

Minutes of T11.1 HIPPI Ad Hoc Working Group

May 13-14, 1999

Milwaukee, WI

1.0 Opening remarks and introductions

The T11.1 Chairman, Roger Ronald of Power Micro Research, opened this meeting at 1 PM and thanked Carl Pick and Don Woelz of Genroco 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 the www.hippi.org web site before the meeting. One item was added under the ST section (7.0) for discussions on DIX Ethernet.

These minutes reflect the approved agenda.

3.0 Review minutes of previous meeting

The minutes of the April 13-14, 1999, HIPPI Working Meeting in Palm Springs, CA were reviewed. Don Tolmie had noted prior to the meeting that the wrong month was associated with some of the text for the June 99 meeting in Minneapolis (this text had already been corrected for the minutes on the web site).

Jeff Young moved and Sean Rieb seconded, to approve the April 13-14, 1999 working meeting minutes as amended. Motion passed unanimously.

4.0 Review old action items

The action items from the April 13-14, 1999 meeting were reviewed.

1. John Gibbon to consider development and documentation of HIPPI-6400 MIBs. (**carryover**)
2. Jean-Michel Pittet to provide the IETF RFC number for "ARP and IP broadcast over HIPPI-800 (GSN)" to Don Tolmie and Roger Ronald. (**carryover**)
3. Marck Doppke to process the HIPPI-800 switch MIB through the IETF. (**carryover**)
4. Greg Chesson to draft an IEEE tutorial on HIPPI-6400 ULA usage and the ULAs special to HIPPI-6400. (**carryover**)
5. Greg Chesson to collect information for an "ST FAQ" html document. (**carryover**)
6. Michael McGowen to collect, tabulate, and document everyone's requirements for HIPPI-800 and HIPPI-6400 translation environments. (**overcome by events**)
7. Greg Chesson to draft a top-level structure for the ST-API document. (**carryover**)
8. Don Tolmie to make the edits to the Scheduled Transfer document (rev 2.9) agreed to at the April 99 meeting. (**done**)
9. Don Tolmie to differentiate buffer size from ST buffer size in the ST document. The group felt that almost every occurrence of "buffer size" should be preceded with "ST". "ST Buffer Size" would be a clear and distinct label for the size that is used for tiling into the users total receive space (the total *user buffer size* clearly would be a greater than ST Buffer size"). (**done**)
10. Don Tolmie to add a definition for "ST Buffer Size" noting that this is the unit of memory addressed by ST for Bux and Offset calculations. (**done**)
11. Jim Pinkerton to investigate and report back on the use of DIX Ethernet. (**overcome by events**)
12. Steph Bailey to update the SCSI over ST text with the changes agreed to at this meeting. (**done**)

13. Jim Pinkerton to write a requirements paragraph on error codes for ST. Send to Don Tolmie ASAP as this may be the only hold up for forwarding the ST document. (**done** by Steph Bailey)
14. Eric Salo to update the level 0 ST API text with the changes agreed to at the March/April meeting. (**done**)
15. Ed Cady to provide the committee with information on the PCI version of the FCI-Berg Micropax connector. (**carryover**)
16. Eric Salo to investigate the proper way to specify field sizes in the ST API and whether the fields should be signed or unsigned integers. (**partially done**, a format was chosen but updates were not complete)
17. Roger Ronald to produce the minutes for the April 99 meeting. (**done**)
18. Roger Ronald to post the HIPPI-800 to GSN bridging document on the web site. (**done**)

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

The draft amendment provided by Don Tolmie is waiting on resolution of cable EMI issues prior to completion.

A call to ask for any further input on the project proposal calling for a an amendment (to change HIPPI-6400-PH with respect to the order of bits across the wire) was made.

5.1 VC2 Usage

During discussions last month, a possible additional change for HIPPI-6400-PH was identified. Because HIPPI-6400 is flow controlled, a slow host talking to a fast host results in both sides operating at the slowest common denominator for the VC in use. If only a single VC can be used, this means that an entire system is only as fast as the slowest client. If more than one VC can be used, the fast system is capable of serving more slow clients.

It was also noted that mis-matched rates would become institutionalized should a GSN-Lite be adopted (supporting GSN at half speed).

Since VC2 is not designated as a channel for any specific traffic and overlaps with VC0 in its size range, most in the group agree that it is underutilized. The proposal that surfaced during the previous meeting would modify the Message size

limitations of VC2 and allow larger Messages (perhaps as large as the VC3 limit of 4GB).

