

INTERNATIONAL STANDARD

ISO/IEC
14776-112

First edition
2002-05

**Information technology –
Small computer system interface (SCSI) –**

**Part 112:
Parallel Interface-2 (SPI-2)**



Reference number
ISO/IEC 14776-112:2002(E)

IECNORM.COM : Click to view the full PDF of ISO/IEC 14776-112:2002

INTERNATIONAL STANDARD

ISO/IEC 14776-112

First edition
2002-05

Information technology – Small computer system interface (SCSI) – Part 112: Parallel Interface-2 (SPI-2)

© ISO/IEC 2002

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 the publisher.

ISO/IEC Copyright Office • Case postale 56 • CH-1211 Genève 20 • Switzerland



PRICE CODE

XC

For price, see current catalogue

Contents

	Page
Introduction	xiv
1 Scope	1
2 Normative references	2
2.1 Normative references	2
2.2 Approved references	2
2.3 References under development	3
3 Definitions, symbols, abbreviations, and conventions	3
3.1 Definitions	3
3.2 Symbols and abbreviations	11
3.3 Keywords	11
3.4 Conventions	12
3.5 Notation for Procedures and Functions	13
4 General	14
4.1 Overview	14
4.1.1 Data transfer modes	14
4.1.2 Cables, Connectors, Signals, Transceivers	14
4.1.3 Physical architecture of bus	14
4.1.4 Physical topology details and definitions	15
4.1.5 Bus loading	17
4.1.6 Termination requirements	18
4.1.7 Device Addressing	18
4.1.8 Protocol	18
5 SCSI parallel interface connectors	21
5.1 SCSI parallel interface connectors overview	21
5.2 Nonshielded connector	21
5.2.1 Nonshielded connector alternative 1 - A cable	21
5.2.2 Nonshielded connector alternative 2 - A cable	21
5.2.3 Nonshielded connector alternative 3 - P cable and Q cable	21
5.2.4 Nonshielded connector alternative 4	22
5.3 Shielded connector	27
5.3.1 Shielded connector alternative 1 - A cable	28
5.3.2 Shielded connector alternative 2 - A cable	28
5.3.3 Shielded connector alternative 3 - P cable and Q cable	28
5.3.4 Shielded connector alternative 4 - P cable and Q cable	28
5.4 Connector contact assignments	37
5.4.1 SE assignments	38
5.4.2 Differential connector contact assignments	42
6 SCSI bus interconnect	49
6.1 SCSI bus interconnect overview	49
6.2 SCSI bus cables	49
6.3 Interconnect characteristics for signals	50
6.4 Decoupling characteristics for TERMPWR and TERMPWRQ lines	52
6.5 Connection requirements for RESERVED lines	53
6.6 Cables used with SE transceivers	53
6.6.1 SE ground offset	54
6.7 Cables used with differential transceivers	54
6.7.1 HVD stub length and spacing	54
6.7.2 LVD stub length and spacing	55
7 SCSI parallel interface electrical characteristics	56
7.1 SCSI parallel interface electrical characteristics overview	56
7.2 SE alternative	57

7.2.1 SE termination	57
7.2.2 SE output characteristics	57
7.2.3 SE input characteristics	60
7.2.4 SE input and output characteristics	61
7.3 LVD alternative	62
7.3.1 LVD termination	62
7.3.2 LVD driver characteristics	66
7.3.3 LVD receiver characteristics	66
7.3.4 LVD capacitive loads	67
7.3.4.1 Management of LVD release glitches	69
7.3.5 SE/HVD transmission mode detection	69
7.3.5.1 LVD DIFFSENS driver	70
7.3.5.2 LVD DIFFSENS receiver	70
7.4 LVD/MSE multimode alternative	72
7.4.1 LVD/MSE multimode termination	72
7.4.2 LVD/MSE multimode transceiver characteristics	73
7.4.3 Power for LVD/MSE multimode transceivers	74
7.4.4 Ground drivers	74
7.5 HVD alternative	75
7.5.1 HVD termination	75
7.5.2 HVD output characteristics	76
7.5.3 HVD input characteristics	78
7.5.4 SE driver protection	79
7.6 Terminator power	80
8 SCSI bus signals	82
8.1 SCSI bus signals overview	82
8.2 Signal descriptions	83
8.3 Parity checking rules	84
8.4 Signal states	84
8.4.1 SE	84
8.4.2 Differential	85
8.5 OR-tied signals	86
8.6 Signal sources	86
9 SCSI parallel bus timing	88
9.1 SCSI parallel bus timing values	88
9.2 Timing description	89
9.2.1 Arbitration delay	89
9.2.2 Bus clear delay	90
9.2.3 Bus free delay	90
9.2.4 Bus set delay	90
9.2.5 Bus settle delay	90
9.2.6 Cable skew	90
9.2.7 Data release delay	90
9.2.8 Disconnection delay	90
9.2.9 Power on to selection	90
9.2.10 Receive assertion period	90
9.2.11 Receive hold time	91
9.2.12 Receive negation period	91
9.2.13 Receive setup time	91
9.2.14 Receive period tolerance	91
9.2.15 Reset hold time	91
9.2.16 Reset to Selection	91
9.2.17 Selection abort time	91
9.2.18 Selection time-out delay	91
9.2.19 Signal Timing Skew	91
9.2.20 System deskew delay	91
9.2.21 Transmit assertion period	92
9.2.22 Transmit hold time	92
9.2.23 Transmit negation period	92

9.2.24 Transmit setup time	92
9.2.25 Transmit period tolerance	92
9.3 Measurement points	92
9.3.1 SE fast-10 data transfer rates	92
9.3.2 SE fast-20 data transfer rates	93
9.3.3 LVD	94
9.3.4 HVD	95
9.4 Setup and hold timings	96
9.4.1 Fast-10 data transfer rates	96
9.4.2 Fast-20 data transfer rates	98
9.4.3 Fast-40 data transfer rates	100
10 Removal and insertion of SCSI devices	102
10.1 Removal and insertion of SCSI devices overview	102
10.2 Case 1 - Power-off during removal or insertion	102
10.3 Case 2 - RST signal asserted continuously during removal or insertion	102
10.4 Case 3 - Current I/O processes not allowed during insertion or removal	102
10.5 Case 4 - Current I/O process allowed during insertion or removal	103
11 Logical characteristics	104
11.1 SCSI bus phases	104
11.1.1 SCSI bus phase overview	104
11.1.2 BUS FREE phase	104
11.1.3 Unexpected bus free	104
11.1.4 ARBITRATION phase	105
11.1.5 SELECTION phase	105
11.1.6 SELECTION time-out procedure	106
11.1.7 RESELECTION phase	106
11.1.7.1 RESELECTION phase overview	106
11.1.7.2 RESELECTION	106
11.1.7.3 RESELECTION time-out procedure	107
11.1.8 Information transfer phases	107
11.1.8.1 Information transfer phases overview	107
11.1.8.2 Asynchronous information transfer	108
11.1.8.3 Synchronous data transfer	108
11.1.8.4 Wide data transfer	109
11.1.9 COMMAND phase	111
11.1.10 Data phase	111
11.1.10.1 Data phase overview	111
11.1.10.2 DATA IN phase	111
11.1.10.3 DATA OUT phase	111
11.1.11 STATUS phase	111
11.1.11.1 STATUS phase overview	111
11.1.11.2 STATUS phase exception condition handling	111
11.1.12 Message phase	111
11.1.12.1 Message phase overview	111
11.1.12.2 MESSAGE IN phase	111
11.1.12.3 MESSAGE IN phase exception condition handling	112
11.1.12.4 MESSAGE OUT phase	112
11.1.12.5 MESSAGE OUT phase exception condition handling	112
11.1.13 Signal restrictions between phases	112
11.2 SCSI bus conditions	113
11.2.1 SCSI bus conditions overview	113
11.2.2 Attention condition	113
11.2.3 Hard reset	113
11.2.4 Reset events	114
11.2.5 Transceiver mode change reset event	114
11.3 SCSI bus phase sequences	114
11.4 SCSI pointers	115
11.5 SCSI interlocked protocol messages	115
11.5.1 SCSI interlocked protocol messages overview	115

11.5.2 Message protocols and formats	116
11.5.2.1 Message protocol rules	116
11.5.2.2 Message formats	116
11.5.2.2.1 Message formats overview	116
11.5.2.2.2 One-byte messages	117
11.5.2.2.3 Two-byte messages	117
11.5.2.2.4 Extended messages	117
11.5.3 Link control messages	118
11.5.3.1 Link control message codes	118
11.5.3.2 CONTINUE TASK	118
11.5.3.3 DISCONNECT	119
11.5.3.4 IDENTIFY	119
11.5.3.5 IGNORE WIDE RESIDUE	120
11.5.3.6 INITIATOR DETECTED ERROR	121
11.5.3.7 MESSAGE PARITY ERROR	121
11.5.3.8 MESSAGE REJECT	121
11.5.3.9 MODIFY DATA POINTER	122
11.5.3.10 NO OPERATION	122
11.5.3.11 RESTORE POINTERS	122
11.5.3.12 SAVE DATA POINTER	122
11.5.3.13 SYNCHRONOUS DATA TRANSFER REQUEST	122
11.5.3.13.1 SYNCHRONOUS DATA TRANSFER REQUEST overview	122
11.5.3.13.2 Target initiated SDTR negotiation	125
11.5.3.13.3 Initiator initiated SDTR negotiation	125
11.5.3.14 TARGET TRANSFER DISABLE	125
11.5.3.15 TASK COMPLETE	126
11.5.3.16 WIDE DATA TRANSFER REQUEST	126
11.5.3.16.1 WIDE DATA TRANSFER REQUEST overview	126
11.5.3.16.2 Target initiated WDTR negotiation	128
11.5.3.16.3 Initiator initiated WDTR negotiation	128
11.5.4 Task attribute messages	129
11.5.4.1 Task attribute messages overview	129
11.5.4.2 ACA	130
11.5.4.3 HEAD OF QUEUE	130
11.5.4.4 LINKED COMMAND COMPLETE	131
11.5.4.5 LINKED COMMAND COMPLETE (WITH FLAG)	131
11.5.4.6 ORDERED	131
11.5.4.7 SIMPLE	131
11.5.5 Task management messages	132
11.5.5.1 Task management message codes	132
11.5.5.2 ABORT TASK	132
11.5.5.3 ABORT TASK SET	132
11.5.5.4 CLEAR ACA	133
11.5.5.5 CLEAR TASK SET	133
11.5.5.6 LOGICAL UNIT RESET	133
11.5.5.7 TARGET RESET	133
11.5.5.8 TERMINATE TASK	133
11.6 Command processing considerations and exception conditions	133
11.6.1 Command processing considerations and exception conditions overview	133
11.6.2 Asynchronous event notification	134
11.6.3 Incorrect initiator connection	134
11.6.4 Unexpected reselection	135
11.7 Use of disconnect-reconnect page parameters	135
11.8 SCSI parallel interface services	136
11.8.1 SCSI parallel interface services overview	136
11.8.2 Procedure terms	136
11.8.3 Application client SCSI command services	138
11.8.3.1 Application client SCSI command services procedure call	138
11.8.3.2 Send SCSI command service	139
11.8.4 Device server SCSI command services	139
11.8.4.1 Device server SCSI command services procedure call	139

11.8.4.2 Data-in delivery service	139
11.8.4.3 Data-out delivery service	140
11.8.5 Task management services	140
11.8.5.1 Task management services procedure call	140
11.8.5.2 Task management function service	140
11.8.5.2.1 Task management function service overview	140
11.8.5.2.2 ABORT TASK	140
11.8.5.2.3 ABORT TASK SET	141
11.8.5.2.4 CLEAR ACA	141
11.8.5.2.5 CLEAR TASK SET	141
11.8.5.2.6 LOGICAL UNIT RESET	141
11.8.5.2.7 RESET SERVICE DELIVERY SUBSYSTEM	141
11.8.5.2.8 TARGET RESET	141
11.8.5.2.9 TERMINATE TASK	141
11.8.5.2.10 WAKEUP	141
Annex A Additional requirements for LVD SCSI drivers and receivers	142
A.1 System level requirements	142
A.2 Driver requirements	142
A.2.1 Differential output voltage, VS	143
A.2.2 Offset (common-mode output) voltage, VCM	145
A.2.3 Short-circuit currents, IO-S and IO+S	146
A.2.4 Open-circuit output voltages, $V_{O-(OC)}$ and $V_{O+(OC)}$	147
A.2.5 Output signal waveform	147
A.2.6 Dynamic output signal balance, VCM(PP)	149
A.3 Receiver characteristics	150
A.3.1 Receiver steady state input voltage requirements	151
A.3.2 Compliance test	151
A.3.3 Receiver setup and hold times	152
A.4 Transceiver characteristics	152
A.4.1 Transceiver output/input currents, II-L and II+L	152
A.4.2 Transceiver maximum input voltages	153
Annex B SCSI configured automatically (SCAM)	154
B.1 Model	154
B.2 Glossary	154
B.3 SCAM requirements	155
B.3.1 Configuration requirements	155
B.3.2 Timing requirements	155
B.3.2.1 SCAM tolerant power-on to selection delay	156
B.3.2.2 SCAM tolerant reset to selection delay	156
B.3.2.3 SCAM tolerant selection response time	156
B.3.2.4 SCAM unassigned ID selection response delay	156
B.3.2.5 SCAM power-on to SCAM selection delay	156
B.3.2.6 SCAM reset to SCAM selection delay	156
B.3.2.7 SCAM selection response time	156
B.3.2.8 Recommended SCAM selection response time	156
B.3.2.9 Wide arbitration time	156
B.3.3 Device requirements	157
B.3.3.1 SCAM tolerant target	157
B.3.3.2 Level 1 SCAM initiator	157
B.3.3.3 Level 1 SCAM target	157
B.3.3.4 Level 2 SCAM initiator	158
B.3.3.5 Level 2 SCAM target	158
B.4 SCAM protocol	159
B.4.1 Initiation	159
B.4.1.1 Transfer cycles	159
B.4.1.2 Wired-OR glitch filtering	160
B.4.1.3 Isolation stage	161

B.4.1.4 Function sequences	164
B.4.1.4.1 Isolate function	165
B.4.1.4.2 Isolate and set priority flag function	166
B.4.1.4.3 Configuration process complete function	166
B.4.1.4.4 Dominant initiator contention function	167
B.5 SCAM operations	167
B.5.1 SCAM initiator	167
B.5.1.1 Dominant SCAM initiator	168
B.5.1.1.1 SCSI ID categorization	168
B.5.1.1.2 SCSI ID assignment	168
B.5.1.2 Subordinate SCAM initiator	168
B.5.2 Level 1 SCAM target	169
B.5.3 Level 2 SCAM target	170
Annex C SCSI bus fairness	172
C.1 Model	172
C.2 Determining fairness by monitoring prior bus activity	172
C.3 Fairness algorithm	172
C.4 Additional comments	173
Annex D SCA-2 Connector pinouts	175
D.1 SCA-2 Signal Definitions	175
D.1.1 VOLTAGE and GROUND signals	175
D.1.2 CHARGE signals	175
D.1.3 SPINDLE SYNC	176
D.1.4 ACTIVE LED OUT	177
D.1.5 Motor Start Controls	177
D.1.6 SCSI ID Selection	178
D.1.7 MATED Signals	179
D.1.7.1 MATED 2/Drive Side	179
D.1.7.2 MATED 2/Backplane Side	180
D.1.7.3 MATED 1/Drive Side	180
D.1.7.4 MATED 1/Backplane Side	180
Annex E Interconnecting buses of different widths	182
Annex F Cabling and cable measurement method recommendations	185
F.1 Cabling	185
F.2 Cable measurement	185
F.2.1 Impedance, TDR, single-ended	185
F.2.2 Impedance, TDR, differential	185
F.2.3 Attenuation, differential	186
F.2.4 Velocity (propagation delay) and skew	186
Annex G Transmission line considerations for fast-20 data transfer rates	187
Annex H Measuring SE pin capacitance	190
Annex I SCSI ICONS	191
Annex J Example of HVD terminator network	193
Annex K SCSI-3 to SCSI-2 terminology mapping	194
Annex L SCA-2 Unshielded Connections	195
L.1 Definitions and Conventions	196
L.1.1 Definitions	196

L.2 General Description	198
Annex M VHDCI Shielded Configurations	200
M.1 Definitions and Conventions	200
M.1.1 Definitions	200
M.2 General Description	202

IECNORM.COM : Click to view the full PDF of ISO/IEC 14776-112:2002

Tables

	Page
1 Transceiver/speed support map	14
2 Cross-reference to A cable contact assignments	37
3 SE contact assignments - A cable	38
4 SE contact assignments - P cable	39
5 SE contact assignments - Q cable	40
6 SE contact assignments - nonshielded alternative 4 connector	41
7 LVD contact assignments - A cable	42
8 LVD contact assignments - P cable	43
9 LVD contact assignments - Q cable	44
10 HVD contact assignments - A cable	45
11 HVD contact assignments - P cable	46
12 HVD contact assignments - Q cable	47
13 HVD and LVD contact assignments - nonshielded alternative 4 connector	48
14 Recommended minimum conductor size	50
15 SE characteristic impedance of cable at maximum indicated data transfer rate	51
16 Differential characteristic impedance of cable at indicated data transfer rate	52
17 SE maximum distance between terminators	53
18 Differential maximum distance between terminators	54
20 Minimum stub connection spacing rules for HVD SCSI devices	55
19 HVD stub length and spacing requirements	55
21 Minimum stub connection spacing rules for LVD SCSI devices	55
22 Electrical input requirements at the device connector	56
23 Input current requirements at the device connector for lines not being driven by the device	57
24 SE steady state output voltage characteristics	58
25 SE input voltage characteristics	61
26 SE input and output electrical characteristics	62
27 I-V requirements for LVD impedance, common mode impedance, and V_{BIAS} tests	64
28 Values for LVD termination balance test	65
29 Values for LVD capacitive loads	68
30 Glitch management requirements for devices using LVD drivers	69
31 LVD DIFFSENS driver specifications	70
32 DIFFSENS receiver operating requirements	71
33 I-V requirements for HVD impedance and common mode impedance tests	76
34 Values for HVD termination balance test	76
35 HVD output voltage characteristics	77
37 Values for HVD capacitive loads	79
36 HVD input voltage characteristics	79
38 Terminator power characteristics at the terminator	81
39 Arbitration priorities by SCSI ID	83
40 Parity checking rules	84
41 Signal sources	87
42 SCSI bus timing values	89
43 Information transfer phases	108
44 Wide SCSI byte order	110
45 Message format	117
46 Extended message format	117
47 Link control message codes	118
48 IDENTIFY message format	119
49 IGNORE WIDE RESIDUE message format	120
50 IGNORE field definition	121
51 MODIFY DATA POINTER message format	122
52 SYNCHRONOUS DATA TRANSFER message format	123
53 TRANSFER PERIOD FACTOR field	123
54 SDTR messages implied agreements	124
55 WIDE DATA TRANSFER message format	126
56 WDTR messages implied agreements	128
57 Task attribute message codes	130
58 ACA message format	130
59 HEAD OF QUEUE message format	130

60 ORDERED message format	131
61 SIMPLE message format	131
62 Task management message codes	132
63 DATA TRANSFER DISCONNECT CONTROL	136
64 SCSI Parallel Interface-2 Standard terms mapped to terms from other SCSI-3 standards	136
65 Procedure terms	138
66 Processing of send SCSI command service procedure	139
67 Processing of data-in delivery service procedure	140
68 Processing of data-out delivery service procedure	140
A.1 System level requirements	142
A.2 Driver steady-state test limits and conditions	143
A.3 Driver switching test circuit parameters	148
A.4 Dynamic output balance limits	150
A.5 Receiver steady state input voltage ranges	151
A.6 Receiver minimum and maximum input voltages.	151
B.1 SCAM timing values	155
B.2 Transfer cycle conditions	162
B.3 Identification string	163
B.4 Type code	163
B.5 MAXIMUM ID CODE	163
B.6 ID VALID	164
B.7 Function codes	165
B.8 Action codes	166
B.9 Dominance preference code	167
D.1 Voltage specification limits	175
D.2 Charge supply to SCSI device	176
D.3 Output characteristics of drive 'ACTIVE LED OUT signal	177
D.4 Definition of motor start controls	178
D.5 Electronic requirements for input controls	178
D.6 SCSI device ID selection signals	179
K.1 SPI-2 to SCSI-2 terminology mapping	194

IECNORM.COM : Click to view the full PDF of ISO/IEC 14776-112:2002

Figures

	Page
1 General Structure of SCSI	2
2 Differential SCSI bus	15
3 SE SCSI bus	15
4 SCSI bus topology details	17
5 SCSI Parallel Interface service reference mode	19
6 Model for a four step confirmed services	19
7 Model for a two step confirmed services	20
8 50/68-contact alternative 1/alternative 3 nonshielded SCSI device connector	23
9 50/68-contact alternative 1/alternative 3 nonshielded mating connector	24
10 50-contact alternative 2 nonshielded SCSI device connector (A cable)	25
11 50-contact alternative 2 nonshielded mating connector (A cable)	26
12 80-contact alternative 4 nonshielded SCSI device connector (P cable)	27
13 80-contact alternative 4 nonshielded contact positions	27
14 50-contact alternative 1 shielded SCSI device connector (A cable)	29
15 50-contact alternative 1 shielded mating connector (A cable)	30
16 50-contact alternative 2 shielded SCSI device connector (A cable)	31
17 50-contact alternative 2 shielded mating connector (A cable)	32
18 68-contact alternative 3 shielded SCSI device connector (P cable/Q cable)	33
19 68-contact alternative 3 shielded mating connector (P cable/Q cable)	34
20 68-contact alternative 4 shielded SCSI device connector (P cable/Q cable)	35
21 68-contact alternative 4 shielded contact positions (P cable/Q cable)	36
22 Terminator decoupling example	53
23 Active negation current vs. voltage	59
24 SE A.C. test circuit	60
25 Differential SPI-2 bus terminator	62
26 Test circuit for terminator differential impedance	63
27 Termination I-V characteristics for differential and common mode impedance tests	63
28 Test circuit for termination common mode impedance test	64
29 Termination balance test configuration	65
30 Termination balance test data definition	65
31 LVD transceiver architecture	66
32 LVD receiver example	67
33 Capacitive loads	68
34 LVD DIFFSENS driver signal definitions	70
35 DIFFSENS receiver function	71
36 LVD DIFFSENS receiver example	72
37 Multimode terminator architecture	73
38 Multimode transceiver architecture	74
39 HVD test circuit	78
40 SE driver protection circuit	80
41 Voltage and current definitions	85
42 Signaling sense	86
43 Fast-10 SE timing measurement points	93
44 Fast-20 SE timing measurement points	94
45 LVD timing measurement points	95
46 HVD timing measurement points	96
47 Fast-10 setup and hold times for SE applications	97
48 Fast-10 setup and hold timing for HVD applications	97
49 Fast-20 setup and hold times for SE applications	99
50 Fast-20 setup and hold timing for HVD applications	99
51 Fast-40 System setup and hold timings (all times in ns)	101
52 Phase sequences	115
A.1 Differential steady-state output voltage test circuit	144
A.2 Domain for driver assertion and negation levels	145
A.3 Driver offset steady-state voltage test circuit	146
A.4 Common mode output voltage test	146
A.5 Driver short-circuit test circuit	147
A.6 Open-circuit output voltage test circuit	147
A.7 Differential output switching voltage test circuit	148

A.8 Driver output signal waveform	149
A.9 Driver offset switching voltage test circuit	150
A.10 Receiver input voltage threshold test circuit	152
A.11 Transceiver off-state output current test circuit	153
B.1 Transfer cycles	160
B.2 Level 1 SCAM target states	169
B.3 Level 2 SCAM target states	170
D.1 Sample circuit for mated indications	181
E.1 Interconnecting single-ended A and P cables	183
E.2 Interconnecting HVD A and P cables	184
F.1 Differential impedance measurement	186
G.1 Minimum device spacing versus bus and device capacitance	189
I.1 SE icon for SCSI	191
I.2 LVD icon for SCSI	191
I.3 SE/LVD multimode icon for SCSI	192
I.4 HVD icon for SCSI	192
J.1 Example of termination network for HVD devices	193
L.1 Mating side gender definition	198
L.2 Contact positioning architecture	199
M.1 Mating side gender definition	202

IECNORM.COM : Click to view the full PDF of ISO/IEC 14776-112:2002

**INFORMATION TECHNOLOGY –
SMALL COMPUTER SYSTEM INTERFACE (SCSI) –
Part 112: Parallel Interface-2 (SPI-2)**

FOREWORD

- 1) ISO (International Organization for Standardization) and IEC (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.
- 2) In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.
- 3) Attention is drawn to the possibility that some of the elements of this International Standard may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

International Standard ISO/IEC 14776-112 was prepared by subcommittee 25: Interconnection of information technology equipment, of ISO/IEC joint technical committee 1: Information technology.

ISO/IEC 14776-112 is intended to be used in conjunction with ISO/IEC 14776-311¹. The resulting interface facilitates the interconnection of computers and intelligent peripherals and thus provides a common interface standard for both system integrators and suppliers of intelligent peripherals.

Annexes A, B, C and D form an integral part of this standard.

Annexes E, F, G, H, I, J, K, L and M are for information only.

¹ Under consideration.

Introduction

The SCSI protocol is designed to provide an efficient peer-to-peer I/O bus with the maximum number of hosts and peripherals determined by the bus width (typically 8 or 16 with 32 allowed). Data may be transferred asynchronously or synchronously at rates that depend primarily on device implementation and cable length.

SCSI is an I/O interface that may be operated over a wide range of media and transfer rates. The objectives of the parallel interface in SCSI are:

- a) To provide host computers with device independence within a class of devices. Thus, different disk drives, tape drives, printers, optical media drives, and other devices may be added to the host computers without requiring modifications to generic system hardware. Vendor unique indications are accommodated. Reserved areas are provided for future standardization.
- b) To provide interoperability with SCSI-2 devices. Devices meeting SCSI-2 and the SCSI Parallel Interface-2 standards co-exist on the same bus. SCSI-3 devices should be permissive of the SCSI-2 or SCSI-3 compliant behavior of other devices including those not implementing optional extensions of the SCSI Parallel Interface-2 Standard.

The interface protocol includes provision for the connection of multiple initiators (SCSI devices capable of initiating a task) and multiple targets (SCSI devices capable of responding to a request to perform a task). Distributed arbitration (i.e., bus-contention logic) is built into the architecture of parallel SCSI. A default priority system awards interface control to the highest priority SCSI device that is contending for use of the bus and an optional fairness algorithm is defined.

This international standard combines the functionality of the SCSI-3 Interlocked Protocol (SIP) standard, the SCSI-3 Parallel Interface (SPI) standard, and the SCSI-3 Fast-20 standard. This international standard is intended to be the choice for new designs over the above standards. In addition to combining the above standards several new features have been added into this international standard. Some of those features are; Low Voltage Differential (LVD), fast-40 data transfers, multimode signal-ended (MSE), and additional connectors.

The SCSI Parallel Interface-2 standard is divided into the following clauses:

- Clause 1 is the scope;
- Clause 2 enumerates the normative references that apply to this standard;
- Clause 3 describes the definitions, symbols, conventions and abbreviations used in this standard;
- Clause 4 describes the SCSI parallel interface model used in this standard;
- Clause 5 describes the connectors;
- Clause 6 describes the cable characteristics;
- Clause 7 describes the electrical characteristics;
- Clause 8 describes the SCSI bus signals;
- Clause 9 describes the bus timing;
- Clause 10 describes the removal and insertion of parallel SCSI devices;
- Clause 11 describes the SCSI parallel protocol characteristics;

Small Computer System Interface - Part 112: Parallel Interface-2 (SPI-2)

1 Scope

This international standard defines the mechanical, electrical, timing, and protocol requirements of the SCSI parallel interface to allow conforming devices to interoperate. The SCSI parallel interface is a local I/O bus that may be operated over a wide range of transfer rates. The objectives of the SCSI parallel interface are

- a) To provide host computers with device independence within a class of devices. Thus, different disk drives, tape drives, printers, optical media drives, and other devices may be added to the host computers without requiring modifications to generic system hardware. Provision is made for the addition of special features and functions through the use of vendor-specific options. Reserved areas are provided for future standardization.
- b) To provide compatibility such that properly conforming SCSI-2 devices may interoperate with SCSI-3 devices given that the systems engineering is correctly done. Properly conforming SCSI-2 devices should respond in an acceptable manner to reject SCSI-3 protocol extensions. SCSI-3 protocol extensions are designed to be permissive of such rejections and thus allow the SCSI-2 devices to continue operation without requiring the use of the extension.

The interface protocol includes provision for the connection of multiple initiators (SCSI devices capable of initiating an I/O process) and multiple targets (SCSI devices capable of responding to a request to perform an I/O process). Distributed arbitration (i.e., bus-contention logic) is built into the architecture of SCSI. A default priority system awards interface control to the highest priority SCSI device that is contending for use of the bus and an optional fairness algorithm is defined.

This standard defines the physical attributes of an input/output bus for interconnecting computers and peripheral devices.

Figure 1 is intended to show the general structure of SCSI standards. The figure is not intended to imply a relationship such as a hierarchy, protocol stack, or system architecture.




Figure 1 - General Structure of SCSI

This international standard is intended as an alternate to the SCSI-3 Parallel Interface Standard. This international standard, in addition to containing all the information in the SCSI-3 Parallel Interface Standard contains information and specifications for LVD and fast-40 along with many other improvements.

2 Normative references

2.1 Normative references

The following normative documents contain provisions which, through reference in this text, constitute provisions of this part of ISO/IEC 14776. For dated references, subsequent amendments to, or revisions of, any of these publications do not apply. However, parties to agreements based on this part of ISO/IEC 14776 are encouraged to investigate the possibility of applying the most recent editions of the normative documents indicated below. For undated references, the latest edition of the normative document referred to applies. Members of IEC and ISO maintain registers of currently valid International Standards.

2.2 Approved references

ISO/IEC 9316:1995, *Small Computer System Interface - 2*

ISO/IEC 8482:1993, *Information technology - Telecommunications and information exchange between*

systems - Twisted pair multipoint interconnections

ISO/IEC 14776-411, 1999, *SCSI-3 Architecture Model standard*

IEC 60512-2:1985, *Electromechanical components for electronic equipment; basic testing procedures and measuring methods - Part 2: General examination, electrical continuity and contact resistance tests, insulation tests and voltage stress tests*

IEC 60512-11-7:1996, *Electromechanical components for electronic equipment - Basic testing procedures and measuring methods - Part 11-7: Climatic tests - Test 11g: Flowing mixed gas corrosion test*

EIA-700A0AE (SP-3651), *Detail Specification for Trapezoidal Connectors with Non-removable Ribbon Contacts on 1,27 mm Pitch Double Row used with Single Connector Attachments (SCA-2)*

EIA-700A0AF (SP-3652), *Detail Specification for Trapezoidal Connector 0,8mm Pitch used with Very High Density Cable Interconnect (VHDCI)*

2.3 References under development

ISO/IEC 14776-412, *SCSI Architecture Model-2 standard*

ISO/IEC 14776-311, *SCSI-3 Primary Commands Standard*

3 Definitions, symbols, abbreviations, and conventions

3.1 Definitions

3.1.1

A cable:

A 50-conductor cable (see and 3.1.51) that provides a primary 8-bit DATA BUS and control signals.

3.1.2

ACKx:

A bus signal that is either the ACK or ACKQ signal.

3.1.3

agent:

Carries out the actions of a requested service following the rules of the protocol.

3.1.4

application client:

An object that is the source of SCSI commands. Further definition of an application client is found in the SCSI-3 Architecture Model standard.

3.1.5

asynchronous event notification:

A procedure used by targets to notify initiators of events that occur when a pending task does not exist for that initiator.

3.1.6

asynchronous transfer:

An information transfer that uses the asynchronous REQ/ACK handshake.

3.1.7**auto-contingent allegiance:**

A condition of a task set following the return of a CHECK CONDITION or COMMAND TERMINATED status. See the SCSI-3 Architecture Model standard for a detailed definition of auto-contingent allegiance.

3.1.8**bus path:**

The electrical path directly between the bus terminators.

3.1.9**byte:**

Indicates an 8-bit construct.

3.1.10**confirmation:**

The last step of a confirmed service informing the upper protocol layer that the requested service has been completed.

3.1.11**confirmed protocol service:**

A service available at the protocol service interface, that requires confirmation of completion.

3.1.12**contact:**

The electrically-conductive portion of a connector associated with a single conductor in a cable.

3.1.13**contingent allegiance:**

One of the conditions of a task set following the return of a CHECK CONDITION or COMMAND TERMINATED status. A detailed definition of contingent allegiance may be found in SCSI-2.

3.1.14**current task:**

A task that is in the process of sending messages, status, transferring data, or transferring command data to or from the initiator.

3.1.15**data bus:**

An 8-bit, 16-bit or 32-bit data bus (see 8.2).

3.1.16**device server:**

An object within the logical unit which executes SCSI tasks according to the rules for task management as described in the SCSI-3 Architecture Model standard.

3.1.17**differential:**

A signalling alternative that employs differential drivers and receivers to improve signal-to-noise ratios and increase maximum cable lengths (also see 3.1.69 SE). Includes both LVD (see 7.3) and HVD (see 7.5) differential drivers and receivers.

3.1.18

disconnect:

The action that occurs when a SCSI device releases control of the SCSI bus, allowing it to go to the BUS FREE phase.

3.1.19

driver:

The circuitry used to control the state of the bus.

3.1.20

exception condition:

Any event that causes a SCSI device to enter an auto-contingent allegiance or contingent allegiance condition.

3.1.21

fast-5:

Negotiated to receive synchronous data at a transfer rate less than or equal to 5 megatransfers per second. Note: Previous SCSI standards referred to this transfer range as slow.

3.1.22

fast-10:

Negotiated to receive synchronous data at a transfer rate greater than 5 megatransfers per second and less than or equal to a transfer rate of 10 megatransfers per second. Note: Previous SCSI standards referred to this transfer range as fast SCSI.

3.1.23

fast-20:

Negotiated to receive synchronous data at a transfer rate greater than 10 megatransfers per second and less than or equal to a transfer rate of 20 megatransfers per second.

3.1.24

fast-40:

Negotiated to receive synchronous data at a transfer rate greater than 20 megatransfers per second and less than or equal to a transfer rate of 40 megatransfers per second.

3.1.25

flag:

An abstraction indicating that the condition will be communicated to the recipient of the flag.

3.1.26

field:

A group of one or more contiguous bits.

3.1.27

indication:

The second step of a four step confirmed service in reply to a request.

3.1.28

initial connection:

An initial connection is the result of a connect. It exists from the assertion of the BSY signal (see 11.1.4) in a SELECTION phase until the next BUS FREE phase.

3.1.29

initiator:

An SCSI device containing application clients that originate device service and task management requests to be processed by a target SCSI device. See the SCSI-3 Architecture Model standard for a detailed definition of an initiator.

3.1.30

I_T nexus:

A nexus which exists between an initiator and a target.

3.1.31

I_T_L nexus:

A nexus which exists between an initiator, a target, and a logical unit. This relationship replaces the prior I_T nexus.

3.1.32

I_T_L_Q nexus:

A nexus between an initiator, a target, a logical unit, and a queue tag following the successful receipt of one of the queue tag messages. This relationship replaces the prior I_T_L nexus.

3.1.33

logical unit:

An externally addressable entity within a target that implements an SCSI device model. See the SCSI-3 Architecture Model standard for a detailed definition of a logical unit.

3.1.34

logical unit number:

An identifier for a logical unit.

3.1.35

I/O process:

An I/O process consists of one initial connection and zero or more reconnections, all pertaining to a single task or group tasks. An I/O process begins with the establishment of a nexus. An I/O process normally ends with the BUS FREE phase (see 11.1.2).

3.1.36

magnitude:

The absolute value of a number or quantity.

3.1.37

mandatory:

The referenced item is required to claim compliance with this standard.

3.1.38

megatransfers per second:

The repetitive rate at which data are transferred across the bus. This is equivalent to megabytes per second on an 8-bit wide bus.

3.1.39

message:

One or more bytes transferred between an initiator and a target to do link control, task management, and to attach task attributes to commands.

3.1.40

multimode single-ended (MSE):

A signalling alternative for multimode SCSI devices that employs MSE (see 7.4) drivers and receivers to allow multimode SCSI devices to operate when SE SCSI devices are present on the bus.

3.1.41

nexus:

A relationship between an initiator and a target that begins with an initial connection and ends with the completion of the associated I/O process. This relationship is formed as the result of a task.

3.1.42

object:

An architectural abstraction that encapsulates data types, services, or other objects that are related in some way.

3.1.43

odd parity:

Odd logical parity, where the parity bit is driven and verified to be that value that makes the number of assertions on the associated data byte plus the parity bit equal to an odd number (1, 3, 5, 7, or 9). See 3.1.19, parity bit.

3.1.44

one:

A true signal value or a true condition of a variable.

3.1.45

P cable:

A 68-conductor cable that provides the primary 16-bit DATA BUS and control signals.

3.1.46

parity bit:

A bit associated with a byte that is used to detect the presence of single-bit errors within the byte. The parity bit is driven such that the number of logical ones in the byte plus the parity bit is odd.

3.1.47

path:

The cable, printed circuit board or other means for providing the conductors and insulators that connect two or more points.

3.1.48

pending task:

A task that is not the current task.

3.1.49

port:

A single attachment to an SCSI bus from an SCSI device.

3.1.50

primary bus:

The bus that contains the first 8 or 16 bits of the data and the control signals used to operate SCSI devices.

3.1.51

Q cable:

A 68-conductor cable that provides the secondary 16-bit DATA BUS. This cable is used in conjunction with the P cable to provide a 32-bit data path.

3.1.52

queue:

The arrangement of tasks within a task set usually according to the temporal order in which they were created.

3.1.53

queue tag:

The parameter associated with a task that uniquely identifies it from other tagged tasks for a logical unit from the same initiator.

3.1.54

receiver:

The circuitry used to detect the state of the bus.

