



112 Gbps Electrical Interfaces: Does rate drive architecture or does architecture enable rate?

Nathan Tracy, Technologist

April 12, 2019



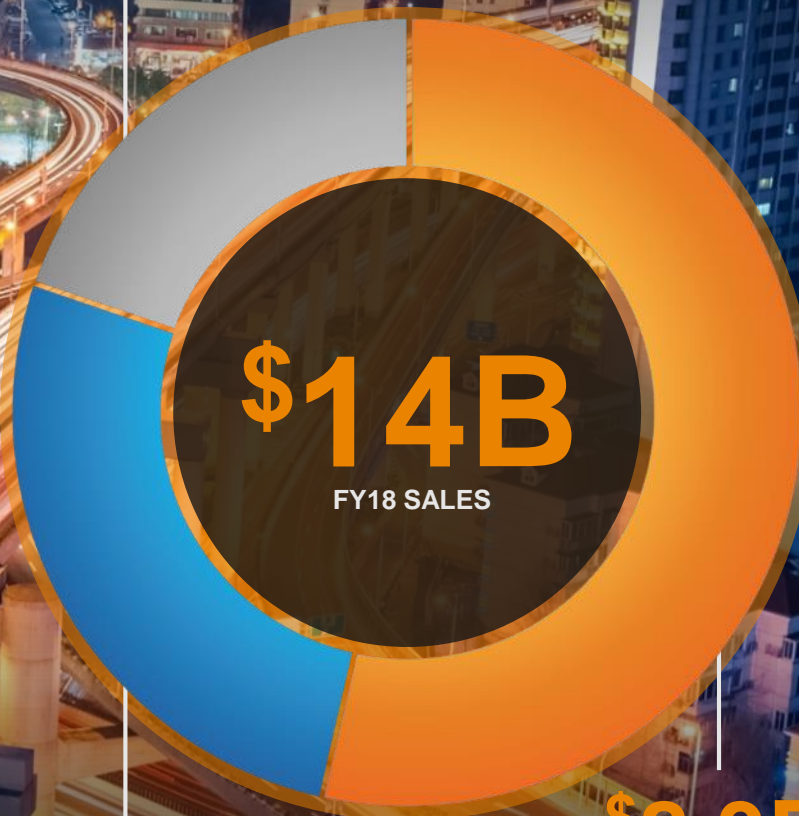
CONNECTING THE WORLD

220B

PRODUCTS MANUFACTURED ANNUALLY

\$1.8B

COMMUNICATIONS
Appliances, Data & Devices



\$3.9B

INDUSTRIAL
Industrial, Aerospace,
Defense & Marine,
Medical, Energy

\$8.3B

TRANSPORTATION
Automotive, Industrial & Commercial
Transportation, Sensors,
Application Tooling




Presentation Outline

- Introduction – Why higher rates?
- Challenges of going to higher rates
- What are the electrical channels and architectures being addressed?
- Impairments...and....Improvements
- Case Study: Considering Impedance, Skew and Reflections
 - Skew Impact on S_{dd21} & S_{cd21}
 - Mating Zone Reflections
 - Measured results of 112 Gbps Copper Cable
 - Impact of skew and reflections
 - Measured results of 112 Gbps Chip-2-Module
 - Impact of skew and reflections
- Conclusions

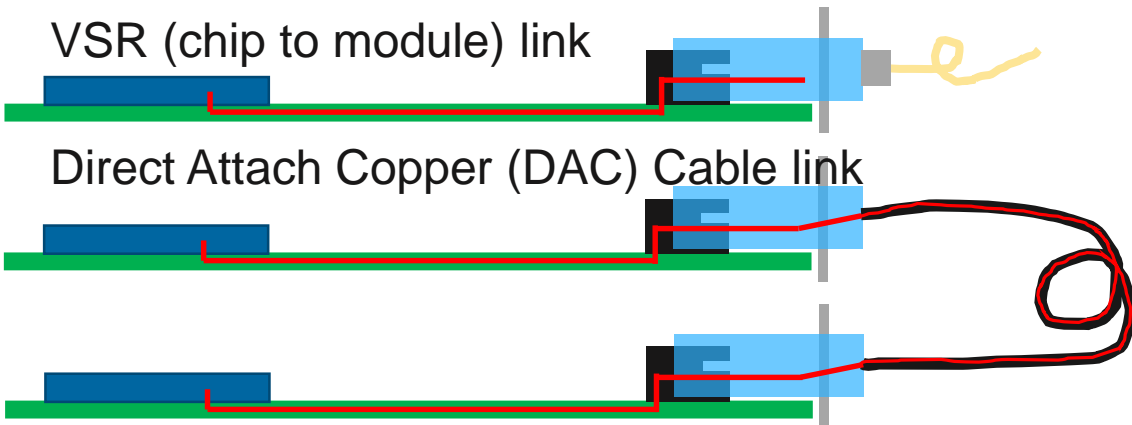
Introduction: Why 100G?

100 Gbps rate enables reasonable port counts and aligns with roadmaps



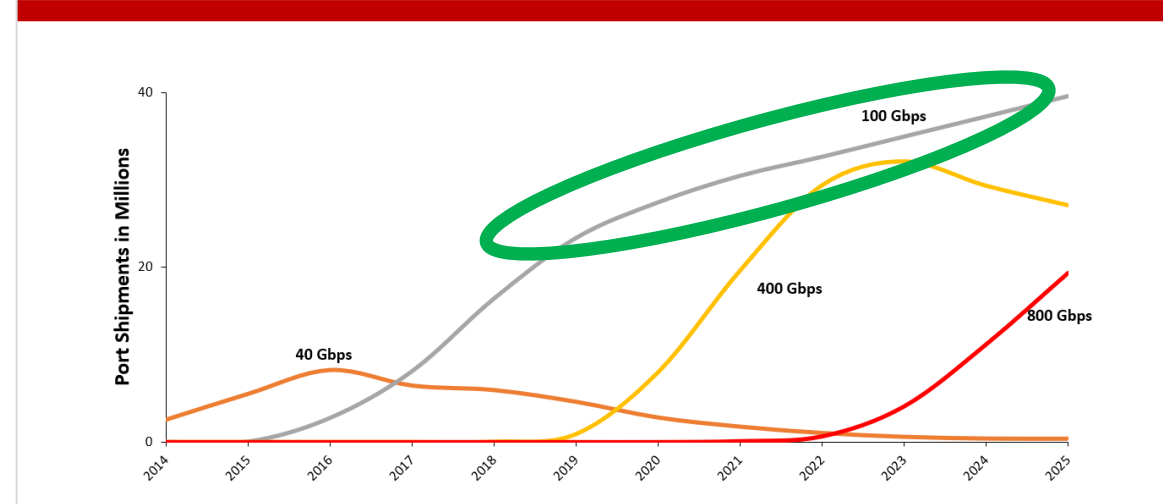
Is 100G electrical difficult?

