Minutes of SAS Physical Working Group teleconference, July 6, 2006 T10/06-314r0

Attendance

Mr. Bryan Kantack	Agilent
Mr. Kevin Marks	Dell
Mr. David Freeman	Finistar
Mr. Rob Elliott	Hewlett Packard Co.
Dr. Bill Ham	Hewlett Packard Co.
Mr. James Rockrorh	IBM
Mr. Harvey Newman	Infineon
Mr. Michael Jenkins	LSI Logic Corp.
Mr. Gabriel Romero	LSI Logic Corp.
Mr. Wei Zhou	Marvell
Mr. Jeff Choun	Marvell
Mr. David Geddes	Marvell
Mr. Yuriy Greshishchev	PMC-Sierra
Mr. Amr Wassal	PMC-Sierra
Mr. Henry Wong	PMC-Sierra
Mr. Alvin Cox	Seagate Technology
Mr. Massimo Pozzoni	ST Microelectronics
Mr. Benoit Mercier	ST Microelectronics
Mr. Doug Loree	Toshiba
Mr. Kevin Witt	Vitesse Semiconductor
Mr. Robert Kembel	

21 People Present

Agenda:

SAS-2 PHYSICAL address frame

http://www.t10.org/ftp/t10/document.06/06-301r0.pdf

Discussed the requirement to support G1 versus G2 or no particular speed. We looked at the speed negotiation sequences in Harvey's presentation for the discussion. It was suggested that maybe the G3 window could be done with the G2 speed rather than G1. I don't think solves anything as the backwards support only moves up one generation. Since G3 is expected to need some training prior to acceptable data transfer in some situations, the window space is defined but needs something besides ALIGN synchronization.

- FCAL has formalized two rates back as a requirement.
- OOB is a key issue that introduces some unique problems. How does it limit future speeds? Is it really a problem for optical transmission?
- Several interfaces support communication back to G1.
- Bill Ham mentioned the FCAL version of speed negotiation that requires independence of transmitter and receiver speed during the initial set-up.
- G2 may not be successful in the G3 window in the 10-meter external cable application.

An alternative is to keep existing frame structure for each speed supported and use the final window for training and locking. Not that much time is lost for the multiple 600uS windows.

Another alternative is to develop a new COMSAS that goes to a different speed negotiation sequence that might use out-of-band to communicate useful data as is trying to be done with the physical address frame.

Actions from July 6 call:

- Rob to update his proposal with some minor changes.
- Harvey and Amr to update proposals to include G1 or G2 negotiated rate in the G3 window.
- Rob to contact STA concerning backwards support and the issue being discussed about the G3 window.

The following items were not discussed on the call separately but are included for consideration at next week's meeting.

SSC

http://www.t10.org/ftp/t10/document.06/06-263r3.pdf

New bullets were not discussed in detail as the information is too new. Plan to have input on the next call.

* Is a 1/(non-power-of-two) ALIGN rate a good idea?

* Separate ALIGN rates to optimize performance or just one rate for all cases?

* What is the best term to replace "clock skew management" in 7.3?

* The average up-spreading cannot exactly match the average down-spreading in an SSC profile. How do we specify the accuracy? Does the ALIGN rate of 1/2048 provide enough extra ALIGNs to account for the inaccuracy?
Comment: The 100 ppm reference clock accuracy has the same affect. Round up for the buffer, but there still needs to be a reasonable limit on the upper side. Rob needs to explain what problem the question is trying to address.

* Make sure everyone agrees with the math on the last page, and that we're not off by a factor of 2 anywhere.

Old questions:

Are definitions for down spreading and center spreading acceptable?

ST indicated that they would rather train on an SSC signal rather than a non-SSC signal. Agere and Vitesse indicated that it is better to train on non-SSC, but there would be a need to verify signal integrity after SSC is turned on. Training with SSC enabled is probably possible, but this has no data to support that it can actually be done. **It was suggested that it be assumed on for now and if it has to be changed later, then it can change.** To change it later could be a significant design impact for drives. Expanders are expected to have independent control of SSC on PHY's already, so the impact would be less for them.
SSC can be enabled and disabled without significant impact to the transmitted signal if done at the zero-crossing but it may take several microseconds to make the change from SSC to non-SSC.

