_Type'Base) return Float_Type'Base;
    function Cos       (X, Cycle : Float_Type'Base) return Float_Type'Base;
    function Tan       (X : Float_Type'Base) return Float_Type'Base;
    function Tan       (X, Cycle : Float_Type'Base) return Float_Type'Base;
    function Cot       (X : Float_Type'Base) return Float_Type'Base;
    function Cot       (X, Cycle : Float_Type'Base) return Float_Type'Base;

```

```

function Arcsin (X : Float_Type'Base) return Float_Type'Base;
function Arcsin (X, Cycle : Float_Type'Base) return Float_Type'Base;
function Arccos (X : Float_Type'Base) return Float_Type'Base;
function Arccos (X, Cycle : Float_Type'Base) return Float_Type'Base;
function Arctan (Y : Float_Type'Base;
                X : Float_Type'Base := 1.0)
return Float_Type'Base;

function Arctan (Y : Float_Type'Base;
                X : Float_Type'Base := 1.0;
                Cycle : Float_Type'Base) return Float_Type'Base;
function Arccot (X : Float_Type'Base;
                Y : Float_Type'Base := 1.0)
return Float_Type'Base;

function Arccot (X : Float_Type'Base;
                Y : Float_Type'Base := 1.0;
                Cycle : Float_Type'Base) return Float_Type'Base;

function Sinh (X : Float_Type'Base) return Float_Type'Base;
function Cosh (X : Float_Type'Base) return Float_Type'Base;
function Tanh (X : Float_Type'Base) return Float_Type'Base;
function Coth (X : Float_Type'Base) return Float_Type'Base;
function Arcsinh (X : Float_Type'Base) return Float_Type'Base;
function Arccosh (X : Float_Type'Base) return Float_Type'Base;
function Arctanh (X : Float_Type'Base) return Float_Type'Base;
function Arccoth (X : Float_Type'Base) return Float_Type'Base;

end Ada.Numerics.Generic_Elementary_Functions;

```

The library package `Numerics.Elementary_Functions` is declared pure and defines the same subprograms as `Numerics.Generic_Elementary_Functions`, except that the predefined type `Float` is systematically substituted for `Float_Type'Base` throughout. Nongeneric equivalents of `Numerics.Generic_Elementary_Functions` for each of the other predefined floating point types are defined similarly, with the names `Numerics.Short_Elementary_Functions`, `Numerics.Long_Elementary_Functions`, etc.

The functions have their usual mathematical meanings. When the `Base` parameter is specified, the `Log` function computes the logarithm to the given base; otherwise, it computes the natural logarithm. When the `Cycle` parameter is specified, the parameter `X` of the forward trigonometric functions (`Sin`, `Cos`, `Tan`, and `Cot`) and the results of the inverse trigonometric functions (`Arcsin`, `Arccos`, `Arctan`, and `Arccot`) are measured in units such that a full cycle of revolution has the given value; otherwise, they are measured in radians.

The computed results of the mathematically multivalued functions are rendered single-valued by the following conventions, which are meant to imply the principal branch:

- The results of the `Sqrt` and `Arccosh` functions and that of the exponentiation operator are nonnegative.
- The result of the `Arcsin` function is in the quadrant containing the point  $(1.0, x)$ , where  $x$  is the value of the parameter `X`. This quadrant is I or IV; thus, the range of the `Arcsin` function is approximately  $-\pi/2.0$  to  $\pi/2.0$  ( $-\text{Cycle}/4.0$  to  $\text{Cycle}/4.0$ , if the parameter `Cycle` is specified).
- The result of the `Arccos` function is in the quadrant containing the point  $(x, 1.0)$ , where  $x$  is the value of the parameter `X`. This quadrant is I or II; thus, the `Arccos` function ranges from  $0.0$  to approximately  $\pi$  ( $\text{Cycle}/2.0$ , if the parameter `Cycle` is specified).
- The results of the `Arctan` and `Arccot` functions are in the quadrant containing the point  $(x, y)$ , where  $x$  and  $y$  are the values of the parameters `X` and `Y`, respectively. This may be any quadrant (I through IV) when the parameter `X` (resp., `Y`) of `Arctan` (resp., `Arccot`) is specified, but it is restricted to quadrants I and IV (resp., I and II) when that parameter is omitted. Thus, the range when that parameter is specified is approximately  $-\pi$  to  $\pi$  ( $-\text{Cycle}/2.0$  to  $\text{Cycle}/2.0$ , if the parameter `Cycle` is specified); when omitted, the range of `Arctan` (resp., `Arccot`) is that of `Arcsin` (resp., `Arccos`), as given above. When the point  $(x, y)$  lies on the negative  $x$ -axis, the result approximates

- $\pi$  (resp.,  $-\pi$ ) when the sign of the parameter  $Y$  is positive (resp., negative), if `Float_Type'Signed_Zeros` is `True`;
- $\pi$ , if `Float_Type'Signed_Zeros` is `False`.

(In the case of the inverse trigonometric functions, in which a result lying on or near one of the axes may not be exactly representable, the approximation inherent in computing the result may place it in an adjacent quadrant, close to but on the wrong side of the axis.)

#### Dynamic Semantics

The exception `Numerics.Argument_Error` is raised, signaling a parameter value outside the domain of the corresponding mathematical function, in the following cases:

- by any forward or inverse trigonometric function with specified cycle, when the value of the parameter `Cycle` is zero or negative;
- by the `Log` function with specified base, when the value of the parameter `Base` is zero, one, or negative;
- by the `Sqrt` and `Log` functions, when the value of the parameter  $X$  is negative;
- by the exponentiation operator, when the value of the left operand is negative or when both operands have the value zero;
- by the `Arcsin`, `Arccos`, and `Arctanh` functions, when the absolute value of the parameter  $X$  exceeds one;
- by the `Arctan` and `Arccot` functions, when the parameters  $X$  and  $Y$  both have the value zero;
- by the `Arccosh` function, when the value of the parameter  $X$  is less than one; and
- by the `Arccoth` function, when the absolute value of the parameter  $X$  is less than one.

The exception `Constraint_Error` is raised, signaling a pole of the mathematical function (analogous to dividing by zero), in the following cases, provided that `Float_Type'Machine_Overflows` is `True`:

- by the `Log`, `Cot`, and `Coth` functions, when the value of the parameter  $X$  is zero;
- by the exponentiation operator, when the value of the left operand is zero and the value of the exponent is negative;
- by the `Tan` function with specified cycle, when the value of the parameter  $X$  is an odd multiple of the quarter cycle;
- by the `Cot` function with specified cycle, when the value of the parameter  $X$  is zero or a multiple of the half cycle; and
- by the `Arctanh` and `Arccoth` functions, when the absolute value of the parameter  $X$  is one.

`Constraint_Error` can also be raised when a finite result overflows (see G.2.4); this may occur for parameter values sufficiently *near* poles, and, in the case of some of the functions, for parameter values with sufficiently large magnitudes. When `Float_Type'Machine_Overflows` is `False`, the result at poles is unspecified.

When one parameter of a function with multiple parameters represents a pole and another is outside the function's domain, the latter takes precedence and `Numerics.Argument_Error` is raised.

#### Implementation Requirements

In the implementation of `Numerics.Generic_Elementary_Functions`, the range of intermediate values allowed during the calculation of a final result shall not be affected by any range constraint of the subtype `Float_Type`.

In the following cases, evaluation of an elementary function shall yield the *prescribed result*, provided that the preceding rules do not call for an exception to be raised:

- When the parameter  $X$  has the value zero, the `Sqrt`, `Sin`, `Arcsin`, `Tan`, `Sinh`, `Arcsinh`, `Tanh`, and `Arctanh` functions yield a result of zero, and the `Exp`, `Cos`, and `Cosh` functions yield a result of one.

- When the parameter *X* has the value one, the Sqrt function yields a result of one, and the Log, Arccos, and Arccosh functions yield a result of zero.
- When the parameter *Y* has the value zero and the parameter *X* has a positive value, the Arctan and Arccot functions yield a result of zero.
- The results of the Sin, Cos, Tan, and Cot functions with specified cycle are exact when the mathematical result is zero; those of the first two are also exact when the mathematical result is  $\pm 1.0$ .
- Exponentiation by a zero exponent yields the value one. Exponentiation by a unit exponent yields the value of the left operand. Exponentiation of the value one yields the value one. Exponentiation of the value zero yields the value zero.

Other accuracy requirements for the elementary functions, which apply only in implementations conforming to the Numerics Annex, and then only in the “strict” mode defined there (see G.2), are given in G.2.4.

When Float\_Type'Signed\_Zeros is True, the sign of a zero result shall be as follows:

- A prescribed zero result delivered *at the origin* by one of the odd functions (Sin, Arcsin, Sinh, Arcsinh, Tan, Arctan or Arccot as a function of *Y* when *X* is fixed and positive, Tanh, and Arctanh) has the sign of the parameter *X* (*Y*, in the case of Arctan or Arccot).
- A prescribed zero result delivered by one of the odd functions *away from the origin*, or by some other elementary function, has an implementation-defined sign.
- A zero result that is not a prescribed result (i.e. one that results from rounding or underflow) has the correct mathematical sign.

#### Implementation Permissions

The nongeneric equivalent packages can be actual instantiations of the generic package for the appropriate predefined type, though that is not required.

### A.5.2 Random Number Generation

Facilities for the generation of pseudo-random floating point numbers are provided in the package Numerics.Float\_Random; the generic package Numerics.Discrete\_Random provides similar facilities for the generation of pseudo-random integers and pseudo-random values of enumeration types. For brevity, pseudo-random values of any of these types are called *random numbers*.

Some of the facilities provided are basic to all applications of random numbers. These include a limited private type each of whose objects serves as the generator of a (possibly distinct) sequence of random numbers; a function to obtain the “next” random number from a given sequence of random numbers (that is, from its generator); and subprograms to initialize or reinitialize a given generator to a time-dependent state or a state denoted by a single integer.

Other facilities are provided specifically for advanced applications. These include subprograms to save and restore the state of a given generator; a private type whose objects can be used to hold the saved state of a generator; and subprograms to obtain a string representation of a given generator state, or, given such a string representation, the corresponding state.

#### Static Semantics

The library package Numerics.Float\_Random has the following declaration:

```
package Ada.Numerics.Float_Random
  with Global => in out synchronized is
    -- Basic facilities
    type Generator is limited private;
```

```

subtype Uniformly_Distributed is Float range 0.0 .. 1.0;
function Random (Gen : Generator) return Uniformly_Distributed
  with Global => overriding in out Gen;

procedure Reset (Gen      : in Generator;
                 Initiator : in Integer)
  with Global => overriding in out Gen;
procedure Reset (Gen      : in Generator)
  with Global => overriding in out Gen;

-- Advanced facilities

type State is private;

procedure Save  (Gen      : in Generator;
                 To_State : out State);
procedure Reset (Gen      : in Generator;
                 From_State : in State)
  with Global => overriding in out Gen;

Max_Image_Width : constant := implementation-defined integer value;

function Image (Of_State : State) return String;
function Value (Coded_State : String) return State;

private
  ... -- not specified by the language
end Ada.Numerics.Float_Random;

```

The type Generator needs finalization (see 10.6).

The generic library package Numerics.Discrete\_Random has the following declaration:

```

generic
  type Result_Subtype is (<>);
package Ada.Numerics.Discrete_Random
  with Global => in out synchronized is

  -- Basic facilities

  type Generator is limited private;

  function Random (Gen : Generator) return Result_Subtype
    with Global => overriding in out Gen;

  function Random (Gen      : Generator;
                 First : Result_Subtype;
                 Last  : Result_Subtype) return Result_Subtype
    with Post => Random'Result in First .. Last,
        Global => overriding in out Gen;

  procedure Reset (Gen      : in Generator;
                 Initiator : in Integer)
    with Global => overriding in out Gen;
  procedure Reset (Gen      : in Generator)
    with Global => overriding in out Gen;

  -- Advanced facilities

  type State is private;

  procedure Save  (Gen      : in Generator;
                 To_State : out State);
  procedure Reset (Gen      : in Generator;
                 From_State : in State)
    with Global => overriding in out Gen;

  Max_Image_Width : constant := implementation-defined integer value;

  function Image (Of_State : State) return String;
  function Value (Coded_State : String) return State;

private
  ... -- not specified by the language
end Ada.Numerics.Discrete_Random;

```

The type Generator needs finalization (see 10.6) in every instantiation of Numerics.Discrete\_Random.

An object of the limited private type Generator is associated with a sequence of random numbers. Each generator has a hidden (internal) state, which the operations on generators use to determine the position in the associated sequence. All generators are implicitly initialized to an unspecified state that does not vary from one program execution to another; they may also be explicitly initialized, or reinitialized, to a time-dependent state, to a previously saved state, or to a state uniquely denoted by an integer value.

An object of the private type State can be used to hold the internal state of a generator. Such objects are only necessary if the application is designed to save and restore generator states or to examine or manufacture them. The implicit initial value of type State corresponds to the implicit initial value of all generators.

The operations on generators affect the state and therefore the future values of the associated sequence. The semantics of the operations on generators and states are defined below.

```
function Random (Gen : Generator) return Uniformly_Distributed;
function Random (Gen : Generator) return Result_Subtype;
```

Obtains the “next” random number from the given generator, relative to its current state, according to an implementation-defined algorithm.

```
function Random (Gen      : Generator;
                  First    : Result_Subtype;
                  Last     : Result_Subtype) return Result_Subtype
with Post => Random'Result in First .. Last;
```

Obtains the “next” random number from the given generator, relative to its current state, according to an implementation-defined algorithm. If the range First .. Last is a null range, Constraint\_Error is raised.

```
procedure Reset (Gen      : in Generator;
                 Initiator : in Integer);
procedure Reset (Gen      : in Generator);
```

Sets the state of the specified generator to one that is an unspecified function of the value of the parameter Initiator (or to a time-dependent state, if only a generator parameter is specified). The latter form of the procedure is known as the *time-dependent Reset procedure*.

```
procedure Save  (Gen      : in Generator;
                 To_State : out State);
procedure Reset (Gen      : in Generator;
                 From_State : in State);
```

Save obtains the current state of a generator. Reset gives a generator the specified state. A generator that is reset to a state previously obtained by invoking Save is restored to the state it had when Save was invoked.

```
function Image (Of_State : State) return String;
function Value (Coded_State : String) return State;
```

Image provides a representation of a state coded (in an implementation-defined way) as a string whose length is bounded by the value of Max\_Image\_Width. Value is the inverse of Image: Value(Image(S)) = S for each state S that can be obtained from a generator by invoking Save.

#### Dynamic Semantics

Instantiation of Numerics.Discrete\_Random with a subtype having a null range raises Constraint\_Error.

#### Bounded (Run-Time) Errors

It is a bounded error to invoke Value with a string that is not the image of any generator state. If the error is detected, Constraint\_Error or Program\_Error is raised. Otherwise, a call to Reset with the resulting state will produce a generator such that calls to Random with this generator will produce a



sequence of values of the appropriate subtype, but which are not necessarily random in character. That is, the sequence of values do not necessarily fulfill the implementation requirements of this subclause.

#### *Implementation Requirements*

Each call of a Random function has a *result range*; this is the range First .. Last for the version of Random with First and Last parameters and the range of the result subtype of the function otherwise.

A sufficiently long sequence of random numbers obtained by consecutive calls to Random that have the same generator and result range is approximately uniformly distributed over the result range.

A Random function in an instantiation of Numerics.Discrete\_Random is guaranteed to yield each value in its result range in a finite number of calls, provided that the number of such values does not exceed  $2^{15}$ .

Other performance requirements for the random number generator, which apply only in implementations conforming to the Numerics Annex, and then only in the “strict” mode defined there (see G.2), are given in G.2.5.

#### *Documentation Requirements*

No one algorithm for random number generation is best for all applications. To enable the user to determine the suitability of the random number generators for the intended application, the implementation shall describe the algorithm used and shall give its period, if known exactly, or a lower bound on the period, if the exact period is unknown. Periods that are so long that the periodicity is unobservable in practice can be described in such terms, without giving a numerical bound.

The implementation also shall document the minimum time interval between calls to the time-dependent Reset procedure that are guaranteed to initiate different sequences, and it shall document the nature of the strings that Value will accept without raising Constraint\_Error.

#### *Implementation Advice*

Any storage associated with an object of type Generator should be reclaimed on exit from the scope of the object.

If the generator period is sufficiently long in relation to the number of distinct initiator values, then each possible value of Initiator passed to Reset should initiate a sequence of random numbers that does not, in a practical sense, overlap the sequence initiated by any other value. If this is not possible, then the mapping between initiator values and generator states should be a rapidly varying function of the initiator value.

NOTE 1 If two or more tasks are to share the same generator, then the tasks have to synchronize their access to the generator as for any shared variable (see 12.10).

NOTE 2 Within a given implementation, a repeatable random number sequence can be obtained by relying on the implicit initialization of generators or by explicitly initializing a generator with a repeatable initiator value. Different sequences of random numbers can be obtained from a given generator in different program executions by explicitly initializing the generator to a time-dependent state.

EXAMPLE 1 *Example of a program that plays a simulated dice game:*

```

with Ada.Numerics.Discrete_Random;
procedure Dice_Game is
  subtype Die is Integer range 1 .. 6;
  subtype Dice is Integer range 2*Die'First .. 2*Die'Last;
  package Random_Die is new Ada.Numerics.Discrete_Random (Die);
  use Random_Die;
  G : Generator;
  D : Dice;
begin
  Reset (G); -- Start the generator in a unique state in each run
  loop
    -- Roll a pair of dice; sum and process the results
    D := Random(G) + Random(G);
    ...
  end loop;
end Dice_Game;

```

EXAMPLE 2 *Example of a program that simulates coin tosses:*

```

with Ada.Numerics.Discrete_Random;
procedure Flip_A_Coin is
  type Coin is (Heads, Tails);
  package Random_Coin is new Ada.Numerics.Discrete_Random (Coin);
  use Random_Coin;
  G : Generator;
begin
  Reset (G); -- Start the generator in a unique state in each run
  loop
    -- Toss a coin and process the result
    case Random(G) is
      when Heads =>
        ...
      when Tails =>
        ...
    end case;
    ...
  end loop;
end Flip_A_Coin;

```

EXAMPLE 3 *Example of a parallel simulation of a physical system, with a separate generator of event probabilities in each task:*

```
with Ada.Numerics.Float_Random;
procedure Parallel_Simulation is
  use Ada.Numerics.Float_Random;
  task type Worker is
    entry Initialize_Generator (Initiator : in Integer);
    ...
  end Worker;
  W : array (1 .. 10) of Worker;
  task body Worker is
    G : Generator;
    Probability_Of_Event : Uniformly_Distributed;
  begin
    accept Initialize_Generator (Initiator : in Integer) do
      Reset (G, Initiator);
    end Initialize_Generator;
    loop
      ...
      Probability_Of_Event := Random(G);
      ...
    end loop;
  end Worker;
begin
  -- Initialize the generators in the Worker tasks to different states
  for I in W'Range loop
    W(I).Initialize_Generator (I);
  end loop;
  ... -- Wait for the Worker tasks to terminate
end Parallel_Simulation;
```

Although each Worker task initializes its generator to a different state, those states will be the same in every execution of the program. The generator states can be initialized uniquely in each program execution by instantiating `Ada.Numerics.Discrete_Random` for the type `Integer` in the main procedure, resetting the generator obtained from that instance to a time-dependent state, and then using random integers obtained from that generator to initialize the generators in each Worker task.

### A.5.3 Attributes of Floating Point Types

#### Static Semantics

The following *representation-oriented attributes* are defined for every subtype *S* of a floating point type *T*.

#### S'Machine\_Radix

Yields the radix of the hardware representation of the type *T*. The value of this attribute is of the type *universal\_integer*.

The values of other representation-oriented attributes of a floating point subtype, and of the “primitive function” attributes of a floating point subtype described later, are defined in terms of a particular representation of nonzero values called the *canonical form*. The canonical form (for the type *T*) is the form

$$\pm \text{mantissa} \cdot T'\text{Machine\_Radix}^{\text{exponent}}$$

where

- *mantissa* is a fraction in the number base *T'Machine\_Radix*, the first digit of which is nonzero, and
- *exponent* is an integer.

#### S'Machine\_Mantissa

Yields the largest value of *p* such that every value expressible in the canonical form (for the type *T*), having a *p*-digit *mantissa* and an *exponent* between *T'Machine\_Emin* and *T'Machine\_Emax*, is a machine number (see 6.5.7) of the type *T*. This attribute yields a value of the type *universal\_integer*.

**S'Machine\_Emin**

Yields the smallest (most negative) value of *exponent* such that every value expressible in the canonical form (for the type *T*), having a *mantissa* of *T'Machine\_Mantissa* digits, is a machine number (see 6.5.7) of the type *T*. This attribute yields a value of the type *universal\_integer*.

**S'Machine\_Emax**

Yields the largest (most positive) value of *exponent* such that every value expressible in the canonical form (for the type *T*), having a *mantissa* of *T'Machine\_Mantissa* digits, is a machine number (see 6.5.7) of the type *T*. This attribute yields a value of the type *universal\_integer*.

**S'Denorm** Yields the value True if every value expressible in the form

$$\pm \textit{mantissa} \cdot T'\textit{Machine\_Radix}^{T'\textit{Machine\_Emin}}$$

where *mantissa* is a nonzero *T'Machine\_Mantissa*-digit fraction in the number base *T'Machine\_Radix*, the first digit of which is zero, is a machine number (see 6.5.7) of the type *T*; yields the value False otherwise. The value of this attribute is of the predefined type Boolean.

The values described by the formula in the definition of S'Denorm are called *denormalized numbers*. A nonzero machine number that is not a denormalized number is a *normalized number*. A normalized number *x* of a given type *T* is said to be *represented in canonical form* when it is expressed in the canonical form (for the type *T*) with a *mantissa* having *T'Machine\_Mantissa* digits; the resulting form is the *canonical-form representation* of *x*.

**S'Machine\_Rounds**

Yields the value True if rounding is performed on inexact results of every predefined operation that yields a result of the type *T*; yields the value False otherwise. The value of this attribute is of the predefined type Boolean.

**S'Machine\_Overflows**

Yields the value True if overflow and divide-by-zero are detected and reported by raising Constraint\_Error for every predefined operation that yields a result of the type *T*; yields the value False otherwise. The value of this attribute is of the predefined type Boolean.

**S'Signed\_Zeros**

Yields the value True if the hardware representation for the type *T* has the capability of representing both positively and negatively signed zeros, these being generated and used by the predefined operations of the type *T* as specified in IEC 559:1989; yields the value False otherwise. The value of this attribute is of the predefined type Boolean.

For every value *x* of a floating point type *T*, the *normalized exponent* of *x* is defined as follows:

- the normalized exponent of zero is (by convention) zero;
- for nonzero *x*, the normalized exponent of *x* is the unique integer *k* such that  $T'\textit{Machine\_Radix}^{k-1} \leq |x| < T'\textit{Machine\_Radix}^k$ .

The following *primitive function attributes* are defined for any subtype *S* of a floating point type *T*.

**S'Exponent** S'Exponent denotes a function with the following specification:

```
function S'Exponent (X : T)
return universal_integer
```

The function yields the normalized exponent of *X*.

**S'Fraction** S'Fraction denotes a function with the following specification:

```
function S'Fraction (X : T)
return T
```

The function yields the value  $X \cdot T'\textit{Machine\_Radix}^{-k}$ , where *k* is the normalized exponent of *X*. A zero result, which can only occur when *X* is zero, has the sign of *X*.

S'Compose S'Compose denotes a function with the following specification:

```
function S'Compose (Fraction : T;  
                    Exponent : universal_integer)  
return T
```

Let  $v$  be the value  $Fraction \cdot T'Machine\_Radix^{Exponent-k}$ , where  $k$  is the normalized exponent of  $Fraction$ . If  $v$  is a machine number of the type  $T$ , or if  $|v| \geq T'Model\_Small$ , the function yields  $v$ ; otherwise, it yields either one of the machine numbers of the type  $T$  adjacent to  $v$ . *Constraint\_Error* is optionally raised if  $v$  is outside the base range of  $S$ . A zero result has the sign of  $Fraction$  when *S'Signed\_Zeros* is True.

S'Scaling S'Scaling denotes a function with the following specification:

```
function S'Scaling (X : T;  
                   Adjustment : universal_integer)  
return T
```

Let  $v$  be the value  $X \cdot T'Machine\_Radix^{Adjustment}$ . If  $v$  is a machine number of the type  $T$ , or if  $|v| \geq T'Model\_Small$ , the function yields  $v$ ; otherwise, it yields either one of the machine numbers of the type  $T$  adjacent to  $v$ . *Constraint\_Error* is optionally raised if  $v$  is outside the base range of  $S$ . A zero result has the sign of  $X$  when *S'Signed\_Zeros* is True.

S'Floor S'Floor denotes a function with the following specification:

```
function S'Floor (X : T)  
return T
```

The function yields the value  $\lfloor X \rfloor$ , i.e. the largest (most positive) integral value less than or equal to  $X$ . When  $X$  is zero, the result has the sign of  $X$ ; a zero result otherwise has a positive sign.

S'Ceiling S'Ceiling denotes a function with the following specification:

```
function S'Ceiling (X : T)  
return T
```

The function yields the value  $\lceil X \rceil$ , i.e. the smallest (most negative) integral value greater than or equal to  $X$ . When  $X$  is zero, the result has the sign of  $X$ ; a zero result otherwise has a negative sign when *S'Signed\_Zeros* is True.

S'Rounding S'Rounding denotes a function with the following specification:

```
function S'Rounding (X : T)  
return T
```

The function yields the integral value nearest to  $X$ , rounding away from zero if  $X$  lies exactly halfway between two integers. A zero result has the sign of  $X$  when *S'Signed\_Zeros* is True.

S'Unbiased\_Rounding

S'Unbiased\_Rounding denotes a function with the following specification:

```
function S'Unbiased_Rounding (X : T)  
return T
```

The function yields the integral value nearest to  $X$ , rounding toward the even integer if  $X$  lies exactly halfway between two integers. A zero result has the sign of  $X$  when *S'Signed\_Zeros* is True.

S'Machine\_Rounding

S'Machine\_Rounding denotes a function with the following specification:

```
function S'Machine_Rounding (X : T)  
return T
```

The function yields the integral value nearest to  $X$ . If  $X$  lies exactly halfway between two integers, one of those integers is returned, but which of them is returned is unspecified. A zero result has the sign of  $X$  when *S'Signed\_Zeros* is True. This function provides access to the rounding behavior which is most efficient on the target processor.

**S'Truncation**

S'Truncation denotes a function with the following specification:

```
function S'Truncation (X : T)
  return T
```

The function yields the value  $\lceil X \rceil$  when  $X$  is negative, and  $\lfloor X \rfloor$  otherwise. A zero result has the sign of  $X$  when S'Signed\_Zeros is True.

**S'Remainder**

S'Remainder denotes a function with the following specification:

```
function S'Remainder (X, Y : T)
  return T
```

For nonzero  $Y$ , let  $v$  be the value  $X - n \cdot Y$ , where  $n$  is the integer nearest to the exact value of  $X/Y$ ; if  $|n - X/Y| = 1/2$ , then  $n$  is chosen to be even. If  $v$  is a machine number of the type  $T$ , the function yields  $v$ ; otherwise, it yields zero. Constraint\_Error is raised if  $Y$  is zero. A zero result has the sign of  $X$  when S'Signed\_Zeros is True.

**S'Adjacent** S'Adjacent denotes a function with the following specification:

```
function S'Adjacent (X, Towards : T)
  return T
```

If  $Towards = X$ , the function yields  $X$ ; otherwise, it yields the machine number of the type  $T$  adjacent to  $X$  in the direction of  $Towards$ , if that machine number exists. If the result would be outside the base range of  $S$ , Constraint\_Error is raised. When  $T$ 'Signed\_Zeros is True, a zero result has the sign of  $X$ . When  $Towards$  is zero, its sign has no bearing on the result.

**S'Copy\_Sign**

S'Copy\_Sign denotes a function with the following specification:

```
function S'Copy_Sign (Value, Sign : T)
  return T
```

If the value of  $Value$  is nonzero, the function yields a result whose magnitude is that of  $Value$  and whose sign is that of  $Sign$ ; otherwise, it yields the value zero. Constraint\_Error is optionally raised if the result is outside the base range of  $S$ . A zero result has the sign of  $Sign$  when S'Signed\_Zeros is True.

**S'Leading\_Part**

S'Leading\_Part denotes a function with the following specification:

```
function S'Leading_Part (X : T;
                        Radix_Digits : universal_integer)
  return T
```

Let  $v$  be the value  $T$ 'Machine\_Radix <sup>$k - Radix\_Digits$</sup> , where  $k$  is the normalized exponent of  $X$ . The function yields the value

- $\lfloor X/v \rfloor \cdot v$ , when  $X$  is nonnegative and  $Radix\_Digits$  is positive;
- $\lceil X/v \rceil \cdot v$ , when  $X$  is negative and  $Radix\_Digits$  is positive.

Constraint\_Error is raised when  $Radix\_Digits$  is zero or negative. A zero result, which can only occur when  $X$  is zero, has the sign of  $X$ .

**S'Machine** S'Machine denotes a function with the following specification:

```
function S'Machine (X : T)
  return T
```

If  $X$  is a machine number of the type  $T$ , the function yields  $X$ ; otherwise, it yields the value obtained by rounding or truncating  $X$  to either one of the adjacent machine numbers of the type  $T$ . Constraint\_Error is raised if rounding or truncating  $X$  to the precision of the machine numbers results in a value outside the base range of  $S$ . A zero result has the sign of  $X$  when S'Signed\_Zeros is True.



The following *model-oriented attributes* are defined for any subtype  $S$  of a floating point type  $T$ .

#### S'Model\_Mantissa

If the Numerics Annex is not supported, this attribute yields an implementation defined value that is greater than or equal to  $\lceil d \cdot \log(10) / \log(T'Machine\_Radix) \rceil + 1$ , where  $d$  is the requested decimal precision of  $T$ , and less than or equal to the value of  $T'Machine\_Mantissa$ . See G.2.2 for further requirements that apply to implementations supporting the Numerics Annex. The value of this attribute is of the type *universal\_integer*.

#### S'Model\_Emin

If the Numerics Annex is not supported, this attribute yields an implementation defined value that is greater than or equal to the value of  $T'Machine\_Emin$ . See G.2.2 for further requirements that apply to implementations supporting the Numerics Annex. The value of this attribute is of the type *universal\_integer*.

#### S'Model\_Epsilon

Yields the value  $T'Machine\_Radix^{1 - T'Model\_Mantissa}$ . The value of this attribute is of the type *universal\_real*.

#### S'Model\_Small

Yields the value  $T'Machine\_Radix^{T'Model\_Emin - 1}$ . The value of this attribute is of the type *universal\_real*.

**S'Model** S'Model denotes a function with the following specification:

```
function S'Model (X : T)
return T
```

If the Numerics Annex is not supported, the meaning of this attribute is implementation defined; see G.2.2 for the definition that applies to implementations supporting the Numerics Annex.

#### S'Safe\_First

Yields the lower bound of the safe range (see 6.5.7) of the type  $T$ . If the Numerics Annex is not supported, the value of this attribute is implementation defined; see G.2.2 for the definition that applies to implementations supporting the Numerics Annex. The value of this attribute is of the type *universal\_real*.

#### S'Safe\_Last

Yields the upper bound of the safe range (see 6.5.7) of the type  $T$ . If the Numerics Annex is not supported, the value of this attribute is implementation defined; see G.2.2 for the definition that applies to implementations supporting the Numerics Annex. The value of this attribute is of the type *universal\_real*.

### A.5.4 Attributes of Fixed Point Types

#### Static Semantics

The following *representation-oriented attributes* are defined for every subtype  $S$  of a fixed point type  $T$ .

#### S'Machine\_Radix

Yields the radix of the hardware representation of the type  $T$ . The value of this attribute is of the type *universal\_integer*.

#### S'Machine\_Rounds

Yields the value True if rounding is performed on inexact results of every predefined operation that yields a result of the type  $T$ ; yields the value False otherwise. The value of this attribute is of the predefined type Boolean.

#### S'Machine\_Overflows

Yields the value True if overflow and divide-by-zero are detected and reported by raising Constraint\_Error for every predefined operation that yields a result of the type  $T$ ; yields the value False otherwise. The value of this attribute is of the predefined type Boolean.

### A.5.5 Big Numbers

Support is provided for integer arithmetic involving values larger than those supported by the target machine, and for arbitrary-precision real numbers.

#### Static Semantics

The library package `Numerics.Big_Numbers` has the following declaration:

```
package Ada.Numerics.Big_Numbers
  with Pure, Nonblocking, Global => null is
  subtype Field is Integer range 0 .. implementation-defined;
  subtype Number_Base is Integer range 2 .. 16;
end Ada.Numerics.Big_Numbers;
```

### A.5.6 Big Integers

#### Static Semantics

The library package `Numerics.Big_Numbers.Big_Integers` has the following declaration:

```
with Ada.Strings.Text_Buffers;
package Ada.Numerics.Big_Numbers.Big_Integers
  with Preelaborate, Nonblocking, Global => in out synchronized is

  type Big_Integer is private
    with Integer_Literal => From_Universal_Image,
        Put_Image => Put_Image;

  function Is_Valid (Arg : Big_Integer) return Boolean
    with Convention => Intrinsic;

  subtype Valid_Big_Integer is Big_Integer
    with Dynamic_Predicate => Is_Valid (Valid_Big_Integer),
        Predicate_Failure => (raise Program_Error);

  function "=" (L, R : Valid_Big_Integer) return Boolean;
  function "<" (L, R : Valid_Big_Integer) return Boolean;
  function "<=" (L, R : Valid_Big_Integer) return Boolean;
  function ">" (L, R : Valid_Big_Integer) return Boolean;
  function ">=" (L, R : Valid_Big_Integer) return Boolean;

  function To_Big_Integer (Arg : Integer) return Valid_Big_Integer;

  subtype Big_Positive is Big_Integer
    with Dynamic_Predicate => (if Is_Valid (Big_Positive)
                              then Big_Positive > 0),
        Predicate_Failure => (raise Constraint_Error);

  subtype Big_Natural is Big_Integer
    with Dynamic_Predicate => (if Is_Valid (Big_Natural)
                              then Big_Natural >= 0),
        Predicate_Failure => (raise Constraint_Error);

  function In_Range (Arg, Low, High : Valid_Big_Integer) return Boolean is
    (Low <= Arg and Arg <= High);

  function To_Integer (Arg : Valid_Big_Integer) return Integer
    with Pre => In_Range (Arg,
                        Low => To_Big_Integer (Integer'First),
                        High => To_Big_Integer (Integer'Last))
    or else raise Constraint_Error;
```

```

generic
  type Int is range <>;
package Signed_Conversions is
  function To_Big_Integer (Arg : Int) return Valid_Big_Integer;
  function From_Big_Integer (Arg : Valid_Big_Integer) return Int
    with Pre => In_Range (Arg,
      Low  => To_Big_Integer (Int'First),
      High => To_Big_Integer (Int'Last))
    or else raise Constraint_Error;
end Signed_Conversions;

generic
  type Int is mod <>;
package Unsigned_Conversions is
  function To_Big_Integer (Arg : Int) return Valid_Big_Integer;
  function From_Big_Integer (Arg : Valid_Big_Integer) return Int
    with Pre => In_Range (Arg,
      Low  => To_Big_Integer (Int'First),
      High => To_Big_Integer (Int'Last))
    or else raise Constraint_Error;
end Unsigned_Conversions;

function To_String (Arg : Valid_Big_Integer;
  Width : Field := 0;
  Base : Number_Base := 10) return String
  with Post => To_String'Result'First = 1;

function From_String (Arg : String) return Valid_Big_Integer;

function From_Universal_Image (Arg : String) return Valid_Big_Integer
  renames From_String;

procedure Put_Image
  (Buffer : in out Ada.Strings.Text_Buffers.Root_Buffer_Type'Class;
  Arg : in Valid_Big_Integer);

function "+" (L : Valid_Big_Integer) return Valid_Big_Integer;
function "-" (L : Valid_Big_Integer) return Valid_Big_Integer;
function "abs" (L : Valid_Big_Integer) return Valid_Big_Integer;
function "+" (L, R : Valid_Big_Integer) return Valid_Big_Integer;
function "-" (L, R : Valid_Big_Integer) return Valid_Big_Integer;
function "*" (L, R : Valid_Big_Integer) return Valid_Big_Integer;
function "/" (L, R : Valid_Big_Integer) return Valid_Big_Integer;
function "mod" (L, R : Valid_Big_Integer) return Valid_Big_Integer;
function "rem" (L, R : Valid_Big_Integer) return Valid_Big_Integer;
function "***" (L : Valid_Big_Integer; R : Natural)
  return Valid_Big_Integer;
function Min (L, R : Valid_Big_Integer) return Valid_Big_Integer;
function Max (L, R : Valid_Big_Integer) return Valid_Big_Integer;

function Greatest_Common_Divisor
  (L, R : Valid_Big_Integer) return Big_Positive
  with Pre => (L /= 0 and R /= 0) or else raise Constraint_Error;

private
  ... -- not specified by the language
end Ada.Numerics.Big_Numbers.Big_Integers;

```

To\_String and From\_String behave analogously to the Put and Get procedures defined in Text\_IO.Integer\_IO (in particular, with respect to the interpretation of the Width and Base parameters) except that Constraint\_Error, not Data\_Error, is propagated in error cases and the result of a call to To\_String with a Width parameter of 0 and a nonnegative Arg parameter does not include a leading blank. Put\_Image calls To\_String (passing in the default values for the Width and Base parameters), prepends a leading blank if the argument is nonnegative, and writes the resulting value to the buffer using Text\_Buffers.Put.

The other functions have their usual mathematical meanings.

The type Big\_Integer needs finalization (see 10.6).

*Dynamic Semantics*

For purposes of determining whether predicate checks are performed as part of default initialization, the type `Big_Integer` is considered to have a subcomponent that has a `default_expression`.

*Implementation Requirements*

No storage associated with a `Big_Integer` object shall be lost upon assignment or scope exit.

**A.5.7 Big Reals***Static Semantics*

The library package `Numerics.Big_Numbers.Big_Reals` has the following declaration:

```

with Ada.Numerics.Big_Numbers.Big_Integers;
use all type Big_Integers.Big_Integer;
with Ada.Strings.Text_Buffers;
package Ada.Numerics.Big_Numbers.Big_Reals
with Preelaborate, Nonblocking, Global => in out synchronized is
  type Big_Real is private
    with Real_Literal => From_Universal_Image,
         Put_Image => Put_Image;
  function Is_Valid (Arg : Big_Real) return Boolean
    with Convention => Intrinsic;
  subtype Valid_Big_Real is Big_Real
    with Dynamic_Predicate => Is_Valid (Valid_Big_Real),
         Predicate_Failure => raise Program_Error;
  function "/" (Num, Den : Big_Integers.Valid_Big_Integer)
    return Valid_Big_Real
    with Pre => Den /= 0
        or else raise Constraint_Error;
  function Numerator
    (Arg : Valid_Big_Real) return Big_Integers.Valid_Big_Integer
    with Post => (if Arg = 0.0 then Numerator'Result = 0);
  function Denominator (Arg : Valid_Big_Real)
    return Big_Integers.Big_Positive
    with Post =>
      (if Arg = 0.0 then Denominator'Result = 1
       else Big_Integers.Greatest_Common_Divisor
        (Numerator (Arg), Denominator'Result) = 1);
  function To_Big_Real (Arg : Big_Integers.Valid_Big_Integer)
    return Valid_Big_Real is (Arg / 1);
  function To_Real (Arg : Integer) return Valid_Big_Real is
    (Big_Integers.To_Big_Integer (Arg) / 1);
  function "=" (L, R : Valid_Big_Real) return Boolean;
  function "<" (L, R : Valid_Big_Real) return Boolean;
  function "<=" (L, R : Valid_Big_Real) return Boolean;
  function ">" (L, R : Valid_Big_Real) return Boolean;
  function ">=" (L, R : Valid_Big_Real) return Boolean;
  function In_Range (Arg, Low, High : Valid_Big_Real) return Boolean is
    (Low <= Arg and Arg <= High);

```

```

generic
  type Num is digits <>;
package Float_Conversions is
  function To_Big_Real (Arg : Num) return Valid_Big_Real;
  function From_Big_Real (Arg : Valid_Big_Real) return Num
    with Pre => In_Range (Arg,
                        Low  => To_Big_Real (Num'First),
                        High => To_Big_Real (Num'Last))
    or else (raise Constraint_Error);
end Float_Conversions;

generic
  type Num is delta <>;
package Fixed_Conversions is
  function To_Big_Real (Arg : Num) return Valid_Big_Real;
  function From_Big_Real (Arg : Valid_Big_Real) return Num
    with Pre => In_Range (Arg,
                        Low  => To_Big_Real (Num'First),
                        High => To_Big_Real (Num'Last))
    or else (raise Constraint_Error);
end Fixed_Conversions;

function To_String (Arg : Valid_Big_Real;
                   Fore : Field := 2;
                   Aft  : Field := 3;
                   Exp  : Field := 0) return String
  with Post => To_String'Result'First = 1;

function From_String (Arg : String) return Valid_Big_Real;

function From_Universal_Image (Arg : String) return Valid_Big_Real
  renames From_String;

function From_Universal_Image (Num, Den : String)
  return Valid_Big_Real is
  (Big_Integers.From_Universal_Image (Num) /
   Big_Integers.From_Universal_Image (Den));

function To_Quotient_String (Arg : Valid_Big_Real) return String is
  (To_String (Numerator (Arg)) & " / " & To_String (Denominator (Arg)));
function From_Quotient_String (Arg : String) return Valid_Big_Real;

procedure Put_Image
  (Buffer : in out Ada.Strings.Text_Buffers.Root_Buffer_Type'Class;
   Arg    : in Valid_Big_Real);

function "+" (L : Valid_Big_Real) return Valid_Big_Real;
function "-" (L : Valid_Big_Real) return Valid_Big_Real;
function "abs" (L : Valid_Big_Real) return Valid_Big_Real;
function "+" (L, R : Valid_Big_Real) return Valid_Big_Real;
function "-" (L, R : Valid_Big_Real) return Valid_Big_Real;
function "*" (L, R : Valid_Big_Real) return Valid_Big_Real;
function "/" (L, R : Valid_Big_Real) return Valid_Big_Real;
function "***" (L : Valid_Big_Real; R : Integer)
  return Valid_Big_Real;
function Min (L, R : Valid_Big_Real) return Valid_Big_Real;
function Max (L, R : Valid_Big_Real) return Valid_Big_Real;

private
  ... -- not specified by the language
end Ada.Numerics.Big_Numbers.Big_Reals;

```

To\_String and From\_String behave analogously to the Put and Get procedures defined in Text\_IO.Float\_IO (in particular, with respect to the interpretation of the Fore, Aft, and Exp parameters), except that Constraint\_Error (not Data\_Error) is propagated in error cases. From\_Quotient\_String implements the inverse function of To\_Quotient\_String; Constraint\_Error is propagated in error cases. Put\_Image calls To\_String, and writes the resulting value to the buffer using Text\_Buffers.Put.

For an instance of `Float_Conversions` or `Fixed_Conversions`, `To_Big_Real` is exact (that is, the result represents exactly the same mathematical value as the argument) and `From_Big_Real` is subject to the same precision rules as a type conversion of a value of type `T` to the target type `Num`, where `T` is a hypothetical floating point type whose model numbers include all of the model numbers of `Num` as well as the exact mathematical value of the argument.

The other functions have their usual mathematical meanings.

The type `Big_Real` needs finalization (see 10.6).

#### Dynamic Semantics

For purposes of determining whether predicate checks are performed as part of default initialization, the type `Big_Real` is considered to have a subcomponent that has a `default_expression`.

#### Implementation Requirements

No storage associated with a `Big_Real` object shall be lost upon assignment or scope exit.

## A.6 Input-Output

Input-output is provided through language-defined packages, each of which is a child of the root package `Ada`. The generic packages `Sequential_IO` and `Direct_IO` define input-output operations applicable to files containing elements of a given type. The generic package `Storage_IO` supports reading from and writing to an in-memory buffer. Additional operations for text input-output are supplied in the packages `Text_IO`, `Wide_Text_IO`, and `Wide_Wide_Text_IO`. Heterogeneous input-output is provided through the child packages `Streams.Stream_IO` and `Text_IO.Text_Streams` (see also 16.13). The package `IO_Exceptions` defines the exceptions used by the predefined input-output packages.

## A.7 External Files and File Objects

#### Static Semantics

Values input from the external environment of the program, or output to the external environment, are considered to occupy *external files*. An external file can be anything external to the program that can produce a value to be read or receive a value to be written. An external file is identified by a string (the *name*). A second string (the *form*) gives further system-dependent characteristics that may be associated with the file, such as the physical organization or access rights. The conventions governing the interpretation of such strings shall be documented.

Input and output operations are expressed as operations on objects of some *file type*, rather than directly in terms of the external files. In the remainder of this clause, the term *file* is always used to refer to a file object; the term *external file* is used otherwise.

Input-output for sequential files of values of a single element type is defined by means of the generic package `Sequential_IO`. In order to define sequential input-output for a given element type, an instantiation of this generic unit, with the given type as actual parameter, has to be declared. The resulting package contains the declaration of a file type (called `File_Type`) for files of such elements, as well as the operations applicable to these files, such as the `Open`, `Read`, and `Write` procedures.

Input-output for direct access files is likewise defined by a generic package called `Direct_IO`. Input-output in human-readable form is defined by the (nongeneric) packages `Text_IO` for `Character` and `String` data, `Wide_Text_IO` for `Wide_Character` and `Wide_String` data, and `Wide_Wide_Text_IO` for `Wide_Wide_Character` and `Wide_Wide_String` data. Input-output for files containing streams of elements representing values of possibly different types is defined by means of the (nongeneric) package `Streams.Stream_IO`.

Before input or output operations can be performed on a file, the file first has to be associated with an external file. While such an association is in effect, the file is said to be *open*, and otherwise the file is said to be *closed*.



The language does not define what happens to external files after the completion of the main program and all the library tasks (in particular, if corresponding files have not been closed). The effect of input-output for access types is unspecified.

An open file has a *current mode*, which is a value of one of the following enumeration types:

```
type File_Mode is (In_File, Inout_File, Out_File);  -- for Direct_IO
```

These values correspond respectively to the cases where only reading, both reading and writing, or only writing are to be performed.

```
type File_Mode is (In_File, Out_File, Append_File);  
-- for Sequential_IO, Text_IO, Wide_Text_IO, Wide_Wide_Text_IO, and Stream_IO
```

These values correspond respectively to the cases where only reading, only writing, or only appending are to be performed.

The mode of a file can be changed.

Several file management operations are common to Sequential\_IO, Direct\_IO, Text\_IO, Wide\_Text\_IO, and Wide\_Wide\_Text\_IO. These operations are described in A.8.2 for sequential and direct files. Any additional effects concerning text input-output are described in A.10.2.

The exceptions that can be propagated by the execution of an input-output subprogram are defined in the package IO\_Exceptions; the situations in which they can be propagated are described following the description of the subprogram (and in A.13). The exceptions Storage\_Error and Program\_Error may be propagated. (Program\_Error can only be propagated due to errors made by the caller of the subprogram.) Finally, exceptions can be propagated in certain implementation-defined situations.

NOTE 1 Each instantiation of the generic packages Sequential\_IO and Direct\_IO declares a different type File\_Type. In the case of Text\_IO, Wide\_Text\_IO, Wide\_Wide\_Text\_IO, and Streams.Stream\_IO, the corresponding type File\_Type is unique.

NOTE 2 A bidirectional device can often be modeled as two sequential files associated with the device, one of mode In\_File, and one of mode Out\_File. An implementation can restrict the number of files that can be associated with a given external file.

## A.8 Sequential and Direct Files

### Static Semantics

Two kinds of access to external files are defined in this subclause: *sequential access* and *direct access*. The corresponding file types and the associated operations are provided by the generic packages Sequential\_IO and Direct\_IO. A file object to be used for sequential access is called a *sequential file*, and one to be used for direct access is called a *direct file*. Access to *stream files* is described in A.12.1.

For sequential access, the file is viewed as a sequence of values that are transferred in the order of their appearance (as produced by the program or by the external environment). When the file is opened with mode In\_File or Out\_File, transfer starts respectively from or to the beginning of the file. When the file is opened with mode Append\_File, transfer to the file starts after the last element of the file.

For direct access, the file is viewed as a set of elements occupying consecutive positions in linear order; a value can be transferred to or from an element of the file at any selected position. The position of an element is specified by its *index*, which is a number, greater than zero, of the implementation-defined integer type Count. The first element, if any, has index one; the index of the last element, if any, is called the *current size*; the current size is zero if there are no elements. The current size is a property of the external file.

An open direct file has a *current index*, which is the index that will be used by the next read or write operation. When a direct file is opened, the current index is set to one. The current index of a direct file is a property of a file object, not of an external file.

### A.8.1 The Generic Package Sequential\_IO

#### Static Semantics

The generic library package Sequential\_IO has the following declaration:

```

with Ada.IO_Exceptions;
generic
  type Element_Type(<>) is private;
package Ada.Sequential_IO
  with Global => in out synchronized is

  type File_Type is limited private;

  type File_Mode is (In_File, Out_File, Append_File);

  -- File management

  procedure Create(File : in out File_Type;
                  Mode : in File_Mode := Out_File;
                  Name : in String := "";
                  Form : in String := "");

  procedure Open  (File : in out File_Type;
                  Mode : in File_Mode;
                  Name : in String;
                  Form : in String := "");

  procedure Close (File : in out File_Type);
  procedure Delete(File : in out File_Type);
  procedure Reset (File : in out File_Type; Mode : in File_Mode);
  procedure Reset (File : in out File_Type);

  function Mode   (File : in File_Type) return File_Mode;
  function Name   (File : in File_Type) return String;
  function Form   (File : in File_Type) return String;

  function Is_Open(File : in File_Type) return Boolean;

  procedure Flush (File : in File_Type)
    with Global => overriding in out File;

  -- Input and output operations

  procedure Read  (File : in File_Type; Item : out Element_Type)
    with Global => overriding in out File;
  procedure Write (File : in File_Type; Item : in Element_Type)
    with Global => overriding in out File;

  function End_Of_File(File : in File_Type) return Boolean;

  -- Exceptions

  Status_Error : exception renames IO_Exceptions.Status_Error;
  Mode_Error   : exception renames IO_Exceptions.Mode_Error;
  Name_Error   : exception renames IO_Exceptions.Name_Error;
  Use_Error    : exception renames IO_Exceptions.Use_Error;
  Device_Error : exception renames IO_Exceptions.Device_Error;
  End_Error    : exception renames IO_Exceptions.End_Error;
  Data_Error   : exception renames IO_Exceptions.Data_Error;

  package Wide_File_Names is
    -- File management

    procedure Create(File : in out File_Type;
                    Mode : in File_Mode := Out_File;
                    Name : in Wide_String := "";
                    Form : in Wide_String := "");

    procedure Open  (File : in out File_Type;
                    Mode : in File_Mode;
                    Name : in Wide_String;
                    Form : in Wide_String := "");

```

```

function Name      (File : in File_Type) return Wide_String;
function Form      (File : in File_Type) return Wide_String;
end Wide_File_Names;
package Wide_Wide_File_Names is
    -- File management

    procedure Create (File : in out File_Type;
                     Mode : in File_Mode := Out_File;
                     Name : in Wide_Wide_String := "";
                     Form : in Wide_Wide_String := "");

    procedure Open   (File : in out File_Type;
                     Mode : in File_Mode;
                     Name : in Wide_Wide_String;
                     Form : in Wide_Wide_String := "");

    function Name     (File : in File_Type) return Wide_Wide_String;
    function Form     (File : in File_Type) return Wide_Wide_String;
end Wide_Wide_File_Names;

private
    ... -- not specified by the language
end Ada.Sequential_IO;

```

The type File\_Type needs finalization (see 10.6) in every instantiation of Sequential\_IO.

## A.8.2 File Management

### Static Semantics

The procedures and functions described in this subclause provide for the control of external files; their declarations are repeated in each of the packages for sequential, direct, text, and stream input-output. For text input-output, the procedures Create, Open, and Reset have additional effects described in A.10.2.

```

procedure Create (File : in out File_Type;
                  Mode : in File_Mode := default_mode;
                  Name : in String := "";
                  Form : in String := "");

```

Establishes a new external file, with the given name and form, and associates this external file with the given file. The given file is left open. The current mode of the given file is set to the given access mode. The default access mode is the mode Out\_File for sequential, stream, and text input-output; it is the mode Inout\_File for direct input-output. For direct access, the size of the created file is implementation defined.

A null string for Name specifies an external file that is not accessible after the completion of the main program (a temporary file). A null string for Form specifies the use of the default options of the implementation for the external file.

The exception Status\_Error is propagated if the given file is already open. The exception Name\_Error is propagated if the string given as Name does not allow the identification of an external file. The exception Use\_Error is propagated if, for the specified mode, the external environment does not support creation of an external file with the given name (in the absence of Name\_Error) and form.

```

procedure Open (File : in out File_Type;
                Mode : in File_Mode;
                Name : in String;
                Form : in String := "");

```

Associates the given file with an existing external file having the given name and form, and sets the current mode of the given file to the given mode. The given file is left open.

The exception `Status_Error` is propagated if the given file is already open. The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an external file; in particular, this exception is propagated if no external file with the given name exists. The exception `Use_Error` is propagated if, for the specified mode, the external environment does not support opening for an external file with the given name (in the absence of `Name_Error`) and form.

**procedure** `Close`(`File` : **in out** `File_Type`);

Severs the association between the given file and its associated external file. The given file is left closed. In addition, for sequential files, if the file being closed has mode `Out_File` or `Append_File`, then the last element written since the most recent open or reset is the last element that can be read from the file. If no elements have been written and the file mode is `Out_File`, then the closed file is empty. If no elements have been written and the file mode is `Append_File`, then the closed file is unchanged.

The exception `Status_Error` is propagated if the given file is not open.

**procedure** `Delete`(`File` : **in out** `File_Type`);

Deletes the external file associated with the given file. The given file is closed, and the external file ceases to exist.

The exception `Status_Error` is propagated if the given file is not open. The exception `Use_Error` is propagated if deletion of the external file is not supported by the external environment.

**procedure** `Reset`(`File` : **in out** `File_Type`; `Mode` : **in** `File_Mode`);

**procedure** `Reset`(`File` : **in out** `File_Type`);

Resets the given file so that reading from its elements can be restarted from the beginning of the external file (for modes `In_File` and `Inout_File`), and so that writing to its elements can be restarted at the beginning of the external file (for modes `Out_File` and `Inout_File`) or after the last element of the external file (for mode `Append_File`). In particular, for direct access this means that the current index is set to one. If a `Mode` parameter is supplied, the current mode of the given file is set to the given mode. In addition, for sequential files, if the given file has mode `Out_File` or `Append_File` when `Reset` is called, the last element written since the most recent open or reset is the last element that can be read from the external file. If no elements have been written and the file mode is `Out_File`, the reset file is empty. If no elements have been written and the file mode is `Append_File`, then the reset file is unchanged.

The exception `Status_Error` is propagated if the file is not open. The exception `Use_Error` is propagated if the external environment does not support resetting for the external file and, also, if the external environment does not support resetting to the specified mode for the external file.

**function** `Mode`(`File` : **in** `File_Type`) **return** `File_Mode`;

Returns the current mode of the given file.

The exception `Status_Error` is propagated if the file is not open.

**function** `Name`(`File` : **in** `File_Type`) **return** `String`;

Returns a string which uniquely identifies the external file currently associated with the given file (and may thus be used in an `Open` operation).

The exception `Status_Error` is propagated if the given file is not open. The exception `Use_Error` is propagated if the associated external file is a temporary file that cannot be opened by any name.

**function** `Form`(`File` : **in** `File_Type`) **return** `String`;

Returns the form string for the external file currently associated with the given file. If an external environment allows alternative specifications of the form (for example, abbreviations

using default options), the string returned by the function should correspond to a full specification (that is, it should indicate explicitly all options selected, including default options).

The exception `Status_Error` is propagated if the given file is not open.

**function** `Is_Open`(`File` : **in** `File_Type`) **return** `Boolean`;

Returns `True` if the file is open (that is, if it is associated with an external file); otherwise, returns `False`.

**procedure** `Flush`(`File` : **in** `File_Type`);

The `Flush` procedure synchronizes the external file with the internal file (by flushing any internal buffers) without closing the file. For a direct file, the current index is unchanged; for a stream file (see A.12.1), the current position is unchanged.

The exception `Status_Error` is propagated if the file is not open. The exception `Mode_Error` is propagated if the mode of the file is `In_File`.

The nested package `Wide_File_Names` provides operations equivalent to the operations of the same name of the outer package except that `Wide_String` is used instead of `String` for the name and form of the external file.

The nested package `Wide_Wide_File_Names` provides operations equivalent to the operations of the same name of the outer package except that `Wide_Wide_String` is used instead of `String` for the name and form of the external file.

#### *Implementation Permissions*

An implementation may propagate `Name_Error` or `Use_Error` if an attempt is made to use an I/O feature that cannot be supported by the implementation due to limitations in the external environment. Any such restriction should be documented.

### **A.8.3 Sequential Input-Output Operations**

#### *Static Semantics*

The operations available for sequential input and output are described in this subclause. The exception `Status_Error` is propagated if any of these operations is attempted for a file that is not open.

**procedure** `Read`(`File` : **in** `File_Type`; `Item` : **out** `Element_Type`);

Operates on a file of mode `In_File`. Reads an element from the given file, and returns the value of this element in the `Item` parameter.

The exception `Mode_Error` is propagated if the mode is not `In_File`. The exception `End_Error` is propagated if no more elements can be read from the given file. The exception `Data_Error` can be propagated if the element read cannot be interpreted as a value of the subtype `Element_Type` (see A.13).

**procedure** `Write`(`File` : **in** `File_Type`; `Item` : **in** `Element_Type`);

Operates on a file of mode `Out_File` or `Append_File`. Writes the value of `Item` to the given file.

The exception `Mode_Error` is propagated if the mode is not `Out_File` or `Append_File`. The exception `Use_Error` is propagated if the capacity of the external file is exceeded.

**function** `End_Of_File`(`File` : **in** `File_Type`) **return** `Boolean`;

Operates on a file of mode `In_File`. Returns `True` if no more elements can be read from the given file; otherwise, returns `False`.

The exception `Mode_Error` is propagated if the mode is not `In_File`.

## A.8.4 The Generic Package Direct\_IO

### Static Semantics

The generic library package Direct\_IO has the following declaration:

```

with Ada.IO_Exceptions;
generic
  type Element_Type is private;
package Ada.Direct_IO
  with Global => in out synchronized is

  type File_Type is limited private;

  type File_Mode is (In_File, Inout_File, Out_File);
  type Count is range 0 .. implementation-defined;
  subtype Positive_Count is Count range 1 .. Count'Last;

  -- File management

  procedure Create(File : in out File_Type;
    Mode : in File_Mode := Inout_File;
    Name : in String := "";
    Form : in String := "");

  procedure Open (File : in out File_Type;
    Mode : in File_Mode;
    Name : in String;
    Form : in String := "");

  procedure Close (File : in out File_Type);
  procedure Delete(File : in out File_Type);
  procedure Reset (File : in out File_Type; Mode : in File_Mode);
  procedure Reset (File : in out File_Type);

  function Mode (File : in File_Type) return File_Mode;
  function Name (File : in File_Type) return String;
  function Form (File : in File_Type) return String;

  function Is_Open(File : in File_Type) return Boolean;

  procedure Flush (File : in File_Type)
    with Global => overriding in out File;

  -- Input and output operations

  procedure Read (File : in File_Type; Item : out Element_Type;
    From : in Positive_Count)
    with Global => overriding in out File;
  procedure Read (File : in File_Type; Item : out Element_Type)
    with Global => overriding in out File;

  procedure Write(File : in File_Type; Item : in Element_Type;
    To : in Positive_Count)
    with Global => overriding in out File;
  procedure Write(File : in File_Type; Item : in Element_Type)
    with Global => overriding in out File;

  procedure Set_Index(File : in File_Type; To : in Positive_Count)
    with Global => overriding in out File;

  function Index(File : in File_Type) return Positive_Count;
  function Size (File : in File_Type) return Count;

  function End_Of_File(File : in File_Type) return Boolean;

```



```

-- Exceptions
Status_Error : exception renames IO_Exceptions.Status_Error;
Mode_Error   : exception renames IO_Exceptions.Mode_Error;
Name_Error   : exception renames IO_Exceptions.Name_Error;
Use_Error    : exception renames IO_Exceptions.Use_Error;
Device_Error : exception renames IO_Exceptions.Device_Error;
End_Error    : exception renames IO_Exceptions.End_Error;
Data_Error   : exception renames IO_Exceptions.Data_Error;

package Wide_File_Names is
    -- File management
    procedure Create(File : in out File_Type;
                    Mode : in File_Mode := Inout_File;
                    Name : in Wide_String := "";
                    Form : in Wide_String := "");

    procedure Open  (File : in out File_Type;
                    Mode : in File_Mode;
                    Name : in Wide_String;
                    Form : in Wide_String := "");

    function Name   (File : in File_Type) return Wide_String;
    function Form   (File : in File_Type) return Wide_String;
end Wide_File_Names;

package Wide_Wide_File_Names is
    -- File management
    procedure Create(File : in out File_Type;
                    Mode : in File_Mode := Inout_File;
                    Name : in Wide_Wide_String := "";
                    Form : in Wide_Wide_String := "");

    procedure Open  (File : in out File_Type;
                    Mode : in File_Mode;
                    Name : in Wide_Wide_String;
                    Form : in Wide_Wide_String := "");

    function Name   (File : in File_Type) return Wide_Wide_String;
    function Form   (File : in File_Type) return Wide_Wide_String;
end Wide_Wide_File_Names;

private
    ... -- not specified by the language
end Ada.Direct_IO;

```

The type `File_Type` needs finalization (see 10.6) in every instantiation of `Direct_IO`.

## A.8.5 Direct Input-Output Operations

### Static Semantics

The operations available for direct input and output are described in this subclause. The exception `Status_Error` is propagated if any of these operations is attempted for a file that is not open.

```

procedure Read(File : in File_Type; Item : out Element_Type;
               From : in Positive_Count);
procedure Read(File : in File_Type; Item : out Element_Type);

```

Operates on a file of mode `In_File` or `Inout_File`. In the case of the first form, sets the current index of the given file to the index value given by the parameter `From`. Then (for both forms) returns, in the parameter `Item`, the value of the element whose position in the given file is specified by the current index of the file; finally, increases the current index by one.

The exception `Mode_Error` is propagated if the mode of the given file is `Out_File`. The exception `End_Error` is propagated if the index to be used exceeds the size of the external file. The exception `Data_Error` can be propagated if the element read cannot be interpreted as a value of the subtype `Element_Type` (see A.13).

```
procedure Write(File : in File_Type; Item : in Element_Type;
                To   : in Positive_Count);
procedure Write(File : in File_Type; Item : in Element_Type);
```

Operates on a file of mode `Inout_File` or `Out_File`. In the case of the first form, sets the index of the given file to the index value given by the parameter `To`. Then (for both forms) gives the value of the parameter `Item` to the element whose position in the given file is specified by the current index of the file; finally, increases the current index by one.

The exception `Mode_Error` is propagated if the mode of the given file is `In_File`. The exception `Use_Error` is propagated if the capacity of the external file is exceeded.

```
procedure Set_Index(File : in File_Type; To : in Positive_Count);
```

Operates on a file of any mode. Sets the current index of the given file to the given index value (which may exceed the current size of the file).

```
function Index(File : in File_Type) return Positive_Count;
```

Operates on a file of any mode. Returns the current index of the given file.

```
function Size(File : in File_Type) return Count;
```

Operates on a file of any mode. Returns the current size of the external file that is associated with the given file.

```
function End_Of_File(File : in File_Type) return Boolean;
```

Operates on a file of mode `In_File` or `Inout_File`. Returns `True` if the current index exceeds the size of the external file; otherwise, returns `False`.

The exception `Mode_Error` is propagated if the mode of the given file is `Out_File`.

NOTE Append\_File mode is not supported for the generic package `Direct_IO`.

## A.9 The Generic Package `Storage_IO`

The generic package `Storage_IO` provides for reading from and writing to an in-memory buffer. This generic package supports the construction of user-defined input-output packages.

### *Static Semantics*

The generic library package `Storage_IO` has the following declaration:

```
with Ada.IO_Exceptions;
with System.Storage_Elements;
generic
  type Element_Type is private;
package Ada.Storage_IO
  with Preelaborate, Global => in out synchronized is

  Buffer_Size : constant System.Storage_Elements.Storage_Count :=
    implementation-defined;
  subtype Buffer_Type is
    System.Storage_Elements.Storage_Array(1..Buffer_Size);

  -- Input and output operations

  procedure Read (Buffer : in Buffer_Type; Item : out Element_Type);
  procedure Write(Buffer : out Buffer_Type; Item : in Element_Type);
```

```
-- Exceptions
Data_Error : exception renames IO_Exceptions.Data_Error;
end Ada.Storage_IO;
```

In each instance, the constant `Buffer_Size` has a value that is the size (in storage elements) of the buffer required to represent the content of an object of subtype `Element_Type`, including any implicit levels of indirection used by the implementation. The `Read` and `Write` procedures of `Storage_IO` correspond to the `Read` and `Write` procedures of `Direct_IO` (see A.8.4), but with the content of the `Item` parameter being read from or written into the specified `Buffer`, rather than an external file.

NOTE A buffer used for `Storage_IO` holds only one element at a time; an external file used for `Direct_IO` holds a sequence of elements.

## A.10 Text Input-Output

### Static Semantics

This subclause describes the package `Text_IO`, which provides facilities for input and output in human-readable form. Each file is read or written sequentially, as a sequence of characters grouped into lines, and as a sequence of lines grouped into pages. The specification of the package is given below in A.10.1.

The facilities for file management given above, in A.8.2 and A.8.3, are available for text input-output. In place of `Read` and `Write`, however, there are procedures `Get` and `Put` that input values of suitable types from text files, and output values to them. These values are provided to the `Put` procedures, and returned by the `Get` procedures, in a parameter `Item`. Several overloaded procedures of these names exist, for different types of `Item`. These `Get` procedures analyze the input sequences of characters based on lexical elements (see Clause 5) and return the corresponding values; the `Put` procedures output the given values as appropriate lexical elements. Procedures `Get` and `Put` are also available that input and output individual characters treated as character values rather than as lexical elements. Related to character input are procedures to look ahead at the next character without reading it, and to read a character “immediately” without waiting for an end-of-line to signal availability.

In addition to the procedures `Get` and `Put` for numeric and enumeration types of `Item` that operate on text files, analogous procedures are provided that read from and write to a parameter of type `String`. These procedures perform the same analysis and composition of character sequences as their counterparts which have a file parameter.

For all `Get` and `Put` procedures that operate on text files, and for many other subprograms, there are forms with and without a file parameter. Each such `Get` procedure operates on an input file, and each such `Put` procedure operates on an output file. If no file is specified, a default input file or a default output file is used.

At the beginning of program execution the default input and output files are the so-called standard input file and standard output file. These files are open, have respectively the current modes `In_File` and `Out_File`, and are associated with two implementation-defined external files. Procedures are provided to change the current default input file and the current default output file.

At the beginning of program execution a default file for program-dependent error-related text output is the so-called standard error file. This file is open, has the current mode `Out_File`, and is associated with an implementation-defined external file. A procedure is provided to change the current default error file.

From a logical point of view, a text file is a sequence of pages, a page is a sequence of lines, and a line is a sequence of characters; the end of a line is marked by a *line terminator*; the end of a page is marked by the combination of a line terminator immediately followed by a *page terminator*; and the end of a file is marked by the combination of a line terminator immediately followed by a page terminator and then a *file terminator*. Terminators are generated during output; either by calls of procedures provided expressly for that purpose; or implicitly as part of other operations, for example, when a bounded line length, a bounded page length, or both, have been specified for a file.

The actual nature of terminators is not defined by the language and hence depends on the implementation. Although terminators are recognized or generated by certain of the procedures that follow, they are not necessarily implemented as characters or as sequences of characters. Whether they are characters (and if so which ones) in any particular implementation is not of concern to a user who neither explicitly outputs nor explicitly inputs control characters. The effect of input (Get) or output (Put) of control characters (other than horizontal tabulation) is not specified by the language.

The characters of a line are numbered, starting from one; the number of a character is called its *column number*. For a line terminator, a column number is also defined: it is one more than the number of characters in the line. The lines of a page, and the pages of a file, are similarly numbered. The current column number is the column number of the next character or line terminator to be transferred. The current line number is the number of the current line. The current page number is the number of the current page. These numbers are values of the subtype *Positive\_Count* of the type *Count* (by convention, the value zero of the type *Count* is used to indicate special conditions).

```
type Count is range 0 .. implementation-defined;
subtype Positive_Count is Count range 1 .. Count'Last;
```

For an output file or an append file, a *maximum line length* can be specified and a *maximum page length* can be specified. If a value to be output cannot fit on the current line, for a specified maximum line length, then a new line is automatically started before the value is output; if, further, this new line cannot fit on the current page, for a specified maximum page length, then a new page is automatically started before the value is output. Functions are provided to determine the maximum line length and the maximum page length. When a file is opened with mode *Out\_File* or *Append\_File*, both values are zero: by convention, this means that the line lengths and page lengths are unbounded. (Consequently, output consists of a single line if the subprograms for explicit control of line and page structure are not used.) The constant *Unbounded* is provided for this purpose.

### A.10.1 The Package Text\_IO

*Static Semantics*

The library package *Text\_IO* has the following declaration:

```
with Ada.IO_Exceptions;
package Ada.Text_IO
  with Global => in out synchronized is
    type File_Type is limited private;
    type File_Mode is (In_File, Out_File, Append_File);
    type Count is range 0 .. implementation-defined;
    subtype Positive_Count is Count range 1 .. Count'Last;
    Unbounded : constant Count := 0; -- line and page length

    subtype Field      is Integer range 0 .. implementation-defined;
    subtype Number_Base is Integer range 2 .. 16;
    type Type_Set is (Lower_Case, Upper_Case);

    -- File Management

    procedure Create (File : in out File_Type;
                     Mode : in File_Mode := Out_File;
                     Name : in String    := "";
                     Form : in String    := "");

    procedure Open  (File : in out File_Type;
                     Mode : in File_Mode;
                     Name : in String;
                     Form : in String := "");

    procedure Close (File : in out File_Type);
    procedure Delete (File : in out File_Type);
    procedure Reset  (File : in out File_Type; Mode : in File_Mode);
    procedure Reset  (File : in out File_Type);
```

```

function Mode      (File : in File_Type) return File_Mode;
function Name      (File : in File_Type) return String;
function Form      (File : in File_Type) return String;

function Is_Open (File : in File_Type) return Boolean;

-- Control of default input and output files

procedure Set_Input (File : in File_Type);
procedure Set_Output (File : in File_Type);
procedure Set_Error (File : in File_Type);

function Standard_Input return File_Type;
function Standard_Output return File_Type;
function Standard_Error return File_Type;

function Current_Input return File_Type;
function Current_Output return File_Type;
function Current_Error return File_Type;

type File_Access is access constant File_Type;

function Standard_Input return File_Access;
function Standard_Output return File_Access;
function Standard_Error return File_Access;

function Current_Input return File_Access;
function Current_Output return File_Access;
function Current_Error return File_Access;

--Buffer control
procedure Flush (File : in File_Type)
  with Global => overriding in out File;
procedure Flush
  with Global => in out all;

-- Specification of line and page lengths

procedure Set_Line_Length (File : in File_Type; To : in Count)
  with Global => overriding in out File;
procedure Set_Line_Length (To : in Count)
  with Global => in out all;

procedure Set_Page_Length (File : in File_Type; To : in Count)
  with Global => overriding in out File;
procedure Set_Page_Length (To : in Count)
  with Global => in out all;

function Line_Length (File : in File_Type) return Count;
function Line_Length return Count
  with Global => in all;

function Page_Length (File : in File_Type) return Count;
function Page_Length return Count
  with Global => in all;

-- Column, Line, and Page Control

procedure New_Line (File : in File_Type;
  Spacing : in Positive_Count := 1)
  with Global => overriding in out File;
procedure New_Line (Spacing : in Positive_Count := 1)
  with Global => in out all;

procedure Skip_Line (File : in File_Type;
  Spacing : in Positive_Count := 1)
  with Global => overriding in out File;
procedure Skip_Line (Spacing : in Positive_Count := 1)
  with Global => in out all;

function End_Of_Line (File : in File_Type) return Boolean;
function End_Of_Line return Boolean;

```

```

procedure New_Page (File : in File_Type)
  with Global => overriding in out File;
procedure New_Page
  with Global => in out all;
procedure Skip_Page (File : in File_Type)
  with Global => overriding in out File;
procedure Skip_Page
  with Global => in out all;

function End_Of_Page(File : in File_Type) return Boolean;
function End_Of_Page return Boolean
  with Global => in all;

function End_Of_File(File : in File_Type) return Boolean;
function End_Of_File return Boolean
  with Global => in all;

procedure Set_Col (File : in File_Type; To : in Positive_Count)
  with Global => overriding in out File;
procedure Set_Col (To : in Positive_Count)
  with Global => in out all;

procedure Set_Line(File : in File_Type; To : in Positive_Count)
  with Global => overriding in out File;
procedure Set_Line(To : in Positive_Count)
  with Global => in out all;

function Col (File : in File_Type) return Positive_Count;
function Col return Positive_Count
  with Global => in all;

function Line(File : in File_Type) return Positive_Count;
function Line return Positive_Count
  with Global => in all;

function Page(File : in File_Type) return Positive_Count;
function Page return Positive_Count
  with Global => in all;

-- Character Input-Output
procedure Get(File : in File_Type; Item : out Character)
  with Global => overriding in out File;
procedure Get(Item : out Character)
  with Global => in out all;

procedure Put(File : in File_Type; Item : in Character)
  with Global => overriding in out File;
procedure Put(Item : in Character)
  with Global => in out all;

procedure Look_Ahead (File : in File_Type;
  Item : out Character;
  End_Of_Line : out Boolean)
  with Global => overriding in out File;
procedure Look_Ahead (Item : out Character;
  End_Of_Line : out Boolean)
  with Global => in out all;

procedure Get_Immediate (File : in File_Type;
  Item : out Character)
  with Global => overriding in out File;
procedure Get_Immediate (Item : out Character)
  with Global => in out all;

```



```

procedure Get_Immediate(File      : in File_Type;
                        Item      : out Character;
                        Available  : out Boolean)
    with Global => overriding in out File;
procedure Get_Immediate(Item      : out Character;
                        Available  : out Boolean)
    with Global => in out all;

-- String Input-Output

procedure Get(File : in File_Type; Item : out String)
    with Global => overriding in out File;
procedure Get(Item : out String)
    with Global => in out all;

procedure Put(File : in File_Type; Item : in String)
    with Global => overriding in out File;
procedure Put(Item : in String)
    with Global => in out all;

procedure Get_Line(File : in File_Type;
                  Item : out String;
                  Last : out Natural)
    with Global => overriding in out File;
procedure Get_Line(Item : out String; Last : out Natural)
    with Global => in out all;

function Get_Line(File : in File_Type) return String
    with Global => overriding in out File;
function Get_Line return String
    with Global => in out all;

procedure Put_Line(File : in File_Type; Item : in String)
    with Global => overriding in out File;
procedure Put_Line(Item : in String)
    with Global => in out all;

-- Generic packages for Input-Output of Integer Types

generic
    type Num is range <>;
package Integer_IO is
    Default_Width : Field := Num'Width;
    Default_Base  : Number_Base := 10;

    procedure Get(File : in File_Type;
                Item : out Num;
                Width : in Field := 0)
        with Global => overriding in out File;
    procedure Get(Item : out Num;
                Width : in Field := 0)
        with Global => in out all;

```

```

procedure Put(File : in File_Type;
              Item : in Num;
              Width : in Field := Default_Width;
              Base : in Number_Base := Default_Base)
    with Global => overriding in out File;
procedure Put(Item : in Num;
              Width : in Field := Default_Width;
              Base : in Number_Base := Default_Base)
    with Global => in out all;
procedure Get(From : in String;
              Item : out Num;
              Last : out Positive)
    with Nonblocking;
procedure Put(To : out String;
              Item : in Num;
              Base : in Number_Base := Default_Base)
    with Nonblocking;
end Integer_IO;

generic
  type Num is mod <>;
package Modular_IO is
  Default_Width : Field := Num'Width;
  Default_Base : Number_Base := 10;

  procedure Get(File : in File_Type;
                Item : out Num;
                Width : in Field := 0)
    with Global => overriding in out File;
  procedure Get(Item : out Num;
                Width : in Field := 0)
    with Global => in out all;

  procedure Put(File : in File_Type;
                Item : in Num;
                Width : in Field := Default_Width;
                Base : in Number_Base := Default_Base)
    with Global => overriding in out File;
  procedure Put(Item : in Num;
                Width : in Field := Default_Width;
                Base : in Number_Base := Default_Base)
    with Global => in out all;
  procedure Get(From : in String;
                Item : out Num;
                Last : out Positive)
    with Nonblocking;
  procedure Put(To : out String;
                Item : in Num;
                Base : in Number_Base := Default_Base)
    with Nonblocking;
end Modular_IO;

-- Generic packages for Input-Output of Real Types
generic
  type Num is digits <>;
package Float_IO is
  Default_Fore : Field := 2;
  Default_Aft : Field := Num'Digits-1;
  Default_Exp : Field := 3;

```

```

procedure Get(File : in File_Type;
               Item : out Num;
               Width : in Field := 0)
    with Global => overriding in out File;
procedure Get(Item : out Num;
               Width : in Field := 0)
    with Global => in out all;

procedure Put(File : in File_Type;
               Item : in Num;
               Fore : in Field := Default_Fore;
               Aft  : in Field := Default_Aft;
               Exp  : in Field := Default_Exp)
    with Global => overriding in out File;
procedure Put(Item : in Num;
               Fore : in Field := Default_Fore;
               Aft  : in Field := Default_Aft;
               Exp  : in Field := Default_Exp)
    with Global => in out all;

procedure Get(From : in String;
               Item : out Num;
               Last : out Positive)
    with Nonblocking;
procedure Put(To   : out String;
               Item : in Num;
               Aft  : in Field := Default_Aft;
               Exp  : in Field := Default_Exp)
    with Nonblocking;
end Float_IO;

generic
    type Num is delta <>;
package Fixed_IO is

    Default_Fore : Field := Num'Fore;
    Default_Aft  : Field := Num'Aft;
    Default_Exp  : Field := 0;

    procedure Get(File : in File_Type;
                  Item : out Num;
                  Width : in Field := 0)
        with Global => overriding in out File;
    procedure Get(Item : out Num;
                  Width : in Field := 0)
        with Global => in out all;

    procedure Put(File : in File_Type;
                  Item : in Num;
                  Fore : in Field := Default_Fore;
                  Aft  : in Field := Default_Aft;
                  Exp  : in Field := Default_Exp)
        with Global => overriding in out File;
    procedure Put(Item : in Num;
                  Fore : in Field := Default_Fore;
                  Aft  : in Field := Default_Aft;
                  Exp  : in Field := Default_Exp)
        with Global => in out all;

```

```

procedure Get(From : in String;
               Item : out Num;
               Last : out Positive)
    with Nonblocking;
procedure Put(To : out String;
               Item : in Num;
               Aft : in Field := Default_Aft;
               Exp : in Field := Default_Exp)
    with Nonblocking;
end Fixed_IO;

generic
    type Num is delta <> digits <>;
package Decimal_IO is

    Default_Fore : Field := Num'Fore;
    Default_Aft  : Field := Num'Aft;
    Default_Exp  : Field := 0;

    procedure Get(File : in File_Type;
                  Item : out Num;
                  Width : in Field := 0)
        with Global => overriding in out File;
    procedure Get(Item : out Num;
                  Width : in Field := 0)
        with Global => in out all;

    procedure Put(File : in File_Type;
                  Item : in Num;
                  Fore : in Field := Default_Fore;
                  Aft  : in Field := Default_Aft;
                  Exp  : in Field := Default_Exp)
        with Global => overriding in out File;
    procedure Put(Item : in Num;
                  Fore : in Field := Default_Fore;
                  Aft  : in Field := Default_Aft;
                  Exp  : in Field := Default_Exp)
        with Global => in out all;

    procedure Get(From : in String;
                  Item : out Num;
                  Last : out Positive)
        with Nonblocking;
    procedure Put(To : out String;
                  Item : in Num;
                  Aft : in Field := Default_Aft;
                  Exp : in Field := Default_Exp)
        with Nonblocking;
end Decimal_IO;

```

— Generic package for Input-Output of Enumeration Types

```

generic
    type Enum is (<>);
package Enumeration_IO is

    Default_Width : Field := 0;
    Default_Setting : Type_Set := Upper_Case;

    procedure Get(File : in File_Type;
                  Item : out Enum)
        with Global => overriding in out File;
    procedure Get(Item : out Enum)
        with Global => in out all;

```

```

procedure Put (File : in File_Type;
               Item : in Enum;
               Width : in Field := Default_Width;
               Set : in Type_Set := Default_Setting)
    with Global => overriding in out File;
procedure Put (Item : in Enum;
               Width : in Field := Default_Width;
               Set : in Type_Set := Default_Setting)
    with Global => in out all;

procedure Get (From : in String;
               Item : out Enum;
               Last : out Positive)
    with Nonblocking;
procedure Put (To : out String;
               Item : in Enum;
               Set : in Type_Set := Default_Setting)
    with Nonblocking;
end Enumeration_IO;

-- Exceptions
Status_Error : exception renames IO_Exceptions.Status_Error;
Mode_Error : exception renames IO_Exceptions.Mode_Error;
Name_Error : exception renames IO_Exceptions.Name_Error;
Use_Error : exception renames IO_Exceptions.Use_Error;
Device_Error : exception renames IO_Exceptions.Device_Error;
End_Error : exception renames IO_Exceptions.End_Error;
Data_Error : exception renames IO_Exceptions.Data_Error;
Layout_Error : exception renames IO_Exceptions.Layout_Error;

package Wide_File_Names is
    -- File management

    procedure Create (File : in out File_Type;
                    Mode : in File_Mode := Out_File;
                    Name : in Wide_String := "";
                    Form : in Wide_String := "");

    procedure Open (File : in out File_Type;
                  Mode : in File_Mode;
                  Name : in Wide_String;
                  Form : in Wide_String := "");

    function Name (File : in File_Type) return Wide_String;
    function Form (File : in File_Type) return Wide_String;
end Wide_File_Names;

package Wide_Wide_File_Names is
    -- File management

    procedure Create (File : in out File_Type;
                    Mode : in File_Mode := Out_File;
                    Name : in Wide_Wide_String := "";
                    Form : in Wide_Wide_String := "");

    procedure Open (File : in out File_Type;
                  Mode : in File_Mode;
                  Name : in Wide_Wide_String;
                  Form : in Wide_Wide_String := "");

    function Name (File : in File_Type) return Wide_Wide_String;
    function Form (File : in File_Type) return Wide_Wide_String;
end Wide_Wide_File_Names;

```

```

private
  ... -- not specified by the language
end Ada.Text_IO;

```

The type `File_Type` needs finalization (see 10.6).

## A.10.2 Text File Management

### Static Semantics

The only allowed file modes for text files are the modes `In_File`, `Out_File`, and `Append_File`. The subprograms given in A.8.2 for the control of external files, and the function `End_Of_File` given in A.8.3 for sequential input-output, are also available for text files. There is also a version of `End_Of_File` that refers to the current default input file. For text files, the procedures have the following additional effects:

- For the procedures `Create` and `Open`: After a file with mode `Out_File` or `Append_File` is opened, the page length and line length are unbounded (both have the conventional value zero). After a file (of any mode) is opened, the current column, current line, and current page numbers are set to one. If the mode is `Append_File`, it is implementation defined whether a page terminator will separate preexisting text in the file from the new text to be written.
- For the procedure `Close`: If the file has the current mode `Out_File` or `Append_File`, has the effect of calling `New_Page`, unless the current page is already terminated; then outputs a file terminator.
- For the procedure `Reset`: If the file has the current mode `Out_File` or `Append_File`, has the effect of calling `New_Page`, unless the current page is already terminated; then outputs a file terminator. The current column, line, and page numbers are set to one, and the line and page lengths to Unbounded. If the new mode is `Append_File`, it is implementation defined whether a page terminator will separate preexisting text in the file from the new text to be written.

The exception `Mode_Error` is propagated by the procedure `Reset` upon an attempt to change the mode of a file that is the current default input file, the current default output file, or the current default error file.

NOTE An implementation can define the `Form` parameter of `Create` and `Open` to control effects including the following:

- the interpretation of line and column numbers for an interactive file, and
- the interpretation of text formats in a file created by a foreign program.

## A.10.3 Default Input, Output, and Error Files

### Static Semantics

The following subprograms provide for the control of the particular default files that are used when a file parameter is omitted from a `Get`, `Put`, or other operation of text input-output described below, or when application-dependent error-related text is to be output.

```

procedure Set_Input (File : in File_Type);

```

Operates on a file of mode `In_File`. Sets the current default input file to `File`.

The exception `Status_Error` is propagated if the given file is not open. The exception `Mode_Error` is propagated if the mode of the given file is not `In_File`.

```

procedure Set_Output (File : in File_Type);
procedure Set_Error (File : in File_Type);

```

Each operates on a file of mode `Out_File` or `Append_File`. `Set_Output` sets the current default output file to `File`. `Set_Error` sets the current default error file to `File`. The exception `Status_Error` is propagated if the given file is not open. The exception `Mode_Error` is propagated if the mode of the given file is not `Out_File` or `Append_File`.



```
function Standard_Input return File_Type;
function Standard_Input return File_Access;
```

Returns the standard input file (see A.10), or an access value designating the standard input file, respectively.

```
function Standard_Output return File_Type;
function Standard_Output return File_Access;
```

Returns the standard output file (see A.10) or an access value designating the standard output file, respectively.

```
function Standard_Error return File_Type;
function Standard_Error return File_Access;
```

Returns the standard error file (see A.10), or an access value designating the standard error file, respectively.

The Form strings implicitly associated with the opening of Standard\_Input, Standard\_Output, and Standard\_Error at the start of program execution are implementation defined.

```
function Current_Input return File_Type;
function Current_Input return File_Access;
```

Returns the current default input file, or an access value designating the current default input file, respectively.

```
function Current_Output return File_Type;
function Current_Output return File_Access;
```

Returns the current default output file, or an access value designating the current default output file, respectively.

```
function Current_Error return File_Type;
function Current_Error return File_Access;
```

Returns the current default error file, or an access value designating the current default error file, respectively.

```
procedure Flush (File : in File_Type);
procedure Flush;
```

The effect of Flush is the same as the corresponding subprogram in Sequential\_IO (see A.8.2). If File is not explicitly specified, Current\_Output is used.

#### *Erroneous Execution*

The execution of a program is erroneous if it invokes an operation on a current default input, default output, or default error file, and if the corresponding file object is closed or no longer exists.

NOTE 1 The standard input, standard output, and standard error files cannot be opened, closed, reset, or deleted, because the parameter File of the corresponding procedures has the mode **in out**.

NOTE 2 The standard input, standard output, and standard error files are different file objects, but not necessarily different external files.

### A.10.4 Specification of Line and Page Lengths

#### *Static Semantics*

The subprograms described in this subclause are concerned with the line and page structure of a file of mode Out\_File or Append\_File. They operate either on the file given as the first parameter, or, in the absence of such a file parameter, on the current default output file. They provide for output of text with a specified maximum line length or page length. In these cases, line and page terminators are output implicitly and automatically when necessary. When line and page lengths are unbounded (that is, when

they have the conventional value zero), as in the case of a newly opened file, new lines and new pages are only started when explicitly called for.

In all cases, the exception `Status_Error` is propagated if the file to be used is not open; the exception `Mode_Error` is propagated if the mode of the file is not `Out_File` or `Append_File`.

```
procedure Set_Line_Length(File : in File_Type; To : in Count);
procedure Set_Line_Length(To : in Count);
```

Sets the maximum line length of the specified output or append file to the number of characters specified by `To`. The value zero for `To` specifies an unbounded line length.

The exception `Use_Error` is propagated if the specified line length is inappropriate for the associated external file.

```
procedure Set_Page_Length(File : in File_Type; To : in Count);
procedure Set_Page_Length(To : in Count);
```

Sets the maximum page length of the specified output or append file to the number of lines specified by `To`. The value zero for `To` specifies an unbounded page length.

The exception `Use_Error` is propagated if the specified page length is inappropriate for the associated external file.

```
function Line_Length(File : in File_Type) return Count;
function Line_Length return Count;
```

Returns the maximum line length currently set for the specified output or append file, or zero if the line length is unbounded.

```
function Page_Length(File : in File_Type) return Count;
function Page_Length return Count;
```

Returns the maximum page length currently set for the specified output or append file, or zero if the page length is unbounded.

### A.10.5 Operations on Columns, Lines, and Pages

#### *Static Semantics*

The subprograms described in this subclause provide for explicit control of line and page structure; they operate either on the file given as the first parameter, or, in the absence of such a file parameter, on the appropriate (input or output) current default file. The exception `Status_Error` is propagated by any of these subprograms if the file to be used is not open.

```
procedure New_Line(File : in File_Type; Spacing : in Positive_Count := 1);
procedure New_Line(Spacing : in Positive_Count := 1);
```

Operates on a file of mode `Out_File` or `Append_File`.

For a Spacing of one: Outputs a line terminator and sets the current column number to one. Then increments the current line number by one, except in the case that the current line number is already greater than or equal to the maximum page length, for a bounded page length; in that case a page terminator is output, the current page number is incremented by one, and the current line number is set to one.

For a Spacing greater than one, the above actions are performed `Spacing` times.

The exception `Mode_Error` is propagated if the mode is not `Out_File` or `Append_File`.

```
procedure Skip_Line(File : in File_Type; Spacing : in Positive_Count := 1);
procedure Skip_Line(Spacing : in Positive_Count := 1);
```

Operates on a file of mode `In_File`.

For a Spacing of one: Reads and discards all characters until a line terminator has been read, and then sets the current column number to one. If the line terminator is not immediately

followed by a page terminator, the current line number is incremented by one. Otherwise, if the line terminator is immediately followed by a page terminator, then the page terminator is skipped, the current page number is incremented by one, and the current line number is set to one.

For a Spacing greater than one, the above actions are performed Spacing times.

The exception Mode\_Error is propagated if the mode is not In\_File. The exception End\_Error is propagated if an attempt is made to read a file terminator.

```
function End_Of_Line(File : in File_Type) return Boolean;
function End_Of_Line return Boolean;
```

Operates on a file of mode In\_File. Returns True if a line terminator or a file terminator is next; otherwise, returns False.

The exception Mode\_Error is propagated if the mode is not In\_File.

```
procedure New_Page(File : in File_Type);
procedure New_Page;
```

Operates on a file of mode Out\_File or Append\_File. Outputs a line terminator if the current line is not terminated, or if the current page is empty (that is, if the current column and line numbers are both equal to one). Then outputs a page terminator, which terminates the current page. Adds one to the current page number and sets the current column and line numbers to one.

The exception Mode\_Error is propagated if the mode is not Out\_File or Append\_File.

```
procedure Skip_Page(File : in File_Type);
procedure Skip_Page;
```

Operates on a file of mode In\_File. Reads and discards all characters and line terminators until a page terminator has been read. Then adds one to the current page number, and sets the current column and line numbers to one.

The exception Mode\_Error is propagated if the mode is not In\_File. The exception End\_Error is propagated if an attempt is made to read a file terminator.

```
function End_Of_Page(File : in File_Type) return Boolean;
function End_Of_Page return Boolean;
```

Operates on a file of mode In\_File. Returns True if the combination of a line terminator and a page terminator is next, or if a file terminator is next; otherwise, returns False.

The exception Mode\_Error is propagated if the mode is not In\_File.

```
function End_Of_File(File : in File_Type) return Boolean;
function End_Of_File return Boolean;
```

Operates on a file of mode In\_File. Returns True if a file terminator is next, or if the combination of a line, a page, and a file terminator is next; otherwise, returns False.

The exception Mode\_Error is propagated if the mode is not In\_File.

The following subprograms provide for the control of the current position of reading or writing in a file. In all cases, the default file is the current output file.

```
procedure Set_Col(File : in File_Type; To : in Positive_Count);
procedure Set_Col(To : in Positive_Count);
```

If the file mode is Out\_File or Append\_File:

- If the value specified by To is greater than the current column number, outputs spaces, adding one to the current column number after each space, until the current column number equals the specified value. If the value specified by To is equal to the current column number, there is no effect. If the value specified by To is less than the current

column number, has the effect of calling `New_Line` (with a spacing of one), then outputs  $(To - 1)$  spaces, and sets the current column number to the specified value.

- The exception `Layout_Error` is propagated if the value specified by `To` exceeds `Line_Length` when the line length is bounded (that is, when it does not have the conventional value zero).

If the file mode is `In_File`:

- Reads (and discards) individual characters, line terminators, and page terminators, until the next character to be read has a column number that equals the value specified by `To`; there is no effect if the current column number already equals this value. Each transfer of a character or terminator maintains the current column, line, and page numbers in the same way as a `Get` procedure (see A.10.6). (Short lines will be skipped until a line is reached that has a character at the specified column position.)
- The exception `End_Error` is propagated if an attempt is made to read a file terminator.

```
procedure Set_Line(File : in File_Type; To : in Positive_Count);
procedure Set_Line(To   : in Positive_Count);
```

If the file mode is `Out_File` or `Append_File`:

- If the value specified by `To` is greater than the current line number, has the effect of repeatedly calling `New_Line` (with a spacing of one), until the current line number equals the specified value. If the value specified by `To` is equal to the current line number, there is no effect. If the value specified by `To` is less than the current line number, has the effect of calling `New_Page` followed, if `To` is greater than 1, by a call of `New_Line` with a spacing equal to  $(To - 1)$ .
- The exception `Layout_Error` is propagated if the value specified by `To` exceeds `Page_Length` when the page length is bounded (that is, when it does not have the conventional value zero).

If the mode is `In_File`:

- Has the effect of repeatedly calling `Skip_Line` (with a spacing of one), until the current line number equals the value specified by `To`; there is no effect if the current line number already equals this value. (Short pages will be skipped until a page is reached that has a line at the specified line position.)
- The exception `End_Error` is propagated if an attempt is made to read a file terminator.

```
function Col(File : in File_Type) return Positive_Count;
function Col return Positive_Count;
```

Returns the current column number.

The exception `Layout_Error` is propagated if this number exceeds `Count'Last`.

```
function Line(File : in File_Type) return Positive_Count;
function Line return Positive_Count;
```

Returns the current line number.

The exception `Layout_Error` is propagated if this number exceeds `Count'Last`.

```
function Page(File : in File_Type) return Positive_Count;
function Page return Positive_Count;
```

Returns the current page number.

The exception `Layout_Error` is propagated if this number exceeds `Count'Last`.

The column number, line number, or page number are allowed to exceed `Count'Last` (as a consequence of the input or output of sufficiently many characters, lines, or pages). These events do not cause any exception to be propagated. However, a call of `Col`, `Line`, or `Page` propagates the exception `Layout_Error` if the corresponding number exceeds `Count'Last`.

NOTE A page terminator is always skipped whenever the preceding line terminator is skipped. An implementation can represent the combination of these terminators by a single character, provided that it is properly recognized on input.

## A.10.6 Get and Put Procedures

### Static Semantics

The procedures Get and Put for items of the type Character, String, numeric types, and enumeration types are described in subsequent subclauses. Features of these procedures that are common to most of these types are described in this subclause. The Get and Put procedures for items of type Character and String deal with individual character values; the Get and Put procedures for numeric and enumeration types treat the items as lexical elements.

All procedures Get and Put have forms with a file parameter, written first. Where this parameter is omitted, the appropriate (input or output) current default file is understood to be specified. Each procedure Get operates on a file of mode In\_File. Each procedure Put operates on a file of mode Out\_File or Append\_File.

All procedures Get and Put maintain the current column, line, and page numbers of the specified file: the effect of each of these procedures upon these numbers is the result of the effects of individual transfers of characters and of individual output or skipping of terminators. Each transfer of a character adds one to the current column number. Each output of a line terminator sets the current column number to one and adds one to the current line number. Each output of a page terminator sets the current column and line numbers to one and adds one to the current page number. For input, each skipping of a line terminator sets the current column number to one and adds one to the current line number; each skipping of a page terminator sets the current column and line numbers to one and adds one to the current page number. Similar considerations apply to the procedures Get\_Line, Put\_Line, and Set\_Col.

Several Get and Put procedures, for numeric and enumeration types, have *format* parameters which specify field lengths; these parameters are of the nonnegative subtype Field of the type Integer.

Input-output of enumeration values uses the syntax of the corresponding lexical elements. Any Get procedure for an enumeration type begins by skipping any leading blanks, or line or page terminators. A *blank* is defined as a space or a horizontal tabulation character. Next, characters are input only so long as the sequence input is an initial sequence of an identifier or of a character literal (in particular, input ceases when a line terminator is encountered). The character or line terminator that causes input to cease remains available for subsequent input.

For a numeric type, the Get procedures have a format parameter called Width. If the value given for this parameter is zero, the Get procedure proceeds in the same manner as for enumeration types, but using the syntax of numeric literals instead of that of enumeration literals. If a nonzero value is given, then exactly Width characters are input, or the characters up to a line terminator, whichever comes first; any skipped leading blanks are included in the count. The syntax used for numeric literals is an extended syntax that allows a leading sign (but no intervening blanks, or line or page terminators) and that also allows (for real types) an integer literal as well as forms that have digits only before the point or only after the point.

Any Put procedure, for an item of a numeric or an enumeration type, outputs the value of the item as a numeric literal, identifier, or character literal, as appropriate. This is preceded by leading spaces if required by the format parameters Width or Fore (as described in later subclauses), and then a minus sign for a negative value; for an enumeration type, the spaces follow instead of leading. The format given for a Put procedure is overridden if it is insufficiently wide, by using the minimum necessary width.

Two further cases arise for Put procedures for numeric and enumeration types, if the line length of the specified output file is bounded (that is, if it does not have the conventional value zero). If the number of characters to be output does not exceed the maximum line length, but is such that they cannot fit on the current line, starting from the current column, then (in effect) New\_Line is called (with a spacing of

one) before output of the item. Otherwise, if the number of characters exceeds the maximum line length, then the exception `Layout_Error` is propagated and nothing is output.

The exception `Status_Error` is propagated by any of the procedures `Get`, `Get_Line`, `Put`, and `Put_Line` if the file to be used is not open. The exception `Mode_Error` is propagated by the procedures `Get` and `Get_Line` if the mode of the file to be used is not `In_File`; and by the procedures `Put` and `Put_Line`, if the mode is not `Out_File` or `Append_File`.

The exception `End_Error` is propagated by a `Get` procedure if an attempt is made to skip a file terminator. The exception `Data_Error` is propagated by a `Get` procedure if the sequence finally input is not a lexical element corresponding to the type, in particular if no characters were input; for this test, leading blanks are ignored; for an item of a numeric type, when a sign is input, this rule applies to the succeeding numeric literal. The exception `Layout_Error` is propagated by a `Put` procedure that outputs to a parameter of type `String`, if the length of the actual string is insufficient for the output of the item.

EXAMPLE 1 In the examples, here and in A.10.8 and A.10.9, the string quotes and the lower case letter b are not transferred: they are shown only to reveal the layout and spaces.

```
N : Integer;
...
Get(N);

--          Characters at input  Sequence input      Value of N

--          bb-12535b          -12535 -12535
--          bb12_535e1b        12_535e1 125350
--          bb12_535e;         12_535e  (none) Data_Error raised
```

EXAMPLE 2 Example of overridden width parameter:

```
Put(Item => -23, Width => 2); -- "-23"
```

## A.10.7 Input-Output of Characters and Strings

*Static Semantics*

For an item of type `Character` the following procedures are provided:

```
procedure Get(File : in File_Type; Item : out Character);
procedure Get(Item : out Character);
```

After skipping any line terminators and any page terminators, reads the next character from the specified input file and returns the value of this character in the out parameter `Item`.

The exception `End_Error` is propagated if an attempt is made to skip a file terminator.

```
procedure Put(File : in File_Type; Item : in Character);
procedure Put(Item : in Character);
```

If the line length of the specified output file is bounded (that is, does not have the conventional value zero), and the current column number exceeds it, has the effect of calling `New_Line` with a spacing of one. Then, or otherwise, outputs the given character to the file.

```
procedure Look_Ahead (File      : in File_Type;
                      Item      : out Character;
                      End_Of_Line : out Boolean);
procedure Look_Ahead (Item      : out Character;
                      End_Of_Line : out Boolean);
```

`Status_Error` is propagated if the file is not open. `Mode_Error` is propagated if the mode of the file is not `In_File`. Sets `End_Of_Line` to `True` if at end of line, including if at end of page or at end of file; in each of these cases the value of `Item` is not specified. Otherwise, `End_Of_Line` is set to `False` and `Item` is set to the next character (without consuming it) from the file.



```

procedure Get_Immediate(File : in File_Type;
                        Item : out Character);
procedure Get_Immediate(Item : out Character);

```

Reads the next character, either control or graphic, from the specified File or the default input file. Status\_Error is propagated if the file is not open. Mode\_Error is propagated if the mode of the file is not In\_File. End\_Error is propagated if at the end of the file. The current column, line and page numbers for the file are not affected.

```

procedure Get_Immediate(File      : in File_Type;
                        Item       : out Character;
                        Available  : out Boolean);
procedure Get_Immediate(Item     : out Character;
                        Available  : out Boolean);

```

If a character, either control or graphic, is available from the specified File or the default input file, then the character is read; Available is True and Item contains the value of this character. If a character is not available, then Available is False and the value of Item is not specified. Status\_Error is propagated if the file is not open. Mode\_Error is propagated if the mode of the file is not In\_File. End\_Error is propagated if at the end of the file. The current column, line and page numbers for the file are not affected.

For an item of type String the following subprograms are provided:

```

procedure Get(File : in File_Type; Item : out String);
procedure Get(Item : out String);

```

Determines the length of the given string and attempts that number of Get operations for successive characters of the string (in particular, no operation is performed if the string is null).

```

procedure Put(File : in File_Type; Item : in String);
procedure Put(Item : in String);

```

Determines the length of the given string and attempts that number of Put operations for successive characters of the string (in particular, no operation is performed if the string is null).

```

function Get_Line(File : in File_Type) return String;
function Get_Line return String;

```

Returns a result string constructed by reading successive characters from the specified input file, and assigning them to successive characters of the result string. The result string has a lower bound of 1 and an upper bound of the number of characters read. Reading stops when the end of the line is met; Skip\_Line is then (in effect) called with a spacing of 1.

Constraint\_Error is raised if the length of the line exceeds Positive'Last; in this case, the line number and page number are unchanged, and the column number is unspecified but no less than it was before the call. The exception End\_Error is propagated if an attempt is made to skip a file terminator.

```

procedure Get_Line(File : in File_Type;
                  Item : out String;
                  Last : out Natural);
procedure Get_Line(Item : out String;
                  Last : out Natural);

```

Reads successive characters from the specified input file and assigns them to successive characters of the specified string. Reading stops if the end of the string is met. Reading also stops if the end of the line is met before meeting the end of the string; in this case Skip\_Line is (in effect) called with a spacing of 1. The values of characters not assigned are not specified.

If characters are read, returns in Last the index value such that Item(Last) is the last character assigned (the index of the first character assigned is Item'First). If no characters are read, returns in Last an index value that is one less than Item'First. The exception End\_Error is propagated if an attempt is made to skip a file terminator.

```
procedure Put_Line(File : in File_Type; Item : in String);
procedure Put_Line(Item : in String);
```

Calls the procedure Put for the given string, and then the procedure New\_Line with a spacing of one.

#### *Implementation Advice*

The Get\_Immediate procedures should be implemented with unbuffered input. For a device such as a keyboard, input should be “available” if a key has already been typed, whereas for a disk file, input should always be available except at end of file. For a file associated with a keyboard-like device, any line-editing features of the underlying operating system should be disabled during the execution of Get\_Immediate.

NOTE 1 Get\_Immediate can be used to read a single key from the keyboard “immediately”; that is, without waiting for an end of line. In a call of Get\_Immediate without the parameter Available, the caller will wait until a character is available.

NOTE 2 In a literal string parameter of Put, the enclosing string bracket characters are not output. Each doubled string bracket character in the enclosed string is output as a single string bracket character, as a consequence of the rule for string literals (see 5.6).

NOTE 3 A string read by Get or written by Put can extend over several lines. An implementation is allowed to assume that certain external files do not contain page terminators, in which case Get\_Line and Skip\_Line can return as soon as a line terminator is read.

### A.10.8 Input-Output for Integer Types

#### *Static Semantics*

The following procedures are defined in the generic packages Integer\_IO and Modular\_IO, which have to be instantiated for the appropriate signed integer or modular type respectively (indicated by Num in the specifications).

Values are output as decimal or based literals, without low line characters or exponent, and, for Integer\_IO, preceded by a minus sign if negative. The format (which includes any leading spaces and minus sign) can be specified by an optional field width parameter. Values of widths of fields in output formats are of the nonnegative integer subtype Field. Values of bases are of the integer subtype Number\_Base.

```
subtype Number_Base is Integer range 2 .. 16;
```

The default field width and base to be used by output procedures are defined by the following variables that are declared in the generic packages Integer\_IO and Modular\_IO:

```
Default_Width : Field := Num'Width;
Default_Base  : Number_Base := 10;
```

The following procedures are provided:

```
procedure Get(File : in File_Type; Item : out Num; Width : in Field := 0);
procedure Get(Item : out Num; Width : in Field := 0);
```

If the value of the parameter Width is zero, skips any leading blanks, line terminators, or page terminators, then reads a plus sign if present or (for a signed type only) a minus sign if present, then reads the longest possible sequence of characters matching the syntax of a numeric literal without a point. If a nonzero value of Width is supplied, then exactly Width characters are input, or the characters (possibly none) up to a line terminator, whichever comes first; any skipped leading blanks are included in the count.

Returns, in the parameter Item, the value of type Num that corresponds to the sequence input.

The exception Data\_Error is propagated if the sequence of characters read does not form a legal integer literal or if the value obtained is not of the subtype Num.

```

procedure Put (File   : in File_Type;
               Item   : in Num;
               Width  : in Field := Default_Width;
               Base   : in Number_Base := Default_Base);

procedure Put (Item   : in Num;
               Width  : in Field := Default_Width;
               Base   : in Number_Base := Default_Base);

```

Outputs the value of the parameter Item as an integer literal, with no low lines, no exponent, and no leading zeros (but a single zero for the value zero), and a preceding minus sign for a negative value.

If the resulting sequence of characters to be output has fewer than Width characters, then leading spaces are first output to make up the difference.

Uses the syntax for decimal literal if the parameter Base has the value ten (either explicitly or through Default\_Base); otherwise, uses the syntax for based literal, with any letters in upper case.

```

procedure Get (From : in String; Item : out Num; Last : out Positive);

```

Reads an integer value from the beginning of the given string, following the same rules as the Get procedure that reads an integer value from a file, but treating the end of the string as a file terminator. Returns, in the parameter Item, the value of type Num that corresponds to the sequence input. Returns in Last the index value such that From (Last) is the last character read.

The exception Data\_Error is propagated if the sequence input does not have the required syntax or if the value obtained is not of the subtype Num.

```

procedure Put (To      : out String;
               Item   : in Num;
               Base   : in Number_Base := Default_Base);

```

Outputs the value of the parameter Item to the given string, following the same rule as for output to a file, using the length of the given string as the value for Width.

Integer\_Text\_IO is a library package that is a nongeneric equivalent to Text\_IO.Integer\_IO for the predefined type Integer:

```

with Ada.Text_IO;
package Ada.Integer_Text_IO is new Ada.Text_IO.Integer_IO(Integer);

```

For each predefined signed integer type, a nongeneric equivalent to Text\_IO.Integer\_IO is provided, with names such as Ada.Long\_Integer\_Text\_IO.

#### Implementation Permissions

The nongeneric equivalent packages can be actual instantiations of the generic package for the appropriate predefined type, though that is not required.

EXAMPLE *Examples of use of an instantiation of Text\_IO.Integer\_IO:*

```

subtype Byte_Int is Integer range -127 .. 127;
package Int_IO is new Integer_IO(Byte_Int); use Int_IO;
-- default format used at instantiation,
-- Default_Width = 4, Default_Base = 10

Put(126);                -- "b126"
Put(-126, 7);            -- "bbb-126"
Put(126, Width => 13, Base => 2); -- "bbb2#1111110#"

```

### A.10.9 Input-Output for Real Types

#### Static Semantics

The following procedures are defined in the generic packages `Float_IO`, `Fixed_IO`, and `Decimal_IO`, which have to be instantiated for the appropriate floating point, ordinary fixed point, or decimal fixed point type respectively (indicated by `Num` in the specifications).

Values are output as decimal literals without low line characters. The format of each value output consists of a `Fore` field, a decimal point, an `Aft` field, and (if a nonzero `Exp` parameter is supplied) the letter `E` and an `Exp` field. The two possible formats thus correspond to:

`Fore . Aft`

and to:

`Fore . Aft E Exp`

without any spaces between these fields. The `Fore` field may include leading spaces, and a minus sign for negative values. The `Aft` field includes only decimal digits (possibly with trailing zeros). The `Exp` field includes the sign (plus or minus) and the exponent (possibly with leading zeros).

For floating point types, the default lengths of these fields are defined by the following variables that are declared in the generic package `Float_IO`:

```
Default_Fore : Field := 2;
Default_Aft  : Field := Num'Digits-1;
Default_Exp  : Field := 3;
```

For ordinary or decimal fixed point types, the default lengths of these fields are defined by the following variables that are declared in the generic packages `Fixed_IO` and `Decimal_IO`, respectively:

```
Default_Fore : Field := Num'Fore;
Default_Aft  : Field := Num'Aft;
Default_Exp  : Field := 0;
```

The following procedures are provided:

```
procedure Get(File : in File_Type; Item : out Num; Width : in Field := 0);
procedure Get(Item : out Num; Width : in Field := 0);
```

If the value of the parameter `Width` is zero, skips any leading blanks, line terminators, or page terminators, then reads the longest possible sequence of characters matching the syntax of any of the following (see 5.4):

- `[+|-]numeric_literal`
- `[+|-]numeral.[exponent]`
- `[+|-]numeral[exponent]`
- `[+|-]base#based_numeral.#[exponent]`
- `[+|-]base#.based_numeral#[exponent]`

If a nonzero value of `Width` is supplied, then exactly `Width` characters are input, or the characters (possibly none) up to a line terminator, whichever comes first; any skipped leading blanks are included in the count.

Returns in the parameter `Item` the value of type `Num` that corresponds to the sequence input, preserving the sign (positive if none has been specified) of a zero value if `Num` is a floating point type and `Num'Signed_Zeros` is `True`.

The exception `Data_Error` is propagated if the sequence input does not have the required syntax or if the value obtained is not of the subtype `Num`.

```

procedure Put (File : in File_Type;
               Item : in Num;
               Fore : in Field := Default_Fore;
               Aft  : in Field := Default_Aft;
               Exp  : in Field := Default_Exp);

```

```

procedure Put (Item : in Num;
               Fore : in Field := Default_Fore;
               Aft  : in Field := Default_Aft;
               Exp  : in Field := Default_Exp);

```

Outputs the value of the parameter *Item* as a decimal literal with the format defined by *Fore*, *Aft* and *Exp*. If the value is negative, or if *Num* is a floating point type where *Num*'s *Signed\_Zeros* is *True* and the value is a negatively signed zero, then a minus sign is included in the integer part. If *Exp* has the value zero, then the integer part to be output has as many digits as are needed to represent the integer part of the value of *Item*, overriding *Fore* if necessary, or consists of the digit zero if the value of *Item* has no integer part.

If *Exp* has a value greater than zero, then the integer part to be output has a single digit, which is nonzero except for the value 0.0 of *Item*.

In both cases, however, if the integer part to be output has fewer than *Fore* characters, including any minus sign, then leading spaces are first output to make up the difference. The number of digits of the fractional part is given by *Aft*, or is one if *Aft* equals zero. The value is rounded; a value of exactly one half in the last place is rounded away from zero.

If *Exp* has the value zero, there is no exponent part. If *Exp* has a value greater than zero, then the exponent part to be output has as many digits as are needed to represent the exponent part of the value of *Item* (for which a single digit integer part is used), and includes an initial sign (plus or minus). If the exponent part to be output has fewer than *Exp* characters, including the sign, then leading zeros precede the digits, to make up the difference. For the value 0.0 of *Item*, the exponent has the value zero.

```

procedure Get (From : in String; Item : out Num; Last : out Positive);

```

Reads a real value from the beginning of the given string, following the same rule as the *Get* procedure that reads a real value from a file, but treating the end of the string as a file terminator. Returns, in the parameter *Item*, the value of type *Num* that corresponds to the sequence input. Returns in *Last* the index value such that *From*(*Last*) is the last character read.

The exception *Data\_Error* is propagated if the sequence input does not have the required syntax, or if the value obtained is not of the subtype *Num*.

```

procedure Put (To   : out String;
               Item : in Num;
               Aft  : in Field := Default_Aft;
               Exp  : in Field := Default_Exp);

```

Outputs the value of the parameter *Item* to the given string, following the same rule as for output to a file, using a value for *Fore* such that the sequence of characters output exactly fills the string, including any leading spaces.

*Float\_Text\_IO* is a library package that is a nongeneric equivalent to *Text\_IO.Float\_IO* for the predefined type *Float*:

```

with Ada.Text_IO;
package Ada.Float_Text_IO is new Ada.Text_IO.Float_IO (Float);

```

For each predefined floating point type, a nongeneric equivalent to *Text\_IO.Float\_IO* is provided, with names such as *Ada.Long\_Float\_Text\_IO*.

#### Implementation Permissions

An implementation may extend *Get* and *Put* for floating point types to support special values such as infinities and NaNs.

The implementation of Put may produce an output value with no greater accuracy than that which is supported for the base subtype. The additional accuracy, if any, of the value produced by Put when the number of requested digits in the integer and fractional parts exceeds the required accuracy is implementation defined.

The nongeneric equivalent packages can be actual instantiations of the generic package for the appropriate predefined type, though that is not required.

NOTE 1 For an item with a positive value, if output to a string exactly fills the string without leading spaces, then output of the corresponding negative value will propagate Layout\_Error.

NOTE 2 The rules for the Value attribute (see 6.5) and the rules for Get are based on the same set of formats.

EXAMPLE *Examples of use of an instantiation of Text\_IO.Float\_IO:*

```
package Real_IO is new Float_IO(Real); use Real_IO;
-- default format used at instantiation, Default_Exp = 3
X : Real := -123.4567; -- digits 8 (see 6.5.7)
Put(X); -- default format "-1.2345670E+02"
Put(X, Fore => 5, Aft => 3, Exp => 2); -- "bbb-1.235E+2"
Put(X, 5, 3, 0); -- "b-123.457"
```

### A.10.10 Input-Output for Enumeration Types

#### Static Semantics

The following procedures are defined in the generic package Enumeration\_IO, which has to be instantiated for the appropriate enumeration type (indicated by Enum in the specification).

Values are output using either upper or lower case letters for identifiers. This is specified by the parameter Set, which is of the enumeration type Type\_Set.

