Accredited Standards Committee*
X3, Information Technology

Doc. No.: X3T10/96-122r0
Date: February 4, 1996
Project: 1142-D
Ref. Doc.:
Reply to: John Lohmeyer

To: Membership of X3T10

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

Subject: Minutes of SPI-2 LVD SCSI Working Group Meeting
Denver, CO -- January 29-30, 1996

Agenda

1. Opening Remarks

2. Approval of Agenda

3. Attendance and Membership

4. Releasing Bus from Active Negation [Uber]
   4.1 Digital Filtering of Negation Release Glitches [Lohmeyer]
   4.2 Symmetric vs. Asymmetric Drivers
   4.3 Timing Budget [Jander]

5. Three-Resistor Terminator Equations [Haynes]
   5.1 Additional Termination Issue (96-125) [Aloisi]

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

7. Voltage Mode Drivers [Bridgewater]

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

9. Test Data (96-127) [Ham]

10. Standing Waves [Bridgewater]

11. Meeting Schedule

12. Adjournment

Results of Meeting

1. Opening Remarks

John Lohmeyer, the X3T10 Chair, called the meeting to order at 9:00 a.m., Monday January 29, 1996. He thanked the host, Symbios Logic, perhaps too profusely.

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

2. Approval of Agenda

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. Wally Bridgewater	V	Adaptec, Inc.	wally@eng.adaptec.com
Mr. Richard Moore	V	Adaptec, Inc.	richard_moore@corp.adaptec.com
Mr. Dennis R. Haynes	O	Burr-Brown Corp.	haynes_dennis@bbrown.com
Mr. Justin McEldowney	V	Burr-Brown Corp.	mceldowney_justin@bbrown.com
Mr. Louis Grantham	P	Dallas Semiconductor	grantham@dalsemi.com
Dr. William Ham	A#	Digital Equipment Corp.	ham@subsys.enet.dec.com
Mr. Dan Colegrove	A#	IBM Corp.	colegrove@vnet.ibm.com
Mr. Richard Greenberg	V	IBM Corp.	richg@vnet.ibm.com
Mr. Dean Wallace	P	Linfinity Micro	75671.3443@compuserve.com
Mr. Ting Li Chan	A	QLogic Corp.	t_chan@qlc.com
Mr. Richard Uber	V	Quantum Corp.	duber@tdh.qntm.com
Mr. John F. Fobel	O	Rancho Technology, Inc.	scsi@tstonramp.com
Mr. Gene Milligan	P	Seagate Technology	Gene_Milligan@notes.seagate.com
Mr. Dave Guss	A	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. Kevin Bruno	V	Symbios Logic Inc.	kevin.bruno@symbios.com
Mr. Kelly Bryant	V	Symbios Logic Inc.	kelly.bryant@symbios.com
Mr. Brian Day	V	Symbios Logic Inc.	brian.day@symbios.com
Mr. Frank Gasparik	V	Symbios Logic Inc.	frank.gasparik@symbios.com
Mr. Mark Jander	V	Symbios Logic Inc.	mark.jander@symbios.com
Mr. Richard Mourn	V	Symbios Logic Inc.	richard.mourn@symbios.com
Mr. Tracy Spitler	V	Symbios Logic Inc.	tracy.spitler@symbios.com
Mr. Kevin Gingerich	V	Texas Instruments	4307725@mcimail.com
Mr. Paul D. Aloisi	P	Unitrode Integrated Circuits	Aloisi@uicc.com
Mr. Tak Asami	A	Western Digital Corporation	asami@dt.wdc.com

27 People Present

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

4. Releasing Bus from Active Negation [Uber]

Dick Uber presented some simulation data (96-126) regarding the release glitches (active negation to high impedance).

Dick also included data on glitches that can occur on an assertion to high impedance transition. There glitches are different from the “wired-OR” glitches already documented in SPI because they can occur when only one device is driving the bus. Bill Ham noted that SPI-2 already requires all non-wired-OR signals to be negated for at least a Bus Settle Delay before being released and therefore the point becomes moot.

Dick also had some data that indicated attenuation puts a limit on how asymmetrical drivers can be.

