

# XSR / USR Interface Analysis including Chord Signaling Options

**David R Stauffer**  
**Margaret Wang Johnston**  
**Andy Stewart**  
**Amin Shokrollahi**

**Kandou Bus SA**

May 12, 2014



**Contribution Number: OIF2014.112**

**Working Group: Physical & Link Layer WG (Electrical Track)**

**Title: XSR / USR Interface Analysis including Chord Signaling Options**

**Source: David R Stauffer**

[david@kandou.com](mailto:david@kandou.com)

**Kandou Bus, S.A.**

**Date: May 12, 2014**

**Abstract: This analysis compares signal integrity and power analysis results for various Chord Signaling codes in CEI-56G Extra Short Reach (XSR) and Ultra Short Reach (USR) channel applications. Codes are compared to an NRZ baseline.**

Notice: This contribution has been created to assist the Optical Internetworking Forum (OIF). This document is offered to the OIF solely as a basis for discussion and is not a binding proposal on the companies listed as resources above. Each company in the source list, and the OIF, reserves the rights to at any time to add, amend, or withdraw statements contained herein.

This Working Text represents work in progress by the OIF, and must not be construed as an official OIF Technical Report. Nothing in this document is in any way binding on the OIF or any of its members. The document is offered as a basis for discussion and communication, both within and without the OIF.

For additional information contact:  
The Optical Internetworking Forum, 48377 Fremont Blvd.,  
Suite 117, Fremont, CA 94538  
510-492-4040 phone ☎ [info@oiforum.com](mailto:info@oiforum.com)



# Disclosure

- Kandou Bus, S.A. discloses that we own intellectual property related to Chord Signaling and other material described in this contribution.
- We are committed to RAND licensing of all of our technologies.
- We are committed to adhering to the bylaws of all standards organizations to which we contribute and maintain membership.
- We are committed to be good corporate citizens.




# Next Generation Switch Chip Considerations

- Alcatel-Lucent outlined problem in oif2014.029:
- Power:
  - 56 Gb/s Serdes will hardly scale power or area.
    - 30% performance improvement is not enough.
  - Advanced modulation schemes also unlikely to scale power sufficiently.
  - CEI 56G solution for XSR must be below 5 pJ/bit.
  - CEI 56G solution for USR must be below 3 pJ/bit.
  - *Observation: These targets are not very aggressive.*
- Area:
  - Beachfront reduction is also needed so that Serdes can fit around periphery of chip.




# XSR Application Space



- Optical engines located on PCB adjacent to switch chip.
- Switch chip is packaged silicon:
  - Serdes edge beachfront is limited by bump pitch.
  - Pin efficiency will drive whether beachfront can be reduced.
- Channels consist of package models, 5 cm of PCB trace.
  - Reflections caused by discontinuities in package models are a significant factor in signal integrity analysis.
- Considerations for signaling technology selection:
  - Min. Tx amplitude and min. signal processing required to meet channel requirements.
  - Compatibility with USR solution.



# USR Application Space



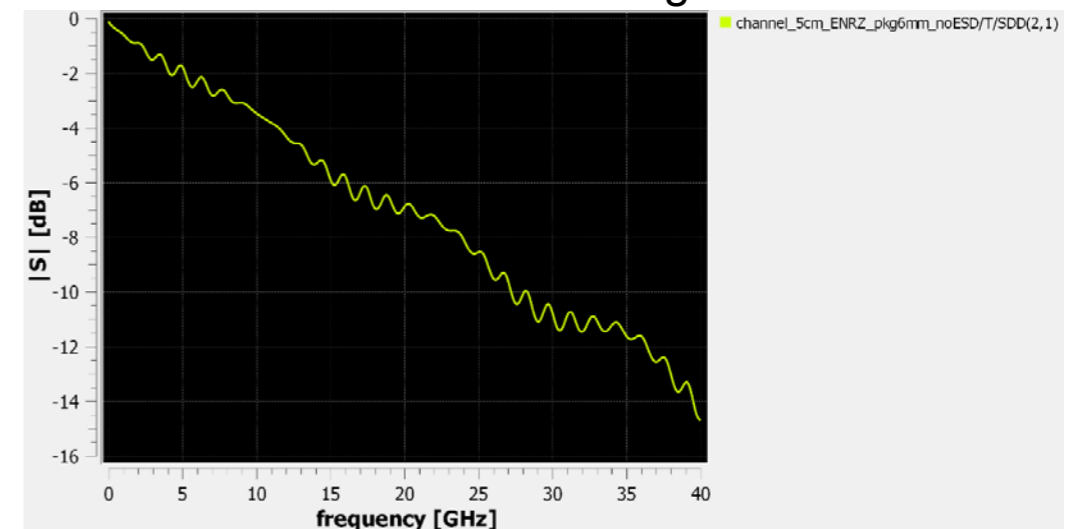
- Optical engines or outboard Serdes located on silicon substrate with switch chip (2.5D) or stacked (3D).
- Channels consist of 1 cm of silicon substrate trace, no packages.
  - Signal integrity is less of a concern when package models are removed from channel.
- Considerations for signaling technology selection:
  - Minimize power to greatest extent possible.
  - Compatibility with XSR solution.




