X3T9.2/87-67

Minutes of IPI/SCSI Higher Speed Working Meeting
April 1-2, 1987
Santa Clara, CA

Dal Allan, vice chair of X3T9.3 chaired the meeting. At the end of the meeting Don Tolmie was asked to write up minutes, and has done so to the best of his recollection. Thanks were expressed to National Semiconductor for hosting the meeting and providing excellent meeting facilities and support. The number of attendees was about three times larger than last time, and much greater than originally expected.

Attendees included:

Dal Allan	ENDL
Ed Balogh	Gould
Bob Bellino	Madison Cable
Bruce Blaylock	NASA Ames
Kent Bossange	Alliant Computer Systems
Dennis Bowman	Martin Marietta
Rich Branning	Dataram
Ed Cardinal	Amdahl Corp
Roger Cummings	CDC Canada
Ed Dacey	IBIS
Randy Davis	National Semiconductor
Gene Dornhoff	Los Alamos National Lab
Mel Genter	Scientific Computer Systems
Ed Gross	Interphase Corp
Robert Grow	GP Associates
John Hawkes	Elxsi
Rick Henderson	National Semiconductor
Steve Johnson	Cray Research
Mitch Kane	MMI
Sam Karunanithi	Hitachi Microsystems
John Kingsolver	DEC
Adam Levinthal	Pixar
Rich Lewis	DEC
Robert Liu	Fujitsu America
Jim Luttrull	Fujitsu America
Rick Mansfield	CDC-MPI
Tom McClendon	Convex Computer Corp
Michael McGowen	Los Alamos National Lab
Fred Meadows	NCR Corp.
Ken Miura	Fujitsu America
Ken Moe	IBM
Robert Morris	Simulex
Masanori Motegi	Amdahl
Gary Murdock	Fairchild Semiconductor
Bob Norman	MMI
Robert Ober	Scientific Computer Systems
Robert Olson	Elxsi
Peter Paluzzi	John von Neumann Nat'l Computer Center
Karen Parker	Fairchild
Vit Patel	CDC-MPI
Newt Perdue	Ultra Corporation
Ronald Perloff	GP Associates
Frank Preston	NASA - Ames
Michael Pugh	Integrated Photonics
Wayne Sanderson	Control Data
Jim Scheussler	National Semiconductor
Craig Scott	Interphase Corp
Kunar Sivasothy	National Semiconductor
Robert Snively	Adaptec
Jeff Stai	Western Digital
Wally St. John	Integrated Photonics
Don Tolmie	Los Alamos National Lab
Horst Truestedt	IBM
Jayshree Ullal	AMD
Sarosh Vesuna	AMD
Bill Wilson	Martin Marietta
Jim Withworth	Priam

HIGHER SPEED SIGNALLING

Bob Bellino of Madison Cable started off with a discussion of skew characteristics of twisted pair cables. By paying careful attention to cable lay, insulation size, and materials Bob felt that we could obtain +/- 2% skew difference between separate pairs in the same cable. On a 50 meter cable this amounts to about +/- 5 ns.

Rick Henderson and Kumar Sivasothy of National Semiconductor presented some improved performance figures for the RS-485 DS3695 transceivers based on constraining the voltage and temperature extremes. Improvements were about 30% - 35% over current specifications. It is not required to use chips from the same lot to obtain the improvements, but screening would be required. It is too soon to know how surface mount chips will compare to DIPs.

Roger Cummings of CDC Canada presented some summary sheets for computing total transmission prameters and their uncertanties. The worst case numbers were filled in on-line with input from the attendees and were based on the cable skew number provided by Bob Bellino, the improved performance transceivers, and 50 meters of cable. After working through the worst case numbers it was concluded that higher speed SCSI and IPI would never work. The major contributors to the uncertainty budget, in order of significance, were (1) intersymbol distortion - 20 ns, (2) receiver propagation delay and receiver rise/fall times - 14.1 ns, (3) cable skew - 10 ns, (4) signal crosstalk and reflections in the cable - 10 ns, (5) receiver jitter - 10 ns, (6) transmitter propagation delay, asymmetry, rise/fall, etc - 9 ns, and (7) clock logic setup/hold times for the transmitter and receiver - 5 ns each. There are also some 1 and 2 ns uncertainties due to board traces and drop cables.

In actual practice things seem to work better than our worst case analysis, what are we doing wrong? Can all of the factors occur at once? Can we use statistical methods to improve things? It was pointed out that a single vendor has tighter control over his system, and hence may not need to allow for the worst case, but in the standards arena where different vendors are interconnecting we do need to consider worst case. It was requested that organizations with experience with high speed signalling on cables help provide some guidance on what really works.

Some of the other possibilities discussed included (1) limiting the distance to less than 50 meters to decrease the cable skew, (2) using ECL drivers and receivers to reduce the

driver/receiver uncertainties, (3) limiting the topology to point-to-point to reduce reflections, (4) using run-length coding to reduce intersymbol distortion, and (5) limiting the higher speed version to co-ax to decrease the high frequencey signal attenuation.

HIGH SPEED OPEN CHANNEL

Don Tolmie of the Los Alamos National Laboratory led a discussion on goals for a much higher bandwidth channel based on some papers distributed before the meeting. The application areas for the channel were stated to be (1) processor to processor, (2) processor to real-time graphics display (video display), and (3) processor to disk system. At this point the intent was to consider only goals, although Los Alamos had also provided a strawman draft specification of a proposed channel.

1. Burst bandwidth - Los Alamos proposed 800 Mbit/s as the bandwidth goal, and there was no disagreement from the attendees.

