112 Gbps Electrical Interfaces – An OIF Update on CEI-112G

Panel Session

Nathan Tracy: OIF President, TE Connectivity
Gary Nicholl: OIF Board Member, Cisco Systems
Cathy Liu: OIF Board Member, Broadcom
Mike Li: OIF Board Member, Intel
Ed Frlan, OIF Technical Committee Chair, Semtech
Steve Sekel: OIF Interop Chair, Keysight Technologies

OFC
IEEE ComSoc IEEE Communications Society
IEEE Photonics Society
OSA The Optical Society | 100 Since 1916

SPEAKERS

Nathan Tracy

OIF President/Board Member,
TE Connectivity Technologist
ntracy@te.com

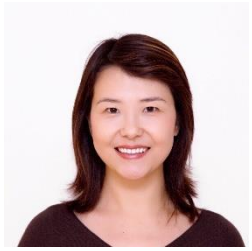

Cathy Liu

OIF Board Member,
Broadcom Inc. SerDes Architect
cathy.liu@broadcom.com

Mike Li

OIF Board Member,
Intel Fellow
peng.mike.li@intel.com

Gary Nicholl

OIF Board Member,
Cisco Principal Engineer
gnicholl@cisco.com

Ed Frlan

OIF Technical Committee Chair,
Semtech Senior System Architect
efrlan@semtech.com

Steve Sekel

OIF PLL Interop WG chair,
Keysight Technologies 400G Solutions Specialist
steve.sekel@keysight.com

What is the OIF?

- Since 1998, OIF has brought together industry groups from the data and optical worlds
- Mission: To foster the development and deployment of interoperable products and services for data switching and routing using optical networking technologies
- Our 100+ member companies represent the entire industry ecosystem:
  - Network operators and network users
  - Component and systems vendors
  - Testing and software companies

OIF CEI (Common Electrical IO) Electrical Implementation Agreements

The Common Electrical IO (CEI) Implementation Agreement (IA) is a clause-based format supporting publication of new clauses over time:

- CEI-1.0: included CEI-6G-SR, CEI-6G-LR, and CEI-11G-SR.
- CEI-2.0: added CEI-11G-LR
- CEI-3.0: added CEI-25G-LR, CEI-28G-SR
- CEI-3.1: added CEI-28G-MR and CEI-28G-VSR
- CEI-4.0: added CEI-56G-USR-NRZ, CEI-56G-XSR-NRZ, CEI-56G-VSR-PAM4, CEI-56G-MR-PAM4, CEI-56G-LR-PAM4, and CEI-56G-LR-ENRZ.

More info at https://www.oiforum.com/technical-work/hot-topics/common-electrical-interface-cei-112g-2/

Existing document available at https://www.oiforum.com/wp-content/uploads/2019/01/OIF-CEI-04.0.pdf

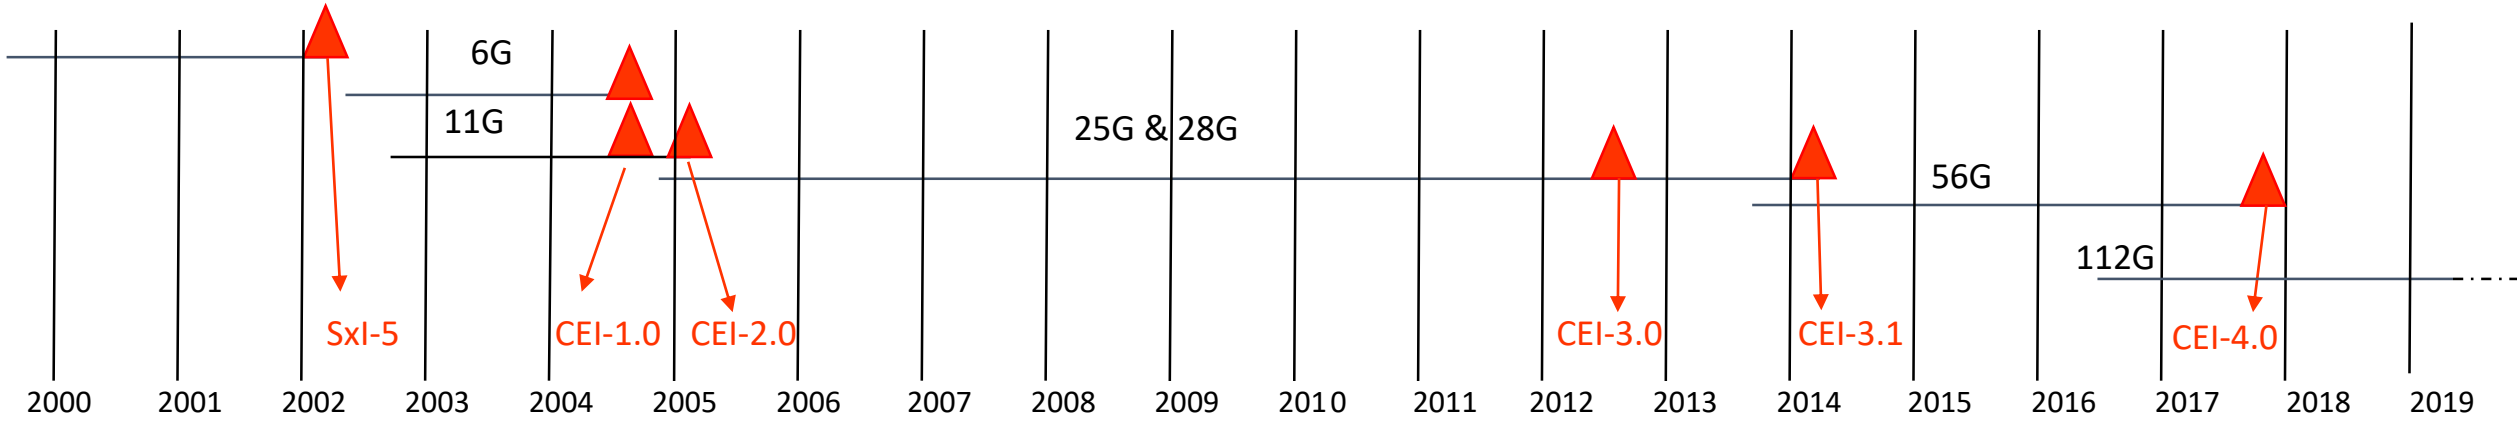

OIF's CEI Work Has Been a Significant Industry Contributor

Name	Rate per pair	Year	Activities that Adopted, Adapted or were influenced by the OIF CEI
CEI-112G	112Gbps	2021 (projected)	Five channel reach projects in progress, IEEE, InfiniBand, T11 (Fibre Channel), Interlaken, ITU.
CEI-56G	56Gbps	2017	IEEE, InfiniBand, T11 (Fibre Channel), Interlaken, ITU
CEI-28G	28 Gbps	2012	InfiniBand EDR, 32GFC, SATA 3.2, SAS-4,100GBASE-KR4, CR4, CAUI4, Interlaken, ITU
CEI-11G	11 Gbps	2008	InfiniBand QDR, 10GBASE-KR, 10GFC, 16GFC, SAS-3, RapidIO v3, Interlaken, ITU
CEI-6G	6 Gbps	2004	4GFC, 8GFC, InfiniBand DDR, SATA 3.0, SAS-2, RapidIO v2, HyperTransport 3.1, Interlaken, ITU
SxI5	3.125 Gbps	2002-3	Interlaken, FC 2G, InfiniBand SDR, XAUI, 10GBASE-KX4, 10GBASE-CX4, SATA 2.0, SAS-1, RapidIO v1, ITU
SPI4, SFI4	1.6 Gbps	2001-2	SPI-4.2, HyperTransport 1.03
SPI3, SFI3	0.800 Gbps	2000	(from PL3)

OIF CEI-112G Development Application Space

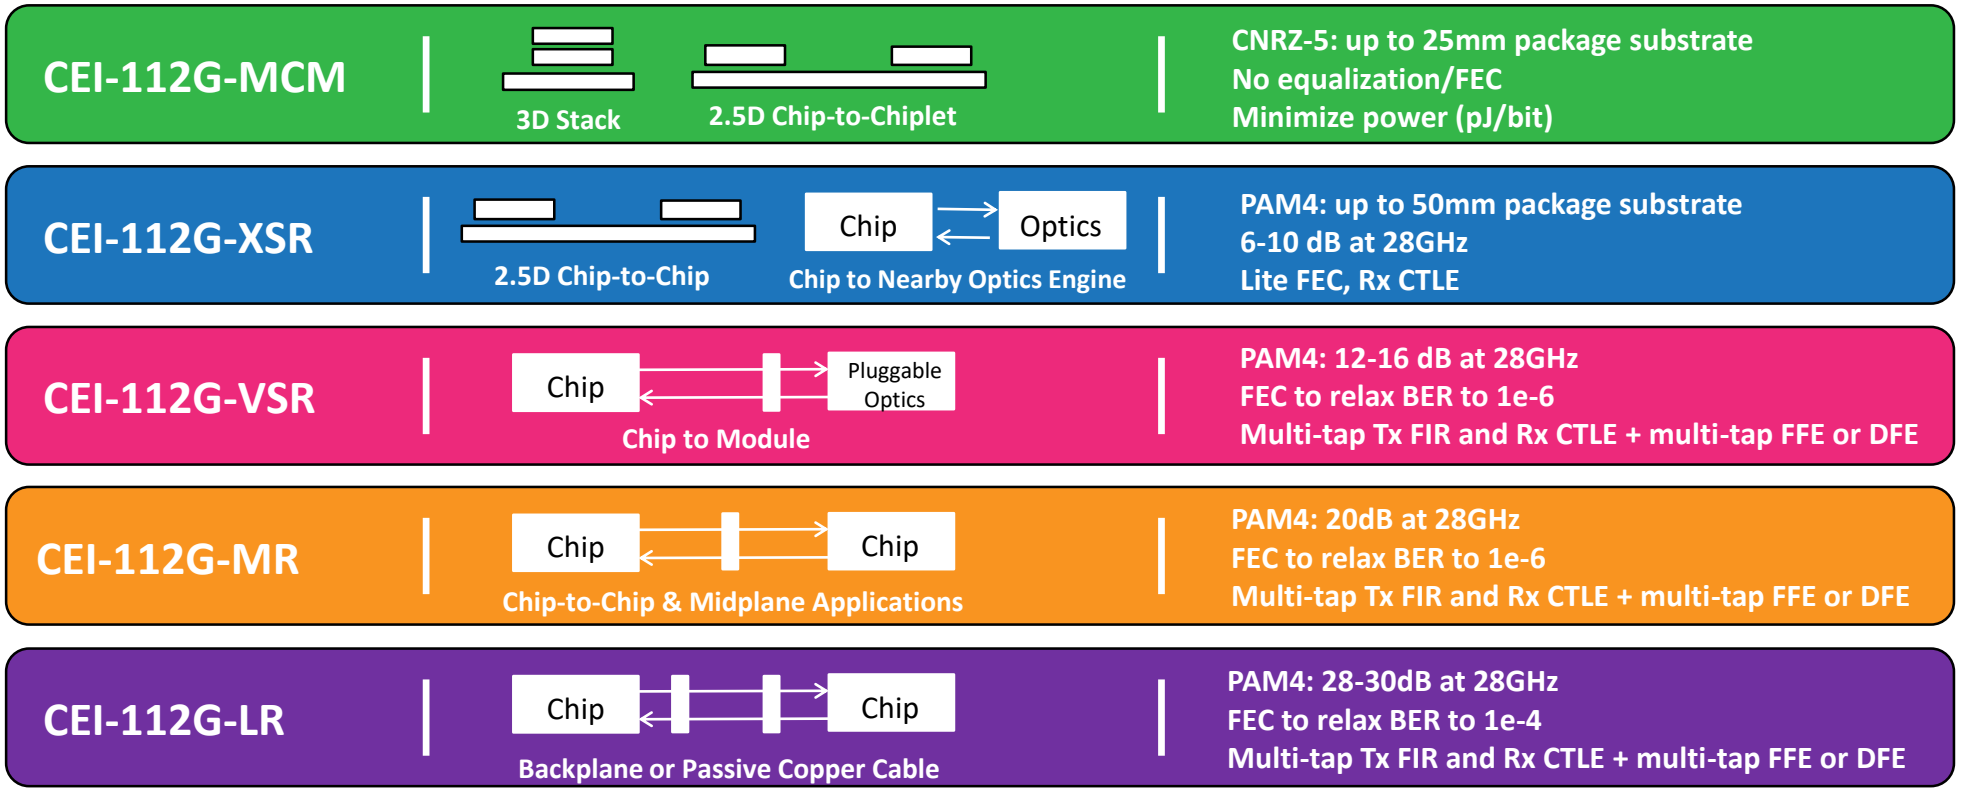

- PAM4 modulation scheme becomes dominant in OIF CEI-112 Gbps interface IA
- One SerDes core might not be able to cover multiple applications from XSR to LR
- For short reach applications, simpler and lower power equalizations are desired