# XSR Channel Definition

- Channel Model:
  - Generated by SystemSI
  - FR4 ( $\epsilon_r=3.7$ ,  $\tan D=0.019$ )
  - Length = 5 cm
  - $Z_{diff}(1-2) = Z_{diff}(3-4) = 99.7$  ohms
- Package Models:
  - COM method
  - Max return loss <2-GHz is  $\sim 11.2$ dB


Sdd21 for Channel+Packages



COM pkg, Cd = 20fF, 6mm trace, Cp = 80fF



# Signal Processing Assumptions (Tx)

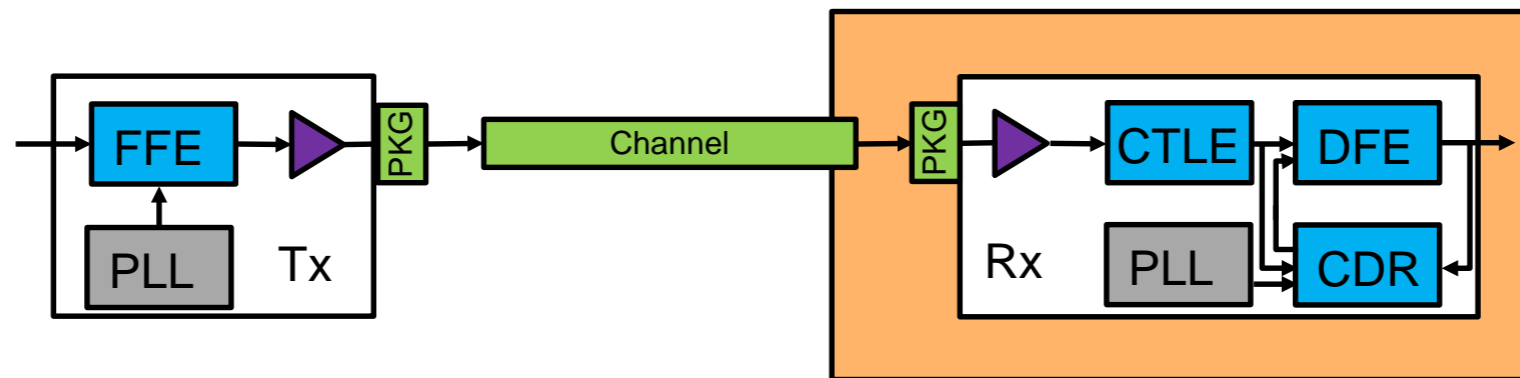


- Tx PLL:
  - Assume PLL for NRZ exists on chip and is shared with other functions.
  - PLL also not needed for EP3L (baud rate is NRZ baud rate divided by 2).
  - Other Chord signaling codes require different frequencies, so PLL is included in analysis for these codes.
- FFE: Assume 1-tap FFE for all options.
  - SI simulations show advantage to having a post cursor tap.
- Driver Amplitude: Assume 200 mVppd where possible.
  - Increase to 400 mVppd where dictated by SI results.





# Signal Processing Assumptions (Rx)



- CTLE: SI simulations do not show advantage for CTLE.
  - Assume no CTLE in XSR/USR applications.
- CDR: Both forwarded clock and CDR options evaluated.
  - Either can be used independent of signaling option.
  - Power savings of eliminating CDR is offset by additional power to drive and receive the clock signal.
  - Forwarded clock adds additional pins on beachfront.
- DFE: Not needed for XSR/USR applications.



# Code Comparison

- A number of chord signaling codes can be applied to XSR/USR applications, including (but not limited to) the 4-wire ENRZ and EP3L codes, and 6-wire Glasswing code.
  - Chord signaling codes have higher code efficiency than NRZ and better SI characteristics than PAM-4.
  - Higher code density can translate to better power efficiency (pJ/bit).
- This presentation evaluates NRZ, ENRZ, and Glasswing codes:
  - NRZ is included as a baseline.
  - ENRZ, EP3L, and Glasswing are evaluated because we have existing power data on these codes.
  - PAM-4 is not evaluated; we do not have access to power data for MLS codes.
- Other Chord Signaling codes may also merit investigation and we may present those in the future.

	NRZ	ENRZ	EP3L	Glasswing	PAM-4
Code Classification	1b2w	3b4w	4b4w	5b6w	2b2w
Code Efficiency	0.5	0.75	1.00	0.83	1.0
ISI Ratio	1	1	2	2	3
Eye Amplitude (Normalized)	1.0	0.67	0.5	0.5	0.33
Baud Rate	56 GBd	37 GBd	28 GBd	22.4 or 44.8 GBd	28 GBd



# ENRZ 3b4w Code Description

- ENRZ is a 3b4w code.
- Code book consists of all permutations of:  
(+1, -1/3, -1/3, -1/3) and (-1, +1/3, +1/3, +1/3).
  - Total of 8 code words used to encode 3 bits of data.
- $V_{CM}$  is a constant (sum of state values for all code words is zero).
- Receiver differentially decodes each sub-channel by combining inputs according to the Hadamard matrix:
  - Row 2:  $(A+C)-(B+D)$
  - Row 3:  $(A+B)-(C+D)$
  - Row 4:  $(A+D)-(B+C)$



# EP3L 4b4w Code Description

- EP3L uses 5 levels on each of four wires (  $\{1, \frac{1}{2}, 0, -\frac{1}{2}, -1\}$ )
  - ENRZ has four levels
  - Optimized to deliver the best vertical opening
- 16 codewords are used out of 18
  - Delivers exactly 4 bits over 4 wires
  - The extra 2 codewords are also available for use
- Receiver is similar to ENRZ, but the output of the ENRZ comparators is followed by PAM-3 slicers
  - Code is a particular variant of PAM-3 over ENRZ
- 4 bits are extracted from the resulting 3 ternary values using a simple decoder
  - Delivers a native 4x sub-multiplexing structure to support 4 x 25 Gb/s optics without additional bit-muxing



# Glasswing 5b6w Code Description

- Glasswing 5b6w code is a ternary code that encodes 5 bits per baud symbol over 6 wires.
- Code book consists of permutations of: (+1, +1, 0, 0, -1, -1)
  - Total of 32 code words used to encode 5 bits of data.
- $V_{CM}$  is a constant (sum of state values for all code words is zero).
- Receiver differentially decodes each sub-channel by combining inputs.
- Decode is performed directly by comparators; no logic decode stage is needed.




