

P1394B Working Group Meeting - Plenary
October 13, 1998
Lahaina, Maui, Hawaii

1. Review Agenda
 2. Review of September Meeting Minutes
 3. Procedures
 - 3.1. Voting
 - 3.2. Price/Pricing
 - 3.3. Call For Patents
 4. Meeting Schedule
 - 4.1. December 7 & 8, Monterey, CA
 - 4.2. February 8 & 9, Huntington Beach, CA
 - 4.3. March 22 & 23, Tempe, AZ
 - 4.4 April (Need Invitation)
 5. Presentations
 - 5.1 1394 & 802 -- Michael Smith
 - 5.2 6.25 micron GOF
 - 5.3 Power Distribution over UTP5 & Fiber
 6. Task Group Reports (As required)
 7. Review of Action Items
 8. Adjournment
-
-

Bill Northey moved to accept August meeting minutes, Colin Whitby-Stevens seconded. Minutes approved without objection. Point of silent order: A correction to the previous, previous meeting (August) minutes: Ron Soderstrom was present at the August meeting, therefore, the voting roster needs to be corrected to show his voting status.

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

More spec. work has to do over e-mail on the reflector. There are insufficient time during the working group meetings and the time between the next few meetings makes time a valuable commodity.

Worked through meeting schedule. Intel will host in Tempe, Arizona (same location as last year) for March.

A host is needed for April and beyond.

Michale Smith ControlNet, Inc.) delivered a presentation on "1394 & 802". The focus of the presentation was to announce 802 plenary meeting (they occur about three or four times a year). He delivered an invitation to all in attendance

to attend the next meeting on November 10th in Albuquerque, New Mexico.. Their goal (as a company) is to use switch 1394 in LAN traffic; Extend media to 100 meters at 400Mbps (and faster); Incorporate Ethernet connectivity.

Their intention is to use 1394b standards (fiber, 100 meter UTP, new "PMII" (PHY/Link Interface), bus protocols); They intent to produce a new UTP PHY based on 802.3y for 100/299/400 (based on 802.3ab for 800/1600)

Their value proposition is "Lower cost, higher speed networking (PC interfaces, switches); QoS and Isochronous support for Ethernet (No RSVP/TCP/IP overhead; very low latency)

ControlNet coined a term for this "EtherWire™". EtherWire™ is a switch (multi-port 1394 bridge); EtherWire™ is a LAN gateway (Ethernet LAN)

EtherWire™ PHY: 100/200/400Mbps IEEE 1394 at 100 meters (up to 100 Mbps Full-Duplex on each of 4-pairs; synchronous signaling; based on 802.3 standard); Uses fast Ethernet copper wiring (certifiable FCC Class B); Single chip transceiver (data cable only, no power)

EtherWire™ - Ethernet over IEEE1394 PHY/Link Interface; 100/200/400MBps; Uses streams for multicast traffic.

<sec. note: Mike shall make a soft copy of his presentation (with the "confidential" stamp removed) available for posting to the P1394b ftp repository.>

In summary, the point is to converge two technologies (Ethernet & Firewire™); ControlNet is looking for partnerships (standards work, development); An invitation has been extended to all IEEE P1394b Working Group attendees to attend the November 10th meeting (an 802.3 tutorial) in Albuquerque. In the 802 Standards (LAN) tree in the IEEE Internet website, there is a section for meetings logistics - this meeting can be found there (air fare, hotels, meeting location, etc.).

Masood Shariff provided a presentation on selecting 62.5 micron multi-mode GOF - all reading this need to know that the handout delivered during the meeting (upon which these minutes are based) had some errors in it. Folks must download from the ftp repository a corrected revision of his presentation (the name of the file containing the corrected revision of the presentation was not available at the time these minutes were created).

Masood (Lucent Technologies) proposes to add 62.5 micron multi mode fiber to 1394-B by including the following specifications. Note that only those parameters that differ from the present 50 micron multi-mode fiber specifications are listed:

Description	Value	Unit
Modal bandwidth at 850 nm (minimum, overfilled launch)	200	MHz * km
Zero dispersion wavelength (λ_0) range	$1330 \leq \lambda_0 \leq 1365$	nm
Zero dispersion slope (S_0) (maximum)	0.11 for $\leq \lambda_0 \leq 1348$ and 0.001 ($1458-\lambda_0$) for $1348 \leq \lambda_0 \leq 1365$	ps/nm ² * km

The international standard reference is IEC 793-2 type A1b.

