CDAUI-8 Chip-to-Module (C2M) System Analysis #3

Ben Smith and Stephane Dallaire, Inphi Corporation

IEEE 802.3bs, Bonita Springs, September 2015

Supporters

- Ali Ghiasi, Ghiasi Quantum LLC
- Marco Mazzini, Cisco
- Andre Szczepanek, Inphi
- John Ewen, Global Foundries

Introduction (1)

- We investigate the merits of various reference receiver architectures for C2M
  - Relative merits are evaluated on the basis of Channel Operating Margin (COM) of the full C2M link (rather than TP1a, which doesn't account for RX package reflections)

- We show that the CAUI-4 C2M Reference CTLE + LFEQ isn't "enough" to close higher loss links
  - **PAM4 is much more sensitive to residual ISI than NRZ**
    - PAM4 is more sensitive to ILD and package reflections
      - We need to do as good as job as possible on the "easily equalizable" part of the signal
    - A (1z,2p) is inadequate for closing higher loss links
    - A 2-tap TXFIR provides significant improvement in margin

Introduction (2)

C2M Link Margins

- Several contributions have been made, each using a different model and a different quantification of performance. Some results seem more optimistic than others—what gives??
- **Eye Height**:
  - Eye Height spec in draft 0.9 (50mV) is quite stringent for high loss channels
- **Transmitter SNR**:
  - At 29 dB (peak-to-rms), transmitter noise is a large impairment
    - But it seems clear that different contributions have made different assumptions about the definition (and modelling) of TX SNR
    - Current 56G VSR OIF draft does not provide a definition of TX SNDR, even though an informative TP0a value is provided
- **Package Model**
  - As seen in several C2C contributions (healey_3bs_01_0315, hegde_3bs_01_0715), the package model has a significant influence on PAM4 margins

System Model

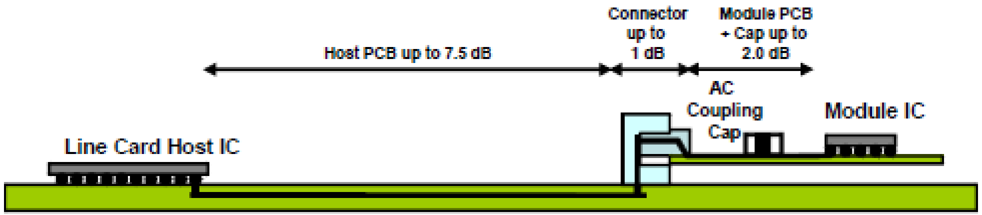

- TX and RX package models (.s4p file) each add ~1dB of IL @ 13.28125 GHz
- Die Termination with 120fF parasitic capacitance
- Module RX model:
  - (1z,1p) low-frequency equalizer (zero & pole ~1GHz)
  - (1z, 2p) reference CTLE (from OIF-VSR-56G PAM-4 and CAUI-4 C2M):

Peaking (dB)	G	P1/2π (GHz)	P2/2π (GHz)	Z1/2π (GHz)
1	0.891	18.6	14.1	8.31
2	0.794	18.6	14.1	7.10
3	0.708	15.6	14.1	5.68
4	0.631	15.6	14.1	4.98
5	0.562	15.6	14.1	4.35
6	0.501	15.6	14.1	3.82
7	0.447	15.6	14.1	3.43
8	0.398	15.6	14.1	3.00
9	0.355	15.6	14.1	2.67

System Model

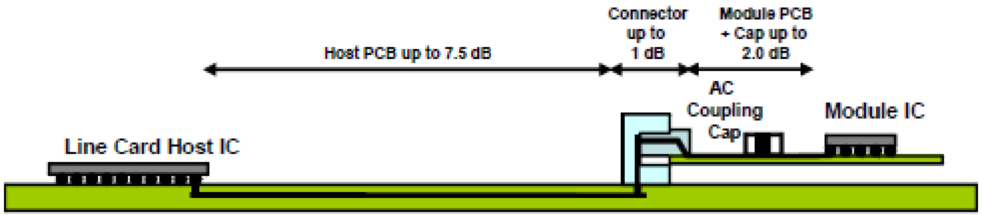

- Host TX model:
  - 750 mV differential peak-to-peak
  - TX SNR = 29 dB (peak-to-rms)
  - RLM = 0.9
  - RJ = 0.01 UIrms
  - DJ = 0.05 UI peak-to-peak
  - 2-tap TXFIR (i.e., pre+cursor)

