Design Techniques for Scalable, Sub-pJ/b Serial I/O Transceivers

Samuel Palermo

spalermo@tamu.edu

Analog & Mixed-Signal Center

Texas A&M University

Outline

- Motivation
- Power-Scalable I/O Techniques
- Low-Power Clocking
- Low-Power Equalizers
- Conclusion

More and More Data …

Human-driven traffic growth

Hi definition video conference

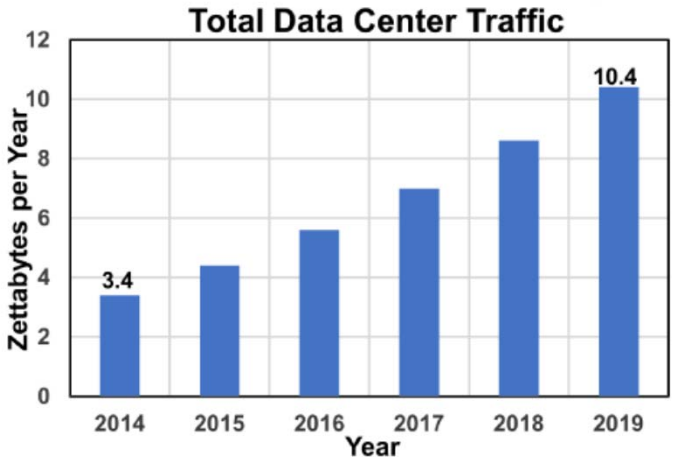

Source: Cisco Global Cloud Index, 2014 - 2019

Machine-driven traffic growth

Cloud service, big data, IoT

Enterprise service

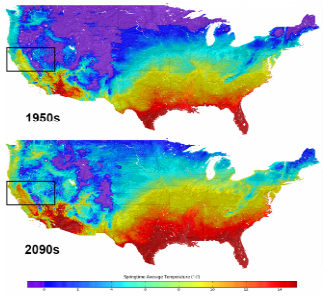

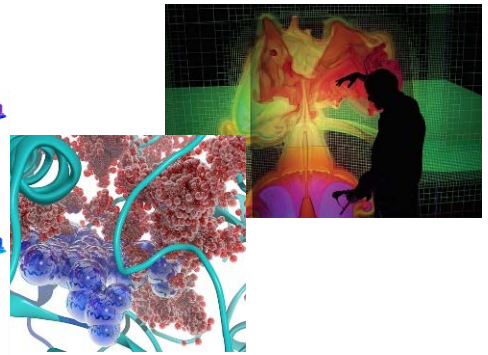

Supercomputer

High-Speed Serial I/O

- Found in applications ranging from high-end computing systems to smart mobile devices
- Typical processor platform
  - Processor-to-memory: DDR3
  - Processor-to-peripheral: PCIe & USB
  - Storage: SATA
  - Network: LAN
- Mobile systems
  - DSI : Display Serial Interface
  - CSI : Camera Serial Interface
  - UniPRO : MIPI Universal Protocol

Intel IvyBridge w/ Chipset

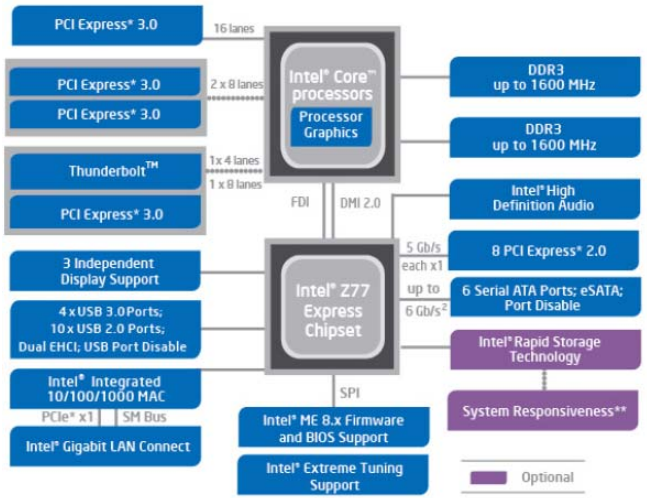

High-Speed Electrical Link System

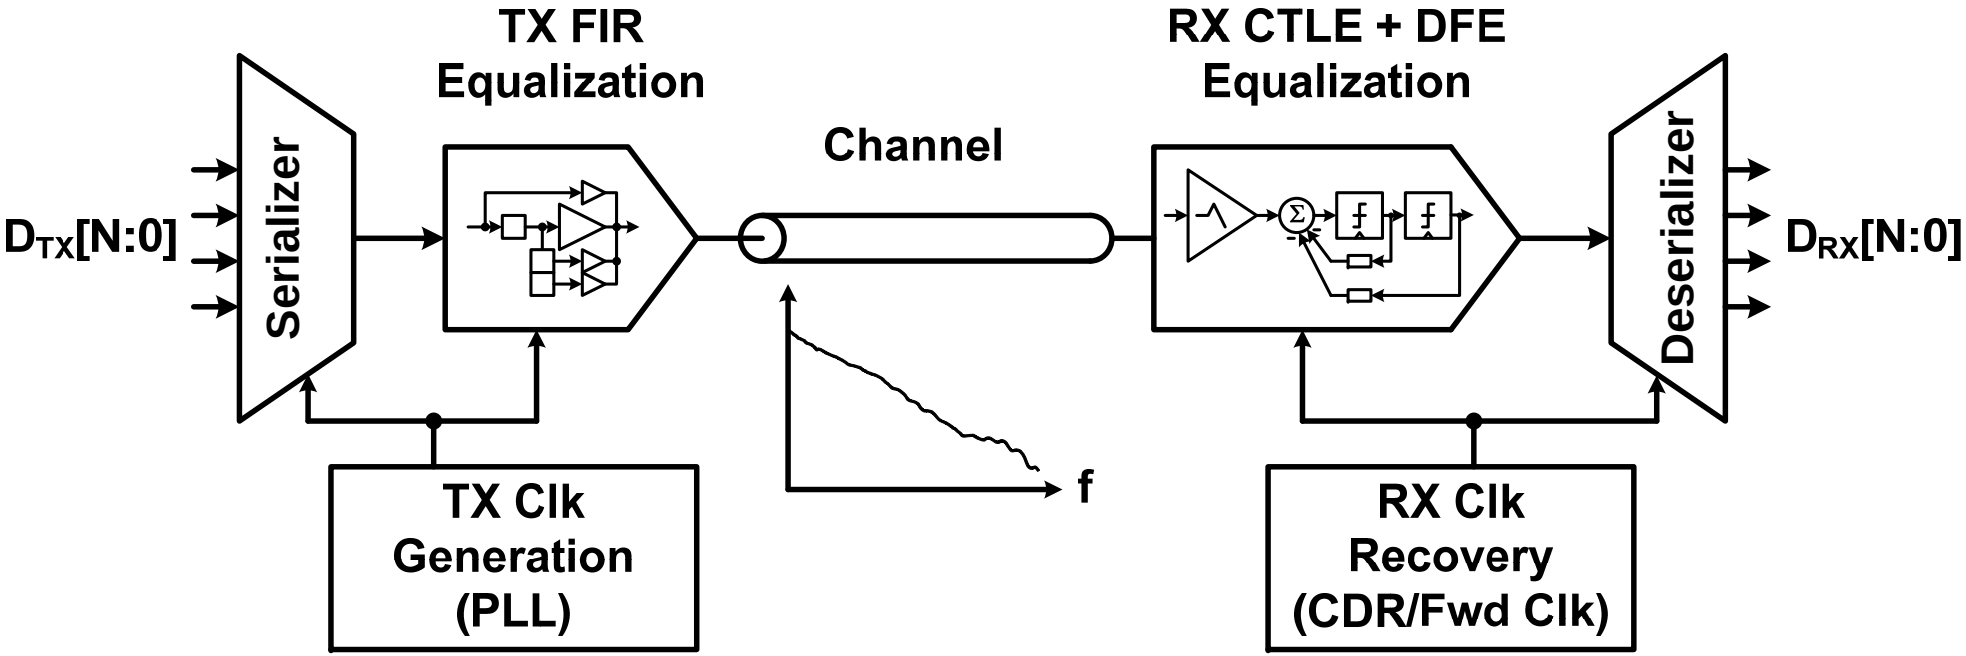

- Data serialization required due to limited I/O channel count
- Future systems demand efficient high-speed drivers, receivers, and clock generation/recovery circuitry
- Equalization circuitry compensates for high frequency channel loss

I/O Energy Efficiency is Paramount

- High-performance processor aggregate I/O bandwidth demands will soon approach 1TB/s
- Typical I/O power budgets are 10W or less
- Energy efficiencies near 1mW/Gbps are necessary

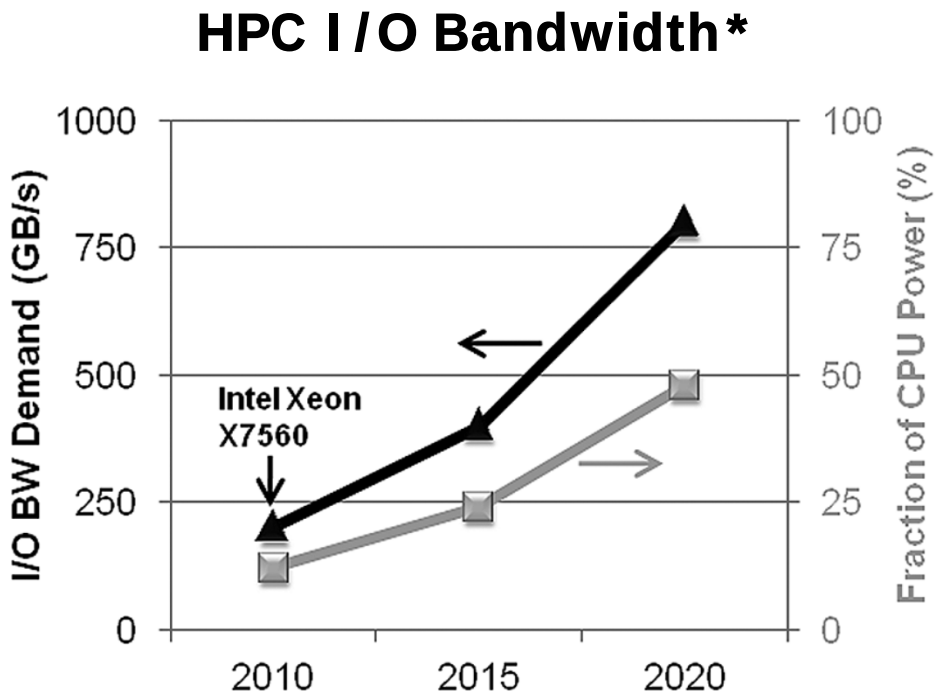

*M. Mansuri *et al*, "A Scalable 0.128–1 Tb/s, 0.8–2.6 pJ/bit, 64-Lane Parallel I/O in 32-nm CMOS," *IEEE JSSC*, Dec. 2013.

Outline

- Motivation
- Power-Scalable I/O Techniques
- Low-Power Clocking
- Low-Power Equalizers
- Conclusion

Scaling Supply with Data Rate

- Adaptive power supply regulation allows the minimum voltage required for a given data rate
- Efficient DC-DC converters driven by a frequency controller generate the supply voltage for the I/O clocking and serialization
- Dramatic energy efficiency improvements possible, particularly as data rates scale down based on I/O bandwidth demand

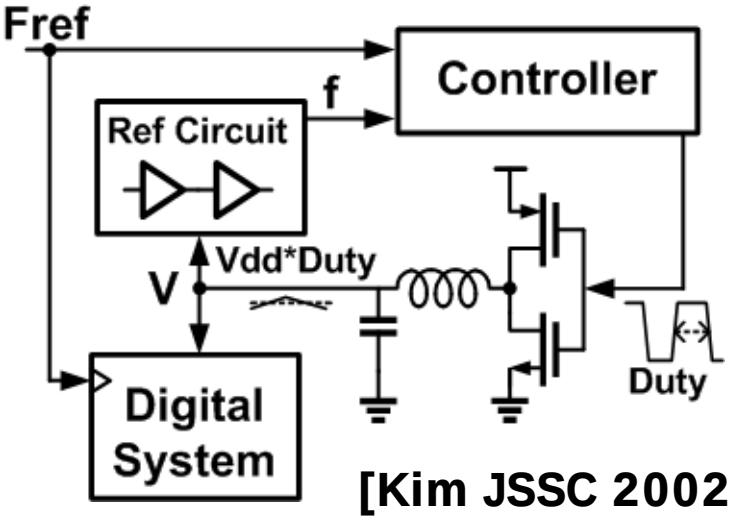

[Kim JSSC 2002]

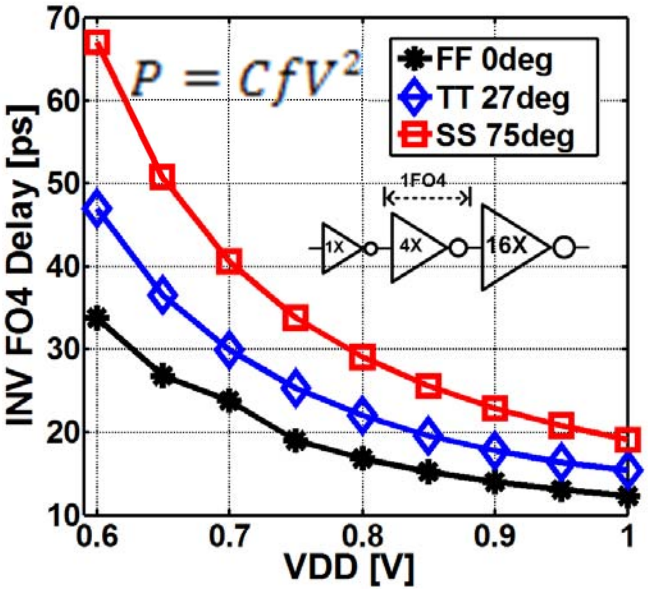

* GP 65nm CMOS Technology

Increasing Data Rate with Parallelism

$$\text{Data Rate} = N \times f_{clk}$$

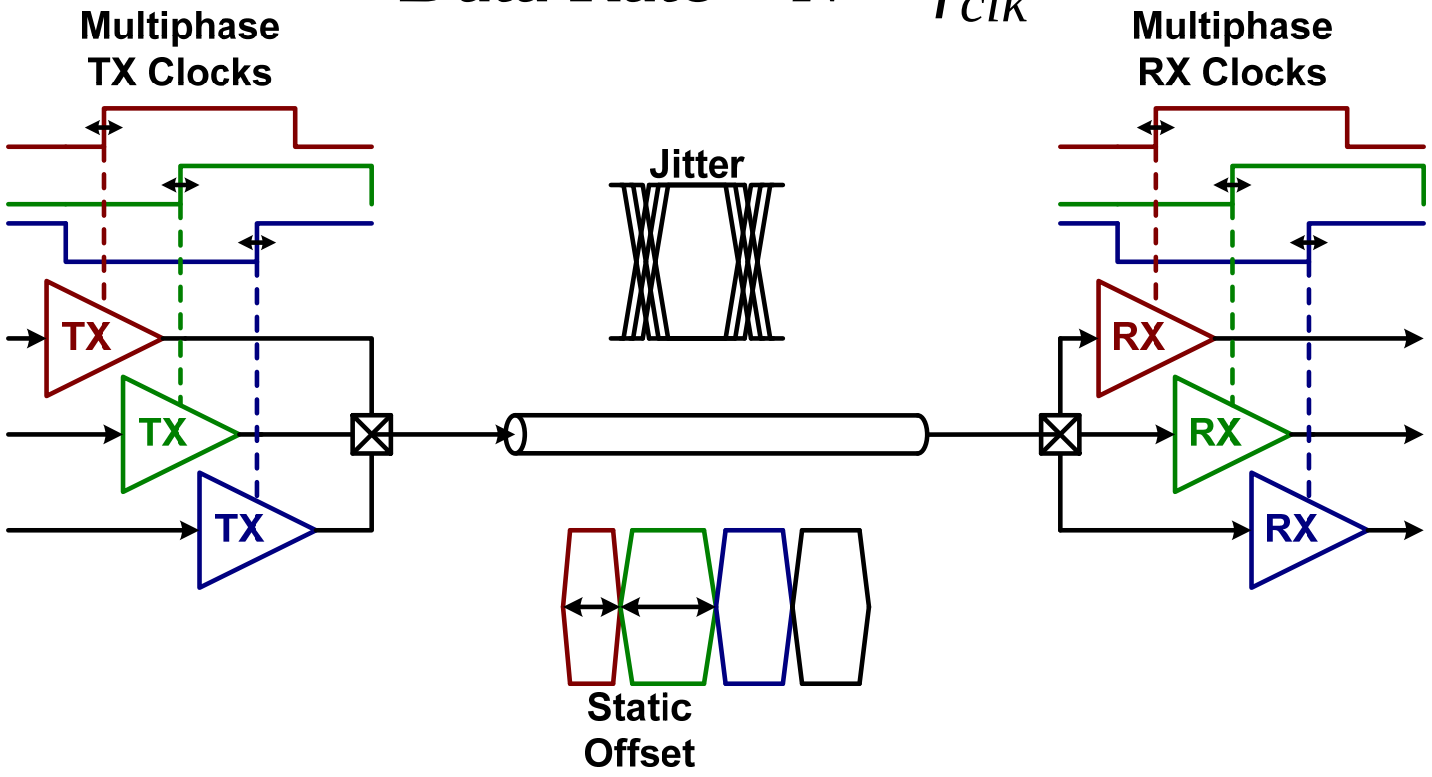

- Utilizing large mux/demux factors allows parallel segments to operate at low clock frequencies and low supply voltages
- Important to minimize jitter and static phase offset of multiple clock phases

Fast Power-State Transitioning

- Efficient system operation demands minimal latency when adjusting the I/O per-channel data rates
- Certain applications, such as memory interfaces, have bursts of data traffic which necessitate rapidly achieving maximum I/O bandwidth
- Techniques must be developed to enable fast power-state transitioning of key I/O circuits