```
type Type_Set is (Lower_Case, Upper_Case);
```

The format (which includes any trailing spaces) can be specified by an optional field width parameter. The default field width and letter case are defined by the following variables that are declared in the generic package Enumeration\_IO:

```
Default_Width : Field := 0;
Default_Setting : Type_Set := Upper_Case;
```

The following procedures are provided:

```
procedure Get(File : in File_Type; Item : out Enum);
procedure Get(Item : out Enum);
```

After skipping any leading blanks, line terminators, or page terminators, reads an identifier according to the syntax of this lexical element (lower and upper case being considered equivalent), or a character literal according to the syntax of this lexical element (including the apostrophes). Returns, in the parameter Item, the value of type Enum that corresponds to the sequence input.

The exception Data\_Error is propagated if the sequence input does not have the required syntax, or if the identifier or character literal does not correspond to a value of the subtype Enum.



```

procedure Put (File : in File_Type;
               Item : in Enum;
               Width : in Field := Default_Width;
               Set : in Type_Set := Default_Setting);

procedure Put (Item : in Enum;
               Width : in Field := Default_Width;
               Set : in Type_Set := Default_Setting);

```

Outputs the value of the parameter Item as an enumeration literal (either an identifier or a character literal). The optional parameter Set indicates whether lower case or upper case is used for identifiers; it has no effect for character literals. If the sequence of characters produced has fewer than Width characters, then trailing spaces are finally output to make up the difference. If Enum is a character type, the sequence of characters produced is as for Enum'Image(Item), as modified by the Width and Set parameters.

```

procedure Get (From : in String; Item : out Enum; Last : out Positive);

```

Reads an enumeration value from the beginning of the given string, following the same rule as the Get procedure that reads an enumeration value from a file, but treating the end of the string as a file terminator. Returns, in the parameter Item, the value of type Enum that corresponds to the sequence input. Returns in Last the index value such that From(Last) is the last character read.

The exception Data\_Error is propagated if the sequence input does not have the required syntax, or if the identifier or character literal does not correspond to a value of the subtype Enum.

```

procedure Put (To : out String;
               Item : in Enum;
               Set : in Type_Set := Default_Setting);

```

Outputs the value of the parameter Item to the given string, following the same rule as for output to a file, using the length of the given string as the value for Width.

Although the specification of the generic package Enumeration\_IO would allow instantiation for an integer type, this is not the intended purpose of this generic package, and the effect of such instantiations is not defined by the language.

NOTE 1 There is a difference between Put defined for characters, and for enumeration values. Thus

```

Ada.Text_IO.Put('A'); -- outputs the character A

package Char_IO is new Ada.Text_IO Enumeration_IO(Character);
Char_IO.Put('A'); -- outputs the character 'A', between apostrophes

```

NOTE 2 The type Boolean is an enumeration type, hence Enumeration\_IO can be instantiated for this type.

### A.10.11 Input-Output for Bounded Strings

The package Text\_IO.Bounded\_IO provides input-output in human-readable form for Bounded\_Strings.

#### Static Semantics

The generic library package Text\_IO.Bounded\_IO has the following declaration:

```

with Ada.Strings.Bounded;
generic
  with package Bounded is
    new Ada.Strings.Bounded.Generic_Bounded_Length (<>);
package Ada.Text_IO.Bounded_IO
  with Global => in out synchronized is

  procedure Put
    (File : in File_Type;
     Item : in Bounded.Bounded_String);

```

```

procedure Put
  (Item : in Bounded.Bounded_String);

procedure Put_Line
  (File : in File_Type;
   Item : in Bounded.Bounded_String);

procedure Put_Line
  (Item : in Bounded.Bounded_String);

function Get_Line
  (File : in File_Type)
  return Bounded.Bounded_String;

function Get_Line
  return Bounded.Bounded_String;

procedure Get_Line
  (File : in File_Type; Item : out Bounded.Bounded_String);

procedure Get_Line
  (Item : out Bounded.Bounded_String);

end Ada.Text_IO.Bounded_IO;

```

For an item of type Bounded\_String, the following subprograms are provided:

```

procedure Put
  (File : in File_Type;
   Item : in Bounded.Bounded_String);

  Equivalent to Text_IO.Put (File, Bounded.To_String(Item));

procedure Put
  (Item : in Bounded.Bounded_String);

  Equivalent to Text_IO.Put (Bounded.To_String(Item));

procedure Put_Line
  (File : in File_Type;
   Item : in Bounded.Bounded_String);

  Equivalent to Text_IO.Put_Line (File, Bounded.To_String(Item));

procedure Put_Line
  (Item : in Bounded.Bounded_String);

  Equivalent to Text_IO.Put_Line (Bounded.To_String(Item));

function Get_Line
  (File : in File_Type)
  return Bounded.Bounded_String;

  Returns Bounded.To_Bounded_String(Text_IO.Get_Line(File));

function Get_Line
  return Bounded.Bounded_String;

  Returns Bounded.To_Bounded_String(Text_IO.Get_Line);

procedure Get_Line
  (File : in File_Type; Item : out Bounded.Bounded_String);

  Equivalent to Item := Get_Line (File);

procedure Get_Line
  (Item : out Bounded.Bounded_String);

  Equivalent to Item := Get_Line;

```

### A.10.12 Input-Output for Unbounded Strings

The package Text\_IO.Unbounded\_IO provides input-output in human-readable form for Unbounded\_Strings.

The library package Text\_IO.Unbounded\_IO has the following declaration:

```

with Ada.Strings.Unbounded;
package Ada.Text_IO.Unbounded_IO
  with Global => in out synchronized is

  procedure Put
    (File : in File_Type;
     Item : in Strings.Unbounded.Unbounded_String);

  procedure Put
    (Item : in Strings.Unbounded.Unbounded_String);

  procedure Put_Line
    (File : in File_Type;
     Item : in Strings.Unbounded.Unbounded_String);

  procedure Put_Line
    (Item : in Strings.Unbounded.Unbounded_String);

  function Get_Line
    (File : in File_Type)
    return Strings.Unbounded.Unbounded_String;

  function Get_Line
    return Strings.Unbounded.Unbounded_String;

  procedure Get_Line
    (File : in File_Type; Item : out Strings.Unbounded.Unbounded_String);

  procedure Get_Line
    (Item : out Strings.Unbounded.Unbounded_String);

end Ada.Text_IO.Unbounded_IO;

```

For an item of type Unbounded\_String, the following subprograms are provided:

```

procedure Put
  (File : in File_Type;
   Item : in Strings.Unbounded.Unbounded_String);

  Equivalent to Text_IO.Put (File, Strings.Unbounded.To_String(Item));

procedure Put
  (Item : in Strings.Unbounded.Unbounded_String);

  Equivalent to Text_IO.Put (Strings.Unbounded.To_String(Item));

procedure Put_Line
  (File : in File_Type;
   Item : in Strings.Unbounded.Unbounded_String);

  Equivalent to Text_IO.Put_Line (File, Strings.Unbounded.To_String(Item));

procedure Put_Line
  (Item : in Strings.Unbounded.Unbounded_String);

  Equivalent to Text_IO.Put_Line (Strings.Unbounded.To_String(Item));

function Get_Line
  (File : in File_Type)
  return Strings.Unbounded.Unbounded_String;

  Returns Strings.Unbounded.To_Unbounded_String(Text_IO.Get_Line(File));

function Get_Line
  return Strings.Unbounded.Unbounded_String;

  Returns Strings.Unbounded.To_Unbounded_String(Text_IO.Get_Line);

procedure Get_Line
  (File : in File_Type; Item : out Strings.Unbounded.Unbounded_String);

  Equivalent to Item := Get_Line (File);

```

```
procedure Get_Line
  (Item : out Strings.Unbounded.Unbounded_String);
```

Equivalent to Item := Get\_Line;

## A.11 Wide Text Input-Output and Wide Wide Text Input-Output

The packages Wide\_Text\_IO and Wide\_Wide\_Text\_IO provide facilities for input and output in human-readable form. Each file is read or written sequentially, as a sequence of wide characters (or wide wide characters) grouped into lines, and as a sequence of lines grouped into pages.

### *Static Semantics*

The specification of package Wide\_Text\_IO is the same as that for Text\_IO, except that in each Get, Look\_Ahead, Get\_Immediate, Get\_Line, Put, and Put\_Line subprogram, any occurrence of Character is replaced by Wide\_Character, and any occurrence of String is replaced by Wide\_String. Nongeneric equivalents of Wide\_Text\_IO.Integer\_IO and Wide\_Text\_IO.Float\_IO are provided (as for Text\_IO) for each predefined numeric type, with names such as Ada.Integer\_Wide\_Text\_IO, Ada.Long\_Integer\_Wide\_Text\_IO, Ada.Float\_Wide\_Text\_IO, Ada.Long\_Float\_Wide\_Text\_IO.

The specification of package Wide\_Wide\_Text\_IO is the same as that for Text\_IO, except that in each Get, Look\_Ahead, Get\_Immediate, Get\_Line, Put, and Put\_Line subprogram, any occurrence of Character is replaced by Wide\_Wide\_Character, and any occurrence of String is replaced by Wide\_Wide\_String. Nongeneric equivalents of Wide\_Wide\_Text\_IO.Integer\_IO and Wide\_Wide\_Text\_IO.Float\_IO are provided (as for Text\_IO) for each predefined numeric type, with names such as Ada.Integer\_Wide\_Wide\_Text\_IO, Ada.Long\_Integer\_Wide\_Wide\_Text\_IO, Ada.Float\_Wide\_Wide\_Text\_IO, Ada.Long\_Float\_Wide\_Wide\_Text\_IO.

The specification of package Wide\_Text\_IO.Wide\_Bounded\_IO is the same as that for Text\_IO.Bounded\_IO, except that any occurrence of Bounded\_String is replaced by Bounded\_Wide\_String, and any occurrence of package Bounded is replaced by Wide\_Bounded. The specification of package Wide\_Wide\_Text\_IO.Wide\_Wide\_Bounded\_IO is the same as that for Text\_IO.Bounded\_IO, except that any occurrence of Bounded\_String is replaced by Bounded\_Wide\_Wide\_String, and any occurrence of package Bounded is replaced by Wide\_Wide\_Bounded.

The specification of package Wide\_Text\_IO.Wide\_Unbounded\_IO is the same as that for Text\_IO.Unbounded\_IO, except that any occurrence of Unbounded\_String is replaced by Unbounded\_Wide\_String, and any occurrence of package Unbounded is replaced by Wide\_Unbounded. The specification of package Wide\_Wide\_Text\_IO.Wide\_Wide\_Unbounded\_IO is the same as that for Text\_IO.Unbounded\_IO, except that any occurrence of Unbounded\_String is replaced by Unbounded\_Wide\_Wide\_String, and any occurrence of package Unbounded is replaced by Wide\_Wide\_Unbounded.

## A.12 Stream Input-Output

The packages Streams.Stream\_IO, Text\_IO.Text\_Streams, Wide\_Text\_IO.Text\_Streams, and Wide\_Wide\_Text\_IO.Text\_Streams provide stream-oriented operations on files.

### A.12.1 The Package Streams.Stream\_IO

The subprograms in the child package Streams.Stream\_IO provide control over stream files. Access to a stream file is either sequential, via a call on Read or Write to transfer an array of stream elements, or positional (if supported by the implementation for the given file), by specifying a relative index for an element. Since a stream file can be converted to a Stream\_Access value, calling stream-oriented attribute subprograms of different element types with the same Stream\_Access value provides heterogeneous input-output. See 16.13 for a general discussion of streams.

The elements of a stream file are stream elements. If positioning is supported for the specified external file, a current index and current size are maintained for the file as described in A.8. If positioning is not supported, a current index is not maintained, and the current size is implementation defined.

The library package Streams.Stream\_IO has the following declaration:

```

with Ada.IO_Exceptions;
package Ada.Streams.Stream_IO
  with Preelaborate, Global => in out synchronized is
  type Stream_Access is access all Root_Stream_Type'Class;
  type File_Type is limited private
    with Preelaborable_Initialization;
  type File_Mode is (In_File, Out_File, Append_File);
  type Count is range 0 .. implementation-defined;
  subtype Positive_Count is Count range 1 .. Count'Last;
  -- Index into file, in stream elements.
  procedure Create (File : in out File_Type;
    Mode : in File_Mode := Out_File;
    Name : in String := "";
    Form : in String := "");
  procedure Open (File : in out File_Type;
    Mode : in File_Mode;
    Name : in String;
    Form : in String := "");
  procedure Close (File : in out File_Type);
  procedure Delete (File : in out File_Type);
  procedure Reset (File : in out File_Type; Mode : in File_Mode);
  procedure Reset (File : in out File_Type);
  function Mode (File : in File_Type) return File_Mode;
  function Name (File : in File_Type) return String;
  function Form (File : in File_Type) return String;
  function Is_Open (File : in File_Type) return Boolean;
  function End_Of_File (File : in File_Type) return Boolean;
  function Stream (File : in File_Type) return Stream_Access;
  -- Return stream access for use with T'Input and T'Output
  -- Read array of stream elements from file
  procedure Read (File : in File_Type;
    Item : out Stream_Element_Array;
    Last : out Stream_Element_Offset;
    From : in Positive_Count)
    with Global => overriding in out File;
  procedure Read (File : in File_Type;
    Item : out Stream_Element_Array;
    Last : out Stream_Element_Offset)
    with Global => overriding in out File;
  -- Write array of stream elements into file
  procedure Write (File : in File_Type;
    Item : in Stream_Element_Array;
    To : in Positive_Count)
    with Global => overriding in out File;
  procedure Write (File : in File_Type;
    Item : in Stream_Element_Array)
    with Global => overriding in out File;
  -- Operations on position within file

```

```

procedure Set_Index(File : in File_Type; To : in Positive_Count)
    with Global => overriding in out File;

function Index(File : in File_Type) return Positive_Count;
function Size (File : in File_Type) return Count;

procedure Set_Mode(File : in out File_Type; Mode : in File_Mode);
procedure Flush(File : in File_Type);

-- exceptions
Status_Error : exception renames IO_Exceptions.Status_Error;
Mode_Error   : exception renames IO_Exceptions.Mode_Error;
Name_Error   : exception renames IO_Exceptions.Name_Error;
Use_Error    : exception renames IO_Exceptions.Use_Error;
Device_Error : exception renames IO_Exceptions.Device_Error;
End_Error    : exception renames IO_Exceptions.End_Error;
Data_Error   : exception renames IO_Exceptions.Data_Error;

package Wide_File_Names is
    -- File management

    procedure Create (File : in out File_Type;
        Mode : in File_Mode := Out_File;
        Name : in Wide_String := "";
        Form : in Wide_String := "");

    procedure Open (File : in out File_Type;
        Mode : in File_Mode;
        Name : in Wide_String;
        Form : in Wide_String := "");

    function Name (File : in File_Type) return Wide_String;
    function Form (File : in File_Type) return Wide_String;
end Wide_File_Names;

package Wide_Wide_File_Names is
    -- File management

    procedure Create (File : in out File_Type;
        Mode : in File_Mode := Out_File;
        Name : in Wide_Wide_String := "";
        Form : in Wide_Wide_String := "");

    procedure Open (File : in out File_Type;
        Mode : in File_Mode;
        Name : in Wide_Wide_String;
        Form : in Wide_Wide_String := "");

    function Name (File : in File_Type) return Wide_Wide_String;
    function Form (File : in File_Type) return Wide_Wide_String;
end Wide_Wide_File_Names;

private
    ... -- not specified by the language
end Ada.Streams.Stream_IO;

```

The type File\_Type needs finalization (see 10.6).

The subprograms given in A.8.2 for the control of external files (Create, Open, Close, Delete, Reset, Mode, Name, Form, Is\_Open, and Flush) are available for stream files.

The End\_Of\_File function:

- Propagates Mode\_Error if the mode of the file is not In\_File;
- If positioning is supported for the given external file, the function returns True if the current index exceeds the size of the external file; otherwise, it returns False;



- If positioning is not supported for the given external file, the function returns True if no more elements can be read from the given file; otherwise, it returns False.

The Set\_Mode procedure sets the mode of the file. If the new mode is Append\_File, the file is positioned to its end; otherwise, the position in the file is unchanged.

The Stream function returns a Stream\_Access result from a File\_Type object, thus allowing the stream-oriented attributes Read, Write, Input, and Output to be used on the same file for multiple types. Stream propagates Status\_Error if File is not open.

The procedures Read and Write are equivalent to the corresponding operations in the package Streams. Read propagates Mode\_Error if the mode of File is not In\_File. Write propagates Mode\_Error if the mode of File is not Out\_File or Append\_File. The Read procedure with a Positive\_Count parameter starts reading at the specified index. The Write procedure with a Positive\_Count parameter starts writing at the specified index. For a file that supports positioning, Read without a Positive\_Count parameter starts reading at the current index, and Write without a Positive\_Count parameter starts writing at the current index.

The Size function returns the current size of the file.

The Index function returns the current index.

The Set\_Index procedure sets the current index to the specified value.

If positioning is supported for the external file, the current index is maintained as follows:

- For Open and Create, if the Mode parameter is Append\_File, the current index is set to the current size of the file plus one; otherwise, the current index is set to one.
- For Reset, if the Mode parameter is Append\_File, or no Mode parameter is given and the current mode is Append\_File, the current index is set to the current size of the file plus one; otherwise, the current index is set to one.
- For Set\_Mode, if the new mode is Append\_File, the current index is set to current size plus one; otherwise, the current index is unchanged.
- For Read and Write without a Positive\_Count parameter, the current index is incremented by the number of stream elements read or written.
- For Read and Write with a Positive\_Count parameter, the value of the current index is set to the value of the Positive\_Count parameter plus the number of stream elements read or written.

If positioning is not supported for the given file, then a call of Index or Set\_Index propagates Use\_Error. Similarly, a call of Read or Write with a Positive\_Count parameter propagates Use\_Error.

#### *Erroneous Execution*

If the File\_Type object passed to the Stream function is later closed or finalized, and the stream-oriented attributes are subsequently called (explicitly or implicitly) on the Stream\_Access value returned by Stream, execution is erroneous. This rule applies even if the File\_Type object was opened again after it had been closed.

### **A.12.2 The Package Text\_IO.Text\_Streams**

The package Text\_IO.Text\_Streams provides a function for treating a text file as a stream.

#### *Static Semantics*

The library package Text\_IO.Text\_Streams has the following declaration:

```
with Ada.Streams;
package Ada.Text_IO.Text_Streams
  with Global => in out synchronized is
    type Stream_Access is access all Streams.Root_Stream_Type'Class;

    function Stream (File : in File_Type) return Stream_Access;
  end Ada.Text_IO.Text_Streams;
```

The Stream function has the same effect as the corresponding function in Streams.Stream\_IO.

NOTE 1 The ability to obtain a stream for a text file allows Current\_Input, Current\_Output, and Current\_Error to be processed with the functionality of streams, including the mixing of text and binary input-output, and the mixing of binary input-output for different types.

NOTE 2 Performing operations on the stream associated with a text file does not affect the column, line, or page counts.

### A.12.3 The Package Wide\_Text\_IO.Text\_Streams

The package Wide\_Text\_IO.Text\_Streams provides a function for treating a wide text file as a stream.

*Static Semantics*

The library package Wide\_Text\_IO.Text\_Streams has the following declaration:

```
with Ada.Streams;
package Ada.Wide_Text_IO.Text_Streams
  with Global => in out synchronized is
    type Stream_Access is access all Streams.Root_Stream_Type'Class;

    function Stream (File : in File_Type) return Stream_Access;
  end Ada.Wide_Text_IO.Text_Streams;
```

The Stream function has the same effect as the corresponding function in Streams.Stream\_IO.

### A.12.4 The Package Wide\_Wide\_Text\_IO.Text\_Streams

The package Wide\_Wide\_Text\_IO.Text\_Streams provides a function for treating a wide wide text file as a stream.

*Static Semantics*

The library package Wide\_Wide\_Text\_IO.Text\_Streams has the following declaration:

```
with Ada.Streams;
package Ada.Wide_Wide_Text_IO.Text_Streams
  with Global => in out synchronized is
    type Stream_Access is access all Streams.Root_Stream_Type'Class;

    function Stream (File : in File_Type) return Stream_Access;
  end Ada.Wide_Wide_Text_IO.Text_Streams;
```

The Stream function has the same effect as the corresponding function in Streams.Stream\_IO.

## A.13 Exceptions in Input-Output

The package IO\_Exceptions defines the exceptions used by the predefined input-output packages.

*Static Semantics*

The library package IO\_Exceptions has the following declaration:

```
package Ada.IO_Exceptions
  with Pure is
    Status_Error : exception;
    Mode_Error   : exception;
    Name_Error   : exception;
    Use_Error     : exception;
    Device_Error : exception;
    End_Error     : exception;
    Data_Error    : exception;
    Layout_Error  : exception;

  end Ada.IO_Exceptions;
```

If more than one error condition exists, the corresponding exception that appears earliest in the following list is the one that is propagated.

The exception `Status_Error` is propagated by an attempt to operate upon a file that is not open, and by an attempt to open a file that is already open.

The exception `Mode_Error` is propagated by an attempt to read from, or test for the end of, a file whose current mode is `Out_File` or `Append_File`, and also by an attempt to write to a file whose current mode is `In_File`. In the case of `Text_IO`, the exception `Mode_Error` is also propagated by specifying a file whose current mode is `Out_File` or `Append_File` in a call of `Set_Input`, `Skip_Line`, `End_Of_Line`, `Skip_Page`, or `End_Of_Page`; and by specifying a file whose current mode is `In_File` in a call of `Set_Output`, `Set_Line_Length`, `Set_Page_Length`, `Line_Length`, `Page_Length`, `New_Line`, or `New_Page`.

The exception `Name_Error` is propagated by a call of `Create` or `Open` if the string given for the parameter `Name` does not allow the identification of an external file. For example, this exception is propagated if the string is improper, or, alternatively, if either none or more than one external file corresponds to the string.

The exception `Use_Error` is propagated if an operation is attempted that is not possible for reasons that depend on characteristics of the external file. For example, this exception is propagated by the procedure `Create`, among other circumstances, if the given mode is `Out_File` but the form specifies an input only device, if the parameter `Form` specifies invalid access rights, or if an external file with the given name already exists and overwriting is not allowed.

The exception `Device_Error` is propagated if an input-output operation cannot be completed because of a malfunction of the underlying system.

The exception `End_Error` is propagated by an attempt to skip (read past) the end of a file.

The exception `Data_Error` can be propagated by the procedure `Read` (or by the `Read` attribute) if the element read cannot be interpreted as a value of the required subtype. This exception is also propagated by a procedure `Get` (defined in the package `Text_IO`) if the input character sequence fails to satisfy the required syntax, or if the value input does not belong to the range of the required subtype.

The exception `Layout_Error` is propagated (in text input-output) by `Col`, `Line`, or `Page` if the value returned exceeds `Count'Last`. The exception `Layout_Error` is also propagated on output by an attempt to set column or line numbers in excess of specified maximum line or page lengths, respectively (excluding the unbounded cases). It is also propagated by an attempt to `Put` too many characters to a string.

These exceptions are also propagated by various other language-defined packages and operations, see the definition of those entities for other reasons that these exceptions are propagated.

#### *Documentation Requirements*

The implementation shall document the conditions under which `Name_Error`, `Use_Error` and `Device_Error` are propagated.

#### *Implementation Permissions*

When the associated check is complex, it is optional to propagate `Data_Error` as part of a procedure `Read` (or the `Read` attribute) when the value read cannot be interpreted as a value of the required subtype.

#### *Erroneous Execution*

If the element read by the procedure `Read` (or by the `Read` attribute) cannot be interpreted as a value of the required subtype, but this is not detected and `Data_Error` is not propagated, then the resulting value can be abnormal, and subsequent references to the value can lead to erroneous execution, as explained in 16.9.1.

## A.14 File Sharing

### Dynamic Semantics

It is not specified by the language whether the same external file can be associated with more than one file object. If such sharing is supported by the implementation, the following effects are defined:

- Operations on one text file object do not affect the column, line, and page numbers of any other file object.
- For direct and stream files, the current index is a property of each file object; an operation on one file object does not affect the current index of any other file object.
- For direct and stream files, the current size of the file is a property of the external file.

All other effects are identical.

## A.15 The Package Command\_Line

The package `Command_Line` allows a program to obtain the values of its arguments and to set the exit status code to be returned on normal termination.

### Static Semantics

The library package `Ada.Command_Line` has the following declaration:

```
package Ada.Command_Line
  with Preelaborate, Nonblocking, Global => in out synchronized is
    function Argument_Count return Natural;
    function Argument (Number : in Positive) return String;
    function Command_Name return String;
    type Exit_Status is implementation-defined integer type;
    Success : constant Exit_Status;
    Failure : constant Exit_Status;
    procedure Set_Exit_Status (Code : in Exit_Status);
private
  ... -- not specified by the language
end Ada.Command_Line;
```

```
function Argument_Count return Natural;
```

If the external execution environment supports passing arguments to a program, then `Argument_Count` returns the number of arguments passed to the program invoking the function. Otherwise, it returns 0. The meaning of “number of arguments” is implementation defined.

```
function Argument (Number : in Positive) return String;
```

If the external execution environment supports passing arguments to a program, then `Argument` returns an implementation-defined value with lower bound 1 corresponding to the argument at relative position `Number`. If `Number` is outside the range `1..Argument_Count`, then `Constraint_Error` is propagated.

```
function Command_Name return String;
```

If the external execution environment supports passing arguments to a program, then `Command_Name` returns an implementation-defined value with lower bound 1 corresponding to the name of the command invoking the program; otherwise, `Command_Name` returns the null string.

**type** Exit\_Status **is** *implementation-defined integer type*;

The type Exit\_Status represents the range of exit status values supported by the external execution environment. The constants Success and Failure correspond to success and failure, respectively.

**procedure** Set\_Exit\_Status (Code : **in** Exit\_Status);

If the external execution environment supports returning an exit status from a program, then Set\_Exit\_Status sets Code as the status. Normal termination of a program returns as the exit status the value most recently set by Set\_Exit\_Status, or, if no such value has been set, then the value Success. If a program terminates abnormally, the status set by Set\_Exit\_Status is ignored, and an implementation-defined exit status value is set.

If the external execution environment does not support returning an exit value from a program, then Set\_Exit\_Status does nothing.

#### *Implementation Permissions*

An alternative declaration is allowed for package Command\_Line if different functionality is appropriate for the external execution environment.

NOTE Argument\_Count, Argument, and Command\_Name correspond to the C language's argc, argv[n] (for n>0) and argv[0], respectively.

### **A.15.1 The Packages Wide\_Command\_Line and Wide\_Wide\_Command\_Line**

The packages Wide\_Command\_Line and Wide\_Wide\_Command\_Line allow a program to obtain the values of its arguments and to set the exit status code to be returned on normal termination.

#### *Static Semantics*

The specification of package Wide\_Command\_Line is the same as for Command\_Line, except that each occurrence of String is replaced by Wide\_String.

The specification of package Wide\_Wide\_Command\_Line is the same as for Command\_Line, except that each occurrence of String is replaced by Wide\_Wide\_String.

### **A.16 The Package Directories**

The package Directories provides operations for manipulating files and directories, and their names.

#### *Static Semantics*

The library package Directories has the following declaration:

```
with Ada.IO_Exceptions;
with Ada.Calendar;
package Ada.Directories
  with Global => in out synchronized is
    -- Directory and file operations:

    function Current_Directory return String;

    procedure Set_Directory (Directory : in String);

    procedure Create_Directory (New_Directory : in String;
                               Form           : in String := "");

    procedure Delete_Directory (Directory : in String);

    procedure Create_Path (New_Directory : in String;
                          Form           : in String := "");

    procedure Delete_Tree (Directory : in String);

    procedure Delete_File (Name : in String);

    procedure Rename (Old_Name, New_Name : in String);
```

```

procedure Copy_File (Source_Name,
                     Target_Name : in String;
                     Form         : in String := "");

-- File and directory name operations:
function Full_Name (Name : in String) return String
  with Nonblocking;

function Simple_Name (Name : in String) return String
  with Nonblocking;

function Containing_Directory (Name : in String) return String
  with Nonblocking;

function Extension (Name : in String) return String
  with Nonblocking;

function Base_Name (Name : in String) return String
  with Nonblocking;

function Compose (Containing_Directory : in String := "";
                  Name                  : in String;
                  Extension              : in String := "") return String
  with Nonblocking;

type Name_Case_Kind is
  (Unknown, Case_Sensitive, Case_Insensitive, Case_Preserving);

function Name_Case_Equivalence (Name : in String) return Name_Case_Kind;

-- File and directory queries:
type File_Kind is (Directory, Ordinary_File, Special_File);
type File_Size is range 0 .. implementation-defined;
function Exists (Name : in String) return Boolean;
function Kind (Name : in String) return File_Kind;
function Size (Name : in String) return File_Size;
function Modification_Time (Name : in String) return Ada.Calendar.Time;

-- Directory searching:
type Directory_Entry_Type is limited private;
type Filter_Type is array (File_Kind) of Boolean;
type Search_Type is limited private;

procedure Start_Search (Search      : in out Search_Type;
                       Directory    : in String;
                       Pattern      : in String;
                       Filter       : in Filter_Type := (others => True));

procedure End_Search (Search : in out Search_Type);

function More_Entries (Search : in Search_Type) return Boolean;

procedure Get_Next_Entry (Search : in out Search_Type;
                         Directory_Entry : out Directory_Entry_Type);

procedure Search (
  Directory : in String;
  Pattern   : in String;
  Filter    : in Filter_Type := (others => True);
  Process   : not null access procedure (
    Directory_Entry : in Directory_Entry_Type))
  with Allows_Exit;

-- Operations on Directory Entries:

function Simple_Name (Directory_Entry : in Directory_Entry_Type)
  return String;

```



```

function Full_Name (Directory_Entry : in Directory_Entry_Type)
    return String;

function Kind (Directory_Entry : in Directory_Entry_Type)
    return File_Kind;

function Size (Directory_Entry : in Directory_Entry_Type)
    return File_Size;

function Modification_Time (Directory_Entry : in Directory_Entry_Type)
    return Ada.Calendar.Time;

Status_Error : exception renames Ada.IO_Exceptions.Status_Error;
Name_Error   : exception renames Ada.IO_Exceptions.Name_Error;
Use_Error    : exception renames Ada.IO_Exceptions.Use_Error;
Device_Error : exception renames Ada.IO_Exceptions.Device_Error;

private
    ... -- not specified by the language
end Ada.Directories;

```

External files may be classified as directories, special files, or ordinary files. A *directory* is an external file that is a container for files on the target system. A *special file* is an external file that cannot be created or read by a predefined Ada input-output package. External files that are not special files or directories are called *ordinary files*.

A *file name* is a string identifying an external file. Similarly, a *directory name* is a string identifying a directory. The interpretation of file names and directory names is implementation defined.

The *full name* of an external file is a full specification of the name of the file. If the external environment allows alternative specifications of the name (for example, abbreviations), the full name should not use such alternatives. A full name typically will include the names of all of the directories that contain the item. The *simple name* of an external file is the name of the item, not including any containing directory names. Unless otherwise specified, a file name or directory name parameter in a call to a predefined Ada input-output subprogram can be a full name, a simple name, or any other form of name supported by the implementation.

A *root directory* is a directory that has no containing directory.

The *default directory* is the directory that is used if a directory or file name is not a full name (that is, when the name does not fully identify all of the containing directories).

A *directory entry* is a single item in a directory, identifying a single external file (including directories and special files).

For each function that returns a string, the lower bound of the returned value is 1.

The following file and directory operations are provided:

```

function Current_Directory return String;

```

Returns the full directory name for the current default directory. The name returned shall be suitable for a future call to Set\_Directory. The exception Use\_Error is propagated if a default directory is not supported by the external environment.

```

procedure Set_Directory (Directory : in String);

```

Sets the current default directory. The exception Name\_Error is propagated if the string given as Directory does not identify an existing directory. The exception Use\_Error is propagated if the external environment does not support making Directory (in the absence of Name\_Error) a default directory.

```

procedure Create_Directory (New_Directory : in String;
                           Form           : in String := "");

```

Creates a directory with name New\_Directory. The Form parameter can be used to give system-dependent characteristics of the directory; the interpretation of the Form parameter is implementation defined. A null string for Form specifies the use of the default options of the

implementation of the new directory. The exception `Name_Error` is propagated if the string given as `New_Directory` does not allow the identification of a directory. The exception `Use_Error` is propagated if the external environment does not support the creation of a directory with the given name (in the absence of `Name_Error`) and form.

**procedure** `Delete_Directory` (`Directory` : **in** `String`);

Deletes an existing empty directory with name `Directory`. The exception `Name_Error` is propagated if the string given as `Directory` does not identify an existing directory. The exception `Use_Error` is propagated if the directory is not empty or the external environment does not support the deletion of the directory with the given name (in the absence of `Name_Error`).

**procedure** `Create_Path` (`New_Directory` : **in** `String`;  
                          `Form` : **in** `String` := "");

Creates zero or more directories with name `New_Directory`. Each nonexistent directory named by `New_Directory` is created. For example, on a typical Unix system, `Create_Path ("/usr/me/my")`; would create directory "me" in directory "usr", then create directory "my" in directory "me". The `Form` parameter can be used to give system-dependent characteristics of the directory; the interpretation of the `Form` parameter is implementation defined. A null string for `Form` specifies the use of the default options of the implementation of the new directory. The exception `Name_Error` is propagated if the string given as `New_Directory` does not allow the identification of any directory. The exception `Use_Error` is propagated if the external environment does not support the creation of any directories with the given name (in the absence of `Name_Error`) and form. If `Use_Error` is propagated, it is unspecified whether a portion of the directory path is created.

**procedure** `Delete_Tree` (`Directory` : **in** `String`);

Deletes an existing directory with name `Directory`. The directory and all of its contents (possibly including other directories) are deleted. The exception `Name_Error` is propagated if the string given as `Directory` does not identify an existing directory. The exception `Use_Error` is propagated if the external environment does not support the deletion of the directory or some portion of its contents with the given name (in the absence of `Name_Error`). If `Use_Error` is propagated, it is unspecified whether a portion of the contents of the directory is deleted.

**procedure** `Delete_File` (`Name` : **in** `String`);

Deletes an existing ordinary or special file with name `Name`. The exception `Name_Error` is propagated if the string given as `Name` does not identify an existing ordinary or special external file. The exception `Use_Error` is propagated if the external environment does not support the deletion of the file with the given name (in the absence of `Name_Error`).

**procedure** `Rename` (`Old_Name`, `New_Name` : **in** `String`);

Renames an existing external file (including directories) with name `Old_Name` to `New_Name`. The exception `Name_Error` is propagated if the string given as `Old_Name` does not identify an existing external file or if the string given as `New_Name` does not allow the identification of an external file. The exception `Use_Error` is propagated if the external environment does not support the renaming of the file with the given name (in the absence of `Name_Error`). In particular, `Use_Error` is propagated if a file or directory already exists with name `New_Name`.

**procedure** `Copy_File` (`Source_Name`,  
                          `Target_Name` : **in** `String`;  
                          `Form` : **in** `String` := "");

Copies the contents of the existing external file with name `Source_Name` to an external file with name `Target_Name`. The resulting external file is a duplicate of the source external file. The `Form` parameter can be used to give system-dependent characteristics of the resulting external file; the interpretation of the `Form` parameter is implementation defined. Exception `Name_Error` is propagated if the string given as `Source_Name` does not identify an existing external ordinary or special file, or if the string given as `Target_Name` does not allow the

identification of an external file. The exception `Use_Error` is propagated if the external environment does not support creating the file with the name given by `Target_Name` and form given by `Form`, or copying of the file with the name given by `Source_Name` (in the absence of `Name_Error`). If `Use_Error` is propagated, it is unspecified whether a portion of the file is copied.

The following file and directory name operations are provided:

**function** `Full_Name` (`Name` : **in** `String`) **return** `String`;

Returns the full name corresponding to the file name specified by `Name`. The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an external file (including directories and special files).

**function** `Simple_Name` (`Name` : **in** `String`) **return** `String`;

Returns the simple name portion of the file name specified by `Name`. The simple name of a root directory is a name of the root itself. The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an external file (including directories and special files).

**function** `Containing_Directory` (`Name` : **in** `String`) **return** `String`;

Returns the name of the containing directory of the external file (including directories) identified by `Name`. (If more than one directory can contain `Name`, the directory name returned is implementation defined.) The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an external file. The exception `Use_Error` is propagated if the external file does not have a containing directory.

**function** `Extension` (`Name` : **in** `String`) **return** `String`;

Returns the extension name corresponding to `Name`. The extension name is a portion of a simple name (not including any separator characters), typically used to identify the file class. If the external environment does not have extension names, then the null string is returned. The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an external file.

**function** `Base_Name` (`Name` : **in** `String`) **return** `String`;

Returns the base name corresponding to `Name`. The base name is the remainder of a simple name after removing any extension and extension separators. The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an external file (including directories and special files).

**function** `Compose` (`Containing_Directory` : **in** `String` := "";  
                   `Name` : **in** `String`;  
                   `Extension` : **in** `String` := "") **return** `String`;

Returns the name of the external file with the specified `Containing_Directory`, `Name`, and `Extension`. If `Extension` is the null string, then `Name` is interpreted as a simple name; otherwise, `Name` is interpreted as a base name. The exception `Name_Error` is propagated if:

- the string given as `Containing_Directory` is not null and does not allow the identification of a directory;
- the string given as `Extension` is not null and is not a possible extension;
- the string given as `Name` is not a possible simple name (if `Extension` is null) or base name (if `Extension` is nonnull); or
- the string given as `Name` is a root directory, and `Containing_Directory` or `Extension` is nonnull.

**function** `Name_Case_Equivalence` (`Name` : **in** `String`) **return** `Name_Case_Kind`;

Returns the file name equivalence rule for the directory containing `Name`. Raises `Name_Error` if `Name` is not a full name. Returns `Case_Sensitive` if file names that differ only in the case of

letters are considered different names. If file names that differ only in the case of letters are considered the same name, then `Case_Preserving` is returned if names have the case of the file name used when a file is created; and `Case_Insensitive` is returned otherwise. Returns `Unknown` if the file name equivalence is not known.

The following file and directory queries and types are provided:

```
type File_Kind is (Directory, Ordinary_File, Special_File);
```

The type `File_Kind` represents the kind of file represented by an external file or directory.

```
type File_Size is range 0 .. implementation-defined;
```

The type `File_Size` represents the size of an external file.

```
function Exists (Name : in String) return Boolean;
```

Returns `True` if an external file represented by `Name` exists, and `False` otherwise. The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an external file (including directories and special files).

```
function Kind (Name : in String) return File_Kind;
```

Returns the kind of external file represented by `Name`. The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an existing external file.

```
function Size (Name : in String) return File_Size;
```

Returns the size of the external file represented by `Name`. The size of an external file is the number of stream elements contained in the file. If the external file is not an ordinary file, the result is implementation defined. The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an existing external file. The exception `Constraint_Error` is propagated if the file size is not a value of type `File_Size`.

```
function Modification_Time (Name : in String) return Ada.Calendar.Time;
```

Returns the time that the external file represented by `Name` was most recently modified. If the external file is not an ordinary file, the result is implementation defined. The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an existing external file. The exception `Use_Error` is propagated if the external environment does not support reading the modification time of the file with the name given by `Name` (in the absence of `Name_Error`).

The following directory searching operations and types are provided:

```
type Directory_Entry_Type is limited private;
```

The type `Directory_Entry_Type` represents a single item in a directory. These items can only be created by the `Get_Next_Entry` procedure in this package. Information about the item can be obtained from the functions declared in this package. A default-initialized object of this type is invalid; objects returned from `Get_Next_Entry` are valid.

```
type Filter_Type is array (File_Kind) of Boolean;
```

The type `Filter_Type` specifies which directory entries are provided from a search operation. If the `Directory` component is `True`, directory entries representing directories are provided. If the `Ordinary_File` component is `True`, directory entries representing ordinary files are provided. If the `Special_File` component is `True`, directory entries representing special files are provided.

```
type Search_Type is limited private;
```

The type `Search_Type` contains the state of a directory search. A default-initialized `Search_Type` object has no entries available (function `More_Entries` returns `False`). Type `Search_Type` needs finalization (see 10.6).

```

procedure Start_Search (Search      : in out Search_Type;
                        Directory    : in String;
                        Pattern      : in String;
                        Filter       : in Filter_Type := (others => True));

```

Starts a search in the directory named by Directory for entries matching Pattern and Filter. Pattern represents a pattern for matching file names. If Pattern is the null string, all items in the directory are matched; otherwise, the interpretation of Pattern is implementation defined. Only items that match Filter will be returned. After a successful call on Start\_Search, the object Search may have entries available, but it may have no entries available if no files or directories match Pattern and Filter. The exception Name\_Error is propagated if the string given by Directory does not identify an existing directory, or if Pattern does not allow the identification of any possible external file or directory. The exception Use\_Error is propagated if the external environment does not support the searching of the directory with the given name (in the absence of Name\_Error). When Start\_Search propagates Name\_Error or Use\_Error, the object Search will have no entries available.

```

procedure End_Search (Search : in out Search_Type);

```

Ends the search represented by Search. After a successful call on End\_Search, the object Search will have no entries available.

```

function More_Entries (Search : in Search_Type) return Boolean;

```

Returns True if more entries are available to be returned by a call to Get\_Next\_Entry for the specified search object, and False otherwise.

```

procedure Get_Next_Entry (Search : in out Search_Type;
                        Directory_Entry : out Directory_Entry_Type);

```

Returns the next Directory\_Entry for the search described by Search that matches the pattern and filter. If no further matches are available, Status\_Error is raised. It is implementation defined as to whether the results returned by this subprogram are altered if the contents of the directory are altered while the Search object is valid (for example, by another program). The exception Use\_Error is propagated if the external environment does not support continued searching of the directory represented by Search.

```

procedure Search (
    Directory : in String;
    Pattern   : in String;
    Filter    : in Filter_Type := (others => True);
    Process   : not null access procedure (
        Directory_Entry : in Directory_Entry_Type))
with Allows Exit;

```

Searches in the directory named by Directory for entries matching Pattern and Filter. The subprogram designated by Process is called with each matching entry in turn. Pattern represents a pattern for matching file names. If Pattern is the null string, all items in the directory are matched; otherwise, the interpretation of Pattern is implementation defined. Only items that match Filter will be returned. The exception Name\_Error is propagated if the string given by Directory does not identify an existing directory, or if Pattern does not allow the identification of any possible external file or directory. The exception Use\_Error is propagated if the external environment does not support the searching of the directory with the given name (in the absence of Name\_Error).

```

function Simple_Name (Directory_Entry : in Directory_Entry_Type)
return String;

```

Returns the simple external name of the external file (including directories) represented by Directory\_Entry. The format of the name returned is implementation defined. The exception Status\_Error is propagated if Directory\_Entry is invalid.



```
function Full_Name (Directory_Entry : in Directory_Entry_Type)
    return String;
```

Returns the full external name of the external file (including directories) represented by Directory\_Entry. The format of the name returned is implementation defined. The exception Status\_Error is propagated if Directory\_Entry is invalid.

```
function Kind (Directory_Entry : in Directory_Entry_Type)
    return File_Kind;
```

Returns the kind of external file represented by Directory\_Entry. The exception Status\_Error is propagated if Directory\_Entry is invalid.

```
function Size (Directory_Entry : in Directory_Entry_Type)
    return File_Size;
```

Returns the size of the external file represented by Directory\_Entry. The size of an external file is the number of stream elements contained in the file. If the external file represented by Directory\_Entry is not an ordinary file, the result is implementation defined. The exception Status\_Error is propagated if Directory\_Entry is invalid. The exception Constraint\_Error is propagated if the file size is not a value of type File\_Size.

```
function Modification_Time (Directory_Entry : in Directory_Entry_Type)
    return Ada.Calendar.Time;
```

Returns the time that the external file represented by Directory\_Entry was most recently modified. If the external file represented by Directory\_Entry is not an ordinary file, the result is implementation defined. The exception Status\_Error is propagated if Directory\_Entry is invalid. The exception Use\_Error is propagated if the external environment does not support reading the modification time of the file represented by Directory\_Entry.

#### *Implementation Requirements*

For Copy\_File, if Source\_Name identifies an existing external ordinary file created by a predefined Ada input-output package, and Target\_Name and Form can be used in the Create operation of that input-output package with mode Out\_File without raising an exception, then Copy\_File shall not propagate Use\_Error.

#### *Implementation Advice*

If other information about a file (such as the owner or creation date) is available in a directory entry, the implementation should provide functions in a child package Directories.Information to retrieve it.

Start\_Search and Search should raise Name\_Error if Pattern is malformed, but not if it can represent a file in the directory but does not actually do so.

Rename should be supported at least when both New\_Name and Old\_Name are simple names and New\_Name does not identify an existing external file.

NOTE 1 The operations Containing\_Directory, Full\_Name, Simple\_Name, Base\_Name, Extension, and Compose operate on file names, not external files. The files identified by these operations do not necessarily exist. Name\_Error is raised only if the file name is malformed and cannot possibly identify a file. Of these operations, only the result of Full\_Name depends on the current default directory; the result of the others depends only on their parameters.

NOTE 2 Using access types, values of Search\_Type and Directory\_Entry\_Type can be saved and queried later. However, another task or application can modify or delete the file represented by a Directory\_Entry\_Type value or the directory represented by a Search\_Type value; such a value can only give the information valid at the time it is created. Therefore, long-term storage of these values is not recommended.

NOTE 3 If the target system does not support directories inside of directories, then Kind will never return Directory and Containing\_Directory will always raise Use\_Error.

NOTE 4 If the target system does not support creation or deletion of directories, then Create\_Directory, Create\_Path, Delete\_Directory, and Delete\_Tree will always propagate Use\_Error.



### A.16.1 The Package Directories.Hierarchical\_File\_Names

The library package Directories.Hierarchical\_File\_Names is an optional package providing operations for file name construction and decomposition for targets with hierarchical file naming.

*Static Semantics*

If provided, the library package Directories.Hierarchical\_File\_Names has the following declaration:

```
package Ada.Directories.Hierarchical_File_Names
  with Nonblocking, Global => in out synchronized is

  function Is_Simple_Name (Name : in String) return Boolean;
  function Is_Root_Directory_Name (Name : in String) return Boolean;
  function Is_Parent_Directory_Name (Name : in String) return Boolean;
  function Is_Current_Directory_Name (Name : in String) return Boolean;
  function Is_Full_Name (Name : in String) return Boolean;
  function Is_Relative_Name (Name : in String) return Boolean;
  function Simple_Name (Name : in String) return String
    renames Ada.Directories.Simple_Name;
  function Containing_Directory (Name : in String) return String
    renames Ada.Directories.Containing_Directory;
  function Initial_Directory (Name : in String) return String;
  function Relative_Name (Name : in String) return String;
  function Compose (Directory      : in String := "";
                  Relative_Name  : in String;
                  Extension      : in String := "") return String;

end Ada.Directories.Hierarchical_File_Names;
```

In addition to the operations provided in package Directories.Hierarchical\_File\_Names, the operations in package Directories can be used with hierarchical file names. In particular, functions Full\_Name, Base\_Name, and Extension provide additional capabilities for hierarchical file names.

```
function Is_Simple_Name (Name : in String) return Boolean;
```

Returns True if Name is a simple name, and returns False otherwise.

```
function Is_Root_Directory_Name (Name : in String) return Boolean;
```

Returns True if Name is syntactically a root (a directory that cannot be decomposed further), and returns False otherwise.

```
function Is_Parent_Directory_Name (Name : in String) return Boolean;
```

Returns True if Name can be used to indicate symbolically the parent directory of any directory, and returns False otherwise.

```
function Is_Current_Directory_Name (Name : in String) return Boolean;
```

Returns True if Name can be used to indicate symbolically the directory itself for any directory, and returns False otherwise.

```
function Is_Full_Name (Name : in String) return Boolean;
```

Returns True if the leftmost directory part of Name is a root, and returns False otherwise.

```
function Is_Relative_Name (Name : in String) return Boolean;
```

Returns True if Name allows the identification of an external file (including directories and special files) but is not a full name, and returns False otherwise.

```
function Initial_Directory (Name : in String) return String;
```

Returns the leftmost directory part in Name. That is, it returns a root directory name (for a full name), or one of a parent directory name, a current directory name, or a simple name (for a

relative name). The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an external file (including directories and special files).

```
function Relative_Name (Name : in String) return String;
```

Returns the entire file name except the `Initial_Directory` portion. The exception `Name_Error` is propagated if the string given as `Name` does not allow the identification of an external file (including directories and special files), or if `Name` has a single part (this includes if any of `Is_Simple_Name`, `Is_Root_Directory_Name`, `Is_Parent_Directory_Name`, or `Is_Current_Directory_Name` are True).

```
function Compose (Directory      : in String := "";
                  Relative_Name  : in String;
                  Extension      : in String := "") return String;
```

Returns the name of the external file with the specified `Directory`, `Relative_Name`, and `Extension`. The exception `Name_Error` is propagated if the string given as `Directory` is not the null string and does not allow the identification of a directory, or if `Is_Relative_Name` (`Relative_Name`) is False, or if the string given as `Extension` is not the null string and is not a possible extension, or if `Extension` is not the null string and `Simple_Name` (`Relative_Name`) is not a base name.

The result of `Compose` is a full name if `Is_Full_Name` (`Directory`) is True; result is a relative name otherwise.

#### *Implementation Advice*

`Directories.Hierarchical_File_Names` should be provided for systems with hierarchical file naming, and should not be provided on other systems.

NOTE 1 These operations operate on file names, not external files. The files identified by these operations do not necessarily exist. `Name_Error` is raised only as specified or if the file name is malformed and cannot possibly identify a file. The result of these operations depends only on their parameters.

NOTE 2 `Containing_Directory` raises `Use_Error` if `Name` does not have a containing directory, including when any of `Is_Simple_Name`, `Is_Root_Directory_Name`, `Is_Parent_Directory_Name`, or `Is_Current_Directory_Name` are True.

### **A.16.2 The Packages `Wide_Directories` and `Wide_Wide_Directories`**

The packages `Wide_Directories` and `Wide_Wide_Directories` provide operations for manipulating files and directories, and their names.

#### *Static Semantics*

The specification of package `Wide_Directories` is the same as for `Directories` (including its optional child packages `Information` and `Hierarchical_File_Names`), except that each occurrence of `String` is replaced by `Wide_String`.

The specification of package `Wide_Wide_Directories` is the same as for `Directories` (including its optional child packages `Information` and `Hierarchical_File_Names`), except that each occurrence of `String` is replaced by `Wide_Wide_String`.

### **A.17 The Package `Environment_Variables`**

The package `Environment_Variables` allows a program to read or modify environment variables. Environment variables are name-value pairs, where both the name and value are strings. The definition of what constitutes an *environment variable*, and the meaning of the name and value, are implementation defined.

#### *Static Semantics*

The library package `Environment_Variables` has the following declaration:

```
package Ada.Environment_Variables
with Preelaborate, Nonblocking, Global => in out synchronized is
```

```

function Value (Name : in String) return String;
function Value (Name : in String; Default : in String) return String;
function Exists (Name : in String) return Boolean;
procedure Set (Name : in String; Value : in String);
procedure Clear (Name : in String);
procedure Clear;
procedure Iterate
  (Process : not null access procedure (Name, Value : in String))
  with Allows_Exit;
end Ada.Environment_Variables;

function Value (Name : in String) return String;

```

If the external execution environment supports environment variables, then Value returns the value of the environment variable with the given name. If no environment variable with the given name exists, then Constraint\_Error is propagated. If the execution environment does not support environment variables, then Program\_Error is propagated.

```

function Value (Name : in String; Default : in String) return String;

```

If the external execution environment supports environment variables and an environment variable with the given name currently exists, then Value returns its value; otherwise, it returns Default.

```

function Exists (Name : in String) return Boolean;

```

If the external execution environment supports environment variables and an environment variable with the given name currently exists, then Exists returns True; otherwise, it returns False.

```

procedure Set (Name : in String; Value : in String);

```

If the external execution environment supports environment variables, then Set first clears any existing environment variable with the given name, and then defines a single new environment variable with the given name and value. Otherwise, Program\_Error is propagated.

If implementation-defined circumstances prohibit the definition of an environment variable with the given name and value, then Constraint\_Error is propagated.

It is implementation defined whether there exist values for which the call Set(Name, Value) has the same effect as Clear (Name).

```

procedure Clear (Name : in String);

```

If the external execution environment supports environment variables, then Clear deletes all existing environment variables with the given name. Otherwise, Program\_Error is propagated.

```

procedure Clear;

```

If the external execution environment supports environment variables, then Clear deletes all existing environment variables. Otherwise, Program\_Error is propagated.

```

procedure Iterate
  (Process : not null access procedure (Name, Value : in String))
  with Allows_Exit;

```

If the external execution environment supports environment variables, then Iterate calls the subprogram designated by Process for each existing environment variable, passing the name and value of that environment variable. Otherwise, Program\_Error is propagated.

If several environment variables exist that have the same name, Process is called once for each such variable.

*Bounded (Run-Time) Errors*

It is a bounded error to call `Value` if more than one environment variable exists with the given name; the possible outcomes are that:

- one of the values is returned, and that same value is returned in subsequent calls in the absence of changes to the environment; or
- `Program_Error` is propagated.

*Erroneous Execution*

Making calls to the procedures `Set` or `Clear` concurrently with calls to any subprogram of package `Environment_Variables`, or to any instantiation of `Iterate`, results in erroneous execution.

Making calls to the procedures `Set` or `Clear` in the actual subprogram corresponding to the `Process` parameter of `Iterate` results in erroneous execution.

*Documentation Requirements*

An implementation shall document how the operations of this package behave if environment variables are changed by external mechanisms (for instance, calling operating system services).

*Implementation Permissions*

An implementation running on a system that does not support environment variables is permitted to define the operations of package `Environment_Variables` with the semantics corresponding to the case where the external execution environment does support environment variables. In this case, it shall provide a mechanism to initialize a nonempty set of environment variables prior to the execution of a partition.

*Implementation Advice*

If the execution environment supports subprocesses, the currently defined environment variables should be used to initialize the environment variables of a subprocess.

Changes to the environment variables made outside the control of this package should be reflected immediately in the effect of the operations of this package. Changes to the environment variables made using this package should be reflected immediately in the external execution environment. This package should not perform any buffering of the environment variables.

### **A.17.1 The Packages `Wide_Environment_Variables` and `Wide_Wide_Environment_Variables`**

The packages `Wide_Environment_Variables` and `Wide_Wide_Environment_Variables` allow a program to read or modify environment variables.

*Static Semantics*

The specification of package `Wide_Environment_Variables` is the same as for `Environment_Variables`, except that each occurrence of `String` is replaced by `Wide_String`.

The specification of package `Wide_Wide_Environment_Variables` is the same as for `Environment_Variables`, except that each occurrence of `String` is replaced by `Wide_Wide_String`.

## **A.18 Containers**

This clause presents the specifications of the package `Containers` and several child packages, which provide facilities for storing collections of elements.

A variety of sequence and associative containers are provided. Each container package defines a *cursor* type as well as a container type. A cursor is a reference to an element within a container. Many operations on cursors are common to all of the containers. A cursor referencing an element in a container is considered to be overlapping only with the element itself.

Some operations of the language-defined child units of Ada.Containers have access-to-subprogram parameters. To ensure such operations are well-defined, they guard against certain actions by the designated subprogram. An action on a container that can add or remove an element is considered to *tamper with cursors*, and these are prohibited during all such operations. An action on a container that can replace an element with one of a different size is considered to *tamper with elements*, and these are prohibited during certain of such operations. The details of the specific actions that are considered to tamper with cursors or elements are defined for each child unit of Ada.Containers.

Several of the language-defined child units of Ada.Containers include a nested package named Stable, which provides a view of a container that prohibits any operations that would tamper with elements. By using a Stable view for manipulating a container, the number of tampering checks performed while performing the operations can be reduced. The details of the Stable subpackage are defined separately for each child unit of Ada.Containers that includes such a nested package.

Within this clause we provide Implementation Advice for the desired average or worst case time complexity of certain operations on a container. This advice is expressed using the Landau symbol  $O(X)$ . Presuming  $f$  is some function of a length parameter  $N$  and  $t(N)$  is the time the operation takes (on average or worst case, as specified) for the length  $N$ , a complexity of  $O(f(N))$  means that there exists a finite  $A$  such that for any  $N$ ,  $t(N)/f(N) < A$ .

If the advice suggests that the complexity should be less than  $O(f(N))$ , then for any arbitrarily small positive real  $D$ , there should exist a positive integer  $M$  such that for all  $N > M$ ,  $t(N)/f(N) < D$ .

When a formal function is used to provide an ordering for a container, it is generally required to define a strict weak ordering. A function " $<$ " defines a *strict weak ordering* if it is irreflexive, asymmetric, transitive, and in addition, if  $x < y$  for any values  $x$  and  $y$ , then for all other values  $z$ ,  $(x < z)$  or  $(z < y)$ . Elements are in a *smallest first* order using such an operator if, for every element  $y$  with a predecessor  $x$  in the order,  $(y < x)$  is false.

#### Static Semantics

Certain subprograms declared within instances of some of the generic packages presented in this clause are said to *perform indefinite insertion*. These subprograms are those corresponding (in the sense of the copying described in 15.3) to subprograms that have formal parameters of a generic formal indefinite type and that are identified as performing indefinite insertion in the subclause defining the generic package.

If a subprogram performs indefinite insertion, then certain run-time checks are performed as part of a call to the subprogram; if any of these checks fail, then the resulting exception is propagated to the caller and the container is not modified by the call. These checks are performed for each parameter corresponding (in the sense of the copying described in 15.3) to a parameter in the corresponding generic whose type is a generic formal indefinite type. The checks performed for a given parameter are those checks explicitly specified in 7.8 that would be performed as part of the evaluation of an initialized allocator whose access type is declared immediately within the instance, where:

- the value of the `qualified_expression` is that of the parameter; and
- the designated subtype of the access type is the subtype of the parameter; and
- finalization of the collection of the access type has started if and only if the finalization of the instance has started.

#### Implementation Requirements

For an indefinite container (one whose type is defined in an instance of a child package of Containers whose `defining_identifier` contains "Indefinite"), each element of the container shall be created when it is inserted into the container and finalized when it is deleted from the container (or when the container object is finalized if the element has not been deleted). For a bounded container (one whose type is defined in an instance of a child package of Containers whose `defining_identifier` starts with "Bounded") that is not an indefinite container, all of the elements of the capacity of the container shall be created and default initialized when the container object is created; the elements shall be

finalized when the container object is finalized. For other kinds of containers, when elements are created and finalized is unspecified.

For an instance *I* of a container package with a container type, the specific type *T* of the object returned from a function that returns an object of an iterator interface, as well as the primitive operations of *T*, shall be nonblocking. The Global aspect specified for *T* and the primitive operations of *T* shall be (**in all, out synchronized**) or a specification that allows access to fewer global objects.

### A.18.1 The Package Containers

The package Containers is the root of the containers subsystem.

#### Static Semantics

The library package Containers has the following declaration:

```
package Ada.Containers
  with Pure is
    type Hash_Type is mod implementation-defined;
    type Count_Type is range 0 .. implementation-defined;
    Capacity_Error : exception;
  end Ada.Containers;
```

Hash\_Type represents the range of the result of a hash function. Count\_Type represents the (potential or actual) number of elements of a container.

Capacity\_Error is raised when the capacity of a container is exceeded.

#### Implementation Advice

Hash\_Type'Modulus should be at least  $2^{32}$ . Count\_Type'Last should be at least  $2^{31}-1$ .

### A.18.2 The Generic Package Containers.Vectors

The language-defined generic package Containers.Vectors provides private types Vector and Cursor, and a set of operations for each type. A vector container allows insertion and deletion at any position, but it is specifically optimized for insertion and deletion at the high end (the end with the higher index) of the container. A vector container also provides random access to its elements.

A vector container behaves conceptually as an array that expands as necessary as items are inserted. The *length* of a vector is the number of elements that the vector contains. The *capacity* of a vector is the maximum number of elements that can be inserted into the vector prior to it being automatically expanded.

Elements in a vector container can be referred to by an index value of a generic formal type. The first element of a vector always has its index value equal to the lower bound of the formal type.

A vector container may contain *empty elements*. Empty elements do not have a specified value.

#### Static Semantics

The generic library package Containers.Vectors has the following declaration:

```
with Ada.Iterator_Interfaces;
generic
  type Index_Type is range <>;
  type Element_Type is private;
  with function "=" (Left, Right : Element_Type)
    return Boolean is <>;
package Ada.Containers.Vectors
  with Preelaborate, Remote_Types,
    Nonblocking, Global => in out synchronized is
```



```

subtype Extended_Index is
  Index_Type'Base range
    Index_Type'First-1 ..
    Index_Type'Min (Index_Type'Base'Last - 1, Index_Type'Last) + 1;
No_Index : constant Extended_Index := Extended_Index'First;

type Vector is tagged private
  with Constant_Indexing => Constant_Reference,
    Variable_Indexing => Reference,
    Default_Iterator => Iterate,
    Iterator_Element => Element_Type,
    Iterator_View => Stable.Vector,
    Aggregate => (Empty => Empty,
      Add_Unnamed => Append,
      New_Indexed => New_Vector,
      Assign_Indexed => Replace_Element),
    Stable_Properties => (Length, Capacity,
      Tampering_With_Cursors_Prohibited,
      Tampering_With_Elements_Prohibited),
    Default_Initial_Condition =>
      Length (Vector) = 0 and then
        (not Tampering_With_Cursors_Prohibited (Vector)) and then
        (not Tampering_With_Elements_Prohibited (Vector)),
    Preelaborable_Initialization;

type Cursor is private
  with Preelaborable_Initialization;

Empty_Vector : constant Vector;
No_Element : constant Cursor;

function Has_Element (Position : Cursor) return Boolean
  with Nonblocking, Global => in all, Use_Formal => null;

function Has_Element (Container : Vector; Position : Cursor)
  return Boolean
  with Nonblocking, Global => null, Use_Formal => null;

package Vector_Iterator_Interfaces is new
  Ada.Iterator_Interfaces (Cursor, Has_Element);

function "=" (Left, Right : Vector) return Boolean;

function Tampering_With_Cursors_Prohibited
  (Container : Vector) return Boolean
  with Nonblocking, Global => null, Use_Formal => null;

function Tampering_With_Elements_Prohibited
  (Container : Vector) return Boolean
  with Nonblocking, Global => null, Use_Formal => null;

function Maximum_Length return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;

function Empty (Capacity : Count_Type := implementation-defined)
  return Vector
  with Pre => Capacity <= Maximum_Length
    or else raise Constraint_Error,
    Post =>
      Capacity (Empty'Result) >= Capacity and then
        not Tampering_With_Elements_Prohibited (Empty'Result) and then
        not Tampering_With_Cursors_Prohibited (Empty'Result) and then
        Length (Empty'Result) = 0;

```

```

function To_Vector (Length : Count_Type) return Vector
  with Pre => Length <= Maximum_Length or else raise Constraint_Error,
  Post =>
    To_Vector'Result.Length = Length and then
    not Tampering_With_Elements_Prohibited (To_Vector'Result)
    and then
    not Tampering_With_Cursors_Prohibited (To_Vector'Result)
    and then
    To_Vector'Result.Capacity >= Length;

function To_Vector
  (New_Item : Element_Type;
   Length : Count_Type) return Vector
  with Pre => Length <= Maximum_Length or else raise Constraint_Error,
  Post =>
    To_Vector'Result.Length = Length and then
    not Tampering_With_Elements_Prohibited (To_Vector'Result)
    and then
    not Tampering_With_Cursors_Prohibited (To_Vector'Result)
    and then
    To_Vector'Result.Capacity >= Length;

function New_Vector (First, Last : Index_Type) return Vector is
  (To_Vector (Count_Type (Last - First + 1)))
  with Pre => First = Index_Type'First;

function "&" (Left, Right : Vector) return Vector
  with Pre => Length (Left) <= Maximum_Length - Length (Right)
    or else raise Constraint_Error,
  Post => Length (Vectors."&"'Result) =
    Length (Left) + Length (Right) and then
    not Tampering_With_Elements_Prohibited
      (Vectors."&"'Result) and then
    not Tampering_With_Cursors_Prohibited
      (Vectors."&"'Result) and then
    Vectors."&"'Result.Capacity >=
      Length (Left) + Length (Right);

function "&" (Left : Vector;
  Right : Element_Type) return Vector
  with Pre => Length (Left) <= Maximum_Length - 1
    or else raise Constraint_Error,
  Post => Vectors."&"'Result.Length = Length (Left) + 1 and then
    not Tampering_With_Elements_Prohibited
      (Vectors."&"'Result) and then
    not Tampering_With_Cursors_Prohibited
      (Vectors."&"'Result) and then
    Vectors."&"'Result.Capacity >= Length (Left) + 1;

function "&" (Left : Element_Type;
  Right : Vector) return Vector
  with Pre => Length (Right) <= Maximum_Length - 1
    or else raise Constraint_Error,
  Post => Length (Vectors."&"'Result) = Length (Right) + 1 and then
    not Tampering_With_Elements_Prohibited
      (Vectors."&"'Result) and then
    not Tampering_With_Cursors_Prohibited
      (Vectors."&"'Result) and then
    Vectors."&"'Result.Capacity >= Length (Right) + 1;

function "&" (Left, Right : Element_Type) return Vector
  with Pre => Maximum_Length >= 2 or else raise Constraint_Error,
  Post => Length ("&"'Result) = 2 and then
    not Tampering_With_Elements_Prohibited
      (Vectors."&"'Result) and then
    not Tampering_With_Cursors_Prohibited
      (Vectors."&"'Result) and then
    Vectors."&"'Result.Capacity >= 2;

```

```

function Capacity (Container : Vector) return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;

procedure Reserve_Capacity (Container : in out Vector;
                           Capacity  : in      Count_Type)
  with Pre  => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error,
        Post => Container.Capacity >= Capacity;

function Length (Container : Vector) return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;

procedure Set_Length (Container : in out Vector;
                     Length      : in      Count_Type)
  with Pre  => (not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error) and then
        (Length <= Maximum_Length
        or else raise Constraint_Error),
        Post => Container.Length = Length and then
        Capacity (Container) >= Length;

function Is_Empty (Container : Vector) return Boolean
  with Nonblocking, Global => null, Use_Formal => null,
        Post => Is_Empty'Result = (Length (Container) = 0);

procedure Clear (Container : in out Vector)
  with Pre  => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error,
        Post => Length (Container) = 0;

function To_Cursor (Container : Vector;
                   Index       : Extended_Index) return Cursor
  with Post => (if Index in
        First_Index (Container) .. Last_Index (Container)
        then Has_Element (Container, To_Cursor'Result)
        else To_Cursor'Result = No_Element),
        Nonblocking, Global => null, Use_Formal => null;

function To_Index (Position : Cursor) return Extended_Index
  with Nonblocking, Global => in all;

function To_Index (Container : Vector;
                   Position   : Cursor) return Extended_Index
  with Pre  => Position = No_Element or else
        Has_Element (Container, Position) or else
        raise Program_Error,
        Post => (if Position = No_Element then To_Index'Result = No_Index
        else To_Index'Result in First_Index (Container) ..
        Last_Index (Container)),
        Nonblocking, Global => null, Use_Formal => null;

function Element (Container : Vector;
                  Index      : Index_Type)
  return Element_Type
  with Pre  => Index in
        First_Index (Container) .. Last_Index (Container)
        or else raise Constraint_Error,
        Nonblocking, Global => null, Use_Formal => Element_Type;

function Element (Position : Cursor) return Element_Type
  with Pre  => Position /= No_Element or else raise Constraint_Error,
        Nonblocking, Global => in all, Use_Formal => Element_Type;

function Element (Container : Vector;
                   Position   : Cursor) return Element_Type
  with Pre => (Position /= No_Element or else
        raise Constraint_Error) and then
        (Has_Element (Container, Position)
        or else raise Program_Error),
        Nonblocking, Global => null, Use_Formal => Element_Type;

```

```

procedure Replace_Element (Container : in out Vector;
                           Index      : in      Index_Type;
                           New_Item   : in      Element_Type)
with Pre => (not Tampering_With_Elements_Prohibited (Container)
             or else raise Program_Error) and then
  (Index in
   First_Index (Container) .. Last_Index (Container)
   or else raise Constraint_Error);

procedure Replace_Element (Container : in out Vector;
                           Position   : in      Cursor;
                           New_item   : in      Element_Type)
with Pre => (not Tampering_With_Elements_Prohibited (Container)
             or else raise Program_Error) and then
  (Position /= No_Element
   or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error);

procedure Query_Element
  (Container : in Vector;
   Index     : in Index_Type;
   Process   : not null access procedure (Element : in Element_Type))
with Pre => Index in
  First_Index (Container) .. Last_Index (Container)
  or else raise Constraint_Error;

procedure Query_Element
  (Position : in Cursor;
   Process   : not null access procedure (Element : in Element_Type))
with Pre => Position /= No_Element or else raise Constraint_Error,
  Global => in all;

procedure Query_Element
  (Container : in Vector;
   Position   : in Cursor;
   Process     : not null access procedure (Element : in Element_Type))
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error);

procedure Update_Element
  (Container : in out Vector;
   Index     : in      Index_Type;
   Process    : not null access procedure
               (Element : in out Element_Type))
with Pre => Index in
  First_Index (Container) .. Last_Index (Container)
  or else raise Constraint_Error;

procedure Update_Element
  (Container : in out Vector;
   Position   : in      Cursor;
   Process     : not null access procedure
               (Element : in out Element_Type))
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error);

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
with Implicit_Dereference => Element,
  Nonblocking, Global => in out synchronized,
  Default_Initial_Condition => (raise Program_Error);

```

```

type Reference_Type (Element : not null access Element_Type) is private
  with Implicit_Dereference => Element,
        Nonblocking, Global => in out synchronized,
        Default_Initial_Condition => (raise Program_Error);

function Constant_Reference (Container : aliased in Vector;
                             Index      : in Index_Type)
  return Constant_Reference_Type
  with Pre    => Index in
        First_Index (Container) .. Last_Index (Container)
        or else raise Constraint_Error,
        Post    => Tampering_With_Cursors_Prohibited (Container),
        Nonblocking, Global => null, Use_Formal => null;

function Reference (Container : aliased in out Vector;
                    Index      : in Index_Type)
  return Reference_Type
  with Pre    => Index in
        First_Index (Container) .. Last_Index (Container)
        or else raise Constraint_Error,
        Post    => Tampering_With_Cursors_Prohibited (Container),
        Nonblocking, Global => null, Use_Formal => null;

function Constant_Reference (Container : aliased in Vector;
                             Position  : in Cursor)
  return Constant_Reference_Type
  with Pre => (Position /= No_Element
        or else raise Constraint_Error) and then
        (Has_Element (Container, Position)
        or else raise Program_Error),
        Post    => Tampering_With_Cursors_Prohibited (Container),
        Nonblocking, Global => null, Use_Formal => null;

function Reference (Container : aliased in out Vector;
                    Position  : in Cursor)
  return Reference_Type
  with Pre => (Position /= No_Element
        or else raise Constraint_Error) and then
        (Has_Element (Container, Position)
        or else raise Program_Error),
        Post    => Tampering_With_Cursors_Prohibited (Container),
        Nonblocking, Global => null, Use_Formal => null;

procedure Assign (Target : in out Vector; Source : in Vector)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
        or else raise Program_Error,
        Post => Length (Source) = Length (Target) and then
        Capacity (Target) >= Length (Target);

function Copy (Source : Vector; Capacity : Count_Type := 0)
  return Vector
  with Pre => Capacity = 0 or else Capacity >= Length (Source)
        or else raise Capacity_Error,
        Post => Length (Copy'Result) = Length (Source) and then
        not Tampering_With_Elements_Prohibited (Copy'Result)
        and then
        not Tampering_With_Cursors_Prohibited (Copy'Result)
        and then
        Copy'Result.Capacity >= (if Capacity = 0 then
        Length (Source) else Capacity);

```

```

procedure Move (Target : in out Vector;
                Source : in out Vector)
    with Pre => (not Tampering_With_Cursors_Prohibited (Target)
                or else raise Program_Error) and then
                (not Tampering_With_Cursors_Prohibited (Source)
                or else raise Program_Error),
    Post => (if not Target'Has_Same_Storage (Source) then
            Length (Target) = Length (Source)'Old and then
            Length (Source) = 0 and then
            Capacity (Target) >= Length (Source)'Old);

procedure Insert_Vector (Container : in out Vector;
                        Before      : in      Extended_Index;
                        New_Item    : in      Vector)
    with Pre => (not Tampering_With_Cursors_Prohibited (Container)
                or else raise Program_Error) and then
                (Before in
                 First_Index (Container) .. Last_Index (Container) + 1
                 or else raise Constraint_Error) and then
                (Length (Container) <= Maximum_Length - Length (New_Item)
                 or else raise Constraint_Error)
    Post => Length (Container)'Old + Length (New_Item) =
            Length (Container) and then
            Capacity (Container) >= Length (Container);

procedure Insert_Vector (Container : in out Vector;
                        Before      : in      Cursor;
                        New_Item    : in      Vector)
    with Pre => (not Tampering_With_Cursors_Prohibited (Container)
                or else raise Program_Error) and then
                (Before = No_Element or else
                 Has_Element (Container, Before)
                 or else raise Program_Error) and then
                (Length (Container) <= Maximum_Length - Length (New_Item)
                 or else raise Constraint_Error),
    Post => Length (Container)'Old + Length (New_Item) =
            Length (Container) and then
            Capacity (Container) >= Length (Container);

procedure Insert_Vector (Container : in out Vector;
                        Before      : in      Cursor;
                        New_Item    : in      Vector;
                        Position    : out Cursor)
    with Pre => (not Tampering_With_Cursors_Prohibited (Container)
                or else raise Program_Error) and then
                (Before = No_Element or else
                 Has_Element (Container, Before)
                 or else raise Program_Error) and then
                (Length (Container) <= Maximum_Length - Length (New_Item)
                 or else raise Constraint_Error),
    Post => Length (Container)'Old + Length (New_Item) =
            Length (Container) and then
            Has_Element (Container, Position) and then
            Capacity (Container) >= Length (Container);

```



```

procedure Insert (Container : in out Vector;
  Before      : in      Extended_Index;
  New_Item    : in      Element_Type;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Before in
    First_Index (Container) .. Last_Index (Container) + 1
    or else raise Constraint_Error) and then
    (Length (Container) <= Maximum_Length - Count
    or else raise Constraint_Error),
  Post => Length (Container)'Old + Count =
    Length (Container) and then
    Capacity (Container) >= Length (Container);

procedure Insert (Container : in out Vector;
  Before      : in      Cursor;
  New_Item    : in      Element_Type;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Before = No_Element or else
    Has_Element (Container, Before)
    or else raise Program_Error) and then
    (Length (Container) <= Maximum_Length - Count
    or else raise Constraint_Error),
  Post => Length (Container)'Old + Count =
    Length (Container) and then
    Capacity (Container) >= Length (Container);

procedure Insert (Container : in out Vector;
  Before      : in      Cursor;
  New_Item    : in      Element_Type;
  Position    : out     Cursor;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Before = No_Element or else
    Has_Element (Container, Before)
    or else raise Program_Error) and then
    (Length (Container) <= Maximum_Length - Count
    or else raise Constraint_Error),
  Post => Length (Container)'Old + Count =
    Length (Container) and then
    Has_Element (Container, Position) and then
    Capacity (Container) >= Length (Container);

procedure Insert (Container : in out Vector;
  Before      : in      Extended_Index;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Before in
    First_Index (Container) .. Last_Index (Container) + 1
    or else raise Constraint_Error) and then
    (Length (Container) <= Maximum_Length - Count
    or else raise Constraint_Error),
  Post => Length (Container)'Old + Count =
    Length (Container) and then
    Capacity (Container) >= Length (Container);

```

```

procedure Insert (Container : in out Vector;
  Before      : in      Cursor;
  Position    :      out Cursor;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Before = No_Element or else
  Has_Element (Container, Before)
  or else raise Program_Error) and then
  (Length (Container) <= Maximum_Length - Count
  or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container)
and then Has_Element (Container, Position) and then
  Capacity (Container) >= Length (Container);

procedure Prepend_Vector (Container : in out Vector;
  New_Item    : in      Vector)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Length (Container) <= Maximum_Length - Length (New_Item)
  or else raise Constraint_Error),
  Post => Length (Container)'Old + Length (New_Item) =
  Length (Container) and then
  Capacity (Container) >= Length (Container);

procedure Prepend (Container : in out Vector;
  New_Item    : in      Element_Type;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Length (Container) <= Maximum_Length - Count
  or else raise Constraint_Error),
  Post => Length (Container)'Old + Count =
  Length (Container) and then
  Capacity (Container) >= Length (Container);

procedure Append_Vector (Container : in out Vector;
  New_Item    : in      Vector)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Length (Container) <= Maximum_Length - Length (New_Item)
  or else raise Constraint_Error),
  Post => Length (Container)'Old + Length (New_Item) =
  Length (Container) and then
  Capacity (Container) >= Length (Container);

procedure Append (Container : in out Vector;
  New_Item    : in      Element_Type;
  Count       : in      Count_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Length (Container) <= Maximum_Length - Count
  or else raise Constraint_Error),
  Post =>
  Length (Container)'Old + Count = Length (Container) and then
  Capacity (Container) >= Length (Container);

procedure Append (Container : in out Vector;
  New_Item    : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Length (Container) <= Maximum_Length - 1
  or else raise Constraint_Error),
  Post => Length (Container)'Old + 1 = Length (Container) and then
  Capacity (Container) >= Length (Container);

```

```

procedure Insert_Space (Container : in out Vector;
                        Before      : in      Extended_Index;
                        Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Before in
     First_Index (Container) .. Last_Index (Container) + 1
     or else raise Constraint_Error) and then
    (Length (Container) <= Maximum_Length - Count
     or else raise Constraint_Error),
    Post => Length (Container)'Old + Count =
        Length (Container) and then
        Capacity (Container) >= Length (Container);

procedure Insert_Space (Container : in out Vector;
                        Before      : in      Cursor;
                        Position    : out Cursor;
                        Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Before = No_Element or else
     Has_Element (Container, Before)
     or else raise Program_Error) and then
    (Length (Container) <= Maximum_Length - Count
     or else raise Constraint_Error),
    Post => Length (Container)'Old + Count =
        Length (Container) and then
        Has_Element (Container, Position) and then
        Capacity (Container) >= Length (Container);

procedure Delete (Container : in out Vector;
                  Index      : in      Extended_Index;
                  Count      : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Index in
     First_Index (Container) .. Last_Index (Container) + 1
     or else raise Constraint_Error),
    Post => Length (Container)'Old - Count <=
        Length (Container);

procedure Delete (Container : in out Vector;
                  Position    : in out Cursor;
                  Count      : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Position /= No_Element
     or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
     or else raise Program_Error),
    Post => Length (Container)'Old - Count <=
        Length (Container) and then
        Position = No_Element;

procedure Delete_First (Container : in out Vector;
                       Count      : in      Count_Type := 1)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error,
    Post => Length (Container)'Old - Count <= Length (Container);

procedure Delete_Last (Container : in out Vector;
                      Count      : in      Count_Type := 1)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error,
    Post => Length (Container)'Old - Count <= Length (Container);

```

```

procedure Reverse_Elements (Container : in out Vector)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error;

procedure Swap (Container : in out Vector;
  I, J : in Index_Type)
  with Pre => (not Tampering_With_Elements_Prohibited (Container)
    or else raise Program_Error) and then
    (I in First_Index (Container) .. Last_Index (Container)
    or else raise Constraint_Error) and then
    (J in First_Index (Container) .. Last_Index (Container)
    or else raise Constraint_Error);

procedure Swap (Container : in out Vector;
  I, J : in Cursor)
  with Pre => (not Tampering_With_Elements_Prohibited (Container)
    or else raise Program_Error) and then
    (I /= No_Element or else Constraint_Error) and then
    (J /= No_Element or else Constraint_Error) and then
    (Has_Element (Container, I)
    or else raise Program_Error) and then
    (Has_Element (Container, J)
    or else raise Program_Error);

function First_Index (Container : Vector) return Index_Type
  with Nonblocking, Global => null, Use_Formal => null,
    Post => First_Index'Result = Index_Type'First;

function First (Container : Vector) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
    then Has_Element (Container, First'Result)
    else First'Result = No_Element);

function First_Element (Container : Vector)
  return Element_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

function Last_Index (Container : Vector) return Extended_Index
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if Length (Container) = 0
    then Last_Index'Result = No_Index
    else Count_Type (Last_Index'Result - Index_Type'First) =
      Length (Container) - 1);

function Last (Container : Vector) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
    then Has_Element (Container, Last'Result)
    else Last'Result = No_Element);

function Last_Element (Container : Vector)
  return Element_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

function Next (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then Next'Result = No_Element);

```

```

function Next (Container : Vector; Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position = No_Element then Next'Result = No_Element
      elsif Has_Element (Container, Next'Result) then
        To_Index (Container, Next'Result) =
          To_Index (Container, Position) + 1
      elsif Next'Result = No_Element then
        Position = Last (Container)
      else False);

procedure Next (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;

procedure Next (Container : in      Vector;
  Position      : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position /= No_Element
      then Has_Element (Container, Position));

function Previous (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element
      then Previous'Result = No_Element);

function Previous (Container : Vector;
  Position      : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position = No_Element
      then Previous'Result = No_Element
      elsif Has_Element (Container, Previous'Result) then
        To_Index (Container, Previous'Result) =
          To_Index (Container, Position) - 1
      elsif Previous'Result = No_Element then
        Position = First (Container)
      else False);

procedure Previous (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;

procedure Previous (Container : in      Vector;
  Position      : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position /= No_Element
      then Has_Element (Container, Position));

function Find_Index (Container : Vector;
  Item      : Element_Type;
  Index     : Index_Type := Index_Type'First)
  return Extended_Index;

```

```

function Find (Container : Vector;
               Item      : Element_Type;
               Position   : Cursor := No_Element)
return Cursor
with Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Find'Result /= No_Element
            then Has_Element (Container, Find'Result));

function Reverse_Find_Index (Container : Vector;
                            Item       : Element_Type;
                            Index      : Index_Type := Index_Type'Last)

return Extended_Index;

function Reverse_Find (Container : Vector;
                      Item       : Element_Type;
                      Position   : Cursor := No_Element)

return Cursor
with Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Reverse_Find'Result /= No_Element
            then Has_Element (Container, Reverse_Find'Result));

function Contains (Container : Vector;
                  Item       : Element_Type) return Boolean;

procedure Iterate
(Container : in Vector;
 Process   : not null access procedure (Position : in Cursor))
with Allows_Exit;

procedure Reverse_Iterate
(Container : in Vector;
 Process   : not null access procedure (Position : in Cursor))
with Allows_Exit;

function Iterate (Container : in Vector)
return Vector_Iterator_Interfaces.Parallel_Reversible_Iterator'Class
with Post => Tampering_With_Cursors_Prohibited (Container);

function Iterate (Container : in Vector; Start : in Cursor)
return Vector_Iterator_Interfaces.Reversible_Iterator'Class
with Pre  => (Start /= No_Element
            or else raise Constraint_Error) and then
            (Has_Element (Container, Start)
            or else raise Program_Error),
    Post   => Tampering_With_Cursors_Prohibited (Container);

generic
with function "<" (Left, Right : Element_Type)
return Boolean is <>;
package Generic_Sorting
with Nonblocking, Global => null is

    function Is_Sorted (Container : Vector) return Boolean;

    procedure Sort (Container : in out Vector)
        with Pre => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error;

```



```

procedure Merge (Target : in out Vector;
                  Source  : in out Vector)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
             (not Tampering_With_Cursors_Prohibited (Source)
             or else raise Program_Error) and then
             (Length (Target) <= Maximum_Length - Length (Source)
             or else raise Constraint_Error) and then
             ((Length (Source) = 0 or else
              not Target'Has_Same_Storage (Source))
             or else raise Program_Error),
    Post => (declare
             Result_Length : constant Count_Type :=
               Length (Source)'Old + Length (Target)'Old;
begin
             (Length (Source) = 0 and then
              Length (Target) = Result_Length and then
              Capacity (Target) >= Result_Length);
end Generic_Sorting;

package Stable is
    type Vector (Base : not null access Vectors.Vector) is
        tagged limited private
        with Constant_Indexing => Constant_Reference,
            Variable_Indexing => Reference,
            Default_Iterator   => Iterate,
            Iterator_Element   => Element_Type,
            Stable_Properties  => (Length, Capacity),
            Global              => null,
            Default_Initial_Condition => Length (Vector) = 0,
            Preelaborable_Initialization;

    type Cursor is private
        with Preelaborable_Initialization;

    Empty_Vector : constant Vector;
    No_Element   : constant Cursor;

    function Has_Element (Position : Cursor) return Boolean
        with Nonblocking, Global => in all, Use_Formal => null;

    package Vector_Iterator_Interfaces is new
        Ada.Iterator_Interfaces (Cursor, Has_Element);

    procedure Assign (Target : in out Vectors.Vector;
                    Source  : in Vector)
        with Post => Length (Source) = Length (Target) and then
            Capacity (Target) >= Length (Target);

    function Copy (Source : Vectors.Vector) return Vector
        with Post => Length (Copy'Result) = Length (Source);

    type Constant_Reference_Type
        (Element : not null access constant Element_Type) is private
        with Implicit_Dereference => Element,
            Nonblocking, Global => null, Use_Formal => null,
            Default_Initial_Condition => (raise Program_Error);

    type Reference_Type
        (Element : not null access Element_Type) is private
        with Implicit_Dereference => Element,
            Nonblocking, Global => null, Use_Formal => null,
            Default_Initial_Condition => (raise Program_Error);

    -- Additional subprograms as described in the text
    -- are declared here.

private

```

```

    ... -- not specified by the language
end Stable;

private
    ... -- not specified by the language
end Ada.Containers.Vectors;

```

The actual function for the generic formal function "=" on `Element_Type` values is expected to define a reflexive and symmetric relationship and return the same result value each time it is called with a particular pair of values. If it behaves in some other manner, the functions defined to use it return an unspecified value. The exact arguments and number of calls of this generic formal function by the functions defined to use it are unspecified.

The type `Vector` is used to represent vectors. The type `Vector` needs finalization (see 10.6).

`Empty_Vector` represents the empty vector object. It has a length of 0. If an object of type `Vector` is not otherwise initialized, it is initialized to the same value as `Empty_Vector`.

`No_Element` represents a cursor that designates no element. If an object of type `Cursor` is not otherwise initialized, it is initialized to the same value as `No_Element`.

The primitive "=" operator for type `Cursor` returns `True` if both cursors are `No_Element`, or designate the same element in the same container.

Execution of the default implementation of the `Input`, `Output`, `Read`, or `Write` attribute of type `Cursor` raises `Program_Error`.

`Vector'Write` for a `Vector` object *V* writes `Length(V)` elements of the vector to the stream. It may also write additional information about the vector.

`Vector'Read` reads the representation of a vector from the stream, and assigns to *Item* a vector with the same length and elements as was written by `Vector'Write`.

`No_Index` represents a position that does not correspond to any element. The subtype `Extended_Index` includes the indices covered by `Index_Type` plus the value `No_Index` and, if it exists, the successor to the `Index_Type'Last`.

Some operations check for "tampering with cursors" of a container because they depend on the set of elements of the container remaining constant, and others check for "tampering with elements" of a container because they depend on elements of the container not being replaced. When tampering with cursors is *prohibited* for a particular vector object *V*, `Program_Error` is propagated by the finalization of *V*, as well as by a call that passes *V* to certain of the operations of this package, as indicated by the precondition of such an operation. Similarly, when tampering with elements is *prohibited* for *V*, `Program_Error` is propagated by a call that passes *V* to certain of the other operations of this package, as indicated by the precondition of such an operation.

```

function Has_Element (Position : Cursor) return Boolean
with Nonblocking, Global => in all, Use_Formal => null;

```

Returns `True` if `Position` designates an element, and returns `False` otherwise.

```

function Has_Element (Container : Vector; Position : Cursor)
return Boolean
with Nonblocking, Global => null, Use_Formal => null;

```

Returns `True` if `Position` designates an element in `Container`, and returns `False` otherwise.

```

function "=" (Left, Right : Vector) return Boolean;

```

If `Left` and `Right` denote the same vector object, then the function returns `True`. If `Left` and `Right` have different lengths, then the function returns `False`. Otherwise, it compares each element in `Left` to the corresponding element in `Right` using the generic formal equality operator. If any such comparison returns `False`, the function returns `False`; otherwise, it returns `True`. Any exception raised during evaluation of element equality is propagated.

```
function Tampering_With_Cursors_Prohibited
  (Container : Vector) return Boolean
with Nonblocking, Global => null, Use_Formal => null;
```

Returns True if tampering with cursors or tampering with elements is currently prohibited for Container, and returns False otherwise.

```
function Tampering_With_Elements_Prohibited
  (Container : Vector) return Boolean
with Nonblocking, Global => null, Use_Formal => null;
```

Always returns False, regardless of whether tampering with elements is prohibited.

```
function Maximum_Length return Count_Type
with Nonblocking, Global => null, Use_Formal => null;
```

Returns the maximum Length of a Vector, based on the index type.

```
function Empty (Capacity : Count_Type := implementation-defined)
return Vector
with Pre => Capacity <= Maximum_Length
      or else raise Constraint_Error,
      Post =>
        Capacity (Empty'Result) >= Capacity and then
        not Tampering_With_Elements_Prohibited (Empty'Result) and then
        not Tampering_With_Cursors_Prohibited (Empty'Result) and then
        Length (Empty'Result) = 0;
```

Returns an empty vector.

```
function To_Vector (Length : Count_Type) return Vector
with Pre => Length <= Maximum_Length or else raise Constraint_Error,
      Post =>
        To_Vector'Result.Length = Length and then
        not Tampering_With_Elements_Prohibited (To_Vector'Result)
          and then
        not Tampering_With_Cursors_Prohibited (To_Vector'Result)
          and then
        To_Vector'Result.Capacity >= Length;
```

Returns a vector with a length of Length, filled with empty elements.

```
function To_Vector
  (New_Item : Element_Type;
   Length   : Count_Type) return Vector
with Pre => Length <= Maximum_Length or else raise Constraint_Error,
      Post =>
        To_Vector'Result.Length = Length and then
        not Tampering_With_Elements_Prohibited (To_Vector'Result)
          and then
        not Tampering_With_Cursors_Prohibited (To_Vector'Result)
          and then
        To_Vector'Result.Capacity >= Length;
```

Returns a vector with a length of Length, filled with elements initialized to the value New\_Item.

```

function "&" (Left, Right : Vector) return Vector
  with Pre => Length (Left) <= Maximum_Length - Length (Right)
    or else raise Constraint_Error,
  Post => Length (Vectors."&"Result) =
    Length (Left) + Length (Right) and then
    not Tampering_With_Elements_Prohibited (Vectors."&"Result)
    and then
    not Tampering_With_Cursors_Prohibited (Vectors."&"Result)
    and then
    Vectors."&"Result.Capacity >=
      Length (Left) + Length (Right);

```

Returns a vector comprising the elements of Left followed by the elements of Right.

```

function "&" (Left : Vector;
  Right : Element_Type) return Vector
  with Pre => Length (Left) <= Maximum_Length - 1
    or else raise Constraint_Error,
  Post => Vectors."&"Result.Length = Length (Left) + 1 and then
    not Tampering_With_Elements_Prohibited (Vectors."&"Result)
    and then
    not Tampering_With_Cursors_Prohibited (Vectors."&"Result)
    and then
    Vectors."&"Result.Capacity >= Length (Left) + 1;

```

Returns a vector comprising the elements of Left followed by the element Right.

```

function "&" (Left : Element_Type;
  Right : Vector) return Vector
  with Pre => Length (Right) <= Maximum_Length - 1
    or else raise Constraint_Error,
  Post => Length (Vectors."&"Result) = Length (Right) + 1 and then
    not Tampering_With_Elements_Prohibited (Vectors."&"Result)
    and then
    not Tampering_With_Cursors_Prohibited (Vectors."&"Result)
    and then
    Vectors."&"Result.Capacity >= Length (Right) + 1;

```

Returns a vector comprising the element Left followed by the elements of Right.

```

function "&" (Left, Right : Element_Type) return Vector
  with Pre => Maximum_Length >= 2 or else raise Constraint_Error,
  Post => Length ("&"Result) = 2 and then
    not Tampering_With_Elements_Prohibited (Vectors."&"Result)
    and then
    not Tampering_With_Cursors_Prohibited (Vectors."&"Result)
    and then
    Vectors."&"Result.Capacity >= 2;

```

Returns a vector comprising the element Left followed by the element Right.

```

function Capacity (Container : Vector) return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;

```

Returns the capacity of Container.

```

procedure Reserve_Capacity (Container : in out Vector;
  Capacity : in Count_Type)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
  Post => Container.Capacity >= Capacity;

```

If the capacity of Container is already greater than or equal to Capacity, then Reserve\_Capacity has no effect. Otherwise, Reserve\_Capacity allocates additional storage as necessary to ensure that the length of the resulting vector can become at least the value Capacity without requiring an additional call to Reserve\_Capacity, and is large enough to hold the current length of Container. Reserve\_Capacity then, as necessary, moves elements into the new storage and

deallocates any storage no longer needed. Any exception raised during allocation is propagated and Container is not modified.

```
function Length (Container : Vector) return Count_Type
with Nonblocking, Global => null, Use_Formal => null;
```

Returns the number of elements in Container.

```
procedure Set_Length (Container : in out Vector;
                     Length    : in    Count_Type)
with Pre  => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
          (Length <= Maximum_Length or else raise Constraint_Error),
      Post => Container.Length = Length and then
          Capacity (Container) >= Length;
```

If Length is larger than the capacity of Container, Set\_Length calls Reserve\_Capacity (Container, Length), then sets the length of the Container to Length. If Length is greater than the original length of Container, empty elements are added to Container; otherwise, elements are removed from Container.

```
function Is_Empty (Container : Vector) return Boolean
with Nonblocking, Global => null, Use_Formal => null,
      Post => Is_Empty'Result = (Length (Container) = 0);
```

Returns True if Container is empty.

```
procedure Clear (Container : in out Vector)
with Pre  => not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error,
      Post => Length (Container) = 0;
```

Removes all the elements from Container. The capacity of Container does not change.

```
function To_Cursor (Container : Vector;
                   Index      : Extended_Index) return Cursor
with Post => (if Index in
             First_Index (Container) .. Last_Index (Container)
             then Has_Element (Container, To_Cursor'Result)
             else To_Cursor'Result = No_Element),
      Nonblocking, Global => null, Use_Formal => null;
```

Returns a cursor designating the element at position Index in Container; returns No\_Element if Index does not designate an element. For the purposes of determining whether the parameters overlap in a call to To\_Cursor, the Container parameter is not considered to overlap with any object (including itself).

```
function To_Index (Position : Cursor) return Extended_Index
with Nonblocking, Global => in all, Use_Formal => null;
```

If Position is No\_Element, No\_Index is returned. Otherwise, the index (within its containing vector) of the element designated by Position is returned.

```
function To_Index (Container : Vector;
                  Position   : Cursor) return Extended_Index
with Pre  => Position = No_Element or else
          Has_Element (Container, Position) or else
          raise Program_Error,
      Post => (if Position = No_Element then To_Index'Result = No_Index
             else To_Index'Result in First_Index (Container) ..
                 Last_Index (Container)),
      Nonblocking, Global => null, Use_Formal => null;
```

Returns the index (within Container) of the element designated by Position; returns No\_Index if Position does not designate an element. For the purposes of determining whether the

parameters overlap in a call to `To_Index`, the `Container` parameter is not considered to overlap with any object (including itself).

```
function Element (Container : Vector;
                  Index      : Index_Type)
return Element_Type
with Pre => Index in First_Index (Container) .. Last_Index (Container)
      or else raise Constraint_Error,
      Nonblocking, Global => null, Use_Formal => Element_Type;
```

Element returns the element at position `Index`.

```
function Element (Position : Cursor) return Element_Type
with Pre => Position /= No_Element or else raise Constraint_Error,
      Nonblocking, Global => in all, Use_Formal => Element_Type;
```

Element returns the element designated by `Position`.

```
function Element (Container : Vector;
                  Position   : Cursor) return Element_Type
with Pre => (Position /= No_Element
      or else raise Constraint_Error) and then
      (Has_Element (Container, Position)
      or else raise Program_Error),
      Nonblocking, Global => null, Use_Formal => Element_Type;
```

Element returns the element designated by `Position` in `Container`.

```
procedure Replace_Element (Container : in out Vector;
                           Index      : in      Index_Type;
                           New_Item   : in      Element_Type)
with Pre => (not Tampering_With_Elements_Prohibited (Container)
      or else raise Program_Error) and then
      (Index in First_Index (Container) .. Last_Index (Container)
      or else raise Constraint_Error);
```

`Replace_Element` assigns the value `New_Item` to the element at position `Index`. Any exception raised during the assignment is propagated. The element at position `Index` is not an empty element after successful call to `Replace_Element`. For the purposes of determining whether the parameters overlap in a call to `Replace_Element`, the `Container` parameter is not considered to overlap with any object (including itself), and the `Index` parameter is considered to overlap with the element at position `Index`.

```
procedure Replace_Element (Container : in out Vector;
                           Position   : in      Cursor;
                           New_Item   : in      Element_Type)
with Pre => (not Tampering_With_Elements_Prohibited (Container)
      or else raise Program_Error) and then
      (Position /= No_Element
      or else raise Constraint_Error) and then
      (Has_Element (Container, Position)
      or else raise Program_Error);
```

`Replace_Element` assigns `New_Item` to the element designated by `Position`. Any exception raised during the assignment is propagated. The element at `Position` is not an empty element after successful call to `Replace_Element`. For the purposes of determining whether the parameters overlap in a call to `Replace_Element`, the `Container` parameter is not considered to overlap with any object (including itself).



```

procedure Query_Element
(Container : in Vector;
Index      : in Index_Type;
Process    : not null access procedure (Element : in Element_Type))
with Pre => Index in First_Index (Container) .. Last_Index (Container)
           or else raise Constraint_Error;

```

Query\_Element calls Process.all with the element at position Index as the argument. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Query_Element
(Position : in Cursor;
Process   : not null access procedure (Element : in Element_Type))
with Pre => Position /= No_Element or else raise Constraint_Error
        Global => in all;

```

Query\_Element calls Process.all with the element designated by Position as the argument. Tampering with the elements of the vector that contains the element designated by Position is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Query_Element
(Container : in Vector;
Position   : in Cursor;
Process    : not null access procedure (Element : in Element_Type))
with Pre => (Position /= No_Element
           or else raise Constraint_Error) and then
           (Has_Element (Container, Position)
           or else raise Program_Error);

```

Query\_Element calls Process.all with the element designated by Position as the argument. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Update_Element
(Container : in out Vector;
Index      : in Index_Type;
Process    : not null access procedure
              (Element : in out Element_Type))
with Pre => Index in First_Index (Container) .. Last_Index (Container)
           or else raise Constraint_Error;

```

Update\_Element calls Process.all with the element at position Index as the argument. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

If Element\_Type is unconstrained and definite, then the actual Element parameter of Process.all shall be unconstrained.

The element at position Index is not an empty element after successful completion of this operation.

```

procedure Update_Element
(Container : in out Vector;
Position   : in      Cursor;
Process    : not null access procedure
              (Element : in out Element_Type))
with Pre => (Position /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Position)
              or else raise Program_Error);

```

Update\_Element calls Process.all with the element designated by Position as the argument. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

If Element\_Type is unconstrained and definite, then the actual Element parameter of Process.all shall be unconstrained.

The element designated by Position is not an empty element after successful completion of this operation.

```

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

type Reference_Type (Element : not null access Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

```

The types Constant\_Reference\_Type and Reference\_Type need finalization.

```

function Constant_Reference (Container : aliased in Vector;
                             Index      : in Index_Type)

return Constant_Reference_Type
with Pre   => Index in First_Index (Container) .. Last_Index (Container)
        or else raise Constraint_Error,
      Post  => Tampering_With_Cursors_Prohibited (Container),
      Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Constant\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read access to an individual element of a vector given an index value.

Constant\_Reference returns an object whose discriminant is an access value that designates the element at position Index. Tampering with the elements of Container is prohibited while the object returned by Constant\_Reference exists and has not been finalized.

```

function Reference (Container : aliased in out Vector;
                    Index      : in Index_Type)

return Reference_Type
with Pre   => Index in First_Index (Container) .. Last_Index (Container)
        or else raise Constraint_Error,
      Post  => Tampering_With_Cursors_Prohibited (Container),
      Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Variable\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read and write access to an individual element of a vector given an index value.

Reference returns an object whose discriminant is an access value that designates the element at position Index. Tampering with the elements of Container is prohibited while the object returned by Reference exists and has not been finalized.

The element at position Index is not an empty element after successful completion of this operation.

```
function Constant_Reference (Container : aliased in Vector;
                             Position  : in Cursor)
return Constant_Reference_Type
with Pre  => (Position /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Position)
              or else raise Program_Error),
    Post  => Tampering_With_Cursors_Prohibited (Container),
    Nonblocking, Global => null, Use_Formal => null;
```

This function (combined with the Constant\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read access to an individual element of a vector given a cursor.

Constant\_Reference returns an object whose discriminant is an access value that designates the element designated by Position. Tampering with the elements of Container is prohibited while the object returned by Constant\_Reference exists and has not been finalized.

```
function Reference (Container : aliased in out Vector;
                    Position  : in Cursor)
return Reference_Type
with Pre  => (Position /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Position)
              or else raise Program_Error),
    Post  => Tampering_With_Cursors_Prohibited (Container),
    Nonblocking, Global => null, Use_Formal => null;
```

This function (combined with the Variable\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read and write access to an individual element of a vector given a cursor.

Reference returns an object whose discriminant is an access value that designates the element designated by Position. Tampering with the elements of Container is prohibited while the object returned by Reference exists and has not been finalized.

The element designated by Position is not an empty element after successful completion of this operation.

```
procedure Assign (Target : in out Vector; Source : in Vector)
with Pre  => not Tampering_With_Cursors_Prohibited (Target)
              or else raise Program_Error,
    Post  => Length (Source) = Length (Target) and then
              Capacity (Target) >= Length (Target);
```

If Target denotes the same object as Source, the operation has no effect. If the length of Source is greater than the capacity of Target, Reserve\_Capacity (Target, Length (Source)) is called. The elements of Source are then copied to Target as for an assignment\_statement assigning Source to Target (this includes setting the length of Target to be that of Source).

```
function Copy (Source : Vector; Capacity : Count_Type := 0)
return Vector
with Pre => Capacity = 0 or else Capacity >= Length (Source)
              or else raise Capacity_Error,
    Post => Length (Copy'Result) = Length (Source) and then
              not Tampering_With_Elements_Prohibited (Copy'Result)
              and then
              not Tampering_With_Cursors_Prohibited (Copy'Result)
              and then
              Copy'Result.Capacity >= (if Capacity = 0 then
              Length (Source) else Capacity);
```

Returns a vector whose elements are initialized from the corresponding elements of Source.

```

procedure Move (Target : in out Vector;
                Source : in out Vector)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
              or else raise Program_Error) and then
              (not Tampering_With_Cursors_Prohibited (Source)
              or else raise Program_Error),
  Post => (if not Target'Has_Same_Storage (Source) then
          Length (Target) = Length (Source)'Old and then
          Length (Source) = 0 and then
          Capacity (Target) >= Length (Source)'Old);

```

If Target denotes the same object as Source, then the operation has no effect. Otherwise, Move first calls Reserve\_Capacity (Target, Length (Source)) and then Clear (Target); then, each element from Source is removed from Source and inserted into Target in the original order.

```

procedure Insert_Vector (Container : in out Vector;
                        Before      : in      Extended_Index;
                        New_Item    : in      Vector)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
              or else raise Program_Error) and then
              (Before in
               First_Index (Container) .. Last_Index (Container) + 1
               or else raise Constraint_Error) and then
              (Length (Container) <= Maximum_Length - Length (New_Item)
               or else raise Constraint_Error),
  Post => Length (Container)'Old + Length (New_Item) =
          Length (Container) and then
          Capacity (Container) >= Length (Container);

```

If Length(New\_Item) is 0, then Insert\_Vector does nothing. Otherwise, it computes the new length  $NL$  as the sum of the current length and Length (New\_Item); if the value of Last appropriate for length  $NL$  would be greater than Index\_Type'Last, then Constraint\_Error is propagated.

If the current vector capacity is less than  $NL$ , Reserve\_Capacity (Container,  $NL$ ) is called to increase the vector capacity. Then Insert\_Vector slides the elements in the range Before .. Last\_Index (Container) up by Length(New\_Item) positions, and then copies the elements of New\_Item to the positions starting at Before. Any exception raised during the copying is propagated.

```

procedure Insert_Vector (Container : in out Vector;
                        Before      : in      Cursor;
                        New_Item    : in      Vector)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
              or else raise Program_Error) and then
              (Before = No_Element or else
               Has_Element (Container, Before)
               or else raise Program_Error) and then
              (Length (Container) <= Maximum_Length - Length (New_Item)
               or else raise Constraint_Error),
  Post => Length (Container)'Old + Length (New_Item) =
          Length (Container) and then
          Capacity (Container) >= Length (Container);

```

If Length(New\_Item) is 0, then Insert\_Vector does nothing. If Before is No\_Element, then the call is equivalent to Insert\_Vector (Container, Last\_Index (Container) + 1, New\_Item); otherwise, the call is equivalent to Insert\_Vector (Container, To\_Index (Before), New\_Item);

```

procedure Insert_Vector (Container : in out Vector;
                        Before      : in      Cursor;
                        New_Item    : in      Vector;
                        Position    : out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Before = No_Element or else
     Has_Element (Container, Before)
     or else raise Program_Error) and then
    (Length (Container) <= Maximum_Length - Length (New_Item)
     or else raise Constraint_Error),
    Post => Length (Container)'Old + Length (New_Item) =
           Length (Container) and then
    Has_Element (Container, Position) and then
    Capacity (Container) >= Length (Container);

```

If Before equals No\_Element, then let  $T$  be Last\_Index (Container) + 1; otherwise, let  $T$  be To\_Index (Before). Insert\_Vector (Container,  $T$ , New\_Item) is called, and then Position is set to To\_Cursor (Container,  $T$ ).

```

procedure Insert (Container : in out Vector;
                Before      : in      Extended_Index;
                New_Item    : in      Element_Type;
                Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Before in
     First_Index (Container) .. Last_Index (Container) + 1
     or else raise Constraint_Error) and then
    (Length (Container) <= Maximum_Length - Count
     or else raise Constraint_Error),
    Post => Length (Container)'Old + Count = Length (Container) and then
    Capacity (Container) >= Length (Container);

```

Equivalent to Insert (Container, Before, To\_Vector (New\_Item, Count));

```

procedure Insert (Container : in out Vector;
                Before      : in      Cursor;
                New_Item    : in      Element_Type;
                Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Before = No_Element or else
     Has_Element (Container, Before)
     or else raise Program_Error) and then
    (Length (Container) <= Maximum_Length - Count
     or else raise Constraint_Error),
    Post => Length (Container)'Old + Count = Length (Container) and then
    Capacity (Container) >= Length (Container);

```

Equivalent to Insert (Container, Before, To\_Vector (New\_Item, Count));

```

procedure Insert (Container : in out Vector;
  Before      : in      Cursor;
  New_Item    : in      Element_Type;
  Position    :      out Cursor;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Before = No_Element or else
  Has_Element (Container, Before)
  or else raise Program_Error) and then
  (Length (Container) <= Maximum_Length - Count
  or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container) and then
  Has_Element (Container, Position) and then
  Capacity (Container) >= Length (Container);

```

Equivalent to Insert (Container, Before, To\_Vector (New\_Item, Count), Position);

```

procedure Insert (Container : in out Vector;
  Before      : in      Extended_Index;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Before in
  First_Index (Container) .. Last_Index (Container) + 1
  or else raise Constraint_Error) and then
  (Length (Container) <= Maximum_Length - Count
  or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container) and then
  Capacity (Container) >= Length (Container);

```

If Count is 0, then Insert does nothing. Otherwise, it computes the new length *NL* as the sum of the current length and Count; if the value of Last appropriate for length *NL* would be greater than Index\_Type'Last, then Constraint\_Error is propagated.

If the current vector capacity is less than *NL*, Reserve\_Capacity (Container, *NL*) is called to increase the vector capacity. Then Insert slides the elements in the range Before .. Last\_Index (Container) up by Count positions, and then inserts elements that are initialized by default (see 6.3.1) in the positions starting at Before.

```

procedure Insert (Container : in out Vector;
  Before      : in      Cursor;
  Position    :      out Cursor;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Before = No_Element or else
  Has_Element (Container, Before)
  or else raise Program_Error) and then
  (Length (Container) <= Maximum_Length - Count
  or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container) and then
  Has_Element (Container, Position) and then
  Capacity (Container) >= Length (Container);

```

If Before equals No\_Element, then let *T* be Last\_Index (Container) + 1; otherwise, let *T* be To\_Index (Before). Insert (Container, *T*, Count) is called, and then Position is set to To\_Cursor (Container, *T*).



```

procedure Prepend_Vector (Container : in out Vector;
                          New_Item  : in      Vector)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (Length (Container) <= Maximum_Length - Length (New_Item)
             or else raise Constraint_Error),
  Post => Length (Container)'Old + Length (New_Item) =
          Length (Container) and then
          Capacity (Container) >= Length (Container);

```

Equivalent to Insert (Container, First\_Index (Container), New\_Item).

```

procedure Prepend (Container : in out Vector;
                  New_Item  : in      Element_Type;
                  Count     : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (Length (Container) <= Maximum_Length - Count
             or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container) and then
          Capacity (Container) >= Length (Container);

```

Equivalent to Insert (Container, First\_Index (Container), New\_Item, Count).

```

procedure Append_Vector (Container : in out Vector;
                        New_Item  : in      Vector)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (Length (Container) <= Maximum_Length - Length (New_Item)
             or else raise Constraint_Error),
  Post => Length (Container)'Old + Length (New_Item) =
          Length (Container) and then
          Capacity (Container) >= Length (Container);

```

Equivalent to Insert (Container, Last\_Index (Container) + 1, New\_Item).

```

procedure Append (Container : in out Vector;
                 New_Item  : in      Element_Type;
                 Count     : in      Count_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (Length (Container) <= Maximum_Length - Count
             or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container) and then
          Capacity (Container) >= Length (Container);

```

Equivalent to Insert (Container, Last\_Index (Container) + 1, New\_Item, Count).

```

procedure Append (Container : in out Vector;
                 New_Item  : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (Length (Container) <= Maximum_Length - 1
             or else raise Constraint_Error),
  Post => Length (Container)'Old + 1 = Length (Container) and then
          Capacity (Container) >= Length (Container);

```

Equivalent to Insert (Container, Last\_Index (Container) + 1, New\_Item, 1).

```

procedure Insert_Space (Container : in out Vector;
                        Before      : in      Extended_Index;
                        Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Before in
     First_Index (Container) .. Last_Index (Container) + 1
     or else raise Constraint_Error) and then
    (Length (Container) <= Maximum_Length - Count
     or else raise Constraint_Error),
    Post => Length (Container)'Old + Count = Length (Container) and then
    Capacity (Container) >= Length (Container);

```

If Count is 0, then Insert\_Space does nothing. Otherwise, it computes the new length *NL* as the sum of the current length and Count; if the value of Last appropriate for length *NL* would be greater than Index\_Type'Last, then Constraint\_Error is propagated.

