

Date: November 04, 2002

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli)

Subject: SPI-5 Letter Ballot Comment Resolution

Summary

23 - HP
03 - LSI
76 - Maxtor
07 - Texas Instruments
01 - TycoElectronics
=====
110 - Total

Comments attached to No ballot from William Ham of Hewlett Packard Co.:

Comment number: 001 (T) Document location: 10.7.4.2.1 Training pattern overview, 4.12.4.6.8 RTI (maybe)

Comment: 00-132r1 included a rule that whenever a target received a PPR, it had to re-run the training sequence regardless of whether it was retaining training information. This is the only way an initiator can force training to occur. This was lost when 00-132r1 was rolled into 99-295r5, which was accepted into SPI-5 revision 0. SPI-5 should mandate this.

Proposed resolution:

In 10.7.4.2.1, change:

If retain training information is disabled a training pattern shall be transferred at the start of the first DT data phase for each data transfer direction after each physical connect and physical reconnect. The training pattern shall not be transferred again until after a physical disconnection occurs.

If the retain training information is enabled a training pattern shall be transferred at the start of the first DT data phase for each data transfer direction after the retain training information is enabled. The SCSI device shall save training configuration values for each I_T nexus that has negotiated to retain training information. The SCSI device shall use the saved training configuration values for all paced transfers. The SCSI target port may retrain an I_T nexus if it determines the training configuration values are invalid, without having to renegotiate the retain training information protocol option.

NOTE 33 - The training configuration values are vendor specific.

If the retain training information is enabled and a port changes from a SCSI initiator port to a SCSI target port that SCSI target port shall retrain if the saved training configuration values were saved while the port was a SCSI initiator port.

to:

The SCSI device shall save paced data transfer training information values for each I_T nexus that has negotiated to retain training information (see 4.12.4.6.8). The SCSI device shall use the saved training configuration values for all paced transfers. The SCSI target port may retrain an I_T nexus if it determines

the training configuration values are invalid, without having to renegotiate the retain training information protocol option.

NOTE 33 - The training configuration values are vendor specific.

If paced data transfer training information is invalid, the SCSI target port shall transfer a training pattern at the start of the first DT data phase for each data transfer direction after each physical connect or physical reconnect (e.g., one training pattern for DT DATA IN, another for DT DATA OUT). The training pattern shall not be transferred again for that data transfer direction until after a physical disconnection occurs.

If paced data transfer training information is valid, the SCSI target port shall not transfer a training pattern and shall use the paced data transfer training information.

A SCSI port shall invalidate its paced data transfer training information for an I_T nexus:

- a) after a PPR negotiation occurs for that I_T nexus; and
- b) after every physical disconnection from that I_T nexus if it has negotiated not to retain training information (see 4.12.4.6.8).

A SCSI target port should not invalidate its paced data transfer training information after any physical disconnection from an I_T nexus if it has negotiated to retain training information, but may do so if needed.

A SCSI initiator port should not invalidate its paced data transfer training information after any physical disconnection from an I_T nexus if it has negotiated to retain training information, but may do so if needed. If so, it shall originate a PPR negotiation to force training again on the next connection for that I_T nexus.

Accept: the concept but not the suggest wording. The following paragraph was added to section 10.7.4.2.1 to fix the problem:

After any PPR negotiation occurs that enables paced transfers, a training pattern shall be transferred at the start of the first DT data phase for each data transfer direction regardless of the negotiated value of the RTI bit (see 4.12.4.6.8).

Comment number: 002 (T) Document location: 18.1 table 76 Comment: "Not allowed" entry for subpage 00h is confusing, since 00h is the value that is placed in the MODE SENSE CDB's SUBPAGE CODE field when reading the short format. Proposed resolution: Change it to "Port Control mode page short format", reference 18.1.4.1, with a note that "SPF shall be set to zero when accessing the short format."

Accept: Added the following footnote to table 75.

The short format for the Port Control mode page is selected when the subpage code field is set to 00h in a MODE SELECT command.

Added the following footnote to table 76.

The subpage code field in a mode page with SPF set to one.

Comment number: 003 (E/T) Document location: table 76 Comment: Move table 76 from 18.1.1 into 18.1.4 (the Port Control mode page subclause). Proposed resolution: implement the comment

Reject: Moving that table would place the overall list of mode pages in two different places. A reader would have no one place to go to see all the possible mode pages that have been defined in the standard. Subpages are just a way to add in more page codes they are no unique or otherwise any different than regular mode pages.

