Feasibility of 200 Gb/s per lane electrical interfaces

Adam Healey

29 March 2021 (r1)

Acknowledgments

- Chi Tu, Broadcom Inc.
- Regee Petaja, Broadcom Inc.

Electrical interface building blocks to support Ethernet rates

Number of lanes		Per-lane data rate		
		50G	100G	200G
Ethernet rate	200G	4	2	1
	400G	8	4	2
	800G *	16	8	4
	1.6T *	32	16	8

* Proposed rates under consideration

Standards established for 50 Gb/s lanes and in development for 100 Gb/s lanes

Look ahead to 200 Gb/s lanes

Considering 200 Gb/s per lane

- Most straight-forward path to higher data rates includes some combination of …
  - Increasing the number of bits per symbol, e.g., PAM-M
  - Increasing the number symbols transmitted per unit time, i.e. signaling rate
- What performance can be achieved with these techniques?
- Consider chip-to-module (C2M) interfaces
- Chip-to-chip (C2C) interface performance can be inferred from these results
- These interfaces enable future-generation optical modules and active cables

Review of existing error correction architectures

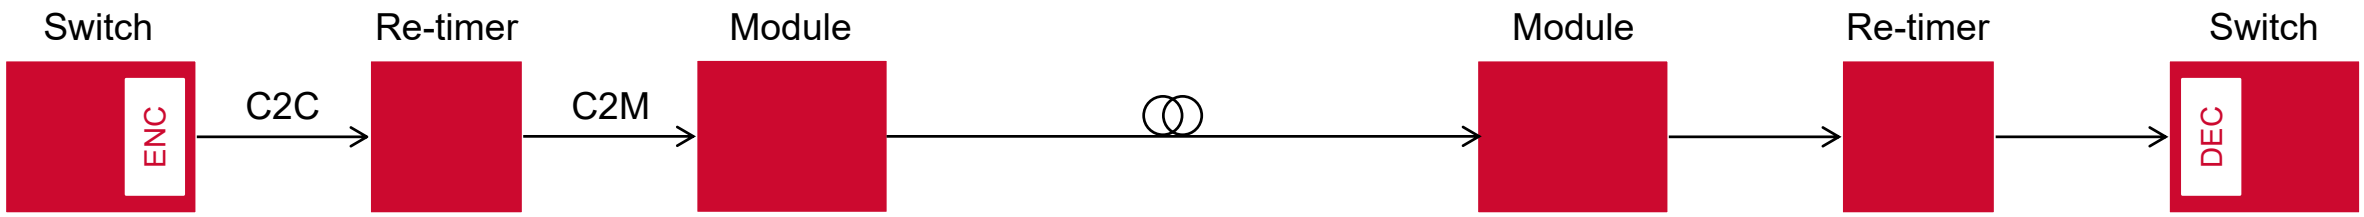

- End-to-end error correction model dedicates the lion's share of the coding gain to the optical link
- Target BER for an electrical interface is reduced so that errors seen by the decoder are dominated by the worst-case optical link

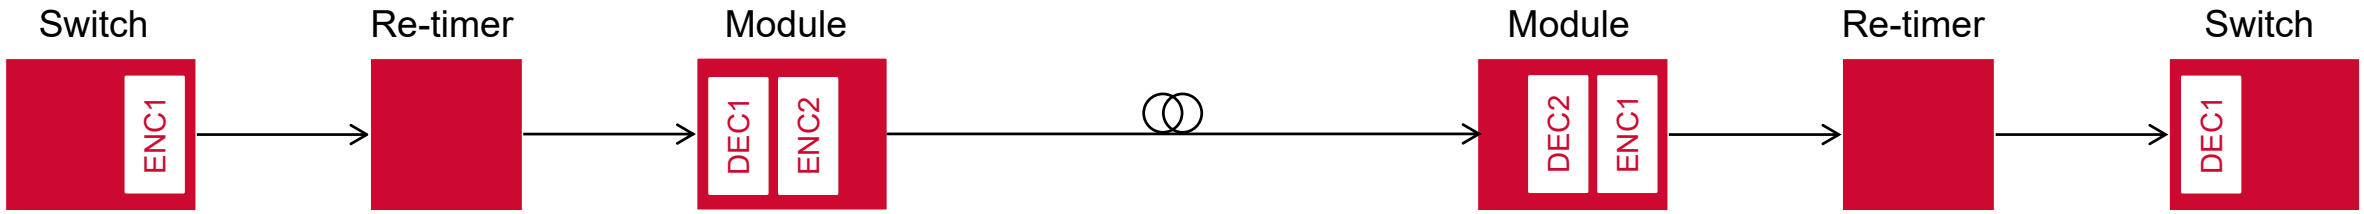

- Optical link has a stronger, dedicated error correcting code (e.g., 100GBASE-ZR, 400GBASE-ZR)
- Electrical link to the module could use weaker code to reduce decoder complexity in module host interface

Architecture, code selection, and link BER target dictate electrical interface BER targets