# Evaluation Cases

- 112 Gb/s Interfaces:
  - 4 x 28 GBd NRZ (baseline)
  - 2 x 56 GBd NRZ with shared CDR (XSR/USR)
  - 2 x 56 GBd NRZ with forwarded clock (XSR/USR)
  - 1 x 37 GBd ENRZ (3b4w) with CDR (XSR/USR)
  - 1 x 37 GBd ENRZ (3b4w) with forwarded clock (XSR/USR)
  - 1 x 28 GBd EP3L (4b4w) with CDR (XSR/USR)
  - 1 x 28 GBd EP3L (4b4w) with forwarded clock (XSR/USR)
  - 1 x 22.4 GBd Glasswing (5b6w) with CDR (XSR chan)
  - 1 x 22.4 GBd Glasswing (5b6w) with forwarded clock (XSR chan)
  - 1 x 22.4 GBd Glasswing (5b6w) with CDR (USR chan)
  - 1 x 22.4 GBd Glasswing (5b6w) with forwarded clock (USR chan)
- 224 Gb/s Interfaces:
  - 8 x 28 GBd NRZ (baseline)
  - 4 x 56 GBd NRZ with shared CDR (XSR/USR)
  - 4 x 56 GBd NRZ with forwarded clock (XSR/USR)
  - 2 x 37 GBd ENRZ (3b4w) with CDR (XSR/USR)
  - 2 x 37 GBd ENRZ (3b4w) with forwarded clock (XSR/USR)
  - 2 x 28 GBd EP3L (4b4w) with CDR (XSR/USR)
  - 2 x 28 GBd EP3L (4b4w) with forwarded clock (XSR/USR)
  - 1 x 44.8 GBd Glasswing (5b6w) with CDR (XSR chan)
  - 1 x 44.8 GBd Glasswing (5b6w) with forwarded clock (XSR chan)
  - 1 x 44.8 GBd Glasswing (5b6w) with CDR (USR chan)
  - 1 x 44.8 GBd Glasswing (5b6w) with forwarded clock (USR chan)



# KEYE Results – NRZ XSR

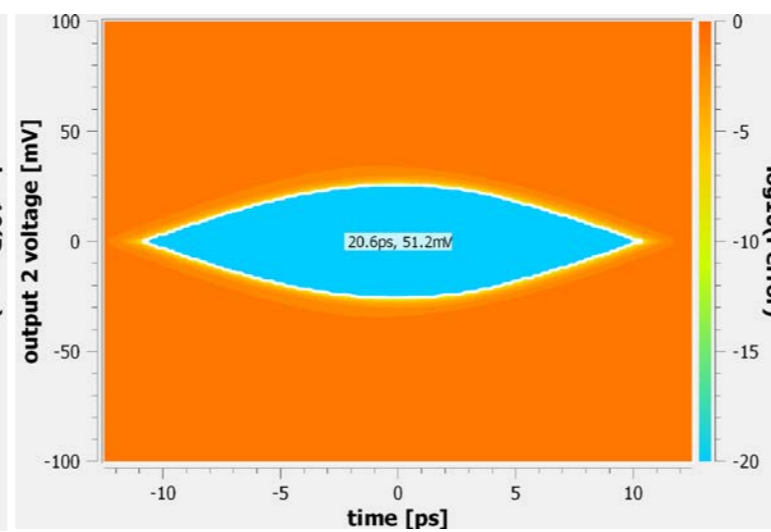
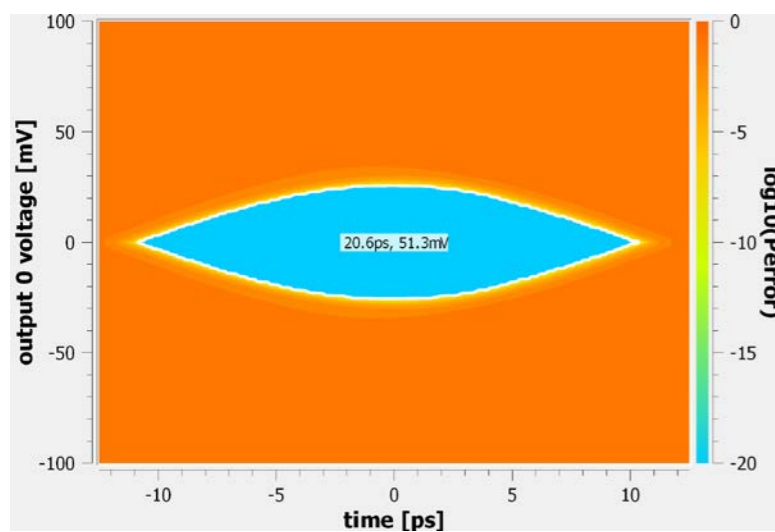



- Simulation conditions:
  - NRZ @56 GBd
  - Tx Launch: 200 mVppd
  - 1-tap FFE
  - No CTLE or DFE.
  - XSR: (5 cm, w/Pkg)
  - BER = 1E-15
- Eye Width/Height:
  - 37.3 mV
  - 0.74 UI
  - Eye open