Comment number: 004 (T) Document location: table 77 and table 79 Comment: change byte 0 bit 6 from "RESERVED" to "SPF (0)" Proposed resolution: implement the comment

Accept

Comment number: 005 (E) Document location: table 81 Comment: change "PROTOCOL IDENTIFIER = 1h" to "PROTOCOL IDENTIFIER (1h)" Proposed resolution: implement the comment

Accept

Comment number: 006 (E) Document location: multiple Comment: Change all mode page (and subpage) names to Mixed Case to match the convention agreed to for SPC-3 and other standards. Proposed resolution: implement the comment

Accept

Comment number: 007 (E) Document location: 18.1.4.1 Comment: "If the parameter data of a MODE SELECT command contains a subpage format page with the SUBPAGE CODE field is zero the SCSI target device shall return a CHECK CONDITION status. " should be "...with the subpage code field set to zero..." Proposed resolution: implement the comment

Accept

Comment number: 008 (E) Document location: Table of contents Comment: all lower level entries should be indented in relation to their respective upper level entries Proposed resolution: indent all lower level entries

Reject: That is not the style that is accepted by ANSI and/or ISO

Comment number: 009 (T) Document location: 3.1.1 Comment: should read "cable assembly" in the definition to ensure that the A cable includes the connectors. Proposed resolution: change "conductor cable" to "conductor cable assembly"

Accept

Comment number: 010 (T) Document location: after 3.1.10 Comment: add definition for cable assembly Proposed resolution: add: "cable assembly: a bulk cable that is connector terminated. A cable that has connectors attached by a manufacturer and is ready for installation in a system."

Accept: The following was added to the standard:

0.0.1 cable assembly:A bulk cable that has connectors attached.

Comment number: 011 (T) Document location: Figures 8 and 9 Comment: change cable to cable assembly and identify the device connectors Proposed resolution: implement the comment

Accept

Comment number: 012 (T) Document location: Figure 10 Comment: the "receiver" identified in this figure is the internal receiver and not the receiver connector where the specifications apply Proposed resolution: clearly identify where the connector for the SCSI device that contains the receiver is in this picture. Alternatively change the term "receiver" to something like "internal circuitry used within the receiving device for detecting the logic state of the incoming signal""

Accept: Changed to "Incoming Signal Detector"

Comment number: 013 (T) Document location: after 3.1.7 Comment: add definition for backplane Proposed resolution: Add: Backplane: a printed circuit board with connectors attached that is used for interconnecting multiple SCSI devices, especially disk drives.

Accept: Added in the following:

0.0.2 backplane:A printed circuit board with connectors attached that is used for interconnecting multiple SCSI devices.

Comment number: 014 (T) Document location: 3.1.64 Comment: Change "P cable: A 68-conductor cable or an 80-conductor connector that provides the 16-bit DATA BUS and control signals" to "P cable: A 68-conductor cable assembly or an interconnect assembly, notably backplane, that uses the 80-conductor SCA-2 connector to provide the 16-bit DATA BUS and control signals." Proposed resolution: implement the comment

Accept: Changed to the following:

0.0.3 P cable:A 68-conductor cable assemble or backplane that uses an 80-conductor connector to provide the 16-bit DATA BUS and control signals.

Comment number: 015 (T) Document location: clause 6.3 Comment: The material in this clause should be either effectively eliminated by referring to PIP or updated to match the definitions and requirements in PIP. The normative requirements should be on the interconnect assembly (the bulk cable or backplane that has the device connectors) where the specifications for signals apply. The present material in SPI-5 requires that the bulk cable meet certain requirements and misses the main point of the document which is to achieve interpretability at the device connectors. While the performance of the bulk cable is an important ingredient in the performance of the interconnect assembly, specification of only the uniform bulk cable is woefully inadequate to guarantee predictable performance for the interconnect. Further, the use of the specified requirements for bulk cable should be optional as it is of value only where multiple sourcing of bulk cable is desired. The normative requirements for interconnect should be only on the interconnect assembly measured under the conditions specified in PIP.

Note that these requirements in PIP specifically call out the signal quality requirements specified in SPI-x and that there is no conflict of requirements. Proposed resolution: Implement the comment using the reference to PIP as the main methodology.

Accept: Added the following wording as the first paragraph in section 6.3.1

The performance of the bulk cable is an important ingredient in the performance of the cable assembly, however 6.3 only specifies the uniform bulk cable which may not define a predictable performance of the interconnect. The SCSI Passive Interconnect Performance standard defines normative requirements for the cable assembly or backplane measured under specific conditions. Any normative requirements defined in the SCSI Passive Interconnect Performance standard that conflict with those in 6.3 shall take precedence over those specified in 6.3.

Comment number: 016 (T) Document location: clause 6.2 Comment: The following material in 6.2 "The following requirements ensure that all SCSI round cables may be used with LVD transceivers: a) In the P cable conductor pairs ACK and REQ shall be in the cable core; b) In the P cable, if there are more than four conductor pairs in the cable core, conductor pairs ACK and REQ shall not be adjacent to each other; c) In the A cable conductor pairs ACK and REQ shall be in the cable core; d) In the A cable, if there are more than three conductor pairs in the cable core, conductor pairs ACK and REQ shall not be adjacent to each other; e) Cable conductor pairs used for the DATA BUS (DBnP1) and P_CRCA shall be in the outer layer of the cable; f) Each cable conductor pair shall consist of the signal return and its associated signal.

Crosstalk noise is minimized by conductor placement (REQ and ACK in the center, data around the periphery) in round, twisted-pair cables and by the pin assignments on the connector on planar cables."

Should be made a design recommendation, not a normative requirement. These specifications are probably reasonable guidance for producing interconnect assemblies that meet the requirements specified in PIP (which include the signal timing and quality specifications in SPI-x) but do not deliver the stated result.

Proposed resolution: reword the referenced material as a recommended design practice rather than as a normative requirement.

