To: GOP@us.IBM.COM @ INTERNET T10/00-118r0
cc: t10@t10.org@internet
From: Gene Milligan
Date: 12/15/99 12:08:33 PM
Subject: Comment Resolution

Please excuse the extra hyphens in the stuff below. They are an artifact of copying from the pdf.

<<Accept: Proposal A wording placed into section 4.8.3.1>>

I can not find that section. Well I guess I did, it is 4.9.3.2.

<<SPI-3: 29.25) 25) (T) Page 192, Annex B:
Replace entire annex with:>>

There does not appear to be a stated resolution of this comment.

In Annex B a reference to 10.5.1 should be given in <<Whenever a requirement for arbitration arises, a SCSI device shall first check to see if its fairness register is clear.>>

<<Therefore, only those lower priority SCSI devices latched into the fairness register at that time arbitrate ahead of this SCSI device.>>

I think this statement is wrong. SCSI devices that have not previously arbitrated may arbitrate at this time and all SCSI devices not using fairness may arbitrate.

<<to detect which SCSI devices are attempting arbitration,>> should be "to detect SCSI devices that are attempting arbitration," or "to detect SCSI devices that are arbitrating,".

<<the data bus must be sampled after a bus set delay>> should be "the data bus shall be sampled after a bus set delay". But on the other hand this seems redundant to the earlier part of the paragraph and probably should be deleted.

Annex B, new I guess, uses <<nsec>>. I assume the rest of the document uses ns and Annex B should use ns.

<<A SCSI device shall implement a lockout delay to prevent devices that stop
arbitrating from causing deadlock.>>

There should be a reference to where the lockout delay is defined and how it is used to prevent deadlock.

<<Interconnection of SCSI devices by means other than cables is allowed (e.g., by backplanes using printed wiring boards) (see annex J). Detailed descriptions of these other means are not part of this standard; how-ever, they are subject to the electrical requirements in sections 6.3.
Accept>>

It is not clear to me what impact this had on maximum lengths.

<<Attach a 122 ohm resistor to the far end of the pair under test.
Accept>>

Shouldn't this be "Attach a 122 ohm resistor across the far end of the pair under test."

<<Board impedance shall be tightly controlled to within
5% of the impedance of the environment.>>

What is the environment or where is it defined?

<<2) Connect a wire (short) to the sockets of the test fixture and perform a “short” calibration as specified by HP.>>

Shouldn't there be a reference to where HP specifies this and how to obtain the specification. Shouldn't an equivalent be allowed. Are we liable to a suit from LeCroy?

<<The text immediately following Figure 16 is incorrectly states as 50W single ended to 150W differential. The 150W should be changed to 122W differential.
Accept>>

The draft is OK but the comment should be ohms.

<<Change L.5 as follows:
The READ BUFFER and WRITE BUFFER commands access physical buffers in the target. Many implementations do not protect the buffer contents if there is an intervening command from any other process. Therefore, the application client should ensure that no other SCSI processes are active while performing tests.

The READ BUFFER and WRITE BUFFER commands include an echo buffer option that provides buffer protection in multi-initiator environ-ments. Other mechanisms that may help prevent buffer corruption in multiple initiator environments are RESERVE/RELEASE and linked commands.
Accept>>

Delete "Many".

<<d.5) After the bus free delay in step (b), SCSI devices with arbitration fairness enabled which are not arbitrating shall wait a bus set delay and start sampling the DATA BUS to determine which SCSI de-vices attempted arbitration, which SCSI device won, and which SCSI devices lost. This sampling shall continue for an arbitration delay after the bus free delay in step (b). Each SCSI device shall update its fairness register with all lower-priority device IDs that lost arbitration.>>

Did you de-which this?

