

# **Working Draft American National Standard**

# **T10 Project 1525D**

Revision 6  
16-Jan.-2003

---

## **Information technology - SCSI Parallel Interface-5 (SPI-5)**

This is a draft proposed American National Standard of Accredited Standards Committee INCITS. As such this is not a completed standard. The T10 Technical Committee may modify this document as a result of comments received during public review and its approval as a standard. Use of the information contained here in is at your own risk.

Permission is granted to members of INCITS, its technical committees, and their associated task groups to reproduce this document for the purposes of INCITS standardization activities without further permission, provided this notice is included. All other rights are reserved. Any commercial or for-profit duplication is strictly prohibited.

T10 Technical Editor:      George O. Penokie  
                                 IBM/Tivoli  
                                 MS 2C6  
                                 3605 Highway 52 N.  
                                 Rochester, MN 55901  
                                 USA  
  
                                 Telephone: 1-507-253-5208  
                                 Facsimile: 1-507-253-2880  
                                 Email:    gop@us.ibm.com

---

Reference number  
ISO/IEC \*\*\*\*\* : 200x  
ANSI NCITS.    - 200x

Printed Thursday, January 16, 2003 5:13 pm

## POINTS OF CONTACT:

### **T10 Chair**

John B. Lohmeyer  
LSI  
4420 ArrowsWest Drive  
Colo Spgs, CO 80907-3444  
Tel: 1-719-533-7560  
Fax: 1-719-593-7183  
Email: lohmeier@t10.org

### **T10 Vice-Chair**

George O. Penokie  
IBM/Tivoli  
3605 Highway 52 N. MS 2C6  
Rochester, MN 55901  
Tel: 1-507-253-5208  
Fax: 1-507-253-2880  
Email: gop@us.ibm.com

### **INCITS Secretariat**

INCITS Secretariat  
1250 Eye Street, NW Suite 200  
Washington, DC 20005

Telephone: 1-202-737-8888  
Facsimile: 1-202-638-4922  
Email: incits@itic.org

**T10 Web Site** <http://www.t10.org/>

**T10 Reflector** Internet address for subscription to the T10 reflector: majordomo@t10.org  
The message body should be: subscribe t10  
Internet address for distribution via T10 reflector: T10@t10.org  
Internet address to unsubscribe from the T10 reflector: majordomo@t10.org  
The message body should be: unsubscribe t10

### **Document Distribution**

INCITS Online Store	<a href="http://www.techstreet.com/incits.html">http://www.techstreet.com/incits.html</a>
managed by Techstreet	Telephone: 1-734-302-7801
1327 Jones Drive	or 1-800-699-9277
Ann Arbor, MI 48105	Facsimile: 1-734-302-7811

or

Global Engineering	<a href="http://global.ihs.com/">http://global.ihs.com/</a>
15 Inverness Way East	Telephone: 1-303-792-2181
Englewood, CO 80112-5704	or 1-800-854-7179
	Facsimile: 1-303-792-2192

American National Standard  
for Information Technology

## **SCSI Parallel Interface-5 (SPI-5)**

Secretariat  
Information Technology Industry Council

Approved mm.dd.yy

American National Standards Institute, Inc.

### **ABSTRACT**

This standard defines mechanical, electrical, timing requirements, command, and task management delivery protocol requirements to transfer commands and data between SCSI devices attached to a SCSI parallel interface. This standard is intended to be used in conjunction with the SCSI command sets. 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.

## American National Standard

Approval of an American National Standard requires verification by ANSI that the requirements for due process, consensus, and other criteria for approval have been met by the standards developer. Consensus is established when, in the judgment of the ANSI Board of Standards Review, substantial agreement has been reached by directly and materially affected interests. Substantial agreement means much more than a simple majority, but not necessarily unanimity. Consensus requires that all views and objections be considered, and that effort be made towards their resolution.

The use of American National Standards is completely voluntary; their existence does not in any respect preclude anyone, whether he has approved the standards or not, from manufacturing, marketing, purchasing, or using products, processes, or procedures not conforming to the standards.

The American National Standards Institute does not develop standards and will in no circumstances give interpretation on any American National Standard. Moreover, no person shall have the right or authority to issue an interpretation of an American National Standard in the name of the American National Standards Institute. Requests for interpretations should be addressed to the secretariat or sponsor whose name appears on the title page of this standard.

**CAUTION NOTICE:** This American National Standard may be revised or withdrawn at any time. The procedures of the American National Standards Institute require that action be taken periodically to reaffirm, revise, or withdraw this standard. Purchasers of American National Standards may receive current information on all standards by calling or writing the American National Standards Institute.

**CAUTION:** The developers of this standard have requested that holders of patents that may be required for the implementation of the standard, disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this standard. As of the date of publication of this standard, following calls for the identification of patents that may be required for the implementation of the standard, no such claims have been made. No further patent search is conducted by the developer or the publisher in respect to any standard it processes. No representation is made or implied that licenses are not required to avoid infringement in the use of this standard.

Published by  
**American National Standards Institute**  
**11 W. 42nd Street, New York, New York 10036**

Copyright © 2002 by Information Technology Industry Council (ITI)  
All rights reserved.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without prior written permission of ITI, 1250 Eye Street NW, Suite 200, Washington, DC 20005.

Printed in the United States of America

## Contents

	Page
Foreword (This foreword is not part of this standard) .....	xxi
Introduction .....	xxiii
1 Scope .....	1
2 Normative references .....	3
2.1 Normative references .....	3
2.2 Approved references .....	3
2.3 References under development .....	4
2.4 Other references .....	4
3 Definitions, symbols, abbreviations, and conventions .....	4
3.1 Definitions .....	4
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 General overview .....	14
4.2 Cables, connectors, signals, transceivers .....	14
4.3 Physical architecture of bus segment .....	14
4.4 Driver-receiver connections .....	15
4.5 Physical topology details and definitions .....	15
4.6 Bus segment loading .....	17
4.7 Termination requirements .....	18
4.8 SCSI device Addressing .....	18
4.9 Clocking methods for data transfers .....	18
4.10 Paced transfer on a SCSI bus .....	21
4.11 Data transfers .....	23
4.11.1 Data transfer modes .....	23
4.11.1.1 Asynchronous transfers .....	23
4.11.1.2 Synchronous transfers .....	23
4.11.1.3 Paced transfers .....	23
4.11.2 ST DATA phase parallel transfers .....	23
4.11.3 DT DATA phase parallel transfers .....	24
4.11.3.1 DT DATA phase parallel transfers format .....	24
4.11.3.2 Data group transfers .....	24
4.11.3.3 Information unit transfers .....	24
4.12 Negotiation .....	25
4.12.1 Negotiation introduction .....	25
4.12.2 Negotiation algorithm .....	25
4.12.3 When to negotiate .....	26
4.12.4 Negotiable fields .....	27
4.12.4.1 Negotiable fields introduction .....	27
4.12.4.2 Transfer agreements .....	28
4.12.4.3 Transfer period factor .....	29
4.12.4.4 REQ/ACK offset .....	31
4.12.4.5 Transfer width exponent .....	32
4.12.4.6 Protocol options .....	32
4.12.4.6.1 Protocol options introduction .....	32
4.12.4.6.2 IU_REQ .....	33
4.12.4.6.3 DT_REQ .....	34
4.12.4.6.4 QAS_REQ .....	34
4.12.4.6.5 HOLD_MCS .....	34

4.12.4.6.6 WR_FLOW .....	35
4.12.4.6.7 RD_STRM .....	35
4.12.4.6.8 RTI .....	35
4.12.4.6.9 PCOMP_EN .....	35
4.12.5 Negotiable field combinations .....	36
4.12.6 Message restrictions .....	38
4.12.7 Negotiation message sequences .....	38
4.12.7.1 Negotiation message sequences overview .....	38
4.12.7.2 SCSI initiator port originated PPR negotiation .....	39
4.12.7.3 SCSI initiator port originated WDTR negotiation .....	41
4.12.7.4 SCSI initiator port originated SDTR negotiation .....	43
4.12.7.5 SCSI target port originated WDTR negotiation .....	45
4.12.7.6 SCSI target port originated SDTR negotiation .....	47
4.13 Protocol .....	49
5 SCSI parallel interface connectors .....	52
5.1 SCSI parallel interface connectors overview .....	52
5.2 Nonshielded connector .....	52
5.2.1 Nonshielded connector alternative 1 - A cable .....	52
5.2.2 Nonshielded connector alternative 2 - A cable .....	52
5.2.3 Nonshielded connector alternative 3 - P cable .....	53
5.2.4 Nonshielded connector alternative 4 .....	53
5.3 Shielded connector .....	59
5.3.1 Shielded connector overview .....	59
5.3.2 Shielded connector alternative 1 - A cable .....	59
5.3.3 Shielded connector alternative 2 - A cable .....	59
5.3.4 Shielded connector alternative 3 - P cable .....	59
5.3.5 Shielded connector alternative 4 - P cable .....	59
5.4 Connector contact assignments .....	67
5.4.1 Connector contact assignments overview .....	67
5.4.2 Differential connector contact assignments .....	68
6 SCSI bus segment interconnect .....	71
6.1 SCSI bus segment interconnect overview .....	71
6.2 SCSI bus segment cables .....	71
6.3 Interconnect characteristics for signals .....	72
6.3.1 SCSI bulk cable .....	72
6.3.2 Minimum conductor size for signals .....	73
6.3.3 Local transmission line impedance .....	73
6.3.4 Extended distance transmission line impedance .....	73
6.3.5 Capacitance .....	73
6.3.6 Differential propagation time and propagation time skew .....	73
6.3.7 Differential attenuation .....	73
6.3.8 Crosstalk .....	74
6.4 Decoupling characteristics for TERMPWR lines .....	74
6.5 Connection requirements for RESERVED lines .....	75
6.6 Cables used with LVD transceivers .....	75
6.7 LVD stub length and spacing .....	76
7 SCSI parallel interface electrical characteristics .....	77
7.1 SCSI parallel interface electrical characteristics overview .....	77
7.2 LVD termination .....	77
7.2.1 LVD termination overview .....	77
7.2.2 LVD driver characteristics .....	82
7.2.3 LVD receiver characteristics .....	85
7.2.4 LVD capacitive loads .....	86
7.2.4.1 Management of LVD release glitches .....	88

7.2.5 SE/HVD transmission mode detection .....	90
7.2.5.1 SE/HVD transmission mode detection overview .....	90
7.2.5.2 LVD DIFFSENS driver .....	90
7.2.5.3 LVD DIFFSENS receiver .....	91
7.3 Terminator power .....	93
8 SCSI bus signals .....	95
8.1 SCSI bus signals overview .....	95
8.2 Signal descriptions .....	95
8.3 LVD Signal states .....	98
8.4 OR-tied signals .....	99
8.5 Signal sources .....	99
9 SCSI parallel bus timing .....	101
9.1 SCSI parallel bus timing values .....	101
9.2 Timing description .....	107
9.2.1 Arbitration delay .....	107
9.2.2 ATN transmit setup time .....	108
9.2.3 ATN receive setup time .....	108
9.2.4 Bus clear delay .....	108
9.2.5 Bus free delay .....	108
9.2.6 Bus set delay .....	108
9.2.7 Bus settle delay .....	109
9.2.8 Cable skew .....	109
9.2.9 Chip noise in receiver .....	109
9.2.10 Clock jitter .....	109
9.2.11 Crosstalk time shift .....	109
9.2.12 Flow control receive hold time .....	109
9.2.13 Flow control receive setup time .....	109
9.2.14 Flow control transmit hold time .....	109
9.2.15 Flow control transmit setup time .....	109
9.2.16 Offset induced time asymmetry .....	110
9.2.17 pCRC receive hold time .....	110
9.2.18 pCRC receive setup time .....	110
9.2.19 pCRC transmit hold time .....	110
9.2.20 pCRC transmit setup time .....	110
9.2.21 Data release delay .....	110
9.2.22 DIFFSENS voltage filter time .....	110
9.2.23 Physical disconnection delay .....	110
9.2.24 Power on to selection .....	110
9.2.25 QAS arbitration delay .....	111
9.2.26 QAS assertion delay .....	111
9.2.27 QAS release delay .....	111
9.2.28 QAS non-DATA phase REQ(ACK) period .....	111
9.2.29 Receive assertion period .....	111
9.2.30 Receive hold time .....	111
9.2.31 Receive internal hold time .....	111
9.2.32 Receive internal setup time .....	111
9.2.33 Receive negation period .....	112
9.2.34 Receive setup time .....	112
9.2.35 Receive REQ(ACK) period tolerance .....	112
9.2.36 Receive REQ assertion period with P_CRCA transitioning .....	112
9.2.37 Receive REQ negation period with P_CRCA transitioning .....	112
9.2.38 Receive Skew Compensation .....	112
9.2.39 Receiver amplitude time skew .....	112
9.2.40 Receiver de-skewed data valid window .....	112
9.2.41 REQ(ACK) period .....	113

9.2.42 Reset delay .....	113
9.2.43 Reset hold time .....	113
9.2.44 Reset to selection .....	113
9.2.45 Residual skew error .....	114
9.2.46 Selection abort time .....	114
9.2.47 Selection time-out delay .....	114
9.2.48 Signal timing skew .....	114
9.2.49 Skew correction range .....	114
9.2.50 Strobe offset tolerance .....	115
9.2.51 System deskew delay .....	115
9.2.52 System noise at launch .....	115
9.2.53 System noise at receiver .....	115
9.2.54 Time asymmetry .....	115
9.2.55 Transmit assertion period .....	115
9.2.56 Transmit hold time .....	116
9.2.57 Transmit ISI Compensation .....	116
9.2.58 Transmit negation period .....	116
9.2.59 Transmit setup time .....	116
9.2.60 Transmit REQ(ACK) period tolerance .....	116
9.2.61 Transmit REQ assertion period with P_CRCA transitioning .....	116
9.2.62 Transmit REQ negation period with P_CRCA transitioning .....	116
9.2.63 Transmitter skew .....	116
9.2.64 Transmitter time asymmetry .....	117
9.3 Measurement points .....	117
9.3.1 Measurement points overview .....	117
9.3.2 LVD measurement points .....	118
9.4 Receiver Masks .....	120
9.4.1 Synchronous transfers and for fast-160 paced transfer clocking signals .....	120
9.4.2 Paced transfers with precompensation enabled on fast-160 .....	121
9.4.3 Paced transfers with precompensation disabled on fast-160 and fast-320 .....	123
9.5 Timing parameters .....	129
9.6 Setup and hold timings .....	129
9.6.1 ST data transfer calculations .....	129
9.6.2 DT data transfer calculations .....	130
10 SCSI bus phases .....	132
10.1 SCSI bus phases overview .....	132
10.2 BUS FREE phase .....	132
10.3 Expected and unexpected bus free phases .....	132
10.4 Arbitration .....	133
10.4.1 Arbitration and QAS overview .....	133
10.4.2 NORMAL ARBITRATION phase .....	133
10.4.3 QAS protocol .....	134
10.4.4 QAS phase .....	135
10.5 SELECTION phase .....	136
10.5.1 Selection Overview .....	136
10.5.2 Selection using attention condition .....	137
10.5.2.1 Starting the SELECTION phase when using attention condition .....	137
10.5.2.2 Information unit transfers disabled .....	137
10.5.2.3 Information unit transfers enabled .....	137
10.5.2.4 Selection using attention condition time-out procedure .....	137
10.5.3 Selection without using attention condition .....	138
10.5.3.1 Information unit transfers disabled or enabled .....	138
10.5.3.2 Selection without using attention condition time-out procedure .....	138
10.6 RESELECTION phase .....	139
10.6.1 RESELECTION phase overview .....	139
10.6.2 Physical reconnection .....	139



10.6.3 Physical reconnection time-out procedure .....	140
10.7 Information transfer phases .....	140
10.7.1 Information transfer phases overview .....	140
10.7.2 Asynchronous transfer .....	141
10.7.3 Synchronous transfer .....	142
10.7.3.1 Synchronous transfer overview .....	142
10.7.3.2 ST synchronous transfer .....	142
10.7.3.3 DT synchronous transfer .....	143
10.7.3.3.1 DT synchronous transfer overview .....	143
10.7.3.3.2 Information unit transfer .....	143
10.7.3.3.3 DT DATA IN phase information unit transfer exception condition handling .....	144
10.7.3.3.4 DT DATA OUT phase information unit transfer exception condition handling .....	145
10.7.3.3.5 Data group data field transfer .....	146
10.7.3.3.6 Data Group Pad field and pCRC field transfer to SCSI initiator port .....	146
10.7.3.3.7 Data Group Pad field and pCRC field transfer to SCSI target port .....	148
10.7.4 Paced transfer .....	149
10.7.4.1 Paced transfer overview .....	149
10.7.4.2 Paced transfer training pattern .....	150
10.7.4.2.1 Training pattern overview .....	150
10.7.4.2.2 DT DATA IN phase training pattern .....	151
10.7.4.2.3 DT DATA OUT phase training pattern .....	152
10.7.4.3 P1 data valid/invalid state transitions .....	153
10.7.4.3.1 P1 data valid/invalid state transitions overview .....	153
10.7.4.3.2 Starting pacing transfers at end of training pattern .....	154
10.7.4.3.3 Starting pacing transfers with no training pattern .....	154
10.7.4.3.4 Ending pacing transfers .....	155
10.7.4.4 Paced information unit transfer .....	155
10.7.4.5 Deskewing .....	156
10.7.5 Wide transfer .....	156
10.8 COMMAND phase .....	157
10.8.1 COMMAND phase description .....	157
10.8.2 COMMAND phase exception condition handling .....	157
10.9 DATA phase .....	158
10.9.1 DATA phase overview .....	158
10.9.2 DT DATA IN phase .....	158
10.9.3 DT DATA OUT phase .....	158
10.9.4 ST DATA IN phase .....	158
10.9.5 ST DATA OUT phase .....	158
10.10 STATUS phase .....	158
10.10.1 STATUS phase description .....	158
10.10.2 STATUS phase exception condition handling .....	159
10.11 MESSAGE phase .....	159
10.11.1 MESSAGE phase overview .....	159
10.11.2 MESSAGE IN phase .....	159
10.11.3 MESSAGE IN phase exception condition handling .....	159
10.11.4 MESSAGE OUT phase .....	159
10.11.5 MESSAGE OUT phase exception condition handling .....	160
10.12 Signal restrictions between phases .....	160
11 DATA BUS protection .....	162
11.1 DATA BUS protection overview .....	162
11.2 ST DATA BUS protection using parity .....	162
11.3 DT DATA BUS protection using CRC .....	162
11.3.1 DT DATA BUS protection using CRC overview .....	162
11.3.2 Error detection capabilities .....	162
11.3.3 Order of bytes in the CRC field .....	163
11.3.4 CRC generation and checking .....	165

11.3.5 Test cases .....	165
12 SCSI bus conditions .....	166
12.1 SCSI bus conditions overview .....	166
12.2 Attention condition .....	166
12.3 Bus reset condition .....	167
12.4 Hard reset .....	167
12.5 Reset events .....	168
12.5.1 Reset events overview .....	168
12.5.2 Bus reset event .....	168
12.5.3 Power on reset event .....	168
12.5.4 Target reset event .....	168
13 SCSI bus phase sequences .....	169
13.1 SCSI bus phase sequences overview .....	169
13.2 Phase sequences with information units disabled .....	170
13.2.1 Phase sequences for physical reconnection or selection using attention condition .....	170
13.2.2 Phase sequences for selection without using attention condition .....	171
13.3 Phase sequences with information unit enabled .....	172
13.3.1 Phase sequences for physical reconnection or selection without using attention condition ..	172
13.3.2 Phase sequences for selection using attention condition .....	173
14 SPI information units .....	174
14.1 SPI information unit overview .....	174
14.2 Information unit transfer logical operations .....	174
14.3 SPI information units .....	180
14.3.1 SPI command information unit .....	180
14.3.2 SPI L_Q information unit .....	184
14.3.3 SPI data information unit .....	187
14.3.4 SPI data stream information unit .....	188
14.3.5 SPI status information unit .....	189
15 SCSI pointers .....	193
16 SCSI messages .....	194
16.1 SCSI messages overview .....	194
16.2 Message protocols and formats .....	194
16.2.1 Message protocol rules .....	194
16.2.2 Message formats .....	194
16.2.3 One-byte messages .....	195
16.2.4 Two-byte messages .....	195
16.2.5 Extended messages .....	195
16.3 Link control messages .....	197
16.3.1 Link control message codes .....	197
16.3.2 DISCONNECT .....	198
16.3.3 IDENTIFY .....	198
16.3.4 IGNORE WIDE RESIDUE .....	199
16.3.5 INITIATOR DETECTED ERROR .....	200
16.3.6 LINKED COMMAND COMPLETE .....	200
16.3.7 MESSAGE PARITY ERROR .....	200
16.3.8 MESSAGE REJECT .....	200
16.3.9 MODIFY DATA POINTER .....	201
16.3.10 MODIFY BIDIRECTIONAL DATA POINTER .....	202
16.3.11 NO OPERATION .....	203
16.3.12 PARALLEL PROTOCOL REQUEST .....	203
16.3.13 QAS REQUEST .....	203
16.3.14 RESTORE POINTERS .....	204

16.3.15 SAVE DATA POINTERS .....	204
16.3.16 SYNCHRONOUS DATA TRANSFER REQUEST .....	204
16.3.17 TASK COMPLETE .....	204
16.3.18 WIDE DATA TRANSFER REQUEST .....	205
16.4 Task attribute messages .....	205
16.4.1 Task attribute message overview and codes .....	205
16.4.2 ACA .....	206
16.4.3 HEAD OF QUEUE .....	207
16.4.4 ORDERED .....	207
16.4.5 SIMPLE .....	208
16.5 Task management messages .....	209
16.5.1 Task management message codes .....	209
16.5.2 ABORT TASK .....	209
16.5.3 ABORT TASK SET .....	210
16.5.4 CLEAR ACA .....	210
16.5.5 CLEAR TASK SET .....	210
16.5.6 LOGICAL UNIT RESET .....	210
16.5.7 TARGET RESET .....	211
17 Command processing considerations and exception conditions .....	212
17.1 Command processing considerations and exception conditions overview .....	212
17.2 Asynchronous event notification .....	212
17.3 Incorrect initiator connection .....	212
17.4 Unexpected RESELECTION phase .....	213
18 SCSI management features for the SCSI parallel interface .....	214
18.1 SCSI mode parameters .....	214
18.1.1 SCSI mode parameter overview and codes .....	214
18.1.2 Disconnect-Reconnect mode page .....	215
18.1.3 Logical Unit Control mode page .....	218
18.1.4 Port Control mode page .....	219
18.1.4.1 Port Control mode page overview .....	219
18.1.4.2 Margin Control mode subpage .....	221
18.1.4.3 Saved Training Configuration Values mode subpage .....	223
18.1.4.4 Negotiated Settings mode subpage .....	225
18.1.4.5 Report Transfer Capabilities mode subpage .....	226
19 SCSI parallel interface services .....	227
19.1 SCSI parallel interface services overview .....	227
19.2 Procedure objects .....	227
19.3 Application client SCSI command services .....	228
19.3.1 Application client SCSI command services overview .....	228
19.3.2 Send SCSI command service .....	229
19.4 Device server SCSI command services .....	229
19.4.1 Device server SCSI command services overview .....	229
19.4.2 Data-in delivery service .....	229
19.4.3 Data-out delivery service .....	230
19.5 Task management services .....	230
19.5.1 Task management functions overview .....	230
19.5.2 Task management functions .....	230
19.5.3 ABORT TASK .....	231
19.5.4 ABORT TASK SET .....	231
19.5.5 CLEAR ACA .....	231
19.5.6 CLEAR TASK SET .....	231
19.5.7 LOGICAL UNIT RESET .....	231
19.5.8 RESET SERVICE DELIVERY SUBSYSTEM .....	231
19.5.9 TARGET RESET .....	231

19.5.10 WAKEUP .....	231
Annex A Additional requirements for LVD SCSI drivers and receivers .....	232
A.1 System level requirements .....	232
A.2 Driver requirements .....	234
A.2.1 Driver requirements overview .....	234
A.2.2 Differential output voltage, VS .....	234
A.2.3 Offset (common-mode output) voltage (VCM) .....	238
A.2.4 Short-circuit currents (IO-S and IO+S) .....	239
A.2.5 Open-circuit output voltages ( $V_{O-(OC)}$ and $V_{O+(OC)}$ ) .....	240
A.2.6 Output signal waveform .....	241
A.2.7 Dynamic output signal balance (VCM(PP)) .....	242
A.3 Receiver characteristics .....	244
A.3.1 Receiver characteristics overview .....	244
A.3.2 Receiver steady state input voltage requirements .....	244
A.3.2.1 Steady state input voltage requirements for fast-160 and slower receivers .....	244
A.3.2.2 Steady state input voltage requirements for fast-320 receivers .....	244
A.3.3 Compliance test .....	245
A.3.4 Receiver setup and hold times .....	246
A.3.5 Fast-160 receiver bandwidth specifications .....	246
A.4 Transceiver characteristics .....	247
A.4.1 Transceiver output/input currents, II-L and II+L .....	248
A.4.2 Transceiver maximum input voltages .....	248
Annex B SCSI bus fairness .....	249
B.1 Model .....	249
B.2 Determining fairness by monitoring prior bus activity .....	249
B.2.1 Fairness for normal arbitration method .....	249
B.2.2 Fairness for QAS .....	250
B.3 Fairness algorithm .....	250
B.3.1 Fairness states overview .....	250
B.3.2 Fairness idle state .....	250
B.3.3 Fairness wait state .....	250
B.3.4 Fairness participate state .....	250
B.3.5 Lockout delay .....	250
B.3.6 Mixed arbitration .....	251
B.4 SCSI initiator port fairness recommendations .....	251
Annex C Nonshielded connector alternative 4 .....	252
C.1 Nonshielded connector alternative 4 signal definitions .....	252
C.2 VOLTAGE and GROUND signals .....	252
C.3 CHARGE signals .....	253
C.4 SPINDLE SYNC .....	253
C.5 ACTIVE LED OUT .....	254
C.6 Motor start controls .....	254
C.7 SCSI ID selection .....	255
C.8 MATED signals .....	256
C.8.1 MATED signals overview .....	256
C.8.2 MATED 2/drive side .....	257
C.8.3 MATED 2/backplane side .....	257
C.8.4 MATED 1/drive side .....	257
C.8.5 MATED 1/backplane side .....	257
Annex D Removal and insertion of SCSI devices .....	259

D.1 Removal and insertion of SCSI devices overview .....	259
D.2 Case 1 - Power off during removal or insertion .....	259
D.3 Case 2 - RST signal asserted continuously during removal or insertion .....	259
D.4 Case 3 - Current I/O processes not allowed during insertion or removal .....	259
D.5 Case 4 - Current I/O process allowed during insertion or removal .....	260
Annex E Simple expander requirements .....	261
E.1 Introduction .....	261
E.2 Glossary .....	261
E.3 Bus segments in a SCSI domain .....	261
E.4 Simple bus expanders .....	262
E.5 Homogeneous type .....	263
E.6 Heterogeneous types .....	263
E.7 SCSI domain examples using simple expanders .....	263
E.8 Rules for SCSI domains using simple expanders .....	264
E.8.1 Rule summary .....	264
E.8.2 Rule 1 .....	265
E.8.3 Rule 2 .....	265
E.8.4 Rule 3 .....	266
E.8.5 Rule 4 .....	267
E.8.5.1 Effects of wired-or glitches .....	267
E.8.5.2 Simple expander propagation delay effects .....	267
E.8.5.3 Sample calculations .....	268
E.8.6 Rule 5 .....	269
E.8.7 Rule 6 .....	269
E.9 Special performance considerations for SCSI domains with simple expanders .....	270
Annex F Expander Communication Protocol .....	273
F.1 Introduction .....	273
F.2 Glossary and Definitions .....	273
F.3 Symbols and abbreviations .....	274
F.4 Enabling ECP .....	274
F.5 Communicative expander function structures .....	274
F.6 Expander functions .....	277
F.6.1 Outbound multiple functions .....	277
F.6.1.1 Outbound multiple function data transfer rules .....	277
F.6.1.2 ASSIGN ADDRESS .....	277
F.6.1.3 MARGIN CONTROL .....	278
F.6.2 Outbound single functions .....	279
F.6.2.1 Outbound single function data transfer rules .....	279
F.6.2.2 CONTROL .....	280
F.6.3 Inbound multiple functions .....	281
F.6.3.1 Inbound multiple function data transfer rules .....	281
F.6.3.2 MARGIN REPORT .....	281
F.6.3.3 REPORT CAPABILITIES .....	282
F.6.3.4 REPORT CURRENT STATUS .....	283
F.6.4 Inbound single functions .....	284
F.6.4.1 Inbound single function data transfer rules .....	284
F.6.4.2 EXPANDER INQUIRY .....	285
F.6.4.3 REPORT SAVED TRAINING CONFIGURATION VALUES .....	286
F.7 Data Transfer Requirements .....	288
Annex G Interconnecting bus segments of different widths .....	289
Annex H SCSI ICONS .....	291

Annex I Backplane construction guidelines .....	292
I.1 Universal backplane construction .....	292
I.1.1 Universal backplane construction overview .....	292
I.1.2 Microstrip .....	292
I.1.3 Embedded Microstrip .....	292
I.1.4 Stripline .....	293
I.1.5 Dual Stripline .....	294
I.1.6 Differential Impedance .....	295
I.1.7 Single ended impedance .....	296
I.1.8 Differential stripline .....	296
I.1.9 Dielectric material selection .....	299
I.1.10 Vias .....	299
I.1.11 Transmission Lines .....	301
Annex J SPI-5 to SCSI-2 terminology mapping .....	306

## Tables

	Page
1 LVD transceiver/speed support map .....	14
2 Negotiable fields and effects of successful negotiation .....	27
3 Responding message requirements .....	28
4 Transfer agreements .....	29
5 Transfer Period Factor .....	30
6 Transfer Period Factor relationships .....	31
7 REQ/ACK offset .....	32
8 Transfer Width Exponent .....	32
9 Protocol options bits .....	33
10 Bus phases resulting from iu_req changes .....	34
11 Valid negotiable field combinations .....	37
12 Cross-reference to A cable contact assignments .....	67
13 LVD contact assignments - A cable .....	68
14 LVD contact assignments - P cable .....	69
15 LVD contact assignments - nonshielded alternative 4 connector .....	70
16 LVD local transmission line impedance .....	73
17 Bulk cable capacitance limits .....	73
18 Attenuation requirements for SCSI bulk cable .....	74
19 LVD maximum bus segment path length between terminators .....	76
20 Minimum stub connection spacing for LVD SCSI devices .....	76
21 Absolute electrical limits at the device connector .....	77
22 Input current requirements at the device connector for lines not being driven by the device .....	77
23 I-V requirements for differential impedance, common mode impedance, and $V_{BIAS}$ tests .....	80
24 Values for LVD termination balance test .....	82
25 Fast-160 precompensation .....	83
26 Values for LVD capacitive loads .....	88
27 Glitch management requirements for SCSI devices using LVD drivers .....	89
28 LVD DIFFSENS driver specifications .....	90
29 DIFFSENS receiver operating requirements .....	92
30 Terminator power characteristics at the terminator .....	94
31 Arbitration priorities by SCSI ID .....	95
32 P_CRC signal usage requirements .....	97
33 Signal sources .....	100
34 SCSI bus control timing values .....	101
35 SCSI bus data & information phase ST timing values .....	102
36 Miscellaneous SCSI bus data & information phase DT timing values .....	103
37 Transmit SCSI bus data & information phase DT timing values .....	104
38 Receive SCSI bus data & information phase DT timing values .....	105
39 SCSI fast-160 and fast-320 non-compensatable timing budget .....	106
40 SCSI fast-160 and fast-320 interconnect timing budget .....	107
41 Receiver Eye Mask Values .....	124
42 Information transfer phases .....	141
43 Wide SCSI byte order .....	157
44 Parity checking rules for SELECTION and RESELECTION phases .....	162
45 SPI information units .....	174
46 SPI command information unit .....	181
47 TASK ATTRIBUTE .....	182
48 TASK MANAGEMENT FUNCTIONS .....	183
49 SPI L_Q information unit .....	185
50 TYPE .....	186
51 BIDI DIRECTION .....	187
52 SPI data information unit .....	188
53 SPI data stream information unit .....	189
54 SPI status information unit .....	190
55 PACKETIZED FAILURES field .....	191

56	PACKETIZED FAILURE CODE .....	192
57	Message format .....	195
58	Extended message format .....	196
59	Extended message codes .....	196
60	Link control message codes .....	197
61	IDENTIFY message format .....	198
62	IGNORE WIDE RESIDUE message format .....	199
63	IGNORE field definition .....	200
64	MODIFY DATA POINTER message format .....	201
65	MODIFY BIDIRECTIONAL DATA POINTER message format .....	202
66	PARALLEL PROTOCOL message format .....	203
67	SYNCHRONOUS DATA TRANSFER message format .....	204
68	WIDE DATA TRANSFER message format .....	205
69	Task attribute message codes .....	206
70	ACA message format .....	207
71	HEAD OF QUEUE message format .....	207
72	ORDERED message format .....	207
73	SIMPLE message format .....	208
74	Task management message codes .....	209
75	Mode page codes for the SCSI parallel interface .....	214
76	Mode subpage codes for the SCSI parallel interface .....	214
77	Disconnect-Reconnect mode page .....	216
78	DTDC .....	218
79	Logical Unit Control mode page .....	219
80	Port Control mode page short format .....	220
81	Port Control mode subpage format .....	221
82	Margin Control mode subpage .....	222
83	Summary of margin control field values .....	223
84	Saved Training Configuration mode subpage .....	224
85	Negotiated Settings mode subpage .....	225
86	Bus mode .....	226
87	Report Transfer Capabilities mode subpage .....	226
88	This standards objects mapped to objects from other SCSI standards .....	227
89	Procedure objects .....	228
90	Processing of send SCSI command service procedure .....	229
91	Processing of data-in delivery service procedure .....	230
92	Processing of data-out delivery service procedure .....	230
A.1	System level requirements .....	233
A.2	Driver steady-state test limits and conditions for non-paced transfers .....	235
A.3	Driver steady-state test limits and conditions for paced transfers .....	235
A.4	Driver switching test circuit parameters .....	242
A.5	Dynamic output balance limits .....	243
A.6	Receiver steady state input voltage ranges .....	244
A.7	Fast-320 receiver steady state input voltage ranges .....	245
A.8	Receiver minimum and maximum input voltages. ....	246
A.9	Properties of receiver pulse train .....	247
C.1	Voltage specification limits .....	252
C.2	Charge supply to SCSI device .....	253
C.3	Output characteristics of drive ACTIVE LED OUT signal .....	254
C.4	Definition of motor start controls .....	255
C.5	Electronic requirements for input controls .....	255
C.6	SCSI device ID selection signals .....	256
E.1	SCSI Domain rules .....	265
E.2	SCSI domain delay calculation .....	269
E.3	Minimum REQ/ACK OFFSET for maximum performance .....	272
F.1	Expander function header .....	275
F.2	Expander functions .....	276



F.3 SEDB format .....	276
F.4 Device class .....	277
F.5 ASSIGN ADDRESS SEDB .....	278
F.6 MARGIN CONTROL SEDB .....	279
F.7 CONTROL data structure .....	280
F.8 FAR_CTL field values .....	281
F.9 MARGIN REPORT SEDB .....	282
F.10 REPORT CAPABILITES SEDB .....	283
F.11 REPORT CURRENT STATUS SEDB .....	284
F.12 EXPANDER INQUIRY expander function header .....	285
F.13 EXPANDER INQUIRY data .....	286
F.14 REPORT SAVED TRAINING CONFIGURATION VALUES data structure .....	287
I.1 Dielectric constants .....	299
J.1 SPI-5 to SCSI-2 terminology mapping .....	306

## Figures

	Page
1 SCSI Document Structure .....	2
2 Differential SCSI bus segment .....	15
3 SCSI bus segment topology details .....	17
4 ST latching data vs. DT latching data .....	19
5 ST synchronous transfer example .....	20
6 DT synchronous transfer example .....	20
7 Paced transfer example .....	21
8 Example of a SCSI bus with fast-160 paced transfers .....	22
9 Example of a SCSI bus with fast-320 paced transfers .....	22
10 Skew compensation and clock shift example .....	23
11 Error-free negotiation message sequences. ....	26
12 SCSI initiator port originated PPR negotiation: SCSI initiator port response .....	39
13 SCSI initiator port originated PPR negotiation: SCSI target port response .....	40
14 SCSI initiator port originated WDTR negotiation: SCSI initiator port response .....	41
15 SCSI initiator port originated WDTR negotiation: SCSI target port response .....	42
16 SCSI initiator port originated SDTR negotiation: SCSI initiator port response .....	43
17 SCSI initiator port originated SDTR negotiation: SCSI target port response .....	44
18 SCSI target port originated WDTR negotiation: SCSI target port response .....	45
19 SCSI target port originated WDTR negotiation: SCSI initiator port response .....	46
20 SCSI target port originated SDTR negotiation: SCSI target port response .....	47
21 SCSI target port originated SDTR negotiation: SCSI initiator port response .....	48
22 SCSI Parallel Interface service reference mode .....	49
23 Model for a four step confirmed service .....	50
24 Model for a two step confirmed service .....	50
25 50/68-contact alternative 1/alternative 3 nonshielded SCSI device connector .....	54
26 50/68-contact alternative 1/alternative 3 nonshielded mating connector .....	55
27 50-contact alternative 2 nonshielded SCSI device connector (A cable) .....	56
28 50-contact alternative 2 nonshielded mating connector (A cable) .....	57
29 80-contact alternative 4 nonshielded SCSI device connector (P cable) .....	58
30 80-contact alternative 4 nonshielded contact positions (P cable) .....	58
31 50-contact alternative 1 shielded SCSI device connector (A cable) .....	60
32 50-contact alternative 1 shielded mating connector (A cable) .....	61
33 50-contact alternative 2 shielded SCSI device connector (A cable) .....	62
34 50-contact alternative 2 shielded mating connector (A cable) .....	63
35 68-contact alternative 3 shielded SCSI device connector (P cable) .....	64
36 68-contact alternative 3 shielded mating connector (P cable) .....	65
37 68-contact alternative 4 shielded SCSI device connector (P cable) .....	66
38 68-contact alternative 4 shielded contact positions (P cable) .....	67
39 Terminator decoupling example .....	75
40 LVD bus segment terminator .....	78
41 Test circuit for terminator differential impedance .....	78
42 Termination I-V characteristics for differential and common mode impedance tests .....	79
43 Test circuit for termination common mode impedance test .....	79
44 Termination balance test configuration .....	81
45 Termination balance test data definition .....	81
46 Examples of fast-160 driver precompensation .....	84
47 LVD transceiver architecture .....	85
48 LVD receiver example .....	86
49 LVD Capacitive loads .....	87
50 LVD DIFFSENS driver signal definitions .....	91
51 DIFFSENS receiver function .....	91
52 LVD DIFFSENS receiver example .....	93
53 Voltage and current definitions .....	98
54 LVD Signaling sense .....	99
55 Timing budget item locations .....	107

56 Receiver de-skew parameters .....	113
57 Skew correction range .....	115
58 Transmitter skew .....	117
59 Transmitter time asymmetry .....	117
60 LVD timing measurement points for ST synchronous transfers .....	118
61 LVD timing measurement points for DT synchronous transfers .....	119
62 LVD timing measurement points for DT paced transfers .....	120
63 LVD receiver mask for synchronous transfers and for fast-160 paced transfer clocking signals .....	121
64 Fast-160 LVD receiver mask 1 for paced transfers with precompensation enabled .....	122
65 Fast-160 LVD receiver mask 2 for paced transfers with precompensation enabled .....	123
66 LVD paced transfer receiver mask non-precompensated signals transitioning at the negotiated transfer period .....	125
67 LVD paced transfer receiver mask non-precompensated minimum amplitude change .....	126
68 LVD paced transfer receiver mask non-precompensated non-clocking isolated assertion pulse .....	127
69 LVD paced transfer receiver mask non-precompensated non-clocking isolated negation pulse .....	128
70 LVD paced transfer eye mask non-precompensated non-clocking pulse .....	129
71 System setup and hold timings for ST data transfers .....	130
72 System setup and hold timings for DT synchronous transfers .....	131
73 Usage of P1 to establish data valid and data invalid states .....	154
74 CRC generation and transmission .....	164
75 Phase sequences for physical reconnection or selection using attention condition with information unit transfers disabled .....	170
76 Phase sequences for selection without using attention condition with information unit transfers disabled .....	171
77 Phase sequences for physical reconnection or selection without using attention condition with information unit transfers enabled .....	172
78 Phase sequences for selection with attention condition with information unit transfers enabled .....	173
79 SPI information unit sequence during initial connection .....	176
80 SPI information unit sequence during data type transfers .....	177
81 SPI information unit sequence during data stream type transfers .....	178
82 SPI information unit sequence during status transfers .....	179
A.1 Differential steady-state output voltage test circuit .....	236
A.2 Domain for driver assertion and negation levels for non-paced transfers .....	237
A.3 Domain for driver assertion and negation levels for paced transfers .....	238
A.4 Driver offset steady-state voltage test circuit .....	239
A.5 Common mode output voltage test .....	239
A.6 Driver short-circuit test circuit .....	240
A.7 Open-circuit output voltage test circuit .....	240
A.8 Differential output switching voltage test circuit .....	241
A.9 Driver output signal waveform .....	242
A.10 Driver offset switching voltage test circuit .....	243
A.11 Input signal offset component ( $V_{RXOFFSET}$ ) .....	245
A.12 Receiver input voltage threshold test circuit .....	246
A.13 Fast-160 receiver pulse train example .....	247
A.14 Transceiver off-state output current test circuit .....	248
C.1 Sample circuit for mated indications .....	258
E.1 A two segment domain using a single simple expander circuit .....	262
E.2 Three ways to couple bus segments together with simple expanders .....	263
E.3 Examples of SCSI domains using simple expanders .....	264
E.4 Intermediate bus segments and performance ranking .....	266
E.5 Two configurations for SCSI domain delay calculations .....	268
E.6 Examples of illegal loops .....	270
G.1 Interconnecting LVD A and P cables .....	290
H.1 LVD icon for SCSI .....	291
I.1 Microstrip geometry .....	292
I.2 Embedded microstrip geometry .....	293
I.3 Stripline geometry .....	293

I.4 Dual stripline geometry .....	294
I.5 Lossless model for differential impedance .....	295
I.6 Edge coupled differential microstrip .....	297
I.7 Edge coupled differential stripline .....	297
I.8 Broadside coupled differential stripline .....	297
I.9 Typical transmission line element .....	301

**Foreword (This foreword is not part of this standard)**

Requests for interpretation, suggestions for improvement and addenda, or defect reports are welcome. They should be sent to the INCITS Secretariat, ITI, 1250 Eye Street, NW, Suite 200, Washington, DC 20005-3922.

This standard was processed and approved for submittal to ANSI by InterNational Committee for Information Technology Standards (INCITS). Committee approval of this standard does not necessarily imply that all committee members voted for approval. At the time it approved this standard, INCITS had the following members:

Karen Higginbottom, Chair  
David Michael, Vice-chair  
Monica Vago, Secretary

(INCITS Membership to be inserted)

Technical Committee T10 on Lower Level Interfaces, that developed this standard, had the following members:

John B. Lohmeyer, Chair  
George O. Penokie, Vice-Chair  
Ralph O. Weber, Secretary

Paul D. Aloisi	Terence J. Nelson	Tim Mackley (Alt)
Charles Binford	Robert H. Nixon	Fabio Maino (Alt)
Tim Bradshaw	Vit Novak	John Majernik (Alt)
Joe Breher	Erich Oetting	A. Bruce Manildi (Alt)
Roger Cummings	David Peterson	Ron Martin (Alt)
Zane Daggett	Ron Roberts	David McFadden (Alt)
James R. (Bob) Davis	Gary S. Robinson	Pete McLean (Alt)
Claudio DeSanti	Cris Simpson	Milan J. Merhar (Alt)
Rob Elliott	Douglas Wagner	Carl Mies (Alt)
Terry Enright	Michael Wingard	Richard Moore (Alt)
Paul Entzel	I. Dal Allan (Alt)	Carl Mozingo (Alt)
Mark Evans	Dennis Appleyard (Alt)	Nathan Obr (Alt)
Jie Fan	Dave Baldwin (Alt)	William Petty (Alt)
Mike Fitzpatrick	Michael Banther (Alt)	Sumi Puri (Alt)
Brian Forbes	David Black (Alt)	Charley Riegger (Alt)
Bill Galloway	Craig W. Carlson (Alt)	John P. Scheible (Alt)
Edward A. Gardner	Greg Casey (Alt)	Robert Sheffield (Alt)
Robert Griswold	Joe Dambach (Alt)	Phil Shelton (Alt)
Tom Hanan	Dan Dawiedczyk (Alt)	Robert Snively (Alt)
Nathan Hastad	Lance Flake (Alt)	Ronald Stockford (Alt)
Emily Hill	Chuck Grant (Alt)	Himmeler Themistocle (Alt)
Kenneth Hirata	Randy Haagens (Alt)	Pete Tobias (Alt)
Gerald Houlder	William Ham (Alt)	Joshua Tseng (Alt)
Peter Johansson	Rob Haydt (Alt)	John Tyndall (Alt)
Skip Jones	Steve Hemmah (Alt)	Rudolf Vitti (Alt)
Tasuku Kasebayashi	Randall C. Hines (Alt)	Dean Wallace (Alt)
James A. Lott, Jr.	Bill Hooper (Alt)	Michael Wamsley (Alt)
Kevin Marks	Titkwan Hui (Alt)	Randy Wasylak (Alt)
Ron Mathews	Jim Jones (Alt)	Steve Wong (Alt)
William P. McFerrin	Jerry Kachlic (Alt)	
Charles Monia	Jim Koser (Alt)	
Dennis Moore	Marjorie Krueger (Alt)	
Jay Neer	Ben-Koon Lin (Alt)	

## Introduction

The SCSI parallel interface 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 (i.e., 8 or 16). Data may be transferred asynchronously or synchronously at rates that depend on implementation.

This 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 SCSI parallel bus timing;
- Clause 10 describes the SCSI bus phases;
- Clause 11 describes the DATA BUS protection;
- Clause 12 describes the SCSI bus conditions;
- Clause 13 describes the SCSI bus phase sequences;
- Clause 14 describes the SPI information unit sequences;
- Clause 15 describes the SCSI pointers;
- Clause 16 describes the SCSI messages;
- Clause 17 describes the Command processing considerations and exception conditions;
- Clause 18 describes the SCSI management features for the SCSI parallel interface;
- Clause 19 describes the SCSI parallel interface services;

Annexes A, B, C, D, E, and F form an integral part of this standard. Annexes G through J are for information purposes only.





# American National Standard for Information Technology -

## SCSI Parallel Interface-5 (SPI-5)

### 1 Scope

This standard defines the mechanical, electrical, timing, and protocol requirements of the SCSI parallel interface to allow conforming SCSI devices to inter-operate. 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 SCSI 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 conforming SPI-2, SPI-3, SPI-4 devices may interoperate with SPI-5 devices given that the systems engineering is correctly done. Conforming SPI-2, SPI-3, and SPI-5 devices should respond in an acceptable manner to reject SPI-5 protocol extensions. SPI-5 protocol extensions are designed to be permissive of such rejections and thus allow SPI-2, SPI-3, and SPI-4 devices to continue operation without requiring the use of the extensions.

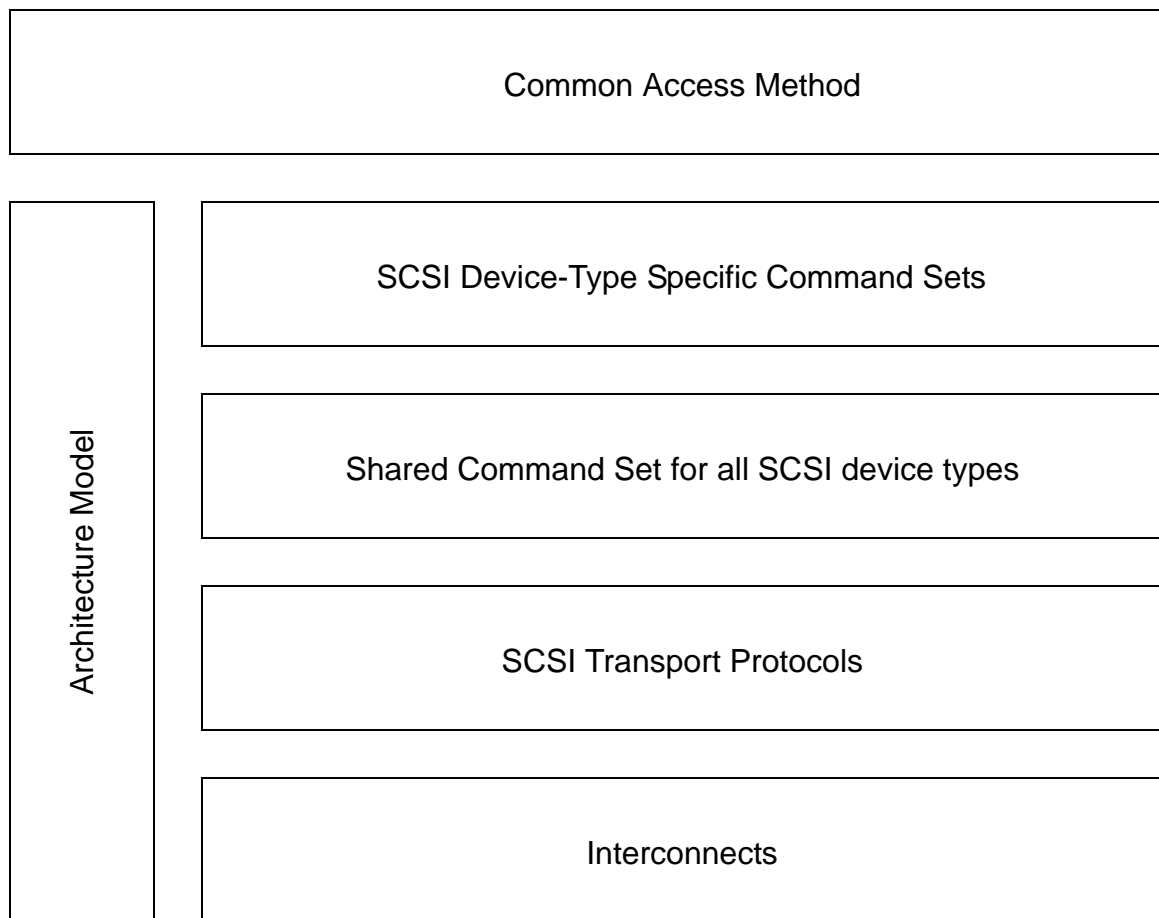
The interface protocol includes provision for the connection of multiple SCSI initiator ports (i.e., SCSI devices capable of initiating an I/O process) and multiple SCSI target ports (i.e., SCSI devices capable of responding to a request to perform an I/O process). Distributed arbitration (i.e., bus-contention logic) is built into the architecture of this standard. 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.

The set of SCSI standards specifies the interfaces, functions, and operations necessary to ensure interoperability between conforming SCSI implementations. This standard is a functional description. Conforming implementations may employ any design technique that does not violate interoperability.

This standard has made obsolete single-ended and multimode signaling alternatives. Implementations that use single-ended or multimode signaling alternatives should reference the SCSI Parallel Interface-2 standard (ISO/IEC 14776-112).

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 - SCSI Document Structure**

At the time this standard was generated examples of the SCSI document structure included:

Physical Interconnects:

- SCSI Parallel Interface - 4 [INCITS.362-200x]
- Serial Storage Architecture Physical Layer 1 [X3.293-1996]
- Serial Storage Architecture Physical Layer 2 [NCITS.307-1998]

SCSI Transport protocols:

- SCSI Parallel Interface - 2 [X3.302-1998]
- Serial Storage Architecture Transport Layer 1 [X3.295-1996]
- SCSI Fibre Channel Protocol - 2 [T10/1144D]
- SCSI Serial Bus Protocol - 2 [NCITS.325-1998]
- Serial Storage Architecture SCSI-3 Protocol [NCITS.309-1998]
- Serial Storage Architecture Transport Layer 2 [NCITS.308-1998]

Shared Command Set:

SCSI Primary Commands-3 standard [T10/1416D]

Device-Type Specific Commands Sets:

SCSI-3 Block Commands [NCITS.306-1998]  
 SCSI-3 Enclosure Services [NCITS.305-1998]  
 SCSI-3 Stream Commands [T10/997D]  
 SCSI-3 Medium Changer Commands [T10/999D]  
 SCSI Controller Commands - 2 [T10/1225D]  
 SCSI Multimedia Command Set - 2 [T10/1228D]

Architecture Model:

SCSI Architecture Model - 2 [T10/1157D]

The term SCSI is used wherever it is not necessary to distinguish between the versions of SCSI. The Small Computer System Interface - 2 (ANSI X3.131-1994) is referred to herein as SCSI-2.

## 2 Normative references

### 2.1 Normative references

The following standards contain provisions that, by reference in the text, constitute provisions of this standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent editions of the standards listed in clause 2.

Copies of the following documents may be obtained from ANSI: approved ANSI standards, approved and draft international and regional standards (e.g., ISO, IEC, CEN/CENELEC, ITUT), and approved and draft foreign standards including BSI, JIS, and DIN. For further information, contact ANSI Customer Service Department at 212-642-4900 (phone), 212-302-1286 (fax) or via the World Wide Web at <http://www.ansi.org>.

Additional availability contact information is provided as needed.

### 2.2 Approved references

ISO/IEC 14776-113, *SCSI Parallel Interface-3 standard*

ISO/IEC 14776-xxx, *SCSI Parallel Interface-4 standard*

EIA-700AOAE (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-700AOAF (SP-3652), *Detail Specification for Trapezoidal Connector 0.8 mm Pitch used with Very High Density Cable Interconnect (VHDCI)*

IEC 60512-2:1985-11, *Low Level Contact Resistance Test Procedure for Electric Connectors*

IEC 60512-11-7:1996-01, *Standard Practice For Conducting Mixed Flowing Gas Environmental Tests*

ISO 129, *Technical Drawings - Dimensioning - General Principals*

ISO 1660, *Technical Drawings - Dimensioning And Tolerancing*

ISO/IEC 14776-112, *SCSI Parallel Interface-2 standard*

## 2.3 References under development

At the time of publication, the following referenced standards were still under development. For information on the current status of the document, or regarding availability, contact the relevant standards body or other organization as indicated.

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

ISO/IEC 14776-xxx, *SCSI Primary Commands-3 standard*

ISO/IEC 14776-xxx, *SCSI Passive Interconnect Performance standard (T10/1439D)*

ISO/IEC 14776-xxx, *SCSI Signal Modeling -2 standard (T10/1514D)*

NOTE 1 - For more information on the current status of the document, contact the INCITS Secretariat at 202-737-8888 (phone), 202-638-4922 (fax) or via Email at [incits@itic.org](mailto:incits@itic.org). To obtain copies of this document, contact Global Engineering at 15 Inverness Way, East Englewood, CO 80112-5704 at 303-792-2181 (phone), 800-854-7179 (phone), or 303-792-2192 (fax).

## 2.4 Other references

For information on the current status of the listed documents, or regarding availability, contact the indicated organization.

SFF-8451, *SCA-2 Unshielded Connections*

NOTE 2 - For more information on the current status of the document, contact the SFF committee at 408-867-6630 (phone), or 408-867-2115 (fax). To obtain copies of this document, contact the SFF committee at 14426 Black Walnut Court, Saratoga, CA 95070 at 408-867-6630 (phone) or 408-741-1600 (fax).

ASTM D-4566, *Standard Test Methods for Electrical Performance Properties of Insulations and Jackets for Telecommunications Wire and Cable*

ASTM B827, *Standard Practice For Conducting Mixed Flowing Gas Environmental Tests*

IEEE 1364, *Verilog® Hardware Description Language*

T10/1378DT, *SCSI Domain Validation technical report*

NOTE 3 - For more information on the current status of the T10 document, contact the INCITS Secretariat at 202-737-8888 (phone), 202-638-4922 (fax) or via Email at [incits@itic.org](mailto:incits@itic.org). To obtain copies of this document, contact Global Engineering at 15 Inverness Way, East Englewood, CO 80112-5704 at 303-792-2181 (phone), 800-854-7179 (phone), or 303-792-2192 (fax).

## 3 Definitions, symbols, abbreviations, and conventions

### 3.1 Definitions

**3.1.1 A cable:** A 50-conductor cable assembly that provides an 8-bit DATA BUS and control signals.

**3.1.2 agent:** Carries out the actions of a requested service following the rules of the protocol.

**3.1.3 application client:** An object that is the source of SCSI commands. Further definition of an

application client is found in the SCSI Architecture Model-2 standard.

**3.1.4 asynchronous event notification:** An optional procedure used by SCSI target devices to notify SCSI initiator devices of events that occur when a pending task does not exist for that SCSI initiator device.

**3.1.5 asynchronous transfer:** An information transfer that uses the REQ/ACK handshake with an offset of zero.

**3.1.6 auto-contingent allegiance (ACA):** An optional condition of a task set following the return of a CHECK CONDITION status. See the SCSI Architecture Model-2 standard for a detailed definition of auto-contingent allegiance.

**3.1.7 backplane:** A printed circuit board with connectors attached that is used for interconnecting multiple SCSI devices.

**3.1.8 bulk cable:** The collection of conductors and associated insulation used between, but not including, the connectors or transition regions in a SCSI bus segment. Bulk cable includes permanent features (e.g., flat regions) designed for purposes of enabling connector attachment.

**3.1.9 bus segment path:** The electrical path between the bus segment terminators.

**3.1.10 byte:** Indicates an 8-bit construct.

**3.1.11 cable assembly:** A bulk cable that has connectors attached.

**3.1.12 confirmation:** A response returned to the application client that signals the completion of a service request (see 4.13).

**3.1.13 confirmed service:** A service available at the protocol service interface, that requires confirmation of completion. The confirmed service consists of the request and confirmation steps and optionally the indication and response steps.

**3.1.14 contact:** The electrically-conductive portion of a connector associated with a single conductor in a cable.

**3.1.15 contingent allegiance (CA):** An optional condition of a task set following the return of a CHECK CONDITION status. A detailed definition of contingent allegiance may be found in the SCSI Architecture Model-2 standard.

**3.1.16 current task:** A task that is in the process (i.e., of sending messages, status, transferring data, or transferring command data to or from the SCSI initiator port).

**3.1.17 cyclic redundancy check (CRC):** An error detecting code used to detect the validity of data.

**3.1.18 DATA BUS:** An 8-bit or 16-bit bus (see 8.2).

**3.1.19 data field:** The portion of a data group that contains data bytes.

**3.1.20 data group:** A sequence of data bytes, any pad bytes, and the four pCRC bytes transmitted during a DT DATA IN phase or a DT DATA OUT phase.

**3.1.21 data group transfer:** Parallel transfers that transfer data and pCRC information using only data groups.

**3.1.22 device server:** An object within the logical unit that executes SCSI tasks according to the rules for task management as described in the SCSI Architecture Model-2 standard.

**3.1.23 differential:** A signaling alternative that uses drivers and receivers with two complementary signals to improve signal-to-noise ratios.

**3.1.24 double transition (DT):** The latching of data on both the assertion edge and the negation edge of the REQ or ACK signals.

**3.1.25 driver:** The circuitry used to control the state of a signal line in a bus segment.

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

**3.1.27 expander:** A device that connects SCSI bus segments together to form a single SCSI domain.

**3.1.28 fast-5:** Negotiated to receive synchronous data at a transfer period that translates into a transfer rate of less than or equal to 5 megatransfers per second.

**3.1.29 fast-10:** Negotiated to receive synchronous data at a transfer period that translates into a transfer rate greater than 5 megatransfers per second and less than or equal to a transfer rate of 10 megatransfers per second.

**3.1.30 fast-20:** Negotiated to receive synchronous data at a transfer period that translates into a transfer rate greater than 10 megatransfers per second and less than or equal to a transfer rate of 20 megatransfers per second.

**3.1.31 fast-40:** Negotiated to receive synchronous data at a transfer period that translates into a transfer rate greater than 20 megatransfers per second and less than or equal to a transfer rate of 40 megatransfers per second.

**3.1.32 fast-80:** Negotiated to receive synchronous data at a transfer period that translates into a transfer rate greater than 40 megatransfers per second and less than or equal to a transfer rate of 80 megatransfers per second.

**3.1.33 fast-160:** Negotiated to receive synchronous data at a transfer period that translates into a transfer rate of 160 megatransfers per second.

**3.1.34 fast-320:** Negotiated to receive synchronous data at a transfer period that translates into a transfer rate of 320 megatransfers per second.

**3.1.35 field:** A group of one or more contiguous bits.

**3.1.36 hard reset:** A SCSI target device action in response to a reset event in which a SCSI target device performs the operations described in 12.4.

**3.1.37 indication:** A signal notifying the SCSI device server or task manager of an event (see 4.13)

**3.1.38 information unit transfer:** Parallel transfers that transfer data, status, commands, task attributes, task management, iuCRC, and nexus information using SPI information units.

**3.1.39 initial connection:** An initial connection is the result of a physical connect. It exists from the assertion of the BSY signal (see 10.5) in a SELECTION phase until the next BUS FREE phase or the next QAS REQUEST message.

**3.1.40 initiator:** Synonymous with SCSI initiator port (see 3.1.97).

**3.1.41 I/O process:** An I/O process consists of one initial connection (or, if information units are enabled, the establishment of a nexus) and zero or more physical or logical reconnections, all pertaining to a single task or a group of tasks. An I/O process begins with the establishment of a nexus. If the SPI information

unit transfers are disabled an I/O process normally ends with a COMMAND COMPLETE message. If information unit transfers are enabled an I/O process normally ends with a SPI L\_Q information unit with the type field set to status and the DATA LENGTH field set to zero.

**3.1.42 interconnect:** The electrical bulk cable, connectors, and passive loads used to connect the TERMPWR, terminators, and SCSI devices in a SCSI bus segment

**3.1.43 I\_T nexus:** A nexus that exists between a SCSI initiator port and a SCSI target port.

**3.1.44 I\_T\_L nexus:** A nexus that exists between a SCSI initiator port, a SCSI target port, and a logical unit. This relationship replaces the prior I\_T nexus.

**3.1.45 I\_T\_L\_Q nexus:** A nexus between a SCSI initiator port, a SCSI target port, a logical unit, and a queue tag following the successful receipt of a queue tag. This relationship replaces the prior I\_T nexus or I\_T\_L nexus.

**3.1.46 intersymbol interference (ISI):** The effect that a transition (or symbol) on a signal line has on transitions before or after that transition (or symbol) received on the same line.

**3.1.47 iuCRC protection:** The use of CRC to detect DT DATA phase data transmission errors during SPI information unit transfers.

**3.1.48 logical connect:** Establishes an I\_T\_L\_Q nexus using SPI L\_Q information units during an initial connection.

**3.1.49 logical disconnect:** Reduces the current I\_T\_L\_Q nexus to an I\_T nexus.

**3.1.50 logical reconnect:** Reestablishes an I\_T\_L\_Q nexus from an I\_T nexus using SPI L\_Q information units.

**3.1.51 logical unit:** An externally addressable entity within a SCSI target device. See the SCSI Architecture Model-2 standard for a detailed definition of a logical unit.

**3.1.52 logical unit number:** An identifier for a logical unit.

**3.1.53 logical unit reset:** A logical unit action in response to a logical unit reset event in which the logical unit performs the operations described in SCSI Architecture Model-2.

**3.1.54 logical unit reset event:** An event that triggers a logical unit reset from a logical unit as described in SCSI Architecture Model-2.

**3.1.55 magnitude:** The absolute value of a number or quantity.

**3.1.56 maximum transfer rate:** The highest transfer rate supported by a SCSI device.

**3.1.57 megatransfers per second:** The repetitive rate that data are transferred across the bus. This is equivalent to megabytes per second on an 8-bit wide bus.

**3.1.58 message:** One or more bytes transferred between a SCSI initiator port and a SCSI target port to perform link control or task management, or to associate task attributes with commands.

**3.1.59 multidrop:** A characteristic of the SCSI bus segment that allows SCSI devices to be connected to the SCSI bus segment without disrupting the continuity of the electrical path between the terminators (see 4.3).

**3.1.60 negotiation required flag:** A flag maintained by a port for each other port indicating when it is required to originate negotiation with that port (see 4.12.3).

**3.1.61 nexus:** A relationship between a SCSI initiator port and a SCSI target port that may extend to a logical unit and a queue tag.

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

**3.1.63 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 (e.g., 1, 3, 5, 7, or 9). See 3.1.70, parity bit. If an even number of asserted bits are detected at the receiver a parity error occurs.

**3.1.64 one :** A true signal value or a true condition of a variable.

**3.1.65 P cable:** A 68-conductor cable assemble or backplane that uses an 80-conductor connector to provide the 16-bit DATA BUS and control signals.

**3.1.66 paced transfer:** Parallel transfers that transfer information using pacing.

**3.1.67 pacing:** Use of the ACK or REQ signal as a continuously running clock in combination with the P1 signal to indicate when data is valid.

**3.1.68 packetized:** A method of transferring information using SPI information units (see 4.11.3.3)

**3.1.69 pad field:** The portion of a data group that contains pad information.

**3.1.70 parity bit:** A bit associated with a byte that is used to detect the presence of an odd number of asserted bits within the byte. The parity bit is driven such that the number of logical ones in the byte plus the parity bit is odd.

**3.1.71 parity error:** When the number of assertions on the associated data byte plus the parity bit equal to an even number (e.g., 0, 2, 4, 6, or 8).

**3.1.72 parallel transfer:** The transfer of information using information transfer phases.

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

**3.1.74 pCRC field:** The portion of a data group that contains pCRC information.

**3.1.75 pCRC protection:** The use of CRC to detect DT DATA phase data transmission errors during data group transfers.

**3.1.76 pending task:** A task that is not the current task.

**3.1.77 physical connect:** The act of establishing an I\_T nexus during a connection.

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

**3.1.79 physical reconnect:** The act of resuming a nexus to continue processing a task. A SCSI target port initiates a physical reconnect when conditions are appropriate for the physical bus to transfer data associated with a nexus between a SCSI initiator port and a SCSI target port.

**3.1.80 physical reconnection:** A physical reconnection is the result of a physical reconnect that exists from the assertion of the BSY signal in a SELECTION or RESELECTION phase. A physical reconnection ends with the BUS FREE phase (see 10.2) or a QAS REQUEST message (see 10.4.3).



- 3.1.81 planar:** A bulk cable construction where the signal wires are side by side. The bulk cable may contain twisted pairs or straight wires or any combination of twisted pairs/straight wires.
- 3.1.82 point-to-point:** A subset of the multidrop architecture (see 4.3) where only two SCSI devices are attached within an allowed stub length of the terminators on a SCSI bus segment.
- 3.1.83 port:** A single attachment to a SCSI bus segment from a SCSI device.
- 3.1.84 power on:** Power being applied.
- 3.1.85 queue:** The arrangement of tasks within a task set usually according to the temporal order that they were created.
- 3.1.86 queue tag:** The parameter associated with a task that uniquely identifies it from other tagged tasks for a logical unit from the same SCSI initiator port.
- 3.1.87 receiver:** The circuitry used to detect the electrical state of the bus segment.
- 3.1.88 request:** A request to the initiator parallel interface agent to invoke a service (see 4.13)
- 3.1.89 response:** A response from the SCSI device server or task manager in reply to an indication (see 4.13).
- 3.1.90 reset event:** An event that triggers a hard reset from a SCSI device as described in 12.5. Reset events defined in this standard are bus reset event (see 12.5.2), power on reset event (see 12.5.3), and target reset event (see 12.5.4).
- 3.1.91 SCSI address:** The decimal representation of the unique address assigned to a SCSI device.
- 3.1.92 SCSI bus:** All the conductors and connectors required to attain signal line continuity between every driver, receiver, and terminator for each signal.
- 3.1.93 SCSI device:** A device containing at least one SCSI port and the means to connect its drivers and receivers to the bus segment.
- 3.1.94 SCSI device port:** Either a SCSI initiator port or a SCSI target port
- 3.1.95 SCSI ID:** The bit-significant representation of the SCSI address.
- 3.1.96 SCSI initiator device:** A SCSI device containing application clients and SCSI initiator ports that originate device service and task management requests to be processed by a SCSI target device. See the SCSI Architecture Model-2 standard for a detailed definition of a SCSI initiator device.
- 3.1.97 SCSI initiator port:** A SCSI initiator device object that acts as the connection between application clients and the service delivery subsystem through which requests and responses are routed. See the SCSI Architecture Model-2 standard for a detailed definition of a SCSI initiator port.
- 3.1.98 SCSI target device:** A SCSI device containing logical units and SCSI target ports that receives device service and task management requests for processing. See the SCSI Architecture Model-2 standard for a detailed definition of a SCSI target device.
- 3.1.99 SCSI target port:** A SCSI target device object that contains a task router and acts as the connection between device servers and task managers and the service delivery subsystem through which requests and responses are routed. See the SCSI Architecture Model-2 standard for a detailed definition of a SCSI target port.
- 3.1.100 SCSI terminator:** The terminator is at each end of a SCSI bus segment. The terminator provides

impedance matching and biasing, holding the bus in a negated state when it is not driven.

**3.1.101 signal assertion:** The act of driving a signal to the true state.

**3.1.102 signal negation:** The act of performing a signal release or of driving a signal to the false state.

**3.1.103 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.

**3.1.104 single transition (ST):** The latching of data only on the assertion edge of the REQ or ACK signals.

**3.1.105 skew:** The maximum difference in propagation time allowed between any two SCSI bus signals measured between two specified positions in the bus segment using a free running clock data pattern.

**3.1.106 source (a signal):** The act of either signal assertion, signal negation, or signal release.

**3.1.107 SPI information unit:** Data structures that encapsulate data, status, command, task attributes, iuCRC, and nexus information into various formats.

**3.1.108 stub:** Any electrical path connected to the bus segment that is not part of the bus segment path.

**3.1.109 synchronous transfer:** An information transfer that uses a REQ/ACK offset other than zero and does not use pacing.

**3.1.110 target:** Synonymous with SCSI target port (see 3.1.99).

**3.1.111 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 physical or logical reconnections, all pertaining to the task.

**3.1.112 task manager:** An agent within the logical unit that processes task management functions.

**3.1.113 task management function:** A task manager service that may be invoked by a task management message or by setting one of the task management functions in a SPI L\_Q information unit to affect the execution of one or more tasks.

**3.1.114 task set:** A group of tasks within a logical unit, whose interaction is dependent on the task management, contingent allegiance and auto-contingent allegiance rules. See the SCSI Architecture Model-2 standard for a detailed definition of a task set.

**3.1.115 transceiver:** A device that implements both the SCSI receiver and SCSI driver functions.

**3.1.116 transfer period:** The negotiated time between edges of REQ or ACK that latch data. For ST the transfer period is measured from an assertion edge of the REQ or ACK signal to the next assertion edge of the signal. For DT the transfer period is measured from a transition edge of the REQ or ACK signal to the next transition edge of the signal.

**3.1.117 transfer rate:** The negotiated megatransfers per second.

**3.1.118 upper level protocol:** Any protocol executed through services provided by a lower level protocol.

**3.1.119 vendor-specific:** Something (e.g., a bit, field, or code value) that is not defined by this standard and may be used differently in various implementations.

**3.1.120 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 3.1.6)
AWG	American wire gauge
CA	Contingent allegiance (see 3.1.15)
CMOS	Complementary metal oxide semiconductor
CRC	Cyclic Redundancy Check (see 3.1.17)
DT	Double transition (see 3.1.24)
DUT	Device under test
EMI	Electromagnetic interference
EMC	Electromagnetic compatibility
ESD	Electrostatic discharge
HVD	High voltage differential
IDC	Insulation displacement contact
ISI	Intersymbol interference
iuCRC	Information unit CRC
LSB	Least significant bit
LUN	Logical unit number
LVD	Low voltage differential
MSB	Most significant bit
NEXT	Near end crosstalk
pCRC	Parallel CRC
PPR	Parallel protocol request
QAS	Quick Arbitration and Selection
SAM-2	SCSI Architecture Model-2
SCSI	Small Computer System Interface
SCSI-2	Small Computer System Interface - 2
SCSI-3	Small Computer System Interface - 3
SDTR	Synchronous data transfer request
SE	Single ended
SPC-3	SCSI Primary Commands-3
SPI-2	SCSI Parallel Interface-2
SPI-3	SCSI Parallel Interface-3
SPI-4	SCSI Parallel Interface-4
ST	Single transition (see 3.1.104)
TDR	Time Domain Reflectometer
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 an 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 may not:** Keywords that indicates flexibility of choice with no implied preference.

**3.3.6 obsolete :** A keyword indicating that an item was defined in prior SCSI standards but has been removed from this standard.

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

**3.3.8 reserved:** A keyword referring to bits, bytes, words, fields and code values that are set aside for future standardization. Their use and interpretation may be specified by future extensions to this or other standards. A reserved bit, byte, word or field shall be set to zero, or in accordance with a future extension to this standard. Recipients are not required to check reserved bits, bytes, words or fields for zero values. Receipt of reserved code values in defined fields shall be reported as an error.

**3.3.9 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.

**3.3.10 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 American National 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).

An alphanumeric list (e.g., a,b,c or A,B,C) of items indicate the items in the list are unordered.

A numeric list (e.g., 1,2,3) of items indicate the items in the list are ordered (i.e., item 1 is required occur or complete before item 2).

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

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

### 3.5 Notation for Procedures and Functions

Procedure Name ([input:1a|input:1b|input:1c][,input:2a+input:2b]...[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:1a input:1b ...	A number of arguments of which only one shall be used in any single procedure
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.

## 4 General

### 4.1 General overview

This standard defines the cables, connectors, signals, transceivers, terminators, and protocol used to interconnect parallel SCSI device ports and the services provided to the application client.

### 4.2 Cables, connectors, signals, transceivers

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

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

**Table 1 - LVD transceiver/speed support map**

Data Latching (ST/DT)	Maximum transfer rate							
	Async.	Fast-5	Fast-10	Fast-20	Fast-40	Fast-80	Fast-160	Fast-320
<b>ST</b>	yes	yes	yes	yes	yes	no	no	no
<b>DT</b>	no	no	yes	yes	yes	yes	paced	paced

Key:  
 yes = Transceiver/speed combination supported by this standard.  
 no = Transceiver/speed combination not supported by this standard.  
 paced = Transceiver/speed combination using paced transfers is supported by this standard (see 4.10).

SCSI devices may connect to the bus segment via 8-bit or 16-bit ports. The 8-bit ports shall connect to a bus segment with an A cable or equivalent (see clause 5). The 16-bit ports shall connect to a bus segment with a P cable or equivalent (see clause 5).

### 4.3 Physical architecture of bus segment

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

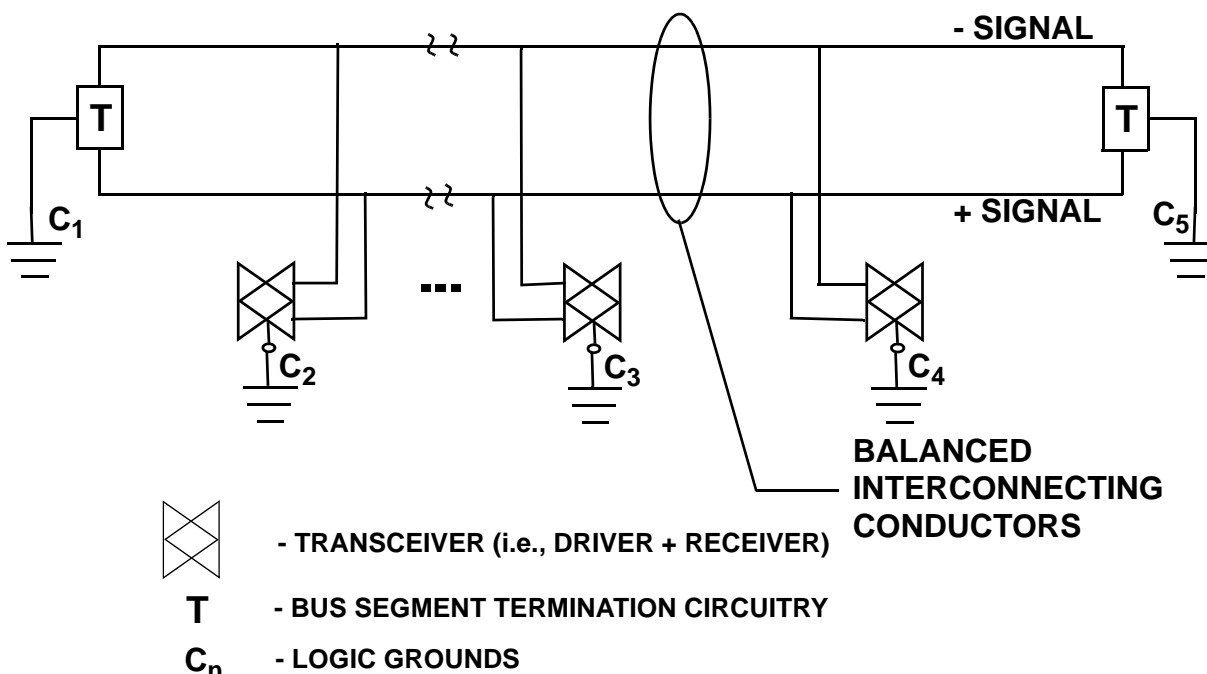


Figure 2 - Differential SCSI bus segment

#### 4.4 Driver-receiver connections

A driver-receiver connection is the connection between a driver and a receiver on a single SCSI bus segment. A SCSI signal may be used by more than one driver-receiver connection (i.e., a SCSI signal connects to all the SCSI devices where a driver-receiver connection only connects the selected SCSI initiator/target port pair). Driver-receiver connections always occur in pairs; one driver-receiver connection drives a SCSI signal from an initiator driver to a target receiver and the other driver-receiver connection drives the same SCSI signal from a target driver to an initiator receiver. Except for or-tied signals (see 8.4) only one end of the connection shall drive the SCSI bus at a time.

Setting programmable electrical properties on one end of the connection shall cause the drivers for all the SCSI signals on that end of the connection to be set to those electrical properties.

Setting programmable electrical properties at one end of the connection shall not automatically cause the driver on the other end of the connection to be set to those electrical properties.

#### 4.5 Physical topology details and definitions

The SCSI bus is a multidrop architecture described in 4.3. 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 between the two terminators forms the bus segment path. Any electrical path that is not part of the bus segment path is a stub. The point where a stub meets the bus segment path is termed the stub connection.

This standard requires the bus signals to be held at a negated state when no drivers are active on the bus. The terminators hold the bus signals to a negated state when there are no active drivers on the bus, and they provide impedance matching.

Figure 3 shows examples of connectors, bus segment paths, stubs, and stub connections.

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

The functional definitions are:

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

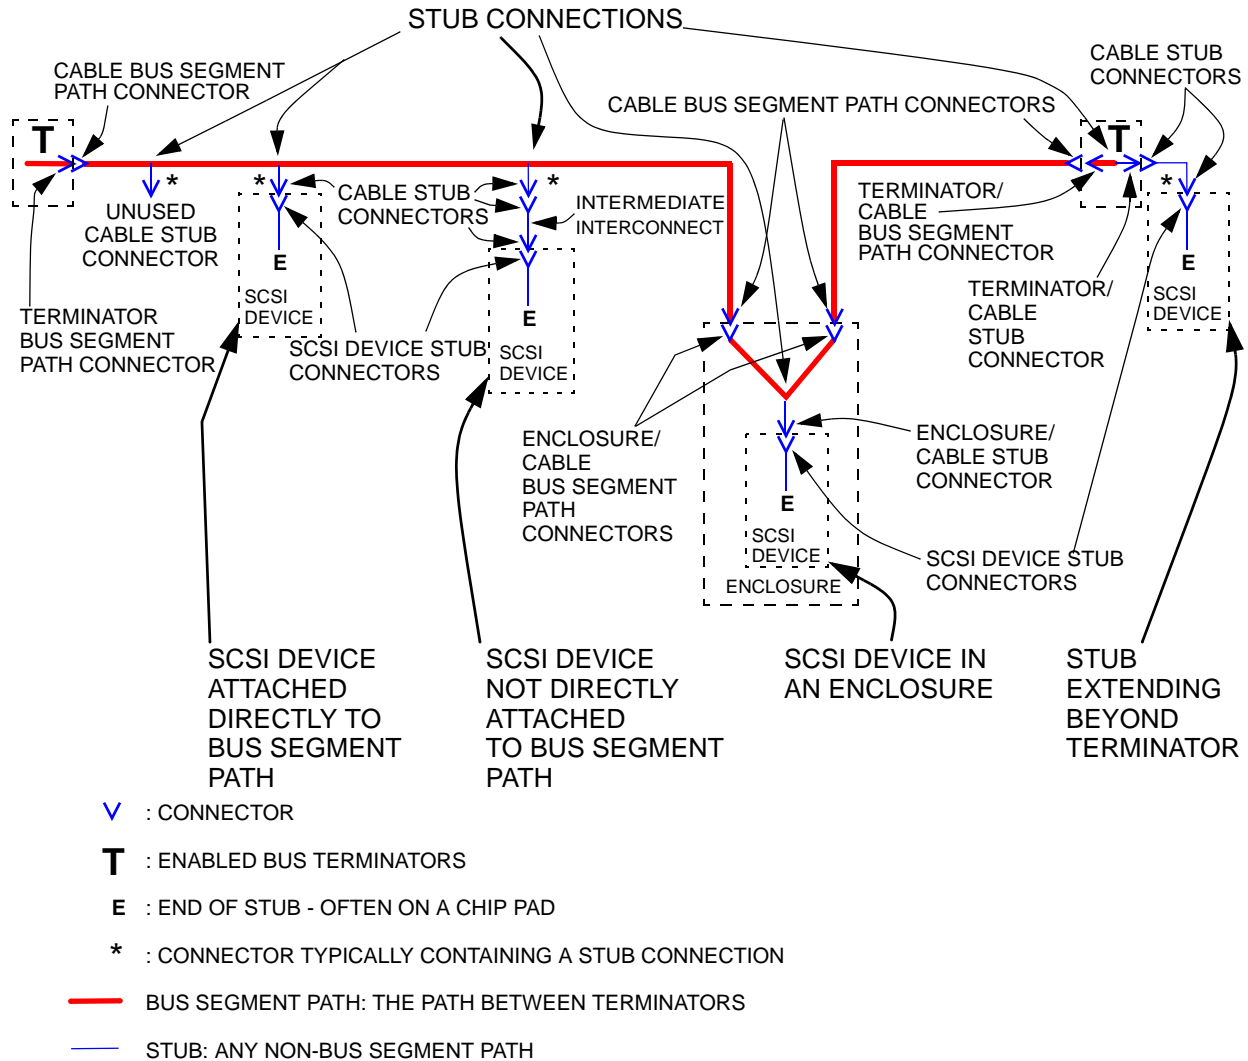
Common physical placement definitions are:

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

SCSI bus connectors (e.g., device stub connector, terminator bus segment 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 3 - SCSI bus segment topology details**

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

NOTE 4 - Any extensions of the connection beyond the terminator as shown in the right side of figure 3 should be minimized or avoided as that extension produces stubs and bus segment loading.

NOTE 5 - In order to support daisy-chain connections, SCSI devices that use shielded connectors should provide two shielded SCSI device connectors on the SCSI 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 SCSI device enclosure is included when calculating the total cable length of the SCSI bus segment. (see figure 3)

## 4.6 Bus segment loading

Bus segment loading is the electrical current flowing through the stub connection for lines that are not

being driven by the attached SCSI device. The bus segment termination circuitry also provides bus segment loading. Bus segment 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 plus and minus 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 plus and minus signals rather than by the value of the current.

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

Bus segment termination loading is separate from bus segment loading. SCSI devices containing enabled bus segment termination shall present loading at the stub connection that is no more than the sum of the maximum allowed termination loading plus the maximum allowed bus segment loading. See 4.7 for requirements of disabled termination circuitry.

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

## 4.7 Termination requirements

The SCSI bus segment termination defines the ends of the SCSI bus segment. Bus segment termination is required to set the negated state when no SCSI device is driving (also called biasing) and to match the impedance to that of the bulk cable. A termination circuit is providing bus segment 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 segment 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 SCSI devices in terms of bus segment loading if they are individually attached to the bus segment. If they are contained within a SCSI device the disabled terminators become part of the SCSI device load budget for that SCSI device.

## 4.8 SCSI device Addressing


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

## 4.9 Clocking methods for data transfers


This standard defines optional methods of latching data from the REQ and ACK signals depending on whether ST DATA phases or DT DATA phases are being used for information transfers as shown in figure 4. Data shall only be latched on the asserting edge of the REQ or ACK signal except in DT DATA phases. When DT DATA phases are used, data shall be latched on both the asserting edge and the negating edge of the REQ or ACK signal.

