

Minutes of T11.1 HIPPI Ad Hoc Working Group

**February 9-10, 1999
Huntington Beach, CA**

1.0 Opening remarks and introductions

The T11.1 acting Chairman, Roger Ronald of Power Micro Research, opened this meeting at 9:00 AM and thanked Skip Jones and Lisa Truswell of QLOGIC for hosting this meeting.

This group is constituted as both the HIPPI Working Group under T11.1, and the High Performance Networking Forum (HNF) - Technical Committee (TC).

Roger lead a round of introductions. The list of attendees is at the end of these minutes. Roger volunteered to produce these minutes.

2.0 Review / modify the draft agenda

Draft agendas were distributed via e-mail before the meeting. Items added were:

Item 6.2, HIPPI-6400-SC new front matter (boilerplate) review.

Item 7.6, Read sequence table change to disallow the use of zero for a RTR I-Id.

Item 9.5 on the ST NOP command was moved up to be item 7.7.

These minutes reflect the approved agenda.

3.0 Review minutes of previous meeting

The minutes of the January 12-13, 1999, HIPPI Working Meeting in Albuquerque, NM, were reviewed and the following changes were made:

- Page 4, the limitation "unless the queue was full" was added to the second bullet under 9.2 discussing blocking writes.
- Page 5, the day and date for the second day of the March working meeting was corrected to be Wednesday, 3/10.

- Page 7, the month of the adjournment was corrected from December to January

Jeff Young moved and Jim Pinkerton seconded, to approve the January 12-13, 1999 working meeting minutes as modified. Motion passed unanimously.

4.0 Review old action items

The action items from the January meeting were reviewed.

1. John Gibbon to consider development and documentation of HIPPI-6400 MIBs. (**Carry-over**)
2. Jean-Michel Pittet to update "IP and ARP over HIPPI-6400 (GSN)" with the changes agreed to at the December meeting, forward to IETF as a RFC, and to Roger Ronald for posting on the HIPPI web page. (**1st part done, but the delivery to Roger Ronald and the web site posting is a carryover**).
3. Don Woelz to document the assignment of a ULA for the HARP server made during the meeting for inclusion in Jean-Michel's documents. (**Carryover**)
4. Jeff Young to research how to use ATM beyond LAN emulation (i.e. specify QOS parameters). (**Carryover**)
5. Nicolas Droux to develop a first draft of an Ethernet to HIPPI bridging ARP document. (**Carryover**)
6. Marck Doppke to process the HIPPI-800 switch MIB through the IETF. (**Carryover**)
7. Greg Chesson to draft an IEEE tutorial on HIPPI-6400 ULA usage and the ULAs special to HIPPI-6400. (**Carryover**)
8. Greg Chesson to collect information for an "ST FAQ" html document. (**Carryover**)
9. Greg Chesson to set up an e-mail reflector for ST API discussions. (**Carryover**)

10. Michael McGowen to develop "profile type" information for an ST over FC mapping placeholder document. **(Carryover)**
11. Michael McGowen to collect, tabulate, and document everyone's requirements for HIPPI-800 and HIPPI-6400 translation environments. **(Carryover)**
12. Joel Darnauer to confer with the HIPPI-6400 optics folks before the February meeting to discuss jitter. **(Carryover)**
13. Greg Chesson to draft a top-level structure for the ST-API document. **(Carryover)**
14. Jim Pinkerton to present material on bridging at the February meeting. **(Done)**
15. Nicolas Droux to present material on bridging HIPPI-800 at the February meeting. **(Done)**
16. Sean Rieb to write up his comments, conclusions, and recommendations on the cable shielding specified in HIPPI-6400-PH. **(Overcome by events)**
17. Roger Ronald to update the HIPPI-6400-SC document boilerplate and create a PDF file with bookmarks. **(Done)**
18. Don Tolmie to make the edits to the Scheduled Transfer document (rev 2.6) agreed to at this meeting. **(Done)**
19. Bob Pearson to provide Don Tolmie with wording for:
 - Limiting concurrent Put operations to a single persistent memory region and
 - Defining how adapters either support out-of-order Blocks or support only in-order Blocks within a set of Put operations. **(Done)**
20. Jim Pinkerton and Stephen Bailey to incorporate the changes agreed to at this meeting in their proposal for a SCSI over ST API (see Paragraph 9.1 of these minutes for a list of the items). **(Carryover)**
21. Stephen Bailey to define error codes and an error code format for adding to ST. **(Carryover)**
22. Everyone to consider if either or both a zero-length RTS and/or a NOP command should be added to ST. **(Done)**
23. Roger Ronald to provide Don Tolmie with wording for disabling the checking of S_Port for certain specified connections. **(Done)**