Accept: All shalls changed to shoulds.

Comment number: 017 (T) Document location: 9.2.8 Comment: Change "excluding any signal distortion skew delays" to "measured with a free running clock data pattern" - signal distortion skew is nebulous, the change suggested makes the definition clear. Proposed resolution: implement the comment

Accept

Comment number: 018 (T) Document location: after 3.1.102 Comment: add a definition for "skew" Proposed resolution: suggest the following: Skew: The maximum difference in propagation time allowed between any two SCSI bus signals measured between two specified positions in the bus segment using a free running clock data pattern.

Accept

Comment number: 019 (T) Document location: figures 66 thru 70 and associated text Comment: The timing references for the non-precomp clock-like signals is the only one where the timing reference for the display is clear to me. By using the average of the signal zero crossings (after removing the d.c. content) to set the bit boundaries is as good (or bad) as other methods in other standards. We could talk for quite a while about why this can be bad but that is not the point of this comment.

Using the signal itself as the timing reference divorces it from any skew issues with respect to the overall clock that is used to create the signals in the driver. I assume that the purpose of this requirement is to ensure that with the proper phase relationship (however produced) that the local amplitude/time relationships are adequate.

When one goes to the non-clock like, non-precomp requirements there is a statement that the bit boundaries are at the same point as for the clock-like signals. On the surface that seems OK until one realizes that we do not have a clock-like signal present when measuring non-clock-like signals. That means that the same external timing reference used to create the clock-like signals must be used for the non-clock-like signals and that there can be no time translations of the display between the clock-like measurements and the non-clock-like measurements. A bit awkward but if one first does the clock-like measurements to set the bit time boundary, uses the same external trigger, and proceeds to do the non clock-like measurement by changing only the data pattern (the longitudinal data pattern on that signal line - not the data fed at the byte level to be sent across the bit lines in parallel - another point that is not well documented) seems like that can work. So the question here: did I capture the intent above? Sure not clear from the material in the standard. Proposed resolution: clearly specify how one is expected to set the timing references in

Accept: Added in the following above figure 63 :

Bit-cell boundaries are determined by the zero crossings of the REQ or ACK signal shifted by the signal timing skew, excluding ISI, between the REQ or ACK signal and the measured signal using the same conditions defined in figure 66. When measuring a data signal in figure 63, figure 64, figure 65, figure 67, figure 68, figure 69, and figure 70 the REQ or ACK signal shall be used as the timing reference for aligning data bit cell boundaries.

Comment number: 020 (T) Document location: Figures 63, 64, 65 and associated text Comment: There is no timing reference specified for the receiver masks for the precomp signals. Further, there is no explicit indication that says what the data pattern is for the requirements. I assume that it applies to an arbitrary (longitudinal) data pattern. I suppose one could use essentially the same methods as used for the non-precomp signals to set the bit boundaries but that is not specified. And there is the extrapolation that both isolated 0's (as shown) and isolated 1's (as not shown) are bound by this mask. As I read the standard, any signal that can be forced between the excluded areas by time translation is compliant. That

sort of scheme can work for clock recovery transmissions but is inviting trouble when a common latching signal is used across many bits.

The timing reference is an intrinsic part of this requirement. Seems to me that this requirement needs to be referenced to the clock in the driver that produced the signals in the first place to be effective. See also next comment relating to skew. Proposed resolution: define the timing reference used for these requirements

Accept: Added in the following above figure 63 and add in cell boundaries in figures 63, 64, and 65.

Bit-cell boundaries are determined by the zero crossings of the REQ or ACK signal shifted by the signal timing skew, excluding ISI, between the REQ or ACK signal and the measured signal using the same conditions defined in figure 66. When measuring a data signal in figure 63, figure 64, figure 65, figure 67, figure 68, figure 69, and figure 70 the REQ or ACK signal shall be used as the timing reference for aligning data bit cell boundaries.

Comment number: 021 (T) Document location: Figure 62 and related discussion Comment: Background for the comment: The propagation time skew for clock like signals (which is the only type of time related property that deserves the name 'skew' - e.g. ISI is clearly jitter) is different for every signal. Within the receiver, the skew is sensed during training so that the receiver can set its skew compensation as required. Only the receiver knows what skew compensation has been applied to each specific signal and the receiver does not share that information with anyone else. Further, signal requirements must be specified independently of any specific receiver. Therefore, it is not possible to know how much skew will be compensated for any given signal line by only observing the signal going into the receiver (unless one knows exactly how that specific receiver does the compensation AND knows exactly what is being presented to the receiver by all the other signal lines during training).

On the other hand one could put some boundaries on the amount skew that the specific signal line under test would expect to have compensated by the receiver if one measured the skew for that specific line (with respect to all others - not just REQ or ACK) any applied some skew compensation algorithm.

And, yes, there is a requirement for at least a certain amount of skew to be compensated by the receiver but there is no requirement that any specific signal line have a certain amount of skew compensation.

The amount of skew compensated can be a very significant portion of the bit time so if one does not know how much skew compensation to allow for any specific signal, measuring the setup and hold time with respect to the observed REQ or ACK seems bogus.