However, during further discussions, the downside to allowing a larger mix of Message sizes become visible.

- Admin responses are sent using VC2. If large Messages are present, blocking of Admin responses may occur. A single unlimited size Message is over five seconds in length, even when the link is full rate and no other VC is active while the time-out for Admin responses is one second.
- One problem with mixed sizes on a VC is that inadvertent priority effects can result in severe bandwidth degradation for clients that use small Message sizes. The current size limit reduces concerns in this area.

While Roger Ronald of PMR had been prepared to propose that HIPPI-6400-PH be changed to allow the same Message sizes on VC2 that are now allowed on VC3, the above listed concerns convinced the group that quick action was unwise.

Further, Genroco noted that flow control problems from adapters only existed with adapters that lacked some sort of suitable on-board buffer that could hold entire Messages. Everyone in attendance agreed that adapters built in this way would provide superior results.

In the absence of a consensus to proceed, the chair accepted an action item to disseminate this information to the HIPPI e-mail reflector and to call for any further proposals in this area.

Don Tolmie presented some text on VC usage recommendations and this information was accepted as valuable and appropriate for the HIPPI-6400-PH amendment. Don accepted an action item to add the text.

5.2 EMI and Cable problems

No input on the cable testing was available from SGI. The chair accepted an action item to collect status from SGI on this topic.

In previous months, it had been indicated through preliminary testing 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 at the February meet-

ing was not reached on whether the "hard-grounding" would be a recommendation or a requirement.

Berg was constructing a cable with a modified shield. The current shield was mylar, aluminized on one side only. Due to the way the shield wrapped around the cable, the ground path was basically helical and this was considered a problem. The new shield will be aluminized on both sides.

The cable, grounded on both ends, but prior to the shield change fails the FCC test criteria by 2db at 850 MHz.

It was also noted that a PCI capable connector/cable assembly was needed and that Berg was supposed to be working on one. However, there was no information presented at this meeting on a PCI compliant cable/connector assembly.

5.3 Draft Amendment Processing Plans

Since the cable issue was not resolved, we will not be able to forward the amendment at the June meeting. August seems a more reasonable date, but even August will not be achieved without an increased pace of progress.

6.0 HIPPI-6400-SC (reference Rev 2.5)

6.1 Processing 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 had to wait for the March 30, 1999 NCITS meeting to complete its management review.

The management review is now complete (it passed) and an NCITS letter ballot will complete on May 13. No new information was available at this meeting. Hopefully, the final procedural steps in the process will complete quickly and HIPPI-6400-SC will become a standard within the next month or two.

7.0 Scheduled Transfer (ref. Rev 3.0)

7.1 Review general document changes

Changes were reviewed and accepted with a few minor edits. Don Tolmie accepted these for inclusion in the next revision of the document.

7.2 Error Code Formats

Jim Pinkerton had previously been given an action item to provide some text for the ST document detailing and describing a format for ST Error Codes.

During this meeting, Steph Bailey stepped in and provided some text to Don Tolmie for error code formats. A general format was defined and specific error code definition was left to the ULPs.

This text was considered acceptable and Don Tolmie included it into the ST document sometime after the meeting on Thursday night and before the meeting on Friday morning. Thus, the group was able to review the included text in its final format. The addition was accepted.

7.3 Other Proposed Changes

Tom Gilbert had provided comments on a couple of points in the ST document. He noted that the document made a statement about assuming transmit buffer size in Gets and argued that the value should be explicitly passed in the connection sequence.

In the ST document, there is no other mention of transmit buffer sizes. The group felt that there was no reason to even imply a transmit buffer size, but felt strongly that data placed using a Put operation should be at the same location as data fetched with a Get operation. The section was reworded to eliminate the offending wording.

Tom also pointed out a spot where some text identifying the "unknown or reserved value" for R-id was needed.

7.4 DIX Ethernet

Arie Van Praag from CERN presented a proposal for carrying ST traffic over DIX Ethernet.

The reception from the group was positive and it was agreed that directly documenting ST on DIX Ethernet was a much better approach than documenting ST on DIX Ethernet on ATM LANE.

In fact, there was consensus that all of the ATM mappings (token ring, 802.3, and DIX) were unneeded and should be removed.

