
**Information technology — Programming
languages — Guide for the use of the Ada
programming language in high integrity
systems**

*Technologies de l'information — Langues de programmation — Guide
pour l'emploi du langage de programmation Ada dans les systèmes de
haute intégrité*

PDF disclaimer

This PDF file may contain embedded typefaces. In accordance with Adobe's licensing policy, this file may be printed or viewed but shall not be edited unless the typefaces which are embedded are licensed to and installed on the computer performing the editing. In downloading this file, parties accept therein the responsibility of not infringing Adobe's licensing policy. The ISO Central Secretariat accepts no liability in this area.

Adobe is a trademark of Adobe Systems Incorporated.

Details of the software products used to create this PDF file can be found in the General Info relative to the file; the PDF-creation parameters were optimized for printing. Every care has been taken to ensure that the file is suitable for use by ISO member bodies. In the unlikely event that a problem relating to it is found, please inform the Central Secretariat at the address given below.

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 15942:2000

© ISO/IEC 2000

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 734 10 79
E-mail copyright@iso.ch
Web www.iso.ch

Printed in Switzerland

Contents

1	Scope	1
1.1	Within the scope	1
1.2	Out of scope	2
2	Verification Techniques	2
2.1	Traceability	2
2.2	Reviews	3
2.3	Analysis	3
2.3.1	Control Flow analysis	4
2.3.2	Data Flow analysis	4
2.3.3	Information Flow analysis	4
2.3.4	Symbolic Execution	4
2.3.5	Formal Code Verification	5
2.3.6	Range Checking	6
2.3.7	Stack Usage analysis	6
2.3.8	Timing Analysis	6
2.3.9	Other Memory Usage analysis	6
2.3.10	Object Code Analysis	7
2.4	Testing	7
2.4.1	Principles	7
2.4.2	Requirements-based Testing	7
2.4.3	Structure-based Testing	8
2.5	Use of Verification Techniques in this Technical Report	8
3	General Language Issues	9
3.1	Writing Verifiable Programs	9
3.1.1	Language Rules to Achieve Predictability	10
3.1.2	Language Rules to Allow Modelling	10
3.1.3	Language Rules to Facilitate Testing	11
3.1.4	Pragmatic Considerations	12
3.1.5	Language Enhancements	12
3.2	The Choice of Language	13
4	Significance of Language Features for High Integrity	14
4.1	Criteria for Assessment of Language Features	14
4.2	How to use this Technical Report	14
5	Assessment of Language Features	15
5.1	Types with Static Attributes	16
5.1.1	Evaluation	17
5.1.2	Notes	17
5.1.3	Guidance	17
5.2	Declarations	17
5.2.1	Evaluation	18
5.2.2	Notes	18
5.2.3	Guidance	18
5.3	Names, including Scope and Visibility	19
5.3.1	Evaluation	19
5.3.2	Notes	19
5.3.3	Guidance	20
5.4	Expressions	20
5.4.1	Evaluation	21
5.4.2	Notes	21
5.4.3	Guidance	22

5.5	Statements	22
5.5.1	Evaluation	23
5.5.2	Notes	23
5.5.3	Guidance	23
5.6	Subprograms	24
5.6.1	Evaluation	24
5.6.2	Notes	24
5.6.3	Guidance	25
5.7	Packages (child and library)	25
5.7.1	Evaluation	26
5.7.2	Notes	26
5.7.3	Guidance	26
5.8	Arithmetic Types	27
5.8.1	Evaluation	27
5.8.2	Notes	27
5.8.3	Guidance	28
5.9	Low Level and Interfacing	29
5.9.1	Evaluation	30
5.9.2	Notes	30
5.9.3	Guidance	31
5.10	Generics	31
5.10.1	Evaluation	32
5.10.2	Notes	32
5.10.3	Guidance	33
5.11	Access Types and Types with Dynamic Attributes	34
5.11.1	Evaluation	34
5.11.2	Notes	34
5.11.3	Guidance	35
5.12	Exceptions	35
5.12.1	Evaluation	36
5.12.2	Notes	36
5.12.3	Guidance	36
5.13	Tasking	37
5.13.1	Evaluation	39
5.13.2	Notes	39
5.13.3	Guidance	39
5.14	Distribution	40
5.14.1	Evaluation	40
5.14.2	Notes	40
5.14.3	Guidance	40
6	Compilers and Run-time Systems	40
6.1	Language issues	41
6.2	Compiler Qualification	41
6.3	Run-Time System	42
7	References	43
7.1	Applicable Documents	43
7.2	Referenced Documents	44

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.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 3.

In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. Draft International Standards adopted by the technical committees are circulated to the member bodies for voting. Publication as an International Standard requires approval by at least 75 % of the member bodies casting a vote.

In exceptional circumstances, when a technical committee has collected data of a different kind from that which is normally published as an International Standard ("state of the art", for example), it may decide by a simple majority vote of its participating members to publish a Technical Report. A Technical Report is entirely informative in nature and does not have to be reviewed until the data it provides are considered to be no longer valid or useful.

Attention is drawn to the possibility that some of the elements of this Technical Report may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

ISO/IEC TR 15942 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 22, *Programming languages, their environments and system software interfaces*.

Introduction

As a society, we are increasingly reliant upon high integrity systems: for safety systems (such as fly-by-wire aircraft), for security systems (to protect digital information) or for financial systems (e.g., cash dispensers). As the complexity of these systems grows, so do the demands for improved techniques for the production of the software components of the system.

These high integrity systems must be shown to be fully predictable in operation and have all the properties required of them. This can only be achieved by analysing the software, in addition to the use of conventional dynamic testing.

There is, currently, no mainstream high level language where all programs in that language are guaranteed to be predictable and analysable. Therefore for any choice of implementation language it is essential to control the ways that the language is used by the application.

The Ada language [ARM] is designed with specific mechanisms for controlling the use of certain aspects of the language. Furthermore,

1. The semantics of Ada programs are well-defined, even in error situations. Specifically, the effect of a program can be predicted from the language definition with few implementation dependencies or interactions between language features.
2. The strong typing within the language can be used to reduce the scope (and cost) of analysis to verify key properties.
3. The Ada language has been successfully used on many high integrity applications. This demonstrates that validated Ada compilers have the quality required for such applications.
4. Guidance can be provided to facilitate the use of the language and to encourage the development of tools for further verification.

Ada is therefore ideally suited for implementing high integrity software and this document provides guidance in the controls that are required on the use of Ada to ensure that programs are predictable and analysable.

All language design balances functionality against integrity. For instance, the ability to control storage allocation directly will impact the need to ensure the integrity of data. An aspect of the integrity of Ada programs is the possibility of avoiding access types (references) completely, whereas in other languages references are linked to array accessing and/or parameter passing, and therefore cannot be excluded.

There are a number of different analysis techniques in use for high integrity software and this document is not prescriptive about which techniques to use. Furthermore, each analysis technique requires different controls on the use of the language. Ada assists analysis: for instance, the modes of Ada parameters, suitably used, provide information for data flow analysis which other languages cannot always provide. This Technical Report, therefore, catalogues specific verification techniques (see 2.5), and classifies the impact that language features have on the use of these techniques (in the tables in Section 5).

It is the user's responsibility to select the analysis techniques for a particular application; this document can then be used to define the full set of controls necessary for using that set of techniques.

The guidance given here first specifies its scope, by reference to the safety and security standards to which high integrity applications may be written.

Section 2 then analyses the verification techniques that are applied in the development of high integrity systems. By this means, the regulatory rules of the standards for safety and security are abstracted to avoid the need to consider each such standard separately.

Section 3 addresses general issues concerning how computer languages must be constructed if programs written in that language are to be fully predictable. These issues are relevant to any restricted language defined through the application of this guidance.

Section 4 provides identification of a three-way classification system used for Ada language features. This classification is based upon the ease with which verification techniques can be applied to a program containing the feature. This classification is needed since while the majority of the core features in Ada assists verification, the use of certain features makes the resulting code difficult or impossible to analyse with the currently available program analysis tools and techniques.

Section 5 provides the main technical material of this Technical Report by classifying Ada language features. Users of this Technical Report can then determine which features of Ada are appropriate to use from the verification techniques that are to be employed. The assessment has shown that the vast majority of the Ada features lend themselves to effective use in the construction of high integrity systems.

The Technical Report concludes, in Section 6, by providing information to aid the choice of a suitable Ada compiler together with its associated run-time system.

References to relevant standards and guides are provided. A detailed analysis of Ada95 for high integrity systems is available in References [CAT1, CAT2] and [CAT3].

A comprehensive index is provided to ease the use of the Technical Report.

Levels of criticality

Many of the Standards to which high integrity software is written use multiple levels to classify the criticality of the software components which make up the system. While the number and nature of the levels vary, the general approach is always the same: the higher the criticality of the system, the more verification techniques need to be used for its assurance. Table 1 relates the various levels of classification used in some well known International Standards.

Table 1: Levels of criticality in some Standards

Standard	Number of levels	Lowest Level	Highest Level
[DO-178B]	4	D	A
[IEC-61508]	4	Safety Integrity Level 1	Safety Integrity Level 4
[ITSEC]	7	E0	E6

This Technical Report emphasizes the higher levels of criticality, for which the more demanding verification techniques are employed and for which Ada provides major benefits.

This Technical Report, however, does not directly use any such levels but focuses on the correlation between the features of the language and the verification techniques to be employed at the higher levels of criticality. The material in [ISO/IEC 15026], [DS 00-56], [ARP 4754] and [ARP 4761] may be useful in determining the criticality of a system if this is not covered by application-specific standards.

Readership

This Technical Report has been written for:

1. Those responsible for coding standards applicable to high integrity Ada software.
2. Those developing high integrity systems in Ada.
3. Vendors marketing Ada compilers, source code generators, and verification tools for use in the development of high integrity systems.
4. Regulators who need to approve high integrity systems containing software written in Ada.
5. Those concerned with high integrity systems who wish to consider the advantages of using the Ada language.

This Technical Report is not a tutorial on the use of Ada or on the development of high integrity software. Developers using this report are assumed to have a working knowledge of the language and an understanding of good Ada style, as in [AQS].

History

When proposals were made that a subset of Ada should be specified for high integrity applications, it was realized that the provision of the Safety and Security Annex in the Ada standard did not satisfy all the requirements of the developers of high integrity systems. In consequence, a group was formed under WG9 to consider what action was needed. This group, called the HRG, proposed and drafted this Technical Report over a three year period.

Conventions

In line with the Ada standard, the main text is in a Roman font. Ada identifiers are set in a sans-serif font, and the Ada keywords in a bold sans-serif font.

Postscript

The guidance provided here reflects the understanding of the issues based mainly on using the previous Ada standard in developing high integrity applications. Over the next few years, the current Ada standard will be used for further high integrity applications which will no doubt need to be reflected in a revision of this guidance. Specifically, further detail can be produced based upon the experience gained.

Instructions for comment submission

Informal comments on this Technical Report may be sent by e-mail to hrg@cise.npl.co.uk. If appropriate, the project editor will document the issue for corrective action.

Comments should use the following format:

!topic: Title which is a summary in one sentence
!reference TR 15942- ss.ss.ss
!from Author, Name, yy-mm-dd
!keywords keywords related to topic
!discussion
text of discussion

where *ss.ss.ss* is the section number, *yy-mm-dd* is the date. If the comment requests a change, a rationale for this and the substance of the actual change proposed would facilitate the processing of the comment.

Information technology - Programming languages - Guide for the use of the Ada programming language in high integrity systems

1 Scope

This Technical Report provides guidance on the use of Ada when producing high integrity systems. In producing such applications it is usually the case that adherence to guidelines or standards has to be demonstrated to independent bodies. These guidelines or standards vary according to the application area, industrial sector or nature of the risk involved.

For safety applications, the international generic standard is [IEC 61508] of which part 3 is concerned with software.

For security systems, the multi-national generic assessment guide is [ISO CD 15408].

For sector-specific guidance and standards there are:

Airborne civil avionics: [DO-178B]

Nuclear power plants: [IEC 880]

Medical systems: [IEC 601-4]

Pharmaceutical: [GAMP]

For national/regional guidance and standards there are the following:

UK Defence: [DS 00-55]

European rail: [EN 50128]

European security: [ITSEC]

US nuclear: [NRC]

UK automotive: [MISRA]

US medical: [FDA]

US space: [NASA]

The above standards and guides are referred to as Standards in this Technical Report. The above list is not exhaustive but indicative of the type of Standard to which this Technical Report provides guidance.

The specific Standards above are not addressed individually but this Technical Report is synthesized from an analysis of their requirements and recommendations.

1.1 Within the scope

This Technical Report assumes that a system is being developed in Ada to meet a standard listed above or one of a similar nature. The primary goal of this Technical Report is to translate general requirements into Ada specific ones. For example, a general standard might require that dynamic testing provides evidence of the execution of all the statements in the code of the application. In the case of generics, this is interpreted by this Technical Report to mean all instantiations of the generic should be executed.

This Technical Report is intended to provide guidance only, and hence there are no 'shall's'. However, this Technical Report identifies verification and validation issues which should be resolved and documented according to the sector-specific standards being employed.

The following topics are within the scope of this Technical Report:

- the choice of features of the language which aid verification and compliance to the standards,
- identification of language features requiring additional verification steps,
- the use of tools to aid design and verification,
- issues concerning qualification of compilers for use on high integrity applications,
- tools, such as graphic design tools, which generate Ada source code which is accessible to users.

Tools which generate Ada source code require special consideration. Where generated code may be modified or extended, verification of the extensions and overall system will be assisted if the guidelines have been taken into account. Even where modification is not planned, inspection and analysis of the generated code may be unavoidable unless the generator is trusted or 'qualified' according to an applicable standard. Finally, even if generated code is neither modified nor inspected, the overall verification process may be made more complicated if the code deviates from guidelines intended to facilitate testing and analysis. Potential users of such tools should evaluate their code generation against the guidance provided in this Technical Report.

1.2 Out of scope

The following topics are considered to be out of scope with respect to this Technical Report:

- Domain-specific standards,
- Application-specific issues,
- Hardware and system-specific issues,
- Human factor issues in the application (as opposed to human factors in the use of the Ada language which is in scope).

2 Verification Techniques

Verification is the confirmation by examination and provision of objective evidence that specified requirements have been fulfilled [ISO 8402: 2.18] .

