



# 448G Signaling for AI


Curtis Ling and Sridhar Ramesh  
April 2025

# Outline

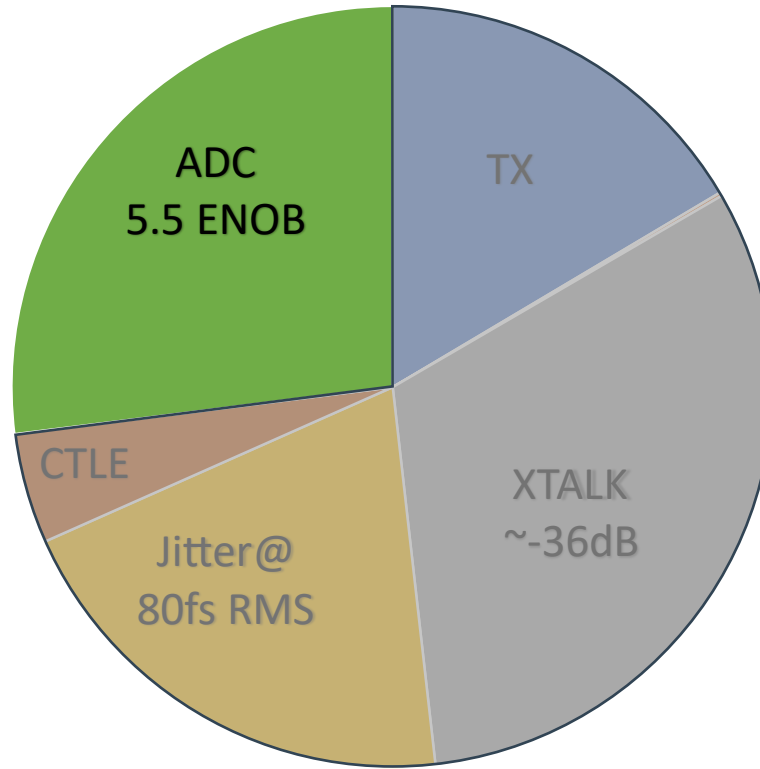
- ADC efficiency trades versus SNDR and Nyquist frequency for PAMx
- FEC options
- Advanced equalization and detection schemes
- Copackaged optics and copper (CPx)

