

Accredited Standards Committee*
National Committee for Information Technology Standards (NCITS)

Doc. No.: T10/97-272r0

Date: November 11, 1997

Project: 1142-D

Ref. Doc.:

Reply to: John Lohmeyer

To: Membership of T10

From: Ralph Weber, Secretary T10
Larry Lamers, Vice-chair T10
John Lohmeyer, Chair T10

Subject: Minutes of Fast-xx Study Group and SPI-2 Working Group
November 3, 1997 -- Palm Springs, CA

Agenda

1. Opening Remarks
2. Approval of Agenda
3. Attendance and Membership
4. Fast-xx Study Group {Monday Morning}
 - 4.1 Low-Attenuation Cables (97-213, 97-228) [Rogers, Bellino]
 - 4.2 Cable Attenuation & Margin Budget Analysis at 80 MHz (97-269) [Gintz]
 - 4.3 Bias Reduction Proposal (97-214) [Bridgewater]
 - 4.4 Dual Clocking Proposal (97-208) [Bastiani]
 - 4.5 Project Proposal for SPI-3 (97-262) [Lamers]
 - 4.6 Fast 80 MHz SCSI Testing (97-276) [Bastiani]
5. SPI-2 Working Group {Monday Afternoon}
 - 5.1 SPI-2 specifications on interconnect (97-235) [Ham]
 - 5.2 0.5 pF balance capacitance for LVD SCSI (Table 29 SPI-2R13) (97-236) [Aloisi/Kapraun]
 - 5.3 Capacitance Measurement at 2 Volts [Wallace]
 - 5.4 Universal Backplane Annex (97-121r0) [Wallace]
 - 5.5 Latching and Counting (97-216) [Guss]
 - 5.6 SPI-2 Revision 16 Review [Penokie]
 - 5.7 SPI-2 Timing Tolerance (97-268) [Asami]
 - 5.8 Load Compensation (97-281) [Novak]
 - 5.9 Proposal to remove clause 10.5 from SPI-2 Rev 16 [Lohmeyer]
6. Meeting Schedule
7. Adjournment

Results of Meeting

1. Opening Remarks

John Lohmeyer, the T10 Chair, called the meeting to order at 9:00 a.m., Monday November 3, 1997. He thanked Norm Harris of Adaptec for hosting the meeting.

As is customary, the people attending introduced themselves and a copy of the attendance list was circulated.

2. Approval of Agenda

The agenda was approved with the addition of Greg Kapraun to item 5.2:

5.2 0.5 pF balance capacitance for LVD SCSI (Table 29 SPI-2R13) (97-236) [Aloisi/Kapraun]

The following agenda items were added during the meeting:

4.6 Fast 80 MHz SCSI Testing (97-276) [Bastiani]

3. Attendance and Membership

Attendance at working group meetings does not count toward minimum attendance requirements for T10 membership. Working group meetings are open to any person or organization directly and materially affected by T10's scope of work. The following people attended the meeting:

Name	S	Organization	Electronic Mail Address
Mr. Lawrence J. Lamers	P	Adaptec, Inc.	ljlammers@ix.netcom.com
Mr. Vincent Bastiani	A#	Adaptec, Inc.	bastiani@corp.adaptec.com
Mr. Tak Asami	V	Adaptec, Inc.	asami@itc.adaptec.com
Mr. Wally Bridgewater	V	Adaptec, Inc.	wally@eng.adaptec.com
Mr. Bill Gintz	V	Adaptec, Inc.	bgintz@corp.adaptec.com
Mr. Tom Schneider	V	Adaptec, Inc.	schneid@itc.adaptec.com
Mr. Charles Brill	P	AMP, Inc.	cebrill@amp.com
Mr. John Laskowsky	V	Berg Electronics	laskowj@bergelect.com
Mr. Domininc K.V. Zarrella	V	Berg Electronics	zarreldv@bergelect.com
Mr. Richard Wagner	P	Cable Design Technologies	rwagner@montrose-cdt.com
Mr. Edward Haske	O	CMD Technology	haske@cmd.com
Mr. Charles Tashbook	P	Dallas Semiconductor	charles.tashbook@dalsemi.com
Mr. Greg McSorley	P	Data General / Clariion	greg_mcsorley@dgc.ceo.dg.com
Dr. William Ham	A#	Digital Equipment Corp.	bill.ham@digital.com
Mr. George Penokie	P	IBM Corp.	gop@us.ibm.com
Mr. Tim Bradshaw	P	Iomega Corp.	bradshat@iomega.com
Mr. Dean Wallace	P	Linfinity Micro	75671.3443@compuserve.com
Mr. Louis Grantham	A	Linfinity Micro	lgrantham@linfinity.com
Mr. Alan Littlewood	P	LSI Logic Corp.	alanl@lsil.com
Mr. Makesh Kothandaraman	V	Lucent Technologies	makesh@lucent.com
Dr. Bernard Morris	V	Lucent Technologies	blmorris@lucent.com
Mr. Brian Davis	V	Mylex Corp.	briand@mylex.com
Mr. Skip Jones	P	QLogic Corp.	sk_jones@qlc.com
Mr. Ting Li Chan	A	QLogic Corp.	t_chan@qlc.com
Mr. Patrick McGarrah	P	Quantum Corp.	pmcgarra@tdh.qntm.com
Mr. Mark Evans	A	Quantum Corp.	mevans@qntm.com
Mr. Bruce Leshay	V	Quantum Corp.	bleshay@tdh.qntm.com
Mr. Richard Uber	V	Quantum Corp.	duber@tdh.qntm.com

