Chip-to-module far-end TX eye measurement proposal

Raj Hegde & Adam Healey

IEEE P802.3bs 400 Gb/s Ethernet Task Force
March 2017 Vancouver, BC, Canada

Background

- In smith_3bs_01a_0915, it was shown that the module transmitter needs to provide fixed pre-cursor de-emphasis to close the link budget for a number of chip-to-module channels
- This need also motivated hegde_3bs_01_0116, hegde_01_042516_elect and hegde_3bs_02_0516. These contributions formed the basis for the content in 120E.3.2.1.1.
- It was assumed that including a loss channel in the far-end eye measurement methodology would capture the pre-cursor component requirement
- **Was this assumption correct? This is the subject of comment i-91.**

BROADCOM

Questions

- Is it possible for the module transmitter to pass the far-end eye requirements without fixed pre-cursor equalization? **Answer: Yes**
- If the far-end eye opening requirements are met, does it matter if pre-cursor equalization is provided or not? **Answer: Yes**
- How might we better enforce the pre-cursor equalization requirement?

BROADCOM

Far-end eye with hypothetical module transmitter

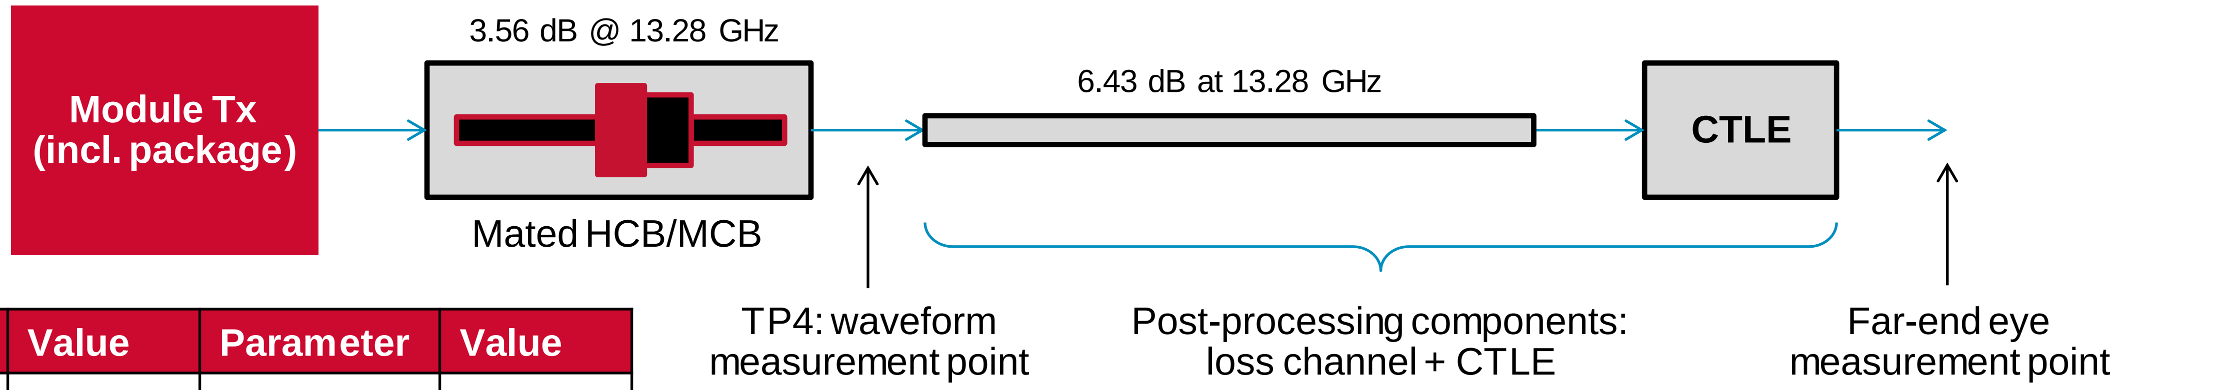

$$VEO = 2A_s\left(1 - 10^{-COM/20}\right)$$

Parameter	Value	Parameter	Value
T_r (β = 2)	13 ps	R_d	45 Ω
A_v	0.45 V	C_d	90 fF
SNR_{TX}	35 dB	z_p	10 mm
R_{LM}	0.98	Z_c	90 Ω
A_{DD}	20 mUI	C_p	90 fF
σ_{RJ}	10 mUI		

