

P1394B Working Group Meeting - Plenary  
September 11, 1998  
Shaumburg, Illinois

1. Review Agenda
  2. Review of Minutes of August meeting
  3. Procedures
    - 4.1. Voting
    - 4.2. Price/Pricing
    - 4.3. Call For Patents
  4. Meeting Schedule
    - 4.1 October 12 & 13, Maui, HI
    - 4.2 December 7 & 8, Monterey, CA
    - 4.3 February 3 & 4, Newport Beach, CA
    - 4.4 March, April (Need invitations)
  5. Abbreviated Task Group Reports
    - 5.1. S3200 in spec.
    - 5.2. Additional IRM Support
    - 5.3. Maximum Packet Size
  6. Task Group Reports (As required)
  7. Review of Action Items
  8. Adjournment
- 
- 

Mike Teener moved to accept August meeting minutes, Len Young seconded. Minutes approved without objection..

Voting, price/pricing, Call for Patents, etc. Yadda, yadda, yadda...

Worked through meeting schedule. No immediate sponsors for March and April.

#### S3200 in Specification

Brad Saunders moved to remove S3200 from the 8B10B (connector side) side of the electrical parameter portion of the specification (anything other than protocol) specifically, clauses: 4, 5, 6, 7 as contained in draft 0.13.

Group consensus during discussion was to NOT loose the content after removal from the spec. Not clear as to how the content will (and where) be preserved.

11 in favor, 4 opposed, 4 abstains

Motion carries.

## MAXIMUM PACKET SIZE

Shall packet size scale with speed? Isochronous? Asynchronous? Shall they track simultaneous or not?

IEEE P1394a draft 2.0 states packet size shall scale with speed.

A motion (if proposed) would be to NOT scale with speed.

John Fuller moved that asynchronous packet data payloads be limited in size to no larger than 2K. Jerry Hauck seconded.

Steve Bard requested a friendly amendment to the motion to replace 2K with 4K bytes. John Fuller accepted the friendly amendment. Jerry Hauck agreed to include the amendment in his second to the motion.

14 in favor, 2 opposed, 6 abstain

Jerry opposed the motion because he has not heard any data which explains why two 2K packets are not acceptable.

Motion Passes.

Colin Whitby-Strevens suggested P1394b liaison with the P1394a Ballot Review Committee to comprehend this change in the IEEE P1394a specification.

John Fuller moved to limit isochronous packet data payload be limited to no larger than 4k bytes. Motion dies for lack of a second.

Steve Bard requested of the chair for open discussion of the merits and demerits of limiting isochronous packet data payload. Chair agreed to a 20 minute discussion.

After 10 minutes of discussion, Colin Whitby-Strevens moved to endorse table 9-1 in IEEE P1394a Draft 2.0. Seconded by Mike Teener.

17 in favor; 0 opposed; 4 abstain. Motion carries.

## Additional (enhanced?) IRM support

Mike Teener presented Richard Churchill's proposal on IRM enhancements... Richard has sent a second "chapter" in the proposal on the reflector - it has been included in these minutes for the sake of completeness:

---

*Subject: Quality of Service and IRM Transactions, Part 2*

*Gentlefolk,*

*This is the next installment of the discussion started with part 1 of this series ... The reader is again referred to the document "Quality of Service and IRM Transactions," located via the p1394b web page, as well as part one of the present discussion contained in the minutes to the August p1394b plenary meeting.*

*Continuing with the discussion of proposal 1, which would require any implementation of the BANDWIDTH\_AVAILABLE and CHANNELS\_AVAILABLE registers in such a manner that requests to these registers could produce unified transactions, and the alternative that these registers be implemented in hardware when present in a node.*

*In previous discussions of the recommendation of rules for implementing the BANDWIDTH\_AVAILABLE and CHANNELS\_AVAILABLE registers, there is a supposition that it is possible for the IRM to respond to a request to either of these registers as a unified transaction ... with the response concatenated to the ack returned by the IRM node in reply to the original request. Though this is the case for -1995 and 1394a compliant nodes, it is not for p1394b nodes in p1394b only environments. This difference results from the efficiencies we anticipate gaining from the use of the "return channel" arbitration schemes that have been discussed at various p1394b and accelerations task group meetings over the last several months. These accelerations should certainly result in the ability to grant the bus to another arbitrating node before a request to an isoch. resource register on the IRM can be fully decoded and a response generated, even if these registers are implemented in hardware, since arbitration for the bus occurs during the transmission of the current packet.*

*So, how does this relate to proposal 1 as it stands? Well, as per the preceding discussion, though improvements in 1394a behaviors (specifically the removal of the requirement for an initial read of the isoch. resource registers), this is beyond our scope, and entirely in the hands of the voters and ballot response committee. When discussing mixed -1995/a/b environments, the behaviors are as yet not well defined, and so are difficult to model and thus to discuss. In pure 'b' environments, the idea of implementing these isoch. resource registers in hardware still applies, and should help reduce the interval between a request and completion of the requested action during which an overlapping request may be received... assuming an implementation that does not insert other hurdles between iso-resource requests and responses... and thereby help decrease the size of the window during which one or more prior requests can block subsequent ones.*