CEI-112G
System Vendor's Perspective

Gary Nicholl (gnicholl@cisco.com)
Principal Engineer, Cisco

OFC 2020, San Jose, March 10-12, 2020

CEI-112G Overview

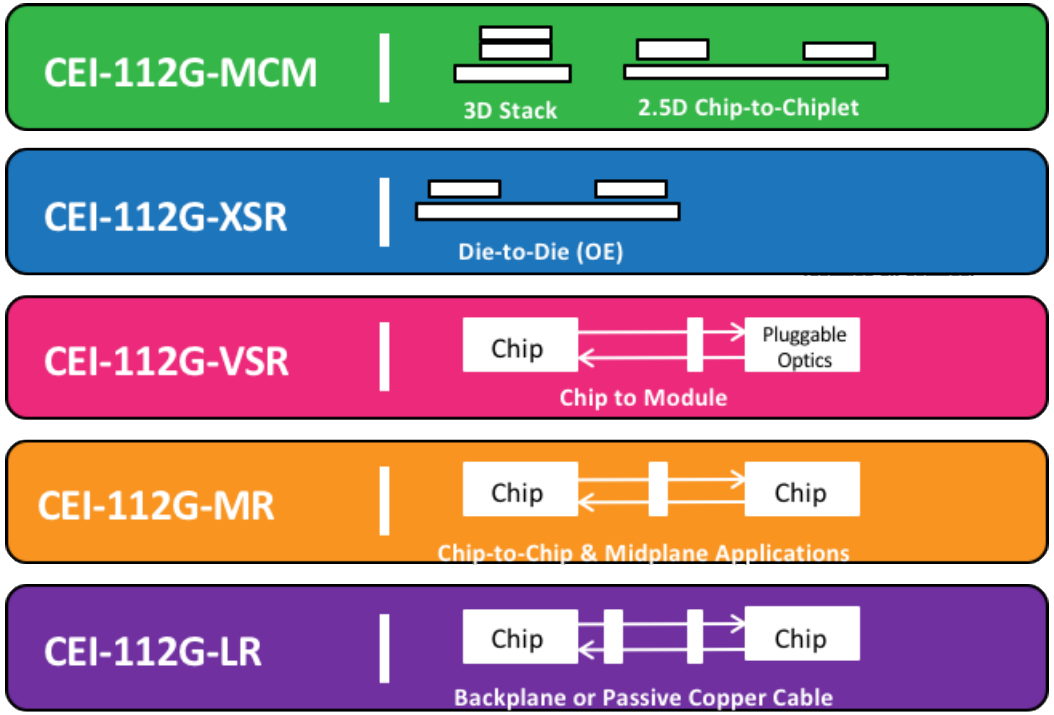

- Where do these fit into real systems ?
- What do they enable ?
- What are some of the key challenges and takeaways ?
- Have we learnt anything from our experiences at 50G ?

System Applications

Fixed (Pizza Box)

Modular (Chassis)

Two primary system form factors:

Fixed (Pizza box):

- 1RU, 32/36 port, up to 12.8T (today)
- Single ASIC architecture (typical)
- Limited flexibility/scalability per box

Modular (Chassis):

- Much larger (up to a full rack)
- Multiple line card (LC) slots (typically 4-18)
- Switch Fabric (provide LC-to-LC connectivity)
- Multi-ASIC architecture
- Very Scalable

Fixed Switch (e.g. 25.6T, 32x800G Pluggable)