Rationale:

The MMF installed base is predominately 62.5 micron core size. A 1996 survey of 107 large business in the United States conducted by Compaq for IEEE 802.3z concluded that 84% of all fibers within buildings is 62.5 micron multi-mode fiber. This specification allows IEEE 1394-B to leverage that installed base. The proposed fiber supports 1394-B up to 2000 Mbaud rate for distances over 100 meters. Refer to link analysis charts for details.

David Wooten expressed his opinion that "1394 as a business network 'sucks'..."

The same connector can be used for 50 micron or 62.5 micron. The transceiver must be one which accepts BOTH 50 micron and 62.5 micron.

Ron Soderstrom stated that we must go look at all of these numbers again to determine if one transceiver can, indeed, serve all purposes. Ron is not sure that there has been enough time to insure these new numbers are exactly correct.

Masood made a motion to re-open the issue of supporting 62.5 micron MMF. Bill Northey seconds this motion. 3 in favor; 4 opposed; 13 abstains

Masood continued, by proposing the following specification for receiver bandwidth (Eric Hannah spoke against the reason for having this specification when there is a bit error rate and eye opening information in the specification already):

Description	Value	Unit
Receiver electrical 3-db upper cutoff frequency (maximum)	2400	MHz
Receiver electrical 3-dB lower cutoff frequency (minimum)	1600	MHz

Rationale: This specification helps ensure interoperability between various manufacturers devices. The upper cutoff frequency limits the jitter-generating effects of differential mode delay. The lower cutoff frequency specification limits the inter-symbol interference contribution by the receiver (see link analysis spread sheet comparing receiver bandwidths of 1000 and 1600 MHz). While it is common practice to design the lower cutoff frequency of 75 to 80% of the baud rate, the lack of a lower cutoff specification opens this standard up to interoperability problems and uncertainty in link analysis calculations.

Ron Soderstrom spoke to the reason for incorporating these numbers which dealt with transceiver interoperability. Eric still believes the current specification should cover this issue. Eric could support the inclusion of these numbers as an informative portion of the specification.

Revote: 3 in favor; 4 against; 12 abstaining (Eric left)

After considering the "strong" opinion of the group on this one, the chair gave the interpretation that this issue will NOT be re-opened.

David Wooten "slipped" in another agenda item (5.3) which, by the way, was not approved when the agenda was first "approved" - but that is O.K. because he is the chair and can do those type of things.

So, in that regard, a discussion ensued in which the allowance of a mixed set of connector types on a platform must be allowed. Take, for example, a sub-PHY which needs to consume power on one port, but, it AC coupled (fiber, say) on the other port and, thereby, is not able to pass power. In addition, a platform which has a six pin port and some sort of AC port (UTP5, Fiber, etc) is like unto a mixture of four port and six port connectors.

Steve Bard has taken an AR to act as liaison with the TA subtask Cable Power Distribution group to have the next revision of CPD address (appropriately) these issues.

TASK GROUP REPORT OUT

UTP 5 - Colin Whitby-Strevens

Colin has provided the secretary with a more complete set of task group notes for inclusion in these minutes. Meanwhile, Colin presented a brief synopsis of the meeting events to the body present.

Upstarts - Colin Whitby-Strevens

Colin has provided the secretary with a more complete set of task group notes for inclusion in these minutes (which follow). Meanwhile, Colin presented a brief synopsis of the meeting events to the body present. The task-group is moving into a monitor and maintenance mode. Essentially, Upstarts is done until folks determine something must be changed or fixed.

B-Port Accelerations - Mike Teener

Four things were done:

Discussion on services list Jerry Hauck did (a reformatted version of what he posted to the reflector). This will be put onto the ftp repository

State machines that Mike thought he would have ready will not be ready until about a week before the next meeting (the end of November)

A couple of other minor issues: We wanted to delete the legacy requests from the end signals and the indication that we want to be able to automatically run in beta mode whenever we can - even when speeds are below 800 megabits. Need to determine if we are in beta mode (quickly) in order to get through a bridge and in order to that the easy way, Jerry proposed we need to discover during the self-ID process the maximum data strobe node which exists on the network (Alistair knows how to do this).

At this stage, this (the state machines and simulations) is the single item in the critical path for B completion. The chair stated that "...anything which is not yet completed is a problem"

Mike Teener will e-mail to the secretary a more complete set of meeting notes for inclusion in this minutes. If the e-mail is not received by Friday, these minutes will be published without the notes.

Copperheads -Bill Northey

Bill has provided the secretary with a more complete set of task group notes for inclusion in these minutes (which follow). Meanwhile, Bill presented a brief synopsis of the meeting events to the body present.

TASK GROUP MINUTES:

UTP 5 - Colin Whitby-Strevens

P1394b UTP5 Task Group Meeting Minutes - October 12th 1998

1 Welcome, introductions, apologies

Introductions were assumed from the immediately previous session. An apology was received from Alistair Coles.

