SNIA DEVELOPER CONFERENCE

SDC 24

BY Developers FOR Developers

September 16-18, 2024
Santa Clara, CA

PCIe® Cabling Solutions for Next Generation AI, HPC, and Storage Applications

Mohiuddin Mazumder
Co-chair, PCI-SIG® Electrical WG and Optical Cable WG
Senior Principal Engineer, Intel Corporation

Agenda

- Need for Cable Solutions
- PCI-SIG®'s Cabling – Journey to CopprLink™ and beyond
- A Brief Overview of PCIe® 5.0/6.0 CopprLink Internal and External Cable Specifications
- Example Methodologies to Quantify Electrical Performance of Cable Solutions
- Sideband Specifications of CopprLink Internal and External Cables
- Update on PCIe Optical Cable Solution for PCIe 6.0 specification and beyond

PCIe® Technology Continues to Deliver Ahead of Industry Requirements

- New specification about every 3 years
- Bandwidth doubles between versions
- Low-latency and high-reliability
- Power efficient
- Maintains backwards compatibility with previous generations
- Data integrity and security

Revision	Max Data Rate	Encoding	Signaling	Nyquist Frequency
PCIe 1.0 (2003)	2.5 GT/s	8b/10b	NRZ	1.25 GHz
PCIe 2.0 (2007)	5.0 GT/s	8b/10b	NRZ	2.5 GHz
PCIe 3.0 (2010)	8.0 GT/s	128b/130b	NRZ	4 GHz
PCIe 4.0 (2017)	16.0 GT/s	128b/130b	NRZ	8 GHz
PCIe 5.0 (2019)	32.0 GT/s	128b/130b	NRZ	16 GHz
PCIe 6.0 (2022)	64.0 GT/s	1b/1b (Flit Mode*)	PAM4	16 GHz
PCIe 7.0 (2025)	128.0 GT/s	1b/1b (Flit Mode*)	PAM4	32 GHz

(*Flit Mode also enabled in other Data Rate with their respective encoding)

PCIe® Speeds/Feeds - Pick Your Bandwidth

- **Flexible to meet needs from handheld/client to server/HPC**
- **~Max Total Bandwidth = Max RX bandwidth + Max TX bandwidth**
- **35 Permutations yielding 11 unique bandwidth profiles**
- **Encoding overhead and header efficiency not included**

Specifications	Lanes				
	x1	x2	x4	x8	x16
2.5 GT/s (PCIe 1.x +)	500 MB/S	1 GB/S	2 GB/S	4 GB/S	8 GB/S
5.0 GT/s (PCIe 2.x +)	1 GB/S	2 GB/S	4 GB/S	8 GB/S	16 GB/S
8.0 GT/s (PCIe 3.x +)	2 GB/S	4 GB/S	8 GB/S	16 GB/S	32 GB/S
16.0 GT/s (PCIe 4.x +)	4 GB/S	8 GB/S	16 GB/S	32 GB/S	64 GB/S
32.0 GT/s (PCIe 5.x +)	8 GB/S	16 GB/S	32 GB/S	64 GB/S	128 GB/S
64.0 GT/s (PCIe 6.x +)	16 GB/S	32 GB/S	64 GB/S	128 GB/S	256 GB/S
128.0 GT/s (PCIe 7.x +)	32 GB/S	64 GB/S	128 GB/S	256 GB/S	512 GB/S

+ = data rate supported by this and subsequent spec revisions.

PCIe® Architecture: One Interconnect – Infinite Opportunities

HPC / Cloud

Data Center / Enterprise Servers

Artificial Intelligence / Machine Learning

Automotive

Internet of Things

Military / Aerospace

Storage

Challenges Over Copper-Based PCB Routing

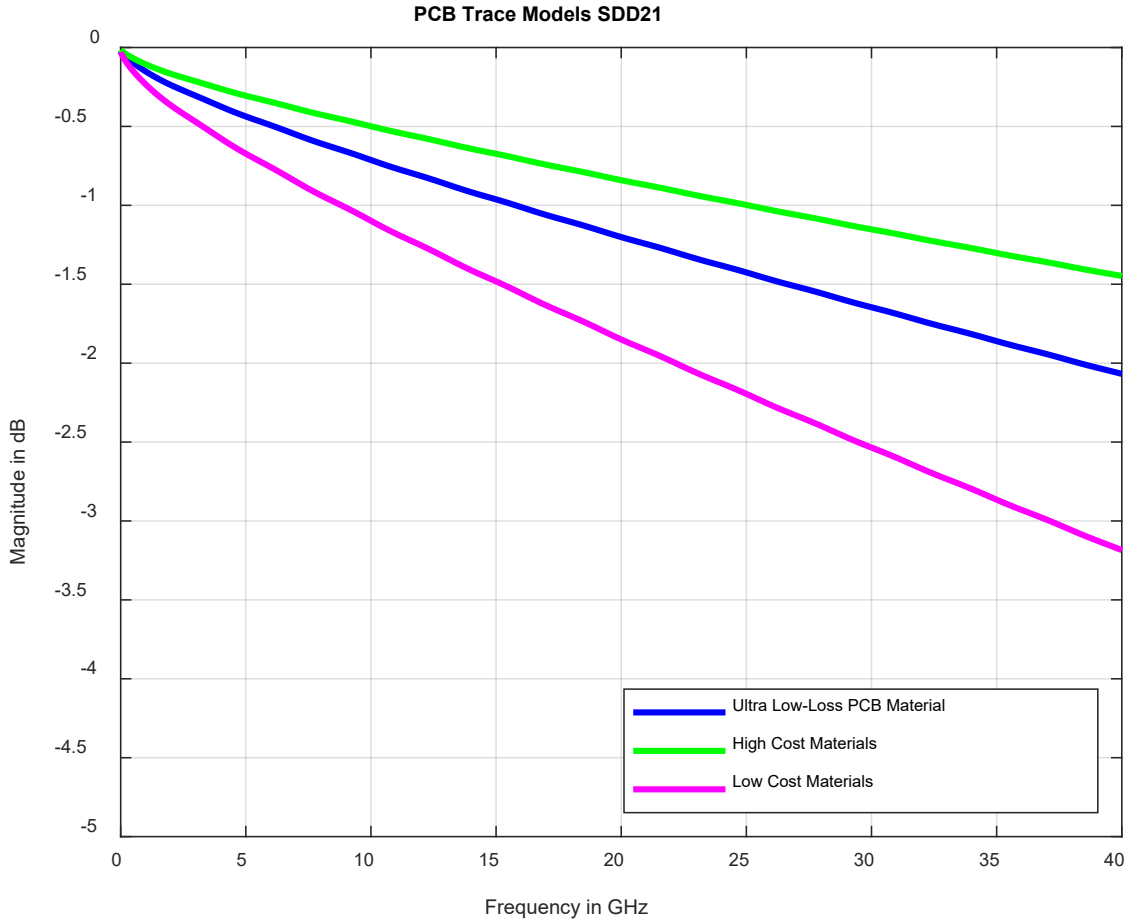

PCB Loss at 16 GHz and 32 GHz limits channel reach and impacts platform architecture choices

