IEEE 1394.1 Bridge Standards Meeting

AGENDA for July 15-16, 1998
Bath, United Kingdom

1. **Administrative:**
   1.1. Minutes of June 9-10, 1998 meeting
   1.2. Vote results
   1.3. Web Page

2. **Technical topics**
   2.1. Virtual Node Ids – David Wooten
   ---- Added during meeting ----
   2.2. Standards Closing Action Table (SCAT) - Morrow
   2.3. Forwarding speeds of asynchronous packets - Morrow
   2.4. Time synchronization – Mike Teener (email)
   2.5. Isochronous time stamp adjustments – David Wooten
   2.6. Device discovery across bridges
   2.7. Time of death for packets held by bridges – Jerry Hauck
   2.8. Fairness rules for bridges

3. **Assignment of Action Items**

Chair:	Dick Scheel	dicks@lsi.sel.sony.com
Editor:	Peter Johansson	pjohansson@aol.com
Secretary:	Subrata Banerjee	sub@philabs.research.philips.com

Minutes of the July 15-16, 1998 Meeting

Administrative

1.1. Hard copies of minutes from the previous meeting (June 9-10, 1998) were made available. Peter noted that not all aspects of his presentation were included in the minutes. He has provided a write-up describing the highlights of his presentation. This write-up is included below. Minutes of June 9-10 meeting will be modified to include this write-up. Dave Wooten moved that the modified minutes be approved. Peter Johansson seconded the motion. The motion passed unanimously.

"Peter Johansson led a discussion on some aspects of virtual node IDs; the slides are available as BR025r00 on the FTP site. The topics included a recap of how bridges would map a virtual (soft) node ID to the 16-bit physical ID used on a local bus. The open issues are:

a) do virtual node IDs work if request(s) and response(s) follow different paths through bridges;
b) how is synchronization of virtual node IDs performed if they cannot vary according to route;
c) how virtual node IDs interact with bus resets (i.e., what are the responsibilities of the bridge portal that observed the reset);
d) the life cycle of a virtual node ID (expiration of stale IDs); and
e) are directory services necessary from a central location."

1.2. Vote Results:
About two weeks ago Dick Scheel had called for votes on the following two topics:
(a) May multiple bridges actively connect a pair of bridges?
(b) Should a request and its response follow the same path?
Last day to vote (via email) was Friday, July 10, 1998. Dick announced the results as follows:
Eligible voters: 23
(a) Yes: 9, No: 6
(b) Yes: 15, No: 0

Peter Johansson (PJ) suggested that before having a ballot we need to have a document on the topic of vote. Dave Wooten felt that there weren't enough details to make a decision. Several people were not sure what they were voting on and they didn't vote.

PJ's motions: (a) Require posting of a document to the ftp/web site two weeks before the vote and let the ballot be open for two weeks. (b) Discard the results of the last week's vote. Dave Wooten seconded. Passed 10 for 0 against. Hence, the vote results were declared void. Note: The person proposing a vote will have to provide a relevant technical document.

1.3. Web page: IEEE Standards Organization (SA) has allocated web space for P1394.1 work group. IEEE1394.1 related contents of the ftp site at ftp.symbios.com will copied to this web space. Once the web page is set-up its URL will be announced on the reflector. Subrata Banerjee noted that IEEE SA wants drafts to be posted on the web in a password protected area. PJ and others would prefer the drafts to be posted in a *non*-password protected area. P1394.1 chair Dick Scheel (DS) will write a letter to IEEE-SA regarding this issue.

- All emails regarding "loops" over reflector during the past two weeks will be captured as a file and posted on the ftp/web site. Author of any relevant email objecting to this should contact DS. No one objected to this in the meeting.

- Attendees felt that current draft version 0.03 is too outdated. Someone not attending the meetings regularly can be misled to think that the current draft represents our best effort to-date. PJ commented that producing an intermediate draft for the benefit of these people is not worthwhile. Brad suggested having a road-map of our standardization effort and to regularly update the draft. PJ suggested organizing all the documents produced/presented at the meetings in order to help others to track our progress.

