

Date: Jan 10, 2000

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: WDTR/SDTR/PPR Interactions

Overview

The below changes are required because the interactions between SDTR and WDTR on the protocol options defined in the PPR message are not clearly defined in SPI-3.

0.0.1 SYNCHRONOUS DATA TRANSFER REQUEST

0.0.1.1 SYNCHRONOUS DATA TRANSFER REQUEST message description

SYNCHRONOUS DATA TRANSFER REQUEST (SDTR) messages (see table 1) are used to negotiate a synchronous data transfer agreement between two SCSI devices.

Table 1 - SYNCHRONOUS DATA TRANSFER message format

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

The TRANSFER PERIOD FACTOR field is defined in table 2.

Table 2 - TRANSFER PERIOD FACTOR field

Code	Description
00h-09h	Reserved (note 1)
0Ah	transfer period equals 25 ns (note 2)
0Bh	transfer period equals 30,3 ns (note 2)
0Ch	transfer period equals 50 ns (note 3)
0Dh-18h	transfer period equals the transfer period factor * 4 (note 3)
19h-31h	transfer period equals the transfer period factor * 4 (note 4)
32h-FFh	transfer period equals the transfer period factor * 4 (note 5)
note: 1 - Faster timings may be allowed by future SCSI parallel interface standards. 2 - Fast-40 data transfer rates that have a period equal to 25 ns or 30,3 ns. 3 - Fast-20 data transfer rates that have a period of less than or equal to 96 ns and greater than or equal to 50 ns. 4 - Fast-10 data transfer rates that have a period of less than or equal to 196 ns and greater than or equal to 100 ns. 5 - Fast-5 data transfer rates that have a period of less than or equal to 1 020 ns and greater than or equal to 200 ns.	

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

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

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

An SDTR message exchange shall be initiated by a target whenever a previously arranged synchronous data transfer agreement may have become invalid. An SDTR message exchange shall be initiated by an initiator if that initiator does not support the PPR message or the initiator has determined the target does not support the PPR message whenever a previously arranged synchronous data transfer agreement may have become invalid. The agreement becomes invalid after any condition that may leave the data transfer agreement in an indeterminate state such as:

- a) after a hard reset;

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

Any condition that leaves the data transfer agreement in an indeterminate state shall cause the SCSI device to enter an asynchronous data transfer mode and any protocol options shall no longer be in effect (see 16.3.10).

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

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

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

The originating SCSI device (the SCSI device that sends the first of the pair of SDTR messages) sets its values according to the rules above to permit it to receive data successfully. If the responding SCSI device is able to also receive data successfully with these values (or smaller transfer periods or larger REQ/ACK offsets or both), it returns the same values in its SDTR message. If it requires a larger transfer period, a smaller REQ/ACK offset, or both in order to receive data successfully, it substitutes values in its SDTR message as required, returning unchanged any value not required to be changed. Each SCSI device when transmitting data shall respect the negotiated limits set by the other's SDTR message, but it is permitted to transfer data with larger transfer periods, smaller synchronous REQ/ACK offsets, or both. The completion of an exchange of SDTR messages implies an agreement as shown in table 3.

Table 3 - SDTR messages implied agreements

Responding SCSI device SDTR response	Implied agreement
Non-zero REQ/ACK offset	Synchronous transfer (i.e., Each SCSI device transmits 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) with ST DATA IN and ST DATA OUT phases. <u>Any protocol options shall no longer be in effect (see 16.3.10).</u>
REQ/ACK offset equal to zero	Asynchronous transfer <u>and any protocol options shall no longer be in effect (see 16.3.10).</u>
MESSAGE REJECT message (as a result of the responding message)	The originating SCSI device shall set asynchronous transfer <u>and any protocol options shall no longer be in effect (see 16.3.10).</u>
Parity error (on responding message)	Asynchronous transfer <u>and any protocol options shall no longer be in effect (see 16.3.10).</u>
Unexpected bus free (as a result of the responding message)	Asynchronous transfer <u>and any protocol options shall no longer be in effect (see 16.3.10).</u>
No response	Asynchronous transfer <u>and any protocol options shall no longer be in effect (see 16.3.10).</u>

If there is an unrecoverable parity error on the initial SDTR message (see 10.12.2 and 10.12.4) the initiating SCSI device shall retain its previous data transfer mode. If there is an unexpected bus free on the

initial SDTR message the initiating SCSI device shall retain its previous data transfer mode.

