IEEE P1394a meeting
Minutes

June 24 & 25, 1997
Bothell, Washington

Peter Johansson, chair, Colin Whitby-Strevens acting
Preshant Kanhere, secretary, Richard Churchill acting

1. Introductions and procedures
2. Review of minutes
3. Review of scope/closing actions table
4. Old Action Items
5. Old business
6. New business
7. Meeting Schedule
   7.1 Working Group Meetings
   7.2 Editorial sessions
8. Review of Action items
9. Wrap-up review of SCAT
10. Adjournment

1. Introductions and procedures
   Introductions around, with presentation of the agenda, which was accepted with out objection.

2. Review of Minutes
   Minutes were accepted without objection or revision.

3. Review of Scope/Closing Actions Table
   Handout of actions items (SCAT -- Scope and Closing Actions Table) was distributed. Colin made some comments on the nature of the table, and proposed (without objection) to defer consideration of the table and associated matters until tomorrow. Some attempt will be made to track SCAT items addressed during the assorted presentations.

4. Old Action Items

   Sony 4-pin Connector

   List of Sony DV related products was presented. showing 9 items on the market.
   A sheet showing radiated emissions for one such product was presented.
   A graph of emissions for the PC-7 was shown.
   Equipment layout used for test was shown (with Japanese annotations.
   [David Wooten asked which cables were 1394 cables. One cable between two cameras (operating at S100) was indicated. Colin asked for the information to be provided in soft format for the reflector. The material presented was confidential, and thus was immediately removed,

with comments by Colin regarding the fact that material presented is considered publicly disclosed. Colin asked if this information is being made public, with notice that presentation placed the material in the public domain. Portions at least were promised for release via the groups reflector, and the presentation continued.]
Sheet showing test of a DVD presented.
Pictures of the tested equipment and configuration were shown.
Diagram of connections shown, with tags/legend of components present.
[Questions regarding the testing 4-pin to 6-pin connections were made, with the answer indicating that they had not. Further questions about presence or absence of a ferrite bead on the cable ("X" means ferrite) in the presented diagrams. The observation that the test used only 2m cable, with the claim by Sony that the FCC test rules require a 2m cable. The observation that anything longer than 2m is allowed to "go off to infinity on the ground plane." This was observed to be an area may need review, with suggestion that we form a study group to make a proposal to the FCC. Question raised again whether the connection was 4-pin-to-4-pin only, in the presented diagram. Colin made observations regarding regulatory objectives, but others raised questions regarding whether we have enough information regarding FCC rules and requirements, with questions raised on the validity of the data. What we are doing is to define a standard that configuration of a system using the connector does not violate regulatory rules. Colin recommended that concerned parties form a working group. Colin asked whether this constitutes a sufficient body of material to satisfy our needs. Questions were also raised regarding whether we should worry about this for sundry reasons. The question was raised (in the context of needing multiple ports for 1394) whether the test results .... <secretary could not hear the speaker> What is the validity of these results for more than one port per device? Observation made that this material did not address the 4-to-6-pin problem, with suggestion that we defer the question.
Taka stated that this is an accomplished fact and that we must support it for various reasons. Question about the "Sony computer video I/F" answered as "we needed something to connect these components ...."
<*** Secretary's editorial comment: We are in no way constrained in our actions by the actions of any business or group of businesses, just as they in turn are not constrained by us, beyond the extent they choose to be. Thus to say that the 4-pin connector is an accomplished fact is irrelevant with respect to this standard. As a group, we should be morally bound to exercise due diligence, and to fully accept or responsibility to conduct peer review of the material covered or mandated by the draft standard. We should therefore feel bound to reject any and all material which we find there to have been inadequate peer review of, and which is otherwise not necessary to the task at hand. Doing so in the case of the 4-pin connector would in no respect a slight to the company or companies using or planning to use it. As Taka has pointed out, the 4-pin connector is an accomplished fact. It is in use despite not being covered under any IEEE 1394 standard. (It is, however, under the ISO 61883 draft, which is nearing ratification.) A company wishing to use it is free to do so. The secretary therefore finds Taka's argue-

ment that we must include the 4-pin connector in the standard simply because it is already in use wholly specious, and that the question of whether the 4-pin connector should and must be incorporated in the draft is answered by our deciding whether we have had adequate time and material to conduct an acceptable peer review, and whether it is necessary for us to be able to meet the objectives of this draft standard. ***> Wooten questioned the presence of a line in one graph, which turned out to be an internal design guideline, lower than FCC. Sony comment that they would not release a product that did not pass this requirement, but that a shipping product operates under different rules. Can we identify 1394's contribution to this test of complex configuration, with so much added noise from equipment other than the 1394 connection? Not yet, though Mike Brown offered to work on it. How did IEC pass this? IEC does not require the same peer review. Intel would like one month to do additional work ... that takes us to the next meeting, in early august.]
Chair does anyone think insufficient material was presented? Churchill indicated yes. Vote to retain consideration for another month -- question reversed to "To remove any futher consideration of the connector from the P1394a standard." Taka -- other companies are working on some products with this ... Colin -- Any other data available now? No votes for removal, 22 against removing.