Mr. John A. Fobel	O	Rancho Technology, Inc.	johnf@rancho.com
Mr. Gene Milligan	P	Seagate Technology	Gene_Milligan@notes.seagate.com
Mr. Gerald Houlder	A	Seagate Technology	Gerry_Houlder@notes.seagate.com
Mr. Daniel (Dan) F. Smith	O	Seagate Technology	daniel_f_smith@notes.seagate.com
Mr. Vit Novak	A	Sun Microsystems, Inc.	vit.novak@sun.com
Mr. John Lohmeyer	P	Symbios Logic Inc.	john.lohmeyer@symbios.com
Mr. Ralph O. Weber	A	Symbios Logic Inc.	roweber@acm.org
Mr. Paul D. Aloisi	P	Unitrode Corporation	aloisi@unitrode.com
Mr. Gregory Kapraun	V	Western Digital Corp.	Gregory.D.Kapraun@wdc.com

37 People Present

Status Key: P - Principal
 A,A# - Alternate
 O - Observer
 L - Liaison
 V - Visitor

For the benefit of those present at this meeting who will not be attending the T10 plenary meeting, John described the pilot membership program developed for T10 by the T10 leadership and the NCITS secretariat. Details of John's presentation can be found in the minutes of the T10 plenary meeting (97-279).

4. Fast-xx Study Group {Monday Morning}

4.1 Low-Attenuation Cables (97-213, 97-228) [Rogers, Bellino]

Neither Kirk Rogers nor Bob Bellino were present. John Lohmeyer mentioned a presentation from Bob Bellino on Lower Attenuation Cables (97-228r0) that was in the July mailing, but not discussed at any meeting. He expressed the belief that the data was beneficial to those concerned with cable characteristics.

4.2 Cable Attenuation & Margin Budget Analysis at 80 MHz (97-269) [Gintz]

Bill Gintz presented a revised timing budget analysis (97-269, revised from 97-207). Bill reported his conclusions first. He described the Fast-80 amplitude and jitter budgets results to be shown as covering: analytical projections, worst case margins 2.5 ns/6.25 ns, multi-drop effect, reflections and ISI, and a hot plug simulation. He described the key features of his empirical results as:

Reflections dominate jitter,
 Multi-drop dominates amplitude,
 Skew and ISI are unwieldy, and
 Hot plug is problematical.

For about an hour, the group discussed Bill's analysis: results, methodology, and assumptions. Bill's summary was as follows:

1. **Dual-edge clock still seems to be the best solution** - The group generally agreed with this conclusion.
2. **Cable skew and differential capacitance dominate offsets** - There was no argument on this point either.
3. **Cable induced ISI still enigmatic – more data is being taken** - Bill Gintz and Bill Ham agreed that testing is important and that the data gathered thus far are insufficient. Bill Gintz expressed the hope that one of the cable companies will do the needed testing. Bill Ham expressed a preference for looking at silicon solutions while keeping the current cable definition.
4. **The next step is ISI dual-edge test silicon to verify margin analysis**

5. **Jitter margin with nominal driver and skew control expected to be 2-6 ns**
6. **Hot-plugging is problematical – more test data is needed**

To this list, Bill Ham proposed the addition of a seventh summary point involving duty cycle distortion (DCD) issues. Bill Gintz argued that the DCD issues are not as serious as Bill Ham was suggesting. After a brief discussion, it appeared that Bill Gintz accepted DCD issues as potentially important.