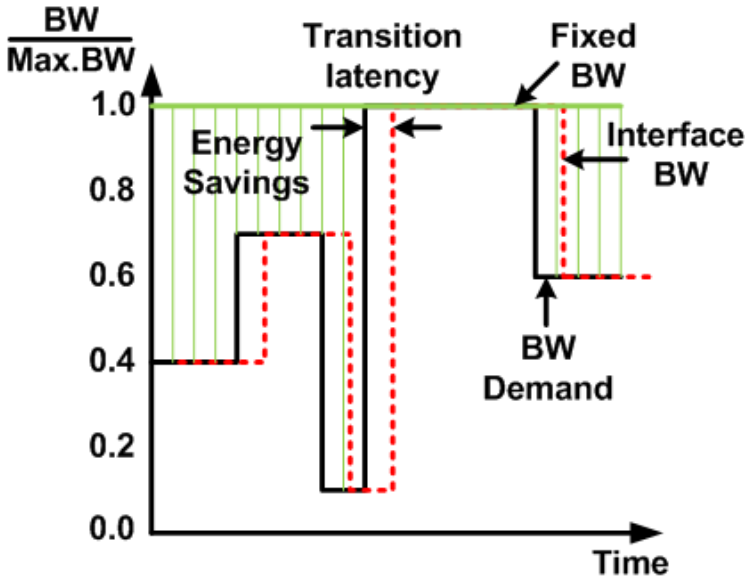

[O'Mahony VLSI-DAT 2009]

Low-Swing TX Driver Comparison

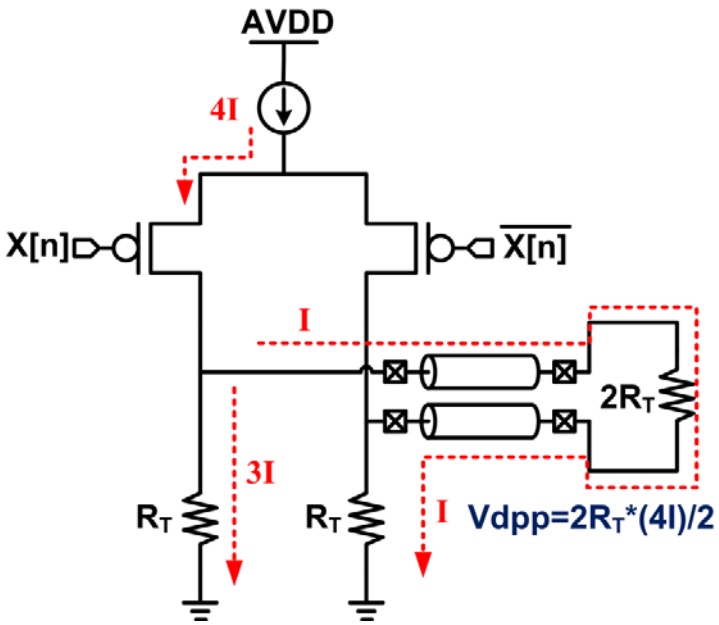

Current-Mode Driver (CM)

- ☺ Norton-equivalent parallel termination
- ☺ High PSRR
- ☺ Low pre-driver complexity
- ☹ High signaling power

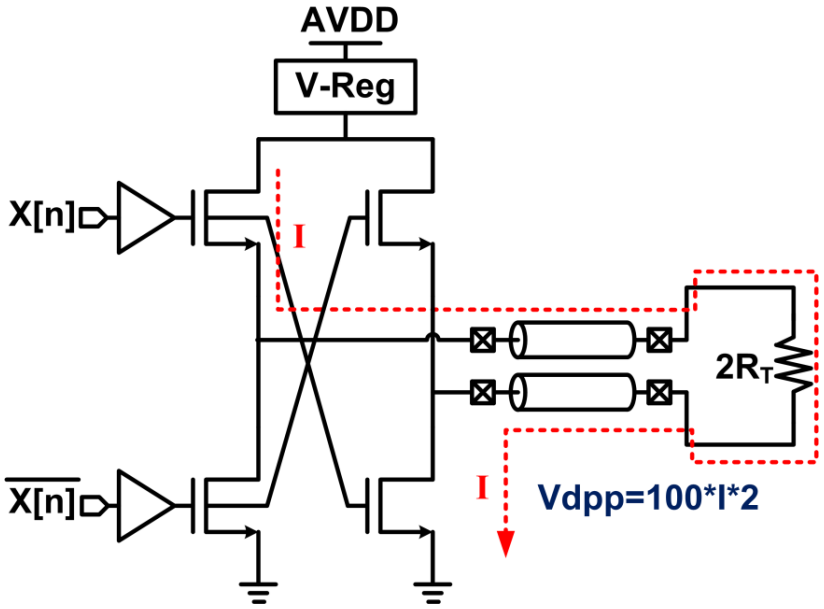

Voltage-Mode Driver (VM)

- ☹ Thevenin-equivalent series termination
- ☹ Voltage-regulator is required
- ☺ ¼ signaling power of CM

VM driver uses 4X less current than CM driver

Low-Voltage Serial I/O Transceiver

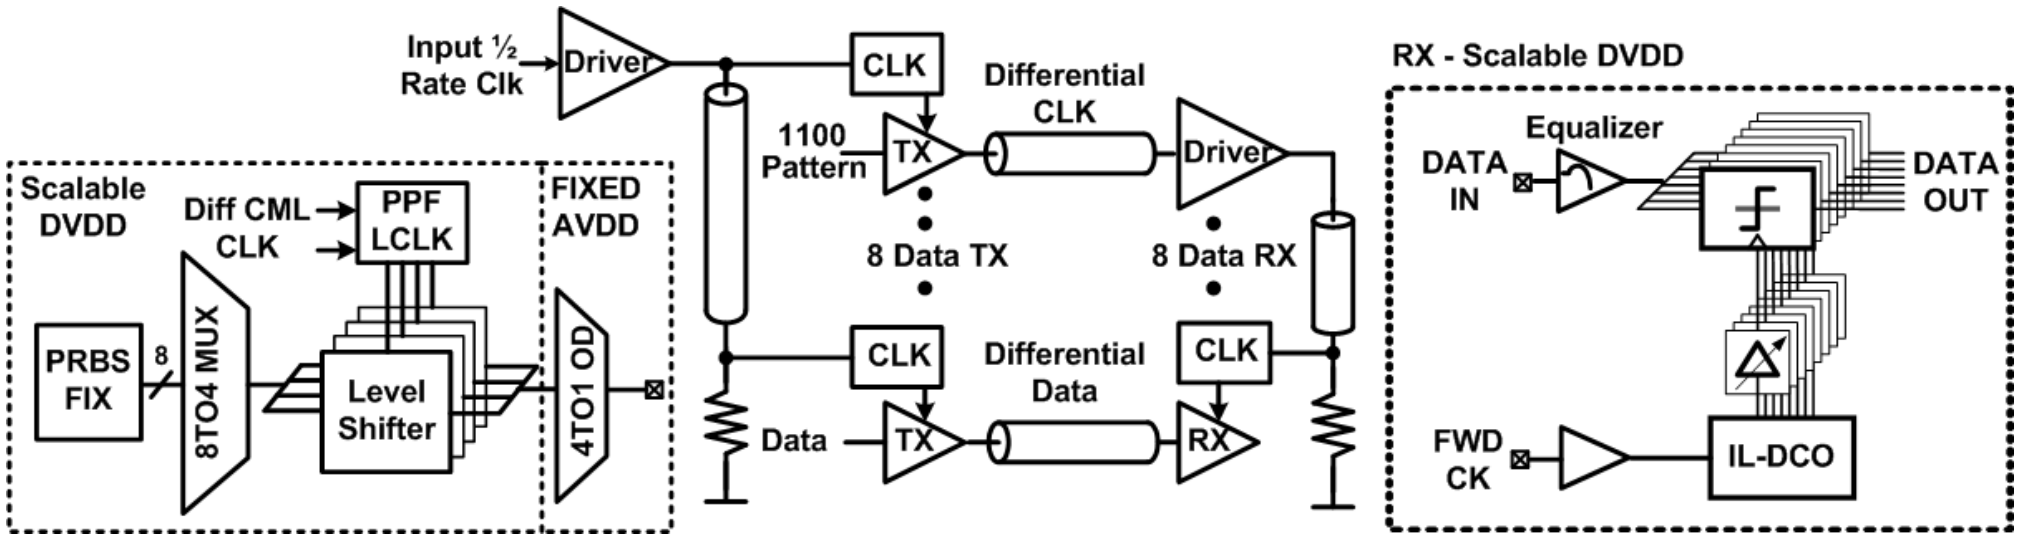

- Utilizes a high TX output multiplexing (4:1) and RX input multiplexing (1:8) factor for low-voltage operation

Y.-H. Song, R. Bai, P. Chiang, and S. Palermo, "A 0.47-0.66pJ/bit, 4.8-8Gb/s I/O Transceiver in 65nm-CMOS," IEEE JSSC, vol. 48, no. 5, pp. 1276-1289, May 2013.

4:1 Output Multiplexing Voltage-Mode TX

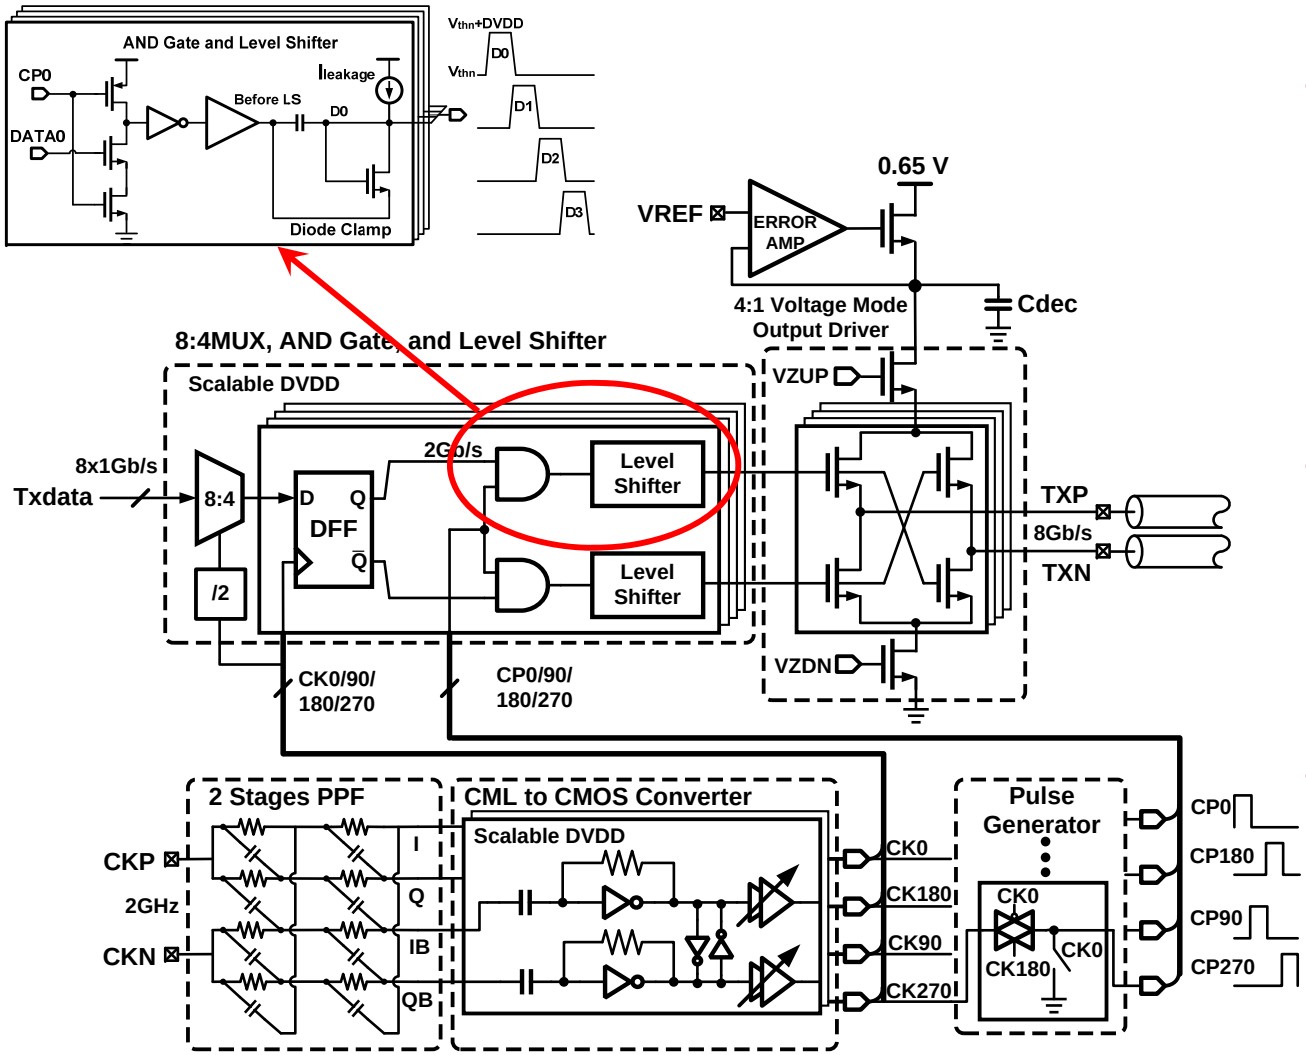

- 4 parallel voltage-mode output segments perform output multiplexing
- Efficient quadrature clock generation with 2-stage poly-phase filter
- Level-shifting pre-driver allows for smaller output transistors

Y.-H. Song, R. Bai, P. Chiang, and S. Palermo, "A 0.47-0.66pJ/bit, 4.8-8Gb/s I/O Transceiver in 65nm-CMOS," IEEE JSSC, vol. 48, no. 5, pp. 1276-1289, May 2013.

1:8 Input De-Multiplexing RX

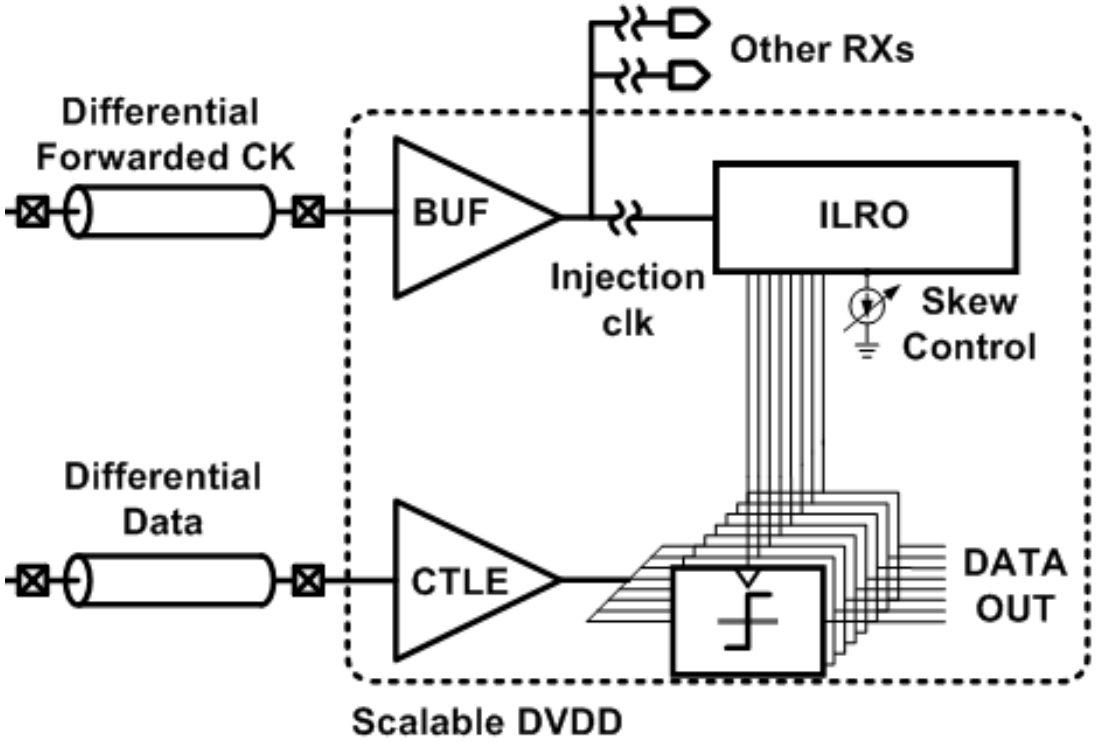

- 1:8 input de-multiplexing allows input comparators to operate at low voltages
- Injection-locked-oscillator is used for efficient multi-phase clock generation and de-skew

Y.-H. Song, R. Bai, P. Chiang, and S. Palermo, "A 0.47-0.66pJ/bit, 4.8-8Gb/s I/O Transceiver in 65nm-CMOS," IEEE JSSC, vol. 48, no. 5, pp. 1276-1289, May 2013.

