

Proposed Agenda
P1394b Working Group
December 8, 1998

1. Review/Approve Agenda
 - 1.1. Introductions
2. Review of October Meeting Minutes
3. Procedures
 - 3.1. Voting
 - 3.2. Price/Pricing
 - 3.3. Call for Patents
4. Meeting Schedule
 - 4.1. February 8 & 9, Huntington Beach, CA
 - 4.2. March 22 & 23, Tempe, AZ
 - 4.3. April, May NEED INVITATIONS!!!
5. Presentations
 - 5.1. Speeds & Feeds - Colin
6. Task Group Reports (As required)
7. New Business
8. SCAT List Creation
9. Review of Action Items
10. Adjournment

"SCAT" discussed: Defined as "Standard Closure Action Table."

Review of agenda and call for agenda additions. Colin requested "Speeds & Feeds" be added. Proposed agenda accepted via lack of objection.

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

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

Meeting schedule reviewed. Call for meeting hosts for April and May.

Floor passed to Colin:

Background information presented by Colin: Upstarts code did not deal with a particular case. Something like a DS only port on a 1394b PHY. The text at the beginning of the upstarts section should be modified, Colin almost did the update, however, upon doing so, he discovered a possible issue and thus, the purpose of his presentation:

A P1394b Port

- a. shall detect and manage all connect/disconnect events; (**except disabled ports**)
- b. shall be capable of operating in Beta mode in a contiguous speed range compatible with the transmission medium (see Table 1-1); the lowest transmission rate in the port's seed range shall be **no higher than** the required speed and the highest transmission rate shall be no higher than the maximum allowed speed;
- c. when brought out to a 1394-1995 (6 ckt) or a 1394a (4 ckt) connector, shall be capable of operating in DS mode **at S100-S400; at all speeds up to the highest Beta speed on the PHY or S400, which ever is less**
- d. shall start up in Beta mode if Beta capable at a speed compatible with the Beta mode capability of the connected port, and shall start up and run (the "operational speed") at the fastest such speed;
- e. in Beta mode shall support all speeds of packet at or below its operational speed through the use of padding

Note: A P1394b PHY may have P1394a or 1394-1995 ports, which are not considered P1394b ports for these purposes.

- f. shall start up in DS mode if Beta mode is not possible, and if both the port and the port to which it is connected are capable of DS operation and a DC connection exists between them;
 - g. shall support suspend, resume, and low power connection signaling
-

P1394b Media-Dependent Beta Speed Requirements

Table 11-1 Media dependent Beta mode speed requirements

Transmission Medium	Required Speed	Maximum allowed speed
1394-1995 6 ckt or 1394a 4 ckt connector and cable	S800	S1600
RJ-45 connector and UTP-5 cable	S100	S100
PN Connector and POF/HPCF fiber	S100	S200
Duplex LC Connector and 50 micrometre MMF	S400	S3200
Other (not defined in this standard)	S800	undefined

Note that when a port is designed to be capable of supporting more than one medium, then the maximum speed may be set by writing to the max_port_speed register or other implementation dependent mechanism.

Jerry pointed out that a two-port implementation used for the purpose converting DS to UTP-5 precludes the requirement for DS to do anything more than S100.

Much discussion ensued with regard to setting the "required speed" in the table to S100 (which, in actuality, would remove the this column from the table).

An alternative would be to make the first entry in the table S100 for Required Speed.

David expressed an opinion against the six circuit implementation of Beta to have a required speed of S100. He believes this will dilute the proliferation of high speed beta as well as increasing the possibility of a larger number of speed traps.

Colin expressed support of going to all S100 because of the "goodness" which has been brought to 1394 via 1394b could be taken advantage of by the lower speed devices without incurring the higher cost which might come with implementation of the higher speeds.

David (as chair) committed to solicit the body for a motion to change the required speed in the table to S100 at the next meeting (February). Between now and then, the body has been asked to consider the arguments for and against the change. [AR: Colin accepted an AR from the chair to take this issue to the 1394b reflector and solicit further discussion.](#)

Port Properties

A P1394b port may be connected directly to a suitable transceiver for long haul connection (thus also providing DC isolation).

