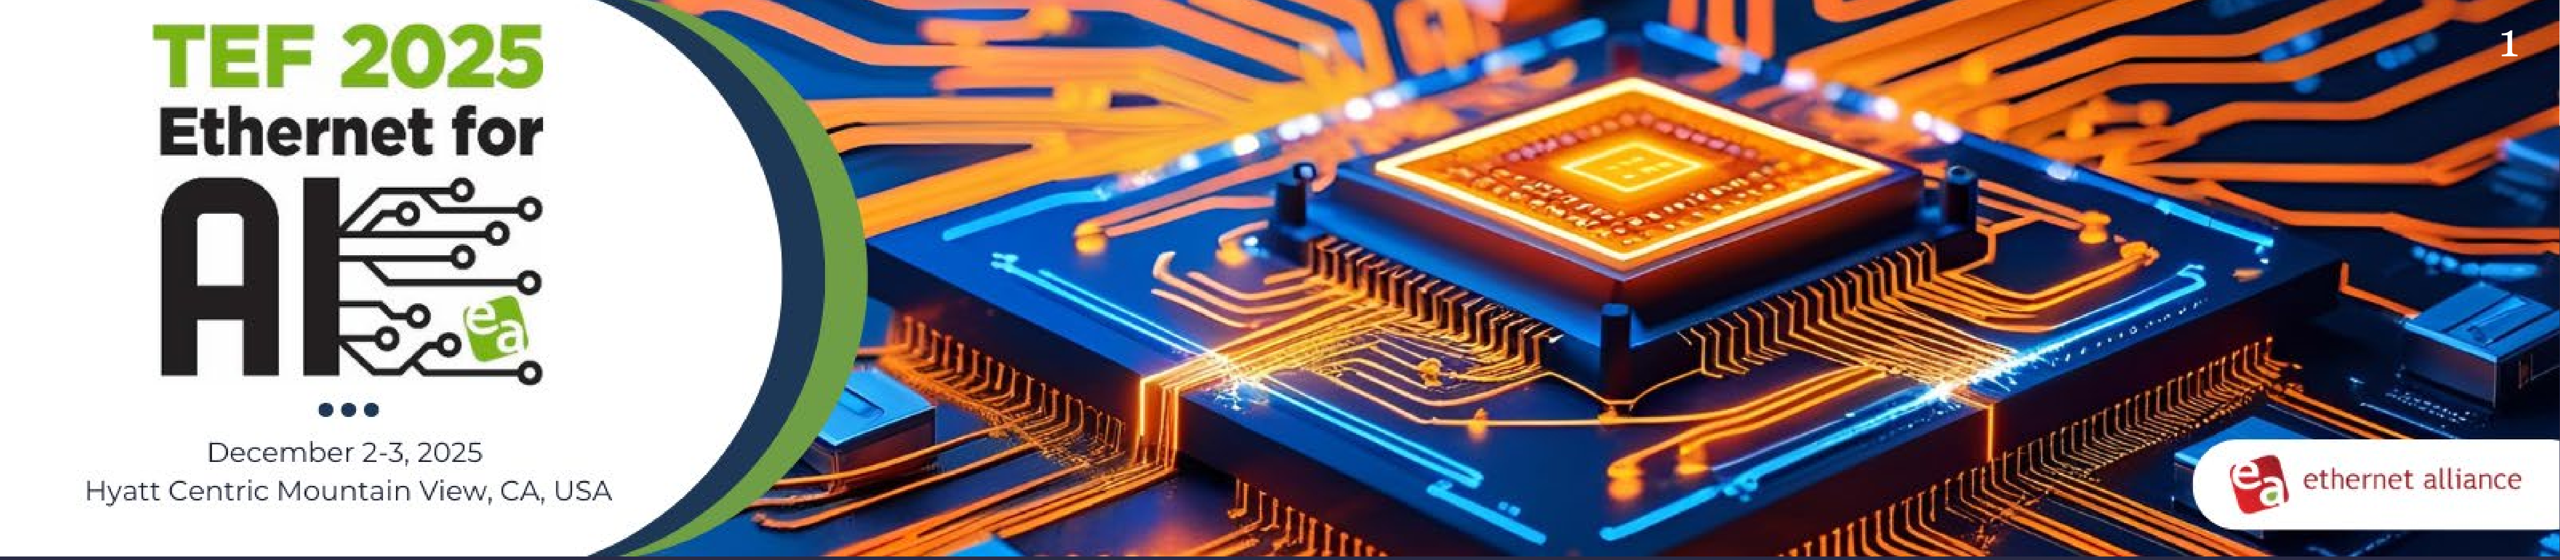

400G Electrical Signaling For AI Networks

December 2-3, 2025

This presentation has been developed within the Ethernet Alliance, and is intended to educate and promote the exchange of information. Opinions expressed during this presentation are the views of the presenters, and should not be considered the views or positions of the Ethernet Alliance

400G Electrical Signaling Panel

Panelists

- **Ayal Shoval, Synopsys**, “Proposed Enhancements to COM: Addressing RLM Sensitivity in Next-Generation Ethernet Designs”
- **Hadrien Louchet, Keysight**, “Advancing 448G Signal Generation and Analysis with Modern Test & Measurement Solutions”
- **Jim Hsieh, MediaTek**, “Unveiling Bottlenecks in Measured Co-Packaged Copper Channels: Sensitivity Analysis at 100 GB+ for PAM6/4”
- **Bijan Mowroozi, Lightmatter**, “Modularity at the Package Edge: Composable Interconnect for 400 Gb/s/Lane AI Systems”

Proposed Enhancements to COM: Addressing RLM Sensitivity in Next-Generation Ethernet Designs

Ayal Shoval, Chief Architect, Synopsys
John Stonick, Fellow, Synopsys

Overview

- PAM6 exhibits higher margin loss due to impairments compared to PAM4, one example is ratio level mismatch (RLM) distortion
- To account for margin loss due to RLM, COM derates the SBR
- For 400Gbs/lane, we recommend updating the RLM derating factor in COM to better align the margin loss with SerDes performance

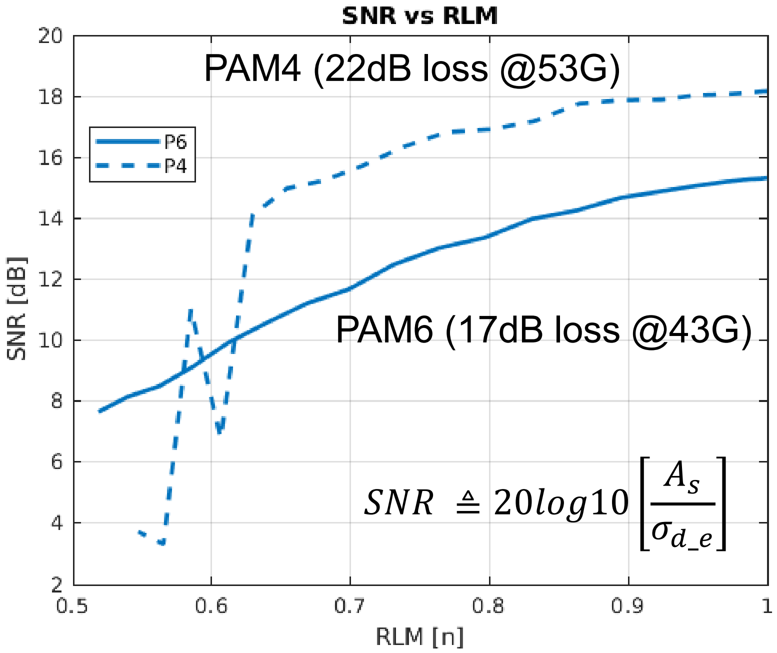

same physical channel for both rates

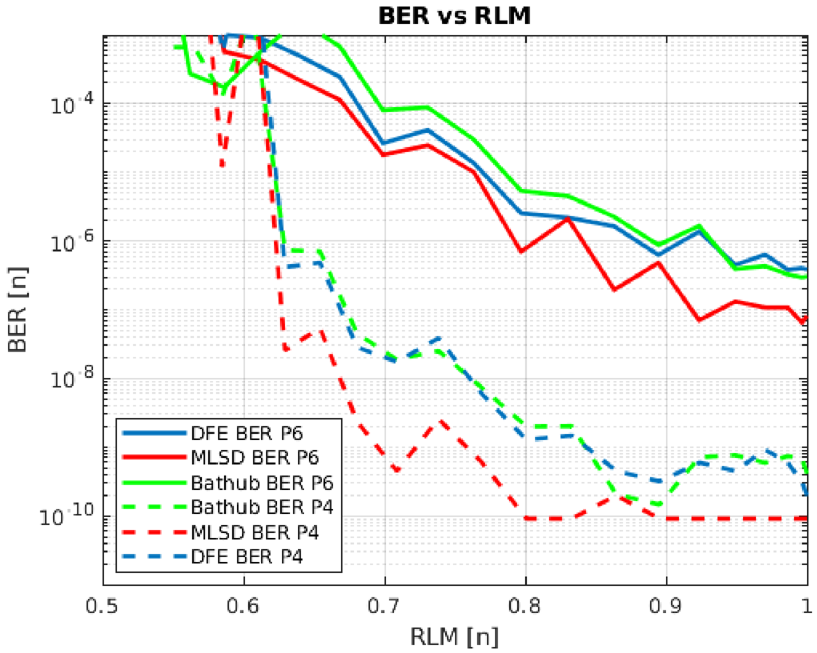

TEF 2025 Ethernet for AI

Effect of RLM Distortion (a review)

- RLM is the actual ratio of PAM levels divided by the ideal ratio of PAM levels
- Here A_x is some amplitude prior to the effect of distortion, and A_y is the observed output amplitude post distortion
- For the two PAM4 amplitudes shown, the effect of distortion as measured by the RLM is (0.57/0.2)/(3/1) = 0.95

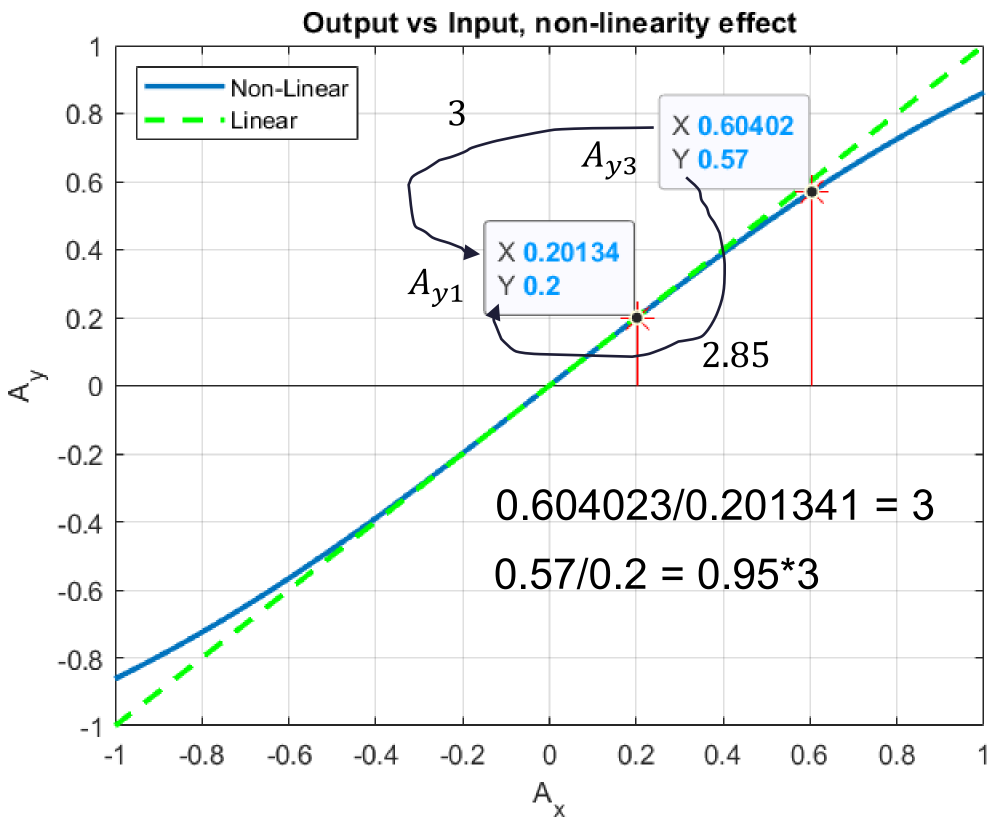

TEF 2025 Ethernet for AI

RLM Margin Loss Derating in COM

- Performance is dominated by the outer EYEs
- COM accounts for RLM through linear scaling of the SBR (it derates A_x by the RLM reduction of the outer EYE)
  - Equivalent to signal amplitude being reduced by RLM
  - $A_y = RLM * A_x$

```
%% 93A.1.6 step c defines A_s %%
cursor = sbr(THIS.cursor_i);
THIS.A_p=sbr(sbr_peak_i);
THIS.A_s = param.R_LM*cursor/(param.levels-1);
```

From COM 4.12.0_beta1

TEF 2025 Ethernet for AI

More Accurate Derating for PAM4

- COM ignores the fact that the reduction of the outer EYEs and inner EYE are affected differently

- Consider PAM4 with desired levels of 1 and 3 and RLM = 0.95
  - Inner level = 1 and outer level = 0.95*3 = 2.85 to satisfy RLM definition
  - Decision threshold for outer EYE is half-way between levels, (1+2.85)/2 = 1.925
  - The error distance from decision threshold to outer level is 2.85-1.925 = 0.925
  - Ratio of outer EYE margin loss to inner EYE is 0.925/1