0.47-0.66pJ/bit, 4.8-8Gb/s GP 65nm CMOS Prototype

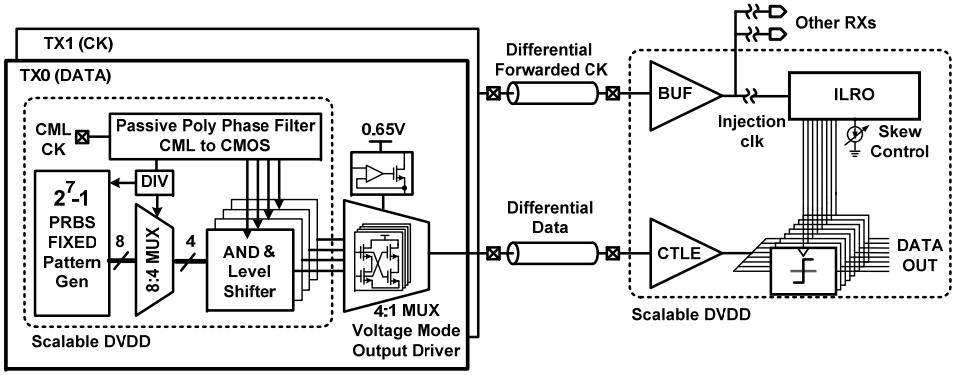

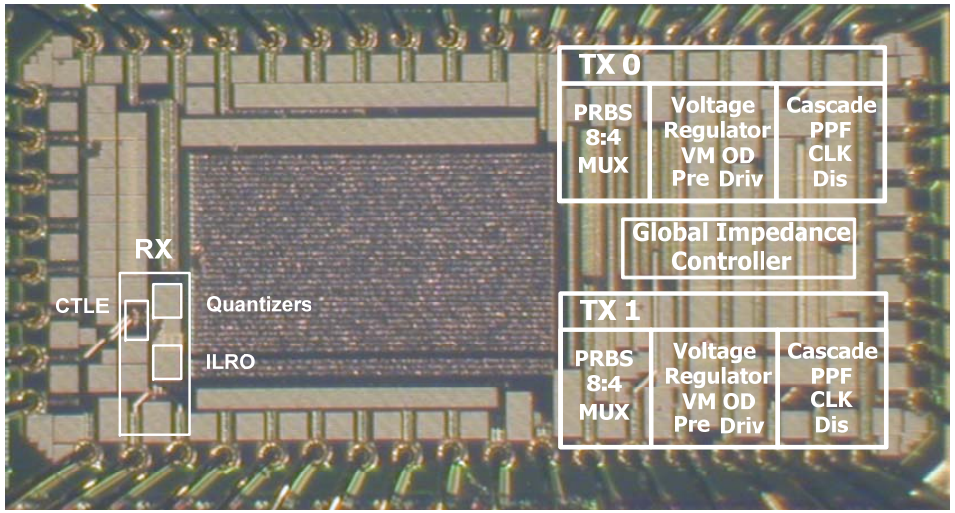

Testing with 20cm FR-4 Channel

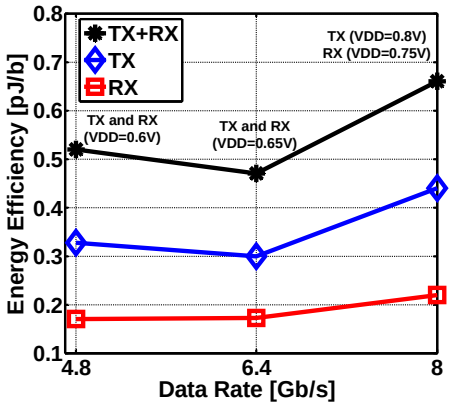

TX Power Breakdown (6.4Gb/s at 0.65V)	
LDO & Output Driver (150mV$_{ppd}$)	793uW
Serializer, Pre-drivers, Clocking	933uW
Global Impedance Control (amortized across 9 TX)	193uW
TX Energy Efficiency	0.3pJ/b
RX Power Breakdown (6.4Gb/s at 0.65V)	
CTLE, Quantizers, ILRO	1.07mW
Clock Distribution	38uW
RX Energy Efficiency	0.17pJ/b
Total Energy Efficiency	**0.47pJ/b**

- Optimal 0.47pJ/b energy efficiency achieved at 6.4Gb/s
  - At low data rates, less amortization of static current
  - At high data rates, higher voltage required for serialization timing

Y.-H. Song, R. Bai, P. Chiang, and S. Palermo, "A 0.47-0.66pJ/bit, 4.8-8Gb/s I/O Transceiver in 65nm-CMOS," IEEE JSSC, vol. 48, no. 5, pp. 1276-1289, May 2013.

Outline

- Motivation
- Power-Scalable I/O Techniques
- Low-Power Clocking
- Low-Power Equalizers
- Conclusion

Low-Power Transmitter Clocking

- Transmitters which utilize voltage-scaling to save power require efficient generation of multi-phase clocks
- Key issue is the extreme phase variations faced with low-voltage operation

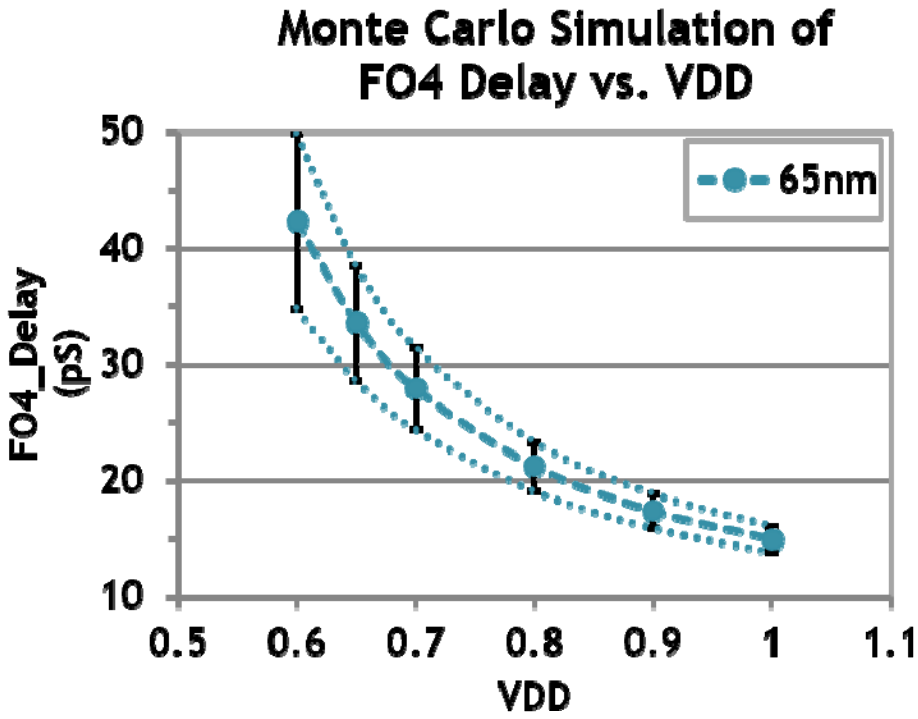

Passive Poly-Phase Filter Clock Generation

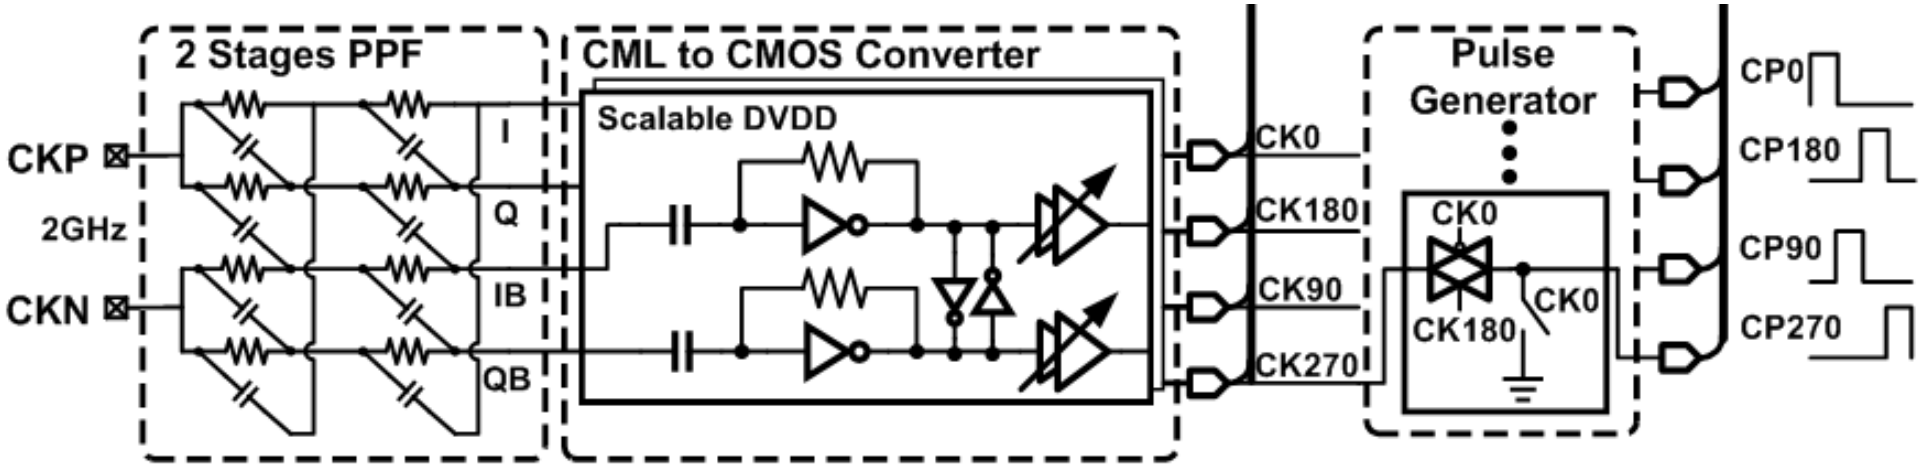

- 2-stage passive poly-phase filter generates 4 clock phases for output multiplexing from low-swing global TX ¼-rate differential clocks

- Requires subsequent CML2CMOS converter to generate TX clocks

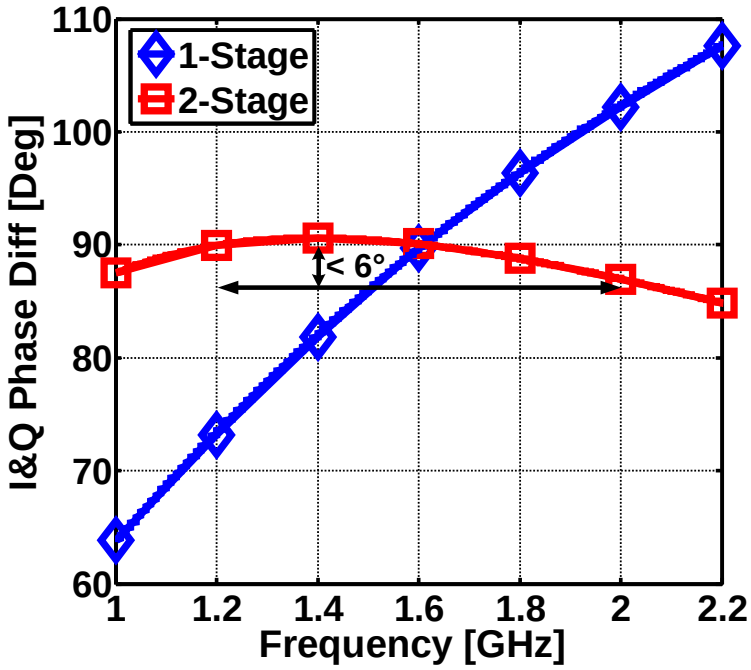

Y.-H. Song, R. Bai, P. Chiang, and S. Palermo, "A 0.47-0.66pJ/bit, 4.8-8Gb/s I/O Transceiver in 65nm-CMOS," IEEE JSSC, vol. 48, no. 5, pp. 1276-1289, May 2013.

Injection-Locked Oscillator (ILO) Clock Generation

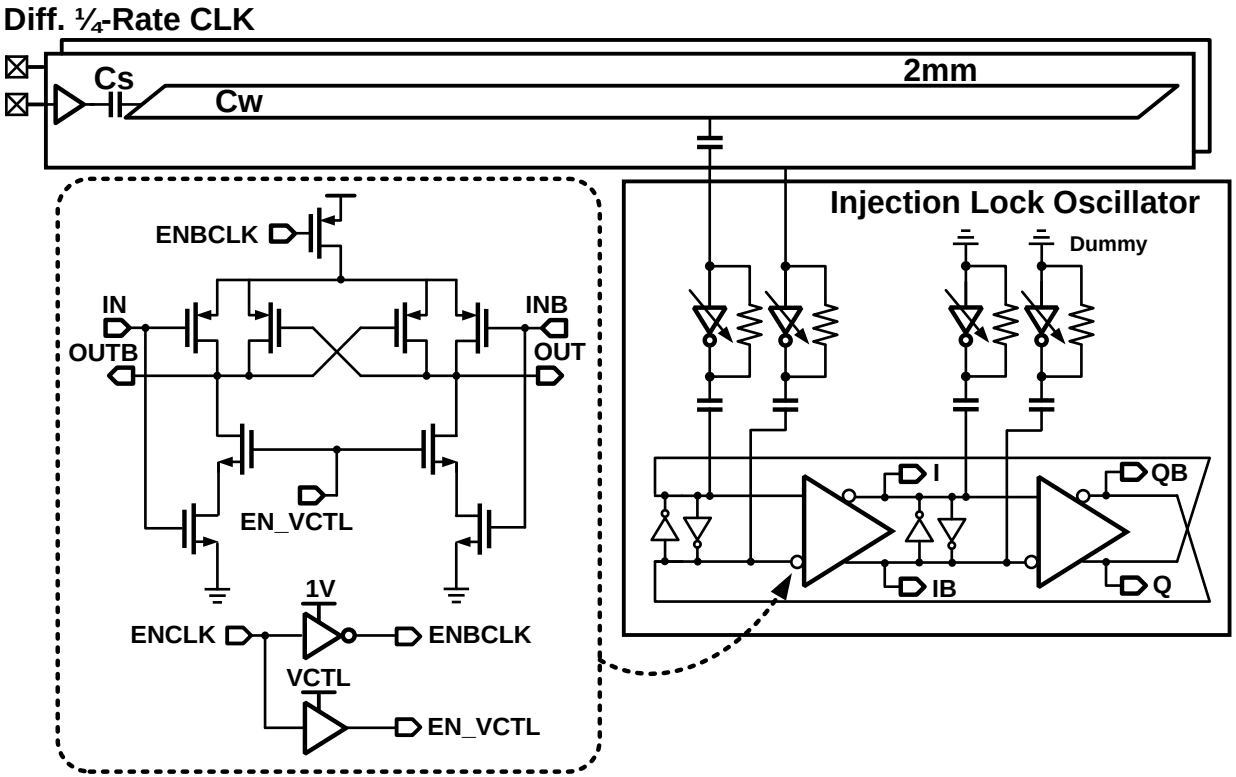

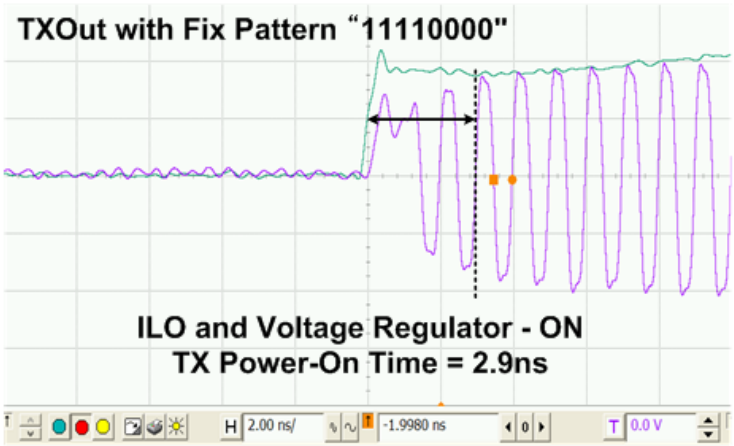

- 4-phase CLK generation by ILO
  - Eliminates CML2CMOS convertor
  - Fine frequency control by EN_VCTL also enables fast power state transition

Y.-H. Song, H.-W. Yang, H. Li, P. Chiang, and S. Palermo, "An 8–16 Gb/s, 0.65–1.05 pJ/b, Voltage-Mode Transmitter With Analog Impedance Modulation Equalization and Sub-3 ns Power-State Transitioning," IEEE JSSC, vol. 49, no. 11, pp. 2631-2643, Nov. 2014.

Async. Sampling Based Phase Calibration

- Compensates for deterministic jitter (DJ) due to duty-cycle distortion (DCD) and phase mismatches of quadrature clocks

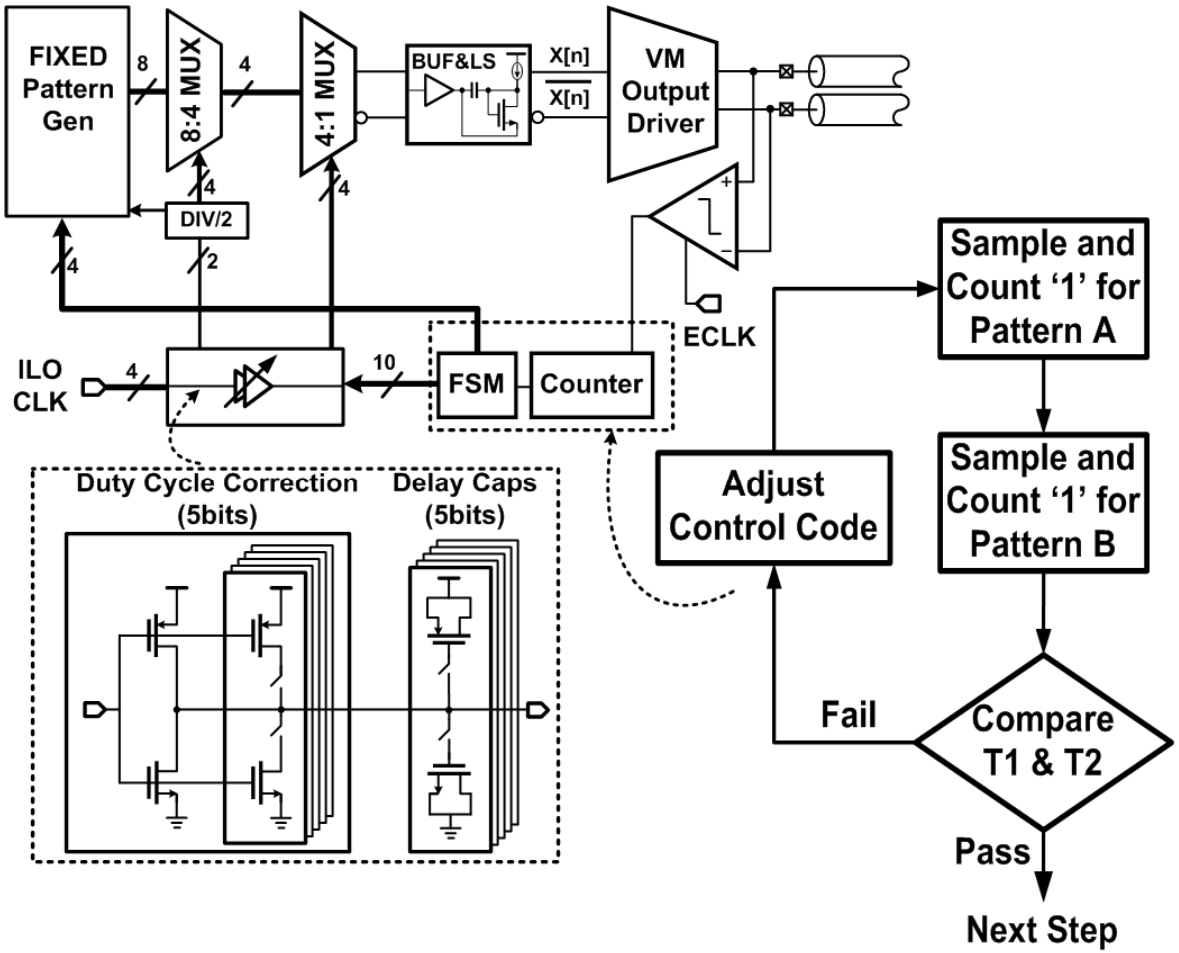

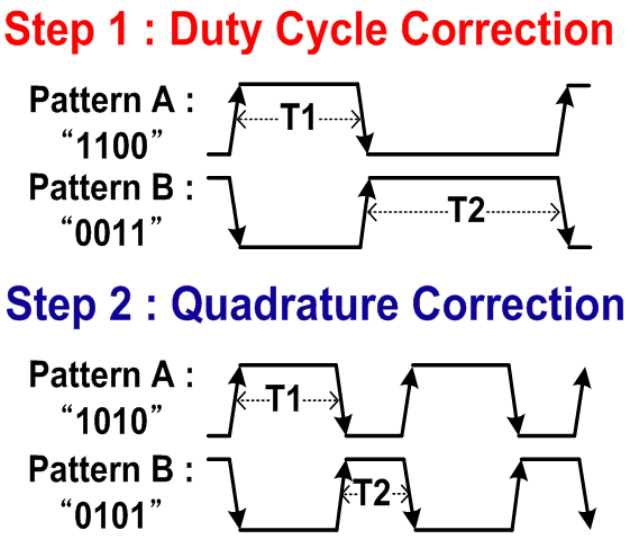

Y.-H. Song, H.-W. Yang, H. Li, P. Chiang, and S. Palermo, "An 8–16 Gb/s, 0.65–1.05 pJ/b, Voltage-Mode Transmitter With Analog Impedance Modulation Equalization and Sub-3 ns Power-State Transitioning," IEEE JSSC, vol. 49, no. 11, pp. 2631-2643, Nov. 2014.