# Electrical Link Budget: 448G PAM4/6/8 Noise

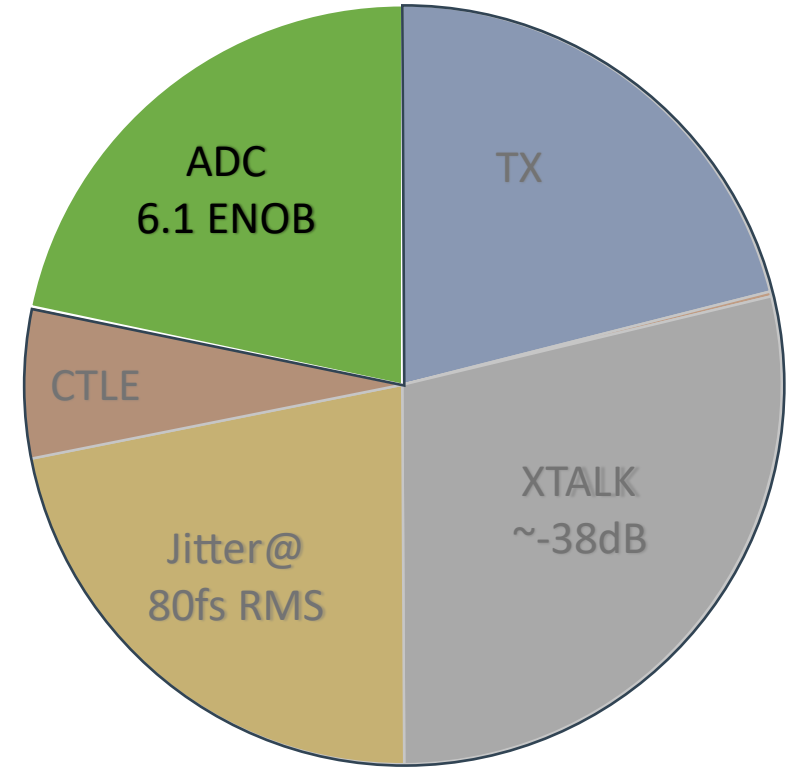
**PAM4, 212 GBaud**



**PAM6, 164 GBaud**

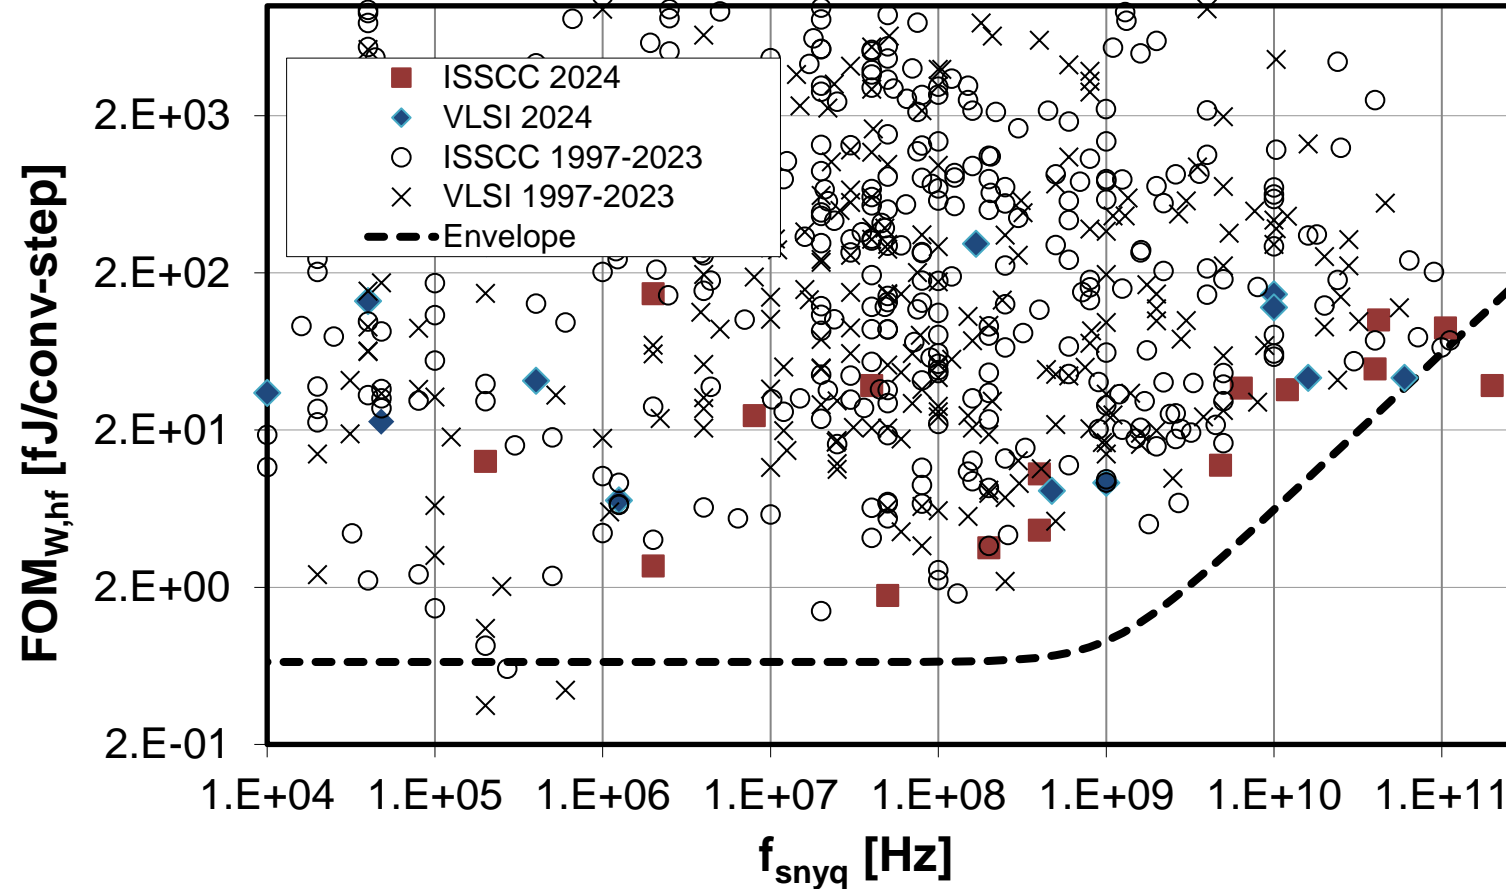


