Accredited Standards Committee*
X3, Information Technology

Doc. No.: X3T10.1/96a101R0
Date: January 21, 1996
Project:
Ref. Doc.:
Reply to: Lawrence J. Lamers

To: Membership of X3T10.1

From: Bill Ham, Project Leader PH2
Lawrence J. Lamers, Chair X3T10.1 (acting)

Subject: Minutes of SSA 2nd Generation Working Group
January 16-17, 1996 : Winchester, UK

Agenda

1. Opening Remarks

2. Attendance and Membership, Introductions

3. Approval of Agenda

4. Document Distribution

5. PH1 Issues [Ham]
   5.1 Line segment termination values []
   5.2 DC voltage drop considerations []

6. Speed negotiation [Kapraun]

7. PH2 Architecture [Ham]
   7.1 SSA 40 Testing results
   7.2 Performance Requirements
   7.3 Termination/Coupling
   7.4 Jitter Transfer/Tolerance
   7.5 Measurement points
   7.6 Optical Specifcations
   7.7 Option Block Functions
   7.8 Device connector

8. Call for Patents

9. Action Items

10. Meeting Schedule

11. Adjournment

Opening Remarks

Larry Lamers convened the meeting at 9:00 am. He thanked Adge Hawes of IBM UK for hosting the meeting. As is customary, the people attending introduced themselves. A copy of the attendance list was circulated for attendance and corrections.

It was stated that the meeting had been authorized by X3T10.1 and would be conducted under the X3 rules. Ad hoc meetings take no final actions, but prepare recommendations for approval by the X3T10.1 task group. The voting rules for the meeting are those of the parent committee, X3T10. For the ad hoc, other than straw votes, the voting rules are: one vote per participating company

The minutes of this meeting will be posted to the X3T10 BBS and the SSA Reflector and will be included in the next X3T10.1 committee mailing.

Larry stated that the X3T10.1 mailings are now part of the X3T10 mailings. Persons that want to receive documents should subscribe to the X3T10 mailings by sending their request to the secretariat. An electronic option should be available during 1996.

Attendance and Membership, Introductions

Attendance at working group meetings does not count toward minimum attendance requirements for X3T10.1 membership. Working group meetings are open to any person or company to attend and to express their opinion on the subjects being discussed.

The following people attended the meeting.

name	company	telephone	email
Lawrence Lamers	Adaptec	408-957-7817	ljlamers@aol.com
Brent Hatfield	AMP	3531 806 6529	
Charles Brill	Amp Inc	717-592-6198	cebrill@amp.com
Lisa A. Huff	Amp Inc	717-986-3143	lisa.huff@amp.com
Bill Ham	Digital Equipment Corp.	508-841-2629	ham@subsys.enet.dec.com
Martin Chesters	IBM Corporation	44-1705-486363	mchesters@vnet.ibm.com
Matt Cordrey-Gale	IBM Corporation	44-1705-486363 x5669	matt@uk.ibm.com
Adge Hawes	IBM Corporation	011-44-1705-486363	adge@vnet.ibm.com
James S. Mason	IBM Corporation	44-1705-486363 x4226	jmason@vnet.ibm.com
Phil Murfet	IBM Corporation	01705 486363	
Chris Parker	IBM Corporation	912-892-2719	cparker@vnet.ibm.com
Vinay Shah	IBM Corporation	011-44-1705-486363	vvshah@vnet.ibm.com
Wolfgang Drichelt	ITT Cannon	49-7151 699 233	
Rob Damjanovic	Molex	49-89-413092 55	
Adrian Stokes	Molex	01420 477070	astokes@molex.com
Rick Born	Symbios Logic Inc.	970-223-5100 x9701	rick.born@symbios.com
Greg Kapraun	Symbios Logic Inc.	970-225-4843	greg.kapraun@symbios.com
Tim Beck	Trimm Technologies	44-115-951-9951	100612.3055@compuserve.com
Graham R. Mills	Trimm Technologies	44-115-951-9951	
Mark Buer	VLSI Technology Inc.	602-752-6451	Mark.Buer@tempe.vlsi.com
John Ciccone	VLSI Technology Inc.	602-752-6233	zack.cicone@tempe.vlsi.com
Stephen Nemetz	VLSI Technology Inc.	602-752-6240	steve.nemetz@tempe.vlsi.com
Sam Sanyal	VLSI Technology Inc.	408-922-5371	sanyal_s@sanjose.vlsi.com
Neil Edmunds	Xyratex	011-44-1705-486363	neiledmunds@uk.xyratex.com
David E. Instone	Xyratex	01705 486363	dinstone@uk.xyratex.com
John Veal	Xyratex	01705 486363	