- Thus, one should derate the SBR by 0.925 for an RLM of 0.95 as opposed to derating by the 0.95 currently used in COM

- For PAM4 modulation, COM is optimistic by 0.95/0.925, or 0.2316dB

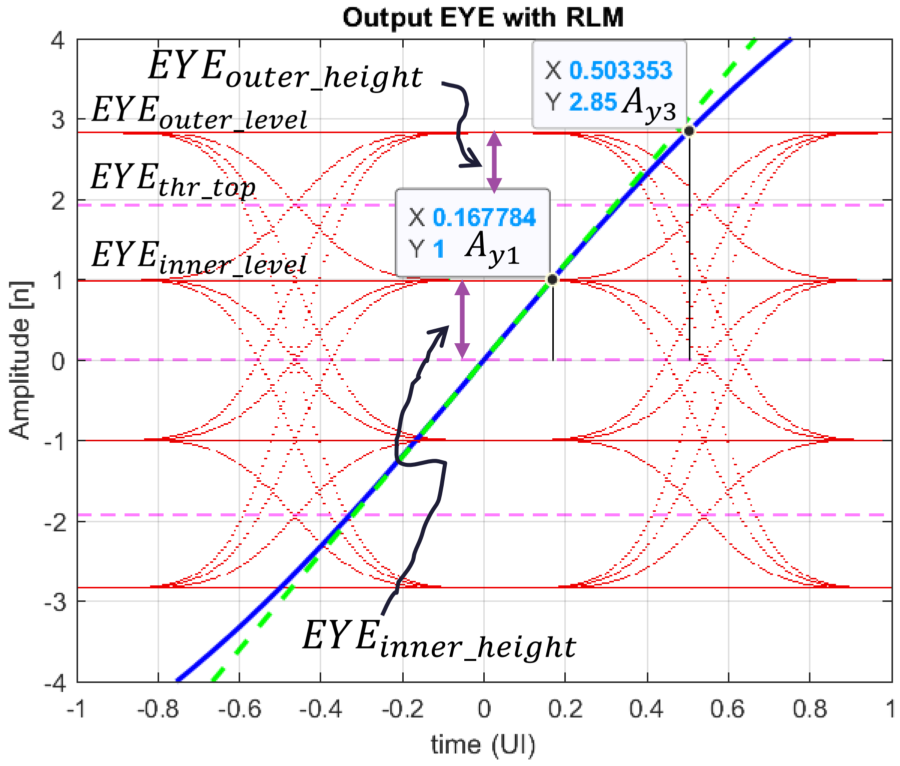

$$\frac{EYE_{outer\_level} - EYE_{thr\_top}}{EYE_{inner\_level} - EYE_{thr\_mid}} = \frac{(A_{y3} - (A_{y3} + A_{y1})/2)}{A_{y1} - 0}$$

TEF 2025 Ethernet for AI

More Accurate Derating for PAMm

- For the general PAMm case, the following equations hold where L ($L = PAMm - 1$), represents the PAMm outer level (L=3 for PAM4) and L-2 represents the next lower level (1 for PAM4)

$$EYE_{thr\_top} = (A_y(L) + A_y(L-2))/2$$

$$EYE_{outer\_height} = (A_y(L) - EYE_{thr\_top})$$

$$updated\_RLM\_derate = \frac{EYE_{outer\_level} - EYE_{thr\_top}}{EYE_{inner\_level} - EYE_{thr\_mid}} = \frac{(A_y(L) - A_y(L-2))/2}{A_y(1) - 0} \quad \text{[Eqn. 1]}$$

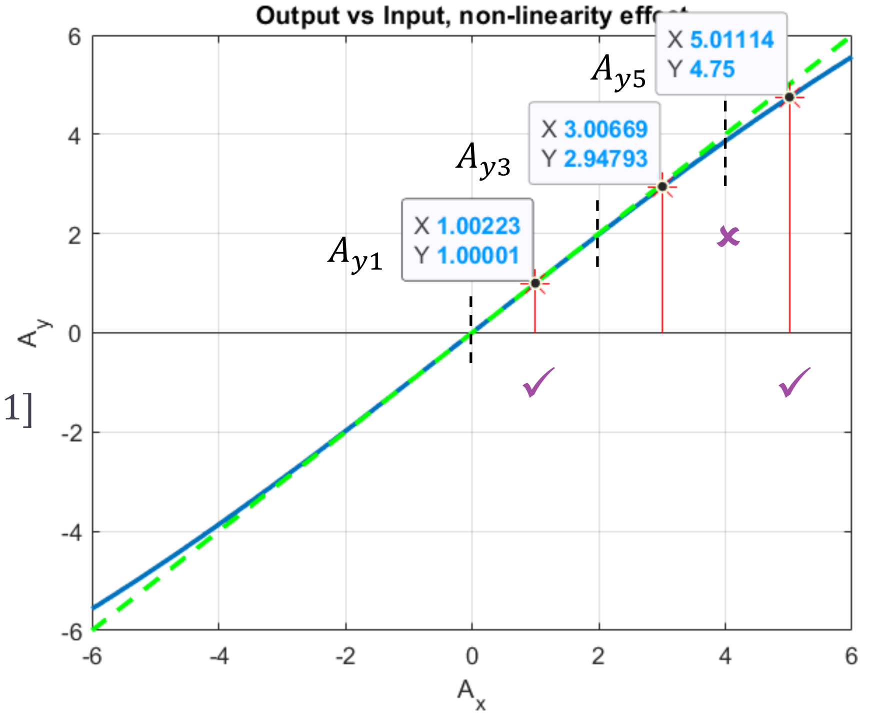

- For the PAM6 case at right, and with RLM = 0.95, one obtains

$$\frac{(A_{y5} - A_{y3})/2}{A_{y1} - 0} = \frac{(4.75 - 2.9479)/2}{1.0 - 0} = 0.9011$$

- One should derate the SBR by 0.9011 as opposed to 0.95 currently used in COM
- For PAM6 modulation, COM is optimistic by 0.95/0.9011, or 0.4594dB
  - At 400G/lane, we need that ~0.5dB of margin!"

TEF 2025 Ethernet for AI

Derating Based on a *tanh* Non-Linearity

- The derating factor is less intuitive to obtain when the modulation is higher than PAM4 because there is an intermediate level which is not directly determined via the RLM definition
- We therefore need to define a non-linear function governing the transmitter input to output relationship
- We propose using a ***tanh*** non-linearity of the form shown below, modelling a transmitter input-output relationship

$$A_y = \frac{1}{\tau} tanh[\tau A_x] \qquad \text{[Eqn. 2]}$$

- RLM is the actual ratio of PAM levels divided by the ideal ratio of PAM levels. From Eqn. 2 one obtains:

$$RLM = \frac{\frac{1}{\tau} tanh[\tau A_x]}{\frac{1}{\tau} tanh[\tau A_x / L]} / \frac{L}{1} \qquad \text{[Eqn. 3]}$$

Derating Based on a *tanh* Non-Linearity

- Solving for the term τA_x in Eqn. 3 using iterative processing (or in Matlab via MMSE optimization), one can solve for the term τA_x for any PAM at any desired RLM value

```
f = @(tauAx) (tanh(tauAx)/(L*tanh(tauAx/L)) - RLM).^2;
tauAx_o = fminsearch(@(tauAx) f(tauAx), 1.0);
```

[Eqn. 4]

- For the chosen ***tanh*** non-linearity model, using Eqn. 1 and Eqn. 2, one can obtain a value for the upper most two PAMm levels from which a derating factor can be computed as below:

$$updated\_RLM\_derate = \frac{\left(tanh\left[\tau A_{x\_o}\right] - tanh\left[\tau A_{x\_o} * (L-2)/L\right]\right)/2}{tanh\left[\tau A_{x\_o}/L\right] - 0}$$ [Eqn. 5]

TEF 2025
Ethernet for AI

Updated Derating vs COM Existing Derating

- At an RLM of 0.95, using Eqn.4 and Eqn. 5, one obtains the results tabulated below for PAM4, PAM6 and PAM8:

- Notice that as the modulation level increases, the existing COM derating factor based on RLM becomes more optimistic

PAMm	Existing A_x derate	Updated A_x derate	Existing to Updated derate ratio in dB	τA_x
4	0.95	0.9250	0.2316	0.4257
6	0.95	0.9011	0.4594	0.4083
8	0.95	0.8894	0.5722	0.4039

Example Through COM Analysis (default config)

- Run COM (V4.12.0_beta1) as is and with updated RLM derating
  - Bit rate set to 212.5Gbs (802.3dj)
  - Channel: host_pkg_top_500mm_max_skew_cable_module_pin_pad (mellitz_3dj_01_2409.pdf)
  - config_com_ieee8023_93a=df_200G_PAM4_fr55_C2M_TP1a_11_2022 .xlsx
- Agreement with above analysis is observed, but why the gap?

Package Type/loss ~53.125G	PAM4 COM (dB) existing	PAM4 COM (dB) updated	Delta (dB)
1/13.65dB	5.3440	5.1429	0.2011
2/17.37dB	5.3322	5.1294	0.2028
3/20.51dB	5.1844	4.9838	0.2006

Package Type/loss ~42.5G	PAM6 COM (dB) existing	PAM6 COM (dB) updated	Delta (dB)
1/10.71dB	2.2629	1.8478	0.4150
2/13.90dB	2.0621	1.6667	0.3954
3/16.48dB	1.9436	1.5455	0.3981

Example Through COM Analysis (with T_O = 0)

- Run COM (V4.12.0_beta1) as is and with updated RLM derating
  - Bit rate set to 212.5Gbs (802.3dj)
  - Channel: host_pkg_top_500mm_max_skew_cable_module_pin_pad (mellitz_3dj_01_2409.pdf)
  - config_com_ieee8023_93a=df_200G_PAM4_fr55_C2M_TP1a_11_2022 .xlsx
- Exact agreement with above analysis is observed

Package Type/loss ~53.125G	PAM4 COM (dB) existing	PAM4 COM (dB) updated	Delta (dB)
1/13.65dB	7.5497	7.3180	0.2316
2/17.37dB	7.3077	7.0761	0.2316
3/20.51dB	7.1314	6.8997	0.2316

Package Type/loss ~42.5G	PAM6 COM (dB) existing	PAM6 COM (dB) updated	Delta (dB)
1/10.71dB	4.1532	3.6942	0.4590
2/13.90dB	4.0140	3.5550	0.4590
3/16.48dB	3.8900	3.4268	0.4632

Matlab Function For Updated RLM Derating

```
function RLM_rescaled = RLM_rescale(pamM,RLM)

% this function rescales the RLM value to reflect the compression of the
% outer eye
%
% pamM is the PAM modulation value: 4,6,8,16,...
%
% RLM is the Ratio Leve Mismatch, mismatch of max level to innermost level

L = pamM-1;

%%%% based on the function Ay = (1/tau)*tanh(tau*Ax)
f = @(tauAx) (tanh(tauAx)/(L*tanh(tauAx/L)) - RLM).^2;      % tanh nonlinearity squared error
tauAx_o = fminsearch(@(tauAx) f(tauAx), 1.0);               % solve for MMSE factor for tanh at start point tauAx = 1.0

% tauAx_o is the tanh argument (tauAx) that yields the specified RLM
% outermost level is (1/tau)*tanh(tauAx), where outermost level
% next level below it is (1/tau)*tanh(tauAx*(L-2)/L)

% distance from the outermost symbol level to its associated threshold decision level
% is: outermost level - (outermost level + next level below)/2
% this is equivlant to (outermost level - next level below)/2

RLM_rescaled = (tanh(tauAx_o) - tanh(tauAx_o*(L-2)/L))/(2*tanh(tauAx_o/L));

% RLM_scaled is ratio of the distance of outermost symbol level to its associated
% threshold level to the distance of the innermost level to its associated
% threshold level (which is 0)

return
```

QUESTIONS?

Advancing 448G Signal Generation and Analysis with Modern Test & Measurement Solutions

Hadrien Louchet, Armands Ostrovskis, Fabio Pittala
Keysight Technologies

TEF 2025 Ethernet for AI

Agenda

- 448G generation & analysis - What is possible today?
- What are we missing to standardize PAM6?