**PAM8, 142 GBaud**



- ENOB alone is not a useful metric, while the pJ/b tradeoff space is complex

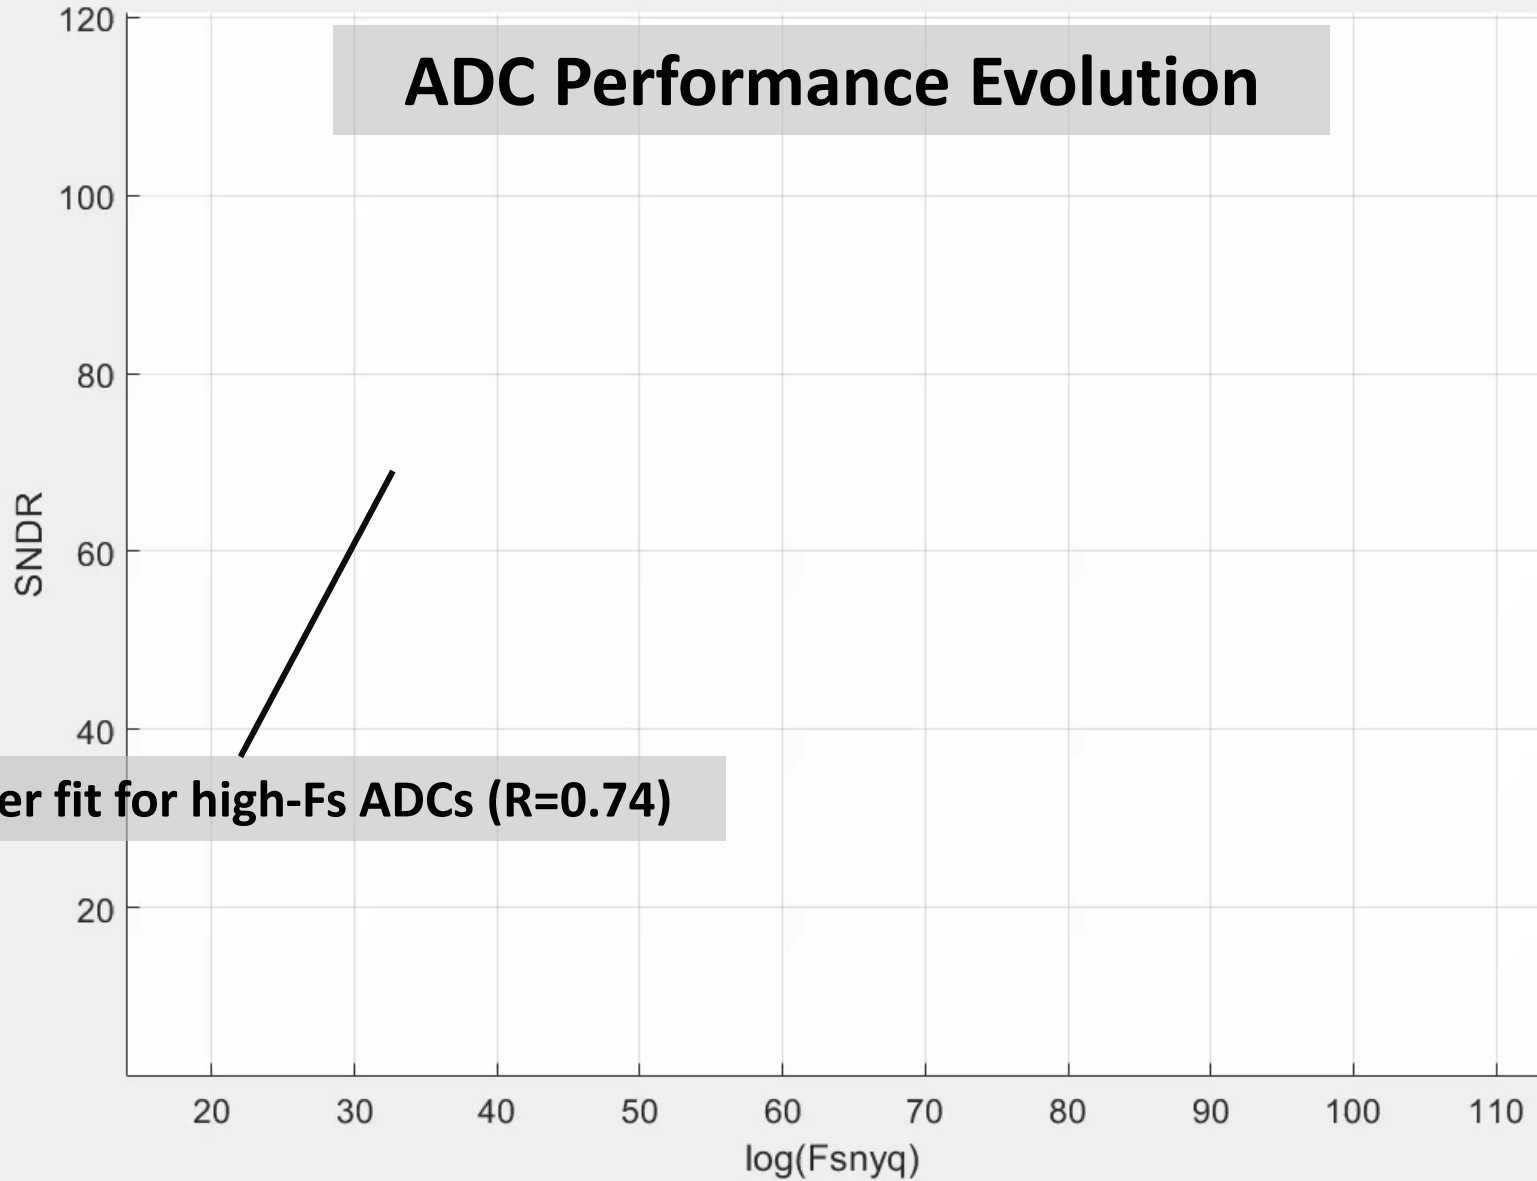
# Converter Tradeoff Data – Murmann ADC Survey[1]



