T10/02-382r0

Voting Results on T10 Letter Ballot 02-381r0 on
Forwarding SPI-5 to First Public Review
Ballot closed: 2002/10/21 12:00 noon MDT

| Organization | Name | S | Vote | Add'l Info |
|---|---|---|---|---|
| Adaptec, Inc. | | | DNV | |
| Amphenol Interconnect | Michael Wingard | P | Yes | |
| Andiamo Systems, Inc. | Claudio DeSanti | P | Yes | |
| BREA Technologies, Inc. | Bill Galloway | P | Yes | |
| Brocade | Brian Forbes | P | Yes | |
| Cisco Systems, Inc. | David Peterson | P | Yes | |
| Congruent Software, Inc. | Peter Johansson | P | Yes | |
| Crossroads Systems, Inc. | John Tyndall | A | Yes | |
| Dallas Semiconductor | James A. Lott, Jr. | P | Yes | |
| Dell Computer Corp. | Kevin Marks | P | Yes | |
| EMC | Gary S. Robinson | P | Yes | |
| Emulex | Robert H. Nixon | P | Yes | |
| ENDL | Ralph O. Weber | P | Yes | |
| Exabyte Corp. | Joe Breher | P | Yes | |
| FCI | Douglas Wagner | P | Yes | |
| Fujitsu | Mike Fitzpatrick | P | Yes | |
| General Dynamics | Nathan Hastad | P | Yes | |
| Hewlett Packard Co. | William Ham | A | No | Cmnts |
| Hitachi Cable Manchester | Zane Daggett | P | Yes | |
| Honda Connectors | | | DNV | |
| IBM Corp. | George O. Penokie | P | Yes | |
| Intel Corp. | Cris Simpson | P | Yes | |
| Iomega Corp. | Tim Bradshaw | P | Yes | |
| KnowledgeTek, Inc. | Dennis Moore | P | Yes | |
| LSI Logic Corp. | William Petty | A | No | Cmnts |
| Maxtor Corp. | Mark Evans | P | Yes | Cmnts |
| Microsoft Corp. | Emily Hill | P | Yes | |
| Molex Inc. | Jay Neer | P | Yes | |
| Network Appliance Inc. | James R. (Bob) Davis | P | Yes | |
| Nishan Systems Inc. | Charles Monia | P | Yes | |
| Ophidian Designs | Edward A. Gardner | P | Abs | Cmnts |
| Panasonic Technologies, Inc | Terence J. Nelson | P | Yes | |
| Philips Electronics | William P. McFerrin | P | Yes | |
| Pirus Networks | Milan J. Merhar | A | Yes | |
| QLogic Corp. | Skip Jones | P | Yes | |
| Quantum Corp. | Paul Entzel | P | Yes | |
| Seagate Technology | Gerald Houlder | P | Yes | |
| Storage Technology Corp. | | | DNV | |
| Sun Microsystems, Inc. | Vit Novak | P | Yes | |
| Texas Instruments | Paul D. Aloisi | P | Yes | Cmnts |
| Toshiba America Elec. Comp. | Tasuku Kasebayashi | P | Yes | |
| TycoElectronics | Jie Fan | P | No | Cmnts |
| UNISYS | Phil Shelton | A | Yes | |
| Veritas Software | Roger Cummings | P | Abs | Cmnts |
| Vixel Corp. | Kenneth Hirata | P | Yes | |
| Western Digital Corporation | | | DNV | |

Ballot totals: (37:3:2:4=46)
 37 Yes
  3 No
  2 Abstain
  4 Organization(s) did not vote
 46 Total voting organizations
  2 Duplicate ballot(s) not counted
  8 Ballot(s) included comments

This 2/3rds majority ballot passed.
37 Yes are more than half the membership eligible to vote minus abstentions
  [greater than 22] AND
37 Yes are at least 27 (2/3rds of those voting, excluding abstentions [40]) AND
37 Yes are equal to or exceed a quorum [15]

Key:

P Voter is principal member
A Voter is alternate member
Abs Abstain vote
DNV Organization did not vote
Cmnts Comments were included with ballot
NoCmnts No comments were included with a vote that requires comments

**********

Comments attached to DUPLICATE Abs ballot from Peter Johansson of
Congruent Software, Inc.:

My abstention is because of a lack of technical expertise in the subject
matter.

**********

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

SPI-5 comments: HP

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.


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."


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


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


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

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

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

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

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"

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."

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

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""

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.

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

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 interoperability 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.

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.

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

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.

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 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

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

Comment number: 021 (T)
Document location: Figure 22 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.

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.

**********

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:
Pacecd transfers shall only be used for a negotiated transfer rate of fast-160 or fast-320

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.

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.

**********

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

Comments are in T10/02-425r0.

**********

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:

Comments are in T10/02-407r0.

**********

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.

**********

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

Not within our organizations scope of expertise

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