*Another method whereby the problems attendant to overlapped requests may be prevented or reduced is to implement IRM capable nodes with sufficient receive buffer space to allow a reasonable number of anticipated transactions to be queued. Well, if some people don't like the idea of adding a fist full of gates in order to implement these registers in hardware, what will they think if we tell them to significantly increase receive buffer sizes?*

*Returning to the allocation example in the previous installment of this discussion, if we assume the requests are dealt with via unified transaction, and use roughly the same sequencing of transactions, three of the four allocating nodes would have successfully done so by the end of the fifth fairness interval, and the fourth would succeed in the next (seventh) interval. If instead we assume that the response to each request occurs quickly enough that only one other request can be received while a prior request is pending, two nodes would have succeeded in completing allocations, with the remaining two following quickly, without interference from the two that are finished. These are clearly better results than the original example, where NONE of the nodes succeeded in allocating bandwidth in six fairness intervals.*

*Now, just in case anyone thinks this is a non-issue, let me add another example where significant delays in completing isoch. reallocations would be undesirable, let me offer another case where such delays would be bad. Consider any environment where 1394 is used for process control or continuous status monitoring. We could achieve a reliable data communication using isochronous channels for an instrument to feed*

updated readings to devices using that data at an 8,000 Hz rate. A failure to reallocate resources for this logical connection in a timely manner could cut off the follow of critical status information for a second or more. A specific example is that there has been discussion of the use of 1394 for connecting avionics units in aircraft. Would you really want to risk the potential effects of a loss of access to avionics status for that much time in flight, particularly during take-off or landing? Before you immediately discard the idea of such a solution, consider that few if any of us would have thought of using the backplane model implementation in satellites, but a few months ago we were informed that at least one company has done just that.

In the final analysis, we need to find some means of keeping the risk of pathological behaviors resulting from allocation efforts failing due to the IRM being busy. In 'a' environments, we can require that the IRM respond to such requests via unified transactions, but not in 'b' environments. Also, some people believe this would mean substantial added complexity and expense for IRM capable nodes. Requiring the registers to be implemented in hardware helps the situation by allowing quicker responses, but does not eliminate the problem in 'b' environments. Therefore, regardless of whether we choose to require some specific method of implementation, we need to strongly encourage designers to consider how to make the handling of isochronous resource allocations as quick as possible.

Now for the next proposal ...

Proposal 2 was that p1212r should define a "bounded\_allocate" lock transaction, which would behave along the lines of "new\_value = (data\_value <= old\_value) ? (old\_value - data\_value) : old\_value;" Such a transaction type would allow the allocation of quantitative resources via a single transaction without having to first determine the amount available ... If enough is available, the transaction completes via allocation of the desired quantity and return of old\_value, etc., in the response.

At the June meeting of the p1212r working group (held in St. Petersburg Beach, FL), a rather better version of this was agreed to, in two forms: big-endian and little endian. The behavior (from memory) is "new\_value = (arg\_value <= old\_value) ? (old\_value - data\_value) : old\_value;" with arg\_value >= data\_value. This allows an application to allocate a quantitative resource while leaving a "reserve" quantity for other applications. (For exact details, consult the p1212r June minutes.)

This change to p1212r sets the stage for decisions on proposals 3, 4, 9 and 10 ...

Proposal 3 is simple. It proposes that we add a requirement that IRM capable p1394b compliant nodes support bounded\_allocate lock transactions upon the BANDWIDTH\_AVAILABLE register, in addition to the presently required quadlet read and compare\_swap lock transactions. This accomplishes at least three things: it eliminates the need to perform a quadlet read, regardless of what happens in the p1394a balloting process, it guarantees that a single transaction will successfully allocate bandwidth, provided sufficient bandwidth remains available and the request is successfully received by the target, and it practically assures that bandwidth allocation will take place as an order N process, instead of an order N<sup>2</sup> one.

In order to achieve an order N<sup>2</sup> process using bounded\_allocate lock transactions without prior reads, the delay between issuance of a request and the corresponding response would need to be long enough to allow nearly all of the allocating nodes to issue requests during that delay. As the window of vulnerability decreases in size we reach an order N process, and if we adopt some form of proposal 1, or any other means of assuring prompt response to requests these requests, we approach simple N transaction process.