[Chair suggested starting at 8:30 tomorrow, with the shuttle starting at 7:00.]

5.3 Suspend/Resume and Power Management Group Report (Claude Cruz)

Power distribution and management documents progressed, with some work on suspend/resume. Attempt to narrow scope to meet P1394a schedule requirements. Will attempt to use TpBias, and be able to detect normal attach/detach events, etc. ... Mike Brown spoke about a way to turn off everything, and shorting TpBias to ground, so that some 10's of uA during suspend from TpB, which may provide a means of detecting disconnect. Mike Sorna at IBM is participating in this effort. The method uses the DC coupled portion of the connection, not the AC coupled portion. There is also a "chirp" for various reasons. Trying to avoid "logic based" schemes. Should have some further information next month. The 1394.2 effort looked at the problem, and arrived at use of chirps <2 us chirp out of every 50 us, with 20% of the power usage of leaving link active.
[Colin points out that we need objectives, and methods, including how to put a node to sleep, coordination between nodes across a connecion for going to sleep, etc. Also need mechanism for putting a node to sleep remotely. Do we need a mechanism for putting a particular port to sleep? There was further discussion of mechanism, with several options discussed. Power Rangers will try provide suitable mechanisms, but needs to come to some conclusion soon. We need to avoid having a large base of incompatible silicon that would delay or preclude a more robust solution. PM group needs to provide the basis to assure con-

sistent implementation, even if the details are worked out and "set in stone" via a standard later. How is this base mechanism to accomplish turning off a single port-to-port connection? Naked PHYs are problematic, and may require special treatment. The list generated (see below) does not mean much without a context, and we may end up going a different direction <Wooten>.
(Claude Cruz drew a diagram ...)
Need to flesh out a requirements and capabilities list that is comprehensive for the August meeting in Honolulu. (Wooten) What you put up there is 31A, not 31 .... PM workgroup reflector (list@p1394pm.org) should be used to continue the discussion. How many would be willing to participate? <Capture a list of those willing to participate. Any interested party should contact Steve Bard for sign-up. This is a separate list from the PM list, and should be only those who will be actively participating. This will entail one or more conference calls in the near future.> Further comments to the task group. Sign-up passed ... result needed the week before the TA meeting (specify that this should be Thursday ...)]

Method to remotely put a port to sleep
Method to wait for suspension
Generate a mechanism for PHY packet to ?????

4.3 Cable Test Procedures (Eric Hannah)

(hard copy available)
1394 Connector and Cable Testing

Presentation Goals
-- Discuss critical parameters for a cable interconnect in P1394a sys
-- Present a series of reasonable tests that system integrators can to validate ...

Critical areas for cable interconnect
-- Differential Impedance profile
-- skew
-- Differential eye diagram
-- ???
-- ???
-- ???

Differential Impedance profile
-- Differential signals are the primary signaling mechanism
-- Variations in differential impedance lead to signal loss and reflections
-- TDR studies of differential impedance can detect bad cable terminations
*** diagram ***

Skew

????

Differential eye diagram
-- The receiver end eye diagram is the fundamental measure of signal quality
*** diagram ***

Common Mode impedance profile
-- some signal (speed signaling) are common mode ...
-- ?????

Common Mode Crosstalk
-- speed signaling at self-ID time can be corrupted by near-end crosstalk
*** diagram *** *** diagram ***

[There was much discussion regarding timing, crosstalk, etc., between speed signals.]

Cable EMI shield effectiveness
-- recommended procedure
   -- test cables in 3 m or equivalent screen romm, with full differential 1394 voltage applied as single-ended signal to cover maximum skews, terminate in 110 Ohms.
   *** diagram ***
   *** table ***

<Make this SCAT item 53, with action item for Peter for wording ... done by concensus on the basis of its being accidentally omitted.>

<Action (by Taka Fujimori) to present further material on crosstalk suspended ...>

