

Minutes of X3T11 HIPPI SWG, and HNF - Technical Committee (TC)
December 4-5, 1995
San Diego, California

1. Opening remarks

The Chairman, Don Tolmie of Los Alamos National Laboratory, opened the Technical Committee portion of the HNF meeting by thanking Elen Bahr Renwick and NetStar for arranging for Monday's meeting room at the San Diego Convention Center, and thanking Dal Allan of ENDL for hosting Tuesday's meeting at the Clarion Bay View Hotel. This group is constituted as both the HIPPI special working group (SWG) under X3T11, and the HIPPI Networking Forum (HNF) - Technical Committee (TC). The meeting attendees are listed at the end of these minutes.

1.a HNF Plenary Business

Jim Toy started the meeting with a short HNF Plenary business meeting. The first item on the agenda was election of a new HNF Board of Directors. The six nominees were Bill Boas of Essential Communications, Simon Fok of Avaika Networks, Francois Gaullier of Hewlett-Packard, John Kerr of PsiTech, John Renwick of NetStar, and Don Tolmie of Los Alamos National Lab. Jim Toy asked if anyone had problems with Francois being nominated while his membership was still being processed; no one objected. Clive Towndrow of PsiTech moved, and Dick Kachelmeyer of NetStar seconded, to accept this slate. A roll call vote of the HNF voting members was recorded as:

AMCC	Not present
Avaika Networks	Yes
Broadband Communications	Yes
Cray Research	Yes
Essential Communications	Yes
IBM	Not present
Los Alamos National Lab	Yes
Maximum Strategy	Yes
Methode Electronics	Yes
Myriad Logic	Yes
NASA Ames	Yes
NetStar	Yes
Network Systems	Yes
PsiTech	Yes
Silicon Graphics	Yes
Triplex Systems	Yes
TRW	Not present

The motion passed: 14 yes, 0 no, and 3 not voting.

Bill Boas of Essential Communications reported on reactions to the HNF booth at Interop in Atlanta.

There seemed to be a general lack of HIPPI name-recognition and HIPPI awareness among the attendees. Bill proposed working with Nick Lippis of Strategic Network Consulting to try to get the HIPPI name and message to a wider audience. A contract with Nick would be on the order of \$20K; this is not a currently sanctioned HNF activity. Nick will be attending SC'95 on Thursday, and Bill invited all interested people to talk with Nick. The HNF will then consider a course of action.

2. Review / modify the draft agenda

A draft agenda for the technical committee portion of the meeting was distributed before the meeting via e-mail, shown at Monday's meeting, and distributed at the Tuesday meeting. No additional items were added. Don Tolmie agreed to take the meeting minutes. These minutes reflect the items on the approved agenda.

3. Next Generation HIPPI-PH Objectives

Don Tolmie started the presentations by showing a set of objectives for a next generation HIPPI-PH. They were:

- Bandwidth (on the order of 1 GByte/s, 6400 Mbit/s)
- Throughput (more significant than bandwidth)
- Latency ($< 1 \mu\text{s}$, the lower the better)
- Predictable (for B/W engineering and modeling ease)
- Distance (full speed over 1 km; copper for low-cost @ short distances)
- Some multiplexing ability (don't let one large message block everything else)
- Data ordering (in-order delivery at full speed)
- Compatibility with HIPPI (at HIPPI-FP level, or possibly HIPPI-PH format)
- Flow control (don't drop messages, don't overrun)
- Error detection (must be robust, need a 10^{-X} goal)
- Error recovery (consider retransmission in hardware)
- Cost (really cost/performance, and looking for a bounded number)
- Network management and auto configuration (we didn't do well on original HIPPI, and don't want a repeat of that experience)
- Simplicity (it worked for HIPPI, and usually complexity \neq speed)

- No options (it worked for HIPPI-PH, and options gave us problems elsewhere)
- Decoupled flow and error control (self explanatory)
- Time frame for availability (we need a basic capability in 1996)
- Expandability (how do we get 4x - 10x speeds in 5 years)
- Availability of parts or silicon designs (not practical with SSI, what chips or aids will be available, and when?)
- Low development risk (we can't afford a failure)
- Network capable (switching and multiple interconnections available)
- Testability (define "stimulus / observable response" sets)

4. Next Generation HIPPI-PH Presentations

Greg Chesson of Silicon Graphics Inc., presented a proposal for HIPPI-6400. Steve Scott of Cray Research presented a proposal for GigaRing. Discussion followed. Both presentations were well received, and both technologies are real in that they have both been implemented. Discussion followed.