Channel Models

CHANNEL	FEXT	NEXT	IL @ 13.28125 GHz (dB)	ILD (dBrms)
From IEEE 802.3bs shanbhag_3bs_14_0623:				
(1) Nelco 4000-13SI Host PCB + next gen 28Gb/s high density SMT IO	5	0	8.7	0.110
(2) EM-888 Host PCB + next gen 28Gb/s press-fit stacked IO	7	0	8.9	0.051
From IEEE 802.3bs shanbhag_3bs_01_1014:				
(3) 4in Megtron6 Host PCB + next gen 28Gb/s high density SMT IO	5	0	4.3	0.110
(4) 10in Megtron6 Host PCB + next gen 28Gb/s high density SMT IO	5	0	8.8	0.106
(5) 4in Megtron6 Host PCB + next gen 28Gb/s press-fit stacked IO	7	0	4.5	0.051
(6) 10in Megtron6 Host PCB + next gen 28Gb/s press-fit stacked IO	7	0	9.0	0.052
Cisco Channels:				
(7) Cisco 2in Stacked	0	0	8.5	0.237
(8) Cisco 5in Stacked	0	0	11.3	0.245

Baseline Results

- Reference CTLE Receiver
  - **No TXFIR, No LFEQ, DER_0=1E-6**

Channel	1	2	3	4	5	6	7	8
COM (dB)	-0.07	-0.04	1.01	-0.45	1.24	-0.13	-1.37	-2.65

- **Only the ~4dB channels have positive margin**

Improvements (1)

- Reference CTLE + LFEQ
  - COM program optimizes LFEQ: $0.5\ \text{GHz} \leq z \leq 2.5\ \text{GHz}$, $0.5\ \text{GHz} \leq p \leq 2.5\ \text{GHz}$
  - **No TXFIR, DER_0=1E-6**

Channel	1	2	3	4	5	6	7	8
CTLE	-0.07	-0.04	1.01	-0.45	1.24	-0.13	-1.37	-2.65
CTLE + LFEQ	0.45	0.50	1.39	-0.14	1.92	0.27	-1.37	-2.49

  - LFEQ improves COM margin by 0.4 to 0.5 dB in most cases

Improvements (2)

- Reference CTLE + TXFIR
  - COM program optimizes TXFIR: $|C_{-1}| \leq 0.15, |C_{-1}| + |C_0| = 1$
  - **No LFEQ, DER_0=1E-6**

Channel	1	2	3	4	5	6	7	8
CTLE	-0.07	-0.04	1.01	-0.45	1.24	-0.13	-1.37	-2.65
CTLE + TXFIR	1.47	1.53	1.43	0.84	2.08	1.35	0.84	0.55

- A 2-tap TXFIR brings significant improvement on higher loss channels
  - Improvement is > 1dB for high loss channels

Improvements (3)

- Reference CTLE + TXFIR + LFEQ
  - COM program optimizes TXFIR and LFEQ; **DER_0=1E-6**

Channel	1	2	3	4	5	6	7	8
CTLE	-0.07	-0.04	1.01	-0.45	1.24	-0.13	-1.37	-2.65
CTLE + TXFIR	1.47	1.53	1.43	0.84	2.08	1.35	0.84	0.55
CTLE + LFEQ	0.45	0.50	1.39	-0.14	1.92	0.27	-1.37	-2.49
CTLE + TXFIR + LFEQ	2.26	2.50	2.13	1.28	2.95	2.14	1.43	0.84

