ENRZ Advanced Modulation for Low Latency Applications

OIF CEI-56G – Signal Integrity to the Forefront

David R Stauffer

Kandou Bus SA

March 22, 2016

KANDOU BUS

CEI Application Space is Evolving

- The “OIF Next Generation Interconnect Framework” white paper lays out a roadmap for CEI-56G serial links.
  - 2.5D and 3D applications are becoming increasingly relevant.
  - Mid-plane architectures are increasingly used to limit channel loss.
  - High function ASICs (such as switch chips) are driving requirements for higher I/O density and lower interface power.

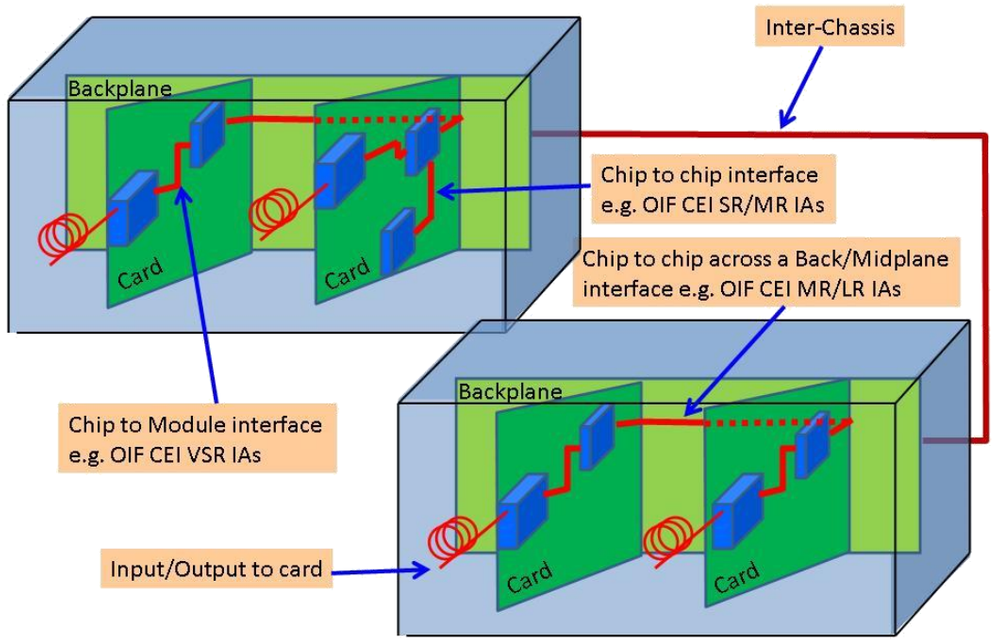

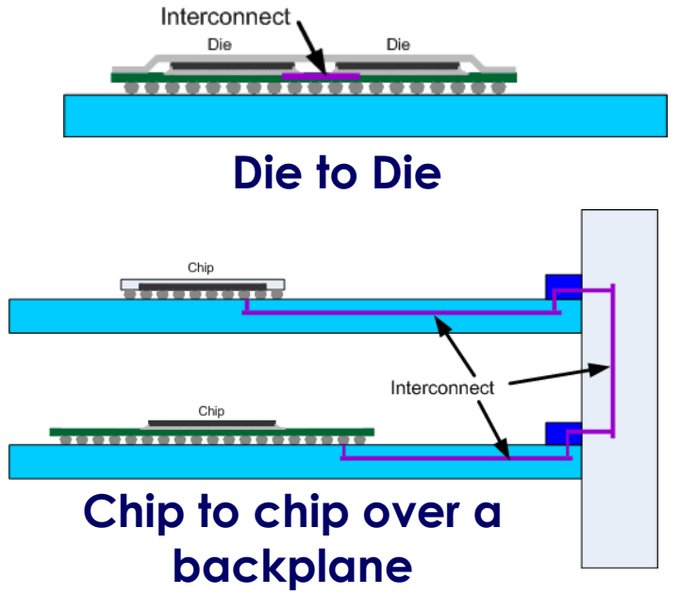

OIF CEI-56G Projects

- CEI-56G Projects are underway for five link reach applications.
- Each reach optimizes the link budget with the goal of providing the lowest possible power dissipation for the application.

