

Date: Oct 03, 1997

To: T10 Committee (SCSI)

From: George Penokie, Jeff Williams

Subject: Packetized SCSI

1 Overview

The packetization of SPI will provides benefits such as:

- a) ability to stream multiple CDBs to a target without having to re Arbitrate for the bus,
- b) reducing of the time the target takes to receive the initial connection information (i.e., Identify, queue tag, CDB, etc.) by approximately 6 microseconds even without streaming,
- c) reducing the disconnection time by approximately 1.5 microseconds,
- d) allowing autosense, and
- e) close to seamless mapping of fibre channel on to SPI.

NOTE 1 - All times assume Fast-40 timings.

All of the above plus backward compatibility is possible if the following proposal is accepted for inclusion into SPI-3.

2 Signal sources

Table 1 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 1 - Signal sources

SCSI bus phase	P cable signals						Q cable signals		
	A cable signals								
	BSY	SEL	C/D I/O MSG REQ	ACK ATN	DB7-0 DB(P)	DB15-8 DB(P1)	REQQ	ACKQ	DB31-16 DB(P2) DB(P3)
INFORMATION UNIT IN	Targ	None	Targ	Init	Targ	Targ	Targ	Init	Targ
INFORMATION UNIT OUT	Targ	None	Targ	Init	Init	Init	Targ	Init	Init
Init: If driven, this signal shall be driven only by the active initiator. Targ: If the signal is driven, it shall be driven only by the active target.									

3 Logical characteristics

3.1 SCSI bus phases

The SCSI architecture includes nine distinct phases:

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

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

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

3.1.0.1 Unexpected bus free

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

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

A BUS FREE phase occurring after any of the following information units shall not be considered an unexpected bus free:

- a) after the last SPI command information unit is successfully received by a target;
- b) after a SPI data information unit is successfully received by or transmitted from a target;
- c) after a SPI status information unit is successfully transmitted from a target.

3.1.1 SELECTION phase

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

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

3.1.1.1 Selection that disables information unit phases

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

3.1.1.2 Selection that enables information unit phases

If a prior IUTR negotiation (see 5.0.1.1) allows information unit phases then the initiator shall set the DATA BUS to a value that is the OR of its SCSI ID bit and the target's SCSI ID bit and it shall not assert the ATN signal (indicating that a INFORMATION UNIT OUT phase is to follow the SELECTION phase). The initiator shall then wait at least two deskew delays and release the BSY signal. The initiator shall then wait at least a bus settle delay before looking for ~~a response~~ an assertion of the BSY signal from the target.

3.1.2 Information transfer phases

The COMMAND, DATA, INFORMATION UNIT, 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 section.

Table 2 - Information transfer phases

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

3.1.2.1 Synchronous data transfer

Synchronous data transfer is optional and is only used in data phases and information unit phases. It shall be used in a data phase or an information unit phase if a synchronous data transfer agreement has been established (see xxx or 5.0.1.1). The agreement specifies the REQx/ACKx offset and the minimum transfer

period.

3.1.2.2 Wide data transfer

Wide data transfer is optional and may be used in data phases and information unit phases only if a non-zero wide data transfer agreement is in effect (see xxx or 5.0.1.1). The agreement specifies the wide mode to be used by both SCSI devices and establishes a data path width to be used during data phases and information unit phases.

3.1.3 Information unit phase

The information unit phase is a term that encompasses both the INFORMATION UNIT IN phase and the INFORMATION UNIT OUT phase.

3.1.3.1 INFORMATION UNIT IN phase

The INFORMATION UNIT IN phase allows the target to request that information units be sent to the initiator from the target.

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

3.1.3.2 INFORMATION UNIT IN phase exception condition handling

The initiator shall not release the ACK for the last byte of any information unit CRC until the CRC has been verified to be correct.

If the initiator detects a parity error on any byte or a CRC error in any information unit it receives while in the INFORMATION UNIT IN phase the initiator shall create an attention condition by asserting the ATN signal before the ACK signal is released for the last byte of CRC. When the target switches to a MESSAGE OUT phase the initiator shall send an INITIATOR DETECTED ERROR message (see xxx) to the target. This message notifies the target that a byte in the INFORMATION UNIT IN phase was invalid.

If the target does not retry the INFORMATION UNIT IN 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 INITIATOR DETECTED ERROR MESSAGE RECEIVED for the task associated with the received INITIATOR DETECTED ERROR message. (48/00)

3.1.3.3 INFORMATION UNIT OUT phase

The INFORMATION UNIT OUT phase allows the target to request that information units be sent from the initiator to the target.

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

3.1.3.4 INFORMATION UNIT OUT phase exception condition handling

If the nexus has been fully identified (i.e., an I_T_L nexus or an I_T_L_Q nexus has been established) and target detects a parity error on any byte or a CRC error in any information unit it receives while in the INFORMATION UNIT OUT phase the target shall return CHECK CONDITION status and set the sense key to ABORTED COMMAND and the additional sense code to SCSI PARITY ERROR for the task associated with the parity error or CRC error. (47/00)

If the nexus is not fully identified (i.e., an I_T nexus exists, but the target is reconnecting for an I_T_L nexus or an I_T_L_Q nexus) and target detects a parity error on any byte or a CRC error in any information unit it receives while in the INFORMATION UNIT OUT phase the target shall cause an unexpected bus free by

generating a BUS FREE phase (see xxx)

3.2 SCSI bus conditions

3.2.1 Attention condition

A target shall respond with MESSAGE OUT phase as follows:

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

3.3 SCSI bus phase sequences

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

A hard reset can abort any phase and is always followed by the BUS FREE phase. Also any other phase can be followed by the BUS FREE phase but many such instances are error conditions (see 3.1.0.1).