If the current vector capacity is less than *NL*, Reserve\_Capacity (Container, *NL*) is called to increase the vector capacity. Then Insert\_Space slides the elements in the range Before .. Last\_Index (Container) up by Count positions, and then inserts empty elements in the positions starting at Before.

```

procedure Insert_Space (Container : in out Vector;
                        Before      : in      Cursor;
                        Position    : out Cursor;
                        Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Before = No_Element or else
     Has_Element (Container, Before)
     or else raise Program_Error) and then
    (Length (Container) <= Maximum_Length - Count
     or else raise Constraint_Error),
    Post => Length (Container)'Old + Count = Length (Container) and then
    Has_Element (Container, Position) and then
    Capacity (Container) >= Length (Container);

```

If Before equals No\_Element, then let *T* be Last\_Index (Container) + 1; otherwise, let *T* be To\_Index (Before). Insert\_Space (Container, *T*, Count) is called, and then Position is set to To\_Cursor (Container, *T*).

```

procedure Delete (Container : in out Vector;
                  Index      : in      Extended_Index;
                  Count      : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Index in
     First_Index (Container) .. Last_Index (Container) + 1
     or else raise Constraint_Error),
    Post => Length (Container)'Old - Count <= Length (Container);

```

If Count is 0, Delete has no effect. Otherwise, Delete slides the elements (if any) starting at position Index + Count down to Index. Any exception raised during element assignment is propagated.

```

procedure Delete (Container : in out Vector;
                  Position  : in out Cursor;
                  Count     : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
            (Position /= No_Element
            or else raise Constraint_Error) and then
            (Has_Element (Container, Position)
            or else raise Program_Error),
    Post => Length (Container)'Old - Count <= Length (Container)
    and then Position = No_Element;

```

Delete (Container, To\_Index (Position), Count) is called, and then Position is set to No\_Element.

```

procedure Delete_First (Container : in out Vector;
                       Count     : in      Count_Type := 1)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error,
    Post => Length (Container)'Old - Count <= Length (Container);

```

Equivalent to Delete (Container, First\_Index (Container), Count).

```

procedure Delete_Last (Container : in out Vector;
                      Count     : in      Count_Type := 1)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error,
    Post => Length (Container)'Old - Count <= Length (Container);

```

If Length (Container) <= Count, then Delete\_Last is equivalent to Clear (Container). Otherwise, it is equivalent to Delete (Container, Index\_Type'Val(Index\_Type'Pos(Last\_Index (Container)) - Count + 1), Count).

```

procedure Reverse_Elements (Container : in out Vector)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error;

```

Reorders the elements of Container in reverse order.

```

procedure Swap (Container : in out Vector;
               I, J       : in      Index_Type)
with Pre => (not Tampering_With_Elements_Prohibited (Container)
            or else raise Program_Error) and then
            (I in First_Index (Container) .. Last_Index (Container)
            or else raise Constraint_Error) and then
            (J in First_Index (Container) .. Last_Index (Container)
            or else raise Constraint_Error);

```

Swap exchanges the values of the elements at positions I and J.

```

procedure Swap (Container : in out Vector;
               I, J       : in      Cursor)
with Pre => (not Tampering_With_Elements_Prohibited (Container)
            or else raise Program_Error) and then
            (I /= No_Element or else Constraint_Error) and then
            (J /= No_Element or else Constraint_Error) and then
            (Has_Element (Container, I)
            or else raise Program_Error) and then
            (Has_Element (Container, J)
            or else raise Program_Error);

```

Swap exchanges the values of the elements designated by I and J.

```

function First_Index (Container : Vector) return Index_Type
with Nonblocking, Global => null, Use_Formal => null,
    Post => First_Index'Result = Index_Type'First;

```

Returns the value Index\_Type'First.

```

function First (Container : Vector) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
      then Has_Element (Container, First'Result)
      else First'Result = No_Element);

```

If Container is empty, First returns No\_Element. Otherwise, it returns a cursor that designates the first element in Container.

```

function First_Element (Container : Vector)
  return Element_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

```

Equivalent to Element (Container, First\_Index (Container)).

```

function Last_Index (Container : Vector) return Extended_Index
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if Length (Container) = 0 then Last_Index'Result = No_Index
      else Count_Type (Last_Index'Result - Index_Type'First) =
        Length (Container) - 1);

```

If Container is empty, Last\_Index returns No\_Index. Otherwise, it returns the position of the last element in Container.

```

function Last (Container : Vector) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
      then Has_Element (Container, Last'Result)
      else Last'Result = No_Element);

```

If Container is empty, Last returns No\_Element. Otherwise, it returns a cursor that designates the last element in Container.

```

function Last_Element (Container : Vector)
  return Element_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

```

Equivalent to Element (Container, Last\_Index (Container)).

```

function Next (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then Next'Result = No_Element);

```

If Position equals No\_Element or designates the last element of the container, then Next returns the value No\_Element. Otherwise, it returns a cursor that designates the element with index To\_Index (Position) + 1 in the same vector as Position.

```

function Next (Container : Vector;
  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position = No_Element then Next'Result = No_Element
      elsif Has_Element (Container, Next'Result) then
        To_Index (Container, Next'Result) =
          To_Index (Container, Position) + 1
      elsif Next'Result = No_Element then
        Position = Last (Container)
      else False);

```

Returns a cursor designating the next element in Container, if any.

```
procedure Next (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;
```

Equivalent to Position := Next (Position).

```
procedure Next (Container : in      Vector;
                Position  : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
  Pre  => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
  Post => (if Position /= No_Element
        then Has_Element (Container, Position));
```

Equivalent to Position := Next (Container, Position).

```
function Previous (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
  Post => (if Position = No_Element
        then Previous'Result = No_Element);
```

If Position equals No\_Element or designates the first element of the container, then Previous returns the value No\_Element. Otherwise, it returns a cursor that designates the element with index To\_Index (Position) - 1 in the same vector as Position.

```
function Previous (Container : Vector;
                  Position  : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
  Pre  => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
  Post => (if Position = No_Element then Previous'Result = No_Element
        elsif Has_Element (Container, Previous'Result) then
          To_Index (Container, Previous'Result) =
          To_Index (Container, Position) - 1
        elsif Previous'Result = No_Element then
          Position = First (Container)
        else False);
```

Returns a cursor designating the previous element in Container, if any.

```
procedure Previous (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;
```

Equivalent to Position := Previous (Position).

```
procedure Previous (Container : in      Vector;
                  Position  : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
  Pre  => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
  Post => (if Position /= No_Element
        then Has_Element (Container, Position));
```

Equivalent to Position := Previous (Container, Position).

```
function Find_Index (Container : Vector;
                    Item       : Element_Type;
                    Index      : Index_Type := Index_Type'First)
  return Extended_Index;
```

Searches the elements of Container for an element equal to Item (using the generic formal equality operator). The search starts at position Index and proceeds towards Last\_Index (Container). If no equal element is found, then Find\_Index returns No\_Index. Otherwise, it returns the index of the first equal element encountered.

```

function Find (Container : Vector;
               Item      : Element_Type;
               Position   : Cursor := No_Element)
return Cursor
with Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Find'Result /= No_Element
            then Has_Element (Container, Find'Result));

```

Find searches the elements of Container for an element equal to Item (using the generic formal equality operator). The search starts at the first element if Position equals No\_Element, and at the element designated by Position otherwise. It proceeds towards the last element of Container. If no equal element is found, then Find returns No\_Element. Otherwise, it returns a cursor designating the first equal element encountered.

```

function Reverse_Find_Index (Container : Vector;
                             Item      : Element_Type;
                             Index     : Index_Type := Index_Type'Last)

return Extended_Index;

```

Searches the elements of Container for an element equal to Item (using the generic formal equality operator). The search starts at position Index or, if Index is greater than Last\_Index (Container), at position Last\_Index (Container). It proceeds towards First\_Index (Container). If no equal element is found, then Reverse\_Find\_Index returns No\_Index. Otherwise, it returns the index of the first equal element encountered.

```

function Reverse_Find (Container : Vector;
                       Item      : Element_Type;
                       Position   : Cursor := No_Element)

return Cursor
with Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Reverse_Find'Result /= No_Element
            then Has_Element (Container, Reverse_Find'Result));

```

Reverse\_Find searches the elements of Container for an element equal to Item (using the generic formal equality operator). The search starts at the last element if Position equals No\_Element, and at the element designated by Position otherwise. It proceeds towards the first element of Container. If no equal element is found, then Reverse\_Find returns No\_Element. Otherwise, it returns a cursor designating the first equal element encountered.

```

function Contains (Container : Vector;
                  Item      : Element_Type) return Boolean;

```

Equivalent to Has\_Element (Find (Container, Item)).

```

procedure Iterate
(Container : in Vector;
 Process  : not null access procedure (Position : in Cursor))
with Allows_Exit;

```

Invokes Process.all with a cursor that designates each element in Container, in index order. Tampering with the cursors of Container is prohibited during the execution of a call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Reverse_Iterate
(Container : in Vector;
 Process  : not null access procedure (Position : in Cursor))
with Allows_Exit;

```

Iterates over the elements in Container as per procedure Iterate, except that elements are traversed in reverse index order.



```

function Iterate (Container : in Vector)
  return Vector_Iterator_Interfaces.Parallel_Reversible_Iterator'Class
  with Post    => Tampering_With_Cursors_Prohibited (Container);

```

Iterate returns an iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each node in Container, starting with the first node and moving the cursor as per the Next function when used as a forward iterator, and starting with the last node and moving the cursor as per the Previous function when used as a reverse iterator, and processing all nodes concurrently when used as a parallel iterator. Tampering with the cursors of Container is prohibited while the iterator object exists (in particular, in the sequence\_of\_statements of the loop\_statement whose iterator\_specification denotes this object). The iterator object needs finalization.

```

function Iterate (Container : in Vector; Start : in Cursor)
  return Vector_Iterator_Interfaces.Reversible_Iterator'Class
  with Pre      => (Start /= No_Element
                   or else raise Constraint_Error) and then
                   (Has_Element (Container, Start)
                   or else raise Program_Error),
  Post      => Tampering_With_Cursors_Prohibited (Container);

```

Iterate returns a reversible iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each node in Container, starting with the node designated by Start and moving the cursor as per the Next function when used as a forward iterator, or moving the cursor as per the Previous function when used as a reverse iterator. Tampering with the cursors of Container is prohibited while the iterator object exists (in particular, in the sequence\_of\_statements of the loop\_statement whose iterator\_specification denotes this object). The iterator object needs finalization.

The actual function for the generic formal function "<" of Generic\_Sorting is expected to return the same value each time it is called with a particular pair of element values. It should define a strict weak ordering relationship (see A.18); it should not modify Container. If the actual for "<" behaves in some other manner, the behavior of the subprograms of Generic\_Sorting are unspecified. The number of times the subprograms of Generic\_Sorting call "<" is unspecified.

```

function Is_Sorted (Container : Vector) return Boolean;

```

Returns True if the elements are sorted smallest first as determined by the generic formal "<" operator; otherwise, Is\_Sorted returns False. Any exception raised during evaluation of "<" is propagated.

```

procedure Sort (Container : in out Vector)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error;

```

Reorders the elements of Container such that the elements are sorted smallest first as determined by the generic formal "<" operator provided. Any exception raised during evaluation of "<" is propagated.

```

procedure Merge (Target  : in out Vector;
                  Source  : in out Vector)
  with Pre => (not Tampering_With_Cursors_Prohibited (Target)
              or else raise Program_Error) and then
              (not Tampering_With_Cursors_Prohibited (Source)
              or else raise Program_Error) and then
              (Length (Target) <= Maximum_Length - Length (Source)
              or else raise Constraint_Error) and then
              ((Length (Source) = 0 or else
               not Target'Has_Same_Storage (Source))
              or else raise Program_Error),
  Post => (declare
          Result_Length : constant Count_Type :=
            Length (Source)'Old + Length (Target)'Old;
          begin
            (Length (Source) = 0 and then
             Length (Target) = Result_Length and then
             Capacity (Target) >= Result_Length);

```

Merge removes elements from Source and inserts them into Target; afterwards, Target contains the union of the elements that were initially in Source and Target, Source is left empty. If Target and Source are initially sorted smallest first, then Target is ordered smallest first as determined by the generic formal "<" operator; otherwise, the order of elements in Target is unspecified. Any exception raised during evaluation of "<" is propagated.

The nested package Vectors.Stable provides a type Stable.Vector that represents a *stable* vector, which is one that cannot grow and shrink. Such a vector can be created by calling the To\_Vector or Copy functions, or by establishing a *stabilized view* of an ordinary vector.

The subprograms of package Containers.Vectors that have a parameter or result of type Vector are included in the nested package Stable with the same specification, except that the following are omitted:

Tampering\_With\_Cursors\_Prohibited, Tampering\_With\_Elements\_Prohibited, Reserve\_Capacity, Assign, Move, Insert, Insert\_Space, Insert\_Vector, Append, Append\_Vector, Prepend, Prepend\_Vector, Clear, Delete, Delete\_First, Delete\_Last, and Set\_Length

The generic package Generic\_Sorting is also included with the same specification, except that Merge is omitted.

The operations of this package are equivalent to those for ordinary vectors, except that the calls to Tampering\_With\_Cursors\_Prohibited and Tampering\_With\_Elements\_Prohibited that occur in preconditions are replaced by False, and any that occur in postconditions are replaced by True.

If a stable vector is declared with the Base discriminant designating a pre-existing ordinary vector, the stable vector represents a stabilized view of the underlying ordinary vector, and any operation on the stable vector is reflected on the underlying ordinary vector. While a stabilized view exists, any operation that tampers with elements performed on the underlying vector is prohibited. The finalization of a stable vector that provides such a view removes this restriction on the underlying ordinary vector (though some other restriction can exist due to other concurrent iterations or stabilized views).

If a stable vector is declared without specifying Base, the object is necessarily initialized. The initializing expression of the stable vector, typically a call on To\_Vector or Copy, determines the Length of the vector. The Length of a stable vector never changes after initialization.

#### Bounded (Run-Time) Errors

Reading the value of an empty element by calling Element, Query\_Element, Update\_Element, Constant\_Reference, Reference, Swap, Is\_Sorted, Sort, Merge, "=", Find, or Reverse\_Find is a bounded error. The implementation may treat the element as having any normal value (see 16.9.1) of the element type, or raise Constraint\_Error or Program\_Error before modifying the vector.

Calling Merge in an instance of Generic\_Sorting with either Source or Target not ordered smallest first using the provided generic formal "<" operator is a bounded error. Either Program\_Error is raised after Target is updated as described for Merge, or the operation works as defined.

It is a bounded error for the actual function associated with a generic formal subprogram, when called as part of an operation of this package, to tamper with elements of any Vector parameter of the operation. Either Program\_Error is raised, or the operation works as defined on the value of the Vector either prior to, or subsequent to, some or all of the modifications to the Vector.

It is a bounded error to call any subprogram declared in the visible part of Containers.Vectors when the associated container has been finalized. If the operation takes Container as an **in out** parameter, then it raises Constraint\_Error or Program\_Error. Otherwise, the operation either proceeds as it would for an empty container, or it raises Constraint\_Error or Program\_Error.

A Cursor value is *ambiguous* if any of the following have occurred since it was created:

- Insert, Insert\_Space, Insert\_Vector, or Delete has been called on the vector that contains the element the cursor designates with an index value (or a cursor designating an element at such an index value) less than or equal to the index value of the element designated by the cursor; or
- The vector that contains the element it designates has been passed to the Sort or Merge procedures of an instance of Generic\_Sorting, or to the Reverse\_Elements procedure.

It is a bounded error to call any subprogram other than "=" or Has\_Element declared in Containers.Vectors with an ambiguous (but not invalid, see below) cursor parameter. Possible results are:

- The cursor may be treated as if it were No\_Element;
- The cursor may designate some element in the vector (but not necessarily the element that it originally designated);
- Constraint\_Error may be raised; or
- Program\_Error may be raised.

#### *Erroneous Execution*

A Cursor value is *invalid* if any of the following have occurred since it was created:

- The vector that contains the element it designates has been finalized;
- The vector that contains the element it designates has been used as the Target of a call to Assign, or as the target of an assignment\_statement;
- The vector that contains the element it designates has been used as the Source or Target of a call to Move; or
- The element it designates has been deleted or removed from the vector that previously contained the element.

The result of "=" or Has\_Element is unspecified if it is called with an invalid cursor parameter. Execution is erroneous if any other subprogram declared in Containers.Vectors is called with an invalid cursor parameter.

Execution is erroneous if the vector associated with the result of a call to Reference or Constant\_Reference is finalized before the result object returned by the call to Reference or Constant\_Reference is finalized.

#### *Implementation Requirements*

No storage associated with a vector object shall be lost upon assignment or scope exit.

The execution of an assignment\_statement for a vector shall have the effect of copying the elements from the source vector object to the target vector object and changing the length of the target object to that of the source object.

*Implementation Advice*

Containers.Vectors should be implemented similarly to an array. In particular, if the length of a vector is  $N$ , then

- the worst-case time complexity of `Element` should be  $O(\log N)$ ;
- the worst-case time complexity of `Append` with `Count=1` when  $N$  is less than the capacity of the vector should be  $O(\log N)$ ; and
- the worst-case time complexity of `Prepend` with `Count=1` and `Delete_First` with `Count=1` should be  $O(N \log N)$ .

The worst-case time complexity of a call on procedure `Sort` of an instance of `Containers.Vectors.Generic_Sorting` should be  $O(N^2)$ , and the average time complexity should be better than  $O(N^2)$ .

`Containers.Vectors.Generic_Sorting.Sort` and `Containers.Vectors.Generic_Sorting.Merge` should minimize copying of elements.

`Move` should not copy elements, and should minimize copying of internal data structures.

If an exception is propagated from a vector operation, no storage should be lost, nor any elements removed from a vector unless specified by the operation.

NOTE 1 All elements of a vector occupy locations in the internal array. If a sparse container is required, a `Hashed_Map` can be used rather than a vector.

NOTE 2 If `Index_Type'Base'First = Index_Type'First` an instance of `Ada.Containers.Vectors` will raise `Constraint_Error`. A value below `Index_Type'First` is required so that an empty vector has a meaningful value of `Last_Index`.

### A.18.3 The Generic Package `Containers.Doubly_Linked_Lists`

The language-defined generic package `Containers.Doubly_Linked_Lists` provides private types `List` and `Cursor`, and a set of operations for each type. A list container is optimized for insertion and deletion at any position.

A doubly-linked list container object manages a linked list of internal *nodes*, each of which contains an element and pointers to the next (successor) and previous (predecessor) internal nodes. A cursor designates a particular node within a list (and by extension the element contained in that node). A cursor keeps designating the same node (and element) as long as the node is part of the container, even if the node is moved in the container.

The *length* of a list is the number of elements it contains.

*Static Semantics*

The generic library package `Containers.Doubly_Linked_Lists` has the following declaration:

```
with Ada.Iterator_Interfaces;
generic
  type Element_Type is private;
  with function "=" (Left, Right : Element_Type)
    return Boolean is <>;
package Ada.Containers.Doubly_Linked_Lists
  with Preelaborate, Remote_Types,
    Nonblocking, Global => in out synchronized is
```

```

type List is tagged private
  with Constant_Indexing => Constant_Reference,
        Variable_Indexing => Reference,
        Default_Iterator  => Iterate,
        Iterator_Element  => Element_Type,
        Iterator_View     => Stable.List,
        Aggregate         => (Empty      => Empty,
                               Add_Unnamed => Append),
        Stable_Properties => (Length,
                               Tampering_With_Cursors_Prohibited,
                               Tampering_With_Elements_Prohibited),
        Default_Initial_Condition =>
          Length (List) = 0 and then
            (not Tampering_With_Cursors_Prohibited (List)) and then
              (not Tampering_With_Elements_Prohibited (List)),
        Preelaborable_Initialization;

type Cursor is private
  with Preelaborable_Initialization;

Empty_List : constant List;

No_Element : constant Cursor;

function Has_Element (Position : Cursor) return Boolean
  with Nonblocking, Global => in all, Use_Formals => null;

function Has_Element (Container : List; Position : Cursor)
  return Boolean
  with Nonblocking, Global => null, Use_Formals => null;

package List_Iterator_Interfaces is new
  Ada.Iterator_Interfaces (Cursor, Has_Element);

function "=" (Left, Right : List) return Boolean;

function Tampering_With_Cursors_Prohibited
  (Container : List) return Boolean
  with Nonblocking, Global => null, Use_Formals => null;

function Tampering_With_Elements_Prohibited
  (Container : List) return Boolean
  with Nonblocking, Global => null, Use_Formals => null;

function Empty return List
  is (Empty_List)
  with Post =>
    not Tampering_With_Elements_Prohibited (Empty'Result) and then
      not Tampering_With_Cursors_Prohibited (Empty'Result) and then
        Length (Empty'Result) = 0;

function Length (Container : List) return Count_Type
  with Nonblocking, Global => null, Use_Formals => null;

function Is_Empty (Container : List) return Boolean
  with Nonblocking, Global => null, Use_Formals => null,
      Post => Is_Empty'Result = (Length (Container) = 0);

procedure Clear (Container : in out List)
  with Pre  => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error,
      Post => Length (Container) = 0;

function Element (Position : Cursor) return Element_Type
  with Pre  => Position /= No_Element or else raise Constraint_Error,
      Nonblocking, Global => in all, Use_Formals => Element_Type;

```

```

function Element (Container : List;
                  Position  : Cursor) return Element_Type
    with Pre => (Position /= No_Element or else
                raise Constraint_Error) and then
                (Has_Element (Container, Position)
                 or else raise Program_Error),
    Nonblocking, Global => null, Use_Formal => Element_Type;

procedure Replace_Element (Container : in out List;
                           Position  : in      Cursor;
                           New_item  : in      Element_Type)
    with Pre  => (not Tampering_With_Elements_Prohibited (Container)
                or else raise Program_Error) and then
                (Position /= No_Element
                 or else raise Constraint_Error) and then
                (Has_Element (Container, Position)
                 or else raise Program_Error);

procedure Query_Element
    (Position : in Cursor;
     Process  : not null access procedure (Element : in Element_Type))
    with Pre  => Position /= No_Element or else raise Constraint_Error,
    Global => in all;

procedure Query_Element
    (Container : in List;
     Position  : in Cursor;
     Process   : not null access procedure (Element : in Element_Type))
    with Pre  => (Position /= No_Element
                or else raise Constraint_Error) and then
                (Has_Element (Container, Position)
                 or else raise Program_Error);

procedure Update_Element
    (Container : in out List;
     Position  : in      Cursor;
     Process   : not null access procedure
                 (Element : in out Element_Type))
    with Pre  => (Position /= No_Element
                or else raise Constraint_Error) and then
                (Has_Element (Container, Position)
                 or else raise Program_Error);

type Constant_Reference_Type
    (Element : not null access constant Element_Type) is private
    with Implicit_Dereference => Element,
    Nonblocking, Global => in out synchronized,
    Default_Initial_Condition => (raise Program_Error);

type Reference_Type (Element : not null access Element_Type) is private
    with Implicit_Dereference => Element,
    Nonblocking, Global => in out synchronized,
    Default_Initial_Condition => (raise Program_Error);

function Constant_Reference (Container : aliased in List;
                             Position  : in Cursor)
    return Constant_Reference_Type
    with Pre  => (Position /= No_Element or else
                raise Constraint_Error) and then
                (Has_Element (Container, Position) or else
                 raise Program_Error),
    Post   => Tampering_With_Cursors_Prohibited (Container),
    Nonblocking, Global => null, Use_Formal => null;

```



```

function Reference (Container : aliased in out List;
                    Position  : in Cursor)
    return Reference_Type
    with Pre => (Position /= No_Element or else
                raise Constraint_Error) and then
                (Has_Element (Container, Position) or else
                 raise Program_Error),
    Post  => Tampering_With_Cursors_Prohibited (Container),
    Nonblocking, Global => null, Use_Formal => null;

procedure Assign (Target : in out List; Source : in List)
    with Pre => not Tampering_With_Cursors_Prohibited (Target)
                or else raise Program_Error,
    Post => Length (Source) = Length (Target);

function Copy (Source : List)
    return List
    with Post =>
        Length (Copy'Result) = Length (Source) and then
        not Tampering_With_Elements_Prohibited (Copy'Result) and then
        not Tampering_With_Cursors_Prohibited (Copy'Result);

procedure Move (Target : in out List;
                Source : in out List)
    with Pre => (not Tampering_With_Cursors_Prohibited (Target)
                or else raise Program_Error) and then
                (not Tampering_With_Cursors_Prohibited (Source)
                 or else raise Program_Error),
    Post => (if not Target'Has_Same_Storage (Source) then
        Length (Target) = Length (Source'Old) and then
        Length (Source) = 0);

procedure Insert (Container : in out List;
                  Before     : in Cursor;
                  New_Item   : in Element_Type;
                  Count      : in Count_Type := 1)
    with Pre => (not Tampering_With_Cursors_Prohibited (Container)
                or else raise Program_Error) and then
                (Before = No_Element or else
                 Has_Element (Container, Before)
                 or else raise Program_Error) and then
                (Length (Container) <= Count_Type'Last - Count
                 or else raise Constraint_Error),
    Post => Length (Container)'Old + Count = Length (Container);

procedure Insert (Container : in out List;
                  Before     : in Cursor;
                  New_Item   : in Element_Type;
                  Position    : out Cursor;
                  Count      : in Count_Type := 1)
    with Pre => (not Tampering_With_Cursors_Prohibited (Container)
                or else raise Program_Error) and then
                (Before = No_Element or else
                 Has_Element (Container, Before)
                 or else raise Program_Error) and then
                (Length (Container) <= Count_Type'Last - Count
                 or else raise Constraint_Error),
    Post => Length (Container)'Old + Count = Length (Container)
        and then Has_Element (Container, Position);

```

```

procedure Insert (Container : in out List;
  Before      : in      Cursor;
  Position    :      out Cursor;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Before = No_Element or else
  Has_Element (Container, Before)
  or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - Count
  or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container)
  and then Has_Element (Container, Position);

procedure Prepend (Container : in out List;
  New_Item    : in      Element_Type;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - Count
  or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container);

procedure Append (Container : in out List;
  New_Item    : in      Element_Type;
  Count       : in      Count_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - Count
  or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container);

procedure Append (Container : in out List;
  New_Item    : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - 1
  or else raise Constraint_Error),
  Post => Length (Container)'Old + 1 = Length (Container);

procedure Delete (Container : in out List;
  Position    : in out Cursor;
  Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error) and then
  (Position /= No_Element
  or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
  or else raise Program_Error),
  Post => Length (Container)'Old - Count <= Length (Container)
  and then Position = No_Element;

procedure Delete_First (Container : in out List;
  Count       : in      Count_Type := 1)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error,
  Post => Length (Container)'Old - Count <= Length (Container);

procedure Delete_Last (Container : in out List;
  Count       : in      Count_Type := 1)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error,
  Post => Length (Container)'Old - Count <= Length (Container);

procedure Reverse_Elements (Container : in out List)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error;

```

```

procedure Swap (Container : in out List;
                 I, J      : in      Cursor)
with Pre => (not Tampering_With_Elements_Prohibited (Container)
             or else raise Program_Error) and then
  (I /= No_Element or else Constraint_Error) and then
  (J /= No_Element or else Constraint_Error) and then
  (Has_Element (Container, I)
   or else raise Program_Error) and then
  (Has_Element (Container, J)
   or else raise Program_Error);

procedure Swap_Links (Container : in out List;
                     I, J      : in      Cursor)
with Pre => (not Tampering_With_Elements_Prohibited (Container)
             or else raise Program_Error) and then
  (I /= No_Element or else Constraint_Error) and then
  (J /= No_Element or else Constraint_Error) and then
  (Has_Element (Container, I)
   or else raise Program_Error) and then
  (Has_Element (Container, J)
   or else raise Program_Error);

procedure Splice (Target   : in out List;
                  Before   : in      Cursor;
                  Source    : in out List)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
  (not Tampering_With_Cursors_Prohibited (Source)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Has_Element (Target, Before)
   or else raise Program_Error) and then
  (Target'Has_Same_Storage (Source) or else
   Length (Target) <= Count_Type'Last - Length (Source)
   or else raise Constraint_Error),
  Post => (if not Target'Has_Same_Storage (Source) then
    (declare
      Result_Length : constant Count_Type :=
        Length (Source)'Old + Length (Target)'Old;
    begin
      Length (Source) = 0 and then
      Length (Target) = Result_Length));

```

```

procedure Splice (Target   : in out List;
                  Before    : in      Cursor;
                  Source     : in out List;
                  Position   : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
(not Tampering_With_Cursors_Prohibited (Source)
 or else raise Program_Error) and then
(Position /= No_Element
 or else raise Constraint_Error) and then
(Has_Element (Source, Position)
 or else raise Program_Error) and then
(Before = No_Element or else
 Has_Element (Target, Before)
 or else raise Program_Error) and then
(Target'Has_Same_Storage (Source) or else
 Length (Target) <= Count_Type'Last - 1
 or else raise Constraint_Error),
Post => (declare
        Org_Target_Length : constant Count_Type :=
            Length (Target)'Old;
        Org_Source_Length : constant Count_Type :=
            Length (Source)'Old;
        begin
            (if Target'Has_Same_Storage (Source) then
                Position = Position'Old
            else
                Length (Source) = Org_Source_Length - 1 and then
                Length (Target) = Org_Target_Length + 1 and then
                Has_Element (Target, Position));
procedure Splice (Container: in out List;
                  Before    : in      Cursor;
                  Position   : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
(Position /= No_Element
 or else raise Constraint_Error) and then
(Has_Element (Container, Position)
 or else raise Program_Error) and then
(Before = No_Element or else
 Has_Element (Container, Before)
 or else raise Program_Error),
Post => Length (Container) = Length (Container)'Old;
function First (Container : List) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
Post => (if not Is_Empty (Container)
        then Has_Element (Container, First'Result)
        else First'Result = No_Element);
function First_Element (Container : List)
return Element_Type
with Pre => (not Is_Empty (Container)
             or else raise Constraint_Error);
function Last (Container : List) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
Post => (if not Is_Empty (Container)
        then Has_Element (Container, Last'Result)
        else Last'Result = No_Element);
function Last_Element (Container : List)
return Element_Type
with Pre => (not Is_Empty (Container)
             or else raise Constraint_Error);

```

```

function Next (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then Next'Result = No_Element);

function Next (Container : List;
  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
    or else raise Program_Error,
    Post => (if Position = No_Element then Next'Result = No_Element
      elsif Next'Result = No_Element then
        Position = Last (Container)
      else Has_Element (Container, Next'Result));

function Previous (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then
      Previous'Result = No_Element);

function Previous (Container : List;
  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
    or else raise Program_Error,
    Post => (if Position = No_Element then
      Previous'Result = No_Element
    elsif Previous'Result = No_Element then
      Position = First (Container)
    else Has_Element (Container, Previous'Result));

procedure Next (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;

procedure Next (Container : in List;
  Position : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
    or else raise Program_Error,
    Post => (if Position /= No_Element
      then Has_Element (Container, Position));

procedure Previous (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;

procedure Previous (Container : in List;
  Position : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
    or else raise Program_Error,
    Post => (if Position /= No_Element then
      Has_Element (Container, Position));

function Find (Container : List;
  Item : Element_Type;
  Position : Cursor := No_Element)
return Cursor
  with Pre => Position = No_Element or else
    Has_Element (Container, Position)
    or else raise Program_Error,
    Post => (if Find'Result /= No_Element
      then Has_Element (Container, Find'Result));

```

```

function Reverse_Find (Container : List;
                        Item       : Element_Type;
                        Position    : Cursor := No_Element)
    return Cursor
with Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Reverse_Find'Result /= No_Element
        then Has_Element (Container, Reverse_Find'Result));

function Contains (Container : List;
                  Item       : Element_Type) return Boolean;

procedure Iterate
    (Container : in List;
     Process   : not null access procedure (Position : in Cursor))
with Allows_Exit;

procedure Reverse_Iterate
    (Container : in List;
     Process   : not null access procedure (Position : in Cursor))
with Allows_Exit;

function Iterate (Container : in List)
    return List_Iterator_Interfaces.Parallel_Reversible_Iterator'Class
with Post => Tampering_With_Cursors_Prohibited (Container);

function Iterate (Container : in List; Start : in Cursor)
    return List_Iterator_Interfaces.Reversible_Iterator'Class
with Pre  => (Start /= No_Element
        or else raise Constraint_Error) and then
        (Has_Element (Container, Start)
        or else raise Program_Error),
    Post  => Tampering_With_Cursors_Prohibited (Container);

generic
    with function "<" (Left, Right : Element_Type)
        return Boolean is <>;
package Generic_Sorting
with Nonblocking, Global => null is

    function Is_Sorted (Container : List) return Boolean;

    procedure Sort (Container : in out List)
        with Pre => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error;

    procedure Merge (Target  : in out List;
                    Source   : in out List)
        with Pre => (not Tampering_With_Cursors_Prohibited (Target)
            or else raise Program_Error) and then
            (not Tampering_With_Elements_Prohibited (Source)
            or else raise Program_Error) and then
            (Length (Target) <= Count_Type'Last - Length (Source)
            or else raise Constraint_Error) and then
            ((Length (Source) = 0 or else
            not Target'Has_Same_Storage (Source))
            or else raise Constraint_Error),
        Post => (declare
            Result_Length : constant Count_Type :=
                Length (Source)'Old + Length (Target)'Old;
            begin
                (Length (Source) = 0 and then
                Length (Target) = Result_Length));

end Generic_Sorting;

package Stable is

```



```

type List (Base : not null access Doubly_Linked_Lists.List) is
  tagged limited private
    with Constant_Indexing => Constant_Reference,
          Variable_Indexing => Reference,
          Default_Iterator  => Iterate,
          Iterator_Element  => Element_Type,
          Stable_Properties => (Length),
          Global => null,
          Default_Initial_Condition => Length (List) = 0,
          Preelaborable_Initialization;

type Cursor is private
  with Preelaborable_Initialization;

Empty_List : constant List;

No_Element : constant Cursor;

function Has_Element (Position : Cursor) return Boolean
  with Nonblocking, Global => in all, Use_Formal => null;

package List_Iterator_Interfaces is new
  Ada.Iterator_Interfaces (Cursor, Has_Element);

procedure Assign (Target : in out Doubly_Linked_Lists.List;
                  Source : in List)
  with Post => Length (Source) = Length (Target);

function Copy (Source : Doubly_Linked_Lists.List) return List
  with Post => Length (Copy'Result) = Length (Source);

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
  with Implicit_Dereference => Element,
        Nonblocking, Global => null, Use_Formal => null,
        Default_Initial_Condition => (raise Program_Error);

type Reference_Type
  (Element : not null access Element_Type) is private
  with Implicit_Dereference => Element,
        Nonblocking, Global => null, Use_Formal => null,
        Default_Initial_Condition => (raise Program_Error);

-- Additional subprograms as described in the text
-- are declared here.

private
  ... -- not specified by the language

end Stable;

private
  ... -- not specified by the language

end Ada.Containers.Doubly_Linked_Lists;

```

The actual function for the generic formal function "=" on Element\_Type values is expected to define a reflexive and symmetric relationship and return the same result value each time it is called with a particular pair of values. If it behaves in some other manner, the functions Find, Reverse\_Find, and "=" on list values return an unspecified value. The exact arguments and number of calls of this generic formal function by the functions Find, Reverse\_Find, and "=" on list values are unspecified.

The type List is used to represent lists. The type List needs finalization (see 10.6).

Empty\_List represents the empty List object. It has a length of 0. If an object of type List is not otherwise initialized, it is initialized to the same value as Empty\_List.

No\_Element represents a cursor that designates no element. If an object of type Cursor is not otherwise initialized, it is initialized to the same value as No\_Element.

The primitive "=" operator for type `Cursor` returns `True` if both cursors are `No_Element`, or designate the same element in the same container.

Execution of the default implementation of the `Input`, `Output`, `Read`, or `Write` attribute of type `Cursor` raises `Program_Error`.

`List'Write` for a `List` object *L* writes `Length(L)` elements of the list to the stream. It may also write additional information about the list.

`List'Read` reads the representation of a list from the stream, and assigns to *Item* a list with the same length and elements as was written by `List'Write`.

Some operations check for "tampering with cursors" of a container because they depend on the set of elements of the container remaining constant, and others check for "tampering with elements" of a container because they depend on elements of the container not being replaced. When tampering with cursors is *prohibited* for a particular list object *L*, `Program_Error` is propagated by the finalization of *L*, as well as by a call that passes *L* to certain of the operations of this package, as indicated by the precondition of such an operation. Similarly, when tampering with elements is *prohibited* for *L*, `Program_Error` is propagated by a call that passes *L* to certain of the other operations of this package, as indicated by the precondition of such an operation.

```
function Has_Element (Position : Cursor) return Boolean
  with Nonblocking, Global => in all, Use_Formal => null;
```

Returns `True` if `Position` designates an element, and returns `False` otherwise.

```
function Has_Element (Container : List; Position : Cursor)
  return Boolean
  with Nonblocking, Global => null, Use_Formal => null;
```

Returns `True` if `Position` designates an element in `Container`, and returns `False` otherwise.

```
function "=" (Left, Right : List) return Boolean;
```

If `Left` and `Right` denote the same list object, then the function returns `True`. If `Left` and `Right` have different lengths, then the function returns `False`. Otherwise, it compares each element in `Left` to the corresponding element in `Right` using the generic formal equality operator. If any such comparison returns `False`, the function returns `False`; otherwise, it returns `True`. Any exception raised during evaluation of element equality is propagated.

```
function Tampering_With_Cursors_Prohibited
  (Container : List) return Boolean
  with Nonblocking, Global => null, Use_Formal => null;
```

Returns `True` if tampering with cursors or tampering with elements is currently prohibited for `Container`, and returns `False` otherwise.

```
function Tampering_With_Elements_Prohibited
  (Container : List) return Boolean
  with Nonblocking, Global => null, Use_Formal => null;
```

Always returns `False`, regardless of whether tampering with elements is prohibited.

```
function Length (Container : List) return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;
```

Returns the number of elements in `Container`.

```
function Is_Empty (Container : List) return Boolean
  with Nonblocking, Global => null, Use_Formal => null,
    Post => Is_Empty'Result = (Length (Container) = 0);
```

Returns `True` if `Container` is empty.

```

procedure Clear (Container : in out List)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = 0;

```

Removes all the elements from Container.

```

function Element (Position : Cursor) return Element_Type
  with Pre => Position /= No_Element or else raise Constraint_Error,
    Nonblocking, Global => in all, Use_Formal => Element_Type;

```

Element returns the element designated by Position.

```

function Element (Container : List;
  Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element or else
    raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error),
    Nonblocking, Global => null, Use_Formal => Element_Type;

```

Element returns the element designated by Position in Container.

```

procedure Replace_Element (Container : in out List;
  Position : in Cursor;
  New_item : in Element_Type)
  with Pre => (not Tampering_With_Elements_Prohibited (Container)
    or else raise Program_Error) and then
    (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error);

```

Replace\_Element assigns the value New\_Item to the element designated by Position. For the purposes of determining whether the parameters overlap in a call to Replace\_Element, the Container parameter is not considered to overlap with any object (including itself).

```

procedure Query_Element
  (Position : in Cursor;
  Process : not null access procedure (Element : in Element_Type))
  with Pre => Position /= No_Element or else raise Constraint_Error,
    Global => in all;

```

Query\_Element calls Process.all with the element designated by Position as the argument. Tampering with the elements of the list that contains the element designated by Position is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Query_Element
  (Container : in List;
  Position : in Cursor;
  Process : not null access procedure (Element : in Element_Type))
  with Pre => (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error);

```

Query\_Element calls Process.all with the element designated by Position as the argument. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Update_Element
(Container : in out List;
 Position  : in      Cursor;
 Process   : not null access procedure
              (Element : in out Element_Type))
with Pre => (Position /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Position)
              or else raise Program_Error);

```

Update\_Element calls Process.all with the element designated by Position as the argument. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

If Element\_Type is unconstrained and definite, then the actual Element parameter of Process.all shall be unconstrained.

```

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

type Reference_Type (Element : not null access Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

```

The types Constant\_Reference\_Type and Reference\_Type need finalization.

```

function Constant_Reference (Container : aliased in List;
                             Position  : in Cursor)
return Constant_Reference_Type
with Pre => (Position /= No_Element or else
              raise Constraint_Error) and then
              (Has_Element (Container, Position) or else
              raise Program_Error),
      Post => Tampering_With_Cursors_Prohibited (Container),
      Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Constant\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read access to an individual element of a list given a cursor.

Constant\_Reference returns an object whose discriminant is an access value that designates the element designated by Position. Tampering with the elements of Container is prohibited while the object returned by Constant\_Reference exists and has not been finalized.

```

function Reference (Container : aliased in out List;
                    Position  : in Cursor)
return Reference_Type
with Pre => (Position /= No_Element or else
              raise Constraint_Error) and then
              (Has_Element (Container, Position) or else
              raise Program_Error),
      Post => Tampering_With_Cursors_Prohibited (Container),
      Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Variable\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read and write access to an individual element of a list given a cursor.

Reference returns an object whose discriminant is an access value that designates the element designated by Position. Tampering with the elements of Container is prohibited while the object returned by Reference exists and has not been finalized.

```

procedure Assign (Target : in out List; Source : in List)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error,
    Post => Length (Source) = Length (Target);

```

If Target denotes the same object as Source, the operation has no effect. Otherwise, the elements of Source are copied to Target as for an `assignment_statement` assigning Source to Target.

```

function Copy (Source : List)
  return List
  with Post =>
    Length (Copy'Result) = Length (Source) and then
    not Tampering_With_Elements_Prohibited (Copy'Result) and then
    not Tampering_With_Cursors_Prohibited (Copy'Result);

```

Returns a list whose elements match the elements of Source.

```

procedure Move (Target : in out List;
                Source : in out List)
  with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
             (not Tampering_With_Cursors_Prohibited (Source)
             or else raise Program_Error),
    Post => (if not Target'Has_Same_Storage (Source) then
            Length (Target) = Length (Source'Old) and then
            Length (Source) = 0);

```

If Target denotes the same object as Source, then the operation has no effect. Otherwise, the operation is equivalent to `Assign (Target, Source)` followed by `Clear (Source)`.

```

procedure Insert (Container : in out List;
                  Before   : in      Cursor;
                  New_Item  : in      Element_Type;
                  Count     : in      Count_Type := 1)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (Before = No_Element or else
              Has_Element (Container, Before)
             or else raise Program_Error) and then
             (Length (Container) <= Count_Type'Last - Count
             or else raise Constraint_Error),
    Post => Length (Container)'Old + Count = Length (Container);

```

Insert inserts Count copies of New\_Item prior to the element designated by Before. If Before equals No\_Element, the new elements are inserted after the last node (if any). Any exception raised during allocation of internal storage is propagated, and Container is not modified.

```

procedure Insert (Container : in out List;
                  Before   : in      Cursor;
                  New_Item  : in      Element_Type;
                  Position  : out      Cursor;
                  Count     : in      Count_Type := 1)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (Before = No_Element or else
              Has_Element (Container, Before)
             or else raise Program_Error) and then
             (Length (Container) <= Count_Type'Last - Count
             or else raise Constraint_Error),
    Post => Length (Container)'Old + Count = Length (Container)
           and then Has_Element (Container, Position);

```

Insert allocates Count copies of New\_Item, and inserts them prior to the element designated by Before. If Before equals No\_Element, the new elements are inserted after the last element (if any). Position designates the first newly-inserted element, or if Count equals 0, then Position is

assigned the value of Before. Any exception raised during allocation of internal storage is propagated, and Container is not modified.

```
procedure Insert (Container : in out List;
                  Before    : in      Cursor;
                  Position   : out Cursor;
                  Count      : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Before = No_Element or else
   Has_Element (Container, Before)
   or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - Count
   or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container)
  and then Has_Element (Container, Position);
```

Insert inserts Count new elements prior to the element designated by Before. If Before equals No\_Element, the new elements are inserted after the last node (if any). The new elements are initialized by default (see 6.3.1). Position designates the first newly-inserted element, or if Count equals 0, then Position is assigned the value of Before. Any exception raised during allocation of internal storage is propagated, and Container is not modified.

```
procedure Prepend (Container : in out List;
                  New_Item   : in      Element_Type;
                  Count      : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - Count
   or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container);
```

Equivalent to Insert (Container, First (Container), New\_Item, Count).

```
procedure Append (Container : in out List;
                  New_Item   : in      Element_Type;
                  Count      : in      Count_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - Count
   or else raise Constraint_Error),
  Post => Length (Container)'Old + Count = Length (Container);
```

Equivalent to Insert (Container, No\_Element, New\_Item, Count).

```
procedure Append (Container : in out List;
                  New_Item   : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - 1
   or else raise Constraint_Error),
  Post => Length (Container)'Old + 1 = Length (Container);
```

Equivalent to Insert (Container, No\_Element, New\_Item, 1).



```

procedure Delete (Container : in out List;
                  Position  : in out Cursor;
                  Count     : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
            (Position /= No_Element
            or else raise Constraint_Error) and then
            (Has_Element (Container, Position)
            or else raise Program_Error),
    Post => Length (Container)'Old - Count <= Length (Container)
    and then Position = No_Element;

```

Delete removes (from Container) Count elements starting at the element designated by Position (or all of the elements starting at Position if there are fewer than Count elements starting at Position). Finally, Position is set to No\_Element.

```

procedure Delete_First (Container : in out List;
                       Count     : in      Count_Type := 1)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error,
    Post => Length (Container)'Old - Count <= Length (Container);

```

If Length (Container) <= Count, then Delete\_First is equivalent to Clear (Container). Otherwise, it removes the first Count nodes from Container.

```

procedure Delete_Last (Container : in out List;
                      Count     : in      Count_Type := 1)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error,
    Post => Length (Container)'Old - Count <= Length (Container);

```

If Length (Container) <= Count, then Delete\_Last is equivalent to Clear (Container). Otherwise, it removes the last Count nodes from Container.

```

procedure Reverse_Elements (Container : in out List)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error;

```

Reorders the elements of Container in reverse order.

```

procedure Swap (Container : in out List;
               I, J       : in      Cursor)
with Pre => (not Tampering_With_Elements_Prohibited (Container)
            or else raise Program_Error) and then
            (I /= No_Element or else Constraint_Error) and then
            (J /= No_Element or else Constraint_Error) and then
            (Has_Element (Container, I)
            or else raise Program_Error) and then
            (Has_Element (Container, J)
            or else raise Program_Error);

```

Swap exchanges the values of the elements designated by I and J.

```

procedure Swap_Links (Container : in out List;
                    I, J       : in      Cursor)
with Pre => (not Tampering_With_Elements_Prohibited (Container)
            or else raise Program_Error) and then
            (I /= No_Element or else Constraint_Error) and then
            (J /= No_Element or else Constraint_Error) and then
            (Has_Element (Container, I)
            or else raise Program_Error) and then
            (Has_Element (Container, J)
            or else raise Program_Error);

```

Swap\_Links exchanges the nodes designated by I and J.

```

procedure Splice (Target    : in out List;
                  Before    : in      Cursor;
                  Source     : in out List)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
             (not Tampering_With_Cursors_Prohibited (Source)
             or else raise Program_Error) and then
             (Before = No_Element or else
              Has_Element (Target, Before)
              or else raise Program_Error) and then
             (Target'Has_Same_Storage (Source) or else
              Length (Target) <= Count_Type'Last - Length (Source)
              or else raise Constraint_Error),
  Post => (if not Target'Has_Same_Storage (Source) then
          (declare
           Result_Length : constant Count_Type :=
             Length (Source)'Old + Length (Target)'Old;
          begin
            Length (Source) = 0 and then
            Length (Target) = Result_Length));

```

If Source denotes the same object as Target, the operation has no effect. Otherwise, Splice reorders elements such that they are removed from Source and moved to Target, immediately prior to Before. If Before equals No\_Element, the nodes of Source are spliced after the last node of Target.

```

procedure Splice (Target    : in out List;
                  Before    : in      Cursor;
                  Source     : in out List;
                  Position   : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
             (not Tampering_With_Cursors_Prohibited (Source)
             or else raise Program_Error) and then
             (Position /= No_Element
             or else raise Constraint_Error) and then
             (Has_Element (Source, Position)
             or else raise Program_Error) and then
             (Before = No_Element or else
              Has_Element (Target, Before)
              or else raise Program_Error) and then
             (Target'Has_Same_Storage (Source) or else
              Length (Target) <= Count_Type'Last - 1
              or else raise Constraint_Error),
  Post => (declare
           Org_Target_Length : constant Count_Type :=
             Length (Target)'Old;
           Org_Source_Length : constant Count_Type :=
             Length (Source)'Old;
          begin
            (if Target'Has_Same_Storage (Source) then
             Position = Position'Old
            else Length (Source) = Org_Source_Length - 1 and then
             Length (Target) = Org_Target_Length + 1 and then
             Has_Element (Target, Position));

```

If Source denotes the same object as Target, then there is no effect if Position equals Before, else the element designated by Position is moved immediately prior to Before, or, if Before equals No\_Element, after the last element. Otherwise, the element designated by Position is removed from Source and moved to Target, immediately prior to Before, or, if Before equals No\_Element, after the last element of Target. Position is updated to represent an element in Target.

```

procedure Splice (Container: in out List;
    Before : in Cursor;
    Position : in Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error) and then
    (Before = No_Element or else
    Has_Element (Container, Before)
    or else raise Program_Error),
    Post => Length (Container) = Length (Container)'Old;

```

If Position equals Before there is no effect. Otherwise, the element designated by Position is moved immediately prior to Before, or, if Before equals No\_Element, after the last element.

```

function First (Container : List) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
    then Has_Element (Container, First'Result)
    else First'Result = No_Element);

```

If Container is empty, First returns No\_Element. Otherwise, it returns a cursor that designates the first node in Container.

```

function First_Element (Container : List)
return Element_Type
with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

```

Equivalent to Element (Container, First\_Index (Container)).

```

function Last (Container : List) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
    then Has_Element (Container, Last'Result)
    else Last'Result = No_Element);

```

If Container is empty, Last returns No\_Element. Otherwise, it returns a cursor that designates the last node in Container.

```

function Last_Element (Container : List)
return Element_Type
with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

```

Equivalent to Element (Last (Container)).

```

function Next (Position : Cursor) return Cursor
with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then Next'Result = No_Element);

```

If Position equals No\_Element or designates the last element of the container, then Next returns the value No\_Element. Otherwise, it returns a cursor that designates the successor of the element designated by Position.

```

function Next (Container : List;
               Position : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Position = No_Element then Next'Result = No_Element
            elsif Next'Result = No_Element then
                Position = Last (Container)
            else Has_Element (Container, Next'Result));

```

Returns a cursor designating the successor of the element designated by Position in Container.

```

function Previous (Position : Cursor) return Cursor
with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then
        Previous'Result = No_Element);

```

If Position equals No\_Element or designates the first element of the container, then Previous returns the value No\_Element. Otherwise, it returns a cursor that designates the predecessor of the element designated by Position.

```

function Previous (Container : List;
                  Position : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Position = No_Element then
        Previous'Result = No_Element
        elsif Previous'Result = No_Element then
            Position = First (Container)
        else Has_Element (Container, Previous'Result));

```

Returns a cursor designating the predecessor of the element designated by Position in Container, if any.

```

procedure Next (Position : in out Cursor)
with Nonblocking, Global => in all, Use_Formal => null;

```

Equivalent to Position := Next (Position).

```

procedure Next (Container : in List;
                Position : in out Cursor)
with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Position /= No_Element
        then Has_Element (Container, Position));

```

Equivalent to Position := Next (Container, Position).

```

procedure Previous (Position : in out Cursor)
with Nonblocking, Global => in all, Use_Formal => null;

```

Equivalent to Position := Previous (Position).

```

procedure Previous (Container : in List;
                   Position : in out Cursor)
with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Position /= No_Element
        then Has_Element (Container, Position));

```

Equivalent to Position := Previous (Container, Position).

```

function Find (Container : List;
               Item      : Element_Type;
               Position   : Cursor := No_Element)
return Cursor
with Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Find'Result /= No_Element
            then Has_Element (Container, Find'Result));

```

Find searches the elements of Container for an element equal to Item (using the generic formal equality operator). The search starts at the element designated by Position, or at the first element if Position equals No\_Element. It proceeds towards Last (Container). If no equal element is found, then Find returns No\_Element. Otherwise, it returns a cursor designating the first equal element encountered.

```

function Reverse_Find (Container : List;
                      Item      : Element_Type;
                      Position   : Cursor := No_Element)
return Cursor
with Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Reverse_Find'Result /= No_Element
            then Has_Element (Container, Reverse_Find'Result));

```

Find searches the elements of Container for an element equal to Item (using the generic formal equality operator). The search starts at the element designated by Position, or at the last element if Position equals No\_Element. It proceeds towards First (Container). If no equal element is found, then Reverse\_Find returns No\_Element. Otherwise, it returns a cursor designating the first equal element encountered.

```

function Contains (Container : List;
                  Item      : Element_Type) return Boolean;

```

Equivalent to Find (Container, Item) /= No\_Element.

```

procedure Iterate
(Container : in List;
 Process   : not null access procedure (Position : in Cursor))
with Allows_Exit;

```

Iterate calls Process.all with a cursor that designates each node in Container, starting with the first node and moving the cursor as per the Next function. Tampering with the cursors of Container is prohibited during the execution of a call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Reverse_Iterate
(Container : in List;
 Process   : not null access procedure (Position : in Cursor))
with Allows_Exit;

```

Iterates over the nodes in Container as per procedure Iterate, except that elements are traversed in reverse order, starting with the last node and moving the cursor as per the Previous function.

```

function Iterate (Container : in List)
return List_Iterator_Interfaces.Parallel_Reversible_Iterator'Class
with Post => Tampering_With_Cursors_Prohibited (Container);

```

Iterate returns an iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each node in Container, starting with the first node and moving the cursor as per the Next function when used as a forward iterator, and starting with the last node and moving the cursor as per the Previous function when used as a reverse iterator, and processing all nodes concurrently when used as a parallel iterator. Tampering with the cursors of

Container is prohibited while the iterator object exists (in particular, in the sequence\_of\_statements of the loop\_statement whose iterator\_specification denotes this object). The iterator object needs finalization.

```
function Iterate (Container : in List; Start : in Cursor)
return List_Iterator_Interfaces.Reversible_Iterator'Class
with Pre    => (Start /= No_Element
                or else raise Constraint_Error) and then
                (Has_Element (Container, Start)
                 or else raise Program_Error),
    Post    => Tampering_With_Cursors_Prohibited (Container);
```

Iterate returns a reversible iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each node in Container, starting with the node designated by Start and moving the cursor as per the Next function when used as a forward iterator, or moving the cursor as per the Previous function when used as a reverse iterator. Tampering with the cursors of Container is prohibited while the iterator object exists (in particular, in the sequence\_of\_statements of the loop\_statement whose iterator\_specification denotes this object). The iterator object needs finalization.

The actual function for the generic formal function "<" of Generic\_Sorting is expected to return the same value each time it is called with a particular pair of element values. It should define a strict weak ordering relationship (see A.18); it should not modify Container. If the actual for "<" behaves in some other manner, the behavior of the subprograms of Generic\_Sorting are unspecified. The number of times the subprograms of Generic\_Sorting call "<" is unspecified.

```
function Is_Sorted (Container : List) return Boolean;
```

Returns True if the elements are sorted smallest first as determined by the generic formal "<" operator; otherwise, Is\_Sorted returns False. Any exception raised during evaluation of "<" is propagated.

```
procedure Sort (Container : in out List)
with Pre    => not Tampering_With_Cursors_Prohibited (Container)
                or else raise Program_Error;
```

Reorders the nodes of Container such that the elements are sorted smallest first as determined by the generic formal "<" operator provided. The sort is stable. Any exception raised during evaluation of "<" is propagated.

```
procedure Merge (Target   : in out List;
                Source   : in out List)
with Pre    => (not Tampering_With_Cursors_Prohibited (Target)
                or else raise Program_Error) and then
                (not Tampering_With_Elements_Prohibited (Source)
                 or else raise Program_Error) and then
                (Length (Target) <= Count_Type'Last - Length (Source)
                 or else raise Constraint_Error) and then
                ((Length (Source) = 0 or else
                 not Target'Has_Same_Storage (Source))
                 or else raise Constraint_Error),
    Post    => (declare
                Result_Length : constant Count_Type :=
                    Length (Source)'Old + Length (Target)'Old;
                begin
                    (Length (Source) = 0 and then
                     Length (Target) = Result_Length));
```

Merge removes elements from Source and inserts them into Target; afterwards, Target contains the union of the elements that were initially in Source and Target; Source is left empty. If Target and Source are initially sorted smallest first, then Target is ordered smallest first as determined by the generic formal "<" operator; otherwise, the order of elements in Target is unspecified. Any exception raised during evaluation of "<" is propagated.



The nested package `Doubly_Linked_Lists.Stable` provides a type `Stable.List` that represents a *stable* list, which is one that cannot grow and shrink. Such a list can be created by calling the `Copy` function, or by establishing a *stabilized view* of an ordinary list.

The subprograms of package `Containers.Doubly_Linked_Lists` that have a parameter or result of type `List` are included in the nested package `Stable` with the same specification, except that the following are omitted:

`Tampering_With_Cursors_Prohibited`, `Tampering_With_Elements_Prohibited`, `Assign`, `Move`, `Insert`, `Append`, `Prepend`, `Clear`, `Delete`, `Delete_First`, `Delete_Last`, `Splice`, `Swap_Links`, and `Reverse_Elements`

The operations of this package are equivalent to those for ordinary lists, except that the calls to `Tampering_With_Cursors_Prohibited` and `Tampering_With_Elements_Prohibited` that occur in preconditions are replaced by `False`, and any that occur in postconditions are replaced by `True`.

If a stable list is declared with the `Base` discriminant designating a pre-existing ordinary list, the stable list represents a stabilized view of the underlying ordinary list, and any operation on the stable list is reflected on the underlying ordinary list. While a stabilized view exists, any operation that tampers with elements performed on the underlying list is prohibited. The finalization of a stable list that provides such a view removes this restriction on the underlying ordinary list (though some other restriction can exist due to other concurrent iterations or stabilized views).

If a stable list is declared without specifying `Base`, the object is necessarily initialized. The initializing expression of the stable list, typically a call on `Copy`, determines the `Length` of the list. The `Length` of a stable list never changes after initialization.

#### *Bounded (Run-Time) Errors*

Calling `Merge` in an instance of `Generic_Sorting` with either `Source` or `Target` not ordered smallest first using the provided generic formal "<" operator is a bounded error. Either `Program_Error` is raised after `Target` is updated as described for `Merge`, or the operation works as defined.

It is a bounded error for the actual function associated with a generic formal subprogram, when called as part of an operation of this package, to tamper with elements of any `List` parameter of the operation. Either `Program_Error` is raised, or the operation works as defined on the value of the `List` either prior to, or subsequent to, some or all of the modifications to the `List`.

It is a bounded error to call any subprogram declared in the visible part of `Containers.Doubly_Linked_Lists` when the associated container has been finalized. If the operation takes `Container` as an **in out** parameter, then it raises `Constraint_Error` or `Program_Error`. Otherwise, the operation either proceeds as it would for an empty container, or it raises `Constraint_Error` or `Program_Error`.

#### *Erroneous Execution*

A `Cursor` value is *invalid* if any of the following have occurred since it was created:

- The list that contains the element it designates has been finalized;
- The list that contains the element it designates has been used as the `Target` of a call to `Assign`, or as the target of an `assignment_statement`;
- The list that contains the element it designates has been used as the `Source` or `Target` of a call to `Move`; or
- The element it designates has been removed from the list that previously contained the element.

The result of `"="` or `Has_Element` is unspecified if it is called with an invalid cursor parameter. Execution is erroneous if any other subprogram declared in `Containers.Doubly_Linked_Lists` is called with an invalid cursor parameter.

Execution is erroneous if the list associated with the result of a call to `Reference` or `Constant_Reference` is finalized before the result object returned by the call to `Reference` or `Constant_Reference` is finalized.

*Implementation Requirements*

No storage associated with a doubly-linked list object shall be lost upon assignment or scope exit.

The execution of an `assignment_statement` for a list shall have the effect of copying the elements from the source list object to the target list object and changing the length of the target object to that of the source object.

*Implementation Advice*

`Containers.Doubly_Linked_Lists` should be implemented similarly to a linked list. In particular, if  $N$  is the length of a list, then the worst-case time complexity of `Element`, `Insert with Count=1`, and `Delete with Count=1` should be  $O(\log N)$ .

The worst-case time complexity of a call on procedure `Sort` of an instance of `Containers.Doubly_Linked_Lists.Generic_Sorting` should be  $O(N^2)$ , and the average time complexity should be better than  $O(N^2)$ .

`Move` should not copy elements, and should minimize copying of internal data structures.

If an exception is propagated from a list operation, no storage should be lost, nor any elements removed from a list unless specified by the operation.

NOTE Sorting a list never copies elements, and is a stable sort (equal elements remain in the original order). This is different than sorting an array or vector, which will often copy elements, and hence is probably not a stable sort.

**A.18.4 Maps**

The language-defined generic packages `Containers.Hashed_Maps` and `Containers.Ordered_Maps` provide private types `Map` and `Cursor`, and a set of operations for each type. A map container allows an arbitrary type to be used as a key to find the element associated with that key. A hashed map uses a hash function to organize the keys, while an ordered map orders the keys per a specified relation.

This subclause describes the declarations that are common to both kinds of maps. See A.18.5 for a description of the semantics specific to `Containers.Hashed_Maps` and A.18.6 for a description of the semantics specific to `Containers.Ordered_Maps`.

*Static Semantics*

The actual function for the generic formal function "=" on `Element_Type` values is expected to define a reflexive and symmetric relationship and return the same result value each time it is called with a particular pair of values. If it behaves in some other manner, the function "=" on map values returns an unspecified value. The exact arguments and number of calls of this generic formal function by the function "=" on map values are unspecified.

The type `Map` is used to represent maps. The type `Map` needs finalization (see 10.6).

A map contains pairs of keys and elements, called *nodes*. Map cursors designate nodes, but also can be thought of as designating an element (the element contained in the node) for consistency with the other containers. There exists an equivalence relation on keys, whose definition is different for hashed maps and ordered maps. A map never contains two or more nodes with equivalent keys. The *length* of a map is the number of nodes it contains.

Each nonempty map has two particular nodes called the *first node* and the *last node* (which may be the same). Each node except for the last node has a *successor node*. If there are no other intervening operations, starting with the first node and repeatedly going to the successor node will visit each node in the map exactly once until the last node is reached. The exact definition of these terms is different for hashed maps and ordered maps.

Some operations check for "tampering with cursors" of a container because they depend on the set of elements of the container remaining constant, and others check for "tampering with elements" of a container because they depend on elements of the container not being replaced. When tampering with cursors is *prohibited* for a particular map object  $M$ , `Program_Error` is propagated by the finalization of  $M$ , as well as by a call that passes  $M$  to certain of the operations of this package, as indicated by the

precondition of such an operation. Similarly, when tampering with elements is *prohibited* for *M*, *Program\_Error* is propagated by a call that passes *M* to certain of the other operations of this package, as indicated by the precondition of such an operation.

*Empty\_Map* represents the empty *Map* object. It has a length of 0. If an object of type *Map* is not otherwise initialized, it is initialized to the same value as *Empty\_Map*.

*No\_Element* represents a cursor that designates no node. If an object of type *Cursor* is not otherwise initialized, it is initialized to the same value as *No\_Element*.

The primitive "=" operator for type *Cursor* returns *True* if both cursors are *No\_Element*, or designate the same element in the same container.

Execution of the default implementation of the *Input*, *Output*, *Read*, or *Write* attribute of type *Cursor* raises *Program\_Error*.

*Map'Write* for a *Map* object *M* writes *Length(M)* elements of the map to the stream. It may also write additional information about the map.

*Map'Read* reads the representation of a map from the stream, and assigns to *Item* a map with the same length and elements as was written by *Map'Write*.

```
function Has_Element (Position : Cursor) return Boolean
with Nonblocking, Global => in all, Use_Formal => null;
```

Returns *True* if *Position* designates an element, and returns *False* otherwise.

```
function Has_Element (Container : Map; Position : Cursor)
return Boolean
with Nonblocking, Global => null, Use_Formal => null;
```

Returns *True* if *Position* designates an element in *Container*, and returns *False* otherwise.

```
function "=" (Left, Right : Map) return Boolean;
```

If *Left* and *Right* denote the same map object, then the function returns *True*. If *Left* and *Right* have different lengths, then the function returns *False*. Otherwise, for each key *K* in *Left*, the function returns *False* if:

- a key equivalent to *K* is not present in *Right*; or
- the element associated with *K* in *Left* is not equal to the element associated with *K* in *Right* (using the generic formal equality operator for elements).

If the function has not returned a result after checking all of the keys, it returns *True*. Any exception raised during evaluation of key equivalence or element equality is propagated.

```
function Tampering_With_Cursors_Prohibited
(Container : Map) return Boolean
with Nonblocking, Global => null, Use_Formal => null;
```

Returns *True* if tampering with cursors or tampering with elements is currently prohibited for *Container*, and returns *False* otherwise.

```
function Tampering_With_Elements_Prohibited
(Container : Map) return Boolean
with Nonblocking, Global => null, Use_Formal => null;
```

Always returns *False*, regardless of whether tampering with elements is prohibited.

```
function Length (Container : Map) return Count_Type
with Nonblocking, Global => null, Use_Formal => null;
```

Returns the number of nodes in *Container*.

```

function Is_Empty (Container : Map) return Boolean
  with Nonblocking, Global => null, Use_Formal => null,
    Post => Is_Empty'Result = (Length (Container) = 0);

```

Returns True if Container is empty.

```

procedure Clear (Container : in out Map)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = 0;

```

Removes all the nodes from Container.

```

function Key (Position : Cursor) return Key_Type
  with Pre => Position /= No_Element
    or else raise Constraint_Error,
    Nonblocking, Global => in all, Use_Formal => Key_Type;

```

Key returns the key component of the node designated by Position.

```

function Key (Container : Map;
  Position : Cursor) return Key_Type
  with Pre => (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error),
    Nonblocking, Global => null, Use_Formal => Key_Type;

```

Key returns the key component of the node designated by Position.

```

function Element (Position : Cursor) return Element_Type
  with Pre => Position /= No_Element
    or else raise Constraint_Error,
    Nonblocking, Global => in all, Use_Formal => Element_Type;

```

Element returns the element component of the node designated by Position.

```

function Element (Container : Map;
  Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error),
    Nonblocking, Global => null, Use_Formal => Element_Type;

```

Element returns the element component of the node designated by Position.

```

procedure Replace_Element (Container : in out Map;
  Position : in Cursor;
  New_item : in Element_Type)
  with Pre => (not Tampering_With_Elements_Prohibited (Container)
    or else raise Program_Error) and then
    (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error);

```

Replace\_Element assigns New\_Item to the element of the node designated by Position. For the purposes of determining whether the parameters overlap in a call to Replace\_Element, the Container parameter is not considered to overlap with any object (including itself).

```

procedure Query_Element
(Position : in Cursor;
 Process  : not null access procedure (Key      : in Key_Type;
                                       Element   : in Element_Type))

with Pre  => Position /= No_Element
      or else raise Constraint_Error,
      Global => in all;

```

Query\_Element calls Process.all with the key and element from the node designated by Position as the arguments. Tampering with the elements of the map that contains the element designated by Position is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Query_Element
(Container : in Map;
 Position  : in Cursor;
 Process   : not null access procedure (Key      : in Key_Type;
                                       Element   : in Element_Type))

with Pre  => (Position /= No_Element
      or else raise Constraint_Error) and then
      (Has_Element (Container, Position)
      or else raise Program_Error);

```

Query\_Element calls Process.all with the key and element from the node designated by Position as the arguments. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Update_Element
(Container : in out Map;
 Position  : in Cursor;
 Process   : not null access procedure (Key      : in Key_Type;
                                       Element   : in out Element_Type))

with Pre  => (Position /= No_Element
      or else raise Constraint_Error) and then
      (Has_Element (Container, Position)
      or else raise Program_Error);

```

Update\_Element calls Process.all with the key and element from the node designated by Position as the arguments. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

If Element\_Type is unconstrained and definite, then the actual Element parameter of Process.all shall be unconstrained.

```

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
  with Implicit_Dereference => Element,
       Nonblocking, Global => in out synchronized,
       Default_Initial_Condition => (raise Program_Error);

type Reference_Type (Element : not null access Element_Type) is private
  with Implicit_Dereference => Element,
       Nonblocking, Global => in out synchronized,
       Default_Initial_Condition => (raise Program_Error);

```

The types Constant\_Reference\_Type and Reference\_Type need finalization.

```

function Constant_Reference (Container : aliased in Map;
                             Position  : in Cursor)
return Constant_Reference_Type
with Pre  => (Position /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Position)
              or else raise Program_Error),
    Post => Tampering_With_Cursors_Prohibited (Container),
    Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Constant\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read access to an individual element of a Map given a cursor.

Constant\_Reference returns an object whose discriminant is an access value that designates the element designated by Position. Tampering with the elements of Container is prohibited while the object returned by Constant\_Reference exists and has not been finalized.

```

function Reference (Container : aliased in out Map;
                    Position  : in Cursor)
return Reference_Type
with Pre  => (Position /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Position)
              or else raise Program_Error),
    Post => Tampering_With_Cursors_Prohibited (Container),
    Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Variable\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read and write access to an individual element of a Map given a cursor.

Reference returns an object whose discriminant is an access value that designates the element designated by Position. Tampering with the elements of Container is prohibited while the object returned by Reference exists and has not been finalized.

```

function Constant_Reference (Container : aliased in Map;
                             Key       : in Key_Type)
return Constant_Reference_Type
with Pre  => Find (Container, Key) /= No_Element
              or else raise Constraint_Error,
    Post => Tampering_With_Cursors_Prohibited (Container),
    Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Constant\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read access to an individual element of a map given a key value.

Equivalent to Constant\_Reference (Container, Find (Container, Key)).

```

function Reference (Container : aliased in out Map;
                    Key       : in Key_Type)
return Reference_Type
with Pre  => Find (Container, Key) /= No_Element
              or else raise Constraint_Error,
    Post => Tampering_With_Cursors_Prohibited (Container),
    Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Variable\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read and write access to an individual element of a map given a key value.

Equivalent to Reference (Container, Find (Container, Key)).



```

procedure Assign (Target : in out Map; Source : in Map)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error,
    Post => Length (Source) = Length (Target);

```

If Target denotes the same object as Source, the operation has no effect. Otherwise, the key/element pairs of Source are copied to Target as for an assignment\_statement assigning Source to Target.

```

procedure Move (Target : in out Map;
                Source : in out Map)
  with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
             (not Tampering_With_Cursors_Prohibited (Source)
             or else raise Program_Error),
    Post => (if not Target'Has_Same_Storage (Source) then
             Length (Target) = Length (Source'Old) and then
             Length (Source) = 0);

```

If Target denotes the same object as Source, then the operation has no effect. Otherwise, the operation is equivalent to Assign (Target, Source) followed by Clear (Source).

```

procedure Insert (Container : in out Map;
                  Key       : in      Key_Type;
                  New_Item  : in      Element_Type;
                  Position   : out Cursor;
                  Inserted   : out Boolean)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (Length (Container) <= Count_Type'Last - 1
             or else raise Constraint_Error),
    Post => (declare
             Original_Length : constant Count_Type :=
               Length (Container)'Old;
             begin
               Has_Element (Container, Position) and then
               (if Inserted then
                 Length (Container) = Original_Length + 1
               else
                 Length (Container) = Original_Length);
             end);

```

Insert checks if a node with a key equivalent to Key is already present in Container. If a match is found, Inserted is set to False and Position designates the element with the matching key. Otherwise, Insert allocates a new node, initializes it to Key and New\_Item, and adds it to Container; Inserted is set to True and Position designates the newly-inserted node. Any exception raised during allocation is propagated and Container is not modified.

```

procedure Insert (Container : in out Map;
                  Key       : in      Key_Type;
                  Position  :      out Cursor;
                  Inserted  :      out Boolean)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - 1
   or else raise Constraint_Error),
  Post => (declare
          Original_Length : constant Count_Type :=
            Length (Container)'Old;
          begin
            Has_Element (Container, Position) and then
            (if Inserted then
              Length (Container) = Original_Length + 1
            else
              Length (Container) = Original_Length));

```

Insert inserts Key into Container as per the five-parameter Insert, with the difference that an element initialized by default (see 6.3.1) is inserted.

```

procedure Insert (Container : in out Map;
                  Key       : in      Key_Type;
                  New_Item  : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - 1
   or else raise Constraint_Error),
  Post => Length (Container) = Length (Container)'Old + 1;

```

Insert inserts Key and New\_Item into Container as per the five-parameter Insert, with the difference that if a node with a key equivalent to Key is already in the map, then Constraint\_Error is propagated.

```

procedure Include (Container : in out Map;
                  Key       : in      Key_Type;
                  New_Item  : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - 1
   or else raise Constraint_Error),
  Post => (declare
          Original_Length : constant Count_Type :=
            Length (Container)'Old;
          begin
            Length (Container)
              in Original_Length | Original_Length + 1);

```

Include inserts Key and New\_Item into Container as per the five-parameter Insert, with the difference that if a node with a key equivalent to Key is already in the map, then this operation assigns Key and New\_Item to the matching node. Any exception raised during assignment is propagated.

```

procedure Replace (Container : in out Map;
                  Key       : in      Key_Type;
                  New_Item  : in      Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error,
  Post => Length (Container) = Length (Container)'Old;

```

Replace checks if a node with a key equivalent to Key is present in Container. If a match is found, Replace assigns Key and New\_Item to the matching node; otherwise, Constraint\_Error is propagated.

```

procedure Exclude (Container : in out Map;
                    Key       : in      Key_Type)
with Pre  => not Tampering_With_Cursors_Prohibited (Container)
          or else raise Program_Error,
      Post => (declare
                Original_Length : constant Count_Type :=
                    Length (Container)'Old;
              begin
                Length (Container)
                  in Original_Length - 1 | Original_Length);

```

Exclude checks if a node with a key equivalent to Key is present in Container. If a match is found, Exclude removes the node from the map.

```

procedure Delete (Container : in out Map;
                  Key       : in      Key_Type)
with Pre  => not Tampering_With_Cursors_Prohibited (Container)
          or else raise Program_Error,
      Post => Length (Container) = Length (Container)'Old - 1;

```

Delete checks if a node with a key equivalent to Key is present in Container. If a match is found, Delete removes the node from the map; otherwise, Constraint\_Error is propagated.

```

procedure Delete (Container : in out Map;
                  Position  : in out Cursor)
with Pre  => (not Tampering_With_Cursors_Prohibited (Container)
          or else raise Program_Error) and then
      (Position /= No_Element
       or else raise Constraint_Error) and then
      (Has_Element (Container, Position)
       or else raise Program_Error),
      Post => Length (Container) = Length (Container)'Old - 1 and then
      Position = No_Element;

```

Delete removes the node designated by Position from the map.

```

function First (Container : Map) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
      Post => (if not Is_Empty (Container)
              then Has_Element (Container, First'Result)
              else First'Result = No_Element);

```

If Length (Container) = 0, then First returns No\_Element. Otherwise, First returns a cursor that designates the first node in Container.

```

function Next (Position : Cursor) return Cursor
with Nonblocking, Global => in all, Use_Formal => null,
      Post => (if Position = No_Element then Next'Result = No_Element);

```

Returns a cursor that designates the successor of the node designated by Position. If Position designates the last node, then No\_Element is returned. If Position equals No\_Element, then No\_Element is returned.

```

function Next (Container : Map;
               Position  : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
      Pre  => Position = No_Element or else
          Has_Element (Container, Position)
          or else raise Program_Error,
      Post => (if Position = No_Element then Next'Result = No_Element
              elsif Next'Result = No_Element then
                Position = Last (Container)
              else Has_Element (Container, Next'Result));

```

Returns a cursor designating the successor of the node designated by Position in Container.

```
procedure Next (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;
```

Equivalent to Position := Next (Position).

```
procedure Next (Container : in Map;
  Position : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
  Pre => Position = No_Element or else
    Has_Element (Container, Position)
  or else raise Program_Error,
  Post => (if Position /= No_Element
    then Has_Element (Container, Position));
```

Equivalent to Position := Next (Container, Position).

```
function Find (Container : Map;
  Key : Key_Type) return Cursor
  with Post => (if Find'Result = No_Element
    then Has_Element (Container, Find'Result));
```

If Length (Container) equals 0, then Find returns No\_Element. Otherwise, Find checks if a node with a key equivalent to Key is present in Container. If a match is found, a cursor designating the matching node is returned; otherwise, No\_Element is returned.

```
function Element (Container : Map;
  Key : Key_Type) return Element_Type;
```

Equivalent to Element (Find (Container, Key)).

```
function Contains (Container : Map;
  Key : Key_Type) return Boolean;
```

Equivalent to Find (Container, Key) /= No\_Element.

```
procedure Iterate
  (Container : in Map;
  Process : not null access procedure (Position : in Cursor))
  with Allows_Exit;
```

Iterate calls Process.all with a cursor that designates each node in Container, starting with the first node and moving the cursor according to the successor relation. Tampering with the cursors of Container is prohibited during the execution of a call on Process.all. Any exception raised by Process.all is propagated.

The nested package Stable provides a type Stable.Map that represents a *stable* map, which is one that cannot grow and shrink. Such a map can be created by calling the Copy function, or by establishing a *stabilized view* of an ordinary map.

The subprograms of the map package that have a parameter or result of type Map are included in the nested package Stable with the same specification, except that the following are omitted:

Tampering\_With\_Cursors\_Prohibited, Tampering\_With\_Elements\_Prohibited, Assign, Move, Insert, Include, Clear, Delete, Exclude, (for Ordered\_Maps) Delete\_First and Delete\_Last, and (for Hashed\_Maps) Reserve\_Capacity

The operations of this package are equivalent to those for ordinary maps, except that the calls to Tampering\_With\_Cursors\_Prohibited and Tampering\_With\_Elements\_Prohibited that occur in preconditions are replaced by False, and any that occur in postconditions are replaced by True.

If a stable map is declared with the Base discriminant designating a pre-existing ordinary map, the stable map represents a stabilized view of the underlying ordinary map, and any operation on the stable map is reflected on the underlying ordinary map. While a stabilized view exists, any operation that tampers with elements performed on the underlying map is prohibited. The finalization of a stable map that provides such a view removes this restriction on the underlying ordinary map (though some other restriction can exist due to other concurrent iterations or stabilized views).

If a stable map is declared without specifying Base, the object is necessarily initialized. The initializing expression of the stable map, typically a call on Copy, determines the Length of the map. The Length of a stable map never changes after initialization.

*Bounded (Run-Time) Errors*

It is a bounded error for the actual function associated with a generic formal subprogram, when called as part of an operation of a map package, to tamper with elements of any map parameter of the operation. Either Program\_Error is raised, or the operation works as defined on the value of the map either prior to, or subsequent to, some or all of the modifications to the map.

It is a bounded error to call any subprogram declared in the visible part of a map package when the associated container has been finalized. If the operation takes Container as an **in out** parameter, then it raises Constraint\_Error or Program\_Error. Otherwise, the operation either proceeds as it would for an empty container, or it raises Constraint\_Error or Program\_Error.

*Erroneous Execution*

A Cursor value is *invalid* if any of the following have occurred since it was created:

- The map that contains the node it designates has been finalized;
- The map that contains the node it designates has been used as the Target of a call to Assign, or as the target of an assignment\_statement;
- The map that contains the node it designates has been used as the Source or Target of a call to Move; or
- The node it designates has been removed from the map that previously contained the node.

The result of "=" or Has\_Element is unspecified if these functions are called with an invalid cursor parameter. Execution is erroneous if any other subprogram declared in Containers.Hashed\_Maps or Containers.Ordered\_Maps is called with an invalid cursor parameter.

Execution is erroneous if the map associated with the result of a call to Reference or Constant\_Reference is finalized before the result object returned by the call to Reference or Constant\_Reference is finalized.

*Implementation Requirements*

No storage associated with a map object shall be lost upon assignment or scope exit.

The execution of an assignment\_statement for a map shall have the effect of copying the elements from the source map object to the target map object and changing the length of the target object to that of the source object.

*Implementation Advice*

Move should not copy elements, and should minimize copying of internal data structures.

If an exception is propagated from a map operation, no storage should be lost, nor any elements removed from a map unless specified by the operation.

## A.18.5 The Generic Package Containers.Hash\_Maps

### Static Semantics

The generic library package Containers.Hash\_Maps has the following declaration:

```

with Ada.Iterator_Interfaces;
generic
  type Key_Type is private;
  type Element_Type is private;
  with function Hash (Key : Key_Type) return Hash_Type;
  with function Equivalent_Keys (Left, Right : Key_Type)
    return Boolean;
  with function "=" (Left, Right : Element_Type)
    return Boolean is <>;
package Ada.Containers.Hash_Maps
  with Preelaborate, Remote_Types,
    Nonblocking, Global => in out synchronized is

  type Map is tagged private
    with Constant_Indexing => Constant_Reference,
      Variable_Indexing => Reference,
      Default_Iterator => Iterate,
      Iterator_Element => Element_Type,
      Iterator_View => Stable.Map,
      Aggregate => (Empty => Empty,
        Add_Named => Insert),
      Stable_Properties => (Length,
        Tampering_With_Cursors_Prohibited,
        Tampering_With_Elements_Prohibited),
      Default_Initial_Condition =>
        Length (Map) = 0 and then
          (not Tampering_With_Cursors_Prohibited (Map)) and then
          (not Tampering_With_Elements_Prohibited (Map)),
      Preelaborable_Initialization;

  type Cursor is private
    with Preelaborable_Initialization;

  Empty_Map : constant Map;
  No_Element : constant Cursor;

  function Has_Element (Position : Cursor) return Boolean
    with Nonblocking, Global => in all, Use_Formal => null;

  function Has_Element (Container : Map; Position : Cursor)
    return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

  package Map_Iterator_Interfaces is new
    Ada.Iterator_Interfaces (Cursor, Has_Element);

  function "=" (Left, Right : Map) return Boolean;

  function Tampering_With_Cursors_Prohibited
    (Container : Map) return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

  function Tampering_With_Elements_Prohibited
    (Container : Map) return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

  function Empty (Capacity : Count_Type := implementation-defined)
    return Map
    with Post =>
      Capacity (Empty'Result) >= Capacity and then
        not Tampering_With_Elements_Prohibited (Empty'Result) and then
        not Tampering_With_Cursors_Prohibited (Empty'Result) and then
        Length (Empty'Result) = 0;

```



```

function Capacity (Container : Map) return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;

procedure Reserve_Capacity (Container : in out Map;
                           Capacity : in Count_Type)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error,
        Post => Container.Capacity >= Capacity;

function Length (Container : Map) return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;

function Is_Empty (Container : Map) return Boolean
  with Nonblocking, Global => null, Use_Formal => null,
        Post => Is_Empty'Result = (Length (Container) = 0);

procedure Clear (Container : in out Map)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error,
        Post => Capacity (Container) = Capacity (Container)'Old and then
        Length (Container) = 0;

function Key (Position : Cursor) return Key_Type
  with Pre => Position /= No_Element
        or else raise Constraint_Error,
        Nonblocking, Global => in all, Use_Formal => Key_Type;

function Key (Container : Map;
              Position : Cursor) return Key_Type
  with Pre => (Position /= No_Element
        or else raise Constraint_Error) and then
        (Has_Element (Container, Position)
        or else raise Program_Error),
        Nonblocking, Global => null, Use_Formal => Key_Type;

function Element (Position : Cursor) return Element_Type
  with Pre => Position /= No_Element
        or else raise Constraint_Error,
        Nonblocking, Global => in all, Use_Formal => Element_Type;

function Element (Container : Map;
                  Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element
        or else raise Constraint_Error) and then
        (Has_Element (Container, Position)
        or else raise Program_Error),
        Nonblocking, Global => null, Use_Formal => Element_Type;

procedure Replace_Element (Container : in out Map;
                          Position : in Cursor;
                          New_item : in Element_Type)
  with Pre => (not Tampering_With_Elements_Prohibited (Container)
        or else raise Program_Error) and then
        (Position /= No_Element
        or else raise Constraint_Error) and then
        (Has_Element (Container, Position)
        or else raise Program_Error);

procedure Query_Element
  (Position : in Cursor;
   Process : not null access procedure (Key : in Key_Type;
                                         Element : in Element_Type))
  with Pre => Position /= No_Element
        or else raise Constraint_Error,
        Global => in all;

```

```

procedure Query_Element
(Container : in Map;
Position   : in Cursor;
Process    : not null access procedure (Key      : in Key_Type;
                                         Element : in Element_Type))

with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error);

procedure Update_Element
(Container : in out Map;
Position   : in Cursor;
Process    : not null access procedure
              (Key      : in Key_Type;
               Element : in out Element_Type))
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error);

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

type Reference_Type (Element : not null access Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

function Constant_Reference (Container : aliased in Map;
                             Position  : in Cursor)
return Constant_Reference_Type
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error),
  Post => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

function Reference (Container : aliased in out Map;
                    Position  : in Cursor)
return Reference_Type
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error),
  Post => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

function Constant_Reference (Container : aliased in Map;
                             Key       : in Key_Type)
return Constant_Reference_Type
with Pre => Find (Container, Key) /= No_Element
             or else raise Constraint_Error,
  Post => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

function Reference (Container : aliased in out Map;
                    Key       : in Key_Type)
return Reference_Type
with Pre => Find (Container, Key) /= No_Element
             or else raise Constraint_Error,
  Post => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

```

```

procedure Assign (Target : in out Map; Source : in Map)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
  Post => Length (Source) = Length (Target) and then
    Capacity (Target) >= Length (Source);

function Copy (Source : Map; Capacity : Count_Type := 0)
  return Map
  with Pre => Capacity = 0 or else Capacity >= Length (Source)
    or else raise Capacity_Error,
  Post =>
    Length (Copy'Result) = Length (Source) and then
    not Tampering_With_Elements_Prohibited (Copy'Result) and then
    not Tampering_With_Cursors_Prohibited (Copy'Result) and then
    Copy'Result.Capacity = (if Capacity = 0 then
      Length (Source) else Capacity);

procedure Move (Target : in out Map;
  Source : in out Map)
  with Pre => (not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error) and then
    (not Tampering_With_Cursors_Prohibited (Source)
    or else raise Program_Error),
  Post => (if not Target'Has_Same_Storage (Source) then
    Length (Target) = Length (Source_Old) and then
    Length (Source) = 0);

procedure Insert (Container : in out Map;
  Key : in Key_Type;
  New_Item : in Element_Type;
  Position : out Cursor;
  Inserted : out Boolean)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
  Post => (declare
    Original_Length : constant Count_Type :=
      Length (Container)'Old;
  begin
    Has_Element (Container, Position) and then
    (if Inserted then
      Length (Container) = Original_Length + 1
    else
      Length (Container) = Original_Length) and then
    Capacity (Container) >= Length (Container);

procedure Insert (Container : in out Map;
  Key : in Key_Type;
  Position : out Cursor;
  Inserted : out Boolean)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
  Post => (declare
    Original_Length : constant Count_Type :=
      Length (Container)'Old;
  begin
    Has_Element (Container, Position) and then
    (if Inserted then
      Length (Container) = Original_Length + 1
    else
      Length (Container) = Original_Length) and then
    Capacity (Container) >= Length (Container);

```

```

procedure Insert (Container : in out Map;
    Key          : in      Key_Type;
    New_Item     : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
    Post => Length (Container) = Length (Container)'Old + 1 and then
    Capacity (Container) >= Length (Container);

procedure Include (Container : in out Map;
    Key          : in      Key_Type;
    New_Item     : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
        begin
            Length (Container)
                in Original_Length | Original_Length + 1) and then
        Capacity (Container) >= Length (Container);

procedure Replace (Container : in out Map;
    Key          : in      Key_Type;
    New_Item     : in      Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = Length (Container)'Old;

procedure Exclude (Container : in out Map;
    Key          : in      Key_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
        begin
            Length (Container)
                in Original_Length - 1 | Original_Length);

procedure Delete (Container : in out Map;
    Key          : in      Key_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = Length (Container)'Old - 1;

procedure Delete (Container : in out Map;
    Position     : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error),
    Post => Length (Container) = Length (Container)'Old - 1 and then
    Position = No_Element;

function First (Container : Map) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
        then Has_Element (Container, First'Result)
        else First'Result = No_Element);

```

```

function Next (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then Next'Result = No_Element);

function Next (Container : Map;
  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position = No_Element then Next'Result = No_Element
      elsif Next'Result = No_Element then
        Position = Last (Container)
      else Has_Element (Container, Next'Result));

procedure Next (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;

procedure Next (Container : in Map;
  Position : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position /= No_Element
      then Has_Element (Container, Position));

function Find (Container : Map;
  Key : Key_Type)
  return Cursor
  with Post => (if Find'Result /= No_Element
    then Has_Element (Container, Find'Result));

function Element (Container : Map;
  Key : Key_Type)
  return Element_Type;

function Contains (Container : Map;
  Key : Key_Type) return Boolean;

function Equivalent_Keys (Left, Right : Cursor)
  return Boolean
  with Pre => (Left /= No_Element and then Right /= No_Element)
    or else raise Constraint_Error,
    Global => in all;

function Equivalent_Keys (Left : Cursor;
  Right : Key_Type)
  return Boolean
  with Pre => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

function Equivalent_Keys (Left : Key_Type;
  Right : Cursor)
  return Boolean
  with Pre => Right /= No_Element or else raise Constraint_Error,
    Global => in all;

procedure Iterate
  (Container : in Map;
  Process : not null access procedure (Position : in Cursor))
  with Allows_Exit;

function Iterate (Container : in Map)
  return Map_Iterator_Interfaces.Parallel_Iterator'Class
  with Post => Tampering_With_Cursors_Prohibited (Container);

package Stable is

```

```

type Map (Base : not null access Hashed_Maps.Map) is
  tagged limited private
    with Constant_Indexing => Constant_Reference,
        Variable_Indexing => Reference,
        Default_Iterator   => Iterate,
        Iterator_Element   => Element_Type,
        Stable_Properties  => (Length),
        Global              => null,
        Default_Initial_Condition => Length (Map) = 0,
        Preelaborable_Initialization;

type Cursor is private
with Preelaborable_Initialization;

Empty_Map : constant Map;

No_Element : constant Cursor;

function Has_Element (Position : Cursor) return Boolean
  with Nonblocking, Global => in all, Use_Formal => null;

package Map_Iterator_Interfaces is new
  Ada.Iterator_Interfaces (Cursor, Has_Element);

procedure Assign (Target : in out Hashed_Maps.Map;
  Source : in Map)
  with Post => Length (Source) = Length (Target);

function Copy (Source : Hashed_Maps.Map) return Map
  with Post => Length (Copy'Result) = Length (Source);

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
  with Implicit_Dereference => Element,
      Nonblocking, Global => null, Use_Formal => null,
      Default_Initial_Condition => (raise Program_Error);

type Reference_Type
  (Element : not null access Element_Type) is private
  with Implicit_Dereference => Element,
      Nonblocking, Global => null, Use_Formal => null,
      Default_Initial_Condition => (raise Program_Error);

-- Additional subprograms as described in the text
-- are declared here.

private
  ... -- not specified by the language
end Stable;

private
  ... -- not specified by the language
end Ada.Containers.Hashed_Maps;

```

An object of type Map contains an expandable hash table, which is used to provide direct access to nodes. The *capacity* of an object of type Map is the maximum number of nodes that can be inserted into the hash table prior to it being automatically expanded.

Two keys *K1* and *K2* are defined to be *equivalent* if Equivalent\_Keys (*K1*, *K2*) returns True.

The actual function for the generic formal function Hash is expected to return the same value each time it is called with a particular key value. For any two equivalent key values, the actual for Hash is expected to return the same value. If the actual for Hash behaves in some other manner, the behavior of this package is unspecified. Which subprograms of this package call Hash, and how many times they call it, is unspecified.

The actual function for the generic formal function Equivalent\_Keys on Key\_Type values is expected to return the same value each time it is called with a particular pair of key values. It should define an equivalence relationship, that is, be reflexive, symmetric, and transitive. If the actual for



Equivalent\_Keys behaves in some other manner, the behavior of this package is unspecified. Which subprograms of this package call Equivalent\_Keys, and how many times they call it, is unspecified.

If the value of a key stored in a node of a map is changed other than by an operation in this package such that at least one of Hash or Equivalent\_Keys give different results, the behavior of this package is unspecified.

Which nodes are the first node and the last node of a map, and which node is the successor of a given node, are unspecified, other than the general semantics described in A.18.4.

```
function Empty (Capacity : Count_Type := implementation-defined)
return Map
with Post =>
    Capacity (Empty'Result) >= Capacity and then
    not Tampering_With_Elements_Prohibited (Empty'Result) and then
    not Tampering_With_Cursors_Prohibited (Empty'Result) and then
    Length (Empty'Result) = 0;
```

Returns an empty map.

```
function Capacity (Container : Map) return Count_Type
with Nonblocking, Global => null, Use_Formal => null;
```

Returns the capacity of Container.

```
procedure Reserve_Capacity (Container : in out Map;
    Capacity : in Count_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Container.Capacity >= Capacity;
```

Reserve\_Capacity allocates a new hash table such that the length of the resulting map can become at least the value Capacity without requiring an additional call to Reserve\_Capacity, and is large enough to hold the current length of Container. Reserve\_Capacity then rehashes the nodes in Container onto the new hash table. It replaces the old hash table with the new hash table, and then deallocates the old hash table. Any exception raised during allocation is propagated and Container is not modified.

```
procedure Clear (Container : in out Map)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Capacity (Container) = Capacity (Container)'Old and then
    Length (Container) = 0;
```

In addition to the semantics described in A.18.4, Clear does not affect the capacity of Container.

```
procedure Assign (Target : in out Map; Source : in Map)
with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
    Post => Length (Source) = Length (Target) and then
    Capacity (Target) >= Length (Source);
```

In addition to the semantics described in A.18.4, if the length of Source is greater than the capacity of Target, Reserve\_Capacity (Target, Length (Source)) is called before assigning any elements.

```

function Copy (Source : Map; Capacity : Count_Type := 0)
  return Map
  with Pre => Capacity = 0 or else Capacity >= Length (Source)
    or else raise Capacity_Error,
    Post =>
      Length (Copy'Result) = Length (Source) and then
      not Tampering_With_Elements_Prohibited (Copy'Result) and then
      not Tampering_With_Cursors_Prohibited (Copy'Result) and then
      Copy'Result.Capacity = (if Capacity = 0 then
        Length (Source) else Capacity);

```

Returns a map whose keys and elements are initialized from the keys and elements of Source.

```

procedure Insert (Container : in out Map;
  Key           : in      Key_Type;
  New_Item      : in      Element_Type;
  Position      : out Cursor;
  Inserted      : out Boolean)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
    Post => (declare
      Original_Length : constant Count_Type :=
        Length (Container)'Old;
      begin
        Has_Element (Container, Position) and then
        (if Inserted then
          Length (Container) = Original_Length + 1
        else
          Length (Container) = Original_Length) and then
          Capacity (Container) >= Length (Container);
      end);

```

In addition to the semantics described in A.18.4, if Length (Container) equals Capacity (Container), then Insert first calls Reserve\_Capacity to increase the capacity of Container to some larger value.

```

function Equivalent_Keys (Left, Right : Cursor)
  return Boolean
  with Pre => (Left /= No_Element and then Right /= No_Element)
    or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Equivalent\_Keys (Key (Left), Key (Right)).

```

function Equivalent_Keys (Left : Cursor;
  Right : Key_Type) return Boolean
  with Pre => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Equivalent\_Keys (Key (Left), Right).

```

function Equivalent_Keys (Left : Key_Type;
  Right : Cursor) return Boolean
  with Pre => Right /= No_Element or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Equivalent\_Keys (Left, Key (Right)).

```

function Iterate (Container : in Map)
  return Map_Iterator_Interfaces.Parallel_Iterator'Class
  with Post => Tampering_With_Cursors_Prohibited (Container);

```

Iterate returns an iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each node in Container, starting with the first node and moving the cursor according to the successor relation when used as a forward iterator, and processing all nodes concurrently when used as a parallel iterator. Tampering with the cursors of Container is

prohibited while the iterator object exists (in particular, in the `sequence_of_statements` of the `loop_statement` whose `iterator_specification` denotes this object). The iterator object needs finalization.

*Implementation Advice*

If  $N$  is the length of a map, the average time complexity of the subprograms `Element`, `Insert`, `Include`, `Replace`, `Delete`, `Exclude`, and `Find` that take a key parameter should be  $O(\log N)$ . The average time complexity of the subprograms that take a cursor parameter should be  $O(1)$ . The average time complexity of `Reserve_Capacity` should be  $O(N)$ .

## A.18.6 The Generic Package Containers.Ordered\_Maps

*Static Semantics*

The generic library package `Containers.Ordered_Maps` has the following declaration:

```
with Ada.Iterator_Interfaces;
generic
  type Key_Type is private;
  type Element_Type is private;
  with function "<" (Left, Right : Key_Type) return Boolean is <>;
  with function "=" (Left, Right : Element_Type) return Boolean is <>;
package Ada.Containers.Ordered_Maps
  with Preelaborate, Remote_Types,
    Nonblocking, Global => in out synchronized is

  function Equivalent_Keys (Left, Right : Key_Type) return Boolean
    is (not ((Left < Right) or (Right < Left)));

  type Map is tagged private
    with Constant_Indexing => Constant_Reference,
      Variable_Indexing => Reference,
      Default_Iterator => Iterate,
      Iterator_Element => Element_Type,
      Iterator_View => Stable.Map,
      Aggregate => (Empty => Empty,
        Add_Named => Insert),
      Stable_Properties => (Length,
        Tampering_With_Cursors_Prohibited,
        Tampering_With_Elements_Prohibited),
      Default_Initial_Condition =>
        Length (Map) = 0 and then
          (not Tampering_With_Cursors_Prohibited (Map)) and then
          (not Tampering_With_Elements_Prohibited (Map)),
      Preelaborable_Initialization;

  type Cursor is private
    with Preelaborable_Initialization;

  Empty_Map : constant Map;

  No_Element : constant Cursor;

  function Has_Element (Position : Cursor) return Boolean
    with Nonblocking, Global => in all, Use_Formal => null;

  function Has_Element (Container : Map; Position : Cursor)
    return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

  package Map_Iterator_Interfaces is new
    Ada.Iterator_Interfaces (Cursor, Has_Element);

  function "=" (Left, Right : Map) return Boolean;

  function Tampering_With_Cursors_Prohibited
    (Container : Map) return Boolean
    with Nonblocking, Global => null, Use_Formal => null;
```

```

function Tampering_With_Elements_Prohibited
  (Container : Map) return Boolean
  with Nonblocking, Global => null, Use_Formal => null;

function Empty return Map
  is (Empty_Map)
  with Post =>
    not Tampering_With_Elements_Prohibited (Empty'Result) and then
    not Tampering_With_Cursors_Prohibited (Empty'Result) and then
    Length (Empty'Result) = 0;

function Length (Container : Map) return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;

function Is_Empty (Container : Map) return Boolean
  with Nonblocking, Global => null, Use_Formal => null,
  Post => Is_Empty'Result = (Length (Container) = 0);

procedure Clear (Container : in out Map)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
  or else raise Program_Error,
  Post => Length (Container) = 0;

function Key (Position : Cursor) return Key_Type
  with Pre => Position /= No_Element or else raise Constraint_Error,
  Nonblocking, Global => in all, Use_Formal => Key_Type;

function Key (Container : Map;
  Position : Cursor) return Key_Type
  with Pre => (Position /= No_Element
  or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
  or else raise Program_Error),
  Nonblocking, Global => null, Use_Formal => Key_Type;

function Element (Position : Cursor) return Element_Type
  with Pre => Position /= No_Element or else raise Constraint_Error,
  Nonblocking, Global => null, Use_Formal => Element_Type;

function Element (Container : Map;
  Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element
  or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
  or else raise Program_Error),
  Nonblocking, Global => null, Use_Formal => Element_Type;

procedure Replace_Element (Container : in out Map;
  Position : in Cursor;
  New_item : in Element_Type)
  with Pre => (not Tampering_With_Elements_Prohibited (Container)
  or else raise Program_Error) and then
  (Position /= No_Element
  or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
  or else raise Program_Error);

procedure Query_Element
  (Position : in Cursor;
  Process : not null access procedure (Key : in Key_Type;
  Element : in Element_Type))
  with Pre => Position /= No_Element
  or else raise Constraint_Error,
  Global => in all;

```

```

procedure Query_Element
(Container : in Map;
Position   : in Cursor;
Process    : not null access procedure (Key      : in Key_Type;
                                         Element : in Element_Type))

with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error);

procedure Update_Element
(Container : in out Map;
Position   : in Cursor;
Process    : not null access procedure
              (Key      : in Key_Type;
               Element : in out Element_Type))

with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error);

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
with Implicit_Dereference => Element,
     Nonblocking, Global => in out synchronized,
     Default_Initial_Condition => (raise Program_Error);

type Reference_Type (Element : not null access Element_Type) is private
with Implicit_Dereference => Element,
     Nonblocking, Global => in out synchronized,
     Default_Initial_Condition => (raise Program_Error);

function Constant_Reference (Container : aliased in Map;
                             Position  : in Cursor)
return Constant_Reference_Type
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error),
  Post => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

function Reference (Container : aliased in out Map;
                    Position   : in Cursor)
return Reference_Type
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error),
  Post => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

function Constant_Reference (Container : aliased in Map;
                             Key       : in Key_Type)
return Constant_Reference_Type
with Pre => Find (Container, Key) /= No_Element
             or else raise Constraint_Error,
  Post => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

function Reference (Container : aliased in out Map;
                    Key       : in Key_Type)
return Reference_Type
with Pre => Find (Container, Key) /= No_Element
             or else raise Constraint_Error,
  Post => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

```

```

procedure Assign (Target : in out Map; Source : in Map)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
  Post => Length (Source) = Length (Target);

function Copy (Source : Map)
  return Map
  with Post =>
    Length (Copy'Result) = Length (Source) and then
    not Tampering_With_Elements_Prohibited (Copy'Result) and then
    not Tampering_With_Cursors_Prohibited (Copy'Result);

procedure Move (Target : in out Map;
  Source : in out Map)
  with Pre => (not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error) and then
    (not Tampering_With_Cursors_Prohibited (Source)
    or else raise Program_Error),
  Post => (if not Target'Has_Same_Storage (Source) then
    Length (Target) = Length (Source'Old) and then
    Length (Source) = 0);

procedure Insert (Container : in out Map;
  Key : in Key_Type;
  New_Item : in Element_Type;
  Position : out Cursor;
  Inserted : out Boolean)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
  Post => (declare
    Original_Length : constant Count_Type :=
      Length (Container)'Old;
  begin
    Has_Element (Container, Position) and then
    (if Inserted then
      Length (Container) = Original_Length + 1
    else
      Length (Container) = Original_Length));

procedure Insert (Container : in out Map;
  Key : in Key_Type;
  Position : out Cursor;
  Inserted : out Boolean)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
  Post => (declare
    Original_Length : constant Count_Type :=
      Length (Container)'Old;
  begin
    Has_Element (Container, Position) and then
    (if Inserted then
      Length (Container) = Original_Length + 1
    else
      Length (Container) = Original_Length));

procedure Insert (Container : in out Map;
  Key : in Key_Type;
  New_Item : in Element_Type)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
  Post => Length (Container) = Length (Container)'Old + 1;

```



```

procedure Include (Container : in out Map;
    Key      : in      Key_Type;
    New_Item : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
    begin
        Length (Container)
            in Original_Length | Original_Length + 1);

procedure Replace (Container : in out Map;
    Key      : in      Key_Type;
    New_Item : in      Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = Length (Container)'Old;

procedure Exclude (Container : in out Map;
    Key      : in      Key_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
    begin
        Length (Container)
            in Original_Length - 1 | Original_Length);

procedure Delete (Container : in out Map;
    Key      : in      Key_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = Length (Container)'Old - 1;

procedure Delete (Container : in out Map;
    Position : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error),
    Post => Length (Container) = Length (Container)'Old - 1 and then
    Position = No_Element;

procedure Delete_First (Container : in out Map)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
    begin
        (if Original_Length = 0 then Length (Container) = 0
        else Length (Container) = Original_Length - 1));

procedure Delete_Last (Container : in out Map)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
    begin
        (if Original_Length = 0 then Length (Container) = 0
        else Length (Container) = Original_Length - 1));

```

```

function First (Container : Map) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
      then Has_Element (Container, First'Result)
      else First'Result = No_Element);

function First_Element (Container : Map) return Element_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

function First_Key (Container : Map) return Key_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

function Last (Container : Map) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
      then Has_Element (Container, Last'Result)
      else Last'Result = No_Element);

function Last_Element (Container : Map) return Element_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

function Last_Key (Container : Map) return Key_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

function Next (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then Next'Result = No_Element);

function Next (Container : Map;
  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position = No_Element then Next'Result = No_Element
      elsif Next'Result = No_Element then
        Position = Last (Container)
      else Has_Element (Container, Next'Result));

procedure Next (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;

procedure Next (Container : in Map;
  Position : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position /= No_Element
      then Has_Element (Container, Position));

function Previous (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then
      Previous'Result = No_Element);

```

```

function Previous (Container : Map;
                  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
    or else raise Program_Error,
    Post => (if Position = No_Element then
      Previous'Result = No_Element
    elsif Previous'Result = No_Element then
      Position = First (Container)
    else Has_Element (Container, Previous'Result));

procedure Previous (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;

procedure Previous (Container : in Map;
                  Position : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
    or else raise Program_Error,
    Post => (if Position /= No_Element
      then Has_Element (Container, Position));

function Find (Container : Map;
               Key : Key_Type) return Cursor
  with Post => (if Find'Result /= No_Element
    then Has_Element (Container, Find'Result));

function Element (Container : Map;
                  Key : Key_Type) return Element_Type;

function Floor (Container : Map;
                Key : Key_Type) return Cursor
  with Post => (if Floor'Result /= No_Element
    then Has_Element (Container, Floor'Result));

function Ceiling (Container : Map;
                  Key : Key_Type) return Cursor
  with Post => (if Ceiling'Result /= No_Element
    then Has_Element (Container, Ceiling'Result));

function Contains (Container : Map;
                  Key : Key_Type) return Boolean;

function "<" (Left, Right : Cursor) return Boolean
  with Pre => (Left /= No_Element and then Right /= No_Element)
    or else raise Constraint_Error,
    Global => in all;

function "=" (Left, Right : Cursor) return Boolean
  with Pre => (Left /= No_Element and then Right /= No_Element)
    or else raise Constraint_Error,
    Global => in all;

function "<" (Left : Cursor; Right : Key_Type) return Boolean
  with Pre => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

function ">" (Left : Cursor; Right : Key_Type) return Boolean
  with Pre => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

function "<" (Left : Key_Type; Right : Cursor) return Boolean
  with Pre => Right /= No_Element or else raise Constraint_Error,
    Global => in all;

function ">" (Left : Key_Type; Right : Cursor) return Boolean
  with Pre => Right /= No_Element or else raise Constraint_Error,
    Global => in all;

```

```

procedure Iterate
  (Container : in Map;
   Process   : not null access procedure (Position : in Cursor))
  with Allows_Exit;

procedure Reverse_Iterate
  (Container : in Map;
   Process   : not null access procedure (Position : in Cursor))
  with Allows_Exit;

function Iterate (Container : in Map)
  return Map_Iterator_Interfaces.Parallel_Reversible_Iterator'Class
  with Post => Tampering_With_Cursors_Prohibited (Container);

function Iterate (Container : in Map; Start : in Cursor)
  return Map_Iterator_Interfaces.Reversible_Iterator'Class
  with Pre  => (Start /= No_Element
               or else raise Constraint_Error) and then
               (Has_Element (Container, Start)
               or else raise Program_Error),
   Post => Tampering_With_Cursors_Prohibited (Container);

package Stable is

  type Map (Base : not null access Ordered_Maps.Map) is
    tagged limited private
    with Constant_Indexing => Constant_Reference,
         Variable_Indexing => Reference,
         Default_Iterator  => Iterate,
         Iterator_Element  => Element_Type,
         Stable_Properties => (Length),
         Global             => null,
         Default_Initial_Condition => Length (Map) = 0,
         Preelaborable_Initialization;

  type Cursor is private
    with Preelaborable_Initialization;

  Empty_Map : constant Map;

  No_Element : constant Cursor;

  function Has_Element (Position : Cursor) return Boolean
    with Nonblocking, Global => in all, Use_Formal => null;

  package Map_Iterator_Interfaces is new
    Ada.Iterator_Interfaces (Cursor, Has_Element);

  procedure Assign (Target : in out Ordered_Maps.Map;
                   Source : in Map)
    with Post => Length (Source) = Length (Target);

  function Copy (Source : Ordered_Maps.Map) return Map
    with Post => Length (Copy'Result) = Length (Source);

  type Constant_Reference_Type
    (Element : not null access constant Element_Type) is private
    with Implicit_Dereference => Element,
         Nonblocking, Global => null, Use_Formal => null,
         Default_Initial_Condition => (raise Program_Error);

  type Reference_Type
    (Element : not null access Element_Type) is private
    with Implicit_Dereference => Element,
         Nonblocking, Global => null, Use_Formal => null,
         Default_Initial_Condition => (raise Program_Error);

  -- Additional subprograms as described in the text
  -- are declared here.

private

  ... -- not specified by the language

```

```

    end Stable;

private
    ... -- not specified by the language

end Ada.Containers.Ordered_Maps;

```

Two keys  $K1$  and  $K2$  are *equivalent* if both  $K1 < K2$  and  $K2 < K1$  return False, using the generic formal "<" operator for keys. Function `Equivalent_Keys` returns True if Left and Right are equivalent, and False otherwise.

The actual function for the generic formal function "<" on `Key_Type` values is expected to return the same value each time it is called with a particular pair of key values. It should define a strict weak ordering relationship (see A.18). If the actual for "<" behaves in some other manner, the behavior of this package is unspecified. Which subprograms of this package call "<" and how many times they call it, is unspecified.

If the value of a key stored in a map is changed other than by an operation in this package such that at least one of "<" or "=" give different results, the behavior of this package is unspecified.

The *first node* of a nonempty map is the one whose key is less than the key of all the other nodes in the map. The *last node* of a nonempty map is the one whose key is greater than the key of all the other elements in the map. The *successor* of a node is the node with the smallest key that is larger than the key of the given node. The *predecessor* of a node is the node with the largest key that is smaller than the key of the given node. All comparisons are done using the generic formal "<" operator for keys.

```

function Copy (Source : Map)
return Map
with Post =>
    Length (Copy'Result) = Length (Source) and then
    not Tampering_With_Elements_Prohibited (Copy'Result) and then
    not Tampering_With_Cursors_Prohibited (Copy'Result);

```

Returns a map whose keys and elements are initialized from the corresponding keys and elements of Source.

```

procedure Delete_First (Container : in out Map)
with Pre  => not Tampering_With_Cursors_Prohibited (Container)
or else raise Program_Error,
Post => (declare
    Original_Length : constant Count_Type :=
        Length (Container)'Old;
begin
    if Original_Length = 0 then Length (Container) = 0
    else Length (Container) = Original_Length - 1);

```

If Container is empty, `Delete_First` has no effect. Otherwise, the node designated by `First (Container)` is removed from Container. `Delete_First` tampers with the cursors of Container.

```

procedure Delete_Last (Container : in out Map)
with Pre  => not Tampering_With_Cursors_Prohibited (Container)
or else raise Program_Error,
Post => (declare
    Original_Length : constant Count_Type :=
        Length (Container)'Old;
begin
    if Original_Length = 0 then Length (Container) = 0
    else Length (Container) = Original_Length - 1);

```

If Container is empty, `Delete_Last` has no effect. Otherwise, the node designated by `Last (Container)` is removed from Container. `Delete_Last` tampers with the cursors of Container.

```
function First_Element (Container : Map) return Element_Type
with Pre => (not Is_Empty (Container)
             or else raise Constraint_Error);
```

Equivalent to Element (First (Container)).

```
function First_Key (Container : Map) return Key_Type
with Pre => (not Is_Empty (Container)
             or else raise Constraint_Error);
```

Equivalent to Key (First (Container)).

```
function Last (Container : Map) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
     Post => (if not Is_Empty (Container)
              then Has_Element (Container, Last'Result)
              else Last'Result = No_Element);
```

Returns a cursor that designates the last node in Container. If Container is empty, returns No\_Element.

```
function Last_Element (Container : Map) return Element_Type
with Pre => (not Is_Empty (Container)
             or else raise Constraint_Error);
```

Equivalent to Element (Last (Container)).

```
function Last_Key (Container : Map) return Key_Type
with Pre => (not Is_Empty (Container)
             or else raise Constraint_Error);
```

Equivalent to Key (Last (Container)).

```
function Previous (Position : Cursor) return Cursor
with Nonblocking, Global => in all, Use_Formal => null,
     Post => (if Position = No_Element then
              Previous'Result = No_Element);
```

If Position equals No\_Element, then Previous returns No\_Element. Otherwise, Previous returns a cursor designating the predecessor node of the one designated by Position. If Position designates the first element, then Previous returns No\_Element.

```
function Previous (Container : Map;
                  Position : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
     Pre => Position = No_Element or else
           Has_Element (Container, Position)
           or else raise Program_Error,
     Post => (if Position = No_Element then
              Previous'Result = No_Element
            elsif Previous'Result = No_Element then
              Position = First (Container)
            else Has_Element (Container, Previous'Result));
```

Returns a cursor designating the predecessor of the node designated by Position in Container, if any.

```
procedure Previous (Position : in out Cursor)
with Nonblocking, Global => in all, Use_Formal => null;
```

Equivalent to Position := Previous (Position).



```

procedure Previous (Container : in      Map;
                    Position  : in out Cursor)
with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
        Has_Element (Container, Position)
    or else raise Program_Error,
    Post => (if Position /= No_Element
            then Has_Element (Container, Position));

```

Equivalent to Position := Previous (Container, Position).

```

function Floor (Container : Map;
                Key       : Key_Type) return Cursor
with Post => (if Floor'Result /= No_Element
            then Has_Element (Container, Floor'Result));

```

Floor searches for the last node whose key is not greater than Key, using the generic formal "<" operator for keys. If such a node is found, a cursor that designates it is returned. Otherwise, No\_Element is returned.

```

function Ceiling (Container : Map;
                 Key       : Key_Type) return Cursor
with Post => (if Ceiling'Result /= No_Element
            then Has_Element (Container, Ceiling'Result));

```

Ceiling searches for the first node whose key is not less than Key, using the generic formal "<" operator for keys. If such a node is found, a cursor that designates it is returned. Otherwise, No\_Element is returned.

```

function "<" (Left, Right : Cursor) return Boolean
with Pre  => (Left /= No_Element and then Right /= No_Element)
    or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Key (Left) < Key (Right).

```

function ">" (Left, Right : Cursor) return Boolean
with Pre  => (Left /= No_Element and then Right /= No_Element)
    or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Key (Right) < Key (Left).

```

function "<" (Left : Cursor; Right : Key_Type) return Boolean
with Pre  => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Key (Left) < Right.

```

function ">" (Left : Cursor; Right : Key_Type) return Boolean
with Pre  => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Right < Key (Left).

```

function "<" (Left : Key_Type; Right : Cursor) return Boolean
with Pre  => Right /= No_Element or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Left < Key (Right).

```

function ">" (Left : Key_Type; Right : Cursor) return Boolean
with Pre  => Right /= No_Element or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Key (Right) < Left.

**procedure** Reverse\_Iterate

```
(Container : in Map;
  Process   : not null access procedure (Position : in Cursor))
with Allows_Exit;
```

Iterates over the nodes in Container as per procedure Iterate, with the difference that the nodes are traversed in predecessor order, starting with the last node.

**function** Iterate (Container : in Map)

```
return Map_Iterator_Interfaces.Parallel_Reversible_Iterator'Class
with Post => Tampering_With_Cursors_Prohibited (Container);
```

Iterate returns an iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each node in Container, starting with the first node and moving the cursor according to the successor relation when used as a forward iterator, and starting with the last node and moving the cursor according to the predecessor relation when used as a reverse iterator, and processing all nodes concurrently when used as a parallel iterator. Tampering with the cursors of Container is prohibited while the iterator object exists (in particular, in the sequence\_of\_statements of the loop\_statement whose iterator\_specification denotes this object). The iterator object needs finalization.

**function** Iterate (Container : in Map; Start : in Cursor)

```
return Map_Iterator_Interfaces.Reversible_Iterator'Class
with Pre  => (Start /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Start)
              or else raise Program_Error),
  Post => Tampering_With_Cursors_Prohibited (Container);
```

Iterate returns a reversible iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each node in Container, starting with the node designated by Start and moving the cursor according to the successor relation when used as a forward iterator, or moving the cursor according to the predecessor relation when used as a reverse iterator. Tampering with the cursors of Container is prohibited while the iterator object exists (in particular, in the sequence\_of\_statements of the loop\_statement whose iterator\_specification denotes this object). The iterator object needs finalization.

*Implementation Advice*

If  $N$  is the length of a map, then the worst-case time complexity of the Element, Insert, Include, Replace, Delete, Exclude, and Find operations that take a key parameter should be  $O((\log N)^2)$  or better. The worst-case time complexity of the subprograms that take a cursor parameter should be  $O(1)$ .

**A.18.7 Sets**

The language-defined generic packages Containers.Hash\_Sets and Containers.Ordered\_Sets provide private types Set and Cursor, and a set of operations for each type. A set container allows elements of an arbitrary type to be stored without duplication. A hashed set uses a hash function to organize elements, while an ordered set orders its element per a specified relation.

This subclause describes the declarations that are common to both kinds of sets. See A.18.8 for a description of the semantics specific to Containers.Hash\_Sets and A.18.9 for a description of the semantics specific to Containers.Ordered\_Sets.

*Static Semantics*

The actual function for the generic formal function "=" on Element\_Type values is expected to define a reflexive and symmetric relationship and return the same result value each time it is called with a particular pair of values. If it behaves in some other manner, the function "=" on set values returns an unspecified value. The exact arguments and number of calls of this generic formal function by the function "=" on set values are unspecified.

The type Set is used to represent sets. The type Set needs finalization (see 10.6).

A set contains elements. Set cursors designate elements. There exists an equivalence relation on elements, whose definition is different for hashed sets and ordered sets. A set never contains two or more equivalent elements. The *length* of a set is the number of elements it contains.

Each nonempty set has two particular elements called the *first element* and the *last element* (which may be the same). Each element except for the last element has a *successor element*. If there are no other intervening operations, starting with the first element and repeatedly going to the successor element will visit each element in the set exactly once until the last element is reached. The exact definition of these terms is different for hashed sets and ordered sets.

Some operations check for “tampering with cursors” of a container because they depend on the set of elements of the container remaining constant and on elements of the container not being replaced. When tampering with cursors is *prohibited* for a particular set object *S*, Program\_Error is propagated by the finalization of *S*, as well as by a call that passes *S* to certain of the operations of this package, as indicated by the precondition of such an operation.

Empty\_Set represents the empty Set object. It has a length of 0. If an object of type Set is not otherwise initialized, it is initialized to the same value as Empty\_Set.

No\_Element represents a cursor that designates no element. If an object of type Cursor is not otherwise initialized, it is initialized to the same value as No\_Element.

The primitive “=” operator for type Cursor returns True if both cursors are No\_Element, or designate the same element in the same container.

Execution of the default implementation of the Input, Output, Read, or Write attribute of type Cursor raises Program\_Error.

Set'Write for a Set object *S* writes Length(*S*) elements of the set to the stream. It may also write additional information about the set.

Set'Read reads the representation of a set from the stream, and assigns to *Item* a set with the same length and elements as was written by Set'Write.