He had built a cable attenuation model for the REQ/ACK pulses that run at 40 MHz. Simple skin effects and relavant harmonics up to 440 MHz were considered. With a worst-case cable and symmetric drivers, 32 mA drivers would be needed. With worst-case cable and asymmetric drivers, an 8/4 mA driver would be required. Dick’s ‘worst-case’ cable consisted of 30 Gage single-strand wire, 110 ohms minimum line impedance, and 12 meters long. The skin effect is stronger where there is less "skin". The use of non-stranded wire provides the minimum surface area for a given copper cross section and makes stranded more appealing for high frequency uses. It will be important not to restrict the lengths of cables arbrtrarily

4.1 Digital Filtering of Negation Release Glitches [Lohmeyer]

John Lohmeyer spoke about the message he had sent to the SCSI reflector identifying where thenegation release glitches would occur. He said that he believes the negation release glitches can be filtered digitally should the group agree on an analog LVD SCSI solution that produces these glitches. SPI-2 wording will need careful review to minimize the places where negation release glitches would be permitted and to clearly identify where digital filtering would be required.

Later in the meeting, it was concluded that negation release glitches may still occur. It will be necessary to specify release glitch filtering in SPI-2. It is also important to consider if the present requirement for actively negating all assertions for a bus settle delay has any protocol impact.

4.2 Symmetric vs. Asymmetric Drivers

Earlier in the month, the SPI-2 LDV SCSI working group started investigating using asymmetric drivers (weaker negation driver than assertion driver). This was in part because some people believed this would help eliminate the negation release glitches (false assertions when going from active negation to released). The SCSI working group had recommended that analog people not go to asymmetric drivers just to remove these glitches -- they believed that the glitches can be filtered digitally. Most of the meeting danced around whether there are good analog reasons to switch to asymmetric drivers. The most important reason that evolved was that considerable power can be saved in the negation driver. This also would deliver a more symmetric signal to the receiver, which would help with setup and hold times.

The eventual solution permitted either symmetric or asymmetric driver implementations. However, asymmetric drivers should consume less power.

4.3 Timing Budget [Jander]

Mark Jander asked that the setup and hold time budgets be re-visited based on the changes to the driver specifications. Bill Ham said that such discussions really need a setup and hold time measurement figure like the one in Annex B of Fast-20 in order to be useful.

Bill drew part of a figure that represented the receiver setup and hold times. He started to draw a similar figure for the driver and ran into some discussion from folks who thought the driver should have different critical voltage levels from the receiver. This indicated that some folks had not carefully considered how the Fast-20 timing diagrams were derived. In the Fast-20 case timing allowance was made for signal degradation during its

journey to the receiver. For example the pulse widths at the driver are required to be significantly longer than at the receiver. Actually this timing difference was first introduced in the SCSI-3 SPI document.

Bill reasserted that this model should be continued for LVD but that the voltage levels would need to be adjusted for the LVD receiver case. There is no voltage level at the driver that can be used as a reference other than the receiver switching levels Bill claimed.

In a related discussion, Gene Miligan querried whether there was a requirement that the signal pattern used during timing measurements be regular (i.e. alternating assertions and negations of equal length). Bill Ham noted that there is presently no requirement on the data pattern for SPI or Fast-20 and that this hole that should be plugged in SPI-2. Bill suggested a model where a regular pattern is used for the timing measurements for the driver and where the receiver timings would include the cable plant and intersymbol interference timing budget as well as that needed by the receiver silicon. (The intersymbol interference arises when non-regular pattern are used.) There may need to be some requirements on the driver concerning how much deviation from "regularity" is allowed.

Bill took the action item to proposed a more detailed timing measurement scheme for LVD for the March meeting.

5. Three-Resistor Terminator Equations [Haynes]

Dennis Haynes presented a foil describing the derivation for his SCSI reflector message describing a three-resistor terminator. The conclusion was that using a single voltage applied directly to one side of a three-resistor totem pole stack terminator (with the other side connected to local ground) will not meet all of the LDV SCSI requirements. Such terminators will need to have two non-ground voltage levels, one applied to each side of the resistor stack, to meet all of the specification requirements. This does not imply a complication for the distribution of TERMPWR for LVD terminators but it does imply that LVD terminators will need two voltage levels that are derived from the TERMPWR line (or a local source) and that "active" or "regulated" LVD terminators will generally be needed.