<<<< Recess for lunch, 12:04 PM, to resume at 1:00 PM, actual at 1:10 PM >>>>

. Fairness Optimizations

A general discussion was held regarding what might be gained by use of either optimization -- Budget or Privilege -- as well as some of the liabilities. Churchill made the point that if you are looking for an absolute upper bound on throughput, little is to be gained from either methods, since either the bandwidth lost to reset gaps is already small compared to data transmission time plus subaction gaps, or traffic was not that intense, and performance so critical anyway. If the problem is viewed somewhat differently (raw throughput is not so critical as responsiveness in dealing with many small packets, such as during system configuration) the modeling is much more difficult, but the size of the reset gap becomes more prominent as compared to subaction gaps and reduced transmission times. Thus the bandwidth in this less "optimal" situation is more significantly affected. After extended

discussion the motion was presented to go with the fairness budget, and a vote was eventually taken, with the fairness budget winning convincingly. <*** The actual voting numbers were lost. The secretary of the group was participating in the discussion, and the alternate did not record the vote. However, the vote was on the order of three to one in favor of the budgatary approach. ***>

. Copy Protection (Brendan Traw, Intel)

Update on Content Protection for the IEEE 1394 Serial Bus
(Brendan Traw)

IEEE 1394 Content Protection
-- Overall Goal is to ensure that license conditions for copyrighted content are enforceable through technical and legal means
-- Key enabling technology for digital transport of copyrighted "Hollywood" content
-- Work is taking place in the Copy Protection Technical Working Group (CPTWG). The CPTWG is the forum for PC, CE and content industries to adopt content protection for emerging technologies including IEEE 1394 and DVD
No digital output of Hollywood content will be allowed without content protection!
*** illustration ***

Content protection chain
-- copyrighted content must be protected during all phases of distribution, transmission and playback.
-- Protection is only as strong as the weakest link.
*** illustration ***

CPTWG Digital Transmission Discussion Group (DTDG) Goals
-- protect copyrighted content traversing ...
  -- can be implemented with reasonable processing overhead in software on a PC
  -- Can be implemented with reasonable cost in consumer electronics devices
-- Content protection should be transparent to users
who comly with content’s copyright
-- Licensable component but no expensive IP
-- No import/export restrictions

Copy Control Information
-- Copy control information (CCI) specifies the conditions under which copyrighted content can be copied
  -- Should include support for at least CGMS, APS and Digital Source bits
  -- Layer with greatest impact on IEEE 1394 standards
-- Three forms of CCI
  -- Exposed CCI

      -- carried in isochronous packet or CIP headers
      -- Integrity not guaranteed
      -- Typically a subset of full CCI
   -- Embedded CCI
      -- Carried with content (may be a watermark)
      -- Integrity guaranteed through encryption, hashing, or other means
   -- Other options are also available for exchanging CCI
      -- Carried "out of band"
      -- Integrity guaranteed

Proposals submitted
-- Hitachi
-- Intel
-- MEI
-- NDS
-- PictureTel
-- Sony
-- TI
-- Toshiba
Disclaimer:
   -- These proposals are moving targets
   -- Some are significantly more detailed than others, so in some cases information is inferred or based on verbal comments

DTDG Proposal Summary Chart
*** table *** showing features of some proposals

What does the DTDG need?
-- Input on the suitability of the proposed techniques from a 1394a perspective
-- Placeholder in 1394a spec for content protection

MEI Proposal
*** diagram ***
-- Reallocate 2 bits of sync field for CCI
-- Encoding: 00 = copy free, 01 = bit-stream copied, 10 = copy once, 11 = copy prohibited.

Sony Proposal
*** diagram ***
-- 2 reserved bits from CIP header used for CCI
-- Encoding: 00 = copy free, 01 = bit-stream, 10 = copy once, 11 = copy prohibited
-- TCode = ’C’ may be used.

Current OHCI Approach
*** diagram ***
-- TCode = ’C’ to indicate copyrighted content
-- CCI information stored in TAG field

Intel's Position
-- Exposed CCI is not required
   -- Information is redundant
   -- Easy to circumvent
   -- Other layers contribute to robustness
-- We recommend no changes to the IEEE 1394-1995 specificaiton for content protection

[There was considerable discussion of the scope and limits within which we (a 1394 standards working group) should be addressing the underlying legal issues, etc. Brendan made the point that the subject is better addressed in the CPTWG, and recommend that those interested should go there. However, the CPTWG and DTDG need feedback regarding the technical issues associated with the proposals, and wants a technical recommendation for one of these proposals. Colin observed that we are being asked to modify more than lies in the understood scope for P1394a, and thus requires a two-thirds vote to consider. Comments were made to the effect that this is within scope, and discussion continued. Recommendation of the CPTWG was NOT to use the OHCI scheme, which was a distant third in the CPTWG's popularity votes, but to use either the MEI or SONY approach. Long discussion regarding the legal issues, and the fact that we are not attorneys. The fact that the CPTWG just wants a technical recommendation]

Moved by John Fuller, seconded by Steve Bard that we (1) reject OHCI approach, and (2) if an exposed CCI is required, we state a preference for MEI's approach.