Regardless of whether ST or DT transfers are enabled the negotiated transfer period sets the maximum rate at which the data is clocked in megatransfers per second. As a result, the time from rising edge to rising edge for REQ and ACK signals for the same transfer rate is twice as long for a DT transfer as it is for an ST transfer. An example of a negotiated transfer period of 25 ns with ST transfers is shown in figure 5. An example of a negotiated transfer period of 25 ns with DT transfers is shown in figure 6. In these examples the rising edge to rising edge time for DT is 50 ns while for ST is 25 ns. In both cases data is transferred at 25 ns intervals. For ST and DT synchronous transfers the clocking signal (i.e., REQ or ACK) occurs when the DATA BUS is in a steady state.

Figure 7 shows an example of paced transfers with a negotiated transfer period of 6,25 ns at the receiving SCSI device's connector. Data being transmitted using paced transfers is latched in the center of a bit cell, however, the relationship between the data and REQ or ACK is required to be shifted in the SCSI devices receiver to align REQ or ACK to the center of the data cell that then matches the synchronous transfers DT Data shown in figure 4. For paced transfers the clocking signal (i.e., REQ or ACK) may occur when the DATA BUS is changing state.

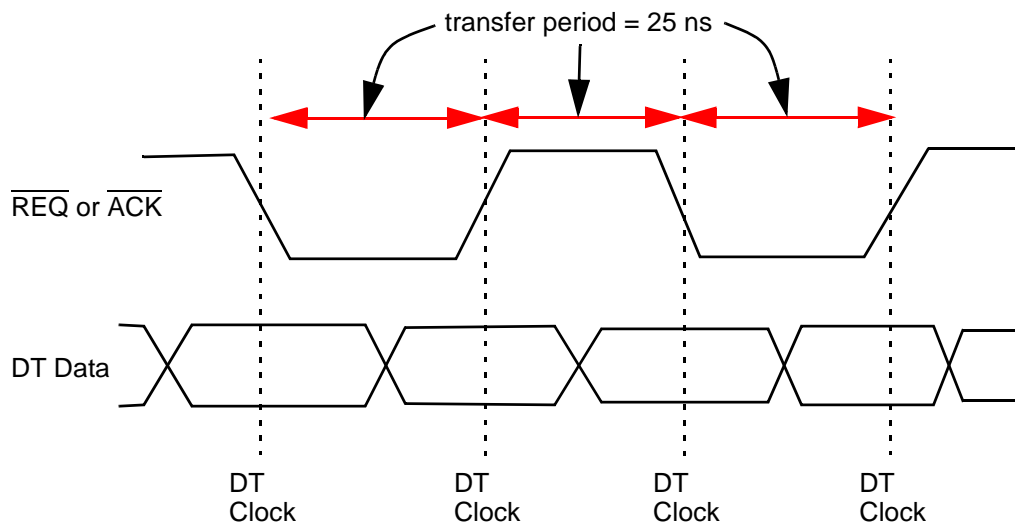


**Figure 4 - ST latching data vs. DT latching data**



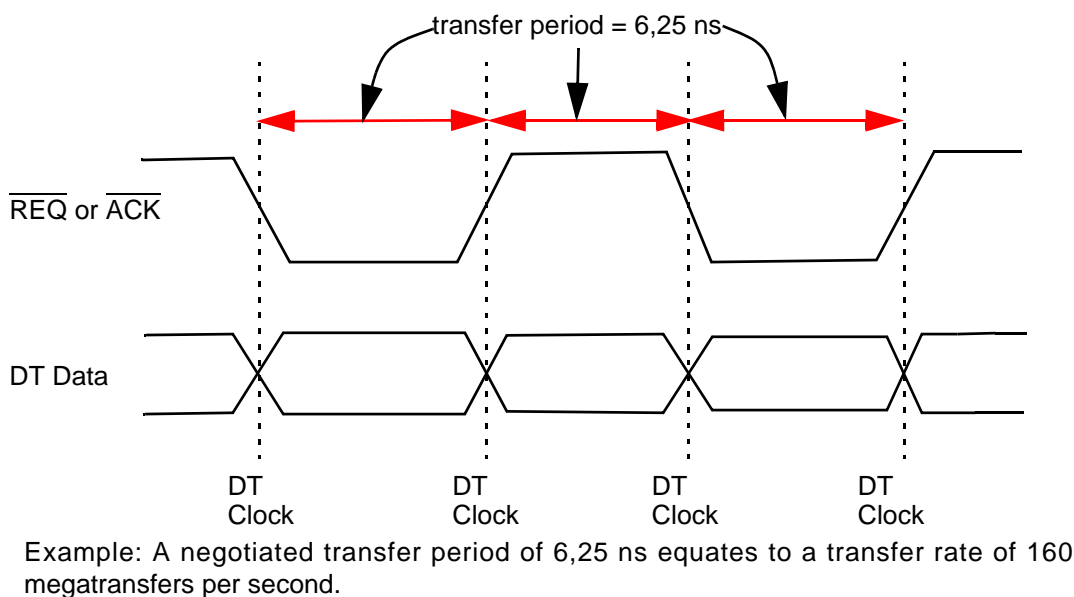
Example: A negotiated transfer period of 25 ns equates to a maximum transfer rate of 40 megatransfers per second.

**Figure 5 - ST synchronous transfer example**



Example: A negotiated transfer period of 25 ns equates to a maximum transfer rate of 40 megatransfers per second.

**Figure 6 - DT synchronous transfer example**



**Figure 7 - Paced transfer example**

#### 4.10 Paced transfer on a SCSI bus

A SCSI bus that supports paced transfers has driver and receiver functions required in addition to those used with synchronous transfers or asynchronous transfers.

For fast-160 these functions include driver precompensation, receiver skew compensation, receiver clock shifting, and an optional receiver signal adjustment (see figure 8). In addition the driver precompensation may be switched out of the data path at the request of the receiving SCSI device. See 7.2.2 for more information on precompensation.

For fast-320 these functions include receiver skew compensation, receiver clock shifting, and a receiver signal adjustment (see figure 9). For fast-320 precompensation shall be disabled.

The receiver skew compensation and clock shifting adjust the timing relationship between the clocking signal (i.e., REQ or ACK) and the signals being clocked (e.g., the data bus signals). That adjustment causes the signals being clocked to align with the middle of the clock signal when those signals enter the receiver (see figure 10). The receiver is then able to use the clock signal to latch valid data.

During paced transfers the clock signal (i.e., REQ or ACK) transitions at the negotiated transfer period. Data is qualified by the clock signal and the phase of the P1 signal (see 10.7.4.3).

Receiver skew compensation is vendor specific and, therefore, not defined in this standard.




Figure 8 - Example of a SCSI bus with fast-160 paced transfers




Figure 9 - Example of a SCSI bus with fast-320 paced transfers

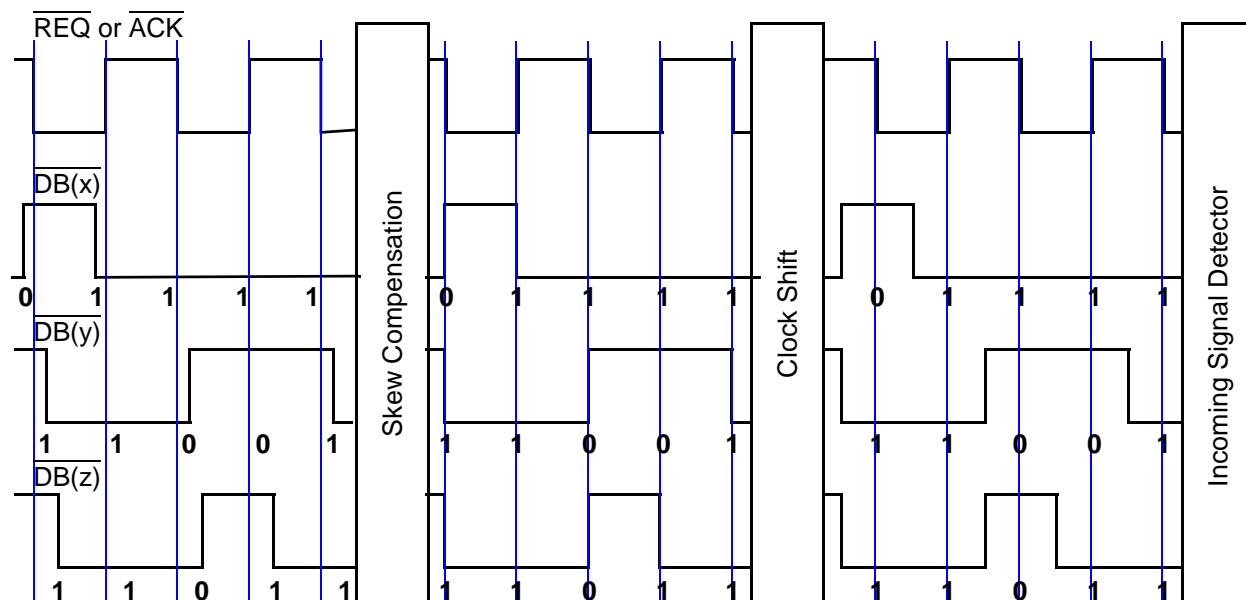


Figure 10 - Skew compensation and clock shift example

## 4.11 Data transfers

### 4.11.1 Data transfer modes

#### 4.11.1.1 Asynchronous transfers

SCSI device ports default to 8-bit asynchronous transfers.

8-bit asynchronous transfers are used for all COMMAND, STATUS, and MESSAGE phases.

ST DATA phases may use 8-bit or 16-bit asynchronous transfers. Asynchronous transfers are not permitted in DT DATA phases.

#### 4.11.1.2 Synchronous transfers

ST DATA phases shall use synchronous transfers when a synchronous transfer agreement is in effect. ST DATA phases may use 8-bit or 16-bit synchronous transfers (see 4.12.4.2).

DT DATA phases shall use synchronous transfers when a synchronous transfer agreement (see 4.12.4.2) is in effect. DT DATA phases shall only use wide transfers.

#### 4.11.1.3 Paced transfers

Paced transfers shall only be used in DT DATA phases when a paced transfer agreement is in effect (see 4.12.4.2). DT DATA phases shall only use wide transfers.

### 4.11.2 ST DATA phase parallel transfers

The format of data transmitted during ST DATA phases consists of data and protection. Parity is used to protect the data (see 11.2).

### 4.11.3 DT DATA phase parallel transfers

#### 4.11.3.1 DT DATA phase parallel transfers format

The format of the data transmitted during DT DATA phases is dependent on the negotiated protocol options. If data group transfers are enabled then all data and protection are transmitted in data groups. If information unit transfers are enabled then all nexus, task management, task attribute, command, data, and protection are transmitted in SPI information units.

#### 4.11.3.2 Data group transfers

Data group transfers are permitted when a synchronous transfer agreement is in effect. Data group transfers are not permitted when an asynchronous transfer agreement or a paced transfer agreement is in effect (see 10.7.5).

When using data group transfers each DT DATA IN phase and DT DATA OUT phase contains one or more data groups. A data group consists of a non-zero length data field containing an even number of bytes (see 16.3.4), followed by a pad field when pad bytes are needed, and then followed by a pCRC field. The number of bytes transferred within a data group shall always be a multiple of four.

If the number of bytes in the data field is not a multiple of four the transmitting SCSI device shall place two pad bytes into the pad field. If the number of bytes in the data field is a multiple of four the transmitting SCSI device shall omit the pad field. Regardless of the number of bytes in the data field the pCRC field shall be the last four bytes of the data group.

The value of the pad bytes within the pad field is vendor specific.

During DT DATA IN phase if the number of bytes in a data field is not a multiple of two bytes, then after sending the pad and pCRC fields the SCSI target port shall change to MESSAGE IN phase and send an IGNORE WIDE RESIDUE message (see 16.3.4) with the NUMBER OF BYTES TO IGNORE field set to 01h.

During DT DATA OUT phase if a SCSI target port requests a pCRC field prior to the last data field of a task, the SCSI initiator port shall transmit an even number of bytes in that data field.

The pCRC shall be used to protect all data group transfers. The SCSI device transmitting data sends the necessary pad field(s) and a pCRC field at a point determined by the SCSI target port.

#### 4.11.3.3 Information unit transfers

Information unit transfers are permitted when a synchronous transfer agreement is in effect. Information unit transfers are mandatory when a paced transfer agreement is in effect. Information unit transfers are not permitted when an asynchronous transfer agreement is in effect (see 10.7.5).

During information unit transfers each DT DATA IN phase and DT DATA OUT phase contains one or more SPI information units. The number of bytes transferred within a SPI information unit shall always be a multiple of four.

If the number of bytes in the SPI information unit is not a multiple of four, the transmitting SCSI device shall transmit one, two, or three pad bytes as is necessary to make the transfer a multiple of four bytes before transmitting an iuCRC. If the number of bytes in the SPI information unit is a multiple of four the transmitting SCSI device shall not transmit any pad bytes. Regardless of the number of bytes in the SPI information unit the last four bytes of the SPI information unit shall be an iuCRC.

The value of the pad bytes is vendor specific.

The iuCRC shall be used to protect all SPI information units. The SCSI device that originates the SPI information unit sends the necessary pad byte(s) and iuCRC field(s).



An iuCRC interval may also be specified. The iuCRC interval specifies the number of bytes transferred before pad bytes (if any) and the iuCRC is transferred within SPI data information units and SPI data stream information units. A SPI data information unit or a SPI data stream information unit may contain zero or more iuCRC intervals depending on the values specified in the SPI L\_Q information unit. At a minimum there shall be at least one iuCRC at the end of each SPI data information unit and SPI data stream information unit regardless of the size of the iuCRC interval. If specified, an iuCRC interval shall begin on the first transfer of each data information unit or data stream information unit.

The iuCRC interval is required to be a multiple of two, however, if it is not a multiple of four then two pad bytes shall be transmitted before the iuCRC is transmitted.

SPI data stream information units may be used to transfer data to or from a SCSI device. Support of data streaming during DT DATA OUT phases, called write streaming, is mandatory. Support of data streaming during DT DATA IN phases, called read streaming, is optional. The use of read streaming is part of the negotiated transfer agreement between two SCSI devices (i.e., the RD\_STRM bit set to one). A SCSI target port is not required to use read streaming even if streaming support is enabled.

A SCSI target port, while streaming data, may give an indication that the stream of SPI data stream information units are about to end while still sending the current SPI data stream information unit. This early warning is called flow control. Support of flow control during DT DATA OUT phases, called write flow control, is optional. Support of flow control during DT DATA IN phases, called read flow control, is mandatory if read streaming is enabled. The use of write flow control is part of the negotiated transfer agreement between two SCSI devices (i.e., the WR\_FLOW bit set to one).

NOTE 6 - Because the read flow control and read streaming require the same SCSI bus signals to be controlled in the same way at the same time this standard makes no further reference to read flow control. However, that functionality is implied by the term read streaming.

## 4.12 Negotiation

### 4.12.1 Negotiation introduction

PARALLEL PROTOCOL REQUEST (PPR) (see 16.3.12), SYNCHRONOUS DATA TRANSFER REQUEST (SDTR) (see 16.3.16), and WIDE DATA TRANSFER REQUEST (WDTR) (see 16.3.18) messages are used to alter the transfer agreement between two ports (see 3.1.83). The transfer agreement defines the protocol used during DATA phases (e.g., transfer period, REQ/ACK offset, transfer width) and agreement on features not affecting DATA phases (e.g., QAS). All other information transfer phases (i.e., COMMAND, MESSAGE, and STATUS) use eight-bit asynchronous data transfers.

PPR, SDTR, and WDTR messages are called negotiation messages. When a SCSI initiator port sends one of them, the message names are PPR OUT, SDTR OUT, and WDTR OUT. When a SCSI target port sends one of them, the message names are PPR IN, SDTR IN, and WDTR IN. A negotiation sequence consists of at least one matching set of negotiation messages (e.g., PPR OUT and PPR IN).

A transfer agreement is maintained by each port for each other port on the SCSI bus. Each port (see 3.1.83) may be used as either a SCSI target port (see 3.1.99) or a SCSI initiator port (see 3.1.97). The same transfer agreement applies whether the port is being used as a SCSI target port or as a SCSI initiator port.

### 4.12.2 Negotiation algorithm

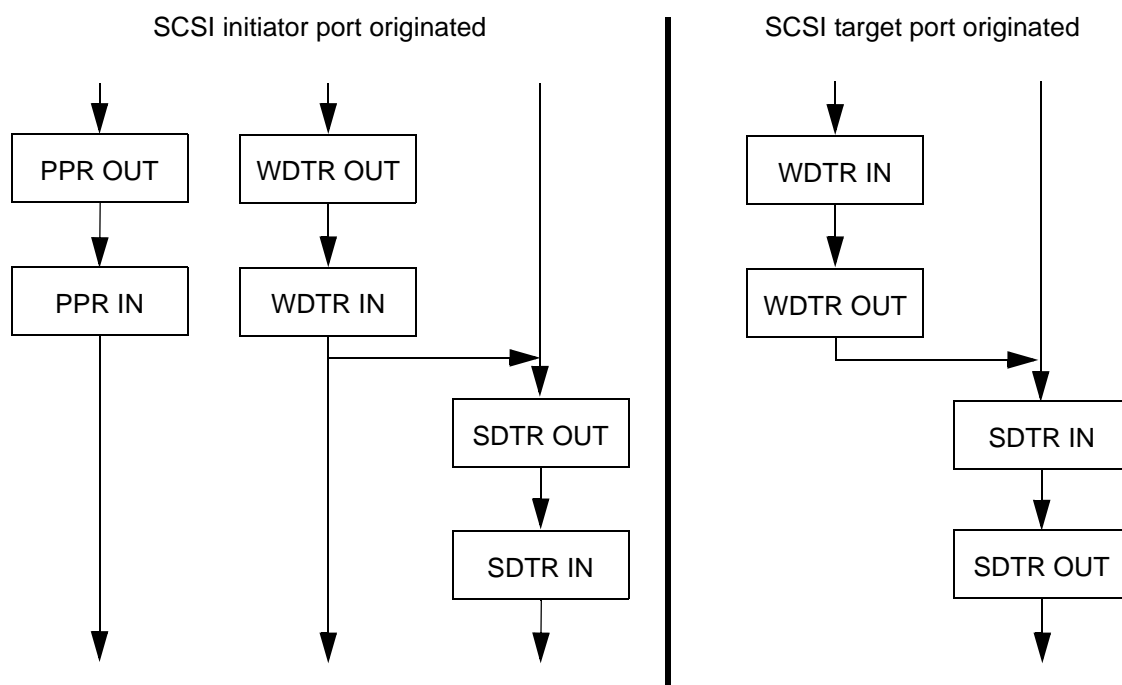
A SCSI initiator port and SCSI target port exchange negotiation messages to perform negotiation. The originating port is the one that sends the first negotiation message and the responding port is the one that replies.

Ports shall not set message fields to values they do not support. The originating port should set the fields in

the originating negotiation message to the maximum values (e.g., fastest transfer period, largest REQ/ACK offset) it supports. If the responding port is able to support the requested values, it shall return the same values in the responding negotiation message. If the responding port requires different values (i.e., a subset of the originating port's request), it shall return those values in the responding negotiation message (e.g., if the originating port asks for a REQ/ACK offset of 32 and the responding port only supports a REQ/ACK offset of 16, then the responding port replies with an offset of 16).

If the responding negotiation message contains values the originating port does not support, the originating port shall respond with a MESSAGE REJECT message.

The valid error-free negotiation message sequences are shown in figure 11. A description of the possible message sequences see 4.12.7.



**Figure 11 - Error-free negotiation message sequences.**

#### 4.12.3 When to negotiate

Each port shall maintain a negotiation required flag for each other port.

A port shall set its negotiation required flags to true for all other ports after a reset event (see 12.5). A port shall set its negotiation required flag to true for a given port after an error occurs while transmitting a responding negotiation message to that port.

A SCSI initiator port shall set its negotiation required flag to true for a SCSI target port after an unexpected COMMAND phase occurs when selecting without using attention condition (i.e., when selecting a SCSI target port with information units enabled) (see 10.5.3).

A logical unit reset (see 16.5.6) has no effect on negotiation required flags or on transfer agreements.

After a reset event a port shall set its transfer agreements for all other ports to the default transfer agreement (see table 4).

A SCSI initiator port shall originate negotiation before sending a command to a SCSI target port whenever its negotiation required flag is true for that SCSI target port. A SCSI target port shall originate negotiation before accepting a command from a SCSI initiator port whenever its negotiation required flag is true for that SCSI initiator port. After successful negotiation or reaching the default transfer agreement, the negotiation required flag shall be set to false.

A port may originate negotiation even if its negotiation required flag is false (e.g., to change the settings, as part of integrity checking procedures, or, for a SCSI initiator port, after a SCSI target port has originated negotiation). Negotiation should not be originated after every selection and reselection as this may impact performance.

NOTE 7 - SCSI target ports may have had their support for originating negotiation after power on disabled to support illegal SCSI initiator device software. If a SCSI initiator port sends a command to a SCSI target device that has been powered on (e.g., after a hot plug) that results in a unit attention condition, the SCSI initiator port determines that negotiation is required and originates negotiation before the next command. However, if the command is INQUIRY, REPORT LUNS, or REQUEST SENSE, a unit attention condition is not created. An invalid data phase may occur if the SCSI target port does not originate negotiation. If the SCSI initiator port always originates negotiation before sending those commands, the data phase runs correctly. When information units are disabled, a SCSI initiator port may originate negotiation with its currently negotiated settings before each INQUIRY, REPORT LUNS, or REQUEST SENSE command to avoid this problem. When information units are enabled, the selection without attention results in an unexpected COMMAND phase that notifies the SCSI initiator port that negotiation before each INQUIRY, REPORT LUNS, or REQUEST SENSE command is not needed.

## 4.12.4 Negotiable fields

### 4.12.4.1 Negotiable fields introduction

Table 2 lists the fields that may be negotiated and the effects of successful negotiation on those fields by each of the different negotiation messages. Ports shall implement a given message if they implement fields that are negotiable with that message.

**Table 2 - Negotiable fields and effects of successful negotiation**

Field Name		Negotiation message pair		
		PPR	WDTR	SDTR
TRANSFER PERIOD FACTOR		negotiated (valid values: 08h-FFh)	No requirement	negotiated (valid values: 0Ah-FFh)
REQ/ACK OFFSET		negotiated	sets to 00h	negotiated
TRANSFER WIDTH EXPONENT		negotiated (valid values: 00h-01h)	negotiated (valid values: 00h-01h)	unchanged
Protocol options	PCOMP_EN	negotiated	sets to 0	sets to 0
	RTI	negotiated	sets to 0	sets to 0
	RD_STRM	negotiated	sets to 0	sets to 0
	WR_FLOW	negotiated	sets to 0	sets to 0
	HOLD_MCS	negotiated	sets to 0	sets to 0
	QAS_REQ	negotiated	sets to 0	sets to 0
	DT_REQ	negotiated	sets to 0	sets to 0
	IU_REQ	negotiated	sets to 0	sets to 0

When negotiating, the responding port shall respond with values that are a subset of the values in the originating message as indicated in table 3 (e.g., if the originating message requests a REQ/ACK offset of 10h, the responding message has a REQ/ACK offset field set to 10h or lower).

**Table 3 - Responding message requirements**

Field Name		Message	Response shall be numerically
TRANSFER PERIOD FACTOR		PPR, SDTR	Greater than or equal
REQ/ACK OFFSET		PPR, SDTR	Less than or equal
TRANSFER WIDTH EXPONENT		PPR, WDTR	Less than or equal
PROTOCOL OPTIONS	PCOMP_EN	PPR	0 or 1
	RTI	PPR	Less than or equal
	RD_STRM	PPR	Less than or equal
	WR_FLOW	PPR	Less than or equal
	HOLD_MCS	PPR	Less than or equal
	QAS_REQ	PPR	Less than or equal
	DT_REQ	PPR	Less than or equal
	IU_REQ	PPR	Less than or equal

#### 4.12.4.2 Transfer agreements

The transfer agreements that are in effect for various combinations of field values are described in table 4.

**Table 4 - Transfer agreements**

<b>Transfer agreement</b>	<b>REQ/ACK OFFSET</b>	<b>TRANSFER PERIOD FACTOR</b>	<b>TRANSFER WIDTH EXPONENT</b>	<b>DT_REQ</b>	<b>IU_REQ</b>	<b>QAS_REQ</b>	<b>All other protocol options</b>
default	00h	any	00h	0	0	0	0
asynchronous	00h	any	any	0	0	any	0
synchronous	GE 01h	GE 09h	any	any	any	any	any
ST synchronous	GE 01h	GE 0Ah	any	0	0	any	any
DT synchronous	GE 01h	GE 09h	01h	1	any	any	any
paced	GE 01h	08h or 07h	01h	1	1	any	any
wide	any	any	01h	any	any	any	any
narrow	any	any	00h	any	any	any	any
data group	GE 01h	any	01h	1	0	any	any
information unit	GE 01h	any	01h	1	1	any	any
ST data	GE 01h	GE 0Ah	any	0	any	any	any
DT data	GE 01h	any	01h	1	any	any	any
<sup>a</sup> Not all combinations are valid. See table 11 for valid field combinations.							

**4.12.4.3 Transfer period factor**

The TRANSFER PERIOD FACTOR field selects the transfer period (see 3.1.116) and determines which transfer rate's timing values in table 35, table 36, table 37, and table 38 shall be honored, provided that REQ/ACK OFFSET is greater than 00h. The field values are defined in table 5.

**Table 5 - Transfer Period Factor**

<b>Value</b>	<b>Description</b>	<b>Message</b>	<b>Transfer rate</b>
00h - 06h	Reserved <sup>c</sup>	N/A	N/A
07h	Transfer period equals 3,125 ns	PPR	Fast-320 <sup>a</sup>
08h	Transfer period equals 6,25 ns	PPR	Fast-160 <sup>a</sup>
09h	Transfer period equals 12,5 ns	PPR	Fast-80 <sup>a</sup>
0Ah	Transfer period equals 25 ns	PPR, SDTR	Fast-40 <sup>b</sup>
0Bh	Transfer period equals 30,3 ns	PPR, SDTR	Fast-40 <sup>b</sup>
0Ch	Transfer period equals 50 ns	PPR, SDTR	Fast-20 <sup>b</sup>
0Dh - 18h	Transfer period equals the TRANSFER PERIOD FACTOR x 4	PPR, SDTR	Fast-20 <sup>b</sup>
19h - 31h	Transfer period equals the TRANSFER PERIOD FACTOR x 4	PPR, SDTR	Fast-10 <sup>b</sup>
32h - FFh	Transfer period equals the TRANSFER PERIOD FACTOR x 4	PPR, SDTR	Fast-5 <sup>b</sup>
<sup>a</sup> See table 36, table 37, and table 38. Table 35 does not apply because DT transfers are required for this transfer period factor. <sup>b</sup> See table 35, table 36, table 37, and table 38. <sup>c</sup> Faster transfer periods may be defined by future standards.			

Table 6 shows which transfer period factors may be used with different types of transfer agreements, provided REQ/ACK OFFSET is greater than 00h.

**Table 6 - Transfer Period Factor relationships**

Value	Transfer agreement					
	Synchronous	Paced	Data group	Information unit	ST data	DT data
00h - 06h	reserved					
07h - 08h	NS	M	NS	M	NS	M
09h	M	NS	O	O	NS	M
0Ah	M	NS	O	O	O	O
0Bh	M	NS	O	O	O	O
0Ch	M	NS	O	O	O	O
0Dh - 18h	M	NS	O	O	O	O
19h - 31h	M	NS	O	O	O	O
32h - FFh	M	NS	O	O	O	O
Key: M=Mandatory: Support for the indicated transfer agreement shall be implemented if the indicated transfer period factor is implemented. O=Optional: Support for the indicated transfer agreement may be implemented if the indicated transfer period factor is implemented. NS=Not supported: The indicated transfer agreement shall not be allowed if the indicated transfer factor is selected.						

Table 11 defines valid combinations of TRANSFER PERIOD FACTOR and other fields.

#### 4.12.4.4 REQ/ACK offset

The REQ/ACK OFFSET field determines the maximum number of REQs allowed to be outstanding before a corresponding ACK is received at the SCSI target port during synchronous or paced transfers.

For ST synchronous transfers the REQ/ACK offset is the number of REQ assertions that may be sent by the SCSI target port in advance of the number of ACK assertions received from the SCSI initiator port.

For DT synchronous transfers the REQ/ACK offset is the number of REQ transitions that may be sent by the SCSI target port in advance of the number of ACK transitions received from the SCSI initiator port.

For paced transfers in DT DATA IN phase the REQ/ACK offset is the number of data valid state REQ assertions (see 10.7.4.3) that may be sent by the SCSI target port in advance of ACK assertions received from the SCSI initiator port.

For paced transfers in DT DATA OUT phase the REQ/ACK offset is the number of REQ assertions that may be sent by the SCSI target port in advance of the number of data valid state ACK assertions (see 10.7.4.3) received from the SCSI initiator port.

See 4.9 for an explanation of the differences between ST and DT data transfers.

The REQ/ACK OFFSET value is chosen to prevent overflow conditions in the port's receive buffer and offset

counter. The REQ/ACK OFFSET values and which timing values shall be selected are defined in table 7.

**Table 7 - REQ/ACK offset**

Value	Description	Timing values
00h	Specifies asynchronous transfer agreement. <sup>a</sup>	Asynch (See table 35)
01h-FEh	Synchronous or paced transfers with specified offset.	Determined by transfer period factor (See table 5)
FFh	Synchronous or paced transfers with unlimited offset.	Determined by transfer period factor (See table 5)
<sup>a</sup> Transfer period factor and protocol options other than QAS_REQ shall be ignored.		

Table 11 defines valid combinations of REQ/ACK OFFSET and other fields.

#### 4.12.4.5 Transfer width exponent

The TRANSFER WIDTH EXPONENT field defines the transfer width to be used during DATA IN and DATA OUT phases. The values are defined in table 8.

If any of the protocol options bits other than QAS\_REQ are set to one, then only wide transfer agreements are valid. If all the protocol options bits other than QAS\_REQ are set to zero, wide transfer agreements and narrow transfer agreements are valid.

**Table 8 - Transfer Width Exponent**

Value	Description
00h	Specifies 8 bit data bus (i.e., narrow transfer agreement).
01h	Specifies 16 bit data bus (i.e., wide transfer agreement).
02h	Obsolete
03h-FFh	Reserved

Table 11 defines valid combinations of TRANSFER WIDTH EXPONENT and other fields.

#### 4.12.4.6 Protocol options

##### 4.12.4.6.1 Protocol options introduction

The protocol options fields affect the protocol used between the ports.

The SCSI target port uses the protocol options bits to indicate to the SCSI initiator port if it agrees to enable the requested protocol options. Except for the PCOMP\_EN bit, the SCSI target port shall not enable any protocol options that were not enabled in the negotiation message received from the SCSI initiator port.

Table 9 lists the protocol options bits.



**Table 9 - Protocol options bits**

<b>Name</b>	<b>Description</b>
PCOMP_EN	Precompensation enable
RTI	Retain training information
RD_STRM	Read streaming and read flow control enable
WR_FLOW	Write flow control enable
HOLD_MCS	Hold margin control settings
QAS_REQ	QAS enable request
DT_REQ	DT clocking enable request
IU_REQ	Information units enable request

**4.12.4.6.2 IU\_REQ**

The SCSI initiator port shall set IU\_REQ to one in the PPR OUT message to request that information unit transfers (see 4.11.3.3) be enabled. In response, the SCSI target port shall set its IU\_REQ to one if it agrees to use information unit transfers or zero if it does not.

The SCSI initiator port shall set IU\_REQ to zero in the PPR OUT message to request that information unit transfers be disabled. In response, the SCSI target port shall set IU\_REQ to zero in the PPR IN message.

If IU\_REQ is set to one, an information unit transfer agreement is in effect. If IU\_REQ is set to zero, an asynchronous, ST synchronous, or data group transfer agreement is in effect.

Table 11 defines valid combinations of IU\_REQ and other fields.

Each SCSI target port shall maintain a bus free required flag. Each time a negotiation is successful that results in the IU\_REQ bit being changed from the previous agreement (i.e., zero to one or one to zero) the SCSI target port shall set its bus free required flag to true. Any intermediate changes (e.g., from multiple successful PPR negotiations) shall be treated as changing IU\_REQ even if the final value equals the initial value.

At the conclusion of the message phases, if the bus free required flag is set to true, the target port shall:

- 1) abort all tasks for the SCSI initiator port;
- 2) set the bus free required flag to false; and
- 3) go to a BUS FREE phase.

At the conclusion of the message phases, if the IU\_REQ bit was changed as the result of a negotiation, the SCSI initiator port shall abort all tasks routed through that target port.

If the IU\_REQ bit was set to one during a previous PPR negotiation and not changed by a subsequent PPR negotiation sequence, the target port shall not request that the task manager abort any tasks for that SCSI initiator port and shall go to the BUS FREE phase after responding with a PPR IN message.

Table 10 describes the bus phases resulting from IU\_REQ changes. See clause 14 for additional requirements.

**Table 10 - Bus phases resulting from IU\_REQ changes**

<b>Initial IU_REQ value</b>	<b>Modified IU_REQ value</b>	<b>Causes</b>	<b>Bus phase following MESSAGE phases</b>
0	0	a) PPR negotiation keeping IU_REQ set to zero; b) WDTR negotiation; or c) SDTR negotiation	COMMAND, DATA, STATUS, or BUS FREE phase
0	1	a) PPR negotiation setting IU_REQ to one	BUS FREE phase
1	0	a) PPR negotiation setting IU_REQ to zero; b) WDTR negotiation; or c) SDTR negotiation	BUS FREE phase
1	1	a) PPR negotiation keeping IU_REQ set to one	BUS FREE phase

**4.12.4.6.3 DT\_REQ**

The SCSI initiator port shall set DT\_REQ to one to request that DT DATA phases (see 4.11.3) be enabled. In response, the SCSI target port shall set DT\_REQ to one if it agrees to use DT DATA phases or zero if it does not.

The SCSI initiator port shall set DT\_REQ to zero to request that information unit transfers be disabled. In response, the SCSI target port shall set DT\_REQ to zero in the PPR IN message.

If DT\_REQ is set to one, a DT data transfer agreement is in effect. If DT\_REQ is set to zero, an asynchronous or ST data transfer agreement is in effect.

Table 11 defines valid combinations of DT\_REQ and other fields.

**4.12.4.6.4 QAS\_REQ**

The SCSI initiator port shall set QAS\_REQ to one to request that QAS (see 10.4.3) be enabled. In response, the SCSI target port shall set QAS\_REQ to one if it supports QAS or zero if it does not.

The SCSI initiator port shall set QAS\_REQ to zero to request that QAS be disabled. In response, the SCSI target port shall set QAS\_REQ to zero in the PPR IN message.

Table 11 defines valid combinations of QAS\_REQ and other fields.

When an initiator port and a target port have negotiated with each other to enable QAS, either of the two ports may participate in QAS arbitrations when attempting to connect to the other port. When an initiator port and target port have negotiated with each other to disable QAS, neither port shall participate in QAS arbitrations when attempting to connect to the other port.

When QAS and information unit transfers are both enabled for a connected SCSI target port, that SCSI target port may issue a QAS REQUEST message to release the bus after a DT DATA phase. When QAS is enabled for and information unit transfers are disabled for a connected SCSI target port, that SCSI target port shall not issue QAS REQUEST messages.

**4.12.4.6.5 HOLD\_MCS**

The SCSI initiator port shall set HOLD\_MCS to one to indicate that the SCSI target port should hold any

margin control settings set with the margin control subpage of the port control mode page (see 18.1.4). In response, the SCSI target port shall set HOLD\_MCS to one if it is capable of retaining the settings and zero if it is not.

The SCSI initiator port shall set HOLD\_MCS to zero to indicate that the SCSI target port shall reset to their default values any margin control settings set with the margin control subpage of the port control mode page (see 18.1.4). In response, the SCSI target port shall set HOLD\_MCS to zero.

Table 11 defines valid combinations of HOLD\_MCS and other fields.

#### **4.12.4.6.6 WR\_FLOW**

The SCSI initiator port shall set WR\_FLOW to one to indicate that the SCSI target port should enable write flow control during write streaming (see table 32, 4.11.3.3 and 8.2). In response, the SCSI target port shall set WR\_FLOW to one if it is capable of write flow control and zero if it is not.

The SCSI initiator port shall set WR\_FLOW to zero to indicate that the SCSI target port shall disable write flow control during write streaming. In response, the SCSI target port shall set WR\_FLOW to zero.

Write streaming and write flow control only occurs during information unit transfers.

Table 11 defines valid combinations of WR\_FLOW and other fields.

#### **4.12.4.6.7 RD\_STRM**

The SCSI initiator port shall set RD\_STRM to one to indicate that the SCSI target port should enable read streaming and read flow control (see table 32, 4.11.3.3, 8.2, and 14.3.4). In response, the SCSI target port shall set RD\_STRM to one if it is capable of read streaming and read flow control and zero if it is not.

The SCSI initiator port shall set RD\_STRM to zero to indicate that the SCSI target port shall disable read streaming and read flow control. In response, the SCSI target port shall set RD\_STRM to zero.

Read streaming and read flow control only occur during information unit transfers.

Table 11 defines valid combinations of RD\_STRM and other fields.

#### **4.12.4.6.8 RTI**

The SCSI initiator port shall set RTI to one to indicate it is capable of saving paced data transfer training information (see 10.7.4.2.1) and to indicate that the SCSI target port does not need to retrain on each connection. In response, the SCSI target port shall set RTI to one if it is capable of saving paced data transfer training information and zero if it is not.

The SCSI initiator port shall set RTI to zero to indicate it is not capable of saving paced data transfer training information and to indicate the SCSI target port shall retrain on each connection. In response, the SCSI target port shall set RTI to zero.

Table 11 defines valid combinations of RTI and other fields. For negotiated transfer periods slower than fast-160 the RTI bit shall be set to zero.

#### **4.12.4.6.9 PCOMP\_EN**

The SCSI initiator port that is negotiating for a fast-160 transfer period shall set PCOMP\_EN to one to indicate that the SCSI target port shall enable precompensation on all signals transmitted during DT DATA phases (see 4.8, 7.2.2, and 10.7.4.1). The SCSI initiator port shall set PCOMP\_EN to zero to indicate that the SCSI target port shall disable precompensation.

The SCSI target port that is negotiating for a fast-160 transfer period shall set PCOMP\_EN to one to indicate that the SCSI initiator port shall enable precompensation on all signals transmitted during DT DATA phases (see 4.8, 7.2.2, and 10.7.4.1). The SCSI target port shall set PCOMP\_EN to zero to indicate that the SCSI initiator port shall disable precompensation.

Table 11 defines valid combinations of PCOMP\_EN and other fields. Ports that have been successfully negotiated to a fast-160 transfer period shall support enabling and disabling precompensation of their drivers. For negotiated transfer periods other than fast-160 the PCOMP\_EN bit shall be set to zero.

NOTE 8 - Unlike other fields and bits in the PPR message the PCOMP\_EN bit is not a negotiated value; instead, it instructs the receiving SCSI device as to whether or not precompensation is to be disabled or enabled. Because of this, precompensation may be enabled on one of the SCSI devices and disabled on the other SCSI device at the completion of a successful PPR negotiation.

#### **4.12.5 Negotiable field combinations**

Not all combinations of the negotiable fields are valid. Only the combinations defined in table 11 shall be allowed. All other combinations of the listed fields are reserved.

Table 11 - Valid negotiable field combinations

TRANSFER PERIOD FACTOR	REQ/ACK OFFSET	TRANSFER WIDTH EXPONENT	Protocol options								Description
			PCOMP_EN	RTI	RD_STRM	WR_FLOW	HOLD_MCS	QAS_REQ	DT_REQ	IU_REQ	
ignore	00h	00h or 01h	0	0	0	0	0	0	0	0	Use ST DATA IN and ST DATA OUT phases to transfer data with asynchronous transfers
ignore	00h	00h or 01h	0	0	0	0	0	1	0	0	Use ST DATA IN and ST DATA OUT phases to transfer data with asynchronous transfers, and participate in QAS arbitrations
0Ah - FFh	01h - FFh	00h or 01h	0	0	0	0	0	0	0	0	Use ST DATA IN and ST DATA OUT phases to transfer data with synchronous transfers
09h - FFh	01h - FFh	01h	0	0	0	0	0	0	1	0	Use DT DATA IN and DT DATA OUT phases with data group transfers
09h - FFh	01h - FFh	01h	0	0	0	0	0	1	1	0	Use DT DATA IN and DT DATA OUT phases with data group transfers, and participate in QAS arbitrations
0Ah - FFh	01h - FFh	00h or 01h	0	0	0	0	0	1	0	0	Use ST DATA IN and ST DATA OUT phases to transfer data with synchronous transfers, and participate in QAS arbitrations
09h - FFh	01h - FFh	01h	0	0	0 or 1	0 or 1	0	0	1	1	Use DT DATA IN and DT DATA OUT phases with synchronous transfers and information unit transfers
07h	01h - FFh	01h	0	0 or 1	0 or 1	0 or 1	0 or 1	0	1	1	Use DT DATA IN and DT DATA OUT phases with paced transfers and information unit transfers
08h	01h - FFh	01h	0 or 1	0 or 1	0 or 1	0 or 1	0 or 1	0	1	1	Use DT DATA IN and DT DATA OUT phases with paced transfers and information unit transfers
09h - FFh	01h - FFh	01h	0	0	0 or 1	0 or 1	0	1	1	1	Use DT DATA IN and DT DATA OUT phases with synchronous transfers and information unit transfers, participate in QAS arbitrations, and issue QAS_REQUEST messages to initiate QAS arbitrations
07h	01h - FFh	01h	0	0 or 1	0 or 1	0 or 1	0 or 1	1	1	1	Use DT DATA IN and DT DATA OUT phases with paced transfers and information unit transfers, participate in QAS arbitrations, and issue QAS_REQUEST messages to initiate QAS arbitrations
08h	01h - FFh	01h	0 or 1	0 or 1	0 or 1	0 or 1	0 or 1	1	1	1	Use DT DATA IN and DT DATA OUT phases with paced transfers and information unit transfers, participate in QAS arbitrations, and issue QAS_REQUEST messages to initiate QAS arbitrations

#### 4.12.6 Message restrictions

PPR may be originated by SCSI initiator ports but shall not be originated by SCSI target ports.

If bus expanders are present, SCSI initiator ports should only use PPR when requesting values not attainable via WDTR and SDTR (e.g., setting any protocol option bits to one). If a SCSI target port responds to PPR only with values that are attainable via WDTR and SDTR (i.e., all protocol option bits set to zero), the SCSI initiator port should repeat negotiation with a WDTR and SDTR negotiation sequence. This ensures that bus expanders that do not support PPR are still able to handle data phases correctly.

WDTR and SDTR may be originated by either SCSI target ports or SCSI initiator ports. Since WDTR resets all the values that SDTR sets (see 4.12.4.1), it shall be sent first if both are needed.

SCSI target ports capable of wide transfer agreements shall originate negotiation with WDTR followed with SDTR.

NOTE 9 - If IU\_REQ was set to one and a successful SCSI target port originated WDTR negotiation occurs, a BUS FREE phase generated because the SCSI target port detected that IU\_REQ was changed is indistinguishable from a BUS FREE phase generated because the target port was detecting parity errors on the WDTR OUT. Following the WDTR negotiation with an SDTR negotiation before the BUS FREE occurs ensures that the SCSI initiator port and SCSI target port both know that IU\_REQ has changed.

#### 4.12.7 Negotiation message sequences

##### 4.12.7.1 Negotiation message sequences overview

A SCSI initiator port originated negotiation sequence contains up to four steps:

- 1) SCSI initiator port's originating message;
- 2) SCSI target port response;
- 3) SCSI initiator port response; and
- 4) SCSI target port second response.

A SCSI target port originated negotiation sequence contains up to four steps:

- 1) SCSI target port's originating message;
- 2) SCSI initiator port response;
- 3) SCSI target port response; and
- 4) SCSI initiator port second response.

If the negotiation fails after a vendor-specific number of retries, the SCSI port originating the negotiation sequence may discontinue communication with the other SCSI port.

#### 4.12.7.2 SCSI initiator port originated PPR negotiation

Figure 12 shows how the SCSI initiator port shall respond to various SCSI target port responses to an originating PPR OUT. The SCSI initiator port shall maintain the previous transfer agreement unless otherwise indicated.

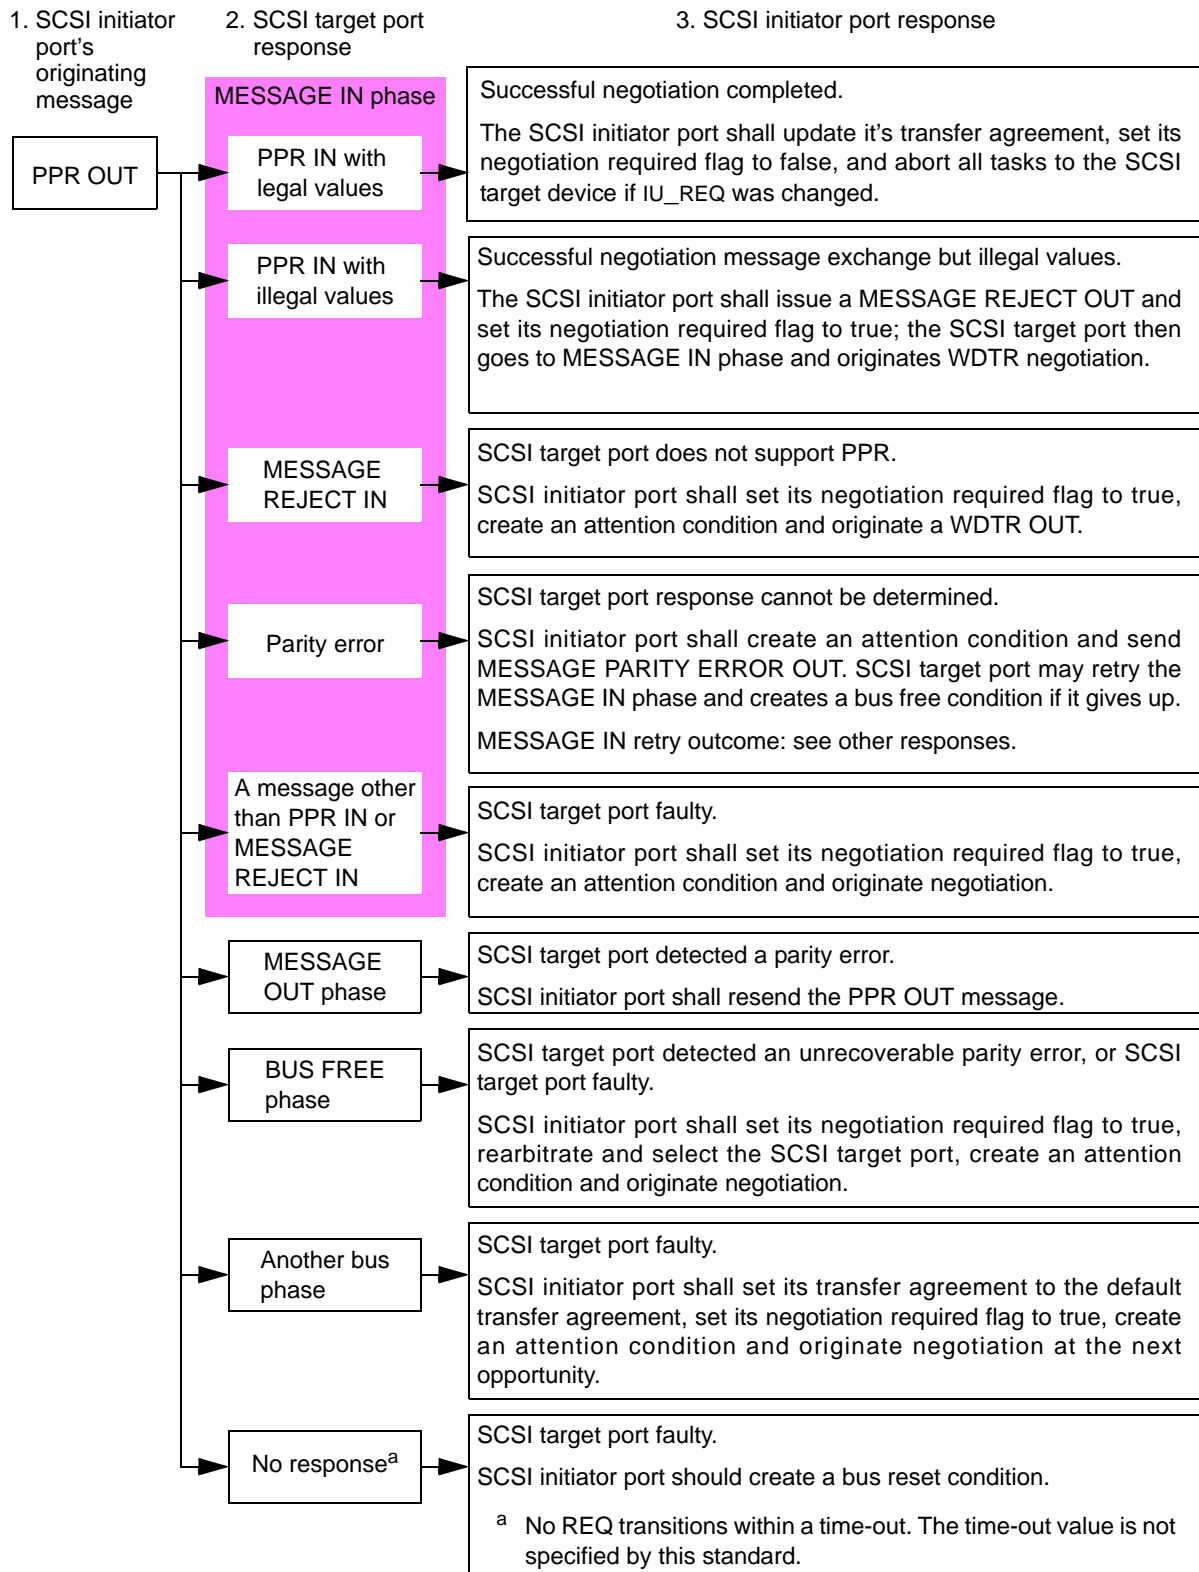


Figure 12 - SCSI initiator port originated PPR negotiation: SCSI initiator port response

Figure 13 shows how the SCSI target port shall respond to various SCSI initiator port responses to a responding PPR IN. The SCSI target port shall maintain the previous transfer agreement unless otherwise indicated.

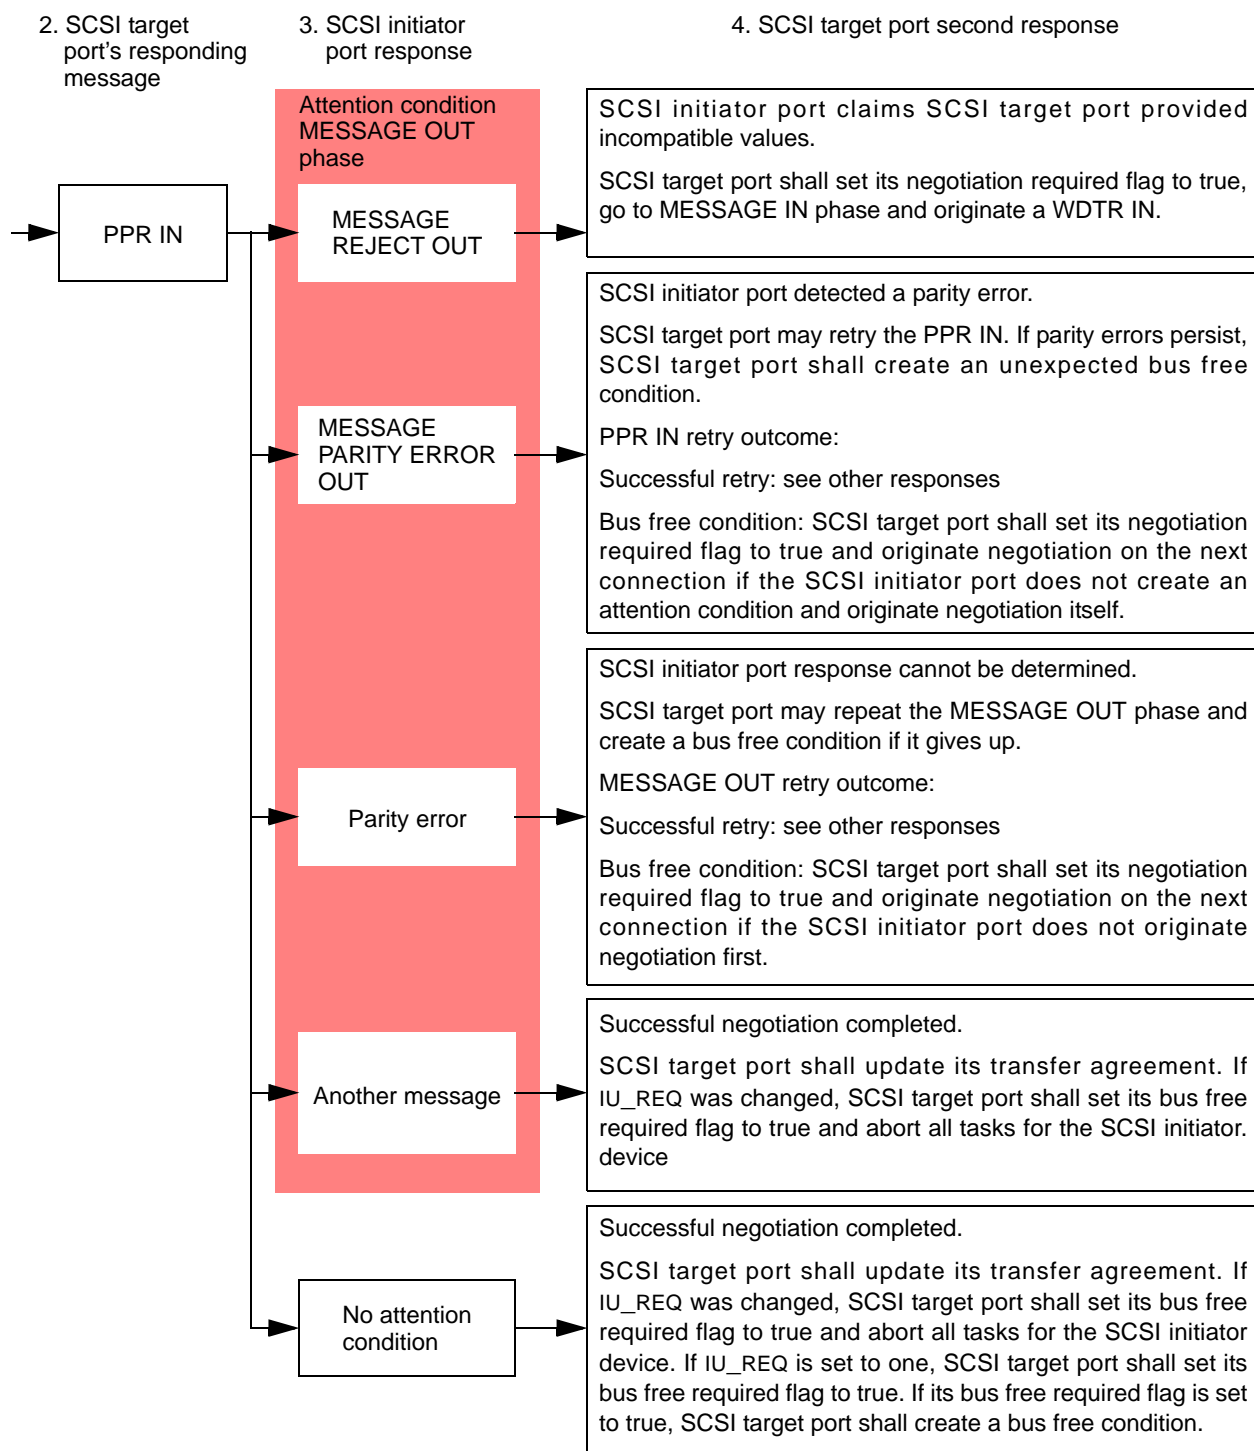


Figure 13 - SCSI initiator port originated PPR negotiation: SCSI target port response



#### 4.12.7.3 SCSI initiator port originated WDTR negotiation

Figure 14 shows how the SCSI initiator port shall respond to various target port responses to an originating WDTR OUT. The SCSI initiator port shall maintain the previous transfer agreement unless otherwise indicated.

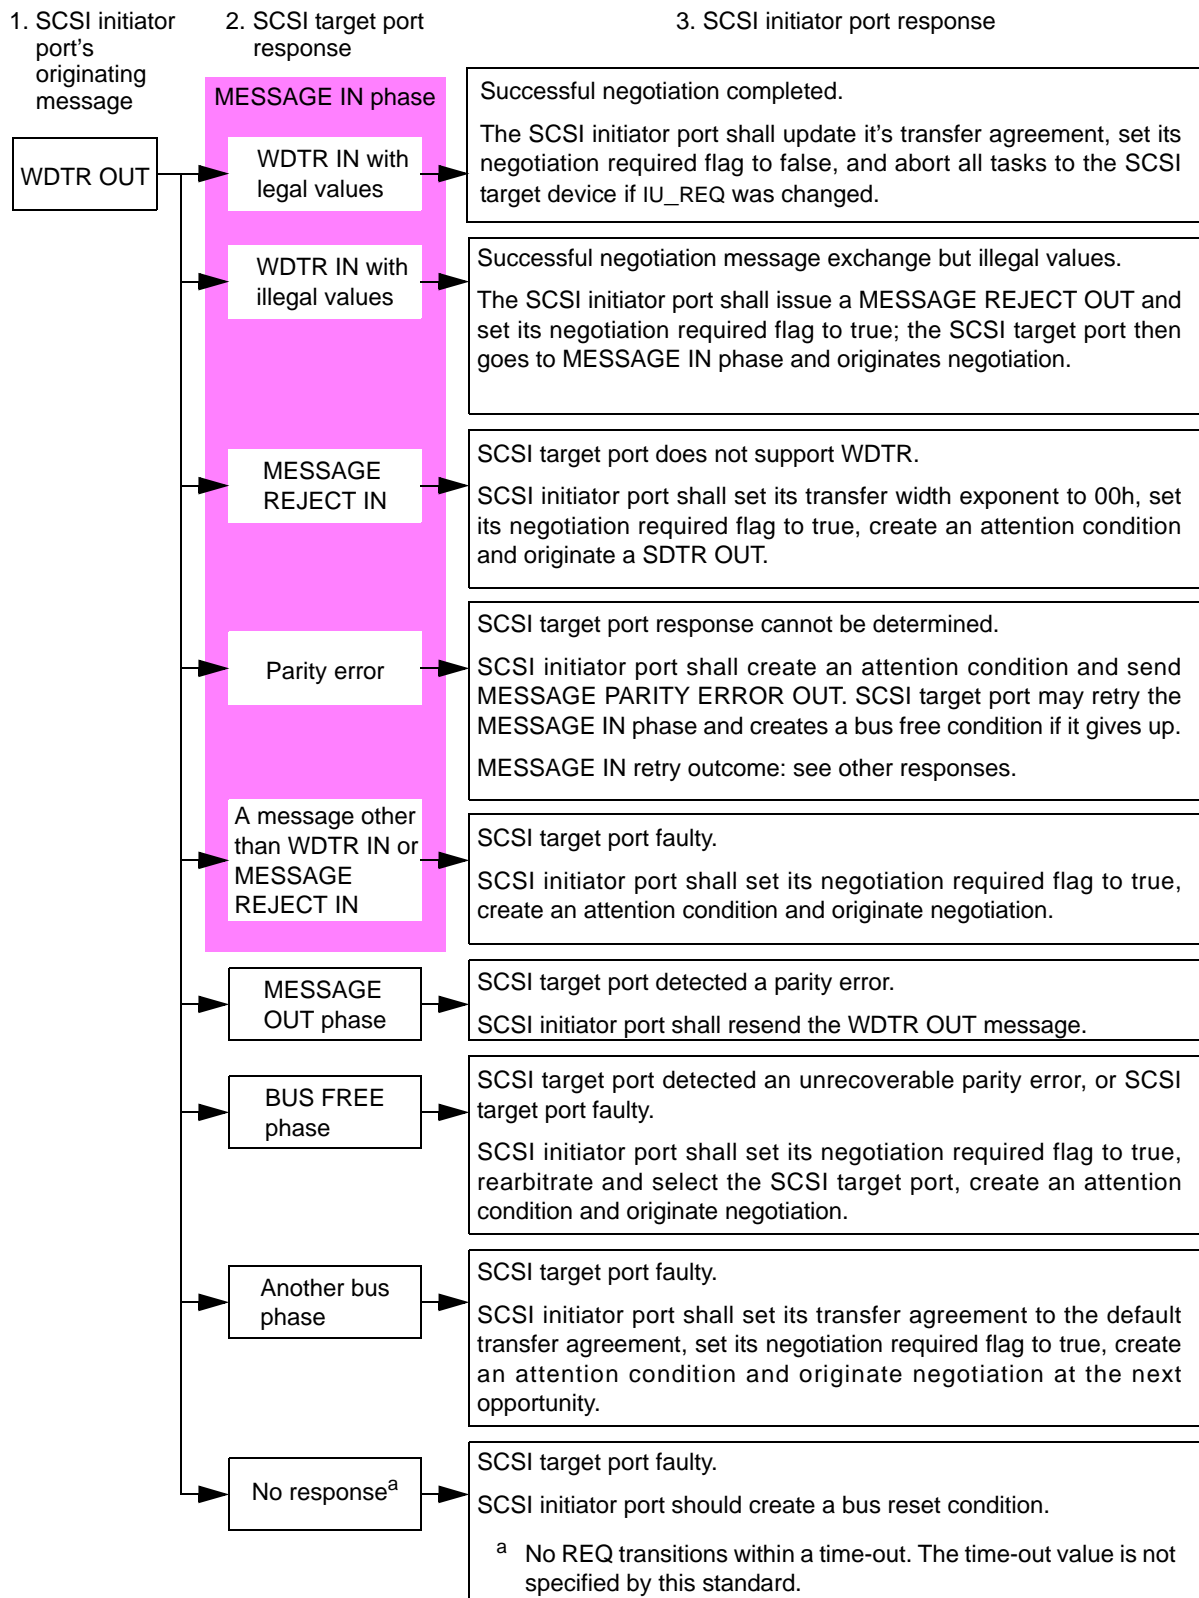


Figure 14 - SCSI initiator port originated WDTR negotiation: SCSI initiator port response

Figure 15 shows how the SCSI target port shall respond to various SCSI initiator port responses to a responding WDTR IN. The SCSI target port shall maintain the previous transfer agreement unless otherwise indicated.

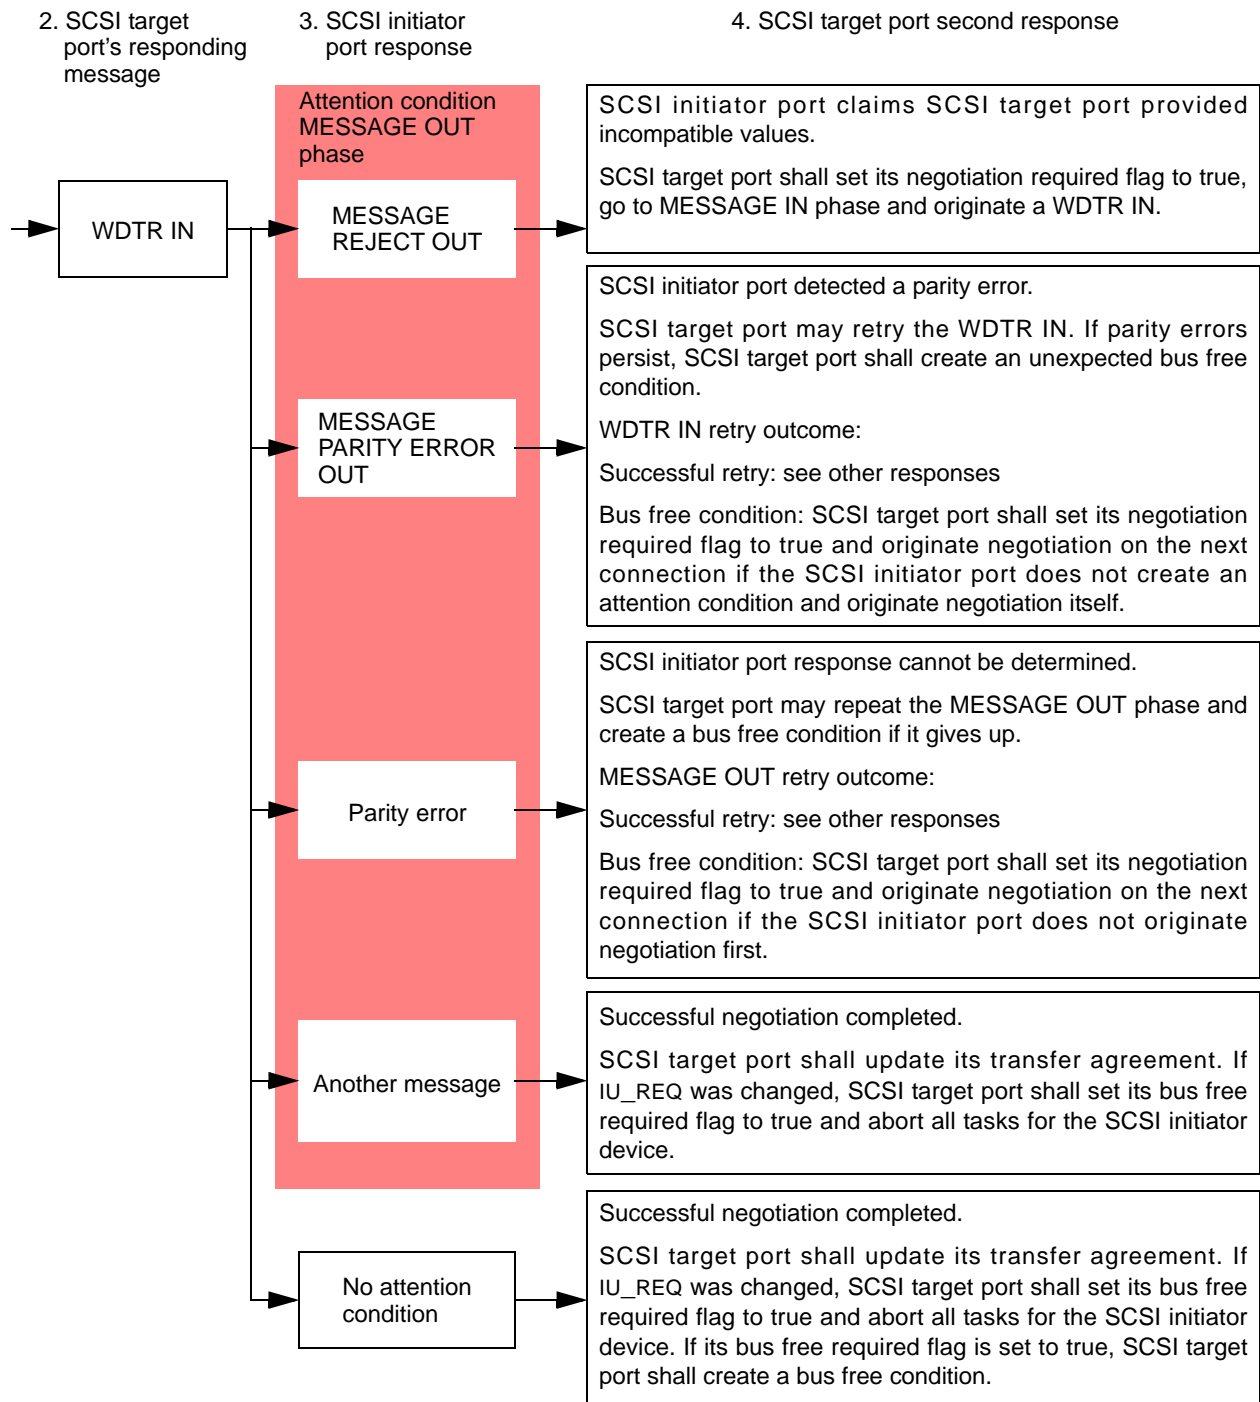


Figure 15 - SCSI initiator port originated WDTR negotiation: SCSI target port response

#### 4.12.7.4 SCSI initiator port originated SDTR negotiation

Figure 16 shows how the SCSI initiator port shall respond to various SCSI target port responses to an originating SDTR OUT. The SCSI initiator port shall maintain the previous transfer agreement unless otherwise indicated.

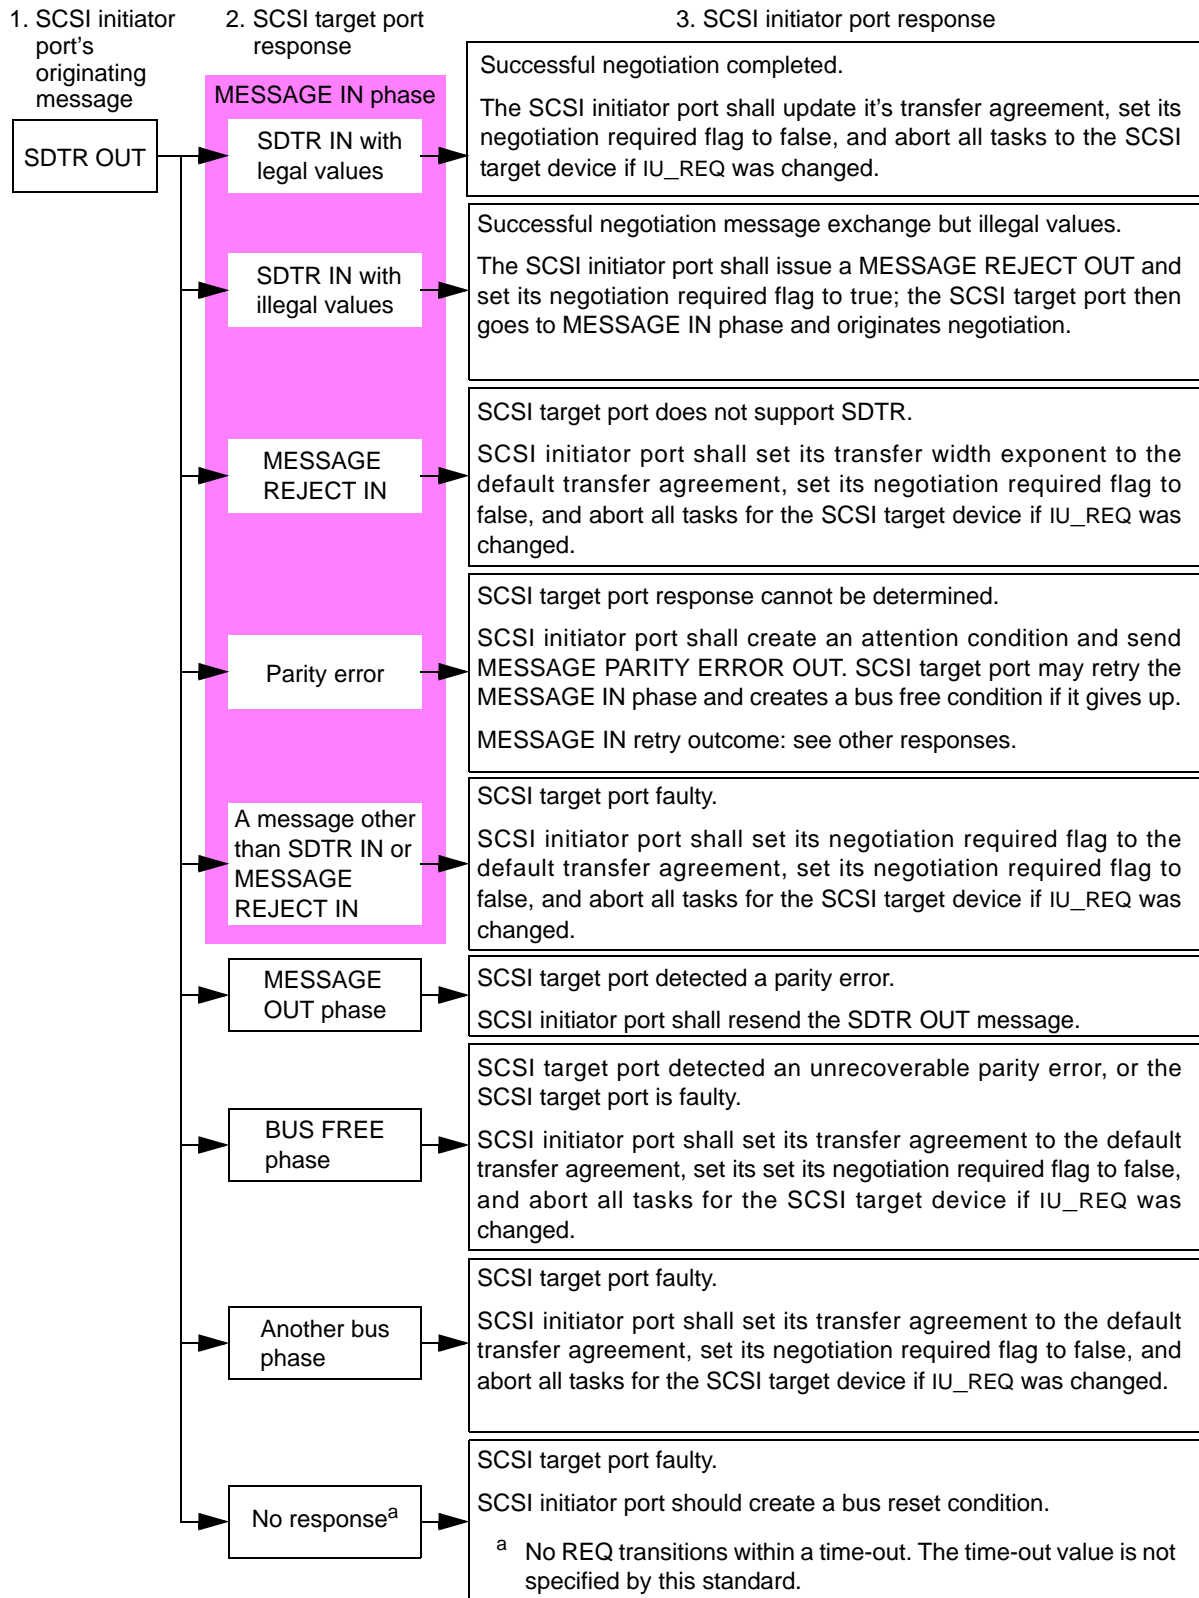


Figure 16 - SCSI initiator port originated SDTR negotiation: SCSI initiator port response

Figure 17 shows how the SCSI target port shall respond to various SCSI initiator port responses to a responding SDTR IN. The SCSI target port shall maintain the previous transfer agreement unless otherwise indicated.

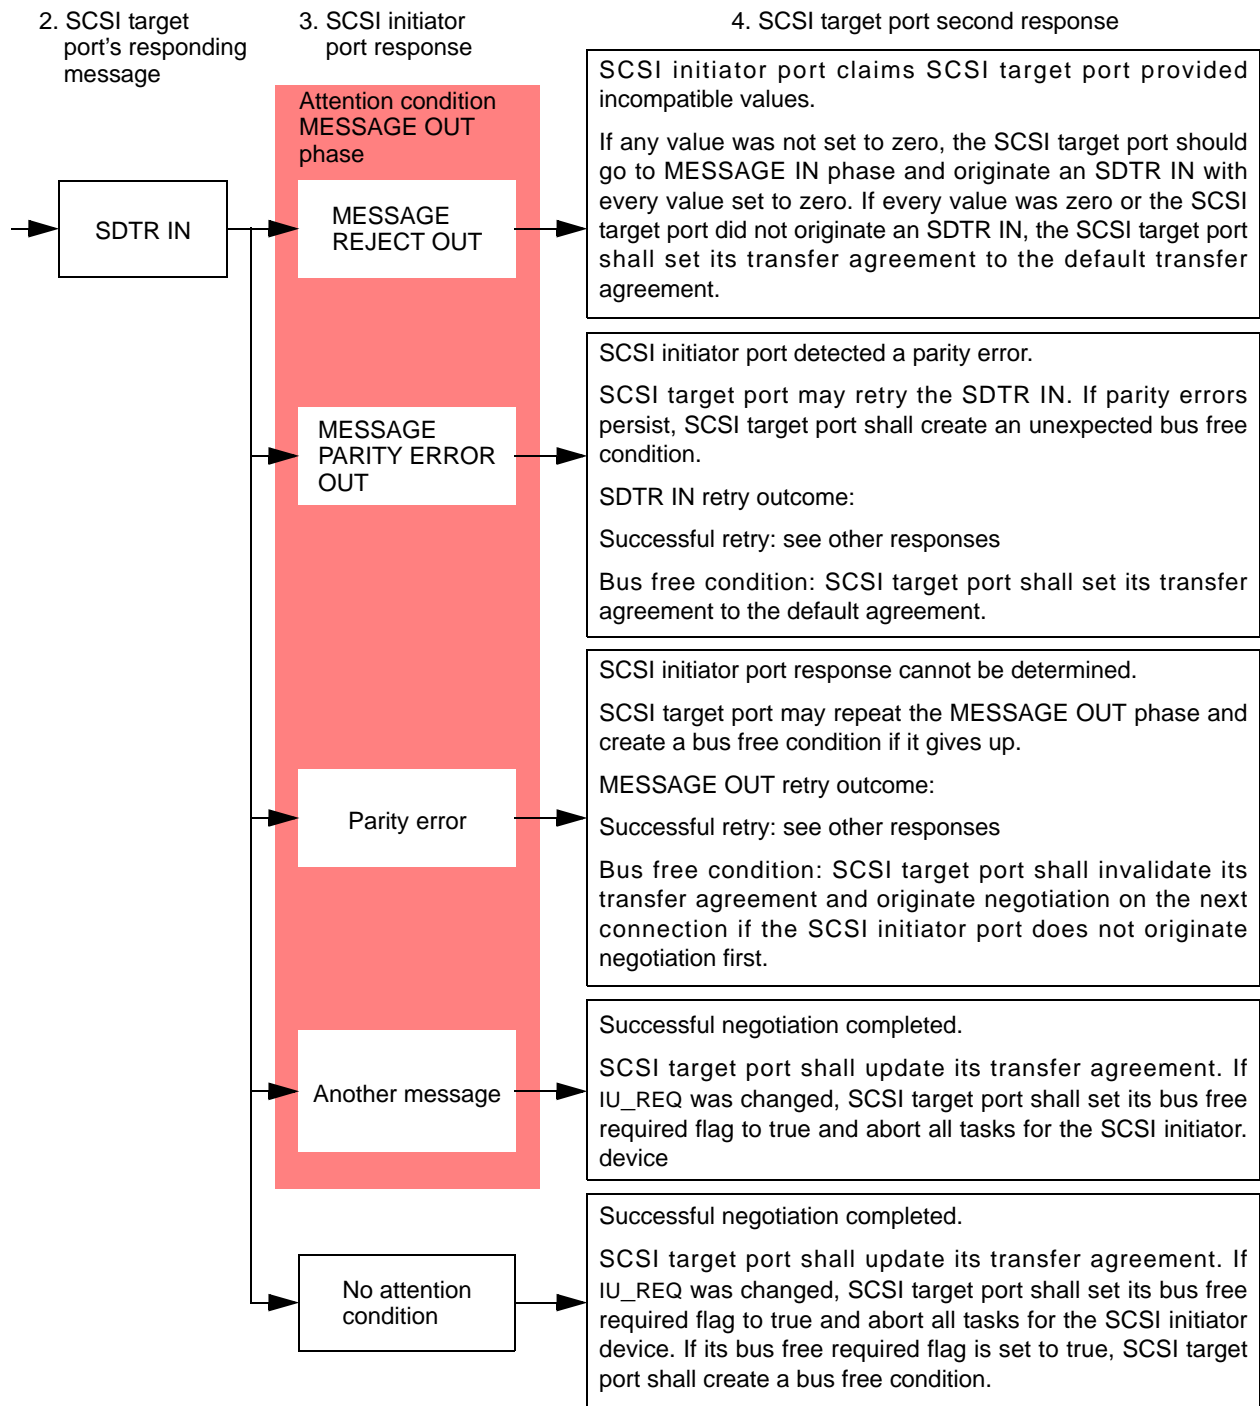


Figure 17 - SCSI initiator port originated SDTR negotiation: SCSI target port response

#### 4.12.7.5 SCSI target port originated WDTR negotiation

Figure 18 shows how the SCSI target port shall respond to various SCSI initiator port responses to an originating WDTR IN. The SCSI target port shall maintain the previous transfer agreement unless otherwise indicated.

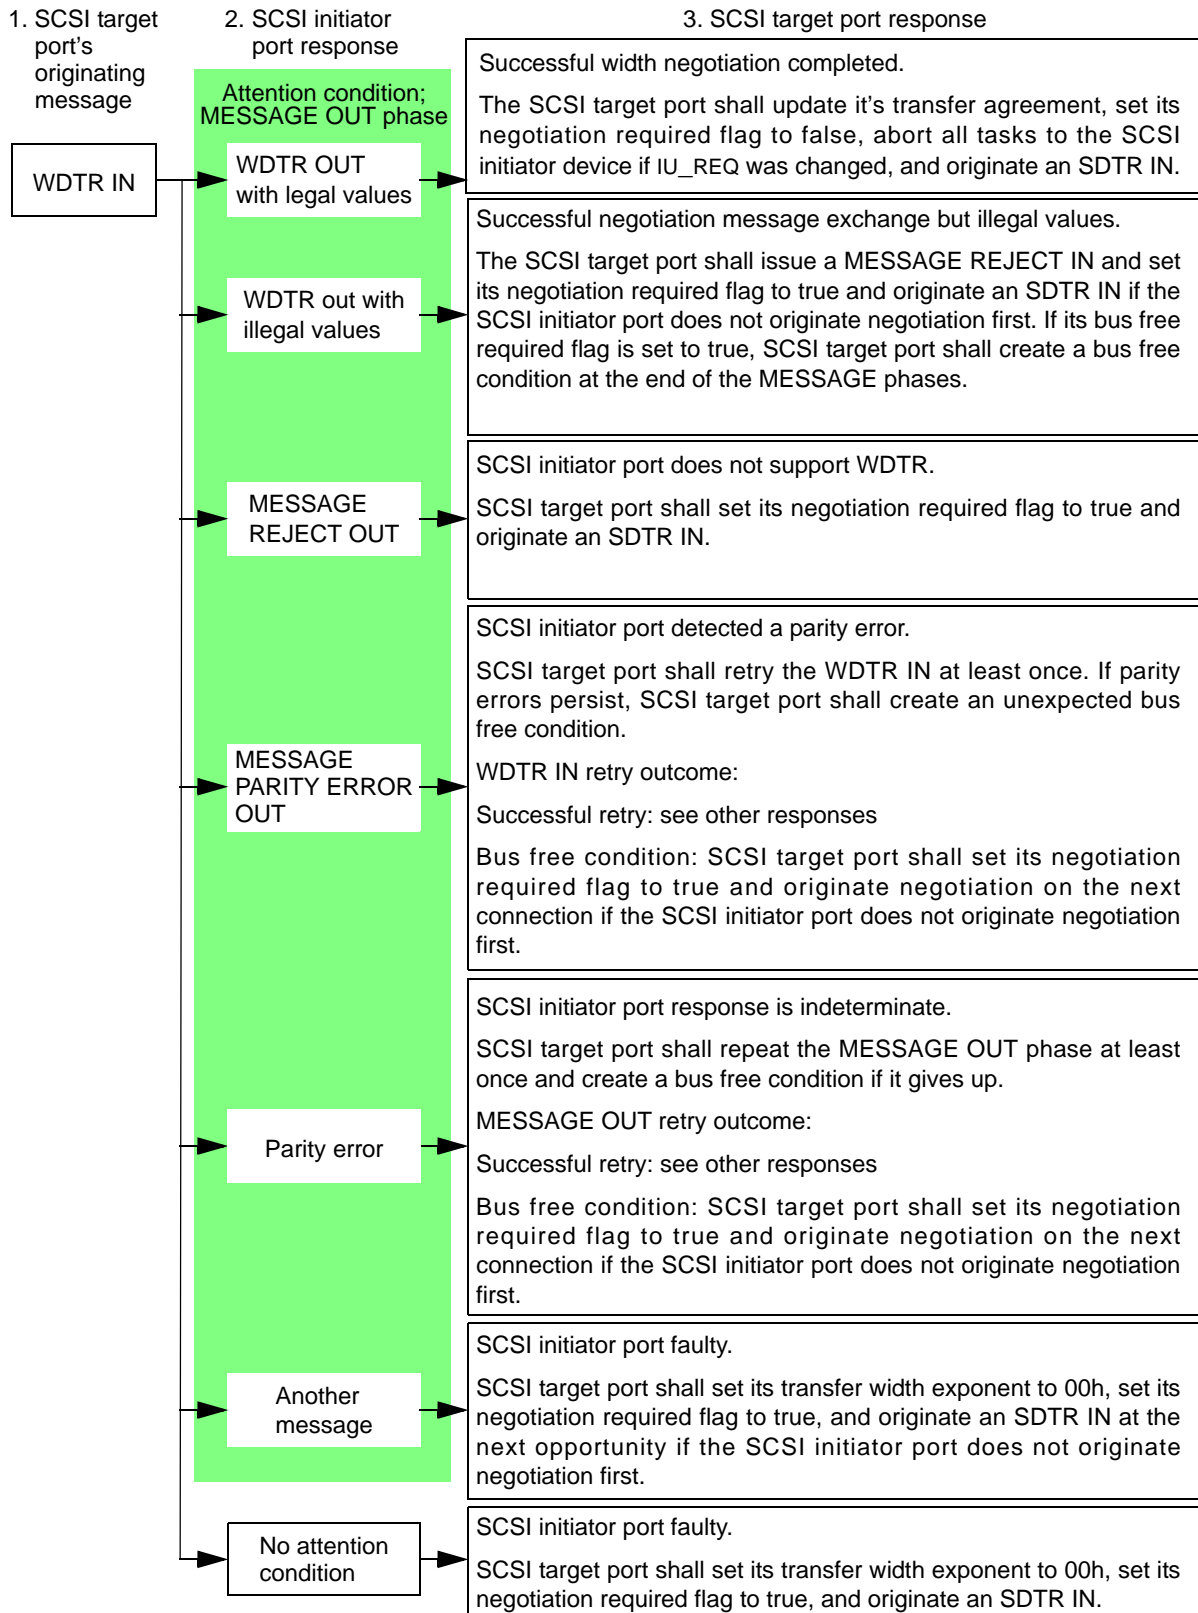


Figure 18 - SCSI target port originated WDTR negotiation: SCSI target port response

Figure 19 shows how the SCSI initiator port shall respond to various SCSI target port responses to a responding WDTR OUT. The SCSI initiator port shall maintain the previous transfer agreement unless otherwise indicated.