5.1 Additional Termination Issue (96-125) [Aloisi]

Paul presented colorful slides (96-125) describing a potential problem if a terminator were designed with a fairly high common-mode resistor of 500 ohms and if the driver is capable of sourcing 24 mA (which is permitted in the EIA/TIA standard). This could result in a 12 v pulse when two drivers are turned off at the same time, which would likely be harmful to the drivers and receivers. Paul proposed that the common-mode resistor range be specified to be in the 100 - 150 ohm range instead of today’s 100 - 500 ohm range.

Wally Bridgewater was concerned about some of the terminators shown during the discussion. They appeared to be quite different loads as seen by the drivers.

Kevin Gingerich noted that having two drivers on at the same time would have the effect of adding another driver to his model. The impedance seen would depend on the driver type. The voltage spike problem is worse with voltage-mode drivers because they have a much lower impedance. It may be necessary to specify a maximum current requirement to limit this problem. The other factor that affects this problem is the amount of common mode voltage on the cable.

After considerable more discussion, the specification on the common-mode resistance was adjusted to be 100 - 300 ohms.

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

Kevin Gingerich presented a model for LVD SCSI (see 96-123). It includes both AC and DC models and assumes a lossless cable. Kevin used the solver feature in Excel to optimize several parameters. From the results, Kevin prepared 96-124, which proposes how SPI-2 should be modified to use his results in specifying the driver and receiver characteristics. This document was reviewed under agenda item 8.

7. Voltage Mode Drivers [Bridgewater]

Wally asked that this topic be combined with the Standing Wave topic (see agenda item 10).

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

Bill Ham lead a discussion of the terminator specifications in Rev 3 of SPI-2. Figures 2 and 3 were revised to use a voltage source instead of a current source. V3 was specified at 1.0 volts and V4 was specified at -1.0 volts. Note that higher voltages may be seen under wired-OR conditions, but these should not matter as the protocol permits a bus settle delay for the wired-OR signals to settle.

Kevin Gingerich questioned why the common-mode test circuit for terminators in Figure 6 specified 0 to 1 MHz instead of a driver-speed pulse, which would have much higher frequency components. The 1MHz number was used from the last meeting as a reasonable guess for the frequency components of externally generated noise. No decision was made at this point and the frequency specification for the terminators was left To Be Defined after we settled on the driver specifications. At the end of the meeting we agreed to use the maximum frequency of 120 Mhz for all the terminator tests in the next draft document. This frequency was chosen as the third harmonic of the fundamental for the Fast 40 REQ/ACK signals. This was viewed as reasonable since the drivers themselves contribute common mode noise due their imperfect balance.

On Tuesday morning, Bill lead a discussion of the driver and receiver sections of SPI-2. Figure 15 was revised to use different terminology. Kevin Gingerich lead a detailed review of his 96-124 document that contains suggested changes to the driver and receiver sections based on the data he developed in 96-123, LVD SCSI Model. Bill will edit all the figures, tables, and text to conform to the new terminology.

There was considerable debate over whether Kevin's model included enough margin for attenuation. Kevin included 70 mV which he claimed was 2 dB. Dick Uber wanted at least 3 dB and he asserted that this should be measured over the entire signal swing instead of the 0 to peak range. After reviewing Bill Ham's lab data (see item 9) and considering that the attenuation does really apply to the entire wavefront it was agreed to drop the use of db's and to increase the budget for attenuation to 130 mV for subsequent work.

Kevin proposed a complete set of tests and test limits based on the assumptions of the last meeting.

A great deal of time was spent reviewing the remainder of Kevin's document. We eventually agreed to wait on specifying the minimum rise and fall times -- these will be discussed on the SCSI reflector. The basic issue is whether the test circuits are an adequate model for designing drivers or whether additional reactive components are needed in the test circuit to better simulate a real cable. The use of a real cable was seen as problematic due the the need to test the extreme ranges of the characterictic impadances and it is not easy to produce precise enough cables so that different test sites would get similar results.

The common-mode voltage specifications need to be revised to 0.7 to 1.8 volts (1.25 +/- 0.55) for the silicon tests. Accomplishing this may require reducing the allowed ground shiflt. It was agreed that the test circuits were OK but that we needed the new numbers for the test conditions and allowed test result limits. Kevin agreed to incorporate all these new agreements into a new version of his model and to send these to Bill for inclusion in the next draft of SPI-2.