Procedure (using COM):

- Measure the waveform at the output of module compliance board
- Apply Bessel-Thomson low-pass response with 33 GHz bandwidth
- Apply “loss channel” (151 mm PCB) and reference CTLE
- Search over all CTLE gain settings to obtain the best eye opening

- With [c(−1), c(0), c(1)] = [0, 1, 0], the vertical eye opening (VEO) is 35.8 mV
- COM-based example supports lab findings (note that eye width and ESMW requirements also met)

Note: Mated HCB/MCB model from http://www.ieee802.org/3/bs/public/channel/mccom/diminico_3bs_01_0516.s4p.

Two transmitters with similar vertical eye opening

Module transmitter A (slide 4)

Parameter	Value	Parameter	Value
T_r ($\beta = 2$)	13 ps	R_d	45 Ω
A_v	0.45 V	C_d	90 fF
SNR_{TX}	35 dB	z_p	10 mm
R_{LM}	0.98	Z_c	90 Ω
A_{DD}	20 mUI	C_p	90 fF
σ_{RJ}	10 mUI		

Module transmitter B

Parameter	Value	Parameter	Value
T_r ($\beta = 2$)	13 ps	R_d	45 Ω
A_v	0.45 V	C_d	210 fF
SNR_{TX}	31 dB	z_p	12 mm
R_{LM}	0.95	Z_c	90 Ω
A_{DD}	20 mUI	C_p	110 fF
σ_{RJ}	10 mUI		

Test case	[c(−1), c(0), c(1)]	Far-end	
		COM, dB	VEO, mV
Transmitter A1	[0, 1, 0]	3.62	35.8
Transmitter A2	[−0.1, 0.9, 0]	8.57	61.3
Transmitter B	[−0.1, 0.9, 0]	4.56	36.3

Same VEO to within 0.5 mV

- Set transmitter B parameters closer to Table 120D-8 values to obtain ~36 mV vertical eye opening with fixed pre-cursor de-emphasis

BROADCOM

Emulate full module-to-host link

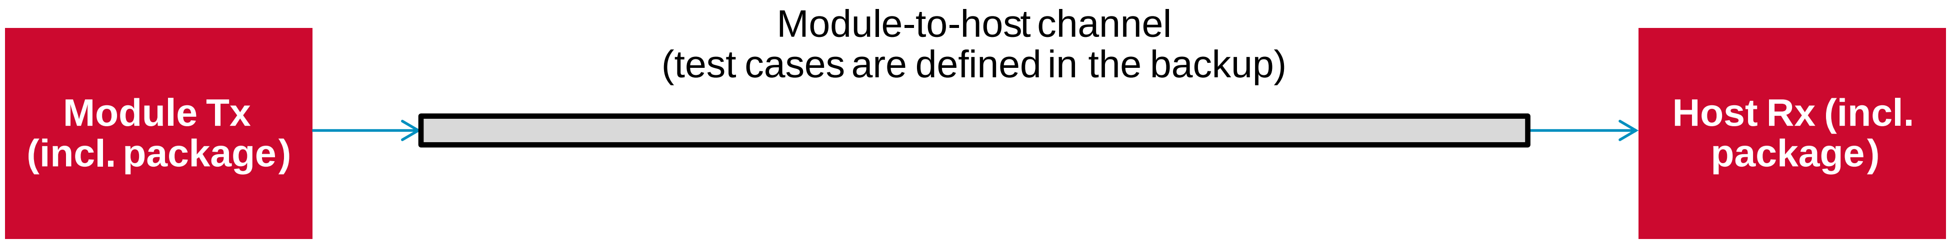

- Module transmitter A1, A2, or B

Parameter	Value	Parameter	Value
R_d	55 Ω	η_0	2.6 x 10^{-8} V2/GHz
C_d	280 fF		
z_p	30 mm		
Z_c	90 Ω		
C_p	110 fF		