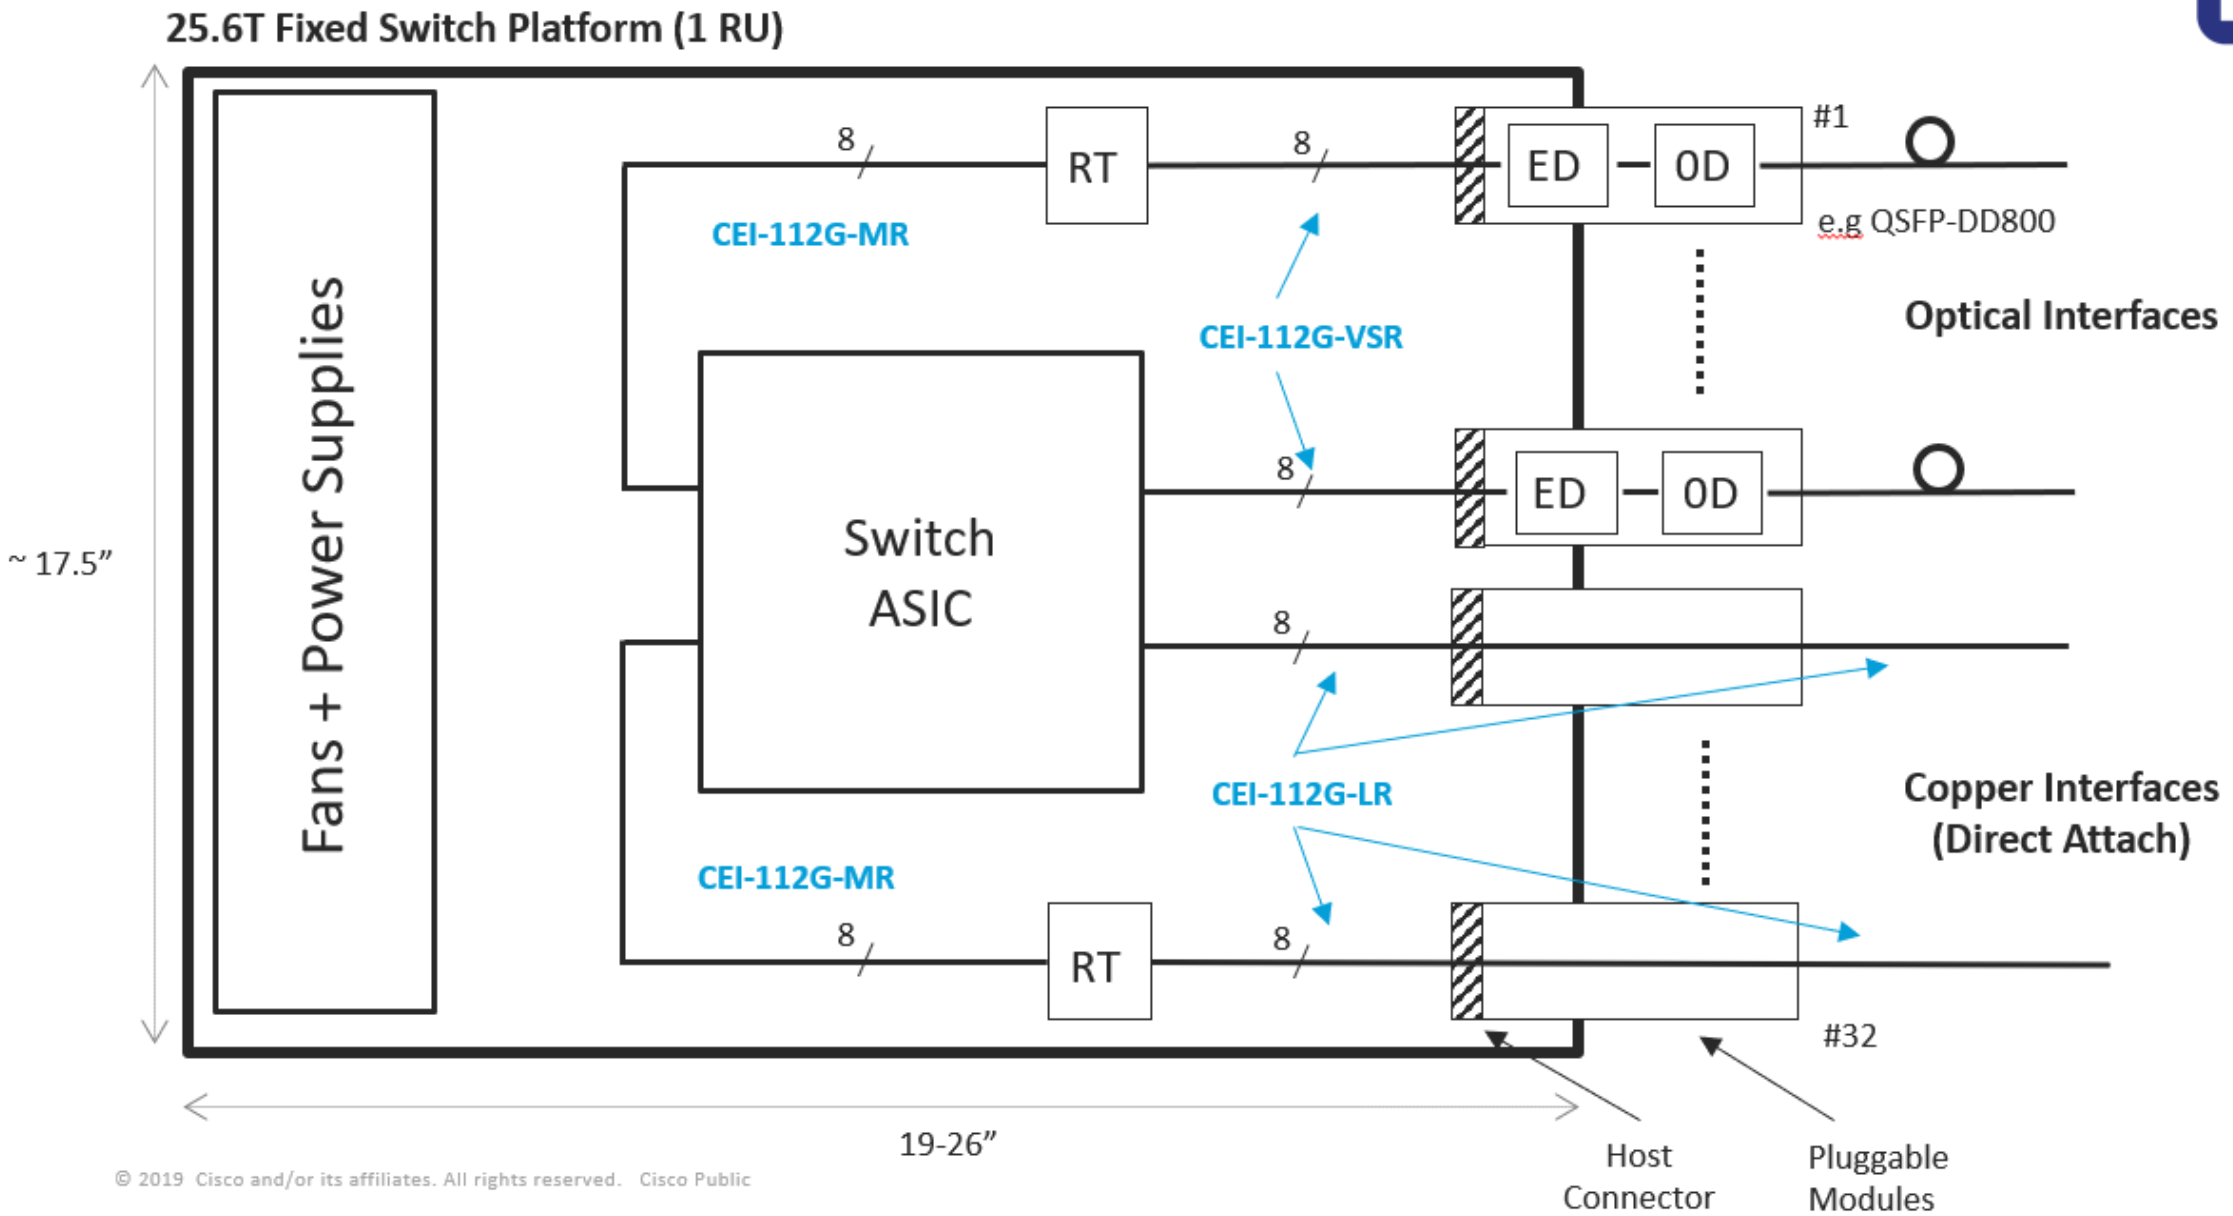

RT=Retimer, ED-Electrical Die, OD=Optical Die

Fixed Switch (more forward looking)

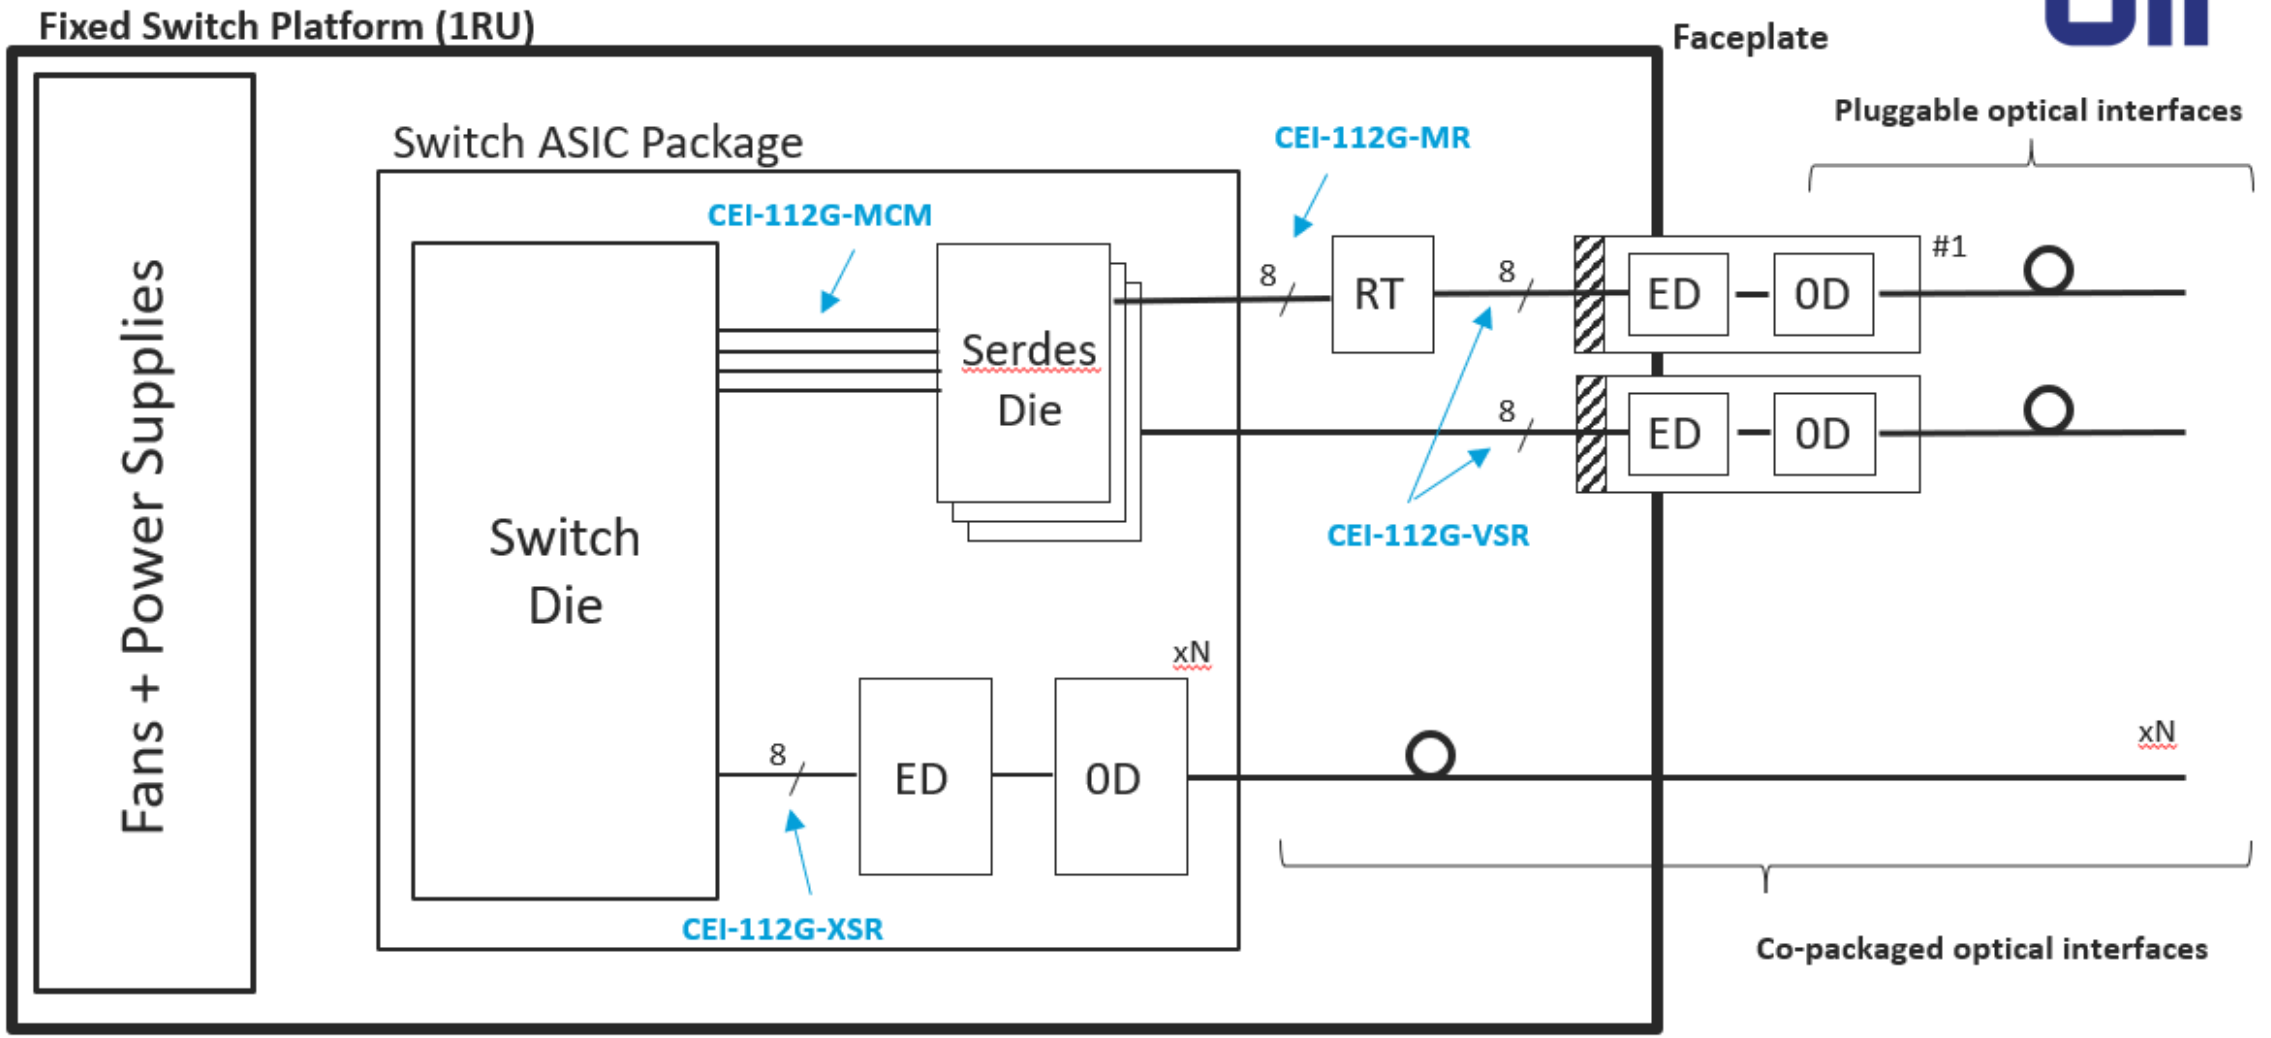

CEI-112G-MCM/XSR typically used within a package (die-to-die)

RT=Retimer, ED-Electrical Die, OD=Optical Die

Modular Switch

Modular Switch Platform

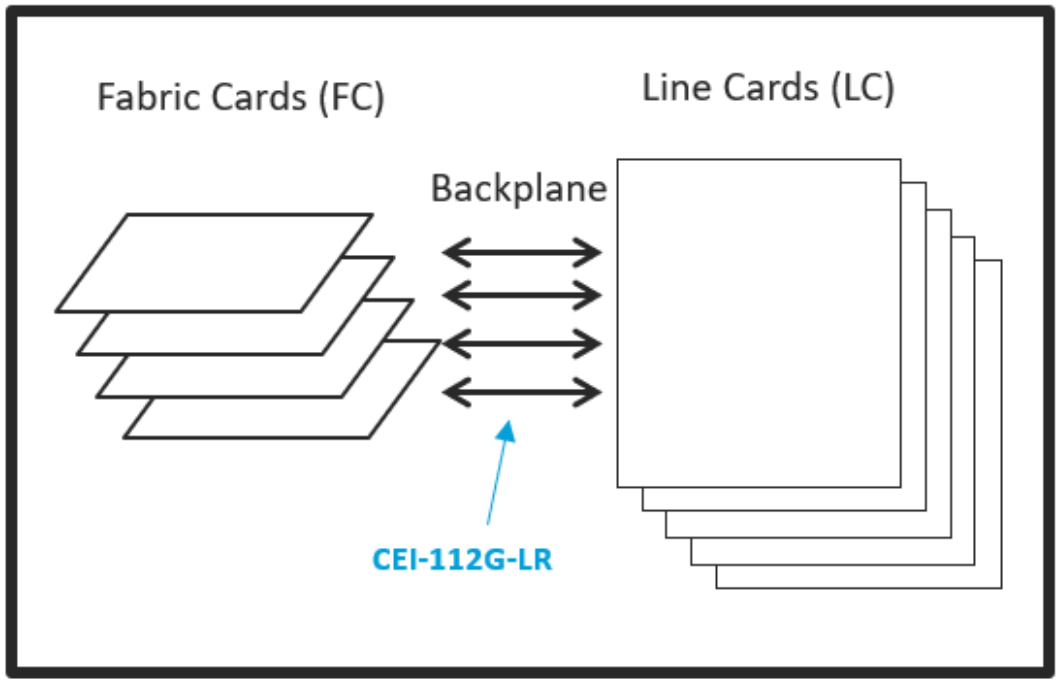

- Front (customer) facing half of Line card is identical to Fixed Switch in terms of electrical interfaces:
  - CEI-112G-VSR/MR/LR
- Main difference is the additional links required to interconnect the Line cards and Fabric cards:
  - CEI-112G-MR/LR