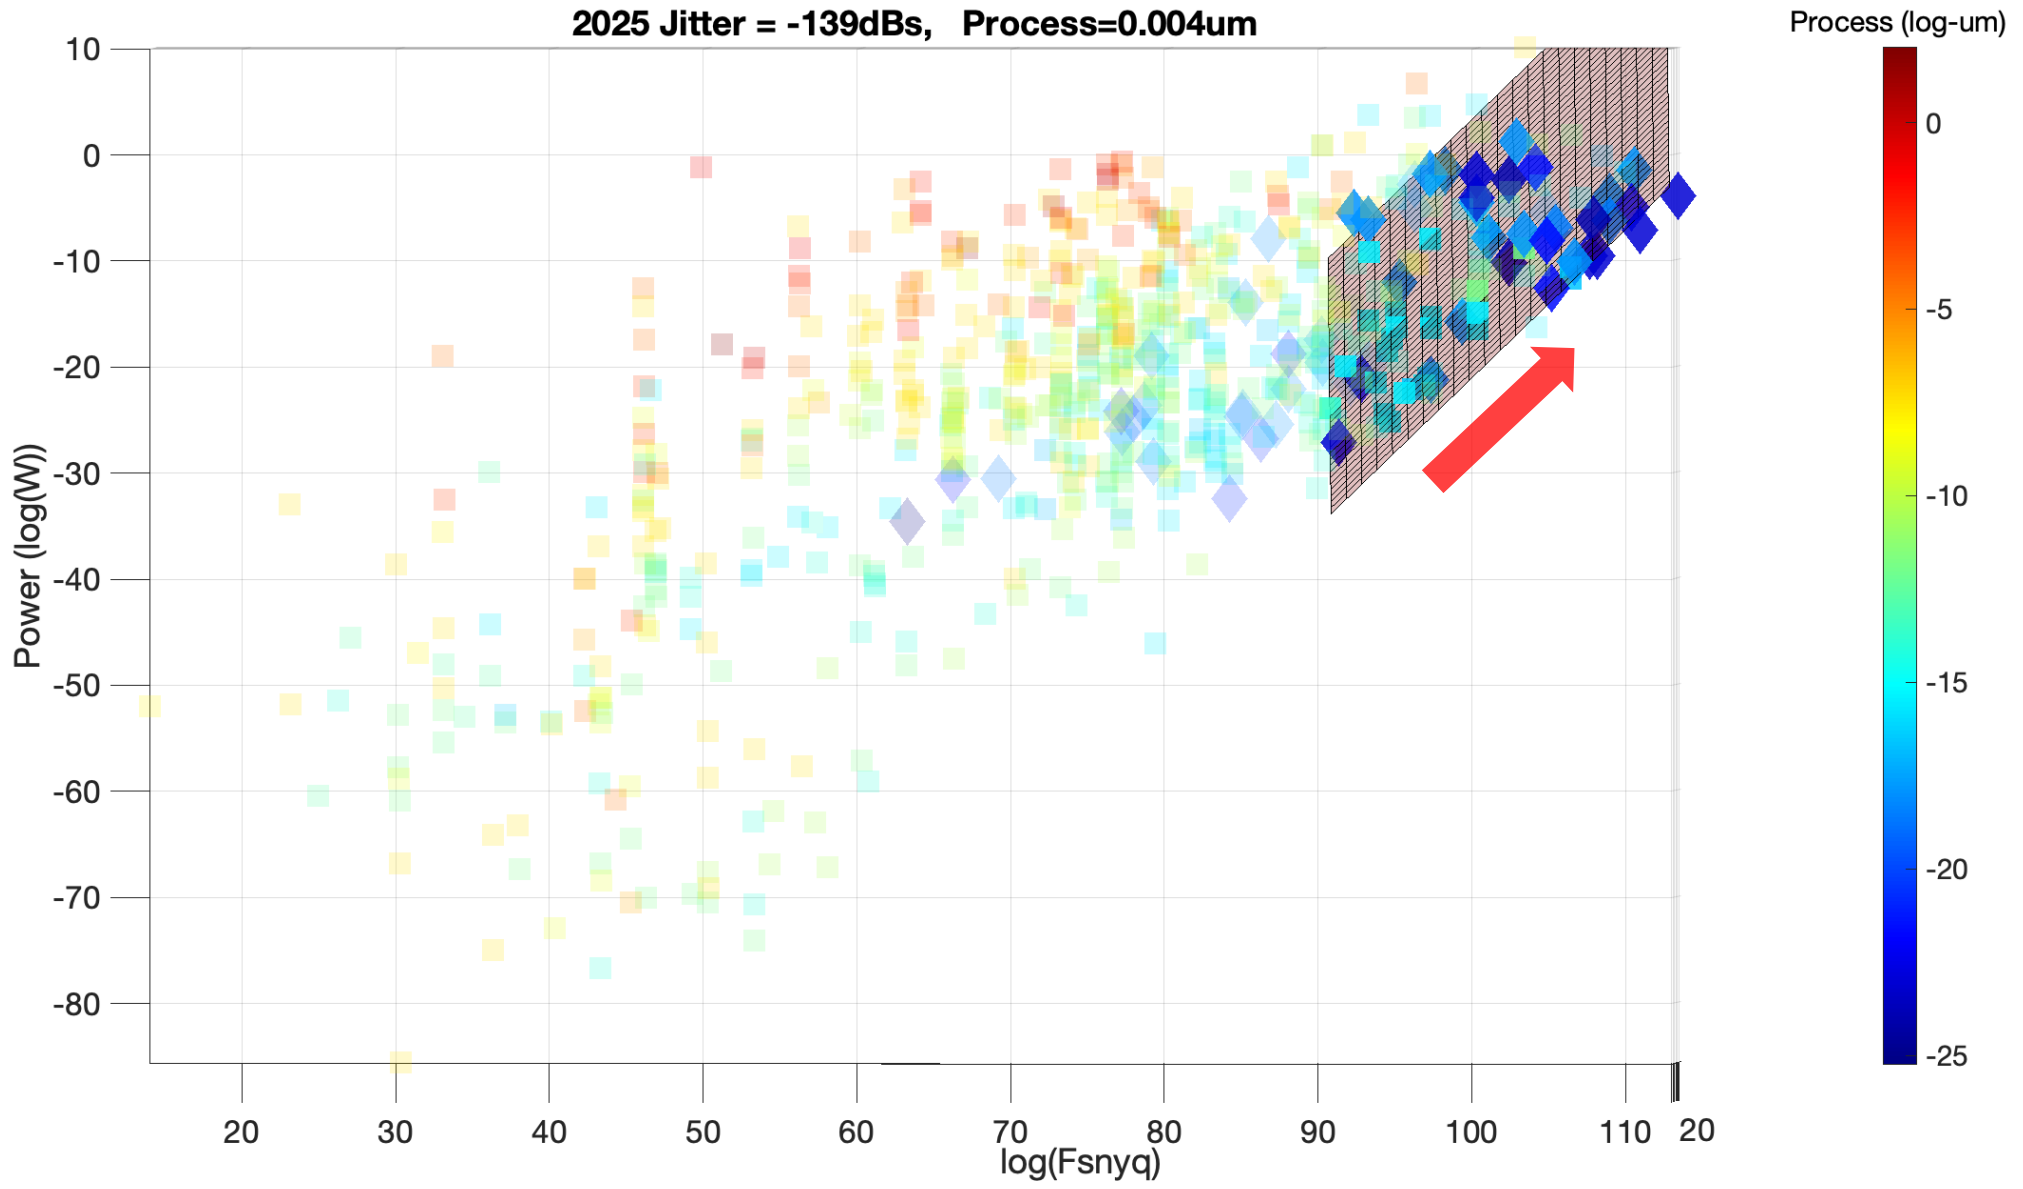
- Survey catalogues all ADCs published in ISSCC and VLSI since 1997
- We can use this data to predict (CMOS) ADC performance trades