A P1394b port may use DC (e.g. capacitive or galvanic) isolation when operating in Beta mode.

TASK GROUPS REPORTS:

Sean Kileen passed on giving a report out from the PHY/Link task group. Sean did, however, call for input, feedback, etc. on some of the issues which came forward during the task group meeting. Colin gave a brief (very brief) report out for UTP and Upstarts. Mike Teener gave a brief (very brief) report out on Accelerations/B-Port. Steve Bard passed on giving a report out from B-Low Power. The reason stated for such brief reports: the greater body present attended all task-groups.

NOTE: Included with these Plenary minutes (immediately after the Plenary attendance list) are the minutes of the Sub-Task groups (at least those for which minutes were submitted minutes to the Plenary Secretary - Steve Bard).

NEW BUSINESS:

David spoke that there have been some concerns expressed concerning many of the issues carried forth from "legacy" 1394 (such as loop healing). If the group is desirous to fix things like that, it might be necessary to move 1394b away from direct interconnect interoperability with "legacy" 1394. This might require changes to Tree-ID, Cycle Start, etc.

Mike Teener expressed his opinion that it is possible to architect a 1394b PHY using a bus initialization procedure which would heal loops - end up with a beta mode operational enumerated. Other things could be done which could prevent node-IDs changing when new devices are attached. Things could also be done to eliminate a significant number of bus resets.

Mike Teener believes it would be easy to architect an Isochronous mode transport mechanism which does not use the exact procedures currently incorporated in the existing method.

David inquired of the body as to what the opinions were with regard to allowing a six circuit beta only connector cable. Jerry Hauck expressed an opinion against and was supported by Rod Sinks. Mike Teener expressed an opinion in favor. It was brought forward that we have connectors today in 1394b which do not interconnect with existing 1394-1995 and P1394a connectors, however, the current 6-pin circuit for Beta is required to support DS mode of operation if a DS only device is connected to it. Tony Foster expressed that a bad user experience issue would ensue if a 1394-1994/P1394a 6-pin circuit was allowed to be used in beta only implementation.

A straw poll was taken. Twenty three folks believe it is a good idea to be able to have a copper beta only connector which supports speeds greater than S100. No one expressed an opinion that it was a bad idea.

A discussion then took place as to the merits (or not) of a system which delivers power. Expressions against delivering power through the cable and connector were based upon current architecture and specifications. Suffice it to say, systems exist today which do successfully deliver power. Further discussion of this issue will take place (again) at a later date.

A question was opened for discussion: Is there a need for a copper connector for use in beta only implementations which is different than the existing 4-pin or 6-pin copper connector? If not, a user could connect a beta only PHY to a DS only PHY and things would not work.

AR: David will develop an action item for the Copperhead task group to go off and do something about this. AR: David will post the words which make up the action item to the reflector.

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.

STANDARD CLOSURE ACTION LIST (SCAT) LIST:

- 1) Fix start-up to align with P1394a when it has been completed
- 2) Verify that root contention time values will work with long-haul
- 3) Complete design of register map (to include speed limit bits in port registers - settable by firmware, as well as other required bits
- 4) proxy self-id packets for "leaf" nodes in standby
- 5) Complete port state machines
- 6) Resolve signal detect and related issues for UTP-5
- 7) Verify functionality of 1995 cables at S800
- 8) Complete UATP electricals
- 9) Finalize jitter budgets and specifications for all speeds and media
- 10)Upstarts complete standby interaction
- 11)Standby encoding signaling

- 12) arbitration token completion
- 13) PHY-Link electricals
- 14) A-Link to B-PHY operating model
- 15) finalize BOSS
- 16) Resolve beta mode media dependent speed
- 17) Able to identify Beta PHY ID at self-ID time
- 18) New beta mode only copper connector
- 19) New Tree ID which will perform loop healing (and other nice features)
- 20) New Cycle start
- 21) 1394.1 liaison as required
- 22) Resolve speed map issues
- 23) Complete DC BIAS specification
- 24) Changes to IRM (band-width & channels allocation)
- 25) PHY test modes
- 26) Delayed error detection & reporting (10.9.3.2)
- 27) Write clause three (Summary)
- 28) Beta only copper cables for speeds greater than S100

David invited those who do not find their presence adds value to generation of additional items to the SCAT list may consider the meeting adjourned. The remaining folk will continue to develop SCAT items.

