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

Stephane Dallaire and Ben Smith, August 24th 2015

Introduction

- We investigate the merits of various reference receiver architectures for 26.5625GBaud PAM4 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 investigate the benefits of a 2-tap TXFIR
  - To avoid auto-negotiation or other TXFIR optimization schemes, we show that a "coarse" 2-tap TXFIR (10%, 5%, or 0% pre) provides most of the benefit of pre-equalization
- Similar to Hedge et al. (who focused on C2C channels):
  - A low-frequency equalizer (LFEQ) is shown to be beneficial
  - Reducing the target PAM4 symbol error rate is beneficial, and well-motivated due to non-bursty nature of error events
- We discuss the choice of the target PAM4 symbol error rate for a CTLE-based RX
  - We shouldn't be too pessimistic!

System Model

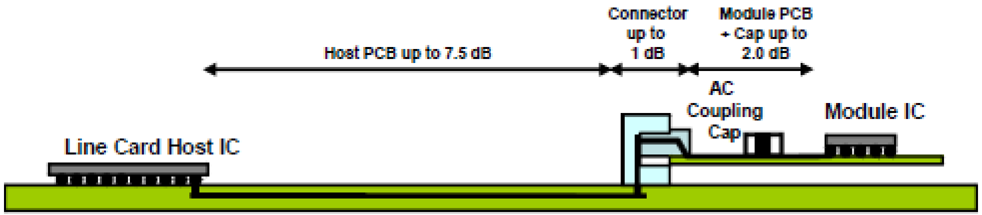

- TX and RX package models (.s4p file) each add ~1dB of IL @ 13.28125 GHz
- Die Termination with 120fF parasitic capacitance
- Module RX model:
  - (1z,1 p) low-frequency equalizer (zero & pole ~1GHz)
  - (1z, 2p) reference CTLE (from OIF-VSR-56G PAM-4 and CAUI-4 C2M):

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

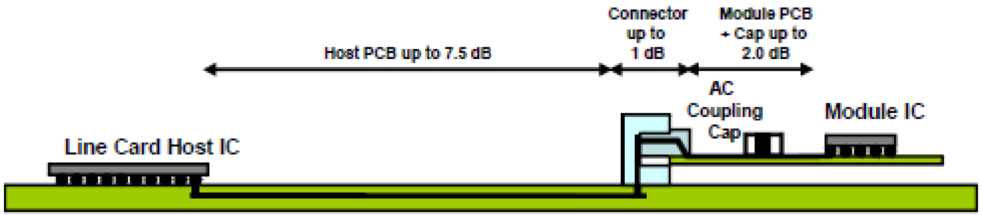

- Host TX model:
  - 750 mV differential peak-to-peak
  - SNDR = 29 dB (peak-to-rms)
  - RLM = 0.9
  - RJ = 0.01 UIrms
  - DJ = 0.05 UI peak-to-peak
  - 2-tap TXFIR (i.e., pre+cursor)

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

Link Margin Calculation

- The COM definition of margin is a quantification of the Vertical Eye Opening (VEO)
  - $\mathrm{COM} \equiv \mathrm{VEO} \triangleq 20\log_{10}\left(\min\left\{\frac{Av_{upp}}{Av_{upp}-v_{upp}}, \frac{Av_{mid}}{Av_{mid}-v_{mid}}, \frac{Av_{low}}{Av_{low}-v_{low}}\right\}\right)$
  - Eye contours are measured for a target symbol error rate **DER_0**

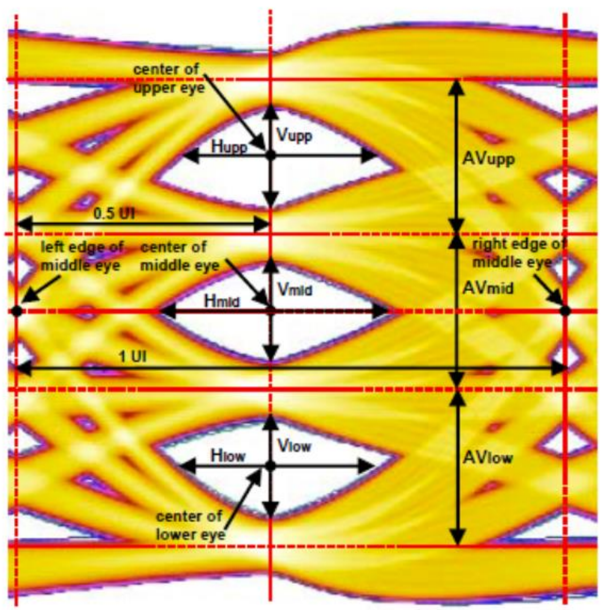

Figure 16-9. TP1a and TP4 jitter and Eye Height parameters

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 + TXFIR
  - COM program optimizes TXFIR: $|C_{-1}| \leq 0.15, |C_{-1}| + |C_0| = 1$
  - **No LFEQ, $\mathbf{DER_0}$=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
  - Residual ISI cancellation is critical to PAM4