3.1.55

reconnect:

The act of resuming a nexus to continue a task. A target reconnects when conditions are appropriate for the physical bus to transfer data associated with a nexus between an initiator and a target.

3.1.56

reconnection:

A reconnection is the result of a reconnect and it exists from the receipt of a selection confirmation with the selection won flag set to one or a reselection confirmation with the reselection won flag set to one until the next bus free indication occurs.

3.1.57

request:

A transaction invoking a service.

3.1.58

REQx:

A bus signal that is either the REQ or REQQ signal.

3.1.59

reselection ID:

The bit-significant representation of the target SCSI address in combination with the initiator SCSI address that is the result of a successful reselection service request.

3.1.60

response:

The third step of a four set confirmed service in reply to an indication.

3.1.61

SCSI address:

The decimal representation of the unique address assigned to an SCSI device.

3.1.62

SCSI bus:

Bus consisting of all the conductors and connectors required to attain signal line continuity between every driver, receiver, and terminator for each signal.

3.1.63

SCSI device:

A device containing at least one SCSI port and the means to connect the drivers and receivers to the bus.

3.1.64

SCSI ID:

The bit-significant representation of the SCSI address.

3.1.65

secondary bus:

The bus that contains 16 bits of the data and is used in conjunction with a primary bus to create a 32-bit data path.

3.1.66

signal assertion:

The act of driving a signal to the true state.

3.1.67

signal negation:

The act of performing a signal release or of driving a signal to the false state.

3.1.68

signal release:

The act of allowing the cable terminators to bias the signal to the false state (by placing the driver in the high impedance condition).

3.1.69

single-ended (SE):

A signalling alternative that employs SE (see 7.2) drivers and receivers to increase circuit density (also see 3.1.17, differential).

3.1.70

source (a signal):

The act of either signal assertion, signal negation, or signal release.

3.1.71**stub:**

Any electrical path connected to the bus that is not part of the bus path.

3.1.72**stubbed path:**

Path with stubs attached.

3.1.73**target:**

An SCSI device that receives SCSI commands and directs such commands to one or more logical units.

3.1.74**task:**

An object within the logical unit representing the work associated with a command or group of linked commands. A task consists of one initial connection and zero or more reconnections, all pertaining to the task. A task causes the nexus to be generated.

3.1.75**task manager:**

A server within the target device which executes task management functions.

3.1.76**task management function:**

A task manager service which can be invoked by an application client to affect the execution of one or more tasks.

3.1.77**task set:**

A group of tasks within a target device, whose interaction is dependent on the task management, contingent allegiance and auto-contingent allegiance rules. See the SCSI-3 Architecture Model standard for a detailed definition of a task set.

3.1.78**transceiver:**

A device that implements both the SCSI bus receiver and driver functions.

3.1.79**upper level protocol:**

Any protocol executed through services provided by a lower level protocol.

3.1.80**unconfirmed protocol service:**

A service available at the protocol service interface, that does not result in a completion confirmation.

3.1.81**vendor-specific:**

Something (i.g., a bit, field, code value) that is not defined by this standard and may be used differently in various implementations.

3.1.82**zero:**

A false signal value or a false condition of a variable.

3.2 Symbols and abbreviations

≠ or NE	not equal
≤ or LE	less than or equal to
±	plus or minus
≈	approximately
x	multiply
+	add
-	subtract
< or LT	less than
= or EQ	equal
> or GT	greater than
≥ or GE	greater than or equal to
ACA	auto-contingent allegiance (see 3.1.7)
CA	Contingent allegiance (see 3.1.13)
CMOS	Complementary metal oxide semiconductor
EMI	Electro-magnetic interference
EMC	Electro-magnetic compatibility
ESD	Electro-static discharge
HVD	High voltage differential
IDC	Insulation displacement contact
LSB	Least significant bit
LUN	Logical unit number
LVD	Low voltage differential
MLVD	Multimode low voltage differential
MSB	Most significant bit
MSE	Multimode single ended
RFI	Radio frequency interface
RIRC	Request indication response confirmation
SCSI	Either SCSI-2 or SCSI-3.
SCSI-2	Small Computer System Interface - 2
SCSI-3	Small Computer System Interface - 3
SDTR	Synchronous data transfer request
SE	Single-ended
WDTR	Wide data transfer request

3.3 Keywords**3.3.1****expected:**

A keyword used to describe the behavior of the hardware or software in the design models assumed by this standard. Other hardware and software design models may also be implemented.

3.3.2**invalid:**

A keyword used to describe an illegal or unsupported bit, byte, word, field or code value. Receipt of an invalid bit, byte, word, field or code value shall be reported as error.

3.3.3**mandatory:**

A keyword indicating an item that is required to be implemented as defined in this standard to claim compliance with this standard.

3.3.4

may:

A keyword that indicates flexibility of choice with no implied preference.

3.3.5

optional:

A keyword that describes features that are not required to be implemented by this standard. However, if any optional feature defined by this standard is implemented, it shall be implemented as defined in this standard.

3.3.6

reserved:

A keyword referring to bits, bytes, words, fields and code values that are set aside for future standardization. A reserved bit, byte, word or field shall be set to zero, or in accordance with a future extension to this standard. Recipients may not check reserved bits, bytes, words or fields for zero values. Receipt of reserved code values in defined fields shall be reported as error.

3.3.7

shall:

A keyword indicating a mandatory requirement. Designers are required to implement all such requirements to ensure interoperability with other products that conform to this standard.

3.3.8

should:

A keyword indicating flexibility of choice with a preferred alternative; equivalent to the phrase "it is recommended".

3.4 Conventions

Certain words and terms used in this International Standard have a specific meaning beyond the normal English meaning. These words and terms are defined either in clause 3 or in the text where they first appear. Names of signals, phases, messages, commands, statuses, sense keys, additional sense codes, and additional sense code qualifiers are in all uppercase (e.g., REQUEST SENSE), names of fields are in small uppercase (e.g., STATE OF SPARE), lower case is used for words having the normal English meaning.

Fields containing only one bit are usually referred to as the name bit instead of the name field.

Numbers that are not immediately followed by lower-case b or h are decimal values.

Numbers immediately followed by lower-case b (xxb) are binary values.

Numbers immediately followed by lower-case h (xxh) are hexadecimal values.

Decimals are indicated with a comma (e.g., two and one half is represented as 2,5).

Decimal numbers having a value exceeding 999 are represented with a space (e.g., 24 255).

In the event of conflicting information the precedence for requirements defined in this standard is:

- 1) text,
- 2) tables, and
- 3) figures.

3.5 Notation for Procedures and Functions

Procedure Name ([input-1a|input-1b|input-1c][,input-2a+input2b]...[input-n])
[output-1][,output-2]...[output -n])

Where:

Procedure Name:	A descriptive name for the function to be performed.
"(...)":	Parentheses enclosing the lists of input and output arguments.
input-1, input-2, ...:	A comma-separated list of names identifying caller-supplied input data objects.
output-1, output-2, ...:	A comma-separated list of names identifying output data objects to be returned by the procedure.
" ":	A separator providing the demarcation between inputs and outputs. Inputs are listed to the left of the separator; outputs, if any, are listed to the right.
"[...]":	Brackets enclosing optional or conditional parameters and arguments.
" ":	A separator providing the demarcation between a number of arguments of which only one shall be used in any single procedure.
"+":	A collection of objects presented to a single object. No ordering is implied.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14776-112:2002

4 General

4.1 Overview

The SCSI Parallel Interface-2 Standard defines the cables, connectors, signals, transceivers, and protocol used to interconnect SCSI devices and the services provided to the application client.

4.1.1 Data transfer modes

SCSI parallel interface devices default to 8-bit asynchronous transfer. The 8-bit asynchronous information transfer mode is always used for all information transfers except DATA IN phases and DATA OUT phases. DATA IN phases and DATA OUT phases may use asynchronous or synchronous transfers that may be 8-bits, 16-bits or 32-bits wide, if a synchronous transfer agreement or a wide transfer agreement is in effect.

4.1.2 Cables, Connectors, Signals, Transceivers

SCSI parallel interface devices may be implemented with either 50, 68, or 80 pin connectors.

Table 1 defines the bus modes and transfer rates supported with the various transceivers defined within this standard.

Table 1 - Transceiver/speed support map

Transceiver	Maximum transfer rate				
	Asynchronous	Fast-5	Fast-10	Fast-20	Fast-40
SE	yes	yes	yes	yes	no
MSE (Note)	yes	yes	yes	yes	no
LVD	yes	yes	yes	yes	yes
HVD	yes	yes	yes	yes	yes
Key: yes = Transceiver/speed combination supported by this standard. No = Transceiver/speed combination not supported by this standard.					
Note-MSE is identical to SE except for the requirements in 7.4 and table 22.					

SCSI devices may connect to the bus via 8-bit, 16-bit, or 32-bit ports. The 8-bit and 16-bit ports shall connect to a primary bus that consists of an A cable or P cable via a single connector. The 32-bit port shall connect to a primary bus that consists of a P cable and a secondary bus that consists of a Q cable via two connectors; one to a P cable and the other to a Q cable. (see 5)

4.1.3 Physical architecture of bus

The position of the drivers, receivers, and terminators for a SE bus are shown in figure 3 and for a differential bus are shown in figure 2. The electrical properties of the drivers and receivers are all measured at the stub connections. Unless otherwise noted, all voltages are with respect to the signal ground of the SCSI device.




Figure 2 - Differential SCSI bus




Figure 3 - SE SCSI bus

4.1.4 Physical topology details and definitions

The SCSI bus is a multidrop architecture described in 4.1.3. Other details important to ensure the proper operation of this topology are described in this subclause.

The SCSI bus consists of all the conductors and connectors required to attain signal line continuity between every driver, receiver, and terminator for each signal. The electrical connection directly between the two terminators forms the bus-path. Any electrical path that is not part of the bus-path is a stub. The

point where a stub meets the bus path is termed the stub connection.

Figure 4 shows examples of connectors, bus paths, stubs, and stub connections.

SCSI bus connectors are any connector, defined within this standard, used to create the SCSI bus. SCSI bus connectors shall be defined by their function and by their physical placement.

The functional definitions are:

- a) connectors used to provide part of the bus-path are bus-path connectors, and
- b) connectors used to provide part of a stub are stub connectors.

Common physical placements are:

- a) connectors physically part of SCSI devices are device connectors,
- b) connectors physically part of cables, backplanes, or other non-device conductors are cable connectors,
- c) connectors physically part of terminators are terminator connectors,
- d) connectors physically part of enclosures are enclosure connectors, and
- e) other physical placements may be used.

SCSI bus connectors (e.g., device stub connector, terminator bus path connector) referred to in this standard use both the functional definition and a physical placement.

The portion of the stub contained within the stub connector that has the stub connection may be ignored.




Figure 4 - SCSI bus topology details

If an intermediate interconnection is added to connect the device to the bus path this additional interconnect (including its connectors) and the device all contribute to the stub and bus loading. In system implementations that use an intermediate interconnect the parameters specified in this standard at the device connector shall apply at the stub connection.

NOTE 1 - Any extensions of the connection beyond the terminator as shown in the right side of figure 4 should be minimized or avoided as that extension produces stubs and bus loading.

NOTE 2 - In order to support daisy-chain connections, SCSI devices that use shielded connectors should provide two shielded device connectors on the device enclosure. Inside the enclosure the cable should be looped from one shielded connector to the other. The loop should pass the connecting point to the transceivers within the enclosure in such a manner that stub lengths are minimized. The length of the cable within the device enclosure is included when calculating the total cable length of the SCSI bus. (see figure 4)

4.1.5 Bus loading

Bus loading is the electrical current flowing through the stub connection for lines that are not being driven by the attached device. The bus termination circuitry also provides bus loading. Bus loading shall appear capacitive to A.C. signals and may also have a D.C. leakage component. The stub capacitance is caused by electrical paths and components within the stub. The leakage is caused by imperfect insulation of + and - signals and by components attached to the paths within the stub. The capacitive current loading is specified by the value of the capacitances at the + and - signals rather than by the value of the current.

Bus termination circuitry bus loading is the capacitance measured at the terminator bus-path connector. Any D.C. leakage within enabled terminators is part of the performance requirements in 7.2.1, 7.3.1, 7.4.1, and 7.5.1 and does not constitute bus loading.

Bus termination loading is separate from bus loading. Devices containing enabled bus termination shall present maximum loading at the stub connection that is the sum of the maximum allowed termination loading and the maximum allowed bus loading. See 4.1.6 for requirements of disabled termination circuitry.

For stub connections within an allowed stub length from enabled bus termination circuitry, the maximum bus loading allowed is the sum of the maximum bus termination loading and the maximum bus loading. If the enabled terminators are within a SCSI device and if either the bus termination loading or the bus loading is less than the maximum allowed, the other entity may increase its loading as long as the total for both entities does not exceed the maximum allowed.

4.1.6 Termination requirements

The SCSI bus termination defines the ends of the SCSI bus. Bus termination is required to set the negation state when no device is driving (also called biasing) and to match the impedance to that of the interconnect media. A termination circuit is providing bus termination only when it is delivering the performance requirements for biasing and impedance matching. Such a termination circuit is said to be enabled when it is providing the bus termination.

Terminator circuits may also be in a disabled state, when they are not providing any of the termination functions of bias and impedance matching. One way of disabling a terminator is to disconnect all the signal lines (optionally including DIFFSENS) by an electronic switch. Such a terminator circuit is called a switchable terminator.

Disabled terminators count as devices in terms of bus loading if they are individually attached to the bus. If they are contained within a device the disabled terminators become part of the device load budget for that device.

4.1.7 Device Addressing

The number of devices that may be addressed depends on the width of the data path on the bus; an 8-bit data path allows up to 8 devices to be addressed, a 16-bit data path allows up to 16 devices to be addressed, and a 32-bit data path allows up to 32 devices to be addressed. However, the number of devices that may be connected to the bus is dependent on several factors (e.g., bus length, data transfer rates, capacitance loading of the SCSI device) that are described throughout this standard.

4.1.8 Protocol

This standard describes a device's behavior in terms of functional levels, service interfaces between levels and peer-to-peer protocols. For a full description of the model used in this standard see the SCSI Architecture Model-2 standard. Figure 5 shows the model as it appears from the point of view of this standard.




Figure 5 - SCSI Parallel Interface service reference mode

Services between service levels are either four step confirmed services or two step confirmed services. A four step confirmed service consists of a service request, indication, response, and confirmation. A two step confirmed service consists of a service request and confirmation.

Figure 6 shows the service and protocol interactions for a four step confirmed service.




Figure 6 - Model for a four step confirmed service

The SCSI parallel interface service interface consists of the following interactions:

- A request to the initiator parallel interface agent to invoke a service;
- An indication from the target parallel interface agent notifying the device server or task manager of an event;
- A response from the device server or task manager in reply to an indication;
- A confirmation from the initiator parallel agent upon service completion.

Only application clients shall request a four step confirmed service be invoked.

Figure 7 shows the service and protocol interactions for a two step confirmed service.




Figure 7 - Model for a two step confirmed service

The SCSI parallel interface service interface consists of the following interactions:

- a) A request to the target parallel interface agent to invoke a service;
- b) A confirmation from the target parallel interface agent upon service completion.

Only device servers shall request a two step confirmed service be invoked.

5 SCSI parallel interface connectors

5.1 SCSI parallel interface connectors overview

Two types of connectors are defined: nonshielded and shielded. The nonshielded connectors are typically used within an enclosure. The shielded connectors are typically used for external applications where electromagnetic compatibility (EMC) and electrostatic discharge (ESD) protection may be required. Either type of connector may be used with the single-ended or differential transceivers.

This standard defines all the supported SCSI device connectors. The 80-contact alternative 4 non-shielded SCSI device connector and the 68-contact alternative 4 shielded SCSI device connector are defined by reference to IEC standards (see 2).

The alternative 1 nonshielded, alternative 3 nonshielded, alternative 1 shielded, and alternative 2 shielded connectors shall have contact geometry and normal force sufficient to pass the following test:

- a) Measure contact resistances of the connectors being evaluated using a test procedure for low-level contact resistance. Use IEC 60512-2 (low-level contact resistance test procedure for electronic connectors) as a reference procedure. Record measurements as initial contact resistances;
- b) Mate and unmate connectors 50 cycles;
- c) Contact resistance is measured in accordance with item a) above (this is an optional step);
- d) Expose mated connectors to mixed flowing gas consisting of 10 parts per billion (ppb) of chlorine, 10 ppb of hydrogen sulfide, 200 ppb of sulfur dioxide, and 200 ppb of nitrogen dioxide for 20 days at 70 % relative humidity and 30 °C. Use IEC 60512-11-7 (standard practice for conducting mixed flowing gas environmental tests) as a reference procedure;
- e) Remove connectors from the mixed flowing gas, remeasure contact resistance in accordance with item a) above. Any contact with an increase of 15 milliohms or greater is a failure.

The resistance shall be measured using a four-point dry-circuit method directly across the mated contact.

5.2 Nonshielded connector

5.2.1 Nonshielded connector alternative 1 - A cable

The alternative 1 nonshielded SCSI device connector for the A cable shall be a 50-conductor connector consisting of two rows of 25 female contacts with adjacent contacts 1,27 mm (0,05 in) apart, as shown in figure 8. The nonmating portion of the connector is shown for reference only.

The alternative 1 nonshielded mating connector for the A cable shall be a 50-conductor connector consisting of two rows of 25 male contacts with adjacent contacts 1,27 mm (0,05 in) apart, as shown in figure 9. The nonmating portion of the connector is shown for reference only.

5.2.2 Nonshielded connector alternative 2 - A cable

The alternative 2 nonshielded SCSI device connector for the A cable shall be a 50-conductor connector consisting of two rows of 25 male contacts with adjacent contacts 2,54 mm (0,1 in) apart as shown in figure 10. A shroud and header body should be used. The non-mating portion of the connector is shown for reference only.

The alternative 2 nonshielded mating connector for the A cable shall be a 50-conductor connector consisting of two rows of 25 female contacts with adjacent contacts 2,54 mm (0,1 in) apart as shown in figure 11. It is recommended that keyed connectors be used.

5.2.3 Nonshielded connector alternative 3 - P cable and Q cable

The alternative 3 nonshielded SCSI device connector for the P cable and Q cable shall be a 68-conductor connector consisting of two rows of 34 female contacts with adjacent contacts 1,27 mm (0,05 in) apart, as shown in figure 8. The nonmating portion of the connector is shown for reference only.


The alternative 3 nonshielded mating connector for the P cable and Q cable shall be a 68-conductor connector consisting of two rows of 34 male contacts with adjacent contacts 1,27 mm (0,05 in) apart, as shown in figure 9. The nonmating portion of the connector is shown for reference only.

5.2.4 Nonshielded connector alternative 4

The alternative 4 nonshielded SCSI device connector shall be a 80-conductor connector consisting of two rows of ribbon contacts spaced 1,27 mm (0,05 in) apart, as shown in figure 12 and figure 13. For the detailed dimensional drawings of this connector see the SCA-2 EIA specification EIA-700A0AE.


The alternative 4 nonshielded mating connector shall be a 80-conductor connector consisting of two rows of ribbon contacts spaced 1,27 mm (0,05 in) apart, as shown in figure 12 and figure 13. For the detailed dimensional drawings of this connector see the SCA-2 EIA specification EIA-700A0AE and annex L.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14776-112:2002




Dimensions	50 Position		68 Position	
	Millimeters	Inches	Millimeters	Inches
B1	34,70	1,366	46,13	1,816
B2	5,54	0,218	5,54	0,218
B3	2,54	0,100	2,54	0,100
B4	1,27	0,050	1,27	0,050
B5	30,48	1,200	41,91	1,650
B6	15°	15°	15°	15°
B7	1,00 R	0,039 R	1,00 R	0,039 R
B8	0,61±0,05	0,024±0,002	0,61±0,05	0,024±0,002
B9	0,15	0,006	0,15	0,006
B10	0,86±0,10	0,034±0,004	0,86±0,10	0,034±0,004
B11	0,15	0,006	0,15	0,006
B12	0,05	0,002	0,05	0,002
B13	5,00±0,13	0,197±0,005	5,00±0,13	0,197±0,005
B14	1,75 MAX	0,069 MAX	1,75 MAX	0,069 MAX
SOCKET X	25		34	
SOCKET Y	26		35	
SOCKET Z	50		68	

Figure 8 - 50/68-contact alternative 1/alternative 3 nonshielded SCSI device connector
(A cable/P cable/Q cable)




Dimensions	50 Position		68 Position	
	Millimeters	Inches	Millimeters	Inches
A1	34,85	1,372	46,28	1,822
A2	5,69	0,224	5,69	0,224
A3	2,54	0,100	2,54	0,100
A4	1,27	0,050	1,27	0,050
A5	30,48	1,200	41,91	1,650
A6	15°	15°	15°	15°
A7	1,04 R	0,041 R	1,04 R	0,041 R
A8	0,396±0,010	0,0156±0,0004	0,396±0,010	0,0156±0,0004
A9	0,23	0,009	0,23	0,009
A10	0,61±0,03	0,024±0,001	0,61±0,03	0,024±0,001
A11	0,23	0,009	0,23	0,009
A12	0,05	0,002	0,05	0,002
A13	5,16±0,15	0,203±0,006	5,16±0,15	0,203±0,006
A14	4,39 MAX	0,173 MAX	4,39 MAX	0,173 MAX
A15	3,02 MIN	0,119 MIN	3,02 MIN	0,119 MIN
A16	1,02±0,25	0,040±0,010	1,02±0,25	0,040±0,010
PIN X	25		34	
PIN Y	26		35	
PIN Z	50		68	

Figure 9 - 50/68-contact alternative 1/alternative 3 nonshielded mating connector
(A cable/P cable/Q cable)



Dimension	mm	in	Comments
<i>D1</i>	2,54	0,100	
<i>D2</i>	82,80	3,260	Reference Only
<i>D3</i>	2,54	0,100	
<i>D4</i>	8,89	0,350	Reference Only
<i>D5</i>	72,64	2,860	Reference Only
<i>D6</i>	78,74	3,100	Reference Only
<i>D7</i>	13,94	0,549	Reference Only
<i>D8</i>	4,19±0,25	0,165±0,010	
<i>D9</i>	6,10	0,240	
<i>D10</i>	6,60	0,260	Reference Only
NOTES			
1 Two rows of twenty five contacts on 2,540 mm (0,100 in) spacing = 60,960 mm (2,400 in).			
2 Tolerances ±0,127 mm (0,005 in) non-cumulative, unless specified otherwise.			

Figure 10 - 50-contact alternative 2 nonshielded SCSI device connector (A cable)



Dimensions	mm	in	Comments
C1	2,54	0,100	
C2	60,96	2,400	
C3	2,54	0,100	
C4	3,30	0,130	
C5	32,38	1,275	
C6	68,07	2,680	
C7	6,10	0,240	
C8	7,62	0,300	Maximum

NOTES

1 Fifty contacts on 1,270 mm (0,050 in) staggered spacing = 62,230 mm (2,450 in) [reference only].

2 Tolerances $\pm 0,127$ mm (0,005) non-cumulative, unless specified otherwise.

3 Connector cover and strain relief are optional.

Figure 11 - 50-contact alternative 2 nonshielded mating connector (A cable)




Figure 12 - 80-contact alternative 4 nonshielded SCSI device connector (P cable)




Figure 13 - 80-contact alternative 4 nonshielded contact positions

5.3 Shielded connector

Two shielded connector alternatives are specified for the A cable, the P cable, and the Q cable.

The D.C. resistance from the cable shield where it attaches to the connector to the enclosure should be less than 10 milliohms.

5.3.1 Shielded connector alternative 1 - A cable

The alternative 1 shielded SCSI device connector for the A cable shall be a 50-conductor connector consisting of two rows of 25 female contacts with adjacent contacts 1,27 mm (0,05 in) apart, as shown in figure 14. The nonmating portion of the connector is shown for reference only.

The alternative 1 shielded mating connector for the A cable shall be a 50-conductor connector consisting of two rows of 25 male contacts with adjacent contacts 1,27 mm (0,05 in) apart, as shown in figure 15. The nonmating portion of the connector is shown for reference only.

5.3.2 Shielded connector alternative 2 - A cable

The alternative 2 shielded SCSI device connector for the A cable shall be a 50-conductor connector consisting of two rows of ribbon contacts spaced 2,16 mm (0,085 in) apart, as shown in figure 16. The non-mating portion of the connector is shown for reference only.

The alternative 2 shielded mating connector for the A cable shall be a 50-conductor connector consisting of two rows of ribbon contacts spaced 2,16 mm (0,085 in) apart, as shown in figure 17. The non-mating portion of the connector is shown for reference only.

5.3.3 Shielded connector alternative 3 - P cable and Q cable

The alternative 3 shielded SCSI device connector for the P cable and Q cable shall be a 68-conductor connector consisting of two rows of 34 female contacts with adjacent contacts 1,27 mm (0,05 in) apart, as shown in figure 18. The nonmating portion of the connector is shown for reference only.

The alternative 3 shielded mating connector for the P cable and Q cable shall be a 68-conductor connector consisting of two rows of 34 male contacts with adjacent contacts 1,27 mm (0,05 in) apart, as shown in figure 19. The nonmating portion of the connector is shown for reference only.

Cable retention shall consist of #2-56 thread jack screws capable of withstanding a minimum torque of 1.2 Nm (11 inch-pounds).

5.3.4 Shielded connector alternative 4 - P cable and Q cable

The alternative 4 shielded SCSI device connector for the P cable and Q cable shall be a 68-conductor connector consisting of two rows of ribbon contacts spaced 0,8 mm (0,0315 in) apart, as shown in figure 20 and figure 21. For the detailed dimensional drawings of this connector see the VHDCI EIA specification EIA-700A0AF.

The alternative 4 shielded mating connector for the P cable and Q cable shall be a 68-conductor connector consisting of two rows of ribbon contacts spaced 0,8 mm (0,0315 in) apart, as shown in figure 20 and figure 21. For the detailed dimensional drawings of this connector see the VHDCI EIA specification EIA-700A0AF and annex M.




Figure 14 - 50-contact alternative 1 shielded SCSI device connector (A cable)




Figure 15 - 50-contact alternative 1 shielded mating connector (A cable)




Figure 16 - 50-contact alternative 2 shielded SCSI device connector (A cable)




Figure 17 - 50-contact alternative 2 shielded mating connector (A cable)





Figure 18 - 68-contact alternative 3 shielded SCSI device connector (P cable/Q cable)



Dimensions	Millimeters	Inches
A1	46,28	1,822
A2	5,69	0,224
A3	2,54	0,100
A4	1,27	0,050
A5	41,91	1,650
A6	15°	15°
A7	1,04 R	0,041 R
A8	0,396±0,010	0,0156±0,0004
A9	0,23	0,009
A10	0,60±0,03	0,024±0,001
A11	0,23	0,009
A12	0,05	0,002
A13	4,90±0,10	0,193±0,004
A14	4,27 max.	0,168 max.
A15	57,91±0,13	2,280±0,005
A16	0,25±0,13	0,010±0,005
A17	3,43±0,15	0,135±0,006
A18	2,64 min.	0,104 min.
A19	1,02±0,25	0,040±0,010

Figure 19 - 68-contact alternative 3 shielded mating connector (P cable/Q cable)




Figure 20 - 68-contact alternative 4 shielded SCSI device connector (P cable/Q cable)




Figure 21 - 68-contact alternative 4 shielded contact positions (P cable/Q cable)

5.4 Connector contact assignments

The connector contact assignments are defined in tables 2 through 13. The signals are defined in subclause 8.2. The items under signal name labelled TERMPWR, TERMPWRQ, and RESERVED are not signals and are not required to meet the cable characteristics for signals in subclause 6.3. See 6.4 for characteristics of TERMPWR and TERMPWRQ. See 6.5 for characteristics of RESERVED lines.

Table 2 - Cross-reference to A cable contact assignments

Connector type	Transmission mode	Connector figure	Contact assignment table	Contact set
Nonshielded alternative 1	SE	8 and 9	3	2
Nonshielded alternative 1	LVD	8 and 9	7	2
Nonshielded alternative 1	HVD	8 and 9	10	2
Nonshielded alternative 2	SE	10 and 11	3	1
Nonshielded alternative 2	LVD	10 and 11	7	1
Nonshielded alternative 2	HVD	10 and 11	10	1
Shielded alternative 1	SE	14 and 15	3	2
Shielded alternative 1	LVD	14 and 15	7	2
Shielded alternative 1	HVD	14 and 15	10	2
Shielded alternative 2	SE	16 and 17	3	1
Shielded alternative 2	LVD	16 and 17	7	1
Shielded alternative 2	HVD	16 and 17	10	1

5.4.1 SE assignments

Table 3 defines the connector contact assignments for a 50 conductor primary bus that uses SE transceivers.

Table 3 - SE contact assignments - A cable

Signal name	Connector contact number		Cable conductor number		Connector contact number		Signal name
	Set 2	Set 1			Set 1	Set 2	
SIGNAL RETURN	1	1	1	2	2	26	-DB(0)
SIGNAL RETURN	2	3	3	4	4	27	-DB(1)
SIGNAL RETURN	3	5	5	6	6	28	-DB(2)
SIGNAL RETURN	4	7	7	8	8	29	-DB(3)
SIGNAL RETURN	5	9	9	10	10	30	-DB(4)
SIGNAL RETURN	6	11	11	12	12	31	-DB(5)
SIGNAL RETURN	7	13	13	14	14	32	-DB(6)
SIGNAL RETURN	8	15	15	16	16	33	-DB(7)
SIGNAL RETURN	9	17	17	18	18	34	-DB(P)
GROUND	10	19	19	20	20	35	GROUND
GROUND	11	21	21	22	22	36	GROUND
RESERVED	12	23	23	24	24	37	RESERVED
OPEN (1)	13	25	25	26	26	38	TERMPWR
RESERVED	14	27	27	28	28	39	RESERVED
GROUND	15	29	29	30	30	40	GROUND
SIGNAL RETURN	16	31	31	32	32	41	-ATN
GROUND	17	33	33	34	34	42	GROUND
SIGNAL RETURN	18	35	35	36	36	43	-BSY
SIGNAL RETURN	19	37	37	38	38	44	-ACK
SIGNAL RETURN	20	39	39	40	40	45	-RST
SIGNAL RETURN	21	41	41	42	42	46	-MSG
SIGNAL RETURN	22	43	43	44	44	47	-SEL
SIGNAL RETURN	23	45	45	46	46	48	-C/D
SIGNAL RETURN	24	47	47	48	48	49	-REQ
SIGNAL RETURN	25	49	49	50	50	50	-I/O
Notes 1 Open lines shall be open in all SCSI devices. 2 The minus sign next to a signal indicates active low. 3 The conductor number refers to the conductor position when using flat-ribbon cable. 4 Two sets of contact assignments are shown. Refer to table 2 to determine which set of contacts applies to each connector.							

Table 4 defines the connector contact assignments for a 68 conductor primary bus that uses SE transceivers.

Table 4 - SE contact assignments - P cable

Signal name	Connector contact number	Cable conductor number		Connector contact number	Signal name
SIGNAL RETURN	1	1	2	35	-DB(12)
SIGNAL RETURN	2	3	4	36	-DB(13)
SIGNAL RETURN	3	5	6	37	-DB(14)
SIGNAL RETURN	4	7	8	38	-DB(15)
SIGNAL RETURN	5	9	10	39	-DB(P1)
SIGNAL RETURN	6	11	12	40	-DB(0)
SIGNAL RETURN	7	13	14	41	-DB(1)
SIGNAL RETURN	8	15	16	42	-DB(2)
SIGNAL RETURN	9	17	18	43	-DB(3)
SIGNAL RETURN	10	19	20	44	-DB(4)
SIGNAL RETURN	11	21	22	45	-DB(5)
SIGNAL RETURN	12	23	24	46	-DB(6)
SIGNAL RETURN	13	25	26	47	-DB(7)
SIGNAL RETURN	14	27	28	48	-DB(P)
GROUND	15	29	30	49	GROUND
GROUND	16	31	32	50	GROUND
TERMPWR	17	33	34	51	TERMPWR
TERMPWR	18	35	36	52	TERMPWR
RESERVED	19	37	38	53	RESERVED
GROUND	20	39	40	54	GROUND
SIGNAL RETURN	21	41	42	55	-ATN
GROUND	22	43	44	56	GROUND
SIGNAL RETURN	23	45	46	57	-BSY
SIGNAL RETURN	24	47	48	58	-ACK
SIGNAL RETURN	25	49	50	59	-RST
SIGNAL RETURN	26	51	52	60	-MSG
SIGNAL RETURN	27	53	54	61	-SEL
SIGNAL RETURN	28	55	56	62	-C/D
SIGNAL RETURN	29	57	58	63	-REQ
SIGNAL RETURN	30	59	60	64	-I/O
SIGNAL RETURN	31	61	62	65	-DB(8)
SIGNAL RETURN	32	63	64	66	-DB(9)
SIGNAL RETURN	33	65	66	67	-DB(10)
SIGNAL RETURN	34	67	68	68	-DB(11)
Notes 1 The minus sign next to a signal indicates active low. 2 The conductor number refers to the conductor position when using flat-ribbon cable.					

Table 5 defines the connector contact assignments for a 68 conductor secondary bus that uses SE transceivers

Table 5 - SE contact assignments - Q cable

Signal name	Connector contact number	Cable conductor number		Connector contact number	Signal name
SIGNAL RETURN	1	1	2	35	-DB(28)
SIGNAL RETURN	2	3	4	36	-DB(29)
SIGNAL RETURN	3	5	6	37	-DB(30)
SIGNAL RETURN	4	7	8	38	-DB(31)
SIGNAL RETURN	5	9	10	39	-DB(P3)
SIGNAL RETURN	6	11	12	40	-DB(16)
SIGNAL RETURN	7	13	14	41	-DB(17)
SIGNAL RETURN	8	15	16	42	-DB(18)
SIGNAL RETURN	9	17	18	43	-DB(19)
SIGNAL RETURN	10	19	20	44	-DB(20)
SIGNAL RETURN	11	21	22	45	-DB(21)
SIGNAL RETURN	12	23	24	46	-DB(22)
SIGNAL RETURN	13	25	26	47	-DB(23)
SIGNAL RETURN	14	27	28	48	-DB(P2)
GROUND	15	29	30	49	GROUND
GROUND	16	31	32	50	GROUND
TERMPWRQ	17	33	34	51	TERMPWRQ
TERMPWRQ	18	35	36	52	TERMPWRQ
RESERVED	19	37	38	53	RESERVED
GROUND	20	39	40	54	GROUND
GROUND	21	41	42	55	TERMINATED
GROUND	22	43	44	56	GROUND
GROUND	23	45	46	57	TERMINATED
SIGNAL RETURN	24	47	48	58	-ACKQ
GROUND	25	49	50	59	TERMINATED
GROUND	26	51	52	60	TERMINATED
GROUND	27	53	54	61	TERMINATED
GROUND	28	55	56	62	TERMINATED
SIGNAL RETURN	29	57	58	63	-REQQ
GROUND	30	59	60	64	TERMINATED
SIGNAL RETURN	31	61	62	65	-DB(24)
SIGNAL RETURN	32	63	64	66	-DB(25)
SIGNAL RETURN	33	65	66	67	-DB(26)
SIGNAL RETURN	34	67	68	68	-DB(27)
Notes 1 The minus sign next to a signal indicates active low. 2 The conductor number refers to the conductor position when using flat-ribbon cable.					

Table 6 defines the connector contact assignments for a 68 conductor primary bus that uses SE transceivers.

Table 6 - SE contact assignments - nonshielded alternative 4 connector

Signal name	Host Pin Length (note 3)	Connector contact number	Connector contact number	Host Pin Length (note 3)	Signal name
12V CHARGE (note 1)	Long	1	41	Long	12V GROUND (note 1)
12V (note 1)	Short	2	42	Long	12V GROUND (note 1)
12V (note 1)	Short	3	43	Long	12V GROUND (note 1)
12V (note 1)	Short	4	44	Short	MATED 1 (note 1)
3,3V (note 1)	Short	5	45	Long	3,3V CHARGE (note 1)
3,3V (note 1)	Short	6	46	Long	GROUND
-DB(11)	Short	7	47	Short	SIGNAL RETURN
-DB(10)	Short	8	48	Short	SIGNAL RETURN
-DB(9)	Short	9	49	Short	SIGNAL RETURN
-DB(8)	Short	10	50	Short	SIGNAL RETURN
-I/O	Short	11	51	Short	SIGNAL RETURN
-REQ	Short	12	52	Short	SIGNAL RETURN
-C/D	Short	13	53	Short	SIGNAL RETURN
-SEL	Short	14	54	Short	SIGNAL RETURN
-MSG	Short	15	55	Short	SIGNAL RETURN
-RST	Short	16	56	Short	SIGNAL RETURN
-ACK	Short	17	57	Short	SIGNAL RETURN
-BSY	Short	18	58	Short	SIGNAL RETURN
-ATN	Short	19	59	Short	SIGNAL RETURN
-DB(P)	Short	20	60	Short	SIGNAL RETURN
-DB(7)	Short	21	61	Short	SIGNAL RETURN
-DB(6)	Short	22	62	Short	SIGNAL RETURN
-DB(5)	Short	23	63	Short	SIGNAL RETURN
-DB(4)	Short	24	64	Short	SIGNAL RETURN
-DB(3)	Short	25	65	Short	SIGNAL RETURN
-DB(2)	Short	26	66	Short	SIGNAL RETURN
-DB(1)	Short	27	67	Short	SIGNAL RETURN
-DB(0)	Short	28	68	Short	SIGNAL RETURN
-DB(P1)	Short	29	69	Short	SIGNAL RETURN
-DB((15)	Short	30	70	Short	SIGNAL RETURN
-DB(14)	Short	31	71	Short	SIGNAL RETURN
-DB(13)	Short	32	72	Short	SIGNAL RETURN
-DB(12)	Short	33	73	Short	SIGNAL RETURN
5V (note 1)	Short	34	74	Short	MATED 2 (note 1)
5V (note 1)	Short	35	75	Long	5V GROUND (note 1)
5V CHARGE (note 1)	Long	36	76	Long	5V GROUND (note 1)
SPINDLE SYNC (note 1)	Long	37	77	Long	ACTIVE LED OUT (note 1)
RMT_START (note 1)	Long	38	78	Long	DLYD_START (note 1)
SCSI ID (0) (note 1)	Long	39	79	Long	SCSI ID (1) (note 1)
SCSI ID (2) (note 1)	Long	40	80	Long	SCSI ID (3) (note 1)
Notes 1 See annex D for the definition of these signals. 2 The minus sign next to a signal indicates active low. 3 The pins identified as being short and long only applies to the host connector and not the connector on the SCSI device. All pins on the SCSI device connector are the same length.					