<<f) After the QAS arbitration delay in step (d), SCSI devices with arbitration fairness enabled which are not arbitrating shall start sampling the DATA BUS to determine which SCSI devices are attempting arbitration, which SCSI device won, and which SCSI devices lost. This sampling shall continue for a bus settle delay plus two system deskew delays. The devices shall update their fairness register with all device IDs that lost in arbitration.
Accept>>

Did you de-which this?

<<Change Verilog to “Verilog® Hardware Description Language (IEEE 1364)”
Accept>>

Did this result in an additional normative reference or an additional informative reference?

<<In the first sentence “...though reference in the text...” should be changed to “...though referenced in the text...”

Accept>>

I think the comment was a misread and the change should be removed. "through reference in the text," is correct. Other alternatives would be OK such as "where referenced in the text" or "by reference in the text".

<<All of the figures in this Annex (except E.6) are at least somewhat blurry (they look even worse in my printed version of the PDF). Is there anything that can be done about this?

The editor can only work with what he receives. This is especially true when it comes to complex figures. However, on review the PDF on may machine the figures do not appear to be all that bad.>>

May should be my. The images are usable on my machine with Fit Visible and just the page showing. They appear to be scanned images and do show some shadowing that makes them less sharp. Apparently the figures are scanned images. I point this out since I just learned today the ISO/IEC central secretariat does not like scanned images due probably to the blurring. They are after me to redraw all the scanned figures in two standards. Pete McLean has redrawn scanned images in ATA/ATAPI and I have an Email in to him to find out how he does that. In my case the source of the originals has disappeared. I suggest before the source evaporates that you obtain the original to avoid second generation copies.

<<2) Do we still like the name of T10 as shown on page 21?

As far as I know that is still what we are called.>>

OK for the ballot but I will propose at the T10 meeting that we get rid of the dumb name Technical Committee T10 on Lower Level Interfaces. T11, T12, and T13 work on the same levels. Perhaps Technical Committee T10 on SCSI.

<<17) The use of <<input-1a and input-1b>> is not explained as the stated convention causes the reader to expect “input-1 and input-2”.
Accept: Changed - to :. But the convention is described right below the expression.>>

Sorry if my comment was not clear. The example uses 1a and the convention uses 1. The explanation for the a, b, c is what I was looking for. It is hard to get 17c from n. Is inout a typo or an odd name? I do not understand Changed - to :.

<<31) Consider using IEC specifications rather than EIA.
Accept: Changed all normative references to ISO in section 2.>>

You went beyond my request but that is OK with me. It may be better to change the SCSI ones back to ANS. If not the resolution should be Changed all normative references to IEC and/or ISO/IEC in section 2. If you do change them back the resolution should either be just plain Accept or Changed all EIA normative references to IEC in section 2.

<<35) The drafting standard used for figures 11-13, 17-19, and 21-22 should be added to the normative refer-ences (preferably the international version).
Reject: If I have such a thing I would. But we have several standards that are now ISO that have these same figures and have not had any requests for a reference to a standard.>>

I reject your rejection and raise you one. I do not think errors in the past should require errors in the future. T10 should ask the authors of these figures to identify the standard so that individual engineers without the knowledge do not have to hire a consultant to find out what they mean.

<<49) In table 13 why is there both note 1 and note 3? Note 1 includes the concept of note 3.

In table 13 and in table 14 add to the remaining note or both notes "If the maximum cable requirements are exceeded, SCSI devices shall not be required to conform to the electrical and timing requirements of this standard.
Accept removing note 3 from table 13; Reject adding additional suggested wording to tables 13 and 14.>>

I recall Bill Ham agreeing to the additional wording in a recent SPI working group (I think it was the Monterey meeting) and thought the working group reached consensus on the additional wording.