VSR (chip to module) link

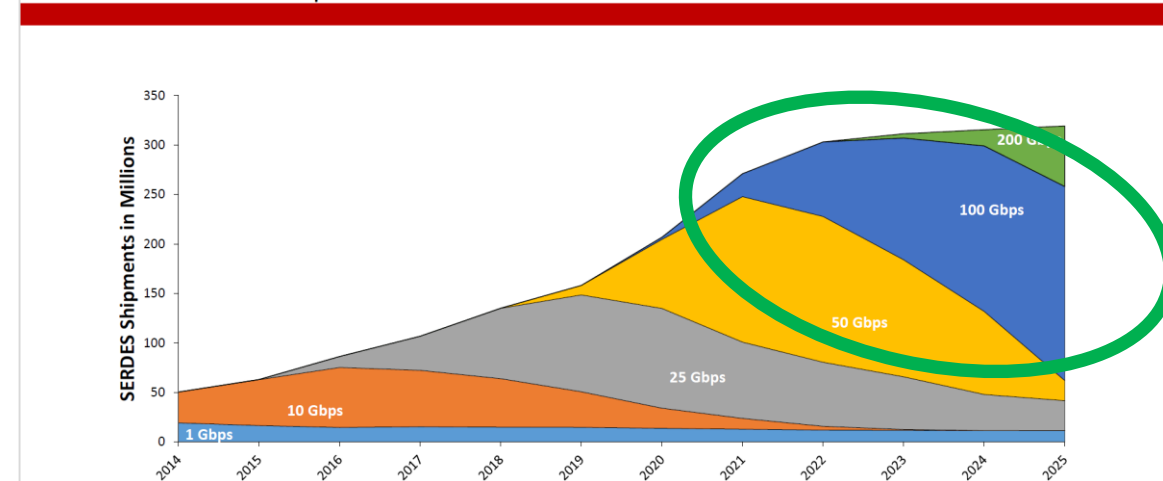


Direct Attach Copper (DAC) Cable link

Ethernet Switch – Data Center: Total Port Shipments



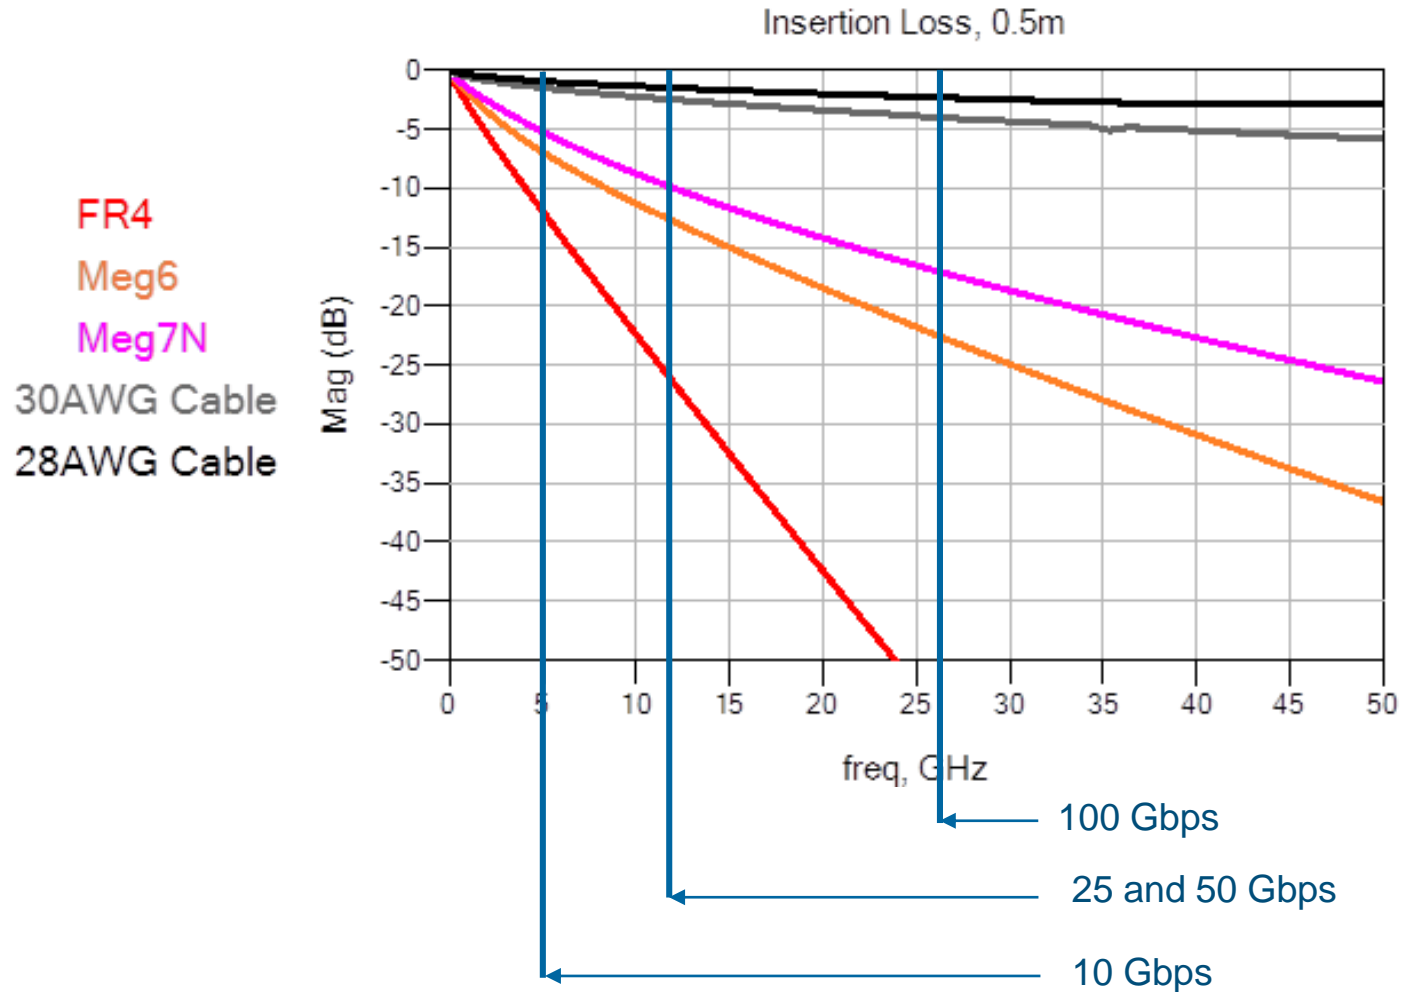
Merchant Silicon – Data Center Switching: Total SERDES Shipments



Charts used with the permission of 650 Group, LLC, Apr 2019



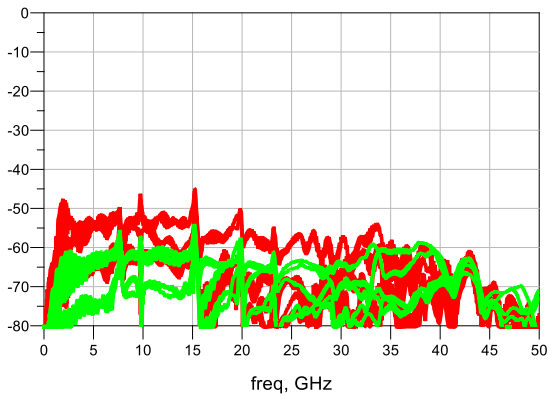
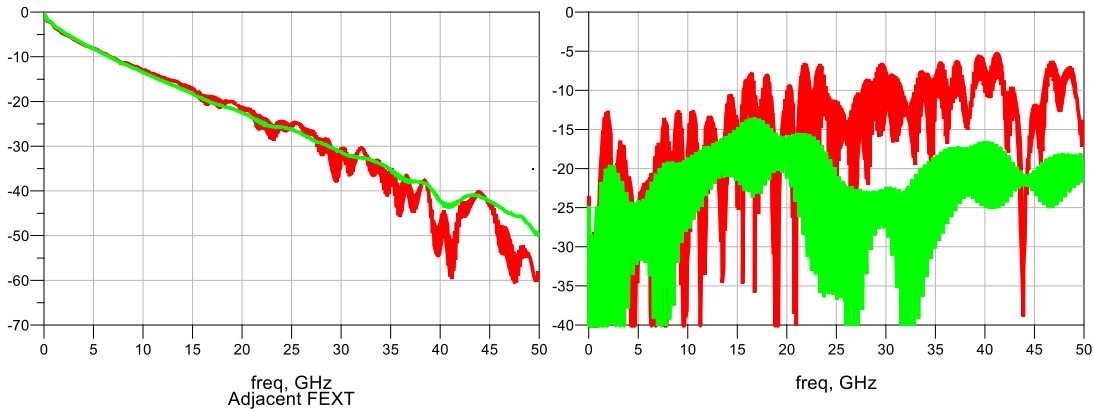
Effect Of Data Rate on Reach (considering loss only)



- Looking at loss per 0.5m length
- New materials help us at each new data, but now the materials are changing (from PCB to cable)
- Brings a new set of challenges

But at 100 Gbps, We Have To Worry About More Than Loss

Consider two 100 Gbps 28dB loss channels, one clean and one with additional impairments:

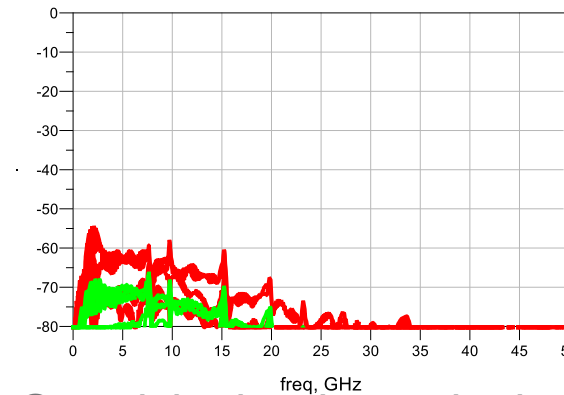
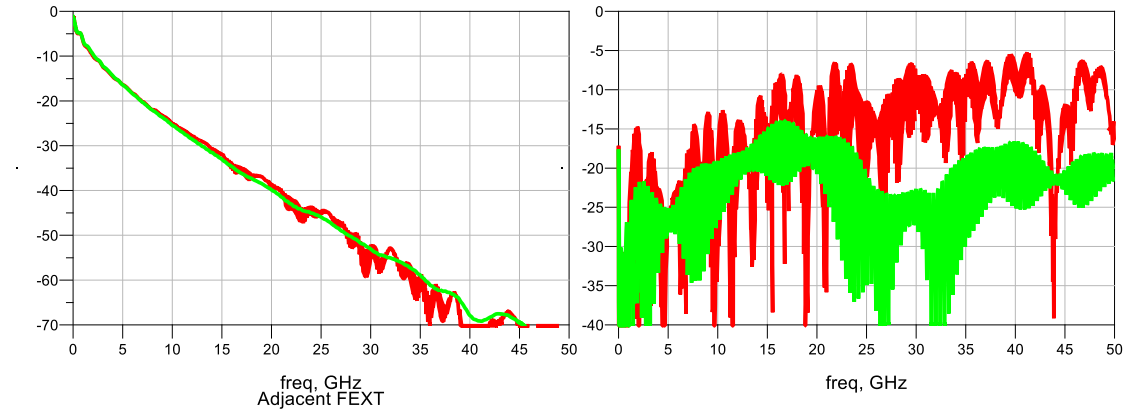


Insertion Loss	COM1	COM2
27.9 dB	5.9	3.5
27.6 dB	1.9	1.1

Not surprisingly, channel 2 fails COM

Now consider the same level of channel impairments at 25 Gbps:

We added more length to get 30dB of loss at 28 Gbps

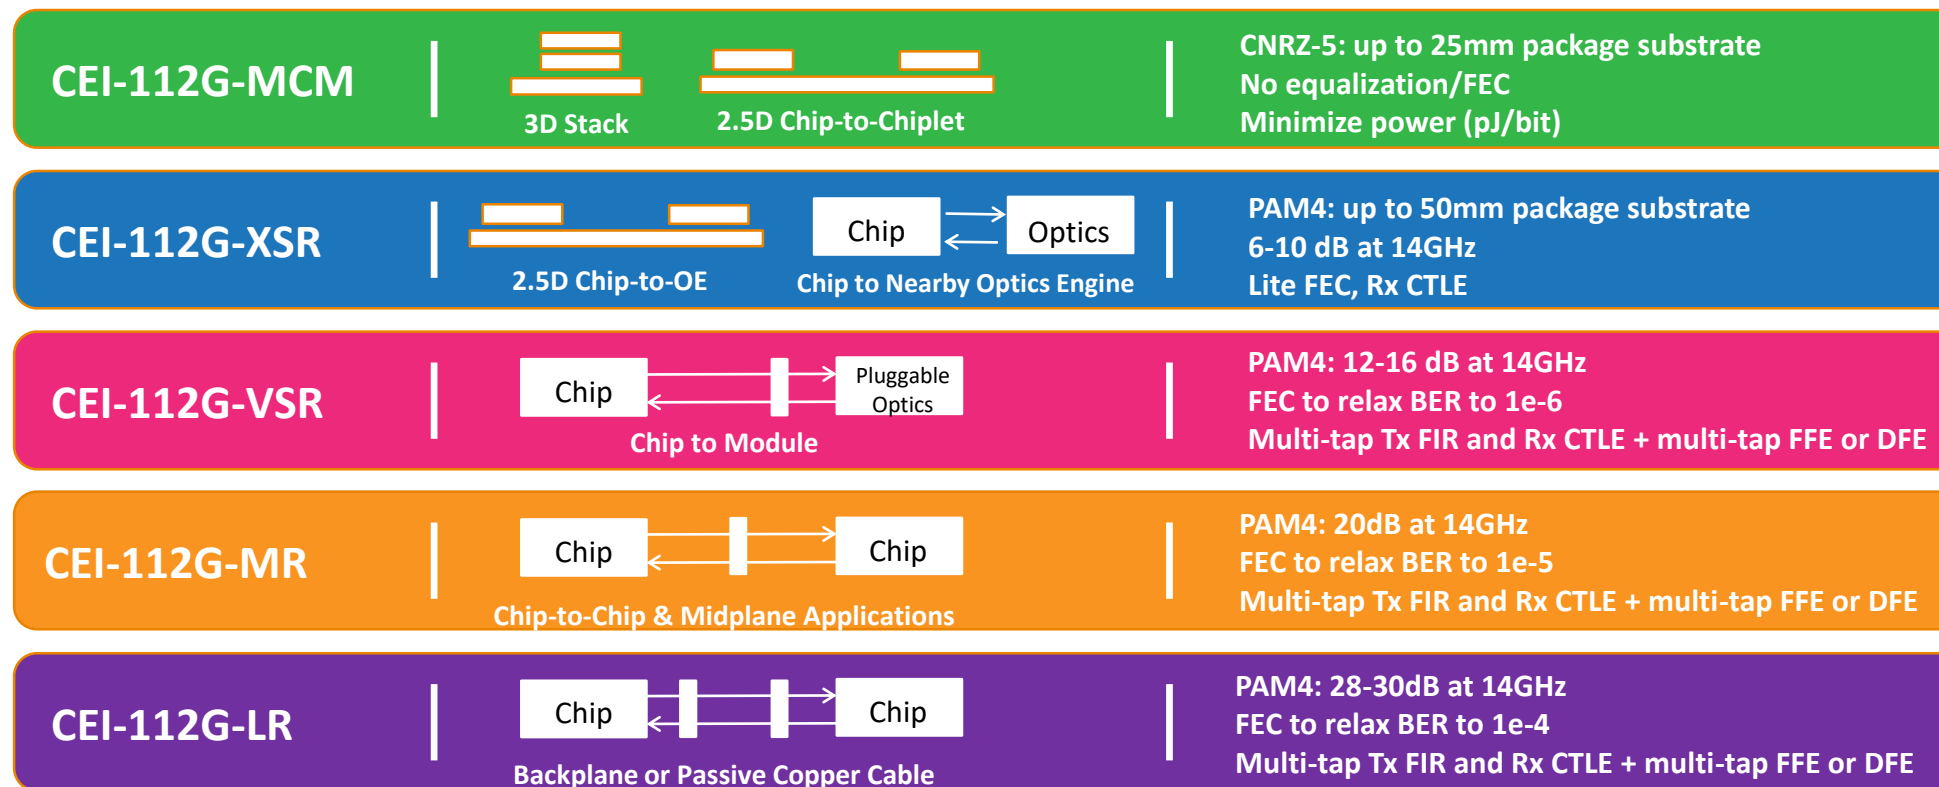


Insertion Loss	COM1	COM2
29.9 dB	6.3	5.2
29.0 dB	5.9	5.1

Surprisingly, channel 2 has good COM performance

Attention to detail on the signal integrity design of every element in the channel is going to be critical at 100 Gbps

What Electrical Channels?




PAM4 modulation scheme becomes dominant in OIF CEI-112 Gbps interface IA

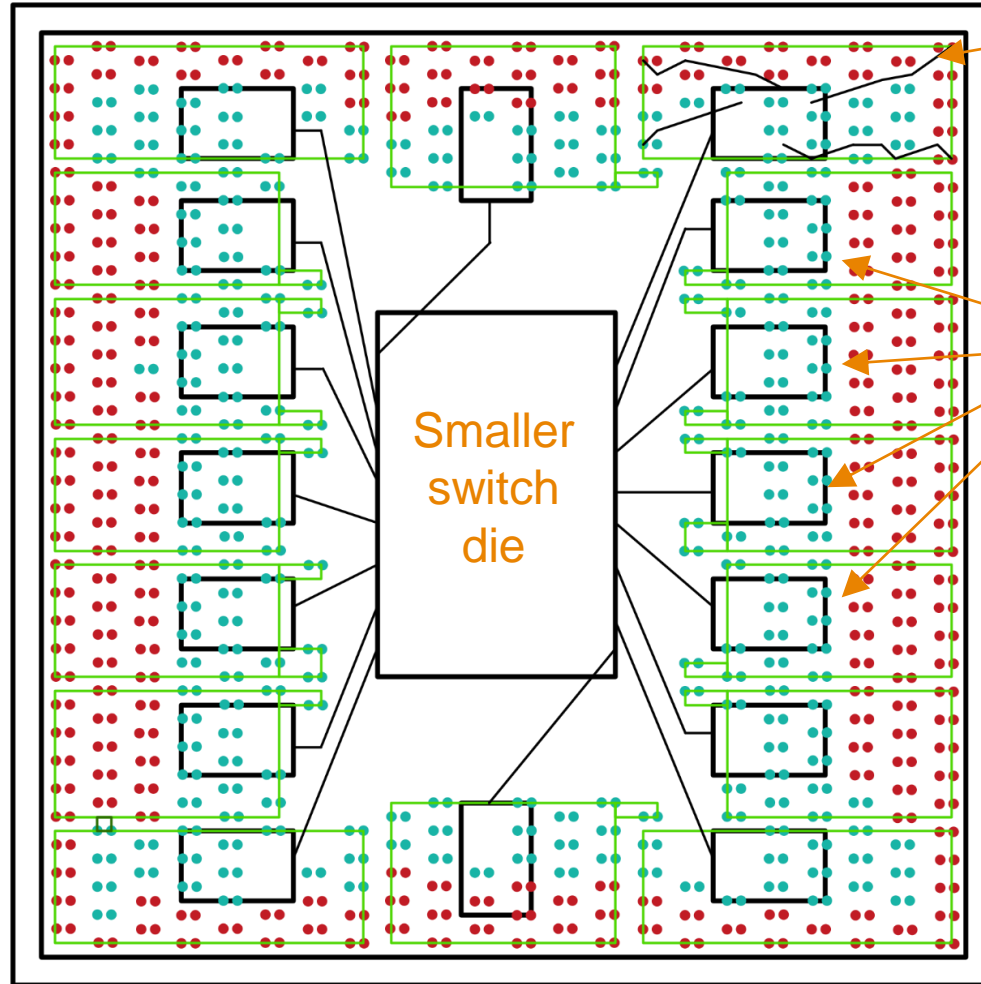
One SerDes core might not be able to cover multiple applications from XSR to LR

For short reach applications, simpler and lower power equalizations are desired

CEI-112G-MCM (Die to Die for massive bandwidth)



High loss trace to optics




Low loss trace to optics

Distributed SerDes die


Enables reduced power and longer reach

CEI-112G-XSR (Die to Die with pair orientation)

Front Plate Pluggable Optics




Co-Packaged Optics



Mount optical engines directly on switch silicon package, replace long lossy traces to pluggable optics
Reduced SerDes power for significant power savings and enable higher bandwidth density