Actual comment:

The setup and hold times for both synchronous and paced transfers are made between the REQ/ACK and SPECIFIC DATA signals one at a time at the device connector. There is a clear reference to the strobe offset (good thing) for paced transfers in Figure 22 but no reference to how to account for the skew that will be compensated by the receiver in paced transfers for the specific signal being measured.

If we do not measure the skew present in the signal line under test and allow a certain portion of that skew for that specific line in the measurement before attempting to measure a setup and hold time the measured setup and hold time can be off by multiple nanoseconds.

I assume that there should be some method provided to include the actual skew in these signal measurements but it does not seem to be stated anywhere in the document.

This issue applies to both precomp and non-precomp signals.

Proposed resolution: assuming that the issues discussed above have been included in the timing requirements and are not exposing a significant technical deficiency, the comment is to add new text and figures

that explain how to account for the skew compensation in much the same way that Figure 22 explains how to account for the strobe offset.

Accept: Added the following wording in 9.2.30

For paced data transfers negative values as measured at the device connector are accommodated by skew compensation in the receiver. Receive hold time measured at the device connector shall not exceed the skew correction range.

and in 9.2.34

For paced data transfers negative values as measured at the device connector are accommodated by skew compensation in the receiver. Receive setup time measured at the device connector shall not exceed the skew correction range.

Comment number: 022 (T) Document location: Figure 22 Comment: why is it OK in SPI-5 (and SPI-4) to not include the effects identified in SPI-3 in setting the timing reference points with respect to the signal crossing. In other words, why is setting the reference at the zero level justified here but not before? Proposed resolution: Offer some explanation for this change that is technically sound or change figure 22 to include the effects that were accounted for in SPI-3.

Accept: Added the following wording to 9.3.1

In paced transfers the timing budget and receiver masks account for the differences between the setup and hold detection thresholds that occur for synchronous transfers.

Comment number 023

When changing to a message phase after a DT DATA IN phase with a free running REQ clock (e.g. to send a QAS REQUEST message), it is possible for the skew between REQ and CD/IO/MSG to confuse both the initiator and other QAS-snooping devices on the bus. They can see CD/IO/MSG change before the last REQ edge. The target itself is not confused because it only makes the phase change after it stops sending REQ.

Corrections needed.

1. Any device monitoring the bus for message phases on CD, IO, and MSG after a DT DATA IN phase shall wait a minimum required filter time before honoring REQ.
2. The target device shall not send the real REQ for the message until a time longer than that filter time. [it already has to wait a deskew delay after driving DB]

Affected sections (no proposed wording changes yet)

10.7.1 Information transfer phases overview

...

Additionally, during the information transfer phases, the SCSI target port shall continuously envelope the REQ or ACK handshakes with the C/D, I/O, and MSG signals in such a manner that these control signals are valid for one bus settle delay before the assertion of the REQ signal of the first handshake and remain valid until after the negation of the ACK signal at the end of the handshake of the last transfer of the phase.

The SCSI target port shall not transition into an information transfer phase unless the REQ and ACK signals are negated. The SCSI target port shall not transition from an information transfer phase into another information transfer phase unless the REQ and ACK signals are negated.

NOTE 29 - After the negation of the ACK signal of the last transfer of the phase, the SCSI target port may prepare for a new phase by asserting or negating the C/D, I/O, and MSG signals. These signals may be changed together or individually. They may be changed in any order and may be changed more than NOTE 30 - A phase is defined as ending when the C/D, I/O, or MSG signals change after the negation of the ACK signal. The time between the end of a phase and the assertion of the REQ signal beginning a new phase is undefined.

10.7.2 Asynchronous transfer

If the I/O signal is true (i.e., transfer to the SCSI initiator port), the SCSI target port shall first drive the DB(7-0,P_CRCA) or DB(15-0,P_CRCA,P1) signals to their values, delay at least one system deskew delay plus one cable skew, then assert the REQ signal.

Accept: Added the following to the list in section 10.3

If paced transfers are enabled the C/D, I/O, MSG, DATA BUS, DB(P_CRCA), and DB(P1) signals shall not change for at least one system deskew delay from the negation of REQ.

Comments attached to No ballot from William Petty of LSI Logic Corp.:

1. 10.7.1 "Paced transfers shall only be used for a negotiated transfer rate of fast-160" Should be: Paced transfers shall only be used for a negotiated transfer rate of fast-160 or fast-320

Accept

2. 10.7.4.1 "data shall not be clocked by the originating SCSI device and" The receiver must ignore the invalid data, therefore this statement in not needed and could conflict with existing SPI-4 designs. This requirement was not part of SPI-4.

Reject: The indicated statement is in the approved version of SPI-4 which is rev 10.

3. 10.12 sub clause c Vs 10.7.4.3.4 The timing relationship of stopping the free-running REQ relative to changing the SCSI phase lines is not defined and assumed to be zero. This leaves no margin for cable skew. If REQ negates simultaneously with a phase change, it is possible for a device to observe REQ still asserted in the new phase.

We should use the standard "wait at least two system deskew delays" from negating the REQ before changing the state of MSG, CD, or IO.

Accept: Added the following to the list in section 10.3

If paced transfers are enabled the C/D, I/O, and MSG signals shall not change for at least one system deskew delay from the negation of the last REQ of a paced transfer data phase.