The only downside for ST on DIX Ethernet was the absence of a length field. Without a length field, fill at the end of a payload is indistinguishable from

payload data. Arie Van Praag accepted an action item to investigate whether the eight bytes of zero fill replacing the LLC/SNAP header could be used to carry a length field.

7.5 Processing plans

ST seems almost finished, but the requested changes this month will prevent the document's forwarding in June. August seems very reasonable as a target.

8.0 HIPPI-LE (reference Rev 3.7)

8.1 Processing Status

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.

According to Jean-Michel Pittet, the IETF RFC number should be provided by the June meeting.

8.2 Other Proposed Changes

A call was made for anyone with requirements to change HIPPI-LE to come forward. No change proposals were made.

9.0 ST-API Mappings

9.1 SCSI on ST API

Steph Bailey reviewed his latest version of SCSI on ST.

Steph walked through the steps required for an "abort". One important step was to guarantee that all Blocks are completed; either by sending all of the data STUs in the Blocks or by sending an in-order STU with the Last bit set (which might be done as a zero length STU if there was a need for the abort).

Don Tolmie asked if the bit definitions in section 2.3.5 were duplications of SCSI documentation. Steph responded that these bits were defined for SCSI on ST and needed to remain.

After some discussions about the lack of progress in moving forward towards a document (during the Level 0 ST API talk), the group agreed that SCSI on ST ought to be a separate project. The chair accepted an action item to write a project proposal.

9.2 SGI Low Level API Proposal.

Eric Salo was present on Thursday afternoon to present his latest Level 0 API document.

Included in the text was a format for defining integers of a specific size (requested as an action item to allow cross platform compatibility). This format was accepted, but the document still needs further updates to assure that all of the variables are in the correct size and format.

Ian Philp discussed how Schedule Headers could be pulled from a single Schedule Header queue that supported many connections. The general consensus was that this should not be a problem. If many Virtual Connections were mapped to a single queue. In fact, it seemed unlikely that the Schedule Headers could be differentiated based on the Virtual Connection from such a combined queue.

Steph Bailey pointed out a potential problem in the ST API (this happened on Friday morning after Eric Salo had left). There is nothing written that prevents a Level 0 API client from trying to use the same MX for a number of transfers and/or persistent memory regions) even though this approach won't work in all implementations. The document needs words describing the limits on MX usage and this wording may have to account for differing implementations with different limitations.

Don Tolmie expressed frustration with the lack of progress towards the goal of a standards document. Spotty attendance and participation by various people who have been involved in the effort has reduced progress even when the meeting time has not been completely utilized. It's hoped that the pending change back to bi-monthly meetings will reduce the standards "attendance fatigue" and encourage participants to treat each meeting as important.

9.3 Document Editing Plans

As noted in the SCSI on ST section of these minutes, a decision was taken to split the documentation effort for the Level 0 ST API away from the SCSI on ST mapping. SCSI on ST has a defined champion, willing to do the needed editing towards a standard. It's not clear if anyone is willing to do what is needed to make the Level 0 ST API into a standard.

Both Jim Pinkerton and Greg Chesson have a boil-erplate document in Framemaker format.

Since Don Tolmie might do some early document editing, it was possible/probable that FrameMaker might not be the correct editor.

10.0 Other HIPPI items

10.1 Bridging to HIPPI-800

At the last meeting, Nicolas Droux gave Roger Ronald an electronic copy of his early work on HIPPI-800 bridging and an action item to post it. The document is available on the web site.

Nicolas was not present, nor was any one else able to provide insight into this work, so there was no further progress at this meeting on this topic.

10.2 ARP and IP Broadcast over HIPPI-800

Jean-Michel forwarded this document to the IETF in December and again in February. 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. Nothing further at this meeting.

10.3 IP and ARP over HIPPI-6400 (GSN)

An electronic copy of this document is still needed for posting on the web site.

During the April, 1999 meeting, Jean-Michel was present to defend his current HARP approach and successfully defended keeping the documents unchanged; at least until they moved forward in the IETF process. He noted that the offending "extra" packet formats disliked by some (requiring HIPPI-6400 hosts to know HIPPI-800 formats) could easily be removed per IETF rules at a later date, but that adding the format back would be difficult if it were really needed.

Nothing further at this meeting.

10.4 HIPPI end-point MIB

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

10.5 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.

10.6 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.