Returning to the previous pathological case study, due to the spacing of the commencement of the allocation process by the nodes, so long as the IRM responds in anything resembling a reasonable interval, allocations would be completed in four fairness intervals, using only four bounded\_allocate lock transactions. However, if we had all four nodes attempting allocation starting in the same fairness interval, with responses requiring to the start of the next fairness interval, we would have an order N<sup>2</sup> process of something like N(N+1)/2 transactions required, worst case. Reducing the response latency for requests increases the chance that more than one allocation would complete during a given fairness

*interval. If each request can be serviced before the next arrives, the four nodes complete allocation in four requests.*

*Proposal 9 is simply the procedural change required to allow use of the bounded\_allocate lock transaction in manipulating the BANDWIDTH\_AVAILABLE register. If we accept proposal 3, we must accept proposal 9 in order to make "compliant" use of it.*

*Proposal 4 is a further simplification of management of the BANDWIDTH\_AVAILABLE register, recommending that we also require nodes with this register to support fetch\_add lock transactions to this register. The benefit of this is admittedly minor. It allows the deallocation of bandwidth to be performed as a single transaction that will always complete, provided the target node isn't busy, and the transaction doesn't fail due to some bus or node fault. In other words it makes the bus slightly more efficient, while reducing the already slim chances that a deallocation might be obstructed by an allocation and vice versa. This change is consistent with the general intent of the whole group of proposals, which is to stream-line resource management and reduce or eliminate opportunities for pathological behavior, and so I believe it to be appropriate to accept it.*

*Proposal 10 corresponds to proposal 9, in that it is simply the procedural end of the changes to the 1394 standards needed to allow vendors and their applications to use the facility provide by proposal 4 and remain explicitly compliant with the standards regarding methods for managing bandwidth.*

*I will continue with discussions of proposals 5, 6, 7 and 8 at a later time...*

*Sincerely,*

*Richard Churchill,  
(281)514-6984,  
richard.churchill@compaq.com*

---

The first proposal was to use "allocate" for the bandwidth available (BW\_AVAIL) register.

The second proposal was to use the already defined mask swap (MASK\_SWAP) for the channels available (CH\_AVAIL). Multiple channels can be allocated with one MASK\_SWAP. If a request for two channels is made and one channel is already taken, no channels will be allocated. Both proposals have the attribute that they will succeed if everything runs right. Note: The suggestion is two use these proposals, however, the COMPARE\_SWAP must still be supported by the IRM.

The third method would be defined channel allocation procedure.

The fourth would support "unified" transactions with IRM registers

David Wooten states that "unified" in quotes does not mean it is required... the intent is to get the IRM registers implemented in the hardware it does not mean you always have to perform a unified transaction.

An "official" straw poll was taken to determine if some folks should be encouraged to go off and write this up and have it included in the spec;

11 in favor, 0 opposed, 7 abstains

Another "official" straw poll: How many are in favor of making the first two items (use "allocate" for BW\_AVAIL and use "MASK\_SWAP" for CH\_AVAIL) a requirement to be included in the 1394b spec for IRMs.

12 favor, 0 opposed, 10 abstains

Team chosen (volunteered?) to draft the words for the spec: Mike Teener, Richard Churchill, Jerry Hauck

### TASK GROUP REPORT OUT

#### ***Copperheads - Max Bassler***

NEXT - Near end cross talk for cable/connector requirement. Colin Whitby-Strevens has requested that Copperheads provide NEXT information for the 6 circuit I/O system. These are needed to understand total electrical budget. Request is based on work done in Fibre Channel.

Based upon the draft electrical spec for B today, Copperheads propose the following changes to meet the NEXT and total electrical budget for connectors and cables:

S800 - 200 picoseconds risetime (250 picoseconds for EMC concerns)

S1600 - 100 picoseconds (same in draft)

All of this is based on -26dB crosstalk, Skew values in P1394b and a Total Budget of 80 millivolt maximum noise (NEXT 40 millivolt, noise 15 millivolt, off transmitter 20 millivolt, margin 5 millivolt)

This will be remanded to the Electrical Task group for dispensation.

A second point:

Cable detect pin in copper connector/cable (S800-1600) has a proposal under review. User feedback has been requested by Alps to allow for high speed operation with legacy performance. David Wooten requested at the August meeting that Copperheads to determine whether a new cable + detect mechanism is needed - results by the October meeting.

Copperhead report on the conclusion is (after much discussion in the task group):

S800 may not need a detect device (to be confirmed by additional testing).  
S1600 will need new cable construction with the modified connector + detect device.

Impact to pinout/function of modified connector/cable + PHY.

Probable decision point: two pins → 4 degrees of freedom.

Pin 1 → 0 or 1

Pin 2 → 0 or 1

Four possible combinations.

This report has been passed to the Upstarts group to insure they do not upstart sooner than the cable will allow them.

**UTP 5 - Colin Whitby-Strevens**

Received an interesting presentation by L. Chen regarding support for a particular equalization scheme. Details (specific task group meeting minutes) can be found later in these minutes.

Motion to delete 50 meters passed. David Wooten made motion to delete requirement to . Seconded by Brad Sanders. Tabled to next meeting (after some discussion).

Liason report from TIA TR-41 regarding bundled fibers. Recommendation is that plenary chair reply to TIA TR-41 that their application does not apply to 1394b.

**Upstarts - Colin Whitby-Strevens**

The most important aspect of the open issues is toning into TpBias. The risk seems to be limited to a scenario in which both ends are just powering up. Further investigation required for confirmation.

Signal\_detect electrical specification proposal has made it into the latest P1394b draft.

**Simulations Task Force - Jerry Hauck (via Colin Whitby-Strevens)**

Only one bug found. Work is well in hand. Useful timing simulations (drawing) are being presented which are much appreciated.