2 Equalisation

D L Chen presented some results of the performance of the receiver equalisation scheme that he had presented at the previous meeting (presentation to be placed on the web site). These results were based on the use of MLT-3 with PRBS encoding. CWS requested that the measurements be repeated using the P1394b NRZ (binary) signalling and the P1394b encoding scheme (which limits the run length).

Eric Hannah indicated that he is developing a technique for a combination of pre-emphasis and receive equalisation which he intends to present at the next meeting.

3 Reach requirements

The previously tabled motion "To delete the requirement that all receive circuits must work over all lengths of cable up to 100m" was considered. CWS read out an email contribution from Alistair Coles. After a short discussion the vote was taken. DW reminded the group that of the attendance requirements for voting. For: 1; Against: 11; Abstain: 7. Motion fails.

4 Review of the current specification and open issues

4.1 Equalisation example.

It was agreed that it would be good to include an informative example of an adaptive equaliser. D L Chen agreed to check with his company (LSI Logic) as to whether the example he had presented can be incorporated (block diagram and descriptive paragraph).

4.2 Electrical specification

Eric Hannah indicated that he is carrying out some simulations which will assist confirming the values which are currently TBD in the specification. His analysis would also cover the jitter budget. CWS agreed to check whether the latest FC knowledge on jitter measurement techniques has been incorporated into the P1394b jitter budget (this applies to all specifications, not just UTP5).

CWS requested further contributions on electrical modelling in order to corroborate these values.

It was agreed that the electrical specification would not be aligned with the short haul - it would retain the lower launch levels for EMC reasons and the more stringent sensitivity.

The current value of BER in the specification was discussed. It was agreed to increase this to 10^{*-12} .

4.3 Testing

It was agreed to add the requirement that installed cable plant should be tested according to TIA TSB-67.

A final request is made for contributions on testing. If none are forthcoming by the next meeting, then no further testing information will be included in the spec.

Present

Dao-Long Chen
Richard Churchill
Michele Dale
John Fuller
John Glenski
Thomas Han
Eric Hannah
Jerry Hauck
Sean Killeen
Steven D Martinez
Angel Mora
Shuji Nagao
Yoshikatsu Nima
Bill Northey
Takayuki Nyu
Ozay Oktay
Bill Prouty
Kevin Ross
Kyoza Saito
Tetsuya Sato
Masood Shariff
Michael Shinkarovsky
Jim Skidmore
Ron Soderstrom
Michael Jonas Teener
Dave Thompson
Satoru Toguchi
Colin Whitby-Strevens (chairman)
David Wooten

Upstarts - Colin Whitby-Strevens

P1394b Upstarts Minutes - October 12th 1998

1. Welcome, Introductions

Those not present at the earlier meetings in the day introduced themselves

2. Simulation task group feedback

None

3. Review of latest code - v9 (in P1394b Rev 0.14)

CWS summarised the modifications made in the latest version

Response to simulation report of September 9th

Bugs:-

1. `receive_speed_indications()` - Centering of speed tone bits modification accepted and implemented. Note that the "chatter" code in `signal_detect_OK()` is removed.

Technical Issues:-

1. The constants are now defined with their appropriate min/max values. In addition, the variables which are supposed to be PHY register bits are now carefully described as such.

2. accepted

3 & 4. Accepted - `received_speed_indications` has been updated to deal with the longest possible valid runlength of zeros in a speed indication (see also 5)

5. To deal with main toning loop speed variations, and to deal with clock speed variations (bug when a port is in suspend - not noticed by Jerry, but related to those he did notice), yet another autonomous routing called "toner" has been invented. This sends the periodic tone when required, and also sends the speed indications between tones when required. Thus all the transmit timing is now localised to this routine, and all paths through it should take the same time!

C Code syntax and nits - all accepted

Editorial suggestions - all accepted, except for the comment on redundant parentheses in an "if" statement - they were not redundant, but there were some missing parentheses elsewhere!

Other points (see Jerry's email "Simulation feedback" of 11th September):-

To deal with "toning into `TpBias`", the `send_tone()` routine now does the necessary check, and simply waits for `TONE_DURATION` if it sees incoming `TpBias`. This enables some other simplifications in other parts of the code, and allows removal of the unnecessary waits.

The `set_beta` routine has been changed as suggested to try to avoid the unnecessary speed exchange (though it is not clear that this will always be the case).

Code has NOT been changed anything to avoid toning once `tried_bias` has failed.