5. Review Next Generation HIPPI-PH proposals and start selection process

It was pointed out that the architectures were quite different. GigaRing was a ring while HIPPI-6400 was switch-based, and the original HIPPI is also switch-based. GigaRing also used DMA transfers between trusted machines, while HIPPI is oriented towards message passing between heterogeneous machines, i.e., machines from different vendors. The present implementation of GigaRing does not provide in-order delivery at full rate, but can provide it at a slower rate. Running the HIPPI protocols over GigaRing would require some translation mechanism, while it more native to HIPPI-6400. HIPPI-6400 used two separate 16-bit CRCs, and a cumulative CRC over a whole message, while GigaRing used a single 16-bit CRC.

It was stated that both architectures made sense, and may well co-exist in systems. The question of where we should put our standardization efforts was raised, i.e., for a specific physical interface, or for the boundaries between the different systems, e.g., between HIPPI-6400 and GigaRing, between HIPPI-6400 and HIPPI-800, etc. Some people felt that it would be desirable to have two separate physical interface standards, one for each proposal.

They also felt that it was desirable to develop both standards within the X3T11 committee.

After some thrashing, Don Tolmie said that he would volunteer to lead a standardization effort based on the HIPPI-6400 proposal, and be the document editor. He also said he felt that the committee did not presently have the resources to develop two standards, i.e., one for each proposal, without some major commitments from someone else to be the technical editor of a GigaRing-based document.

The discussion then turned to what features we should commit to include as we write the SD-3 Project Proposal. Don Tolmie was actioned to draft an SD-3.

6. Review Minutes of Previous Meeting

The minutes of the X3T11 HIPPI SWG / HNF - TC meeting of October 3, 1995, in Toronto, Ontario, Canada were reviewed.

Steve Quan of NASA Ames moved, and Roger Ronald of E-Systems seconded, to approve the minutes as written. Passed unanimously.

Don Tolmie noted that previous meeting's minutes are available on the HNF WWW page at <http://esscom.com/hnf/>. They are also available in Postscript format via anonymous ftp from <ftp.network.com/hippi/minutes>.

7. Review of old action items

The action items from the August, 1995, meeting were reviewed for the current status.

1. Michael McGowen of Essential Communications - Generate a HIPPI switch MIB. (In process, carryover.)
2. Everyone to consider changes to HIPPI-SC, and submit drafts of their proposed changes. (In process, see item 11.)
3. Michael McGowen of Essential Communications - Update HIPPI-AC to work with HIPPI-SC and its recent changes. (In process, carryover.)
4. Don Tolmie of Los Alamos - Forward the revised SD-3 Project Proposal for HIPPI-AC to X3T11 for further processing. (Done, passed, see 10.2.)
5. Everyone to propose changes to HIPPI-LE. (Done, see item 15.)

6. Don Tolmie of Los Alamos - Issue a call via e-mail for proposed changes to HIPPI-LE. (Done, see 15.1.)
7. Jim Toy of BCP - Check with Tera Computing to see if they are using the HIPPI-Serial copper variant. (Done, see 8.1.)
8. Don Tolmie of Los Alamos - Issue a warning over e-mail that the copper variant of HIPPI-Serial will be removed unless support for it is received within a specified reasonable time. (Done, see 8.1.)
9. Don Tolmie of Los Alamos - Update HIPPI-Serial Rev 1.5 to answer the X3T11 Letter Ballot comments. (Done, see 8.1.)
10. Don Tolmie of Los Alamos - Generate replies to the X3T11 Letter Ballot commentors. (Done, see 8.2.)

8. HIPPI-Serial

8.1 Review Rev 2.0 changes

HIPPI-Serial Rev 2.0, and a change list, were distributed at the meeting, and are also available via anonymous ftp from <ftp.network.com/X3T11/hippi>. The changes to Rev 1.5 to make Rev 2.0 were extensive.