IEEE P1394b has not reached the 100% perspiration level - there is still quite a bit of inspiration to be completed. It will be much easier to determine a standards completion date (goal/target, etc.) when the group reaches the 100% perspiration level. It is not clear when that point will be reached.

[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.](#)

Meeting officially adjourned 4:49 PM

✓	V	Name	Company	Email	Phone
✓		Baker, Richard	Solectron, Texas	rtb@artb.itg.ti.com	512-425-6205
✓	x	Bard, Steve	Intel	steve.bard@intel.com	503-264-2923
✓	x	Brunker, Dave	Molex	dbrunker@molex.com	630-527-2622
✓	x	Chen, Dao-Long	LSI Logic	dao-long.chen@lsil.com	970-223-5100 x5461
✓		Clee, Jim	Lucent	jclee@lucent.com	(610) 712-2848
✓	x	Coletta, Mike	Harris	mcoletta@harris.com	949-707-1143
✓		Conroy, Keith	Pulse	keithconroy@pulseeng.com	(215) 781-6400
✓		Feller, Scott	Hewlett-Packard	scott_feller@hp.com	(650) 857-7759
✓		Foster, Tony	Hewlett-Packard	tony_foster@hp.com	(916) 785-1092
✓	x	Fuller, John	Microsoft	jfuller@microsoft.com	425-703-3863
✓	x	Furuya, Nobuo	NEC	nobuo_furuya@el.nec.com	(408) 969-2479
✓	x	Hannah, Eric	Intel	eric.hannah@inbtel.com	408-765-4441
✓	x	Hauck, Jerry	Zayante, Inc.	jhasuck@zayante.com	510-668-1006
✓	x	Heilmann, Keith	IBM	kheilman@us.ibm.com	914-892-2413
✓	x	Hiraoka, Daisuke	Sony	dai@sm.sony.co.jp	+81 3 5448 4603
✓	x	Killeen, Sean	SSL	sean.killeen@ssl.ie	+353 1 402 5700
✓		Le, Thang	Hewlett-Packard	tl@rose.hp.com	(916) 785-4667
✓		Liburdi, Francesco	Hollingsworth	liburdif@hollingsworth.com	607-748-0025
✓		Lopata, John	Molex	jlopata@molex.com	(630) 579-4110
✓		Lu, Jinghui	Rocketchips	jlu@rocketchips.com	512-685-4429
✓		Ma, Te Khac	Lucent Technologies	tma@lucent.com	(610) 712-2195
✓		McWilliams, Bob	Initio	bobm@initio.com	(408) 577-1919
✓		Nave, James	Texas Instruments	jim-nave@ti.com	972-480-3468
✓		Nelson, Jim	Northrop Grumman	nelsoja5@mail.northgrum.com	562-948-8357
✓	x	Northey, Bill	FCI	northewa@bergelect.com	717-938-2119
✓	x	Oktay, Ozay	Canon	ozay_oktay@cissc.canon.com	714-856-7180
✓		Rose, Bonnie	Amphenol	bonnier@asmphenolcanada.com	(416) 291-4401
✓	x	Saito, Kyozo	Alps	kyozo_s@gw3.alps.co.jp	+81 229 23 5111
✓	x	Saito, Tomoki	NEC	saito@ccm.cl.nec.co.jp	+81 44 856 2082
✓		Sato, Hideaki	Canon	hsato@cns.canon.co.jp	+81 3 3757 8985
✓		Sinks, Rod	Omneon Video Net	rsinks@omneon.com	(408) 558-2107
✓	x	Skidmore, Jim	Texas Instruments	j_skidmore@ti.com	972-480-2094
✓		Smith, David	Texas Instruments	desmith@ti.com	(972) 480-6345
✓	x	Smith, Michael	ControlNet	michael@controlnet.com	(408) 341-1428
✓	x	Teener, Michael Johas	Zayante	mike@zayante.com	831-461-4901
✓		Thatcher, Tom	Hewlett-Packard	thomas@dtc.hp.com	(650) 857-4039
✓	x	Whitby-Strevens, Colin	ST Microelectronics	colin.whitby-strevens@st.com	+44 1454 611500
✓	x	Wooten, David	Compaq	david.wooten@compaq.com	281-518-7231
✓	x	Yu, Patrick	NEC	patrick_yu@el.nec.com	408-588-5436