Approval of Agenda

The agenda was approved.

Document Distribution

Larry Lamers stated that the next mailing deadline is Wednesday, January 16, 1996.
Documents are available on the SCSI BBS (714) 574-0424 and at ftp.symbios.com.

PH1 Issues [Ham]

41 Line segment termination values []

The objective is to acheive a consensus on what is needed for line segment termination. The terminology was precisely defined at the last meeting in San Jose.

Bill Ham gave a presentation illustrating the need for specifing termination values (see 96a102). The test results are based on the IBM test board and ten meters of cable. Capacitance as small as 1 pf and very short stubs did have significant effect on the eye pattern. The point is that the termination values need to be carefully specified as to inductive and reactive components, including parasitic capacitance.

Source termination has value for reducing reflections at the cost of increasing the driver current and introducing current paths other than through the device connector. This produces some concern for silicon suppliers who want the same silicon to serve multiple applications.

David Instone made a presentation on waveform test results. The trigger point can affect the measured eye; any common mode noise also affects this measurement. The tests were done with 25 m and 33 m cables. David's data was taken using a regularly repeating K28.5 pattern. As a result he always triggered on the same edge and did not see the effects of intersymbol interference on the jitter of the trigger. When other edges were used as the trigger points there was an increase in the jitter as expected and as is required in the present PH1 document.

He also made a presentation on TDR measurements of 5 devices from two manufacturers. The two single-ended signals taken as a differential using CSA 803 pulses. The dip is caused by the differential input capacitance of the silicon.

James Mason commented that the packaging and ESD parasitics are the main components. Chris Parker questioned why the filtering wasn't having more effect on the test results. It came about that Instone's data was taken with a 250 ps filtering instead of the presently required 750 ps. This produced a much greater "dip". Even accounting for the increased dip there was such a large excursion that the present specification could not be met.

Phil Murfet had some data that supported the results shown by David Instone. Phil suggested allowing a single excursion outside the specification. This had been suggested before but rejected on the basis that controlling the rise time of the TDR pulse covered the issue. Nevertheless, since it may not be possible to meet the present specifications with real silicon, the nearest match to the ideal was sought by seeking the least possible time excursion below the lower impedance limit. For the devices tested it appeared that we need to allow for a single excursion below 120 ohms differential for up to 1.5 ns.

The lowest level reached by the TDR trace during this excursion is not specified.

It is not known quantitatively how much signal distortion will result from this revised specification nor what the effect will be on faster edges. For the moment this the best specification that is possible for SSA 20.

Phil made a presentation to support the need for near-end, or source end, termination for 40 MB/sec. The alternative is higher error rates, shorter cables, more expense in devices. The near-end termination reduces the size of the source end reflection that reaches the receiver. The negative is that the current path may not be exactly balanced. Adding in capacitors to get rid of the DC path changes the common mode by -356 mV. An improvement in the common mode input definition needs to be addressed.

Bill stated that we need to ascertain the slowest slew rate that is workable for 40.

Chris Parker showed the effects of using capacitive termination at the source. Going to 90 ohm single ended impedance on the lines near the terminating resistor improves signal quality. Using a 225 ohm termination at the near end is also an improvement.

Adge Hawes stated that the metrics in PH1 should be adjusted so that PH2 devices do not violate PH1.

