

## Agenda

### P1394b Editorial Working Group

#### (Draft Specification Page Turner)

October 28<sup>th</sup> & 29<sup>th</sup> 1999

*Beginning of Meeting SCAT COUNTDOWN: 3 - OPEN; 9 - WIP; 9 - AIP; 33 - Done; 11 - Closed; 65 total database entries.*

*End of Meeting: SCAT COUNTDOWN: 3 - OPEN; 9 - WIP; 9 - AIP; 33 - Done; 11 - Closed; 61 total database entries.*

Those in attendance: David Wooten (Compaq), Victoria Teng (NEC), Jerry Hauck (Zayante), Michael Teener (Zayante), Steve Bard (Intel), Colin Whitby-Strevens (Zayante), Eric Hannah (Intel)

Action Items from previous meeting:

#1) Steve Bard has an AR to add text to clause 11.51 standby state transitions. [*STATUS: 10/28/99: Much thanks to Colin Whitby-Strevens who took the initiative to craft the words for this and insert them into the draft.*]

#2) Mike Teener will write informative text for section 3.3 that will explain where in the draft each of the boxes in figure 3.1 is specified. [*STATUS: 10/28/99 - Mike has not yet completed this task. His plans are to discuss the task with Eric and have it completed and in the draft prior to the editorial page turner on December 8<sup>th</sup>.*]

#3) An Action Item to the Editor: Change all occurrences of lower case "l" in "Legacy" to an upper case "L" (like in clause 4.2.2); [*STATUS: 10/28/99 - Not yet Complete*]

#4) Colin to develop a proposal for definitions of the terms outlined in the table included below. [*STATUS: 10/28/99 - Completed and incorporated in the 0.87 draft.*]

#5) Steve Bard will deliver a list of all folks whom have attended any P1394b meeting to David. - [*STATUS: 10/28/99 - Done*]

### **Meeting Notes:**

There was discussion of what needs to be done to the draft before a true "editorial" page-turner. There is a good amount of technical work that must still be done in the draft. Clarification: the technical work is not NEW technical content, rather, correctly integrating agreed on technical content into the draft. This is to include correcting numerous errors in the current 'C-Code' as well as correcting text in the draft to align it with those portions of 'C-code' that are correct. Also to be updated to incorporate correct technical content is the PIL/FOP section of the PHY/Link chapter (clause 13.11 through clause and In ADDition, Is it needed to have one or more face-to-face "technical" editing sessions - sessions in which the draft is examined for technical correctness for all of the agreed in principle and agreed in concept SCAT items as well as a critical review of all of the technical content that the SCAT issues table shows as "DONE". This includes review of sections of text on B.O.S.S. (for example) and making sure it is technically correct and in alignment with the 'C-Code". Jerry pointed out that the 'C-Code' in draft 0.091 has numerable content errors, i.e. the 'C-Code' does not contain the correct algorithms to process the technical content of 139b that the working group has adopted and agreed upon.

Five items for discussion from Jerry:

- 1) Cancelled Requests
- 2) Cycle Start starvation
- 3) Arb state timeout
- 4) request mapping for point to point protocol packet unexpected idle ( the issue really is unavailable control tokens and/or overloading of existing tokens)
- 5) What about the long resume time required for Sony Camcorder (the 1 second TpBias assertion issue)

2) Cycle Start starvation: Jerry believed that there was a solution developed over dinner Monday evening in Rennes, France, however, during an open forum discussion in the plenary on Tuesday, David Wooten brought up the issue without the previously agreed upon solution. Jerry wanted to close on this issue and make certain there truly is a solution that should be incorporated into the draft.

Jerry reviewed the solution that was discussed at the dinner session in Rennes.

Colin has crafted the 'C-Code' that implements the solution Jerry outlined. That 'C-Code' is in the draft.

-----

A tangent discussion took place as a result of the lunch interruption... Dealing with Plug-present, a DC connection and to tone or not to tone and what about state preservation during suspend, when does a loop free build take place, etc. **AR Steve Bard:** [Search through e-mail archives for the thread in which the decision was made to always tone and not rely upon the existence of a DC connection and include that "decision rationale" in these minutes.](#)

Steve found the e-mail from Colin. The e-mail has been inserted (below):

*From: Colin Whitby-Strevens [colin@zayante.com]  
Sent: Saturday, September 11, 1999 6:55 PM  
To: Upstarts (E-mail)  
Subject: Connection debounce vs DC connections*

*In our last SCAT meeting, I was asked to make the connection debounce more robust.*

*This has caused me to reveal a problem*

*Currently, the code makes use of the 1394a-style DC connection circuit. If a DC connection is detected, then it uses this circuit rather than toning to maintain connection status during suspend and disable.*

