Minutes of the Ultra SCSI Special Subject Working Group (SSWG) of the Small Form Factor (SFF) Committee.

Held March 2nd, 1994, at Milpitas, California.

Chaired by: Jim McGrath, Quantum Corp.
Hosted by: Jim McGrath, Quantum Corp.
Minutes recorded by: Robbie Shergill, National Semiconductor Corp.

Jim McGrath brought the meeting to order at 9:30 AM and asked each person in attendance to introduce themselves. An attendance sheet was also circulated.

ATTENDANCE LIST:

Name	Company	e-mail
Wally Bridgewater	Adaptec	wally@adaptec.com
Mike Dreitlein	Adaptec	ljlamers@adaptec.com
Norm Harris	Adaptec	nharris@adaptec.com
Shahe Krakirian	Adaptec	shahe@adaptec.com
Larry Lamers	Adaptec	ljlamers@adaptec.com
Craig Stuber	Adaptec	stuber@adaptec.com
Mark Knecht	AMD	mark.knecht@amd.com
Joseph Chen	Cirrus Logic	chen@cirrus.com
Bill Gintz	Conner	n/a
Louis Grantham	Dallas Semiconductor	grantham@dalsemi.com
Jens Paetau	Dallas Semiconductor	n/a
Michael Smith	Dallas Semiconductor	msmith@dalsemi.com
Bill Ham	DEC	ham@shwsys.enet.dec.com
Ed Kavetsky	DPT	kavetsky@dpt.com
Ron Roberts	Maxtor	ron_roberts@maxtor.com
Mike Farhat	National Semiconductor	farhat@rockie.nsc.com
John Goldie	National Semiconductor	cjfgsc@tevm2.nsc.com
Steve Ichikawa	National Semiconductor	tsmisc@eccc.nsc.com
Todd Nelson	National Semiconductor	ctwnsc@tevm2.nsc.com
Robbie Shergill	National Semiconductor	rss@berlioz.nsc.com
Mark Jander	NCR	mark.jander@colospgs.ncr.com
John Lohmeyer	NCR	john.lohmeyer@ftcollinsco.ncr.com
Harry Mason	NCR	n/a
Ting Chan	Q-Logic	t_chan@qlc.com
Jim McGrath	Quantum	jmcgrath@qntm.com
Brian Davis	Seagate	(to Gene Milligan)
Steve Finch	SSI	5723283@mcimail.com
Vit Novak	Sun	vit.novak@sun.com
Pete Tobias	Tandem	tobias_pete@tandem.com
Kasem Elkhalid	TI	4307725@mcimail.com
Kevin Gingerich	TI	4307725@mcimail.com
Paul Aloisi	Unitrode	aloisi@uicc.com
Tak Asami	Western Digital	asami@dt.wdc.com

Jim McGrath briefly described the SFF Committee's objective and its workings. He then presented the proposed agenda for this meeting. The agenda was discussed briefly and adopted as follows:

AGENDA:

Introductions	Mcgrath

Review of SFF Rules	Mcgrath
Objective of this SSWG	Mcgrath
Date and Place of Next Meeting	Mcgrath
Ultra SCSI (1) Concept	Ham
Media Results	Ham
Timing Budget	Ham

OBJECTIVE OF THIS SSWG:

Jim McGrath reiterated his needs from the Ultra-SCSI specification as previously stated in his Ultra-SCSI SSWG document. Jim has customer that need next higher level of transfer rates, and they are primarily interested in short distance, inside the box connections. So, the cable length can be traded for speed. Jim knows that there are also other people that want faster speeds over longer distances, outside the box. Such out-of-box applications have to be differential; but the short distance, inside-the-box applications should be kept single-ended for power and cost reasons.

John Lohmeyer raised the issue that he doesn't want this work continuing in SFF. He pointed out that this work started in X3T10; so why not keep it there?
Jim McGrath said that he has no problem with that - perhaps the next meeting in Newport Beach can be a joint SFF/X3T10 meeting. John said that he doesn't see a need for even that as we should have basic agreemnet today.

Bill Ham pointed out that the SPI-2 project proposal is being prepared as we speak and Ultra-SCSI transfer rate can be part of it. John Lohmeyer added that the SPI spec has only one public comment against it and it is from Lohmeyer on behalf of the committee. It is to include SCAM into SPI. SPI-2 project proposal will be posted to reflector within a few days.