Other topics discussed included:

1. **Skew compensation or CRC for data signals.** Bill Gintz and Vince Bastiani extolled the virtues of HIPPI electronic skew compensation. Then, Bill Gintz noted that CRC could solve the problems too. Bill Ham noted that skew compensation fixes the problem, but CRC only detects it. The consensus appeared to be that electronic skew compensation would be preferable, but it will most likely not be needed until we attempt much higher data rates than Fast-80.
2. **Case 4 hot plugging.** Vince asked what had happened to the proposal to provide staged contact resistance to a hot plugging connector. Bill Ham said AMP was still investigating the design. The group also discussed whether LVD is robust enough to do case 4 hot plugging without special connectors; whether today's parity and REQ/ACK counting were sufficient, or whether CRC should be employed. (See 5.9 for more on this topic.)

4.3 Bias Reduction Proposal (97-214) [Bridgewater]

Wally Bridgewater asked that discussion of this item be deferred to the next meeting.

4.4 Dual Clocking Proposal (97-208) [Bastiani]

Vince Bastiani presented his discussion of dual-edge clocking as a review of time interval testing for 80 MHz (see 4.6). He asked that this topic or one very similar to it be included on the next agenda.

4.5 Project Proposal for SPI-3 (97-262) [Lamers]

The first topic for argument was the name and the group agreed to change "SCSI-4 Parallel Interconnect - 3" to "SCSI Parallel Interconnect - 3". The group continued wordsmithing the proposal at length.

Dan Smith proposed creation of a group to create and review the electrical design data needed to reach the kind of transfer rates proposed by the group. The needs for and procedures of such a group were discussed.

In the absence of any objections, the group unanimously recommended that the T10 plenary approve the revised SPI-3 project proposal (97-262r2).

4.6 Fast 80 MHz SCSI Testing (97-276) [Bastiani]

Vince Bastiani presented measured data analyzing setup and hold timings (see 97-276) based on full 80 MHz signals (as opposed to dual edge clocking). A much-discussed component of the data measurement equipment was a HP 3310A Timing Interval Analyzer. Bill Ham noted a very large jitter in the base-line case. After some discussion, Gene Milligan led the group to the conclusion that the jitter could have been induced by the pseudo-random data that was being driven on the other data lines.

After a lengthy review of the data and discussion, Vince concluded that there would be sufficient margin to detect bits if dual-edge clocking is used.

5. SPI-2 Working Group {Monday Afternoon}

5.1 SPI-2 specifications on interconnect (97-235) [Ham]

Bill's presentation was deferred to item 5.6, since George Penokie has already incorporated the changes in SPI-2 revision 16.

5.2 0.5 pF balance capacitance for LVD SCSI (Table 29 SPI-2R13) (97-236) [Aloisi/Kapraun]

After a brief introduction from Paul Aloisi, Greg Kapraun presented new data (97-277) with capacitance imbalances introduced in each of the thirteen drives in the test configuration. The group was reminded that Greg's previous data (where only one drive was modified to have a capacitance imbalance) could be found in 97-236.

Greg's conclusion was that the 0.5 pF requirement should not be relaxed. However, questions were raised about the degree to which the data presented supported the conclusion. Greg said that he will conduct more detailed tests, including measuring the actual capacitance after the altering hardware has been installed.

Wally Bridgewater asked at what frequency the capacitance tests were to be performed. Paul Aloisi said that the testing method was specified in one of the source documents for SPI-2. After reading SPI-2, Wally suggested that a change in the frequency might be in order, especially for Fast-xx.

5.3 Capacitance Measurement at 2 Volts [Wallace]

Dean Wallace asked that this item be dropped from this and future agendas.

5.4 Universal Backplane Annex (97-121r0) [Wallace]

Dean Wallace asked that this topic be deferred to the next working group meeting.

5.5 Latching and Counting (97-216) [Guss]

Dave Guss was not present. Bill Ham expressed concern that the latching and counting topic doesn't have a home. There was a brief argument between SPI-2 and SPI-3 as potential homes. The consensus appeared to be that the material should be integrated in SPI-3.

Additional discussion, based on new document revisions if any, was deferred to the next working group meeting.

5.6 SPI-2 Revision 16 Review [Penokie]

George Penokie reviewed the SPI-2 changes between revisions 15 and 16. George noted several technical changes, for which he will request plenary votes to approve making the changes in SPI-2 revision 17. In the absence of any objections, the working group recommended that the plenary approve the technical changes made during its review.