*Firstly, an observation on the behavior of the DC connect detect comparator, as defined in 1394a. The RC constant for discharging the capacitor for the "connected" state is of the order of 1.5 milliseconds. Once in this state, the RC constant for charging the capacitor (depending on the circuit in the local PHY) to transition into the "disconnected" state is of the order of 100 milliseconds.*

*As I understand it, this means that "short" disconnects due to connection bounce won't be reported to the connection\_status() logic. The connection\_status() logic look as as if the timer is restarted for each "bounce", but, in effect, CONNECT\_TIMEOUT is simply a time to allow the connector bouncing to settle down. I'm not sure if that is really the intention, but so be it.*

*Sidebar:- I'm using this as an excuse not to try to put any clever connect debounce logic in whilst toning (just keep on toning during a CONNECT\_TIMEOUT interval, but don't check that every tone is received).*

*In order to increase the connection debounce robustness, I've tried to incorporate the logic from 1394a in the case of a DC connection being detected. The problem is that if a DC connection is made with a peer port which is toning (which is likely) then each tone will come with its own bias which will make the DC connect comparator look false (i.e. "disconnected"). With a CONNECT\_TIMEOUT of 330ms, and a tone interval of 41 ms, then this looks all too likely.*

*It seems to me that either*

*1) we change the logic to stop toning when we detect a putative DC connection, let the DC connection debounce, and then start toning again.*

or

2) we don't try to use the DC comparator at all as soon as we hear a tone, and we don't try to maintain DC connected status in Beta mode.

*The first option seems to be a big change in the connection algorithm to me, and rather risky. What if the far end has not heard my tone yet? What if he's only connected but not powered on? Its probably doable, but there seems to be a lot of corner cases to work through.*

*The second option seems easier, but requires a couple of changes:-*

*i) if the port is in Beta mode then we'd use toning during suspend even if physically DC connected. Given that we do this if AC connected, it seems to me not to be a big change.*

*ii) currently, if we hear a tone, but we fail to negotiate a speed (my max is less than your min, or vice versa), but we're DC connected, then we drop back to DS mode. This facility now has to be withdrawn. I'm not sure why it was ever there - perhaps in the days when there was the possibility of using legacy connectors - or is it needed for Device bay???*

*The second option also has the advantage that Beta-mode only ports need not implement the DC connect detect circuitry at all.*

*I'm implementing option 2, as it is the smallest change to the current code, whilst any further discussion takes place.*

*With best wishes*

*Colin*

--

*Colin Whitby-Strevens      Email: colin@zayante.com  
Zayante Inc                      Web: www.zayante.com  
269, Mt Hermon Rd, Suite 201 Tel: +1 831 461 4948  
Scotts Valley                      Mobile: +1 831 818 4028  
CA 95066 USA                      Fax: +1 831 461 1394*

-----  
Continuing with Cycle start starvation: IEEE P394b will do it just like 1394a. ENABL\_ACCEL in 1394b PHY's definition shall be the same as 1394a. If the link sets ENABL\_ACCEL it shall provide notification of cycle start due to the PHY. This behavior is the same as that which occurs in 1394a - If ENABL\_ACCEL is set, a1394b link shall send CYCLE\_START\_EXPECTED. The PHY knows a cycle start has been received when it receives or originates a cycle start token - it may then start accelerations again.

3) Arb state timeout - Jerry said a separate has nothing to do with the minutes. Don't even raise that specter. Colin pointed out that a naïve implementation might use a separate timer for Arb state timeout.

**GROUP AR: Upon Reviewing Chapter 13 (PHY/Link Chapter) The group is to make certain the link shall properly tag the end of a packet that marks the end of a sub-action.**

In a hybrid bus BOSS restart time is established by the gap count and is the same value as sub-action gap. The senior border node is responsible for performing the BOSS restart time.

In a beta only bus, the responsibility of BOSS restart time is the last node transmitting BOSS and the time shall be a function of gap count. In a beta only bus, the gap count might be considered always fixed - as opposed to a gap count based upon bus topology. In a beta only bus there is a local notion of GAP count. The local notion is that it is either always 63<sub>10</sub> or it is whatever the gap count was set to.

A PHY that receives an S400 packet on one port will send data null out any and all ports whose speed is less than S400.

**A moment of "education" for those in need:**

**One over BASE\_RATE is, always has been and ever will be approximately 10 nanoseconds. Gap count is measured in units of (16/BASE\_RATE).**

**AR Steve Bard: Contact Dave Johnson request bus holders information to be sent to SCAT NOW! To avoid Ballot Review Comment.** *[STATUS 10/28/99: E-mail has been sent with urgent request to try to have info sent to SCAT before meeting end this week... STATUS: 10/29/99: Dave has responded and will follow-through by providing the appropriate material to the SCAT for inclusion as another implementation option for solution in the PHY/Link section. The material will not arrive prior to closure of this SCAT meeting.]*

4) Point-to-Point Packet Protocol: David Wooten will write proposed text to place in the draft that will outline how the point to point protocol packet will be protected. The Packet\_prefix symbol has been placed back into the symbol pool.