5.4.2 Differential connector contact assignments

Table 7 defines the connector contact assignments for a 50 conductor primary bus that uses LVD transceivers.

Table 7 - LVD contact assignments - A cable

Signal name	Connector contact number		Cable conductor number		Connector contact number		Signal name
	Set 2	Set 1			Set 1	Set 2	
+DB(0)	1	1	1	2	2	26	-DB(0)
+DB(1)	2	3	3	4	4	27	-DB(1)
+DB(2)	3	5	5	6	6	28	-DB(2)
+DB(3)	4	7	7	8	8	29	-DB(3)
+DB(4)	5	9	9	10	10	30	-DB(4)
+DB(5)	6	11	11	12	12	31	-DB(5)
+DB(6)	7	13	13	14	14	32	-DB(6)
+DB(7)	8	15	15	16	16	33	-DB(7)
+DB(P)	9	17	17	18	18	34	-DB(P)
GROUND	10	19	19	20	20	35	GROUND
DIFFSENS	11	21	21	22	22	36	GROUND
RESERVED	12	23	23	24	24	37	RESERVED
TERMPWR	13	25	25	26	26	38	TERMPWR
RESERVED	14	27	27	28	28	39	RESERVED
GROUND	15	29	29	30	30	40	GROUND
+ATN	16	31	31	32	32	41	-ATN
GROUND	17	33	33	34	34	42	GROUND
+BSY	18	35	35	36	36	43	-BSY
+ACK	19	37	37	38	38	44	-ACK
+RST	20	39	39	40	40	45	-RST
+MSG	21	41	41	42	42	46	-MSG
+SEL	22	43	43	44	44	47	-SEL
+C/D	23	45	45	46	46	48	-C/D
+REQ	24	47	47	48	48	49	-REQ
+I/O	25	49	49	50	50	50	-I/O
Notes 1 The conductor number refers to the conductor position when using flat-ribbon cable. 2 Two sets of contact assignments are shown. Refer to table 2 to determine which set of contacts applies to each connector.							

Table 8 defines the connector contact assignments for a 68 conductor primary bus that uses LVD transceivers.

Table 8 - LVD contact assignments - P cable

Signal name	Connector contact number	Cable conductor number		Connector contact number	Signal name
+DB(12)	1	1	2	35	-DB(12)
+DB(13)	2	3	4	36	-DB(13)
+DB(14)	3	5	6	37	-DB(14)
+DB(15)	4	7	8	38	-DB(15)
+DB(P1)	5	9	10	39	-DB(P1)
+DB(0)	6	11	12	40	-DB(0)
+DB(1)	7	13	14	41	-DB(1)
+DB(2)	8	15	16	42	-DB(2)
+DB(3)	9	17	18	43	-DB(3)
+DB(4)	10	19	20	44	-DB(4)
+DB(5)	11	21	22	45	-DB(5)
+DB(6)	12	23	24	46	-DB(6)
+DB(7)	13	25	26	47	-DB(7)
+DB(P)	14	27	28	48	-DB(P)
GROUND	15	29	30	49	GROUND
DIFFSENS	16	31	32	50	GROUND
TERMPWR	17	33	34	51	TERMPWR
TERMPWR	18	35	36	52	TERMPWR
RESERVED	19	37	38	53	RESERVED
GROUND	20	39	40	54	GROUND
+ATN	21	41	42	55	-ATN
GROUND	22	43	44	56	GROUND
+BSY	23	45	46	57	-BSY
+ACK	24	47	48	58	-ACK
+RST	25	49	50	59	-RST
+MSG	26	51	52	60	-MSG
+SEL	27	53	54	61	-SEL
+C/D	28	55	56	62	-C/D
+REQ	29	57	58	63	-REQ
+I/O	30	59	60	64	-I/O
+DB(8)	31	61	62	65	-DB(8)
+DB(9)	32	63	64	66	-DB(9)
+DB(10)	33	65	66	67	-DB(10)
+DB(11)	34	67	68	68	-DB(11)
Notes 1 The conductor number refers to the conductor position when using flat-ribbon cable.					

Table 9 defines the connector contact assignments for a 68 conductor secondary bus that uses LVD transceivers.

Table 9 - LVD contact assignments - Q cable

Signal name	Connector contact number	Cable conductor number		Connector contact number	Signal name
+DB(28)	1	1	2	35	-DB(28)
+DB(29)	2	3	4	36	-DB(29)
+DB(30)	3	5	6	37	-DB(30)
+DB(31)	4	7	8	38	-DB(31)
+DB(P3)	5	9	10	39	-DB(P3)
+DB(16)	6	11	12	40	-DB(16)
+DB(17)	7	13	14	41	-DB(17)
+DB(18)	8	15	16	42	-DB(18)
+DB(19)	9	17	18	43	-DB(19)
+DB(20)	10	19	20	44	-DB(20)
+DB(21)	11	21	22	45	-DB(21)
+DB(22)	12	23	24	46	-DB(22)
+DB(23)	13	25	26	47	-DB(23)
+DB(P2)	14	27	28	48	-DB(P2)
GROUND	15	29	30	49	GROUND
DIFFSENS	16	31	32	50	GROUND
TERMPWRQ	17	33	34	51	TERMPWRQ
TERMPWRQ	18	35	36	52	TERMPWRQ
RESERVED	19	37	38	53	RESERVED
GROUND	20	39	40	54	GROUND
TERMINATED	21	41	42	55	TERMINATED
GROUND	22	43	44	56	GROUND
TERMINATED	23	45	46	57	TERMINATED
+ACKQ	24	47	48	58	-ACKQ
TERMINATED	25	49	50	59	TERMINATED
TERMINATED	26	51	52	60	TERMINATED
TERMINATED	27	53	54	61	TERMINATED
TERMINATED	28	55	56	62	TERMINATED
+REQQ	29	57	58	63	-REQQ
TERMINATED	30	59	60	64	TERMINATED
+DB(24)	31	61	62	65	-DB(24)
+DB(25)	32	63	64	66	-DB(25)
+DB(26)	33	65	66	67	-DB(26)
+DB(27)	34	67	68	68	-DB(27)
Notes					
1 The conductor number refers to the conductor position when using flat-ribbon cable.					

Table 10 defines the connector contact assignments for a 50 conductor primary bus that uses HVD transceivers.

Table 10 - HVD contact assignments - A cable

Signal name	Connector contact number		Cable conductor number		Connector contact number		Signal name
	Set 2	Set 1			Set 1	Set 2	
GROUND	1	1	1	2	2	26	GROUND
+DB(0)	2	3	3	4	4	27	-DB(0)
+DB(1)	3	5	5	6	6	28	-DB(1)
+DB(2)	4	7	7	8	8	29	-DB(2)
+DB(3)	5	9	9	10	10	30	-DB(3)
+DB(4)	6	11	11	12	12	31	-DB(4)
+DB(5)	7	13	13	14	14	32	-DB(5)
+DB(6)	8	15	15	16	16	33	-DB(6)
+DB(7)	9	17	17	18	18	34	-DB(7)
+DB(P)	10	19	19	20	20	35	-DB(P)
DIFFSENS	11	21	21	22	22	36	GROUND
RESERVED	12	23	23	24	24	37	RESERVED
TERMPWR	13	25	25	26	26	38	TERMPWR
RESERVED	14	27	27	28	28	39	RESERVED
+ATN	15	29	29	30	30	40	-ATN
GROUND	16	31	31	32	32	41	GROUND
+BSY	17	33	33	34	34	42	-BSY
+ACK	18	35	35	36	36	43	-ACK
+RST	19	37	37	38	38	44	-RST
+MSG	20	39	39	40	40	45	-MSG
+SEL	21	41	41	42	42	46	-SEL
+C/D	22	43	43	44	44	47	-C/D
+REQ	23	45	45	46	46	48	-REQ
+I/O	24	47	47	48	48	49	-I/O
GROUND	25	49	49	50	50	50	GROUND
Notes 1 The conductor number refers to the conductor position when using flat-ribbon cable. 2 Two sets of contact assignments are shown. Refer to table 2 to determine which set of contacts applies to each connector.							

Table 11 defines the connector contact assignments for a 68 conductor primary bus that uses HVD transceivers.

Table 11 - HVD contact assignments - P cable

Signal name	Connector contact number	Cable conductor number		Connector contact number	Signal name
+DB(12)	1	1	2	35	-DB(12)
+DB(13)	2	3	4	36	-DB(13)
+DB(14)	3	5	6	37	-DB(14)
+DB(15)	4	7	8	38	-DB(15)
+DB(P1)	5	9	10	39	-DB(P1)
GROUND	6	11	12	40	GROUND
+DB(0)	7	13	14	41	-DB(0)
+DB(1)	8	15	16	42	-DB(1)
+DB(2)	9	17	18	43	-DB(2)
+DB(3)	10	19	20	44	-DB(3)
+DB(4)	11	21	22	45	-DB(4)
+DB(5)	12	23	24	46	-DB(5)
+DB(6)	13	25	26	47	-DB(6)
+DB(7)	14	27	28	48	-DB(7)
+DB(P)	15	29	30	49	-DB(P)
DIFFSENS	16	31	32	50	GROUND
TERMPWR	17	33	34	51	TERMPWR
TERMPWR	18	35	36	52	TERMPWR
RESERVED	19	37	38	53	RESERVED
+ATN	20	39	40	54	-ATN
GROUND	21	41	42	55	GROUND
+BSY	22	43	44	56	-BSY
+ACK	23	45	46	57	-ACK
+RST	24	47	48	58	-RST
+MSG	25	49	50	59	-MSG
+SEL	26	51	52	60	-SEL
+C/D	27	53	54	61	-C/D
+REQ	28	55	56	62	-REQ
+I/O	29	57	58	63	-I/O
GROUND	30	59	60	64	GROUND
+DB(8)	31	61	62	65	-DB(8)
+DB(9)	32	63	64	66	-DB(9)
+DB(10)	33	65	66	67	-DB(10)
+DB(11)	34	67	68	68	-DB(11)
Notes 1 The conductor number refers to the conductor position when using flat-ribbon cable.					

Table 12 defines the connector contact assignments for a 68 conductor secondary bus that uses HVD transceivers.

Table 12 - HVD contact assignments - Q cable

Signal name	Connector contact number	Cable conductor number		Connector contact number	Signal name
+DB(28)	1	1	2	35	-DB(28)
+DB(29)	2	3	4	36	-DB(29)
+DB(30)	3	5	6	37	-DB(30)
+DB(31)	4	7	8	38	-DB(31)
+DB(P3)	5	9	10	39	-DB(P3)
GROUND	6	11	12	40	GROUND
+DB(16)	7	13	14	41	-DB(16)
+DB(17)	8	15	16	42	-DB(17)
+DB(18)	9	17	18	43	-DB(18)
+DB(19)	10	19	20	44	-DB(19)
+DB(20)	11	21	22	45	-DB(20)
+DB(21)	12	23	24	46	-DB(21)
+DB(22)	13	25	26	47	-DB(22)
+DB(23)	14	27	28	48	-DB(23)
+DB(P2)	15	29	30	49	-DB(P2)
DIFFSENS	16	31	32	50	GROUND
TERMPWRQ	17	33	34	51	TERMPWRQ
TERMPWRQ	18	35	36	52	TERMPWRQ
RESERVED	19	37	38	53	RESERVED
TERMINATED	20	39	40	54	TERMINATED
GROUND	21	41	42	55	GROUND
TERMINATED	22	43	44	56	TERMINATED
+ACKQ	23	45	46	57	-ACKQ
TERMINATED	24	47	48	58	TERMINATED
TERMINATED	25	49	50	59	TERMINATED
TERMINATED	26	51	52	60	TERMINATED
TERMINATED	27	53	54	61	TERMINATED
+REQQ	28	55	56	62	-REQQ
TERMINATED	29	57	58	63	TERMINATED
GROUND	30	59	60	64	GROUND
+DB(24)	31	61	62	65	-DB(24)
+DB(25)	32	63	64	66	-DB(25)
+DB(26)	33	65	66	67	-DB(26)
+DB(27)	34	67	68	68	-DB(27)
Notes 1 The conductor number refers to the conductor position when using flat-ribbon cable.					

Table 13 defines the connector contact assignments for a 68 conductor primary bus that uses HVD or LVD transceivers.

Table 13 - HVD and LVD contact assignments - nonshielded alternative 4 connector

Signal name	Host Pin Length (note 2)	Connector contact number	Connector contact number	Host Pin Length (note 2)	Signal name
12V CHARGE (note 1)	Long	1	41	Long	12V GROUND (note 1)
12V (note 1)	Short	2	42	Long	12V GROUND (note 1)
12V (note 1)	Short	3	43	Long	12V GROUND (note 1)
12V (note 1)	Short	4	44	Short	MATED 1 (note 1)
3,3V (note 1)	Short	5	45	Long	3,3V CHARGE (note 1)
3,3V (note 1)	Short	6	46	Long	DIFFSNS
-DB(11)	Short	7	47	Short	+DB(11)
-DB(10)	Short	8	48	Short	+DB(10)
-DB(9)	Short	9	49	Short	+DB(9)
-DB(8)	Short	10	50	Short	+DB(8)
-I/O	Short	11	51	Short	+I/O
-REQ	Short	12	52	Short	+REQ
-C/D	Short	13	53	Short	+C/D
-SEL	Short	14	54	Short	+SEL
-MSG	Short	15	55	Short	+MSG
-RST	Short	16	56	Short	+RST
-ACK	Short	17	57	Short	+ACK
-BSY	Short	18	58	Short	+BSY
-ATN	Short	19	59	Short	+ATN
-DB(P)	Short	20	60	Short	+DB(P)
-DB(7)	Short	21	61	Short	+DB(7)
-DB(6)	Short	22	62	Short	+DB(6)
-DB(5)	Short	23	63	Short	+DB(5)
-DB(4)	Short	24	64	Short	+DB(4)
-DB(3)	Short	25	65	Short	+DB(3)
-DB(2)	Short	26	66	Short	+DB(2)
-DB(1)	Short	27	67	Short	+DB(1)
-DB(0)	Short	28	68	Short	+DB(0)
-DB(P1)	Short	29	69	Short	+DB(P1)
-DB((15)	Short	30	70	Short	+DB((15)
-DB(14)	Short	31	71	Short	+DB(14)
-DB(13)	Short	32	72	Short	+DB(13)
-DB(12)	Short	33	73	Short	+DB(12)
5V (note 1)	Short	34	74	Short	MATED 2 (note 1)
5V (note 1)	Short	35	75	Long	5V GROUND (note 1)
5V CHARGE (note 1)	Long	36	76	Long	5V GROUND (note 1)
SPINDLE SYNC (note 1)	Long	37	77	Long	ACTIVE LED OUT (note 1)
RMT_START (note 1)	Long	38	78	Long	DLYD_START (note 1)
SCSI ID (0) (note 1)	Long	39	79	Long	SCSI ID (1) (note 1)
SCSI ID (2) (note 1)	Long	40	80	Long	SCSI ID (3) (note 1)
Notes 1 See annex D for the definition of these signals. 2 The pins identified as being short and long only applies to the host connector and not the connector on the SCSI device. All pins on the SCSI device connector are the same length.					

6 SCSI bus interconnect

6.1 SCSI bus interconnect overview

This clause defines the characteristics of interconnects used to connect SCSI parallel interface devices. These interconnects are part of the SCSI path. See figure 4 for the topology of the interconnects that make up a SCSI path. Examples of types of Interconnects are:

- a) unshielded flat-ribbon cable;
- b) unshielded flat twisted-pair ribbon cable;
- c) unshielded round twisted-pair cables;
- d) shielded round twisted-pair cables;
- e) backplanes.

The interconnect is defined as the electrical media (including connectors and passive loads) used to connect the TERMPWR, terminators, and SCSI devices in a SCSI bus.

The function of the interconnect is to:

- a) carry the signals,
- b) carry the terminator power from TERMPWR sources to the terminators, and
- c) to provide continuity between reserved pins and ground pins between devices and terminators.

Minimum driven signals shall result in received signals that meet the requirements contained in clause 7. Signals for this requirement include DB(0) through DB(31), DB(P), DB(P1), DB(P2), DB(P3), C/D, I/O, MSG, BSY, SEL, ATN, REQ, ACK, DIFFSENS, and RESET. At least minimum TERMPWR shall be delivered to the terminator from minimum sources per the requirements in clause 7.6.

The requirements on interconnect components in this clause are intended to produce interchangeable components while achieving the desired signal transmission properties.

6.2 SCSI bus cables

If twisted-pair cables are used, the twisted pairs in the cable shall be wired to physically opposing contacts in the connector.

The items under signal name labelled TERMPWR and TERMPWRQ are not signals and are not required to meet the cable characteristics for signals in 6.3. See 6.4 for characteristics of TERMPWR and TERMPWRQ.

See 6.5 for characteristics of RESERVED lines.

Interconnection of SCSI devices by means other than cables is allowed (e.g., by backplanes using printed wiring boards). Detailed descriptions of these other means are not part of this standard; however, they are subject to the electrical characteristics presented in this standard. Examples of these electrical characteristics are:

- a) characteristic impedance (see 6.3);
- b) propagation delay (see 6.3);
- c) cumulative length (see 6.6 and 6.7); and
- d) signal attenuation (see 6.3).

A primary bus carries an 8-bit or 16-bit data bus and the signals used to move information between SCSI devices. A primary bus provides an 8-bit or 16-bit data transfer path. A secondary bus carries an additional 16-bit data bus that, used in conjunction with a 16-bit primary bus, provides a 32-bit data transfer path.

The signals shall not be internally connected together within the connectors or cables. See 8.2 for signal definitions.

6.3 Interconnect characteristics for signals

The minimum conductor size for signals should be as specified in table 14.

Table 14 - Recommended minimum conductor size

	Recommended minimum conductor size				
	SE		LVD		HVD
	3 meter cable	6 meter cable (note)	12 meter cable	>12 meter cable	25 meter cable
All conductors except terminator power and reserved lines	0,050 92 mm ²	0,080 42 mm ²	0,032 4 mm ²	0,080 42 mm ²	0,050 92 mm ²
Single terminator power line	0,080 42 mm ²	0,080 42 mm ²	N/A	N/A	N/A
Multiple terminator power lines	0,050 92 mm ²	0,050 92 mm ²	0,032 4 mm ²	0,080 42 mm ²	0,050 92 mm ²
Note: 1 SE 6 meter cable applies to SCSI devices running at or below fast-5 data transfer rates. 2 For interconnects supporting LVD/SE multimode the larger conductor size should be used. 3 The conductor size recommendations assume voltage drops only due to the cables (i.e., does not include voltage drops due to connectors, ground effects, etc.)					

The SE characteristic impedance of the cable is defined in table 15 and the differential characteristic impedance of the cable is defined in table 16. Two measurement techniques may be used to determine the impedance. The SE measurement technique is applicable to cables used with SE transceivers. The differential measurement technique is applicable to cables used with differential transceivers. See annex F for measurement techniques

Table 15 - SE characteristic impedance of cable at maximum indicated data transfer rate

Description	SE characteristic impedance		
	Fast-5 (note 1) (ohms)	Fast-10 (ohms)	Fast-20 (ohms)
Maximum for REQ and ACK signals	96	96	96
Nominal for REQ and ACK signals	84	84	90
Minimum for REQ and ACK signals	72	72	84
Maximum, all other signals	96	96	100
Nominal, all other signals	84	84	90
Minimum, all other signals	72	72	80
Maximum difference between any two signals in the same cable	12	12	12
Note: 1 The SE characteristic impedance of fast-5 is equivalent to the SE characteristic impedance of fast-10 because of different measurement techniques used in previous standards. 2 Cables meeting the requirements for higher data transfer rates improve margins for lower data transfer rates. 3 Values specified by the transfer rate corresponding to the maximum rate supported by the interconnect shall apply even if a slower transfer rate is negotiated.			

Table 16 - Differential characteristic impedance of cable at indicated data transfer rate

Description	Differential (ohms)		
	Fast-10	Fast-20	Fast-40
Maximum for REQ and ACK signals	160	160	135
Nominal for REQ and ACK signals	122	122	N/A
Minimum for REQ and ACK signals	115	115	110
Maximum, all other signals	160	160	135
Nominal, all other signals	122	122	N/A
Minimum, all other signals	115	115	110
Maximum difference between any two signals in the same cable	20	20	N/A
Note: 1 Cables meeting the requirements for higher data transfer rates improve margins for lower data transfer rates. 2 Values specified by the transfer rate corresponding to the maximum rate supported by the interconnect shall apply even if a slower transfer rate is negotiated.			

The maximum propagation delay of any signal on SCSI cables shall be 5,4 ns/m or a maximum propagation delay of 135 ns for the entire bus.

The maximum sine wave signal attenuation shall be 0,095 dB maximum per meter at 5 MHz, measured differentially or a maximum sine wave signal attenuation of 1,41 dB at 5 MHz for the entire bus measured differentially.

6.4 Decoupling characteristics for TERMPWR and TERMPWRQ lines

The TERMPWR and TERMPWRQ lines should be decoupled at each terminator with at least 2,2 μF and not greater than 10 μF bypass capacitor.

See 7.6 and 10 for additional information.




Figure 22 - Terminator decoupling example

6.5 Connection requirements for RESERVED lines

The RESERVED lines shall be left open in the bus terminator assemblies and in the SCSI devices. The RESERVED lines shall have continuity from one end of the SCSI bus to the other end.

6.6 Cables used with SE transceivers

The maximum distance between terminators when using SE transceivers shall be as defined in table 17.

Table 17 - SE maximum distance between terminators

Number of attached SCSI devices	Maximum distance between terminators (meters) (note 1)		
	Fast-5	Fast-10	Fast-20 (note 2)
1 to 4 maximum capacitance (25 pF) SCSI devices	6	3	3
5 to 8 maximum capacitance (25 pF) SCSI devices	6	3	1.5
9 to 16 maximum capacitance (25 pF) SCSI devices	6	3	note 3
Note: 1 For environments where all elements of the bus (cables, device interfaces, environmental noise and other values) are controlled to be better than minimally required, it may be possible to extend the path length and device count (see note 7 in 7.2.4). 2 It is recommended that the devices be uniformly spaced between terminators with the end devices located as close as possible to the terminators. 3 Extending the device count beyond sixteen requires restrictions beyond the those specified in this standard. 4 The negotiated transfer rate shall have no effect on the SE maximum distance between terminators. 5 If the SCSI device(s) have less than 25 pF capacitance it may be possible to extend the path length and device count (see note 7 in 7.2.4).			

The stub length when using SE transceivers shall not exceed 0,1 meter. The stub length is measured from

the transceiver to the connection of the SCSI bus path (see figure 4). The spacing of devices on the SCSI bus path should be at least three times the stub length to avoid stub clustering (See annex G).

The following requirements based on the connector contact assignments in 5.4.1 ensure that all SCSI round cables may be used with either SE or differential transceivers:

- a) In the P cable and Q cable conductor pairs #47-48 (ACK, ACKQ) and #57-58 (REQ, REQQ) shall be in the cable core;
- b) In the P cable and Q cable, if there are more than three conductor pairs in the cable core, conductor pairs #47-48 (ACK, ACKQ) and #57-58 (REQ, REQQ) shall not be adjacent to each other;
- c) In the A cable conductor pairs #37-38 (ACK) and #47-48 (REQ) shall be in the cable core;
- d) In the A cable, if there are more than three conductor pairs in the cable core, conductor pairs #37-38 (ACK) and #47-48 (REQ) shall not be adjacent to each other;
- e) Cable conductor pairs used for the DATA BUS (DBnPx) shall be in the outer layer of the cable;
- f) Each cable conductor pair shall consist of the signal return and its associated signal.

See annex E for information on interconnecting busses of different widths.

6.6.1 SE ground offset

The magnitude of the ground offset voltage between logic grounds on any two device connectors shall be less than 50 mV at any frequency detectable by the receiver.

6.7 Cables used with differential transceivers

Balanced interconnect media (e.g., twisted-flat, discrete wire twisted pairs, matched printed circuit board races) should be used with differential transceivers.

The maximum distance between terminators when using differential transceivers shall be as defined in table 18.

Table 18 - Differential maximum distance between terminators

Interconnect	Maximum distance between terminators (meters) (note 1)			
	Fast-5	Fast-10	Fast-20	Fast-40
Point-to-point interconnect	25	25	25	25 (note 2)
Multidrop interconnect	25	25	25	12
Note: 1 For environments where all elements of the bus (cables, device interfaces, environmental noise and other values) are controlled to be better than minimally required, it may be possible to extend the path length and device count (see note 7 in 7.2.4). 2 SCSI devices shall be within a stub length of the termination (see figure 4). 3 Values specified by the maximum distance between terminators shall apply even if a slower transfer rate is negotiated.				

6.7.1 HVD stub length and spacing

The stub length and spacing when using HVD transceivers shall be as defined in table 19.

Table 19 - HVD stub length and spacing requirements

Maximum transfer rate	HVD Stub length and spacing requirements																	
Fast-5 Fast-10 Fast-20	The stub length shall not exceed 0,2 meter. The stub length is measured from the transceiver to the connection of the SCSI bus path. The spacing of devices on the SCSI bus path should be at least three times the stub length to avoid stub clustering (See annex G).																	
Fast-40	The stub length shall not exceed 0,1 meter. The difference in stub length shall be less than 1.27 cm for the REQ, ACK, DATA, and PARITY signals. Stub length differences on the + and - signals of the same differential line should be minimized. The stub length is measured from the transceiver to the connection of the SCSI bus path. The spacing of devices on the SCSI bus path shall be as indicated in table 20. Table 20 - Minimum stub connection spacing rules for HVD SCSI devices <table><tr><th rowspan="2">Transmission mode</th><th colspan="5">Minimum spacing between stub connectors (meters)</th></tr><tr><th>40 pF/m (note)</th><th>65 pF/m (note)</th><th>90 pF/m (note)</th><th>115 pF/m (note)</th><th>140 pF/m (note)</th></tr><tr><td>HVD</td><td>0,20</td><td>0,12</td><td>0,09</td><td>0,07</td><td>0,06</td></tr></table> Note: Media capacitance with no devices attached measured between a signal conductor and ground when all other conductors in the path are connected to ground.	Transmission mode	Minimum spacing between stub connectors (meters)					40 pF/m (note)	65 pF/m (note)	90 pF/m (note)	115 pF/m (note)	140 pF/m (note)	HVD	0,20	0,12	0,09	0,07	0,06
Transmission mode	Minimum spacing between stub connectors (meters)																	
	40 pF/m (note)	65 pF/m (note)	90 pF/m (note)	115 pF/m (note)	140 pF/m (note)													
HVD	0,20	0,12	0,09	0,07	0,06													

6.7.2 LVD stub length and spacing

The stub length when using LVD transceivers shall not exceed 0,1 meter. The difference in stub length shall be less than 1,27 cm for the REQ, ACK, DATA, and PARITY signals. Stub length differences on the + and - signals of the same differential line should be minimized. The stub length is measured from the stub connection (see 4.1.4) to the end of the stub. The spacing of devices on the SCSI bus path shall be as indicated in table 21.

Table 21 - Minimum stub connection spacing rules for LVD SCSI devices

Transmission mode	Minimum spacing between stub connections (meters)				
	40 pF/m (note)	65 pF/m (note)	90 pF/m (note)	115 pF/m (note)	140 pF/m (note)
LVD	0,36	0,22	0,16	0,13	0,10

Note: Media capacitance with no devices attached measured between a signal conductor and ground when all other conductors in the path are connected to ground.

7 SCSI parallel interface electrical characteristics

7.1 SCSI parallel interface electrical characteristics overview

The SCSI parallel interface may use the following transmitter implementations:

- a) SE passive negation;
- b) SE active negation;
- c) HVD;
- d) LVD.

LVD/MSE transmitters support SE active negation. LVD/MSE device may be damaged if exposed to SE passive negation voltages.

If a transceiver fully complies with the requirements of more than one of the above transmitter implementations then it may interoperate with those transceiver types.

The SE and HVD alternatives are electrically incompatible and cannot operate on the same bus.

The LVD and HVD alternatives are mutually exclusive and damage may occur to SCSI devices if a HVD SCSI device is plugged into a bus that contains LVD transceivers (see 7.5).

For each transmitter implementation one or more different receiver and capacitance specifications may apply.

For measurements in this clause, SCSI bus termination is assumed to be external to the SCSI device. See 6.5 for the terminating requirements for the RESERVED lines. SCSI devices may have provision for allowing optional internal termination provided the internal termination conforms with 7.2.1, 7.3.1, 7.4.1 or 7.5.1 when enabled and the SCSI device, including the disabled termination, conforms with 7.2.4, 7.3.4, or 7.5.3 when the internal termination is disabled.

In addition to the device electrical requirements defined in the remaining subclauses of this clause devices shall meet the requirements specified in table 22 and table 23 at the device connector.

Table 22 - Electrical input requirements at the device connector

Value	Minimum	Maximum	Notes
SE (passive negation) input voltage	-0,5 V D.C.	5,5 V D.C.	Absolute maximum at all operating conditions, including the DIFFSENS connection for SE (passive negation) SCSI devices meeting the fast-5 and fast-10 requirements in table 24.
SE (active negation) and MSE input voltage	-0,5 V D.C.	4,1 V D.C.	Absolute maximum at all operating conditions, including the DIFFSENS connection for SE (active negation), and LVD/MSE SCSI devices, for SCSI devices meeting the fast-20 requirements in table 24.
LVD input voltage (D.C. V + or - signal to local ground)	-0,5 V D.C.	4,1 V D.C.	Absolute maximum at all operating conditions all signals except DIFFSENS.
HVD input voltage (D.C. V + or - signal to local ground)			See the ISO/IEC 8482:1993-12 for the absolute maximum at all operating conditions.
HVD input voltage DIFFSENS	-7 V	17,25 V	Absolute maximum at all operating conditions.

Table 23 - Input current requirements at the device connector for lines not being driven by the device

Value	Maximum	Notes
MSE current magnitude	20 μ A D.C.	Measured from + or - signal V_{IN} 0 to 4,1 V to local ground each signal pin.
LVD current magnitude	20 μ A D.C.	Measured from + or - signal $V_{IN} < 2,5$ V to local ground each signal pin.

7.2 SE alternative

7.2.1 SE termination

All SCSI bus signals are common among all devices connected to the bus. All signal lines shall be terminated at both ends with a terminator that is compatible with the type of transceivers used in the SCSI devices. The termination points define the ends of the bus. These termination points may be internal to an SCSI device.

NOTE 3 - If the termination is provided and enabled within an SCSI device that device should not be removed from the SCSI bus while the bus is in use.

All SE conductors not defined as RESERVED, SIGNAL RETURN, or TERMPWR shall be terminated exactly once at each end of the bus. The termination of each signal shall meet the following requirements:

- the terminators shall be powered by the TERMPWR line and may receive additional power from other sources but shall not require such additional power for proper operation (see 7.6);
- each terminator shall source current to the signal line whenever its terminal voltage is below 2,5 V D.C. and this current shall not exceed 22,4 mA for any line voltage above 0,5 V D.C. and 25,4 mA for any line voltage between 0,5 V D.C. and 0,2 V D.C. even when all other signal lines are driven at 4.0 V D.C.;
- the voltage on all released signal lines shall be at least 2,5 V D.C.;
- these conditions shall be met with any conforming configuration of targets and initiators as long as at least one device is supplying TERMPWR;
- the terminator at each end of the SCSI bus (see 7.2.4) shall add a maximum of 25 pF capacitance to each signal.
- the terminator shall not source current to the signal line whenever its terminal voltage is above 3,24 V D.C. except terminators may source current when the voltage is above 3,24 V D.C. in applications where the bus is less than 0,3 m.

Terminators employing a 220 ohm resistor to 5 volts and a 330 ohm resistor to ground shall not be used.

7.2.2 SE output characteristics

SE signals shall use either passive-negation or, except OR-tied signals (see 8.5), active-negation drivers. Passive-negation drivers shall only be used in SCSI devices that support transfer rates no greater than fast-10. Passive-negation drivers have two states, asserted and high-impedance. Passive-negation drivers are usually implemented using an open-collector or an open-drain circuit. Active-negation drivers have three states: asserted, negated, and high-impedance. Each signal sourced by an SCSI device shall have the steady state D.C. output characteristics defined in table 24 measured at the SCSI device's connector.

Table 24 - SE steady state output voltage characteristics

Driver Type	Maximum transfer rate	SE steady state output voltage characteristics
Passive negation	Fast-5, Fast-10	a) V_{OL} (low-level output voltage) = 0,0 V D.C. to 0,5 V D.C. at $I_{OL} = 48$ mA (signal asserted); b) V_{OH} (high-level output voltage) = 2,5 V D.C. to 5,25 V D.C. (signal negated)
Active negation	Fast-5, Fast-10, Fast-20	a) V_{OL} (low-level output voltage) = 0,0 V D.C. to 0,5 V D.C. at $I_{OL} = 48$ mA (signal asserted); b) V_{OH} (high-level output voltage) = 2,5 V D.C. to 3,7 V D.C. (signal negated)
Note: SE steady state output voltage characteristics specified by transfer rate shall apply even if a slower transfer rate is negotiated.		

NOTE 4 - Passive-negation drivers do not source current to achieve the V_{OH} voltage level. They enter the high-impedance state and rely on the terminator to source the current.

NOTE 5 - It is recommended that drivers not source current above 3,7 V D.C.

The output characteristics (signal negated) for active-negation drivers shall be constrained to operate in the non-shaded areas of figure 23 for fast-20 devices and recommend for all others.




Figure 23 - Active negation current vs. voltage

NOTE 6 - Figure 23 shows the allowed domains for the D.C. output characteristics of an active-negation driver when negated. It is not intended to show A.C. output characteristics, which may be different due to other requirements such as slew rate specifications. To measure the actual device D.C. output characteristics, it is necessary to vary the device load, so the test circuit shown in figure 24 is not applicable to this measurement.

While active-negation drivers may be used on any non-OR-tied signal (see 8.5), their usage is particularly valuable on the ACK, REQ, ACKQ, and REQQ signals, because these signals are vulnerable to glitches on the true-to-false transition that could lead to false ACK/ACKQ or REQ/REQQ detection. Additional benefit may be achieved by using active-negation drivers on the DATA BUS and parity signals when operating in fast synchronous data transfer mode by reducing the skews between the first group of signals (ACK, REQ, ACKQ, and REQQ) and the DATA BUS and parity signals.

All SE drivers shall maintain the high-impedance state during power-on and power-off cycles.

SCSI devices should meet the following specifications for all signals when measured on the test circuit

shown in figure 24 with a load capacitor (C_L) of $15 \text{ pF} \pm 5\%$:

- a) t_{rise} (rise rate) = 520 mv per ns maximum (0,7 V D.C. to 2,3 V D.C.);
- b) t_{fall} (fall rate) = 520 mv per ns maximum (2,3 V D.C. to 0,7 V D.C.).

The slew rates specified above are requirements for a driver when using the SE test circuit in figure 24 they are not the observed rise rate or fall rate that would be observed on an actual SCSI bus.

All other A.C. output timing specifications shall be measured with the test circuit shown in figure 24 with a load capacitor (C_L) of $200 \text{ pF} \pm 5\%$. The driver output timing using the load in the test circuit represents the timings defined in figure 43 and figure 44.




Figure 24 - SE A.C. test circuit

7.2.3 SE input characteristics

SCSI devices with power-on shall meet the requirements in 9.3.1 and 9.3.2 and the electrical characteristics in table 25 on each signal (including both receivers and disabled drivers).

Table 25 - SE input voltage characteristics

Maximum transfer rate	SE input voltage characteristics
Fast-5	a) V_{IL} (low-level input voltage) = 0,0 V D.C. to 0,8 V D.C. (signal true); b) V_{IH} (high-level input voltage) = 2,0 V D.C. to 5,25 V D.C. (signal false) c) I_{IL} (low-level input current) = - 0,4 mA to 0,0 mA at $V_I = 0,5$ V D.C.; d) I_{IH} (high-level input current) = 0,0 mA to 0,1 mA at $V_I = 2,7$ V D.C.; e) Minimum input hysteresis = 0,2 V D.C.
Fast-10	a) V_{IL} (low-level input voltage) = 0,0 V D.C. to 0,8 V D.C. (signal true); b) V_{IH} (high-level input voltage) = 2,0 V D.C. to 5,25 V D.C. (signal false) c) I_{IL} (low-level input current) = ± 20 μ A at $V_I = 0,5$ V D.C.; d) I_{IH} (high-level input current) = ± 20 μ A at $V_I = 2,7$ V D.C.; e) Minimum input hysteresis = 0,3 V D.C.
Fast-20	a) V_{IL} (low-level input voltage) = 1,0 V D.C. maximum (signal true); b) V_{IH} (high-level input voltage) = 1,9 V D.C. minimum (signal false) c) I_{IL} (low-level input current) = ± 20 μ A at $V_I = 0,5$ V D.C.; d) I_{IH} (high-level input current) = ± 20 μ A at $V_I = 2,7$ V D.C.; e) Minimum input hysteresis = 0,3 V D.C.
Note: 1 SE input voltage characteristics specified by the maximum transfer rate shall apply even if a slower transfer rate is negotiated. 2 Due to the tighter voltage thresholds for fast-20, the power supply should have a maximum $\pm 5\%$ tolerance of the nominal voltage. 3 All values apply to both active negation and passive negation devices.	