# KEYE Results – ENRZ XSR

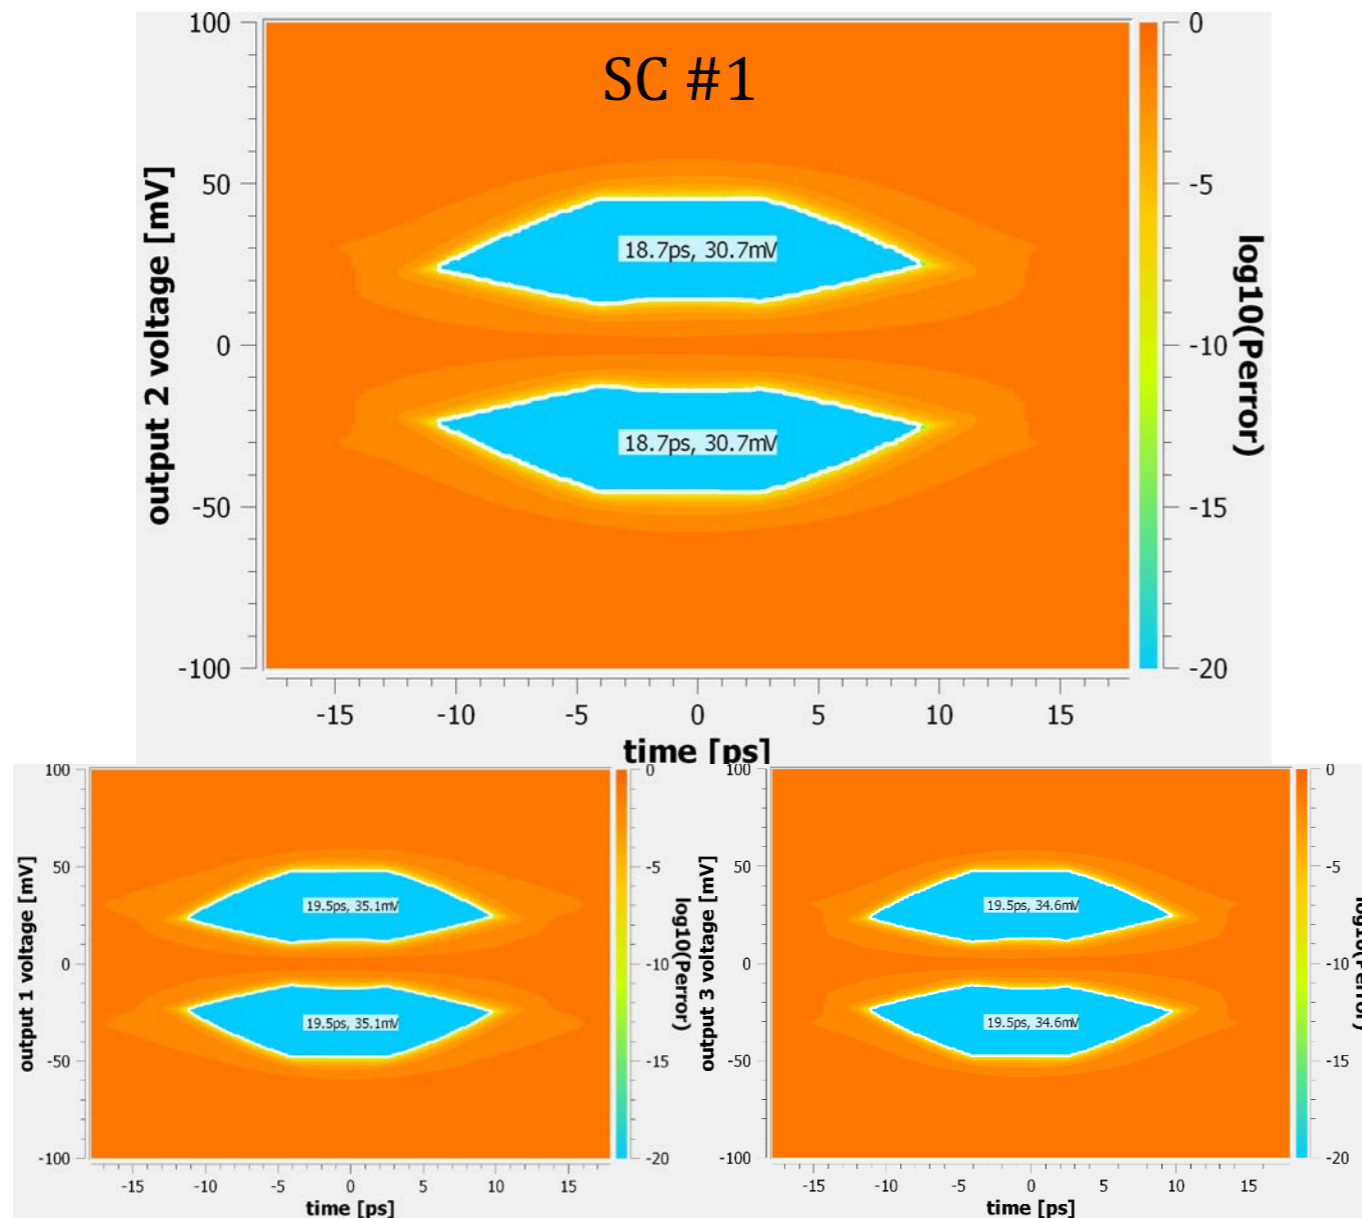


- Simulation conditions:
  - ENRZ @40 GBd
  - Tx Launch: 200 mVppd
  - 1-tap FFE
  - No CTLE or DFE
  - XSR: (5 cm, w/Pkg)
  - BER = 1E-15
- Eye Width/Height:
  - SC#1 is weakest eye
  - 48.3 mV
  - 0.80 UI
  - Eye open