General test setup

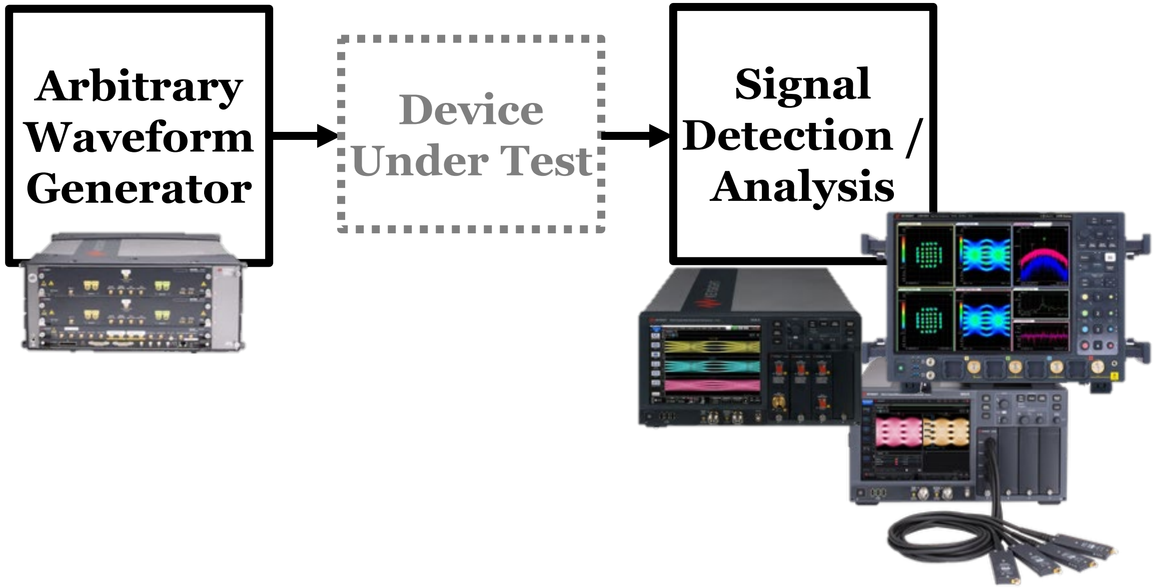

448G Generation – Enhancing SNR

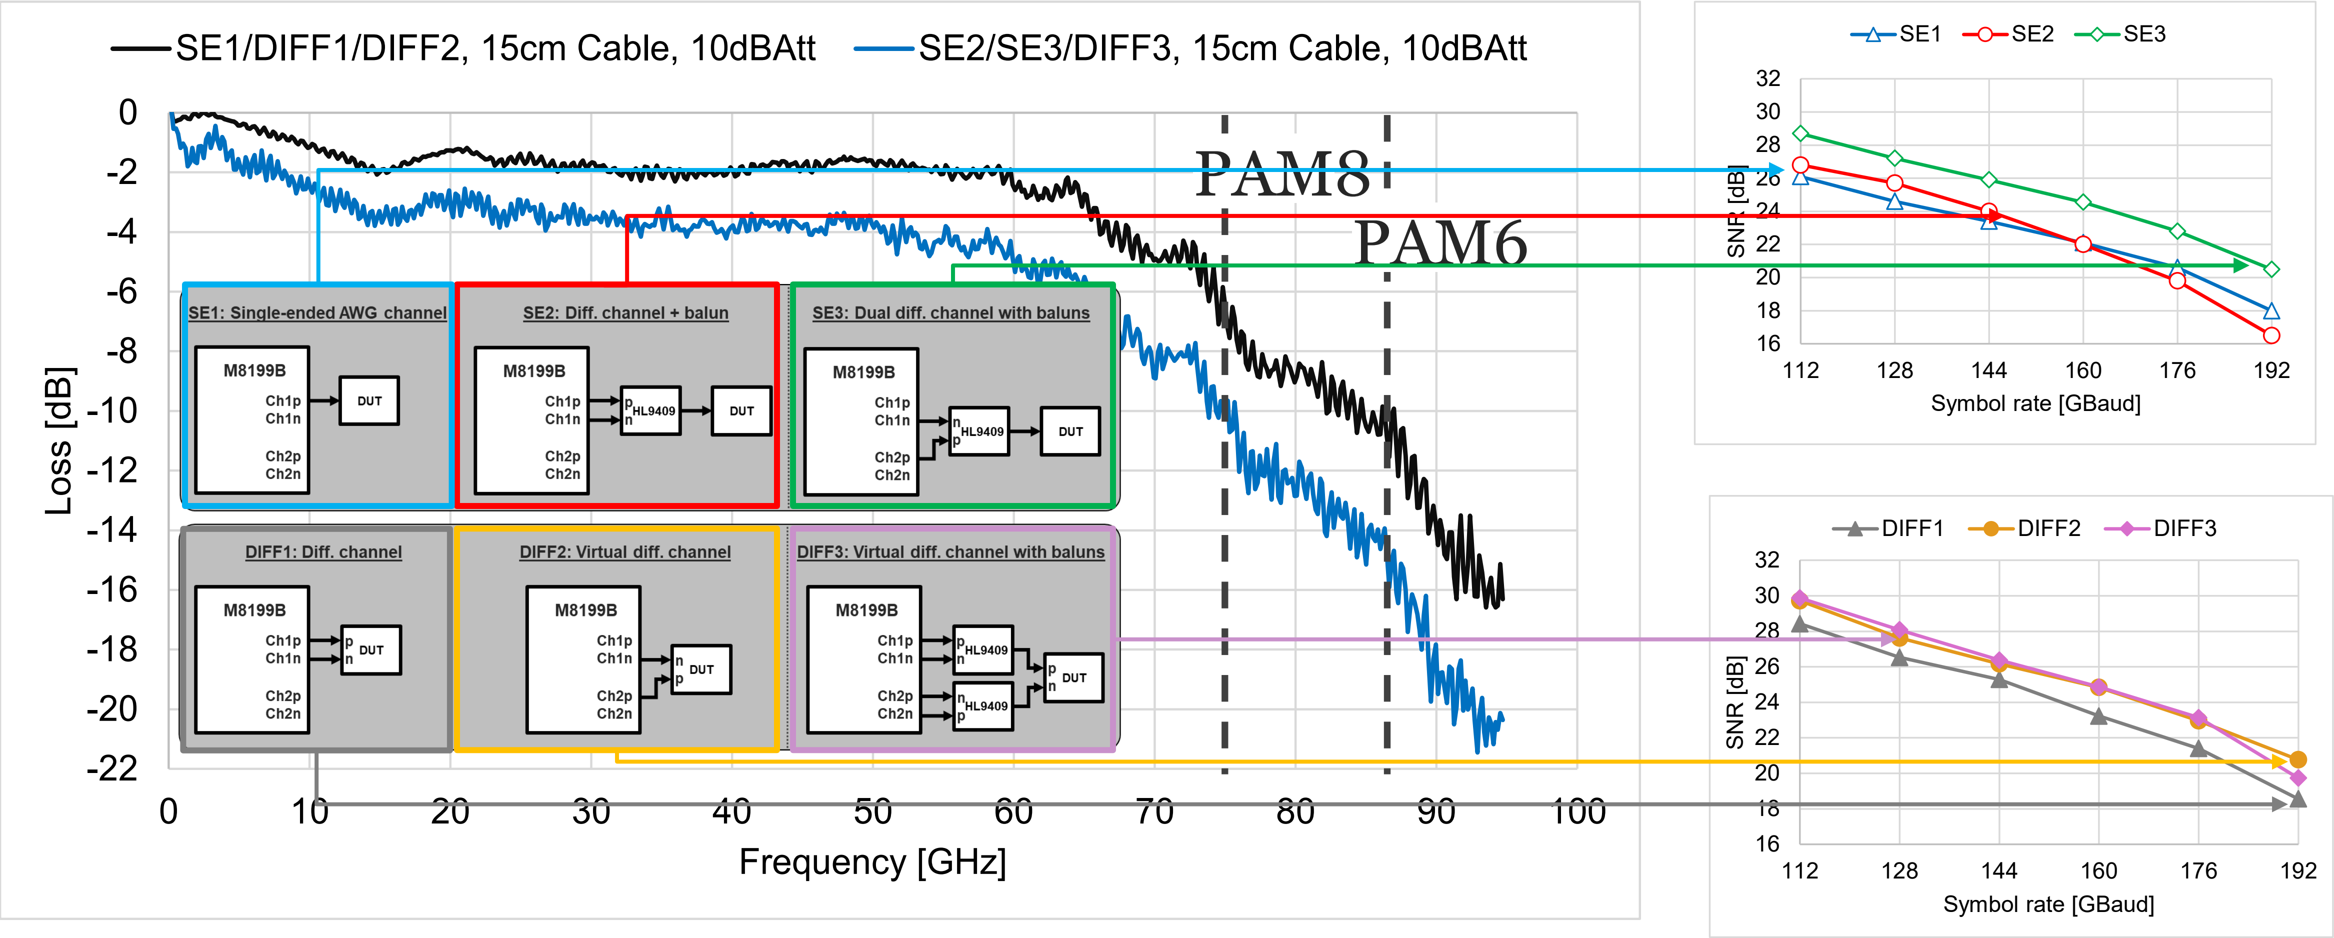

448G PAM4 Generation using Freq. Domain Interleaving

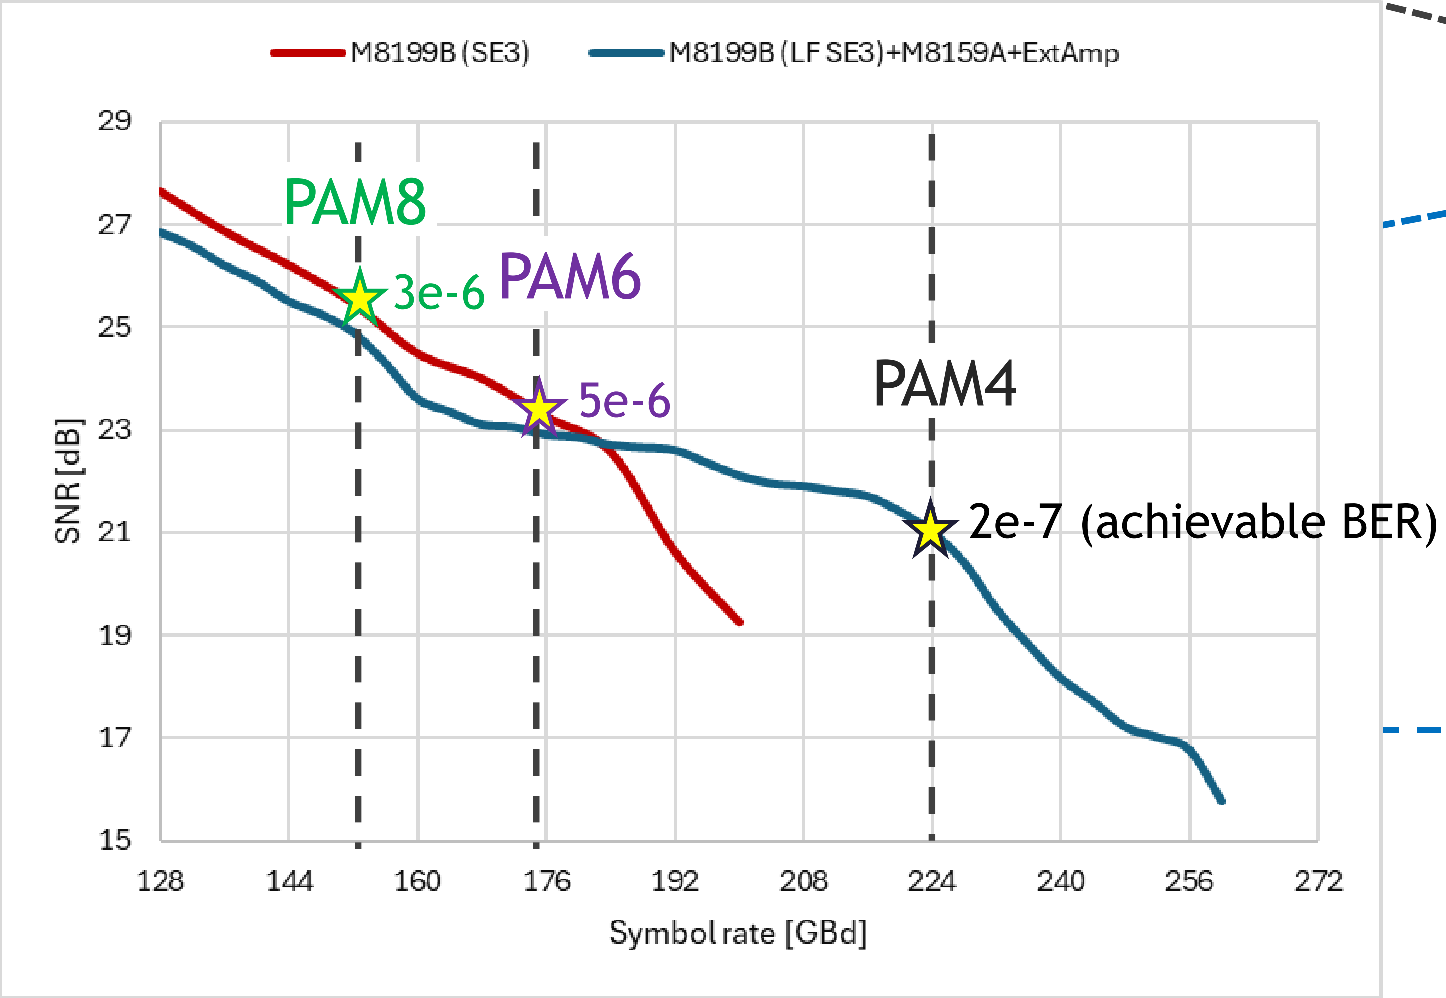

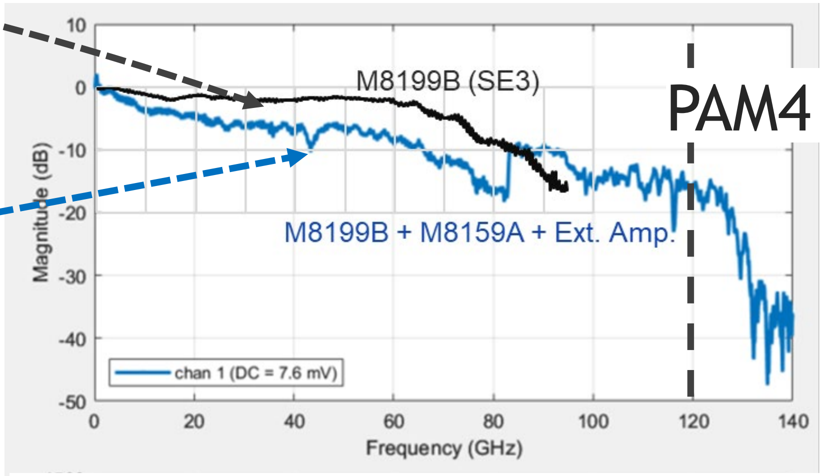

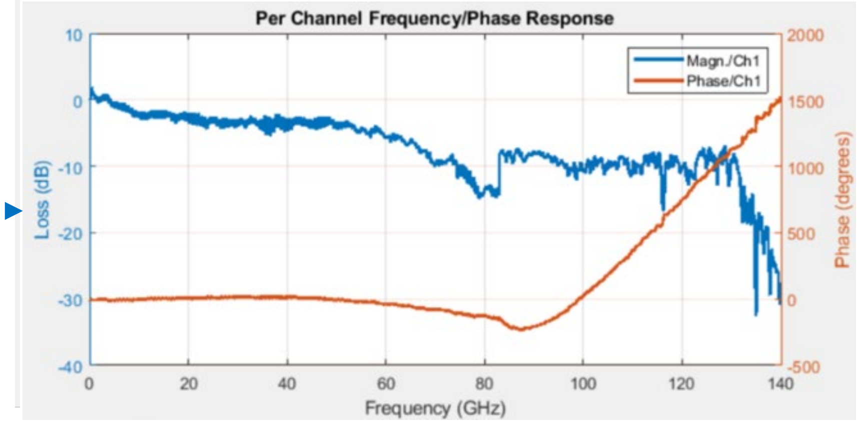

TEF 2
Etherne
AI

425G Electrical Link Performance

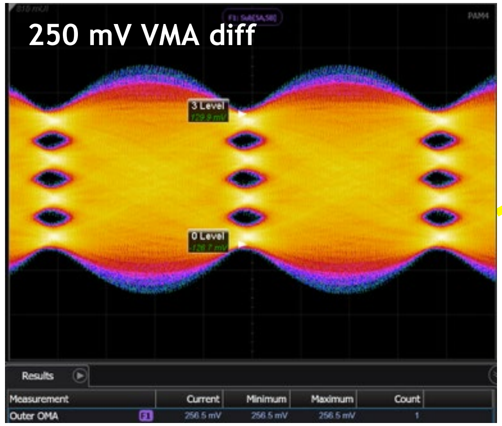

Measured: Waveform (425 Gbps)
SE+Balun or "true diff"

Ref Rx

- Input referred noise (eta_0= 1e-8 V²/GHz)

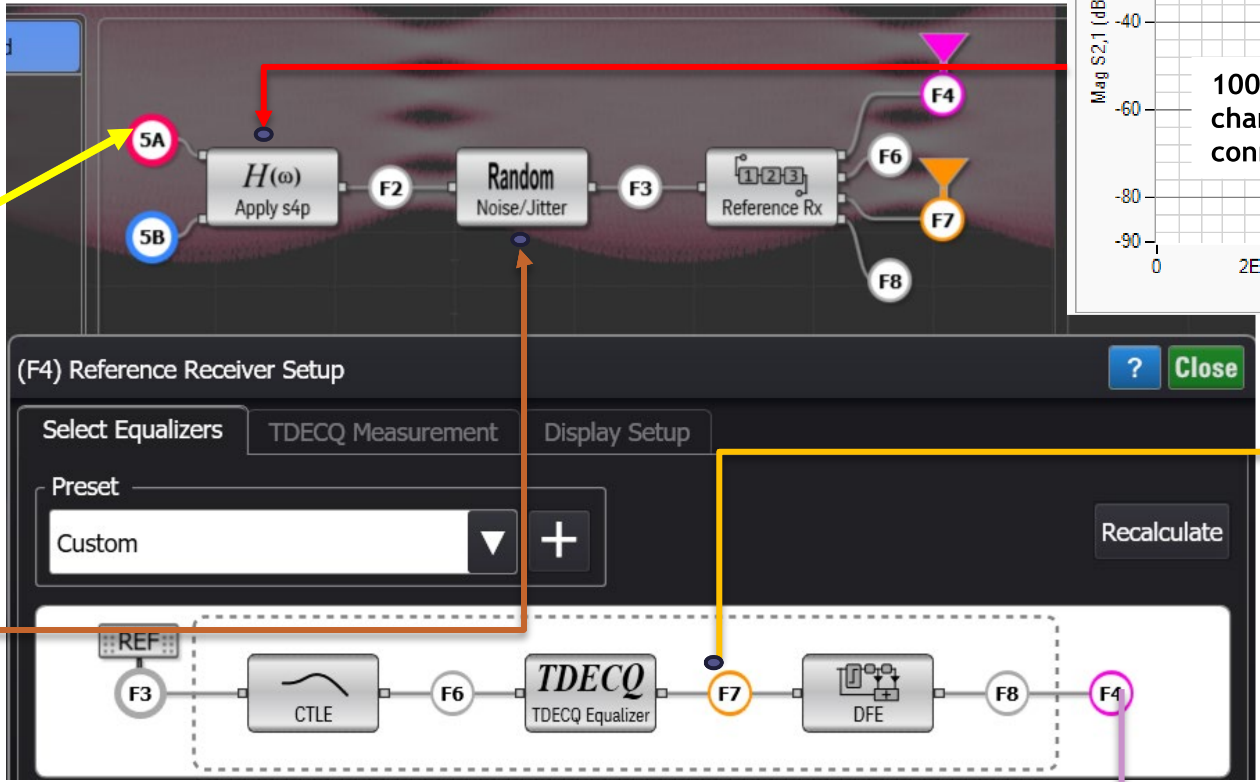

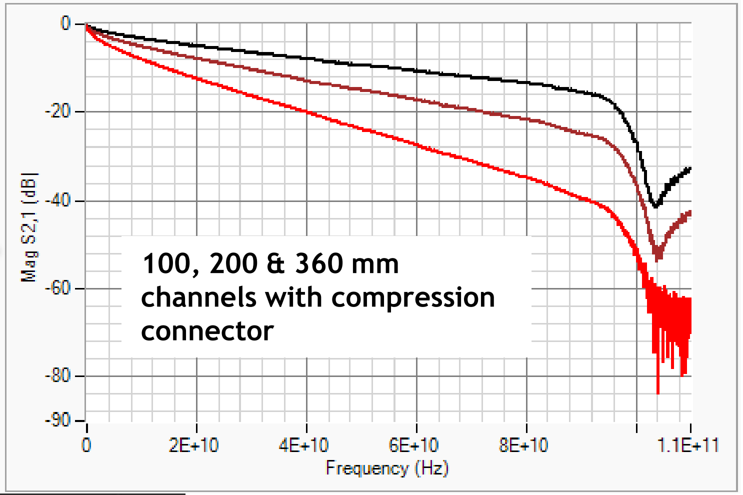

Channel S-Parameters to 110GHz

Matlab post processing

- PR-MLSD
- 7 tpbs pre-shapin)
- 1 symbol transistion memory
- 10 symbols trellis