PCI-SIG® Cabling – Journey to CopprLink™

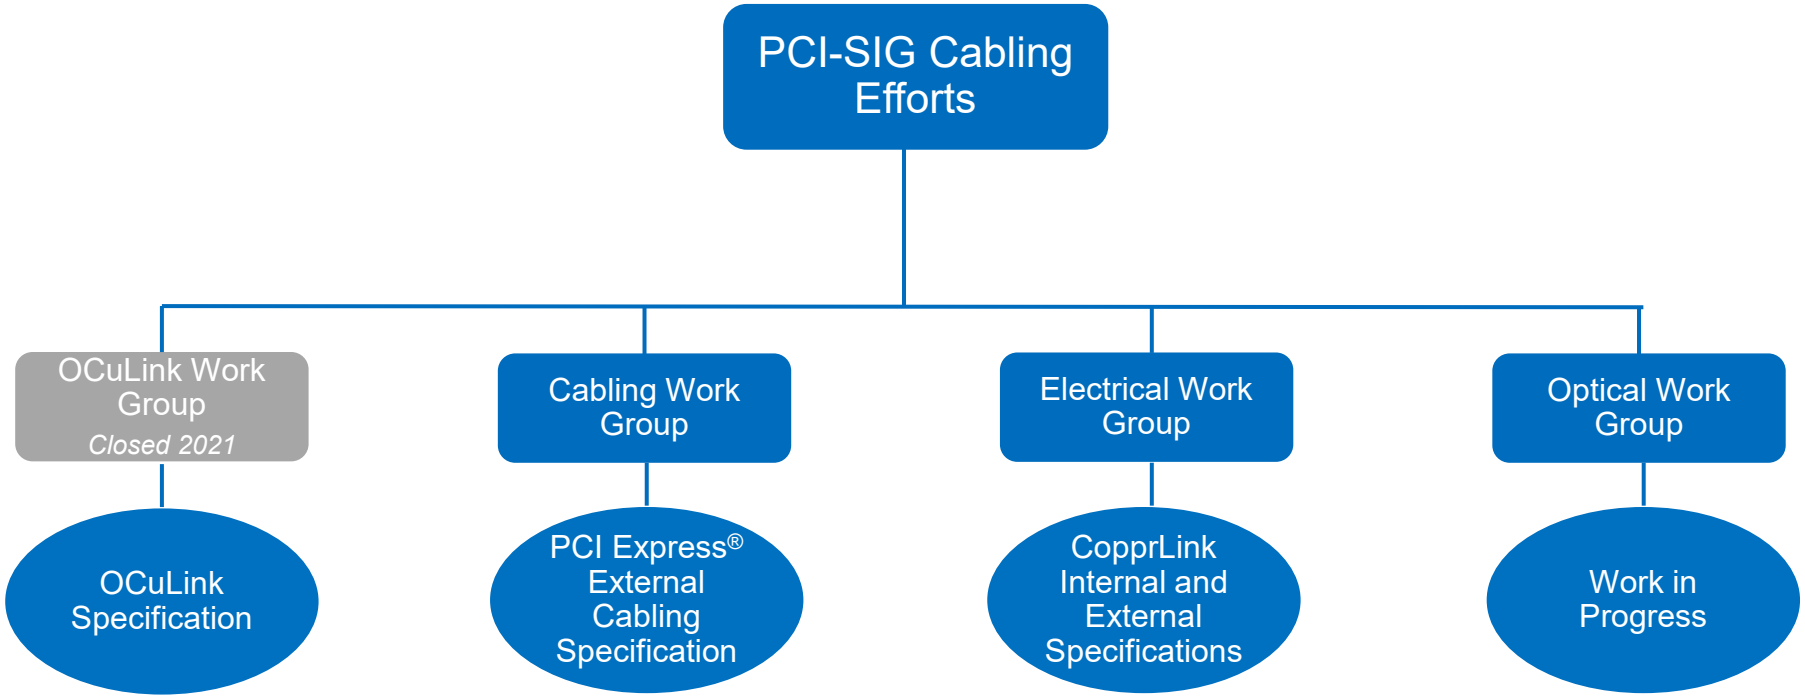

PCI-SIG® Cabling Timeline

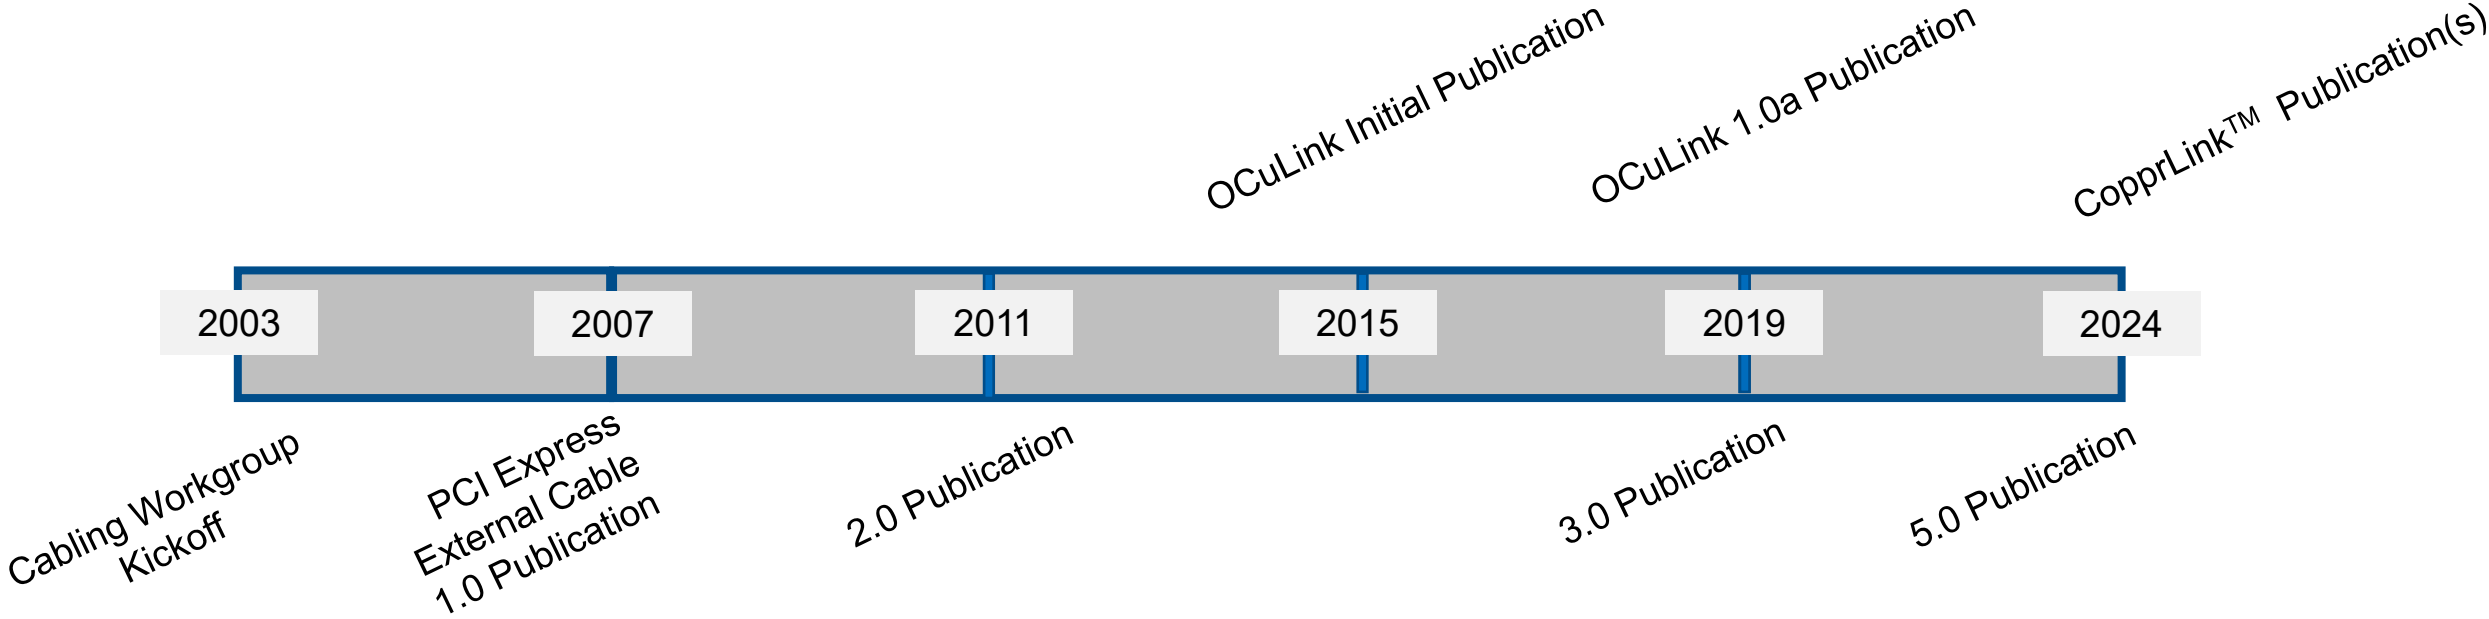

- Cabling efforts have been part of PCI-SIG for 20+ years
- In the last five years, cabling interest reinvigorated with the CopprLink™ effort(s)
- Cabling solutions must be aligned with architecture for platform and form factor

PCI Express® External Cabling Specification

- About the specification:
  - Supports PCIe® 1.0 thru 5.0 technology signaling up to at 32.0 GT/s
  - Includes the SNIA SFF-8614 connector form factor (for Revision 3.0 and beyond)
  - Maximum of 2m reach within a single system
  - Example implementations include motherboard-to-add-in-card, add-in-card-to-add-in-card, and Host Bus Adapter (HBA)
  - Target application is traditional storage, an alternative to SAS/SATA deployments

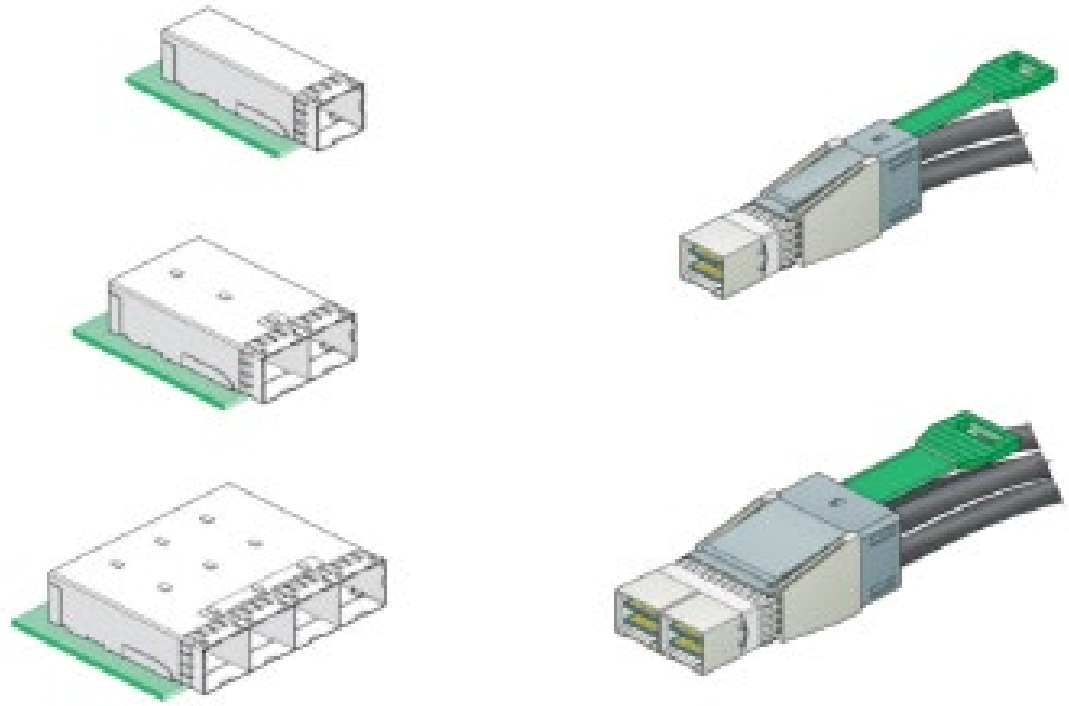

PCI-SIG® Cabling Initiatives

Name	Date	Specification	Targeted Speeds	Internal or External	Connector Form Factor	Reach	Target Applications
OCuLink	Oct. 2015	PCIe® 3.0	2.5, 5.0, 8.0 GT/s	Both	Reference the OCuLink specification for more information	Up to 2m (8.0 GT/s)	Storage, data center
CopprLink™ Internal Cable	Mar. 2024	PCIe 5.0, PCIe 6.0	32.0, 64.0 GT/s	Internal	SNIA SFF-TA-1016	1m within a single system	Storage, data center
CopprLink™ External Cable	Apr. 2024	PCIe 5.0, PCIe 6.0	32.0, 64.0 GT/s	External	SNIA SFF-TA-1032	Up to 2m reach in within rack and rack-to-rack connections	AI and ML
PCI Express® External Cabling	2H 2024	PCIe 5.0	8.0, 16.0, 32.0 GT/s	External	SNIA SFF-8614	3m (8.0 GT/s), 2.5m (16.0GT/s), 2m (32.0 GT/s)	Traditional storage
Optical	WG est. Aug. 2023	WIP	WIP	WIP	WIP	WIP	Cloud and Quantum Computing, Hyperscale Data Centers, HPC

Agenda

- Need for Cable Solutions
- PCI-SIG®'s Cabling – Journey to CopprLink™ and beyond
- A Brief Overview of PCIe® 5.0/6.0 CopprLink Internal and External Cable Specifications
- Example Methodologies to Quantify Electrical Performance of Cable Solutions
- Sideband Specifications of CopprLink Internal and External Cables
- Update on PCIe Optical Cable Solution for PCIe 6.0 specification and beyond

PCI Express® 5.0/6.0 CopprLink Cable Specifications

Motivation for New Cable Specifications

- High-bandwidth applications require longer channel reach and topological flexibility
- At 32.0 and 64.0 GT/s, PCB loss reduces channel reach and limits platform flexibility
- Low loss is not enough, and minimization of reflection, crosstalk, and skew is critical for cable solutions
- Cable is another option to give flexibility to system designers and boost innovation

Internal Cable Usages

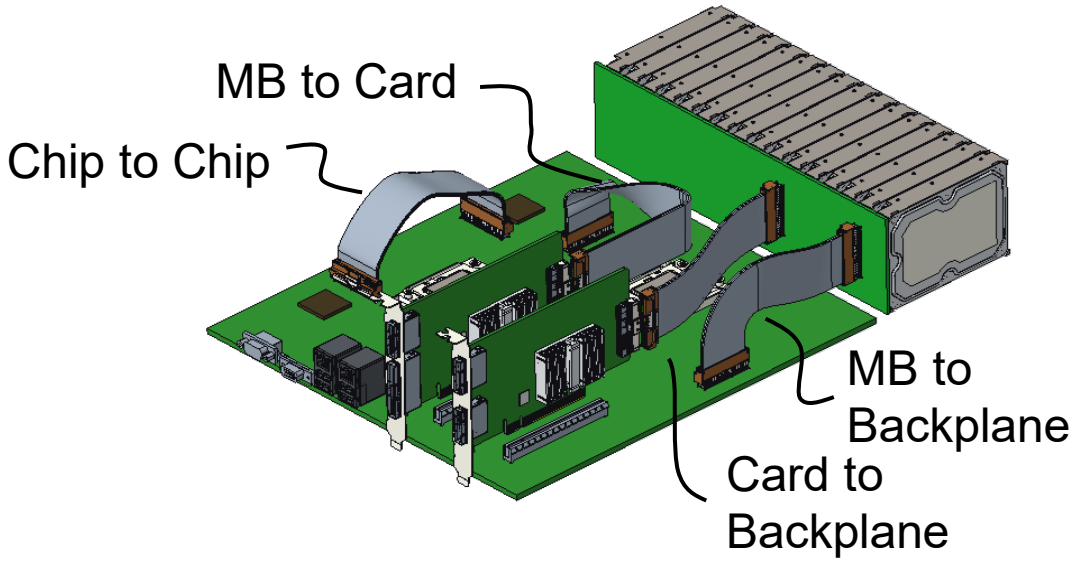

External Cable Usages

- Connecting boards within a rack
- Rack-to-rack

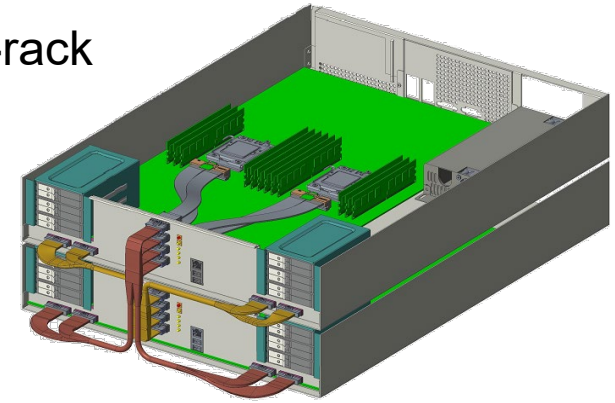

Board to Board

PCI Express® 5.0/6.0 CopprLink Cable Specifications

Feature	Internal	External
Delivers 32.0 GT/s and 64.0 GT/s data rates	✓	✓
Form Factor	SFF-TA-1016	SFF-TA-1032
Maximum cable length	1 meter	2 meters
Pulse Amplitude Modulation with 4 levels (PAM4) and NRZ signaling	✓	✓
Supports x4, x8, x16 configurations, cable bifurcation	✓	✓
Side band specifications	✓	✓
Targeted for data center applications, servers, storage, networking and accelerators	✓	✓

Objectives of PCIe® 5.0/6.0 CopprLink Internal and External Cable Specifications

- Electrical specifications for 32 GT/s and 64 GT/s for mated cable assembly and mated connector based on SFF-TA-1016 and SFF-TA-1032 Specifications
- Specifications of sideband functions
- Pinout for the high-speed and sideband pins
- Guidelines for 32 GT/s and 64 GT/s electrical spec compliance testing

The SFF-TA-1016 and SFF-TA-1032 Specification documents and references therein must be used for Mechanical, Material, Environmental, Reliability, and other Electrical specifications.

SDC 24

Agenda

- Need for Cable Solutions
- PCI-SIG®'s Cabling – Journey to CopprLink™ and beyond
- A Brief Overview of PCIe® 5.0/6.0 CopprLink Internal and External Cable Specifications
- Example Methodologies to Quantify Electrical Performance of Cable Solutions
- Sideband Specifications of CopprLink Internal and External Cables
- Update on PCIe Optical Cable Solution for PCIe 6.0 specification and beyond

Example Topologies for Signal Integrity Study

PCB Loss at 16 GHz: ~ 1 dB/in
Cable Loss at 16 GHz: ~0.1875 dB/in

Cable loss can be 5-10x lower than PCB loss

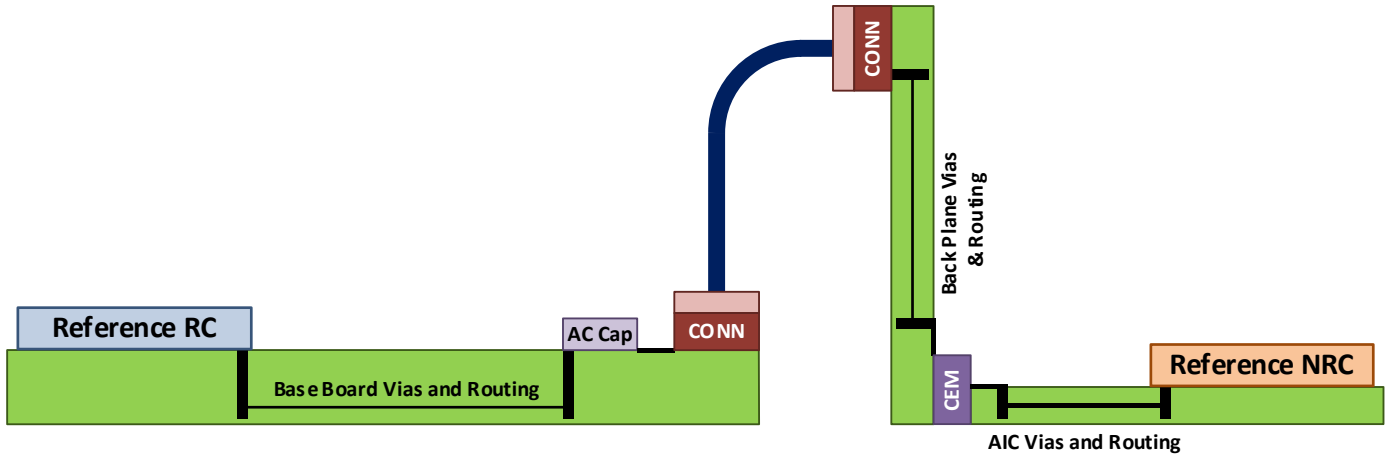

A Typical Internal Cable Topology
(e.g., connecting a Riser/Backplane to the system board)