*Included here are the comments from Jerry Hauck regarding the simulations discussions. This has been included in the minutes for the sake of completeness:*

*1) set\_beta sends the speed code before checking the received speed code. Consequently, we often see an unnecessary additional speed exchange. (Sequence: local node sends speed and then waits for the*

next interval. In the meantime, the peer node sends its speed which is recognized by the local port. However, the local port starts the next interval by sending the speed tone and then looks at the received speed. If we checked the received speed first, we could immediately go into the acknowledgement or downshift phase.)

2) The current code reverts to continuous toning after it's one time trial of bias. Since we now rely on `dc_connected` to tell us if we have a dc path to the peer node, this continuous toning is no longer required. (It was once required when we didn't have a dc connection indication and had to handle passive ac connects. Note that we are absolutely relying on the `dc_connected` signal telling us we have a dc path all the way to the peer port.) After toning into a DC connection and then trying BIAS, we know the far end is powered-off. We can then go silent and when the peer node becomes powered, it will initiate communication with either a tone or `TpBIAS`. Dave Wooten pointed this out, but then had an argument (which I didn't capture) for why he didn't care if we wasted the power toning into the powered-down device. As a result, the code may stay as is.

3) We had a long discussion about toning into `TpBias`. The current code has instances in which we tone for 666.67 us over `TpBias`. This is unavoidable since we can't hear bias when we are toning ... so it is bound to happen. One concern is that we may cause shifts in the peer's `TpBias` and, if it is an old PHY with a single `TpBias` generator, this may corrupt traffic repeating through that PHY. It was observed, however, that we could limit the problem to the "rising edge" of the remote bias assertion by diligently checking for BIAS before going into toning. Since the rising edge only occurs when the remote PHY is just powering up, there should be no traffic corruption. In summary, we need to check carefully for BIAS before starting to tone. This includes making sure our BIAS comparator is stable after our own POR before trying the tones.

The current C code needs to be updated ... we currently look for bias and then call `receive_signal_ok()` to look for a tone which consumes ~666 us. If no tone is heard, we immediately send a tone. So there is a 666 us window in which `tpBias` could appear and we simply aren't listening.

### **PHY/Link Interface - Sean Killeen**

Bulk of time looking at a proposal on a pin out for signaling. PHY designers have an outstanding AR to examine the feasibility of S3200. Three conclusions for the appropriate data path and clocking scheme for S800 and S1600. Work will be coordinated with the accelerations and b-port group to come up with the procedures and protocols to be used for the interface.

Detailed task group minutes have been included in these minutes.

Based upon a question from Colin Whitby-Strevens, the plenary chair tasked Sean to prepare a matrix of the actions the task group has chosen to take for the purpose of ratification by the plenary. Sean presented the following for

ratification by the plenary for these decisions in developing: data path, clock rate, edge, usage models:

S800	8-bit data	100 MHz	Single edge
S1600	8-bit data	200 MHz	Single edge
S3200	8-bit data	200 MHz	Double edge

Brad Saunders proposed an amendment to the task group model which would eliminate S3200 - accepted for consideration by plenary chair.

Group straw poll taken for consideration of the proposed ammendment:

1 in favor; 9 opposed; 2 abstained; Amendment proposal fails.

Group straw poll taken for reaffirmation by plenary of the chosen path taken by the task group:

11 in favor, 0 opposed, 2 abstain. Plenary ratifies task group direction.

### **ATM Liaison Letter - David Wooten**

David Wooten presented a liaison letter from ATM forum regarding their desire to route FUNI over 1394. They sought the opinion the IEEE 1394b working group regarding this topic. It was decided that the chair will contact them and suggest they may find it more appropriate to execute a liaison with the 1394 Trade Association (the Technical Officer specifically) as the IEEE Group does not find the subject of any specific concern or interest. The letter has been included in these minutes for the sake of completeness.

---

The following are Task Group Minutes submitted to the secretary for inclusion in the plenary minutes (these) as well as items of an interest (or not) which were discussed during the plenary session:

ATMF/98L-059

**The ATM Forum****September 9, 1998****Subject: Native ATM Services over ATU-R**

**David Wooten, Chairman IEEE P1394B Working Group  
Compaq Computer Corporation  
PO Box 692000 MS080701  
Houston, TX 77269-2000**

**Dear David,**

The ATM Forum Technical Committee has started a new work item to investigate using the FUNI protocol across Ethernet, USB, and IEEE 1394 links. This is being done to extend Native ATM Services to a personal computer, in the context of attachment to an ADSL access network, via an ADSL modem external to the personal computer.

We invite your comments about this work item. Our next meetings are scheduled for October 5-9 and November 30 - December 4.

Should you desire additional information, please feel free to contact me (tel +1 732 332 5174, fax: +1 732 949 1196, email: gratta@lucent.com).

Sincerely,

**G. Ratta, Chairman – The ATM Forum Technical Committee**

---

This document includes information on work in progress within The ATM Forum. The information contained is subject to change after more study. It is offered for discussion in the interest of maximizing the availability of interoperable equipment.