24. Roger Ronald to contact Debra Donovan to get an ANSI style guide for Jim Pinkerton. **(Done)**-email sent on 1/26/99
25. Roger Ronald to produce the minutes for this meeting. **(Done)**

5.0 HIPPI-6400-PH (ref: Rev 2.3)

5.1 Final processing Status

HIPPI-6400-PH Rev 2.3 was approved as ANSI standard "ANSI NCITS 323-1998, Information Technology – High-Performance Parallel Interface – 6400 Mbit/s Physical Layer (HIPPI-6400-PH)" on November 20, 1998. It has been forwarded to ANSI for final editing and publication. No further information has been received on when the standard will be published

5.2 Cable EMI problems

Joel Darnauer of SGI showed up to provide status on work being done to overcome reported RFI/EMI problems with the HIPPI-6400 cable. Joel has measured certain frequencies exceeding the FCC limit by 30 db.

A double shielded cable was being built for testing that would have two shields, one grounded at each end. However, preliminary results seem to indicate that only "hard-grounding" the shield at both ends will reduce emissions sufficiently for passing the test.

The group agreed that this was the correct technical approach, but consensus was not reached on whether the "hard-grounding" would be a recommendation or a requirement. It was agreed that the participants will ponder the question and try to make a decision at the March meeting.

Joel accepted an action item to provide his presentation for the web site.

5.3 Amendment to the HIPPI-6400-PH standard

Although Don Tolmie was not present, he provided a draft of his amendment to the HIPPI-6400-PH standard. Beyond changing the bit order, other changes include changing Berg's name to reflect their purchase by FCI. It's also fairly certain that there are RFI/EMI problems that will result in a need for change.

The draft amendment provided by Don Tolmie was reviewed. One change was noted, the word

“National” in the “caution notice” should be capitalized.

A call to ask for any further input on the project proposal calling for a an amendment (to change the order of bits across the 6400 interface) was made. No comments were received.

The project proposal to authorize the amendment to HIPPI-6400-PH will be formally voted upon at the T11.1 plenary session that will be held on Wednesday the 10th.

6.0 HIPPI-6400-SC (reference Rev 2.5)

6.1 Public Review status

HIPPI-6400-SC Rev 2.5 passed its second public review on November 24, 1998, with no comments. Unfortunately, it missed the NCITS meeting cycle and will have to wait for the March 30, 1999 NCITS meeting to complete its management review. No problems are foreseen, it will just take some time before we have an approved standard.

6.2 HIPPI-6400-SC Front Matter (boilerplate) Review

The new boilerplate for HIPPI-6400-SC with updated people and addresses/numbers was reviewed and accepted as correct.

Roger Ronald also noted that six older documents had been updated by Don Tolmie in a similar fashion and were on the web site, but that these changes would not be reviewed by the group (unless someone noticed any changes that were needed).

7.0 Scheduled Transfer (ref. Rev 2.7)

7.1 Review general document changes

Changes were reviewed and accepted with a few minor edits. Roger Ronald will pass these on to Don Tolmie for inclusion in the next revision of the document.

7.2 Changes for Put Flow Control

At the Albuquerque meeting, Roger Ronald had noted that Data operations on a single persistent memory region shared the same Mx and that efficient state utilization mandated that only one Put to a given region be received at a time. The group

agreed with this restriction and the need for a stated requirement in the document.

It was further agreed that implementations capable of out-of-order Block processing needed to be able to support filling in “holes” (by remembering the good, but out-of-sequence Blocks received) when good Blocks were received after bad Blocks. The Implementations that can't deal with out-of-order Blocks only need to keep the B-Num of the highest consecutive good Block received.

The new wording in the document for these items was accepted as written.

7.3 Changes for alternative connection creation methods

SGI had noted that state space “explosion” occurred if a unique connection was required in all cases. Roger Ronald provided some text to Don Tolmie that S_Port checking needed to be turned off for some connections. This text was reviewed and was accepted as written.

