Accredited Standards Committee*
X3, Information Technology

Doc. No.: X3T10/96-157r0
Date: May 14, 1996
Project: 1142-D
Ref. Doc.:
Reply to: John Lohmeyer

To: Membership of X3T10

From: Bill Ham, SPI-2 Technical Editor
Larry Lamers, Vice Chair X3T10
John Lohmeyer, Chair X3T10

Subject: Minutes of SPI-2 Working Group
April 22-23, 1996 -- Milpitas, CA

Agenda

1. Opening Remarks

2. Approval of Agenda

3. Attendance and Membership

4. Test Data (96-141) [Ham]

5. LVD SCSI Driver Specification Presentation (96-145r0) [Moore]

6. Standing Waves [Bridgewater]

7. Glitch Issues
   7.1 Releasing Bus from Active Negation [Uber]
   7.2 Digital Filtering of Negation Release Glitches (96-158r0) [Lohmeyer]
   7.3 LVD Release Glitches Proposal (96-160) [Kakirian]
   7.4 Glitch Filtering [Ham]

8. Transmission Line Model (96-123 & -124) [Gingerich]

9. Timing Budget [Ham]

10. Hot Plugging [Ham]

11. SPI-2 Document Review (X3T10/1142D) [Ham]

12. Summary of Meeting Results

13. Meeting Schedule

14. Adjournment

Results of Meeting

1. Opening Remarks

John Lohmeyer, the X3T10 Chair, called the meeting to order at 9:00 a.m., Monday April 22, 1996. 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 following additions:

Asymmetrical driver specifications (actually covered under agenda item 5)
Hot Plugging

3. Attendance and Membership

Attendance at working group meetings does not count toward minimum attendance requirements for X3T10 membership. Working group meetings are open to any person or organization directly and materially affected by X3T10's scope of work. The following people attended the meeting:

Name	S	Organization	Electronic Mail Address
Mr. Norm Harris	P	Adaptec, Inc.	nharris@eng.adaptec.com
Mr. Lawrence J. Lamers	A	Adaptec, Inc.	ljlamers@aol.com
Mr. Tak Asami	A#	Adaptec, Inc.	asami@itc.adaptec.com
Mr. Wally Bridgewater	V	Adaptec, Inc.	wally@eng.adaptec.com
Mr. Richard Moore	V	Adaptec, Inc.	richard_moore@corp.adaptec.com
Mr. Tom Schneider	V	Adaptec, Inc.	schneid@itc.adaptec.com
Mr. Chris Burns	V	Adaptec, Inc.	chrisb@eng.adaptec.com
Mr. Shahe Krakirian	V	Adaptec, Inc.	shahek@eng.adaptec.com
Mr. Louis Grantham	P	Dallas Semiconductor	grantham@dalsemi.com
Mr. Siegfried Schmalz	V	Dallas Semiconductor	schmalz@dalsemi.com
Dr. William Ham	A#	Digital Equipment Corp.	ham@subsys.enet.dec.com
Mr. Richard Greenberg	V	IBM Corp.	richg@vnet.ibm.com
Mr. Dean Wallace	P	Linfinity Micro	75671.3443@compuserve.com
Mr. Wayne E. Werner	O	Lucent Technologies	wew@aluxpo.lucent.com
Mr. Ting Li Chan	A	QLogic Corp.	t_chan@qlc.com
Mr. Richard Uber	V	Quantum Corp.	duber@tdh.qntm.com
Mr. Gene Milligan	P	Seagate Technology	Gene_Milligan@notes.seagate.com
Mr. James Whitworth	A#	Seagate Technology	james.whitworth@conner.com
Mr. Brian N. Davis	A#	Seagate Technology	brian_davis@notes.seagate.com
Mr. William C. Gintz	V	Seagate Technology	bill.gintz@conner.com
Mr. Dave Guss	P	Silicon Systems, Inc.	dave.guss@tus.ssi1.com
Mr. Vit Novak	A	Sun Microsystems, Inc.	vit.novak@sun.com
Mr. John Lohmeyer	P	Symbios Logic Inc.	john.lohmeyer@symbios.com
Mr. Frank Gasparik	V	Symbios Logic Inc.	frank.gasparik@symbios.com
Mr. Pete Tobias	A	Tandem Computers	tobias_pete@tandem.com

Mr. Matt Kaltenbach	V	The Panda Project	mattk1@msn.com
Mr. Paul D. Aloisi	P	Unitrode Integrated Circuits	Aloisi@uicc.com