- Chip-to-module receiver
  - Reference CTLE defined in 120E.3.1.7 (plus 33 GHz Bessel-Thomson low-pass filter)
- Chip-to-chip receiver
  - 2-stage CTLE and decision feedback equalizer (DFE) as defined in 120D.4

Compare transmitters A1 and A2 with chip-to-module receiver

* [c(-1), c(0), c(1)]

Test case		1	2	3	4	5	6	7	8	9
IL at 13.28 GHz, dB		8.74	8.94	4.29	8.81	4.5	9.01	9.28	10.29	11.61
COM, dB	A1: [0, 1, 0] *	0.89	1.11	2.14	0.39	3.25	0.6	−0.22	−0.77	−1.42
	A2: [−0.1, 0.9, 0]	3.27	4.39	4.27	3.11	5.05	4.19	2.84	2.56	1.25
Penalty for no pre-cursor, dB		2.4	3.3	2.1	2.7	1.8	3.6	3.1	3.3	2.7
VEO, mV	A1: [0, 1, 0]	8.1	10.2	28.4	3.7	39.5	5.6	0.0	0.0	0.0
	A2: [−0.1, 0.9, 0]	23.8	30.2	47.1	22.8	54.0	28.8	21.1	17.3	8.3
Loss of VEO for no pre-cursor, mV		15.7	20.1	18.7	19.1	14.5	23.3	21.1	17.3	8.3

- Transmitter A2 has the fixed pre-cursor emphasis but transmitter A1 does not
- Significant reduction in margin without the fixed pre-cursor de-emphasis
- Test cases 1 to 6 have low FEXT and 7 to 9 have no crosstalk

Compare transmitters A1 and A2 with chip-to-chip receiver

* [c(-1), c(0), c(1)]

Test case		1	2	3	4	5	6	7	8	9
IL at 13.28 GHz, dB		8.74	8.94	4.29	8.81	4.5	9.01	9.28	10.29	11.61
COM, dB	A1: [0, 1, 0] *	4.83	4.76	4.72	4.99	5.04	5.17	4.02	4.18	4.22
	A2: [−0.1, 0.9, 0]	8.39	8.9	6.8	7.63	7.5	8.17	6.99	7.73	7.36
Penalty for no pre-cursor, dB		3.56	4.14	2.08	2.64	2.46	3	2.97	3.55	3.14
VEO, mV	A1: [0, 1, 0]	31.5	28.4	60.1	32.5	51.6	30.6	24.5	22.3	18.9
	A2: [−0.1, 0.9, 0]	65.1	55.0	86.0	62.2	91.9	69.7	61.6	48.8	40.5
Loss of VEO for no pre-cursor, mV		33.7	26.6	25.8	29.7	40.3	39.1	37.1	26.5	21.6

- Significant reduction in margin without the fixed pre-cursor de-emphasis
- Without the de-emphasis, many cases with VEO near or below 30 mV

BROADCOM

Compare transmitters A1 and B with chip-to-chip receiver

* [c(-1), c(0), c(1)]

Test case		1	2	3	4	5	6	7	8	9
IL at 13.28 GHz, dB		8.74	8.94	4.29	8.81	4.5	9.01	9.28	10.29	11.61
COM, dB	Transmitter A1	4.83	4.76	4.72	4.99	5.04	5.17	4.02	4.18	4.22
	Transmitter B	5.94	6.23	4.46	5.37	4.80	5.92	4.80	5.27	5.19
Penalty for no pre-cursor, dB		1.11	1.47	−0.26	0.38	−0.24	0.75	0.78	1.09	0.97
VEO, mV	Transmitter A1	31.5	28.4	60.1	32.5	51.6	30.6	24.5	22.3	18.9
	Transmitter B	45.0	46.5	59.5	49.5	58.1	48.6	40.7	37.6	26.5
Loss of VEO for no pre-cursor, mV		13.5	18.1	−0.6	17.0	6.6	18.0	16.2	15.3	7.6

- Performance is significantly better for transmitter B (with pre-cursor equalization) than transmitter A1 (without pre-cursor equalization) over higher-loss hosts even though similar VEO
- Implies that not all far-end eyes are created equal

BROADCOM

Is there a better way to enforce the requirement?