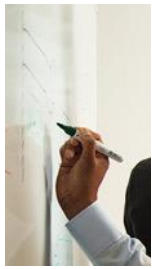

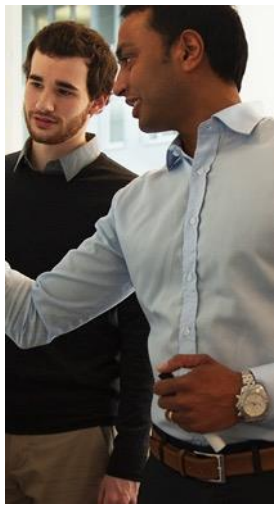

CEI-112G: Status Update on VSR, and LR Channels

Nathan Tracy, OIF President and TE Connectivity Technologist

EVERY CONNECTION COUNTS

What Happens When We Double The Data Rate?

Going from 50 Gbps to 100 Gbps electrical rates brings a host of challenges including noise (cross talk), reflections, mode conversion, etc., but especially insertion loss (reach), because the equipment is the same size as it was at 50 Gbps.

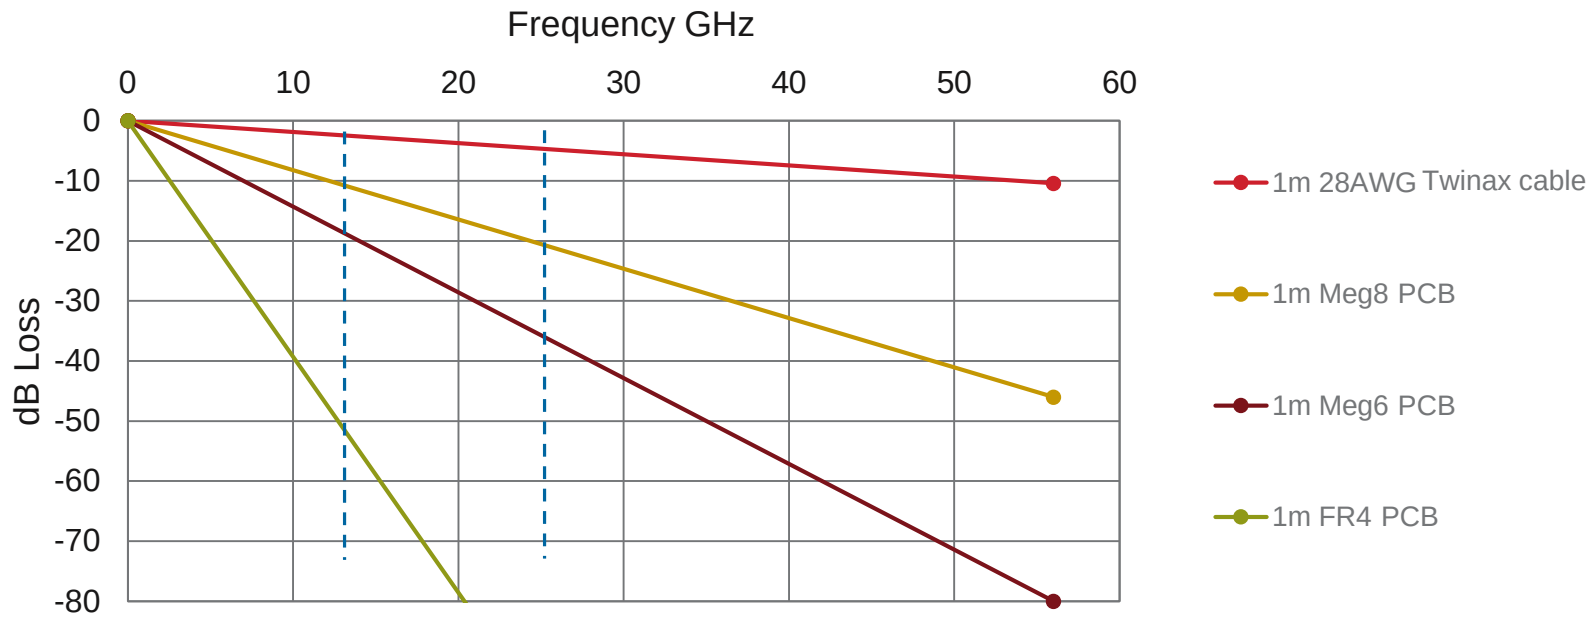

As we consider channels for 100 Gbps, we need lower loss architectures

CEI-112G-VSR-PAM4

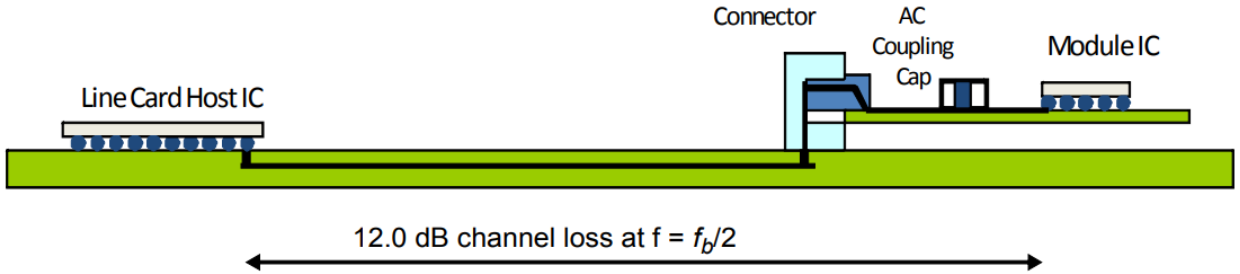

Interface for chip to module applications with a channel reach of at least 10 cm of host trace, one connector and 2 cm of module trace

Defines compliance test methodology

VSR specifies Tx and Rx parameters at the connector compliance point since host assemblies and modules are made by different parties

Alternate implementations are allowed

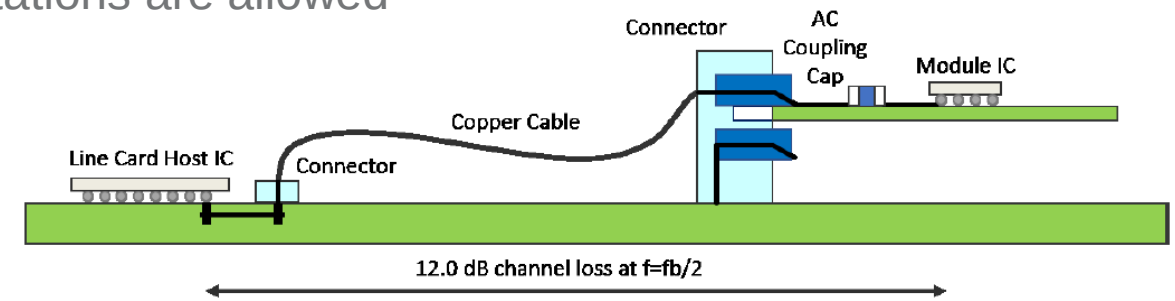

VSR Channel, Test vs Model Results, Existing 50Gbps Connector

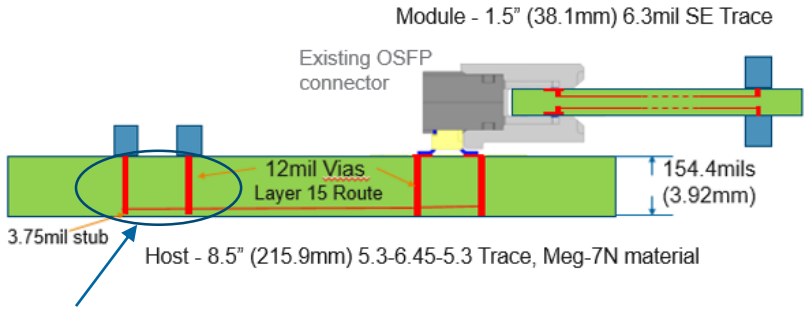

Note: Measured channel includes second set of vias to test point, modeled channels do not include the second set of vias.