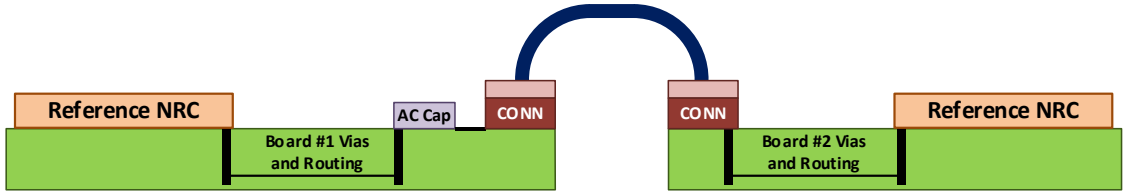

A Typical External Cable Topology
(e.g., connecting two boards within a rack)

Cable mitigates PCB loss limitations but increases reflection and xtalk impact

3-Conn Cable Topology – A Highly Reflective Channel Example

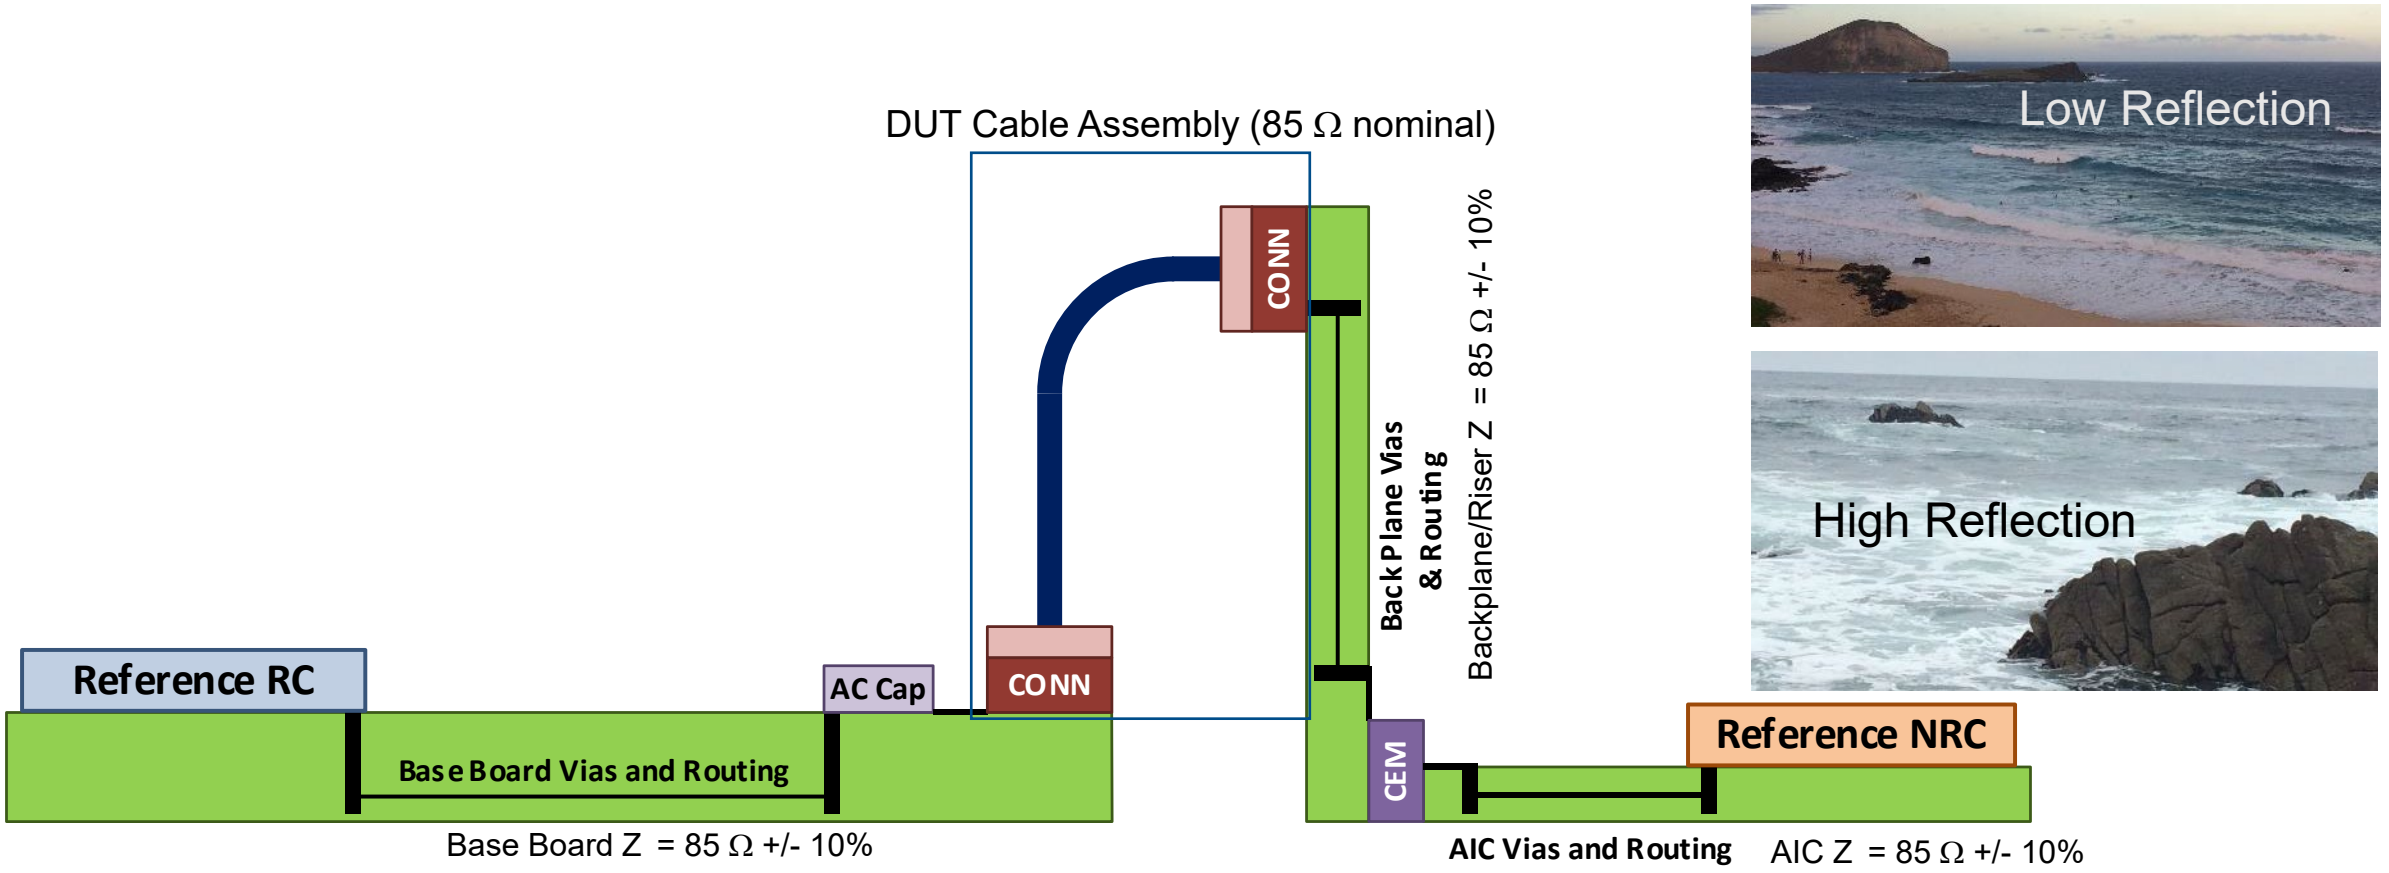

Reflection impact increases with the number of interfaces, the impedance mismatch between interfaces, and the adjacency of the interfaces

A Few Cases for Reflection Impact Study

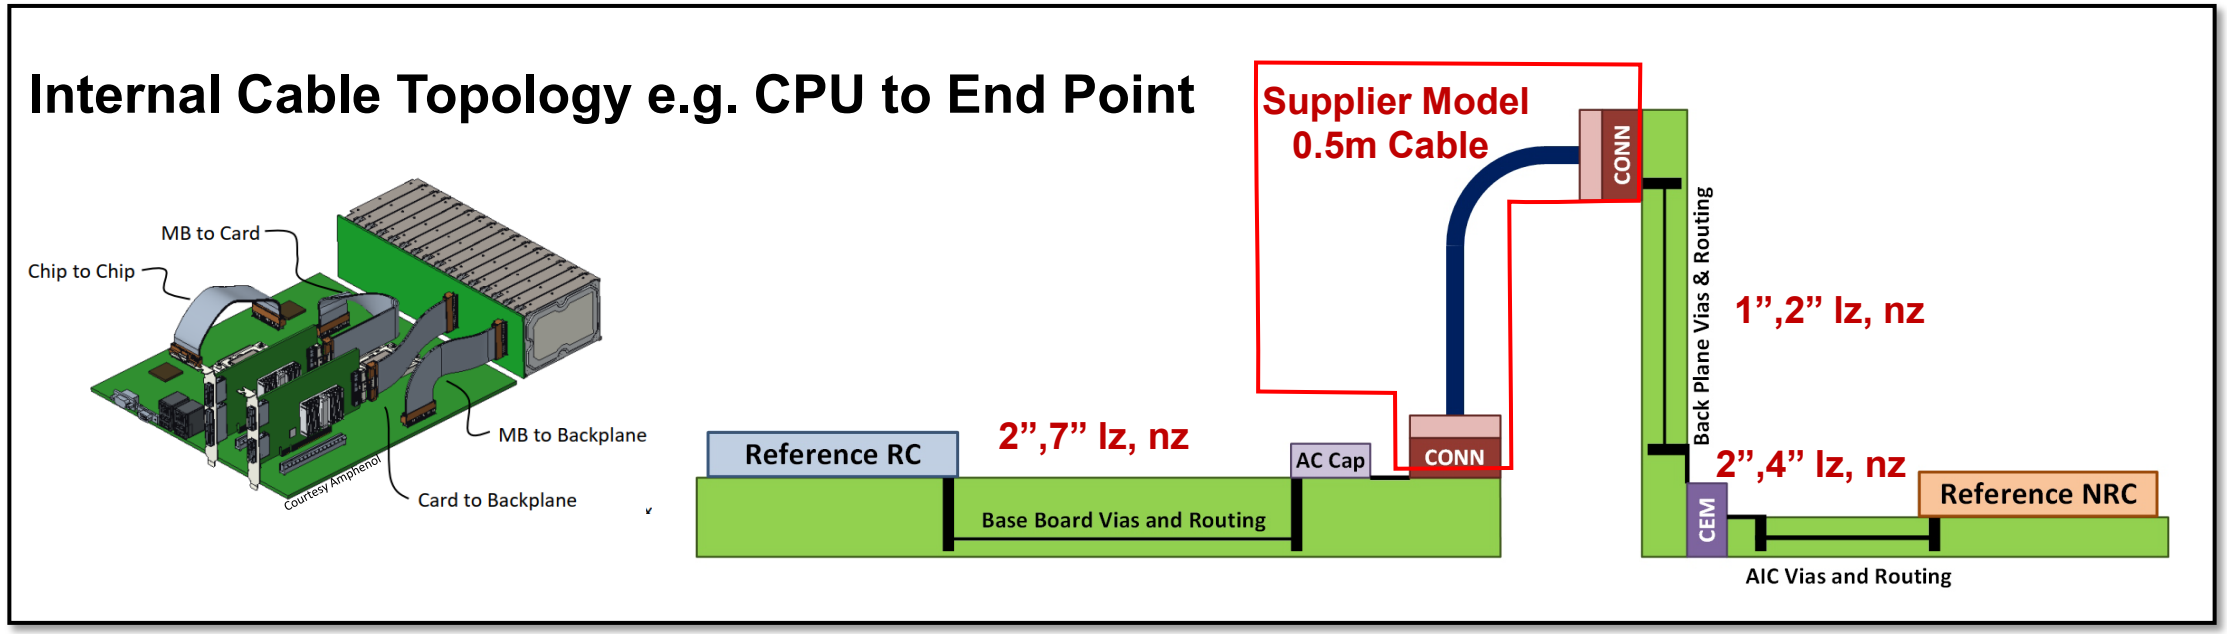

Cable length = 0.5m
Case 1: short channel and Z of bb1, hsbp/riser, bb2 = all low, all nom
BB1 = 2", hsbp/riser = 1", aic3 = 2"

Case 2: mid-channel and Z of bb1, hsbp, bb2 = all low, all nom
BB1 = 7", hsbp/riser = 1", aic3 = 2"

Case 3: long channel and Z of bb1, hsbp, bb2 = all low, all nom
BB1 = 7", hsbp/riser = 2", aic3 = 4"

Review of Integrated Return Loss (iRL) Definition

$$iRL = dB\left(\sqrt{\frac{1}{N}\sum_{i=1}^{N} W(f_i)\, RL_{avg}^2(f_i)}\right)$$

1. $RL_{11}(f)$, $RL_{22}(f)$= cable assembly/connector return loss
2. $RL_{avg}(f_i) = (|RL_{11}(f_i)| + |RL_{22}(f_i)|)/2$
3. Weighting function $W(f_i) = sinc^2(f_i/f_b)\,\frac{1}{1+(\frac{f_i}{f_t})^4}\,\frac{1}{1+(\frac{f_i}{f_r})^8}$
4. f_b, Symbol Rate = 32 GBaud
5. T_r = Rise Time = 25 ps
6. f_t = 9.46 GHz, f_r =24 GHz (where f_t = 0.2365/ T_r; T_r = 25 ps, and f_r = 1.5*Nyquist)
7. N = Number of samples, length of frequency array, usually in 10 MHz step