```
function Has_Element (Position : Cursor) return Boolean
with Nonblocking, Global => in all, Use_Formal => null;
```

Returns True if Position designates an element, and returns False otherwise.

```
function Has_Element (Container : Set; Position : Cursor)
return Boolean
with Nonblocking, Global => null, Use_Formal => null;
```

Returns True if Position designates an element in Container, and returns False otherwise.

```
function "=" (Left, Right : Set) return Boolean;
```

If Left and Right denote the same set object, then the function returns True. If Left and Right have different lengths, then the function returns False. Otherwise, for each element *E* in Left, the function returns False if an element equal to *E* (using the generic formal equality operator) is not present in Right. If the function has not returned a result after checking all of the elements, it returns True. Any exception raised during evaluation of element equality is propagated.

```
function Equivalent_Sets (Left, Right : Set) return Boolean;
```

If Left and Right denote the same set object, then the function returns True. If Left and Right have different lengths, then the function returns False. Otherwise, for each element *E* in Left, the function returns False if an element equivalent to *E* is not present in Right. If the function has not returned a result after checking all of the elements, it returns True. Any exception raised during evaluation of element equivalence is propagated.

```
function Tampering_With_Cursors_Prohibited
  (Container : Set) return Boolean
  with Nonblocking, Global => null, Use_Formal => null;
```

Returns True if tampering with cursors is currently prohibited for Container, and returns False otherwise.

```
function To_Set (New_Item : Element_Type) return Set
  with Post => Length (To_Set'Result) = 1 and then
    not Tampering_with_Cursors_Prohibited (To_Set'Result);
```

Returns a set containing the single element New\_Item.

```
function Length (Container : Set) return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;
```

Returns the number of elements in Container.

```
function Is_Empty (Container : Set) return Boolean
  with Nonblocking, Global => null, Use_Formal => null,
    Post => Is_Empty'Result = (Length (Container) = 0);
```

Returns True if Container is empty.

```
procedure Clear (Container : in out Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = 0;
```

Removes all the elements from Container.

```
function Element (Position : Cursor) return Element_Type
  with Pre => Position /= No_Element or else raise Constraint_Error,
    Nonblocking, Global => in all, Use_Formal => Element_Type;
```

Element returns the element designated by Position.

```
function Element (Container : Set;
  Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error),
    Nonblocking, Global => null, Use_Formal => Element_Type;
```

Element returns the element designated by Position.

```
procedure Replace_Element (Container : in out Set;
  Position : in Cursor;
  New_item : in Element_Type)
  with Pre => (not Tampering_With_Elements_Prohibited (Container)
    or else raise Program_Error) and then
    (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error);
```

Replace\_Element assigns New\_Item to the element designated by Position. Any exception raised by the assignment is propagated. For the purposes of determining whether the parameters overlap in a call to Replace\_Element, the Container parameter is not considered to overlap with any object (including itself).

```

procedure Query_Element
(Position : in Cursor;
Process  : not null access procedure (Element : in Element_Type))
with Pre  => Position /= No_Element
           or else raise Constraint_Error,
      Global => in all;

```

Query\_Element calls Process.all with the element designated by Position as the argument. Tampering with the elements of the set that contains the element designated by Position is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Query_Element
(Container : in Set;
Position  : in Cursor;
Process   : not null access procedure (Element : in Element_Type))
with Pre  => (Position /= No_Element
           or else raise Constraint_Error) and then
           (Has_Element (Container, Position)
           or else raise Program_Error);

```

Query\_Element calls Process.all with the key and element from the node designated by Position as the arguments. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

type Constant_Reference_Type
(Element : not null access constant Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

```

The type Constant\_Reference\_Type needs finalization.

```

function Constant_Reference (Container : aliased in Set;
                             Position  : in Cursor)
return Constant_Reference_Type
with Pre  => (Position /= No_Element
           or else raise Constraint_Error) and then
           (Has_Element (Container, Position)
           or else raise Program_Error),
      Post => Tampering_With_Cursors_Prohibited (Container),
      Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Constant\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read access to an individual element of a set given a cursor.

Constant\_Reference returns an object whose discriminant is an access value that designates the element designated by Position. Tampering with the cursors of Container is prohibited while the object returned by Constant\_Reference exists and has not been finalized.

```

procedure Assign (Target : in out Set; Source : in Set)
with Pre  => not Tampering_With_Cursors_Prohibited (Target)
           or else raise Program_Error,
      Post => Length (Source) = Length (Target);

```

If Target denotes the same object as Source, the operation has no effect. Otherwise, the elements of Source are copied to Target as for an assignment\_statement assigning Source to Target.

```

procedure Move (Target : in out Set;
                Source : in out Set)
  with Pre => (not Tampering_With_Cursors_Prohibited (Target)
              or else raise Program_Error) and then
              (not Tampering_With_Cursors_Prohibited (Source)
              or else raise Program_Error),
  Post => (if not Target'Has_Same_Storage (Source) then
          Length (Target) = Length (Source'Old) and then
          Length (Source) = 0);

```

If Target denotes the same object as Source, then the operation has no effect. Otherwise, the operation is equivalent to Assign (Target, Source) followed by Clear (Source).

```

procedure Insert (Container : in out Set;
                 New_Item  : in      Element_Type;
                 Position   : out Cursor;
                 Inserted   : out Boolean)
  with Pre => (not Tampering_With_Elements_Prohibited (Container)
              or else raise Program_Error) and then
              (Length (Container) <= Count_Type'Last - 1
              or else raise Constraint_Error),
  Post => (declare
          Original_Length : constant Count_Type :=
            Length (Container)'Old;
          begin
            Has_Element (Container, Position) and then
            (if Inserted then
              Length (Container) = Original_Length + 1
            else
              Length (Container) = Original_Length));

```

Insert checks if an element equivalent to New\_Item is already present in Container. If a match is found, Inserted is set to False and Position designates the matching element. Otherwise, Insert adds New\_Item to Container; Inserted is set to True and Position designates the newly-inserted element. Any exception raised during allocation is propagated and Container is not modified.

```

procedure Insert (Container : in out Set;
                 New_Item  : in      Element_Type)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
              or else raise Program_Error) and then
              (Length (Container) <= Count_Type'Last - 1
              or else raise Constraint_Error),
  Post => Length (Container) = Length (Container)'Old + 1;

```

Insert inserts New\_Item into Container as per the four-parameter Insert, with the difference that if an element equivalent to New\_Item is already in the set, then Constraint\_Error is propagated.

```

procedure Include (Container : in out Set;
                  New_Item  : in      Element_Type)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
              or else raise Program_Error) and then
              (Length (Container) <= Count_Type'Last - 1
              or else raise Constraint_Error),
  Post => (declare
          Original_Length : constant Count_Type :=
            Length (Container)'Old;
          begin
            Length (Container)
              in Original_Length | Original_Length + 1);

```

Include inserts New\_Item into Container as per the four-parameter Insert, with the difference that if an element equivalent to New\_Item is already in the set, then it is replaced. Any exception raised during assignment is propagated.



```

procedure Replace (Container : in out Set;
                   New_Item  : in      Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
           or else raise Program_Error,
       Post => Length (Container) = Length (Container)'Old;

```

Replace checks if an element equivalent to New\_Item is already in the set. If a match is found, that element is replaced with New\_Item; otherwise, Constraint\_Error is propagated.

```

procedure Exclude (Container : in out Set;
                   Item       : in      Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
           or else raise Program_Error,
       Post => (declare
                 Original_Length : constant Count_Type :=
                     Length (Container)'Old;
               begin
                 Length (Container) in
                     Original_Length - 1 | Original_Length);

```

Exclude checks if an element equivalent to Item is present in Container. If a match is found, Exclude removes the element from the set.

```

procedure Delete (Container : in out Set;
                  Item       : in      Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
           or else raise Program_Error,
       Post => Length (Container) = Length (Container)'Old - 1;

```

Delete checks if an element equivalent to Item is present in Container. If a match is found, Delete removes the element from the set; otherwise, Constraint\_Error is propagated.

```

procedure Delete (Container : in out Set;
                  Position  : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
           or else raise Program_Error) and then
           (Position /= No_Element
           or else raise Constraint_Error) and then
           (Has_Element (Container, Position)
           or else raise Program_Error),
       Post => Length (Container) = Length (Container)'Old - 1 and then
           Position = No_Element;

```

Delete removes the element designated by Position from the set.

```

procedure Union (Target : in out Set;
                 Source  : in      Set)
with Pre => not Tampering_With_Cursors_Prohibited (Target)
           or else raise Program_Error,
       Post => Length (Target) <= Length (Target)'Old + Length (Source);

```

Union inserts into Target the elements of Source that are not equivalent to some element already in Target.

```

function Union (Left, Right : Set) return Set
with Post => Length (Union'Result) <=
             Length (Left) + Length (Right) and then
             not Tampering_With_Cursors_Prohibited (Union'Result);

```

Returns a set comprising all of the elements of Left, and the elements of Right that are not equivalent to some element of Left.

```

procedure Intersection (Target : in out Set;
                        Source : in Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
        or else raise Program_Error,
        Post => Length (Target) <= Length (Target)'Old + Length (Source);

```

Intersection deletes from Target the elements of Target that are not equivalent to some element of Source.

```

function Intersection (Left, Right : Set) return Set
  with Post => Length (Intersection'Result) <=
        Length (Left) + Length (Right) and then
        not Tampering_With_Cursors_Prohibited (Intersection'Result);

```

Returns a set comprising all the elements of Left that are equivalent to the some element of Right.

```

procedure Difference (Target : in out Set;
                     Source : in Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
        or else raise Program_Error,
        Post => Length (Target) <= Length (Target)'Old + Length (Source);

```

If Target denotes the same object as Source, then Difference clears Target. Otherwise, it deletes from Target the elements that are equivalent to some element of Source.

```

function Difference (Left, Right : Set) return Set
  with Post => Length (Difference'Result) <= Length (Left) +
        Length (Right) and then
        not Tampering_With_Cursors_Prohibited (Difference'Result);

```

Returns a set comprising the elements of Left that are not equivalent to some element of Right.

```

procedure Symmetric_Difference (Target : in out Set;
                               Source : in Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
        or else raise Program_Error,
        Post => Length (Target) <= Length (Target)'Old + Length (Source);

```

If Target denotes the same object as Source, then Symmetric\_Difference clears Target. Otherwise, it deletes from Target the elements that are equivalent to some element of Source, and inserts into Target the elements of Source that are not equivalent to some element of Target.

```

function Symmetric_Difference (Left, Right : Set) return Set
  with Post => Length (Symmetric_Difference'Result) <=
        Length (Left) + Length (Right) and then
        not Tampering_With_Cursors_Prohibited (
            Symmetric_Difference'Result);

```

Returns a set comprising the elements of Left that are not equivalent to some element of Right, and the elements of Right that are not equivalent to some element of Left.

```

function Overlap (Left, Right : Set) return Boolean;

```

If an element of Left is equivalent to some element of Right, then Overlap returns True. Otherwise, it returns False.

```

function Is_Subset (Subset : Set;
                   Of_Set : Set) return Boolean;

```

If an element of Subset is not equivalent to some element of Of\_Set, then Is\_Subset returns False. Otherwise, it returns True.

```

function First (Container : Set) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
      then Has_Element (Container, First'Result)
      else First'Result = No_Element);

```

If Length (Container) = 0, then First returns No\_Element. Otherwise, First returns a cursor that designates the first element in Container.

```

function Next (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then Next'Result = No_Element);

```

Returns a cursor that designates the successor of the element designated by Position. If Position designates the last element, then No\_Element is returned. If Position equals No\_Element, then No\_Element is returned.

```

function Next (Container : Set;
  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position = No_Element then Next'Result = No_Element
      elsif Next'Result = No_Element then
        Position = Last (Container)
      else Has_Element (Container, Next'Result));

```

Returns a cursor designating the successor of the node designated by Position in Container.

```

procedure Next (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;

```

Equivalent to Position := Next (Position).

```

procedure Next (Container : in Set;
  Position : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position /= No_Element
      then Has_Element (Container, Position));

```

Equivalent to Position := Next (Container, Position).

```

function Find (Container : Set;
  Item : Element_Type) return Cursor
  with Post => (if Find'Result /= No_Element
    then Has_Element (Container, Find'Result));

```

If Length (Container) equals 0, then Find returns No\_Element. Otherwise, Find checks if an element equivalent to Item is present in Container. If a match is found, a cursor designating the matching element is returned; otherwise, No\_Element is returned.

```

function Contains (Container : Set;
  Item : Element_Type) return Boolean;

```

Equivalent to Find (Container, Item) /= No\_Element.

```

procedure Iterate
  (Container : in Set;
  Process : not null access procedure (Position : in Cursor))
  with Allows_Exit;

```

Iterate calls Process.all with a cursor that designates each element in Container, starting with the first element and moving the cursor according to the successor relation. Tampering with the

cursors of Container is prohibited during the execution of a call on **Process.all**. Any exception raised by **Process.all** is propagated.

Both `Containers.Hash_Set` and `Containers.Ordered_Set` declare a nested generic package `Generic_Keys`, which provides operations that allow set manipulation in terms of a key (typically, a portion of an element) instead of a complete element. The formal function `Key` of `Generic_Keys` extracts a key value from an element. It is expected to return the same value each time it is called with a particular element. The behavior of `Generic_Keys` is unspecified if `Key` behaves in some other manner.

A key is expected to unambiguously determine a single equivalence class for elements. The behavior of `Generic_Keys` is unspecified if the formal parameters of this package behave in some other manner.

```
function Key (Position : Cursor) return Key_Type
with Pre => Position /= No_Element or else raise Constraint_Error,
      Global => in all;
```

Equivalent to `Key (Element (Position))`.

```
function Key (Container : Set;
              Position : Cursor) return Key_Type
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
      (Has_Element (Container, Position)
       or else raise Program_Error);
```

Equivalent to `Key (Element (Container, Position))`.

The subprograms in package `Generic_Keys` named `Contains`, `Find`, `Element`, `Delete`, and `Exclude`, are equivalent to the corresponding subprograms in the parent package, with the difference that the `Key` parameter is used to locate an element in the set.

```
procedure Replace (Container : in out Set;
                  Key       : in Key_Type;
                  New_Item  : in Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
           or else raise Program_Error,
      Post => Length (Container) = Length (Container)'Old;
```

Equivalent to `Replace_Element (Container, Find (Container, Key), New_Item)`.

```
procedure Update_Element_Preserving_Key
(Container : in out Set;
 Position  : in Cursor;
 Process   : not null access procedure
              (Element : in out Element_Type))
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
      (Has_Element (Container, Position)
       or else raise Program_Error);
```

`Update_Element_Preserving_Key` uses `Key` to save the key value *K* of the element designated by `Position`. `Update_Element_Preserving_Key` then calls **Process.all** with that element as the argument. Tampering with the cursors of `Container` is prohibited during the execution of the call on **Process.all**. Any exception raised by **Process.all** is propagated. After **Process.all** returns, `Update_Element_Preserving_Key` checks if *K* determines the same equivalence class as that for the new element; if not, the element is removed from the set and `Program_Error` is propagated.

If `Element_Type` is unconstrained and definite, then the actual `Element` parameter of **Process.all** shall be unconstrained.

```
type Reference_Type (Element : not null access Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);
```

The type `Reference_Type` needs finalization.

```

function Reference_Preserving_Key (Container : aliased in out Set;
                                   Position  : in Cursor)
return Reference_Type
with Pre  => (Position /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Position)
              or else raise Program_Error),
      Post => Tampering_With_Cursors_Prohibited (Container);

```

This function (combined with the Implicit\_Dereference aspect) provides a convenient way to gain read and write access to an individual element of a set given a cursor.

Reference\_Preserving\_Key uses Key to save the key value *K*; then returns an object whose discriminant is an access value that designates the element designated by Position. Tampering with the cursors of Container is prohibited while the object returned by Reference\_Preserving\_Key exists and has not been finalized. When the object returned by Reference\_Preserving\_Key is finalized, a check is made if *K* determines the same equivalence class as that for the new element; if not, the element is removed from the set and Program\_Error is propagated.

```

function Constant_Reference (Container : aliased in Set;
                             Key       : in Key_Type)
return Constant_Reference_Type
with Pre  => Find (Container, Key) /= No_Element
              or else raise Constraint_Error,
      Post => Tampering_With_Cursors_Prohibited (Container);

```

This function (combined with the Implicit\_Dereference aspect) provides a convenient way to gain read access to an individual element of a set given a key value.

Equivalent to Constant\_Reference (Container, Find (Container, Key)).

```

function Reference_Preserving_Key (Container : aliased in out Set;
                                   Key       : in Key_Type)
return Reference_Type
with Pre  => Find (Container, Key) /= No_Element
              or else raise Constraint_Error,
      Post => Tampering_With_Cursors_Prohibited (Container);

```

This function (combined with the Implicit\_Dereference aspect) provides a convenient way to gain read and write access to an individual element of a set given a key value.

Equivalent to Reference\_Preserving\_Key (Container, Find (Container, Key)).

The nested package Stable provides a type Stable.Set that represents a *stable* set, which is one that cannot grow and shrink. Such a set can be created by calling the Copy function, or by establishing a *stabilized view* of an ordinary set.

The subprograms of the set package that have a parameter or result of type Set are included in the nested package Stable with the same specification, except that the following are omitted:

Tampering\_With\_Cursors\_Prohibited, Assign, Move, Insert, Include, Clear, Delete, Exclude, Replace, Replace\_Element, procedures Union, Intersection, Difference, and Symmetric\_Difference, (for Ordered\_sets) Delete\_First and Delete\_Last, and (for Hashed\_sets) Reserve\_Capacity

The operations of this package are equivalent to those for ordinary sets, except that the calls to Tampering\_With\_Cursors\_Prohibited that occur in preconditions are replaced by False, and any that occur in postconditions are replaced by True.

If a stable set is declared with the Base discriminant designating a pre-existing ordinary set, the stable set represents a stabilized view of the underlying ordinary set, and any operation on the stable set is reflected on the underlying ordinary set. While a stabilized view exists, any operation that tampers with cursors performed on the underlying set is prohibited. The finalization of a stable set that provides such a view removes this restriction on the underlying ordinary set (though some other restriction can exist due to other concurrent iterations or stabilized views).

If a stable set is declared without specifying Base, the object is necessarily initialized. The initializing expression of the stable set, typically a call on Copy, determines the Length of the set. The Length of a stable set never changes after initialization.

*Bounded (Run-Time) Errors*

It is a bounded error for the actual function associated with a generic formal subprogram, when called as part of an operation of a set package, to tamper with elements of any set parameter of the operation. Either Program\_Error is raised, or the operation works as defined on the value of the set either prior to, or subsequent to, some or all of the modifications to the set.

It is a bounded error to call any subprogram declared in the visible part of a set package when the associated container has been finalized. If the operation takes Container as an **in out** parameter, then it raises Constraint\_Error or Program\_Error. Otherwise, the operation either proceeds as it would for an empty container, or it raises Constraint\_Error or Program\_Error.

*Erroneous Execution*

A Cursor value is *invalid* if any of the following have occurred since it was created:

- The set that contains the element it designates has been finalized;
- The set that contains the element it designates has been used as the Target of a call to Assign, or as the target of an assignment\_statement;
- The set that contains the element it designates has been used as the Source or Target of a call to Move; or
- The element it designates has been removed from the set that previously contained the element.

The result of "=" or Has\_Element is unspecified if these functions are called with an invalid cursor parameter. Execution is erroneous if any other subprogram declared in Containers.Hashed\_Sets or Containers.Ordered\_Sets is called with an invalid cursor parameter.

Execution is erroneous if the set associated with the result of a call to Reference or Constant\_Reference is finalized before the result object returned by the call to Reference or Constant\_Reference is finalized.

*Implementation Requirements*

No storage associated with a set object shall be lost upon assignment or scope exit.

The execution of an assignment\_statement for a set shall have the effect of copying the elements from the source set object to the target set object and changing the length of the target object to that of the source object.

*Implementation Advice*

Move should not copy elements, and should minimize copying of internal data structures.

If an exception is propagated from a set operation, no storage should be lost, nor any elements removed from a set unless specified by the operation.



## A.18.8 The Generic Package Containers.Hash\_Sets

*Static Semantics*

The generic library package Containers.Hash\_Sets has the following declaration:

```

with Ada.Iterator_Interfaces;
generic
  type Element_Type is private;
  with function Hash (Element : Element_Type) return Hash_Type;
  with function Equivalent_Elements (Left, Right : Element_Type)
    return Boolean;
  with function "=" (Left, Right : Element_Type) return Boolean is <>;
package Ada.Containers.Hash_Sets
  with Preelaborate, Remote_Types,
    Nonblocking, Global => in out synchronized is

  type Set is tagged private
    with Constant_Indexing => Constant_Reference,
      Default_Iterator    => Iterate,
      Iterator_Element    => Element_Type,
      Iterator_View       => Stable.Set,
      Aggregate           => (Empty      => Empty,
                             Add_Unnamed => Include),
      Stable_Properties  => (Length,
                             Tampering_With_Cursors_Prohibited),
      Default_Initial_Condition =>
        Length (Set) = 0 and then
          (not Tampering_With_Cursors_Prohibited (Set)),
      Preelaborable_Initialization;

  type Cursor is private
    with Preelaborable_Initialization;

  Empty_Set : constant Set;

  No_Element : constant Cursor;

  function Has_Element (Position : Cursor) return Boolean
    with Nonblocking, Global => in all, Use_Formal => null;

  function Has_Element (Container : Set; Position : Cursor)
    return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

  package Set_Iterator_Interfaces is new
    Ada.Iterator_Interfaces (Cursor, Has_Element);

  function "=" (Left, Right : Set) return Boolean;

  function Equivalent_Sets (Left, Right : Set) return Boolean;

  function Tampering_With_Cursors_Prohibited
    (Container : Set) return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

  function Empty (Capacity : Count_Type := implementation-defined)
    return Set
    with Post =>
      Capacity (Empty'Result) >= Capacity and then
        not Tampering_With_Cursors_Prohibited (Empty'Result) and then
          Length (Empty'Result) = 0;

  function To_Set (New_Item : Element_Type) return Set
    with Post => Length (To_Set'Result) = 1 and then
      not Tampering_With_Cursors_Prohibited (To_Set'Result);

  function Capacity (Container : Set) return Count_Type
    with Nonblocking, Global => null, Use_Formal => null;

```

```

procedure Reserve_Capacity (Container : in out Set;
                           Capacity : in Count_Type)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error,
        Post => Container.Capacity >= Capacity;

function Length (Container : Set) return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;

function Is_Empty (Container : Set) return Boolean
  with Nonblocking, Global => null, Use_Formal => null,
        Post => Is_Empty'Result = (Length (Container) = 0);

procedure Clear (Container : in out Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error,
        Post => Capacity (Container) = Capacity (Container)'old and then
        Length (Container) = 0;

function Element (Position : Cursor) return Element_Type
  with Pre => Position /= No_Element or else raise Constraint_Error,
        Nonblocking, Global => in all, Use_Formal => Element_Type;

function Element (Container : Set;
                  Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element
        or else raise Constraint_Error) and then
        (Has_Element (Container, Position)
        or else raise Program_Error),
        Nonblocking, Global => null, Use_Formal => Element_Type;

procedure Replace_Element (Container: in out Set;
                          Position : in Cursor;
                          New_item : in Element_Type)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error) and then
        (Position /= No_Element
        or else raise Constraint_Error) and then
        (Has_Element (Container, Position)
        or else raise Program_Error);

procedure Query_Element
  (Position : in Cursor;
   Process : not null access procedure (Element : in Element_Type))
  with Pre => Position /= No_Element or else raise Constraint_Error,
        Global => in all;

procedure Query_Element
  (Container : in Set;
   Position : in Cursor;
   Process : not null access procedure (Element : in Element_Type))
  with Pre => (Position /= No_Element
        or else raise Constraint_Error) and then
        (Has_Element (Container, Position)
        or else raise Program_Error);

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
  with Implicit_Dereference => Element,
        Nonblocking, Global => in out synchronized,
        Default_Initial_Condition => (raise Program_Error);

```

```

function Constant_Reference (Container : aliased in Set;
                             Position  : in Cursor)
    return Constant_Reference_Type
    with Pre  => (Position /= No_Element
                 or else raise Constraint_Error) and then
                 (Has_Element (Container, Position)
                  or else raise Program_Error),
    Post => Tampering_With_Cursors_Prohibited (Container),
    Nonblocking, Global => null, Use_Formal => null;

procedure Assign (Target : in out Set; Source : in Set)
    with Pre  => not Tampering_With_Cursors_Prohibited (Target)
                 or else raise Program_Error,
    Post => Length (Source) = Length (Target) and then
            Capacity (Target) >= Length (Source);

function Copy (Source : Set; Capacity : Count_Type := 0)
    return Set
    with Pre  => Capacity = 0 or else Capacity >= Length (Source)
                 or else raise Capacity_Error,
    Post =>
        Length (Copy'Result) = Length (Source) and then
        not Tampering_With_Cursors_Prohibited (Copy'Result) and then
        Copy'Result.Capacity = (if Capacity = 0 then
                                Length (Source) else Capacity);

procedure Move (Target : in out Set;
                Source : in out Set)
    with Pre  => (not Tampering_With_Cursors_Prohibited (Target)
                 or else raise Program_Error) and then
                 (not Tampering_With_Cursors_Prohibited (Source)
                  or else raise Program_Error),
    Post => (if not Target'Has_Same_Storage (Source) then
            Length (Target) = Length (Source'Old) and then
            Length (Source) = 0);

procedure Insert (Container : in out Set;
                  New_Item  : in Element_Type;
                  Position   : out Cursor;
                  Inserted   : out Boolean)
    with Pre  => (not Tampering_With_Cursors_Prohibited (Container)
                 or else raise Program_Error) and then
                 (Length (Container) <= Count_Type'Last - 1
                  or else raise Constraint_Error),
    Post => (declare
            Original_Length : constant Count_Type :=
                Length (Container)'Old;
            begin
                Has_Element (Container, Position) and then
                (if Inserted then
                    Length (Container) = Original_Length + 1
                else
                    Length (Container) = Original_Length) and then
                Capacity (Container) >= Length (Container);

procedure Insert (Container : in out Set;
                  New_Item  : in Element_Type)
    with Pre  => (not Tampering_With_Cursors_Prohibited (Container)
                 or else raise Program_Error) and then
                 (Length (Container) <= Count_Type'Last - 1
                  or else raise Constraint_Error),
    Post => Length (Container) = Length (Container)'Old + 1 and then
            Capacity (Container) >= Length (Container);

```

```

procedure Include (Container : in out Set;
    New_Item : in Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
    begin
        Length (Container)
            in Original_Length | Original_Length + 1) and then
        Capacity (Container) >= Length (Container);

procedure Replace (Container : in out Set;
    New_Item : in Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = Length (Container)'Old;

procedure Exclude (Container : in out Set;
    Item : in Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
    begin
        Length (Container) in
            Original_Length - 1 | Original_Length);

procedure Delete (Container : in out Set;
    Item : in Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = Length (Container)'Old - 1;

procedure Delete (Container : in out Set;
    Position : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error),
    Post => Length (Container) = Length (Container)'Old - 1 and then
    Position = No_Element;

procedure Union (Target : in out Set;
    Source : in Set)
with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
    Post => Length (Target) <= Length (Target)'Old + Length (Source);

function Union (Left, Right : Set) return Set
    with Post => Length (Union'Result) <=
        Length (Left) + Length (Right) and then
    not Tampering_With_Cursors_Prohibited (Union'Result);

function "or" (Left, Right : Set) return Set renames Union;

procedure Intersection (Target : in out Set;
    Source : in Set)
with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
    Post => Length (Target) <= Length (Target)'Old + Length (Source);

```

```

function Intersection (Left, Right : Set) return Set
  with Post =>
    Length (Intersection'Result) <=
      Length (Left) + Length (Right) and then
      not Tampering_With_Cursors_Prohibited (Intersection'Result);

function "and" (Left, Right : Set) return Set renames Intersection;

procedure Difference (Target : in out Set;
  Source : in Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
    Post => Length (Target) <= Length (Target)'Old + Length (Source);

function Difference (Left, Right : Set) return Set
  with Post =>
    Length (Difference'Result) <=
      Length (Left) + Length (Right) and then
      not Tampering_With_Cursors_Prohibited (Difference'Result);

function "-" (Left, Right : Set) return Set renames Difference;

procedure Symmetric_Difference (Target : in out Set;
  Source : in Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
    Post => Length (Target) <= Length (Target)'Old + Length (Source);

function Symmetric_Difference (Left, Right : Set) return Set
  with Post =>
    Length (Symmetric_Difference'Result) <=
      Length (Left) + Length (Right) and then
      not Tampering_With_Cursors_Prohibited (
        Symmetric_Difference'Result);

function "xor" (Left, Right : Set) return Set
  renames Symmetric_Difference;

function Overlap (Left, Right : Set) return Boolean;

function Is_Subset (Subset : Set;
  Of_Set : Set) return Boolean;

function First (Container : Set) return Cursor
  with Nonblocking, Global => null, Use_Format => null,
    Post => (if not Is_Empty (Container)
      then Has_Element (Container, First'Result)
      else First'Result = No_Element);

function Next (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Format => null,
    Post => (if Position = No_Element then Next'Result = No_Element);

function Next (Container : Set;
  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Format => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position = No_Element then Next'Result = No_Element
      elsif Next'Result = No_Element then
        Position = Last (Container)
      else Has_Element (Container, Next'Result));

procedure Next (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Format => null;

```

```

procedure Next (Container : in      Set;
                Position  : in out Cursor)
with Nonblocking, Global => null, Use_Formal => null,
    Pre  => Position = No_Element or else
        Has_Element (Container, Position)
    or else raise Program_Error,
    Post => (if Position /= No_Element
        then Has_Element (Container, Position));

function Find (Container : Set;
               Item       : Element_Type)
return Cursor
with Post => (if Find'Result /= No_Element
    then Has_Element (Container, Find'Result));

function Contains (Container : Set;
                  Item       : Element_Type) return Boolean;

function Equivalent_Elements (Left, Right : Cursor)
return Boolean
with Pre  => (Left /= No_Element and then Right /= No_Element)
    or else raise Constraint_Error,
    Global => in all;

function Equivalent_Elements (Left  : Cursor;
                             Right : Element_Type)
return Boolean
with Pre  => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

function Equivalent_Elements (Left  : Element_Type;
                             Right : Cursor)
return Boolean
with Pre  => Right /= No_Element or else raise Constraint_Error,
    Global => in all;

procedure Iterate
(Container : in Set;
 Process   : not null access procedure (Position : in Cursor))
with Allows_Exit;

function Iterate (Container : in Set)
return Set_Iterator_Interfaces.Parallel_Iterator'Class
with Post => Tampering_With_Cursors_Prohibited (Container);

generic
type Key_Type (<>) is private;
with function Key (Element : Element_Type) return Key_Type;
with function Hash (Key : Key_Type) return Hash_Type;
with function Equivalent_Keys (Left, Right : Key_Type)
return Boolean;

package Generic_Keys
with Nonblocking, Global => null is

    function Key (Position : Cursor) return Key_Type
        with Pre  => Position /= No_Element or else raise Constraint_Error,
        Global => in all;

    function Key (Container : Set;
                 Position  : Cursor) return Key_Type
        with Pre  => (Position = No_Element
            or else raise Constraint_Error) and then
            (Has_Element (Container, Position)
            or else raise Program_Error);

    function Element (Container : Set;
                     Key        : Key_Type)
return Element_Type;

```



```

procedure Replace (Container : in out Set;
                   Key       : in     Key_Type;
                   New_Item  : in     Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
           or else raise Program_Error,
      Post => Length (Container) = Length (Container)'Old;

procedure Exclude (Container : in out Set;
                   Key       : in     Key_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
           or else raise Program_Error,
      Post => (declare
                Original_Length : constant Count_Type :=
                    Length (Container)'Old;
              begin
                Length (Container)
                  in Original_Length - 1 | Original_Length);

procedure Delete (Container : in out Set;
                   Key       : in     Key_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
           or else raise Program_Error,
      Post => Length (Container) = Length (Container)'Old - 1;

function Find (Container : Set;
                Key       : Key_Type)
return Cursor
with Post => (if Find'Result = No_Element
              then Has_Element (Container, Find'Result));

function Contains (Container : Set;
                   Key       : Key_Type)
return Boolean;

procedure Update_Element_Preserving_Key
(Container : in out Set;
 Position  : in     Cursor;
 Process   : not null access procedure
              (Element : in out Element_Type))
with Pre => (Position /= No_Element or else
              raise Constraint_Error) and then
              (Has_Element (Container, Position) or else
              raise Program_Error);

type Reference_Type
  (Element : not null access Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

function Reference_Preserving_Key (Container : aliased in out Set;
                                Position  : in Cursor)
return Reference_Type
with Pre => (Position /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Position)
              or else raise Program_Error),
      Post => Tampering_With_Cursors_Prohibited (Container);

function Constant_Reference (Container : aliased in Set;
                             Key       : in Key_Type)
return Constant_Reference_Type
with Pre => Find (Container, Key) /= No_Element
              or else raise Constraint_Error,
      Post => Tampering_With_Cursors_Prohibited (Container);

```

```

function Reference_Preserving_Key (Container : aliased in out Set;
                                   Key       : in Key_Type)
    return Reference_Type
    with Pre  => Find (Container, Key) /= No_Element
        or else raise Constraint_Error,
        Post => Tampering_With_Cursors_Prohibited (Container);
end Generic_Keys;

package Stable is
    type Set (Base : not null access Hashed_Sets.Set) is
        tagged limited private
        with Constant_Indexing => Constant_Reference,
            Default_Iterator  => Iterate,
            Iterator_Element  => Element_Type,
            Stable_Properties => (Length),
            Global             => null,
            Default_Initial_Condition => Length (Set) = 0,
            Preelaborable_Initialization;

    type Cursor is private
        with Preelaborable_Initialization;

    Empty_Set : constant Set;

    No_Element : constant Cursor;

    function Has_Element (Position : Cursor) return Boolean
        with Nonblocking, Global => in all, Use_Formal => null;

    package Set_Iterator_Interfaces is new
        Ada.Iterator_Interfaces (Cursor, Has_Element);

    procedure Assign (Target : in out Hashed_Sets.Set;
                    Source : in Set)
        with Post => Length (Source) = Length (Target);

    function Copy (Source : Hashed_Sets.Set) return Set
        with Post => Length (Copy'Result) = Length (Source);

    type Constant_Reference_Type
        (Element : not null access constant Element_Type) is private
        with Implicit_Dereference => Element,
            Nonblocking, Global => null, Use_Formal => null,
            Default_Initial_Condition => (raise Program_Error);

    -- Additional subprograms as described in the text
    -- are declared here.

    private
        ... -- not specified by the language

    end Stable;

private
    ... -- not specified by the language

end Ada.Containers.Hashed_Sets;

```

An object of type *Set* contains an expandable hash table, which is used to provide direct access to elements. The *capacity* of an object of type *Set* is the maximum number of elements that can be inserted into the hash table prior to it being automatically expanded.

Two elements *E1* and *E2* are defined to be *equivalent* if *Equivalent\_Elements* (*E1*, *E2*) returns *True*.

The actual function for the generic formal function *Hash* is expected to return the same value each time it is called with a particular element value. For any two equivalent elements, the actual for *Hash* is expected to return the same value. If the actual for *Hash* behaves in some other manner, the behavior of this package is unspecified. Which subprograms of this package call *Hash*, and how many times they call it, is unspecified.

The actual function for the generic formal function `Equivalent_Elements` is expected to return the same value each time it is called with a particular pair of `Element` values. It should define an equivalence relationship, that is, be reflexive, symmetric, and transitive. If the actual for `Equivalent_Elements` behaves in some other manner, the behavior of this package is unspecified. Which subprograms of this package call `Equivalent_Elements`, and how many times they call it, is unspecified.

If the actual function for the generic formal function `"=`" returns `True` for any pair of nonequivalent elements, then the behavior of the container function `"=`" is unspecified.

If the value of an element stored in a set is changed other than by an operation in this package such that at least one of `Hash` or `Equivalent_Elements` give different results, the behavior of this package is unspecified.

Which elements are the first element and the last element of a set, and which element is the successor of a given element, are unspecified, other than the general semantics described in A.18.7.

```
function Empty (Capacity : Count_Type := implementation-defined)
return Set
with Post =>
    Capacity (Empty'Result) >= Capacity and then
    not Tampering_With_Cursors_Prohibited (Empty'Result) and then
    Length (Empty'Result) = 0;
```

Returns an empty set.

```
function Capacity (Container : Set) return Count_Type
with Nonblocking, Global => null, Use_Formal => null;
```

Returns the capacity of Container.

```
procedure Reserve_Capacity (Container : in out Set;
                           Capacity : in Count_Type)
with Pre  => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error,
    Post => Container.Capacity >= Capacity;
```

`Reserve_Capacity` allocates a new hash table such that the length of the resulting set can become at least the value `Capacity` without requiring an additional call to `Reserve_Capacity`, and is large enough to hold the current length of `Container`. `Reserve_Capacity` then rehashes the elements in `Container` onto the new hash table. It replaces the old hash table with the new hash table, and then deallocates the old hash table. Any exception raised during allocation is propagated and `Container` is not modified.

```
procedure Clear (Container : in out Set)
with Pre  => not Tampering_With_Cursors_Prohibited (Container)
        or else raise Program_Error,
    Post => Capacity (Container) = Capacity (Container)'Old and then
    Length (Container) = 0;
```

In addition to the semantics described in A.18.7, `Clear` does not affect the capacity of `Container`.

```
procedure Assign (Target : in out Set; Source : in Set)
with Pre  => not Tampering_With_Cursors_Prohibited (Target)
        or else raise Program_Error,
    Post => Length (Source) = Length (Target) and then
    Capacity (Target) >= Length (Source);
```

In addition to the semantics described in A.18.7, if the length of `Source` is greater than the capacity of `Target`, `Reserve_Capacity (Target, Length (Source))` is called before assigning any elements.

```

function Copy (Source : Set; Capacity : Count_Type := 0)
  return Set
  with Pre => Capacity = 0 or else Capacity >= Length (Source)
    or else raise Capacity_Error,
    Post =>
      Length (Copy'Result) = Length (Source) and then
      not Tampering_With_Cursors_Prohibited (Copy'Result) and then
      Copy'Result.Capacity = (if Capacity = 0 then
        Length (Source) else Capacity);

```

Returns a set whose elements are initialized from the elements of Source.

```

procedure Insert (Container : in out Set;
  New_Item : in Element_Type;
  Position : out Cursor;
  Inserted : out Boolean)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
    or else raise Constraint_Error),
    Post => (declare
      Original_Length : constant Count_Type :=
        Length (Container)'Old;
      begin
        Has_Element (Container, Position) and then
        (if Inserted then
          Length (Container) = Original_Length + 1
        else
          Length (Container) = Original_Length) and then
        Capacity (Container) >= Length (Container);

```

In addition to the semantics described in A.18.7, if Length (Container) equals Capacity (Container), then Insert first calls Reserve\_Capacity to increase the capacity of Container to some larger value.

```

function First (Container : Set) return Cursor;

```

If Length (Container) = 0, then First returns No\_Element. Otherwise, First returns a cursor that designates the first hashed element in Container.

```

function Equivalent_Elements (Left, Right : Cursor)
  return Boolean
  with Pre => (Left /= No_Element and then Right /= No_Element)
    or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Equivalent\_Elements (Element (Left), Element (Right)).

```

function Equivalent_Elements (Left : Cursor;
  Right : Element_Type) return Boolean
  with Pre => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Equivalent\_Elements (Element (Left), Right).

```

function Equivalent_Elements (Left : Element_Type;
  Right : Cursor) return Boolean
  with Pre => Right /= No_Element or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Equivalent\_Elements (Left, Element (Right)).

```

function Iterate (Container : in Set)
  return Set_Iterator_Interfaces.Parallel_Iterator'Class
  with Post => Tampering_With_Cursors_Prohibited (Container);

```

Iterate returns an iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each element in Container, starting with the first element and moving the cursor according to the successor relation when used as a forward iterator, and processing all nodes concurrently when used as a parallel iterator. Tampering with the cursors of Container is prohibited while the iterator object exists (in particular, in the `sequence_of_statements` of the `loop_statement` whose `iterator_specification` denotes this object). The iterator object needs finalization.

For any element  $E$ , the actual function for the generic formal function `Generic_Keys.Hash` is expected to be such that  $\text{Hash}(E) = \text{Generic\_Keys.Hash}(\text{Key}(E))$ . If the actuals for `Key` or `Generic_Keys.Hash` behave in some other manner, the behavior of `Generic_Keys` is unspecified. Which subprograms of `Generic_Keys` call `Generic_Keys.Hash`, and how many times they call it, is unspecified.

For any two elements  $E1$  and  $E2$ , the boolean values `Equivalent_Elements(E1, E2)` and `Equivalent_Keys(Key(E1), Key(E2))` are expected to be equal. If the actuals for `Key` or `Equivalent_Keys` behave in some other manner, the behavior of `Generic_Keys` is unspecified. Which subprograms of `Generic_Keys` call `Equivalent_Keys`, and how many times they call it, is unspecified.

#### Implementation Advice

If  $N$  is the length of a set, the average time complexity of the subprograms `Insert`, `Include`, `Replace`, `Delete`, `Exclude`, and `Find` that take an element parameter should be  $O(\log N)$ . The average time complexity of the subprograms that take a cursor parameter should be  $O(1)$ . The average time complexity of `Reserve_Capacity` should be  $O(N)$ .

### A.18.9 The Generic Package Containers.Ordered\_Sets

#### Static Semantics

The generic library package `Containers.Ordered_Sets` has the following declaration:

```

with Ada.Iterator_Interfaces;
generic
  type Element_Type is private;
  with function "<" (Left, Right : Element_Type) return Boolean is <>;
  with function "=" (Left, Right : Element_Type) return Boolean is <>;
package Ada.Containers.Ordered_Sets
  with Preelaborate, Remote_Types,
    Nonblocking, Global => in out synchronized is

  function Equivalent_Elements (Left, Right : Element_Type) return Boolean;

  type Set is tagged private
    with Constant_Indexing => Constant_Reference,
      Default_Iterator    => Iterate,
      Iterator_Element    => Element_Type,
      Iterator_View       => Stable.Set,
      Aggregate           => (Empty      => Empty,
                             Add_Unnamed => Include),
      Stable_Properties   => (Length,
                             Tampering_With_Cursors_Prohibited),
      Default_Initial_Condition =>
        Length (Set) = 0 and then
          (not Tampering_With_Cursors_Prohibited (Set)),
      Preelaborable_Initialization;

  type Cursor is private
    with Preelaborable_Initialization;

  Empty_Set : constant Set;
  No_Element : constant Cursor;

```

```

function Has_Element (Position : Cursor) return Boolean
  with Nonblocking, Global => in all, Use_Formal => null;

function Has_Element (Container : Set; Position : Cursor)
  return Boolean
  with Nonblocking, Global => null, Use_Formal => null;

package Set_Iterator_Interfaces is new
  Ada.Iterator_Interfaces (Cursor, Has_Element);

function "=" (Left, Right : Set) return Boolean;

function Equivalent_Sets (Left, Right : Set) return Boolean;

function Tampering_With_Cursors_Prohibited
  (Container : Set) return Boolean
  with Nonblocking, Global => null, Use_Formal => null;

function Empty return Set
  is (Empty_Set)
  with Post =>
    not Tampering_With_Cursors_Prohibited (Empty'Result) and then
      Length (Empty'Result) = 0;

function To_Set (New_Item : Element_Type) return Set
  with Post => Length (To_Set'Result) = 1 and then
    not Tampering_With_Cursors_Prohibited (To_Set'Result);

function Length (Container : Set) return Count_Type
  with Nonblocking, Global => null, Use_Formal => null;

function Is_Empty (Container : Set) return Boolean
  with Nonblocking, Global => null, Use_Formal => null,
  Post => Is_Empty'Result = (Length (Container) = 0);

procedure Clear (Container : in out Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
  Post => Length (Container) = 0;

function Element (Position : Cursor) return Element_Type
  with Pre => Position /= No_Element or else raise Constraint_Error,
  Nonblocking, Global => in all, Use_Formal => Element_Type;

function Element (Container : Set;
  Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error),
  Nonblocking, Global => null, Use_Formal => Element_Type;

procedure Replace_Element (Container : in out Set;
  Position : in Cursor;
  New_item : in Element_Type)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error);

procedure Query_Element
  (Position : in Cursor;
  Process : not null access procedure (Element : in Element_Type))
  with Pre => Position /= No_Element
    or else raise Constraint_Error,
  Global => in all;

```



```

procedure Query_Element
  (Container : in Set;
   Position  : in Cursor;
   Process   : not null access procedure (Element : in Element_Type))
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error);

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
with Implicit_Dereference => Element,
     Nonblocking, Global => in out synchronized,
     Default_Initial_Condition => (raise Program_Error);

function Constant_Reference (Container : aliased in Set;
                             Position  : in Cursor)
return Constant_Reference_Type
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error),
  Post => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

procedure Assign (Target : in out Set; Source : in Set)
with Pre => not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error,
  Post => Length (Source) = Length (Target);

function Copy (Source : Set) return Set
with Post => Length (Copy'Result) = Length (Source) and then
             not Tampering_With_Cursors_Prohibited (Copy'Result);

procedure Move (Target : in out Set;
                Source : in out Set)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
  (not Tampering_With_Cursors_Prohibited (Source)
   or else raise Program_Error),
  Post => (if not Target'Has_Same_Storage (Source) then
           Length (Target) = Length (Source'Old) and then
           Length (Source) = 0);

procedure Insert (Container : in out Set;
                  New_Item  : in Element_Type;
                  Position   : out Cursor;
                  Inserted   : out Boolean)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Length (Container) <= Count_Type'Last - 1
   or else raise Constraint_Error),
  Post => (declare
           Original_Length : constant Count_Type :=
             Length (Container)'Old;
           begin
             Has_Element (Container, Position) and then
             (if Inserted then
              Length (Container) = Original_Length + 1
             else
              Length (Container) = Original_Length));

```

```

procedure Insert (Container : in out Set;
                  New_Item  : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
     or else raise Constraint_Error),
    Post => Length (Container) = Length (Container)'Old + 1;

procedure Include (Container : in out Set;
                  New_Item  : in      Element_Type)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
    (Length (Container) <= Count_Type'Last - 1
     or else raise Constraint_Error),
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
        begin
            Length (Container)
                in Original_Length | Original_Length + 1);

procedure Replace (Container : in out Set;
                  New_Item  : in      Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error,
    Post => Length (Container) = Length (Container)'Old;

procedure Exclude (Container : in out Set;
                  Item       : in      Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error,
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
        begin
            Length (Container)
                in Original_Length - 1 | Original_Length);

procedure Delete (Container : in out Set;
                  Item       : in      Element_Type)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error,
    Post => Length (Container) = Length (Container)'Old - 1;

procedure Delete (Container : in out Set;
                  Position  : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
    (Position /= No_Element
     or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
     or else raise Program_Error),
    Post => Length (Container) = Length (Container)'Old - 1 and then
    Position = No_Element;

procedure Delete_First (Container : in out Set)
with Pre => not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error,
    Post => (declare
        Original_Length : constant Count_Type :=
            Length (Container)'Old;
        begin
            (if Original_Length = 0 then Length (Container) = 0
             else Length (Container) = Original_Length - 1));

```

```

procedure Delete_Last (Container : in out Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
  Post => (declare
    Original_Length : constant Count_Type :=
      Length (Container)'Old;
    begin
      (if Original_Length = 0 then Length (Container) = 0
      else Length (Container) = Original_Length - 1));

procedure Union (Target : in out Set;
  Source : in Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
  Post => Length (Target) <= Length (Target)'Old + Length (Source);

function Union (Left, Right : Set) return Set
  with Post => Length (Union'Result) <=
    Length (Left) + Length (Right) and then
    not Tampering_With_Cursors_Prohibited (Union'Result);

function "or" (Left, Right : Set) return Set renames Union;

procedure Intersection (Target : in out Set;
  Source : in Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
  Post => Length (Target) <= Length (Target)'Old + Length (Source);

function Intersection (Left, Right : Set) return Set
  with Post =>
    Length (Intersection'Result) <=
      Length (Left) + Length (Right) and then
      not Tampering_With_Cursors_Prohibited (Intersection'Result);

function "and" (Left, Right : Set) return Set renames Intersection;

procedure Difference (Target : in out Set;
  Source : in Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
  Post => Length (Target) <= Length (Target)'Old + Length (Source);

function Difference (Left, Right : Set) return Set
  with Post =>
    Length (Difference'Result) <=
      Length (Left) + Length (Right) and then
      not Tampering_With_Cursors_Prohibited (Difference'Result);

function "-" (Left, Right : Set) return Set renames Difference;

procedure Symmetric_Difference (Target : in out Set;
  Source : in Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Target)
    or else raise Program_Error,
  Post => Length (Target) <= Length (Target)'Old + Length (Source);

function Symmetric_Difference (Left, Right : Set) return Set
  with Post =>
    Length (Symmetric_Difference'Result) <=
      Length (Left) + Length (Right) and then
      not Tampering_With_Cursors_Prohibited (
        Symmetric_Difference'Result);

function "xor" (Left, Right : Set) return Set renames
  Symmetric_Difference;

function Overlap (Left, Right : Set) return Boolean;

function Is_Subset (Subset : Set;
  Of_Set : Set) return Boolean;

```

```

function First (Container : Set) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container)
      then Has_Element (Container, First'Result)
      else First'Result = No_Element);

function First_Element (Container : Set)
  return Element_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

function Last (Container : Set) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Post => (if not Is_Empty (Container) then
      Has_Element (Container, Last'Result)
      else Last'Result = No_Element);

function Last_Element (Container : Set)
  return Element_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

function Next (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then Next'Result = No_Element);

function Next (Container : Set;
  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position = No_Element then Next'Result = No_Element
      elsif Next'Result = No_Element then
        Position = Last (Container)
      else Has_Element (Container, Next'Result));

procedure Next (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;

procedure Next (Container : in Set;
  Position : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position /= No_Element
      then Has_Element (Container, Position));

function Previous (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element then
      Previous'Result = No_Element);

function Previous (Container : Set;
  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
      Has_Element (Container, Position)
      or else raise Program_Error,
    Post => (if Position = No_Element then
      Previous'Result = No_Element
      elsif Previous'Result = No_Element then
        Position = First (Container)
      else Has_Element (Container, Previous'Result));

procedure Previous (Position : in out Cursor)
  with Nonblocking, Global => in all,
    Use_Formal => null;

```

```

procedure Previous (Container : in      Set;
                    Position  : in out Cursor)
with Nonblocking, Global => null, Use_Formal => null,
    Pre  => Position = No_Element or else
        Has_Element (Container, Position)
    or else raise Program_Error,
    Post => (if Position /= No_Element
        then Has_Element (Container, Position));

function Find (Container : Set;
               Item       : Element_Type) return Cursor
with Post => (if Find'Result /= No_Element
    then Has_Element (Container, Find'Result));

function Floor (Container : Set;
               Item       : Element_Type) return Cursor
with Post => (if Floor'Result /= No_Element
    then Has_Element (Container, Floor'Result));

function Ceiling (Container : Set;
               Item       : Element_Type) return Cursor
with Post => (if Ceiling'Result /= No_Element
    then Has_Element (Container, Ceiling'Result));

function Contains (Container : Set;
                  Item       : Element_Type) return Boolean;

function "<" (Left, Right : Cursor) return Boolean
with Pre  => (Left /= No_Element and then Right /= No_Element)
    or else raise Constraint_Error,
    Global => in all;

function ">" (Left, Right : Cursor) return Boolean
with Pre  => (Left /= No_Element and then Right /= No_Element)
    or else raise Constraint_Error,
    Global => in all;

function "<" (Left : Cursor; Right : Element_Type) return Boolean
with Pre  => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

function ">" (Left : Cursor; Right : Element_Type) return Boolean
with Pre  => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

function "<" (Left : Element_Type; Right : Cursor) return Boolean
with Pre  => Right /= No_Element or else raise Constraint_Error,
    Global => in all;

function ">" (Left : Element_Type; Right : Cursor) return Boolean
with Pre  => Right /= No_Element or else raise Constraint_Error,
    Global => in all;

procedure Iterate
(Container : in Set;
 Process  : not null access procedure (Position : in Cursor))
with Allows_Exit;

procedure Reverse_Iterate
(Container : in Set;
 Process  : not null access procedure (Position : in Cursor))
with Allows_Exit;

function Iterate (Container : in Set)
return Set_Iterator_Interfaces.Parallel_Reversible_Iterator'Class
with Post => Tampering_With_Cursors_Prohibited (Container);

```

```

function Iterate (Container : in Set; Start : in Cursor)
  return Set_Iterator_Interfaces.Reversible_Iterator'Class
  with Pre => (Start /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Start)
    or else raise Program_Error),
    Post => Tampering_With_Cursors_Prohibited (Container);

generic
  type Key_Type (<>) is private;
  with function Key (Element : Element_Type) return Key_Type;
  with function "<" (Left, Right : Key_Type)
    return Boolean is <>;
package Generic_Keys
with Nonblocking, Global => null is

  function Equivalent_Keys (Left, Right : Key_Type)
    return Boolean;

  function Key (Position : Cursor) return Key_Type
    with Pre => Position /= No_Element or else raise Constraint_Error,
    Global => in all;

  function Key (Container : Set;
    Position : Cursor) return Key_Type
    with Pre => (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error);

  function Element (Container : Set;
    Key : Key_Type)
    return Element_Type;

  procedure Replace (Container : in out Set;
    Key : in Key_Type;
    New_Item : in Element_Type)
    with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = Length (Container)'Old;

  procedure Exclude (Container : in out Set;
    Key : in Key_Type)
    with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => (declare
      Original_Length : constant Count_Type :=
        Length (Container)'Old;
    begin
      Length (Container) in
        Original_Length - 1 | Original_Length);

  procedure Delete (Container : in out Set;
    Key : in Key_Type)
    with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Length (Container) = Length (Container)'Old - 1;

  function Find (Container : Set;
    Key : Key_Type) return Cursor
    with Post => (if Find'Result /= No_Element
    then Has_Element (Container, Find'Result));

  function Floor (Container : Set;
    Key : Key_Type) return Cursor
    with Post => (if Floor'Result /= No_Element
    then Has_Element (Container, Floor'Result));

```



```

function Ceiling (Container : Set;
                  Key       : Key_Type) return Cursor
  with Post => (if Ceiling'Result /= No_Element
              then Has_Element (Container, Ceiling'Result));

function Contains (Container : Set;
                  Key       : Key_Type) return Boolean;

procedure Update_Element_Preserving_Key
(Container : in out Set;
 Position  : in      Cursor;
 Process   : not null access procedure
            (Element : in out Element_Type))
  with Pre  => (Position /= No_Element
              or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
     or else raise Program_Error);

type Reference_Type
  (Element : not null access Element_Type) is private
  with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

function Reference_Preserving_Key (Container : aliased in out Set;
                                Position  : in Cursor)

  return Reference_Type
  with Pre  => (Position /= No_Element
              or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
     or else raise Program_Error),
    Post => Tampering_With_Cursors_Prohibited (Container);;

function Constant_Reference (Container : aliased in Set;
                            Key       : in Key_Type)

  return Constant_Reference_Type
  with Pre  => Find (Container, Key) /= No_Element
              or else raise Constraint_Error,
    Post => Tampering_With_Cursors_Prohibited (Container);;

function Reference_Preserving_Key (Container : aliased in out Set;
                                Key       : in Key_Type)

  return Reference_Type
  with Pre  => Find (Container, Key) /= No_Element
              or else raise Constraint_Error,
    Post => Tampering_With_Cursors_Prohibited (Container);;

end Generic_Keys;

package Stable is

  type Set (Base : not null access Ordered_Sets.Set) is
    tagged limited private
    with Constant_Indexing => Constant_Reference,
        Default_Iterator  => Iterate,
        Iterator_Element  => Element_Type,
        Stable_Properties => (Length),
        Global            => null,
        Default_Initial_Condition => Length (Set) = 0,
        Preelaborable_Initialization;

  type Cursor is private
    with Preelaborable_Initialization;

  Empty_Set : constant Set;

  No_Element : constant Cursor;

  function Has_Element (Position : Cursor) return Boolean
    with Nonblocking, Global => in all, Use_Formal => null;

```

```

package Set_Iterator_Interfaces is new
  Ada.Iterator_Interfaces (Cursor, Has_Element);

procedure Assign (Target : in out Ordered_Sets.Set;
                  Source : in Set)
  with Post => Length (Source) = Length (Target);

function Copy (Source : Ordered_Sets.Set) return Set
  with Post => Length (Copy'Result) = Length (Source);

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
  with Implicit_Dereference => Element,
        Nonblocking, Global => null, Use_Formal => null,
        Default_Initial_Condition => (raise Program_Error);

-- Additional subprograms as described in the text
-- are declared here.

private
  ... -- not specified by the language

end Stable;

private
  ... -- not specified by the language

end Ada.Containers.Ordered_Sets;

```

Two elements  $E1$  and  $E2$  are *equivalent* if both  $E1 < E2$  and  $E2 < E1$  return False, using the generic formal "<" operator for elements. Function Equivalent\_Elements returns True if Left and Right are equivalent, and False otherwise.

The actual function for the generic formal function "<" on Element\_Type values is expected to return the same value each time it is called with a particular pair of key values. It should define a strict weak ordering relationship (see A.18). If the actual for "<" behaves in some other manner, the behavior of this package is unspecified. Which subprograms of this package call "<" and how many times they call it, is unspecified.

If the actual function for the generic formal function "=" returns True for any pair of nonequivalent elements, then the behavior of the container function "=" is unspecified.

If the value of an element stored in a set is changed other than by an operation in this package such that at least one of "<" or "=" give different results, the behavior of this package is unspecified.

The *first element* of a nonempty set is the one which is less than all the other elements in the set. The *last element* of a nonempty set is the one which is greater than all the other elements in the set. The *successor* of an element is the smallest element that is larger than the given element. The *predecessor* of an element is the largest element that is smaller than the given element. All comparisons are done using the generic formal "<" operator for elements.

```

function Copy (Source : Set) return Set
  with Post => Length (Copy'Result) = Length (Source) and then
        not Tampering_With_Cursors_Prohibited (Copy'Result);

```

Returns a set whose elements are initialized from the corresponding elements of Source.

```

procedure Delete_First (Container : in out Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
  Post => (declare
    Original_Length : constant Count_Type :=
      Length (Container)'Old;
  begin
    (if Original_Length = 0 then Length (Container) = 0
    else Length (Container) = Original_Length - 1));

```

If Container is empty, Delete\_First has no effect. Otherwise, the element designated by First (Container) is removed from Container. Delete\_First tampers with the cursors of Container.

```

procedure Delete_Last (Container : in out Set)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
  Post => (declare
    Original_Length : constant Count_Type :=
      Length (Container)'Old;
  begin
    (if Original_Length = 0 then Length (Container) = 0
    else Length (Container) = Original_Length - 1));

```

If Container is empty, Delete\_Last has no effect. Otherwise, the element designated by Last (Container) is removed from Container. Delete\_Last tampers with the cursors of Container.

```

function First_Element (Container : Set) return Element_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

```

Equivalent to Element (First (Container)).

```

function Last (Container : Set) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
  Post => (if not Is_Empty (Container) then
    Has_Element (Container, Last'Result)
  else Last'Result = No_Element);

```

Returns a cursor that designates the last element in Container. If Container is empty, returns No\_Element.

```

function Last_Element (Container : Set) return Element_Type
  with Pre => (not Is_Empty (Container)
    or else raise Constraint_Error);

```

Equivalent to Element (Last (Container)).

```

function Previous (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
  Post => (if Position = No_Element then
    Previous'Result = No_Element);

```

If Position equals No\_Element, then Previous returns No\_Element. Otherwise, Previous returns a cursor designating the predecessor element of the one designated by Position. If Position designates the first element, then Previous returns No\_Element.

```

function Previous (Container : Set;
                  Position : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Position = No_Element then
        Previous'Result = No_Element
        elsif Previous'Result = No_Element then
            Position = First (Container)
        else Has_Element (Container, Previous'Result));

```

Returns a cursor designating the predecessor of the node designated by Position in Container, if any.

```

procedure Previous (Position : in out Cursor)
with Nonblocking, Global => in all, Use_Formal => null;

```

Equivalent to Position := Previous (Position).

```

procedure Previous (Container : in Set;
                  Position : in out Cursor)
with Nonblocking, Global => null, Use_Formal => null,
    Pre => Position = No_Element or else
        Has_Element (Container, Position)
        or else raise Program_Error,
    Post => (if Position /= No_Element
        then Has_Element (Container, Position));

```

Equivalent to Position := Previous (Container, Position).

```

function Floor (Container : Set;
                Item : Element_Type) return Cursor
with Post => (if Floor'Result /= No_Element
        then Has_Element (Container, Floor'Result));

```

Floor searches for the last element which is not greater than Item. If such an element is found, a cursor that designates it is returned. Otherwise, No\_Element is returned.

```

function Ceiling (Container : Set;
                 Item : Element_Type) return Cursor
with Post => (if Ceiling'Result /= No_Element
        then Has_Element (Container, Ceiling'Result));

```

Ceiling searches for the first element which is not less than Item. If such an element is found, a cursor that designates it is returned. Otherwise, No\_Element is returned.

```

function "<" (Left, Right : Cursor) return Boolean
with Pre => (Left /= No_Element and then Right /= No_Element)
        or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Element (Left) < Element (Right).

```

function ">" (Left, Right : Cursor) return Boolean
with Pre => (Left /= No_Element and then Right /= No_Element)
        or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Element (Right) < Element (Left).

```

function "<" (Left : Cursor; Right : Element_Type) return Boolean
with Pre => Left /= No_Element or else raise Constraint_Error,
    Global => in all;

```

Equivalent to Element (Left) < Right.

```

function ">" (Left : Cursor; Right : Element_Type) return Boolean
  with Pre    => Left /= No_Element or else raise Constraint_Error,
        Global => in all;

```

Equivalent to Right < Element (Left).

```

function "<" (Left : Element_Type; Right : Cursor) return Boolean
  with Pre    => Right /= No_Element or else raise Constraint_Error,
        Global => in all;

```

Equivalent to Left < Element (Right).

```

function ">" (Left : Element_Type; Right : Cursor) return Boolean
  with Pre    => Right /= No_Element or else raise Constraint_Error,
        Global => in all;

```

Equivalent to Element (Right) < Left.

```

procedure Reverse_Iterate
(Container : in Set;
 Process   : not null access procedure (Position : in Cursor))
with Allows_Exit;

```

Iterates over the elements in Container as per procedure Iterate, with the difference that the elements are traversed in predecessor order, starting with the last element.

```

function Iterate (Container : in Set)
return Set_Iterator_Interfaces.Parallel_Reversible_Iterator'Class
with Post => Tampering_With_Cursors_Prohibited (Container);

```

Iterate returns an iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each element in Container, starting with the first element and moving the cursor according to the successor relation when used as a forward iterator, and starting with the last element and moving the cursor according to the predecessor relation when used as a reverse iterator, and processing all nodes concurrently when used as a parallel iterator. Tampering with the cursors of Container is prohibited while the iterator object exists (in particular, in the sequence\_of\_statements of the loop\_statement whose iterator\_specification denotes this object). The iterator object needs finalization.

```

function Iterate (Container : in Set; Start : in Cursor)
return Set_Iterator_Interfaces.Reversible_Iterator'Class
with Pre  => (Start /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Start)
               or else raise Program_Error),
        Post => Tampering_With_Cursors_Prohibited (Container);

```

Iterate returns a reversible iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each element in Container, starting with the element designated by Start and moving the cursor according to the successor relation when used as a forward iterator, or moving the cursor according to the predecessor relation when used as a reverse iterator. Tampering with the cursors of Container is prohibited while the iterator object exists (in particular, in the sequence\_of\_statements of the loop\_statement whose iterator\_specification denotes this object). The iterator object needs finalization.

For any two elements *E1* and *E2*, the boolean values (*E1* < *E2*) and (Key(*E1*) < Key(*E2*)) are expected to be equal. If the actuals for Key or Generic\_Keys."<" behave in some other manner, the behavior of this package is unspecified. Which subprograms of this package call Key and Generic\_Keys."<", and how many times the functions are called, is unspecified.

In addition to the semantics described in A.18.7, the subprograms in package Generic\_Keys named Floor and Ceiling, are equivalent to the corresponding subprograms in the parent package, with the difference that the Key subprogram parameter is compared to elements in the container using the Key and "<" generic formal functions. The function named Equivalent\_Keys in package Generic\_Keys returns

True if both  $Left < Right$  and  $Right < Left$  return False using the generic formal "<" operator, and returns True otherwise.

*Implementation Advice*

If  $N$  is the length of a set, then the worst-case time complexity of the Insert, Include, Replace, Delete, Exclude, and Find operations that take an element parameter should be  $O((\log N)^2)$  or better. The worst-case time complexity of the subprograms that take a cursor parameter should be  $O(1)$ .

### A.18.10 The Generic Package Containers.Multiway\_Trees

The language-defined generic package Containers.Multiway\_Trees provides private types Tree and Cursor, and a set of operations for each type. A multiway tree container is well-suited to represent nested structures.

A multiway tree container object manages a tree of *nodes*, consisting of a *root node* and a set of *internal nodes*; each internal node contains an element and pointers to the parent, first child, last child, next (successor) sibling, and previous (predecessor) sibling internal nodes. A cursor designates a particular node within a tree (and by extension the element contained in that node, if any). A cursor keeps designating the same node (and element) as long as the node is part of the container, even if the node is moved within the container.

A *subtree* is a particular node (which *roots the subtree*) and all of its child nodes (including all of the children of the child nodes, recursively). The root node is always present and has neither an associated element value nor any parent node; it has pointers to its first child and its last child, if any. The root node provides a place to add nodes to an otherwise empty tree and represents the base of the tree.

A node that has no children is called a *leaf node*. The *ancestors* of a node are the node itself, its parent node, the parent of the parent node, and so on until a node with no parent is reached. Similarly, the *descendants* of a node are the node itself, its child nodes, the children of each child node, and so on.

The nodes of a subtree can be visited in several different orders. For a *depth-first order*, after visiting a node, the nodes of its child list are each visited in depth-first order, with each child node visited in natural order (first child to last child).

*Static Semantics*

The generic library package Containers.Multiway\_Trees has the following declaration:

```
with Ada.Iterator_Interfaces;
generic
  type Element_Type is private;
  with function "=" (Left, Right : Element_Type) return Boolean is <>;
package Ada.Containers.Multiway_Trees
  with Preelaborate, Remote_Types,
  Nonblocking, Global => in out synchronized is
  type Tree is tagged private
    with Constant_Indexing => Constant_Reference,
         Variable_Indexing => Reference,
         Default_Iterator   => Iterate,
         Iterator_Element   => Element_Type,
         Iterator_View       => Stable.Tree,
         Stable_Properties => (Node_Count,
                               Tampering_With_Cursors_Prohibited,
                               Tampering_With_Elements_Prohibited),
         Default_Initial_Condition =>
           Node_Count (Tree) = 1 and then
             (not Tampering_With_Cursors_Prohibited (Tree)) and then
             (not Tampering_With_Elements_Prohibited (Tree)),
         Preelaborable_Initialization;
  type Cursor is private
    with Preelaborable_Initialization;
  Empty_Tree : constant Tree;
```



```

No_Element : constant Cursor;

function Equal_Element (Left, Right : Element_Type)
    return Boolean renames "=";

function Has_Element (Position : Cursor) return Boolean
    with Nonblocking, Global => in all, Use_Formal => null;

function Has_Element (Container : Tree; Position : Cursor)
    return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

package Tree_Iterator_Interfaces is new
    Ada.Iterator_Interfaces (Cursor, Has_Element);

function Equal_Subtree (Left_Position : Cursor;
    Right_Position: Cursor) return Boolean;

function "=" (Left, Right : Tree) return Boolean;

function Tampering_With_Cursors_Prohibited
    (Container : Tree) return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

function Tampering_With_Elements_Prohibited
    (Container : Tree) return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

function Empty return Tree
    is (Empty_Tree)
    with Post =>
        not Tampering_With_Elements_Prohibited (Empty'Result) and then
        not Tampering_With_Cursors_Prohibited (Empty'Result) and then
        Node_Count (Empty'Result) = 1;

function Is_Empty (Container : Tree) return Boolean
    with Nonblocking, Global => null, Use_Formal => null,
    Post => Is_Empty'Result = (Node_Count (Container) = 1);

function Node_Count (Container : Tree) return Count_Type
    with Nonblocking, Global => null, Use_Formal => null;

function Subtree_Node_Count (Position : Cursor) return Count_Type
    with Nonblocking, Global => in all, Use_Formal => null;

function Subtree_Node_Count (Container : Tree; Position : Cursor)
    return Count_Type
    with Pre => Meaningful_For (Container, Position)
        or else raise Program_Error,
        Nonblocking, Global => null, Use_Formal => null;

function Depth (Position : Cursor) return Count_Type
    with Nonblocking, Global => in all, Use_Formal => null;

function Depth (Container : Tree; Position : Cursor)
    return Count_Type
    with Pre => Meaningful_For (Container, Position)
        or else raise Program_Error,
        Nonblocking, Global => null, Use_Formal => null;

function Is_Root (Position : Cursor) return Boolean
    with Nonblocking, Global => in all, Use_Formal => null;

function Is_Root (Container : Tree; Position : Cursor)
    return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

function Is_Leaf (Position : Cursor) return Boolean
    with Nonblocking, Global => in all, Use_Formal => null;

```

```

function Is_Leaf (Container : Tree; Position : Cursor)
  return Boolean
  with Pre => Meaningful_For (Container, Position)
    or else raise Program_Error,
    Nonblocking, Global => null, Use_Format => null;

function Is_Ancessor_Of (Container : Tree;
  Parent : Cursor;
  Position : Cursor) return Boolean
  with Pre => (Meaningful_For (Container, Position)
    or else raise Program_Error) and then
    (Meaningful_For (Container, Parent)
    or else raise Program_Error),
    Nonblocking, Global => null, Use_Format => null;

function Root (Container : Tree) return Cursor
  with Nonblocking, Global => null, Use_Format => null,
  Post => Root'Result /= No_Element and then
    not Has_Element (Container, Root'Result);

function Meaningful_For (Container : Tree; Position : Cursor)
  return Boolean is
    (Position = No_Element or else
    Is_Root (Container, Position) or else
    Has_Element (Container, Position))
  with Nonblocking, Global => null, Use_Format => null;

procedure Clear (Container : in out Tree)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
  Post => Node_Count (Container) = 1;

function Element (Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element or else
    raise Constraint_Error) and then
    (Has_Element (Position) or else raise Program_Error),
  Nonblocking, Global => in all, Use_Format => Element_Type;

function Element (Container : Tree;
  Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element or else
    raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error),
  Nonblocking, Global => null, Use_Format => Element_Type;

procedure Replace_Element (Container : in out Tree;
  Position : in Cursor;
  New_item : in Element_Type)
  with Pre => (not Tampering_With_Elements_Prohibited (Container)
    or else raise Program_Error) and then
    (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error);

procedure Query_Element
  (Position : in Cursor;
  Process : not null access procedure (Element : in Element_Type))
  with Pre => (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Position) or else raise Program_Error),
  Global => in all;

```

```

procedure Query_Element
(Container : in Tree;
Position  : in Cursor;
Process   : not null access procedure (Element : in Element_Type))
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
             (Has_Element (Container, Position)
             or else raise Program_Error);

procedure Update_Element
(Container : in out Tree;
Position  : in Cursor;
Process   : not null access procedure
            (Element : in out Element_Type))
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
             (Has_Element (Container, Position)
             or else raise Program_Error);

type Constant_Reference_Type
(Element : not null access constant Element_Type) is private
with Implicit_Dereference => Element,
     Nonblocking, Global => in out synchronized,
     Default_Initial_Condition => (raise Program_Error);

type Reference_Type (Element : not null access Element_Type) is private
with Implicit_Dereference => Element,
     Nonblocking, Global => in out synchronized,
     Default_Initial_Condition => (raise Program_Error);

function Constant_Reference (Container : aliased in Tree;
                             Position  : in Cursor)
return Constant_Reference_Type
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
             (Has_Element (Container, Position)
             or else raise Program_Error),
     Post => Tampering_With_Cursors_Prohibited (Container),
     Nonblocking, Global => null, Use_Formal => null;

function Reference (Container : aliased in out Tree;
                    Position  : in Cursor)
return Reference_Type
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
             (Has_Element (Container, Position)
             or else raise Program_Error),
     Post => Tampering_With_Cursors_Prohibited (Container),
     Nonblocking, Global => null, Use_Formal => null;

procedure Assign (Target : in out Tree; Source : in Tree)
with Pre => not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error,
     Post => Node_Count (Source) = Node_Count (Target);

function Copy (Source : Tree) return Tree
with Post =>
     Node_Count (Copy'Result) = Node_Count (Source) and then
     not Tampering_With_Elements_Prohibited (Copy'Result) and then
     not Tampering_With_Cursors_Prohibited (Copy'Result);

```

```

procedure Move (Target : in out Tree;
                Source : in out Tree)
    with Pre => (not Tampering_With_Cursors_Prohibited (Target)
                or else raise Program_Error) and then
                (not Tampering_With_Cursors_Prohibited (Source)
                or else raise Program_Error),
    Post => (if not Target'Has_Same_Storage (Source) then
            Node_Count (Target) = Node_Count (Source'Old) and then
            Node_Count (Source) = 1);

procedure Delete_Leaf (Container : in out Tree;
                      Position : in out Cursor)
    with Pre => (not Tampering_With_Cursors_Prohibited (Container)
                or else raise Program_Error) and then
                (Position /= No_Element
                or else raise Constraint_Error) and then
                (Has_Element (Container, Position)
                or else raise Program_Error) and then
                (Is_Leaf (Container, Position)
                or else raise Constraint_Error),
    Post =>
        Node_Count (Container)'Old = Node_Count (Container)+1 and then
        Position = No_Element;

procedure Delete_Subtree (Container : in out Tree;
                        Position : in out Cursor)
    with Pre => (not Tampering_With_Cursors_Prohibited (Container)
                or else raise Program_Error) and then
                (Position /= No_Element
                or else raise Constraint_Error) and then
                (Has_Element (Container, Position)
                or else raise Program_Error),
    Post => Node_Count (Container)'Old = Node_Count (Container) +
            Subtree_Node_Count (Container, Position)'Old and then
            Position = No_Element;

procedure Swap (Container : in out Tree;
                I, J : in Cursor)
    with Pre => (not Tampering_With_Cursors_Prohibited (Container)
                or else raise Program_Error) and then
                (I /= No_Element or else Constraint_Error) and then
                (J /= No_Element or else Constraint_Error) and then
                (Has_Element (Container, I)
                or else raise Program_Error) and then
                (Has_Element (Container, J)
                or else raise Program_Error);

function Find (Container : Tree;
               Item : Element_Type)
    return Cursor
    with Post => (if Find'Result /= No_Element
                then Has_Element (Container, Find'Result));

function Find_In_Subtree (Position : Cursor;
                        Item : Element_Type)
    return Cursor
    with Pre => Position /= No_Element or else raise Constraint_Error,
    Post => (if Find_In_Subtree'Result = No_Element
                then Has_Element (Find_In_Subtree'Result)),
    Global => in all;

```

```

function Find_In_Subtree (Container : Tree;
                          Position  : Cursor;
                          Item      : Element_Type)

    return Cursor
    with Pre => (Position /= No_Element
                or else raise Constraint_Error) and then
                (Meaningful_For (Container, Position)
                or else raise Program_Error),
    Post => (if Find_In_Subtree'Result /= No_Element
            then Has_Element (Container, Find_In_Subtree'Result));

function Ancestor_Find (Position : Cursor;
                        Item      : Element_Type)

    return Cursor
    with Pre => Position /= No_Element or else raise Constraint_Error,
    Post => (if Ancestor_Find'Result = No_Element
            then Has_Element (Ancestor_Find'Result)),
    Global => in all;

function Ancestor_Find (Container : Tree;
                        Position  : Cursor;
                        Item      : Element_Type)

    return Cursor
    with Pre => (Position /= No_Element
                or else raise Constraint_Error) and then
                (Meaningful_For (Container, Position)
                or else raise Program_Error),
    Post => (if Ancestor_Find'Result = No_Element
            then Has_Element (Container, Ancestor_Find'Result));

function Contains (Container : Tree;
                  Item      : Element_Type) return Boolean;

procedure Iterate
    (Container : in Tree;
     Process   : not null access procedure (Position : in Cursor))
    with Allows_Exit;

procedure Iterate_Subtree
    (Position  : in Cursor;
     Process   : not null access procedure (Position : in Cursor))
    with Allows_Exit,
    Pre => Position /= No_Element or else raise Constraint_Error,
    Global => in all;

procedure Iterate_Subtree
    (Container : in Tree;
     Position  : in Cursor;
     Process   : not null access procedure (Position : in Cursor))
    with Allows_Exit,
    Pre => (Position /= No_Element
            or else raise Constraint_Error) and then
            (Meaningful_For (Container, Position)
            or else raise Program_Error);

function Iterate (Container : in Tree)
    return Tree_Iterator_Interfaces.Parallel_Iterator'Class
    with Post => Tampering_With_Cursors_Prohibited (Container);

function Iterate_Subtree (Position : in Cursor)
    return Tree_Iterator_Interfaces.Parallel_Iterator'Class
    with Pre   => Position /= No_Element or else raise Constraint_Error,
    Global => in all;

```

```

function Iterate_Subtree (Container : in Tree; Position : in Cursor)
  return Tree_Iterator_Interfaces.Parallel_Iterator'Class
  with Pre => (Position /= No_Element
    or else raise Constraint_Error) and then
    (Meaningful_For (Container, Position)
    or else raise Program_Error),
    Post => Tampering_With_Cursors_Prohibited (Container);

function Child_Count (Parent : Cursor) return Count_Type
  with Post => (if Parent = No_Element then Child_Count'Result = 0),
  with Nonblocking, Global => in all, Use_Formal => null;

function Child_Count (Container : Tree; Parent : Cursor)
  return Count_Type
  with Pre => Meaningful_For (Container, Parent)
    or else raise Program_Error,
    Post => (if Parent = No_Element then Child_Count'Result = 0),
    Nonblocking, Global => null, Use_Formal => null;

function Child_Depth (Parent, Child : Cursor) return Count_Type
  with Pre => (Parent = No_Element and then Child = No_Element)
    or else raise Constraint_Error,
  with Nonblocking, Global => in all, Use_Formal => null;

function Child_Depth (Container : Tree; Parent, Child : Cursor)
  return Count_Type
  with Pre => ((Parent = No_Element and then Child = No_Element)
    or else raise Constraint_Error) and then
    (Meaningful_For (Container, Parent)
    or else raise Program_Error) and then
    (Meaningful_For (Container, Child)
    or else raise Program_Error),
    Nonblocking, Global => null, Use_Formal => null;

procedure Insert_Child (Container : in out Tree;
  Parent      : in      Cursor;
  Before      : in      Cursor;
  New_Item    : in      Element_Type;
  Count       : in      Count_Type := 1)
  with Pre => (not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error) and then
    (Parent /= No_Element
    or else raise Constraint_Error) and then
    (Meaningful_For (Container, Parent)
    or else raise Program_Error) and then
    (Meaningful_For (Container, Before)
    or else raise Program_Error) and then
    (Before = No_Element or else
    Container.Parent (Before) = Parent
    or else raise Constraint_Error),
    Post => Node_Count (Container) =
    Node_Count (Container)'Old + Count;

```



```

procedure Insert_Child (Container : in out Tree;
                        Parent    : in      Cursor;
                        Before    : in      Cursor;
                        New_Item  : in      Element_Type;
                        Position  : out      Cursor;
                        Count     : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
(Parent /= No_Element
 or else raise Constraint_Error) and then
(Meaningful_For (Container, Parent)
 or else raise Program_Error) and then
(Meaningful_For (Container, Before)
 or else raise Program_Error) and then
(Before = No_Element or else
 Container.Parent (Before) = Parent
 or else raise Constraint_Error),
Post => (Node_Count (Container) =
         Node_Count (Container)'Old + Count) and then
Has_Element (Container, Position);

procedure Insert_Child (Container : in out Tree;
                        Parent    : in      Cursor;
                        Before    : in      Cursor;
                        Position  : out      Cursor;
                        Count     : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
(Parent /= No_Element
 or else raise Constraint_Error) and then
(Meaningful_For (Container, Parent)
 or else raise Program_Error) and then
(Meaningful_For (Container, Before)
 or else raise Program_Error) and then
(Before = No_Element or else
 Container.Parent (Before) = Parent
 or else raise Constraint_Error),
Post => (Node_Count (Container) =
         Node_Count (Container)'Old + Count) and then
Has_Element (Container, Position);

procedure Prepend_Child (Container : in out Tree;
                        Parent    : in      Cursor;
                        New_Item  : in      Element_Type;
                        Count     : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
(Parent /= No_Element
 or else raise Constraint_Error) and then
(Meaningful_For (Container, Parent)
 or else raise Program_Error),
Post => Node_Count (Container) =
         Node_Count (Container)'Old + Count;

procedure Append_Child (Container : in out Tree;
                        Parent    : in      Cursor;
                        New_Item  : in      Element_Type;
                        Count     : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
(Parent /= No_Element
 or else raise Constraint_Error) and then
(Meaningful_For (Container, Parent)
 or else raise Program_Error),
Post => Node_Count (Container) =
         Node_Count (Container)'Old + Count;

```

```

procedure Delete_Children (Container : in out Tree;
                           Parent    : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Container, Parent)
   or else raise Program_Error),
  Post => (Node_Count (Container) = Node_Count (Container)'Old -
          Child_Count (Container, Parent)'Old) and then
  Child_Count (Container, Parent) = 0;

procedure Copy_Subtree (Target    : in out Tree;
                       Parent     : in      Cursor;
                       Before     : in      Cursor;
                       Source     : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
  (Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Target, Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Target, Before)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Target.Parent (Before) = Parent
   or else raise Constraint_Error) and then
  (not Is_Root (Source)
   or else raise Constraint_Error),
  Post => Node_Count (Target) =
          Node_Count (Target)'Old + Subtree_Node_Count (Source),
  Global => in all;

procedure Copy_Local_Subtree (Target    : in out Tree;
                              Parent     : in      Cursor;
                              Before     : in      Cursor;
                              Source     : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
  (Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Target, Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Target, Before)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Target.Parent (Before) = Parent
   or else raise Constraint_Error) and then
  (Meaningful_For (Target, Source)
   or else raise Program_Error) and then
  (not Is_Root (Source)
   or else raise Constraint_Error),
  Post => Node_Count (Target) = Node_Count (Target)'Old +
          Subtree_Node_Count (Target, Source);

```

```

procedure Copy_Subtree (Target    : in out Tree;
                        Parent    : in      Cursor;
                        Before    : in      Cursor;
                        Source    : in      Tree;
                        Subtree   : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
or else raise Program_Error) and then
(Parent /= No_Element
or else raise Constraint_Error) and then
(Meaningful_For (Target, Parent)
or else raise Program_Error) and then
(Meaningful_For (Target, Before)
or else raise Program_Error) and then
(Before = No_Element or else
Target.Parent (Before) = Parent
or else raise Constraint_Error) and then
(Meaningful_For (Source, Subtree)
or else raise Program_Error) and then
(not Is_Root (Source, Subtree)
or else raise Constraint_Error),
Post => Node_Count (Target) = Node_Count (Target)'Old +
Subtree_Node_Count (Source, Subtree);

procedure Splice_Subtree (Target    : in out Tree;
                          Parent    : in      Cursor;
                          Before    : in      Cursor;
                          Source    : in out Tree;
                          Position  : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
or else raise Program_Error) and then
(not Tampering_With_Cursors_Prohibited (Source)
or else raise Program_Error) and then
(Parent /= No_Element
or else raise Constraint_Error) and then
(Meaningful_For (Target, Parent)
or else raise Program_Error) and then
(Meaningful_For (Target, Before)
or else raise Program_Error) and then
(Before = No_Element or else
Target.Parent (Before) /= Parent
or else raise Constraint_Error) and then
(Position /= No_Element
or else raise Constraint_Error) and then
(Has_Element (Source, Position)
or else raise Program_Error) and then
(Target'Has_Same_Storage (Source) or else
Position = Before or else
Is_Ancessor_Of (Target, Position, Parent)
or else raise Constraint_Error),
Post => (declare
Org_Sub_Count renames
Subtree_Node_Count (Source, Position)'Old;
Org_Target_Count renames Node_Count (Target)'Old;
begin
(if not Target'Has_Same_Storage (Source) then
Node_Count (Target) = Org_Target_Count +
Org_Sub_Count and then
Node_Count (Source) = Node_Count (Source)'Old -
Org_Sub_Count and then
Has_Element (Target, Position)
else
Target.Parent (Position) = Parent and then
Node_Count (Target) = Org_Target_Count));

```

```

procedure Splice_Subtree (Container: in out Tree;
                          Parent   : in      Cursor;
                          Before   : in      Cursor;
                          Position : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Container, Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Container, Before)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Container.Parent (Before) /= Parent
   or else raise Constraint_Error) and then
  (Position /= No_Element
   or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error) and then
  (Position = Before or else
   Is_Ancessor_Of (Container, Position, Parent)
   or else raise Constraint_Error);
  Post => (Node_Count (Container) =
          Node_Count (Container)'Old and then
          Container.Parent (Position) = Parent);

procedure Splice_Children (Target      : in out Tree;
                           Target_Parent : in      Cursor;
                           Before        : in      Cursor;
                           Source        : in out Tree;
                           Source_Parent : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
  (not Tampering_With_Cursors_Prohibited (Source)
   or else raise Program_Error) and then
  (Target_Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Target, Target_Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Target, Before)
   or else raise Program_Error) and then
  (Source_Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Source, Source_Parent)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Parent (Target, Before) /= Target_Parent
   or else raise Constraint_Error) and then
  (Target'Has_Same_Storage (Source) or else
   Target_Parent = Source_Parent or else
   Is_Ancessor_Of (Target, Source_Parent, Target_Parent)
   or else raise Constraint_Error);
  Post => (declare
          Org_Child_Count renames
            Child_Count (Source, Source_Parent)'Old;
          Org_Target_Count renames Node_Count (Target)'Old;
begin
    (if not Target'Has_Same_Storage (Source) then
     Node_Count (Target) = Org_Target_Count +
       Org_Child_Count and then
     Node_Count (Source) = Node_Count (Source)'Old -
       Org_Child_Count
    else
     Node_Count (Target) = Org_Target_Count));

```

```

procedure Splice_Children (Container      : in out Tree;
                          Target_Parent  : in      Cursor;
                          Before         : in      Cursor;
                          Source_Parent  : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Target_Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Container, Target_Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Container, Before)
   or else raise Program_Error) and then
  (Source_Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Container, Source_Parent)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Parent (Container, Before) /= Target_Parent
   or else raise Constraint_Error) and then
  (Target_Parent = Source_Parent or else
   Is_Ancessor_Of (Container, Source_Parent, Target_Parent)
   or else raise Constraint_Error),
  Post => Node_Count (Container) = Node_Count (Container)'Old;

function Parent (Position : Cursor) return Cursor
with Nonblocking, Global => in all, Use_Formal => null,
  Post => (if Position = No_Element or else
           Is_Root (Position) then Parent'Result = No_Element);

function Parent (Container : Tree;
                 Position  : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
  Pre  => Meaningful_For (Container, Position)
       or else raise Program_Error,
  Post => (if Position = No_Element or else
           Is_Root (Container, Position)
           then Parent'Result = No_Element
           else Has_Element (Container, Parent'Result));

function First_Child (Parent : Cursor) return Cursor
with Nonblocking, Global => in all, Use_Formal => null,
  Pre  => Parent /= No_Element or else raise Constraint_Error;

function First_Child (Container : Tree;
                     Parent     : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
  Pre  => (Parent /= No_Element
       or else raise Constraint_Error) and then
  (Meaningful_For (Container, Parent)
   or else raise Program_Error),
  Post => First_Child'Result = No_Element or else
       Has_Element (Container, First_Child'Result);

function First_Child_Element (Parent : Cursor) return Element_Type
with Nonblocking, Global => in all, Use_Formal => Element_Type,
  Pre  => (Parent /= No_Element and then
       Last_Child (Parent) /= No_Element)
       or else raise Constraint_Error;

function First_Child_Element (Container : Tree;
                             Parent     : Cursor) return Element_Type
with Nonblocking, Global => null, Use_Formal => Element_Type,
  Pre  => (Parent /= No_Element
       or else raise Constraint_Error) and then
  (Meaningful_For (Container, Parent)
   or else raise Program_Error) and then
  (First_Child (Container, Parent) /= No_Element
   or else raise Constraint_Error);

```

```

function Last_Child (Parent : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Pre => Parent /= No_Element or else raise Constraint_Error;

function Last_Child (Container : Tree;
  Parent : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => (Parent /= No_Element
      or else raise Constraint_Error) and then
      (Meaningful_For (Container, Parent)
      or else raise Program_Error),
    Post => Last_Child'Result = No_Element or else
      Has_Element (Container, Last_Child'Result);

function Last_Child_Element (Parent : Cursor) return Element_Type
  with Nonblocking, Global => in all, Use_Formal => Element_Type,
    Pre => (Parent /= No_Element and then
      Last_Child (Parent) /= No_Element)
      or else raise Constraint_Error;

function Last_Child_Element (Container : Tree;
  Parent : Cursor) return Element_Type
  with Nonblocking, Global => null, Use_Formal => Element_Type,
    Pre => (Parent /= No_Element
      or else raise Constraint_Error) and then
      (Meaningful_For (Container, Parent)
      or else raise Program_Error) and then
      (Last_Child (Container, Parent) /= No_Element
      or else raise Constraint_Error);

function Next_Sibling (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element
      then Next_Sibling'Result = No_Element);

function Next_Sibling (Container : Tree;
  Position : Cursor) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Meaningful_For (Container, Position)
      or else raise Program_Error,
    Post => (if Next_Sibling'Result = No_Element then
      Position = No_Element or else
      Is_Root (Container, Position) or else
      Last_Child (Container, Parent (Container, Position))
      = Position
      else Has_Element (Container, Next_Sibling'Result));

procedure Next_Sibling (Position : in out Cursor)
  with Nonblocking, Global => in all, Use_Formal => null;

procedure Next_Sibling (Container : in Tree;
  Position : in out Cursor)
  with Nonblocking, Global => null, Use_Formal => null,
    Pre => Meaningful_For (Container, Position)
      or else raise Program_Error,
    Post => (if Position /= No_Element
      then Has_Element (Container, Position));

function Previous_Sibling (Position : Cursor) return Cursor
  with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element
      then Previous_Sibling'Result = No_Element);

```



```

function Previous_Sibling (Container : Tree;
                           Position  : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
     Pre  => Meaningful_For (Container, Position)
     or else raise Program_Error,
     Post => (if Previous_Sibling'Result = No_Element then
               Position = No_Element or else
               Is_Root (Container, Position) or else
               First_Child (Container, Parent (Container, Position))
               = Position
             else Has_Element (Container, Previous_Sibling'Result));

procedure Previous_Sibling (Position : in out Cursor)
with Nonblocking, Global => in all, Use_Formal => null;

procedure Previous_Sibling (Container : in Tree;
                           Position  : in out Cursor)
with Nonblocking, Global => null, Use_Formal => null,
     Pre  => Meaningful_For (Container, Position)
     or else raise Program_Error,
     Post => (if Position /= No_Element
             then Has_Element (Container, Position));

procedure Iterate_Children
(Parent : in Cursor;
 Process : not null access procedure (Position : in Cursor))
with Allows_Exit,
     Pre  => Parent /= No_Element or else raise Constraint_Error,
     Global => in all, Use_Formal => null;

procedure Iterate_Children
(Container : in Tree;
 Parent    : in Cursor;
 Process   : not null access procedure (Position : in Cursor))
with Allows_Exit,
     Pre  => (Parent /= No_Element
             or else raise Constraint_Error) and then
             (Meaningful_For (Container, Parent)
             or else raise Program_Error);

procedure Reverse_Iterate_Children
(Parent : in Cursor;
 Process : not null access procedure (Position : in Cursor))
with Allows_Exit,
     Pre  => Parent /= No_Element or else raise Constraint_Error,
     Global => in all, Use_Formal => null;

procedure Reverse_Iterate_Children
(Container : in Tree;
 Parent    : in Cursor;
 Process   : not null access procedure (Position : in Cursor))
with Allows_Exit,
     Pre  => (Parent /= No_Element
             or else raise Constraint_Error) and then
             (Meaningful_For (Container, Parent)
             or else raise Program_Error);

function Iterate_Children (Container : in Tree; Parent : in Cursor)
return Tree_Iterator_Interfaces.Parallel_Reversible_Iterator'Class
with Pre  => (Parent /= No_Element
             or else raise Constraint_Error) and then
             (Meaningful_For (Container, Parent)
             or else raise Program_Error),
     Post => Tampering_With_Cursors_Prohibited (Container);

package Stable is

```

```

type Tree (Base : not null access Multiway_Trees.Tree) is
  tagged limited private
    with Constant_Indexing => Constant_Reference,
          Variable_Indexing => Reference,
          Default_Iterator => Iterate,
          Iterator_Element => Element_Type,
          Stable_Properties => (Node_Count),
          Global => null,
          Default_Initial_Condition => Node_Count (Tree) = 1,
          Preelaborable_Initialization;

type Cursor is private
  with Preelaborable_Initialization;

Empty_Tree : constant Tree;
No_Element : constant Cursor;

function Has_Element (Position : Cursor) return Boolean
  with Nonblocking, Global => in all, Use_Formal => null;

package Tree_Iterator_Interfaces is new
  Ada.Iterator_Interfaces (Cursor, Has_Element);

procedure Assign (Target : in out Multiway_Trees.Tree;
                  Source : in Tree)
  with Post => Node_Count (Source) = Node_Count (Target);

function Copy (Source : Multiway_Trees.Tree) return Tree
  with Post => Node_Count (Copy'Result) = Node_Count (Source);

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
  with Implicit_Dereference => Element,
        Nonblocking, Global => null,
        Default_Initial_Condition => (raise Program_Error);

type Reference_Type
  (Element : not null access Element_Type) is private
  with Implicit_Dereference => Element,
        Nonblocking, Global => null,
        Default_Initial_Condition => (raise Program_Error);

-- Additional subprograms as described in the text
-- are declared here.

private
  ... -- not specified by the language
end Stable;

private
  ... -- not specified by the language
end Ada.Containers.Multiway_Trees;

```

The actual function for the generic formal function "=" on Element\_Type values is expected to define a reflexive and symmetric relationship and return the same result value each time it is called with a particular pair of values. If it behaves in some other manner, the functions Find, Reverse\_Find, Equal\_Subtree, and "=" on tree values return an unspecified value. The exact arguments and number of calls of this generic formal function by the functions Find, Reverse\_Find, Equal\_Subtree, and "=" on tree values are unspecified.

The type Tree is used to represent trees. The type Tree needs finalization (see 10.6).

Empty\_Tree represents the empty Tree object. It contains only the root node (Node\_Count (Empty\_Tree) returns 1). If an object of type Tree is not otherwise initialized, it is initialized to the same value as Empty\_Tree.

No\_Element represents a cursor that designates no element. If an object of type Cursor is not otherwise initialized, it is initialized to the same value as No\_Element.

The primitive "=" operator for type Cursor returns True if both cursors are No\_Element, or designate the same element in the same container.

Execution of the default implementation of the Input, Output, Read, or Write attribute of type Cursor raises Program\_Error.

Tree'Write for a Tree object *T* writes Node\_Count(*T*) - 1 elements of the tree to the stream. It may also write additional information about the tree.

Tree'Read reads the representation of a tree from the stream, and assigns to *Item* a tree with the same elements and structure as was written by Tree'Write.

Some operations check for "tampering with cursors" of a container because they depend on the set of elements of the container remaining constant, and others check for "tampering with elements" of a container because they depend on elements of the container not being replaced. When tampering with cursors is *prohibited* for a particular tree object *T*, Program\_Error is propagated by the finalization of *T*, as well as by a call that passes *T* to certain of the operations of this package, as indicated by the precondition of such an operation. Similarly, when tampering with elements is *prohibited* for *T*, Program\_Error is propagated by a call that passes *T* to certain of the other operations of this package, as indicated by the precondition of such an operation.

```
function Has_Element (Position : Cursor) return Boolean
with Nonblocking, Global => in all, Use_Formal => null;
```

Returns True if Position designates an element, and returns False otherwise. In particular, Has\_Element returns False if the cursor designates a root node or equals No\_Element.

```
function Has_Element (Container : Tree; Position : Cursor)
return Boolean
with Nonblocking, Global => null, Use_Formal => null;
```

Returns True if Position designates an element in Container, and returns False otherwise. In particular, Has\_Element returns False if the cursor designates a root node or equals No\_Element.

```
function Equal_Subtree (Left_Position : Cursor;
                       Right_Position: Cursor) return Boolean;
```

If Left\_Position or Right\_Position equals No\_Element, propagates Constraint\_Error. If the number of child nodes of the element designated by Left\_Position is different from the number of child nodes of the element designated by Right\_Position, the function returns False. If Left\_Position designates a root node and Right\_Position does not, the function returns False. If Right\_Position designates a root node and Left\_Position does not, the function returns False. Unless both cursors designate a root node, the elements are compared using the generic formal equality operator. If the result of the element comparison is False, the function returns False. Otherwise, it calls Equal\_Subtree on a cursor designating each child element of the element designated by Left\_Position and a cursor designating the corresponding child element of the element designated by Right\_Position. If any such call returns False, the function returns False; otherwise, it returns True. Any exception raised during the evaluation of element equality is propagated.

```
function "=" (Left, Right : Tree) return Boolean;
```

If Left and Right denote the same tree object, then the function returns True. Otherwise, it calls Equal\_Subtree with cursors designating the root nodes of Left and Right; the result is returned. Any exception raised during the evaluation of Equal\_Subtree is propagated.

```
function Tampering_With_Cursors_Prohibited
(Container : Tree) return Boolean
with Nonblocking, Global => null, Use_Formal => null;
```

Returns True if tampering with cursors or tampering with elements is currently prohibited for Container, and returns False otherwise.

```
function Tampering_With_Elements_Prohibited
(Container : Tree) return Boolean
with Nonblocking, Global => null, Use_Formal => null;
```

Always returns False, regardless of whether tampering with elements is prohibited.

```
function Is_Empty (Container : Tree) return Boolean
with Nonblocking, Global => null, Use_Formal => null,
Post => Is_Empty'Result = (Node_Count (Container) = 1);
```

Returns True if Container is empty.

```
function Node_Count (Container : Tree) return Count_Type
with Nonblocking, Global => null, Use_Formal => null;
```

Node\_Count returns the number of nodes in Container.

```
function Subtree_Node_Count (Position : Cursor) return Count_Type
with Nonblocking, Global => in all, Use_Formal => null;
```

If Position is No\_Element, Subtree\_Node\_Count returns 0; otherwise, Subtree\_Node\_Count returns the number of nodes in the subtree that is rooted by Position.

```
function Subtree_Node_Count (Container : Tree; Position : Cursor)
return Count_Type
with Pre => Meaningful_For (Container, Position)
or else raise Program_Error,
Nonblocking, Global => null, Use_Formal => null;
```

If Position is No\_Element, Subtree\_Node\_Count returns 0; otherwise, Subtree\_Node\_Count returns the number of nodes in the subtree of Container that is rooted by Position.

```
function Depth (Position : Cursor) return Count_Type
with Nonblocking, Global => in all, Use_Formal => null;
```

If Position equals No\_Element, Depth returns 0; otherwise, Depth returns the number of ancestor nodes of the node designated by Position (including the node itself).

```
function Depth (Container : Tree; Position : Cursor)
return Count_Type
with Pre => Meaningful_For (Container, Position)
or else raise Program_Error,
Nonblocking, Global => null, Use_Formal => null;
```

If Position equals No\_Element, Depth returns 0; otherwise, Depth returns the number of ancestor nodes of the node of Container designated by Position (including the node itself).

```
function Is_Root (Position : Cursor) return Boolean
with Nonblocking, Global => in all, Use_Formal => null;
```

Is\_Root returns True if the Position designates the root node of some tree; and returns False otherwise.

```
function Is_Root (Container : Tree; Position : Cursor)
return Boolean
with Nonblocking, Global => null, Use_Formal => null;
```

Is\_Root returns True if the Position designates the root node of Container; and returns False otherwise.

```
function Is_Leaf (Position : Cursor) return Boolean
with Nonblocking, Global => in all, Use_Formal => null;
```

Is\_Leaf returns True if Position designates a node that does not have any child nodes; and returns False otherwise.

```

function Is_Leaf (Container : Tree; Position : Cursor)
  return Boolean
  with Pre => Meaningful_For (Container, Position)
    or else raise Program_Error,
    Nonblocking, Global => null, Use_Formal => null;

```

Is\_Leaf returns True if Position designates a node in Container that does not have any child nodes; and returns False otherwise.

```

function Is_Ancessor_Of (Container : Tree;
  Parent : Cursor;
  Position : Cursor) return Boolean
  with Pre => (Meaningful_For (Container, Position)
    or else raise Program_Error) and then
    (Meaningful_For (Container, Parent)
    or else raise Program_Error),
    Nonblocking, Global => null, Use_Formal => null;

```

Is\_Ancessor\_Of returns True if Parent designates an ancestor node of Position (including Position itself), and returns False otherwise.

```

function Root (Container : Tree) return Cursor
  with Nonblocking, Global => null, Use_Formal => null,
    Post => Root'Result /= No_Element and then
    not Has_Element (Container, Root'Result);

```

Root returns a cursor that designates the root node of Container.

```

procedure Clear (Container : in out Tree)
  with Pre => not Tampering_With_Cursors_Prohibited (Container)
    or else raise Program_Error,
    Post => Node_Count (Container) = 1;

```

Removes all the elements from Container.

```

function Element (Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element or else
    raise Constraint_Error) and then
    (Has_Element (Position) or else raise Program_Error),
    Nonblocking, Global => in all, Use_Formal => Element_Type;

```

Element returns the element designated by Position.

```

function Element (Container : Tree;
  Position : Cursor) return Element_Type
  with Pre => (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error),
    Nonblocking, Global => null, Use_Formal => Element_Type;

```

Element returns the element designated by Position in Container.

```

procedure Replace_Element (Container : in out Tree;
  Position : in Cursor;
  New_item : in Element_Type)
  with Pre => (not Tampering_With_Elements_Prohibited (Container)
    or else raise Program_Error) and then
    (Position /= No_Element
    or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
    or else raise Program_Error);

```

Replace\_Element assigns the value New\_Item to the element designated by Position. For the purposes of determining whether the parameters overlap in a call to Replace\_Element, the Container parameter is not considered to overlap with any object (including itself).

```

procedure Query_Element
(Position : in Cursor;
 Process  : not null access procedure (Element : in Element_Type))
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
    (Has_Element (Position)
     or else raise Program_Error),
    Global => in all;

```

Query\_Element calls Process.all with the element designated by Position as the argument. Tampering with the elements of the tree that contains the element designated by Position is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Query_Element
(Container : in Tree;
 Position  : in Cursor;
 Process   : not null access procedure (Element : in Element_Type))
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
     or else raise Program_Error);

```

Query\_Element calls Process.all with the element designated by Position as the argument. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Update_Element
(Container : in out Tree;
 Position  : in Cursor;
 Process   : not null access procedure
             (Element : in out Element_Type))
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
    (Has_Element (Container, Position)
     or else raise Program_Error);

```

Update\_Element calls Process.all with the element designated by Position as the argument. Tampering with the elements of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

If Element\_Type is unconstrained and definite, then the actual Element parameter of Process.all shall be unconstrained.

```

type Constant_Reference_Type
(Element : not null access constant Element_Type) is private
with Implicit_Dereference => Element,
    Nonblocking, Global => in out synchronized,
    Default_Initial_Condition => (raise Program_Error);

type Reference_Type (Element : not null access Element_Type) is private
with Implicit_Dereference => Element,
    Nonblocking, Global => in out synchronized,
    Default_Initial_Condition => (raise Program_Error);

```

The types Constant\_Reference\_Type and Reference\_Type need finalization.



```

function Constant_Reference (Container : aliased in Tree;
                             Position  : in Cursor)
return Constant_Reference_Type
with Pre  => (Position /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Position)
              or else raise Program_Error),
  Post   => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Constant\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read access to an individual element of a tree given a cursor.

Constant\_Reference returns an object whose discriminant is an access value that designates the element designated by Position. Tampering with the elements of Container is prohibited while the object returned by Constant\_Reference exists and has not been finalized.

```

function Reference (Container : aliased in out Tree;
                    Position  : in Cursor)
return Reference_Type
with Pre  => (Position /= No_Element
              or else raise Constraint_Error) and then
              (Has_Element (Container, Position)
              or else raise Program_Error),
  Post   => Tampering_With_Cursors_Prohibited (Container),
  Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Variable\_Indexing and Implicit\_Dereference aspects) provides a convenient way to gain read and write access to an individual element of a tree given a cursor.

Reference returns an object whose discriminant is an access value that designates the element designated by Position. Tampering with the elements of Container is prohibited while the object returned by Reference exists and has not been finalized.

```

procedure Assign (Target : in out Tree; Source : in Tree)
with Pre  => not Tampering_With_Cursors_Prohibited (Target)
              or else raise Program_Error,
  Post => Node_Count (Source) = Node_Count (Target);

```

If Target denotes the same object as Source, the operation has no effect. Otherwise, the elements of Source are copied to Target as for an assignment\_statement assigning Source to Target.

```

function Copy (Source : Tree) return Tree
with Post =>
  Node_Count (Copy'Result) = Node_Count (Source) and then
  not Tampering_With_Elements_Prohibited (Copy'Result) and then
  not Tampering_With_Cursors_Prohibited (Copy'Result);

```

Returns a tree with the same structure as Source and whose elements are initialized from the corresponding elements of Source.

```

procedure Move (Target : in out Tree;
                Source : in out Tree)
with Pre  => (not Tampering_With_Cursors_Prohibited (Target)
              or else raise Program_Error) and then
              (not Tampering_With_Cursors_Prohibited (Source)
              or else raise Program_Error),
  Post => (if not Target'Has_Same_Storage (Source) then
          Node_Count (Target) = Node_Count (Source'Old) and then
          Node_Count (Source) = 1);

```

If Target denotes the same object as Source, then the operation has no effect. Otherwise, Move first calls Clear (Target). Then, the nodes other than the root node in Source are moved to

Target (in the same positions). After Move completes, Node\_Count (Target) is the number of nodes originally in Source, and Node\_Count (Source) is 1.

```
procedure Delete_Leaf (Container : in out Tree;
                      Position  : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (Position /= No_Element
              or else raise Constraint_Error) and then
             (Has_Element (Container, Position)
              or else raise Program_Error) and then
             (Is_Leaf (Container, Position)
              or else raise Constraint_Error),
  Post =>
    Node_Count (Container)'Old = Node_Count (Container) + 1 and then
    Position = No_Element;
```

Delete\_Leaf removes (from Container) the element designated by Position, and Position is set to No\_Element.

```
procedure Delete_Subtree (Container : in out Tree;
                        Position  : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (Position /= No_Element
              or else raise Constraint_Error) and then
             (Has_Element (Container, Position)
              or else raise Program_Error),
  Post => Node_Count (Container)'Old = Node_Count (Container) +
          Subtree_Node_Count (Container, Position)'Old and then
          Position = No_Element;
```

Delete\_Subtree removes (from Container) the subtree designated by Position (that is, all descendants of the node designated by Position including the node itself), and Position is set to No\_Element.

```
procedure Swap (Container : in out Tree;
               I, J      : in Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
             (I /= No_Element or else Constraint_Error) and then
             (J /= No_Element or else Constraint_Error) and then
             (Has_Element (Container, I)
              or else raise Program_Error) and then
             (Has_Element (Container, J)
              or else raise Program_Error);
```

Swap exchanges the values of the elements designated by I and J.

```
function Find (Container : Tree;
              Item       : Element_Type)
return Cursor
with Post => (if Find'Result /= No_Element
             then Has_Element (Container, Find'Result));
```

Find searches the elements of Container for an element equal to Item (using the generic formal equality operator). The search starts at the root node. The search traverses the tree in a depth-first order. If no equal element is found, then Find returns No\_Element. Otherwise, it returns a cursor designating the first equal element encountered.

```

function Find_In_Subtree (Position : Cursor;
                        Item       : Element_Type)
return Cursor
with Pre => Position /= No_Element or else raise Constraint_Error,
      Post => (if Find_In_Subtree'Result = No_Element
              then Has_Element (Find_In_Subtree'Result)),
      Global => in all;

```

Find\_In\_Subtree searches the subtree rooted by Position for an element equal to Item (using the generic formal equality operator). The search starts at the element designated by Position. The search traverses the subtree in a depth-first order. If no equal element is found, then Find returns No\_Element. Otherwise, it returns a cursor designating the first equal element encountered.

```

function Find_In_Subtree (Container : Tree;
                        Position   : Cursor;
                        Item       : Element_Type)

return Cursor
with Pre => (Position /= No_Element
            or else raise Constraint_Error) and then
            (Meaningful_For (Container, Position)
            or else raise Program_Error),
      Post => (if Find_In_Subtree'Result = No_Element
              then Has_Element (Container, Find_In_Subtree'Result));

```

Find\_In\_Subtree searches the subtree of Container rooted by Position for an element equal to Item (using the generic formal equality operator). The search starts at the element designated by Position. The search traverses the subtree in a depth-first order. If no equal element is found, then Find returns No\_Element. Otherwise, it returns a cursor designating the first equal element encountered.

```

function Ancestor_Find (Position : Cursor;
                       Item       : Element_Type)

return Cursor
with Pre => Position /= No_Element or else raise Constraint_Error,
      Post => (if Ancestor_Find'Result = No_Element
              then Has_Element (Container, Ancestor_Find'Result)),
      Global => in all;

```

Ancestor\_Find searches for an element equal to Item (using the generic formal equality operator). The search starts at the node designated by Position, and checks each ancestor proceeding toward the root of the subtree. If no equal element is found, then Ancestor\_Find returns No\_Element. Otherwise, it returns a cursor designating the first equal element encountered.

```

function Ancestor_Find (Container : Tree;
                       Position   : Cursor;
                       Item       : Element_Type)

return Cursor
with Pre => (Position /= No_Element
            or else raise Constraint_Error) and then
            (Meaningful_For (Container, Position)
            or else raise Program_Error),
      Post => (if Ancestor_Find'Result = No_Element
              then Has_Element (Container, Ancestor_Find'Result));

```

Ancestor\_Find searches for an element equal to Item (using the generic formal equality operator). The search starts at the node designated by Position in Container, and checks each ancestor proceeding toward the root of the subtree. If no equal element is found, then Ancestor\_Find returns No\_Element. Otherwise, it returns a cursor designating the first equal element encountered.

```
function Contains (Container : Tree;
                  Item      : Element_Type) return Boolean;
```

Equivalent to Find (Container, Item) /= No\_Element.

```
procedure Iterate
(Container : in Tree;
 Process   : not null access procedure (Position : in Cursor))
with Allows_Exit;
```

Iterate calls Process.all with a cursor that designates each element in Container, starting from the root node and proceeding in a depth-first order. Tampering with the cursors of Container is prohibited during the execution of a call on Process.all. Any exception raised by Process.all is propagated.

```
procedure Iterate_Subtree
(Position : in Cursor;
 Process  : not null access procedure (Position : in Cursor))
with Allows_Exit,
    Pre => Position /= No_Element or else raise Constraint_Error,
    Global => in all;
```

Iterate\_Subtree calls Process.all with a cursor that designates each element in the subtree rooted by the node designated by Position, starting from the node designated by Position and proceeding in a depth-first order. Tampering with the cursors of the tree that contains the element designated by Position is prohibited during the execution of a call on Process.all. Any exception raised by Process.all is propagated.

```
procedure Iterate_Subtree
(Container : in Tree;
 Position  : in Cursor;
 Process   : not null access procedure (Position : in Cursor))
with Allows_Exit,
    Pre => (Position /= No_Element
            or else raise Constraint_Error) and then
            (Meaningful_For (Container, Position)
            or else raise Program_Error);
```

Iterate\_Subtree calls Process.all with a cursor that designates each element in the subtree rooted by the node designated by Position in Container, starting from the node designated by Position and proceeding in a depth-first order. Tampering with the cursors of the tree that contains the element designated by Position is prohibited during the execution of a call on Process.all. Any exception raised by Process.all is propagated.

```
function Iterate (Container : in Tree)
return Tree_Iterator_Interfaces.Parallel_Iterator'Class
with Post => Tampering_With_Cursors_Prohibited (Container);
```

Iterate returns an iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each element in Container, starting from the root node and proceeding in a depth-first order when used as a forward iterator, and processing all nodes concurrently when used as a parallel iterator. Tampering with the cursors of Container is prohibited while the iterator object exists (in particular, in the sequence\_of\_statements of the loop\_statement whose iterator\_specification denotes this object). The iterator object needs finalization.