The transient leakage current that may occur (e.g. with some ESD protection circuits) at the time of physical insertion of an SCSI device is an exponentially decaying current that does not exceed the following specifications:

- a) $I_{IH,HP}$ (hot-plug high-level input current peak value excluding the first 10ns) = +1,5 mA at $V_I = 2,7$ V D.C.;
- b) T_{HP} (transient current duration to 10 % of peak value) = 20 μ s maximum.

SCSI devices with power-off should meet the above I_{IL} and I_{IH} electrical characteristics on each signal, except at time of physical insertion, when $I_{IH,HP}$ and T_{HP} prevail.

The nominal switching threshold should be 1,4 V D.C. to achieve maximum noise immunity and to ensure proper operation with complex cable configurations.

SCSI devices should incorporate a glitch filter function on REQ/REQQ and ACK/ACKQ signals to reduce or eliminate the effect of glitch pulses.

If implemented, the glitch filter period shall not be so long as to mask out the subsequent valid transition edges of the incoming REQ/REQQ and ACK/ACKQ signals.

7.2.4 SE input and output characteristics

The SE signals shall have the characteristics defined in table 26 when measured at the SCSI device's connector.

Table 26 - SE input and output electrical characteristics

Maximum transfer rate	SE input and output electrical characteristics
Fast-5	a) Maximum signal capacitance = 25 pF, measured at the beginning of the stub (see figure 4).
Fast-10	a) I_L (Leakage current) = -20 μ A to + 20 μ A at V_I = 0,0 V D.C. to 5,25 V D.C. (high-impedance state); b) Maximum signal capacitance = 25 pF, measured at the beginning of the stub (see figure 4).
Fast-20	a) I_L (Leakage current) = -20 μ A to + 20 μ A at V_I = 0,0 V D.C. to 4,1 V D.C. (high-impedance state); b) Maximum signal capacitance = 25 pF, measured at the beginning of the stub (see figure 4).
Note: 1 SE input and output voltage characteristics specified by the maximum transfer rate shall apply even if a slower transfer rate is negotiated. 2 All values apply to both active negation and passive negation devices.	

NOTE 7 - It is practical to design devices with a lumped capacitance of less than 16 pF. Devices without a switchable terminator may reduce this node capacitance even further. A decrease in lumped capacitance of the node and a uniform increase of the impedance along the SCSI bus towards an optimum value improves the margin and may allow for a greater number of attached devices. Backplane designs give the implementor the possibility of increasing the margins and connecting a greater number of devices to the bus.

7.3 LVD alternative

7.3.1 LVD termination

The electrical characteristics of LVD bus termination shall be as specified in this subclause. Figure 25 shows the V_B and V_A measurement points, referenced to local ground, for the differential SPI-2 bus terminator.


**Figure 25 - Differential SPI-2 bus terminator**

Figure 25 through figure 30 define the measurement points for the LVD terminators. Electrical characteristics shall meet the requirements in table 27 and table 28.




Figure 26 - Test circuit for terminator differential impedance

The requirements on the LVD bus termination that relate to differential impedance are specified in figure 26 and table 27. Figure 27 and table 27 show the allowed ranges for I_{diff} and V_{diff} in figure 26. The requirements that relate to differential impedance are specified in figure 27 and table 27. Table 27 specifies the allowed ranges for I_{diff} and V_{diff} in figure 26. The terminator bias voltage V_{BIAS} (V_{BIAS} is the voltage measured when $I=0$ in figure 27) shall have the values measured between V_1 and V_2 as measured at V_{diff} in figure 26 with the range values defined in table 27 in the LVD impedance and V_{BIAS} tests column.





Figure 27 - Termination I-V characteristics for differential and common mode impedance tests



VOLTAGE IS DRIVEN, CURRENT IS MEASURED

Figure 28 - Test circuit for termination common mode impedance test

The requirements that relate to common mode impedance are specified in figure 27 and table 27. Table 27 specifies the allowed ranges for I_{CM} and V_{CM} in figure 28. The terminator bias voltage V_{BIAS} (V_{BIAS} is the voltage measured when $I = 0$ in figure 27) shall have the values measured between V_1 and V_2 as measured at V_{CM} in figure 28 with the range values defined in table 27 in the common mode impedance and V_{BIAS} tests column.

Table 27 - I-V requirements for LVD impedance, common mode impedance, and V_{BIAS} tests

Values (figure 27)	LVD impedance and V_{BIAS} tests (note) (figure 26)	Common mode impedance and V_{BIAS} tests (figure 28)
V_1 (mV)	100	1125
V_2 (mV)	125	1375
V_3 (V)	1,0	2,0
V_4 (V)	-1,0	0,5
I_{max} (mA)	9,00	N/A
I_{min} (mA)	-11.25	N/A
S_1 (ohms)	100	100
S_2 (ohms)	110	300
Measurement	D.C.	D.C.
Note: $V_A + V_B = 2,5 \pm 0,2$ V (figure 26)		

The requirements on termination that relate to electrical balance are specified in figure 29, figure 30, and table 28. The voltage V_{test1} in figure 29 is varied over frequencies of 0 to 40 MHz with amplitude varied over the range V_{MIN} to V_{MAX} specified in table 28 while the voltage named ΔV in figure 29 is measured. The maximum difference between values of ΔV (ΔV in figure 29) measured during this test shall be as specified in table 28.




Figure 29 - Termination balance test configuration




Figure 30 - Termination balance test data definition

Table 28 - Values for LVD termination balance test

Value	LVD
V_{MIN} (V_{test1} peak)	0,7
V_{MAX} (V_{test1} peak)	1,8
R (ohms)	$100 \pm 0,01\%$
δV	20 mV max
Note: ΔV - Input impedance for instrumentation > 10 Kohms V_{test1} swept through all values between V_{MIN} and V_{MAX}	

7.3.2 LVD driver characteristics

The LVD driver shall provide balanced asymmetrical sources that provide current from positive supply voltage to one signal line while sinking the same current to ground from the other signal line as shown in figure 31. Diagonally opposite sources operate together to produce a signal assertion or a signal negation. An assertion is produced when positive supply voltage current is sourced from source 4 to the +signal line and source 2 sinks the same current from the -signal line to ground. A negation is produced when positive supply voltage current is sourced from source 1 to the -signal line and source 3 sinks the same current from the +signal line to ground.




Figure 31 - LVD transceiver architecture

Balanced transmissions occur when the changes in +SIGNAL current and the changes in the -SIGNAL current precisely cancel each other. The balance is important to reduce EMI and common mode signals. Asymmetry occurs when the intensity of the source 2 and 4 assertion pair is different from the source 1 and 3 negation pair. To compensate for the negation biasing effect of the terminators, the 2 and 4 assertion pair is stronger than the 1 and 3 negation pair.

LVD drivers shall meet the specifications in annex A.

7.3.3 LVD receiver characteristics

LVD receivers shall be connected to the +signal and -signal as shown in figure 31. An example of an LVD receiver is shown in figure 32. LVD receivers shall meet the requirements in annex A.




Figure 32 - LVD receiver example

7.3.4 LVD capacitive loads

Capacitive loads on differential SCSI busses shall meet the requirements specified in this subclause.

As shown in figure 33 there are three components to differential SCSI bus capacitive loading:

- a) -signal to local ground (C1);
- b) +signal to local ground (C2); and
- c) -signal to +signal (C3).

The values C1, C2, and C3 represent measurements between the indicated points and do not represent discrete capacitors. Capacitance measurements shall be made with a nominal 1 MHz source with the same nominal D.C. level on the +signal and the -signal as specified in table 29. The driving source from the instrumentation shall apply an A.C. signal level less than 100 mV rms.




Figure 33 - Capacitive loads

Table 29 - Values for LVD capacitive loads

Capacitance measurement	Maximum (1)	Notes
C1 (pF)	20	@V=0,7 to 1,8 V D.C. -sig/gnd
C2 (pF)	20	@V=0,7 to 1,8 V D.C. +sig/gnd
C3 (pF)	10	@V=0,7 to 1,8 V D.C. both - and +sig/gnd V is the same for both sigs $\pm 100\text{mV}$
$ C1 - C2 $ (pF)	1,5	REQ, ACK, DATA and PARITY (same signal)
$ C1 - C2 $ (pF)	3	all other signals (same signal)
$ C1(i) - C1(\text{REQ}) $ (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1
$ C2(i) - C2(\text{REQ}) $ (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1
$ C1(i) - C1(\text{ACK}) $ (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1
$ C2(i) - C2(\text{ACK}) $ (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1
note: It is recommended that implementors design capacitive loads to be as small as practical.		

Devices containing the enabled bus termination shall have maximum values 1,5 times the maximums listed in table 29. Differential bus termination circuitry that is not part of a device shall have maximum values 0,5 times the maximums listed in table 29.

7.3.4.1 Management of LVD release glitches

Under some conditions, an LVD signal that transitions from actively negated to released may cause brief pulses to the true signal state. These pulses are called "release glitches" and may last up to a bus settle delay. Requirements are defined in this subclause to avoid adverse affects from release glitches.

SCSI devices shall incorporate the requirements specified in table 30 when using LVD drivers and optionally may incorporate the requirements when using other drivers. The usage of active negation increases cross talk noise margin and increases the true-to-false transition speed as compared to passive negation.

Table 30 - Glitch management requirements for devices using LVD drivers

Signals	Mode	Active negation	Transmitting device	Receiving device
BSY, SEL, RST	I,T	P		
ACK, ATN	I	R	The initiator shall wait for a BUS FREE phase (note) before releasing the ACK and ATN signals from the actively negated state.	Starting no later than a Bus Settle Delay after releasing the BSY signal, the target shall ignore the ACK and ATN signals until a subsequent connection.
REQ	T	R	The target shall wait 2,5 x (Bus Settle Delay) after releasing the BSY signal before releasing the REQ signal from the actively negated state.	The initiator shall begin to ignore the REQ signal within 1,5 x (Bus Settle Delay) of the transition of the BSY signal from true to false
C/D, I/O, MSG	T	R	After a selection or reselection phase, these signals shall not be released until the BSY signal is released.	No glitch management required.
DATA BUS (SELECTION and RSELECTION phases)	I,T	P	The transmitting device shall release all false data bits during these phases.	No glitch management required.
DATA BUS (During information transfers)	I,T	R	No glitch management required.	No glitch management required.
Key: I = initiator; P = prohibited; R = required; T = target				
Note: BUS FREE phase starts a Bus Settle Delay after the BSY and SEL signals are both false.				

7.3.5 SE/HVD transmission mode detection

Transmission mode detection by LVD SCSI devices of SE and HVD SCSI devices is accomplished through the use of the DIFFSENS line. Requirements for devices and terminators for DIFFSENS are not the same as for "signal" lines because DIFFSENS is driven and detected using its own SE transmission and detection scheme.

LVD termination shall drive the DIFFSENS line as specified in 7.3.5.1 and LVD SCSI devices shall sense the DIFFSENS signal as specified in 7.3.5.2.

Devices and terminators connected to the DIFFSENS line shall comply with the requirements in table 22 and table 23.

7.3.5.1 LVD DIFFSENS driver

The LVD DIFFSENS driver sets a voltage level on the DIFFSENS line that uniquely defines a LVD transmission mode. LVD terminators and multimode terminators (see 7.4.1) shall provide a LVD DIFFSENS driver according to the specifications in table 31.

Table 31 - LVD DIFFSENS driver specifications

Value	max.	nominal	min.	notes
V_O (volts) when $I_O = 0$ (shorted to ground) to 5 mA	1,4	1,3	1,2	
I_{OS} (mA)	15	5		With TERMPWR at operational levels and $V_O = 0$.
Input current D.C. (μA)	10			With terminator disabled.
Input sink current D.C. (μA) at $V_O = 2,75V$	200		20	Required to prevent the line from floating and to ensure the HVD DIFFSENS drivers dominate the LVD DIFFSENS drivers.
Note: - All requirements apply at the terminator bussing connection (see figure 4). - All measurements per figure 34.				




Figure 34 - LVD DIFFSENS driver signal definitions

7.3.5.2 LVD DIFFSENS receiver

LVD SCSI devices shall incorporate a LVD DIFFSENS receiver that detects the voltage level on the DIFFSENS line for purposes of informing the device of the transmission mode being used by the bus. The LVD DIFFSENS receiver shall be capable of detecting SE, LVD, and HVD SCSI devices. Table 32 and figure 35 define the receiver input levels for each of the three modes.




Figure 35 - DIFFSENS receiver function

Table 32 - DIFFSENS receiver operating requirements

V_{in} range	Sensed differential driver type
-0,35V to +0,5V	SE
> +0,5V to < 0,7V	indeterminate for detecting a driver type
0,7V to 1,9V	LVD
> 1,9V to < 2,4V	indeterminate for detecting a driver type
2,4V to 5,5V	HVD
Note: Input resistance (V_{in}/I_{in}) shall be 200 Kohms to 250 Kohms @ $V_{in} < 2,7V$ under all conditions of power supply (i.e., power on, power off, power transients) All voltages measured at the device connector with respect to local ground.	

The input resistance requirement is for purposes of providing ground reference if no DIFFSENS drivers are connected to the bus and to ensure that the DIFFSENS receivers do not load the DIFFSENS drivers excessively and to ensure that SE mode is detected.

Devices shall not allow the signal drivers to leave the high impedance state during initial power on until both of the following conditions are satisfied:

- a) the device is capable of logical operation for at least 100 ms, and

NOTE 8 - Note: The 100 ms delay allows time for the DIFFSENS pin to connect after the initial power connection (in the case of insertion of a device into an active system), or allows time for the power distribution system to settle.

- b) the DIFFSENS mode detected has remained stable for an additional 100 ms after a) is achieved.

A device shall not change its present signal driver or receiver mode based on the DIFFSENS voltage level unless a new mode is sensed continuously for at least 100 ms.

An example implementation of an LVD DIFFSENS receiver is shown in figure 36. The reference voltage tolerance is greater on the higher voltage reference. This allows a simple resistor divider between V_{CC} and ground for the references. The 100 ms requirement is implemented in logic in this example and is not

shown in the figure 36.




Figure 36 - LVD DIFFSENS receiver example

7.4 LVD/MSE multimode alternative

7.4.1 LVD/MSE multimode termination

Figure 37 shows the architecture of the multimode terminator.

Multimode terminators sense the DIFFSENS line while sourcing the DIFFSENS signal with the LVD levels (see table 31). The DIFFSENS line being grounded indicates that one or more SE SCSI devices or SE terminators are attached to the bus. A multimode terminator shall switch to the termination mode that is appropriate for the bus based on the value of the DIFFSENS input voltage. The appropriate mode is indicated in table 32.

When operating in the LVD mode the requirements in 7.3 and LVD requirements in table 22 shall apply. When operating in the MSE mode SE termination requirements in 7.2 and MSE requirements in table 22 shall apply.

Multimode terminators are required to provide a ground driver (similar to that described for multimode transceivers) for purposes of establishing a ground reference for the SE transmission lines. The ground driver shall turn on and remain on while the DIFFSENS line indicates SE operation. When turned on ground drivers shall appear resistive with the following performance requirements:

$$< 0,5\text{V} @ +5 \text{ mA}, > -0,5\text{V} @ -5 \text{ mA}.$$

The requirements for the multimode terminator ground driver are different from those for the multimode transceiver ground driver (see 7.4.4) because the devices provide the bulk of the grounding. Devices may be located far from the ends of the bus where ground references are more important.

NOTE 9 - There is at least one hard ground on the +SIGNAL line when operating in MSE mode (caused by the devices and/or terminators that are single-ended). This hard ground provides a return path for any low frequency currents in the +SIGNAL line.

NOTE 10 - The disabled ground driver capacitance should match the capacitance of the disabled assertion and negation drivers.

When operating in an HVD environment the voltage on a termination contact may reach as high as 15 V above local ground due to allowed common mode transients for HVD. Multimode termination is not recommended for HVD environments unless the common mode voltages in the environment are controlled to safe levels for SE and LVD SCSI devices (see table 22).

NOTE 11 - When using only the SCA-2 connector (see 5.2.4) the SE, LVD, and HVD connector contact numbers allow switching between all three modes. In this case the terminator may switch to HVD mode if so indicated by the DIFFSENS line.




Figure 37 - Multimode terminator architecture

7.4.2 LVD/MSE multimode transceiver characteristics

The architecture for the multimode transceiver is shown in figure 38.

The contact assignments in table 8, table 7, table 6, and table 13 provide compatible alignment of the pins between SE and LVD for all connector alternates. This alignment allows a single interface to supply MSE and LVD transceivers within the same device. It is not the intent to define a dynamically changing transmission mode but rather to prevent incompatible devices from attempting to interoperate.

When operating in an HVD environment the voltage on a transceiver contact may reach as high as 15 V above local ground due to high allowed common mode transients for HVD. LVD/MSE multimode transceivers are not recommended for exposure to HVD environments unless the common mode voltages

in the environment are controlled to safe levels for SE, LVD, and LVD/SE multimode SCSI devices (see table 22).

The LVD/MSE multimode transceiver shall meet the requirements in 7.3.5.2.

LVD/MSE multimode transceivers shall be set to the appropriate mode by sensing the output of the DIFFSENS receiver. If the DIFFSENS receiver indicates SE the LVD/MSE multimode transceiver shall follow the SE requirements in 7.2 and MSE requirements in table 22. If the DIFFSENS line indicates LVD mode the LVD/MSE multimode transceiver shall follow the requirements in 7.3. If HVD operation is indicated by the DIFFSENS receiver all signals (except DIFFSENS) shall be set to a high impedance state ($> 100K$ ohms to the local ground) (see 7.5.4).

NOTE 12 - Protocol chips may offer SE fast-40 signals to drive separate LVD or HVD transceivers. These fast-40 SE signals are not specified for direct connection to a SCSI bus.




Figure 38 - Multimode transceiver architecture

7.4.3 Power for LVD/MSE multimode transceivers

Power for the differential transceivers may be supplied from a different source than the power for the MSE transceivers. This enables the MSE voltage levels while allowing the differential transceivers to have lower voltages to reduce chip power or for compatibility with low voltage logic levels.

7.4.4 Ground drivers

When using the multimode driver architecture described in figure 38 a MSE driver is required for the ground side of the driver. The ground driver provides the connection to ground for the SE ground line associated with the SE -signal line. When the MSE alternate is implemented the signal ground connections

should be through the ground drivers to allow switching to LVD mode.

When turned on, ground drivers shall appear resistive with the following performance requirements: $<0,5V$ @ $+25$ mA, $> -0.5V$ @ -25 mA. Ground drivers shall remain on for the entire time the device is powered and is sensing a SE transmission mode from the DIFFSENS receiver. Ground drivers are not required to implement any slew rate controls.

7.5 HVD alternative

7.5.1 HVD termination

All SCSI bus signals are common among all devices connected to the bus.

All HVD signals consist of two lines denoted +SIGNAL and -SIGNAL. A signal is true when +SIGNAL is more positive than -SIGNAL, and a signal is false when -SIGNAL is more positive than +SIGNAL. All assigned HVD signals described in table 10, table 11, table 12, and table 13 except TERMPWR, DIFFSENS, RESERVED, and GROUND shall be terminated at each end of the SCSI bus with HVD termination that is compatible with terminators defined in this subclause. Termination points may be internal to an SCSI device.

NOTE 13 - If the termination is provided within an SCSI device, that device should not be removed from the SCSI bus while the bus is in use.

The electrical characteristics of HVD bus termination shall be as specified in this subclause. Figure 25 shows the placement of the termination circuitry between the signals and the local ground.

Electrical characteristics shall meet the requirements in figure 25 through figure 30, table 33, and table 34.

The requirements on the HVD bus termination that relate to differential impedance are specified in figure 26 and table 33. Figure 27 and table 33 show the allowed ranges for I_{diff} and V_{diff} in figure 26. The requirements that relate to differential impedance are specified in figure 27 and table 33. Table 33 specifies the allowed ranges for I_{diff} and V_{diff} in figure 26. The terminator bias voltage V_{BIAS} (V_{BIAS} is the voltage measured when $I = 0$ in figure 27) shall have the values measured between V_1 and V_2 as measured at V_{diff} in figure 26 with the range values defined in table 33 in the HVD impedance and V_{BIAS} tests column.

The requirements that relate to common mode impedance are specified in figure 27 and table 33. Table 33 specifies the allowed ranges for I_{CM} and V_{CM} in figure 28. The terminator bias voltage V_{BIAS} (V_{BIAS} is the voltage measured when $I = 0$ in figure 27) shall have the values measured between V_1 and V_2 as measured at V_{CM} in figure 28 with the range values defined in table 33 in the common mode impedance and V_{BIAS} tests column.

For an example of a terminator network that meets requirements in this clause see Annex J.

Table 33 - I-V requirements for HVD impedance and common mode impedance tests

Value (figure 27)	HVD impedance and V_{BIAS} tests (note) (figure 26)	Common mode impedance and V_{BIAS} tests (figure 28)
V1 (mV)	650	1920
V2 (mV)	1050	2860
V3 (V)	6,0	12,0
V4 (V)	-6,0	-7,0
I_{max} (mA)	46,50	
I_{min} (mA)	-61,30	
S1 (ohms)	115	145
S2 (ohms)	130	200
Measurement	D.C.	D.C.
Note: $V_A + V_B = 2,5 \pm 0,2$ V (figure 26)		

The requirements on termination that relate to electrical balance are specified in figure 29, figure 30, and table 34. The voltage V in figure 29 is varied over frequencies of 0 to 40 MHz with amplitude varied over the range V_{MIN} to V_{MAX} specified in table 34 while the voltage named ΔV in figure 29 is measured. The maximum difference between values of ΔV (ΔV in figure 29) measured during this test shall be as specified in table 34.

Table 34 - Values for HVD termination balance test

Value	HVD
V_{MIN} (V_{test1} peak)	-7,0
V_{MAX} (V_{test1} peak)	12,0
R (ohms)	$100 \pm 0,01\%$
δV	20 mV max
Note: ΔV - Input impedance for instrumentation > 10 Kohms V_{test1} swept through all values between V_{MIN} and V_{MAX}	

7.5.2 HVD output characteristics

Each HVD signal sourced by an SCSI device shall have the output characteristics defined in table 35 when measured at the SCSI device's connector.

All HVD drivers shall maintain the high-output impedance during power-on and power-off cycles.

Table 35 - HVD output voltage characteristics

Maximum transfer rate	HVD output voltage characteristics
Fast-5	a) V_{OL} (low-level output voltage) = 1,7 maximum at I_{OL} (low-level output current = 55 mA); b) V_{OH} (high-level output voltage) = 2,7 minimum at I_{OH} (high-level output current = -55 mA); c) V_{OD} (HVD output voltage) = 1,0 minimum; d) shall conform to ISO/IEC 8482:1993-12. V_{OL} and V_{OH} shall be as measured between the output terminal and the SCSI device's logic ground reference. The test circuit for testing these characteristics is shown in figure 39.
Fast-10	a) V_{OD} (HVD output voltage) = 1,0 V minimum; b) shall conform to ISO/IEC 8482:1993-12. The test circuit for testing these characteristics is shown in figure 39.
Fast-20	a) V_{OD} (HVD output voltage) = 1,0 V minimum; b) shall conform to ISO/IEC 8482:1993-12. The test circuit for testing these characteristics is shown in figure 39.
Fast-40	a) V_{OD} (HVD output voltage) = 1,0 V minimum; b) shall conform to ISO/IEC 8482:1993-12. HVD drivers for fast-40 need not comply with the ISO/IEC 8482:1993-12 requirement that the difference in differential output voltages magnitudes between opposite binary states be less than 0,2 V. Note: The additional requirements for driver asymmetry do not affect the operation of HVD drivers at fast-20 and lower data transfer rates. In the test circuit for testing these characteristics is shown in figure 39. The magnitude of the steady state differential output voltage of a fast-40 HVD driver shall comply with the following requirements: $1\text{ V} \leq V_A \text{ or } V_N \leq 6\text{ V and } 0,42 V_N + 0,69\text{ V} \leq V_A \leq 2,49 V_N - 1,59\text{ V where } V_A \text{ and } V_N \text{ are defined as } V_{OD} \text{ in figure 39.}$
Note: HVD output voltage characteristics specified by the maximum transfer rate shall apply even if a slower transfer rate is negotiated.	




Figure 39 - HVD test circuit

7.5.3 HVD input characteristics

Each HVD signal shall have the input characteristics defined in table 36 when measured at the SCSI device's connector (including both receivers and disabled drivers).

Table 36 - HVD input voltage characteristics

Max. transfer rate	HVD input voltage characteristics																														
Fast-5	SCSI devices shall meet the following electrical characteristics on each signal (including both receivers and passive drivers): a) I_I (input current on either input) = $\pm 2,0$ mA maximum; b) maximum input capacitance = 25 pF; c) input characteristics shall conform to ISO/IEC 8482:1993-12. The I_I requirement shall be met with the input voltage varying between -7 V D.C. and +12V D.C., with power on or off, and with the hysteresis equaling 35 mV, minimum.																														
Fast-10	a) shall conform to ISO/IEC 8482:1993-12; b) exhibit at least 35 mV of hysteresis; c) measure no more than 25 pF of capacitance.																														
Fast-20	a) shall conform to ISO/IEC 8482:1993-12; b) exhibit at least 35 mV of hysteresis; c) measure no more than 25 pF of capacitance.																														
Fast-40	a) shall conform to ISO/IEC 8482:1993-12; b) exhibit at least 35 mV of hysteresis c) measure the capacitance loading see 7.3.4 and figure 33 d) see table 37 for capacitance loading e) Devices containing the enabled bus termination shall have maximum values 1,5 times the maximums listed in table 37. Differential bus termination circuitry that is not part of a device shall have maximum values 0,5 times the maximums listed in table 37; Table 37 - Values for HVD capacitive loads <table><tr><th>Capacitance</th><th>Max.</th><th>Notes</th></tr><tr><td>C1 (pF)</td><td>30</td><td>@V= -7 to 12 V D.C. -sig/gnd</td></tr><tr><td>C2 (pF)</td><td>30</td><td>@V=-7 to 12 V D.C. +sig/gnd</td></tr><tr><td>C3 (pF)</td><td>15</td><td>@V=-7 to 12 V D.C. both - and +sig/gnd V is the same for both sigs ± 100mV</td></tr><tr><td> C1 - C2 (pF)</td><td>1,5</td><td>REQ, ACK, DATA and PARITY (same signal)</td></tr><tr><td> C1 - C2 (pF)</td><td>3</td><td>all other signals (same signal)</td></tr><tr><td> C1(i) - C1(REQ) (pF)</td><td>2</td><td>For DATA(i) i = 0-15 and PARITY (i) i = 0,1</td></tr><tr><td> C2(i) - C2(REQ) (pF)</td><td>2</td><td>For DATA(i) i = 0-15 and PARITY (i) i = 0,1</td></tr><tr><td> C1(i) - C1(ACK) (pF)</td><td>2</td><td>For DATA(i) i = 0-15 and PARITY (i) i = 0,1</td></tr><tr><td> C2(i) - C2(ACK) (pF)</td><td>2</td><td>For DATA(i) i = 0-15 and PARITY (i) i = 0,1</td></tr></table>	Capacitance	Max.	Notes	C1 (pF)	30	@V= -7 to 12 V D.C. -sig/gnd	C2 (pF)	30	@V=-7 to 12 V D.C. +sig/gnd	C3 (pF)	15	@V=-7 to 12 V D.C. both - and +sig/gnd V is the same for both sigs ± 100 mV	C1 - C2 (pF)	1,5	REQ, ACK, DATA and PARITY (same signal)	C1 - C2 (pF)	3	all other signals (same signal)	C1(i) - C1(REQ) (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1	C2(i) - C2(REQ) (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1	C1(i) - C1(ACK) (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1	C2(i) - C2(ACK) (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1
Capacitance	Max.	Notes																													
C1 (pF)	30	@V= -7 to 12 V D.C. -sig/gnd																													
C2 (pF)	30	@V=-7 to 12 V D.C. +sig/gnd																													
C3 (pF)	15	@V=-7 to 12 V D.C. both - and +sig/gnd V is the same for both sigs ± 100 mV																													
C1 - C2 (pF)	1,5	REQ, ACK, DATA and PARITY (same signal)																													
C1 - C2 (pF)	3	all other signals (same signal)																													
C1(i) - C1(REQ) (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1																													
C2(i) - C2(REQ) (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1																													
C1(i) - C1(ACK) (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1																													
C2(i) - C2(ACK) (pF)	2	For DATA(i) i = 0-15 and PARITY (i) i = 0,1																													
Note: 1 HVD input voltage characteristics specified by the maximum transfer rate shall apply even if a slower transfer rate is negotiated. 2 It is recommended that implementors design capacitive loads to be as small as practical.																															

7.5.4 SE driver protection

Transmission mode detection of SE SCSI devices by HVD SCSI devices is accomplished through the use of the DIFFSENS line. The DIFFSENS signal is a single-ended signal that is used as an active high enable

for the differential drivers. If a SE SCSI device or terminator is inadvertently connected, this line is grounded, disabling the differential drivers (see figure 40). Devices and terminators that only support SE shall connect the DIFFSENS line to local ground.

Requirements for devices, terminators, and interconnect media for DIFFSENS are not the same as for "signal" lines because DIFFSENS is driven and detected using its own single-ended transmission and detection scheme.

HVD SCSI devices shall implement a DIFFSENS driver and receiver that have equivalent performance to the circuit shown in figure 40. HVD terminators are not involved with DIFFSENS except that they shall not ground the line.




Figure 40 - SE driver protection circuit

Devices and terminators connected to the DIFFSENS line shall comply with the requirements in table 22.

NOTE 14 - LVD SCSI devices are not detected by HVD SCSI devices not complying to this standard. On all HVD SCSI devices except the ones that have a SCA-2 connector the signals do not match. As a result some drivers are shorted to ground. Therefore damage may occur to SCSI devices if a HVD SCSI device is plugged into a bus that contains LVD transceivers.

7.6 Terminator power

Provision shall be made to provide power from one or more sources to the TERMPWR lines of the SCSI bus. Terminator power shall be supplied through a low forward drop diode or similar semiconductor that prevents backflow of power if one of the sources of TERMPWR is powered-off.

Bus terminators shall be powered from at least one source of termination power (TERMPWR). The TERMPWR lines in the cable are available for distribution of termination power. Direct connection between the TERMPWR source and the individual terminators without using the TERMPWR line is also allowed.

If the TERMPWR source is connected to the cable TERMPWR line, the source shall be isolated in a manner that prevents sinking of current from the TERMPWR line into the TERMPWR source (for example if the TERMPWR source voltage falls below the voltage existing on the TERMPWR line, the TERMPWR source will sink current unless unidirectional isolation is present in the TERMPWR source).

Regulatory agencies may require limiting maximum (short circuit) current to the TERMPWR lines. These

requirements generally mandate the use of current limiting circuits and may restrict the number of sources provided for TERMPWR.

The terminator power characteristics, for each terminator, at the terminator shall be as defined in table 38.

Table 38 - Terminator power characteristics at the terminator

Terminator power characteristics	Terminator type						
	SE			LVD	SE/LVD (Multimode)	HVD	
	A cable	P cable/Q cable				A cable	P cable/ Q cable
		0,2 V dropout regulator					
$I_{\min}$ (A) @ $V_{\min}$	0,35	0,6	0,6	0,5	0,65	0,6	1,0
$V_{\min}$ (V) @ $I_{\min}$	4,0	2,7	4,0	3,0	3,0	4,0	4,0
$V_{\max}$ (V) @ all conditions	5,25	5,25	5,25	5,25	5,25	5,25	5,25
Note: The recommended TERMPWR current limiting is 2,0 amps.							

NOTE 15 - SCSI devices connected with a SE A cable (table 3) are not able to meet the source current requirements in table 38 unless the TERMPWR conductor size is 0,080 98 mm² minimum because the SE A cable contains only one TERMPWR line.

NOTE 16 - It is recommended that an SCSI device connected with the nonshielded alternative 2 connectors (see 5.2.2) that provide terminator power use keyed connectors to prevent accidental grounding or the incorrect connection of terminator power.

It is recommended that the terminator power lines be decoupled at each terminator with a bypass capacitor of at least 2,2 μ F, to improve signal quality, but not greater than 10 μ F. (see 6.4)

The TERMPWR lines may be used for distribution of power for purposes other than for SCSI bus termination as long as the voltage delivered to the SCSI bus terminators remains adequate to supply the requirements of the terminators under all conditions of SCSI bus operation and under all conditions of other loading.

8 SCSI bus signals

8.1 SCSI bus signals overview

Information transfer on the SCSI bus is allowed between only two SCSI devices at any given time. The maximum number of SCSI devices is determined by the width of the data path implemented. The SCSI devices may be any combination of initiators and targets, provided there is at least one of each.

Each SCSI device has an SCSI address and a corresponding SCSI ID bit assigned to it. When two SCSI devices communicate on the SCSI bus, one acts as an initiator and the other acts as a target. An initiator originates an I/O process and the target performs the I/O process.

NOTE 17 - An SCSI device is usually fixed as an initiator or target, but some SCSI devices may be able to assume either role.

Table 39 shows the relationship between SCSI Addresses, SCSI IDs, and arbitration priority. In table 39 a hyphen ("-") represents a logical zero bit.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14776-112:2002

Table 39 - Arbitration priorities by SCSI ID

SCSI address	D B 31	D B 24	D B 23	D B 16	D B 15	D B 8	D B 7	D B 0	Priority
7	-	-	-	-	-	-	1	-	1
6	-	-	-	-	-	-	-	1	2
5	-	-	-	-	-	-	-	1	3
4	-	-	-	-	-	-	-	1	4
3	-	-	-	-	-	-	-	1	5
2	-	-	-	-	-	-	-	1	6
1	-	-	-	-	-	-	-	1	7
0	-	-	-	-	-	-	-	1	8
15	-	-	-	-	1	-	-	-	9
14	-	-	-	-	-	1	-	-	10
13	-	-	-	-	-	-	1	-	11
12	-	-	-	-	-	-	-	1	12
11	-	-	-	-	-	-	-	1	13
10	-	-	-	-	-	-	-	1	14
9	-	-	-	-	-	-	-	1	15
8	-	-	-	-	-	-	-	1	16
23	-	-	1	-	-	-	-	-	17
22	-	-	-	1	-	-	-	-	18
21	-	-	-	-	1	-	-	-	19
20	-	-	-	-	-	1	-	-	20
19	-	-	-	-	-	-	1	-	21
18	-	-	-	-	-	-	-	1	22
17	-	-	-	-	-	-	-	1	23
16	-	-	-	-	-	-	-	1	24
31	1	-	-	-	-	-	-	-	25
30	-	1	-	-	-	-	-	-	26
29	-	-	1	-	-	-	-	-	27
28	-	-	-	1	-	-	-	-	28
27	-	-	-	-	1	-	-	-	29
26	-	-	-	-	-	1	-	-	30
25	-	-	-	-	-	-	1	-	31
24	-	-	-	-	-	-	-	1	32

8.2 Signal descriptions

BSY (BUSY). An "OR-tied" signal that indicates that the SCSI bus is in use.

SEL (SELECT). An "OR-tied" signal used by an initiator to select a target or by a target to reselect an initiator.

C/D (CONTROL/DATA). A signal sourced by a target that indicates whether control or data information is on the DATA BUS. True indicates CONTROL.

I/O (INPUT/OUTPUT). A signal sourced by a target that controls the direction of data movement on the DATA BUS with respect to an initiator. True indicates INPUT. This signal is also used to distinguish between SELECTION and RESELECTION phases.

MSG (MESSAGE). A signal sourced by a target to indicate the MESSAGE phase. True indicates MESSAGE.

REQ (REQUEST). A signal sourced by a target to indicate a request for an information transfer on the primary SCSI bus's data path.

REQQ (REQUEST). A signal sourced by a target to indicate a request for an information transfer on the secondary SCSI bus's data path.

ACK (ACKNOWLEDGE). A signal sourced by an initiator to respond with an acknowledgment of an information transfer on the primary SCSI bus's data path.

ACKQ (ACKNOWLEDGE). A signal sourced by an initiator to respond with an acknowledgment of an information transfer on the secondary SCSI bus's data path.

ATN (ATTENTION). A signal sourced by an initiator to indicate the ATTENTION condition.

RST (RESET). An "OR-tied" signal that indicates the RESET condition.

DB(7-0,P) (8-bit DATA BUS). Eight data-bit signals, plus a parity-bit signal that form the 8-bit DATA BUS. DB(P) shall contain odd parity for DB(7-0). Bit significance and priority during arbitration are shown in table 39.

DB(15-0,P,P1) (16-bit DATA BUS). Sixteen data-bit signals, plus two parity-bit signals that form the 16-bit DATA BUS. DB(P,P1) shall contain odd parity for DB(7-0) and DB(15-8), respectively. Bit significance and priority during arbitration are shown in table 39.

DB(31-0,P,P1,P2,P3) (32-bit DATA BUS). Thirty-two data-bit signals, plus four parity-bit signals that form the 32-bit DATA BUS. DB(P,P1,P2,P3) shall contain odd parity for DB(7-0), DB(15-8), DB(23-16), and DB(31-24), respectively. Bit significance and priority during arbitration are shown in table 39.

8.3 Parity checking rules

Valid parity is determined by rules in table 40.

Table 40 - Parity checking rules

Check for odd parity on:	If at least one bit is active on:
DB(7-0),P	DB(31-0,P,P1,P2,P3)
DB(15-8),P1	DB(31-8,P1,P2,P3)
DB(23-16),P2	DB(31-16,P2,P3)
DB(31-24),P3	DB(31-16,P2,P3)

NOTE 18 - These rules are necessary to permit interoperability of SCSI devices with different DATA BUS widths. For example, if a 16-bit SCSI device selects a 32-bit SCSI device, the 32-bit SCSI device will observe invalid parity on the upper 16 bits of the DATA BUS.

8.4 Signal states

8.4.1 SE

Signals may be in a true (asserted) or false (negated) state. Signals that are asserted are actively driven to the true state. Signals that are negated may either be actively driven to the false state or released to the false state. A signal that is released goes to the false state because the bias of the terminator pulls the signal false. OR-tied signals shall not be actively driven false.