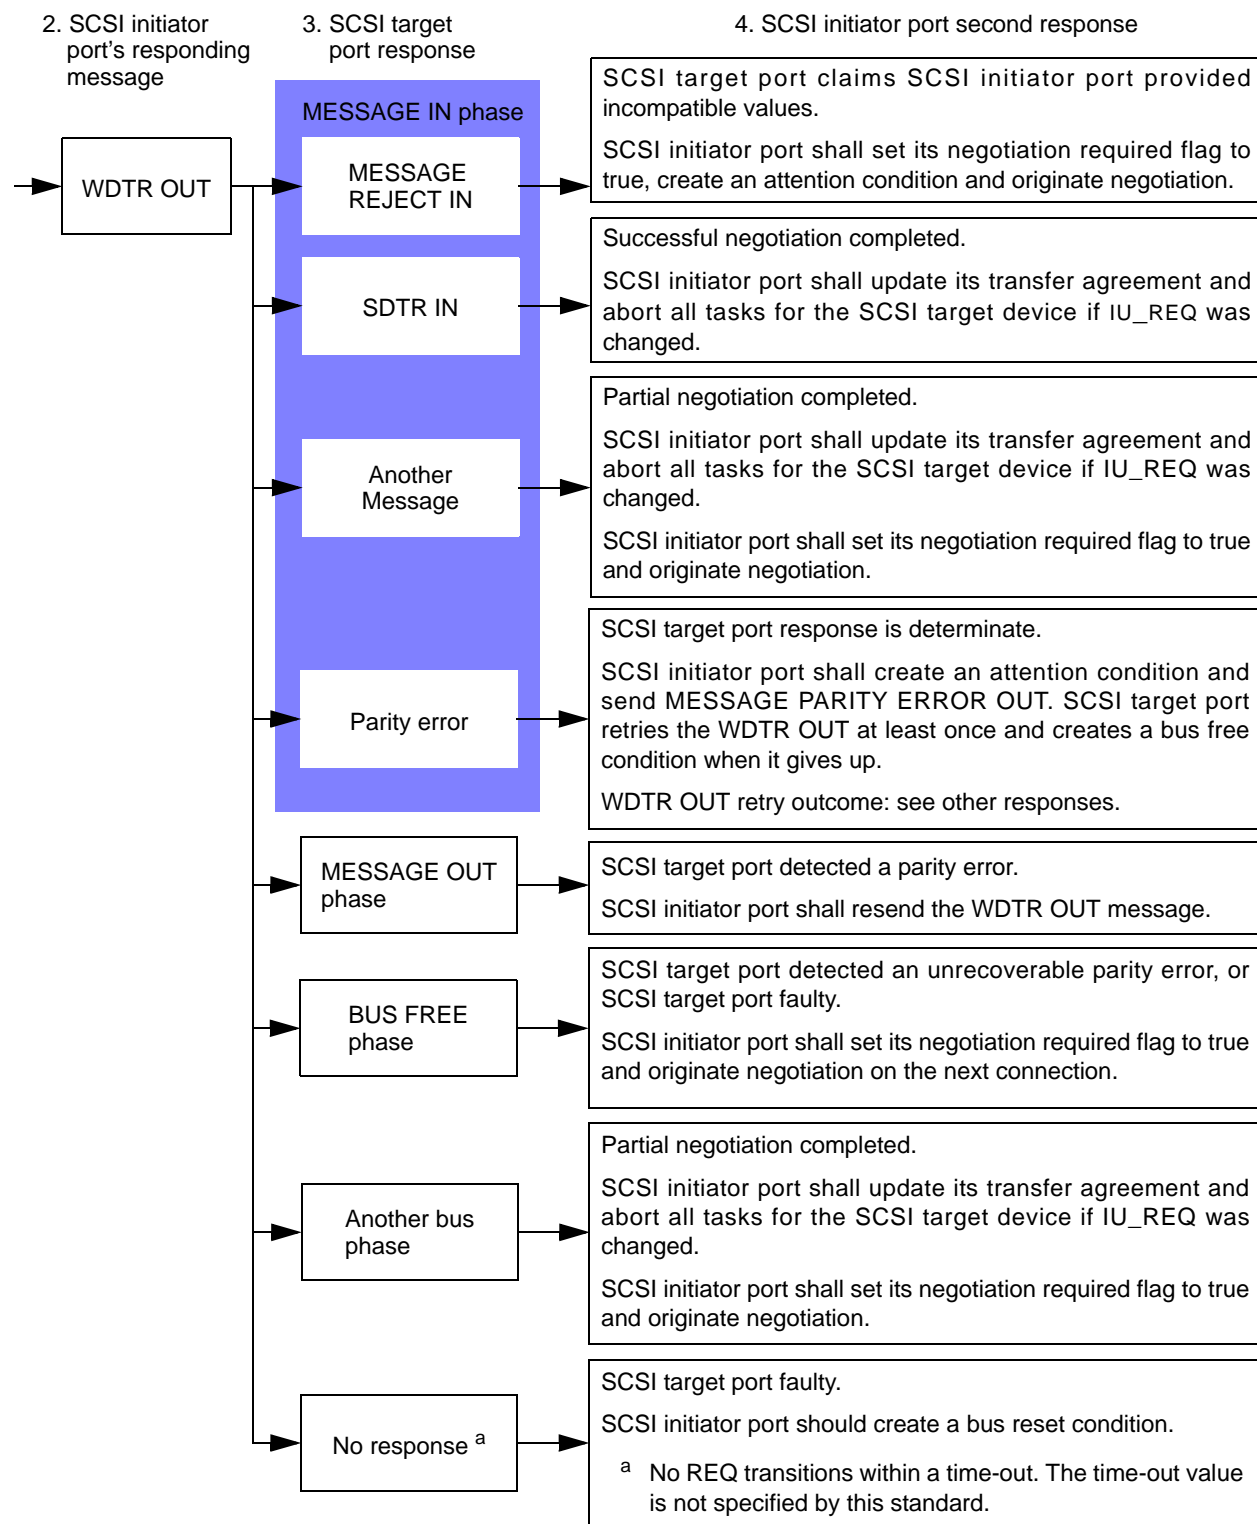


Figure 19 - SCSI target port originated WDTR negotiation: SCSI initiator port response

#### 4.12.7.6 SCSI target port originated SDTR negotiation

Figure 20 shows how the SCSI target port shall respond to various SCSI initiator port responses to an originating SDTR IN. The SCSI target port shall maintain the previous transfer agreement unless otherwise indicated.

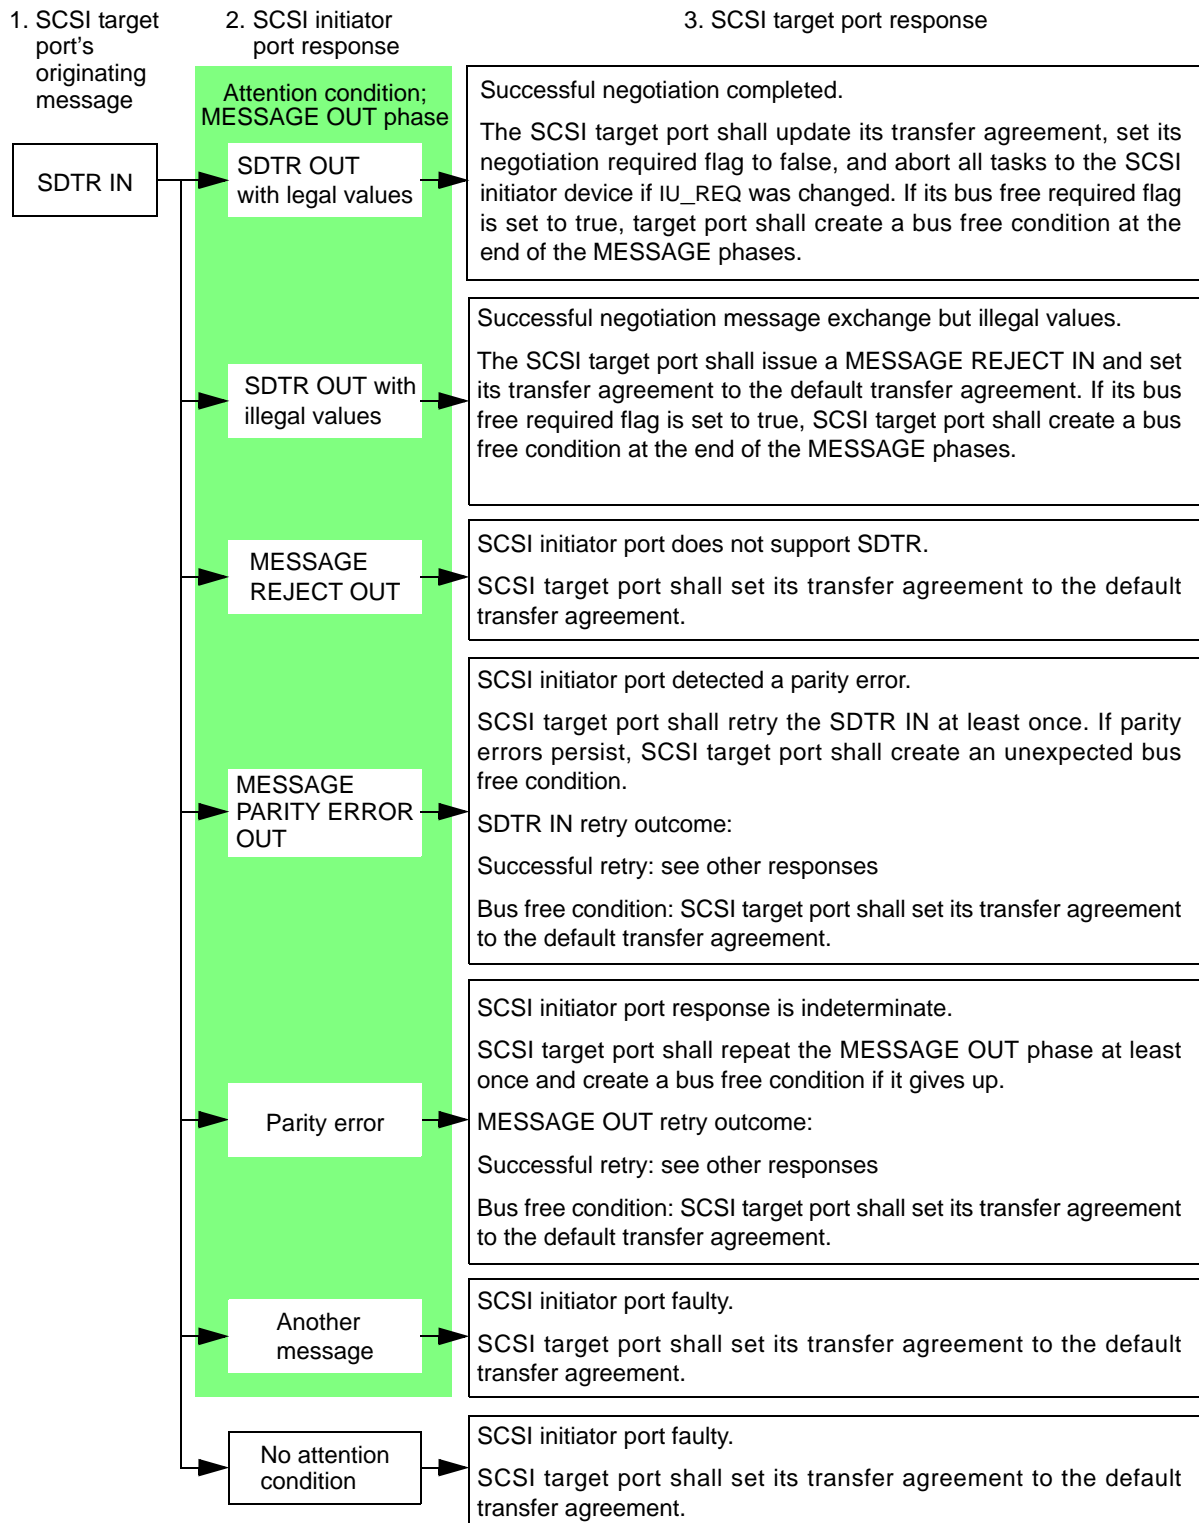


Figure 20 - SCSI target port originated SDTR negotiation: SCSI target port response

Figure 21 shows how the SCSI initiator port shall respond to various SCSI target port responses to a responding SDTR OUT. The SCSI initiator port shall maintain the previous transfer agreement unless otherwise indicated.

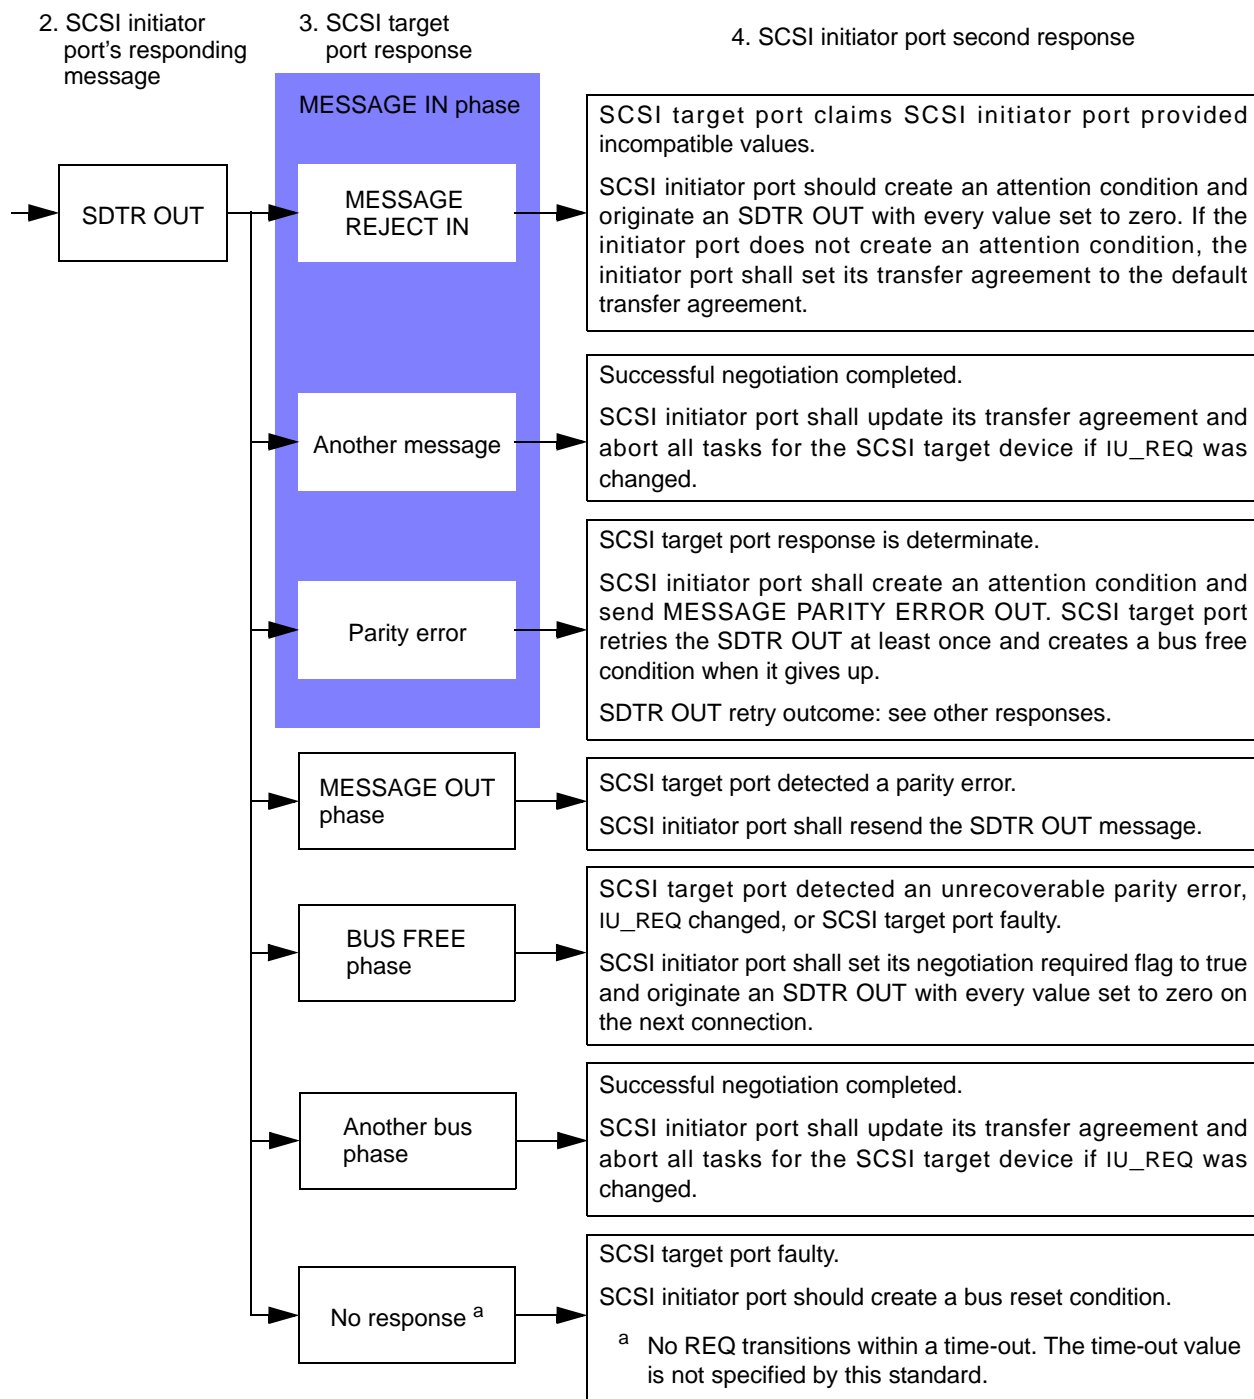
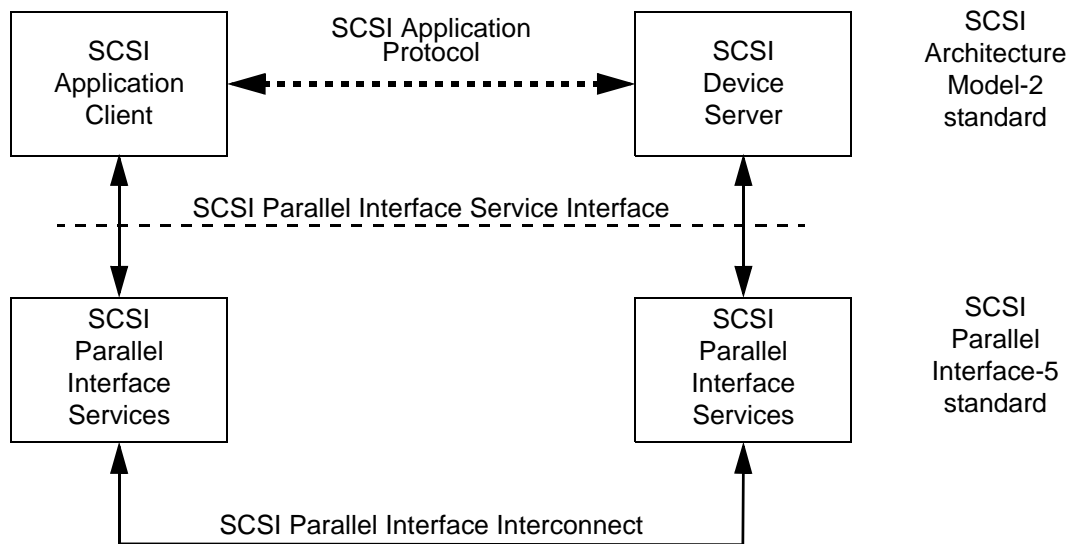


Figure 21 - SCSI target port originated SDTR negotiation: SCSI initiator port response



### 4.13 Protocol

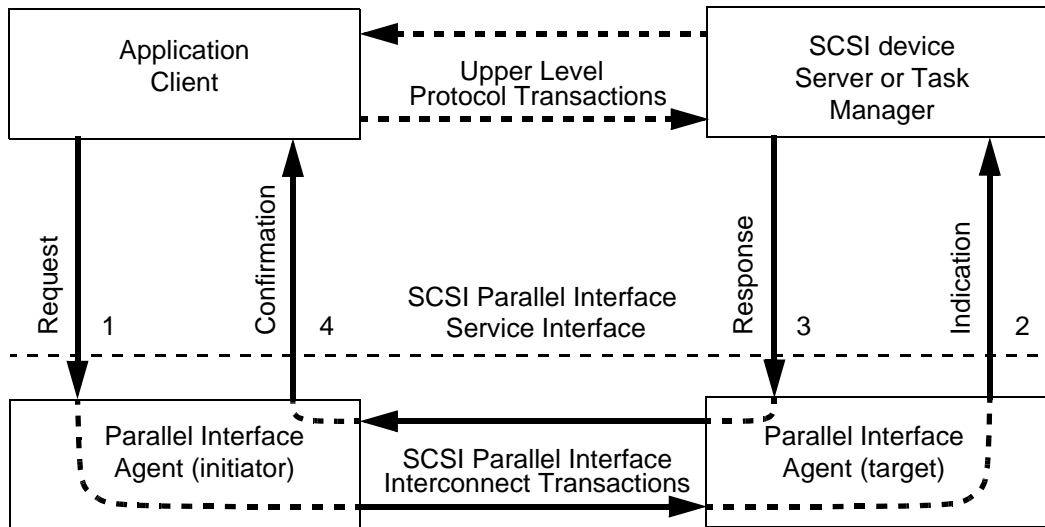
This standard describes a SCSI device's behavior in terms of functional levels, service interfaces between levels and peer-to-peer protocols. For a full description of the model used in this standard see the SCSI Architecture Model-2 standard. Figure 22 shows the model as it appears from the point of view of this standard.



**Figure 22 - 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 23 shows the service and protocol interactions for a four step confirmed service.



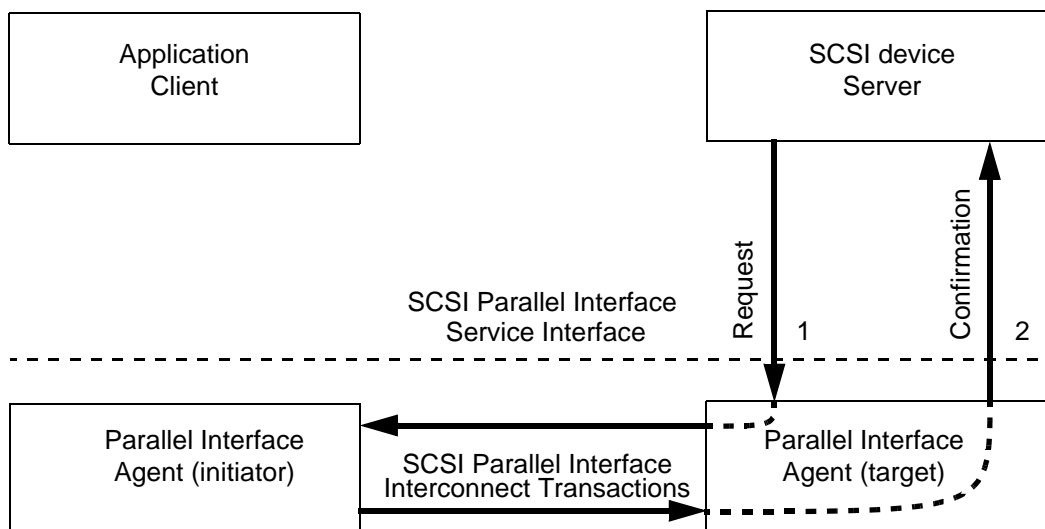
**Figure 23 - Model for a four step confirmed service**

The SCSI parallel interface four step confirmed service protocol consists of the following interactions:

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

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

Figure 24 shows the service and protocol interactions for a two step confirmed service.



**Figure 24 - Model for a two step confirmed service**

The SCSI parallel interface two step confirmed 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 SCSI 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.

The connector mechanical drawings conform to the ISO 1660, technical drawings - dimensioning and tolerancing standard.

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 clause 2).

The alternative 1 nonshielded, alternative 3 nonshielded, alternative 1 shielded, and alternative 3 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 512-2, 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) Measure contact resistance in accordance with item a) above. This is an optional step;
- d) Expose mated connectors to mixed flowing gas consisting of  $10 \times 10^{-9}$  parts of chlorine,  $10 \times 10^{-9}$  parts of hydrogen sulfide,  $200 \times 10^{-9}$  parts of sulfur dioxide, and  $200 \times 10^{-9}$  parts of nitrogen dioxide for 20 days at 70 % relative humidity and 30 °C. Use IEC 512-11-7, 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 mΩ 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 (i.e., 0,05 in) apart, as shown in figure 25. 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 (i.e., 0,05 in) apart, as shown in figure 26. 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 (i.e., 0,1 in) apart as shown in figure 27. 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 (i.e., 0,1 in) apart as shown in figure 28. It is recommended that keyed connectors be used.

### 5.2.3 Nonshielded connector alternative 3 - P cable


The alternative 3 nonshielded SCSI device connector for the P cable shall be a 68-conductor connector consisting of two rows of 34 female contacts with adjacent contacts 1,27 mm (i.e., 0,05 in) apart, as shown in figure 25. The nonmating portion of the connector is shown for reference only.

The alternative 3 nonshielded mating connector for the P cable shall be a 68-conductor connector consisting of two rows of 34 male contacts with adjacent contacts 1,27 mm (i.e., 0,05 in) apart, as shown in figure 26. 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 for the P-cable shall be an 80-conductor connector consisting of two rows of ribbon contacts spaced 1,27 mm (i.e., 0,05 in) apart, as shown in figure 29 and figure 30. For the detailed dimensional drawings of this connector see the SCA-2 EIA standard EIA-700AOAE.

The alternative 4 nonshielded mating connector for the P-cable shall be an 80-conductor connector consisting of two rows of ribbon contacts spaced 1,27 mm (i.e., 0,05 in) apart, as shown in figure 29 and figure 30. For the detailed dimensional drawings of this connector see the SCA-2 EIA standard EIA-700AOAE and SCA-2 Unshielded Connections SFF-8451.




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 25 - 50/68-contact alternative 1/alternative 3 nonshielded SCSI device connector**  
(A cable/P cable)



**Figure 26 - 50/68-contact alternative 1/alternative 3 nonshielded mating connector**  
(A cable/P 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 27 - 50-contact alternative 2 nonshielded SCSI device connector (A cable)





Dimensions	mm	in	Comments
<i>C1</i>	2,54	0,100	Maximum
<i>C2</i>	60,96	2,400	
<i>C3</i>	2,54	0,100	
<i>C4</i>	3,30	0,130	
<i>C5</i>	32,38	1,275	
<i>C6</i>	68,07	2,680	
<i>C7</i>	6,10	0,240	
<i>C8</i>	7,62	0,300	
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 28 - 50-contact alternative 2 nonshielded mating connector (A cable)**




Figure 29 - 80-contact alternative 4 nonshielded SCSI device connector (P cable)




Figure 30 - 80-contact alternative 4 nonshielded contact positions (P cable)

## 5.3 Shielded connector

### 5.3.1 Shielded connector overview

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

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

### 5.3.2 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 (i.e., 0,05 in) apart, as shown in figure 31. 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 (i.e., 0,05 in) apart, as shown in figure 32. The nonmating portion of the connector is shown for reference only.

### 5.3.3 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 (i.e., 0,085 in) apart, as shown in figure 33. 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 (i.e., 0,085 in) apart, as shown in figure 34. The non-mating portion of the connector is shown for reference only.

### 5.3.4 Shielded connector alternative 3 - P cable

The alternative 3 shielded SCSI device connector for the P cable shall be a 68-conductor connector consisting of two rows of 34 female contacts with adjacent contacts 1,27 mm (i.e., 0,05 in) apart, as shown in figure 35. The nonmating portion of the connector is shown for reference only.


The alternative 3 shielded mating connector for the P cable shall be a 68-conductor connector consisting of two rows of 34 male contacts with adjacent contacts 1,27 mm (i.e., 0,05 in) apart, as shown in figure 36. 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).

### 5.3.5 Shielded connector alternative 4 - P cable


The alternative 4 shielded SCSI device connector for the P cable shall be a 68-conductor connector consisting of two rows of ribbon contacts spaced 0,8 mm (i.e., 0,0315 in) apart, as shown in figure 37 and figure 38. For the detailed dimensional drawings of this connector see the VHDCI EIA standard EIA-700AOAF.

The alternative 4 shielded mating connector for the P cable shall be a 68-conductor connector consisting of two rows of ribbon contacts spaced 0,8 mm (i.e., 0,0315 in) apart, as shown in figure 37 and figure 38. For the detailed dimensional drawings of this connector see the VHDCI EIA standard EIA-700AOAF and VHDCI Shielded Configurations SFF-8441.



Dimensions	Millimeters	Inches
B1	34,70	1,366
B2	5,54	0,218
B3	2,54	0,100
B4	1,27	0,050
B5	30,48	1,200
B6	15°	15°
B7	1,00 R	0,039 R
B8	0,61±0,05	0,024±0,002
B9	0,15	0,006
B10	0,86±0,10	0,034±0,004
B11	0,15	0,006
B12	0,05	0,002
B13	5,10±0,05	0,201±0,002
B14	5,00±0,13	0,197±0,005
B15	1,85 max.	0,073 max.
B16	1,50±0,03	0,059±0,001
B17	42,29±0,10	1,665±0,004

Figure 31 - 50-contact alternative 1 shielded SCSI device connector (A cable)



Dimensions	Millimeters	Inches
A1	34,85	1,372
A2	5,69	0,224
A3	2,54	0,100
A4	1,27	0,050
A5	30,48	1,200
A6	15°	15°
A7	1,04 R	0,041 R
A8	0,40±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	2,64 min.	0,104 min.
A16	0,25±0,13	0,010±0,005

**Figure 32 - 50-contact alternative 1 shielded mating connector (A cable)**


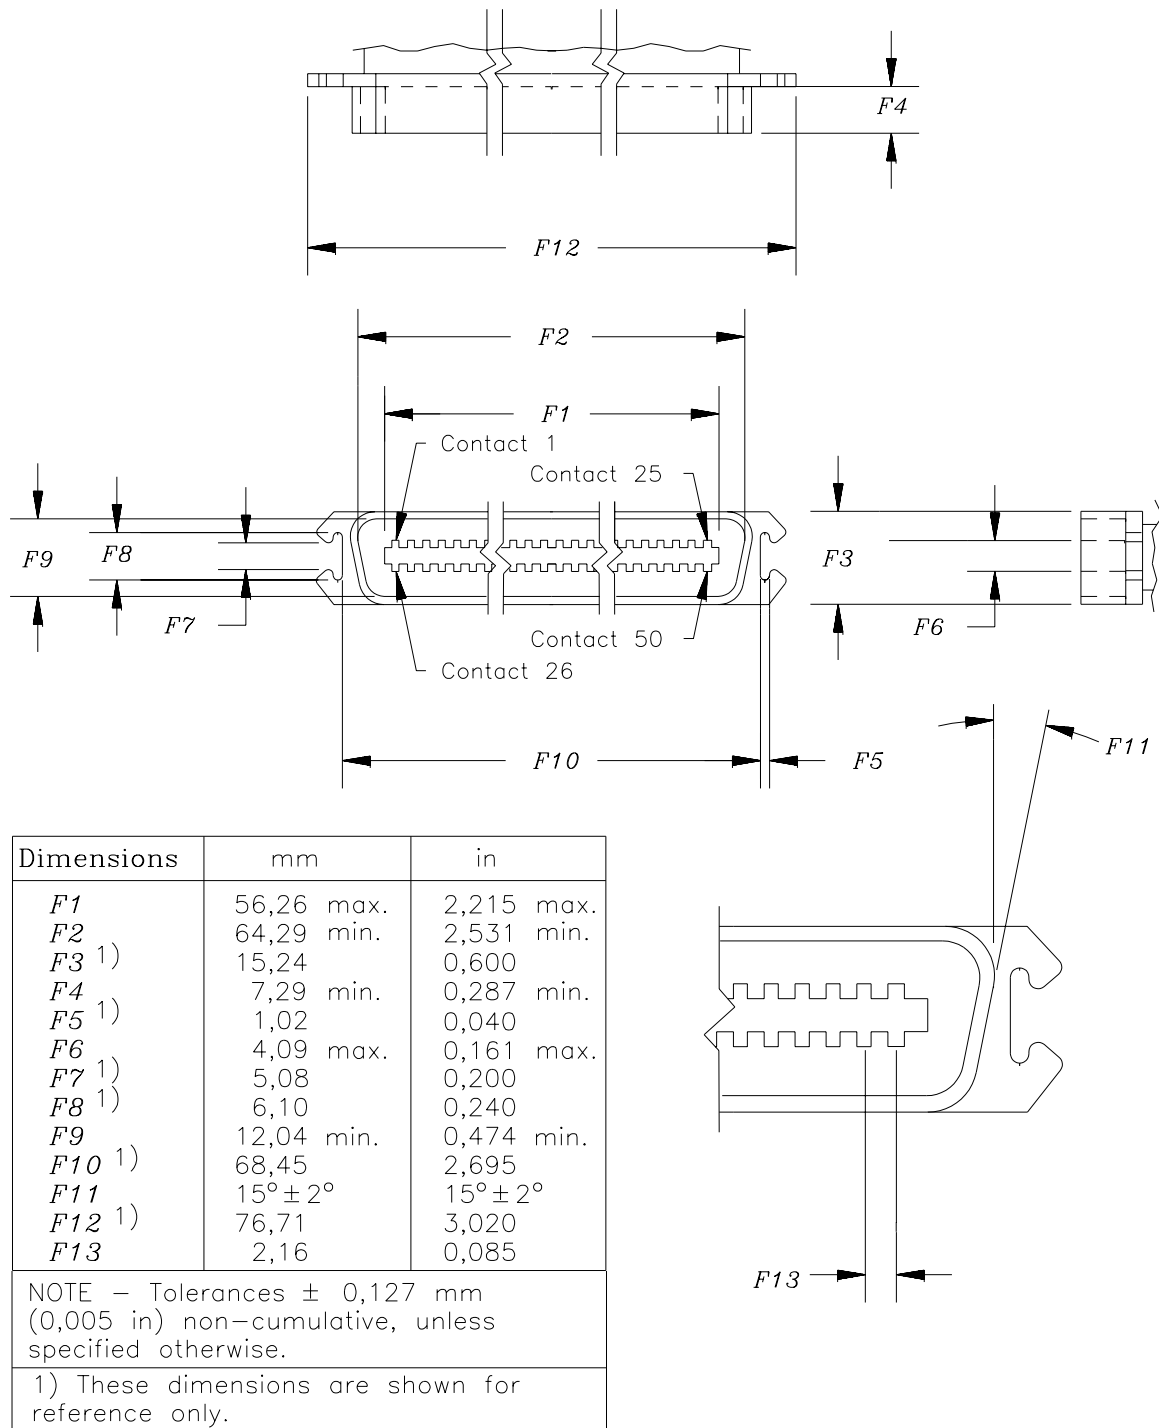


Figure 33 - 50-contact alternative 2 shielded SCSI device connector (A cable)



**Figure 34 - 50-contact alternative 2 shielded mating connector (A cable)**






Figure 35 - 68-contact alternative 3 shielded SCSI device connector (P 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 36 - 68-contact alternative 3 shielded mating connector (P cable)



**Figure 37 - 68-contact alternative 4 shielded SCSI device connector (P cable)**



**Figure 38 - 68-contact alternative 4 shielded contact positions (P cable)**

## 5.4 Connector contact assignments

### 5.4.1 Connector contact assignments overview

The connector contact assignments are defined in tables 12 through 15. See 8.2 for the definitions of the signals. The items under signal name labelled TERMPWR, and RESERVED are not signals and are not required to meet the cable characteristics for signals in 6.3. See 6.4 for characteristics of TERMPWR. See 6.5 for characteristics of RESERVED lines.

**Table 12 - Cross-reference to A cable contact assignments**

Connector type	Transmission mode	Connector figure	Contact assignment table	Contact set
Nonshielded alternative 1	LVD	25 and 26	13	2
Nonshielded alternative 2	LVD	27 and 28	13	1
Shielded alternative 1	LVD	31 and 32	13	2
Shielded alternative 2	LVD	33 and 34	13	1

### 5.4.2 Differential connector contact assignments

Table 13 defines the connector contact assignments for a 50 conductor bus segment that uses LVD transceivers.

**Table 13 - 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)
+P_CRCA	9	17	17	18	18	34	-P_CRCA
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
<p>The conductor number refers to the conductor position when using planar bulk cable.</p> <p>Two sets of contact assignments are shown, Refer to table 12 to determine which set of contacts applies to each connector.</p>							

Table 14 defines the connector contact assignments for a 68 conductor bus segment that uses LVD transceivers.

**Table 14 - 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)
+P_CRCA	14	27	28	48	-P_CRCA
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)
The conductor number refers to the conductor position when using planar bulk cable.					

Table 15 defines the connector contact assignments for an 80 conductor bus segment that uses LVD transceivers.

**Table 15 - LVD contact assignments - nonshielded alternative 4 connector**

Signal name	Host Pin Length <sup>b</sup>	Connector contact number	Connector contact number	Host Pin Length <sup>b</sup>	Signal name
12 V CHARGE <sup>a</sup>	Long	1	41	Long	12 V GROUND <sup>a</sup>
12 V <sup>a</sup>	Short	2	42	Long	12 V GROUND <sup>a</sup>
12 V <sup>a</sup>	Short	3	43	Long	12 V GROUND <sup>a</sup>
12 V <sup>a</sup>	Short	4	44	Short	MATED 1 <sup>a</sup>
3,3 V <sup>a</sup>	Short	5	45	Long	3,3 V CHARGE <sup>a</sup>
3,3 V <sup>a</sup>	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
-P_CRCA	Short	20	60	Short	+P_CRCA
-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)
5 V <sup>a</sup>	Short	34	74	Short	MATED 2 <sup>a</sup>
5 V <sup>a</sup>	Short	35	75	Long	5 V GROUND <sup>a</sup>
5 V CHARGE <sup>a</sup>	Long	36	76	Long	5 V GROUND <sup>a</sup>
SPINDLE SYNC <sup>a</sup>	Long	37	77	Long	ACTIVE LED OUT <sup>a</sup>
RMT_START <sup>a</sup>	Long	38	78	Long	DLYD_START <sup>a</sup>
SCSI ID (0) <sup>a</sup>	Long	39	79	Long	SCSI ID (1) <sup>a</sup>
SCSI ID (2) <sup>a</sup>	Long	40	80	Long	SCSI ID (3) <sup>a</sup>

<sup>a</sup> See Annex C for the definition of these signals.

<sup>b</sup> The pins identified as being short and long only applies to the cable/backplane-side connector and not the SCSI device-side connector. All pins on the SCSI device-side connector are the same length.

## 6 SCSI bus segment interconnect

### 6.1 SCSI bus segment interconnect overview

Clause 6 defines the characteristics of interconnects used to connect devices on the parallel SCSI bus. These interconnects are part of the SCSI path. See figure 3 for the topology of the interconnects that make up a SCSI path. Examples of types of Interconnects are:

- a) unshielded planar bulk cable;
- b) unshielded round twisted-pair bulk cable;
- c) shielded round twisted-pair bulk cable;
- d) backplanes.

The interconnect is defined as the electrical bulk cable, connectors, and passive loads used to connect the TERMPWR, terminators, expanders, and SCSI devices in a SCSI bus segment.

The functions of the interconnects are 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 SCSI devices, expanders, and terminators.

The interconnect shall meet the specified requirements for bulk cable and shall transport compliant worst-case transmitted signals in a manner that results in received signals that meet the requirements in clause 7. Signals for this requirement include DB(0) through DB(15), P\_CRCA, DB(P1), C/D, I/O, MSG, BSY, SEL, ATN, REQ, ACK, DIFFSENS, and RST. At least minimum TERMPWR shall be delivered to the terminator from minimum sources per the requirements in subclause 7.3.

The requirements on interconnect components in clause 6 are intended to aid in producing interchangeable components while achieving the required signal transmission properties.

### 6.2 SCSI bus segment cables

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

The following are recommendations that apply to SCSI round cables used with LVD transceivers:

- a) In the P cable conductor pairs ACK and REQ should be in the cable core;
- b) In the P cable, if there are more than four conductor pairs in the cable core, conductor pairs ACK and REQ should not be adjacent to each other;
- c) In the A cable conductor pairs ACK and REQ should be in the cable core;
- d) In the A cable, if there are more than three conductor pairs in the cable core, conductor pairs ACK and REQ should not be adjacent to each other;
- e) Cable conductor pairs used for the DATA BUS (DBnP1) and P\_CRCA should be in the outer layer of the cable;
- f) Each cable conductor pair should consist of the signal return and its associated signal.

Crosstalk noise is minimized by conductor placement (REQ and ACK in the center, data around the periphery) in round, twisted-pair cables and by the pin assignments on the connector on planar cables.

See Annex G for information on interconnecting bus segments of different widths.

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

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) (see Annex I). Detailed descriptions of these other means are not part of this standard; however, they are subject to the electrical requirements in clause 6.

A SCSI bus segment carries an 8-bit or 16-bit DATA BUS and the signals used to move information between SCSI devices.

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**

### **6.3.1 SCSI bulk cable**

The performance of the bulk cable is an important ingredient in the performance of the cable assembly, however 6.3 only specifies the uniform bulk cable which may not define a predictable performance of the interconnect. The SCSI Passive Interconnect Performance standard defines normative requirements for the cable assembly or backplane measured under specific conditions. Any normative requirements defined in the SCSI Passive Interconnect Performance standard that conflict with those in 6.3 shall take precedence over those specified in 6.3.

The test fixtures and test procedures for the measurements specified in clause 6 see the SCSI Passive Interconnect Performance standard.

The requirements in clause 6 apply to uniform bulk cable. Non-uniform bulk cable is bulk cable that contains dissimilar sections for purposes of enabling connector attachment.

Non-uniformities (e.g., a planar section created for connector attachment) are considered to be part of a cable assembly or harness whose performance is affected by the attached connectors, even if they are unused, as well as by the non-uniformity in the bulk cable.

Implementors using non-uniform bulk cable may construct special uniform test bulk cable using manufacturing processes similar to that used for the non-uniform bulk cable for purposes of measuring the properties of the bulk cable between the connector attachment areas (e.g., the twisted regions in a twisted/straight planar construction).

Only total, end to end, requirements are specified. Bulk cable intended to be used in cable assemblies that may be concatenated with other cable assemblies in the same bus segment shall meet the per unit length specifications. This reduces the risk that the concatenated cable assembly combination does not meet the signal transport requirements. Any bulk cable not meeting the per unit length requirements shall be labeled in a manner indicating that it may not be suitable for use in cable assemblies that are used in a concatenated manner.

Meeting the SCSI signal requirements in complete SCSI segments may require allowances beyond the uniform bulk cable requirements specified in clause 6. See the SCSI Passive Interconnect Performance standard.

The requirements in clause 6 apply to all the SCSI signals in the bulk cable except where otherwise specified.



### 6.3.2 Minimum conductor size for signals

The minimum conductor size for signals should be as specified in table 18.

### 6.3.3 Local transmission line impedance

The the differential transmission line impedance of the cable is defined in table 16.

**Table 16 - LVD local transmission line impedance**

Cable construction	Local differential transmission line impedance	
	Minimum	Maximum
All	110 $\Omega$	135 $\Omega$
All values are measured by time domain reflectometry		

### 6.3.4 Extended distance transmission line impedance

The swept frequency differential impedance limits used in extended distance transmission lines shall be a maximum peak-to-peak variation of 60  $\Omega$  over the frequency range 30 MHz to 600 MHz on a 30 m cable.

### 6.3.5 Capacitance

The capacitance limits for bulk cable are as shown in table 17.

**Table 17 - Bulk cable capacitance limits**

Minimum capacitance	Maximum capacitance	Frequency
26 pF/m	46 pF/m	100 kHz and 1 MHz

The dielectric constant variation in the material forming the insulation directly in contact with the conductors in the bulk cable between 300 kHz and 600 MHz shall be the less than 1,10 when the maximum dielectric constant in the frequency range divided by the minimum dielectric constant in the frequency range.

### 6.3.6 Differential propagation time and propagation time skew

The differential propagation time shall be no greater than 5,4 ns/m within the bulk cable and no greater than 135 ns from termination to terminator.

The differential propagation time for pair to pair skew shall be no greater than 82 ps/m within the bulk cable and no greater than 2,0 ns from terminator to terminator.

### 6.3.7 Differential attenuation

The attenuation requirements for differential attenuation are specified in table 18.

Table 18 - Attenuation requirements for SCSI bulk cable

Distance between SCSI bus segment terminators (m)	Attenuation per m maximum (dB) D.C. to 200 MHz	Attenuation of length equivalent to terminator to terminator distance maximum (dB) D.C. to 200 MHz	Distances are consistent with these minimum size conductors when used with high quality dielectrics	Notes
0 to 9 <sup>a</sup>	0,63 <sup>c</sup>	6 <sup>c</sup>	0,032 4 mm <sup>2</sup> (32 AWG) solid/ 0,050 92 mm <sup>2</sup> (30 AWG) stranded	multiple loads allowed
0 to 12 <sup>d</sup>	0,48 <sup>c</sup>	6 <sup>c</sup>	0,050 92 mm <sup>2</sup> (30 AWG) solid/ 0,080 42 mm <sup>2</sup> (28 AWG) stranded	multiple loads allowed
>12 to 25 <sup>b</sup>	0,48	12	0,050 92 mm <sup>2</sup> (30 AWG) solid/ 0,080 42 mm <sup>2</sup> (28 AWG) stranded	point to point only
<sup>a</sup> Twist and flat cable for fast-320 is restricted to 2 meters (limited by crosstalk) or additional testing using the Passive Interconnect Performance standard (see PIP). <sup>b</sup> Fast-320 restricts the length to 20 meters or additional testing (see PIP). <sup>c</sup> These number apply to speeds up to and including fast-160, for a definition of multiple loads interconnect testing for fast-320 see PIP. <sup>d</sup> Fast-320 restricts the lengths to 10 meters.				

### 6.3.8 Crosstalk

The maximum NEXT on REQ or ACK is 3,0 % of the 1,0 ns rise time aggressor signal amplitude. Crosstalk percent is calculated as follows:

$$\%NEXT = \frac{\sum \text{peak absolute differential induced voltages on REQ or ACK}}{\text{peak-to-peak differential aggressor voltage}}$$

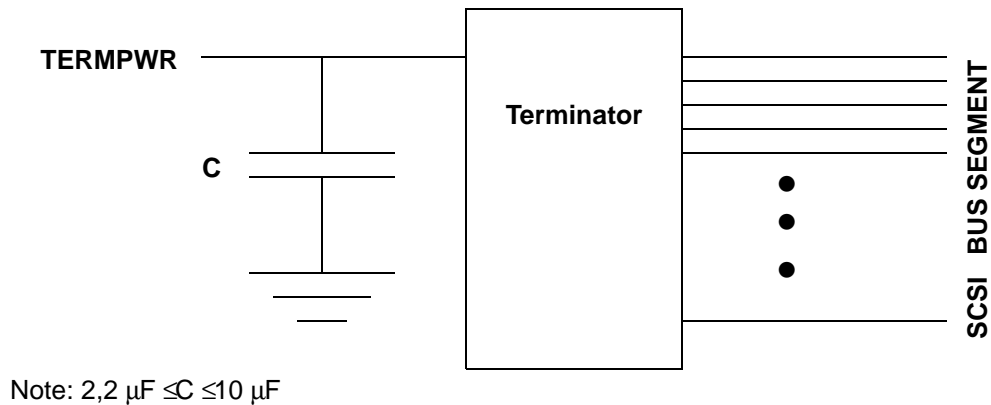
The aggressor signals are each of the DB(0-15), P\_CRCA, DB(P1), and REQ or ACK pairs. If REQ is the victim line DB(0-15), P\_CRCA, DB(P1), and ACK shall constitute the set of aggressor signals. If ACK is the victim line DB(0-15), P\_CRCA, DB(P1), and REQ constitute the set of aggressor signals. Each aggressor signal is separately excited, the induced absolute peak noise (i.e., deviation from zero differential) on the victim pair (see SCSI Passive Interconnect Performance standard).

NOTE 10 - 3.0 % NEXT yields 58,9 mV peak max at 1 963 mV peak-to-peak aggressor signal amplitude (i.e., 135  $\Omega$  maximum cable impedance at 7,3 mA max driver current). The crosstalk requirement is based only on percentage as that is all the cable influences.

### 6.4 Decoupling characteristics for TERMPWR lines

The TERMPWR lines should be decoupled at each terminator with at least a 2,2  $\mu$ F and not greater than a 10  $\mu$ F bypass capacitor.

See 7.3 and Annex D for additional information.



**Figure 39 - Terminator decoupling example**

## 6.5 Connection requirements for RESERVED lines

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

## 6.6 Cables used with LVD transceivers

Balanced interconnect conductors (e.g., twisted-planar, discrete wire twisted pairs, matched printed circuit board traces) should be used with LVD transceivers.

NOTE 11 - Use of unbalanced conductors such as planar untwisted construction typically produces higher crosstalk than balanced constructions but may be used if all electrical requirements are met.

The maximum distance between terminators when using LVD transceivers shall be as defined in table 19.

**Table 19 - LVD maximum bus segment path length between terminators**

Interconnect	Maximum bus segment path length between terminators (m) <sup>a, b</sup>						
	Fast-5	Fast-10	Fast-20	Fast-40	Fast-80	Fast-160	Fast-320
Point-to-point interconnect	25	25	25	25	25	25	20
Point-to-point twist and flat cable	d	d	d	d	d	d	3 <sup>c</sup>
Multidrop interconnect <sup>e</sup>	12	12	12	12	12	12	10
Multidrop twist and flat cable	d	d	d	d	d	d	2 <sup>c</sup>
<sup>a</sup> For environments where all elements of the bus segment (e.g., 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. <sup>b</sup> The maximum bus segment path lengths are achievable only if the receiver mask requirements are met (see 9.4). <sup>c</sup> Twist and flat cable technology at the time this standard was published. <sup>d</sup> Not defined in a previous SCSI standard. <sup>e</sup> Cable based.							

## 6.7 LVD stub length and spacing

The stub length when using LVD transceivers shall not exceed 0,1 m. The difference in stub length shall be less than 12,7 mm for the REQ, ACK, DB(15,0), P\_CRCA and DB(P1) signals. Stub length differences on the plus and minus signals of the same differential line should be minimized. The stub length is measured from the stub connection (see 4.5) to the end of the stub. It is recommended that the spacing of SCSI devices on the SCSI bus segment path be as indicated in table 20.

**Table 20 - Minimum stub connection spacing for LVD SCSI devices**

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

## 7 SCSI parallel interface electrical characteristics

### 7.1 SCSI parallel interface electrical characteristics overview

The SCSI parallel interface may use LVD transmitter implementations.

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

For measurements in clause 7, SCSI bus segment termination is assumed to be external to the SCSI device. See 6.5 for the termination requirements for the RESERVED lines. SCSI devices may have provision for allowing optional internal termination provided the internal termination conforms with 7.2.1 when enabled and the SCSI device, including the disabled termination, conforms with 7.2.4 when the internal termination is disabled.

In addition to the SCSI device electrical requirements defined in clause 7 SCSI devices shall meet the requirements specified in table 21 and table 22 at the device connector. The requirements in table 22 shall apply to both powered and unpowered SCSI devices.

**Table 21 - Absolute electrical limits at the device connector**

Mode	Minimum	Maximum	Notes
DIFFSENS for LVD/MSE input voltage	-0,5 V D.C.	4,1 V D.C.	Absolute maximum at all operating conditions for the DIFFSENS connection.
A combination of LVD SCSI device ports and HVD SCSI device ports on the same bus segment may result in damage to the DIFFSENS input of the LVD SCSI device port.			

**Table 22 - Input current requirements at the device connector for lines not being driven by the device**


Transceiver Type	Maximum	Notes
LVD	$\pm 20 \mu\text{A}$ D.C.	Measured from signal line $0 < V_{\text{IN}} < 2,5 \text{ V}$ to local ground for each signal pin.

### 7.2 LVD termination

#### 7.2.1 LVD termination overview


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.

The electrical characteristics of LVD bus segment termination shall be as specified in this subclause. Figure 40 shows the  $V_B$  and  $V_A$  measurement points, referenced to local ground, for the LVD bus segment terminator.



**Figure 40 - LVD bus segment terminator**

Figure 40 through figure 45 define the measurement points for the LVD terminators. Electrical characteristics shall meet the requirements in table 23 and table 24.



**CURRENT IS DRIVEN, VOLTAGE IS MEASURED**

**Figure 41 - Test circuit for terminator differential impedance**

The requirements that relate to differential impedance are specified in figure 42 and table 23. Table 23 specifies the allowed ranges for  $I_{diff}$  and  $V_{diff}$  in figure 41. The terminator bias voltage  $V_{BIAS}$  (i.e., the voltage measured when  $I = 0$  in figure 42) shall have the values measured between  $V_1$  and  $V_2$  as measured at  $V_{diff}$  in figure 41 with the range values defined in table 23 in the LVD impedance and  $V_{BIAS}$  tests column.




Figure 42 - Termination I-V characteristics for differential and common mode impedance tests




Figure 43 - Test circuit for termination common mode impedance test

The requirements that relate to common mode impedance are specified in figure 42 and table 23. Table 23 specifies the allowed ranges for  $I_{CM}$  and  $V_{CM}$  in figure 43. The terminator bias voltage  $V_{BIAS}$  (i.e., the voltage measured when  $I = 0$  in figure 42) shall have the values measured between  $V_1$  and  $V_2$  as measured at  $V_{CM}$  in figure 43 with the range values defined in table 23 in the common mode impedance and  $V_{BIAS}$  tests column.

**Table 23 - I-V requirements for differential impedance, common mode impedance, and  $V_{BIAS}$  tests**

Values (see figure 42)	Differential impedance and $V_{BIAS}$ tests <sup>a</sup> (see figure 41)	Common mode impedance and $V_{BIAS}$ tests (see figure 43)
$V_1$ (mV)	n/a	1125
$V_2$ (mV)	n/a	1375
$I_1$ (mA) <sup>d</sup>	1,0	n/a
$I_2$ (mA) <sup>e</sup>	1,1	n/a
$V_3$ (V)	1,0	2,0
$V_4$ (V)	-1,0	0,5
$I_{max}$ (mA)	9,00 <sup>f</sup>	N/A
$I_{min}$ (mA)	-11,25 <sup>f</sup>	N/A
$S_1$ ( $\Omega$ )	100 <sup>b</sup>	75 to 100 <sup>c</sup>
$S_2$ ( $\Omega$ )	110 <sup>b</sup>	300 to 400 <sup>c</sup>
Measurement	D.C.	D.C.
<sup>a</sup> $V_A + V_B = 2,5 \pm 0,2$ V (see figure 41) <sup>b</sup> The differential impedances of $S_1$ and $S_2$ is the open bus segment path value that shall be set to any value from 55 $\Omega$ to 130 $\Omega \pm 14\%$ for closed bus segment path. The difference between $S_1$ and $S_2$ shall not be greater than 10 $\Omega$ across the 27 lines. An open bus segment path is a segment that may be constructed by end users. A closed bus segment path is constructed by the manufacturer. <sup>c</sup> The common mode $S_1$ and $S_2$ impedances change with differential impedance changes such that the nominal $S_1$ (i.e., 100 $\Omega$ ) and nominal $S_2$ (i.e., 110 $\Omega$ ) differential is an $S_1$ common mode of 100 $\Omega$ and an $S_2$ common mode of 300 $\Omega$ <sup>d</sup> The $I_1$ value is calculated from the $V_1$ and $S_1$ values in figure 42 as follows: $I_1 = V_1 / S_1$ <sup>e</sup> The $I_2$ value is calculated from the $V_2$ and $S_2$ values in figure 42 as follows: $I_2 = V_2 / S_2$ <sup>f</sup> $I_{max}$ and $I_{min}$ are measured at the nominal differential impedance where $S_1$ is 100 $\Omega$ and $S_2$ is 110 $\Omega$		

Programmable terminators shall be adjustable from the nominal value of 105  $\Omega$  to a range between 55  $\Omega$  and 130  $\Omega$

The requirements on termination that relate to electrical balance are specified in figure 44, figure 45, and table 24. The voltage  $V_{test1}$  in figure 44 is varied over frequencies of 0 to 40 MHz with amplitude varied



over the range  $V_{\text{MIN}}$  to  $V_{\text{MAX}}$  specified in table 24 while the voltage named  $\Delta V$  in figure 44 is measured. The maximum difference between values of  $\Delta V$  (see figure 44) measured during this test shall be as specified in table 24.

NOTE 12 - The + signal line and - signal line capacitance should be balanced on disabled terminators.




Figure 44 - Termination balance test configuration




Figure 45 - Termination balance test data definition

**Table 24 - Values for LVD termination balance test**

Parameter	Value
$V_{\text{MIN}}$ ( $V_{\text{test1}}$ peak)	0,7
$V_{\text{MAX}}$ ( $V_{\text{test1}}$ peak)	1,8
$R$ ( $\Omega$ )	$100 \pm 0,01\%$
$\delta V$	20 mV max
$\Delta V$ - Input impedance for instrumentation > 10 K $\Omega$ $V_{\text{test1}}$ swept through all values between $V_{\text{MIN}}$ and $V_{\text{MAX}}$	

### 7.2.2 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 47. 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.

During paced transfers for fast-160 with precompensation enabled, the driver source current shall either be strong or weak to comply with precompensation requirements in table 25. The strong driver current is sourced to or from the SCSI data bus signal for the first data transfer time after a change in the data bus line state from asserted to negated or from negated to asserted. The weak driver current is sourced to or from the SCSI data bus signal after a data transfer time from the last data bus signal transition (see A.2.1). The rules for fast-160 precompensation shall apply to the REQ, ACK, P1, P\_CRCA and DB(15-0) signals. See figure 46 for an example of signals with and without precompensation.

Table 25 - Fast-160 precompensation

	Bit pattern and driver strength							
<b>First data bits</b>	1	0	1	0	1	0	1	0
<b>Driver strength</b>	Weak	Strong	Strong	Strong	Strong	Strong	Strong	Strong
<b>First data bits</b>	1	1	0	1	1	0	1	1
<b>Driver strength</b>	Weak	Weak	Strong	Strong	Weak	Strong	Strong	Weak
<b>First data bits</b>	1	1	1	0	0	0	1	1
<b>Driver strength</b>	Weak	Weak	Weak	Strong	Weak	Weak	Strong	Weak
<b>First data bits</b>	0	1	0	1	0	1	0	1
<b>Driver strength</b>	Strong	Strong	Strong	Strong	Strong	Strong	Strong	Strong
<b>First data bits</b>	0	0	1	0	0	1	0	0
<b>Driver strength</b>	Strong	Weak	Strong	Strong	Weak	Strong	Strong	Weak
<b>First data bits</b>	0	0	0	1	1	1	0	0
<b>Driver strength</b>	Strong	Weak	Weak	Strong	Weak	Weak	Strong	Weak
<p>In all the examples in this table the bit before the start of the bit pattern is a 1b.</p> <p>The weak driver is used anytime a bit value does not change.</p> <p>The level of a SCSI data bus signal, after training shall follow the rules in this subclause regardless of whether or not the level is qualified as valid or invalid by the P1 signal (see 10.7.4.3).</p> <p>See 10.7.4.2 to determine when these precompensation rules apply to the training pattern.</p>								

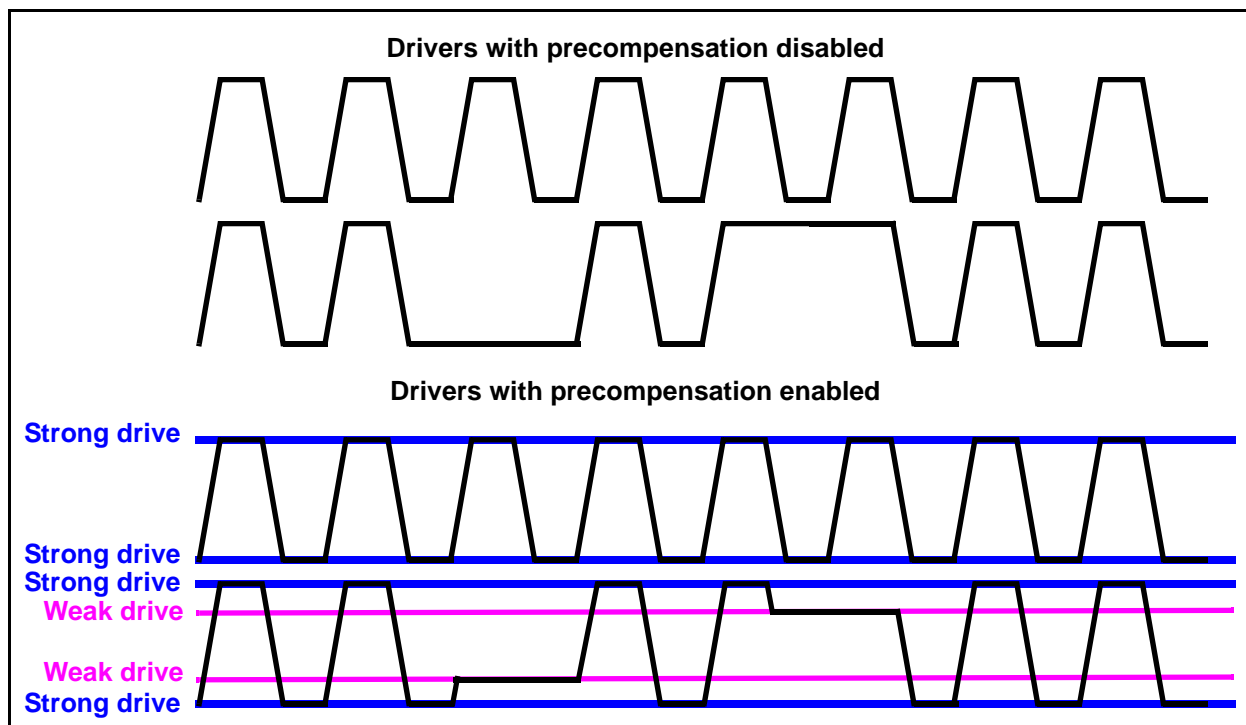
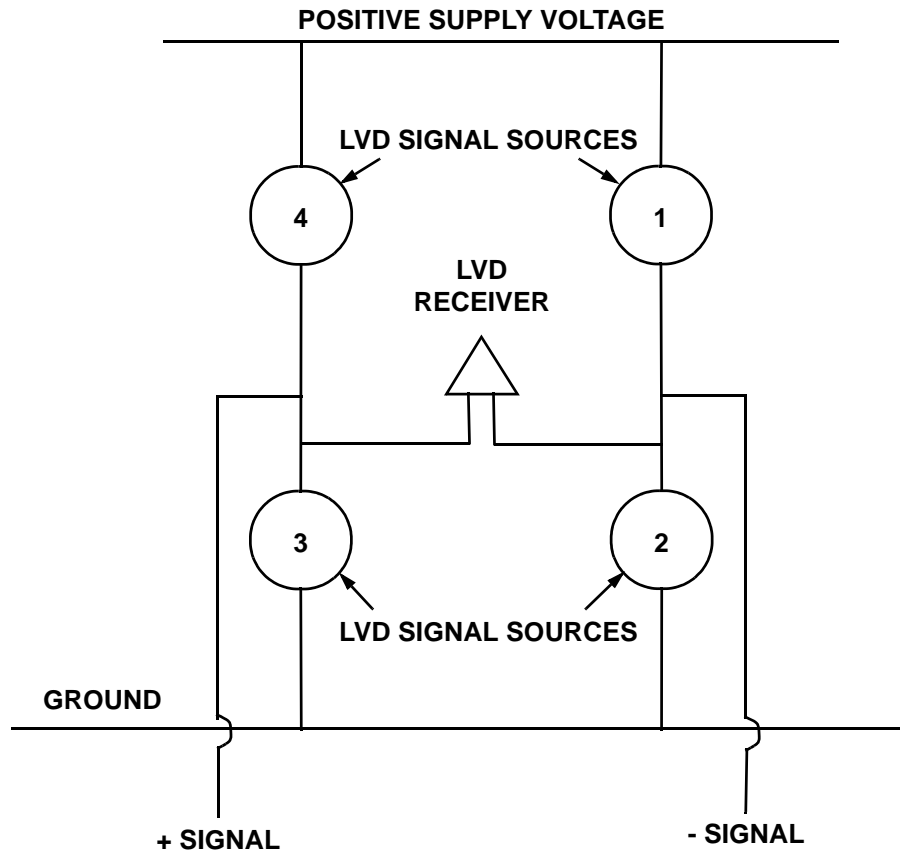


Figure 46 - Examples of fast-160 driver precompensation



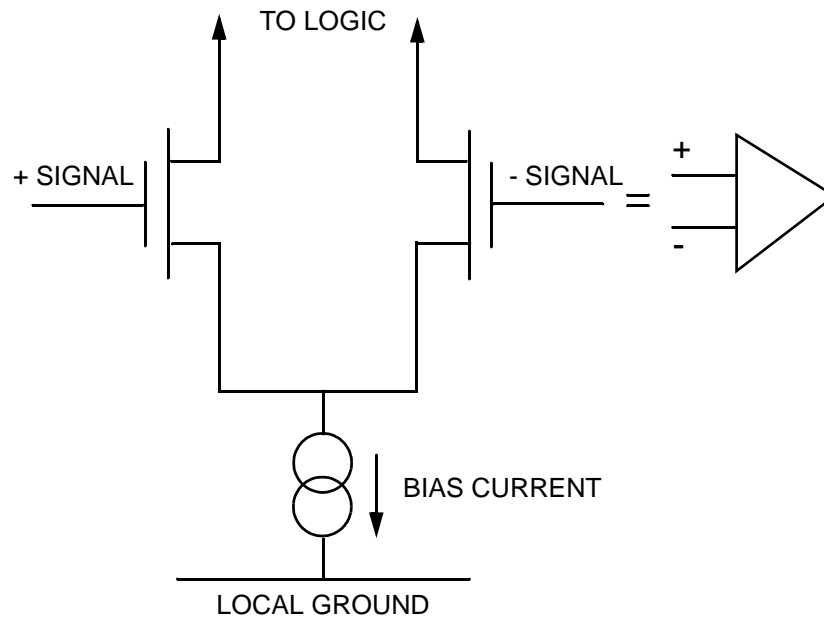
**Figure 47 - LVD transceiver architecture**

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 requirements in Annex A.

### 7.2.3 LVD receiver characteristics

LVD receivers shall be connected to the + signal and - signal as shown in figure 47. An example of an LVD receiver is shown in figure 48. LVD receivers shall meet the requirements in Annex A.



**Figure 48 - LVD receiver example**

#### 7.2.4 LVD capacitive loads

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

There are three components to differential SCSI bus segment capacitive loading: - signal to local ground (C1), + signal to local ground (C2), and - signal to + signal (C3) as shown in figure 49. 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 26. The driving source from the instrumentation shall apply an A.C. signal level less than 100 mV rms.

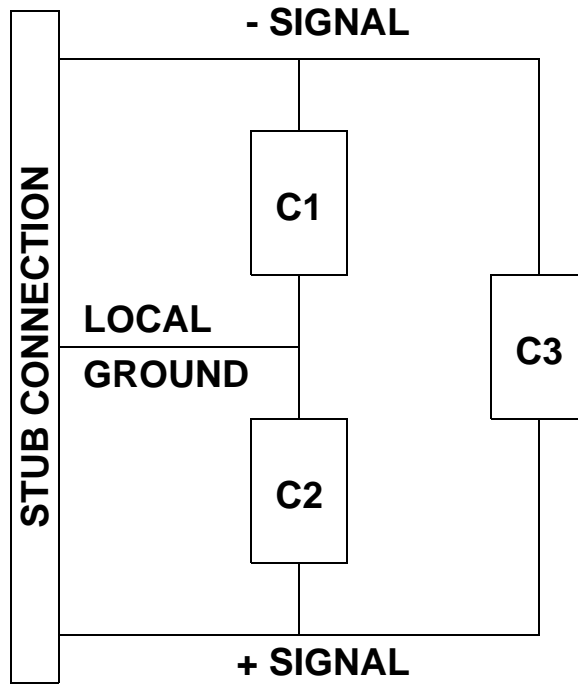


Figure 49 - LVD Capacitive loads

Table 26 - Values for LVD capacitive loads

Capacitance measurement	Maximum <sup>a</sup>	Description
C1 (pF)	15	@V = 0,7 to 1,8 V D.C. -sig/gnd REQ, ACK, DB(15-0), P_CRCA, and DB(P1)
C2 (pF)	15	@V = 0,7 to 1,8 V D.C. +sig/gnd REQ, ACK, DB(15-0), P_CRCA and DB(P1)
C3 (pF)	8	@V = 0,7 to 1,8 V D.C. both - and +sig/gnd V is the same for both sigs $\pm 100\text{mV}$ REQ, ACK, DB(15-0), P_CRCA and DB(P1)
C1 (pF)	25	@V = 0,7 to 1,8 V D.C. -sig/gnd all other signals
C2 (pF)	25	@V = 0,7 to 1,8 V D.C. +sig/gnd all other signals
C3 (pF)	13	@V = 0,7 to 1,8 V D.C. both - and +sig/gnd V is the same for both sigs $\pm 100\text{mV}$ all other signals
C1 - C2   (pF)	1,5	REQ, ACK, DB(15-0), P_CRCA and DB(P1) (within the signal pair)
C1 - C2   (pF)	3	all other signals (within the signal pair)
C1(i) - C1(REQ)   (pF)	2	For DATA(i) i = 0-15 and DB(P1) and P_CRCA
C2(i) - C2(REQ)   (pF)	2	For DATA(i) i = 0-15 and DB(P1) and P_CRCA
C1(i) - C1(ACK)   (pF)	2	For DATA(i) i = 0-15 and DB(P1) and P_CRCA
C2(i) - C2(ACK)   (pF)	2	For DATA(i) i = 0-15 and DB(P1) and P_CRCA
<sup>a</sup> It is recommended that implementors design capacitive loads to be balanced and as small as practical.		

SCSI devices containing the enabled bus segment termination shall have maximum values 1,5 times the maximums listed in table 26. Differential bus segment termination circuitry that is not part of a SCSI device shall have maximum values 0,5 times the maximums listed in table 26.

#### 7.2.4.1 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 (see 9.2.7). Requirements are defined in this subclause to avoid adverse affects from release glitches.

SCSI devices shall incorporate the requirements specified in table 27 when using LVD drivers and may



incorporate the requirements when using other drivers. The use of active negation increases crosstalk noise margin and increases the true-to-false transition speed as compared to passive negation.

**Table 27 - Glitch management requirements for SCSI devices using LVD drivers**

Signals	Mode	Active negation	Transmitting device	Receiving device
BSY, SEL, RST	I,T	P	No glitch management required.	No glitch management required.
ACK, ATN during transitions to BUS FREE phase	I	R	The SCSI initiator port shall wait for a BUS FREE phase <sup>a</sup> 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 SCSI target port shall ignore the ACK and ATN signals until a subsequent physical connect.
ACK, ATN during QAS	I	R	The SCSI initiator port shall wait until two system deskew delays after it detects C/D, I/O, and MSG false before releasing the ACK and ATN signals from the actively negated state.	Starting no later than two system deskew delays after negating C/D, I/O, and MSG, the SCSI target port shall ignore the ACK and ATN signals until a subsequent physical connect.
REQ during transitions to BUS FREE phase	T	R	The SCSI target port shall wait 2,5 times bus settle delay after releasing the BSY signal before releasing the REQ signal from the actively negated state.	The SCSI initiator port shall begin to ignore the REQ signal within 1,5 times bus settle delay of the transition of the BSY signal from true to false
REQ during QAS	T	R	The SCSI target port shall wait two system deskew delays after negating C/D, I/O, and MSG before releasing the REQ signal from the actively negated state.	Starting no later than two system deskew delays (see 9.2.51) after detecting C/D, I/O, and MSG false, the SCSI initiator port shall ignore the REQ signal until a subsequent physical connect.
C/D, I/O, MSG except during SELECTION and RESELECTION phases of QAS	T	R	After a SELECTION or RESELECTION phase, these signals shall not be released until the BSY signal is released.	No glitch management required.
C/D, I/O, MSG during SELECTION and RESELECTION phases of QAS	T	P	After detecting SEL true following QAS arbitration, the SCSI target port shall release these signals within a QAS release delay.	No glitch management required
DATA BUS during SELECTION and RESELECTION 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 = SCSI initiator port; P = prohibited; R = required; T = SCSI target port				
<sup>a</sup> BUS FREE phase starts one bus settle delay after the BSY and SEL signals are both false.				

## 7.2.5 SE/HVD transmission mode detection

### 7.2.5.1 SE/HVD transmission mode detection overview

Neither SE nor HVD is defined in this standard. For information on SE SCSI device implementation see the SCSI Parallel Interface-4 standard. For information on HVD SCSI device implementation see the SCSI Parallel Interface-2 standard.

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

LVD termination shall drive the DIFFSENS line as specified in 7.2.5.2 and LVD SCSI devices shall sense the DIFFSENS signal as specified in 7.2.5.3.

SCSI devices and terminators connected to the DIFFSENS line shall comply with the requirements in table 21 and table 22.

### 7.2.5.2 LVD DIFFSENS driver

The LVD DIFFSENS driver sets a voltage level on the DIFFSENS line that uniquely defines an LVD transmission mode. LVD terminators shall provide an LVD DIFFSENS driver according to the specifications in table 28.

**Table 28 - LVD DIFFSENS driver specifications**

Value	max.	nominal	min.	notes
$V_O$ when $I_O = 0$ 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.   ( $\mu A$ )	10			With terminator disabled.
Input sink current D.C. ( $\mu A$ ) 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.
All requirements apply at the terminator bus segment path connection (see figure 3). All measurements per figure 50. $I_{OS} = I_O$ short circuit, when SE SCSI device is attached.				

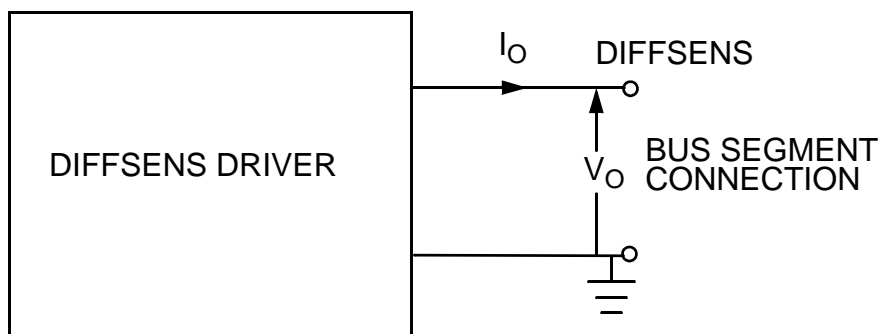


Figure 50 - LVD DIFFSENS driver signal definitions

### 7.2.5.3 LVD DIFFSENS receiver

LVD SCSI devices shall incorporate an LVD DIFFSENS receiver that detects the voltage level on the DIFFSENS line for purposes of informing the SCSI device of the transmission mode being used by the bus segment. The LVD DIFFSENS receiver shall be capable of detecting SE, LVD, and HVD SCSI devices. Table 29 and figure 51 define the receiver input levels for each of the three modes.

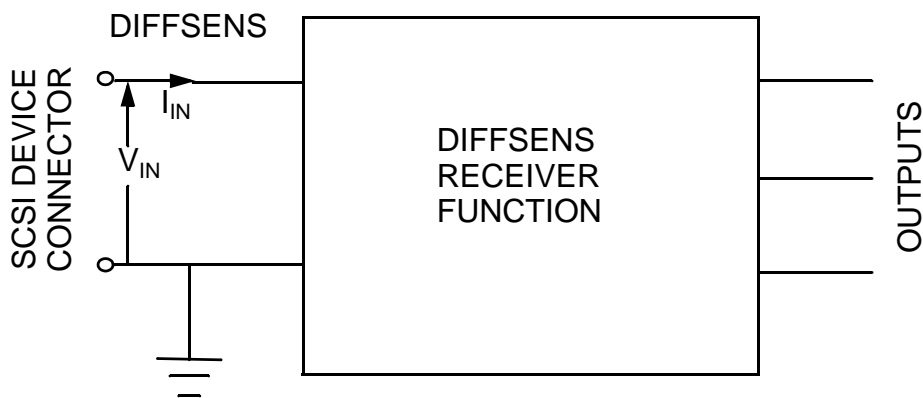


Figure 51 - DIFFSENS receiver function

**Table 29 - DIFFSENS receiver operating requirements**

<b>V<sub>in</sub> range</b>	<b>Sensed differential driver type</b>
-0,35 V to +0,5 V	SE
> +0,5 V to < 0,7 V	indeterminate for detecting SE and LVD driver type
0,7 V to 1,9 V	LVD
> 1,9 V to < 2,4 V	indeterminate for detecting LVD and HVD driver type
2,4 V to 5,5 V	HVD
Input resistance (V <sub>in</sub> /lin) shall be 200 K $\Omega$ to 250 K $\Omega$ @ V <sub>in</sub> < 2,7 V under all conditions of power supply (i.e., powering on, powering off, power transients) All voltages measured at the device connector with respect to local ground. A combination of LVD/MSE devices and HVD devices on the same bus segment may result in damage to the DIFFSENS input of the LVD/MSE device.	

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

SCSI devices shall not allow the + signal line or - signal line drivers to leave the high impedance state during initial power on until both of the following conditions are satisfied:

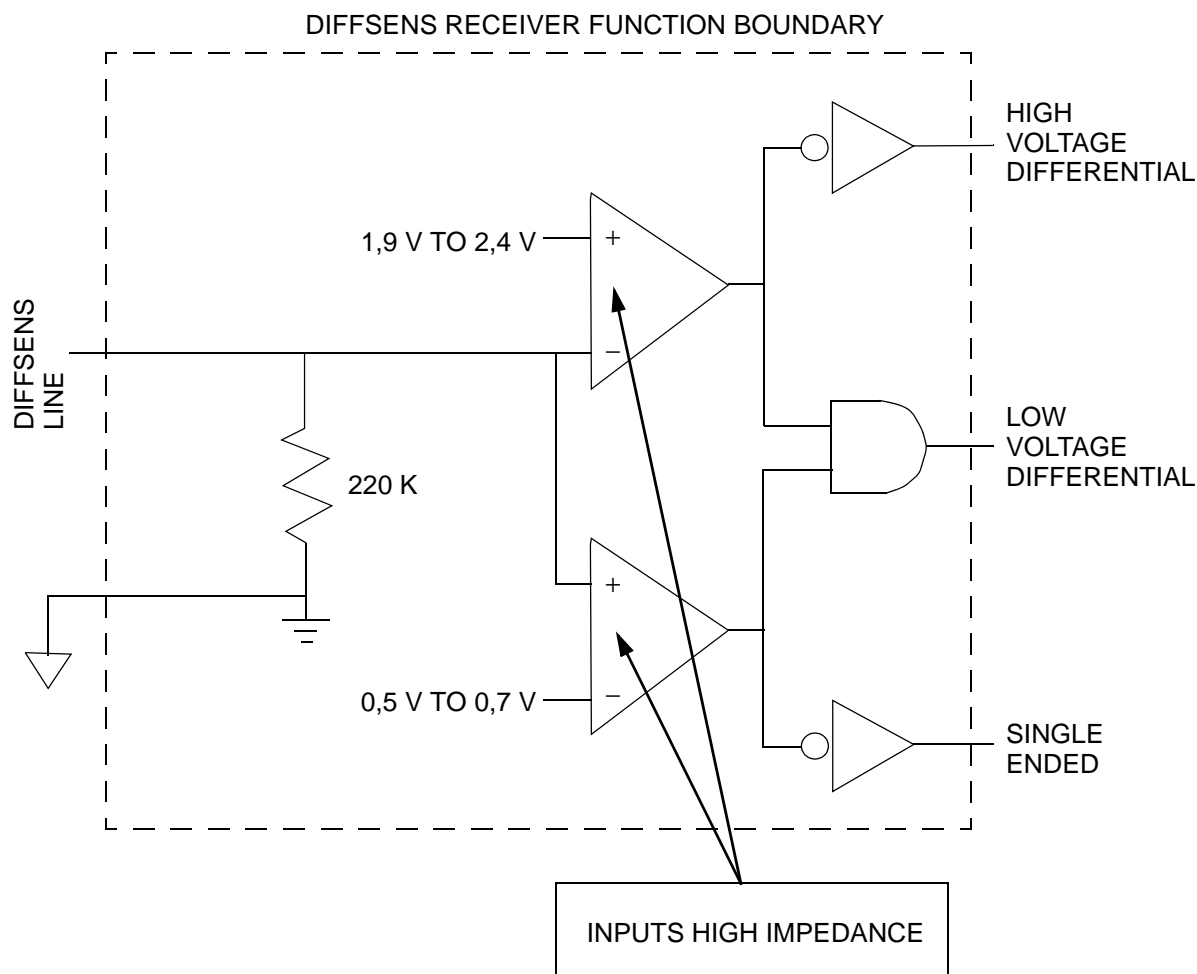
- a) the SCSI device is capable of logical operation for at least a DIFFSENS voltage filter time (see 9.2.22), and

NOTE 13 - The DIFFSENS voltage filter time delay allows time for the DIFFSENS pin to settle after the initial power connection (e.g., in the case of insertion of a SCSI 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 SCSI 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 a DIFFSENS voltage filter time. A multimode SCSI device shall change to the new signal driver or receiver mode based on the DIFFSENS voltage level within 400 ms of the last DIFFSENS voltage change regardless of the DIFFSENS voltage filter time. SCSI devices not capable of the new mode shall release the SCSI bus and remain in the high impedance state after the DIFFSENS voltage filter time.

An example implementation of an LVD DIFFSENS receiver is shown in figure 52. The reference voltage tolerance is greater on the higher voltage reference. This allows a simple resistor divider between V<sub>CC</sub> and ground for the references. The DIFFSENS voltage filter time requirement is implemented in logic in this example and is not shown in the figure 52.



**Figure 52 - LVD DIFFSENS receiver example**

### 7.3 Terminator power

Provision shall be made to provide power from one or more sources to the TERMPWR lines of the SCSI bus segment. 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 segment 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 (e.g., if the TERMPWR source voltage falls below the voltage existing on the TERMPWR line, the TERMPWR source sinks current unless unidirectional isolation is present in the TERMPWR source).

Regulatory agencies may require limiting maximum (i.e., 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, as measured at the terminator bus segment path connector shall be as defined in table 30.