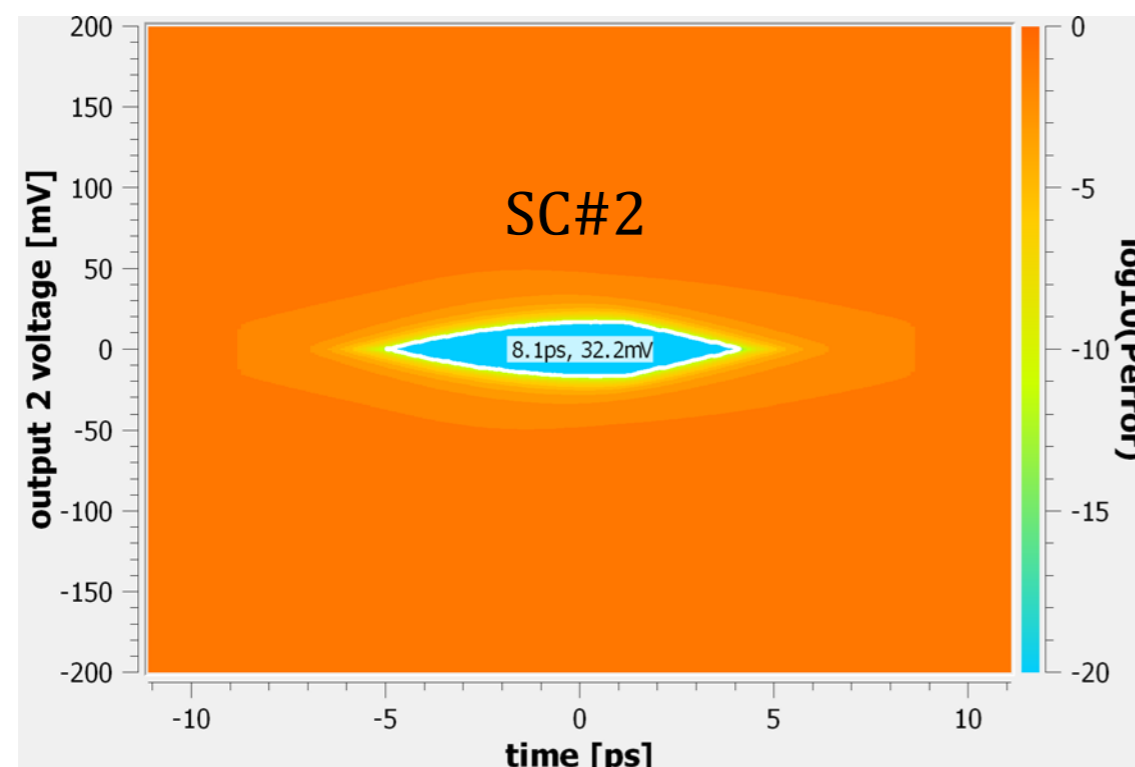

# KEYE Results – EP3L XSR



- Simulation conditions:
  - EP3L @28 GBd
  - Tx Launch: 200 mVppd
  - 1-tap FFE (1 post-tap)
  - No CTLE or DFE
  - XSR: (5 cm, w/Pkg)
  - BER = 1E-15
- Eye Width/Height:
  - SC#1 is weakest eye
  - 31.3 mV
  - 0.525 UI
  - Eye open




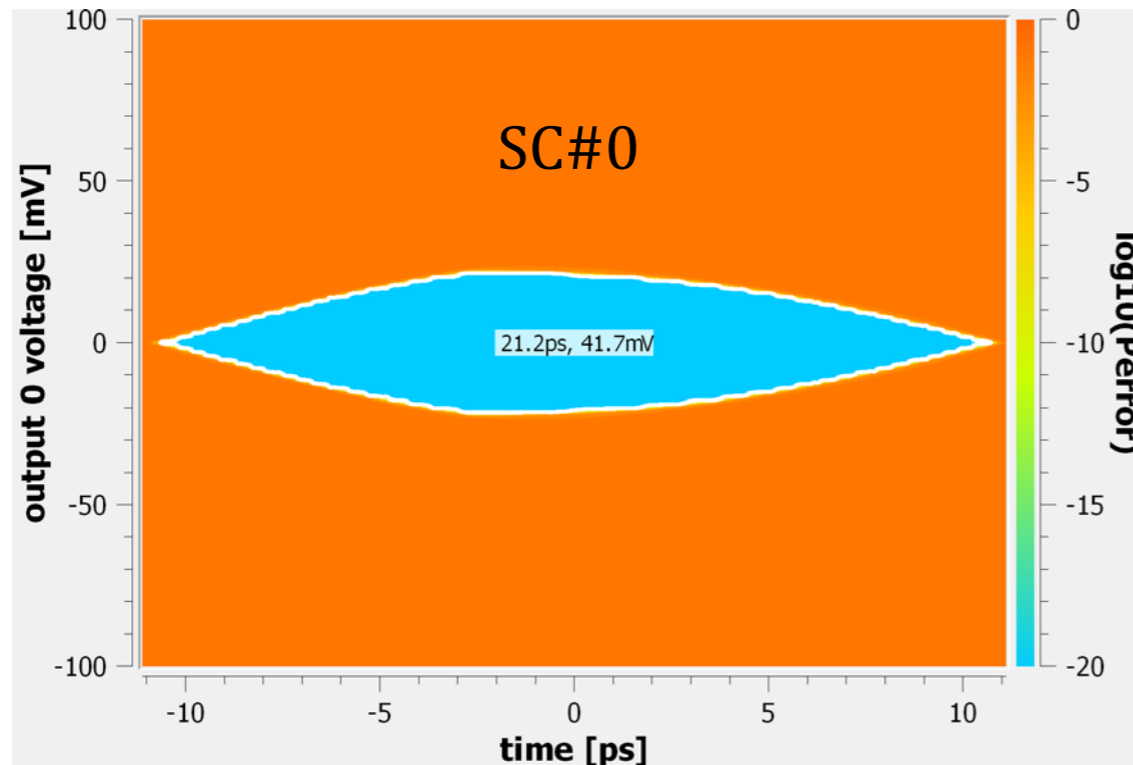
# KEYE Results – GW XSR (224 Gb/s i/f)