---

**The ATM Forum**

Worldwide Headquarters  
2570 West El Camino Real, Suite 304  
Mountain View, CA 94040-1313 USA  
Phone: +1.650.949.6700  
Fax: =1.650.949.6705

**The ATM Forum**

Europe Office Asia  
Avenue de Tervueren 402  
1150 Brussels, Belgium  
Phone: +32.2.761.6677  
Fax: =32.2.761.6679

**The ATM Forum**

Pacific Office  
Hamamatsucho Suzuki Building 3F  
1-2-11, Hamamatsucho, Minato-ku  
Tokyo 105-0013, Japan  
Phone: +81.3.3438.3694  
Fax: =81.3.3438.3698

PHY Link Interface Task Group  
September 10th 1998, Chicago  
Minutes

Executive summary:-

Voting in the meeting selected the following data-path width and clocking schemes for each maximum speed:-

S800	8 bit datapath	100MHz clock	Single edge clocking
S1600	8 bit datapath	200MHz clock	Single edge clocking
S3200	8 bit datapath	200MHz clock	Double edge clocking

Agenda

1. Welcome and Sign-up
2. Previous Minutes
3. Previous Actions
4. Presentation of pinout proposal & discussion
5. Electrical Issues
6. Roadmap
7. Homework - Quad-clocking, S3200 (& beyond)

1. Welcome & sign-up
2. Previous Minutes  
As distributed
3. Previous Actions

Designers were asked to review whether 400MT/s over 8 bits is reasonable with 100MHz clocking.

Conclusion - 400Mt/s over 8 bits is not a good way to go.

All were asked to consider whether to retain a S3200 definition.

Conclusion - Keep at least a S3200 specification.

4. Presentation of pinout proposal

SK introduced a new document (to be placed on the P1394b web site) on this topic.

There was extensive discussion on the table on page 2 on datapath width, speeds and clocking regime. This table provides a range of options for each maximum speed.

It was agreed that the standard should specify a single mechanism for each maximum speed.

For S800: DW proposed, MT seconded that we adopt Interface clock speed 100MHz, Single Edge clocking, 8 bit Datapath width.

Accepted without opposition

For S1600: DW proposed, DC Sessions seconded 8 bit datapath width, 200MHz Single edge clocking

Before voting, it was agreed to take a 3-way straw poll between three options

8 bits - 100MHz - Double edge clocking:- 2

8 bits - 200MHz - Single edge clocking:- 9

16 bits - 100MHz - Single edge clocking:- 1

Voting on the motion For: 13 Opposed: 1 Abstain: 5

Eric Hannah proposed and Dave Wooten seconded: To strike all the 32 bit solutions

Approved unanimously.

The S3200 option for 16 bits - 100MHz - Double edge clocking was eliminated by consensus.

Mike Teener proposed, Dave Wooten seconded: Strike the remaining S6400 specification.

Brad Saunders moved to table, Sean Killeen seconded

In favour: 6 ; 7 opposed

Motion to table fails

The question was called

For striking 6400 - 14, Opposed, 0, 3 abstentions

There was then discussion on the options for S3200 - including consideration of 400MHz clocking

Accept 8bits, 200MHz, Double edge clocking

Keith Heilman proposed, JH seconded to accept for S3200 8bits, 200MHz, Double edge clocking

Before taking the vote, it was agreed to take a 3-way Straw poll  
- including the 8bit datapath - 400MHz - Single Edge Clocking option

8 bit datapath - 200 MHz clock - Double edge clocking: 11  
16 bit datapath - 200 MHz clock - Single edge clocking: 1  
8 bit datapath - 400 MHz clock - Single edge clocking: 2

Voting on the motion: For: 14, Opposed: 1, Abstain: 3

Dave Thompson recorded that his opposition was because he believed it to be technically infeasible

Sean Killeen then introduced the proposed pin/signal description

Areas of concern which were identified included  
- number of bits of Sctl[0:1]  
- issue of LinkOn  
- use of Pstatus and Lstatus (maybe wrap these back onto the data lines).

In further discussion it was noted:-

We need a list of services and indications that are required before we can proceed much further with the details of the interface.

The Bport/accelerations group will try to provide one in the near future (in time to be worked on before the next meeting)

We need to specify a way of the PHY distinguishing between a 1394-1995/1394a link and a 1394b link

There was discussion on the clocks - the Link clock is frequency locked to the PHY clock, the phase relationship is undefined but static (except for jitter).

An issue was identified on interface length: time of flight vs symbol rate, and danger of reflections if there are multiple symbols in flight.