The additional allowable sequences for a selection that disables information unit phases shall be as shown in figure 1.

The normal progression for selection that disables information unit phases (see 3.1.1.1) is from the BUS FREE phase to ARBITRATION, from ARBITRATION to SELECTION or RESELECTION, and from SELECTION or RESELECTION to one or more of the information transfer phases, excluding information unit phases (i.e., COMMAND, DATA, STATUS, or MESSAGE). The final information transfer phase is normally the MESSAGE IN phase where a DISCONNECT, or COMMAND COMPLETE message is transferred, followed by the BUS FREE phase.

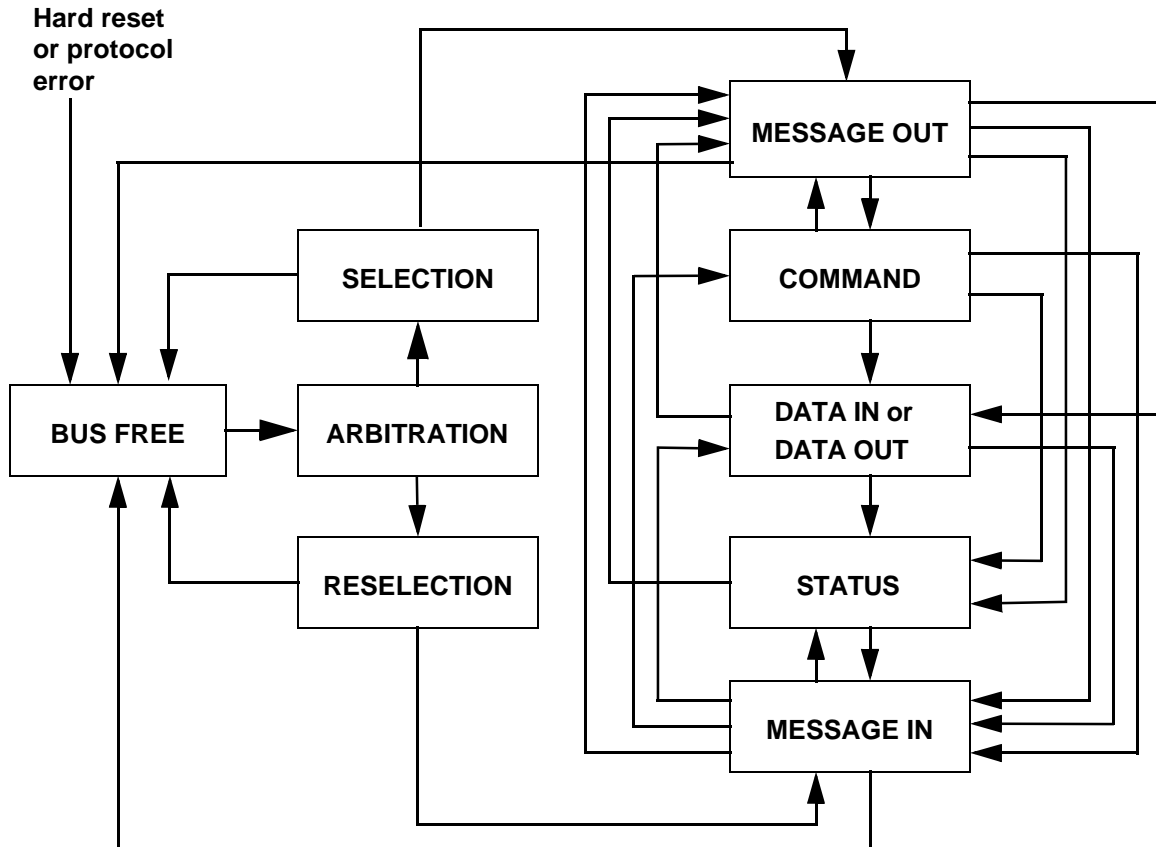


Figure 1 - Phase sequences for selection that disables information unit phases

The additional allowable sequences for a selection that enables information unit phases shall be as shown in figure 2.

If a data transfer agreement is in effect that enables information unit phases (see 5.0.1.1), the normal progression for selection that enables information unit phases (see 3.1.1.2) is from the BUS FREE phase to ARBITRATION, from ARBITRATION to SELECTION or RESELECTION, and from SELECTION or RESELECTION to one or more information unit phases (INFORMATION UNIT OUT or INFORMATION UNIT IN). The final information unit phase is followed by the BUS FREE phase.