- Capture the PRBS13Q waveform and calculate the linear fit pulse as defined in 120D.3.1.3
- The linear fit pulse should include the impact of the loss channel and [optimized] CTLE
- Find the amplitude of the pulse peak
- Find the magnitude of the pulse response 1 UI prior to the peak. Call this the pre-cursor value.
- Define the pre-cursor ratio as the pre-cursor value divided by the pulse peak

BROADCOM

Two transmitters evaluated with pre-cursor ratio

Module transmitter A

Parameter	Value	Parameter	Value
T_r ($\beta = 2$)	13 ps	R_d	45 Ω
A_v	0.45 V	C_d	90 fF
SNR_{TX}	35 dB	z_p	10 mm
R_{LM}	0.98	Z_c	90 Ω
A_{DD}	20 mUI	C_p	90 fF
σ_{RJ}	10 mUI		

Module transmitter B

Parameter	Value	Parameter	Value
T_r ($\beta = 2$)	13 ps	R_d	45 Ω
A_v	0.45 V	C_d	210 fF
SNR_{TX}	31 dB	z_p	12 mm
R_{LM}	0.95	Z_c	90 Ω
A_{DD}	20 mUI	C_p	110 fF
σ_{RJ}	10 mUI		

Test case	[c(−1), c(0), c(1)]	Far-end		
		COM, dB	VEO, mV	PCR, %
Transmitter A1	[0, 1, 0]	3.62	35.8	9.2
Transmitter A2	[−0.1, 0.9, 0]	8.57	61.3	2.1
Transmitter B	[−0.1, 0.9, 0]	4.56	36.3	0.85

Same VEO to within 0.5 mV but different pre-cursor ratio (PCR)

- Performance is significantly better for transmitter B (with pre-cursor equalization) than transmitter A1 (without pre-cursor equalization) over higher-loss hosts
- Pre-cursor ratio can help tell the difference between eyes with similar openings.

BROADCOM

Proposed draft text

Add the following row to Table 120E-3 after "Far-end eye height, differential (min.)":

Parameter	Reference	Value	Units
Far-end pre-cursor ratio (max.)	120E.3.2.2	2.5	%

Add new subclause after 120E.3.2.1.1:

120E.3.2.2 Far-end pre-cursor ratio

Capture the PRBS13Q waveform corresponding to the far-end eye (see 120E.3.2.1.1) and calculate the linear fit pulse using the procedure defined in 120D.3.1.3. The setting of the reference CTLE is the same used to measure eye width and height.

The peak amplitude of the linear fit pulse is p_{max}. The pre-cursor p_{pre} is the absolute value of the linear fit pulse 1 UI prior to the time of the pulse peak. The pre-cursor ratio is p_{pre} / p_{max}.

Insert the following paragraph in 120E.3.3.2.1 prior to the paragraph beginning with "The pattern is then changed…":

The far-end pre-cursor ratio is measured using the method defined in 120E.3.2.2 and it shall meet the specification in Table 120E-3. Pre-emphasis capability is likely to be required in the pattern generator to meet this requirement.

