

Proposed Agenda
P1394b Working Group
February 9, 1999

1. Review Agenda
2. Review of December Meeting Minutes
 - 2.1. Fix “Ding” to Doing” [Bard]
3. Procedures
 - 3.1. Voting
 - 3.2. Price/Pricing
 - 3.3. Call for Patents
 - 3.4. Review Action items from December Meeting
4. Meeting Schedule
 - 4.1. March 22 &23, Tempe, AZ
 - 4.2. April – Relying on 1394.1 to find venue
 - 4.3. May NEED INVITATIONS!!!
5. Presentations
 - 5.1. Dave James
6. Task Group Reports (As required)
7. SCAT
8. New Issues
 - 8.1. Off Cycle Meetings
 - 8.2. IP Terms and Conditions
 - 8.3. Beta-Only Support
 - 8.4. How Easy Should it Be to Use Legacy Devices?
 - 8.5. Apple free-b

9. Review of Action Items From this Meeting

10. Adjournment

MEETING MINUTES

Review of agenda and call for agenda additions. David James needs time.
Added review of actions items from December meeting.

Accept previous meeting minutes: Colin Whitby-Strevens move to accept, Jerry Hauck seconded. Passed without objection.

Voting, price/pricing, Call for Patents - all presented in the usual, professional manner by the group chair.

Review of Action Items from December meeting:

AR: Colin accepted an AR from the chair to take this issue to the 1394b reflector and solicit further discussion. 2/9/99 Status: Done. There was no significant reflector traffic resulting. Further discussion is needed on this matter, however, at Colin's suggestion, further discussion will be tabled to a future meeting.

AR: David will develop an action item for the Copperhead task group to go off and do something about this. 2/9/99 Status: Done

AR: David will post the words which make up the action item to the reflector. 2/9/99 Status: Done Copperhead Group will report out their activity on this matter.

AR: David will generate a letter and have it delivered via registered postal service to Apple for the purpose of soliciting their position on any Apple intellectual property which may be contained within the current IEEE P1394b draft specification. 02/09/99 Status: David wrote and sent (unregistered) a letter to Apple. A response was received. The response is compliant with IEEE terms (e.g. "reasonable and nondiscriminatory"). A second "short" letter was sent to Apple. The nature of the short letter was not disclosed. A response for the second letter has NOT yet been received.

AR: Steve Bard to develop the SCAT list into a table and include perspiration and inspiration fields as well as status, dispensation and ownership fields. This table will be posted to the ftp repository. 02/09/99 Status: Steve created a SCAT list. It has not been posted to the ftp site. A soft copy (on diskette) was given to the Chair. A hard copy was not brought to the meeting.

Meeting schedule reviewed.

March hosted by Intel. Everybody MUST send in a registration form – attendance optional.

P1394.1 tasked with obtaining April meeting location and host.

Call for meeting hosts for May. John Fuller indicated Microsoft could host a meeting the second or third week in May. Gerald Marazas was solicited for a host in May. Microsoft could also do it in June. Gerald said North Carolina is nice in June.

David James – 2100 – ShuttleBus p2100 Overview

<ftp://ftp.SCIzzL.com/P2100>

David's presentation does not represent any opinion from Sony.

David James will provide a softcopy of the presentation to the P1394b Secretary for posting to the P1394b ftp site.

Task Group Reports

B-Low Power: Task group believe works is complete with resolution of restoring data to the restoring node. Steve to write white-paper and post before March meeting and, if no objections, submit to Plenary for acceptance followed by inclusion in the draft standard. [Bard]

Upstarts: – no meeting this time and never again (hopefully). Things are being polished in the existing code (through simulations, etc). One issue which did come up today (during lunch) concerns the treatment of a disabled port. As P13194a has, in fact, made a provisional determination as to what should be done with regard to disconnect/disconnect on a disabled port. With an accepted change to the 1394a BRC solution (e.g. disable connect-detect circuit when port is disabled and INT_ENABLE is NOT set to one), a disabled beta port will behave similarly. If INT_ENABLE is set to one, a disabled port will continue to perform connection toning. If INT_ENABLE is set to zero, the beta port will perform no such connect toning. [Colin Whitby-Stevens]

UTP-5 [Colin Whitby-Stevens]: – The main thing that came out of yesterday's discussions, more perfect numbers will be plagiarized from 100Base-T (who stole it from FDDI, etc. etc.).