With respect to dealing with a P1394a PHY powering up after we have tried bias once (email from Jerry of September 10th - Bias detection handling in `connect_status()`), CWS expressed a preference for JH's final possibility. JH agreed to review the code to see if this simplification is possible. Note that the current code functions correctly.

Other other points:-

1) There's now a PHY register bit defined to say that there's a DC connection.

2) There is now a cable speed variable, set by hardware.

In discussion it was agreed that if software sets the maximum port speed to a speed (S400 or below) at which the port is not designed to support Beta mode, then the port should attempt to establish a D/S connection. If the port is designed to support Beta mode at this speed, then indeed the port should attempt to negotiate Beta mode.

A need is therefore identified for a concept of "minimum Beta speed". If the incoming speed in a speed negotiation is less than this, then the port will not accept the beta speed negotiation, but drop back to D/S. This is a physical property of the implementation and is made available to software via a read-only PHY register.

4. Other presentations

None

5. Review of open issues

- Code update to be consistent with final version of 1394a
 - monitor connection status on a beta_mode disabled port?
- Experience of Fibre Channel start-up
 - high impedance or static transmitter may produce valid-looking signal at the far end
 - NEXT issues - FC indicates NEXT up to 10%

Post meeting feedback from the connector group. It was agreed that the 6-pin connector will not support "local-plug-present". The corresponding PHY pin should which therefore should be strapped to false, but the algorithm changed so that toning is tried should a DC connection be detected, even though plug present is not set.

There was discussion on the possibility of having a simple converter (dongle) hard wired to a 6-pin connector which converts to UTP5. This would require a form of the 6-pin connector which would be able to detect S100 only. It would also be necessary to have some form of keying to prevent this special 6-pin connector being plugged into existing 6-pin receptacles. The implication for upstarts that "avoid toning once Tried-Bias has failed" is not the desired behaviour - the dongle would indicate a DC connection, but there would be no connection to the "other" end. So toning would have to continue to allow a connection to be subsequently established. However, consideration of the proliferation of connector types results the "sub-PHY" solution being preferred, where the dongle contains a PHY.

Present

Steve Bard
Dave Brunker
Dao-Long Chen
Richard Churchill
Michele Dale
Steve Finch
Noburo Furuya
John Glenski
Eric Hannah
Jerry Hauck
Sean Killeen

Steven D Martinez
Angel Mora
Shuji Nagao
Yoshikatsu Nima
Toshihiro Niitsu
Bill Northey
Takayuki Nyu
Ozay Oktay
Bill Prouty
Matt Rooke
Kevin Ross
Tetsuya Sato
Masood Shariff
Michael Shinkarovsky
Jim Skidmore
Michael Smith
Ron Soderstrom
Michael Jonas Teener
Dave Thompson
Satoru Toguchi
Colin Whitby-Stevens (chairman)
David Wooten
Patrick Yu

B-Port Accelerations - Mike Teener

Copperheads -Bill Northey

IEEE P1394b Connector Task Group "Copperheads" P1394b- Hawaii October 13,1998

Max Bassler (Could not attend)
Chairman
mbassler@molex.com

Bill Northey (Acting Chair)
Secretary
NortheyWA@bergelect.com

Agenda

- Review and approve of last meeting's minutes
 - *Approved*
- Review of past presentations
 - Switched 1394 Header
 - Needed for S1600- 2 pins proposed
 - Geometry changes to contacts by BERG
 - New material
 - *Waiting for mechanical test report – Bill Northey*
 - EMI/RFI Testing
 - New version posted

New items

- Presentations
 - S1600 cable speed capability – Bill Northey
- Discussion on electrical requirements for S800-1600
- Develop Ongoing Action Items

Summary of October Minutes

- Proposal for bulk cable testing methodology passed
 - Forwarded to the secretary to put into Framemaker
 - Action Bill Northey to draft this section into Framemaker.
 - Complete- under review-copies at the meeting
 - Revised working proposed about equipment used
- Proposal made for cable detect pin in 6 ckt I/O using an inexpensive switch
 - Voted and reported to Plenary that a switched version will be necessary
- Ongoing proposal to change shape/geometry of existing contacts in the socket
 - Follow-up presentation required on durability and reliability numbers- Pending for next meeting
- Chuck Brill to propose an EMI/RFI test for review
 - Information presented, details under review
 - Chuck to provide more legible copy for review
- Colin Whitby- Strevens has requested that this group provide NEXT information for the 6 ckt I/O system. These are needed to understand total electrical budget. Requested is based on work done in Fibre Channel
 - Based on the draft electrical chapter today we propose the followings changes to meet the NEXT and total electrical budget for connectors and cables:
 - S800 - 200ps risetime (250ps for EMC concerns)
 - S1600 - 100ps (under review)
 - Based on -26Db crosstalk, Skew values in P1394b
 - Total Budget- 80 mv maximum launch (NEXT 40mv, noise 15mv, off transmitter 20mv, margin 5mv)