**Table 30 - Terminator power characteristics at the terminator**

Terminator power characteristics	LVD
$I_{\min}$ (A) @ $V_{\min}$	0,5
$V_{\min}$ (V) @ $I_{\min}$	3,0
$V_{\max}$ (V) @ all conditions	5,25
The recommended TERMPWR current limiting is 2,0 A.	

NOTE 14 - It is recommended that a 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  $\mu\text{F}$ , to improve signal quality, but not greater than 10  $\mu\text{F}$ . (see 6.4)

Usage of the terminator power lines for distribution of power for purposes other than for SCSI bus segment termination is outside the scope of this standard.

## 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 except during MESSAGE IN phase when QAS is enabled (see 4.12.4.6.4). All SCSI devices that have QAS enabled are required to monitor messages during a MESSAGE IN phase for a QAS REQUEST message. The maximum number of SCSI devices is determined by the width of the data path implemented and restrictions in clause 7. The SCSI devices may be any combination of SCSI initiator ports and SCSI target ports, provided there is at least one of each.

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

NOTE 15 - A port is usually fixed as a SCSI initiator port or SCSI target port, but some ports may be able to assume either role.

Table 31 shows the relationship between SCSI Addresses, SCSI IDs, and arbitration priority. In table 31 a hyphen ("-") represents a logical zero bit resulting from the data bus bit being released.

**Table 31 - Arbitration priorities by SCSI ID**

SCSI address	DB 15	DB 8	DB 7	DB 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
Key: - = a logical zero bit resulting from the data bus bit being released					

### 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 a SCSI initiator port to select a SCSI target port or by a SCSI target port to reselect a SCSI initiator port.

**RST (RESET).** An "OR-tied" signal that indicates the bus reset condition (see 12.3).

**C/D (CONTROL/DATA).** A signal sourced by a SCSI target port that indicates whether control or DATA phase information is on the DATA BUS. Asserted indicates control (i.e., COMMAND, STATUS, and MESSAGE phases).

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

**MSG (MESSAGE).** A signal sourced by a SCSI target port to indicate the MESSAGE phase or a DT DATA phase depending on whether C/D is true or false. Asserted indicates MESSAGE or DT DATA.

**REQ (REQUEST).** A signal sourced by a SCSI target port to indicate a request for an information transfer on the SCSI bus.

**ACK (ACKNOWLEDGE).** A signal sourced by a SCSI initiator port to respond with an acknowledgment of an information transfer on the SCSI bus.

**ATN (ATTENTION).** A signal sourced by a SCSI initiator port to indicate the attention condition.

**P\_CRCA (PARITY/CRC AVAILABLE).** A signal indicating either parity or CRC available based on bus phase and negotiated settings.

During the SELECTION phase, RESELECTION phase, ST DATA phase, COMMAND phase, MESSAGE phase, and STATUS phase, this signal is referred to as DB(P\_CRCA) and is sourced by the SCSI device port driving the DATA BUS. The DB(P\_CRCA) signal is associated with the DB(7-0) signals and is used to detect the presence of an odd number of bit errors within the byte. The DB(P\_CRCA) bit is driven such that the number of logical ones in the byte plus the parity bit is odd.

During DT DATA phases when data group transfers are enabled (see 4.12.4.6.3) this signal is referred to as P\_CRCA and is sourced by the SCSI target port to control whether a data group field is a pad field, pCRC field, or data field (see 10.7.3.3.5). When asserted the data group field shall be pad or pCRC fields that shall not be transferred to the application client. When negated the data group field shall be a data field that shall be transferred to the application client.

During DT DATA phases when information unit transfers are enabled this signal is referred to as P\_CRCA and is sourced by the SCSI target port. Depending on the negotiated condition of read streaming and write flow control the SCSI initiator port and SCSI target port usage for P\_CRCA is different. When information unit transfers are enabled the SCSI target port and SCSI initiator port shall use the P\_CRCA signal as indicated in table 32.



Table 32 - P\_CRC signal usage requirements

Write flow control <sup>b</sup>	Read streaming <sup>a</sup>	DT DATA phase	SCSI initiator port response to P_CRCA	SCSI target port usage of P_CRCA
disabled	disabled	All	Ignore	Continuously negated
enabled	disabled	DT DATA IN	Ignore	Continuously negated
		DT DATA OUT	Monitor	Asserts to indicate when the current SPI data stream information unit is the last SPI data stream information unit of the current write stream.
disabled	enabled	DT DATA IN	Monitor	Asserts to indicate when the current SPI data stream information unit is the last SPI data stream information unit of the current read stream.
		DT DATA OUT	Ignore	Continuously negated
enabled	enabled	DT DATA IN	Monitor	Asserts to indicate when the current SPI data stream information unit is the last SPI data stream information unit of the current read stream.
		DT DATA OUT	Monitor	Asserts to indicate when the current SPI data stream information unit is the last SPI data stream information unit of the current write stream.

<sup>a</sup> A SCSI device is not required to use read streaming even if it is enabled.

<sup>b</sup> A SCSI device is not required to use write flow control even if it is enabled.

**P1 (PARITY 1).** A signal normally sourced by the SCSI device driving the DATA BUS. The P1 signal is associated with the DB(15-8) signals and is used to detect the presence of an odd number of bit errors within the byte. The P1 bit is driven such that the number of logical ones in the byte plus the P1 bit is odd.

During the ST DATA phase with transfer length set for narrow transfers, COMMAND phase, MESSAGE phase, and STATUS phase the P1 signal shall not be driven by any SCSI device.

During the SELECTION phase and the RESELECTION phase on a 16-bit wide bus segment the P1 signal shall be sourced by the SCSI device driving the DATA BUS.

During DT DATA phases when data group transfers are enabled (see 4.12.4.6.3) the P1 signal shall be continuously negated by the SCSI device driving the DB(15-0) signals and shall be ignored by the SCSI device receiving the DB(15-0) signals during DT DATA phases.

During DT DATA phases when information unit transfers and DT synchronous transfers are enabled the P1 signal shall be continuously negated by the SCSI device driving the DB(15-0) signals and shall be ignored by the SCSI device receiving the DB(15-0) signals during DT DATA phases.

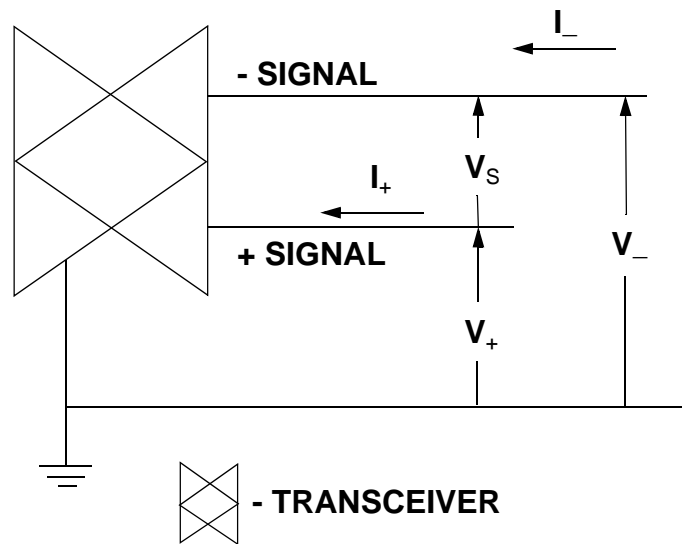
During DT DATA phases when information unit transfers and paced transfers are enabled the P1 signal shall be sourced by the SCSI device driving the DATA BUS. The P1 signal is used to indicate the data valid or data invalid state during paced transfers (see 10.7.4.3).

**DB(7-0) (8-bit DATA BUS).** Eight data-bit signals that form the 8-bit DATA BUS. Bit significance and priority during arbitration are shown in table 31.

**DB(15-0) (16-bit DATA BUS).** Sixteen data-bit signals that form the 16-bit DATA BUS. Bit significance and priority during arbitration are shown in table 31.

### 8.3 LVD Signal states

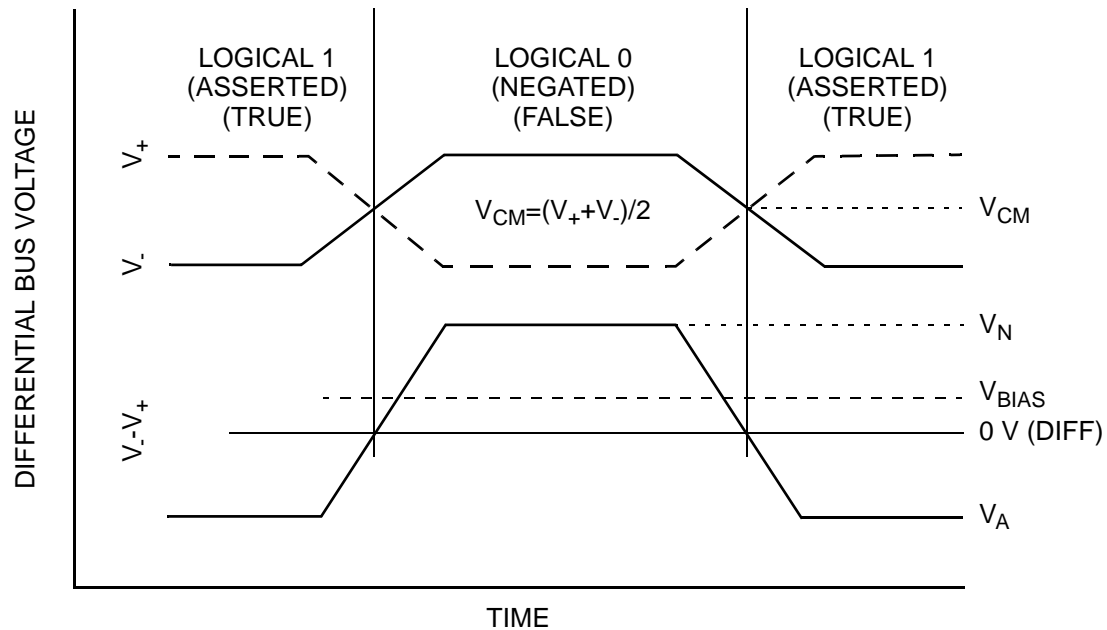
Figure 53 defines the voltage and current definitions. A signal that is released goes to the false state because the bias of the terminator pulls the signal false.



**Figure 53 - Voltage and current definitions**

Figure 54 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 54 - LVD Signaling sense**

NOTE 16 - For a description of  $V_{BIAS}$  see 7.2.1.

## 8.4 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, DB(P\_CRCA), and DB(P1) are simultaneously driven by two or more SCSI devices. DB(P\_CRCA), and DB(P1) shall not be driven false during the ARBITRATION phase but may be driven false in other phases.

## 8.5 Signal sources

Table 33 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 33 - Signal sources

SCSI bus phase	P cable signals						
	A cable signals						DB15-8 DB(P1)
	BSY	SEL	C/D I/O MSG REQ	ACK ATN	DB(7-0)	P_CRCA	
BUS FREE	None	None	None	None	None	None	None
Normal ARBITRATION	All	Win	None	None	S ID	S ID	S ID
QAS ARBITRATION	PT	Win	None	None	S ID	S ID	S ID
SELECTION	I&T	Init	None	Init	Init	Init	Init
RESELECTION	I&T	Targ	Targ	Init	Targ	Targ	Targ
COMMAND	Targ	None	Targ	Init	Init	Init	None
ST DATA IN	Targ	None	Targ	Init	Targ	Targ	Targ
ST DATA OUT	Targ	None	Targ	Init	Init	Init	Init
DT DATA IN	Targ	Targ	Targ	Init	Targ	Targ	Targ
DT DATA OUT	Targ	Targ	Targ	Init	Init	Targ	Init
STATUS	Targ	None	Targ	Init	Targ	Targ	None
MESSAGE IN	Targ	None	Targ	Init	Targ	Targ	None
MESSAGE OUT	Targ	None	Targ	Init	Init	Init	None
<p>Key:</p> <p>All: The signal shall be driven by all SCSI devices that are actively arbitrating.</p> <p>S ID: A unique data bit representing 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 P_CRCA and DB(P1) bits may be released or driven to the true state, but shall not be driven to the false state during this phase.</p> <p>I&amp;T: The signal shall be driven by the SCSI initiator port, SCSI target port, or both, as specified in the SELECTION phase and RESELECTION phase.</p> <p>Init: If driven, the signal shall be driven only by the active SCSI initiator port.</p> <p>None: The signal shall be released; that is, not driven by any SCSI device. The bias circuitry of the bus segment terminators pulls the signal to the false state.</p> <p>Win: The signal shall be driven by the one SCSI device that wins arbitration (see 10.4).</p> <p>Targ: If driven, the signal shall be driven only by the active SCSI target port.</p> <p>PT: The signal shall be driven by the SCSI target port that initiated the QAS arbitration (see 10.4).</p>							

## 9 SCSI parallel bus timing

### 9.1 SCSI parallel bus timing values

See table 34, table 35, table 36, table 37, and table 38 for the SCSI bus timing values. Unless otherwise indicated, the timing measurements for each SCSI device, shown in table 34, shall be calculated from signal conditions existing at that SCSI device port. The timing characteristics of each signal are described in the following paragraphs. Timing requirements relating to LVD release glitches are defined in 7.2.4.1.

Table 39 and table 40 list the timing budget for paced transfers. Figure 55 gives the measurement reference point location for the items that make up the timing budget. The accessibility of the measurement points is outside the scope of this standard.

**Table 34 - SCSI bus control timing values**

Subclause	Timing description	Type	Timing values
9.2.1	Arbitration delay	minimum	2,4 $\mu$ s
9.2.4	Bus clear delay	maximum	800 ns
9.2.5	Bus free delay	minimum	800 ns
9.2.6	Bus set delay	maximum	1,6 $\mu$ s
9.2.7	Bus settle delay	minimum	400 ns
9.2.8	Cable skew <sup>a</sup>	maximum	4 ns
9.2.21	Data release delay	maximum	400 ns
9.2.22	DIFFSENS voltage filter time	minimum	100 ms
9.2.23	Physical disconnection delay	minimum	200 $\mu$ s
9.2.24	Power on to selection <sup>b</sup>	maximum	10 s
9.2.25	QAS arbitration delay	minimum	1000 ns
9.2.26	QAS assertion delay	maximum	200 ns
9.2.27	QAS release delay	maximum	200 ns
9.2.28	QAS non-DATA phase REQ(ACK) period	minimum	50 ns
9.2.42	Reset delay	minimum	200 ns
9.2.43	Reset hold time	minimum	25 $\mu$ s
9.2.44	Reset to selection <sup>b</sup>	maximum	250 ms
9.2.46	Selection abort time	maximum	200 $\mu$ s
9.2.47	Selection time-out delay <sup>b</sup>	minimum	250 ms
9.2.51	System deskew delay	minimum	45 ns
<sup>a</sup> Cable Skew is measured at each SCSI device connection within the same bus segment with the transmitted skew subtracted from the received skew. <sup>b</sup> This is a recommended time. It is not mandatory.			

Table 35 - SCSI bus data &amp; information phase ST timing values

Subclause	Timing description	Type	Timing values for negotiated transfer rate <sup>d</sup>				
			Asynch	Fast-5	Fast-10	Fast-20	Fast-40
9.2.2	ATN transmit setup time	min	90 ns	33 ns	33 ns	21,5 ns	19,25 ns
9.2.3	ATN receive setup time	min	45 ns	17 ns	17 ns	8,5 ns	6,75 ns
9.2.8	Cable skew <sup>a</sup>	max	4 ns	4 ns	4 ns	3 ns	2,5 ns
9.2.29	Receive assertion period <sup>b</sup>	min	N/A	70 ns	22 ns	11 ns	6,5 ns
9.2.30	Receive hold time <sup>b, c</sup>	min	N/A	25 ns	25 ns	11,5 ns	4,75 ns
9.2.33	Receive negation period <sup>b</sup>	min	N/A	70 ns	22 ns	11 ns	6,5 ns
9.2.34	Receive setup time <sup>b, c</sup>	min	N/A	15 ns	15 ns	6,5 ns	4,75 ns
9.2.29	Receive REQ(ACK) period tolerance	min	N/A	1,1 ns	1,1 ns	1,1 ns	1,1 ns
9.2.48	Signal timing skew	max	8 ns	8 ns	8 ns	5 ns	4,5 ns
9.2.41	REQ(ACK) period	nominal	N/A	200 ns	100 ns	50 ns	25 ns
9.2.55	Transmit assertion period <sup>b</sup>	min	N/A	80 ns	30 ns	15 ns	8 ns
9.2.56	Transmit hold time <sup>b, c</sup>	min	N/A	53 ns	33 ns	16,5 ns	9,25 ns
9.2.58	Transmit negation period <sup>b</sup>	min	N/A	80 ns	30 ns	15 ns	8 ns
9.2.59	Transmit setup time <sup>b, c</sup>	min	N/A	23 ns	23 ns	11,5 ns	9,25 ns
9.2.60	Transmit REQ(ACK) period tolerance	max	N/A	1 ns	1 ns	1 ns	1 ns
<sup>a</sup> Cable skew is measured at each SCSI device connection within the same bus segment with the transmitted skew subtracted from the received skew. <sup>b</sup> See 9.3 for measurement points for the timing specifications. <sup>c</sup> See 9.6 for examples of how to calculate setup and hold timing. <sup>d</sup> SCSI bus timing values specified by the maximum transfer rate for the given range shall apply even if a slower transfer rate within the given range is negotiated.							

**Table 36 - Miscellaneous SCSI bus data & information phase DT timing values**

Sub-clause	Timing description	Type	Timing values for negotiated transfer rate <sup>b</sup>					
			Fast-10	Fast-20	Fast-40	Fast-80	Fast-160	Fast-320
9.2.8	Cable skew <sup>a</sup>	max	4 ns	3 ns	2,5 ns	2,5 ns	2,5 ns	2,5 ns
9.2.41	REQ(ACK) period	nominal	200 ns	100 ns	50 ns	25 ns	12,5 ns	6,25 ns
9.2.45	Residual Skew Error <sup>c</sup>	max	N/A	N/A	N/A	N/A	±0,15 ns	±0,20 ns
9.2.49	Skew correction range <sup>d</sup>	min	N/A	N/A	N/A	N/A	±3,45 ns <sup>e</sup>	±3,45 ns <sup>e</sup>
9.2.48	Signal timing skew	max	26,8 ns	13,4 ns	6,7 ns	3,35 ns	4,85 ns	4,85 ns
9.2.50	Strobe Offset Tolerance	max	N/A	N/A	N/A	N/A	±0,125 ns	±0,20 ns
Fast-160 and fast-320 SCSI devices shall not change timing parameters between training (see 10.7.4.2) or reset events (see 12.5).								
<sup>a</sup> Cable skew is measured at each SCSI device connection within the same bus segment with the transmitted skew subtracted from the received skew. <sup>b</sup> SCSI bus timing values specified by the maximum transfer rate for the given range shall apply even if a slower transfer rate within the given range is negotiated. <sup>c</sup> Calculated assuming timing budget shown in table 39 and table 40. <sup>d</sup> Measured at the receiving devices connector using clean input signals with 500 mV peak amplitude and 1 ns rise and fall time between 20 % and 80 % of the signal. <sup>e</sup> Relative to the REQ(ACK) clocking signal.								

Table 37 - Transmit SCSI bus data &amp; information phase DT timing values

Sub-clause	Timing description	Type	Timing values for negotiated transfer rate <sup>c</sup>					
			Fast-10	Fast-20	Fast-40	Fast-80	Fast-160	Fast-320
9.2.2	ATN transmit setup time	min	48,4 ns	29,2 ns	19,6 ns	14,8 ns	14 ns	14 ns
9.2.14	Flow control transmit hold time	min	38,4 ns	19,2 ns	9,6 ns	4,8 ns	14 ns	14 ns
9.2.15	Flow control transmit setup time	min	48,4 ns	29,2 ns	19,6 ns	14,8 ns	14 ns	14 ns
9.2.19	pCRC transmit hold time	min	38,4 ns	19,2 ns	9,6 ns	4,8 ns	N/A	N/A
9.2.20	pCRC transmit setup time	min	48,4 ns	29,2 ns	19,6 ns	14,8 ns	N/A	N/A
9.2.55	Transmit assertion period <sup>a</sup>	min	92 ns	46 ns	23 ns	11,5 ns	5,69 ns	2,565 ns
9.2.56	Transmit hold time <sup>a, b</sup>	min	38,4 ns	19,2 ns	9,6 ns	4,8 ns	4,77 ns	0,365 ns
9.2.57	Transmit ISI Compensation	max	<sup>d</sup>	<sup>d</sup>	<sup>d</sup>	<sup>d</sup>	1,0 ns	0 ns
9.2.58	Transmit negation period <sup>a</sup>	min	92 ns	46 ns	23 ns	11,5 ns	5,69 ns	2,565 ns
9.2.60	Transmit REQ(ACK) period tolerance	max	0,6 ns	0,6 ns	0,6 ns	0,6 ns	0,06 ns	0,06 ns
9.2.61	Transmit REQ assertion period with P_CRCA transitioning	min	97,5 ns	54 ns	35,5 ns	24 ns	N/A	N/A
9.2.62	Transmit REQ negation period with P_CRCA transitioning	min	97,5 ns	54 ns	35,5 ns	24 ns	N/A	N/A
9.2.59	Transmit setup time <sup>a, b</sup>	min	38,4 ns	19,2 ns	9,6 ns	4,8 ns	-1,48 ns	-2,76 ns
9.2.63	Transmitter skew	max	N/A	N/A	N/A	N/A	±0,75 ns	±0,75 ns
9.2.64	Transmitter time asymmetry	max	N/A	N/A	N/A	N/A	±0,25 ns	±0,25 ns
Fast-160 and fast-320 SCSI devices shall not change timing parameters between training (see 10.7.4.2) or reset events (see 12.5).								
<sup>a</sup> See 9.3 for measurement points for the timing specifications. <sup>b</sup> See 9.6 for examples of how to calculate setup and hold timing. <sup>c</sup> SCSI bus timing values specified by the maximum transfer rate for the given range shall apply even if a slower transfer rate within the given range is negotiated. <sup>d</sup> Calculated without any ISI compensation.								



Table 38 - Receive SCSI bus data &amp; information phase DT timing values

Sub-clause	Timing description	Type	Timing values for negotiated transfer rate <sup>c</sup>					
			Fast-10	Fast-20	Fast-40	Fast-80	Fast-160	Fast-320
9.2.3	ATN receive setup time	min	13,6 ns	7,8 ns	4,9 ns	3,45 ns	3 ns	3 ns
9.2.12	Flow control receive hold time	min	11,6 ns	5,8 ns	2,9 ns	1,45 ns	3 ns	3 ns
9.2.13	Flow control receive setup time	min	18,6 ns	12,8 ns	9,9 ns	8,45 ns	3 ns	3 ns
9.2.17	pCRC receive hold time	min	11,6 ns	5,8 ns	2,9 ns	1,45 ns	N/A	N/A
9.2.18	pCRC receive setup time	min	18,6 ns	12,8 ns	9,9 ns	8,45 ns	N/A	N/A
9.2.29	Receive assertion period <sup>a</sup>	min	80 ns	40 ns	20 ns	8,5 ns	4,74 ns	1,615 ns
9.2.30	Receive hold time <sup>a, b</sup>	min	11,6 ns	5,8 ns	2,9 ns	1,45 ns	-0,08 ns	-6,635 ns
9.2.33	Receive negation period <sup>a</sup>	min	80 ns	40 ns	20 ns	8,5 ns	4,74 ns	1,615 ns
9.2.34	Receive setup time <sup>a, b</sup>	min	11,6 ns	5,8 ns	2,9 ns	1,45 ns	-6,33 ns	-9,76 ns
9.2.35	Receive REQ(ACK) period tolerance	min	0,7 ns	0,7 ns	0,7 ns	0,7 ns	0,06 ns	0,06 ns
9.2.36	Receive REQ assertion period with P_CRCA transitioning	min	85,5 ns	48 ns	32,5 ns	21 ns	N/A	N/A
9.2.37	Receive REQ negation period with P_CRCA transitioning	min	85,5 ns	48 ns	32,5 ns	21 ns	N/A	N/A
9.2.38	Receive Skew Compensation	max	N/A	N/A	N/A	N/A	4,4 ns	4,4 ns
9.2.31	Receive Internal Hold Time <sup>d</sup>	min	N/A	N/A	N/A	N/A	0,345 ns	0,032 ns
9.2.32	Receive Internal Setup Time <sup>d</sup>	min	N/A	N/A	N/A	N/A	0,345 ns	0,032 ns
9.2.40	Receiver de-skewed data valid window <sup>d</sup>	min	N/A	N/A	N/A	N/A	±2,1 ns	±1,1 ns
Fast-160 and fast-320 SCSI devices shall not change timing parameters between training (see 10.7.4.2) or reset events (see 12.5).								
<sup>a</sup> See 9.3 for measurement points for the timing specifications. <sup>b</sup> See 9.6 for examples of how to calculate setup and hold timing. <sup>c</sup> SCSI bus timing values specified by the maximum transfer rate for the given range shall apply even if a slower transfer rate within the given range is negotiated. <sup>d</sup> Calculated assuming timing budget shown in table 39 and table 40.								

Table 39 - SCSI fast-160 and fast-320 non-compensatable timing budget

Item	Location	Fast-160		Fast-320		Comments
<b>Nominals:</b>						
REQ(ACK) period		12,5 ns		6,25 ns		from table 36
Transfer period		6,25 ns		3,125 ns		REQ(ACK) period / 2
<b>Transmitter and receiver errors:</b>		<b>Total Errors</b>	<b>Post Compensation</b>	<b>Total Errors</b>	<b>Post Compensation</b>	<b>Worse case total of + and - time shift except where noted</b>
<b>Transmitter errors:</b>						
REQ(ACK) period tolerance / 2	A	0,06 ns	0,06 ns	0,06 ns	0,06 ns	Tolerance of transmitter plus measurement error <sup>b</sup>
Clock jitter	A	0,25 ns	0,25 ns	0,25 ns	0,25 ns	
System noise at transmitter	B	0,25 ns	0,25 ns	0,2 ns	0,2 ns	Time impact
Transmitter chip skew	A	0,75 ns	0 ns	0,75 ns	0 ns	
Transmitter trace skew	B	0,2 ns	0 ns	0,2 ns	0 ns	
Transmit time asymmetry	A	0,50 ns	0,50 ns	0,50 ns	0 ns	Compensated for on fast-320 SCSI devices
<b>Total transmitter error budget:</b>		<b>2,01 ns</b>	<b>1,06 ns</b>	<b>1,96 ns</b>	<b>0,51 ns</b>	
<b>Receiver errors:</b>						
System noise at receiver	G	0,25 ns	0,25 ns	0,2 ns	0,2 ns	Time impact
Chip noise in receiver	G	0,2 ns	0,2 ns	0,2 ns	0,2 ns	Time impact
Receiver chip skew	G	0,75 ns	0 ns	0,75 ns	0 ns	
Receiver trace skew	D	0,2 ns	0 ns	0,2 ns	0 ns	
Receiver time asymmetry	G	0,35 ns	0,35 ns	0,35 ns	0 ns	
Residual Skew error	F	0 ns	0,3 ns	0 ns	0,2 ns	After skew compensation
Strobe offset tolerance	F	0 ns	0,5 ns	0 ns	0,2 ns	Accuracy of centering strobe
Offset induced time asymmetry	G	N/A <sup>a</sup>	N/A <sup>a</sup>	0,8 ns	0,2 ns	Time impact from cumulative D.C. offsets at receiver
Receiver amplitude time skew	G	0,2 ns	0,2 ns	0,2 ns	0,05 ns	
<b>Total receiver error budget:</b>		<b>1,95 ns</b>	<b>1,8 ns</b>	<b>2,7 ns</b>	<b>1,05 ns</b>	
<b>Total transmitter + receiver error budget:</b>		<b>3,96 ns</b>	<b>2,86 ns</b>	<b>4,44 ns</b>	<b>1,56 ns</b>	
<b>Total timing error budget for interconnect and system margin:</b>		<b>2,29 ns</b>	<b>3,39 ns</b>	<b>-1,535 ns</b>	<b>1,565 ns</b>	Transfer period - (total transmitter + receiver error budget)
<sup>a</sup> Timing budgets in some previous standards neglected asymmetry & detection ambiguity and lumps chip noise, clock jitter, crosstalk time shift, noise, ISI and receiver amplitude skew into other terms (e.g., signal distortion skew) and/or ignores the effects.						
<sup>b</sup> Tolerance adjusted for half cycle (i.e., transfer period)						

Table 40 - SCSI fast-160 and fast-320 interconnect timing budget

Item	Loca- tion	Fast-160		Fast-320		Comments
Nominals:						
REQ(ACK) period		12,5 ns		6,25 ns		from table 36
Transfer period		6,25 ns		3,125 ns		REQ(ACK) period / 2
Transmitter and receiver errors:		Total Errors	Post Compensa- tion	Total Errors	Post Compensa- tion	Worse case total of + and - time shift except where noted
Interconnect errors:						
Cable skew <sup>a b</sup>	C	2,5 ns	0 ns	2,5 ns	0 ns	
Crosstalk time shift	C	0,7 ns	0,7 ns	0,5 ns	0,5 ns	Time impact
ISI of data	C	4,0 ns	2,0 ns	4,0 ns	1,0 ns	Worse case pattern
Total interconnect budget:		7,2 ns	2,7 ns	7,0 ns	1,5 ns	
<sup>a</sup> See 9.2.8.						
<sup>b</sup> The residual deskew error is included in the receiver error budget.						
<sup>c</sup> For more information on interconnect errors see the SCSI Passive Interconnect Performance standard.						

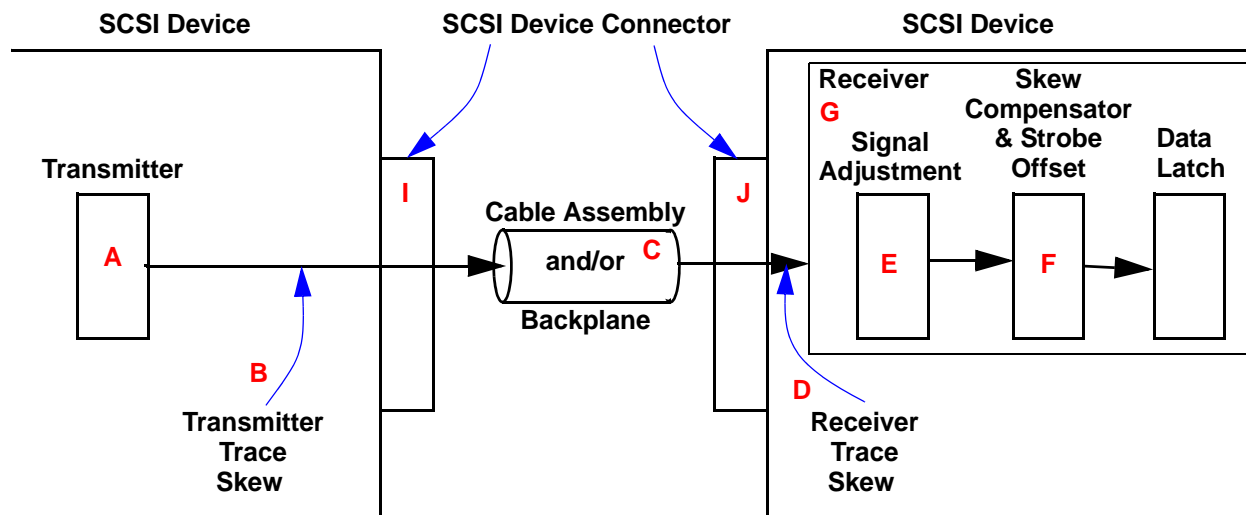


Figure 55 - Timing budget item locations

## 9.2 Timing description

### 9.2.1 Arbitration delay

The minimum time a SCSI device shall wait from asserting the BSY signal for arbitration until the DATA BUS is examined to see if arbitration has been won (see 10.4). There is no maximum time.

### 9.2.2 ATN transmit setup time

When information unit transfers are not being used, the ATN transmit setup time is the minimum time provided by the transmitter between the assertion of the ATN signal and the last negation of the ACK signal in any phase.

When information unit transfers are being used with synchronous transfers, the ATN transmit setup time is the minimum time provided by the transmitter between the assertion of the ATN signal and the negation of the ACK signal corresponding to the last iuCRC transfer of an information unit.

When information unit transfers are being used with paced transfers, the ATN transmit setup time is the minimum time provided by the transmitter between the assertion of the ATN signal and the assertion of the ACK signal corresponding to the last iuCRC transfer of an information unit.

Specified to provide the increased ATN receive setup time, subject to intersymbol interference, cable skew, and other distortions.

### 9.2.3 ATN receive setup time

When information unit transfers are not being used, the ATN receive setup time is the minimum time required at the receiver between the assertion of the ATN signal and the last negation of the ACK signal in any phase to recognize the assertion of an attention condition.

When information unit transfers are being used with synchronous transfers, the ATN receive setup time is the minimum time required at the receiver between the assertion of the ATN signal and the negation of the ACK signal corresponding to the last iuCRC transfer of an information unit to recognize the assertion of an attention condition.

When information unit transfers are being used with paced transfers, the ATN receive setup time is the minimum time required at the receiver between the assertion of the ATN signal and the assertion of the ACK signal corresponding to the last iuCRC transfer of an information unit to recognize the assertion of an attention condition.

### 9.2.4 Bus clear delay

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

- a) the BUS FREE phase is detected (i.e., 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 a SCSI device to release all SCSI bus signals is 1200 ns from the BSY and SEL signals first becoming both false. If a 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.

### 9.2.5 Bus free delay

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

### 9.2.6 Bus set delay

The maximum time for a 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.

### 9.2.7 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.

Provides time for a signal transition to propagate from the driver to the terminator and back to the driver.

### 9.2.8 Cable skew

The maximum difference in propagation time allowed between any two SCSI bus signals measured between any two SCSI devices while transmitting a free running clock data pattern.

### 9.2.9 Chip noise in receiver

The maximum transition to transition time shift due to the internal physics of the receiving SCSI device circuitry.

### 9.2.10 Clock jitter

The maximum transition to transition time shift of SCSI bus signals caused by short term variations in the transmitting SCSI device's clock.

### 9.2.11 Crosstalk time shift

The peak-to-peak time shift error on DB(0-15), P\_CRCA, or DB(P1) caused by transitions on all other DB(0-15), P\_CRCA, or DB(P1) signals.

### 9.2.12 Flow control receive hold time

The maximum time required by the SCSI initiator port between the assertion of the REQ signal corresponding to the last iuCRC transfer of a SPI data stream information unit and the changing of the P\_CRCA signal.

### 9.2.13 Flow control receive setup time

The maximum time required by the SCSI initiator port between the assertion of the P\_CRCA signal and the assertion of the REQ signal corresponding to the last iuCRC transfer of a SPI data stream information unit. Also, the maximum time required by the SCSI initiator port between the negation of the P\_CRCA signal and the assertion of the REQ signal corresponding to any valid data transfer of a SPI L\_Q information unit.

### 9.2.14 Flow control transmit hold time

The minimum time provided by the SCSI target port between the assertion of the REQ signal corresponding to the last iuCRC transfer of a SPI data stream information unit and the changing of the P\_CRCA signal.

Specified to provide the increased P\_CRCA receive setup time, subject to intersymbol interference, cable skew, and other distortions.

### 9.2.15 Flow control transmit setup time

The minimum time provided by the SCSI target port between the assertion of the P\_CRCA signal and the assertion of the REQ signal corresponding to the last iuCRC transfer of a SPI data stream information unit. Also, the minimum time provided by the SCSI target port between the negation of the P\_CRCA signal and the assertion of the REQ signal corresponding to any valid data transfer of a SPI L\_Q information unit.

Specified to provide the increased P\_CRCA receive setup time, subject to intersymbol interference, cable

skew, and other distortions.

#### **9.2.16 Offset induced time asymmetry**

Time symmetry error created by the cumulative sum of all offset errors seen by the receiver. This includes non-symmetrical transmitter drive plus terminator current mismatch, receiver offset, and voltage drop due to resistance in the interconnect within the cable or backplane.

#### **9.2.17 pCRC receive hold time**

The minimum time required at the receiver between the transition of the REQ signal and the transition of the P\_CRCA signal during data group transfers.

#### **9.2.18 pCRC receive setup time**

The minimum time required at the receiver between the transition of the P\_CRCA signal and the transition of the REQ signal during data group transfers.

Specified to ease receiver timing requirements and ensure that this signal, that is outside CRC protection, is received correctly.

#### **9.2.19 pCRC transmit hold time**

The minimum time provided by the transmitter between the transition of the REQ signal and the transition of the P\_CRCA signal during data group transfers.

#### **9.2.20 pCRC transmit setup time**

The minimum time provided by the transmitter between the transition of the P\_CRCA signal and the transition of the REQ signal during data group transfers.

Specified to provide the increased receive setup time, subject to intersymbol interference, cable skew, and other distortions.

#### **9.2.21 Data release delay**

The maximum time for a SCSI initiator port to release the DATA BUS, DB(P\_CRCA), and/or DB(P1) signals, following the transition of the I/O signal from false to true.

#### **9.2.22 DIFFSENS voltage filter time**

The minimum time DIFFSENS voltage shall be sensed continuously within the voltage range of a valid SCSI bus mode.

#### **9.2.23 Physical disconnection delay**

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

#### **9.2.24 Power on to selection**

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

**9.2.25 QAS arbitration delay**

The minimum time a SCSI device with QAS enabled (see 4.12.4.6.4) shall wait from the detection of the MSG, C/D, and I/O signals being false to start QAS until the DATA BUS is examined to see if QAS has been won (see 10.4).

**9.2.26 QAS assertion delay**

The maximum time allowed for a SCSI device to assert certain signals during QAS.

**9.2.27 QAS release delay**

The maximum time allowed for a SCSI device to release certain signals during QAS.

**9.2.28 QAS non-DATA phase REQ(ACK) period**

The minimum time a QAS-capable SCSI initiator port shall ensure the REQ and ACK signals are asserted and that data is valid during COMMAND, MESSAGE, and STATUS phases.

**9.2.29 Receive assertion period**

The minimum time provided at a SCSI device receiving a REQ signal for the signal to be asserted while using synchronous transfers or paced transfers, provided P\_CRCA is not transitioning during data group transfers. Also, the minimum time required at a SCSI device receiving an ACK signal for the signal to be asserted while using synchronous transfers or paced transfers. See figure 60 and figure 61 for signal measurement points.

**9.2.30 Receive hold time**

For ST data transfers the minimum time provided at the receiving SCSI device between the assertion of the REQ signal or the ACK signals and the changing of the DATA BUS, DB(P\_CRCA), and/or DB(P1) signals while using synchronous transfers, provided P\_CRCA is not transitioning during data group transfers.

For DT data transfers the minimum time required at the receiving SCSI device between the transition (i.e. assertion or negation) of the REQ signal or the ACK signals and the changing of the DATA BUS, DB(P\_CRCA), and/or DB(P1) signals.

For paced data transfers negative values as measured at the device connector are accommodated by skew compensation in the receiver. Receive hold time measured at the device connector shall not exceed the skew correction range.

**9.2.31 Receive internal hold time**

The minimum time provided for hold time in the receive data detector after allowance for timing errors and timing compensation from all sources measured from the worse case bit (i.e., data or parity) to the compensated offset strobe.

NOTE 17 - This time may not be observable to other than the SCSI device designer.

**9.2.32 Receive internal setup time**

The minimum time provided for setup time in the receive data detector after allowance for timing errors and timing compensation from all sources measured from the worse case bit (i.e., data or parity) to the compensated offset strobe.

NOTE 18 - This time may not be observable to other than the SCSI device designer.

**9.2.33 Receive negation period**

The minimum time required at a SCSI device receiving a REQ signal for the signal to be negated while using synchronous transfers or paced transfers. Also, the minimum time required at a SCSI device receiving an ACK signal for the signal to be asserted while using synchronous transfers or paced transfers. See figure 60 and figure 61 for signal measurement points.

**9.2.34 Receive setup time**

For ST data transfers the minimum time provided at the receiving SCSI device between the changing of DATA BUS, DB(P\_CRCA), and/or DB(P1) signals and the assertion of the REQ signal or the ACK signal while using synchronous transfers.

For DT data transfers the minimum time required at the receiving SCSI device between the changing of DATA BUS, DB(P\_CRCA), and/or DB(P1) signals and the transition of the REQ signal or the ACK signal.

For paced data transfers negative values as measured at the device connector are accommodated by skew compensation in the receiver. Receive setup time measured at the device connector shall not exceed the skew correction range.

**9.2.35 Receive REQ(ACK) period tolerance**

The minimum tolerance that a SCSI device shall allow to be subtracted from the REQ(ACK) period. The tolerance comprises the transmit REQ(ACK) tolerance plus a measurement error due to noise.

**9.2.36 Receive REQ assertion period with P\_CRCA transition**

The minimum time required at a SCSI device receiving a REQ signal for the signal to be asserted while P\_CRCA is transitioning during data group transfers.

Specified to ensure that the assertion period is longer than the receive hold time plus the receive setup time.

**9.2.37 Receive REQ negation period with P\_CRCA transition**

The minimum time required at a SCSI device receiving a REQ signal for the signal to be negated while P\_CRCA is transitioning during data group transfers.

Specified to ensure that the negation period is longer than the receive hold time plus the receive setup time.

**9.2.38 Receive Skew Compensation**

The effective reduction in worst case timing skew of data, parity, and strobe signals provided by the receiving SCSI device but not directly observable at the receiving SCSI device connector.

**9.2.39 Receiver amplitude time skew**


The maximum time shift of SCSI bus signals caused by the difference in receiver switching delay of a minimum amplitude signal versus a maximum amplitude signal.

**9.2.40 Receiver de-skewed data valid window**

The minimum difference in time allowed between the rising or falling edge of a "1010..." pattern on the DATA BUS or DB(P1) and its clocking signal on the ACK or REQ signal as measured at their zero-crossing points after skew compensation is applied by the receiver without allowing any error in the received data (see figure 56). The receiver de-skewed data valid window shall be equal to:



$\pm [(\text{data transfer period}) - (\text{residual skew error}) - (\text{strobe offset tolerance}) - (\text{receiver asymmetry}) - (\text{chip noise at receiver}) - (\text{system noise at receiver})] / 2$ .



**Figure 56 - Receiver de-skew parameters**

#### 9.2.41 REQ(ACK) period

The REQ(ACK) period during synchronous transfers or paced transfers, specified in table 35 for ST DATA phases and in table 36 for DT DATA phases, is the nominal time between adjacent assertion edges of the REQ or ACK signal for the fastest negotiated transfer rate. For the purpose of calculating the actual REQ(ACK) period tolerance the REQ(ACK) period should be measured without interruptions (e.g., offsets pauses). To minimize the impact of crosstalk time shift and ISI the measurements should be made by averaging the time between edges during long (i.e., greater than 512 bytes) all zero or all ones data transfers and by ignoring the first and last 10 transitions.

In DT DATA phases the negotiated transfer period for data is half of the REQ(ACK) period since data is qualified on both the assertion and negation edges of the REQ or ACK signal. In ST DATA phases the negotiated transfer period for data is equal to the REQ(ACK) period during synchronous transfers since data is only qualified on the assertion edge of the REQ or ACK signal.

#### 9.2.42 Reset delay

The minimum time that the RST signal shall be continuously true before the SCSI device shall initiate a hard reset (see 12.4).

#### 9.2.43 Reset hold time

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

#### 9.2.44 Reset to selection

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

**9.2.45 Residual skew error**

The maximum timing error between the deskewed data and REQ or ACK internal to the receiving SCSI device after skew compensation.

**9.2.46 Selection abort time**

The maximum time that a 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 SCSI target port or SCSI initiator port does not assert the BSY signal after a SELECTION or RESELECTION phase has been aborted.

**9.2.47 Selection time-out delay**

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

**9.2.48 Signal timing skew**


The maximum signal timing skew occurs when transferring random data and in combination with interruptions of the REQ or ACK signal transitions (e.g., pauses caused by offsets). The signal timing skew includes cable skew that is measured with 0101... patterns and signal distortion skew caused by random data patterns and transmission line reflections as shown in figure 60 and figure 61.

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)

NOTE 19 - For timing budget purposes the value stated in table 36 is calculated without the benefit of skew compensation.

**9.2.49 Skew correction range**

The minimum skew correction capability of the receiver of a signal on the DATA BUS or DB(P1) relative to the ACK or REQ signal as measured at the receiver's connector. The skew correction range shall be equal to  $\pm [(transmitter\ chip\ skew) + (cable\ skew) + (trace\ skew)]$  relative to the corresponding ACK or REQ clock signal for that transition (see figure 57). Receiver chip skew and the receiver connector to receiver trace skew is not included, as it is internal to the receiving SCSI device.



**Figure 57 - Skew correction range**

#### 9.2.50 Strobe offset tolerance

The time tolerance of centering the compensated REQ or ACK strobe in the transfer period during the training pattern.

#### 9.2.51 System deskew delay

The minimum time that a SCSI device should wait after receiving a SCSI signal to ensure that with asynchronous transfers at the same time are valid. The system deskew delay shall not be applied to the synchronous transfers or paced transfers.

#### 9.2.52 System noise at launch

The maximum time shift of SCSI bus signals caused by system noise at the transmitter (e.g., noise caused by current changes in the voice coil) measured at the transmitting SCSI device connector.

#### 9.2.53 System noise at receiver

The maximum time shift of SCSI bus signals caused by system noise at the receiver (e.g., noise caused by current changes in the voice coil) measured at the receiving SCSI device connector not including the time shift from the system noise at launch.

#### 9.2.54 Time asymmetry

The maximum time difference between the asserted and negated signal for data, REQ, or ACK transitions that are intended to be equidistant.

#### 9.2.55 Transmit assertion period

The minimum time that a SCSI target port shall assert the REQ signal while using synchronous transfers or paced transfers, provided it is not transitioning P\_CRCA during data group transfers. Also, the minimum time that a SCSI initiator port shall assert the ACK signal while using synchronous transfers or paced

transfers.

#### **9.2.56 Transmit hold time**

For ST data transfers the minimum time provided by the transmitting SCSI device between the assertion of the REQ signal or the ACK signal and the changing of the DATA BUS, DB(P\_CRCA), and/or DB(P1) signals while using synchronous transfers. For DT data transfers the minimum time provided by the transmitting SCSI device between the transition of the REQ signal or the ACK signal and the changing of the DATA BUS, DB(P\_CRCA), and/or DB(P1) signals.

#### **9.2.57 Transmit ISI Compensation**

The effective reduction in worst case ISI timing shift provided by the transmitting SCSI device as seen at the receiving SCSI device connector.

#### **9.2.58 Transmit negation period**

The minimum time that a SCSI target port shall negate the REQ signal while using synchronous transfers or paced transfers, provided it is not transitioning P\_CRCA during data group transfers. Also, the minimum time that a SCSI initiator port shall negate the ACK signal while using synchronous transfers or paced transfers.

#### **9.2.59 Transmit setup time**

For ST data transfers the minimum time provided by the transmitting SCSI device between the changing of DATA BUS, DB(P\_CRCA), and/or DB(P1) signals and the assertion of the REQ signal or the ACK signal while using synchronous transfers. For DT data transfers the minimum time provided by the transmitting SCSI device between the changing of DATA BUS, DB(P\_CRCA), and/or DB(P1) signals and the transition of the REQ signal or the ACK signal.

#### **9.2.60 Transmit REQ(ACK) period tolerance**

The maximum tolerance that a SCSI device may subtract from the REQ(ACK) period.

#### **9.2.61 Transmit REQ assertion period with P\_CRCA transitioning**

The minimum time that a SCSI target port shall assert the REQ signal while transitioning P\_CRCA during data group transfers.

Specified to provide the increased receive REQ assertion period, subject to loss on the interconnect.

#### **9.2.62 Transmit REQ negation period with P\_CRCA transitioning**

The minimum time that a SCSI target port shall negate the REQ signal while transitioning P\_CRCA during data group transfers.

Specified to provide the increased receive REQ negation period, subject to loss on the interconnect.

#### **9.2.63 Transmitter skew**

The maximum difference in time allowed between the rising or falling edge of a “1010...” pattern on the DATA BUS or DB(P1) signal and its clocking signal on the ACK or REQ signal as measured at their zero-crossing points (see figure 58). The signals for the output waveforms shall be measured at the connector of the transmitting device under the test conditions as described in A.2.6.




Figure 58 - Transmitter skew

### 9.2.64 Transmitter time asymmetry

The maximum time on DATA BUS, DB(P1), ACK, or REQ signal from any transition edge to the subsequent transition edge during a “1010...” pattern, as measured at their zero-crossing points, minus the data transfer period (see figure 59). The signals for the output waveforms shall be measured at the connector of the transmitting device under the test conditions as described in A.2.6.




Figure 59 - Transmitter time asymmetry

## 9.3 Measurement points

### 9.3.1 Measurement points overview

The measurements points for LVD ACK, REQ, DATA, P\_CRCA, and PARITY signals are defined in 9.3. All measurements are at the SCSI connector.

When paced transfers are enabled the timing shall be measured relative to the zero crossing of the

differential signal. In paced transfers the timing budget and receiver masks account for the differences between the setup and hold detection thresholds that occur for synchronous transfers.

### 9.3.2 LVD measurement points

When transferring data using ST DATA phases LVD SCSI devices shall use the measurement points defined in figure 60 for the measurement of the timing values. When transferring data using DT DATA phases LVD SCSI devices using synchronous transfers shall use the measurement points defined in figure 61 for the measurement of the timing values. When transferring data using DT DATA phases LVD SCSI devices using paced transfers shall use the measurement points defined in figure 62 for the measurement of the timing values. The rise and fall times for the LVD REQ and ACK signals shall be nominally the same as for the LVD DATA, P\_CRCA, and DB(P1) signals.

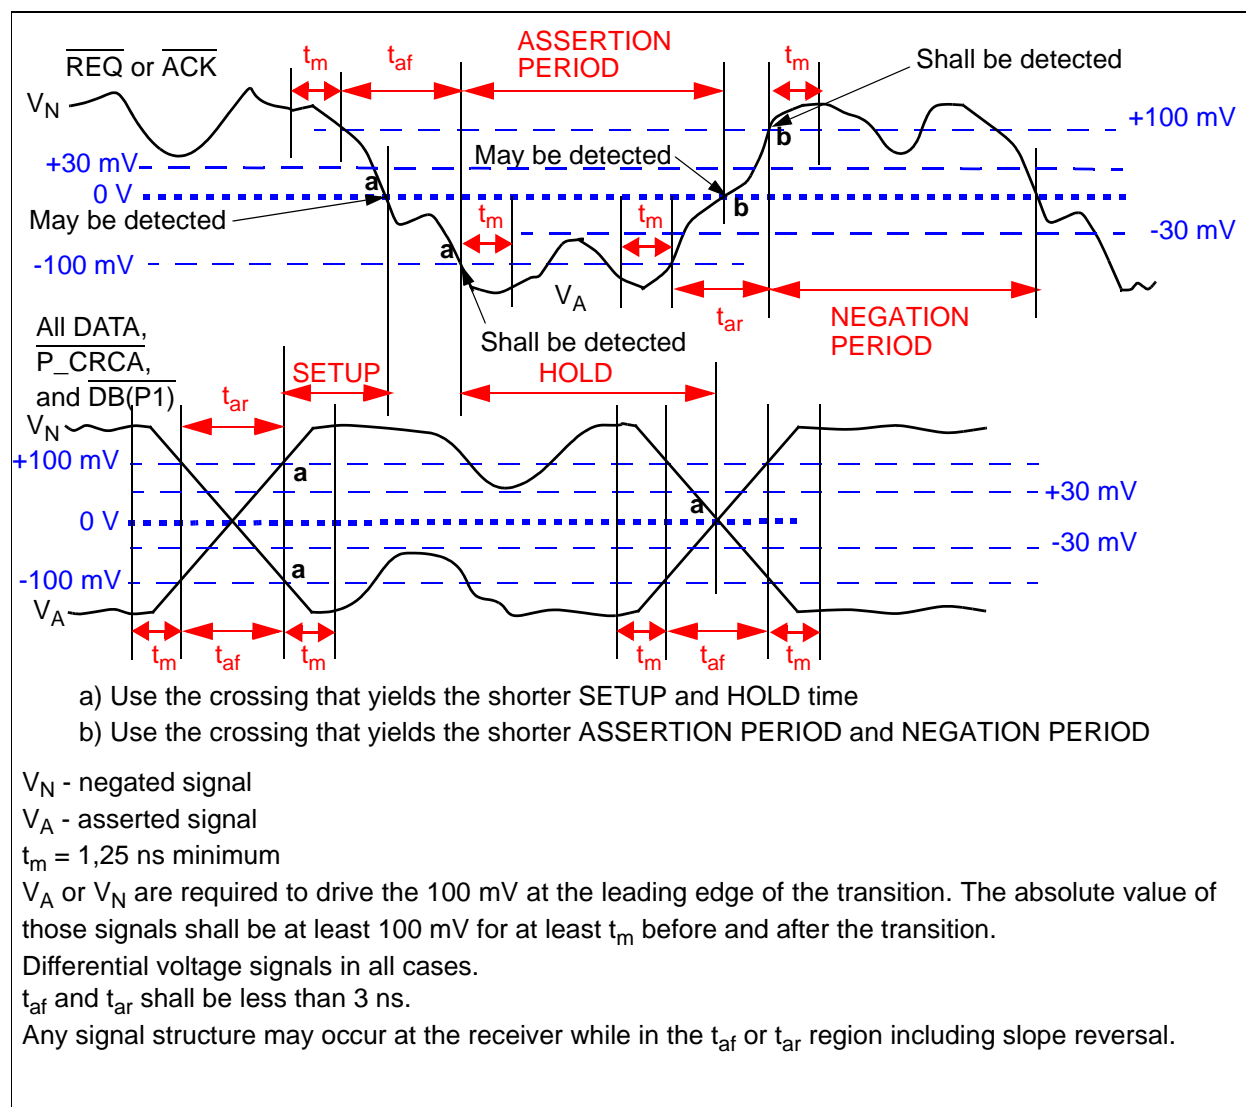



Figure 60 - LVD timing measurement points for ST synchronous transfers



**Figure 61 - LVD timing measurement points for DT synchronous transfers**

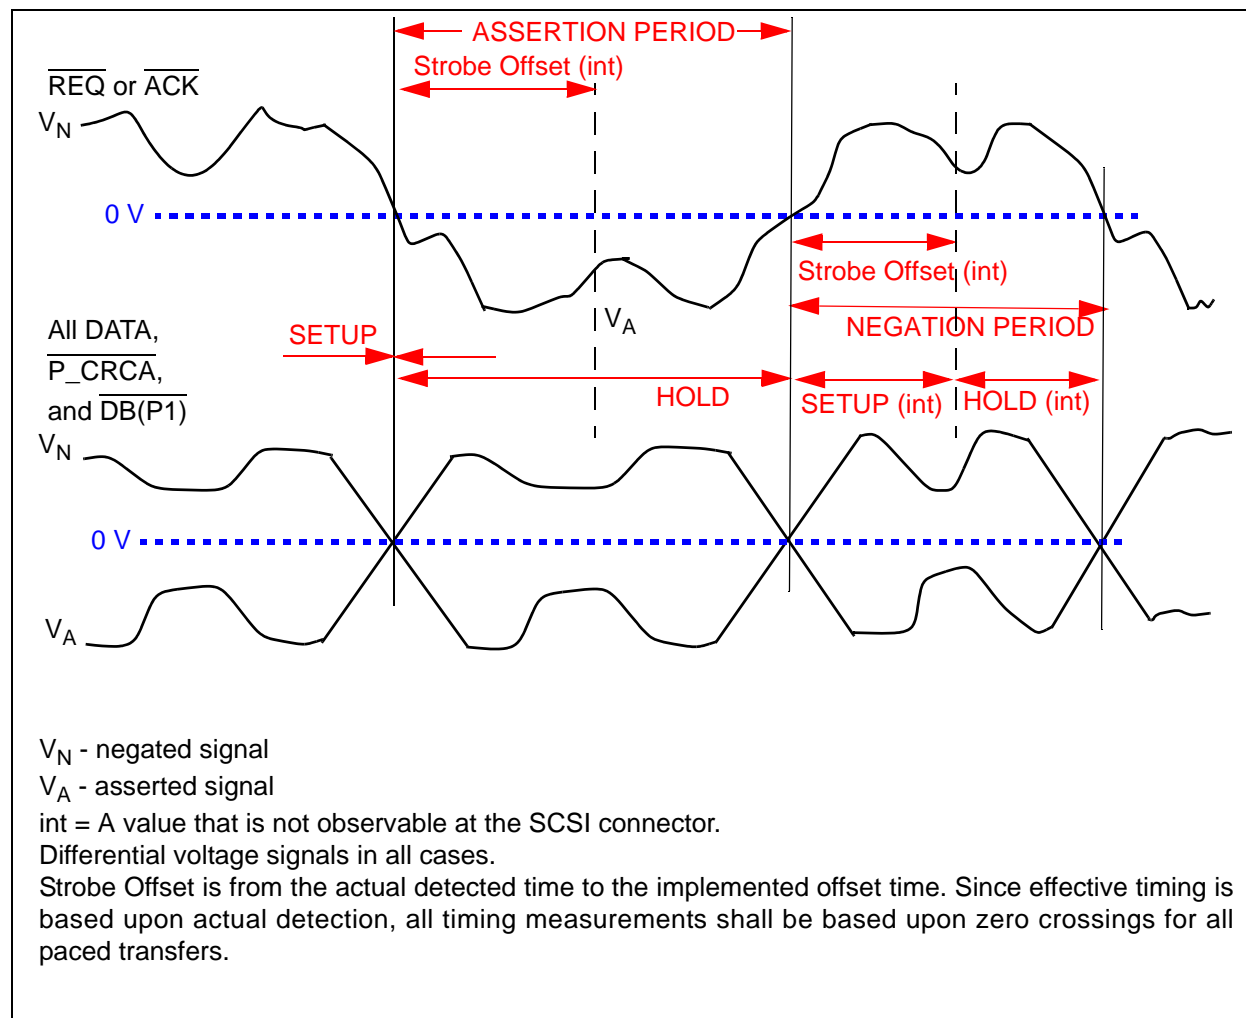


Figure 62 - LVD timing measurement points for DT paced transfers

## 9.4 Receiver Masks

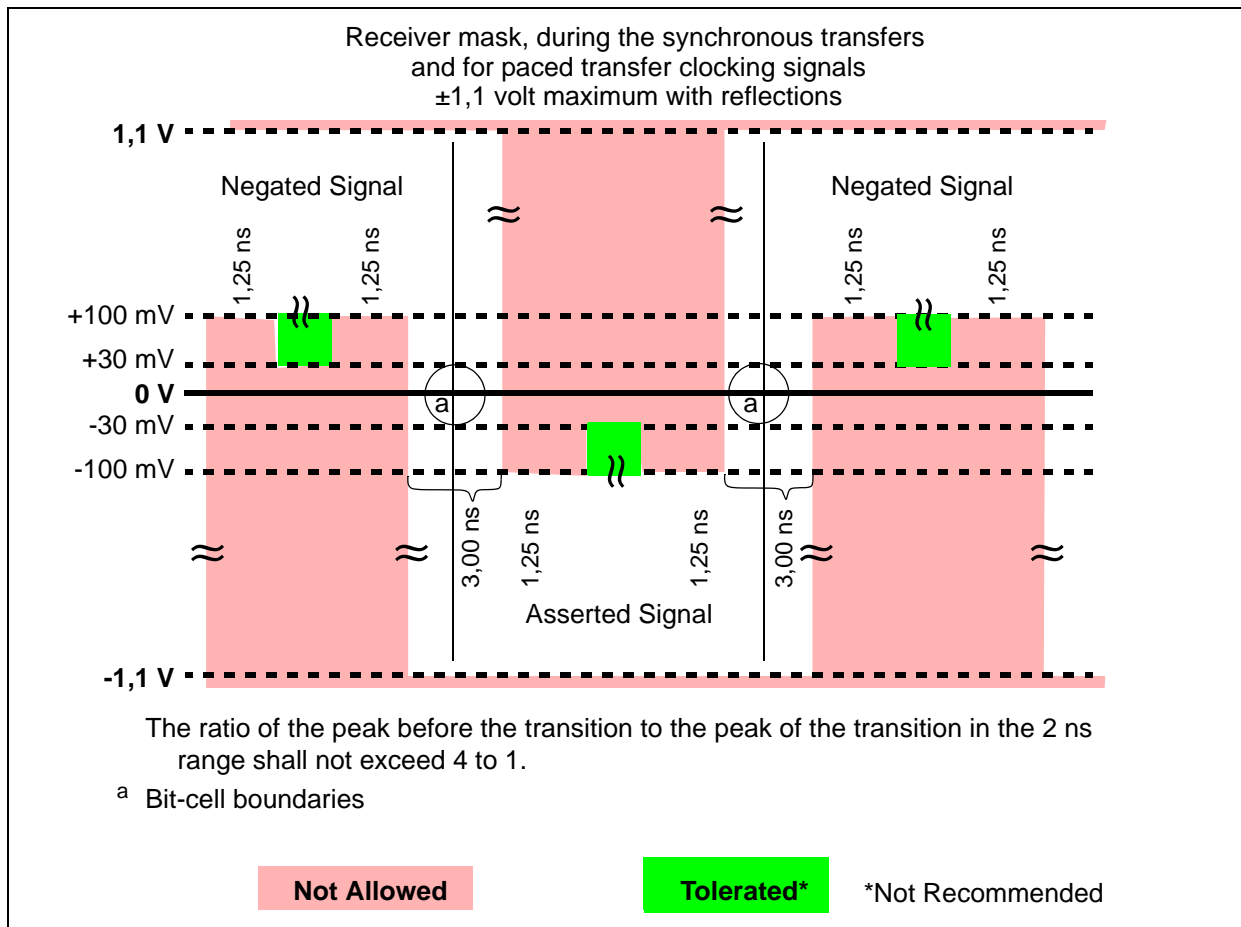
### 9.4.1 Synchronous transfers and for fast-160 paced transfer clocking signals

Figure 63 shows the LVD signal requirements at the receiving SCSI device with synchronous transfers and fast-160 paced transfer clocking signals.

For synchronous transfers and paced transfer clocking signals the signals shall transition from -100 to +100 mV or +100 to -100 mV in 0 to 3 ns, the waveform between -100 and +100 mV is not otherwise specified. The absolute value of the signals shall remain above the 100 mV level for 1,25 ns at each end of the transition (see figure 63). The absolute value of the signals shall not drop below 30 mV except during the transitions (see 7.2.2).

Bit-cell boundaries are determined by the zero crossings of the REQ or ACK signal shifted by the signal timing skew, excluding ISI, between the REQ or ACK signal and the measured signal using the same conditions defined in figure 66. When measuring a data signal in figure 63, figure 64, figure 65, figure 67, figure 68, figure 69, and figure 70 the REQ or ACK signal shall be used as the timing reference for aligning data bit cell boundaries.





**Figure 63 - LVD receiver mask for synchronous transfers and for fast-160 paced transfer clocking signals**

#### 9.4.2 Paced transfers with precompensation enabled on fast-160

During paced transfers with a negotiated transfer rate of fast-160 and precompensation enabled the clocking signals shall meet all the requirements specified in 9.4.1.

For non-clocking signals the SCSI devices shall operate with signals at the receiving SCSI device meeting either figure 64, figure 65, or both. Mask 2 is applicable to signals that have more timing margin than those for mask 1 and allows less amplitude margin than does mask 1. The lower amplitude margin of mask 2 may result in timing margin loss internal to the receiver but is accounted for in the timing budget. The higher amplitude margin of mask 1 should result in less timing margin loss internal to the receiver. For the cases where the data signal remains at a particular bit state without transitions for more than one transfer period, these masks include a variable asserted or negated period as a function of  $n$  where  $n$  is the number of transfer periods with adjacent data at the same state.

A non-clocking signal shall transition from -130 to +130 mV or +130 to -130 mV in 0 to 4,25 ns, the waveform between -130 and +130 mV is not otherwise specified. The absolute value of a non-clocking signal shall remain above the 130 mV level for 1 ns at before and after each transition (see figure 64). The absolute value of a non-clocking signal shall not drop below 30 mV except during the transitions (see 7.2.2).

Alternatively, a non-clocking signal shall transition from -80 to +80 mV or +80 to -80 mV in 0 to 3,25 ns, the waveform between -80 and +80 mV is not otherwise specified. The absolute value of a non-clocking

signal shall remain above the 80 mV level for 1,25 ns before and after each transition (see figure 65). The absolute value of a non-clocking signal shall not drop below 30 mV except during the transitions (see 7.2.2).

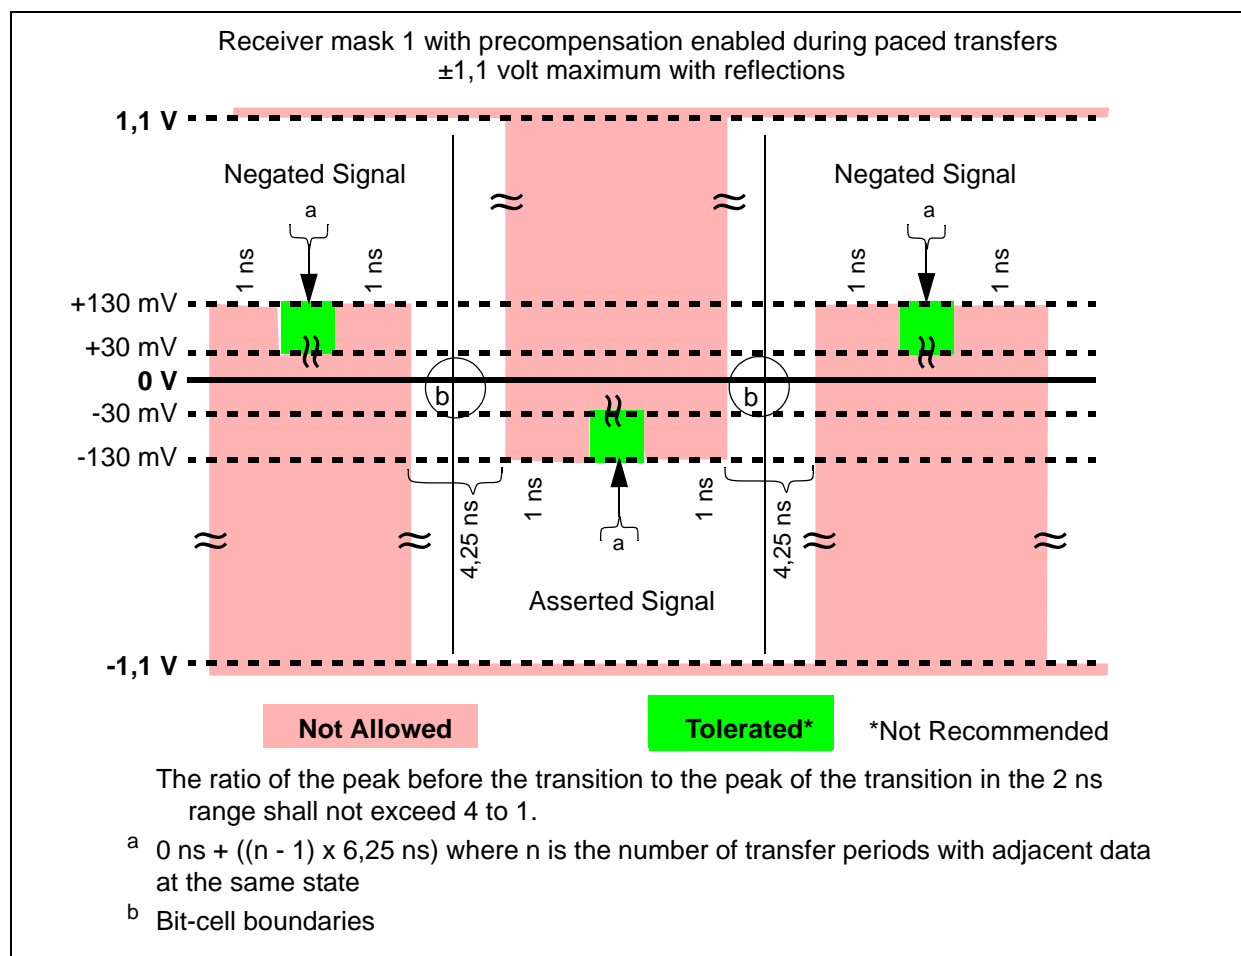


Figure 64 - Fast-160 LVD receiver mask 1 for paced transfers with precompensation enabled




Figure 65 - Fast-160 LVD receiver mask 2 for paced transfers with precompensation enabled

#### 9.4.3 Paced transfers with precompensation disabled on fast-160 and fast-320

During paced transfers with a negotiated transfer rate of fast-160 and precompensation disabled, or a negotiated transfer rate of fast-320 receiving SCSI devices shall operate with signals that meet all the following requirements at the receiver.

- A change in amplitude for all transitions shall be greater than 100 mV (see figure 67).
- The differential low frequency signal amplitude not shall exceed the following limits:

$$600 \text{ mV} < (|V_A| + |V_N|) < 1,7 \text{ V peak-to-peak}$$

Where:

- $V_A$  and  $V_N$  shall be measured after 100 ns continuous assertion or negation of the signal; and
- the low frequency is 2,5 MHz or less.

NOTE 20 - The section A of the training pattern meets the low frequency requirements (see 10.7.4.2).

- The differential offset component of the signal shall not exceed the following limits.

For fast-160:

$$-50 \text{ mV} < [(|V_A| - |V_N|) / 2] < 50 \text{ mV peak-to-peak}$$

For fast-320:

$$-75 \text{ mV} < [(|V_A| - |V_N|) / 2] < 75 \text{ mV peak-to-peak}$$

Where:

- a)  $V_A$  and  $V_N$  shall be measured after 100 ns continuous assertion or negation of the signal.

NOTE 21 - The section A of the training pattern meets the low frequency requirements (see 10.7.4.2).

- d) The amplitude of any signal transitioning at the negotiated transfer period (e.g., clocking signals, any data signal with a 1010 bit pattern) shall be at least 240 mV derived as follows:

$$240 \text{ mV} < (|V_A| \text{ min} + |V_N| \text{ min})$$

Where:

- a)  $V_A$  and  $V_N$  shall not be measured until at least 20 transitions have occurred on the signal being measured;  
 b) for fast-160,  $V_A$  and  $V_N$  shall be measured within a 2 ns interval centered on the assertion or negation pulse (see figure 66); and  
 c) for fast-320,  $V_A$  and  $V_N$  shall be measured within a 1 ns interval centered on the assertion or negation pulse (see figure 66).

NOTE 22 - The section A of the training pattern meets the requirements described in this subclause (see 10.7.4.2).


- e) All non-clocking signals shall meet the requirements defined for the isolated negation pulse receiver mask (see figure 69) and the isolated assertion pulse receiver mask (see figure 68).

The receiver input masks only measures the A.C components of the input waveforms. Any D.C. differential offset component shall be removed before applying the mask.


Receiver input eye masks in figure 66, figure 67, figure 68, figure 69, and figure 70 define the requirements described in this subclause using the parameters in table 41.

**Table 41 - Receiver Eye Mask Values**

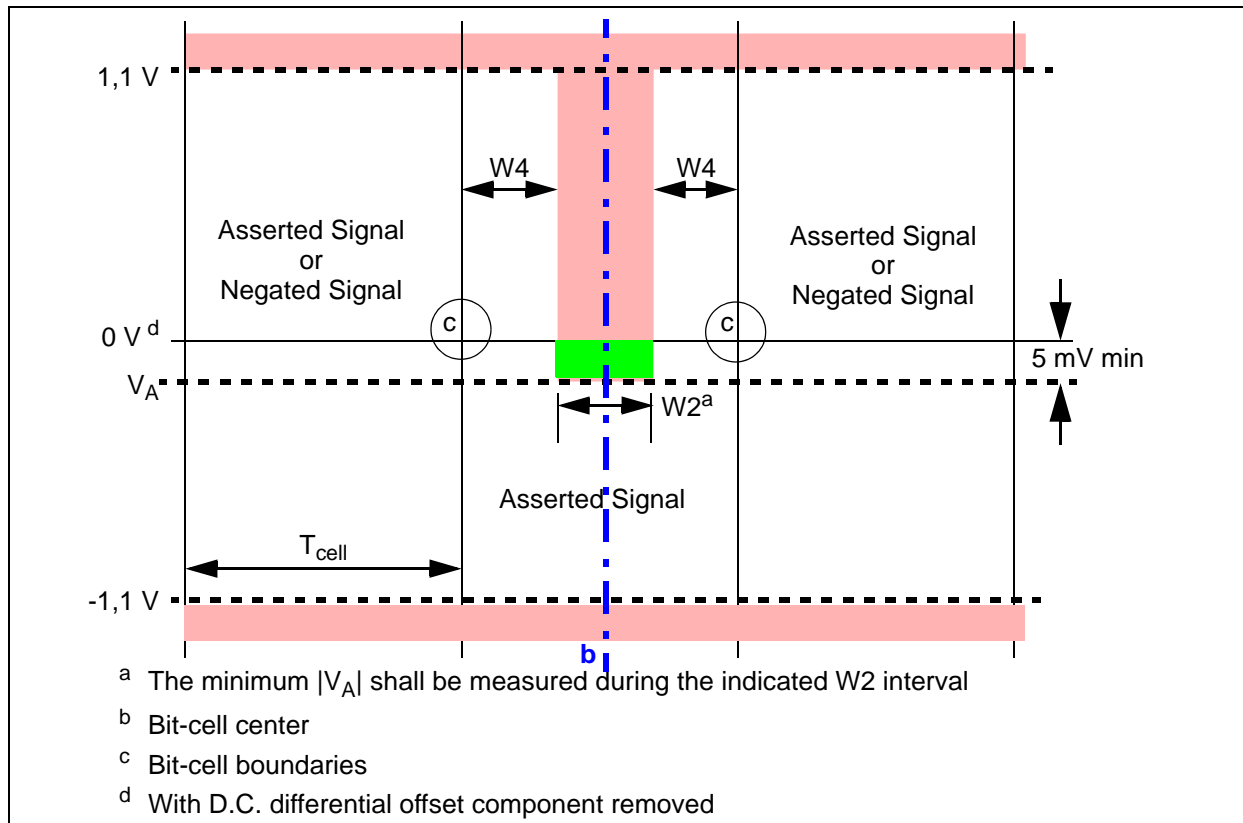
Parameter	Fast-160	Fast-320
$T_{\text{cell}}$	6,25 ns	3,125 ns
W1	2,0 ns	1,0 ns
W2	1,50 ns	0,75 ns
W3	2,125 ns	1,063 ns
W4	2,375 ns	1,188 ns



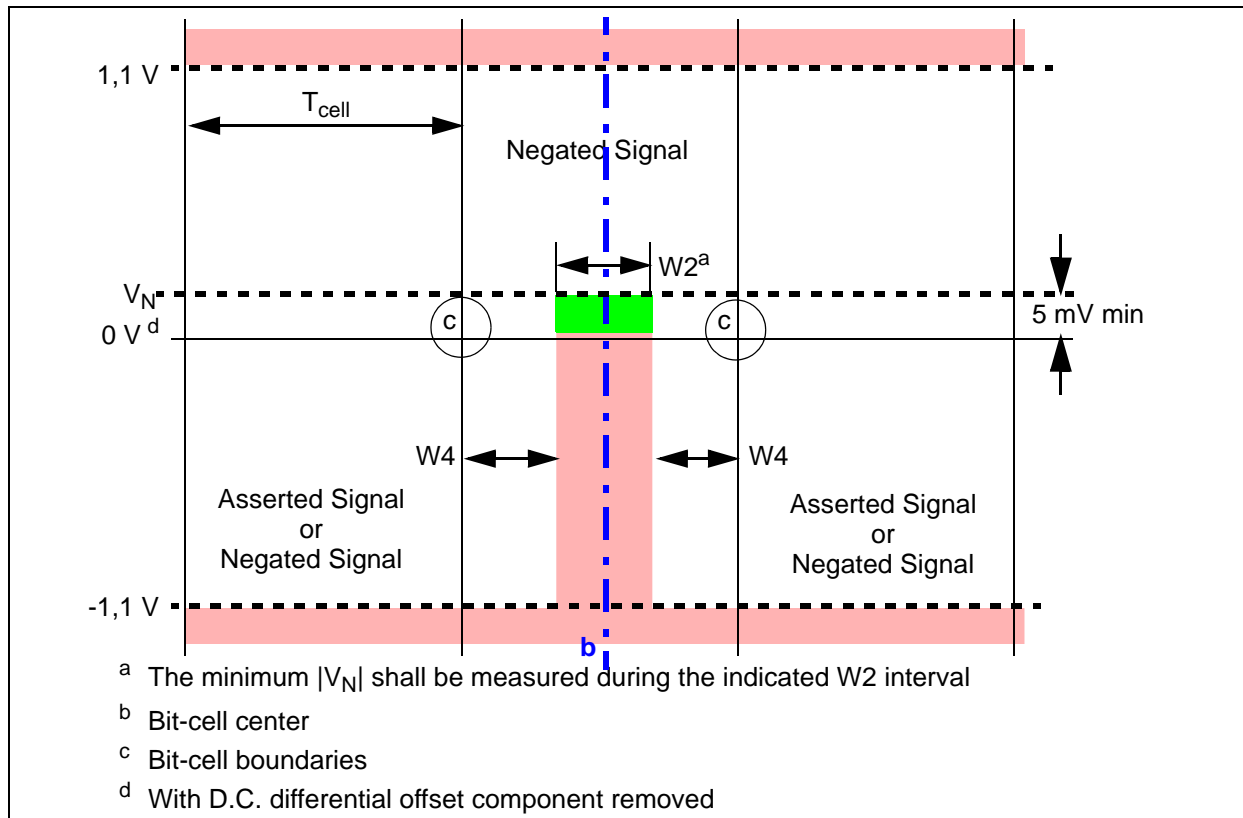
**Figure 66 - LVD paced transfer receiver mask non-precompensated signals transitioning at the negotiated transfer period**



**Figure 67 - LVD paced transfer receiver mask non-precompensated minimum amplitude change**




**Figure 68 - LVD paced transfer receiver mask non-precompensated non-clocking isolated assertion pulse**



**Figure 69 - LVD paced transfer receiver mask non-precompensated non-clocking isolated negation pulse**

The eye diagram is formed by superimposing the waveforms of all bit cells of the input data pattern. For fast-160 the input data pattern shall be open for 1,5 ns centered on the bit-cell center and for fast- 320 the input data pattern shall be open for 0,75 ns centered on the bit-cell center (see figure 70). The bit-cell center location is determined by the clocking signal after at least 20 transitions on the clocking signal as in figure 66. The receiver input eye only measures the A.C components of the input waveforms. Any D.C. differential offset component shall be removed before applying the mask.





**Figure 70 - LVD paced transfer eye mask non-precompensated non-clocking pulse**

## 9.5 Timing parameters

Driver timing parameters applied at the connector of the transmitting SCSI device are defined in figure 60, figure 61, and figure 62. The driver timing parameters shall be measured using the circuit and test conditions defined in A.2.6. Receiver timing parameters applied at the connector of the receiving SCSI device are defined in figure 60, figure 61 and figure 62. The receiver timing parameters include the effects of arbitrary data patterns and REQ/ACK pauses.

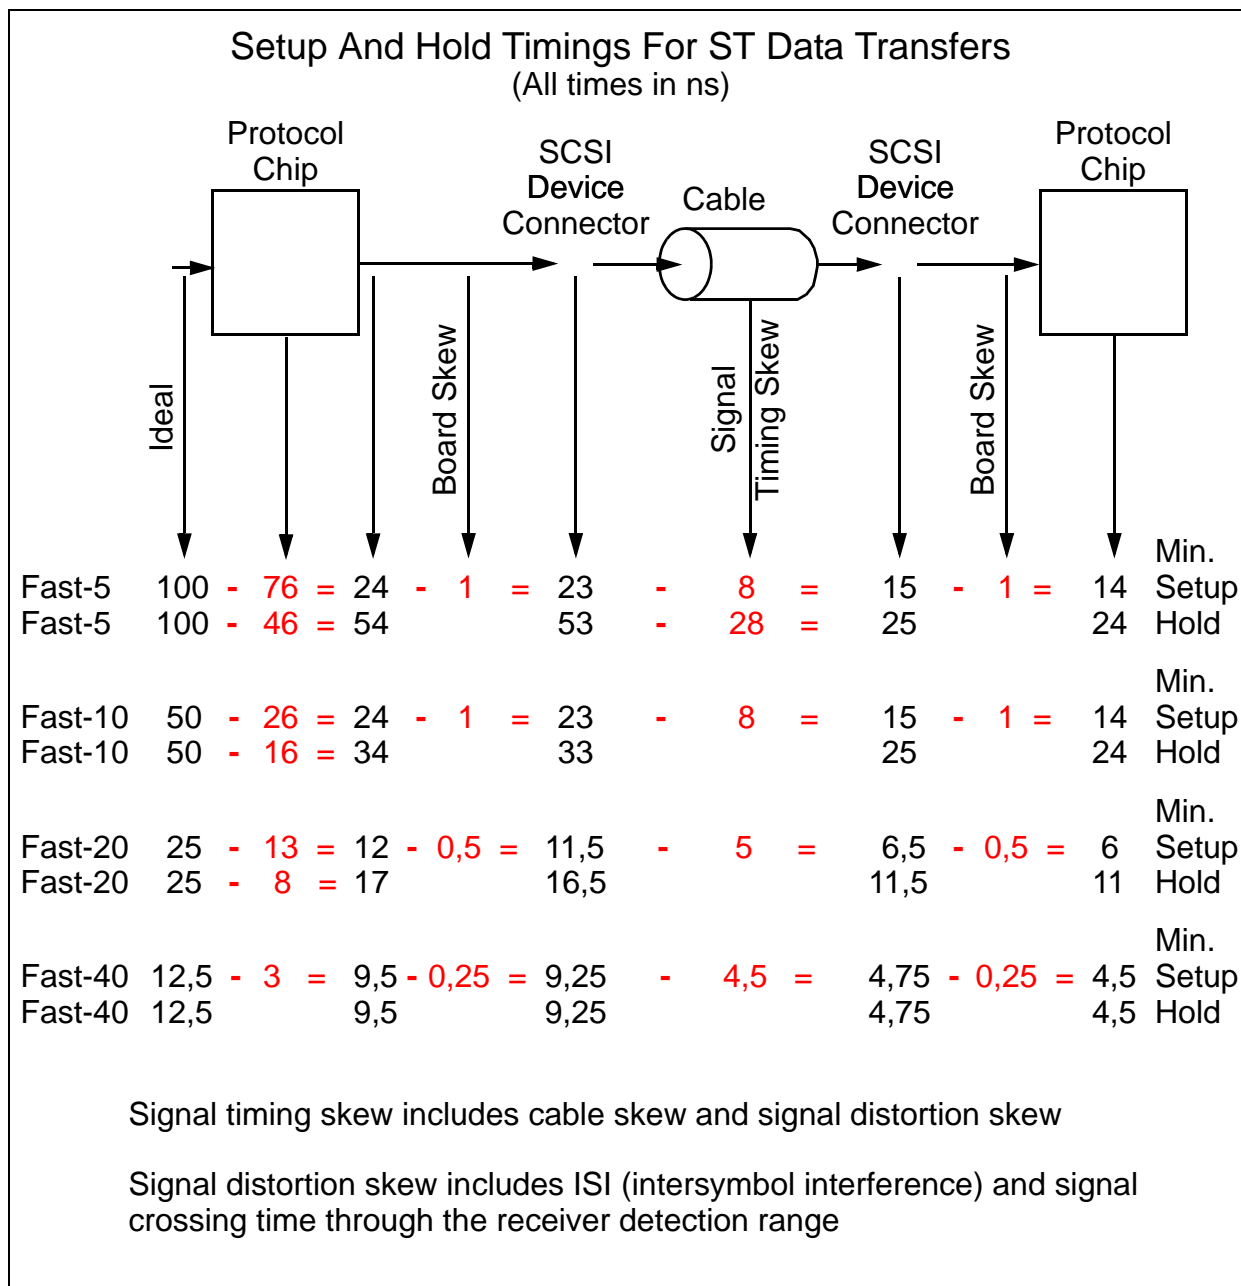
## 9.6 Setup and hold timings

### 9.6.1 ST data transfer calculations

Figure 71 shows how the setup and hold times are calculated for various physical configurations on SCSI devices that support ST data transfers. The minimum set up and hold timings specified in figure 71 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 of attenuation need to be covered:

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

Table 35 specifies setup and hold times at the device connector. Figure 71 illustrates a possible timing budget behind the device connector, with time apportioned to board skew and to the protocol chip.