Technical Topics

2.2 SCAT (Standard Closing Action Table)
Presenter: *Neil Morrow*
Neil presented a list of eight open items that were brought up at the previous meeting.
1. Multiple bridges may actively connect a pair of buses (configured loop topology).
2. Asynchronous response & request packets shall follow the same path.
3. Sub-net architecture
4. Setting-up of the routing table (vs. routing bounds as in draft v0.3)
5. Virtual Node Ids
6. Clock synchronization via go-fast, go-slow mechanism
7. Quarantine bit.

Items Added:
9. Device/Node discovery mechanism

2.3 Forwarding speeds of asynchronous packets
Presenter: *Neil Morrow*
Neil made a presentation on forwarding asynchronous packets with speed being a function of the PRI field.

DS noted that PRI field is used in backplanes. PJ noted that there are two parts of the problem. The first part is to *discover* that there is a remote device to which a node wants to talk to, and the second part is to figure out at what *speed* they should communicate. There was some discussion on how to figure out a forwarding speed. It was noted that initial speed (from the source to its local bus) need not be the same as the final speed (from the local bus to the destination), or the intermediate speeds (from bridge to bridge). One obvious way to discover forwarding speed is to try various speeds - missing responses can be used to find out which speeds don't work. It was noted that SBP2 provides a source node with 16bit destination node address, speed, and max_rec to talk to another node. IEC61883 defines three roles: Talker, Listener, and Agent. Something similar to an agent may be used to discover the necessary communications parameters for the talker and the listener. The group felt that a device discovery mechanism needs to be studied first before addressing this issue.

2.4 Discussion Topic: **Mike Teener's Email to reflector on Async. Clock Synchronization** (Dated 13 July'98)
David Wooten described the proposal (Document Number BR030r00).
Clock adjustment information has to be delivered in time. Asynchronous write may not work due to the high latency involved. Anything other than ±1 tick adjustment requires averaging and hence extra hardware. Both solutions proposed by Mike Teener (MT) will violate the letters of 1394-1995 standard but they may still be compatible with the existing (and planned) 1394 capable devices. PJ questioned if we know for sure that go-fast/go-slow commands via asynchronous transactions won't work. Possibility of FM modulation problem was mentioned. *It appears that the committee has (unofficially) ruled out asynchronous arbitration for `add a tick` and `subtract a tick` mechanism for clock synchronization.* PJ suggested using a Tcode to identify 'add a tick' and 'subtract a tick' commands. DS encouraged people to give this issue some serious thought and he invited discussions on this topic on the reflector. PJ will write a concrete proposal on this topic so that people not attending this meeting can participate in the discussions.

2.5 Talk Title: **Isochronous Time-Stamp Adjustments**
Presenter: ***David Wooten***
Three methods were proposed for adjusting absolute timestamps contained in a packet while it passes through a bridge.
1. *Constant Offset*

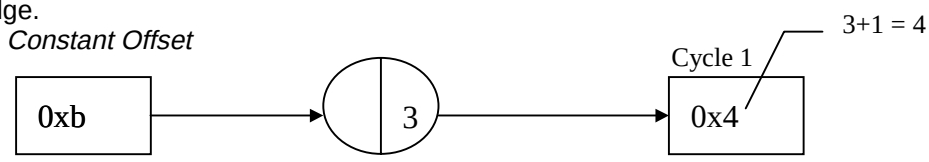

2. Variable Offset within Packet

In this approach the timestamps are modified as follows:

New timestamp = (<timestamp in packet> - <incoming bus cycle number> + <outgoing bus cycle number>) $_{\text{mod <mod value>}}$ = 0x8 – 0x2 + 0x6 = 0xC

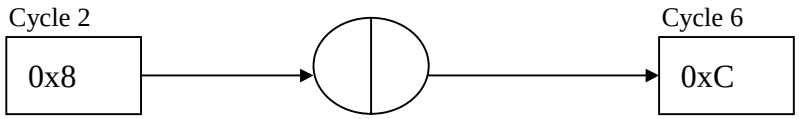

Only the *cycle time* field needs to be adjusted (and not of the *offset* field). In this approach the offset value becomes absolute with respect to a future cycle when the packet will be consumed. Takashi Sato (TS) said that offset field should be adjusted as well. PJ noted that Rudi Bloks did a careful review of this issue in IEC61883 and SBP2 and concluded that offset value need *not* be modified.