NOTE 19 - The advantage of actively negating signals false during information transfer is that the noise margin is higher than if the signal is simply released. This facilitates reliable data transfer at high transfer rates.

Bits of the DATA BUS are defined as one when the signal is true, and defined as zero when the signal is false.

8.4.2 Differential

Figure 41 defines the voltage and current definitions. A signal that is released goes to the false state because the bias of the terminator pulls the signal false.




Figure 41 - Voltage and current definitions

Figure 42 defines the signaling sense of the voltages appearing on the - signal and + signal lines as follows:

- The - signal terminal of the driver shall be negative with respect to the + signal terminal for an asserted state.
- The - signal terminal of the driver shall be positive with respect to the + signal terminal for a negated state.




Figure 42 - Signaling sense

NOTE 20 - For a description of V_{BIAS} see 7.3.1.

8.5 OR-tied signals

The BSY, SEL, and RST signals shall be OR-tied.

BSY and RST signals may be simultaneously driven true by several SCSI devices. No signals other than BSY, SEL, RST, and DB(P,P1,P2,P3) are simultaneously driven by two or more SCSI devices. Parity bits shall not be driven false during the ARBITRATION phase but may be driven false in other phases.

8.6 Signal sources

Table 41 indicates the type of SCSI device allowed to source each signal. No attempt is made to show if the source is driving asserted, driving negated, or is released. All SCSI device drivers that are not active sources shall be in the high-impedance state. The RST signal may be asserted by any SCSI device at any time.

Table 41 - Signal sources

SCSI bus phase	P cable signals						Q cable signals		
	A cable signals					DB15-8 DB(P1)	REQQ	ACKQ	DB31-16 DB(P2) DB(P3)
	BSY	SEL	C/D I/O MSG REQ	ACK ATN	DB7-0 DB(P)				
BUS FREE	None	None	None	None	None	None	None	None	None
ARBITRATION	All	Win	None	None	S ID	S ID	None	None	S ID
SELECTION	I&T	Init	None	Init	Init	Init	None	None	Init
RESELECTION	I&T	Targ	Targ	Init	Targ	Targ	None	None	Targ
COMMAND	Targ	None	Targ	Init	Init	None	None	None	None
DATA IN	Targ	None	Targ	Init	Targ	Targ	Targ	Init	Targ
DATA OUT	Targ	None	Targ	Init	Init	Init	Targ	Init	Init
STATUS	Targ	None	Targ	Init	Targ	None	None	None	None
MESSAGE IN	Targ	None	Targ	Init	Targ	None	None	None	None
MESSAGE OUT	Targ	None	Targ	Init	Init	None	None	None	None
All: The signal shall be driven by all SCSI devices that are actively arbitrating. S ID: A unique data bit (the SCSI ID) shall be driven by each SCSI device that is actively arbitrating; the other data bits shall be released (i.e., not driven) by this SCSI device. The parity bit(s) may be released or driven to the true state, but shall not be driven to the false state during this phase. I&T: The signal shall be driven by the initiator, target, or both, as specified in the SELECTION phase and RESELECTION phase. Init: If driven, this signal shall be driven only by the active initiator. None: The signal shall be released; that is, not driven by any SCSI device. The bias circuitry of the bus terminators pulls the signal to the false state. Win: The signal shall be driven by the one SCSI device that wins arbitration. Targ: If the signal is driven, it shall be driven only by the active target.									

9 SCSI parallel bus timing

9.1 SCSI parallel bus timing values

Unless otherwise indicated, the delay-time measurements for each SCSI device, shown in table 42, shall be calculated from signal conditions existing at that SCSI device's port. Thus, these measurements (except cable skew) may be made without considering delays in the cable. The timing characteristics of each signal are described in the following paragraphs. Timing requirements relating to glitches are defined in clause 7.3.4.1.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14776-112:2002

Table 42 - SCSI bus timing values

Timing description	Timing values				
	Asynch	Fast-5	Fast-10	Fast-20	Fast-40
Arbitration Delay	2,4 µs	2,4 µs	2,4 µs	2,4 µs	2,4 µs
Bus Clear Delay	800 ns	800 ns	800 ns	800 ns	800 ns
Bus Free Delay	800 ns	800 ns	800 ns	800 ns	800 ns
Bus Set Delay	1,6 µs	1,6 µs	1,6 µs	1,6 µs	1,6 µs
Bus Settle Delay	400 ns	400 ns	400 ns	400 ns	400 ns
Cable Skew (note 1)	4 ns	4 ns	4 ns	3 ns	2,5 ns
Data Release Delay	400 ns	400 ns	400 ns	400 ns	400 ns
Disconnection Delay	200 µs	200 µs	200 µs	200 µs	200 µs
Power on to Selection (note 4)	10 s	10 s	10 s	10 s	10 s
Receive Assertion Period (note 2)	N/A	70 ns	22 ns	11 ns	6,5 ns
Receive Hold Time (note 2 and note 3)	N/A	25 ns	25 ns	11,5 ns	4,75 ns
Receive Negation Period (note 2)	N/A	70 ns	22 ns	11 ns	6,5 ns
Receive Setup Time (note 2 and note 3)	N/A	15 ns	15 ns	6,5 ns	4,75 ns
Receive Period Tolerance	N/A	1,1 ns	1,1 ns	1,1 ns	1,1 ns
Reset Hold Time	25 µs	25 µs	25 µs	25 µs	25 µs
Reset to Selection (note 4)	250 ms	250 ms	250 ms	250 ms	250 ms
Selection Abort Time	200 µs	200 µs	200 µs	200 µs	200 µs
Selection Time-out Delay (note 4)	250 ms	250 ms	250 ms	250 ms	250 ms
Signal Timing Skew (note 1)	8 ns	8 ns	8 ns	5 ns	4,5 ns
System Deskew Delay	45 ns	45 ns	20 ns	15 ns	8 ns
Transfer Period during Synchronous Data Transfer Phases (note 5)	N/A	200 ns	100 ns	50 ns	25 ns
Transmit Assertion Period (note 2)	N/A	80 ns	30 ns	15 ns	8 ns
Transmit Hold Time (note 2 and note 3)	N/A	53 ns	33 ns	16,5 ns	9,25 ns
Transmit Negation Period (note 2)	N/A	80 ns	30 ns	15 ns	8 ns
Transmit Setup Time (note 2 and note 3)	N/A	23 ns	23 ns	11,5 ns	9,25 ns
Transmit Period Tolerance	N/A	1 ns	1 ns	1 ns	1 ns
Notes: 1 This is measured at the device connection. 2 See 9.3 for measurement points for the timing specifications. 3 See 9.4 for examples of how to calculate setup and hold timing. 4 This is a recommended time. It is not mandatory. 5 The transfer period is measured from an assertion edge of REQ/REQQ (ACK/ACKQ) signal to the next assertion edge of the signal. 6 SCSI bus timing values specified by the maximum transfer rate shall apply even if a slower transfer rate is negotiated.					

9.2 Timing description

9.2.1 Arbitration delay

The minimum time an SCSI device shall wait from asserting the BSY signal for arbitration until the DATA BUS is examined to see if arbitration has been won. There is no maximum time.

9.2.2 Bus clear delay

The maximum time for an SCSI device to release all SCSI bus signals after:

- a) the BUS FREE phase is detected (the BSY and SEL signals are both false for a bus settle delay);
- b) the SEL signal is received from another SCSI device during the ARBITRATION phase;
- c) the transition of the RST signal to true.

For item a) above, the maximum time for an SCSI device to release all SCSI bus signals is 1200 ns from the BSY and SEL signals first becoming both false. If an SCSI device requires more than a bus settle delay to detect BUS FREE phase, it shall release all SCSI bus signals within a bus clear delay minus the excess time.

9.2.3 Bus free delay

The minimum time that an SCSI device shall wait from its detection of the BUS FREE phase (BSY and SEL both false for a bus settle delay) until its assertion of the BSY signal in preparation for entering the ARBITRATION phase.

9.2.4 Bus set delay

The maximum time for an SCSI device to assert the BSY signal and its SCSI ID after it detects a BUS FREE phase for the purpose of entering the ARBITRATION phase.

9.2.5 Bus settle delay

The minimum time to wait for the bus to settle after changing certain control signals as called out in the protocol definitions.

9.2.6 Cable skew

The maximum difference in propagation time allowed between any two SCSI bus signals measured between any two SCSI devices excluding any signal distortion skew delays.

9.2.7 Data release delay

The maximum time for an initiator to release the DATA BUS signals following the transition of the I/O signal from false to true.

9.2.8 Disconnection delay

The minimum time that a target shall wait after releasing BSY before participating in an ARBITRATION phase when honoring a DISCONNECT message from the initiator.

9.2.9 Power on to selection

The recommended maximum time from power application until an SCSI target is able to respond with appropriate status and sense data to the TEST UNIT READY, INQUIRY, and REQUEST SENSE commands (See SCSI-3 Primary Commands Standard).

9.2.10 Receive assertion period

The minimum time required at a SCSI device receiving a REQ or REQQ signal for the signal to be asserted while using synchronous data transfers. Also, the minimum time required at a SCSI device receiving an ACK or ACKQ signal for the signal to be asserted while using synchronous data transfers. For SE fast-5 and fast-10 operation, the time period is measured at the 0,8 V level. For SE fast-20 operation the period is measured at the 1,0 V level. The timings for the REQQ and ACKQ signals only apply to wide data transfers.

9.2.11 Receive hold time

The minimum time required at the receiving SCSI device between the assertion of the REQ or REQQ signal or the ACK or ACKQ signals and the changing of the DATA BUS while using synchronous data transfers. The timings for the REQQ and ACKQ signals only apply to 32-bit-wide data transfers.

9.2.12 Receive negation period

The minimum time required at a SCSI device receiving a REQ or REQQ signal for the signal to be negated while using synchronous data transfers. Also, the minimum time required at a SCSI device receiving an ACK or ACKQ signal for the signal to be asserted while using synchronous data transfers. For SE fast-5 and fast-10 operation, the time period is measured at the 2,0 V level. For SE fast-20 operation the period is measured at the 1,9 V level. The timings for the REQQ and ACKQ signals only apply to wide data transfers.

9.2.13 Receive setup time

The minimum time required at the receiving SCSI device between the changing of DATA BUS and the assertion of the REQ or REQQ signal or the ACK or ACKQ signal while using synchronous data transfers. The timings for the REQQ and ACKQ signals only apply to 32-bit-wide data transfers.

9.2.14 Receive period tolerance

The minimum tolerance that an SCSI device shall allow to be subtracted from the negotiated synchronous period.

9.2.15 Reset hold time

The minimum time that the RST signal is asserted. There is no maximum time.

9.2.16 Reset to Selection

The recommended maximum time from after a reset condition until an SCSI target is able to respond with appropriate status and sense data to the TEST UNIT READY, INQUIRY, and REQUEST SENSE commands (See SCSI-3 Primary Commands Standard).

9.2.17 Selection abort time

The maximum time that an SCSI device shall take from its most recent detection of being selected or reselected until asserting the BSY signal in response. This time-out is required to ensure that a target or initiator does not assert the BSY signal after a SELECTION or RESELECTION phase has been aborted.

9.2.18 Selection time-out delay

The minimum time an initiator or target should wait for an assertion of the BSY signal in response during the SELECTION or RESELECTION phase before starting the time-out procedure. Note that this is only a recommended time period.

9.2.19 Signal Timing Skew

The maximum signal timing skew occurs when transferring random data. The signal timing skew includes cable skew (measured with 0101... patterns) and signal distortion skew caused by random data patterns and transmission line reflections as shown in figure 43, figure 44, figure 45, and figure 46.

The receiver detection range is the part of the signal between the 'may detect' level and the 'shall detect' level on either edge. (see 9.3)

9.2.20 System deskew delay

The minimum time that a SCSI device should wait after receiving an SCSI signal to ensure that any signals transmitted at the same time are valid.

9.2.21 Transmit assertion period

The minimum time that a target shall assert the REQ or REQQ signal while using synchronous data transfers. Also, the minimum time that an initiator shall assert the ACK or ACKQ signal while using synchronous data transfers. The timing for the REQQ and ACKQ signals only apply to 32-bit-wide data transfers.

9.2.22 Transmit hold time

The minimum time provided by the transmitting SCSI device between the assertion of the REQ or REQQ signal or the ACK or ACKQ signal and the changing of the DATA BUS while using synchronous data transfers. The timings for the REQQ and ACKQ signals only apply to 32-bit wide data transfers.

9.2.23 Transmit negation period

The minimum time that a target shall negate the REQ or REQQ signal while using synchronous data transfers. Also, the minimum time that an initiator shall negate the ACK or ACKQ signal while using synchronous data transfers. The timing for the REQQ and ACKQ signals only apply to 32-bit -wide data transfers.

9.2.24 Transmit setup time

The minimum time provided by the transmitting SCSI device between the changing of DATA BUS and the assertion of the REQ or REQQ signal or the ACK or ACKQ signal while using synchronous data transfers. The timings for the REQQ and ACKQ signals only apply to 32-bit -wide data transfers.

9.2.25 Transmit period tolerance

The maximum tolerance that an SCSI device may subtract from the negotiated synchronous period.

9.3 Measurement points

The measurements points for SE and differential ACK, REQ, DATA, and PARITY signals are defined in this clause.

9.3.1 SE fast-10 data transfer rates

SE SCSI devices with data transfer rates up to and including fast-10 shall use the measurement points defined in figure 43 for the measurement of the timing values. The rise and fall times for the SE REQ/ACK signals shall be nominally the same as for the SE DATA/PARITY signals.




Figure 43 - Fast-10 SE timing measurement points

9.3.2 SE fast-20 data transfer rates

SE SCSI devices with data transfer rates up to and including fast-20 shall use the measurement points defined in figure 44 for the measurement of the timing values. The rise and fall times for the SE REQ/ACK signals shall be nominally the same as for the SE DATA/PARITY signals.

SE fast-20 timing measurement points shall apply even if a slower transfer rate is negotiated.





Figure 44 - Fast-20 SE timing measurement points

9.3.3 LVD

LVD SCSI devices shall use the measurement points defined in figure 45 for the measurement of the timing values. The rise and fall times for the LVD REQ/ACK signals shall be nominally the same as for the LVD DATA/PARITY signals.



* Use the crossing that yields the shorter SETUP and HOLD time

$$V_{TNA} = - \max [60 \text{ mV or } (0,25 \times V_N)]$$

$$V_{TAN} = \max [60 \text{ mV or } (-0,25 \times V_A)]$$

$$1,0 \text{ V} \geq V_N \geq 60 \text{ mV}$$

$$-1,0 \text{ V} \leq V_A \leq -60 \text{ mV}$$

Differential voltage signals in all cases

Notes:

1 V_N and V_A are measured at the start of a transition.

2 V_{TNA} - Voltage threshold negation to assertion.

3 V_{TAN} - Voltage threshold assertion to negation.


Figure 45 - LVD timing measurement points

Figure 45 illustrates that receivers may require a larger LVD signal to overcome a strongly asserted or negated state than required for a weakly asserted or negated state. With the maximum assertion level of V_A it requires a signal that crosses the zero differential level by at least $0,25 \times V_A$ (but at least by 60 mV in all cases) to guarantee detection of a negation for fast signals. The same relationship applies to the maximum negated level V_N . Conditions exist with longer loaded SCSI busses and irregular REQ and ACK pulse widths where long assertions or negations produce a much larger signal than short assertions or negations. This sets up an environment where the short REQ or ACK pulses may not have adequate timing margin unless the definitions in figure 45 are used in the measurement of timing parameters.

Measurement of driver timing parameters shall be performed using the circuit and test conditions defined in A.2.5 applied to the device connector. Receiver timing parameters are defined by the waveforms existing at the connector of the receiving SCSI device. The receiver timing parameters include the effects of data pattern. The receiver data pattern is therefore not defined.

9.3.4 HVD

HVD SCSI devices shall use the measurement points defined in figure 46 for the measurement of the timing values. The rise and fall times for the HVD REQ/ACK signals shall be nominally the same as for the HVD DATA/PARITY signals.



- * Use the crossing that yields the shorter SETUP and HOLD time
- $$V_{TNA} = -\max [200 \text{ mV or } (0,25 \times V_N)]$$
- $$V_{TAN} = \max [200 \text{ mV or } (-0,25 \times V_A)]$$
- $$1,0 \text{ V} \geq V_N \geq 200\text{mV}$$
- $$-1,0 \text{ V} \leq V_A \leq -200 \text{ mV}$$
- Differential voltage signals in all cases

Notes:

- 1 V_N and V_A are measured at the start of a transition.
- 2 V_{TNA} - Voltage threshold negation to assertion.
- 3 V_{TAN} - Voltage threshold assertion to negation.

Figure 46 - HVD timing measurement points

Figure 46 illustrates that receivers may require a larger HVD signal to overcome a strongly asserted or negated state than required for a weakly asserted or negated state. With the maximum assertion level of V_A it requires a signal that crosses the zero differential level by at least $0.25 \times V_A$ (but at least by 200 mV in all cases) to guarantee detection of a negation for fast signals. The same relationship applies for the maximum negated level V_N . Conditions exist with longer loaded SCSI busses and irregular REQ and ACK pulse widths where long assertions or negations produce a much larger signal than short assertions or negations. This sets up an environment where the short REQ or ACK pulses may not have adequate timing margin unless the definitions in figure 46 are used in the measurement of timing parameters.

9.4 Setup and hold timings

9.4.1 Fast-10 data transfer rates

Figure 47 and figure 48 show how the setup and hold times are calculated for various physical configurations on SCSI devices that support up to and including fast-10 data transfers. SCSI timing is specified at the SCSI connectors. To calculate the setup and hold timings for SCSI protocol chips, the following examples are provided




Figure 47 - Fast-10 setup and hold times for SE applications




Figure 48 - Fast-10 setup and hold timing for HVD applications

The receiver skew is the maximum difference in propagation time between any two receivers on the REQ, REQQ, ACK, ACKQ, DATA BUS or parity signals of the same SCSI bus when external receivers are used.

The transmitter skew is the maximum difference in propagation time between any two transmitters on the REQ, REQQ, ACK, ACKQ, DATA BUS or parity signals of the same SCSI bus when external transmitters are used.

The cable skew plus signal distortion skew is equal to the signal timing skew (see 9.2.19).

In systems with external transceivers, the total skew budget is 27 ns.

Transmitter chip	32 ns setup/42 ns hold	
Foil	1 ns	} 27 ns
External driver	recommend to be 8 ns	
----- TX connector -----		
25 meter cable	4 ns	
Distortion	4 ns	
Tolerance	0 ns	
----- RX connector -----		
External receiver	9 ns	
Foil	1 ns	
Receiver chip	5 ns setup/15 ns hold	

At its connector, the transmitting SCSI device should

- a) drive data no less than 23 ns before asserting the REQ_x or ACK_x signal;
- b) keep that data valid for no less than 33 ns following the assertion of the REQ_x or ACK_x signal.

The receiving SCSI device shall be able to latch the data at its connector when

- a) data is valid no more than 15 ns prior to the false-to-true transition of the REQ_x or ACK_x signal;
- b) data is valid no more than 25 ns following the false-to-true transition of REQ_x or ACK_x signal.

When 9 ns is added to the transmit device timing for transmitter skew and skew due to foil delays, the transmitting SCSI chip setup and hold timings are 32 ns and 42 ns, respectively. Similarly, when 10 ns is subtracted from the skew budget of the receiving SCSI device, 5 ns and 15 ns are left for receive chip setup and hold, respectively.

In the case of fast timing with no external transceivers over a 25 m cable, the total skew budget is 10 ns, compared to 27 ns. The 17 ns difference is used to relax the timing at the SCSI protocol chips (8 ns for the transmitting chip, and 9 ns for the receiving chip).

NOTE 21 - Component vendors may require that differential drivers and receivers be operated within restricted voltage and temperature differences to achieve the specified transmitter and receiver skew values.

9.4.2 Fast-20 data transfer rates

Figure 49 and figure 50 show how the setup and hold times are calculated for various physical configurations on SCSI devices that support up to and including fast-20 data transfers. SCSI timing is specified at the SCSI connectors. To calculate the setup and hold timings for SCSI protocol chips, the following examples are provided.

Fast-20 setup and hold times shall apply even if a slower transfer rate is negotiated.




Figure 49 - Fast-20 setup and hold times for SE applications





Figure 50 - Fast-20 setup and hold timing for HVD applications

The receiver skew is the maximum difference in propagation time between any two receivers on the REQ, REQQ, ACK, ACKQ, DATA BUS, or parity signals of the same SCSI bus when external receivers are used.

The transmitter skew is the maximum difference in propagation time between any two transmitters on the REQ, REQQ, ACK, ACKQ, DATA BUS, or parity signals of the same SCSI bus when external transmitters are used.

The cable skew plus signal distortion skew is equal to the signal timing skew (see 9.2.19).

The cable skew includes skew caused by non-uniform SCSI bus loading in mixed width applications, path length differences between different signals, and other factors affecting propagation time differences between the connectors, and between different signals.



At its connector, the transmitting SCSI device should:

- a) drive data no less than 11,5 ns before asserting the REQx or ACKx signal;
- b) keep that data valid for no less than 16,5 ns following the assertion of the REQx or ACKx signal.

The receiving SCSI device shall be able to latch the data at its connector when:

- a) data is valid no more than 6,5 ns prior to the false-to-true transition of the REQx or ACKx signal;
- b) data is valid no more than 11,5 ns following the false-to-true transition of REQx or ACKx signal.

When 4,5 ns is added to the transmit device timing for transmitter skew and skew due to foil delays, the transmitting SCSI chip setup and hold timings are 16 ns and 21 ns, respectively. Similarly, when 5,5 ns is subtracted from the skew budget of the receiving SCSI device, 1 ns and 6 ns are left for the receiving SCSI chip setup and hold, respectively.

In the case of fast-20 timing with no external transceivers over a 3 m signal path, the total skew budget is 6 ns, compared to 15 ns. The 9 ns difference is used to relax the timing at the SCSI protocol chips (4 ns for the transmitting chip, and 5 ns for the receiving chip).


NOTE 22 - Component vendors may require that differential drivers and receivers be operated within restricted voltage and temperature differences to achieve the specified transmitter and receiver skew values.

9.4.3 Fast-40 data transfer rates

Figure 51 shows how the setup and hold times are calculated for various physical configurations on SCSI devices that support up to and including fast-40 data transfers. The minimum set up and hold timings specified in figure 51 shall be used. Note that these values are different for the driver and the receiver but that the receiver sensitivity provides the threshold points for both. This is required because both extreme cases need to be covered:

- a) receivers connected to drivers with very short interconnect, and
- b) receivers connected to drivers through worst case interconnect.

Fast-40 setup and hold times shall apply even if a slower transfer rate is negotiated.



INT T/R (SEP T/R) indicates integrated (separate) protocol and transceiver chip

Board skew includes the separate transceiver and trace skew

Signal timing skew includes cable skew and signal distortion skew

Distortion skew includes ISI (intersymbol interference) and signal crossing time through the receiver detection range

Figure 51 - Fast-40 System setup and hold timings (all times in ns)

10 Removal and insertion of SCSI devices

10.1 Removal and insertion of SCSI devices overview

This clause defines the physical requirements for removal and insertion of SCSI devices on the SCSI bus. The issues related to the logical configuration of the SCSI bus and characteristics of the SCSI devices when a replacement occurs are beyond the scope of this standard. It should also be noted that the cases listed are distinguished for compatibility reasons and in most cases describe a system environment independent of this standard.

Four cases are addressed. The cases are differentiated by the state of the SCSI bus when the removal or insertion occurs.

10.2 Case 1 - Power-off during removal or insertion

- a) All SCSI devices are powered-off during physical reconfiguration.

10.3 Case 2 - RST signal asserted continuously during removal or insertion

- a) RST signal shall be asserted continuously by the initiator during removal or insertion.
- b) The system shall be designed such that the SCSI device being inserted shall make its power ground and logic ground connections at least 1 ms prior to the connection of any device connector contact to the SCSI bus. The ground connections shall be maintained during and after the connection of the SCSI device to the SCSI bus;
- c) The system shall be designed such that the SCSI device being removed shall maintain its power ground and logic ground prior to, during, and for at least 1 ms after the disconnection of any device connector contact from the SCSI bus.

NOTE 23 - The translation of the 1 ms time to mechanical provisions is vendor specific.

10.4 Case 3 - Current I/O processes not allowed during insertion or removal

- a) All I/O processes for all SCSI devices shall be quiesced;
- b) The system shall be designed such that the SCSI device being inserted shall make its power ground and logic ground connections at least 1 ms prior to the connection of any device connector contact to the SCSI bus. The ground connections shall be maintained during and after the connection of the SCSI device to the SCSI bus;
- c) The system shall be designed such that the SCSI device being removed shall maintain its power ground and logic ground prior to, during, and for at least 1 ms after the disconnection of any device connector contact from the SCSI bus;
- d) The SCSI device being removed or inserted shall employ transceivers that conform to the applicable requirements in 7.2.2, 7.3.2, 7.4.2, and 7.5.2 for glitch-free power on/off. The SCSI device shall maintain the high-impedance state at the device connector contacts during a power cycle until the transceivers are enabled. Power cycling includes on-board TERMPWR cycling caused by plugging, and SCSI device power cycling caused by plugging and switching;

NOTE 24 - Any on board switchable terminators as well as device transceivers may affect the impedance state at the device connector contacts.

- e) The SCSI device power may be simultaneously switched with the SCSI bus contacts if the power distribution system is able to maintain adequate power stability to other SCSI devices during the transition and the grounding requirements in items (b) and (c) above are met;
- f) The SCSI bus termination shall be external to the SCSI device being inserted or removed.
- g) Resumption of I/O processes is vendor specific but shall not occur sooner than 200 milliseconds after the completion of the insertion or removal event.
- h) Bypassing capacitors connecting to the TERMPWR line on the SCSI device being inserted or removed shall not exceed 10 μ F. For single-ended applications, SCSI bus terminations shall use voltage regulation.

NOTE 25 - In a multimode environment any insertion or removal that changes the SCSI bus mode causes

a transceiver mode change reset event (see 11.2.5).

10.5 Case 4 - Current I/O process allowed during insertion or removal

- a) All I/O processes for the SCSI device being inserted or removed shall be quiesced prior to removal.
- b) A SCSI device being inserted shall make its power ground and logic ground connection at least 1 ms prior to the connection of any device connector contact to the SCSI bus. The ground connections shall be maintained during and after the connection of the SCSI device to the SCSI bus;
- c) A SCSI device being removed shall maintain its power ground and logic ground prior to, during, and for at least 1 ms after the disconnection of any device connector contact from the SCSI bus;
- d) The SCSI device being removed or inserted shall employ transceivers that conform to the applicable requirements in 7.2.2, 7.3.2, and 7.5.2 for glitch-free power on/off. The SCSI device shall maintain the high-impedance state at the device connector contacts during a power cycle until the transceivers are enabled. Power cycling includes on board TERMPWR cycling caused by plugging, and SCSI device power cycling caused by plugging and switching;

NOTE 26 - Any on-board switchable terminators as well as SCSI device transceivers may affect the impedance state at the device connector contacts.

- e) The SCSI device power may be simultaneously switched with the SCSI bus contacts if the power distribution system is able to maintain adequate power stability to other SCSI devices during the transition and the grounding requirements in items (b) and (c) above are met;
- f) The SCSI bus termination shall be external to the SCSI device being inserted or removed;
- g) Initiation or resumption of I/O processes for a newly inserted or removed SCSI device is vendor specific but shall not occur sooner than 200 milliseconds after the completion of the insertion or removal event.
- h) Bypassing capacitors connecting to the TERMPWR line on the SCSI device being inserted or removed shall not exceed 10 μ F. For single-ended applications, SCSI bus terminations shall use voltage regulation.

NOTE 27 - In a multimode environment any insertion or removal that changes the bus mode causes a transceiver mode change reset event (see 11.2.5).

NOTE 28 - LVD SCSI devices may require more stringent system design to tolerate transients that occur during case 4 insertion or removal.

11 Logical characteristics

11.1 SCSI bus phases

11.1.1 SCSI bus phase overview

The SCSI architecture includes eight distinct phases:

- a) BUS FREE phase,
- b) ARBITRATION phase,
- c) SELECTION phase,
- d) RESELECTION phase,
- e) COMMAND phase,
- f) DATA phase,
- g) STATUS phase, and
- h) MESSAGE phase.

The COMMAND phase, DATA phase, STATUS phase, and MESSAGE phase are collectively termed the information transfer phases.

The SCSI bus phases are defined such that the SCSI bus is never in more than one phase at any given time. In the following descriptions, signals that are not mentioned shall not be asserted.

11.1.2 BUS FREE phase

The BUS FREE phase indicates that there is no current task and that the SCSI bus is available for a connection or reconnection.

SCSI devices shall detect the BUS FREE phase after the SEL and BSY signals are both false for at least a bus settle delay.

SCSI devices shall release all SCSI bus signals within a bus clear delay after the BSY and SEL signals become continuously false for a bus settle delay. If an SCSI device requires more than a bus settle delay to detect the BUS FREE phase then it shall release all SCSI bus signals within a bus clear delay minus the excess time to detect the BUS FREE phase. The total time to clear the SCSI bus shall not exceed a bus settle delay plus a bus clear delay.

During normal operation the BUS FREE phase is entered when a target releases the BSY signal.

11.1.3 Unexpected bus free

An unexpected bus free occurs when an initiator detects a BUS FREE phase that is not expected. Initiators only expect a BUS FREE phase to occur after one of the following occurs:

- a) after a hard reset is detected;
- b) after an ABORT TASK message is successfully received by a target;
- c) after an ABORT TASK SET message is successfully received by a target;
- d) after a CLEAR TASK SET message is successfully received by a target;
- e) after a LOGICAL UNIT RESET message is successfully received by a target;
- f) after a TARGET RESET message is successfully received by a target;
- g) after a TERMINATE TASK message is successfully received by a target;
- h) after a DISCONNECT message is successfully transmitted from a target;
- i) after a TASK COMPLETE message is successfully transmitted from a target;
- j) after a RELEASE RECOVERY message is successfully received by a target;
- k) after the release of the SEL signal after a SELECTION or RESELECTION phase time-out;
- l) after a transceiver mode change.

The target uses an unexpected bus free to inform the initiator of a protocol error. The target may switch to a BUS FREE phase at any time, except during an ARBITRATION phase, independent of the state of the ATN signal.

The target shall terminate the task that was the current task before the BUS FREE phase by clearing all data and status for that task. The target may optionally prepare sense data that may be retrieved by a REQUEST SENSE command. However, an unexpected bus free shall not create an exception condition.

The initiator shall terminate the task that was the current task before the BUS FREE phase occurred and shall manage this condition as an exception condition.

11.1.4 ARBITRATION phase

The ARBITRATION phase allows one SCSI device to gain control of the SCSI bus to allow that SCSI device to initiate or resume a task.

The procedure for an SCSI device to obtain control of the SCSI bus is as follows:

- a) The SCSI device shall first wait for the BUS FREE phase to occur. The BUS FREE phase is detected whenever both the BSY and SEL signals are simultaneously and continuously false for a minimum of a bus settle delay.

NOTE 29 - This bus settle delay is necessary because a transmission line phenomenon known as a wired-OR glitch may cause the BSY signal to briefly appear false, even though it is being driven true.

- b) The SCSI device shall wait a minimum of a bus free delay after detection of the BUS FREE phase (i.e. after the BSY and SEL signals are both false for a bus settle delay) before driving any signal.
- c) Following the bus free delay in step (b), the SCSI device may arbitrate for the SCSI bus by asserting both the BSY signal and its own SCSI ID, however the SCSI device shall not arbitrate (i.e. assert the BSY signal and its SCSI ID) if more than a bus set delay has passed since the BUS FREE phase was last observed.

NOTE 30 - There is no maximum delay before asserting the BSY signal and the SCSI ID following the bus free delay in step (b) as long as the bus remains in the BUS FREE phase. However, SCSI devices that delay longer than a bus settle delay plus a bus set delay from the time when the BSY and SEL signals first become false may fail to participate in arbitration when competing with faster SCSI devices.

- d) After waiting at least an arbitration delay (measured from its assertion of the BSY signal) the SCSI device shall examine the DATA BUS. If a higher priority SCSI ID bit is true on the DATA BUS (see table 39 for the SCSI ID arbitration priorities), then the SCSI device has lost the arbitration and the SCSI device may release its signals and return to step (a). If no higher priority SCSI ID bit is true on the DATA BUS, then the SCSI device has won the arbitration and it shall assert the SEL signal. Any SCSI device other than the winner has lost the arbitration and shall release the BSY signal and its SCSI ID bit within a bus clear delay after the SEL signal becomes true. An SCSI device that loses arbitration may return to step (a).

NOTE 31 - Step d) above requires any device that begins arbitration phase to complete the arbitration phase to the point of SEL being asserted if it begins the arbitration phase as stated in step c). This precludes the possibility of the bus being hung.

NOTE 32 - It is recommended that implementations wait for the SEL signal to be asserted before releasing the BSY signal and SCSI ID bit when arbitration is lost.

- e) The SCSI device that wins arbitration shall wait at least a bus clear delay plus a bus settle delay after asserting the SEL signal before changing any signals.

The SCSI ID bit is a single bit on the DATA BUS that corresponds to the SCSI device's unique SCSI address. All other DATA BUS bits shall be released by the SCSI device. Parity is not valid during the ARBITRATION phase. During the ARBITRATION phase, DB(P), DB(P1) (if present), and DB(P2) (if present) may be released or asserted, but shall not be actively driven false.

11.1.5 SELECTION phase

The SELECTION phase allows an initiator to select a target for the purpose of initiating some target function (e.g., READ or WRITE command). During the SELECTION phase the I/O signal is negated to distinguish this phase from the RESELECTION phase.

The SCSI device that won the arbitration has both the BSY and SEL signals asserted and has delayed at least a bus clear delay plus a bus settle delay before ending the ARBITRATION phase. The SCSI device that won the arbitration becomes an initiator by not asserting the I/O signal.

The initiator shall set the DATA BUS to a value that is the OR of its SCSI ID bit, the target's SCSI ID bit, and the appropriate parity bit(s) and it shall assert the ATN signal (indicating that a MESSAGE OUT phase is to follow the SELECTION phase). The initiator shall then wait at least two system deskew delays and release the BSY signal. The initiator shall then wait at least a bus settle delay before looking for an assertion of the BSY signal from the target.

The target shall determine that it is selected when the SEL signal and its SCSI ID bit are true and the BSY and I/O signals are false for at least a bus settle delay. The selected target may examine the DATA BUS in order to determine the SCSI ID of the selecting initiator. The selected target shall then assert the BSY signal within a selection abort time of its most recent detection of being selected; this is required for correct operation of the selection time-out procedure.

The target shall not respond to a selection if bad parity is detected (see 8.3). Also, if more than two SCSI ID bits are on the DATA BUS, the target shall not respond to selection.

No less than two system deskew delays after the initiator detects the BSY signal is true, it shall release the SEL signal and may change the DATA BUS. The target shall wait until the SEL signal is false before asserting the REQ signal to enter an information transfer phase.

11.1.6 SELECTION time-out procedure

Two optional selection time-out procedures are specified for clearing the SCSI bus if the initiator waits a minimum of a selection time-out delay and there has been no BSY signal response from the target:

- a) Optionally, the initiator shall assert the RST signal (see 11.2.3);
- b) Optionally, the initiator shall continue asserting the SEL and ATN signals and shall release the DATA BUS. If the initiator has not detected the BSY signal to be true after at least a selection abort time plus two system deskew delays, the initiator shall release the SEL and ATN signals allowing the SCSI bus to go to the BUS FREE phase. SCSI devices shall ensure that when responding to selection that the selection was still valid within a selection abort time of their assertion of the BSY signal. Failure to comply with this requirement could result in an improper selection (two targets connected to the same initiator, wrong target connected to an initiator, or a target connected to no initiator).

11.1.7 RESELECTION phase

11.1.7.1 RESELECTION phase overview

RESELECTION is a phase that allows a target to reconnect to an initiator for the purpose of continuing some operation that was previously started by the initiator but was suspended by the target, (i.e., the target disconnected by allowing a BUS FREE phase to occur before the operation was complete).

11.1.7.2 RESELECTION

Upon completing the ARBITRATION phase, the winning SCSI device has both the BSY and SEL signals asserted and has delayed at least a bus clear delay plus a bus settle delay. The winning SCSI device becomes a target by asserting the I/O signal. The winning SCSI device shall also set the DATA BUS to a value that is the logical OR of its SCSI ID bit and the initiator's SCSI ID bit and the appropriate parity bit(s). The target shall wait at least two system deskew delays and release the BSY signal. The target shall then wait at least a bus settle delay before looking for a response from the initiator.

The initiator shall determine that it is reselected when the SEL and I/O signals and its SCSI ID bit are true and the BSY signal is false for at least a bus settle delay. The reselected initiator may examine the DATA BUS in order to determine the SCSI ID of the reselecting target. The reselected initiator shall then assert the BSY signal within a selection abort time of its most recent detection of being reselected; this is required for correct operation of the time-out procedure. The initiator shall not respond to a RESELECTION phase if bad parity is detected (see 8.3). Also, the initiator shall not respond to a RESELECTION phase if other

than two SCSI ID bits are on the DATA BUS.

After the target detects the BSY signal is true, it shall also assert the BSY signal and wait at least two system deskew delays and then release the SEL signal. The target may then change the I/O signal and the DATA BUS. After the reselected initiator detects the SEL signal is false, it shall release the BSY signal. The target shall continue asserting the BSY signal until it relinquishes the SCSI bus.

NOTE 33 - When the target is asserting the BSY signal, a transmission line phenomenon known as a wired-OR glitch may cause the BSY signal to appear false for up to a round-trip propagation delay following the release of the BSY signal by the initiator. This is the reason why the BUS FREE phase is recognized only after both the BSY and SEL signals are continuously false for a minimum of a bus settle delay. For more information on glitches see 7.2.3 and 7.3.4.1.