**Figure 71 - System setup and hold timings for ST data transfers**


### 9.6.2 DT data transfer calculations

Figure 72 shows how the setup and hold times are calculated for various physical configurations using synchronous transfers on SCSI devices that support DT data transfers. The minimum set up and hold timings specified in figure 72 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 of attenuation need to be covered:

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

Table 36 specifies setup and hold times at the device connector. Figure 72 illustrates possible timing

budgets behind the device connector, with time apportioned to board skew and to the protocol chip.



**Figure 72 - System setup and hold timings for DT synchronous transfers**

## 10 SCSI bus phases

### 10.1 SCSI bus phases overview

The SCSI architecture includes the following 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.

### 10.2 BUS FREE phase

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

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

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

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

### 10.3 Expected and unexpected bus free phases

SCSI target ports shall create a BUS FREE phase after any of the following:

- a) after any bus reset event (see 12.5.2);
- b) after an ABORT TASK task management function is successfully received by a SCSI target port (see 16.5.2 and 14.3.1);
- c) after an ABORT TASK SET task management function is successfully received by a SCSI target port (see 16.5.3 and 14.3.1);
- d) after a CLEAR TASK SET task management function is successfully received by a SCSI target port (see 16.5.5 and 14.3.1);
- e) after a LOGICAL UNIT RESET task management function is successfully received by a SCSI target port (see 16.5.6 and 14.3.1);
- f) after a TARGET RESET task management function is successfully received by a SCSI target port (see 16.5.7 and 14.3.1);
- g) after a CLEAR ACA task management function is successfully received by a SCSI target port (see 16.5.4 and 14.3.1);

- h) after a DISCONNECT message is successfully transmitted from a SCSI target port (see 16.3.2);
- i) after a TASK COMPLETE message is successfully transmitted from a SCSI target port (see 16.3.17);
- j) after a DISCONNECT message is successfully received by a SCSI target port when information unit transfers are enabled (see 16.3.2);
- k) after the release of the SEL signal after a SELECTION or RESELECTION phase time-out;
- l) after a PPR negotiation in response to a selection using attention condition when information unit transfers are enabled (see 4.12); or
- m) after any successful negotiation that causes information unit transfers to be enabled (see 4.12) or disabled (see 4.12).

SCSI target ports may create a BUS FREE phase after one of the following:

- a) after the last SPI command information unit is successfully received by a SCSI target port;
- b) after a SPI data information unit is successfully received by or transmitted from a SCSI target port;
- c) after a SPI status information unit is successfully transmitted from a SCSI target port;
- d) after a SPI L\_Q information unit if the SPI L\_Q information unit DATA LENGTH field is set to zero; or
- e) during a QAS phase.

An unexpected bus free occurs when a SCSI initiator port detects a BUS FREE phase that it does not expect.

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

The SCSI target device shall terminate the task that was the current task before the BUS FREE phase by clearing all data and status for that task. The SCSI target device 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 SCSI initiator device shall terminate the task that was the current task before the BUS FREE phase occurred and shall manage this condition as an exception condition.

## 10.4 Arbitration

### 10.4.1 Arbitration and QAS overview

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

There are two methods that a SCSI device may use to arbitrate for the SCSI bus: normal arbitration and QAS. Normal arbitration is mandatory and requires the detection of a BUS FREE phase on the SCSI bus before starting. QAS is optional and, when enabled (see 4.12.4.6.4), requires the detection of a QAS REQUEST message (see 4.12.4.6.4) before starting.

SCSI devices with arbitration fairness enabled shall maintain a fairness register that records the SCSI IDs of SCSI device ports that need a chance to arbitrate (see Annex B). Fairness in normal arbitration is enabled in SCSI target ports by the Disconnect-Reconnect mode page (see 18.1.2). Fairness is always enabled when QAS is enabled.

### 10.4.2 NORMAL ARBITRATION phase

The procedure for a SCSI target port, either with QAS disabled or with QAS enabled and information unit transfers disabled, to indicate it wants to release the bus is by going to the BUS FREE state.

The procedure for a SCSI device to obtain control of the SCSI bus via normal arbitration is as follows:

- 1) 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 one bus settle delay.

NOTE 23 - 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.

- 2) The SCSI device shall wait a minimum of one bus free delay after detection of the BUS FREE phase (i.e., after the BSY and SEL signals are both false for one bus settle delay) before driving any signal.
- 3) Following the bus free delay in step 2), 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) during this NORMAL ARBITRATION phase if more than one bus set delay has passed since the BUS FREE phase was last observed. If arbitration fairness is enabled, the SCSI device shall not arbitrate until its fairness register is cleared (see Annex B).

NOTE 24 - There is no maximum delay before asserting the BSY signal and the SCSI ID following the bus free delay in step 2) as long as the bus remains in the BUS FREE phase. However, SCSI devices that delay longer than one bus settle delay plus one 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 and may not be ensured fair arbitration by the arbitration fairness algorithm.

- 4) After waiting at least one arbitration delay, measured from its assertion of the BSY signal, the SCSI device shall examine the DATA BUS.
  - 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.
  - B) If a higher priority SCSI ID bit is true on the DATA BUS (see table 31 for the SCSI ID arbitration priorities), then the SCSI device has lost the arbitration and the SCSI device shall release the BSY signal and the SCSI ID after the SEL signal becomes true, within one bus clear delay after the SEL signal becomes true. Any losing SCSI devices may return to step 1).

NOTE 25 - Step 4) above requires any SCSI device that begins NORMAL ARBITRATION phase to complete the NORMAL ARBITRATION phase to the point of SEL being asserted if it begins the NORMAL ARBITRATION phase as stated in step 3). This precludes the possibility of the bus being hung.

- 5) After the bus free delay in step 2), SCSI devices with arbitration fairness enabled that are not arbitrating shall wait one bus set delay and start sampling the DATA BUS to determine the SCSI devices that attempted arbitration, the SCSI device that won, and the SCSI devices that lost. This sampling shall continue for an arbitration delay after the bus free delay in step 2). Each SCSI device shall update its fairness register with all lower-priority device IDs that lost arbitration.

NOTE 26 - For ease of implementation, this sampling may begin when BSY is true following BUS FREE and end when SEL is true.

- 6) The SCSI device that wins arbitration shall wait at least one bus clear delay plus one 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. During the NORMAL ARBITRATION phase, DB(P\_CRCA), and DB(P1) (if present) may be released or asserted, but shall not be actively driven false.

#### 10.4.3 QAS protocol

QAS allows a SCSI target port with an information unit transfer agreement in effect and QAS enabled (see

4.12.4.6.4) that is currently connected to a SCSI initiator port that has QAS enabled to transfer control of the bus to another SCSI device that has QAS enabled without an intervening BUS FREE phase. SCSI devices that participate in QAS arbitration shall report that capability in the INQUIRY command.

Before a SCSI initiator port may use QAS that SCSI initiator port shall negotiate, using the PPR message, the use of the QAS phase with each SCSI target port that has indicated support of QAS. Any time a SCSI initiator port's negotiation required flag is true (see 4.12) that SCSI initiator port shall renegotiate to enable QAS. SCSI devices that support QAS shall implement the fairness algorithm (see Annex B) during all QAS arbitrations. SCSI devices shall negotiate the use of QAS with a particular SCSI device before using QAS to select or reselect that SCSI device. Also, SCSI target ports shall have negotiated the use of QAS with a particular SCSI initiator port before using QAS REQUEST message to do a physical disconnect from that SCSI initiator port, and SCSI initiator ports shall have negotiated the use of QAS with a particular SCSI target port before accepting a QAS REQUEST message from that SCSI target port. If a SCSI initiator port receives a QAS REQUEST message from a SCSI target port that has not negotiated the use of QAS, then the SCSI initiator port shall create an attention condition for the QAS REQUEST message, and shall report MESSAGE REJECT on the following MESSAGE OUT phase.

In an environment where some SCSI devices have QAS enabled and other SCSI devices do not, it is possible for the SCSI devices that have QAS enabled to prevent SCSI devices that do not have QAS enabled from arbitrating for the bus. This occurs when SCSI devices that have QAS enabled never go to a BUS FREE phase.

A QAS SCSI initiator port may interrupt a sequence of QAS cycles to force a normal arbitration with the following procedure:

- 1) Perform a QAS arbitration;
- 2) On winning QAS arbitration, continue driving the SCSI initiator port's ID on the DATA BUS instead of asserting SEL to enter selection phase;
- 3) Wait until the SCSI target port transitions to BUS FREE (i.e., after two QAS arbitration delays);
- 4) After detecting BSY false, release the DATA BUS; and
- 5) After one bus settle delay from when the SCSI target port drove BSY false, the bus is in BUS FREE phase. The SCSI initiator port may then arbitrate using normal arbitration and perform a selection if it wins.

#### 10.4.4 QAS phase

The procedure for a SCSI target port with both information unit transfers and QAS enabled to indicate it wants to release the bus after a DT DATA phase is as follows:

- 1) The SCSI target port shall change to a MESSAGE IN phase and issue a single QAS REQUEST message (see 10.11.2) and wait for ACK to be true.
- 2) After detection of the ACK signal being false and if the SCSI initiator port did not create an attention condition, the SCSI target port shall release all SCSI signals except the BSY, MSG, C/D, I/O, and REQ signals. Then the SCSI target port shall negate the MSG, C/D, and I/O signals within two system deskew delays. The SCSI target port shall wait two system deskew delays after negating the C/D, I/O, and MSG signals before releasing the REQ signal.
- 3) If the SCSI initiator port did not create an attention condition then the SCSI initiator port shall release all SCSI signals except ACK and ATN within two system deskew delays after detecting MSG, C/D, and I/O signals false. The ACK and ATN signals shall follow the timing specified in table 27.
- 4) If the SCSI initiator port creates an attention condition then the SCSI target port shall go to a MESSAGE OUT phase, receive all the message bytes, and cause an unexpected bus free by generating a BUS FREE phase (see 10.3).
- 5) If the SCSI target port detects the SEL signal being true, the SCSI target port shall release the BSY, MSG, C/D, and I/O signals within one QAS release delay.
- 6) After waiting at least one QAS arbitration delay from negating the SCSI MSG, C/D, and I/O signals in step 2), if there are no SCSI ID bits true the SCSI target port shall transition to the BUS FREE

phase.

- 7) After waiting at least one QAS arbitration delay from negating the MSG, C/D, and I/O signals in step 2), if there are any SCSI ID bits true the SCSI target port shall wait at least a second QAS arbitration delay. If the SEL signal is not true by the end of the second QAS arbitration delay the SCSI target port shall transition to the BUS FREE phase.

NOTE 27 - The release of MSG, C/D, and I/O may cause release glitches; Step 5) above ensures these glitches occur at a time when no connection is established on the bus so that they do not interfere with proper operation.

The procedure for a SCSI device with QAS enabled to obtain control of the SCSI bus via QAS is as follows:

- 1) The SCSI device shall first wait for MESSAGE IN phase to occur following a DT DATA phase with a single QAS REQUEST message. When the SCSI device detects the ACK signal being false for the QAS REQUEST message and the attention condition is cleared it shall begin the QAS phase.
- 2) The SCSI device shall wait a minimum of two system deskew delays after detection of the MSG, C/D, and I/O signals being false before driving any signal.
- 3) Following the delay in step 2), the SCSI device may arbitrate for the SCSI bus by asserting its own SCSI ID within one QAS assertion delay from detection of the MSG, C/D, and I/O signals being false. If arbitration fairness is enabled, the SCSI device shall not arbitrate until its fairness register is cleared.
- 4) After waiting at least one QAS arbitration delay, measured from the detection of the MSG, C/D, and I/O signals being negated, the SCSI device shall examine the DATA BUS.
  - A) If no higher priority SCSI ID bit is true on the DATA BUS and the fairness algorithm allowed the SCSI device to participate, then the SCSI device has won the arbitration and it shall assert the SEL signal.
  - B) If a higher priority SCSI ID bit is true on the DATA BUS (see table 31 for the SCSI ID arbitration priorities) or the fairness algorithm prevented the SCSI device from participating in QAS arbitration, then the SCSI device has lost the arbitration.
  - C) Any SCSI device other than the winner has lost the arbitration and shall release its SCSI ID bit after two system deskew delays and within one QAS release delay after detection of the SEL signal being true. A SCSI device that loses arbitration may return to step 1).
- 5) The SCSI device that wins arbitration shall wait at least one QAS arbitration delay after asserting the SEL signal before changing any signals.
- 6) After the QAS arbitration delay in step 4), SCSI devices with arbitration fairness enabled that are not arbitrating shall start sampling the DATA BUS to determine the SCSI devices that are attempting arbitration, the SCSI device that won, and the SCSI devices that lost. This sampling shall continue for one bus settle delay plus two system deskew delays. The SCSI devices shall update their fairness register with all device IDs that lost arbitration.

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. The DB(P\_CRCA) and DB(P1) are not valid during the QAS phase. During the QAS phase, DB(P\_CRCA), and DB(P1) may be released or asserted, but shall not be actively driven false.

## 10.5 SELECTION phase

### 10.5.1 Selection Overview

The SELECTION phase allows a SCSI initiator port to select a SCSI target port for the purpose of initiating some SCSI target port 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 a normal arbitration has both the BSY and SEL signals asserted and has delayed at least one bus clear delay plus one bus settle delay before ending the NORMAL ARBITRATION phase.



The SCSI device that won QAS has the SEL signal asserted and has delayed at least one QAS arbitration delay before ending the QAS phase.

The SCSI device that won the arbitration identifies itself as a SCSI initiator port by not asserting the I/O signal.

## **10.5.2 Selection using attention condition**

### **10.5.2.1 Starting the SELECTION phase when using attention condition**

The SCSI initiator port shall set the DATA BUS to a value that is the OR of its SCSI ID bit, the SCSI target port's SCSI ID bit, and the appropriate parity bit(s) (i.e., DB(P\_CRCA), and/or DB(P1)). The SCSI initiator port shall create an attention condition, indicating that a MESSAGE OUT phase is to follow the SELECTION phase.

If the arbitration was a normal arbitration then the SCSI initiator port shall wait at least two system deskew delays and release the BSY signal. The SCSI initiator port shall then wait at least one bus settle delay before attempting to detect an assertion of the BSY signal from the SCSI target port.

If QAS was used for arbitration then the SCSI initiator port shall wait at least one bus settle delay before attempting to detect an assertion of the BSY signal from the SCSI target port.

The SCSI target port shall detect 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 one bus settle delay. The selected SCSI target port may examine the DATA BUS in order to determine the SCSI ID of the selecting SCSI initiator port. The selected SCSI target port shall then assert the BSY signal within one selection abort time of its most recent detection of being selected; this is required for correct operation of the selection time-out procedure.

The SCSI target port shall not respond to a selection if a parity error is detected (see 11.2). Also, if more or less than two SCSI ID bits are on the DATA BUS, the SCSI target port shall not respond to selection.

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

### **10.5.2.2 Information unit transfers disabled**

If an information unit transfer agreement is not in effect for the connecting SCSI initiator port device the SCSI target port shall follow the phase sequences defined in 13.2.1.

### **10.5.2.3 Information unit transfers enabled**

If an information unit transfer agreement is in effect for the connecting SCSI initiator port the SCSI target port shall proceed to a MESSAGE OUT phase (see 13.3.2). On detecting the MESSAGE OUT phase the SCSI initiator port shall begin a PPR negotiation. On completion of the PPR negotiation the SCSI target port shall proceed to a BUS FREE phase. If the first message received by the SCSI target port during the MESSAGE OUT phase is not a TARGET RESET message or a PPR message the SCSI target port shall change to a MESSAGE IN phase and issue a MESSAGE REJECT message then originate WDTR negotiation (see 4.12.7.5) with the TRANSFER WIDTH EXPONENT field set to 00h.

### **10.5.2.4 Selection using attention condition time-out procedure**

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

- a) Optionally, the SCSI initiator port shall assert the RST signal (see 12.4);

- b) Optionally, the SCSI initiator port shall continue asserting the SEL and ATN signals and shall release DATA BUS, DB(P\_CRCA), and/or DB(P1). If the SCSI initiator port has not detected the BSY signal to be true after at least one selection abort time plus two system deskew delays, the SCSI initiator port shall release the SEL and ATN signals allowing the SCSI bus to go to the BUS FREE phase. When responding to selection, SCSI devices shall ensure that the selection was still valid within one selection abort time of their assertion of the BSY signal. Failure to comply with this requirement may result in an improper selection (e.g., two SCSI target ports connected to the same SCSI initiator port, wrong SCSI target port connected to a SCSI initiator port, or a SCSI target port connected to no SCSI initiator port).

### **10.5.3 Selection without using attention condition**

#### **10.5.3.1 Information unit transfers disabled or enabled**

The SCSI initiator port shall set the DATA BUS to a value that is the OR of its SCSI ID bit, the SCSI target port's SCSI ID bit, and the appropriate parity bit(s) (i.e., DB(P\_CRCA), and/or DB(P1)) and it shall clear the attention condition, indicating that a INFORMATION UNIT OUT phase is to follow the SELECTION phase.

If the arbitration was a normal arbitration then the SCSI initiator port shall wait at least two system deskew delays and release the BSY signal. The SCSI initiator port shall then wait at least one bus settle delay before attempting to detect an assertion of the BSY signal from the SCSI target port.

If QAS was used for arbitration then the SCSI initiator port shall wait at least one bus settle delay before attempting to detect an assertion of the BSY signal from the SCSI target port.

The SCSI target port shall detect 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 one bus settle delay. The selected SCSI target port may examine the DATA BUS in order to determine the SCSI ID of the selecting SCSI initiator port. The selected SCSI target port shall then assert the BSY signal within one selection abort time of its most recent detection of being selected; this is required for correct operation of the selection time-out procedure.

The SCSI target port shall not respond to a selection if a parity error is detected (see 11.2). Also, if more or less than two SCSI ID bits are on the DATA BUS, the SCSI target port shall not respond to selection.

The SCSI initiator port shall wait at least two system deskew delays after detecting that the BSY signal is true. The SCSI initiator port shall then release the SEL signal and may change the DATA BUS signals. The SCSI target port shall wait until the SEL signal is false before asserting the REQ signal to enter an information transfer phase.

If an information unit transfer agreement is in effect for the connecting SCSI initiator port the SCSI target port shall follow the phase sequences defined in 13.3.1.

If an information unit transfer agreement is not in effect for the connecting SCSI initiator port the SCSI target port shall follow the phase sequences defined in 13.2.1.

If a SCSI initiator port, when selecting without using an attention condition, detects an unexpected COMMAND phase, it shall set its transfer agreement to the default transfer agreement and set its negotiation required flag to true, create an attention condition, and on the corresponding MESSAGE OUT phase shall issue an ABORT TASK message. On the next selection of the SCSI target port that received the ABORT TASK message the SCSI initiator port shall do a selection using the attention condition and should negotiate to enable information unit transfers.

#### **10.5.3.2 Selection without using attention condition time-out procedure**

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

- a) Optionally, the SCSI initiator port shall assert the RST signal (see 12.4);
- b) Optionally, the SCSI initiator port shall continue asserting the SEL signal and shall release the DATA BUS, DB(P\_CRCA), or DB(P1). If the SCSI initiator port has not detected the BSY signal to be true after at least one selection abort time plus two system deskew delays, the SCSI initiator port shall release the SEL signal allowing the SCSI bus to go to the BUS FREE phase. When responding to selection, SCSI devices shall ensure that the selection was still valid within one selection abort time of their assertion of the BSY signal. Failure to comply with this requirement may result in an improper selection (e.g., two SCSI target ports connected to the same SCSI initiator port, wrong SCSI target port connected to a SCSI initiator port, or a SCSI target port connected to no SCSI initiator port).

## 10.6 RESELECTION phase

### 10.6.1 RESELECTION phase overview

The RESELECTION phase allows a SCSI target port to physically reconnect to a SCSI initiator port for the purpose of continuing some operation that was previously started by the SCSI initiator device but was suspended by the SCSI target port, (i.e., the SCSI target port physically disconnected by allowing a BUS FREE phase to occur or issued a QAS REQUEST message before the operation was complete). During the RESELECTION phase the I/O signal is asserted to distinguish this phase from the SELECTION phase.

### 10.6.2 Physical reconnection

The SCSI device that won a normal arbitration shall assert both the BSY and SEL signals and wait at least one bus clear delay plus one bus settle delay before ending the NORMAL ARBITRATION phase.

The SCSI device that won a QAS shall assert the SEL signal and has waited at least one QAS arbitration delay before ending the QAS phase.

The SCSI device that won the arbitration shall identify itself as a SCSI target port 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 SCSI initiator port's SCSI ID bit and the appropriate parity bit(s) (i.e., DB(P\_CRCA), and/or DB(P1)).

If the arbitration was a normal arbitration then SCSI target port shall wait at least two system deskew delays and release the BSY signal. The SCSI target port shall then wait at least one bus settle delay before attempting to detect an assertion of the BSY signal by the SCSI initiator port.

If QAS was used for arbitration then the SCSI target port shall wait at least one bus settle delay before attempting to detect an assertion of the BSY signal from the SCSI initiator port.

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

The SCSI initiator port shall not respond to a physical reconnection if a parity error is detected (see 11.2). Also, if more than or less than two SCSI ID bits are on the DATA BUS, the SCSI initiator port shall not respond to a physical reconnection.

After the SCSI target port detects the assertion of the BSY signal, it shall also assert the BSY signal and wait at least two system deskew delays and then release the SEL signal. The SCSI target port may then

change the I/O signal and the DATA BUS. After the physically reconnected SCSI initiator port detects the SEL signal is false, it shall release the BSY signal. The SCSI target port shall continue asserting the BSY signal until it relinquishes the SCSI bus.

NOTE 28 - When the SCSI target port 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 SCSI initiator port. 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 one bus settle delay. For more information on glitches see 7.2.4.1.

### 10.6.3 Physical reconnection time-out procedure

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

- a) Optionally, the SCSI target port shall assert the RST signal (see 12.4);
- b) Optionally, the SCSI target port shall continue asserting the SEL and I/O signals and shall release all DATA BUS, DB(P\_CRCA), and/or DB(P1) signals. If the SCSI target port has not detected the BSY signal to be true after at least one selection abort time plus two system deskew delays, the SCSI target port 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 physical reconnection was still valid within one selection abort time of their assertion of the BSY signal. Failure to comply with this requirement may result in an improper physical reconnection (e.g., two SCSI initiator ports connected to the same SCSI target port or the wrong SCSI initiator port connected to a SCSI target port).

## 10.7 Information transfer phases

### 10.7.1 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 subclause.

The C/D, I/O, and MSG signals are used to distinguish between the different information transfer phases (see table 42). The SCSI target port drives these three signals and therefore controls all changes from one phase to another. The SCSI initiator port requests a MESSAGE OUT phase creating an attention condition. The SCSI target port causes the BUS FREE phase by releasing the MSG, C/D, I/O, and BSY signals.

The information transfer phases use one or more REQ or ACK handshakes to control the information transfer. Each REQ/ACK handshake allows the transfer of 8- or 16-bits of information depending on the negotiated transfer width (see 4.12.4.5). During the information transfer phases the BSY signal shall remain true and the SEL signal shall remain false, except to indicate a paced transfer training pattern is going to occur (see 10.7.4.2). Additionally, during the information transfer phases, the SCSI target port shall continuously envelope the REQ or ACK handshakes with the C/D, I/O, and MSG signals in such a manner that these control signals are valid for one bus settle delay before the assertion of the REQ signal of the first handshake and remain valid until after the negation of the ACK signal at the end of the handshake of the last transfer of the phase.

The SCSI target port shall not transition into an information transfer phase unless the REQ and ACK signals are negated. The SCSI target port shall not transition from an information transfer phase into another information transfer phase unless the REQ and ACK signals are negated.

NOTE 29 - After the negation of the ACK signal of the last transfer of the phase, the SCSI target port 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 signal is asserted for the first byte of the new phase.

NOTE 30 - A phase is defined as ending when the C/D, I/O, or MSG signals change after the negation of the ACK signal. The time between the end of a phase and the assertion of the REQ signal beginning a new phase is undefined.

There are three methods of transferring data using information transfers:

- a) asynchronous transfers (see 10.7.2);
- b) synchronous transfers (see 10.7.3); and
- c) paced transfers (see 10.7.4).

Synchronous transfers shall only be used for negotiated transfer rates less than or equal to fast-80.

Paced transfers shall only be used for a negotiated transfer rate of fast-160 or fast-360.

**Table 42 - Information transfer phases**

Signal			Phase	Direction of transfer	Comment	
C/D	MSG	I/O				
0	0	0	ST DATA OUT	SCSI initiator port to SCSI target port	ST DATA phase	DATA phase
0	0	1	ST DATA IN	SCSI initiator port from SCSI target port		
0	1	0	DT DATA OUT	SCSI initiator port to SCSI target port	DT DATA phase	
0	1	1	DT DATA IN	SCSI initiator port from SCSI target port		
1	0	0	COMMAND	SCSI initiator port to SCSI target port		
1	0	1	STATUS	SCSI initiator port from SCSI target port		
1	1	0	MESSAGE OUT	SCSI initiator port to SCSI target port	MESSAGE phase	
1	1	1	MESSAGE IN	SCSI initiator port from SCSI target port		
Key: 0 = False; 1 = True						

### 10.7.2 Asynchronous transfer

The SCSI target port 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 SCSI target port to the SCSI initiator port. When

the I/O signal is false, information shall be transferred from the SCSI initiator port to the SCSI target port.

If the I/O signal is true (i.e., transfer to the SCSI initiator port), the SCSI target port shall first drive the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals to their values, delay at least one system deskew delay plus one cable skew, then assert the REQ signal. The DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals shall remain valid until the ACK signal is true at the SCSI target port. The SCSI initiator port shall read the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals after the REQ signal is true, then indicate its acceptance of the data by asserting the ACK signal. When the ACK signal becomes true at the SCSI target port, the SCSI target port may change or release the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals and shall negate the REQ signal. After the REQ signal is false the SCSI initiator port shall then negate the ACK signal. After the ACK signal is false the SCSI target port may continue the transfer by driving the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals and asserting the REQ signal, as described in this subclause.

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

### **10.7.3 Synchronous transfer**

#### **10.7.3.1 Synchronous transfer overview**

Synchronous transfer is optional and is only used in DATA phases. It shall be used in a DATA phase if a synchronous transfer agreement has been established (see 4.12). The transfer agreement specifies the REQ/ACK offset and the transfer period.

When synchronous transfers are being used data may be transferred using ST data transfers or, optionally, DT data transfers. DT data transfers shall only be used on 16 bit wide buses that transmit and receive data using LVD transceivers.

Implementors shall not use this clause for timing requirements. For timing requirements see 9.2.

#### **10.7.3.2 ST synchronous transfer**

When a ST data transfer agreement has been established the SCSI target port shall only use the ST DATA IN phase and ST DATA OUT phase for data transfers.

The REQ/ACK offset specifies the maximum number of REQ assertions that shall be sent by the SCSI target port in advance of the number of ACK assertions received from the SCSI initiator port. If the number of REQ assertions exceeds the number of ACK assertions by the REQ/ACK offset, the SCSI target port shall not assert the REQ signal until after the next ACK assertion is received. For successful completion of the ST DATA phase the number of ACK and REQ assertions shall be equal.

For the timing requirements of the negotiated transfer period see 9.2.

If the I/O signal is true (i.e., transfer to the SCSI initiator port), the SCSI target port shall first drive the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals to their values, wait at least one transmit setup time, then assert the REQ signal. The DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals shall be held valid for a minimum of one transmit hold time after the assertion of the REQ signal. The SCSI target port shall assert the REQ signal for a minimum of one transmit assertion period. The SCSI target port may then

negate the REQ signal and change or release the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals. The SCSI initiator port shall read the value on the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals within one receive hold time of the transition of the REQ signal to true. The SCSI initiator port shall then respond with an ACK assertion.

If the I/O signal is false (i.e., transfer to the SCSI target port), the SCSI initiator port, after detecting a REQ assertion, shall first drive the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals to their values, delay at least one transmit setup time, then assert the ACK signal. The SCSI initiator port shall hold the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals valid for at least one transmit hold time after the assertion of the ACK signal. The SCSI initiator port shall assert the ACK signal for a minimum of one transmit assertion period. The SCSI initiator port may then negate the ACK signal and may change or release the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals. The SCSI target port shall read the value of the DB(7-0,P\_CRCA) or DB(15-0,P\_CRCA,P1) signals within one receive hold time of the transition of the ACK signal to true.

### 10.7.3.3 DT synchronous transfer

#### 10.7.3.3.1 DT synchronous transfer overview

When a DT data transfer agreement has been established the SCSI target port shall only use the DT DATA IN phase and DT DATA OUT phase for data transfers.

During DT data transfers data shall be clocked on both the assertion and negation of the REQ and ACK signal lines. References to REQ or ACK transitions in this clause refer to either an assertion or a negation of the REQ or ACK signal.

The REQ/ACK offset specifies the maximum number of REQ transitions that shall be sent by the SCSI target port in advance of the number of ACK transitions received from the SCSI initiator port. If the number of REQ transitions exceeds the number of ACK transitions by the REQ/ACK offset, the SCSI target port shall not transition the REQ signal until after the next ACK transition is received. For successful completion of the DT DATA phase the number of ACK and REQ transitions shall be equal and both REQ and ACK shall be negated.

For the timing requirements of the negotiated transfer period see 9.2.

#### 10.7.3.3.2 Information unit transfer

When an information unit transfer agreement has been established (see 4.12.4.6.2);

- a) Information units shall be transferred on the DT DATA OUT phase and the DT DATA IN phase, and
- b) the information units' embedded iuCRC shall be used to detect information unit data errors.

If the I/O signal is true (i.e., transfer to the SCSI initiator port), to transfer SPI information units the SCSI target port:

- 1) Shall drive the DB(15-0) signals to their values;
- 2) shall wait at least one transmit setup time from DB(15-0) being driven with valid data;
- 3) shall transition the REQ signal;
- 4) shall hold the DB(15-0) signals valid for a minimum of one transmit hold time;
- 5) may change or release the DB(15-0) signals; and
- 6) shall not change the REQ signal for a minimum of one transmit assertion period.

If the I/O signal is true (i.e., transfer to the SCSI initiator port), to receive SPI information units the SCSI initiator port shall:

- 1) Read the value on the DB(15-0) signals within one receive hold time of the transition of the REQ

- signal; and
- 2) respond with an ACK transition.

If the I/O signal is false (i.e., transfer to the SCSI target port), to transfer SPI information units the SCSI initiator port:

- 1) Shall wait until detecting a REQ transition;
- 2) shall drive the DB(15-0) signals to their values;
- 3) shall delay at least one transmit setup time;
- 4) shall transition the ACK signal;
- 5) shall hold the DB(15-0) signals valid for at least one transmit hold time;
- 6) shall not change the ACK signal for a minimum of one transmit assertion period; and
- 7) may then change or release the DB(15-0) signals.

If the I/O signal is false (i.e., transfer to the SCSI target port), to receive SPI information units the SCSI target port:

- 1) Shall read the value of the DB(15-0) signals within one receive hold time of the transition of the ACK; and
- 2) the REQ/ACK offset shall be zero at the beginning and end of all SPI information units except for consecutive SPI data stream information units.

As a result of a SPI information unit always being an even number of transfers, the REQ and ACK signals are negated both before and after the transmission of the SPI information unit.

#### **10.7.3.3.3 DT DATA IN phase information unit transfer exception condition handling**

The SCSI initiator port shall not negate the ACK for the last byte of the last iuCRC in an information unit until the entire information unit has been verified and any required attention condition has been established.

If an I\_T\_L\_Q nexus has been established and the SCSI initiator port detects an iuCRC error in any information unit it receives, other than a SPI status information unit, then the SCSI initiator port shall create an attention condition on or before acknowledging the last iuCRC of the failed information unit. When the SCSI target port switches to a MESSAGE OUT phase the SCSI initiator port should send an INITIATOR DETECTED ERROR message (see 16.3.5) to the SCSI target port. This message notifies the SCSI target device that data in the information unit was invalid.

If a SCSI initiator port detects an iuCRC error in a SPI status information unit the SCSI initiator port shall create an attention condition on or before the last iuCRC of the information unit is acknowledged. If the SCSI target port detects an attention condition it shall switch to a MESSAGE OUT phase and the SCSI initiator port shall send an INITIATOR DETECTED ERROR message (see 16.3.5) or an ABORT TASK message to the SCSI target device. These messages notify the SCSI target device that the SPI status information unit was invalid.

If an I\_T\_L\_Q nexus has been established and the information unit that failed was not a SPI status information unit and the message received from the SCSI initiator port was an INITIATOR DETECTED ERROR message then the SCSI target port shall send a SPI L\_Q/SPI status information unit pair to the SCSI initiator port with a CHECK CONDITION status and a sense key set to ABORTED COMMAND and an additional sense code set to INITIATOR DETECTED ERROR MESSAGE RECEIVED for the task associated with the received INITIATOR DETECTED ERROR message.

If the information unit that failed was a SPI status information unit and the message received was an INITIATOR DETECTED ERROR message then the SCSI target port shall retry transferring the SPI L\_Q/SPI status information unit pair to the SCSI initiator port with the original status information.

If the information unit that failed was a SPI status information unit and the message received was an



ABORT TASK message then the SCSI target port shall cause a bus free by generating a BUS FREE phase.

If the SCSI initiator port is receiving a SPI L\_Q information unit and the SCSI initiator port detects an iuCRC error (i.e., the nexus identification fails) while in the DT DATA IN phase the SCSI initiator port shall create the attention condition on or before the acknowledgment of the last iuCRC. When the SCSI target port switches to a MESSAGE OUT phase the SCSI initiator port shall send an INITIATOR DETECTED ERROR message (see 16.3.5) to the SCSI target port. This message notifies the SCSI target device that the nexus identification failed. The SCSI target port shall then cause a bus free by generating a BUS FREE phase. The SCSI target device shall retry the task associated with the failed SPI L\_Q information unit.

If the SCSI initiator port receives a SPI L\_Q information unit with a type code that is not defined in table 50 that SCSI initiator port shall create an attention condition on or before the acknowledgment of the last iuCRC. When the SCSI target port switches to a MESSAGE OUT phase the SCSI initiator port shall send an ABORT TASK message (see 16.5.2) to the SCSI target port. The message notifies the SCSI target device that the SPI L\_Q information units type code was rejected. The SCSI target port shall then cause a bus free by generating a BUS FREE phase.

#### **10.7.3.3.4 DT DATA OUT phase information unit transfer exception condition handling**

The SCSI target port shall only respond to an iuCRC error after all the data in an information unit has been received.

If an I\_T\_L\_Q nexus has been established and the SCSI target port detects an iuCRC error in any multiple command SPI command information unit the SCSI target port shall, before transitioning a REQ signal to receive the next SPI L\_Q information unit, switch to a DT DATA IN phase and send a SPI L\_Q/SPI status information unit pair to the SCSI initiator port with a CHECK CONDITION status and a sense key set to ABORTED COMMAND and the additional sense code set to INFORMATION UNIT iuCRC ERROR DETECTED for the task associated with the iuCRC error.

If an I\_T\_L\_Q nexus has been established and the SCSI target port detects an iuCRC error in any SPI data information unit, SPI data stream information unit, or last command SPI command information unit it receives while in the DT DATA OUT phase the SCSI target port shall switch to a DT DATA IN phase and send a SPI L\_Q/SPI status information unit pair to the SCSI initiator port, before sending any other SPI L\_Q information unit. The status information shall indicate a CHECK CONDITION status and a sense key set to ABORTED COMMAND and the additional sense code set to INFORMATION UNIT iuCRC ERROR DETECTED for the task associated with the iuCRC error.

If the SCSI target port detects an iuCRC error on an iuCRC interval that is not at the end of a SPI information unit the SCSI target port shall not respond to the error until all the bytes of the SPI information unit in which the error occurred have been transferred. The SCSI target port may discard the transmitted information.

If the SCSI target port is receiving a SPI L\_Q information unit and the SCSI target port detects an iuCRC error (i.e., the nexus identification fails) the SCSI target port shall cause an unexpected bus free by generating a BUS FREE phase (see 10.3).

If a SCSI target port receives a SPI L\_Q information unit with a type code that is not defined in table 50 that SCSI target port shall transfer all the bytes indicated in the DATA LENGTH field and iuCRC INTERVAL field and shall discard the transmitted information for the information unit that follows the SPI L\_Q information unit. After transferring all the bytes the SCSI target port shall change to a DT DATA IN phase and transmit a SPI status information unit with a RSPVALID bit of one and the packetized failure code set to INVALID TYPE CODE RECEIVED IN SPI L\_Q INFORMATION UNIT.

If a SCSI target port receives a SPI L\_Q information unit with an illegal data length (see 14.3.2) the SCSI target port shall transfer all the bytes indicated by the data length and iuCRC interval and shall discard the transmitted information. After transferring all the bytes the SCSI target port shall change to a DT DATA IN

phase and transmit a SPI status information unit with a RSPVALID bit of one and the packetized failure code set to ILLEGAL REQUEST RECEIVED IN SPI L\_Q INFORMATION UNIT.

#### 10.7.3.3.5 Data group data field transfer

When the SCSI target port is transferring consecutive data groups, it shall not transition the REQ signal while the P\_CRCA signal is asserted for the current data group until the SCSI initiator port has acknowledged the entire previous data group.

NOTE 31 - The requirement in the preceding paragraph ensures the SCSI initiator port is not required to maintain more than one simultaneous pCRC calculation in different data groups.

If the I/O signal is true (i.e., transfer to the SCSI initiator port), to transfer the data field the SCSI target port:

- 1) Shall drive the DB(15-0) signals to their values and shall negate the P\_CRCA signal;
- 2) shall wait at least the longer of one pCRC transmit setup time from the negation of P\_CRCA or one transmit setup time from DB(15-0) being driven with valid data;
- 3) shall transition the REQ signal;
- 4) shall hold the DB(15-0) signals valid for a minimum of one transmit hold time and shall hold the P\_CRCA signal for a minimum of one pCRC transmit hold time;
- 5) may change or release the DB(15-0) and P\_CRCA signals; and
- 6) shall not change the REQ signal for at least one transmit assertion period if asserted or one transmit negation period if negated.

If the I/O signal is true (i.e., transfer to the SCSI initiator port), to receive the data field the SCSI initiator port shall:

- 1) Read the value on the DB(15-0) signals within one receive hold time of the transition of the REQ signal;
- 2) Read the value of the P\_CRCA signal within one pCRC receive hold time of the transition of the REQ signal; and
- 3) respond with an ACK transition.

If the I/O signal is false (i.e., transfer to the SCSI target port), to transfer the data field the SCSI initiator port:

- 1) Shall wait until detecting a REQ transition with P\_CRCA negated;
- 2) shall drive the DB(15-0) signals to their values;
- 3) shall delay at least one transmit setup time;
- 4) shall transition the ACK signal;
- 5) shall hold the DB(15-0) signals valid for at least one transmit hold time;
- 6) may then change or release the DB(15-0) signals; and
- 7) shall not change the ACK signal for at least one transmit assertion period if asserted or one transmit negation period if negated.

If the I/O signal is false (i.e., transfer to the SCSI target port), to receive the data field the SCSI target port:

- 1) Shall read the value of the DB(15-0) signals within one receive hold time of the transition of the ACK.

#### 10.7.3.3.6 Data Group Pad field and pCRC field transfer to SCSI initiator port

If the SCSI target port determines a pad field is required, has completed the data field transfer of the current data group, the I/O signal is true (i.e., transfer to the SCSI initiator port), and the REQ signal is asserted, the SCSI target port shall:

- 1) Wait at least one pCRC transmit hold time since the last REQ assertion to assert P\_CRCA;

- 2) wait at least one transmit hold time since the last REQ assertion to assert the DB(15-0) signals to their pad values;
- 3) wait at least the longer of one pCRC transmit setup time from the assertion of P\_CRCA or one transmit setup time from DB(15-0) being driven with valid pad data;
- 4) wait until the SCSI initiator port has responded with all ACK transitions for the previous data group;
- 5) wait at least one transmit REQ assertion period with P\_CRCA transitioning since the last REQ assertion;
- 6) negate the REQ signal without waiting for the ACK transition corresponding to the previous REQ transition unless the negotiated offset would be exceeded;
- 7) hold the DB(15-0) signals valid for a minimum of one transmit hold time and hold the REQ signal negated for a minimum of one transmit negation period;
- 8) drive the DB(15-0) signals to their pCRC values;
- 9) wait at least one transmit setup time;
- 10) assert the REQ signal without waiting for the ACK transition corresponding to the previous REQ transition unless the negotiated offset would be exceeded;
- 11) hold the DB(15-0) signals for a minimum of one transmit hold time and hold the REQ signal asserted for a minimum of one transmit assertion period;
- 12) drive the DB(15-0) signals to their pCRC values;
- 13) wait at least one transmit setup time;
- 14) negate the REQ signal without waiting for the ACK transition corresponding to the previous REQ transition unless the negotiated offset would be exceeded;
- 15) hold the DB(15-0) signals for a minimum of one transmit hold time and hold the P\_CRCA signal asserted for at least one pCRC transmit hold time; and
- 16) hold the REQ signal negated for at least one transmit REQ negation period with P\_CRCA transitioning since the last REQ negation.

NOTE 32 - The above requirements in steps 6), 10), and 14) to not wait for the ACK transition corresponding to the previous REQ transition were not present in the SPI-3 standard. For compatibility with old designs SCSI initiator ports should generate ACK transitions for all received REQ transitions.

If the SCSI target port determines that a pad field is not required, has completed the data field transfer of the current data group, the I/O signal is true (i.e., transfer to the SCSI initiator port), and the REQ signal is negated, the SCSI target port shall:

- 1) Wait at least one pCRC transmit hold time since the last REQ negation to assert P\_CRCA;
- 2) wait at least one transmit hold time since the last REQ negation to assert the DB(15-0) signals to their pCRC values;
- 3) wait at least the longer of one pCRC transmit setup time from the assertion of P\_CRCA or one transmit setup time from DB(15-0) being driven with valid pCRC data;
- 4) wait until the SCSI initiator port has responded with all ACK transitions for the previous data group;
- 5) wait at least one transmit REQ negation period with P\_CRCA transitioning since the last REQ negation;
- 6) assert the REQ signal;
- 7) hold the DB(15-0) signals for a minimum of one transmit hold time and hold the REQ signal asserted for a minimum of one transmit assertion period;
- 8) drive the DB(15-0) signals to their pCRC values;
- 9) wait at least one transmit setup time;
- 10) negate the REQ signal; and
- 11) hold the DB(15-0) signals for a minimum of one transmit hold time and hold the P\_CRCA signal asserted for a minimum of one pCRC transmit hold time.
- 12) hold the REQ signal negated for at least one transmit REQ negation period with P\_CRCA transitioning since the last REQ negation

After either of the above sequences described in this subclause is complete the SCSI target port has ended a data group transfer.

The SCSI initiator port shall read the value on the DB(15-0) signals within one receive hold time of the

transition of the REQ signal. The SCSI initiator port shall then respond with an ACK transition.

The SCSI initiator port shall continue to use the pad bytes, if any, for checking against the computed pCRC for the current data group. Upon receipt of the last byte of the pCRC field, the received pCRC and computed pCRC shall be compared. If they do match (i.e., no pCRC error). then the SCSI initiator port shall negate the ACK signal.

If received pCRC and computed pCRC do not match (i.e., a pCRC error is detected), or if an improperly formatted data group is transferred, then the SCSI initiator port shall create an attention condition or before the last transfer of the data group. When the SCSI target port switches to a MESSAGE OUT phase the SCSI initiator port should send an INITIATOR DETECTED ERROR message (see 16.3.5) to the SCSI target port. This message notifies the SCSI target device that data contained within the data group was invalid.

If the SCSI target port does not retry transferring the information transfer or it exhausts its retry limit the SCSI target port shall go into a STATUS phase and send a CHECK CONDITION status with a sense key set to ABORTED COMMAND and an additional sense code set to INITIATOR DETECTED ERROR MESSAGE RECEIVED for the task associated with the received INITIATOR DETECTED ERROR message.

#### **10.7.3.3.7 Data Group Pad field and pCRC field transfer to SCSI target port**

If the I/O signal is false (i.e., transfer to the SCSI target port), the SCSI initiator port determines the data field transfer is complete by detecting an assertion of the P\_CRCA signal. If the REQ signal is asserted (i.e., pad field required) the SCSI initiator port shall first transfer the two pad bytes, then the four pCRC bytes. If the REQ signal is negated (i.e., no pad field required) the SCSI initiator port shall transfer the four pCRC bytes.

Pad field data and pCRC field data are transferred using the same negotiated transfer period as the data field data.

The SCSI target port may continue to send REQs, up to the negotiated offset, for the next data group. The SCSI target port shall not transition REQ with P\_CRCA asserted until the SCSI initiator port has responded with all ACK transitions for the previous data group.

When the SCSI initiator port detects an assertion of the P\_CRCA signal and the REQ signal is asserted (i.e., pad field required) it shall then:

- 1) Transfer data bytes for all outstanding REQs received prior to the REQ that had the P\_CRCA signal asserted;
- 2) drive the DB(15-0) signals to their pad values;
- 3) delay at least one transmit setup time;
- 4) negate the ACK signal;
- 5) hold the DB(15-0) signals valid for a minimum of one transmit hold time and hold the ACK signal negated for a minimum of one transmit assertion period;
- 6) drive the DB(15-0) signals to their pCRC values;
- 7) delay at least one transmit setup time;
- 8) assert the ACK signal;
- 9) hold the DB(15-0) signals valid for a minimum of one transmit hold time and hold the ACK signal asserted for a minimum of one transmit assertion period;
- 10) drive the DB(15-0) signals to their pCRC values;
- 11) delay at least one transmit setup time;
- 12) negate the ACK signal; and
- 13) hold the DB(15-0) signals valid for a minimum of one transmit hold time and hold the ACK signal negated for a minimum of one transmit assertion period.

When the SCSI initiator port detects an assertion of the P\_CRCA signal and the REQ signal is negated

(i.e., no pad field required) it shall then:

- 1) Transfer data bytes for all outstanding REQs received prior to the REQ that had the P\_CRCA signal asserted;
- 2) drive the DB(15-0) signals to their pCRC values;
- 3) delay at least one transmit setup time;
- 4) assert the ACK signal;
- 5) hold the DB(15-0) signals valid for a minimum of one transmit hold time and hold the ACK signal asserted for a minimum of one transmit assertion period;
- 6) drive the DB(15-0) signals to their pCRC values;
- 7) delay at least one transmit setup time;
- 8) negate the ACK signal; and
- 9) hold the DB(15-0) signals valid for a minimum of one transmit hold time and hold the ACK signal negated for a minimum of one transmit assertion period.

After either of the above sequences described in this subclause is complete the SCSI target port has ended a data group transfer.

As a result of a data group always being an even number of transfers, the REQ and ACK signals are negated both before and after the transmission of the data group.

The SCSI target port shall read the value of the DB(15-0) signals within one receive hold time of the transition of the ACK signal.

The SCSI initiator port shall use the pad bytes, if any, in the generation of the transmitted pCRC. The SCSI target port shall then use those pad bytes, if any, for checking against the computed pCRC for the current data group. Upon receipt of the last byte of the pCRC field, the received pCRC and computed pCRC shall be compared.

If received pCRC and computed pCRC do not match (i.e., a pCRC error is detected), or if an improperly formatted data group is transferred, then the associated data group shall be considered invalid.

If the SCSI target port does not retry transferring the information transfer or it exhausts its retry limit the SCSI target device shall go into a STATUS phase and send a CHECK CONDITION status with a sense key set to ABORTED COMMAND and an additional sense code set to SCSI PARITY ERROR for the task associated with the pCRC error.

## **10.7.4 Paced transfer**

### **10.7.4.1 Paced transfer overview**

If a paced transfer agreement has been established it shall be used in DT DATA phase and information unit transfers shall be used. The transfer agreement also specifies the REQ/ACK offset and the transfer period (see 4.12).

When paced transfers are being used data shall be transferred using DT data transfers on 16-bit wide buses that transmit and receive data using LVD transceivers.

If driver precompensation is enabled at the SCSI device, that SCSI device shall apply driver precompensation to all the data, parity, REQ, and ACK signals.

During paced DT data transfers, if the phase of the P1 signal indicates data is valid (see 10.7.4.3) on REQ or ACK assertions, data shall be clocked by the originating SCSI device by both the assertion and negation of the REQ or ACK signal lines. The receiving SCSI device shall clock DT data on both the assertion and negation of the REQ or ACK signal line after having been processed by the receiving SCSI device. If the phase of the P1 signal indicates data is invalid on REQ or ACK assertions, data shall not be clocked by the originating SCSI device and shall be ignored by the receiving SCSI device. If driver precompensation is

enabled at the originating SCSI device, the originating SCSI device shall apply driver precompensation to all the data signals, the P\_CRCA signal, the P1 signal, and the REQ, and or ACK signal.

For paced DT DATA IN phases the REQ/ACK offset specifies the maximum number of data valid state REQ assertions (see 10.7.4.3) that shall be sent by the SCSI target port in advance of the number of ACK assertions received from the SCSI initiator port. If the number of data valid state REQ assertions exceeds the number of ACK assertions by the REQ/ACK offset, the SCSI target port shall change P1 to enable the data invalid state prior to the next assertion of REQ and shall not change P1 to enable a data valid state until after the next ACK assertion is received. For successful completion of a paced DT DATA IN phase the number of data valid state REQ assertions and ACK assertions shall be equal. Each assertion indicates a single 32-bit data transfer.

For paced DT DATA OUT phases the REQ/ACK offset specifies the maximum number of REQ assertions that shall be sent by the SCSI target port in advance of the number of data valid state ACK assertions received from the SCSI initiator port. If the number of REQ assertions exceeds the number of data valid state ACK assertions by the REQ/ACK offset, the SCSI target port shall not assert REQ until after the next data valid state ACK assertion is received. For successful completion of a paced DT DATA OUT phase the number of REQ assertions and data valid state ACK assertions shall be equal. Each assertion indicates a single 32-bit data transfer.

Implementors shall not use the following subclauses for timing requirements. For timing requirements see 9.2.

#### **10.7.4.2 Paced transfer training pattern**

##### **10.7.4.2.1 Training pattern overview**

After any PPR negotiation occurs that enables paced transfers, a training pattern shall be transferred at the start of the first DT data phase for each data transfer direction regardless of the negotiated value of the RTI bit (see 4.12.4.6.8).

If retain training information is disabled (see 4.12.4.6.8) a training pattern shall be transferred at the start of the first DT data phase for each data transfer direction after each physical connect and physical reconnect. The training pattern shall not be transferred again until after a physical disconnection occurs.

If the retain training information is enabled (see 4.12.4.6.8) a training pattern shall be transferred at the start of the first DT data phase for each data transfer direction after the retain training information is enabled. The SCSI device shall save training configuration values for each I\_T nexus that has negotiated to retain training information. The SCSI device shall use the saved training configuration values for all paced transfers. The SCSI target port may retrain an I\_T nexus if it determines the training configuration values are invalid, without having to renegotiate the retain training information protocol option.

NOTE 33 - The training configuration values are vendor specific.

If the retain training information is enabled and a port changes from a SCSI initiator port to a SCSI target port that SCSI target port shall retrain if the saved training configuration values were saved while the port was a SCSI initiator port.

The training pattern for a DATA IN phase shall conform to 10.7.4.2.2. The training pattern for a DATA OUT phase shall conform to 10.7.4.2.3. The receiving SCSI device shall use some or all elements of the training pattern to achieve deskewing (see 10.7.4.5). The transmitting SCSI device shall not make an intentional shift in relative timing between the data bus signals and the REQ or ACK signal during the DT data phase.

NOTE 34 - The requirement to not intentionally change relative timing does not include the effects of ISI, noise, or jitter.

The training pattern consists of three sections; A, B, and C. Each section contains a different pattern that

may be used to train circuits within a receiver.

#### 10.7.4.2.2 DT DATA IN phase training pattern

The SCSI target port shall indicate a training pattern is going to occur on a DT DATA IN phase by:

- 1) Releasing SEL for a minimum of two system deskew delays;
- 2) asserting the SEL signal a minimum of two system deskew delays; and
- 3) then asserting the REQ signal.

The SCSI target port shall begin the section A of its training pattern only after all the signal restrictions between information transfers phases listed in 10.12 or the signal restrictions between a RESELECTION phase and a DT DATA IN phase listed in 10.6.2 are met.

For fast-160, the SCSI target port shall transmit the training pattern as described by section A, section B, and section C in this subclause. For fast-320 the SCSI target port shall transmit the training pattern as described by section A, section B, and section C in this subclause except that the polarity of DB(0, 1, 4, 5, 9, 10, 13, 14, and P\_CRCA) shall be inverted during transmission of section A, section B, and section C (i.e., where it is specified that these signals shall be asserted, they shall be negated, and where it is specified that these signals shall be negated, they shall be asserted). These signals shall return to their normal polarity after completion of training pattern transmission.

Start of section A:

- 1) if precompensation is enabled then set the drivers to the strong driver state;
- 2) simultaneously assert REQ, P1, P\_CRCA, and DB(15-0) signals;
- 3) wait 200 ns (e.g., the equivalent of 32 transfer periods at fast-160);
- 4) simultaneously negate REQ, P1, P\_CRCA, and DB(15-0) signals;
- 5) wait 200 ns (e.g., the equivalent of 32 transfer periods at fast-160);
- 6) set precompensation to negotiated state;
- 7) negate SEL signal;
- 8) simultaneously assert and negate REQ, P1, P\_CRCA, and DB(15-0) signals at the negotiated transfer period for 800 ns (e.g., the equivalent of 128 transfer periods at fast-160);

Start of section B:

- 1) wait 1200 ns from the first assertion of REQ in step 2 of section A (e.g., the equivalent of 192 transfer periods at fast-160);
- 2) keep the P1, P\_CRCA, and DB(15-0) signals negated while continuing to assert and negate REQ at the negotiated transfer period for 50 ns (e.g., the equivalent of 8 transfer periods at fast-160);
- 3) keep the P1, P\_CRCA, DB(15-0), and REQ signals negated for an additional 50 ns;
- 4) simultaneously assert and negate P1, P\_CRCA, and DB(15-0) signals at twice the negotiated transfer period (i.e., simultaneously repeat a 1100b bit pattern on each signal) while asserting and negating REQ for 300 ns at the negotiated transfer period (e.g., the equivalent of 48 transfer periods at fast-160); and

Start of section C:

- 1) assert and negate REQ at the negotiated transfer period for 800 ns and at the same time assert and negate P1 at twice the negotiated transfer period while repeating a 0000010011111011b bit pattern on each of the P\_CRCA and DB(15-0) signals (e.g., the equivalent of 128 transfer periods at fast-160).

The SCSI initiator port shall begin its training pattern independent of the start of the SCSI target ports training pattern if it detects the SEL, MSG, and I/O true and C/D false on the first assertion of the REQ signal. The SCSI initiator port shall transmit the following training pattern:

- 1) assert ACK signal within 200 ns of the first REQ assertion;
- 2) if precompensation is enabled then set the drivers to the strong driver state;
- 3) wait 200 ns (e.g., the equivalent of 32 transfer periods at fast-160);
- 4) negate ACK signal;
- 5) wait 200 ns (e.g., the equivalent of 32 transfer periods at fast-160);

- 6) set precompensation to negotiated state; and
- 7) assert and negate ACK signal at the negotiated transfer period for 400 ns (e.g., the equivalent of 64 transfer periods at fast-160).

At the completion of its training pattern the SCSI target port continues asserting and negating the REQ signal at the negotiated transfer period (e.g., 6,25 ns at fast-160) and the P1 signal at twice the negotiated transfer period (e.g., 12,5 ns at fast-160). When the SCSI target port is ready to transfer data it shall reverse the phase of P1 (see 10.7.4.3).

#### 10.7.4.2.3 DT DATA OUT phase training pattern

The SCSI target port shall request a training pattern on a DT DATA OUT phase by asserting the SEL signal a minimum of two system deskew delays before asserting the REQ signal.

The SCSI target port shall begin its training pattern only after all the signal restrictions between information transfer phases listed in 10.12 or the signal restrictions between a SELECTION phase and a DT DATA OUT phase listed in 10.5.3 are met. The SCSI target port shall transmit the following training pattern:

- 1) if precompensation is enabled then set the drivers to the strong driver state;
- 2) simultaneously assert REQ and P\_CRCA signals;
- 3) wait 200 ns (e.g., the equivalent of 32 transfer periods at fast-160);
- 4) simultaneously negate REQ and P\_CRCA signals;
- 5) wait 200 ns (e.g., the equivalent of 32 transfer periods at fast-160);
- 6) set precompensation to negotiated state;
- 7) negate SEL signal;
- 8) simultaneously assert and negate REQ and P\_CRCA signals for 400 ns (e.g., the equivalent of 64 transfer periods at fast-160);
- 9) negate REQ and P\_CRCA for at least 100 ns (e.g., the equivalent of 16 transfer periods at fast-160); and
- 10) the SCSI target port shall begin asserting and negating REQ to indicate to the SCSI initiator port valid data may be sent. The number of REQ assertions shall not exceed the negotiated REQ/ACK offset.

The SCSI initiator port shall begin the section A of its training pattern independent of the start of the SCSI target ports training pattern if it detects the SEL and MSG true, and C/D and I/O false on the first assertion of the REQ signal. The SCSI initiator port shall transmit the training pattern described by section A, section B, and section C in this subclause:

For fast-320 the SCSI initiator port shall transmit the training pattern described by section A, section B, and section C in this subclause except that the polarity of DB(0, 1, 4, 5, 9, 10, 13, and 14) shall be inverted during transmission of section A, section B, and section C (i.e., where it is specified that these signals shall be asserted, they shall be negated, and where it is specified that these signals shall be negated, they shall be asserted). These signals shall return to their normal polarity after completion of training pattern transmission.

Start of section A:

- 1) if precompensation is enabled then set the drivers to the strong driver state;
- 2) simultaneously assert ACK, P1, and DB(15-0) signals within 200 ns of the first REQ assertion (e.g., the equivalent of 32 transfer periods at fast-160);
- 3) wait 200 ns (e.g., the equivalent of 32 transfer periods at fast-160);
- 4) simultaneously negate ACK, P1, and DB(15-0) signals;
- 5) wait 200 ns;
- 6) set precompensation to negotiated state;
- 7) simultaneously assert and negate ACK, P1, and DB(15-0) signals at the negotiated transfer period for 800 ns (e.g., the equivalent of 128 transfer periods at fast-160);

Start of section B:

- 1) wait the 1200 ns from the first assertion of ACK in step 2 of section A (e.g., equivalent of 192



- transfer periods at fast-160);
- 2) keep the P1, and DB(15-0) signals negated while continuing to assert and negate ACK at the negotiated transfer period for 50 ns (e.g., the equivalent of 8 transfer periods at fast-160);
  - 3) keep the P1, DB(15-0), and ACK signals negated for an additional 50 ns;
  - 4) simultaneously assert and negate P1 and DB(15-0) signals at twice the negotiated transfer period (i.e., simultaneously repeat a 1100b bit pattern on each signal) while asserting and negating ACK 300 ns (e.g., the equivalent of 48 transfer periods at fast-160); and

Start of section C:

- 1) assert and negate ACK at the negotiated transfer period for 800 ns and at the same time assert and negate P1 at twice the negotiated transfer period while repeating a 0000010011111011b bit pattern on each of the DB(15-0) signals (e.g., the equivalent of 128 transfer periods at fast-160).

At the completion of its training pattern the SCSI initiator port continues asserting and negating the ACK signal at the negotiated transfer period (e.g., 6,25 ns transfer period at fast-160) and the P1 signal at twice the negotiated transfer period (e.g., 12,5 ns transfer period at fast-160). When the SCSI initiator port is ready to transfer data and the REQ/ACK offset value is not zero it shall reverse the phase of P1 (see 10.7.4.3).

#### **10.7.4.3 P1 data valid/invalid state transitions**

##### **10.7.4.3.1 P1 data valid/invalid state transitions overview**


The transmitting SCSI device port shall indicate the start of a data valid state by reversing the phase of the P1 signal coincident with a REQ or ACK assertion. This is accomplished by withholding the next transition of P1 at the start of the first two transfer periods of valid data. Beginning with the third valid data word, P1 shall be toggled every two transfer periods, coincident with a REQ or ACK assertion. The minimum duration of the data valid state is two transfer periods, and the data valid state shall consist of an even number of transfer periods.

Anytime the sending SCSI device port pauses the sending of data, it shall reverse the phase of P1 by withholding the next transition of P1 at the start of the first two transfer periods that have invalid data. Beginning with the third transfer period with invalid data, P1 shall be toggled every two transfer periods until valid data is sent. The data invalid state shall have at least one transition of P1 before changing states. The minimum data invalid time is four transfer periods. This ensures a maximum run length of three cycles for P1. The data invalid state shall last an even number of transfer periods.

From the data invalid state, the sending SCSI device port may resume sending data by reversing the phase of P1 again.

P1 has the same transmit setup and hold time requirements as data and shall always be detected by the receiving device on the assertion edge of the delayed clocked REQ or ACK signal.

See figure 73 for examples of how the P1 signal is used to determine when the REQ or ACK transition clocks valid data.



**Figure 73 - Usage of P1 to establish data valid and data invalid states**

#### 10.7.4.3.2 Starting pacing transfers at end of training pattern

See 10.7.4.2 for the description of starting a data valid state after a training pattern.

#### 10.7.4.3.3 Starting pacing transfers with no training pattern

Before starting the DT DATA IN phase the SCSI target port shall wait at least two system deskew delays after the SEL signal is negated before the first assertion of the REQ signal.

The DT DATA IN phase without training starts on the first assertion of REQ if the SEL is not asserted.

The SCSI target port shall begin pacing transfers only after meeting all the following:

- a) signal restrictions between information transfer phases listed in 10.12;
- b) the signal restrictions between a RESELECTION phase and a DT DATA IN phase listed in 10.6.2; and
- c) the signal restrictions between a SELECTION phase and a DT DATA OUT phase listed in 10.5.3.

The SCSI target port shall begin pacing transfers by:

- 1) simultaneously with the assertion of REQ the SCSI target port shall begin asserting and negating P1 at twice the negotiated transfer period (e.g., 12,5 ns for fast-160);
- 2) SCSI target port shall assert and negate P1 for at least 100 ns (e.g., the equivalent of 16 transfer periods for fast-160); and
- 3) the SCSI target port may establish a data valid state as described in 10.7.4.3.1.

The DT DATA OUT phase without training starts on the first assertion of REQ if the SEL is not asserted.

The SCSI target port shall begin pacing transfers only after meeting all the following:

- a) signal restrictions between information transfer phases listed in 10.12;
- b) the signal restrictions between a RESELECTION phase and a DT DATA IN phase listed in 10.6.2; and
- c) the signal restrictions between a SELECTION phase and a DT DATA OUT phase listed in 10.5.3.

The SCSI initiator port shall begin pacing transfers by:

- 1) simultaneously with the assertion of ACK the SCSI initiator port shall begin asserting and negating P1 at twice the negotiated transfer period (e.g., 12,5 ns for fast-160);
- 2) SCSI initiator port shall assert and negate P1 for at least 100 ns (e.g., the equivalent of 16 transfer periods for fast-160); and
- 3) the SCSI initiator port may establish a data valid state as described in 10.7.4.3.1.

#### 10.7.4.3.4 Ending pacing transfers

After transmitting the last data word of a DT DATA IN phase the SCSI target port shall end pacing by waiting for all REQs to be responded to by ACKs then negate the REQ and P1 signals. After the SCSI target port stops asserting and negating REQ it shall not assert REQ again until the requirements in 10.12 are met.

After transmitting the last data word of a DT DATA OUT phase the SCSI initiator port shall:

- a) continue asserting and negating the ACK and P1 signals until it detects a change to the C/D, I/O, or MSG signals; and
- b) negate the ACK and P1 signals within 200 ns of detecting a change to the C/D, I/O, or MSG signals.

When the SCSI target port changes from a DT DATA OUT phase to any other phase it shall wait at least a bus settle delay plus a data release delay before asserting REQ and shall ignore any ACK transitions for at least a bus settle delay plus a data release delay after transitioning the C/D, I/O, or MSG signals.

#### 10.7.4.4 Paced information unit transfer

Information units shall be transferred on the DT DATA OUT phase and the DT DATA IN phase, and the information units' embedded iuCRC shall be used to detect information unit data errors.

If the I/O signal is true (i.e., transfer to the SCSI initiator port) and the phase of the P1 signal indicates data is valid, to transfer SPI information units the SCSI target port:

- 1) shall drive the DB(15-0) signals to their values simultaneous with the next REQ signal assertion;
- 2) shall hold the DB(15-0) signals valid for a minimum of one transmit hold time;
- 3) shall drive the DB(15-0) signals to their values simultaneous with the next REQ signal negation; and
- 4) shall hold the DB(15-0) signals valid for a minimum of one transmit hold time.

If the I/O signal is true (i.e., transfer to the SCSI initiator port), to receive SPI information units the SCSI initiator port shall:

- 1) Read the value on the DB(15-0) signals within one receive hold time of the transition of the REQ signal; and
- 2) respond with an ACK signal assertion after each REQ assertion/negation pair.

If the I/O signal is false (i.e., transfer to the SCSI target port) and the phase of the P1 signal indicates data is valid, to transfer SPI information units the SCSI initiator port:

- 1) Shall wait until detecting a REQ assertion;
- 2) shall drive the DB(15-0) signals to their values simultaneous with the next ACK signal assertion;
- 3) shall hold the DB(15-0) signals valid for a minimum of one transmit hold time;
- 4) shall drive the DB(15-0) signals to their values simultaneous with the next ACK signal negation;  
and
- 5) shall hold the DB(15-0) signals valid for a minimum of one transmit hold time.

If the I/O signal is false (i.e., transfer to the SCSI target port), to receive SPI information units the SCSI target port:

- 1) Shall read the value of the DB(15-0) signals within one receive hold time of the transition of the ACK signal.

If write flow control is enabled and the current SPI data stream information unit is the last SPI data stream information unit of the stream then:

- 1) The SCSI target port shall assert the P\_CRCA signal a minimum of a flow control transmit setup time before the end of the last information unit and shall keep the P\_CRCA signal asserted for a flow control transmit hold time;
- 2) The SCSI target port shall not assert the P\_CRCA signal until a minimum of a flow control hold time after the end of the previous information unit; and
- 3) The SCSI target port shall negate the P\_CRCA signal a minimum of a flow control transmit setup time before the start of the next information unit.

NOTE 35 - The earlier in a SPI data stream information unit that the SCSI target port asserts the P\_CRCA signal, the better the SCSI initiator port may manage data pre-fetch.

As a result of a SPI information unit always being an even number of transfers, the REQ and ACK signals are negated both before and after the transmission of the SPI information unit.

Paced information unit transfers exception handling conditions are defined in 10.7.3.3.3, and 10.7.3.3.4.

#### 10.7.4.5 Deskewing

The deskewing technique used in the receiving SCSI device is vendor specific. Any technique that works with the specified training pattern and complies with the specified receive skew compensation timing requirement is allowed. Deskewing shall only be enabled for paced transfers.

#### 10.7.5 Wide transfer

Wide data transfers shall be used for DT DATA phases. Wide data transfer may be used in the ST DATA phase only if a wide transfer agreement is in effect (see 4.12).

All SCSI devices shall support narrow data transfers.

During narrow transfers, all information shall be transferred in bytes across the DB(7-0) and DB(P\_CRCA) signals on the SCSI bus. At the receiving SCSI device the DB(15-8) (if present) and DB(P1) (if present) signals are undefined.

During wide transfers, the first and second information bytes for each DATA phase shall be transferred across the DB(7-0) and DB(15-8) signals, respectively, on the SCSI bus. Subsequent pairs of information bytes are likewise transferred in parallel across the SCSI bus (see table 43).

The IGNORE WIDE RESIDUE message may be used to indicate that the last byte of a data field or the last data byte of information unit is undefined.

**Table 43 - Wide SCSI byte order**

Transfer number	SCSI Bus		Data transfer width
	15.....8	7.....0	
1	N/A	W	8-bit
2	N/A	X	
3	N/A	Y	
4	N/A	Z	
1	X	W	16-bit
2	Z	Y	

When transferring consecutive bytes W, X, Y, and Z across the buses, they are transferred as shown in this table.

This table does not necessarily represent how these bytes are stored in memory.

If the last information byte transferred does not fall on the DB(15-8) signals for a 16-bit wide transfer then the values of the remaining higher-numbered bits are undefined. However, when using parity protection the DB(P1) signal for this undefined byte shall be valid for whatever data is placed on the bus.

## 10.8 COMMAND phase

### 10.8.1 COMMAND phase description

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

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

A QAS-capable SCSI initiator port shall wait a minimum of a QAS non-DATA phase REQ(ACK) period to assert ACK after detecting the assertion of REQ.

A QAS-capable SCSI initiator port shall assert ACK for a minimum of a QAS non-DATA phase REQ(ACK) period and shall keep the command data valid until the negation of ACK.

### 10.8.2 COMMAND phase exception condition handling

If the SCSI target port detects one or more parity errors on the command bytes received, it may retry the command by switching to the MESSAGE IN phase and sending a RESTORE POINTERS message. The SCSI target port shall then switch to the COMMAND phase to receive the original command.

If the SCSI target port does not retry the COMMAND phase or it exhausts its retry limit it shall return CHECK CONDITION status and set the sense key to ABORTED COMMAND and the additional sense code to SCSI PARITY ERROR.

## 10.9 DATA phase

### 10.9.1 DATA phase overview

DATA phase is a term that encompasses both the ST DATA phases and the DT DATA phases. ST DATA phase is a term that encompasses both the ST DATA IN phase and ST DATA OUT phase. DT DATA phase is a term that encompasses both the DT DATA IN phase, and the DT DATA OUT phase.

### 10.9.2 DT DATA IN phase

The DT DATA IN phase allows the SCSI target device to request that data be sent to the SCSI initiator device from the SCSI target device using DT data transfers.

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

### 10.9.3 DT DATA OUT phase

The DT DATA OUT phase allows the SCSI target device to request that data be sent from the SCSI initiator device to the SCSI target device using DT data transfers.

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

### 10.9.4 ST DATA IN phase

The ST DATA IN phase allows the SCSI target device to request that data be sent to the SCSI initiator device from the SCSI target device using ST data transfers.

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

### 10.9.5 ST DATA OUT phase

The ST DATA OUT phase allows the SCSI target device to request that data be sent from the SCSI initiator device to the SCSI target device using ST data transfers.

The SCSI target port shall negate the C/D, I/O, and MSG signals during the REQ/ACK handshakes of this phase.

## 10.10 STATUS phase

### 10.10.1 STATUS phase description

The STATUS phase allows the SCSI target device to request that a status byte be sent from the SCSI target device to the SCSI initiator device.

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

A QAS-capable SCSI initiator port shall wait a minimum of one QAS non-DATA phase REQ(ACK) period to assert ACK after detecting the assertion of REQ.

A QAS-capable SCSI initiator port shall assert ACK for a minimum of one QAS non-DATA phase REQ(ACK) period.

### 10.10.2 STATUS phase exception condition handling

If the SCSI initiator port detects a parity error on the status byte the SCSI initiator port shall create an attention condition. When the SCSI target port switches to a MESSAGE OUT phase the SCSI initiator port should send an INITIATOR DETECTED ERROR message (see 16.3.5) to the SCSI target port. This message notifies the SCSI target device that the status byte was invalid.

## 10.11 MESSAGE phase

### 10.11.1 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.

### 10.11.2 MESSAGE IN phase

The MESSAGE IN phase allows the SCSI target port to request that messages be sent to the SCSI initiator port from the SCSI target port.

The SCSI target port shall assert the C/D, I/O, and MSG signals during the REQ/ACK handshakes of this phase.

A QAS-capable SCSI initiator port shall wait a minimum of one QAS non-DATA phase REQ(ACK) period to assert ACK after detecting the assertion of REQ.

A QAS-capable SCSI initiator port shall assert ACK for a minimum of one QAS non-DATA phase REQ(ACK) period.

### 10.11.3 MESSAGE IN phase exception condition handling

If the SCSI initiator port detects a parity error on any message byte it receives the SCSI initiator port shall create an attention condition. When the SCSI target port switches to a MESSAGE OUT phase the SCSI initiator port shall send a MESSAGE PARITY ERROR message (see 16.3.7) to the SCSI target port. This message notifies the SCSI target device that the message in byte was invalid.

### 10.11.4 MESSAGE OUT phase

The MESSAGE OUT phase allows the SCSI target port to request that messages be sent from the SCSI initiator port to the SCSI target port. The SCSI target port invokes this phase in response to the attention condition created by the SCSI initiator port (see 12.2).

The SCSI target port shall assert the C/D and MSG signals and negate the I/O signal during the REQ/ACK handshakes of this phase. The SCSI target port shall handshake bytes in this phase until the attention condition is cleared, except when rejecting a message.

A QAS-capable SCSI initiator port shall wait a minimum of one QAS non-DATA phase REQ(ACK) period to assert ACK after detecting the assertion of REQ.

A QAS-capable SCSI initiator port shall assert ACK for a minimum of one QAS non-DATA phase REQ(ACK) period and shall keep the message data valid until the negation of ACK.

If the SCSI target port receives all of the message bytes successfully (i.e., no parity errors), it shall indicate that no retry is being attempted by changing to any information transfer phase other than the MESSAGE OUT phase and transferring at least one byte. The SCSI target port may also indicate it has successfully

received the message bytes by changing to the BUS FREE phase (e.g., after receiving ABORT TASK SET, LOGICAL UNIT RESET, or TARGET RESET messages).

### 10.11.5 MESSAGE OUT phase exception condition handling

If the SCSI target port detects one or more parity errors on the message bytes received, it may request a retry of the messages by asserting the REQ signal after detecting that the attention condition has been cleared and prior to changing to any other phase. The SCSI initiator port, upon detecting this condition, shall resend all of the previous message bytes in the same order as previously sent during this phase. When resending more than one message byte, the SCSI initiator port shall re-establish the attention condition as described in 12.2.

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

## 10.12 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 and ACK signals shall not change.
- b) If paced transfers are disabled the SEL signal shall not change.
- c) The REQ signal shall not change until it is asserted to qualify the start of a new phase.
- d) The C/D, I/O, MSG, DATA BUS, DB(P\_CRCA), and DB(P1) signals may change.
- e) If paced transfers are enabled the C/D, I/O, and MSG signals shall not change for at least one system deskew delay from the negation of the last REQ of a paced transfer data phase.
- f) If paced transfers are enabled the SEL signal may change.
- g) When switching the DATA BUS or DB(P\_CRCA) signal direction from out (i.e., SCSI initiator port driving) to in (i.e., SCSI target port driving), the SCSI target port shall delay driving the DATA BUS, DB(P\_CRCA), and/or DB(P1) by at least one data release delay plus one bus settle delay after asserting the I/O signal and the SCSI initiator port shall release the DATA BUS, DB(P\_CRCA), and/or DB(P1) no later than one data release delay after the transition of the I/O signal to true. When switching the DATA BUS, DB(P\_CRCA), and/or DB(P1) direction from in (i.e., SCSI target port driving) to out (i.e., SCSI initiator port driving), the SCSI target port shall release the DATA BUS, DB(P\_CRCA), and/or DB(P1) no later than one system deskew delay after negating the I/O signal. The SCSI initiator port shall drive the DATA BUS, DB(P\_CRCA), and/or DB(P1) no sooner than one system deskew delay after the detection of the negation of the I/O signal.
- h) The DB(P\_CRCA) signal direction may switch direction while the DATA BUS and/or DB(P1) does not (e.g., changing from COMMAND phase to DT DATA OUT phase). When switching the DB(P\_CRCA) signal direction from out (i.e., SCSI initiator port driving) to in (i.e., SCSI target port driving), the SCSI target port shall delay driving the DB(P\_CRCA) by at least one data release delay plus one bus settle delay after negating the C/D signal and the SCSI initiator port shall release the DB(P\_CRCA) signal no later than one data release delay after the transition of the C/D signal to false. When switching the DB(P\_CRCA) signal direction from in (i.e., SCSI target port driving) to out (i.e., SCSI initiator port driving), the SCSI target port shall release the DB(P\_CRCA) signal no later than one system deskew delay after asserting the C/D signal. The SCSI initiator port shall negate the DB(P\_CRCA) signal no sooner than one system deskew delay after the detection of the assertion of the C/D signal.



- i) The ATN and RST signals may change as defined under the descriptions for the attention condition (see 12.2) and bus reset condition (see 12.3).

## 11 DATA BUS protection

### 11.1 DATA BUS protection overview

The DB(P\_CRCA) signal and the DB(P1) signal are used to generate parity or control the transfer of pCRC information on the DATA BUS.

### 11.2 ST DATA BUS protection using parity

For ARBITRATION phase the DB(P\_CRCA) and DB(P1) signals shall not be checked for parity errors.

For SELECTION and RESELECTION phases valid parity is determined by rules in table 44.

**Table 44 - Parity checking rules for SELECTION and RESELECTION phases**

Action	Condition
<b>Check for odd parity on:</b>	<b>If at least one bit is active on:</b>
DB(7-0,P_CRCA)	DB(15-0,P_CRCA,P1)
DB(15-8,P1)	DB(15-8,P1)

NOTE 36 - These rules are necessary to permit interoperation of SCSI devices with different DATA BUS widths. For example, if an 8-bit SCSI device selects a 16-bit SCSI device, the 16-bit SCSI device observes invalid parity on the upper 8 bits of the DATA BUS.

For COMMAND, MESSAGE, and STATUS phases the DB(P\_CRCA) signal shall indicate odd parity for DB(7-0). The DB(P1) signal shall not be checked.

For ST DATA phases the DB(P\_CRCA) signal shall indicate odd parity for DB(7-0). If narrow transfers are enabled the DB(P1) signal shall not be checked. If wide transfers are enabled the DB(P1) signal shall indicate odd parity for DB(15-8). If wide transfers are enabled and the last information byte transferred does not fall on the DB(15-8) signals the DB(P1) signal shall be valid for whatever data is placed on the bus.

Parity protection is not enabled during DT DATA phases.

### 11.3 DT DATA BUS protection using CRC

#### 11.3.1 DT DATA BUS protection using CRC overview

When pCRC protection or iuCRC protection are enabled the error detecting code is a 32-bit (i.e., four byte) Cyclic Redundancy Check (CRC), referred to as CRC-32. It is also used by several other device I/O standards. Four CRC bytes are transferred with data to increase the reliability of data transfers

#### 11.3.2 Error detection capabilities

The CRC detects all single bit errors, any two bits in error, or any combination of errors within a single 32-bit range.

### 11.3.3 Order of bytes in the CRC field

Figure 74 shows how transmitted data is used to calculate the CRC and how the CRC information is then transmitted.

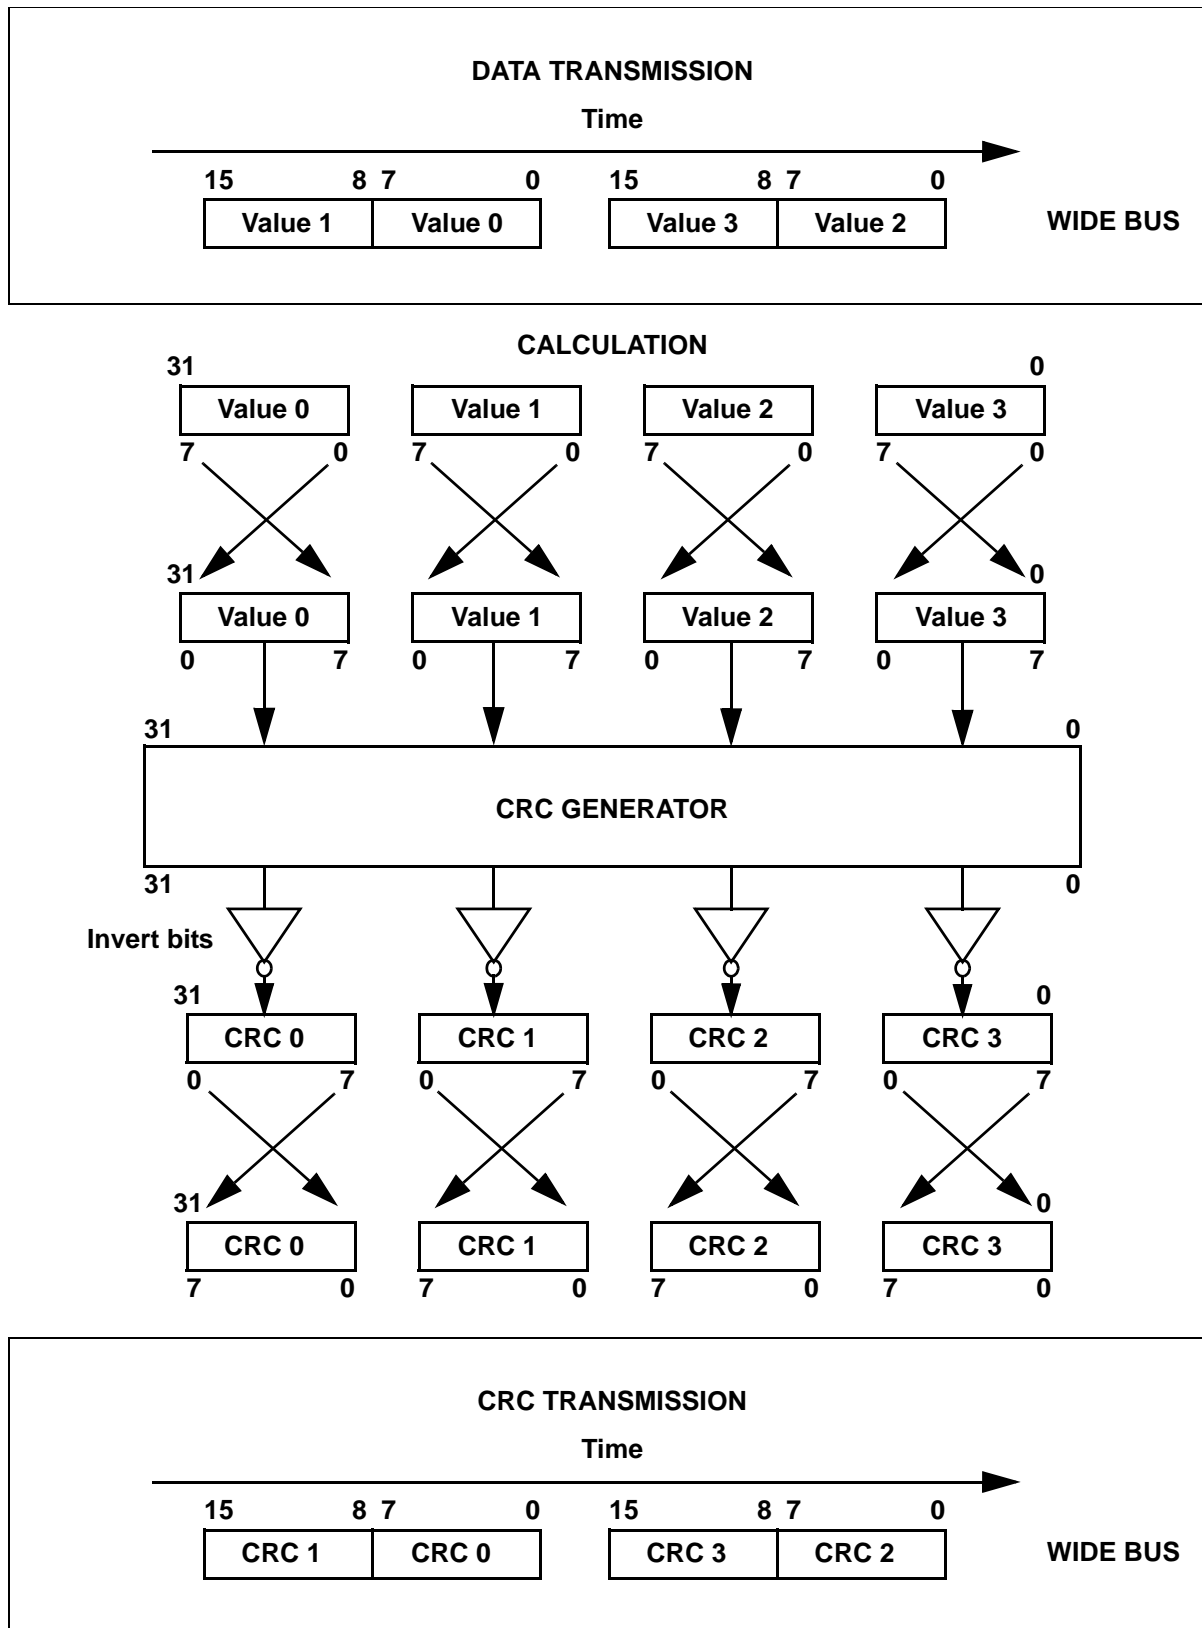


Figure 74 - CRC generation and transmission

### 11.3.4 CRC generation and checking

The 32-bit generator polynomial used is:

$$x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$

This equals 104C11DB7h.