Automatic Phase Correction

- Eye diagrams without and with phase calibration

8Gb/s **16Gb/s**

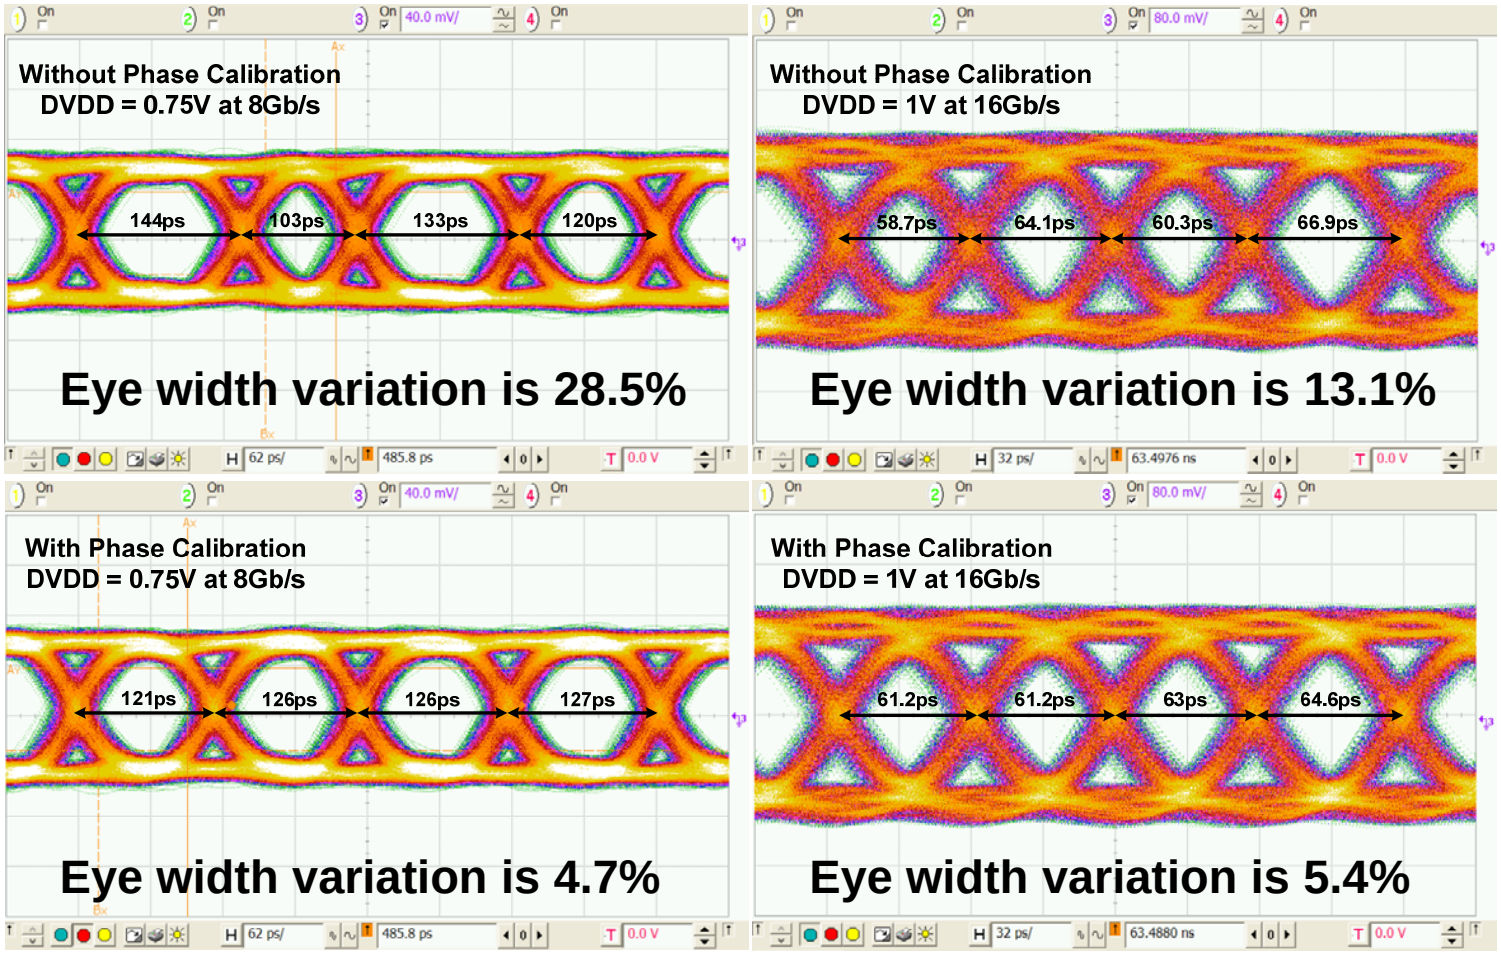

Y.-H. Song, H.-W. Yang, H. Li, P. Chiang, and S. Palermo, "An 8–16 Gb/s, 0.65–1.05 pJ/b, Voltage-Mode Transmitter With Analog Impedance Modulation Equalization and Sub-3 ns Power-State Transitioning," IEEE JSSC, vol. 49, no. 11, pp. 2631-2643, Nov. 2014.

RX-Forwarded Clock I/O De-Skew

- "Coherent" clocking allows jitter tracking, but still need to employ per-channel de-skew to maximize timing margins

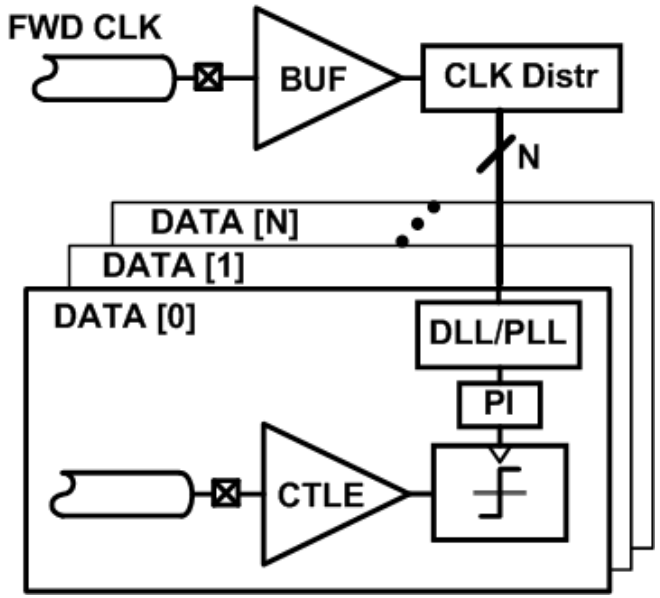

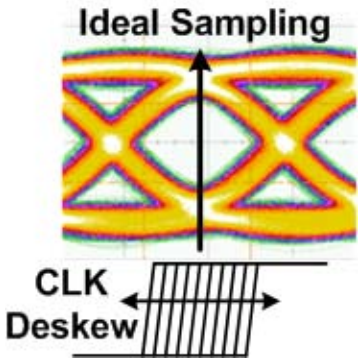

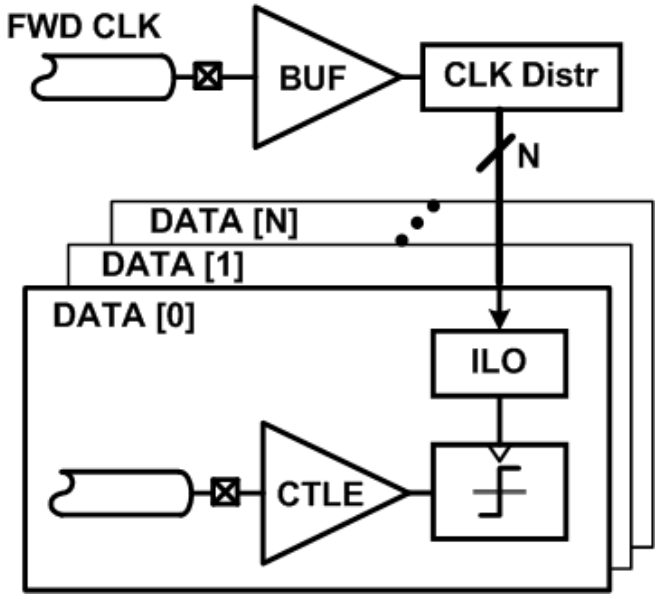

DLL/PLL + Phase Interpolator (PI)

- DLL can have jitter amplification, while PLL can have jitter accumulation
- Both circuits can occupy significant area

Injection-Locked Oscillator (ILO)

- Compact low-power implementation
- High jitter tracking bandwidth

ILO-Based De-Skew

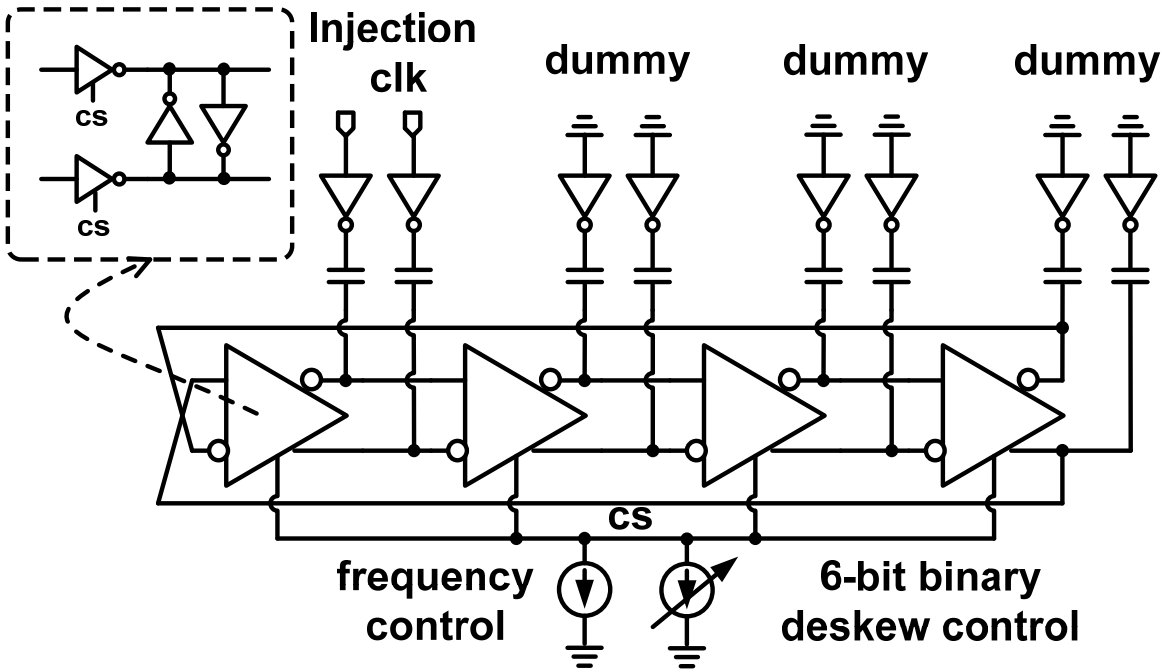

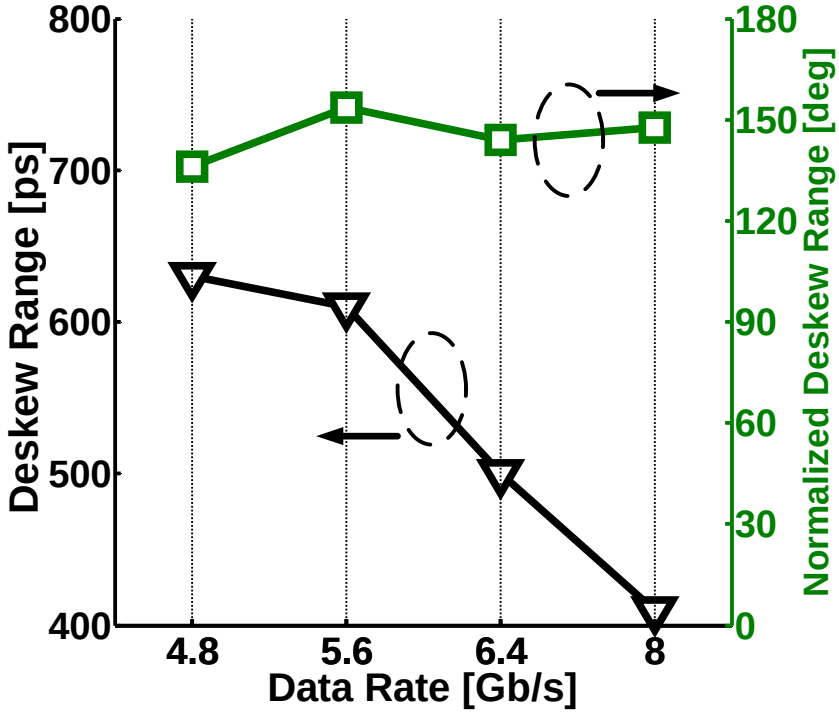

- Current-starved inverter-based ILO produces the multiple clock phases necessary for the receiver samplers
- Fine de-skew control by 6-bit binary current mirror which changes ILO free-running frequency

Y.-H. Song, R. Bai, P. Chiang, and S. Palermo, "A 0.47-0.66pJ/bit, 4.8-8Gb/s I/O Transceiver in 65nm-CMOS," IEEE JSSC, vol. 48, no. 5, pp. 1276-1289, May 2013.

Phase Drifts with ILO-Based Clocking

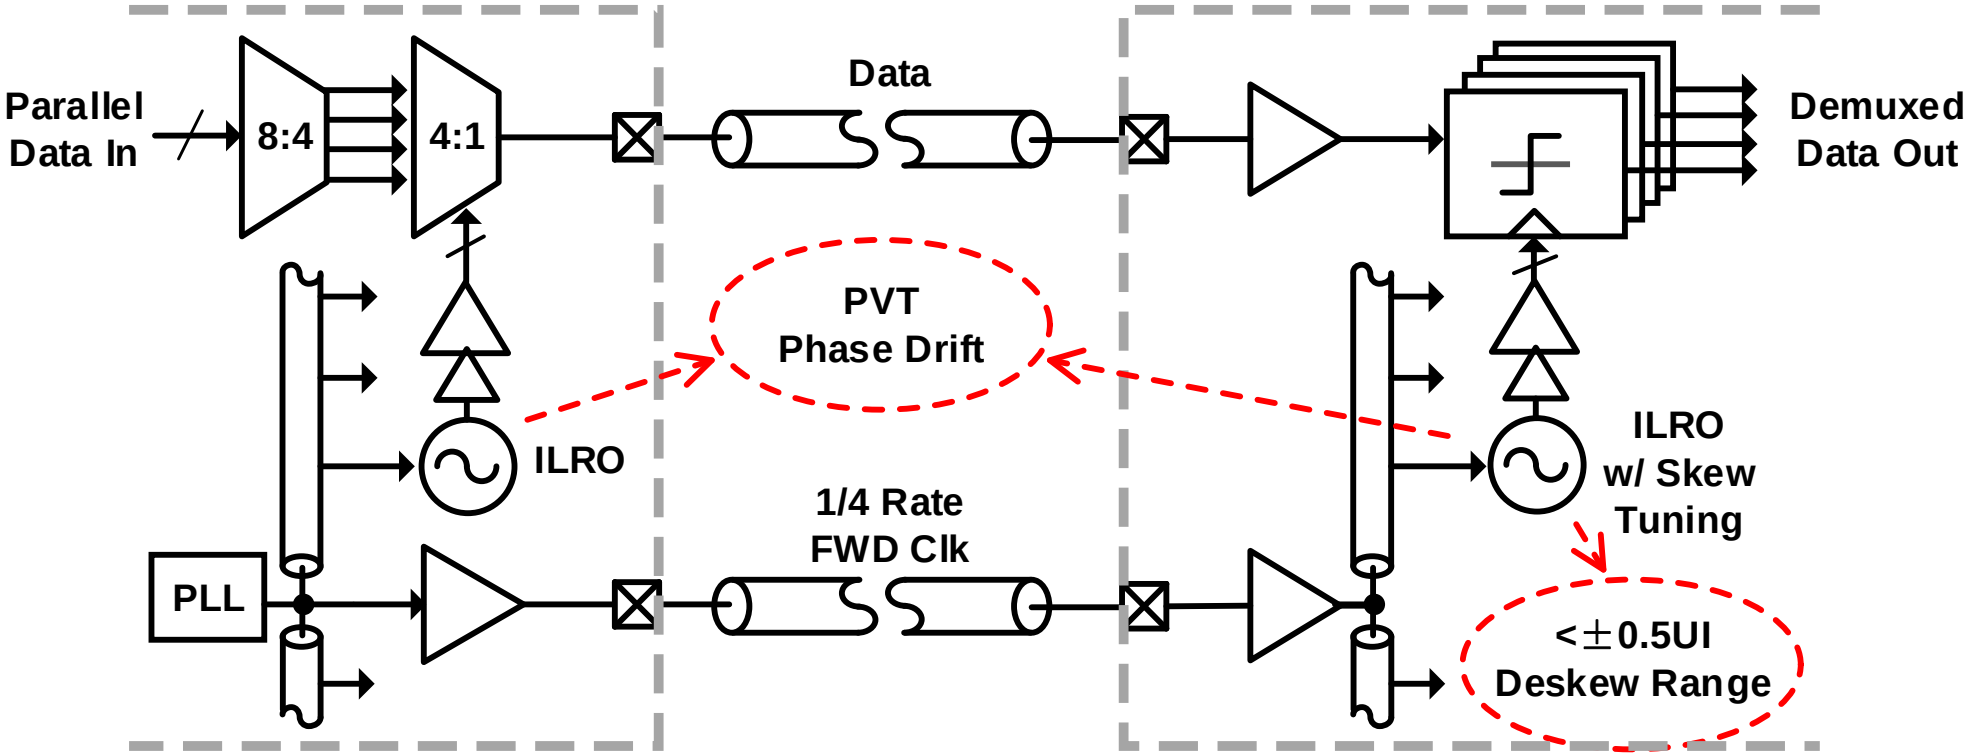

- Voltage and temperature variations can cause the TX/RX ILOs' free running frequency to change, and thus the phase relationship can drift with time