BROADCOM

Summary

- The far-end eye opening is not sufficient to ensure that the module transmitter provides fixed pre-cursor compensation
- Pre-cursor equalization can have a significant impact on link performance
- A supplemental test is proposed to ensure that pre-cursor ISI is compensated
- The proposed test is relatively simple and based on established techniques
- Draft text has been provided

Backup slides

Test case descriptions

Test case	Channel	No. of FEXT	No. of NEXT	IL at 13.28 GHz, dB
From shanbhag_3bs_14_0623				
1	Nelco 4000-13SI Host PCB + next gen 28Gb/s high density SMT IO	5	0	8.74
2	EM-888 Host PCB + next gen 28Gb/s press-fit stacked IO	7	0	8.94
From shanbhag_3bs_01_1014				
3	Next-gen 28Gb/s high density SMT IO + 4 inch host	5	0	4.29
4	Next-gen 28Gb/s high density SMT IO + 10 inch host	5	0	8.81
5	Next-gen 28Gb/s press-fit stacked IO + 4 inch host	7	0	4.5
6	Next-gen 28Gb/s press-fit stacked IO + 10 inch host	7	0	9.01
Cisco channels				
7	HCB_MCB + 3” passive	0	0	9.28
8	HCB_MCB + 4” passive	0	0	10.29
9	HCB_MCB + 5” passive	0	0	11.61