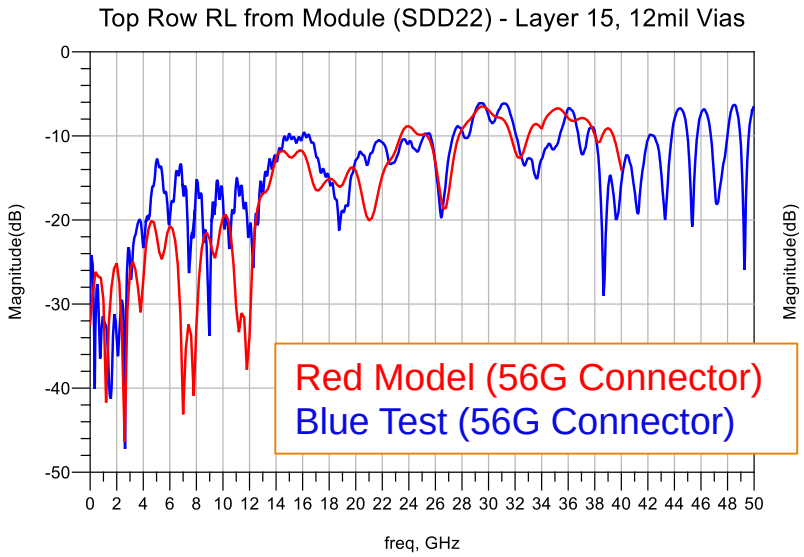

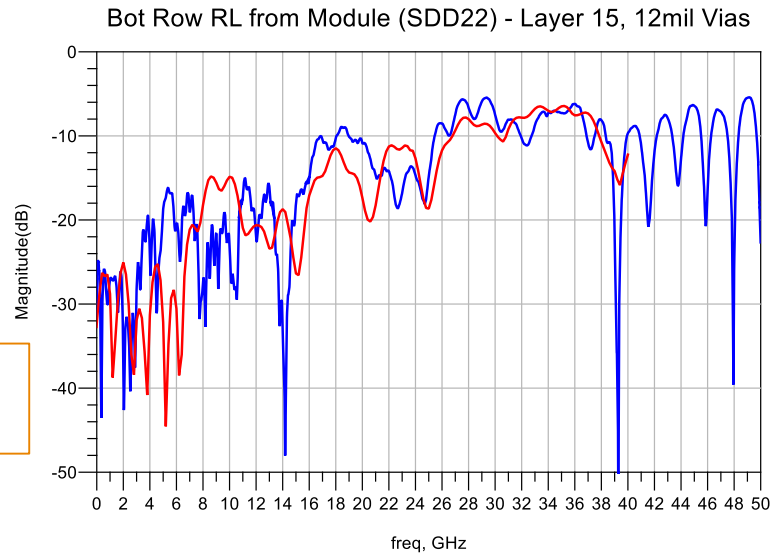

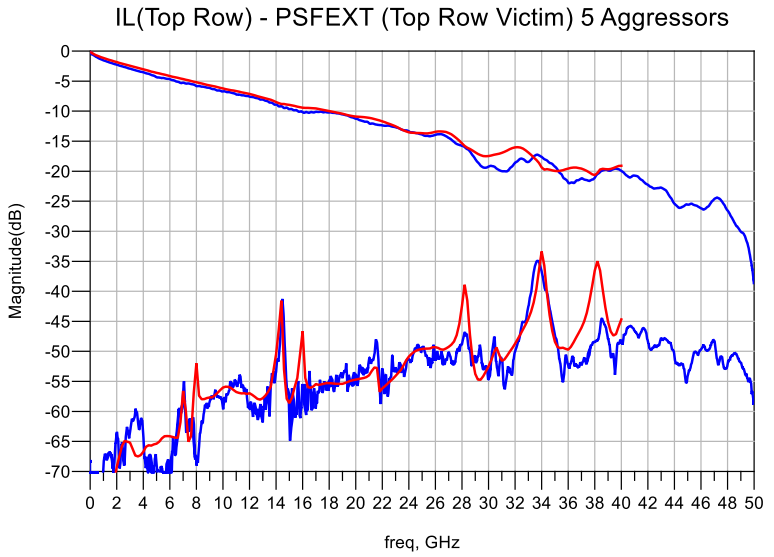

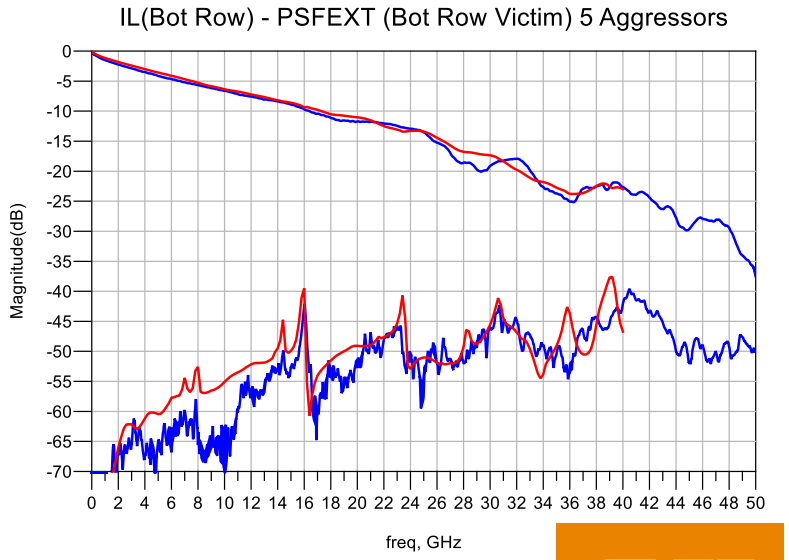

100G Connector Improvements (in the same 15dB channel)

Efforts were focused on insertion loss and return loss optimizations

Improved 100Gbps OSFP connector

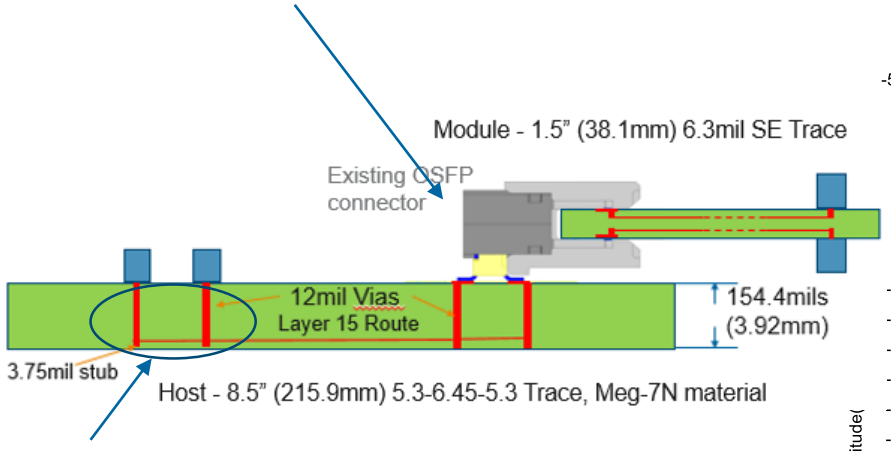

Note: Measured channel includes second set of vias to test point, modeled channels do not include the second set of vias.

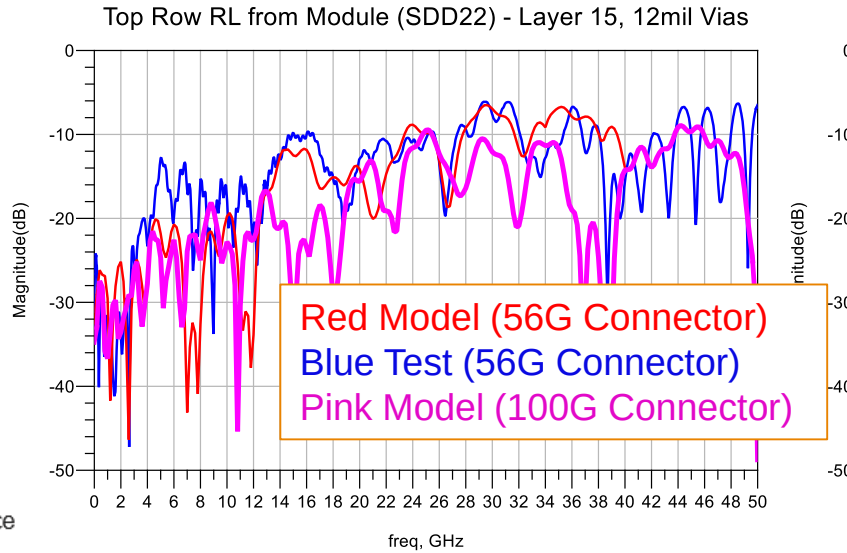

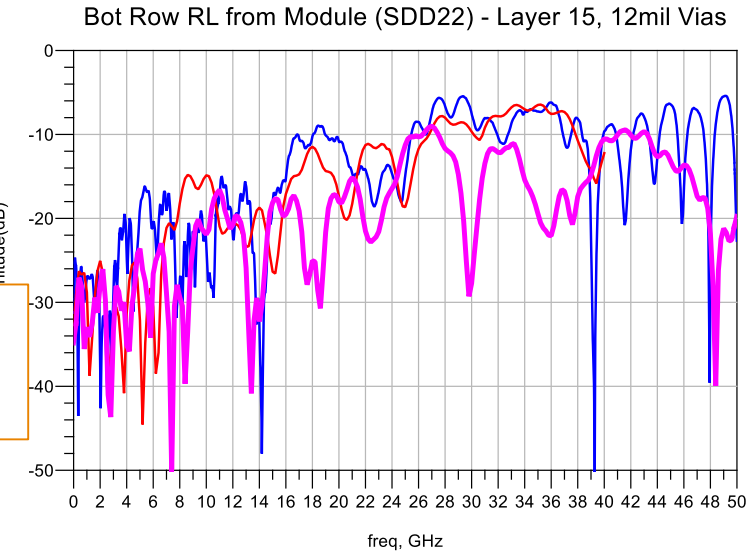

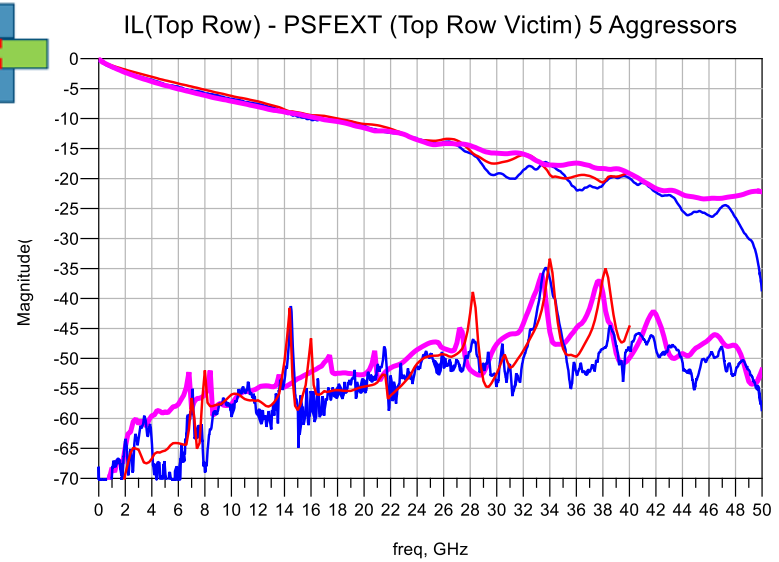

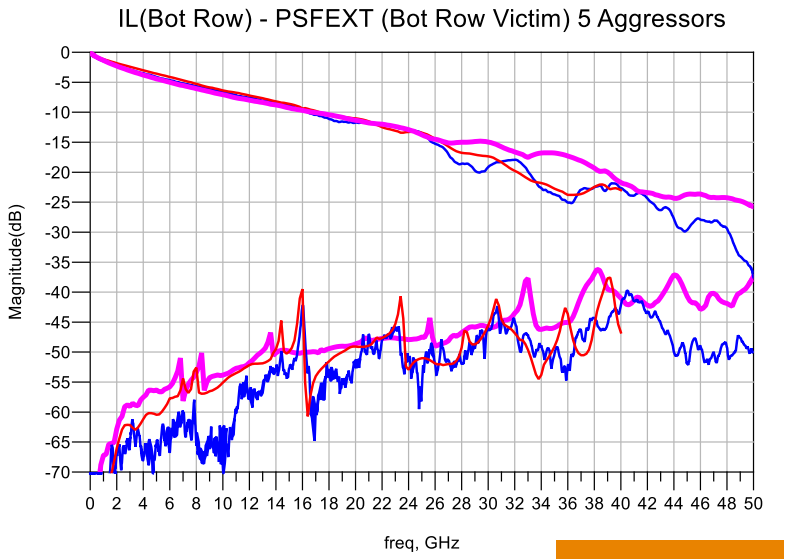

CEI-112G-LR-PAM4

Interface for backplane applications with a channel reach of 100 cm over 1 or 2 mated connectors and PCB and/or twinax cables.

Compliance point will be semiconductor ball.