There are currently four approaches required by applicable standards and guidelines (see Section 7.1) to support the verification of software:

1. traceability,
2. reviews,
3. analysis, and
4. testing.

Each one of these is discussed below. Where appropriate, language-specific techniques that support each approach are discussed. Finally, these techniques are grouped into categories that can form a basis for the analysis of Ada language features. This analytical approach forms the basis for the assessment presented in Section 5.

2.1 Traceability

Traceability is required to establish that the implementation is complete, and to identify new derived requirements. It occurs throughout the life cycle, e.g., there needs to be traceability from:

- lower level (decomposed) requirements to higher level requirements;
- test procedures to requirements, design, or code;
- object code to source code.

While traceability is not language specific, certain attributes of design or coding styles can aid in (or detract from) accomplishing this objective. For example, consider a single module that implements a single low level requirement, which has an associated single test procedure, then the method to support traceability is straightforward. On the other hand, if there is a many-to-many relationship between the various decomposed levels of software (because of design choices or test procedures), traceability can become very complicated. Deduction of completeness of implementation (without extraneous code) may therefore be difficult or impossible.

Additionally, the use of some of the more sophisticated language features of high-level languages that require extensive compiler generated code may detract from the straightforward translation into, and hence traceability of, the object code.

2.2 Reviews

Reviews are an important part of the verification process. They can be carried out on requirements, design, code, test procedures, or analysis reports. Reviews are conducted by humans and may be undertaken 'formally' such as in a Fagan inspection [11] or 'informally' such as in desk checks. Typically, reviews are done by an 'independent' person i.e., the producer of the artefact is different from the reviewer. This independence is a mandatory requirement of safety-critical software standards.

Coding standards and avoidance of certain language features of high-level languages are essential for high integrity systems in order to facilitate reviews. These aspects become important since the 'independent' code review may at times be conducted by an expert in the application domain who may not have detailed insight into language constructs and their interactions.

There is a distinct tradeoff between the use of language features to capture design intent and support abstraction versus the need for simplicity and predictability. This document leaves to coding guidelines, such as the Ada Quality and Style Guide [AQS], software engineering issues such as those concerned with readability and reusability. Those issues that are associated with quantifiable analysis are considered in the following sections.

2.3 Analysis

This Technical Report distinguishes between analysis (i.e., static analysis) and testing (i.e., dynamic analysis). Analysis supplements testing to establish that the requirements are correctly implemented.

Analysis can be performed on requirements, design, or code; the major emphasis of this Technical Report is the analysis of the design and code.

Described below are ten analysis methods which are required in different combinations by various standards.

1. Control Flow
2. Data Flow
3. Information Flow
4. Symbolic Execution
5. Formal Code Verification
6. Range Checking
7. Stack Usage
8. Timing Analysis
9. Other Memory Usage
10. Object Code Analysis

2.3.1 Control Flow analysis

Control Flow analysis is conducted to

1. ensure that code is executed in the right sequence,
2. ensure that code is well structured,
3. locate any syntactically unreachable code, and
4. highlight the parts of the code where termination needs to be considered, i.e., loops and recursion.

Call tree analysis, an example of one of the many Control Flow analysis techniques available, is used to verify that the sequencing stated by the design is correctly implemented. Also, call tree analysis can help detect direct and indirect recursion, which are prohibited by most high integrity standards. Furthermore, if a system is partitioned into critical and non-critical parts, then the call tree analysis can confirm that the design rules for partitioning have been followed.

Ada is rich in facilities for program flow control. Language rules, such as prohibition of modification of **for** loop control variables, make it difficult to produce poorly structured Ada code. If the **goto** statement is not used and relatively minor restrictions are made on placement of **exit** and **return** statements, Ada code becomes inherently well structured.

2.3.2 Data Flow analysis

The objective of Data Flow analysis is to show that there is no execution path in the software that would access a variable that has not been set a value. Data Flow analysis uses the results of Control Flow Analysis in conjunction with the read or write access to variables to perform the analysis. Data Flow analysis can also detect other code anomalies such as multiple writes without intervening reads.

In most general-purpose languages, Data Flow analysis is a complex activity, mainly because global variables can be accessed from anywhere, and because subprogram parameters do not support **out**-only modes. The job can be made significantly easier in Ada which has packages to contain potentially shared data, and **out** mode parameters on subprograms.

2.3.3 Information Flow analysis

Information Flow analysis identifies how execution of a unit of code creates dependencies between the inputs to and outputs from that code. For example:

```
X := A+B;
Y := D-C;
if X>0 then
  Z:=(Y+1);
end if;
```

Here, X depends on A and B, Y depends on C and D, and Z depends on A, B, C, and D (and implicitly on its own initial value).

These dependencies can be verified against the dependencies in the specification to ensure that all the required dependencies are implemented and no incorrect ones are established. It can be performed either internal to a module (i.e., a procedure or a function), across modules, or across the entire software (or system). This analysis can be particularly appropriate for a critical output that can be traced back all the way to the inputs of the hardware/software interface.

2.3.4 Symbolic Execution

The objective of Symbolic Execution is to verify properties of a program by algebraic manipulation of the source text without requiring a formal specification. This technique is typically applied using tools that also undertake Control Flow, Data Flow and Information Flow analysis.

Symbolic Execution is a technique where the program is 'executed' by performing back-substitution; in essence, the right hand side of each assignment is substituted for the left hand side variable in its subsequent uses. This converts the sequential logic into a set of parallel assignments in which output values are expressed in terms of input values. Conditional branches are represented as conditions under which the relevant expression gives the values of the outputs from the inputs. To undertake this

computation, it is assumed that no aliasing is taking place, i.e., two variables X and Y do not refer to the same entity and that functions have no side-effects. Tools that provide support for symbolic execution may or may not check for these conditions.

Using the fragment of Ada code which illustrated Information Flow analysis gives:

```

A+B ≤ 0 :
  X = A+B
  Y = D-C
  Z = not defined on this path (retains initial value)
A+B > 0 :
  X = A+B
  Y = D-C
  Z = D-C+1

```

These algebraic expressions give the output in terms of the input and can be compared (manually) with the specification of a subprogram to verify the code.

Symbolic Execution can also be used to assist with reasoning that run-time errors will not occur (e.g., Range Checking). The Symbolic Execution model is extended to include expressions indicating the conditions under which a run-time error may occur. If these expressions are mutually contradictory for a particular execution path then that path is free from potential run-time errors. In the above example, if Z has an upper declared bound of 10, then we have the condition that if $A+B > 0$, then $D-C+1 < 11$ to avoid raising the exception on the upper bound.

2.3.5 Formal Code Verification

Formal Code Verification is the process of proving that the code of a program is correct with respect to the formal specification of its requirements. The objective is to explore all possible program executions, which is infeasible by dynamic testing alone.

Each program unit is verified separately, against those parts of the specification that apply to it. For instance, Formal Code Verification of a subprogram involves proving that its code is consistent with its formally-stated post-condition (specifying the intended relationships between variables on termination), given its pre-condition (specifying the conditions which must apply when the subprogram is called). A more restricted proof aimed at demonstrating a particular safety/security property can also be constructed.

The verification is usually performed in two stages:

1. Generation of Verification Conditions (VCs). These theorems are proof obligations, whose truth implies that if the pre-condition holds initially, and execution of the code terminates, then the post-condition holds on termination. These VCs are usually generated mechanically.
2. Proof of VCs. Machine assistance in the form of a suitable proof tool can be used to discharge verification conditions.

The process outlined above establishes *partial correctness*. To establish *total correctness* it is also necessary to prove *termination* of all loops when the stated pre-condition holds and termination of any recursion. Recursion is not normally permitted in high integrity systems. Termination is usually demonstrated by exhibiting a *variant* expression for every loop, and showing that this expression gives a non-negative number that decreases on each iteration. *Termination conditions* can be generated and proved, similarly to the generation and proof of verification conditions.

The value of Formal Code Verification depends on the availability of a specification expressed in a suitable form such as results from formal specification methods. Formal methods involve the use of formal logic, discrete mathematics, and computer-readable languages to improve the specification of software.

Proof of absence of run-time errors

In some real-time high integrity systems, occurrence of run-time errors is not acceptable. An example is the flight-control system of a dynamically unstable aircraft, in which there would not be time to recover from such an error. The techniques of Formal Code Verification described above can be used to prove that (with appropriate language constraints) certain classes of run-time errors, e.g., range constraint violations, cannot arise in any execution.

To perform such verifications, the object type and variable declarations are used to construct pre-conditions on the ranges of initial values, and at each place in the source code where a run-time check would be produced, an assertion formally describing

the check is generated. From these pre-conditions, assertions and the program code, verification conditions are mechanically produced.

These verifications (or 'proof obligations') are numerous, but for the most part simple enough to be proved automatically. Full formal requirement specifications are not needed to apply this technique.

2.3.6 Range Checking

The objective of this analysis is to verify that the data values lie within the specified ranges as well as maintain specified accuracy. These forms of analysis include, but are not limited to,

1. overflow and underflow analysis,
2. rounding errors,
3. range checking, and
4. array bounds.

For discrete types, the static bounds placed upon variables often allow many cases to be checked automatically. When enumerated types are used instead of integer types, these checks are more effective. For real types, the need to show the absence of overflow is more demanding than the analysis of operations on discrete types.

Since the semantics of Ada remain defined even in error conditions, the necessary checks can be explicitly specified. Furthermore, the 'Valid' attribute makes it straightforward to check that scalar data, especially where it is obtained from sources external to the program, is a legal Ada value, without the risk of run-time exceptions being generated.

2.3.7 Stack Usage analysis

The stack is a part of the memory shared by different subprograms and used for storing data local to the subprogram, temporary data and return addresses generated by the compiler. Stack Usage analysis is a particular form of shared resource analysis that establishes the maximum possible size of the stack required by the system and whether there is sufficient physical memory to support this maximum stack size. Also, some compilers use multiple stacks and this analysis needs to be conducted for each stack.

Another aspect of Stack Usage analysis is to ensure that there is no stack-heap collision at run-time. This analysis is avoided when dynamic heap allocation is prohibited.

Stack Usage analysis is made simpler for a programming language such as Ada where subprogram calls and return semantics are unambiguous, and where there is a clean distinction between static and dynamic types. Compilers supporting the Safety and Security Annex provide the necessary information to undertake this analysis, see [ARM: H.3.1(15)].

2.3.8 Timing Analysis

The overall objective of this analysis is to establish temporal properties of the input/output dependencies. A common and important aspect of this analysis is the worst-case execution time for the correct behaviour of the overall system.

Certain programming language features or design approaches make timing analysis difficult, e.g., loops without static upper bounds and the manipulation of dynamic data structures.

The static typing of Ada and the unambiguous semantics of its control structures facilitate these analyses. Also, the **pragma Reviewable** and **pragma Inspection_Point** ensure that there is traceability from the source code to the object code to facilitate timing analysis.

2.3.9 Other Memory Usage analysis

This analysis is required for any resource that is shared between different 'partitions' of software. These forms of analysis include, but are not limited to, memory (heap), I/O ports, and special purpose hardware, that perform specific computations or watchdog timer functions.

Other Memory Usage analysis will show the absence of interference between Ada and other components such as low-level and hardware device drivers and resource managers. In particular heap memory should usually be avoided and IO devices rigorously partitioned. Ada is particularly useful when doing such analysis since the **pragma Restrictions (No_Allocators)** can

be used to ensure no explicit use of the heap and **pragma** Restrictions (No_Implicit_Heap_Allocation) to ensure no implicit usage of the heap.

2.3.10 Object Code Analysis

The purpose of Object Code Analysis is to demonstrate that object code is a correct translation of source code and that errors have not been introduced as a consequence of a compiler failure.

This analysis is sometimes undertaken by manual inspection of the machine code generated by the compiler. The compiler vendor may provide details of the mapping from the source code to object code so that manual checks are simpler to undertake. Unfortunately, it is not currently within the state of the art to formally verify the equivalence of source code and the generated object code.

The Ada **pragma** Reviewable provides basic information to assist in tracing from source code to object code. **Pragma** Inspection_Point can be used to determine the exact status of variables at specific points. For the requirements of such an analysis, see [15].

2.4 Testing

2.4.1 Principles

Testing (sometimes known as dynamic analysis) is the execution of software on a digital computer, which is often the target machine on which the final application runs. Testing has the advantage of providing tangible, auditable, evidence of software execution behaviour [14].

There are many testing techniques and new ones are being invented continually. This section is limited to those procedures that are required by various software standards. It is not intended to be an exhaustive encyclopaedia of the various testing techniques known at the present time.

Testing can be performed at various levels of software (and system):

- Software module level (individual procedures or functions),
- Software integration testing (i.e., module integration testing),
- Hardware/Software integration testing, and
- System testing.

The testing procedures described below focus on the first two aspects, i.e., module and module integration testing, since the choice of programming language has a direct impact on the ease or difficulty of testing. Within this framework, there are two basic forms of testing:

- Requirements-based (or black-box) Testing, and
- Structure-based (or white-box) Testing.

Since exhaustive testing is infeasible for any realistic program, one approach is to partition the data domain into equivalence classes and their boundary values in order to limit the number of test cases.

2.4.2 Requirements-based Testing

The Requirements-based Testing methods aim to show that the actual behaviour of the program is in accordance with its requirements. For this reason, these methods are sometimes also called 'functional testing' or 'black-box testing'. This is to highlight the fact that the program structure is not taken into account. There are two common methods for conducting Requirements-based Testing:

Equivalence class testing

The inputs and outputs of the component [BS 7925-1: 3.42] under test are divided into equivalence classes in which the values within one class can reasonably be expected to be treated by the component in the same way. The equivalence class for numeric data is a range having the same sign or zero. For data of an enumerated type, each value usually forms a class, since each value could be expected to be treated differently. For composite types, the equivalence classes are obtained by combining

the classes derived from the components of the type. Testing is then undertaken using a sample from each equivalence class [BS 7925-2: 5.1].

Boundary value testing

This approach enhances the equivalence class testing by requiring testing with values at the boundary of the specified range. Additionally, 'stress' testing or 'robustness' testing may be undertaken outside the specified range if required by the application domain standard [BS 7925-2: 5.2].

2.4.3 Structure-based Testing

The objective of these testing methods is to increase the confidence in software by exercising the program beyond the Requirements-based Testing mentioned above.

