

40Gbe Backplane Considerations

Andre Szczepanek

Texas Instruments

a-szczepanek@ti.com

11 March 2008

1

Supporters

- Arthur Marris – Cadence
- Ilango Ganga - Intel

What Backplane wants

- Maximum Re-use of 10GBASE-KR and associated IP
 - AN, FEC, TxEQ training

- Simple PCS that leverages existing 10GBASE-R IP
 - Proposed PCS solutions all appear overkill for backplane use
 - Would prefer simpler solution like a 64B/66B sync based alignment

10GBASE-KR reuse

- Extend Clause 73 Auto-Negotiation to 40Gbe
 - Straightforward, just need to add a technology ability bit
 - More an editorial/administrative task than a technical one
 - Use 10GBASE-KX4 AN as model.
 - Single channel (0) used to negotiate
 - Similar to 10GBASE-KR, make AN mandatory at 40G to avoid use of parallel detection
- Transmit equalizer training
 - Re-use 10GBASE-KR training
 - Mandate inter-lane randomness
 - The definition of training frame PRBS seeding already does this
 - Link will not be up until all lanes complete equalizer training
- FEC
 - The FEC is designed as a generic sub-layer below the 10GBASE-R PCS
 - Provides agnostic transport of 64B/66B code words
 - Can easily be used with either direct PCS output or striped 64B/66B
 - Error correction could be difficult to implement at non-striped rates
 - An Error-correction disabled option would allow FEC coding to be used for alignment without the penalty of FEC correction latency

64B/66B Sync based alignment

- The 64B/66B Sync bits can deskew 32UI
- Backplane link skew sources
 - Transmit SERDES
 - Backplane including connectors
 - Receive SERDES
 - Receive re-alignment
- This scheme could also be used as the basis of a XLAUI or CAUI extender interface
 - XLAUI/CAUI will be less demanding than backplane
 - 10inch, 1 connector vs 40inch 2 connectors

Transmit SERDES skew

- With appropriate design SERDES inter-lane transmit skew can be limited to analog skews
 - $<1UI$
 - Need to sync digital stages between lanes
- Where this has not been done we have seen up to $4UI$ of digital skew (in addition to the analog skew)

Backplane skew

- Backplane skew is caused by differences in channel length between the channels on
 - The line card
 - The connectors
 - The Backplane
- Total backplane channel length is ≤ 40 inches (1m)
- At 10Gbps 1inch of FR4 has <200 ps or ~ 2 UI of delay
 - 12UI of backplane skew budget would allow for 6inches of channel length variation between channels (15%)
 - Typical backplane interchannel skew is <1 ns (10UI)
 - Need input from more backplane manufacturers

Receive SERDES skew

- Analog skews
 - $<1\text{UI}$
- Quantization skew
 - Each lane deserializes data into words, each with its own word clock
 - Even if the word clocks are closely aligned, Sync bits may fall into different words due to other skews
 - This is seen as an effective skew of the word width
 - For wide buses, this can be significant (32bit) or prohibitive (64bit)

Alignment skew

- SERDES receive data must be gearboxed to the 66bit symbol width and aligned to symbol boundaries
 - If the gearbox and aligner are not combined then there can be up to 66 bits of quantization skew
 - If gearbox and aligner are combined the quantization skew is limited to the SERDES word size

Avoiding Quantization skew

- SERDES with bus width = $66/n$, and “jog”
 - Aligning the bus width with the symbol width and allowing the deserializer position to be “jogged” removes quantization skew
 - Commonly used for 8B/10B with a 10bit or 20bit word width.
 - Unfortunately 11,22,33,or 66bits are not common SERDES bus widths!
- Asynchronous Gearbox/Aligner
 - If the SERDES bus to 66bit bus gearbox is built as an asynchronous bit-wide circular FIFO with built in alignment mux, the quantization skew will be removed.
 - The gearbox is effectively a bit wide FIFO
 - loaded SERDES bus width bits at a time
 - unloaded 66bits at a time at a 0 to 65bit offset
 - Because the data has effectively been re-serialized the quantization is removed
 - The 66bit output word is offset (based on 10GBASE-R lock SM) to align to the Sync bits

Is 32UI enough for backplane applications ?

- YES, 64B/66B sync based alignment is sufficient for Backplane applications
 - 32UI of budget is adequate as long as quantization skew is managed
 - Receive path MUST be designed to limit quantization skew
 - Note that although 64B/66B has the ability to remove 32UI in either direction from a mid-skew position, we can't control skew distribution so individual skews must be limited to 32UI

Proposed 64B/66B alignment skew budget

- Transmit SERDES (Analog)
 - 2UI
- Backplane including connectors
 - 12UI
- Receive SERDES (Analog)
 - 2UI
- Receive gearbox & re-alignment (quantization skew)
 - 16UI
 - ====
 - 32UI
 - ====
- Assumptions
 - 16bit SERDES bus @644Mhz is practical in modern Si processes
 - Same as XSBI rate
 - SERDES with data widths wider than 16bits will have to limit quantization skew to 16 bits

Standardizing 64B/66B alignment

- Similar standards only specify electrical (backplane) skew
 - cf SFI4.2
 - They do not specify how implementations allocate rest of available budget
 - This allows maximum implementation flexibility
- Propose we require 12UI of skew tolerance at receiver
- We should also consider specifying maximum inter-lane skew for the channels
 - Although there is no normative 802.3ap channel, differential pair skew was specified.

Conclusions

- Initial skew numbers indicate 64B/66B alignment is sufficient for Backplane skew
 - And therefore also sufficient for XLAUI & CAUI
- 64B/66B Sync alignment allows a simple PCS that leverages existing 10GBASE-R IP
 - Simple 64B/66B word distribution
 - No lane markers or alignment symbols
- 64B/66B Sync alignment is a sub-set of MLD
 - MLD distribution without alignment markers