The need for source end termination directly arises from poor far end or receiver end termination. It is still an open question as to whether one needs to require source end termination for PH2.

Agreements:

.A line segment termination specification is required

.PH1 and PH2 will have the same line segment termination specifications

.The maximum driver slew rate will be a number for PH1 that allows maximum overlap for PH2 allowed range.

.Present method of specifing line segment termination in PH1 is ok for PH2

.Line segment termination test shall represent operating conditions

.The line segment termination test shall use 750 ps minimum rise time and produce a max differential impedance of 180 ohms and min differential impedance of 120 ohms as presently specified.

.A single excursion below 120 ohms for max of 1.5 ns is allowed for line segment termination when measured with a 750 ps rise time pulse.

42 DC voltage drop considerations []

There is an assumption in the existing working draft that there is no dc voltage drop in the cable because there is no DC data transmitted. The lowest frequency is a K28.5 character at 20 Mbits per second. Lisa Huff reported that 28 AWG is 0.065 ohms/foot; 24 AWG is 0.025 ohms/foot.

Equalization circuits reduce jitter but consume power causing further voltage drop. These circuits are not specifically budgeted for in the existing specifications. Design efforts are under way to reduce the voltage drop from what occurs in a typical Fibre Channel equalizer.

For the time being no special budget for DC voltage for equalizers or cable wire gauge will be developed. The requirements to meet the minimum receiver eye signal cover most of the issues without need for change.

James Mason stated that SSA 40 will go 20 meters. If longer distances are required a larger wire diameter in the cable; passive equalization may get you to 40 meters; an active equalizer could correct the attenuation; or an optical extender could be used.

Speed negotiation [Kapraun]

Greg Kapraun presented the speed negotiation proposal. This proposal allows devices to work at 20 even if they are capable of 40 in the event that the error rate is excessive or the cable medium is damaged. The proposal also allows for support of 16 speeds and devices with different supported speeds.

Adage - recap of issues with Ian's proposal.
- need to work at lower speed than max transmitter speed
- supporting non-harmonic speeds
- transmitter and receiver is much easier if at same frequency

There is a need to find a better "off-hook" for the modem model. Need to address the DC balance effects. Could a 10 Mhz signal be substituted for the 'no transition'. Yes,

The spec needs to be linearly responsive to frequency so that a 20 cable works at 40 and 20 and a 40 cable works at 20 and 40.

Change 'logic zero' to 'no transitions'. Bill Ham 'The outbound line is held stable and the receiver looks for no transitions.'

The times need to be re examined. (10 us, 1 us for 16 speeds ?). Change state diagram for balanced lines.

The best near-end termination for speed negotion appears to be the one using a capacitor and resistor at the driver. The time constant for transitions is guessed to be 10 us.

The idea of using a 10 Mhz signal instead of a DC signal seemed to be very attractive since it was consistent with the use of AC coupling and optical interconnect.

This whole issue will be revisited but appears to be accepted by the study group with the use of the 10 Mhz signaling. Dave I pointed out the need for the driver to be able to transmit such a signal and that a frequency tolerance would be needed.

PH2 Architecture [Ham]

61 SSA 40 Testing results

Matt Cordrey-Gale reviewed the testing results from IBM UK on SSA 40.

62 Performance Requirements

Bill H reiterated the agreements attained at the Ithaca meeting where SSA 40 and hopefully faster would use the same interconnect as SSA 20 (with a possible length reduction). As the length tends to be dominated by the attenuation and data dependent jitter one could reasonable expect to have approximately the same length for SSA 40 as SSA 20 if the receiver were to be twice as sensitive for SSA 40 as for SSA 20.

The driver eye for SSA 40 has 2 options: easy and aggressive (we will specify the easy version as required but encourage the agressive to increase noise margin). The same maximum edge rates will be used for SSA 40 as for SSA 20 in order to preserve the same line segment termination as SSA 20. The easy version preserves the general shape of the SSA 20 mask by halving the jitter spec, keeping the maximum slew rate the same, and retaining the peak amplitude. The agressive version allows the maximum slew rate all the way to the maximum amplitude.