5.7 SPI-2 Timing Tolerance (97-268) [Asami]

After a lengthy discussion of the tolerance issues raised by Tak Asami, the group almost reached an agreement that the Receive Period Tolerance row be removed from table 42 in SPI-2 revision 16. But, that agreement was short-lived when the group discussed additional statements that should be in the standard. Another lengthy discussion followed. Much of the debate centered on a concern that relaxing the specifications on the time interval significantly would impinge on setup and hold times. Also, the transfer period tolerance table entries were added in response to a problem that perhaps is better solved with the other clarifications that were made to the SDTR negotiation.

Bruce Leshay noted that the transfer-to-transfer jitter does not scale linearly with transfer rate; 0,25% may be okay for Fast-10, but it is certainly too small for Fast-40.

Eventually, the group agreed to recommend to the plenary that:

- 1) every entry in the Transmit Period Tolerance row in Table 42 of SPI-2 revision 17 be changed to 1 ns; and
- 2) every entry in the Receive Period Tolerance row in Table 42 of SPI-2 revision 17 be changed to 1.1 ns.

5.8 Load Compensation (97-281) [Novak]

Vit Novak expressed concern that his presentation concerned single-ended SCSI. He wondered whether the group would want to spend time on this topic, in view of the late hour. George Penokie asked that topics not intended for SPI-2 should be deferred. Vit presented an executive summary of his work.

Basically, Vit's presentation contained a review of experiments he performed to compensate for device capacitance presented to the bus by inserting an inductor and resistor on both sides of each signal for each device. The circuit reduces the disturbance to the bus signals caused by the device.

John Lohmeyer stated that this topic would be placed under SPI-3 in agenda for the next meeting.

5.9 Proposal to remove clause 10.5 from SPI-2 Rev 16 [Lohmeyer]

John Lohmeyer proposed removing clause 10.5 (Hot Plugging) from SPI-2. Gene Milligan asked what really goes away if clause 10.5 is removed. George Penokie said that the restriction against case 4 hot plugging goes away for LVD and SE/LVD multi-mode. John said that prohibiting case 4 hot plugging for LVD and multi-mode is overly restrictive. All testing and analysis done to date has shown that case 4 hot plugging with real drivers and receivers does indeed work. While this does not guarantee that all LVD implementations will work, the standard should not prohibit those that do work.

After a little research, Gene Milligan reported that removing clause 10.5 has no effect, since equivalent words are present in clause 10.4. At this point, John proposed amending clause 10.4. The group proceeded to discuss the quality of the hardware implementation that is necessary to produce a system capable of case 4 hot plugging.

Gerry Houlder raised concerns about drive vendors incorrectly being blamed for not conforming to the "hot plugging rule". John Lohmeyer responded that the market demands are forcing case 4 hot plugging with LVD products and customers will go to non-SCSI products if the demands are not met.

The group then considered the case 4 hot plugging conditions that might produce undetected errors. Bill Ham described problems with latching and counting that can result in undetected errors (data corruption).

Gene Milligan moved that clause 10.5 be removed; that the second sentence from bullet a) be removed from clause 10.4; and that a warning regarding undetected errors be added to clause 10.4. Before a second was made, the group debated the "undetected" wording in Gene's motion. John Lohmeyer seconded the motion. The real focus of the debate that followed was the exact wording of the note to be added to 10.4.

After four attempts, the group agreed on the following wording: "LVD may require more stringent system design to tolerate transients that occur during insertion and removal events."

The motion passed 13:2.

6. Meeting Schedule

The next meeting of SPI-2 and SPI-3 Working Group associated with a T10 meeting week will be March 16, 1998, in San Diego, CA at the Hyatt Islandia Hotel (619-224-1234), hosted by QLogic. Several other meetings were discussed, but the T10 plenary later authorized the following meetings:

SPI-2 Editors Meeting	Dec 3, 1997	Irvine, CA / QLogic
SPI-3 Protocol Meeting	Dec 4, 1997	Irvine, CA / QLogic
SPI-3 Physical Meeting	Jan 21, 1998	Irvine, CA / QLogic
SPI-2 LB Comments Resolution	Jan 22, 1998	Irvine, CA / QLogic

(STA has also planned SCSI Harbor meetings for the afternoon of Dec 2nd and Jan 20th.)

7. Adjournment

The meeting was adjourned at 6:47 p.m. on Monday November 3, 1997.