A call for interest in keeping the copper variant in HIPPI-Serial did not turn up anyone interested in keeping it. Tera Computer is using a copper version of HIPPI-Serial within their machine, but has no desire to connect to other vendors' hardware. Due to lack of interest in keeping it, the copper variant and its related text was removed from HIPPI-Serial. Also, an effort was made to make the document more of a lean and mean specification rather than a description of an implementation. Comments were made that Rev 2.0 is a real improvement over Rev 1.5.

Rev 2.0 was reviewed, and suggestions for changes were received. The changes will be incorporated into Rev 2.1. Any further suggestions for changes should be passed to Don Tolmie, the Technical Editor, as soon as possible. Don took an action item to update HIPPI-Serial Rev 2.0 with the agreed upon changes. The plan is to have a last look at the document at the February meeting, and forwarding it to X3T11 at that meeting.

8.2 Review responses to X3T11 letter ballot comments

X3T11 letter ballot No votes, and comments, were received from Ed Grivna of Cypress Semiconductor, and from Roger Cummings of StorageTek. Don Tolmie replied to their comments based on the changes made to the document, and the responses were distributed.

9. IETF related items

9.1 IP over HIPPI, RFC 1374

John Renwick of NetStar previously reported that RFC 1374 has been broken into two RFC drafts, one titled "IP over HIPPI", and the other "ARP over HIPPI". Both drafts were revised and have been submitted to IETF as internet drafts. They are available via anonymous ftp from the Internet-Drafts repositories as:

```
draft-renwick-hippiip-01.txt
draft-renwick-hippiarp-01.txt
```

John said that some the revised draft removed some options, resulting in just one way to construct a packet -- it is solidifying what people are actually doing. John solicited other comments. The committee agreed that John should put out an e-mail "last call" for changes. Lacking substantive comments, John is tasked with furthering the processing within IETF. The next IETF step is for the document to receive an RFC number (different from 1374), and process the document as an IETF Draft Standard.

9.2 ARP over HIPPI

It was reported that Hewlett-Packard, Essential Communications, and NetStar have implementations. Interoperability is being tested at the Supercomputing'95 show. It was agreed that this document should also be furthered within IETF, i.e., process it as a Proposed Standard and get a RFC number assigned.

9.3 HIPPI end-point MIB

John Renwick of NetStar previously reported that the HIPPI end-point MIB has been revised, and is available via anonymous ftp from the Internet-Drafts repositories as:

```
draft-renwick-hippimib-01.txt
```

Mark Kelley of Cray Research said that he will soon be doing an implementation. Based on his

experience, we will consider document modifications before further processing in IETF.

9.4 HIPPI switch MIB

Michael McGowen of Essential Communications took the action item to develop this document. Since Michael was not present for this agenda item nothing new was reported.

10. HIPPI-AC

10.1 HIPPI-AC document status

The document is in development within Essential Communications, and was not available at the meeting.

10.2 HIPPI-AC SD-3 project proposal status

Don Tolmie reported that this project proposal passed its X3T11 letter ballot on November 22 by a vote of 52 Yes, 0 No, and 9 not voting. X3T11 will now forward the project proposal to OMC for further processing.

11. HIPPI-SC

11.1 Distribution and review of HIPPI-SC Rev 2.9

Don Tolmie distributed copies of HIPPI-SC Rev 2.9. It is also available via anonymous ftp from ftp.network.com/X3T11/hippi.

11.2 Include point-to-point address self-discovery ?

John Renwick of NetStar had previously noted that the address self-discovery procedures in Rev 2.9 work well for HIPPI hosts attached to a HIPPI-SC compliant switch, but do not work when HIPPI hosts are attached point-to-point without an intermediate switch. It was felt that we should do whatever we can to make connecting HIPPI equipment as automatic as possible, and hence we should try to add whatever it takes to cover the point-to-point case. In addition, it would be useful if we could also signal the upper-layer protocols that on a point-to-point connection it is not necessary to drop the connection between packets, e.g., between 64 KByte IP packets as specified in the IP over HIPPI RFC.

A possible method was developed at the meeting. John took an action item to draft text and forward it to Don Tolmie, the document editor, for

incorporation in HIPPI-SC. John also said that he would look at ways to inform the HIPPI hosts that they did not need to drop the connection between packets -- this may require HIPPI-LE changes.

Ken Powell of Silicon Graphics requested that clauses D.1.2 and D.1.3, describing dropping READY indications, be included in HIPPI-SC. No objections were raised.