- Simulation conditions:
  - Glasswing @45 GBd (for 224 Gb/s i/f)
  - Tx Launch: 400 mVppd
  - 1-tap FFE
  - No CTLE or DFE
  - XSR: (5 cm, w/Pkg)
  - BER = 1E-15
- Eye Width/Height:
  - 31.4 mV (SC#0)
  - 0.36 UI (SC#2,#3)
  - Eye open




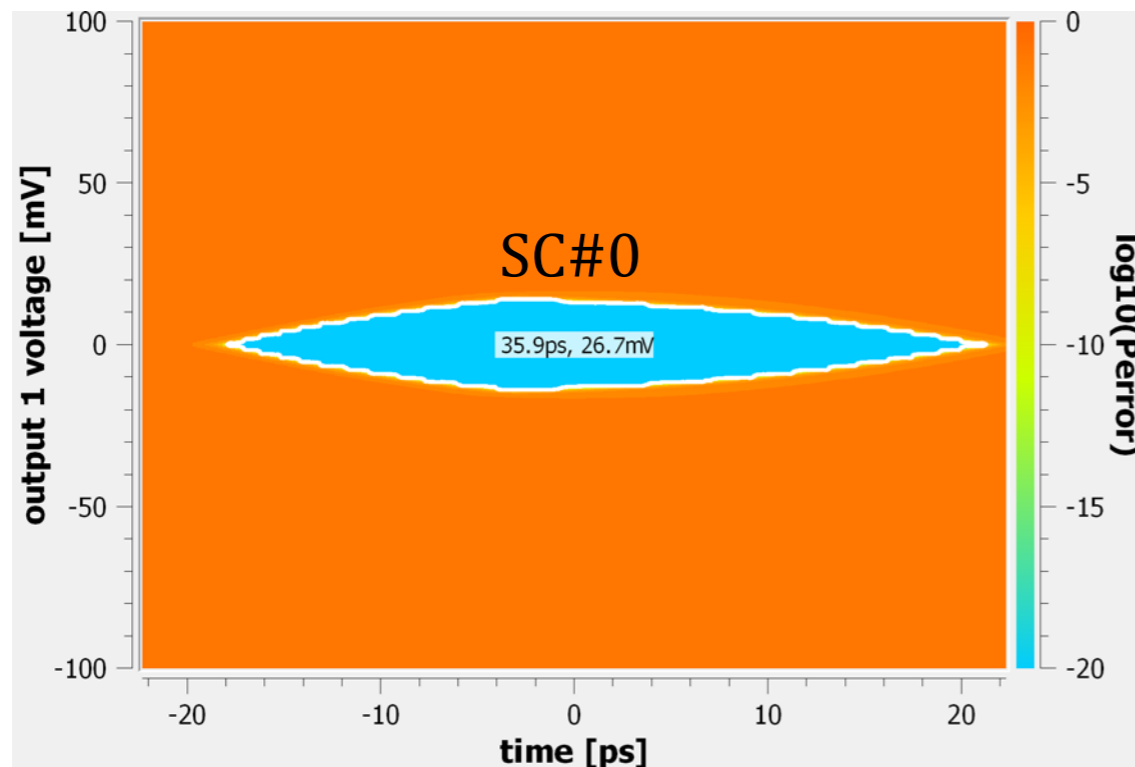
# KEYE Results – GW USR (224 Gb/s i/f)



- Simulation conditions:
  - Glasswing @45 GBd (for 224 Gb/s i/f)
  - Tx Launch: 200 mVppd
  - 1-tap FFE
  - No CTLE or DFE
  - USR: (1 cm, no Pkg)
  - BER = 1E-15
- Eye Width/Height:
  - 41.7 mV (SC#0)
  - 0.68 UI
  - Eye open



# KEYE Results – GW XSR (112 Gb/s i/f)



- Simulation conditions:
  - Glasswing @22 GBd (for 112 Gb/s i/f)
  - Tx Launch: 200 mVppd
  - 1-tap FFE
  - No CTLE or DFE
  - XSR: (5 cm, w/Pkg)
  - BER = 1E-15
- Eye Width/Height:
  - 26.7 mV (SC#0)
  - 0.58 UI (SC#2,#3)
  - Eye open



# Signal Integrity Conclusions

- NRZ @56 GBd: Open eye exists on XSR channels with 200 mVppd launch.
- ENRZ @37 GBd: Open eye exists on XSR channels with 200 mVppd launch.
- EP3L @28 GBd: Open eye exists on XSR channels with 200 mVppd launch.
- Glasswing @45 GBd
  - Open eye on XSR channels with 400 mVppd launch.
  - Open eye on USR channels with 200 mVppd launch.  
(Allowing additional power savings.)