Examples of the methods used, and the definitions implied, are:

Statement Coverage

The application of test cases such that every statement in the program has been invoked at least once.

Decision (Branch) Coverage

The application of test cases such that every point of entry and exit in the program has been invoked at least once, and every decision in the program has taken all possible outcomes at least once.

Modified Condition/Decision Coverage

The application of test cases such that every point of entry and exit in the program has been invoked at least once, every condition in a decision in the program has taken all possible outcomes at least once, every decision in the program has taken all possible outcomes at least once, and each condition in a decision has been shown to independently affect that decision's outcome. A condition is shown to independently affect a decision's outcome by varying just that condition while holding fixed all other possible conditions.

These Structure-based Testing approaches apply only to source code of high-level languages such as Ada. In addition, standards for the development of safety-critical software often require that the set of tests are shown to execute every instruction in the object code of the software under test.

2.5 Use of Verification Techniques in this Technical Report

When considering the previous verification techniques, certain groupings can be formed. The verification techniques within a group have similar properties with respect to the complexity of that technique when applied to a particular language feature. The following table shows the groups, and abbreviations, used in the Assessment of Language Features (Section 5).

Table 2 : Verification Techniques

Approach	Group Name	Technique
Analysis	Flow Analysis (FA)	Control Flow
		Data Flow
		Information Flow
	Symbolic Analysis (SA)	Symbolic Execution
		Formal Code Verification
	Range Checking (RC)	Range Checking
	Stack Usage (SU)	Stack Usage
	Timing Analysis (TA)	Timing Analysis
	Other Memory Usage (OMU)	Other Memory Usage
Testing	Object Code Analysis (OCA)	Object Code Analysis
	Requirements-based Testing (RT)	Equivalence Class
		Boundary Value
	Structure-based Testing (ST)	Statement Coverage
		Branch Coverage
		Modified Condition/Decision Coverage

For convenience, some techniques are grouped together as follows:

Flow Analysis

The three methods under this group name are Control Flow, Data Flow and Information Flow (Analysis). The basic methods of control flow and data flow are combined with information flow analysis.

Symbolic Analysis

Two methods are grouped together here: Formal Code Verification and Symbolic Execution. Both these methods share a mathematical foundation and typically use tools which analyse the source text of a program algebraically.

Requirements-based Testing

The Equivalence Class and Boundary Value are closely related in that Boundary Value testing extends the set of test cases required by Equivalence Class testing.

Structure-based Testing

Three testing methods are combined here, which in order of reducing severity are: Modified Condition/Decision Coverage, Branch Coverage and Statement Coverage. These methods are closely related and are often tested with the help of the same tool.

3 General Language Issues

3.1 Writing Verifiable Programs

Standards and guidelines for high integrity software demand assurances of its fitness for purpose, based on its static analysis (see Section 2.3), and also testing (see Section 2.4).

The choice of the verification techniques to be used for a software development is one part of the definition of the overall development process for that software. Coding standards and language subsets can only be defined after the analysis methods have been selected, to ensure that none of the analysis methods are compromised by the use of inappropriate language features.

In the context of high integrity systems, fitness for purpose depends on the behaviour of a system being predictable, and demonstrably conforming to a precise specification of its requirement. Historically, demonstration of conformance to requirement specifications by static analysis and testing has usually been attempted retrospectively, after a system has been developed. This approach has rarely proved a cost-effective way of producing the desired results, for several reasons. Firstly, because the design and implementation processes have not been chosen with verification in mind. Secondly, the static analysis or testing technologies have been difficult and costly to apply, and were not very revealing. And finally, where such analyses have revealed errors, these were detected so late that their correction was extremely costly.

The high cost of repeated testing as mistakes are uncovered and rectified, and improvements in the technology of static analysis, have led to a realization that there are significant cost benefits in performing static code analysis as the code is produced. This accords well with the movement, as specification methods have improved, towards 'Correctness by Construction' - the 'lean engineering' view of software development. Here static analysis methods, where they are relevant, feature as an integral part of the development process. Their purpose is to establish the validity (in particular respects) of intermediate products before they are used to develop further ones. This constructive approach also contrasts with retrospective verification in matters of style: where certain language features are positively required in order to support a particular model. The retrospective approach leads to a list of language prohibitions. Although the constructive approach is more demanding, conceptually, it uses language and tools to better effect, improving technical quality as well as reducing costs.

With motivations such as these for performing verification, it can be seen that there are four different reasons for needing or rejecting particular language features within this context:

1. Language rules to achieve predictability,
2. Language rules to allow modelling,
3. Language rules to facilitate testing,
4. Pragmatic considerations.

In addition, there are reasons to enhance the language by adding forms of annotations.

In the following sections these matters are considered in detail.

3.1.1 Language Rules to Achieve Predictability

Language rules to achieve predictability are independent of the analysis methods to be used. It is a requirement, in all high integrity system development, that the program source code is unambiguous. All possible forms of language ambiguity must be prevented in some way.

The ways in which programs written in a high-level language can be ambiguous are well-known. A few cases are discussed below.

Side-effects

It is generally accepted that side-effects in functions are undesirable. Side-effects can make order of evaluation of an expression a significant issue, and can make repeated calls to a function return different results. Many programming languages permit functions to have side-effects.

Effects of elaboration order

Programs in languages that support default initialization or elaboration will be dependent on the order of evaluation within a unit and the order in which units are evaluated.

Effects of parameter-passing mechanisms

Parameters to a subprogram in a high level language can be passed as a copy or passed as a reference. The results of an execution may depend upon which method is chosen by an implementation (where a choice is permitted).

Such ambiguities make the formal analysis of a program difficult. The uncertainties can be reduced by writing programs that do not depend on these types of issues.

There is often an attempt in high integrity systems to eliminate offending language features to remove the uncertainties or implementation dependencies. If taken to its extreme this can lead to a language that is 'safe' but ineffective. It will lack the expressive power required to tackle the application requirements. The language may be 'safe' but the programming process becomes more error-prone as complex designs must be produced to circumvent the inadequacies of the language.

3.1.2 Language Rules to Allow Modelling

All static analysis methods involve the construction of a model of program source text (or object code), and the application of algorithms or reasoning processes to this model, to check whether it has some particular properties. It is possible that if a particular feature of a language is used at all, or in a particular way, the model would be unable to capture those aspects of the program text that are important to the analysis to be performed. As a consequence, to ensure validity of a static analysis method some language features must be disallowed or their use rigorously controlled.

Assuming that language rules, or other devices, exist to ensure a program is meaningful, this section considers the requirements for the various forms of modelling and analysis methods to be well founded and, at least in principle, applicable.

The simpler kinds of static analysis, in particular Control Flow, Data Flow, and Information Flow analysis, essentially involve the study of paths (corresponding to execution histories) on a directed graph model of a program (effectively a flowchart). The arcs of the program graph have labels which capture some syntactic and semantic properties of the program statements. For Control Flow analysis only the graph structure is required, for data flow analysis each arc label specifies which variables the corresponding statement reads and updates, for Information Flow analysis the dependencies between read and updated variables must be defined more precisely. Conceptually, Formal Code Verification (and proof of absence of run-time errors) is also based on paths (or execution histories) on such a graph, but here the arc labels are partial functions of the set of program variables, defining the conditions for traversing the arcs and the transformations of variables performed in doing so. These static analysis techniques are all applied to the same structural model, but form a progressively richer description of the program objects and elements.

The inclusion of exception flows and run-time dispatching significantly complicates the graph structure and tends, in the general case, towards the analysis becoming intractable; the **goto** statement also introduces difficulties. Otherwise, there are no obstacles to Control Flow analysis, which only requires and provides relatively limited information.

Data Flow analysis begins to take an interest in the data objects of a program. When using this method difficulties of distinguishing between dynamically-selected components of composite objects are first encountered. Two issues are: arrays must be treated as entire objects (although their components, selected by dynamically-computed indices, can reappear in Formal Code Verification), and dynamic creation of objects in the course of execution (using access types) cannot be represented. Information Flow analysis essentially requires the same constraints.

Such static models of programs only apply to sequential code. If the program contains concurrent threads (e.g., tasks), then each thread must be analysed independently in terms of its flow, with the concurrency aspects being addressed by different models and reasoning processes. With a concurrent system the key objective of this additional modelling is to prove that the system cannot enter undesirable states such as deadlock. The main static techniques for achieving this are finite state automata [3], petri-nets [4, 5] and process algebras [6-8] using model checking or mathematical proof. All try to construct the set of all feasible histories for the system to show that unsafe conditions cannot be reached. The degree to which thread synchronization impacts on the flow analysis of individual threads determines the complexity of the overall verification problem. The more asynchronous the model, the more straightforward the flow analysis of each thread. If there is only asynchronous data passing between threads then each thread can be verified in isolation. Moreover, the complete timing/scheduling analysis of the entire concurrent system becomes tractable.

Surprisingly perhaps, the domain to which Formal Code Verification is applicable in principle is no more restricted; indeed, it is somewhat larger, as composite objects can be considered more generally. There is a large gap, however, between what is analysable in principle and what can be analysed in practice and a number of strategic choices arise. If the dynamic semantics of the language are made as simple as reasonably possible, the burden of verification is placed as far as possible on static analysis tools. So, for example, if a program is deemed exception free by formal proof, coverage testing is significantly simplified.

In conclusion, the adequacy of modelling for different analysis methods gives rise to some language constraints. But just as important in determining how a programming language should be used, are the precise reasons for using the analysis methods and the ways in which they are to be employed — in other words, the choice of software development process.

3.1.3 Language Rules to Facilitate Testing

Demonstration of a program's correct dynamic behaviour, by testing, is an important part of its verification. In practice the development of high integrity systems usually involves two different forms of testing: investigative, informal or debugging testing; and formal testing of a kind or kinds required by various standards. As noted above, the detection and elimination of errors by static analysis methods early in the project life cycle can greatly reduce testing costs by shifting the emphasis of testing from debugging to demonstrating correctness.

The constraints placed on programming languages to facilitate testing are generally less significant than those demanded by static analysis techniques. Although language features do have some impact on testing, the use of static analysis techniques as part of the development process will ensure that language features which might complicate such testing will be avoided.

Language features that obstruct coverage testing are those such as dispatching which introduces very dynamic control flow; **goto** which complicates control flow; and features which complicate the view taken of data, for example: variant records, unchecked conversions and dynamic pointer usage. Predefined exceptions present a particular difficulty for coverage testing because it is often not possible to reach all the paths these introduce by external stimulation of unmodified code.

Language features that assist testing are those which constrain data values by strong and static typing and those which assist in locating errors by indicating their presence close to the point where they arise; here predefined exceptions are beneficial.

There is a general tension between the dictates of good software engineering, which encourage information (or more accurately "detail") hiding, and testing which often seeks to monitor the values of internal system state. Language features which enable these conflicting requirements to be reconciled are of particular value and avoid distorting the design just to achieve testability. Ada has particular strengths in this area which often allow test harnesses to have access to "hidden" data in non-intrusive ways. Examples include:

- A child package may be used to monitor state in the private part of its parent without requiring any change to that parent package.
- Subunits may be used to place the source for an embedded subprogram in a separate file thus allowing the construction of a suitable test driver for it. Test point subprograms in packages can also usefully be placed in a subunit and replaced with a null subprogram in delivered code.
- Parameterless functions may be used to return the value of "trimming variables" which are often found in control systems. These variables behave as constants at the program level but can be dynamically adjusted by direct memory access during rig testing. The use of parameterless functions to return their value reconciles the software's view that they are constants with the testing need to adjust their values.

- The use of a package to contain "test point" variables. These are "write-only" variables used solely for test monitoring purposes. At the point where the (hidden) value to be monitored is generated a copy can be passed to this package and used to update its state; this is preferable to distorting the design by making the data directly visible and emphasizes the distinction between state needed by the software for its correct behaviour and that introduced for testing purposes only.

3.1.4 Pragmatic Considerations

The considerations of side-effects and elaboration order described in section 3.1.1 impose conditions essential to the validity and technical relevance of static analysis methods. Further linguistic issues are important to the tractability and eventually the economic viability of their application.

The questions of what conditions make the static analysis methods easy to apply, and what makes their results meaningful and useful, have essentially the same answers: a program should be well-designed.

All the static analysis and testing methods and construction techniques are most effective when software is well-structured, with every module having a single entry point and a single exit point. Although data flow analysis can still be applied to 'spaghetti code', it will be less efficient in finding data flow errors and anomalies. Similarly, information flow analysis can be extended to 'spaghetti code', for which the unsurprising outcome is that almost every variable may depend on every other. In constructing proofs, it is preferable to have building blocks with pre-conditions on their single entry points and post-conditions defining state at their single exits. For all these reasons, language subsets for high integrity programming invariably exclude **goto** statements, return statements from procedure subprograms and apart from at the top most level, exception handlers are also often avoided.

Proving correctness of programs, like reasoning about them informally, is about relating fragments of code to fragments of specification. Decomposition of functionality is very important. Verifiability depends on simplicity of contextual information, i.e., scope and visibility should be as limited as reasonably possible. So it is important to pay attention to architecture, the location of state, access to state, and embedding.

For such reasoning to be practical, the use of two-valued logic is preferred, where predicates are either true or false but not undefined. To achieve this it is useful for all variables to have valid values. This implies that there should not be undefined variables, a condition that can be met in three different ways:

- by language constraints, e.g., imposing initialization of all variables in their declarations;
- by using a data flow analyser, to ensure that no variable is ever read before being initialized;
- by a tool-assisted software development process that ensures proper initializations.

Of these options, the latter two are preferred, since it is usually better for all initializations to be meaningful.

The conclusion is that for high integrity software development, the language employed is more than a pick-and-mix selection of language features. It must be built as a coherent whole as a means of expressing a chosen development philosophy that is compatible with tools used to implement it.

3.1.5 Language Enhancements

A programming language should be used in a style that enhances its fitness for purpose. The rules discussed above for predictability and modelling are restrictions on the language: the elimination of syntactic forms or other, more subtle features. Fitness can also be enhanced by adding elements to the standard language, to facilitate derivation or verification of a program given its specification.

Tools and techniques that relate the source code to a formal specification, such as Formal Code Verification, often need additional information to the program source in order to work efficiently. Such information describes program properties that follow from the laws of the language and that the programmer intuitively believes but are not expressed directly in the language syntax. Examples include loop invariants, relationships between formal parameters, and state hidden inside other service-providing packages¹. Capturing this information enables proofs that are otherwise computationally hard.