BROADCOM

Pulse amplitude modulation (PAM-M) trade-offs

Data rate, Gb/s	106.25	212.5					
Number of levels, M	**4**	**4**	**5**	**6**	**7**	**8**	**16**
Bits per symbol [1]	2	2	2.25	2.5	2.75	3	4
Signaling rate, Gbaud	**53.13**	**106.25**	**94.44**	**85**	**77.27**	**70.83**	**53.13**
Unit interval, ps	18.82	9.41	10.59	11.76	12.94	14.12	18.82
Fundamental frequency, GHz	26.56	53.13	47.22	42.5	38.64	35.42	26.56
Required SNR at slicer, dB [2]	19.46	19.46	21.46	23.06	24.41	25.57	31.53
SNR penalty, dB	**0**	**0**	**2**	**3.61**	**4.95**	**6.11**	**12.08**
Jitter for 1 dB penalty, mUI RMS [3]	**21.76**	**21.76**	**16.99**	**13.69**	**11.15**	**9.16**	**3.34**

[1] Includes mapping overhead assuming a 4D constellation for PAM-5 and PAM-7 and a 2D constellation for PAM-6.

[2] For BER = 1e−5 (required SNR is about 1 dB higher for BER = 1e−6 and about 1.3 dB lower for BER = 1e−4).

[3] For BER = 1e−5 and 20% excess bandwidth.

Expand the constellation size to reduce bandwidth only when the SNR can support it

BROADCOM

Channel model overview: Host-to-module direction

Frequency, GHz	IL, dB/inch
28	0.87
42.5	1.13
56	1.33

Connector model courtesy of Amphenol

NOTE — Channel models include dominant near-end aggressor and dominant far-end aggressor.