Home network people are concerned about class-B compliance: We should be able to use the sorts of signal levels others have used – this assumes they were compliant.

PHY-Link [Sean Killeen]: 75-80% Complete. No issues with incorporating feedback from yesterday's meeting. Details in their meeting minutes (attached to the end of these minutes)

Copperheads – [Bill Northey]: Dave Bruncker made a presentation on evaluation of A & B speck with regard to electrical parameters and identified where there are some differences. This will be posted to the web site. The other item of interest that came out of the meeting was a brief discussion of alternative connectors. Low cost is of utmost concern. David Wooten would like to fine tune the requirements for this new connector. Copperheads is waiting for that input.

Beta Only “B” [David Wooten]: David Wooten took the floor to explain that the new connector is to part and parcel to making Beta-B only as attractive as possible. Especially the consumer electronic (CE) folks. The cable and connector is of real interest to the CE folks. It would be very good if the new Beta only connector is a much better solution that the current cable and connector being used by the CE folk. Bottom line: We need to do whatever it takes to move the CE folks to beta only as soon as possible.

Accelerations/B-Port [Mike Teener]: Trying to get the new mode for accelerated arbitration worked out and also the isochronous arbitration into an even/odd versus current/next context.

Trying to get holes punched out (thanks Jerry) so “C-Code” state machines can be published in March.

SCAT [David Wooten]

AR: David Wooten will review the SCAT list created and fill in the assigned/ownership fields (as appropriate) and return it to the secretary some time next week. The secretary will ping the chair for said SCAT list next week.

New Issues [David Wooten]

Off Cycle Meetings

Need more meetings to get the work done. David will be calling off cycle meetings of the task group chairs to further the work faster.

IP Terms and Conditions

David has been in contact with the IEEE regarding the latitude a working group may take with respect to requesting a company to divulge terms and conditions to the working group. In short words: It is not appropriate for a working group to determine what is and is not reasonable terms and conditions if the company has complied with IEEE requirements by providing a letter stating “reasonable and non-discriminatory” terms and conditions.

AR: Eric Anderson (Apple) will provide the secretary with name of an individual at Apple that the members of the P1394b working group can contact to obtain an answer to the question: "What are Apple's terms and conditions for use of Apple's IP in IEEE P1394b". 02/09/99 Status: Larry Lowe, llowe@apple.com; (408) 974-9463

Beta-only Support & How Easy Should it Be to Use Legacy Devices?

Mike Teener stated that it is about three times as complicated to have a Beta capable PHY support 1394a as well when compared to a Beta only PHY (a bi-lingual PHY).

A discussion of what it means to have a Beta PHY function in a network which contains 1394a nodes was discussed. It is, essentially, a "nit" to include "things" in a Beta-only PHY to allow it to support a 1394a PHY some place out there on the network.

A Beta PHY which will make use of either the existing 4-pin or 6-pin copper circuit connectors shall be DS capable (the ports connected to the 4-pin or 6-pin connectors are bi-lingual).

A Beta only PHY shall not use the existing 4-pin or 6-pin copper circuit connector.

The IEEE P1394b PAR states that interoperability with 1394-1995 and 1394a is a requirement. It does not state HOW interoperability shall be accomplished (e.g. plug compatibility is not a requirement in the PAR).

Apple free-b

Firooz Farhoomand read the EE Times article about "Apple IP Free P1394b" to the group. Firooz posed the question to the general body: "What is the opinion of the group regarding this position?"

Firooz is really interested in discovering whether the IEEE P1394b working group is, indeed, taking steps to develop an "Apple IP free" specification.

The question is one of whether plug compatibility and interoperability with "legacy" nodes is desirable or not.

As a general statement, it is NOT appropriate for an IEEE Working Group to select a particular company's intellectual property for removal from a standard specification providing the company has complied with IEEE IP requirements.