Joe Chen asked about expectations over time. Jim McGrath said that he doesn't expect complete agreement today because people would want to go home and analyze. But he does expects this spec to be completed in 3 months' time.

NEXT MEETING:

The next meeting of this SSWG will be held on March 31, 1994. It will be held at this same location (Crown Sterling Suites Hotel in Milpitas, CA) and will be hosted by Quantum Corp. (Note that the next discussion on Ultra-SCSI will take place in the SCSI working group in Newport Beach, CA, on March 16, 1994).

ORIGINAL CONCEPT AND MEDIA RESULTS:

Bill Ham (DEC) took floor next and quickly went through his original presentation to X3T10. Key points of his proposal are:
- intended as informative annex to SPI.
- fully backward compatible with current SPI.
- uses exactly same physical connection.

Bill then presented the results of DEC's lab work. Bills' tester can go up to 30 mega transfers/sec (MT/s) and uses separate transceiver chips for both differential (National, DS36954) and single-ended operation (TI: 9-channel, BICMOS, SPI-compatible with 24-mA active-negation and termination on-board). The system sweeps across various data-to-clock time values and measures error rates.
ting chan (q-logic) pointed out that the worst case is when the recvr is very close to the xmtr but the bus is trmntd at the far end. bill agrees, his system can do this, but it appears that he has done his msrmts with rcvr at the far end.

Bill's differential system operated error-free (millions of transfers) with data-to-clock times of about 6ns to 40ns at 20 MT/s; and about 6ns to 27ns at 30 MT/s.

More discussion took place on the single-ended results. With a shielded cable, 5 ft long, signals look good at 20 MT/s with little jitter and data-to-clock times of about 7 to 44 ns allowed error-free operation over millions of transfers. With 1.5m flat-cable, thirteen wires-only stubs on 4" centers, and driving and receiving at mid-bus, times degrade only to about 8 to 43ns at 20 MT/s. However, when Bill mixed 1.5m round cable with 1.5m flat cable with its associated stubs, the signal got very bad. The Vol level was only about 1v and jitter was high. data-to-clock times went to about 14 to 36ns.

Bill concluded that short distance and controlled rise and fall times appear to provide a very friendly environment for 20 Mt/s operation. Key question is "can we control the silicon well enough ?"
(Bill will provide hard copies of his results at the Newport Beach meeting).

TIMING BUDGET:

Bill Ham outlined a set of "ground zero" assumptions for further discussion:

1. The Ultra-SCSI device will maintain chip timing and other specs that are as good or better than the present (SPI) spec when operating in a non-ultra mode.

Larry Lamers asked if we are going to allow old devices (pre-SPI) on the bus? Jim McGrath and Bill Ham are willing to go along with this limitation if needed.

2. We have almost a clean sheet of paper once the two devices are in Ultra-SCSI mode (data transfer).

3. specific areas to look at:
   - window around slew rates
   - receiver input spec
     Larry Lamers mentioned that glich filter requirements have to be relaxed.
   - max assertion level (Vol)
     Wally Bridgewater (Adaptec) stated that perhaps the Vol should be spec'd in an AC way. paul Aloisi (Unitrode) doesn't want to put more energy into the bus. He and Bill Gintz (Conner) asserted that the focus should be on the receiver.
   - setup and hold times
   - active negation parameters
   - clarifying how timings are defined.

John Lohmeyer asked about any possible emi issues.Bill Ham said that he hasn't looked at it yet, but older data says that data traffic is not where emi peaks are. Ting Chan said that it may not be a concern at 20MT/s but should be considered for rates above that.

4. Specify narrower chip (system) operational parameters that are needed in normal office/computer room environment.

5. Chip specs that can be met by at least two siliocn suppliers using established cost structures - ie, no BC process.

6. Stay with present packaging technology, but pinout changes are okay.

With the above list agreed to, Kevin Gingerich (TI) asked that we narrow down the system operational parameters before we get into details of silicon. This was agreed to and Bil Ham prioritized the above list in order of discussion.

However, the discussion soon moved into skew budget issues.

Bill Ham said that at least two cable manufacturers are telling him that cable skew (currently 4ns) can be cut to 2ns. Taking 2ns each from cable and distortion skew, we can gain 4ns in the budget.