Channel insertion loss at the fundamental frequency

PAM-M: 4 5 6 7 8

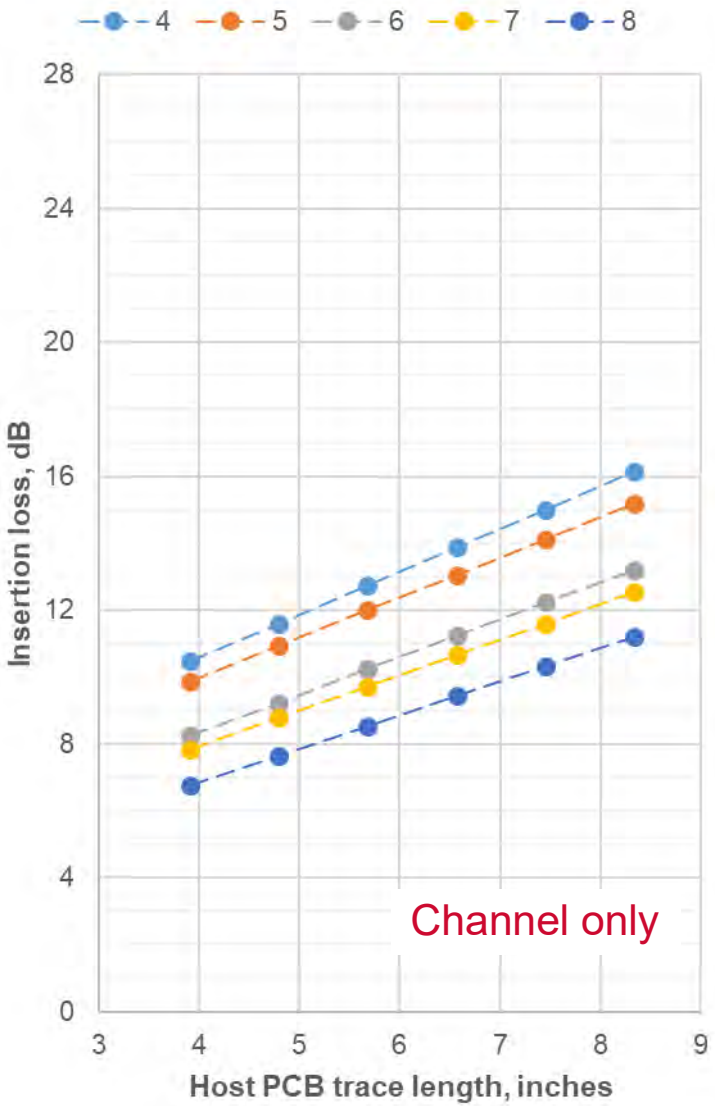

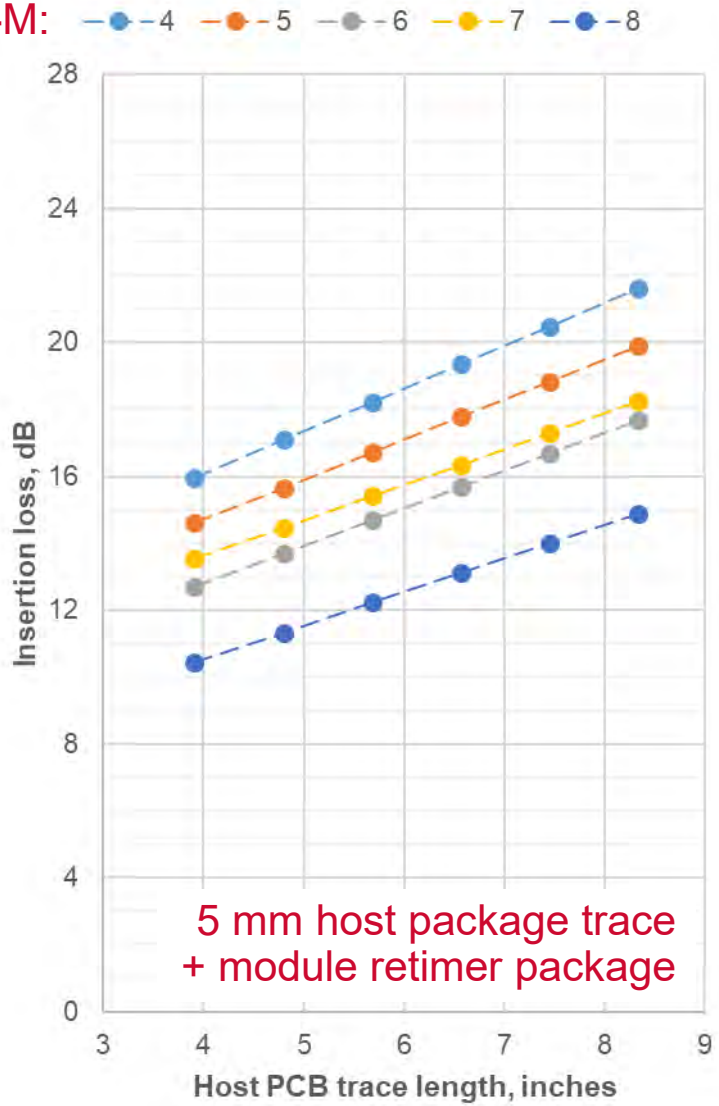

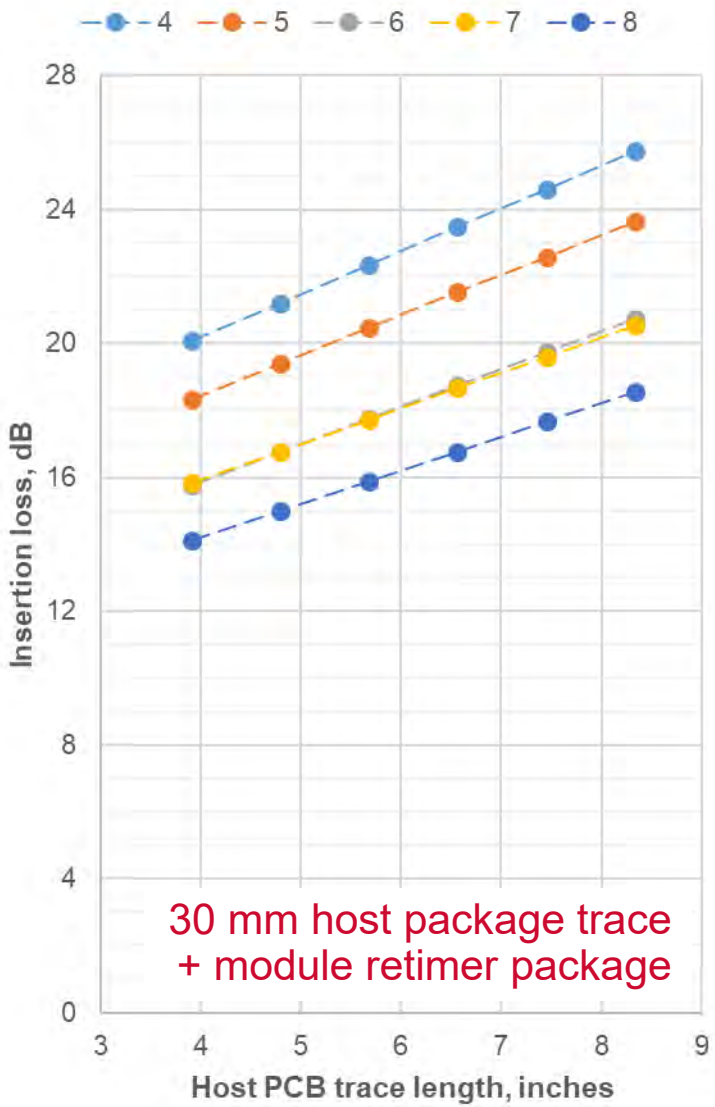

BROADCOM

Link model for Salz SNR calculations

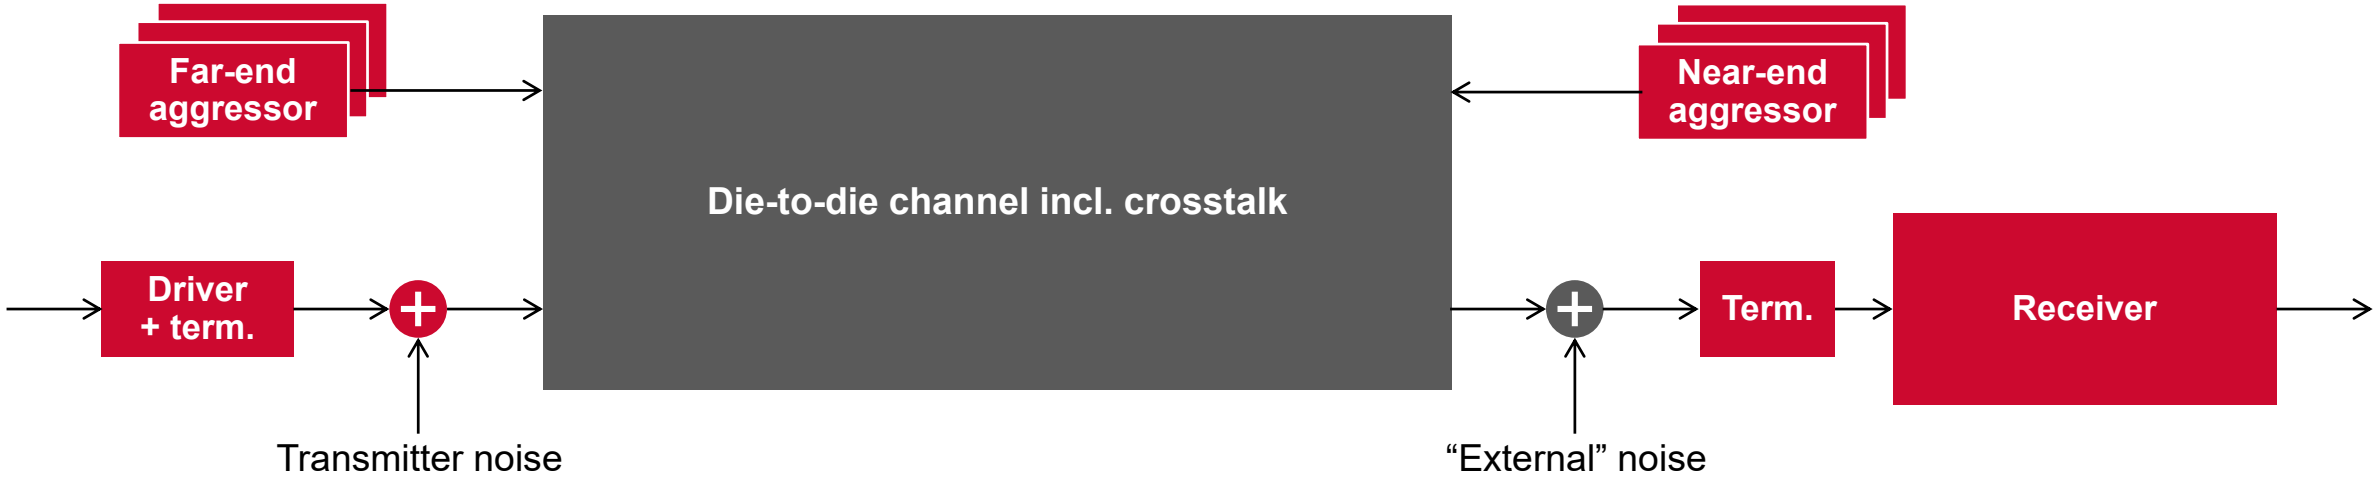

Modeled impairments

Parameter	Value
Driver differential output amplitude (peak), V	0.4
Driver rise/fall times (20-80%), ps	5.5
Transmitter uncorrelated jitter, fs RMS	140
Transmitter signal-to-noise ratio, dB	34
Far-end aggressor output amplitude (peak), V	0.4
Near-end aggressor output amplitude (peak), V	0.4
External noise spectral density (2-sided), dBm/Hz	−164.8
Minimum implementation margin, dB	6

Decision-point SNR for ideal DFE (a.k.a. Salz SNR)

$$SNR_{MMSE-DFE} = \frac{1}{2\pi}\int_{-\pi}^{\pi} 10\log_{10}(F(\theta)+1)d\theta$$

…where $F(\theta)$ is the folded SNR at frequency θ.

$$F(\theta) = \frac{1}{T}\sum_{m}\left|S\left(\frac{\theta+2\pi m}{2\pi T}\right)\right|^2 \quad -\pi < \theta \leq \pi$$

…and $S(f)$ is the frequency-dependent ratio of signal power to noise power.

BROADCOM

Consider equalizer complexity

- Salz SNR does not account for practical constraints on receiver complexity
- Lower complexity is desirable for chip-to-module interfaces
- To account for the impact of limited complexity, reflections beyond a specified delay or “post-cursor span” are treated as noise
- Salz SNR will be reduced for reflections beyond the post-cursor span
- Result converges to the “traditional” Salz SNR value as the post-cursor span increases

