P1394A Working Group Meeting Minutes
2/5/97 - 2/6/97

The meeting was hosted by John Fuller of Microsoft at their offices in Bothell, Washington. Chair Peter Johansson called the meeting to order at 9:00 am and set forth a full agenda that called for a continuation of Bill Duckwall's PHY enhancement presentation from the last meeting followed by a discussion on power management by Dave Wooten. This was followed by a detailed review of the latest P13994a draft standard. The review was continued the next day followed by a discussion on some additional PHY enhancements, a discussion on a compilation of corrections/clarifications presented by Steve Finch, P1394a features requested by OpenHCI and some additional miscellaneous topics.

PHY ENHANCEMENTS:

ACCELERATED ARBITRATION: Bill Duckwall started off by reviewing the accelerated arbitration features. He noted that link will now have to keep track of the arbitration fairness status for concatenating fair requests/responses.

MULTI-SPEED CONCATENATED PACKETS: One of the action items from the last meeting was to determine how the current PHYs will handle the proposed multi-speed concatenated packets. The main question was if they transmit speed signals for the concatenated packets and how they treat the absence of speed signals on received concatenated packets. Three PHY vendors provided details of operation of their existing devices. The existing PHYs were found to have the following flavors:

1. Transmitting PHY sends speed signal for each (same speed) concatenated packet.
2. Transmitting PHY does not send speed signal for each (same speed) concatenated packets.
3. Receiving PHY assumes that a concatenated packet received with no associated speed signal will be interpreted as a 100mb/s packet.
4. Receiving PHY assumes that a concatenated packet received with no associated speed signal will be interpreted as having the same speed as the first packet in the concatenated sequence.

The interoperability problem occurs when a flavor 2 PHY transmits 200mb/s (or higher) concatenated packets to a flavor 3 PHY. The transmitting PHY does not transmit the

speed signal for concatenated packets and the receiving PHY will interpret these as 100mb/s packets.

For the P1394a PHYs two options were suggested:

Option A:

1. Transmitting PHYs may concatenate multi-speed packets with one exception: they may not concatenate 100mb/s packets after a 200mb/s or higher speed packets. (100mb/s packets may only be concatenated after a 100mb/s packet.
2. Transmitting PHYs must send a speed signal before every concatenated packet.
3. If a receiving PHY sees a concatenated packet with no associated speed signal, (as would happen with some existing PHY's) it should interpret the packet's speed to be the same as the previous packet's speed.

Option B:

1. Transmitting PHY's may concatenate multi-speed packets with one exception: they may not concatenate 100mb/s packets after a 200mb/s or higher speed packets. (100mb/s packets may only be concatenated after a 100mb/s packet.
2. Transmitting PHY's must send a speed signal before every concatenated packet.
3. If a receiving PHY sees a concatenated packet without an associated speed signal, (as would happen with some existing PHY's) it should receive the packet at 100mb/s.

A straw poll was taken on these two options with the following results: Option A: 10 Option B: 0

Peter Johansson took the action item to generate a brief description of the PHY interoperability problem and post it on the TA web pages.

TOKEN STYLE ARBITRATION: Bill explained how this arbitration scheme works and how it benefits an application like a multimedia server. The current arbitration scheme grants control to the root first and then down the chain all the way to the leaf nodes. However, in order to take full advantage of the proposed accelerated arbitration techniques, there needs to be a way of ensuring that during the isochronous cycle, control is granted in sequence to the leaf node first and then on up the chain with the root node being the last. One of the conclusion reached was that if token style arbitration is in use, a device that does not understand token style arbitration will not see any difference in bus behavior. As such, given that this scheme will most likely be used in a closed environment (such as a

media server) should this scheme be included in P1394a supplement? If so, should it be informative or normative? A straw poll was taken as follows:

1. Include in the P1394a document and perform the work in 'a' working group? yes 11 no 1
2. Normative if it does not fall behind schedule ? yes 10 no 0

PER PORT SOFTWARE DISCONNECT: Bill explained that on a software 're-connect' the controlling PHY turns on the Tp bias and the 're-connected' PHY should follow the same debounce mechanism as specified for the reset detect. The main issues here were centered around whether a reset should be generated after port connect or disconnect. The conclusion was:

1. After a software disconnect no reset should be generated since the controlling node could always generate a bus reset if it needed one after the disconnect.
2. A software 're-connect' should be treated as a new connection. If Tp bias is detected, a reset should be generated. (The re-connected node will generate a reset as soon as it sees the Tp bias on its ports)

Also, Bill wants to require some sort of visual indication that the software has disconnected a port. It was concluded that this will be hard to describe as a requirement but could be included as an implementer's note.