The discussion moved into actual operational parameters such as loading and power supply range until Larry Lamers asserted that we're getting too deep in; let's back off and look at the skew budget.

At this point, John Lohmeyer presented NCR's proposal on skew budget.

Single-Ended:
  total cable skew = 4ns
  tx chip s/h = 12/17
  board skew still 1ns at each end
  rx chip s/h = 6/11ns

Differential:
  total cable skew = 4ns
  tx driver skew = 4ns
  rx receiver skew = 5ns
  board skew still = 1ns each
  tx chip s/h = 16/21
  rx chip s/h = 1/6

John asked if differential transceiver skew geting cut in half is okay with TI and National as this was a hotly debated SPI spec. Kevin Gingerich (TI) said that this is okay if the environment parameters are okay. John Goldie (National) said that John's numbers look feasible.

Larry Lamers showed Adaptec's proposed skew budget numbers. Adaptec shows assertion time degrading by only 2ns from tx to rx side. John Lohmeyer observed that Adpatec numbers will have trouble when applied to the differential case.

General sense of the group was that NCR and Adaptec proposals are close enough that any differences can be worked out after further analysis. NCR and Adaptec will bring revised proposals to the Newport Beach meeting.

Larry wanted to discuss the receiver switching points before dissecting the skew numbers any further. In particular, Adaptec wants to increase hysteresis to 0.4v and redefine spec measurement points with hysteresis considerd. Adaptec also wants to eliminate the glitch filtering requirement if the group agrees to increase hysteresis to 0.4 volts. No one else offered any objections on these points.

Steve Finch (SSI) drew some timing diagrams that served as basis for further discussion on timing measurement points. After significant discussion, the following measurement points were agreed to:

At the receiver and at the driver:

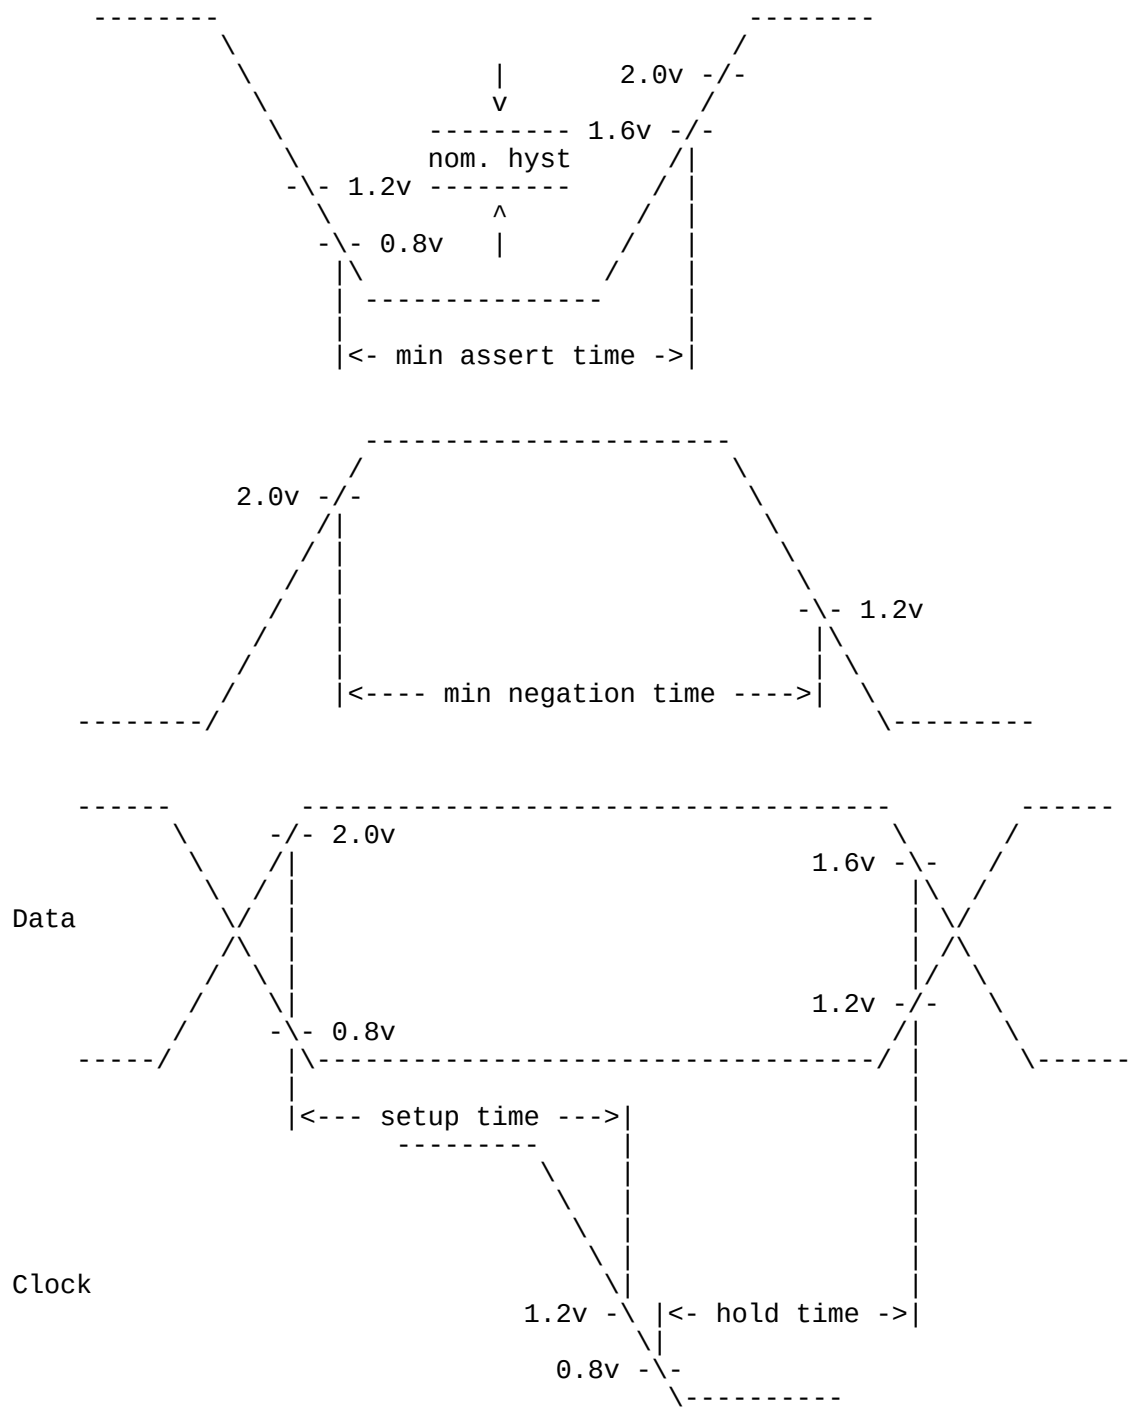

Rise and fall times were discussed next. Should the spec remain in terms of 10% and 90% points, or should it be changed to specific voltage levels such as 0.5 volts and 2.5 volts? There was agreement that the current SPI specification needs to be enhanced in this this area as the 0% and 100% ("full amplitude") voltage levels are not clearly defined.

Paul Aloisi (Unitrode) pointed out that maximum high voltage has to be specified, if 10/90 % points are used, so that your tr/tf absolute numbers don't kill the budget.

Robbie Shergill (National) pointed out that the good thing about 10% and 90% points is that the full transition range, where the greatest slewing takes place, is covered. Many people agreed that the best thing to do is to specify the maximum slew rate. Finally, the group agreed on 400mV/ns maximum on the test circuit. This is roughly the same as the current SPI spec.

Tak Asami (WD) wanted the present load circuit to be changed. All agreed that this needs to be looked at.

Bill Gintz (Conner) wanted receiver test conditions also to be specified. All agreed that this needs to be looked at.

Robbie Shergill wanted the System Deskew Delay time to be defined. John Lohmeyer clarified that this time is defined as half of total skew in system contributed by the two nodes (other than cable plant).

It was also mentioned that ground voltage offset spec. for the single ended case should be defined as well and the current RS-422 based spec for the differential case should be refined.

In the end, the group came up with the following wish-list for further lab work on Bill Ham's tester:

a. stubs actually loaded with devices (all SPI-compliant, 16-bit devices).
b. receiver a minimum distance (4") away from driver, which is at the end.
c. cable distance of 1.5m and 3m.
d. loaded with up to 16 devices.

The meeting was adjourned at 3:45 PM.