7.4 Other Proposed Changes

Jim Pinkerton and Steph Bailey believed that there were some flaws with the termination of transfers in progress. They believed that the current specification could not guarantee the reception of all sent data if a termination occurred (using End and End_Ack).

Some other options (using the optional payload field or using the Request_Status were suggested.

Jim and Steph accepted these options as useful and will report back if they change their opinion.

7.5 Processing plans

The main incentive for forwarding is to get some outside review and comments. However, there are still quite a few items getting comments and internally generated changes so it appears that forwarding will not be possible until at least April (and probably not before June).

7.6 Read Sequence Change

Ian Philp noted that it was not possible to tell a read sequence RTS (sent in response to a RTR) from a write sequence RTS if zero was the selected D-Id. Thus it was decided that all FFFFs was required in the D-Id for a write sequence.

Ian agreed to provide text for this change to Don Tolmie and was granted an action item for his kindness.

7.7 NOP Command

Don Tolmie's separate NOP command proposal was reviewed and generally accepted with a few changes:

- There was a desire to be able to use all free fields in the command and Don had not included the Param field in the opaque data.
- The group felt that the op-code for a NOP should not be in the connection space (where it was with a value of zero), but should be given a value of "10".

8.0 HIPPI-LE (reference Rev 3.7)

The main gating item before forwarding HIPPI-LE is inclusion of the RFC number for the "ARP and IP Broadcast over HIPPI-800" RFC in clauses 4, 6.1.1 twice, and Annex B. No action occurred at this meeting.

9.0 ST-API Mappings

9.1 SCSI over ST API

Steph Bailey was at this meeting to further the proposal. Steph did not present a new version of the document but lead a general discussion of some open issues. A new document version was promised for the next meeting.

There were a number of items that had previously been discussed that were carried further in discussions as well as some new topics:

- Only some portion of the I-id was really available for local ULP use in identifying transactions. At least some of the I-id was still needed to prevent duplicate operation aliasing. It was generally agreed that ULPs could have the freedom to use the ID fields in any manner that they wanted, if they were willing to tackle the aliasing issues on their own. SCSI will take this approach and guarantee proper termination of sequences so that it can effectively use a smaller ID name-space.
- It seemed worthwhile to at least provide an option that would allow for a security authentication handshake prior to allowing the execution of disk operations (otherwise, any VIA client could

access any disk connected in this fashion). Steph agreed to add an "authentication required" bit, although it was unclear if he would actually define the authentication handshake.

- It needed to be noted that the ST specified keep-alive option was required for this ULP but just for the target.
- No change would be made to allow use of different data channels with a single persistent memory region. If different data channels were needed, multiple RMRs would occur.
- Error codes still need to be defined
- A zero length transfer was needed for this ULP. All agreed that using a NOP command was a better approach (see Paragraph 7.7).
- Steph discussed that the nature of SCSI demanded proper treatment of "residuals" (i.e. overruns and underruns). This was discussed at length and the tools available with ST seemed sufficient. The leaning seemed to be that the optional payload might be extended carry some residual length information
- Virtual Connection tear down would be handled by tearing down all transfers in progress prior to tearing down the VC.

9.2 SGI Proposal

SGI's proposal was presented by Eric Salo. Eric did not have an updated document, but promised one for the next meeting.

Eric went through the call sequences he expected to set up for a data transfer. These included:

- Reserve Bufx (and de-reserve)
- Allocate MX (and deallocate)
- Prepare (pin) buffer (and deallocate)

Ian Philp and Sunlung Suen expressed some doubts about the static nature of Bufx assignment, but those with (or building adapters) were steadfast that Bufxs were a scarce resource and could not easily be dynamically assigned.

9.3 Bridging

Nicolas Droux of Essential/ODS made a short presentation on bridging with HIPPI-800 and Ethernet. There were some questions on whether the HARP server in bridged environment also needed to sup-

port 802.1d, but the general consensus was that the function was separate.

The chair commented that reviewing a draft specification was important to insure that the work and discussions performed by the group was not "lost", but captured over time. Nicolas is aware of his action item to provide a draft document on bridging.

9.4 Error Codes

There was nothing new for this meeting on error codes.

9.5 Document Plans

Roger Ronald agreed to send a boilerplate document in Framemaker format to Jim Pinkerton. All in attendance claimed to agree with the chair that writing our work down was important and that the time had come to start the process.