0.0.1.2 Target initiated SDTR negotiation

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

The initiator shall create an attention condition on the last byte of the SDTR message from the target, and the initiator shall respond with its SDTR message, MESSAGE PARITY ERROR message, or with a MESSAGE REJECT message.

If an abnormal condition prevents the initiator from responding with a SDTR message or with a MESSAGE REJECT message then both SCSI devices shall return to asynchronous data transfer mode and any protocol options shall no longer be in effect (see 16.3.10) for data transfers between the two SCSI devices.

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

If the target does not support any of the initiator's responding SDTR message's values the target shall switch to a MESSAGE IN phase and the first message shall be a MESSAGE REJECT message. In this case the implied agreement shall be considered to be negated and both SCSI devices shall use the asynchronous data transfer mode and any protocol options shall no longer be in effect (see 16.3.10) for data transfers between the two SCSI devices.

If a parity error occurs, the implied agreement shall be reinstated if a retransmission of a subsequent pair of messages is successfully accomplished. After a vendor-specific number of retry attempts (greater than zero), if the target continues to receive parity errors, it shall terminate the retry activity. This is done by the target causing an unexpected bus free. The initiator shall accept such action as aborting the SDTR negotiation, and both SCSI devices shall go to asynchronous data transfer mode and any protocol options shall no longer be in effect (see 16.3.10) for data transfers between the two SCSI devices.

0.0.1.3 Initiator initiated SDTR negotiation

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

If an abnormal condition prevents the target from responding with a SDTR message or with a MESSAGE REJECT message then both SCSI devices shall go to asynchronous data transfer mode and any protocol options shall no longer be in effect (see 16.3.10) for data transfers between the two SCSI devices.

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

- a) the initiator receives the last byte of the SDTR message and parity is valid; and
- b) the target does not detect an attention condition on the last byte of the SDTR message.

If the initiator does not support the target's responding SDTR message's values the initiator shall create an attention condition and the first message shall be a MESSAGE REJECT message.

If during the SDTR message the initiator creates an attention condition and the first message out is either MESSAGE PARITY ERROR or MESSAGE REJECT the synchronous operation shall be considered to be negated by both the initiator and the target. In this case, both SCSI devices shall go to asynchronous data transfer mode and any protocol options shall no longer be in effect (see 16.3.10) for data transfers between the two SCSI devices.

0.0.2 TASK COMPLETE

The TASK COMPLETE message is sent from a target to an initiator to indicate that a task has completed and that valid status has been sent to the initiator when information unit transfers are disabled.

After successfully sending this message the target shall go to the BUS FREE phase by releasing the BSY signal. The target shall consider the message transmission to be successful when there is no attention condition on the TASK COMPLETE message.

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

0.0.3 WIDE DATA TRANSFER REQUEST

0.0.3.1 WIDE DATA TRANSFER REQUEST message description

WIDE DATA TRANSFER REQUEST (WDTR) messages (see table 4) are used to negotiate a wide data transfer agreement between two SCSI devices.

Table 4 - WIDE DATA TRANSFER message format

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

The TRANSFER WIDTH EXPONENT field defines the transfer width to be used during ST DATA IN phases and ST DATA OUT phases. The transfer width that is established applies to both SCSI devices. Valid transfer widths are 8 bits (m=00h) and 16 bits (m=01h). A TRANSFER WIDTH EXPONENT field value of 02h is obsolete and values greater than 02h are reserved.

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

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

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

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

Any condition that leaves the data transfer agreement in an indeterminate state shall cause the SCSI

device to enter an eight-bit wide data transfer mode.

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

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

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

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

- a) if either WDTR message is rejected with a MESSAGE REJECT message the prior synchronous data transfer agreement and any protocol option bits (see 16.3.10) shall remain intact;
- b) If either WDTR message fails for any other reason the prior synchronous data transfer agreement and any protocol option bits (see 16.3.10) shall remain intact; or
- c) if both WDTR messages are not rejected with a MESSAGE REJECT message the WDTR message shall cause a reset of the synchronous data transfer agreement to asynchronous mode and any protocol options shall no longer be in effect (see 16.3.10).