Differential Return Loss: Cable Assembly

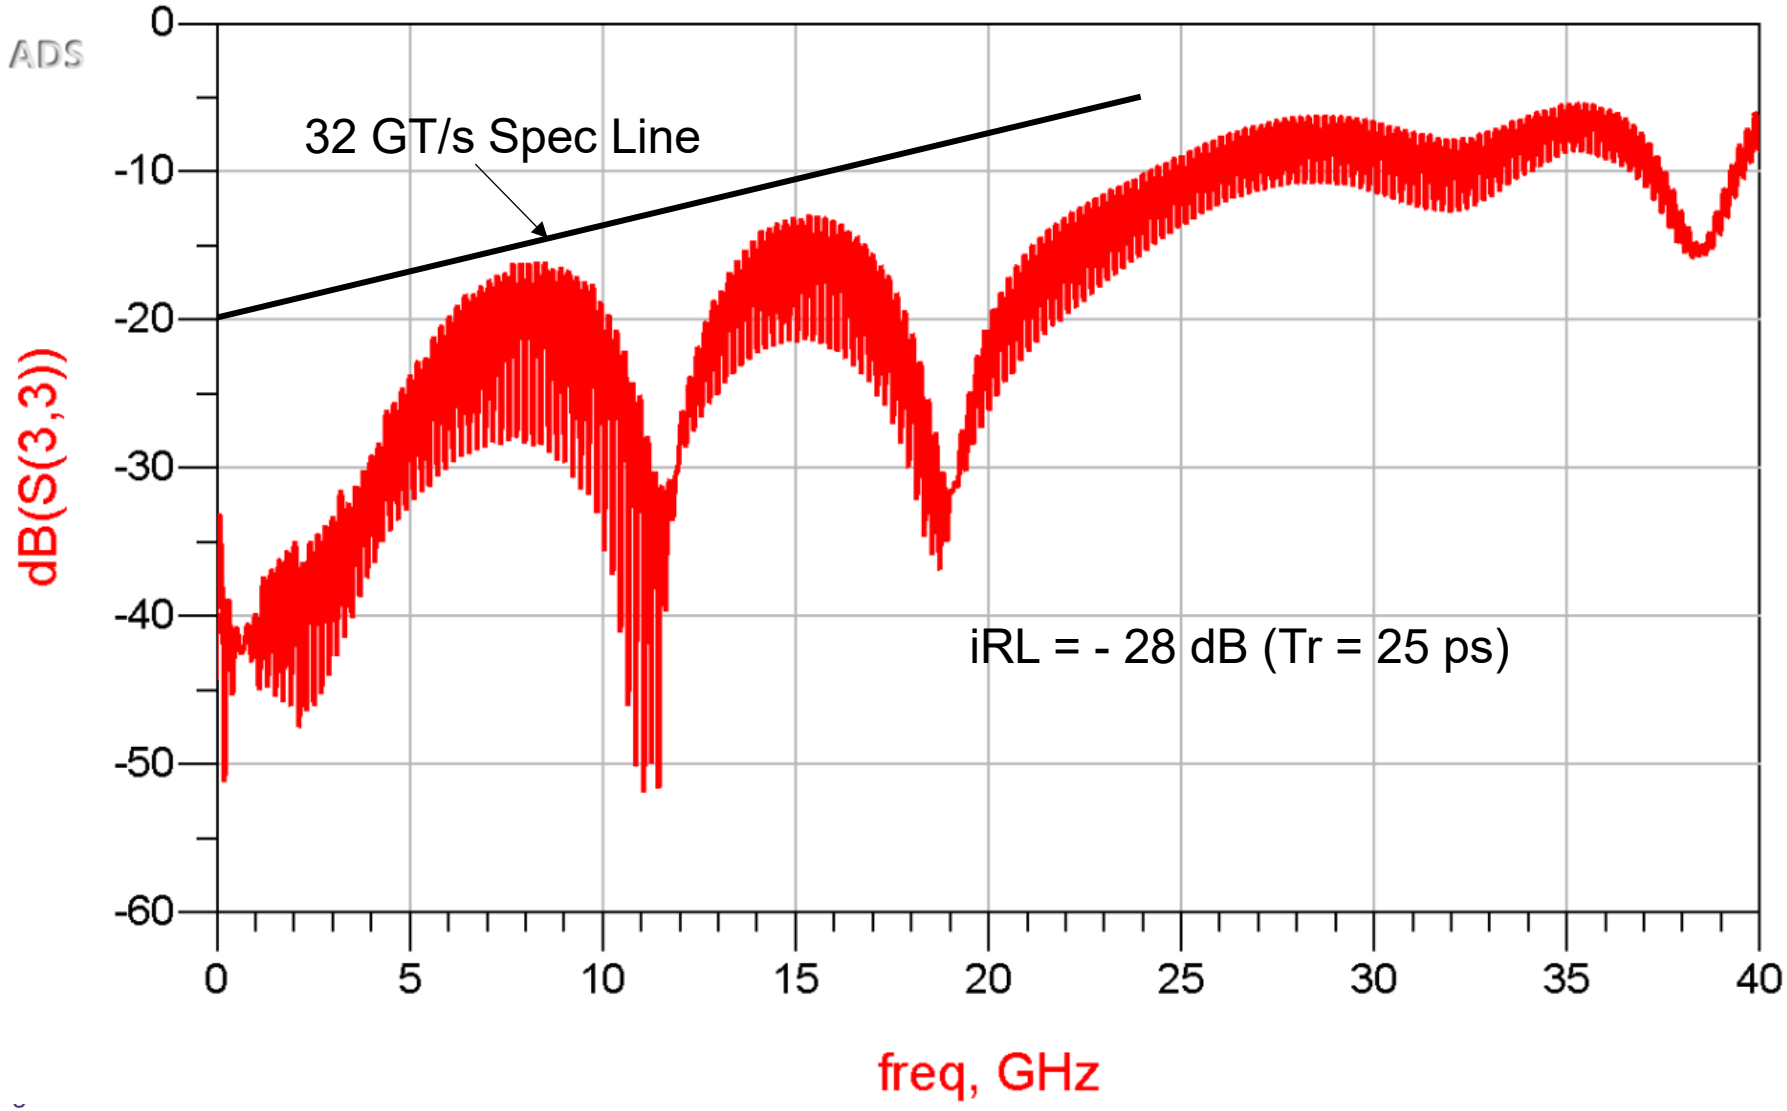

Scaled Differential Return Loss S-Parameters

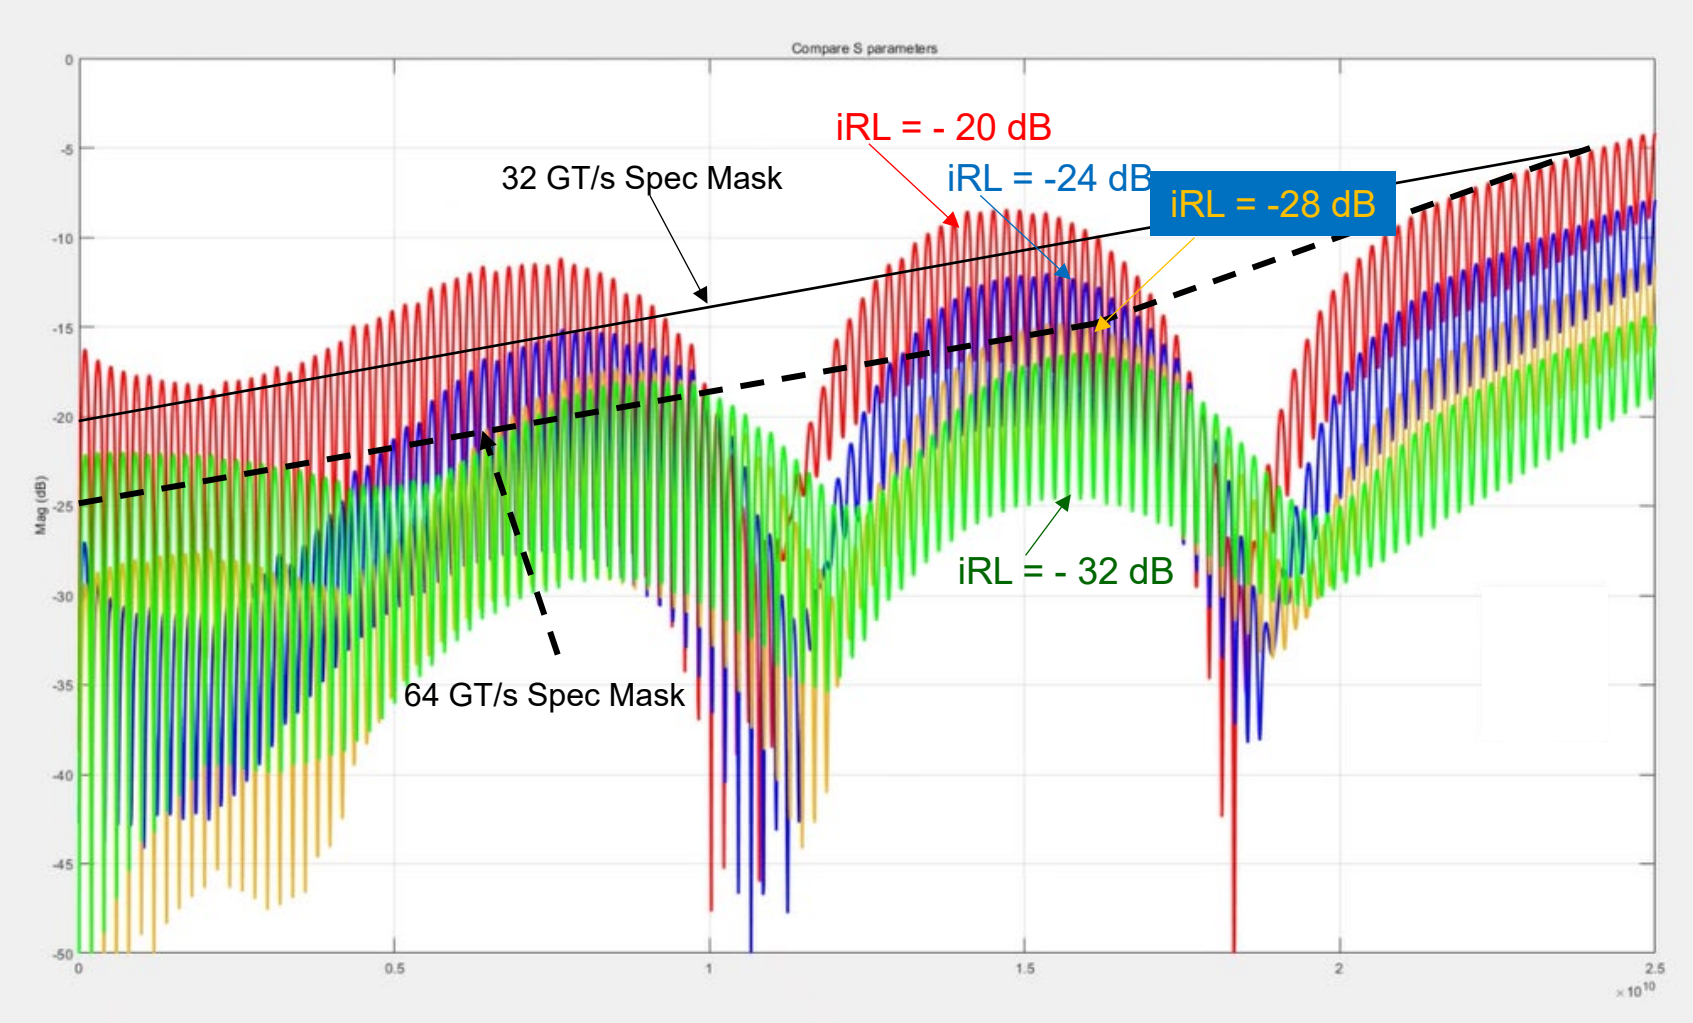

Short Channel, All Low Impedance Case – Return Loss Impact on Eye Margin at 64 GT/s