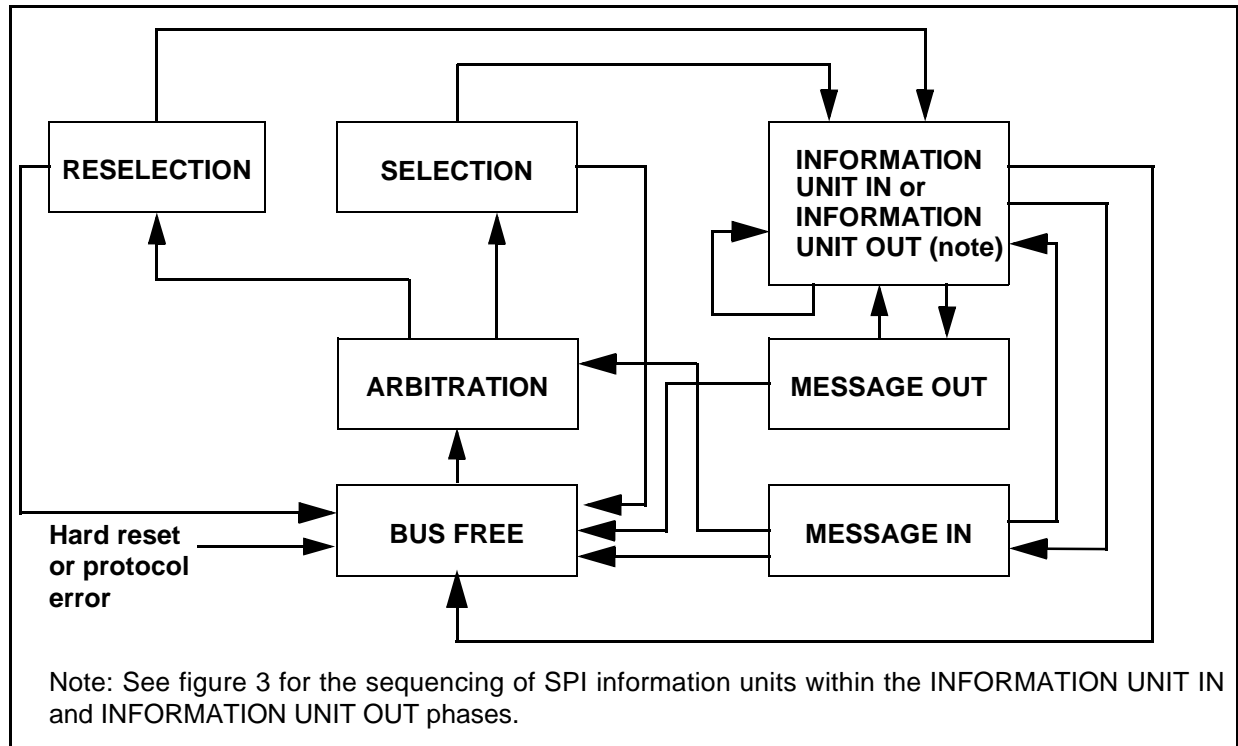


Figure 2 - Phase sequences for selection disables information unit phases

4 SPI information unit

This clause describes the SPI information units.

4.1 SPI information unit sequences

The information unit phase transfers data in SPI information units. The order in which SPI information units are transferred within the information unit phase follows a prescribed sequence.

The allowable SPI information unit sequences shall be as shown in figure 3, figure 4, and figure 5. See figure 2 for the sequencing rules between the INFORMATION UNIT IN or INFORMATION UNIT OUT phases and the other phases.

The normal progression is from one or more SPI L_Q information unit/SPI command information unit pair(s), to one or more SPI L_Q information unit/SPI data information unit pair(s), to a SPI L_Q information unit/SPI status information unit pair(s).

NOTE 2 - An initiator may force a BUS FREE phase by asserting the ATN signal and sending a DISCONNECT message on the corresponding MESSAGE OUT phase. This allows an initiator to break up a long sequence of SPI L_Q information unit/SPI data information unit pairs into smaller sequences.

The target shall not start an information unit transfer until all REQ(s)/REQQ(s) have been responded to by an equal number of ACK(s)/ACKQ(s) (i.e., any information related to an offset has been moved into storage).