Comments attached to Yes ballot from Mark Evans of Maxtor Corp.:

Page 5

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 2:11:45 PM

Was this page left blank intentionally? If not, it should be deleted.

Page 13

T10/02-452 revision 2

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:04:05 AM

This should be changed from subscript to regular font.

Page 21

Comment 3; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:05:55 AM

The page number should be aligned properly.

Comment 4; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:06:54 AM

This should be changed from subscript to regular font.

Page 26

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:08:31 AM

I recommend that "transport protocol" be changed to "parallel interface."

Page 27

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:10:44 AM

I recommend that this sentence be changed to read, "This standard has made obsolete the single-ended and multimode signaling alternatives. Implementations that use single-ended or multimode signaling should reference the SCSI Parallel Interface-2 standard (ISO/IEC 14776-112)."

Page 30

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:13:30 AM

The definition for AAF should be deleted, as it is not used anywhere in the document.

Comment 3; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 2:11:16 PM

I recommend that SAM-3 be included, as well. I also think that references later in the document should be to SAM-3 as the latest version of this is the most up-to-date.

Reject: SAM-3 is going to leave parallel in the dust where SAM-2 has a parallel bent to it.

Page 31

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:21:55 AM

I recommend that, "asynchronous" be deleted as it is redundant.

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:31:35 AM

I recommend that this be changed to the following, "A response returned to the application client that signals the completion of a service request (see 4.13)."

Comment 3; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:23:48 AM

For proper alphabetical order, "cyclic redundancy check" should be moved after "current task".

Comment 4; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:25:46 AM

I recommend that this be changed to, "The task that is being executed. Execution may include sending messages, status, transferring data, or transferring command data to or from the SCSI initiator port.

Accept: The following change was made:

0.0.4 current task:A task that is in the process (i.e., of sending messages, status, transferring data, or transferring command data to or from the SCSI initiator port).

Comment 5; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:28:06 AM

I recommend that this read, "data group: A sequence of data bytes, any pad bytes, and the four pCRC bytes transmitted during a DT DATA IN phase or a DT DATA OUT phase."

Page 32

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:31:15 AM

I recommend that this be changed to the following, "A signal notifying the SCSI device server or task manager of an event (see 4.13)."

Page 33

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 9:53:39 AM

I recommend changing "highest" to "fastest".

Page 34

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:15:31 AM

I recommend that this sentence be changed to, "The act of resuming a nexus to continue executing a task."

Accept: But changed executing to processing.

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:16:02 AM

The word "substantially" should be deleted.

Page 35

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:32:39 AM

I recommend that this be changed to the following, "A request to the initiator parallel interface agent to invoke a service (see 4.13)."

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:33:12 AM

I recommend that this be changed to the following, "A response from the SCSI device server or task manager in reply to an indication (see 4.13)."

Comment 3; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:34:24 AM

The word. "match" should be changed to "matching".

Page 36

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:37:55 AM

"...from assertion edge..." should be changed to, "...from an assertion edge..."

Page 37

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:38:57 AM

"AAF" should be deleted as it is not used in the document.

Page 40

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:42:47 AM

Since there is only one transceiver type defined in this standard, I recommend that the column "Transceiver Type" be deleted and the table renamed. "LVD transceiver speed support map."

Page 41

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 10:51:45 AM

The word "directly" should be deleted.

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:21:00 AM

This phrase should be, "...on the other end..."

Page 44

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:22:46 AM

I recommend that "...entity..." and "...entities..." be replaced with "...parameter..." and "...parameters..."

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:21:49 AM

entities

Page 45

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 11:00:00 AM

The extra "at" in this sentence should be deleted.

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:24:58 AM

I recommend that the phrase, "...ns while ST is 25 ns." be replaced with, "...ns while for ST it is 25 ns."

Comment 3; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:26:37 AM

I'm unsure what this sentence is supposed to mean and recommend that it be clarified.

Accept: Change the sentence to: Data being transmitted using paced transfers is latched in the center of a bit cell, however, the relationship between the data and REQ or ACK is required to be shifted in the SCSI devices receiver to align REQ or ACK to the center of the data cell that then matches the synchronous transfers DT Data shown in figure 4.

Page 47

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:28:33 AM

I recommend that this sentence be changed to, "A SCSI bus that supports paced transfers has driver and receiver functions required in addition to those used with synchronous transfers or asynchronous transfers."

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:30:14 AM

I think that, in most cases, the signals being clocked are aligned to the clock signals, and that this sentence should be changed to say that.

Accept: This the new wording: That adjustment causes the signals being clocked to align with the middle of the clock signal when those signals enter the receiver (see figure 10).

Page 51

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 1:15:58 PM

Isn't it possible for an initiator/target port to negotiate one set of elements as an initiator and another as a target? I don't see any reason why these couldn't be different for each I_T nexus.

Reject: In reality when a port is an initiator it would be using one set of negotiation values for the targets it is talking and if that same port switched to a target it would now be talking to a set of different initiators so it would have to have a different set of negotiation values. But in the weird case where the addresses of the I_T stayed the same (i.e., both the I and the T changed but kept the same address) then the negotiation values would stay the same.