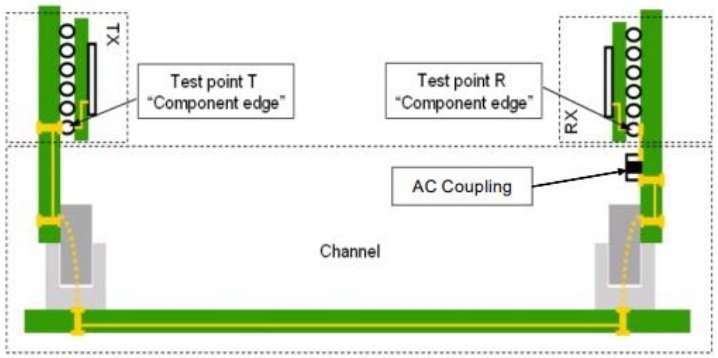

Traditional Backplane Architecture

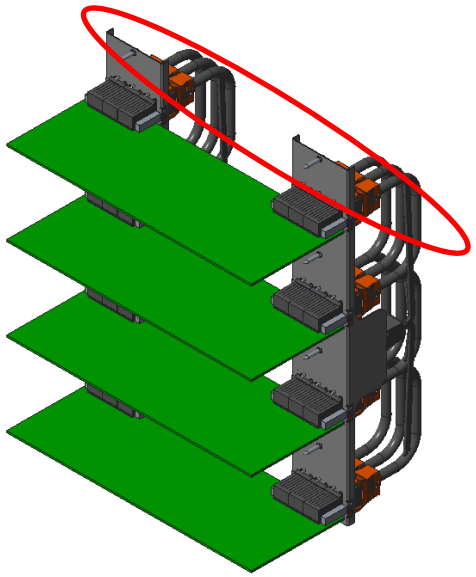

Cabled Backplane Architecture

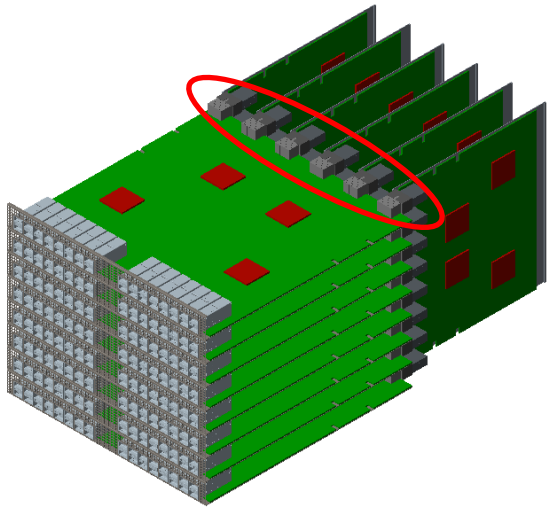

Orthogonal Backplane Architecture

CEI-112G-LR Interface

Cathy Liu

CEI-112G-LR - Channels

CEI-112G-LR

Backplane or Passive Copper Cable

PAM4: Backplane or Cu cable interface
IL: Up to 28dB at 28GHz
FEC to relax BER to 1e-4

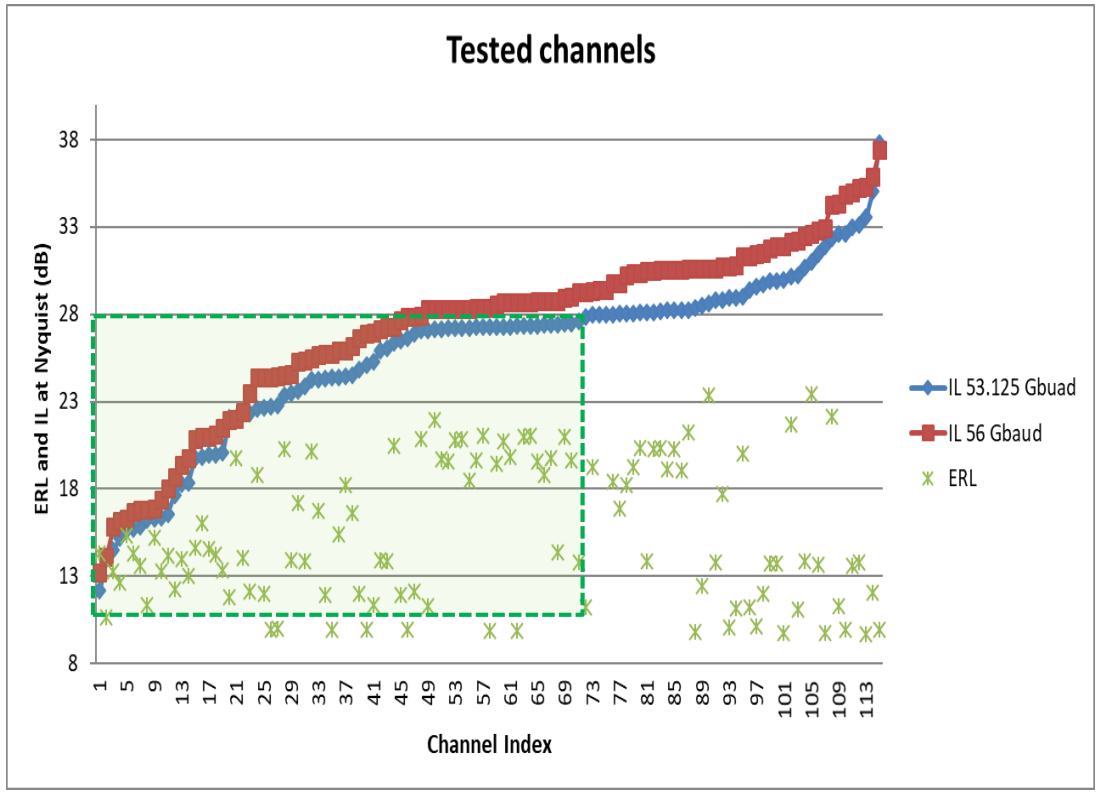

- Differentiations between CEI-112G and 802.3ck
  - CEI covers more interfaces
  - CEI covers wider range of baud rates
  - CEI allows low latency and low cost FEC for VSR and XSR interfaces
- CEI-112G-LR COM and ERL define the channel compliance
  - Insertion Loss target: 28dB at Nyquist
  - ERL specifies channel effective return loss: 10.5dB
- Trade-offs between:
  - Channel loss
  - BER target
  - Link rates
  - Power and cost

CEI-112G-LR – Reference Equalization

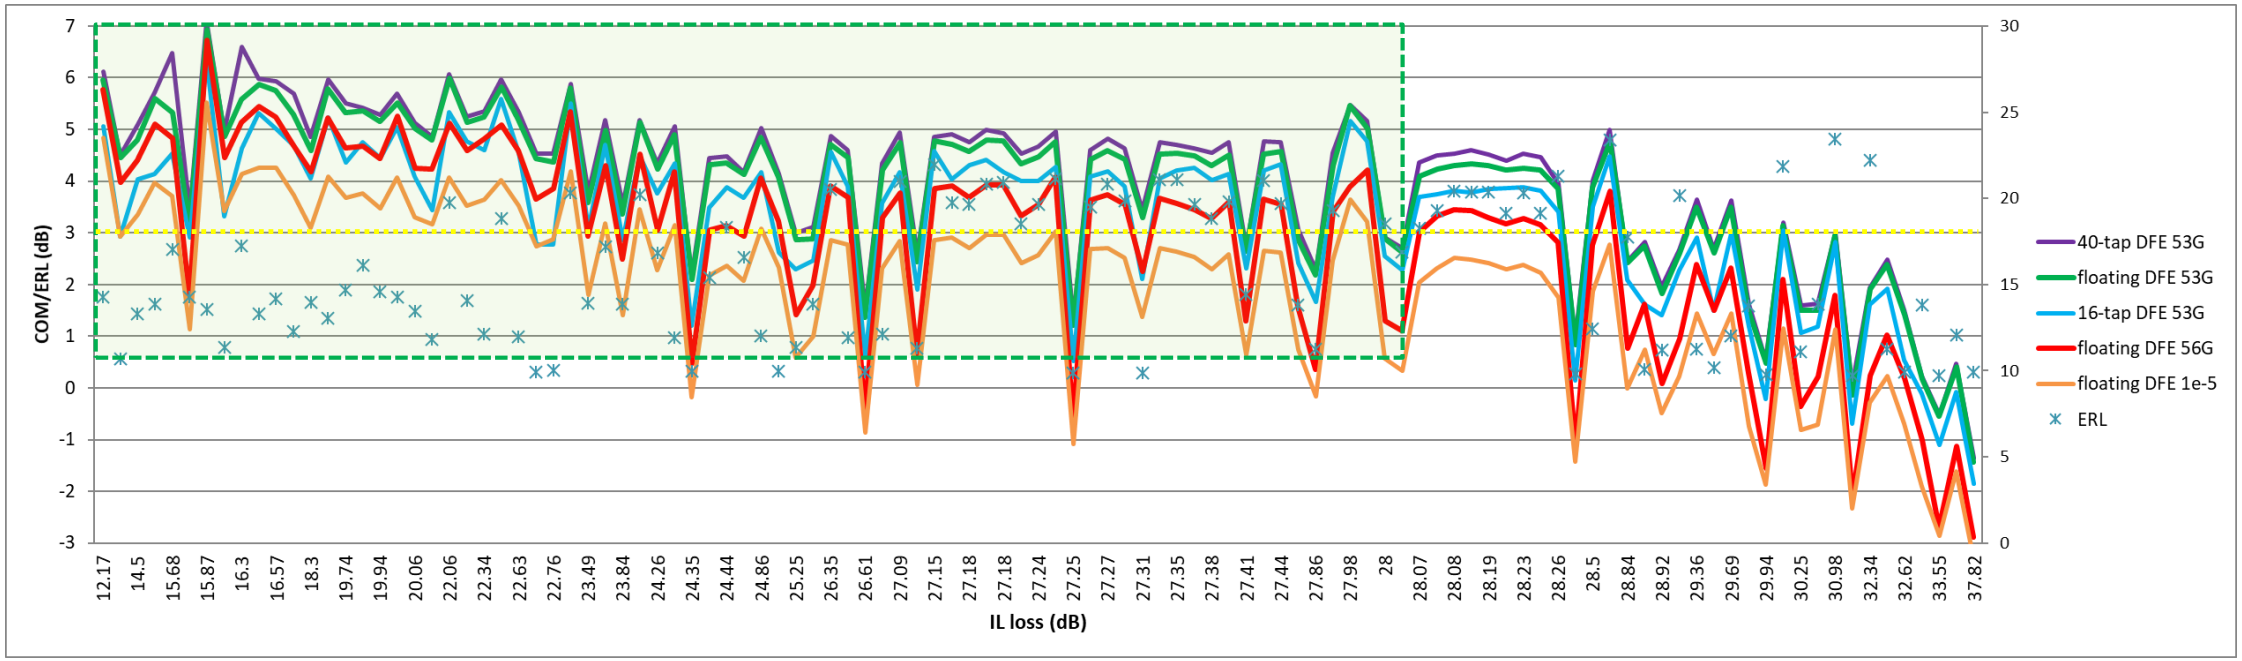

- Latest CEI-112G-LR IA adopts
  - 5-tap Tx FIR
    - c(-3), c(-2), c(-1), c(0), and c(1)
  - CTLE
    - 2-zero and 3-pole
  - DFE:
    - Floating tap DFE 12 fixed taps, 3 banks of 3 floating taps with 40UI span
    - Coefficient constraints bmax(1)=0.85, bmax(2-3)=0.3, bmax(4-12)=0.2, bmax(floating)=0.05, bmin=-0.05

DFE schemes	% passing channels (among green zone)
Floating DFE/53G/1e-4	97
16-tap DFE/53G/1e-4	86.6
Floating DFE/56G/1e-4	87.5
Floating DFE/56G/1e-5	41.3

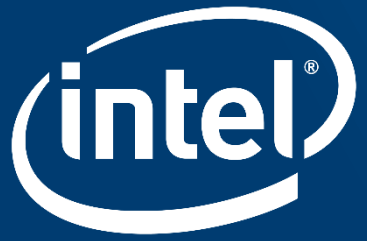

CEI-112G-XSR-PAM4 Interface

Mike Peng Li

Intel

CEI-112G-XSR-PAM4 Link Requirements

26.1 Requirements

1. Support serial baud rates (f_b) within the range from 36 Gsym/s to 58 Gsym/s as specified for the device using PAM4 coding. A CEI implementation complies to the specifications of this clause over the range of baud rates stated for the implementation within this range.
2. Capable of achieving a raw Bit Error Ratio (BER) of 10^{-9} or better per lane. FEC is assumed to be used in the system to achieve corrected BER of 10^{-15}or better per lane. The baud rate includes the overhead required for FEC. The definition of FEC is outside the scope of this IA (See Appendix 26.A for FEC guidance).
3. Capable of driving up to 50 mm of package substrate trace, or package (one or two package of the link) on package, with no connector.
4. Shall support DC-coupled operation. AC coupling can be supported by capacitors implemented on die.
5. Shall allow multi-lanes (1 to n).