Agenda for October

- Presentations planned:
 1. Update on modified electrical contact design
 - In testing for mechanical performance
 2. Continued discussion on electrical requirements for S800-1600
 - S1600 cable presentation summary by Bill Northey (Full presentation posted to web site separately)
 1. 6 cables tested @1600Mb/s (3 @4.5 meters, 3 short cables).
 2. Cables were commercially available and made by several manufacturers
 3. All but one cable supported S1600 speeds
 4. 100% cable testing would be required to find failed cables, at least until the manufacturing process were determined to be consistently good
 5. S800 speeds would certainly not be a problem with the tested cables
 - *Question was raised regarding data run length. Bill Northey to determine and report*
 3. Open discussion
 - D. Wooten questioned whether currently defined masks have enough margin to accommodate losses between TP-3 and TP-4. Response is out of scope, but if modified masks are provided, we could accommodate. *D. Wooten to question editor on this point*
 - Each signaling rate should reflect a specification level consistent with its' respective performance level (cost issue)
 - Discussion regarding switches in 6-pin connector

- Connector presence switch NOT needed on 6-pin connector, but is necessary in optical interface (6-pin copper connector can support DC presence test)
- Cable ID requirement unconfirmed @ S800; confirmed @ S1600. Further empirical/theoretical evaluation required for S800 – *Action : Northey/Brunker. Dave B. volunteered to consider the many variables that could be adjusted to make a cable that barely meets the 1394-95 performance requirements. Bill N. indicated that he would be willing to test this cable on the test connectors used to gather data for his presentation today.*
- Colin Whitby-Strevens indicated his feeling that S400 cables probably perform at S800 because of the move from common mode to differential, DC balanced signaling and 8B-10B coding.

Ongoing Action Items

- EMI/RFI Baseline test method
 1. Presentation posted on web.
 2. Call for test method input and review of details.
 3. New version posted on the web, reported by Mike Fogg that there is no IP involved with this technology
- Cable detect pin in copper connector
 - Proposal under review
 - After much discussion it was concluded:
 - S800 may not need a detect device (to be confirm by additional testing)
 - S1600 will need new cable/mod. connector + detect device
 - Impact to pinout/function of connector/cable + PHY
- Bulk cable specification methodology
 - Passed- To be added to P1394b document
 - Document in Framemaker and under review
 - Complete and given to editor
- Matrix of physical interfaces
 - Updated from last meeting
 - Post next meeting
- Connector contact design modifications
 - Awaiting durability/reliability data
 - Additional data to be presented in Hawaii

PHY-Link Interface - Sean Killeen

Though a task group report was not given for the PHY-Link Interface task group, their minutes have been included here for the sake of completeness.

P1394b Phy-Link Interface Task Group Minutes, October 12/13 1998

Agenda

1. Review of previous Minutes
2. Summary of current draft
3. Readthrough of draft material
4. Outstanding Questions/issues
5. Work to be done
6. Electrical update
7. AOB

1. Review of previous Minutes

Taken as agreed

2. Summary of current draft

SK reported that this is not yet included in draft P1394b spec. He presented a proposal for Clause 12 of the P1394b for review. This document was distributed at the meeting and on email immediately prior to the meeting (and is also now on the web site).

It was agreed that for the Monday meeting, SK would present a summary of the draft, and that the Tuesday meeting would then review it (after people had had a chance to read it overnight).

SK identified the following highlights in the draft proposal

Interface pinout

The document proposes Interface Phases to provide extra visibility, but needing no extra control pins

There are new link Reqs

- short form (for common actions)
- extended form - allows unlimited amount of information

Status transfer is relocated and is now not on the data lines

- independent operation (of ongoing packet in either direction)
- short form
- extended form (register reads etc)

So at all times there can be a packet transfer, an Lreq and a status transfer

Interface handover

- safe, synchronized
- "registered" operation
- times to drive the handover phases - note that this is much longer than we are used to, due to use of source clocking (5 to 7 cycles).

There was extensive discussion on this, with an analysis showing why 5 cycles are needed - Grant, Wait, Return, plus two cycles of phase delay. SK will write up a rationale of this.

Interface clocking changes

PHY-Link comes up in a base mode, and then negotiates a transition.

TX/RX packet formats - description of how the padding mechanisms will work.

3. Readthrough of draft material

It was agreed that everyone would have the chance to read the document and go through it at the continuation session on Tuesday.