5. Electrical Issues  
Deferred

6. Roadmap  
Deferred

7. Homework - Quad-clocking, S3200 (& beyond)

Covered in previous agenda items

Present

Rich Bowers  
Dao-Long Chen  
Dave Eicher  
John Fuller  
Thomas Hamilton  
Eric Hannah  
Jerry Hauck  
Keith W Heilmann  
Sean Killeen  
Joe Kryzak  
Ozay Oktay  
Brad Saunders  
DC Sessions  
Michael A Smith  
John Smolka  
Michael Jonas Teener  
Dave Thompson  
Colin Whitby-Stevens  
David Wooten

## P1394b UTP5 Task Group Minutes September 10 1998

### Agenda

1. Welcome, Introductions
2. New ideas on UTP5  
Presentation by Dao-Long Chen
3. Motion to delete the 50m Option  
Proposed by Ozay Oktay, Seconded by Dave James  
Tabled at the last two meetings
4. Review of the current specification
5. Liaison from TIA TR-41 on bundled fibres

### 1. Welcome, Introductions

Those present introduced themselves. Apologies were received from Eric Hannah and Alistair Coles.

### 2. New ideas on UTP5

CWS summarised the proposal for pre-emphasis (equalisation at the receiver), and referred to the email from Alistair Coles on the relative merits of equalisation at the transmitter and receiver. CWS stated that the current situation is that P1394b requires the receiver to compensate for frequency-dependent attenuation in a worst case cable - the specification is similar in this respect to 100BASE-T ethernet and requires little further work. The issue under discussion is whether to stick to this, or investigate other compensation schemes, which would require much more work to complete. It was reported that the chairman of 802.3 has indicated that the 100BASE-T was a compromise of various considerations, and that alternative approaches may well prove just as suitable or more so.

There was extensive discussion.

Dao-Long Chen presented an implementation of an adaptive equaliser in the receiver (presentation shortly to appear on the P1394b web site). This showed a complexity of 5K to 7K gates, plus a 5 bit ADC (possibly 4 bit for the P1394b binary encoding). CWS reported that he had received information on two commercial designs, where adaptive equalization occupied about 10% of the port in terms of area.

DC Sessions pointed out that some very simple pre-emphasis may provide enough signal conditioning to allow passive equalisation at the receiver.

Dave Wooten reminded the group that our working assumption is that

the standard electrical (short haul) transceivers would be used. In particular, he was concerned if the UTP specification required specialised and incompatible silicon.

After further discussion, a straw poll was taken on "We welcome further contributions on alternative equalisation schemes" This passed, 8 For, 4 Opposed.

### 3. Motion to delete the 50m option

In discussion on the motion, Dave Wooten suggested that perhaps it would be better to allow the length supported to be implementation dependent.

The motion passed For: 18; Opposed: 0; Abstain: 5

Dave Wooten proposed the motion:

To delete the requirement that all receive circuits must work over all lengths of cable up to 100m. This was seconded by Brad Saunders.

After discussion, it was agreed by consensus to table the motion until the next meeting, and that further discussion should take place on the reflector.

### 4. Draft review

This item was deferred due to lack of time and absence of the editor. Note that the latest specification is not in the latest P1394b draft, but is separately available on the web site.

### 5. Liaison from TIA TR-41 on Bundled Fibres

Masood Shariff clarified that the requirement limited cross-talk between bundles to 30mV pk-pk. CWS pointed out that currently our total noise margin is 30mV pk-pk (15 mV amplitude in 1394 terminology). It was agreed to liaise back to TR-41 that this is not suitable for our application, as we need noise margin for noise from other sources within a total 30mV budget.

Present

Steve Bard  
Max Bassler  
Rich Bowers  
Charles Brill  
Dao-Long Chen

Dan Colegrove  
Michael Coletta  
Thomas Hamilton  
Jerry Hauck  
Keith W Heilmann  
Daisuke Hiraoka  
Sean Killeen  
Bill Northey  
Ozay Oktay  
Mark Richman  
Tomoki Saito  
Brad Saunders  
D.C. Sessions  
Masood Shariff  
Michael A Smith  
John Smolka  
Michael Jonas Teener  
David Thompson  
Colin Whitby-Stevens  
Davie Wooten  
Len Young  
Frank Zhao

## P1394b Upstarts Task Group Minutes September 10 1998

### Agenda

1. Welcome, Introductions
2. Simulation Task Group feedback
3. Review of latest code - v8 (in P1394b Rev 0.13)
4. Review of open issues

1. Welcome, Introductions  
From the immediately preceding UTP5 meeting

2. Simulation Task Group feedback  
Jerry Hauck presented the latest results  
(document to be placed on the web site)  
This comprised 1 bug, 5 new technical issues, various editorial issues,  
and simulation with speed negotiation.  
The bug concerns centering in `receive_speed_indications()`.  
The technical issues were
  - specification of min/max values
  - `signal_detect_OK` now simply incurs an unneeded delay
  - gap detection is not futureproof, and could be improved
  - the interval between start bits (keepalive tone) varies slightly  
depending on circumstances.

CWS agreed to review the code to correct the bugs or raise concerns,  
and make the editorial improvements.

The simulation results reveal one or two areas where further  
optimisation is possible.

It also reveals an occurrence of toning into `TpBias`. It was agreed  
that this particular circumstance is an artefact of the simulation.  
Concern was raised that when toning into `TpBias` does  
occur, then this could have an impact on other ports on the peer  
device (as the `TpBias` level would, in general, change, which could  
affect, for example, speed signalling).