11.1.7.3 RESELECTION time-out procedure

Two optional RESELECTION time-out procedures are specified for clearing the SCSI bus during a RESELECTION phase if the target waits a minimum of a selection time-out delay and there has been no BSY signal response from the initiator:

- a) Optionally, the target shall assert the RST signal (see 11.2.3);
- b) Optionally, the target shall continue asserting the SEL and I/O signals and shall release all DATA BUS signals. If the target has not detected the BSY signal to be true after at least a selection abort time plus two system deskew delays, the target shall release the SEL and I/O signals allowing the SCSI bus to go to the BUS FREE phase. SCSI devices shall ensure that the reselection was still valid within a selection abort time of their assertion of the BSY signal. Failure to comply with this requirement could result in an improper reselection (two initiators connected to the same target or the wrong initiator connected to a target).

11.1.8 Information transfer phases

11.1.8.1 Information transfer phases overview

The COMMAND, DATA, STATUS, and MESSAGE phases are all grouped together as the information transfer phases because they are all used to transfer data or control information via the DATA BUS. The actual content of the information is beyond the scope of this standard.

The C/D, I/O, and MSG signals are used to distinguish between the different information transfer phases (see table 43). The target drives these three signals and therefore controls all changes from one phase to another. The initiator requests a MESSAGE OUT phase by asserting the ATN signal, while the target causes the BUS FREE phase by releasing the MSG, C/D, I/O, and BSY signals.

When the data bus width is 32-bits, the REQQ and ACKQ signals are used to transfer the upper bytes of the word over the secondary SCSI bus. The REQQ and ACKQ signals are used in the same fashion as the REQ and ACK signals, but due to propagation time differences, the timing between the two sets of signals may not be identical. In the timing descriptions in this subclause, references to the REQ_x and ACK_x signals refer to the REQ and ACK signals for all 8/16 bit data transfers and the REQQ and ACKQ signals for 32-bit transfers.

The information transfer phases use one or more REQ_x/ACK_x handshakes to control the information transfer. Each REQ_x/ACK_x handshake allows the transfer of 8-, 16-, or 32-bits of information depending on the negotiated data transfer width (see 11.5.3.16). During the information transfer phases the BSY signal shall remain true and the SEL signal shall remain false. Additionally, during the information transfer phases, the target shall continuously envelope the REQ_x/ACK_x handshake(s) with the C/D, I/O, and MSG signals in such a manner that these control signals are valid for a bus settle delay before the assertion of the REQ signal of the first handshake and remain valid until after the negation of the ACK_x signal at the end of the handshake of the last transfer of the phase.

NOTE 34 - After the negation of the ACK_x signal of the last transfer of the phase, the target may prepare for a new phase by asserting or negating the C/D, I/O, and MSG signals. These signals may be changed together or individually. They may be changed in any order and may be changed more than once. It is desirable that each line change only once. A new phase does not begin until the REQ_x signal is asserted for the first byte of the new phase.

NOTE 35 - A phase is defined as ending when the C/D, I/O, or MSG signals change after the negation of the ACKx signal. The time between the end of a phase and the assertion of the REQx signal beginning a new phase is undefined.

Table 43 - Information transfer phases

Signal			Phase	Direction of transfer	Comment
MSG	C/D	I/O			
0	0	0	DATA OUT	Initiator to target	Data phase
0	0	1	DATA IN	Initiator from target	
0	1	0	COMMAND	Initiator to target	
0	1	1	STATUS	Initiator from target	
1	0	0	Reserved		
1	0	1	Reserved		
1	1	0	MESSAGE OUT	Initiator to target	Message phase
1	1	1	MESSAGE IN	Initiator from target	
Key: 0 = False; 1 = True					

11.1.8.2 Asynchronous information transfer

The target shall control the direction of information transfer by means of the I/O signal. When the I/O signal is true, information shall be transferred from the target to the initiator. When the I/O signal is false, information shall be transferred from the initiator to the target.

If the I/O signal is true (transfer to the initiator), the target shall first drive the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals to their desired values, delay at least one system deskew delay plus a cable skew, then assert the REQx signal. The DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals shall remain valid until the ACKx signal is true at the target. The initiator shall read the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals after the REQx signal is true, then indicate its acceptance of the data by asserting the ACKx signal. When the ACKx signal becomes true at the target, the target may change or release the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals and shall negate the REQx signal. After the REQx signal is false the initiator shall then negate the ACKx signal. After the ACKx signal is false the target may continue the transfer by driving the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals and asserting the REQx signal, as described above.

If the I/O signal is false (transfer to the target) the target shall request information by asserting the REQx signal. The initiator shall drive the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals to their desired values, delay at least one system deskew delay plus a cable skew and assert the ACKx signal. The initiator shall continue to drive the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals until the REQx signal is false. When the ACKx signal becomes true at the target, the target shall read the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals then negate the REQx signal. When the REQx signal becomes false at the initiator, the initiator may change or release the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals and shall negate the ACKx signal. After the ACKx signal is false the target may continue the transfer by asserting the REQx signal, as described above.

11.1.8.3 Synchronous data transfer

Synchronous data transfer is optional and is only used in data phases. It shall be used in a data phase if a

synchronous data transfer agreement has been established (see 11.5.3.13). The agreement specifies the REQx/ACKx offset and the minimum transfer period.

The REQx/ACKx offset specifies the maximum number of REQx assertions that shall be sent by the target in advance of the number of ACKx assertions received from the initiator, establishing a pacing mechanism. If the number of REQx assertions exceeds the number of ACKx assertions by the REQx/ACKx offset, the target shall not assert the REQx signal until after the next ACKx assertion is received. For successful completion of the data phase the number of ACKx and REQx assertions shall be equal.

The target shall assert the REQx signal for a minimum of an transmit assertion period. The target shall then wait at least the greater of a transfer period from the last transition of the REQx signal to true or a minimum of a transmit negation period from the last transition of the REQx signal to false before again asserting the REQx signal.

The initiator shall assert the ACKx signal for each REQx assertion received. The ACKx signal may be asserted as soon as the corresponding REQx assertion has been received. The initiator shall assert the ACKx signal for a minimum of an transmit assertion period. The initiator shall wait at least the greater of a transfer period from the last transition of the ACKx signal to true or for a minimum of a transmit negation period from the last transition of the ACKx signal to false before asserting the ACKx signal.

If the I/O signal is true (transfer to the initiator), the target shall first drive the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals to their desired values, wait at least one transmit setup time, then assert the REQx signal. The DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals shall be held valid for a minimum of one transmit hold time after the assertion of the REQx signal. The target shall assert the REQx signal for a minimum of an assertion period. The target may then negate the REQx signal and change or release the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals. The initiator shall read the value on the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals within one receive hold time of the transition of the REQx signal to true. The initiator shall then respond with an ACKx assertion.

If the I/O signal is false (transfer to the target), the initiator shall transfer one byte for each REQx assertion received. After detecting a REQx assertion, the initiator shall first drive the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals to their desired values, delay at least one transmit setup time, then assert the ACKx signal. The initiator shall hold the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals valid for at least one transmit hold time after the assertion of the ACKx signal. The initiator shall assert the ACKx signal for a minimum of an transmit assertion period. The initiator may then negate the ACKx signal and may change or release the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals. The target shall read the value of the DB(7-0,P), DB(15-0,P,P1), or DB(31-0,P,P1,P2,P3) signals within one receive hold time of the transition of the ACKx signal to true.

11.1.8.4 Wide data transfer

Wide data transfer is optional and may be used in the data phase only if a non-zero wide data transfer agreement is in effect (see WIDE DATA TRANSFER REQUEST message, 11.5.3.16). The messages determine the use of wide mode by both SCSI devices and establish a data path width to be used during the data phase.

Wide data transfers of 16- or 32-bits may be established. Targets and initiators that support 32-bit wide transfers should also support 16-bit wide transfers. All SCSI devices shall support 8-bit data transfers.

During 8-bit data transfers, the logical data byte for each data phase shall be transferred across the DB(7-0,P) signals on the primary SCSI bus. At the receiving device the DB(15-8,P1) (if present) signals are undefined and parity may not be valid. Subsequent data bytes are likewise transferred using DB(7-0,P).

During 16-bit wide data transfers, the first and second logical data bytes for each data phase shall be transferred across the DB(7-0,P) and DB(15-8,P1) signals, respectively, on the primary SCSI bus. Subsequent pairs of data bytes are likewise transferred in parallel across the primary SCSI bus (see table 44).

Table 44 - Wide SCSI byte order

Transfer number	Secondary bus		Primary bus		Data transfer width
	31....24	23....16	15.....8	7.....0	
1	N/A	N/A	N/A	W	8-bit
2	N/A	N/A	N/A	X	
3	N/A	N/A	N/A	Y	
4	N/A	N/A	N/A	Z	
1	N/A	N/A	X	W	16-bit
2	N/A	N/A	Z	Y	
1	Z	Y	X	W	32-bit
Note: When transferring consecutive bytes W, X, Y, and Z across the buses, they are transferred as shown above. This table does not necessarily represent how these bytes are stored in device memory.					

During 32-bit wide data transfers, the first and second logical data bytes shall be transferred across the DB(7-0,P) and DB(15-8,P1) signals, respectively. The third and fourth logical data bytes shall be transferred across the DB(23-16,P2) and DB(31-24,P3) signals, respectively. Subsequent sets of four data bytes are likewise transferred in parallel (see table 44).

If the last logical data byte transferred does not fall on the DB(15-8) signals for a 16-bit wide transfer or the DB(31-24) signals for a 32-bit wide transfer, then the values of the remaining higher-numbered bits are undefined. However, parity bits for these undefined bytes shall be valid for whatever data is placed on the bus.

To ensure proper data integrity on a 32-bit wide data bus, certain sequence requirements shall be met between the primary SCSI bus and the secondary SCSI bus:

- The REQQ and ACKQ signals shall only be asserted during DATA IN and DATA OUT phases. These signals shall not be asserted during other phases;
- The same information transfer mode (asynchronous or synchronous) shall be used for both the primary and secondary SCSI busses. If synchronous data transfer mode is in effect, the same REQ/ACK offset and transfer period shall be used for both busses;
- The target shall ensure that the number of REQx/ACKx handshakes on both the primary and secondary SCSI busses in a data phase are equal before it changes to another phase. The target shall not have a new phase, until the ACK and ACKQ signals have both become false for the last REQx/ACKx handshake.

If any violations of these rules are detected by the target, the target may attempt to end the data phase and return CHECK CONDITION status. If it is impossible to correctly terminate the data phase, the target may abnormally terminate the task by an unexpected bus free. If any violations of these rules are detected by the initiator, the initiator may attempt to send an INITIATOR DETECTED ERROR message to the target. If the initiator is unable to terminate the task normally, it may generate a hard reset.

The only pacing mechanism available for a target to manage the timing relationship between the signals on the two busses is its management of the REQ and REQQ signals. Similarly, the only pacing mechanism for the initiator to manage the timing between the two busses is its management of the ACK and ACKQ signals.

11.1.9 COMMAND phase

The COMMAND phase allows the target to request command information from the initiator.

The target shall assert the C/D signal and negate the I/O and MSG signals during the REQ/ACK handshake(s) of this phase.

11.1.10 Data phase

11.1.10.1 Data phase overview

The data phase is a term that encompasses both the DATA IN phase and the DATA OUT phase.

11.1.10.2 DATA IN phase

The DATA IN phase allows the target to request that data be sent to the initiator from the target.

The target shall assert the I/O signal and negate the C/D and MSG signals during the REQx/ACKx handshake(s) of this phase.

11.1.10.3 DATA OUT phase

The DATA OUT phase allows the target to request that data be sent from the initiator to the target.

The target shall negate the C/D, I/O, and MSG signals during the REQx/ACKx handshake(s) of this phase.

11.1.11 STATUS phase

11.1.11.1 STATUS phase overview

The STATUS phase allows the target to request that status information be sent from the target to the initiator.

The target shall assert the C/D and I/O signals and negate the MSG signal during the REQ/ACK handshake of this phase.

11.1.11.2 STATUS phase exception condition handling

If the initiator detects a parity error on the status byte the initiator shall create an attention condition. When the target switches to a MESSAGE OUT phase the initiator shall send an INITIATOR DETECTED ERROR message (see 11.5.3.6) to the target. This message notifies the target that the status byte was invalid.

11.1.12 Message phase

11.1.12.1 Message phase overview

The message phase is a term that references either a MESSAGE IN, or a MESSAGE OUT phase. Multiple messages may be sent during either phase. The first byte transferred in either of these phases shall be either a single-byte message or the first byte of a multiple-byte message. Multiple-byte messages shall be wholly contained within a single message phase.

11.1.12.2 MESSAGE IN phase

The MESSAGE IN phase allows the target to request that message(s) be sent to the initiator from the target.

The target shall assert the C/D, I/O, and MSG signals during the REQ/ACK handshake(s) of this phase.

11.1.12.3 MESSAGE IN phase exception condition handling

If the initiator detects a parity error on any message byte it receives the initiator shall create an attention condition. When the target switches to a MESSAGE OUT phase the initiator shall send a MESSAGE PARITY ERROR message (see 11.5.3.7) to the target. This message notifies the target that the message in byte was invalid.

11.1.12.4 MESSAGE OUT phase

The MESSAGE OUT phase allows the target to request that message(s) be sent from the initiator to the target. The target invokes this phase in response to the attention condition created by the initiator (see 11.2.2).

The target shall assert the C/D and MSG signals and negate the I/O signal during the REQ/ACK handshake(s) of this phase. The target shall handshake byte(s) in this phase until the ATN signal is negated, except when rejecting a message.

If the target receives all of the message byte(s) successfully (i.e. no parity errors), it shall indicate that it does not wish to retry by changing to any information transfer phase other than the MESSAGE OUT phase and transfer at least one byte. The target may also indicate that it has successfully received the message byte(s) by changing to the BUS FREE phase (e.g. ABORT TASK SET or TARGET RESET messages).

11.1.12.5 MESSAGE OUT phase exception condition handling

If the target detects one or more parity error(s) on the message byte(s) received, it may indicate its desire to retry the message(s) by asserting the REQ signal after detecting the ATN signal has gone false and prior to changing to any other phase. The initiator, upon detecting this condition, shall resend all of the previous message byte(s) in the same order as previously sent during this phase. When resending more than one message byte, the initiator shall assert the ATN signal at least two system deskew delays prior to asserting the ACK signal on the first byte and shall maintain the ATN signal asserted until the last byte is sent as described in 11.2.2.

If the target does not retry the MESSAGE OUT phase or it exhausts its retry limit it may;

- a) return CHECK CONDITION status and set the sense key to ABORTED COMMAND and the additional sense code to MESSAGE ERROR or;
- b) indicate a protocol error by performing an unexpected bus free.

The target may act on messages as received as long as no parity error is detected and may ignore all remaining messages sent under one ATN condition after a parity error is detected. When a sequence of messages is resent by an initiator because of a target detected parity error, the target shall not act on any message which it acted on the first time received.

11.1.13 Signal restrictions between phases

When the SCSI bus is between two information transfer phases, the following restrictions shall apply to the SCSI bus signals:

- a) The BSY, SEL, REQ, REQQ, ACK and ACKQ signals shall not change.
- b) The C/D, I/O, MSG, and DATA BUS signals may change. When switching the DATA BUS direction from out (initiator driving) to in (target driving), the target shall delay driving the DATA BUS by at least a data release delay plus a bus settle delay after asserting the I/O signal and the initiator shall release the DATA BUS no later than a data release delay after the transition of the I/O signal to true. When switching the DATA BUS direction from in (target driving) to out (initiator driving), the target shall release the DATA BUS no later than a system deskew delay after negating the I/O signal. The initiator shall assert the DATA BUS no sooner than a deskew delay after the detection of the negation of the I/O signal.
- c) The ATN and RST signals may change as defined under the descriptions for the attention condition (see 11.2.2) and hard reset (see 11.2.3).

11.2 SCSI bus conditions

11.2.1 SCSI bus conditions overview

The SCSI bus has two asynchronous conditions; the attention condition and the hard reset. These conditions cause the SCSI device to perform certain actions that may alter the phase sequence.

Furthermore, SCSI devices may not all be powered-on at the same time. This standard does not address power sequencing issues. However, each SCSI device, as it is powered on, should perform appropriate internal reset operations and internal test operations. Following a power-on to selection time after power-on, SCSI targets should be able to respond with appropriate status and sense data to the TEST UNIT READY, INQUIRY, and REQUEST SENSE commands.

11.2.2 Attention condition

The attention condition allows an initiator to inform a target that the initiator has a message ready. The target may get this message by performing a MESSAGE OUT phase.

The initiator creates the attention condition by asserting ATN at any time except during the ARBITRATION or BUS FREE phases.

The initiator shall negate the ATN signal at least two system deskew delays before asserting the ACK signal while transferring the last byte of the messages indicated with a Yes in tables 47, 57, and 62. If the target detects that the initiator failed to meet this requirement, then the target shall go to BUS FREE phase (see 11.1.2).

The initiator shall assert the ATN signal at least two system deskew delays before negating the ACK signal for the last byte transferred in a bus phase for the attention condition to be honored before transition to a new bus phase. Asserting the ATN signal later might not be honored until a later bus phase and then may not result in the expected action.

A target shall respond with MESSAGE OUT phase as follows:

- a) If the ATN signal becomes true during a COMMAND phase, the target shall enter MESSAGE OUT phase after transferring part or all of the command descriptor block bytes.
- b) If the ATN signal becomes true during a DATA phase, the target shall enter MESSAGE OUT phase at the target's earliest convenience (often, but not necessarily on a logical block boundary). The initiator shall continue REQ/ACK handshakes until it detects the phase change.
- c) If the ATN signal becomes true during a STATUS phase, the target shall enter MESSAGE OUT phase after the status byte has been acknowledged by the initiator.
- d) If the ATN signal becomes true during a MESSAGE IN phase, the target shall enter MESSAGE OUT phase before it sends another message. This permits a MESSAGE PARITY ERROR message from the initiator to be associated with the appropriate message.
- e) If the ATN signal becomes true during a SELECTION phase and before the initiator releases the BSY signal, the target shall enter MESSAGE OUT phase immediately after that SELECTION phase.
- f) If the ATN signal becomes true during a RESELECTION phase, the target shall enter MESSAGE OUT phase after the target has sent its IDENTIFY message for that RESELECTION phase.

NOTE 36 - The initiator should only assert the ATN signal during a RESELECTION phase to transmit a TARGET RESET or DISCONNECT message. Other uses may result in ambiguities concerning the nexus.

The initiator shall keep the ATN signal asserted if more than one byte is to be transferred. The initiator may negate the ATN signal at any time except it shall not negate the ATN signal while the ACK signal is asserted during a MESSAGE OUT phase. Normally, the initiator negates the ATN signal while the REQ signal is true and the ACK signal is false during the last REQ/ACK handshake of the MESSAGE OUT phase.

11.2.3 Hard reset

The hard reset is used to immediately clear all SCSI devices from the bus. This condition shall take precedence over all other phases and conditions. Any SCSI device may create the reset condition by

asserting the RST signal for a minimum of a reset hold time.

All SCSI devices shall release all SCSI bus signals (except the RST signal) within a bus clear delay of the transition of the RST signal to true. The BUS FREE phase always follows the reset condition.

The effect of the hard reset on tasks that have not completed, SCSI device reservations, and SCSI device operating modes is defined in the SCSI-3 Architecture Model standard.

In addition to the hard reset characteristics defined in the SCSI-3 Architecture Model standard this standard has the following hard reset characteristics:

- a) may or may not change the SCSI ID.

NOTE 37 - Environmental conditions (e.g. static discharge) may generate brief glitches on the RST signal. It is recommended that SCSI devices not react to these glitches. The manner of rejecting glitches is vendor-specific. The bus clear delay following a RST signal transition to true is measured from the original transition of the RST signal, not from the time that the signal has been confirmed. This limits the time to confirm the RST signal to a maximum of a bus clear delay.

11.2.4 Reset events

When an SCSI device detects a reset event it shall only initiate an internal hard reset (i.e., the SCSI device shall not assert the RST signal).

11.2.5 Transceiver mode change reset event

When an SCSI device that contains multimode transceivers detects a transceiver mode change from LVD mode to MSE mode it shall cause a reset event. In response to the transceiver mode change reset event, a target shall create a unit attention condition for all initiators. The unit attention condition sense key shall be set to UNIT ATTENTION, and the additional sense code set to TRANSCEIVER MODE CHANGED TO SE.

When an SCSI device that contains multimode transceivers detects a transceiver mode change from MSE mode to LVD mode it shall cause a reset event. In response to the transceiver mode change reset event, a target shall create a unit attention condition for all initiators. The unit attention condition sense key shall be set to UNIT ATTENTION, and the additional sense code set to TRANSCEIVER MODE CHANGED TO LVD.

Any SCSI device that detects a transceiver mode change shall:

- a) set the data transfer width to eight-bit transfer mode
- b) set the data transfer mode to asynchronous data transfer mode, and
- c) targets switch to a BUS FREE phase.

11.3 SCSI bus phase sequences

The order in which phases are used on the SCSI bus follows a prescribed sequence.

A hard reset aborts any phase and is always followed by the BUS FREE phase. Also any phase may be followed by the BUS FREE phase but many such instances are exception conditions (see 11.1.3).

The additional allowable sequences shall be as shown in figure 52. The normal progression is from the BUS FREE phase to ARBITRATION, from ARBITRATION to SELECTION or RESELECTION, and from SELECTION or RESELECTION to one or more of the information transfer phases (COMMAND, DATA, STATUS, or MESSAGE). The final information transfer phase is normally the MESSAGE IN phase where a DISCONNECT, or COMMAND COMPLETE message is transferred, followed by the BUS FREE phase.




Figure 52 - Phase sequences

11.4 SCSI pointers

The initiator provides for a set of three pointers for each task, called the saved pointers. The set of three pointers consists of one for the command, one for the data, and one for the status. When a send command service is received from an application client, the task's three saved pointers are copied into the initiator's set of three active pointers. There is only one set of active pointers in each initiator. The active pointers point to the next command, data, or status byte to be transferred between the initiator and the target. The saved and active pointers reside in the initiator.

The saved command pointer always points to the start of the command descriptor block for the task. The saved status pointer always points to the start of the status area for the task. The saved data pointer points to the start of the data area until the target sends a SAVE DATA POINTER message for the task.

In response to the SAVE DATA POINTER message, the initiator stores the value of the active data pointer into the saved data pointer for that task. The target may restore the active pointers to the saved pointer values for the current task by sending a RESTORE POINTERS message to the initiator. The initiator then copies the set of saved pointers into the set of active pointers. Whenever a target disconnects from the bus, only the set of saved pointers are retained. The set of active pointers is restored from the set of saved pointers upon reconnection of the task.

Since the data pointer value may be modified by the target before the task ends, it should not be used to test for actual transfer length because the value may no longer be valid.

11.5 SCSI interlocked protocol messages

11.5.1 SCSI interlocked protocol messages overview

SCSI interlocked protocol messages allow communication between an initiator and a target for the

purpose of link management. The link management messages used for this purpose are defined within this standard and their use is confined to this standard. Other SCSI interlocked protocol messages allow communication between the application client and the task manager for the purpose of task management. The task management messages are defined in the SCSI-3 Architecture Model standard, however, their binary values for the SCSI Parallel Interface-2 Standard are defined by this standard.

11.5.2 Message protocols and formats

11.5.2.1 Message protocol rules

One or more messages may be sent during a single MESSAGE phase, but a message may not be split between multiple MESSAGE phases.

The first message sent by the initiator after a successful SELECTION phase shall be an IDENTIFY, ABORT TASK SET (see 11.5.5.3), or TARGET RESET message. If a target receives any other message it shall switch to a BUS FREE phase (see 11.1.3).

If the first message is an IDENTIFY message, then it may be immediately followed by other messages, such as the first of a pair of SYNCHRONOUS DATA TRANSFER REQUEST messages. With tagged queuing a task attribute shall immediately follow the IDENTIFY message, then more messages may immediately follow. The IDENTIFY message establishes a logical connection between the initiator and the specified logical unit within the target known as an I_T_L nexus. After the RESELECTION phase, the target's first message shall be IDENTIFY. This allows the I_T_L nexus to be re-established. Only one logical unit shall be identified for any connection or reconnection; if a target receives a second IDENTIFY message with a different logical unit number during a connection or reconnection, it shall cause an unexpected bus free by generating a BUS FREE phase (see 11.1.3).

All initiators shall implement the mandatory messages tabulated in the "Initiator" column of tables 47, 57, and 62. All targets shall implement the mandatory messages tabulated in the "Target" column of tables 47, 57, and 62.

The initiator is required to end the MESSAGE OUT phase (by negating ATN) when it sends certain messages identified in tables 47, 57, and 62. These messages are identified by a "Yes" entry in the column headed "Negate ATN before last ACK".

Whenever an I_T_L nexus is established by an initiator that is allowing disconnection, the initiator shall ensure that the active pointers are equal to the saved pointers for that particular logical unit. An implied restore pointers operation shall occur as a result of a RESELECTION phase.

11.5.2.2 Message formats

11.5.2.2.1 Message formats overview

One-byte, Two-byte, and Extended message formats are defined. The first byte of the message determines the format as defined in table 45.

Table 45 - Message format

Code	Message format
00h	One-byte message (TASK COMPLETE)
01h	Extended messages
02h - 13h	One-byte messages
14h - 15h	Reserved One-byte messages
16h - 17h	One-byte messages
18h - 1Fh	Reserved One-byte messages
20h - 24h	Two-byte messages
25h - 2Fh	Reserved Two-byte messages
30h - 7Fh	Reserved
80h - FFh	One-byte message (IDENTIFY)

11.5.2.2.2 One-byte messages

One-byte messages consist of a single byte transferred during a MESSAGE IN phase or a MESSAGE OUT phase. The code of the byte determines which message is to be performed as defined in tables 47, 57, and 62.

11.5.2.2.3 Two-byte messages

Two-byte messages consist of two consecutive bytes transferred during two consecutive MESSAGE IN phases or two consecutive MESSAGE OUT phases. The code of the first byte determines which message is to be performed as defined in tables 47, 57, and 62. The second byte is a parameter byte which is used as defined in the message description.

11.5.2.2.4 Extended messages

A value of 01h in the first byte of a message indicates the beginning of a multiple-byte extended message. The minimum number of bytes sent for an extended message is three. All of the extended message bytes shall be transferred in consecutive MESSAGE IN phases or consecutive MESSAGE OUT phases. The extended message format is shown in table 46.

Table 46 - Extended message format

Bit Byte	7	6	5	4	3	2	1	0
0	EXTENDED MESSAGE (01h)							
1	EXTENDED MESSAGE LENGTH (n-1)							
2	EXTENDED MESSAGE CODE (y)							
3-n	EXTENDED MESSAGE ARGUMENTS							

The EXTENDED MESSAGE LENGTH specifies the length in bytes of the EXTENDED MESSAGE CODE plus the extended message arguments to follow. Therefore, the total length of the message is equal to the EXTENDED MESSAGE LENGTH plus two. A value of zero for the EXTENDED MESSAGE LENGTH indicates 256 bytes follow.

The EXTENDED MESSAGE CODES are listed in table 47.

The EXTENDED MESSAGE ARGUMENTS are specified within the extended message descriptions (see 11.5.3.9, 11.5.3.13, and 11.5.3.16).

11.5.3 Link control messages

11.5.3.1 Link control message codes

Table 47 - Link control message codes

Code	Support		Message Name	Direction		Negate ATN before last ACK
	Initiator	Target				
12h	O	O	CONTINUE TASK		Out	Yes
04h	O	O	DISCONNECT	In		n/a
04h	O	O	DISCONNECT		Out	Yes
80h+	M	O	IDENTIFY	In		n/a
80h+	M	M	IDENTIFY		Out	Not required
23h	O	O	IGNORE WIDE RESIDUE	In		n/a
05h	M	M	INITIATOR DETECTED ERROR		Out	Yes
09h	M	M	MESSAGE PARITY ERROR		Out	Yes
07h	M	M	MESSAGE REJECT	In	Out	Yes
***	O	O	MODIFY DATA POINTER	In		n/a
08h	M	M	NO OPERATION		Out	Yes
03h	O	O	RESTORE POINTERS	In		n/a
02h	O	O	SAVE DATA POINTER	In		n/a
***	O	O	SYNCHRONOUS DATA TRANSFER REQUEST	In	Out	Yes
13h	O	O	TARGET TRANSFER DISABLE		Out	Yes
00h	M	M	TASK COMPLETE	In		n/a
***	O	O	WIDE DATA TRANSFER REQUEST	In	Out	Yes
Key: M=Mandatory support, O=Optional support In=Target to initiator, Out=Initiator to target Yes=Initiator shall negate ATN before last ACK of message. Not required=Initiator may or may not negate ATN before last ACK of message (see 11.2.2). n/a=Not applicable ***=Extended message 80h+=Codes 80h through FFh are used for IDENTIFY messages						

11.5.3.2 CONTINUE TASK

The CONTINUE TASK message is sent from the initiator to the target to reconnect to a task. This message shall be sent as one of the messages within the consecutive MESSAGE OUT phases sent after the IDENTIFY message.

Thus the messages within the consecutive MESSAGE OUT phases following a SELECTION phase consists of the IDENTIFY, task attribute (if any), and CONTINUE TASK messages.

The purpose of the CONTINUE TASK message is to distinguish a valid initiator attempt at a reconnection phase from an incorrect initiator reconnection phase.

If the target expects a significant delay before it will be ready to continue processing the reconnected task, it may attempt to free the SCSI bus by sending a DISCONNECT message to the initiator. The initiator may reject the disconnection attempt by responding with MESSAGE REJECT message.

If the CONTINUE TASK message occurs on an initial connection then the target shall switch to a BUS FREE phase.

If the CONTINUE TASK message occurs on a subsequent reconnection then the target may either treat this as a dynamic head-of-queue request or it may reject the message with a MESSAGE REJECT message.

An initiator that gets rejected should assert the ATN signal and send an ABORT TAG message on the resulting MESSAGE OUT phase. Otherwise, the target may treat the reconnection as an incorrect initiator connection (see 11.6.3).

Initiators should avoid sending this message to targets that have not implemented this message. Such targets may not respond as described in this subclause. An application client may determine whether a device server implements this message by examining the TRANDIS bit in the standard INQUIRY data (see SCSI-3 Primary Commands Standard). The application client shall inform the initiator to use the CONTINUE TASK message by issuing a TARGET TRANSFER DISABLE link control function in the send SCSI command phase.

11.5.3.3 DISCONNECT

The DISCONNECT message is sent from a target to inform an initiator that the target plans to disconnect by releasing the BSY signal, and that a later reconnect will be required in order to complete the current task. This message shall not cause the initiator to save the data pointer. After successfully sending this message, the target shall go to the BUS FREE phase by releasing the BSY signal. The target shall consider the message transmission to be successful when it detects the negation of the ACK signal for the DISCONNECT message with the ATN signal false.

Targets that break data transfers into one or more reconnections shall end each successful data transfer (except possibly the last) with a SAVE DATA POINTER - DISCONNECT message sequence.

This message may also be sent from an initiator to a target to instruct the target to disconnect from the SCSI bus. If this option is supported, and after the DISCONNECT message is received, the target shall switch to MESSAGE IN phase, send the DISCONNECT message to the initiator (possibly preceded by SAVE DATA POINTER message), and then disconnect by releasing BSY. After releasing the BSY signal, the target shall not participate in another ARBITRATION phase for at least a disconnection delay or the time limit specified in the DISCONNECT TIME LIMIT mode parameter (see 11.7) whichever is greater. If this option is not supported or the target is not able to disconnect at the time when it receives the DISCONNECT message from the initiator, the target shall respond by sending a MESSAGE REJECT message to the initiator.

11.5.3.4 IDENTIFY

The IDENTIFY message (see table 48) is sent by either the initiator or the target to establish an I_T_L nexus.

Table 48 - IDENTIFY message format

Bit Byte	7	6	5	4	3	2	1	0
0	IDENTIFY	DISCPRIV	LUN					

The IDENTIFY bit shall be set to one to specify that this is an IDENTIFY message.

A disconnect privilege (DISCPRIV) bit of one specifies that the initiator has granted the target the privilege of disconnecting. A DISCPRIV bit of zero specifies that the target shall not disconnect. This bit is not defined and shall be set to zero when an IDENTIFY message is sent by a target.

The target shall generate a BUSY status (see SCSI-3 Architecture Model standard) for a task not granting disconnect privilege (DISCPRIV bit set to zero) in the IDENTIFY message if:

- a) there are any pending tasks, and
- b) the target determines that reconnection of one or more pending tasks is required before the current task is completed.

The LUN field specifies a logical unit number.

Only one logical unit number shall be identified per task. The initiator may send one or more IDENTIFY messages during a task. A second IDENTIFY message with a different value in the LUN field shall not be issued before a BUS FREE phase has occurred; if a target receives a second IDENTIFY message with a different value in this field, it shall cause an unexpected bus free (see 11.1.3) by generating a BUS FREE phase. Thus an initiator may change the DISCPRIV bit, but may not attempt to switch to another task. (See the DTDC field of the disconnect/reconnect mode page in the SCSI-3 Primary Commands Standard for additional controls over disconnection.)

An implied RESTORE POINTERS message shall be performed by the initiator following successful identification of the nexus during the MESSAGE IN phase of a reconnection.

Identification is considered successful during an initial connection or an initiator reconnect when the target detects no error during the transfer of the IDENTIFY message and an optional task attribute message in the MESSAGE OUT phase immediately following the SELECTION phase. See 11.5.4 for the ordering of the IDENTIFY and task attribute messages. See 11.1.12.5 for handling target detected errors during the MESSAGE OUT phase.

Identification is considered successful during a target reconnect when the ATN signal is not asserted during the transfer of either the IDENTIFY message or the SIMPLE message for an I_T_L_Q nexus in the MESSAGE IN phase immediately following the RESELECTION phase. See the 11.5.4 for the ordering of the IDENTIFY and task attribute messages. See 11.2.2, item d), for handling target detected errors during the MESSAGE IN phase.

11.5.3.5 IGNORE WIDE RESIDUE

The IGNORE WIDE RESIDUE message (see table 49) shall be sent from a target to indicate that the number of valid bytes sent in the last REQ/ACK handshake and REQQ/ACKQ handshake data of a DATA IN phase is less than the negotiated transfer width. This message shall be sent immediately following that DATA IN phase and prior to any other messages.

If the residual byte(s) contain(s) valid data then the IGNORE WIDE RESIDUE message should not be sent.

Table 49 - IGNORE WIDE RESIDUE message format

Bit Byte	7	6	5	4	3	2	1	0
0	MESSAGE CODE (23h)							
1	NUMBER OF BYTES TO IGNORE (01h, 02h, 03h)							

The NUMBER OF BYTES TO IGNORE field indicates the number of invalid data bytes transferred. See table 50 for a definition of the IGNORE field codes.

NOTE 38 - More than one IGNORE WIDE RESIDUE message may occur during a task.

Table 50 - IGNORE field definition

Codes	Invalid data bits	
	32-bit transfers	16-bit transfers
00h	Reserved	Reserved
01h	DB(31-24)	DB(15-8)
02h	DB(31-16)	Reserved
03h	DB(31-8)	Reserved
04h-FFh	Reserved	Reserved

11.5.3.6 INITIATOR DETECTED ERROR

The INITIATOR DETECTED ERROR message is sent from an initiator to inform a target that an error has occurred that does not preclude the target from retrying the task. The source of the error may either be related to previous activities on the SCSI bus or may be internal to the initiator and unrelated to any previous SCSI bus activity. Although the integrity of the currently active pointers is not assured, a RESTORE POINTERS message or a disconnect followed by a reconnect, shall cause the pointers to be restored to their defined prior state.

11.5.3.7 MESSAGE PARITY ERROR

The MESSAGE PARITY ERROR message is sent from the initiator to the target to indicate that it received a message byte with a parity error (see 11.1.12.5).

In order to indicate its intentions of sending this message, the initiator shall assert the ATN signal prior to its release of the ACK signal for the REQ/ACK handshake of the message byte that has the parity error. This provides an interlock so that the target is able to determine which message byte has the parity error. If the target receives this message under any other circumstance, it shall signal a catastrophic error condition by releasing the BSY signal without any further information transfer attempt (see 11.1.2).

If after receiving the MESSAGE PARITY ERROR message the target returns to the MESSAGE IN phase before switching to some other phase, the target shall resend the entire message that had the parity error.

11.5.3.8 MESSAGE REJECT

The MESSAGE REJECT message is sent from either the initiator or target to indicate that the last message or message byte it received was inappropriate or has not been implemented.

In order to indicate its intentions of sending this message, the initiator shall assert the ATN signal prior to its release of the ACK signal for the REQ/ACK handshake of the message byte that is to be rejected. If the target receives this message under any other circumstance, it shall reject this message.

When a target sends this message, it shall change to MESSAGE IN phase and send this message prior to requesting additional message bytes from the initiator. This provides an interlock so that the initiator is able to determine which message byte is rejected.

After a target sends a MESSAGE REJECT message and if the ATN signal is still asserted, then it shall return to the MESSAGE OUT phase. The subsequent MESSAGE OUT phase shall begin with the first byte of a message.

11.5.3.9 MODIFY DATA POINTER

The MODIFY DATA POINTER message (see table 51) is sent from the target to the initiator and requests that the signed ARGUMENT be added (two's complement) to the value of the current data pointer. The Enable Modify Data Pointer (EMDP) bit in the Disconnect-Reconnect mode page (see 11.7) indicates whether or not the target is permitted to issue the MODIFY DATA POINTER message.

Table 51 - MODIFY DATA POINTER message format

Bit Byte	7	6	5	4	3	2	1	0
0	EXTENDED MESSAGE (01h)							
1	EXTENDED MESSAGE LENGTH (05h)							
2	MODIFY DATA POINTER (00h)							
3	(MSB)							
4	ARGUMENT							
5								
6								
	(LSB)							

11.5.3.10 NO OPERATION

The NO OPERATION message is sent from an initiator in response to a target's request for a message when the initiator does not currently have any other valid message to send.