IEEE P1394b
B-Low Power Task Group
Minutes
7th December 1998
Double Tree In
Monterey, California

Agenda presented:

1) Introductions

2) Accept volunteer to write minutes

3) Approve Previous Meeting Minutes

4) Review

4.1 Characteristics of "Standby"

4.2 Issues

4.3 Others -

5) Discussion

5.1 Self-ID Proxy Register Contents

5.1.1 Node-ID

5.1.2 Power Class

5.1.3 Port present

5.2 Multi-Port "Leaf" Nodes

5.2.1 Only one connected port

5.2.1 Only once active port, 1 or more suspend ports, 1 or more disconnected ports

5.3 Extended Control Code for Standby

6) Review AR's

7) Opens

8) Adjourn

Minutes:

The sub-chair wishes to express his appreciation to Jerry Hauck for volunteering to take minutes for the meeting.

- Previous minutes (October meeting) approved unanimously.
- For the benefit of newcomers, the working group assumptions were reviewed. Nodes that are the target of standby need to understand if they have been commanded to suspend versus going into standby. For example, an SBP2 disk drive reacts differently: suspend may cause a logout while standby doesn't. Since the target node must discern between the two, it is not possible to preserve legacy silicon (PHY and Link) for use in a target by using TX_Suspend to mean either suspend or standby. Furthermore, the Data/Strobe arbitration state codeings are exhausted; consequently, standby will be a feature only supported across a beta connection.

- Given that new silicon is required both on the target side and the parent side of a “standby” connection, the working group reaffirmed that the command to enter standby will be addressed to the child side (leaf) of the connection rather than to the parent side. This provides the most flexibility in defining what state a node must be in to allow entry into the standby state. The child has the benefit of all of its local state information in making this determination; the parent does not. The additional flexibility of having the child make the determination allowed the working group to consider a looser interpretation of a “leaf” node.
- The working group had a lengthy discussion on the definition of a “leaf” node. When a given device is in standby, another node (i.e. the parent) is responsible for acting as a proxy and synthesizing the self-id packet of the sleeping device. To limit the complexity and buffer requirements in the proxy, only “leaf” nodes can enter standby. (Thus, the parent only needs to synthesize a single self-id sequence for the sleeping child.)

Steve Bard raised the question of whether a multi-port PHY with only one active connection (to the parent) should be considered a “leaf”. Typical applications could be when a product used a 3-port PHY but only added one connector. Additionally, it was argued that there may be some utility in allowing a node with more than one connector to participate in standby as long as the other connections were not active (i.e., disabled, disconnected, or suspended). Since such inactive connections don’t appear in the self-id sequence, the parent still wouldn’t need to proxy anymore than a single node’s self-id sequence.

The working group decided that a node will only enter standby if the parent port is the only active port. If, while in standby, a child port has an event, the sleeping node will come out of standby and will force a bus reset to push the topology change out to the active network.