4. Outstanding Questions/issues

To be discussed on Tuesday

5. Work to be done:-

Fill out the Lreqs
Document status
P1394a operation
PHY register map
Electrical issues
Lreq cancellations
Reset and initialization

6. Electrical update

None - SK to check on status with DC Sessions

7. AOB

None

=====

Continuation on Tuesday 13th October

Review of Clause 12 Proposal.

The document was reviewed page by page, and various issues discussed. The main points of discussion and agreed resolution (where appropriate) is described below.

Page 1

Reference the transfer speeds as S100, S200, etc (as against Mbps)

Discussion of a link with b PHY and vice versa.

Requirement is for an autonomous transfer of status from PHY to link (low latency, and non-interruptible, in contrast to current PHY/Link interface).

Requirement for reset of the interface when the link goes away (c/f LPS in P1394a).

Page 2

Make LPS and LinkOn mandatory on the PHY but not the link (follow the example of P1394a for mandatory and optional signals on the PHY and the Link - of course, for P1394b, it is always an 8 bit data bus).

Do we need 8 pins for a S100 only PHY??? We could save pins here, like 1394-1995. This is an issue. Note that the Data path usage may be more than 2 pins for grant type etc.

Do people like the PINT pin - everyone to review it. (Initial reaction positive, no negatives)

Align names like SDATA, SCTL to 1394a where the function is the same. Perhaps the different names should be BXXXX?

Page 3

Interface reset is not yet defined - suggest use the P1394a definition as a basis (deviate only where necessary). LSCLK issue - propose that the Link keeps the clock to itself, and reception of LSCLK be one of

the requirements for the PHY to consider reset of the interface complete.

Note that it is intended that the clock speeds should be constant between resets (at least).

Lreq is synchronous with LSCLK.

Clock frequency - This is a system design issue (board design). Propose that the link determines the speed and instructs the PHY. The system designer ensures that the selected PHY can accept the speed. If necessary the system designer arranges for implementation dependent mechanism to set the speed as controlled by the link to a speed which the PHY can accept.

LSCLK shall be derived from the PSCLK

There needs to be a spec defining the jitter on the clocks, and the maximum instantaneous frequency difference (or whatever the correct term is).

Page 4

Isolation barrier is required for all signals (not just LINKON).

LMODE, maybe we could save a pin by using LSCLK input for this (say strapped high to indicate a on start-up). Take care on isolated interface.

PHY-Link phase proposal - this seems contentious - so discuss it on the reflector.

Concatenated packets - current proposal skips the Grant, but this requires the Link to know all about the PHY timing. It is better to have some form of Grant for each concatenated packet. Maybe send a grant signal over PINT.

Page 8 - Link requests

Page 9 Lreq format

Maybe three bits of info will be sufficient

Don't need the short/long bit - just use the continue bit for this

This looks overly complex - is this amount of flexibility needed? Suggest to enumerate the number of Lreqs that are really required. Prefer a fixed format Lreq, or at most two formats (as now!).

Page 11

PHY/Link interface handover

SK concluded the meeting with a general encouragement to make comments

He anticipated at least two new drafts before the next meeting.

Present (Monday sign-up list)

David Brunker
Dao-Long Chen
Michele Dale
Steve Finch
Noburo Furuya
Eric Hannah
Jerry Hauck
Sean Killeen (chairman)
Steven D Martinez
Angel Mora
Yoshikatsu Nima
Toshihiro Niitsu
Takayuki Nyu
Ozay Oktay
Matt Rooke
Kevin Ross
Tetsuya Sato
Masood Shariff
Michael Shinkarovsky
Jim Skidmore
Michael Jonas Teener
Dave Thompson
Colin Whitby-Strevens
David Wooten
Patrick Yu

B-Low Power - Steve Bard

Though a task group report was not given for the B-Low Power task group, their minutes have been included here for the sake of completeness.

IEEE P1394b B-Low Power Task Group

MINUTES
October 13, 1998
Sheraton
Maui, Hawaii

1) *Introductions*

Those present introduced themselves

2) *Accept volunteer to write minutes*

Colin Whitby-Strevens was accepted as official note taker for this meeting.

Appropriate expressions of appreciation were bestowed.

3) Approve Previous Meeting Minutes

Accepted without opposition.

4) Presentations

The following presentations are already on the web site

4.1 "D1" Low Power State (Review from Previous Meeting) - [Steve Bard]
Discussion on the requirements for a parent to synthesize a Self-ID packet on behalf of a child leaf in D1. Exactly what goes into this synthesized packet? There may also be issues concerning 1394.1 "second self-ID packet" which will have to be consistent.