For example, if the target does not respond to the attention condition until a later phase and at that time the original message is no longer valid the initiator may send the NO OPERATION message when the target switches to a MESSAGE OUT phase.

11.5.3.11 RESTORE POINTERS

The RESTORE POINTERS message is sent from a target to direct the initiator to copy the most recently saved command, data, and status pointers for the task to the corresponding active pointers. The command and status pointers shall be restored to the beginning of the present command and status areas. The data pointer shall be restored to the value at the beginning of the data area in the absence of a SAVE DATA POINTER message or to the value at the point at which the last SAVE DATA POINTER message occurred for that task.

11.5.3.12 SAVE DATA POINTER

The SAVE DATA POINTER message is sent from a target to direct the initiator to copy the current data pointer to the saved data pointer for the current task.

11.5.3.13 SYNCHRONOUS DATA TRANSFER REQUEST

11.5.3.13.1 SYNCHRONOUS DATA TRANSFER REQUEST overview

SYNCHRONOUS DATA TRANSFER REQUEST (SDTR) messages (see table 52) are used to negotiate a synchronous data transfer agreement between two SCSI devices.

Table 52 - SYNCHRONOUS DATA TRANSFER message format

Bit Byte	7	6	5	4	3	2	1	0
0	EXTENDED MESSAGE (01h)							
1	EXTENDED MESSAGE LENGTH (03h)							
2	SYNCHRONOUS DATA TRANSFER REQUEST (01h)							
3	TRANSFER PERIOD FACTOR							
4	REQ/ACK OFFSET							

The TRANSFER PERIOD FACTOR field is defined in table 53.

Table 53 - TRANSFER PERIOD FACTOR field

Code	Description
00h-09h	Reserved (note 1)
0Ah	transfer period equals 25 ns (note 2)
0Bh	transfer period equals 30,3 ns (note 2)
0Ch	transfer period equals 50 ns (note 3)
0Dh-18h	transfer period equals the transfer period factor * 4 (note 3)
19h-31h	transfer period equals the transfer period factor * 4 (note 4)
32h-FFh	transfer period equals the transfer period factor * 4 (note 5)
note: 1 - Faster timings may be allowed by future SCSI parallel interface standards. 2 - Fast-40 data transfer rates that have a period equal 25 ns or 30,3 ns. 3 - Fast-20 data transfer rates that have a period of less than or equal 96 ns and greater than or equal to 50 ns. 4 - Fast-10 data transfer rates that have a period of less than or equal 196 ns and greater than or equal 100 ns. 5 - Fast-5 data transfer rates that have a period of less than or equal 1020 ns and greater than or equal to 200 ns.	

The REQ/ACK OFFSET is the maximum number of REQ assertions allowed to be outstanding before a corresponding ACK assertion is received at the target. The size of a data transfer may be 1, 2, or 4 bytes depending on what values, if any, have been previously negotiated through an exchange of WIDE DATA TRANSFER REQUEST messages. The REQ/ACK OFFSET value is chosen to prevent overflow conditions in the device's reception buffer and offset counter. A REQ/ACK OFFSET value of zero shall indicate asynchronous data transfer mode and that the TRANSFER PERIOD FACTOR field shall be ignored; a value of FFh shall indicate unlimited REQ/ACK offset.

An SDTR agreement applies to all logical units of the two SCSI devices that negotiated agreement. That is, if SCSI device A, acting as an initiator negotiates a synchronous data transfer agreement with SCSI device B (a target), then the same data transfer agreement applies to SCSI devices A and B even if SCSI

device B changes to an initiator.

A synchronous data transfer agreement only applies to the two SCSI devices that negotiate the agreement. Separate synchronous data transfer agreements are negotiated for each pair of SCSI devices. The synchronous data transfer agreement only applies to data phases.

An SDTR message exchange shall be initiated by an SCSI device whenever a previously arranged synchronous data transfer agreement may have become invalid. The agreement becomes invalid after any condition which may leave the data transfer agreement in an indeterminate state such as:

- a) after a hard reset;
- b) after a TARGET RESET message;
- c) after a power cycle and;
- d) after a change in the transceiver mode (e.g., LVD mode to MSE mode).

Any condition that leaves the data transfer agreement in an indeterminate state shall cause the SCSI device to enter an asynchronous data transfer mode.

An SCSI device may initiate an SDTR message exchange whenever it is appropriate to negotiate a new data transfer agreement (either synchronous or asynchronous). SCSI devices that are capable of synchronous data transfers shall not respond to an SDTR message with a MESSAGE REJECT message.

Renegotiation after every selection is not recommended, since a significant performance impact is likely.

The SDTR message exchange establishes the permissible transfer periods and the REQ/ACK offsets for all logical units on the two SCSI devices. This agreement only applies to DATA IN phases and DATA OUT phases. All other phases shall use asynchronous transfers.

The originating SCSI device (the SCSI device that sends the first of the pair of SDTR messages) sets its values according to the rules above to permit it to receive data successfully. If the responding SCSI device is able to also receive data successfully with these values (or smaller transfer periods or larger REQ/ACK offsets or both), it returns the same values in its SDTR message. If it requires a larger transfer period, a smaller REQ/ACK offset, or both in order to receive data successfully, it substitutes values in its SDTR message as required, returning unchanged any value not required to be changed. Each SCSI device when transmitting data shall respect the limits set by the other's SDTR message, but it is permitted to transfer data with larger transfer periods, smaller REQ/ACK offsets, or both than specified in the other's SDTR message. The completion of an exchange of SDTR messages implies an agreement as shown in table 54.

Table 54 - SDTR messages implied agreements

Responding agent SDTR response	Implied agreement
Non-zero REQ/ACK offset	Synchronous transfer (i.e., Each SCSI device transmit data with a transfer period equal to or greater than and a REQ/ACK offset equal to or less than the values received in the other device's SDTR message).
REQ/ACK offset equal to zero	Asynchronous transfer
MESSAGE REJECT message	Asynchronous transfer
Parity error (on responding message)	Asynchronous transfer
Unexpected bus free (as a result of the responding message)	Asynchronous transfer
No response	Asynchronous transfer

If there is an unrecoverable parity error on the initial SDTR message (see 11.1.12.2 and 11.1.12.4) the initiating SCSI device shall retain its previous data transfer mode. If there is an unexpected bus free on the initial SDTR message the initiating SCSI device shall retain its previous data transfer mode.

11.5.3.13.2 Target initiated SDTR negotiation

If the target recognizes that SDTR negotiation is required, it sends an SDTR message to the initiator.

Prior to releasing the ACK signal on the last byte of the SDTR message from the target, the initiator shall assert the ATN signal and respond with its SDTR message, MESSAGE PARITY ERROR message, or with a MESSAGE REJECT message.

If an abnormal condition prevents the initiator from responding with a SDTR message or with a MESSAGE REJECT message then both devices shall return to asynchronous data transfer mode for data transfers between the two SCSI devices.

Following an initiator's responding SDTR message, an implied agreement for synchronous operation shall not be considered to exist until the target leaves MESSAGE OUT phase, indicating that the target has accepted the SDTR negotiation.

If the target does not support any of the initiator's responding SDTR message's values the target shall switch to a MESSAGE IN phase and the first message shall be a MESSAGE REJECT message. In this case the implied agreement shall be considered to be negated and both SCSI devices shall use the asynchronous data transfer mode for data transfers between the two SCSI devices.

If a parity error occurs, the implied agreement shall be reinstated if a retransmission of a subsequent pair of messages is successfully accomplished. After a vendor-specific number of retry attempts (greater than zero), if the target continues to receive parity errors, it shall terminate the retry activity. This is done by the target causing an unexpected bus free. The initiator shall accept such action as aborting the SDTR negotiation, and both SCSI devices shall go to asynchronous data transfer mode for data transfers between the two SCSI devices.

11.5.3.13.3 Initiator initiated SDTR negotiation

If the initiator recognizes that SDTR negotiation is required, it asserts the ATN signal and sends a SDTR message to begin the negotiating process. After successfully completing the MESSAGE OUT phase, the target shall respond with the SDTR message or a MESSAGE REJECT message.

If an abnormal condition prevents the target from responding with a SDTR message or with a MESSAGE REJECT message then both SCSI devices shall go to asynchronous data transfer mode for data transfers between the two SCSI devices.

Following a target's responding SDTR message, an implied agreement for synchronous data transfers shall not be considered to exist until;

- a) the initiator receives the last byte of the SDTR message and parity is valid; and
- b) the target does not detect an assertion of the ATN signal before the ACK signal is released on the last byte of the SDTR message.

If the initiator does not support the target's responding SDTR message's values the initiator shall assert ATN and the first message shall be a MESSAGE REJECT message.

If during the SDTR message the initiator asserts the ATN signal and the first message out is either MESSAGE PARITY ERROR or MESSAGE REJECT the synchronous operation shall be considered to be negated by both the initiator and the target. In this case, both SCSI devices shall go to asynchronous data transfer mode for data transfers between the two SCSI devices.

11.5.3.14 TARGET TRANSFER DISABLE

The TARGET TRANSFER DISABLE message is sent from an initiator to a target to request that subsequent reconnections for data transfer be done by the initiator instead of the target. The target might reconnect for other purposes, but shall not switch to any DATA IN phase or DATA OUT phases after a target reconnection. SCSI devices that implement this message shall also implement the CONTINUE TASK message.

This message shall be sent as the last message of the series of consecutive messages of an initial connection. In the case where the initiator initiates a SDTR and/or a WDTR the initiator shall assert ATN during the last SDTR or WDTR negotiation within the same initial connection and send the target transfer disable message in response to the corresponding MESSAGE OUT phase. The target may continue the task, including any DATA OUT phases on the initial connection, until the target would normally disconnect, but the target shall not reconnect to transfer data. That is, the target shall not switch to a DATA IN phase on the initial connection and the target shall not switch to any DATA IN or DATA OUT phases on any subsequent target reconnection for the task.

When the target is ready to transfer data for a disconnected task for which a TARGET TRANSFER DISABLE message has been sent, the target shall reconnect to the initiator for the task (via a RESELECTION phase and consecutive MESSAGE IN phases containing an IDENTIFY message, and an optional SIMPLE TASK message), send a DISCONNECT message, and, if the initiator does not respond with a MESSAGE REJECT message, switch to a BUS FREE phase. This reconnection serves to notify the initiator that the task is ready for data transfer. If the initiator rejects the DISCONNECT message, the target may switch to a DATA IN or DATA OUT phase; otherwise, the initiator may reconnect to the task as described in the CONTINUE TASK message to perform the data transfer.

Initiators should avoid sending the TARGET TRANSFER DISABLE message to targets that have not implemented this message. Such targets may not respond as described in this subclause. An application client may determine whether a SCSI device implements this message by examining the (TRANDIS) bit in the standard INQUIRY data (see SCSI-3 Primary Commands Standard). The application client shall request the initiator use the CONTINUE TASK message by issuing a TARGET TRANSFER DISABLE link control function in the send SCSI command phase.

11.5.3.15 TASK COMPLETE

The TASK COMPLETE message is sent from a target to an initiator to indicate that a task has completed and that valid status has been sent to the initiator. After successfully sending this message, the target shall go to the BUS FREE phase by releasing the BSY signal. The target shall consider the message transmission to be successful when it detects the negation of ACK for the TASK COMPLETE message with the ATN signal false.

The task may have completed successfully or unsuccessfully as indicated in the status.

11.5.3.16 WIDE DATA TRANSFER REQUEST

11.5.3.16.1 WIDE DATA TRANSFER REQUEST overview

WIDE DATA TRANSFER REQUEST (WDTR) messages (see table 55) are used to negotiate a wide data transfer agreement between two SCSI devices.

Table 55 - WIDE DATA TRANSFER message format

Bit Byte	7	6	5	4	3	2	1	0
0	EXTENDED MESSAGE (01h)							
1	EXTENDED MESSAGE LENGTH (02h)							
2	WIDE DATA TRANSFER REQUEST (03h)							
3	TRANSFER WIDTH EXPONENT (m)							

The TRANSFER WIDTH EXPONENT field defines the transfer width to be used during DATA IN phases and DATA OUT phases. The transfer width that is established applies to all logical units on both SCSI devices. Valid transfer widths are 8 bits (m=00h), 16 bits (m=01h), and 32 bits (m=02h). Values of m greater than 02h are reserved.

A WDTR agreement applies to all logical units of the two SCSI devices that negotiated agreement. That is, if SCSI device A, acting as an initiator negotiates a wide data transfer agreement with SCSI device B (a target), then the same transfer width agreement applies to SCSI devices A and B even if SCSI device B changes to an initiator.

A wide data transfer agreement only applies to the two SCSI devices that negotiate the agreement. Separate wide transfer agreements are negotiated for each pair of SCSI devices. The wide data transfer agreement only applies to data phases.

A WDTR message exchange shall be initiated by an SCSI device whenever a previously arranged wide transfer agreement may have become invalid. The agreement becomes invalid after any condition which may leave the wide transfer agreement in an indeterminate state such as:

- a) after a hard reset;
- b) after a TARGET RESET message; and
- c) after a power cycle;
- d) after a change in the transceiver mode (e.g., LVD mode to MSE mode).

Any condition that leaves the data transfer agreement in an indeterminate state shall cause the SCSI device to enter an eight-bit wide data transfer mode.

An SCSI device may initiate a WDTR message exchange whenever it is appropriate to negotiate a new wide transfer agreement. SCSI devices that are capable of wide data transfers (greater than 8 bits) shall not respond to a WDTR message with a MESSAGE REJECT message.

Renegotiation after every selection is not recommended, since a significant performance impact is likely.

The WDTR message exchange establishes an agreement between the two SCSI devices on the width of the data path to be used for data phase transfers between two SCSI devices. This agreement only applies to DATA IN phases and DATA OUT phases. All other information transfer phases shall use an eight-bit data path.

If a SCSI device implements both wide data transfer option and synchronous data transfer option, then it shall negotiate the wide data transfer agreement prior to negotiating the synchronous data transfer agreement. If a synchronous data transfer agreement is in effect, then:

- a) if a WDTR message is rejected with a MESSAGE REJECT message the prior synchronous data transfer agreement shall remain intact;
- b) If a WDTR message fails for any other reason the prior synchronous data transfer agreement shall remain intact; or
- c) if a WDTR message is not rejected with a MESSAGE REJECT message a WDTR message shall reset the synchronous data transfer agreement to asynchronous mode.

The originating SCSI device (the SCSI device that sends the first of the pair of WDTR messages) sets its transfer width value to the maximum data path width it elects to accommodate. If the responding SCSI device is able to also accommodate this transfer width, it returns the same value in its WDTR message. If it requires a smaller transfer width, it substitutes the smaller value in its WDTR message. The successful completion of an exchange of WDTR messages implies an agreement as shown in table 56.

Table 56 - WDTR messages implied agreements

Responding agent WDTR response	Implied agreement
Non-zero TRANSFER WIDTH EXPONENT	Wide transfer (i.e., the initiator and the target transmit data with a transfer width equal to the responding device's transfer width). If the initiating SCSI device does not support the responding SCSI device's TRANSFER WIDTH EXPONENT then the initiating SCSI device shall MESSAGE REJECT the WDTR message (see 11.5.3.16.2 and 11.5.3.16.3).
TRANSFER WIDTH equal to zero	Eight-bit data transfer
MESSAGE REJECT message	Eight-bit data transfer
Parity error (on responding message)	Eight-bit data transfer
Unexpected bus free (as a result of the responding message)	Eight-bit data transfer
No response	Eight-bit data transfer

If there is an unrecoverable parity error on the initial WDTR message (see 11.1.12.2 and 11.1.12.4) the initiating SCSI device shall retain its previous data transfer mode. If there is an unexpected bus free on the initial WDTR message the initiating SCSI device shall retain its previous data transfer mode.

11.5.3.16.2 Target initiated WDTR negotiation

If the target recognizes that WDTR negotiation is required, it sends a WDTR message to the initiator.

Prior to releasing the ACK signal on the last byte of the WDTR message from the target, the initiator shall assert the ATN signal and respond with its WDTR message MESSAGE PARITY ERROR message, or with a MESSAGE REJECT message.

If an abnormal condition prevents the initiator from responding with a WDTR message with a MESSAGE REJECT message then both SCSI devices shall go to eight-bit data transfer mode for data transfers between the two SCSI devices.

Following an initiator's responding WDTR message, an implied agreement for wide data transfers operation shall not be considered to exist until the target leaves the MESSAGE OUT phase, indicating that the target has accepted the negotiation.

If the target does not support the initiator's responding TRANSFER WIDTH EXPONENT the target shall switch to a MESSAGE IN phase and the first message shall be a MESSAGE REJECT message. In this case the implied agreement shall be considered to be negated and both SCSI devices shall use the eight-bit data transfer mode for data transfers between the two SCSI devices. Any prior synchronous data transfer agreement shall remain intact.

If a parity error occurs, the implied agreement shall be reinstated if a retransmission of a subsequent pair of messages is successfully accomplished. After a vendor-specific number of retry attempts (greater than zero), if the target continues to receive parity errors, it shall terminate the retry activity. This is done by the target causing an unexpected bus free. The initiator shall accept such action as aborting the WDTR negotiation, and both SCSI devices shall go to eight-bit data transfer mode for data transfers between the two SCSI devices. Any prior synchronous data transfer agreement shall remain intact.

11.5.3.16.3 Initiator initiated WDTR negotiation

If the initiator recognizes that WDTR negotiation is required, it asserts the ATN signal and sends a WDTR message to begin the negotiating process. After successfully completing the MESSAGE OUT phase, the target shall respond with a WDTR message or a MESSAGE REJECT message.

If an abnormal condition prevents the target from responding with a WDTR message or with a MESSAGE REJECT message then both SCSI devices shall go to eight-bit transfer mode for data transfers between the two SCSI devices.

Following a target's responding WDTR message, an implied agreement for wide data transfers shall not be considered to exist until;

- a) the initiator receives the last byte of the WDTR message and parity is valid; and
- b) the target does not detect an assertion of the ATN signal before the ACK signal is released on the last byte of the WDTR message.

If the initiator does not support the target's responding TRANSFER WIDTH EXPONENT the initiator shall assert ATN and the first message shall be a MESSAGE REJECT message.

If during the WDTR message the initiator asserts the ATN signal and the first message of the MESSAGE OUT phase is either a MESSAGE PARITY ERROR or MESSAGE REJECT message the wide data transfers shall be considered to be negated by both SCSI devices. In this case, both SCSI devices shall use the eight-bit data transfer mode for data transfers between the two devices.

11.5.4 Task attribute messages

11.5.4.1 Task attribute messages overview

Two byte task attribute messages are used to specify an identifier, called a tag, for a task which establishes the I_T_L_Q nexus. The TAG field is an 8-bit unsigned integer assigned by the application client and sent to the initiator in the send SCSI command phase. The tag for every task for each I_T_L_Q nexus shall be uniquely assigned by the application client. There is no requirement for the task manager to check whether a tag is currently in use for another I_T_L_Q nexus. If the task manager check and receives a tag that is currently in use for the I_T_L_Q nexus, then it shall respond as defined in 11.6.3. A tag becomes available for reassignment when the task ends. The numeric value of a tag is arbitrary, providing there are no outstanding duplicates, and has no effect on the order of execution.

For each logical unit on each target, each application client has up to 256 tags to assign to tasks. Thus a target with eight logical units could have up to 14 336 tasks concurrently in existence if there were seven initiators on the bus.

Whenever an initiator connects to a target, the appropriate task attribute message shall be sent immediately following the IDENTIFY message to establish the I_T_L_Q nexus for the task. Only one I_T_L_Q nexus may be established during an initial connection or reconnection. If a task attribute message is not sent, then only an I_T_L_Q nexus is established for the task (i.e., an untagged command).

Whenever a target reconnects to an initiator to continue a tagged task, the SIMPLE QUEUE message shall be sent immediately following the IDENTIFY message to resume the I_T_L_Q nexus for the task. Only one I_T_L_Q nexus may occur during a reconnection. If the SIMPLE TAG message is not sent, then only an I_T_L_Q nexus occurs for the task (i.e., an untagged command).

If a target attempts to reconnect using an invalid tag, then the initiator should assert the ATN signal. After the corresponding MESSAGE OUT phase the initiator shall respond with an ABORT TASK message.

If a target does not implement tagged queuing and a queue tag message is received the target shall switch to a MESSAGE IN phase with a MESSAGE REJECT message and accept the task as if it were untagged provided there are no outstanding untagged tasks from that initiator.

See SCSI-3 Architecture Model standard for the task set management rules.

Table 57 - Task attribute message codes

Code	Support		Message Name	Direction		Negate ATN before last ACK
	Initiator	Target				
24h	O	O	ACA		Out	Not required
21h	Q	Q	HEAD OF QUEUE		Out	Not required
0Ah	O	O	LINKED COMMAND COMPLETE	In		n/a
0Bh	O	O	LINKED COMMAND COMPLETE (WITH FLAG)	In		n/a
22h	Q	Q	ORDERED		Out	Not required
20h	Q	Q	SIMPLE	In	Out	Not required
Key: M=Mandatory support, O=Optional support, Q=Mandatory if tagged queuing is implemented In=Target to initiator, Out=Initiator to target Yes=Initiator shall negate ATN before last ACK of message. Not required=Initiator may or may not negate ATN before last ACK of message (see 11.2.2). n/a=Not applicable ***=Extended message						

11.5.4.2 ACA

See table 58 for the format of the ACA message.

Table 58 - ACA message format

Bit Byte	7	6	5	4	3	2	1	0
0	MESSAGE CODE (24h)							
1	TAG (00h-FFh)							

The ACA message specifies that the task shall be placed in the task set as an ACA task. The rules used by the task manager to handle ACA tasks within a task set are defined in the SCSI-3 Architecture Model standard.

11.5.4.3 HEAD OF QUEUE

See table 59 for the format of the HEAD OF QUEUE message.

Table 59 - HEAD OF QUEUE message format

Bit Byte	7	6	5	4	3	2	1	0
0	MESSAGE CODE (21h)							
1	TAG (00h-FFh)							

The HEAD OF QUEUE message specifies that the task shall be placed in the task set as a HEAD OF QUEUE task. The rules used by the device server to handle HEAD OF QUEUE tasks within a task set are

defined in the SCSI-3 Architecture Model standard.

11.5.4.4 LINKED COMMAND COMPLETE

The LINKED COMMAND COMPLETE message is sent from a target to an initiator to indicate that a linked command has completed and that status has been sent. The initiator shall then set the pointers to the initial state for the next linked command.

11.5.4.5 LINKED COMMAND COMPLETE (WITH FLAG)

The LINKED COMMAND COMPLETE (WITH FLAG) message is sent from a target to an initiator to indicate that a linked command (with the flag bit set to one) has completed and that status has been sent. The initiator shall then set the pointers to the initial state of the next linked command. Typically this message would be used to cause an interrupt in the application client between two linked commands.

11.5.4.6 ORDERED

See table 60 for the format of the ORDERED message.

Table 60 - ORDERED message format

Bit Byte	7	6	5	4	3	2	1	0
0	MESSAGE CODE (22h)							
1	TAG (00h-FFh)							

The ORDERED message specifies that the task shall be placed in the task set as an ORDERED task. The rules used by the task manager to handle ORDERED tasks within a task set are defined in the SCSI-3 Architecture Model standard.

11.5.4.7 SIMPLE

See table 61 for the format of the SIMPLE message.

Table 61 - SIMPLE message format

Bit Byte	7	6	5	4	3	2	1	0
0	MESSAGE CODE (20h)							
1	TAG (00h-FFh)							

The SIMPLE message specifies that the task shall be placed in the task set as a SIMPLE task. The rules used by the task manager to handle SIMPLE tasks within a task set are defined in the SCSI-3 Architecture Model standard.

11.5.5 Task management messages

11.5.5.1 Task management message codes

Table 62 - Task management message codes

Code	Support		Message Name	Direction		Negate ATN before last ACK
	Initiator	Target				
0Dh	Q	Q	ABORT TASK		Out	Yes
06h	O	M	ABORT TASK SET		Out	Yes
16h	O	O	CLEAR ACA		Out	Not required
0Eh	Q	Q	CLEAR TASK SET		Out	Yes
17h	O	O	LOGICAL UNIT RESET (Note)		Out	Yes
0Ch	O	M	TARGET RESET		Out	Yes
11h	O	O	TERMINATE TASK		Out	Yes
Key: M=Mandatory support, O=Optional support, Q=Mandatory if tagged queuing is implemented In=Target to initiator, Out=Initiator to target Yes=Initiator shall negate ATN before last ACK of message. Not required=Initiator may or may not negate ATN before last ACK of message (see 11.2.2). n/a=Not applicable ***=Extended message						
Note-The LOGICAL UNIT RESET message is mandatory if hierarchical addressing (see SCSI Controller Command Standard) is implemented by the target.						

11.5.5.2 ABORT TASK

The ABORT TASK message is defined in the SCSI-3 Architecture Model standard.

In addition to the requirements in the SCSI-3 Architecture Model standard the target shall go to the BUS FREE phase following the successful receipt of the ABORT TASK message.

If only an I_T nexus has been established, the target shall go to the BUS FREE phase. No status or message shall be sent for the current task and no other task shall be affected.

NOTE 39 - The ABORT TASK message in the case of only an I_T nexus is useful to an initiator that is not able to get an IDENTIFY message through to the target due to parity errors and just needs to end the current task. Any pending data, status, or tasks for the I_T nexus are not affected. It is not possible to abort an I_T nexus on a reconnection because of item f) in 11.2.2.

On a reconnection, the ABORT TASK message aborts the current task if it is fully identified. If the current task is not fully identified (i.e. an I_T_L nexus exists, but the target is reconnecting for an I_T_L_Q nexus), then the current task is not aborted and the target goes to the BUS FREE phase.

NOTE 40 - A nexus is not fully identified on a reconnection if the ATN signal is asserted during or prior to the IDENTIFY message and the target only has tagged tasks for that initiator on that logical unit.

It is not an error to issue this message to an I_T_x nexus that does not have any pending tasks.

11.5.5.3 ABORT TASK SET

The ABORT TASK SET message is defined in the SCSI Architecture Model-2 standard.

In addition to the requirements in the SCSI-3 Architecture Model standard the target shall go to the BUS FREE phase following the successful receipt of the ABORT TASK SET message.

If only an I_T nexus has been established, the target shall switch to a BUS FREE phase. No status or message shall be sent for the current task and no other task shall be affected.

The ABORT TASK SET message in the case of only an I_T nexus is useful to an initiator that is not able to receive an IDENTIFY message through to the target due to parity errors and just needs to end the current task or task management function. No pending data, status, or tasks are affected.

It is not an error to issue this message to an I_T_L nexus that does not have any pending or current tasks.

11.5.5.4 CLEAR ACA

The CLEAR ACA message is defined in the SCSI-3 Architecture Model standard.

On receipt of a CLEAR ACA message the task manager, in addition to clearing the ACA condition, shall continue processing the current task.

It is not an error to issue a CLEAR ACA message when no ACA condition is in effect.

11.5.5.5 CLEAR TASK SET

The CLEAR TASK SET message is defined in the SCSI-3 Architecture Model standard.

In addition to the requirements in the SCSI-3 Architecture Model standard the target shall go to the BUS FREE phase following the successful receipt of the CLEAR TASK SET message.

11.5.5.6 LOGICAL UNIT RESET

The LOGICAL UNIT RESET message is defined in the SCSI-3 Architecture Model standard.

If only an I_T nexus has been established the LOGICAL UNIT RESET shall be performed as if it was a TARGET RESET.

In addition to the requirements in the SCSI-3 Architecture Model standard the target shall go to the BUS FREE phase following the successful receipt of the LOGICAL UNIT RESET message.

11.5.5.7 TARGET RESET

The TARGET RESET message is defined in the SCSI-3 Architecture Model standard.

In addition to the requirements in the SCSI-3 Architecture Model standard the target, following the successful receipt of the TARGET RESET message:

- a) shall go to the BUS FREE phase, and
- b) shall not change any SCSI IDs.

11.5.5.8 TERMINATE TASK

The TERMINATE TASK message is defined in the SCSI-3 Architecture Model standard.

If the task manager does not support this message or is unable to stop the current task, it shall send a MESSAGE REJECT message to the initiator and continue the task in a normal manner.

11.6 Command processing considerations and exception conditions

11.6.1 Command processing considerations and exception conditions overview

The following subclauses describe some aspects of command processing, including exception conditions and error handling which are specific to the SCSI Parallel Interface-2 Standard.

11.6.2 Asynchronous event notification

Notification of an asynchronous event is performed using the SEND command with the AER bit set to one. The information identifying the condition being reported shall be returned during the data out delivery phase of the SEND command (see SCSI-3 Primary Commands Standard).

An error condition or unit attention condition shall be reported once per occurrence of the event causing it. The target may choose to use an asynchronous event notification or to return CHECK CONDITION status on a subsequent command, but not both. Notification of command-related error conditions shall be sent only to the initiator that requested the task.

The asynchronous event notification protocol may be used to notify processor devices that a system resource has become available. If a target chooses to use this method, the sense key in the sense data sent to the processor device shall be set to UNIT ATTENTION.

The asynchronous event notification protocol shall be used only with SCSI devices that return processor device type with an AERC bit of one in response to an INQUIRY command. The INQUIRY command should be sent to logical unit zero of each SCSI device responding to selection. This procedure shall be conducted prior to the first asynchronous event notification and shall be repeated whenever the device deems it appropriate or when an event occurs that may invalidate the current information. (See SYNCHRONOUS DATA TRANSFER REQUEST message (11.5.3.13) for examples of these events.)

Each SCSI device that returns processor device type with an AERC bit of one shall be issued a TEST UNIT READY command to determine that the SCSI device is ready to receive an asynchronous event notification. A SCSI device returning CHECK CONDITION status is issued a REQUEST SENSE command. This clears any pending unit attention condition. A SCSI device that returns processor device type with an AERC bit of one and returns GOOD status when issued a TEST UNIT READY command shall accept a SEND command with an AER bit of one.

NOTE 41 - A SCSI device that uses asynchronous event notification at initialization time should provide means to defeat these notifications. This may be done with a switch or jumper wire. Devices which implement saved parameters may alternatively save the asynchronous event notification permissions either on a per SCSI device basis or as a system wide option. In any case, a device conducts a survey with INQUIRY commands to be sure that the devices on the SCSI bus are appropriate destinations for SEND commands with an AER bit of one. (The devices on the bus or the SCSI ID assignments may have changed.)

See asynchronous event reporting in the SCSI-3 Architecture Model standard for more information on asynchronous event notification.

11.6.3 Incorrect initiator connection

An incorrect initiator connection occurs on a reconnection if an initiator attempts to reconnect to a task, and then the initiator does not send an ABORT TASK SET, ABORT TASK, TARGET RESET, CLEAR TASK SET, CONTINUE TASK, or TERMINATE TASK message as one of the messages within the MESSAGE OUT phase after the IDENTIFY message.

A task manager that detects an incorrect initiator connection shall abort all tasks for the initiator and the associated logical unit and shall return CHECK CONDITION status for the task that caused the incorrect initiator connection. The sense key shall be set to ABORTED COMMAND and the additional sense code shall be set to OVERLAPPED COMMANDS ATTEMPTED with the additional sense code qualifier set to the value of the duplicate tag.

NOTE 42 - An incorrect initiator connection may be indicative of a serious error and, if not detected, could result in a task operating with a wrong set of pointers. This is considered a catastrophic failure on the part of the initiator. Therefore, vendor-specific error recovery procedures may be required to guarantee the data integrity on the medium. The target may return additional sense data to aid in this error recovery procedure (e.g., sequential-access devices may return the residue of blocks remaining to be written or read at the time the second command was received).

NOTE 43 - Some targets may not detect an incorrect initiator connection until after the command descriptor block has been received.

11.6.4 Unexpected reselection

An unexpected reselection occurs if a target attempts to reconnect to a task for which a nexus does not exist. An initiator should respond to an unexpected reselection by sending an ABORT TASK message.

11.7 Use of disconnect-reconnect page parameters

The disconnect-reconnect page (see SCSI-3 Primary Commands Standard) provides the application client the means to tune the performance of the SCSI parallel interface. The following subclause defines the fields in the disconnect-reconnect mode page of the MODE SENSE or MODE SELECT command that are used by targets.

The application client passes the fields used to control the SCSI parallel interface to a device server by means of a MODE SELECT command. The device server then communicates the field values to the target. The field values are communicated from the device server to the target in a vendor specific manner.

SCSI-3 parallel devices shall only use disconnect-reconnect page parameter fields defined below. If any other fields within the disconnect-reconnect page of the MODE SELECT command contain a non-zero value, the device server shall return CHECK CONDITION status for that MODE SELECT command. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

The BUFFER FULL RATIO field and BUFFER EMPTY RATIO FIELD are used as described in the SCSI-3 Primary Commands Standard.

The BUS INACTIVITY LIMIT field indicates the maximum time in 100 μ s increments that the target is permitted to assert the BSY signal without a REQ/ACK handshake. If the bus inactivity limit is exceeded the target shall attempt to disconnect (see 11.5.3.3) if the initiator has granted the disconnect privilege (see 11.5.3.4) and it is not restricted by DTDC. This value may be rounded as defined in the SCSI-3 Primary Commands Standard. A value of zero indicates that there is no bus inactivity limit.

The DISCONNECT TIME LIMIT field indicates the minimum time in 100 μ s increments that the target shall wait after releasing the SCSI bus before attempting reselection. This value may be rounded as defined in the SCSI-3 Primary Commands Standard. A value of zero indicates that there is no disconnect time limit.

The CONNECT TIME LIMIT field indicates the maximum time in 100 μ s increments that the target is allowed to use the SCSI bus before disconnecting, if the initiator has granted the disconnect privilege (see 11.5.3.4) and it is not restricted by DTDC. This value may be rounded as defined in the SCSI-3 Primary Commands Standard. A value of zero indicates that there is no connect time limit.

The MAXIMUM BURST SIZE field indicates the maximum amount of data that the target shall transfer during a data phase before disconnecting if the initiator has granted the disconnect privilege (see 11.5.3.4). This value is expressed in increments of 512 bytes (e.g. a value of one means 512 bytes, two means 1024 bytes, etc.). A value of zero indicates there is no limit on the amount of data transferred per burst.

The enable modify data pointers (EMDP) bit indicates whether or not the initiator allows the MODIFY DATA POINTERS message to be issued by the target. If the EMDP bit is a zero, the target shall not issue the MODIFY DATA POINTERS message. If the EMDP bit is a one, the target is allowed to issue MODIFY DATA POINTERS messages.

If the EMDP bit is a one and the initiator responds to a MODIFY DATA POINTER message with a MESSAGE REJECT, then the target shall return a CHECK CONDITION. The sense key shall be set to ABORTED COMMAND and the sense code shall be set to INVALID MESSAGE ERROR.

A disconnect immediate (DIMM) bit of zero indicates that the target may request DATA IN or DATA OUT phases following a COMMAND phase without attempting a disconnect (see 11.5.3.3). A DIMM bit of one indicates that the target shall attempt a disconnect (see 11.5.3.3) after a COMMAND phase and before a subsequent DATA IN or DATA OUT phase. The DIMM bit only applies when the initiator has granted the disconnect privilege (see 11.5.3.4).

The DATA TRANSFER DISCONNECT CONTROL (DTDC) field (see table 63) defines further restrictions on when a

disconnect is permitted.

TABLE 63 - DATA TRANSFER DISCONNECT CONTROL

DTDC	Description
000b	DATA TRANSFER DISCONNECT CONTROL is not used. Disconnect is controlled by the other fields in this page.
001b	A target shall not attempt to disconnect once the data transfer of a command has started until all data the command is to transfer has been transferred. The connect time limit and bus inactivity limit are ignored during the data transfer.
010b	Reserved
011b	A target shall not attempt to disconnect once the data transfer of a command has started, until the command is complete. The connect time limit and bus inactivity limit are ignored once data transfer has started.
100b-111b	Reserved

If DTDC is non-zero and the maximum burst size is non-zero, the target shall return a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

11.8 SCSI parallel interface services

11.8.1 SCSI parallel interface services overview

SCSI parallel interface services are provided by the initiator enabling the application client to accomplish tasks and task management functions (see SCSI-3 Architecture Model standard) and by the target enabling the device server to receive commands and move data to/from an application client. The SCSI parallel interface services are described in terms of the services the initiator and target provide. Each SCSI parallel interface service causes sequence phases to be generated by the addressed SCSI devices. figure 52 shows all the valid phase sequences.

11.8.2 Procedure terms

See table 64 for the mapping of the procedure terms used in this standard to the equivalent procedure terms used in the SCSI-3 Architecture Model standard.