IEEE P1394b Working Group Plenary Attendees (February 9, 1999):

Name	Company	Email	Phone
Anderson, Eric	Apple	ewa@apple.com	408-974-8187
Baker, Richard	Solectron, Texas	richard.baker@tx.slr.com	512-425-6205
Bard, Steve	Intel	steve.bard@intel.com	503-264-2923
Barton, Amy	Circuit Assembly	amyb@circuitassembly.com	949-855-7887
Brill, Charles	AMP	cebrill@amp.com -or- cebrill@ix.netcom.com	717-810-4642 -or- 717- 533-1275
Brunker, Dave	Molex	dbrunker@molex.com	630-527-2622
Chander, Navin	Texas Instruments	navim@ti.com	972-480-2430
Coles, Alistair	HP	anc@hplb.hpl.hp.com	+44 117 922 8750
Coletta, Mike	Harris	mcoletta@harris.com	949-707-1143
Dorsey, Chris	ST Microelectronics	christopher.dorsey@st.com	972-466-7850
Farhoomand, Firooz	Panasonic	farhoomandf@panasonic.com	408-653-4059
Fasano, Lou	IBM	fasano@us.ibm.com	914-892-8904
Feller, Scott	Hewlett-Packard	scott_feller@hp.com	(650) 857-7759
Foster, Tony	Hewlett-Packard	tony_foster@hp.com	(916) 785-1092
Fuller, John	Microsoft	jfuller@microsoft.com	425-703-3863
Furuya, Nobuo	NEC	nobuo_furuya@el.nec.com	(408) 969-2479
Gannon, Bob	C&M Corp	r.gannon@cm.corp.com	860-779-4249
Hannah, Eric	Intel	eric.hannah@intel.com	408-765-4441
Hauck, Jerry	Zayante, Inc.	jhauck@zayante.com	510-668-1006
Inoue, Tatsuo	Arch Tech	ted@archtech.co.jp	+81 3 5545 7813
James, David	Sony	davej@lsi.sel.sony.com	650-494-0926
Johnson, David	Texas Instruments	dkjohnson@ti.com	972-480-3632
Kadison, Eric	Harmonic Technologies	e_kadison@earthlink.net	949-470-4725
Killeen, Sean	SSL	sean.killeen@ssl.ie	+353 1 402 5700
Le, Thang	Hewlett-Packard	tl@rose.hp.com	(916) 785-4667
Lohmeyer, John	LSI Logic	lohmeier@lx.netcom.com	719-533-7560
Lopata, John	Molex	jlopata@molex.com	(630) 579-4110
Marazas, Gerald	IBM	marazas@us.ibm.com	919-543-6892
McDonnell, Edward	HP Labs	emcd@hplb.hpl.hp.com	117-922-8942 (UK)
Momeni, Cyrus	Cirrus Logic	cyrus@corp.cirrus.com	
Morgan, Michael	Quantum Corporation	mike.morgan@quantum.com	408-894-4315
Myers, James	Picolight, Inc.	jmyers@picolight.com	303-530-3189 x246
Nelson, Jim	Northrop Grumman	nelsoja5@mail.northgrum.com	562-948-8357
Nordby, Fritz	Silicon Engineering	fritz@sei.com	831-438-5331
Northey, Bill	FCI	northewa@bergelect.com	717-938-2119
Nyu, Takayuki	NEC	new@optsys.cl.nec.co.jp	+81 44 856 2082
Oktay, Ozay		oktay@earthlink.net	949-651-1139
Saito, Kyoza	Alps	kyozo_s@gw3.alps.co.jp	+81 229 23 5111

Name	Company	Email	Phone
Saunders, Brad	Xircom	brad_saunders@xircom.com	805-376-6686
Sessions, D. C.	VLSI	dc.sessions@vlsi.com	602-752-6545
Shergill, Robbie	National Semiconductor	robbie.shergill@nsc.com	408-721-7959
Soderstrom, Ron	IBM	v2rons@us.ibm.com	507-253-6290
Sorna, Michael	IBM	sorna@vnet.ibm.com	914-892-3425
Teener, Michael Johas	Zayante	mike@zayante.com	831-461-4901
Teng, Victoria	NEC	victoria_teng@el.nec.com	408-969-2861
Thatcher, Tom	Hewlett-Packard	thomas@dtc.hp.com	(650) 857-4039
Thompson, David	Lucent	davethompson@lucent.com	610-712-2730
Vrhovnik, David	Gennum Corporation	dave_v@gennum.com	905-632-2999 x3004
Wakai, Hirosha	Sharp	wakai@slab.tnr.sharp.co.jp	650-638-7273
Whitby-Stevens, Colin	Zayante	colin@zayante.com	831-461-4948
Yoshikatsu, Niwa	Sony	niwa@sm.sony.co.jp	+81 3 5448 4603
Young, Len	Corning	younglg@corning.com	607-974-8676
Yu, Patrick	NEC	patrick_yu@el.nec.com	408-588-5436
Zchikawa, Yujz	Sharp	ichikawa@slab.tnr.sharp.co.jp	+81-743-65-4529

IEEE P1394b
B-Low Power Task Group

AGENDA
8th February 1999
Hilton Waterfront Hotel
Huntington Beach, California

1) Introductions

2) Accept volunteer to write minutes

Appreciation to Jerry Hauck for accepting a request to take meeting minutes.