```
function Iterate_Subtree (Position : in Cursor)
return Tree_Iterator_Interfaces.Parallel_Iterator'Class
with Pre   => Position /= No_Element or else raise Constraint_Error,
    Global => in all;
```

Iterate\_Subtree returns an iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each element in the subtree rooted by the node designated by Position, starting from the node designated by Position and proceeding in a depth-first order

when used as a forward iterator, and processing all nodes in the subtree concurrently when used as a parallel iterator. Tampering with the cursors of the container that contains the node designated by `Position` is prohibited while the iterator object exists (in particular, in the `sequence_of_statements` of the `loop_statement` whose `iterator_specification` denotes this object). The iterator object needs finalization.

```
function Iterate_Subtree (Container : in Tree; Position : in Cursor)
return Tree_Iterator_Interfaces.Parallel_Iterator'Class
with Pre => (Position /= No_Element
             or else raise Constraint_Error) and then
             (Meaningful_For (Container, Position)
             or else raise Program_Error),
    Post => Tampering_With_Cursors_Prohibited (Container);
```

`Iterate_Subtree` returns an iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each element in the subtree rooted by the node designated by `Position` in `Container`, starting from the node designated by `Position` and proceeding in a depth-first order when used as a forward iterator, and processing all nodes in the subtree concurrently when used as a parallel iterator. Tampering with the cursors of the container that contains the node designated by `Position` is prohibited while the iterator object exists (in particular, in the `sequence_of_statements` of the `loop_statement` whose `iterator_specification` denotes this object). The iterator object needs finalization.

```
function Child_Count (Parent : Cursor) return Count_Type
with Post => (if Parent = No_Element then Child_Count'Result = 0),
    Nonblocking, Global => in all, Use_Formal => null;
```

`Child_Count` returns the number of child nodes of the node designated by `Parent`.

```
function Child_Count (Container : Tree; Parent : Cursor)
return Count_Type
with Pre => Meaningful_For (Container, Parent)
             or else raise Program_Error,
    Post => (if Parent = No_Element then Child_Count'Result = 0),
    Nonblocking, Global => null, Use_Formal => null;
```

`Child_Count` returns the number of child nodes of the node designated by `Parent` in `Container`.

```
function Child_Depth (Parent, Child : Cursor) return Count_Type
with Pre => (Parent /= No_Element and then Child /= No_Element)
             or else raise Constraint_Error,
    Nonblocking, Global => in all, Use_Formal => null;
```

`Child_Depth` returns the number of ancestor nodes of `Child` (including `Child` itself), up to but not including `Parent`; `Program_Error` is propagated if `Parent` is not an ancestor of `Child`.

```
function Child_Depth (Container : Tree; Parent, Child : Cursor)
return Count_Type
with Pre => ((Parent /= No_Element and then Child /= No_Element)
             or else raise Constraint_Error) and then
             (Meaningful_For (Container, Parent)
             or else raise Program_Error) and then
             (Meaningful_For (Container, Child)
             or else raise Program_Error),
    Nonblocking, Global => null, Use_Formal => null;
```

`Child_Depth` returns the number of ancestor nodes of `Child` within `Container` (including `Child` itself), up to but not including `Parent`; `Program_Error` is propagated if `Parent` is not an ancestor of `Child`.

```

procedure Insert_Child (Container : in out Tree;
                        Parent      : in      Cursor;
                        Before      : in      Cursor;
                        New_Item    : in      Element_Type;
                        Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Container, Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Container, Before)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Container.Parent (Before) = Parent
   or else raise Constraint_Error),
  Post => Node_Count (Container) =
         Node_Count (Container)'Old + Count;

```

Insert\_Child allocates Count nodes containing copies of New\_Item and inserts them as children of Parent. If Parent already has child nodes, then the new nodes are inserted prior to the node designated by Before, or, if Before equals No\_Element, the new nodes are inserted after the last existing child node of Parent. Any exception raised during allocation of internal storage is propagated, and Container is not modified.

```

procedure Insert_Child (Container : in out Tree;
                        Parent      : in      Cursor;
                        Before      : in      Cursor;
                        New_Item    : in      Element_Type;
                        Position    : out      Cursor;
                        Count       : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Container, Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Container, Before)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Container.Parent (Before) = Parent
   or else raise Constraint_Error),
  Post => (Node_Count (Container) =
         Node_Count (Container)'Old + Count) and then
         Has_Element (Container, Position);

```

Insert\_Child allocates Count nodes containing copies of New\_Item and inserts them as children of Parent. If Parent already has child nodes, then the new nodes are inserted prior to the node designated by Before, or, if Before equals No\_Element, the new nodes are inserted after the last existing child node of Parent. Position designates the first newly-inserted node, or if Count equals 0, then Position is assigned the value of Before. Any exception raised during allocation of internal storage is propagated, and Container is not modified.



```

procedure Insert_Child (Container : in out Tree;
                        Parent   : in      Cursor;
                        Before   : in      Cursor;
                        Position :      out Cursor;
                        Count    : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Parent /= No_Element
     or else raise Constraint_Error) and then
    (Meaningful_For (Container, Parent)
     or else raise Program_Error) and then
    (Meaningful_For (Container, Before)
     or else raise Program_Error) and then
    (Before = No_Element or else
     Container.Parent (Before) = Parent
     or else raise Constraint_Error),
    Post => (Node_Count (Container) =
            Node_Count (Container)'Old + Count) and then
            Has_Element (Container, Position);

```

Insert\_Child allocates Count nodes, the elements contained in the new nodes are initialized by default (see 6.3.1), and the new nodes are inserted as children of Parent. If Parent already has child nodes, then the new nodes are inserted prior to the node designated by Before, or, if Before equals No\_Element, the new nodes are inserted after the last existing child node of Parent. Position designates the first newly-inserted node, or if Count equals 0, then Position is assigned the value of Before. Any exception raised during allocation of internal storage is propagated, and Container is not modified.

```

procedure Prepend_Child (Container : in out Tree;
                        Parent   : in      Cursor;
                        New_Item : in      Element_Type;
                        Count    : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Parent /= No_Element
     or else raise Constraint_Error) and then
    (Meaningful_For (Container, Parent)
     or else raise Program_Error),
    Post => Node_Count (Container) =
            Node_Count (Container)'Old + Count;

```

Equivalent to Insert\_Child (Container, Parent, First\_Child (Container, Parent), New\_Item, Count).

```

procedure Append_Child (Container : in out Tree;
                        Parent   : in      Cursor;
                        New_Item : in      Element_Type;
                        Count    : in      Count_Type := 1)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
            or else raise Program_Error) and then
    (Parent /= No_Element
     or else raise Constraint_Error) and then
    (Meaningful_For (Container, Parent)
     or else raise Program_Error),
    Post => Node_Count (Container) =
            Node_Count (Container)'Old + Count;

```

Equivalent to Insert\_Child (Container, Parent, No\_Element, New\_Item, Count).

```

procedure Delete_Children (Container : in out Tree;
                           Parent    : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
    (Parent /= No_Element
     or else raise Constraint_Error) and then
    (Meaningful_For (Container, Parent)
     or else raise Program_Error),
    Post => (Node_Count (Container) = Node_Count (Container)'Old -
            Child_Count (Container, Parent)'Old) and then
            Child_Count (Container, Parent) = 0;

```

Delete\_Children removes (from Container) all of the descendants of Parent other than Parent itself.

```

procedure Copy_Subtree (Target    : in out Tree;
                       Parent     : in      Cursor;
                       Before      : in      Cursor;
                       Source      : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
    (Parent /= No_Element
     or else raise Constraint_Error) and then
    (Meaningful_For (Target, Parent)
     or else raise Program_Error) and then
    (Meaningful_For (Target, Before)
     or else raise Program_Error) and then
    (Before = No_Element or else
     Target.Parent (Before) = Parent
     or else raise Constraint_Error) and then
    (not Is_Root (Source)
     or else raise Constraint_Error),
    Post => Node_Count (Target) =
            Node_Count (Target)'Old + Subtree_Node_Count (Source),
    Global => in all;

```

If Source is equal to No\_Element, then the operation has no effect. Otherwise, the subtree rooted by Source (which can be from any tree; it does not have to be a subtree of Target) is copied (new nodes are allocated to create a new subtree with the same structure as the Source subtree, with each element initialized from the corresponding element of the Source subtree) and inserted into Target as a child of Parent. If Parent already has child nodes, then the new nodes are inserted prior to the node designated by Before, or, if Before equals No\_Element, the new nodes are inserted after the last existing child node of Parent. The parent of the newly created subtree is set to Parent, and the overall count of Target is incremented by Subtree\_Node\_Count (Source). Any exception raised during allocation of internal storage is propagated, and Container is not modified.

```

procedure Copy_Local_Subtree (Target    : in out Tree;
                             Parent    : in      Cursor;
                             Before    : in      Cursor;
                             Source    : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
  (Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Target, Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Target, Before)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Target.Parent (Before) = Parent
   or else raise Constraint_Error) and then
  (Meaningful_For (Target, Source)
   or else raise Program_Error) and then
  (not Is_Root (Source)
   or else raise Constraint_Error),
  Post => Node_Count (Target) = Node_Count (Target)'Old +
    Subtree_Node_Count (Target, Source);

```

If Source is equal to No\_Element, then the operation has no effect. Otherwise, the subtree rooted by Source in Target is copied (new nodes are allocated to create a new subtree with the same structure as the Source subtree, with each element initialized from the corresponding element of the Source subtree) and inserted into Target as a child of Parent. If Parent already has child nodes, then the new nodes are inserted prior to the node designated by Before, or, if Before equals No\_Element, the new nodes are inserted after the last existing child node of Parent. The parent of the newly created subtree is set to Parent. Any exception raised during allocation of internal storage is propagated, and Container is not modified.

```

procedure Copy_Subtree (Target    : in out Tree;
                       Parent    : in      Cursor;
                       Before    : in      Cursor;
                       Source    : in      Tree;
                       Subtree   : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
  (Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Target, Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Target, Before)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Target.Parent (Before) = Parent
   or else raise Constraint_Error) and then
  (Meaningful_For (Source, Subtree)
   or else raise Program_Error) and then
  (not Is_Root (Source, Subtree)
   or else raise Constraint_Error),
  Post => Node_Count (Target) = Node_Count (Target)'Old +
    Subtree_Node_Count (Source, Subtree);

```

If Subtree is equal to No\_Element, then the operation has no effect. Otherwise, the subtree rooted by Subtree in Source is copied (new nodes are allocated to create a new subtree with the same structure as the Subtree, with each element initialized from the corresponding element of the Subtree) and inserted into Target as a child of Parent. If Parent already has child nodes, then the new nodes are inserted prior to the node designated by Before, or, if Before equals No\_Element, the new nodes are inserted after the last existing child node of Parent. The parent of the newly created subtree is set to Parent. Any exception raised during allocation of internal storage is propagated, and Container is not modified.

```

procedure Splice_Subtree (Target    : in out Tree;
                          Parent    : in      Cursor;
                          Before    : in      Cursor;
                          Source    : in out Tree;
                          Position  : in out Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
             (not Tampering_With_Cursors_Prohibited (Source)
             or else raise Program_Error) and then
             (Parent /= No_Element
             or else raise Constraint_Error) and then
             (Meaningful_For (Target, Parent)
             or else raise Program_Error) and then
             (Meaningful_For (Target, Before)
             or else raise Program_Error) and then
             (Before = No_Element or else
              Target.Parent (Before) /= Parent
             or else raise Constraint_Error) and then
             (Position /= No_Element
             or else raise Constraint_Error) and then
             (Has_Element (Source, Position)
             or else raise Program_Error) and then
             (Target'Has_Same_Storage (Source) or else
              Position = Before or else
              Is_Ancessor_Of (Target, Position, Parent)
             or else raise Constraint_Error),
             Post => (declare
                      Org_Sub_Count renames
                        Subtree_Node_Count (Source, Position)'Old;
                      Org_Target_Count renames Node_Count (Target)'Old;
                      begin
                        (if not Target'Has_Same_Storage (Source) then
                          Node_Count (Target) = Org_Target_Count +
                            Org_Sub_Count and then
                          Node_Count (Source) = Node_Count (Source)'Old -
                            Org_Sub_Count and then
                          Has_Element (Target, Position)
                        else
                          Target.Parent (Position) = Parent and then
                          Node_Count (Target) = Org_Target_Count));

```

If Source denotes the same object as Target, then: if Position equals Before there is no effect; otherwise, the subtree rooted by the element designated by Position is moved to be a child of Parent. If Parent already has child nodes, then the moved nodes are inserted prior to the node designated by Before, or, if Before equals No\_Element, the moved nodes are inserted after the last existing child node of Parent. In each of these cases, Position and the count of Target are unchanged, and the parent of the element designated by Position is set to Parent.

Otherwise (if Source does not denote the same object as Target), the subtree designated by Position is removed from Source and moved to Target. The subtree is inserted as a child of Parent. If Parent already has child nodes, then the moved nodes are inserted prior to the node designated by Before, or, if Before equals No\_Element, the moved nodes are inserted after the last existing child node of Parent. In each of these cases, the count of Target is incremented by Subtree\_Node\_Count (Position), and the count of Source is decremented by Subtree\_Node\_Count (Position), Position is updated to represent an element in Target.

```

procedure Splice_Subtree (Container: in out Tree;
                          Parent   : in      Cursor;
                          Before    : in      Cursor;
                          Position  : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Container, Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Container, Before)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Container.Parent (Before) /= Parent
   or else raise Constraint_Error) and then
  (Position /= No_Element
   or else raise Constraint_Error) and then
  (Has_Element (Container, Position)
   or else raise Program_Error) and then
  (Position = Before or else
   Is_Ancessor_Of (Container, Position, Parent)
   or else raise Constraint_Error),
  Post => (Node_Count (Container) =
           Node_Count (Container)'Old and then
           Container.Parent (Position) = Parent);

```

If Position equals Before, there is no effect. Otherwise, the subtree rooted by the element designated by Position is moved to be a child of Parent. If Parent already has child nodes, then the moved nodes are inserted prior to the node designated by Before, or, if Before equals No\_Element, the moved nodes are inserted after the last existing child node of Parent. The parent of the element designated by Position is set to Parent.

```

procedure Splice_Children (Target           : in out Tree;
                          Target_Parent      : in      Cursor;
                          Before              : in      Cursor;
                          Source              : in out Tree;
                          Source_Parent      : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Target)
             or else raise Program_Error) and then
  (not Tampering_With_Cursors_Prohibited (Source)
   or else raise Program_Error) and then
  (Target_Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Target, Target_Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Target, Before)
   or else raise Program_Error) and then
  (Source_Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Source, Source_Parent)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Parent (Target, Before) /= Target_Parent
   or else raise Constraint_Error) and then
  (Target'Has_Same_Storage (Source) or else
   Target_Parent = Source_Parent or else
   Is_Ancessor_Of (Target, Source_Parent, Target_Parent)
   or else raise Constraint_Error),
  Post => (declare
           Org_Child_Count renames
             Child_Count (Source, Source_Parent)'Old;
           Org_Target_Count renames Node_Count (Target)'Old;
begin
  (if not Target'Has_Same_Storage (Source) then
   Node_Count (Target) = Org_Target_Count +
   Org_Child_Count and then
   Node_Count (Source) = Node_Count (Source)'Old -
   Org_Child_Count
  else
   Node_Count (Target) = Org_Target_Count));

```

If Source denotes the same object as Target, then:

- if Target\_Parent equals Source\_Parent there is no effect; else
- the child elements (and the further descendants) of Source\_Parent are moved to be child elements of Target\_Parent. If Target\_Parent already has child elements, then the moved elements are inserted prior to the node designated by Before, or, if Before equals No\_Element, the moved elements are inserted after the last existing child node of Target\_Parent. The parent of each moved child element is set to Target\_Parent.

Otherwise (if Source does not denote the same object as Target), the child elements (and the further descendants) of Source\_Parent are removed from Source and moved to Target. The child elements are inserted as children of Target\_Parent. If Target\_Parent already has child elements, then the moved elements are inserted prior to the node designated by Before, or, if Before equals No\_Element, the moved elements are inserted after the last existing child node of Target\_Parent. In each of these cases, the overall count of Target is incremented by Subtree\_Node\_Count (Source\_Parent)-1, and the overall count of Source is decremented by Subtree\_Node\_Count (Source\_Parent)-1.



```

procedure Splice_Children (Container      : in out Tree;
                           Target_Parent  : in      Cursor;
                           Before         : in      Cursor;
                           Source_Parent  : in      Cursor)
with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
  (Target_Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Container, Target_Parent)
   or else raise Program_Error) and then
  (Meaningful_For (Container, Before)
   or else raise Program_Error) and then
  (Source_Parent /= No_Element
   or else raise Constraint_Error) and then
  (Meaningful_For (Container, Source_Parent)
   or else raise Program_Error) and then
  (Before = No_Element or else
   Parent (Container, Before) /= Target_Parent
   or else raise Constraint_Error) and then
  (Target_Parent = Source_Parent or else
   Is_Ancessor_Of (Container, Source_Parent, Target_Parent)
   or else raise Constraint_Error),
  Post => Node_Count (Container) = Node_Count (Container)'Old;

```

If Target\_Parent equals Source\_Parent there is no effect. Otherwise, the child elements (and the further descendants) of Source\_Parent are moved to be child elements of Target\_Parent. If Target\_Parent already has child elements, then the moved elements are inserted prior to the node designated by Before, or, if Before equals No\_Element, the moved elements are inserted after the last existing child node of Target\_Parent. The parent of each moved child element is set to Target\_Parent.

```

function Parent (Position : Cursor) return Cursor
with Nonblocking, Global => in all, Use_Formal => null,
  Post => (if Position = No_Element or else
           Is_Root (Position) then Parent'Result = No_Element);

```

Returns a cursor designating the parent node of the node designated by Position.

```

function Parent (Container : Tree;
                 Position   : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
  Pre  => Meaningful_For (Container, Position)
   or else raise Program_Error,
  Post => (if Position = No_Element or else
           Is_Root (Container, Position)
           then Parent'Result = No_Element
           else Has_Element (Container, Parent'Result));

```

Returns a cursor designating the parent node of the node designated by Position in Container.

```

function First_Child (Parent : Cursor) return Cursor
with Nonblocking, Global => in all, Use_Formal => null,
  Pre  => Parent /= No_Element or else raise Constraint_Error;

```

First\_Child returns a cursor designating the first child node of the node designated by Parent; if there is no such node, No\_Element is returned.

```

function First_Child (Container : Tree;
                    Parent      : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
    Pre  => (Parent /= No_Element
            or else raise Constraint_Error) and then
            (Meaningful_For (Container, Parent)
            or else raise Program_Error),
    Post => First_Child'Result = No_Element or else
            Has_Element (Container, First_Child'Result);

```

First\_Child returns a cursor designating the first child node of the node designated by Parent in Container; if there is no such node, No\_Element is returned.

```

function First_Child_Element (Parent : Cursor) return Element_Type
with Nonblocking, Global => in all, Use_Formal => Element_Type,
    Pre  => (Parent /= No_Element and then
            Last_Child (Parent) /= No_Element)
            or else raise Constraint_Error;

```

Equivalent to Element (First\_Child (Parent)).

```

function First_Child_Element (Container : Tree;
                    Parent      : Cursor) return Element_Type
with Nonblocking, Global => null, Use_Formal => Element_Type,
    Pre  => (Parent /= No_Element
            or else raise Constraint_Error) and then
            (Meaningful_For (Container, Parent)
            or else raise Program_Error) and then
            (First_Child (Container, Parent) /= No_Element
            or else raise Constraint_Error);

```

Equivalent to Element (Container, First\_Child (Container, Parent)).

```

function Last_Child (Parent : Cursor) return Cursor
with Nonblocking, Global => in all, Use_Formal => null,
    Pre  => Parent /= No_Element or else raise Constraint_Error;

```

Last\_Child returns a cursor designating the last child node of the node designated by Parent; if there is no such node, No\_Element is returned.

```

function Last_Child (Container : Tree;
                    Parent      : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
    Pre  => (Parent /= No_Element
            or else raise Constraint_Error) and then
            (Meaningful_For (Container, Parent)
            or else raise Program_Error),
    Post => Last_Child'Result = No_Element or else
            Has_Element (Container, Last_Child'Result);

```

Last\_Child returns a cursor designating the last child node of the node designated by Parent in Container; if there is no such node, No\_Element is returned.

```

function Last_Child_Element (Parent : Cursor) return Element_Type
with Nonblocking, Global => in all, Use_Formal => Element_Type,
    Pre  => (Parent /= No_Element and then
            Last_Child (Parent) /= No_Element)
            or else raise Constraint_Error;

```

Equivalent to Element (Last\_Child (Parent)).

```

function Last_Child_Element (Container : Tree;
                             Parent      : Cursor) return Element_Type
with Nonblocking, Global => null, Use_Formal => Element_Type,
    Pre => (Parent /= No_Element
            or else raise Constraint_Error) and then
    (Meaningful_For (Container, Parent)
     or else raise Program_Error) and then
    (Last_Child (Container, Parent) /= No_Element
     or else raise Constraint_Error);

```

Equivalent to Element (Container, Last\_Child (Container, Parent)).

```

function Next_Sibling (Position : Cursor) return Cursor
with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element
             then Next_Sibling'Result = No_Element);

```

If Position equals No\_Element or designates the last child node of its parent, then Next\_Sibling returns the value No\_Element. Otherwise, it returns a cursor that designates the successor (with the same parent) of the node designated by Position.

```

function Next_Sibling (Container : Tree;
                       Position  : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
    Pre => Meaningful_For (Container, Position)
            or else raise Program_Error,
    Post => (if Next_Sibling'Result = No_Element then
             Position = No_Element or else
             Is_Root (Container, Position) or else
             Last_Child (Container, Parent (Container, Position))
             = Position
            else Has_Element (Container, Next_Sibling'Result));

```

Next\_Sibling returns a cursor that designates the successor (with the same parent) of the node designated by Position in Container.

```

function Previous_Sibling (Position : in out Cursor)
with Nonblocking, Global => in all, Use_Formal => null,
    Post => (if Position = No_Element
             then Previous_Sibling'Result = No_Element);

```

If Position equals No\_Element or designates the first child node of its parent, then Previous\_Sibling returns the value No\_Element. Otherwise, it returns a cursor that designates the predecessor (with the same parent) of the node designated by Position.

```

function Previous_Sibling (Container : Tree;
                           Position  : Cursor) return Cursor
with Nonblocking, Global => null, Use_Formal => null,
    Pre => Meaningful_For (Container, Position)
            or else raise Program_Error,
    Post => (if Previous_Sibling'Result = No_Element then
             Position = No_Element or else
             Is_Root (Container, Position) or else
             First_Child (Container, Parent (Container, Position))
             = Position
            else Has_Element (Container, Previous_Sibling'Result));

```

Previous\_Sibling returns a cursor that designates the predecessor (with the same parent) of the node designated by Position in Container.

```

procedure Next_Sibling (Position : in out Cursor)
with Nonblocking, Global => in all, Use_Formal => null;

```

Equivalent to Position := Next\_Sibling (Position);

```

procedure Next_Sibling (Container : in      Tree;
                        Position  : in out Cursor)
with Nonblocking, Global => null, Use_Formal => null,
    Pre  => Meaningful_For (Container, Position)
    or else raise Program_Error,
    Post => (if Position /= No_Element
            then Has_Element (Container, Position));

```

Equivalent to Position := Next\_Sibling (Container, Position);

```

procedure Previous_Sibling (Position : in out Cursor)
with Nonblocking, Global => in all, Use_Formal => null;

```

Equivalent to Position := Previous\_Sibling (Position);

```

procedure Previous_Sibling (Container : in      Tree;
                        Position  : in out Cursor)
with Nonblocking, Global => null, Use_Formal => null,
    Pre  => Meaningful_For (Container, Position)
    or else raise Program_Error,
    Post => (if Position /= No_Element
            then Has_Element (Container, Position));

```

Equivalent to Position := Previous\_Sibling (Container, Position);

```

procedure Iterate_Children
    (Parent  : in Cursor;
     Process : not null access procedure (Position : in Cursor))
with Allows_Exit,
    Pre  => Parent /= No_Element or else raise Constraint_Error,
    Global => in all, Use_Formal => null;

```

Iterate\_Children calls Process.all with a cursor that designates each child node of Parent, starting with the first child node and moving the cursor as per the Next\_Sibling function.

Tampering with the cursors of the tree containing Parent is prohibited during the execution of a call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Iterate_Children
    (Container : in Tree;
     Parent    : in Cursor;
     Process   : not null access procedure (Position : in Cursor))
with Allows_Exit,
    Pre  => (Parent /= No_Element
            or else raise Constraint_Error) and then
            (Meaningful_For (Container, Parent)
            or else raise Program_Error);

```

Iterate\_Children calls Process.all with a cursor that designates each child node of Container and Parent, starting with the first child node and moving the cursor as per the Next\_Sibling function.

Tampering with the cursors of the tree containing Parent is prohibited during the execution of a call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Reverse_Iterate_Children
    (Parent  : in Cursor;
     Process : not null access procedure (Position : in Cursor))
with Allows_Exit,
    Pre  => Parent /= No_Element or else raise Constraint_Error,
    Global => in all, Use_Formal => null;

```

Reverse\_Iterate\_Children calls Process.all with a cursor that designates each child node of Parent, starting with the last child node and moving the cursor as per the Previous\_Sibling function.

Tampering with the cursors of the tree containing Parent is prohibited during the execution of a call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Reverse_Iterate_Children
  (Container : in Tree;
   Parent    : in Cursor;
   Process    : not null access procedure (Position : in Cursor))
with Allows_Exit,
   Pre => (Parent /= No_Element
          or else raise Constraint_Error) and then
      (Meaningful_For (Container, Parent)
       or else raise Program_Error);

```

Reverse\_Iterate\_Children calls Process.all with a cursor that designates each child node of Container and Parent, starting with the last child node and moving the cursor as per the Previous\_Sibling function.

Tampering with the cursors of the tree containing Parent is prohibited during the execution of a call on Process.all. Any exception raised by Process.all is propagated.

```

function Iterate_Children (Container : in Tree; Parent : in Cursor)
return Tree_Iterator_Interfaces.Parallel_Reversible_Iterator_Class
with Pre => (Parent /= No_Element
          or else raise Constraint_Error) and then
      (Meaningful_For (Container, Parent)
       or else raise Program_Error),
   Post => Tampering_With_Cursors_Prohibited (Container);

```

Iterate\_Children returns an iterator object (see 8.5.1) that will generate a value for a loop parameter (see 8.5.2) designating each child node of Parent. When used as a forward iterator, the nodes are designated starting with the first child node and moving the cursor as per the function Next\_Sibling; when used as a reverse iterator, the nodes are designated starting with the last child node and moving the cursor as per the function Previous\_Sibling; when used as a parallel iterator, processing all child nodes concurrently. Tampering with the cursors of Container is prohibited while the iterator object exists (in particular, in the sequence\_of\_statements of the loop statement whose iterator\_specification denotes this object). The iterator object needs finalization.

The nested package Multiway\_Trees.Stable provides a type Stable.Tree that represents a *stable* tree, which is one that cannot grow and shrink. Such a tree can be created by calling the Copy function, or by establishing a *stabilized view* of an ordinary tree.

The subprograms of package Containers.Multiway\_Trees that have a parameter or result of type tree are included in the nested package Stable with the same specification, except that the following are omitted:

Tampering\_With\_Cursors\_Prohibited, Tampering\_With\_Elements\_Prohibited, Assign, Move, Clear, Delete\_Leaf, Insert\_Child, Delete\_Children, Delete\_Subtree, Copy\_Subtree, Copy\_Local\_Subtree, Splice\_Subtree, and Splice\_Children

The operations of this package are equivalent to those for ordinary trees, except that the calls to Tampering\_With\_Cursors\_Prohibited and Tampering\_With\_Elements\_Prohibited that occur in preconditions are replaced by False, and any that occur in postconditions are replaced by True.

If a stable tree is declared with the Base discriminant designating a pre-existing ordinary tree, the stable tree represents a stabilized view of the underlying ordinary tree, and any operation on the stable tree is reflected on the underlying ordinary tree. While a stabilized view exists, any operation that tampers with elements performed on the underlying tree is prohibited. The finalization of a stable tree that provides such a view removes this restriction on the underlying ordinary tree (though some other restriction can exist due to other concurrent iterations or stabilized views).

If a stable tree is declared without specifying Base, the object is necessarily initialized. The initializing expression of the stable tree, typically a call on Copy, determines the Node\_Count of the tree. The Node\_Count of a stable tree never changes after initialization.

*Bounded (Run-Time) Errors*

It is a bounded error for the actual function associated with a generic formal subprogram, when called as part of an operation of this package, to tamper with elements of any Tree parameter of the operation. Either Program\_Error is raised, or the operation works as defined on the value of the Tree either prior to, or subsequent to, some or all of the modifications to the Tree.

It is a bounded error to call any subprogram declared in the visible part of Containers.Multiway\_Trees when the associated container has been finalized. If the operation takes Container as an **in out** parameter, then it raises Constraint\_Error or Program\_Error. Otherwise, the operation either proceeds as it would for an empty container, or it raises Constraint\_Error or Program\_Error.

*Erroneous Execution*

A Cursor value is *invalid* if any of the following have occurred since it was created:

- The tree that contains the element it designates has been finalized;
- The tree that contains the element it designates has been used as the Source or Target of a call to Move;
- The tree that contains the element it designates has been used as the Target of a call to Assign or the target of an assignment\_statement;
- The element it designates has been removed from the tree that previously contained the element.

The result of "=" or Has\_Element is unspecified if it is called with an invalid cursor parameter. Execution is erroneous if any other subprogram declared in Containers.Multiway\_Trees is called with an invalid cursor parameter.

Execution is erroneous if the tree associated with the result of a call to Reference or Constant\_Reference is finalized before the result object returned by the call to Reference or Constant\_Reference is finalized.

*Implementation Requirements*

No storage associated with a multiway tree object shall be lost upon assignment or scope exit.

The execution of an assignment\_statement for a tree shall have the effect of copying the elements from the source tree object to the target tree object and changing the node count of the target object to that of the source object.

*Implementation Advice*

Containers.Multiway\_Trees should be implemented similarly to a multiway tree. In particular, if  $N$  is the overall number of nodes for a particular tree, then the worst-case time complexity of Element, Parent, First\_Child, Last\_Child, Next\_Sibling, Previous\_Sibling, Insert\_Child with Count=1, and Delete should be  $O(\log N)$ .

Move should not copy elements, and should minimize copying of internal data structures.

If an exception is propagated from a tree operation, no storage should be lost, nor any elements removed from a tree unless specified by the operation.

### A.18.11 The Generic Package Containers.Indefinite\_Vectors

The language-defined generic package Containers.Indefinite\_Vectors provides a private type Vector and a set of operations. It provides the same operations as the package Containers.Vectors (see A.18.2), with the difference that the generic formal Element\_Type is indefinite.

*Static Semantics*

The declaration of the generic library package Containers.Indefinite\_Vectors has the same contents and semantics as Containers.Vectors except:

- The generic formal Element\_Type is indefinite.



- The procedures with the profiles:

```

procedure Insert (Container : in out Vector;
                  Before   : in      Extended_Index;
                  Count     : in      Count_Type := 1);

procedure Insert (Container : in out Vector;
                  Before     : in      Cursor;
                  Position    : out     Cursor;
                  Count       : in      Count_Type := 1);

```

are omitted.

- The actual Element parameter of access subprogram Process of Update\_Element may be constrained even if Element\_Type is unconstrained.
- The operations "&", Append, Insert, Prepend, Replace\_Element, and To\_Vector that have a formal parameter of type Element\_Type perform indefinite insertion (see A.18).
- The description of Tampering\_With\_Elements\_Prohibited is replaced by:
 

Returns True if tampering with elements is prohibited for Container, and False otherwise.
- Tampering\_With\_Cursors\_Prohibited is replaced by Tampering\_With\_Elements\_Prohibited in the postcondition for the operations Reference and Constant\_Reference.
- The operations Replace\_Element, Reverse\_Elements, and Swap, and the nested generic unit Generic\_Sorting are omitted from the nested package Stable.

#### A.18.12 The Generic Package Containers.Indefinite\_Doubly\_Linked\_Lists

The language-defined generic package Containers.Indefinite\_Doubly\_Linked\_Lists provides private types List and Cursor, and a set of operations for each type. It provides the same operations as the package Containers.Doubly\_Linked\_Lists (see A.18.3), with the difference that the generic formal Element\_Type is indefinite.

##### Static Semantics

The declaration of the generic library package Containers.Indefinite\_Doubly\_Linked\_Lists has the same contents and semantics as Containers.Doubly\_Linked\_Lists except:

- The generic formal Element\_Type is indefinite.
- The procedure with the profile:

```

procedure Insert (Container : in out List;
                  Before     : in      Cursor;
                  Position    : out     Cursor;
                  Count       : in      Count_Type := 1);

```

is omitted.

- The actual Element parameter of access subprogram Process of Update\_Element may be constrained even if Element\_Type is unconstrained.
- The operations Append, Insert, Prepend, and Replace\_Element that have a formal parameter of type Element\_Type perform indefinite insertion (see A.18).
- The description of Tampering\_With\_Elements\_Prohibited is replaced by:
 

Returns True if tampering with elements is prohibited for Container, and False otherwise.
- Tampering\_With\_Cursors\_Prohibited is replaced by Tampering\_With\_Elements\_Prohibited in the postcondition for the operations Reference and Constant\_Reference.
- The operations Replace\_Element and Swap are omitted from the nested package Stable.

### A.18.13 The Generic Package Containers.Indefinite\_Hashed\_Maps

The language-defined generic package Containers.Indefinite\_Hashed\_Maps provides a map with the same operations as the package Containers.Hashed\_Maps (see A.18.5), with the difference that the generic formal types Key\_Type and Element\_Type are indefinite.

#### Static Semantics

The declaration of the generic library package Containers.Indefinite\_Hashed\_Maps has the same contents and semantics as Containers.Hashed\_Maps except:

- The generic formal Key\_Type is indefinite.
- The generic formal Element\_Type is indefinite.
- The procedure with the profile:
 

```

procedure Insert (Container : in out Map;
                  Key       : in      Key_Type;
                  Position  :         out Cursor;
                  Inserted  :         out Boolean);
      
```

 is omitted.
- The actual Element parameter of access subprogram Process of Update\_Element may be constrained even if Element\_Type is unconstrained.
- The operations Include, Insert, Replace, and Replace\_Element that have a formal parameter of type Element\_Type perform indefinite insertion (see A.18).
- The description of Tampering\_With\_Elements\_Prohibited is replaced by:
 

Returns True if tampering with elements is prohibited for Container, and False otherwise.
- Tampering\_With\_Cursors\_Prohibited is replaced by Tampering\_With\_Elements\_Prohibited in the postcondition for the operations Reference and Constant\_Reference.
- The operations Replace and Replace\_Element are omitted from the nested package Stable.

### A.18.14 The Generic Package Containers.Indefinite\_Ordered\_Maps

The language-defined generic package Containers.Indefinite\_Ordered\_Maps provides a map with the same operations as the package Containers.Ordered\_Maps (see A.18.6), with the difference that the generic formal types Key\_Type and Element\_Type are indefinite.

#### Static Semantics

The declaration of the generic library package Containers.Indefinite\_Ordered\_Maps has the same contents and semantics as Containers.Ordered\_Maps except:

- The generic formal Key\_Type is indefinite.
- The generic formal Element\_Type is indefinite.
- The procedure with the profile:
 

```

procedure Insert (Container : in out Map;
                  Key       : in      Key_Type;
                  Position  :         out Cursor;
                  Inserted  :         out Boolean);
      
```

 is omitted.
- The actual Element parameter of access subprogram Process of Update\_Element may be constrained even if Element\_Type is unconstrained.
- The operations Include, Insert, Replace, and Replace\_Element that have a formal parameter of type Element\_Type perform indefinite insertion (see A.18).
- The description of Tampering\_With\_Elements\_Prohibited is replaced by:

Returns True if tampering with elements is prohibited for Container, and False otherwise.

- Tampering\_With\_Cursors\_Prohibited is replaced by Tampering\_With\_Elements\_Prohibited in the postcondition for the operations Reference and Constant\_Reference.
- The operations Replace and Replace\_Element are omitted from the nested package Stable.

#### A.18.15 The Generic Package Containers.Indefinite\_Hashed\_Sets

The language-defined generic package Containers.Indefinite\_Hashed\_Sets provides a set with the same operations as the package Containers.Hashed\_Sets (see A.18.8), with the difference that the generic formal type Element\_Type is indefinite.

##### Static Semantics

The declaration of the generic library package Containers.Indefinite\_Hashed\_Sets has the same contents and semantics as Containers.Hashed\_Sets except:

- The generic formal Element\_Type is indefinite.
- The actual Element parameter of access subprogram Process of Update\_Element\_Preserving\_Key may be constrained even if Element\_Type is unconstrained.
- The operations Include, Insert, Replace, Replace\_Element, and To\_Set that have a formal parameter of type Element\_Type perform indefinite insertion (see A.18).

#### A.18.16 The Generic Package Containers.Indefinite\_Ordered\_Sets

The language-defined generic package Containers.Indefinite\_Ordered\_Sets provides a set with the same operations as the package Containers.Ordered\_Sets (see A.18.9), with the difference that the generic formal type Element\_Type is indefinite.

##### Static Semantics

The declaration of the generic library package Containers.Indefinite\_Ordered\_Sets has the same contents and semantics as Containers.Ordered\_Sets except:

- The generic formal Element\_Type is indefinite.
- The actual Element parameter of access subprogram Process of Update\_Element\_Preserving\_Key may be constrained even if Element\_Type is unconstrained.
- The operations Include, Insert, Replace, Replace\_Element, and To\_Set that have a formal parameter of type Element\_Type perform indefinite insertion (see A.18).

#### A.18.17 The Generic Package Containers.Indefinite\_Multiway\_Trees

The language-defined generic package Containers.Indefinite\_Multiway\_Trees provides a multiway tree with the same operations as the package Containers.Multiway\_Trees (see A.18.10), with the difference that the generic formal Element\_Type is indefinite.

##### Static Semantics

The declaration of the generic library package Containers.Indefinite\_Multiway\_Trees has the same contents and semantics as Containers.Multiway\_Trees except:

- The generic formal Element\_Type is indefinite.
- The procedure with the profile:

```
procedure Insert_Child (Container : in out Tree;
                       Parent    : in      Cursor;
                       Before    : in      Cursor;
                       Position  :          out Cursor;
                       Count     : in      Count_Type := 1);
```

is omitted.

- The actual Element parameter of access subprogram Process of Update\_Element may be constrained even if Element\_Type is unconstrained.
- The operations Append\_Child, Insert\_Child, Prepend\_Child, and Replace\_Element that have a formal parameter of type Element\_Type perform indefinite insertion (see A.18).
- The description of Tampering\_With\_Elements\_Prohibited is replaced by:
 

Returns True if tampering with elements is prohibited for Container, and False otherwise.
- Tampering\_With\_Cursors\_Prohibited is replaced by Tampering\_With\_Elements\_Prohibited in the postcondition for the operations Reference and Constant\_Reference.
- The operations Replace\_Element and Swap are omitted from the nested package Stable.

### A.18.18 The Generic Package Containers.Indefinite\_Holders

The language-defined generic package Containers.Indefinite\_Holders provides a private type Holder and a set of operations for that type. A holder container holds a single element of an indefinite type.

A holder container allows the declaration of an object that can be used like an uninitialized variable or component of an indefinite type.

A holder container may be *empty*. An empty holder does not contain an element.

#### Static Semantics

The generic library package Containers.Indefinite\_Holders has the following declaration:

```

generic
  type Element_Type (<>) is private;
  with function "=" (Left, Right : Element_Type) return Boolean is <>;
package Ada.Containers.Indefinite_Holders
  with Preelaborate, Remote_Types,
    Nonblocking, Global => in out synchronized is

  type Holder is tagged private
    with Stable_Properties => (Is_Empty,
      Tampering_With_The_Element_Prohibited),
    Default_Initial_Condition => Is_Empty (Holder),
    Preelaborable_Initialization;

  Empty_Holder : constant Holder;

  function Equal_Element (Left, Right : Element_Type) return Boolean
    renames "=";

  function "=" (Left, Right : Holder) return Boolean;

  function Tampering_With_The_Element_Prohibited
    (Container : Holder) return Boolean
    with Nonblocking, Global => null, Use_Formal => null;

  function Empty return Holder
    is (Empty_Holder)
    with Post =>
      not Tampering_With_The_Element_Prohibited (Empty'Result)
      and then Is_Empty (Empty'Result);

  function To_Holder (New_Item : Element_Type) return Holder
    with Post => not Is_Empty (To_Holder'Result);

  function Is_Empty (Container : Holder) return Boolean
    with Global => null, Use_Formal => null;

  procedure Clear (Container : in out Holder)
    with Pre => not Tampering_With_The_Element_Prohibited (Container)
      or else raise Program_Error,
      Post => Is_Empty (Container);
  
```

```

function Element (Container : Holder) return Element_Type
  with Pre => not Is_Empty (Container) or else raise Constraint_Error,
        Global => null, Use_Formal => Element_Type;

procedure Replace_Element (Container : in out Holder;
                           New_Item : in Element_Type)
  with Pre => not Tampering_With_The_Element_Prohibited (Container)
        or else raise Program_Error,
        Post => not Is_Empty (Container);

procedure Query_Element
  (Container : in Holder;
   Process   : not null access procedure (Element : in Element_Type))
  with Pre => not Is_Empty (Container) or else raise Constraint_Error;

procedure Update_Element
  (Container : in out Holder;
   Process   : not null access procedure (Element : in out Element_Type))
  with Pre => not Is_Empty (Container) or else raise Constraint_Error;

type Constant_Reference_Type
  (Element : not null access constant Element_Type) is private
  with Implicit_Dereference => Element,
        Nonblocking, Global => in out synchronized,
        Default_Initial_Condition => (raise Program_Error);

type Reference_Type
  (Element : not null access Element_Type) is private
  with Implicit_Dereference => Element,
        Nonblocking, Global => in out synchronized,
        Default_Initial_Condition => (raise Program_Error);

function Constant_Reference (Container : aliased in Holder)
  return Constant_Reference_Type
  with Pre => not Is_Empty (Container)
        or else raise Constraint_Error,
        Post => Tampering_With_The_Element_Prohibited (Container),
        Nonblocking, Global => null, Use_Formal => null;

function Reference (Container : aliased in out Holder)
  return Reference_Type
  with Pre => not Is_Empty (Container)
        or else raise Constraint_Error,
        Post => Tampering_With_The_Element_Prohibited (Container),
        Nonblocking, Global => null, Use_Formal => null;

procedure Assign (Target : in out Holder; Source : in Holder)
  with Post => (Is_Empty (Source) = Is_Empty (Target));

function Copy (Source : Holder) return Holder
  with Post => (Is_Empty (Source) = Is_Empty (Copy'Result));

procedure Move (Target : in out Holder; Source : in out Holder)
  with Pre => (not Tampering_With_The_Element_Prohibited (Target)
        or else raise Program_Error) and then
        (not Tampering_With_The_Element_Prohibited (Source)
        or else raise Program_Error),
        Post => (if not Target'Has_Same_Storage (Source) then
        Is_Empty (Source) and then (not Is_Empty (Target))));

procedure Swap (Left, Right : in out Holder)
  with Pre => (not Tampering_With_The_Element_Prohibited (Left)
        or else raise Program_Error) and then
        (not Tampering_With_The_Element_Prohibited (Right)
        or else raise Program_Error),
        Post => Is_Empty (Left) = Is_Empty (Right)'Old and then
        Is_Empty (Right) = Is_Empty (Left)'Old;

private
  ... -- not specified by the language

```

```
end Ada.Containers.Indefinite_Holders;
```

The actual function for the generic formal function "=" on `Element_Type` values is expected to define a reflexive and symmetric relationship and return the same result value each time it is called with a particular pair of values. If it behaves in some other manner, the function "=" on holder values returns an unspecified value. The exact arguments and number of calls of this generic formal function by the function "=" on holder values are unspecified.

The type `Holder` is used to represent holder containers. The type `Holder` needs finalization (see 10.6).

`Empty_Holder` represents an empty holder object. If an object of type `Holder` is not otherwise initialized, it is initialized to the same value as `Empty_Holder`.

Some operations check for "tampering with the element" of a container because they depend on the element of the container not being replaced. When tampering with the element is *prohibited* for a particular holder object *H*, `Program_Error` is propagated by the finalization of *H*, as well as by a call that passes *H* to certain of the operations of this package, as indicated by the precondition of such an operation.

```
function "=" (Left, Right : Holder) return Boolean;
```

If `Left` and `Right` denote the same holder object, then the function returns `True`. Otherwise, it compares the element contained in `Left` to the element contained in `Right` using the generic formal equality operator, returning the result of that operation. Any exception raised during the evaluation of element equality is propagated.

```
function Tampering_With_The_Element_Prohibited  
(Container : Holder) return Boolean  
with Nonblocking, Global => null, Use_Formal => null;
```

Returns `True` if tampering with the element is currently prohibited for `Container`, and returns `False` otherwise.

```
function To_Holder (New_Item : Element_Type) return Holder  
with Post => not Is_Empty (To_Holder'Result);
```

Returns a nonempty holder containing an element initialized to `New_Item`. `To_Holder` performs indefinite insertion (see A.18).

```
function Is_Empty (Container : Holder) return Boolean  
with Global => null, Use_Formal => null;
```

Returns `True` if `Container` is empty, and `False` if it contains an element.

```
procedure Clear (Container : in out Holder)  
with Pre => not Tampering_With_The_Element_Prohibited (Container)  
          or else raise Program_Error,  
      Post => Is_Empty (Container);
```

Removes the element from `Container`.

```
function Element (Container : Holder) return Element_Type  
with Pre => not Is_Empty (Container) or else raise Constraint_Error,  
      Global => null, Use_Formal => Element_Type;
```

Returns the element stored in `Container`.

```
procedure Replace_Element (Container : in out Holder;  
                          New_Item : in Element_Type)  
with Pre  => not Tampering_With_The_Element_Prohibited (Container)  
          or else raise Program_Error,  
      Post => not Is_Empty (Container);
```

`Replace_Element` assigns the value `New_Item` into `Container`, replacing any preexisting content of `Container`; `Replace_Element` performs indefinite insertion (see A.18).



```

procedure Query_Element
(Container : in Holder;
Process   : not null access procedure (Element : in Element_Type))
with Pre => not Is_Empty (Container) or else raise Constraint_Error,
      Global => null, Use_Formal => null;

```

Query\_Element calls Process.all with the contained element as the argument. Tampering with the element of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

procedure Update_Element
(Container : in out Holder;
Process   : not null access procedure (Element : in out Element_Type))
with Pre => not Is_Empty (Container) or else raise Constraint_Error;

```

Update\_Element calls Process.all with the contained element as the argument. Tampering with the element of Container is prohibited during the execution of the call on Process.all. Any exception raised by Process.all is propagated.

```

type Constant_Reference_Type
(Element : not null access constant Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

type Reference_Type (Element : not null access Element_Type) is private
with Implicit_Dereference => Element,
      Nonblocking, Global => in out synchronized,
      Default_Initial_Condition => (raise Program_Error);

```

The types Constant\_Reference\_Type and Reference\_Type need finalization.

```

function Constant_Reference (Container : aliased in Holder)
return Constant_Reference_Type
with Pre  => not Is_Empty (Container) or else raise Constraint_Error,
      Post => Tampering_With_The_Element_Prohibited (Container),
      Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Implicit\_Dereference aspect) provides a convenient way to gain read access to the contained element of a holder container.

Constant\_Reference returns an object whose discriminant is an access value that designates the contained element. Tampering with the element of Container is prohibited while the object returned by Constant\_Reference exists and has not been finalized.

```

function Reference (Container : aliased in out Holder)
return Reference_Type
with Pre  => not Is_Empty (Container) or else raise Constraint_Error,
      Post => Tampering_With_The_Element_Prohibited (Container),
      Nonblocking, Global => null, Use_Formal => null;

```

This function (combined with the Implicit\_Dereference aspects) provides a convenient way to gain read and write access to the contained element of a holder container.

Reference returns an object whose discriminant is an access value that designates the contained element. Tampering with the element of Container is prohibited while the object returned by Reference exists and has not been finalized.

```

procedure Assign (Target : in out Holder; Source : in Holder)
with Post => (Is_Empty (Source) = Is_Empty (Target));

```

If Target denotes the same object as Source, the operation has no effect. If Source is empty, Clear (Target) is called. Otherwise, Replace\_Element (Target, Element (Source)) is called.

```
function Copy (Source : Holder) return Holder
  with Post => (Is_Empty (Source) = Is_Empty (Copy'Result));
```

If Source is empty, returns an empty holder container; otherwise, returns To\_Holder (Element (Source)).

```
procedure Move (Target : in out Holder; Source : in out Holder)
  with Pre  => (not Tampering_With_The_Element_Prohibited (Target)
    or else raise Program_Error) and then
    (not Tampering_With_The_Element_Prohibited (Source)
    or else raise Program_Error),
  Post => (if not Target'Has_Same_Storage (Source) then
    Is_Empty (Source) and then (not Is_Empty (Target)));
```

If Target denotes the same object as Source, then the operation has no effect. Otherwise, the element contained by Source (if any) is removed from Source and inserted into Target, replacing any preexisting content.

```
procedure Swap (Left, Right : in out Holder)
  with Pre  => (not Tampering_With_The_Element_Prohibited (Left)
    or else raise Program_Error) and then
    (not Tampering_With_The_Element_Prohibited (Right)
    or else raise Program_Error),
  Post => Is_Empty (Left) = Is_Empty (Right)'Old and then
    Is_Empty (Right) = Is_Empty (Left)'Old;
```

If Left denotes the same object as Right, then the operation has no effect. Otherwise, operation exchanges the elements (if any) contained by Left and Right.

#### *Bounded (Run-Time) Errors*

It is a bounded error for the actual function associated with a generic formal subprogram, when called as part of an operation of this package, to tamper with the element of any Holder parameter of the operation. Either Program\_Error is raised, or the operation works as defined on the value of the Holder either prior to, or subsequent to, some or all of the modifications to the Holder.

It is a bounded error to call any subprogram declared in the visible part of Containers.Indefinite\_Holders when the associated container has been finalized. If the operation takes Container as an **in out** parameter, then it raises Constraint\_Error or Program\_Error. Otherwise, the operation either proceeds as it would for an empty container, or it raises Constraint\_Error or Program\_Error.

#### *Erroneous Execution*

Execution is erroneous if the holder container associated with the result of a call to Reference or Constant\_Reference is finalized before the result object returned by the call to Reference or Constant\_Reference is finalized.

#### *Implementation Requirements*

No storage associated with a holder object shall be lost upon assignment or scope exit.

The execution of an assignment\_statement for a holder container shall have the effect of copying the element (if any) from the source holder object to the target holder object.

#### *Implementation Advice*

Move and Swap should not copy any elements, and should minimize copying of internal data structures.

If an exception is propagated from a holder operation, no storage should be lost, nor should the element be removed from a holder container unless specified by the operation.

### A.18.19 The Generic Package Containers.Bounded\_Vectors

The language-defined generic package Containers.Bounded\_Vectors provides a private type Vector and a set of operations. It provides the same operations as the package Containers.Vectors (see A.18.2), with the difference that the maximum storage is bounded.

#### Static Semantics

The declaration of the generic library package Containers.Bounded\_Vectors has the same contents and semantics as Containers.Vectors except:

- The aspect Preelaborate is replaced with aspect Pure. Aspect Global is deleted.
- The type Vector is declared with a discriminant that specifies the capacity:

```
type Vector (Capacity : Count_Type) is tagged private...
```

- The aspect\_definition for Preelaborable\_Initialization for type Vector is changed to:

```
Preelaborable_Initialization =>
  Element_Type'Preelaborable_Initialization
```

- The type Vector needs finalization if and only if type Element\_Type needs finalization.
- Capacity is omitted from the Stable\_Properties of type Vector.
- In function Empty, the postcondition is altered to:

```
Post =>
  Empty'Result.Capacity = Capacity and then
  not Tampering_With_Elements_Prohibited (Empty'Result) and then
  not Tampering_With_Cursors_Prohibited (Empty'Result) and then
  Length (Empty'Result) = 0;
```

- In function Copy, the postcondition is altered to:

```
Post => Length (Copy'Result) = Length (Source) and then
  (if Capacity > Length (Source) then
    Copy'Result.Capacity = Capacity
  else Copy'Result.Capacity >= Length (Source));
```

- The description of Reserve\_Capacity is replaced with:

```
procedure Reserve_Capacity (Container : in out Vector;
  Capacity : in Count_Type)
with Pre => Capacity <= Container.Capacity
or else raise Capacity_Error;
```

This operation has no effect, other than checking the precondition.

- The portion of the postcondition checking the capacity is omitted from subprograms Set\_Length, Assign, Insert, Insert\_Space, Prepend, Append, and Delete.
- For procedures Insert, Insert\_Space, Prepend, and Append, the part of the precondition reading:

```
(<some length> <= Maximum_Length - <some other length>
or else raise Constraint_Error)
```

is replaced by:

```
(<some length> <= Maximum_Length - <some other length>
or else raise Constraint_Error) and then
(<some length> <= Container.Capacity - <some other length>
or else raise Capacity_Error)
```

#### Bounded (Run-Time) Errors

It is a bounded error to assign from a bounded vector object while tampering with elements or cursors of that object is prohibited. Either Program\_Error is raised by the assignment, execution proceeds with the target object prohibiting tampering with elements or cursors, or execution proceeds normally.

*Erroneous Execution*

When a bounded vector object *V* is finalized, if tampering with cursors is prohibited for *V* other than due to an assignment from another vector, then execution is erroneous.

*Implementation Requirements*

For each instance of Containers.Vectors and each instance of Containers.Bounded\_Vectors, if the two instances meet the following conditions, then the output generated by the Vector'Output or Vector'Write subprograms of either instance shall be readable by the Vector'Input or Vector'Read of the other instance, respectively:

- the Element\_Type parameters of the two instances are statically matching subtypes of the same type; and
- the output generated by Element\_Type'Output or Element\_Type'Write is readable by Element\_Type'Input or Element\_Type'Read, respectively (where Element\_Type denotes the type of the two actual Element\_Type parameters); and
- the preceding two conditions also hold for the Index\_Type parameters of the instances.

*Implementation Advice*

Bounded vector objects should be implemented without implicit pointers or dynamic allocation. The implementation advice for procedure Move to minimize copying does not apply.

### A.18.20 The Generic Package Containers.Bounded\_Doubly\_Linked\_Lists

The language-defined generic package Containers.Bounded\_Doubly\_Linked\_Lists provides a private type List and a set of operations. It provides the same operations as the package Containers.Doubly\_Linked\_Lists (see A.18.3), with the difference that the maximum storage is bounded.

*Static Semantics*

The declaration of the generic library package Containers.Bounded\_Doubly\_Linked\_Lists has the same contents and semantics as Containers.Doubly\_Linked\_Lists except:

- The aspect Preelaborate is replaced with aspect Pure. Aspect Global is deleted.
- The type List is declared with a discriminant that specifies the capacity (maximum number of elements) as follows:

```
type List (Capacity : Count_Type) is tagged private...
```

- The aspect definition for Preelaborable\_Initialization for type List is changed to:

```
Preelaborable_Initialization =>
  Element_Type'Preelaborable_Initialization
```

- The type List needs finalization if and only if type Element\_Type needs finalization.
- The function Empty is replaced by:

```
function Empty (Capacity : Count_Type := implementation-defined)
return List
with Post =>
  Empty'Result.Capacity = Capacity and then
  not Tampering_With_Elements_Prohibited (Empty'Result) and then
  not Tampering_With_Cursors_Prohibited (Empty'Result) and then
  Length (Empty'Result) = 0;
```

- For procedures Insert, Prepend, Append, Merge, and the three-parameter Splice whose parameter Source has type List, the part of the precondition reading:

```
(<some length> <= Count_Type'Last - <some other length>
or else raise Constraint_Error)
```

is replaced by:

```

(<some length> <= Count_Type'Last - <some other length>
  or else raise Constraint_Error) and then
(<some length> <= Container.Capacity - <some other length>
  or else raise Capacity_Error)

```

- In procedure Assign, the precondition is altered to:

```

Pre => (not Tampering_With_Cursors_Prohibited (Target)
  or else raise Program_Error) and then
  (Length (Source) <= Target.Capacity
    or else raise Capacity_Error),

```

- The function Copy is replaced with:

```

function Copy (Source : List; Capacity : Count_Type := 0)
  return List
  with Pre => Capacity = 0 or else Capacity >= Length (Source)
    or else raise Capacity_Error,
  Post =>
    Length (Copy'Result) = Length (Source) and then
    not Tampering_With_Elements_Prohibited (Copy'Result) and then
    not Tampering_With_Cursors_Prohibited (Copy'Result) and then
    Copy'Result.Capacity = (if Capacity = 0 then
      Length (Source) else Capacity);

```

Returns a list whose elements have the same values as the elements of Source.

- In the four-parameter procedure Splice, the precondition is altered to:

```

Pre => (not Tampering_With_Cursors_Prohibited (Target)
  or else raise Program_Error) and then
  (not Tampering_With_Cursors_Prohibited (Source)
    or else raise Program_Error) and then
  (Position /= No_Element
    or else raise Constraint_Error) and then
  (Has_Element (Source, Position)
    or else raise Program_Error) and then
  (Before = No_Element or else Has_Element (Target, Before)
    or else raise Program_Error) and then
  (Target'Has_Same_Storage (Source) or else
    Length (Target) /= Count_Type'Last
    or else raise Constraint_Error) and then
  (Target'Has_Same_Storage (Source) or else
    Length (Target) /= Target.Capacity
    or else raise Capacity_Error),

```

#### *Bounded (Run-Time) Errors*

It is a bounded error to assign from a bounded list object while tampering with elements or cursors of that object is prohibited. Either Program\_Error is raised by the assignment, execution proceeds with the target object prohibiting tampering with elements or cursors, or execution proceeds normally.

#### *Erroneous Execution*

When a bounded list object *L* is finalized, if tampering with cursors is prohibited for *L* other than due to an assignment from another list, then execution is erroneous.

#### *Implementation Requirements*

For each instance of Containers.Doubly\_Linked\_Lists and each instance of Containers.Bounded\_Doubly\_Linked\_Lists, if the two instances meet the following conditions, then the output generated by the List'Output or List'Write subprograms of either instance shall be readable by the List'Input or List'Read of the other instance, respectively:

- the Element\_Type parameters of the two instances are statically matching subtypes of the same type; and

- the output generated by `Element_Type'Output` or `Element_Type'Write` is readable by `Element_Type'Input` or `Element_Type'Read`, respectively (where `Element_Type` denotes the type of the two actual `Element_Type` parameters).

*Implementation Advice*

Bounded list objects should be implemented without implicit pointers or dynamic allocation.

The implementation advice for procedure `Move` to minimize copying does not apply.

### A.18.21 The Generic Package Containers.Bounded\_Hashed\_Maps

The language-defined generic package `Containers.Bounded_Hashed_Maps` provides a private type `Map` and a set of operations. It provides the same operations as the package `Containers.Hashed_Maps` (see A.18.5), with the difference that the maximum storage is bounded.

*Static Semantics*

The declaration of the generic library package `Containers.Bounded_Hashed_Maps` has the same contents and semantics as `Containers.Hashed_Maps` except:

- The aspect `Preelaborate` is replaced with aspect `Pure`. Aspect `Global` is deleted.
- The type `Map` is declared with discriminants that specify both the capacity (number of elements) and modulus (number of distinct hash values) of the hash table as follows:

```
type Map (Capacity : Count_Type;
          Modulus   : Hash_Type) is tagged private...
```

- The aspect `_definition` for `Preelaborable_Initialization` for type `Map` is changed to:

```
Preelaborable_Initialization =>
  Element_Type'Preelaborable_Initialization
  and
  Key_Type'Preelaborable_Initialization
```

- The type `Map` needs finalization if and only if type `Key_Type` or type `Element_Type` needs finalization.
- In function `Empty`, the postcondition is altered to:

```
Post =>
  Empty'Result.Capacity = Capacity and then
  Empty'Result.Modulus = Default_Modulus (Capacity) and then
  not Tampering_With_Elements_Prohibited (Empty'Result) and then
  not Tampering_With_Cursors_Prohibited (Empty'Result) and then
  Length (Empty'Result) = 0;
```

- The description of `Reserve_Capacity` is replaced with:

```
procedure Reserve_Capacity (Container : in out Map;
                           Capacity  : in      Count_Type)
  with Pre => Capacity <= Container.Capacity
       or else raise Capacity_Error;
```

This operation has no effect, other than checking the precondition.

- An additional operation is added immediately following `Reserve_Capacity`:

```
function Default_Modulus (Capacity : Count_Type) return Hash_Type;
```

`Default_Modulus` returns an implementation-defined value for the number of distinct hash values to be used for the given capacity (maximum number of elements).

- For procedures `Insert` and `Include`, the part of the precondition reading:

```
(<some length> <= Count_Type'Last - <some other length>
  or else raise Constraint_Error)
```

is replaced by:



```

(<some length> <= Count_Type'Last - <some other length>
  or else raise Constraint_Error) and then
(<some length> > Container.Capacity - <some other length>
  or else raise Capacity_Error)

```

— In procedure Assign, the precondition is altered to:

```

Pre => (not Tampering_With_Cursors_Prohibited (Target)
  or else raise Program_Error) and then
  (Length (Source) <= Target.Capacity
  or else raise Capacity_Error),

```

— The function Copy is replaced with:

```

function Copy (Source : Map;
  Capacity : Count_Type := 0;
  Modulus : Hash_Type := 0) return Map
with Pre => Capacity = 0 or else Capacity >= Length (Source)
  or else raise Capacity_Error,
  Post =>
    Length (Copy'Result) = Length (Source) and then
    not Tampering_With_Elements_Prohibited (Copy'Result) and then
    not Tampering_With_Cursors_Prohibited (Copy'Result) and then
    Copy'Result.Capacity = (if Capacity = 0 then
      Length (Source) else Capacity) and then
    Copy'Result.Modulus = (if Modulus = 0 then
      Default_Modulus (Capacity) else Modulus);

```

Returns a map with key/element pairs initialized from the values in Source.

#### *Bounded (Run-Time) Errors*

It is a bounded error to assign from a bounded map object while tampering with elements or cursors of that object is prohibited. Either Program\_Error is raised by the assignment, execution proceeds with the target object prohibiting tampering with elements or cursors, or execution proceeds normally.

#### *Erroneous Execution*

When a bounded map object *M* is finalized, if tampering with cursors is prohibited for *M* other than due to an assignment from another map, then execution is erroneous.

#### *Implementation Requirements*

For each instance of Containers.Hashing\_Maps and each instance of Containers.Bounded\_Hashing\_Maps, if the two instances meet the following conditions, then the output generated by the Map'Output or Map'Write subprograms of either instance shall be readable by the Map'Input or Map'Read of the other instance, respectively:

- the Element\_Type parameters of the two instances are statically matching subtypes of the same type; and
- the output generated by Element\_Type'Output or Element\_Type'Write is readable by Element\_Type'Input or Element\_Type'Read, respectively (where Element\_Type denotes the type of the two actual Element\_Type parameters); and
- the preceding two conditions also hold for the Key\_Type parameters of the instances.

#### *Implementation Advice*

Bounded hashed map objects should be implemented without implicit pointers or dynamic allocation. The implementation advice for procedure Move to minimize copying does not apply.

### **A.18.22 The Generic Package Containers.Bounded\_Ordered\_Maps**

The language-defined generic package Containers.Bounded\_Ordered\_Maps provides a private type Map and a set of operations. It provides the same operations as the package Containers.Ordered\_Maps (see A.18.6), with the difference that the maximum storage is bounded.

## Static Semantics

The declaration of the generic library package Containers.Bounded\_Ordered\_Maps has the same contents and semantics as Containers.Ordered\_Maps except:

- The aspect Preelaborate is replaced with aspect Pure. Aspect Global is deleted.
- The type Map is declared with a discriminant that specifies the capacity (maximum number of elements) as follows:

```
type Map (Capacity : Count_Type) is tagged private...
```

- The aspect\_definition for Preelaborable\_Initialization for type Map is changed to:

```
Preelaborable_Initialization =>
  Element_Type'Preelaborable_Initialization
  and
  Key_Type'Preelaborable_Initialization
```

- The type Map needs finalization if and only if type Key\_Type or type Element\_Type needs finalization.

- The function Empty is replaced by:

```
function Empty (Capacity : Count_Type := implementation-defined)
return Map
with Post =>
  Empty'Result.Capacity = Capacity and then
  not Tampering_With_Elements_Prohibited (Empty'Result) and then
  not Tampering_With_Cursors_Prohibited (Empty'Result) and then
  Length (Empty'Result) = 0;
```

- For procedures Insert and Include, the part of the precondition reading:

```
(<some length> <= Count_Type'Last - <some other length>
  or else raise Constraint_Error)
```

is replaced by:

```
(<some length> <= Count_Type'Last - <some other length>
  or else raise Constraint_Error) and then
(<some length> <= Container.Capacity - <some other length>
  or else raise Capacity_Error)
```

- In procedure Assign, the precondition is altered to:

```
Pre => (not Tampering_With_Cursors_Prohibited (Target)
  or else raise Program_Error) and then
  (Length (Source) <= Target.Capacity
  or else raise Capacity_Error),
```

- The function Copy is replaced with:

```
function Copy (Source : Map;
  Capacity : Count_Type := 0) return Map
with Pre => Capacity = 0 or else Capacity >= Length (Source)
  or else raise Capacity_Error,
  Post =>
  Length (Copy'Result) = Length (Source) and then
  not Tampering_With_Elements_Prohibited (Copy'Result) and then
  not Tampering_With_Cursors_Prohibited (Copy'Result) and then
  Copy'Result.Capacity = (if Capacity = 0 then
    Length (Source) else Capacity);
```

Returns a map with key/element pairs initialized from the values in Source.

## Bounded (Run-Time) Errors

It is a bounded error to assign from a bounded map object while tampering with elements or cursors of that object is prohibited. Either Program\_Error is raised by the assignment, execution proceeds with the target object prohibiting tampering with elements or cursors, or execution proceeds normally.

*Erroneous Execution*

When a bounded map object *M* is finalized, if tampering with cursors is prohibited for *M* other than due to an assignment from another map, then execution is erroneous.

*Implementation Requirements*

For each instance of `Containers.Ordered_Maps` and each instance of `Containers.Bounded_Ordered_Maps`, if the two instances meet the following conditions, then the output generated by the `Map'Output` or `Map'Write` subprograms of either instance shall be readable by the `Map'Input` or `Map'Read` of the other instance, respectively:

- the `Element_Type` parameters of the two instances are statically matching subtypes of the same type; and
- the output generated by `Element_Type'Output` or `Element_Type'Write` is readable by `Element_Type'Input` or `Element_Type'Read`, respectively (where `Element_Type` denotes the type of the two actual `Element_Type` parameters); and
- the preceding two conditions also hold for the `Key_Type` parameters of the instances.

*Implementation Advice*

Bounded ordered map objects should be implemented without implicit pointers or dynamic allocation. The implementation advice for procedure `Move` to minimize copying does not apply.

### A.18.23 The Generic Package `Containers.Bounded_Hashed_Sets`

The language-defined generic package `Containers.Bounded_Hashed_Sets` provides a private type `Set` and a set of operations. It provides the same operations as the package `Containers.Hashed_Sets` (see A.18.8), with the difference that the maximum storage is bounded.

*Static Semantics*

The declaration of the generic library package `Containers.Bounded_Hashed_Sets` has the same contents and semantics as `Containers.Hashed_Sets` except:

- The aspect `Preelaborate` is replaced with aspect `Pure`. Aspect `Global` is deleted.
- The type `Set` is declared with discriminants that specify both the capacity (number of elements) and modulus (number of distinct hash values) of the hash table as follows:

```
type Set (Capacity : Count_Type;
          Modulus   : Hash_Type) is tagged private...
```

- The aspect\_definition for `Preelaborable_Initialization` for type `Set` is changed to:

```
Preelaborable_Initialization =>
  Element_Type'Preelaborable_Initialization
```

- The type `Set` needs finalization if and only if type `Element_Type` needs finalization.
- In function `Empty`, the postcondition is altered to:

```
Post =>
  Empty'Result.Capacity = Capacity and then
  Empty'Result.Modulus = Default_Modulus (Capacity) and then
  not Tampering_With_Cursors_Prohibited (Empty'Result) and then
  Length (Empty'Result) = 0;
```

- The description of `Reserve_Capacity` is replaced with:

```
procedure Reserve_Capacity (Container : in out Set;
                           Capacity  : in      Count_Type)
with Pre => Capacity <= Container.Capacity
or else raise Capacity_Error;
```

This operation has no effect, other than checking the precondition.

- An additional operation is added immediately following `Reserve_Capacity`:

```
function Default_Modulus (Capacity : Count_Type) return Hash_Type;
```

Default\_Modulus returns an implementation-defined value for the number of distinct hash values to be used for the given capacity (maximum number of elements).

- For procedures Insert and Include, the part of the precondition reading:

```
(<some length> <= Count_Type'Last - <some other length>
  or else raise Constraint_Error)
```

is replaced by:

```
(<some length> <= Count_Type'Last - <some other length>
  or else raise Constraint_Error) and then
(<some length> <= Container.Capacity - <some other length>
  or else raise Capacity_Error)
```

- In procedure Assign, the precondition is altered to:

```
Pre => (not Tampering_With_Cursors_Prohibited (Target)
  or else raise Program_Error) and then
(Length (Source) <= Target.Capacity
  or else raise Capacity_Error),
```

- The function Copy is replaced with:

```
function Copy (Source : Set;
  Capacity : Count_Type := 0;
  Modulus : Hash_Type := 0) return Map
with Pre => Capacity = 0 or else Capacity >= Length (Source)
  or else raise Capacity_Error,
  Post =>
    Length (Copy'Result) = Length (Source) and then
    not Tampering_With_Cursors_Prohibited (Copy'Result) and then
    Copy'Result.Capacity = (if Capacity = 0 then
      Length (Source) else Capacity) and then
    Copy'Result.Modulus = (if Modulus = 0 then
      Default_Modulus (Capacity) else Modulus);
```

Returns a set with key/element pairs initialized from the values in Source.

#### *Bounded (Run-Time) Errors*

It is a bounded error to assign from a bounded set object while tampering with elements or cursors of that object is prohibited. Either Program\_Error is raised by the assignment, execution proceeds with the target object prohibiting tampering with elements or cursors, or execution proceeds normally.

#### *Erroneous Execution*

When a bounded set object *S* is finalized, if tampering with cursors is prohibited for *S* other than due to an assignment from another set, then execution is erroneous.

#### *Implementation Requirements*

For each instance of Containers.Hash\_Sets and each instance of Containers.Bounded\_Hash\_Sets, if the two instances meet the following conditions, then the output generated by the Set'Output or Set'Write subprograms of either instance shall be readable by the Set'Input or Set'Read of the other instance, respectively:

- the Element\_Type parameters of the two instances are statically matching subtypes of the same type; and
- the output generated by Element\_Type'Output or Element\_Type'Write is readable by Element\_Type'Input or Element\_Type'Read, respectively (where Element\_Type denotes the type of the two actual Element\_Type parameters).

#### *Implementation Advice*

Bounded hashed set objects should be implemented without implicit pointers or dynamic allocation.

The implementation advice for procedure Move to minimize copying does not apply.

#### A.18.24 The Generic Package Containers.Bounded\_Ordered\_Sets

The language-defined generic package Containers.Bounded\_Ordered\_Sets provides a private type Set and a set of operations. It provides the same operations as the package Containers.Ordered\_Sets (see A.18.9), with the difference that the maximum storage is bounded.

##### Static Semantics

The declaration of the generic library package Containers.Bounded\_Ordered\_Sets has the same contents and semantics as Containers.Ordered\_Sets except:

- The aspect Preelaborate is replaced with aspect Pure. Aspect Global is deleted.
- The type Set is declared with a discriminant that specifies the capacity (maximum number of elements) as follows:

```
type Set (Capacity : Count_Type) is tagged private...
```

- The aspect\_definition for Preelaborable\_Initialization for type Set is changed to:

```
Preelaborable_Initialization =>
  Element_Type'Preelaborable_Initialization
```

- The type Set needs finalization if and only if type Element\_Type needs finalization.
- The function Empty is replaced by:

```
function Empty (Capacity : Count_Type := implementation-defined)
return Set
with Post =>
  Empty'Result.Capacity = Capacity and then
  not Tampering_With_Cursors_Prohibited (Empty'Result) and then
  Length (Empty'Result) = 0;
```

- For procedures Insert and Include, the part of the precondition reading:

```
(<some length> <= Count_Type.Last - <some other length>
or else raise Constraint_Error)
```

is replaced by:

```
(<some length> <= Count_Type'Last - <some other length>
or else raise Constraint_Error) and then
(<some length> <= Container.Capacity - <some other length>
or else raise Capacity_Error)
```

- In procedure Assign, the precondition is altered to:

```
Pre => (not Tampering_With_Cursors_Prohibited (Target)
or else raise Program_Error) and then
Length (Source) <= Target.Capacity
or else raise Capacity_Error),
```

- The function Copy is replaced with:

```
function Copy (Source : Set;
  Capacity : Count_Type := 0) return Map
with Pre => Capacity = 0 or else Capacity >= Length (Source)
  or else raise Capacity_Error,
  Post =>
    Length (Copy'Result) = Length (Source) and then
    not Tampering_With_Cursors_Prohibited (Copy'Result) and then
    Copy'Result.Capacity = (if Capacity = 0 then
      Length (Source) else Capacity);
```

Returns a set with key/element pairs initialized from the values in Source.

*Bounded (Run-Time) Errors*

It is a bounded error to assign from a bounded set object while tampering with elements or cursors of that object is prohibited. Either Program\_Error is raised by the assignment, execution proceeds with the target object prohibiting tampering with elements or cursors, or execution proceeds normally.

*Erroneous Execution*

When a bounded set object *S* is finalized, if tampering with cursors is prohibited for *S* other than due to an assignment from another set, then execution is erroneous.

*Implementation Requirements*

For each instance of Containers.Ordered\_Sets and each instance of Containers.Bounded\_Ordered\_Sets, if the two instances meet the following conditions, then the output generated by the Set'Output or Set'Write subprograms of either instance shall be readable by the Set'Input or Set'Read of the other instance, respectively:

- the Element\_Type parameters of the two instances are statically matching subtypes of the same type; and
- the output generated by Element\_Type'Output or Element\_Type'Write is readable by Element\_Type'Input or Element\_Type'Read, respectively (where Element\_Type denotes the type of the two actual Element\_Type parameters).

*Implementation Advice*

Bounded ordered set objects should be implemented without implicit pointers or dynamic allocation.

The implementation advice for procedure Move to minimize copying does not apply.

### A.18.25 The Generic Package Containers.Bounded\_Multiway\_Trees

The language-defined generic package Containers.Bounded\_Multiway\_Trees provides a private type Tree and a set of operations. It provides the same operations as the package Containers.Multiway\_Trees (see A.18.10), with the difference that the maximum storage is bounded.

*Static Semantics*

The declaration of the generic library package Containers.Bounded\_Multiway\_Trees has the same contents and semantics as Containers.Multiway\_Trees except:

- The aspect Preelaborate is replaced with aspect Pure. Aspect Global is deleted.
- The type Tree is declared with a discriminant that specifies the capacity (maximum number of elements) as follows:

```
type Tree (Capacity : Count_Type) is tagged private...
```

- The aspect\_definition for Preelaborable\_Initialization for type Tree is changed to:

```
Preelaborable_Initialization =>
  Element_Type'Preelaborable_Initialization
```

- The type Tree needs finalization if and only if type Element\_Type needs finalization.
- The function Empty is replaced by:

```
function Empty (Capacity : Count_Type := implementation-defined)
return Tree
with Post =>
  Empty'Result.Capacity = Capacity and then
  not Tampering_With_Elements_Prohibited (Empty'Result) and then
  not Tampering_With_Cursors_Prohibited (Empty'Result) and then
  Node_Count (Empty'Result) = 1;
```

- For procedures Insert\_Child, Prepend\_Child, and Append\_Child, the initial subexpression of the precondition is replaced with:



```

with Pre => (not Tampering_With_Cursors_Prohibited (Container)
             or else raise Program_Error) and then
(Node_Count (Container) - 1 <= Container.Capacity - Count
             or else raise Capacity_Error)

```

- In procedure Assign, the precondition is altered to:

```

Pre => (not Tampering_With_Cursors_Prohibited (Target)
        or else raise Program_Error) and then
(Node_Count (Source) - 1 <= Target.Capacity
        or else raise Capacity_Error),

```

- Function Copy is declared as follows:

```

function Copy (Source : Tree; Capacity : Count_Type := 0)
return Tree
with Pre => Capacity = 0 or else Capacity >= Node_Count (Source) - 1
             or else raise Capacity_Error,
Post =>
Node_Count (Copy'Result) = Node_Count (Source) and then
not Tampering_With_Elements_Prohibited (Copy'Result) and then
not Tampering_With_Cursors_Prohibited (Copy'Result) and then
Copy'Result.Capacity = (if Capacity = 0 then
Node_Count (Source) - 1 else Capacity);

```

Returns a list whose elements have the same values as the elements of Source.

- In the four-parameter procedure Copy\_Subtree, the last **or else** of the precondition is replaced by:

```

(not Is_Root (Source)
 or else raise Constraint_Error) and then
(Node_Count (Target) - 1 + Subtree_Node_Count (Source) <=
Target.Capacity
 or else raise Capacity_Error),

```

- In the five-parameter procedure Copy\_Subtree, the last **or else** of the precondition is replaced by:

```

(not Is_Root (Source, Subtree)
 or else raise Constraint_Error) and then
(Node_Count (Target) - 1 +
Subtree_Node_Count (Source, Subtree) <= Target.Capacity
 or else raise Capacity_Error),

```

- In Copy\_Local\_Subtree, the last **or else** of the precondition is replaced by:

```

(not Is_Root (Source, Subtree)
 or else raise Constraint_Error) and then
(Node_Count (Target) - 1 +
Subtree_Node_Count (Target, Source) <= Target.Capacity
 or else raise Capacity_Error),

```

- In the five-parameter procedure Splice\_Subtree, the penultimate **or else** of the precondition is replaced by:

```

(Has_Element (Source, Position)
 or else raise Program_Error) and then
(Target'Has_Same_Storage (Source) or else
Node_Count (Target) - 1 +
Subtree_Node_Count (Source, Position) <= Target.Capacity
 or else raise Capacity_Error) and then

```

- In the five-parameter procedure Splice\_Children, the penultimate **elsif** of the precondition is replaced by:

```

(Before = No_Element or else
Parent (Target, Before) /= Target_Parent
 or else raise Constraint_Error) and then
(Target'Has_Same_Storage (Source) or else
Node_Count (Target) - 1 +
Child_Count (Source, Source_Parent) <= Target.Capacity
 or else raise Capacity_Error) and then

```

*Bounded (Run-Time) Errors*

It is a bounded error to assign from a bounded tree object while tampering with elements or cursors of that object is prohibited. Either Program\_Error is raised by the assignment, execution proceeds with the target object prohibiting tampering with elements or cursors, or execution proceeds normally.

*Erroneous Execution*

When a bounded tree object *T* is finalized, if tampering with cursors is prohibited for *T* other than due to an assignment from another tree, then execution is erroneous.

*Implementation Requirements*

For each instance of Containers.Multiway\_Trees and each instance of Containers.Bounded\_Multiway\_Trees, if the two instances meet the following conditions, then the output generated by the Tree'Output or Tree'Write subprograms of either instance shall be readable by the Tree'Input or Tree'Read of the other instance, respectively:

- the Element\_Type parameters of the two instances are statically matching subtypes of the same type; and
- the output generated by Element\_Type'Output or Element\_Type'Write is readable by Element\_Type'Input or Element\_Type'Read, respectively (where Element\_Type denotes the type of the two actual Element\_Type parameters).

*Implementation Advice*

Bounded tree objects should be implemented without implicit pointers or dynamic allocation.

The implementation advice for procedure Move to minimize copying does not apply.

### A.18.26 Array Sorting

The language-defined generic procedures Containers.Generic\_Array\_Sort, Containers.Generic\_Constrained\_Array\_Sort, and Containers.Generic\_Sort provide sorting on arbitrary array types.

*Static Semantics*

The generic library procedure Containers.Generic\_Array\_Sort has the following declaration:

```
generic
  type Index_Type is (<>);
  type Element_Type is private;
  type Array_Type is array (Index_Type range <>) of Element_Type;
  with function "<" (Left, Right : Element_Type)
    return Boolean is <>;
procedure Ada.Containers.Generic_Array_Sort (Container : in out Array_Type)
  with Pure, Nonblocking, Global => null;
```

Reorders the elements of Container such that the elements are sorted smallest first as determined by the generic formal "<" operator provided. Any exception raised during evaluation of "<" is propagated.

The actual function for the generic formal function "<" of Generic\_Array\_Sort is expected to return the same value each time it is called with a particular pair of element values. It should define a strict weak ordering relationship (see A.18); it should not modify Container. If the actual for "<" behaves in some other manner, the behavior of the instance of Generic\_Array\_Sort is unspecified. The number of times Generic\_Array\_Sort calls "<" is unspecified.

The generic library procedure `Containers.Generic_Constrained_Array_Sort` has the following declaration:

```
generic
  type Index_Type is (<>);
  type Element_Type is private;
  type Array_Type is array (Index_Type) of Element_Type;
  with function "<" (Left, Right : Element_Type)
    return Boolean is <>;
procedure Ada.Containers.Generic_Constrained_Array_Sort
  (Container : in out Array_Type)
  with Pure, Nonblocking, Global => null;
```

Reorders the elements of Container such that the elements are sorted smallest first as determined by the generic formal "<" operator provided. Any exception raised during evaluation of "<" is propagated.

The actual function for the generic formal function "<" of `Generic_Constrained_Array_Sort` is expected to return the same value each time it is called with a particular pair of element values. It should define a strict weak ordering relationship (see A.18); it should not modify Container. If the actual for "<" behaves in some other manner, the behavior of the instance of `Generic_Constrained_Array_Sort` is unspecified. The number of times `Generic_Constrained_Array_Sort` calls "<" is unspecified.

The generic library procedure `Containers.Generic_Sort` has the following declaration:

```
generic
  type Index_Type is (<>);
  with function Before (Left, Right : Index_Type) return Boolean;
  with procedure Swap (Left, Right : in Index_Type);
procedure Ada.Containers.Generic_Sort
  (First, Last : Index_Type'Base)
  with Pure, Nonblocking, Global => null;
```

Reorders the elements of an indexable structure, over the range First .. Last, such that the elements are sorted in the ordering determined by the generic formal function Before; Before should return True if Left is to be sorted before Right. The generic formal Before compares the elements having the given indices, and the generic formal Swap exchanges the values of the indicated elements. Any exception raised during evaluation of Before or Swap is propagated.

The actual function for the generic formal function Before of `Generic_Sort` is expected to return the same value each time it is called with index values that identify a particular pair of element values. It should define a strict weak ordering relationship (see A.18); it should not modify the elements. The actual function for the generic formal Swap should exchange the values of the indicated elements. If the actual for either Before or Swap behaves in some other manner, the behavior of `Generic_Sort` is unspecified. The number of times the `Generic_Sort` calls Before or Swap is unspecified.

#### *Implementation Advice*

The worst-case time complexity of a call on an instance of `Containers.Generic_Array_Sort` or `Containers.Generic_Constrained_Array_Sort` should be  $O(N^{**2})$  or better, and the average time complexity should be better than  $O(N^{**2})$ , where  $N$  is the length of the Container parameter.

`Containers.Generic_Array_Sort` and `Containers.Generic_Constrained_Array_Sort` should minimize copying of elements.

The worst-case time complexity of a call on an instance of `Containers.Generic_Sort` should be  $O(N^{**2})$  or better, and the average time complexity should be better than  $O(N^{**2})$ , where  $N$  is the difference between the Last and First parameters plus 1.

`Containers.Generic_Sort` should minimize calls to the generic formal Swap.

### A.18.27 The Generic Package Containers.Synchronized\_Queue\_Interfaces

The language-defined generic package Containers.Synchronized\_Queue\_Interfaces provides interface type Queue, and a set of operations for that type. Interface Queue specifies a first-in, first-out queue.

#### Static Semantics

The generic library package Containers.Synchronized\_Queue\_Interfaces has the following declaration:

```

generic
  type Element_Type is private;
package Ada.Containers.Synchronized_Queue_Interfaces
  with Pure, Nonblocking, Global => null is

  type Queue is synchronized interface;

  procedure Enqueue
    (Container : in out Queue;
     New_Item  : in      Element_Type) is abstract
    with Synchronization => By_Entry,
         Nonblocking => False,
         Global'Class=> in out synchronized;

  procedure Dequeue
    (Container : in out Queue;
     Element   : out Element_Type) is abstract
    with Synchronization => By_Entry,
         Nonblocking => False,
         Global'Class=> in out synchronized;

  function Current_Use (Container : Queue) return Count_Type is abstract
    with Nonblocking, Global'Class=> null, Use_Formal => null;
  function Peak_Use (Container : Queue) return Count_Type is abstract
    with Nonblocking, Global'Class=> null, Use_Formal => null,
         Post'Class=> Peak_Use'Result >= Current_Use (Container);

end Ada.Containers.Synchronized_Queue_Interfaces;

```

The subprogram behavior descriptions given below are the semantics for the corresponding callable entities found in the language-defined generic packages that have a formal package named Queue\_Interfaces.

```

procedure Enqueue
  (Container : in out Queue;
   New_Item  : in      Element_Type) is abstract
  with Synchronization => By_Entry
       Nonblocking => False,
       Global'Class=> in out synchronized;

```

A queue type that implements this interface is allowed to have a bounded *capacity*. If the queue object has a bounded capacity, and the number of existing elements equals the capacity, then Enqueue blocks until storage becomes available; otherwise, Enqueue does not block. In any case, it then copies New\_Item onto the queue.

```

procedure Dequeue
  (Container : in out Queue;
   Element   : out Element_Type) is abstract
  with Synchronization => By_Entry
       Nonblocking => False,
       Global'Class=> in out synchronized;

```

If the queue is empty, then Dequeue blocks until an item becomes available. In any case, it then assigns the element at the head of the queue to Element, and removes it from the queue.

```

function Current_Use (Container : Queue) return Count_Type is abstract
  with Nonblocking, Global'Class=> null, Use_Formal => null;

```

Returns the number of elements currently in the queue.

```

function Peak_Use (Container : Queue) return Count_Type is abstract
  with Nonblocking, Global'Class=> null, Use_Formal => null,
    Post'Class => Peak_Use'Result >= Current_Use (Container);

```

Returns the maximum number of elements that have been in the queue at any one time.

NOTE Unlike other language-defined containers, there are no queues whose element types are indefinite. Elements of an indefinite type can be handled by defining the element of the queue to be a holder container (see A.18.18) of the indefinite type, or to be an explicit access type that designates the indefinite type.

### A.18.28 The Generic Package Containers.Unbounded\_Synchronized\_Queues

*Static Semantics*

The language-defined generic package Containers.Unbounded\_Synchronized\_Queues provides type Queue, which implements the interface type Containers.Synchronized\_Queue\_Interfaces.Queue.

```

with System;
with Ada.Containers.Synchronized_Queue_Interfaces;
generic
  with package Queue_Interfaces is
    new Ada.Containers.Synchronized_Queue_Interfaces (<>);
    Default_Ceiling : System.Any_Priority := System.Priority'Last;
package Ada.Containers.Unbounded_Synchronized_Queues
  with Preelaborate,
    Nonblocking, Global => in out synchronized is

  package Implementation is
    ... -- not specified by the language
  end Implementation;

  protected type Queue
    (Ceiling : System.Any_Priority := Default_Ceiling)
    with Priority => Ceiling is
    new Queue_Interfaces.Queue with

    overriding
    entry Enqueue (New_Item : in Queue_Interfaces.Element_Type);
    overriding
    entry Dequeue (Element : out Queue_Interfaces.Element_Type);

    overriding
    function Current_Use return Count_Type
      with Nonblocking, Global => null, Use_Formal => null;
    overriding
    function Peak_Use return Count_Type
      with Nonblocking, Global => null, Use_Formal => null;

  private
    ... -- not specified by the language
  end Queue;

private
  ... -- not specified by the language
end Ada.Containers.Unbounded_Synchronized_Queues;

```

The type Queue is used to represent task-safe queues.

The capacity for instances of type Queue is unbounded.

### A.18.29 The Generic Package Containers.Bounded\_Synchronized\_Queues

*Static Semantics*

The language-defined generic package Containers.Bounded\_Synchronized\_Queues provides type Queue, which implements the interface type Containers.Synchronized\_Queue\_Interfaces.Queue.

```

with System;
with Ada.Containers.Synchronized_Queue_Interfaces;
generic
  with package Queue_Interfaces is
    new Ada.Containers.Synchronized_Queue_Interfaces (<>);
    Default_Capacity : Count_Type;
    Default_Ceiling   : System.Any_Priority := System.Priority'Last;
package Ada.Containers.Bounded_Synchronized_Queues
  with Preelaborate,
    Nonblocking, Global => in out synchronized is

package Implementation is
  ... -- not specified by the language
end Implementation;

protected type Queue
  (Capacity : Count_Type := Default_Capacity;
   Ceiling   : System.Any_Priority := Default_Ceiling)
  with Priority => Ceiling is
    new Queue_Interfaces.Queue with

    overriding
    entry Enqueue (New_Item : in Queue_Interfaces.Element_Type);
    overriding
    entry Dequeue (Element : out Queue_Interfaces.Element_Type);

    overriding
    function Current_Use return Count_Type
      with Nonblocking, Global => null, Use_Formal => null;
    overriding
    function Peak_Use return Count_Type
      with Nonblocking, Global => null, Use_Formal => null;

private
  ... -- not specified by the language
end Queue;

private
  ... -- not specified by the language
end Ada.Containers.Bounded_Synchronized_Queues;

```

The semantics are the same as for Unbounded\_Synchronized\_Queues, except:

- The capacity for instances of type Queue is bounded and specified by the discriminant Capacity.

*Implementation Advice*

Bounded queue objects should be implemented without implicit pointers or dynamic allocation.

### A.18.30 The Generic Package Containers.Unbounded\_Priority\_Queues

*Static Semantics*

The language-defined generic package Containers.Unbounded\_Priority\_Queues provides type Queue, which implements the interface type Containers.Synchronized\_Queue\_Interfaces.Queue.



```

with System;
with Ada.Containers.Synchronized_Queue_Interfaces;
generic
  with package Queue_Interfaces is
    new Ada.Containers.Synchronized_Queue_Interfaces (<>);
  type Queue_Priority is private;
  with function Get_Priority
    (Element : Queue_Interfaces.Element_Type) return Queue_Priority is <>;
  with function Before
    (Left, Right : Queue_Priority) return Boolean is <>;
  Default_Ceiling : System.Any_Priority := System.Priority'Last;
package Ada.Containers.Unbounded_Priority_Queues
  with Preelaborate,
    Nonblocking, Global => in out synchronized is

  package Implementation is
    ... -- not specified by the language
  end Implementation;

  protected type Queue
    (Ceiling : System.Any_Priority := Default_Ceiling)
      with Priority => Ceiling is
    new Queue_Interfaces.Queue with

    overriding
      entry Enqueue (New_Item : in Queue_Interfaces.Element_Type);
    overriding
      entry Dequeue (Element : out Queue_Interfaces.Element_Type);

    not overriding
      procedure Dequeue_Only_High_Priority
        (At_Least : in Queue_Priority;
         Element : in out Queue_Interfaces.Element_Type;
         Success : out Boolean);

    overriding
      function Current_Use return Count_Type
        with Nonblocking, Global => null, Use_Formal => null;
    overriding
      function Peak_Use return Count_Type
        with Nonblocking, Global => null, Use_Formal => null;

  private
    ... -- not specified by the language
  end Queue;

private
  ... -- not specified by the language
end Ada.Containers.Unbounded_Priority_Queues;

```

The type `Queue` is used to represent task-safe priority queues.

The capacity for instances of type `Queue` is unbounded.

Two elements *E1* and *E2* are equivalent if `Before(Get_Priority(E1), Get_Priority(E2))` and `Before(Get_Priority(E2), Get_Priority(E1))` both return `False`.

The actual functions for `Get_Priority` and `Before` are expected to return the same value each time they are called with the same actuals, and should not modify their actuals. `Before` should define a strict weak ordering relationship (see A.18). If the actual functions behave in some other manner, the behavior of `Unbounded_Priority_Queues` is unspecified.

`Enqueue` inserts an item according to the order specified by the `Before` function on the result of `Get_Priority` on the elements; `Before` should return `True` if `Left` is to be inserted before `Right`. If the queue already contains elements equivalent to `New_Item`, then it is inserted after the existing equivalent elements.

For a call on `Dequeue_Only_High_Priority`, if the head of the nonempty queue is  $E$ , and the function `Before(At_Least, Get_Priority(E))` returns `False`, then  $E$  is assigned to `Element` and then removed from the queue, and `Success` is set to `True`; otherwise, `Success` is set to `False` and `Element` is unchanged.

### A.18.31 The Generic Package `Containers.Bounded_Priority_Queues`

#### Static Semantics

The language-defined generic package `Containers.Bounded_Priority_Queues` provides type `Queue`, which implements the interface type `Containers.Synchronized_Queue_Interfaces.Queue`.

```

with System;
with Ada.Containers.Synchronized_Queue_Interfaces;
generic
  with package Queue_Interfaces is
    new Ada.Containers.Synchronized_Queue_Interfaces (<>);
  type Queue_Priority is private;
  with function Get_Priority
    (Element : Queue_Interfaces.Element_Type) return Queue_Priority is <>;
  with function Before
    (Left, Right : Queue_Priority) return Boolean is <>;
  Default_Capacity : Count_Type;
  Default_Ceiling : System.Any_Priority := System.Priority'Last;
package Ada.Containers.Bounded_Priority_Queues
  with Preelaborate,
    Nonblocking, Global => in out synchronized is

  package Implementation is
    ... -- not specified by the language
  end Implementation;

  protected type Queue
    (Capacity : Count_Type := Default_Capacity;
     Ceiling : System.Any_Priority := Default_Ceiling)
    with Priority => Ceiling is
    new Queue_Interfaces.Queue with

    overriding
    entry Enqueue (New_Item : in Queue_Interfaces.Element_Type);
    overriding
    entry Dequeue (Element : out Queue_Interfaces.Element_Type);

    not overriding
    procedure Dequeue_Only_High_Priority
      (At_Least : in Queue_Priority;
       Element : in out Queue_Interfaces.Element_Type;
       Success : out Boolean);

    overriding
    function Current_Use return Count_Type
      with Nonblocking, Global => null, Use_Formal => null;
    overriding
    function Peak_Use return Count_Type
      with Nonblocking, Global => null, Use_Formal => null;

  private
    ... -- not specified by the language
  end Queue;

private
  ... -- not specified by the language
end Ada.Containers.Bounded_Priority_Queues;

```

The semantics are the same as for `Unbounded_Priority_Queues`, except:

- The capacity for instances of type `Queue` is bounded and specified by the discriminant `Capacity`.

*Implementation Advice*

Bounded priority queue objects should be implemented without implicit pointers or dynamic allocation.

### A.18.32 The Generic Package Containers.Bounded\_Indefinite\_Holders

The language-defined generic package Containers.Bounded\_Indefinite\_Holders provides a private type Holder and a set of operations for that type. It provides the same operations as the package Containers.Indefinite\_Holders (see A.18.18), with the difference that the maximum storage is bounded.

*Static Semantics*

The declaration of the generic library package Containers.Bounded\_Indefinite\_Holders has the same contents and semantics as Containers.Indefinite\_Holders except:

- The following is added to the context clause:

```
with System.Storage_Elements; use System.Storage_Elements;
```

- An additional generic parameter follows Element\_Type:

```
Max_Element_Size_in_Storage_Elements : Storage_Count;
```

- The aspect\_definition for Preelaborable\_Initialization for type Holder is changed to:

```
Preelaborable_Initialization =>  
  Element_Type'Preelaborable_Initialization
```

- Add to the precondition of To\_Holder and Replace\_Element:

```
and then (New_Item'Size <=  
  Max_Element_Size_in_Storage_Elements * System.Storage_Unit  
or else raise Program_Error)
```

*Bounded (Run-Time) Errors*

It is a bounded error to assign from a bounded holder object while tampering with elements of that object is prohibited. Either Program\_Error is raised by the assignment, execution proceeds with the target object prohibiting tampering with elements, or execution proceeds normally.

*Implementation Requirements*

For each instance of Containers.Indefinite\_Holders and each instance of Containers.Bounded\_Indefinite\_Holders, if the two instances meet the following conditions, then the output generated by the Holder'Output or Holder'Write subprograms of either instance shall be readable by the Holder'Input or Holder'Read of the other instance, respectively:

- the Element\_Type parameters of the two instances are statically matching subtypes of the same type; and
- the output generated by Element\_Type'Output or Element\_Type'Write is readable by Element\_Type'Input or Element\_Type'Read, respectively (where Element\_Type denotes the type of the two actual Element\_Type parameters).

*Implementation Advice*

Bounded holder objects should be implemented without dynamic allocation and any finalization should be trivial unless Element\_Type needs finalization.

The Implementation Advice about the Move and Swap operations is deleted for bounded holders; these operations can copy elements as necessary.

### A.18.33 Example of Container Use

EXAMPLE The following example is an implementation of Dijkstra's shortest path algorithm in a directed graph with positive distances. The graph is represented by a map from nodes to sets of edges.

```

with Ada.Containers.Vectors;
with Ada.Containers.Doubly_Linked_Lists;
use Ada.Containers;
generic
  type Node is range <>;
package Shortest_Paths is
  type Distance is new Float range 0.0 .. Float'Last;
  type Edge is record
    To, From : Node;
    Length   : Distance;
  end record;

  package Node_Maps is new Vectors (Node, Node);
  -- The algorithm builds a map to indicate the node used to reach a given
  -- node in the shortest distance.

  package Adjacency_Lists is new Doubly_Linked_Lists (Edge);
  use Adjacency_Lists;

  package Graphs is new Vectors (Node, Adjacency_Lists.List);
  package Paths is new Doubly_Linked_Lists (Node);

  function Shortest_Path
    (G : Graphs.Vector; Source : Node; Target : Node) return Paths.List
    with Pre => G (Source) /= Adjacency_Lists.Empty_List;
end Shortest_Paths;

package body Shortest_Paths is
  function Shortest_Path
    (G : Graphs.Vector; Source : Node; Target : Node) return Paths.List
  is
    use Node_Maps, Paths, Graphs;
    Reached : array (Node) of Boolean := (others => False);
    -- The set of nodes whose shortest distance to the source is known.

    Reached_From : array (Node) of Node;
    So_Far : array (Node) of Distance := (others => Distance'Last);
    The_Path : Paths.List := Paths.Empty_List;
    Nearest_Distance : Distance;
    Next : Node;
  begin
    So_Far(Source) := 0.0;

    while not Reached(Target) loop
      Nearest_Distance := Distance'Last;

      -- Find closest node not reached yet, by iterating over all nodes.
      -- A more efficient algorithm uses a priority queue for this step.
      Next := Source;
      for N in Node'First .. Node'Last loop
        if not Reached(N)
          and then So_Far(N) < Nearest_Distance then
          Next := N;
          Nearest_Distance := So_Far(N);
        end if;
      end loop;

      if Nearest_Distance = Distance'Last then
        -- No next node found, graph is not connected
        return Paths.Empty_List;
      else
        Reached(Next) := True;
      end if;

      -- Update minimum distance to newly reachable nodes.
    end loop;
  end Shortest_Path;
end package body Shortest_Paths;

```

```

    for E of G (Next) loop
        if not Reached(E.To) then
            Nearest_Distance := E.Length + So_Far(Next);

            if Nearest_Distance < So_Far(E.To) then
                Reached_From(E.To) := Next;
                So_Far(E.To) := Nearest_Distance;
            end if;
        end if;
    end loop;
end loop;

-- Rebuild path from target to source.

declare
    N : Node := Target;
begin
    Prepend (The_Path, N);
    while N /= Source loop
        N := Reached_From(N);
        Prepend (The_Path, N);
    end loop;
end;

return The_Path;
end;
end Shortest_Paths;

```

Note that the effect of the Constant\_Indexing aspect (on type Vector) and the Implicit\_Dereference aspect (on type Reference\_Type) is that

G (Next)

is a convenient shorthand for

G.Constant\_Reference (Next).Element.all

Similarly, the effect of the loop:

```

for E of G (Next) loop
    if not Reached(E.To) then
        ...
    end if;
end loop;

```

is the same as:

```

for C in G (Next).Iterate loop
    declare
        E : Edge renames G (Next) (C);
    begin
        if not Reached(E.To) then
            ...
        end if;
    end;
end loop;

```

which is the same as:

```

declare
  L : Adjacency_Lists.List renames G (Next);
  C : Adjacency_Lists.Cursor := L.First;
begin
  while Has_Element (C) loop
    declare
      E : Edge renames L(C);
    begin
      if not Reached(E.To) then
        ...
      end if;
    end;
    C := L.Next (C);
  end loop;
end;

```

## A.19 The Package Locales

A *locale* identifies a geopolitical place or region and its associated language, which can be used to determine other internationalization-related characteristics.

*Static Semantics*

The library package Locales has the following declaration:

```

package Ada.Locales
  with Preelaborate, Remote_Types is
    type Language_Code is new String (1 .. 3)
      with Dynamic_Predicate =>
        (for all E of Language_Code => E in 'a' .. 'z');
    type Country_Code is new String (1 .. 2)
      with Dynamic_Predicate =>
        (for all E of Country_Code => E in 'A' .. 'Z');
    Language_Unknown : constant Language_Code := "und";
    Country_Unknown : constant Country_Code := "ZZ";
    function Language return Language_Code;
    function Country return Country_Code;
end Ada.Locales;

```

The *active locale* is the locale associated with the partition of the current task.

Function `Language` returns the code of the language associated with the active locale. If the `Language_Code` associated with the active locale cannot be determined from the environment, then `Language` returns `Language_Unknown`. Otherwise, the result shall be a lower-case string representation of an ISO 639-3:2007 alpha-3 code that identifies a language.

Function `Country` returns the code of the country associated with the active locale. If the `Country_Code` associated with the active locale cannot be determined from the environment, then `Country` returns `Country_Unknown`. Otherwise, the result shall be an upper-case string representation of an ISO 3166-1:2006 alpha-2 code that identifies a country.



## Annex B (normative)

### Interface to Other Languages

This Annex describes features for writing mixed-language programs. General interface support is presented first; then specific support for C, COBOL, and Fortran is defined, in terms of language interface packages for each of these languages.

#### *Implementation Requirements*

Support for interfacing to any foreign language is optional. However, an implementation shall not provide any optional aspect, attribute, library unit, or pragma having the same name as an aspect, attribute, library unit, or pragma (respectively) specified in the subclauses of this Annex unless the provided construct is either as specified in those subclauses or is more limited in capability than that required by those subclauses. A program that attempts to use an unsupported capability of this Annex shall either be identified by the implementation before run time or shall raise an exception at run time.

#### B.1 Interfacing Aspects

An *interfacing* aspect is a representation aspect that is one of the aspects Import, Export, Link\_Name, External\_Name, or Convention.

Specifying the Import aspect to have the value True is used to import an entity defined in a foreign language into an Ada program, thus allowing a foreign-language subprogram to be called from Ada, or a foreign-language variable to be accessed from Ada. In contrast, specifying the Export aspect to have the value True is used to export an Ada entity to a foreign language, thus allowing an Ada subprogram to be called from a foreign language, or an Ada object to be accessed from a foreign language. The Import and Export aspects are intended primarily for objects and subprograms, although implementations are allowed to support other entities. The Link\_Name and External\_Name aspects are used to specify the link name and external name, respectively, to be used to identify imported or exported entities in the external environment.

The Convention aspect is used to indicate that an Ada entity should use the conventions of another language. It is intended primarily for types and “callback” subprograms. For example, “**with** Convention => Fortran” on the declaration of an array type Matrix implies that Matrix should be represented according to the conventions of the supported Fortran implementation, namely column-major order.

A pragma Linker\_Options is used to specify the system linker parameters necessary when a given compilation unit is included in a partition.

#### *Syntax*

The form of a pragma Linker\_Options is as follows:

```
pragma Linker_Options(string_expression);
```

A pragma Linker\_Options is allowed only at the place of a *declarative\_item*.

#### *Name Resolution Rules*

The Import and Export aspects are of type Boolean.

The Link\_Name and External\_Name aspects are of type String.

The expected type for the *string\_expression* in pragma Linker\_Options is String.

## Legality Rules

The aspect Convention shall be specified by a *convention\_identifier* which shall be the name of a *convention*. The convention names are implementation defined, except for certain language-defined ones, such as Ada and Intrinsic, as explained in 9.3.1. Additional convention names generally represent the calling conventions of foreign languages, language implementations, or specific run-time models. The convention of a callable entity is its *calling convention*.

If *L* is a *convention\_identifier* for a language, then a type *T* is said to be *compatible with convention L*, (alternatively, is said to be an *L-compatible type*) if any of the following conditions are met:

- *T* is declared in a language interface package corresponding to *L* and is defined to be *L-compatible* (see B.3, B.3.1, B.3.2, B.4, B.5),
- Convention *L* has been specified for *T*, and *T* is *eligible for convention L*; that is:
  - *T* is an enumeration type such that all internal codes (whether assigned by default or explicitly) are within an implementation-defined range that includes at least the range of values 0 .. 2\*\*15-1;
  - *T* is an array type with either an unconstrained or statically-constrained first subtype, and its component type is *L-compatible*,
  - *T* is a record type that has no discriminants and that only has components with statically-constrained subtypes, and each component type is *L-compatible*,
  - *T* is an access-to-object type, its designated type is *L-compatible*, and its designated subtype is not an unconstrained array subtype,
  - *T* is an access-to-subprogram type, and its designated profile's parameter and result types are all *L-compatible*.
- *T* is derived from an *L-compatible* type,
- *T* is an anonymous access type, and *T* is eligible for convention *L*,
- The implementation permits *T* as an *L-compatible* type.

If the Convention aspect is specified for a type, then the type shall either be compatible with or eligible for the specified convention.

If convention *L* is specified for a type *T*, for each component of *T* that has an anonymous access type, the convention of the anonymous access type is *L*. If convention *L* is specified for an object that has an anonymous access type, the convention of the anonymous access type is *L*.

Notwithstanding any rule to the contrary, a declaration with a True Import aspect shall not have a completion.

An entity with a True Import aspect (or Export aspect) is said to be *imported* (respectively, *exported*). An entity shall not be both imported and exported.

The declaration of an imported object shall not include an explicit initialization expression. Default initializations are not performed.

The type of an imported or exported object shall be compatible with the specified Convention aspect, if any.

For an imported or exported subprogram, the result and parameter types shall each be compatible with the specified Convention aspect, if any.

The *aspect\_definition* (if any) used to directly specify an Import, Export, External\_Name, or Link\_Name aspect shall be a static expression. The *string\_expression* of a pragma Linker\_Options shall be static. An External\_Name or Link\_Name aspect shall be specified only for an entity that is either imported or exported.

*Static Semantics*

The Convention aspect represents the calling convention or representation convention of the entity. For an access-to-subprogram type, it represents the calling convention of designated subprograms. In addition:

- A True Import aspect indicates that the entity is defined externally (that is, outside the Ada program). This aspect is never inherited; if not directly specified, the Import aspect is False.
- A True Export aspect indicates that the entity is used externally. This aspect is never inherited; if not directly specified, the Export aspect is False.
- For an entity with a True Import or Export aspect, an external name, link name, or both may also be specified.

An *external name* is a string value for the name used by a foreign language program either for an entity that an Ada program imports, or for referring to an entity that an Ada program exports.

A *link name* is a string value for the name of an exported or imported entity, based on the conventions of the foreign language's compiler in interfacing with the system's linker tool.

The meaning of link names is implementation defined. If neither a link name nor the Address attribute of an imported or exported entity is specified, then a link name is chosen in an implementation-defined manner, based on the external name if one is specified.

Pragma Linker\_Options has the effect of passing its string argument as a parameter to the system linker (if one exists), if the immediately enclosing compilation unit is included in the partition being linked. The interpretation of the string argument, and the way in which the string arguments from multiple Linker\_Options pragmas are combined, is implementation defined.

*Dynamic Semantics*

Notwithstanding what this document says elsewhere, the elaboration of a declaration with a True Import aspect does not create the entity. Such an elaboration has no other effect than to allow the defining name to denote the external entity.

*Erroneous Execution*

It is the programmer's responsibility to ensure that the use of interfacing aspects does not violate Ada semantics; otherwise, program execution is erroneous. For example, passing an object with mode **in** to imported code that modifies it causes erroneous execution. Similarly, calling an imported subprogram that is not pure from a pure package causes erroneous execution.

*Implementation Advice*

If an implementation supports Export for a given language, then it should also allow the main subprogram to be written in that language. It should support some mechanism for invoking the elaboration of the Ada library units included in the system, and for invoking the finalization of the environment task. On typical systems, the recommended mechanism is to provide two subprograms whose link names are "adainit" and "adafinal". Adainit should contain the elaboration code for library units. Adafinal should contain the finalization code. These subprograms should have no effect the second and subsequent time they are called.

Automatic elaboration of preelaborated packages should be provided when specifying the Export aspect as True is supported.

For each supported convention *L* other than Intrinsic, an implementation should support specifying the Import and Export aspects for objects of *L*-compatible types and for subprograms, and the Convention aspect for *L*-eligible types and for subprograms, presuming the other language has corresponding features. Specifying the Convention aspect should be supported for enumeration types whose internal codes fall within the range 0 .. 2\*\*15-1, but no recommendation is made for other scalar types.

NOTE 1 Implementations can place restrictions on interfacing aspects; for example, requiring each exported entity to be declared at the library level.

NOTE 2 The Convention aspect in combination with the Import aspect indicates the conventions for accessing external entities. It is possible that the actual entity is written in assembly language, but reflects the conventions of a particular language. For example, **with** Convention => Ada can be used to interface to an assembly language routine that obeys the Ada compiler's calling conventions.

NOTE 3 To obtain "call-back" to an Ada subprogram from a foreign language environment, the Convention aspect can be specified both for the access-to-subprogram type and the specific subprogram(s) to which 'Access is applied.

NOTE 4 Machine code insertions can also be relevant for interfacing; see 16.8.

NOTE 5 If both External\_Name and Link\_Name are specified for a given entity, then the External\_Name is ignored.

EXAMPLE *Example of interfacing aspects:*

```
package Fortran_Library is
  function Sqrt (X : Float) return Float
    with Import => True, Convention => Fortran;
  type Matrix is array (Natural range <>, Natural range <>) of Float
    with Convention => Fortran;
  function Invert (M : Matrix) return Matrix
    with Import => True, Convention => Fortran;
end Fortran_Library;
```

## B.2 The Package Interfaces

Package Interfaces is the parent of several library packages that declare types and other entities useful for interfacing to foreign languages. It also contains some implementation-defined types that are useful across more than one language (in particular for interfacing to assembly language).

### Static Semantics

The library package Interfaces has the following skeletal declaration:

```
package Interfaces
  with Pure is

  type Integer_n is range -2**(n-1) .. 2**(n-1) - 1; --2's complement
  type Unsigned_n is mod 2**n;

  function Shift_Left (Value : Unsigned_n; Amount : Natural)
    return Unsigned_n;
  function Shift_Right (Value : Unsigned_n; Amount : Natural)
    return Unsigned_n;
  function Shift_Right_Arithmetic (Value : Unsigned_n; Amount : Natural)
    return Unsigned_n;
  function Rotate_Left (Value : Unsigned_n; Amount : Natural)
    return Unsigned_n;
  function Rotate_Right (Value : Unsigned_n; Amount : Natural)
    return Unsigned_n;
  ...
end Interfaces;
```

### Implementation Requirements

An implementation shall provide the following declarations in the visible part of package Interfaces:

- Signed and modular integer types of  $n$  bits, if supported by the target architecture, for each  $n$  that is at least the size of a storage element and that is a factor of the word size. The names of these types are of the form `Integer_n` for the signed types, and `Unsigned_n` for the modular types;
- For each such modular type in Interfaces, shifting and rotating subprograms as specified in the declaration of Interfaces above. These subprograms are Intrinsic. They operate on a bit-by-bit basis, using the binary representation of the value of the operands to yield a binary representation for the result. The Amount parameter gives the number of bits by which to shift or rotate. For shifting, zero bits are shifted in, except in the case of `Shift_Right_Arithmetic`, where one bits are shifted in if Value is at least half the modulus.

- Floating point types corresponding to each floating point format fully supported by the hardware.

#### Implementation Permissions

An implementation may provide implementation-defined library units that are children of Interfaces, and may add declarations to the visible part of Interfaces in addition to the ones defined above.

A child package of package Interfaces with the name of a convention may be provided independently of whether the convention is supported by the Convention aspect and vice versa. Such a child package should contain any declarations that would be useful for interfacing to the language (implementation) represented by the convention. Any declarations useful for interfacing to any language on the given hardware architecture should be provided directly in Interfaces.

#### Implementation Advice

An implementation supporting an interface to C, COBOL, or Fortran should provide the corresponding package or packages described in the following subclauses.

### B.3 Interfacing with C and C++

The facilities relevant to interfacing with the C language and the corresponding subset of the C++ language are the package Interfaces.C and its children, and support for specifying the Convention aspect with *convention\_identifiers* C, C\_Pass\_By\_Copy, and any of the C\_Variadic\_n conventions described below.

The package Interfaces.C contains the basic types, constants, and subprograms that allow an Ada program to pass scalars and strings to C and C++ functions. When this subclause mentions a C entity, the reference also applies to the corresponding entity in C++.