The remainder is generated by dividing the bytes in the data and pad fields by the generator polynomial, modulo two. The remainder is generated 32 bits at a time.

The remainder is initialized to all ones (i.e., FFFF FFFFh). This is the seed value. It is reloaded at the beginning of each DT DATA phase and after each CRC is generated/checked.

The data transferred on the asserting edge of a REQ or an ACK is bit-reversed and becomes the most significant 16 bits of the CRC generator's input. The data transferred on the negating edge of a REQ or an ACK is bit-reversed and becomes the least significant 16 bits of the CRC generator's input.

The most significant 16 bits of the CRC generator's output are bit reversed and the one's complement of this result forms the portion of the CRC field that is transferred on the asserting edge of the CRC REQ or ACK. The least significant 16 bits of the CRC generator's output are bit reversed and the one's complement of this result forms the portion of the CRC field that is transferred on the negating edge of the CRC REQ or ACK.

A unique remainder is generated by an error free data group. The unique remainder polynomial of an error free group is:

$$x^{31} + x^{30} + x^{26} + x^{25} + x^{24} + x^{18} + x^{15} + x^{14} + x^{12} + x^{11} + x^{10} + x^8 + x^6 + x^5 + x^4 + x^3 + x + 1$$

This equals C704DD7Bh.

### 11.3.5 Test cases

Using the polynomial described in 11.3.4, the CRC calculated for a 32-byte transfer of all 00h is: 55ADh, 190Ah.

Using the polynomial described in 11.3.4, the CRC calculated for a 32-byte transfer of all FFh: AB0Bh, FF6Ch.

Using the polynomial described in 11.3.4, the CRC calculated for a 32-byte transfer of an incrementing pattern from 00h to 1Fh is: 7E8Ah, 9126h.

## 12 SCSI bus conditions

### 12.1 SCSI bus conditions overview

The SCSI bus has asynchronous conditions that 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 powering on, SCSI target devices should be able to respond with appropriate status and sense data to the TEST UNIT READY, INQUIRY, and REQUEST SENSE commands.

### 12.2 Attention condition

The attention condition allows a SCSI initiator port to inform a SCSI target port that the SCSI initiator port has a message ready. The SCSI target port shall honor all valid attention conditions by performing a MESSAGE OUT phase.

The SCSI initiator port may create an attention condition during the SELECTION phase and during all information transfer phases.

To create an attention condition during the SELECTION phase following normal arbitration, the SCSI initiator port shall assert the ATN signal at least two system deskew delays before releasing the BSY signal.

To create an attention condition during the SELECTION phase following a QAS, the SCSI initiator port shall assert the ATN signal at least two system deskew delays before asserting the SCSI target port's ID on the bus.

To create an attention condition during an information transfer phase, the SCSI initiator port shall assert the ATN signal at least an ATN transmit setup time before the specified transition of the ACK signal. To re-establish an attention condition during a multi-byte MESSAGE OUT retry, the SCSI initiator port shall assert the ATN signal two system deskew delays before asserting the ACK signal on the first message byte. To clear an attention condition during an information transfer phase, the SCSI initiator port shall negate the ATN signal at least two system deskew delays before asserting the ACK signal. The SCSI initiator port shall not negate the ATN signal while the ACK signal is asserted during a MESSAGE OUT phase.

The SCSI initiator port shall create the attention condition on or before the last information transfer in a bus phase or information unit, for the attention condition to be honored before transition to a new bus phase or information unit. If the SCSI initiator port does not meet the attention condition setup time, the SCSI target port may not honor the attention condition until a later bus phase or information unit, possibly resulting in an unexpected action. The SCSI initiator port shall keep the ATN signal asserted until the SCSI target port responds to the attention condition.

Once the SCSI target port has responded to the attention condition by going to MESSAGE OUT phase, the SCSI initiator port shall keep the attention condition set if more than one message byte is to be transferred. The SCSI initiator port shall clear the attention condition on the last message byte to be sent. The SCSI initiator port shall clear the attention condition while transferring the last byte of the messages indicated with a Yes in tables 60, 69, and 74. If the SCSI target port detects that the SCSI initiator port failed to meet this requirement, then the SCSI target port shall go to BUS FREE phase (see 10.2).

A SCSI target port shall respond to an attention condition with MESSAGE OUT phase as follows:

- a) If an attention condition is detected during a COMMAND phase, the SCSI target port shall enter

MESSAGE OUT phase after transferring part or all of the command descriptor block.

- b) If an attention condition is detected during a DATA phase, the SCSI target port shall enter MESSAGE OUT phase at the SCSI target port's earliest convenience (e.g., on a logical block boundary). The SCSI initiator port shall continue REQ/ACK handshakes until it detects a change in the C/D, I/O, or MSG signals.
- c) If an attention condition is detected during a STATUS phase, the SCSI target port shall enter MESSAGE OUT phase after the status byte has been acknowledged by the SCSI initiator port.
- d) If an attention condition is detected during a MESSAGE IN phase, the SCSI target port shall enter MESSAGE OUT phase before it sends another message. This permits a MESSAGE PARITY ERROR message from the SCSI initiator port to be associated with the appropriate message.
- e) If an attention condition is detected during a SELECTION phase the SCSI target port shall enter MESSAGE OUT phase after that SELECTION phase.
- f) If SPI information unit transfers are disabled and an attention condition is detected during a RESELECTION phase, the SCSI target port shall enter MESSAGE OUT phase after the SCSI target port has sent its IDENTIFY message for that RESELECTION phase.
- g) If the attention condition is detected during an information unit transfer, other than a SPI data stream information unit, the SCSI target port shall enter MESSAGE OUT phase at the completion of the current SPI information unit (i.e., after receiving all the ACKs from the SCSI initiator port for the current SPI information unit).
- h) If the attention condition is detected during the transfer of a SPI data stream information unit, the SCSI target port shall terminate the current stream by entering into MESSAGE OUT phase at the end of any SPI data stream information unit in the current stream.
- i) If the attention condition is detected between SPI information units the SCSI target port shall enter MESSAGE OUT phase at the completion of the next SPI information unit.

During a RESELECTION phase the SCSI initiator port should only create an attention condition to transmit an ABORT TASK, ABORT TASK SET, CLEAR TASK SET, DISCONNECT, LOGICAL UNIT RESET, NO OPERATION, or TARGET RESET message. Other uses may result in ambiguities concerning the nexus.

The SCSI initiator port shall keep the ATN signal asserted throughout the MESSAGE OUT phase if more than one byte is to be transferred. Unless otherwise specified, the SCSI initiator port may negate the ATN signal at any time, that does not violate the specified setup and hold times, except it shall not negate the ATN signal while the ACK signal is asserted during a MESSAGE OUT phase. Normally, the SCSI initiator port 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.

### 12.3 Bus reset condition

The bus reset condition is used to create a hard reset (see 12.4) for all SCSI devices on the bus and change the bus to a BUS FREE phase. This condition shall take precedence over all other phases and conditions. Any SCSI device may create the bus reset condition by asserting the RST signal for a minimum of a reset hold time.

Environmental conditions (e.g., static discharge) may generate brief glitches on the RST signal. SCSI devices shall not react to glitches on the RST signal that are less than a reset delay. 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. This limits the time to confirm the RST signal to a maximum of a bus clear delay.

### 12.4 Hard reset

A SCSI device detecting a reset event shall release all SCSI bus signals within a bus clear delay of the transition of the RST signal to true, except the RST signal if it is asserting RST. The BUS FREE phase always follows the hard reset condition. The SCSI device shall not assert the RST signal in response to a reset event on the same bus segment.

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 Architecture Model-2 standard.

Any SCSI device that detects a hard reset shall also set its transfer agreement to the default transfer agreement (see 4.12.4.2).

## **12.5 Reset events**

### **12.5.1 Reset events overview**

When a SCSI device detects a reset event it shall initiate hard reset (see 12.4).

### **12.5.2 Bus reset event**

When a SCSI device detects a bus reset condition by detecting RST true for a reset delay, it shall cause a reset event. In response to a bus reset event, a SCSI target port shall create a unit attention condition for all SCSI initiator ports. The sense key shall be set to UNIT ATTENTION with the additional sense code set to either SCSI BUS RESET OCCURRED or POWER ON, RESET, OR BUS DEVICE RESET OCCURRED.

### **12.5.3 Power on reset event**

When a SCSI device is powered on, it shall cause a reset event. In response to a power on reset event, the SCSI target port shall create a unit attention condition for all SCSI initiator ports. The sense key shall be set to UNIT ATTENTION with the additional sense code set to either POWER ON OCCURRED or POWER ON, RESET, OR BUS DEVICE RESET OCCURRED.

### **12.5.4 Target reset event**

When a SCSI device successfully receives a TARGET RESET message, it shall cause a reset event. In response to a target reset event, the SCSI target port shall create a unit attention condition for all SCSI initiator ports. The sense key shall be set to UNIT ATTENTION with the additional sense code set to either BUS DEVICE RESET FUNCTION OCCURRED or POWER ON, RESET, OR BUS DEVICE RESET OCCURRED.



## 13 SCSI bus phase sequences

### 13.1 SCSI bus phase sequences overview

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

During DT DATA phases the SCSI target port shall not change phases except at data group boundaries or SPI information unit boundaries. If a SCSI initiator port detects a change to the C/D, I/O, or MSG signals within a data group or information unit it shall consider any data transferred for that data group or information unit to have been transferred incorrectly. The SCSI initiator port shall consider this condition a protocol error and respond accordingly.

A hard reset (see 12.4) 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 for SCSI initiator ports (see 10.3).

## 13.2 Phase sequences with information units disabled

### 13.2.1 Phase sequences for physical reconnection or selection using attention condition

The allowable sequences for either physical reconnection or selection using attention condition while an information unit transfer agreement is not in effect shall be as shown in figure 75.

The normal progression for selection using attention condition (see 10.5.2) is:

- 1) from BUS FREE to ARBITRATION;
- 2) from ARBITRATION to SELECTION or RESELECTION; and
- 3) from SELECTION or RESELECTION to one or more of the information transfer phases (i.e., COMMAND, DATA, STATUS, or MESSAGE).

The final information transfer phase is normally the MESSAGE IN phase where a DISCONNECT, or TASK COMPLETE message is transferred, followed by the BUS FREE phase.

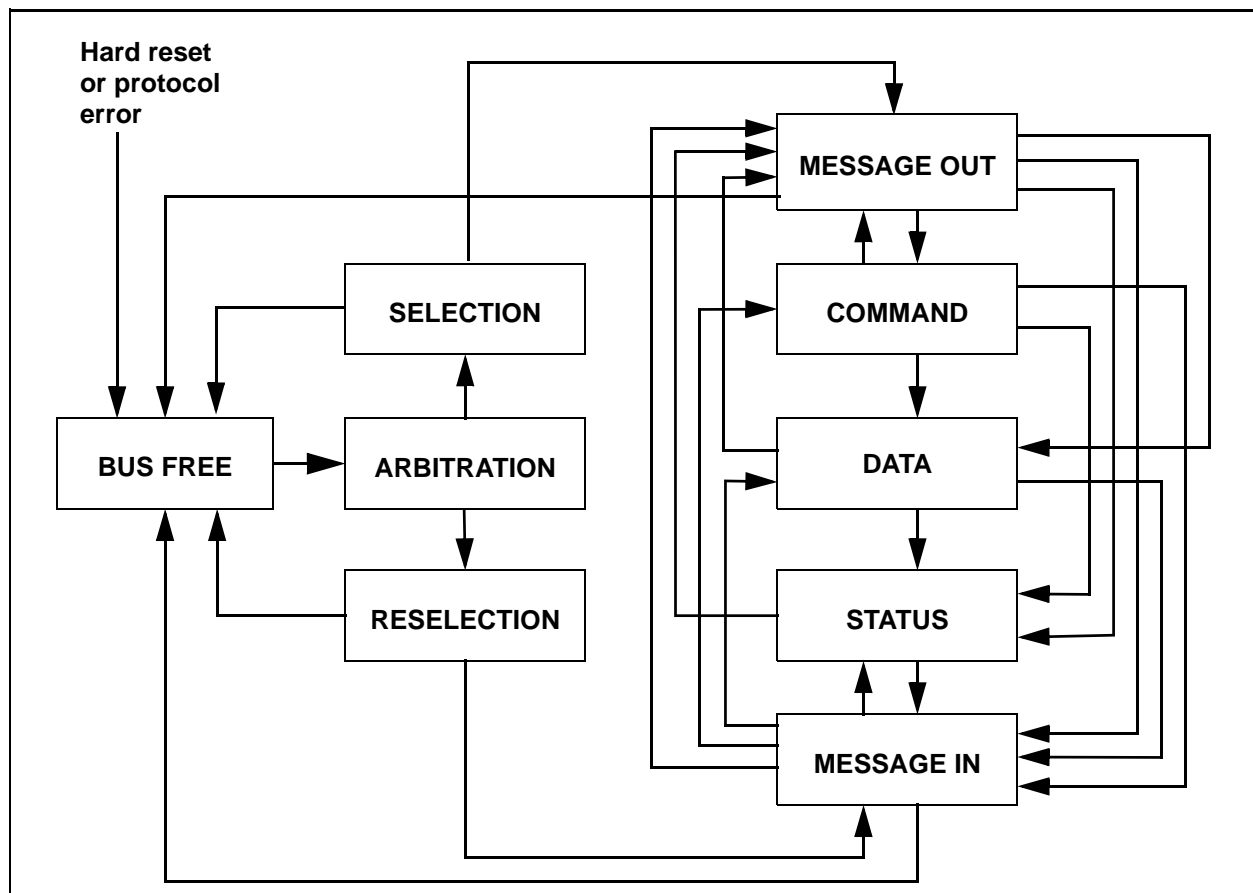


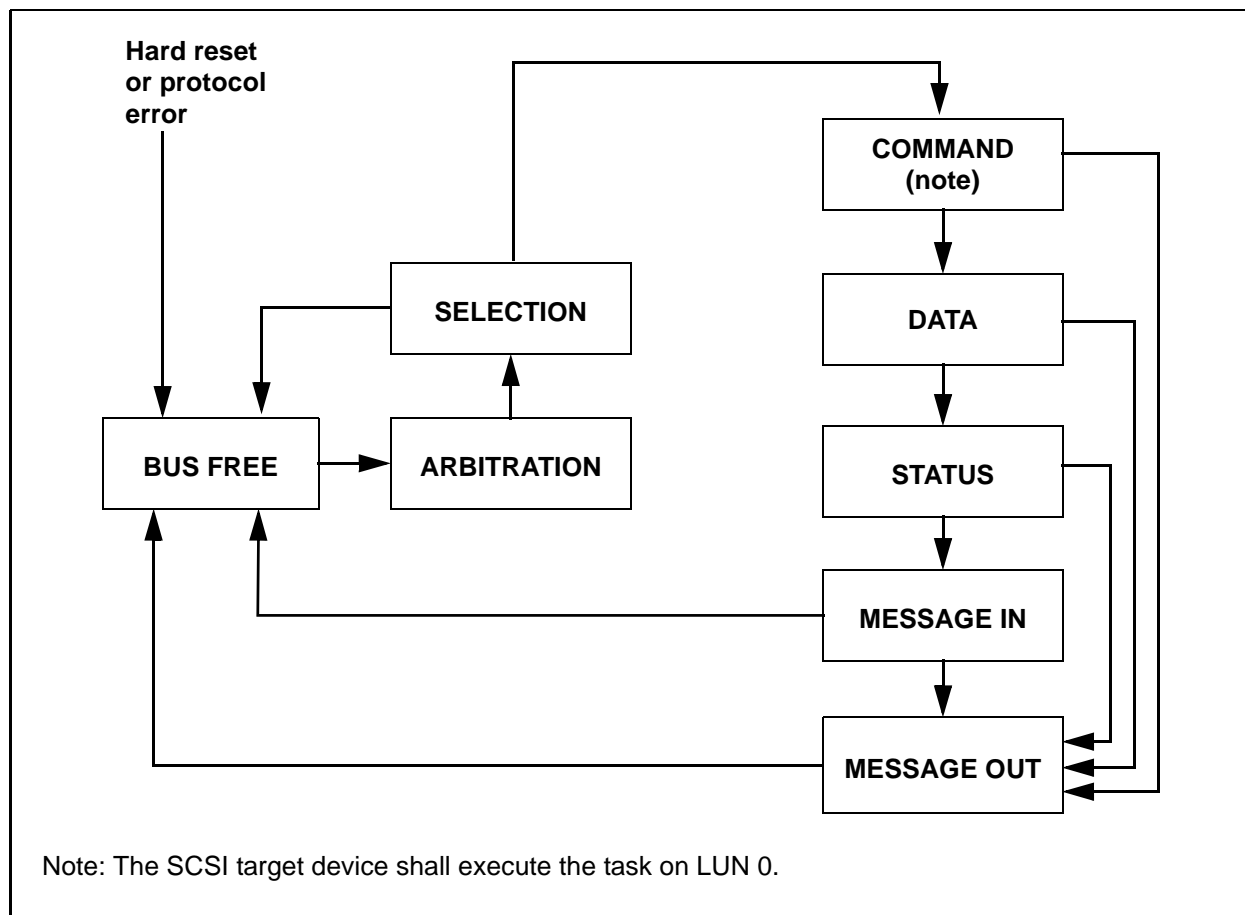
Figure 75 - Phase sequences for physical reconnection or selection using attention condition with information unit transfers disabled

### 13.2.2 Phase sequences for selection without using attention condition

The additional sequences for selection without using attention condition while an information unit transfer agreement is not in effect shall be as shown in figure 76.

The normal progression for selection without using attention condition (see 10.5.3) is;

- 1) from BUS FREE to ARBITRATION;
- 2) from ARBITRATION to SELECTION;
- 3) from SELECTION to COMMAND;
- 4) from COMMAND to DATA;
- 5) from DATA to STATUS;
- 6) from STATUS to MESSAGE IN where a TASK COMPLETE message is transferred; and
- 7) from MESSAGE IN to BUS FREE phase.



**Figure 76 - Phase sequences for selection without using attention condition with information unit transfers disabled**

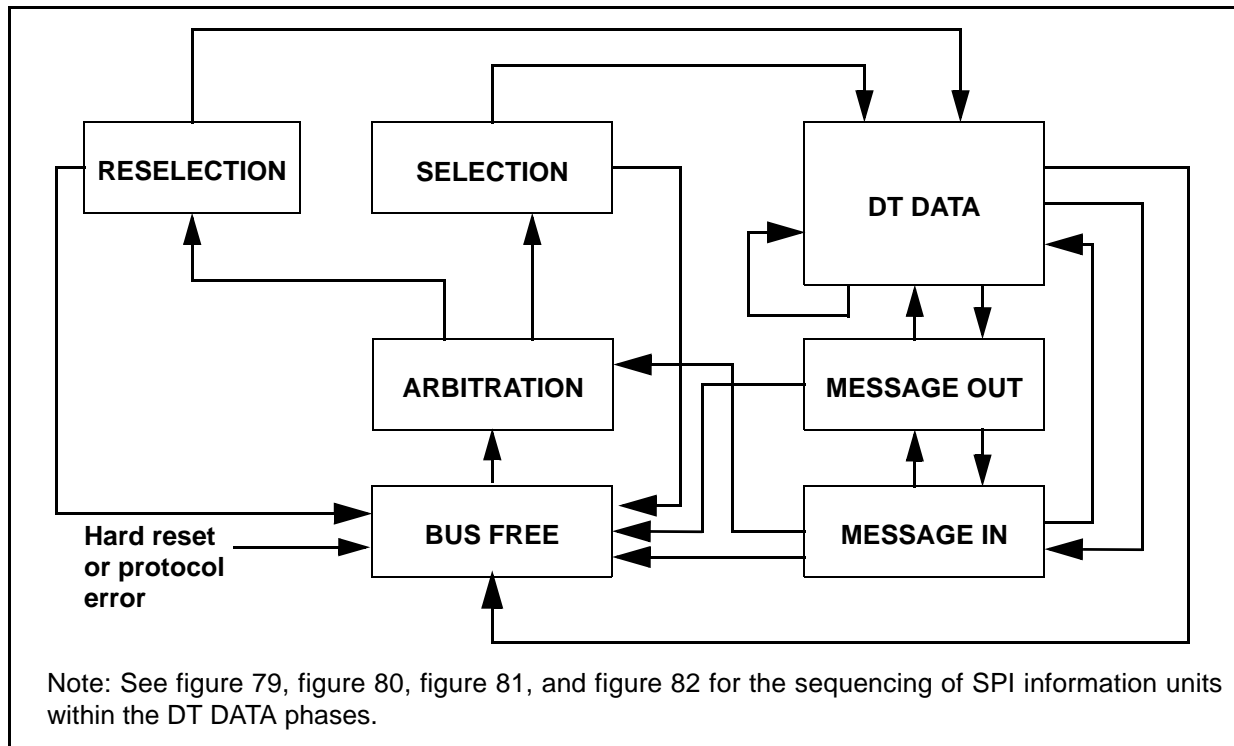
### 13.3 Phase sequences with information unit enabled

#### 13.3.1 Phase sequences for physical reconnection or selection without using attention condition

The sequences for physical reconnection or selection without using attention condition while an information unit transfer agreement is in effect shall be as shown in figure 77.

The normal progression for selection without using attention condition (see 10.5.3) if QAS is disabled is:

- 1) from BUS FREE to ARBITRATION;
- 2) from ARBITRATION to SELECTION or RESELECTION; and
- 3) from SELECTION or RESELECTION to one or more DT DATA phases; and
- 4) from the final DT DATA phase to BUS FREE.



**Figure 77 - Phase sequences for physical reconnection or selection without using attention condition with information unit transfers enabled**

### 13.3.2 Phase sequences for selection using attention condition

The sequences for a selection with attention condition while an information unit transfer agreement is in effect shall be as shown in figure 78.

The normal progression for selection using attention condition (see 10.5.2.3) if QAS is disabled is:

- 1) from BUS FREE to ARBITRATION;
- 2) from ARBITRATION to SELECTION;
- 3) from SELECTION to MESSAGE OUT;
- 4) from MESSAGE OUT to MESSAGE IN; and
- 5) from MESSAGE IN to BUS FREE.

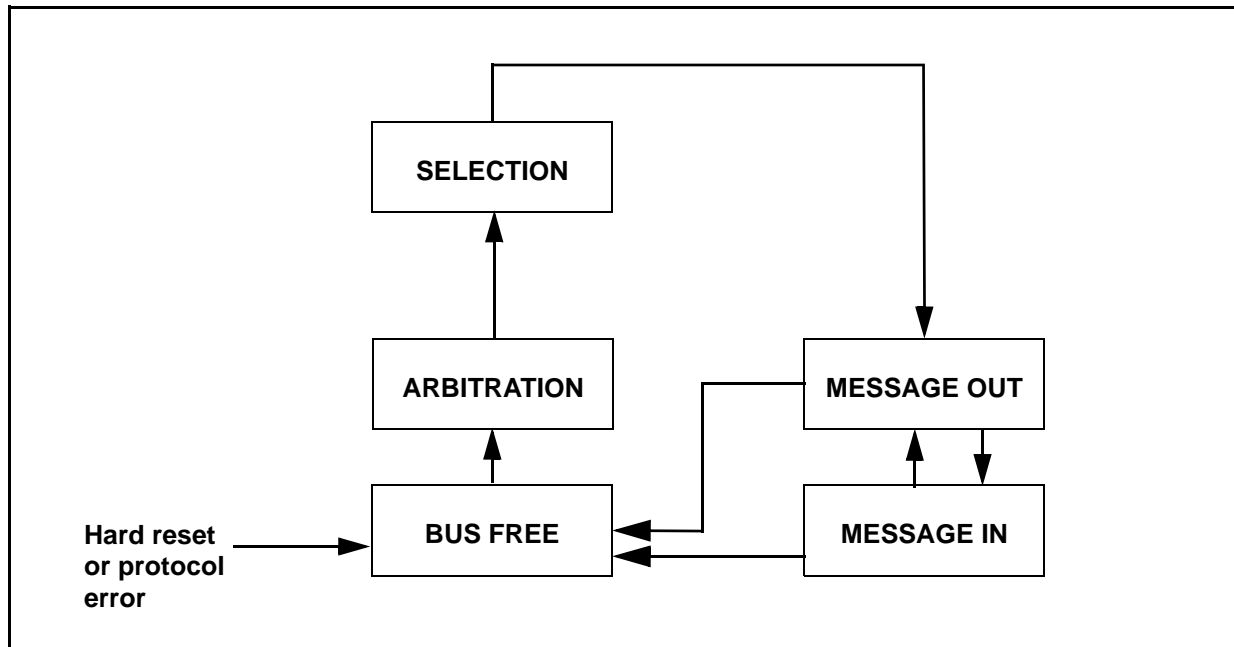


Figure 78 - Phase sequences for selection with attention condition with information unit transfers enabled

## 14 SPI information units

### 14.1 SPI information unit overview

An information unit transfer transfers data in SPI information units. The order in which SPI information units are transferred within an information unit transfer follows a prescribed sequence. When information unit transfers are enabled only SPI information units shall be transferred within the DT DATA OUT phase and DT DATA IN phase.

See table 45 for the list of SPI information units defined by this standard.

**Table 45 - SPI information units**

SPI information unit	Reference
SPI command information unit	14.3.1
SPI L_Q information unit	14.3.2
SPI data information unit	14.3.3
SPI data stream information unit	14.3.4
SPI status information unit	14.3.5

The SPI information unit sequences shall be as shown in figure 79, figure 80, figure 81, and figure 82. See figure 77 and figure 76 for the sequencing rules between the DT DATA IN or DT DATA OUT phases and the other phases.

For information on information unit exception handling see 10.7.3.3.3 and 10.7.3.3.4.

The normal progression is from SPI L\_Q information unit/SPI command information unit pairs, to SPI L\_Q information unit/SPI data information unit pairs, to SPI L\_Q information unit/SPI status information unit pairs.

NOTE 37 - A SCSI initiator port may request a BUS FREE phase by creating an attention condition and sending a DISCONNECT message on the corresponding MESSAGE OUT phase. This allows a SCSI initiator port to request the SCSI target port break up a long sequence of SPI L\_Q information unit/SPI data information unit pairs into smaller sequences.

After message phases complete that contain any negotiation (i.e., PPR or WDTR or SDTR) that results in IU\_REQ being changed, the SCSI target device shall abort all tasks, except the current task, for the SCSI initiator port participating in the negotiation and the SCSI initiator device shall abort all tasks, except the current task, for the SCSI target device. (see 4.12.4.6.2).

When an information unit transfer agreement is in effect there is no option equivalent to the "physical disconnect without sending a SAVE DATA POINTERS message." The SCSI initiator port shall save the data pointers as soon as the last byte of the last iuCRC for a SPI information unit is transferred. The save shall occur even if the SCSI initiator port detects an error in the SPI data information unit.

The SCSI target port shall not start a new information unit transfer until all previous REQ(s) have been responded to by an equal number of ACK(s) except during a sequence of SPI data stream information units (see 14.3.4).

### 14.2 Information unit transfer logical operations

SCSI devices using information unit transfers may transfer SPI information units for any number of I/O

processes by using logical connects, logical disconnects, and logical reconnects.

If there are no phase changes to a MESSAGE OUT phase or a MESSAGE IN phase then logical disconnects shall only occur at the completion of:

- a) each SPI command information unit;
- b) each SPI status information unit;
- c) each SPI data information unit;
- d) any SPI L\_Q information unit if the SPI L\_Q information unit DATA LENGTH field is set to zero; and
- e) the last SPI data stream information unit.

At completion of those SPI information units the I\_T\_L\_Q nexus becomes an I\_T nexus. The I\_T nexus remains in place until the SCSI target port does a physical disconnect or an I\_T\_L\_Q nexus is reestablished by the SCSI target port transmitting a SPI L\_Q information unit.

Logical reconnections occur on the successful SCSI target port transmission and SCSI initiator port receipt of a SPI L\_Q information unit for an existing I/O process. The logical reconnection reestablishes the I\_T\_L\_Q nexus for that I/O process.

SCSI devices using information unit transfers may receive several commands during an initial connection. This occurs when a SCSI initiator port uses the multiple command option in the SPI L\_Q information unit. For each SPI L\_Q received with a multiple command type or a last command type a logical connection occurs and an I\_T\_L\_Q nexus is formed.

If there is a phase change to a MESSAGE OUT phase or a MESSAGE IN phase then there is no logical disconnect and the I\_T\_L\_Q nexus remains in place. If a DT DATA phase follows the message phase then the L\_Q portion of the current I\_T\_L\_Q nexus shall be replaced with the L\_Q in the next SPI L\_Q information unit.

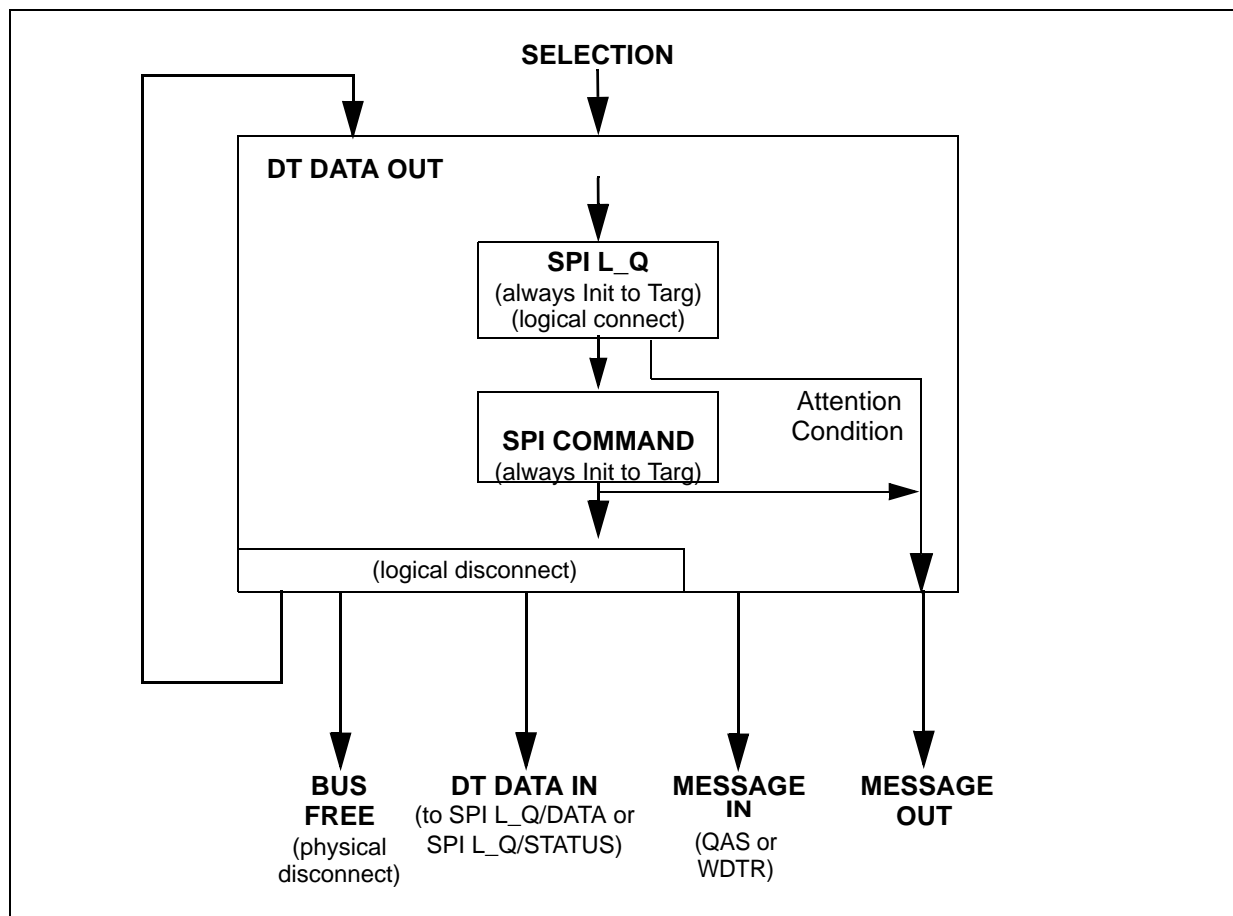


Figure 79 - SPI information unit sequence during initial connection



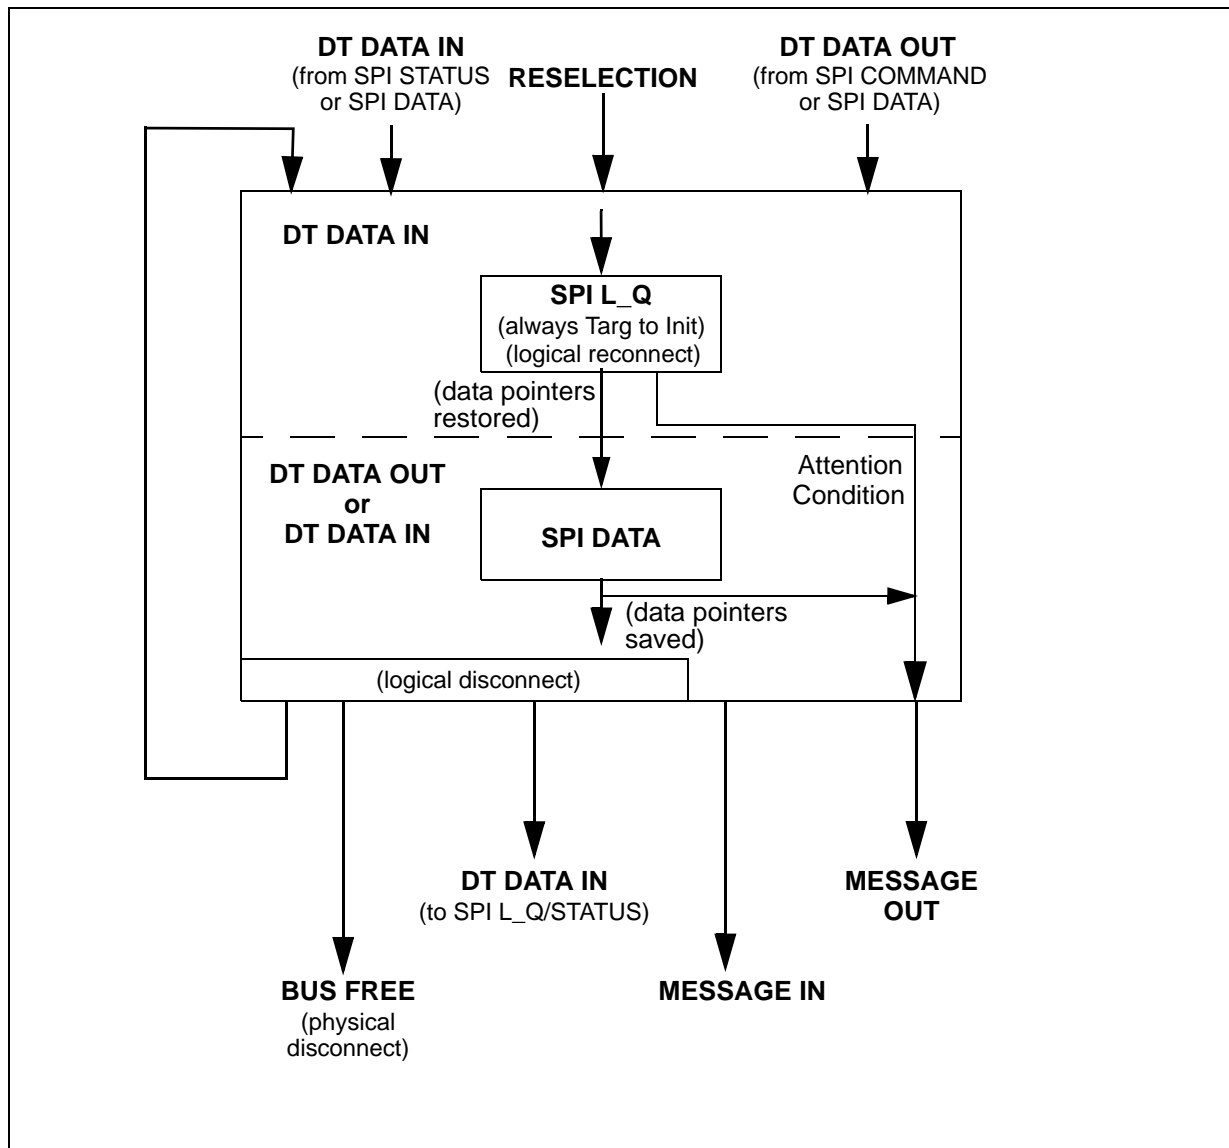


Figure 80 - SPI information unit sequence during data type transfers

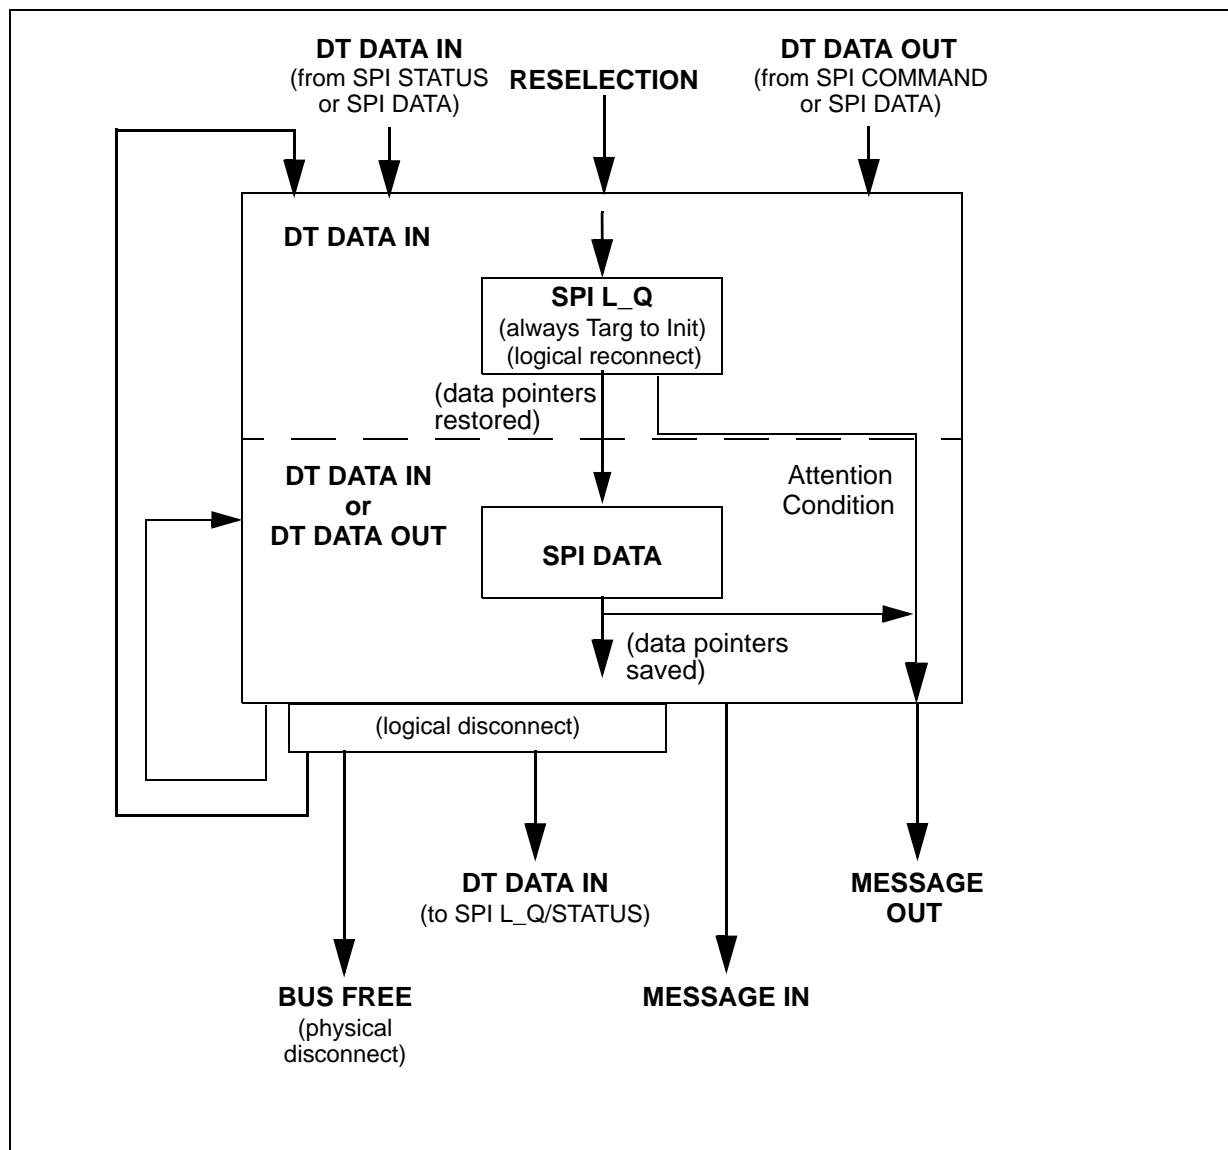


Figure 81 - SPI information unit sequence during data stream type transfers

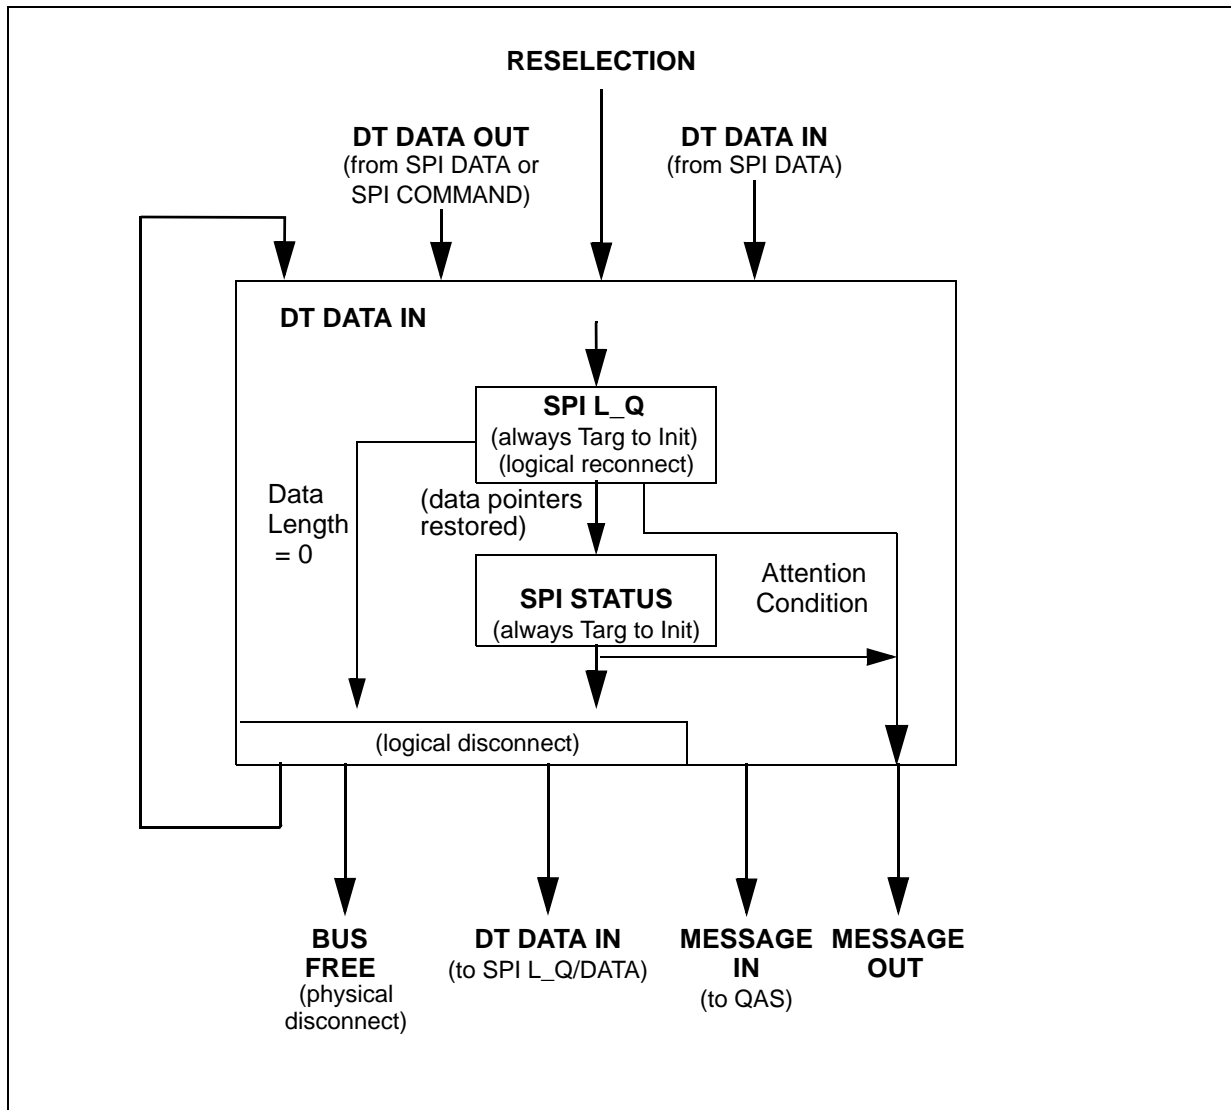


Figure 82 - SPI information unit sequence during status transfers

## 14.3 SPI information units

### 14.3.1 SPI command information unit

The SPI command information unit (see table 46) transfers CDBs, task attributes, and task management requests to be performed by a device server.

A SCSI initiator port shall consider a BUS FREE phase after the transfer of a SPI command information unit to be equivalent to receiving a DISCONNECT message.

If the SCSI target device terminates a SPI L\_Q/SPI command information unit pair for one of the following reasons:

- a) TASK SET FULL status,
- b) BUSY status,
- c) CHECK CONDITION due to a SPI command information unit iuCRC error, or
- d) a bus free due to a SPI L\_Q information unit iuCRC error

it shall have no effect on any other SPI L\_Q/SPI command information unit pair except those effects caused by any task management functions contained within the last SPI L\_Q/SPI command information unit pair.

Table 46 - SPI command information unit

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED							
1	RESERVED					TASK ATTRIBUTE		
2	TASK MANAGEMENT FUNCTIONS							
3	ADDITIONAL CDB LENGTH = (number of additional CDB bytes)/4						RDDATA	WRDATA
4	CDB							
19								
20	ADDITIONAL CDB							
n								
n+1	IUCRC							
n+2								
n+3								
n+4								
	(MSB)							(LSB)

The TASK ATTRIBUTE field is defined in table 47.

**Table 47 - TASK ATTRIBUTE**

<b>Codes</b>	<b>Description</b>
000b	Requests that the task be managed according to the rules for a simple task attribute. (See the SCSI Architecture Model-2 standard)
001b	Requests that the task be managed according to the rules for a head of queue task attribute. (See the SCSI Architecture Model-2 standard)
010b	Requests that the task be managed according to the rules for an ordered attribute. (See the SCSI Architecture Model-2 standard)
011b	Reserved
100b	Requests that the task be managed according to the rules for an automatic contingent allegiance task attribute. (See the SCSI Architecture Model-2 standard)
101b-111b	Reserved

The TASK MANAGEMENT FUNCTIONS field is defined in table 48. If a task management function fails the task manager shall terminate the task with a GOOD status. The packetized failure code shall be set to TASK MANAGEMENT FUNCTION FAILED.

If a SCSI target port does not have the resources to accept a SPI command information unit and the TASK MANAGEMENT FUNCTIONS field equals 00h the SCSI target port shall transfer all the bytes of the current SPI command information unit and shall discard the transmitted information. After transferring all the SPI command information unit bytes the SCSI target port shall transmit a SPI L\_Q/SPI status information unit pair with the status defined in the SCSI Architecture Model-2 standard for this condition. This SPI status information unit may be transferred in the same or a subsequent connection. If the SCSI initiator device has more commands to send to the SCSI target port, the SCSI initiator device shall wait at least until the next selection before those remaining commands may be sent.

If the TASK MANAGEMENT FUNCTIONS field is a supported value not equal to 00h the SCSI target device shall perform the selected task management function before processing any further SPI information units regardless of the command type. On completion of a supported task management function the SCSI target port shall go to a BUS FREE phase. No SPI status information unit shall be reported for the task management function. If the TASK MANAGEMENT FUNCTIONS field is not a supported value then the task manager shall terminate the task with a GOOD status and the packetized failure code shall be set to TASK MANAGEMENT FUNCTION NOT SUPPORTED. If a task management function fails the task manager shall terminate the task with a GOOD status. The packetized failure code shall be set to TASK MANAGEMENT FUNCTION FAILED.

**Table 48 - TASK MANAGEMENT FUNCTIONS**

<b>Codes</b>	<b>Description</b>
00h	Indicates no task management requests for the current task.
01h	The task manager shall abort the task as defined in the ABORT TASK message (see 16.5.2).
02h	The task manager shall abort the task set as defined in the ABORT TASK SET message (see 16.5.3).
04h	The task manager shall clear the task set as defined in the CLEAR TASK SET message (see 16.5.5).
08h	The task manager shall perform a logical unit reset of the selected logical unit as defined in the LOGICAL UNIT RESET message (see 16.5.6).
20h	The task manager shall perform a target reset as defined in the TARGET RESET message (see 16.5.7).
40h	The task manager shall perform a clear ACA as defined in the CLEAR ACA message (see 16.5.4).
All other values reserved	The task manager shall terminate the task with a GOOD status. The packetized failure code shall be set to TASK MANAGEMENT FUNCTION NOT SUPPORTED.

The ADDITIONAL CDB LENGTH field contains the length in 4-byte words of the ADDITIONAL CDB field.

The write data (WRDATA) bit and read data (RDDATA) bit are defined in other SCSI transport protocol standards and shall be ignored by this standard.

The CDB field contains the actual CDB to be interpreted by the addressed logical unit. The CDB field and the task attribute field is not valid and is ignored if the TASK MANAGEMENT FUNCTIONS field is not zero. Any bytes between the end of a 6 byte CDB, 10 byte CDB, or 12 byte CDB and the end of the CDB field shall be reserved.

The ADDITIONAL CDB field contains any CDB bytes beyond those contained within the standard 16 byte CDB field.

The CDB field, ADDITIONAL CDB field, and TASK ATTRIBUTE field are not valid and are ignored if the TASK MANAGEMENT FUNCTIONS field is not zero.

The contents of the CDB and ADDITIONAL CDB fields shall be as defined in the SCSI command standards.

The IUCRC field shall use the algorithm defined in 11.3.

#### **14.3.2 SPI L\_Q information unit**

The SPI L\_Q information unit (see table 49) contains L\_Q nexus information for the information unit that follows, the type of information unit that follows, and the length of information unit that follows. A SPI L\_Q information unit shall precede all SPI command information units, SPI multiple command information units, SPI data information units, SPI status information units, and the first of an uninterrupted sequence of SPI data stream information units.

The receipt of an error free (i.e., no iuCRC error) SPI L\_Q information unit by a SCSI initiator port shall cause the SCSI initiator port to restore the data pointers.



Table 49 - SPI L\_Q information unit

Bit Byte	7	6	5	4	3	2	1	0
0	TYPE							
1	RESERVED							
2	(MSB) TAG (LSB)							
3								
4	(MSB) LOGICAL UNIT NUMBER (LSB)							
11								
12	RESERVED							
13	(MSB) DATA LENGTH (LSB)							
14								
15								
16	BIDI DIRECTION		RESERVED					
17	RESERVED							
18	(MSB) IUCRC INTERVAL (LSB)							
19								
20	(MSB) IUCRC (LSB)							
21								
22								
23								

The TYPE field is defined in table 50. If a SCSI initiator port receives a type code that is not defined in table 50 that SCSI initiator port shall follow the procedures defined in 10.7.3.3.3. If a SCSI target port receives a type code that is not defined in table 50 that SCSI target port shall follow the procedures defined in 10.7.3.3.4.

**Table 50 - TYPE**

<b>Codes</b>	<b>Type</b>	<b>Description</b>
01h	Last Command	Sent by a SCSI initiator device to indicate a SPI command information unit shall follow this SPI L_Q information unit. Indicates the SCSI initiator device shall not send any more SPI command information units during the current connection. The value of the DATA LENGTH field shall be greater than or equal to 14h and less than or equal to 90h. The IUCRC INTERVAL field shall be set to zero by the SCSI initiator device and ignored by the SCSI target device. The BIDI DIRECTION field shall be set to zero by the SCSI initiator device and ignored by the SCSI target device.
02h	Multiple Command	Sent by a SCSI initiator device to indicate a SPI command information unit shall follow this SPI L_Q information unit. Indicates the SCSI initiator device has another SPI L_Q information unit and SPI command information unit during the current connection. The value of the DATA LENGTH field shall be greater than or equal to 14h and less than or equal to 90h. The IUCRC INTERVAL field shall be set to zero by the SCSI initiator device and ignored by the SCSI target device. The BIDI DIRECTION field shall be set to zero by the SCSI initiator device and ignored by the SCSI target device.
04h	Data	Sent by a SCSI target device to indicate a SPI data information unit shall follow this SPI L_Q information unit. The DATA LENGTH field shall not be set to zero. For a bidirectional command, the direction of the SPI data information unit shall be indicated in the BIDI DIRECTION field of the SPI L_Q information unit as defined in table 51
05h	Data Stream	Sent by a SCSI target device to indicate an unspecified number of SPI data stream information unit shall follow this SPI L_Q information unit. The DATA LENGTH field shall not be set to zero. For a bidirectional command, the direction of the SPI data stream information units shall be indicated in the BIDI DIRECTION field of the SPI L_Q information unit as defined in table 51.
08h	Status	Sent by a SCSI target device to indicate a SPI status information unit may follow this SPI L_Q information unit. A length of zero in the DATA LENGTH field shall indicate no SPI status information unit shall follow the SPI L_Q information unit (see 14.3.5). The IUCRC INTERVAL field shall be set to zero by the SCSI target device and ignored by the SCSI initiator device. The BIDI DIRECTION field shall be set to zero by the SCSI target device and ignored by the SCSI initiator device.
F0h - FFh		Vendor specific
All others		Reserved

The TAG field is an 16-bit integer assigned by the application client and sent to the SCSI initiator port in the send SCSI command request (see 19.3.2). For more details on tags see 16.4.1.

The LOGICAL UNIT NUMBER field specifies the address of the logical unit of the I\_T\_L\_Q nexus for the current task. The structure of the logical unit number field shall be as defined in the SCSI Architecture Model-2

standard. If the addressed logical unit does not exist, the task manager shall follow the SCSI rules for selection of invalid logical units as defined in the SCSI Primary Commands-3 standard.

The DATA LENGTH field contains the length in bytes of the following information units. For SPI data stream information units the data length field contains the length in bytes of each SPI data stream information unit that follows (i.e., the total number of bytes transferred would equal the data length times the number of SPI data stream information units transferred). The data length shall not include any of the 4 byte iuCRC nor any transmitted pad bytes (e.g., a data length of 509 with a iuCRC interval of zero or greater than 509 would transfer 509 bytes of data plus 3 bytes of pad plus 4 bytes of iuCRC for a total transfer of 516 bytes). The SCSI target device shall not set the data length to a value that exceeds the maximum burst size as defined in the Disconnect-Reconnect mode page (see 18.1.2).

The BIDI DIRECTION field determines the data direction if the command is a bidirectional command and the type code is data or data stream. The code values for the BIDI DIRECTION field are defined in table 51.

**Table 51 - BIDI DIRECTION**

<b>Codes</b>	<b>Description</b>
00b	A unidirectional command or a type code other than data or data stream (see table 50).
01b	A bidirectional command transferring data from the SCSI initiator device to the SCSI target device.
10b	A bidirectional command transferring data from the SCSI target device to the SCSI initiator device.
11b	Reserved

The IUCRC INTERVAL field contains the length in bytes of the data to be sent before a iuCRC is transferred. The iuCRC interval length shall not include the 4 byte iuCRC nor any transmitted pad bytes (e.g., an iuCRC interval length of 510 transfer 510 bytes of data plus 2 bytes of pad plus 4 bytes of iuCRC for a total transfer of 516 bytes). The iuCRC interval shall be a multiple of two (i.e., odd numbers are not allowed). If the iuCRC interval is equal to zero or is greater than or equal to the data length only one iuCRC shall occur at the end of the SPI information unit.

The IUCRC field shall use the algorithm defined in 11.3.

#### **14.3.3 SPI data information unit**

The SPI data information unit (see table 52) contains data.

The detection of a BUS FREE phase following a SPI data information unit by a SCSI initiator port shall be equivalent to the SCSI initiator port receiving a DISCONNECT message.

The detection of a QAS REQUEST message following a SPI data information unit by a SCSI initiator port shall be equivalent to the SCSI initiator port receiving a DISCONNECT message.

**Table 52 - SPI data information unit**

Bit Byte	7	6	5	4	3	2	1	0
0	DATA							
n								
n+1	(MSB)							
n+2	IUCRC							
n+3								
n+4								
	(LSB)							

The DATA field may contain any type of information (e.g., parameter lists, mode pages, user data).

The IUCRC field shall be calculated using the algorithm defined in 11.3. If the IUCRC INTERVAL field of the SPI L\_Q information unit contains a value greater than zero and less than the data length then there is an IUCRC field at each iucrc interval in addition to the iucrc shown in table 52. These additional IUCRC fields are not shown in table 52.

#### 14.3.4 SPI data stream information unit

The SPI data stream information unit (see table 53) contains data.

All the SPI data stream information units transferred after a SPI L\_Q information unit with a type of data stream shall be the size indicated in the DATA LENGTH field of the SPI L\_Q information unit.

If the data transfer size is not a multiple of the data length, the SCSI target port shall end the stream at a data length boundary and shall send a new SPI L\_Q with a smaller data length to finish the data transfer. The new SPI L\_Q may or may not be sent during the current physical connection.

During write streaming the sequence of SPI data stream information units shall end with any change to the C/D, I/O, or MSG signals on a SPI data stream information unit boundary. If during write streaming SPI data stream information units a SCSI initiator port detects a REQ transition after transmitting the last iucrc for a SPI data stream information unit that SCSI initiator port shall transmit the next SPI data stream information unit (see 8.2 and table 32).

During read streaming the SCSI target port shall end a sequence of SPI data stream information units by performing one of the following:

- should assert the P\_CRCA signal before the end of the current SPI data stream information unit boundary (see 8.2 and table 32); or
- may change the C/D, I/O, or MSG signals on a SPI data stream information unit boundary.

If during the last SPI data stream information unit, of a read stream, the P\_CRCA signal was not asserted and a SCSI initiator port detects a REQ transition after receiving the last iucrc for a SPI data stream information unit that SCSI initiator port shall receive the next SPI data stream information unit. If during the

last SPI data stream information unit the P\_CRCA signal was asserted and a SCSI initiator port detects a REQ transition after receiving the last iuCRC for a SPI data stream information unit that SCSI initiator port shall logically disconnect from the current I\_T\_L\_Q nexus.

If during a sequence of SPI data stream information units a SCSI initiator port detects any change to the C/D, I/O, or MSG signals after transmitting or receiving the last iuCRC for a SPI data stream information unit that SCSI initiator port shall consider the current I/O process to be logically disconnected or in the case of detecting a BUS FREE phase or a MESSAGE IN phase to be physically disconnected.

The detection of a BUS FREE phase following a SPI data stream information unit by a SCSI initiator port shall be equivalent to the SCSI initiator port receiving a DISCONNECT message.

The detection of a QAS REQUEST message following a SPI data stream information unit by a SCSI initiator port shall be equivalent to the SCSI initiator port receiving a DISCONNECT message.

To end a sequence of SPI data stream information units a SCSI initiator port may request a disconnect by establishing an attention condition. The SCSI initiator port shall continue to transfer or receive data, pad bytes (if any), and iuCRC(s) until the SCSI target port changes to the MESSAGE OUT phase.

During a sequence of SPI data stream information units the offset count is not required to go to zero at the boundary of any SPI data stream information unit if the next SPI information unit is a SPI data stream information unit.

**Table 53 - SPI data stream information unit**

Bit Byte	7	6	5	4	3	2	1	0
0	DATA							
n								
n+1	(MSB)							
n+2	IUCRC							
n+3								
n+4								
	(LSB)							

The DATA field may contain any type of information (e.g., parameter lists, mode pages, user data).

The IUCRC field shall be calculated using the algorithm defined in 11.3. If the IUCRC INTERVAL field of the SPI L\_Q information unit contains a value greater than zero and less than the data length then there is an IUCRC field at each iuCRC interval in addition to the iuCRC shown in table 53. These additional IUCRC fields are not shown in table 53.

#### 14.3.5 SPI status information unit

The SPI status information unit (see table 54) contains the completion status of the task indicated by the preceding SPI L\_Q information unit. The SCSI target port shall consider the SPI status information unit

transmission to be successful when there is no attention condition on the transfer of the information unit.

If a task completes with a GOOD status, a SNSVALID bit of zero, and a RSPVALID bit of zero then the SCSI target port shall set the DATA LENGTH field in the SPI L\_Q information unit (see 14.3.2) to zero.

**Table 54 - SPI status information unit**

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED							
1	RESERVED							
2	RESERVED FOR FCP						SNSVALID	RSPVALID
3	STATUS							
4	(MSB)	SENSE DATA LIST LENGTH (n-m)						
7								(LSB)
8	(MSB)	PACKETIZED FAILURES LIST LENGTH (m-11)						
11								(LSB)
12	(MSB)	PACKETIZED FAILURES						
m								(LSB)
1+m	SENSE DATA							
n								
n+1	(MSB)	IUCRC						
n+2								
n+3								
n+4								(LSB)

A sense data valid bit (SNSVALID) of zero indicates the sense data list length shall be ignored and no sense data is provided. A SNSVALID bit of one indicates the SENSE DATA LIST LENGTH field specifies the number of bytes in the SENSE DATA field. If the STATUS field contains a CHECK CONDITION status the SNSVALID bit shall be set to one.

If sense data is provided, SNSVALID shall be set to one and the SENSE DATA LIST LENGTH field shall specify the number of bytes in the SENSE DATA field. The SENSE DATA LIST LENGTH field shall only contain even lengths greater than zero and shall not be set to a value greater than 252.

If no sense data is provided, SNSVALID shall be set to zero. The SCSI initiator port shall ignore the SENSE DATA LIST LENGTH field and shall assume a length of zero.

If packetized failure data is provided, the packetized failures valid bit (RSPVALID) shall be set to one and the PACKETIZED FAILURES LIST LENGTH field shall specify the number of bytes in the PACKETIZED FAILURES field. The PACKETIZED FAILURES LIST LENGTH field shall contain a length of 4. Other lengths are reserved for future standardization.

If no packetized failure data is provided, RSPVALID shall be set to zero. The SCSI initiator port shall ignore the PACKETIZED FAILURES LIST LENGTH field and shall assume a length of zero.

The STATUS field contains the status of a task that completes. See the SCSI Architecture Model-2 standard for a list of status codes.

The PACKETIZED FAILURES field (see table 55) contains information describing the packetized failures detected during the execution of a task. The PACKETIZED FAILURES field shall contain valid information if the SCSI target device detects any of the conditions described by the packetized failure code (see table 56).

**Table 55 - PACKETIZED FAILURES field**

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED							
1	RESERVED							
2	RESERVED							
3	PACKETIZED FAILURE CODE							

The PACKETIZED FAILURE CODE field is defined in table 56.

**Table 56** - PACKETIZED FAILURE CODE

<b>Codes</b>	<b>Description</b>
00h	NO FAILURE
01h	Reserved
02h	SPI COMMAND INFORMATION UNIT FIELDS INVALID
03h	Reserved
04h	TASK MANAGEMENT FUNCTION NOT SUPPORTED
05h	TASK MANAGEMENT FUNCTION FAILED
06h	INVALID TYPE CODE RECEIVED IN SPI L_Q INFORMATION UNIT
07h	ILLEGAL REQUEST RECEIVED IN SPI L_Q INFORMATION UNIT
08h-FFh	Reserved

The SENSE DATA field contains the information specified by the SCSI Primary Commands-3 standard for presentation by the REQUEST SENSE command. The proper sense data shall be presented when a SCSI status byte of CHECK CONDITION is presented as specified by the SCSI Primary Commands-3 standard.

The IUCRC field shall be calculated using the algorithm defined in 11.3.



## 15 SCSI pointers

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

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-out pointer points to the start of the data-out area until the SCSI target port sends a SAVE DATA POINTERS message for the task or after the SCSI initiator port successfully transmits a SPI data information unit. The saved data-in pointer points to the start of the data-in area until the SCSI target port sends a SAVE DATA POINTERS message for the task or after the SCSI initiator port successfully receives a SPI data information unit.

In response to the SAVE DATA POINTERS message or successful receipt or transmission of a SPI data information unit, the SCSI initiator port stores the value of the current data-out and data-in pointers into the saved data-out and data-in pointers for that task. If information units are disabled the SCSI target port may restore the active pointers to the saved pointer values for the current task by sending a RESTORE POINTERS message to the SCSI initiator port. If information units are disabled the SCSI initiator port then copies the set of saved pointers into the set of active pointers. Whenever a SCSI target port does a physical disconnect 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 a physical reconnection of the task or a successful receipt of a SPI L\_Q information unit.

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

If information units are disabled, the SCSI target port shall do a physical disconnect and physical reconnect to change the data transfer direction.

## 16 SCSI messages

### 16.1 SCSI messages overview

SCSI transport protocol messages allow communication between a SCSI initiator port and a SCSI target port 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 transport protocol messages allow communication between the application client and the task manager for the purpose of task management. The task management functions are defined in the SCSI Architecture Model-2 standard. Messages that convey the task management functions are defined by this standard.

### 16.2 Message protocols and formats

#### 16.2.1 Message protocol rules

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

If an information unit agreement is not in effect, the first message sent by the SCSI initiator port after a successful SELECTION phase with an attention condition shall be an IDENTIFY, ABORT TASK SET (see 16.5.3), or TARGET RESET message. If a SCSI target port receives any other message it shall cause an unexpected bus free by generating a BUS FREE phase (see 10.3).

If the first message is an IDENTIFY message, then it may be followed by other messages, such as the first of a pair of SYNCHRONOUS DATA TRANSFER REQUEST messages. With tagged queuing a task attribute shall follow the IDENTIFY message, then more messages may follow. The IDENTIFY message establishes a logical connection between the SCSI initiator port and the specified logical unit within the SCSI target device known as an I\_T\_L nexus.

If an information unit agreement is not in effect, after the RESELECTION phase, the SCSI target port'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 physical connection or physical reconnection; if a SCSI target port receives a second IDENTIFY message with a different logical unit number during a physical connection or physical reconnection, it shall cause an unexpected bus free by generating a BUS FREE phase (see 10.3).

If an information unit agreement is in effect the SCSI target port enters a DT DATA phase after the RESELECTION phase as described in figure 77.

All SCSI initiator ports shall implement the mandatory messages tabulated in the "Initiator" column of table 60, table 69, and table 74. All SCSI target ports shall implement the mandatory messages tabulated in the "Target" column of table 60, table 69, and table 74.

Whenever an I\_T\_L nexus is established by a SCSI initiator port that is allowing physical disconnection, the SCSI initiator port 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 or a successful receipt of a SPI L\_Q information unit.

#### 16.2.2 Message formats

One-byte, Two-byte, and Extended message formats are defined. The first byte of the message determines the format as defined in table 57.

**Table 57 - Message format**

<b>Code</b>	<b>Message format</b>
00h	One-byte message
01h	Extended messages
02h - 0Ah	One-byte messages
0Bh	Obsolete One-byte messages
0Ch - 0Eh	One-byte messages
0Fh - 10h	Reserved One-byte messages
11h - 13h	Obsolete 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 - 54h	Reserved
55h	One-byte message
56h - 7Fh	Reserved
80h - FFh	One-byte message (IDENTIFY)

### 16.2.3 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 the message that is to be performed as defined in table 60, table 69, and table 74.

### 16.2.4 Two-byte messages

Two-byte messages consist of two consecutive bytes transferred during a MESSAGE IN phase or a MESSAGE OUT phase. The code of the first byte determines the message that is to be performed as defined in table 60, table 69, and table 74. The second byte is a parameter byte that is used as defined in the message description.

### 16.2.5 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 58.

**Table 58 - 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 59.

**Table 59 - Extended message codes**

Code	Extended message
00h	MODIFY DATA POINTER
01h	SYNCHRONOUS DATA TRANSFER REQUEST
02h	Reserved
03h	WIDE DATA TRANSFER REQUEST
04h	PARALLEL PROTOCOL REQUEST
05h	MODIFY BIDIRECTIONAL DATA POINTER
06h - FFh	Reserved

The EXTENDED MESSAGE ARGUMENTS are specified within the extended message descriptions (see 16.3.9, 16.3.10, 16.3.12, 16.3.16, and 16.3.18).

## 16.3 Link control messages

### 16.3.1 Link control message codes

Table 60 - Link control message codes

Code	Support				Message Name	Direction		Clear Attention Condition <sup>a</sup>
	IU Transfers Disabled		IU Transfers Enabled					
	Init	Targ	Init	Targ				
04h	O	O	NS	NS	DISCONNECT	In		n/a
04h	O	O	O	O	DISCONNECT		Out	Yes
80h+	M	O	NS	NS	IDENTIFY <sup>b</sup>	In		n/a
80h+	M	M	NS	NS	IDENTIFY <sup>b</sup>		Out	Not required
23h	O	O	NS	NS	IGNORE WIDE RESIDUE <sup>b</sup>	In		n/a
05h	M	M	M	M	INITIATOR DETECTED ERROR		Out	Yes
0Ah	O	O	NS	NS	LINKED COMMAND COMPLETE <sup>b</sup>	In		n/a
09h	M	M	M	M	MESSAGE PARITY ERROR		Out	Yes
07h	M	M	M	M	MESSAGE REJECT	In	Out	Yes
01h,05h,00h	O	O	NS	NS	MODIFY DATA POINTER	In		n/a
01h,09h,05h	O	O	NS	NS	MODIFY BIDIRECTIONAL DATA POINTER	In		n/a
08h	M	M	M	M	NO OPERATION		Out	Yes
01h,06h,04h	M	M	M	M	PARALLEL PROTOCOL REQUEST	In	Out	Yes
55h	NS	O	O	O	QAS REQUEST	In		n/a
03h	O	O	NS	NS	RESTORE POINTERS	In		n/a
02h	O	O	NS	NS	SAVE DATA POINTERS	In		n/a
01h,03h,01h	O	O	M	M	SYNCHRONOUS DATA TRANSFER REQUEST	In	Out	Yes
00h	M	M	NS	NS	TASK COMPLETE <sup>b</sup>	In		n/a
01h,02h,03h	O	O	M	M	WIDE DATA TRANSFER REQUEST	In	Out	Yes
Key: M=Mandatory support, O=Optional support, NS=Not supported In=SCSI target port to SCSI initiator port, Out=SCSI initiator port to SCSI target port Yes=SCSI initiator port shall clear the attention condition before last ACK of the MESSAGE OUT phase. Not required=SCSI initiator port may or may not clear the attention condition before last ACK of the MESSAGE OUT phase (see 12.2). NS=Not supported as a message. The receiving SCSI device shall reject this message. n/a=Not applicable. Init=SCSI initiator port, Targ=SCSI target port 80h+=Codes 80h through FFh are used for IDENTIFY messages.								
<sup>a</sup> The clear attention condition only applies during MESSAGE OUT phase <sup>b</sup> When IU transfers are enabled, the equivalent functions are implemented in the SPI L_Q information unit (see 14.3.2).								

### 16.3.2 DISCONNECT

The DISCONNECT message is sent from a SCSI target port to inform a SCSI initiator port that the SCSI target port plans to do a physical disconnect by releasing the BSY signal, and that a later physical reconnect is going to be required in order to complete the current task. This message shall not cause the SCSI initiator port to save the data pointers. The SCSI target port shall consider the message transmission to be successful when there is no attention condition on the DISCONNECT message.

After successfully sending this message the SCSI target port shall go to the BUS FREE phase by releasing the BSY signal.

If an information unit transfer agreement is not in effect any SCSI target port that breaks data transfers into one or more physical reconnections shall end each successful data transfer (except possibly the last) with a SAVE DATA POINTERS - DISCONNECT message sequence.

If an information unit transfer agreement is in effect SCSI target ports shall not transmit a DISCONNECT message.

This message may also be sent from a SCSI initiator port to a SCSI target port to instruct the SCSI target port to do a physical disconnect. If this option is enabled and a DISCONNECT message is received the SCSI target port shall either:

- a) if an information unit transfer agreement is not in effect switch to MESSAGE IN phase, send the DISCONNECT message to the SCSI initiator port (possibly preceded by SAVE DATA POINTERS message), and then do a physical disconnect by releasing BSY; or
- b) if an information unit transfer agreement is in effect, regardless of the QAS mode, do a physical disconnect by releasing BSY.

After releasing the BSY signal, the SCSI target port shall not participate in another ARBITRATION phase for at least a disconnection delay or the time limit specified in the PHYSICAL DISCONNECT TIME LIMIT mode parameter (see 18.1.2) whichever is greater. If this option is disabled or the SCSI target port is not able to do a physical disconnect at the time when it receives the DISCONNECT message from the SCSI initiator port, the SCSI target port shall respond by sending a MESSAGE REJECT message to the SCSI initiator port.

### 16.3.3 IDENTIFY

The IDENTIFY message (see table 61) is sent by either the SCSI initiator port or the SCSI target port to establish an I\_T\_L nexus when information unit transfers are disabled.

**Table 61 - 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 SCSI initiator port has granted the SCSI target port the privilege of doing physical disconnects. A DISCPRIV bit of zero specifies that the SCSI target port shall not do physical disconnects. This bit is not defined and shall be set to zero when an IDENTIFY message is sent by a SCSI target port.

The SCSI target port shall generate a BUSY status (see SCSI Architecture Model-2 standard) for a task not granting a physical disconnect privilege (i.e., DISCPRIV bit set to zero) in the IDENTIFY message if:

- a) there are any pending tasks, and
- b) the SCSI target device determines that a physical reconnection of one or more pending tasks is required before the current task may be 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; if a SCSI target port receives a second IDENTIFY message with a different value in this field, it shall cause an unexpected bus free (see 10.3) by generating a BUS FREE phase. Thus a SCSI initiator port may change the DISCPRIV bit, but shall not attempt to switch to another task. See the DTDC field of the physical Disconnect-Reconnect mode page in the 18.1.2 for additional controls over physical disconnection.

An implied RESTORE POINTERS message shall be performed by the SCSI initiator port following successful identification of the nexus during the MESSAGE IN phase of a physical reconnection or a successful receipt of a SPI L\_Q information unit.

Identification is considered successful during a SCSI initiator port's initial connection or a physical reconnection when the SCSI target port detects no error during the transfer of the IDENTIFY message and an optional task attribute message in the MESSAGE OUT phase following the SELECTION phase. See 16.4 for the ordering of the IDENTIFY and task attribute messages. See 10.11.5 for handling SCSI target port detected errors during the MESSAGE OUT phase.

Identification is considered successful during a SCSI target port's physical reconnect when there is no attention condition on either the IDENTIFY message or the SIMPLE message for an I\_T\_L\_Q nexus in the MESSAGE IN phase following the RESELECTION phase. See the 16.4 for the ordering of the IDENTIFY and task attribute messages. See 12.2, item d), for handling SCSI target port detected errors during the MESSAGE IN phase.

#### 16.3.4 IGNORE WIDE RESIDUE

The IGNORE WIDE RESIDUE message (see table 62) shall be sent from a SCSI target port to indicate that the number of valid bytes sent in the last REQ/ACK handshake data of a DATA IN phase is less than the negotiated transfer width. When information unit transfers are disabled the IGNORE WIDE RESIDUE message shall be sent following that DATA IN phase and prior to any other messages.

If the residual byte contains valid data then the IGNORE WIDE RESIDUE message should not be sent.

**Table 62 - 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)							

The NUMBER OF BYTES TO IGNORE field indicates the number of invalid data bytes transferred. See table 63 for a definition of the IGNORE field codes.

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

**Table 63 - IGNORE field definition**

Codes	Invalid data bits
	wide transfers
00h	Reserved
01h	DB(15-8)
02h	Obsolete
03h	Obsolete
04h-FFh	Reserved

### 16.3.5 INITIATOR DETECTED ERROR

The INITIATOR DETECTED ERROR message is sent from a SCSI initiator port to inform a SCSI target port that an error has occurred that does not preclude the SCSI target port 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 SCSI initiator port 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 physical disconnect followed by a reconnect shall cause the pointers to be restored to their defined prior state.

### 16.3.6 LINKED COMMAND COMPLETE

The LINKED COMMAND COMPLETE message is sent from a SCSI target device to a SCSI initiator device to indicate that a linked command has completed and that status has been sent. The SCSI initiator port shall then set the pointers to the initial state for the next linked command.

### 16.3.7 MESSAGE PARITY ERROR

The MESSAGE PARITY ERROR message is sent from the SCSI initiator port to the SCSI target port to indicate that it received a message byte with a parity error (see 10.11.5).

In order to indicate its intentions of sending this message, the SCSI initiator port shall create an attention condition on the message byte that has the parity error. This provides an interlock so that the SCSI target port is able to determine which message byte has the parity error. If the SCSI target port receives this message under any other circumstance, it shall signal a catastrophic error condition by going to a BUS FREE phase without any further information transfer attempt (see 10.2).

If the SCSI target port attempts a retry after receiving the MESSAGE PARITY ERROR message the SCSI target port shall return to the MESSAGE IN phase before switching to some other phase, the SCSI target port shall resend the entire message that had the parity error.

### 16.3.8 MESSAGE REJECT

The MESSAGE REJECT message is sent from either the SCSI initiator port or SCSI target port 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 SCSI initiator port shall create an attention condition on the message byte that is to be rejected. If the SCSI target port receives this message under any other circumstance, it shall reject this message.

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

After a SCSI target port sends a MESSAGE REJECT message and if the attention condition is still set, then it shall return to the MESSAGE OUT phase. The subsequent MESSAGE OUT phase shall begin with the first byte of a message.

### 16.3.9 MODIFY DATA POINTER

The MODIFY DATA POINTER message (see table 64) is sent from the SCSI target port to the SCSI initiator port and requests that the signed ARGUMENT be added using two's complement arithmetic to the value of the current data pointer. The data pointer is whichever of the data-out or data-in pointers is being used by the command. The enable modify data pointer (EMDP) bit in the Disconnect-Reconnect mode page (see 18.1.2) indicates whether or not the SCSI target port is permitted to issue the MODIFY DATA POINTER message. The SCSI target port shall only issue the MODIFY DATA POINTER message during a unidirectional command.

It is recommended that the SCSI target port not attempt to move the data pointer outside the range addressed by the command. SCSI initiator ports may or may not place further restrictions on the acceptable values. Should the SCSI target port send an ARGUMENT value that is not supported by the SCSI initiator port, the SCSI initiator port may reject the value by responding with the MESSAGE REJECT message. In this case, the data pointer is not changed from its value prior to the rejected MODIFY DATA POINTER message.

If an information unit transfer agreement is in effect SCSI target ports shall not transmit a MODIFY DATA POINTER message.

**Table 64 - 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) _____  _____ ARGUMENT _____  _____ (LSB)							
4								
5								
6								

### 16.3.10 MODIFY BIDIRECTIONAL DATA POINTER

The MODIFY BIDIRECTIONAL DATA POINTER message (see table 65) is sent from the SCSI target port to the SCSI initiator port and requests that the signed DATA-OUT ARGUMENT be added using two's complement arithmetic to the value of the current data-out pointer and signed DATA-IN ARGUMENT be added using two's complement arithmetic to the value of the current data-in pointer. The enable modify data pointer (EMDP) bit in the Disconnect-Reconnect mode page (see 18.1.2) indicates whether or not the SCSI target port is permitted to issue the MODIFY BIDIRECTIONAL DATA POINTER message. The SCSI target port shall only issue the MODIFY BIDIRECTIONAL DATA POINTER message during a bidirectional command.

It is recommended that the SCSI target port not attempt to move the data-out pointer or the data-in pointer outside the range addressed by the command. SCSI initiator ports may or may not place further restrictions on the acceptable values. Should the SCSI target port send a DATA-OUT ARGUMENT or a DATA-IN ARGUMENT value that is not supported by the SCSI initiator port, the SCSI initiator port may reject the value by responding with the MESSAGE REJECT message. In this case, both the data-out pointer and the data-in pointer are not changed from their values prior to the rejected MODIFY BIDIRECTIONAL DATA POINTER message.

If an information unit transfer agreement is in effect SCSI target ports shall not transmit a MODIFY BIDIRECTIONAL DATA POINTER message.

**Table 65 - MODIFY BIDIRECTIONAL DATA POINTER message format**

Bit Byte	7	6	5	4	3	2	1	0
0	EXTENDED MESSAGE (01h)							
1	EXTENDED MESSAGE LENGTH (09h)							
2	MODIFY BIDIRECTIONAL DATA POINTER (05h)							
3	(MSB)	DATA-OUT ARGUMENT						
4								
5								
6								(LSB)
7	(MSB)	DATA-IN ARGUMENT						
8								
9								
10								(LSB)

### 16.3.11 NO OPERATION

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

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

### 16.3.12 PARALLEL PROTOCOL REQUEST

PARALLEL PROTOCOL REQUEST messages (see table 66) are used to negotiate a synchronous transfer agreement, a wide data transfer agreement, and set the protocol options between two SCSI devices.

**Table 66 - PARALLEL PROTOCOL message format**

Bit Byte	7	6	5	4	3	2	1	0
0	EXTENDED MESSAGE (01h)							
1	EXTENDED MESSAGE LENGTH (06h)							
2	PARALLEL PROTOCOL REQUEST (04h)							
3	TRANSFER PERIOD FACTOR							
4	RESERVED							
5	REQ/ACK OFFSET							
6	TRANSFER WIDTH EXPONENT							
7	PCOMP_EN	RTI	RD_STRM	WR_FLOW	HOLD_MCS	QAS_REQ	DT_REQ	IU_REQ

PPR messages shall be supported by ports supporting transfer period factors less than 0Ah or supporting any of the protocol options. PPR messages shall be supported by SCSI target ports with a CLOCKING field indicating DT support, IUS set to one, or QAS set to one in the INQUIRY page of all their logical units (see SCSI Primary Commands-3 standard).

If information units are enabled and a selection with attention occurs with a PPR message the SCSI target port shall create a bus free condition. (see 10.3)

Usage of this message is defined in 4.12. Fields are defined in 4.12.4.

### 16.3.13 QAS REQUEST

The QAS REQUEST message is sent from a SCSI target port that has both information unit transfers and

QAS enabled to begin a QAS phase after a DT DATA phase (see 10.4.3).

#### 16.3.14 RESTORE POINTERS

The RESTORE POINTERS message is sent from a SCSI target port to direct the SCSI initiator port 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 pointers shall be restored to either the values at the beginning of the data areas in the absence of a SAVE DATA POINTERS message or to the values at the point at which the last SAVE DATA POINTERS message occurred for that task.

If an information unit transfer agreement is in effect SCSI target ports shall not transmit a RESTORE POINTERS message.

#### 16.3.15 SAVE DATA POINTERS

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

#### 16.3.16 SYNCHRONOUS DATA TRANSFER REQUEST

SYNCHRONOUS DATA TRANSFER REQUEST (SDTR) messages (see table 67) are used to negotiate a synchronous transfer agreement between two SCSI devices.

**Table 67 - 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							

SDTR messages shall be supported by devices supporting synchronous data transfers (i.e., non-zero REQ/ACK offsets). SDTR messages shall be supported by SCSI target ports with SYNC set to one in the INQUIRY page of all their logical units (see SCSI Primary Commands-3 standard).

Usage of this message is defined in 4.12. Fields are defined in 4.12.4.

#### 16.3.17 TASK COMPLETE

The TASK COMPLETE message is sent from a SCSI target device to a SCSI initiator device to indicate that a task has completed and that valid status has been sent to the SCSI initiator port when information unit transfers are disabled.

After successfully sending this message the SCSI target port shall go to the BUS FREE phase by releasing

the BSY signal. The SCSI target port shall consider the message transmission to be successful when there is no attention condition on the TASK COMPLETE message.

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

### 16.3.18 WIDE DATA TRANSFER REQUEST

WIDE DATA TRANSFER REQUEST (WDTR) messages (see table 68) are used to negotiate a wide data transfer agreement between two SCSI devices.

**Table 68 - 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							

WDTR messages shall be supported by ports supporting wide data transfers (i.e., non-zero transfer width exponents). WDTR messages shall be supported by SCSI target ports with WBUS16 set to one in the INQUIRY page of all their logical unit (see SCSI Primary Commands-3 standard).

Usage of this message is defined in 4.12. Fields are defined in 4.12.4.