More agonizing discussion to no particular benefit .... Fuller commented that an use of the sync field is covered under 1394, since this is specified for use by the application anyway.

Friendly amendment to split the questions accepted.

Question 1: For 21, against 1, abstaining 2. (Comment of Wooten who voted against the motion is that this is premature to vote on, as we have not had time to consider adequately.)

Question 2:
Move to table the motion by Wooten, seconded by Churchill
failed (vote lost but decisive).
Eight prefer to remain silent on the issue, 11 wish to express a recommendation in a straw poll.
motion to table til tomorrow (offered by Fuller, seconded by Wooten) passes (vote lost due to rush), and thus tabled until Wednesday, June 25.

4.1 Dual-phase retry (SCAT 33) (Johansson)
Peter was not present to deal with this item.

[Discussion inaudible to the secretary ensued.]
[Do we need to modify the standard? ... Questions of two interval versus three or four interval busy ... Opinion expressed that there is no way to fix the problems short of changing silicon. <Fuller>]

?.? Token Style Arbitration (SCAT 51)
[We apparently agreed that this is informative.]
Declared informative without objection, pending Peter's further elucidation.

4.4 CPTWG letter (Johansson)
Left as an open issue

4.5 Direct connection drawing (Wooten)
[Peter apparently just wanted picture of logic and chasis grounds wired together. Arose from a discussion in Eindhoven wherein a presenter suggested this, which seems to be a bad idea according to PC manufacturers.
Citation of minutes from Eindhoven meeting elucidated the action item, resulting from comments by Mike Brown and Mike Teener over a suggestion by Max Bassler.
We have questions regarding coupling of sources of noise to logic ground, and we should not add more instability to the grounding scheme. Diagram could add little or nothing but the point that this introduces instabilities into the system, and thus is quite problematic.
Observation raised that the diagram was for didactic purposes, to show that this is a bad idea. This applies to the Annex A discussion, which should indicate that PHY-LINK isolation is not required.]

4.6 Annex A (Eric Hannah)
We will indicate that the typical situation is that all nodes will operate on the same ground/green-wire, but that if this is not the case the vendor should exercise good engineering practice.

4.7 PHY/LINK Interface reset via LPS (Baker)
4.8 Power distribution safety issues (Wooten)
4.9 Configuration ROM "generation" bit (Johansson)
4.11
4.12

<<<< Recess for night, to reconvene around 8:45 Wednesday. >>>>

<<<< Reconvened at roughly 8:45 >>>>

. Copy Protection (Brendan Traw) continued

Motion to amend Question 2 to state that we reject all methods proposed. Moved by Paul ????, Seconded by Bill ?????

Vote on amendment Yea 9, Nay 10, rejecting the amendment.

[There was cosiderable discussion regarding the appropriateness of the mechanisms proposed, along with questions regarding the hardware changes that would be required by the various proposals to OHCI and LINKs. David Wooten observed

Question 2 is called, with vote 3 yeas, 19 nays, and 8 abstentions.

Straw poll proposed asking whether we should say nothing -- remain silent -- regarding options proposed.

Moved that we comment that this does not fall within the scope of the P1394a task, etc., died for lack of a second.

Moved that we reply that we regard the options proposed as being unsuitable for copy protection, as they do not remove the PC as a threat. (Designate as Question 3) Moved by David Wooten, seconded by Richard Churchill. Friendly amendment proposed (by Traw) to add "may also lead to compatibility problems." Rejected by Wooten, and dies for lack of a second. Friendly amendment to insert "and other user reprogrammable devices," accepted.

Question called, with 17 yeas, 4 nays, and 13 abstentions, carrying the motion.

4.7 PHY-LINK reset by LPS (SCAT 32) (Neil Morrow?)

LPS Reset
-- PHY-LINK interface reset mechanism
   -- power intialization
   -- recover from "hung" PHY-LINK i/f
-- Can be used for:
   -- host power reset
   -- software initiated reset
-- LPS well behaved
   -- does not generate 1394 bus reset

Link Power Status (LPS) Need
-- Turns off PHY-LINK I/F when LINK is powered down
   -- OFF
      -- protect I/Os from excessive current
      -- reduces PHY power consumption
      -- PHY ignores unknown states of PHY-LINK inputs
   -- ON
      -- resets PHY-LINK interface when LINK is powered up

1394 LPS Proposal
-- LPS communicates LINK power status
   -- if LINK power is OFF, LPS = 0

   -- if LINK power is ON, LPS = 1, or switching waveform
-- LPS resets the PHY-LINK interface by signaling a power cycle seq.
-- "Switching Waveform"
   -- works in isolated environments

1394 LPS Proposal - PHY
-- PHY senses LPS active when: -LPS - 1 (immediately)
-- PHY senses LPS inactive when: -LPS = 0 for more than 2.75 us
-- When PHY senses LPS inactive
   -- All PHY-LINK outputs are driven to 0, including clk, control Lreq and data signals.