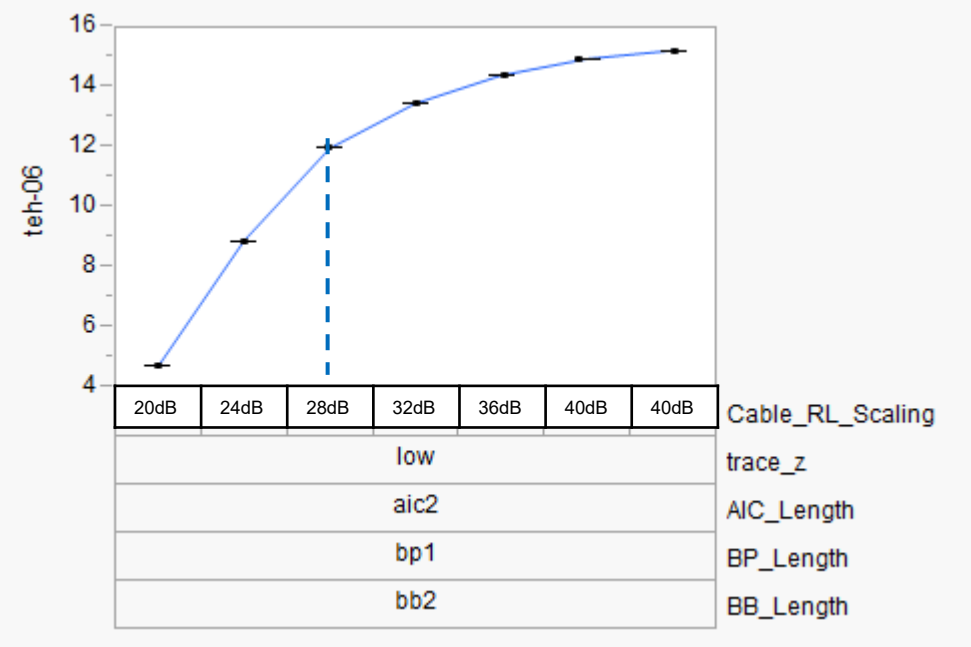

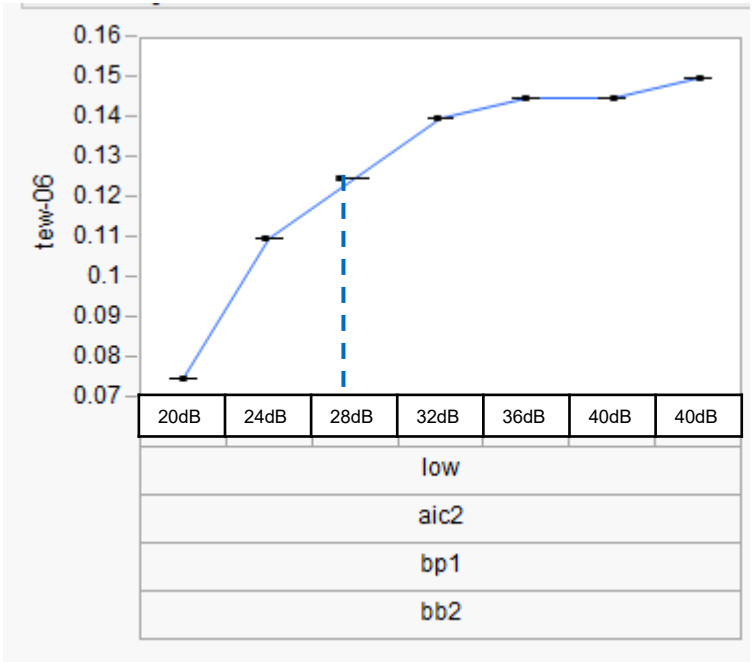

Significant degradation in eye height and eye width with increase of reflection – negating cable low loss benefit

An Example Topology for Crosstalk Impact Assessment at 32 GT/s

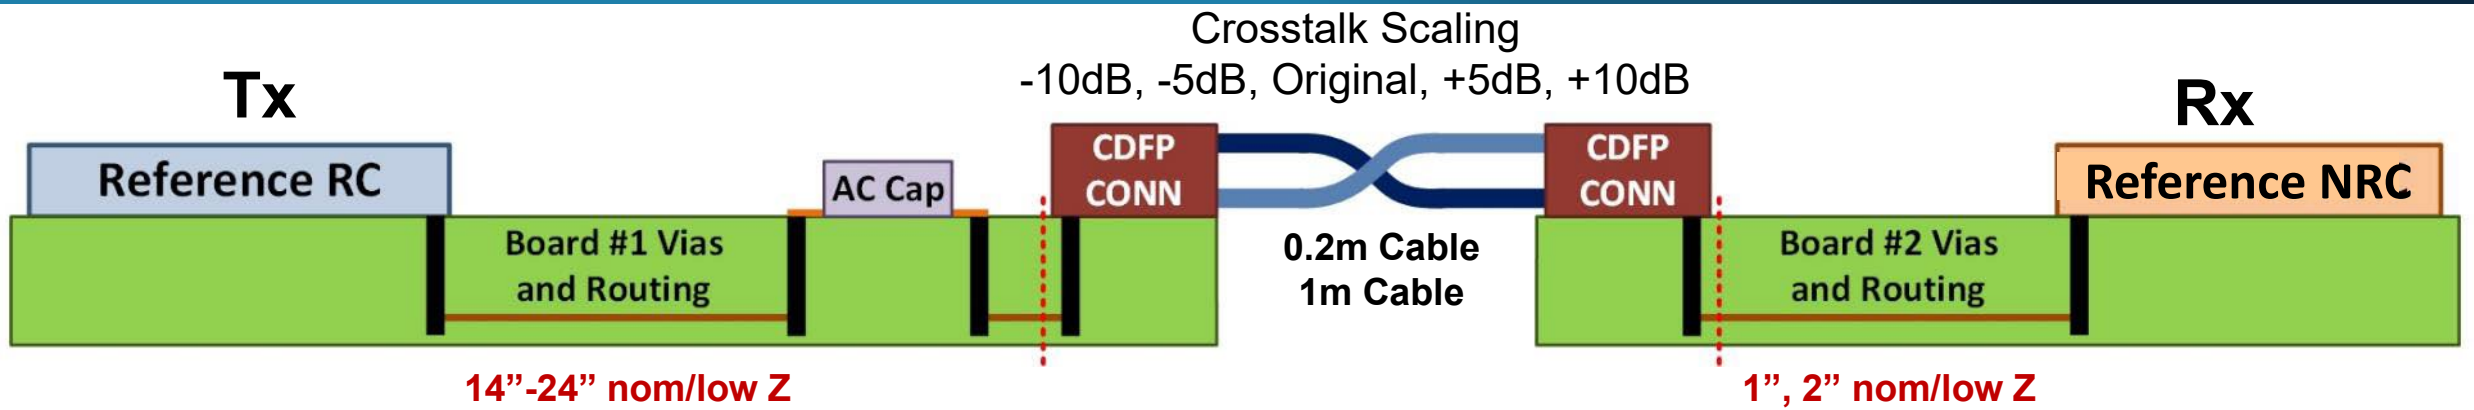

- Used all available CDFP (SFF-TA-1032) crosstalk aggressors in S32P model
  - 3 FEXT 0.8V, 4 NEXT 1.0V - included in simulation
  - Port 2/4 as Victim after evaluation of all Ports 2,4,6,8
- No Crosstalk in any component except of Cable assembly. Pre- and Post-Channel lanes are replicated and the same to ensure we see only the Cable Assembly behavior
- PCIe® Gen6 RC/NRC Reference Package Models – RC to NRC
- Ultra low-loss PCB (~ -1dB/inch @ 16GHz), Low and Nom Z
- 0.2m and 1m CDFP Cable Assembly Models from TE received May 12th and 18th 2022
  (TE_Mated_CDFP_CableAssembly_1p0m_20220518.s32p, TE_Mated_CDFP_CableAssembly_0p2m_20220512.s32p)

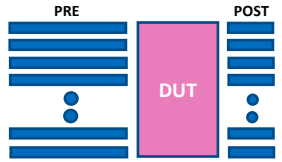

CDFP (SFF-TA-1032) CopprLink™ External Mated Cable Assembly Port Assignment

Seq	Pair	Port	Signal	Position		Position	Signal	Port	Pair	Seq
1	1	1	Tx5p	A20		C20	Rx5p	2	2	9
		3	Tx5n	A21		C21	Rx5n	4		
2	3	5	Tx3p	A23		C23	Rx3p	6	4	10
		7	Tx3n	A24		C24	Rx3n	8		
3	5	9	Tx4p	B20		D20	Rx4p	10	6	11
		11	Tx4n	B21		D21	Rx4n	12		
4	7	13	Tx2p	B23		D23	Rx2p	14	8	12
		15	Tx2n	B24		D24	Rx2n	16		
5	9	17	Rx5p	C20		A20	Tx5p	18	10	13
		19	Rx5n	C21		A21	Tx5n	20		
6	11	21	Rx3p	C23		A23	Tx3p	22	12	14
		23	Rx3n	C24		A24	Tx3n	24		
7	13	25	Rx4p	D20		B20	Tx4p	26	14	15
		27	Rx4n	D21		B21	Tx4n	28		
8	15	29	Rx2p	D23		B23	Tx2p	30	16	16
		31	Rx2n	D24		B24	Tx2n	32		

Port 2/4 as Victim

FEXT 3,5,7 for Vic 2
FEXT 1,5,7 for Vic 4

NEXT 10,12,14,16

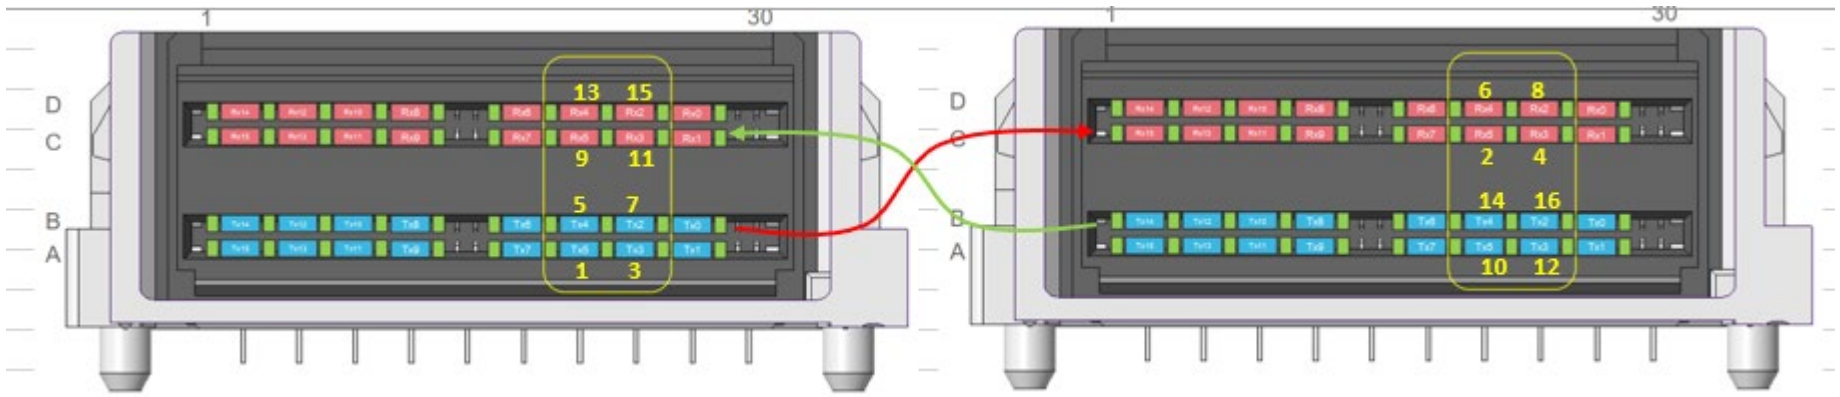

Crosstalk Scaling for Crosstalk Sensitivity Studies

- Used MATLAB script to scale touchstone crosstalk path by -10dB, -5dB, +5dB, +10dB. Red is original model
- Only FEXT, NEXT S(x,y) crosstalk path are scaled by same amount. All other S-Par paths untouched
- Scaled models ports 2,4 as Victims. SQC verified for Passivity/Causality