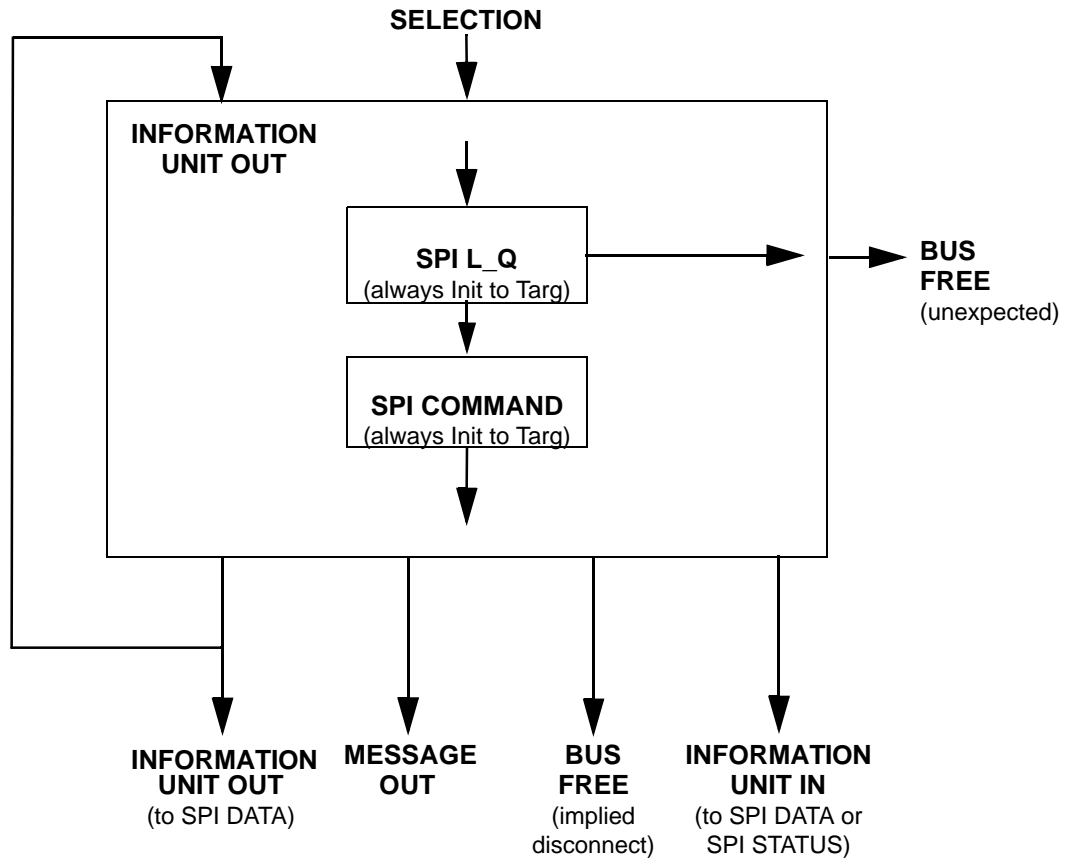


Figure 3 - SPI information unit sequence during initial connection

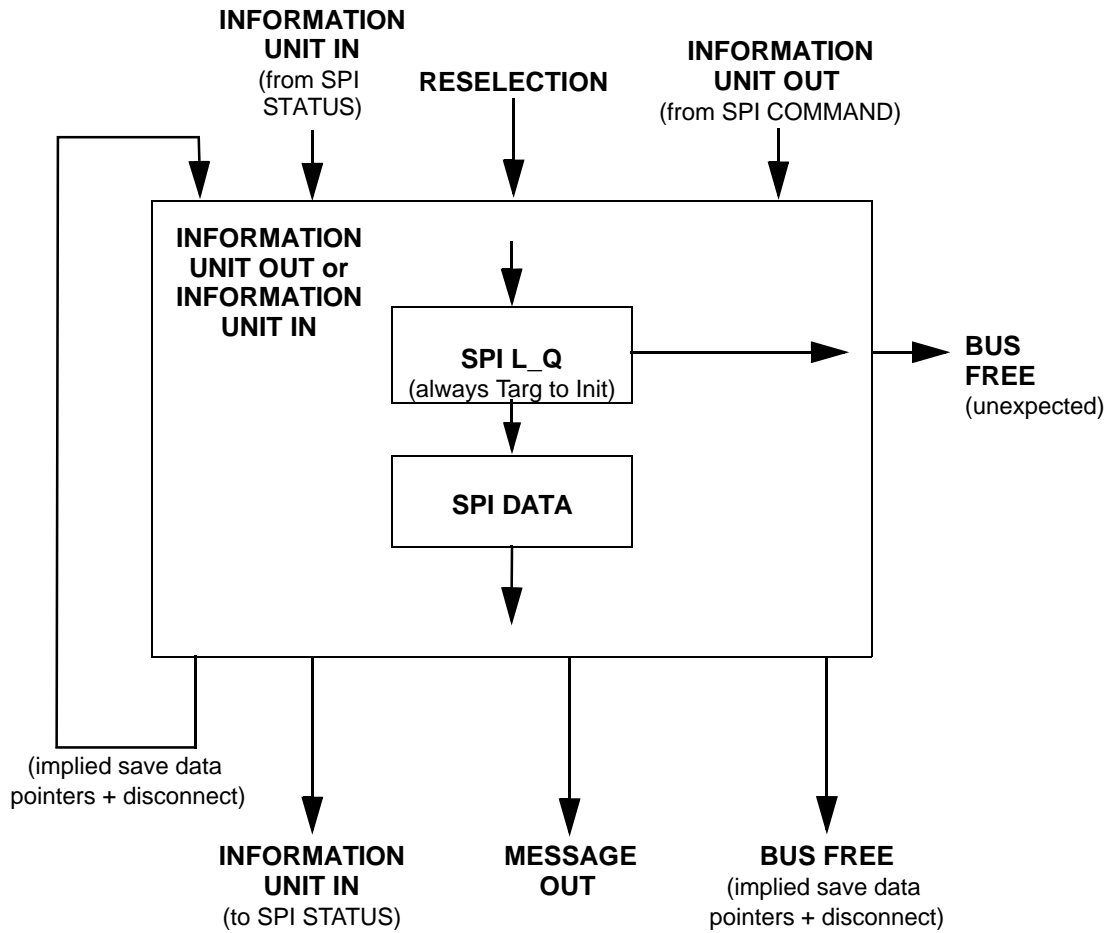


Figure 4 - SPI information unit sequence during data transfers

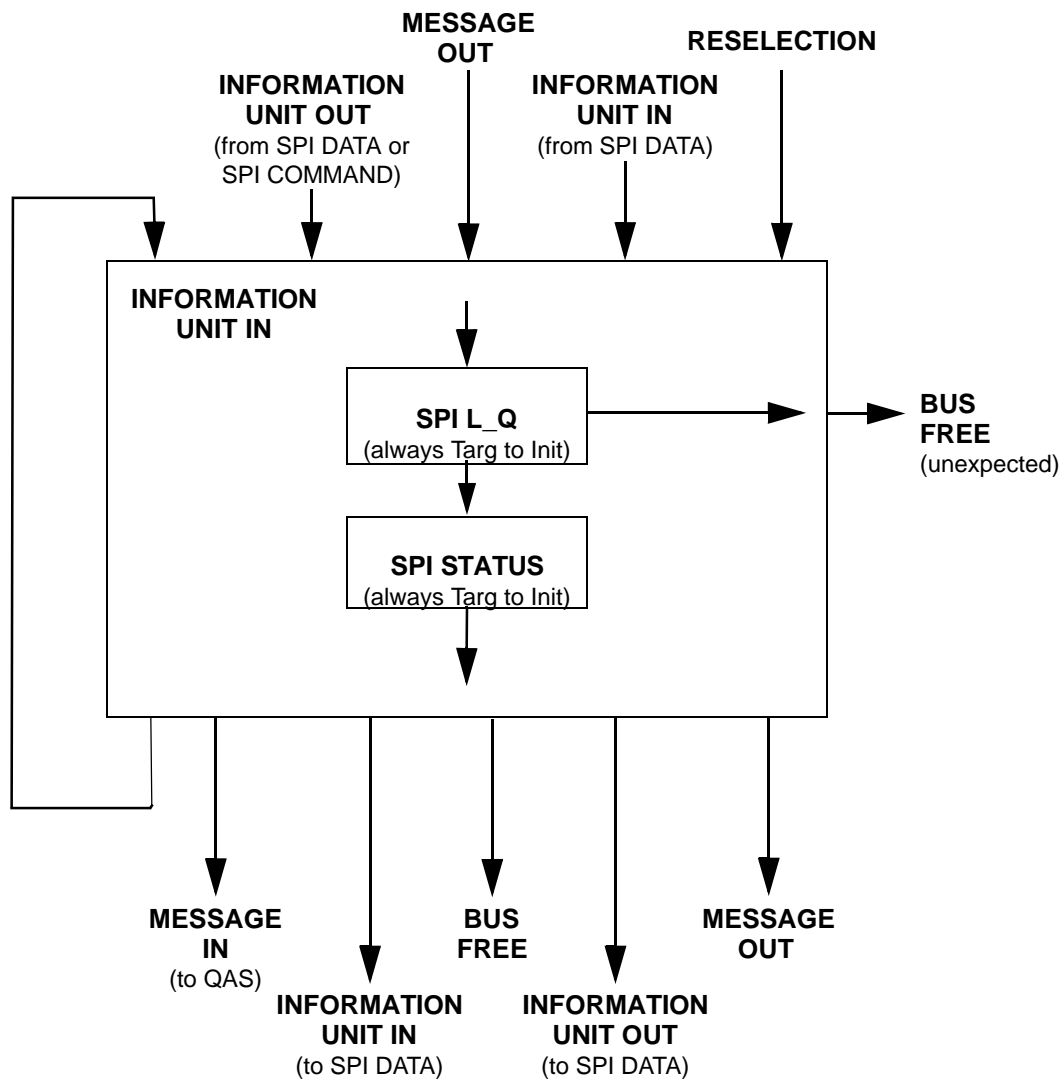


Figure 5 - SPI information unit sequence during status transfers

4.2 SPI command information unit

The SPI command information unit (see table 3) carries CDBs, task attributes, and task management requests to be performed by a device server.

An initiator shall consider a BUS FREE phase after the transfer of a SPI command information unit to be equivalent to receiving a DISCONNECT message.

Table 3 - SPI command information unit

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) LOGICAL UNIT NUMBER (LSB)							
7								
8	RESERVED							
9	RESERVED				TASK ATTRIBUTE			
10	TASK MANAGEMENT FLAGS							
11	RESERVED					RDDATA	WRDATA	
12	(MSB) CDB (LSB)							
27								
28	MSB							
29								
30	DATA LENGTH							
31	LSB							
32	MSB							
33								
34	CRC							
35	LSB							

The LOGICAL UNIT NUMBER field shall be ignored. See 4.3 for the logical unit that shall be used to establish the I_T_L nexus.

The TASK ATTRIBUTE field is defined in table 4.