¹ It is a design goal of Ada that a client routine should not access state variables in another package (to prevent the client abusing them, among other good reasons). However knowledge of the state is needed to reason about the behaviour of the whole program.

One approach to verification is to capture these properties by embedding annotations (formal comments, with their own syntax and semantics) in the source code, together with static design rules that relate the annotations to program objects. An annotation can express an invariant as an equivalence between program objects, and can reveal hidden state in packages without compromising their integrity.

The literate programming paradigm [17] provides another way of relating source code to specification. In literate programming, program and program fragments are embedded within a larger document, which also contains natural language description and other formal notations. One can relate and check objects in different fragments, and assemble a complete program for compilation.

Other development paradigms involve automated generation of package templates from a specification. As well as the package code, the toolkit may supply additional information to the programmer, to exclude unwanted interactions across module boundaries.

All the above approaches facilitate correctness by construction, by narrowing the gap between a package or component and its specification, and by enabling its verification in a stand-alone manner at any time during development.

3.2 The Choice of Language

Any language that is to be used for the implementation of high integrity systems should:

- be strongly typed,
- support a range of static types,
- have a consistent semantics that is defined in an international standard,
- support abstractions and information hiding,
- have available validated compilers.

Ada is unique in its compliance to all these attributes. Nevertheless it may be necessary to restrict or prevent the use of certain features to achieve full predictability, and to allow all the forms of static analysis and testing considered important. However verification poses few problems with the main part of the language [10].

Rules of Ada usage are determined by considerations of the elimination of ambiguity, the feasibility of modelling and analysis, and the constructive use of static analysis. In some cases it is found that a language feature is undesirable (e.g., **goto** statements), or that it renders some kind of static analysis intractable (e.g., predefined exceptions), or infeasible (e.g., full tasking). More frequently the difficulty stems from the use of several features in combination in a particular manner. The most satisfactory way of applying appropriate restrictions in this case is through annotations, and rules between these and the Ada code. The annotation system requires very careful design if it is to be secure. The most important concern is with ways of using Ada that provide good implementations of specifications and designs, and that can to a large extent be rigorously verified as they are constructed. Experience indicates that when these conditions are met, testing is also greatly simplified.

Whether the restrictions are based on language-defined or implementation-defined restrictions, tool-based analysis or annotation-based analysis, the result is a subset of the entire language, or possibly a collection of subsets of the language. It would be ideal if a single subset could be defined which would satisfy all requirements. The reality of different user communities, each with their own regulatory and commercial pressures, and multiple levels of criticality make this ideal impossible to achieve. Since a single set of restrictions cannot be defined, this Technical Report provides detailed guidance that assists users in constructing their own restrictions based upon the verification techniques that they require.

Consideration of the language requirements for correctness by construction is a matter of positive choices. It may lead to some restrictions on the use of Ada, but these should be based on sound engineering judgement of the most appropriate combination of language features. Here the aim is to use Ada as the vehicle to support the chosen design and development paradigms, in the course of code production — which may for example suggest a particular way of using packages to implement object templates.

To perform all necessary constraint checks as program components are constructed, it is necessary to utilize specification and design information, and relate it to the program architecture and code. To do this, tools such as [16, 18] use a system of annotations ('formal Ada comments'), as discussed in 3.1.5. Introduction of an annotation system may also bring with it some stylistic restrictions on the use of Ada, for instance to reduce overloading and limit visibility. Annotations may also be necessary to fully support modular development. For example, to reason about the process performed by a module, knowledge

is required of precisely what information it is accessing, and what are its effects, both direct and indirect. Specifically, its use of, or action on, a global variable cannot be ignored.

The correctness by construction approach, with correctness checking module-by-module, may also require rules to avoid unwanted properties arising from the incremental development. For instance, it must be possible to prevent recursion (usually frowned upon in safety-critical applications) by rules applicable at the module level, rather than by checking for recursion when a system is complete.

It is reiterated that Ada is currently the only viable language with sufficient industrial heritage that can provide the framework for static analysis and correctness by construction. As indicated above some restrictions on the use of the language are necessary. Section 5 of this Technical Report provides a detailed assessment of these features.

4 Significance of Language Features for High Integrity

Language features are classified to facilitate the cost-effective application of the required verification techniques.

4.1 Criteria for Assessment of Language Features

Control of the language features used in the implementation of high integrity systems is one essential ingredient to the use of a high order language. Ada was expressly designed to facilitate this effort. Yet, like all modern general purposes languages, Ada offers a broader variety of features than strictly required for any specific application domain. Hence, important design decisions need to be made to determine which language features best suit the implementation and verification requirements that emanate from the specific application domain.

The reasoning behind this can be shown by an example. High integrity control systems often handle physical data that continuously varies. The processing involves the input/output of these quantities from/to digital to analogue converters. Within the program, the natural way to handle such data is by means of an Ada real type (fixed or floating point). The verification process for fixed point is slightly different to that for floating point. Hence, a design decision should be made as to which form of real type should be used. In doing so, one should also consider that real types are only suitable for modelling continuously varying quantities and, therefore, are not needed at all in some types of applications.

The choice of language features is determined by the nature and criticality of the application and the verification techniques to be employed. In practice, however, the availability of software tools may further constrain the choice of language features. For instance, if formal verification with proof is required, then the use of any real types might not be viable so that the application might have to be coded using integer types only.

This Technical Report rates language features using three categories as follows:

Included

A feature is 'included' if it is directly amenable to the designated verification technique. In fact, most Ada features are rated 'included' for most verification techniques. Included features enable the analysis to be undertaken and directly support the production of high integrity code.

Allowed

A feature is 'allowed' if the designated verification step is not straightforward, but is still achievable; or if the use of the feature is necessary and the use of the problematic verification technique can be effectively circumvented.

Excluded

A feature is 'excluded' if there is no current cost effective way of undertaking the designated verification technique. Assurance of exclusion requires some form of verification.

Even without the excluded and allowed features, Ada remains a rich language of great expressive power. In particular, all the features needed to support large scale, effective software engineering practice such as abstraction, encapsulation and concurrency are retained.

4.2 How to use this Technical Report

The user of this Technical Report should proceed in four steps, as follows:

1. Determine the verification techniques required from the relevant application specific standards or guidelines.
2. Identify and understand the objectives to be satisfied by each of the verification techniques.
3. Using the tables in Section 5, determine the actual rating of the language features.
4. Confirm that the resulting choice of subset and the additional verification steps for any allowed features can satisfy the programming and verification requirements. This step should take into account available tools.

In some situations a verification technique is hard to apply because of the interaction of two or more language features. When such an interaction occurs, it is addressed under the feature that is primarily responsible for the problem. For example, the use of representation clauses has impact on a number of features; these issues are all considered in the low level programming section (see 5.9).

5 Assessment of Language Features

The Ada features are split into fourteen groups closely related to chapters of the ARM. These are:

- types with static attributes;
- declarations;
- names, including scope and visibility;
- expressions;
- statements;
- subprograms;
- packages (child and library);
- arithmetic types;
- low level and interfacing;
- generics;
- access types and types with dynamic attributes;
- exceptions;
- tasking;
- distribution.

Each group of features is assessed separately. For each group of features assessed, this Technical Report includes one evaluation table and one textual section providing constructive guidance to use of the designated features.

The table indicates how a given language feature performs with respect to each of the nine groups of verification techniques identified in Table 2 (reproduced below for ease of reference). The rating 'included' is denoted by the abbreviation 'Inc' (plain faced) in the relevant entry; the rating 'allowed' by the abbreviation 'Alld' (bold faced); and the rating 'excluded' by the abbreviation 'Exc' (bold faced). Some of the ratings, including all marked as allowed or excluded, are accompanied by explanatory notes providing the rationale for the assessment. If it is necessary to make a global statement about the feature then a note is attached to the feature itself. Where a feature has the rating 'excluded' against a number of verification techniques then we assume it will not be used in high integrity systems. The assessments presented are thus simplified by not including all possible interactions with such features.

Table 2 (copied for reference) : Verification Techniques

Approach	Group Name	Technique
Analysis	Flow Analysis (FA)	Control Flow
		Data Flow
		Information Flow
	Symbolic Analysis (SA)	Symbolic Execution
		Formal Code Verification
	Range Checking (RC)	Range Checking
	Stack Usage (SU)	Stack Usage
	Timing Analysis (TA)	Timing Analysis
	Other Memory Usage (OMU)	Other Memory Usage
Testing	Object Code Analysis (OCA)	Object Code Analysis
	Requirements-based Testing (RT)	Equivalence Class
		Boundary Value
	Structure-based Testing (ST)	Statement Coverage
		Branch Coverage
		Modified Condition/Decision Coverage

Predefined Language Environment

The Technical Report gives no explicit consideration to the predefined language environment. If components of the environment are used in a high integrity application, then three situations arise:

1. The components are written in Ada. In this case, the guidance in Section 5 applies.
2. The components are part of the run-time system. In this case, the guidance in Section 6 applies.
3. If neither of the above applies, then no guidance is provided.

5.1 Types with Static Attributes

The strong typing mechanism in Ada is a significant contribution to software engineering. All values are associated with the type that is appropriate to their domain, and the type definition covers precisely the set of applicable values. This section is only concerned with types and subtypes that have static bounds and attributes, and that are statically allocated. Moreover, this section ignores interactions with Representation Clauses, which are dealt with in Section 5.9.

Ada95 has added the notion of statically matching subtypes to Ada83. When a compiler determines that a subprogram completion, such as a subprogram body or a renames of a subprogram, does not match the static subtype profile of the specification, it reports an error. In general, the concept of static expressions and subtypes enables many erroneous programs to be detected at compile-time rather than run-time. For example the component types of array types must statically match during conversion. Ada 95 enhances the already-rigorous static-matching rules of Ada 83.

The rules for the derivation of types, explicit and implicit conversion between types, and mechanisms for the derivation, extension and overriding of primitive subprograms are complete, consistent and beneficial to programming of high integrity systems. Values associated with one type must be explicitly converted to other types. This forces designers to be explicit about conversions and makes the conversions visible in the source code. In addition, conversions can occur only between types that are closely related (for example all numeric types, or types that have been derived from the same parent type). The name-based typing and derivation support software engineering principles and prevent many classes of errors.

Abstract types and subprograms provide significant capability to Ada at virtually no run-time cost. A tagged type that is declared abstract prevents objects of that type from being declared, although type derivation (or extension) is permitted, and these new types may then have objects and actual primitive operations. When a tagged type is extended, abstract primitive operations ensure that the extender of the type provides *real* code for each primitive operation, and does not accidentally use default or incomplete operations. Abstract subprograms of non-tagged types allow the definer of the type to make unavailable certain operations that would ordinarily be available for the type.

5.1.1 Evaluation

Table 3 : Types with Static Attributes

Feature	Group Name								
	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Anonymous Types	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Subtypes ¹	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Enumerated Types	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Character	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Boolean	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Integer	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Derived	Inc	Alld ²	Inc	Inc	Inc	Inc	Alld ²	Inc	Inc
Arrays	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Records	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Discriminated Records	Inc	Exc ³	Inc	Inc	Inc	Exc ³	Alld ³	Inc	Exc ³
Tagged Types without 'Class	Inc	Alld ²	Inc	Inc	Inc	Inc	Alld ²	Inc	Inc
Class Wide Operations	Exc ⁴	Exc ⁴	Inc	Inc	Alld ⁴	Inc	Alld ⁴	Inc	Exc ⁴
Abstract Types & Subprograms	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc

5.1.2 Notes

1. The use of 'Image for any subtype should be avoided. 'Image returns an unconstrained string which may require unbounded memory techniques for implementation.
2. Derivation of a type causes an automatic overloading of all primitive subprograms of the type. This makes a complete analysis harder, but still tractable since the effects are limited to the unit that does the derivation, plus units that depend on the original unit.
3. Discriminants on records can be used to create unconstrained objects, to make some components inaccessible in some variants, and to define indefinite generic formal parameters and private types. This leads to significant analysis complexity, potential execution-time exceptions, and can cause the use of dynamic memory techniques.
4. The analysis of class-wide calls involves dispatching and hence must consider every candidate object type and analyse the associated subprogram. This can be a difficult undertaking because types declared in library-level package specifications can be extended anywhere in the partition.

5.1.3 Guidance

Integer types should always be declared with explicit ranges, instead of deriving from one of the predefined integer types.

Avoiding Run-Time Dispatching

User invoked dispatching only occurs if 'Class is used. Prevention of dispatching can be enforced by the use of **pragma Restrictions(No_Dispatch)**.

The use of derived types or tagged types with 'Class requires that all operations on the types are checked to ensure that the operations called according to the language rules are the ones the application design requires.

5.2 Declarations

A declaration associates a name with an entity and describes some characteristics of that entity. The Ada language uses these characteristics both *statically* (at compile time) and *dynamically* (at run-time) to ensure that entities are not used in inappropriate ways.

5.2.1 Evaluation

Table 4 : Declarations

Feature	Group Name								
	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Named numbers	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Constants	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc ¹	Inc
Variables	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Unconstrained Object ²	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Initialization ³	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Aliased Object or Component – Simple Case	Alld ⁴	Alld ⁴	Inc	Inc	Inc	Inc	Alld ⁴	Inc	Inc
Aliased Object or Component – Complex Case	Exc ⁵	Exc ⁵	Inc	Inc	Exc ⁵	Inc	Exc ⁵	Inc	Inc
Declarative Part in Block Statement	Alld ⁶	Alld ⁶	Inc	Inc	Inc	Inc	Inc	Inc	Alld ⁶

5.2.2 Notes

1. Boundary value testing is more problematic if a constant is assigned a non-static value on scope entry.
2. The use of discriminants in the creation of unconstrained objects buys high integrity systems more problems than the advantages it may procure. The use of objects of a discriminated subtype may force the compiler to defer to run-time important decisions on the management of the object (e.g., allocation, access). Those cases may prove difficult to analyse. The entries in Table 4 have been assigned assuming that discriminants are *not* used for the creation of unconstrained objects.
3. The initialization of variable objects in high integrity systems should always be static and explicit. The former property ensures that the initializing expression can be verified at compile time. The latter ensures that the initialization is deliberate.
4. Simple definitions of aliased objects are ones that don't change any properties of the object or component from those defined by the initial type. Aliased objects can be accessed by generalized access types, and hence are subject to the same issues that affect pointers. Flow Analysis and Symbolic Analysis may become intractable if views of the object are exported beyond the module being analysed.
5. Complex definitions of aliased objects or components occur when properties of the object may be inconsistent with non-aliased objects of the same type. The language rules in these cases are complex, making analysis unlikely. Examples of this occur when the original type is indefinite, unconstrained, or modified by representation clauses.
6. Ada allows block statements to enclose a declarative part. Whereas the use of this feature is one of the language means for encapsulation and hierarchical program structuring, it also presents some drawbacks to Flow Analysis and Symbolic Analysis as well as to structural coverage. Block statements can be avoided by replacing them by a parameterless subprogram.

