Proposal for Parallel SCSI: Domain Validation

To: T10 Technical committee
From: Jim McGrath
Quantum Corporation
500 McCarthy Boulevard
Milpitas, CA USA 95035
Phone: 408-894-4504
Fax: 408-952-3620
Email: jim.mcgrath@quantum.com
Date: June 9, 1998

0 Background

Rev 3

This revision is based on a phone conference call help on June 8, 1998. Representatives from Adaptec, Compaq, Qlogic, Quantum, Seagate, Symbios Logic and others attended.

IDVM and TDVM are at least as large as the negotiated offset.

A domain validation now consists of multiple DATA IN and DATA OUT phases, delineated by messages. One device (the one that initiated the PPR message exchange) does all of its validation first, and then the other device does its validation.

The DT protocol validation has been rewritten to be more explicit about behavior during various error conditions.

Rev 2

This revision notes the changes from revision 1. It is an outgrowth of the discussions at the T10 working group meeting on May 22, 1998 in Chicago. It will be reviewed on a conference call on June 5, with one or more subsequent revisions generated for the next working group meeting on June 19 in Irvine.

Name changes: DEC to DT, double-edged to double transition, dual-edged to dual transition, IUTR to PPR.

Domain validation is explicitly done only if DT/CRC is negotiated (although a note is inserted indicating that could be a discussion point).

A domain validation is bracketed with DOMAIN VALIDATION START and DOMAIN VALIDATION END messages (previously only the START message was used).

A note indicates that failures in testing may be expected, and do not necessarily invalidate the domain, since testing to failure is a permitted technique (the testing device, who knows what to do with the test results, is responsible for keeping track of whether the domain was invalidated).

During general testing it was explicitly noted that CRC and parity tests were to be made.

For DT/CRC testing, provision was made for the initiator to start it. Note was made that either device can declare an invalidation of the domain (since the test does not involve margining, and so a failure is a domain invalidation).

I rewrote the termination section along the lines we discussed, although requiring a more explicit sequence of testing might be easier.

Under exception conditions I made it clear that CRC checking is done as per SPI-3 rules.

Finally, I added a section on messages to keep track of the new ones in the document. Note that they could all be accommodated with a single 2-byte message with appropriate bits set.

Rev 1

This proposal is the outgrowth of discussions on the topic of expanders at the last T10 working group meeting. I have circulated drafts for comment to various people over the last few days, and I'd like to thank those who have taken the time to review the drafts and give me feedback. As usual, all errors are mine.

Since the key intent of this proposal is the creation of an architecture for domain validation testing, and to resolve the issue of detecting old expanders which do not understand double transition clocking, I have deliberately left out some more controversial items that we have discussed previously. If we can get agreement on the key items here, then we can discuss other possible enhancements to the standard later. Note that many of those items can be done with just this architecture in the standard (i.e. the architecture gives implementers a lot of freedom to innovate in this area and still maintain product interoperability).

After the meeting it was brought to my attention that the initiator should have the opportunity to conduct its own domain validation tests. This required the addition of a message to get the target's attention – which in turn could also be used by the target to clearly identify a domain validation. The sequence is still terminated with an affirming PPR negotiation (or a new negotiation in the case of a failure).

This proposal is not structured as an explicit modification to SPI. Since there is still some disagreement on the overall content, I felt it was best to focus on that with this revision. If we can get enough agreement, the next revision will be structured as a formal modification to the SPI-3 draft document.

1 Introduction

This document describes a proposal for the SCSI parallel interface to:

1) Create an architecture to test elements of the SCSI domain (principally the cable plant) in order to determine if the data transfer parameters negotiated between the target and the initiator can be supported in their domain;
2) Specify one test for double-transition clocking and CRC;

The methods defined in this proposal allow initiators and targets to be backward compatible with SPI-2 compatible initiators and targets. This proposal may be implemented in hardware, firmware, or a combination of both hardware and firmware.

2 Definitions

2.1 CRC (Cyclic Redundancy Check): An error detecting code used to detect the validity of data that has been transferred during the current DT DATA IN or DT DATA OUT phase.

2.2 DT (double-transition or dual-transition clocking): The method for transferring data into a register or latch on both polarity edges of the clock signal. Double-transition clocking is used in the DT DATA phases to perform data transfers on both edges of the REQx or ACKx. Double-transition clocking shall occur only when the DT DATA OUT and DT DATA IN phases are in effect. The DT DATA OUT phase is in effect when MSG is asserted, C/D is negated and I/O is negated. The DT DATA IN phase is in effect when MSG is asserted, C/D is negated and I/O is asserted.

