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 21May 21, 1998

0 Background

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 DOMIANDOMAIN 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 edge 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 IUTR 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-transitionedged 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 DTEC DATA IN or DTEC DATA OUT phase.

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

negated. The DTEC 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 transitionedge 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 PPRIUTR negotiation process (specifedspecified in another proposal).]

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 PPRIUTR negotiation process (specifedspecified in another proposal).]

3 Initiating Ddomain Vvalidation Ssequence

Another proposal documents the process by which new protocols and speeds are negotiated using the PPRIUTR message. After these protocols and speeds have been negotiated, a sequence of one or more domain validations may be performed if the DT/CRC protocol has been negotiated and the Domain Validation bit in the PPRIUTR message was set.

[There was some discussion of allowing Ddomain Vvalidation in the ST (single transition – current DATA phase) protocol, but I do not remember if there was a consensus. Similarly, there was some discussion that if the domain validation was confined to DT/CRC, then it could be required and the Domain Validation bit in the PPR eliminated. Once again, I do not remember a consensus on this.]

In this case the target and initiator shall have also indicated through the PPRIUTR negotiation the amount of high-speedhigh-speed memory (i.e. capable of being filled or emptied at the negotiated transfer rate) available for domain validation.

Domain Vvalidation may be performed for the following items:

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

4 Domain Vvalidation

Either the initiator or the target may initiate a Ddomain Vvalidation. A Ddomain Vvalidation is shall be initiated by the sending of a DOMAIN VALIDATION START message. This uniquely identifies the start of each Ddomain Vvalidation. It isThis shall then be followed by a DTEC DATA IN phase, or a pair of DTEC DATA IN and DTEC DATA OUT phases. It isThe domain validation shall be terminated by aeither another DOMAIN VALIDATION ENDSTART message. (to start another Domain Validation) or another IUTR negotiation.

Note that multiple Ddomain Vvalidations may have to be executed in order to validate a domain. A Ddomain Vvalidation may fail, 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 Ddomain Vvalidation is shall be responsible for making the decision as to whether the domain is valid or invalid.

4.1 Target Iinitiated Ddomain Vvalidation

The following is the sequence for target initiated Ddomain Vvalidation. Note that the DTEC/CRC protocol validation is unique in that only steps A, B, and C are performed (see 4.3.1 for details).

a) The target shall initiates a domain validation by going into MESSAGE IN phase and transferring a DOMAIN VALIDATION START message;
b) the target shall enters a DTEC DATA IN phase;
c) The target shall transfers the IDVM or TDVM number of data transfers, whichever is fewer, to the initiator (the data transferred is shall be vendor unique and CRC information is 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 enters a DTEC DATA OUT phase;
e) The initiator shall transfers the IDVM or TDVM number of data transfers, whichever is fewer, to the target (the data transferred is shall be the data received by the initiator from the target in the previous DTEC DATA IN phase and is shall be transferred in the order received by the initiator; once again, the CRC information is 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 compares the data received from the initiator with the data sent by the target during the previous DTEC DATA IN phase;
g) The target shall ends the domain validation by going into MESSAGE IN phase and transferring a DOMAIN VALIDATION END message;
h) Based on the outcome of the datathis comparison (if any), previous Ddomain Vvalidations, and any exception conditions (see 6) encountered, the domain may be validated or invalidated for the item being tested (see clause 4.3 for specific domain validations) by the target.

4.2 Initiator Iinitiated Ddomain Vvalidation

The following is the sequence for initiator initiated Ddomain Vvalidation.

a) The initiator shall initiates 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 enters a DTEC DATA OUT phase;
c) The initiator shall transfers the IDVM or TDVM number of data transfers to the target, whichever is fewer, (the data transferred is shall be vendor unique and CRC information is 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 enters the DTEC DATA IN phase;
e) The target shall transfers the IDVM or TDVM number of data transfers, whichever is fewer, of data to the initiator (the data transferred is shall be the data received by the target from the initiator in the previous DTEC DATA OUT phase and is transferred in the order received by the target; once again, the CRC information is 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 compares the data received from the target with the data sent by the initiator during the previous DTEC DATA OUT phase;
g) The initiator ends the domain validation by creating an ATTENTION condition and transferring a DOMAIN VALIDATION END message in the subsequent MESSAGE OUT phase;

h) Based on the outcome of the datathis comparison (if any), previous Ddomain Vvalidations, and any exception conditions (see 6) encountered, the domain may be validated or invalidated for the item being tested (see clause 4.3 for specific domain validations) by the initiator.