[1] <https://github.com/bmurmann/ADC-survey>


# ADC Performance Evolution



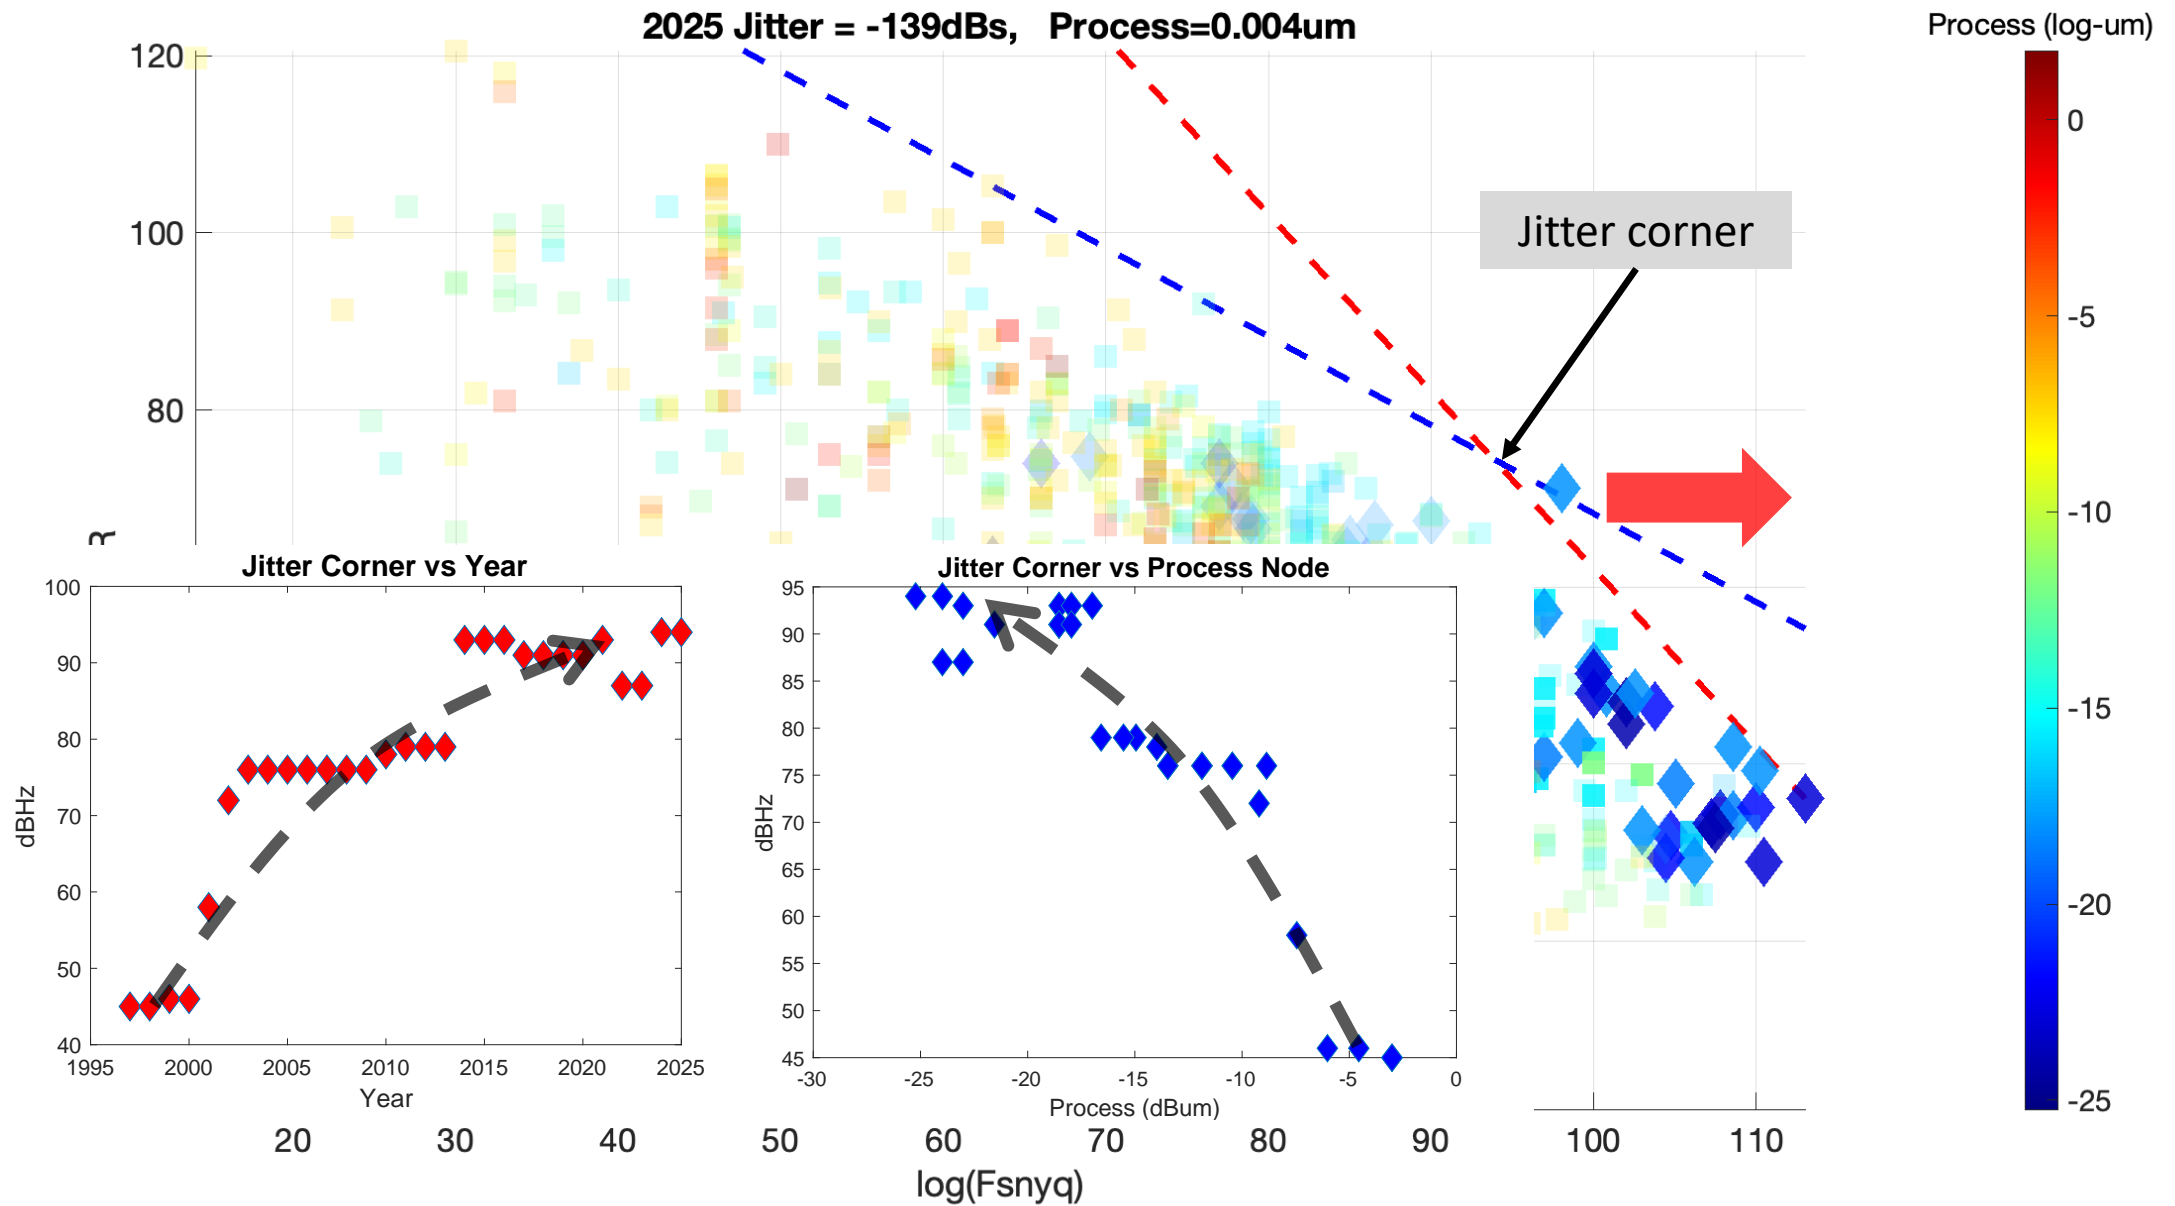
First-order fit for high-Fs ADCs ( $R=0.74$ )



# ADC power vs Fs (constant SNDR): 1.4dBW/dBHz (0.1dB $\sigma$ )



## ADC power vs SNDR (constant Fs): 0.6dBW/dB (0.07dB $\sigma$ )



**Jitter corner improvement ~ 1dB/yr, 2dBHz/dBum**



# ADC Performance Trade: PAM4 vs PAM6 vs PAM8

- Power vs frequency and SNDR tradeoff can be projected for 400G/lane



	Mean	+3 $\sigma$	-3 $\sigma$
Power vs Fs (dB/dB)	1.4	1.7	1.1
Power vs SNDR (dB/dB)	0.59	0.79	0.38

- ADC efficiency Summary

	PAM4	PAM4+3S	PAM4-3S	PAM6	PAM6 +3S	PAM6-3S	PAM8	PAM8+3S	PAM8-3S
pJ/bit	<b>0.63</b>	0.54	<b>0.74</b>	<b>0.48</b>	<b>0.42</b>	<b>0.55</b>	0.56	<b>0.55</b>	0.56

