Outstanding Issues in Document 98-177: Proposal for Parallel SCSI: Increase Transfer Rate and Improve Error Detection

To: T10 Technical committee
From: Jim McGrath
Quantum Corporation
500 McCarthy Boulevard
Milpitas, CA USA 95035
Phone: 408-894-4019
Fax: 408-952-3620
Email: jim.mcgrath@quantum.com
Date: September 11,1998

0 Overview

This proposal deals with the issues identified in 98-177 at the conference call to review revision 4 of that document, held on August 27, 1998. Representatives from Adaptec, Compaq, IBM, QLogic, Quantum, Seagate, Symbios, and WD attended. The goal of that conference call was to identify issues that needed resolution in order to allow the approval of 98-177 for inclusion into SPI-3 at the September plenary meeting.

When approval of 98-177 comes up at the working group, and debate/voting is needed to resolve the issues, then I am suggesting that a straightforward process of identifying the issues, stating the possible options, limited debate (since many of these options have been debated before), and the a vote on the options with the aim of driving the issue to closure. If the group is divided, then if necessary we should discuss it at the plenary, but I hope that we will get some significant support for a single position on each issue. To facilitate this I asked the group on the con call to identify the outstanding issues, with the promise that I would turn the document one more time to provide appropriate wording for the alternatives mentioned. In that way, we will not have to waste time wording on an exact wording change at the meeting, but rather will have "pre packaged" wording from the author for the various positions, allowing the committee to focus on the substance of the issues.

The specific changes in 98-177 that enact any of the resolutions to the issues raised is to be found in revision 5 of that document (please see the revision notes of that

document for a description of the changes and how to identify them). This document outlines the substansive issues raised by these issues.

The issues are as follows:

A. Degree of support for 8 bit transfers
B. Degree of support for single ended
C. Use of the P1 signal
D. Value of the Pad bytes

1 8 bit Mode

A) 8 bit mode: the draft currently explicitly allows implementation for both 8 bit and 16 bit data transfer widths for the DT data phases. While there is agreement that the proposal provides enough information to implement either 8 bit or 16 bit, there is a dispute over the desirability of implementing 8 bit. This appears to be a "policy," not a technical issue per se. Since there are people with decided opinions on this topic, I suggest we vote on the following options:

A1) 32 and 16 bit only, 8 bit forbidden
A2) 32 and 16 bit only, silence on 8 bit
A3) 32 and 16 bit recommended but and 8 bit allowed
A4) 32 and 16 bit and 8 bit allowed

These options span from the most 8 bit unfriendly to the most 8 bit friendly. Specifically, the changes in 98-177 are to include this statement, and to make other minor textual changes that are necessary:

A1)The DT DATA IN and DT DATA OUT phases are valid when using 16 or 32 bit wide buses. These phases are invalid when using 8 bit wide buses.

A2) The DT DATA IN and DT DATA OUT phases are valid when using 16 or 32 bit wide buses.

A3) The DT DATA IN and DT DATA OUT phases are recommended only when 16 or 32 bit buses. These phases are valid for 8 bit wide buses.

A4) The DT DATA IN and DT DATA OUT phases are valid when using either 8, 16, or 32 bit buses.

If this range is acceptable then I suggest we start with a quick straw poll to see where the group stands. After suitable debate, I suggest we vote on all four again, and use successive elimination (on the next vote we drop the option with the lowest vote, and vote among the remaining three; then we drop the option with the lowest vote, and vote between the remaining two). While this may sound overly complicated, past experience convinces me that it is the quickest way to resolution.

Based on conversation with people, I would recommend the second or third options. I think that the first and fourth options would result in something that some people could not feel comfortable with, while the second and third options minimize this "backing into the corner."

2 Single Ended Mode

B) Single ended mode: the draft currently implicitly allows implementation for both single ended and LVD signaling technology for the DT data phases by being silent on the topic of transceivers. While not as hotly debated as the previous issue, I think we should make our position clear on this topic as well. This appears to be a "policy," not a technical issue per se. Since there are people with decided opinions on this topic, I suggest we vote on the following options:

B1) LVD only, single ended forbidden
B2) LVD only, silence on single ended
B3) LVD recommended but and single ended allowed
B4) LVD and single ended allowed

These options span from the most single ended unfriendly to the most single ended friendly. . Specifically, the changes in 98-177 are to include this statement:

B1) The DT DATA IN and DT DATA OUT phases are valid when using LVD signaling. These phases are invalid when using single ended or HVD signaling.

B2) The DT DATA IN and DT DATA OUT phases are valid when using LVD signaling. These phases are invalid when using HVD signaling.

B3) The DT DATA IN and DT DATA OUT phases are recommended only when using LVD signaling. These phases are valid for single ended signaling. These phases are invalid when using HVD signaling.

B4) The DT DATA IN and DT DATA OUT phases are valid when using either LVD or single ended signaling. These phases are invalid when using HVD signaling.

If this range is acceptable then I suggest we start with a quick straw poll to see where the group stands. After suitable debate, I suggest we vote on all four again, and use successive elimination (on the next vote we drop the option with the lowest vote, and vote among the remaining three; then we drop the option with the lowest vote, and vote between the remaining two). While this may sound overly complicated, past experience convinces me that it is the quickest way to resolution.

Based on conversation with people, I would recommend the fourth option. I do not have a strong opinion on this matter, but do not see how forbidding single ended (running at a maximum of Fast-20 speeds) makes anyone's implementation any easier).

3 P1 line in DT data phase

C) P1 line in DT data phase: the draft currently has this being parity for all 16 bits, as per very early discussion of the working group. During the con call it was determined that the group was in flux on this issue, and it needed a firm resolution. My reading of the group indicates a decline in support for using this line as a 16 bit line, and instead specifying it as reserved for future use. I am proposing two options for consideration:

C1) the current wording indicating that it is 16 bit parity
C2) new wording indicating that the line is reserved for future use

There is some confusion over what "reserved" means for this line (the current SCSI SPI-3 definition of leaving it unconnected does not work, since it is needed for parity in the old ST data transfer mode). So I propose wording that during the DT phase, the device that did not previously drive the signal in the ST (old SCSI) phase continue to refrain from driving that signal. That is, in DT DATA OUT the target cannot drive the signal; in DT DATA IN the initiator cannot drive the signal. The other device (e.g. target in DATA IN, initiator in DATA OUT), may drive the signal (but no faster than the negotiated speed for the data lines), but that the signal shall be considered invalid and the value ignored by the receiving device. This wording would appear to maintain maximum compatibility with previous uses of the signal in the ST mode, and is a minimum burden on implementers. By making the value invalid (instead of, say, 16 bit parity) it prevents the users from starting to rely on this signal for a particular function when in reality we will be redefining this function in the future (most likely in Ultra4). This is consistent with a future definition of this line as a free running clock (previously discussed in the working group). While I have not discussed this with everyone, this would appear to be acceptable to a number of people.

Note that the driving device may assert, deassert, or tristate the line. One possible additional vote could be taken to outlaw tri-stating if people felt uncomfortable about having a line "float" so close to the high speed REQ/ACK and data lines.

Based on conversation with people, I would recommend the reserved option.

4 Pad Byte values

Pad byte values: we still have no resolution over whether the pad bytes should be required to have a certain value. Two suggestions have been offered in the past:

D1) pad bytes are 00
D2) pad bytes may be anything

In neither case is the receiver required to test for any specific value. I suggest we simply decide between these two once and for all.

Based on conversation with people, I would recommend the first option.