10.7 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 Plenary week, June 7-11, Minneapolis, MN

The June Plenary week location is in Minneapolis, MN. The hotel information has been linked to the www.hippi.org web site. The tentative schedule is as follows:

Tuesday 6/8 9 AM - 6 PM HIPPI ad hoc

Tuesday 6/8 6 PM - 9 PM HIPPI-6400-OPT

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

Wednesday 6/9 6 PM - 8 PM T11.1 Plenary

11.2 Interim meeting, July 13-14 in Seattle, WA

The next interim working meeting will be hosted by Sean Rieb and Tera Computer in Seattle, Washington. The meeting information is linked to the www.hippi.org web site. The interim meeting schedule will be:

Tuesday 7/13 1 PM - 9 PM HIPPI ad hoc

Wednesday 7/14 9 AM - 5 PM HIPPI ad hoc

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.

The group continues to feel that interim meetings after July will not be required.

1999

Aug 2-6 Plenary Rochester, 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 provide the IETF RFC number for "ARP and IP broadcast over HIPPI-800 (GSN)" to Don Tolmie and Roger Ronald.
3. Marck Doppke to process the HIPPI-800 switch MIB through the IETF.
4. Greg Chesson to draft an IEEE tutorial on HIPPI-6400 ULA usage and the ULAs special to HIPPI-6400.
5. Greg Chesson to collect information for an "ST FAQ" html document.
6. Greg Chesson to draft a top-level structure for the ST-API document.
7. Don Tolmie to make the edits to the Scheduled Transfer document (rev 3.0) agreed to at this May, 1999 meeting.

8. Steph Bailey to update the SCSI on ST text with the changes agreed to at this May, 1999 meeting.
9. Eric Salo to update the level 0 ST API text with the changes agreed to at this May 1999 meeting.
10. Ed Cady to provide the committee with information on the PCI version of the FCI-Berg Micropax connector.
11. Eric Salo to specify field sizes in the ST API in accordance with the approach selected at the May, 1999 meeting.
12. The Chair (Roger Ronald) to summarize the VC2 size issue on the HIPPI reflector and ask for comments and/or proposals for change in this area.
13. The Chair (Roger Ronald) to contact Berg and SGI to collect status on the cable and connector problems (EMI/RFI and PCI form factor compliance).
14. The Chair (Roger Ronald) to prepare a project proposal for the SCSI on ST Mapping, including a proposal for liaison with ANSI NCITS T10.
15. Arie Van Praag to investigate why exactly eight bytes of zeros are required in the ST on DIX Ethernet mapping versus six bytes of zeros or a length field.
16. Ian Philp to contact Eric Salo by email to discuss the open issues in the ST API.
17. Don Tolmie to add text on VC usage recommendations to the HIPPI-6400-PH amendment.
18. The Chair (Roger Ronald) to produce the minutes for the May, 1999 meeting.

13.0 Adjournment

Motion to adjourn by Jeff Young.

Seconded by Don Woelz.

The meeting was adjourned around 11 AM on the 14th of May. There were no objections.

Attendance

Name	Status	Company	Phone	E_mail
Arie Van Praag	O	CERN	41-22-7675034	a.van.praag@cern.ch
Stephen Bailey	O	Genroco	414-644-8700	steph@genroco.com
Carl Pick	P	Genroco Inc.	414-644-2500	carl@genroco.com
Larry Beine	O	Genroco, Inc.	414-644-8700	larry@genroco.com
Brian Breuer	O	Genroco, Inc.	414-644-8700	brian@genroco.com
Chris Good	O	Genroco, Inc.	414-644-8700	chris@genroco.com
Gary James	O	Genroco, Inc.	414-644-8700	gary@genroco.com
Donald Woelz	A	Genroco, Inc.	414-644-2505	don@genroco.com
Tan Truong	P	IBM	512-838-7813	truongtm@us.ibm.com
Ian Philp	P	Los Alamos National Lab	505-667-4305	philp@lanl.gov
Roger Ronald	P	Power Micro Research	972-437-9461	rronald@pmr.com
Eric Salo	O	SGI	651-683-5607	salo@sgi.com
Jeff Young	A	SGI/Cray Research	651-683-5536	jsy@cray.com
Sean M Rieb	P	Tera Computer Co.	206-701-2164	srieb@tera.com
Don Tolmie	O	Tolmie Enterprises	505-662-7110	tolmie@concentric.net