CEI-56G-USR

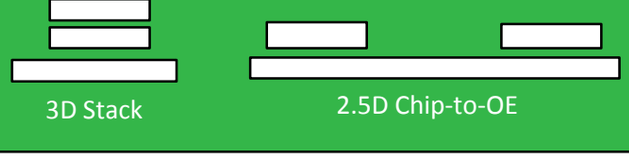

USR: 2.5D/3D die-to-die applications
- 1 cm, no connectors, no packages

CEI-56G-XSR

XSR: Chip to nearby optics engine or LR driver chip
- 5 cm, no connectors
- 5-10 dB loss at 28 GHz

CEI-56G-VSR

VSR: Chip-to-Module interfaces
- 10 cm, 1 connector
- 10-20 dB loss at 28 GHz

CEI-56G-MR

MR: Interface for chip-to-chip and midrange backplane
- 50 cm, 1 connector
- 15-25 dB loss at 14 GHz
- 20-50 dB loss at 28 GHz

CEI-56G-LR

LR: Interface for chip-to-chip over a backplane
- 100 cm, 2 connectors
- 35 dB loss at 14 GHz

OIF OPTICAL INTERNETWORKING FORUM

Optimal Power for the Application

- Power has become the key cost driver in system design.
- Power requirements are driving standards to optimize link budgets for each application space.
  - Driver amplitude
  - Signal processing (FIR, DFE, FEC)
  - Clocking (CDR vs. Common or Forwarded Clock)
- Past practice of using one or two SerDes designs across a wide range of application spaces is no longer feasible.
  - Number of links on switch chips may preclude using LR Serdes.
  - Using USR/XSR interfaces to connect switch chips to off-board Optics Engines or LR Repeater Chips reduces power on the switch chip.

CEI-56G-USR	CEI-56G-XSR	CEI-56G-VSR	CEI-56G-MR	CEI-56G-LR
<< 1 pJ/bit	< 1.5 pJ/bit	< 2.5 pJ/bit	< 5 pJ/bit (incl. FEC)	< 7 pJ/bit (incl. FEC)

- Kandou has presented papers on very low power die-to-die interfaces using advanced modulation: “A Pin-Efficient 20.83Gb/s/wire 0.94 pJ/bit Forwarded Clock CNRZ-5-Coded SerDes up to 12mm for MCM Packages in 28nm CMOS”, Shokrollahi, et al., ISSCC 2016 Session 10.

KANDOU BUS

OIF OPTICAL INTERNETWORKING FORUM

HPC and Networking Applications Diverging

- CEI has been successful as the central specification for SerDes.
  - CEI-based SerDes are available in all major FPGA and ASIC flows.
  - Many public and proprietary CPU-CPU, CPU-I/O, and CPU-Memory interconnects are based on CEI SerDes.
- These interconnects typically use small SRAM or register-based transaction buffers.
  - Credit based flow control is used to avoid overflows.
  - CRC error detection and retry is used to handle errors.
- Credit based flow control is sensitive to latency.
  - Typical range of the bandwidth-delay product is 20-200 bytes.
  - Latency of Ethernet RS FECs exceed these limits and would force redesign to include larger buffers.
  - For some protocols buffers would need to be larger than the limits supported by the protocol.
- **Conclusion: HPC applications are creating a demand for alternative standards that are not dependent on FEC to achieve the link budget.**

KANDOU BUS

OIF OPTICAL INTERNETWORKING FORUM

CEI-56G Electrical Modulation Variants

- OIF is pursuing multiple modulation variants for several reach applications.
- Networking applications (802.3, T11.2) include FEC. Most PAM-4 variants assume FEC is implemented and optimize power/cost based on this assumption.
- Data center applications using token-based protocols cannot tolerate latency associated with FEC. NRZ variants support reasonable BER without utilizing FEC.
- ENRZ provides a "no FEC" option for higher loss applications where NRZ does not work.