CEI-112G-VSR Channels (Chip to Module links)




Desire to support many faceplate optical and copper cable modules connected to a high bandwidth switch chip

Number of modules determines the PCB trace length between the silicon and the modules

Strong desire to have a single SerDes that can drive optical modules and passive copper cables

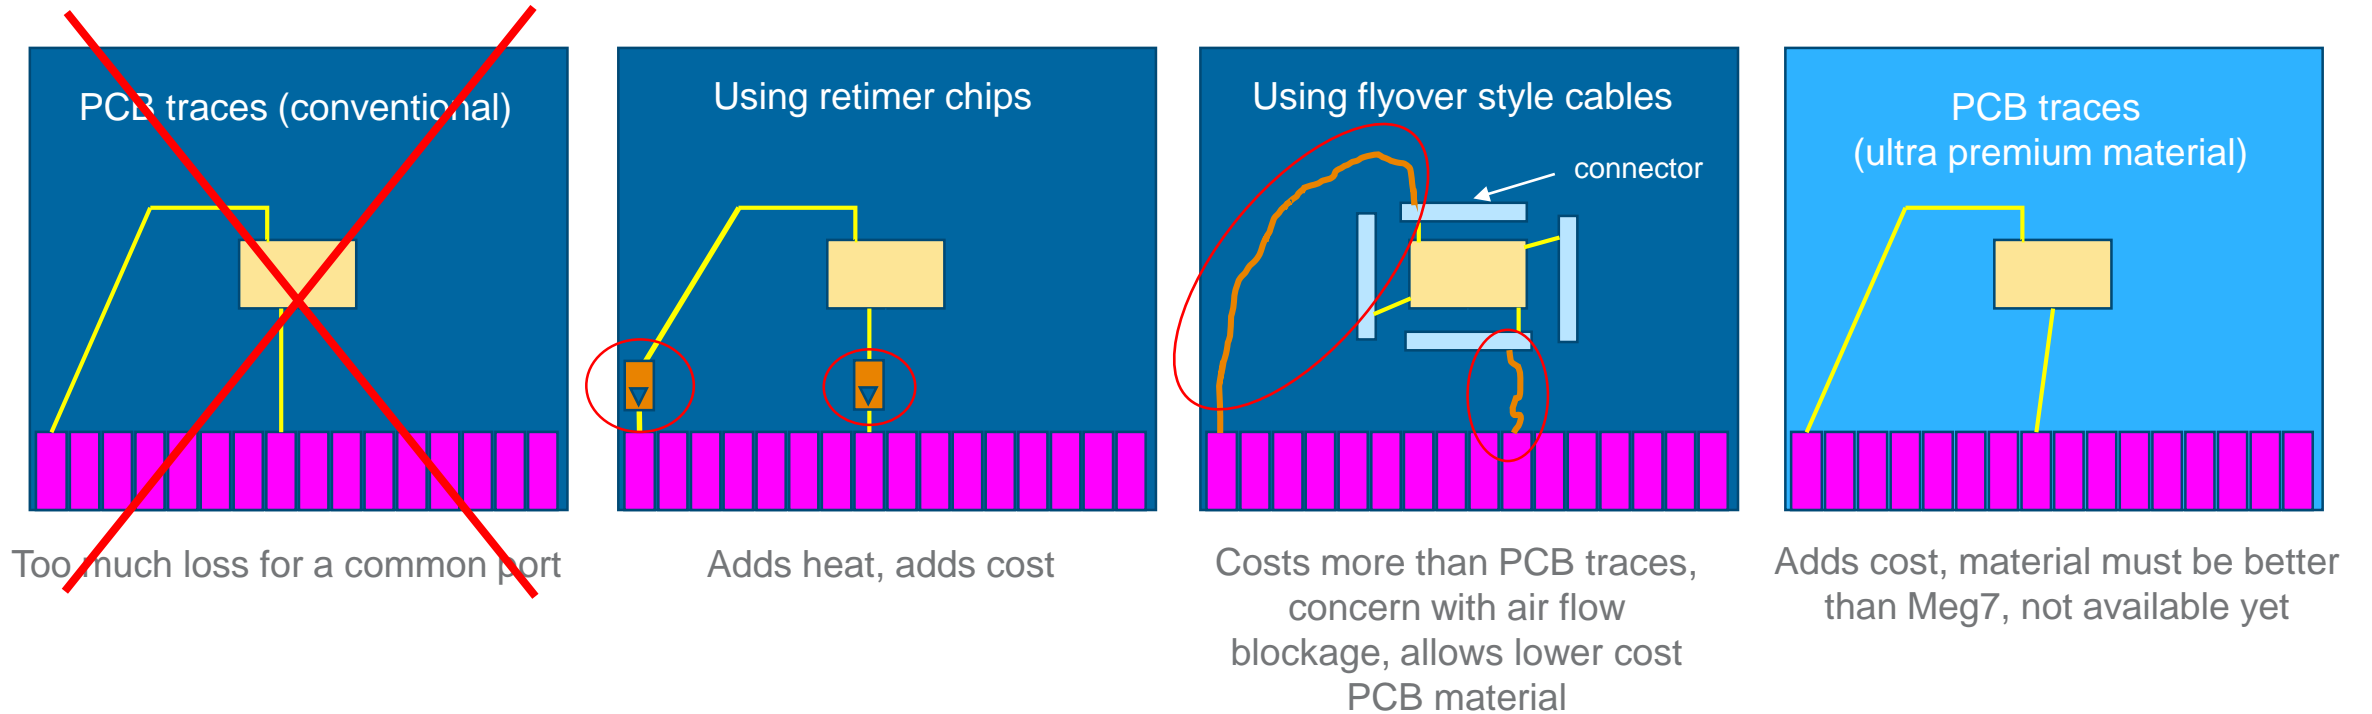
At 112Gbps, it may be impossible to have a common channel for both optics and copper cables

Loss Discussions for 100 Gbps (all numbers are “placeholders”)





- IEEE is discussing a limit of 28dB for cable assembly channels to achieve 2m length
- Loss required on the host portions of the channel for copper cables is on the order of approx. 7dB
- 7dB on the host is too low for chip to optical module links on a 32 port line card, something more like 15 dB will be required
- Therefore, it is possible that we won't be able to have a “common” port that works for both optical and copper

Host PCB Configurations to Enable a Common Port



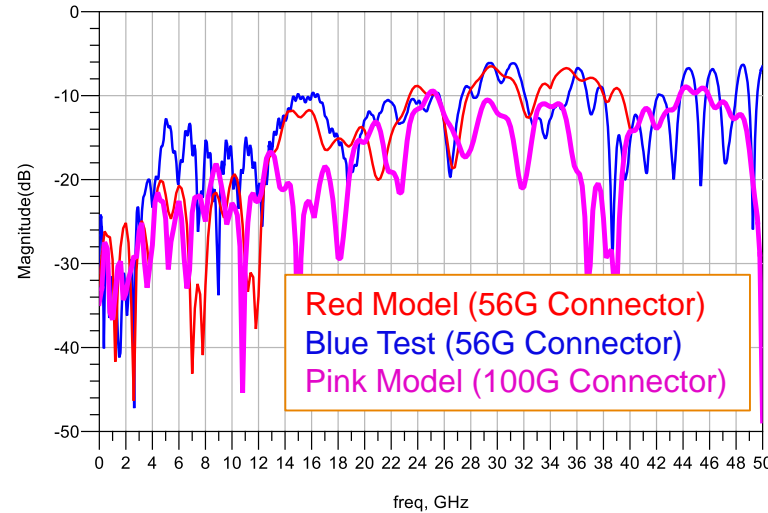
- Retimers, flyover style cables and further PCB improvements may enable equipment designs to restore ports to being “common” for optics and copper cables

VSR Channel, Modeled and Measured With 50 Gbps Connector, Modeled With 100 Gbps Connector Improvements

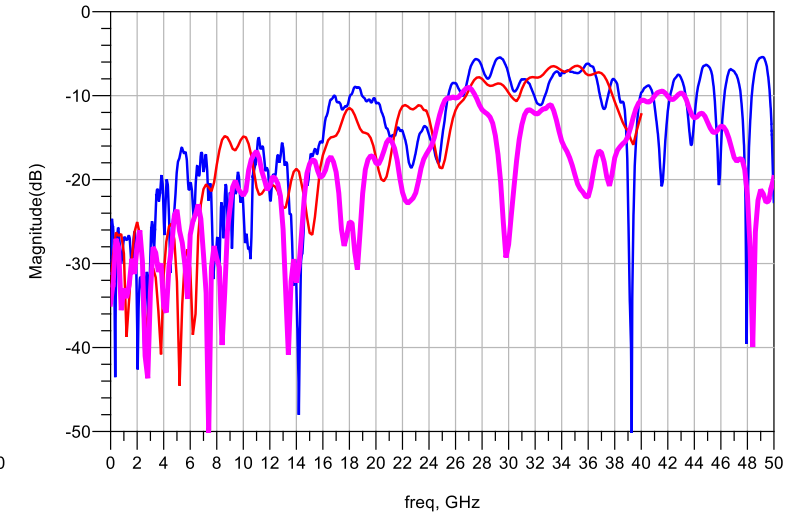


Note: Measured channel includes second set of vias to test point, modeled channels do not include the second set of vias.