TABLE 4 - TASK ATTRIBUTE

Codes	Description
000b	Requests that the task be managed according to the rules for a simple task attribute. (See SAM)
001b	Requests that the task be managed according to the rules for a head of queue task attribute. (See SAM)
010b	Requests that the task be managed according to the rules for a ordered attribute. (See SAM)
011b	Reserved
100b	Requests that the task be managed according to the rules for a automatic contingent allegiance task attribute. (See SAM)
101b	Requests that the task be managed according to the rules for an untagged task. (See SAM).
110b-111b	Reserved

The TASK MANAGEMENT FLAGS field is defined in table 5. If a task management function fails the task manager shall terminate the task with a GOOD status. The protocol failure code shall be set to task management function failed.

TABLE 5 - TASK MANAGEMENT FLAGS

Codes	Description
00h	Indicates no task management requests for the current task.
01h	The task manger shall abort the task as defined in the ABORT TASK message (see xxx).
02h	The task manger shall abort the task set as defined in the ABORT TASK SET message (see xxx).
04h	The task manger shall clear the task set as defined in the CLEAR TASK SET message (see xxx).
08h	The task manager shall perform a hard reset to the selected logical unit as defined in the LOGICAL UNIT RESET message (see xxx).
20h	The task manger shall perform a hard reset as defined in the TARGET RESET message (see xxx).
40h	The task manger shall perform a clear ACA as defined in the CLEAR ACA message (see xxx).
80h	The task manger shall perform a terminate task as defined in the TERMINATE TASK message (see xxx).
others	The task manager shall terminate the task with a GOOD status. The protocol failure code shall be set to task management function not supported.

The write data bit (WRDATA) and read data bit (RDDATA) shall be ignored.

The CDB field contains the actual CDB to be interpreted by the addressed logical unit. The maximum CDB length is 16 bytes. The CDB field is not valid and is ignored if any task management flag except the clear ACA flag (40h) is set to one. The contents of the CDB shall be as defined in the command standards (see xxx).

The DATA LENGTH field shall be ignored.

The crc field shall use the same algorithm as defined in FC-PH.

4.3 SPI L_Q information unit

The SPI L_Q information unit (see table 6) contains L_Q nexus information for the information unit to immediately follow, the type of information unit to immediately follow, and the length of information unit to immediately follow. A SPI L_Q information unit shall immediately precede all SPI command information units, SPI multiple command information units, SPI data information units, and SPI status information units.

Table 6 - 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	VBYTE	RESERVED							
13	MSB	DATA LENGTH							LSB
14									
15									
16	MSB	CRC							LSB
17									
18									
19									

The TYPE field is defined in table 7.

TABLE 7 - TYPE

Codes	Type	Description
01h	Command	Sent by an initiator to indicate a SPI command information unit shall immediately follow this SPI L_Q information unit. Indicates the initiator will not send any more SPI command information units during the current connection. The DATA LENGTH field shall be set to 24h.
02h	Multiple Command	Sent by an initiator to indicate a SPI command information unit shall immediately follow this SPI L_Q information unit. Indicates the initiator will send another SPI L_Q information unit and SPI command information unit during the current connection. The DATA LENGTH field shall be set to 24h.
04h	Data	Sent by a target to indicate a SPI data information unit shall immediately follow this SPI L_Q information unit. A length of zero in the DATA LENGTH field shall not be considered an error.
08h	Status	Sent by a target to indicate a SPI status information unit shall immediately follow this SPI L_Q information unit.
All others		Reserved

See xxx for a description of the TAG field.

The LOGICAL UNIT NUMBER field specifies the address of the logical unit of the I_T_L nexus for the current task. The structure of the logical unit number field shall be as defined in SAM-2. 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 SPC.

A valid byte bit (VBYTE) of zero indicates the last byte of the next information unit contains a pad byte of 00h. A VBYTE of one indicates the last byte of the next information unit contains valid information.

The DATA LENGTH field contains the length in 2-byte words of the following information unit. The data length shall not include 4 bytes of CRC (e.g., a data length of 256 would transfer 512 bytes of data plus 4 bytes of CRC for a total transfer of 516 bytes). The target shall not set the data length to a value that would exceed the maximum burst size as defined in the disconnect-reconnect page (see xxx).

The crc field shall use the same algorithm as defined in FC-PH.

4.4 SPI data information unit

The SPI data information unit (see table 8) contains data.

An initiator shall consider the successful transfer of a SPI data information unit to be equivalent to receiving a SAVE DATA POINTERS message and a DISCONNECT message.

Table 8 - SPI data information unit

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
n-4	DATA							
n-3	MSB							
n-2	CRC							
n-1								
n	LSB							

The DATA field may contain any type of information (e.g., parameter lists, mode pages, user data, etc.). The DATA field shall contain an even number of bytes. If the amount of data to transfer has an odd number of bytes the sending SCSI device shall pad the last byte with a 00h and set the VBYTE to zero (see 4.3) if no valid data is available.

The crc field shall use the same algorithm as defined in FC-PH.

4.5 SPI status information unit

The SPI status information unit (see table 9) contains the completion status of the task indicated by the preceding SPI L_Q information unit. The target shall consider the message transmission to be successful when it detects the negation of ACK for the last byte of the information unit with the ATN signal false.