Low-Overhead CDR w/ ILO-Based De-Skew

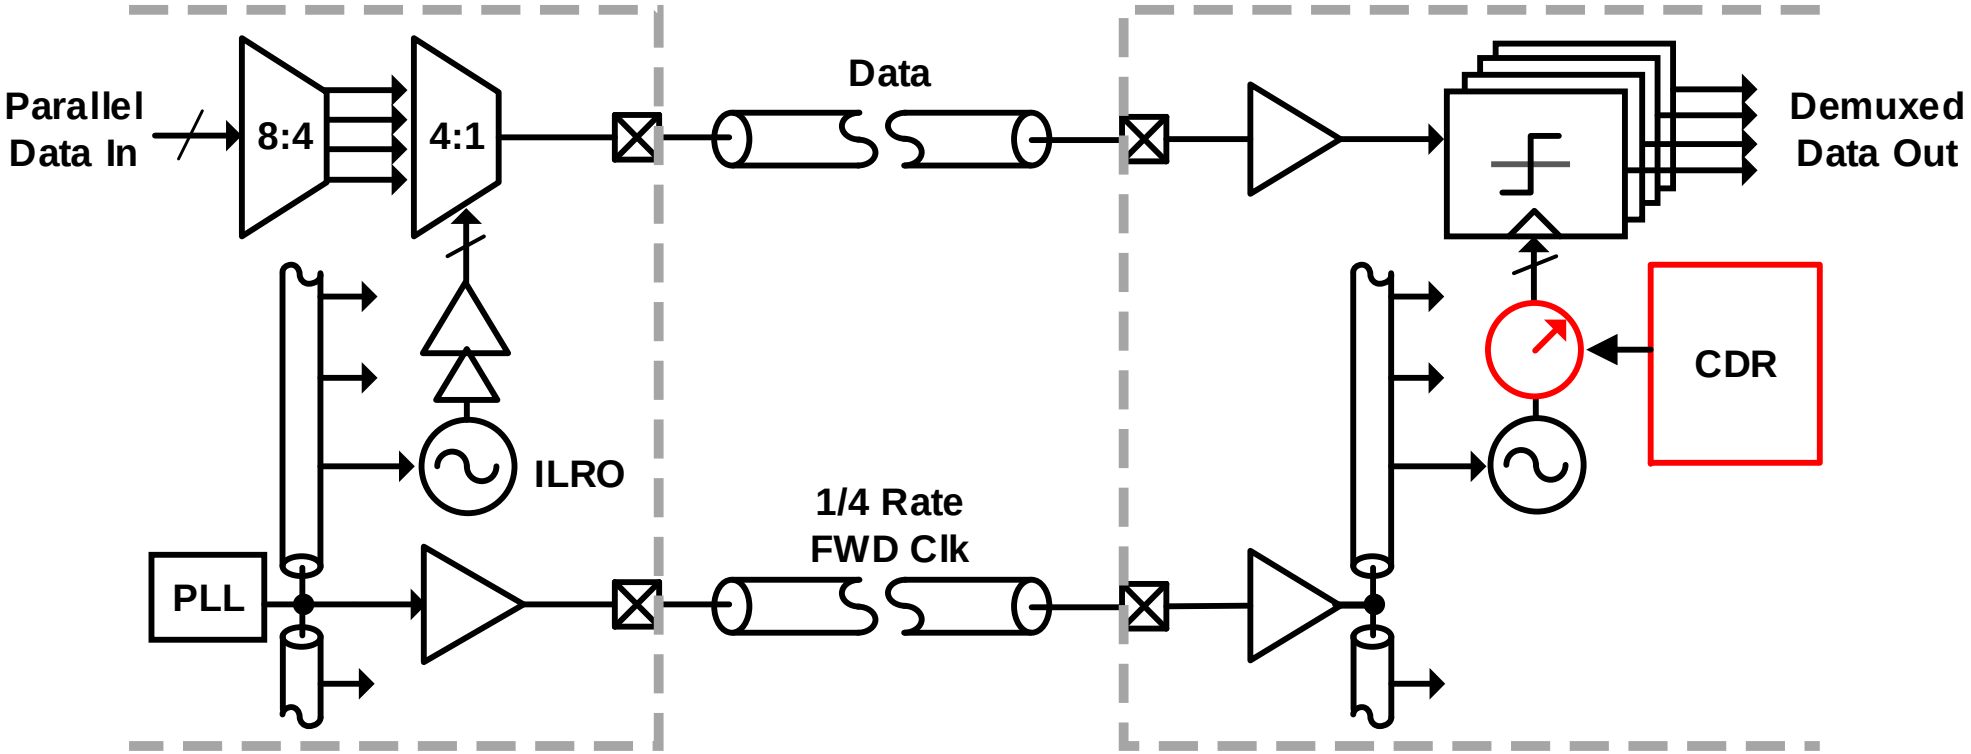

- Introducing a low-overhead CDR into a forwarded-clock system allows tracking of low-frequency phase drifts, while maintaining correlated jitter tracking

Multi-Phase Errors at Low VDD

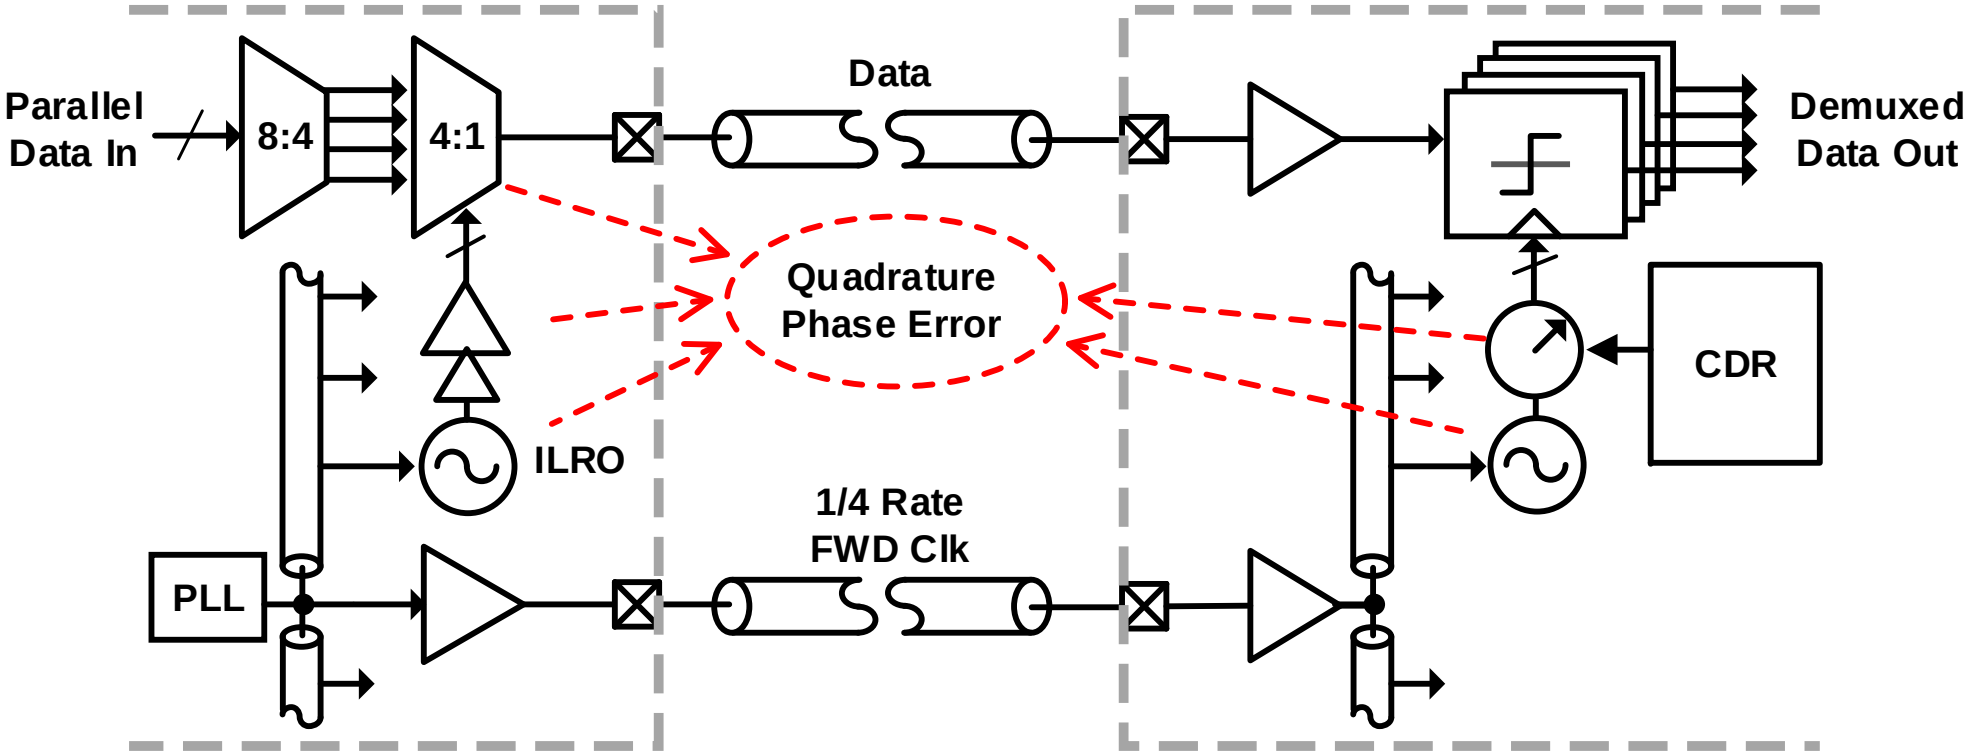

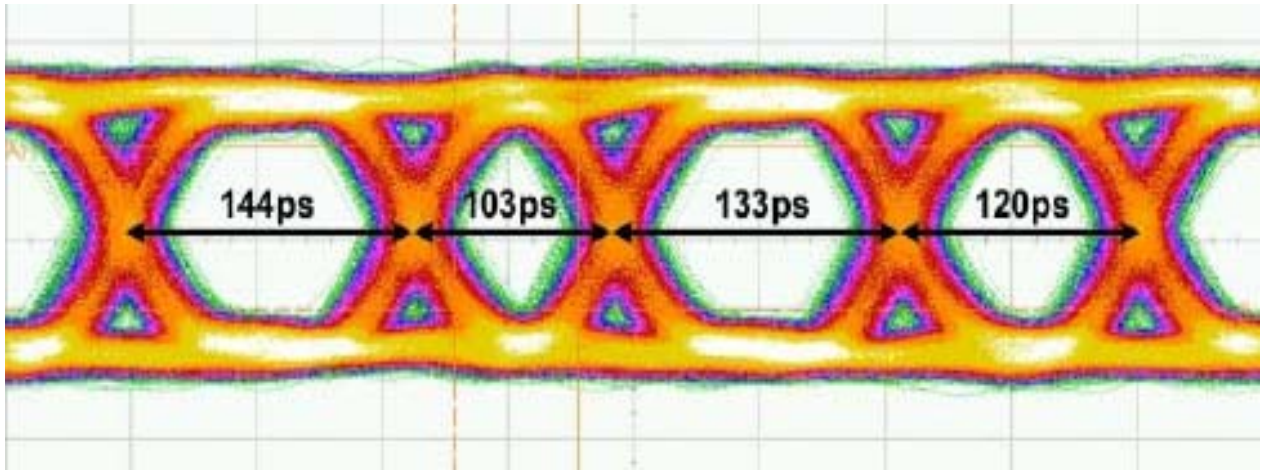

Edge-Rotating 5/4X Sub-Rate CDR

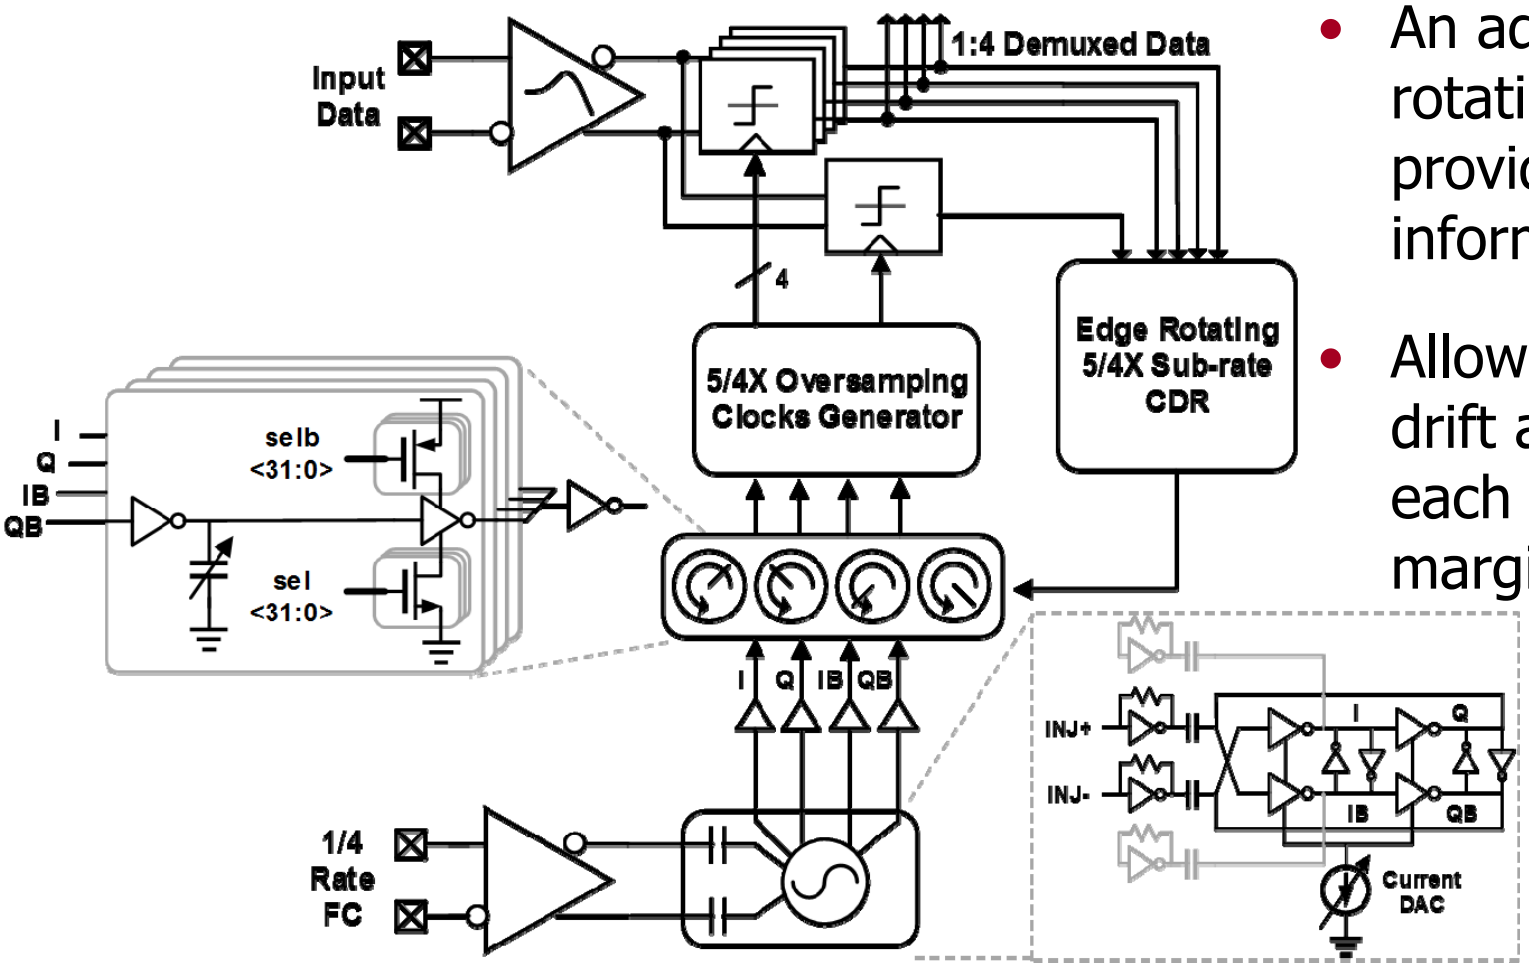

- An additional periodically rotating edge sampler provides the 4-eye phase information to CDR logic
- Allows tracking of phase drift and optimization of each sampler timing margin

H. Li, S. Chen, L. Yang, R. Bai, W. Hu, F. Zhong, S. Palermo, and P. Chiang, "A 0.8V, 560fJ/bit, 14Gb/s Injection-Locked Receiver with Input Duty-Cycle Distortion Tolerable Edge-Rotating 5/4X Sub-Rate CDR in 65nm CMOS," VLSI Symp., June 2014.