5.2.3 Guidance

Declarations should be used to encapsulate the program design in as exact and static a manner as possible. For instance, constraints should be as tight as possible and also static so that analysis can be used to verify them.

Use of Named Numbers

Named numbers are beneficial in that they denote the value of a static expression which is evaluated, with full precision, at compile time. This makes them the natural means of incorporating static numeric expressions within the program and thereby eliminates the possibility that `Constraint_Error` may be raised at run-time during the evaluation of such expressions.

For the same reason, named numbers should be used at all times in the executable code in the place of numeric literals.

Initialization of Variables

All variables should be given a meaningful value before use. Failure to do so may raise a predefined exception or cause a bounded error at run-time.

Initial values may be given by:

1. Associating an explicit initialization expression with the variable at the point of its declaration.
2. Making an assignment to the variable that will be executed prior to references to it.

For state variables in packages, assignments may also be made in the package elaboration part. A consistent approach to the initialization of package state variables should be adopted.

In all cases, Data Flow analysis should be used to confirm that every object has been assigned a value before it is used. The effectiveness of the analysis is undermined if variables are initialized unnecessarily (sometimes called 'junk initialization'). Compilers supporting Annex H provide information on the initialization status of variables, see [ARM:H.3.1(8)].

Use of Aliased Objects

The strong typing of Ada facilitates the achievement of type-safe access to variables through access values. Furthermore, the accessibility rules of the language help in creating objects without the need for general access types. Static tools can also be defined using, for example, Ada Semantic Interface Specification [ASIS], to complement the power of the compiler in the determination of 'unsafe' use of aliased objects. The use of Ada aliased objects is therefore generally safer than the use of pointer or reference objects in other languages.

5.3 Names, including Scope and Visibility

Entities are denoted by names controlled by the rules for scoping and visibility.

Name de-referencing in Ada is usually quite straightforward and determinable at compile time with a few notable exceptions. Renamings can be used to introduce short names for use in a restricted scope. Object renaming declarations can be used to provide a short name for a component of an object. Overloading of names can enhance the readability of a program if applied judiciously.

Nesting of packages inside other packages provides information hiding and containment, but does not increase the scope levels (i.e., cause the information to be nested at deeper levels on the program stack or task stack). There is a tension between nesting and non-nesting of subprograms inside other subprograms or tasks. Many opportunities arise to simplify algorithms by placing some of the logic inside a local function or procedure. On the other hand, high integrity software processes usually enforce the isolated testing of all executable units.

5.3.1 Evaluation

Table 5 : Names, including Scope and Visibility

Feature	Group Name								
	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Names	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Static renaming	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Renaming – complex or dynamic evaluation	Exc ¹	Exc ¹	Exc ¹	Inc	Inc	Inc	Exc ¹	Inc	Inc
Overloading	All ²	All ²	Inc	Inc	Inc	Inc	All ²	Inc	Inc
Nesting package spec	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Nesting package body, subprogram	Inc	Inc	All ³	Inc	Inc	Inc	Inc	Inc	All ³

5.3.2 Notes

1. The use of complex forms of renaming (i.e., those which require run-time evaluation of bounds or object components; or those which extend component lifetime), is excluded because it hinders Symbolic Analysis, Flow Analysis and Range Checking; and complicates Object Code Analysis as it embeds run-time code that has no associated visible source code.
2. If multiple subprograms of the same name are used (whether overloaded, overridden or homographic), it should be shown that the intended subprogram is in fact called.
3. The textual inclusion of packages or generics inside a package body, private part of a package, subprogram body, or task creates challenges for coverage-based testing. External tests that stimulate the enclosing unit can rarely exercise all of the branches, conditions, or statements in the enclosed unit. Similarly Range Checking becomes problematic. Therefore the use of such generics should be restricted to avoid these problems.

5.3.3 Guidance

Names are always resolvable at compile-time, but may involve range checks. Two aspects involving names are considered below.

Renaming

Renaming can improve readability but runs the risk of making aliasing hard to detect. Hence reviews, perhaps supported by tools, are needed to ensure correct usage. Such reviews are greatly simplified if continued use of the original name of a renamed entity is avoided.

The subtype indication in a renaming declaration should statically match the subtype of the renamed object. Similarly, the profile of a subprogram renaming declaration should be subtype-conformant with the profile of the renamed subprogram. A subprogram can be renamed in order to furnish new default parameter values. A renaming-as-body can be used in cases where a subprogram can be implemented directly by calling some other subprogram.

Nesting

The textual inclusion of units inside package specifications and private part of library units provides information containment without sacrificing accessibility of the module for testing. Nested units in the private part can be accessed by child packages of the parent unit.

Packages should be library-level units, visible subunits, or child library packages. Where standards demand unit testing of all subprograms it may be necessary to avoid the declaration of subprograms locally within other subprograms. Accessibility for unit testing can also be achieved, without the need to compromise program structure, by making local subprograms subunits and placing their source in a separate file.

5.4 Expressions

An expression is a formula defining how a value is to be calculated. The value of an expression is determined by evaluation of the formula using the current values of the operands that appear in the formula. Operands are any names, literals, function calls, allocators, type conversion, qualified expressions, or aggregates.

Every expression has a fixed type, and this is the type of the value resulting from its evaluation. The type is determined either from the types of the operands within the expression or from the context of the expression. In an imperative language like Ada expressions are readily available for review and analysis. This facility is supported by the strong typing of expression results.

Static expressions are constant values that are completely determined at compile-time. They can be used to determine the value of numeric constants, enumeration literals, string constants and bounds for ranges and arrays. Static expressions have the following properties: they enhance review and analysis; they move the checks on array bounds and ranges from run-time to compile-time; and they eliminate run-time complexity in the calculation of such values.

Numeric type conversions are straightforward since the underlying numerical value is represented in the new type according to precise, deterministic rules [ARM: 4.6 (29-33)] unless `Constraint_Error` is raised.

5.4.1 Evaluation

Table 6 : Expressions

Feature	Group Name								
	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Operators with Composite Operands ¹	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Logical Operators	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc ²
Short-circuit control forms	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc ²
Relational operators	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Membership Tests	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Slices	Inc	All ³	Inc	Inc	All ³	Inc	All ³	Inc	Inc
Qualified Expressions	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Aggregates ⁴	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Static Expressions.	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Type Conversion Numeric	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Type Conversion Other	Inc	All ⁵	Inc	Inc	Inc ⁵	Inc	All ⁵	Inc	Inc
Indexing	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc

5.4.2 Notes

1. If one or more of the operands is a composite object then the compiler may insert loops into the object code that are not in the source code. Examples where this may occur are relational operations on arrays and logic operations on boolean arrays. If the size of these objects is not static then the timing analysis must be based on the worst case (which is probably, but not necessarily, when the objects have the largest possible size).
2. When testing decisions with multiple conditions, the value of every condition will need to be considered. If the short circuit forms are used then the values of some conditions will be insignificant in some test cases and need not be considered in those test cases. Achieving full coverage of compound decisions is more difficult when the conditions within the decision are not independent. If short circuit forms are used then the effect of dependencies between conditions is reduced.
3. The bounds of the range that defines the slice must be of the type of the index of the array. Two or more slices of the same array may overlap (so that components of the array appear in more than one slice). This may make understanding of the code more difficult, although the language definition ensures that the effects of overlapping slices are well defined. Slices will introduce loops into the object code that are not visible in the source and this may make timing analysis more difficult. Slices of packed arrays will introduce further timing complexity. Slices defined by ranges that are static subtypes will be simpler to analyse.
4. If an aggregate is to be assigned to an object and the aggregate references that object then the object code will be more complex than it might otherwise be:
 - The enforced use of a temporary object makes the code more obscure and compiler dependent.
 - The compiler may require heap space for the storage of the temporary object but heap management algorithms may not be present.
 - Data values are much more difficult to trace, as they are transferred through the hidden temporary object.

Temporary objects may also be introduced by any use of non-static aggregates. The above notes, therefore, apply.

5. There are a wide variety of type conversions for which the underlying system must be shown to implement the proper behaviour of the program. In particular conversions of composite and non-static objects generate additional code that must be traced back to source code. Conversion of composite objects usually requires a temporary object, heap storage, and unpredictable timing. View conversion (e.g., between tagged types, or applied to actual **in out** and **out** parameters), requires dynamic checks as does conversion between generalized access types. Conversion between access types with different storage pools is potentially erroneous.

5.4.3 Guidance

For multi-termed expressions it is advisable to constrain the compiler's actions by fully parenthesizing the expression. The programmer should have a clear view of the required order of evaluation of sub-expressions and should enforce this view. This reduction in non-determinacy will significantly simplify verification and improve readability. For example, proof of exception free evaluation is obviously easier if there is only one order of evaluation.

Aggregates should be used as actual parameters and function return values in preference to temporary variables for the individual components, and as the right-hand side of assignments to composite objects in preference to separate assignments to the individual components. This makes it easier to verify that all individual components have been set, makes explicit any dependencies on the composite object assigned to, and makes maintenance more robust should the type definition change. Aggregates used to initialize objects should be static wherever possible, to avoid the need to build a temporary object (this is a specific implementation permission). Named aggregates should be preferred over positional except for large, static arrays, and explicitly named choices should be preferred over **others** (viz '1..n => ...'). Named aggregates are more likely to require compiler-generated code to fill holes and **others** to require dynamic storage.

When two types are not readily convertible, or fall into one of the cases noted below, the application should be recoded to avoid such conversions using, for instance, look-up tables, function calls or case statements:

- composite objects;
- tagged and indefinite types;
- **in out** or **out** parameters in a procedure call;
- access and generalized access types.

Complete test coverage may be difficult to achieve for expressions that create implicit loops in the object code (e.g., aggregates, expressions with composite operands). For example it may not be possible to create operands for which a loop is executed zero times, and this omission will need justification.

5.5 Statements

Statements are the basic commands, both simple and compound, that make up the Ada programming language. This section is concerned only with those statements that are not specific to a particular area of the language. For example, the accept statement is not considered as it relates only to tasking (see 5.13).

5.5.1 Evaluation

Table 7 : Statements

Feature	Group Name								
	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Null	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Assignment	Inc	Inc	Inc	Inc	Inc ¹	Inc	Inc ¹	Inc	Inc
Block	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Goto ²	Exc²	Exc²	Inc	Inc	Exc²	Inc	Inc	Inc	Exc²
For loop	Inc	Inc	Inc	Inc	Inc ³	Inc	Inc	Inc	Inc
While loop	Inc	Inc	Inc	Inc	All⁵	Inc	All⁴	Inc	All⁴
Simple loop with exit	Inc	Inc	Inc	Inc	All⁵	Inc	Inc	Inc	Inc
case	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
If	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc

5.5.2 Notes

1. There may be an impact on Timing Analysis when assignment statements are between non-scalar values, e.g., when assigning values to arrays. This impact should be predictable, but must be considered when the analysis takes place.
2. The **goto** statement is exceptional because its use is contrary to all principles of structured programming. There are no circumstances in which **goto** can be used where the use of some other construct is not preferable on grounds of good practice, readability, and aesthetics. Given this, the use of **goto** within high integrity systems is almost not an issue, and the reasons for not using it is based on the applicability of the set of analysis techniques (although compelling) are almost irrelevant.
3. The Timing Analysis of a **for** loop gives precise results if the range through which the iteration takes place is statically determinable. The analysis is more difficult, and the results less useful, if the range is dependent on one or more variables. Timing analysis is complicated significantly by the use of an **exit** statement within a **for** loop because the exact behaviour of the loop is harder to predict.
4. When generating the object code for a **while** loop, a compiler may generate more than one test for exit from the loop, i.e., a single statement in source code may map to multiple pieces of object code. This makes Object Code Analysis more difficult, and also makes it harder to ensure that all branches and statements have been fully exercised when testing Structural Coverage.
5. Timing Analysis is difficult for **while** loops, and any other loops where a condition is required to be satisfied before exit. It is not usually possible to determine accurately how many times the loop will be traversed before the exit condition is satisfied. The use of annotations to capture maximum loop counts is recommended.

5.5.3 Guidance

The use of statements of the type described in this section is fundamental to any structured programming language; these are the basic tools that build the underlying program structure. They allow the use of loops and conditional branching, essential features of the vast majority of meaningful programs. Because of their fundamental nature, the behaviour of most of these constructs is well-defined. That does not, however, necessarily mean that it is always absolutely predictable. For this reason, the ease of applying a specific analysis technique to a type of statement may depend on how that statement is used.

Apart from **goto**, and to a small extent loops, there is no need to restrict the use of these basic statements (unless there are difficult-to-resolve timing issues), and to do so would place a great burden on the programmer. There is a need for caution, however, and good programming style should be used at all times. **loops**, **case**, and **if**, in particular determine the main structure of a program (or subprogram), and using these in an effective, well-structured, manner can make the whole analysis and testing process much simpler to perform.

Functional correctness of nested loops is facilitated if the following are true: an exit statement must only transfer control to the level above, and the exit condition must be tested on each iteration.

5.6 Subprograms

Procedures are a basic unit of abstraction for statements and are an essential element of any imperative programming language. Ada subprogram specifications allow the mode and subtype of each parameter to be specified, allowing both compile-time type checks and run-time constraint checks on parameters in a call. Ada's compilation environment (library) requirements, and strong type checking eliminate most forms of incorrect invocation of a subprogram.