Channel-only results: 8.3 inch host trace

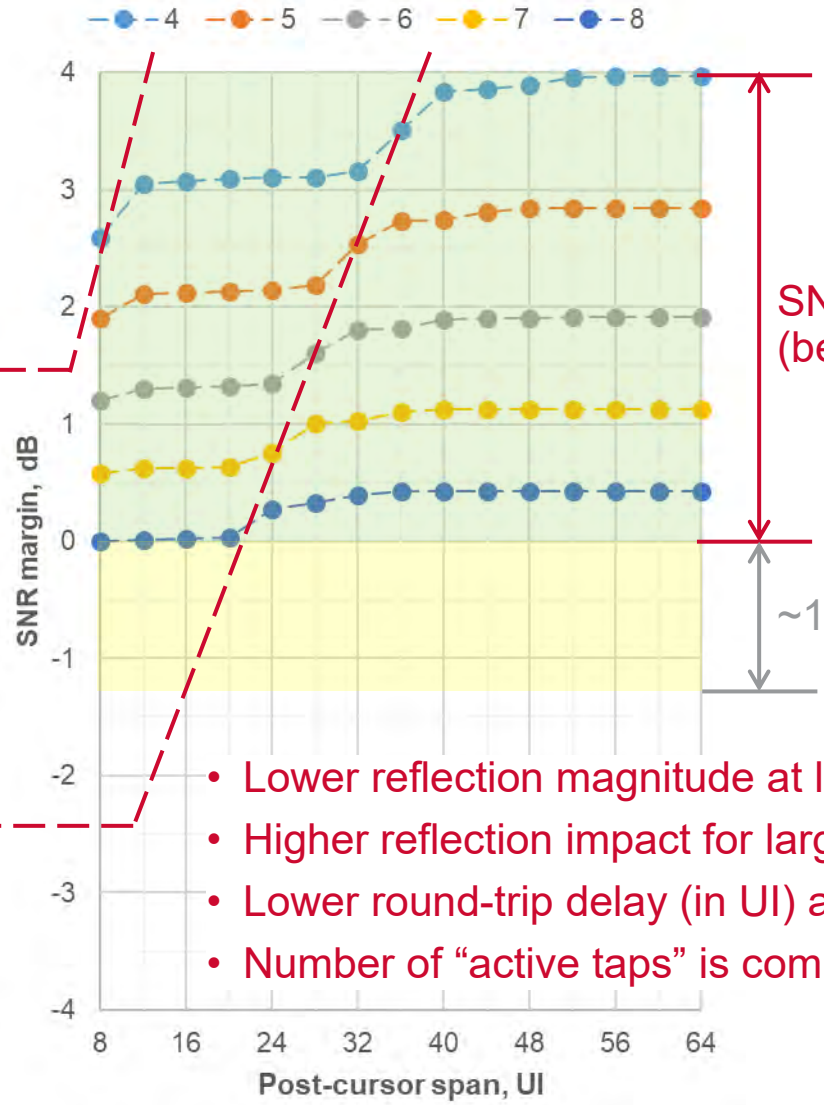

Reflection internal to connector

Reflection between connector and module BGA pad

SNR margin relative BER = 1e−5 (beyond 6 dB implementation margin)

~1.3 dB additional margin for BER = 1e−4

- Lower reflection magnitude at lower bandwidth
- Higher reflection impact for larger constellation
- Lower round-trip delay (in UI) at lower signaling rate
- Number of “active taps” is comparable (e.g., for “floating tap” equalizers)