10.0 Other HIPPI items

10.1 ARP and IP Broadcast over HIPPI-800

Jean-Michel forwarded this document to the IETF in December. We are waiting to get an RFC number for this RFC, so that it can be added to the HIPPI-LE document as a reference.

10.2 IP and ARP over HIPPI-6400 (GSN)

Jim Pinkerton reported that Jean-Michel Pittet forwarded this document to the IETF, but a copy still is needed for posting on the web site.

Jeff Young and others had a few comments on the existing approach for HARP and broadcast, but Jim Pinkerton was able to convince the group to defer discussion of the topic until Jean-Michel attended at the next meeting.

10.3 HIPPI end-point MIB

If we want to pursue this MIB further then we need a champion for it. Nothing new at this meeting.

10.4 HIPPI switch MIB

Marck Doppke of Essential Communications has a draft document out for comment. Michael McGowen previously offered Marck's services for processing the HIPPI-800 switch MIB through IETF,

and Marck has acknowledged this assignment. Nothing new at this meeting.

10.5 HIPPI-6400 MIB

Von Welch of NCSA has a draft document, based on HIPPI-6400-PH Rev 1.4, out for comment. Von was not at this meeting and nothing new was reported.

John Gibbon had previously stated that Essential Communications may be interested in picking up MIB work on the HIPPI-6400 switch and he volunteered to have this action item reassigned to himself.

10.6 Tutorial for HIPPI-6400 ULA use

Greg Chesson has obtained the format material from Bob Snively of Sun. Drafting of the actual text is pending. Greg verified that this item is not holding up forwarding ST.

11.0 Future meeting schedule

11.1 Interim meeting, March 9-10 in Richardson (Dallas airports), TX

The next interim working meeting will be hosted by Roger Ronald and Power Micro Research in Richardson, TX. The location will be the PMR facilities, Suite 1100, 1411 East Campbell Road, Richardson, Texas. See the meeting notice on the HIPPI Standards Activities web page at www.hippi.org for details.

Tuesday 3/9 1 PM - 9 PM HIPPI ad hoc

Wednesday 3/10 8 AM - 5 PM HIPPI ad hoc

11.2 Plenary week, April 12-16, Palm Springs, CA

The April Plenary week location is in Palm Springs, CA. The hotel information has been linked to the www.hippi.org web site. The tentative schedule is as follows:

Tuesday 4/13 9 AM - 6 PM HIPPI ad hoc

Tuesday 4/13 6 PM - 9 PM HIPPI-6400-OPT

Wednesday 4/14 9 PM - 6 PM HIPPI ad hoc

Wednesday 4/14 6 PM - 8 PM T11.1 Plenary

11.3 Future meeting dates and locations

The T11.1 (i.e., HIPPI), Plenary meeting will be on Wednesday evening of the T11 Plenary week, following the HIPPI working meetings. The HIPPI and T11.1 meeting days are not specified within the Plenary weeks.

Interim working meetings have been scheduled for May 13-14 (originally May 11-12) in Milwaukee, WI (with Genroco as the host) and July 13-14 in Seattle, WA (with Tera as the host).

1999

May 13-14	Interim	Milwaukee, WI	Genroco
Jun 7-11	Plenary	Minneapolis, MN	Ancor
Jul 13-14	Interim	Seattle, WA	Tera
Aug 2-6	Plenary	Minneapolis, MN	ENDL
Oct 4-8	Plenary	Miami Beach, FL	Jaycor
Dec 6-10	Plenary	Reno, NV	Solution

2000

Feb 7-11	Plenary	Huntington Beach, CA	Qlogic
Apr 3-7	Plenary	Palm Springs, CA	Brocade
Jun 5-9	Plenary	Boise, ID	HP
Aug 7-11	Plenary	** open **	?
Oct 2-6	Plenary	Manchester, NH	Hitachi
Dec 4-8	Plenary	Austin, TX	Crossroads

2001

Feb 5-9	Plenary	So. California	Qlogic
Apr 9-13	Plenary	Palm Springs, CA	Brocade

12.0 Review action items

The following action items were carryovers or were added at this meeting.

1. John Gibbon to consider development and documentation of HIPPI-6400 MIBs.
2. Jean-Michel Pittet to forward "IP and ARP over HIPPI-6400 (GSN)" to Roger Ronald for posting on the HIPPI web page.