CEI-112G-XSR-PAM4 Channel Requirements

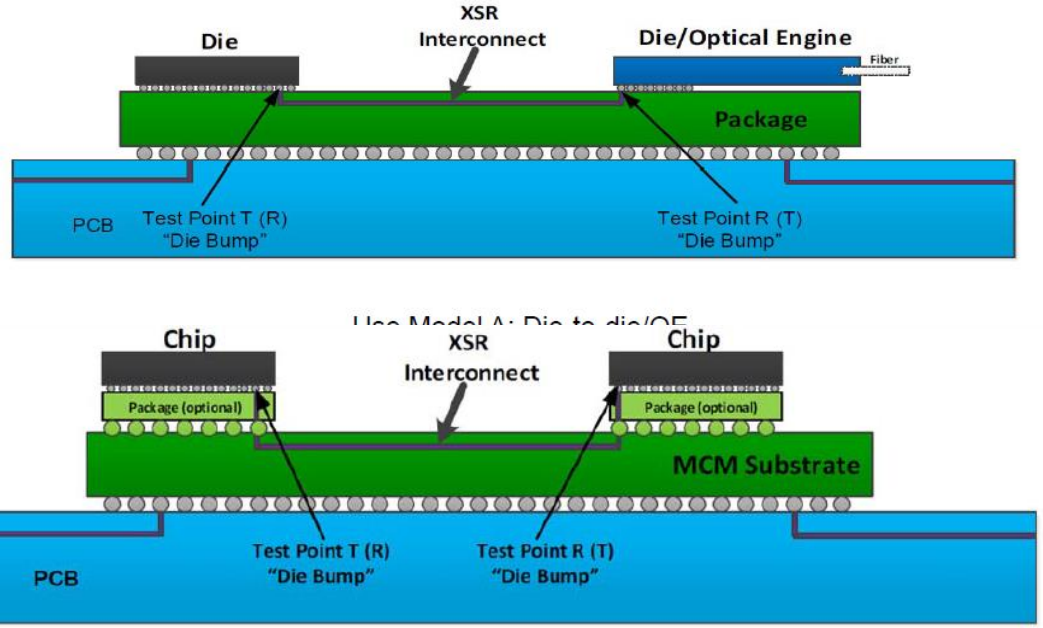

Figure 26-2. Channel Insertion Loss Limit for 58.0 Gsym/s

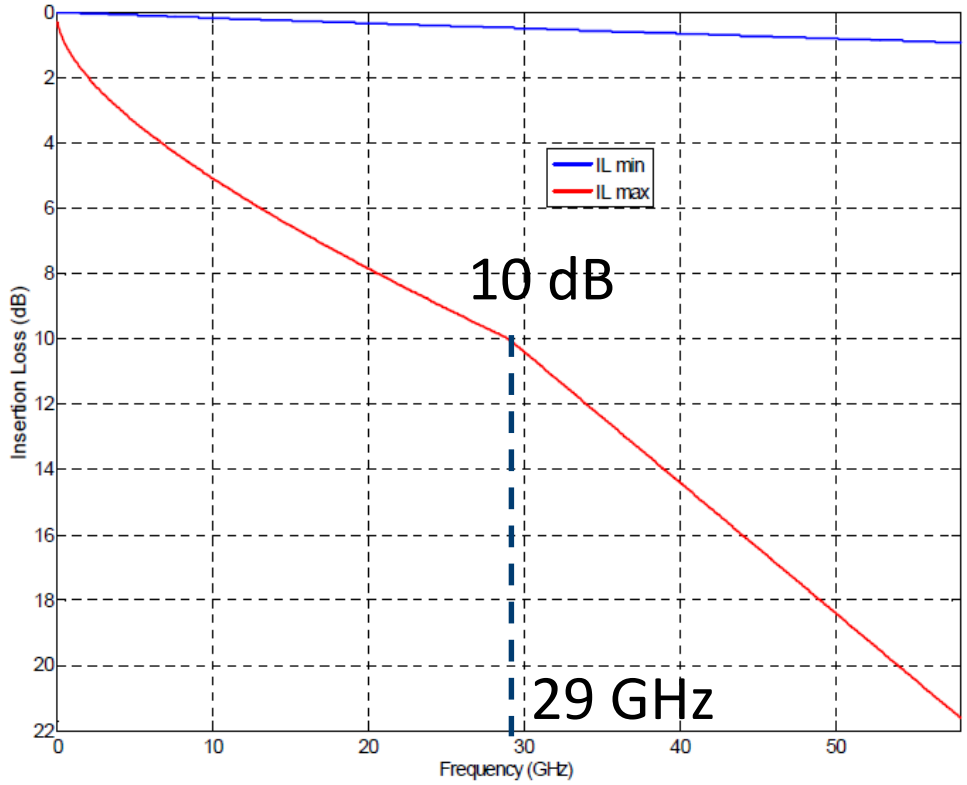

CEI-112G-XSR-PAM4 COM Highlights

Ref TX

- 3-tap FIR: 1-pre, 1-post

Ref RX

- CTLE (8 dB)

BER

- 1e-9 (raw)

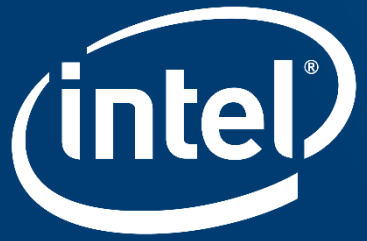

CEI-112G-MR-PAM4 Interface

Mike Peng Li

Intel

CEI-112G-MR-PAM4 Link Requirements

1. Support serial baud rates (f_b) within the range from 36.0 Gsym/s to 58.0 Gsym/s as specified for the device using PAM4 coding. A CEI implementation complies to the specifications of this clause over the range of baud rates stated for the implementation within this range.
2. Capable of achieving a raw Bit Error Ratio (BER) of 10^{-6} or better per lane. FEC is assumed to be used in the system to achieve corrected BER of 10^{-15} or better per lane. The baud rate includes the overhead required for FEC. The definition of FEC is outside the scope of this IA (see Appendix 16.D).
3. Capable of driving up to 500 mm of backplane and up to 1 connector.
4. Shall support AC-coupled operation.
5. Shall allow multi-lanes (1 to n).
6. Shall support hot plug.

CEI-112G-MR-PAM4 Channel Requirements

Figure 25-1.CEI-112G-MR-PAM4 Reference Model

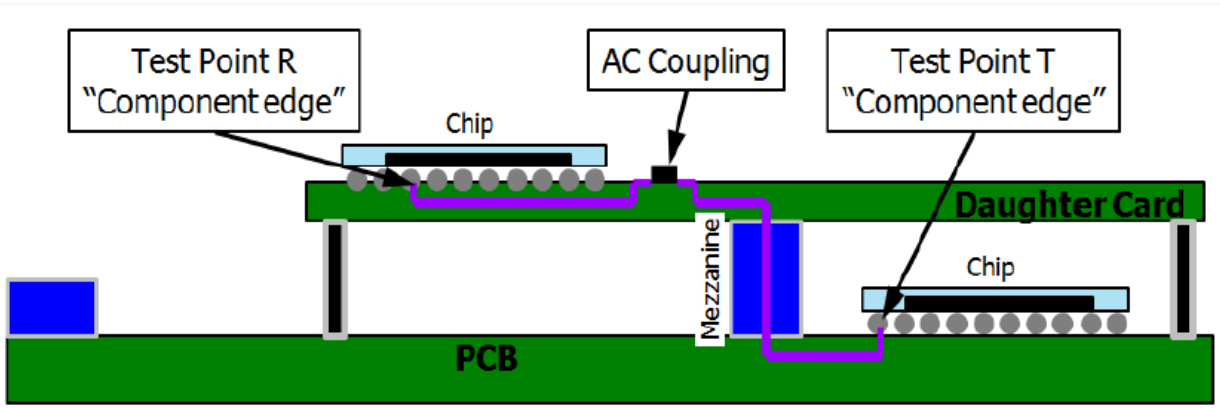

CEI-112G-MR-PAM4 COM Highlights

Ref TX

- 5-tap FIR: 3-pre, 1-post

Ref RX

- CTLE (20 dB), and 14-tap DFE

Ref package length

- 7 mm to 31 mm, TX/RX

BER

- 1e-6 (raw)

25.2.4.3 Informative Channel Insertion Loss

Figure 25-2.Channel Insertion Loss Limit for 58.0 Gsym/s

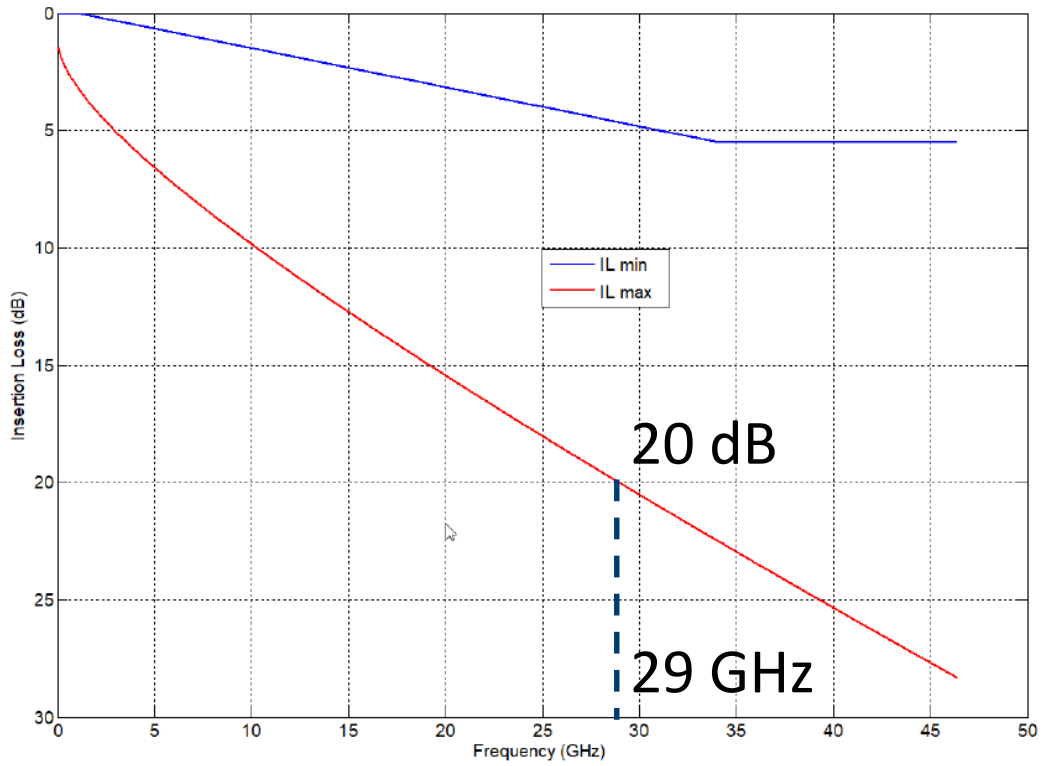

CEI-112G MCM and VSR Interfaces

Ed Frlan

Principal Product Definition Specialist, Semtech

CEI-112G-MCM Overview

MCM applications based on interconnected chiplets:

- Combination of many dies into large packages

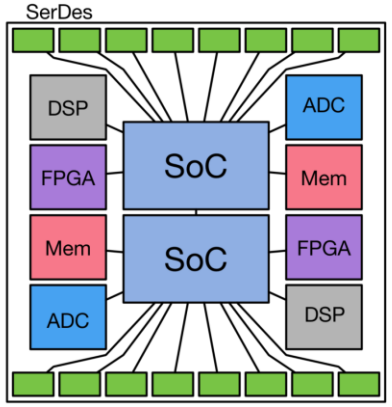

Non-interposer MCMs can easily use 20 or more dies

- 2.5 cm can accommodate a 70 mm package
- Improvement in yield and cost over a monolithic IC
- Enables multi-vendor ecosystem
- I/O subsystem dies contain SerDes placed around the perimeter, creating smaller virtual packages

CEI-112G-MCM key features:

- 2.5 cm bump-to-bump (up to 6 dB loss)
- CNRZ-5 modulation
- DC coupled path
- Silicon to silicon
- Clock forwarded
- Fat-pipe applications (i.e., between logic chips)
- System raw BER ≤ 10-15

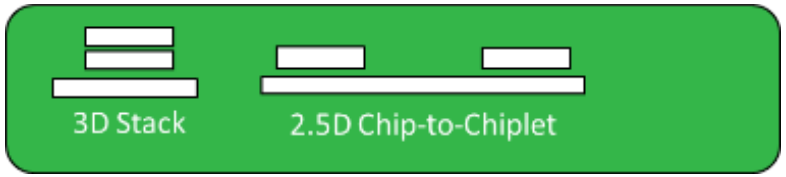

CEI-112G-VSR Overview

112G VSR application:

- For next-generation higher density pluggable modules (e.g., those based on 4x 100G or 8x 100G electrical interfaces)
- Still a traditional chip-to-module interconnect with one connector

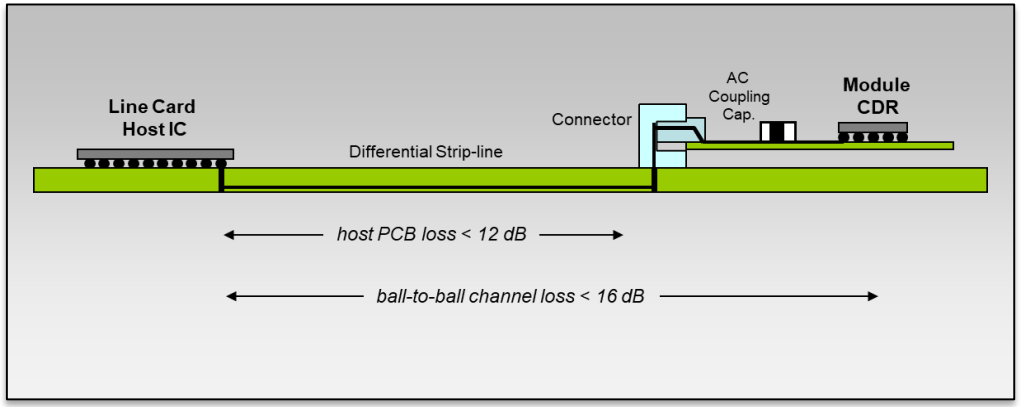

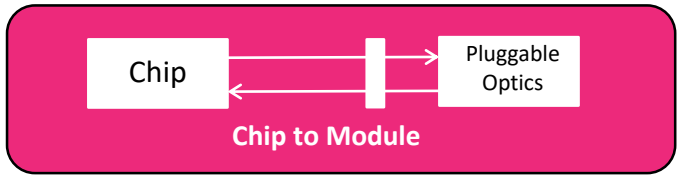

CEI-112G-VSR key features:

- 16 dB channel insertion loss at 29 GHz
- PAM4 modulation with allowable range of 36 – 58 GBd
- System pre-FEC BER ≤ 10^{-6}
- AC coupled channel with one connector
- Clocking based on per lane CDR function
- Receivers required to be self-adaptive and autonomous
- Acceptable system transmitter performance based upon reference receiver oscilloscope measurements via mated compliance boards at connector interfaces
- Reference receiver architecture is CTLE with 4-tap DFE

CEI-112G-VSR Challenges

Tradeoff between equalizer capability and expected channel performance

- Especially important for VSR and XSR types of interfaces where transceiver pJ/bit is paramount
- 112G VSR transceiver efficiencies < 3 pJ/bit likely to be enabled with simple equalization schemes
- Actual implementations may be Analog or DSP

VSR application space is larger than IEEE 802.3ck chip-to-module

- OIF BER < 10^{-6} (cf. IEEE 802.3ck BER < 10^{-5})
- OIF baud rate < 58 GBd (cf. IEEE baud rate = 53.1 GBd)

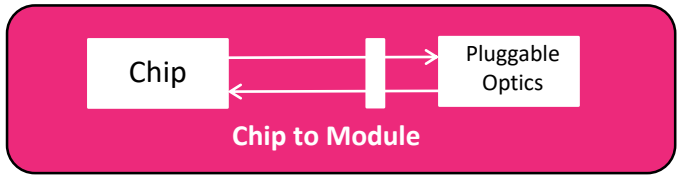

Sensitivity of PAM4 modulation to channel discontinuities, especially for lower loss interconnects

- Short channels exhibit significant sensitivity to present package model discontinuities at the PCB interface

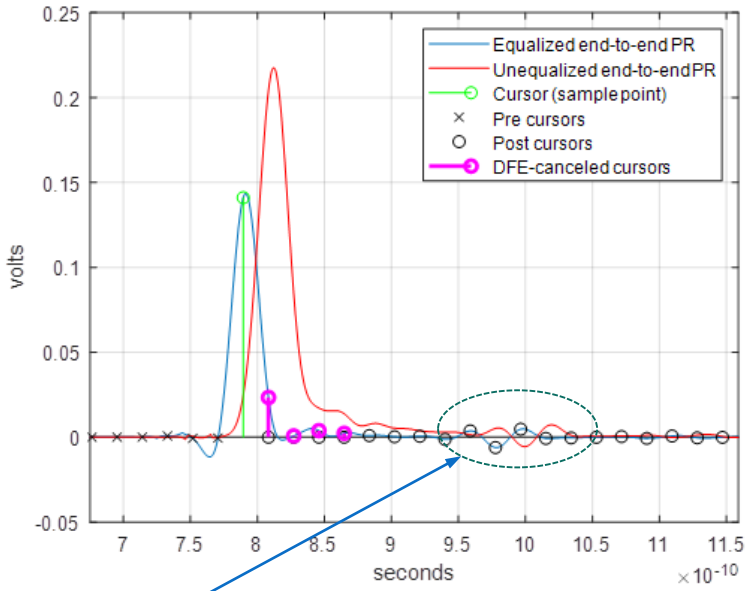

These reflections result in significant system penalties

OIF CEI-112G Measurement Update

Steve Sekel	11 MAR 2020
400G Solutions Specialist Keysight Technologies	

CEI-112G Compliance Measurements

EVOLUTIONARY FROM CEI-56G BASE

- Unlike CEI-53G Revolutionary shift from NRZ to PAM4, CEI-112G is more Evolutionary
- Refinements to measurements used in CEI-56G - No major new measurements added
- Assumption that receivers will be implemented with ADCs
- Margins are much tighter – artifacts which were considered “too small” now need consideration
  - Have we identified them all? - How close will actual measurements match COM simulations?
- More complex reference receivers

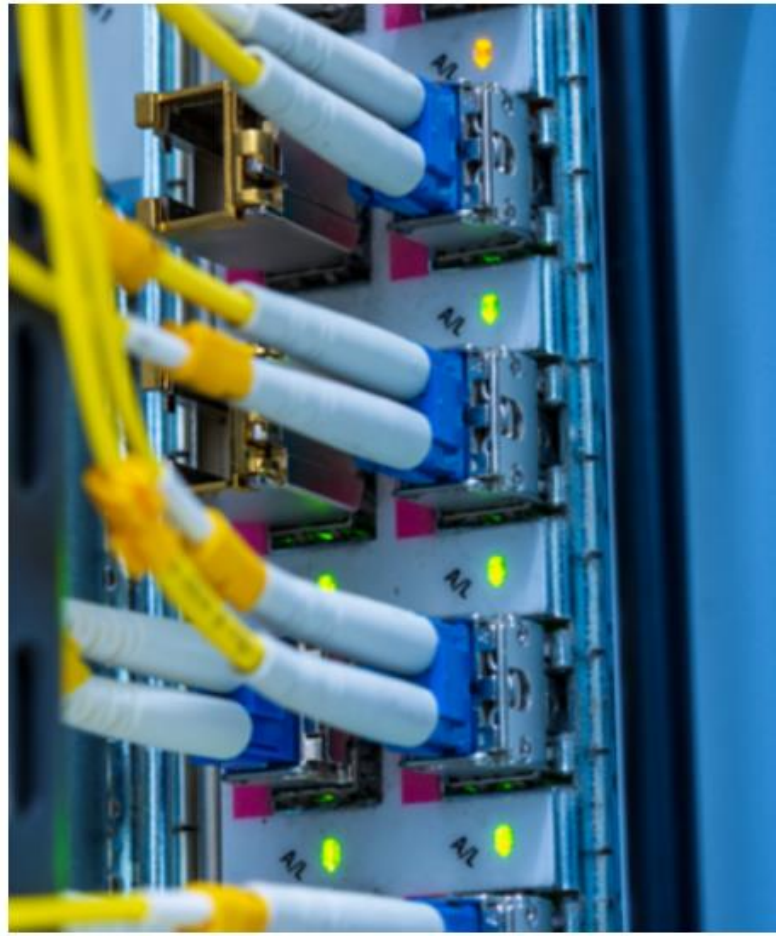

Impacts on shrinking margins

"SMALL" ARTIFACTS NO LONGER TOO SMALL

- Receiver front end analog distortion and noise
- ADC sampling clock jitter
- ADC quantization noise
- Equalizer coefficient quantization effects
- Actual termination impedance versus nominal
- Differential skew errors

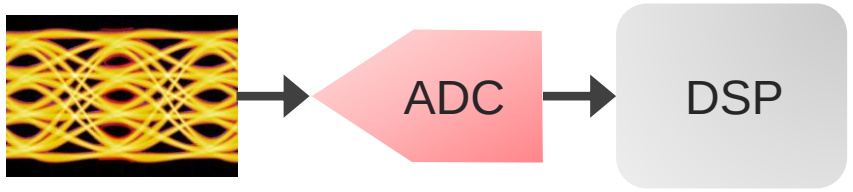

Work to be done

- Complete Reference Receiver implementation / optimization in both sampling and RT oscilloscopes
- Lab measurements with real SerDes outputs
- Compare with simulation
- Compare results measured with different instruments (Sampling and Real Time)
- Compare results measured with different vendors instruments
- In addition to obtaining repeatability, work will be required to speed measurement acquisition and processing to a reasonable level
  - (Measuring true EW-6*EH-6 product for optimization is much too slow!)

112 Gbps Electrical Interfaces: An OIF Update On CEI-112G

QUESTIONS?