BROADCOM

Add device packages and terminations: 8.3 inch host trace

Enhanced by reflection internal to host package

Interaction between module package and connector

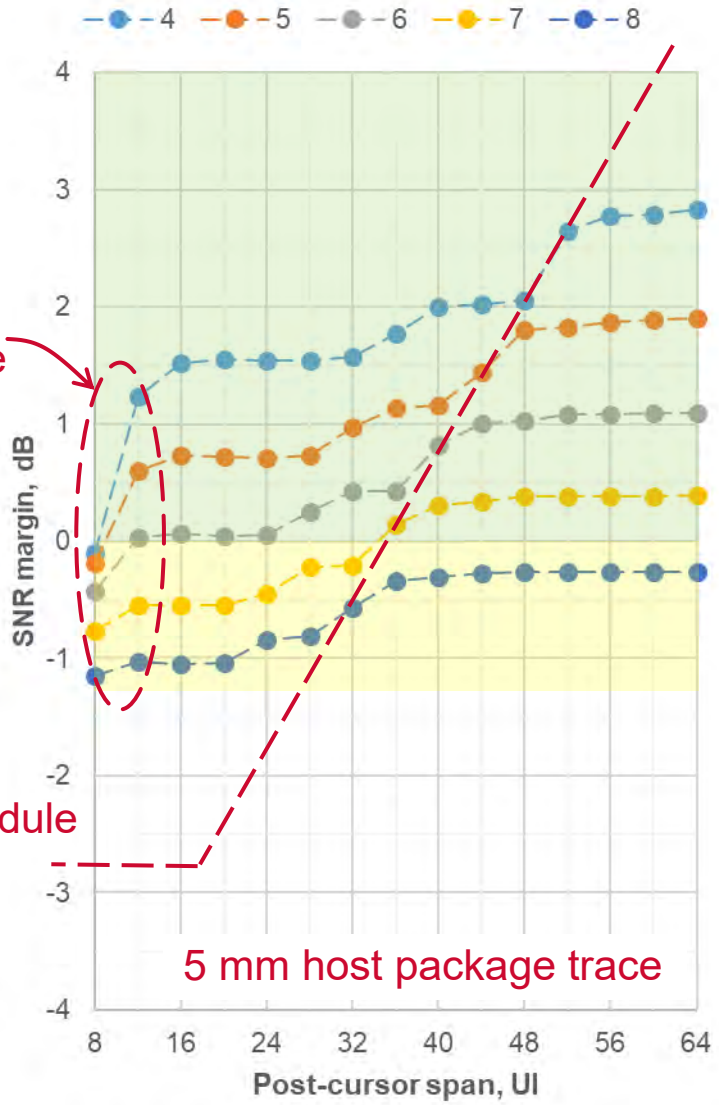

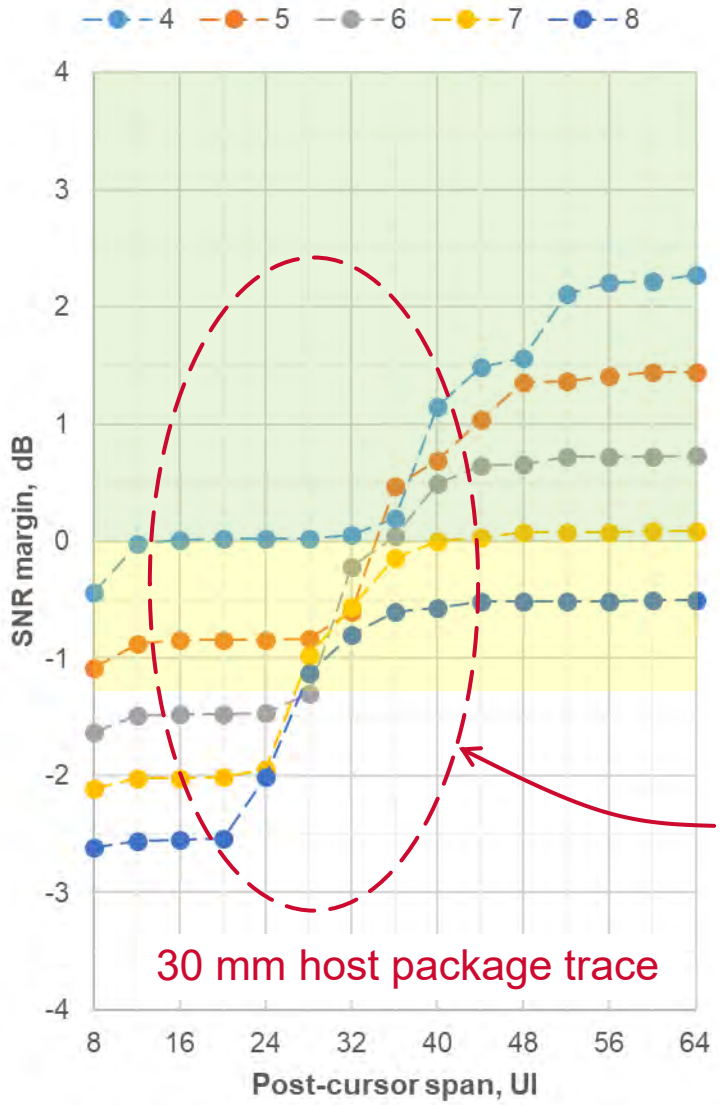

Enhanced by reflection internal to host package

BROADCOM

Performance vs. host trace length (12 UI post-cursor span)