Table 9 - SPI status information unit

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) RESERVED (LSB)							
7	RESERVED							
8	RESERVED							
9	RESERVED							
10	RESERVED			UNDER	OVER	SNSVALID	RSPVALID	
11	STATUS							
12	(MSB) SENSE DATA LIST LENGTH (n-m) (LSB)							
15	RESERVED							
16	(MSB) PROTOCOL FAILURES LIST LENGTH (m-19) (LSB)							
19	RESERVED							
20	(MSB) PROTOCOL FAILURES (LSB)							
m	RESERVED							
1+m	(MSB) SENSE DATA (LSB)							
n	RESERVED							
n+1	MSB							
n+2	RESERVED							
n+3	CRC							
n+4	LSB							

The OVER field and UNDER field shall be set to zero by the target and ignored by the initiator.

A sense data valid bit (SNSVALID) of zero indicates the sense data list length is not valid 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.

A protocol failures valid bit (RSPVALID) of zero indicates the protocol failures list length is not valid and no protocol failure information is provided. A RSPVALID bit of one indicates the PROTOCOL FAILURES LIST LENGTH field specifies the number of bytes in the PROTOCOL FAILURES field.

If the RSPVALID bit is one, the PROTOCOL FAILURES LIST LENGTH field contains the length in bytes of the PROTOCOL FAILURES field. If there is no protocol failure information the RSPVALID bit and the PROTOCOL FAILURES LIST LENGTH field shall be set to zero. The PROTOCOL FAILURES LIST LENGTH field shall only contain lengths of 0, or 4. Other values or length are reserved for future standardization.

If the SNSVALID bit is one, the SENSE DATA LIST LENGTH field contains the length in bytes of the SENSE DATA

field. If there is no sense data the SNSVALID bit and the SENSE DATA LIST LENGTH field shall be set to zero. The SENSE DATA LIST LENGTH field shall only contain even lengths (i.e., 0,2,4, etc.).

The PROTOCOL FAILURES field (see table 10) contains information describing the protocol failures detected during the execution of a task. The PROTOCOL FAILURES field shall contain valid information if the target detects any of the conditions described by the protocol failure code (see table 11).

Table 10 - PROTOCOL FAILURES field

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED							
1	RESERVED							
2	RESERVED							
3	PROTOCOL FAILURE CODE							

The protocol failure code field is defined in table 11.

TABLE 11 - PROTOCOL FAILURE CODE

Codes	Description
00h	Indicates no failure or task management function complete.
01h	Reserved
02h	SPI command information unit fields invalid.
03h	Reserved
04h	The task management function not supported.
05h	The task management function failed.
06h-FFh	Reserved

The SENSE DATA field contains the information specified by the SCSI-3 Primary Commands Standard for presentation by the REQUEST SENSE command. The proper sense data shall be presented when a SCSI status byte of CHECK CONDITION or COMMAND TERMINATED is presented as specified by the SCSI-3 Primary Commands Standard.

The crc field shall use the same algorithm as defined in FC-PH.

5 Enabling mechanism

5.0.1 Link control messages

Table 12 - Link control message codes

Code	Support		Message Name	Direction		Negate ATN before last ACK
	Initiator	Target		In	Out	
***	O	O	INFORMATION UNIT TRANSFER REQUEST	In	Out	Yes

Key: M=Mandatory support, O=Optional support
 In=Target to initiator, Out=Initiator to target
 Yes=Initiator shall negate ATN before last ACK of message.
 ***=Extended message

5.0.1.1 INFORMATION UNIT TRANSFER REQUEST

INFORMATION UNIT TRANSFER REQUEST (IUTR) messages (see table 13) are used to negotiate a synchronous data transfer agreement, a wide data transfer agreement, and information unit phase usage between two SCSI devices.

Table 13 - INFORMATION UNIT TRANSFER message format

Bit Byte	7	6	5	4	3	2	1	0
0	EXTENDED MESSAGE (01h)							
1	EXTENDED MESSAGE LENGTH (06h)							
2	INFORMATION UNIT TRANSFER REQUEST (04h)							
3	RESERVED							
4	TRANSFER PERIOD FACTOR							
5	REQ/ACK OFFSET							
6	TRANSFER WIDTH EXPONENT (m)							
7	RESERVED							ENABLEIU

he TRANSFER PERIOD FACTOR field is defined in table 14.

Table 14 - TRANSFER PERIOD FACTOR field

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

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

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

An enable information unit phases (ENABLEIU) bit of zero indicates the SCSI device does not support information unit phases. An ENABLEIU bit of one indicates the SCSI device shall use information unit phases when requested.

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

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

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

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

Any condition that leaves the data transfer agreement in an indeterminate state shall cause the SCSI device to enter an asynchronous, eight-bit wide data transfer mode with the information unit phases disabled.

In addition, an SCSI device may initiate a IUTR message exchange whenever it is appropriate to negotiate a data transfer agreement. SCSI devices that are capable of supporting information unit phases shall not respond to a IUTR message with a MESSAGE REJECT message.