4.2 "D1" Standby - Extended Command of PY Command Packet - [Steve Bard]

4.2.1 Entering into "Standby" - [Steve Bard]
Probably need something like TX-SUSPEND from the parent to the leaf node, but it has to be different, so that the resume will NOT assert reset.

4.2.2 Restoring from Standby - [Steve Bard]

4.3 Issues

4.3.1 Do we need a way of precluding a parent port from initiating D1 to inappropriate child ports. (Note, assume that if the Child is Beta capable then it is D1 capable). Jerry Hauck says yes, it should be in silicon because s/w may want to try. **But decided that we will not preclude it. Software must do the right thing.** Considered whether it is a better idea to send the Standby command to the child, and let the parent snoop on this. **Decided that the D1 mechanism will work only over Beta mode connections. Decided that the standby command will be sent to the parent port. Decided that Beta mode only leaf nodes do not need to do Suspend/Resume, they only need to support D1.**

4.3.2 How and where to save the initial Leaf Self-ID

4.3.3 How to update the node-ID of the stored Leaf Self-ID subsequent to a bus reset whilst Leaf is in Standby

4.3.4 How does the parent node synthesize the self-ID of a Leaf node in Standby

4.3.5 Determining parent ports in a node which have child Leaf nodes in Standby

4.3.6 Multi-port with only one connection - is this allowed to be placed in Standby - Answer: No - software shall not issue a Standby command implicating such a PHY. A PHY is a multi-port PHY if it happens to be a leaf node now, but it is possible to be other than a leaf node, then it does not support D1. But, for example if a multi-port PHY is implemented in a target device such that only one port is connectable, then it is a leaf node. Revised proposal:- A port which is currently a leaf node can be placed in D1. If, whilst it is in D1, a connect event occurs so that it is not a leaf node any more, then things come up, a bus reset occurs, and it is no longer a leaf node so it cannot now be placed in D1.

4.3.7 If 4.3.6 is allowed, shall a new connection cause "leaf" port to restore from D1 to D0?

4.3.8 Further, would multi-port nodes with only one active port and all other ports suspended "qualify" for "D1" state? (the suspended ports do not require a self-ID to be proxied).

5 Review AR's

5.1 Steve Bard to write up issue list and post to reflector.

5.2 Steve Bard to bring to the reflector for discussion, the circumstance under which a port on a multi-port node may be placed into Standby.

5.3 Steve Bard to write up resolution proposals for remaining open issues and present to group via reflector e-mail. The goal is to have as many issues resolved by the December (Monterey) meeting as possible.

6 Opens

6.1 Add: how to deal with second self-ID packet for 1394.1

7 Adjourn

IEEE P1394b B-Low Power Task Group
ATTENDANCE LIST: September 11th, 1998, Schaumburg, Illinois

✓	NAME	COMPANY	TELEPHONE	E-MAIL ADDRESS
✓	Bard, Steve	Intel	(503) 264-2923	steve.bard@intel.com
✓	Brunker, Dave	Molex	(630) 527-2622	dbrunker@mlex.com
✓	Chen, Dao-Long	Symbios	970-223-5100 x9461	dao-long.chen@symbios.com
✓	Churchill, Richard	Compaq	(281) 514-6984	richard.churchill@compaq.com
✓	Dale, Michele	Lucent	(610) 712-3294	mzdale@lucent.com
✓	Finch, Stephen	Silicon Systems, Inc.	714-573-6808	steve.finch@tus.ssi1.com
✓	Fuller, John	Microsoft	425-703-3863	jfuller@microsoft.com
✓	Furuya, Nobuo	NEC	(408) 969-2479	nobuo_furuya@el.nec.com
✓	Hamilton, Thomas	AMD	(512) 602-9658	thomas.Hamilton@amd.com
✓	Hannah, Eric	Intel	408-765-4441	eric.hannah@intel.com
✓	Hauck, Jerry	Macro Designs	510-668-1006	jhauck@zayante.com
✓	Killeen, Sean	SSL	+353 1 402 5700	sean.killeen@ssl.ie
✓	Martinez, Steven D.	Honeywell	505-828-5433	steven.martinez@das.honeywell.com
✓	Mora, Angel	CTR	+43 4242 5630 0210	angel.mora@ctr.ac.at
✓	Niitsu, Toshihiro	Molex, Japan	+81 462 61 4500	tniitsu@molex.com
✓	Oktay, Ozay	Canon Information Systems	(949) 856-7180	ozay_aktay@cissc.canon.com
✓	Ross, Kevin	Philips	(408) 991-4836	kevin.ross@sv.sc.philips.com
✓	Saito, Kyozo	Alps	+81 229 23 5111	kyozo_s@gw3.alps.co.jp
✓	Sato, Tetsuya	Sony	+81 3 5448 4603	tetsuya@sm.sony.co.jp
✓	Skidmore, Jim	Texas Instruments	972-480-2094	j_skidmore@ti.com
✓	Smith, Michael A.	Control Net	(408) 341-1428	michael@controlnet.com
✓	Soderstrom, Ron	IBM	(507) 253-6290	y2rons@us.ibm.com
✓	Teener, Michael Johas	Zayante	(831) 461-4901	mike@zayante.com
✓	Thompson, David	Lucent	(610) 712-2730	davethompson@lucent.com
✓	Whitby-Strevens, Colin	ST Microelectronics	+44 1454 611500	colin.whitby-strevens@st.com
✓	Wooten, David	Compaq	(281) 518-7231	david.wooten@compaq.com
✓	Yoshikatsu, Niwa	Sony	+81 3 5448 4603	niwa@sm.sony.co.jp
✓	Yu, Patrick	NEC	408-588-5436	patrick_yu@el.nec.com