DFE (for reference)

425G PAM4 Electrical Link Performance

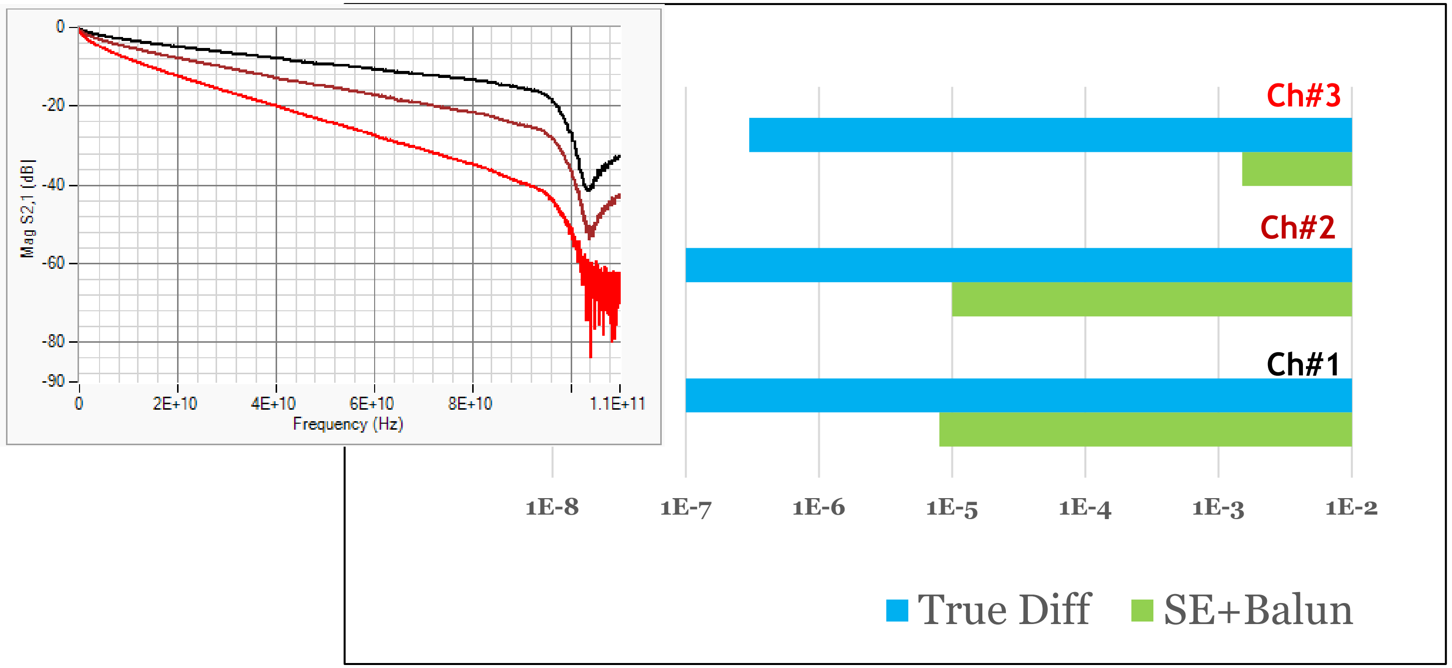

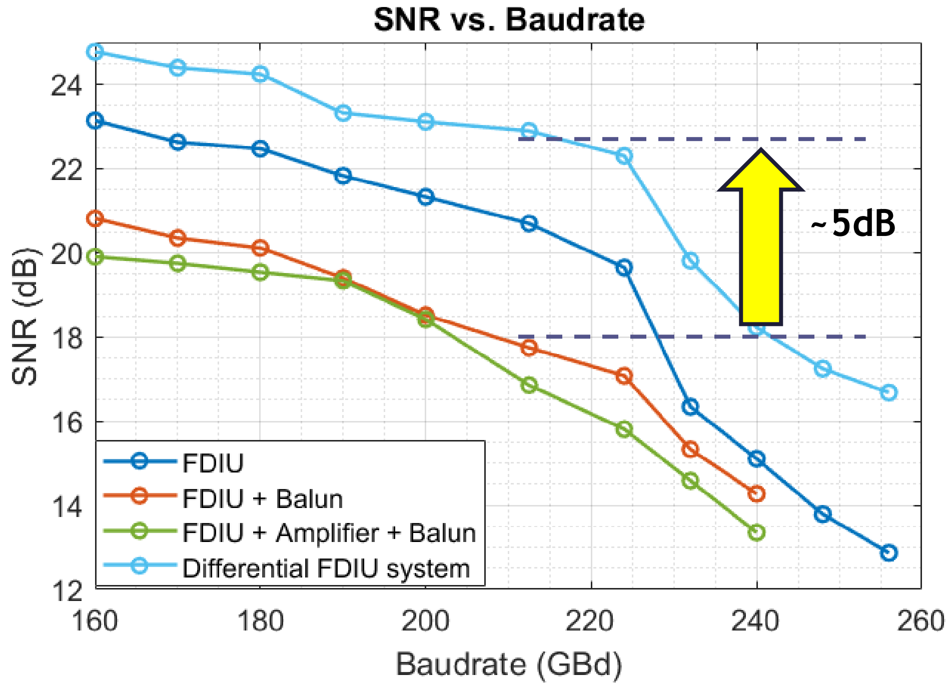

Differential FDIU system

- 2xFDI units
- 3x M8199B AWG

Balun-based system

- 1xFDI unit
- 2x M8199B AWG
- 1 RF amplifier
- 1 Balun

TEF 2025
Ethernet for AI

448G Analysis – PAM6 Toolkit

✓Equalization

✓BER/SER

✓SNDR

Possible (same as PAM4), but no reference pattern defined so far

○ Jitter measurement

Next step: „phase only“ for J_{nu} and J_{rms} to remove the impact of the channel on the jitter measurement[1]

174Gbaud PAM6 after 1.6T MTF (1mm)

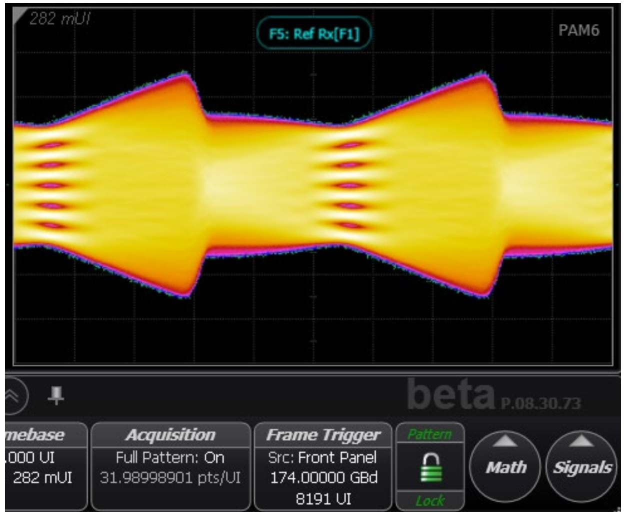

Measured: Channel S-Parameters to 110GHz (100, 200 & 400 mm with compression connector)