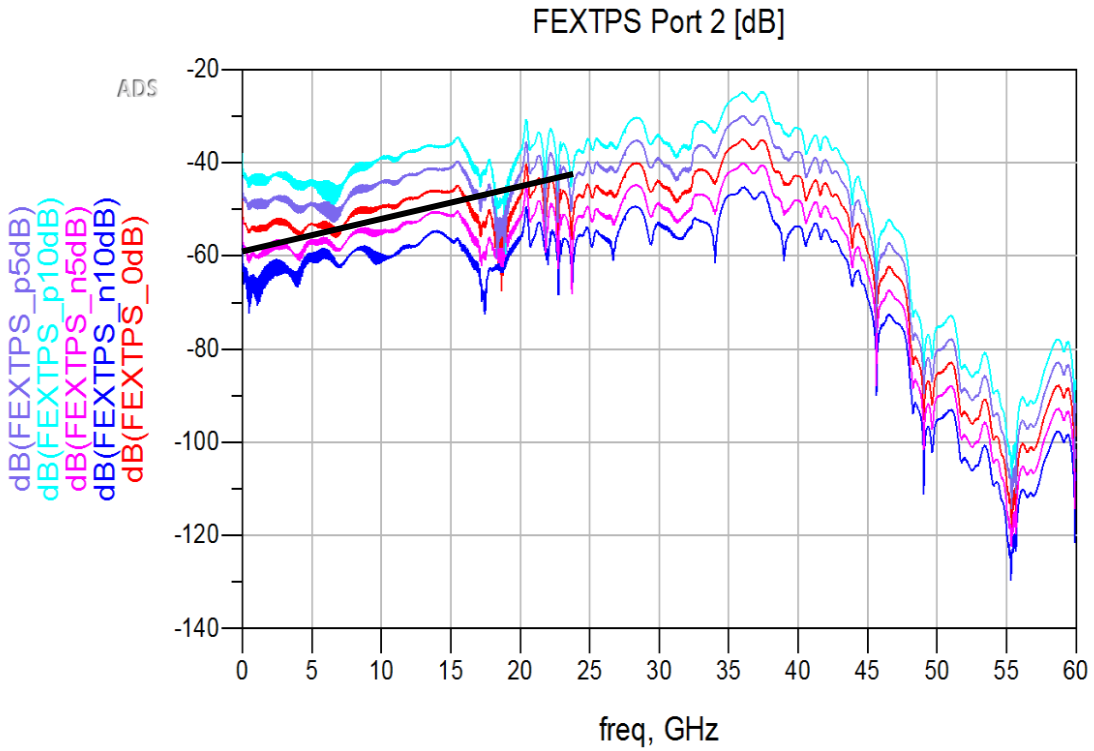

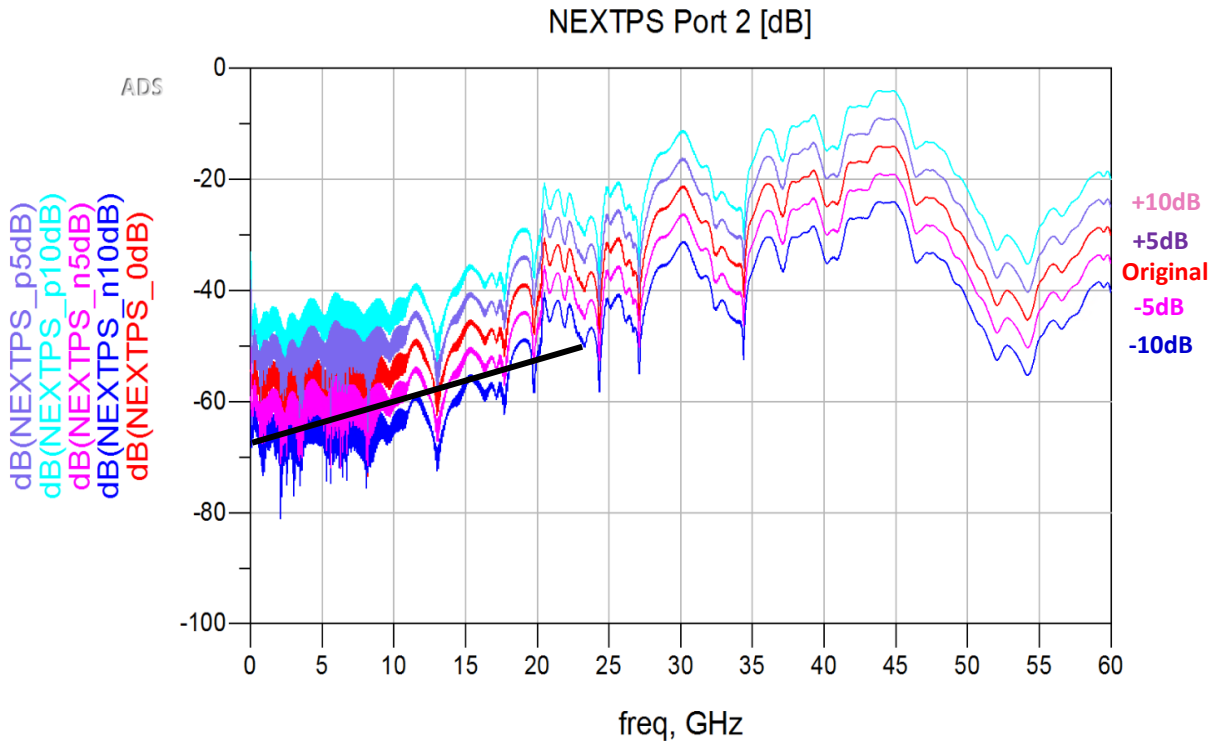

CDFP (SFF-TA-1032) 1m Cable Link Margin Degradation vs. Scaled Crosstalk

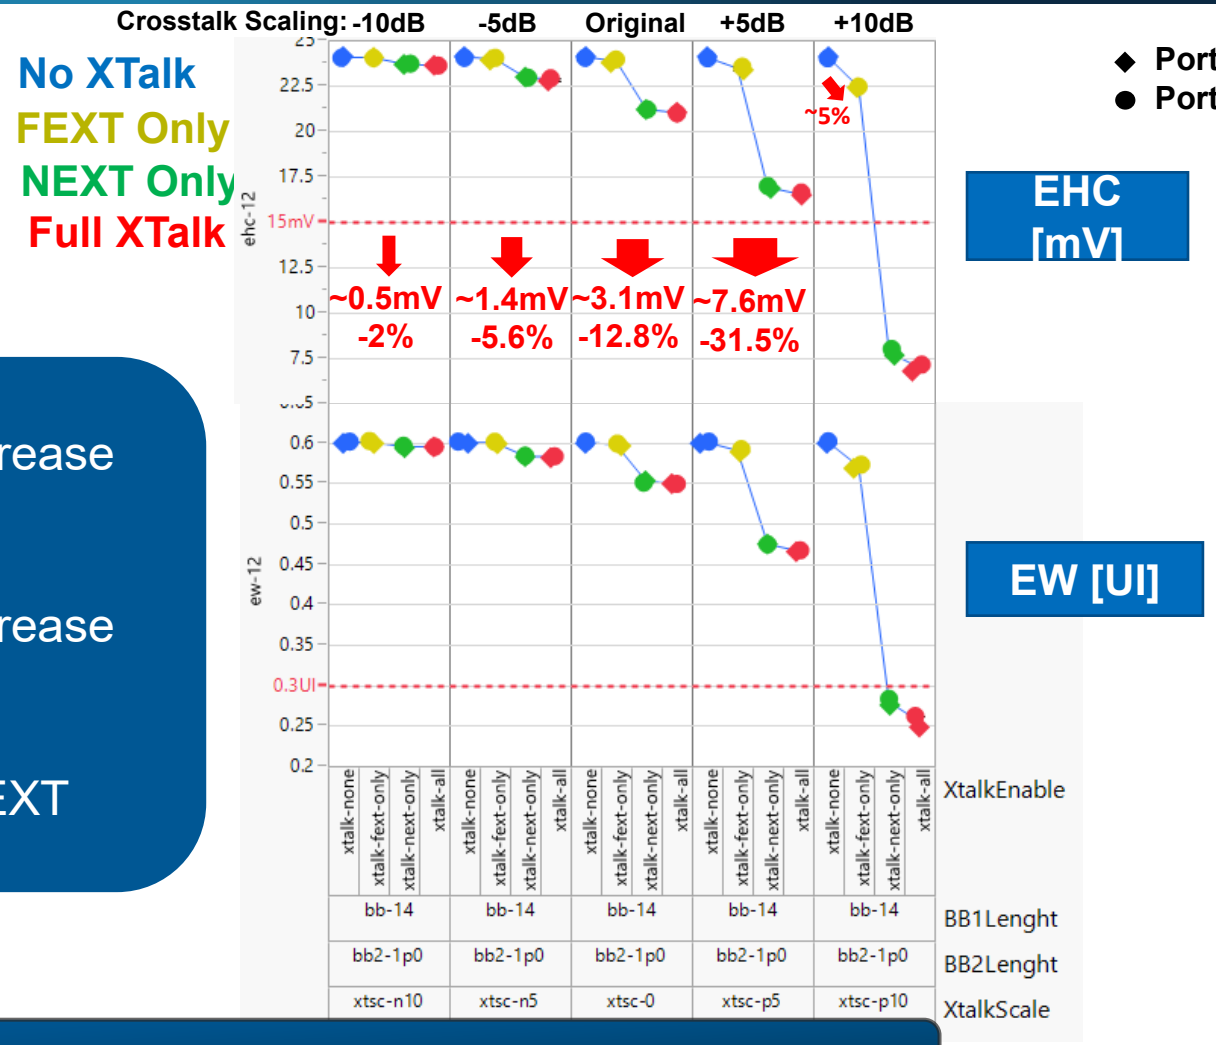

- Relaxation of NEXT and FEXT by +5dB will increase EH degradation by ~ 2.5x (12.8% to 31.5%)
- Tightening of NEXT and FEXT by -5dB will decrease EH degradation by ~2.3x (12.8% to 5.6%)
- FEXT impact comparably smaller relative to NEXT

Xtalk scaling methodology quantifies the component and system-level trade-offs

Agenda

- Need for Cable Solutions
- PCI-SIG®'s Cabling – Journey to CopprLink™ and beyond
- A Brief Overview of PCIe® 5.0/6.0 CopprLink Internal and External Cable Specifications
- Example Methodologies to Quantify Electrical Performance of Cable Solutions
- Sideband Specifications of CopprLink Internal and External Cables
- Update on PCIe Optical Cable Solution for PCIe 6.0 specification and beyond

Signal Types and Their Functions: Internal Cable

Signal	Input/Output (From Root Complex End)	Description
PET[p/n]	Output	Differential PCI Express Transmitter Lanes
PER[p/n]	Input	Differential PCI Express Receiver Lanes
2WCL	Output	Management Bus Clock
2WDA	Input/Output	Management Bus Data
PERST#	Output	Active low, push-pull at source. A discrete functional reset to the endpoint device as defined by the PCI Express Base specification.
PRPE	Input/Output	Bidirectional signal (**PR**esence/**PE**STI) used to indicate the attachment of a cable assembly and a module to a port. Optional use as PESTI (see Section 12.3.3 of [PCIe Base])[1]. If PESTI is not supported, the Root-Complex and Non-Root-Complex must support a Presence Indicator with the Root-Complex having a 100K (+/- 5%) pull-up to 3.3V and the Non-Root-Complex with a 2.27K or less (must be greater than 200 ohms) pull-down following power-up.
RCLKp / RCLKn	Output	PCIe Reference Clock Signal (100 MHz)
FLEXIO[0:9]		Flexible sideband pins. The discovery mechanisms for the specific usages of the pins have been specified in Section 12.3.3 of [PCIe Base][1]
USB2p/USB2n	Input/Output	Optional Universal Serial Bus 2.0.

- PCIe® 5.0/6.0 CopprLink™ Internal Cable is a passive cable with no power pins
- Each end of the mated cable has a fixed connector with two sides – Side A and Side B
- The cable is a full crossover cable, and it connects A side of the pins at one end of the cable to the corresponding B side of the pins at the other end of the cable

Internal Cable Sideband Sets: Baseline and Extended

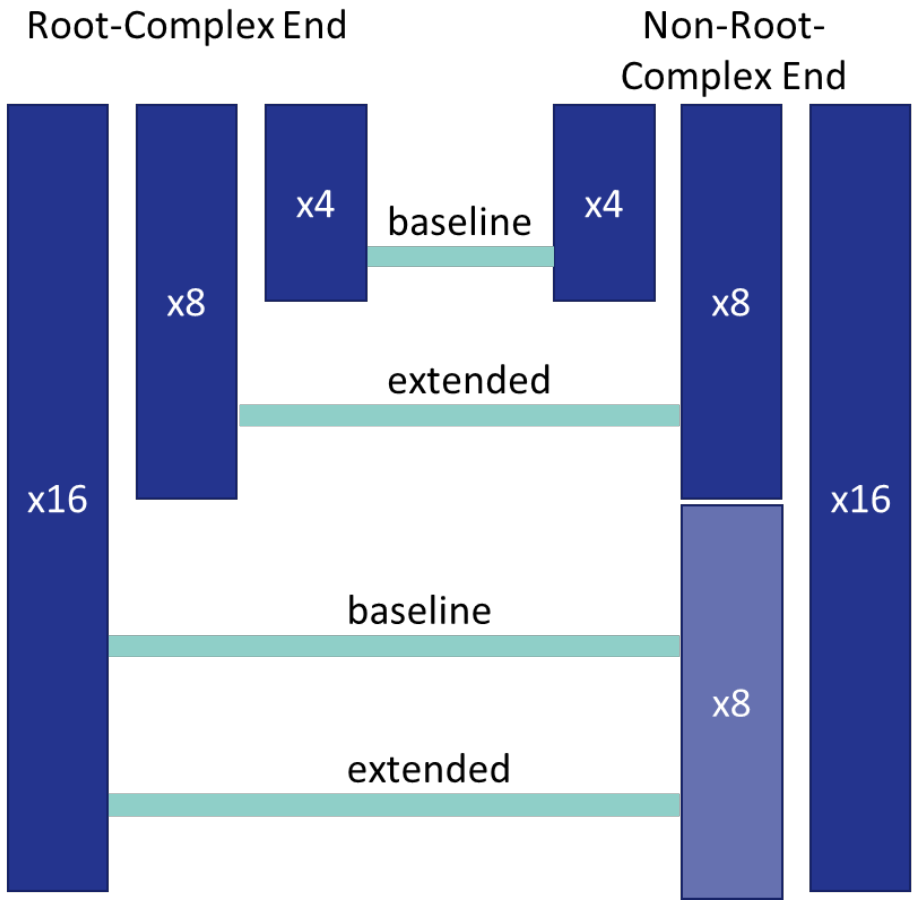

38 pin: x4 with 10 sidebands

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
G	S	S	G	S	S	G	SB	SB	SB	SB	SB	G	S	S	G	S	S	G
G	S	S	G	S	S	G	SB	SB	SB	SB	SB	G	S	S	G	S	S	G
20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38

74 pin: x8 with 20 sidebands

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37
G	S	S	G	S	S	G	SB	SB	SB	SB	SB	G	S	S	G	S	S	G	S	S	G	S	S	G	SB	SB	SB	SB	SB	G	S	S	G	S	S	G
G	S	S	G	S	S	G	SB	SB	SB	SB	SB	G	S	S	G	S	S	G	S	S	G	S	S	G	SB	SB	SB	SB	SB	G	S	S	G	S	S	G
38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74

Baseline and extended set of sideband pins provide an array of system manageability functions

Internal Cable: Baseline Set of Sideband Pins

Root-Complex End		Direction	Non-Root-Complex End	
Pin	Signal		Signal	Pin
A7	GND	----------	GND	B7
A8	FLEXIO9_A	<-------->	FLEXIO9_A	B8
A9	FLEXIO0_A	<-------->	FLEXIO0_A	B9
A10	GND	----------	GND	B10
A11	RCLK_Ap	<-------->	RCLK_Ap	B11
A12	RCLK_An	<-------->	RCLK_An	B12
A13	GND	----------	GND	B13

Root-Complex End		Direction	Non-Root-Complex End	
Pin	Signal		Signal	Pin
B7	GND	----------	GND	A7
B8	2WCL_A	--------->	2WCL_A	A8
B9	2WDA_A	<-------->	2WDA_A	A9
B10	GND	----------	GND	A10
B11	PERST#_A	--------->	PERST#_A	A11
B12	PRPE_A	<-------->	PRPE_A	A12
B13	GND	----------	GND	A13

A total of 10 pins is allocated for the Baseline set of sideband pins

Internal Cable: Extended Set of Sideband Pins

Root-Complex End		Direction	Non-Root-Complex End	
Pin	Signal		Signal	Pin
A25	GND	----------	GND	B25
A26	FLEXIO5_A	<--------->	FLEXIO5_A	B26
A27	FLEXIO6_A	<--------->	FLEXIO6_A	B27
A28	GND	----------	GND	B28
A29	FLEXIO7_A	<--------->	FLEXIO7_A	B29
A30	FLEXIO8_A	<--------->	FLEXIO8_A	B30
A31	GND	----------	GND	B31
B25	GND	----------	GND	A25
B26	FLEXIO3_A	<--------->	FLEXIO3_A	A26
B27	FLEXIO4_A	<--------->	FLEXIO4_A	A27
B28	GND	----------	GND	A28
B29	FLEXIO1_A	<--------->	FLEXIO1_A	A29
B30	FLEXIO2_A	<--------->	FLEXIO2_A	A30
B31	GND	----------	GND	A31

- If the FLEXIO pins 1-4 are used for x1 PCIe® technology
  - FLEXIO3-FLEXIO4 = p-n of the Tx at the Root-Complex End
  - FLEXIO3-FLEXIO4 = p-n of the Rx at the Non-Root-Complex End
  - FLEXIO5-FLEXIO6 = p-n of the Rx at the Root-Complex End
  - FLEXIO5-FLEXIO6 = p-n of the Tx at the Non-Root-Complex End
- The differential pair FLEXIO7-FLEXIO8 is allocated for optional use for USB2p-USB2n. The pins FLEXIO1 and FLEXIO2 can be used as single-ended pins or as a differential pair.