<<64) Delete <<NOTE 14 - When using only the SCA-2 connector (see 5.2.4) the SE, LVD, and HVD connec-tor contact numbers allow switching between all three modes. In this case the terminator may switch to HVD mode if so indicated by the DIFFSENS line.>> or change it to "NOTE 14 - When using only the SCA-2 connector (see 5.2.4) the SE, LVD, and HVD connector contact numbers allow switching between all three modes. In this case the terminator may switch to HVD mode if so indicated by the DIFFSENS line and the SCSI devices should be replaced with HVD SCSI devices."
Reject: We have no control over irrational behavior.>>

I think you have missed the central point of the comment. <<the SE, LVD, and HVD connector contact numbers allow switching between all three modes. >> There are no SCSI devices that can switch between these three modes.

<<In this case the terminator may switch to HVD mode if so indicated by the DIFFSENS line.>> Are there any terminators that switch to HVD mode? As far as I know HVD termination is not switchable between modes.

My suggestion to delete was a good one. My other suggestion was flawed. If an alternative is needed to deletion I now suggest as an alternative "NOTE 14 - When using only the SCA-2 connector (see 5.2.4) the SE, LVD, and HVD connector contact numbers accommodate all three modes. Switching between the LVD and the SE modes may occur with multi-mode devices if so indicated by the DIFFSENS line. HVD mode requires all devices to be HVD devices."

I have used devices rather than SCSI devices to encompass both SCSI devices and terminators and do not recall if SPI has a better term for that.

<<67) Regarding <<The TERMPWR lines may be used for distribution of power for purposes other than for SCSI bus termination as long as the voltage delivered to the SCSI bus terminators remains adequate to supply the requirements of the terminators under all conditions of SCSI bus operation and under all conditions of other loading.>> This should be deleted or a current limit added. I will open the bidding at 0 amps.
Reject - Not change made>>

I reject the rejection and raise you one. The rejection should be No change made, but I still reject it. The issue is the standard does not state what entity is responsible for making the voltage adequate. If the requirements for achieving this undefined configuration are not given, the idea should also not be given. This is strictly a closed system definition and is not needed in the standard. Perhaps it could be added to SPI-4 if the requirements were proposed and agreed to making it an open system requirement. I use open system in terms of multi-vendor and not in the sense of the open system model.

<<The proposal, that was accepted for DT added the definition of REQ (ACK) period. It included the mea-surement concept of using a regular signal (without offsets) to measure the time to avoid ISI. That concept seems to have disappeared and should be restored.
Reject for now: The change in question is the change from the REQ (ACK) period being an average vs. a minimum as is now specified. This change was made as a result of a SPI-3 rev 4 review in the April 1999 working group meeting.>>

The problem is the change made, apparently without a proposal failed to take into account the original accepted proposal that added the definition and the need for the definition. SPI-2 defined and specified a REQ(ACK) Period tolerance without having defined a REQ(ACK) Period. The tolerance is used in the timing budget. The tolerance was to account in the budget for differences in the transfer period as a result of clock accuracy versus the negotiated transfer period. The tolerance was not intended to include ISI, distortion, or noise effects. A regular REQ(ACK) (without offsets) and averaging is needed to make the measurement while eliminating or at least minimizing these other unwanted effects. The other effects are separately specified in the budget.

<<80) <<however the SCSI device shall not arbitrate (i.e. assert the BSY signal and its SCSI ID) if more than a bus set delay has passed since the BUS FREE phase was last observed.>> This sounds like a deadly embrace. It should probably be changed to (comma for sure) "however the SCSI device shall not arbitrate (i.e., assert the BSY signal and its SCSI ID) if more than a bus set delay has passed since the BUS FREE phase was last observed and the SCSI bus is not BUS FREE."
Reject: (Except the comma). The may appear to be bus free but it may not be because of propaga-tion delays on the singles. If we make this change the max timing (i.e., bus set delay) is meaning-less.>>

Suggest The bus may appear, but what allows the SCSI device to arbitrate again if the SCSI BUS is BUS FREE for a fortnight?