3) Approve Previous Meeting Minutes

Unanimously agreed.

4) Review

4.1 Characteristics of "Standby" [Steve Bard]

5) Discussion

5.1 Expected Response when accessing node in Standby

No expressed opposition to accept ACK_MISSING as the expected response when endeavoring to access a node in Standby. Only nodes aware of the standby state of the leaf would, realistically, be trying to access the leaf; therefore, it would be expected that those nodes would command the parent node to restore the leaf from standby prior to accessing the leaf node.

5.2 Control Code for Standby (Upstarts/B-Port?)

Talked about reusing suspend/resume code as control code to tell parent node that leaf has gone into standby. This, however, requires the parent node to know if the child was a "leaf" or not so that it can get the proper context for the suspend indication (one case requires bus resets, etc. and the other requires proxying of self-id's, etc.).

One "code" should be sufficient to enter standby. The parent PHY stores (for each port) the limited set of bits (12 by the current count) from each self-id sequence - needed to proxy self-id packets for it's children. When a standby code is received from a child, the parent already has the proxy information needed. [Alistair took the action item to see if a control code is available, or whether we should look at a PHY packet from the child to the parent.](#)

The PHY packet for standby doesn't look promising! Fritz pointed out: the leaf

node knows which port is its parent port, but it doesn't know the PHY ID of its parent node, therefore, endeavoring to address a "Standby PHY packet" to its parent node is somewhat of a problem – the "Control code" seems to be the mechanism of choice.

Seeking resolution to the problem of delivering node-ID, gap count and gap count "stick-bit" status to the leaf node on restore, two similar schemes were proposed.

First: When a leaf node is restoring, its parent node arbitrates for the bus and then sends a PHY packet addressed to itself (the parent node). The restoring leaf node decodes the first PHY packet it receives and to grab the relevant information (node-ID, etc.). The parent node must service one and only one restoring child leaf at a time. Since nodes won't be considered active until the parent delivers the PHY packet, there is no danger that other parent nodes in the network which are in the process of restoring child leaf nodes will forward the PHY packet erroneously to one of their restoring leaf nodes.

5.3 Steve Bard took an AR to write-up the standby feature/function and present it to the next Plenary for ratification.

6) Review AR's

Alistair Coles: Investigate and report availability and value of a Control Code for Standby.

Steve Bard: Write up standby feature for delivery to next Plenary.

7) Opens


None

8) Adjourn

Discussion:

The group held a discussion about how a device enters standby on its own. There was a suggestion from Fuller (generally accepted by the group) that a leaf node entering into Standby on its own accord will transmit a PHY command packet (7/1) addressed to itself – allowing other nodes on the bus which may be interested (such as the Power Policy Owner of the leaf node), to observe the command packet. This procedure will appear to the link exactly the same as it would if the Power Manager or the Power Policy Owner commanded the leaf to enter into standby.

When the parent node of a leaf node in Standby goes into suspend, the parent node will issue a bus reset to the leaf node when the leaf node restores from standby as the result of the parent node resuming from suspend. The following drawing was used in the discussion:



When the NPC is in standby and the MPN is in suspend, the MPN will issue a bus reset when it resumes.

Attendance Roster for "B" Low-Power (2/9/99):