#### Static Semantics

The library package Interfaces.C has the following declaration:

```
package Interfaces.C
with Pure is

  -- Declarations based on C's <limits.h>
  CHAR_BIT   : constant := implementation-defined;    -- typically 8
  SCHAR_MIN  : constant := implementation-defined;    -- typically -128
  SCHAR_MAX  : constant := implementation-defined;    -- typically 127
  UCHAR_MAX  : constant := implementation-defined;    -- typically 255

  -- Signed and Unsigned Integers
  type int     is range implementation-defined;
  type short  is range implementation-defined;
  type long   is range implementation-defined;

  type signed_char is range SCHAR_MIN .. SCHAR_MAX;
  for signed_char'Size use CHAR_BIT;

  type unsigned      is mod implementation-defined;
  type unsigned_short is mod implementation-defined;
  type unsigned_long  is mod implementation-defined;

  type unsigned_char is mod (UCHAR_MAX+1);
  for unsigned_char'Size use CHAR_BIT;

  subtype plain_char is implementation-defined;

  type ptrdiff_t is range implementation-defined;

  type size_t is mod implementation-defined;

  -- Boolean Type
  type C_bool is new Boolean;

  -- Floating Point
  type C_float   is digits implementation-defined;
```

```

type double      is digits implementation-defined;
type long_double is digits implementation-defined;
-- Characters and Strings
type char is <implementation-defined character type>;
nul : constant char := implementation-defined;
function To_C    (Item : in Character) return char;
function To_Ada (Item : in char) return Character;
type char_array is array (size_t range <>) of aliased char
  with Pack;
for char_array'Component_Size use CHAR_BIT;
function Is_Nul_Terminated (Item : in char_array) return Boolean;
function To_C    (Item      : in String;
                  Append_Nul : in Boolean := True)
  return char_array;
function To_Ada (Item      : in char_array;
                  Trim_Nul  : in Boolean := True)
  return String;
procedure To_C (Item      : in String;
                Target    : out char_array;
                Count     : out size_t;
                Append_Nul : in Boolean := True);
procedure To_Ada (Item      : in char_array;
                  Target    : out String;
                  Count     : out Natural;
                  Trim_Nul  : in Boolean := True);
-- Wide Character and Wide String
type wchar_t is <implementation-defined character type>;
wide_nul : constant wchar_t := implementation-defined;
function To_C    (Item : in Wide_Character) return wchar_t;
function To_Ada (Item : in wchar_t) return Wide_Character;
type wchar_array is array (size_t range <>) of aliased wchar_t
  with Pack;
function Is_Nul_Terminated (Item : in wchar_array) return Boolean;
function To_C    (Item      : in Wide_String;
                  Append_Nul : in Boolean := True)
  return wchar_array;
function To_Ada (Item      : in wchar_array;
                  Trim_Nul  : in Boolean := True)
  return Wide_String;
procedure To_C (Item      : in Wide_String;
                Target    : out wchar_array;
                Count     : out size_t;
                Append_Nul : in Boolean := True);
procedure To_Ada (Item      : in wchar_array;
                  Target    : out Wide_String;
                  Count     : out Natural;
                  Trim_Nul  : in Boolean := True);
-- ISO/IEC 10646 compatible types.
type char16_t is <implementation-defined character type>;
char16_nul : constant char16_t := implementation-defined;

```



```

function To_C (Item : in Wide_Character) return char16_t;
function To_Ada (Item : in char16_t) return Wide_Character;

type char16_array is array (size_t range <>) of aliased char16_t
  with Pack;

function Is_Nul_Terminated (Item : in char16_array) return Boolean;
function To_C (Item      : in Wide_String;
               Append_Nul : in Boolean := True)
  return char16_array;

function To_Ada (Item      : in char16_array;
                 Trim_Nul  : in Boolean := True)
  return Wide_String;

procedure To_C (Item      : in Wide_String;
                Target    : out char16_array;
                Count     : out size_t;
                Append_Nul : in Boolean := True);

procedure To_Ada (Item      : in char16_array;
                  Target    : out Wide_String;
                  Count     : out Natural;
                  Trim_Nul  : in Boolean := True);

type char32_t is <implementation-defined character type>;
char32_nul : constant char32_t := implementation-defined;

function To_C (Item : in Wide_Wide_Character) return char32_t;
function To_Ada (Item : in char32_t) return Wide_Wide_Character;

type char32_array is array (size_t range <>) of aliased char32_t
  with Pack;

function Is_Nul_Terminated (Item : in char32_array) return Boolean;
function To_C (Item      : in Wide_Wide_String;
               Append_Nul : in Boolean := True)
  return char32_array;

function To_Ada (Item      : in char32_array;
                 Trim_Nul  : in Boolean := True)
  return Wide_Wide_String;

procedure To_C (Item      : in Wide_Wide_String;
                Target    : out char32_array;
                Count     : out size_t;
                Append_Nul : in Boolean := True);

procedure To_Ada (Item      : in char32_array;
                  Target    : out Wide_Wide_String;
                  Count     : out Natural;
                  Trim_Nul  : in Boolean := True);

Terminator_Error : exception;
end Interfaces.C;

```

Each of the types declared in Interfaces.C is C-compatible.

The types `int`, `short`, `long`, `unsigned`, `ptrdiff_t`, `size_t`, `double`, `char`, `wchar_t`, `char16_t`, and `char32_t` correspond respectively to the C types having the same names. The types `signed_char`, `unsigned_short`, `unsigned_long`, `unsigned_char`, `C_bool`, `C_float`, and `long_double` correspond respectively to the C types `signed char`, `unsigned short`, `unsigned long`, `unsigned char`, `bool`, `float`, and `long double`.

The type of the subtype `plain_char` is either `signed_char` or `unsigned_char`, depending on the C implementation.

```

function To_C (Item : in Character) return char;
function To_Ada (Item : in char) return Character;

```

The functions `To_C` and `To_Ada` map between the Ada type `Character` and the C type `char`.

```
function Is_Nul_Terminated (Item : in char_array) return Boolean;
```

The result of Is\_Nul\_Terminated is True if Item contains nul, and is False otherwise.

```
function To_C (Item : in String; Append_Nul : in Boolean := True)
return char_array;
```

```
function To_Ada (Item : in char_array; Trim_Nul : in Boolean := True)
return String;
```

The result of To\_C is a char\_array value of length Item'Length (if Append\_Nul is False) or Item'Length+1 (if Append\_Nul is True). The lower bound is 0. For each component Item(I), the corresponding component in the result is To\_C applied to Item(I). The value nul is appended if Append\_Nul is True. If Append\_Nul is False and Item'Length is 0, then To\_C propagates Constraint\_Error.

The result of To\_Ada is a String whose length is Item'Length (if Trim\_Nul is False) or the length of the slice of Item preceding the first nul (if Trim\_Nul is True). The lower bound of the result is 1. If Trim\_Nul is False, then for each component Item(I) the corresponding component in the result is To\_Ada applied to Item(I). If Trim\_Nul is True, then for each component Item(I) before the first nul the corresponding component in the result is To\_Ada applied to Item(I). The function propagates Terminator\_Error if Trim\_Nul is True and Item does not contain nul.

```
procedure To_C (Item      : in String;
                Target    : out char_array;
                Count     : out size_t;
                Append_Nul : in Boolean := True);
```

```
procedure To_Ada (Item      : in char_array;
                  Target    : out String;
                  Count     : out Natural;
                  Trim_Nul  : in Boolean := True);
```

For procedure To\_C, each element of Item is converted (via the To\_C function) to a char, which is assigned to the corresponding element of Target. If Append\_Nul is True, nul is then assigned to the next element of Target. In either case, Count is set to the number of Target elements assigned. If Target is not long enough, Constraint\_Error is propagated.

For procedure To\_Ada, each element of Item (if Trim\_Nul is False) or each element of Item preceding the first nul (if Trim\_Nul is True) is converted (via the To\_Ada function) to a Character, which is assigned to the corresponding element of Target. Count is set to the number of Target elements assigned. If Target is not long enough, Constraint\_Error is propagated. If Trim\_Nul is True and Item does not contain nul, then Terminator\_Error is propagated.

```
function Is_Nul_Terminated (Item : in wchar_array) return Boolean;
```

The result of Is\_Nul\_Terminated is True if Item contains wide\_nul, and is False otherwise.

```
function To_C (Item : in Wide_Character) return wchar_t;
function To_Ada (Item : in wchar_t) return Wide_Character;
```

To\_C and To\_Ada provide the mappings between the Ada and C wide character types.

```

function To_C (Item      : in Wide_String;
                Append_Nul : in Boolean := True)
    return wchar_array;

function To_Ada (Item      : in wchar_array;
                Trim_Nul   : in Boolean := True)
    return Wide_String;

procedure To_C (Item      : in Wide_String;
                Target     : out wchar_array;
                Count      : out size_t;
                Append_Nul : in Boolean := True);

procedure To_Ada (Item      : in wchar_array;
                Target     : out Wide_String;
                Count      : out Natural;
                Trim_Nul   : in Boolean := True);

```

The To\_C and To\_Ada subprograms that convert between Wide\_String and wchar\_array have analogous effects to the To\_C and To\_Ada subprograms that convert between String and char\_array, except that wide\_nul is used instead of nul.

```

function Is_Nul_Terminated (Item : in char16_array) return Boolean;

```

The result of Is\_Nul\_Terminated is True if Item contains char16\_nul, and is False otherwise.

```

function To_C (Item : in Wide_Character) return char16_t;
function To_Ada (Item : in char16_t ) return Wide_Character;

```

To\_C and To\_Ada provide mappings between the Ada and C 16-bit character types.

```

function To_C (Item      : in Wide_String;
                Append_Nul : in Boolean := True)
    return char16_array;

function To_Ada (Item      : in char16_array;
                Trim_Nul   : in Boolean := True)
    return Wide_String;

procedure To_C (Item      : in Wide_String;
                Target     : out char16_array;
                Count      : out size_t;
                Append_Nul : in Boolean := True);

procedure To_Ada (Item      : in char16_array;
                Target     : out Wide_String;
                Count      : out Natural;
                Trim_Nul   : in Boolean := True);

```

The To\_C and To\_Ada subprograms that convert between Wide\_String and char16\_array have analogous effects to the To\_C and To\_Ada subprograms that convert between String and char\_array, except that char16\_nul is used instead of nul.

```

function Is_Nul_Terminated (Item : in char32_array) return Boolean;

```

The result of Is\_Nul\_Terminated is True if Item contains char32\_nul, and is False otherwise.

```

function To_C (Item : in Wide_Wide_Character) return char32_t;
function To_Ada (Item : in char32_t ) return Wide_Wide_Character;

```

To\_C and To\_Ada provide mappings between the Ada and C 32-bit character types.

```

function To_C (Item          : in Wide_Wide_String;
               Append_Nul    : in Boolean := True)
  return char32_array;

function To_Ada (Item          : in char32_array;
               Trim_Nul       : in Boolean := True)
  return Wide_Wide_String;

procedure To_C (Item          : in Wide_Wide_String;
               Target         : out char32_array;
               Count          : out size_t;
               Append_Nul     : in Boolean := True);

procedure To_Ada (Item          : in char32_array;
               Target         : out Wide_Wide_String;
               Count          : out Natural;
               Trim_Nul       : in Boolean := True);

```

The To\_C and To\_Ada subprograms that convert between Wide\_Wide\_String and char32\_array have analogous effects to the To\_C and To\_Ada subprograms that convert between String and char\_array, except that char32\_nul is used instead of nul.

The Convention aspect with *convention\_identifier* C\_Pass\_By\_Copy shall only be specified for a type.

The eligibility rules in B.1 do not apply to convention C\_Pass\_By\_Copy. Instead, a type T is eligible for convention C\_Pass\_By\_Copy if T is an unchecked union type or if T is a record type that has no discriminants and that only has components with statically constrained subtypes, and each component is C-compatible.

If a type is C\_Pass\_By\_Copy-compatible, then it is also C-compatible.

The identifiers C\_Variadic\_0, C\_Variadic\_1, C\_Variadic\_2, and so on are *convention\_identifiers*. These conventions are said to be *C\_Variadic*. The convention C\_Variadic\_n is the calling convention for a variadic C function taking n fixed parameters and then a variable number of additional parameters. The C\_Variadic\_n convention shall only be specified as the convention aspect for a subprogram, or for an access-to-subprogram type, having at least n parameters. A type is compatible with a C\_Variadic convention if and only if the type is C-compatible.

#### Implementation Requirements

An implementation shall support specifying aspect Convention with a C *convention\_identifier* for a C-eligible type (see B.1). An implementation shall support specifying aspect Convention with a C\_Pass\_By\_Copy *convention\_identifier* for a C\_Pass\_By\_Copy-eligible type.

#### Implementation Permissions

An implementation may provide additional declarations in the C interface packages.

An implementation is not required to support specifying the Convention aspect with *convention\_identifier* C in the following cases:

- for a subprogram that has a parameter of an unconstrained array subtype, unless the Import aspect has the value True for the subprogram;
- for a function with an unconstrained array result subtype;
- for an object whose nominal subtype is an unconstrained array subtype.

#### Implementation Advice

The constants nul, wide\_nul, char16\_nul, and char32\_nul should have a representation of zero.

An implementation should support the following interface correspondences between Ada and C.

- An Ada procedure corresponds to a void-returning C function.

- An Ada function corresponds to a non-void C function.
- An Ada enumeration type corresponds to a C enumeration type with corresponding enumeration literals having the same internal codes, provided the internal codes fall within the range of the C int type.
- An Ada **in** scalar parameter is passed as a scalar argument to a C function.
- An Ada **in** parameter of an access-to-object type with designated type T is passed as a t\* argument to a C function, where t is the C type corresponding to the Ada type T.
- An Ada **access** T parameter, or an Ada **out** or **in out** parameter of an elementary type T, is passed as a t\* argument to a C function, where t is the C type corresponding to the Ada type T. In the case of an elementary **out** or **in out** parameter, a pointer to a temporary copy is used to preserve by-copy semantics.
- An Ada parameter of a (record) type T of convention C\_Pass\_By\_Copy, of mode **in**, is passed as a t argument to a C function, where t is the C struct corresponding to the Ada type T.
- An Ada parameter of a record type T, other than an **in** parameter of a type of convention C\_Pass\_By\_Copy, is passed as a t\* argument to a C function, with the const modifier if the Ada mode is **in**, where t is the C struct corresponding to the Ada type T.
- An Ada parameter of an array type with component type T is passed as a t\* argument to a C function, with the const modifier if the Ada mode is **in**, where t is the C type corresponding to the Ada type T.
- An Ada parameter of an access-to-subprogram type is passed as a pointer to a C function whose prototype corresponds to the designated subprogram's specification.
- An Ada parameter of a private type is passed as specified for the full view of the type.
- The rules of correspondence given above for parameters of mode **in** also apply to the return object of a function.

An implementation should provide unsigned long\_long and long\_long as 64-bit modular and signed integer types (respectively) in package Interfaces.C if the C implementation supports unsigned long long and long long as 64-bit types.

NOTE 1 Values of type char\_array are not implicitly terminated with nul. If a char\_array is to be passed as a parameter to an imported C function requiring nul termination, it is the programmer's responsibility to obtain this effect.

NOTE 2 To obtain the effect of C's sizeof(item\_type), where Item\_Type is the corresponding Ada type, evaluate the expression: size\_t(Item\_Type Size/CHAR\_BIT).

NOTE 3 A variadic C function can correspond to several Ada subprograms, taking various specific numbers and types of parameters.

EXAMPLE *Example of using the Interfaces.C package:*

```
--Calling the C Library Functions strcpy and printf
with Interfaces.C;
procedure Test is
  package C renames Interfaces.C;
  use type C.char_array;
  -- Call <string.h>strcpy:
  -- C definition of strcpy: char *strcpy(char *s1, const char *s2);
  -- This function copies the string pointed to by s2 (including the
terminating null character)
  -- into the array pointed to by s1. If copying takes place between
objects that overlap,
  -- the behavior is undefined. The strcpy function returns the value of
s1.
```

```

-- Note: since the C function's return value is of no interest, the Ada
interface is a procedure
procedure Strcpy (Target : out C.char_array;
                  Source : in C.char_array)
    with Import => True, Convention => C, External_Name => "strcpy";

-- Call <sdtio.h>printf:
-- C definition of printf: int printf ( const char * format, ... );
-- This function writes the C string pointed by format to the standard
output (stdout).
-- If format includes format specifiers (subsequences beginning with
%), the additional
-- arguments following format are formatted and inserted in the
resulting string
-- replacing their respective specifiers. If the number of arguments
does not match
-- the number of format specifiers, or if the types of the arguments do
not match
-- the corresponding format specifier, the behaviour is undefined. On
success, the
-- printf function returns the total number of characters written to
the standard output.
-- If a writing error occurs, a negative number is returned.

-- Note: since the C function's return value is of no interest, the Ada
interface is a procedure
procedure Printf (Format : in C.char_array;
                  Param1 : in C.char_array;
                  Param2 : in C.int)
    with Import => True, Convention => C_Variadic_1, External_Name =>
"printf";

Chars1 : C.char_array(1..20);
Chars2 : C.char_array(1..20);

begin
    Chars2(1..6) := "qwert" & C.nul;
    Strcpy(Chars1, Chars2);
    -- Now Chars1(1..6) = "qwert" & C.Nul
    Printf("The String=%s, Length=%d", Chars1, Chars1'Length);
end Test;

```

### B.3.1 The Package Interfaces.C.Strings

The package Interfaces.C.Strings declares types and subprograms allowing an Ada program to allocate, reference, update, and free C-style strings. In particular, the private type chars\_ptr corresponds to a common use of "char \*" in C programs, and an object of this type can be passed to a subprogram to which **with** Import => True, Convention => C has been specified, and for which "char \*" is the type of the argument of the C function.

#### Static Semantics

The library package Interfaces.C.Strings has the following declaration:

```

package Interfaces.C.Strings
    with Preelaborate, Nonblocking, Global => in out synchronized is

    type char_array_access is access all char_array;

    type chars_ptr is private
        with Preelaborable_Initialization;

    type chars_ptr_array is array (size_t range <>) of aliased chars_ptr;

    Null_Ptr : constant chars_ptr;

```



```

function To_Chars_Ptr (Item      : in char_array_access;
                       Nul_Check : in Boolean := False)
    return chars_ptr;

function New_Char_Array (Chars   : in char_array) return chars_ptr;

function New_String (Str : in String) return chars_ptr;

procedure Free (Item : in out chars_ptr);

Dereference_Error : exception;

function Value (Item : in chars_ptr) return char_array;

function Value (Item : in chars_ptr; Length : in size_t)
    return char_array;

function Value (Item : in chars_ptr) return String;

function Value (Item : in chars_ptr; Length : in size_t)
    return String;

function Strlen (Item : in chars_ptr) return size_t;

procedure Update (Item   : in chars_ptr;
                  Offset : in size_t;
                  Chars   : in char_array;
                  Check   : in Boolean := True);

procedure Update (Item   : in chars_ptr;
                  Offset : in size_t;
                  Str     : in String;
                  Check   : in Boolean := True);

Update_Error : exception;

private
    ... -- not specified by the language
end Interfaces.C.Strings;

```

The type `chars_ptr` is C-compatible and corresponds to the use of C's "char \*" for a pointer to the first char in a char array terminated by nul. When an object of type `chars_ptr` is declared, its value is by default set to `Null_Ptr`, unless the object is imported (see B.1).

```

function To_Chars_Ptr (Item      : in char_array_access;
                       Nul_Check : in Boolean := False)
    return chars_ptr;

```

If `Item` is **null**, then `To_Chars_Ptr` returns `Null_Ptr`. If `Item` is not **null**, `Nul_Check` is **True**, and `Item.all` does not contain nul, then the function propagates `Terminator_Error`; otherwise, `To_Chars_Ptr` performs a pointer conversion with no allocation of memory.

```

function New_Char_Array (Chars   : in char_array) return chars_ptr;

```

This function returns a pointer to an allocated object initialized to `Chars(Chars'First .. Index) & nul`, where

- `Index = Chars'Last` if `Chars` does not contain nul, or
- `Index` is the smallest `size_t` value `I` such that `Chars(I+1) = nul`.

`Storage_Error` is propagated if the allocation fails.

```

function New_String (Str : in String) return chars_ptr;

```

This function is equivalent to `New_Char_Array(To_C(Str))`.

```

procedure Free (Item : in out chars_ptr);

```

If `Item` is `Null_Ptr`, then `Free` has no effect. Otherwise, `Free` releases the storage occupied by `Value(Item)`, and resets `Item` to `Null_Ptr`.

**function** Value (Item : **in** chars\_ptr) **return** char\_array;

If Item = Null\_Ptr, then Value propagates Dereference\_Error. Otherwise, Value returns the prefix of the array of chars pointed to by Item, up to and including the first nul. The lower bound of the result is 0. If Item does not point to a nul-terminated string, then execution of Value is erroneous.

**function** Value (Item : **in** chars\_ptr; Length : **in** size\_t)  
**return** char\_array;

If Item = Null\_Ptr, then Value propagates Dereference\_Error. Otherwise, Value returns the shorter of two arrays, either the first Length chars pointed to by Item, or Value(Item). The lower bound of the result is 0. If Length is 0, then Value propagates Constraint\_Error.

**function** Value (Item : **in** chars\_ptr) **return** String;

Equivalent to To\_Ada(Value(Item), Trim\_Nul=>True).

**function** Value (Item : **in** chars\_ptr; Length : **in** size\_t)  
**return** String;

Equivalent to To\_Ada(Value(Item, Length) & nul, Trim\_Nul=>True).

**function** Strlen (Item : **in** chars\_ptr) **return** size\_t;

Returns Val'Length-1 where Val = Value(Item); propagates Dereference\_Error if Item = Null\_Ptr.

**procedure** Update (Item : **in** chars\_ptr;  
Offset : **in** size\_t;  
Chars : **in** char\_array;  
Check : Boolean := True);

If Item = Null\_Ptr, then Update propagates Dereference\_Error. Otherwise, this procedure updates the value pointed to by Item, starting at position Offset, using Chars as the data to be copied into the array. Overwriting the nul terminator, and skipping with the Offset past the nul terminator, are both prevented if Check is True, as follows:

- Let N = Strlen(Item). If Check is True, then:
  - If Offset+Chars'Length>N, propagate Update\_Error.
  - Otherwise, overwrite the data in the array pointed to by Item, starting at the char at position Offset, with the data in Chars.
- If Check is False, then processing is as above, but with no check that Offset+Chars'Length>N.

**procedure** Update (Item : **in** chars\_ptr;  
Offset : **in** size\_t;  
Str : **in** String;  
Check : **in** Boolean := True);

Equivalent to Update(Item, Offset, To\_C(Str, Append\_Nul => False), Check).

#### *Erroneous Execution*

Execution of any of the following is erroneous if the Item parameter is not null\_ptr and Item does not point to a nul-terminated array of chars.

- a Value function not taking a Length parameter,
- the Free procedure,
- the Strlen function.

Execution of Free(X) is also erroneous if the chars\_ptr X was not returned by New\_Char\_Array or New\_String.

Reading or updating a freed char\_array is erroneous.

Execution of Update is erroneous if Check is False and a call with Check equal to True would have propagated Update\_Error.

NOTE New\_Char\_Array and New\_String can be implemented either through the allocation function from the C environment ("malloc") or through Ada dynamic memory allocation ("new"). The key points are

- the returned value (a chars\_ptr) is represented as a C "char \*" so that it can be passed to C functions;
- the allocated object can be freed by the programmer via a call of Free, rather than by calling a C function.

### B.3.2 The Generic Package Interfaces.C.Pointers

The generic package Interfaces.C.Pointers allows the Ada programmer to perform C-style operations on pointers. It includes an access type Pointer, Value functions that dereference a Pointer and deliver the designated array, several pointer arithmetic operations, and "copy" procedures that copy the contents of a source pointer into the array designated by a destination pointer. As in C, it treats an object Ptr of type Pointer as a pointer to the first element of an array, so that for example, adding 1 to Ptr yields a pointer to the second element of the array.

The generic allows two styles of usage: one in which the array is terminated by a special terminator element; and another in which the programmer keeps track of the length.

#### Static Semantics

The generic library package Interfaces.C.Pointers has the following declaration:

```
generic
  type Index is (<>);
  type Element is private;
  type Element_Array is array (Index range <>) of aliased Element;
  Default_Terminator : Element;
package Interfaces.C.Pointers
  with Preelaborate, Nonblocking, Global => in out synchronized is

  type Pointer is access all Element;

  function Value(Ref          : in Pointer;
                 Terminator : in Element := Default_Terminator)
    return Element_Array;

  function Value(Ref          : in Pointer;
                 Length      : in ptrdiff_t)
    return Element_Array;

  Pointer_Error : exception;

  -- C-style Pointer arithmetic

  function "+" (Left : in Pointer;   Right : in ptrdiff_t) return Pointer
    with Convention => Intrinsic;
  function "+" (Left : in ptrdiff_t; Right : in Pointer)   return Pointer
    with Convention => Intrinsic;
  function "-" (Left : in Pointer;   Right : in ptrdiff_t) return Pointer
    with Convention => Intrinsic;
  function "-" (Left : in Pointer;   Right : in Pointer)   return ptrdiff_t
    with Convention => Intrinsic;

  procedure Increment (Ref : in out Pointer)
    with Convention => Intrinsic;
  procedure Decrement (Ref : in out Pointer)
    with Convention => Intrinsic;

  function Virtual_Length (Ref          : in Pointer;
                           Terminator : in Element := Default_Terminator)
    return ptrdiff_t;
```

```

procedure Copy_Terminated_Array
  (Source      : in Pointer;
   Target      : in Pointer;
   Limit       : in ptrdiff_t := ptrdiff_t'Last;
   Terminator  : in Element := Default_Terminator);

procedure Copy_Array (Source : in Pointer;
                      Target  : in Pointer;
                      Length  : in ptrdiff_t);

end Interfaces.C.Pointers;

```

The type `Pointer` is C-compatible and corresponds to one use of C's "Element \*". An object of type `Pointer` is interpreted as a pointer to the initial `Element` in an `Element_Array`. Two styles are supported:

- Explicit termination of an array value with `Default_Terminator` (a special terminator value);
- Programmer-managed length, with `Default_Terminator` treated simply as a data element.

```

function Value(Ref      : in Pointer;
               Terminator : in Element := Default_Terminator)
return Element_Array;

```

This function returns an `Element_Array` whose value is the array pointed to by `Ref`, up to and including the first `Terminator`; the lower bound of the array is `Index'First`. `Interfaces.C.Strings.Dereference_Error` is propagated if `Ref` is **null**.

```

function Value(Ref      : in Pointer;
               Length    : in ptrdiff_t)
return Element_Array;

```

This function returns an `Element_Array` comprising the first `Length` elements pointed to by `Ref`. The exception `Interfaces.C.Strings.Dereference_Error` is propagated if `Ref` is **null**.

The "+" and "-" functions perform arithmetic on `Pointer` values, based on the `Size` of the array elements. In each of these functions, `Pointer_Error` is propagated if a `Pointer` parameter is **null**.

```

procedure Increment (Ref : in out Pointer);

```

Equivalent to `Ref := Ref+1`.

```

procedure Decrement (Ref : in out Pointer);

```

Equivalent to `Ref := Ref-1`.

```

function Virtual_Length (Ref      : in Pointer;
                        Terminator : in Element := Default_Terminator)
return ptrdiff_t;

```

Returns the number of `Elements`, up to the one just before the first `Terminator`, in `Value(Ref, Terminator)`.

```

procedure Copy_Terminated_Array
  (Source      : in Pointer;
   Target      : in Pointer;
   Limit       : in ptrdiff_t := ptrdiff_t'Last;
   Terminator  : in Element := Default_Terminator);

```

This procedure copies `Value(Source, Terminator)` into the array pointed to by `Target`; it stops either after `Terminator` has been copied, or the number of elements copied is `Limit`, whichever occurs first. `Dereference_Error` is propagated if either `Source` or `Target` is **null**.

```

procedure Copy_Array (Source : in Pointer;
                      Target  : in Pointer;
                      Length  : in ptrdiff_t);

```

This procedure copies the first `Length` elements from the array pointed to by `Source`, into the array pointed to by `Target`. `Dereference_Error` is propagated if either `Source` or `Target` is **null**.

*Erroneous Execution*

It is erroneous to dereference a Pointer that does not designate an aliased Element.

Execution of Value(Ref, Terminator) is erroneous if Ref does not designate an aliased Element in an Element\_Array terminated by Terminator.

Execution of Value(Ref, Length) is erroneous if Ref does not designate an aliased Element in an Element\_Array containing at least Length Elements between the designated Element and the end of the array, inclusive.

Execution of Virtual\_Length(Ref, Terminator) is erroneous if Ref does not designate an aliased Element in an Element\_Array terminated by Terminator.

Execution of Copy\_Terminated\_Array(Source, Target, Limit, Terminator) is erroneous in either of the following situations:

- Execution of both Value(Source, Terminator) and Value(Source, Limit) are erroneous, or
- Copying writes past the end of the array containing the Element designated by Target.

Execution of Copy\_Array(Source, Target, Length) is erroneous if either Value(Source, Length) is erroneous, or copying writes past the end of the array containing the Element designated by Target.

EXAMPLE *Example of Interfaces.C.Pointers:*

```

with Interfaces.C.Pointers;
with Interfaces.C.Strings;
procedure Test_Pointers is
  package C renames Interfaces.C;
  package Char_Ptrs is
    new C.Pointers (Index      => C.size_t,
                    Element    => C.char,
                    Element_Array => C.char_array,
                    Default_Terminator => C.nul);

  use type Char_Ptrs.Pointer;
  subtype Char_Star is Char_Ptrs.Pointer;

  procedure Strcpy (Target_Ptr, Source_Ptr : Char_Star) is
    Target_Temp_Ptr : Char_Star := Target_Ptr;
    Source_Temp_Ptr : Char_Star := Source_Ptr;
    Element : C.char;
  begin
    if Target_Temp_Ptr = null or Source_Temp_Ptr = null then
      raise C.Strings.Dereference_Error;
    end if;

    loop
      Element := Source_Temp_Ptr.all;
      Target_Temp_Ptr.all := Element;
      exit when C."="(Element, C.nul);
      Char_Ptrs.Increment(Target_Temp_Ptr);
      Char_Ptrs.Increment(Source_Temp_Ptr);
    end loop;
  end Strcpy;
begin
  ...
end Test_Pointers;

```

### B.3.3 Unchecked Union Types

Specifying aspect Unchecked\_Union to have the value True defines an interface correspondence between a given discriminated type and some C union. The aspect requires that the associated type shall be given a representation that allocates no space for its discriminant(s).

*Static Semantics*

For a discriminated record type having a `variant_part`, the following language-defined representation aspect may be specified:

**Unchecked\_Union**

The type of aspect `Unchecked_Union` is Boolean. If directly specified, the `aspect_definition` shall be a static expression. If not specified (including by inheritance), the aspect is False.

*Legality Rules*

A type for which aspect `Unchecked_Union` is True is called an *unchecked union type*. A subtype of an unchecked union type is defined to be an *unchecked union subtype*. An object of an unchecked union type is defined to be an *unchecked union object*.

All component subtypes of an unchecked union type shall be C-compatible.

If a component subtype of an unchecked union type is subject to a per-object constraint, then the component subtype shall be an unchecked union subtype.

Any name that denotes a discriminant of an object of an unchecked union type shall occur within the declarative region of the type or as the `selector_name` of an aggregate, and shall not occur within a `record_representation_clause`.

The type of a component declared in a `variant_part` of an unchecked union type shall not need finalization. In addition to the places where Legality Rules normally apply (see 15.3), this rule also applies in the private part of an instance of a generic unit. For an unchecked union type declared within the body of a generic unit, or within the body of any of its descendant library units, no part of the type of a component declared in a `variant_part` of the unchecked union type shall be of a formal private type or formal private extension declared within the formal part of the generic unit.

The completion of an incomplete or private type declaration having a `known_discriminant_part` shall not be an unchecked union type.

An unchecked union subtype shall only be passed as a generic actual parameter if the corresponding formal type has no known discriminants or is an unchecked union type.

*Static Semantics*

An unchecked union type is eligible for convention C.

All objects of an unchecked union type have the same size.

Discriminants of objects of an unchecked union type are of size zero.

Any check which would require reading a discriminant of an unchecked union object is suppressed (see 14.5). These checks include:

- The check performed when addressing a variant component (i.e. a component that was declared in a variant part) of an unchecked union object that the object has this component (see 7.1.3).
- Any checks associated with a type or subtype conversion of a value of an unchecked union type (see 7.6). This includes, for example, the check associated with the implicit subtype conversion of an assignment statement.
- The subtype membership check associated with the evaluation of a qualified expression (see 7.7) or an uninitialized allocator (see 7.8).

*Dynamic Semantics*

A view of an unchecked union object (including a type conversion or function call) has *inferable discriminants* if it has a constrained nominal subtype, unless the object is a component of an enclosing unchecked union object that is subject to a per-object constraint and the enclosing object lacks inferable discriminants.



An expression of an unchecked union type has inferable discriminants if it is either a name of an object with inferable discriminants or a qualified expression whose `subtype_mark` denotes a constrained subtype.

Program\_Error is raised in the following cases:

- Evaluation of the predefined equality operator for an unchecked union type if either of the operands lacks inferable discriminants.
- Evaluation of the predefined equality operator for a type which has a subcomponent of an unchecked union type whose nominal subtype is unconstrained.
- Evaluation of an individual membership test if the `subtype_mark` (if any) denotes a constrained unchecked union subtype and the *tested\_simple\_expression* lacks inferable discriminants.
- Conversion from a derived unchecked union type to an unconstrained non-unchecked-union type if the operand of the conversion lacks inferable discriminants.
- Execution of the default implementation of the Write or Read attribute of an unchecked union type.
- Execution of the default implementation of the Output or Input attribute of an unchecked union type if the type lacks default discriminant values.

NOTE The use of an unchecked union to obtain the effect of an unchecked conversion results in erroneous execution (see 14.5). Execution of the following example is erroneous even if `Float'Size = Integer'Size`:

```

type T (Flag : Boolean := False) is
  record
    case Flag is
      when False =>
        F1 : Float := 0.0;
      when True =>
        F2 : Integer := 0;
    end case;
  end record
  with Unchecked_Union;

X : T;
Y : Integer := X.F2; -- erroneous

```

## B.4 Interfacing with COBOL

The facilities relevant to interfacing with the COBOL language are the package `Interfaces.COBOL` and support for specifying the Convention aspect with *convention\_identifier* COBOL.

The COBOL interface package supplies several sets of facilities:

- A set of types corresponding to the native COBOL types of the supported COBOL implementation (so-called “internal COBOL representations”), allowing Ada data to be passed as parameters to COBOL programs
- A set of types and constants reflecting external data representations such as can be found in files or databases, allowing COBOL-generated data to be read by an Ada program, and Ada-generated data to be read by COBOL programs
- A generic package for converting between an Ada decimal type value and either an internal or external COBOL representation

### Static Semantics

The library package `Interfaces.COBOL` has the following declaration:

```

package Interfaces.COBOL
  with Preelaborate, Nonblocking, Global => in out synchronized is

  -- Types and operations for internal data representations

  type Floating      is digits implementation-defined;
  type Long_Floating is digits implementation-defined;

```

```

type Binary is range implementation-defined;
type Long_Binary is range implementation-defined;

Max_Digits_Binary : constant := implementation-defined;
Max_Digits_Long_Binary : constant := implementation-defined;

type Decimal_Element is mod implementation-defined;
type Packed_Decimal is array (Positive range <>) of Decimal_Element
    with Pack;

type COBOL_Character is implementation-defined character type;

Ada_To_COBOL : array (Character) of COBOL_Character := implementation-
defined;

COBOL_To_Ada : array (COBOL_Character) of Character := implementation-
defined;

type Alphanumeric is array (Positive range <>) of COBOL_Character
    with Pack;

function To_COBOL (Item : in String) return Alphanumeric;
function To_Ada (Item : in Alphanumeric) return String;

procedure To_COBOL (Item : in String;
                    Target : out Alphanumeric;
                    Last : out Natural);

procedure To_Ada (Item : in Alphanumeric;
                  Target : out String;
                  Last : out Natural);

type Numeric is array (Positive range <>) of COBOL_Character
    with Pack;

-- Formats for COBOL data representations

type Display_Format is private;

Unsigned : constant Display_Format;
Leading_Separate : constant Display_Format;
Trailing_Separate : constant Display_Format;
Leading_Nonseparate : constant Display_Format;
Trailing_Nonseparate : constant Display_Format;

type Binary_Format is private;

High_Order_First : constant Binary_Format;
Low_Order_First : constant Binary_Format;
Native_Binary : constant Binary_Format;

type Packed_Format is private;

Packed_Unsigned : constant Packed_Format;
Packed_Signed : constant Packed_Format;

-- Types for external representation of COBOL binary data

type Byte is mod 2**COBOL_Character'Size;
type Byte_Array is array (Positive range <>) of Byte
    with Pack;

Conversion_Error : exception;

generic
    type Num is delta <> digits <>;
package Decimal_Conversions is
    -- Display Formats: data values are represented as Numeric

    function Valid (Item : in Numeric;
                  Format : in Display_Format) return Boolean;

    function Length (Format : in Display_Format) return Natural;

    function To_Decimal (Item : in Numeric;
                      Format : in Display_Format) return Num;

```

```

function To_Display (Item    : in Num;
                     Format  : in Display_Format) return Numeric;

-- Packed Formats: data values are represented as Packed_Decimal
function Valid (Item    : in Packed_Decimal;
               Format  : in Packed_Format) return Boolean;

function Length (Format : in Packed_Format) return Natural;

function To_Decimal (Item    : in Packed_Decimal;
                   Format  : in Packed_Format) return Num;

function To_Packed (Item    : in Num;
                   Format  : in Packed_Format) return Packed_Decimal;

-- Binary Formats: external data values are represented as Byte_Array
function Valid (Item    : in Byte_Array;
               Format  : in Binary_Format) return Boolean;

function Length (Format : in Binary_Format) return Natural;
function To_Decimal (Item    : in Byte_Array;
                   Format  : in Binary_Format) return Num;

function To_Binary (Item    : in Num;
                   Format  : in Binary_Format) return Byte_Array;

-- Internal Binary formats: data values are of type Binary or
Long_Binary
function To_Decimal (Item : in Binary) return Num;
function To_Decimal (Item : in Long_Binary) return Num;

function To_Binary      (Item : in Num) return Binary;
function To_Long_Binary (Item : in Num) return Long_Binary;

end Decimal_Conversions;

private
... -- not specified by the language
end Interfaces.COBOL;

```

Each of the types in Interfaces.COBOL is COBOL-compatible.

The types Floating and Long\_Floating correspond to the native types in COBOL for data items with computational usage implemented by floating point. The types Binary and Long\_Binary correspond to the native types in COBOL for data items with binary usage, or with computational usage implemented by binary.

Max\_Digits\_Binary is the largest number of decimal digits in a numeric value that is represented as Binary. Max\_Digits\_Long\_Binary is the largest number of decimal digits in a numeric value that is represented as Long\_Binary.

The type Packed\_Decimal corresponds to COBOL's packed-decimal usage.

The type COBOL\_Character defines the run-time character set used in the COBOL implementation. Ada\_To\_COBOL and COBOL\_To\_Ada are the mappings between the Ada and COBOL run-time character sets.

Type Alphanumeric corresponds to COBOL's alphanumeric data category.

Each of the functions To\_COBOL and To\_Ada converts its parameter based on the mappings Ada\_To\_COBOL and COBOL\_To\_Ada, respectively. The length of the result for each is the length of the parameter, and the lower bound of the result is 1. Each component of the result is obtained by applying the relevant mapping to the corresponding component of the parameter.

Each of the procedures To\_COBOL and To\_Ada copies converted elements from Item to Target, using the appropriate mapping (Ada\_To\_COBOL or COBOL\_To\_Ada, respectively). The index in Target of the last element assigned is returned in Last (0 if Item is a null array). If Item'Length exceeds Target'Length, Constraint\_Error is propagated.

Type Numeric corresponds to COBOL's numeric data category with display usage.

The types Display\_Format, Binary\_Format, and Packed\_Format are used in conversions between Ada decimal type values and COBOL internal or external data representations. The value of the constant Native\_Binary is either High\_Order\_First or Low\_Order\_First, depending on the implementation.

```
function Valid (Item    : in Numeric;
                Format   : in Display_Format) return Boolean;
```

The function Valid checks that the Item parameter has a value consistent with the value of Format. If the value of Format is other than Unsigned, Leading\_Separate, and Trailing\_Separate, the effect is implementation defined. If Format does have one of these values, the following rules apply:

- Format=Unsigned: if Item comprises one or more decimal digit characters, then Valid returns True, else it returns False.
- Format=Leading\_Separate: if Item comprises a single occurrence of the plus or minus sign character, and then one or more decimal digit characters, then Valid returns True, else it returns False.
- Format=Trailing\_Separate: if Item comprises one or more decimal digit characters and finally a plus or minus sign character, then Valid returns True, else it returns False.

```
function Length (Format : in Display_Format) return Natural;
```

The Length function returns the minimal length of a Numeric value sufficient to hold any value of type Num when represented as Format.

```
function To_Decimal (Item    : in Numeric;
                    Format   : in Display_Format) return Num;
```

Produces a value of type Num corresponding to Item as represented by Format. The number of digits after the assumed radix point in Item is Num'Scale. Conversion\_Error is propagated if the value represented by Item is outside the range of Num.

```
function To_Display (Item    : in Num;
                    Format   : in Display_Format) return Numeric;
```

This function returns the Numeric value for Item, represented in accordance with Format. The length of the returned value is Length(Format), and the lower bound is 1. Conversion\_Error is propagated if Num is negative and Format is Unsigned.

```
function Valid (Item    : in Packed_Decimal;
                Format   : in Packed_Format) return Boolean;
```

This function returns True if Item has a value consistent with Format, and False otherwise. The rules for the formation of Packed\_Decimal values are implementation defined.

```
function Length (Format : in Packed_Format) return Natural;
```

This function returns the minimal length of a Packed\_Decimal value sufficient to hold any value of type Num when represented as Format.

```
function To_Decimal (Item    : in Packed_Decimal;
                    Format   : in Packed_Format) return Num;
```

Produces a value of type Num corresponding to Item as represented by Format. Num'Scale is the number of digits after the assumed radix point in Item. Conversion\_Error is propagated if the value represented by Item is outside the range of Num.

```
function To_Packed (Item    : in Num;
                   Format   : in Packed_Format) return Packed_Decimal;
```

This function returns the Packed\_Decimal value for Item, represented in accordance with Format. The length of the returned value is Length(Format), and the lower bound is 1. Conversion\_Error is propagated if Num is negative and Format is Packed\_Unsigned.

```
function Valid (Item    : in Byte_Array;
                Format   : in Binary_Format) return Boolean;
```

This function returns True if Item has a value consistent with Format, and False otherwise.

```
function Length (Format : in Binary_Format) return Natural;
```

This function returns the minimal length of a Byte\_Array value sufficient to hold any value of type Num when represented as Format.

```
function To_Decimal (Item    : in Byte_Array;
                    Format   : in Binary_Format) return Num;
```

Produces a value of type Num corresponding to Item as represented by Format. Num'Scale is the number of digits after the assumed radix point in Item. Conversion\_Error is propagated if the value represented by Item is outside the range of Num.

```
function To_Binary (Item    : in Num;
                   Format   : in Binary_Format) return Byte_Array;
```

This function returns the Byte\_Array value for Item, represented in accordance with Format. The length of the returned value is Length(Format), and the lower bound is 1.

```
function To_Decimal (Item : in Binary)      return Num;
```

```
function To_Decimal (Item : in Long_Binary) return Num;
```

These functions convert from COBOL binary format to a corresponding value of the decimal type Num. Conversion\_Error is propagated if Item is too large for Num.

```
function To_Binary      (Item : in Num) return Binary;
```

```
function To_Long_Binary (Item : in Num) return Long_Binary;
```

These functions convert from Ada decimal to COBOL binary format. Conversion\_Error is propagated if the value of Item is too large to be represented in the result type.

#### *Implementation Requirements*

An implementation shall support specifying aspect Convention with a COBOL *convention\_identifier* for a COBOL-eligible type (see B.1).

#### *Implementation Permissions*

An implementation may provide additional constants of the private types Display\_Format, Binary\_Format, or Packed\_Format.

An implementation may provide further floating point and integer types in Interfaces.COBOL to match additional native COBOL types, and may also supply corresponding conversion functions in the generic package Decimal\_Conversions.

#### *Implementation Advice*

An Ada implementation should support the following interface correspondences between Ada and COBOL.

- An Ada **access** T parameter is passed as a “BY REFERENCE” data item of the COBOL type corresponding to T.
- An Ada **in** scalar parameter is passed as a “BY CONTENT” data item of the corresponding COBOL type.
- Any other Ada parameter is passed as a “BY REFERENCE” data item of the COBOL type corresponding to the Ada parameter type; for scalars, a local copy is used if necessary to ensure by-copy semantics.

NOTE 1 An implementation is not required to support specifying aspect Convention for access types, nor is it required to support specifying aspects Import, Export, or Convention for functions.

NOTE 2 If an Ada subprogram is exported to COBOL, then a call from COBOL call can specify either "BY CONTENT" or "BY REFERENCE".

EXAMPLE *Examples of Interfaces.COBOL:*

```

with Interfaces.COBOL;
procedure Test_Call is

    -- Calling a foreign COBOL program
    -- Assume that a COBOL program PROG has the following declaration
    -- in its LINKAGE section:
    -- 01 Parameter-Area
    --     05 NAME      PIC X(20).
    --     05 SSN       PIC X(9).
    --     05 SALARY PIC 99999V99 USAGE COMP.
    -- The effect of PROG is to update SALARY based on some algorithm

package COBOL renames Interfaces.COBOL;

type Salary_Type is delta 0.01 digits 7;

type COBOL_Record is
    record
        Name      : COBOL.Numeric(1..20);
        SSN       : COBOL.Numeric(1..9);
        Salary    : COBOL.Binary;  -- Assume Binary = 32 bits
    end record
    with Convention => COBOL;

procedure Prog (Item : in out COBOL_Record)
    with Import => True, Convention => COBOL;

package Salary_Conversions is
    new COBOL.Decimal_Conversions(Salary_Type);

    Some_Salary : Salary_Type := 12345.67;
    Some_Record : COBOL_Record :=
        (Name => "Johnson, John",
         SSN  => "111223333",
         Salary => Salary_Conversions.To_Binary(Some_Salary));

begin
    Prog (Some_Record);
    ...
end Test_Call;

with Interfaces.COBOL;
with COBOL_Sequential_IO; -- Assumed to be supplied by implementation
procedure Test_External_Formats is

    -- Using data created by a COBOL program
    -- Assume that a COBOL program has created a sequential file with
    -- the following record structure, and that we want
    -- process the records in an Ada program
    -- 01 EMPLOYEE-RECORD
    --     05 NAME      PIC X(20).
    --     05 SSN       PIC X(9).
    --     05 SALARY PIC 99999V99 USAGE COMP.
    --     05 ADJUST PIC S999V999 SIGN LEADING SEPARATE.
    -- The COMP data is binary (32 bits), high-order byte first

package COBOL renames Interfaces.COBOL;

type Salary_Type      is delta 0.01 digits 7;
type Adjustments_Type is delta 0.001 digits 6;

```



```

type COBOL_Employee_Record_Type is -- External representation
  record
    Name      : COBOL.Alphanumeric(1..20);
    SSN       : COBOL.Alphanumeric(1..9);
    Salary    : COBOL.Byte_Array(1..4);
    Adjust    : COBOL.Numeric(1..7); -- Sign and 6 digits
  end record
with Convention => COBOL;

package COBOL_Employee_IO is
  new COBOL_Sequential_IO(COBOL_Employee_Record_Type);
use COBOL_Employee_IO;

COBOL_File : File_Type;

type Ada_Employee_Record_Type is -- Internal representation
  record
    Name      : String(1..20);
    SSN       : String(1..9);
    Salary    : Salary_Type;
    Adjust    : Adjustments_Type;
  end record;

COBOL_Record : COBOL_Employee_Record_Type;
Ada_Record   : Ada_Employee_Record_Type;

package Salary_Conversions is
  new COBOL.Decimal_Conversions(Salary_Type);
use Salary_Conversions;

package Adjustments_Conversions is
  new COBOL.Decimal_Conversions(Adjustments_Type);
use Adjustments_Conversions;

begin
  Open (COBOL_File, Name => "Some_File");

  loop
    Read (COBOL_File, COBOL_Record);

    Ada_Record.Name := COBOL.To_Ada(COBOL_Record.Name);
    Ada_Record.SSN  := COBOL.To_Ada(COBOL_Record.SSN);
    Ada_Record.Salary :=
      To_Decimal(COBOL_Record.Salary, COBOL.High_Order_First);
    Ada_Record.Adjust :=
      To_Decimal(COBOL_Record.Adjust, COBOL.Leading_Separate);
    ... -- Process Ada_Record
  end loop;
exception
  when End_Error => ...
end Test_External_Formats;

```

## B.5 Interfacing with Fortran

The facilities relevant to interfacing with the Fortran language are the package Interfaces.Fortran and support for specifying the Convention aspect with *convention\_identifier* Fortran.

The package Interfaces.Fortran defines Ada types whose representations are identical to the default representations of the Fortran intrinsic types Integer, Real, Double Precision, Complex, Logical, and Character in a supported Fortran implementation. These Ada types can therefore be used to pass objects between Ada and Fortran programs.

### Static Semantics

The library package Interfaces.Fortran has the following declaration:

```

with Ada.Numerics.Generic_Complex_Types; -- see G.1.1
pragma Elaborate_All(Ada.Numerics.Generic_Complex_Types);
package Interfaces.Fortran
  with Pure is

```

```

type Fortran_Integer is range implementation-defined;
type Real is digits implementation-defined;
type Double_Precision is digits implementation-defined;
type Logical is new Boolean;
package Single_Precision_Complex_Types is
    new Ada.Numerics.Generic_Complex_Types (Real);
type Complex is new Single_Precision_Complex_Types.Complex;
subtype Imaginary is Single_Precision_Complex_Types.Imaginary;
i : Imaginary renames Single_Precision_Complex_Types.i;
j : Imaginary renames Single_Precision_Complex_Types.j;
package Double_Precision_Complex_Types is
    new Ada.Numerics.Generic_Complex_Types (Double_Precision);
type Double_Complex is new Double_Precision_Complex_Types.Complex;
subtype Double_Imaginary is Double_Precision_Complex_Types.Imaginary;
type Character_Set is implementation-defined character type;
type Fortran_Character is array (Positive range <>) of Character_Set
    with Pack;

function To_Fortran (Item : in Character) return Character_Set;
function To_Ada (Item : in Character_Set) return Character;

function To_Fortran (Item : in String) return Fortran_Character;
function To_Ada (Item : in Fortran_Character) return String;

procedure To_Fortran (Item : in String;
                     Target : out Fortran_Character;
                     Last : out Natural);

procedure To_Ada (Item : in Fortran_Character;
                  Target : out String;
                  Last : out Natural);

end Interfaces.Fortran;

```

The types Fortran\_Integer, Real, Double\_Precision, Logical, Complex, Double\_Complex, Character\_Set, and Fortran\_Character are Fortran-compatible.

The To\_Fortran and To\_Ada functions map between the Ada type Character and the Fortran type Character\_Set, and also between the Ada type String and the Fortran type Fortran\_Character. The To\_Fortran and To\_Ada procedures have analogous effects to the string conversion subprograms found in Interfaces.COBOL.

#### Implementation Requirements

An implementation shall support specifying aspect Convention with a Fortran *convention\_identifier* for a Fortran-eligible type (see B.1).

#### Implementation Permissions

An implementation may add additional declarations to the Fortran interface packages. For example, declarations are permitted for the character types corresponding to Fortran character kinds 'ascii' and 'iso\_10646', which in turn correspond to ISO/IEC 646:1991 and to UCS-4 as specified in ISO/IEC 10646:2020.

#### Implementation Advice

An Ada implementation should support the following interface correspondences between Ada and Fortran:

- An Ada procedure corresponds to a Fortran subroutine.
- An Ada function corresponds to a Fortran function.

- An Ada parameter of an elementary, array, or record type  $T$  is passed as a  $T_F$  argument to a Fortran procedure, where  $T_F$  is the Fortran type corresponding to the Ada type  $T$ , and where the INTENT attribute of the corresponding dummy argument matches the Ada formal parameter mode; the Fortran implementation's parameter passing conventions are used. For elementary types, a local copy is used if necessary to ensure by-copy semantics.
- An Ada parameter of an access-to-subprogram type is passed as a reference to a Fortran procedure whose interface corresponds to the designated subprogram's specification.

NOTE 1 An object of a Fortran-compatible record type, declared in a library package or subprogram, can correspond to a Fortran common block; the type also corresponds to a Fortran “derived type”.

NOTE 2 For Fortran facilities not addressed by this subclause, consider using the Fortran to C interoperability features defined in ISO/IEC 1594-1:2018 along with the C interfacing features defined in B.3.

EXAMPLE *Example of Interfaces.Fortran:*

```
with Interfaces.Fortran;
use Interfaces.Fortran;
procedure Ada_Application is
  type Fortran_Matrix is
    array (Fortran_Integer range <>,
           Fortran_Integer range <>) of Double_Precision
    with Convention => Fortran;
    -- stored in Fortran's
    -- column-major order
  procedure Invert (Rank : in Fortran_Integer; X : in out Fortran_Matrix)
    with Import => True, Convention => Fortran; -- a Fortran subroutine
  Rank      : constant Fortran_Integer := 100;
  My_Matrix : Fortran_Matrix (1 .. Rank, 1 .. Rank);
  Precision : constant := 6;
  type Standard_Deviation is digits Precision
    with Convention => Fortran;
  Deviation : Standard_Deviation;
  -- Declarations to match the following Fortran declarations:
  --   integer, parameter :: precision = selected_real_kind(p=6)
  --   real(precision) :: deviation
begin
  ...
  My_Matrix := ...;
  ...
  Invert (Rank, My_Matrix);
  ...
  Deviation := ...;
  ...
end Ada_Application;
```

## Annex C (informative)

### Systems Programming

The Systems Programming Annex specifies additional capabilities provided for low-level programming. These capabilities are also required in many real-time, embedded, distributed, and information systems.

#### C.1 Access to Machine Operations

This subclause specifies rules regarding access to machine instructions from within an Ada program.

##### *Implementation Requirements*

The implementation shall support machine code insertions (see 16.8) or intrinsic subprograms (see 9.3.1) (or both). The implementation shall allow the use of Ada entities as operands for such machine code insertions or intrinsic subprograms.

##### *Implementation Advice*

The machine code or intrinsics support should allow access to all operations normally available to assembly language programmers for the target environment, including privileged instructions, if any.

The support for interfacing aspects (see Annex B) should include interface to assembler; the default assembler should be associated with the convention identifier Assembler.

If an entity is exported to assembly language, then the implementation should allocate it at an addressable location, and should ensure that it is retained by the linking process, even if not otherwise referenced from the Ada code. The implementation should assume that any call to a machine code or assembler subprogram is allowed to read or update every object that is specified as exported.

##### *Documentation Requirements*

The implementation shall document the overhead associated with calling machine-code or intrinsic subprograms, as compared to a fully-inlined call, and to a regular out-of-line call.

The implementation shall document the types of the package `System.Machine_Code` usable for machine code insertions, and the attributes to be used in machine code insertions for references to Ada entities.

The implementation shall document the subprogram calling conventions associated with the convention identifiers available for use with the Convention aspect (Ada and Assembler, at a minimum), including register saving, exception propagation, parameter passing, and function value returning.

For exported and imported subprograms, the implementation shall document the mapping between the `Link_Name` string, if specified, or the Ada designator, if not, and the external link name used for such a subprogram.

##### *Implementation Advice*

The implementation should ensure that little or no overhead is associated with calling intrinsic and machine-code subprograms.

It is recommended that intrinsic subprograms be provided for convenient access to any machine operations that provide special capabilities or efficiency and that are not otherwise available through the language constructs. Examples of such instructions include:

- Atomic read-modify-write operations — e.g. test and set, compare and swap, decrement and test, enqueue/dequeue.
- Standard numeric functions — e.g. *sin*, *log*.
- String manipulation operations — e.g. translate and test.
- Vector operations — e.g. compare vector against thresholds.
- Direct operations on I/O ports.

## C.2 Required Representation Support

This subclause specifies minimal requirements on the support for representation items and related features.

### Implementation Requirements

The implementation shall support at least the functionality defined by the recommended levels of support in Clause 16.

## C.3 Interrupt Support

This subclause specifies the language-defined model for hardware interrupts in addition to mechanisms for handling interrupts.

### Dynamic Semantics

An *interrupt* represents a class of events that are detected by the hardware or the system software. Interrupts are said to occur. An *occurrence* of an interrupt is separable into generation and delivery. *Generation* of an interrupt is the event in the underlying hardware or system that makes the interrupt available to the program. *Delivery* is the action that invokes part of the program as response to the interrupt occurrence. Between generation and delivery, the interrupt occurrence (or interrupt) is *pending*. Some or all interrupts may be *blocked*. When an interrupt is blocked, all occurrences of that interrupt are prevented from being delivered. Certain interrupts are *reserved*. The set of reserved interrupts is implementation defined. A reserved interrupt is either an interrupt for which user-defined handlers are not supported, or one which already has an attached handler by some other implementation-defined means. Program units can be connected to nonreserved interrupts. While connected, the program unit is said to be *attached* to that interrupt. The execution of that program unit, the *interrupt handler*, is invoked upon delivery of the interrupt occurrence.

While a handler is attached to an interrupt, it is called once for each delivered occurrence of that interrupt. While the handler executes, the corresponding interrupt is blocked.

While an interrupt is blocked, all occurrences of that interrupt are prevented from being delivered. Whether such occurrences remain pending or are lost is implementation defined.

Each interrupt has a *default treatment* which determines the system's response to an occurrence of that interrupt when no user-defined handler is attached. The set of possible default treatments is implementation defined, as is the method (if one exists) for configuring the default treatments for interrupts.

An interrupt is delivered to the handler (or default treatment) that is in effect for that interrupt at the time of delivery.

An exception propagated from a handler that is invoked by an interrupt has no effect.

If the Ceiling\_Locking policy (see D.3) is in effect, the interrupt handler executes with the active priority that is the ceiling priority of the corresponding protected object.

*Implementation Requirements*

The implementation shall provide a mechanism to determine the minimum stack space that is necessary for each interrupt handler and to reserve that space for the execution of the handler. This space should accommodate nested invocations of the handler where the system permits this.

If the hardware or the underlying system holds pending interrupt occurrences, the implementation shall provide for later delivery of these occurrences to the program.

If the Ceiling Locking policy is not in effect, the implementation shall provide means for the application to specify whether interrupts are to be blocked during protected actions.

*Documentation Requirements*

The implementation shall document the following items:

1. For each interrupt, which interrupts are blocked from delivery when a handler attached to that interrupt executes (either as a result of an interrupt delivery or of an ordinary call on a procedure of the corresponding protected object).
2. Any interrupts that cannot be blocked, and the effect of attaching handlers to such interrupts, if this is permitted.
3. Which run-time stack an interrupt handler uses when it executes as a result of an interrupt delivery; if this is configurable, what is the mechanism to do so; how to specify how much space to reserve on that stack.
4. Any implementation- or hardware-specific activity that happens before a user-defined interrupt handler gets control (e.g. reading device registers, acknowledging devices).
5. Any timing or other limitations imposed on the execution of interrupt handlers.
6. The state (blocked/unblocked) of the nonreserved interrupts when the program starts; if some interrupts are unblocked, what is the mechanism a program can use to protect itself before it can attach the corresponding handlers.
7. Whether the interrupted task is allowed to resume execution before the interrupt handler returns.
8. The treatment of interrupt occurrences that are generated while the interrupt is blocked; i.e. whether one or more occurrences are held for later delivery, or all are lost.
9. Whether predefined or implementation-defined exceptions are raised as a result of the occurrence of any interrupt, and the mapping between the machine interrupts (or traps) and the predefined exceptions.
10. On a multi-processor, the rules governing the delivery of an interrupt to a particular processor.

*Implementation Permissions*

If the underlying system or hardware does not allow interrupts to be blocked, then no blocking is required as part of the execution of subprograms of a protected object for which one of its subprograms is an interrupt handler.

In a multi-processor with more than one interrupt subsystem, it is implementation defined whether (and how) interrupt sources from separate subsystems share the same Interrupt\_Id type (see C.3.2). In particular, the meaning of a blocked or pending interrupt may then be applicable to one processor only.

Implementations are allowed to impose timing or other limitations on the execution of interrupt handlers.

Other forms of handlers are allowed to be supported, in which case the rules of this subclause should be adhered to.

The active priority of the execution of an interrupt handler is allowed to vary from one occurrence of the same interrupt to another.



*Implementation Advice*

If the Ceiling\_Locking policy is not in effect, the implementation should provide means for the application to specify which interrupts are to be blocked during protected actions, if the underlying system allows for finer-grained control of interrupt blocking.

NOTE 1 The default treatment for an interrupt can be to keep the interrupt pending or to deliver it to an implementation-defined handler. Examples of actions that an implementation-defined handler can perform include aborting the partition, ignoring (i.e. discarding occurrences of) the interrupt, or queuing one or more occurrences of the interrupt for possible later delivery when a user-defined handler is attached to that interrupt.

NOTE 2 It is a bounded error to call Task\_Identification.Current\_Task (see C.7.1) from an interrupt handler.

NOTE 3 The rule that an exception propagated from an interrupt handler has no effect is modeled after the rule about exceptions propagated out of task bodies.

**C.3.1 Protected Procedure Handlers***Static Semantics*

For a parameterless protected procedure, the following language-defined representation aspects may be specified:

**Interrupt\_Handler**

The type of aspect Interrupt\_Handler is Boolean. If directly specified, the aspect\_definition shall be a static expression. This aspect is never inherited; if not directly specified, the aspect is False.

**Attach\_Handler**

The aspect Attach\_Handler is an expression, which shall be of type Interrupts.Interrupt\_Id. This aspect is never inherited.

*Legality Rules*

If either the Attach\_Handler or Interrupt\_Handler aspect are specified for a protected procedure, the corresponding protected\_type\_declaration or single\_protected\_declaration shall be a library-level declaration and shall not be declared within a generic body. In addition to the places where Legality Rules normally apply (see 15.3), this rule also applies in the private part of an instance of a generic unit.

*Dynamic Semantics*

If the Interrupt\_Handler aspect of a protected procedure is True, then the procedure may be attached dynamically, as a handler, to interrupts (see C.3.2). Such procedures are allowed to be attached to multiple interrupts.

The expression specified for the Attach\_Handler aspect of a protected procedure *P* is evaluated as part of the creation of the protected object that contains *P*. The value of the expression identifies an interrupt. As part of the initialization of that object, *P* (the *handler* procedure) is attached to the identified interrupt. A check is made that the corresponding interrupt is not reserved. Program\_Error is raised if the check fails, and the existing treatment for the interrupt is not affected.

If the Ceiling\_Locking policy (see D.3) is in effect, then upon the initialization of a protected object that contains a protected procedure for which either the Attach\_Handler aspect is specified or the Interrupt\_Handler aspect is True, a check is made that the initial ceiling priority of the object is in the range of System.Interrupt\_Priority. If the check fails, Program\_Error is raised.

When a protected object is finalized, for any of its procedures that are attached to interrupts, the handler is detached. If the handler was attached by a procedure in the Interrupts package or if no user handler was previously attached to the interrupt, the default treatment is restored. If the Attach\_Handler aspect was specified and the most recently attached handler for the same interrupt is the same as the one that was attached at the time the protected object was initialized, the previous handler is restored.

When a handler is attached to an interrupt, the interrupt is blocked (subject to the Implementation Permission in C.3) during the execution of every protected action on the protected object containing the handler.

If restriction `No_Dynamic_Attachment` is in effect, then a check is made that the interrupt identified by an `Attach_Handler` aspect does not appear in any previously elaborated `Attach_Handler` aspect; `Program_Error` is raised if this check fails.

#### *Erroneous Execution*

If the `Ceiling_Locking` policy (see D.3) is in effect and an interrupt is delivered to a handler, and the interrupt hardware priority is higher than the ceiling priority of the corresponding protected object, the execution of the program is erroneous.

If the handlers for a given interrupt attached via aspect `Attach_Handler` are not attached and detached in a stack-like (LIFO) order, program execution is erroneous. In particular, when a protected object is finalized, the execution is erroneous if any of the procedures of the protected object are attached to interrupts via aspect `Attach_Handler` and the most recently attached handler for the same interrupt is not the same as the one that was attached at the time the protected object was initialized.

#### *Metrics*

The following metric shall be documented by the implementation:

- The worst-case overhead for an interrupt handler that is a parameterless protected procedure, in clock cycles. This is the execution time not directly attributable to the handler procedure or the interrupted execution. It is estimated as  $C - (A+B)$ , where A is how long it takes to complete a given sequence of instructions without any interrupt, B is how long it takes to complete a normal call to a given protected procedure, and C is how long it takes to complete the same sequence of instructions when it is interrupted by one execution of the same procedure called via an interrupt.

#### *Implementation Permissions*

When the aspects `Attach_Handler` or `Interrupt_Handler` are specified for a protected procedure, the implementation is allowed to impose implementation-defined restrictions on the corresponding `protected_type_declaration` and `protected_body`.

An implementation may use a different mechanism for invoking a protected procedure in response to a hardware interrupt than is used for a call to that protected procedure from a task.

Notwithstanding what this subclause says elsewhere, the `Attach_Handler` and `Interrupt_Handler` aspects are allowed to be used for other, implementation defined, forms of interrupt handlers.

#### *Implementation Advice*

Whenever possible, the implementation should allow interrupt handlers to be called directly by the hardware.

Whenever practical, the implementation should detect violations of any implementation-defined restrictions before run time.

NOTE 1 The `Attach_Handler` aspect can provide static attachment of handlers to interrupts if the implementation supports preelaboration of protected objects. (See C.4.)

NOTE 2 For a protected object that has a (protected) procedure attached to an interrupt, the correct ceiling priority is at least as high as the highest processor priority at which that interrupt will ever be delivered.

NOTE 3 Protected procedures can also be attached dynamically to interrupts via operations declared in the predefined package `Interrupts`.

NOTE 4 An example of a possible implementation-defined restriction is disallowing the use of the standard storage pools within the body of a protected procedure that is an interrupt handler.

### C.3.2 The Package Interrupts

#### Static Semantics

The following language-defined packages exist:

```

with System;
with System.Multiprocessors;
package Ada.Interrupts
  with Nonblocking, Global => in out synchronized is
    type Interrupt_Id is implementation-defined;
    type Parameterless_Handler is
      access protected procedure
        with Nonblocking => False;

    function Is_Reserved (Interrupt : Interrupt_Id)
      return Boolean;

    function Is_Attached (Interrupt : Interrupt_Id)
      return Boolean;

    function Current_Handler (Interrupt : Interrupt_Id)
      return Parameterless_Handler;

    procedure Attach_Handler
      (New_Handler : in Parameterless_Handler;
       Interrupt    : in Interrupt_Id);

    procedure Exchange_Handler
      (Old_Handler : out Parameterless_Handler;
       New_Handler : in Parameterless_Handler;
       Interrupt    : in Interrupt_Id);

    procedure Detach_Handler
      (Interrupt : in Interrupt_Id);

    function Reference (Interrupt : Interrupt_Id)
      return System.Address;

    function Get_CPU (Interrupt : Interrupt_Id)
      return System.Multiprocessors.CPU_Range;

  private
    ... -- not specified by the language
end Ada.Interrupts;

package Ada.Interrupts.Names
  with Nonblocking, Global => null is
    implementation-defined : constant Interrupt_Id :=
      implementation-defined;
    . . .
    implementation-defined : constant Interrupt_Id :=
      implementation-defined;
end Ada.Interrupts.Names;

```

#### Dynamic Semantics

The `Interrupt_Id` type is an implementation-defined discrete type used to identify interrupts.

The `Is_Reserved` function returns `True` if and only if the specified interrupt is reserved.

The `Is_Attached` function returns `True` if and only if a user-specified interrupt handler is attached to the interrupt.

The `Current_Handler` function returns a value that represents the attached handler of the interrupt. If no user-defined handler is attached to the interrupt, `Current_Handler` returns **null**.

The `Attach_Handler` procedure attaches the specified handler to the interrupt, overriding any existing treatment (including a user handler) in effect for that interrupt. If `New_Handler` is **null**, the default treatment is restored. If `New_Handler` designates a protected procedure for which the aspect

Interrupt\_Handler is False, Program\_Error is raised. In this case, the operation does not modify the existing interrupt treatment.

The Exchange\_Handler procedure operates in the same manner as Attach\_Handler with the addition that the value returned in Old\_Handler designates the previous treatment for the specified interrupt. If the previous treatment is not a user-defined handler, **null** is returned.

The Detach\_Handler procedure restores the default treatment for the specified interrupt.

For all operations defined in this package that take a parameter of type Interrupt\_Id, with the exception of Is\_Reserved and Reference, a check is made that the specified interrupt is not reserved. Program\_Error is raised if this check fails.

If, by using the Attach\_Handler, Detach\_Handler, or Exchange\_Handler procedures, an attempt is made to detach a handler that was attached statically (using the aspect Attach\_Handler), the handler is not detached and Program\_Error is raised.

The Reference function returns a value of type System.Address that can be used to attach a task entry via an address clause (see I.7.1) to the interrupt specified by Interrupt. This function raises Program\_Error if attaching task entries to interrupts (or to this particular interrupt) is not supported.

The function Get\_CPU returns the processor on which the handler for Interrupt is executed. If the handler can execute on more than one processor the value System.Multiprocessors.Not\_A\_Specific\_CPU is returned.

#### *Implementation Requirements*

At no time during attachment or exchange of handlers shall the current handler of the corresponding interrupt be undefined.

#### *Documentation Requirements*

The implementation shall document, when the Ceiling\_Locking policy (see D.3) is in effect, the default ceiling priority assigned to a protected object that contains a protected procedure that specifies either the Attach\_Handler or Interrupt\_Handler aspects, but does not specify the Interrupt\_Priority aspect. This default can be different for different interrupts.

#### *Implementation Advice*

If implementation-defined forms of interrupt handler procedures are supported, such as protected procedures with parameters, then for each such form of a handler, a type analogous to Parameterless\_Handler should be specified in a child package of Interrupts, with the same operations as in the predefined package Interrupts.

NOTE The package Interrupts.Names contains implementation-defined names (and constant values) for the interrupts that are supported by the implementation.

EXAMPLE *Example of interrupt handlers:*

```
Device_Priority : constant
array (Ada.Interrupts.Interrupt_Id range 1..5) of
  System.Interrupt_Priority := ( ... );
protected type Device_Interface
  (Int_Id : Ada.Interrupts.Interrupt_Id)
    with Interrupt_Priority => Device_Priority(Int_Id) is
    procedure Handler
      with Attach_Handler => Int_Id;
  ...
end Device_Interface;
...
Device_1_Driver : Device_Interface(1);
...
Device_5_Driver : Device_Interface(5);
...
```

## C.4 Preelaboration Requirements

This subclause specifies additional implementation and documentation requirements for the Preelaborate aspect (see 13.2.1).

### Implementation Requirements

The implementation shall not incur any run-time overhead for the elaboration checks of subprograms and `protected_bodies` declared in preelaborated library units.

The implementation shall not execute any memory write operations after load time for the elaboration of constant objects declared immediately within the declarative region of a preelaborated library package, so long as the subtype and initial expression (or default initial expressions if initialized by default) of the `object_declaration` satisfy the following restrictions. The meaning of *load time* is implementation defined.

- Any `subtype_mark` denotes a statically constrained subtype, with statically constrained subcomponents, if any;
- no `subtype_mark` denotes a controlled type, a private type, a private extension, a generic formal private type, a generic formal derived type, or a descendant of such a type;
- any `constraint` is a static constraint;
- any `allocator` is for an access-to-constant type;
- any uses of predefined operators appear only within static expressions;
- any `primaries` that are names, other than `attribute_references` for the Access or Address attributes, appear only within static expressions;
- any name that is not part of a static expression is an expanded name or `direct_name` that statically names some entity;
- any `discrete_choice` of an `array_aggregate` is static;
- no language-defined check associated with the elaboration of the `object_declaration` can fail.

### Documentation Requirements

The implementation shall document any circumstances under which the elaboration of a preelaborated package causes code to be executed at run time.

The implementation shall document whether the method used for initialization of preelaborated variables allows a partition to be restarted without reloading.

### Implementation Advice

It is recommended that preelaborated packages be implemented in such a way that there should be little or no code executed at run time for the elaboration of entities not already covered by the Implementation Requirements.

## C.5 Aspect Discard\_Names

Specifying the aspect `Discard_Names` can be used to request a reduction in storage used for the names of entities with runtime name text.

### Static Semantics

An entity with *runtime name text* is a nonderived enumeration first subtype, a tagged first subtype, or an exception.

For an entity with runtime name text, the following language-defined representation aspect may be specified:

## Discard\_Names

The type of aspect Discard\_Names is Boolean. If directly specified, the `aspect_definition` shall be a static expression. If not specified (including by inheritance), the aspect is False.

### Syntax

The form of a `pragma Discard_Names` is as follows:

**pragma** Discard\_Names[[On => ] `local_name`];

A `pragma Discard_Names` is allowed only immediately within a `declarative_part`, immediately within a `package_specification`, or as a configuration pragma.

### Legality Rules

The `local_name` (if present) shall denote an entity with runtime name text. The pragma specifies that the aspect Discard\_Names for the type or exception has the value True. Without a `local_name`, the pragma specifies that all entities with runtime name text declared after the pragma, within the same declarative region have the value True for aspect Discard\_Names. Alternatively, the pragma can be used as a configuration pragma. If the configuration pragma Discard\_Names applies to a compilation unit, all entities with runtime name text declared in the compilation unit have the value True for the aspect Discard\_Names.

### Static Semantics

If a `local_name` is given, then a `pragma Discard_Names` is a representation pragma.

If the aspect Discard\_Names is True for an enumeration type, then the semantics of the `Wide_Wide_Image` and `Wide_Wide_Value` attributes are implementation defined for that type; the semantics of `Image`, `Wide_Image`, `Value`, and `Wide_Value` are still defined in terms of `Wide_Wide_Image` and `Wide_Wide_Value`. In addition, the semantics of `Text_IO Enumeration_IO` are implementation defined. If the aspect Discard\_Names is True for a tagged type, then the semantics of the `Tags.-Wide_Wide_Expanded_Name` function are implementation defined for that type; the semantics of `Tags.-Expanded_Name` and `Tags.Wide_Expanded_Name` are still defined in terms of `Tags.Wide_Wide_Expanded_Name`. If the aspect Discard\_Names is True for an exception, then the semantics of the `Exceptions.Wide_Wide_Exception_Name` function are implementation defined for that exception; the semantics of `Exceptions.Exception_Name` and `Exceptions.Wide_Exception_Name` are still defined in terms of `Exceptions.Wide_Wide_Exception_Name`.

### Implementation Advice

If the aspect Discard\_Names is True for an entity, then the implementation should reduce the amount of storage used for storing names associated with that entity.

## C.6 Shared Variable Control

This subclause defines representation aspects that control the use of shared variables.

### Static Semantics

For an `object_declaration`, a `component_declaration`, a `full_type_declaration`, or a `formal_complete_type_declaration`, the following representation aspects may be specified:

**Atomic** The type of aspect Atomic is Boolean.

**Independent** The type of aspect Independent is Boolean.

**Volatile** The type of aspect Volatile is Boolean.

**Full\_Access\_Only**

The type of aspect Full\_Access\_Only is Boolean.



For a `full_type_declaration` of an array type, an `object_declaration` for an object of an anonymous array type, or the `formal_complete_type_declaration` of a formal array type, the following representation aspects may be specified:

#### Atomic\_Components

The type of aspect `Atomic_Components` is Boolean.

#### Volatile\_Components

The type of aspect `Volatile_Components` is Boolean.

For a `full_type_declaration` of a composite type, an `object_declaration` for an object of an anonymous composite type, or the `formal_complete_type_declaration` of a formal composite type, the following representation aspect may be specified:

#### Independent\_Components

The type of aspect `Independent_Components` is Boolean.

If any of these aspects are directly specified, the `aspect_definition` shall be a static expression. If not specified for a type (including by inheritance), the `Atomic`, `Atomic_Components`, and `Full_Access_Only` aspects are False. If any of these aspects are specified True for a type, then the corresponding aspect is True for all objects of the type. If the `Atomic` aspect is specified True, then the aspects `Volatile`, `Independent`, and `Volatile_Component` (if defined) are True; if the `Atomic_Components` aspect is specified True, then the aspects `Volatile`, `Volatile_Components`, and `Independent_Components` are True. If the `Volatile` aspect is specified True, then the `Volatile_Components` aspect (if defined) is True, and vice versa. When not determined by one of the other aspects, or for an object by its type, the `Volatile`, `Volatile_Components`, `Independent`, and `Independent_Components` aspects are False.

An *atomic* type is one for which the aspect `Atomic` is True. An *atomic* object (including a component) is one for which the aspect `Atomic` is True, or a component of an array for which the aspect `Atomic_Components` is True for the associated type, or any object of an atomic type, other than objects obtained by evaluating a slice.

A *volatile* type is one for which the aspect `Volatile` is True. A *volatile* object (including a component) is one for which the aspect `Volatile` is True, or a component of an array for which the aspect `Volatile_Components` is True for the associated type, or any object of a volatile type. In addition, every atomic type or object is also defined to be volatile. Finally, if an object is volatile, then so are all of its subcomponents (the same does not apply to atomic).

When True, the aspects `Independent` and `Independent_Components` *specify as independently addressable* the named object or component(s), or in the case of a type, all objects or components of that type. All atomic objects and aliased objects are considered to be specified as independently addressable.

The `Full_Access_Only` aspect shall not be specified unless the associated type or object is volatile (or atomic). A *full access* type is any atomic type, or a volatile type for which the aspect `Full_Access_Only` is True. A *full access* object (including a component) is any atomic object, or a volatile object for which the aspect `Full_Access_Only` is True for the object or its type. A `Full_Access_Only` aspect is illegal if any subcomponent of the object or type is a full access object or is of a generic formal type.

#### Legality Rules

If aspect `Independent_Components` is specified for a `full_type_declaration`, the declaration shall be that of an array or record type.

It is illegal to specify either of the aspects `Atomic` or `Atomic_Components` to have the value True for an object or type if the implementation cannot support the indivisible and independent reads and updates required by the aspect (see below).

It is illegal to specify the `Size` attribute of an atomic object, the `Component_Size` attribute for an array type with atomic components, or the layout attributes of an atomic component, in a way that prevents the implementation from performing the required indivisible and independent reads and updates.

If an atomic object is passed as a parameter, then the formal parameter shall either have an atomic type or allow pass by copy. If an atomic object is used as an actual for a generic formal object of mode **in out**, then the type of the generic formal object shall be atomic. If the `prefix` of an `attribute_reference` for an Access attribute denotes an atomic object (including a component), then the designated type of the resulting access type shall be atomic. Corresponding rules apply to volatile objects and to full access objects.

If a nonatomic subcomponent of a full access object is passed as an actual parameter in a call then the formal parameter shall allow pass by copy (and, at run time, the parameter shall be passed by copy). A nonatomic subcomponent of a full access object shall not be used as an actual for a generic formal of mode **in out**. The `prefix` of an `attribute_reference` for an Access attribute shall not denote a nonatomic subcomponent of a full access object.

If the Atomic, Atomic\_Components, Volatile, Volatile\_Components, Independent, Independent\_Components, or Full\_Access\_Only aspect is True for a generic formal type, then that aspect shall be True for the actual type. If an atomic type is used as an actual for a generic formal derived type, then the ancestor of the formal type shall be atomic. A corresponding rule applies to volatile types and similarly to full access types.

If a type with volatile components is used as an actual for a generic formal array type, then the components of the formal type shall be volatile. Furthermore, if the actual type has atomic components and the formal array type has aliased components, then the components of the formal array type shall also be atomic. A corresponding rule applies when the actual type has volatile full access components.

If an aspect Volatile, Volatile\_Components, Atomic, or Atomic\_Components is directly specified to have the value True for a stand-alone constant object, then the aspect Import shall also be specified as True for it.

It is illegal to specify the aspect Independent or Independent\_Components as True for a component, object or type if the implementation cannot provide the independent addressability required by the aspect (see 12.10).

It is illegal to specify a representation aspect for a component, object or type for which the aspect Independent or Independent\_Components is True, in a way that prevents the implementation from providing the independent addressability required by the aspect.

#### *Dynamic Semantics*

For an atomic object (including an atomic component) all reads and updates of the object as a whole are indivisible.

All tasks of the program (on all processors) that read or update volatile variables see the same order of updates to the variables. A use of an atomic variable or other mechanism may be necessary to avoid erroneous execution and to ensure that access to nonatomic volatile variables is sequential (see 12.10).

Two actions are sequential (see 12.10) if each is the read or update of the same atomic object.

If a type is atomic or volatile and it is not a by-copy type, then the type is defined to be a by-reference type. If any subcomponent of a type is atomic or volatile, then the type is defined to be a by-reference type.

If an actual parameter is atomic or volatile, and the corresponding formal parameter is not, then the parameter is passed by copy.

All reads of or writes to any nonatomic subcomponent of a full access object are performed by reading and/or writing all of the nearest enclosing full access object.

#### *Implementation Requirements*

The external effect of a program (see 4.2) is defined to include each read and update of a volatile or atomic object. The implementation shall not generate any memory reads or updates of atomic or volatile objects other than those specified by the program. However, there may be target-dependent cases where reading or writing a volatile but nonatomic object (typically a component) necessarily

involves reading and/or writing neighboring storage, and that neighboring storage can overlap a volatile object.

*Implementation Permissions*

Within the body of an instance of a generic unit that has a formal type  $T$  that is not atomic and an actual type that is atomic, if an object  $O$  of type  $T$  is declared and explicitly specified as atomic, the implementation may introduce an additional copy on passing  $O$  to a subprogram with a parameter of type  $T$  that is normally passed by reference. A corresponding permission applies to volatile parameter passing.

*Implementation Advice*

A load or store of a volatile object whose size is a multiple of `System.Storage_Unit` and whose alignment is nonzero, should be implemented by accessing exactly the bits of the object and no others, except in the case of a volatile but nonatomic subcomponent of an atomic object.

A load or store of an atomic object should, where possible, be implemented by a single load or store instruction.

NOTE 1 An imported volatile or atomic constant behaves as a constant (i.e. read-only) with respect to other parts of the Ada program, but can still be modified by an “external source”.

NOTE 2 Specifying the Pack aspect cannot override the effect of specifying an Atomic or Atomic\_Components aspect.

NOTE 3 When mapping an Ada object to a memory-mapped hardware register, the Ada object can be declared atomic to ensure that the compiler will read and write exactly the bits of the register as specified in the source code and no others.

### C.6.1 The Package `System.Atomic_Operations`

The language-defined package `System.Atomic_Operations` is the parent of a set of child units that provide facilities for manipulating objects of atomic types and for supporting lock-free synchronization. The subprograms of this subsystem are Intrinsic subprograms (see 9.3.1) in order to provide convenient access to machine operations that can provide these capabilities if they are available in the target environment.

*Static Semantics*

The library package `System.Atomic_Operations` has the following declaration:

```
package System.Atomic_Operations
  with Pure, Nonblocking is
end System.Atomic_Operations;
```

`System.Atomic_Operations` serves as the parent of other language-defined library units that manipulate atomic objects; its declaration is empty.

A call to a subprogram is said to be *lock-free* if the subprogram is guaranteed to return from the call while keeping the processor of the logical thread of control busy for the duration of the call.

In each child package, a function `Is_Lock_Free(...)` is provided to check whether the operations of the child package can all be provided lock-free for a given object. `Is_Lock_Free` returns `True` if operations defined in the child package are lock-free when applied to the object denoted by `Item`, and `Is_Lock_Free` returns `False` otherwise.

### C.6.2 The Package `System.Atomic_Operations.Exchange`

The language-defined generic package `System.Atomic_Operations.Exchange` provides the following operations:

- To atomically compare the value of two atomic objects, and update the first atomic object with a desired value if both objects were found to be equal, or otherwise update the second object with the value of the first object.

- To atomically update the value of an atomic object, and then return the value that the atomic object had just prior to the update.

#### Static Semantics

The generic library package `System.Atomic_Operations.Exchange` has the following declaration:

```
generic
  type Atomic_Type is private with Atomic;
package System.Atomic_Operations.Exchange
  with Pure, Nonblocking is

    function Atomic_Exchange (Item : aliased in out Atomic_Type;
                             Value : Atomic_Type) return Atomic_Type
      with Convention => Intrinsic;

    function Atomic_Compare_And_Exchange
      (Item      : aliased in out Atomic_Type;
       Prior     : aliased in out Atomic_Type;
       Desired   : Atomic_Type) return Boolean
      with Convention => Intrinsic;

    function Is_Lock_Free (Item : aliased Atomic_Type) return Boolean
      with Convention => Intrinsic;

end System.Atomic_Operations.Exchange;
```

`Atomic_Exchange` atomically assigns the value of `Value` to `Item`, and returns the previous value of `Item`.

`Atomic_Compare_And_Exchange` first evaluates the value of `Prior`. `Atomic_Compare_And_Exchange` then performs the following steps as part of a single indivisible operation:

- evaluates the value of `Item`;
- compares the value of `Item` with the value of `Prior`;
- if equal, assigns `Item` the value of `Desired`;
- otherwise, makes no change to the value of `Item`.

After these steps, if the value of `Item` and `Prior` did not match, `Prior` is assigned the original value of `Item`, and the function returns `False`. Otherwise, `Prior` is unaffected and the function returns `True`.

**EXAMPLE** *Example of a spin lock using `Atomic_Exchange`:*

```
type Atomic_Boolean is new Boolean with Atomic;
package Exchange is new
  Atomic_Operations.Exchange (Atomic_Type => Atomic_Boolean);

Lock : aliased Atomic_Boolean := False;

...

begin -- Some critical section, trying to get the lock:
  -- Obtain the lock
  while Exchange.Atomic_Exchange (Item => Lock, Value => True) loop
    null;
  end loop;

  ... -- Do stuff

  Lock := False; -- Release the lock
end;
```

### C.6.3 The Package `System.Atomic_Operations.Test_and_Set`

The language-defined package `System.Atomic_Operations.Test_And_Set` provides an operation to atomically set and clear an atomic flag object.

*Static Semantics*

The library package `System.Atomic_Operations.Test_And_Set` has the following declaration:

```
package System.Atomic_Operations.Test_And_Set
with Pure, Nonblocking is

  type Test_And_Set_Flag is mod implementation-defined
    with Atomic, Default_Value => 0, Size => implementation-defined;

  function Atomic_Test_And_Set
    (Item : aliased in out Test_And_Set_Flag) return Boolean
    with Convention => Intrinsic;

  procedure Atomic_Clear
    (Item : aliased in out Test_And_Set_Flag)
    with Convention => Intrinsic;

  function Is_Lock_Free
    (Item : aliased Test_And_Set_Flag) return Boolean
    with Convention => Intrinsic;

end System.Atomic_Operations.Test_And_Set;
```

`Test_And_Set_Flag` represents the state of an atomic flag object. An atomic flag object can either be considered to be set or cleared.

`Atomic_Test_And_Set` performs an atomic test-and-set operation on `Item`. `Item` is set to some implementation-defined nonzero value. The function returns `True` if the previous contents were nonzero, and otherwise returns `False`.

`Atomic_Clear` performs an atomic clear operation on `Item`. After the operation, `Item` contains 0. This call should be used in conjunction with `Atomic_Test_And_Set`.

#### C.6.4 The Package `System.Atomic_Operations.Integer_Arithmetic`

The language-defined generic package `System.Atomic_Operations.Integer_Arithmetic` provides operations to perform arithmetic atomically on objects of integer types.

*Static Semantics*

The generic library package `System.Atomic_Operations.Integer_Arithmetic` has the following declaration:

```
generic
  type Atomic_Type is range <> with Atomic;
package System.Atomic_Operations.Integer_Arithmetic
with Pure, Nonblocking is

  procedure Atomic_Add (Item : aliased in out Atomic_Type;
    Value : Atomic_Type)
    with Convention => Intrinsic;

  procedure Atomic_Subtract (Item : aliased in out Atomic_Type;
    Value : Atomic_Type)
    with Convention => Intrinsic;

  function Atomic_Fetch_And_Add
    (Item : aliased in out Atomic_Type;
    Value : Atomic_Type) return Atomic_Type
    with Convention => Intrinsic;

  function Atomic_Fetch_And_Subtract
    (Item : aliased in out Atomic_Type;
    Value : Atomic_Type) return Atomic_Type
    with Convention => Intrinsic;

  function Is_Lock_Free (Item : aliased Atomic_Type) return Boolean
    with Convention => Intrinsic;

end System.Atomic_Operations.Integer_Arithmetic;
```

The operations of this package are defined as follows:

```
procedure Atomic_Add (Item  : aliased in out Atomic_Type;
                      Value : Atomic_Type)
with Convention => Intrinsic;

    Atomically performs: Item := Item + Value;

procedure Atomic_Subtract (Item  : aliased in out Atomic_Type;
                           Value : Atomic_Type)
with Convention => Intrinsic;

    Atomically performs: Item := Item - Value;

function Atomic_Fetch_And_Add
(Item  : aliased in out Atomic_Type;
 Value : Atomic_Type) return Atomic_Type
with Convention => Intrinsic;

    Atomically performs: Tmp := Item; Item := Item + Value; return Tmp;

function Atomic_Fetch_And_Subtract
(Item  : aliased in out Atomic_Type;
 Value : Atomic_Type) return Atomic_Type
with Convention => Intrinsic;

    Atomically performs: Tmp := Item; Item := Item - Value; return Tmp;
```

### C.6.5 The Package System.Atomic\_Operations.Modular\_Arithmetic

The language-defined generic package System.Atomic\_Operations.Modular\_Arithmetic provides operations to perform arithmetic atomically on objects of modular types.

#### Static Semantics

The generic library package System.Atomic\_Operations.Modular\_Arithmetic has the following declaration:

```
generic
  type Atomic_Type is mod <> with Atomic;
package System.Atomic_Operations.Modular_Arithmetic
with Pure, Nonblocking is

  procedure Atomic_Add (Item  : aliased in out Atomic_Type;
                       Value : Atomic_Type)
    with Convention => Intrinsic;

  procedure Atomic_Subtract (Item  : aliased in out Atomic_Type;
                            Value : Atomic_Type)
    with Convention => Intrinsic;

  function Atomic_Fetch_And_Add
    (Item  : aliased in out Atomic_Type;
     Value : Atomic_Type) return Atomic_Type
    with Convention => Intrinsic;

  function Atomic_Fetch_And_Subtract
    (Item  : aliased in out Atomic_Type;
     Value : Atomic_Type) return Atomic_Type
    with Convention => Intrinsic;

  function Is_Lock_Free (Item : aliased Atomic_Type) return Boolean
    with Convention => Intrinsic;

end System.Atomic_Operations.Modular_Arithmetic;
```

The operations of this package are defined as follows:



```

procedure Atomic_Add (Item : aliased in out Atomic_Type;
                      Value : Atomic_Type)
  with Convention => Intrinsic;

  Atomically performs: Item := Item + Value;

procedure Atomic_Subtract (Item : aliased in out Atomic_Type;
                           Value : Atomic_Type)
  with Convention => Intrinsic;

  Atomically performs: Item := Item - Value;

function Atomic_Fetch_And_Add
  (Item : aliased in out Atomic_Type;
   Value : Atomic_Type) return Atomic_Type
  with Convention => Intrinsic;

  Atomically performs: Tmp := Item; Item := Item + Value; return Tmp;

function Atomic_Fetch_And_Subtract
  (Item : aliased in out Atomic_Type;
   Value : Atomic_Type) return Atomic_Type
  with Convention => Intrinsic;

  Atomically performs: Tmp := Item; Item := Item - Value; return Tmp;

```

## C.7 Task Information

This subclause describes operations and attributes that can be used to obtain the identity of a task. In addition, a package that associates user-defined information with a task is defined. Finally, a package that associates termination procedures with a task or set of tasks is defined.

### C.7.1 The Package Task\_Identification

*Static Semantics*

The following language-defined library package exists:

```

package Ada.Task_Identification
  with Preelaborate, Nonblocking, Global => in out synchronized is
  type Task_Id is private
    with Preelaborable_Initialization;
  Null_Task_Id : constant Task_Id;
  function "=" (Left, Right : Task_Id) return Boolean;

  function Image (T : Task_Id) return String;
  function Current_Task return Task_Id;
  function Environment_Task return Task_Id;
  procedure Abort_Task (T : in Task_Id)
    with Nonblocking => False;

  function Is_Terminated (T : Task_Id) return Boolean;
  function Is_Callable (T : Task_Id) return Boolean;
  function Activation_Is_Complete (T : Task_Id) return Boolean;
private
  ... -- not specified by the language
end Ada.Task_Identification;

```

*Dynamic Semantics*

A value of the type Task\_Id identifies an existent task. The constant Null\_Task\_Id does not identify any task. Each object of the type Task\_Id is default initialized to the value of Null\_Task\_Id.

The function "=" returns True if and only if Left and Right identify the same task or both have the value Null\_Task\_Id.

The function `Image` returns an implementation-defined string that identifies `T`. If `T` equals `Null_Task_Id`, `Image` returns an empty string.

The function `Current_Task` returns a value that identifies the calling task.

The function `Environment_Task` returns a value that identifies the environment task.

The effect of `Abort_Task` is the same as the `abort_statement` for the task identified by `T`. In addition, if `T` identifies the environment task, the entire partition is aborted, see E.1.

The functions `Is_Terminated` and `Is_Callable` return the value of the corresponding attribute of the task identified by `T`.

The function `Activation_Is_Complete` returns `True` if the task identified by `T` has completed its activation (whether successfully or not). It returns `False` otherwise. If `T` identifies the environment task, `Activation_Is_Complete` returns `True` after the elaboration of the `library_items` of the partition has completed.

For a prefix `T` that is of a task type (after any implicit dereference), the following attribute is defined:

`T'Identity` Yields a value of the type `Task_Id` that identifies the task denoted by `T`.

For a prefix `E` that denotes an `entry_declaration`, the following attribute is defined:

`E'Caller` Yields a value of the type `Task_Id` that identifies the task whose call is now being serviced. Use of this attribute is allowed only inside an `accept_statement`, or `entry_body` after the `entry_barrier`, corresponding to the `entry_declaration` denoted by `E`.

`Program_Error` is raised if a value of `Null_Task_Id` is passed as a parameter to `Abort_Task`, `Activation_Is_Complete`, `Is_Terminated`, and `Is_Callable`.

#### *Bounded (Run-Time) Errors*

It is a bounded error to call the `Current_Task` function from an `entry_body`, interrupt handler, or finalization of a task attribute. `Program_Error` is raised, or an implementation-defined value of the type `Task_Id` is returned.

#### *Erroneous Execution*

If a value of `Task_Id` is passed as a parameter to any of the operations declared in this package (or any language-defined child of this package), and the corresponding task object no longer exists, the execution of the program is erroneous.

#### *Documentation Requirements*

The implementation shall document the effect of calling `Current_Task` from an entry body or interrupt handler.

NOTE 1 This package is intended for use in writing user-defined task scheduling packages and constructing server tasks. `Current_Task` can be used in conjunction with other operations requiring a task as an argument such as `Set_Priority` (see D.5).

NOTE 2 The function `Current_Task` and the attribute `Caller` can return a `Task_Id` value that identifies the environment task.

## C.7.2 The Package Task\_Attributes

### Static Semantics

The following language-defined generic library package exists:

```
with Ada.Task_Identification; use Ada.Task_Identification;
generic
  type Attribute is private;
  Initial_Value : in Attribute;
package Ada.Task_Attributes
  with Nonblocking, Global => in out synchronized is
    type Attribute_Handle is access all Attribute;
    function Value(T : Task_Id := Current_Task)
      return Attribute;
    function Reference(T : Task_Id := Current_Task)
      return Attribute_Handle;
    procedure Set_Value(Val : in Attribute;
      T : in Task_Id := Current_Task);
    procedure Reinitialize(T : in Task_Id := Current_Task);
end Ada.Task_Attributes;
```

### Dynamic Semantics

When an instance of Task\_Attributes is elaborated in a given active partition, an object of the actual type corresponding to the formal type Attribute is implicitly created for each task (of that partition) that exists and is not yet terminated. This object acts as a user-defined attribute of the task. A task created previously in the partition and not yet terminated has this attribute from that point on. Each task subsequently created in the partition will have this attribute when created. In all these cases, the initial value of the given attribute is Initial\_Value.

The Value operation returns the value of the corresponding attribute of T.

The Reference operation returns an access value that designates the corresponding attribute of T.

The Set\_Value operation performs any finalization on the old value of the attribute of T and assigns Val to that attribute (see 8.2 and 10.6).

The effect of the Reinitialize operation is the same as Set\_Value where the Val parameter is replaced with Initial\_Value.

For all the operations declared in this package, Tasking\_Error is raised if the task identified by T is terminated. Program\_Error is raised if the value of T is Null\_Task\_Id.

After a task has terminated, all of its attributes are finalized, unless they have been finalized earlier. When the master of an instantiation of Ada.Task\_Attributes is finalized, the corresponding attribute of each task is finalized, unless it has been finalized earlier.

### Bounded (Run-Time) Errors

If the package Ada.Task\_Attributes is instantiated with a controlled type and the controlled type has user-defined Adjust or Finalize operations that in turn access task attributes by any of the above operations, then a call of Set\_Value of the instantiated package constitutes a bounded error. The call may perform as expected or may result in forever blocking the calling task and subsequently some or all tasks of the partition.

### Erroneous Execution

It is erroneous to dereference the access value returned by a given call on Reference after a subsequent call on Reinitialize for the same task attribute, or after the associated task terminates.

If a value of Task\_Id is passed as a parameter to any of the operations declared in this package and the corresponding task object no longer exists, the execution of the program is erroneous.

An access to a task attribute via a value of type `Attribute_Handle` is erroneous if executed concurrently with another such access or a call of any of the operations declared in package `Task_Attributes`. An access to a task attribute is erroneous if executed concurrently with or after the finalization of the task attribute.

#### *Implementation Requirements*

For a given attribute of a given task, the implementation shall perform the operations declared in this package atomically with respect to any of these operations of the same attribute of the same task. The granularity of any locking mechanism necessary to achieve such atomicity is implementation defined.

After task attributes are finalized, the implementation shall reclaim any storage associated with the attributes.

#### *Documentation Requirements*

The implementation shall document the limit on the number of attributes per task, if any, and the limit on the total storage for attribute values per task, if such a limit exists.

In addition, if these limits can be configured, the implementation shall document how to configure them.

#### *Metrics*

The implementation shall document the following metrics: A task calling the following subprograms shall execute at a sufficiently high priority as to not be preempted during the measurement period. This period shall start just before issuing the call and end just after the call completes. If the attributes of task `T` are accessed by the measurement tests, no other task shall access attributes of that task during the measurement period. For all measurements described here, the `Attribute` type shall be a scalar type whose size is equal to the size of the predefined type `Integer`. For each measurement, two cases shall be documented: one where the accessed attributes are of the calling task (that is, the default value for the `T` parameter is used), and the other, where `T` identifies another, nonterminated, task.

The following calls (to subprograms in the `Task_Attributes` package) shall be measured:

- a call to `Value`, where the return value is `Initial_Value`;
- a call to `Value`, where the return value is not equal to `Initial_Value`;
- a call to `Reference`, where the return value designates a value equal to `Initial_Value`;
- a call to `Reference`, where the return value designates a value not equal to `Initial_Value`;
- a call to `Set_Value` where the `Val` parameter is not equal to `Initial_Value` and the old attribute value is equal to `Initial_Value`;
- a call to `Set_Value` where the `Val` parameter is not equal to `Initial_Value` and the old attribute value is not equal to `Initial_Value`.

#### *Implementation Permissions*

An implementation can avoid actually creating the object corresponding to a task attribute until its value is set to something other than that of `Initial_Value`, or until `Reference` is called for the task attribute. Similarly, when the value of the attribute is to be reinitialized to that of `Initial_Value`, the object may instead be finalized and its storage reclaimed, to be recreated when needed later. While the object does not exist, the function `Value` may simply return `Initial_Value`, rather than implicitly creating the object.

An implementation is allowed to place restrictions on the maximum number of attributes a task may have, the maximum size of each attribute, and the total storage size allocated for all the attributes of a task.

*Implementation Advice*

Some implementations are targeted to domains in which memory use at run time has to be completely deterministic. For such implementations, it is recommended that the storage for task attributes will be pre-allocated statically and not from the heap. This can be accomplished by either placing restrictions on the number and the size of the attributes of a task, or by using the pre-allocated storage for the first N attribute objects, and the heap for the others. In the latter case, N should be documented.

Finalization of task attributes and reclamation of associated storage should be performed as soon as possible after task termination.

NOTE 1 An attribute always exists (after instantiation), and has the initial value. An implementation can avoid using memory to store the attribute value until the first operation that changes the attribute value. The same holds true after Reinitialize.

NOTE 2 The result of the Reference function is always safe to use in the task body whose attribute is being accessed. However, when the result is being used by another task, the programmer will want to make sure that the task whose attribute is being accessed is not yet terminated. Failing to do so can make the program execution erroneous.

### C.7.3 The Package Task\_Termination

*Static Semantics*

The following language-defined library package exists:

```
with Ada.Task_Identification;
with Ada.Exceptions;
package Ada.Task_Termination
  with Preelaborate, Nonblocking, Global => in out synchronized is
  type Cause_Of_Termination is (Normal, Abnormal, Unhandled_Exception);
  type Termination_Handler is access protected procedure
    (Cause : in Cause_Of_Termination;
     T      : in Ada.Task_Identification.Task_Id;
     X      : in Ada.Exceptions.Exception_Occurrence);
  procedure Set_Dependents_Fallback_Handler
    (Handler: in Termination_Handler);
  function Current_Task_Fallback_Handler return Termination_Handler;
  procedure Set_Specific_Handler
    (T      : in Ada.Task_Identification.Task_Id;
     Handler : in Termination_Handler);
  function Specific_Handler (T : Ada.Task_Identification.Task_Id)
    return Termination_Handler;
end Ada.Task_Termination;
```

*Dynamic Semantics*

The type `Termination_Handler` identifies a protected procedure to be executed by the implementation when a task terminates. Such a protected procedure is called a *handler*. In all cases T identifies the task that is terminating. If the task terminates due to completing the last statement of its body, or as a result of waiting on a terminate alternative, and the finalization of the task completes normally, then Cause is set to Normal and X is set to Null\_Occurrence. If the task terminates because it is being aborted, then Cause is set to Abnormal; X is set to Null\_Occurrence if the finalization of the task completes normally. If the task terminates because of an exception raised by the execution of its `task_body`, then Cause is set to Unhandled\_Exception; X is set to the associated exception occurrence if the finalization of the task completes normally. Independent of how the task completes, if finalization of the task propagates an exception, then Cause is either Unhandled\_Exception or Abnormal, and X is an exception occurrence that identifies the Program\_Error exception.

Each task has two termination handlers, a *fall-back handler* and a *specific handler*. The specific handler applies only to the task itself, while the fall-back handler applies only to the dependent tasks of the task.

A handler is said to be *set* if it is associated with a nonnull value of type `Termination_Handler`, and *cleared* otherwise. When a task is created, its specific handler and fall-back handler are cleared.

The procedure `Set_Dependents_Fallback_Handler` changes the fall-back handler for the calling task: if Handler is **null**, that fall-back handler is cleared; otherwise, it is set to be `Handler.all`. If a fall-back handler had previously been set it is replaced.

The function `Current_Task_Fallback_Handler` returns the fall-back handler that is currently set for the calling task, if one is set; otherwise, it returns **null**.

The procedure `Set_Specific_Handler` changes the specific handler for the task identified by T: if Handler is **null**, that specific handler is cleared; otherwise, it is set to be `Handler.all`. If a specific handler had previously been set it is replaced.

The function `Specific_Handler` returns the specific handler that is currently set for the task identified by T, if one is set; otherwise, it returns **null**.

As part of the finalization of a `task_body`, after performing the actions specified in 10.6 for finalization of a master, the specific handler for the task, if one is set, is executed. If the specific handler is cleared, a search for a fall-back handler proceeds by recursively following the master relationship for the task. If a task is found whose fall-back handler is set, that handler is executed; otherwise, no handler is executed.

For `Set_Specific_Handler` or `Specific_Handler`, `Tasking_Error` is raised if the task identified by T has already terminated. `Program_Error` is raised if the value of T is `Ada.Task_Identification.Null_Task_Id`.

An exception propagated from a handler that is invoked as part of the termination of a task has no effect.

#### *Erroneous Execution*

For a call of `Set_Specific_Handler` or `Specific_Handler`, if the task identified by T no longer exists, the execution of the program is erroneous.



## Annex D

### (informative)

## Real-Time Systems

This Annex specifies additional characteristics of Ada implementations intended for real-time systems software. To conform to this Annex, an implementation shall also conform to Annex C.

#### *Metrics*

The metrics are documentation requirements; an implementation shall document the values of the language-defined metrics for at least one configuration of hardware or an underlying system supported by the implementation, and shall document the details of that configuration.

The metrics do not necessarily yield a simple number. For some, a range is more suitable, for others a formula dependent on some parameter is appropriate, and for others, it may be more suitable to break the metric into several cases. Unless specified otherwise, the metrics in this annex are expressed in processor clock cycles. For metrics that require documentation of an upper bound, if there is no upper bound, the implementation shall report that the metric is unbounded.

NOTE 1 The specification of the metrics makes a distinction between upper bounds and simple execution times. Where something is just specified as “the execution time of” a piece of code, this leaves one the freedom to choose a nonpathological case. This kind of metric is of the form “there exists a program such that the value of the metric is V”. Conversely, the meaning of upper bounds is “there is no program such that the value of the metric is greater than V”. This kind of metric can only be partially tested, by finding the value of V for one or more test programs.

NOTE 2 The metrics do not cover the whole language; they are limited to features that are specified in Annex C and in this Annex. The metrics are intended to provide guidance to potential users as to whether a particular implementation of such a feature is going to be adequate for a particular real-time application. As such, the metrics are aimed at known implementation choices that can result in significant performance differences.

NOTE 3 The purpose of the metrics is not necessarily to provide fine-grained quantitative results or to serve as a comparison between different implementations on the same or different platforms. Instead, their goal is rather qualitative; to define a standard set of approximate values that can be measured and used to estimate the general suitability of an implementation, or to evaluate the comparative utility of certain features of an implementation for a particular real-time application.

### D.1 Task Priorities

This subclause specifies the priority model for real-time systems. In addition, the methods for specifying priorities are defined.

#### *Static Semantics*

For a task type (including the anonymous type of a `single_task_declaration`), protected type (including the anonymous type of a `single_protected_declaration`), or subprogram, the following language-defined representation aspects may be specified:

**Priority** The aspect `Priority` is an `expression`, which shall be of type `Integer`.

**Interrupt\_Priority**

The aspect `Interrupt_Priority` is an `expression`, which shall be of type `Integer`.

#### *Legality Rules*

If the `Priority` aspect is specified for a subprogram, the `expression` shall be static, and its value shall be in the range of `System.Priority`.

At most one of the `Priority` and `Interrupt_Priority` aspects may be specified for a given entity.

Neither of the Priority or Interrupt\_Priority aspects shall be specified for a synchronized interface type.

#### Static Semantics

The following declarations exist in package System:

```

subtype Any_Priority is Integer range implementation-defined;
subtype Priority is Any_Priority
    range Any_Priority'First .. implementation-defined;
subtype Interrupt_Priority is Any_Priority
    range Priority'Last+1 .. Any_Priority'Last;

Default_Priority : constant Priority := (Priority'First + Priority'Last)/2;

```

The full range of priority values supported by an implementation is specified by the subtype Any\_Priority. The subrange of priority values that are high enough to require the blocking of one or more interrupts is specified by the subtype Interrupt\_Priority. The subrange of priority values below System.Interrupt\_Priority'First is specified by the subtype System.Priority.

#### Dynamic Semantics

The Priority aspect has no effect if it is specified for a subprogram other than the main subprogram; the Priority value is not associated with any task.

A *task priority* is an integer value that indicates a degree of urgency and is the basis for resolving competing demands of tasks for resources. Unless otherwise specified, whenever tasks compete for processors or other implementation-defined resources, the resources are allocated to the task with the highest priority value. The *base priority* of a task is the priority with which it was created, or to which it was later set by Dynamic\_Priorities.Set\_Priority (see D.5). At all times, a task also has an *active priority*, which generally is its base priority unless it inherits a priority from other sources. *Priority inheritance* is the process by which the priority of a task or other entity (for example, a protected object; see D.3) is used in the evaluation of another task's active priority.

The effect of specifying a Priority or Interrupt\_Priority aspect for a protected type or single\_protected\_declaration is discussed in D.3.

The expression specified for the Priority or Interrupt\_Priority aspect of a task type is evaluated each time an object of the task type is created (see 12.1). For the Priority aspect, the value of the expression is converted to the subtype Priority; for the Interrupt\_Priority aspect, this value is converted to the subtype Any\_Priority. The priority value is then associated with the task object.

Likewise, the priority value is associated with the environment task if the aspect is specified for the main subprogram.

The initial value of a task's base priority is specified by default or by means of a Priority or Interrupt\_Priority aspect. After a task is created, its base priority can be changed only by a call to Dynamic\_Priorities.Set\_Priority (see D.5). The initial base priority of a task in the absence of an aspect is the base priority of the task that creates it at the time of creation (see 12.1). If the aspect Priority is not specified for the main subprogram, the initial base priority of the environment task is System.Default\_Priority. The task's active priority is used when the task competes for processors. Similarly, the task's active priority is used to determine the task's position in any queue when Priority\_Queueing is specified (see D.4).

At any time, the active priority of a task is the maximum of all the priorities the task is inheriting at that instant. For a task that is not held (see D.11), its base priority is a source of priority inheritance unless otherwise specified for a particular task dispatching policy. Other sources of priority inheritance are specified under the following conditions:

- During activation, a task being activated inherits the active priority that its activator (see 12.2) had at the time the activation was initiated.
- During rendezvous, the task accepting the entry call inherits the priority of the entry call (see 12.5.3 and D.4).

- While starting a protected action on a protected object when the FIFO\_Spinning admission policy is in effect, a task inherits the ceiling priority of the protected object (see 12.5, D.3, and D.4.1).
- While a task executes a protected action on a protected object, the task inherits the ceiling priority of the protected object (see 12.5 and D.3).

In all of these cases, the priority ceases to be inherited as soon as the condition calling for the inheritance no longer exists.

#### *Implementation Requirements*

The range of System.Interrupt\_Priority shall include at least one value.

The range of System.Priority shall include at least 30 values.

NOTE 1 The priority expression can include references to discriminants of the enclosing type.

NOTE 2 It is a consequence of the active priority rules that at the point when a task stops inheriting a priority from another source, its active priority is re-evaluated. This is in addition to other instances described in this Annex for such re-evaluation.

NOTE 3 An implementation can provide a nonstandard mode in which tasks inherit priorities under conditions other than those specified above.

## **D.2 Priority Scheduling**

This subclause describes the rules that determine which task is selected for execution when more than one task is ready (see Clause 12).

### **D.2.1 The Task Dispatching Model**

The task dispatching model specifies task scheduling, based on conceptual priority-ordered ready queues.

#### *Static Semantics*

The following language-defined library package exists:

```
package Ada.Dispatching
with Preelaborate, Nonblocking, Global => in out synchronized is

  procedure Yield
    with Nonblocking => False;

  Dispatching_Policy_Error : exception;
end Ada.Dispatching;
```

Dispatching serves as the parent of other language-defined library units concerned with task dispatching.

For a noninstance subprogram (including a generic formal subprogram), a generic subprogram, or an entry, the following language-defined aspect may be specified with an `aspect_specification` (see 16.1.1):

**Yield** The type of aspect Yield is Boolean.

If directly specified, the `aspect_definition` shall be a static expression. If not specified (including by inheritance), the aspect is False.

If a Yield aspect is specified True for a primitive subprogram *S* of a type *T*, then the aspect is inherited by the corresponding primitive subprogram of each descendant of *T*.

#### *Legality Rules*

If the Yield aspect is specified for a dispatching subprogram that inherits the aspect, the specified value shall be confirming.

If the Nonblocking aspect (see 12.5) of the associated callable entity is statically True, the Yield aspect shall not be specified as True. For a callable entity that is declared within a generic body, this rule is checked assuming that any nonstatic Nonblocking attributes in the expression of the Nonblocking aspect of the entity are statically True.

In addition to the places where Legality Rules normally apply (see 15.3), these rules also apply in the private part of an instance of a generic unit.

#### Dynamic Semantics

A task can become a *running task* only if it is ready (see Clause 12) and the execution resources required by that task are available. Processors are allocated to tasks based on each task's active priority.

It is implementation defined whether, on a multiprocessor, a task that is waiting for access to a protected object keeps its processor busy.

*Task dispatching* is the process by which a logical thread of control associated with a ready task is selected for execution on a processor. This selection is done during the execution of such a logical thread of control, at certain points called *task dispatching points*. Such a logical thread of control reaches a task dispatching point whenever it becomes blocked, and when its associated task terminates. Other task dispatching points are defined throughout this Annex for specific policies. Below we talk in terms of tasks, but in the context of a parallel construct, a single task can be represented by multiple logical threads of control, each of which can appear separately on a ready queue.

*Task dispatching policies* are specified in terms of conceptual *ready queues* and task states. A ready queue is an ordered list of ready tasks. The first position in a queue is called the *head of the queue*, and the last position is called the *tail of the queue*. A task is *ready* if it is in a ready queue, or if it is running. Each processor has one ready queue for each priority value. At any instant, each ready queue of a processor contains exactly the set of tasks of that priority that are ready for execution on that processor, but are not running on any processor; that is, those tasks that are ready, are not running on any processor, and can be executed using that processor and other available resources. A task can be on the ready queues of more than one processor.

Each processor also has one *running task*, which is the task currently being executed by that processor. Whenever a task running on a processor reaches a task dispatching point it goes back to one or more ready queues; a task (possibly the same task) is then selected to run on that processor. The task selected is the one at the head of the highest priority nonempty ready queue; this task is then removed from all ready queues to which it belongs.

A call of Yield and a `delay_statement` are task dispatching points for all language-defined policies.

If the Yield aspect has the value True, then a call to procedure Yield is included within the body of the associated callable entity, and invoked immediately prior to returning from the body if and only if no other task dispatching points were encountered during the execution of the body.

#### Implementation Permissions

An implementation is allowed to define additional resources as execution resources, and to define the corresponding allocation policies for them. Such resources may have an implementation-defined effect on task dispatching.

An implementation may place implementation-defined restrictions on tasks whose active priority is in the Interrupt\_Priority range.

Unless otherwise specified for a task dispatching policy, an implementation may add additional points at which task dispatching may occur, in an implementation-defined manner.

NOTE 1 Clause 12 specifies under which circumstances a task becomes ready. The ready state is affected by the rules for task activation and termination, delay statements, and entry calls. When a task is not ready, it is said to be blocked.

NOTE 2 An example of a possible implementation-defined execution resource is a page of physical memory, which must be loaded with a particular page of virtual memory before a task can continue execution.

NOTE 3 The ready queues are purely conceptual; there is no requirement that such lists physically exist in an implementation.

NOTE 4 While a task is running, it is not on any ready queue. Any time the task that is running on a processor is added to a ready queue, a new running task is selected for that processor.

NOTE 5 In a multiprocessor system, a task can be on the ready queues of more than one processor. At the extreme, if several processors share the same set of ready tasks, the contents of their ready queues is identical, and so they can be viewed as sharing one ready queue, and can be implemented that way. Thus, the dispatching model covers multiprocessors where dispatching is implemented using a single ready queue, as well as those with separate dispatching domains.