It was speculated that toning into `TpBias` can only occur when both  
units are powering up, in which case it is more benign. It was agreed  
that the code should be changed to reduce the chances of toning into  
`TpBias`, and that further checks should be made as to whether it can  
occur at times other than when both ends are powering on.

3. Review of latest code - v8 (in P1394b Rev 0.13)

Resolved issues now in the code:-

- Resume latency (constant RESUME\_SAMPLING\_INTERVAL = 1.33 ms)
- Tone frequency between 48MHz and 64MHz
- Centering during speed detection
- Gap detection for start of speed signal corrected and parameterized
- Chatter code (see also below)
  - Updated signal\_detect spec proposal in P1394b Draft 0.13
- Use connect\_detect comparator during suspend if there's a DC connection
  - indicate resume with a continuous tone
  - disconnect only if no tone when the connect\_detect comparator indicates no connection
- ST reports power consumption likely to be the same either way
  - but other inputs required
- Reflector discussion confirming that we need to engage in one more round of toning after detecting incoming TpBias
- Miscellaneous tidy-up

#### 4. Review of open issues

CWS briefly reviewed the 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%
- Toning into TpBIAS (issue for electrical group)
- Signal\_detect electrical spec
  - Updated signal\_detect spec proposal in P1394b Draft 0.13
  - True if Vinput above min receiver sensitivity and receiving a tone or valid 8B10B characters
  - False if Vinput < 80mV
  - Max Off transmitter should be 20mV
  - Unspecified otherwise, but filtered for chatter
  - 100 microsecond response time
  - Concerns with respect to use in optical connection??

Present

- Steve Bard
- Max Bassler
- Rich Bowers
- Charles Brill

Dao-Long Chen  
Dan Colegrove  
Michael Coletta  
Thomas Hamilton  
Jerry Hauck  
Keith W Heilmann  
Daisuke Hiraoka  
Sean Killeen  
Bill Northey  
Ozay Oktay  
Mark Richman  
Tomoki Saito  
Brad Saunders  
D.C. Sessions  
Masood Shariff  
Michael A Smith  
John Smolka  
Michael Jonas Teener  
David Thompson  
Colin Whitby-Strevens  
Davie Wooten  
Len Young  
Frank Zhao