[Does this guarantee the reset of the VLU? Add, "and go to idle ctrl state." as bullet following "All PHY-LINK ..." bullet. Question regarding whether the intention is to reset the state machines, though the inactive case also turns off the clock. Should this turn off the clock? Does this overload this signal? Further discussion of how this will work in reseting state machines and minimums, etc. ...]

1394 LPS proposal - LINK power sensing
-- LINK drives LPS active/switching when power is valid
-- LPS inactive when power is not valid

1394 LPS proposal - LINK reset
-- Link drives LPS inactive for longer than 2.75 us to reset the PHY-LINK I/F
   -- PHY drives CLK to 0 when LPS is inactive
--

1394 LPS proposal - timing
*** diagram ***
-- LPS active
   -- Tpwh = 90 ns min
   -- Tpwl - 2.25 us max
-- LPS inactive
   -- Tl = 2.75 us min

[Questions raised about timing, and min/max times. Link must hold low for some additional time for PHY slop ...

PHY-LINK LPS Timing reset (at PHY)
*** diagrams ***

PHY-LINK LPS Timing reset (at LINK)
*** diagrams ***

Isolated PHY-LINK Interface
*** schematic ***

Direct PHY-LINK Interface
*** schematic ***

[Discussion of timings continued, with questions of overloading a signal, continued ... Presenter stated that he actually thinks this is not necessary, but that Richard ??? of TI, for whom he is presenting this material does.] << Much dialog was lost due to the inability of the secretary to hear the discussion well enough to record it. >> [Further comments regarding asynch requirements, etc. Whether R's concern in one area was the overloading of a capacitor, with clamping of input latches ... LINK side of control lines could come up 1-1, so how do you deal with this? ... Should all this be informative? No answer ... Further discussion ...]

PHY-LINK initialization:
When and only when there is an isolation barrier, the LINK shall wait until it sees SCLK, and then initialize the interface by driving C/D/LREQ low for 2 SCLK cycles.

SCAT 32 resolutions
1. LPS -> 0 => clk, ctl, dta -> 0 (Concensus)
2. Initializations of PHY-LINK -- reflector
   Action item for Neil Morrow to elucidate the problem state via the reflector, etc., plus proposal for timeouts, etc., by Joe ??? and Jerry Hauck

. LREQ Issues ()

LREQ Issues
-- What does "in isoch phase" mean for iso LREQ?
   -- Not only receiving and transmitting isoch packets
   -- Also includes receiving/sending CycleStart packet
      -- LINK must send LREQ before entire CycleStart received
-- Proposal
   -- State "in iso phase or during CycleStart reception ..." for
   -- Add language that cycle start needs to be predicted
      -- If not cycle start (ex. bad CRC), when granted, release bus.

[This may be bad stuff since the Hamming distance between some of these ??? is only one ...]
[IF sent cycle sync, next packet with tCode = 8, assume cycle start and enter iso phase. breadcast (stuff in 1st 2 quadlets ...)]

Moved that "CycleStart always at S100." by Jerry Hauck, seconded by J. Bennett.

[It was questioned whether it is appropriate to specify that we send CycleStart at S100 if we do not have a problem that we are definitively fixing by so specifying. Further discussion regarding timings of LREQ vs. CycleStart reception was held, concluding that there is a real

problem, but the one addressed by the question is not that problem.]

Motion withdrawn ...
Action item: Jerry Hauck will investigate maximum delay through PHY to ensure iso request is seen.

[Further discussion regarding when to send acks, etc., when the CRC is verified relative, etc. ... Two issues -- We know we must arbitrate before we get packet ... if a packet comes in before we know it is a cycle start ... One other thing said was that we could hold onto the bus for however long we want ... No, but we can hold the bus for quite a while ... There is a constraint. ...]

LREG Issues
-- Multi-speed Concatenation
   -- When downshifting to S100, must issue another LREQ
   -- In links with multiple transmit data streams, may not be able to transmit LREQ while Xmitting current packet
      -- Packets concatenated into a transmit FIFO
   -- Can LINK use delay thorugh PHY to its advantage?
      -- PHY is asserting data end for .24 us (12 SCLKs)
      -- This occurs while PHY-LINK ingerface is IDLE
-- Propose change to LREQ Table for Isochronous
   -- "... while CTL[0:1] is in receive or transmit and up to 10 SCLKs after last isochronous transmit."