- The combination of the CTLE, LFEQ and 2-tap TXFIR provides substantial improvement over a CTLE-only system
  - CTLE+TXFIR or CTLE+LFEQ do not provide sufficient margin
  - For high loss channels, adding TXFIR and LFEQ improves COM margin by 2dB or more

An Improved Reference RX/TX

- **The following (crudely) improved reference RX/TX provides nearly all of the gain:**

TX FIR	LFEQ: (Z1,P1) (GHz)	CTLE: (Z1,P1,P2) (GHz)
[-0.05,0.95]	(1,1.2)	(8.31,14.1,18.6)
[-0.05,0.95]	(1,1.2)	(7.10,14.1,18.6)
[-0.05,0.95]	(1,1.2)	(5.68,14.1,15.6)
[-0.05,0.95]	(1,1.2)	(4.98,14.1,15.6)
[-0.1,0.9]	(1,1.2)	(4.35,14.1,15.6)
[-0.1,0.9]	(1,1.2)	(3.82,14.1,15.6)
[-0.1,0.9]	(1,1.2)	(3.43,14.1,15.6)
[-0.1,0.9]	(1,1.2)	(3.00,14.1,15.6)
[-0.1,0.9]	(1,1.2)	(2.67,14.1,15.6)

Channel	1	2	3	4	5	6	7	8
CTLE	-0.07	-0.04	1.01	-0.45	1.24	-0.13	-1.37	-2.65
CTLE + TXFIR	1.47	1.53	1.43	0.84	2.08	1.35	0.84	0.55
CTLE + LFEQ	0.45	0.50	1.39	-0.14	1.92	0.27	-1.37	-2.49
CTLE + TXFIR + LFEQ	2.26	2.50	2.13	1.28	2.95	2.14	1.43	0.84
Reference RX/TX	2.22	2.47	2.13	1.28	2.95	2.14	1.18	0.19

- The degradation on 8 is due to insufficient pre-cursor equalization in the reference TX FIR

A Universal TXFIR?

- **Is there a single fixed setting of the TXFIR that is “good enough” for all channels?**

Channel	1	2	3	4	5	6	7	8
0% pre-cursor	0.45	0.50	1.39	-0.45	1.92	0.27	-1.41	-2.65
2.5% pre-cursor	1.19	1.36	1.98	0.57	2.61	0.98	-0.50	-1.95
5% pre-cursor	1.94	2.11	2.19	1.19	2.95	1.86	0.15	-1.20
7.5% pre-cursor	1.98	2.35	1.99	1.22	2.87	1.96	0.87	-0.52
10% pre-cursor	2.26	2.50	1.78	1.28	2.35	2.14	1.39	0.36
12.5% pre-cursor	2.14	2.21	1.31	1.19	1.85	1.81	1.43	0.60

- For channels 1 through 6:
  - 5%-pre, 7.5%-pre and 10%-pre each provide a reasonable performance
- For Channels 7 and 8:
  - 10%-pre and 12.5%-pre provide the best performance
- 10%-pre is best single fixed setting option,
  - However only 4 channels meet >2dB margin

CTLE+LFEQ+10%-Pre: Margin versus TX SNR & SER

Channel	1	2	3	4	5	6	7	8	3”	4”	5”
(29 dB, 1E-6)	2.26	2.50	1.78	1.28	2.35	2.14	1.39	0.36	2.12	1.65	1.27
(31 dB, 1E-6)	3.29	3.60	2.68	2.09	3.37	3.13	2.20	1.04	3.11	2.53	2.11
(29 dB, 1E-5)	3.15	3.39	2.67	2.15	3.27	3.03	2.29	1.18	3.01	2.54	2.11
(31 dB, 1E-5)	4.15	4.46	3.55	2.94	4.26	4.00	3.08	1.84	3.97	3.39	2.93

- 3”, 4” and 5” correspond to Cisco Rev 4
  - MXP cables de-embedded
  - IL: 3” ~ 9.3dB @13.3GHz; 4” ~ 10.3dB; 5” ~ 11.6dB
  - ILD (dBrms; as per 28G-VSR calculation): 3” ~ 0.16; 4” ~ 0.14; 5” ~ 0.097