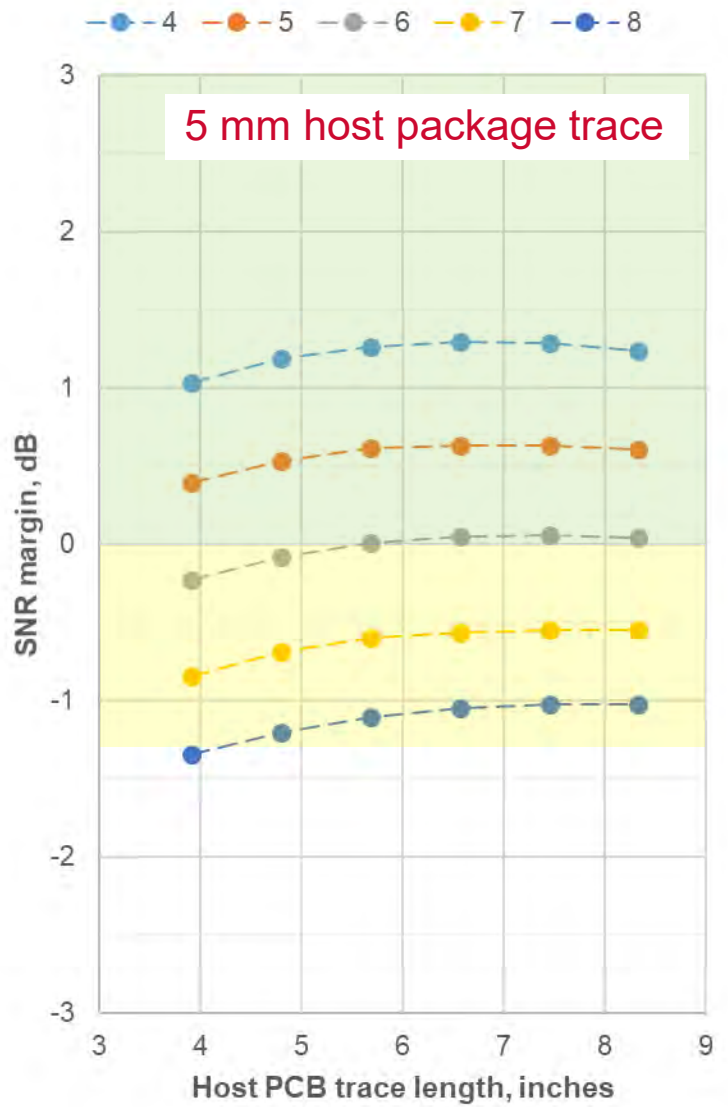

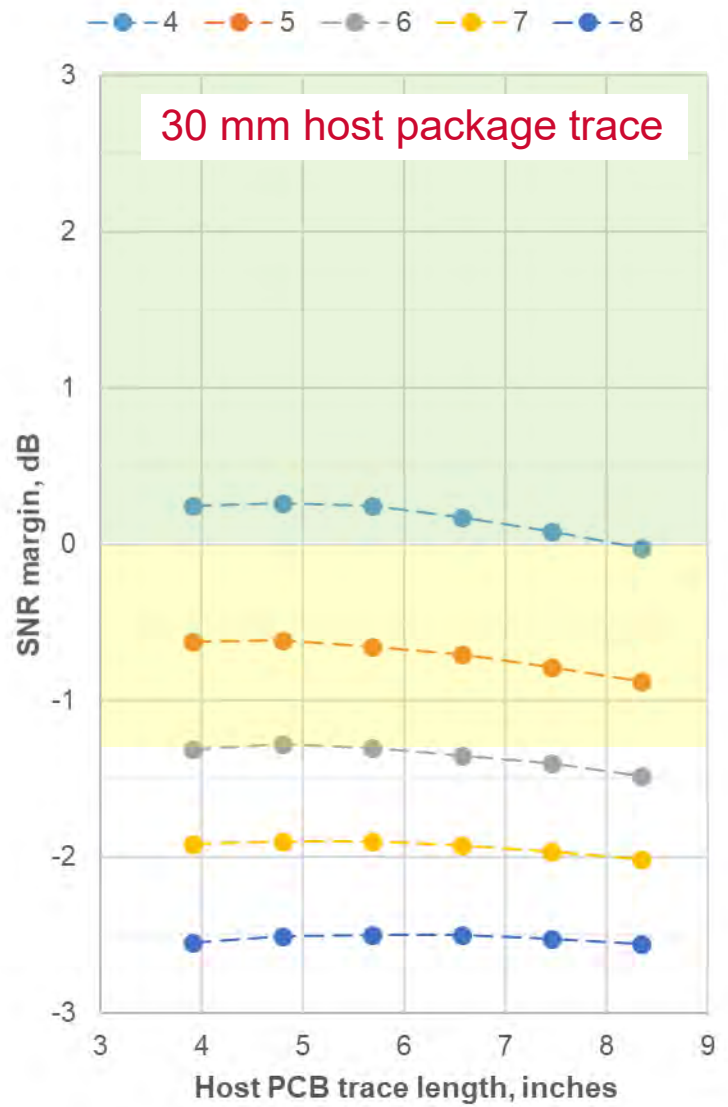

BROADCOM

Crosstalk sensitivity (12 UI post-cursor span)