SLEEP MODE: There was considerable discussion on how the proposed "per port software disconnect" feature ties in with the sleep mode. Dave Wooten conducted a discussion on power management and sleep mode. He described the process of putting leaf nodes to 'sleep' as follows: (Note: there are several open issues on how to put busses to sleep)

1. leaf in minimum PM mode
2. disable port @ parent (turn off Tp bias @ parent)
3. leaf sees loss of bias on TpB but "knows" that it is still connected. (don't know how the node "knows" as yet! work in progress)
4. leaf turns off Tp bias to acknowledge sleep
5. if parent sees Tp bias from leaf after N ms, this is a wake event from leaf.
6. generation of bus resets on Tp bias change is TBD

There was quite a bit of discussion on whether the mechanism of how the PHY "knows" about the sleep mode. Dave mentioned that the whole issue of power management is being discussed in a different forum.

INCREMENTAL BUS RECONFIGURATION: Bill continued the presentation on his proposed PHY enhancements with an explanation of the incremental bus reconfiguration feature. John Fuller indicated that this feature was proposed before the short arbitrated resets and short resets solve almost all the problems that incremental topology re-configuration solve. Richard from Compaq suggested that before discarding this issue we should try and come up with viable scenarios that may be benefited by this feature. The conclusion was that we will defer the discussion to a future meeting.

1394A DRAFT SPEC. REVIEW:
After Bill's presentation Peter Johansson started the review of the latest draft of P1394a spec.

PHY REGISTER MAP: Proposed PHY register map was reviewed with the following conclusions/discussions.

1. Bill Duckwall suggested adding bits to indicate maximum speed capability on a particular port. Consensus was that these can be added. John suggested a per port speed limit as well. This will be useful when the PHY is capable of a higher speed but the cable is not able to handle that speed. Peter suggested that we need to look at the impact of this on the rest of the architecture and this was deferred to a future discussion.

2. John Fuller suggested that the field marked "medium" should be returned to "reserved". There were no objections to this.
3. 'Multi' bit was removed.
4. The Enab_Token bit was kept for now until further discussion.
5. There was some discussion on the link power status and link enable bit. What is the power-up value for the 'L' bit? When LPS is off the PHY tri-states its outputs on the PHY-link interface. The consensus was that L bit will LPS && L.

PHY-LINK INTERFACE
John Fuller proposed a mechanism to reset the PHY-link interface if it gets caught in an unknown state for some reason. One suggestion made was to use a new LREQ request to send such a reset. But what if the request logic is also 'hosed'? Should the PHY-link interface also generate an automatic bus reset? The consensus looks like no. Another proposal was to use LPS (Link power on) to reset the PHY-link interface. When LPS is off the PHY link interface gets reset. However when LPS is off, all PHY-link signals (including the clock) coming from the PHY are tri-stated.

Will this work? This needs some more discussion in future meetings.

BACKPLANE PHY: Peter asked if anyone present was interested in discussing the backplane PHY and no one seemed to be so this issue will be discussed with a wider audience. (on the reflector)

SPEED BITS IN SELF-ID PACKETS: Dave Wooten was concerned about using the old ‘del’ and ‘sp’ bits to create the new 4 bit speed field. There were a lot of questions on this field. This issue seemed to be more complex than was initially apparent and could be deferred to the high-speed PHY task group. Another proposal for the speed bits was to have a "caboose" packet at the end of the self-id packet with its ‘n’ field set to 3. The group moved to defer the issue for now and have two write ups on this issue. The current proposal + the caboose proposal. John Fuller took the action item for sending the ‘caboose’ proposal on the reflectors.

RESET STATE MACHINE: A discussion on reset state machine and C code followed. Minor corrections were made and the draft will be updated.

ARBITRATION STATE MACHINE REVIEW: Arbitration priority between a root and its children may be significant for a token style arbitration. Peter asked if we should discuss the details of the state machine in a meeting like this or on the reflector. Consensus was to discuss this on the reflector.

The second day started off with more discussion on PHY enhancements.

SELF ID COUNTER: John Fuller proposed the following rule for the self ID counter:

1. Self ID counter never increments past 63.
2. If the id is 63 the PHY
   - should not respond to link-on.
   - should not respond to any requests or PHY packets.
   - should not generate any requests.