[This may bread old PHYs, though it should work for "a" PHYs. Why? This was never spec'ed ... Why is this an "a" PHY issue? Proposed that this be in the next draft, subject to confirmation ... with no objections, proposal accepted, with material to be presented on the group's reflector.]

LREQ Issues
-- Background
   -- LINK has Xmitted priority requrest for enhanced arb
   -- Did so while waiting for ack for previous transmit
   -- LINK determines it then must send Cycle Sync
      -- Does it wait for priority requrest to be servced or ...?
      -- Does it send Cycle Sync IMMEDIATELY?
-- Proposal
   -- LINK sends CycleSync LREQ
   -- PHY does not cancel Priority request

<<<< Recess for lunch at 12:06, to reconvene at 1:06 >>>>

Proposed by Chair that we continue with the old business until 2:00, then review briefly the several "beyond scope" items.

. Arbitration enhancements (Ganesh Murthy)

-- Is enab_accel adequate ...

-- enab_accel is adequate to handle all arbitration enhancements
   -- Both ack accelerated and fly-by accelerations occur only during the asynchronous phase. No need for separate bits.
   -- CycleStarts disable ack accelerated and fly-by accelerations
-- enab_accel should power up off or disabled
   -- Leaving it turned on may cause cyclestarvation, since legacy link can't disable it
   -- Not optimal but acceptable for P1394a LINKs.
   -- Power up defaults of all PHY registers needed in draft 0.09
-- enab_accel should be turned on if and only if -
   -- P1394a PHY is a root
   -- Attached link can generate CycleSync LREQ.
   -- P1394a Link or P1394a aware higher layerswill turn enab_accel bit ON.

Action Item:
Colin Whitby-Strevens
Check that there are 1/80 PHYs that prefer child to Link, rephrased as 1/80 PHY which root grants a child before subaction gap, whilst it is

There is a better way?
-- Let the P1394a PHY learn when to turn on enab_accel
-- enab_accel bit powers up dounble
-- P1394a PHY detects it is root and automaically turns on enab_accel
-- P1394a PHY gets a Cycle Sync LREQ, "know" it is attached to a P1394a LINK and so it turns on enab_accel (time (120 us periods of no enhancements)
-- Lets P1394a PHY work with a 1394-1995 link if it is root
-- Higher level layers are not allowed to write to enab_accel!
-- Bus reset resets enab_accel

[Colin asked, "Do we want to do this?" Much discussion follows to no particular point (that can be heard) ... Presenter requested to clarify this material in light of comments.]

Material accepted provisionally and subject to revisions as per the action item.

<<<< Old business partition reached at 2:08 -- New busy only from here >>>>

6.3 Power Class in the Caboose packet (pp Johansson)

Request to extend scope -- Accept as within scope of P1394a draft mechanisms agreed upon in the Power Specification which may affect P1394a, provided that such agreed upon mechanisms are submitted to the P1394a Working Group for consideration prior to the closure of the Draft Specification (August '97).

So moved by Richard Churchill, second by David Wooten, with a friendly amendment accepting the caboose packet proposal as part of this. Votes were (of 30 voting members present in the room) 27 yeas, carrying the motion irrespective of no and abstention votes.

6.5 PHY Version register

P1394a is enabling a brave new world of mix and match links and phys. And while software can identify the link based on its host bus interface, we have no way to tell one phy from another. Yes, I know they are all supposed to be the same, but even if they are there will be extensions to this standard in the future that may require phy changes.

I propose that we define one value of Page_Select to be the following read-only registers:

1000b	Compliance level (00 = P1394a)
1001b	reserved
1010b	manufacturer’s OUI (MSB)

Wooten moved, Jerry Hauck seconded, with 25 yeas and one nay from the 26 voting members present.

6.6

6.8 Lock Transactions (Young)

Opinion expressed that further explanation of how lock transactions work, and operation in relationship with various operations/registers especially as regards the behavior of reserved bits and fields in the behavior of compare operations.

[Reference to the IEEE 1212 standard and David James, where this may be dealt with, and from whom further explanation may be obtained.]

Tabled pending clarification from David James, with note that this is not really a scope issue, being explanatory in nature, rather than technical.

6.12 Port and Line speed reporting

Tabled pending concern, by unanimous concensus.

5.* PHY Pinging

Unfinished.