## Those attending the IEEE P1394b Plenary session:

✓	Name	Company	E-Mail address	Telephone
✓	Bard, Steve	Intel	<a href="mailto:steve.bard@intel.com">steve.bard@intel.com</a>	(503) 264-2923
✓	Bassler, Max	Molex	<a href="mailto:mbassler@molex.com">mbassler@molex.com</a>	(630) 527-4490
	Bommji, Anuradha	Sand Microelectronics, Inc.	<a href="mailto:anuradha@sandmicro.com">anuradha@sandmicro.com</a>	408-321-8200 x24
	Bowers, Rich	Quality Semi	<a href="mailto:rich_bowers@qualitysemi.com">rich_bowers@qualitysemi.com</a>	408-450-4562
	Brill, Charles	AMP	<a href="mailto:cebrill@amp.com">cebrill@amp.com</a>	(717) 810-4642
✓	Chen, Dao-Long	Symbios	<a href="mailto:dao-long.chen@symbios.com">dao-long.chen@symbios.com</a>	970-223-5100 x9461
✓	Colegrove, Dan	IBM	<a href="mailto:colegrov@us.ibm.com">colegrov@us.ibm.com</a>	(408) 256-1978
	Coles, Alistair	HP	<a href="mailto:anc@hplb.hpl.hp.com">anc@hplb.hpl.hp.com</a>	+44 117 922 8750
✓	Coletta, Mike	Harris	<a href="mailto:mcoletta@harris.com">mcoletta@harris.com</a>	949-707-1143
	Domon, Wataru	NEC	<a href="mailto:domon@ccm.cl.nec.co.jp">domon@ccm.cl.nec.co.jp</a>	+81 44 856 2082
	Farhoomand, Firooz	Panasonic	<a href="mailto:farhoomandf@panasonic.com">farhoomandf@panasonic.com</a>	408-653-4059
	Finch, Stephen	Silicon Systems, Inc.	<a href="mailto:steve.finch@tus.ssi1.com">steve.finch@tus.ssi1.com</a>	714-573-6808
✓	Fuller, John	Microsoft	<a href="mailto:jfuller@microsoft.com">jfuller@microsoft.com</a>	425-703-3863
✓	Hamilton, Thomas	AMD	<a href="mailto:thomas.hamilton@amd.com">thomas.hamilton@amd.com</a>	(512) 602-9658
✓	Hannah, Eric	Intel	<a href="mailto:ehannah@mipos2.sc.intel.com">ehannah@mipos2.sc.intel.com</a>	408-765-4441
✓	Hauck, Jerry	Macro Designs	<a href="mailto:jerry@macrodesigns.com">jerry@macrodesigns.com</a>	510-668-1006
✓	Heilmann, Keith	IBM	<a href="mailto:kheilman@us.ibm.com">kheilman@us.ibm.com</a>	914-892-2413
✓	Hiraoka, Daisuke	Sony	<a href="mailto:dai@sm.sony.co.jp">dai@sm.sony.co.jp</a>	+81 3 5448 4603
✓	Killeen, Sean	SSL	<a href="mailto:sean.killeen@ssl.ie">sean.killeen@ssl.ie</a>	+353 1 02 5700
	Lin, Jinghui	Rocketchips	<a href="mailto:jl@rocketchips.com">jl@rocketchips.com</a>	512-685-4429
	Martinez, Steven D.	Honeywell	<a href="mailto:steven.martinez@das.honeywell.com">steven.martinez@das.honeywell.com</a>	505-828-5433
	Nave, James	Texas Instruments	<a href="mailto:jim-nave@ti.com">jim-nave@ti.com</a>	972-480-3468
	Northey, Bill	Berg Electrnics	<a href="mailto:northewa@bergelect.com">northewa@bergelect.com</a>	717-938-2119
✓	Oktay, Ozay	Canon	<a href="mailto:ozay_aktay@cissc.canon.com">ozay_aktay@cissc.canon.com</a>	(714) 856-7180
	Nyu, Takayuki	NEC	<a href="mailto:new@optsys.cl.nec.co.jp">new@optsys.cl.nec.co.jp</a>	+81 44 856 2082
	Pasternak, Jeff	Intel	<a href="mailto:jeff.pasternak@intel.com">jeff.pasternak@intel.com</a>	503-677-6753
	Polishuk, Paul	Infor. Gatekeepers, Inc.	<a href="mailto:ngaboska@aol.com">ngaboska@aol.com</a> (?)	617-738-8088
✓	Prouty, Bill	HP	<a href="mailto:bprouty@hp.com">bprouty@hp.com</a>	916-785-4631
✓	Richman, Mark	Lucent	<a href="mailto:richman@lucent.com">richman@lucent.com</a>	(610) 712-7456
	Rooke, Matt	Fujitsu (FCPA)	<a href="mailto:mrooke@fcpa.fujitsu.com">mrooke@fcpa.fujitsu.com</a>	408-894-3826
✓	Saito, Kyoza	Alps	<a href="mailto:kyozo_s@gw3.alps.co.jp">kyozo_s@gw3.alps.co.jp</a>	+81 229 23 5111
✓	Saito, Tomoki	NEC	<a href="mailto:saito@ccm.cl.nec.co.jp">saito@ccm.cl.nec.co.jp</a>	+81 44 856 2082
	Sato, Tetsuya	Sony	<a href="mailto:tetsuya@sm.sony.co.jp">tetsuya@sm.sony.co.jp</a>	+81 3 5448 4603
✓	Saunders, Brad	Rockwell	<a href="mailto:bradley.saunders@rss.rockwell.com">bradley.saunders@rss.rockwell.com</a>	(714) 221-6513
✓	Sessions, D.C.	VLSI	<a href="mailto:dc.sessions@vlsi.com">dc.sessions@vlsi.com</a>	602-752-6545
✓	Shariff, Masood	Lucent	<a href="mailto:mshariff@lucent.com">mshariff@lucent.com</a>	(732) 957-5479
	Skidmore, Jim	Texas Instruments	<a href="mailto:j_skidmore@ti.com">j_skidmore@ti.com</a>	972-480-2094
✓	Smolka, John	TI	<a href="mailto:rows@vnet.ibm.com">rows@vnet.ibm.com</a>	(507) 253-6290
✓	Teener, Michael Johas	Zayante	<a href="mailto:mike@zayante.com">mike@zayante.com</a>	(831) 461-4901
✓	Thompson, David	Lucent	<a href="mailto:davethompson@lucent.com">davethompson@lucent.com</a>	(610) 712-2730
	Wall, Duncan	3Com	<a href="mailto:duncan_wall@3com.com">duncan_wall@3com.com</a>	+44 1442 438197
✓	Whitby-Strevens, Colin	ST Microelectronics	<a href="mailto:colin.whitby-strevens@st.com">colin.whitby-strevens@st.com</a>	+44 1454 611500
	Wiener, Paul	Raychem	<a href="mailto:pwiener@raychem.com">pwiener@raychem.com</a>	650-361-2341
✓	Wooten, David	Compaq	<a href="mailto:david.wooten@compaq.com">david.wooten@compaq.com</a>	(281) 518-7231
	Yamazaki, Shuntaro	NEC	<a href="mailto:yamazaki@ccm.cl.nec.co.jp">yamazaki@ccm.cl.nec.co.jp</a>	+81 44 856 2082
✓	Young, Len	Corning	<a href="mailto:younglg@corning.com">younglg@corning.com</a>	607-974-8676
	Yu, Patrick	NEC	<a href="mailto:patrick_yu@el.nec.com">patrick_yu@el.nec.com</a>	408-588-5436
	Zalatimo, David	Filanet	<a href="mailto:davidz@filanet.com">davidz@filanet.com</a>	650-968-9711