106.25Gbd PAm4, 31dB Channel

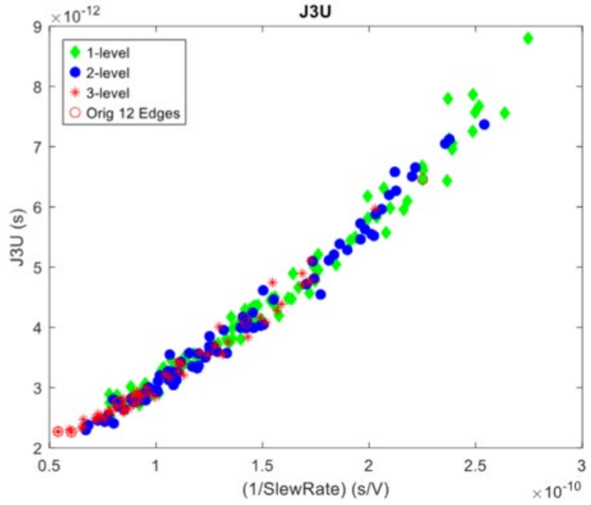

[1] D. Gines, IEEE 802.3dj workgroup

448G Analysis – PAM6 Toolkit

PAM6 Coding

- log2(6)≈2.585 bits/symbol
- 5B2S (2.5bit/ symbol)
- 18B7S (2.57 bits/ symbols)
  - (De)coding too complex

→ Multiple "Gray coding" approaches possible for **5B2S** [2]

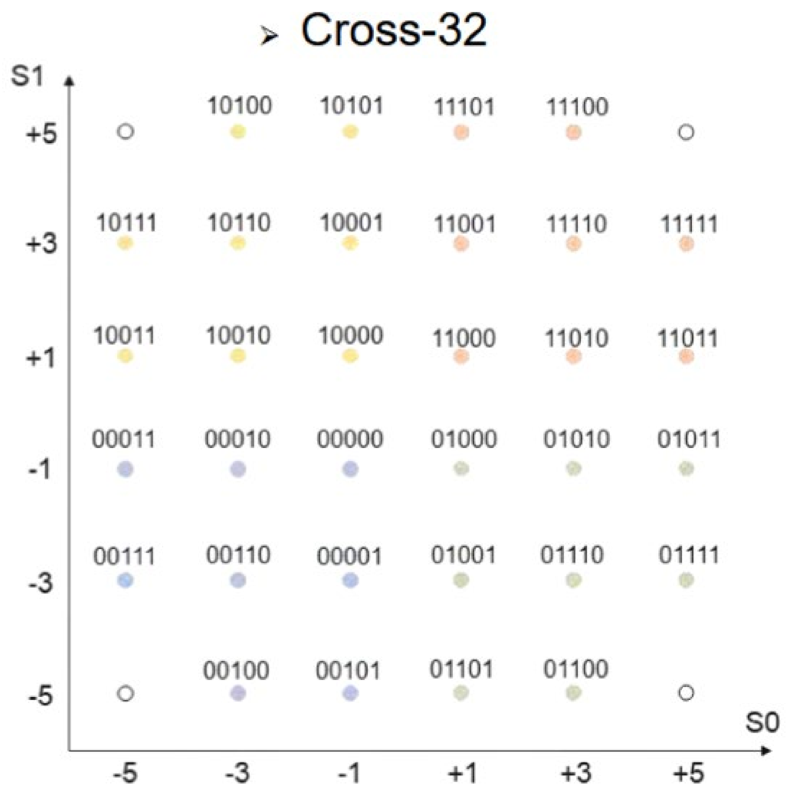

[2] Yu Xuet al, IEEE for AI, xu_e4ai_01_250624.pdf

TEF 2025 Ethernet for AI

448G Analysis – PAM6 Toolkit

Pseudo-random PAM6 sequences (PRPM6S)

- LFSR-based generator
  - Not possible as Galois Field, GF(n) exists only for $n = p^k$ with *p* prime and *k* integer (6=3x2)
  - PRBS-based PAM6 sequences lose some key characteristics

- Bruijn sequence B(k,n)
  - B(k,n) contains all n-tuples of an alphabet of size k
  - Requires non-linear feedback structure → complex to implement in ASICs – use (short) memory-based sequences instead

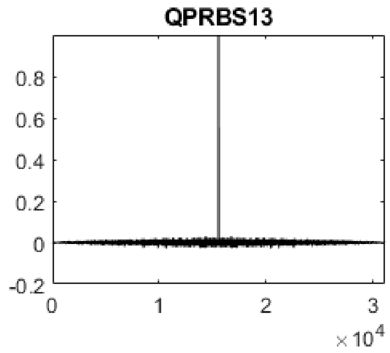

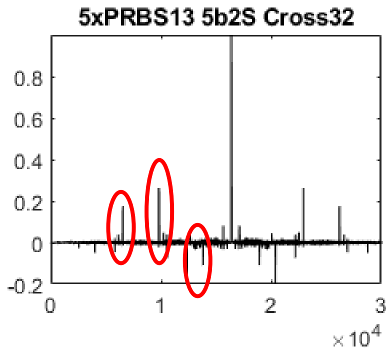

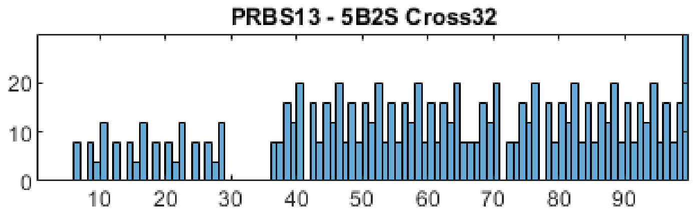

5-tuples of PAM6

Takeaways

- Current T&M solutions can support 448G pathfinding
  - Target SNR can be achieved for PAM4 (differential FDIU)
  - Improvement required for PAM6 and PAM8
- 448G electrical link performance
  - PAM4: Sub-Nyquist transmission possible (DFE, MLSD) but with penalty
  - PAM6 can accommodate “legacy” 1.6T channels (C2M)
  - Realistic scenario
    - PAM4 for CPO (XSR-like interface)
    - PAM6 for CPC (interface tbd)
- Industry to agree on PAM6 coding schemes & test patterns

QUESTIONS?

Unveiling Bottlenecks in Measured Co-Packaged Copper Channels: Sensitivity Analysis at ~~75–90~~ 100+ GBd for PAM6/4 ~~8~~

Jim Hsieh, Tobey P.-R. Li, ZZ Wu, Francis Lin, Howard Yin, Jarris Kuo - Mediatek

TEF 2025
Ethernet for AI

Agenda

- Background
- Channel Setup and Characteristic
- Simulation and Sensitivity Result
- Conclusion

Background

- Early 2025:
  - Modulations under discussion are PAM8, PAM6, PAM4 and bi-directional…
  - Channels under discussion are PCB, NPC(Near Package Copper) and CPC(Co-packaged copper)
- Near end of 2025:
  - Beginning to see early BER simulation result based on real CPC channel only.
  - 2025 OCP showcased a real co-packaged copper channel with 100Ghz+ BW.
- With 100Ghz+ CPC channel on the table. Now we should talk PAM4 vs PAM6.

TEF 2025 Ethernet for AI

Channel Setup Overview

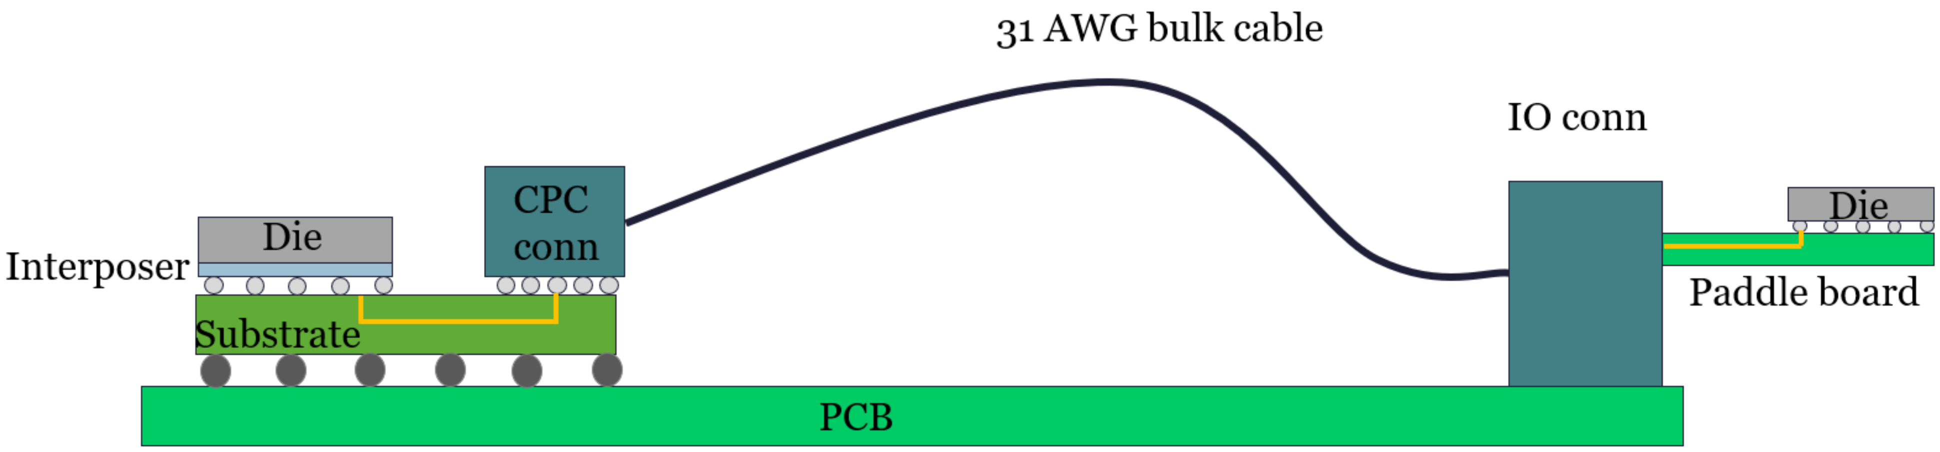

CH1: 20mm PKG + CPC conn[1][2] + 200mm 31AWG bulk cable[1] + IO conn[2][3] + 1 inch paddle board
CH2: 40mm PKG + CPC conn[1][2] + 600mm 31AWG bulk cable [1] + IO conn[2][3] + 1 inch paddle board

[1] Simulation model validated with correlated measurement data.
[2] Crosstalk included. 3 FEXT + 3 NEXT
[3] Simulation model. Measurement result are in-progress.

Courtesy of LUXSHARE-ICT, the channel models used in the following simulation results were provided by LUXSHARE-ICT

Channel Characteristic Of CPC

CH1

CH2

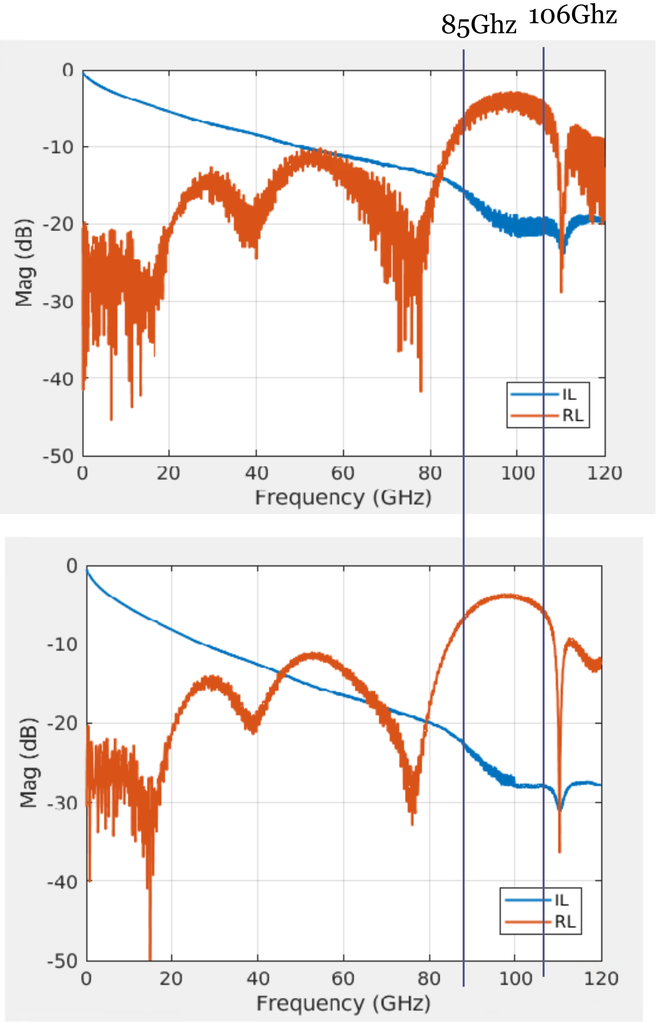

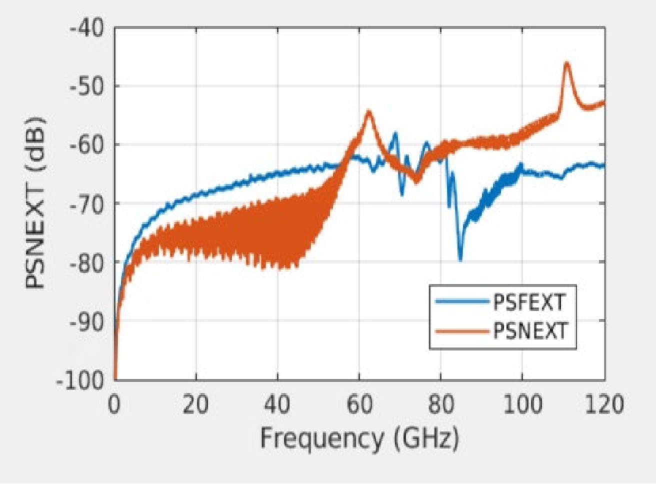

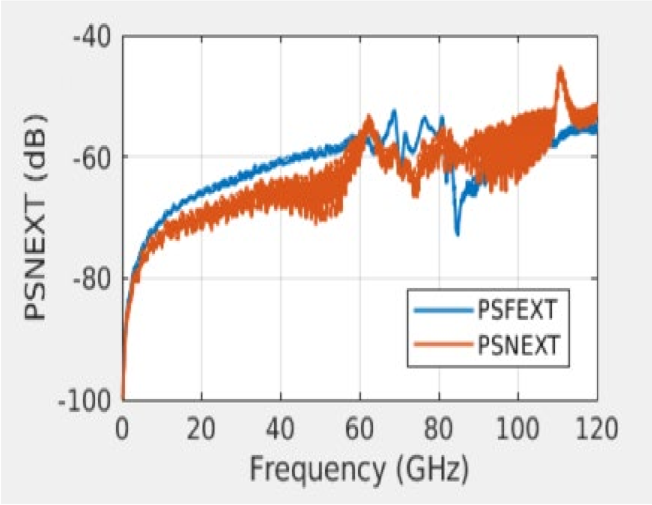

Both substrate and paddle board haven't been included.