11.3 Are we ready to forward HIPPI-SC ?

With the inclusion of these changes, it was felt that we could forward HIPPI-SC to X3T11, for an X3T11 letter ballot, at the February meeting.

12. HIPPI API

12.1 Status of current work

Steve Poole of Performance Group previously offered to work on an expanded API, both for a driver level, and at a higher level. Ken Morris of Essential Communications will be the document editor with Steve Poole as the technical lead.

Steve announced that he had set up an e-mail reflector to discuss this work (contact Steve at sp@beta.lanl.gov to be added to this reflector). About a year ago Lansing Sloan of Lawrence Livermore sent out a questionnaire asking about which options are being used, and Steve has added to this list. Steve is proposing an API that is: high level, independent, simple, and expandable. This API would be above the driver level, and have only a few simple calls. Francois Gaullier of Hewlett-Packard, and Steve Quan of NASA Ames suggested that Steve Poole look at the Fore ATM API as a model. Rami El Sebeiti of Hewlett-Packard proposed including a lower level, i.e., more driver level, API (largely extracted from the HIPPI-FP primitives) as an aid to driver writers. It was agreed that this was a good idea.

13. HIPPI-ATM

13.1 Status of HIPPI-ATM implementations vis-a - vis Rev 1.5x

John Renwick reported that the NetStar implementation was running successfully. He wished that there was another implementation to connect to for an interoperability test, but we do not know of any other implementations running at the OC-3c rate. Rather than wait forever, it was

agreed not to hold the document any longer for interoperability testing.

Don Tolmie was actioned to issue a "last call" for changes over e-mail, with a notice that we planned to forward at the February meeting unless substantive changes were requested.

14. Proposals for higher speeds by striping

14.1 Further consideration of August '95 proposals

At the August meeting Clive Towndrow of PsiTech, and Don Tolmie of Los Alamos, presented proposals for striping a single ULP PDU over multiple HIPPI physical layers. Clive's proposal striped each n'th 64-bit word of a higher-layer PDU across a different HIPPI physical layer. Don's proposal broke the higher-layer PDU up into larger blocks with the blocks transmitted across different HIPPI physical layers.

There has been no further input since the August meeting, and the committee requested that Don flesh out the Los Alamos proposal as a draft standard.

15. HIPPI-LE

15.1 Call for changes to HIPPI-LE

In the process of updating HIPPI-SC with address self-discovery procedures, John Renwick of NetStar noticed that HIPPI-LE has some address self-discovery text in 7.1 and 7.2.3. In light of the HIPPI-SC changes, he suggested that this should be deprecated in HIPPI-LE; this can probably also be handled in the draft revision of RFC 1374, IP over HIPPI. Francois Gaullier and Rami El Sebeiti of Hewlett-Packard said that HP was using the HIPPI-LE text, and requested that it remain -- the committee agreed to leave it as is.

No other changes were suggested for HIPPI-LE at this time. This item will be dropped from future agendas.

16. HIPPI-FP

16.1 New ULP-id for SCSI

Steve Poole of Performance Group reported that he is interested in running the SCSI protocol over HIPPI, and considered requesting a new specific HIPPI-FP ULP-id for SCSI. Lansing Sloan of

Lawrence Livermore warned him of problems with this approach. Steve talked to Dal Allan of ENDL, and Roger Cummings of StorageTek, and they suggested using FCP (the SCSI protocol over Fibre Channel) and HIPPI-FC (the mapping of Fibre Channel upper-layer protocols to run over a HIPPI-FP link). These protocols are current work and should provide a well defined path. Other alternative were GPP (General Packetized Protocol for SCSI) but this was reported to have some performance problems, or to start from scratch.

Hence, no new ULP-id for SCSI is currently needed.

17. Other items

17.1 Document processing in ISO

Don Tolmie reported that the ISO versions of HIPPI-FP, HIPPI-LE and HIPPI-SC have recently passed the vote to become International Standards. The ISO editor's requested changes have been included in revised documents and the camera-ready copies have been forwarded to the ISO Secretariat for publishing. HIPPI-PH was completed a similar process a few months earlier. The names and document numbers are:

ISO/IEC 11518-1:1995 - Information Technology - High-Performance Parallel Interface - Part 1: Mechanical, electrical, and signalling protocol specification (HIPPI-PH).