# Power Analysis Methodology

- Purpose: Benchmark power of chord signaling options to an equivalent NRZ reference design.
- Methodology applied:
  - Kandou Wasp chip used as reference design for Serdes circuit and logic blocks (TSMC 28 nm, 28 GBd).
  - Spice simulations used to determine power for circuit blocks of the reference design. Logic block power estimated based on synthesis results.
  - Reduce block functionality (and remove power) consistent with short reach applications:
    - Lower Tx launch amplitude.
    - 1-tap FFE, no DFE, no CTLE
    - Share forwarded clock or CDR, DLLs, etc. across all lanes.
  - Determine rules for scaling each block to other frequencies.
  - Scale baseline power of each block to TSMC 16 nm process.
  - Equivalent circuit architecture assumptions are used for all codes at all baud rates. (This avoids biasing results with architecture choices.)

Benchmark: NRZ is used as a benchmark for the analysis methodology.



# 112 Gb/s Interface Power

	Drivers	Receivers	Clock Trees	Logic	Width	Power
4 x 28G NRZ (baseline)	LVDS, 200 mVppd, TSMC 28nm	CDR (shared)	Baseline design	Baseline design	4	2.46 pJ/bit
2 x 56G NRZ, CDR	LVDS, 200 mVppd	CDR (shared)	1X Area	Assume similar	2	2.28 pJ/bit
2 x 56G NRZ, Fwd Clk	LVDS, 200 mVppd, 2 chan. + clock	Clk Rx (shared)	1X Area	Assume similar	2	2.72 pJ/bit
1 x 37G ENRZ, CDR	LVDS, 200 mVppd, 4 wire, 4 level	4-wire, 3-comp, CDR	2X Area	3X #chan	1	1.98 pJ/bit or 1.61 pJ/bit w/o Tx PLL
1 x 37G ENRZ, Fwd Clk	LVDS, 200 mVppd, 4 wire, 4 lvl, + clock	4-wire, 3-comp, plus diff. clock	2X Area	3X #chan	1	2.12 pJ/bit
1 x 28G EP3L, CDR	LVDS, 200 mVppd, 4 wire, 5 level	4-wire, 3-comp, CDR	2X Area	4X #chan	1	1.71 pJ/bit or 1.39 pJ/bit w/o Tx PLL
1 x 28G EP3L, Fwd Clk	LVDS, 400 mVppd, 4 wire, 5 lvl + clock	4-wire, 3-comp, plus diff. clock	2X Area	4X #chan	1	1.74 pJ/bit
1 x 22.4 GW, CDR, XSR or USR	CML, 200 mVppd, 6 wire, 3 level	6-wire, 5-comp., CDR	3X Area	5X #chan	1	1.13 pJ/bit
1 x 22.4 GW, Fwd Clk, XSR or USR	CML, 200 mVppd, 6 wire, 3 lvl, + clock	6-wire, 5-comp. plus diff. clock	3X Area	5X #chan	1	1.01 pJ/bit or 0.72 pJ/bit w/o Tx PLL





# 224 Gb/s Interface Power

	Drivers	Receivers	Clock Trees	Logic	Width	Power
8 x 28G NRZ (baseline)	LVDS, 200 mVppd, TSMC 28nm	CDR (shared)	Baseline design	Baseline design	8	2.41 pJ/bit
4 x 56G NRZ, CDR	LVDS, 200 mVppd	CDR (shared)	Assume same area	Assume similar	4	2.21 pJ/bit
4 x 56G NRZ, Fwd Clk	LVDS, 200 mVppd, 2 chan. + clock	Clk Rx (shared)	Assume same area	Assume similar	4	2.36 pJ/bit
2 x 37G ENRZ, CDR	LVDS, 200 mVppd, 4 wire, 4 level	4-wire, 3-comp, CDR	2X Area	3X #chan	2	1.74 pJ/bit
2 x 37G ENRZ, Fwd Clk	LVDS, 200 mVppd, 4 wire, 4 lvl, + clock	4-wire, 3-comp, plus diff. clock	2X Area	3X #chan	2	1.71 pJ/bit or 1.52 pJ/bit w/o Tx PLL
2 x 28G EP3L, CDR	LVDS, 200 mVppd, 4 wire, 5 level	4-wire, 3-comp, CDR	2X Area	4X #chan	2	1.51 pJ/bit
2 x 28G EP3L, Fwd Clk	LVDS, 200 mVppd, 4 wire, 5 lvl, + clock	4-wire, 3-comp, plus diff. clock	2X Area	4X #chan	2	1.41 pJ/bit or 1.25 pJ/bit w/o Tx PLL
1 x 44.8 GW, CDR, XSR	CML, 400 mVppd, 6 wire, 3 level	6-wire, 5-comp., CDR	3X Area	5X #chan	1	1.13 pJ/bit
1 x 44.8 GW, Fwd Clk, XSR	CML, 400 mVppd, 6 wire, 3 lvl, + clock	6-wire, 5-comp. plus diff. clock	3X Area	5X #chan	1	1.00 pJ/bit or 0.80 pJ/bit w/o Tx PLL
1 x 44.8 GW, CDR, USR	CML, 200 mVppd, 6 wire, 3 level	6-wire, 5-comp., CDR	3X Area	5X #chan	1	1.06 pJ/bit
1 x 44.8 GW, Fwd Clk, USR	CML, 200 mVppd, 6 wire, 3 lvl, + clock	6-wire, 5-comp. plus diff. clock	3X Area	5X #chan	1	0.93 pJ/bit or 0.73 pJ/bit w/o Tx PLL