Location of the ping timer is in the LINK.
(SCAT 48?)
Assumption that timer is located in the LINK,
and agreed by unanimous concensus.

5.* Asynchronous Streams

Moved we agree by Richard Churchill, and unanaimously agreed
by concensus

. Power distribution (SCAT 35)

Replacing 4.4.??? in the -1995 standard ....
Closure is planned in the Power Management/Distribution meeting
on July 16. Issues remain open. Contents of 6.1? are correct
to the extent covered, but are incomplete.

. LREQ Stop bits

Previously agreed, but awaiting addition of text by the editor.
[Much discussion regarding the nature of and reasons for the
motion passed in a previous meeting occured.]

5.7 Data length in Request/Response (SCAT 50)

Deals with 9.1?, and agreed by concensus.

5.8 Data length and max_rec (SCAT 29)

Value should be the same for both read and write operations, and
so state, rather than reference to -1995 for definition (p.224).
Moved that we agree and agreed unanimously.

5.9 SCLK availability -- previously covered

5.10 Alternate cable/connector -- previously covered

5.11 Cable PHY enhancements - "Caboose" packet. (SCAT ??)

Moved we retain caboose packet as presently defined, by R. Churchill,
seconded by Steve Bard.
Votes were 7 yeas, 1 no, passing.

6.4 Interoperability of 1394 LINKs and PHYs

Sure ... Continue on

6.7 SPEED_SIGNAL_LENGTH ...

Make these consistent ...

6.* 1394a Isoch connection management (Jerry Hauck)

1394a isoch connection management
-- Two Sections
   -- register locations and specifications
   -- setup/teardown of isoch connections
-- Questions
   -- Is each section normative or informative?
   -- What is required of a 1394a isoch node to comply with this?
   -- What is required of -1995 nodes in this area
   -- What kind of support will be provided in bus driver? What needs to be implemented in hardware?
   -- Do we need more critical peer review? Are all the state machine transitions well understood? Have all setup/tear-down scenarios been considered.

Moved that we restrict clause 8 to allocation of addresses for PCRs, plus an informative cross reference to IEC 1883, to be confirmed at the next meeting, by David Wooten, and seconded by Jerry Hauck. Yeas: 12, 2 abstentions, and no nays, passing the motion.

6.10 Link to check cycle start (Bennett?)

Link should look at whole of 1st quadlet, not just the tCode of ’8’. [What we may need is clarification. Do we need to check all the fields, or is it sufficient that they be reserved, usw. ... There may be reasons to use tCode ’8’ for other purposes someday. Suggested that wording be prepared for review by editor, and by the whole group at the next meeting.]

6.11 Cycle start timing (Shergill)

Concern is that enough delay exists to risk causing ????

Note in 1394a that the jitter on delivery of cycle start is up to 340 ns, peak to peak, and link imlementations should be designed appropriately. (20 ns per hop + 25 ns of clock jitter, 16 hops)

6.13 PHY-LINK handover (Sean Killeen)

This is concerned with who drives the bus when, and addresses when the PHY hands control to the LINK. Material was submitted on 6/23, yielding inadequate time for review. Consideration of any changes this might require are deferred until the next meeting, when review time will have been greater.

6.14 Bandwidth available at speeds greater than S400

The formula for bandwidth allocation needs to be changed to deal

with speed greater than S400 adequately. This is an editorial item only.

7. Meeting Schedule

7.1 Working Group Meetings
Augest 4-5, Honolulu, HI
Septemter 25-26, Boston, MA

7.2 Editorial sessions
July 28-29, San Jose, CA
September 24, Boston, MA

8. Review of Action items

9. Wrap-up review of SCAT

10. Adjournment

Action items:

1. Notification of DTDG regarding decisions of the P1394a WG
   Brendan Traw
2. SCAT 32
   Examine LPS reset proposal to elucidate the problem statement, accommodating comments and suggestions.
   Neil Morrow
3. SCAT 32
   Proposal for timeouts, etc., for LPS; PHY-link clean-up from hang as an alternative to LPS
   Joe Bennett and Jerry Hauck
4. Clarification of the enab_accel management in light of comments.
   Ganesh Murthy
5. Clarify dates of forth-coming meetings
   Richard Churchill
6. Update SCAT
   Colin Whitby-Strevens
7. Check on number of old PHY issues that might preclude accelerated arb
   Mike Eneboe
8. To check on required timings for issuing immediate and iso requests
   Jerry Hauck

Attendees:

Kazuyuki Abe	(415) 528-5904	kabe@sj-pceg.ccgw.nec.com
Steven Bard	(503) 264-2923	steve_bard@ccm.jf.intel.com
Joe Bennett	(916) 356-3722	Joseph_A_Bennett@ccm.fm.intel.com
Vilas Bhade	(408) 248-7882	vilas@wipro.com
Mike Brown	(602) 554-3713	mike_brown@ccm.ch.intel.com