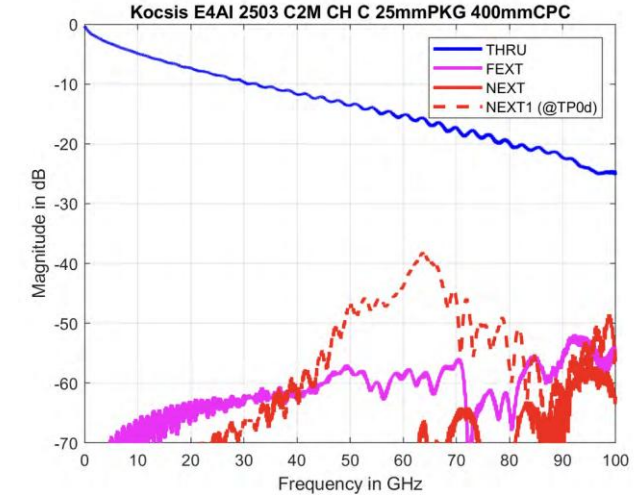
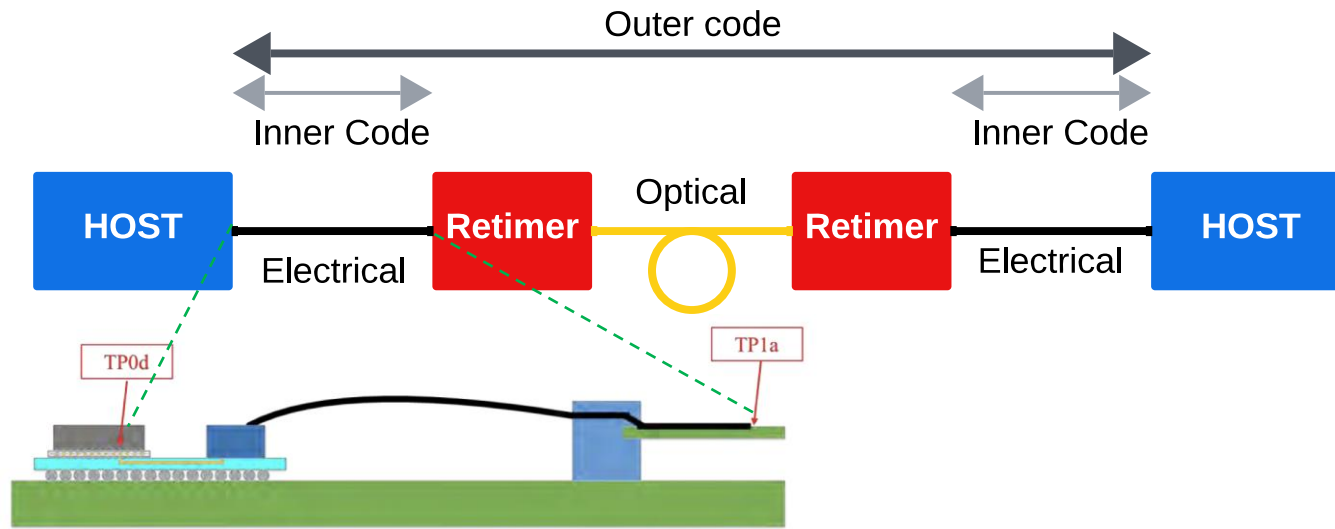
- Accounts for differences in Fs, target SNDR, and information per sample with a code rate of 0.94
  - Requires modestly improved efficiencies from 100G, but 4x faster
  - The Goldilocks choice is PAM6 from an ADC efficiency perspective, with PAM8 surprisingly close
- The analysis is based on Channel “C” model – therefore optimistic for PAM4
- This represents a data-driven best-guess, not reality
  - Actual ADC designs are likely to improve upon these numbers

# FEC improvements – Option 1) RS 1020, 960

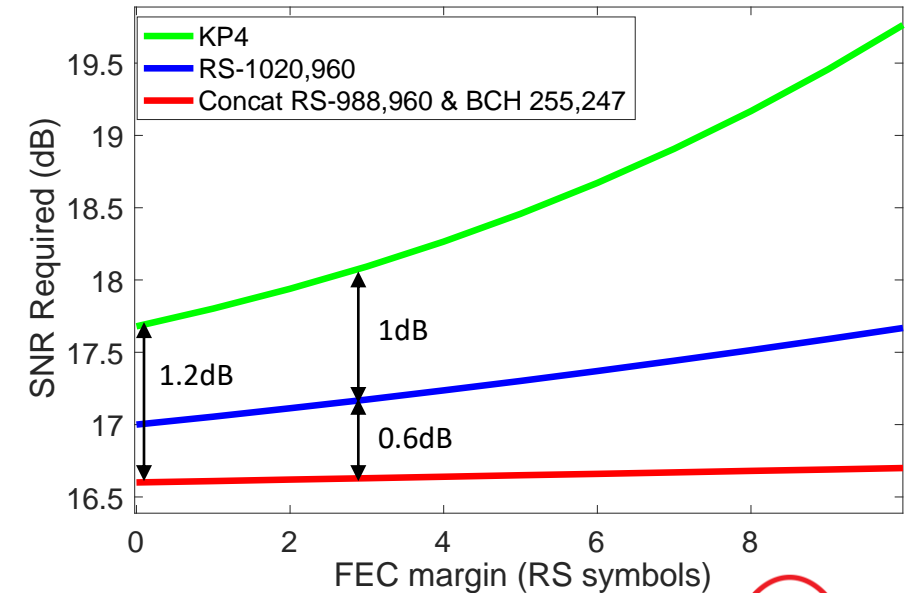


- Simple upgrade to existing KP4 entailing  $O(n)$  increase in complexity
  - › Analysis holds code rate, baud, and Nyquist frequency constant
  - › Offers  $\sim 0.7\text{dB}$  gain over KP4 at  $1\text{e-}15$  post-FEC BER, improving with FEC margin
- $\sim 2\text{x}$  latency of KP4, partly offset by higher baud, parallelism and ASIC clock