Interface	Mod.	Max. Data Rate	IL @Nyquist	Clock Arch.	Elec. BER
CEI-56G-XSR-PAM4	PAM-4	58.0 Gb/s	4.25 dB	Fwd Clk	10^{-15}
CEI-56G-VSR-PAM4	PAM-4	58.0 Gb/s	10 dB	CDR	10^{-6}
CEI-56G-MR-PAM4	PAM-4	58.0 Gb/s	19.67 dB	CDR	10^{-6}
CEI-56G-LR-PAM4	PAM-4	60.0 Gb/s	28.45 dB	CDR	3×10^{-4}
PRELIMINARY – Subject to Change					

Interface	Mod.	Max. Data Rate	IL @Nyquist	Clock Arch.	Elec. BER
CEI-56G-USR-NRZ	NRZ	58.0 Gb/s	2 dB	Fwd Clk	10^{-15}
CEI-56G-XSR-NRZ	NRZ	58.0 Gb/s	8 dB	Fwd Clk	10^{-15}
CEI-56G-VSR-NRZ	NRZ	56.0 Gb/s	20 dB	CDR	10^{-15}
CEI-56G-MR-NRZ	NRZ	56.0 Gb/s	30 dB	CDR	10^{-15}
CEI-56G-LR-ENRZ	ENRZ	112.4 Gb/s (4 wires)	33.59 dB	CDR	10^{-15}
PRELIMINARY – Subject to Change					

KANDOU BUS

OIF OPTICAL INTERNETWORKING FORUM

ENRZ Multiwire Code

ENRZ is a 3-bit over 4-wire Chord™ signaling code that fills the space between single-ended and differential signaling

- Bandwidth per wire is higher than differential NRZ at similar baud rate.
- Inter-symbol interference is lower than PAM-4/8 interfaces
- Noise rejection characteristics are similar to differential signals

Because of signal integrity advantages, ENRZ does not require a FEC as is required for other LR variants.