## 16.4 Task attribute messages

### 16.4.1 Task attribute message overview and codes

Two byte task attribute messages are used to specify an identifier, called a tag, for a task that establishes the I\_T\_L\_Q nexus. The TAG field is an 8-bit integer assigned by the application client and sent to the SCSI initiator port in the send SCSI command request (see 19.3.2). The tag for every task for each I\_T\_L 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 nexus. If the task manager checks the tag value and receives a tag that is currently in use for the I\_T\_L nexus, then it shall abort all tasks for the SCSI initiator device and the associated logical unit and shall return CHECK CONDITION status for the task that caused the overlapped tag. The sense key shall be set to ABORTED COMMAND and the additional sense code shall be set to OVERLAPPED COMMANDS ATTEMPTED (see 17.3). 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 shall not affect the order of execution.

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

Whenever a SCSI initiator port does a physical connection to a SCSI target port, the appropriate task attribute message shall be sent 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 physical reconnection. If

a task attribute message is not sent, then only an I\_T\_L nexus is established for the task (i.e., an untagged command).

Whenever a SCSI target port does a physical reconnection to a SCSI initiator port to continue a tagged task, the SIMPLE QUEUE message shall be sent 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 physical reconnection. If the SIMPLE TAG message is not sent, then only an I\_T\_L nexus occurs for the task (i.e., an untagged command).

If a SCSI target port attempts to do a physical reconnection using a tag not associated with a task in the task set, then the SCSI initiator port should create an attention condition. After the corresponding MESSAGE OUT phase the SCSI initiator port shall respond with an ABORT TASK message.

If a SCSI target device does not implement tagged queuing and a queue tag message is received the SCSI target port 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 SCSI initiator port.

See SCSI Architecture Model-2 standard for the task set management rules.

**Table 69 - Task attribute message codes**

Code	Support				Message Name	Direction		Clear Attention Condition <sup>a</sup>
	IU Transfers Disabled		IU Transfers Enabled					
	Init	Targ	Init	Targ				
24h	O	O	NS	NS	ACA <sup>b</sup>		Out	Not required
21h	Q	Q	NS	NS	HEAD OF QUEUE <sup>b</sup>		Out	Not required
22h	Q	Q	NS	NS	ORDERED <sup>b</sup>		Out	Not required
20h	Q	Q	NS	NS	SIMPLE <sup>b</sup>	In	Out	Not required
Key: O=Optional support Q=Mandatory if tagged queuing is implemented In=SCSI target port to SCSI initiator port, Out=SCSI initiator port to SCSI target port Not required=SCSI initiator port may or may not clear the attention condition before last ACK of the MESSAGE OUT phase (see 12.2). NS=Not supported as a message. The receiving SCSI device shall reject this message. Init=SCSI initiator port, Targ=SCSI target port								
<sup>a</sup> The clear attention condition only applies during MESSAGE OUT phase.								
<sup>b</sup> When IU transfers are enabled, the equivalent task attributes are implemented in the SPI command information unit (see 14.3.1).								

#### 16.4.2 ACA

See table 70 for the format of the ACA message.

**Table 70 - 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 Architecture Model-2 standard.

#### 16.4.3 HEAD OF QUEUE

See table 71 for the format of the HEAD OF QUEUE message.

**Table 71 - 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 Architecture Model-2 standard.

#### 16.4.4 ORDERED

See table 72 for the format of the ORDERED message.

**Table 72 - 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 Architecture Model-2 standard.

### 16.4.5 SIMPLE

See table 73 for the format of the SIMPLE message.

**Table 73 - 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 Architecture Model-2 standard.



## 16.5 Task management messages

### 16.5.1 Task management message codes

Table 74 - Task management message codes

Code	Support				Message Name	Direction		Clear Attention Condition
	IU Transfers Disabled		IU Transfers Enabled					
	Init	Targ	Init	Targ				
0Dh	Q	Q	M	M	ABORT TASK <sup>a</sup>		Out	Yes
06h	O	M	NS	NS	ABORT TASK SET <sup>a</sup>		Out	Yes
16h	O	O	NS	NS	CLEAR ACA <sup>a</sup>		Out	Not required
0Eh	Q	Q	NS	NS	CLEAR TASK SET <sup>a</sup>		Out	Yes
17h	M	O	NS	NS	LOGICAL UNIT RESET <sup>a</sup>		Out	Yes
0Ch	O	M	O	M	TARGET RESET <sup>a</sup>		Out	Yes
Key: M=Mandatory support, O=Optional support Q=Mandatory if tagged queuing is implemented Out=SCSI initiator port to SCSI target port Yes=SCSI initiator port shall clear the attention condition before last ACK of the MESSAGE OUT phase. Not required=SCSI initiator port may or may not clear the attention condition before last ACK of the MESSAGE OUT phase (see 12.2). Init=SCSI initiator port, Targ=SCSI target port NS=Not supported as a message. The receiving SCSI device shall reject this message.								
<sup>a</sup> When IU transfers are enabled, the equivalent task management functions are implemented in the SPI command information unit (see 14.3.1).								

### 16.5.2 ABORT TASK

The ABORT TASK message requests the ABORT TASK task management function defined in the SCSI Architecture Model-2 standard.

In addition to the requirements in the SCSI Architecture Model-2 standard the SCSI target port 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 SCSI target port shall go to the BUS FREE phase. No status or message shall be sent for the current task and no pending data, status, or tasks are affected.

NOTE 39 - The ABORT TASK message in the case of only an I\_T nexus is useful to a SCSI initiator port that is not able to get an IDENTIFY message through to the SCSI target port 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 physical reconnection because of item f) in 12.2.

On a physical reconnection, the ABORT TASK message causes the current task to be aborted if an

I\_T\_L\_Q nexus has been fully identified. If an I\_T\_L nexus exists, but the SCSI target port is doing a physical reconnecting for an I\_T\_L\_Q nexus, then the current task is not aborted and the SCSI target port goes to the BUS FREE phase.

NOTE 40 - A nexus may not be fully identified on a physical reconnection if an attention condition is created during the IDENTIFY message and the SCSI target device has any tagged tasks for that SCSI initiator device on that logical unit.

It is not an error to issue this message to an I\_T\_L or I\_T\_L\_Q nexus that does not have any pending tasks.

### **16.5.3 ABORT TASK SET**

The ABORT TASK SET message requests the ABORT TASK SET task management function defined in the SCSI Architecture Model-2 standard.

In addition to the requirements in the SCSI Architecture Model-2 standard the SCSI target port 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 SCSI target port shall switch to a BUS FREE phase. No status or message shall be sent for the current task and no pending data, status, or tasks are affected.

The ABORT TASK SET message in the case of only an I\_T nexus is useful to a SCSI initiator port that is not able to send an IDENTIFY message through to the SCSI target port due to parity errors and just needs to end the current task or task management function.

It is not an error to issue this message to an I\_T\_L nexus that does not have any pending or current tasks.

### **16.5.4 CLEAR ACA**

The CLEAR ACA message requests the CLEAR ACA task management function defined in the SCSI Architecture Model-2 standard.

The CLEAR ACA message shall only be sent by a SCSI initiator port during an initial connection. If the SCSI target port receives the CLEAR ACA message at any other time the SCSI target port shall switch to a MESSAGE IN phase and issue a MESSAGE REJECT message. The SCSI target device shall then continue processing the task that was in process when the CLEAR ACA message was received.

On receipt of a CLEAR ACA message the task manager, in addition to clearing the ACA condition, shall go to the BUS FREE phase following the successful receipt of the CLEAR ACA message.

It is not an error to issue a CLEAR ACA message when no ACA condition is in effect.

### **16.5.5 CLEAR TASK SET**

The CLEAR TASK SET message requests the CLEAR TASK SET task management function defined in the SCSI Architecture Model-2 standard.

In addition to the requirements in the SCSI Architecture Model-2 standard the SCSI target port shall go to the BUS FREE phase following the successful receipt of the CLEAR TASK SET message.

### **16.5.6 LOGICAL UNIT RESET**

The LOGICAL UNIT RESET message requests the LOGICAL UNIT RESET task management function defined in the SCSI Architecture Model-2 standard.

Receipt of the LOGICAL UNIT RESET message after an I\_T\_L nexus has been established is a logical

unit reset event as defined in the SCSI Architecture Model-2 standard.

In addition to the requirements in the SCSI Architecture Model-2 standard the SCSI target port shall go to the BUS FREE phase following the successful receipt of the LOGICAL UNIT RESET message.

A logical unit reset has no effect on the transfer agreement.

#### **16.5.7 TARGET RESET**

The TARGET RESET message requests the TARGET RESET task management function defined in the SCSI Architecture Model-2 standard.

Successful receipt of a TARGET RESET message is a target reset event (see 12.5.4).

In addition to the requirements in the SCSI Architecture Model-2 standard the SCSI target port shall go to the BUS FREE phase following the successful receipt of the TARGET RESET message.

## 17 Command processing considerations and exception conditions

### 17.1 Command processing considerations and exception conditions overview

This clause describes some aspects of command processing, including exception conditions and error handling that are specific to this standard.

### 17.2 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 Primary Commands-3 standard).

An error condition or unit attention condition shall be reported once per occurrence of the event causing it. The SCSI target device 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 SCSI initiator port 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 SCSI target device 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 SCSI device requires or when an event occurs that may invalidate the current information.

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. SCSI devices that 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 SCSI device conducts a survey with INQUIRY commands to be sure that the SCSI devices on the SCSI bus are appropriate destinations for SEND commands with an AER bit of one. The SCSI devices on the bus or the SCSI ID assignments may have changed.

See asynchronous event reporting in the SCSI Architecture Model-2 standard for more information on asynchronous event notification.

### 17.3 Incorrect initiator connection

An incorrect initiator connection occurs during an initial connection if a SCSI initiator port creates a nexus that already exists and does not send an ABORT TASK SET, ABORT TASK, CLEAR TASK SET, DISCONNECT, LOGICAL UNIT RESET, or TARGET RESET, message as one of the messages of the MESSAGE OUT phase or as one of the task management functions in the SPI command information unit.

A task manager that detects an incorrect initiator connection shall abort all tasks for the SCSI initiator port 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 (see 16.4).

NOTE 42 - An incorrect initiator connection may be indicative of a serious error and, if not detected, may result in a task operating with a wrong set of pointers. This is considered a catastrophic failure on the part of the SCSI initiator port. Therefore, vendor-specific error recovery procedures may be required to guarantee the data integrity on the medium. The SCSI target device 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).

## **17.4 Unexpected RESELECTION phase**

An unexpected RESELECTION phase occurs if a SCSI target port attempts to do a physical reconnect to a task for which a nexus does not exist. A SCSI initiator port should respond to an unexpected RESELECTION phase by sending an ABORT TASK message.

## 18 SCSI management features for the SCSI parallel interface

### 18.1 SCSI mode parameters

#### 18.1.1 SCSI mode parameter overview and codes

This subclause describes the block descriptors and the mode pages used with MODE SELECT and MODE SENSE commands that influence, control and report the behavior of the SCSI parallel interface. All mode parameters not defined in this standard shall influence the behavior of the SCSI devices as specified in the appropriate command set document. The mode pages are addressed to the device server of a logical unit. The mode pages and mode subpages associated with the SCSI parallel interface are listed in table 75 and table 76.

**Table 75 - Mode page codes for the SCSI parallel interface**

Page code	Description	Reference
02h	Disconnect-Reconnect mode page	18.1.2
18h	Logical Unit Control mode page	18.1.3
19h	Port Control mode page <sup>a</sup>	18.1.4
<sup>a</sup> The short format for the Port Control mode page is selected when the SUBPAGE CODE field is set to 00h in a MODE SELECT command.		

**Table 76 - Mode subpage codes for the SCSI parallel interface**

Page code	Subpage code <sup>a</sup>	Description	Reference
19h	00h	Not allowed	
19h	01h	Margin Control mode subpage	18.1.4.2
19h	02h	Saved Training Configuration value mode subpage	18.1.4.3
19h	03h	Negotiated Settings mode subpage	18.1.4.4
19h	04h	Report Transfer Capabilities mode subpage	18.1.4.5
19h	05h - DFh	Reserved	
19h	E0h - FEh	Vendor specific	
19h	FFh	Return all mode subpages for the Port Control mode page	SPC-3
<sup>a</sup> The SUBPAGE CODE field in a mode page with SPF set to one.			

### 18.1.2 Disconnect-Reconnect mode page

The Disconnect-Reconnect mode page (see table 77) 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 SCSI target devices.

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 SCSI target port. The field values are communicated from the device server to the SCSI target port in a vendor specific manner.

SPI SCSI devices shall only use Disconnect-Reconnect mode page parameter fields defined in this subclause. If any other fields within the Disconnect-Reconnect mode 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.

Table 77 - Disconnect-Reconnect mode page

Bit Byte	7	6	5	4	3	2	1	0
0	PS	SPF (0)	PAGE CODE (02h)					
1	PAGE LENGTH (0Eh)							
2	BUFFER FULL RATIO							
3	BUFFER EMPTY RATIO							
4	(MSB)	BUS INACTIVITY LIMIT						
5								(LSB)
6	(MSB)	PHYSICAL DISCONNECT TIME LIMIT						
7								(LSB)
8	(MSB)	CONNECT TIME LIMIT						
9								(LSB)
10	(MSB)	MAXIMUM BURST SIZE						
11								(LSB)
12	EMDP	FAIR ARBITRATION			DIMM	DTDC		
13	RESERVED							
14	RESERVED							
15								

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  $\mu$ s increments that the SCSI target port is permitted to assert the BSY signal without a REQ/ACK handshake. If the bus inactivity limit is exceeded the SCSI target port shall attempt to do a physical disconnect (see 16.3.2) if the SCSI initiator port has granted the physical disconnect privilege (see 16.3.3) and it is not restricted by DTDC. This value may be rounded as defined in the SCSI Primary Commands-3 standard. A value of zero indicates that there is no bus inactivity limit.



The PHYSICAL DISCONNECT TIME LIMIT field indicates the minimum time in 100  $\mu$ s increments that the SCSI target port shall wait after releasing the SCSI bus before attempting a physical reconnection. This value may be rounded as defined in the SCSI Primary Commands-3 standard. A value of zero indicates that there is no physical disconnect time limit.

The CONNECT TIME LIMIT field indicates the maximum time in 100  $\mu$ s increments that the SCSI target port is allowed to use the SCSI bus before doing a physical disconnect, if the SCSI initiator port has granted the physical disconnect privilege (see 16.3.3) and it is not restricted by DTDC. This value may be rounded as defined in the SCSI Primary Commands-3 standard. A value of zero indicates that there is no connect time limit.

If information unit transfers are disabled the MAXIMUM BURST SIZE field indicates the maximum amount of data that the SCSI target port shall transfer during a DATA phase before doing a physical disconnect if the SCSI initiator port has granted the physical disconnect privilege (see 16.3.3).

If information unit transfer are enabled the MAXIMUM BURST SIZE field indicates the maximum amount of data that the SCSI target port shall transfer in a single SPI data information unit.

The maximum burst size is expressed in increments of 512 bytes (e.g., a value of one means 512 bytes, two means 1 024 bytes, etc.). A value of zero indicates there is no limit on the amount of data transferred per burst.

The enable modify data pointer (EMDP) bit indicates whether or not the SCSI initiator port allows the MODIFY DATA POINTER and MODIFY BIDIRECTIONAL DATA POINTERS messages to be issued by the SCSI target port. If the EMDP bit is a zero or an information unit transfer agreement is in effect, the SCSI target port shall not issue the MODIFY DATA POINTER or MODIFY BIDIRECTIONAL DATA POINTERS messages. If the EMDP bit is a one and an information unit transfer agreement is not in effect, the SCSI target port is allowed to issue MODIFY DATA POINTER and MODIFY BIDIRECTIONAL DATA POINTERS messages.

If the MODIFY DATA POINTER and MODIFY BIDIRECTIONAL DATA POINTERS messages are allowed and the SCSI initiator port responds to a MODIFY DATA POINTER or MODIFY BIDIRECTIONAL DATA POINTERS message with a MESSAGE REJECT, then the SCSI target device 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.

If the FAIR ARBITRATION field is set to 000b, the SCSI target port shall not use arbitration fairness during normal arbitration. If this field is set to a nonzero value, the SCSI target port shall use arbitration fairness during normal arbitration (see Annex B).

Regardless of the value in the FAIR ARBITRATION field the SCSI target port shall use arbitration fairness during QAS.

A disconnect immediate (DIMM) bit of zero indicates that the SCSI target port may request DATA IN or DATA OUT phases following a COMMAND phase without attempting a physical disconnect (see 16.3.2). A DIMM bit of one indicates that the SCSI target port shall attempt a physical disconnect (see 16.3.2) after a COMMAND phase and before a subsequent DATA IN or DATA OUT phase. The DIMM bit only applies when the SCSI initiator port has granted the physical disconnect privilege (see 16.3.3).

The data transfer disconnect control (DTDC) field (see table 78) defines further restrictions on when a physical disconnect is permitted.

**Table 78 - DTDC**

<b>DTDC</b>	<b>Description</b>
000b	DTDC is not used. Physical disconnect is controlled by the other fields in this mode page.
001b	A SCSI target port shall not attempt to do a physical 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 SCSI target port shall not attempt to do a physical 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 SCSI target device 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.

### **18.1.3 Logical Unit Control mode page**

The Logical Unit Control mode page (see table 79) contains those parameters that select logical unit operation options. This mode page is not currently defined for SPI SCSI devices. The implementation of any parameter and its associated functions is optional. The mode page follows the MODE SENSE / MODE SELECT rules specified by the SCSI Primary Commands-3 standard.

**Table 79 - Logical Unit Control mode page**

Bit Byte	7	6	5	4	3	2	1	0
0	PS	SPF (0)	PAGE CODE (18h)					
1	PAGE LENGTH (06h)							
2	RESERVED				PROTOCOL IDENTIFIER (1h)			
3	RESERVED							
4	RESERVED							
5	RESERVED							
6	RESERVED							
7	RESERVED							

The PROTOCOL IDENTIFIER field of 1h indicates the protocol that this mode page applies to a SPI SCSI device. See the SCSI Primary Commands-3 standard for other Port Control mode page protocol identifiers.

#### 18.1.4 Port Control mode page

##### 18.1.4.1 Port Control mode page overview

The Port Control mode page (see table 80 and table 81) contains those parameters that affect SCSI target port operation options. The mode page shall be implemented by LUN 0 of all SPI SCSI target devices. The mode page shall not be implemented by logical units other than LUN 0 and W-LUNs. The implementation of any bit and its associated functions is optional. The mode page follows the MODE SENSE / MODE SELECT rules specified by SCSI Primary Commands-3 standard.

Each SCSI target port shall maintain an independent set of Port Control mode page parameters for each SCSI initiator port. The parameters saveable bit in the mode page format header returned with MODE SENSE command shall be set to zero if the mode subpage format is being used (i.e., SPF bit set to one), indicating the parameters are not saved through resets.

After a MODE SELECT command, parameter settings shall remain in effect until either:

- settings are changed by another MODE SELECT command,
- a logical unit reset of LUN 0 occurs,
- an SDTR negotiation successfully completes,
- a WDTR negotiation successfully completes, or
- a PPR negotiation successfully completes with the HOLD\_MCS bit set to zero.

**Table 80 - Port Control mode page short format**

Bit Byte	7	6	5	4	3	2	1	0
0	PS	SPF (0)	PAGE CODE (19h)					
1	PAGE LENGTH (06h)							
2	RESERVED				PROTOCOL IDENTIFIER (1h)			
3	RESERVED							
4	(MSB)							
5	SYNCHRONOUS TRANSFER TIMEOUT (LSB)							
6	RESERVED							
7	RESERVED							

The PROTOCOL IDENTIFIER field of 1h indicates the protocol that this mode page applies to a SPI SCSI device. See the SCSI Primary Commands-3 standard for other Port Control mode page protocol identifiers.

The SYNCHRONOUS TRANSFER TIMEOUT field indicates the maximum amount of time in 1 ms increments that the SCSI target port shall wait before generating an error by doing an unexpected bus free (see 10.3). The SCSI target port shall only go to a BUS FREE phase if one of the following events causes the timer, once started, to not reset or reload before expiring.

- If there is a REQ transition when there are no outstanding REQs waiting for an ACK then load and start the timer.
- If there is a REQ transition when there are any outstanding REQs waiting for an ACK then there is no effect on the timer.
- If there is an ACK transition when there are outstanding REQs waiting for an ACK then load and start the timer.
- If after an ACK transition there are no outstanding REQs waiting for an ACK then stop the timer.

A SYNCHRONOUS TRANSFER TIMEOUT field value of 0000h indicates that the function is disabled. A value of FFFFh indicates an unlimited period.

**Table 81 - Port Control mode subpage format**

Bit Byte	7	6	5	4	3	2	1	0
0	PS	SPF (1)	PAGE CODE (19h)					
1	SUBPAGE CODE							
2	(MSB)							
3	PAGE LENGTH (n-3)							
	(LSB)							
4	RESERVED							
5	RESERVED				PROTOCOL IDENTIFIER (1h)			
6	PROTOCOL SPECIFIC MODE PARAMETERS							
n								

The SUBPAGE CODE field indicates which mode subpage is being accessed. Mode subpage code values are listed in table 76. If the parameter data of a MODE SELECT command contains a mode subpage formatted mode page with the SUBPAGE CODE field is set to zero the SCSI target device 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.

The PAGE LENGTH field specifies the length in bytes of the mode subpage protocol specific mode page parameters after the PAGE LENGTH field.

The PROTOCOL IDENTIFIER field of 1h indicates the protocol that this mode page applies to a SPI SCSI device port. See the SCSI Primary Commands-3 standard for other Port Control mode page protocol identifiers.

#### 18.1.4.2 Margin Control mode subpage

The Margin Control mode subpage (see table 82) contains parameters that set and report margin control values for usage between the SCSI initiator/target port pair on subsequent synchronous and paced transfers.

A MODE SENSE command shall return the current settings for the SCSI initiator/target port pair. Fields that are not implemented shall be reported as zero.

**Table 82 - Margin Control mode subpage**

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED							
1	DRIVER STRENGTH				RESERVED			
2	DRIVER ASYMMETRY				DRIVER PRECOMPENSATION			
3	DRIVER SLEW RATE				RESERVED			
4	RESERVED							
5	RESERVED							
6	RESERVED							
7	VENDOR SPECIFIC							
8	RESERVED							
9	RESERVED							
10	RESERVED							
11	RESERVED							
12	RESERVED							
13	RESERVED							
14	RESERVED							
15	RESERVED							

The DRIVER STRENGTH field indicates the relative amount of driver source current used by the driver (see 7.2.2 and A.2.2). The DRIVER STRENGTH field affects both the strong and weak drivers. A larger value indicates more driver source current.

The DRIVER PRECOMPENSATION field indicates the relative difference between the weak driver and the strong driver amplitudes when precompensation is enabled (see A.2.1). A larger value indicates a larger difference between the weak and strong amplitudes.

The DRIVER ASYMMETRY field indicates the relative difference between the amplitudes of asserted and negated signals launched from the driver (see A.2 and figure A.2). A larger value indicates a relatively

stronger asserted signal compared to the negated signal.

The DRIVER SLEW RATE field indicates the relative difference between the assertion and negation magnitudes divided by the rise or fall time (see figure A.2.6 and figure A.2.7). A larger value indicates a faster slew rate.

The default value of each margin control field should be 0000b.

The margin control fields indicate absolute conditions centered around their default values. Absolute conditions means that the previous history of the parameter has no relevance to the value of the parameter. The fields are two's complement values as shown in table 83. The maximum supported setting for each field is 0111b and the minimum supported setting for each field is 1000b. Up to 16 distinct values are available for each field, representing monotonically changing device response. Devices that support fewer than 16 distinct values for a field should round non-supported settings to a supported value.

The actual response of a SCSI device to a field value is vendor specific and calibration of the actual minimum and maximum responses to different field values is not defined in this standard. Margin control settings should not cause the driver to violate this standards' electrical limits. Margin control settings should affect only the REQUEST, ACKNOWLEDGE, DATA BUS, P\_CRCA, and DB(P1) signals and should affect all of these signals driven by the device by the same amount.

**Table 83 - Summary of margin control field values**

Value		Parameter values
Binary	Decimal	
0111b	7	maximum setting
0110b	6	
0101b	5	
0100b	4	
0011b	3	
0010b	2	
0001b	1	
0000b	0	recommended default value
1111b	-1	
1110b	-2	
1101b	-3	
1100b	-4	
1011b	-5	
1010b	-6	
1001b	-7	
1000b	-8	minimum setting

#### 18.1.4.3 Saved Training Configuration Values mode subpage

The Saved Training Configuration Values mode subpage is used to report the SCSI device's saved training configuration values. These vendor specific values are maintained by the SCSI device when the retain training information option is enabled (see 4.12.4.6.8). The fields are listed in table 84 however the content of the fields is vendor specific.

Only values for the current I\_T nexus are reported.

**Table 84 - Saved Training Configuration mode subpage**

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED							
3								
4	(MSB)	DB(0) VALUE						(LSB)
7								
		.						
		.						
		.						
64	(MSB)	DB(15) VALUE						(LSB)
67								
68	(MSB)	P_CRCA VALUE						(LSB)
71								
72	(MSB)	P1 VALUE						(LSB)
75								
76	(MSB)	BSY VALUE						(LSB)
79								
80	(MSB)	SEL VALUE						(LSB)
83								
84	(MSB)	RST VALUE						(LSB)
87								
88	(MSB)	REQ VALUE						(LSB)
91								
92	(MSB)	ACK VALUE						(LSB)
95								
96	(MSB)	ATN VALUE						(LSB)
99								
100	(MSB)	C/D VALUE						(LSB)
103								
104	(MSB)	I/O VALUE						(LSB)
107								
108	(MSB)	MSG VALUE						(LSB)
111								
112		RESERVED						
227								



#### 18.1.4.4 Negotiated Settings mode subpage

The Negotiated Settings mode subpage, shown in table 85, is used to report the negotiated settings of a SCSI target port for the current I\_T nexus.

**Table 85 - Negotiated Settings mode subpage**

Bit Byte	7	6	5	4	3	2	1	0
0	TRANSFER PERIOD FACTOR							
1	RESERVED							
2	REQ/ACK OFFSET							
3	TRANSFER WIDTH EXPONENT							
4	RESERVED	PROTOCOL OPTIONS BITS						
5	RESERVED				TRANSCIEVER MODE		SENT PCOMP_EN	RECEIVED PCOMP_EN
6	RESERVED							
7	RESERVED							

The TRANSFER PERIOD FACTOR field indicates the negotiated transfer period factor (see 4.12.4.3) for the current I\_T nexus.

The REQ/ACK OFFSET field indicates the negotiated REQ/ACK offset (see 4.12.4.4) for the current I\_T nexus.

The TRANSFER WIDTH EXPONENT field indicates the negotiated transfer width exponent (see 4.12.4.5) for the current I\_T nexus.

The PROTOCOL OPTIONS BITS field contain the negotiated protocol options (see 4.12.4.6) for the current I\_T nexus.

The RECEIVED PCOMP\_EN bit contains the value of the PCOMP\_EN bit (see 4.12.4.6.9) received by the SCSI target port for the current I\_T nexus.

The SENT PCOMP\_EN bit contains the value of the PCOMP\_EN bit (see 4.12.4.6.9) sent by the SCSI target port for the current I\_T nexus.

The TRANSCIEVER MODE field specifies the current bus mode of the SCSI target port as defined in table 86.

**Table 86 - Bus mode**

<b>Code</b>	<b>Bus mode</b>
00b	Unknown (e.g., device not capable of reporting bus mode)
01b	Single ended
10b	Low Voltage Differential
11b	High Voltage Differential

**18.1.4.5 Report Transfer Capabilities mode subpage**

The Report Transfer Capabilities mode subpage, shown in table 87, is used to report the transfer capabilities for the SCSI target port. The values in this mode subpage are not changeable via a MODE SELECT command.

**Table 87 - Report Transfer Capabilities mode subpage**

<b>Bit Byte</b>	<b>7</b>	<b>6</b>	<b>5</b>	<b>4</b>	<b>3</b>	<b>2</b>	<b>1</b>	<b>0</b>
<b>0</b>	MINIMUM TRANSFER PERIOD FACTOR							
<b>1</b>	RESERVED							
<b>2</b>	MAXIMUM REQ/ACK OFFSET							
<b>3</b>	MAXIMUM TRANSFER WIDTH EXPONENT							
<b>4</b>	PROTOCOL OPTIONS BITS SUPPORTED							
<b>5</b>	RESERVED							
<b>6</b>	RESERVED							
<b>7</b>	RESERVED							

The MINIMUM TRANSFER PERIOD FACTOR field shall be set to the smallest value of the transfer period factor (see 4.12.4.3) supported by the SCSI target port.

The MAXIMUM REQ/ACK OFFSET shall be set to the largest value of the REQ/ACK offset (see 4.12.4.4) supported by the SCSI target port.

The MAXIMUM TRANSFER WIDTH EXPONENT shall be set to the largest value of the transfer width exponent (see 4.12.4.5) supported by the SCSI target port.

The SCSI target port shall set the bits in the PROTOCOL OPTIONS BITS SUPPORTED field (see 4.12.4.6) to indicate the protocol options supported by the SCSI target port.

## 19 SCSI parallel interface services

### 19.1 SCSI parallel interface services overview

SCSI parallel interface services are provided by the SCSI initiator port enabling the application client to accomplish tasks and task management functions (see SCSI Architecture Model-2 standard) and by the SCSI target device 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 SCSI initiator port and SCSI target port provide. Each SCSI parallel interface service causes a sequence phases to be generated by the addressed SCSI devices. Figure 75, figure 76, and figure 77 show all the valid phase sequences.

### 19.2 Procedure objects

See table 88 for the mapping of the procedure objects used in this standard to the equivalent procedure objects used in the SCSI Architecture Model-2 standard.

**Table 88 - This standards objects mapped to objects from other SCSI standards**

<b>Objects</b>	<b>Equivalent SCSI Architecture Model-2 standard objects</b>
initiator SCSI ID	initiator identifier
target SCSI ID	target identifier

See table 89 for a list of the procedure objects used when passing services across the SCSI parallel interface service interface. See table 89 for the definitions of the names used within this standard and the equivalent SCSI Architecture Model-2 standard names of the procedure objects, the name of the standard where the objects are defined, the standard where the binary contents of the objects are defined, and the routing of the objects. 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.

Table 89 - Procedure objects

This standard's objects	Standard where term defined	Standard where binary contents of term defined	Term routing
application client buffer offset	SAM-2	SAM-2	DS → targ → init
data-out buffer size	SAM-2	SAM-2	AC → init
data-in buffer size	SAM-2	SAM-2	AC → init
command descriptor block	SAM-2	SAM-2/cmd <sup>b</sup>	AC → init → targ → DS
data-in buffer	SAM-2	cmd <sup>c</sup>	DS → targ → init → AC
data-out buffer	SAM-2	cmd <sup>c</sup>	AC → init → targ → DS
device server buffer	SAM-2	cmd <sup>c</sup>	DS → targ → init
I_T_L_x nexus	SAM-2	this standard	AC → init → targ → DS or AC → init → targ → TM or DS → targ → init
request byte count	SAM-2	SAM-2	DS → targ
service response	SAM-2	this standard <sup>d</sup>	DS → targ → init → AC or targ → DS
service response <sup>a</sup>	SAM-2	this standard <sup>d</sup>	init → AC
autosense request	SAM-2	SAM-2	AC → init → targ → DS
sense data	SAM-2	SPC-3	DS → targ → init → AC
status	SAM-2	SAM-2	DS → targ → init → AC
task attribute	SAM-2	this standard	AC → init → targ → DS
Key: AC=application client, cmd=SCSI command standards, DS=device server, init=SCSI initiator port, TM=task manager, targ=SCSI target port			
<sup>a</sup> Only occurs when unexpected bus free (see 10.3) is detected by the SCSI initiator port. <sup>b</sup> The portions not defined in the SCSI Architecture Model-2 standard are defined in the SCSI command standards (e.g., SCSI-3 Block Commands standard, SCSI Primary Commands-3 standard). <sup>c</sup> Parameter lists are defined within one of the SCSI command standards (e.g., SCSI-3 Block Commands standard, SCSI Primary Commands-3 standard). SCSI standards do not define non-parameter list information. <sup>d</sup> The SERVICE DELIVERY OR TARGET FAILURE value of the service response is not defined in SCSI.			

## 19.3 Application client SCSI command services

### 19.3.1 Application client SCSI command services overview

The SCSI command services shall be requested by the application client using a procedure call defined as:

Execute Command (IN (I\_T\_L\_x nexus, command descriptor block, [task attribute], [data-in buffer size], [data-out buffer], [data-out buffer size], [autosense request]), OUT ([data-in buffer], [sense data], status, service response))

### 19.3.2 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 90.

**Table 90 - Processing of send SCSI command service procedure**

Step	Protocol service name	SCSI Transport protocol Service Interface procedure call
request	send SCSI command request	Send SCSI command (IN (I_T_L_x nexus, command descriptor block, [task attribute], [data-in buffer size], [data-out buffer], [data-out buffer size], [autosense request]))
indication	send SCSI command indication	SCSI command received (IN (I_T_L_x nexus, command descriptor block, [task attribute], [autosense request]))
response	send SCSI command response	Send command complete (IN (I_T_L_x nexus, [sense data], status, service response))
confirmation	send SCSI command confirmation	Command complete received (IN (I_T_L_x nexus, [data-in buffer], [sense data], status, service response))

## 19.4 Device server SCSI command services

### 19.4.1 Device server SCSI command services overview

The SCSI data buffer movement services shall be requested from the device server using a procedure call defined as:

Move data buffer (IN (I\_T\_L\_x nexus, device server buffer, application client buffer offset, request byte count)).

Either data-in delivery, data-out delivery, both data-in and data-out delivery, or neither data delivery may be used while processing one command. If both are used, the device server shall combine the data-in and data-out service responses into one service response.

### 19.4.2 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 a SCSI initiator port.

Processing the execute command procedure call for a data-in delivery service shall be composed of the 2 step confirmed service shown in table 91.

**Table 91 - Processing of data-in delivery service procedure**

Step	Protocol service name	SCSI Transport Protocol Service Interface procedure call
request	data-in delivery request	Send data-in (IN (I_T_L_x nexus, device server buffer, application client buffer offset, request byte count))
confirmation	data-in delivery confirmation	Data-In delivered (IN (I_T_L_x nexus))

#### 19.4.3 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 a SCSI initiator port 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 92.

**Table 92 - Processing of data-out delivery service procedure**

Step	Protocol service name	SCSI Transport Protocol Service Interface procedure call
request	data-out delivery request	Receive data-out (IN (I_T_L_x nexus, application client buffer offset, request byte count, device server buffer))
confirmation	data-out delivery confirmation	Data-out received (IN (I_T_L_x nexus))

### 19.5 Task management services

#### 19.5.1 Task management functions overview

The task management services shall be requested from the application client using a procedure call defined as:

Function name (IN (nexus), service response)

#### 19.5.2 Task management functions

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 Architecture Model-2 standard. This standard defines the actions taken by the SCSI parallel interface service to carry out the requested task management functions.

### **19.5.3 ABORT TASK**

The SCSI parallel interface services request the SCSI initiator port issue an ABORT TASK message (see 16.5.2 and 14.3.1) to the selected I\_T\_L\_Q nexus.

### **19.5.4 ABORT TASK SET**

The SCSI parallel interface services request the SCSI initiator port issue an ABORT TASK SET message (see 16.5.3 and 14.3.1) to the selected I\_T\_L nexus.

### **19.5.5 CLEAR ACA**

The SCSI parallel interface services request the SCSI initiator port issue a CLEAR ACA message (see 16.5.4 and 14.3.1) to the selected I\_T\_L nexus.

### **19.5.6 CLEAR TASK SET**

The SCSI parallel interface services request the SCSI initiator port issue a CLEAR TASK SET message (see 16.5.5 and 14.3.1) to the selected I\_T\_L nexus.

### **19.5.7 LOGICAL UNIT RESET**

The SCSI parallel interface services request the SCSI initiator port issue a LOGICAL UNIT RESET message (see 16.5.6 and 14.3.1) to the selected I\_T\_L nexus.

### **19.5.8 RESET SERVICE DELIVERY SUBSYSTEM**

The SCSI parallel interface services request the SCSI initiator port create a bus reset condition (see 12.3) on the selected SCSI bus containing the selected I\_T nexus.

### **19.5.9 TARGET RESET**

The SCSI parallel interface services request the SCSI initiator port issue a TARGET RESET message (see 16.5.7 and 14.3.1) to the selected I\_T nexus.

### **19.5.10 WAKEUP**

The SCSI parallel interface services request the SCSI initiator port create a bus reset condition (see 12.3) on the selected SCSI bus containing the selected I\_T nexus.

## **Annex A**

(normative)

### **Additional requirements for LVD SCSI drivers and receivers**

#### **A.1 System level requirements**

The requirements for LVD SCSI 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 segment loading conditions and trade-offs between competing parameters.



Table A.1 - System level requirements

Parameter	Minimum	Maximum	Cross-reference
$V_A$ (synchronous transfers and paced transfers clocking signal) (except OR-tied signals) <sup>a, e</sup>	-1 V	-100 mV	N/A
$V_A$ (paced transfers noncompensated transition) (except OR-tied signals) <sup>f</sup>	-1 V	-5 mV	N/A
$V_A$ (paced transfers compensated transition) (except OR-tied signals) <sup>a</sup>	-1 V	-80 mV	N/A
$V_N$ (synchronous transfers and paced transfers clocking signal) (except OR-tied signals) <sup>a, e</sup>	100 mV	1 V	N/A
$V_N$ (paced transfers noncompensated transition) (except OR-tied signals) <sup>f</sup>	5 mV	1 V	N/A
$V_N$ (paced transfers compensated transition) (except OR-tied signals) <sup>a</sup>	80 mV	1 V	N/A
$V_A$ (OR-tied signals) <sup>a</sup>	-3,6 V	-100 mV	N/A
$V_N$ (OR-tied signals) <sup>a, b</sup>	80 mV	145 mV	N/A
loaded impedance ( $\Omega$ ) <sup>c</sup>	85	135	N/A
unloaded impedance ( $\Omega$ )	110	135	subclause 6.3
terminator bias (mV)	100	125	subclause 7.2.1
terminator impedance ( $\Omega$ )	100	110	subclause 7.2.1
device leakage ( $\mu A$ )	-20	20	table 22
number of SCSI devices	2	16	subclause 4.8
ground offset level (mV) <sup>d</sup>	-355	355	N/A
<sup>a</sup> These are the signal levels at the receiver, the system allows 60 mV crosstalk for calculating the minimum driver level. <sup>b</sup> Prior versions of the standard did not account for leakage. <sup>c</sup> Caused by the addition of device capacitive load (see table 16 for calculations). <sup>d</sup> This is the difference in voltage signal commons for SCSI devices on the bus segment (see figure 2). <sup>e</sup> Clocking signals: toggling signal after the first transition shall be the 100 mV limits specified by the fast-80 or slower speeds. <sup>f</sup> Weak driver is disabled.			

## A.2 Driver requirements

### A.2.1 Driver requirements overview

The fundamental requirement for an LVD driver is the generation of a first-step differential output voltage magnitude at the driver connections to the balanced load to achieve required minimum differential signals at every receiver connection to the bus segment. If precompensation is enabled, the weak driver amplitude shall be a minimum of 50 % to a maximum of 75 % of the strong driver amplitude after the first bit of a series of adjacent ones or adjacent zeros. 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.

NOTE 43 - If a weak driver is driving with the minimum amplitude specified in table A.2, then the 370 mV weak driver translates to a strong driver of 493 mV for the 75 % case ranging up to 740 mV for the 50 % case.

The driver requirements are defined in terms of the voltages and currents depicted in figure 53.

### A.2.2 Differential output voltage, $V_S$

This subclause does not specify requirements for drivers with source impedances less than 1 000  $\Omega$

To assure sufficient voltage to define a valid logic state at any device connection on a fully loaded LVD bus segment at least a minimum differential output voltage shall be generated. This value shall be large enough that, after allowance for attenuation, reflections, terminator bias difference, and differential noise coupling,  $V_S$  meets the specified requirements at the device connector to the LVD SCSI bus segment.

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.

For non-paced transfers, with the test circuit of figure A.1 and the test conditions V1 and V2 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 320 mV and less than or equal to 800 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 320 mV and less than or equal to 800 mV. The relationship between  $V_A$  and  $V_N$  specified in table A.2 and shown graphically in figure A.2 shall be maintained.

For paced transfers, with the test circuit of figure A.1 and the test conditions V1 and V2 in table A.3 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 370 mV and less than or equal to 800 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 370 mV and less than or equal to 800 mV. The relationship between  $V_A$  and  $V_N$  specified in table A.3 and shown graphically in figure A.3 shall be maintained. The strong driver relationship between  $V_A$  and  $V_N$  specified in table A.2 strong driver and shown graphically in figure A.3 for the strong 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 and figure A.3 because the applied V1 and V2 simulate the effects of the bus segment termination bias.

**Table A.2 - Driver steady-state test limits and conditions for non-paced transfers**

Test parameter	Test conditions <sup>a</sup> (figure A.1)	Minimum (mV) <sup>b</sup>	Maximum (mV)
V <sub>A</sub>   Differential output voltage magnitude (asserted)	V <sub>1</sub> = 1,056 V V <sub>2</sub> = 0,634 V	320	800
	V <sub>1</sub> = 1,866 V V <sub>2</sub> = 1,444 V	320	800
V <sub>N</sub>   Differential output voltage magnitude (negated)	V <sub>1</sub> = 1,056 V V <sub>2</sub> = 0,634 V	320	800
	V <sub>1</sub> = 1,866 V V <sub>2</sub> = 1,444 V	320	800
V <sub>A</sub>   Differential output voltage magnitude (asserted)	All four conditions shown in the previous rows of this table	0,69 x   V <sub>N</sub>   + 50	1,45 x   V <sub>N</sub>   - 65
The test limits shall be within the shaded area of figure A.2.			
<sup>a</sup> The test circuit (figure A.1) is approximately equivalent to two terminators creating the normal system bias. <sup>b</sup> Including the weak output.			

**Table A.3 - Driver steady-state test limits and conditions for paced transfers**

Test parameter	Test conditions <sup>a</sup> (figure A.1)	Minimum (mV) <sup>b</sup>	Maximum (mV)
V <sub>A</sub>   Differential output voltage magnitude (asserted)	V <sub>1</sub> = 1,056 V V <sub>2</sub> = 0,634 V	370	800
	V <sub>1</sub> = 1,866 V V <sub>2</sub> = 1,444 V	370	800
V <sub>N</sub>   Differential output voltage magnitude (negated)	V <sub>1</sub> = 1,056 V V <sub>2</sub> = 0,634 V	370	800
	V <sub>1</sub> = 1,866 V V <sub>2</sub> = 1,444 V	370	800
V <sub>A</sub>   Differential output voltage magnitude (asserted)	All four conditions shown in the previous rows of this table	0,90 x   V <sub>N</sub>   - 23	1,11 x   V <sub>N</sub>   + 26
The test limits shall be within the shaded area of figure A.3.			
<sup>a</sup> The test circuit (figure A.1) is approximately equivalent to two terminators creating the normal system bias. <sup>b</sup> Including the weak output.			

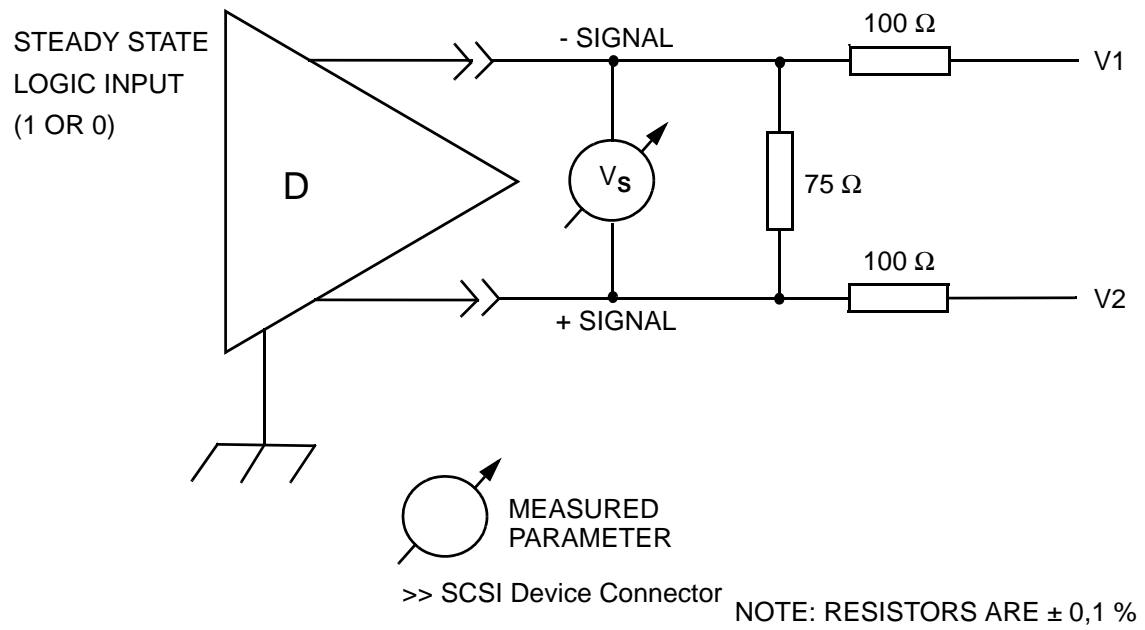


Figure A.1 - Differential steady-state output voltage test circuit





Figure A.2 - Domain for driver assertion and negation levels for non-paced transfers



**Figure A.3 - Domain for driver assertion and negation levels for paced transfers**

### **A.2.3 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.4 shall be greater than or equal to 0,845 V and less than or equal to 1,655 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,  $\Delta V_{CM}$ , shall be 120 mV or less for all  $V_{applied}$  in the range:  $0,845 \leq V_{applied} \leq 1,655$ . See figure A.5.




Figure A.4 - Driver offset steady-state voltage test circuit


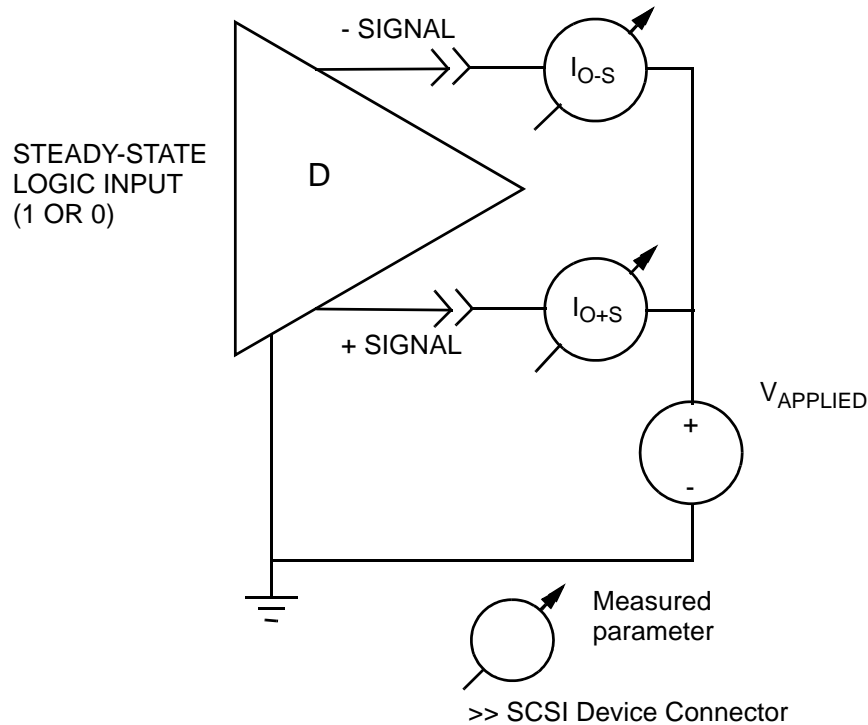


Figure A.5 - Common mode output voltage test

#### A.2.4 Short-circuit currents ( $I_{O-S}$ and $I_{O+S}$ )

Since an LVD bus segment allows multiple drivers, the possibility of contention requires a restriction on the power that may be sourced to the bus segment by a SCSI 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  $0 \leq V_{\text{applied}} \leq 5$  V. (see figure A.6).



**Figure A.6 - Driver short-circuit test circuit**

#### A.2.5 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 ground shall be between 0 V and 3,6 V when measured in accordance with figure A.7. 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.7 - Open-circuit output voltage test circuit**



### A.2.6 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 (e.g., one - zero, zero - one, one - off, off - one, zero - off, off - zero), the differential voltage measured with the test circuit of figure A.8 and table A.4, 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.9 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.9) between 0,2 and 0,8 of  $V_{SS}$  shall be greater than or equal to 1 ns. For paced transfers the output signal rise or fall times (see  $t_r$  in figure A.9) between 0,2 and 0,8 of  $V_{SS}$  shall be greater than or equal to 1 ns and less than or equal to 2,5 ns.

The rise and fall times specified in this subclause are requirements for a driver when using the LVD test circuit in figure A.8. They are not the observed rise or fall rates on an actual SCSI bus segment.

Measurement equipment used for rise and fall rate testing shall provide a bandwidth of 2 GHz minimum.

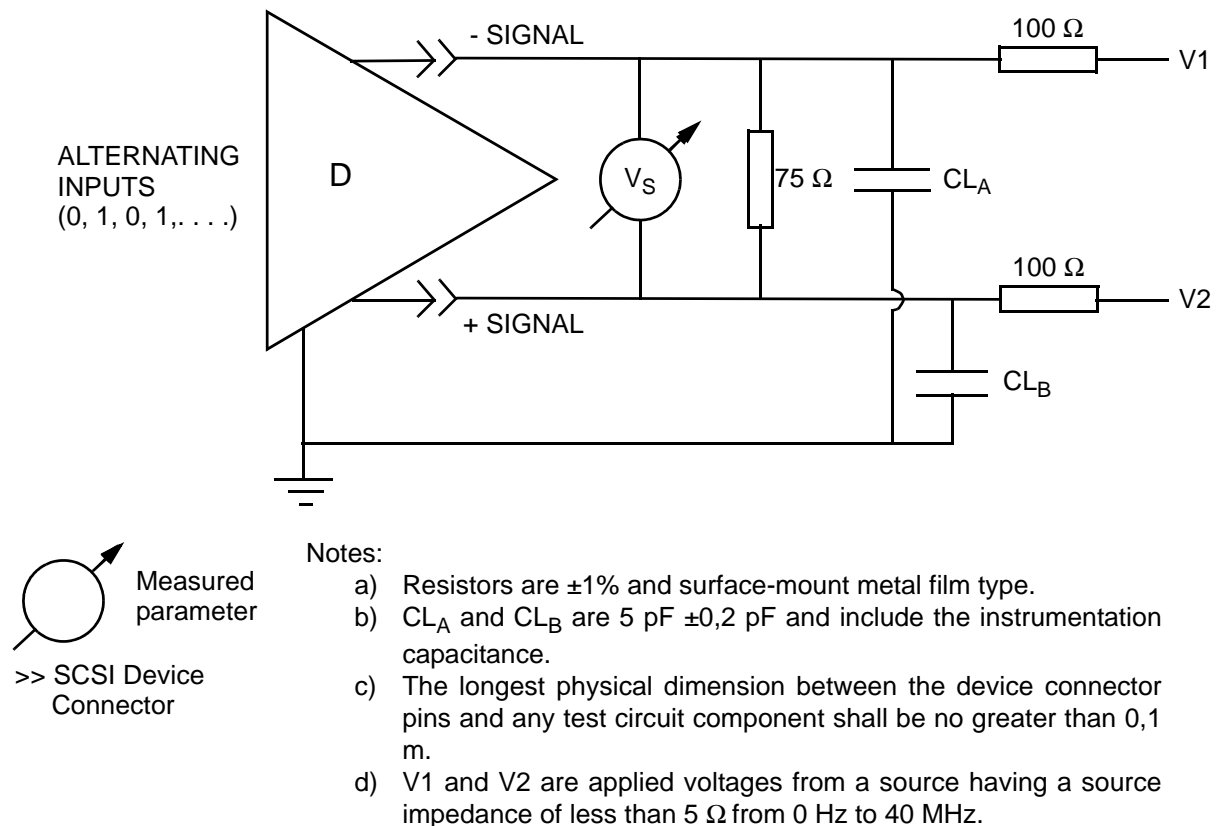



Figure A.8 - Differential output switching voltage test circuit

**Table A.4 - Driver switching test circuit parameters**

Test condition (see figure A.8)	V1	V2
Low common-mode voltage	1,311 V	0,889 V
High common-mode voltage	1,611 V	1,189 V


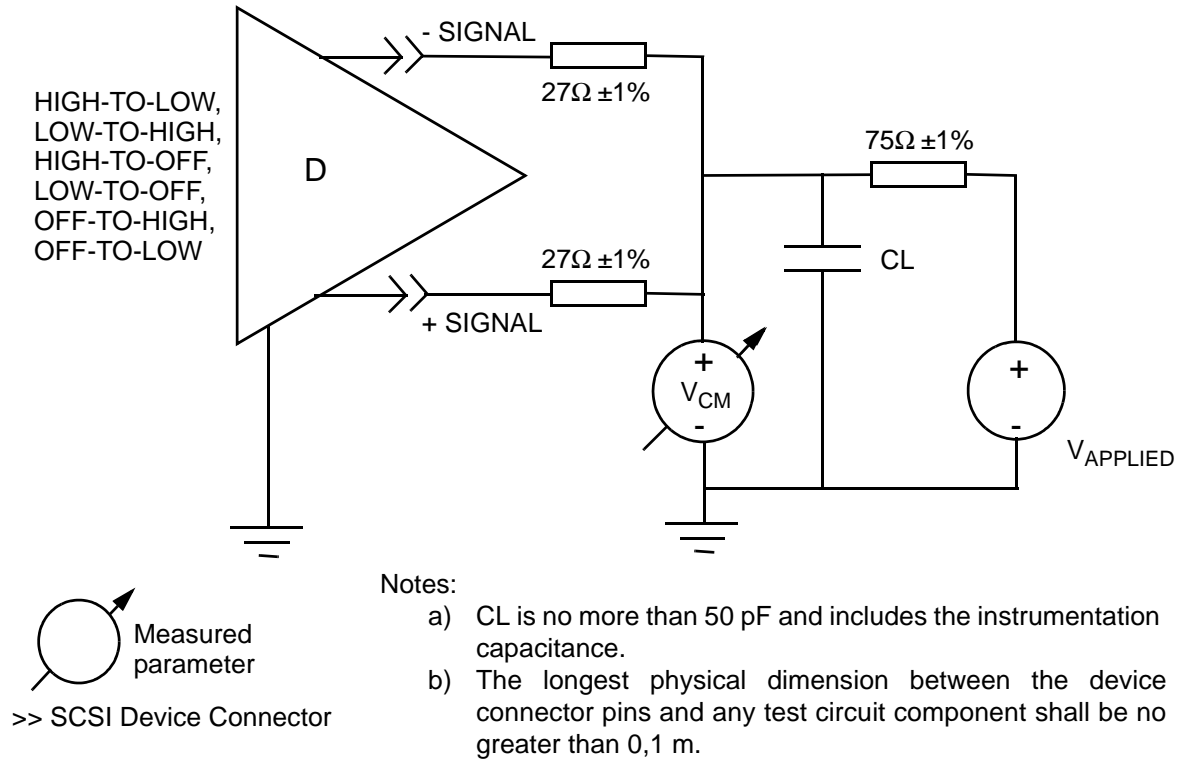
The signal voltage shall comply with the requirements shown in figure A.9.

**Figure A.9 - Driver output signal waveform**

#### **A.2.7 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 bus segment, 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.10, shall not vary more than specified in table A.5 as  $V_{applied}$  is varied over the range  $0,845 \leq V_{applied} \leq 1,655$ . 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.10 - Driver offset switching voltage test circuit**

Lower values of  $V_{CM(PP)}$  have lower EMI risk.

**Table A.5 - 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.3.1 Receiver characteristics overview

A receiver determines the logical state of the LVD bus segment as defined by the differential voltage that exists at the SCSI 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.

SCSI devices should incorporate a glitch filter function on REQ and ACK 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 and ACK signals.

### A.3.2 Receiver steady state input voltage requirements

#### A.3.2.1 Steady state input voltage requirements for fast-160 and slower receivers

Table A.6 defines the voltages for the requirements in this subclause.

Within the common-mode input voltage range ( $V_{CM}$ ), (figure 54)  $0,845\text{ V} \leq V_{CM} \leq 1,655\text{ V}$  an LVD receiver shall indicate the logical states shown in table A.6 with  $V_{IN}$  within the ranges shown in table A.6.

**Table A.6 - Receiver steady state input voltage ranges**

Differential 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.3.2.2 Steady state input voltage requirements for fast-320 receivers

Table A.7 and figure A.11 define the voltages for the requirements in this subclause. This subclause places requirements on receivers only and does not change input signal requirements.

The differential offset component of the receiver input voltage is defined as:

$$V_{RXOFFSET} = ( |V_A| - |V_N| ) / 2$$


where:

- a)  $|V_A|$  and  $|V_N|$  are the steady state differential asserted and negated signal levels applied to the receiver, measured after 100 ns assertion or negation of the signal.

Within the common-mode input voltage range ( $V_{CM}$ ), of  $0,845\text{ V} < V_{CM} < 1,655\text{ V}$  (see table A.8) and after calibration on a training pattern containing a differential offset voltage component,  $V_{RXOFFSET}$ , a fast-320 receiver shall detect assertion and negation with  $V_{IN}$  in the ranges shown in table A.7, and over an input signal offset range of  $-75\text{ mV} < V_{RXOFFSET} < +75\text{ mV}$  (see figure A.11 for the test circuit for  $V_{IN}$ ).

Table A.7 - Fast-320 receiver steady state input voltage ranges

Parameter	Minimum	Maximum
Correctable input signal offset ( $V_{RXOFFSET}$ ) range	n/a	$\pm 75$ mV
$V_{IN}$ to detect assertion <sup>a</sup>	-3,6 V	$V_{RXOFFSET} - 20$ mV
$V_{IN}$ to detect negation <sup>a</sup>	$V_{RXOFFSET} + 20$ mV	3,6 V
<sup>a</sup> after training		


Figure A.11 - Input signal offset component ( $V_{RXOFFSET}$ )

### A.3.3 Compliance test

Compliance to the requirements in A.3.2 shall be verified with the input voltages of table A.8 and the circuit of figure A.12.

Table A.8 - Receiver minimum and maximum input voltages.

Applied voltages (input voltage referenced to circuit ground) (see figure A.12)		Resulting differential input voltage	Resulting common-mode input voltage
$V_-$	$V_+$	$V_{IN}$	$V_{CM}$
0,860	0,830	0,030	0,845
0,830	0,860	-0.030	0,845
1,670	1,640	0,030	1,655
1,640	1,670	-0.030	1,655
3,310	0,000	3,310	1,655
0,000	3,310	-3,310	1,655
3,665	-0,355	4,020	1,655
-0,355	3,665	-4,020	1,655



&gt;&gt; SCSI Device Connector

Figure A.12 - Receiver input voltage threshold test circuit

### A.3.4 Receiver setup and hold times

Figure 60 and figure 61 define the receiver setup and hold times.

NOTE 44 - Dynamic testing is required to verify these timings.

### A.3.5 Fast-160 receiver bandwidth specifications


Receivers in SCSI device ports capable of operation at fast-160 shall not respond to the positive going or

the negative going pulse train shown in figure A.13 for any signal.

The pulse trains shall have the properties listed in table A.9:

**Table A.9 - Properties of receiver pulse train**

Parameter	Value
pulse amplitude	in the range 0 mV to 60 mV
pulse width at 0 mV	in the range 0 ns to 0,2 ns
repetition time	greater than 12,5 ns
starting level	$\pm 60$ mV relative to 0 mV
slew rate	in the range 400 mV/ns to 1000 mV/ns
A.C. common mode	0 mV
D.C. common mode	1,25 V



**Figure A.13 - Fast-160 receiver pulse train example**

#### A.4 Transceiver characteristics

#### A.4.1 Transceiver output/input currents, $I_{I-L}$ and $I_{I+L}$

The requirements in this subclause apply as a test method to ensure compliance with the LVD parameters table 21 and table 22. With the transceiver in an off condition (i.e., not transmitting) and the + and - signals connected to a variable voltage source,  $V_{\text{applied}}$ , the output leakage currents  $I_{I-L}$  and  $I_{I+L}$  shall not exceed the applicable LVD values in table 22 over the range  $0,00 \text{ V} \leq V_{\text{applied}} \leq 3,01 \text{ V}$  (see figure A.14). The maximum LVD applicable current from table 22 is  $I_{\text{max}}$ .

These measurements apply with the transceiver's power supply in both the powered on and powered off conditions.

$$|I_{I-L}| < I_{\text{max}}$$

$$|I_{I+L}| < I_{\text{max}}$$




Figure A.14 - Transceiver off-state output current test circuit

#### A.4.2 Transceiver maximum input voltages

See table 21 and table 22.



## Annex B

(normative)

### SCSI bus fairness

#### B.1 Model

Implementation of SCSI bus fairness is optional, however, if implemented, the SCSI bus fairness protocol shall conform to this annex.

A SCSI device determines fairness by monitoring prior arbitration attempts by other SCSI devices. It shall postpone arbitration for itself until all lower priority SCSI devices that previously lost arbitration either win a subsequent arbitration or discontinue their arbitration attempts (e.g., as in the case where the SCSI initiator port aborted an outstanding command thus removing the need to re-arbitrate).

When a SCSI device is not arbitrating for the SCSI bus, it shall monitor the arbitration attempts of the other SCSI devices and refresh the fairness register with the SCSI IDs of any lower priority SCSI devices that lost arbitration.

Whenever a requirement for arbitration arises, a SCSI device shall first check to see if its fairness register is clear (see 10.4.1). If it is clear, then no lower priority SCSI device has attempted and lost the previous arbitration and therefore, this SCSI device may now participate in arbitration. If the fairness register is not clear, the SCSI device shall postpone arbitration until all lower priority SCSI IDs have been cleared from the fairness register. Lower SCSI IDs are cleared as those SCSI devices win arbitration. SCSI IDs shall also be cleared if a SCSI device discontinues arbitration (e.g., as a result of an ABORT TASK message, ABORT TASK SET message, CLEAR TASK SET message, or logical unit reset).

The fairness register may be refreshed, updated or cleared. The fairness register is refreshed by copying the SCSI IDs of any lower priority SCSI devices that lost arbitration into the fairness register. A refresh of the fairness register completely replaces the previous contents of the fairness register. The fairness register is updated by removing the SCSI IDs of any lower priority devices that win arbitration or discontinue arbitration. The fairness register is cleared by setting all of its bits to zero. SCSI IDs may only be added to the fairness register by a refresh but may be subtracted by a refresh, update or clear.

Since the fairness register is only refreshed when a SCSI device is not arbitrating for itself, the fairness register is effectively frozen by a SCSI device prior to a requirement for its own arbitration arising. Other lower priority SCSI devices that were not latched shall not be added to the fairness register until this SCSI device has successfully arbitrated.

Arbitration fairness in SCSI target devices is controlled with the disconnect-reconnect mode page (see 18.1.2).

#### B.2 Determining fairness by monitoring prior bus activity

##### B.2.1 Fairness for normal arbitration method

For the normal arbitration method this standard requires that within a bus set delay of when BSY was first asserted, the SCSI ID of all arbitrating SCSI devices shall appear on the bus. A SCSI device shall sample the bus after this time, to detect SCSI devices that are arbitrating, which SCSI device won, and which SCSI devices lost. Since the lower priority SCSI IDs are released after an arbitration delay from the assertion of BSY, the data bus shall be sampled after a bus set delay but before an arbitration delay.

NOTE 45 - For ease of implementation, the sample window may begin when BSY is asserted to begin arbitration and extending until SEL is asserted. Sampling of the bus during this time should occur at a high enough rate to ensure multiple samples within this window.

### **B.2.2 Fairness for QAS**

For QAS, after detection of a valid QAS REQUEST message, this standard requires that between 1 000 ns (i.e., QAS arbitration delay) and 1 490 ns (i.e., QAS arbitration delay + bus settle delay + 2 system deskew delays) after detection of the MSG, C/D, and I/O signals being false, the SCSI ID for all arbitrating SCSI devices shall appear on the bus. The SCSI device shall sample the bus during this time to detect which SCSI devices are attempting arbitration, which SCSI device won, and which SCSI devices lost. Since the lower priority SCSI IDs are released at 1 490 ns, a continuous sampling of the data bus between 1 000 ns and 1 490 ns is required.

## **B.3 Fairness algorithm**

### **B.3.1 Fairness states overview**

A SCSI device shall be in one of three fairness states. A SCSI device shall be in fairness wait state if it is waiting for a clear fairness register to participate in arbitration. A SCSI device shall be in the fairness participate state if it is participating in arbitration. A SCSI device shall be in the fairness idle state for all other conditions. A SCSI device shall enter the fairness idle state after any reset event. A SCSI device shall implement a lockout delay to prevent devices that stop arbitrating from causing deadlock.

### **B.3.2 Fairness idle state**

A SCSI device shall be in fairness idle state if it is not arbitrating for the SCSI bus. The fairness register shall be refreshed during each arbitration while in the fairness idle state. The fairness register shall then contain the SCSI IDs of any lower priority arbitration losers. A SCSI device shall transition to the fairness wait state to arbitrate only if the fairness register is not clear. A SCSI device shall transition to the fairness participate state to arbitrate only if the fairness register is clear.

### **B.3.3 Fairness wait state**

A SCSI device shall be in fairness wait state to arbitrate for the SCSI bus if the fairness register is not clear. The fairness register shall be updated during each arbitration while in the fairness wait state. The updates subtract out arbitration winners and non-participants. The updates shall not add any new arbitration participants. A SCSI device shall transition to the fairness participate state once all of the bits in the fairness register have been cleared. A SCSI device shall transition to the fairness idle state if it stops arbitration.

### **B.3.4 Fairness participate state**

A SCSI device shall be in fairness participate state to participate in arbitration. A SCSI device is only allowed to be in fairness participate state if the fairness register is clear. If the SCSI device wins the arbitration it shall refresh the fairness register. If the SCSI device loses arbitration it shall keep the fairness register clear and remain in the fairness participate state. A SCSI device shall transition to the fairness idle state if it stops arbitration.

### **B.3.5 Lockout delay**

A SCSI device shall implement a lockout delay to prevent deadlock during normal arbitration. The lockout delay shall be at least 2 us and shall start once bus free has been detected. Each SCSI device shall clear its fairness register if no SCSI device starts arbitration before the lockout delay expires.

NOTE 46 - The lockout delay is required because all the SCSI devices may be waiting for other SCSI devices to start the arbitration process. Although rare, the following example is valid and may occur.

Assume a SCSI initiator port at SCSI ID 7 starts tasks in SCSI devices at SCSI IDs 0 and 1. After a while, SCSI devices 0 and 7 begin arbitration, 7 wins and SCSI device 0 is recorded in the fairness register of all SCSI devices. Assume that SCSI device 7, the arbitration winner, aborts its task on SCSI device 0. Both SCSI device 7 and 1 is now waiting on SCSI device 0 to arbitrate but SCSI device 0 stops arbitration since its task has been aborted. When the lockout delay expires, SCSI devices 7 or 1 may begin arbitration.

NOTE 47 - The lockout delay is not required for QAS arbitration because the bus shall go bus free after a QAS arbitration delay. The fairness register shall be cleared because of mixed arbitration (see B.3.6).

### **B.3.6 Mixed arbitration**

The fairness register shall be cleared by all SCSI devices that detect a change in arbitration mode. If the previous arbitration was an QAS arbitration then the fairness register shall be cleared if bus free is detected. If the previous arbitration was a normal arbitration then all devices that detect the QAS arbitration shall clear the fairness register.

NOTE 48 - SCSI devices that do not support QAS may not detect the QAS arbitration and may not clear their fairness register.

## **B.4 SCSI initiator port fairness recommendations**

Generally the SCSI initiator port is the highest priority SCSI device on the bus. This guarantees the SCSI initiator port wins arbitration and may overlap commands to multiple SCSI devices. To maintain this capability, the SCSI initiator port should not implement fairness towards lower priority SCSI target ports.

In the case of a SCSI multi-initiator port system, the SCSI initiator ports should be the highest priority SCSI devices. However, in order to implement fairness between them, the higher priority SCSI initiator port may implement fairness with the lower priority SCSI initiator ports only. This would require a second register in which a bit is enabled for each lower priority SCSI device to which a higher priority SCSI device shall be fair.

## Annex C

(normative)

### Nonshielded connector alternative 4

#### C.1 Nonshielded connector alternative 4 signal definitions

For the physical descriptions and usage guidelines for the nonshielded connector alternative 4 see Single Connector Attachments (SCA-2), EIA-700AOAE and SCA-2 Unshielded Connections, SFF-8451.

#### C.2 VOLTAGE and GROUND signals

Three voltage supplies and corresponding ground return signals are provided by the backplane connector to the SCSI device. Table C.1 provides the specifications for each of the voltage supplies.

NOTE 49 - The details of the actual SCSI device supply requirements is tailored for each SCSI device and enclosure combination to ensure appropriate supply voltage to the SCSI device.

**Table C.1 - Voltage specification limits**

Voltage	Number of pins	Number of grounds	Requirements on supply at the nonshielded connector alternative 4 connector	Current capability average/peak
12 V	3	3	12 V D.C.+5 % / -7 %	0/0 to 2,5 / 5 A
5 V	2	2 <sup>a</sup>	5 V D.C.±5 %	0/0 to 2 / 2,5 A
Opt 3.3 V	2	2 <sup>a</sup>	3,3 V D.C.±5 %	0/0 to 3 / 3 A
<sup>a</sup> The two logic level grounds are shared between +5 V D.C.and +3,3 V D.C.				

The peak current capability is measured during operation or initialization after voltages have stabilized at the operating level. Inrush currents are managed by the power supply during normal power on and by the CHARGE signals during hot plugging.

For each voltage, the current supplied to the SCSI device should be distributed as evenly as possible among the connecting pins.

The backplane power supplies are required to operate correctly and maintain regulation from zero current to the peak current. SCSI Device sequencing provisions may be required to avoid overloading power supplies during SCSI device spin-up sequencing. Voltage dips to -10 % are allowed on the 12 V D.C.supply during spin up.

For each voltage, an appropriate number of current return GROUND signal pins have been assigned.

- a) The GROUND signal pins for all voltages shall be tied together in the SCSI device.
- b) The GROUND signals in the backplane may be tied together or connected separately to the power

supplies as required by the particular subsystem.

- c) The logic level grounds, GROUND (5V/3,3V) are shared between the currents provided by the 5 V D.C. and 3,3 V D.C. signals. The sum of the 5 V D.C. and 3,3 V D.C. currents shall not exceed 3 A.

### C.3 CHARGE signals

Three charge signals, one for each of the power supply voltages, provide controlled precharging of the SCSI device's internal circuits to avoid excessive surge currents during hot plugging.

The precharge pin mates early to allow the precharge to take place before the voltage pins make contact. If required, the precharge control circuits are located on the backplane side of the connector. The backplane should assume that the VOLTS signals for each voltage are shorted together with the corresponding CHARGE signal on the SCSI device. Systems without a hot-plug capability or with an alternative hot plugging mechanism are not required to implement the precharge control circuit and are not required to use long and short pins on the backplane connector.

After the SCSI device capacitance is charged, but before the MATED signal indicates that the power signals are seated, the SCSI device shall not use more than 1 A on the precharge voltage pin. This is required to protect the precharge pin from over-current damage and to provide additional flexibility in the design of the precharge circuit. The voltage provided by the precharge circuitry is as specified by table C.2. Any circuitry on the SCSI device that uses the CHARGE voltage for executing initialization operations shall operate within the current and voltage constraints specified for the CHARGE signals.

**Table C.2 - Charge supply to SCSI device**

<b>CHARGE signal</b>	<b>Requirements on supply at the nonshielded connector alternative 4 connector for backplane after CHARGE complete</b>	<b>Maximum surge at SCSI device</b>	<b>Maximum continuous required by SCSI device</b>
12 V	12 V D.C.+5 % / -12 %	6 A	1 A
5 V	5 V D.C.+5 % / -17 %	6 A	1 A
OPT 3,3 V	3,3 V D.C.+5 % / -24 %	6 A	1 A
It may not be possible to meet these parameters with passive components.			

After precharge is complete and the SCSI device is mated, there is no guarantee that the precharge signal provides any current to the SCSI device and the SCSI device should not depend on such current for operation.

The system designer should assume that the VOLTS signals and the corresponding CHARGE signal are shorted together on the SCSI device.

### C.4 SPINDLE SYNC

The spindle synch is assigned a single pin, SPINDLE SYNC. The synchronization protocol and the electronic requirements for the SPINDLE SYNC signal are vendor specific.

Spindle synchronization is managed by the SCSI command set. The signal current requirements shall not exceed 100 mA and the signal voltage shall not be higher than 5,25 V D.C. or lower than -0,25V. The minimum driver capability required by the SPINDLE SYNC signal shall be sufficient to drive the receivers on 30 identical SCSI devices.

The SPINDLE SYNC signal when driving should be capable of driving a minimum of 30 identical SCSI devices.

NOTE 50 - The SPINDLE SYNC signal is a source for noise and may be affected by noise. The design of the SPINDLE SYNC signal interconnections should take this into account by properly laying out the SPINDLE SYNC signals on the backplane or motherboard. Proper layout considers routing relative to other signals, the proper line impedance, and termination if necessary. The selection of the electronic transceiver should also take into account the possibility of noise. The signal levels, signal rise time, receiver thresholds, and receiver hysteresis should be considered as part of that selection.

## C.5 ACTIVE LED OUT

The ACTIVE LED OUT signal is driven by the SCSI device when a SCSI operation is being performed. The ACTIVE LED OUT signal shall be implemented and is used to indicate that the SCSI device is operating on a command. Other optional indications may be provided by flashing the LED. The host system is not required to generate any visual output when the ACTIVE LED OUT signal is raised, but if such a visual output is provided, it shall be white or green to indicate that normal activity is being performed.

The ACTIVE LED OUT signal is designed to pull down the cathode of an LED. The anode shall be attached to the proper +5 V D.C. supply through an appropriate current limiting resistor. The LED and the current limiting resistor are external to the SCSI device.

See table C.3 for the output characteristics of the ACTIVE LED OUT signal

**Table C.3 - Output characteristics of drive ACTIVE LED OUT signal**

State	Current drive available	Output voltage
Drive LED off	$0 < I_{OH} < 100 \mu A$	Not defined by this standard
Drive LED on	$I_{OL} > 30 \text{ mA}$	$0 < V_{OL} < 0,8V$

## C.6 Motor start controls

The method of starting the SCSI device's motor is established by the signals RMT\_START and DLYD\_START, as described in table C.4. The state of these signals may either be wired into the backplane socket or driven by logic on the backplane.

Each SCSI device location should have these signals supplied independently to ensure proper operation. If the signals were bussed, a SCSI device with a power failure may clamp the signals in a condition that caused operational SCSI devices to behave incorrectly.

- If the GROUND state is implemented for RMT\_START, bussing between SCSI devices is permissible.
- If the GROUND state is implemented for DLYD\_START, bussing between SCSI devices is permissible.
- If the OPEN state is implemented for RMT\_START, this signal shall not be bussed between SCSI devices.

- d) If the OPEN state is implemented for DLYD\_START, this signal shall not be bussed between SCSI devices.

**Table C.4 - Definition of motor start controls**

Case	DLYD_START	RMT_START	Motor Spin Function
1	open	open	Motor spins up at D.C. power on.
2	open	ground	Motor spins up only when START UNIT command is received.
3	ground	open	Motor spins up at D.C. power on after a delay in seconds 12 times <sup>b</sup> the value of the numeric SEL_ID for the SCSI device.
4	ground	ground	Reserved. SCSI devices not implementing this option shall execute power control according to the rules of case 2 (see Annex D).
<sup>b</sup> This value may be reduced by SCSI device suppliers to reflect the worst-case time duration of peak current drains at the 12 V D.C. or 5 V D.C. source or both during motor spin up. In no case should the delay exceed 12 s.			

The OPEN and GROUND states are established as described in table C.5.

**Table C.5 - Electronic requirements for input controls**

State	Current	Voltage
open	$0 \mu\text{A} < I_{IH} < 100 \mu\text{A}$	$2,4 \text{ V D.C.} < V_{IH} < V_{CC} + 0,5\text{V}$
ground	$-3 \text{ mA} < I_{IL} < 0 \text{ mA}$	$-0,5 \text{ V D.C.} < V_{IL} < 0,4\text{V}$
The SCSI device provides the voltage source for the open signal state.		

## C.7 SCSI ID selection

The SCSI device address of the attached SCSI device shall be determined by the state of the signals SCSI ID(0-3). Table C.6 indicates the relationship between the level of the SCSI ID signals and the selected SCSI device address.

**Table C.6 - SCSI device ID selection signals**

Address	ID (0)	ID (1)	ID (2)	ID (3)
0	open	open	open	open
1	ground	open	open	open
2	open	ground	open	open
3	ground	ground	open	open
4	open	open	ground	open
5	ground	open	ground	open
6	open	ground	ground	open
7	ground	ground	ground	open
8 <sup>a</sup>	open	open	open	ground
9 <sup>a</sup>	ground	open	open	ground
10 <sup>a</sup>	open	ground	open	ground
11 <sup>a</sup>	ground	ground	open	ground
12 <sup>a</sup>	open	open	ground	ground
13 <sup>a</sup>	ground	open	ground	ground
14 <sup>a</sup>	open	ground	ground	ground
15 <sup>a</sup>	ground	ground	ground	ground
<sup>a</sup> Addresses in the range of 8 to 15 are only supported by SCSI devices implementing a 16-bit DATA BUS (see 8.2).				

The OPEN and GROUND states are established as specified in table C.5.

## C.8 MATED signals

### C.8.1 MATED signals overview

If MATED 1 and MATED 2 signals are not mated then one or more short pins are not mated.

If MATED 1 and MATED 2 signals are mated then the mated condition of the short pins is indeterminate.

The MATED 1 and MATED 2 signals may indicate to the SCSI device that the SCSI device is seated in an nonshielded connector alternative 4 connector and it may begin power on processing. The circuit



described in figure C.1 or a similar circuit is used to implement the MATED function. The signal requirements are indicated in C.8, but may be met by the example circuit or by similar circuits.

### **C.8.2 MATED 2/drive side**

The signal is attached to signal ground on the SCSI device side.

### **C.8.3 MATED 2/backplane side**

The signal is attached either directly or through optional logic in such a manner that the MATED 1 signal is held to a ground level when the MATED 2 connection is completed. The SCSI device shall sink no more than 100 mA to ground through the MATED 2 pin if optional logic is used.

### **C.8.4 MATED 1/drive side**

The MATED 1 signal shall be sensed by the SCSI device. When the MATED 1 connection is determined to be at a ground level, the SCSI device may detect that the SCSI device has been partially mated. Assuming the mating process continues uninterrupted until completion, including sensing of the SCSI ID Selection signals and the motor start controls, then normal power on procedures may begin 250 ms after the MATED 1 signal is observed to transition to the ground level. When the MATED 1 connection is determined to be at the open level, the SCSI device is not mated. The MATED 1 signal is tied up to a TTL positive level when the SCSI device is not installed.

If the SCSI device is mated and operating, it may optionally detect the open level of MATED 1 as an indication that the SCSI device is partially unmated and may be about to be removed.

If the SCSI device supports detection of the open level of MATED 1 to prepare itself for power removal or for physical removal from the enclosure, the detection shall occur within 1 s from the time that the Mated 1 open level occurs at the SCSI device.

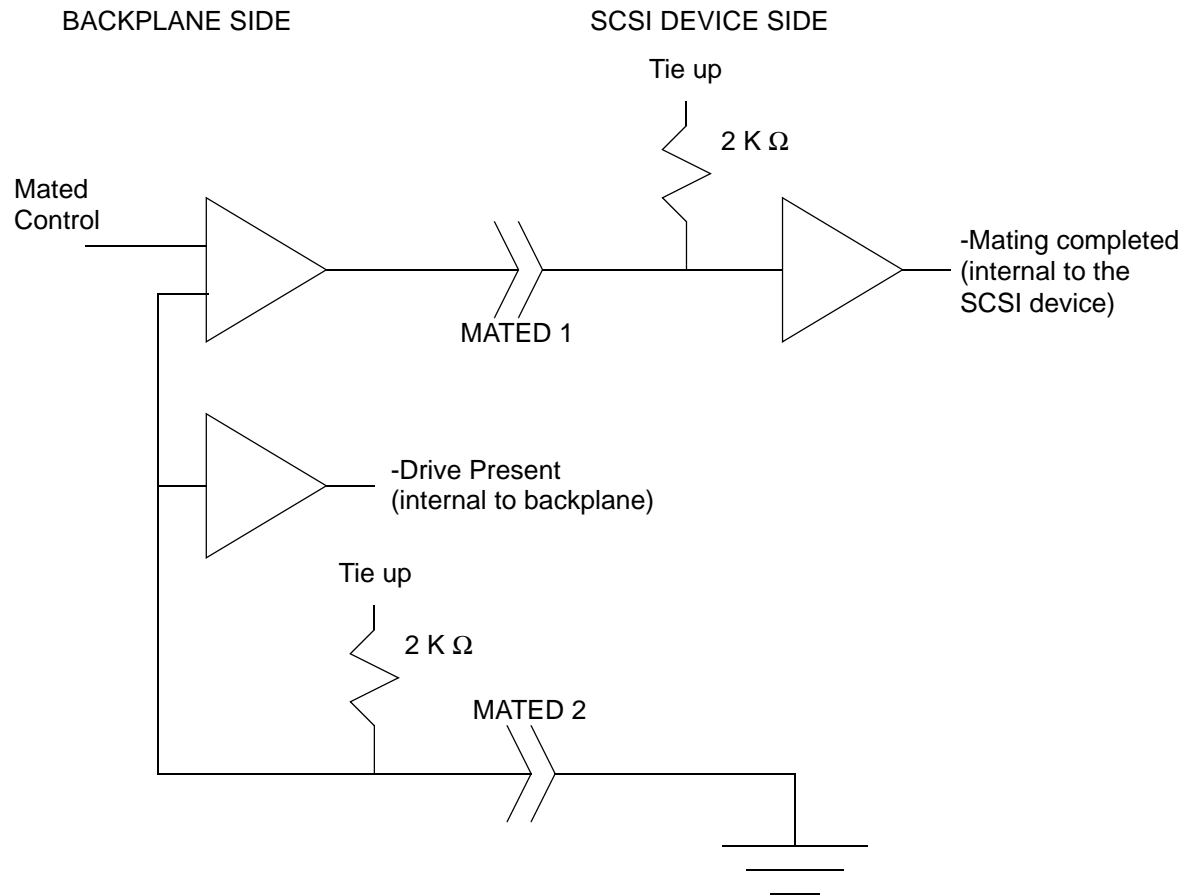
The following SCSI device behaviors are defined when a SCSI device detects the open level of MATED 1:

- a) The SCSI device may optionally perform a spin-down operation.
- b) The SCSI device may optionally transfer any cached information to the storage media.

### **C.8.5 MATED 1/backplane side**

The signal shall be held to a ground level when the MATED 2 connection is completed. The MATED 1 signal shall be held to the open level when the MATED 2 connection is not completed. The ground and open levels are defined by table C.5.

The enclosure may optionally control the MATED 1 signal to indicate that the SCSI device is about to be removed.

**Figure C.1 - Sample circuit for mated indications**

## Annex D

(normative)

### Removal and insertion of SCSI devices

#### D.1 Removal and insertion of SCSI devices overview

This annex defines the physical requirements for removal and insertion of SCSI devices on the SCSI bus segment. 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. 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.

#### D.2 Case 1 - Power off during removal or insertion

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

#### D.3 Case 2 - RST signal asserted continuously during removal or insertion

- a) RST signal shall be asserted continuously by the SCSI initiator port 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 SCSI device connector contact to the SCSI bus segment. The ground connections shall be maintained during and after the connection of the SCSI device to the SCSI bus segment;
- 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 SCSI device connector contact from the SCSI bus segment.

NOTE 51 - The translation of the 1 ms time to mechanical provisions is vendor specific.

#### D.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 SCSI device connector contact to the SCSI bus segment. The ground connections shall be maintained during and after the connection of the SCSI device to the SCSI bus segment;
- 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 SCSI device connector contact from the SCSI bus segment;
- d) The SCSI device being removed or inserted shall employ transceivers that conform to the applicable requirements in 7.2.5.3 for glitch-free powering on and off. The SCSI device shall maintain the high-impedance state at the SCSI 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 52 - Any on board switchable terminators as well as SCSI device transceivers may affect the impedance state at the SCSI device connector contacts.