2. Distance - It was agreed that 50 meters was a minimum distance. It was also agreed that extenders should not be precluded so that distances of up to 2 kilometers could be supported. The 50 meter distance limitation was based on a copper implementation, and after the earlier discussions copper made people nervous. Cray Research is currently offering an 800 Mbit/s channel, 64 bits wide, but limited to 75 feet. Cray Research is unsure at this time whether or not they would want to propose their channel as a standard.

3. Flow control - Sufficient hardware level flow control must be provided to keep the sustained bandwidth at at least X% of the burst bandwidth when using the maximum length cable. The attendees agreed on the need for this requirement but did not set an absolute value during the meeting.

4. Parallel/Serial - It was felt that this was an implementation issue and should not be decided now.

5. Double width allowed - There wasn't much discussion here, but I did get the feeling that people would look favorably on a modular approach, especially allowing narrower versions to run at slower rates, e.g., 200 Mbit/s.

6. Point-to-point/Multi-drop - It was felt that this was also an implementation issue and should be deferred. If our experience with the higher speed signalling for IPI shows that multi-drop is feasible, then it may be included. Point-to-point had been proposed as a way to improve the signal quality on the cable.

7. Half-duplex/Full-duplex - After much wrangling with semantics, it was concluded that what we wanted was a simplex channel. If full duplex is desired, then use two of the simplex channels. It also allows for a different channel in the reverse direction, for example a much slower channel in the reverse direction for a predominately output only device such as a graphics screen.

   Along with this item is the desire to have all signals be unidirectional for signal quality reasons as well as to make it easier to route through a crossbar switch.

8. Packet size - It was agreed that a large packet size is needed to keep the sustained bandwidth high while keeping the packets-per-second to a reasonable level. You do not want too large a packet size since that has a tendency to hog a channel and possibly freeze out other messages for the same destination.

   Don Tolmie proposed a max packet size of at least 128 KBytes based on extrapolation from the Ethernet packet size and data rate. It would be desirable to have larger packets for some applications, for example a 1 MByte video frame as one packet.

   As part of the packet size discussion the question of checksums was raised. The Los Alamos draft proposal used parallel Hamming check bits on a per word basis (ECC for SECDED), but did not include any serial checksums over multiple words. Since CRCs are length dependent, and channel repeaters that cannot buffer a full packet may be used, it was felt that the serial checking should be done on a smaller segment of the packet. One possible solution was a checksum at the end of each burst (every 1 KBytes) in the Los Alamos draft proposal. The attendees felt strongly that some sort of serial checking was required.

   No agreement was reached on a minimum packet size. A minimum of zero is desirable, but a minimum of 16 bytes or so may be preferred for ease of hardware implementation.

9. Addressing - Questions were raised as to why addressing was needed on a point-to-point link. Los Alamos felt the need for addressing since they planned to use the channel in a network environment. Further discussion concerned that if addressing was needed, should it be inside the packet proper, or as an out-of-band component. No resolution was forthcoming.

10. Immediate Status Response - Los Alamos proposed omitting the current IPI immediate status response to commands and data and using higher level protocols to report errors. The attendees felt strongly that an immediate error status line should be included so that retransmission can be triggered as soon as possible.

11. Hardware level reset - Los Alamos proposed a separate reset, or initialize, line to force the interfaces into a "ground zero" state at power on or error time. The attendees gave a mixed reaction, and the item was left as "recommended".

12. Protocol - The attendees felt that for a number of reasons, including addressing defined by discrete signalling sequences, protocol up to layer 4 was also required.

13. Network (crossbar) capability - Los Alamos has plans to use a crossbar to switch multiple channels for computer networking. The attendees felt that Los Alamos had not described their requirements sufficiently for them to judge what channel requirements were necessary to interface with a crossbar switch.

14. Dal Allan presented a viewgraph showing similarities in signal usage and data transfer sequences between the Los Alamos proposal and the streaming mode of IPI.

15. Name - It was suggested that the word "Open" be dropped from the Los Alamos proposal to avoid any confusion with OSI issues since the channel is intended for use in a closed rather than public system. Henceforth, the name will be "High Speed Channel" (HSC).

TAXI CHIPS

Jayshree Ullal and Sarosh Vesuna of AMD gave a presentation describing the Am7968/AM7969 TAXI (Transparent Asynchronous Xmitter - Receiver Interface) chips. They are FDDI compatible, provide 4B/5B encoding-decoding, clock recovery, parallel/serial conversion, and data rates up to 100 Mbit/s. Sarosh presented several examples showing how TAXI chips could be paralleled for wider data busses or higher bandwidths. The chips are available now at a cost of about $50 per chip set in 100-1000 quantities.

FUTURE MEETINGS

The next plenary meeting of X3T9.3 will be on April 27-28, 1987, at the Red Lion Inn in San Jose California. It was pointed out that the not much technical work gets done at the plenary meeting, it is mainly for coordinating activities between the different committees and formalizing the working committee results.

The next working group meeting will be held May 14-15, 1987, at National Semiconductor, 2900 Semiconductor Drive, Santa Clara, CA, and hosted by Balu Balakrishnan (408) 721-4283. No arrangements have been made for hotels, but for people coming into town the Sunnyvale Hilton is suggested as being convient (408) 738-4888.

ACTION ITEMS FOR NEXT WORKING MEETING

1. National Semiconductor to work on tightening up the specs on RS485 drivers and receivers for use in SCSI and IPI systems.

2. Everyone possible to look at ECL cable drivers and receivers, and the best specifications possible.

3. Vendors using high speed switching on existing channels to provide details and experiences on the cable and cable driving components.

4. Los Alamos to provide proposal and scope for the HSC physical and logical entities.

5. Additonal High Speed Channel proposals are solicited, especially ones that have already been implemented.

The meeting adjourned at 3 pm on Thursday, April 2.