27 People Present

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

4. Test Data (96-141) [Ham]

Bill Ham presented data that indicated a 15 meter LVD segment with 15 20pf/20pf/10pf loads on 8-inch centers is possible. A possible alternative is 10pf/10pf/5pf loads on 4-inch centers. He also proposed a set of Fast-40 configuration rules that suggested 25 meters if the loads are spaced at least 1 meter apart and 35 meters point to point. The importance of using asymmetrical drivers that compensate for the biasing effects of the terminators was clearly shown. For this reason an agenda item was added to see if the degree of driver asymmetry could be tightened. See also the following item.

A discussion of ways that might be used to decrease the intersymbol interference suggested that there may be a significant difference between the behavior of current mode drivers vs voltage mode drivers - Paul Aloisi of Unitrode noted that a lower impedance terminator could reduce the height of the long pulses. This would give a better match to the loaded cable impedance ~80ohms vs 110ohms for the present terminator. Wally Bridgewater of Adaptec floated the idea of having the driver reduce the amplitude of long pulses after the initial bit. This would reduce the starting point for the isolated bit transition and result in a larger detectable signal. There were a couple of points that were not explored: (1) how does the driver know how large the long pulse is due to manufacturing tolerances and (2) what sort of algorithm should be used?

Bill's summary: The real lever is the transmission line properties and appropriate driver asymmetry. Rules: 8-inch spacing for cables @20pf/ft; 4-inch spacing for PCB @40pf/ft; devices with 20pf/20pf/10pf loading and drivers that compensate for the terminators biasing. One cannot compensate for poor transmission line properties by shortening the length of the bus segment because the reflections are actually worse in short lengths due to the lack of edge roll off. These capacitance value for the cable capacitance are very approximate -- Bill will do more careful measurement of the details of this property. The rules above are based on the cables used in the testing.

5. LVD SCSI Driver Specification Presentation (96-145r0) [Moore]

Richard Moore presented an analysis of Kevin Gingerich's earlier model. Richard's main point was that the requirements need to be adjusted to more nearly represent the situation for real drivers that need to accommodate manufacturing tolerances and other features. Details were given in his slides.

Bill noted that his direction was what he favored and that even further consideration needed to be given to the receiver overdrive requirements and the tolerances on the asymmetry. The presentation showed a serious commitment on the part of Adaptec to support LVD standards development.

It was shown that part of the problem with some of the driver tolerances were due to the tolerances on the terminators. This caused a revisiting of the terminator specifications and resulted in an agreement to adjust the bias range to 100-125 mV. Also the resistive slope range was narrowed to 100-110 ohms. This gave some much needed relief to the silicon side.

The question was raised concerning whether the common mode may need to be adjusted to 100-150 ohms. This point was not resolved and the standard will remain as is for common mode.

Bill noted that the asymmetry should be specified more tightly than the gray area in Richard's revised region of operation. Asymmetry exists even if asymmetrical drivers are not used. Bill noted that the asymmetry observed in the lab is caused mostly by intersymbol interference.

Richard promised to revise a proposal for the upcoming meeting in Florida.

6. Standing Waves [Bridgewater]

Wally Bridgewater stated that he plans to have test results at the meeting in June. To be continued. Bill noted that he hasn't observed anything in the lab that indicates a problem, but the possibility exists.

7. Glitch Issues

7.1 Releasing Bus from Active Negation [Uber]

No activity, remove from future agendas

7.2 Digital Filtering of Negation Release Glitches (96-158r0) [Lohmeyer]

John Lohmeyer presented a note on where release glitches might occur. He promised to have a proposal that covers all these areas for the May meeting.

7.3 LVD Release Glitches Proposal (96-160) [Kakirian]

Shahe Krakirian of Adaptec made a presentation on LVD release glitches. The presentation initially proposed a new parameter, Bus Round Trip Delay, of 150 ns. This was viewed as unacceptable because it would have the effect of forcing the total SCSI domain lengths to approximately one third of the presently allowed 400 ns. On the other hand, one does need to accommodate release glitches by ensuring that they are not acted on by real devices.

A new scheme was devised that fully restored the 400 ns and protects against these glitches. The problem case is REQ after a data phase. Details of this proposal are in the revised presentation. A complete map of the requirements needed to accommodate these glitches was created. The consensus was to incorporate the revised map from the proposal into SPI-2. The timing diagram that clarifies this will not go into SPI-2 but may be useful for the ENDL SCSI Bench Reference or other similar documents.