- The working group also discussed the minimal amount of storage required to properly synthesize the sleeping node’s self-id sequence. By allowing a multi-port PHY to enter standby, the parent may need to proxy as many as three full self-id packets (for example, when the one active port is port 15).
[AR: Jim Skidmore took an action item to publish \(on the reflector\) his idea of simply memorizing the one active port of a child and, when called upon to proxy, mark all other ports as one of the not-active forms.](#)

IEEE P1394b B-Low Power Task Group
ATTENDANCE LIST: 7th December 1998, Monterey, California

✓	V	Name	Company	Email	Phone
✓		Akgul, Jeff	Texas Instruments	jakgul@ti.com	(972) 480-1974
✓		Baker, Richard	Solectron, Texas	rtb@artb.itg.ti.com	512-425-6205
✓	x	Bard, Steve	Intel	steve.bard@intel.com	503-264-2923
✓	x	Chen, Dao-Long	Symbios	dao-long.chen@symbios.com	970-223-5100 x9461
✓		Clee, Jim	Lucent	jclee@lucent.com	(610) 712-2848
✓		Conroy, Keith	Pulse	keithconroy@pulseeng.com	(215) 781-6400
✓		Feller, Scott	Hewlett-Packard	scott_feller@hp.com	(650) 857-7759
✓		Foster, Tony	Hewlett-Packard	tony_foster@hp.com	(916) 785-1092
✓	x	Furuya, Nobuo	NEC	nobuo_furuya@el.nec.com	408-588-5414
✓	x	Hannah, Eric	Intel	ehannah@mipos2.sc.intel.com	408-765-4441
✓	x	Hauck, Jerry	Macro Designs	jerry@macrodesigns.com	510-668-1006
✓	x	Heilmann, Keith	IBM	kheilman@us.ibm.com	914-892-2413
✓	x	Hiraoka, Daisuke	Sony	dai@sm.sony.co.jp	+81 3 5448 4603
✓	x	Killeen, Sean	SSL	sean.killeen@ssl.ie	+353 1 402 5700
✓		Le, Thang	Hewlett-Packard	tl@rose.hp.com	(916) 785-4667
✓		Lin, Cjhang-Chi	Vishali	ccliu@controlnet.com	
✓		Lu, Jinghui	Rocketchips	jlu@rocketchips.com	512-685-4429
✓		Ma, Te Khac	Lucent Technologies	tma@lucent.com	(610) 712-2195
✓		Marazas, Gerald	IBM	marazas@vnet.ibm.com	919-543-6892
✓		Nave, James	Texas Instruments	jim-nave@ti.com	972-480-3468
✓		Nelson, Jim	Northrop Grumman	nelson@laya.srtc.northop.com	562-948-8357
✓	x	Oktay, Ozay	Canon	ozay_oktay@cissc.canon.com	714-856-7180
✓	x	Saito, Kyoza	Alps	kyozo_s@gw3.alps.co.jp	+81 229 23 5111
✓		Saito, Tomoki	NEC	saito@ccm.cl.nec.co.jp	+81 44 856 2082
✓		Sato, Hideaki	Canon, Inc.	hsato@cns.canon.co.jp	+81 3 3757 8985
✓		Sinks, Rod	Omneon Video Networks	rsinks@omneon.com	(408) 558-2107
✓	x	Skidmore, Jim	Texas Instruments	j_skidmore@ti.com	972-480-2094
✓		Smith, David	Texas Instruments	desmith@ti.com	(972) 480-6345
✓	x	Smith, Michael	ControlNet	michael@controlnet.com	(408) 341-1428
✓	x	Teener, Michael Johas	Zayante	mike@zayante.com	831-461-4901
✓		Thatcher, Tom	Hewlett-Packard	thomas@dtc.hp.com	(650) 857-4039
✓	x	Whitby-Strevens, Colin	ST Microelectronics	colin.whitby-strevens@st.com	+44 1454 611500
✓	x	Wooten, David	Compaq	david.wooten@compaq.com	281-518-7231
✓	x	Yu, Patrick	NEC	patrick_yu@el.nec.com	408-588-5436

**PHY/Link Interface Task Group
December 7th 1998
Monterey, California**

Minutes

Sean Killeen conducted a review of the latest draft (in Draft 0.15 of P1394b). The following comments and discussions were noted.

Page 1

add requirement to support interoperation between adjacent electrical supply generations (3.3, 2.5, 1.8, 1.5)
add requirement to support a reach of 9 inches, further would be good

Page 2

Add title: Figure 12-n PHY/Link logical interface
Add note to say that this diagram shows logical connections, and is not intended to imply a DC connection.

reword - LinkOn and LPS are mandatory for the PHY

Page 3

issue - do we need PINT, or can we use a mechanism as in P1394a. It is suggested that we need this for coherency (see further discussion on Lreqs below)

Page 6

bring the LReq stuff up here

Page 7

discussion - do we want to have a way for the PHY to cancel or truncate an operation in the case of an errant link (packet overrun) - with an appropriate interrupt to the link. Propose a PINT to abort a transmit.