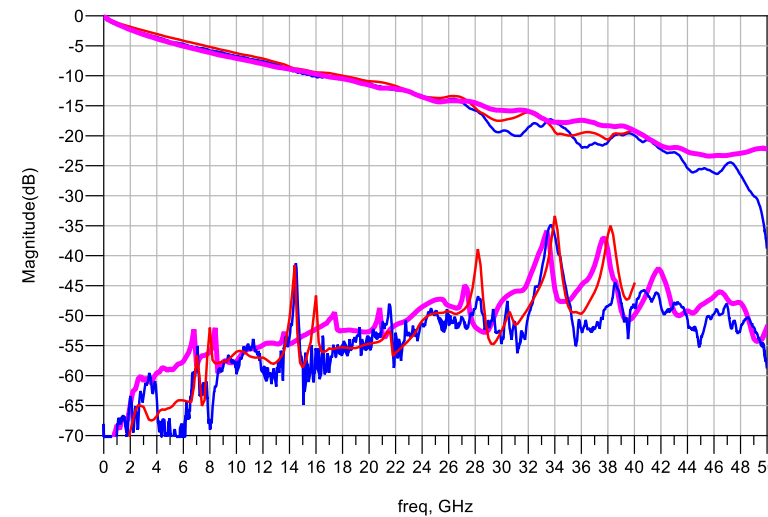
Top Row RL from Module (SDD22) - Layer 15, 12mil Vias



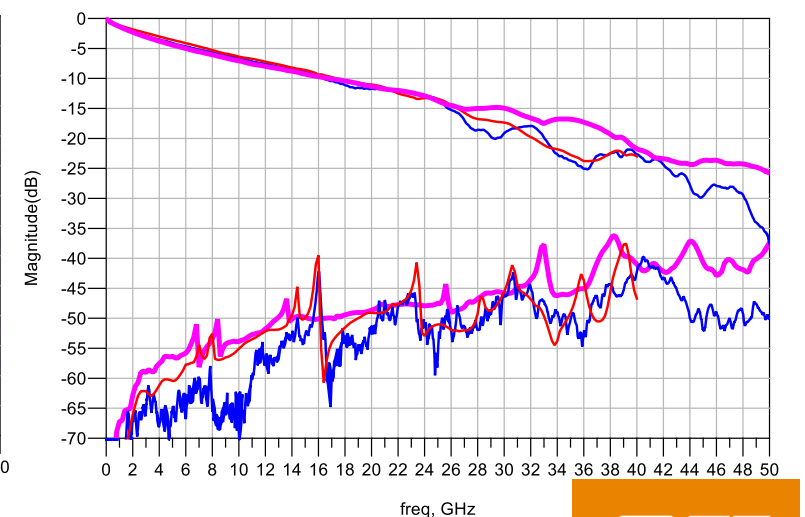
Bot Row RL from Module (SDD22) - Layer 15, 12mil Vias



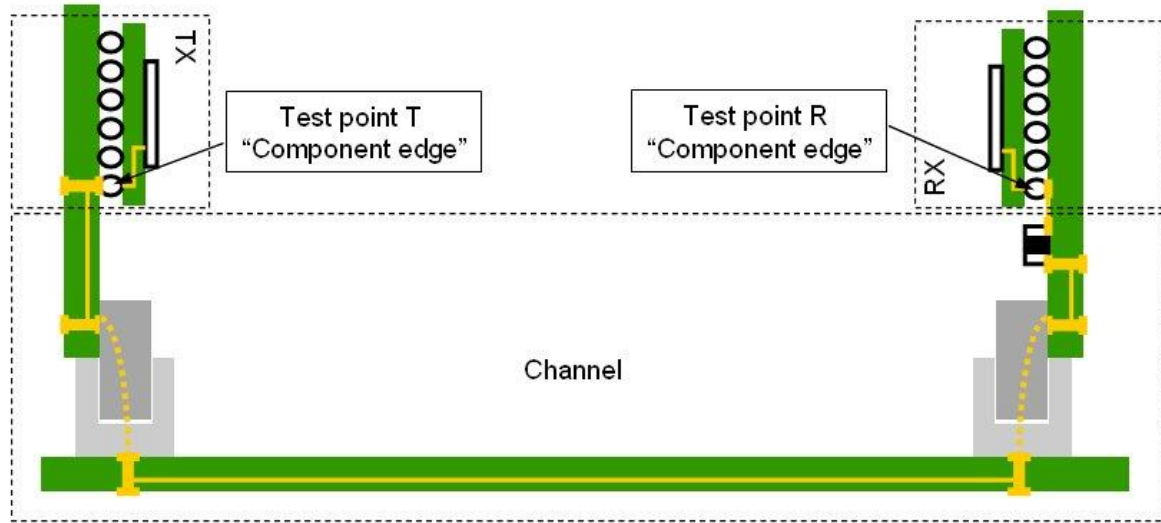
IL(Top Row) - PSFEXT (Top Row Victim) 5 Aggressors




IL(Bot Row) - PSFEXT (Bot Row Victim) 5 Aggressors




CEI-112G-LR Channels (Backplane or “line card to line card” channels)




Long Reach Channel
(Backplane)



Conventional Backplane

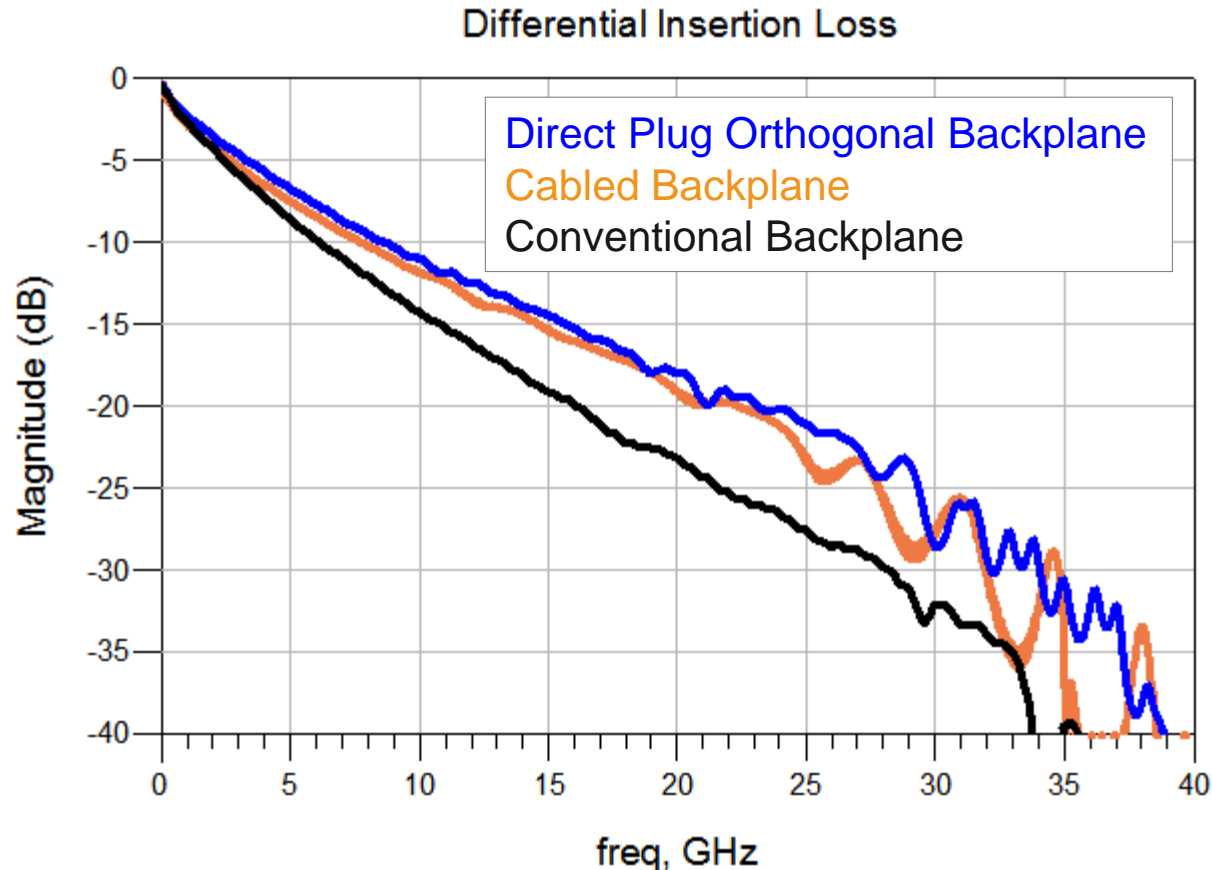


Cabled Backplane




Orthogonal Backplane

Comparison of Insertion Loss for Backplane Applications




Direct Plug Orthogonal Backplane:

Two 9 inch daughter cards; 18 inches total reach




Cabled Backplane:

Two 6 inch daughter cards plus 1m 30AWG cable; 52 inches total reach




Conventional Backplane:

Two 6 inch daughter cards plus 16 inch backplane; 24 inches total reach



Conventional backplanes create challenging channels due to high loss, thick backplanes result in noise from plated through-holes and via stubs so they are not anticipated to be highly deployed in 100 Gbps applications