Page 59

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 1:43:20 PM

The phrase, "the target port shall not request tasks to be aborted for that SCSI initiator port..." should be changed to, "the target port shall abort any tasks for that SCSI initiator port..."

Accept: The following shows the change to the paragraph in question.

If the IU_REQ bit was set to one during a previous PPR negotiation and not changed by a subsequent PPR negotiation sequence, the target port shall not request that the task manager abort any tasks for that SCSI initiator port and shall go to the BUS FREE phase after responding with a PPR IN message.

Page 65

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 7:54:04 AM

The terms "bad parity" and "parity error" are used in many places in the document. However, the closest these come to being defined is in 3.1.62 odd parity. I recommend that the following be included in the Definitions: "Parity error: where the number of assertions on the associated data byte plus the parity bit equal to an even number (e.g., 0, 2, 4, 6, or 8)." Then, "bad parity" should be replaced with "a parity error" (or a similar phrase using "parity error" instead of "bad parity") throughout the document.

Page 78

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 2:16:19 PM

I recommend that the last sentence of this paragraph be deleted as this standard does not define single-ended transceivers.

Page 97

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 2:38:51 PM

The term "SPI devices" is not defined in this document. Therefore, I recommend that this be replaced with, "device on the parallel SCSI bus."

Page 102

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/16/2002 2:45:54 PM

I recommend that, "Not previously defined." be changed to something like, "Not defined in a previous SCSI standard."

Page 116

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 6:00:09 AM

"SE" should be included in the Acronym list (as is "SE").

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 6:02:57 AM

This should be changed to, "Neither SE nor HVD is defined in this standard. For information on SE SCSI device implementation see the SCSI Parallel Interface-4 standard. For information on HVD SCSI device implementation see the SCSI Parallel Interface-2 standard."

Page 149

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:37:02 AM

I recommend that the formatting of this list be adjusted.

Page 150

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:37:34 AM

I recommend that the formatting of this list be adjusted.

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:37:44 AM

I recommend that the formatting of this list be adjusted.

Page 151

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:39:47 AM

"cell" in "Tcell" n figures 66, 67, 69, 70 should be subscript.

Page 165

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 8:02:21 AM

I recommend that these paragraphs be reworded to something like: The SCSI device that won a normal arbitration shall assert both the BSY and SEL signals and wait at least one bus clear delay plus one bus settle delay before ending the NORMAL ARBITRATION phase.

The SCSI device that won a QAS shall assert the SEL signal and has wait at least one QAS arbitration delay before ending the QAS phase.

The SCSI device that won the arbitration shall identify itself as a SCSI target port by asserting the I/O signal.

Page 169

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:42:12 AM

This phrase should be, "When an information unit transfer agreement has..."

Page 170

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 8:34:09 AM

I recommend that the phrases, "If the nexus has been fully identified (i.e., an I_T_L_Q nexus has been established)..." and, "the nexus has been fully identified..." be replaced by, "If an I_T_L_Q nexus has been established..." (four places).

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 8:32:19 AM

If a nexus has been fully identified

Page 171

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 8:32:26 AM

If the nexus has been fully identified (i.e., an I_T_L_Q nexus has been established)

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 8:32:32 AM

If the nexus has been fully identified (i.e., an I_T_L_Q nexus has been established)

Page 172

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 8:56:14 AM

I think that this sentence is cumbersome. I recommend that it be changed to something like, "If the SCSI target port determines that a pad field is required, has completed the data field transfer of the current data group, the I/O signal is true (i.e., transfer to the SCSI initiator port), and the REQ signal is asserted, the SCSI target port shall:"

Page 173

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 8:57:05 AM

I think that this sentence is cumbersome. I recommend that it be changed to something like, "If the SCSI target port determines that a pad field is not required, has completed the data field transfer of the current data group, the I/O signal is true (i.e., transfer to the SCSI initiator port), and the REQ signal is negated, the SCSI target port shall:"

Page 177

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:43:37 AM

The list after, "Start of section C" should begin with a "1)".

Page 179

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:44:08 AM

The list after, "Start of section C" should begin with a "1)".

Page 190

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 9:46:49 AM

The wording here is obtuse. I recommend that these three paragraphs be replaced with something like:

Using the polynomial described above, the CRC calculated for 512 bytes containing 00h is: 55ADh, 190Ah.

Using the polynomial described above, the CRC calculated for 512 bytes containing FFh is: AB0Bh, FF6Ch.

Using the polynomial described above, the CRC calculated for 512 bytes containing values incrementing from 00h to 1Fh is: 7E8Ah, 9126h.

Accept: The following replaced those three paragraphs:

Using the polynomial described in 11.3.4, the CRC calculated for a 32-byte transfer of all 00h is: 55ADh, 190Ah.

Using the polynomial described in 11.3.4, the CRC calculated for a 32-byte transfer of all FFh: AB0Bh, FF6Ch.

Using the polynomial described in 11.3.4, the CRC calculated for a 32-byte transfer of an incrementing pattern from 00h to 1Fh is: 7E8Ah, 9126h.