4.3 Specific Ddomain Vvalidations

4.3.1 Validating the Ddomain for DTEC/CRC Pprotocol

[It was noted at the working group that the final edit of this document may want to put this section before the previous section – but since that would complicate editorial revision marking, the decision to do this has been postponed and may just be left to the editor.]

The purpose of this validation is to detect elements of the domain that do not recognize the DTEC/CRC protocol (e.g., old expanders). This test willshould reliably detect old expanders when data is transferred at 40 or 80 MT/s 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.

The Ddomain Vvalidation shall be performed as specified in clauses 4.1 or 4.2, except that only the DTEC DATA IN portion is done .done. Specifically:

a) Either tThe target shall initiates a domain validation by going into MESSAGE IN phase and transferring a DOMAIN VALIDATION START message; or the initiator shall initiates 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 enters a DTEC DATA IN phase;
c) The target shall transfers 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 performs the normal CRC and parity (if a 16 or 32 bit transfer) check on the data received from the target;
e) The initiator shall ends the domain validation by creating an ATTENTION condition and transferring a DOMAIN VALIDATION END message in the subsequent MESSAGE OUT phase;
f) Based on the outcome of this comparison and any exception conditions (see 6) encountered, the domain is may be validated or invalidated for the DTEC/CRC protocol. Note that either the target or the initiator may invalidate the domain for the DT/CRC protocol.

This Ddomain Vvalidation shall be initiated by either at least one of the devices (only one domain validation is required) whenever the DTEC/CRC protocol has been negotiated. It This domain validation shall be done before any other Ddomain Vvalidation.

5 Terminating Ddomain Vvalidation Ssequence

When the device which did not initiate the PPR negotiation has completed all of its domain validation testing, then it shall send either a DOMAIN VALIDATED or a DOMAIN INVALIDATED message to the other device. This device, whichthat previously initiated the PPR negotiation, may complete its own domain validations. Based on its own validation results and the message received form the other device, it shall determine eif the domain has been validated or invalidated. If validated, it shall notify the other device by sending a DOMAIN VALIDATION COMPLETE message. After that point, normal operations shall resume. If the domain has been invalidated, then the device shall start another PPR negotiation process (excluding the invalidated item from the negotiation).

[This can be simplified if we require that one device perform all of its testing first, then the other device. For instance, if the device that did not initiate negotiation first did all of its tests, then send its DOMAIN VALIDATED/INVALIDATED message, then the initiating device performed its tests, followed by the DOMAIN VALIDATION COMPLETE or another PPR, then things might be easier. Feedback is appreciated.]

Any device which determines that the domain has been invalidated for an item agreed to during the previous IUTR negotiation shall consider the results of the previous IUTR negotiation to be invalidated at that point, and shall initiate another IUTR negotiation. In this case the device shall exclude from consideration during this negotiation the item for which the domain was invalidated.

The device that initiated the IUTR negotiation may terminate Domain Validation. This is done by initiating another IUTR negotiation. If both devices agree that all of the items requiring Domain Validation have been validated, then the Domain Validation bit shall be cleared to zero and the target shall go into whatever SCSI phase is appropriate for the command being executed after the IUTR negotiation. Otherwise the Domain Validation bit shall be set to 1 and another Domain Validation Sequence will begin.

The time to execute all of the PPRIUTR negotiations and Ddomain Vvalidation Ssequences during a given I-T connection shall not exceed 100 ms.

6 Exception Cconditions during Ddomain Vvalidation

6.1 CRC/PartiyParity 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 Ddomain Vvalidation 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 Ssequence Ttimeout

The time to execute the Ddomain Vvalidation Ssequence shall not exceeds 100 ms. If the time exceeds 100 ms, then the domain shall be considered invalidated and the Ddomain Vvalidation Ssequence shall be terminated.

6.5 Unexpected Pprotocol

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

If the target goes to the BUS FREE phase for any reason before the Ddomain Vvalidation Ssequence 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 VALIDATED
DOMAIN INVALIDATED
DOAMINDOMAIN 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

Finally, the PPR negotiation message is needed.