Orthogonal Backplane Channel




- 18" PCB Trace Total
 - 9" Trace per board
 - 6/6/6 trace geometry
 - Meg7N Laminates
 - HVLP Foils
- 140mil (3.56mm) Thick PCBs
 - Victim pair uses layer 2 routing
 - Victim pair: 15mil Stub w/ Shallow EON Technology
 - Aggressor Pairs are thru board to bottom layer
- Next-Gen STRADA Whisper Connector Model
 - Direct-Plug Orthogonal
 - Stub resonance has been addressed
 - Additional noise control features

Orthogonal Backplane Channel Results


TX/RX 8-Agressor PowerSum Crosstalk



Differential Insertion Loss




Orthogonal Backplane Channel Crosstalk



- Pair G11/12 is the central victim pair.
- Near-End and Far-End Crosstalk available in a typical TX/RX Pattern
- 0-60GHz in 10MHz steps


Cabled Backplane Channel




- 12" PCB Trace Total
 - 6" Trace per Board
 - 6/6/6 Geometry
 - Meg7N Laminates
 - HVLP Foils
- 140mil (3.56mm) Thick Footprints
 - Victim pair uses layer 2 routing
 - Victim pair: 15mil stub w/ shallow EON technology
 - Aggressor Pairs are thru board to bottom layer
- Next-Gen STRADA Whisper Connector Model
 - Cabled header to R/A receptacle
 - Additional noise control features
 - Stub resonance addressed
- 1m Cable Length
 - 30AWG TurboTwin twinax cable

Cabled Backplane Channel Results

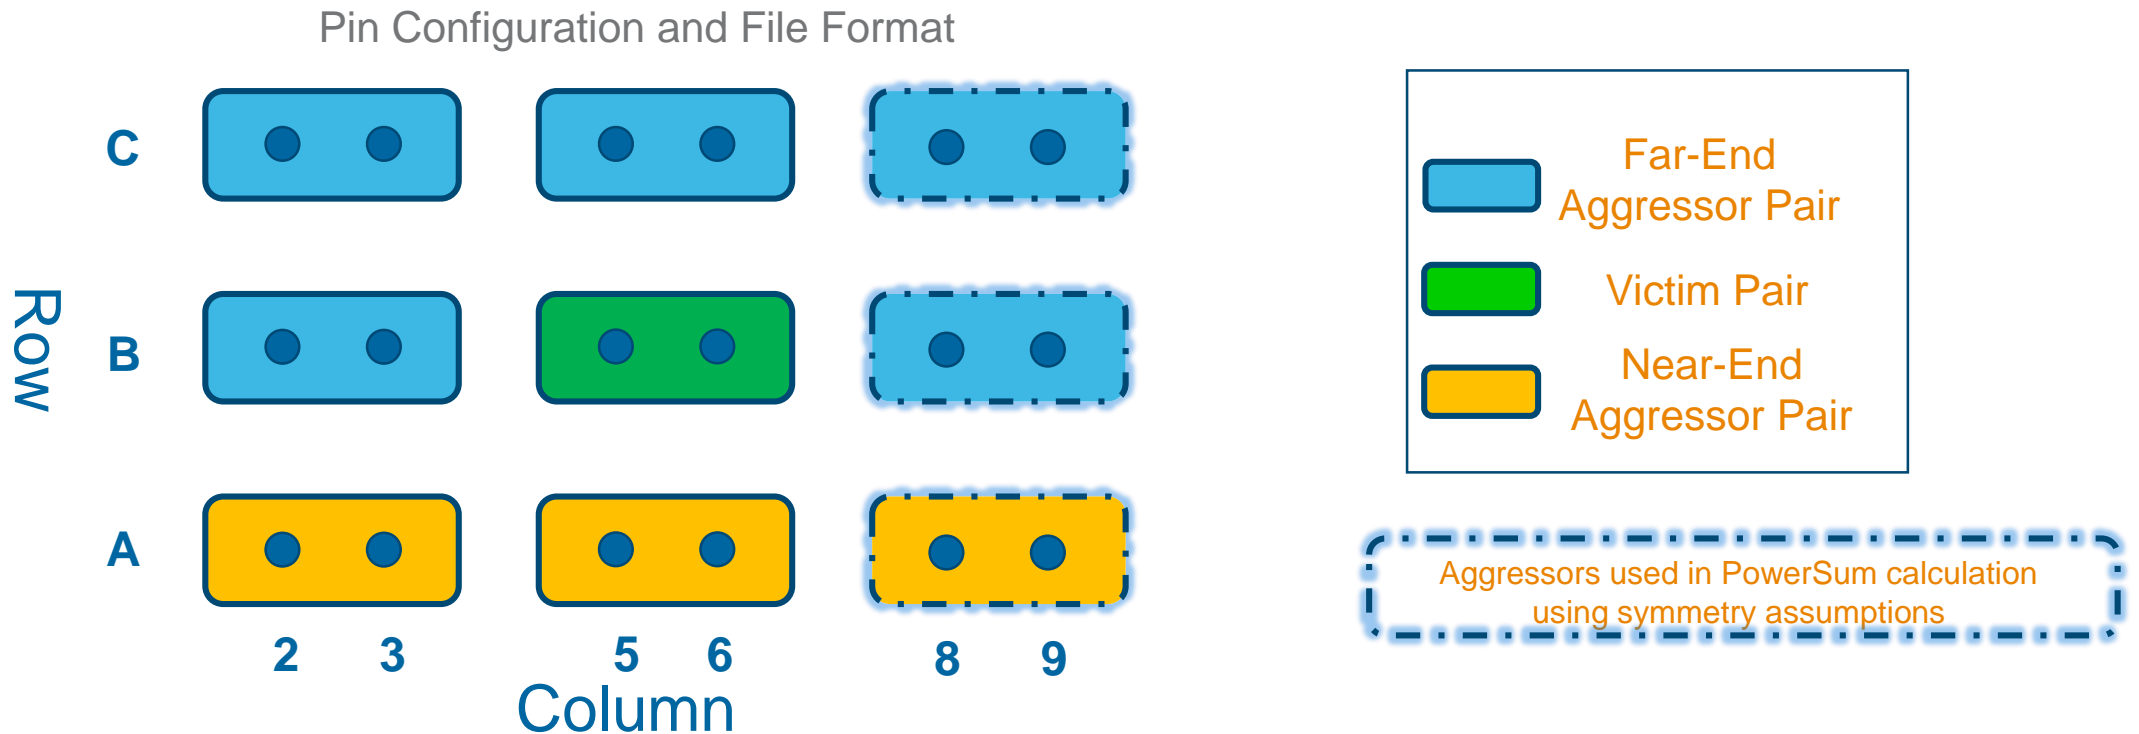
TX/RX 8-Aggressor PowerSum Crosstalk



Differential Insertion Loss



Cabled Backplane Channel Crosstalk




- Pair B5/6 is the central victim pair. Near-End and Far-End Crosstalk available in a typical TX/RX Pattern
- Test vehicle has 6 pairs, 3 more aggressors are added by symmetry
- 0-60GHz in 10MHz steps

Case Study: Considering Impedance, Skew and Reflections


Looking at the 100 Gbps electrical interface for three popular form factors for pluggable optical modules and copper cables currently used on channels with 50Gbps PAM4 electrical signaling:

- QSFP: 4 channels
- OSFP: 8 channels
- QSFP-DD: 8 channels




Mating Zone Impedance (Mating Zone Reflections)



a.) QSFP28



b.) QSFP-DD



c.) OSFP




- OSFP, lowest pad width-to-pitch ratio = 0.63
- 2D field solver used to calculate characteristic impedance of **ONLY** the mating pads acting as microstrip traces
- 10% reduction in pad width b/w QSFP28 & QSFP-DD
 - Yields 10 ohm improvement

Form Factor	Impedance (Ohms)
QSFP28	82
QSFP-DD	92
OSFP	101


Mating Zone Impedance (Mating Zone Reflections)

Form Factor	Modeled Z_o without Rcpt (Ω)	Measured Z_{diff} with Rcpt (Ω)	Δ
QSFP28	82	72	10
QSFP-DD	92	77	15
OSFP	101	86	15

- Largest impact on mating interface impedance comes from contact lead-in and PCB pad stub from contact wipe


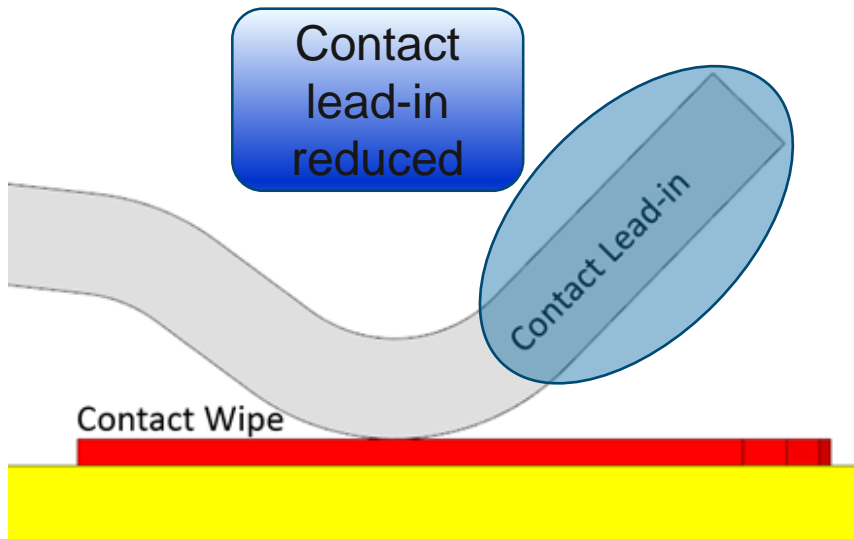


Measured results of mating zone (Plug-in card mated to Receptacle)