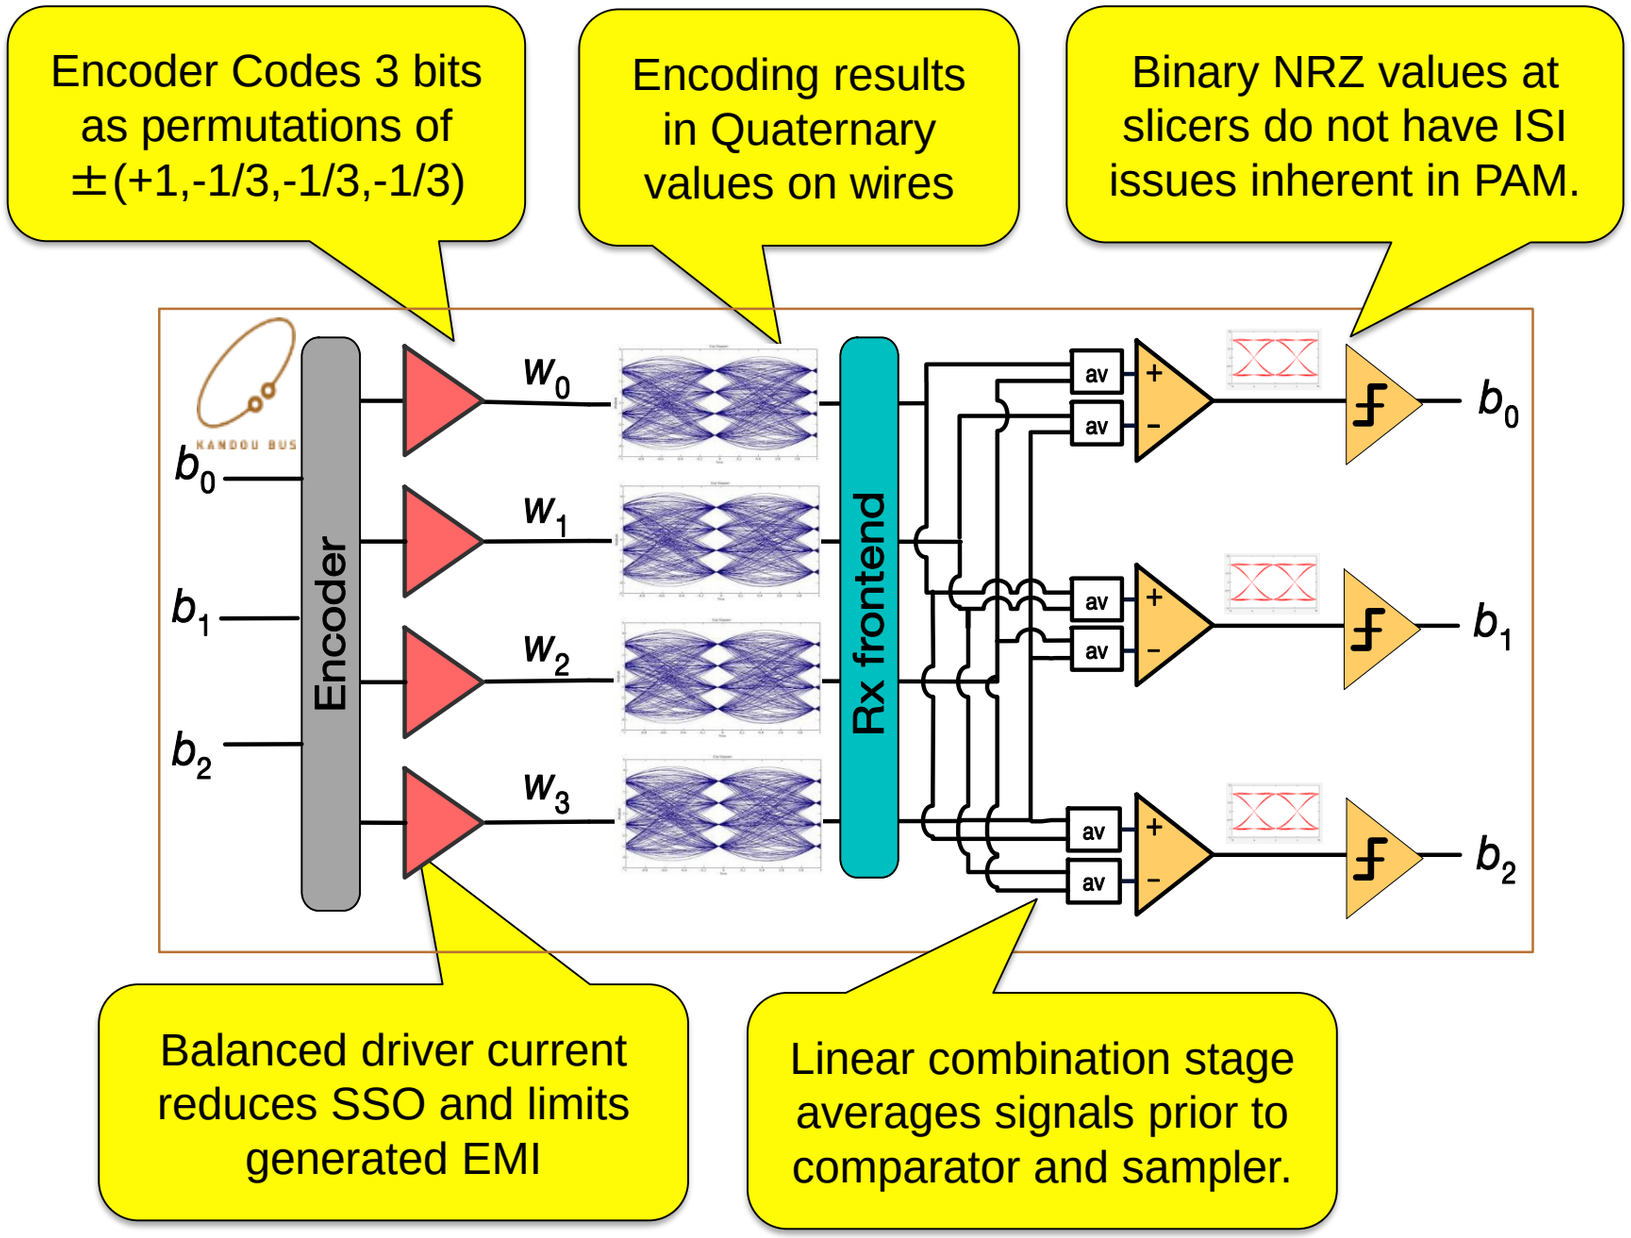

LR Channel Simulation Using ENRZ without FEC

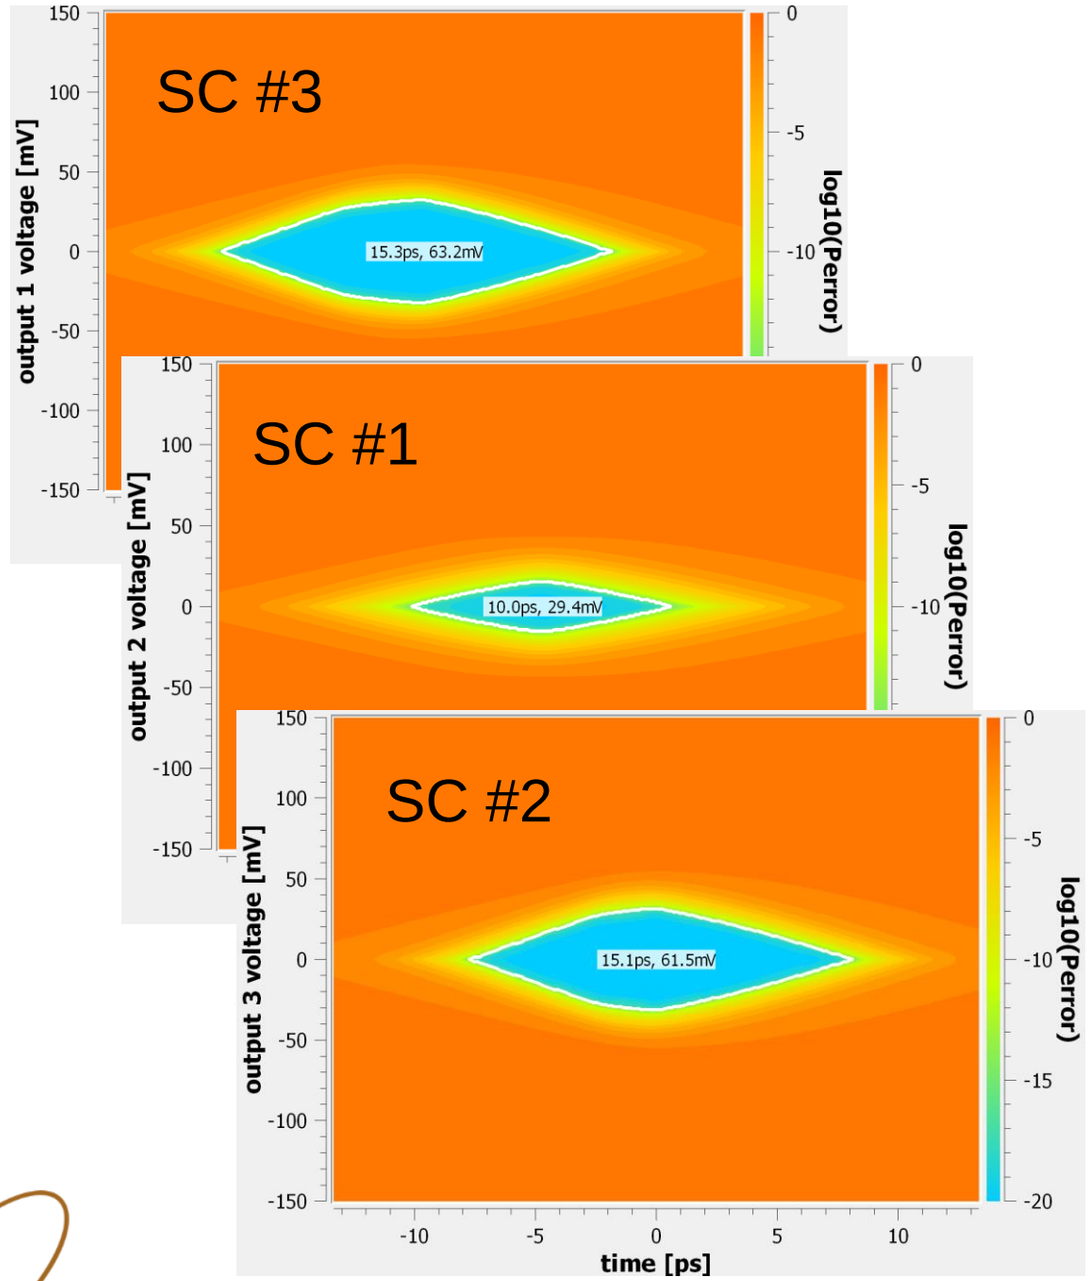

- Simulation conditions:
  - ENRZ @37.5 GBd
  - Tx Launch: 1000 mVppd
  - FFE (3-tap), VGA (10 dB), CTLE (4 dB), DFE (20-tap)
  - BER = 1E-15 (no FEC)
- Results:
  - EH: 29.5 mV (SC#2)
  - EW: 0.375 UI (SC#2)
- Conclusion:
  - SC#2 has sufficient eye opening.
  - No FEC is required.