> Note: These two phase cases were reserved in SPI-2. By using previously reserved cases, low level hardware such as expanders and data capture tools can determine whether double transition clocking is being used and in what direction the DB(Px) line should be driven.

2.3 Initiator Domain Validation Memory (IDVM): The amounts of memory for the initiator, expressed in transfers on the SCSI bus (bytes for an 8-bit bus, words for a 16-bit bus, and double words for a 32-bit bus). [Note: this could be in bytes, but specified as transfers it made the Domain Validation section easier to write. The value is obtained from the initiator as part of the PPR negotiation process (specified in another proposal), and shall be at least as large as the negotiated synchronous offset, expressed in transfers.]

2.4 Target Domain Validation Memory (TDVM): The amount of memory for the target, expressed in transfers on the SCSI bus (bytes for an 8-bit bus, words for a 16-bit bus, and double words for a 32-bit bus). [Note: this could be in bytes, but specified as transfers it made the Domain Validation section easier to write. The value is obtained from the initiator as part of the PPR negotiation process (specified in another proposal), and shall be at least as large as the negotiated synchronous offset, expressed in transfers.]

3 Initiating domain validation sequence

Another proposal documents the process by which new protocols and speeds are negotiated using the PPR message. After these protocols and speeds have been negotiated, a sequence of one or more domain validations may be performed.

- If the PPR message has been used to negotiate the DT protocol, then domain validation shall be attempted by the device that initiated the PPR message exchange, and the other device shall engage in the domain validation.

- If the PPR message has been used to negotiate the ST protocol, then domain validation shall be attempted by the device that initiated the PPR message exchange, and the other device may engage in the domain validation. If it does not engage in domain validation it shall signify this by sending a MESSAGE REJECT message in response to receiving the DOMAIN VALIDATION START message.

- If the PPR message has not been used during negotiation, then domain validation shall not be attempted by either device.

If a PPR message exchange had taken place, then the target and initiator shall have also indicated through the PPR negotiation the amount of high-speed memory (i.e. capable of being filled or emptied at the negotiated transfer rate) available for domain validation.

Domain validation may be performed for the following items:

- Protocol: to validate that double transition clocking and CRC protocol negotiated between the target and the initiator is supported by the domain containing the target and the initiator.

4 Domain validation

A domain validation shall be initiated by sending a DOMAIN VALIDATION START message. This uniquely identifies the start of each domain validation. The domain validation shall be terminated by a DOMAIN VALIDATION END or a DOMAIN VALIDATION TERMINATED message. Each device may initiate multiple domain validations. The last domain validation that a device initiates shall end with a DOMAIN VALIDATION TERMINATED message.

Domain validations shall immediately follow an exchange of PPR message. The device which initiated the PPR message exchange shall first initiated a domain validation. It shall continue to initiate domain validations until it completes by sending the DOMAIN VALIDATION TERMINATED message. The other device may then initiate a domain validation. It shall continue to initiate domain validations until it completes by sending the DOMAIN VALIDATION TERMINATED message.

Note that multiple domain validations may have to be executed in order to validate a domain. Failures may occur during a domain validation, and yet still result in the validation of a domain (in this case the devices may be deliberately testing the domain to failure to identify the degree of margin available). The device that initiates the domain validation shall be responsible for making the decision as to whether the domain is valid or invalid as a result of the tests performed for that domain validation.

4.3 Validating the domain for DT protocol

4.3.1 Purpose and Application

The purpose of this domain validation is to detect elements of the domain that do not recognize the DT protocol (e.g., old expanders). This test is only performed if the DT protocol has been negotiated. In that case, it shall be the first domain validation performed.

This domain validation shall reliably detect old expanders when data is transferred at 40 or 80 MT/s (DT) by failing to transfer data reliably on both edges of the REQx and ACKx signal. At lower transfer rates the old expanders may occasionally work, and so this test is not recommended in environments with old expanders and 10 or 20 MT/s negotiated transfer rates. If a device desires to operate at these lower speeds (perhaps for extra margin or better overall performance), then it should negotiate for 40 or 80 MT/s, validate the domain for DT/CRC protocol, and then renegotiate for the slower transfer rate. During the test, each data pattern transferred shall be transferred at the full negotiated transfer rate. The data patterned shall be designed so that if every other transfer is latched, a CRC error will be generated.