# Power Summary

	112G I/F	224G I/F
28G NRZ Baseline (28 nm)	2.46 pJ/bit	2.21 pJ/bit
56G NRZ	2.28 – 2.72 pJ/bit	2.21 – 2.36 pJ/bit
37G ENRZ	1.61 – 2.12 pJ/bit	1.52 – 1.74 pJ/bit
28G EP3L	1.39 – 1.74 pJ/bit	1.25 – 1.51 pJ/bit
22.4 / 44.8 Glasswing - XSR	0.72 – 1.13 pJ/bit	0.80 – 1.13 pJ/bit
22.4 / 44.8 Glasswing - USR	0.72 – 1.13 pJ/bit	0.73 – 1.06 pJ/bit

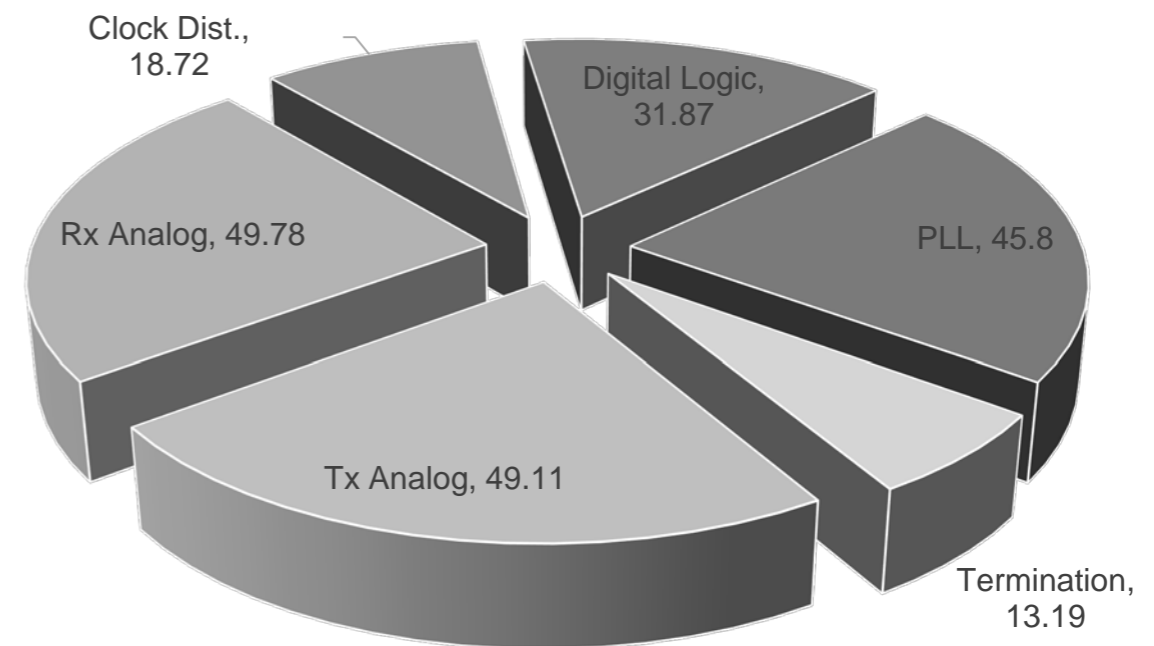
- NRZ power may be lower than existing designs:
  - Assumed shared CDR, etc to avoid bias toward multiwire codes.
  - Excluded PLL from NRZ analysis.
- Forwarded clock requires offers some power savings but is offset by additional driver and termination power.
- Glasswing offers potential for extremely low power:
  - XSR Interfaces: 1.12 pJ/bit (or 0.80 pJ/bit w/o Tx PLL)
  - USR Interfaces: 0.93 pJ/bit (or 0.73 pJ/bit w/o Tx PLL)



# Glasswing 224 Gb/s USR Power Estimate

- Typical Power is 0.93 pJ/bit @44.8 Gbd.
- Power estimate based on:
  - Single Glasswing 5b/6w channel (full duplex) plus forwarded clock
  - USR channel
  - 44.8 GBd
  - 16 nm process
- Architecture features to minimize power for USR applications:
  - Reduced Tx amplitude (200 mVppd)
  - Forwarded Clock
  - 1 tap FFE, no DFE

Power Breakdown (mW)



<b>Power Total (mW)</b>		208.49
<b>Data Throughput (Gb/s)</b>		224.00
<b>Energy per Bit (pJ/bit)</b>		0.93



- XSR / USR Requirements from oif2014.029:
  - CEI 56G XSR must be below 5 pJ/bit.
  - CEI 56G USR must be below 3 pJ/bit.
- Multiple codes exist which are lower power than NRZ baseline.
  - EP3L 4b4w code does not require gearbox when used in data paths that are multiples of 4 bits wide.
- Glasswing offers potential for extremely low power:
  - XSR / USR interfaces on the order of 1.00 pJ/bit.
  - 224 Gb/s USR i/f using clock forwarding: 0.93 pJ/bit.
    - 0.73 pJ/bit excluding Tx PLL
  - Glasswing roadmap exists to support the next generation bits/wire interfaces.



# KANDOU

reinventing the

# BUS