The driver operating specifications for on state and off state output dc currents is TBD pending the results of the requirements for source end termination

Receiver requirements are exactly half in both amplitude and time from that used in SSA 20. Bill H questioned whether yet another halving was possible for SSA 80 or SSA 100. There was no clear answer at this meeting.

Tests of the AMP connector presented by Lisa Huff indicate that an edge rate needs to specified for the measurement of cable assembly impedance. Bill Ham suggested a 187.5 ps edge rate to allow two more doublings beyond that needed for SSA 20 at 750 ps. The AMP connector meets the impedance requirements for cable assemblies (150 +- 10 ohms) at 150 ps rates according to her data.

The number will be resolved at the February meeting in Stateline, NV. The connector folks are requested to provide TDR data at various rise times for that meeting.

63 Termination/Coupling

This item was covered in detail in previous sections. It was pointed out that we will have to specify source termination performance requirements if we choose to require it. Any choice that eliminates the possibility of optical or AC coupling was seen as inferior due to limits in intracab applications.

The April meeting is the target to chose a path here for SSA 40 and beyond.

64 Jitter Transfer/Tolerance

Assumed to be a long term issue; short term needs input from further testing. A more rigorous jitter specification will need to consider what frequencies the receiver PLL tracks. Dividing jitter into random and deterministic parts

will also be required. Today it is mainly deterministic (coming from the data patterns) because of copper only implementations at relatively slow data rates. Fibre Channel is having to deal with these issues and will provide a model to consider.

65 Measurement points

Bill Ham pointed out that the working draft currently specifies driver/receiver/termination characteristics at the port connector; there is no budget to divide up the requirements for silicon, connectors, cables, enclosures. It appears that a budget for enclosures would be most useful. There are additional issues surrounding LRC. This issue will be re-examined at the next plenary.

66 Optical Specifcations

Today these are included as a possible internal path in a port connection coupler. See minutes from San Jose and the last mailing for details. There is no means specified for implementing an interoperable low cost optical scheme for SSA. Therefore, this will remain a longer term task for the moment.

A recurring question in the meeting was the ability of optical schemes to reliably transmit DC levels. Bill H will report at the next meeting on this subject.

67 Option Block Functions

Take to sub-system study group in February.

68 Device connector

Due to recurring confusion about the desirability of using the Fibre Channel 40 pin SCA connector the study group considered whether it should have an alternative for PH2. There was no support for changing from the present multibay device connector to the SCA style. This was due to the ability of the present connector to support cables as well as backpanels and to the bifurcation of the SSA market that would result if alternate connector styles were to be offered.

Call for Patents

Larry Lamers requested that anyone aware of any patents required for the proposals be disclosed in accordance with the ANSI patent policy. Refer to the minutes of prior meetings for items already identified.

Action Items

10) none.

Meeting Schedule

Based on input from Chris Parker and Phil Murfet regarding the availability of test data the PH2 working group is scheduled for Monday and Tuesday of SSA week in April 1996. There is a small possiblity that a specific working goup may be held in early April.

The next working group meeting of X3T10.1 is scheduled for February 26, 1996 at Stateline NV (Harrah's)

The long-term SSA week of meetings are scheduled as follows:
Week of February 26, 1996 in Lake Tahoe, CA hosted by Samsung.

Week of April 29, 1996 in Burlington, VT hosted by IBM.
Week of June 24, 1996 in St. Petersburg Beach, FL hosted by AMP.
Week of August 26, 1996 in Ft. Collins, CO hosted by Symbios Logic, Inc.
Week of October 30, 1996 in San Jose hosted by Adaptec *
Week of December 9, 1996 in Hawaii hosted by IBM. *

* = Tentative locations

Please note that changes to this schedule may occur. All changes to meeting dates, locations, and agendas will be posted to the SSA reflector.

Adjournment

The meeting adjourned at 6:30 p.m. on Thursday Jan 18, 1996