[The above is subject to change based on feedback from expander manufacturers]

4.3.2 Protocol

The domain validation shall be performed as follows:

a) Either the target shall initiate a domain validation by going into MESSAGE IN phase and transferring a DOMAIN VALIDATION START message; or the initiator shall initiate a domain validation by creating an ATTENTION CONDITION and transferring a DOMAIN VALIDATION START message in the subsequent MESSAGE OUT phase.
b) The target shall enter a DT DATA IN phase.
c) The target shall transfer IDVM or TDVM number of data transfers to the initiator, whichever is fewer (the data transferred is vendor unique and CRC data is not counted towards this total).
d) The initiator shall perform the normal CRC and/or parity check on the data received from the target.
e) If the initiator detects a CRC or a parity error or any other exception indicating a failure, it shall create an ATTENTION CONDITION.
f) If the target detects an ATTENTION CONDITION it shall go into MESSAGE OUT phase. The initiator shall then send an INITIATOR DETECTED ERROR message to the target.
   - If the initiator initiated the domain validation, it shall then send a DOMAIN INVALIDATED message, followed by a DOMAIN VALIDATION TERMINATED message.
   - If the target had initiated the domain validation, then it shall then go into MESSAGE IN phase and send a DOMAIN INVALIDATED message, followed by a DOMAIN VALIDATION TERMINATED message.
   - In either case, the domain validation shall be terminated, and the domain shall be considered invalid.
g) If the target does not detect an ATTENTION CONDITION it shall go into DATA OUT phase. The initiator shall then transfer the data previously transferred from the target back to the target.
h) The target shall perform the normal CRC and/or parity check on the data received from the initiator.
i) If the target detects a CRC or a parity error or any other exception indicating a failure, it shall go into MESSAGE IN phase. The target shall then send a TARGET DETECTED ERROR message to the initiator.
   - If the target had initiated the domain validation, then it send a DOMAIN INVALIDATED message, followed by a DOMAIN VALIDATION TERMINATED message.
   - If the initiator initiated the domain validation, then the target shall go into the MESSAGE OUT phase. The initiator shall then send a DOMAIN INVALIDATED message, followed by a DOMAIN VALIDATION TERMINATED message.
   - In either case, the domain validation shall be terminated, and the domain shall be considered invalid.
j) If the target did not detect a CRC or a parity error or other exception indicating a failure, then:
   - If the target had initiated the domain validation, then it send a DOMAIN VALIDATED message, followed by a DOMAIN VALIDATION END message.
   - If the initiator initiated the domain validation, then the target shall go into the MESSAGE OUT phase. The initiator shall then send a DOMAIN VALIDATED message, followed by a DOMAIN VALIDATION END message.
   - In either case, if this was the last domain validation to be initiated by the device, then it shall send the DOMAIN VALIDATION TERMINATED message rather than the DOMAIN VALIDATION END message.

[The above DATA OUT sequence may be eliminated]

4.4 Target initiated domain validation

The following is the general sequence for target initiated domain validation.

a) The target shall initiate a domain validation by going into MESSAGE IN phase and transferring a DOMAIN VALIDATION START message;
b) the target shall enter a DT DATA IN phase;
c) The target shall transfer the IDVM or TDVM number of data transfers, whichever is fewer, to the initiator (the data transferred shall be vendor unique and CRC information shall not be counted towards this total), with the initiator performing the usual CRC and parity (if a 16 or 32 bit bus) checks;
d) The target shall enter a DT DATA OUT phase;
e) The initiator shall transfer the IDVM or TDVM number of data transfers, whichever is fewer, to the target (the data transferred shall be the data received by the initiator from the target in the previous DT DATA IN phase and shall be transferred in the order received by the initiator; the CRC information shall not be counted towards this total), with the target performing the usual CRC and parity (if a 16 or 32 bit bus) checks;
f) The target may then compare the data received from the initiator with the data sent by the target during the previous DT DATA IN phase;
g) The target may loop back to step b), otherwise it shall go to step h)
h) The target shall end the domain validation by going into MESSAGE IN phase.
   - If the domain was validated, it shall send a DOMAIN VALIDATED message
   - If the domain was invalidated, it shall send a DOMAIN INVALIDATED message
   - The target shall send a DOMAIN VALIDATION END message, or a DOMAIN VALIDATION TERMINATED message if all domain validations initiated by the target have been completed.

4.2 Initiator initiated domain validation

The following is the sequence for initiator initiated domain validation.