Concern here was that root will end up with a PHY id of 63 and thus will be in-capable of solving this problem. Dave Wooten proposed that to prevent the bus from hanging if you have more than 63 devices then have the root use physical ID 0 by going first in the self-ID process.

SLEEP MODE: This topic had been discussed the first day and since there were a lot of questions Dave Wooten took the action item to write up a paper on this issue and make it available to those interested.

DUAL PHASE RETRY SCHEME: There seems to be a lot confusion on how this works. No one present at the meeting knew this well enough that they could explain it to the group. The discussion was deferred to a later meeting when we could have more opinions.

CORRECTIONS/CLARIFICATIONS:
Steve Finch presented a compilation of correction/clarifications for the spec.

Isolation: Annex A in IEEE 1394-1995 is normative and requires isolation. Conclusion: We will replace Annex A with a new section in P1394a. Peter took the action item on starting this new section but he needs help from experts in this area.

Power Sourcing on the cable:
Max DC Current: Section 4.2.2.7 changed maximum DC current to 50mA.

Straw Poll on Maximum Output Voltage: Should this limit be changed to 33V from the current 40V: yes 9 no 0.

Straw Poll on Minimum Output Voltage: Should this limit be changed to 20V from the current 8V. Yes 7 No 2. (Don Tonn voted NO due to insufficient information on the subject. Second NO vote was concerned about the mobile PC platforms.) A lengthy discussion followed regarding the reasons for the power minimum/maximum voltage values. It was observed that a lot of people do not know the assumptions behind some of the power specifications and it would be a lot easier if more information was available. A task group (IEEE) was proposed to resolve the issues/clarification on power distribution topic.

There was also some discussion on whether a 1394a compliant PHY self id indication also means that the new power rules will also be complied with? No apparent consensus was reached.

Separate packet queues for requests, responses and isochronous data: Peter suggested that perhaps the 1394 TA should put out an educational white paper on the reasoning behind these rules. Peter took an action item to clarify the retry state protocol in a new section in P1394a.

Dribble bits: No major clarifications required. The C code may have some typographical errors that will be fixed.

Resynch buffer/Transmitter LINK under-run: No corrections required. Seems to be defined properly in the standard.

Pseudo C code errors: Will be fixed in the new draft.

When should speed signal be sampled: Jim Skidmore took an action item to research/discuss this issue with PHY vendors.

Jitter and skew budget: Eric Hannah mentioned that he has been experimenting with some existing cables and connectors on his test setup and is concerned that it may not work at 400mb/s. Current connector(s) may not be compliant to the spec. There was considerable discussion on this topic and in the end this issue was deferred to the next meeting.

Overshoot restrictions: Deferred to the next meeting.

Test Specification and Procedures: Are existing test specifications and procedures described well enough? Eric Hannah is working on the test procedures for high-speed PHY's and will have the work completed in a couple of months. We could include these procedures in P1394a. The procedures will be generic to P1394a and any high-speed PHY extensions with different parameters for the two.

PHY-Link Protocol State Diagrams: Many people commented that the stated diagrams in Annex J are not accurate. Peter suggested that the state diagrams in Annex J do not convey useful information and should be removed. A notation will be made saying that these are wrong in the current standard.

Maximum number of hops: IEEE 1394-1995 suggests that the maximum number of hops should be 16. Should we explicitly state this requirement in the new spec? Is this the right thing to specify given that long cable lengths are also being proposed? A long discussion followed. The bus manager could handle this automatically by using PHY pinging. However there is some uncertainty involved with PHY pinging. What about environments where there is no bus manager present? The general consensus was that the gap count should not be changed in an unmanaged environment.

MISCELLANEOUS DISCUSSIONS:

After Steve's presentation Peter continued with the following miscellaneous issues:

INCREMENTAL RE-CONFIGURATION: A straw poll was taken on whether this feature should be removed from the draft in light of the discussion on this topic on the first day: 8 yes 1 no (Richard Churchill voted NO since he needed more time to examine the proposal.)

HOW IS AN ACK PACKAGE RECOGNIZED: How do we determine whether a given package is an ACK package so we can use ACK accelerated arbitration? Consensus was that any packet with exactly 8 bits with the last four bits being the 1's complement of the first four is an ACK packet—the position and timing of the packet relative to other packets is not considered important in the PHY's determination of whether or not it is an ACK packet.