NOTE 6 The priority of a task is determined by rules specified in this subclause, and under D.1, D.3 and D.5.

NOTE 7 The setting of a task's base priority as a result of a call to `Set_Priority` does not always take effect immediately when `Set_Priority` is called. The effect of setting the task's base priority is deferred while the affected task performs a protected action.

## D.2.2 Task Dispatching Pragmas

This subclause allows a single task dispatching policy to be defined for all priorities, or the range of priorities to be split into subranges that are assigned individual dispatching policies.

### Syntax

The form of a `pragma Task_Dispatching_Policy` is as follows:

```
pragma Task_Dispatching_Policy(policy_identifier);
```

The form of a `pragma Priority_Specific_Dispatching` is as follows:

```
pragma Priority_Specific_Dispatching (  
  policy_identifier, first_priority_expression, last_priority_expression);
```

### Name Resolution Rules

The expected type for *first\_priority\_expression* and *last\_priority\_expression* is Integer.

### Legality Rules

The *policy\_identifier* used in a `pragma Task_Dispatching_Policy` shall be the name of a task dispatching policy.

The *policy\_identifier* used in a `pragma Priority_Specific_Dispatching` shall be the name of a task dispatching policy.

Both *first\_priority\_expression* and *last\_priority\_expression* shall be static expressions in the range of `System.Any_Priority`; *last\_priority\_expression* shall have a value greater than or equal to *first\_priority\_expression*.

### Static Semantics

`Pragma Task_Dispatching_Policy` specifies the single task dispatching policy.

`Pragma Priority_Specific_Dispatching` specifies the task dispatching policy for the specified range of priorities. Tasks with base priorities within the range of priorities specified in a `Priority_Specific_Dispatching` pragma have their active priorities determined according to the specified dispatching policy. Tasks with active priorities within the range of priorities specified in a `Priority_Specific_Dispatching` pragma are dispatched according to the specified dispatching policy.

If a partition contains one or more `Priority_Specific_Dispatching` pragmas, the dispatching policy for priorities not covered by any `Priority_Specific_Dispatching` pragmas is `FIFO_Within_Priorities`.

*Post-Compilation Rules*

A `Task_Dispatching_Policy` pragma is a configuration pragma. A `Priority_Specific_Dispatching` pragma is a configuration pragma.

The priority ranges specified in more than one `Priority_Specific_Dispatching` pragma within the same partition shall not be overlapping.

If a partition contains one or more `Priority_Specific_Dispatching` pragmas it shall not contain a `Task_Dispatching_Policy` pragma.

*Dynamic Semantics*

A *task dispatching policy* specifies the details of task dispatching that are not covered by the basic task dispatching model. These rules govern when tasks are inserted into and deleted from the ready queues. A single task dispatching policy is specified by a `Task_Dispatching_Policy` pragma. Pragma `Priority_Specific_Dispatching` assigns distinct dispatching policies to subranges of `System.Any_Priority`.

If neither pragma applies to any of the program units comprising a partition, the task dispatching policy for that partition is unspecified.

If a partition contains one or more `Priority_Specific_Dispatching` pragmas, a task dispatching point occurs for the currently running task of a processor whenever there is a nonempty ready queue for that processor with a higher priority than the priority of the running task.

A task that has its base priority changed may move from one dispatching policy to another. It is immediately subject to the new dispatching policy.

*Implementation Requirements*

An implementation shall allow, for a single partition, both the locking policy (see D.3) to be specified as `Ceiling_Locking` and also one or more `Priority_Specific_Dispatching` pragmas to be given.

*Implementation Permissions*

Implementations are allowed to define other task dispatching policies, but are not required to support specifying more than one task dispatching policy per partition.

An implementation is not required to support pragma `Priority_Specific_Dispatching` if it is infeasible to support it in the target environment.

**D.2.3 Preemptive Dispatching**

This subclause defines a preemptive task dispatching policy.

*Static Semantics*

The *policy identifier* `FIFO_Within_Priorities` is a task dispatching policy.

*Dynamic Semantics*

When `FIFO_Within_Priorities` is in effect, modifications to the ready queues occur only as follows:

- When a blocked task becomes ready, it is added at the tail of the ready queue for its active priority.
- When the active priority of a ready task that is not running changes, or the setting of its base priority takes effect, the task is removed from the ready queue for its old active priority and is added at the tail of the ready queue for its new active priority, except in the case where the active priority is lowered due to the loss of inherited priority, in which case the task is added at the head of the ready queue for its new active priority.
- When the setting of the base priority of a running task takes effect, the task is added to the tail of the ready queue for its active priority.
- When a task executes a `delay_statement` that does not result in blocking, it is added to the tail of the ready queue for its active priority.



Each of the events specified above is a task dispatching point (see D.2.1).

A task dispatching point occurs for the currently running task of a processor whenever there is a nonempty ready queue for that processor with a higher priority than the priority of the running task. The currently running task is said to be *preempted* and it is added at the head of the ready queue for its active priority.

#### Implementation Requirements

An implementation shall allow, for a single partition, both the task dispatching policy to be specified as FIFO\_Within\_Priorities and also the locking policy (see D.3) to be specified as Ceiling\_Locking.

#### Documentation Requirements

*Priority inversion* is the duration for which a task remains at the head of the highest priority nonempty ready queue while the processor executes a lower priority task. The implementation shall document:

- The maximum priority inversion a user task can experience due to activity of the implementation (on behalf of lower priority tasks), and
- whether execution of a task can be preempted by the implementation processing of delay expirations for lower priority tasks, and if so, for how long.

NOTE 1 If the active priority of a running task is lowered due to loss of inherited priority (as it is on completion of a protected operation) and there is a ready task of the same active priority that is not running, the running task continues to run (provided that there is no higher priority task).

NOTE 2 Setting the base priority of a ready task causes the task to move to the tail of the queue for its active priority, regardless of whether the active priority of the task actually changes.

### D.2.4 Non-Preemptive Dispatching

This subclause defines a non-preemptive task dispatching policy.

#### Static Semantics

The *policy\_identifier* Non\_Preemptive\_FIFO\_Within\_Priorities is a task dispatching policy.

The following language-defined library package exists:

```
package Ada.Dispatching.Non_Preemptive
with Preelaborate, Nonblocking, Global => in out synchronized is
procedure Yield_To_Higher;
procedure Yield_To_Same_Or_Higher renames Yield;
end Ada.Dispatching.Non_Preemptive;
```

A call of Yield\_To\_Higher is a task dispatching point for this policy. If the task at the head of the highest priority ready queue has a higher active priority than the calling task, then the calling task is preempted.

#### Legality Rules

Non\_Preemptive\_FIFO\_Within\_Priorities shall not be specified as the *policy\_identifier* of pragma Priority\_Specific\_Dispatching (see D.2.2).

#### Dynamic Semantics

When Non\_Preemptive\_FIFO\_Within\_Priorities is in effect, modifications to the ready queues occur only as follows:

- When a blocked task becomes ready, it is added at the tail of the ready queue for its active priority.
- When the active priority of a ready task that is not running changes, or the setting of its base priority takes effect, the task is removed from the ready queue for its old active priority and is added at the tail of the ready queue for its new active priority.

- When the setting of the base priority of a running task takes effect, the task is added to the tail of the ready queue for its active priority.
- When a task executes a `delay_statement` that does not result in blocking, it is added to the tail of the ready queue for its active priority.

For this policy, blocking or termination of a task, a `delay_statement`, a call to `Yield_To_Higher`, and a call to `Yield_To_Same_Or_Higher` or `Yield` are the only task dispatching points (see D.2.1).

#### Implementation Requirements

An implementation shall allow, for a single partition, both the task dispatching policy to be specified as `Non_Preemptive_FIFO_Within_Priorities` and also the locking policy (see D.3) to be specified as `Ceiling_Locking`.

#### Implementation Permissions

Since implementations are allowed to round all ceiling priorities in subrange `System.Priority` to `System.Priority'Last` (see D.3), an implementation may allow a task of a partition using the `Non_Preemptive_FIFO_Within_Priorities` policy to execute within a protected object without raising its active priority provided the associated protected unit does not contain any subprograms with aspects `Interrupt_Handler` or `Attach_Handler` specified, nor does the unit have aspect `Interrupt_Priority` specified. When the locking policy (see D.3) is `Ceiling_Locking`, an implementation taking advantage of this permission shall ensure that a call to `Yield_to_Higher` that occurs within a protected action uses the ceiling priority of the protected object (rather than the active priority of the task) when determining whether to preempt the task.

### D.2.5 Round Robin Dispatching

This subclause defines the task dispatching policy `Round_Robin_Within_Priorities` and the package `Round_Robin`.

#### Static Semantics

The *policy\_identifier* `Round_Robin_Within_Priorities` is a task dispatching policy.

The following language-defined library package exists:

```
with System;
with Ada.Real_Time;
package Ada.Dispatching.Round_Robin
  with Nonblocking, Global => in out synchronized is
    Default_Quantum : constant Ada.Real_Time.Time_Span :=
      implementation-defined;
    procedure Set_Quantum (Pri : in System.Priority;
                          Quantum : in Ada.Real_Time.Time_Span);
    procedure Set_Quantum (Low, High : in System.Priority;
                          Quantum : in Ada.Real_Time.Time_Span);
    function Actual_Quantum (Pri : System.Priority)
      return Ada.Real_Time.Time_Span;
    function Is_Round_Robin (Pri : System.Priority) return Boolean;
end Ada.Dispatching.Round_Robin;
```

When task dispatching policy `Round_Robin_Within_Priorities` is the single policy in effect for a partition, each task with priority in the range of `System.Interrupt_Priority` is dispatched according to policy `FIFO_Within_Priorities`.

#### Dynamic Semantics

The procedures `Set_Quantum` set the required `Quantum` value for a single priority level `Pri` or a range of priority levels `Low .. High`. If no quantum is set for a Round Robin priority level, `Default_Quantum` is used.

The function `Actual_Quantum` returns the actual quantum used by the implementation for the priority level `Pri`.

The function `Is_Round_Robin` returns `True` if priority `Pri` is covered by task dispatching policy `Round_Robin_Within_Priorities`; otherwise, it returns `False`.

A call of `Actual_Quantum` or `Set_Quantum` raises exception `Dispatching.Dispatching_Policy_Error` if a predefined policy other than `Round_Robin_Within_Priorities` applies to the specified priority or any of the priorities in the specified range.

For `Round_Robin_Within_Priorities`, the dispatching rules for `FIFO_Within_Priorities` apply with the following additional rules:

- When a task is added or moved to the tail of the ready queue for its base priority, it has an execution time budget equal to the quantum for that priority level. This will also occur when a blocked task becomes executable again.
- When a task is preempted (by a higher priority task) and is added to the head of the ready queue for its priority level, it retains its remaining budget.
- While a task is executing, its budget is decreased by the amount of execution time it uses. The accuracy of this accounting is the same as that for execution time clocks (see D.14).
- When a task has exhausted its budget and is without an inherited priority (and is not executing within a protected operation), it is moved to the tail of the ready queue for its priority level. This is a task dispatching point.

#### *Implementation Requirements*

An implementation shall allow, for a single partition, both the task dispatching policy to be specified as `Round_Robin_Within_Priorities` and also the locking policy (see D.3) to be specified as `Ceiling_Locking`.

#### *Documentation Requirements*

An implementation shall document the quantum values supported.

An implementation shall document the accuracy with which it detects the exhaustion of the budget of a task.

NOTE 1 Due to implementation constraints, the quantum value returned by `Actual_Quantum` can differ from that set with `Set_Quantum`.

NOTE 2 A task that executes continuously with an inherited priority will not be subject to round robin dispatching.

### **D.2.6 Earliest Deadline First Dispatching**

The deadline of a task is an indication of the urgency of the task; it represents a point on an ideal physical time line. The deadline can affect how resources are allocated to the task.

This subclause presents `Dispatching.EDF`, a package for representing the deadline of a task and a dispatching policy that defines Earliest Deadline First (EDF) dispatching. The `Relative_Deadline` aspect is provided to assign an initial deadline to a task. A configuration pragma `Generate_Deadlines` is provided to specify that a task's deadline is recomputed whenever it is made ready.

#### *Static Semantics*

The `policy_identifier` `EDF_Within_Priorities` is a task dispatching policy.

The following language-defined library package exists:

```

with Ada.Real_Time;
with Ada.Task_Identification;
package Ada.Dispatching.EDF
with Nonblocking, Global => in out synchronized is
  subtype Deadline is Ada.Real_Time.Time;
  subtype Relative_Deadline is Ada.Real_Time.Time_Span;
  Default_Deadline : constant Deadline :=
    Ada.Real_Time.Time_Last;
  Default_Relative_Deadline : constant Relative_Deadline :=
    Ada.Real_Time.Time_Span_Last;
  procedure Set_Deadline
    (D : in Deadline;
     T : in Ada.Task_Identification.Task_Id :=
       Ada.Task_Identification.Current_Task);
  function Get_Deadline
    (T : Ada.Task_Identification.Task_Id :=
       Ada.Task_Identification.Current_Task) return Deadline;
  procedure Set_Relative_Deadline
    (D : in Relative_Deadline;
     T : in Ada.Task_Identification.Task_Id :=
       Ada.Task_Identification.Current_Task);
  function Get_Relative_Deadline
    (T : Ada.Task_Identification.Task_Id :=
       Ada.Task_Identification.Current_Task)
    return Relative_Deadline;
  procedure Delay_Until_And_Set_Deadline
    (Delay_Until_Time : in Ada.Real_Time.Time;
     Deadline_Offset : in Ada.Real_Time.Time_Span)
    with Nonblocking => False;
  function Get_Last_Release_Time
    (T : Ada.Task_Identification.Task_Id :=
       Ada.Task_Identification.Current_Task)
    return Ada.Real_Time.Time;
end Ada.Dispatching.EDF;

```

For a subprogram, a task type (including the anonymous type of a single\_task\_declaration), or a protected type (including the anonymous type of a single\_protected\_declaration), the following language-defined representation aspect may be specified:

Relative\_Deadline

The aspect Relative\_Deadline is an expression, which shall be of type Real\_Time.Time\_Span.

The form of pragma Generate\_Deadlines is as follows:

**pragma** Generate\_Deadlines;

The Generate\_Deadlines pragma is a configuration pragma.

#### *Legality Rules*

The Relative\_Deadline aspect shall not be specified on a task or protected interface type. If the Relative\_Deadline aspect is specified for a subprogram, the aspect\_definition shall be a static expression.

#### *Post-Compilation Rules*

If the EDF\_Within\_Priorities policy is specified for a partition, then the Ceiling\_Locking policy (see D.3) shall also be specified for the partition.

If the EDF\_Within\_Priorities policy appears in a Priority\_Specific\_Dispatching pragma (see D.2.2) in a partition, then the Ceiling\_Locking policy (see D.3) shall also be specified for the partition.

The `Relative_Deadline` aspect has no effect if it is specified for a subprogram other than the main subprogram.

If pragma `Generate_Deadlines` is in effect, the deadline of a task is recomputed each time it becomes ready. The new deadline is the value of `Real_Time.Clock` at the time the task is added to a ready queue plus the value returned by `Get_Relative_Deadline`.

The initial absolute deadline for a task with a specified `Relative_Deadline` is the result of adding the value returned by a call of `Real_Time.Clock` to the value of the `expression` specified as the `Relative_Deadline` aspect, where this entire computation, including the call of `Real_Time.Clock`, is performed between task creation and the start of its activation. If the aspect `Relative_Deadline` is not specified, then the initial absolute deadline of a task is the value of `Default_Deadline` (`Ada.Real_Time.Time_Last`). The environment task is also given an initial deadline by this rule, using the value of the `Relative_Deadline` aspect of the main subprogram (if any).

The effect of specifying a `Relative_Deadline` aspect for a protected type or `single_protected_declaration` is discussed in D.3.

A task has both an *active* and a *base* absolute deadline. These are the same except when the task is inheriting a relative deadline during activation or a rendezvous (see below) or within a protected action (see D.3). The procedure `Set_Deadline` changes the (base) absolute deadline of the task to `D`. The function `Get_Deadline` returns the (base) absolute deadline of the task.

The procedure `Set_Relative_Deadline` changes the relative deadline of the task to `D`. The function `Get_Relative_Deadline` returns the relative deadline of the task.

The function `Get_Last_Release_Time` returns the time, as provided by `Real_Time.Clock`, when the task was last made ready (that is, was added to a ready queue).

The procedure `Delay_Until_And_Set_Deadline` delays the calling task until time `Delay_Until_Time`. When the task becomes ready again it will have deadline `Delay_Until_Time + Deadline_Offset`.

On a system with a single processor, the setting of the deadline of a task to the new value occurs immediately at the first point that is outside the execution of a protected action. If the task is currently on a ready queue it is removed and re-entered onto the ready queue determined by the rules defined below.

When `EDF_Within_Priorities` is specified for a priority, the ready queue for that priority is ordered by deadline. The task at the head of a queue is the one with the earliest deadline.

A task dispatching point occurs for the currently running task *T* to which policy `EDF_Within_Priorities` applies:

- when a change to the base (absolute) deadline of *T* occurs;
- there is a nonempty ready queue for that processor with a higher priority than the active priority of the running task;
- there is a ready task with the same priority as *T* but with an earlier absolute deadline.

In these cases, the currently running task is said to be preempted and is returned to the ready queue for its active priority, at a position determined by its active (absolute) deadline.

When the setting of the base priority of a ready task takes effect and the new priority is specified as `EDF_Within_Priorities`, the task is added to the ready queue, at a position determined by its active deadline.

For all the operations defined in `Dispatching.EDF`, `Tasking_Error` is raised if the task identified by *T* has terminated. `Program_Error` is raised if the value of *T* is `Null_Task_Id`.

If two tasks with priority designated as `EDF_Within_Priorities` rendezvous then the deadline for the execution of the accept statement is the earlier of the deadlines of the two tasks.

During activation, a task being activated inherits the deadline that its activator (see 12.2) had at the time the activation was initiated.

#### *Erroneous Execution*

If a value of `Task_Id` is passed as a parameter to any of the subprograms of this package and the corresponding task object no longer exists, the execution of the program is erroneous.

#### *Documentation Requirements*

On a multiprocessor, the implementation shall document any conditions that cause the completion of the setting of the deadline of a task to be delayed later than what is specified for a single processor.

NOTE If two distinct priorities are specified to have policy `EDF_Within_Priorities`, then tasks from the higher priority always run before tasks of the lower priority, regardless of deadlines.

### D.3 Priority Ceiling Locking

This subclause specifies the interactions between priority task scheduling and protected object ceilings. This interaction is based on the concept of the *ceiling priority* of a protected object.

#### *Syntax*

The form of a `pragma Locking_Policy` is as follows:

**pragma** `Locking_Policy`(*policy\_identifier*);

#### *Legality Rules*

The *policy\_identifier* shall either be `Ceiling_Locking` or an implementation-defined *identifier*.

#### *Post-Compilation Rules*

A `Locking_Policy` pragma is a configuration pragma.

#### *Dynamic Semantics*

A locking policy specifies the details of protected object locking. All protected objects have a priority. The locking policy specifies the meaning of the priority of a protected object, and the relationships between these priorities and task priorities. In addition, the policy specifies the state of a task when it executes a protected action, and how its active priority is affected by the locking. The *locking policy* is specified by a `Locking_Policy` pragma. For implementation-defined locking policies, the meaning of the priority of a protected object is implementation defined. If no `Locking_Policy` pragma applies to any of the program units comprising a partition, the locking policy for that partition, as well as the meaning of the priority of a protected object, are implementation defined.

The *expression* specified for the `Priority` or `Interrupt_Priority` aspect (see D.1) is evaluated as part of the creation of the corresponding protected object and converted to the subtype `System.Any_Priority` or `System.Interrupt_Priority`, respectively. The value of the expression is the initial priority of the corresponding protected object. If no `Priority` or `Interrupt_Priority` aspect is specified for a protected object, the initial priority is specified by the locking policy.

There is one predefined locking policy, `Ceiling_Locking`; this policy is defined as follows:

- Every protected object has a *ceiling priority*, which is determined by either a `Priority` or `Interrupt_Priority` aspect as defined in D.1, or by assignment to the `Priority` attribute as described in D.5.2. The ceiling priority of a protected object (or ceiling, for short) is an upper bound on the active priority a task can have when it calls protected operations of that protected object.
- The initial ceiling priority of a protected object is equal to the initial priority for that object.
- If an `Interrupt_Handler` or `Attach_Handler` aspect (see C.3.1) is specified for a protected subprogram of a protected type that does not have either the `Priority` or `Interrupt_Priority` aspect specified, the initial priority of protected objects of that type is implementation defined, but in the range of the subtype `System.Interrupt_Priority`.



- If neither aspect *Priority* nor *Interrupt\_Priority* is specified for a protected type, and no protected subprogram of the type has aspect *Interrupt\_Handler* or *Attach\_Handler* specified, then the initial priority of the corresponding protected object is *System.Priority'Last*.
- While a task executes a protected action, it inherits the ceiling priority of the corresponding protected object.
- When a task calls a protected operation, a check is made that its active priority is not higher than the ceiling of the corresponding protected object; *Program\_Error* is raised if this check fails.

If the task dispatching policy specified for the ceiling priority of a protected object is *EDF\_Within\_Priorities*, the following additional rules apply:

- Every protected object has a *relative deadline*, which is determined by a *Relative\_Deadline* aspect as defined in D.2.6, or by assignment to the *Relative\_Deadline* attribute as described in D.5.2. The relative deadline of a protected object represents a lower bound on the relative deadline a task may have when it calls a protected operation of that protected object.
- If aspect *Relative\_Deadline* is not specified for a protected type then the initial relative deadline of the corresponding protected object is *Ada.Real\_Time.Time\_Span\_Zero*.
- While a task executes a protected action on a protected object *P*, it inherits the relative deadline of *P*. In this case, let *DF* be 'now' ('now' is obtained via a call on *Ada.Real\_Time.Clock* at the start of the action) plus the deadline floor of *P*. If the active deadline of the task is later than *DF*, its active deadline is reduced to *DF*; the active deadline is unchanged otherwise.
- When a task calls a protected operation, a check is made that its active deadline minus its last release time is not less than the relative deadline of the corresponding protected object; *Program\_Error* is raised if this check fails.

#### *Bounded (Run-Time) Errors*

Following any change of priority, it is a bounded error for the active priority of any task with a call queued on an entry of a protected object to be higher than the ceiling priority of the protected object. In this case one of the following applies:

- at any time prior to executing the entry body, *Program\_Error* is raised in the calling task;
- when the entry is open, the entry body is executed at the ceiling priority of the protected object;
- when the entry is open, the entry body is executed at the ceiling priority of the protected object and then *Program\_Error* is raised in the calling task; or
- when the entry is open, the entry body is executed at the ceiling priority of the protected object that was in effect when the entry call was queued.

#### *Implementation Permissions*

The implementation is allowed to round all ceilings in a certain subrange of *System.Priority* or *System.Interrupt\_Priority* up to the top of that subrange, uniformly.

Implementations are allowed to define other locking policies, but are not required to support specifying more than one locking policy per partition.

Since implementations are allowed to place restrictions on code that runs at an interrupt-level active priority (see C.3.1 and D.2.1), the implementation may implement a language feature in terms of a protected object with an implementation-defined ceiling, but the ceiling shall be no less than *Priority'Last*.

#### *Implementation Advice*

The implementation should use names that end with “\_Locking” for implementation-defined locking policies.

NOTE 1 While a task executes in a protected action, it can be preempted only by tasks whose active priorities are higher than the ceiling priority of the protected object.

NOTE 2 If a protected object has a ceiling priority in the range of `Interrupt_Priority`, certain interrupts are blocked while protected actions of that object execute. In the extreme, if the ceiling is `Interrupt_Priority'Last`, all blockable interrupts are blocked during that time.

NOTE 3 As described in C.3.1, whenever an interrupt is handled by one of the protected procedures of a protected object, a check is made that its ceiling priority is in the `Interrupt_Priority` range.

NOTE 4 When specifying the ceiling of a protected object, a correct value is one that is at least as high as the highest active priority at which tasks can be executing when they call protected operations of that object. In determining this value the following factors, which can affect active priority, are relevant: the effect of `Set_Priority`, nested protected operations, entry calls, task activation, and other implementation-defined factors.

NOTE 5 Attaching a protected procedure whose ceiling is below the interrupt hardware priority to an interrupt causes the execution of the program to be erroneous (see C.3.1).

NOTE 6 On a single processor implementation, the ceiling priority rules guarantee that there is no possibility of deadlock involving only protected subprograms (excluding the case where a protected operation calls another protected operation on the same protected object).

## D.4 Entry Queuing Policies

This subclause specifies a mechanism for a user to choose an entry *queuing policy*. It also defines three such policies. Other policies are implementation defined.

### Syntax

The form of a `pragma Queuing_Policy` is as follows:

**pragma** `Queuing_Policy`(*policy\_identifier*);

### Legality Rules

The *policy\_identifier* shall be either `FIFO_Queueing`, `Ordered_FIFO_Queueing`, `Priority_Queueing` or an implementation-defined *identifier*.

### Post-Compilation Rules

A `Queuing_Policy` pragma is a configuration pragma.

### Dynamic Semantics

A *queuing policy* governs the order in which tasks are queued for entry service, and the order in which different entry queues are considered for service. The queuing policy is specified by a `Queuing_Policy` pragma.

Three queuing policies, `FIFO_Queueing`, `Ordered_FIFO_Queueing`, and `Priority_Queueing`, are language defined. If no `Queuing_Policy` pragma applies to any of the program units comprising the partition, the queuing policy for that partition is `FIFO_Queueing`. The rules for the `FIFO_Queueing` policy are specified in 12.5.3 and 12.7.1.

The `Ordered_FIFO_Queueing` policy is defined as follows:

- Calls are selected on a given entry queue in order of arrival.
- When more than one condition of an `entry_barrier` of a protected object becomes `True`, and more than one of the respective queues is nonempty, the call that arrived first is selected.
- If the expiration time of two or more open `delay_alternatives` is the same and no other `accept_alternatives` are open, the sequence of statements of the `delay_alternative` that is first in textual order in the `selective_accept` is executed.
- When more than one alternative of a `selective_accept` is open and has queued calls, the alternative whose queue has the call that arrived first is selected.

The `Priority_Queueing` policy is defined as follows:

- The calls to an entry (including a member of an entry family) are queued in an order consistent with the priorities of the calls. The *priority of an entry call* is initialized from the active priority of

the calling task at the time the call is made, but can change later. Within the same priority, the order is consistent with the calling (or requeuing, or priority setting) time (that is, a FIFO order).

- After a call is first queued, changes to the active priority of a task do not affect the priority of the call, unless the base priority of the task is set while the task is blocked on an entry call.
- When the base priority of a task is set (see D.5), if the task is blocked on an entry call, and the call is queued, the priority of the call is updated to the new active priority of the calling task. This causes the call to be removed from and then reinserted in the queue at the new active priority.
- When more than one `condition` of an `entry_barrier` of a protected object becomes True, and more than one of the respective queues is nonempty, the call with the highest priority is selected. If more than one such call has the same priority, the call that is queued on the entry whose declaration is first in textual order in the `protected_definition` is selected. For members of the same entry family, the one with the lower family index is selected.
- If the expiration time of two or more open `delay_alternatives` is the same and no other `accept_alternatives` are open, the `sequence_of_statements` of the `delay_alternative` that is first in textual order in the `selective_accept` is executed.
- When more than one alternative of a `selective_accept` is open and has queued calls, an alternative whose queue has the highest-priority call at its head is selected. If two or more open alternatives have equal-priority queued calls, then a call on the entry in the `accept_alternative` that is first in textual order in the `selective_accept` is selected.

#### Implementation Permissions

Implementations are allowed to define other queuing policies, but are not required to support specifying more than one queuing policy per partition.

Implementations are allowed to defer the reordering of entry queues following a change of base priority of a task blocked on the entry call if it is not practical to reorder the queue immediately.

#### Implementation Advice

The implementation should use names that end with “\_Queuing” for implementation-defined queuing policies.

#### Static Semantics

For a task type (including the anonymous type of a `single_task_declaration`), protected type (including the anonymous type of a `single_protected_declaration`), or an `entry_declaration`, the following language-defined representation aspect may be specified:

##### Max\_Entry\_Queue\_Length

The type of aspect `Max_Entry_Queue_Length` is Integer.

If directly specified, the `aspect_definition` shall be a static expression no less than -1. If not specified, the aspect has value -1 (representing no additional restriction on queue length).

#### Legality Rules

If the `Max_Entry_Queue_Length` aspect for a type has a nonnegative value, the `Max_Entry_Queue_Length` aspect for every individual entry of that type shall not be greater than the value of the aspect for the type. The `Max_Entry_Queue_Length` aspect of a type is nonoverridable (see 16.1.1).

#### Post-Compilation Rules

If a restriction `Max_Entry_Queue_Length` applies to a partition, any value specified for the `Max_Entry_Queue_Length` aspect specified for the declaration of a type or entry in the partition shall not be greater than the value of the restriction.

*Dynamic Semantics*

If a nonconfirming value is specified for `Max_Entry_Queue_Length` for a type, and an entry call or requeue would cause the queue for any entry of the type to become longer than the specified value, then `Program_Error` is raised at the point of the call or requeue.

If a nonconfirming value is specified for `Max_Entry_Queue_Length` for an entry, and an entry call or requeue would cause the queue for an entry to become longer than the specified value, then `Program_Error` is raised at the point of the call or requeue.

**D.4.1 Admission Policies**

This subclause specifies a mechanism for a user to choose an admission policy. It also defines one such policy. Other policies are implementation defined.

*Syntax*

The form of a pragma `Admission_Policy` is as follows:

**pragma** `Admission_Policy` (*policy\_identifier*);

*Legality Rules*

The *policy\_identifier* shall be either `FIFO_Spinning` or an implementation-defined identifier.

*Post-Compilation Rules*

An `Admission_Policy` pragma is a configuration pragma.

*Dynamic Semantics*

An admission policy governs the order in which competing tasks are evaluated for acquiring the execution resource associated with a protected object. The admission policy is specified by an `Admission_Policy` pragma.

One admission policy, `FIFO_Spinning`, is language defined. If `FIFO_Spinning` is in effect, and starting a protected action on a protected object involves busy-waiting, then calls are selected for acquiring the execution resource of the protected object in the order in which the busy-wait was initiated; otherwise the `FIFO_Spinning` policy has no effect. If no `Admission_Policy` pragma applies to any of the program units in the partition, the admission policy for that partition is implementation defined.

*Implementation Permissions*

Implementations are allowed to define other admission policies, but are not required to support specifying more than one admission policy per partition.

**D.5 Dynamic Priorities**

This subclause describes how the priority of an entity can be modified or queried at run time.

**D.5.1 Dynamic Priorities for Tasks**

This subclause describes how the base priority of a task can be modified or queried at run time.

*Static Semantics*

The following language-defined library package exists:

```
with System;
with Ada.Task_Identification; -- See C.7.1
package Ada.Dynamic_Priorities
  with Preelaborate, Nonblocking, Global => in out synchronized is
  procedure Set_Priority(Priority : in System.Any_Priority;
                        T : in Ada.Task_Identification.Task_Id :=
                        Ada.Task_Identification.Current_Task);
```

```

function Get_Priority (T : Ada.Task_Identification.Task_Id :=
                      Ada.Task_Identification.Current_Task)
return System.Any_Priority;

end Ada.Dynamic_Priorities;

```

#### *Dynamic Semantics*

The procedure `Set_Priority` sets the base priority of the specified task to the specified `Priority` value. `Set_Priority` has no effect if the task is terminated.

The function `Get_Priority` returns `T`'s current base priority. `Tasking_Error` is raised if the task is terminated.

`Program_Error` is raised by `Set_Priority` and `Get_Priority` if `T` is equal to `Null_Task_Id`.

On a system with a single processor, the setting of the base priority of a task *T* to the new value occurs immediately at the first point when *T* is outside the execution of a protected action.

#### *Erroneous Execution*

If any subprogram in this package is called with a parameter `T` that specifies a task object that no longer exists, the execution of the program is erroneous.

#### *Documentation Requirements*

On a multiprocessor, the implementation shall document any conditions that cause the completion of the setting of the priority of a task to be delayed later than what is specified for a single processor.

#### *Metrics*

The implementation shall document the following metric:

- The execution time of a call to `Set_Priority`, for the nonpreempting case, in processor clock cycles. This is measured for a call that modifies the priority of a ready task that is not running (which cannot be the calling one), where the new base priority of the affected task is lower than the active priority of the calling task, and the affected task is not on any entry queue and is not executing a protected operation.

NOTE 1 Setting a task's base priority affects task dispatching. First, it can change the task's active priority. Second, under the `FIFO_Within_Priorities` policy, it always causes the task to move to the tail of the ready queue corresponding to its active priority, even if the new base priority is unchanged.

NOTE 2 Under the priority queuing policy, setting a task's base priority has an effect on a queued entry call if the task is blocked waiting for the call. That is, setting the base priority of a task causes the priority of a queued entry call from that task to be updated and the call to be removed and then reinserted in the entry queue at the new priority (see D.4), unless the call originated from the `triggering_statement` of an `asynchronous_select`.

NOTE 3 The effect of two or more `Set_Priority` calls executed in parallel on the same task is defined as executing these calls in some serial order.

NOTE 4 The rule for when `Tasking_Error` is raised for `Set_Priority` or `Get_Priority` is different from the rule for when `Tasking_Error` is raised on an entry call (see 12.5.3). In particular, the priority of a completed or an abnormal task can be queried, so long as the task is not yet terminated, and the priority of a task can be set in any task state (including for terminated tasks).

NOTE 5 Changing the priorities of a set of tasks can be performed by a series of calls to `Set_Priority` for each task separately. This can be done reliably within a protected operation that has high enough ceiling priority to guarantee that the operation completes without being preempted by any of the affected tasks.

### **D.5.2 Dynamic Priorities for Protected Objects**

This subclause specifies how the priority of a protected object can be modified or queried at run time.

*Static Semantics*

The following attributes are defined for a prefix *P* that denotes a protected object:

**P'Priority** Denotes a non-aliased component of the protected object *P*. This component is of type *System.Any\_Priority* and its value is the priority of *P*. *P'Priority* denotes a variable if and only if *P* denotes a variable. A reference to this attribute shall appear only within the body of *P*.

**P'Relative\_Deadline**

Denotes a non-aliased component of the protected object *P*. This component is of type *Ada.Real\_Time.Time\_Span* and its value is the relative deadline of *P*. *P'Relative\_Deadline* denotes a variable if and only if *P* denotes a variable. A reference to this attribute shall appear only within the body of *P*.

The initial value of the attribute *Priority* is determined by the initial value of the priority of the protected object (see D.3), and can be changed by an assignment. The initial value of the attribute *Relative\_Deadline* is determined by the initial value of the relative deadline of the protected object (see D.3), and can be changed by an assignment.

*Dynamic Semantics*

If the locking policy *Ceiling\_Locking* (see D.3) is in effect, then the ceiling priority of a protected object *P* is set to the value of *P'Priority* at the end of each protected action of *P*.

If the locking policy *Ceiling\_Locking* is in effect, then for a protected object *P* with either an *Attach\_Handler* or *Interrupt\_Handler* aspect specified for one of its procedures, a check is made that the value to be assigned to *P'Priority* is in the range *System.Interrupt\_Priority*. If the check fails, *Program\_Error* is raised.

*Metrics*

The implementation shall document the following metric:

- The difference in execution time of calls to the following procedures in protected object *P*:

```
protected P is
  procedure Do_Not_Set_Ceiling (Pr : System.Any_Priority);
  procedure Set_Ceiling (Pr : System.Any_Priority);
end P;

protected body P is
  procedure Do_Not_Set_Ceiling (Pr : System.Any_Priority) is
  begin
    null;
  end;
  procedure Set_Ceiling (Pr : System.Any_Priority) is
  begin
    P'Priority := Pr;
  end;
end P;
```

NOTE Since *P'Priority* is a normal variable, the value following an assignment to the attribute immediately reflects the new value even though its impact on the ceiling priority of *P* is postponed until completion of the protected action in which it is executed.

## D.6 Preemptive Abort

This subclause specifies requirements on the immediacy with which an aborted construct is completed.

*Dynamic Semantics*

On a system with a single processor, an aborted construct is completed immediately at the first point that is outside the execution of an abort-deferred operation.



*Documentation Requirements*

On a multiprocessor, the implementation shall document any conditions that cause the completion of an aborted construct to be delayed later than what is specified for a single processor.

*Metrics*

The implementation shall document the following metrics:

- The execution time, in processor clock cycles, that it takes for an `abort_statement` to cause the completion of the aborted task. This is measured in a situation where a task T2 preempts task T1 and aborts T1. T1 does not have any finalization code. T2 shall verify that T1 has terminated, by means of the Terminated attribute.
- On a multiprocessor, an upper bound in seconds, on the time that the completion of an aborted task can be delayed beyond the point that it is required for a single processor.
- An upper bound on the execution time of an `asynchronous_select`, in processor clock cycles. This is measured between a point immediately before a task T1 executes a protected operation `Pr.Set` that makes the condition of an `entry_barrier` `Pr.Wait` True, and the point where task T2 resumes execution immediately after an entry call to `Pr.Wait` in an `asynchronous_select`. T1 preempts T2 while T2 is executing the abortable part, and then blocks itself so that T2 can execute. The execution time of T1 is measured separately, and subtracted.
- An upper bound on the execution time of an `asynchronous_select`, in the case that no asynchronous transfer of control takes place. This is measured between a point immediately before a task executes the `asynchronous_select` with a nonnull abortable part, and the point where the task continues execution immediately after it. The execution time of the abortable part is subtracted.

*Implementation Advice*

Even though the `abort_statement` is included in the list of potentially blocking operations (see 12.5.1), it is recommended that this statement be implemented in a way that never requires the task executing the `abort_statement` to block.

On a multi-processor, the delay associated with aborting a task on another processor should be bounded; the implementation should use periodic polling, if necessary, to achieve this.

NOTE 1 Abortion does not change the active or base priority of the aborted task.

NOTE 2 Abortion cannot be more immediate than is allowed by the rules for deferral of abortion during finalization and in protected actions.

## D.7 Tasking Restrictions

This subclause defines restrictions that can be used with a `pragma Restrictions` (see 16.12) to facilitate the construction of highly efficient tasking run-time systems.

*Static Semantics*

A scalar expression within a protected unit is said to be *pure-barrier-eligible* if it is one of the following:

- a static expression;
- a `name` that statically names (see 7.9) a scalar subcomponent of the immediately enclosing protected unit;
- a `Count` attribute `reference` whose `prefix` statically denotes an entry declaration of the immediately enclosing unit;
- a call to a predefined relational operator or boolean logical operator (**and**, **or**, **xor**, **not**), where each operand is pure-barrier-eligible;

- a membership test whose *tested\_simple\_expression* is pure-barrier-eligible, and whose *membership\_choice\_list* meets the requirements for a static membership test (see 7.9);
- a short-circuit control form both of whose operands are pure-barrier-eligible;
- a *conditional\_expression* all of whose conditions, *selecting\_expressions*, and *dependent\_expressions* are pure-barrier-eligible; or
- a pure-barrier-eligible expression enclosed in parentheses.

The following *restriction\_identifiers* are language defined:

#### No\_Task\_Hierarchy

No task depends on a master other than the library-level master.

#### No\_Nested\_Finalization

Objects of a type that needs finalization (see 10.6) are declared only at library level. If an access type does not have library-level accessibility, then there are no allocators of the type where the type determined by the *subtype\_mark* of the *subtype\_indication* or *qualified\_expression* needs finalization.

#### No\_Abort\_Statements

There are no *abort\_statements*, and there is no use of a name denoting *Task\_Identification.Abort\_Task*.

#### No\_Terminate\_Alternatives

There are no *selective\_accepts* with *terminate\_alternatives*.

#### No\_Task\_Allocators

There are no allocators for task types or types containing task subcomponents.

In the case of an initialized allocator of an access type whose designated type is class-wide and limited, a check is made that the specific type of the allocated object has no task subcomponents. *Program\_Error* is raised if this check fails.

#### No\_Implicit\_Heap\_Allocations

There are no operations that implicitly require heap storage allocation to be performed by the implementation. The operations that implicitly require heap storage allocation are implementation defined.

#### No\_Dynamic\_Priorities

There are no semantic dependences on the package *Dynamic\_Priorities*, and no occurrences of the attribute *Priority*.

#### No\_Dynamic\_Attachment

There is no use of a name denoting any of the operations defined in package *Interrupts* (*Is\_Reserved*, *Is\_Attached*, *Current\_Handler*, *Attach\_Handler*, *Exchange\_Handler*, *Detach\_Handler*, and *Reference*).

#### No\_Dynamic\_CPU\_Assignment

No task has the CPU aspect specified to be a non-static expression. Each task (including the environment task) that has the CPU aspect specified as *Not\_A\_Specific\_CPU* will be assigned to a particular implementation-defined CPU. The same is true for the environment task when the CPU aspect is not specified. Any other task without a CPU aspect will activate and execute on the same processor as its activating task.

#### No\_Local\_Protected\_Objects

Protected objects are declared only at library level.

#### No\_Local\_Timing\_Events

*Timing\_Events* are declared only at library level.

No\_Protected\_Type\_Allocators

There are no allocators for protected types or types containing protected type subcomponents.

In the case of an initialized allocator of an access type whose designated type is class-wide and limited, a check is made that the specific type of the allocated object has no protected subcomponents. Program\_Error is raised if this check fails.

No\_Relative\_Delay

There are no delay\_relative\_statements, and there is no use of a name that denotes the Timing\_Events.Set\_Handler subprogram that has a Time\_Span parameter.

No\_Requeue\_Statements

There are no requeue\_statements.

No\_Select\_Statements

There are no select\_statements.

No\_Specific\_Termination\_Handlers

There is no use of a name denoting the Set\_Specific\_Handler and Specific\_Handler subprograms in Task\_Termination.

No\_Tasks\_Unassigned\_To\_CPU

The CPU aspect is specified for the environment task. No CPU aspect is specified to be statically equal to Not\_A\_Specific\_CPU. If aspect CPU is specified (dynamically) to the value Not\_A\_Specific\_CPU, then Program\_Error is raised. If Set\_CPU or Delay\_Until\_And\_Set\_CPU are called with the CPU parameter equal to Not\_A\_Specific\_CPU, then Program\_Error is raised.

Pure\_Barriers

The Boolean expression in each protected entry barrier is pure-barrier-eligible.

Simple\_Barriers

The Boolean expression in each entry barrier is either a static expression or a name that statically names (see 7.9) a subcomponent of the enclosing protected object.

The following *restriction\_parameter\_identifiers* are language defined:

Max\_Select\_Alternatives

Specifies the maximum number of alternatives in a selective\_accept.

Max\_Task\_Entries

Specifies the maximum number of entries per task. The bounds of every entry family of a task unit shall be static, or shall be defined by a discriminant of a subtype whose corresponding bound is static. A value of zero indicates that no rendezvous are possible.

Max\_Protected\_Entries

Specifies the maximum number of entries per protected type. The bounds of every entry family of a protected unit shall be static, or shall be defined by a discriminant of a subtype whose corresponding bound is static.

*Dynamic Semantics*

The following *restriction\_identifier* is language defined:

No\_Task\_Termination

All tasks are nonterminating. It is implementation-defined what happens if a task attempts to terminate. If there is a fall-back handler (see C.7.3) set for the partition it should be called when the first task attempts to terminate.

The following *restriction\_parameter\_identifiers* are language defined:

Max\_Storage\_At\_Blocking

Specifies the maximum portion (in storage elements) of a task's Storage\_Size that can be retained by a blocked task. If an implementation chooses to detect a violation of this

restriction, `Storage_Error` should be raised; otherwise, the behavior is implementation defined.

#### Max\_Asynchronous\_Select\_Nesting

Specifies the maximum dynamic nesting level of `asynchronous_selects`. A value of zero prevents the use of any `asynchronous_select` and, if a program contains an `asynchronous_select`, it is illegal. If an implementation chooses to detect a violation of this restriction for values other than zero, `Storage_Error` should be raised; otherwise, the behavior is implementation defined.

#### Max\_Tasks

Specifies the maximum number of task creations that may be executed over the lifetime of a partition, not counting the creation of the environment task. A value of zero prevents any task creation and, if a program contains a task creation, it is illegal. If an implementation chooses to detect a violation of this restriction, `Storage_Error` should be raised; otherwise, the behavior is implementation defined.

#### Max\_Entry\_Queue\_Length

`Max_Entry_Queue_Length` defines the maximum number of calls that are queued on an entry. Violation of this restriction results in the raising of `Program_Error` at the point of the call or requeue.

#### No\_Standard\_Allocators\_After\_Elaboration

Specifies that an allocator using a standard storage pool (see 16.11) shall not occur within a parameterless library subprogram, nor within the `handled_sequence_of_statements` of a task body. For the purposes of this rule, an allocator of a type derived from a formal access type does not use a standard storage pool.

At run time, `Storage_Error` is raised if an allocator using a standard storage pool is evaluated after the elaboration of the `library_items` of the partition has completed.

It is implementation defined whether the use of pragma Restrictions results in a reduction in executable program size, storage requirements, or execution time. If possible, the implementation should provide quantitative descriptions of such effects for each restriction.

#### Implementation Advice

When feasible, the implementation should take advantage of the specified restrictions to produce a more efficient implementation.

NOTE The above `Storage_Checks` can be suppressed with pragma Suppress.

## D.8 Monotonic Time

This subclause specifies a high-resolution, monotonic clock package.

#### Static Semantics

The following language-defined library package exists:

```
package Ada.Real_Time
  with Nonblocking, Global => in out synchronized is

  type Time is private;
  Time_First : constant Time;
  Time_Last : constant Time;
  Time_Unit : constant := implementation-defined-real-number;

  type Time_Span is private;
  Time_Span_First : constant Time_Span;
  Time_Span_Last : constant Time_Span;
  Time_Span_Zero : constant Time_Span;
  Time_Span_Unit : constant Time_Span;
```

```

Tick : constant Time_Span;
function Clock return Time;

function "+" (Left : Time; Right : Time_Span) return Time;
function "+" (Left : Time_Span; Right : Time) return Time;
function "-" (Left : Time; Right : Time_Span) return Time;
function "-" (Left : Time; Right : Time) return Time_Span;

function "<" (Left, Right : Time) return Boolean;
function "<=" (Left, Right : Time) return Boolean;
function ">" (Left, Right : Time) return Boolean;
function ">=" (Left, Right : Time) return Boolean;

function "+" (Left, Right : Time_Span) return Time_Span;
function "-" (Left, Right : Time_Span) return Time_Span;
function "-" (Right : Time_Span) return Time_Span;
function "*" (Left : Time_Span; Right : Integer) return Time_Span;
function "*" (Left : Integer; Right : Time_Span) return Time_Span;
function "/" (Left, Right : Time_Span) return Integer;
function "/" (Left : Time_Span; Right : Integer) return Time_Span;

function "abs" (Right : Time_Span) return Time_Span;

function "<" (Left, Right : Time_Span) return Boolean;
function "<=" (Left, Right : Time_Span) return Boolean;
function ">" (Left, Right : Time_Span) return Boolean;
function ">=" (Left, Right : Time_Span) return Boolean;

function To_Duration (TS : Time_Span) return Duration;
function To_Time_Span (D : Duration) return Time_Span;

function Nanoseconds (NS : Integer) return Time_Span;
function Microseconds (US : Integer) return Time_Span;
function Milliseconds (MS : Integer) return Time_Span;
function Seconds (S : Integer) return Time_Span;
function Minutes (M : Integer) return Time_Span;

type Seconds_Count is range implementation-defined;

procedure Split(T : in Time; SC : out Seconds_Count; TS : out Time_Span);
function Time_Of(SC : Seconds_Count; TS : Time_Span) return Time;

private
    ... -- not specified by the language
end Ada.Real_Time;

```

In this Annex, *real time* is defined to be the physical time as observed in the external environment. The type *Time* is a *time type* as defined by 12.6; values of this type may be used in a *delay\_until\_statement*. Values of this type represent segments of an ideal time line. The set of values of the type *Time* corresponds one-to-one with an implementation-defined range of mathematical integers.

The *Time* value *I* represents the half-open real time interval that starts with  $E + I \cdot \text{Time\_Unit}$  and is limited by  $E + (I + 1) \cdot \text{Time\_Unit}$ , where *Time\_Unit* is an implementation-defined real number and *E* is an unspecified origin point, the *epoch*, that is the same for all values of the type *Time*. It is not specified by the language whether the time values are synchronized with any standard time reference. For example, *E* can correspond to the time of system initialization or it can correspond to the epoch of some time standard.

Values of the type *Time\_Span* represent length of real time duration. The set of values of this type corresponds one-to-one with an implementation-defined range of mathematical integers. The *Time\_Span* value corresponding to the integer *I* represents the real-time duration  $I \cdot \text{Time\_Unit}$ .

*Time\_First* and *Time\_Last* are the smallest and largest values of the *Time* type, respectively. Similarly, *Time\_Span\_First* and *Time\_Span\_Last* are the smallest and largest values of the *Time\_Span* type, respectively.

A value of type *Seconds\_Count* represents an elapsed time, measured in seconds, since the epoch.

*Dynamic Semantics*

Time\_Unit is the smallest amount of real time representable by the Time type; it is expressed in seconds. Time\_Span\_Unit is the difference between two successive values of the Time type. It is also the smallest positive value of type Time\_Span. Time\_Unit and Time\_Span\_Unit represent the same real time duration. A *clock tick* is a real time interval during which the clock value (as observed by calling the Clock function) remains constant. Tick is the average length of such intervals.

The function To\_Duration converts the value TS to a value of type Duration. Similarly, the function To\_Time\_Span converts the value D to a value of type Time\_Span. For To\_Duration, the result is rounded to the nearest value of type Duration (away from zero if exactly halfway between two values). If the result is outside the range of Duration, Constraint\_Error is raised. For To\_Time\_Span, the value of D is first rounded to the nearest integral multiple of Time\_Unit, away from zero if exactly halfway between two multiples. If the rounded value is outside the range of Time\_Span, Constraint\_Error is raised. Otherwise, the value is converted to the type Time\_Span.

To\_Duration(Time\_Span\_Zero) returns 0.0, and To\_Time\_Span(0.0) returns Time\_Span\_Zero.

The functions Nanoseconds, Microseconds, Milliseconds, Seconds, and Minutes convert the input parameter to a value of the type Time\_Span. NS, US, MS, S, and M are interpreted as a number of nanoseconds, microseconds, milliseconds, seconds, and minutes respectively. The input parameter is first converted to seconds and rounded to the nearest integral multiple of Time\_Unit, away from zero if exactly halfway between two multiples. If the rounded value is outside the range of Time\_Span, Constraint\_Error is raised. Otherwise, the rounded value is converted to the type Time\_Span.

The effects of the operators on Time and Time\_Span are as for the operators defined for integer types.

The function Clock returns the amount of time since the epoch.

The effects of the Split and Time\_Of operations are defined as follows, treating values of type Time, Time\_Span, and Seconds\_Count as mathematical integers. The effect of Split(T,SC,TS) is to set SC and TS to values such that  $T \cdot \text{Time\_Unit} = \text{SC} \cdot 1.0 + \text{TS} \cdot \text{Time\_Unit}$ , and  $0.0 \leq \text{TS} \cdot \text{Time\_Unit} < 1.0$ . The value returned by Time\_Of(SC,TS) is the value T such that  $T \cdot \text{Time\_Unit} = \text{SC} \cdot 1.0 + \text{TS} \cdot \text{Time\_Unit}$ .

*Implementation Requirements*

The range of Time values shall be sufficient to uniquely represent the range of real times from program start-up to 50 years later. Tick shall be no greater than 1 millisecond. Time\_Unit shall be less than or equal to 20 microseconds.

The value of Time\_Span\_First in seconds shall be no greater than -3600, and the value of Time\_Span\_Last in seconds shall be no less than 3600.

A *clock jump* is the difference between two successive distinct values of the clock (as observed by calling the Clock function). There shall be no backward clock jumps.

*Documentation Requirements*

The implementation shall document the values of Time\_First, Time\_Last, Time\_Span\_First, Time\_Span\_Last, Time\_Span\_Unit, and Tick.

The implementation shall document the properties of the underlying time base used for the clock and for type Time, such as the range of values supported and any relevant aspects of the underlying hardware or operating system facilities used.

The implementation shall document whether or not there is any synchronization with external time references, and if such synchronization exists, the sources of synchronization information, the frequency of synchronization, and the synchronization method applied.

The implementation shall document any aspects of the external environment that can interfere with the clock behavior as defined in this subclause.



*Metrics*

For the purpose of the metrics defined in this subclause, real time is defined to be the International Atomic Time (TAI).

The implementation shall document the following metrics:

- An upper bound on the real-time duration of a clock tick. This is a value  $D$  such that if  $t_1$  and  $t_2$  are any real times such that  $t_1 < t_2$  and  $\text{Clock}_{t_1} = \text{Clock}_{t_2}$  then  $t_2 - t_1 \leq D$ .
- An upper bound on the size of a clock jump.
- An upper bound on the *drift rate* of Clock with respect to real time. This is a real number  $D$  such that
 
$$E \cdot (1 - D) \leq (\text{Clock}_{t+E} - \text{Clock}_t) \leq E \cdot (1 + D)$$
 provided that:  $\text{Clock}_t + E \cdot (1 + D) \leq \text{Time\_Last}$ .
- where  $\text{Clock}_t$  is the value of Clock at time  $t$ , and  $E$  is a real time duration not less than 24 hours. The value of  $E$  used for this metric shall be reported.
- An upper bound on the execution time of a call to the Clock function, in processor clock cycles.
- Upper bounds on the execution times of the operators of the types Time and Time\_Span, in processor clock cycles.

*Implementation Permissions*

Implementations targeted to machines with word size smaller than 32 bits may omit support for the full range and granularity of the Time and Time\_Span types.

*Implementation Advice*

When appropriate, implementations should provide configuration mechanisms to change the value of Tick.

It is recommended that Calendar.Clock and Real\_Time.Clock be implemented as transformations of the same time base.

It is recommended that the “best” time base which exists in the underlying system be available to the application through Clock. “Best” may mean highest accuracy or largest range.

NOTE 1 The rules in this subclause do not imply that the implementation can protect the user from operator or installation errors that can result in the clock being set incorrectly.

NOTE 2 Time\_Unit is the granularity of the Time type. In contrast, Tick represents the granularity of Real\_Time.Clock. There is no requirement that these be the same.

## D.9 Delay Accuracy

This subclause specifies performance requirements for the `delay_statement`. The rules apply both to `delay_relative_statement` and to `delay_until_statement`. Similarly, they apply equally to a simple `delay_statement` and to one which appears in a `delay_alternative`.

*Dynamic Semantics*

The effect of the `delay_statement` for Real\_Time.Time is defined in terms of Real\_Time.Clock:

- If  $C_1$  is a value of Clock read before a task executes a `delay_relative_statement` with duration  $D$ , and  $C_2$  is a value of Clock read after the task resumes execution following that `delay_statement`, then  $C_2 - C_1 \geq D$ .
- If  $C$  is a value of Clock read after a task resumes execution following a `delay_until_statement` with Real\_Time.Time value  $T$ , then  $C \geq T$ .

A simple `delay_statement` with a negative or zero value for the expiration time does not cause the calling task to be blocked; it is nevertheless a potentially blocking operation (see 12.5.1).

When a `delay_statement` appears in a `delay_alternative` of a `timed_entry_call` the selection of the entry call is attempted, regardless of the specified expiration time. When a `delay_statement` appears in a `select_alternative`, and a call is queued on one of the open entries, the selection of that entry call proceeds, regardless of the value of the delay expression.

#### Documentation Requirements

The implementation shall document the minimum value of the delay expression of a `delay_relative_statement` that causes the task to actually be blocked.

The implementation shall document the minimum difference between the value of the delay expression of a `delay_until_statement` and the value of `Real_Time.Clock`, that causes the task to actually be blocked.

#### Metrics

The implementation shall document the following metrics:

- An upper bound on the execution time, in processor clock cycles, of a `delay_relative_statement` whose requested value of the delay expression is less than or equal to zero.
- An upper bound on the execution time, in processor clock cycles, of a `delay_until_statement` whose requested value of the delay expression is less than or equal to the value of `Real_Time.Clock` at the time of executing the statement. Similarly, for `Calendar.Clock`.
- An upper bound on the *lateness* of a `delay_relative_statement`, for a positive value of the delay expression, in a situation where the task has sufficient priority to preempt the processor as soon as it becomes ready, and can proceed without waiting for any other execution resources. The upper bound is expressed as a function of the value of the delay expression. The *lateness* is obtained by subtracting the value of the delay expression from the *actual duration*. The *actual duration* is measured from a point immediately before a task executes the `delay_statement` to a point immediately after the task resumes execution following this statement.
- An upper bound on the *lateness* of a `delay_until_statement`, in a situation where the value of the requested expiration time is after the time the task begins executing the statement, the task has sufficient priority to preempt the processor as soon as it becomes ready, and it can proceed without waiting for any other execution resources. The upper bound is expressed as a function of the difference between the requested expiration time and the clock value at the time the statement begins execution. The *lateness* of a `delay_until_statement` is obtained by subtracting the requested expiration time from the real time that the task resumes execution following this statement.

## D.10 Synchronous Task Control

This subclause describes a language-defined private semaphore (suspension object), which can be used for *two-stage suspend* operations and as a simple building block for implementing higher-level queues.

#### Static Semantics

The following language-defined package exists:

```
package Ada.Synchronous_Task_Control
with Preelaborate, Nonblocking, Global => in out synchronized is
```

```

type Suspension_Object is limited private;
procedure Set_True(S : in out Suspension_Object);
procedure Set_False(S : in out Suspension_Object);
function Current_State(S : Suspension_Object) return Boolean;
procedure Suspend_Until_True(S : in out Suspension_Object)
    with Nonblocking => False;
private
    ... -- not specified by the language
end Ada.Synchronous_Task_Control;

```

The type `Suspension_Object` is a by-reference type.

The following language-defined package exists:

```

with Ada.Real_Time;
package Ada.Synchronous_Task_Control.EDF
    with Nonblocking, Global => in out synchronized is
    procedure Suspend_Until_True_And_Set_Deadline
        (S : in out Suspension_Object;
         TS : in Ada.Real_Time.Time_Span)
        with Nonblocking => False;
end Ada.Synchronous_Task_Control.EDF;

```

#### *Dynamic Semantics*

An object of the type `Suspension_Object` has two visible states: `True` and `False`. Upon initialization, its value is set to `False`.

The operations `Set_True` and `Set_False` are atomic with respect to each other and with respect to `Suspend_Until_True`; they set the state to `True` and `False` respectively.

`Current_State` returns the current state of the object.

The procedure `Suspend_Until_True` blocks the calling task until the state of the object `S` is `True`; at that point the task becomes ready and the state of the object becomes `False`.

`Program_Error` is raised upon calling `Suspend_Until_True` if another task is already waiting on that suspension object.

The procedure `Suspend_Until_True_And_Set_Deadline` blocks the calling task until the state of the object `S` is `True`; at that point the task becomes ready with a deadline of `Ada.Real_Time.Clock + TS`, and the state of the object becomes `False`. `Program_Error` is raised upon calling `Suspend_Until_True_And_Set_Deadline` if another task is already waiting on that suspension object.

#### *Bounded (Run-Time) Errors*

It is a bounded error for two or more tasks to call `Suspend_Until_True` on the same `Suspension_Object` concurrently. For each task, `Program_Error` can be raised, the task can proceed without suspending, or the task can suspend, potentially indefinitely. The state of the suspension object can end up either `True` or `False`.

#### *Implementation Requirements*

The implementation is required to allow the calling of `Set_False` and `Set_True` during any protected action, even one that has its ceiling priority in the `Interrupt_Priority` range.

NOTE More complex schemes, such as setting the deadline relative to when `Set_True` is called, can be programmed using a protected object.

### **D.10.1 Synchronous Barriers**

This subclause introduces a language-defined package to synchronously release a group of tasks after the number of blocked tasks reaches a specified count value.

*Static Semantics*

The following language-defined library package exists:

```

package Ada.Synchronous_Barriers
  with Preelaborate, Nonblocking, Global => in out synchronized is
    subtype Barrier_Limit is Positive range 1 .. implementation-defined;
    type Synchronous_Barrier (Release_Threshold : Barrier_Limit) is limited
private;
    procedure Wait_For_Release (The_Barrier : in out Synchronous_Barrier;
                               Notified    : out Boolean)
      with Nonblocking => False;
private
  -- not specified by the language
end Ada.Synchronous_Barriers;

```

Type Synchronous\_Barrier needs finalization (see 10.6).

*Dynamic Semantics*

Each call to Wait\_For\_Release blocks the calling task until the number of blocked tasks associated with the Synchronous\_Barrier object is equal to Release\_Threshold, at which time all blocked tasks are released. Notified is set to True for one of the released tasks, and set to False for all other released tasks.

The mechanism for determining which task sets Notified to True is implementation defined.

Once all tasks have been released, a Synchronous\_Barrier object may be reused to block another Release\_Threshold number of tasks.

As the first step of the finalization of a Synchronous\_Barrier, each blocked task is unblocked and Program\_Error is raised at the place of the call to Wait\_For\_Release.

It is implementation defined whether an abnormal task which is waiting on a Synchronous\_Barrier object is aborted immediately or aborted when the tasks waiting on the object are released.

*Bounded (Run-Time) Errors*

It is a bounded error to call Wait\_For\_Release on a Synchronous\_Barrier object after that object is finalized. If the error is detected, Program\_Error is raised. Otherwise, the call proceeds normally, which may leave a task blocked forever.

## D.11 Asynchronous Task Control

This subclause introduces a language-defined package to do asynchronous suspend/resume on tasks. It uses a conceptual *held* priority value to represent the task's *held* state.

*Static Semantics*

The following language-defined library package exists:

```

with Ada.Task_Identification;
package Ada.Asynchronous_Task_Control
  with Preelaborate, Nonblocking, Global => in out synchronized is
    procedure Hold(T : in Ada.Task_Identification.Task_Id);
    procedure Continue(T : in Ada.Task_Identification.Task_Id);
    function Is_Held(T : Ada.Task_Identification.Task_Id)
      return Boolean;
end Ada.Asynchronous_Task_Control;

```

*Dynamic Semantics*

After the Hold operation has been applied to a task, the task becomes *held*. For each processor there is a conceptual *idle* task, which is always ready. The base priority of the idle task is below System.Any\_

Priority'First. The *held priority* is a constant of the type Integer whose value is below the base priority of the idle task.

For any priority below System.Any\_Priority'First, the task dispatching policy is FIFO\_Within\_Priorities.

The Hold operation sets the state of T to held. For a held task, the active priority is reevaluated as if the base priority of the task were the held priority.

The Continue operation resets the state of T to not-held; its active priority is then reevaluated as determined by the task dispatching policy associated with its base priority.

The Is\_Held function returns True if and only if T is in the held state.

As part of these operations, a check is made that the task identified by T is not terminated. Tasking\_Error is raised if the check fails. Program\_Error is raised if the value of T is Null\_Task\_Id.

#### *Erroneous Execution*

If any operation in this package is called with a parameter T that specifies a task object that no longer exists, the execution of the program is erroneous.

#### *Implementation Permissions*

An implementation may omit support for Asynchronous\_Task\_Control if it is infeasible to support it in the target environment.

NOTE 1 It is a consequence of the priority rules that held tasks cannot be dispatched on any processor in a partition (unless they are inheriting priorities) since their priorities are defined to be below the priority of any idle task.

NOTE 2 The effect of calling Get\_Priority and Set\_Priority on a Held task is the same as on any other task.

NOTE 3 Calling Hold on a held task or Continue on a non-held task has no effect.

NOTE 4 The rules affecting queuing are derived from the above rules, in addition to the normal priority rules:

- When a held task is on the ready queue, its priority is so low as to never reach the top of the queue as long as there are other tasks on that queue.
- If a task is executing in a protected action, inside a rendezvous, or is inheriting priorities from other sources (e.g. when activated), it continues to execute until it is no longer executing the corresponding construct.
- If a task becomes held while waiting (as a caller) for a rendezvous to complete, the active priority of the accepting task is not affected.
- If a task becomes held while waiting in a selective\_accept, and an entry call is issued to one of the open entries, the corresponding accept\_alternative executes. When the rendezvous completes, the active priority of the accepting task is lowered to the held priority (unless it is still inheriting from other sources), and the task does not execute until another Continue.
- The same holds if the held task is the only task on a protected entry queue whose barrier becomes open. The corresponding entry body executes.

## **D.12 Other Optimizations and Determinism Rules**

This subclause describes various requirements for improving the response and determinism in a real-time system.

#### *Implementation Requirements*

If the implementation blocks interrupts (see C.3) not as a result of direct user action (e.g. an execution of a protected action) there shall be an upper bound on the duration of this blocking.

The implementation shall recognize entry-less protected types. The overhead of acquiring the execution resource of an object of such a type (see 12.5.1) shall be minimized. In particular, there should not be any overhead due to evaluating entry\_barrier conditions.

Unchecked\_Deallocation shall be supported for terminated tasks that are designated by access types, and shall have the effect of releasing all the storage associated with the task. This includes any run-time system or heap storage that has been implicitly allocated for the task by the implementation.

#### Documentation Requirements

The implementation shall document the upper bound on the duration of interrupt blocking caused by the implementation. If this is different for different interrupts or interrupt priority levels, it should be documented for each case.

#### Metrics

The implementation shall document the following metric:

- The overhead associated with obtaining a mutual-exclusive access to an entry-less protected object. This shall be measured in the following way:

For a protected object of the form:

```
protected Lock is
  procedure Set;
  function Read return Boolean;
private
  Flag : Boolean := False;
end Lock;

protected body Lock is
  procedure Set is
  begin
    Flag := True;
  end Set;
  function Read return Boolean
  begin
    return Flag;
  end Read;
end Lock;
```

The execution time, in processor clock cycles, of a call to Set. This shall be measured between the point just before issuing the call, and the point just after the call completes. The function Read shall be called later to verify that Set was indeed called (and not optimized away). The calling task shall have sufficiently high priority as to not be preempted during the measurement period. The protected object shall have sufficiently high ceiling priority to allow the task to call Set.

For a multiprocessor, if supported, the metric shall be reported for the case where no contention (on the execution resource) exists from tasks executing on other processors.

### D.13 The Ravenscar and Jorvik Profiles

This subclause defines the Ravenscar and Jorvik profiles.

#### Legality Rules

The *profile\_identifier* Ravenscar and *profile\_identifier* Jorvik are usage profiles (see 16.12). For usage profiles Ravenscar and Jorvik, there shall be no *profile\_pragma\_argument\_-* associations.

#### Static Semantics

The usage profile Ravenscar is equivalent to the following set of pragmas:



```

pragma Task_Dispatching_Policy (FIFO_Within_Priorities);
pragma Locking_Policy (Ceiling_Locking);
pragma Detect_Blocking;
pragma Restrictions (
    No_Abort_Statements,
    No_Dynamic_Attachment,
    No_Dynamic_CPU_Assignment,
    No_Dynamic_Priorities,
    No_Implicit_Heap_Allocations,
    No_Local_Protected_Objects,
    No_Local_Timing_Events,
    No_Protected_Type_Allocators,
    No_Relative_Delay,
    No_Requeue_Statements,
    No_Select_Statements,
    No_Specific_Termination_Handlers,
    No_Task_Allocators,
    No_Task_Hierarchy,
    No_Task_Termination,
    Simple_Barriers,
    Max_Entry_Queue_Length => 1,
    Max_Protected_Entries => 1,
    Max_Task_Entries => 0,
    No_Dependence => Ada.Asynchronous_Task_Control,
    No_Dependence => Ada.Calendar,
    No_Dependence => Ada.Execution_Time.Group_Budgets,
    No_Dependence => Ada.Execution_Time.Timers,
    No_Dependence => Ada.Synchronous_Barriers,
    No_Dependence => Ada.Task_Attributes,
    No_Dependence => System.Multiprocessors.Dispatching_Domains);

```

The usage profile Jorvik is equivalent to the following set of pragmas:

```

pragma Task_Dispatching_Policy (FIFO_Within_Priorities);
pragma Locking_Policy (Ceiling_Locking);
pragma Detect_Blocking;
pragma Restrictions (
    No_Abort_Statements,
    No_Dynamic_Attachment,
    No_Dynamic_CPU_Assignment,
    No_Dynamic_Priorities,
    No_Local_Protected_Objects,
    No_Local_Timing_Events,
    No_Protected_Type_Allocators,
    No_Requeue_Statements,
    No_Select_Statements,
    No_Specific_Termination_Handlers,
    No_Task_Allocators,
    No_Task_Hierarchy,
    No_Task_Termination,
    Pure_Barriers,
    Max_Task_Entries => 0,
    No_Dependence => Ada.Asynchronous_Task_Control,
    No_Dependence => Ada.Execution_Time.Group_Budgets,
    No_Dependence => Ada.Execution_Time.Timers,
    No_Dependence => Ada.Task_Attributes,
    No_Dependence => System.Multiprocessors.Dispatching_Domains);

```

#### *Implementation Advice*

On a multiprocessor system, an implementation should support a fully partitioned approach if one of these profiles is specified. Each processor should have separate and disjoint ready queues.

NOTE 1 For the Ravenscar profile, the effect of the restriction `Max_Entry_Queue_Length => 1` applies only to protected entry queues due to the accompanying restriction `Max_Task_Entries => 0`. The restriction `Max_Entry_Queue_Length` is not applied by the Jorvik profile.

NOTE 2 When the Ravenscar or Jorvik profile is in effect (via the effect of the No\_Dynamic\_CPU\_Assignment restriction), all of the tasks in the partition will execute on a single CPU unless the programmer explicitly uses aspect CPU to specify the CPU assignments for tasks. The use of multiple CPUs requires care, as many guarantees of single CPU scheduling no longer apply.

NOTE 3 Any unit that meets the requirements of the Ravenscar profile also meets the requirements of the Jorvik profile.

## D.14 Execution Time

This subclause describes a language-defined package to measure execution time.

### Static Semantics

The following language-defined library package exists:

```

with Ada.Task_Identification;
with Ada.Real_Time; use Ada.Real_Time;
package Ada.Execution_Time
  with Nonblocking, Global => in out synchronized is

  type CPU_Time is private;
  CPU_Time_First : constant CPU_Time;
  CPU_Time_Last  : constant CPU_Time;
  CPU_Time_Unit  : constant := implementation-defined-real-number;
  CPU_Tick       : constant Time_Span;

  function Clock
    (T : Ada.Task_Identification.Task_Id
     := Ada.Task_Identification.Current_Task)
    return CPU_Time;

  function "+" (Left : CPU_Time; Right : Time_Span) return CPU_Time;
  function "+" (Left : Time_Span; Right : CPU_Time) return CPU_Time;
  function "-" (Left : CPU_Time; Right : Time_Span) return CPU_Time;
  function "-" (Left : CPU_Time; Right : CPU_Time)   return Time_Span;

  function "<" (Left, Right : CPU_Time) return Boolean;
  function "<=" (Left, Right : CPU_Time) return Boolean;
  function ">" (Left, Right : CPU_Time) return Boolean;
  function ">=" (Left, Right : CPU_Time) return Boolean;

  procedure Split
    (T : in CPU_Time; SC : out Seconds_Count; TS : out Time_Span);

  function Time_Of (SC : Seconds_Count;
                  TS : Time_Span := Time_Span_Zero) return CPU_Time;

  Interrupt_Clocks_Supported : constant Boolean := implementation-defined;
  Separate_Interrupt_Clocks_Supported : constant Boolean :=
    implementation-defined;

  function Clock_For_Interrupts return CPU_Time;
private
  ... -- not specified by the language
end Ada.Execution_Time;

```

The *execution time* or CPU time of a given task is defined as the time spent by the system executing that task, including the time spent executing run-time or system services on its behalf. The mechanism used to measure execution time is implementation defined. The Boolean constant Interrupt\_Clocks\_Supported is set to True if the implementation separately accounts for the execution time of interrupt handlers. If it is set to False it is implementation defined which task, if any, is charged the execution time that is consumed by interrupt handlers. The Boolean constant Separate\_Interrupt\_Clocks\_Supported is set to True if the implementation separately accounts for the execution time of individual interrupt handlers (see D.14.3).

The type CPU\_Time represents the execution time of a task. The set of values of this type corresponds one-to-one with an implementation-defined range of mathematical integers.

The `CPU_Time` value  $I$  represents the half-open execution-time interval that starts with  $I \times \text{CPU\_Time\_Unit}$  and is limited by  $(I+1) \times \text{CPU\_Time\_Unit}$ , where `CPU_Time_Unit` is an implementation-defined real number. For each task, the execution time value is set to zero at the creation of the task.

`CPU_Time_First` and `CPU_Time_Last` are the smallest and largest values of the `CPU_Time` type, respectively.

The execution time value for the function `Clock_For_Interrupts` is initialized to zero.

#### *Dynamic Semantics*

`CPU_Time_Unit` is the smallest amount of execution time representable by the `CPU_Time` type; it is expressed in seconds. A *CPU clock tick* is an execution time interval during which the clock value (as observed by calling the `Clock` function) remains constant. `CPU_Tick` is the average length of such intervals.

The effects of the operators on `CPU_Time` and `Time_Span` are as for the operators defined for integer types.

The function `Clock` returns the current execution time of the task identified by `T`; `Tasking_Error` is raised if that task has terminated; `Program_Error` is raised if the value of `T` is `Task_Identification.Null_Task_Id`.

The effects of the `Split` and `Time_Of` operations are defined as follows, treating values of type `CPU_Time`, `Time_Span`, and `Seconds_Count` as mathematical integers. The effect of `Split` (`T`, `SC`, `TS`) is to set `SC` and `TS` to values such that  $T \times \text{CPU\_Time\_Unit} = SC \times 1.0 + TS \times \text{CPU\_Time\_Unit}$ , and  $0.0 \leq TS \times \text{CPU\_Time\_Unit} < 1.0$ . The value returned by `Time_Of`(`SC`, `TS`) is the execution-time value `T` such that  $T \times \text{CPU\_Time\_Unit} = SC \times 1.0 + TS \times \text{CPU\_Time\_Unit}$ .

The function `Clock_For_Interrupts` returns the total cumulative time spent executing within all interrupt handlers. This time is not allocated to any task execution time clock. If `Interrupt_Clocks_Supported` is set to `False` the function raises `Program_Error`.

#### *Erroneous Execution*

For a call of `Clock`, if the task identified by `T` no longer exists, the execution of the program is erroneous.

#### *Implementation Requirements*

The range of `CPU_Time` values shall be sufficient to uniquely represent the range of execution times from the task start-up to 50 years of execution time later. `CPU_Tick` shall be no greater than 1 millisecond.

#### *Documentation Requirements*

The implementation shall document the values of `CPU_Time_First`, `CPU_Time_Last`, `CPU_Time_Unit`, and `CPU_Tick`.

The implementation shall document the properties of the underlying mechanism used to measure execution times, such as the range of values supported and any relevant aspects of the underlying hardware or operating system facilities used.

#### *Metrics*

The implementation shall document the following metrics:

- An upper bound on the execution-time duration of a clock tick. This is a value  $D$  such that if  $t_1$  and  $t_2$  are any execution times of a given task such that  $t_1 < t_2$  and  $\text{Clock}_{t_1} = \text{Clock}_{t_2}$  then  $t_2 - t_1 \leq D$ .
- An upper bound on the size of a clock jump. A clock jump is the difference between two successive distinct values of an execution-time clock (as observed by calling the `Clock` function with the same `Task_Id`).
- An upper bound on the execution time of a call to the `Clock` function, in processor clock cycles.

- Upper bounds on the execution times of the operators of the type CPU\_Time, in processor clock cycles.

#### Implementation Permissions

Implementations targeted to machines with word size smaller than 32 bits may omit support for the full range and granularity of the CPU\_Time type.

#### Implementation Advice

When appropriate, implementations should provide configuration mechanisms to change the value of CPU\_Tick.

### D.14.1 Execution Time Timers

This subclause describes a language-defined package that provides a facility for calling a handler when a task has used a defined amount of CPU time.

#### Static Semantics

The following language-defined library package exists:

```
with System;
package Ada.Execution_Time.Timers
  with Nonblocking, Global => in out synchronized is
  type Timer (T : not null access constant
               Ada.Task_Identification.Task_Id) is
    tagged limited private;
  type Timer_Handler is
    access protected procedure (TM : in out Timer)
    with Nonblocking => False;
  Min_Handler_Ceiling : constant System.Any_Priority :=
    implementation-defined;
  procedure Set_Handler (TM : in out Timer;
                        In_Time : in Time_Span;
                        Handler : in Timer_Handler);
  procedure Set_Handler (TM : in out Timer;
                        At_Time : in CPU_Time;
                        Handler : in Timer_Handler);
  function Current_Handler (TM : Timer) return Timer_Handler;
  procedure Cancel_Handler (TM : in out Timer;
                          Cancelled : out Boolean);
  function Time_Remaining (TM : Timer) return Time_Span;
  Timer_Resource_Error : exception;
private
  ... -- not specified by the language
end Ada.Execution_Time.Timers;
```

The type Timer represents an execution-time event for a single task and is capable of detecting execution-time overruns. The access discriminant T identifies the task concerned. The type Timer needs finalization (see 10.6).

An object of type Timer is said to be *set* if it is associated with a nonnull value of type Timer\_Handler and *cleared* otherwise. All Timer objects are initially cleared.

The type Timer\_Handler identifies a protected procedure to be executed by the implementation when the timer expires. Such a protected procedure is called a *handler*.

#### Dynamic Semantics

When a Timer object is created, or upon the first call of a Set\_Handler procedure with the timer as parameter, the resources required to operate an execution-time timer based on the associated

execution-time clock are allocated and initialized. If this operation would exceed the available resources, `Timer_Resource_Error` is raised.

The procedures `Set_Handler` associate the handler `Handler` with the timer `TM`: if `Handler` is **null**, the timer is cleared; otherwise, it is set. The first procedure `Set_Handler` loads the timer `TM` with an interval specified by the `Time_Span` parameter. In this mode, the timer `TM` *expires* when the execution time of the task identified by `TM.T.all` has increased by `In_Time`; if `In_Time` is less than or equal to zero, the timer expires immediately. The second procedure `Set_Handler` loads the timer `TM` with the absolute value specified by `At_Time`. In this mode, the timer `TM` expires when the execution time of the task identified by `TM.T.all` reaches `At_Time`; if the value of `At_Time` has already been reached when `Set_Handler` is called, the timer expires immediately.

A call of a procedure `Set_Handler` for a timer that is already set replaces the handler and the (absolute or relative) execution time; if `Handler` is not **null**, the timer remains set.

When a timer expires, the associated handler is executed, passing the timer as parameter. The initial action of the execution of the handler is to clear the event.

The function `Current_Handler` returns the handler associated with the timer `TM` if that timer is set; otherwise, it returns **null**.

The procedure `Cancel_Handler` clears the timer if it is set. `Cancelled` is assigned `True` if the timer was set prior to it being cleared; otherwise, it is assigned `False`.

The function `Time_Remaining` returns the execution time interval that remains until the timer `TM` would expire, if that timer is set; otherwise, it returns `Time_Span.Zero`.

The constant `Min_Handler_Ceiling` is the minimum ceiling priority required for a protected object with a handler to ensure that no ceiling violation will occur when that handler is invoked.

As part of the finalization of an object of type `Timer`, the timer is cleared.

For all the subprograms defined in this package, `Tasking_Error` is raised if the task identified by `TM.T.all` has terminated, and `Program_Error` is raised if the value of `TM.T.all` is `Task_Identification.Null_Task_Id`.

An exception propagated from a handler invoked as part of the expiration of a timer has no effect.

#### *Erroneous Execution*

For a call of any of the subprograms defined in this package, if the task identified by `TM.T.all` no longer exists, the execution of the program is erroneous.

#### *Implementation Requirements*

For a given `Timer` object, the implementation shall perform the operations declared in this package atomically with respect to any of these operations on the same `Timer` object. The replacement of a handler by a call of `Set_Handler` shall be performed atomically with respect to the execution of the handler.

When an object of type `Timer` is finalized, the system resources used by the timer shall be deallocated.

#### *Implementation Permissions*

Implementations may limit the number of timers that can be defined for each task. If this limit is exceeded, then `Timer_Resource_Error` is raised.

NOTE A `Timer_Handler` can be associated with several `Timer` objects.

### **D.14.2 Group Execution Time Budgets**

This subclause describes a language-defined package to assign execution time budgets to groups of tasks.

## Static Semantics

The following language-defined library package exists:

```

with System;
with System.Multiprocessors;
package Ada.Execution_Time.Group_Budgets
  with Nonblocking, Global => in out synchronized is
  type Group_Budget (CPU : System.Multiprocessors.CPU :=
    System.Multiprocessors.CPU'First)
    is tagged limited private;
  type Group_Budget_Handler is access
    protected procedure (GB : in out Group_Budget)
    with Nonblocking => False;
  type Task_Array is array (Positive range <>) of
    Ada.Task_Identification.Task_Id;
  Min_Handler_Ceiling : constant System.Any_Priority :=
    implementation-defined;
  procedure Add_Task (GB : in out Group_Budget;
    T : in Ada.Task_Identification.Task_Id);
  procedure Remove_Task (GB: in out Group_Budget;
    T : in Ada.Task_Identification.Task_Id);
  function Is_Member (GB : Group_Budget;
    T : Ada.Task_Identification.Task_Id) return Boolean;
  function Is_A_Group_Member
    (T : Ada.Task_Identification.Task_Id) return Boolean;
  function Members (GB : Group_Budget) return Task_Array;
  procedure Replenish (GB : in out Group_Budget; To : in Time_Span);
  procedure Add (GB : in out Group_Budget; Interval : in Time_Span);
  function Budget_Has_Expired (GB : Group_Budget) return Boolean;
  function Budget_Remaining (GB : Group_Budget) return Time_Span;
  procedure Set_Handler (GB : in out Group_Budget;
    Handler : in Group_Budget_Handler);
  function Current_Handler (GB : Group_Budget)
    return Group_Budget_Handler;
  procedure Cancel_Handler (GB : in out Group_Budget;
    Cancelled : out Boolean);
  Group_Budget_Error : exception;
private
  -- not specified by the language
end Ada.Execution_Time.Group_Budgets;

```

The type `Group_Budget` represents an execution time budget to be used by a group of tasks. The type `Group_Budget` needs finalization (see 10.6). A task can belong to at most one group. Tasks of any priority can be added to a group.

An object of type `Group_Budget` has an associated nonnegative value of type `Time_Span` known as its *budget*, which is initially `Time_Span_Zero`. The type `Group_Budget_Handler` identifies a protected procedure to be executed by the implementation when the budget is *exhausted*, that is, reaches zero. Such a protected procedure is called a *handler*.

An object of type `Group_Budget` also includes a handler, which is a value of type `Group_Budget_Handler`. The handler of the object is said to be *set* if it is not null and *cleared* otherwise. The handler of all `Group_Budget` objects is initially cleared.

## Dynamic Semantics

The procedure `Add_Task` adds the task identified by `T` to the group `GB`; if that task is already a member of some other group, `Group_Budget_Error` is raised.



The procedure `Remove_Task` removes the task identified by `T` from the group `GB`; if that task is not a member of the group `GB`, `Group_Budget_Error` is raised. After successful execution of this procedure, the task is no longer a member of any group.

The function `Is_Member` returns `True` if the task identified by `T` is a member of the group `GB`; otherwise, it returns `False`.

The function `Is_A_Group_Member` returns `True` if the task identified by `T` is a member of some group; otherwise, it returns `False`.

The function `Members` returns an array of values of type `Task_Identification.Task_Id` identifying the members of the group `GB`. The order of the components of the array is unspecified.

The procedure `Replenish` loads the group budget `GB` with `To` as the `Time_Span` value. The exception `Group_Budget_Error` is raised if the `Time_Span` value `To` is nonpositive. Any execution on CPU of any member of the group of tasks results in the budget counting down, unless exhausted. When the budget becomes exhausted (reaches `Time_Span_Zero`), the associated handler is executed if the handler of group budget `GB` is set. Nevertheless, the tasks continue to execute.

The procedure `Add` modifies the budget of the group `GB`. A positive value for `Interval` increases the budget. A negative value for `Interval` reduces the budget, but never below `Time_Span_Zero`. A zero value for `Interval` has no effect. A call of procedure `Add` that results in the value of the budget going to `Time_Span_Zero` causes the associated handler to be executed if the handler of the group budget `GB` is set.

The function `Budget_Has_Expired` returns `True` if the budget of group `GB` is exhausted (equal to `Time_Span_Zero`); otherwise, it returns `False`.

The function `Budget_Remaining` returns the remaining budget for the group `GB`. If the budget is exhausted it returns `Time_Span_Zero`. This is the minimum value for a budget.

The procedure `Set_Handler` associates the handler `Handler` with the `Group_Budget` `GB`: if `Handler` is **null**, the handler of `Group_Budget` is cleared; otherwise, it is set.

A call of `Set_Handler` for a `Group_Budget` that already has a handler set replaces the handler; if `Handler` is not **null**, the handler for `Group_Budget` remains set.

The function `Current_Handler` returns the handler associated with the group budget `GB` if the handler for that group budget is set; otherwise, it returns **null**.

The procedure `Cancel_Handler` clears the handler for the group budget if it is set. `Cancelled` is assigned `True` if the handler for the group budget was set prior to it being cleared; otherwise, it is assigned `False`.

The constant `Min_Handler_Ceiling` is the minimum ceiling priority required for a protected object with a handler to ensure that no ceiling violation will occur when that handler is invoked.

The precision of the accounting of task execution time to a `Group_Budget` is the same as that defined for execution-time clocks from the parent package.

As part of the finalization of an object of type `Group_Budget` all member tasks are removed from the group identified by that object.

If a task is a member of a `Group_Budget` when it terminates, then as part of the finalization of the task it is removed from the group.

For all the operations defined in this package, `Tasking_Error` is raised if the task identified by `T` has terminated, and `Program_Error` is raised if the value of `T` is `Task_Identification.Null_Task_Id`.

An exception propagated from a handler invoked when the budget of a group of tasks becomes exhausted has no effect.

#### *Erroneous Execution*

For a call of any of the subprograms defined in this package, if the task identified by `T` no longer exists, the execution of the program is erroneous.

*Implementation Requirements*

For a given Group\_Budget object, the implementation shall perform the operations declared in this package atomically with respect to any of these operations on the same Group\_Budget object. The replacement of a handler, by a call of Set\_Handler, shall be performed atomically with respect to the execution of the handler.

NOTE 1 Clearing or setting of the handler of a group budget does not change the current value of the budget. Exhaustion or loading of a budget does not change whether the handler of the group budget is set or cleared.

NOTE 2 A Group\_Budget\_Handler can be associated with several Group\_Budget objects.

### D.14.3 Execution Time of Interrupt Handlers

This subclause describes a language-defined package to measure the execution time of interrupt handlers.

*Static Semantics*

The following language-defined library package exists:

```
with Ada.Interrupts;
package Ada.Execution_Time.Interrupts
  with Nonblocking, Global => in out synchronized is
    function Clock (Interrupt : Ada.Interrupts.Interrupt_Id)
      return CPU_Time;
    function Supported (Interrupt : Ada.Interrupts.Interrupt_Id)
      return Boolean;
  end Ada.Execution_Time.Interrupts;
```

The execution time or CPU time of a given interrupt Interrupt is defined as the time spent by the system executing interrupt handlers identified by Interrupt, including the time spent executing run-time or system services on its behalf. The mechanism used to measure execution time is implementation defined. Time spent executing interrupt handlers is distinct from time spent executing any task.

For each interrupt, the execution time value is initially set to zero.

*Dynamic Semantics*

The function Clock returns the current cumulative execution time of the interrupt identified by Interrupt. If Separate\_Interrupt\_Clocks\_Supported is set to False the function raises Program\_Error.

The function Supported returns True if the implementation is monitoring the execution time of the interrupt identified by Interrupt; otherwise, it returns False. For any Interrupt\_Id Interrupt for which Supported(Interrupt) returns False, the function Clock(Interrupt) will return a value equal to Ada.Execution\_Time.Time\_Of(0).

### D.15 Timing Events

This subclause describes a language-defined package to allow user-defined protected procedures to be executed at a specified time without the use of a task or a delay statement.

*Static Semantics*

The following language-defined library package exists:

```
package Ada.Real_Time.Timing_Events
  with Nonblocking, Global => in out synchronized is
    type Timing_Event is tagged limited private;
    type Timing_Event_Handler
      is access protected procedure (Event : in out Timing_Event)
      with Nonblocking => False;
```

```

procedure Set_Handler (Event    : in out Timing_Event;
                       At_Time  : in Time;
                       Handler   : in Timing_Event_Handler);
procedure Set_Handler (Event    : in out Timing_Event;
                       In_Time   : in Time_Span;
                       Handler   : in Timing_Event_Handler);
function Current_Handler (Event : Timing_Event)
    return Timing_Event_Handler;
procedure Cancel_Handler (Event      : in out Timing_Event;
                          Cancelled   : out Boolean);

function Time_Of_Event (Event : Timing_Event) return Time;

private
    ... -- not specified by the language
end Ada.Real_Time.Timing_Events;

```

The type `Timing_Event` represents a time in the future when an event is to occur. The type `Timing_Event` needs finalization (see 10.6).

An object of type `Timing_Event` is said to be *set* if it is associated with a nonnull value of type `Timing_Event_Handler` and *cleared* otherwise. All `Timing_Event` objects are initially cleared.

The type `Timing_Event_Handler` identifies a protected procedure to be executed by the implementation when the timing event occurs. Such a protected procedure is called a *handler*.

#### Dynamic Semantics

The procedures `Set_Handler` associate the handler `Handler` with the event `Event`: if `Handler` is **null**, the event is cleared; otherwise, it is set. The first procedure `Set_Handler` sets the execution time for the event to be `At_Time`. The second procedure `Set_Handler` sets the execution time for the event to be `Real_Time.Clock + In_Time`.

A call of a procedure `Set_Handler` for an event that is already set replaces the handler and the time of execution; if `Handler` is not **null**, the event remains set.

As soon as possible after the time set for the event, the handler is executed, passing the event as parameter. The handler is only executed if the timing event is in the set state at the time of execution. The initial action of the execution of the handler is to clear the event.

If the Ceiling\_Locking policy (see D.3) is in effect when a procedure `Set_Handler` is called, a check is made that the ceiling priority of `Handler.all` is `Interrupt_Priority'Last`. If the check fails, `Program_Error` is raised.

If a procedure `Set_Handler` is called with zero or negative `In_Time` or with `At_Time` indicating a time in the past, then the handler is executed as soon as possible after the completion of the call of `Set_Handler`.

The function `Current_Handler` returns the handler associated with the event `Event` if that event is set; otherwise, it returns **null**.

The procedure `Cancel_Handler` clears the event if it is set. `Cancelled` is assigned `True` if the event was set prior to it being cleared; otherwise, it is assigned `False`.

The function `Time_Of_Event` returns the time of the event if the event is set; otherwise, it returns `Real_Time.Time_First`.

As part of the finalization of an object of type `Timing_Event`, the `Timing_Event` is cleared.

If several timing events are set for the same time, they are executed in FIFO order of being set.

An exception propagated from a handler invoked by a timing event has no effect.

#### Implementation Requirements

For a given `Timing_Event` object, the implementation shall perform the operations declared in this package atomically with respect to any of these operations on the same `Timing_Event` object. The

replacement of a handler by a call of Set\_Handler shall be performed atomically with respect to the execution of the handler.

#### Metrics

The implementation shall document the following metric:

- An upper bound on the lateness of the execution of a handler. That is, the maximum time between the time specified for the event and when a handler is actually invoked assuming no other handler or task is executing during this interval.

#### Implementation Advice

The protected handler procedure should be executed directly by the real-time clock interrupt mechanism.

NOTE 1 Since a call of Set\_Handler is not a potentially blocking operation, it can be called from within a handler.

NOTE 2 A Timing\_Event\_Handler can be associated with several Timing\_Event objects.

## D.16 Multiprocessor Implementation

This subclause allows implementations on multiprocessor platforms to be configured.

#### Static Semantics

The following language-defined library package exists:

```
package System.Multiprocessors
  with Preelaborate, Nonblocking, Global => in out synchronized is
    type CPU_Range is range 0 .. implementation-defined;
    Not_A_Specific_CPU : constant CPU_Range := 0;
    subtype CPU is CPU_Range range 1 .. CPU_Range'Last;
    function Number_Of_CPUs return CPU;
end System.Multiprocessors;
```

A call of Number\_Of\_CPUs returns the number of processors available to the program. Within a given partition, each call on Number\_Of\_CPUs will return the same value.

For a task type (including the anonymous type of a single\_task\_declaration), protected type (including the anonymous type of a single\_protected\_declaration), or subprogram, the following language-defined representation aspect may be specified:

CPU      The aspect CPU is an expression, which shall be of type System.Multiprocessors.CPU\_Range.

#### Legality Rules

If the CPU aspect is specified for a subprogram, the expression shall be static.

The CPU aspect shall not be specified on a task or protected interface type.

#### Dynamic Semantics

The expression specified for the CPU aspect of a task or protected type is evaluated each time an object of the corresponding type is created (see 12.1 and 12.4). The CPU value is then associated with the object.

The CPU aspect has no effect if it is specified for a subprogram other than the main subprogram; the CPU value is not associated with any task.

The CPU value is associated with the environment task if the CPU aspect is specified for the main subprogram. If the CPU aspect is not specified for the main subprogram it is implementation defined on which processor the environment task executes.

For a task, the CPU value determines the processor on which the task will activate and execute; the task is said to be assigned to that processor. If the CPU value is `Not_A_Specific_CPU`, then the task is not assigned to a processor. A task without a CPU aspect specified will activate and execute on the same processor as its activating task if the activating task is assigned a processor. If the CPU value is not in the range of `System.Multiprocessors.CPU_Range` or is greater than `Number_Of_CPUs` the task is defined to have failed, and it becomes a completed task (see 12.2).

For a protected type, the CPU value determines the processor on which calling tasks will execute; the protected object is said to be assigned to that processor. If the CPU value is `Not_A_Specific_CPU`, then the protected object is not assigned to a processor. A call to a protected object that is assigned to a processor from a task that is not assigned a processor or is assigned a different processor raises `Program_Error`.

*Implementation Advice*

Starting a protected action on a protected object statically assigned to a processor should be implemented without busy-waiting.

### D.16.1 Multiprocessor Dispatching Domains

This subclause allows implementations on multiprocessor platforms to be partitioned into distinct dispatching domains during program startup.

*Static Semantics*

The following language-defined library package exists:

```
with Ada.Real_Time;
with Ada.Task_Identification;
package System.Multiprocessors.Dispatching_Domains
  with Nonblocking, Global => in out synchronized is
  Dispatching_Domain_Error : exception;
  type Dispatching_Domain (<>) is limited private;
  System_Dispatching_Domain : constant Dispatching_Domain;
  function Create (First : CPU; Last : CPU_Range) return Dispatching_Domain;
  function Get_First_CPU (Domain : Dispatching_Domain) return CPU;
  function Get_Last_CPU (Domain : Dispatching_Domain) return CPU_Range;
  type CPU_Set is array (CPU range <>) of Boolean;
  function Create (Set : CPU_Set) return Dispatching_Domain;
  function Get_CPU_Set (Domain : Dispatching_Domain) return CPU_Set;
  function Get_Dispatching_Domain
    (T : Ada.Task_Identification.Task_Id :=
      Ada.Task_Identification.Current_Task)
    return Dispatching_Domain;
  procedure Assign_Task
    (Domain : in out Dispatching_Domain;
     CPU : in CPU_Range := Not_A_Specific_CPU;
     T : in Ada.Task_Identification.Task_Id :=
       Ada.Task_Identification.Current_Task);
  procedure Set_CPU
    (CPU : in CPU_Range;
     T : in Ada.Task_Identification.Task_Id :=
       Ada.Task_Identification.Current_Task);
  function Get_CPU
    (T : Ada.Task_Identification.Task_Id :=
      Ada.Task_Identification.Current_Task)
    return CPU_Range;
```

```

procedure Delay_Until_And_Set_CPU
  (Delay_Until_Time : in Ada.Real_Time.Time; CPU : in CPU_Range);

private
  ... -- not specified by the language
end System.Multiprocessors.Dispatching_Domains;

```

A *dispatching domain* represents a set of processors on which a task may execute. Each processor is contained within exactly one dispatching domain. An object of type `Dispatching_Domain` identifies a dispatching domain. `System.Dispatching_Domain` identifies a domain that contains the processor or processors on which the environment task executes. At program start-up all processors are contained within this domain.

For a task type (including the anonymous type of a `single_task_declaration`), the following language-defined representation aspect may be specified:

#### Dispatching\_Domain

The value of aspect `Dispatching_Domain` is an expression, which shall be of type `Dispatching_Domains.Dispatching_Domain`. This aspect is the domain to which the task (or all objects of the task type) are assigned.

#### Legality Rules

The `Dispatching_Domain` aspect shall not be specified for a task interface.

#### Dynamic Semantics

The expression specified for the `Dispatching_Domain` aspect of a task type is evaluated each time an object of the task type is created (see 12.1). If the identified dispatching domain is empty, then `Dispatching_Domain_Error` is raised; otherwise the newly created task is assigned to the domain identified by the value of the expression.

If a task is not explicitly assigned to any domain, it is assigned to that of the activating task. A task always executes on some CPU in its domain.

If both the dispatching domain and CPU are specified for a task, and the CPU value is not contained within the set of processors for the domain (and is not `Not_A_Specific_CPU`), the activation of the task is defined to have failed, and it becomes a completed task (see 12.2).

The function `Create` with `First` and `Last` parameters creates and returns a dispatching domain containing all the processors in the range `First .. Last`. The function `Create` with a `Set` parameter creates and returns a dispatching domain containing the processors for which `Set(I)` is `True`. These processors are removed from `System.Dispatching_Domain`. A call of `Create` will raise `Dispatching_Domain_Error` if any designated processor is not currently in `System.Dispatching_Domain`, or if the system cannot support a distinct domain over the processors identified, or if a processor has a task assigned to it, or if the allocation would leave `System.Dispatching_Domain` empty. A call of `Create` will raise `Dispatching_Domain_Error` if the calling task is not the environment task, or if `Create` is called after the call to the main subprogram.

The function `Get_First_CPU` returns the first CPU in `Domain`, or `CPU'First` if `Domain` is empty; `Get_Last_CPU` returns the last CPU in `Domain`, or `CPU_Range'First` if `Domain` is empty. The function `Get_CPU_Set(D)` returns an array whose low bound is `Get_First_CPU(D)`, whose high bound is `Get_Last_CPU(D)`, with `True` values in the `Set` corresponding to the CPUs that are in the given `Domain`.

The function `Get_Dispatching_Domain` returns the dispatching domain on which the task is assigned.

A call of the procedure `Assign_Task` assigns task `T` to the CPU within the dispatching domain `Domain`. Task `T` can now execute only on CPU, unless CPU designates `Not_A_Specific_CPU` in which case it can execute on any processor within `Domain`. The exception `Dispatching_Domain_Error` is propagated if `Domain` is empty, `T` is already assigned to a dispatching domain other than `System.Dispatching_Domain`, or if CPU is not one of the processors of `Domain` (and is not `Not_A_Specific_CPU`). A call of `Assign_Task` is a task dispatching point for task `T` unless `T` is inside of a protected action, in which case the effect on task `T` is delayed until its next task dispatching point. If `T` is



the `Current_Task` the effect is immediate if `T` is not inside a protected action, otherwise the effect is as soon as practical. Assigning a task already assigned to `System_Dispatching_Domain` to that domain has no effect.

A call of procedure `Set_CPU` assigns task `T` to the CPU. Task `T` can now execute only on CPU, unless CPU designates `Not_A_Specific_CPU`, in which case it can execute on any processor within its dispatching domain. The exception `Dispatching_Domain_Error` is propagated if CPU is not one of the processors of the dispatching domain on which `T` is assigned (and is not `Not_A_Specific_CPU`). A call of `Set_CPU` is a task dispatching point for task `T` unless `T` is inside of a protected action, in which case the effect on task `T` is delayed until its next task dispatching point. If `T` is the `Current_Task` the effect is immediate if `T` is not inside a protected action, otherwise the effect is as soon as practical.

The function `Get_CPU` returns the processor assigned to task `T`, or `Not_A_Specific_CPU` if the task is not assigned to a processor.

A call of `Delay_Until_And_Set_CPU` delays the calling task for the designated time and then assigns the task to the specified processor when the delay expires. The exception `Dispatching_Domain_Error` is propagated if `P` is not one of the processors of the calling task's dispatching domain (and is not `Not_A_Specific_CPU`).

#### *Implementation Requirements*

The implementation shall perform the operations `Assign_Task`, `Set_CPU`, `Get_CPU` and `Delay_Until_And_Set_CPU` atomically with respect to any of these operations on the same dispatching domain, processor or task.

Any task that belongs to the system dispatching domain can execute on any CPU within that domain, unless the assignment of the task has been specified.

#### *Implementation Advice*

Each dispatching domain should have separate and disjoint ready queues.

#### *Documentation Requirements*

The implementation shall document the processor(s) on which the clock interrupt is handled and hence where delay queue and ready queue manipulations occur. For any `Interrupt_Id` whose handler can execute on more than one processor the implementation shall also document this set of processors.

#### *Implementation Permissions*

An implementation may limit the number of dispatching domains that can be created and raise `Dispatching_Domain_Error` if an attempt is made to exceed this number.

The implementation may defer the effect of a `Set_CPU` or an `Assign_Task` operation until the specified task leaves an ongoing parallel construct.

## Annex E (informative)

### Distributed Systems

This Annex defines facilities for supporting the implementation of distributed systems using multiple partitions working cooperatively as part of a single Ada program.

#### Post-Compilation Rules

A *distributed system* is an interconnection of one or more *processing nodes* (a system resource that has both computational and storage capabilities), and zero or more *storage nodes* (a system resource that has only storage capabilities, with the storage addressable by one or more processing nodes).

A *distributed program* comprises one or more partitions that execute independently (except when they communicate) in a distributed system.

The process of mapping the partitions of a program to the nodes in a distributed system is called *configuring the partitions of the program*.

#### Implementation Requirements

The implementation shall provide means for explicitly assigning library units to a partition and for the configuring and execution of a program consisting of multiple partitions on a distributed system; the means are implementation defined.

#### Implementation Permissions

An implementation may require that the set of processing nodes of a distributed system be homogeneous.

NOTE 1 The partitions comprising a program can be executed on differently configured distributed systems or on a nondistributed system without requiring recompilation. A distributed program can be partitioned differently from the same set of library units without recompilation. The resulting execution is semantically equivalent.

NOTE 2 A distributed program retains the same type safety as the equivalent single partition program.

#### E.1 Partitions

The partitions of a distributed program are classified as either active or passive.

#### Post-Compilation Rules

An *active partition* is a partition as defined in 13.2. A *passive partition* is a partition that has no thread of control of its own, whose library units are all preelaborated, and whose data and subprograms are accessible to one or more active partitions.

A passive partition shall include only `library_items` that either are declared pure or are shared passive (see 13.2.1 and E.2.1).

An active partition shall be configured on a processing node. A passive partition shall be configured either on a storage node or on a processing node.

The configuration of the partitions of a program onto a distributed system shall be consistent with the possibility for data references or calls between the partitions implied by their semantic dependences. Any reference to data or call of a subprogram across partitions is called a *remote access*.

*Dynamic Semantics*

A `library_item` is elaborated as part of the elaboration of each partition that includes it. If a normal library unit (see E.2) has state, then a separate copy of the state exists in each active partition that elaborates it. The state evolves independently in each such partition.

An active partition *terminates* when its environment task terminates. A partition becomes *inaccessible* if it terminates or if it is *aborted*. An active partition is aborted when its environment task is aborted. In addition, if a partition fails during its elaboration, it becomes inaccessible to other partitions. Other implementation-defined events can also result in a partition becoming inaccessible.

For a `prefix` `D` that denotes a library-level declaration, excepting a declaration of or within a declared-pure library unit, the following attribute is defined:

`D'Partition_Id`

Denotes a value of the type *universal\_integer* that identifies the partition in which `D` was elaborated. If `D` denotes the declaration of a remote call interface library unit (see E.2.3) the given partition is the one where the body of `D` was elaborated.

*Bounded (Run-Time) Errors*

It is a bounded error for there to be cyclic elaboration dependences between the active partitions of a single distributed program. The possible effects, in each of the partitions involved, are deadlock during elaboration, or the raising of `Communication_Error` or `Program_Error`.

*Implementation Permissions*

An implementation may allow multiple active or passive partitions to be configured on a single processing node, and multiple passive partitions to be configured on a single storage node. In these cases, the scheduling policies, treatment of priorities, and management of shared resources between these partitions are implementation defined.

An implementation may allow separate copies of an active partition to be configured on different processing nodes, and to provide appropriate interactions between the copies to present a consistent state of the partition to other active partitions.

In an implementation, the partitions of a distributed program may be loaded and elaborated at different times; they may be loaded and elaborated one at a time over an extended period of time. An implementation may provide facilities to abort and reload a partition during the execution of a distributed program.

An implementation may allow the state of some of the partitions of a distributed program to persist while other partitions of the program terminate and are later reinvoked.

NOTE 1 Library units are grouped into partitions after compile time, but before run time. At compile time, only the relevant library unit properties are identified using categorization aspects.

NOTE 2 The value returned by the `Partition_Id` attribute can be used as a parameter to implementation-provided subprograms in order to query information about the partition.

## E.2 Categorization of Library Units

Library units can be categorized according to the role they play in a distributed program. Certain restrictions are associated with each category to ensure that the semantics of a distributed program remain close to the semantics for a nondistributed program.

A *categorization aspect* restricts the declarations, child units, or semantic dependences of the library unit to which it applies. A *categorized library unit* is a library unit that has a categorization aspect that is `True`.

The aspects `Shared_Passive`, `Remote_Types`, and `Remote_Call_Interface` are categorization aspects. In addition, for the purposes of this Annex, the aspect `Pure` (see 13.2.1) is considered a categorization aspect.

A library package or generic library package is called a *shared passive* library unit if the `Shared_Passive` aspect of the unit is `True`. A library package or generic library package is called a *remote types* library unit if the `Remote_Types` aspect of the unit is `True`. A library unit is called a *remote call interface* if the `Remote_Call_Interface` aspect of the unit is `True`. A *normal library unit* is one for which no categorization aspect is `True`.

The various categories of library units and the associated restrictions are described in this and the following subclauses. The categories are related hierarchically in that the library units of one category can depend semantically only on library units of that category or an earlier one in the hierarchy, except that the body of a remote types or remote call interface library unit is unrestricted, the declaration of a remote types or remote call interface library unit may depend on preelaborated normal library units that are mentioned only in private with clauses, and all categories can depend on limited views.

The overall hierarchy (including declared pure) is as follows, with a lower-numbered category being “earlier in the hierarchy” in the sense of the previous paragraph:

1. Declared Pure
2. Shared Passive
3. Remote Types
4. Remote Call Interface
5. Normal (no restrictions)

Declared pure and shared passive library units are preelaborated. The declaration of a remote types or remote call interface library unit is required to be preelaborable.

### E.2.1 Shared Passive Library Units

A shared passive library unit is used for managing global data shared between active partitions. The restrictions on shared passive library units prevent the data or tasks of one active partition from being accessible to another active partition through references implicit in objects declared in the shared passive library unit.

#### Legality Rules

When the library unit aspect (see 16.1.1) `Shared_Passive` of a library unit is `True`, the library unit is a *shared passive library unit*. The following restrictions apply to such a library unit:

- it shall be preelaborable (see 13.2.1);
- it shall depend semantically only upon declared pure or shared passive `library_items`;
- it shall not contain a library-level declaration:
  - of an access type that designates a class-wide type;
  - of a type with a part that is of a task type;
  - of a type with a part that is of a protected type with `entry_declarations`; or
  - that contains a name that denotes a type declared within a declared-pure package, if that type has a part that is of an access type; for the purposes of this rule, the parts considered include those of the full views of any private types or private extensions.

Notwithstanding the definition of accessibility given in 6.10.2, the declaration of a library unit P1 is not accessible from within the declarative region of a shared passive library unit P2, unless the shared passive library unit P2 depends semantically on P1.

#### Static Semantics

A shared passive library unit is preelaborated.

#### Post-Compilation Rules

A shared passive library unit shall be assigned to at most one partition within a given program.

Notwithstanding the rule given in 13.2, a compilation unit in a given partition does not *need* (in the sense of 13.2) the shared passive library units on which it depends semantically to be included in that same partition; they will typically reside in separate passive partitions.

## E.2.2 Remote Types Library Units

A remote types library unit supports the definition of types intended for use in communication between active partitions.

### Legality Rules

When the library unit aspect (see 16.1.1) `Remote_Types` of a library unit is `True`, the library unit is a *remote types library unit*. The following restrictions apply to such a library unit:

- it shall be preelaborable;
- it shall depend semantically only on declared pure `library_items`, shared passive library units, other remote types library units, or preelaborated normal library units that are mentioned only in private with clauses;
- it shall not contain the declaration of any variable within the visible part of the library unit;
- the full view of each type declared in the visible part of the library unit that has any available stream attributes shall support external streaming (see 16.13.2).

A named access type declared in the visible part of a remote types or remote call interface library unit is called a *remote access type*. Such a type shall be:

- an access-to-subprogram type, or
- a general access type that designates a class-wide limited private type, a class-wide limited interface type, or a class-wide private extension all of whose ancestors are either private extensions, limited interface types, or limited private types.

A type that is derived from a remote access type is also a remote access type.

A remote access-to-subprogram type shall not be nonblocking (see 12.5).

The following restrictions apply to the use of a remote access-to-subprogram type:

- A value of a remote access-to-subprogram type shall be converted only to or from another (subtype-conformant) remote access-to-subprogram type;
- The `prefix` of an `Access_attribute_reference` that yields a value of a remote access-to-subprogram type shall statically denote a (subtype-conformant) remote subprogram.

The following restrictions apply to the use of a remote access-to-class-wide type:

- The primitive subprograms of the corresponding specific type shall only have access parameters if they are controlling formal parameters. The primitive functions of the corresponding specific type shall only have an access result if it is a controlling access result. Each noncontrolling formal parameter and noncontrolling result type shall support external streaming (see 16.13.2);
- The corresponding specific type shall not have a primitive procedure with the Synchronization aspect specified unless the `synchronization_kind` is `Optional` (see 12.5);
- A value of a remote access-to-class-wide type shall be explicitly converted only to another remote access-to-class-wide type;
- A value of a remote access-to-class-wide type shall be dereferenced (or implicitly converted to an anonymous access type) only as part of a dispatching call to a primitive operation of the designated type where the value designates a controlling operand of the call (see E.4);
- A controlling access result value for a primitive function with any controlling operands of the corresponding specific type shall either be explicitly converted to a remote access-to-class-wide type or be part of a dispatching call where the value designates a controlling operand of the call;

- The `Storage_Pool` attribute is not defined for a remote access-to-class-wide type; the expected type for an `allocator` shall not be a remote access-to-class-wide type. A remote access-to-class-wide type shall not be an actual parameter for a generic formal access type. The `Storage_Size` attribute of a remote access-to-class-wide type yields 0. The `Storage_Pool` and `Storage_Size` aspects shall not be specified for a remote access-to-class-wide type.

#### *Erroneous Execution*

Execution is erroneous if some operation (other than the initialization or finalization of the object) modifies the value of a constant object declared in the visible part of a remote types package.

NOTE 1 A remote types library unit is not necessarily pure, and the types it defines can include levels of indirection implemented by using access types. User-specified Read and Write attributes (see 16.13.2) provide for sending values of such a type between active partitions, with Write marshalling the representation, and Read unmarshalling any levels of indirection.

NOTE 2 The value of a remote access-to-class-wide limited interface can designate an object of a nonlimited type derived from the interface.

NOTE 3 A remote access type can designate a class-wide synchronized, protected, or task interface type.

### **E.2.3 Remote Call Interface Library Units**

A remote call interface library unit can be used as an interface for remote procedure calls (RPCs) (or remote function calls) between active partitions.

#### *Legality Rules*

When the library unit aspect (see 16.1.1) `Remote_Call_Interface` of a library unit is `True`, the library unit is a *remote call interface (RCI)*. A subprogram declared in the visible part of such a library unit, or declared by such a library unit, is called a *remote subprogram*.

The declaration of an RCI library unit shall be preelaborable (see 13.2.1), and shall depend semantically only upon declared pure `library_items`, shared passive library units, remote types library units, other remote call interface library units, or preelaborated normal library units that are mentioned only in private with clauses.

In addition, the following restrictions apply to an RCI library unit:

- its visible part shall not contain the declaration of a variable;
- its visible part shall not contain the declaration of a limited type;
- its visible part shall not contain a nested `generic_declaration`;
- it shall not be, nor shall its visible part contain, the declaration of a subprogram for which aspect `Inline` is `True`;
- it shall not be, nor shall its visible part contain, the declaration of a subprogram that is nonblocking (see 12.5);
- it shall not be, nor shall its visible part contain, a subprogram (or access-to-subprogram) declaration whose profile has a parameter or result of a type that does not support external streaming (see 16.13.2);
- any public child of the library unit shall be a remote call interface library unit.

Specification of a stream-oriented attribute is illegal in the specification of a remote call interface library unit. In addition to the places where Legality Rules normally apply (see 15.3), this rule applies also in the private part of an instance of a generic unit.

`All_Calls_Remote` is a library unit aspect. If the `All_Calls_Remote` aspect of a library unit is `True`, the library unit shall be a remote call interface.



*Post-Compilation Rules*

A remote call interface library unit shall be assigned to at most one partition of a given program. A remote call interface library unit whose parent is also an RCI library unit shall be assigned only to the same partition as its parent.

Notwithstanding the rule given in 13.2, a compilation unit in a given partition that semantically depends on the declaration of an RCI library unit, *needs* (in the sense of 13.2) only the declaration of the RCI library unit, not the body, to be included in that same partition. Therefore, the body of an RCI library unit is included only in the partition to which the RCI library unit is explicitly assigned.

*Implementation Requirements*

If aspect `All_Calls_Remote` is `True` for a given RCI library unit, then the implementation shall route any of the following calls through the Partition Communication Subsystem (PCS); see E.5:

- A direct call to a subprogram of the RCI unit from outside the declarative region of the unit;
- An indirect call through a remote access-to-subprogram value that designates a subprogram of the RCI unit;
- A dispatching call with a controlling operand designated by a remote access-to-class-wide value whose tag identifies a type declared in the RCI unit.

*Implementation Permissions*

An implementation may omit support for the `Remote_Call_Interface` aspect or the `All_Calls_Remote` aspect. Explicit message-based communication between active partitions can be supported as an alternative to RPC.