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

The IUTR message exchange establishes an agreement between the two SCSI devices;

- a) on the permissible transfer periods and the REQ/ACK offsets for all logical units on the two SCSI devices. This agreement only applies to DATA IN phases, DATA OUT phases, INFORMATION UNIT IN phases, and INFORMATION UNIT OUT phases. All other phases shall use asynchronous transfers;
- b) on the width of the data path to be used for data phase transfers between two SCSI devices. This agreement only applies to DATA IN phases, DATA OUT phases, INFORMATION UNIT IN phases, and INFORMATION UNIT OUT phases. All other information transfer phases shall use an eight-bit data path; and
- c) on enabling information unit phases.

If a prior data transfer agreement is in effect, then:

- a) if a IUTR message is rejected with a MESSAGE REJECT message any prior data transfer agreement shall remain intact; or
- b) if a IUTR message fails for any other reason any prior data transfer agreement shall remain intact.

The originating SCSI device (the SCSI device that sends the first of the pair of IUTR messages) sets its values according to the rules above to permit it to receive data successfully. If the responding SCSI device can also receive data successfully with these values (or smaller transfer periods or larger REQ/ACK offsets or both), it returns the same values in its IUTR message. If it requires a larger transfer period, a smaller REQ/ACK offset, a smaller transfer width, or does not support the selected information transfer phase option in order to receive data successfully, it substitutes values in its IUTR message as required, returning unchanged any value not required to be changed. Each SCSI device when transmitting data shall respect the limits set by the other's IUTR message, but it is permitted to transfer data with larger transfer periods, smaller REQ/ACK offsets, or both than specified in the other's IUTR message. The completion of an exchange of IUTR messages implies an agreement as shown in table 15.

Table 15 - IUTR messages implied agreements

Responding agent IUTR response	Implied agreement
Non-zero REQ/ACK offset	Synchronous transfer (i.e., Each SCSI device transmit data with a transfer period equal to or greater than and a REQ/ACK offset equal to or less than the values received in the other device's SDTR message).
REQ/ACK offset equal to zero	Asynchronous transfer
Non-zero TRANSFER WIDTH EXPONENT	Wide transfer (i.e., the initiator and the target transmit data with a transfer width equal to the responding device's transfer width). If the initiating SCSI device does not support the responding SCSI device's TRANSFER WIDTH EXPONENT then the initiating SCSI device shall MESSAGE REJECT the IUTR message (see 5.0.1.1.1 and 5.0.1.1.2).
TRANSFER WIDTH equal to zero	Eight-bit data
ENABLEIU equal to one	Information unit phases enabled
ENABLEIU equal to zero	Information unit phases disabled
MESSAGE REJECT message	Eight-bit/asynchronous data transfer with information unit phases disabled
Parity error	Eight-bit/asynchronous data transfer with information unit phases disabled
No response	Eight-bit/asynchronous data transfer with information unit phases disabled

If a parity error occurs, the implied agreement shall be reinstated if a retransmission of a subsequent pair of messages is successfully accomplished. After a vendor-specific number of retry attempts (greater than zero), if the target continues to receive parity errors, it shall terminate the retry activity. This may be done either by switching to any other information transfer phase and transferring at least one byte of information or by going to the BUS FREE phase (see xxx). The initiator shall accept such action as aborting the IUTR negotiation, and both SCSI devices shall go to eight-bit/asynchronous data transfer mode with information unit phases disabled for data transfers between the two SCSI devices.

5.0.1.1.1 Target initiated IUTR negotiation

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

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

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

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

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

5.0.1.1.2 Initiator initiated IUTR negotiation

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

If an abnormal condition prevents the target from responding with a IUTR message or with a MESSAGE REJECT message then both SCSI devices shall go to eight-bit/asynchronous data transfer mode with information unit phases disabled between the two SCSI devices.

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

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

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

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

6 Open issues

- 1) A requirement needs to be added that states: The target shall not start an information unit transfer until all REQ(s)/REQQ(s) have been responded to by an equal number of ACK(s)/ACKQ(s) (i.e., any information related to an offset has been moved into storage). (DONE)
- 2) There is a question about how to limit the length of the data information unit. The current maximum is FFFFFFFFh two-byte words or 8GBytes. Some suggestions are to:
 - a) Reduce the size of the field,
 - b) Negotiate a maximum size using IUTR,
 - c) Define a way for the target to disconnect in the middle of an information unit,
 - d) Put a recommendation into the standard that suggests the size should be no larger than x, or
 - e) Use the disconnect/reconnect mode page maximum burst size as the maximum data information unit length. (DONE: OPTION E)
- 3) How should linked commands be handled? (OPEN)
 - a) Add in a new type code called linked command.
 - b) Do nothing - State that linked commands are not supported in packetized SCSI.
- 4) A note needs to be added that states: The initiator may force a bus free phase by asserting attention and sending to the target in the corresponding message out phase a disconnect message. (DONE)
- 5) A statement needs to be added: An implied save data pointers message occurs at the end of each data information unit transfer. (DONE)
- 6) Does anything needs to be said about restore data pointers? (NO - IT WORKS THE SAME WAY AS NON-PACKETIZED SCSI.)
- 7) How does packetized SCSI handle the modify data pointers message?
 - a) Add it to the L_Q information unit,
 - b) Make a new information unit just for modify data pointers, or
 - c) Just use it the same way as it is today (i.e., the target switches to message in and then sends the modify data pointers message). (OPTION E IS THE DESIRED CHOICE)