# FEC improvements – Option 2) Concatenated Code



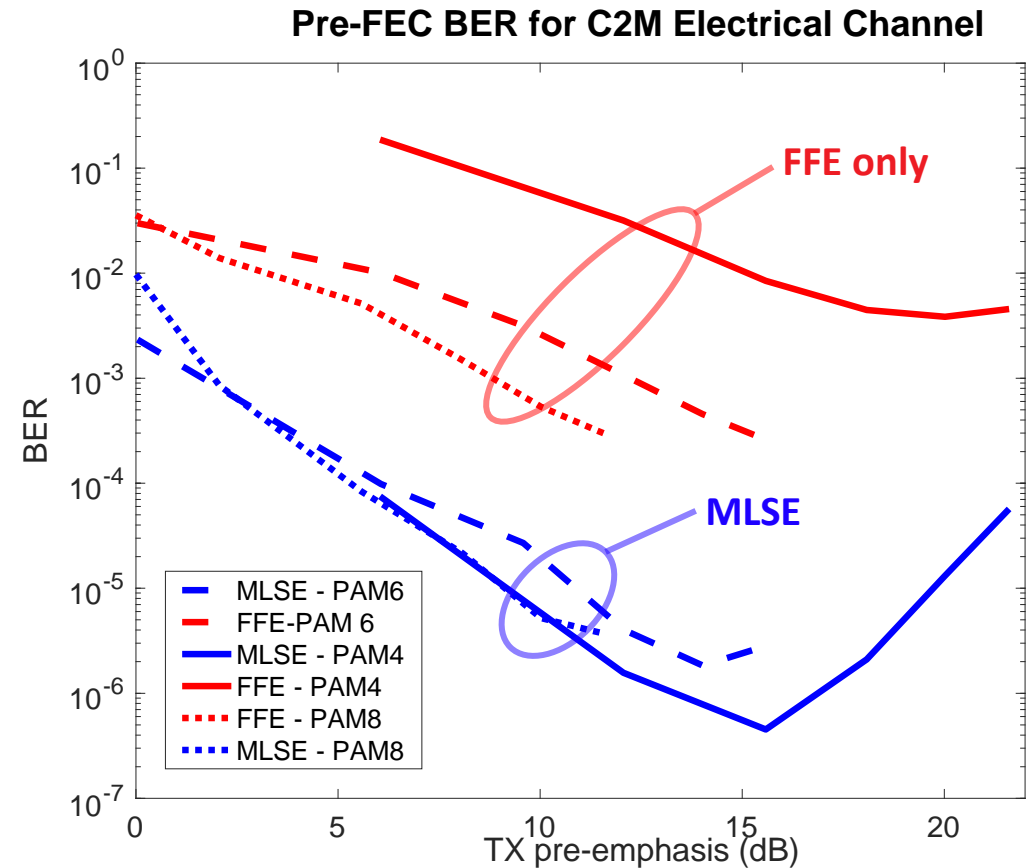
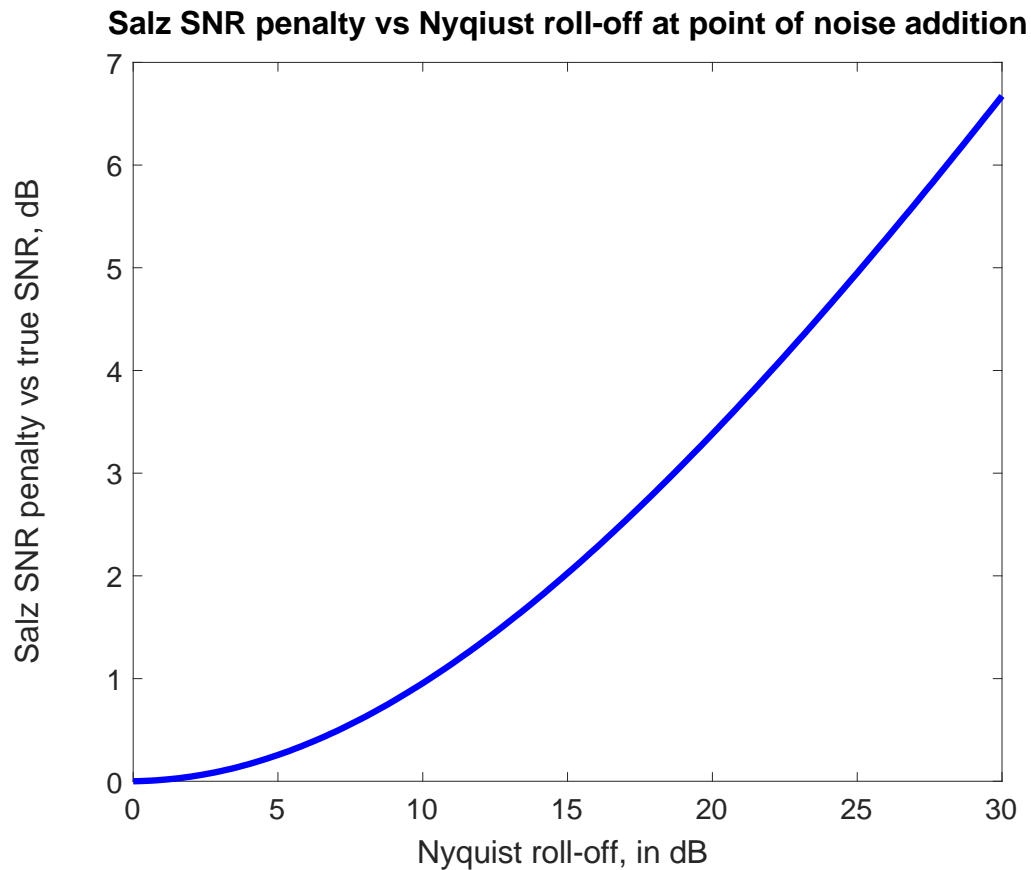
FEC Options: KP4 vs RS-1020,960 vs BCH



- Higher level of protection on electrical links
  - › Soft information enables higher code performance
- > 1dB coding gain vs KP4 @1e-4 pre-FEC BER
- With very little additional SNR, retains full error correction capacity of the link

[2] [https://www.ieee802.org/3/ad\\_hoc/E4AI/public/25\\_0327/kocsis\\_e4ai\\_01\\_250327.pdf](https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0327/kocsis_e4ai_01_250327.pdf)


# Advanced Equalization and Detection



- >20 dB Nyquist loss in channel incurs heavy Salz SNR penalty
- This can be mitigated via ~25dB CTLE peaking and 10-15 dB TX equalization
- MLSE delivers 2-3 decades BER improvement over FFE only


# Evolution toward CPx (x=optical, copper)

Retimed [3]