Page 11

request to everyone to review the PHY/Link interface handover

Page 12

In the timing diagram - would like to see the phase - transmit/receive - perhaps the actual values.

Keith Heilmann/Jim Skidmore - problem of determining the bus mode is not that difficult - perhaps don't really require to support this - perhaps not worth the cost.

Page 14

Note that the PHY only queues one isoch and one asynch, which is tagged for the current or next interval (but see subsequent discussion on Lreqs, below)

Maybe have a current, even, odd Lreqs for asynch requests.
In the other direction (PHY to Link) - we have indications of entering odd,
entering even
Then we don't need an acknowledge the request, and the link can just wait
for the grant.

If the link thinks we're in the even interval, the link is allowed to issue
current and "odd". If the PHY moves up to odd, and receives "odd", it treats
it as current.

Discussion on cancellation. Jerry points out that you can't just use any old
grant, as there is a speed issue - the grant has to be for the correct
speed.

So have a "grant current", grant even and grant odd (note strictly, there's
no ambiguity between the latter two, only one is possible at any one time,
but it is cleaner for symmetry and a bit more robust).

We'll queue two deep in the PHY.

Still need a cancellation policy when the link receives a grant.

Question - should we send the speed with the data rather than with the
LReq - for the Beta only requests. So we need a beta-only grant.

Need various types of grant as it is necessary for the link to be able to
determine which grant was granted, as it needs to know what speed to send.

Idea for solving the speed code problem - the Link puts out the speed with
the packet (rather than with the LReq) - if it is or might be for a legacy
domain, then the link puts out the speed code followed by 140ns. If the link
knows it is Beta only then the link just puts out the speed code followed by
data.

Jerry points out that he thinks there are some optimisations that can be
applied by the PHY
concerning whether a packet ever goes to a legacy domain - the PHY can
determine this information from self-ID. The PHY makes the decision based on
packet speed (e.g. the PHY may know that all legacy packets are S200 and
below).

Move the Lreq section to earlier in the document (just before the response
info)

Isoch period, not fairness interval (description of isoch requests)

Suggest make "isoch next" sticky - it is issued once and then lasts for
ever. If a PHY wants more than one request then it can issue more current
isoch interval requests
- so we need a request to make it unsticky

Need to check that the timing for ACK requests works - that the request can
be got out in time.

Maybe there is a requirement to transmit subaction end notification
synchronously with the data at the end of the packet (analogously to the
current hold protocol). Note that it will be necessary to release the bus

for a concatenated packet as the PHY will still have to ensure the timings etc. Maybe concatenation as a mechanism is not now required. It does not happen on the bus.

Page 17

Everyone to review the status transfers mechanism

Page 22

Electrical spec

All to look at the FC-HSPI spec - www.t11.org - document 98-587v0

Need a service from the PHY - cycle start arrived (cycle start token received)

Need a service to the PHY to transmit a cycle start (entering isoch phase).

Attendance

Jeff Akgul
Richard Baker
Steve Bard
Dao-Long Chen
Jim Clee
Keith Conroy
Scott Feller
Tony Foster
Nobuo Furuta
Eric Hannah
Jerry Hauck
Keith W Heilmann
Daisuke Hiraoka
John L Hizz
Claude Huss
Sean Killeen
Thang Le
Jinghui Lu
Te Khoc Ma
James Nave
Jim Nelson
Ozay Oktay
Kyoza Saito
Tomoki Saito
Robbie Shergill
Rod Sinks
Jim Skidmore
David Smith
Michael A Smith
Michael Jonas Teener
Tom Thatcher
Colin Whitby-Strevens
David Wooten

**P1394b UTP5 Task Group
December 7th 1998
Monterey, California**

Minutes

1. Welcome, introductions, apologies

Introductions were assumed from the previous meetings on the same day. Additions were made to the morning's sign-up sheet. CWS conveyed apologies from Alistair Coles.