C2M Link Margins: Eye Height

- In 802.3bj, a COM margin of 3 dB was considered sufficient for channel compliance
- In 802.3bm, a COM margin of 2dB was considered sufficient
- In Draft 0.9, Eye Height is set to 50mV
  - This is **stringent** for high loss channels, corresponding to a COM much larger than 3dB
    - Example 1:
      - TX Output: 900 mV pk-to-pk; R_{LM}=0.9; PAM levels: (+/-180 mV,+/-450 mV)
      - Equalization of 10dB channel loss (plus TX package losses) scales TX levels by factor of ~2.5
      - Received levels (with perfect TX linearity): (+/- 72, +/- 180)
        - A 50 mV eye opening corresponds to a COM of $20\log_{10}\frac{54}{54-25} = 5.4$ dB
- For reference, the same calculation for 28G-VSR results in a COM of $20\log_{10}\frac{180}{180-47.5} = 2.7$ dB

C2M Link Margins: SNR

- TX SNR is one of the largest impairments, but it has not even been defined for C2M (or for 56G VSR)
  - KP4 COM
    - At the transmitter output, TX **SNR** is defined as ratio of peak transmitter level to rms noise at transmitter output; noise is modelled as purely Gaussian
    - PSD of noise/distortion is not explicitly constrained
      - COM assumes that this noise is “passed through” to the slicer, in the sense that it is modelled as a slicer-referred peak-to-rms noise
        - This is reasonable for CTLE-based systems, as long as the bandwidth of the noise at the TX output is approximately limited to the RX bandwidth, and the receiver approximately inverts the channel
  - 802.3bj/bm
    - TX **SNDR** is based on TP0a measurement, and includes various contributions: (linear) residual ISI, distortion, Gaussian noise
      - It may be argued that it is pessimistic to model this via a purely Gaussian distribution (with the same variance), but TX SNDR and TX SNR were always set to the same value in KP4, KR4, and CAUI-4 C2C
        - Furthermore, the interference tolerance test directly uses the measured value in COM

C2M Link Margins: TX SNR

- For the previous model (i.e., an effective slicer-referred noise), a 29dB SNR results in ~50% eye closure @1E-6 for PAM4, in absence of other impairments
  - Calculation:
    - Normalized PAM levels = [+/-1/3,+/-1]
    - RMS noise = 10^(-29/20) = 0.0355
    - 1E-6 contour is approximately 4.75-sigma of a Gaussian
    - Relative Eye Opening = 1- (2*4.75*0.0355)/(2/3) = 0.49
- Semtech results (frlan_01_082415_elect) showed EH6 > 50mV in several cases, but seemingly used a different model (or definition) for TX noise and distortion
    - For example, Slide 16 shows eye opening of ~75mV, which is well beyond the 50% opening for the stated TX/RX parameters, without even accounting for contribution of residual ISI
      - The same conclusion can be made for the other Semtech results, where residual ISI is an additional significant contributor to eye closure
    - Note that Semtech results assumed perfect eye linearity and no xtalk

Recommendations

- LFEQ+CTLE is **not enough** to close the link for higher loss channels in our simulations
  - However It is possible to achieve >2dB margin on revised channel set by a combination of:
    - Using a 2-tap TX FIR with a fixed/universal 10% pre-emphasis tap
    - Tightening informative TX SNR to 29dB
    - Relax the symbol error rate target, to reflect DFE-less receiver, to 1E-5

- We are proposing:
  - Reference Receiver: Draft 0.9 CTLE + Fixed LFEQ
  - TX FIR is an implementation option that may be required to meet TP1a & TP4 specifications
    - TX FIR on module transmitter should be fixed and not require configuration
  - Relax the symbol error rate target to reflect DFE-less receiver
  - Tighten informative TX SNR
  - Limit ILD
- We need an agreed upon definition and model for TX SNR