Example Pinout: x4 Fixed Connector

Root-Complex End		Direction	Non-Root-Complex End	
Pin	Signal		Signal	Pin
A1	GND		GND	B1
A2	PERp0	←	PETp0	B2
A3	PERn0	←	PETn0	B3
A4	GND		GND	B4
A5	PERp1	←	PETp1	B5
A6	PERn1	←	PETn1	B6
A7	GND		GND	B7
A8	FLEXIO9_A	↔	FLEXIO9_A	B8
A9	FLEXIO0_A	↔	FLEXIO0_A	B9
A10	GND		GND	B10
A11	RCLK_Ap	→	RCLK_Ap	B11
A12	RCLK_An	→	RCLK_An	B12
A13	GND		GND	B13
A14	PERp2	←	PETp2	B14
A15	PERn2	←	PETn2	B15
A16	GND		GND	B16
A17	PERp3	←	PETp3	B17
A18	PERn3	←	PETn3	B18
A19	GND		GND	B19

Root-Complex End		Direction	Non-Root-Complex End	
Pin	Signal		Signal	Pin
B1	GND		GND	A1
B2	PETp0	→	PERp0	A2
B3	PETn0	→	PERn0	A3
B4	GND		GND	A4
B5	PETp1	→	PERp1	A5
B6	PETn1	→	PERn1	A6
B7	GND		GND	A7
B8	2WCL_A	→	2WCL_A	A8
B9	2WDA_A	↔	2WDA_A	A9
B10	GND		GND	A10
B11	PERST#_A	→	PERST#_A	A11
B12	PRPE_A	↔	PRPE_A	A12
B13	GND		GND	A13
B14	PETp2	→	PERp2	A14
B15	PETn2	→	PERn2	A15
B16	GND		GND	A16
B17	PETp3	→	PERp3	A17
B18	PETn3	→	PERn3	A18
B19	GND		GND	A19

Signal Types and Their Functions: External Cable

Signal	Input/Output (From Root Complex End)	Description
PET[p/n]	Output	Differential PCI Express Transmitter Lanes
PER[p/n]	Input	Differential PCI Express Receiver Lanes
2WCL	Output	Management bus (Clock/Data: 2WCL/2WDA) for remote two wire interface. The Root-Complex and Non-Root Complex must provide isolation on the bias for these pins until local power is good to avoid backfeed/leakage. Remote bus must traverse the cable assembly.
2WDA	Input/Output	
SCL	Output	Local management bus (Clock/Data: SCL/SDA) interface between the host and the cable using I2C protocol. A memory device should be connected on each side of the cable assembly via VCC3p3V and the local management bus. The memory device must be compliant with the CMIS, supporting 8-bit addressing and access to 256 bytes of available space to support identification and management functions. Local bus must not traverse the cable assembly.
SDA	Input/Output	
PERST#	Output	Active low, push-pull at source. A discrete functional reset to the endpoint device as defined by the PCI Express Base specification.
PRPE	Input/Output	Bidirectional signal (**PR**esence/**PE**STI) used to indicate the attachment of a cable assembly and a module to a port. Optional use as PESTI (see Chapter 12 of [PCIe Base])[1]. If PESTI is not supported, the Root-Complex and Non-Root-Complex must support a Presence Indicator with the Root-Complex having a 100K (+/- 5%) pull-up to 3.3V and the Non-Root-Complex with a 2.27K or less (must be greater than 200 ohms) pull-down following power-up.
VCC3p3V	Output	Required local power for cable management only – must not be wired through the cable
VCC12V	Output	Optional local power for any special cable that requires power – must not be wired through the cable. If not used, must be left as not connected.
FLEXIO[1:8]	Input/Output	Flexible sideband pins. The discovery mechanisms for the specific usages of these pins have been specified in Chapter 12 of [PCIe Base][1], and the cable assembly must advertise the supported application features in the designated EEPROM locations, see Table 3-9.
RCLKp / RCLKn	Output	Optional PCIe Reference Clock Signal (100 MHz). If implemented in x4, x8, and x16 configurations, FLEXIO3 is RCLKp and FLEXIO4 is RCLKn.
USB2p/USB2n	Input/Output	Optional Universal Serial Bus 2.0. If implemented in x4 and x8 configurations, FLEXIO1 is USB2p and FLEXIO2 is USB2n. If implemented in x16 configuration, FLEXIO7 is USB2p and FLEXIO8 is USB2n.

- PCIe® 5.0/6.0 External Cable is primarily a passive cable with power for cable management
- Each end of the mated cable has a fixed connector with four sides – Side A, B, C, and D
- The cable has full crossover pairs for high-speed signals and extended sideband signals, and direct point to point connections for baseline side band signals

Example Signal Positions: x4 Fixed Connector

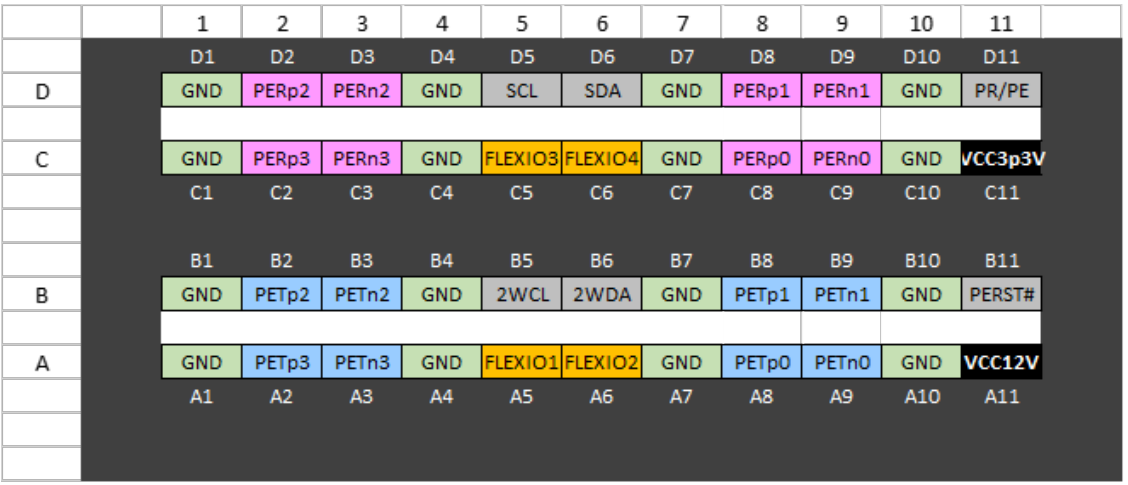

	1	2	3	4	5	6	7	8	9	10	11
	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11
D	GND	PERp2	PERn2	GND	SCL	SDA	GND	PERp1	PERn1	GND	PR/PE
C	GND	PERp3	PERn3	GND	FLEXIO3	FLEXIO4	GND	PERp0	PERn0	GND	VCC3p3V
	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11
	B1	B2	B3	B4	B5	B6	B7	B8	B9	B10	B11
B	GND	PETp2	PETn2	GND	2WCL	2WDA	GND	PETp1	PETn1	GND	PERST#
A	GND	PETp3	PETn3	GND	FLEXIO1	FLEXIO2	GND	PETp0	PETn0	GND	VCC12V
	A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11

Signal positions at the Root Complex End

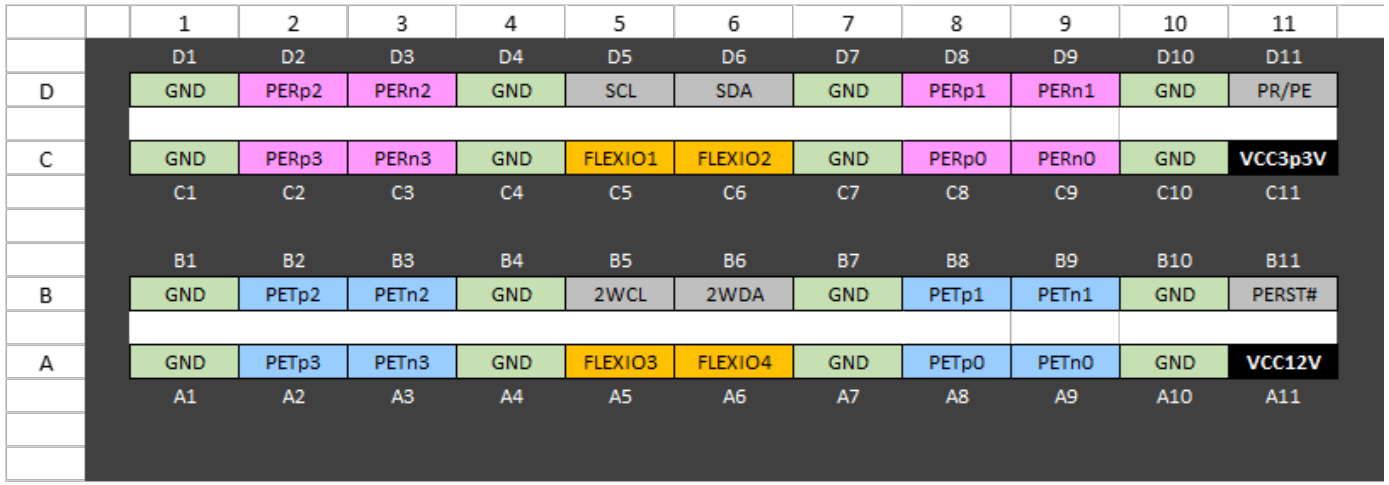

	1	2	3	4	5	6	7	8	9	10	11
	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11
D	GND	PERp2	PERn2	GND	SCL	SDA	GND	PERp1	PERn1	GND	PR/PE
C	GND	PERp3	PERn3	GND	FLEXIO1	FLEXIO2	GND	PERp0	PERn0	GND	VCC3p3V
	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11
	B1	B2	B3	B4	B5	B6	B7	B8	B9	B10	B11
B	GND	PETp2	PETn2	GND	2WCL	2WDA	GND	PETp1	PETn1	GND	PERST#
A	GND	PETp3	PETn3	GND	FLEXIO3	FLEXIO4	GND	PETp0	PETn0	GND	VCC12V
	A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11

Signal positions at the Non-Root Complex End

External Cable Sideband Sets: Baseline and Extended

- X4, x8, and x16 cables provide a Baseline Set of eight sideband signals
- X4 and x8 cables provide an Extended Set of four sideband signals and x16 cables provide an Extended Set of eight sideband signals

Baseline and extended set of sideband pins provide an array of system manageability functions

Baseline Set for x4 External Cable Connector

Root-Complex End		Direction	Non-Root-Complex End	
Pin	Signal		Signal	Pin
A10	GND	———	GND	A10
A11	VCC12V	N/C	VCC12V	A11
B4	GND		GND	B4
B5	2WCL	——→	2WCL	B5
B6	2WDA	←——→	2WDA	B6
B7	GND		GND	B7
B10	GND	———	GND	B10
B11	PERST#	——→	PERST#	B11
C10	GND	———	GND	C10
C11	VCC3p3V	N/C	VCC3p3V	C11
D4	GND	———	GND	D4
D5	SCL	N/C	SCL	D5
D6	SDA	N/C	SDA	D6
D7	GND	———	GND	D7
D10	GND	———	GND	D10
D11	PRPE	←——→	PRPE	D11

Extended Set for x4 External Cable Connector

Root-Complex End		Direction	Non-Root-Complex End	
Pin	Signal		Signal	Pin
A4	GND	⟷	GND	C4
A5	FLEXIO1_RC	⟷	FLEXIO1_NRC	C5
A6	FLEXIO2_RC	⟷	FLEXIO2_NRC	C6
A7	GND	⟷	GND	C7
C4	GND	⟷	GND	A4
C5	FLEXIO3_RC	⟷	FLEXIO3_NRC	A5
C6	FLEXIO4_RC	⟷	FLEXIO4_NRC	A6
C7	GND	⟷	GND	A7

- If the FLEXIO pins 1-4 are used for x1 PCIe® technology
  - FLEXIO1_RC-FLEXIO2_RC = p-n of the Tx at the Root-Complex End
  - FLEXIO1_NRC-FLEXIO2_NRC = p-n of the Rx at the Non-Root-Complex End
  - FLEXIO3_RC-FLEXIO4_RC = p-n of the Rx at the Root-Complex End
  - FLEXIO3_NRC-FLEXIO4_NRC = p-n of the Tx at the Non-Root-Complex End
- The differential pair FLEXIO1-FLEXIO2 is allocated for optional use for USB2p-USB2n. The differential pair FLEXIO3-FLEXIO4 is allocated for optional use for RCLKp-RCLKn.