3. Fixed offset per channel: Absolute timestamps in channel *k* packets going through a bridge will be modified as follows: new timestamp = constant for channel *k* + cycle number of the outgoing bus. This approach is useful if different channels need different offsets for some reason (none of those present could think of a reason why this might be needed).

Committee preferred: "***Variable Offset within Packet***" method of adjusting timestamps. PJ noted that equivalent text can be borrowed from another document (SBP-2) for our draft specification. Note that both the source packet header timestamp and syt timestamp in CIP header may need to be adjusted, depending on values of certain CIP header fields. We may need *C*-code to elaborate this operation. Current documentation in the draft is not enough.

2.6 Device Discovery (John Fuller): After initialization (including in bridges) how can an ordinary node find other nodes in the network?

Possible methods considered:

1. Brute force (exhaustive enumeration)
2. Central Directory – requires Bridge Manager or some other server to hold the directory
3. Distributed Directory
4. Brute force light (David James): Distributed presence directory, bridge portals contain directory of globally visible nodes on their attached bus.

- John proposed that *pros and cons of centralized vs. distributed directory services methods be discussed on the reflector.*

Devices that do not want to be accessed from outside the local bus should not be listed in the global directory.

Directory Entries Considered:
Virtual ID (10 bits)
EUI-64
Bus Info. Block
Function Class of objects: *E.g.,* printers, disk drives, *etc.*
Speed Information

Open Issue: How are directory entries filled in and maintained?

TS proposed that packet size be limited to 512 bytes to solve some of the speed related problems. Dave Wooten encouraged discussion on the reflector on *pros and cons of limiting size of all packets going through a bridge to 512 bytes.*

TS made a presentation on **Isochronous Time Stamp Adjustment.** TS agreed that adjustments of *cycle* count field only (and not the offset field as well) seems sufficient

if (Δ between time-stamp and local cycle count when the packet is transmitted)
DOES NOT change throughout the life of the packet.

Discussion on Virtual Node Ids:
A list of virtual node id properties known to the WG so far was created (see Attachment 1).

2.7 Time of Death for packets held by bridges: (see Attachment 2)

Jerry Hauck led a discussion of "Time of Death" for packets in transit through a bridged network.

If a packet is encountering 'too much' end-to-end transfer delay then the packet should be discarded. It was noted that in 1394a, a t-label can be reused after an *extra* split-time-out amount of time. If the remote requester is bridge-aware then it will wait for twice the split-timeout period before reusing a t-label. PJ asked if we need a *remote* time-out value for remote requests. Note that even when the network works fine, the software process at the responder can be slow and it can consume a significant part of the split-time-out time thereby delaying the response beyond the total allowed split-time-out value. If a requester times-out waiting for a response then that requester cannot generate another request with the same t-label for certain a guard band time-out period. This is required in case the response comes back late. Neil suggested that the first bridge on the path of a request keep track of the t-labels used and do not release that t-label until the response to any outstanding request using the same t-label comes back. New requests with same t-label will get quarantined. The request will be held up until the t-label is released or until the remote_timeout timer expires. (Bridge returns ack_conflict error if a request comes in with a t-label that is already in use.)

There was a suggestion to return an ack_complete packet (or a response packet on behalf of the destination node) for any AV/C write transaction request. The request may then be completed sometime later based on resource availability. WG thought that this violates the meaning of "ack_complete."

Use of time-of-death field to determine when to discard the request/response was considered. This will require timing the forwarding delay inside a bridge. It can be implemented in a bridge by tagging extra information to incoming packets and stripping it before the packet is transmitted to the outbound bus. DS suggested using an extra quadlet added to each packet to carry inter-bridge control information. John Fuller agreed that a 'wrapper' would simplify several tasks.

It was noted that *pre*-1394.1 devices that should work over bridges include isochronous source and isochronous sink. However, pre-1394.1 asynchronous requesters may not work over bridges. Note that all AV/C devices, even those that are only the target of AV/C commands, are async requesters for status information. The AV/C reply is sent with an async write transaction back to the AV/C "commander". Jerry suggested the idea of proxy device for AV/C devices. DS wondered if a general proxy service for non 1394.1 compliant devices can be defined. PJ noted that HAVi group is also thinking of such a proxy service.