<<91) Regarding <<10.5.2.2.2 Data Group data field transfer>> and wherever this style of describing data transfer is used, I question if using a minimum of transmit setup and transmit hold time is a good way to describe the mechanism. If a designer actually transferred with this timing zero SCSI device skew would be required. The transfer time should be targeted at the ideal setup time (half of the ideal bit time). Perhaps this could be described as half of the negotiated bit period or half of the transfer period.
Reject - Because there is already a statement that states these timing requirements shall not be used.>>

Where?

<<107) What is the meaning of both SCSI devices in <<The transfer width that is established applies to all logical units on both SCSI devices.>>?
Accept: That can from WDTR. But it is not right there either. The ’to all logical unit’ has been removed from PPR and WDTR.>>

OK although I am having trouble reading your response.

<<111) Change <<A PARALLEL PROTOCOL REQUEST message exchange shall be initiated by a SCSI device whenever a previously arranged parallel protocol agreement may have become invalid.>> to "A PARALLEL PROTOCOL REQUEST message exchange shall be initiated by an initiator whenever a previ-ously arranged parallel protocol agreement may have become invalid."
The present wording is equivalent to the objectionable SDTR requirement that most targets are required to violate.
Reject: Your new wording prevents targets from initiating PPR. This could cause improper commu-nications to occur between the SCSI devices. The fact that some poorly designed initiators cannot handle a target initiated PPR is not reason to remove it from the standard.>>

I understand that this is no longer rejected and that the resultant change is still being massaged.

<<Accept: Add the follow to the PPR message:>>

Try following.

<<The current way it is specified was the way to was agreed to be changed to. No change.>>

the way it was agreed to

<<130) Why include <<NOTE 45 - The details of the actual SCSI device supply requirements need to be studied for each SCSI device and enclosure combination.>>?
Note needed to reduce risk of problems in systems having 3.3 and 5 volts available. Note change to:
NOTE 1 - The details of the actual SCSI device supply requirements need to be studied for each SCSI device and enclosure combination to ensure appropriate supply voltage to the SCSI device.>>

Well I prefer deleting the note. Studying does no good if you just study and do nothing. If T10 insists on a worthless note perhaps it could have imaginary value by changing it to "NOTE 45 - The details of the actual SCSI device supply needs to be tailored for each SCSI device and enclosure combination to ensure appropriate supply voltage to the SCSI device and terminators."

Feel free to change tailored to adjusted if you prefer.

<<132) Regarding <<NOTE 46 - Industry practice presently requires that SCSI devices interconnected for synchronization be the same or equivalent models.>> Actually industry practice is that spindle synchronization is obsolete. Consider obsoleting this clause.
Accept: Removed the note but not the clause. Although if someone was to make a motion to able the spindle sync line obsolete I would not object!>>

So moved. In the mean time it should be motion to make.

140) In E.7.1 I think input and output are not used in the same sense as in the rest of the standard where input would be found at the receiving SCSI device and the output at the transmitting SCSI device. I think a
definition should be added.
Reject: The definition of input and output is defined in the first line of section E.7.1.

<<Attenuation is calculated from the ratio of output signal level to input signal level through the DUT and is a measure of the losses experienced when transmitting a signal through the DUT.>> That does not taste like a definition of input and output. Perhaps it should be "Reject: You will be able to recognize input and output when you see it. Sort of like porn.

<<151) Regarding J.1.6 what is <<odd mode>>?
There is a definition in the first paragraph of J.1.6.
The two conductors are symmetrical therefore the two self inductance's are equal and the two capacitors to ground are equal (Z 11 =Z 21). The voltage applied between the two lines may be thought of as a superposition of two voltages, a common mode voltage V c (even mode) and a differential mode voltage V d (odd mode) where: V c = (V 1 +V 2)/2 and V d = (V 1 -V 2)/2.

There wasn't but I gather there is now in the revised J.1.6 to be found in the new second paragraph in the form of the explanation above. Thanks.

Well we are much closer now but you will have to hold off on the cigar,

Happy holidays and best of luck in the new year (with the significant new numeric) waiting quite a while for the new century and new millennium,

Gene