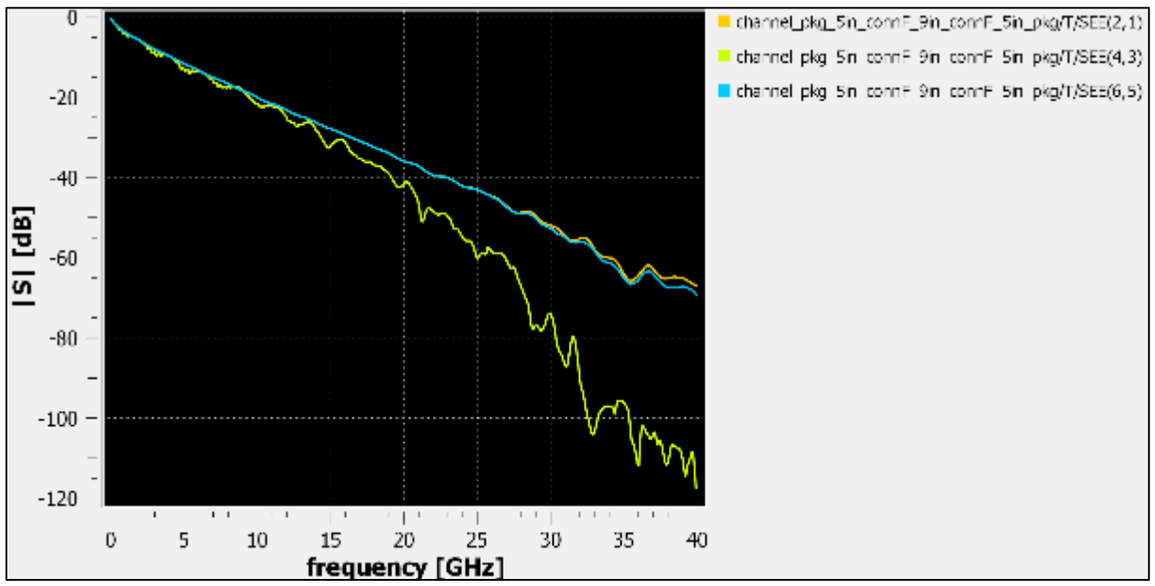

LR Channel Simulation Using PAM4 with FEC

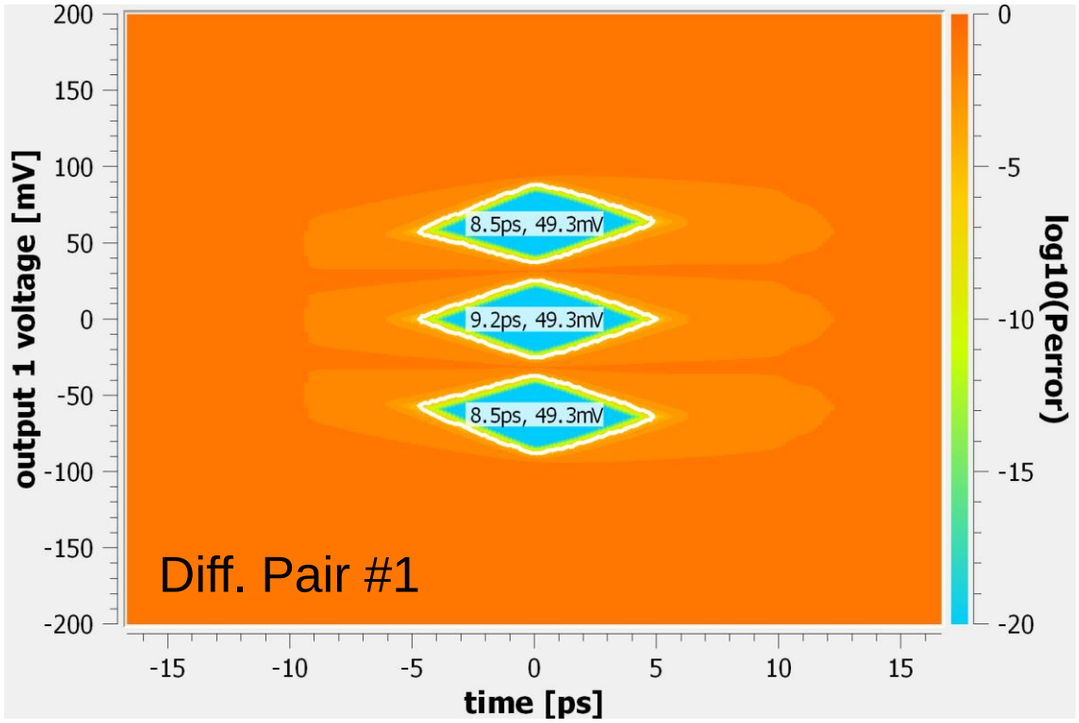

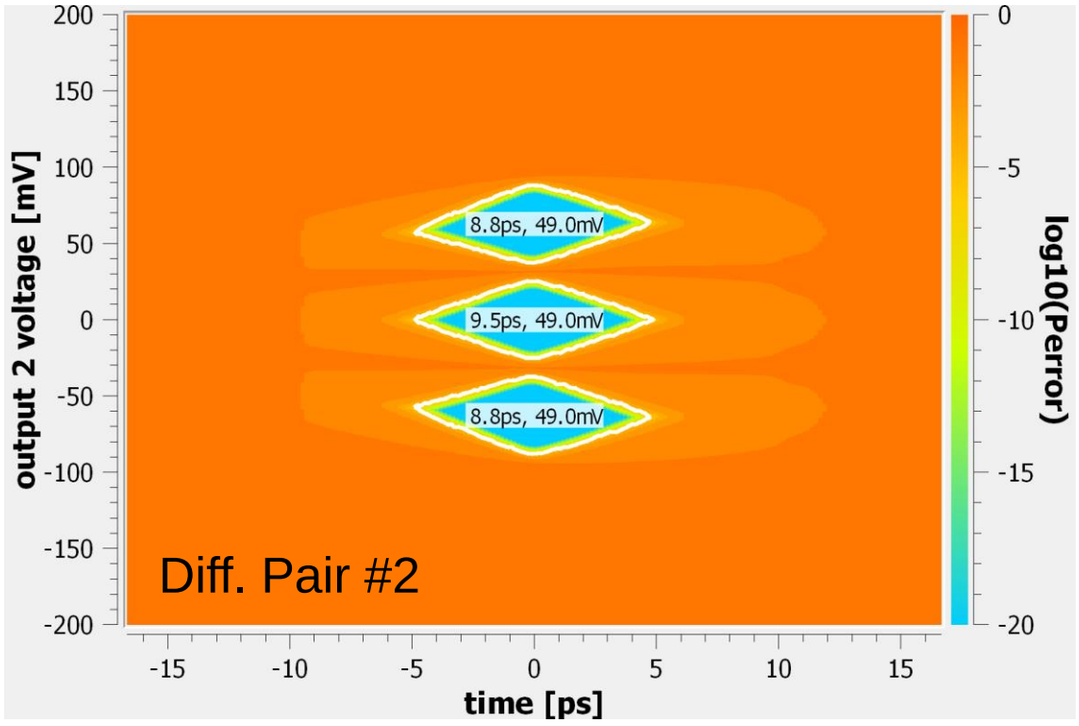

- Simulation conditions:
  - PAM-4 @30 GBd
  - Tx Launch: 1000 mVppd
  - FFE (3-tap), VGA (10 dB), CTLE (12 dB), DFE (20-tap)
  - BER = 1E-6 (assumes FEC)
- Results:
  - EH: 49.0 mV (pair #2)
  - EW: 0.255 UI (pair #1)
- Conclusion:
  - PAM-4 is a viable option, but requires FEC.

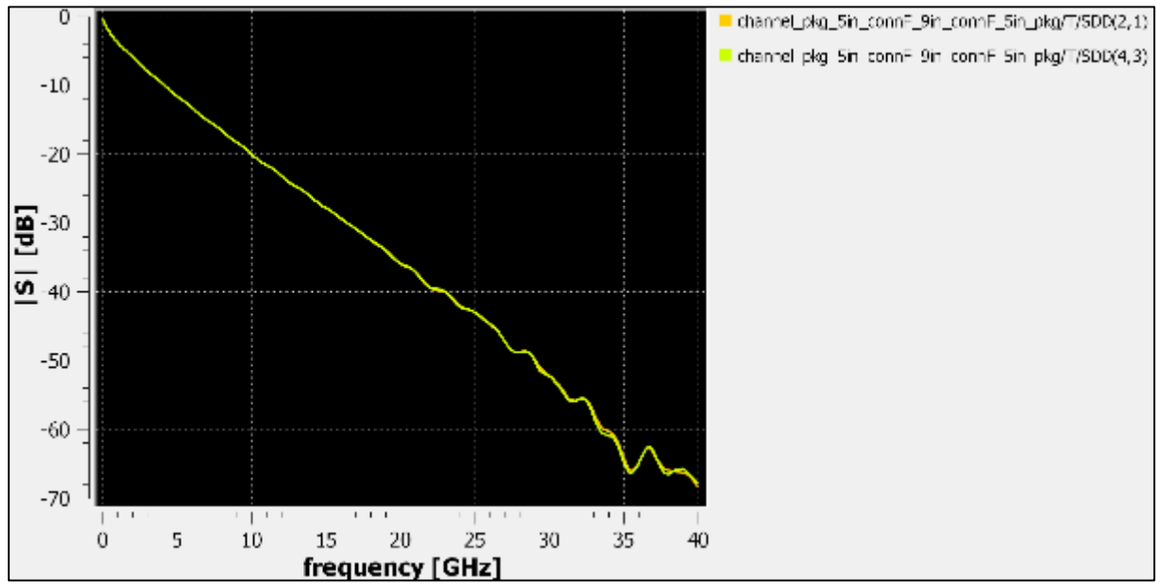

Summary and Conclusions

- Power requirements are driving standards to optimize link budgets for each application space.
- The selection of PAM-4 modulation (which is dependent on FEC) by networking applications has forced a divergence between interface standards for networking and HPC interface applications.
- NRZ modulation can handle interface reaches up to MR without requiring FEC.
- ENRZ provides a “No FEC” solution for LR interfaces.
- **OIF 100G Serial and Beyond Workshop on March 24th will explore next generation CEI-112G interfaces.**
  - **Kandou Bus will be presenting Chord™ signaling architectures for 112G and 224G in this workshop.**

KANDOU

reinventing the

BUS