Should a minimum SSC range be specified? Rob indicated he wants a minimum specified, but most who voiced an opinion did not support a minimum setting. SATA initially had problems with SSC but these seem to be getting better over time. One comment made indicated that the SATA ranges seen by that person typically from 1000 ppm to 3000 ppm rather than the full 5000 ppm

allowed. Many thought it best to be a purchase specification requirement rather than a standard if a minimum value is desired.

Several comments were made regarding EMI and the SSC pattern. The pattern requirements still need some sort of clarification so that an issue of overrunning buffers is not caused, but also that the pattern is effective in reducing EMI. The “area under the curve” approach was mentioned. But it in itself can permit a square wave implementation that would cause a buffer overrun issue.

Speed negotiation sequence

http://www.t10.org/ftp/t10/document.06/06-295r1.pdf
http://www.t10.org/ftp/t10/document.05/05-397r4.pdf

- Reviewed the new presentations and had some concerns about the final negotiation window RCDT. Do expanders with many PHY’s need more than 300uS to process the information? Is 500uS enough?
- Should there be a fixed value or just start sending training pattern when ready?
- How should the configuration data be sent? Should it be a 32-byte packet, handled by new primitives, or some other option?
- What information should be included?

New questions:

How do we know that the address frame was correctly received and processed? Can there be some sort of handshake to verify?

Do we need to specify how options are downgraded after a failed speed negotiation? Alvin suggested that the system determine what to change rather than the target device since the system is more likely to be aware of what can be changed for improvement. It was mentioned that turning off SSC is one possible option, but if this made a system non-compliant, then the device should not be trying to turn off SSC for that application. There should be some sort of control specified, as if both ends are trying to make changes, a conflict may result.

Page 9 of 05-397r4 illustrates a case where one PHY does not support G1 speeds, however, the G3 frame requires support of G1 to transfer the payload of the address frame. This is a contradictory situation.

Additional items needing investigation/comment:

- Should SSC be on or off during receiver equalization setting?

Today’s comments indicated that there are both advantages and disadvantages to having SSC active during the initial setting process. Having it on while setting equalization is an untested item, but is probably possible. If setting is done while it is off, then the signal reception needs to be verified after it is turned on.

- Is it viable to make a drive have independent SSC control on the transmitters of its two ports? Independence is required to set the receiver equalization without SSC since one port may be operating prior to the other one performing speed negotiation.

It is possible in some designs, but an alternate suggestion of turning off SSC at a zero-crossing for both PHY's was proposed as an alternative. There may be some timing issues with a smooth disable of SSC.

- In the beginning of the final speed negotiation window, does there need to be an idle time or can both devices immediately start transmitting the training pattern? It is assumed that if G2 is required, the sequence would follow the SAS 1.1 standard.

A 300uS window was suggested in the 06-295. Since this is close to the existing RCDT of the other windows and minimal compared to the training interval maximum time, it was suggested to just use the existing RCDT time. Some indicated they would like to go ahead and start the training pattern when they were ready rather than at a set time. Additional feedback is needed regarding this. It was also stated that some expanders with many PHY's may have a problem getting the information processed in that amount of time if all the PHY's were trying to communicate with the processor at the same time.

- It is assumed that all expanders and initiators are capable of receiving downspread SSC. Are there any know exceptions?

LSI, Vitesse, Marvell: None

Still needs to be answered with a positive or negative response.

- Will an initiator or expander accept downspreading from a SAS device running at G1 or G2 speed?
-

LSI, Vitesse, Marvell: Yes

Still needs to be answered with a positive or negative response.

- If a phy transmitter has SSC disabled or is using downspread SSC only, it could get away with inserting fewer ALIGNs - 1/128 (the SATA ratio) would cover sending to either SAS or SATA phys with downspread SSC. Is that complication worth a 0.8% performance improvement? (e.g. at 6 Gbps, this is 4.77 MBps). Use 1/64 for all?

LSI, Vitesse, Marvell: Presently don't care. (May change)