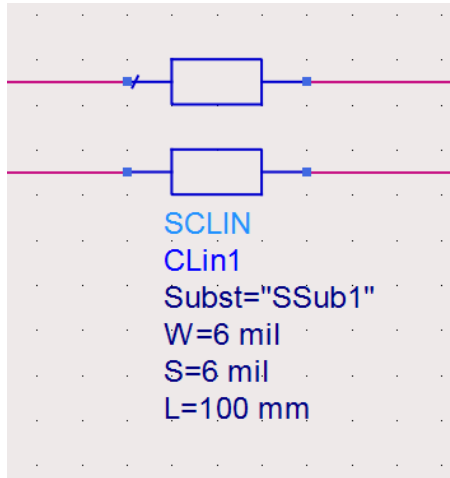


Mating Zone Impedance (Mating Zone Reflections)



- 56 Gbps OSFP connector modified by reducing contact lead-in (referred to as 112 Gbps connector moving forward)
- 7 ohm improvement observed when contact lead-in is reduced
- Following slides show impact of improved mating zone for 112 Gbps channels



Skew Impact on S_{dd21} & S_{cd21}




- Skew inserted as phase delay
- Higher frequencies impacted more than lower frequencies as skew increases
- Skew impacts 112 Gbps channels more than 56 Gbps channels



Skew

- 0 ps (Pink)
- 3 ps (Red)
- 6 ps (Blue)
- 9 ps (Black)

Impact of Skew & Reflections, OSFP 112 Gbps Copper Cable, Test Set-up




- Victim
- FEXT Aggressor
- NEXT Aggressor


OSFP Pin Map

Pin #	60	59	58	57	56	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32	31
	G	Tx1+	Tx1-	G	Tx3+	Tx3-	G	Tx5+	Tx5-	G	Tx7+	Tx7-	G	SB	SB	SB	SB	G	Rx8-	Rx8+	G	Rx6-	Rx6+	G	Rx4-	Rx4+	G	Rx2-	Rx2+	G
	G	Tx2+	Tx2-	G	Tx4+	Tx4-	G	Tx6+	Tx6-	G	Tx8+	Tx8-	G	SB	SB	SB	SB	G	Rx7-	Rxy+	G	Rx5-	Rx5+	G	Rx3-	Rx3+	G	Rx1-	Rx1+	G
Pin #	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

Pin #	60	59	58	57	56	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32	31
	G	Tx1+	Tx1-	G	Tx3+	Tx3-	G	Tx5+	Tx5-	G	Tx7+	Tx7-	G	SB	SB	SB	SB	G	Rx8-	Rx8+	G	Rx6-	Rx6+	G	Rx4-	Rx4+	G	Rx2-	Rx2+	G
	G	Tx2+	Tx2-	G	Tx4+	Tx4-	G	Tx6+	Tx6-	G	Tx8+	Tx8-	G	SB	SB	SB	SB	G	Rx7-	Rxy+	G	Rx5-	Rx5+	G	Rx3-	Rx3+	G	Rx1-	Rx1+	G
Pin #	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



Impact of Reflections, OSFP 112 Gbps Copper Cable, Measurement Results

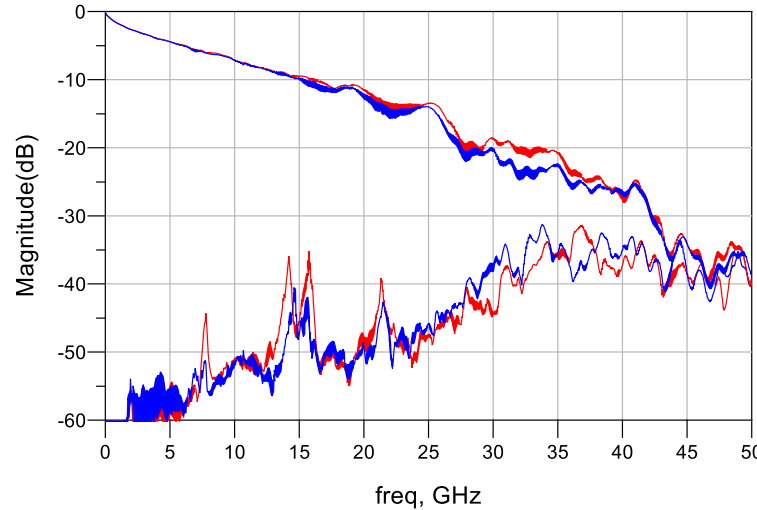


Results at 26.56 GHz


	Conn	S _{dd21}	PSXT
1.0 m 30 AWG	112	-16.38	-45.06
1.0 m 30 AWG	56	-18.26	-44.00
1.5 m 28 AWG	112	-17.73	-42.22
1.5 m 28 AWG	56	-19.41	-41.31

- Measurement results shown for
 - 1.0m 30 AWG Tx8 lane
 - 1.5m 28 AWG Tx8 lane
- 112 Gbps connector improves differential insertion loss at higher frequencies
- Test fixture skew included in measurement results


Sdd21/PSXT - 1.0 m 30 AWG Tx8




Sdd21/PSXT - 1.5 m 28 AWG Tx8



Skew vs Frequency - 1.0 m 30 AWG



Skew vs Frequency - 1.5 m 28 AWG Tx8





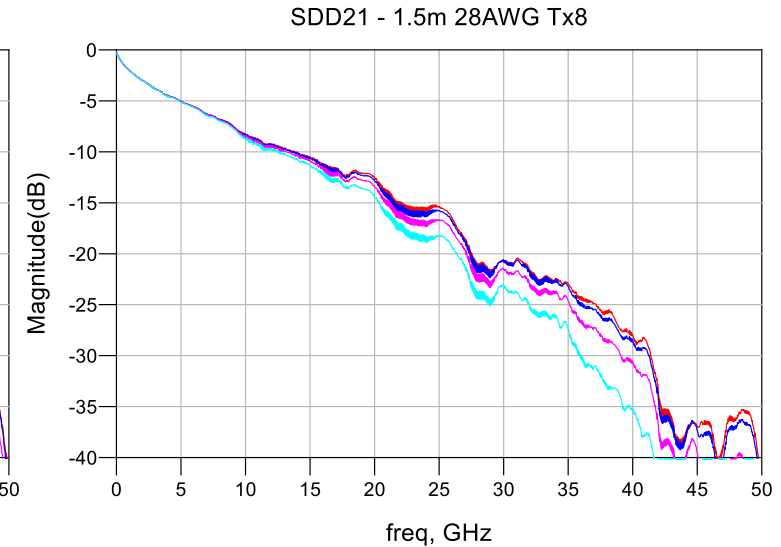
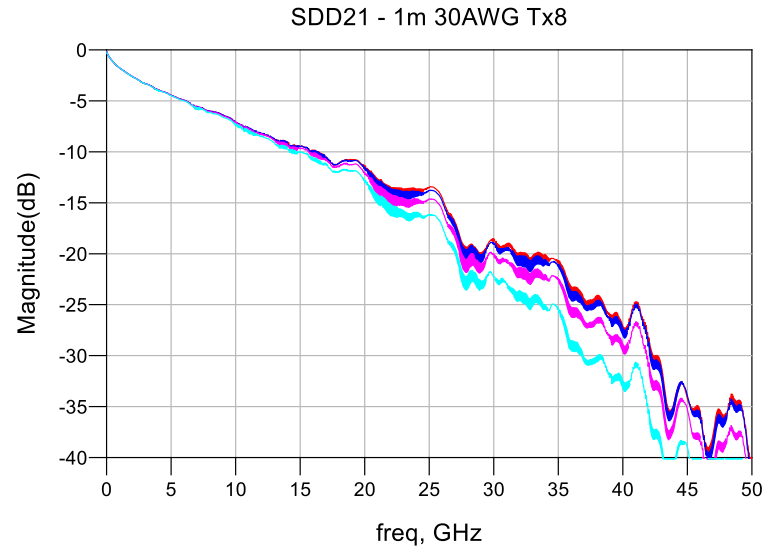
Measured Channel w/ 56 Gbps Conn (Blue)
 Measured Channel w/ 112 Gbps Conn (Red)



Impact of Skew, OSFP 112 Gbps Copper Cable

	Added Skew	S_{dd21} (dB)	S_{cd21} (dB)
1.0 m 30 AWG	0 ps	-16.38	-49.10
1.0 m 30 AWG	3 ps	-16.66	-28.39
1.0 m 30 AWG	6 ps	-17.53	-22.71
1.0 m 30 AWG	9 ps	-19.12	-19.68
1.5 m 28 AWG	0 ps	-17.73	-37.08
1.5 m 28 AWG	3 ps	-17.98	-29.36
1.5 m 28 AWG	6 ps	-18.84	-23.99
1.5 m 28 AWG	9 ps	-20.40	-21.00