14Gb/s GP 65nm CMOS Prototype

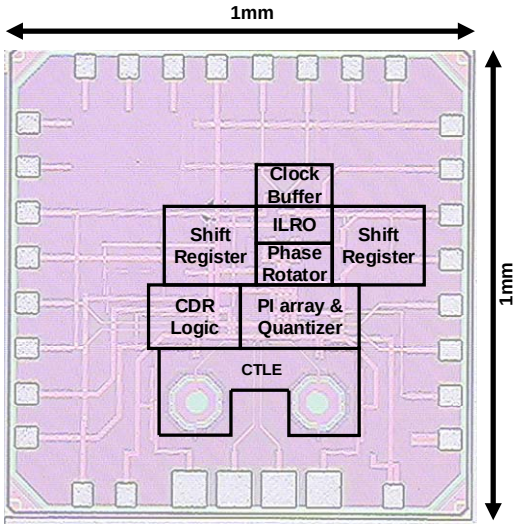

Tracking Non-Uniform Eyes

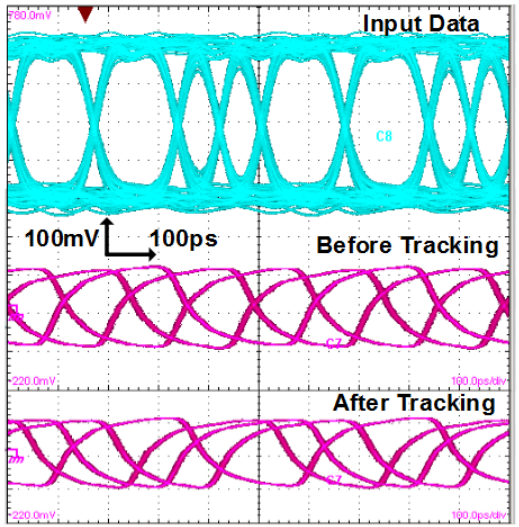

Correlated Jitter Tolerance

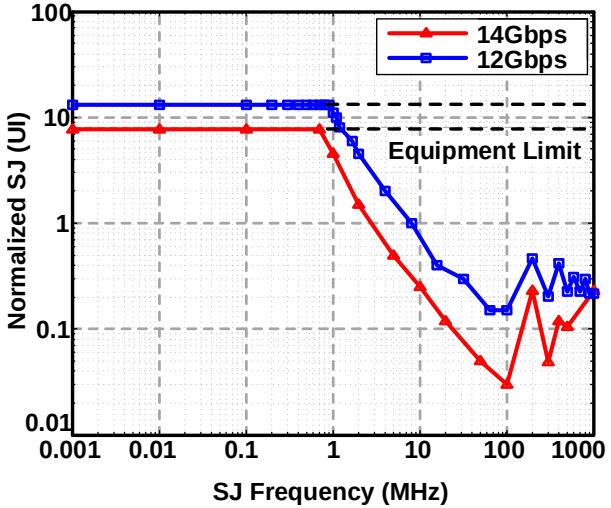

Uncorrelated Jitter Tolerance

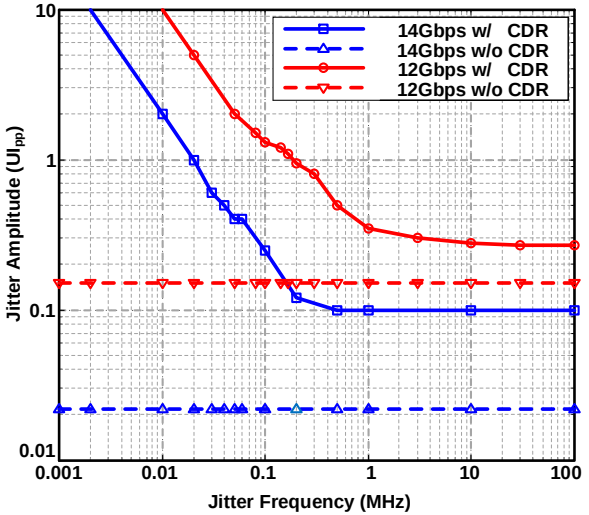

	[2]	[3]	[4]	This Work
V_{DD}, Process	1.0V/65nm	1.25V/90nm	1.08V/32nm	0.8V/65nm
Data Rate	6.4Gb/s	8Gb/s	16Gb/s	14Gb/s
Clock Rate	3.2GHz	2GHz	4GHz	3.5GHz
Clock Arch.	FC	Embedded	FC	FC
Multi-phase Gen.	PLL/PI	DLL	ILO/PI	ILO/PI
RX FOM *(Excludes PLL)	**3.9pJ/bit***	**1.59pJ/bit***	**1.02pJ/bit**	**0.56pJ/bit**

H. Li, S. Chen, L. Yang, R. Bai, W. Hu, F. Zhong, S. Palermo, and P. Chiang, "A 0.8V, 560fJ/bit, 14Gb/s Injection-Locked Receiver with Input Duty-Cycle Distortion Tolerable Edge-Rotating 5/4X Sub-Rate CDR in 65nm CMOS," VLSI Symp., June 2014.

Outline

- Motivation
- Power-Scalable I/O Techniques
- Low-Power Clocking
- Low-Power Equalizers
- Conclusion

Link with Equalization

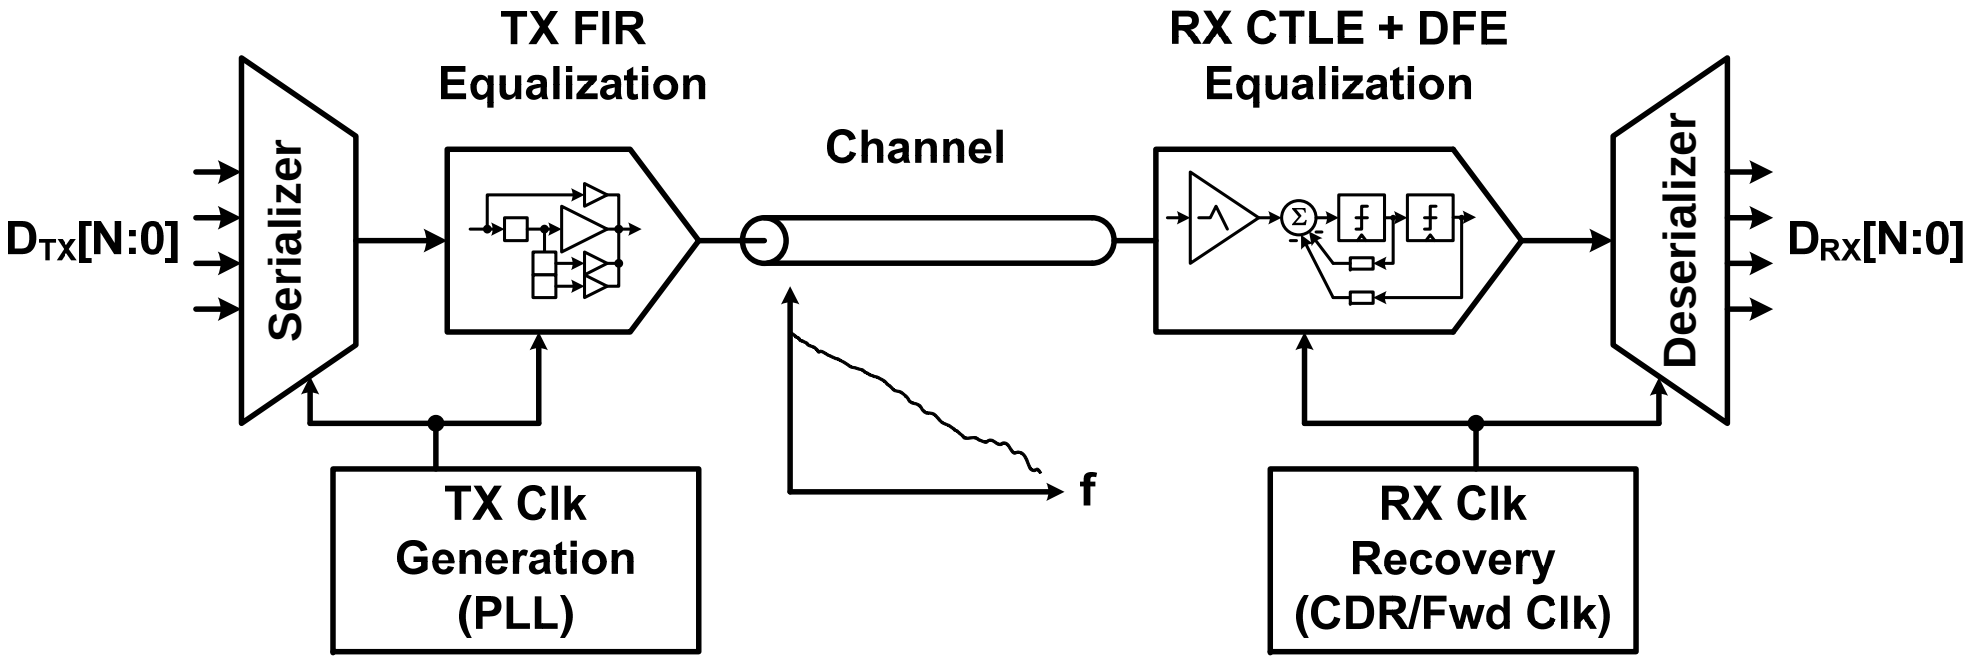

- Equalization goal is to flatten the frequency response out to the Nyquist frequency and remove time-domain ISI

TX-FIR Equalizer Comparisons

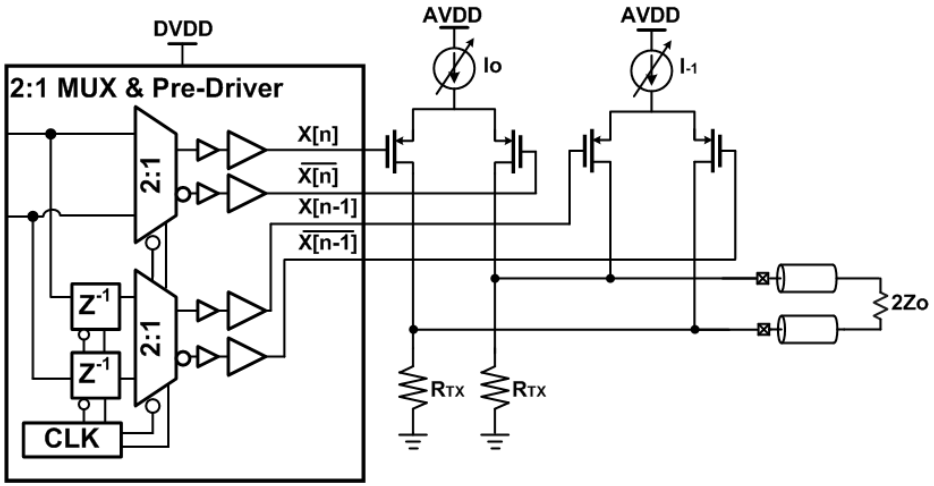

Current-Mode Driver (CM)

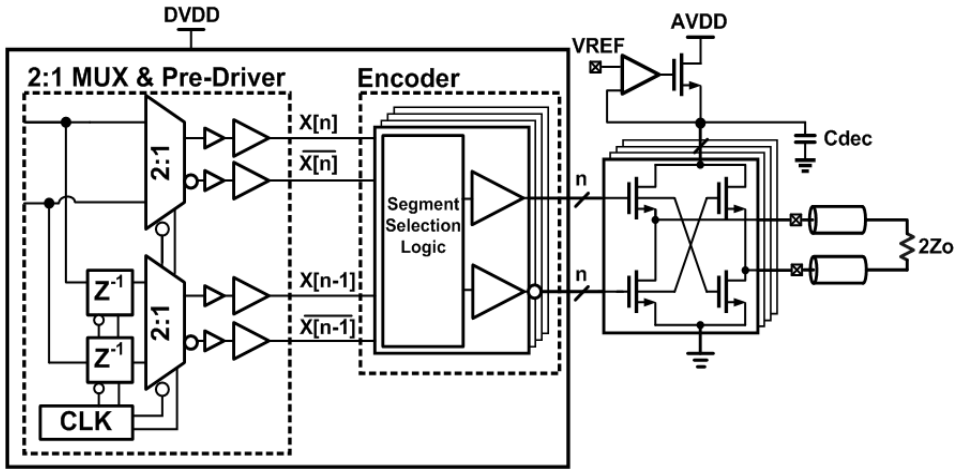

Voltage-Mode Driver (VM)

- FIR equalization can easily be implemented in a current-mode driver by summing tap currents on the termination resistors
- More difficult to implement in voltage-mode drivers due to the series impedance

VM Equalization w/ Shunt Voltage Divider (1)

[Wong JSSC 2004]

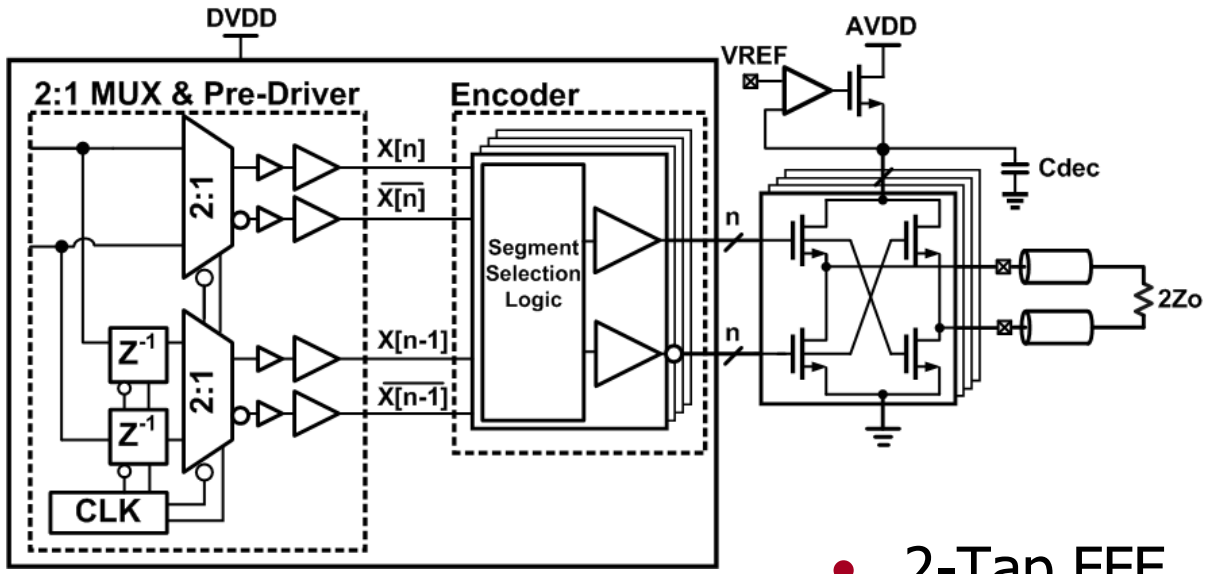

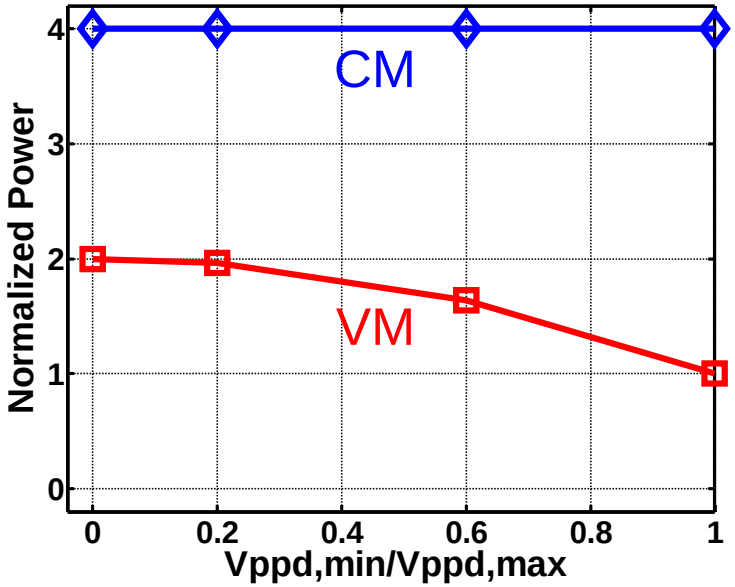

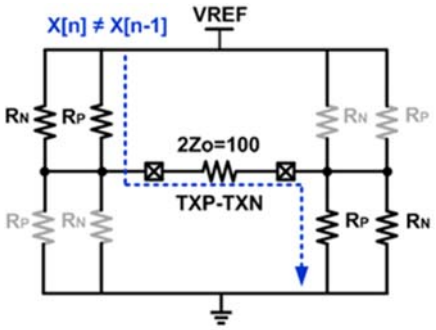

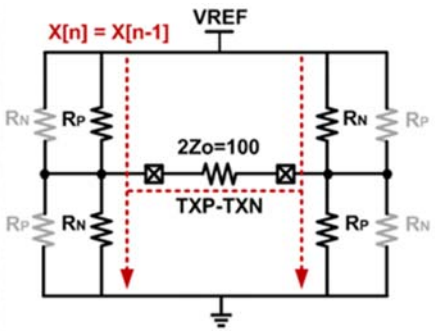

- 2-Tap FFE
  - Parallel combination $\left(R_P = \frac{Zo}{1-\alpha},\ R_N = \frac{Zo}{\alpha} \right)$
- Z-termination
  - $Z_o = R_P \parallel R_N$
- More current for de-emp. voltage swing

$$I_{sig} = \frac{V_{ref}}{R_T}\left[\frac{1}{2} - \left(\frac{V_{de-em}}{V_{ref}}\right)^2\right]$$