5) Resume issue on Sony Camcorders (One Second TpBias Resume):  
MAX\_BIAS\_HANDSHAKE shall be defined to be (12\*RESET\_DETECT) - i.e.  
MAX\_BIAS\_HANDSHAKE shall result in TpBias to be asserted for the same time duration as it is in 1394a.

Meeting adjourned: 7:20 PM

Meeting convened n Friday somewhere around 9:00 AM...

After random discussion of the work David did on Section 13 last evening, the group began to examine the SCAT issues table.

SCAT Issue #	RESOLUTION
65	"B" PHY-s shall implement this rule: They will never pass a grant received from a junior port to either its link or another junior port when the following condition is true: 1) the bus is in the async phase && 2) cycle master is outside of the local cloud (see glossary in the draft spec for definition of "cloud") && 3) a cycle start is expected. The PHY needs help from the link for requirement #3. If there is no active link or the ENABLE_ACCEL flags is set to zero, the PHY always assumes a cycle start is expected. If there is an active link and the ENABL_ACCEL flag is set to one the link is responsible for sending CYCLE_ST_DUE indication to the PHY to mark the condition of an expected cycle start - i.e. requirement #3 becomes true. When the PHY generates or receives a cycle start token requirement #3 is false. To be clear on requirement #2, the question is how far does the cloud extend, i.e. are border nodes part of the B cloud. The issue is that requirement #2 is true if cycle master is senior border. Alternatively, it is preferred that the condition for requirement #2 is established by observing the self-ID packets that came from the root.
63	Colin put out onto the reflector some words which are to be put into the Connection Management Chapter. Colin has the words in his copy and needs to send them to Eric for inclusion in the draft. If it sees loss of sync before S1 or S2 then it has been told to enter into the untested state.
13	Z-CLOSED
32	Based upon SCAT Issue #30, change this to Y-DONE All work is complete and it is in the draft.
36	Y-DONE All work is complete and it is in the draft.
53	Colin to put into the 'C-Code' the resolution for #53 as documented in these meeting. This becomes C-AIP
55	Change to AIP because it is in the 'C-Code' Mike Teener must document this is the text portion of the draft.

SCAT Issue #	RESOLUTION
57	No Action taken at this meeting
58	Y-DONE All work is complete and it is in the draft.
59	Z-CLOSED
62	Y-DONE All work is complete and it is in the draft.
6	Y-DONE All work is complete and it is in the draft.
18	Y-DONE All work is complete and it is in the draft.
19	Y-DONE All work is complete and it is in the draft.
27	Bard & Wooten have not completed their action items. David must submit PIL/FOP summary and Steve must submit standby summary to Eric ASAP.
40	Status C-AIP. Based on the SCAT Meeting minutes in August 1999, there will be a minimum of four symbol times between packets. The Statement in the SCAT "Next Actions" that 320 nanoseconds between packets is incorrect. What is guaranteed is that there will be four symbol times (at the maximum speed of the link) between any packet (PHY or other wise).

That which the group accomplished yesterday (regarding cycle start starvation) doesn't work when the PHY/Link interface is reset or disabled. Therefore, we really do need a learning bit so the PHY can learn that the link is cycle timer aware. The reason for this is that in 1394a an ENABL\_ACCEL only affects requests coming from the local link where now accelerations affect requests coming from elsewhere as well. So we have a learning bit that is cleared on power-on, power reset, on a PHY/Link interface reset (or disable). The learning bit gets set when the link issues CYCLE\_ST\_DUE or a decelerate request (in the case of a 1394a link). The learning bit controls the use of 1394a style accelerations instead of ENAB\_ACCEL and controls whether the PHY believes a cycle start is due. See the SCTA resolution for issue #65. It controls that together with the accelerate bit.

Jerry's #1: Withdrawn (or cancelled) requests: On a hybrid bus, when you grant a senior port don't deny junior ports and leave DS ports idle. Only time this can happen is at the end of a packet. When you grant junior ports deny all others by sending data NULL or data prefix on DS ports. This can happen out of the blue or at the end of a packet. If a grant is received from a senior port and has nothing to do then generate a null packet (get everybody into receive state and back out again). This applies either at the end of a packet or out of the blue. Then BOSS end of packet actions will grant the senior (if there is still nothing to do) and by the new rule (above) will not deny the junior ports. If the grant is received from a junior port (at the end of a packet) and there is nothing to do then send it up to senior border. At the senior border he stays BOSS but before he can do anything he must arbitrate into the "A" cloud if the root is in the "A" cloud. A BOSS shall release DS ports if nothing to do by sending a Legacy format NULL packet. This is the resolution for SCAT Issue #53.

Meeting Adjourned at 1:00 PM