Page 205

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 10:04:05 AM

These two paragraphs describe a condition based on a value in a field that is described later. So, I recommend moving them to after Table 48 - TASK MANAGEMENT FUNCTIONS.

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 10:08:25 AM

I think this means, "Termination of a SPI L_Q/SPI command information unit pair by a SCSI target device shall have no effect on any other SPI L_Q/SPI command information unit pair except those effects caused by any task management functions contained within the last SPI L_Q/SPI command information unit pair." If I am correct, this should be changed.

Accept: You are correct and it has been change to:

If the SCSI target device terminates a SPI L_Q/SPI command information unit pair for one of the following reasons:

- a) TASK SET FULL status,
- b) BUSY status,
- c) CHECK CONDITION due to a SPI command information unit iuCRC error, or
- d) a bus free due to a SPI L_Q information unit iuCRC error

it shall have no effect on any other SPI L_Q/SPI command information unit pair except those effects caused by any task management functions contained within the last SPI L_Q/SPI command information unit pair.

Page 213

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 10:19:36 AM

This should be changed to, "The value in the IUCRC field shall be calculated using the algorithm defined in 11.3.

Page 214

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 10:19:57 AM

This should be changed to, "The value in the IUCRC field shall be calculated using the algorithm defined in 11.3.

Page 217

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 10:21:49 AM

This should be changed to, "The value in the IUCRC field shall be calculated using the algorithm defined in 11.3.

Page 218

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 11:01:21 AM

This is another obtuse paragraph. I think that this is supposed to mean:

The SCSI initiator port provides a set of pointers for each send command service request received from an application client. The set of pointers for a unidirectional command consists of one pointer for the data (either data-in data or data-out data) and one pointer for the status. The set of pointers for a bidirectional command consists of one pointer for data-out data, one pointer for data-in data, and one pointer for status. When the set of pointers is for a task that is not active, the set is called the saved pointers for the task. When the set of pointers is for a task that is active, the set is called the active pointers for the task. When a task becomes active, the saved pointers for the task become the active pointers.

Accept partially: Changed the paragraph to:

The SCSI initiator port provides for a set of pointers for each task, called the saved pointers. The set of pointers for a unidirectional command consists of one pointer for the data (either data-in data or data-out data) and one pointer for the status. The set of pointers for a bidirectional command consists of one pointer for data-out data, one pointer for data-in data, and one pointer for status. When a send command service is received from an application client, the task's saved pointers are copied into the SCSI initiator port's set of active pointers. There is only one set of active pointers in each SCSI initiator port. The active pointers point to the next command, data-out, data-in, or status byte to be transferred between the SCSI initiator port and the SCSI target port. The saved and active pointers reside in the SCSI initiator port.

Page 224

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 11:13:12 AM

I recommend this sentence be changed to, "Identification is considered successful during a SCSI initiator port's initial connection or a physical reconnect when the SCSI target port detects no error during the

transfer of the IDENTIFY message and an optional task attribute message in the MESSAGE OUT phase following the SELECTION phase.

Page 231

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 11:25:10 AM

I recommend changing this to, "...a tag not associated with a task in the task set,"

Page 234

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/17/2002 11:34:32 AM

This is an unclear sentence that should be replaced with something like, "On a physical reconnection, the ABORT TASK message causes the current task to be aborted if an I_T_L_Q nexus exists. If only an I_T_L nexus exists, but the SCSI target port is performing a physical reconnect for an I_T_L_Q nexus, then the current task is not aborted and the SCSI target port changes to the BUS FREE phase."

Accept: Changed the paragraph to:

On a physical reconnection, the ABORT TASK message causes the current task to be aborted if an I_T_L_Q nexus has been fully identified. If an I_T_L nexus exists, but the SCSI target port is doing a physical reconnecting for an I_T_L_Q nexus, then the current task is not aborted and the SCSI target port goes to the BUS FREE phase.

Page 260

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 6:41:50 AM

The subscript for the "N" in VN isn't very distinct in four places in tables A2 and A3.

Page 269

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 6:43:40 AM

I think that this should be changed to, "A receiver determines..."

Page 274

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 6:46:02 AM

This should be changed to, "Implementation of SCSI bus fairness...", i.e., the word "the" should be deleted.

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 6:52:55 AM

There are several places in this annex where the phrases, "...need to..." and "...needs to..." are used. I recommend that these be deleted. An example of a recommended replacement is: "When a SCSI device arbitrates for the SCSI bus..."

Accept: But it was no as simple as removing the needs to and need. This needs to be looked at to make sure I did not mess things up.

Comment 3; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 6:57:41 AM

I recommend that, "...begin to disappear..." be replaced by, "...are released..."

Comment 7; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 6:59:38 AM

I recommend that a comma be added between "false" and "the".

Comment 8; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:00:15 AM

I recommend that the comma be deleted between "time" and "to".

Comment 9; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 6:58:48 AM

I recommend that, "...begin to disappear..." be replaced by, "...are released..."

Page 278

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:01:49 AM

I recommend that, "...defined in the SCSI device's specification." be changed to "...are vendor specific."

Page 294

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:06:06 AM

Later in this annex is a phrase, "...a coupling device called a balun...". If that clarification is necessary, it should be included here (where "balun" first appears).