Simulation Result

- Simulation based on COM v4.10, with BW and impairments scaled to match the baud rate. See table below for details.
- CH1 has sufficient margin to support both PAM4 and PAM6.
- CH2 requires 0.99dB and 1.56dB inner code to achieve a 3dB COM margin.
- PAM6 performs slightly better than PAM4 on both CH1 and CH2.

	Modulation	IL (dB)	COM (dB) Spec > 3dB
CH1	PAM4	26.2	4.28
	PAM6	20.5	4.51
CH2	PAM4	41.4	1.44
	PAM6	32.3	2.01

	Symbol	Unit	PAM6	PAM4
Target error ratio	DER_0		2e-5	2e-5
Baud Rate	f_b	GBd	170	212.5
Device	Scaled by Baud Rate			
PKG	Not scaled. Keep the same 802.3dj setting			
Transition Time	T_r	ps	2.5	2
SNDR	SNR_TX	dB	35	35
Jitter	A_DD	UI	0.02	0.02
	sigma_RJ	UI	0.01	0.01
CTLE	Scaled by Baud Rate			
Noise	eta_0	V^2/GHz	4.69e-9	3.75e-9
Quantization Bit	N_qb	-	7	7
Equalization			8 pre-tap 12 post-tap 3*4 floating-tap Span up to 80	8 pre-tap 16 post-tap 4*4 floating-tap Span up to 100
MLSD			Enabled	Enabled

Sensitivity Result Of PAM6 On CH1

- Starting point(leftmost) of x-axis corresponds to 802.3dj 200G baseline.
- Under the assumption that BW and impairments are scaled to match the baud rate, a COM value of 4.51dB is achieved for PAM6 on CH1.
- Example: at 425Gbps, a A_DD is expected to decrease from 0.02UI to 0.008UI. With COM value improving from 2.8dB to 4.51dB.
- Bandwidth, jitter and TX SNDR plays a significant roles in further scaling to support 400G.

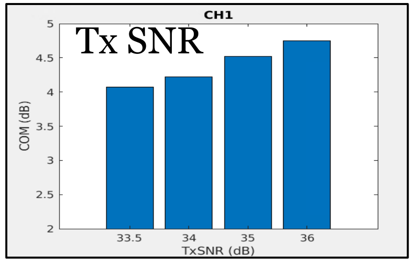

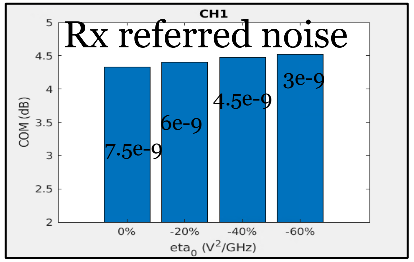

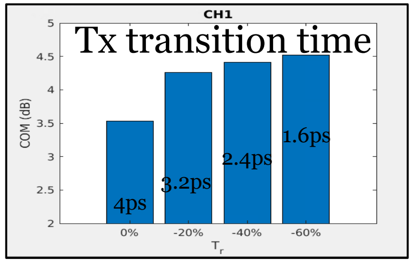

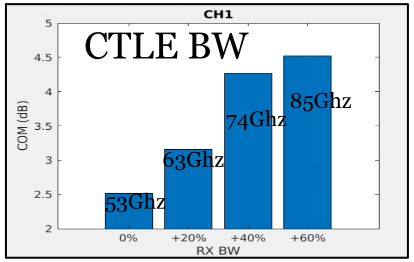

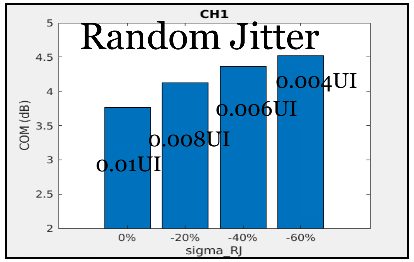

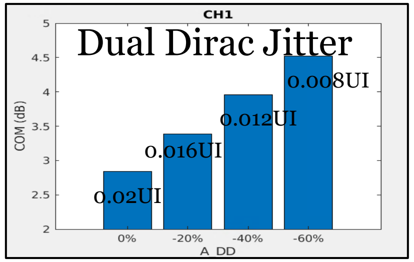

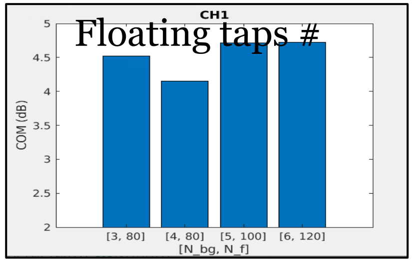

Conclusion

- 100G+ BW real CPC channel make PAM4 a viable candidate now.
- Inner FEC is required for longer channel setups.
- PAM6 still performs slightly better on both channel setups.
- Sensitivity results indicate bandwidth, jitter and TX SNDR plays a significant roles in further scaling to support 400G.
- More channel data is needed to determine whether to use PAM4 or PAM6.

QUESTIONS?

Modularity at the Package Edge: Composable Interconnect for 400 Gb/s/Lane AI Systems

Bijan Nowroozi, Head of Ecosystem Development, Lightmatter

TEF 2025
Ethernet for AI

Physics of 400G Pushes Back on System Design

- 224G+ lane rates exceed PCB reach limits
- LR SerDes dominate power + die area
- Shoreline bottleneck restricts port growth
- AI fabrics require exponential bandwidth scaling

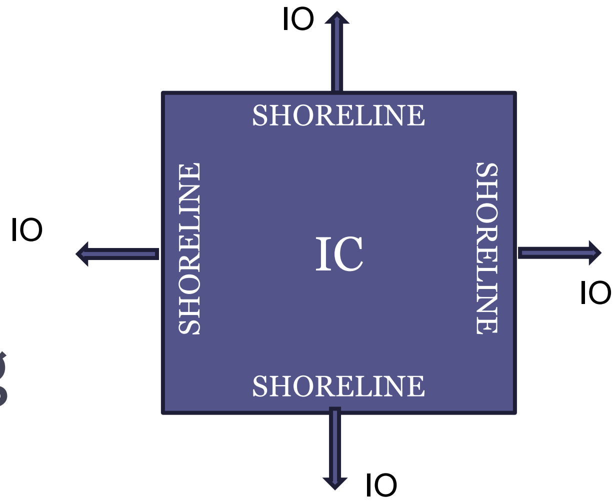

Road to Silicon Photonics : Technology Roadmap

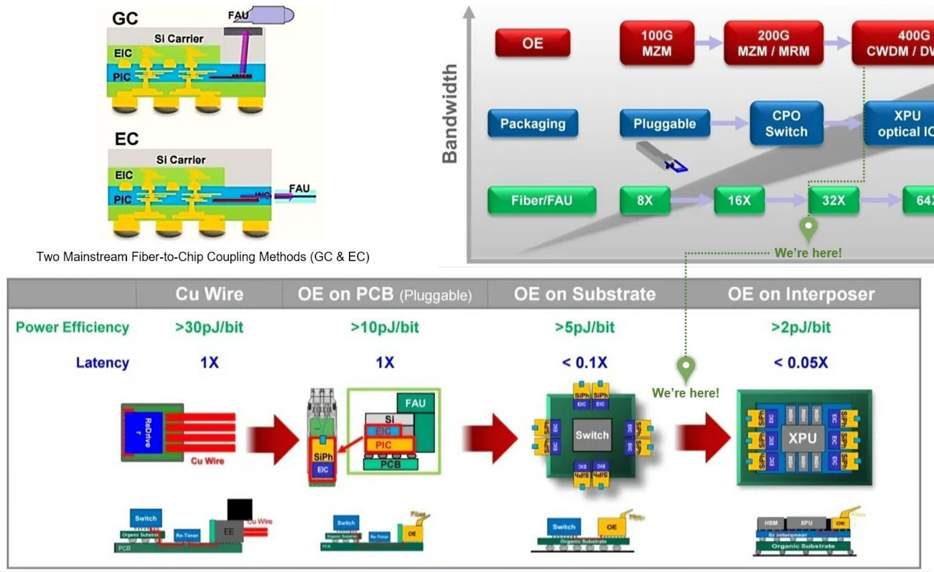

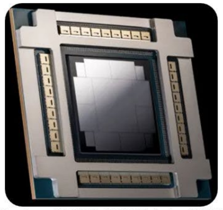

Spectrum-X® "SiPh SiP" (Shipping 2H'26)

Michael Liu

Source: TSMC, Nvidia EIC/PIC: Electronic/Photonic IC GC/EC: Grating/Edge Coupling OE: Optical Engine CPO: Co-Packaged Optics FAU: Fiber Array Unit

Silicon Photonics is promising to help with reach, and being shown everywhere, but the implementation?

TEF 2025 Ethernet for AI

Challenge

- Interconnect is evolving
- AI systems are evolving primarily around ASIC/xPU
- Ecosystem fracturing
- How to put evolutionary trends on same path, grow ecosystem?