Mike Bryan	(303) 682-8407	mike_bryan@notes.seagate.com
Jim Busse	(415) 528-3810	jimb@ccgate.sj.nec.com
Carissa Cheung	(916) 785-3119	carissac@hprnd.rose.com
Richard Churchill	(713) 514-6984	richardc@bangate.compaq.com
Chris Cone	(505) 822-7394	chris.cone@abq.sc.philips.com
Claude A. Cruz	(503) 264-7109	claude_cruz@ccm.jf.intel.com
Jim Dente	(303) 678-2184	jim_dente@maxtor.com
David Doman	(512) 891-2922	daved@oakhill.sps.mot.com
Mike Eneboe	(408) 974-0615	eneboe@apple.com
Mark Evans	(408) 894-4019	mevans@qntm.com
Firooz Farhoomand	(408) 945-5600	firoozf@ix.netcom.com
Tony Foster	(916) 785-1092	tfoster@hprnd.rose.hp.com
Taka Fujimori	81-3-5488-6353	fujimori@cv.sony.co.jp
John Fuller	(425) 703-3863	jfuller@microsoft.com
Nobuo Furuya	(408) 588-5414	nobuo_furuya@el.nec.com
James Gay	(512) 891-2218	jimg@oakhill.sps.mot.com
Sreekanth Godey	(408) 235-8500	sreekant@sandmicro.com
John Grant	(508) 653-8804	j.l.grant@ieee.org
Eric Hannah	(408) 765-4441	ehannah@mipos2.sc.intel.com
John Harimoto	(916) 785-8946	johnha@rosemail.rose.hp.com
Yasumasa Hasegawa	81-22-347-1128	yasumasa@ffm.fujifilm.co.jp
Jerry Hauck	(408) 765-5528	jerry_hauck@ccm.sc.intel.com
Joe Herbst	(408) 969-4607	jherbst@issiusa.com
Kenji Keino	81-462-30-6209	keino@lsis.crl.sony.co.jp
Mark Knecht	(408) 570-1097	mark_knecht@phoenix.com
Larry Kopp	(717) 592-6267	lskopp@amp.com
David LaFollette	(408) 765-2587	dlafolle@mipos2.sc.intel.com
Paul S. Levy	(602) 752-6382	paul.levy@tempe.vlsi.com
Paul Lind	(602) 554-3774	plind@sedona.intel.com
Takashi Matsui	81-76-274-2440	takashi_matsui@ibm.net
Palanisamy Mohanraj	(602) 554-4243	palanisamy_mohanraj@ccm.ch.intel.com
Cyrus Momeni	(510) 624-7061	cyrus@corp.cirrus.com
Neil Morrow	(903) 868-5481	nmorrow@ti.com
Ron Mosgrove	(503) 264-2229	ron_mosgrove@ccm.jf.intel.com
Ganesh Murthy	(503) 264-8737	ganesh_murthy@ccm.jf.intel.com
Karl Nakamura	(408) 433-4516	karln@lsil.com
Bill Northey	(717) 938-2119	northewa@bergelect.com
Farrell Ostler	(505) 822-7791	farrell.ostler@abq.sc.philips.com
Kugao Ouchi	(408) 588-5503	kugao_ohuchi@el.nec.com
Dan Paley	(408) 570-1256	dan_paley@phoenix.com
Bill Prouty	(916) 785-4631	billp@hprnd.rose.hp.com
Ron Roberts	(916) 677-5714	rkroberts@aol.com
Bradley Saunders	(714) 221-6513	bradley.saunders@nb.rockwell.com
David Scott	(503) 264-1615	davidx_j_scott@ccm.jf.intel.com
Robbie Shergill	(510) 721-7959	rss@berioz.nsc.com
Hisato Shima	(408) 955-5524	shima@ssl.sel.sony.com
Sam Takahashi	(408)342-8612	sam@masca.com
Dong Tieu	(408) 765-3944	dong_tieu@ccm.sc.intel.com
C. Brendan S. Traw	(503) 264-8987	brendan_traw@ccm.jf.intel.com
Colin Whitby-Strevens	44-1454-611500	colinws@bristol.st.com

David Wooten	(281) 518-7231	davidw@bangate.compaq.com
Roy Yasoshima	(415) 858-1000	yaso@masca.com
Takao Yasuda	81-3-5488-6353	yasudat@cv.sony.co.jp
Phil Young	44-1908-837247	youngp@euk.nec.co.uk