Page 301

Comment 1; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:08:49 AM

"Balun" is used in many places earlier in this annex. If this clarification phrase is needed, it should be added where first used. One way or the other, the phrase "...a coupling device called a balun..." should be deleted here and replaced with "balun".

Comment 2; Label: Maxtor Corporation, Mark Evans; Date: 10/21/2002 7:09:45 AM

The word "effectively" should be deleted.

Comments attached to Abs ballot from Edward A. Gardner of Ophidian Designs:

I have little knowledge of or experience with recent versions of parallel SCSI and do not feel competent to evaluate this standard.

Comments attached to Yes ballot from Paul D. Aloisi of Texas Instruments:

1. (Clarification) Doesn't this section of the scope have to change without Multimode to exclude the electrical interface compatibility with SCSI-2?

Section 1 Scope

b) To provide compatibility such that conforming SCSI-2, SPI-2, SPI-3, SPI-4 devices may interoperate with SPI-5 devices given that the systems engineering is correctly done. Conforming SCSI-2, SPI-2, SPI-3, and SPI-5 devices should respond in an acceptable manner to reject SPI-5 protocol extensions. SPI-5 protocol extensions are designed to be permissive of such rejections and thus allow SCSI-2, SPI-2, SPI-3, and SPI-4 devices to continue operation without requiring the use of the extensions.

Add

b) To provide compatibility such that conforming SCSI-2 (**Through a Single ended to LVD SCSI expander**), SPI-2, SPI-3, SPI-4 devices may interoperate with SPI-5 devices given that the systems engineering is correctly done. Conforming SCSI-2, SPI-2, SPI-3, and SPI-5 devices should respond in an acceptable manner to reject SPI-5 protocol extensions. SPI-5 protocol extensions are designed to be permissive of such rejections and thus allow SCSI-2, SPI-2, SPI-3, and SPI-4 devices to continue operation without requiring the use of the extensions.

Accept: Removed SCSI-2 from the paragraph.

2. (Editorial?) Figure 22, Shouldn't **SCSI Parallel Interface —4 Standard** be changed to **-5**

Accept

3.(Editorial) Drop the last sentence in 5.1, this standard doesn't support single ended or multimode.

5.1 SCSI parallel interface connectors overview

Two types of connectors are defined: nonshielded and shielded. The nonshielded connectors are typically used within an enclosure. The shielded connectors are typically used for external applications where electromagnetic compatibility (EMC) and electrostatic discharge (ESD) protection may be required. ~~Either type of connector may be used with the single ended or differential transceivers.~~

Accept

4. (Clarification) Clauses 6.3.4 through 6.3.8 should reference PIP to test fixtures and procedures. The one general reference to annex E is not adequate in 6.3.1

Add the following sentence in 6.3.4 through 6.3.8

The test fixtures and test procedures for the measurements are found in the PIP standard.

Accept: Added the following paragraph to section 6.3.1:

The test fixtures and test procedures for the measurements specified in clause 6 see the SCSI Passive Interconnect Performance standard.

5. (Editorial) This needs to be reworded, since there is no way to disable precomp in Fast-320

9.4.3 Paced transfers with precompensation disabled on fast-160 and fast-320

During paced transfers with a negotiated transfer rate of fast-160 ~~or fast-320~~ and precompensation disabled ~~or fast-320~~ receiving SCSI devices shall operate with signals that meet all the following requirements at the receiver.

Accept

6. (Editorial) **There are several references for distance called out in cm, the international standards for mechanical measurement should be in mm. This avoids confusion and errors, like 0.5 cm should be 5 mm.**

Accept

7. (Clarification) Table 1 needs to be expanded to give a better understanding of the changes in technology.

Transceiver Type	Data Latching (ST/DT)	Maximum transfer rate							
		Asyn c	Fast-5	Fast-10	Fast-20	Fast-40	Fast-80	Fast-160	Fast-320
SE	ST	yes	yes	yes	yes	no	no	no	no
MSE	ST	yes	yes	yes	yes	no	no	no	no
LVD	ST	yes	yes	yes	yes	yes	no	no	no
	DT	no	no	SPI-3	SPI-3	SPI-3	yes	Paced	Paced

SE = Single ended transceivers, MSE = multimode single ended transceivers,
LVD = Low Voltage differential transceivers.

ST = Single Transition clocking on the falling edge of REQ or ACK.

DT = Double Transition clocking on both the falling and rising edge of REQ or ACK.

SPI-3 and beyond allows DT clocking for speeds of Fast-10 or greater.

Paced = Paced data transfer mode reduces the protocol overhead with packetized commands and messages, Quick Arbitration and selection, Free running clock, training pattern, skew compensation and ISI compensation.

Reject most: We have eliminated SE and MSE from this standard so it should not only be placed into this table. I did add in the paced retries and a footnote about that term.

Comments attached to No ballot from Jie Fan of TycoElectronics:

We would like to delay the public review of this standard until all related documents such as PIP are settled down. Changes may be required made to this standard.

Reject: We are not going to tie SPI-5 to the adoption on another standard.

Comments attached to Abs ballot from Roger Cummings of Veritas Software:

Not within our organizations scope of expertise

***** End of Ballot Report *****