We need to add something for selection timeout for data and parity. This is John Lohmeyer's action item.

7.4 Glitch Filtering [Ham]

This topic related to the use of a data, parity, REQ, ACK filter that would force the receiver to "stop listening" for a period of time after a signal transition had been recognized. Bill Ham showed data that illustrated a glitch in the REQ signal that could cause false detection of REQ in illegal configurations produced by excessive loading. He recommended a 6 ns glitch filter exactly like that required in SPI for negation transitions. The silicon folks (Tak Asami of Adaptec) pointed out that the 10 ns filter required in SPI was not too difficult to implement because the upper bound on the filter could be much larger than the 10 ns minimum. In the Fast-40 case, where the assertion/ negation period may be as short as 6.5 ns, the glitch filter would have to be much more precise. In addition, the situation becomes even worse for Fast-80. The consensus was not to require a glitch filter but to point out that it could offer valuable protection in illegal configurations. Implementors may wish to incorporate such a filter if possible.

8. Transmission Line Model (96-123 & -124) [Gingerich]

This item was not addressed because of the absence of Kevin from this meeting.

9. Timing Budget [Ham]

The timing budget was reviewed in light of the changes in the measurement method for timing caused by the receiver overdrive requirements. Rev 6 of SPI-2 contains these revised requirements. Time must be allowed for the signal to go from the top to the bottom of the receiver switching range. In addition, the assertion and negation times for Fast-80 needed to be included.

As the transition time is highly dependent on the actual slope of the differential waveform and this slope is determined mainly by the properties of the transmission environment this additional time was added to the pulse distortion skew to give a total of 3 ns. This effectively pushed a tighter requirement onto the drivers.

At the urging of Frank Gasparik of Symbios Logic the receiver sensitivity for use in the timing measurements was increased to 60 mV. This was needed because the present 30 mV level is too close to the internal receiver offsets and will cause too much internal skew in the receivers. On the other hand we need to keep the 30 mV level for the d. c. performance receiver requirements to preserve the cross talk margin for quiescent signals. In light of the receiver overdrive requirements this change is almost not noticeable because the 0.25*Vtan dominates the timing measurements.

There was some debate around the exclusion of single-ended Fast-40. Bill stated that the argument for hot-plugging might legitimize revisiting this topic (see related topic below). Bill proposed removing the SE/DIFF designation from Fig 14 in the last two rows. The assumption is that all devices will have on-chip transceivers.

It was attempted to create a budget that could be used for separate transceivers. The first attempt failed. However, Richard Moore lead a process that actually produced a possibly workable situation where the separate drivers could have 1.5 ns skew budget and the separate receivers could have 2.0 ns skew budget. This latter condition will be incorporated into the next rev of SPI-2.

The following limits were agreed on the assertion and negation periods: Fast-40 8 ns at the driver, 6.5 ns at the receiver; Fast-80 3 ns at the driver, 2.5 ns at the receiver. Fast-80 will have the same driver slew rate requirements as Fast-40.

The leakage current of existing SCSI is 10 uA, 20 uA would make the universal driver circuit easier. We will change this in the next rev.

10. Hot Plugging [Ham]

Bill stated that his analysis indicates that LVD may not work properly with hot plugging, especially the case 4 scenario. The use of LRC and intelligent use of REQ/ACK offset counters could mitigate the concerns. Another suggestion is to use a signal to quiesce the SCSI bus segment to allow an assertion/removal. Another is to use repeaters or single-end Fast-40 with expanders. A list of options was generated for consideration at the next meeting.

The message to marketing people is that case 3 is the general-accepted approach to take. Case 4 can be done with the addition of isolation interfaces (add some added cost).

11. SPI-2 Document Review (X3T10/1142D) [Ham]

The tactical issue of what is in SPI-2 is an item for the plenary. The working group deferred this discussion to the Florida meeting.

12. Summary of Meeting Results

13. Meeting Schedule

The next meeting of SPI-2 Working Group will be May 6, 1996 in Ft. Lauderdale, FL. Another SPI-2 Working Group is scheduled for Thursday-Friday June 6-7, 1996, in Colorado Springs, CO at the Embassy Suites Hotel (719-599-9100) hosted by Symbios Logic.

14. Adjournment

The meeting was adjourned at 4:00 p.m. on Tuesday April 23, 1996.