For the first time we have a model that is complete enough to use for making tradeoffs.

Bill called for an Annex to be created that contains the details of the model. This was generally agreed as useful but it was not clear how we could document the Excel program. Gene Miligan suggest a "C" code listing as one possibility.

The specification of LVD SCSI bus loading was reviewed with the result that the sections in Rev 3 are OK.

9. Test Data (96-127) [Ham]

This item was covered early on Monday. Bill Ham presented 96-127, which contains the results of his testing of the TI LVDS drivers (for the EIA/TIA standard) with biasing terminators. Bill's previous data had not used biasing terminators and was therefore very optimistic. While the drivers do not exactly match the current SPI-2 LVD SCSI specifications, they are close enough to yield very useful data. The data show that bus signals resulting from symmetric drivers used with biasing terminators have little assertion margin (due to fighting the termination bias) and great (excessive?) negation margin. These drivers were calculated to have approximately 6 mA symmetrical current for both assertion and negation and therefore are a very close match to the earlier LVD SCSI driver specification.

With symmeteric drivers Bill's data suggests that 27 meters of heavily loaded bus is possible for Fast 40. He also showed that Fast 80 could not reach these lengths unless steps were taken to improve the regularity of the REQ and ACK signals. The attenuation levels reached were those using multistranded primary copper conductors from the present SCSI 30 guage P cables. With solid conductors one would expect to experience higher attenuation. This whole area needs more work to arrive at the tradeoffs.

A major effect on the bias delivered from the terminators was noted due to driver loading. For the drivers used a single driver reduced the bias from 112 mV to 82 mV. This would clearly be bad news if many drivers were present on the bus at the same time (in a HIZ state). Kevin called back to his office and found that these drivers actually have nearly 0.5 mA "leakage" caused by an internal circuit intended to pull the bus back to its original condition when the driver is not driving. As these drivers were developed for use only in point to point applications this was not a serious issue. Actual LVD SCSI drivers would have less than 20 uA leakage and 15 would therefore cause 0.3 mA total for all drivers. This is 60% of that for the single LVDS driver used for these tests. One can expect no more than 18 mV reduction in terminator bias from a maximally leaky LVD system. The test data exactly matches this case for a terminator delivering an unloaded bias of 100 mV (the present lower limit). This data confirms earlier projections that leakage would have some effect. Fortunately, this level of leakage is easily accommodated.

This means that LVD SCSI is set up for using the newest silicon technologies that are expected to have higher input leakage.

Bill’s conclusions were that asymmetric drivers would significantly benefit noise margins without causing any increase in power consumption. They are also a key for operating at longer lengths.

Bill passed around a new test board developed at Digital that will be used for the creation of precise LVD SCSI bus loading conditions. This test board will allow adjustment of the LVD loading over a broad range and will enable the testing of simultaneous lines. Bill may have data for the March meeting.

Bill also showed a nearly comprehensive set of data describing what happens during negation to HIZ and other HIZ related transitions. This data showed evidence of false assertions significantly after the HIZ transitions caused by reflections. The general conclusion was that software based glitch filtering should be required (for a bus settle delay).

For the next meeting Bill expects to have some data with asymmetric drivers.

10. Standing Waves [Bridgewater]

Wally Bridgewater actually had three items:

1. Can the REQ/ACK Offset be greater than 16? Yes, the REQ/ACK Offset can be up to 254.

2. Are the single-ended cables different than LVD cables because the impedances are specified differently? No, the impedances are measured differently, so the cables can be the same.

3. He has noticed that single-ended cables exhibit standing waves under certain condtions. He plans to do some testing to see if differential cables also exhibit standing waves. Bill Ham will also review his earlier data where some evidence of anomolous frequency dependence was seen on the single ended signals but not on the differential signals. This could be a very important area of concern if it turns out that the differential signals are affected.

11. Meeting Schedule

The next meeting of SPI-2 Working Group will be Monday March 11, 1996, in San Diego, CA at the Hyatt Islandia (619) 224-1234, hosted by QLogic Corp. Another SPI-2 Working Group meeting was tentatively scheduled for April 22-23, 1996 in the San Jose area.

12. Adjournment

The meeting was adjourned at 4:25 p.m. on Tuesday January 30, 1996.