ISO/IEC 11518-2:1995 - Information Technology - High-Performance Parallel Interface - Part 2: Framing Protocol (HIPPI-FP).

ISO/IEC 11518-3:1995 - Information Technology - High-Performance Parallel Interface - Part 3: Encapsulation of ISO/IEC 8802-2 Logical link control protocol data units (HIPPI-LE).

ISO/IEC 11518-6:1995 - Information Technology - High-Performance Parallel Interface - Part 6: Physical Switch Control (HIPPI-SC).

Note that the projects associated with parts 4 and 5 were canceled, i.e., HIPPI-MI and HIPPI-IPI-3.

17.2 HIPPI Users Group in USA

It was felt that the HIPPI Users Group meeting at the Supercomputing'94 show in Washington was very useful for collecting user desires and problems. A HIPPI Users Group was recently started in Europe, and there is interest in starting one in Japan, but the

one in the USA has languished as Jim MacDonald of the University of Minnesota has been pulled off to other things. Steve Quan of NASA Ames agreed to set up an e-mail reflector for a HIPPI Users Group.

18. Future meeting schedule

The next meeting of the X3T11 HIPPI SWG / HNF Technical Committee, will be Tuesday, February 6, 1996, from 9 AM to 5 PM, at the Catamaran Resort Hotel, 3999 Mission Boulevard, San Diego, California, phone 619-488-1081 or reservations by 800-288-0770 in the US or 800-233-8172 from Canada. The rate is \$125 single, \$140 double (tax included) + \$5 parking. The closing date for reservations is January 5. The group name is Vitro/ANSI, and the host is Vitro Corporation.

Other 1996 meetings are currently scheduled for:

Apr 10	Palm Beach, CA	Western Digital
June 11	Santa Fe, NM	Los Alamos
Aug 6	Honolulu, Hawaii	Hitachi
Oct 8	St. Petersburg Beach, FL	AMP
Dec 3	Rochester, MN (?) or Monterey, CA (?)	IBM Adaptec

NOTE - The meeting day for the April, 1996 meeting was changed from Tuesday, April 9, to Wednesday, April 10.

The 1997 meeting dates selected by X3T11, and the preliminary hosts, are listed below. Other hosts are being solicited.

Feb 3-7	??	??
Apr 7-11	Palm Springs, CA	Western Digital
Jun 9-13	Seattle, WA (?)	Boeing (?)
Aug 11-15	Honolulu, Hawaii	? Hitachi (?)
Oct 6-10	Tucson, AZ (?)	FSI (?)
Dec 1-5	??	??

19. Review action items

1. Michael McGowen of Essential Communications - Generate a HIPPI switch MIB.
2. Michael McGowen of Essential Communications - Update HIPPI-AC to work with HIPPI-SC and its recent changes.
3. John Renwick of NetStar - Issue a "last call for changes" for the IP over HIPPI RFC. Lacking substantive changes, further the processing towards an IETF standard.
4. Everyone to consider changes to HIPPI-SC, and submit drafts of their proposed changes.
5. John Renwick of NetStar - Draft text for HIPPI-SC address self-discovery in a host-to-host configuration without an intermediate switch.

6. Don Tolmie of Los Alamos - Integrate HIPPI-SC text and get the document ready for forwarding at the February meeting.
7. Don Tolmie of Los Alamos - Generate a draft SD-3 Project Proposal for a HIPPI-6400 project.
8. Don Tolmie of Los Alamos - Update HIPPI-Serial with the approved changes and get it ready to forward at the February meeting.
9. Don Tolmie of Los Alamos - Issue a "last call" for changes to HIPPI-ATM Rev 1.5x over e-mail, with a notice that we plan to forward at the February meeting unless substantive changes were requested.
10. Don Tolmie of Los Alamos - Flesh out the Los Alamos proposal for striping over multiple HIPPI channels.
11. Steve Quan of NASA Ames - Set up an e-mail reflector for a HIPPI Users Group.

20. Adjournment

The meeting adjourned at 4:00 PM. The Chairman was gratified with the meeting, having been previously worried that the close proximity of Supercomputing'95 would siphon people away for "booth duty" or other activities.