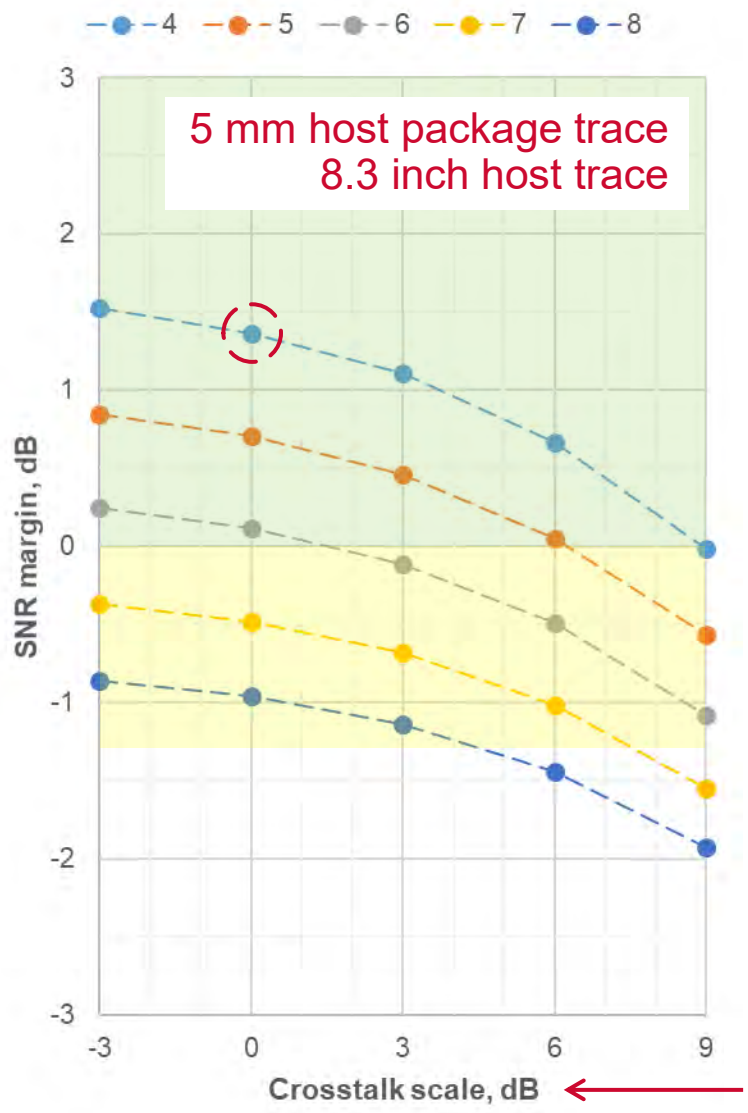

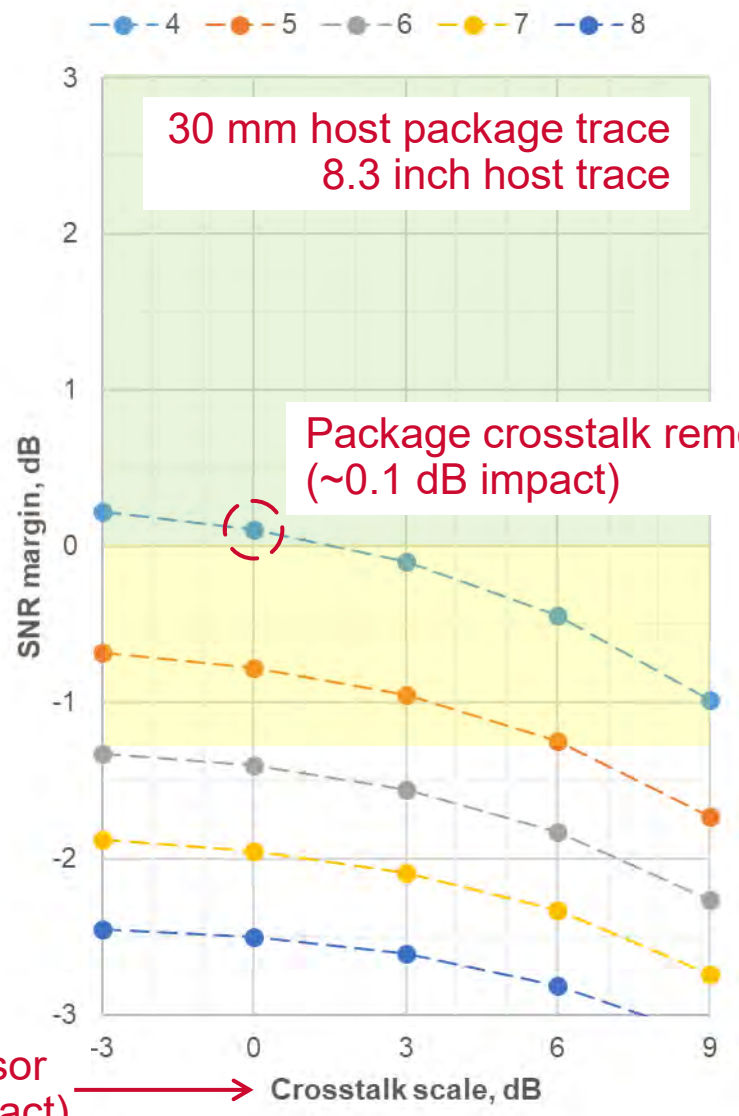

Scale FEXT aggressor
(NEXT has lower impact)

BROADCOM

Will more customization per application become necessary?

Label	Application examples	50G/lane	100G/lane	200G/lane
—	Interface to co-packaged optics	PAM-4	PAM-4	???
C2M	Interface to modular optics, active cables	PAM-4	PAM-4	???
C2C	Re-timers, reach extenders	PAM-4	PAM-4	???
KR, CR	Backplane, mid-plane, passive copper cables	PAM-4	PAM-4	???
—	Direct-detect optics	PAM-4	PAM-4	???

- Common modulation with application-specific performance tuning facilitates multi-purpose designs, design re-use, and a simpler interface to direct-detect optics
- Can this continue (and does it need to continue) for 200G/lane?
- Should we strive for commonality between direct-detect optics, C2M, and C2C?
- Should we maximize copper reach (KR and CR) at the expense of commonality?

Summary

- Analysis suggests that 200G/lane chip-to-module (C2M) interfaces are feasible
- Implies that 200G/lane chip-to-chip (C2C) interfaces are also feasible (given a similar loss budget)
- Important interfaces for future generations of optical modules and active cables
- Potential trade-offs and areas for improvement identified
- Opportunities to leverage existing error correction architecture(s)
- More sophisticated receivers may be required
- Choice of modulation for direct-detect optical links will have an influence

Recommendation

- If objective(s) for 200G optical lanes are added, then objective(s) that address 200G/lane electrical interfaces should also be added
- Link BER target, error correction architecture, encoding overhead, and optical modulation choices impact (or could be influenced by) the electrical interfaces
- A holistic view is needed to optimize the end-to-end link