### E.3 Consistency of a Distributed System

This subclause defines attributes and rules associated with verifying the consistency of a distributed program.

*Static Semantics*

For a `prefix` `P` that statically denotes a program unit, the following attributes are defined:

`P'Version` Yields a value of the predefined type `String` that identifies the version of the compilation unit that contains the declaration of the program unit.

`P'Body_Version`

Yields a value of the predefined type `String` that identifies the version of the compilation unit that contains the body (but not any subunits) of the program unit.

The *version* of a compilation unit changes whenever the compilation unit changes in a semantically significant way. This document does not define the exact meaning of "semantically significant". It is unspecified whether there are other events (such as recompilation) that result in the version of a compilation unit changing.

If `P` is not a library unit, and `P` has no completion, then `P'Body_Version` returns the `Body_Version` of the innermost program unit enclosing the declaration of `P`. If `P` is a library unit, and `P` has no completion, then `P'Body_Version` returns a value that is different from `Body_Version` of any version of `P` that has a completion.

*Bounded (Run-Time) Errors*

In a distributed program, a library unit is *consistent* if the same version of its declaration is used throughout. It is a bounded error to elaborate a partition of a distributed program that contains a compilation unit that depends on a different version of the declaration of a shared passive or RCI library unit than that included in the partition to which the shared passive or RCI library unit was assigned. As a result of this error, `Program_Error` can be raised in one or both partitions during elaboration; in any case, the partitions become inaccessible to one another.

## E.4 Remote Subprogram Calls

A *remote subprogram call* is a subprogram call that invokes the execution of a subprogram in another (active) partition. The partition that originates the remote subprogram call is the *calling partition*, and the partition that executes the corresponding subprogram body is the *called partition*. Some remote procedure calls are allowed to return prior to the completion of subprogram execution. These are called *asynchronous remote procedure calls*.

There are three different ways of performing a remote subprogram call:

- As a direct call on a (remote) subprogram explicitly declared in a remote call interface;
- As an indirect call through a value of a remote access-to-subprogram type;
- As a dispatching call with a controlling operand designated by a value of a remote access-to-class-wide type.

The first way of calling corresponds to a *static* binding between the calling and the called partition. The latter two ways correspond to a *dynamic* binding between the calling and the called partition.

Remote types library units (see E.2.2) and remote call interface library units (see E.2.3) define the remote subprograms or remote access types used for remote subprogram calls.

### Legality Rules

In a dispatching call with two or more controlling operands, if one controlling operand is designated by a value of a remote access-to-class-wide type, then all shall be.

A nonblocking program unit shall not contain, other than within nested units with Nonblocking specified as statically False, a dispatching call with a controlling operand designated by a value of a remote access-to-class-wide type.

### Dynamic Semantics

For the execution of a remote subprogram call, subprogram parameters (and later the results, if any) are passed using a stream-oriented representation (see 16.13.1) which is suitable for transmission between partitions. This action is called *marshalling*. *Unmarshalling* is the reverse action of reconstructing the parameters or results from the stream-oriented representation. Marshalling is performed initially as part of the remote subprogram call in the calling partition; unmarshalling is done in the called partition. After the remote subprogram completes, marshalling is performed in the called partition, and finally unmarshalling is done in the calling partition.

A *calling stub* is the sequence of code that replaces the subprogram body of a remotely called subprogram in the calling partition. A *receiving stub* is the sequence of code (the “wrapper”) that receives a remote subprogram call on the called partition and invokes the appropriate subprogram body.

Remote subprogram calls are executed at most once, that is, if the subprogram call returns normally, then the called subprogram's body was executed exactly once.

The task executing a remote subprogram call blocks until the subprogram in the called partition returns, unless the call is asynchronous. For an asynchronous remote procedure call, the calling task can become ready before the procedure in the called partition returns.

If a construct containing a remote call is aborted, the remote subprogram call is *cancelled*. Whether the execution of the remote subprogram is immediately aborted as a result of the cancellation is implementation defined.

If a remote subprogram call is received by a called partition before the partition has completed its elaboration, the call is kept pending until the called partition completes its elaboration (unless the call is cancelled by the calling partition prior to that).

If an exception is propagated by a remotely called subprogram, and the call is not an asynchronous call, the corresponding exception is reraised at the point of the remote subprogram call. For an

asynchronous call, if the remote procedure call returns prior to the completion of the remotely called subprogram, any exception is lost.

The exception `Communication_Error` (see E.5) is raised if a remote call cannot be completed due to difficulties in communicating with the called partition.

All forms of remote subprogram calls are potentially blocking operations (see 12.5).

In a remote subprogram call with a formal parameter of a class-wide type, a check is made that the tag of the actual parameter identifies a tagged type declared in a declared-pure or shared passive library unit, or in the visible part of a remote types or remote call interface library unit. `Program_Error` is raised if this check fails. In a remote function call which returns a class-wide type, the same check is made on the function result.

In a dispatching call with two or more controlling operands that are designated by values of a remote access-to-class-wide type, a check is made (in addition to the normal `Tag_Check` — see 14.5) that all the remote access-to-class-wide values originated from `Access attribute_references` that were evaluated by tasks of the same active partition. `Constraint_Error` is raised if this check fails.

#### *Implementation Requirements*

The implementation of remote subprogram calls shall conform to the PCS interface as defined by the specification of the language-defined package `System.RPC` (see E.5). The calling stub shall use the `Do_RPC` procedure unless the remote procedure call is asynchronous in which case `Do_APC` shall be used. On the receiving side, the corresponding receiving stub shall be invoked by the RPC-receiver.

With respect to shared variables in shared passive library units, the execution of the corresponding subprogram body of a synchronous remote procedure call is considered to be part of the execution of the calling task. The execution of the corresponding subprogram body of an asynchronous remote procedure call proceeds in parallel with the calling task and does not signal the next action of the calling task (see 12.10).

NOTE 1 A given active partition can both make and receive remote subprogram calls. Thus, an active partition can act as both a client and a server.

NOTE 2 If a given exception is propagated by a remote subprogram call, but the exception does not exist in the calling partition, the exception can be handled by an **others** choice or be propagated to and handled by a third partition.

### **E.4.1 Asynchronous Remote Calls**

This subclause introduces the aspect `Asynchronous` which can be specified to allow a remote subprogram call to return prior to completion of the execution of the corresponding remote subprogram body.

#### *Static Semantics*

For a remote procedure, the following language-defined representation aspect may be specified:

**Asynchronous**

The type of aspect `Asynchronous` is Boolean. If directly specified, the `aspect_definition` shall be a static expression. If not specified, the aspect is False.

For a remote access type, the following language-defined representation aspect may be specified:

**Asynchronous**

The type of aspect `Asynchronous` is Boolean. If directly specified, the `aspect_definition` shall be a static expression. If not specified (including by inheritance), the aspect is False.

#### *Legality Rules*

If aspect `Asynchronous` is specified for a remote procedure, the formal parameters of the procedure shall all be of mode **in**.

If aspect Asynchronous is specified for a remote access type, the type shall be a remote access-to-class-wide type, or the type shall be a remote access-to-procedure type with the formal parameters of the designated profile of the type all of mode **in**.

#### Dynamic Semantics

A remote call is *asynchronous* if it is a call to a procedure, or a call through a value of an access-to-procedure type, for which aspect Asynchronous is True. In addition, if aspect Asynchronous is True for a remote access-to-class-wide type, then a dispatching call on a procedure with a controlling operand designated by a value of the type is asynchronous if the formal parameters of the procedure are all of mode **in**.

#### Implementation Requirements

Asynchronous remote procedure calls shall be implemented such that the corresponding body executes at most once as a result of the call.

### E.4.2 Example of Use of a Remote Access-to-Class-Wide Type

EXAMPLE *Example of using a remote access-to-class-wide type to achieve dynamic binding across active partitions:*

```

package Tapes
  with Pure is
    type Tape is abstract tagged limited private;
    -- Primitive dispatching operations where
    -- Tape is controlling operand
    procedure Copy (From, To : access Tape;
                    Num_Recs : in Natural) is abstract;
    procedure Rewind (T : access Tape) is abstract;
    -- More operations
  private
    type Tape is ...
  end Tapes;

with Tapes;
package Name_Server
  with Remote_Call_Interface is
    -- Dynamic binding to remote operations is achieved
    -- using the access-to-limited-class-wide type Tape_Ptr
    type Tape_Ptr is access all Tapes.Tape'Class;
    -- The following statically bound remote operations
    -- allow for a name-server capability in this example
    function Find (Name : String) return Tape_Ptr;
    procedure Register (Name : in String; T : in Tape_Ptr);
    procedure Remove (T : in Tape_Ptr);
    -- More operations
  end Name_Server;

package Tape_Driver is
  -- Declarations are not shown, they are irrelevant here
end Tape_Driver;

```

```

with Tapes, Name_Server;
package body Tape_Driver is
  type New_Tape is new Tapes.Tape with ...
  overriding
  procedure Rewind (T : access New_Tape);
  overriding
  procedure Copy
    (From, To : access New_Tape; Num_Recs: in Natural) is
  begin
    . . .
  end Copy;
  procedure Rewind (T : access New_Tape) is
  begin
    . . .
  end Rewind;
  -- Objects remotely accessible through use
  -- of Name_Server operations
  Tape1, Tape2 : aliased New_Tape;
begin
  Name_Server.Register ("NINE-TRACK", Tape1'Access);
  Name_Server.Register ("SEVEN-TRACK", Tape2'Access);
end Tape_Driver;

with Tapes, Name_Server;
-- Tape_Driver is not needed and thus not mentioned in the with_clause
procedure Tape_Client is
  T1, T2 : Name_Server.Tape_Ptr;
begin
  T1 := Name_Server.Find ("NINE-TRACK");
  T2 := Name_Server.Find ("SEVEN-TRACK");
  Tapes.Rewind (T1);
  Tapes.Rewind (T2);
  Tapes.Copy (T1, T2, 3);
end Tape_Client;

```

#### Discussion of the example:

- The package Tapes provides the necessary declarations of the type and its primitive operations.
- Name\_Server is a remote call interface package and is elaborated in a separate active partition to provide the necessary naming services (such as Register and Find) to the entire distributed program through remote subprogram calls.
- Tape\_Driver is a normal package that is elaborated in a partition configured on the processing node that is connected to the tape device(s). The abstract operations are overridden to support the locally declared tape devices (Tape1, Tape2). The package is not visible to its clients, but it exports the tape devices (as remote objects) through the services of the Name\_Server. This allows for tape devices to be dynamically added, removed or replaced without requiring the modification of the clients' code.
- The Tape\_Client procedure references only declarations in the Tapes and Name\_Server packages. Before using a tape for the first time, it will query the Name\_Server for a system-wide identity for that tape. From then on, it can use that identity to access the tape device.
- Values of remote access type Tape\_Ptr include the necessary information to complete the remote dispatching operations that result from dereferencing the controlling operands T1 and T2.

## E.5 Partition Communication Subsystem

The *Partition Communication Subsystem* (PCS) provides facilities for supporting communication between the active partitions of a distributed program. The package System.RPC is a language-defined interface to the PCS.

## Static Semantics

The following language-defined library package exists:

```

with Ada.Streams; -- see 16.13.1
package System.RPC
  with Nonblocking => False, Global => in out synchronized is

  type Partition_Id is range 0 .. implementation-defined;

  Communication_Error : exception;

  type Params_Stream_Type (
    Initial_Size : Ada.Streams.Stream_Element_Count) is new
    Ada.Streams.Root_Stream_Type with private;

  procedure Read(
    Stream : in out Params_Stream_Type;
    Item : out Ada.Streams.Stream_Element_Array;
    Last : out Ada.Streams.Stream_Element_Offset);

  procedure Write(
    Stream : in out Params_Stream_Type;
    Item : in Ada.Streams.Stream_Element_Array);

  -- Synchronous call
  procedure Do_RPC(
    Partition : in Partition_Id;
    Params : access Params_Stream_Type;
    Result : access Params_Stream_Type);

  -- Asynchronous call
  procedure Do_APC(
    Partition : in Partition_Id;
    Params : access Params_Stream_Type);

  -- The handler for incoming RPCs
  type RPC_Receiver is access procedure(
    Params : access Params_Stream_Type;
    Result : access Params_Stream_Type);

  procedure Establish_RPC_Receiver(
    Partition : in Partition_Id;
    Receiver : in RPC_Receiver);

private
  ... -- not specified by the language
end System.RPC;

```

A value of the type `Partition_Id` is used to identify a partition.

An object of the type `Params_Stream_Type` is used for identifying the particular remote subprogram that is being called, as well as marshalling and unmarshalling the parameters or result of a remote subprogram call, as part of sending them between partitions.

The `Read` and `Write` procedures override the corresponding abstract operations for the type `Params_Stream_Type`.

## Dynamic Semantics

The `Do_RPC` and `Do_APC` procedures send a message to the active partition identified by the `Partition` parameter.

After sending the message, `Do_RPC` blocks the calling task until a reply message comes back from the called partition or some error is detected by the underlying communication system in which case `Communication_Error` is raised at the point of the call to `Do_RPC`.

`Do_APC` operates in the same way as `Do_RPC` except that it is allowed to return immediately after sending the message.



Upon normal return, the stream designated by the Result parameter of Do\_RPC contains the reply message.

The procedure System.RPC.Establish\_RPC\_Receiver is called once, immediately after elaborating the library units of an active partition (that is, right after the *elaboration of the partition*) if the partition includes an RCI library unit, but prior to invoking the main subprogram, if any. The Partition parameter is the Partition\_Id of the active partition being elaborated. The Receiver parameter designates an implementation-provided procedure called the *RPC-receiver* which will handle all RPCs received by the partition from the PCS. Establish\_RPC\_Receiver saves a reference to the RPC-receiver; when a message is received at the called partition, the RPC-receiver is called with the Params stream containing the message. When the RPC-receiver returns, the contents of the stream designated by Result is placed in a message and sent back to the calling partition.

If a call on Do\_RPC is aborted, a cancellation message is sent to the called partition, to request that the execution of the remotely called subprogram be aborted.

#### *Implementation Requirements*

The implementation of the RPC-receiver shall be reentrant, thereby allowing concurrent calls on it from the PCS to service concurrent remote subprogram calls into the partition.

An implementation shall not restrict the replacement of the body of System.RPC. An implementation shall not restrict children of System.RPC. The related implementation permissions in the introduction to Annex A do not apply.

If the implementation of System.RPC is provided by the user, an implementation shall support remote subprogram calls as specified.

#### *Documentation Requirements*

The implementation of the PCS shall document whether the RPC-receiver is invoked from concurrent tasks. If there is an upper limit on the number of such tasks, this limit shall be documented as well, together with the mechanisms to configure it (if this is supported).

#### *Implementation Permissions*

The PCS is allowed to contain implementation-defined interfaces for explicit message passing, broadcasting, etc. Similarly, it is allowed to provide additional interfaces to query the state of some remote partition (given its partition ID) or of the PCS itself, to set timeouts and retry parameters, to get more detailed error status, etc. These additional interfaces should be provided in child packages of System.RPC.

A body for the package System.RPC is not required to be supplied by the implementation.

An alternative declaration is allowed for package System.RPC as long as it provides a set of operations that is substantially equivalent to the specification defined in this subclause.

#### *Implementation Advice*

Whenever possible, the PCS on the called partition should allow for multiple tasks to call the RPC-receiver with different messages and should allow them to block until the corresponding subprogram body returns.

The Write operation on a stream of type Params\_Stream\_Type should raise Storage\_Error if it runs out of space trying to write the Item into the stream.

NOTE The package System.RPC is not designed for direct calls by user programs. It is instead designed for use in the implementation of remote subprograms calls, being called by the calling stubs generated for a remote call interface library unit to initiate a remote call, and in turn calling back to an RPC-receiver that dispatches to the receiving stubs generated for the body of a remote call interface, to handle a remote call received from elsewhere.

## Annex F (informative)

### Information Systems

This Annex provides a set of facilities relevant to Information Systems programming. These fall into several categories:

- an attribute definition clause specifying `Machine_Radix` for a decimal subtype;
- the package `Decimal`, which declares a set of constants defining the implementation's capacity for decimal types, and a generic procedure for decimal division; and
- the child packages `Text_IO Editing`, `Wide_Text_IO Editing`, and `Wide_Wide_Text_IO Editing`, which support formatted and localized output of decimal data, based on "picture String" values.

Other relevant facilities can be found in 6.5.9, 6.5.10, 7.6, 16.3, A.10.9, Clause B.3, Clause B.4 and Annex G.

The character and string handling packages in Annex A are also relevant for Information Systems.

#### *Implementation Advice*

If COBOL (respectively, C) is widely supported in the target environment, implementations supporting the Information Systems Annex should provide the child package `Interfaces.COBOL` (respectively, `Interfaces.C`) specified in Annex B and should support a *convention\_identifier* of COBOL (respectively, C) for the Convention aspect (see Annex B), thus allowing Ada programs to interface with programs written in that language.

#### F.1 Machine\_Radix Attribute Definition Clause

##### *Static Semantics*

The representation attribute `Machine_Radix` may be specified for a decimal first subtype (see 6.5.9) via an *attribute\_definition\_clause*; the expression of such a clause shall be static, and its value shall be 2 or 10. A value of 2 implies a binary base range; a value of 10 implies a decimal base range.

##### *Implementation Advice*

Packed decimal should be used as the internal representation for objects of subtype S when `S'Machine_Radix = 10`.

EXAMPLE *Example of Machine\_Radix attribute definition clause:*

```
type Money is delta 0.01 digits 15;
for Money'Machine_Radix use 10;
```

#### F.2 The Package Decimal

##### *Static Semantics*

The library package `Decimal` has the following declaration:

```
package Ada.Decimal
with Pure is

Max_Scale : constant := implementation-defined;
Min_Scale : constant := implementation-defined;

Min_Delta : constant := 10.0**(-Max_Scale);
Max_Delta : constant := 10.0**(-Min_Scale);
```

```
Max_Decimal_Digits : constant := implementation-defined;
```

```
generic
```

```
  type Dividend_Type is delta <> digits <>;
```

```
  type Divisor_Type is delta <> digits <>;
```

```
  type Quotient_Type is delta <> digits <>;
```

```
  type Remainder_Type is delta <> digits <>;
```

```
  procedure Divide (Dividend : in Dividend_Type;
                    Divisor   : in Divisor_Type;
                    Quotient  : out Quotient_Type;
                    Remainder : out Remainder_Type)
```

```
    with Convention => Intrinsic;
```

```
end Ada.Decimal;
```

Max\_Scale is the largest N such that  $10.0^{*(-N)}$  is allowed as a decimal type's delta. Its type is *universal\_integer*.

Min\_Scale is the smallest N such that  $10.0^{*(-N)}$  is allowed as a decimal type's delta. Its type is *universal\_integer*.

Min\_Delta is the smallest value allowed for *delta* in a decimal\_fixed\_point\_definition. Its type is *universal\_real*.

Max\_Delta is the largest value allowed for *delta* in a decimal\_fixed\_point\_definition. Its type is *universal\_real*.

Max\_Decimal\_Digits is the largest value allowed for *digits* in a decimal\_fixed\_point\_definition. Its type is *universal\_integer*.

#### Static Semantics

The effect of Divide is as follows. The value of Quotient is Quotient\_Type(Dividend/Divisor). The value of Remainder is Remainder\_Type(Intermediate), where Intermediate is the difference between Dividend and the product of Divisor and Quotient; this result is computed exactly.

#### Implementation Requirements

Decimal.Max\_Decimal\_Digits shall be at least 18.

Decimal.Max\_Scale shall be at least 18.

Decimal.Min\_Scale shall be at most 0.

NOTE The effect of division yielding a quotient with control over rounding versus truncation is obtained by applying either the function attribute Quotient\_Type'Round or the conversion Quotient\_Type to the expression Dividend/Divisor.

### F.3 Edited Output for Decimal Types

The child packages Text\_IO Editing, Wide\_Text\_IO Editing, and Wide\_Wide\_Text\_IO Editing provide localizable formatted text output, known as *edited output*, for decimal types. An edited output string is a function of a numeric value, program-specifiable locale elements, and a format control value. The numeric value is of some decimal type. The locale elements are:

- the currency string;
- the digits group separator character;
- the radix mark character; and
- the fill character that replaces leading zeros of the numeric value.

For Text\_IO Editing the edited output and currency strings are of type String, and the locale characters are of type Character. For Wide\_Text\_IO Editing their types are Wide\_String and Wide\_Character, respectively. For Wide\_Wide\_Text\_IO Editing their types are Wide\_Wide\_String and Wide\_Wide\_Character, respectively.

Each of the locale elements has a default value that can be replaced or explicitly overridden.

A format-control value is of the private type *Picture*; it determines the composition of the edited output string and controls the form and placement of the sign, the position of the locale elements and the decimal digits, the presence or absence of a radix mark, suppression of leading zeros, and insertion of particular character values.

A *Picture* object is composed from a String value, known as a *picture String*, that serves as a template for the edited output string, and a Boolean value that controls whether a string of all space characters is produced when the number's value is zero. A picture String comprises a sequence of one- or two-Character symbols, each serving as a placeholder for a character or string at a corresponding position in the edited output string. The picture String symbols fall into several categories based on their effect on the edited output string:

|                   |      |     |     |     |      |      |
|-------------------|------|-----|-----|-----|------|------|
| Decimal Digit:    | '9'  |     |     |     |      |      |
| Radix Control:    | '.'  | 'V' |     |     |      |      |
| Sign Control:     | '+'  | '-' | '<' | '>' | "CR" | "DB" |
| Currency Control: | '\$' | '#' |     |     |      |      |
| Zero Suppression: | 'Z'  | '*' |     |     |      |      |
| Simple Insertion: | '_'  | 'B' | '0' | '/' |      |      |

The entries are not case-sensitive. Mixed- or lower-case forms for "CR" and "DB", and lower-case forms for 'V', 'Z', and 'B', have the same effect as the upper-case symbols shown.

An occurrence of a '9' Character in the picture String represents a decimal digit position in the edited output string.

A radix control Character in the picture String indicates the position of the radix mark in the edited output string: an actual character position for '.', or an assumed position for 'V'.

A sign control Character in the picture String affects the form of the sign in the edited output string. The '<' and '>' Character values indicate parentheses for negative values. A Character '+', '-', or '<' appears either singly, signifying a fixed-position sign in the edited output, or repeated, signifying a floating-position sign that is preceded by zero or more space characters and that replaces a leading 0.

A currency control Character in the picture String indicates an occurrence of the currency string in the edited output string. The '\$' Character represents the complete currency string; the '#' Character represents one character of the currency string. A '\$' Character appears either singly, indicating a fixed-position currency string in the edited output, or repeated, indicating a floating-position currency string that occurs in place of a leading 0. A sequence of '#' Character values indicates either a fixed- or floating-position currency string, depending on context.

A zero suppression Character in the picture String allows a leading zero to be replaced by either the space character (for 'Z') or the fill character (for '\*').

A simple insertion Character in the picture String represents, in general, either itself (if '/' or '0'), the space character (if 'B'), or the digits group separator character (if '\_'). In some contexts it is treated as part of a floating sign, floating currency, or zero suppression string.

An example of a picture String is "<###Z\_ZZ9.99>". If the currency string is "kr", the separator character is ',', and the radix mark is '.' then the edited output string values for the decimal values 32.10 and -5432.10 are "bbkrbbb32.10b" and "(bkr5,432.10)", respectively, where 'b' indicates the space character.

The generic packages `Text_IO.Decimal_IO`, `Wide_Text_IO.Decimal_IO`, and `Wide_Wide_Text_IO.Decimal_IO` (see A.10.9) provide text input and nonedited text output for decimal types.

NOTE A picture String is of type `Standard.String`, for all of `Text_IO Editing`, `Wide_Text_IO Editing`, and `Wide_Wide_Text_IO Editing`.

### F.3.1 Picture String Formation

A *well-formed picture String*, or simply *picture String*, is a String value that conforms to the syntactic rules, composition constraints, and character replication conventions specified in this subclause.

#### Dynamic Semantics

```

picture_string ::=
    fixed_$_picture_string
  | fixed_#_picture_string
  | floating_currency_picture_string
  | non_currency_picture_string

fixed_$_picture_string ::=
    [fixed_LHS_sign] fixed_$_char {direct_insertion} [zero_suppression]
    number [RHS_sign]

  | [fixed_LHS_sign {direct_insertion}] [zero_suppression]
    number fixed_$_char {direct_insertion} [RHS_sign]

  | floating_LHS_sign number fixed_$_char {direct_insertion} [RHS_sign]

  | [fixed_LHS_sign] fixed_$_char {direct_insertion}
    all_zero_suppression_number {direct_insertion} [RHS_sign]

  | [fixed_LHS_sign {direct_insertion}] all_zero_suppression_number {direct_i
insertion}
    fixed_$_char {direct_insertion} [RHS_sign]

  | all_sign_number {direct_insertion} fixed_$_char {direct_insertion} [RHS_si
gn]

fixed_#_picture_string ::=
    [fixed_LHS_sign] single_#_currency {direct_insertion}
    [zero_suppression] number [RHS_sign]

  | [fixed_LHS_sign] multiple_#_currency {direct_insertion}
    zero_suppression number [RHS_sign]

  | [fixed_LHS_sign {direct_insertion}] [zero_suppression]
    number fixed_#_currency {direct_insertion} [RHS_sign]

  | floating_LHS_sign number fixed_#_currency {direct_insertion} [RHS_sign]

  | [fixed_LHS_sign] single_#_currency {direct_insertion}
    all_zero_suppression_number {direct_insertion} [RHS_sign]

  | [fixed_LHS_sign] multiple_#_currency {direct_insertion}
    all_zero_suppression_number {direct_insertion} [RHS_sign]

  | [fixed_LHS_sign {direct_insertion}] all_zero_suppression_number {direct_i
insertion}
    fixed_#_currency {direct_insertion} [RHS_sign]

```

```
|all_sign_number {direct_insertion} fixed_#_currency {direct_insertion} [RH
S_sign]
```

```
floating_currency_picture_string ::=
    [fixed_LHS_sign] {direct_insertion} floating_$_currency number [RHS_sign]
| [fixed_LHS_sign] {direct_insertion} floating_#_currency number [RHS_sign]
| [fixed_LHS_sign] {direct_insertion} all_currency_number {direct_insertion}
[RHS_sign]
```

```
non_currency_picture_string ::=
    [fixed_LHS_sign {direct_insertion}] zero_suppression number [RHS_sign]
| [floating_LHS_sign] number [RHS_sign]
| [fixed_LHS_sign {direct_insertion}] all_zero_suppression number {direct_i
nsertion}
    [RHS_sign]
| all_sign_number {direct_insertion}
| fixed_LHS_sign direct_insertion {direct_insertion} number [RHS_sign]
```

```
fixed_LHS_sign ::= LHS_Sign
```

```
LHS_Sign ::= + | - | <
```

```
fixed_$_char ::= $
```

```
direct_insertion ::= simple_insertion
```

```
simple_insertion ::= _ | B | 0 | /
```

```
zero_suppression ::= Z {Z | context_sensitive_insertion} | fill_string
```

```
context_sensitive_insertion ::= simple_insertion
```

```
fill_string ::= * { * | context_sensitive_insertion }
```

```
number ::=
```

```
    fore_digits [radix [aft_digits] {direct_insertion}]
| radix aft_digits {direct_insertion}
```

```
fore_digits ::= 9 {9 | direct_insertion}
```

```
aft_digits ::= {9 | direct_insertion} 9
```

```
radix ::= . | V
```

```
RHS_sign ::= + | - | > | CR | DB
```

```
floating_LHS_sign ::=
```

```
    LHS_Sign {context_sensitive_insertion} LHS_Sign {LHS_Sign | context_sensit
ive_insertion}
```

```
single_#_currency ::= #
```



```

multiple_#_currency ::= ## {#}

fixed_#_currency ::= single_#_currency | multiple_#_currency

floating_$_currency ::=
  ${context_sensitive_insertion} $ { $ | context_sensitive_insertion }

floating_#_currency ::=
  #{context_sensitive_insertion} # { # | context_sensitive_insertion }

all_sign_number ::= all_sign_fore [radix [all_sign_aft]] [>]
all_sign_fore ::=
  sign_char { context_sensitive_insertion } sign_char { sign_char | context_sen
sitive_insertion }
all_sign_aft ::= { all_sign_aft_char } sign_char
all_sign_aft_char ::= sign_char | context_sensitive_insertion
sign_char ::= + | - | <

all_currency_number ::= all_currency_fore [radix [all_currency_aft]]
all_currency_fore ::=
  currency_char { context_sensitive_insertion }
  currency_char { currency_char | context_sensitive_insertion }
all_currency_aft ::= { all_currency_aft_char } currency_char
all_currency_aft_char ::= currency_char | context_sensitive_insertion
currency_char ::= $ | #

all_zero_suppression_number ::= all_zero_suppression_fore [radix [all_zero_
suppression_aft]]
all_zero_suppression_fore ::=
  zero_suppression_char { zero_suppression_char | context_sensitive_inserti
on }
all_zero_suppression_aft ::= { all_zero_suppression_aft_char } zero_suppress
ion_char
all_zero_suppression_aft_char ::= zero_suppression_char | context_sensitive
_insertion
zero_suppression_char ::= Z | *

```

The following composition constraints apply to a picture String:

- A floating\_LHS\_sign does not have occurrences of different LHS\_Sign Character values.
- If a picture String has '<' as fixed\_LHS\_sign, then it has '>' as RHS\_sign.
- If a picture String has '<' in a floating\_LHS\_sign or in an all\_sign\_number, then it has an occurrence of '>'.

- If a picture String has '+' or '-' as `fixed_LHS_sign`, in a `floating_LHS_sign`, or in an `all_sign_number`, then it has no `RHS_sign` or '>' character.
- An instance of `all_sign_number` does not have occurrences of different `sign_char` Character values.
- An instance of `all_currency_number` does not have occurrences of different `currency_char` Character values.
- An instance of `all_zero_suppression_number` does not have occurrences of different `zero_suppression_char` Character values, except for possible case differences between 'Z' and 'z'.

A *replicable Character* is a Character that, by the above rules, can occur in two consecutive positions in a picture String.

A *Character replication* is a String

`char & '(' & spaces & count_string & ')'`

where *char* is a replicable Character, *spaces* is a String (possibly empty) comprising only space Character values, and *count\_string* is a String of one or more decimal digit Character values. A Character replication in a picture String has the same effect as (and is said to be *equivalent to*) a String comprising *n* consecutive occurrences of *char*, where  $n = \text{Integer'Value}(\text{count\_string})$ .

An *expanded picture String* is a picture String containing no Character replications.

NOTE Although a sign to the left of the number can float, a sign to the right of the number is in a fixed position.

### F.3.2 Edited Output Generation

#### Dynamic Semantics

The contents of an edited output string are based on:

- A value, Item, of some decimal type Num,
- An expanded picture String Pic\_String,
- A Boolean value, Blank\_When\_Zero,
- A Currency string,
- A Fill character,
- A Separator character, and
- A Radix\_Mark character.

The combination of a True value for Blank\_When\_Zero and a '\*' character in Pic\_String is inconsistent; no edited output string is defined.

A layout error is identified in the rules below if leading nonzero digits of Item, character values of the Currency string, or a negative sign would be truncated; in such cases no edited output string is defined.

The edited output string has lower bound 1 and upper bound N where  $N = \text{Pic\_String'Length} + \text{Currency\_Length\_Adjustment} - \text{Radix\_Adjustment}$ , and

- $\text{Currency\_Length\_Adjustment} = \text{Currency'Length} - 1$  if there is some occurrence of '\$' in Pic\_String, and 0 otherwise.
- $\text{Radix\_Adjustment} = 1$  if there is an occurrence of 'V' or 'v' in Pic\_Str, and 0 otherwise.

Let the magnitude of Item be expressed as a base-10 number  $I_p \cdots I_1.F_1 \cdots F_q$ , called the *displayed magnitude* of Item, where:

- $q = \text{Min}(\text{Max}(\text{Num'Scale}, 0), n)$  where *n* is 0 if Pic\_String has no radix and is otherwise the number of digit positions following radix in Pic\_String, where a digit position corresponds to an occurrence of '9', a `zero_suppression_char` (for an `all_zero_suppression_number`), a

currency\_char (for an all\_currency\_number), or a sign\_char (for an all\_sign\_number).

—  $I_p \neq 0$  if  $p > 0$ .

If  $n < \text{Num'Scale}$ , then the above number is the result of rounding (away from 0 if exactly midway between values).

If Blank\_When\_Zero = True and the displayed magnitude of Item is zero, then the edited output string comprises all space character values. Otherwise, the picture String is treated as a sequence of instances of syntactic categories based on the rules in F.3.1, and the edited output string is the concatenation of string values derived from these categories according to the following mapping rules.

Table F.1 shows the mapping from a sign control symbol to a corresponding character or string in the edited output. In the columns showing the edited output, a lower-case 'b' represents the space character. If there is no sign control symbol but the value of Item is negative, a layout error occurs and no edited output string is produced.

**Table F.1 — Edited Output for Sign Control Symbols**

| Sign Control Symbol | Edited Output for Nonnegative Number | Edited Output for Negative Number |
|---------------------|--------------------------------------|-----------------------------------|
| '+'                 | '+'                                  | '_'                               |
| '_'                 | 'b'                                  | '_'                               |
| '<'                 | 'b'                                  | '('                               |
| '>'                 | 'b'                                  | ')'                               |
| "CR"                | "bb"                                 | "CR"                              |
| "DB"                | "bb"                                 | "DB"                              |

An instance of `fixed_LHS_sign` maps to a character as shown in Table F.1.

An instance of `fixed_$_char` maps to Currency.

An instance of `direct_insertion` maps to Separator if `direct_insertion = '_'`, and to the `direct_insertion` Character otherwise.

An instance of `number` maps to a string *integer\_part* & *radix\_part* & *fraction\_part* where:

— The string for *integer\_part* is obtained as follows:

1. Occurrences of '9' in `fore_digits` of `number` are replaced from right to left with the decimal digit character values for  $I_1, \dots, I_p$ , respectively.
2. Each occurrence of '9' in `fore_digits` to the left of the leftmost '9' replaced according to rule 1 is replaced with '0'.
3. If  $p$  exceeds the number of occurrences of '9' in `fore_digits` of `number`, then the excess leftmost digits are eligible for use in the mapping of an instance of `zero_suppression`, `floating_LHS_sign`, `floating_$_currency`, or `floating_#_currency` to the left of `number`; if there is no such instance, then a layout error occurs and no edited output string is produced.

— The *radix\_part* is:

- "" if `number` does not include a radix, if `radix = 'V'`, or if `radix = 'v'`
- Radix\_Mark if `number` includes '.' as radix

— The string for *fraction\_part* is obtained as follows:

1. Occurrences of '9' in *aft\_digits* of *number* are replaced from left to right with the decimal digit character values for  $F_1, \dots, F_n$ .
2. Each occurrence of '9' in *aft\_digits* to the right of the rightmost '9' replaced according to rule 1 is replaced by '0'.

An instance of *zero\_suppression* maps to the string obtained as follows:

1. The rightmost 'Z', 'z', or '\*' Character values are replaced with the excess digits (if any) from the *integer\_part* of the mapping of the *number* to the right of the *zero\_suppression* instance,
2. A *context\_sensitive\_insertion* Character is replaced as though it were a *direct\_insertion* Character, if it occurs to the right of some 'Z', 'z', or '\*' in *zero\_suppression* that has been mapped to an excess digit,
3. Each Character to the left of the leftmost Character replaced according to rule 1 above is replaced by:
  - the space character if the zero suppression Character is 'Z' or 'z', or
  - the Fill character if the zero suppression Character is '\*'.
4. A layout error occurs if some excess digits remain after all 'Z', 'z', and '\*' Character values in *zero\_suppression* have been replaced via rule 1; no edited output string is produced.

An instance of *RHS\_sign* maps to a character or string as shown in Table F.1.

An instance of *floating\_LHS\_sign* maps to the string obtained as follows.

1. Up to all but one of the rightmost *LHS\_Sign* Character values are replaced by the excess digits (if any) from the *integer\_part* of the mapping of the *number* to the right of the *floating\_LHS\_sign* instance.
2. The next Character to the left is replaced with the character given by the entry in Table F.1 corresponding to the *LHS\_Sign* Character.
3. A *context\_sensitive\_insertion* Character is replaced as though it were a *direct\_insertion* Character, if it occurs to the right of the leftmost *LHS\_Sign* character replaced according to rule 1.
4. Any other Character is replaced by the space character.
5. A layout error occurs if some excess digits remain after replacement via rule 1; no edited output string is produced.

An instance of *fixed\_#\_currency* maps to the *Currency* string with *n* space character values concatenated on the left (if the instance does not follow a *radix*) or on the right (if the instance does follow a *radix*), where *n* is the difference between the length of the *fixed\_#\_currency* instance and *Currency'Length*. A layout error occurs if *Currency'Length* exceeds the length of the *fixed\_#\_currency* instance; no edited output string is produced.

An instance of *floating\_\$\_currency* maps to the string obtained as follows:

1. Up to all but one of the rightmost '\$' Character values are replaced with the excess digits (if any) from the *integer\_part* of the mapping of the *number* to the right of the *floating\_\$\_currency* instance.
2. The next Character to the left is replaced by the *Currency* string.
3. A *context\_sensitive\_insertion* Character is replaced as though it were a *direct\_insertion* Character, if it occurs to the right of the leftmost '\$' Character replaced via rule 1.
4. Each other Character is replaced by the space character.

5. A layout error occurs if some excess digits remain after replacement by rule 1; no edited output string is produced.

An instance of `floating_#_currency` maps to the string obtained as follows:

1. Up to all but one of the rightmost '#' Character values are replaced with the excess digits (if any) from the *integer\_part* of the mapping of the number to the right of the `floating_#_currency` instance.
2. The substring whose last Character occurs at the position immediately preceding the leftmost Character replaced via rule 1, and whose length is `Currency'Length`, is replaced by the Currency string.
3. A `context_sensitive_insertion` Character is replaced as though it were a `direct_insertion` Character, if it occurs to the right of the leftmost '#' replaced via rule 1.
4. Any other Character is replaced by the space character.
5. A layout error occurs if some excess digits remain after replacement rule 1, or if there is no substring with the required length for replacement rule 2; no edited output string is produced.

An instance of `all_zero_suppression_number` maps to:

- a string of all spaces if the displayed magnitude of Item is zero, the `zero_suppression_char` is 'Z' or 'z', and the instance of `all_zero_suppression_number` does not have a radix at its last character position;
- a string containing the Fill character in each position except for the character (if any) corresponding to radix, if `zero_suppression_char = '*'` and the displayed magnitude of Item is zero;
- otherwise, the same result as if each `zero_suppression_char` in `all_zero_suppression_aft` were '9', interpreting the instance of `all_zero_suppression_number` as either zero suppression number (if a radix and `all_zero_suppression_aft` are present), or as zero suppression otherwise.

An instance of `all_sign_number` maps to:

- a string of all spaces if the displayed magnitude of Item is zero and the instance of `all_sign_number` does not have a radix at its last character position;
- otherwise, the same result as if each `sign_char` in `all_sign_number_aft` were '9', interpreting the instance of `all_sign_number` as either `floating_LHS_sign` number (if a radix and `all_sign_number_aft` are present), or as `floating_LHS_sign` otherwise.

An instance of `all_currency_number` maps to:

- a string of all spaces if the displayed magnitude of Item is zero and the instance of `all_currency_number` does not have a radix at its last character position;
- otherwise, the same result as if each `currency_char` in `all_currency_number_aft` were '9', interpreting the instance of `all_currency_number` as `floating_$_currency` number or `floating_#_currency` number (if a radix and `all_currency_number_aft` are present), or as `floating_$_currency` or `floating_#_currency` otherwise.

EXAMPLE Examples of use of edited output; in the result string values shown below, 'b' represents the space character:

|           |                             |                                                                                                        |
|-----------|-----------------------------|--------------------------------------------------------------------------------------------------------|
| Item:     | Picture and Result Strings: |                                                                                                        |
| 123456.78 | Picture:                    | "-###** *** **9.99"                                                                                    |
|           | Result:                     | "bbb\$***123,456.78"<br>"bbFF***123.456,78" (currency = "FF",<br>separator = '.',<br>radix mark = ',') |

|           |                             |                                                                              |
|-----------|-----------------------------|------------------------------------------------------------------------------|
| 123456.78 | Picture: "-\$** ***.99"     |                                                                              |
|           | Result: "b\$**123,456.78"   |                                                                              |
|           |                             | "bFF**123.456,78" (currency = "FF",<br>separator = '.',<br>radix mark = ',') |
| 0.0       | Picture: "-\$\$\$\$\$. \$"  |                                                                              |
|           | Result: "bbbbbbbbbb"        |                                                                              |
| 0.20      | Picture: "-\$\$\$\$\$. \$"  |                                                                              |
|           | Result: "bbbbbb\$.20"       |                                                                              |
| -1234.565 | Picture: "<<<< <<<. <<###>" |                                                                              |
|           | Result: "bb(1,234.57DMb)"   | (currency = "DM")                                                            |
| 12345.67  | Picture: "###_###_##.99"    |                                                                              |
|           | Result: "bbCHF12,345.67"    | (currency = "CHF")                                                           |

### F.3.3 The Package Text\_IO.Editing

The package Text\_IO.Editing provides a private type Picture with associated operations, and a generic package Decimal\_Output. An object of type Picture is composed from a well-formed picture String (see F.3.1) and a Boolean item indicating whether a zero numeric value will result in an edited output string of all space characters. The package Decimal\_Output contains edited output subprograms implementing the effects defined in F.3.2.

#### Static Semantics

The library package Text\_IO.Editing has the following declaration:

```

package Ada.Text_IO.Editing
  with Nonblocking, Global => in out synchronized is

  type Picture is private;

  function Valid (Pic_String : in String;
    Blank_When_Zero : in Boolean := False) return Boolean;

  function To_Picture (Pic_String : in String;
    Blank_When_Zero : in Boolean := False)
    return Picture;

  function Pic_String (Pic : in Picture) return String;
  function Blank_When_Zero (Pic : in Picture) return Boolean;

  Max_Picture_Length : constant := implementation_defined;
  Picture_Error : exception;

  Default_Currency : constant String := "$";
  Default_Fill : constant Character := '*';
  Default_Separator : constant Character := ',';
  Default_Radix_Mark : constant Character := '.';

  generic
    type Num is delta <> digits <>;
    Default_Currency : in String := Text_IO.Editing.Default_Currency;
    Default_Fill : in Character := Text_IO.Editing.Default_Fill;
    Default_Separator : in Character :=
      Text_IO.Editing.Default_Separator;
    Default_Radix_Mark : in Character :=
      Text_IO.Editing.Default_Radix_Mark;

  package Decimal_Output is
    function Length (Pic : in Picture;
      Currency : in String := Default_Currency)
      return Natural;

    function Valid (Item : in Num;
      Pic : in Picture;
      Currency : in String := Default_Currency)
      return Boolean;

```



```

function Image (Item      : in Num;
                Pic       : in Picture;
                Currency   : in String := Default_Currency;
                Fill       : in Character := Default_Fill;
                Separator   : in Character := Default_Separator;
                Radix_Mark  : in Character := Default_Radix_Mark)
    return String;

procedure Put (File      : in File_Type;
               Item      : in Num;
               Pic       : in Picture;
               Currency   : in String := Default_Currency;
               Fill       : in Character := Default_Fill;
               Separator   : in Character := Default_Separator;
               Radix_Mark  : in Character := Default_Radix_Mark)
    with Nonblocking => False;

procedure Put (Item      : in Num;
               Pic       : in Picture;
               Currency   : in String := Default_Currency;
               Fill       : in Character := Default_Fill;
               Separator   : in Character := Default_Separator;
               Radix_Mark  : in Character := Default_Radix_Mark)
    with Nonblocking => False;

procedure Put (To        : out String;
               Item      : in Num;
               Pic       : in Picture;
               Currency   : in String := Default_Currency;
               Fill       : in Character := Default_Fill;
               Separator   : in Character := Default_Separator;
               Radix_Mark  : in Character := Default_Radix_Mark);

end Decimal_Output;
private
    ... -- not specified by the language
end Ada.Text_IO.Editing;

```

The exception `Constraint_Error` is raised if the `Image` function or any of the `Put` procedures is invoked with a null string for `Currency`.

```

function Valid (Pic_String : in String;
                Blank_When_Zero : in Boolean := False) return Boolean;

```

`Valid` returns `True` if `Pic_String` is a well-formed picture String (see F.3.1) the length of whose expansion does not exceed `Max_Picture_Length`, and if either `Blank_When_Zero` is `False` or `Pic_String` contains no `'*`.

```

function To_Picture (Pic_String : in String;
                    Blank_When_Zero : in Boolean := False)
    return Picture;

```

`To_Picture` returns a result `Picture` such that the application of the function `Pic_String` to this result yields an expanded picture String equivalent to `Pic_String`, and such that `Blank_When_Zero` applied to the result `Picture` is the same value as the parameter `Blank_When_Zero`. `Picture_Error` is raised if not `Valid(Pic_String, Blank_When_Zero)`.

```

function Pic_String (Pic : in Picture) return String;

```

```

function Blank_When_Zero (Pic : in Picture) return Boolean;

```

If `Pic` is `To_Picture(String_Item, Boolean_Item)` for some `String_Item` and `Boolean_Item`, then:

- `Pic_String(Pic)` returns an expanded picture String equivalent to `String_Item` and with any lower-case letter replaced with its corresponding upper-case form, and
- `Blank_When_Zero(Pic)` returns `Boolean_Item`.

If `Pic_1` and `Pic_2` are objects of type `Picture`, then `"="(Pic_1, Pic_2)` is `True` when

- $\text{Pic\_String}(\text{Pic}_1) = \text{Pic\_String}(\text{Pic}_2)$ , and
- $\text{Blank\_When\_Zero}(\text{Pic}_1) = \text{Blank\_When\_Zero}(\text{Pic}_2)$ .

```
function Length (Pic      : in Picture;
                  Currency : in String := Default_Currency)
return Natural;
```

Length returns  $\text{Pic\_String}(\text{Pic})\text{'Length} + \text{Currency\_Length\_Adjustment} - \text{Radix\_Adjustment}$  where

- $\text{Currency\_Length\_Adjustment} =$ 
  - $\text{Currency}\text{'Length} - 1$  if there is some occurrence of '\$' in  $\text{Pic\_String}(\text{Pic})$ , and
  - 0 otherwise.
- $\text{Radix\_Adjustment} =$ 
  - 1 if there is an occurrence of 'V' or 'v' in  $\text{Pic\_Str}(\text{Pic})$ , and
  - 0 otherwise.

```
function Valid (Item      : in Num;
                Pic       : in Picture;
                Currency   : in String := Default_Currency)
return Boolean;
```

Valid returns True if  $\text{Image}(\text{Item}, \text{Pic}, \text{Currency})$  does not raise `Layout_Error`, and returns False otherwise.

```
function Image (Item      : in Num;
                Pic       : in Picture;
                Currency   : in String := Default_Currency;
                Fill       : in Character := Default_Fill;
                Separator   : in Character := Default_Separator;
                Radix_Mark  : in Character := Default_Radix_Mark)
return String;
```

Image returns the edited output String as defined in F.3.2 for  $\text{Item}$ ,  $\text{Pic\_String}(\text{Pic})$ ,  $\text{Blank\_When\_Zero}(\text{Pic})$ ,  $\text{Currency}$ ,  $\text{Fill}$ ,  $\text{Separator}$ , and  $\text{Radix\_Mark}$ . If these rules identify a layout error, then Image raises the exception `Layout_Error`.

```
procedure Put (File      : in File_Type;
               Item      : in Num;
               Pic       : in Picture;
               Currency   : in String := Default_Currency;
               Fill       : in Character := Default_Fill;
               Separator   : in Character := Default_Separator;
               Radix_Mark  : in Character := Default_Radix_Mark);
```

```
procedure Put (Item      : in Num;
               Pic       : in Picture;
               Currency   : in String := Default_Currency;
               Fill       : in Character := Default_Fill;
               Separator   : in Character := Default_Separator;
               Radix_Mark  : in Character := Default_Radix_Mark);
```

Each of these Put procedures outputs  $\text{Image}(\text{Item}, \text{Pic}, \text{Currency}, \text{Fill}, \text{Separator}, \text{Radix\_Mark})$  consistent with the conventions for Put for other real types in case of bounded line length (see A.10.6).

```
procedure Put (To        : out String;
               Item      : in Num;
               Pic       : in Picture;
               Currency   : in String := Default_Currency;
               Fill       : in Character := Default_Fill;
               Separator   : in Character := Default_Separator;
               Radix_Mark  : in Character := Default_Radix_Mark);
```

Put copies Image(Item, Pic, Currency, Fill, Separator, Radix\_Mark) to the given string, right justified. Otherwise, unassigned Character values in To are assigned the space character. If To'Length is less than the length of the string resulting from Image, then Layout\_Error is raised.

*Implementation Requirements*

Max\_Picture\_Length shall be at least 30. The implementation shall support currency strings of length up to at least 10, both for Default\_Currency in an instantiation of Decimal\_Output, and for Currency in an invocation of Image or any of the Put procedures.

### F.3.4 The Package Wide\_Text\_IO Editing

*Static Semantics*

The child package Wide\_Text\_IO Editing has the same contents as Text\_IO Editing, except that:

- each occurrence of Character is replaced by Wide\_Character,
- each occurrence of Text\_IO is replaced by Wide\_Text\_IO,
- the subtype of Default\_Currency is Wide\_String rather than String, and
- each occurrence of String in the generic package Decimal\_Output is replaced by Wide\_String.

NOTE Each of the functions Wide\_Text\_IO Editing.Valid, To\_Picture, and Pic\_String has String (versus Wide\_String) as its parameter or result subtype, since a picture String is not localizable.

### F.3.5 The Package Wide\_Wide\_Text\_IO Editing

*Static Semantics*

The child package Wide\_Wide\_Text\_IO Editing has the same contents as Text\_IO Editing, except that:

- each occurrence of Character is replaced by Wide\_Wide\_Character,
- each occurrence of Text\_IO is replaced by Wide\_Wide\_Text\_IO,
- the subtype of Default\_Currency is Wide\_Wide\_String rather than String, and
- each occurrence of String in the generic package Decimal\_Output is replaced by Wide\_Wide\_String.

NOTE Each of the functions Wide\_Wide\_Text\_IO Editing.Valid, To\_Picture, and Pic\_String has String (versus Wide\_Wide\_String) as its parameter or result subtype, since a picture String is not localizable.

## Annex G (informative)

### Numerics

The Numerics Annex specifies

- features for complex arithmetic, including complex I/O;
- a mode (“strict mode”), in which the predefined arithmetic operations of floating point and fixed point types and the functions and operations of various predefined packages have to provide guaranteed accuracy or conform to other numeric performance requirements, which the Numerics Annex also specifies;
- a mode (“relaxed mode”), in which there are no accuracy or other numeric performance requirements to be satisfied, as for implementations not conforming to the Numerics Annex;
- models of floating point and fixed point arithmetic on which the accuracy requirements of strict mode are based;
- the definitions of the model-oriented attributes of floating point types that apply in the strict mode; and
- features for the manipulation of real and complex vectors and matrices.

#### *Implementation Advice*

If Fortran (respectively, C) is widely supported in the target environment, implementations supporting the Numerics Annex should provide the child package Interfaces.Fortran (respectively, Interfaces.C) specified in Annex B and should support a *convention\_identifier* of Fortran (respectively, C) for the Convention aspect (see Annex B), thus allowing Ada programs to interface with programs written in that language.

### G.1 Complex Arithmetic

Types and arithmetic operations for complex arithmetic are provided in `Generic_Complex_Types`, which is defined in G.1.1. Implementation-defined approximations to the complex analogs of the mathematical functions known as the “elementary functions” are provided by the subprograms in `Generic_Complex_Elementary_Functions`, which is defined in G.1.2. Both of these library units are generic children of the predefined package Numerics (see A.5). Nongeneric equivalents of these generic packages for each of the predefined floating point types are also provided as children of Numerics.

#### G.1.1 Complex Types

##### *Static Semantics*

The generic library package `Numerics.Generic_Complex_Types` has the following declaration:

```
generic
  type Real is digits <>;
package Ada.Numerics.Generic_Complex_Types
  with Pure, Nonblocking is
    type Complex is
      record
        Re, Im : Real'Base;
      end record;
```

```

type Imaginary is private
  with Preelaborable_Initialization;

i : constant Imaginary;
j : constant Imaginary;

function Re (X : Complex) return Real'Base;
function Im (X : Complex) return Real'Base;
function Im (X : Imaginary) return Real'Base;

procedure Set_Re (X : in out Complex;
                  Re : in Real'Base);
procedure Set_Im (X : in out Complex;
                  Im : in Real'Base);
procedure Set_Im (X : out Imaginary;
                  Im : in Real'Base);

function Compose_From_Cartesian (Re, Im : Real'Base) return Complex;
function Compose_From_Cartesian (Re : Real'Base) return Complex;
function Compose_From_Cartesian (Im : Imaginary) return Complex;

function Modulus (X : Complex) return Real'Base;
function "abs" (Right : Complex) return Real'Base renames Modulus;

function Argument (X : Complex) return Real'Base;
function Argument (X : Complex;
                  Cycle : Real'Base) return Real'Base;

function Compose_From_Polar (Modulus, Argument : Real'Base)
return Complex;
function Compose_From_Polar (Modulus, Argument, Cycle : Real'Base)
return Complex;

function "+" (Right : Complex) return Complex;
function "-" (Right : Complex) return Complex;
function Conjugate (X : Complex) return Complex;

function "+" (Left, Right : Complex) return Complex;
function "-" (Left, Right : Complex) return Complex;
function "*" (Left, Right : Complex) return Complex;
function "/" (Left, Right : Complex) return Complex;

function "***" (Left : Complex; Right : Integer) return Complex;

function "+" (Right : Imaginary) return Imaginary;
function "-" (Right : Imaginary) return Imaginary;
function Conjugate (X : Imaginary) return Imaginary renames "-";
function "abs" (Right : Imaginary) return Real'Base;

function "+" (Left, Right : Imaginary) return Imaginary;
function "-" (Left, Right : Imaginary) return Imaginary;
function "*" (Left, Right : Imaginary) return Real'Base;
function "/" (Left, Right : Imaginary) return Real'Base;

function "***" (Left : Imaginary; Right : Integer) return Complex;

function "<" (Left, Right : Imaginary) return Boolean;
function "<=" (Left, Right : Imaginary) return Boolean;
function ">" (Left, Right : Imaginary) return Boolean;
function ">=" (Left, Right : Imaginary) return Boolean;

function "+" (Left : Complex; Right : Real'Base) return Complex;
function "+" (Left : Real'Base; Right : Complex) return Complex;
function "-" (Left : Complex; Right : Real'Base) return Complex;
function "-" (Left : Real'Base; Right : Complex) return Complex;
function "*" (Left : Complex; Right : Real'Base) return Complex;
function "*" (Left : Real'Base; Right : Complex) return Complex;
function "/" (Left : Complex; Right : Real'Base) return Complex;
function "/" (Left : Real'Base; Right : Complex) return Complex;

```

```

function "+" (Left : Complex; Right : Imaginary) return Complex;
function "+" (Left : Imaginary; Right : Complex) return Complex;
function "-" (Left : Complex; Right : Imaginary) return Complex;
function "-" (Left : Imaginary; Right : Complex) return Complex;
function "*" (Left : Complex; Right : Imaginary) return Complex;
function "*" (Left : Imaginary; Right : Complex) return Complex;
function "/" (Left : Complex; Right : Imaginary) return Complex;
function "/" (Left : Imaginary; Right : Complex) return Complex;

function "+" (Left : Imaginary; Right : Real'Base) return Complex;
function "+" (Left : Real'Base; Right : Imaginary) return Complex;
function "-" (Left : Imaginary; Right : Real'Base) return Complex;
function "-" (Left : Real'Base; Right : Imaginary) return Complex;
function "*" (Left : Imaginary; Right : Real'Base) return Imaginary;
function "*" (Left : Real'Base; Right : Imaginary) return Imaginary;
function "/" (Left : Imaginary; Right : Real'Base) return Imaginary;
function "/" (Left : Real'Base; Right : Imaginary) return Imaginary;

private

type Imaginary is new Real'Base;
i : constant Imaginary := 1.0;
j : constant Imaginary := 1.0;

end Ada.Numerics.Generic_Complex_Types;

```

The library package Numerics.Complex\_Types is declared pure and defines the same types, constants, and subprograms as Numerics.Generic\_Complex\_Types, except that the predefined type Float is systematically substituted for Real'Base throughout. Nongeneric equivalents of Numerics.Generic\_Complex\_Types for each of the other predefined floating point types are defined similarly, with the names Numerics.Short\_Complex\_Types, Numerics.Long\_Complex\_Types, etc.

Complex is a visible type with Cartesian components.

Imaginary is a private type; its full type is derived from Real'Base.

The arithmetic operations and the Re, Im, Modulus, Argument, and Conjugate functions have their usual mathematical meanings. When applied to a parameter of pure-imaginary type, the "imaginary-part" function Im yields the value of its parameter, as the corresponding real value. The remaining subprograms have the following meanings:

- The Set\_Re and Set\_Im procedures replace the designated component of a complex parameter with the given real value; applied to a parameter of pure-imaginary type, the Set\_Im procedure replaces the value of that parameter with the imaginary value corresponding to the given real value.
- The Compose\_From\_Cartesian function constructs a complex value from the given real and imaginary components. If only one component is given, the other component is implicitly zero.
- The Compose\_From\_Polar function constructs a complex value from the given modulus (radius) and argument (angle). When the value of the parameter Modulus is positive (resp., negative), the result is the complex value represented by the point in the complex plane lying at a distance from the origin given by the absolute value of Modulus and forming an angle measured counterclockwise from the positive (resp., negative) real axis given by the value of the parameter Argument.

When the Cycle parameter is specified, the result of the Argument function and the parameter Argument of the Compose\_From\_Polar function are measured in units such that a full cycle of revolution has the given value; otherwise, they are measured in radians.

The computed results of the mathematically multivalued functions are rendered single-valued by the following conventions, which are meant to imply the principal branch:

- The result of the Modulus function is nonnegative.
- The result of the Argument function is in the quadrant containing the point in the complex plane represented by the parameter X. This may be any quadrant (I through IV); thus, the range of the Argument function is approximately  $-\pi$  to  $\pi$  ( $-\text{Cycle}/2.0$  to  $\text{Cycle}/2.0$ , if the parameter Cycle is



specified). When the point represented by the parameter  $X$  lies on the negative real axis, the result approximates

- $\pi$  (resp.,  $-\pi$ ) when the sign of the imaginary component of  $X$  is positive (resp., negative), if `Real'Signed_Zeros` is `True`;
- $\pi$ , if `Real'Signed_Zeros` is `False`.
- Because a result lying on or near one of the axes may not be exactly representable, the approximation inherent in computing the result may place it in an adjacent quadrant, close to but on the wrong side of the axis.

#### Dynamic Semantics

The exception `Numerics.Argument_Error` is raised by the `Argument` and `Compose_From_Polar` functions with specified cycle, signaling a parameter value outside the domain of the corresponding mathematical function, when the value of the parameter `Cycle` is zero or negative.

The exception `Constraint_Error` is raised by the division operator when the value of the right operand is zero, and by the exponentiation operator when the value of the left operand is zero and the value of the exponent is negative, provided that `Real'Machine_Overflows` is `True`; when `Real'Machine_Overflows` is `False`, the result is unspecified. `Constraint_Error` can also be raised when a finite result overflows (see G.2.6).

#### Implementation Requirements

In the implementation of `Numerics.Generic_Complex_Types`, the range of intermediate values allowed during the calculation of a final result shall not be affected by any range constraint of the subtype `Real`.

In the following cases, evaluation of a complex arithmetic operation shall yield the *prescribed result*, provided that the preceding rules do not call for an exception to be raised:

- The results of the `Re`, `Im`, and `Compose_From_Cartesian` functions are exact.
- The real (resp., imaginary) component of the result of a binary addition operator that yields a result of complex type is exact when either of its operands is of pure-imaginary (resp., real) type.
- The real (resp., imaginary) component of the result of a binary subtraction operator that yields a result of complex type is exact when its right operand is of pure-imaginary (resp., real) type.
- The real component of the result of the `Conjugate` function for the complex type is exact.
- When the point in the complex plane represented by the parameter  $X$  lies on the nonnegative real axis, the `Argument` function yields a result of zero.
- When the value of the parameter `Modulus` is zero, the `Compose_From_Polar` function yields a result of zero.
- When the value of the parameter `Argument` is equal to a multiple of the quarter cycle, the result of the `Compose_From_Polar` function with specified cycle lies on one of the axes. In this case, one of its components is zero, and the other has the magnitude of the parameter `Modulus`.
- Exponentiation by a zero exponent yields the value one. Exponentiation by a unit exponent yields the value of the left operand. Exponentiation of the value one yields the value one. Exponentiation of the value zero yields the value zero, provided that the exponent is nonzero. When the left operand is of pure-imaginary type, one component of the result of the exponentiation operator is zero.

When the result, or a result component, of any operator of `Numerics.Generic_Complex_Types` has a mathematical definition in terms of a single arithmetic or relational operation, that result or result component exhibits the accuracy of the corresponding operation of the type `Real`.

Other accuracy requirements for the `Modulus`, `Argument`, and `Compose_From_Polar` functions, and accuracy requirements for the multiplication of a pair of complex operands or for division by a complex operand, all of which apply only in the strict mode, are given in G.2.6.

The sign of a zero result or zero result component yielded by a complex arithmetic operation or function is implementation defined when `Real'Signed_Zeros` is `True`.

*Implementation Permissions*

The nongeneric equivalent packages can be actual instantiations of the generic package for the appropriate predefined type, though that is not required.

Implementations may obtain the result of exponentiation of a complex or pure-imaginary operand by repeated complex multiplication, with arbitrary association of the factors and with a possible final complex reciprocation (when the exponent is negative). Implementations are also permitted to obtain the result of exponentiation of a complex operand, but not of a pure-imaginary operand, by converting the left operand to a polar representation; exponentiating the modulus by the given exponent; multiplying the argument by the given exponent; and reconverting to a Cartesian representation. Because of this implementation freedom, no accuracy requirement is imposed on complex exponentiation (except for the prescribed results given above, which apply regardless of the implementation method chosen).

*Implementation Advice*

Because the usual mathematical meaning of multiplication of a complex operand and a real operand is that of the scaling of both components of the former by the latter, an implementation should not perform this operation by first promoting the real operand to complex type and then performing a full complex multiplication. In systems that, in the future, support an Ada binding to ISO/IEC 60559:2020, the latter technique will not generate the required result when one of the components of the complex operand is infinite. (Explicit multiplication of the infinite component by the zero component obtained during promotion yields a NaN that propagates into the final result.) Analogous advice applies in the case of multiplication of a complex operand and a pure-imaginary operand, and in the case of division of a complex operand by a real or pure-imaginary operand.

Likewise, because the usual mathematical meaning of addition of a complex operand and a real operand is that the imaginary operand remains unchanged, an implementation should not perform this operation by first promoting the real operand to complex type and then performing a full complex addition. In implementations in which the `Signed_Zeros` attribute of the component type is `True` (and which therefore conform to ISO/IEC 60559:2020 in regard to the handling of the sign of zero in predefined arithmetic operations), the latter technique will not generate the required result when the imaginary component of the complex operand is a negatively signed zero. (Explicit addition of the negative zero to the zero obtained during promotion yields a positive zero.) Analogous advice applies in the case of addition of a complex operand and a pure-imaginary operand, and in the case of subtraction of a complex operand and a real or pure-imaginary operand.

Implementations in which `Real'Signed_Zeros` is `True` should attempt to provide a rational treatment of the signs of zero results and result components. As one example, the result of the `Argument` function should have the sign of the imaginary component of the parameter `X` when the point represented by that parameter lies on the positive real axis; as another, the sign of the imaginary component of the `Compose_From_Polar` function should be the same as (resp., the opposite of) that of the `Argument` parameter when that parameter has a value of zero and the `Modulus` parameter has a nonnegative (resp., negative) value.

## G.1.2 Complex Elementary Functions

*Static Semantics*

The generic library package `Numerics.Generic_Complex_Elementary_Functions` has the following declaration:

```

with Ada.Numerics.Generic_Complex_Types;
generic
  with package Complex_Types is
    new Ada.Numerics.Generic_Complex_Types (<>);
  use Complex_Types;
package Ada.Numerics.Generic_Complex_Elementary_Functions
  with Pure, Nonblocking is

  function Sqrt (X : Complex) return Complex;
  function Log (X : Complex) return Complex;
  function Exp (X : Complex) return Complex;
  function Exp (X : Imaginary) return Complex;
  function "***" (Left : Complex; Right : Complex) return Complex;
  function "***" (Left : Complex; Right : Real'Base) return Complex;
  function "***" (Left : Real'Base; Right : Complex) return Complex;

  function Sin (X : Complex) return Complex;
  function Cos (X : Complex) return Complex;
  function Tan (X : Complex) return Complex;
  function Cot (X : Complex) return Complex;

  function Arcsin (X : Complex) return Complex;
  function Arccos (X : Complex) return Complex;
  function Arctan (X : Complex) return Complex;
  function Arccot (X : Complex) return Complex;

  function Sinh (X : Complex) return Complex;
  function Cosh (X : Complex) return Complex;
  function Tanh (X : Complex) return Complex;
  function Coth (X : Complex) return Complex;

  function Arcsinh (X : Complex) return Complex;
  function Arccosh (X : Complex) return Complex;
  function Arctanh (X : Complex) return Complex;
  function Arccoth (X : Complex) return Complex;

end Ada.Numerics.Generic_Complex_Elementary_Functions;

```

The library package `Numerics.Complex_Elementary_Functions` is declared pure and defines the same subprograms as `Numerics.Generic_Complex_Elementary_Functions`, except that the predefined type `Float` is systematically substituted for `Real'Base`, and the `Complex` and `Imaginary` types exported by `Numerics.Complex_Types` are systematically substituted for `Complex` and `Imaginary`, throughout. Nongeneric equivalents of `Numerics.Generic_Complex_Elementary_Functions` corresponding to each of the other predefined floating point types are defined similarly, with the names `Numerics.Short_Complex_Elementary_Functions`, `Numerics.Long_Complex_Elementary_Functions`, etc.

The overloading of the `Exp` function for the pure-imaginary type is provided to give the user an alternate way to compose a complex value from a given modulus and argument. In addition to `Compose_From_Polar(Rho, Theta)` (see G.1.1), the programmer may write `Rho * Exp(i * Theta)`.

The imaginary (resp., real) component of the parameter `X` of the forward hyperbolic (resp., trigonometric) functions and of the `Exp` function (and the parameter `X`, itself, in the case of the overloading of the `Exp` function for the pure-imaginary type) represents an angle measured in radians, as does the imaginary (resp., real) component of the result of the `Log` and inverse hyperbolic (resp., trigonometric) functions.

The functions have their usual mathematical meanings. However, the arbitrariness inherent in the placement of branch cuts, across which some of the complex elementary functions exhibit discontinuities, is eliminated by the following conventions:

- The imaginary component of the result of the `Sqrt` and `Log` functions is discontinuous as the parameter `X` crosses the negative real axis.
- The result of the exponentiation operator when the left operand is of complex type is discontinuous as that operand crosses the negative real axis.

- The imaginary component of the result of the Arcsin, Arccos, and Arctanh functions is discontinuous as the parameter  $X$  crosses the real axis to the left of  $-1.0$  or the right of  $1.0$ .
- The real component of the result of the Arctan and Arcsinh functions is discontinuous as the parameter  $X$  crosses the imaginary axis below  $-i$  or above  $i$ .
- The real component of the result of the Arccot function is discontinuous as the parameter  $X$  crosses the imaginary axis below  $-i$  or above  $i$ .
- The imaginary component of the Arccosh function is discontinuous as the parameter  $X$  crosses the real axis to the left of  $1.0$ .
- The imaginary component of the result of the Arccoth function is discontinuous as the parameter  $X$  crosses the real axis between  $-1.0$  and  $1.0$ .

The computed results of the mathematically multivalued functions are rendered single-valued by the following conventions, which are meant to imply that the principal branch is an analytic continuation of the corresponding real-valued function in `Numerics.Generic_Elementary_Functions`. (For Arctan and Arccot, the single-argument function in question is that obtained from the two-argument version by fixing the second argument to be its default value.)

- The real component of the result of the Sqrt and Arccosh functions is nonnegative.
- The same convention applies to the imaginary component of the result of the Log function as applies to the result of the natural-cycle version of the Argument function of `Numerics.Generic_Complex_Types` (see G.1.1).
- The range of the real (resp., imaginary) component of the result of the Arcsin and Arctan (resp., Arcsinh and Arctanh) functions is approximately  $-\pi/2.0$  to  $\pi/2.0$ .
- The real (resp., imaginary) component of the result of the Arccos and Arccot (resp., Arccoth) functions ranges from  $0.0$  to approximately  $\pi$ .
- The range of the imaginary component of the result of the Arccosh function is approximately  $-\pi$  to  $\pi$ .

In addition, the exponentiation operator inherits the single-valuedness of the Log function.

#### *Dynamic Semantics*

The exception `Numerics.Argument_Error` is raised by the exponentiation operator, signaling a parameter value outside the domain of the corresponding mathematical function, when the value of the left operand is zero and the real component of the exponent (or the exponent itself, when it is of real type) is zero.

The exception `Constraint_Error` is raised, signaling a pole of the mathematical function (analogous to dividing by zero), in the following cases, provided that `Complex_Types.Real'Machine_Overflows` is `True`:

- by the Log, Cot, and Coth functions, when the value of the parameter  $X$  is zero;
- by the exponentiation operator, when the value of the left operand is zero and the real component of the exponent (or the exponent itself, when it is of real type) is negative;
- by the Arctan and Arccot functions, when the value of the parameter  $X$  is  $\pm i$ ;
- by the Arctanh and Arccoth functions, when the value of the parameter  $X$  is  $\pm 1.0$ .

`Constraint_Error` can also be raised when a finite result overflows (see G.2.6); this may occur for parameter values sufficiently *near* poles, and, in the case of some of the functions, for parameter values having components of sufficiently large magnitude. When `Complex_Types.Real'Machine_Overflows` is `False`, the result at poles is unspecified.

#### *Implementation Requirements*

In the implementation of `Numerics.Generic_Complex_Elementary_Functions`, the range of intermediate values allowed during the calculation of a final result shall not be affected by any range constraint of the subtype `Complex_Types.Real`.

In the following cases, evaluation of a complex elementary function shall yield the *prescribed result* (or a result having the prescribed component), provided that the preceding rules do not call for an exception to be raised:

- When the parameter  $X$  has the value zero, the Sqrt, Sin, Arcsin, Tan, Arctan, Sinh, Arcsinh, Tanh, and Arctanh functions yield a result of zero; the Exp, Cos, and Cosh functions yield a result of one; the Arccos and Arccot functions yield a real result; and the Arccoth function yields an imaginary result.
- When the parameter  $X$  has the value one, the Sqrt function yields a result of one; the Log, Arccos, and Arccosh functions yield a result of zero; and the Arcsin function yields a real result.
- When the parameter  $X$  has the value  $-1.0$ , the Sqrt function yields the result
  - $i$  (resp.,  $-i$ ), when the sign of the imaginary component of  $X$  is positive (resp., negative), if Complex\_Types.Real'Signed\_Zeros is True;
  - $i$ , if Complex\_Types.Real'Signed\_Zeros is False;
- When the parameter  $X$  has the value  $-1.0$ , the Log function yields an imaginary result; and the Arcsin and Arccos functions yield a real result.
- When the parameter  $X$  has the value  $\pm i$ , the Log function yields an imaginary result.
- Exponentiation by a zero exponent yields the value one. Exponentiation by a unit exponent yields the value of the left operand (as a complex value). Exponentiation of the value one yields the value one. Exponentiation of the value zero yields the value zero.

Other accuracy requirements for the complex elementary functions, which apply only in the strict mode, are given in G.2.6.

The sign of a zero result or zero result component yielded by a complex elementary function is implementation defined when Complex\_Types.Real'Signed\_Zeros is True.

#### Implementation Permissions

The nongeneric equivalent packages can be actual instantiations of the generic package with the appropriate predefined nongeneric equivalent of Numerics.Generic\_Complex\_Types, though that is not required; if they are, then the latter shall have been obtained by actual instantiation of Numerics.Generic\_Complex\_Types.

The exponentiation operator may be implemented in terms of the Exp and Log functions. Because this implementation yields poor accuracy in some parts of the domain, no accuracy requirement is imposed on complex exponentiation.

The implementation of the Exp function of a complex parameter  $X$  is allowed to raise the exception Constraint\_Error, signaling overflow, when the real component of  $X$  exceeds an unspecified threshold that is approximately  $\log(\text{Complex\_Types.Real'Safe\_Last})$ . This permission recognizes the impracticality of avoiding overflow in the marginal case that the exponential of the real component of  $X$  exceeds the safe range of Complex\_Types.Real but both components of the final result do not. Similarly, the Sin and Cos (resp., Sinh and Cosh) functions are allowed to raise the exception Constraint\_Error, signaling overflow, when the absolute value of the imaginary (resp., real) component of the parameter  $X$  exceeds an unspecified threshold that is approximately  $\log(\text{Complex\_Types.Real'Safe\_Last}) + \log(2.0)$ . This permission recognizes the impracticality of avoiding overflow in the marginal case that the hyperbolic sine or cosine of the imaginary (resp., real) component of  $X$  exceeds the safe range of Complex\_Types.Real but both components of the final result do not.

#### Implementation Advice

Implementations in which Complex\_Types.Real'Signed\_Zeros is True should attempt to provide a rational treatment of the signs of zero results and result components. For example, many of the complex elementary functions have components that are odd functions of one of the parameter components; in these cases, the result component should have the sign of the parameter component at the origin. Other complex elementary functions have zero components whose sign is opposite that of a parameter component at the origin, or is always positive or always negative.



### G.1.3 Complex Input-Output

The generic package `Text_IO.Complex_IO` defines procedures for the formatted input and output of complex values. The generic actual parameter in an instantiation of `Text_IO.Complex_IO` is an instance of `Numerics.Generic_Complex_Types` for some floating point subtype. Exceptional conditions are reported by raising the appropriate exception defined in `Text_IO`.

#### Static Semantics

The generic library package `Text_IO.Complex_IO` has the following declaration:

```
with Ada.Numerics.Generic_Complex_Types;
generic
  with package Complex_Types is
    new Ada.Numerics.Generic_Complex_Types (<>);
package Ada.Text_IO.Complex_IO
  with Global => in out synchronized is
  use Complex_Types;

  Default_Fore : Field := 2;
  Default_Aft  : Field := Real'Digits - 1;
  Default_Exp  : Field := 3;

  procedure Get (File : in File_Type;
                Item : out Complex;
                Width : in Field := 0);
  procedure Get (Item : out Complex;
                Width : in Field := 0);

  procedure Put (File : in File_Type;
                Item : in Complex;
                Fore : in Field := Default_Fore;
                Aft  : in Field := Default_Aft;
                Exp  : in Field := Default_Exp);
  procedure Put (Item : in Complex;
                Fore : in Field := Default_Fore;
                Aft  : in Field := Default_Aft;
                Exp  : in Field := Default_Exp);

  procedure Get (From : in String;
                Item : out Complex;
                Last : out Positive)
    with Nonblocking;
  procedure Put (To : out String;
                Item : in Complex;
                Aft : in Field := Default_Aft;
                Exp : in Field := Default_Exp)
    with Nonblocking;
end Ada.Text_IO.Complex_IO;
```

The library package `Complex_Text_IO` defines the same subprograms as `Text_IO.Complex_IO`, except that the predefined type `Float` is systematically substituted for `Real`, and the type `Numerics.Complex_Types.Complex` is systematically substituted for `Complex` throughout. Nongeneric equivalents of `Text_IO.Complex_IO` corresponding to each of the other predefined floating point types are defined similarly, with the names `Short_Complex_Text_IO`, `Long_Complex_Text_IO`, etc.

The semantics of the `Get` and `Put` procedures are as follows:

```
procedure Get (File : in File_Type;
              Item : out Complex;
              Width : in Field := 0);
procedure Get (Item : out Complex;
              Width : in Field := 0);
```

The input sequence is a pair of optionally signed real literals representing the real and imaginary components of a complex value. These components have the format defined for the corresponding `Get` procedure of an instance of `Text_IO.Float_IO` (see A.10.9) for the base



subtype of `Complex_Types.Real`. The pair of components may be separated by a comma or surrounded by a pair of parentheses or both. Blanks are freely allowed before each of the components and before the parentheses and comma, if either is used. If the value of the parameter `Width` is zero, then

- line and page terminators are also allowed in these places;
- the components shall be separated by at least one blank or line terminator if the comma is omitted; and
- reading stops when the right parenthesis has been read, if the input sequence includes a left parenthesis, or when the imaginary component has been read, otherwise.

If a nonzero value of `Width` is supplied, then

- the components shall be separated by at least one blank if the comma is omitted; and
- exactly `Width` characters are read, or the characters (possibly none) up to a line terminator, whichever comes first (blanks are included in the count).

Returns, in the parameter `Item`, the value of type `Complex` that corresponds to the input sequence.

The exception `Text_IO.Data_Error` is raised if the input sequence does not have the required syntax or if the components of the complex value obtained are not of the base subtype of `Complex_Types.Real`.

```

procedure Put (File : in File_Type;
               Item : in Complex;
               Fore : in Field := Default_Fore;
               Aft  : in Field := Default_Aft;
               Exp  : in Field := Default_Exp);
procedure Put (Item : in Complex;
               Fore : in Field := Default_Fore;
               Aft  : in Field := Default_Aft;
               Exp  : in Field := Default_Exp);

```

Outputs the value of the parameter `Item` as a pair of decimal literals representing the real and imaginary components of the complex value, using the syntax of an aggregate. More specifically,

- outputs a left parenthesis;
- outputs the value of the real component of the parameter `Item` with the format defined by the corresponding `Put` procedure of an instance of `Text_IO.Float_IO` for the base subtype of `Complex_Types.Real`, using the given values of `Fore`, `Aft`, and `Exp`;
- outputs a comma;
- outputs the value of the imaginary component of the parameter `Item` with the format defined by the corresponding `Put` procedure of an instance of `Text_IO.Float_IO` for the base subtype of `Complex_Types.Real`, using the given values of `Fore`, `Aft`, and `Exp`;
- outputs a right parenthesis.

```

procedure Get (From : in String;
               Item : out Complex;
               Last : out Positive);

```

Reads a complex value from the beginning of the given string, following the same rule as the `Get` procedure that reads a complex value from a file, but treating the end of the string as a file terminator. Returns, in the parameter `Item`, the value of type `Complex` that corresponds to the input sequence. Returns in `Last` the index value such that `From`(`Last`) is the last character read.

The exception `Text_IO.Data_Error` is raised if the input sequence does not have the required syntax or if the components of the complex value obtained are not of the base subtype of `Complex_Types.Real`.

```

procedure Put (To    : out String;
               Item  : in  Complex;
               Aft   : in  Field := Default_Aft;
               Exp   : in  Field := Default_Exp);

```

Outputs the value of the parameter *Item* to the given string as a pair of decimal literals representing the real and imaginary components of the complex value, using the syntax of an aggregate. More specifically,

- a left parenthesis, the real component, and a comma are left justified in the given string, with the real component having the format defined by the *Put* procedure (for output to a file) of an instance of *Text\_IO.Float\_IO* for the base subtype of *Complex\_Types.Real*, using a value of zero for *Fore* and the given values of *Aft* and *Exp*;
- the imaginary component and a right parenthesis are right justified in the given string, with the imaginary component having the format defined by the *Put* procedure (for output to a file) of an instance of *Text\_IO.Float\_IO* for the base subtype of *Complex\_Types.Real*, using a value for *Fore* that completely fills the remainder of the string, together with the given values of *Aft* and *Exp*.

The exception *Text\_IO.Layout\_Error* is raised if the given string is too short to hold the formatted output.

#### *Implementation Permissions*

Other exceptions declared (by renaming) in *Text\_IO* may be raised by the preceding procedures in the appropriate circumstances, as for the corresponding procedures of *Text\_IO.Float\_IO*.

### **G.1.4 The Package *Wide\_Text\_IO.Complex\_IO***

#### *Static Semantics*

Implementations shall also provide the generic library package *Wide\_Text\_IO.Complex\_IO*. Its declaration is obtained from that of *Text\_IO.Complex\_IO* by systematically replacing *Text\_IO* by *Wide\_Text\_IO* and *String* by *Wide\_String*; the description of its behavior is obtained by additionally replacing references to particular characters (commas, parentheses, etc.) by those for the corresponding wide characters.

### **G.1.5 The Package *Wide\_Wide\_Text\_IO.Complex\_IO***

#### *Static Semantics*

Implementations shall also provide the generic library package *Wide\_Wide\_Text\_IO.Complex\_IO*. Its declaration is obtained from that of *Text\_IO.Complex\_IO* by systematically replacing *Text\_IO* by *Wide\_Wide\_Text\_IO* and *String* by *Wide\_Wide\_String*; the description of its behavior is obtained by additionally replacing references to particular characters (commas, parentheses, etc.) by those for the corresponding wide wide characters.

## **G.2 Numeric Performance Requirements**

#### *Implementation Requirements*

Implementations shall provide a user-selectable mode in which the accuracy and other numeric performance requirements detailed in the following subclauses are observed. This mode, referred to as the *strict mode*, may or may not be the default mode; it directly affects the results of the predefined arithmetic operations of real types and the results of the subprograms in children of the *Numerics* package, and indirectly affects the operations in other language defined packages. Implementations shall also provide the opposing mode, which is known as the *relaxed mode*.

#### *Implementation Permissions*

Either mode may be the default mode.

The two modes can be one and the same.

### G.2.1 Model of Floating Point Arithmetic

In the strict mode, the predefined operations of a floating point type shall satisfy the accuracy requirements specified here and shall avoid or signal overflow in the situations described. This behavior is presented in terms of a model of floating point arithmetic that builds on the concept of the canonical form (see A.5.3).

#### Static Semantics

Associated with each floating point type is an infinite set of model numbers. The model numbers of a type are used to define the accuracy requirements that have to be satisfied by certain predefined operations of the type; through certain attributes of the model numbers, they are also used to explain the meaning of a user-declared floating point type declaration. The model numbers of a derived type are those of the parent type; the model numbers of a subtype are those of its type.

The *model numbers* of a floating point type  $T$  are zero and all the values expressible in the canonical form (for the type  $T$ ), in which *mantissa* has  $T$ 'Model\_Mantissa digits and *exponent* has a value greater than or equal to  $T$ 'Model\_Emin. (These attributes are defined in G.2.2.)

A *model interval* of a floating point type is any interval whose bounds are model numbers of the type. The *model interval* of a type  $T$  associated with a value  $v$  is the smallest model interval of  $T$  that includes  $v$ . (The model interval associated with a model number of a type consists of that number only.)

#### Implementation Requirements

The accuracy requirements for the evaluation of certain predefined operations of floating point types are as follows.

An *operand interval* is the model interval, of the type specified for the operand of an operation, associated with the value of the operand.

For any predefined arithmetic operation that yields a result of a floating point type  $T$ , the required bounds on the result are given by a model interval of  $T$  (called the *result interval*) defined in terms of the operand values as follows:

- The result interval is the smallest model interval of  $T$  that includes the minimum and the maximum of all the values obtained by applying the (exact) mathematical operation to values arbitrarily selected from the respective operand intervals.

The result interval of an exponentiation is obtained by applying the above rule to the sequence of multiplications defined by the exponent, assuming arbitrary association of the factors, and to the final division in the case of a negative exponent.

The result interval of a conversion of a numeric value to a floating point type  $T$  is the model interval of  $T$  associated with the operand value, except when the source expression is of a fixed point type with a *small* that is not a power of  $T$ 'Machine\_Radix or is a fixed point multiplication or division either of whose operands has a *small* that is not a power of  $T$ 'Machine\_Radix; in these cases, the result interval is implementation defined.

For any of the foregoing operations, the implementation shall deliver a value that belongs to the result interval when both bounds of the result interval are in the safe range of the result type  $T$ , as determined by the values of  $T$ 'Safe\_First and  $T$ 'Safe\_Last; otherwise,

- if  $T$ 'Machine\_Overflows is True, the implementation shall either deliver a value that belongs to the result interval or raise Constraint\_Error;
- if  $T$ 'Machine\_Overflows is False, the result is implementation defined.

For any predefined relation on operands of a floating point type  $T$ , the implementation may deliver any value (i.e. either True or False) obtained by applying the (exact) mathematical comparison to values arbitrarily chosen from the respective operand intervals.

The result of a membership test is defined in terms of comparisons of the operand value with the lower and upper bounds of the given range or type mark (the usual rules apply to these comparisons).

*Implementation Permissions*

If the underlying floating point hardware implements division as multiplication by a reciprocal, the result interval for division (and exponentiation by a negative exponent) is implementation defined.

## G.2.2 Model-Oriented Attributes of Floating Point Types

In implementations that support the Numerics Annex, the model-oriented attributes of floating point types shall yield the values defined here, in both the strict and the relaxed modes. These definitions add conditions to those in A.5.3.

*Static Semantics*

For every subtype  $S$  of a floating point type  $T$ :

**S'Model\_Mantissa**

Yields the number of digits in the mantissa of the canonical form of the model numbers of  $T$  (see A.5.3). The value of this attribute shall be greater than or equal to

$$\lceil d \cdot \log(10) / \log(T^{\text{Machine\_Radix}}) \rceil + g$$

where  $d$  is the requested decimal precision of  $T$ , and  $g$  is 0 if  $T^{\text{Machine\_Radix}}$  is a positive power of 10 and 1 otherwise. In addition,  $T^{\text{Model\_Mantissa}}$  shall be less than or equal to the value of  $T^{\text{Machine\_Mantissa}}$ . This attribute yields a value of the type *universal\_integer*.

**S'Model\_Emin**

Yields the minimum exponent of the canonical form of the model numbers of  $T$  (see A.5.3). The value of this attribute shall be greater than or equal to the value of  $T^{\text{Machine\_Emin}}$ . This attribute yields a value of the type *universal\_integer*.

**S'Safe\_First**

Yields the lower bound of the safe range of  $T$ . The value of this attribute shall be a model number of  $T$  and greater than or equal to the lower bound of the base range of  $T$ . In addition, if  $T$  is declared by a *floating\_point\_definition* or is derived from such a type, and the *floating\_point\_definition* includes a *real\_range\_specification* specifying a lower bound of  $lb$ , then the value of this attribute shall be less than or equal to  $lb$ ; otherwise, it shall be less than or equal to  $-10.0^{4 \cdot d}$ , where  $d$  is the requested decimal precision of  $T$ . This attribute yields a value of the type *universal\_real*.

**S'Safe\_Last**

Yields the upper bound of the safe range of  $T$ . The value of this attribute shall be a model number of  $T$  and less than or equal to the upper bound of the base range of  $T$ . In addition, if  $T$  is declared by a *floating\_point\_definition* or is derived from such a type, and the *floating\_point\_definition* includes a *real\_range\_specification* specifying an upper bound of  $ub$ , then the value of this attribute shall be greater than or equal to  $ub$ ; otherwise, it shall be greater than or equal to  $10.0^{4 \cdot d}$ , where  $d$  is the requested decimal precision of  $T$ . This attribute yields a value of the type *universal\_real*.

**S'Model**

Denotes a function (of a parameter  $X$ ) whose specification is given in A.5.3. If  $X$  is a model number of  $T$ , the function yields  $X$ ; otherwise, it yields the value obtained by rounding or truncating  $X$  to either one of the adjacent model numbers of  $T$ . *Constraint\_Error* is raised if the resulting model number is outside the safe range of  $S$ . A zero result has the sign of  $X$  when *S'Signed\_Zeros* is True.

Subject to the constraints given above, the values of *S'Model\_Mantissa* and *S'Safe\_Last* are to be maximized, and the values of *S'Model\_Emin* and *S'Safe\_First* minimized, by the implementation as follows:

- First, S'Model\_Mantissa is set to the largest value for which values of S'Model\_Emin, S'Safe\_First, and S'Safe\_Last can be chosen so that the implementation satisfies the strict-mode requirements of G.2.1 in terms of the model numbers and safe range induced by these attributes.
- Next, S'Model\_Emin is set to the smallest value for which values of S'Safe\_First and S'Safe\_Last can be chosen so that the implementation satisfies the strict-mode requirements of G.2.1 in terms of the model numbers and safe range induced by these attributes and the previously determined value of S'Model\_Mantissa.
- Finally, S'Safe\_First and S'Safe\_Last are set (in either order) to the smallest and largest values, respectively, for which the implementation satisfies the strict-mode requirements of G.2.1 in terms of the model numbers and safe range induced by these attributes and the previously determined values of S'Model\_Mantissa and S'Model\_Emin.

### G.2.3 Model of Fixed Point Arithmetic

In the strict mode, the predefined arithmetic operations of a fixed point type shall satisfy the accuracy requirements specified here and shall avoid or signal overflow in the situations described.

#### Implementation Requirements

The accuracy requirements for the predefined fixed point arithmetic operations and conversions, and the results of relations on fixed point operands, are given below.

The operands of the fixed point adding operators, absolute value, and comparisons have the same type. These operations are required to yield exact results, unless they overflow.

Multiplications and divisions are allowed between operands of any two fixed point types; the result has to be (implicitly or explicitly) converted to some other numeric type. For purposes of defining the accuracy rules, the multiplication or division and the conversion are treated as a single operation whose accuracy depends on three types (those of the operands and the result). For decimal fixed point types, the attribute T'Round may be used to imply explicit conversion with rounding (see 6.5.10).

When the result type is a floating point type, the accuracy is as given in G.2.1. For some combinations of the operand and result types in the remaining cases, the result is required to belong to a small set of values called the *perfect result set*; for other combinations, it is required merely to belong to a generally larger and implementation-defined set of values called the *close result set*. When the result type is a decimal fixed point type, the perfect result set contains a single value; thus, operations on decimal types are always fully specified.

When one operand of a fixed-fixed multiplication or division is of type *universal\_real*, that operand is not implicitly converted in the usual sense, since the context does not determine a unique target type, but the accuracy of the result of the multiplication or division (i.e. whether the result has to belong to the perfect result set or merely the close result set) depends on the value of the operand of type *universal\_real* and on the types of the other operand and of the result.

For a fixed point multiplication or division whose (exact) mathematical result is  $v$ , and for the conversion of a value  $v$  to a fixed point type, the perfect result set and close result set are defined as follows:

- If the result type is an ordinary fixed point type with a *small* of  $s$ ,
  - if  $v$  is an integer multiple of  $s$ , then the perfect result set contains only the value  $v$ ;
  - otherwise, it contains the integer multiple of  $s$  just below  $v$  and the integer multiple of  $s$  just above  $v$ .

The close result set is an implementation-defined set of consecutive integer multiples of  $s$  containing the perfect result set as a subset.

- If the result type is a decimal type with a *small* of  $s$ ,
  - if  $v$  is an integer multiple of  $s$ , then the perfect result set contains only the value  $v$ ;



- otherwise, if truncation applies, then it contains only the integer multiple of  $s$  in the direction toward zero, whereas if rounding applies, then it contains only the nearest integer multiple of  $s$  (with ties broken by rounding away from zero).

The close result set is an implementation-defined set of consecutive integer multiples of  $s$  containing the perfect result set as a subset.

- If the result type is an integer type,
  - if  $v$  is an integer, then the perfect result set contains only the value  $v$ ;
  - otherwise, it contains the integer nearest to the value  $v$  (if  $v$  lies equally distant from two consecutive integers, the perfect result set contains the one that is further from zero).

The close result set is an implementation-defined set of consecutive integers containing the perfect result set as a subset.

The result of a fixed point multiplication or division shall belong either to the perfect result set or to the close result set, as described below, if overflow does not occur. In the following cases, if the result type is a fixed point type, let  $s$  be its *small*; otherwise, i.e. when the result type is an integer type, let  $s$  be 1.0.

- For a multiplication or division neither of whose operands is of type *universal\_real*, let  $l$  and  $r$  be the *smalls* of the left and right operands. For a multiplication, if  $(l \cdot r) / s$  is an integer or the reciprocal of an integer (the *smalls* are said to be “compatible” in this case), the result shall belong to the perfect result set; otherwise, it belongs to the close result set. For a division, if  $l / (r \cdot s)$  is an integer or the reciprocal of an integer (i.e. the *smalls* are compatible), the result shall belong to the perfect result set; otherwise, it belongs to the close result set.
- For a multiplication or division having one *universal\_real* operand with a value of  $v$ , note that it is always possible to factor  $v$  as an integer multiple of a “compatible” *small*, but the integer multiple may be “too big”. If there exists a factorization in which that multiple is less than some implementation-defined limit, the result shall belong to the perfect result set; otherwise, it belongs to the close result set.

A multiplication  $P * Q$  of an operand of a fixed point type  $F$  by an operand of type Integer, or vice versa, and a division  $P / Q$  of an operand of a fixed point type  $F$  by an operand of type Integer, are also allowed. In these cases, the result has the type of  $F$ ; explicit conversion of the result is never required. The accuracy required in these cases is the same as that required for a multiplication  $F(P * Q)$  or a division  $F(P / Q)$  obtained by interpreting the operand of the integer type to have a fixed point type with a *small* of 1.0.

The accuracy of the result of a conversion from an integer or fixed point type to a fixed point type, or from a fixed point type to an integer type, is the same as that of a fixed point multiplication of the source value by a fixed point operand having a *small* of 1.0 and a value of 1.0, as given by the foregoing rules. The result of a conversion from a floating point type to a fixed point type shall belong to the close result set. The result of a conversion of a *universal\_real* operand to a fixed point type shall belong to the perfect result set.

The possibility of overflow in the result of a predefined arithmetic operation or conversion yielding a result of a fixed point type  $T$  is analogous to that for floating point types, except for being related to the base range instead of the safe range. If all of the permitted results belong to the base range of  $T$ , then the implementation shall deliver one of the permitted results; otherwise,

- if  $T$ .Machine\_Overflows is True, the implementation shall either deliver one of the permitted results or raise Constraint\_Error;
- if  $T$ .Machine\_Overflows is False, the result is implementation defined.

#### G.2.4 Accuracy Requirements for the Elementary Functions

In the strict mode, the performance of Numerics.Generic\_Elementary\_Functions shall be as specified here.



## Implementation Requirements

When an exception is not raised, the result of evaluating a function in an instance *EF* of *Numerics.Generic\_Elementary\_Functions* belongs to a *result interval*, defined as the smallest model interval of *EF.Float\_Type* that contains all the values of the form  $f \cdot (1.0 + d)$ , where  $f$  is the exact value of the corresponding mathematical function at the given parameter values,  $d$  is a real number, and  $|d|$  is less than or equal to the function's *maximum relative error*. The function delivers a value that belongs to the result interval when both of its bounds belong to the safe range of *EF.Float\_Type*; otherwise,

- if *EF.Float\_Type*'Machine\_Overflows is True, the function either delivers a value that belongs to the result interval or raises *Constraint\_Error*, signaling overflow;
- if *EF.Float\_Type*'Machine\_Overflows is False, the result is implementation defined.

The maximum relative error exhibited by each function is as follows:

- $2.0 \cdot EF.Float\_Type'Model\_Epsilon$ , in the case of the Sqrt, Sin, and Cos functions;
- $4.0 \cdot EF.Float\_Type'Model\_Epsilon$ , in the case of the Log, Exp, Tan, Cot, and inverse trigonometric functions; and
- $8.0 \cdot EF.Float\_Type'Model\_Epsilon$ , in the case of the forward and inverse hyperbolic functions.

The maximum relative error exhibited by the exponentiation operator, which depends on the values of the operands, is  $(4.0 + |Right \cdot \log(Left)| / 32.0) \cdot EF.Float\_Type'Model\_Epsilon$ .

The maximum relative error given above applies throughout the domain of the forward trigonometric functions when the Cycle parameter is specified. When the Cycle parameter is omitted, the maximum relative error given above applies only when the absolute value of the angle parameter *X* is less than or equal to some implementation-defined *angle threshold*, which shall be at least  $EF.Float\_Type'Machine\_Radix^{\lfloor EF.Float\_Type'Machine\_Mantissa/2 \rfloor}$ . Beyond the angle threshold, the accuracy of the forward trigonometric functions is implementation defined.

The prescribed results specified in A.5.1 for certain functions at particular parameter values take precedence over the maximum relative error bounds; effectively, they narrow to a single value the result interval allowed by the maximum relative error bounds. Additional rules with a similar effect are given by Table G.1 for the inverse trigonometric functions, at particular parameter values for which the mathematical result is possibly not a model number of *EF.Float\_Type* (or is, indeed, even transcendental). In each table entry, the values of the parameters are such that the result lies on the axis between two quadrants; the corresponding accuracy rule, which takes precedence over the maximum relative error bounds, is that the result interval is the model interval of *EF.Float\_Type* associated with the exact mathematical result given in the table.

The last line of the table is meant to apply when *EF.Float\_Type*'Signed\_Zeros is False; the two lines just above it, when *EF.Float\_Type*'Signed\_Zeros is True and the parameter *Y* has a zero value with the indicated sign.

**Table G.1 — Tightly Approximated Elementary Function Results**

| Function | Value of X | Value of Y | Exact Result when Cycle Specified | Exact Result when Cycle Omitted |
|----------|------------|------------|-----------------------------------|---------------------------------|
| Arcsin   | 1.0        | n.a.       | Cycle/4.0                         | $\pi/2.0$                       |
| Arcsin   | -1.0       | n.a.       | -Cycle/4.0                        | $-\pi/2.0$                      |
| Arccos   | 0.0        | n.a.       | Cycle/4.0                         | $\pi/2.0$                       |
| Arccos   | -1.0       | n.a.       | Cycle/2.0                         | $\pi$                           |

|                   |          |          |                     |            |
|-------------------|----------|----------|---------------------|------------|
| Arctan and Arccot | 0.0      | positive | Cycle/4.0           | $\pi/2.0$  |
| Arctan and Arccot | 0.0      | negative | $-\text{Cycle}/4.0$ | $-\pi/2.0$ |
| Arctan and Arccot | negative | +0.0     | Cycle/2.0           | $\pi$      |
| Arctan and Arccot | negative | -0.0     | $-\text{Cycle}/2.0$ | $-\pi$     |
| Arctan and Arccot | negative | 0.0      | Cycle/2.0           | $\pi$      |

The amount by which the result of an inverse trigonometric function is allowed to spill over into a quadrant adjacent to the one corresponding to the principal branch, as given in A.5.1, is limited. The rule is that the result belongs to the smallest model interval of *EF.Float\_Type* that contains both boundaries of the quadrant corresponding to the principal branch. This rule also takes precedence over the maximum relative error bounds, effectively narrowing the result interval allowed by them.

Finally, the following specifications also take precedence over the maximum relative error bounds:

- The absolute value of the result of the Sin, Cos, and Tanh functions never exceeds one.
- The absolute value of the result of the Coth function is never less than one.
- The result of the Cosh function is never less than one.

#### Implementation Advice

The versions of the forward trigonometric functions without a Cycle parameter should not be implemented by calling the corresponding version with a Cycle parameter of  $2.0 \times \text{Numerics.Pi}$ , since this will not provide the required accuracy in some portions of the domain. For the same reason, the version of Log without a Base parameter should not be implemented by calling the corresponding version with a Base parameter of *Numerics.e*.

### G.2.5 Performance Requirements for Random Number Generation

In the strict mode, the performance of *Numerics.Float\_Random* and *Numerics.Discrete\_Random* shall be as specified here.

#### Implementation Requirements

Two different calls to the time-dependent Reset procedure shall reset the generator to different states, provided that the calls are separated in time by at least one second and not more than fifty years.

The implementation's representations of generator states and its algorithms for generating random numbers shall yield a period of at least  $2^{31}-2$ ; much longer periods are desirable but not required.

The implementations of *Numerics.Float\_Random.Random* and *Numerics.Discrete\_Random.Random* shall pass at least 85% of the individual trials in a suite of statistical tests. For *Numerics.Float\_Random*, the tests are applied directly to the floating point values generated (i.e. they are not converted to integers first), while for *Numerics.Discrete\_Random* they are applied to the generated values of various discrete types. Each test suite performs 6 different tests, with each test repeated 10 times, yielding a total of 60 individual trials. An individual trial is deemed to pass if the chi-square value (or other statistic) calculated for the observed counts or distribution falls within the range of values corresponding to the 2.5 and 97.5 percentage points for the relevant degrees of freedom (i.e. it shall be neither too high nor too low). For the purpose of determining the degrees of freedom, measurement categories are combined whenever the expected counts are fewer than 5.

### G.2.6 Accuracy Requirements for Complex Arithmetic

In the strict mode, the performance of *Numerics.Generic\_Complex\_Types* and *Numerics.Generic\_Complex\_Elementary\_Functions* shall be as specified here.

When an exception is not raised, the result of evaluating a real function of an instance *CT* of `Numerics.Generic_Complex_Types` (i.e. a function that yields a value of subtype *CT.Real*'Base or *CT.Imaginary*) belongs to a result interval defined as for a real elementary function (see G.2.4).

When an exception is not raised, each component of the result of evaluating a complex function of such an instance, or of an instance of `Numerics.Generic_Complex_Elementary_Functions` obtained by instantiating the latter with *CT* (i.e. a function that yields a value of subtype *CT.Complex*), also belongs to a *result interval*. The result intervals for the components of the result are either defined by a *maximum relative error* bound or by a *maximum box error* bound. When the result interval for the real (resp., imaginary) component is defined by maximum relative error, it is defined as for that of a real function, relative to the exact value of the real (resp., imaginary) part of the result of the corresponding mathematical function. When defined by maximum box error, the result interval for a component of the result is the smallest model interval of *CT.Real* that contains all the values of the corresponding part of  $f \cdot (1.0 + d)$ , where  $f$  is the exact complex value of the corresponding mathematical function at the given parameter values,  $d$  is complex, and  $|d|$  is less than or equal to the given maximum box error. The function delivers a value that belongs to the result interval (or a value both of whose components belong to their respective result intervals) when both bounds of the result interval(s) belong to the safe range of *CT.Real*; otherwise,

- if *CT.Real*'Machine\_Overflows is True, the function either delivers a value that belongs to the result interval (or a value both of whose components belong to their respective result intervals) or raises `Constraint_Error`, signaling overflow;
- if *CT.Real*'Machine\_Overflows is False, the result is implementation defined.

The error bounds for particular complex functions are tabulated in Table G.2. In the table, the error bound is given as the coefficient of *CT.Real*'Model\_Epsilon.

**Table G.2 — Error Bounds for Particular Complex Functions**

| Function or Operator        | Nature of Result | Nature of Bound | Error Bound |
|-----------------------------|------------------|-----------------|-------------|
| Modulus                     | real             | max. rel. error | 3.0         |
| Argument                    | real             | max. rel. error | 4.0         |
| Compose_From_Polar          | complex          | max. rel. error | 3.0         |
| "*" (both operands complex) | complex          | max. box error  | 5.0         |
| "/" (right operand complex) | complex          | max. box error  | 13.0        |
| Sqrt                        | complex          | max. rel. error | 6.0         |
| Log                         | complex          | max. box error  | 13.0        |
| Exp (complex parameter)     | complex          | max. rel. error | 7.0         |
| Exp (imaginary parameter)   | complex          | max. rel. error | 2.0         |
| Sin, Cos, Sinh, and Cosh    | complex          | max. rel. error | 11.0        |
| Tan, Cot, Tanh, and Coth    | complex          | max. rel. error | 35.0        |

|                       |         |                 |      |
|-----------------------|---------|-----------------|------|
| inverse trigonometric | complex | max. rel. error | 14.0 |
| inverse hyperbolic    | complex | max. rel. error | 14.0 |

The maximum relative error given above applies throughout the domain of the Compose\_From\_Polar function when the Cycle parameter is specified. When the Cycle parameter is omitted, the maximum relative error applies only when the absolute value of the parameter Argument is less than or equal to the angle threshold (see G.2.4). For the Exp function, and for the forward hyperbolic (resp., trigonometric) functions, the maximum relative error given above likewise applies only when the absolute value of the imaginary (resp., real) component of the parameter X (or the absolute value of the parameter itself, in the case of the Exp function with a parameter of pure-imaginary type) is less than or equal to the angle threshold. For larger angles, the accuracy is implementation defined.

The prescribed results specified in G.1.2 for certain functions at particular parameter values take precedence over the error bounds; effectively, they narrow to a single value the result interval allowed by the error bounds for a component of the result. Additional rules with a similar effect are given below for certain inverse trigonometric and inverse hyperbolic functions, at particular parameter values for which a component of the mathematical result is transcendental. In each case, the accuracy rule, which takes precedence over the error bounds, is that the result interval for the stated result component is the model interval of CT.Real associated with the component's exact mathematical value. The cases in question are as follows:

- When the parameter X has the value zero, the real (resp., imaginary) component of the result of the Arccot (resp., Arccoth) function is in the model interval of CT.Real associated with the value  $\pi/2.0$ .
- When the parameter X has the value one, the real component of the result of the Arcsin function is in the model interval of CT.Real associated with the value  $\pi/2.0$ .
- When the parameter X has the value  $-1.0$ , the real component of the result of the Arcsin (resp., Arccos) function is in the model interval of CT.Real associated with the value  $-\pi/2.0$  (resp.,  $\pi$ ).

The amount by which a component of the result of an inverse trigonometric or inverse hyperbolic function is allowed to spill over into a quadrant adjacent to the one corresponding to the principal branch, as given in G.1.2, is limited. The rule is that the result belongs to the smallest model interval of CT.Real that contains both boundaries of the quadrant corresponding to the principal branch. This rule also takes precedence over the maximum error bounds, effectively narrowing the result interval allowed by them.

Finally, the results allowed by the error bounds are narrowed by one further rule: The absolute value of each component of the result of the Exp function, for a pure-imaginary parameter, never exceeds one.

#### *Implementation Advice*

The version of the Compose\_From\_Polar function without a Cycle parameter should not be implemented by calling the corresponding version with a Cycle parameter of  $2.0 * \text{Numerics.Pi}$ , since this will not provide the required accuracy in some portions of the domain.

### G.3 Vector and Matrix Manipulation

Types and operations for the manipulation of real vectors and matrices are provided in Generic\_Real\_Arrays, which is defined in G.3.1. Types and operations for the manipulation of complex vectors and matrices are provided in Generic\_Complex\_Arrays, which is defined in G.3.2. Both of these library units are generic children of the predefined package Numerics (see A.5). Nongeneric equivalents of these packages for each of the predefined floating point types are also provided as children of Numerics.

### G.3.1 Real Vectors and Matrices

#### Static Semantics

The generic library package Numerics.Generic\_Real\_Arrays has the following declaration:

```

generic
  type Real is digits <>;
package Ada.Numerics.Generic_Real_Arrays
  with Pure, Nonblocking is

  -- Types

  type Real_Vector is array (Integer range <>) of Real'Base;
  type Real_Matrix is array (Integer range <>, Integer range <>)
                                of Real'Base;

  -- Subprograms for Real_Vector types
  -- Real_Vector arithmetic operations

  function "+" (Right : Real_Vector) return Real_Vector;
  function "-" (Right : Real_Vector) return Real_Vector;
  function "abs" (Right : Real_Vector) return Real_Vector;

  function "+" (Left, Right : Real_Vector) return Real_Vector;
  function "-" (Left, Right : Real_Vector) return Real_Vector;

  function "*" (Left, Right : Real_Vector) return Real'Base;
  function "abs" (Right : Real_Vector) return Real'Base;

  -- Real_Vector scaling operations

  function "*" (Left : Real'Base; Right : Real_Vector)
    return Real_Vector;
  function "*" (Left : Real_Vector; Right : Real'Base)
    return Real_Vector;
  function "/" (Left : Real_Vector; Right : Real'Base)
    return Real_Vector;

  -- Other Real_Vector operations

  function Unit_Vector (Index : Integer;
                       Order : Positive;
                       First : Integer := 1) return Real_Vector;

  -- Subprograms for Real_Matrix types
  -- Real_Matrix arithmetic operations

  function "+" (Right : Real_Matrix) return Real_Matrix;
  function "-" (Right : Real_Matrix) return Real_Matrix;
  function "abs" (Right : Real_Matrix) return Real_Matrix;
  function Transpose (X : Real_Matrix) return Real_Matrix;

  function "+" (Left, Right : Real_Matrix) return Real_Matrix;
  function "-" (Left, Right : Real_Matrix) return Real_Matrix;
  function "*" (Left, Right : Real_Matrix) return Real_Matrix;

  function "*" (Left, Right : Real_Vector) return Real_Matrix;

  function "*" (Left : Real_Vector; Right : Real_Matrix)
    return Real_Vector;
  function "*" (Left : Real_Matrix; Right : Real_Vector)
    return Real_Vector;

  -- Real_Matrix scaling operations

  function "*" (Left : Real'Base; Right : Real_Matrix)
    return Real_Matrix;
  function "*" (Left : Real_Matrix; Right : Real'Base)
    return Real_Matrix;
  function "/" (Left : Real_Matrix; Right : Real'Base)
    return Real_Matrix;

```

```

-- Real_Matrix inversion and related operations
function Solve (A : Real_Matrix; X : Real_Vector) return Real_Vector;
function Solve (A, X : Real_Matrix) return Real_Matrix;
function Inverse (A : Real_Matrix) return Real_Matrix;
function Determinant (A : Real_Matrix) return Real'Base;

-- Eigenvalues and vectors of a real symmetric matrix
function Eigenvalues (A : Real_Matrix) return Real_Vector;

procedure Eigensystem (A          : in Real_Matrix;
                      Values     : out Real_Vector;
                      Vectors    : out Real_Matrix);

-- Other Real_Matrix operations
function Unit_Matrix (Order      : Positive;
                     First_1, First_2 : Integer := 1)
                     return Real_Matrix;

end Ada.Numerics.Generic_Real_Arrays;

```

The library package Numerics.Real\_Arrays is declared pure and defines the same types and subprograms as Numerics.Generic\_Real\_Arrays, except that the predefined type Float is systematically substituted for Real'Base throughout. Nongeneric equivalents for each of the other predefined floating point types are defined similarly, with the names Numerics.Short\_Real\_Arrays, Numerics.Long\_Real\_Arrays, etc.

Two types are defined and exported by Numerics.Generic\_Real\_Arrays. The composite type Real\_Vector is provided to represent a vector with components of type Real; it is defined as an unconstrained, one-dimensional array with an index of type Integer. The composite type Real\_Matrix is provided to represent a matrix with components of type Real; it is defined as an unconstrained, two-dimensional array with indices of type Integer.

The effect of the various subprograms is as described below. In most cases the subprograms are described in terms of corresponding scalar operations of the type Real; any exception raised by those operations is propagated by the array operation. Moreover, the accuracy of the result for each individual component is as defined for the scalar operation unless stated otherwise.

In the case of those operations which are defined to *involve an inner product*, Constraint\_Error may be raised if an intermediate result is outside the range of Real'Base even though the mathematical final result would not be.

```

function "+" (Right : Real_Vector) return Real_Vector;
function "-" (Right : Real_Vector) return Real_Vector;
function "abs" (Right : Real_Vector) return Real_Vector;

```

Each operation returns the result of applying the corresponding operation of the type Real to each component of Right. The index range of the result is Right'Range.

```

function "+" (Left, Right : Real_Vector) return Real_Vector;
function "-" (Left, Right : Real_Vector) return Real_Vector;

```

Each operation returns the result of applying the corresponding operation of the type Real to each component of Left and the matching component of Right. The index range of the result is Left'Range. Constraint\_Error is raised if Left'Length is not equal to Right'Length.

```

function "*" (Left, Right : Real_Vector) return Real'Base;

```

This operation returns the inner product of Left and Right. Constraint\_Error is raised if Left'Length is not equal to Right'Length. This operation involves an inner product.

```

function "abs" (Right : Real_Vector) return Real'Base;

```

This operation returns the L2-norm of Right (the square root of the inner product of the vector with itself).



```
function "*" (Left : Real_Base; Right : Real_Vector) return Real_Vector;
```

This operation returns the result of multiplying each component of Right by the scalar Left using the "\*" operation of the type Real. The index range of the result is Right'Range.

```
function "*" (Left : Real_Vector; Right : Real_Base) return Real_Vector;
function "/" (Left : Real_Vector; Right : Real_Base) return Real_Vector;
```

Each operation returns the result of applying the corresponding operation of the type Real to each component of Left and to the scalar Right. The index range of the result is Left'Range.

```
function Unit_Vector (Index : Integer;
                     Order : Positive;
                     First : Integer := 1) return Real_Vector;
```

This function returns a *unit vector* with Order components and a lower bound of First. All components are set to 0.0 except for the Index component which is set to 1.0. Constraint\_Error is raised if Index < First, Index > First + Order - 1 or if First + Order - 1 > Integer'Last.

```
function "+" (Right : Real_Matrix) return Real_Matrix;
function "-" (Right : Real_Matrix) return Real_Matrix;
function "abs" (Right : Real_Matrix) return Real_Matrix;
```

Each operation returns the result of applying the corresponding operation of the type Real to each component of Right. The index ranges of the result are those of Right.

```
function Transpose (X : Real_Matrix) return Real_Matrix;
```

This function returns the transpose of a matrix X. The first and second index ranges of the result are X'Range(2) and X'Range(1) respectively.

```
function "+" (Left, Right : Real_Matrix) return Real_Matrix;
function "-" (Left, Right : Real_Matrix) return Real_Matrix;
```

Each operation returns the result of applying the corresponding operation of the type Real to each component of Left and the matching component of Right. The index ranges of the result are those of Left. Constraint\_Error is raised if Left'Length(1) is not equal to Right'Length(1) or Left'Length(2) is not equal to Right'Length(2).

```
function "*" (Left, Right : Real_Matrix) return Real_Matrix;
```

This operation provides the standard mathematical operation for matrix multiplication. The first and second index ranges of the result are Left'Range(1) and Right'Range(2) respectively. Constraint\_Error is raised if Left'Length(2) is not equal to Right'Length(1). This operation involves inner products.

```
function "*" (Left, Right : Real_Vector) return Real_Matrix;
```

This operation returns the outer product of a (column) vector Left by a (row) vector Right using the operation "\*" of the type Real for computing the individual components. The first and second index ranges of the result are Left'Range and Right'Range respectively.

```
function "*" (Left : Real_Vector; Right : Real_Matrix) return Real_Vector;
```

This operation provides the standard mathematical operation for multiplication of a (row) vector Left by a matrix Right. The index range of the (row) vector result is Right'Range(2). Constraint\_Error is raised if Left'Length is not equal to Right'Length(1). This operation involves inner products.

```
function "*" (Left : Real_Matrix; Right : Real_Vector) return Real_Vector;
```

This operation provides the standard mathematical operation for multiplication of a matrix Left by a (column) vector Right. The index range of the (column) vector result is Left'Range(1). Constraint\_Error is raised if Left'Length(2) is not equal to Right'Length. This operation involves inner products.