It was noted that to work in a bridged 1394 network environment, devices that are not P1394.1 aware should be single-threaded. Two timers may be maintained (by the portal that first accepts the request) for each such device:

T1 – response timeout = local bus split timeout
T2 – requester quarantine timeout = TBD.

A node will be released from Quarantine state if a response comes back before T1 expires or if timer T2 expires. Any response packet received at the bridge *after* T1 expires and *before* T2 expires will be discarded.

|---------- T1 ----------|----------T2----------|

Multithreaded P1394.1 aware devices should be aware of the quarantine interval.

2.8 Fairness Rules for Bridges
The WG noted that the following issues regarding fairness policy of bridges need to be discussed.

1) Do we mandate dual-phase retries for bridges?
2) Should bridges' unfairness be limited (probably by priority budget)?
3) Should bridges use 1394a concept of allowing the responses to be unfair?
4) Can requests pass requests on bridges? (need to be careful about compromising forward progress)
5) Can responses pass responses?

Jerry suggested having longer meetings. Next meeting in Portland will be 2 full days instead of 1.5 days.

Meeting was adjourned at 12:05pm.

Attendees of July 15-16, 1998 IEEE P1394.1 Bridge Working Group Meeting

Name	Company	Phone	Email
1. Abraham Katz	Rafael	+972-4-9980418	Avikatz@rafael.co.il
2. Alistair Coles	HP Labs	+44-117-922-8750	Anc@hplb.hpl.hp.com
3. Angel Mora	Carinthian Tech	+43-4242-56300-210	Angel.mora@ctr.ac.at
4. Bill Prouty	HP	+1-916-785-4631	Bprouty@hp.com
5. Brad Saunders	Rockwell	+1-949-221-6513	Bradley.saunders@rss.rockwell.com
6. Charles Bill	AMP	+1-717-810-4642	Cebrill@amp.com
7. Daisuke Hiraoka	Sony	+81-3-5448-4603	Dai@sm.sony.co.jp
8. David James	Sony	+1-408-955-6295	Davej@lsi.sel.sony.com
9. David Wooten	Compaq	+1-281-518-7231	David.Wooten@compaq.com
10. Dick Scheel	Sony	+1-408-982-5834	dicks@lsi.sel.sony.com
11. Eric Deliot	HP Labs	+44-117-922-9539	Jhwon@secns.sec.samsung.cp.kr
12. Hans Bjorklund	ABB Power Systems	+46-240-782563	Hans.bjorklund@sepow.mail.abb.com
13. Jerry Hauck	Macro Designs	+1-510-668-1006	jerry@macrodesign.com
14. Johann Zipperer	Texas Instruments	+49-8161-804061	j-zipperer@ti.com
15. John Fuller	Microsoft	+1-425-703-3863	jfuller@microsoft.com
16. John Smolka	Texas Instruments	+1-972-480-1049	jsmolka@ti.com
17. Jun-Ichi Matsuda	NEC Corporation	+81-44-856-2082	Matsuda@ccm.cl.nec.co.jp
18. Legallais Yuon	Thomson Multimedia	+33-0299-273798	Legallaisy@lhmulti.com
19. Masatoshi Ueno	Sony	+81-3-5448-5615	mike@zayante.com
20. Matthew Baker	Philips Research	+44-1293-815287	Bakermp@prl.research.philips.com
21. Neil Morrow	Texas Instruments	+1-972-480-6211	Nmarrow@ti.com
22. Patrick Yu	NEC	+1-408-588-5436	Patrick_yu@el.nec.com
23. Peter Johansson	Congruent Software	+1-510-531-5472	pjohansson@aol.com
24. Philippe Boucachard	Canon CRF	+33-2-99-87-68-73	boucachard@crf.canon.fr
25. Steve Bard	Intel	+1-503-264-2923	Steve.bard@intel.com
26. Subrata Banerjee	Philips Research	+1-914-945-6129	sub@philabs.research.philips.com
27. Takashi Sato	Philips Research	+1-914-945-6099	txs@philabs.research.philips.com
28. Tomoki Saito	NEC	+81-44-856-2082	New@ccm.cl.nec.co.jp
29. Yun Wang	Alps Electric USA	+1-408-432-6413	Yun.wang@alps.com

ATTACHMENT 1 (to Minutes of July 15-16, 1998 Meeting)

This document was created by the WG during the July 15-16, 1998 meeting.