Name	Company	Email	Phone
Barton, Amy	Circuit Assembly	amyb@circuitassembly.com	949-855-7887
Brunker, Dave	Molex	dbrunker@molex.com	630-527-2622
Chander, Navin	Texas Instruments	navim@ti.com	972-480-2430
Chen, Dao-Long	LSI Logic	dao-long.chen@lsil.com	970-206-5461
Coles, Alistair	HP	anc@hplb.hpl.hp.com	+44 117 922 8750
Coletta, Mike	Harris	mcoletta@harris.com	949-707-1143
Dorsey, Chris	ST Microelectronics	christopher.dorsey@st.com	972-466-7850
Farhoomand, Firooz	Panasonic	farhoomandf@panasonic.com	408-653-4059
Feller, Scott	Hewlett-Packard	scott_feller@hp.com	(650) 857-7759
Fuller, John	Microsoft	jfuller@microsoft.com	425-703-3863
Furuya, Nobuo	NEC	nobuo_furuya@el.nec.com	408-969-2479
Gannon, Bob	C&M Corp	r.gannon@cm.corp.com	860-779-4249
Hauck, Jerry	Zayante, Inc.	jhauck@zayante.com	510-668-1006
Inoue, Tatsuo	Arch Tech	ted@archtech.co.jp	+81 3 5545 7813
Kadison, Eric	Harmonic Technologies	e_kadison@earthlink.net	949-470-4725
Killeen, Sean	SSL	sean.killeen@ssl.ie	+353 1 402 5700
Lopata, John	Molex	jlopata@molex.com	(630) 579-4110
McDonnell, Edward	HP Labs	emcd@hplb.hpl.hp.com	117-922-8942 (UK)
Morgan, Michael	Quantum Corporation	mike.morgan@quantum.com	408-894-4315
Niwa, Yoshikatsu	Sony	niwa@sm.sony.co.jp	+81 3 5448 4603
Nordby, Fritz	Silicon Engineering	fritz@sei.com	831-438-5331
Northey, Bill	Berg Electronics	northewa@bergelect.com	717-938-2119
Nyu, Takayuki	NEC	new@optsys.cl.nec.co.jp	+81 44 856 2082
Saito, Kyozo	Alps	kyozo_s@gw3.alps.co.jp	+81 229 23 5111
Saunders, Brad	Xircom	brad_saunders@xircom.com	805-376-6686
Sessions, D. C.	VLSI	dc.sessions@vlsi.com	602-752-6545
Shergill, Robbie	National Semiconductor	robbie.shergill@nsc.com	408-721-7959
Teng, Victoria	NEC	victoria_teng@el.nec.com	408-969-2861
Thatcher, Tom	Hewlett-Packard	thomas@dtc.hp.com	(650) 857-4039
Thompson, David	Lucent	davethompson@lucent.com	610-712-2730
Wakai, Hirosha	Sharp	wakai@slab.tnr.sharp.co.jp	650-638-7273
Whitby-Stevens, Colin	ST Microelectronics	colin.whitby-stevens@st.com	+44 1454 611500
Wooten, David	Compaq	david.wooten@compaq.com	281-518-7231
Yu, Patrick	NEC	patrick_yu@el.nec.com	408-588-5436
Zchikawa, Yujz	Sharp	ichikawa@slab.tnr.sharp.co.jp	+81-743-65-4529

PHY-Link Task Group 8th February 1999

Chairman: Sean Killeen

Introductions were assumed from the previous session

Agenda

1. Review the spec
2. Registers
3. Electrical spec
4. Review of open issues

1. Review the spec

Soft copies and hard copies handed out (Also previously posted on the web)

Clarify the issue of "minimum" distance. The spec allows implementations to be up to 9 inches. Express in metric, round down to 20cm.

Discussion on clocking scheme. The current spec does not allow a S800 link to be connected to a S1600 link. Is this OK?

Discussion on how speeds are selected by writing to registers, and how this may affect the operating speed of the PHY/Link interface.

Issue, this sort of configuration may be done through configuration pins or via software.

Is there a value in having a way of being able to change the clocking speed on the fly? If so, then it should be standardised. Possibly people will find a way of doing this anyway. Maybe want to be able to support changing the speed "not on the fly".

What are the problems in having all links support S3200?

Note that LReq is always single edge, but at the clock rate

NB problem of resets during PHY/Link interface reset in P1394a has not been solved yet.

The requirements of any form of mode switch need to cope with the fact that there are times when the Link is not able to receive update information from the PHY, and the PHY needs to communicate this later.

Note that it is not a requirement that both PHY and Link come up at S800, but there needs to be a way of at least configuring a common-denominator speed.

Is there a requirement to be able to reconfigure other than by configuration pins at power-on. Yes, one example is mobile, where you may wish to save power by running at slower speeds. But this could still be done by strapping option. But, if the pins are optional, some people may prefer a programming option rather than a pin option.

One option could be to write to a programmable register with the new speed, drop LPS, bring up LPS and the interface comes up at the new speed.

We need a coherent story in how to bring up the ports at a software determined speed.

Request proposed solutions.