5.6.1 Evaluation

Table 8 : Subprograms

Feature	Group Name								
	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Procedures ¹⁰	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Functions ¹⁰	Inc	Inc ¹	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Default Expression	Inc	All²	Inc	Inc	Inc	Inc	All²	Inc	Inc
Indefinite Formal Parameters	Inc	Inc	Inc	All³	All³	All³	Inc	Inc	Inc
Complex return types	Inc	Inc	Inc	Exc⁴	Exc⁴	Exc⁴	Inc	Inc	Inc
Inline expansion	Inc	Inc	Inc	Inc	Inc	Inc	All⁵	Inc	All⁵
Return in procedures	All⁶	All⁶	Inc	Inc	Inc	Inc	All⁶	Inc	Inc
Parameter aliasing	Exc⁷	Exc⁷	Inc	Inc	Inc	Inc	Inc	Inc	All⁷
Access Parameters	Exc⁸	Exc⁸	Inc	Exc⁹	Inc	Exc⁹	Inc	Inc	Inc

5.6.2 Notes

1. If functions have side-effects, order of evaluation can become a significant issue, making Symbolic Analysis difficult.
2. The use of default expressions for some parameters is a convenience, but comes at a cost. Implementations may differ in how they support default parameters, making time and space analysis more difficult. The traceability of source code to object code is more difficult and variation between implementations is more marked if the default expression is not static. If default expressions are allowed, then the above issues could be addressed by a tool or code reviews.
3. Procedures and functions can have parameters of indefinite types. This may lead to a requirement for dynamic storage, depending upon the implementation and the operations applied to the parameters.
4. The following types should not be used in function returns because they require dynamic storage techniques:
 - Indefinite types – types with unconstrained or unknown discriminants or class-wide types.
 - Unconstrained types – such as string, and
 - Tagged types – since the returned type and actual type can differ.
5. Inline expansion of a subprogram call can be used to eliminate parameter passing overhead and may reduce the execution time of a program. Code size, however, can be increased, and the tracing of object code to source code can be more difficult.
6. Return statements can make the natural flow of control more apparent but returns from deeply nested structures can be obscure and cause difficulties for Flow Analysis, Object Code Analysis etc. Only allowing returns at the outermost scope is an effective restriction.
7. When parameters are aliased (to non-locals or other parameters) then program proof based upon substitution will be incorrect. Similarly, informal reasoning can easily be in error. Hence if Formal Code Verification or Symbolic Execution is

being used as an analysis technique, the absence of aliasing is required. The absence of aliasing should be determined by the use of tools and code reviews.

8. Aliasing occurs when two distinct Ada identifiers actually refer to the same object. This can be between parameters or between parameters and non-locals. Aliasing will cause Flow Analysis and Symbolic Analysis to malfunction. Similar problems arise with access types when two such values point to the same entity.
9. The lack of accessibility checks on access parameters (i.e., parameters that explicitly use the keyword **access**) makes it very difficult to undertake the analysis of memory usage.
10. For the implications of recursion, see Section 5.6.3.

5.6.3 Guidance

The following restrictions on parameters are advised to eliminate the problems caused by indefinite formal parameters: no concatenation of one-dimensional arrays and no unconstrained records. Subtype conversions in subprogram calls must also be used with caution.

Static analysis must be used to ensure parameters with **out** mode are assigned a value in all execution paths, and return statements are always encountered in functions. Side-effects from parameter evaluation should always be avoided. Similarly, it is usually advisable to prevent functions from having side-effects (although there are some circumstances in which side-effects are unavoidable, but these must be verified separately). The use of non-local variables should always be documented.

Recursion, and mutual recursion, is usually prohibited for flow analysis, stack analysis and timing analysis. Direct recursion is easy to detect. Mutual recursion is harder to detect in general; however, if **pragma Elaborate_Body** is applied to all library units, and tagged types and generic units are not used, then mutual recursion can occur only between subprograms in a single library unit. This is more readily analysed. **Pragma Restrictions (No Recursion)** does not prevent recursion from occurring, and if recursion does occur the execution is erroneous. Implementations might check for violations of the restriction, and may generate somewhat simpler code.

Overloading can make the program reader's job more difficult. The advice of the Ada Quality and Style guide [AQS] seems reasonable: use overloading judiciously, for widely used subprograms that perform similar actions on arguments of different types, and preserve the conventional meaning of overloaded operators.

5.7 Packages (child and library)

Packages are Ada's basic unit of modularity. Therefore packages are fundamental to the creation of any Ada program. Packages allow the partitioning of a program into parts that interact using well-defined interfaces. This can facilitate the analysis of a program by limiting the interactions between its parts.

High integrity programs can sometimes be structured so that the code that deals directly with some critical aspect of the system can be encapsulated in a package body. This is ideal, as Ada's language rules then guarantee that this code is called only using the interface defined in the package specification. Furthermore, any local data used in this critical code is protected from tampering.

Packages define three different levels of isolation. Entities declared in the public part of a package specification are visible wherever the package itself is visible. Entities declared in the private part are visible to the package body, and also within any child packages (if the package is a library unit). Finally, entities declared in the body are visible only within the body. These levels of isolation permit designers and implementers to implement stand-alone service packages, subsystems of cooperating packages, or packages that export all significant items. Subsystems can therefore be built with exactly the proper amount of visibility and security for the systems being designed. Issues to do with nested packages are considered in Section 5.3.

Child library packages provide a mechanism for building subsystems. Because it is closely related to subunits there is little compiler impact.

5.7.1 Evaluation

Table 9 : Packages (child, and libraries)

Feature	Group Name								
	FA	SA	RC	ST	TM	OMU	OCA	RT	ST
Specifications	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Bodies	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Initialization ¹	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
With Clause	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Private ^{2,3}	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc ²
Use Clause ⁴	Inc	Inc	Inc	Inc	Inc	Inc	All ⁴	Inc	Inc
Use Type ⁴	Inc	Inc	Inc	Inc	Inc	Inc	All ⁴	Inc	Inc
Child Private	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Child Public	Inc	Inc ⁵	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Separate	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Elaborate Body	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc

5.7.2 Notes

- Ratings apply when elaboration order dependencies have been avoided by restricting access to indirectly accessed objects and subprograms during package elaboration. The following restrictions are required:
 - Package initialization code should never update objects in packages.
 - Initialization code in package bodies should not directly or indirectly reference subprograms or objects other than static constraints declared outside the package.
 - Expressions initializing objects declared in a package specification may reference objects and subprograms from remote packages provided all such remote packages, whether the reference is direct or indirect, are pure, preelaborated or contain a **pragma Elaborate Body**.
- Operations and objects defined in private parts cannot be explicitly verified from units having visibility only to the specification.
- Derivations from private types and public types with private primitive operations create problems for flow analysis, symbolic analysis, and human understanding. Derivations or extensions of such types allows one to create subprograms with exactly the same signature as the private operations. Which call was executed would depend upon scope visibility. The ratings here assume that these capabilities have been excluded.
- The use clause brings excessive overloading and possibly hiding of names. Human understanding and code review are made difficult. Use type clauses only bring into view the operations and literals of the type named in the clause. This provides a constrained view of the operations on a type.
- The ratings assume that child library units do not **with** units that **with** any parent units. Child library units can create dependencies on units that are dependent on the parent. Multiple views of a type and its primitive operations may be created. Ada language rules are designed to make only one path legal, but there may be confusion on the part of human designers and programmers.

5.7.3 Guidance

Any complex initializations can be done by procedures or assignments instead of during elaboration. The use of such initializing procedures for library-level packages may be preferred, as long as they are called only after the main subprogram has commenced. Then the elaboration of the package body variables and subprograms will have all occurred, but it must still be shown that the call to the initialization subprogram occurs before any use of the package — a non-trivial exercise, especially if any concurrency exists in the program.

The type defined private in Ada has at least two views — the public partial view and the complete private view. Language rules permit some differences between these definitions, such as adding or removing discriminants, adding limited, tagged, or aliased in the full view, and placing primitive subprograms in the private part. It is recommended that derivations from a type that has a

partial view only occur where the full view (and hence all primitive subprograms) is visible, i.e., in the private part, in private child packages and in package bodies.

Use Clauses

Most large software development projects and high integrity software development projects place significant restrictions on use clauses. Those restrictions, however, should not discourage the educated use of the **use type** and **renames** clause. These clauses, in fact, provide alternatives that make primitive operations available without the wholesale import of another package's name space.

Child packages

Child packages permit the aggregation of packages into hierarchies of packages, and allow a subsystem to be extended without forcing the basic definition of the system (in the parent unit) to be modified or recompiled. Tools that analyse coverage must take into consideration all of the ways that child packages can be included in a program.

Child packages should not **with** units that may **with** the parent, and should not derive or extend types that have been declared in the private type of a parent unit.

5.8 Arithmetic Types

The ability in Ada, in contrast to most other languages, to define types and subtypes with static bounds substantially aids the review, analysis and verification process. For instance, simple tools can often affirm that uses of integer variables cannot raise an exception, say in an array indexing operation.

5.8.1 Evaluation

Table 10: Arithmetic Types

	Group Name								
Feature	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Integer types	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Modular types	Inc	All ^{d1}	All ^{d1}	Inc	Inc	Inc	Inc	Inc	Inc
Float types	Inc	All ^{d2}	All ^{d3}	Inc	All ^{d4}	Inc	Inc ⁵	All ^{d2}	Inc
Fixed types	Inc	All ^{d2,10}	All ^{d3,7}	Inc	All ^{d9}	Inc	All ^{d6,9}	All ^{d2,7}	Inc
Dec/fixed types	Inc	All ^{d8}	Inc	Inc	All ^{d8}	Inc	All ^{d8}	All ^{d8}	Inc

5.8.2 Notes

1. The predefined operations on modular integer types are not those of classical mathematics (for instance, operations are non-associative for modular types when the modulus is not a power of 2). In consequence, care is needed to ensure that the operations perform the intended function. The same reservations apply to the use of the **Modulus** attribute. Since the predefined operations in modular types do not raise **Constraint_Error**, one must ensure that the semantic effect is as specified in the application.
2. Due to the potential machine-dependence of the rounding of real operations, the computation of boundary value test cases can be awkward (see Equivalence Class Testing). Similarly, the rounding presents problems with Symbolic Analysis.
3. If the program is required to be exception-free, then an analysis is required of the range of values that variables can have to demonstrate that each numeric operation will indeed be exception-free. This work is typically labour-intensive and is often more difficult with fixed point since the range of values is more restricted.
4. The timing of floating point operations are typically data-dependent. Hence computing the worst-case execution time can be awkward if an analysis of the data is required.
5. Floating point requires complex (hardware/software) support. Such complexity may require specific verification measures.
6. Many fixed point operations require run-time support from the compiler (for instance re-scaling). This support is context dependent and can be quite complex. The effort in showing the compiler does handle every operation correctly is significant and has led some developments to use floating point instead. Errors have been observed in validated compilers in this area, but special tests are available to check implementations.

7. The rounding of some fixed point operations is not defined in the language and the implementation need not provide any information on the actual rounding performed. This implies that the rounding could be context-dependent or even change with a new release of the compiler.
8. Support for decimal fixed point is typically only provided by compilers supporting the Information Systems Annex [ARM: Annex F]. The predefined operations on decimal fixed point types that do not give an exact result are defined to truncate. In consequence, it may be easier to verify programs which use decimal fixed point types than those which use ordinary fixed point types. In other respects, the verification issues are the same as for ordinary fixed point types.
9. Fixed point operations require compiler support for which timing and object code analysis is more complex.
10. Type conversion between fixed point types whose delta is not a power of 2 may introduce additional rounding errors. Hence such types should be avoided. Placing a representation attribute clause for 'Small to match the delta is recommended for all fixed point types.

5.8.3 Guidance

As noted in 5.2.3, named numbers should be used when the application permits compile-time evaluation.

Continuously varying quantities should be modelled by means of Ada real types. It is not usually appropriate to use integer types since there is no simple way of multiplying integer values and re-scaling the result. Hence an important design decision is to use either ordinary fixed point types or floating point types (or less likely, both).

The inherent problem with real types is that rounding is performed which introduces a degree of implementation defined behaviour into the program which does not occur with integer types.

Analysis of real expressions requires analysis of the ranges within which the result of an expression is guaranteed to lie. A floating point expression has a value in a range which has a size relative to the value of the expression and a fixed point expression has a value in a range which has an absolute size.

When using real arithmetic:

- Be aware of whether the implementation supports the Numerics Annex [ARM:Annex G] as the accuracy of real arithmetic is guaranteed only for implementations that support this annex.
- For maximum portability of code, only use implementations that support the Numerics Annex, and ensure that the strict mode is used.
- For implementations that support the Numerics Annex, the accuracy of the predefined operators is defined in terms of model numbers of the relevant type. Although the set of model numbers is implementation defined, it is straightforward for a user to define a set of model numbers that will be provided by all implementations of interest (provided only they have a common radix for number representation). By analysing the accuracy of real arithmetic in terms of this set of model numbers all implementations are guaranteed to provide the analysed accuracy or better.
- If future implementations are unknown then the characteristics assumed by the accuracy analysis should be recorded and it is then easy to determine whether a new implementation conforms (and so whether the previous analysis still applies or needs to be repeated).
- Currently, most of the major processor chips directly support floating point. The older chips do not, and neither do some specialized chips. Hence in some cases, the use of floating point is not viable due to the lack of hardware support and because software support is too slow. If the systems design would allow for either floating point or fixed point, then the choice is critical since it influences the coding, the testing and the qualification of the compiler. If it is decided not to use floating point, then the **pragma Restrictions (No_Floating_Point)** can be used to enforce no explicit use. However, the compiler may still make implicit use of floating point (say, for complex numeric conversions), which may require compiler options to remove.

Not specifically listed in the above table is the use of attributes. There are numerous numeric attributes which can conveniently be divided into two classes:

- Those whose use provides no special problems:

Adjacent,	Aft,	Ceiling,
Compose,	Copy_Sign,	Delta,
Denorm,	Digits,	Exponent,
Floor,	Fore,	Fraction,
Leading_Part,	Round,	Rounding,
Scale,	Small,	Truncation,
Unbiased_Rounding.		

- Those which are low-level and must be used with care to ensure portability:

Machine,	Machine_Emax,	Machine_Emin,
Machine_Mantissa,	Machine_Overflows,	Machine_Radix,
Machine_Rounds,	Model,	Model_Emin,
Model_Epsilon,	Model_Mantissa,	Model_Small,
Remainder,	Safe_First,	Safe_Last,
Scaling,	Signed_Zeros.	