- Signal integrity challenge
- High power

Non-Retimed



“Analog CPx”  
Current approach


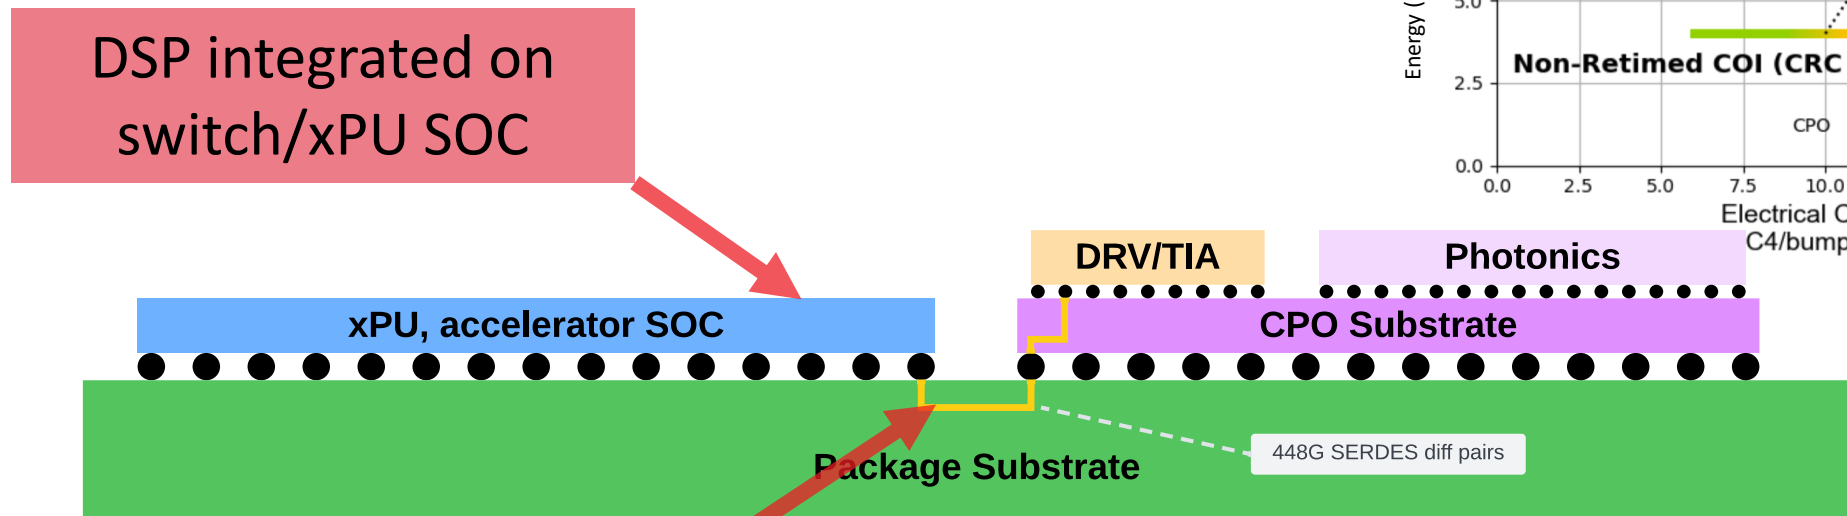
Slow Wide I/F



“Digital CPx”  
The way forward?

[3] Figures from “Considerations for next generation AI compute interconnect”, Hutchins, Jeff, presentation given at the Technology Exploration Forum (TEF), in Santa Clara, CA on October 22-23, 2024, <https://ethernetalliance.org/ethernet-in-the-age-of-ai-agenda/>


# Non-Retimed Analog CPx




- Even at 224G,  $\sim > 7.5\text{dB}$  channel loss[1] imposes energy penalty of  $\sim 1\text{pJ/b}$ 
  - › This interface does not scale easily to future CPx generations
- DSP/SERDES beachfront costs  $\sim 4\text{mm}$  per Rx/Tx 8-lane port
- Interoperability: each ASIC-CPx combo needs reoptimization & requalification
- No standard interface; vendors are not easily interchangeable

[3] Graph adapted from “Considerations for next generation AI compute interconnect”, Hutchins, Jeff, presentation given at the Technology Exploration Forum (TEF), in Santa Clara, CA on October 22-23, 2024, <https://ethernetalliance.org/ethernet-in-the-age-of-ai-agenda/>

# Slow-wide UCle Digital CPx



64 Tx/Rx UCle Module, 12GT/s (768Gbps bidirectional) [4]

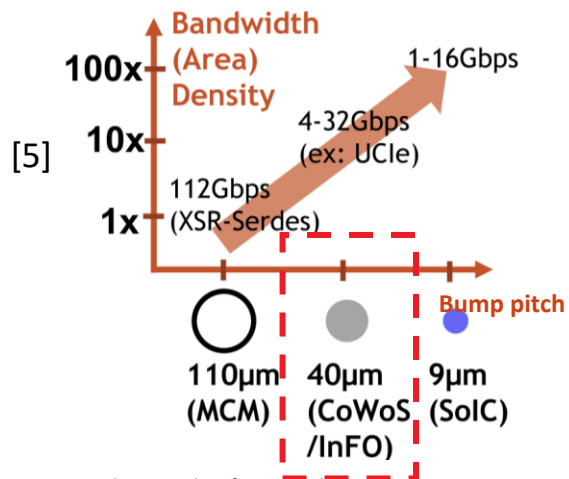


- High beachfront density
  - › 3.94T/mm at 12G/lane on 2mm reach
  - › More than twice the 448G SERDES density
- High energy efficiency
  - › 0.25pJ/b with advanced packaging
- Spreads SERDES heat over ~5x larger area
- Lane tracking and repair for reliability & yield

[4] UCle Specification Revision 2.0, Version 1.0 August 6, 2024

# Slow-wide Digital vs Fast-narrow Analog CPx

	Digital CPx (2mm reach)	Analog CPx
Energy (pJ/b)	0.25	~1 (driver, TIA, DAC/ADC delta)
Beachfront density (Tb/mm)	3.95 (at 12GTps)	1.8 [6]
Beachfront required for 32 ports x 8x448G bidirectional	59mm	128mm
LPO Link Budget Impact	None	>7.5dB loss
ASIC die area per port	~1.8mm <sup>2</sup> including bumps	~4mm <sup>2</sup> excluding bumps [6]
Latency impact	~<2ns	<1ns
SERDES power density [7]	20%	100%



[5] Figure from "High-Bandwidth Chiplet Interconnects for Advanced Packaging Technologies in AI/ML Applications: Challenges and Solutions," Lin, Mu-Shan et al., DOI: 10.1109/OJSSCS.2024.3506694

[6] Assumes die area is comparable to 224G transceiver implementations in 3nm

[7] This is calculated as SERDES power divided by the additional footprint available to dissipate heat, due to fan-out



# More Reasons for Slow-Wide Digital CPx

- Removes linear host-CPx channel → scalable for many generations
  - › No board redesign, link optimization or box requal required with each CPx vendor
- Host compatibility with any CPx device
- Reuse of one CPx product class across many host types (e.g. xPU)
  - › Important for CPx economies of scale
- Host ASICs can evolve at a rate/process node which is independent of CPx
  - › Accelerates the innovation and deployment of new host ASICs
- Supply chain resilience: *allows CPx vendors to be interchangeable*

# Closing statement

- PAM6 is a sensible choice from an ADC efficiency and SI standpoint
- Concatenated codes are well worth a look for resilience and performance
- Advanced equalization and detection keep front panel pluggables in play
- Slow-wide CPx interface offers big efficiency, density, ecosystem benefits

Thank You