- e) The SCSI device power may be simultaneously switched with the SCSI bus segment 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 segment 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 ms 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  $\mu$ F.

## D.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 SCSI device connector contact to the SCSI bus segment. The ground connections shall be maintained during and after the connection of the SCSI device to the SCSI bus segment;
- 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 SCSI device connector contact from the SCSI bus segment;
- d) The SCSI device being removed or inserted shall employ transceivers that conform to the applicable requirements in 7.2.5.3 for glitch-free powering on and off. The SCSI device shall maintain the high-impedance state at the SCSI 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 53 - Any on-board switchable terminators as well as SCSI device transceivers may affect the impedance state at the SCSI device connector contacts.

- e) The SCSI device power may be simultaneously switched with the SCSI bus segment 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 segment 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 ms 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  $\mu$ F.

NOTE 54 - LVD SCSI devices may require more stringent system design to tolerate transients that occur during case 4 insertion or removal.

## Annex E

(normative)

### Simple expander requirements

#### E.1 Introduction

This annex describes extended SCSI configurations that incorporate simple expanders to separate the SCSI domain into electrically isolated parts by using active circuits. The requirements for these active circuits are described in this annex. These active circuits are specified to produce few additional requirements beyond those defined within this standard for SCSI device ports. SCSI device ports that are implemented in accordance to requirements in this standard do not produce any behavior that causes extended configurations defined in this annex to fail provided that the extended configuration meets the requirements in this annex.

#### E.2 Glossary

**E.2.1 SCSI domain:** A SCSI domain is a logical bus with at least one bus segment, at least one SCSI initiator port, and at least one SCSI target port. Domains with multiple bus segments are enabled through the use of bus expanders. Domains consist of the set of SCSI devices that are addressable from SCSI device ports.

**E.2.2 Simple expander:** Devices that couple bus segments together without using SCSI ID's in the device and are "invisible" to SCSI device ports.

#### E.3 Bus segments in a SCSI domain

SCSI bus segments are based on the assumption that there is a single electrically conducting path between bus segment terminators for each signal and that a SCSI domain contains all the devices between these two terminators. This electrical path is assumed to pass signals in both directions without delay other than that caused by the propagation delay of the transmission line associated with the path. It is assumed that there are no intervening active components in the path between the bus segment terminators.

An alternative to this would be to build SCSI domains that use more complex physical implementations where there may be active electrical components between SCSI devices.

A building block for these alternative implementations is a bus segment that is defined as two bus segment terminators and the associated single electrically conducting path between these terminators, for each signal, that satisfies the assumptions in the first paragraph of E.7. Multiple bus segments may be functionally connected together by the active circuits described in this annex.

Each bus segment shall:

- a) have TERMPWR sources and TERMPWR distribution parameters;
- b) use the same transceiver type (e.g., LVD) within the segment, however, a SCSI domain may contain segments that use different transceiver types; and
- c) follow the same rules that are described in this standards.

Although each bus segment does not extend any SCSI properties the combination of multiple bus segments into a single SCSI domain allows:

- a) increased SCSI device counts;
- b) increased physical length limits;

- c) controlling ground voltage shifts;
- d) dynamic removal and replacement of portions of SCSI domains; and
- e) mixing of SCSI device transceiver types (e.g., LVD).

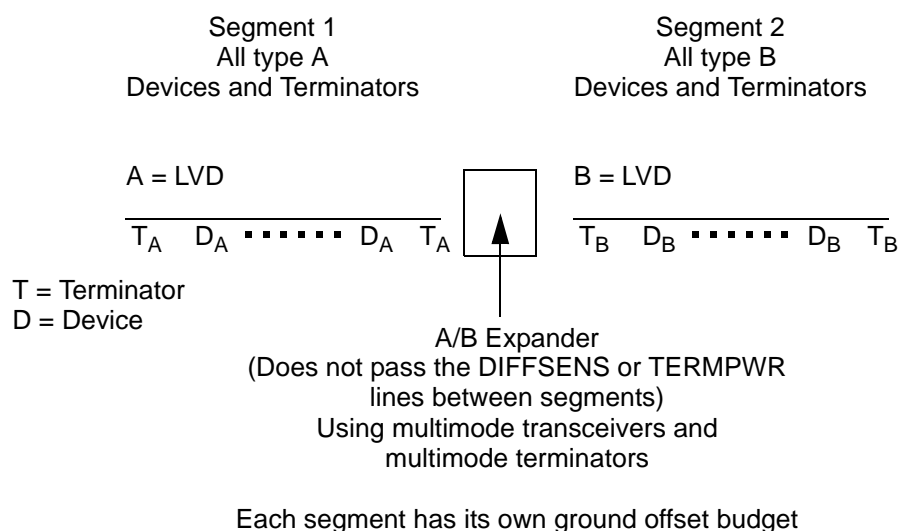
## E.4 Simple bus expanders

The following features shall be the required properties of simple bus expanders:

- a) No SCSI IDs used;
- b) No ARBITRATION phases initiated;
- c) No messages originating with the simple expander;
- d) Retransmitted signals from the simple expander that meet the same requirements as any other SCSI device at the expander boundary. The simple expander boundary may be at a separable connector;
- e) Simple expander receivers shall operate error free with the most degraded signal allowed for any SCSI device receiver at the expander boundary. The simple expander boundary that may be at a separable connector;
- f) Simple expanders shall not interfere with the REQ/ACK OFFSET count in any SCSI port device in the SCSI domain other than that caused by the propagation time through the expander;
- g) Simple expanders shall retransmit RST signal assertions from one bus segment to the other regardless of the state of any other SCSI signals on either side;
- h) Simple expanders shall operate with any arbitrary placement of the SCSI initiator ports and SCSI target ports with respect to the simple expander (e.g., all SCSI target ports and SCSI initiator ports may be on the same side of the expander or there may be SCSI initiator ports and SCSI target ports on both sides of the expander);
- i) TERMPWR shall not be connected in the expander between the bus segments;
- j) DIFFSENS shall not be electrically or logically connected between bus segments; and
- k) Transceiver mode changes (e.g., LVD) on one bus segment shall cause the simple expander to create a bus reset condition (see 12.3) on the another bus segment.

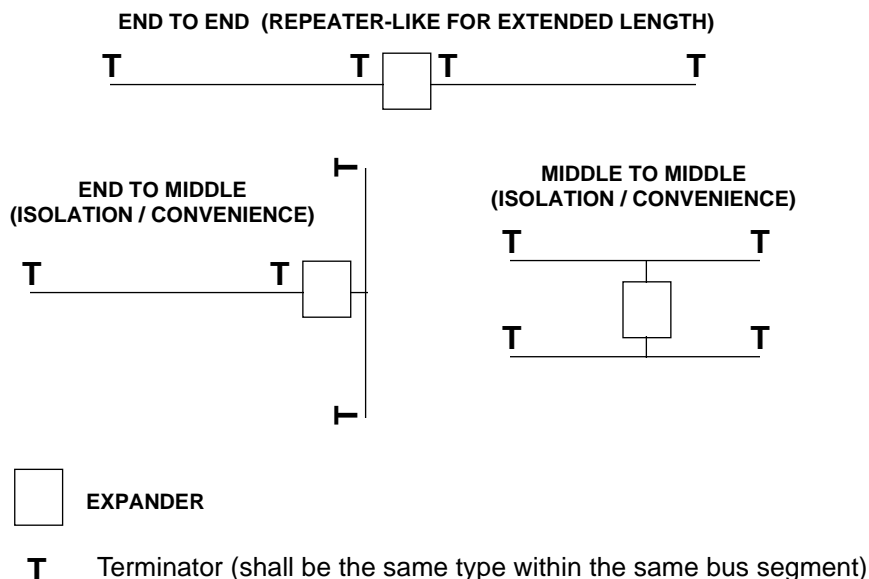
NOTE 55 - Simple expanders should consume minimal propagation time during arbitration so that the end to end SCSI domain propagation time budget may be used primarily for longer physical length connections (see E.8.3).

Figure E.1 shows a single simple expander between two bus segments.



**Figure E.1 - A two segment domain using a single simple expander circuit**

Figure E.2 shows three ways that simple expanders may be used to connect bus segments.



**Figure E.2 - Three ways to couple bus segments together with simple expanders**

## E.5 Homogeneous type

If a simple expander has the same type of bus segment on both sides it is termed a homogeneous expander. The homogeneous expander does not do transceiver type conversions.

NOTE 56 - This kind of simple expander may be useful in existing systems where domain length increases may be achieved by inserting a single homogeneous expander in the right place. By placing a homogeneous expander near the backplane one creates a short, heavily loaded backplane bus segment and a point to point bus segment to the SCSI initiator port. Increases in overall SCSI domain physical length may be achieved because the point to point bus segment length limit is longer than the multi-drop bus segment length.

## E.6 Heterogeneous types

Simple expanders that have different transceiver types on each side are heterogeneous expanders.

## E.7 SCSI domain examples using simple expanders

Figure E.3 shows two examples of SCSI domains built using only simple expanders.




Figure E.3 - Examples of SCSI domains using simple expanders

## E.8 Rules for SCSI domains using simple expanders

### E.8.1 Rule summary

The rules are summarized in table E.1 followed by detailed discussion for each in the subsequent subclauses.

Valid SCSI domains shall conform with the rules in table E.1.



**Table E.1 - SCSI Domain rules**

<b>Rule</b>	<b>Description</b>	<b>Subclause</b>
1	All bus segments in the domain shall comply with their individual bus segment length limits and other bus segment related requirements.	E.8.2
2	Any bus segment between two other bus segments shall support the highest performance level that is capable of being negotiated between the two other bus segments. (e.g., two wide LVD fast-40 bus segments shall not be separated by a bus segment that does not support both wide and fast-40. See figure E.4 for examples.	E.8.3
3	The expander between two bus segments shall support the maximum performance levels supported on each SCSI interface of the simple expander.	E.8.4
4	The maximum propagation time between any two SCSI device ports in the SCSI domain shall not exceed 400 ns.	E.8.5
5	The number of addressable SCSI devices in the SCSI domain shall not exceed the addressability of the SCSI devices in the SCSI domain.	E.8.6
6	Loop topologies are not allowed.	E.8.7

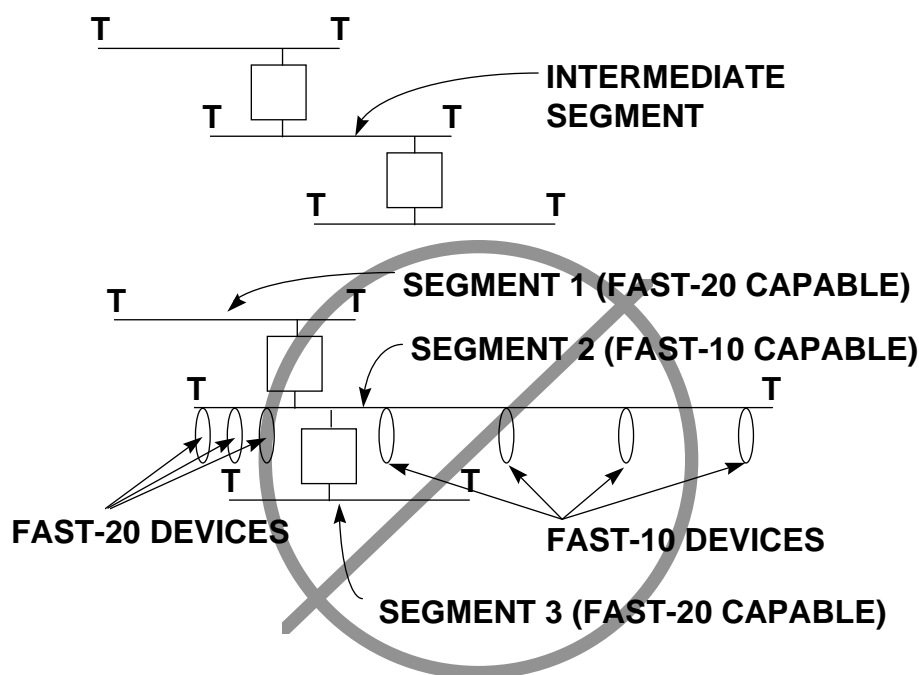
**E.8.2 Rule 1**

Requirements for SCSI domains consisting of a single LVD bus segment are specified in detail in other clauses of this document. Every bus segment in a multi-segment SCSI domain shall conform to the requirements for single bus segment SCSI domains of the same transceiver type.

**E.8.3 Rule 2**

Rule 2 applies to intermediate bus segments that only exist in SCSI domains of at least three bus segments. The bus segment between the two other bus segments is the intermediate bus segment. The intermediate bus segment shall be wide if both other bus segments are wide. The intermediate bus segment shall support the lowest common transfer rate between the other bus segments.

An example of a rule 2 configuration violation is shown in figure E.4.



**Figure E.4 - Intermediate bus segments and performance ranking**

The configuration in figure E.4 is valid only if the transfer rate is limited to fast-10 for any DATA phase transfers between bus segment 1 and bus segment 2, bus segment 2 and bus segment 3, or bus segment 1 and bus segment 3. Even though fast-20 capable SCSI devices in bus segment 2 are located close to the simple expanders and the distance between the simple expanders is small, the bus segment length is defined by the distance between the terminators, not by the distance to the simple expander connection or to the SCSI device ports. The intermediate bus segment is not fast-20 capable and may not be used for fast-20 transfer rates between bus segment 1 and bus segment 3. Bus segment 2 shall not be used for fast-20 transfer rates within bus segment 2. Fast-20 transfer rates are allowed between SCSI device ports in bus segment 1 or between SCSI device ports in bus segment 3.

The intermediate bus segment in this example receives signals at the higher transfer rate on the DATA BUS, DB(P\_CRCA), and/or DB(P1) signals, but since the devices in the intermediate bus segment are not participating in the higher transfer rate and are waiting for the next BUS FREE phase or some other phase they are unaffected by the higher transfer rates.

For multimode bus segments, any dynamic transceiver mode change is treated as a fault and the simple expander shall create a bus reset condition (see 12.3) on the segment opposite the one that experienced the transceiver type change. The simple expander shall detect this change by sensing the DIFFSENS line. This scheme ensures that the SCSI initiator ports on the other bus segments are aware of the transceiver mode change and may reassess whether this new transceiver type is consistent with the performance requirements for the bus segments and the overall parameters for the SCSI domain before allowing information transfers to resume. Once a bus reset event (see 12.5) is detected, SCSI initiator ports shall renegotiate with all SCSI target ports in the SCSI domain.

#### **E.8.4 Rule 3**

Homogeneous expanders between two bus segments shall support the maximum performance levels supported on each bus segment of the homogeneous expander. If one bus segment of the homogeneous expander is wide both bus segments of a homogeneous expander shall be wide. Both bus segments of the homogeneous expander shall support the same maximum transfer rate.

## E.8.5 Rule 4

### E.8.5.1 Effects of wired-or glitches

Wired-or glitches occur when two or more drivers are asserting the same signal line and one subsequently ceases to drive the line. This condition happens frequently during the ARBITRATION phase on the BSY signal and may happen on other wired-or signals. This change in the number of asserted drivers causes a redistribution of current in the segment, with resulting voltage glitches, and may cause false detection of a BUS FREE phase and other errors. The worst-case condition is when two SCSI device ports near a bus segment terminator are involved. In this case, after the SCSI device port stopped asserting the line, it requires a full bus segment length round trip time before the line is again stable. If this condition applies, the round trip time allowed is 400 ns. The one way time is 200 ns.

Waiting the entire domain round trip time may be avoided by ensuring that wired-or glitches do not pass through the simple expander. This standard does not describe how simple expanders implement this capability. If wired-or glitches are not propagated through simple expanders, then the maximum round trip SCSI domain signal propagation time is 800 ns and the one way SCSI domain propagation time is 400 ns.

If a simple expander does not implement the wired-or glitch filter it shall be labeled indicating that it allows propagation of wired-or glitches. This standard assumes that wired-or glitch blocking simple expanders are used and that the maximum SCSI domain round trip time of 800 ns is available.

### E.8.5.2 Simple expander propagation delay effects

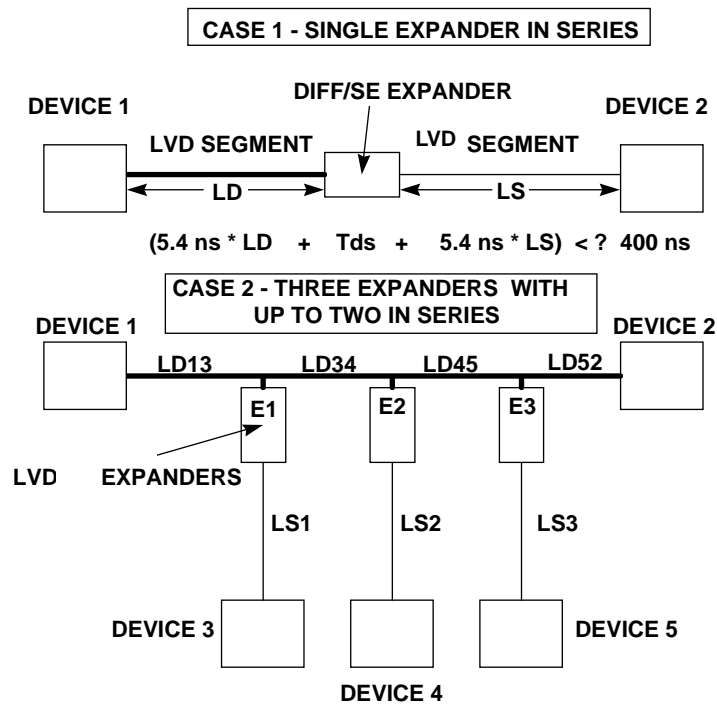
The simple expander is in series with SCSI initiator ports and/or SCSI target ports when the path between the SCSI initiator ports and/or SCSI target ports goes through an expander. In this case the propagation delay through the simple expander shall be counted as part of the 400 ns budget between those SCSI device ports.

The delay varies depending on the implementations. When two simple expanders are in series the delay across the pair may be much less than twice the individual delays. This is because the “direction” change that consumes much of the propagation delay during the ARBITRATION phase only applies to one of the simple expanders at a time. The single simple expander delay,  $T_{ds}$  and the expander series pair delay,  $T_{dp}$  should be specified.

If the simple expander is attached to a bus segment (e.g., as in case of the SCSI device enclosures in the bottom part of figure E.3) it is only in series between the SCSI device ports in the enclosure and other SCSI device ports in the SCSI domain.

The propagation time through the differential transceivers of SCSI initiator ports and SCSI target ports does not need to be separately accounted for if the wired-or glitches do not propagate through the simple expander. The differential transceiver delay effects are confined to the differential bus segments. Using simple expanders that do not pass the wired-or glitch prevents one bus segment’s delays from being passed on to the next.

## E.8.5.3 Sample calculations



**Figure E.5 - Two configurations for SCSI domain delay calculations**

Figure E.5 shows two sample SCSI domain configurations. In figure E.5 parameters whose first letter is “L” are physical lengths, “D” refers to differential bus segments and “S” refers to differential bus segments. In case 1 the delay calculations are shown in figure E.5. For the more complex case 2 all the possible combinations between any two SCSI device ports shall be considered. These calculations are shown in table E.2. The SCSI device port pair that has the largest combination of expander propagation time and interconnect propagation time determines if this configuration meets the 400 ns SCSI device port to SCSI device port maximum propagation time requirement.

Table E.2 - SCSI domain delay calculation

SCSI Device Port Pair	Path Between SCSI Device Ports	Simple Expanders Delay (ns)	Interconnect Delay (ns)
1-2	LD13,LD34,LD45,LD52	0	$5.4 \cdot (LD13 + LD34 + LD45 + LD52)$
1-3	LD13,E1,LS1	$T_{ds}$	$5.4 \cdot (LD13 + LS1)$
1-4	LD13,LD34,E2,LS2	$T_{ds}$	$5.4 \cdot (LD13 + LD34 + LS2)$
1-5	LD13,LD34,LD45,E3,LS3	$T_{ds}$	$5.4 \cdot (LD13 + LD34 + LD45 + LS3)$
2-3	LD52,LD45,LD34,E1,LS1	$T_{ds}$	$5.4 \cdot (LD52 + LD45 + LD34 + LS1)$
2-4	LD52,LD45,E2,LS2	$T_{ds}$	$5.4 \cdot (LD52 + LD45 + LS2)$
2-5	LD52,E3,LS3	$T_{ds}$	$5.4 \cdot (LD52 + LS3)$
3-4	LS1,E1,LD34,E2,LS2	$T_{dp}$	$5.4 \cdot (LS1 + LD34 + LS2)$
3-5	LS1,E1,LD34,LD45,E3,LS3	$T_{dp}$	$5.4 \cdot (LS1 + LD34 + LD45 + LS3)$
4-5	LS2,E2,LD45,E3,LS3	$T_{dp}$	$5.4 \cdot (LS2 + LD45 + LS3)$

**E.8.6 Rule 5**

Since simple expanders have no SCSI ID's the maximum number of addressable SCSI devices in the SCSI domain is not increased or decreased by the use of simple expanders.

**E.8.7 Rule 6**

Loop topologies in any form shall not be configured within a SCSI domain. Using simple expanders connected in a loop it is possible to create conditions where both a simple expander and a SCSI target port or SCSI initiator port are asserting the same line. Under these conditions the line shall not return to the negated state when the SCSI initiator port or SCSI target port releases the line since it shall continue to be driven by the simple expander. The logic state of the line does not change and a lock up condition exists.

Figures E.6 shows some examples of loops.




Figure E.6 - Examples of illegal loops

### E.9 Special performance considerations for SCSI domains with simple expanders

With simple expanders there is an additional consideration relating to the REQ/ACK OFFSET value. Since the round trip domain propagation time may be as large as 800 ns when using simple expanders, the REQ/ACK OFFSET value negotiated between any two SCSI device ports should be larger than used for SCSI domains that do not use expanders. If the REQ/ACK OFFSET is not sufficient to accommodate the round trip time between the SCSI device ports the domain experiences a performance degradation. This minimum REQ/ACK OFFSET value increases with increasing data transfer rates. The minimum REQ/ACK OFFSET value to avoid performance degradation for a variety of conditions are shown in table E.3.

The SCSI device port REQ/ACK OFFSET counter is set to zero before the DATA phase begins. When a REQ transition occurs the offset counter is incremented. When an ACK transition occurs the counter is decremented. After the DATA phase is completed the offset counter should again be at zero since the number of REQ transitions and the number of ACK transitions should be the same.

When the SCSI target port sends the first REQ transition there is a minimum of one round trip time before the first ACK transition may be received from the SCSI initiator port. This round trip time includes the data processing time at the SCSI initiator port. The SCSI target port may continue to issue REQ transitions until the offset counter in the SCSI target port reaches the maximum REQ/ACK OFFSET value that was negotiated.

If the maximum offset value in the SCSI target port is reached, the SCSI target port waits until it receives a decrementing ACK transition before issuing another REQ transition. The time spent waiting for a decrementing ACK transition is lost.

The receiving SCSI device port is required to accept up to at least the negotiated REQ/ACK OFFSET value of ACK or REQ transitions.

Negotiated REQ/ACK OFFSETS do not affect the operation of simple expanders.

The minimum recommended offset value for synchronous transfers is given by:

$$[\{2 \text{ times one way SCSI domain propagation time}\} / \{\text{ACK(REQ) period}\}] + \text{processing overhead}$$

The minimum recommended offset value for paced transfers is given by:

$$[\{2 \text{ times one way SCSI domain propagation time}\} / \{2 \times \text{ACK(REQ) period}\}] + \text{processing overhead}$$

Table E.3 gives some representative values from the minimum recommended offset value equation in this subclause for round trip SCSI domain propagation times greater than 400 ns assuming the processing overhead to be two ACK(REQ) periods in all cases.

**Table E.3 - Minimum REQ/ACK OFFSET for maximum performance**

SCSI domain round trip propagation time (ns)	DATA phase speed	ACK(REQ) period (nominal min)	Minimum REQ/ACK OFFSET to avoid performance degradation (assuming two overhead periods in all cases)
500	Fast-10	100	7
600	Fast-10	100	8
700	Fast-10	100	9
800	Fast-10	100	10
500	Fast-20	50	12
600	Fast-20	50	14
700	Fast-20	50	16
800	Fast-20	50	18
500	Fast-40	25	22
600	Fast-40	25	26
700	Fast-40	25	30
800	Fast-40	25	34
500	Fast-80	12.5	42
600	Fast-80	12.5	50
700	Fast-80	12.5	58
800	Fast-80	12.5	66
500	Fast-160	6.25	42 <sup>a</sup>
600	Fast-160	6.25	50 <sup>a</sup>
700	Fast-160	6.25	58 <sup>a</sup>
800	Fast-160	6.25	66 <sup>a</sup>
500	Fast-320	3.125	82
600	Fast-320	3.125	98
700	Fast-320	3.125	114
800	Fast-320	3.125	130
<sup>a</sup> The minimum REQ/ACK OFFSET for fast-160 is the same as fast-80 because the REQ/ACK OFFSET definition for paced transfers (see 4.12.4.4).			



## Annex F

(normative)

### Expander Communication Protocol

#### F.1 Introduction

This annex describes a method of expander communication and topology discovery called Expander Communication Protocol (ECP). This protocol permits application clients to detect expanders that support the protocol. It also permits the application client to pass parameter settings to expanders and permits expanders to report settings and status information. No new hardware features are required of SCSI device ports to implement this protocol.

ECP depends on the expander being able to monitor the data transfers associated with WRITE BUFFER and READ BUFFER commands (see SCSI Primary Commands-3 standard) and to alter specific portions of the data transferred as it passes through the expander in either direction. To simplify the expander implementation requirements, ECP is restricted to 8-bit asynchronous transfers.

For additional guidelines for using ECP see the SCSI Domain Validation technical report.

#### F.2 Glossary and Definitions

**F.2.1 Communicative expander:** A simple expander (see Annex E) that has the additional capability to support the requirements of this annex and is capable of transmitting information beyond that received on its ports to specific other entities in the domain. In this annex, unless stated otherwise, the term expander means communicative expander.

**F.2.2 Expander function signature:** A specific sequence of data bytes that identifies an ECP function in the data phase of a WRITE BUFFER or READ BUFFER command.

**F.2.3 Far port:** For the current I/O process, an expander port that is not the near port.

**F.2.4 First expander:** In a series expander set, the expander that couples the bus segment containing the SCSI initiator port to the next bus segment on the pathway to the SCSI target port.

**F.2.5 Last expander:** In a series expander set, the expander that couples the bus segment containing the SCSI target port to the next bus segment on the pathway to the SCSI initiator port.

**F.2.6 Near port:** For the current I/O process, the expander port connected directly to the SCSI initiator port through a bus segment or connected to the SCSI initiator port through other expanders and bus segments.

**F.2.7 Non-target port:** A far port that is not a SCSI target port (i.e., all far ports that do not include the SCSI target port for this I/O process).

**F.2.8  $n^{\text{th}}$  expander:** In a series expander set, the  $n^{\text{th}}$  expander between the SCSI initiator port and the SCSI target port.

**F.2.9 pathway:** The set of all bus segments and expanders between a SCSI initiator port and a SCSI target port.

**F.2.10 Series expander set:** The set of one or more expanders that couple the bus segment containing a SCSI initiator port to the bus segment containing a SCSI target port.

**F.2.11 SCSI target port:** For the current I/O process, a far port that is connected directly to the SCSI

target device port through a bus segment or connected to the SCSI target device port through other expanders and bus segments.

### F.3 Symbols and abbreviations

ECP	Expander Communication Protocol
SEDB	Short Expander Descriptor Block
LEDB	Long Expander Descriptor Block

### F.4 Enabling ECP

Following a hard reset (see 12.4), a communicative expander shall function as a simple expander for each SCSI initiator port until the SCSI initiator port enables ECP as follows:

- 1) negotiate the default transfer agreement with some SCSI target port; and
- 2) issue a WRITE BUFFER command to that same SCSI target device with the MODE field set to echo buffer plus enable expander communications protocol (see SCSI Primary Commands-3 standard).

The SCSI initiator port may disable ECP by:

- 1) negotiating to the default agreement with some SCSI target port; and
- 2) issuing a WRITE BUFFER command to that same SCSI target device with the MODE field bit set to disable expander communications protocol (see SCSI Primary Commands-3 standard).

This enabling and disabling of ECP is done on a SCSI initiator port basis (i.e., each SCSI initiator port issues a single WRITE BUFFER command to enable or disable ECP for that SCSI initiator port for all communicative expanders on the bus). The enabling or disabling of ECP occurs regardless of the SCSI device server's response to the WRITE BUFFER command.

While the SCSI initiator port has ECP enabled, it is responsible for not issuing WRITE BUFFER commands with the EXPANDER FUNCTION SIGNATURE in the first 7 bytes of the data buffer unless it is an expander function header (see F.5).

### F.5 Communicative expander function structures

Communicative expander functions consist of outbound and inbound functions. The outbound functions are contained in the data of a WRITE BUFFER command with the MODE field set to write data mode, echo buffer mode or echo buffer plus enable ECP mode. The inbound functions return information in the data of a READ BUFFER command with the mode field set to data mode, echo buffer mode or echo buffer plus enable ECP mode.

The SCSI initiator port shall not enable disconnects for these WRITE BUFFER and READ BUFFER commands (i.e., the DISCPRIV bit in the IDENTIFY message (see 16.3.3) is set to 0).

The outbound and inbound functions are further divided into multiple and single functions.

For multiple functions, the data is transferred in a 176-byte data structure consisting of a 16-byte expander function header followed by ten 16-byte SEDBs. For single functions, the data is transferred in a data structure whose length depends on the selected expander function code. This data structure consists of a 16-byte expander function header followed by one LEDB. In either case, the first 16 bytes of the data structure contain an expander function header as shown in table F.1.

**Table F.1 - Expander function header**

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) EXPANDER FUNCTION SIGNATURE (B7 3384 B850 8F27h) (LSB)							
6								
7	INITIATOR SCSI ADDRESS							
8	EXPANDER FUNCTION CODE							
9	(MSB) Function specific (LSB)							
15								

The EXPANDER FUNCTION SIGNATURE contains a code of B7 3384 B850 8F27h that signifies this WRITE BUFFER or READ BUFFER data is an expander function.

The application client shall set the INITIATOR SCSI ADDRESS field to the SCSI address of the SCSI initiator port through which the command is to be sent (i.e., the I of the I\_T nexus). The expander shall check that this field matches the SCSI address for the SCSI initiator port of the current I/O process.

If both the expander function signature is correct and the INITIATOR SCSI ADDRESS field matches the SCSI initiator port's SCSI address, then this WRITE BUFFER or READ BUFFER data is an expander function that shall be processed by the expander. Otherwise, the WRITE BUFFER or READ BUFFER command shall be ignored by the communicative expanders. In all cases the communicative expanders shall repeat the WRITE BUFFER or READ BUFFER data.

The EXPANDER FUNCTION CODES are described in table F.2.

The function specific bytes are documented for single functions in F.6.2 and F.6.4. The function specific bytes are not presently used for multiple functions.

**Table F.2 - Expander functions**

Type	Code	Expander function	Subclause
Outbound multiple function	00h	ASSIGN ADDRESS	F.6.1.2
	01h	MARGIN CONTROL	F.6.1.3
	02h - 2Fh	Reserved	
	30h - 3Fh	Vendor specific	
Outbound single function	40h	CONTROL	F.6.2.2
	41h - 6Fh	Reserved	
	70h - 7Fh	Vendor specific	
Inbound multiple function	80h	Reserved	
	81h	MARGIN REPORT	F.6.3.2
	82h	REPORT CAPABILITIES	F.6.3.3
	83h	REPORT CURRENT STATUS	F.6.3.4
	84h - AFh	Reserved	
	B0h - BFh	Vendor specific	
Inbound single function	C0h	EXPANDER INQUIRY	F.6.4.2
	C1h	REPORT SAVED TRAINING CONFIGURATION VALUES	F.6.4.3
	C2h - EFh	Reserved	
	F0h - FFh	Vendor specific	

For outbound and inbound multiple functions, the expander function header is followed by ten SEDBs (see table F.3).

For outbound single functions (see F.6.2) and inbound single functions (see F.6.4), the expander function header is followed by one long expander descriptor block.

**Table F.3 - SEDB format**

Bit Byte	7	6	5	4	3	2	1	0
0	USED	RESERVED				D_CLASS		
1	Function specific fields							
15								

The USED bit is defined in F.6.1.1, F.6.2.1, F.6.3.1, and F.6.4.1.

The D\_CLASS field identifies the device class that set the USED bit to one as defined in table F.4:

**Table F.4 - Device class**

Code	Device class
000b	Reserved
001b	Communicative expander device
010b	SCSI initiator port
011b - 111b	Reserved

The remaining fields in the short expander descriptor block are specific to the expander function and are defined in F.6.

## F.6 Expander functions

### F.6.1 Outbound multiple functions

#### F.6.1.1 Outbound multiple function data transfer rules

Outbound multiple functions shall be performed during a WRITE BUFFER command with the MODE field set to one of the values specified in F.5. The SCSI initiator port shall transfer an expander function header followed by ten SEDBs. The application client may use an SEDB to indicate the SCSI initiator port's parameters or settings. If the application client uses an SEDB, it shall use the first SEDB and shall set the USED bit in the first SEDB to 1. The application client shall set the USED bit to 0 in all unused SEDBs.

Each expander in the SCSI domain shall repeat the entire data structure without alteration to its SCSI target port, if any, except the expander shall alter the first SEDB encountered with a USED bit of 0. In this SEDB, the expander shall set the USED bit to 1, shall set the D\_CLASS field as described in table F.4, and shall output a zero in the reserved field of the first byte. The remaining 15 bytes of this SEDB shall be repeated without alteration. The expander shall interpret the other fields of this altered SEDB as defined in the ASSIGN ADDRESS (see F.6.1.2), and MARGIN CONTROL (see F.6.1.3) expander function codes.

Each expander in the domain shall either repeat the data received on the near port or shall repeat the data output to the SCSI target port on all far ports that are not part of the I\_T nexus.

An expander that receives a reserved or unimplemented vendor specific multiple expander function code shall follow all of the rules in this subclause, but shall ignore the contents of the altered SEDB.

#### F.6.1.2 ASSIGN ADDRESS

The ASSIGN ADDRESS expander function is used to assign an expander address to one or more expanders for this SCSI initiator port. The expander address is specific to the assigning SCSI initiator port. The expander address assigned by one SCSI initiator port has no affect on the expander addresses assigned by other SCSI initiator ports. The SEDB for this expander function is shown in table F.5.

**Table F.5 - ASSIGN ADDRESS SEDB**

Bit Byte	7	6	5	4	3	2	1	0
0	USED	RESERVED				D_CLASS		
1	ASSIGN	EXPANDER ADDRESS						
2	RESERVED							
15								

An ASSIGN bit of 1 indicates that the expander shall respond to the expander address specified in the EXPANDER ADDRESS field for single functions from this SCSI initiator port. The address assignment shall remain in effect until changed by another ASSIGN ADDRESS function from this SCSI initiator port or by a hard reset (see 12.4). An ASSIGN bit of 0 indicates that the expander shall not change its expander address assignment for this SCSI initiator port.

Assigning the expander address 0000000b to an expander shall indicate that it has no expander address assigned for this SCSI initiator port. The application client shall not use expander address 0000000b in single functions.

### **F.6.1.3 MARGIN CONTROL**

The MARGIN CONTROL expander function sets parameter settings in the SCSI initiator port or expander for usage between the I\_T nexus on subsequent synchronous transfers and paced transfers. These parameter settings shall remain in effect until changed by another MARGIN CONTROL expander function or by a hard reset (see 12.4).

The MARGIN CONTROL SEDB is shown in table F.6.

**Table F.6 - MARGIN CONTROL SEDB**

Bit Byte	7	6	5	4	3	2	1	0
0	USED	RESERVED				D_CLASS		
1	DRIVER STRENGTH (near port)				RESERVED			
2	DRIVER ASYMMETRY (near port)				DRIVER PRECOMPENSATION (near port)			
3	DRIVER SLEW RATE (near port)				RESERVED			
4	RESERVED				RESERVED			
5	RESERVED				RESERVED			
6	RESERVED				RESERVED			
7	VENDOR SPECIFIC (near port)				VENDOR SPECIFIC (near port)			
8	RESERVED							
9	DRIVER STRENGTH (far port)				RESERVED			
10	DRIVER ASYMMETRY (far port)				DRIVER PRECOMPENSATION (far port)			
11	DRIVER SLEW RATE (far port)				RESERVED			
12	RESERVED				RESERVED			
13	RESERVED				RESERVED			
14	RESERVED				RESERVED			
15	VENDOR SPECIFIC (far port)				VENDOR SPECIFIC (far port)			

Two duplicate sets of margin control fields (i.e., DRIVER STRENGTH, DRIVER ASYMMETRY, DRIVER PRECOMPENSATION, and DRIVER SLEW RATE) are provided, one set for the near port and another set for the far port. If the first SEDB is used for the SCSI initiator port settings, only the far port fields are used and the near port fields are reserved.

For a description of the DRIVER STRENGTH, DRIVER ASYMMETRY, DRIVER PRECOMPENSATION, and DRIVER SLEW RATE fields see 18.1.4.2.

## **F.6.2 Outbound single functions**

### **F.6.2.1 Outbound single function data transfer rules**

Outbound single functions shall be performed during a WRITE BUFFER command with the MODE field set to one of the values specified in F.5. The SCSI initiator port shall transfer an expander function header followed by a LEDB.

Each expander in the SCSI domain shall repeat the entire data structure without alteration to its SCSI target port, if any, except if the USED bit is 0 and the contents of the EXPANDER ADDRESS field in the first byte of the LEDB matches the expanders currently assigned expander address for this SCSI initiator port (see F.6.1.2). In this LEDB the expander shall change the USED bit to 1 and shall output its currently assigned expander address for this SCSI initiator port in the EXPANDER ADDRESS field. The expander shall interpret the other fields of this altered LEDB as defined in the CONTROL (see F.6.2.2) expander function code.

Each expander in the domain shall either repeat the data received on the near port or shall repeat the data output to the SCSI target port on all far ports that are not part of the I\_T nexus.

An expander that receives a reserved or unimplemented vendor specific single expander function code shall follow all of the rules in this subclause, but shall ignore the contents of the altered LEDB.

### F.6.2.2 CONTROL

The CONTROL function is used to set or clear parameters on the addressed expander. The function specific bytes in the expander function header are reserved for the CONTROL function. The LEDB for the CONTROL function is shown in table F.7.

**Table F.7 - CONTROL data structure**

Bit Byte	7	6	5	4	3	2	1	0
0	USED	EXPANDER ADDRESS						
1	TARGET_ADRS							
2	RESERVED					FAR_CTL		
3	RESERVED							
15								

The TARGET\_ADRS field shall be set to the SCSI address of a SCSI target port connected to one of the expander's far ports. This identifies the far port to be controlled on expanders that have multiple far ports. If the specified TARGET\_ADRS does not match a known SCSI target port SCSI address, then the expander shall perform no far port control action on any port.

The FAR\_CTL field is defined as shown in table F.8.



**Table F.8 - FAR\_CTL field values**

Code	Far port control action	Description
000b	No operation	Shall have no effect on the specified far port
001b	Disable far port	Shall cause the expander to stop repeating signals to the specified far port and shall cause the expander to ignore signals from the specified far port upon the next BUS FREE phase.
010b	Enable far port	Shall cause the expander to resume repeating signals to the specified far port and shall cause the expander to resume responding to signals from the specified far port upon the next BUS FREE phase.
100b	Reset far port	Shall cause the expander to create a bus reset condition (see 12.3) on the specified far port upon the next BUS FREE phase on the near port (i.e., the expander creates a pulse on the RST signal). The expander shall not propagate this bus reset condition to any other of its ports.
all other codes	Reserved	

### F.6.3 Inbound multiple functions

#### F.6.3.1 Inbound multiple function data transfer rules

The application client shall set the USED bit in each of the SEDBs to zero. The data structure containing an inbound multiple function is then placed in the SCSI target device's buffer using a WRITE BUFFER command with the MODE field set to one of the values specified in F.5. Expanders shall not alter the data structure during the WRITE BUFFER command if the EXPANDER FUNCTION CODE is 80h to FFh. The inbound multiple function is then performed during a subsequent READ BUFFER command with the MODE field set to one of the values specified in F.5.

During the data transfer phase of the READ BUFFER command, each expander with a SCSI target port shall repeat the entire data structure without alteration to its near port, except the expander shall alter the first SEDB encountered with a USED bit set to zero. In this SEDB, the expander shall change the USED bit to one, shall set the D\_CLASS field as described in table F.4, and shall output zero in the reserved field of the first byte. The remaining 15 bytes of this SEDB shall be output as described in the MARGIN REPORT (see F.6.3.2), REPORT CAPABILITIES (see F.6.3.3), and REPORT CURRENT STATUS (see F.6.3.4) expander function codes.

Each expander in the domain shall either repeat the data received on the SCSI target port or shall repeat the data output to the near port on all far ports that are not part of the I\_T nexus.

An expander that receives a reserved or unimplemented vendor specific multiple EXPANDER FUNCTION CODE shall follow all of the rules in this subclause, but shall output 00h in bytes 1 through 15 of the altered SEDB.

#### F.6.3.2 MARGIN REPORT

The MARGIN REPORT expander function is used to report the current margin settings for the SCSI initiator port or expander. The MARGIN REPORT SEDB is shown in table F.9.

**Table F.9 - MARGIN REPORT SEDB**

Bit Byte	7	6	5	4	3	2	1	0
0	USED	RESERVED				D_CLASS		
1	DRIVER STRENGTH (near port)				RESERVED			
2	DRIVER ASYMMETRY (near port)				DRIVER PRECOMPENSATION (near port)			
3	SLEW RATE (near port)				RESERVED			
4	RESERVED				RESERVED			
5	RESERVED				RESERVED			
6	RESERVED				RESERVED			
7	VENDOR SPECIFIC (near port)				VENDOR SPECIFIC (near port)			
8	RESERVED							
9	DRIVER STRENGTH (far port)				RESERVED			
10	DRIVER ASYMMETRY (far port)				DRIVER PRECOMPENSATION (far port)			
11	SLEW RATE (far port)				RESERVED			
12	RESERVED				RESERVED			
13	RESERVED				RESERVED			
14	RESERVED				RESERVED			
15	VENDOR SPECIFIC (far port)				VENDOR SPECIFIC (far port)			

Two duplicate sets of margin report fields (i.e., DRIVER STRENGTH, DRIVER ASYMMETRY, DRIVER PRECOMPENSATION, and SLEW RATE) are provided, one set for the near port and another set for the far port. If the last SEDB is used to communicate SCSI initiator port settings to the application client, only the far port fields are used while the near port fields are reserved. In this case, the initiator should set the USED bit to 1 in this SEDB before returning the data buffer to the application client.

The margin report fields shall return the current settings for the I\_T nexus. Fields that are not implemented shall be reported as 0000b. Otherwise, the current setting for the field, possibly rounded as described in F.6.1.3, shall be returned.

### **F.6.3.3 REPORT CAPABILITES**

The REPORT CAPABILITES function may be used to determine domain topology and report expander characteristics. The REPORT CAPABILITES SEDB is shown in table F.10.

**Table F.10 - REPORT CAPABILITIES SEDB**

Bit Byte	7	6	5	4	3	2	1	0
0	USED	RESERVED				D_CLASS		
1	(MSB) _____ FAR SCSI ID LIST _____ (LSB)							
2								
3	MINIMUM TRANSFER PERIOD FACTOR							
4	RESERVED							
5	MAXIMUM REQ/ACK OFFSET							
6	MAXIMUM TRANSFER WIDTH EXPONENT							
7	RESERVED	PROTOCOL OPTION BITS SUPPORTED						
8	PORTS				RESERVED			
9								
15	RESERVED _____							

The FAR SCSI ID LIST field contains the inclusive-or of all SCSI IDs known to be located on the SCSI target port of the expander (e.g., if SCSI devices with IDs 0, 1, and 12 were previously accessed on the SCSI target port, the expander sets this field to 1003h).

The MINIMUM TRANSFER PERIOD FACTOR field shall be set to the smallest value of the TRANSFER PERIOD FACTOR (see 4.12.4.3) supported by the expander.

The MAXIMUM REQ/ACK OFFSET field shall be set to the largest value of the REQ/ACK OFFSET (see 4.12.4.4) supported by the expander.

The MAXIMUM TRANSFER WIDTH EXPONENT field shall be set to the largest value of the TRANSFER WIDTH EXPONENT (see 4.12.4.5) supported by the expander.

The PROTOCOL OPTIONS BITS SUPPORTED field shall set the corresponding bit to one for each supported protocol option bit in byte 7 of the PPR message (see 4.12.4.6).

The PORTS field shall contain the number of ports on the expander not including the near port. A value of 0 in this field indicates that the expander is not capable of reporting this information.

#### **F.6.3.4 REPORT CURRENT STATUS**

The REPORT CURRENT STATUS function is used to report the negotiated settings of the near port transceiver and far port transceiver for the current I\_T nexus. The REPORT CURRENT STATUS SEDB is shown in table F.11.

**Table F.11 - REPORT CURRENT STATUS SEDB**

Bit Byte	7	6	5	4	3	2	1	0
0	USED	RESERVED				D_CLASS		
1	RESERVED				NEAR TRANSCEIVER MODE		NEAR SENT PCOMP_EN	NEAR RECEIVED PCOMP_EN
2	RESERVED				FAR TRANSCEIVER MODE		FAR SENT PCOMP_EN	FAR RECEIVED PCOMP_EN
3	RESERVED							
15								

The NEAR RECEIVED PCOMP\_EN bit and the FAR RECEIVED PCOMP\_EN contain the last received value for PCOMP\_EN bit on the corresponding ports. The NEAR SENT PCOMP\_EN bit and the FAR SENT PCOMP\_EN bit contain the last sent values for PCOMP\_EN bit on the corresponding ports. For SCSI initiator ports, only the far port values are defined. The values returned are from the most recent PPR negotiation that resulted in a synchronous or paced data transfer agreement.

The NEAR TRANSCEIVER MODE field and the FAR TRANSCEIVER MODE field specify the current bus mode for the corresponding ports as defined in table 86.

#### **F.6.4 Inbound single functions**

##### **F.6.4.1 Inbound single function data transfer rules**

The application client shall set the USED bit in the LEDB to 0. The data structure containing an inbound single function is then placed in the SCSI target device's buffer using a WRITE BUFFER command with the MODE field set to one of the values specified in F.5. Expanders shall not alter the data structure during the WRITE BUFFER command if the EXPANDER FUNCTION CODE is 80h to FFh. The inbound single function is then performed during a subsequent READ BUFFER command with the MODE field set to one of the values specified in F.5.

During the data transfer phase of the READ BUFFER command, each expander with a SCSI target port shall repeat the entire data structure without alteration to its near port, except if the USED bit is set to zero and the EXPANDER ADDRESS field in the first byte of the LEDB matches its currently assigned expander address (see F.6.1.2). In this LEDB, the expander shall change the USED bit to one and shall output its currently assigned expander address in the EXPANDER ADDRESS field. The expander shall output the remaining bytes of this LEDB as described for the EXPANDER INQUIRY (see F.6.4.2) and REPORT SAVED TRAINING CONFIGURATION VALUES (see F.6.4.3) expander function codes.

Each expander in the domain shall either repeat the data received on the SCSI target port or shall repeat the data output to the near port on all far ports that are not part of the I\_T nexus.

An expander that receives a reserved or unimplemented vendor specific single EXPANDER FUNCTION CODE shall follow all of the rules in this subclause, but shall output 00h in remaining bytes of the altered LEDB.

### F.6.4.2 EXPANDER INQUIRY

The EXPANDER INQUIRY function is used to report information about the specified expander in a manner similar to the SCSI INQUIRY command (see SCSI Primary Commands-3 standard). The expander function header for this function shall include function specific fields as described in table F.12.

**Table F.12 - EXPANDER INQUIRY expander function header**

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) EXPANDER FUNCTION SIGNATURE (B7 3384 B850 8F27h) (LSB)							
6								
7	INITIATOR SCSI ADDRESS							
8	EXPANDER INQUIRY FUNCTION CODE (C0h)							
9	RESERVED							EVPD
10	PAGE CODE							
11	RESERVED							
12	ALLOCATION LENGTH							
13								
15	RESERVED							

An enable vital product data (EVPD) bit of one specifies that the expander shall return the optional vital product data specified by the PAGE CODE field (see SCSI Primary Commands-3 standard). If the expander does not support the optional vital product data, then it shall return zero in all the bytes specified by the allocation length. A EVPD bit of zero specifies the expander shall return EXPANDER INQUIRY data as described in table F.13.

**Table F.13 - EXPANDER INQUIRY data**

Bit Byte	7	6	5	4	3	2	1	0
0	USED	ADDRESS						
1	RESERVED							
3								
4	ADDITIONAL LENGTH (33h)							
5	RESERVED							
7								
8	(MSB)	VENDOR IDENTIFICATION						
15								(LSB)
16	(MSB)	PRODUCT IDENTIFICATION						
31								(LSB)
32	(MSB)	PRODUCT REVISION LEVEL						
35								(LSB)
36	VENDOR SPECIFIC							
55								

The VENDOR IDENTIFICATION, PRODUCT IDENTIFICATION, and PRODUCT REVISION LEVEL fields shall return the STANDARD INQUIRY data format (see SCSI Primary Commands-3 standard).

#### **F.6.4.3 REPORT SAVED TRAINING CONFIGURATION VALUES**

The REPORT SAVED TRAINING CONFIGURATION VALUES function is used to report the SCSI device port's saved training configuration values. These vendor specific values are maintained by the SCSI device port when the retain training information option is enabled (see 4.12.4.6.8).

For expanders implementing this function, both near port and far port values shall be included. For SCSI initiator ports implementing this function, only the far port values shall be included. The near port values are used during DATA OUT phases, while the far port values are used during DATA IN phases.

Only values for the current I\_T nexus are reported.

The meaning of the fields is vendor specific.

The data structure for this function shall include functions specific fields described in table F.14.

Table F.14 - REPORT SAVED TRAINING CONFIGURATION VALUES data structure (Sheet 1 of 2)

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	EXPANDER FUNCTION SIGNATURE						(LSB)
6		(B7 3384 B850 8F27h)						
7		INITIATOR SCSI ADDRESS						
8		REPORT SAVED TRAINING CONFIGURATION VALUES FUNCTION CODE (C1h)						
9		RESERVED						
11								
12	(MSB)	NEAR PORT DB(0) VALUE						(LSB)
15								
		.						
72	(MSB)	NEAR PORT DB(15) VALUE						(LSB)
75								
76	(MSB)	NEAR PORT P_CRCA VALUE						(LSB)
79								
80	(MSB)	NEAR PORT P1 VALUE						(LSB)
83								
84	(MSB)	NEAR PORT BSY VALUE						(LSB)
87								
88	(MSB)	NEAR PORT SEL VALUE						(LSB)
91								
92	(MSB)	NEAR PORT RST VALUE						(LSB)
95								
96	(MSB)	NEAR PORT REQ VALUE						(LSB)
99								
100	(MSB)	NEAR PORT ACK VALUE						(LSB)
103								
104	(MSB)	NEAR PORT ATN VALUE						(LSB)
107								
108	(MSB)	NEAR PORT C/D VALUE						(LSB)
111								
112	(MSB)	NEAR PORT I/O VALUE						(LSB)
115								
116	(MSB)	NEAR PORT MSG VALUE						(LSB)
119								
120		RESERVED						
123								

Table F.14 - REPORT SAVED TRAINING CONFIGURATION VALUES data structure (Sheet 2 of 2)

Bit Byte	7	6	5	4	3	2	1	0	
124	(MSB)	FAR PORT DB(0) VALUE							(LSB)
127									
		.							
		.							
		.							
184	(MSB)	FAR PORT DB(15) VALUE							(LSB)
187									
188	(MSB)	FAR PORT P_CRCA VALUE							(LSB)
191									
192	(MSB)	FAR PORT P1 VALUE							(LSB)
195									
196	(MSB)	FAR PORT BSY VALUE							(LSB)
199									
200	(MSB)	FAR PORT SEL VALUE							(LSB)
203									
204	(MSB)	FAR PORT RST VALUE							(LSB)
207									
208	(MSB)	FAR PORT REQ VALUE							(LSB)
211									
212	(MSB)	FAR PORT ACK VALUE							(LSB)
215									
216	(MSB)	FAR PORT ATN VALUE							(LSB)
219									
220	(MSB)	FAR PORT C/D VALUE							(LSB)
223									
224	(MSB)	FAR PORT I/O VALUE							(LSB)
227									
228	(MSB)	FAR PORT MSG VALUE							(LSB)
231									
232		RESERVED							
235									

## F.7 Data Transfer Requirements

The communicative expander functions shall only be performed when the default transfer agreement is in effect. For any other transfer agreement, the communicative expander shall operate as a simple expander.

When altering data, communicative expanders shall construct correct parity for the altered data on the outgoing port.



## **Annex G**

(informative)

### **Interconnecting bus segments of different widths**

A problem may occur when mixing devices with ports that support this standard on the parallel SCSI bus ports with SCSI-2 device ports. The TERMPWR requirements (see table 30) of SPI have been increased to support a 16-bit DATA BUS. SCSI-2 devices may not supply sufficient TERMPWR. An additional source of TERMPWR (e.g., a target device) may be required.

When bus segments of dissimilar width are adapted to one another as shown in figure G.1 for LVD, the DATA BUS signals from the wider of the two bus segments that end at the adapter should be terminated at the adapter. The connectors are designed such that A and P shielded connectors do not intermate directly.

Two of the RESERVED lines (i.e., A cable contact numbers 23 and 24) and the OPEN line (i.e., A cable contact number 25) on the A cable are TERMPWR lines on the P cable (i.e., P cable contact numbers 33, 34, and 35).

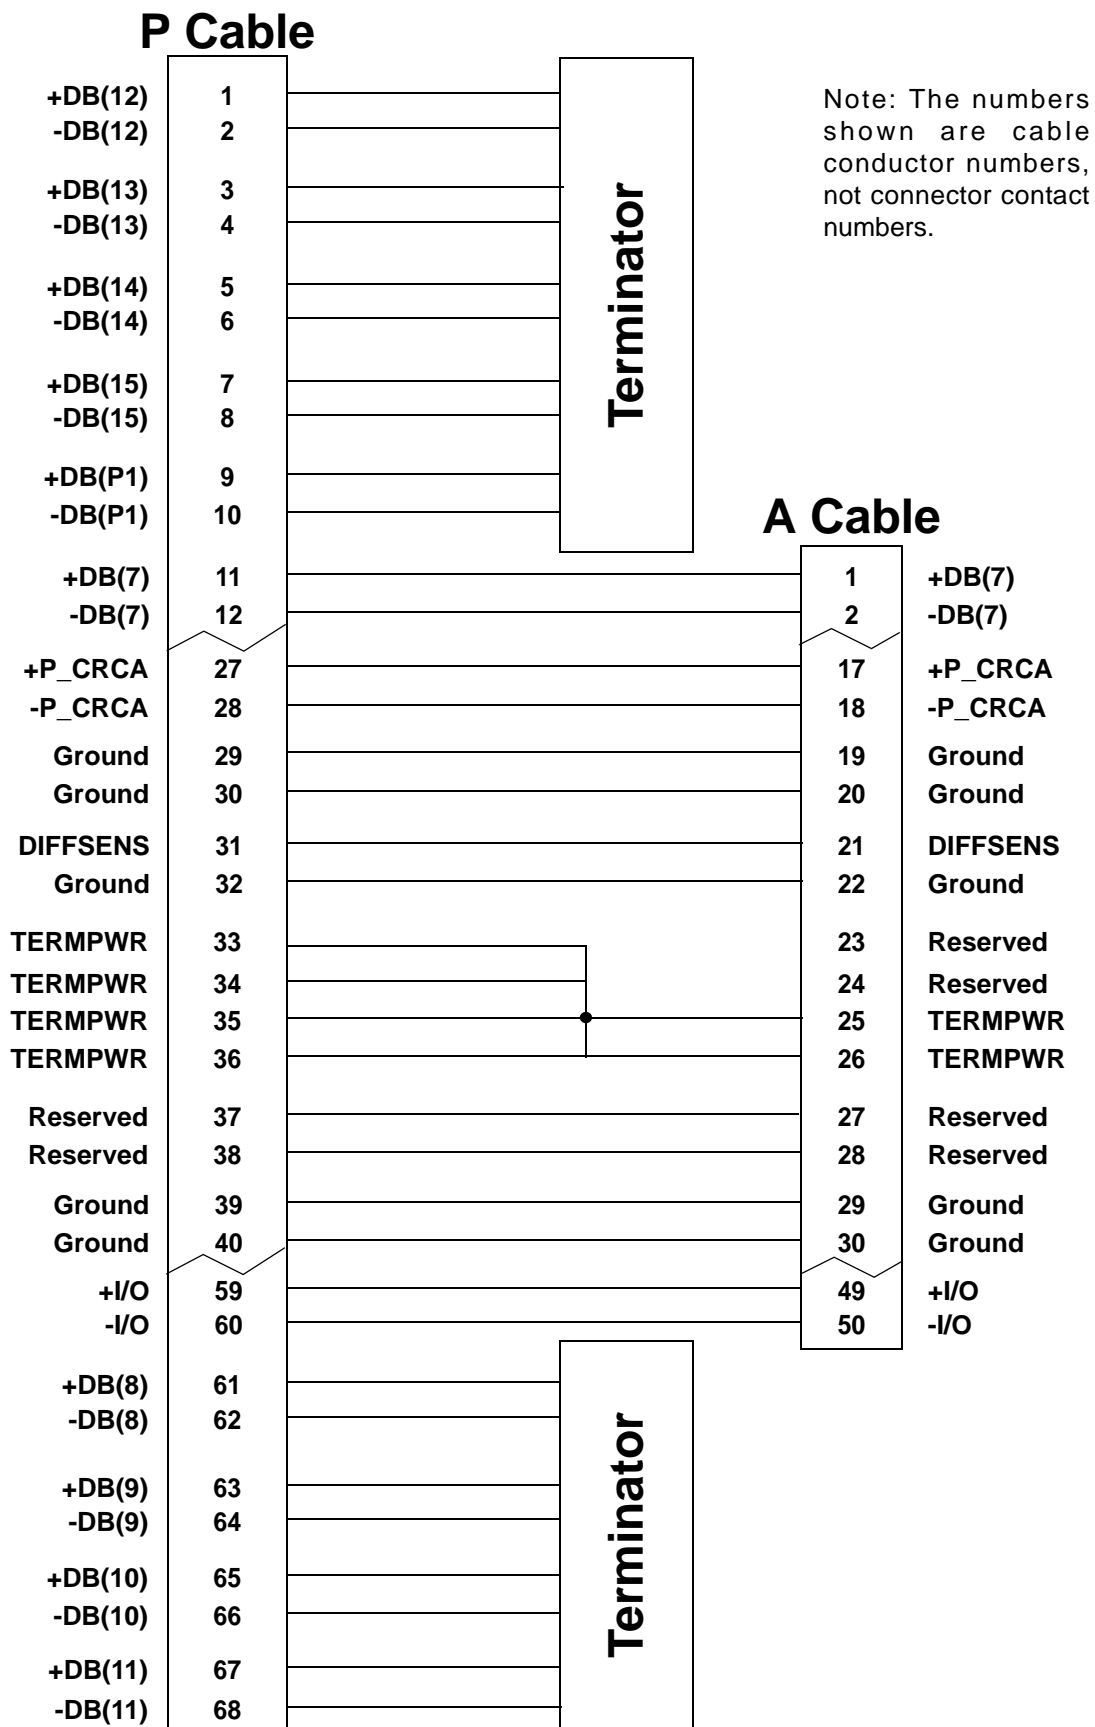


Figure G.1 - Interconnecting LVD A and P cables

## Annex H

(informative)

### SCSI ICONS

These icons are provided as symbols to identify a SCSI port and to indicate whether the port is using LVD transceivers (see figure H.1).

The icons illustrated in figure H.1 may be enlarged or reduced as needed for the application. The text and graphic may be used together or separately. The text font and size may also be adjusted as required.




Figure H.1 - LVD icon for SCSI

## Annex I

(informative)

### Backplane construction guidelines


#### I.1 Universal backplane construction

##### I.1.1 Universal backplane construction overview

Printed circuit boards are constructed using either microstrip or stripline or a combination of both for routing signals. The important electrical characteristics may be determined from the geometry and the material properties of the microstrip or stripline.

##### I.1.2 Microstrip

See figure I.1 for the microstrip geometry.



**Figure I.1 - Microstrip geometry**

The characteristic impedance ( $Z_0$ ) for the microstrip geometry shown in figure I.1 is defined by the following equation:

$$Z_0 = \frac{87}{(e_r + 1, 41)^{1/2}} \times \ln\left(\frac{5, 98h}{0, 83w + t}\right) \Omega$$


The propagation delay ( $T_{pd}$ ) for the microstrip geometry shown in figure I.1 is defined in the following equation:

$$T_{pd} = 3, 336(0, 475e_r + 0, 67)^{1/2} \text{ ps/mm}$$

Where  $e_r$  is equivalent to the relative dielectric constant.

##### I.1.3 Embedded Microstrip

See figure I.2 for the embedded microstrip geometry.



**Figure I.2 - Embedded microstrip geometry**

The characteristic impedance ( $Z_0$ ) for the embedded microstrip geometry shown in figure I.2 is defined by the following equation:

$$Z_0 = \frac{K}{(0,805e_r + 2)^{1/2}} \times \ln\left(\frac{5,98h}{0,8w + t}\right) \Omega$$

The propagation delay ( $T_{pd}$ ) for the microstrip geometry shown in figure I.2 is defined by the following equation:


$$T_{pd} = 3,336(0,475e_r + 0,67)^{1/2} \text{ ps/mm}$$

Where:

- a)  $e_r$  is equivalent to the relative permittivity; and
- b)  $K$  is a constant between 60 and 65 that varies with the thickness of the dielectric covering the conductor. If the thickness is below 380  $\mu\text{m}$   $K = 65$  and at 508  $\mu\text{m}$   $K = 60$ .

#### I.1.4 Stripline

See figure I.3 for the stripline geometry.



**Figure I.3 - Stripline geometry**

The characteristic impedance ( $Z_0$ ) for the stripline geometry shown in figure I.3 is defined by the following

equation:

$$Z_0 = \frac{60}{(e_r)^{1/2}} \times \ln \left( \frac{4h}{0, 67\pi w(0, 8 + \frac{t}{w})} \right) \Omega$$


The propagation delay ( $T_{pd}$ ) for the stripline shown in figure I.3 is defined by the following equation:

$$T_{pd} = 3, 336(e_r)^{1/2} \text{ ps/mm}$$

Where  $e_r$  is equivalent to the relative permittivity.

### I.1.5 Dual Stripline

See figure I.4 for the dual stripline geometry.



**Figure I.4 - Dual stripline geometry**

The characteristic impedance ( $Z_0$ ) for the dual stripline geometry shown in figure I.4 is defined by the following equations:

$$Z_0 = \left( \frac{2F_1 F_2}{F_1 + F_2} \right) \Omega$$

Where:

$$a) \quad F_1 = \frac{60}{(e_r)^{1/2}} \times \ln \left( \frac{8a}{0, 67\pi w(0, 8 + \frac{t}{w})} \right) \Omega; \text{ and}$$

$$b) F_2 = \frac{60}{(e_r)^{1/2}} \times \ln \left( \frac{8(a+d)}{0,67\pi w(0,8 + \frac{t}{w})} \right) \Omega.$$


The propagation delay ( $T_{pd}$ ) for the stripline shown in figure I.3 is defined by the following equation:

$$T_{pd} = 3,336(e_r)^{1/2} \text{ ps/mm}$$

Where  $e_r$  is equivalent to the relative permittivity.

### I.1.6 Differential Impedance

The lossless model (i.e., doesn't include shunt conductance and series resistance) for the differential transmission line contains a series inductance, mutual inductance, capacitance to ground and a differential capacitance (see figure I.5).



**Figure I.5 - Lossless model for differential impedance**

The two conductors are symmetrical therefore the two self inductance's are equal and the two capacitors to ground are equal ( $Z_{11} = Z_{21}$ ). The voltage applied between the two lines may be thought of as a superposition of two voltages, a common mode voltage  $V_c$  (even mode) and a differential mode voltage  $V_d$  (odd mode) where:  $V_c = (V_1 + V_2)/2$  and  $V_d = (V_1 - V_2)/2$ .

When a differential voltage is applied, equal and opposite currents flow and the voltage drop across the mutual inductance is in a direction opposite that of the self inductance. In the differential voltage case the mutual inductance acts to reduce the inductance seen by the differential mode signal to  $L_s - L_m$ . The signals are moving in opposite directions therefore the effect of the mutual capacitance  $C_d$  is doubled. The total capacitance seen by the differential signal is  $C + C_d$ . For this case the odd mode impedance ( $Z_d$ ) is calculated by the following equation:

$$Z_d = \frac{V_d}{I_1} = \left( \frac{L_s - L_m}{C + C_d} \right)^{1/2} \Omega$$

When a common mode voltage source drives the bus segment the current flows in the same direction the odd mode impedance ( $Z_c$ ) is calculated by the following equation:

$$Z_c = \frac{V_c}{I_1} = \left( \frac{L_s + L_m}{C - C_d} \right)^{1/2} \Omega$$

The differential mode impedance is two times the odd mode impedance ( $Z_{diff} = 2Z_d$ ).

The differential impedance may also be derived from the matrix of a two port network as shown in the following equation:

$$\begin{bmatrix} V_1 \\ V_2 \end{bmatrix} = \begin{bmatrix} Z_{11} & Z_{12} \\ Z_{21} & Z_{22} \end{bmatrix} \times \begin{bmatrix} I_1 \\ I_2 \end{bmatrix}$$

Because in differential mode  $I_1 = -I_2$  the matrix of a two port network the differential impedance ( $Z_d$ ) is calculated by the following equation:

$$Z_d = \frac{V_d}{I_1} = \frac{V_1 - V_2}{I_1} \Omega$$

Therefore, from the matrix for the two port network the following equation is derived:

$$\frac{V_1 - V_2}{I_1} = (Z_{11} + Z_{22}) - (Z_{12} - Z_{21}) \Omega$$

For a passive black box  $Z_{12} = Z_{21}$ . For a symmetrical system  $Z_{11} = Z_{21}$ . Therefore the differential impedance may be defined as  $Z_{diff} = 2(Z_{11} - Z_{12}) \Omega$  and the common mode impedance may be defined as  $Z_{cm} = 2(Z_{11} + Z_{12}) \Omega$ .

If the traces are loosely coupled then  $Z_{12}$  is negligible.

### I.1.7 Single ended impedance

The odd mode impedance approaches the single ended impedance as the trace spacing becomes larger than the height above the ground plane. A practical rule for zero coupling between traces is; if the trace to trace spacing is three times the trace to ground plane spacing then the trace to trace coupling is negligible.

If  $V_2$  in figure I.5 is grounded then the single ended impedance is determined by the self inductance and the sum of the differential capacitance plus the single ended capacitance. In this case the single ended impedance ( $Z_0$ ) is calculated by the following equation:

$$Z_0 = \left( \frac{L_s}{C} \right)^{1/2} \Omega$$

### I.1.8 Differential stripline

For differential stripline there are two choices for routing, edge coupled (side by side) figure I.6 and figure I.7, and broadside coupled (stacked) figure I.8.






Figure I.6 - Edge coupled differential microstrip




Figure I.7 - Edge coupled differential stripline




Figure I.8 - Broadside coupled differential stripline

Edge coupled is the most common. Broadside has better coupling characteristics but may be difficult to bring to a common plane and requires a thicker dielectric to keep from coupling to one of the reference planes. Therefore this annex does not define the broadside coupled differential stripline impedance characteristics.

The differential impedance is determined primarily by the conductor spacing and dielectric thickness. First the single-ended impedance ( $Z_0$ ) is calculated (see I.1.7) then it is used along with conductor spacing and dielectric thickness to calculate the differential impedance. The calculations assume the traces of the differential pairs have the same physical cross section dimensions.

The practical range of  $Z_0$  and  $Z_{diff}$  is from 20  $\Omega$  to 150  $\Omega$  with a typical range being between 50  $\Omega$  and 110  $\Omega$ . This impedance depends on the trace width and distance to ground. If the trace is wide and close to ground it is more capacitive and has a lower impedance. If the trace is narrow and the spacing from the ground plane is large, the trace is more inductive and has a higher impedance. Controlled impedance boards with the impedance's matching within several Ohms usually have a characteristic impedance in the 50  $\Omega$  to 80  $\Omega$  range. This is due to manufacturing constraints such as maximum dielectric thickness and minimum trace widths.

Using figure I.6 as an example the edge coupled microstrip differential impedance ( $Z_{diff}$ ) is calculated by the following equation:

$$Z_{diff} = 2Z_0(1 - 0,48e^{-(0,96s)/h})\Omega$$

Using figure I.7 as an example the edge coupled stripline differential impedance ( $Z_{diff}$ ) is calculated by the following equation:

$$Z_{diff} = 2Z_0(1 - 0,374e^{(-2,9s/h)})\Omega$$

The trace capacitance and inductance are important parameters for calculating the transmission line parameters. The method used to calculate these parameters is as follows.

The edge coupled microstrip trace capacitance ( $C_0$ ) and inductance ( $L_0$ ) are calculated by the following equations:

$$C_0 = \frac{T_d}{Z_0} = \frac{3,336(0,475e_r + 0,67)^{1/2}}{Z_0} \text{ pF/mm}$$

$$L_0 = Z_0 T_d = 3,336Z_0(0,475e_r + 0,67)^{1/2} \times 10^{-9} \text{ nH/mm}$$

Where:

- a)  $T_d$  is the propagation delay; and
- b)  $e_r$  is equivalent to the relative permittivity.

The edge coupled stripline trace capacitance ( $C_0$ ) and inductance ( $L_0$ ) are calculated by the following equations:

$$C_0 = \frac{T_d}{Z_0} = \frac{3,336(e_r)^{1/2}}{Z_0} \text{ pF/mm}$$

$$L_0 = Z_0 T_d = 3,336 Z_0 (e_r)^{1/2} \times 10^{-3} \quad \text{nH/mm}$$

Where:

- a)  $T_d$  is the propagation delay; and
- b)  $e_r$  is equivalent to the relative permittivity.

As demonstrated in this subclause the characteristic impedance for the backplane is primarily determined by:

- a) Width and thickness of the conductors;
- b) dielectric constant of the substrate material; and
- c) the substrate material thickness between the conductor and reference planes.

### I.1.9 Dielectric material selection

Proper selection of the dielectric material is very important for high speed PCB's. Two key parameters are the dielectric constant of the material and the loss tangent. The dielectric constant relates to the material's ability to hold charge and the loss tangent refers to how much of the energy is lost in the material due to dissipation. The ideal materials have small numbers. Table J.1 gives a sample of some materials.

**Table I.1 - Dielectric constants**

Material	Dielectric Constant	Loss Tangent
Air	1,0	0
PTFE (polytetrafluoroethylene)	2,1 - 2,5	0,0002 - 0,002
BT resin	2,9 - 3,9	0,003 - 0,12
Polyimide	2,8 - 3,5	0,004 - 0,02
Silica	3,8 - 4,2	0,0006 - 0,005
Polyimide/Glass	3,8 - 4,5	0,003 - 0,01
Epoxy/Glass (FR-4)	4,1 - 5,3	0,002 - 0,02

The most frequently used dielectric are a glass-epoxy (G-10) and a derivative, FR-4. The FR-4 material has acceptable performance for differential.

For higher speeds, materials like PTFE should be considered, although they are much more expensive. PCB manufacturers publish specifications with their boards, among them should be the dielectric constant, loss tangent, and other electrical properties. The board tolerances should also be specified. For example in a FR-4 PCB the dielectric constant may change by as much as 10 % on a single board. These changes affect the propagation velocity and lead to skew.

### I.1.10 Vias

Vias have a parasitic capacitance and inductance that may affect signal quality. The size and density of vias also affects how many traces that may be run between them on inner layers.

The parasitic capacitance (C) to ground of a via is calculated by the following equation:

$$C = \frac{3,58 \epsilon_r T D_1}{D_2 - D_1} \text{ pF}$$

Where:

- a)  $D_2$  is the diameter of the clearance hole in mm;
- b)  $D_1$  is the diameter of pad surrounding the via in mm; and
- c)  $T$  is the thickness of the printed circuit board in mm.

A typical via capacitance is in the 0,2 pF to 0,8 pF range. For critical differential signals the via mismatch should be kept to a minimum.

The parasitic inductance (L) of a via is calculated by the following equation:

$$L = (129h) \left( \ln \left( \frac{4h}{d} \right) + 1 \right) \text{ nH}$$

Where:

- a)  $h$  is the length of the via in mm; and
- b)  $d$  is the diameter of the via in mm.

The via inductance is most detrimental when they are in series with termination power bypass capacitors.

The goal of the universal backplane is to have the LVD backplane impedance to match the LVD backplane impedance of the cable. One of the aspects of this is to determine if the PCB traces are long enough to require the traces to be treated as transmission lines. If the signal electrical length is greater than 1/2 of the rising edge ( $T_{\text{rise}}$ ), then the PCB exhibits transmission line effects. This length may be expressed as:

$$\text{Length} = \frac{T_{\text{rise}}}{20(LC)^{1/2}} \text{ mm}$$

If the round trip time for the switching waveform is greater than the rise or fall time of the driving SCSI transceiver, the settling of the transmission line effects are not hidden during the rise and fall time of the driving SCSI transceiver. In other words, in order to be a transmission line  $2xT_{\text{pd}} > T_{\text{rise}}$  or  $T_{\text{fall}}$  (the minimum of the two) where  $T_{\text{pd}}$  is the one way propagation delay.

Specifically for FR-4 the two following formulas may be used.

- a)  $BW > 23,6/d$  where  $d$  is in mm and the bandwidth is  $BW = 0,35/T_{\text{rise}}$ ; and
- b)  $F_{\text{clock}} > 4,7/d$ .

Complete formulas for the maximum length without termination or controlled impedance may also be used.

The microstrip maximum length ( $\text{Length}_{\text{max}}$ ) is calculated as shown in the following equation:

$$\text{Length}_{\text{max}} = \frac{((C + Z_0)^2 + (12,3 \times 10^6))^{1/2} - (C_t Z_0)}{6,67(0,475 \epsilon_r + 0,67)^{1/2}} \text{ mm}$$

The stripline maximum length ( $\text{Length}_{\max}$ ) is calculated as shown in the following equation:

$$\text{Length}_{\max} = \frac{((C + Z_0)^2 + (12, 3 \times 10^6))^{1/2} - (C_t Z_0)}{6, 67(e_r)^{1/2}} \text{mm}$$


Where:

- a)  $Z_0$  is the unloaded impedance;
- b)  $C_t$  is the total load capacitance;
- c)  $e_r$  is equivalent to the relative permittivity; and
- d)  $C$  is the unloaded characteristic capacitance.

Depending on how conservative the calculation is PCB lengths in the four inch range or greater should probably be treated as transmission lines for differential risetimes in the 1 ns range.

### I.1.11 Transmission Lines

A typical transmission line element (see figure I.9) is broken into four parts, a series resistance, series inductance, shunt conductance, and shunt capacitance.



**Figure I.9 - Typical transmission line element**

When using this model the minimum wavelength should be much longer than the individual section to make it appear as a distributed model. The non-linearities are usually ignored. The general form for the impedance ( $Z_0$ ) is shown in the following equations:

$$Z_0 = \left( \frac{10l(R + jwL)}{2} + \frac{1}{2} \left( (10l)^2 (R + jwL)^2 + 4 \left( \frac{R + jwL}{G + jwC} \right) \right)^{1/2} \right) \Omega$$

or

$$Z_0 = \left( \frac{10l(R + jwL)}{2} - \frac{1}{2} \left( (10l)^2 (R + jwL)^2 + 4 \left( \frac{R + jwL}{G + jwC} \right) \right)^{1/2} \right) \Omega$$

Where:

- a)  $j$  denotes the complex number  $j \times j = -1$ ;
- b)  $w$  is omega or  $2 \times \pi \times$  frequency; and
- c)  $l$  is the section length in mm.

If you make the section length ( $l$ ) of the segment small enough, the result is the distributed model. Where the impedance ( $Z_0$ ) is calculated as shown in the following equation:

$$Z_0 = \left( \frac{R + j\omega L}{G + j\omega C} \right)^{1/2} \Omega$$

If the transmission rate is high (i.e.,  $\omega/2\pi > 100$  KHz) then  $\omega L$  and  $\omega C$  are much larger than  $R$  and  $G$  and the impedance becomes the more commonly used form of;  $Z_0 = (L/C)^{1/2} \Omega$ .

The other extreme is if  $\omega/2\pi \leq 1$  KHz then;  $Z_0 = (R/G)^{1/2} \Omega$ .

Another important factor is the propagation velocity and propagation delay. The propagation delay per unit length ( $T_{\text{delay}}$ ) is calculated as shown in the following equation;

$$T_{\text{delay}} = \frac{T_1}{10^9} [(R + j\omega L) + (G + j\omega C)]^{1/2} \text{ ns/mm}$$

However, if the line is short then  $T_{\text{delay}}$  is calculated as shown in the following equation:

$$T_{\text{delay}} = j\omega l (LC)^{1/2} \text{ ns}$$

Therefore the propagation velocity is  $V = (10^9)/(LC)^{1/2} \text{ mm/ns}$  for lossless lines.

The time delay is  $T = 1/V = (LC)^{1/2} \text{ ns}$  that gives a total propagation delay of  $T = 0, 1/(LC)^{1/2} \text{ ns/mm}$ . The two equations most commonly used for hand calculations are:

- a) for characteristic impedance:  $Z = (L/C)^{1/2} \Omega$ ; and
- b) for propagation delay:  $T = 0, 1/(LC)^{1/2} \text{ ns/mm}$ .

If these values are not specified they may be calculated from the cross section geometry and dielectric material used for the PCB. The equation using the geometry and material are based on the characteristic impedance and propagation delay equations, but should also be treated as if they were operating in the transverse electromagnetic mode.

Another factor that should be taken into account is attenuation due to the skin effect. The skin depth is;

$$\delta = 500(\rho/f)^{1/2} \text{ m.}$$

Where:

- a)  $\rho$  resistivity in  $\Omega/\text{m}$ ;
- b)  $f$  frequency in hertz; and
- c)  $\delta$  thickness in m.

The low loss attenuation per length is;  $\alpha = 4, 34 \left( \frac{R_1}{Z_0} + G_1 Z_0 \right) \text{ dB/length}$ .

The loss due to the dielectric is;  $\alpha = 0,09f(\tan(\delta) \times e_r^{1/2})$  dB/mm where f is in GHz.

The loss due to the metal for  $t < \delta$  is;  $\alpha = 17,1\rho / (100twZ_0)$  dB/mm where:

- a)  $\rho$  resistivity in  $\Omega/\text{m}$ ;
- b)  $t$  thickness in mm;
- c)  $w$  the width in mm; and
- d)  $Z_0$  the single ended impedance.

The loss due to the metal for  $t > \delta$  is;  $\alpha = \frac{0,22(\rho f)^{1/2}}{10w}$  dB/mm where:

- a) f frequency in GHz;
- b)  $\rho$  resistivity in  $\Omega/\text{m}$ ; and
- c)  $w$  the width in mm.

When calculating the impedance and propagation delay, the capacitive loads have to be accounted for. Capacitive loading decreases the impedance and increases the propagation delay.

The loaded propagation delay ( $T'_{pd}$ ) is calculated as shown in the following equation:

$$T'_{pd} = T_{pd} \left( 1 + \left( \frac{C_D}{C_0} \right) \right)^{1/2} \text{ ns}$$

or

$$T'_{pd} = (L(C_D + C_0))^{1/2} \text{ ns}$$

The loaded characteristic impedance ( $Z'_0$ ) is calculated as shown in the following equation:

$$Z'_0 = \frac{Z_0}{\left( 1 + \left( \frac{C_D}{C_0} \right) \right)^{1/2}} \Omega$$

or

$$Z'_0 = \left( \frac{L}{C_D + C_0} \right)^{1/2} \Omega$$

Where:

- a)  $C_0$  is the intrinsic capacitance;
- b)  $C_D$  is the load capacitance;
- c)  $Z_0$  is the single ended impedance; and
- d)  $T_{pd}$  is the propagation delay.

Besides reducing the impedance and increasing the propagation delay, a heavily loaded trace also slows the rise and fall times of the drivers and filters (RC filter) out some high frequency components. The loaded propagation delay ( $T'_{pd}$ ) should be used when deciding whether or not to treat the trace as a transmission line.

When adding the load capacitance; sockets, connectors, vias, and IC's add to the distributed capacitance. Using traces with a higher intrinsic capacitance reduce the effects of the loading (e.g., microstrip is faster than stripline, but is affected more by loading since it has a lower characteristic capacitance.)

The complete forms of loaded propagation delay ( $T'_{pd}$ ) and impedance ( $Z'_0$ ) for microstrip are calculated as shown in the following equations:

$$T'_{pd} = 0,5776(15,85e_r + 25,35 + C_D Z_0(0,475e_r + 0,67)^{1/2})^{1/2} \text{ ps/mm}$$

$$Z'_0 = \frac{Z_0}{\left( \frac{C_D Z_0}{33,36(0,475e_r + 0,67)^{1/2}} + 1 \right)^{1/2}} \Omega$$

The complete forms of loaded propagation delay ( $T'_{pd}$ ) and impedance ( $Z'_0$ ) for stripline are calculated as shown in the following equations:

$$T'_{pd} = 0,5776(33,36e_r + C_D Z_0(e_r)^{1/2})^{1/2} \text{ ps/mm}$$

$$Z'_0 = \frac{Z_0}{\left( \frac{C_D Z_0}{33,36(e_r)^{1/2}} + 1 \right)^{1/2}} \Omega$$

The impedance mismatches between loads, sources, connectors, cables, and traces may cause transmission line effects such as ringing, stair step effects, and long bus settle delays. These are caused by reflections at the impedance discontinuities, the reflection coefficients are:

a) at load:  $\Gamma_1 = \frac{Z_1 - Z_0}{Z_1 + Z_0}$  and

b) at source  $\Gamma_s = \frac{Z_s - Z_0}{Z_s + Z_0}$ .

These reflections, if large enough, may cause false transitions. They may also cause standing waves.

$$VSWR = \frac{1 + \Gamma}{1 - \Gamma}$$

Where VSWR is the voltage standing wave ratio.

The standing waves are at quarter wavelengths ( $\lambda/4$ ) where the wavelength is calculated by the following equation:

$$\lambda = \frac{1}{f(LC)^{1/2}}$$

To capture all the harmonic content of the square wave the frequency used should be  $f=0,5/t_{rise}$ . The signal's magnitude at time  $t$  is  $t = (\Gamma_1 \Gamma_s)^{t/T}$  where  $T$  is the one way propagation delay.

It is important to remember that the driver rise and fall times are very important in the behavior of the



transmission line and that reflections may cause long bus segment settle times.

Some general guidelines for laying out PCB's are:

- a) Evenly distribute loads along the trace, this gives a distributed load and reduces reflections from discontinuities;
- b) avoid T's for critical signals;
- c) add sockets and vias into capacitance calculations;
- d) keep trace lengths as short as possible;
- e) use as low a dielectric constant material as possible;
- f) if possible use a controlled impedance PCB, so that the behavior is predictable;
- g) balance path lengths to reduce skew; and
- h) minimize length through the connector.

Crosstalk is always an issue and all TTL/CMOS signal paths should be isolated from differential signal paths. Since crosstalk is proportional to  $dv/dt$  crosstalk may easily occur if lines with large voltage swings are near differential lines. To isolate the lines, either increase the separation, run ground traces between them, or isolate them by using different planes. Since crosstalk is caused by the capacitive coupling between signals and the mutual inductance some general observations are:

- a) Crosstalk scales with signal amplitude;
- b) crosstalk is proportional to slew rate (Voltage magnitude of output waveform/ $T_{rise}$ ) with slower rise/fall times yielding less crosstalk; and
- c) far end crosstalk width is equivalent to the rise/fall time. The crosstalk width is related to the slew rate in b). The faster the slew rate the larger the crosstalk

## Annex J

(informative)

### SPI-5 to SCSI-2 terminology mapping

This annex contains a mapping of terminology used in SCSI-2 to the terminology used in this standard (see table J.1).

**Table J.1 - SPI-5 to SCSI-2 terminology mapping**

<b>SPI-5 equivalent term</b>	<b>SCSI-2 term</b>
abort task	abort tag
abort task set	abort
cable skew	cable skew delay
clear task set	clear queue
head of queue	head of queue tag
ordered	ordered queue tag
simple	simple queue tag
target reset	bus device reset
task	I/O process
task complete	command complete
task set	queue