Improvements (2)

- Reference CTLE + TXFIR + LFEQ
  - COM program optimizes TXFIR and LFEQ :
    $0.5\text{ GHz} \leq z \leq 2.5\text{ GHz}, 0.5\text{ GHz} \leq p \leq 2.5\text{ GHz}$
  - **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 + TXFIR + LFEQ	2.26	2.50	1.99	1.28	2.95	2.14	1.43	0.84

- LFEQ approximately 0.5 dB to 1 dB of margin improvement in most cases
  - LFEQ is useful to minimize residual ISI in the neighborhood of post-cursors approximately 2 to 5 bauds the cursor

Improvements (3)

- Reference CTLE + TXFIR + LFEQ + (**DER_0 = 1E-5**)
  - DER_0=1E-6 is unnecessarily stringent
  - Errors for proposed receiver are approximately independent identically distributed (or, at least, significantly less bursty than DFE-based RX).
  - From anslow_3bs_03_0515 (slide 17), for electrical sub-links with random errors, **bit**-error-rate ~= 8.2E-4 **per** link, for a 0.1dB penalty in the optical link
    - DER_0 = 1E-5 corresponds to BER=5E-6, so still a conservative target

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 + TXFIR + LFEQ (1E-6)	2.26	2.50	1.99	1.28	2.95	2.14	1.43	0.84
CTLE + TXFIR + LFEQ (1E-5)	3.15	3.39	2.89	2.15	3.87	3.03	2.33	1.72

Practical Considerations

CTLE+TXFIR+LFEQ

- This implies a (potentially) complex optimization…
  - But most of the gain be obtained with fixed settings:
    - **High Loss**: TXFIR = [-0.1,0.9]; LFEQ: (z = 1 GHz, p = 1.2 GHz)
    - **Medium Loss**: TXFIR = [-0.05,0.95]; LFEQ: (z = 1 GHz, p = 1.2 GHz)

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 + TXFIR + LFEQ (1E-6)	2.26	2.50	1.99	1.28	2.95	2.14	1.43	0.84
CTLE + TXFIR + LFEQ (1E-5)	3.15	3.39	2.89	2.15	3.87	3.03	2.33	1.72
CTLE + TXFIR**[-0.1,0.9]** + LFEQ(1G,1.2G) + (1E-5)	3.10	3.35	1.31	2.15	2.34	3.02	2.06	1.00
CTLE + TXFIR**[-0.05,0.95]** + LFEQ(1G,1.2G) + (1E-5)	2.22	2.51	3.03	1.42	3.87	2.15	0.85	-0.51

Conclusions

- A 2-tap TXFIR can be set coarsely, as a function of channel loss, with little loss in performance
  - High Loss (~10dB): [-0.1, 0.9]
  - Medium Loss (~5dB): [-0.05, 0.95]
  - Low Loss: [0,1]
- A reference receiver with an LFEQ provides additional ISI cancellation
- The target PAM4 symbol error strongly influences the system margin
  - Relative to a DFE-based system (e.g., C2C), the target should be decreased, due to absence of long bursts
- For channels with similar IL @ 13.28125 GHz, margin is influenced by ILD
  - Cisco channels have ILD (dBrms) approximately twice as large as the worst of the other six channels

Backup Slides

Improvements (4)

- Increasing the reference receiver bandwidth provides only a minor improvement in link margin, in a few cases
  - Reference CTLE poles and zero scaled by 10%
    - **Scaled** Reference CTLE + TXFIR + LFEQ, DER_0=1E-5

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 + TXFIR + LFEQ (1E-6)	2.26	2.50	1.99	1.28	2.95	2.14	1.43	0.84
CTLE + TXFIR + LFEQ (1E-5)	3.15	3.39	2.89	2.15	3.87	3.03	2.33	1.72
Scaled CTLE + TXFIR + LFEQ (1E-5)	3.40	3.64	3.11	2.49	3.75	2.88	2.36	1.87

Comments on EV6 and EH6

- In 802.3bj, a COM margin of 3 dB was considered sufficient for channel compliance
- In current OIF draft, EH6 is set to 50mV
  - This is quite stringent for high loss channels, corresponding to a COM larger than 3dB
    - Example 1:
      - TX Output: 750mV peak-to-peak; PAM levels: (+/-125 mV,+/-375 mV)
      - Equalization of 10dB channel loss (plus TX package losses) scales TX levels by factor of ~2.5
      - Received levels (with perfect TX linearity): (+/- 41.67, +/- 125)
      - A 50 mV eye opening corresponds to a COM of $20\log_{10}\frac{41.67}{41.67-25} = 8$ dB
    - Example 2 (Same Channel & EQ as Example 1):
      - TX Output: **1V** peak-to-peak, **R_{LM}=0.9**; PAM levels: (+/-200 mV,+/-500 mV)
        - Received levels (with perfect TX linearity): (+/- 80, +/- 200)
        - A 50 mV eye opening corresponds to a COM of $20\log_{10}\frac{60}{60-25} = 4.7$ dB