Review of link requests

Discussion on how to implement fairness, or how to give priority to responses over requests (where the speed may be different). Note that the speed is in the grant. Perhaps it would be better to have a further asynch request to supersede a previous request. Then there is a race condition - either the old or the new request could be granted. If you see a grant come back for the new speed, then you know that you have to re-issue the previous request. If you see the grant come back for the old request, then the link cancels, the new request will be still outstanding, or perhaps still send the old speed packet. In this case, is the request still outstanding - yes - but perhaps not - discussion. Agreed the overall behaviour. But the link can re-issue a request anyway, if the link is not sure whether the request is cancelled or not.

Also there's an issue of the contexts. For example, if a request is issued for next, and then there's an arb reset gap, is the request now considered for the current interval. Can I make a new "next" request. What happens if I now give a current request? Does it supersede the previous "current" request. Dissent.

DW proposal - only one outstanding asynch request, period. Any further asynch supersedes the extant one. If you've issued a "next" (i.e. EVEN or ODD), then you can still issue a current, and this supersedes the "next".

You still need a next request type. But from the PHY point of view, you don't need to keep two queues, you have a single outstanding asynch request (ignoring cycle start and ACKs), this is tagged as current or EVEN or ODD, and whether it is legacy or beta mode.

Agreed.

How does the link interpret the grant coming back - either supply the info in the grant, or perhaps by timing inferences. Supply the info in the Grant (agreed). SK will add the necessary info.

Issue of legacy requests. If you issue a Beta request, and then issue a Legacy request (for the same speed), and see a grant comes back, then you need to know which - same solution, the info will be encoded in the grant.

Cycle starts - OK Note that the Grant indicates a grant for cycle starts

Isoch

- add a notion of next isoch request (ODD/EVEN etc)? Long discussion. Why not just use the standing request - well this solved the problem when we used current/next, but it does not provide sufficient control of which isoch cycle is used in some applications. Note that you have to wait for a cycle start token identifying the start of even before you can issue a request for ODD (have to be sure that that everything knows that the current "ODD" is indeed over).

Is there a usefulness in being able to queue?? Issue for a late request for the current period - good luck - you can't count on it!!!!

Note, there is no "current" isoch request, just ODD and EVEN.

But if you re transmitting the final packet, and you need to get a NEXT out, it will take longer to transmit the LReq than the isoch packet (at S3200).

Perhaps use a token at the end of the packet to say "keep being boss" (that's already available in asynch), so perhaps use this mechanism for isoch as well.

Want to allow use of single FIFOs in Links, so won't know until (nearly) the end whether there's another packet. For example, in a, a clock after the last cycle, you send a something. Perhaps require something to do the same in b.

Indicate at the end of an isoch packet you've got more (to give time for an Lreq to get across) - perhaps "encode" the LReq at the end of the data.

Another way to do it is to say that the PHY can't give out a grant until so many clocks after the end of the previous packet.

Assuming a quadlet FIFO architecture, you need a quadlet clock to see if the FIFO is empty.

Do it the a way, unless there is a good reason for doing it another way. Note this is a tighter requirement than the requirement on 1394a - in a you could use the LReq, so long as you are fast enough. In b, you run out of time. This "hold" is an implied LReq. So there will be a grant from the PHY in the usual way. The only thing we're saving is the LReq time, not the LReq grant time.

Can you keep on trucking, or do you have to turn the bus round. A: you have to turn the bus round, or it gets far too complicated for the Link to track the bus timings (particularly when there's a legacy node in the system).

Interface turnaround time

Propose that the turnaround time not be precisely specified. In DC mode you can turn round quickly (use weak pull-downs to hold the control to 0) - there's an upper bound in the PHY waiting for the turn round. Removes (largely) the distance limitation due to protocol (limited only by electricals). Perhaps set the max turnaround to 10 (or 12) cycles (need about 7 cycles).

So this is simplified. Seems OK. May have an impact on PHY pinging - may care, or may not care any more because of improved arbitration.

PHY register read/write requests

Same rules (don't do write while read is pending) as in P1394a

LReqs

Are Beta-only requests necessary?? Good to be able to support the ability for a link to try Beta-only to see if it can get through to a particular node by using Beta only format. This allows the link to send Beta only requests and the PHY can then humbly comply. Sometimes the PHY can determine whether to use Beta, but this is on a global basis.

Acceleration control requests are not needed.

Issues concerning cycle start notification. Current spec seems correct.