	1	2	3	4	5	6	7	8	9	10	11
	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11
D	GND	PERp2	PERn2	GND	SCL	SDA	GND	PERp1	PERn1	GND	PR/PE
C	GND	PERp3	PERn3	GND	FLEXIO3	FLEXIO4	GND	PERp0	PERn0	GND	VCC3p3V
	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11
	B1	B2	B3	B4	B5	B6	B7	B8	B9	B10	B11
B	GND	PETp2	PETn2	GND	2WCL	2WDA	GND	PETp1	PETn1	GND	PERST#
A	GND	PETp3	PETn3	GND	FLEXIO1	FLEXIO2	GND	PETp0	PETn0	GND	VCC12V
	A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11

Signal positions at the Root Complex End

	1	2	3	4	5	6	7	8	9	10	11
	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11
D	GND	PERp2	PERn2	GND	SCL	SDA	GND	PERp1	PERn1	GND	PR/PE
C	GND	PERp3	PERn3	GND	FLEXIO1	FLEXIO2	GND	PERp0	PERn0	GND	VCC3p3V
	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11
	B1	B2	B3	B4	B5	B6	B7	B8	B9	B10	B11
B	GND	PETp2	PETn2	GND	2WCL	2WDA	GND	PETp1	PETn1	GND	PERST#
A	GND	PETp3	PETn3	GND	FLEXIO3	FLEXIO4	GND	PETp0	PETn0	GND	VCC12V
	A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11

Signal positions at the Non-Root Complex End

Example Pinout: x4 Fixed Connector

Root-Complex End		Direction	Non-Root-Complex End	
Pin	Signal		Signal	Pin
A1	GND	----------	GND	C1
A2	PETp3	---------->	PERp3	C2
A3	PETn3	---------->	PERn3	C3
A4	GND	----------	GND	C4
A5	FLEXIO1_RC	<--------->	FLEXIO1_NRC	C5
A6	FLEXIO2_RC	<--------->	FLEXIO2_NRC	C6
A7	GND	----------	GND	C7
A8	PETp0	---------->	PERp0	C8
A9	PETn0	---------->	PERn0	C9
A10	GND	----------	GND	C10
A11	VCC12V	N/C	VCC12V	A11

Pin	Signal		Signal	Pin
C1	GND	----------	GND	A1
C2	PERp3	<----------	PETp3	A2
C3	PERn3	<----------	PETn3	A3
C4	GND	----------	GND	A4
C5	FLEXIO3_RC	<--------->	FLEXIO3_NRC	A5
C6	FLEXIO4_RC	<--------->	FLEXIO4_NRC	A6
C7	GND	----------	GND	A7
C8	PERp0	<----------	PETp0	A8
C9	PERn0	<----------	PETn0	A9
C10	GND	----------	GND	A10
C11	VCC3p3V	N/C	VCC3p3V	C11

Root-Complex End		Direction	Non-Root-Complex End	
Pin	Signal		Signal	Pin
B1	GND	----------	GND	D1
B2	PETp2	---------->	PERp2	D2
B3	PETn2	---------->	PERn2	D3
B4	GND	----------	GND	D4
B5	2WCL	---------->	2WCL	B5
B6	2WDA	<--------->	2WDA	B6
B7	GND	----------	GND	D7
B8	PETp1	---------->	PERp1	D8
B9	PETn1	---------->	PERn1	D9
B10	GND	----------	GND	D10
B11	PERST#	---------->	PERST#	B11

Pin	Signal		Signal	Pin
D1	GND	----------	GND	B1
D2	PERp2	<----------	PETp2	B2
D3	PERn2	<----------	PETn2	B3
D4	GND	----------	GND	B4
D5	SCL	N/C	SCL	D5
D6	SDA	N/C	SDA	D6
D7	GND	----------	GND	B7
D8	PERp1	<----------	PETp1	B8
D9	PERn1	<----------	PETn1	B9
D10	GND	----------	GND	B10
D11	PRPE	<--------->	PRPE	D11

Compatibility and Future Flexibility

- CMIS provides a flexible and extendable Memory Map
- Advertisements for:
  - Form Factor ID
  - PCI-SIG® Vendor ID
  - Host Interface ID
    - Far-End Lane Configuration
  - Media Interface ID
- PCIe® technology-specific implementations:
  - FLEXIO definition
  - Form Factor sideband management support
  - Placeholders for future PCIe specific needs
    - Optical Link Type, Latency, etc.
- Extendable to Upper Pages, as necessary
- Liaison relationship exists between PCI-SIG and Optical Interconnect Forum (OIF)

Table 3-8 CopprLink Memory Map

Byte [HEX]	Byte [DEC]	Field Description
0	0	SFF-8024 Form Factor Identifier (lower page), per CMIS
1	1	CMIS Revision Number, per CMIS
2	2	Advertisement of flat page memory (for passive copper cables assemblies) Reserved for CMIS functionality beyond the scope of this specification
3	3	Advertisement of interrupt not asserted, and ModuleReady Reserved for CMIS functionality beyond the scope of this specification
4-29	4-41	Reserved for CMIS functionality beyond the scope of this specification
2A-3F	42-63	Reserved for future CMIS definition
40-54	64-84	Used for CopprLink advertisement for PCIe specific implementations (i.e., FLEXIO supported connectivity, RCLK, USB, and 12V support, etc.) See Table 3-9. Reserved for Custom implementations
55-75	85-117	SFF-8024 Host/Media Advertisement, per CMIS
76-7F	118-127	Reserved for CMIS functionality beyond the scope of this specification
80	128	SFF-8024 Form Factor identifier (upper page), per CMIS
81-C7	129-199	Vendor Information
C8	200	ModulePowerClass advertisement, "Class 1", per CMIS
C9	201	Reserved for CMIS functionality beyond the scope of this specification
CA	202	Length advertisement, as constructed, per CMIS
CB-CF	203-207	Connector type and Attenuation advertisement, per CMIS
D0-D1	208-209	Reserved for future CMIS definition
D2-D4	210-212	Media Supported, FarEndConfiguration and Media Interface advertisement, per CMIS
D5-DC	213-220	Reserved for future CMIS definition
DD-FF	221-255	Reserved for Custom implementations

Opportunities for PCIe® Cabling

Storage Applications
Alternative to SAS/SATA

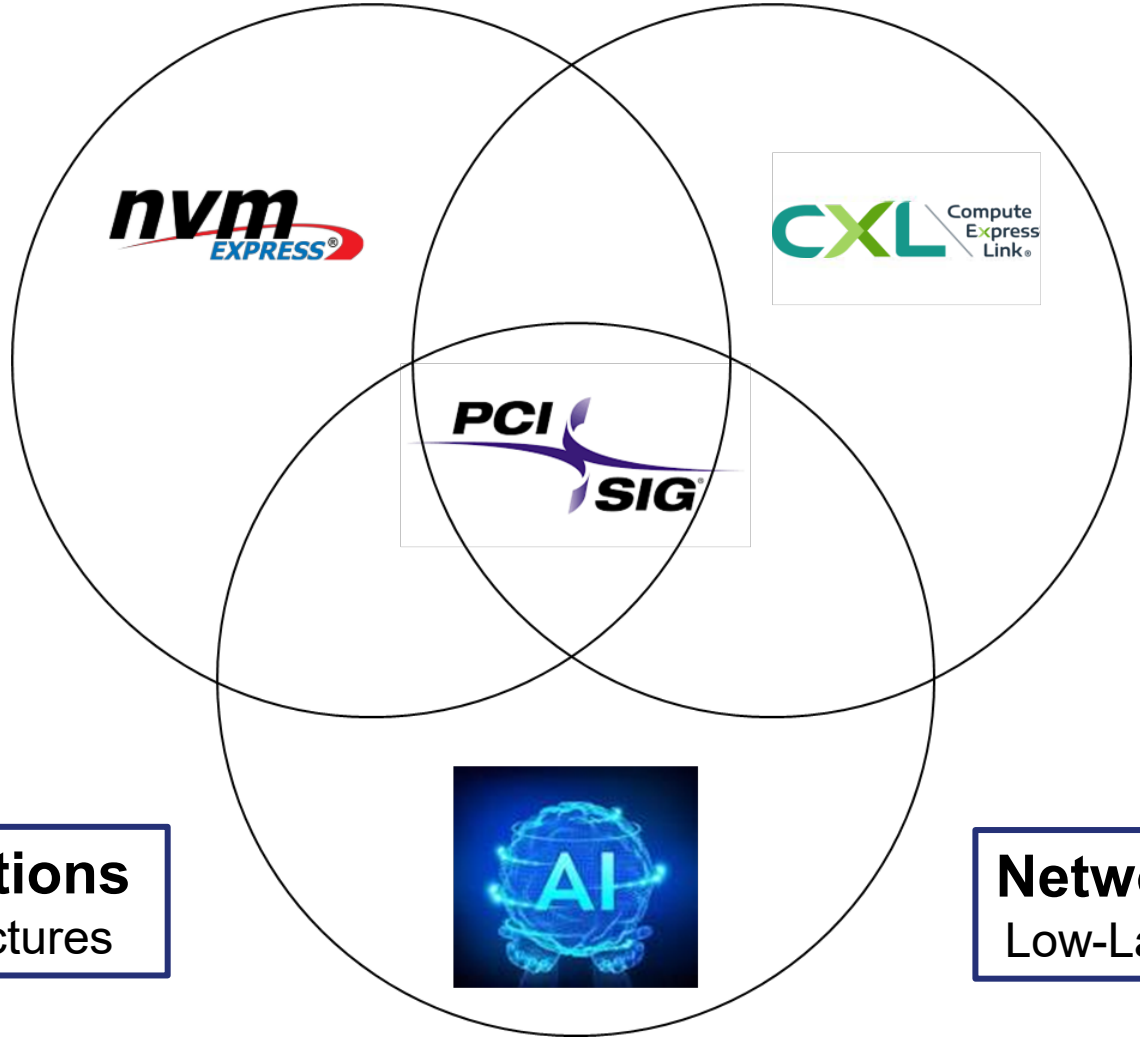

Pooled Memory Applications
CXL 3.0

CPU-GPU Applications
Disaggregated Architectures

Networked Accelerators
Low-Latency Compute Fabrics

SDC 24

Agenda

- Need for Cable Solutions
- PCI-SIG®'s Cabling – Journey to CopprLink™ and beyond
- A Brief Overview of PCIe® 5.0/6.0 CopprLink Internal and External Cable Specifications
- Example Methodologies to Quantify Electrical Performance of Cable Solutions
- Sideband Specifications of CopprLink Internal and External Cables
- Update on PCIe Optical Cable Solution for PCIe 6.0 specification and beyond

PCI-SIG® Optical Work Group

On August 2, 2023, PCI-SIG established a new work group to seek industry feedback in developing an optical interconnect

- The PCI-SIG Optical Workgroup intends to be optical technology-agnostic, supporting a wide range of optical technologies, while potentially developing technology-specific form factors
    - Potential form factors include pluggable optical transceivers, on-board optics, co-packaged optics and optical I/O
- Existing PCI-SIG workgroups will continue their generational march towards a 128GT/s data rate in the PCIe® 7.0 specification, while this new optical workgroup will work to make the PCIe architecture more optical-friendly
- Planned PCIe specification updates:
    - PCI-SIG will be using the same components with minimal changes (e.g., plan to use the same Flit Mode, the same Link training etc.). The potential spec enhancements deal with aspects like coordinating speed transitions to match with optical side, making side-band signals in-band, make the specification more power-efficient with a longer reach, etc.

PCI-SIG® Optical Status

- Optical has the promise of high bandwidth density and reach across a Rack/ Pod
  - Use case: Resource Pooling/ Sharing; Using PCIe® technology for developing composable systems with fabric topologies
  - Pros: Small (unlike copper cables that occupy more space) and reach (order of tens of meters vs Cu 1m)
  - Cons: multiple technologies, cost vs copper – let this play out
- Define optical extension in an optical technology neutral manner:
  - Work across optical technology choices – premature to pick one. E.g., VCSEL (vertical cavity surface emitting laser), TFLN (thin film lithium niobate), SiP (Si Photonics), WDM (wave division multiplexing), uLED, etc
  - PHY Logical enhancements (comprehend sideband, mapping of mainband bit stream), including a Retimer-based approach
  - Form-factor, if needed
- <u>Assumptions:</u>
  - Same type of Retimer is in both ends (optical technology neutrality) – starting position
  - EIC-PIC interface does not need to be defined
  - Complementary to copper cable work (different reach optimization)

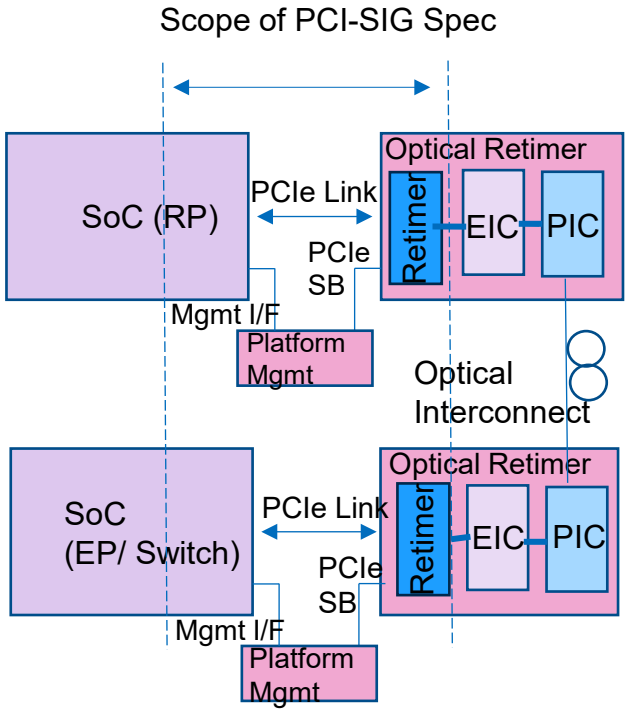

Q&A

Please take a moment to rate this session.

Your feedback is important to us.