Specific problems noted in the table imply that modular types are rarely an appropriate choice instead of integer types.

5.9 Low Level and Interfacing

Low level constructs and interfacing exist to allow an Ada program to interact with:

- elements of the machine, e.g., memory addresses;
- other hardware elements of the system, e.g., tape and screen devices;
- other software elements of the system, e.g., databases, GUI; and
- other languages, e.g., C, Fortran, Machine Code.

Ada is very useful in providing a mechanism for interchanging information between different languages. It is also useful in allowing these mechanisms to be performed in a very constrained way. Many of the interactions that cause difficulties occur between standard Ada facilities and the features discussed here. Nevertheless the feature, or combination of features, can often be used successfully if the encapsulation guidelines are followed, and if multiple combinations of low-level features are not applied simultaneously.

All of the low level features described in the following table must be encapsulated in program units that clearly isolate their behaviour from the rest of the program. The ability of Ada to support this encapsulation and isolation is a key advantage obtained from the use of Ada.

The use of low level constructs and interfacing can be seen to be a mechanism for dealing with elements outside of the Ada system under consideration.

5.9.1 Evaluation

Table 11 : Low level and interfacing

Feature	Group Name								
	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Representation Clause	Inc	All ^{d1}	Inc	All ^{d1,2}	All ^{d1,2}	Inc	All ^{d2}	All ^{d2}	All ^{d2}
Representation Attributes ³	All ^d	All ^d	All ^d	All ^d	All ^d	All ^d	All ^d	All ^d	All ^d
Package System ⁴	All ^d	All ^d	All ^d	All ^d	All ^d	All ^d	All ^d	All ^d	All ^d
Machine Code Insertion ⁵	All ^d	All ^d	All ^d	All ^d	All ^d	All ^d	Inc	All ^d	All ^d
Unchecked Conversion	All ^{d7}	All ^{d6}	All ^{d6}	Inc	Inc	Inc	Inc	All ^{d6}	Inc
Unchecked Access ⁸	Exc	Exc	Exc	Exc	All ^d	Exc	All ^d	Exc	Exc
Streams ⁹	Exc	Exc	Exc	Exc	Exc	Exc	Exc	Exc	Exc
Interfacing Pragmas ¹⁰	All ^d	All ^d	All ^d	Inc	All ^d	All ^d	All ^{d11}	All ^d	All ^d
Address Clause ¹²	Inc	Inc	Inc	Inc	Inc	All ^{d13}	All ^{d14}	Inc	Inc

5.9.2 Notes

- Representation clauses can have interactions with other features that make their semantics, storage use, or timing properties difficult to predict.
- Representation clauses may lead to object code for initialization or unpacking. Also, such clauses can require that the compiler emits very different machine code than is usually the case, leading to difficulties in tracing from source to object code and a risk of incorrect object code.
- These ratings are for representation attributes in expressions (and not their use in representation clauses); they are then, typically, implementation-dependent compile-time constants.
- Package System, including its children, is compiler-dependent. The facilities it provides should be evaluated by comparison with equivalent features described elsewhere in Section 5. For example, constants exported by package system are included (see 5.2.1).
- Machine code is not constrained in its effect and must, therefore, be used with extreme care and encapsulated in the smallest unit possible.
- No constraint check is performed on the assignment of a result from an unchecked conversion. However, it is possible to use 'Valid for scalar results to ensure the value is a valid member of the appropriate subtype. In the case of conversion to a non-scalar type, extreme care is needed to ensure that a valid result is produced, perhaps by examination of the machine code produced (since 'Valid cannot be applied)
- Unchecked conversion of structured types will result in loss of some flow analysis information. It will not be possible to deduce the flow of information from components in the first view to particular components of the converted view.
- Unchecked access value creation, and subsequent use of the access value, can lead to dangling references or corruption of data.
- Streams require class wide types and access parameters, and are therefore difficult to analyse.
- The effect of an imported subprogram needs to be documented in such a way that its callers can be analysed.
- There may be object code generated to account for differences in calling convention.
- Address clauses should not be used to introduce aliasing by overlapping or superimposing variables in memory.

13. Analysis of memory use must consider the impact of the location of objects at specific addresses.
14. Object Code Analysis must ensure that the specified address of a variable is directly generated for all accesses of that variable and that the variable has not been optimized to a register, for example. Use of **pragma Volatile** may be appropriate.

5.9.3 Guidance

If the system has a defined integrity level and the Ada code in the system is required to interact with some other element, then assurances should be obtained that the other element has the same integrity as the Ada code itself. This Technical Report considers COTS software to be another element that ranges from complete software packages to code segments that have not explicitly been created for the current system.

If 'allowed' low-level features are used, it is recommended that their use be encapsulated within a small package body or subprogram. This reduces the possibility of an interaction with some other feature used in a different part of a program, and can facilitate analysis. Careful human analysis of the generated object code of these features will usually be required to replace more traditional analysis forms.

Representation clauses on enumeration types should be avoided, except when confined to a unit body to interface to foreign systems. A type with a representation clause should not be used as a source of a type derivation, nor should it be used as a loop parameter, case expression, array index, in a 'Succ or 'Pred clause, or in a type conversion.

If a record definition is subject to representation clauses, such as packing, alignment, layout or ordering clauses, then the declaration should be restricted to a compilation unit body. Care is required to ensure that there is not a conflict between representations applied to the record, components of the record, and objects of the record (e.g., **pragma Pack** and 'Access may be mutually incompatible). Assignment to or from such a record should be component-by-component to avoid timing issues and unanalysed object code.

Representation Clauses

Such clauses in general change the machine-code emitted by the compiler. They should be used solely to enable such changes and not to 'confirm' an existing representation. An enumeration representation clause with non-contiguous values will cause the code for many operations on the enumerated type to be obscure and hard to relate to the source text; hence such operations should be avoided. Similarly, record representation clauses can make operations such as assignment, equality and conversion non-trivial and hence every operation should be used sparingly and in a localized fashion.

5.10 Generics

Generics provide a powerful mechanism for constructing large-scale programs through the parameterization of program units (packages and subprograms) with types, objects, packages and subprograms. Abstract algorithms and data types can thereby be constructed in terms of parametric types and operations, specifying only as much as necessary. A generic instantiated with different sets of parameters yields different program units that share the same algorithmic structure.

Generic units can be used to:

- define abstract data types;
- parameterize a procedure by another procedure (e.g., a loop iterator);
- replicate a unit.

Generics can appear at library level or nested within other units. Generic units can be children or parameters to other generics, enabling the construction of whole parametric software subsystems.

Use of generics enhances program reliability in several ways. It facilitates reuse, eases maintenance, reduces source code size and helps avoid human replication error. In principle, it also assists with static analysis and testing; if a generic algorithm is verified once, all instances of the generic can be considered automatically verified. In practice, the complexity of the instantiation process mitigates against this ideal.

Analysability of a generic feature depends not only on the semantics and behaviour of the feature, but also on how it is compiled: whether by macro-expansion or code-sharing. In general shared code is highly parametric, because small changes to

actual parameters make dramatic differences to efficiency (e.g., composite vs. elementary types). Code coverage is difficult to achieve, because there are rarely sufficient instances to test all options. The complexity of mapping from the generic unit directly to object code impedes verification as well as robustness and reliability. In practice a majority of compilers adopt macro-expansion, or a hybrid where only simple cases share code. It is recommended that generics compiled by code-sharing are excluded for high integrity systems.

Analysis of a generic feature might be undertaken on the generic unit, or on its instances. The former offers the possibility of 'once-and-for-all' verification: if the generic is shown to have some property, then that property is inherited by all of its instances. In principle it is possible to perform some analyses, notably Symbolic and Flow Analysis, on a generic unit alone, although tool support is presently weak. This is a corollary of Ada's contract model, which states that the formal specification contains sufficient information to determine the legality of the generic body, whatever actual parameters are supplied. Verification conditions at the point of instantiation must also be satisfied. This might involve work, or the imposition of restrictions, e.g., to enforce range constraints on formal parameters.

Alternatively, the analyses can be performed on the individual instances. Other techniques (Object Code Analysis, testing) must be applied this way. Success criteria might be specified individually for each instance, on the generic itself and derived ('instantiated') for the instances, or both (provided consistency is maintained). Annotations and test points needed to express the criteria can be attached to the source code, a compiler intermediate form such as ASIS (where one exists), or onto the object code via a debugger. This last is highly effective for dynamic testing and has been used with success in high integrity systems. It is presently unclear what static analysis annotations might apply to generics or instances, or how to attach them; this presents a challenge to tool builders. However there is no reason to believe that static analysis techniques cannot evolve to incorporate generics in some form.

Analysis of instances offers more immediate potential than analysis of whole generic units. Although it has to be repeated for each individual instance, the tool requirements are much lighter and the complex step of verifying the instantiation is avoided.

5.10.1 Evaluation

The table below assumes a macro-expansion implementation, with analysis tools acting on an intermediate expanded form.

Table 12: Generics

Feature	Group Name								
	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Generic units (packages & subprograms)	All ^{d1}	All ^{d1}	Inc	Inc	Inc	Inc	All ^{d2}	Inc	Inc
Formal subprograms	All ^{d3}	All ^{d3}	All ^{d4}	Inc	Inc	Inc	All ^{d2}	Inc	Inc
in objects	Inc	Inc	Inc	All ^{d5}	Inc	Inc	All ^{d2}	Inc	Inc
in out objects	All ^{d3}	All ^{d3}	Exc ^{3,6}	Inc	Inc	Inc	All ^{d3,6}	Inc	Inc
type parameters ⁷	Inc	Inc	Inc	Inc	Inc	Inc	All ^{d7}	Inc	Inc
Default parameters	Inc	Inc	Inc	Inc	Inc	Inc	All ^{d8}	Inc	Inc
Formal packages	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}
Generic children	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}	All ^{d9}
non-library and nested generic units	Exc ¹⁰	Inc	Inc	Inc	Inc	Inc	All ^{d2}	Inc	Exc ¹¹
non-library and nested instances	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Exc ¹¹	Exc ¹¹

5.10.2 Notes

1. Tool support is not available at the present time. There are potential name capture problems when an actual parameter shares a name with a local of the generic unit.
2. The mapping between source and object code is complex for generics. In particular, implicit operations to perform assignment, equality, constraint checking, etc., are not obvious in the source code and obstruct understanding. This comment applies to object-code analysis of all aspects of generics.

3. Subtype constraints on generic formal **in out** objects and on parameters to generic formal subprograms are disregarded. This means that constraint exceptions cannot be predicted looking at the generic alone. When an **in out** formal object is aliased at instantiation by providing an actual that also corresponds to another formal object, then program proof based upon substitution may be incorrect. Similarly, informal reasoning may be in error.
4. Subtype constraints on parameters to formal subprograms are disregarded: those pertaining to the actual are applied at run-time. To simplify the semantics, formal subprograms should statically match the actual.
5. For non library-level generic instantiations, each instance of an **in** object is non-static and requires an object of potentially unknown size to be created in the stack frame of the instantiating unit. Unlike library level instantiations which only allocate the storage once at program elaboration time, stack-based elaboration may change each time the enclosing unit is called..
6. An instance of a generic **in out** object is an alias of the actual object, with constraints those of the actual, not the formal. The alias complicates the code mapping.
7. 'type parameters' covers all kinds of type: private, limited private, scalar, derived, tagged, array and access. None are excluded, but in general the more closely-defined the type (i.e., smaller value-space), the harder it is to relate source and object code, because more primitive operations occur. If the formal is a derived type, re-export of new derivations from the generic causes confusion between primitive operations. For further analysis, refer to [CAT3, AQS].
8. When binding defaulted formal parameters, different visibility rules apply if the box <> version is used. The mechanism is prone to human error.
9. Formal packages and generic children provide a very powerful mechanism for encapsulating software components and subsystems for reuse, thereby increasing the reliability and maintainability of software. They are high level features: all semantic issues are resolved at instantiation and there is no impact on code generation. Therefore it seems unlikely that they will impede analyses at the instance level. However there is presently little field experience in their use and compilers have not been exercised in this area. Therefore they cannot be recommended for high integrity Ada at the present time.
10. Non-library and nested instances of generics give rise to problems of testing which do not apply to code which can be tested separately from its use. In consequence, achieving coverage of such generics may be difficult or impossible..
11. Non-library generics or instances cannot be breakpointed on a per-instance basis, see also 5.3.2, Note 3.

5.10.3 Guidance

Generics offer great potential for improving program reliability and maintainability. Against this, engineering experience and support in the high integrity field is limited. Therefore, generics should only be allowed where compilers are considered reliable, users have experience, and there exist support tools appropriate to the application.

It is recommended that only macro-expanding compilers be used, and that analysis is performed on each instance in preference to analysis of the generic unit.

There are several issues with generics that obstruct analysis, comprehension, verification or maintenance. These should be excluded from high integrity systems:

- Exclusions to enable analysis:
 - generics compiled by code-sharing,
 - generics and instances not declared at library-level,
 - nested generics,
 - formal **in out** objects.
- Exclusions for reasons of human comprehension and ease of maintenance:
 - formal tagged and derived types, or formal types with unknown discriminants,
 - default parameters to generic units,
 - subprogram parameters with constraints that do not statically match the actual.

In addition features which involve other excluded areas of the language, such as dynamic types, should not be available to the user.

5.11 Access Types and Types with Dynamic Attributes

Access types provide pointers to objects whose memory is allocated from address space regions which are predictable (global and stack) or from regions which are hard to predict (heap). As these pointers are typed, they provide an additional level of security over direct object addresses. When access types are used in ways that create aliases and export aliases, then data analysis and formal analysis of the system that includes these objects becomes quite difficult.

Most unconstrained Ada types can only have constrained objects as their members, with well-known techniques to implement the bounds. Parameters of unconstrained types always have actual objects associated with them. With proper use of preconditions or similar techniques, absolute bounds can be established for such types.

Heap management raises problems with consumption of time and storage. A storage pool may be logically equivalent to a stack of the objects through usage conventions and implementation. The allocation and deallocation in such pools can be made predictable.

Types whose size depends on run-time values make the bounds of storage use difficult to predict. Indefinite subtypes do not have enough information from the type itself to create objects. Unconstrained record objects cause problems with analysis and resource use as a change in their discriminants will change the usage of the storage allocated to the objects during program execution.