Notes from X3T11 Plenary following the HNF-TC

The X3T11 Plenary meet the next day, i.e., December 6. HIPPI related items are reported here for your convenience, the definitive record is the X3T11 minutes.

In a surprise announcement, Roger Cummings, the X3T11 Chairman, said that he was caught in a StorageTek down-sizing, and would be leaving the company by February 6. He said that he should know at the February meeting if he is going to continue his X3T11 participation, and if he will keep the Chair. StorageTek has agreed to fund him to attend the X3T11 meetings through the June meeting. Roger also said that Carl Zeitler of IBM Austin, X3T11 Vice Chairman, and Neil Wanamaker of Amdahl, X3T11 Secretary, are being pressured by their companies to cut back their time associated with X3T11 business. Stay tuned for further developments.

Meeting minutes of X3T11 HIPPI SWG / HNF-TC - December 4-5, 1995

Attendance

Attendance at the Monday meeting is marked with an "M", at the Tuesday meeting with a "T".

M	Ching-Ching Ganley	Abba Technologies	T	Bob Snively	Sun Microsystems
M	Bob Peach	Abba Technologies	M	Isao Morishita	Sumisho Electronics
M	Phil Duclos	Array Tech	M	Joe Jozen	Tokyo Electron America
M	Dick Wilson	Avaika Networks	MT	Ed Yoshino	Tokyo Electron Ltd.
M	Jim Toy	Broadband Communications	M	Bud Binck	Tri-Plex Systems
M	Joe Vaughan	Computing Devices Int'l	M	John Sobolewski	University of NM
M	Phil Kemp	Convex	T	Alan Iguchi	Vixel Corp
M	Robert Halford	Cray Research	MT	Stan Swirhun	Vixel Corp
M	Steve Johnson	Cray Research	M	Ken Pedersen	
MT	Mark Kelley	Cray Research	M	Wes Williams	
MT	BJ Kowalski	Cray Research			
M	Steve Scott	Cray Research			
T	Jeff Young	Cray Research			
T	Gary Demos	Demografx			
M	Dal Allan	ENDL			
M	Bill Boas	Essential Communications			
M	Dan Frigard	Essential Communications			
MT	Michael McGowen	Essential Communications			
M	Eirc Shieh	Essential Communications			
MT	Randy Leibow	E-Systems			
T	Sam Locke	E-Systems			
MT	Roger Ronald	E-Systems			
T	James Schroeder	Harris Corp			
T	Rami El Sebeiti	Hewlett-Packard			
MT	Francois Gaullier	Hewlett-Packard			
M	Steve Joiner	Hewlett-Packard OCD			
M	Dennis Ellis	HNSX Supercomputers			
M	Gary Grider	IBM			
M	Horst Truestedt	IBM			
T	Brian Smith	Infinity Commstor			
MT	Scott Hahn	Intel SSD			
M	Barry Wightman	Lion Rock			
M	Dale Malaney	Loral Defense Systems			
MT	Chris Olson	Loral Defense Systems			
MT	Don Tolmie	Los Alamos National Lab			
M	Chris Wood	Maximum Strategy			
MT	Tim Clay	Methode Electronics			
M	Pat Gilliland	Methode Electronics			
M	Jim McGinley	Methode Electronics			
MT	Richard O'Connell	Myriad Logic			
T	Mark Uebel	Myriad Logic			
M	Shigehi Hosoi	NetOne Systems			
M	Dick Kachelmeyer	NetStar			
M	Thomas Reiners	NetStar			
MT	John Renwick	NetStar			
M	Larry Schermer	NetStar			
M	Will Leslie	NEC Systems Lab			
MT	Jonathan Hahn	NASA Ames Research Ctr			
MT	Stephen Quan	NASA Ames Research Ctr			
M	Joe Parker	Optivision			
MT	Steve Poole	Performance Group			
M	Peter Lawthers	Prisa Networks			
M	John Kerr	PsiTech Inc.			
M	Clive Towndrow	PsiTech Inc.			
MT	Greg Chesson	Silicon Graphics			
M	Bob Newhall	Silicon Graphics			
MT	Ken Powell	Silicon Graphics			
MT	Don Sanders	Silicon Graphics			
M	Thomas Skibo	Silicon Graphics			
M	Steve Nelson	Steve Nelson & Assoc.			
M	Roger Cummings	StorageTek			