The originating SCSI device (the SCSI device that sends the first of the pair of WDTR messages) sets its transfer width value to the maximum data path width it elects to accommodate. If the responding SCSI device is able to also accommodate this transfer width, it returns the same value in its WDTR message. If it requires a smaller transfer width, it substitutes the smaller value in its WDTR message. The successful completion of an exchange of WDTR messages implies an agreement as shown in table 5.

Table 5 - WDTR messages implied agreements

Responding SCSI device WDTR response	Implied wide data transfer agreement
TRANSFER WIDTH EXPONENT equal to 1	16-bit data and any protocol options shall no longer be in effect (see 16.3.10).
TRANSFER WIDTH equal to zero	Eight-bit data transfer and any protocol options shall no longer be in effect (see 16.3.10)
MESSAGE REJECT message (as a result of the responding message)	The originating SCSI device shall set eight-bit data transfer and any protocol options shall no longer be in effect (see 16.3.10)
Parity error (on responding message)	Eight-bit data transfer and any protocol options shall no longer be in effect (see 16.3.10)
Unexpected bus free (as a result of the responding message)	Eight-bit data transfer and any protocol options shall no longer be in effect (see 16.3.10).
No response	Eight-bit data transfer and any protocol options shall no longer be in effect (see 16.3.10)

If there is an unrecoverable parity error on the initial WDTR message (see 10.12.2 and 10.12.4) the initiating SCSI device shall retain its previous data transfer mode. If there is an unexpected bus free on the initial WDTR message the initiating SCSI device shall retain its previous data transfer mode.

0.0.3.2 Target initiated WDTR negotiation

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

The initiator shall create an attention condition on the last byte of the WDTR message from the target, and the initiator shall respond with its WDTR message, MESSAGE PARITY ERROR message, or with a MESSAGE REJECT message.

If an abnormal condition prevents the initiator from responding with a WDTR message or with a MESSAGE REJECT message then both SCSI devices shall go to eight-bit data transfer mode ~~and any protocol options shall no longer be in effect (see 16.3.10)~~ for data transfers between the two SCSI devices.

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

If the target does not support the initiator's responding TRANSFER WIDTH EXPONENT the target shall switch to a MESSAGE IN phase and the first message shall be a MESSAGE REJECT message. In this case the implied agreement shall be considered to be negated and both SCSI devices shall use the eight-bit data transfer mode ~~and any protocol options shall no longer be in effect (see 16.3.10)~~ for data transfers between the two SCSI devices. Any prior synchronous data transfer agreement shall remain intact.

If a parity error occurs, the implied agreement shall be reinstated if a retransmission of a subsequent pair of messages is successfully accomplished. After a vendor-specific number of retry attempts (greater than zero), if the target continues to receive parity errors, it shall terminate the retry activity. This is done by the

target causing an unexpected bus free. The initiator shall accept such action as aborting the WDTR negotiation, and both SCSI devices shall go to eight-bit data transfer mode ~~and any protocol options shall no longer be in effect (see 16.3.10)~~ for data transfers between the two SCSI devices. Any prior synchronous data transfer agreement shall remain intact.

0.0.3.3 Initiator initiated WDTR negotiation

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

If an abnormal condition prevents the target from responding with a WDTR message or with a MESSAGE REJECT message then both SCSI devices shall go to eight-bit transfer mode ~~and any protocol options shall no longer be in effect (see 16.3.10)~~ for data transfers between the two SCSI devices.

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

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

If the initiator does not support the target's responding transfer width exponent the initiator shall create an attention condition and the first message shall be a MESSAGE REJECT message.

If during the responding WDTR message the initiator creates an attention condition and the first message of the MESSAGE OUT phase is either a MESSAGE PARITY ERROR or MESSAGE REJECT message the wide data transfers shall be considered to be negated by both SCSI devices. In this case, both SCSI devices shall use the eight-bit data transfer mode ~~and any protocol options shall no longer be in effect (see 16.3.10)~~ for data transfers between the two devices.