The following is the general sequence for target initiated domain validation.

a) The initiator shall initiate a domain validation by creating an ATTENTION CONDITION transferring a DOMAIN VALIDATION START message during the subsequent MESSAGE OUT phase.
b) the target shall enter a DT DATA OUT phase;
c) The initiator shall transfer the IDVM or TDVM number of data transfers, whichever is fewer, to the target (the data transferred shall be vendor unique and CRC information shall not be counted towards this total), with the target performing the usual CRC and parity (if a 16 or 32 bit bus) checks;
d) The target shall enter a DT DATA IN phase;
e) The target shall transfer the IDVM or TDVM number of data transfers, whichever is fewer, to the initiator (the data transferred shall be the data received by the initiator from the target in the previous DT DATA OUT phase and shall be transferred in the order received by the target; the CRC information shall not be counted towards this total), with the initiator performing the usual CRC and parity (if a 16 or 32 bit bus) checks;
f) The initiator may then compare the data received from the target with the data sent by the initiator during the previous DT DATA IN phase;
g) Ithe initiator shall signal its intention to end the domain validation by creating an ATTENTION CONDITION and going to step h), otherwise the devices shall go back to step b)
h) The target shall go into MESSAGE OUT phase.
   - If the domain was validated, the initiator shall send a DOMAIN VALIDATED message
   - If the domain was invalidated, the initiator shall send a DOMAIN INVALIDATED message

- The initiator shall send a DOMAIN VALIDATION END message, or a DOMAIN VALIDATION TERMINATED message if all domain validations initiated by the initiator have been completed.

5 Terminating domain validation sequence

Once each device has sent a DOMAIN VALIDATION TERMINATED message the domain shall be consider invalidated if any DOMAIN INVALIDATION message was sent/received by either device for any domain validation. If only DOMAIN VALIDATED messages were sent/received, then the domain shall considered to be validated.

Idf the domain is invalidated, then the device which initiated the previous PPR message exchange shall initiation another PPR message exchange.

6 Exception conditions during domain validation

6.1 CRC/Parity Error

A CRC or a parity error (for a 16 bit or 32 bit DATA BUS) (see xxx protection section of SPI-3).

6.2 Data Pattern Mismatch

A mismatch between the data that the device that initiated the domain validation sent to the other device, and the data received from that other device.

6.3 Non-zero REQ/ACK Offset

[Note that this is a general source of errors in SCSI, and may want to be treated more comprehensively elsewhere in SPI-3.]

If the target detects a stable non-zero REQ/ACK offset and no data transfer for at least 1 ms, then the device shall clear its REQ/ACK offset to 0 and shall go into MESSAGE IN phase. On the subsequent MESSAGE IN phase, the target shall send a REQ/ACK OFFSET MISMATCH DETECTED message.

If the initiator detects that the target has switched phases when the REQ/ACK offset is not zero, then it shall clear its REQ/ACK offset to 0 and shall create an ATTENTION CONDITION. On the subsequent MESSAGE OUT phase, the initiator shall send a PHASE CHANGE DETECTED WITH A NON-ZERO REQ/ACK OFFSET message.

6.4 Domain validation sequence timeout

The time to execute all of the domain validation initiated by a given device shall for a I-T connection shall not exceed 100 ms. If the time exceeds 100 ms, then the domain shall be considered invalidated and the domain validation sequence shall be terminated.

6.5 Unexpected protocol

If the target receives any message other than those specified in this section, then the domain shall be considered invalidated and the domain validation sequence shall be terminated.

If the initiator detects any phase change other than those specified in this section, then the domain shall be considered invalidated and the domain validation sequence shall be terminated.

If the target goes to the BUS FREE phase for any reason before the domain validation sequence has been terminated, then the domain shall be considered invalid.

7.0 Message Additions

The following message need to be allocated message codes under this proposal:

DOMAIN VALIDATION START
DOMAIN VALIDATION END
DOMAIN VALIDATION TERMINATED
DOMAIN VALIDATED
DOMAIN INVALIDATED
DOMAIN VALIDATION COMPLETE

The above message may be individual messages or a 2 byte message with different bits set in the second byte. [The group should discuss this.]

In addition, two messages are needed for the REQx/ACKx hang situations:

REQ/ACK OFFSET MISMATCH DETECTED
PHASE CHANGE DETECTED WITH A NON-ZERO REQ/ACK OFFSET
TARGET DETECTED ERROR

Finally, the PPR negotiation message is needed.