MINIMUM SEPARATION BETWEEN SOFTWARE INITIATED RESETS: Peter took the action item to write up the first draft and circulate it.

PHY LINK INTERFACE: Prashant Kanhere mentioned that there may be an interoperability issue between existing PHY's and links in their handling of the aborted status transfer. He took the action item to poll vendors to find out how existing links respond to an aborted status transfer.

DOCUMENT NAMING STANDARD: Peter suggested a document naming standard PnnnRnnn.pdf (P1394a*NN*.pdf)

OpenHCI REQUESTS FOR 1394A:

John Fuller brought up features that are important for the 1394 OpenHCI. Most of the issues had been resolved during the discussions in this meeting but some remain open:

Per port disable for PHY: Resolved except for sleep mode interaction.
PHY ID stops at 63: Resolved
Connection hysteresis (debounce): Resolved
Arbitrated short resets: Resolved
Cycle too long: If link event indication of cycle too long cycle start packets should be suppressed. Open
3-bit speed codes in SPEED_MAP and TOPOLOGY_MAP: Open awaiting 'caboose' proposal.
Response and fair protocol: Response packets are PRIORITY requests. Do not follow fair protocol. Open

"**Generation" bit in BUS_INFO_BLOCK:** This is really an OS request. OpenHCI Operates without it.

The meeting was adjourned at 2:15pm. Peter mentioned that the working group needs a place to hold the San Jose meeting on March 17th and 18th. If anyone can sponsor this meeting in the San Jose area please contact Peter. The action items from this meeting are summarized below.

ACTION ITEMS:

1. Peter Johansson took the action item to generate a brief description of the PHY interoperability problem and post it to the reflector and other appropriate sites.
2. John Fuller took the action item for sending the 'caboose' proposal on the reflectors.
3. Dave Wooten took the action item to write up a paper on sleep mode and make it available to those interested.
4. Peter Johansson took the action item to replace Annex A with a new section in P1394a.
5. Peter Johansson took an action item to research whether or not the single- and dual-phase retry protocols are documented correctly in IEEE 1394-1995 and (if they are problems) to correct or clarify them in P1394a.
6. Jim Skidmore took an action item to research/discuss the issue of speed signaling and when the speed signal is latched with PHY vendors.
7. Peter Johansson took the action item to circulate a write up on the minimum separation between software initiated resets.
8. Prashant Kanhere took the action item to poll vendors re how the existing links respond to an aborted status transfer.

ATTENDANCE

Kazuyuki Abe	(415) 528-5904	kabe@sj-pceg.ccgw.nec.com
Mike Brown	(602) 554-3713	mike_brown@ccm.ch.intel.com
Richard Churchill	(713) 514-6984	richardc@bangate.compaq.com
Ricardo Dominguez	(817) 774-5386	rvd@ti.com
Bill Duckwall	(408) 461-4902	duck@fireflyinc.com
Lou Fasano	(914) 892-8904	lfasano@vnet.ibm.com
Steve Finch	(714) 573-6808	steve.finch@tus.ssi1.com
John Fuller	(206) 703-3863	jfuller@microsoft.com
Eric Hannah	(408) 765-4441	eric_hannah@ccm.sc.intel.com
Yasumasa Hasegawa	81-22-347-1128	yasumasa@ffm.fujifilm.co.jp
Shinichi Hatae	81-3-3723-2111	shatae@cvp.canon.co.jp
Jerry Hauck	(408) 765-5528	jerry_hauck@ccm.sc.intel.com
Peter Johansson	(510) 531-5472	pjohansson@aol.com

Prashant Kanhere	(408) 725-7515	prashant@pacbell.net
David LaFollette	(408) 765-2587	dlafolle@mipos2.sc.intel.com
Hirokazu Mamezaki	81-22-347-1191	mamezaki@ffm.fujifilm.co.jp
Cyrus Momeni	(510) 624-7061	cyrus@copr.cirrus.com
Karl Nakamura	(408) 433-4516	karln@lsil.com
Bill Prouty	(916) 785-4631	billp@hprnd.rose.hp.com
James Skidmore	(972) 480-2094	jskd@msg.ti.com
Peter Teng	(408) 588-8111	pteng@mail.com
Don Tonn	(408) 957-2110	dontonn@eng.adaptec.com
David Wooten	(281) 518-7231	davidw@bangate.compaq.com