www.ethernetalliance.org

2019: OCP Open Accelerator Specification

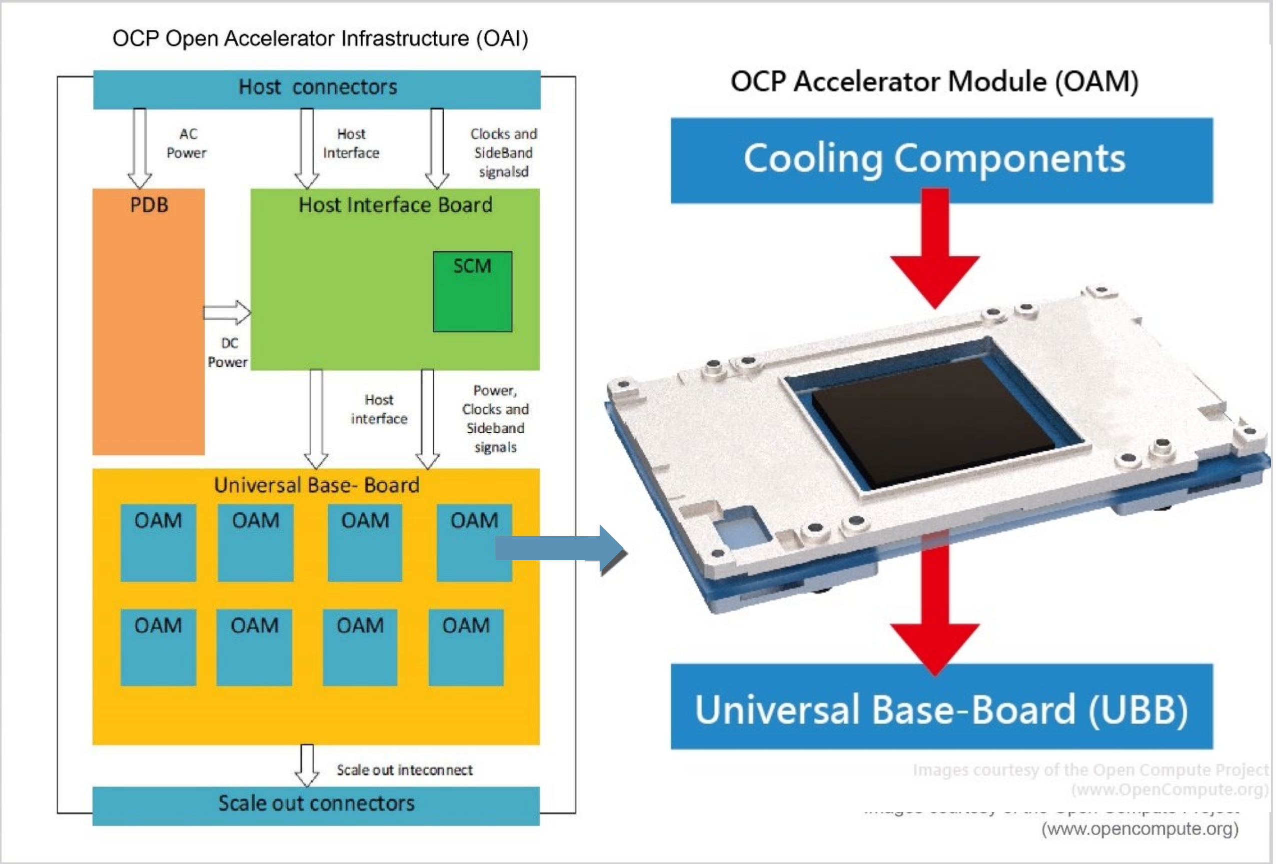

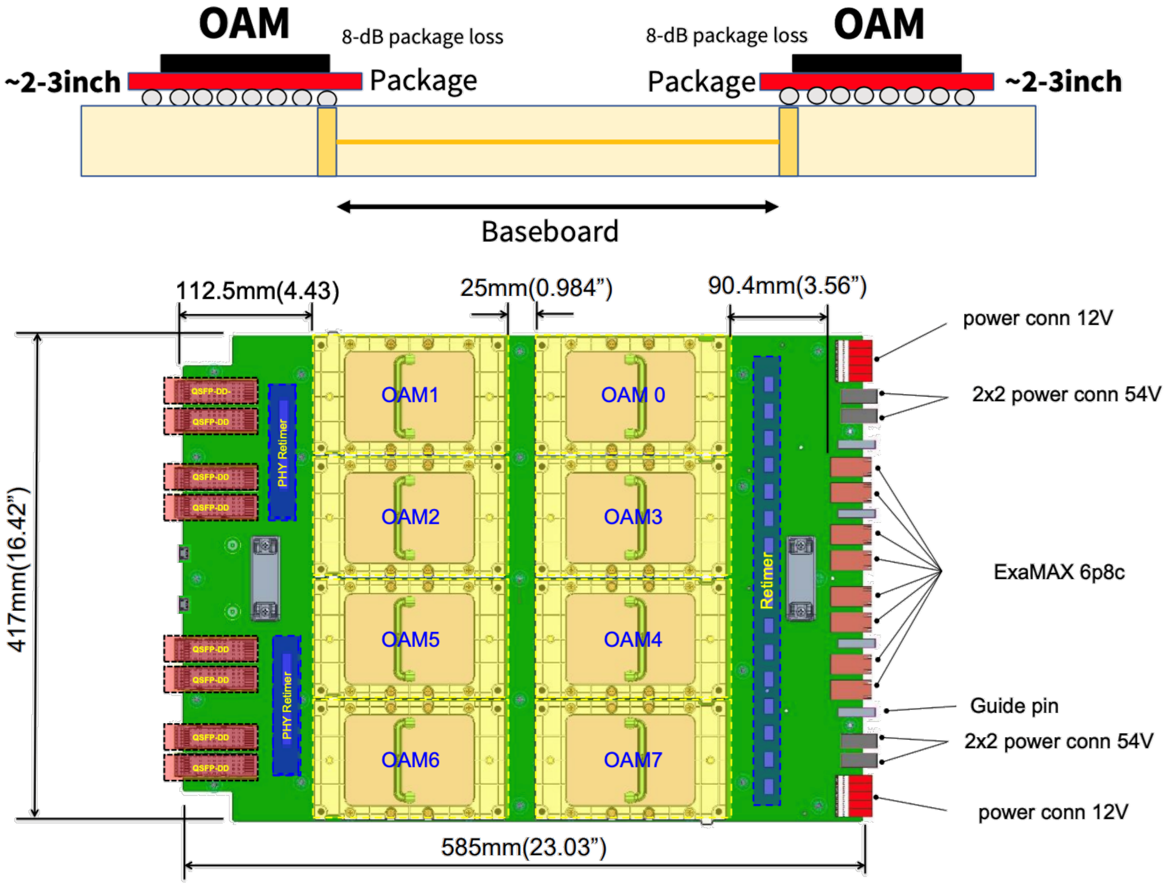

OCP Accelerator Module Spec Detail

- 102mm x 165mm Module Size
- With two high-speed Mirror Mezz connectors (MPN: 2093111115)
- 12V and 48V input DC Power
- Up to 350w (12V) and up to 700w (48V) TDP
  - Up to 440W (air-cooled) and 700W (liquid-cooled)
- Support single or multiple ASIC(s) per Module
- Up to eight x16 Links (Host + inter-module Links)
  - Support one or two x16 High speed link(s) to Host
  - Up to seven x16 high speed interconnect links
- System management and debug interfaces

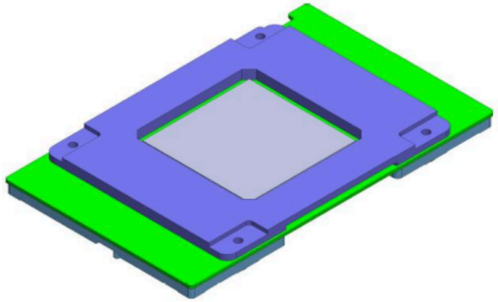

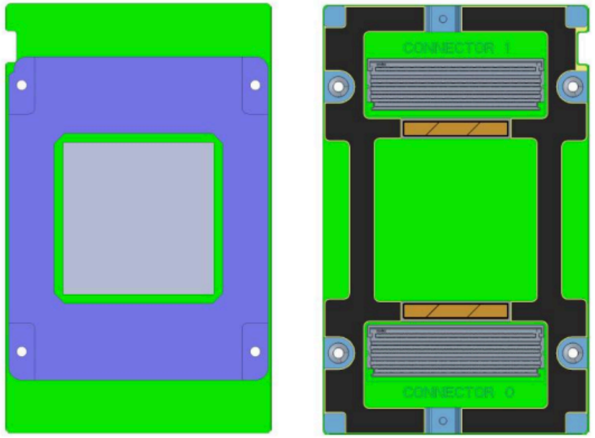

OAM Based Systems Seeded the Market for AI

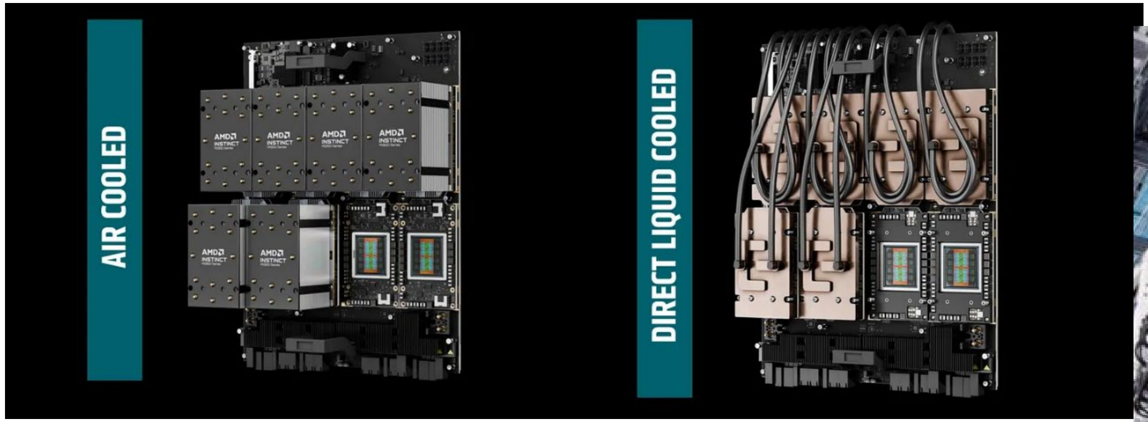

AMD Mi350 OCP OAI Systems

NVIDIA HGX

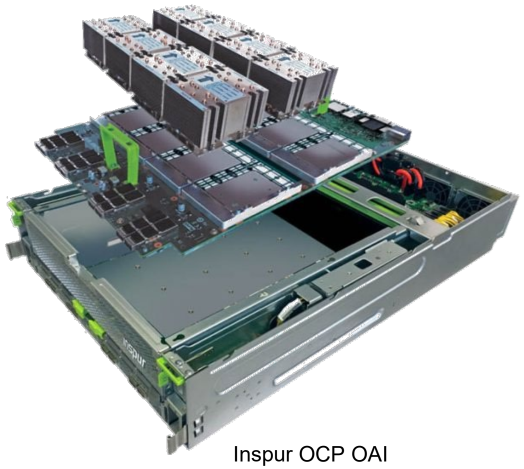

Inspur OCP OAI

Supermicro OCP OAI H100 / xAI Colossus

www.ethernetalliance.org

System Divergence

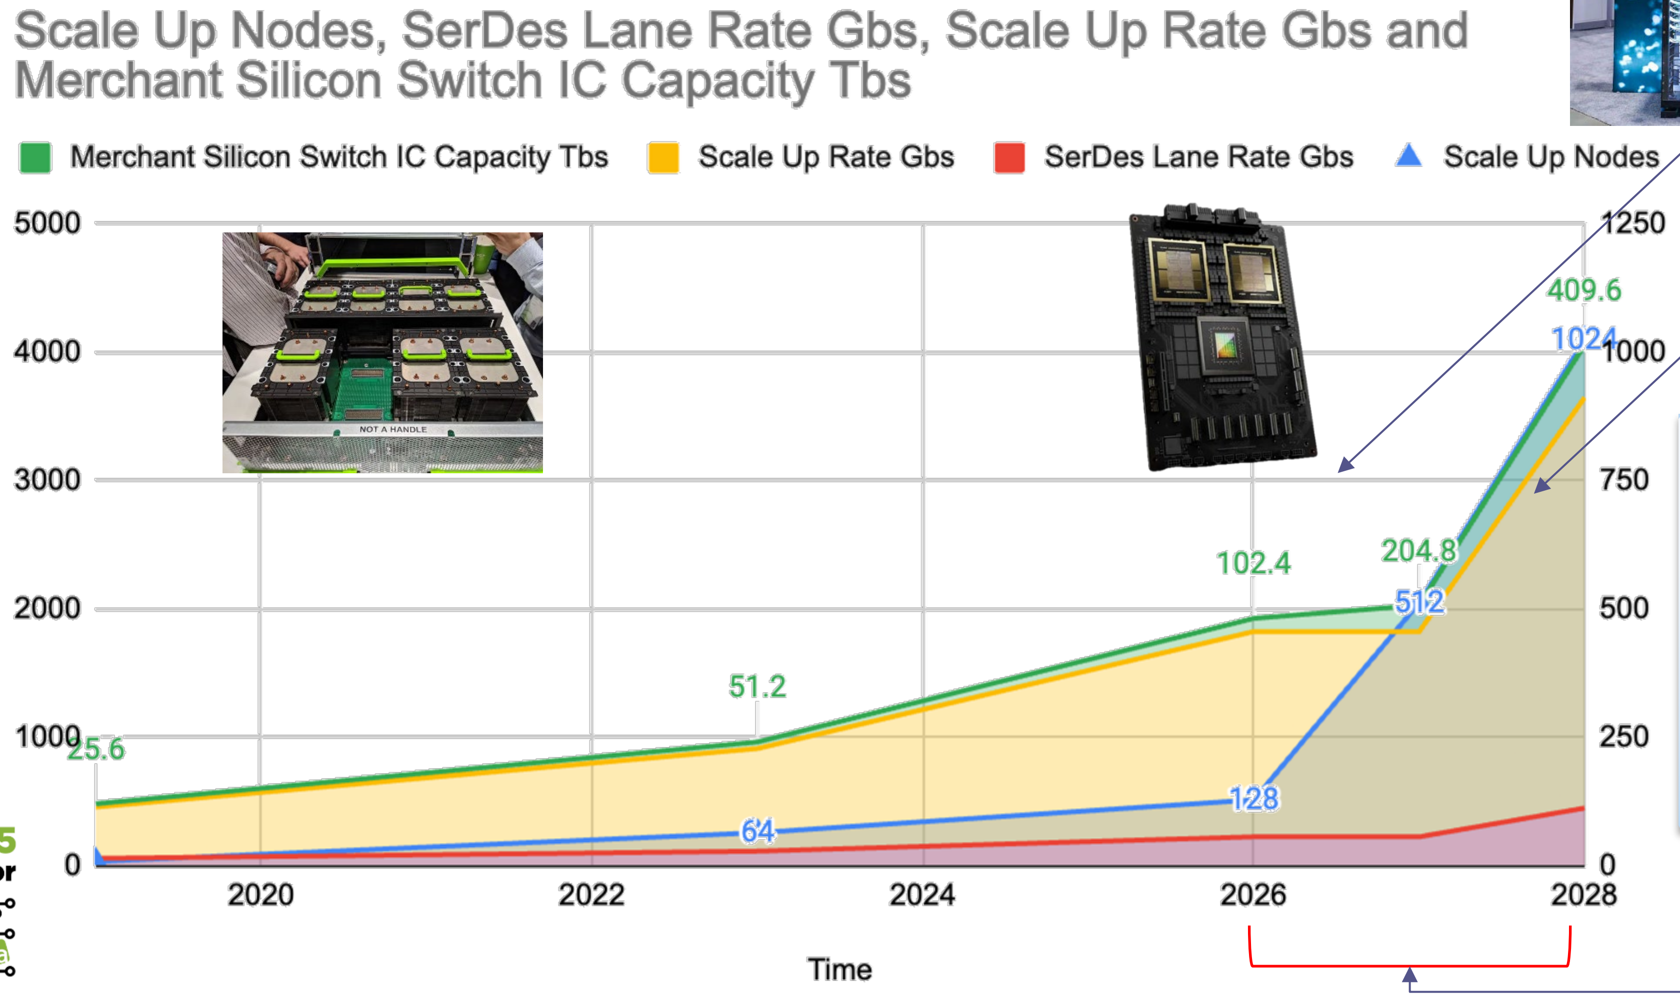

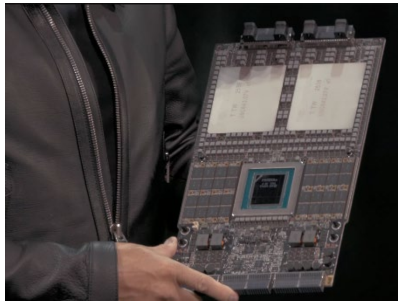

New Accelerators

New Interconnect standards

New Interconnect modalities

TEF 2025
Ethernet for AI

Meeting 400G Today: Composable IO

Define a stable boundary at the ASIC/package edge for IO composability:

- **Utility Plane** for power, control, telemetry
- **Data Plane** for high-bandwidth media (copper → optics)

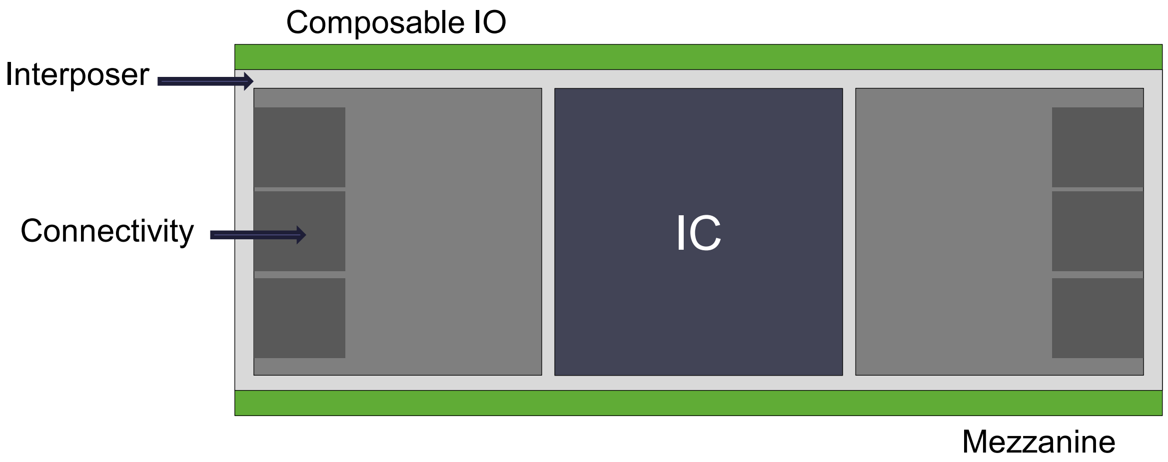

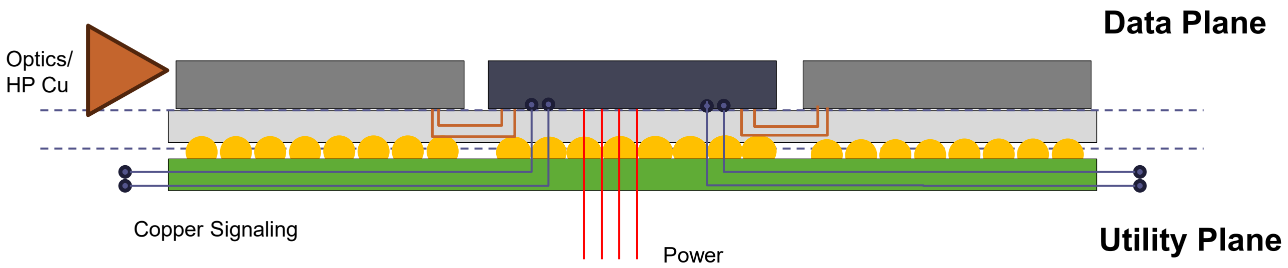

Architecture: Multi Plane Stack

The NG OAM defined not as a PCB, but as a Hybrid 3D Stack:

- **Layer 1** (Base): The Utility RDL
  - Function: Power Delivery, Ground, Sideband (I2C/GPIO), and PCIe Control.
- Routing: Adapts specific ASIC bump maps to a standard OAM pinout.
- **Layer 2** (Middle): Active High Rate Data Interposer
  - Function: Handles all >200G Data Traffic.
  - Size/space specified for modularity
  - Routing: Any-to-Any optical mesh.
- **Layer 3** (Top): Compute & Memory
  - Function: Logic (ASIC) and HBM.

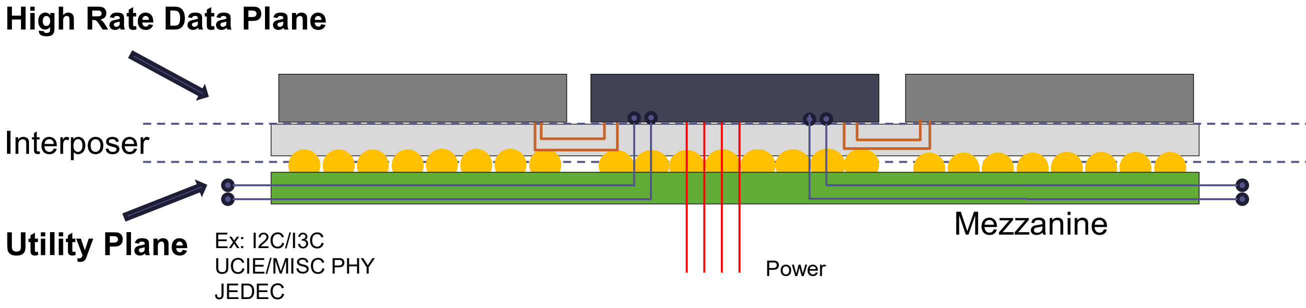

TEF 2025 Ethernet for AI

Physically: Building the NG OAM

- Lane Group Socket Organization
- Compute tile (single-multi reticle)
- Active/passive interconnect substrate
- Blind-mate optical/copper egress
- Utility layer mating to main board
- Thermal partitioning flexibility
- Pre integrated. Suppliers meet space/power etc Requirements

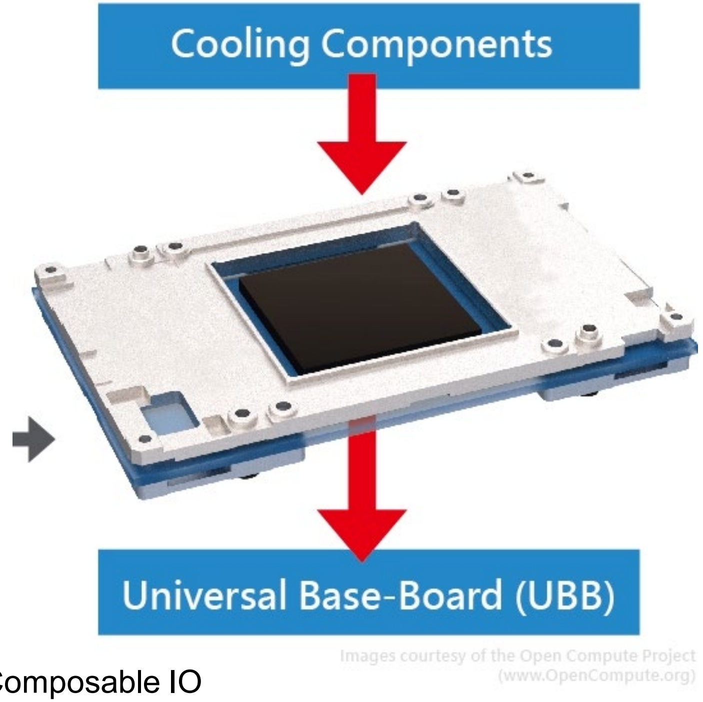

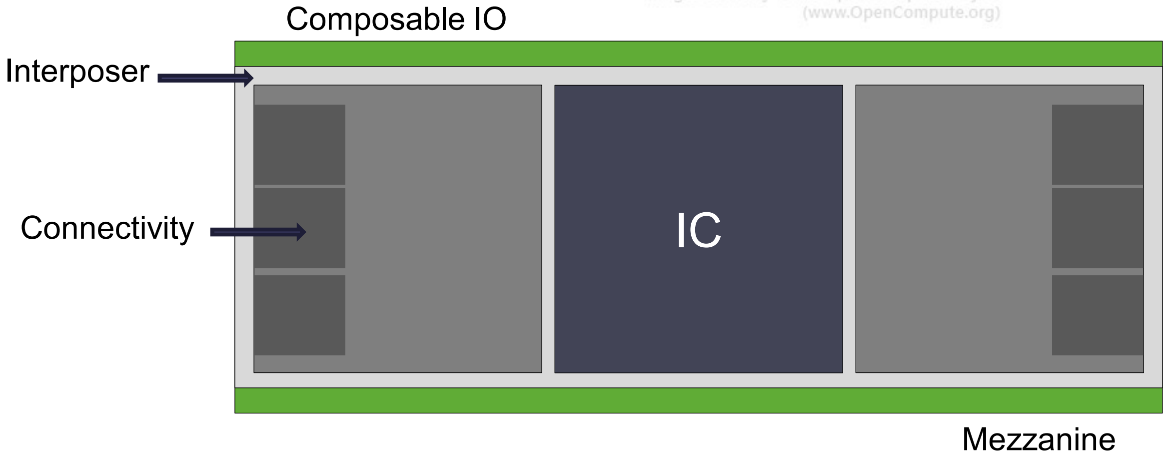

TEF 2025 Ethernet for AI

Abstraction Gain: Data and Utility Plane

Utility Plane

- RDL-based power & control layer
- Standardizable, slow-evolving
- LGA power delivery
- Health monitoring + sideband control

Data (HR IO) Plane

- Media-agnostic high-speed layer
- XSR/UCIe short reach
- Eliminates LR SerDes (20–30% die area reclaimed)
- Copper today, optical substrates tomorrow

Allows system developers to define system boards, services instead of focusing on silicon design.
Pathway to a single/few SKUs to support range of connectivity and ICs via composability.

Host Interface (The "Socket")

To support this architecture, system trays Universal Baseboard (UBB) is defined:

- Connectorless Power/Control: High-speed copper pins are removed. The module utilizes a high-density LGA (Land Grid Array) field for power and low-speed control only.
- Mechanical Clamping: A unified bolster plate and latching mechanism provides the necessary compression force for the LGA and thermal interface.
- Blind-Mate (Connectivity) Egress: The specification defines a "North/South" Keep-Out Zone (KOZ) at the module edge for floating connectors (e.g., MT ferrule arrays) that blind-mate when the module is clamped.

TEF 2025 Ethernet for AI

Benefits & Physics

- Adopting OAM NG yields immediate architectural advantages:

Metric	OAM v2.x (Current Copper)	OAM NG (Multi Plane)
PHY Technology	Long-Reach (LR) SerDes	**XSR / UCIe / VSR**
Shoreline Density	~0.5 Tbps/mm	**> 2.0 Tbps/mm (Reclaimed)**
Motherboard Material	Ultra-Low Loss (Expensive)	**Standard FR4 / Mid-Loss**
Power Overhead	High (PCB + Retimers)	**Low (Direct-Drive Optics)**

RAS & Testability

- Allows close following of OIF CMIS/OCP RAS guidelines
- Real-time telemetry via Utility Plane
- Connectors for field level maintainability
- Cartridge-style replacement → low MTTR

Industry Impact

- Stabilizes ASIC interfaces for a decade
- Reduces redesign cycles
- Enables media evolution without silicon changes
- Aligns OCP, {Ethernet Scaling SDOs), IEEE, OIF, TEF
- Stability: Abstraction boundary remains constant across generations

Closing

- **NG OAM socket:**
  - Composable IO allows choice through modularity
  - Defined as a scalable, open, media-agnostic foundation for next-gen AI infrastructure.
    - Build with 400Gbs/lane fabric at the system level, while avoiding silicon development risk
    - Design common system trays and elements
  - Enables a vibrant ecosystem

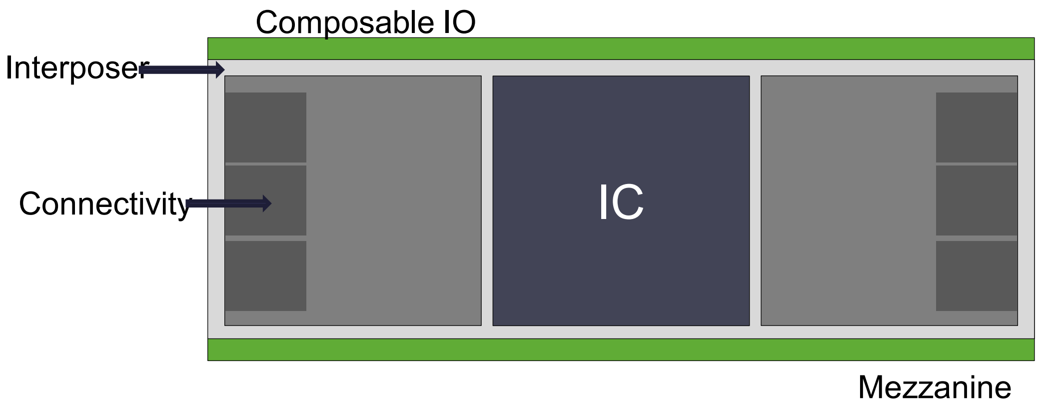

Call to Action

We invite the ecosystem to collaborate on defining the future composable IO socket Define the Split-Plane mechanical boundary

- Define requirements power/size/speeds and feeds
  - Thermal Impedance Targets
  - Keep out zones etc
- Standardize LGA + pin maps
- Define copper + optical data-plane classes
- Form joint OCP/TEF/OIF/SDO super alignment group

QUESTIONS?