Virtual node Ids

1) Bridge portals [somehow] automatically assign ID to each device on local bus that "wants" one.
2) Portal has a map of virtual ID to PHY_ID. This must be updated on each bus reset (either by tracking topology or reading EUI-64s).
3) Virtual IDs stored in a requester time out after t3 time since last access from that requester. (How does the requester know if it is a physical ID or a virtual ID? Answer: The bus_id is <> the requester's own bus_id.)
4) Virtual IDs no longer used (because portal sees the device is no longer present) shall not be returned to the free pool until time t4 since the last access. t4 > t3.
5) Virtual IDs on a bus must be the same as seen through all bridges attached to the bus.
6) Assume that mapping is done at the input of the outbound queue, and at the output of the inbound queue.
7) On bus reset, attached portals:
   a) Freeze inbound asynchronous queues (packets addressed to that bus) (may discard, but not required)
   b) Reexamine the topology of the reset bus (by means of the self-id packets).
      i) Disconnected nodes – place the virtual node-id (if it had existed) into "limbo" (don't re-use until time-out t4 expires, reject inbound packets with some error code)
      ii) Newly connected nodes – assign them a new virtual node-id (as soon as one is available)
   c) If a new bridge is connected, synchronize virtual IDs
   d) Thaw the inbound queues
8) (not specific to virtual node-ids: Assuming t-labels are re-used by requester immediately subsequent to a bus reset, then the portals on the reset bus must quarantine the requester on the reset bus (for twice the split time-out) to prevent sending new requests and acceptance of a response from another bus that duplicates a t-label.)
9) The portal immediately adjacent to the packet originator virtualizes the source-id.
10) The portal immediately adjacent to the packet destination un-virtualizes the destination-id.
11) TBD: Evaluate all cases of connecting two bus segments for correctness.
12) TBD: Method to establish & synchronize virtual_ids (initialization).

ATTACHMENT 2 (to Minutes of July 15-16, 1998 Meeting)

This document was created by the WG during the July 15-16, 1998 meeting.

A. Split time-out protection for non-P1394.1-aware devices

1) Portal has effectively one timer per local bus device
2) Portal knows if each local device is P1394.1-aware (from scan of configuration ROM)
3) For each non-.1-aware device, the portal will perform the following sequence:
   a) the portal sees an outbound off-bus request
   b) the portal starts the timer for the device
   c) at this point we are in "outbound quarantine" mode
      i) "outbound quarantine" mode ends when either a response is received, or the timer reaches the outbound quarantine value (T1)
      ii) during outbound quarantine, the device is blocked from sending new remote requests (by replying with ack_pending)
   d) if a response arrives before the local bus split time-out, it is forwarded to the device
   e) if a response arrives after the local bus split time-out and before the quarantine mode ends (T1), it is discarded

---

B. How do P1394.1-aware devices handle asynchronous requests?

1. Add a remote split transaction time-out CSR (in addition to local split transaction time-out)
2. Responder always must respond within time specified by local split transaction time-out CSR
3. Requester cannot reuse a t-label until remote time-out expires, or the response is received

How do portals handle P1394.1-aware devices' transactions?
Option 1:

1. A "fixed" interval within which a bridge must forward a packet (drop the packet if it cannot be forwarded in this interval) ("fixed" may be configurable)
2. forwarded (as used above) means bridge gets ack code other than ack_busy*, or ack_missing
3. (look in the future at all ack codes in detail)
4. yada-yada-yada

Option 2:

1. <use IP>

Option 3:

1. Put time of death in pri field (either 3a time or 3b hop count)
2. Mandate common time in all buses if 3a
3. Limits backplane bus to leaf

Option 4:

1. Some non-IP wrapper
2. Add a quadlet at the front of the data part of the packet, in path from first bridge to last bridge (added in the first bridge, stripped in the last)
3. Store time to death, decrement as packet is forwarded
4. Cheat on limit on async packet length (for example, allow 516 bytes with wrapper at S100)

Options 1 & 4 had some support. The others had at most one person who "liked" them.