VM Equalization w/ Added Parallel Path (2)

[Dettloff ISSCC 2010]

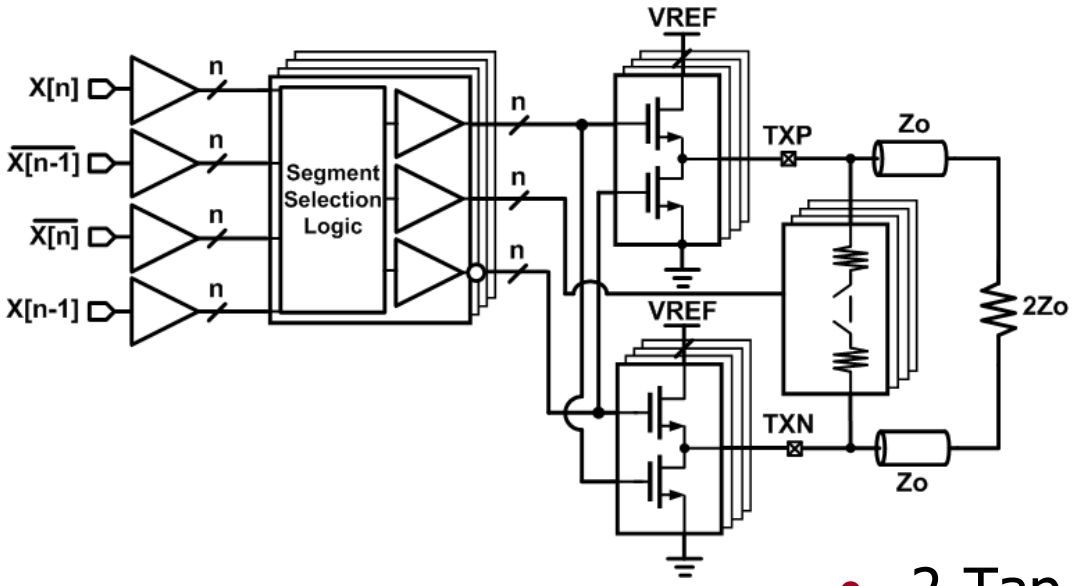

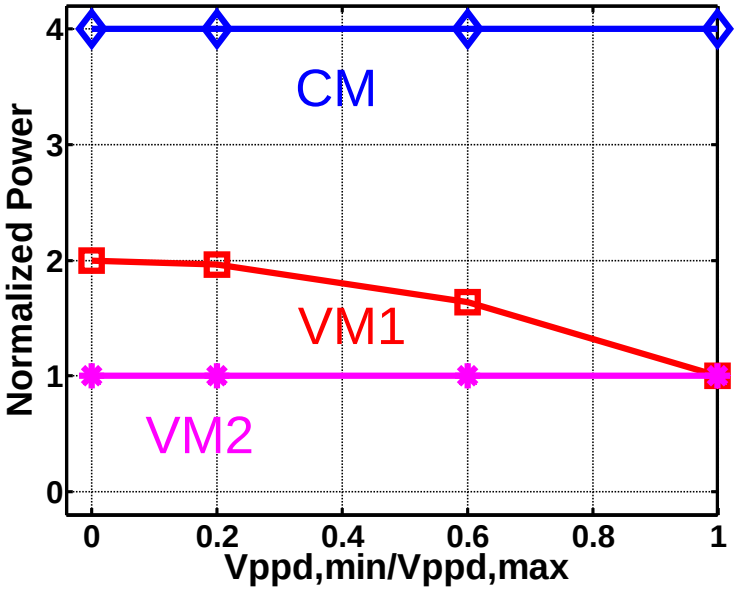

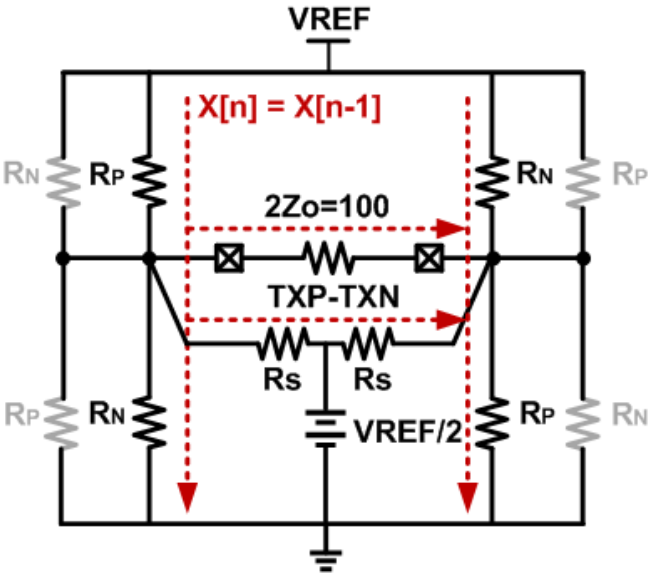

- 2-Tap FFE
  - Extra series-connected path
  - Constant current path
- Z-termination ($Z_O = R_P \parallel R_N \parallel R_S$)
- Constant signaling power for all V_{SW}
- Non-linear impedance mapping
  - Decoding/pre-driver complexity

VM Equalization w/ Impedance Modulation (3)

[Sredojevic JSSC 2011]

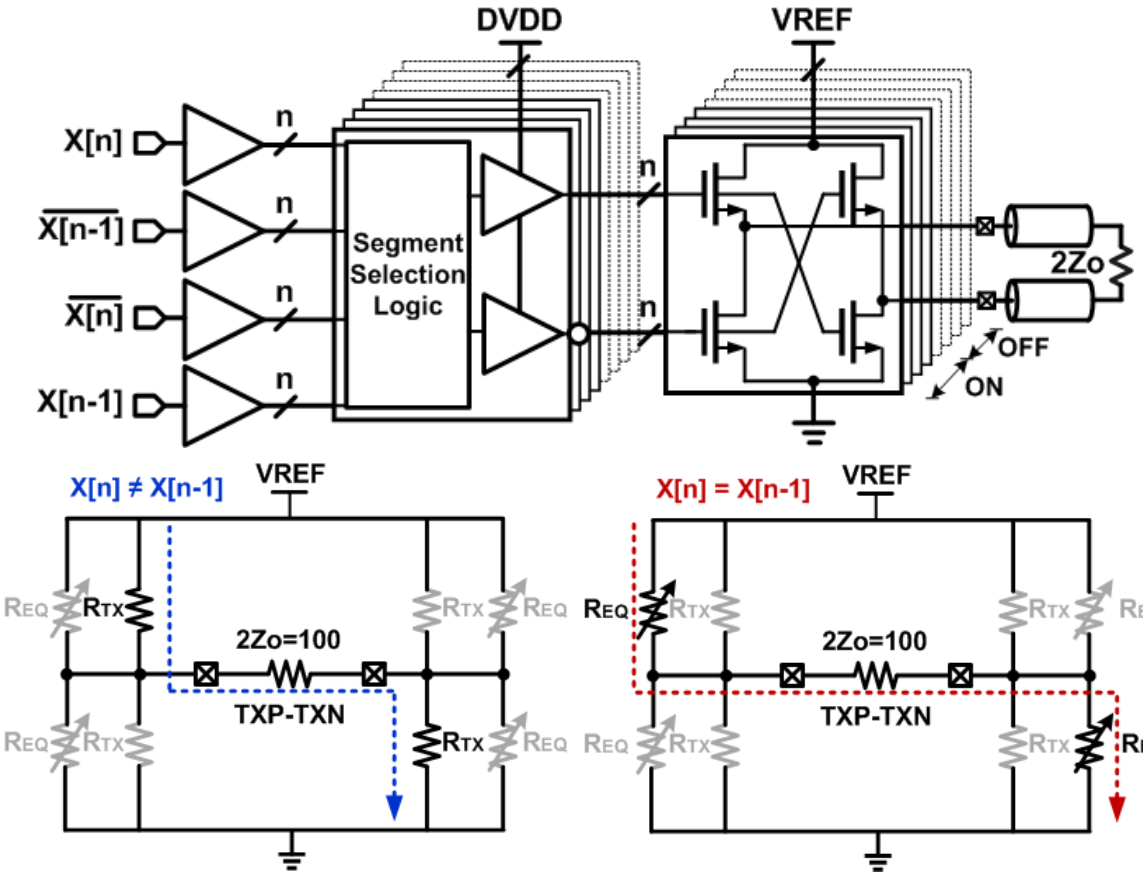

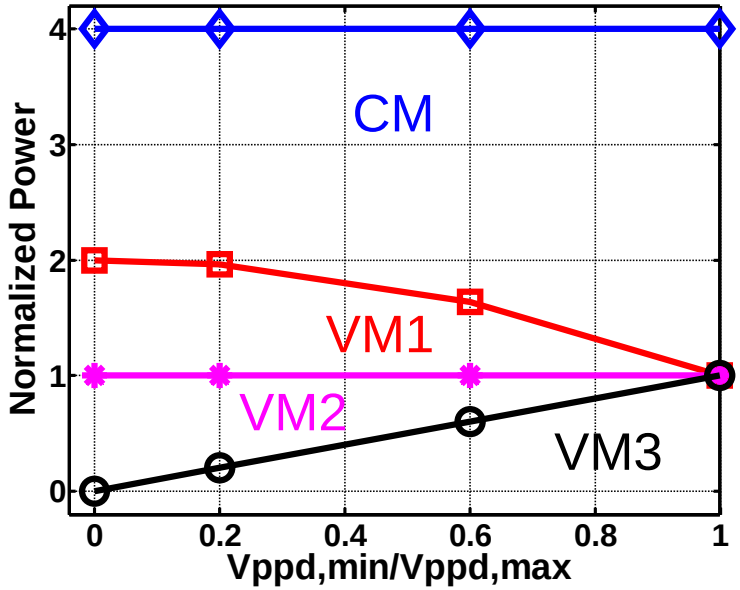

$$I_{sig} = \frac{V_{ppd,\max}}{4R_T}(1-2\alpha)$$

$$R_{TX} = Z_O,\ R_{EQ} = \frac{1+2\alpha}{1-2\alpha}R_{TX}$$

- 2-Tap FFE: Z-modulation (For de-emphasis, higher TX impedance)
- Signaling power $\propto$ Vppd,min / Vdd,max
- Sacrificing the output termination
- High digital power

VM Equalization w/ Analog Impedance Modulation

Digitally-Controlled Segmented Output

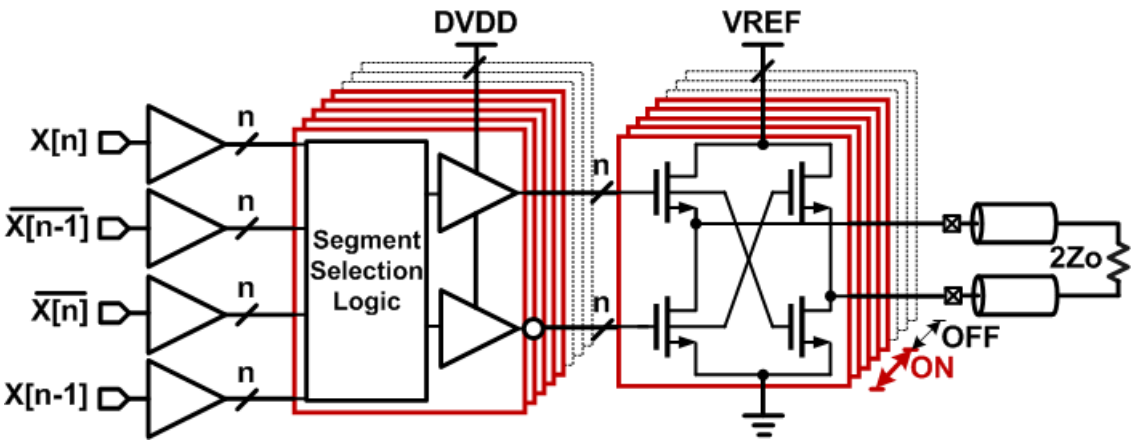

Analog-Controlled Non-Segmented Output

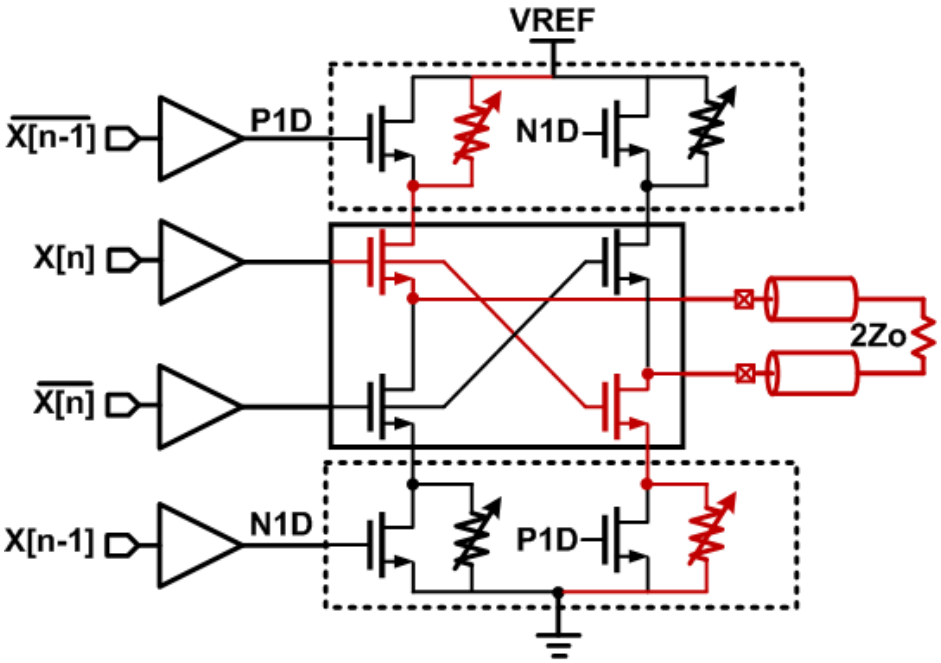

- Segmented pre-driver and output driver significantly increases dynamic power consumption with increased equalization resolution
- Analog tap control obviates output stage segmentation