- Skew added using Keysight ADS tool
- Added skew degrades S_{dd21} and S_{cd21}
- Test fixture skew included in measurement results




0ps Added Skew (Red), 3ps Added Skew (Blue)
 6ps Added Skew (Pink) 9ps Added Skew (Light Blue)





Impact of Skew & Reflections, OSFP 112 Gbps Copper Cable, COM Calculation

- COM calculated using Version 2.51 COM script developed for IEEE 802.3ck specification
- Results for 32 lanes (Tx7 & Tx8 lanes/30 & 28 AWG/ 56 & 112 Gbps Connector/ 0, 3, 6, and 9 ps added skew)
- Failures for COM Case 2, 56 Gbps conn w/ 9 ps of skew & 2 instances of 112 Gbps conn w/ 9 ps of skew
- Channels with increased differential insertion loss exhibit lower COM value
- Minimizing mating zone reflections made channels more tolerant of skew




Impact of Reflections, OSFP 112G VSR (Chip-2-Module) Channel



- 8.5" of 5.3-6.45-5.3 diff traces on host
- 1.5" of 6.3mil SE traces on module
- MEGTRON 7N HVLP material used
- Host uses 12mil differential vias down to a Layer 15 route-out

Results at 26.56 GHz

Conn	S _{dd21}	PSXT
112	-15.61	-47.74
56	-16.51	-48.90



OSFP Pin Map

Pin #	60	59	58	57	56	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32	31
	G	Tx1+	Tx1-	G	Tx3+	Tx3-	G	Tx5+	Tx5-	G	Tx7+	Tx7-	G	SB	SB	SB	SB	G	Rx8-	Rx8+	G	Rx6-	Rx6+	G	Rx4-	Rx4+	G	Rx2-	Rx2+	G
	G	Tx2+	Tx2-	G	Tx4+	Tx4-	G	Tx6+	Tx6-	G	Tx8+	Tx8-	G	SB	SB	SB	SB	G	Rx7-	Rx7+	G	Rx5-	Rx5+	G	Rx3-	Rx3+	G	Rx1-	Rx1+	G
Pin #	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


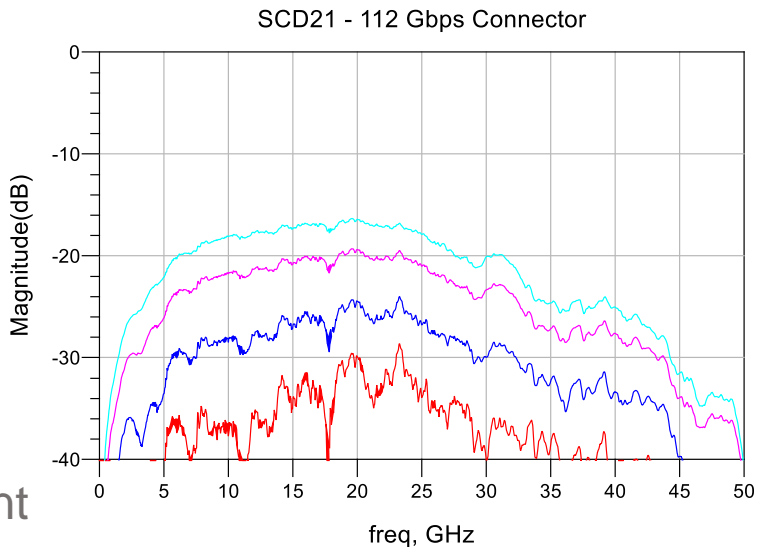


Victim
 FEXT Aggressor



Impact of Skew, 112G VSR (Chip-2-Module) Channel

	Added Skew	S_{dd21} (dB)	S_{cd21} (dB)
112 Gbps Conn	0 ps	-15.61	-34.74
112 Gbps Conn	3 ps	-15.84	-27.64
112 Gbps Conn	6 ps	-16.70	-21.92
112 Gbps Conn	9 ps	-18.29	-18.87
56 Gbps Conn	0 ps	-16.51	-41.47
56 Gbps Conn	3 ps	-16.78	-28.60
56 Gbps Conn	6 ps	-17.64	-22.87
56 Gbps Conn	9 ps	-19.22	-19.82

- Skew added using Keysight ADS tool
- Added skew degrades S_{dd21} and S_{cd21}
- Test fixture skew included in measurement results

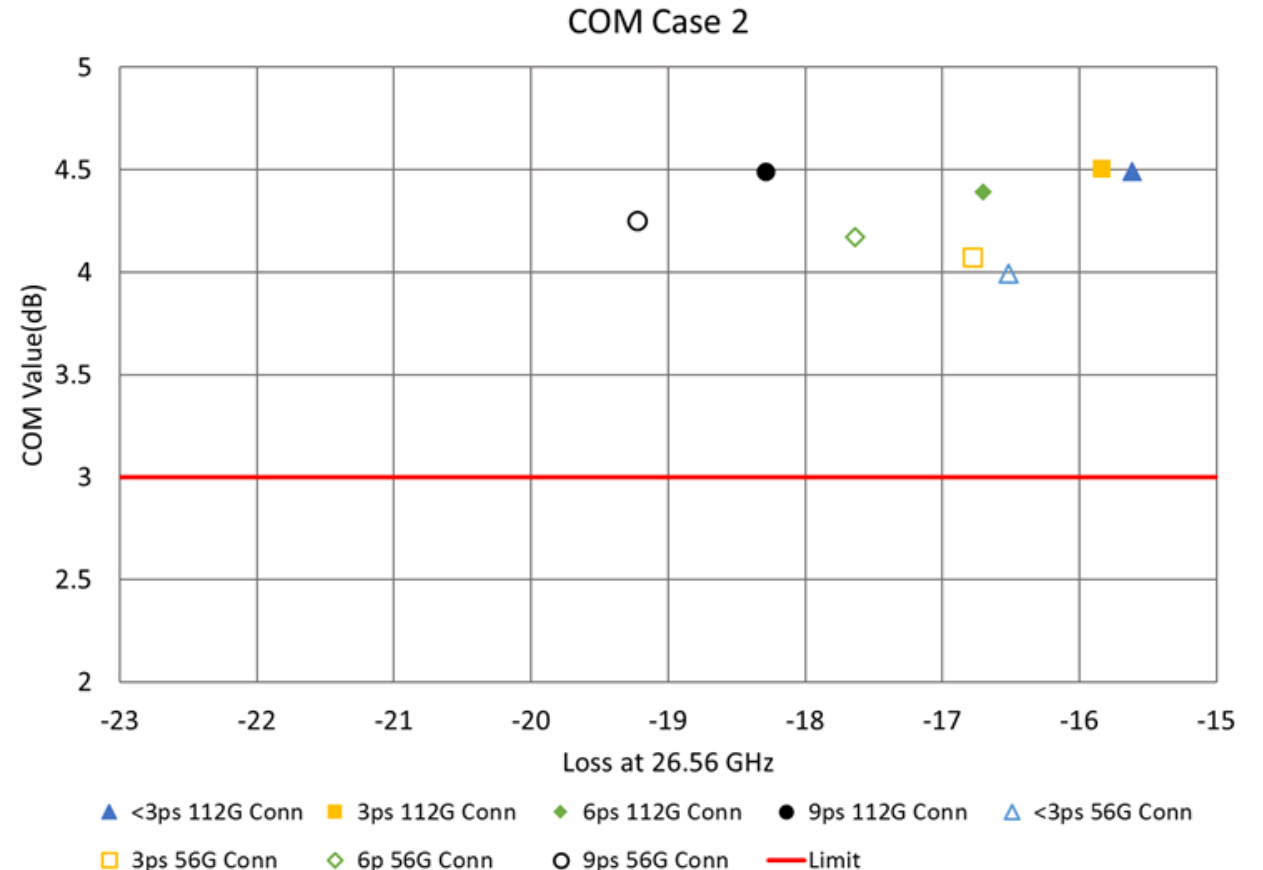


0ps Added Skew (Red), 3ps Added Skew (Blue)
 6ps Added Skew (Pink) 9ps Added Skew (Light Blue)



Impact of Skew & Reflections, 112G VSR (Chip-2-Module), COM Calculation

- COM calculated using Version 2.41 COM script developed for IEEE 802.3ck specification
- Results for 8 lanes (56 & 112 Gbps Connector/ 0, 3, 6, and 9 ps added skew)
- All instances pass COM Case 2
- Unlike copper cable channel, C2M does not exhibit differential insertion loss sensitivity



Conclusions

- As we look to 112 Gbps performance, reach is a critical factor
- Architectural changes can enable reach, including cables replacing PCB traces and orthogonal chassis
- Two of the main drivers in extending reach are skew and mating zone reflections
- Skew, differential insertion loss, and mode conversion are interrelated
- If mating zone reflections can be minimized, a channel can be more tolerant of skew from a differential insertion loss standpoint
- Minimizing both skew and mating zone reflections are strong drivers in being able to maximize channel reach for 112G.
- Focusing on the SI design details of every circuit element is critical at 100 Gbps.

Acknowledgements: Thank you to the team at TE who contributed significant work to these slides including: Bruce Champion, Justin Pickel, Linda Shields, Megha Shanbhag and many others