3. Jean-Michel Pittet to provide the IETF RFC number for "ARP and IP broadcast over HIPPI-800 (GSN)" to Don Tolmie and Roger Ronald.
4. Don Woelz to document the assignment of a ULA for the HARP server made during the meeting for inclusion in Jean-Michel's documents.
5. Jeff Young to research how to use ATM beyond LAN emulation (i.e. specify QOS parameters).
6. Nicolas Droux to develop a first draft of an Ethernet to HIPPI bridging ARP document.
7. Marck Doppke to process the HIPPI-800 switch MIB through the IETF.
8. Greg Chesson to draft an IEEE tutorial on HIPPI-6400 ULA usage and the ULAs special to HIPPI-6400.
9. Greg Chesson to collect information for an "ST FAQ" html document.
10. Greg Chesson to set up an e-mail reflector for ST API discussions.
11. Michael McGowen to develop "profile type" information for an ST over FC mapping placeholder document.
12. Michael McGowen to collect, tabulate, and document everyone's requirements for HIPPI-800 and HIPPI-6400 translation environments.
13. Joel Darnauer to confer with the HIPPI-6400 optics folks before the April meeting to discuss jitter.
14. Greg Chesson to draft a top-level structure for the ST-API document.
15. Don Tolmie to make the edits to the Scheduled Transfer document (rev 2.8) agreed to at this meeting including the addition of the NOP command.
16. Jim Pinkerton and Stephen Bailey to incorporate the changes agreed to at this meeting in their proposal for a SCSI over ST API (see Paragraph 9.1 of these minutes for a list of the items).
17. Jim Pinkerton and Steph Bailey to continue to evolve the concepts of SCSI over ST, particularly with respect to sequence termination, and make a presentation of their recommendations at the March meeting.
18. Stephen Bailey to define error codes for ST.

19. Roger Ronald to provide Don Tolmie with wording for how "id's" should be chosen, replacing the current requirement for "monotonically increasing" values.
20. Eric Salo to write down the concepts he has presented at the last few meetings on a level 1 ST API.
21. Joel Darnhauer to provide a copy of his slides for posting on the web site (**done**)
22. Ed Cady to provide the committee with information on the PCI version of the FCI-Berg Micropax connector.
23. Ian Philp to provide Don Tolmie with text for updating ST to prevent D-Ids in Read and Write sequence RTS operations from being confused (all FFFFs will be required in a Write sequence).
24. Roger Ronald to provide Jim Pinkerton with a Framemaker template for a standards document.
25. Roger Ronald to produce the minutes for this meeting. (**done**)

13.0 Adjournment

Jeff Young moved and Don Woelz seconded to adjourn at 12:42 PM on February 10th. There were no objections.

Attendance

Name	Company	Phone	E_mail
Terri Donaghey	Circuit Assembly	949-855-7887	
Christine Stoneking	Circuit Assembly	949-855-7887	chriss@circuitassembly.com
Nicolas Droux	Essential / ODS	505-344-0080	droux@esscom.com
Ed Cady	FCI-Berg	503-359-4556	edcady@aol.com
Stephen Bailey	Genroco	414-644-8700	steph@genroco.com
Donald Woelz	Genroco, Inc.	414-644-2505	don@genroco.com
Tan Truong	IBM	512-838-7813	truongtm@us.ibm.com
Ian Philp	Los Alamos National Lab	505-667-4305	philp@lanl.gov
Sunlung Suen	Los Alamos National Lab	505-661-7284	sunlung@lanl.gov
Roger Ronald	Power Micro Research	972-437-9461	rronald@pmr.com
Walter R. Smith	Raytheon Systems Co.	972-205-5378	wrsmith@gar.esys.com
Jeff Young	SGI/Cray Research	651-683-5536	jsy@cray.com
Joel Darnauer	Silicon Graphics	650-933-7682	joeld@engr.sgi.com
James Pinkerton	Silicon Graphics	650-933-4943	jimp@sgi.com
Eric Salo	Silicon Graphics	651-683-5607	salo@sgi.com
Albert F. Kelley	Tensolite Co.	904-829-5600 x281	akelley@tensolite.com
Sean M Rieb	Tera Computer Co.	206-701-2164	srieb@tera.com
Herb Van Deusen	W. L. Gore and Associates	302-368-2575 x5086	hvandeus@wlgore.com
Larry Green	xtp.com	805-969-2385	green@xtp.com