2. Equalization

2.1 Alternative equalization schemes

It was agreed to issue a final call for contributions on this topic, which will be removed from further consideration at the next meeting in the absence of any further contributions.

3. Other presentations

None

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

4.1 Equalization example

Dao-Lon Chen offered apologies that LSI Logic are now not able to provide an informative example, due to personnel reassignments.

The request for an informative example remains open.

4.2 Electrical spec

rise/fall times are ??3-5ns TBD (Action EH)
transmit mask coordinates are TBD (Action EH)
Jitter budget- need simulations and proposal

Eric Hannah stated that he is carrying out modelling in order to provide suitable numbers, which he hopes to have at the next meeting. He indicated that he is using an analytical cable model for this. CWS asked about general availability of cable models for UTP5, and the following possible contacts were mentioned: Level 1 (Sailesh Rao), Broadcom (Andrew Castellano). Eric Hannah also indicated that he was considering preparing a white paper which would include the model he is using.

CWS action to check the jitter measurement methods in P1394b.
CWS reported that he had a preliminary presentation which would be made later.

4.3 Testing - currently deleted from the draft

Statement of conformance with TIA TSV-67, BER of 10⁻¹²

Final call for further contributions on testing

There were no further contributions, and so this item is dropped.

5. BER values

CWS reported that the decision in the last meeting to raise the BER to 10⁻¹² had been questioned by Alistair Coles and Masood Shariff on the reflector, primarily on the basis that current components and technology are designed to a figure of 10⁻¹⁰.

There was extensive discussion.

It was pointed out the 1394 has quality of service characteristics which are significantly different from, say, ethernet. In particular, isochronous transport is a unique capability which is assumed to work reliably for extensive periods of time. A BER of 10^{-10} implies an error rate of one every 100 seconds, which would not be very suitable for isochronous service.

A straw poll was taken to keep the 10^{-12} value (17 votes) or to revert to the 10^{-10} value (7 votes), or don't care (4 votes).

After the straw poll, Keith Conroy from Pulse indicated that they have components which have been measured at ATM rates (155 Mbps) over 150m which have a much better error rate than 10^{-12} . The consensus was to keep the 10^{-12} value.

6 Signal detect values

CWS drew attention to the fact that there is no signal detect value, and examination of a possible budget had also shown up possible problems with some of the other electrical values. He also reported that Alistair Coles has already taken an action to consider these issues.

A possible "no signal" budget is

Off Transmitter 20mV
NEXT 26mV
Noise 15mV
Margin 5mV

Total 66mV

Anything lower than this should be considered as no signal.

An alternative is to use the short-haul copper budget - which totals 80mV.

However, the RX sensitivity is currently set at 50mV, which seems odd to have this lower than the off signal threshold, and Transmitter Off is currently set at 162mV, which is far too high.

Keith Conroy pointed out that if the RX sensitivity is set higher than 50mV, this implies that the launch amplitude will also need increasing to compensate for the cable attenuation. This in turn will make it more difficult to pass FCC Class B.

Consideration of this issue is open.

End of meeting

Postlude

Given the available time, CWS presented his jitter review (see presentation on web site)

Attendance

Jeff Akgul
Richard Baker
Steve Bard
Dao-Long Chen
Jim Clee
Keith Conroy
Scott Feller
Tony Foster
Nobuo Furuta
Eric Hannah
Jerry Hauck
Keith W Heilmann
Daisuke Hiraoka
John L Hizz
Claude Huss
Sean Killeen
Thang Le
Chang-Chi Liu
Jinghui Lu
Te Khoc Ma
James Nave
Jim Nelson
Ozay Oktay
Bonnie Rose
Kyoza Saito
Tomoki Saito
Robbie Shergill
Rod Sinks
Jim Skidmore
David Smith
Michael A Smith
Michael Jonas Teener
Tom Thatcher
Colin Whitby-Strevens
David Wooten
Patrick Yu