Unbounded storage is incompatible with high integrity systems since the occurrence of `Storage_Error` is unacceptable. This implies that types with dynamic attributes are either excluded or should be used with extreme care.

The requirement for staticness excludes variant records (which require discriminants, see 5.1) and run-time dispatching.

5.11.1 Evaluation

Table 13 : Types with Dynamic Attributes

Feature	Group Name								
	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Unconstrained array types – including strings ¹	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Full access types	Exc ²	Exc ²	Inc	Exc ²	Exc ²	Exc ²	Inc	Inc	Inc
Restricted storage pools ³	All ⁴	Exc ⁴	Inc	Inc	Inc	Inc	Inc	Inc	Inc
General access types	All ⁴	All ⁴	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Access to subprogram	Exc ⁵	Exc ⁵	Inc	Inc	All ⁵	Inc	All ⁵	Inc	Inc
Controlled types including unrestricted storage pools	Exc ⁶	Exc ⁶	Inc	Inc	Inc	Inc	All ⁶	Inc	Exc ⁶
Indefinite objects ⁷	All ⁴	All ⁴	All ⁴	Exc	Exc	Exc	Exc	Inc	Exc
Non-static array objects ⁸	Inc	All ⁴	All ⁴	All ⁴	All ⁴	All ⁴	Inc	All ⁴	Inc

5.11.2 Notes

1. Note that the concatenate function returns a type of an unconstrained array. Refer to Section 5.6 for more information.
2. Full access types employ the run-time system to allocate from the heap and other storage areas, making memory use unpredictable, timing analysis problematic, heap exhaustion and fragmentation a significant risk. It can also create unbounded aliasing problems.
3. Pool-specific access types use storage similarly to stack-based data types. However, they require careful implementation and use to ensure the algorithms are predictable.

4. Pools and general access types permit aliasing of data. See restrictions in Section 5.2.
5. Control flow analysis of access-to-subprogram types is intractable in the general case. At present there is no known model of this capability to support Symbolic Analysis. Timing Analysis is intractable in general but can be bounded by annotations and manual checks. Object Code Analysis is difficult but not intractable. This exclusion can be enforced by the **pragma Restrictions(No_Access_Subprograms)**. When used with static locations and linker tools, they can be used as a means of system reconfiguration.
6. Controlled types introduce hidden control flows due to user-defined initialization, assignment and, especially, finalization. These are hard to review, analyse or test, particularly in error conditions.
7. Indefinite objects consume time and storage in ways which are difficult, if not impossible, to predict. Their dependence on run-time values complicates analysis. Consequently, these objects should not be used in high integrity systems.
8. Arrays with bounds which are not static complicate analysis of resources used. Time and memory used depends on dynamic bounds. Analysis of data access based on array indexing is further complicated if the bounds are unknown until run-time.

5.11.3 Guidance

As noted in the introduction to this section, the use of dynamic mechanisms is to be minimized in high integrity systems. Appropriate enforcements can be provided by the use of **pragma Restrictions (No_Implicit_Heap_Allocation)**, **pragma Restrictions(No_Allocators)**, **pragma Restrictions (No_Access_Subprograms)**.

Although the evaluation table has 'Inc' against nearly all the testing based verification techniques, it should be noted that the effectiveness of these techniques may well be reduced. For example, a problem arising from the inappropriate use of aliasing may well be difficult to find during Requirements-based Testing and Structure-based Testing. It is also true that code inspection techniques will be made more complex (and hence error prone) by the use of these dynamic features.

5.12 Exceptions

Ada has well-defined semantics even under error conditions. The language allows the user to detect error conditions and to specify the required behaviour under such conditions at run-time.

Predefined exceptions concern error conditions detected by the run-time environment. The implicit raising of predefined exceptions poses some problems where high integrity software is concerned because the location in the program and the exact time at which exceptions are actually raised cannot be easily predicted.

A predefined exception is raised automatically when an associated constraint is violated. In contrast, user exceptions provide the means to specify error conditions whose occurrences must be explicitly detected.

The exception mechanism leads the designer to an intellectual dilemma:

- its use makes the verification more difficult, so the use of exceptions should perhaps be prohibited;
- its use allows residual errors to be detected and handled, so the exception features are potentially a key part of a language for high integrity applications [NRC] [13].

Solutions to this conflict will be proposed in the guidance subsection below.

If the use of exceptions is prohibited this does not in itself prevent predefined exceptions from being raised. The use of the **pragma Restrictions(No_Exceptions)** is possible, but the program will become erroneous if a run-time exception does occur.

The evaluation table below has been produced with the assumption that exceptions are to be used.

5.12.1 Evaluation

Table 14 : Exceptions

Feature	Group Name								
	FA	SA	RC	SU	TA	OMU	OCA	RT	ST
Predefined exceptions ¹	All ^{d2}	All ^{d2}	Inc	Inc	Inc	Inc	All ^{d3}	Inc	All ^{d4}
Declaration (user)	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc	Inc
Raising (user)	All ^{d5}	All ^{d5}	Inc	Inc	Inc	Inc	Inc	Inc	All ^{d5}
Handler (predefined)	All ^{d6}	All ^{d6}	Inc	Inc ⁷	Inc ⁸	Inc	Inc	Inc	All ^{d9}
Handler (user)	All ^{d5}	All ^{d5}	Inc	Inc	Inc	Inc	All ^{d5}	Inc	Inc
Propagation	Exc ¹⁰	Exc ¹⁰	Inc	Inc	All ^{d10}	Inc	All ^{d10}	Inc	All ^{d10}

5.12.2 Notes

1. Predefined exceptions should not be raised explicitly, because such a raising cannot be distinguished from implicitly raised exceptions.
2. Symbolic Analysis and Flow Analysis are problematic, as the locations at which predefined exceptions are raised are not generally known.
3. Object Code Analysis is difficult because detection is either explicit or implicit via a hardware mechanism.
4. It is difficult to find a test sequence that will bring about the error conditions necessary to raise the predefined exception.
5. Raise statements from deeply nested scopes can be obscure and cause difficulties for Flow Analysis and Object Code Analysis.
6. The state immediately prior to the execution of the handler cannot be established making Symbolic Analysis intractable. Similarly, Flow Analysis is intractable since the point of the raising of a predefined exception cannot be established. With careful coding, an exception handler can be made to establish a recovery state that is independent of the implementation-dependent state when the handler is invoked.
7. The stack is used or not to reach the handler depending on the technique used to implement the exception mechanism: conventionally, the Static Mapping which uses a table or the Dynamic Tracking which uses the stack.
8. Timing Analysis depends on the technique used to implement the exception mechanism.
9. It was previously mentioned that it is not possible to control the program from its inputs to raise a predefined exception and thus to execute the associated handler.
10. Once an exception has been propagated, then neither Flow Analysis nor Symbolic Analysis can be undertaken and other forms of verification become difficult. The propagation of an exception cannot guarantee the current machine state and state of the subprogram which raised the exception, as well as possibly global or accessed objects. These conditions make future Flow Analysis and Symbolic Analysis of the program intractable and other analysis difficult.

5.12.3 Guidance

In the introduction, the conflict concerning the use of the exception mechanism was highlighted. On one hand, the previous section dealing with evaluation shows the difficulties induced by this mechanism when applying verification techniques. So, its use seems to be prohibited when high integrity systems are concerned. On the other hand, it is not possible to say that its non-use just reduces the language expression capabilities as it occurs for the other features. Indeed, when exceptions are not used, the errors cannot be handled but their existence is not avoided [13]. To solve this issue, the proposed guidelines concern three design strategies.

A first design choice (called 'exclusion strategy') consists of excluding the use of exceptions. This may be obtained by means of the **pragma Restrictions** (No_Exceptions) written in the Ada programs or by using compiler options provoking the same effect. However, the absence of erroneous states usually leading to an exception raising must be demonstrated. This is easily

obtained concerning potential user exceptions as their raising is explicit. The proof that no predefined exceptions can be raised is more complex, but may still be tractable (see 2.3.5).

The second design choice (called ‘belt-and-braces strategy’) seeks to avoid dependency on the exception mechanism but recognizes that a predefined exception may nevertheless occur for some unforeseen reason. The code is designed to be exception-free, perhaps demonstrated by suitable proof (see exclusion strategy) but an ‘others’ exception handler is introduced at the outermost scope level which does a very simple restart/reset or halt operation (if such an operation is compatible with the application). As signalled in the notes, the actual state of the variables is not well defined when the exception handler is reached. So, this strategy requires the implementation of a mechanism ensuring the recovery of a well-known program state at program resumption time (for instance, a ‘recovery cache mechanism’). It should be noted that the implementation of such a strategy may be very complex when the program includes tasks. In particular, resumption of the task in which an exception raising occurs is frequently not acceptable when synchronizations exist. In this case, a more complex resumption policy must be considered to handle the phenomenon called ‘domino effect’.

The third design choice (called ‘containment strategy’) authorizes the use of exceptions in a simple way. In particular, the following guidelines must be considered to make easier the verification techniques applications. Exception mechanisms can only be used (if required) to handle errors occurring at run-time and not rare events, such as ‘end of file is reached’. Predefined exceptions should not be raised explicitly. Predefined and user exceptions must be handled close to the raising location. In particular, propagation phenomenon should be avoided. Here again, the exception handling must guarantee that the program state is well-defined.

5.13 Tasking

High integrity systems traditionally do not make use of high-level language features such as concurrency. With Ada, these language features can be prohibited by the use of **pragma Restrictions (Max_Tasks=0)**, **pragma Restrictions (No_Protected_Types)** and **pragma Restrictions (No_Delay)**. The view that tasking should not be used is despite the fact that such systems are inherently concurrent. Concurrency is viewed as a ‘systems’ issue. It is visible during design and in the construction of the cyclic executive that implements the separate code fragments, but it is not addressed within the software production phases. Notwithstanding this approach, the existence of an extensive range of concurrency features within Ada does allow concurrency to be expressed at the language level with the resulting benefits of having a standard analysable approach that can be checked by the compiler and supported by other tools.

The requirement to analyse both the functional and temporal behaviour of high integrity systems imposes a number of restrictions on the concurrency model that can be employed. These restrictions then impact on the language features that are needed to support the model. Typical features of the concurrency model are as follows.

- a. A fixed number of tasks.
- b. Each task has a single invocation event, but has a potentially unbounded number of invocations. The invocation event can either be temporal (for a time-triggered task) or a signal from either another task or the environment. A high integrity application may restrict itself to only time-triggered tasks.
- c. Tasks only interact via the use of shared data. Updates to any shared data must be atomic.

These constraints furnish a model that can be implemented using fixed priority scheduling (either pre-emptive or non pre-emptive) and analysed in a number of ways:

- a. The functional behaviour of each task can be verified using the techniques appropriate for sequential code. Shared data is viewed as just environmental input when analysing a task. Timing analysis can ensure that such data is appropriately initialized and temporally valid.
- b. Following the assignment of temporal attributes to each task (period, deadline, priority etc), the system-wide timing behaviour can be verified using the proven techniques in fixed priority analysis [1, 2].

To implement this concurrency model in Ada requires only a small selection of the available tasking features. At the Eighth International Real-Time Ada Workshop (1997) the following profile (called the *Ravenscar profile*) was defined for high integrity, efficient, real-time systems [9].

The Ravenscar Profile is defined by the following.

- a. Task type and object declarations at the library level (that is, no hierarchy of tasks).

- b. No unchecked deallocation of protected and task objects.
- c. No dynamic allocation of task or protected objects (this is not part of the profile but is included here to be consistent with an overall approach to dynamic allocation — see 3.1.2).
- d. Tasks are assumed to be non-terminating.
- e. Library level protected objects with no entries (to ensure atomic updates to shared data).
- f. Library level protected objects with a single entry (for invocation signalling). This entry has a barrier consisting of a single boolean variable; moreover only a single task may queue on this entry.
- g. 'Real-Time' package.
- h. Atomic and volatile pragmas.
- i. **delay until** statements
- j. Count attribute for protected entries (but not within entry barriers).
- k. Task identifiers.
- l. Task discriminants.
- m. Protected procedures as interrupt handlers.

It follows that the following tasking features are not included in the profile: task types and objects other than at the library level, task hierarchies, unchecked de-allocation of protected and task objects, requeue, ATC, abort task entries, dynamic priorities, calendar, relative delays, protected types other than at the library level, protected entries with barriers other than a single Boolean variable declared within the same protected type, attempts to queue more than one task on a single protected entry, locking policies other than ceiling locking, scheduling policies other than FIFO within priorities, all forms of select statement, and user-defined task attributes.

The inclusion of protected entries allows event based scheduling to be used. For many high integrity systems only time triggered actions are employed, hence such entries and their associated interrupt handlers are not required.

The profile defines dispatching to be *FIFO within priority* with protected objects having *Ceiling Locking*. However it also allows a non pre-emptive policy to be defined. Co-operative scheduling (that is, non pre-emption between well defined system calls such as **delay until** or the call of a protected object) can reduce the cost of testing as pre-emption can only occur at well-defined points in the code. It can also reduce the size of the run-time.

With either dispatching policy, the Ravenscar profile can be supported by a relatively small run-time. It is reasonable to assume that a purpose-built run-time (supporting only the profile) would be efficient and 'certifiable' (i.e., built with the evidence necessary for its use in a certified system). An equivalent run-time for a constrained Ada83 tasking model has already been used in a certified application. Static checks are possible to ensure any program conforms to the profile.

Not only does the use of Ada increase the effectiveness of verification of the concurrency aspects of the application, it also facilitates a more flexible approach to the system's timing requirements. The commonly used cyclic executive approach imposes strict constraints on the range and granularity of periodic activities. The Ravenscar profile will support any range and a fine level of granularity. So, for example, tasks with periods of 50ms and 64ms can be supported together. Moreover, changes to the timing attributes of activities only require a re-evaluation of the timing analysis. Cyclic executives are hard to maintain and changes can lead to complete reconstruction. Finally, note that the inclusion of a small number of event triggered activities does not fundamentally change the structure of the concurrent program or the timing analysis, but it does impose significant problems for the cyclic executive. Polling for 'events' is a common approach in high integrity systems; but if the 'event' is rare and the deadline for dealing with the event is short then the time triggered approach is very resource intensive. The event triggered approach will work with much less resources. The guidelines are not intended to imply that event triggering is better than time triggering. The point here, is that the Ravenscar profile deals with both approaches and the migration from one to the other.