Grants

Beta only grants needed? - see this mornings discussion

Packet receive operation

Min IDLE time?? Good to do if we can. Depends on whether there is one which can be guaranteed on the bus. Open issue

Arb-reset gap notification

First cycle of status transfer should be taken as the cycle at which the gap is indicated. (Note, there is probably no subaction gap indication). In fact the timing on Arb-reset gap is not critical./

PHY-Link Interface Data Phases

Richard's scheme incorporated into current draft.

PHY Status indications

Note that REQ_IGNORED is for debugging only - and it is optional.

P1394a link support

Clarify that LCLK is not provided on a P1394a link, and will be held 0 at the PHY.

2. PHY register map

At the PHY register map - make all the speed encodings consistent.

Extra registers:-

Port error register - byte - incremental counter of errors, see p 109 in the current draft
Disable scrambler test mode

Do we need a bit to re-do the connectivity management?
Action on CWS to check how to do this.

3. Electrical issues

Perhaps the amount of overlap may not be so much.

Perhaps we'll go to HSTL (1.5V) and have done.

4. Open issues

Should the PHY/Link interface speed be configurable, or negotiated, and if so, how?
How does this interact with the desire for power saving (operating at lower speeds)
How does this interact with the need to set port speeds?
Min IDLE time?? Good to do if we can. Depends on whether there is one which can be guaranteed on the bus.
Do we need a bit to re-do the connectivity management? (Action on CWS to check how to do this.)

Present

Takayuki
Eric Anderson
Richard Baker
Steve Bard
Amy Barton
Charles Brill
Navin Chander
Dao-Long Chen
Alistair Coles
Doug Day
Chris Dorsey
Firooz Farhoomand
Lou Fasano
Scott Feller
Tony Foster
John Fuller
Nobuo Furuya
Bob Gannon
Eric Hannah
Jerry Hauck
Lee Hu
Yuji Ichikawa
Ted Inoue

David James
Sean Killeen
Michael Johas Teener
David K Johnson
Eric Kadison
Thang Le
Ed McDonnell
Cyrus Momeni
Michael Morgan
Yoshikatsu Niwa
F Nordby
Ozay Oktay
Brad Saunders
D.C. Sessions
Mike Sorna
Victoria Teng
Thomas Thatcher
David Thompson
David Vrhovnik
Hirohisa Wakai
Colin Whitby-Stevens
David Wooten
Patrick Yu

UTP 5 meeting Minutes - 8th February 1999

Agenda

1. Welcome, introductions, apologies
2. Equalization
 - 2.1 Alternative equalization schemes - Closure on final call
3. Presentations
 - Cat 5 Jitter and Transmitter Mask Study - Eric
 - UTP Signal budget - Alistair
4. Review of the current specification and open issues:-
 - 4.1 Equalization example

The request for an informative example remains open.
 - 4.2 Electrical spec

rise/fall times are ??3-5ns TBD (Action EH)

transmit mask co-ordinates are TBD (Action EH)

Jitter budget- need simulations and proposal

1. Welcome, introductions, apologies

Taken from immediately previous meeting

2. Equalization
 - 2.1 Alternative equalization schemes - Closure on final call

No contributions - issue closed
3. Presentations
 - Cat 5 Jitter and Transmitter Mask Study - Eric

See contributions on the web site

Decide that TX jitter should not be specified by the transmitter mask

Propose that jitter is specified as edges not exceeding 1ns (value proposed by Eric H) more often than one bit time in 10^{12} .

Rise/fall times - decide to adopt the 100BASE-T values (probably means that edges will have to be slowed).

Transmit mask - propose to use the 100BASE-T values, Eric H to re-simulate without incorporating jitter, see if the results fit. Purpose is to bound overshoot and undershoot etc. Again, adopt values from 100BASE-T

Discussed the rise/fall times of the single-ended components, and the need to control the cross-points - decided to do no more than 100BASE-T
 - UTP Signal budget - Alistair

See contributions on the web site

Suggest use time averaging for signal detect, then there's not a problem, as this will compensate for ingress noise. Note that the timings will have to be different, in order to meet the P1394b toning timings. Agreed that Alistair will incorporate appropriate text (based largely on text grand-fathered from other standards)

4. Review of the current specification and open issues:-

4.1 Equalization example

The request for an informative example remains open.

4.2 Electrical spec - see above for provisional resolution on all these

rise/fall times

transmit mask

Jitter budget

Attendance

No attendance list was taken. See previous meetings on same day.