BROADCOM

Note the differences between host and module tests

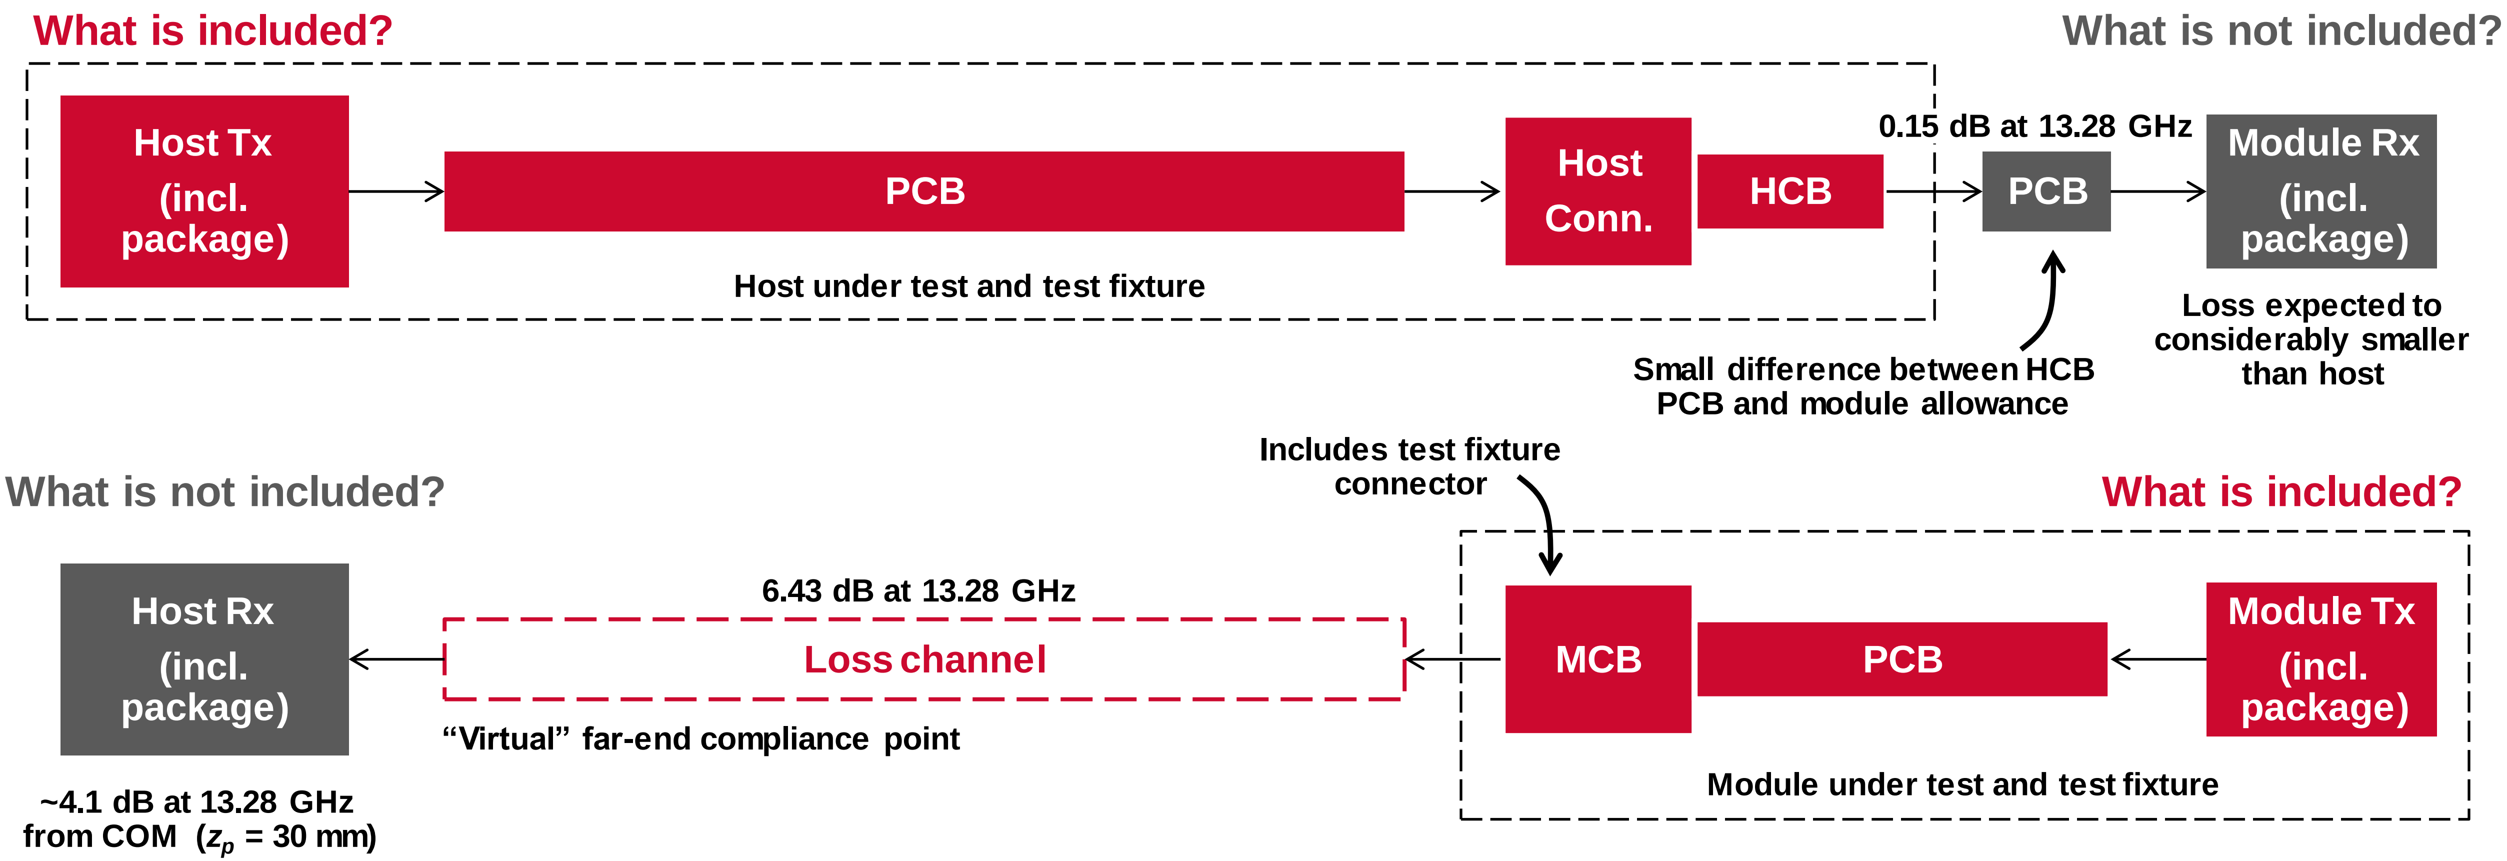

- Module test includes more "idealized" components and does not account for the possibility of a large host package
- The addition of a relatively simple test can significantly close this gap

BROADCOM