Table 64 - SCSI Parallel Interface-2 Standard terms mapped to terms from other SCSI-3 standards

SCSI Parallel Interface-2 Standard terms	Equivalent SCSI-3 Architecture Model standard terms
initiator SCSI ID	initiator identifier
target SCSI ID	target identifier
initiator SCSI ID+target SCSI ID+logical unit number[+tag]	task identifier
target SCSI ID+logical unit number[+tag]	task address
target SCSI ID or target SCSI ID+logical unit number or initiator SCSI ID+target SCSI ID+logical unit number+tag	object identifier
target SCSI ID or target SCSI ID+logical unit number or target SCSI ID+logical unit number+tag	object address
target SCSI ID+initiator SCSI ID	none

See table 65 for a list of the procedure terms used when passing services across the SCSI parallel

interface service interface. See table 65 for the definitions of the SCSI Parallel Interface-2 Standard names and the equivalent SCSI-3 Architecture Model standard names of the procedure terms, the name of the standard where the terms are defined, the standard where the binary contents of the terms are defined, and the routing of the terms. The routing shows:

- a) the originating object of the term,
- b) the object that is the final destination of the term, and
- c) the objects that the term moves through to reach the final destination object.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14776-112:2002

Table 65 - Procedure terms

SCSI Parallel Interface-2 Standard terms	Standard where term defined	Standard where binary contents of term defined	Term routing
application client buffer offset	SAM	SAM	DS → targ → init
command byte count	SAM	SAM	AC → init
command descriptor block	SAM	SAM/cmd (note 2)	AC → init → targ → DS
data-in buffer	SAM	cmd (note 3)	DS → targ → init → AC
data-out buffer	SAM	cmd (note 3)	AC → init → targ → DS
device server buffer	SAM	cmd (note 3)	DS → targ → init
initiator SCSI ID	SAM	this standard	DS → targ or TM → targ
link control function	this standard	this standard	AC → init → targ
logical unit number	SAM	this standard	AC → init → targ → DS or AC → init → targ → TM or DS → targ → init
request byte count	SAM	SAM	DS → targ
service response	SAM	this standard (note 4)	DS → targ → init → AC or targ → DS
service response (note 1)	SAM	this standard (note 4)	init → AC
status	SAM	SAM	DS → targ → init → AC
tag	SAM	this standard	AC → init → targ → DS or AC → init → targ → TM or DS → targ → init
target SCSI ID	SAM	this standard	AC → init → targ → DS or AC → init → targ → TM or DS → targ
target SCSI ID+initiator SCSI ID	this standard	this standard	targ → DS or targ → TM
task attribute	SAM	this standard	AC → init → targ → DS
Key: AC=application client, cmd=SCSI command standards, DS=device server, init=initiator, SAM=SCSI-3 Architecture Model standard, TM=task manager, targ=target			
Notes 1) Only occurs when unexpected bus free (see 11.1.3) is detected by the initiator. 2) The portions not defined in the SCSI-3 Architecture Model standard are defined in the SCSI command standards (e.g., SCSI-3 Block Commands Standard, SCSI-3 Primary Commands Standard). 3) Parameter lists are defined within one of the SCSI command standards (e.g., SCSI-3 Block Commands Standard, SCSI-3 Primary Commands Standard). SCSI standards do not define non-parameter list information. 4) The SERVICE DELIVERY OR TARGET FAILURE value of the service response is not defined in SCSI-3.			

11.8.3 Application client SCSI command services

11.8.3.1 Application client SCSI command services procedure call

The SCSI command services shall be requested by the application client using a procedure call defined as:

Service response = execute command (target SCSI ID+logical unit number[+tag], command descriptor block, [task attribute], [link control function], [data-out buffer], [command byte count] || [data-in buffer], status, service response).

11.8.3.2 Send SCSI command service

The send SCSI command service is a four step confirmed service that provides the means to transfer a command data block to a device server.

Processing the execute command procedure call for a send SCSI command service shall be composed of the 4 step confirmed service shown in table 66.

Table 66 - Processing of send SCSI command service procedure

Step	Protocol service name	SCSI Interlocked Protocol Service Interface procedure call
request	send SCSI command request	send SCSI command (target SCSI ID+logical unit number[+tag], command descriptor block, [task attribute], [link control function], [data-out buffer], [command byte count]).
indication	send SCSI command indication	SCSI command received (target SCSI ID+initiator SCSI ID+logical unit number[+tag], [command descriptor block], [task attribute],).
response	send SCSI command response	send command complete (target SCSI ID+initiator SCSI ID+logical unit number[+tag], [status], [service response],).
confirmation	send SCSI command confirmation	command complete received (target SCSI ID+logical unit number[+tag], [data-in buffer], [status], service response).

11.8.4 Device server SCSI command services

11.8.4.1 Device server SCSI command services procedure call

The SCSI data buffer movement services shall be requested from the device server using a procedure call defined as:

Service response = move data buffer (target SCSI ID+initiator SCSI ID+logical unit number [+tag], device server buffer, application client buffer offset, request byte count ||).

Only one type of data buffer movement procedure call shall be used while processing one command, either data-in delivery or data-out delivery.

11.8.4.2 Data-in delivery service

The data-in delivery service is a two step confirmed service that provides the means to transfer a parameter list or data from a device server to an initiator.

Processing the execute command procedure call for a data-in delivery service shall be composed of the 2 step confirmed service shown in table 67.

Table 67 - Processing of data-in delivery service procedure

Step	Protocol service name	SCSI Interlocked Protocol Service Interface procedure call
request	data-in delivery request	send data-in (target SCSI ID+initiator SCSI ID+logical unit number[+tag], device server buffer, application client buffer offset, request byte count).
confirmation	data-in delivery confirmation	data delivered (target SCSI ID+initiator SCSI ID+logical unit number[+tag], service response).

11.8.4.3 Data-out delivery service

The data-out delivery service is a two step confirmed service that provides the means to transfer a parameter list or data from an initiator to a device server.

Processing the execute command procedure call for a data-out delivery service shall be composed of the 2 step confirmed service shown in table 68.

Table 68 - Processing of data-out delivery service procedure

Step	Protocol service name	SCSI Interlocked Protocol Service Interface procedure call
request	data-out delivery request	receive data-out (target SCSI ID+initiator SCSI ID+logical unit number[+tag], application client buffer offset, request byte count, device server buffer).
confirmation	data-out delivery confirmation	data-out received (target SCSI ID+initiator SCSI ID+logical unit number [+tag] service response)

11.8.5 Task management services**11.8.5.1 Task management services procedure call**

The task management services shall be requested from the application client using a procedure call defined as:

Service response = task management function (target SCSI ID|target SCSI ID+logical unit number|target SCSI ID+logical unit number+tag, service delivery failure flag || service response).

11.8.5.2 Task management function service**11.8.5.2.1 Task management function service overview**

This standard handles task management functions as a four step confirmed service that provides the means to transfer task management functions to a task manager.

The task management functions are defined in the SCSI-3 Architecture Model standard. This standard defines the actions taken by the SCSI parallel interface service to carry out the requested task management functions.

11.8.5.2.2 ABORT TASK

The SCSI parallel interface services request the initiator issue an ABORT TASK message (see 11.5.5.2) to the selected SCSI device.

11.8.5.2.3 ABORT TASK SET

The SCSI parallel interface services request the initiator issue an ABORT TASK SET message (see 11.5.5.3) to the selected SCSI device.

11.8.5.2.4 CLEAR ACA

The SCSI parallel interface services request the initiator issue a CLEAR ACA message (see 11.5.5.4) to the selected SCSI device.

11.8.5.2.5 CLEAR TASK SET

The SCSI parallel interface services request the initiator issue a CLEAR TASK SET message (see 11.5.5.5) to the selected SCSI device.

11.8.5.2.6 LOGICAL UNIT RESET

The SCSI parallel interface services request the initiator issue a LOGICAL UNIT RESET message (see 11.5.5.6) to the selected SCSI device.

11.8.5.2.7 RESET SERVICE DELIVERY SUBSYSTEM

The SCSI parallel interface services request the initiator issue a hard reset (see 11.2.3) to the selected SCSI device.

11.8.5.2.8 TARGET RESET

The SCSI parallel interface services request the initiator issue a TARGET RESET message (see 11.5.5.7) to the selected SCSI device.

11.8.5.2.9 TERMINATE TASK

The SCSI parallel interface services request the initiator issue a TERMINATE TASK message (see 11.5.5.8) to the selected SCSI device.

11.8.5.2.10 WAKEUP

The SCSI parallel interface services request the initiator issue a hard reset (see 11.2.3) to the selected SCSI device.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14776-112:2002

Annex A

(normative)

Additional requirements for LVD SCSI drivers and receivers

A.1 System level requirements

The requirements for LVD drivers and receivers in this annex are based on the system level requirements stated in table A.1. Some of these requirements are specifically called out in other subclauses while others are derived from bus loading conditions and trade-offs between competing parameters.

Table A.1 - System level requirements

Parameter	Minimum	Maximum	Cross-reference
V_A (mV)	-115		note 1
V_N (mV)		115	note 1
differential receiver input (mV)	$-V_N \times 0,25$	$-V_A \times 0,25$	figure 45
attenuation (%)		15	note 2
loaded media impedance (ohms)	85	135	note 3
unloaded media impedance (ohms)	110	135	subclause 6.3
terminator bias (mV)	100	125	subclause 7.3.1
terminator impedance (ohms)	100	110	subclause 7.3.1
device leakage (μ A)	-20	20	table 23
number of devices	2	16	subclause 4.1.7
ground offset level (mV)	-355	355	note 4
Note: 1 -These limits allow 60 mV base A.C. level and a maximum 55 mV for crosstalk and other non-common mode noise. 2 -Measured from the driver to the farthest receiver. 3 -Caused by the addition of device capacitive load (see annex G for calculations). 4 -This is the difference in voltage signal common for devices on the bus (see figure 2).			

A.2 Driver requirements

The fundamental requirement for a LVD driver is the generation of a first-step differential output voltage magnitude at the driver connections to the balanced media to achieve required minimum differential signals at every receiver connection to the bus. Other characteristics that affect overall noise margin are the common-mode output voltage, the maximum differential output voltage, the driver output impedance, and the output signal wave shape.

The driver requirements are defined in terms of the voltages and currents depicted in figure 41.

A.2.1 Differential output voltage, V_S

This subclause does not specify requirements for drivers with source impedances less than 1000 ohms.

To assure sufficient voltage to define a valid logic state at any device connection on a fully loaded LVD bus at least a minimum differential output voltage shall be generated. This value shall be large enough that, after allowance for attenuation, reflections, and differential noise coupling, V_S is at least ± 60 mV at the device connector to the LVD bus.

The SCSI device shall also comply with the upper limits for the differential output voltages and to the symmetry of the differential output voltage magnitudes between logic states in order to assure a first-step transition to the opposite logic state.

With the test circuit of figure A.1 and the test conditions V_1 and V_2 in table A.2 applied, the steady-state magnitude of the differential output voltage, V_S , for an asserted state (V_A), shall be greater than or equal to 270 mV and less than or equal to 780 mV. For the negated state, the polarity of V_S shall be reversed (V_N) and the differential voltage magnitude shall be greater than or equal to 260 mV and less than or equal to 640 mV. The relationship between V_A and V_N specified in table A.2 and shown graphically in figure A.2 shall be maintained.

The assertion drivers and negation drivers require different strengths to achieve the near equality in V_A and V_N shown in figure A.2 because the applied V_1 and V_2 simulate the effects of the bus termination bias.

Figure A.2 shall only apply to drivers with source impedances greater than 1000 ohms.

Table A.2 - Driver steady-state test limits and conditions

Test parameter	Test conditions (figure A.1)	Minimum (mV)	Maximum (mV)
V_A Differential output voltage magnitude (asserted) (note)	$V_1 = 0,957$ V $V_2 = 0,535$ V	270	780
	$V_1 = 1,949$ V $V_2 = 1,527$ V	270	780
V_N Differential output voltage magnitude (negated) (note)	$V_1 = 0,957$ V $V_2 = 0,535$ V	260	640
	$V_1 = 1,949$ V $V_2 = 1,527$ V	260	640
V_A Differential output voltage magnitude (asserted)	All four above conditions	$0,69 \times V_N + 50$	$1,45 \times V_N - 65$
Note: The test circuit (figure A.1) is approximately equivalent to two terminators creating the normal system bias.			




Figure A.1 - Differential steady-state output voltage test circuit




Figure A.2 - Domain for driver assertion and negation levels

A.2.2 Offset (common-mode output) voltage, V_{CM}

The steady-state magnitude of the driver offset voltage (V_{CM}), measured with the test load of figure A.3 shall be greater than or equal to 0,700 V and less than or equal to 1,800 V for either binary state. The steady-state magnitude of the difference of V_{CM} for one logical state and for the opposite logical state, ΔV_{CM} , shall be 120 mV or less for all $V_{applied}$ in the range: $0,700 \leq V_{applied} \leq 1,800$. See figure A.4.




Figure A.3 - Driver offset steady-state voltage test circuit




Figure A.4 - Common mode output voltage test

A.2.3 Short-circuit currents, I_{O-S} and I_{O+S}

Since an LVD bus allows multiple drivers, the possibility of contention requires a restriction on the power that may be sourced to the bus by a device. This is accomplished with a maximum allowable current from the driver.

With the driver output terminals short-circuited to a variable voltage source, the magnitudes of the currents (I_{O-S} and I_{O+S}) shall not exceed 24 mA for either logical state over the range $0 \leq V_{applied} \leq 2,5$ V. (see figure A.5).




Figure A.5 - Driver short-circuit test circuit

A.2.4 Open-circuit output voltages, $V_{O-(OC)}$ and $V_{O+(OC)}$

To limit the maximum steady-state voltage at any device connector, the voltage between each output terminal of the driver circuit and its common shall be between 0 V and 3,6 V when measured in accordance with figure A.6. This requirement shall be met in all logical or high impedance states ($0 \text{ V} \leq V_{O-(OC)} \leq 3,6 \text{ V}$ and $0 \text{ V} \leq V_{O+(OC)} \leq 3,6 \text{ V}$). The highest output voltage occurs with no output current.




Figure A.6 - Open-circuit output voltage test circuit

A.2.5 Output signal waveform

The differential output rise or fall time of a driver is specified since they influence the timing measurements and stub lengths of an LVD interface. Excessive over and under shoot of the output signal may cause electromagnetic emissions or false logic state changes.

During transitions of the driver output between alternating logical states (one - zero, zero - one, one - off, off - one, zero - off, off - zero), the differential voltage measured with the test circuit of figure A.7 and table A.3, shall be such that the voltage monotonically changes between 0,2 and 0,8 of the steady-state output, V_{SS} . V_{SS} is defined as the voltage difference between the two steady-state values of the driver output ($V_{SS} = |V_A| + |V_N|$) (See figure A.8 and table A.2). V_{SS} is expected to be different for different transitions.

The output signal rise or fall times (see t_r in figure A.8) between 0,2 and 0,8 of V_{SS} shall be greater than or equal to 1 ns.

The slew rates specified above are requirements for a driver when using the LVD test circuit in figure A.7. They are not the observed rise or fall rates on an actual SCSI bus.

Measurement equipment used for rise and fall rate testing shall provide a bandwidth of 2 GHz minimum.




Figure A.7 - Differential output switching voltage test circuit

Table A.3 - Driver switching test circuit parameters

Test condition (see figure A.7)	V1	V2
Low common-mode voltage	1,375 V	0,807 V
High common-mode voltage	1,693 V	1,125 V

The signal voltage shall comply with the requirements shown in figure A.8.




Figure A.8 - Driver output signal waveform

A.2.6 Dynamic output signal balance, $V_{CM(PP)}$

A mismatch in the magnitude of rate at which the voltage changes at the - signal and + signal connector pins, results in a common-mode AC signal. This may cause electromagnetic emissions from the media, excursions outside the receivers' common-mode input voltage range, and/or differential noise.

During transitions of the driver output between any state transition of high-to-low, low-to-high, high-to-off, off-to-high, low-to-off, or off-to-low, the voltage (V_{CM}) measured with the test circuit shown in figure A.9, shall not vary more than specified in table A.4 as $V_{applied}$ is varied over the range $0,700 \leq V_{applied} \leq 1,800$. Measurement equipment used for dynamic signal output balance testing shall provide a bandwidth of 400 MHz minimum. The requirements in this subclause apply only to the applicable state transitions.




Figure A.9 - Driver offset switching voltage test circuit

Lower values of V_{CM(PP)} have lower EMI risk.

Table A.4 - Dynamic output balance limits

Transition	V _{CM(PP)} mV max
high-low	120
low-high	120
high-off	400
low-off	400
off-high	400
off-low	400

A.3 Receiver characteristics

A receiver indicates the logical state of the LVD bus as defined by the differential voltage that exists at the device connector. A minimum steady state differential voltage defines the logic state. The receiver shall

detect this difference over the allowable common-mode input voltage range as determined by the driver and terminator output offsets and ground difference voltages.

Table A.5 defines the voltages and currents for the requirements in this subclause.

A.3.1 Receiver steady state input voltage requirements

Within the common-mode input voltage range (V_{CM}), (figure 42) $0,700\text{ V} \leq V_{CM} \leq 1,800\text{ V}$ a LVD receiver shall indicate the logical states shown in table A.5 with V_{IN} within the ranges shown in table A.5.

Table A.5 - Receiver steady state input voltage ranges


Input voltage range steady state (V_{IN})	Receiver detects
$-3,6\text{ V} \leq V_{IN} \leq -0,030\text{ V}$	1
$0,030\text{ V} \leq V_{IN} \leq 3,6\text{ V}$	0

A.3.2 Compliance test

Compliance to the requirements in A.3.1 shall be verified with the input voltages of table A.6 and the circuit of figure A.10.

Table A.6 - Receiver minimum and maximum input voltages.

Applied voltages (input voltage referenced to circuit common) (see figure A.10)		Resulting differential input voltage	Resulting common-voltage input voltage
V_-	V_+	V_{IN}	V_{CM}
0,715	0,685	0,030	0,700
0,685	0,715	-0,030	0,700
1,815	1,785	0,030	1,800
1,785	1,815	-0,030	1,800
3,600	0,000	3,600	1,800
0,000	3,600	-3,600	1,800
3,955	-0,355	4,310	1,800
-0,355	3,955	-4,310	1,800



>> Device Connector

Figure A.10 - Receiver input voltage threshold test circuit

A.3.3 Receiver setup and hold times

Figure 45 and figure 51 define the receiver setup and hold times.

NOTE 44 - Dynamic testing is required to verify these timings.

A.4 Transceiver characteristics

A.4.1 Transceiver output/input currents, I_{I-L} and I_{I+L}

The requirements in this clause apply as a test method to ensure compliance with table 22 and table 23. With the transceiver in an off condition (i.e., not transmitting) and the + and - signals connected to a variable voltage source, V_{applied} , the output leakage currents I_{I-L} and I_{I+L} shall not exceed the applicable values in table 23 over the range $0,00 \text{ V} \leq V_{\text{applied}} \leq 3,6 \text{ V}$ (see figure A.11). The maximum applicable current from table 23 is I_{max} .

These measurements apply with the transceiver's power supply in both power-on and power-off conditions.

$$|I_{I-L}| < I_{\text{max}}$$

$$|I_{I+L}| < I_{\text{max}}$$




Figure A.11 - Transceiver off-state output current test circuit

A.4.2 Transceiver maximum input voltages

See table 22 and table 23.

Annex B

(normative)

SCSI configured automatically (SCAM)

B.1 Model

Implementation of the SCAM protocol is optional, however, if implemented, the SCSI SCAM protocol shall conform to this annex.

The SCAM protocol is defined by this annex. SCAM is defined to ease user problems with the configuration of SCSI IDs on an SCSI bus. Level 2 SCAM defines all the hardware and software requirements necessary to implement all the functionality described in this annex. Level 1 SCAM defines a subset that does not support all of the advanced features of level 2 (e.g., hot plugging). It is intended to solve most configuration problems common to the single-user system.

B.2 Glossary

For the purposes of this annex, the following definitions apply in addition to those in the body of the standard.

B.2.1

assigned ID:

An SCSI ID which has been inherently (in the case of SCAM tolerant devices), explicitly (by SCAM protocol) or implicitly (by SCSI selection of a minimum duration) established for an SCSI device. When an ID is assigned it also becomes the current ID. Similarly, once an ID is assigned any change to the device's current ID by non-SCAM means (e.g., a MODE SELECT command) also changes the assigned ID.

B.2.2

current ID:

The SCSI ID that is available to an SCSI port. It may originate from jumpers, switches, mode parameters or some other source.

B.2.3

SCAM device:

Any SCSI device, initiator or target, that implements the SCAM protocol defined in this annex.

B.2.4

SCAM initiator:

A SCAM device that is capable of initiating SCAM selection and performing the normal functions of an SCSI initiator. These capabilities permit a SCAM initiator to scan an SCSI bus to discriminate between SCAM tolerant and SCAM devices and assign ID's to the SCAM devices.

B.2.5

SCAM target:

A SCAM device that is capable of recognizing and responding to SCAM selection. This capability permits a SCAM target to receive an ID assignment from a SCAM initiator. A SCAM target shall have a current ID, even when it is unassigned.

B.2.6

SCAM tolerant:

An SCSI device which does not implement the SCAM protocol but which complies with certain requirements specified by this annex. SCAM tolerant devices may be detected by a SCAM initiator and may be intermixed with SCAM devices.

B.2.7

unassigned ID:

The current SCSI ID which is available to the device but has not yet been assigned to the device.

B.3 SCAM requirements

B.3.1 Configuration requirements

SCAM configuration requirements permit SCAM tolerant, level 1 SCAM and level 2 SCAM devices to operate on the same SCSI bus. These requirements are:

- a) SCAM intolerant devices (i.e., legacy SCSI devices which are not SCAM tolerant) are not permitted on the bus;
- b) Any SCSI initiator on the bus shall be a SCAM initiator;
- c) No more than one level 1 SCAM initiator is permitted on the bus;
- d) Multiple level 2 SCAM initiators are permitted on the bus, which they may share with up to one level 1 SCAM initiator;
- e) All SCAM tolerant and level 1 SCAM targets on the bus shall be powered on before or concurrently with a SCAM initiator;
- f) If the only SCAM initiator is a level 1 SCAM initiator, all devices should be powered on before or concurrently with the level 1 SCAM initiator. The SCAM protocol does not provide the capability for SCAM level 1 initiators to detect level 2 SCAM targets powered on after the level 1 SCAM initiator has completed SCAM protocol, until a subsequent reset indication.

Some of these configuration requirements may be overcome by means outside the scope of this annex.

B.3.2 Timing requirements

Unless otherwise indicated, the time measurements for each SCAM or SCAM tolerant device, shown in table B.1, shall be measured for signal conditions existing at that SCSI device's own SCSI bus connection.

Table B.1 - SCAM timing values

Description	Value
SCAM tolerant power-on to selection delay	5 s
SCAM tolerant reset to selection delay	250 ms
SCAM tolerant selection response time	1 ms
SCAM unassigned ID selection response delay	4 ms
SCAM power-on to SCAM selection delay	1 s
SCAM reset to SCAM selection delay	250 ms
SCAM selection response time	250 ms
Recommended SCAM selection response time	1 ms
Wide arbitration time	7,2 μ s

B.3.2.1 SCAM tolerant power-on to selection delay

The maximum time a SCAM tolerant device may delay after power-on before enabling its response to selection.

B.3.2.2 SCAM tolerant reset to selection delay

A SCAM tolerant device shall enable its response to selection within this time limit, measured from the release of the RST signal.

A SCAM initiator shall wait at least a SCAM tolerant reset to selection delay before starting SCSI ID categorization (see B.5.1.1.1).

B.3.2.3 SCAM tolerant selection response time

A SCAM tolerant device shall respond to selection of its current ID within this time limit, provided that both the SCAM tolerant power-on to selection and reset to selection delays have been satisfied.

A SCAM initiator should use a minimum selection time-out of a SCAM tolerant selection response time plus two bus settle delays when scanning the bus for SCAM tolerant devices.

B.3.2.4 SCAM unassigned ID selection response delay

The minimum time a SCAM device shall delay before responding to selection of its current ID, provided that the SCAM device has not been assigned an ID since the last power-on or reset.

A SCAM initiator should use a maximum selection time-out less a SCAM unassigned ID selection response delay when scanning the bus for SCAM tolerant devices.

B.3.2.5 SCAM power-on to SCAM selection delay

A SCAM device shall enable its response to SCAM protocol initiation within this time limit.

A SCAM initiator or level 2 SCAM target shall wait at least a SCAM power-on to SCAM selection delay before initiating SCAM protocol.

B.3.2.6 SCAM reset to SCAM selection delay

The minimum time, measured from the bus free that immediately follows a reset, a SCAM device shall delay after a reset before initiating SCAM protocol.

B.3.2.7 SCAM selection response time

The maximum time a SCAM device shall require to detect and respond to SCAM selection. This is also the minimum time a SCAM initiator should maintain SCAM selection in situations where a slow response by other SCAM devices is anticipated (e.g. firmware SCAM implementations).

B.3.2.8 Recommended SCAM selection response time

The minimum time a SCAM device should maintain SCAM selection in situations where a rapid response by other SCAM devices is anticipated (e.g. hardware SCAM implementations). This is also the recommended maximum time a SCAM device should require to detect and respond to SCAM selection.

B.3.2.9 Wide arbitration time

The maximum time after the assertion of BSY within which a SCAM or SCAM tolerant device with an ID greater than 7 shall assert the SEL signal.

NOTE 45 - This requirement is necessary for arbitration without an ID to work on mixed width buses. It is based on the assumption that all wide SCSI devices implement arbitration logic in hardware and therefore relied on to assert the SEL signal quickly if they win arbitration.

B.3.3 Device requirements

The following subclauses define the operational requirements of SCAM and SCAM tolerant devices that may be configured on the same SCSI bus.

In addition, all SCAM devices shall disable active negation of SCSI bus signals during SCAM protocol.

B.3.3.1 SCAM tolerant target

A SCAM tolerant target:

- a) shall enable its response to selection within a SCAM tolerant power-on to selection delay after the device is powered-on.
- b) shall enable its response to selection within a SCAM tolerant reset to selection delay after a reset.
- c) shall recognize a valid selection of the device's current ID whether or not an initiator ID is present during the SELECTION phase.
- d) shall, once selection response is enabled, recognize a valid selection of its current ID and assert the BSY signal no later than a SCAM tolerant selection response time after the start of the SELECTION phase. The current ID and the assigned ID are the same for SCAM tolerant devices.

NOTE 46 - It is recommended that initiators clear the DISCPRIV bit in the IDENTIFY message if selection is performed without an initiator ID.

NOTE 47 - The requirement for rapid response to selection by SCAM tolerant devices and delayed response to selection by SCAM devices that do not have assigned ID's permits SCAM initiators to distinguish between the two. A SCAM initiator may use a relatively short selection time-out (SCAM tolerant selection response time plus two bus settle delays) to scan the bus for SCAM tolerant devices without causing the assignment of an ID.

B.3.3.2 Level 1 SCAM initiator

A level 1 SCAM initiator:

- a) shall recognize reset at all times.

NOTE 48 - SCAM implementations, whether in firmware or hardware, are expected to monitor the RST signal even while engaged in SCAM protocol.

- b) shall be capable of initiating SCAM protocol and utilizing SCAM function sequences to assign ID's to SCAM devices. Level 1 SCAM initiators are not required to detect or respond to SCAM selection.
- c) shall be capable of detecting a dominant initiator contention function code and subsequently participate in the isolation stage for the dominant initiator.
- d) It is recommended that level 1 SCAM initiators perform dominant initiator contention each time SCAM protocol is initiated.
- e) shall have an assigned ID.
- f) shall be able to operate with a selection time-out greater than the SCAM tolerant selection response time and less than the SCAM unassigned ID selection response delay. A level 1 SCAM initiator shall also be able to operate with a selection time-out greater than the SCAM unassigned ID selection response delay.
- g) shall not assert the RST signal upon a selection time-out.

B.3.3.3 Level 1 SCAM target

A level 1 SCAM target:

- a) shall recognize reset at all times.

NOTE 49 - SCAM implementations, whether in firmware or hardware, are expected to monitor the RST signal even while engaged in SCAM protocol.

- b) shall enable its response to SCAM selection within a SCAM power-on to SCAM selection delay after the device is powered-on.
- c) shall enable its response to SCAM selection within a SCAM reset to SCAM selection delay after a

reset.

- d) shall, once SCAM selection response is enabled and provided that its device ID is unassigned, recognize and respond to SCAM selection within a SCAM selection response time.
- e) shall not, while its ID remains unassigned, start a SELECTION phase unless the SEL signal and the SCSI ID bit that encodes the unassigned ID are true and the BSY and I/O signals are false for at least a SCAM unassigned ID selection response delay. Response to selection shall cause the device to have an assigned ID equal to its current ID.
- f) shall, once assigned an ID, behave as a SCAM tolerant device until a subsequent power-on or reset.

NOTE 50 - SCAM devices with assigned ID's neither recognize, respond to, nor initiate SCAM selection.

- g) shall not assert the RST signal a selection time-out.
- h) shall not implement the soft reset alternative as defined in SCSI-2.

B.3.3.4 Level 2 SCAM initiator

A level 2 SCAM initiator:

- a) shall recognize reset at all times.

NOTE 51 - SCAM implementations, whether in firmware or hardware, are expected to monitor the RST signal even while engaged in SCAM protocol.

- b) shall be capable of initiating SCAM protocol and utilizing SCAM function sequences to assign ID's to SCAM devices. Level 2 SCAM initiators are also required to detect and respond to SCAM selection initiated by other SCAM devices.
- c) shall perform dominant initiator contention each time SCAM protocol is initiated.
- d) shall have either an assigned ID or be able to arbitrate without an ID.
- e) shall be able to operate with a selection time-out greater than the SCAM tolerant selection response time and less than the SCAM unassigned ID selection response delay. A level 2 SCAM initiator shall also be able to operate with a selection time-out greater than the SCAM unassigned ID selection response delay.
- f) shall not assert the RST signal upon a selection time-out.
- g) shall, provided an assigned or current ID is available, satisfy the requirements for a SCAM tolerant device.

NOTE 52 - A level 2 SCAM initiator without a current ID may receive an assigned ID by one of two methods: either it assigns itself an ID or, by means of SCAM protocol functions, is assigned an ID by another SCAM initiator. A level 2 SCAM initiator that has a current ID may receive an assigned ID by either of these two methods or its current ID may become its assigned ID if a selection indication for the SCAM initiator's current ID if a SELECTION phase for the current ID is continuously valid for at least a SCAM unassigned ID selection response delay.

- h) shall, once SCAM selection response is enabled and provided that its device ID is unassigned, recognize and respond to SCAM selection within a SCAM selection response time.
- i) shall not, while its ID remains unassigned, start a SELECTION phase unless the SEL signal and the SCSI ID bit that encodes the unassigned ID are true and the BSY and I/O signals are false for at least a SCAM unassigned ID selection response delay. The end of the SELECTION phase shall cause the device to have an assigned ID equal to its current ID.

B.3.3.5 Level 2 SCAM target

A level 2 SCAM target:

- a) shall recognize reset at all times.

NOTE 53 - SCAM implementations, whether in firmware or hardware, are expected to monitor the RST signal even while engaged in SCAM protocol.

- b) shall enable its response to SCAM selection within a SCAM power-on to SCAM selection delay after the device is powered-on.
- c) shall enable its response to SCAM selection within a SCAM reset to SCAM selection delay after a

reset.

d) shall, once selection response is enabled and provided that the device ID is unassigned, recognize and respond to SCAM selection within a SCAM selection response time.

e) shall not, while its ID remains unassigned, start a SELECTION phase unless the SEL signal and the SCSI ID bit that encodes the unassigned ID are true and the BSY and I/O signals are false for at least a SCAM unassigned ID selection response delay. Response to selection shall cause the device to have an assigned ID equal to its current ID.

f) shall, once assigned an ID, behave as a SCAM tolerant device until a subsequent power-on or reset.

NOTE 54 - SCAM devices with assigned ID's neither recognize, respond to nor initiate SCAM selection.

g) shall not assert the RST signal upon a selection time-out.

h) shall not implement the soft reset alternative as defined in SCSI-2.

i) shall be capable of arbitration without an ID. Subsequent to power-on, a level 2 SCAM target shall initiate SCAM protocol provided that the device does not have an assigned ID and no reset has occurred.

B.4 SCAM protocol

SCAM is a distributed algorithm collectively executed by a group of participating SCAM devices. The communication is accomplished by shared (wired-OR) SCSI bus signals that may be asserted or released by the SCAM devices, but which shall not be negated by any participating device. Any SCAM device which is capable of active negation of SCSI bus signals shall disable active negation during SCAM protocol.

B.4.1 Initiation

A device initiates the SCAM protocol by first winning bus arbitration, then performing SCAM selection. The device may arbitrate using its current ID or it may arbitrate without an ID. After winning arbitration the device has the BSY and SEL signals asserted. It shall release all DATA BUS signals and assert the MSG signal, then wait at least two system deskew delays and release the BSY signal. It shall maintain this pattern of the SEL and MSG signals asserted with the BSY signal released for a minimum of a recommended SCAM selection response time, then release the MSG signal. After releasing the MSG signal the device shall wait, using wired-OR glitch filtering (see 7.2.3 and 7.3.2), until the MSG signal has been released by all other devices.

Level 2 SCAM initiators and SCAM targets that have not yet been assigned an ID recognize SCAM selection if a pattern of the SEL and MSG true and the BSY signal false is detected. After a variable delay, devices responding to SCAM selection release the MSG signal, then wait, using wired-OR glitch filtering, until the MSG signal has been released by all devices. SCAM targets should release the MSG signal quickly, perhaps never asserting it at all. SCAM initiators should wait a SCAM selection response time before releasing the MSG signal.

After wired-OR glitch filtering is used to detect the MSG signal false, each SCAM device asserts the BSY signal, waits at least two system deskew delays, then asserts several other signals. SCAM initiators assert the BSY signal followed by the I/O, C/D, DB(6) and DB(7) signals. SCAM targets assert the BSY signal followed by the I/O, DB(6) and DB(7) signals. After asserting its signals each device waits at least two more system deskew delays, then releases the SEL signal and waits, using wired-OR glitch filtering, until the SEL signal has been released by all devices.

After detecting that the SEL signal has been released by all devices, SCAM devices release the DB(6) signal and examine the bus signals. If the C/D signal is false, then there are no SCAM initiators participating and SCAM targets shall release all signals. If the C/D signal is true, each SCAM device waits, using wired-OR glitch filtering, for the DB(6) signal to be released by all devices and then asserts the SEL signal. Initiation of the SCAM protocol is complete after the SEL signal has been asserted.


B.4.1.1 Transfer cycles

The SCAM protocol functions through sequences of transfer cycles. During each transfer cycle certain devices broadcast data to all participating SCAM devices. The actual data received is the logical-OR of the data broadcast by all the sending devices. Each transfer cycle is fully interlocked in the same sense that

asynchronous data transfers are interlocked. Completion of each step of the transfer is explicitly acknowledged, and the transfer rate adapts automatically to the μ A of the SCAM devices involved.

Transfer cycles use the DB(7-5) signals as handshake lines and the DB(4-0) signals as data lines. At the beginning and end of each transfer cycle the DB(7) signal is asserted while the DB(6) and DB(5) signals are released. As shown in figure B.1 each device repeats the following steps for each transfer cycle:

- 1) Assert data on the DB(4-0) signals, if the device is broadcasting data. Devices that have no data to broadcast release these signals. All devices assert the DB(5) signal.
- 2) All devices release the DB(7) signal.
- 3) Wait, using wired-OR glitch filtering, until the DB(7) signal is released by all devices.
- 4) Read and latch data from the DB(4-0) signals. All devices assert the DB(6) signal.
- 5) All devices release the DB(5) signal.
- 6) Wait, using wired-OR glitch filtering, until the DB(5) signal is released by all devices.
- 7) Release or change the DB(4-0) signals. All devices assert the DB(7) signal.
- 8) All devices release the DB(6) signal.
- 9) Wait, using wired-OR glitch filtering, until the DB(6) signal is released by all devices.



Note - Signals are shown asserted low.

Figure B.1 - Transfer cycles

The SCAM protocol continues through successive transfer cycles until the dominant SCAM initiator chooses to terminate it by releasing the C/D signal and all other signals. SCAM targets shall note the release of the C/D signal and release all signals.

B.4.1.2 Wired-OR glitch filtering

Many of the SCSI signals used by SCAM protocol are asserted by more than one SCAM device. Consequently, when one of these signals is released by a SCAM device there may be a transient period in which the signal is observed to be false even though it is still asserted by other SCAM devices. In order to eliminate these incorrect observations (and the consequent malfunction of the SCAM protocol), all SCAM devices shall perform wired-OR glitch filtering on the shared SCSI signals.

Wired-OR glitch filters may be designed into the hardware. In this case the hardware shall be capable of detecting that a signal has remained continuously false for at least a bus settle delay.

NOTE 55 - This continuous observation of a signal requires dedicated hardware; it is not able to be performed by software alone.

As an alternative, SCAM devices may implement wired-OR glitch filtering in software if the procedure described below is used. The algorithm used relies on the fact that a single device that releases a wired-OR signal may cause the signal to be incorrectly observed false for a maximum of a bus settle delay.

- 1) Determine the iteration count to be used in the software polling loop. This is typically 8, 16 or 32 if the SCAM device has knowledge (by means beyond the scope of this annex) that the SCSI bus supports the maximum number of devices. SCAM devices that have no means to determine the width of the SCSI bus should use an iteration count of 32.
- 2) If the signal is observed to be true, reset the iteration count to the initial value determined above. If the signal is observed to be false, decrement the iteration count. If the iteration count is zero, the signal has been released by all devices.
- 3) Wait sufficient time to guarantee that at least a bus settle delay elapses. Continue with the preceding step.

As an alternative to waiting at least a bus settle delay between samples, the implementor may wish to calculate the iteration count as follows. If the minimum sample interval is known, calculate N equal to a bus settle delay divided by the minimum sample interval. Round the number obtained up to the next higher integer. If the iteration count is set to N times the maximum number of devices on the SCSI bus, the algorithm above may be used without the need to wait between successive samples.

B.4.1.3 Isolation stage

Many SCAM function sequences require an isolation stage, which is used to isolate or identify an individual device to perform some action. During an isolation stage each participating device sends an identification string bit serially. As it sends its identification string, each participating device compares its own identification string with the strings of other devices. If a device observes a numerically higher string than its own identification string, it defers for the remainder of the function sequence and participates in subsequent function sequences only after a synchronization pattern is observed. After the isolation stage completes, only the device with the numerically highest identification string is still participating, and that device performs whatever action is specified by the function code (or subsequent action code). Identification strings are sent starting with the most significant bit of the most significant byte and ending with the least significant bit of the least significant byte.

During each transfer cycle of an isolation stage, the devices that are still participating assert the DB(0) signal if the next bit of their identification string is zero, the DB(1) signal if the next bit of their identification string is one, or release the DB(4-0) signals if they have reached the end of their identification string. SCAM initiators may assert the DB(4) signal to terminate the isolation stage prematurely. Each participating device reads the data transferred during each transfer cycle and acts on the conditions defined in table B.2.