Those attending the IEEE P1394b Plenary session:

✓	Name	Company	E-Mail address	Telephone
✓	Bard, Steve	Intel	steve.bard@intel.com	(503) 264-2923
✓	Brunker, Dave	Molex	dbrunker@molex.com	(630) 527-2622
✓	Churchill, Richard	Compaq	richard.churchill@compaq.com	(281) 514-6984
✓	Dale, Michele	Lucent	mzdale@lucent.com	(610) 712-3294
✓	Finch, Stephen	Silicon Systems, Inc.	steve.finch@tus.ssi1.com	714-573-6808
✓	Fuller, John	Microsoft	jfuller@microsoft.com	425-703-3863
✓	Furuya, Nobuo	NEC	nobuo_furuya@el.nec.com	(408) 969-2479
✓	Glenski, John	Rockwell Collins	jfglensk@collins.rockwell.com	(319) 295-9684
✓	Hamilton, Thomas	AMD	thomas.hamilton@amd.com	(512) 602-9658
✓	Hannah, Eric	Intel	eric.hannah@intel.com	408-765-4441
✓	Hauck, Jerry	Macro Designs	jhauck@zayante.com	510-668-1006
✓	Killeen, Sean	SSL	sean.killeen@ssl.ie	+353 1 02 5700
✓	Martinez, Steven D.	Honeywell	steven.martinez@das.honeywell.com	505-828-5433
✓	Mora, Angel	CTR	angel.mora@ctr.ac.af	+43 4242 5630 0210
✓	Nagao, Shuji	HITS	nagao@hits.co.jp	+81 462 77 9467
✓	Niitsu, Toshihiro	Molex, Japan	tniitsu@molex.com	+81 462 61 4500
✓	Northey, Bill	Berg Electrnics	northewa@bergelect.com	717-938-2119
✓	Nyu, Takayuki	NEC	new@optsys.cl.nec.co.jp	+81 44 856 2082
✓	Prouty, Bill	HP	bprouty@hp.com	916-785-4631
✓	Rooke, Matt	Fujitsu (FCPA)	mrooke@fcpa.fujitsu.com	408-894-3826
✓	Ross, Kevin	Philips	kevin.ross@sv.sc.philips.com	(408) 991-4836
✓	Saito, Kyoza	Alps	kyozo_s@gw3.alps.co.jp	+81 229 23 5111
✓	Sato, Tetsuya	Sony	tetsuya@sm.sony.co.jp	+81 3 5448 4603
✓	Shariff, Masood	Lucent	mshariff@lucent.com	(732) 957-5479
✓	Shinkarovsky, Michael	Lucent	mshnkarovsky@lucent.com	(610) 712-2938
✓	Skidmore, Jim	Texas Instruments	j_skidmore@ti.com	972-480-2094
✓	Smith, Michael	ControlNet	michael@controlnet.com	(408) 341-1428
✓	Soderstrom, Ron	IBM	rv2rons@us.ibm.com	(507) 253-6290
✓	Teener, Michael Johas	Zayante	mike@zayante.com	(831) 461-4901
✓	Thompson, David	Lucent	davethompson@lucent.com	(610) 712-2730
✓	Toguchi, Satoru	HITS	wing@hits.co.jp	+81 462 77 9460
✓	Whitby-Stevens, Colin	ST Microelectronics	colin.whitby-stevens@st.com	+44 1454 611500
✓	Wooten, David	Compaq	david.wooten@compaq.com	(281) 518-7231
✓	Yoshikatsu, Niwa	Sony	niwa@sm.sony.co.jp	+81 3 5448 4603
✓	Yu, Patrick	NEC	patrick_yu@el.nec.com	408-588-5436