VM Equalization w/ Analog Impedance Modulation

- Maximum transmitter output swing during a transition bit

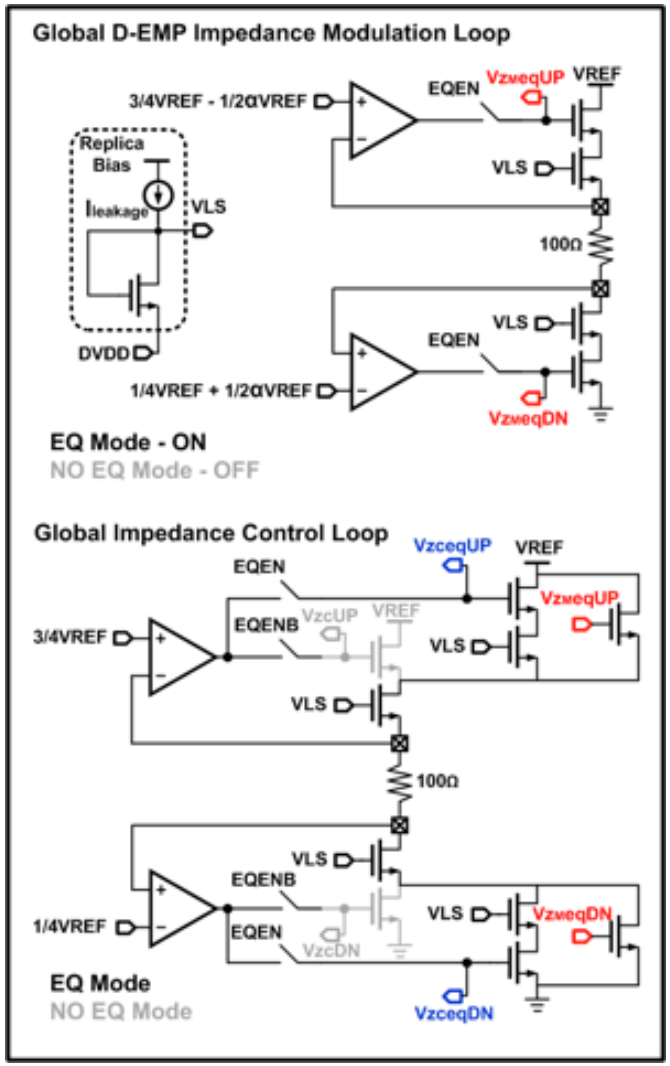

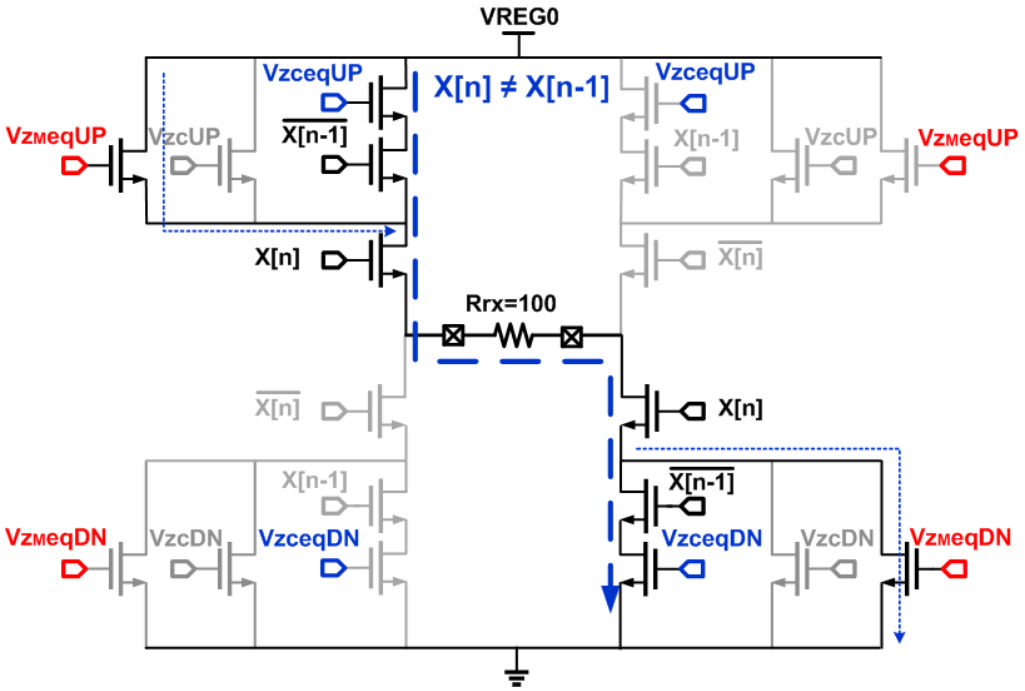

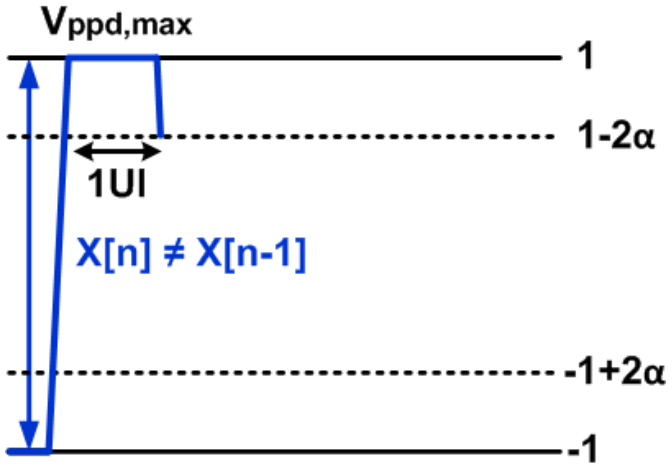

VM Equalization w/ Analog Impedance Modulation

- De-emphasis transmitter output swing for run-length > 1

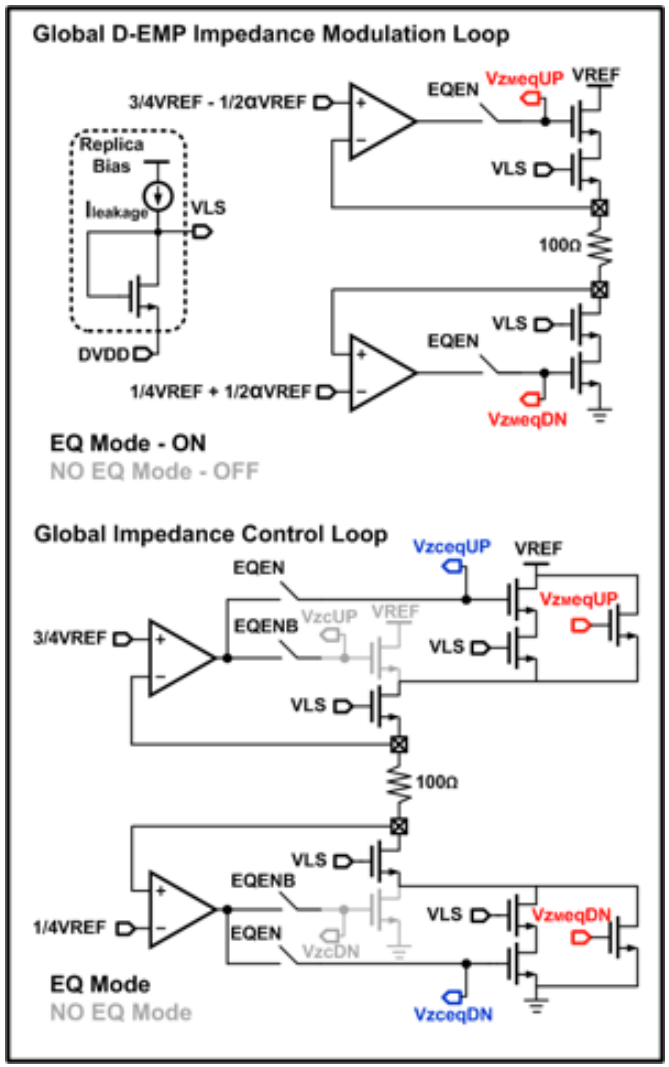

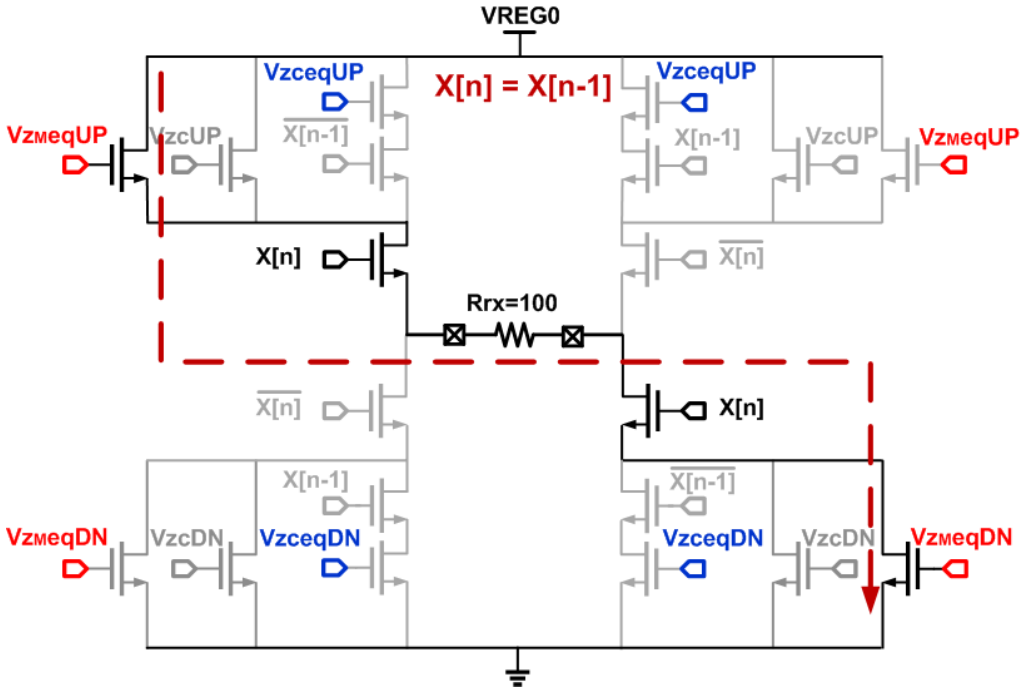

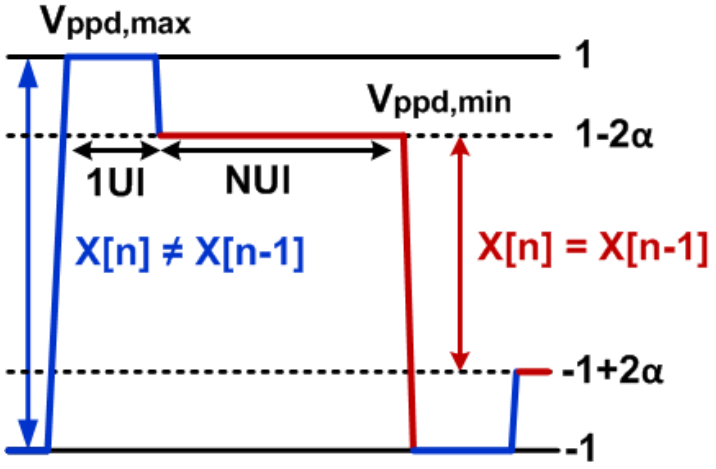

16Gb/s Operation

- 5.8 inch FR4 + 0.6m SMA cable -15.5dB loss at 8GHz

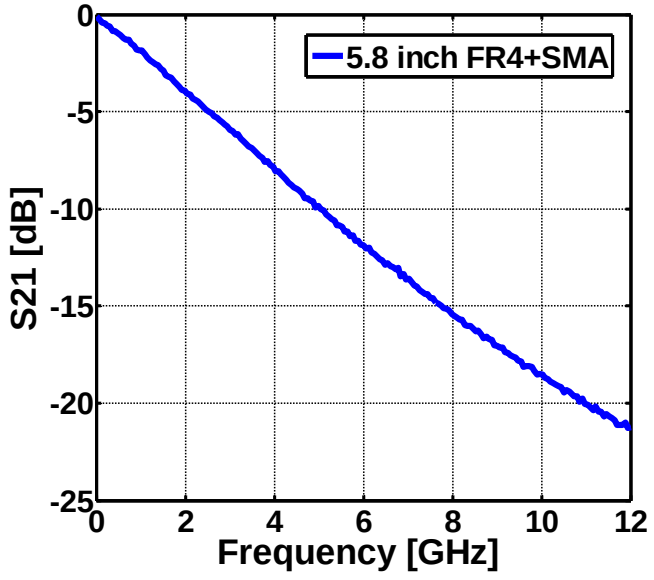

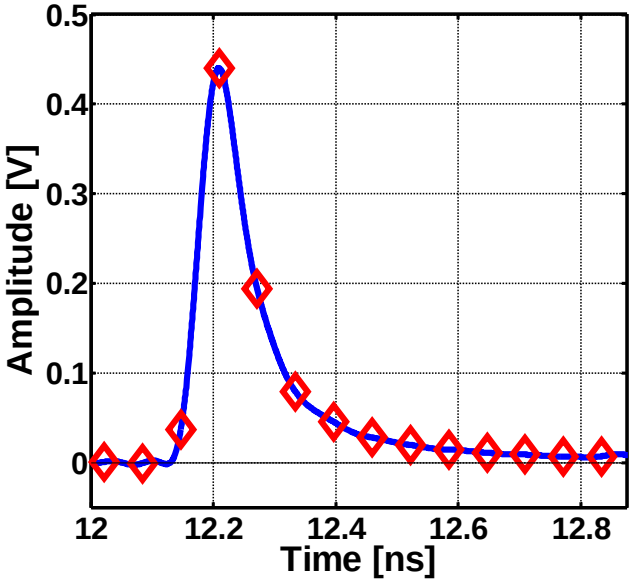

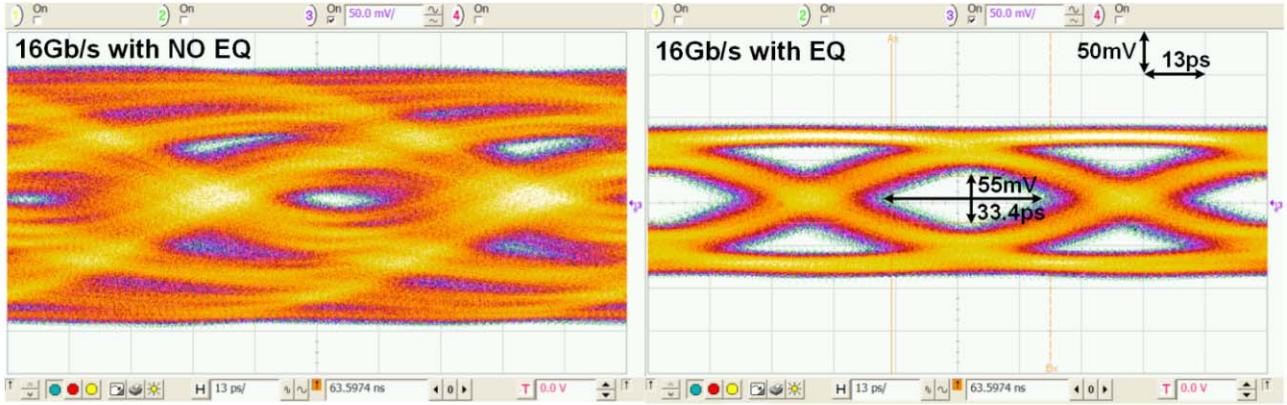

Y.-H. Song, H.-W. Yang, H. Li, P. Chiang, and S. Palermo, "An 8–16 Gb/s, 0.65–1.05 pJ/b, Voltage-Mode Transmitter With Analog Impedance Modulation Equalization and Sub-3 ns Power-State Transitioning," IEEE JSSC, vol. 49, no. 11, pp. 2631-2643, Nov. 2014.

Low-Voltage DFE w/ Charge-Based Latches

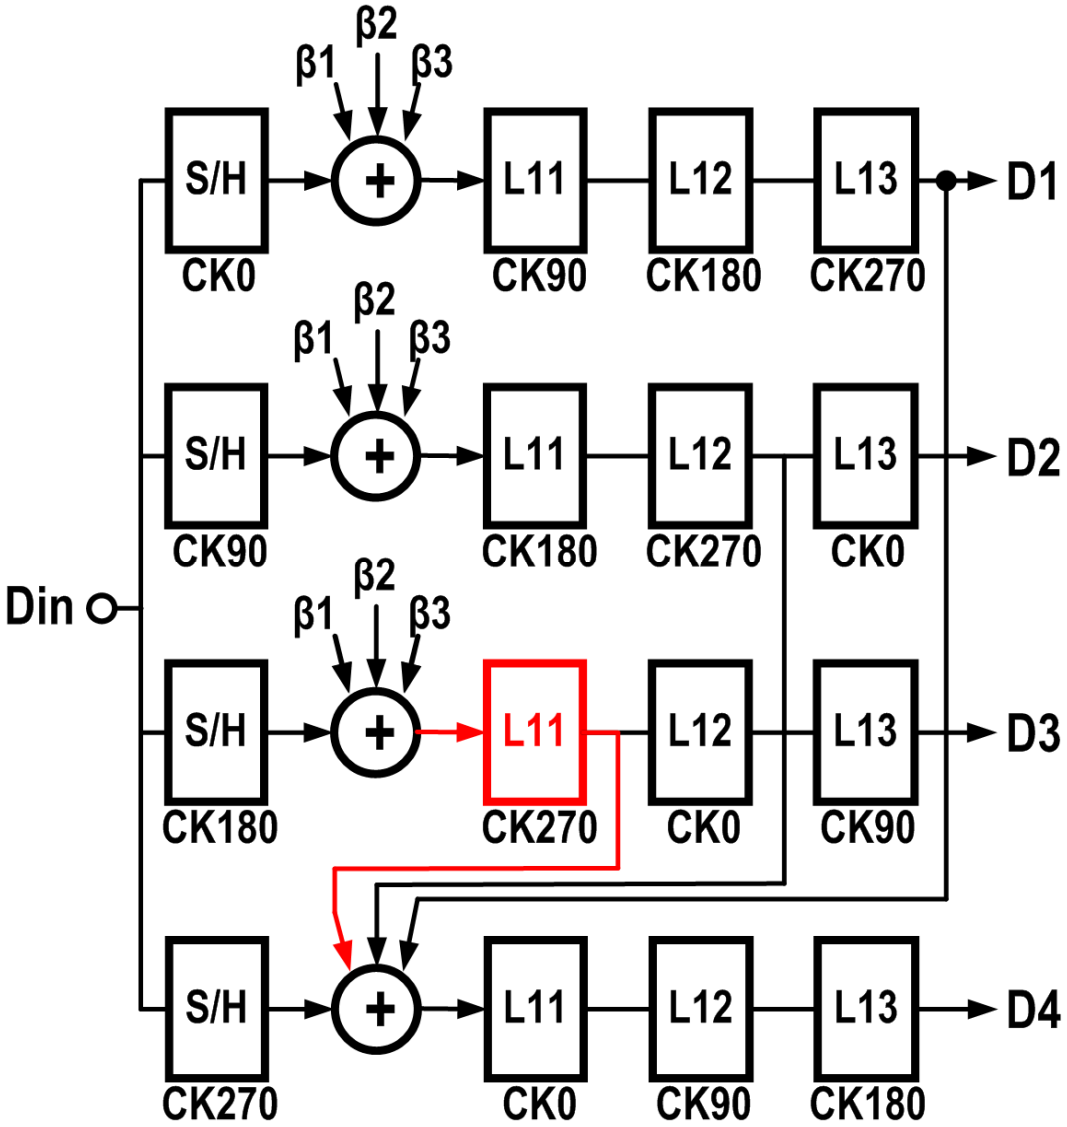

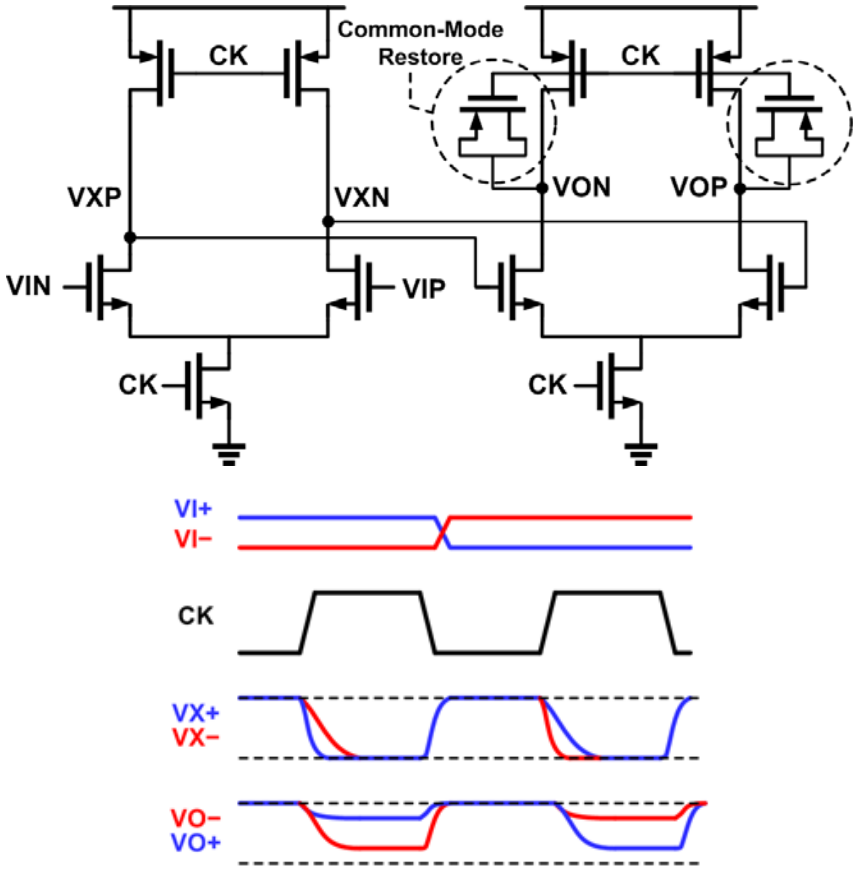

- First stage has small aperture time
- Second stage has small delay to quantized output

R. Bai, S. Palermo, and P. Chiang, "A 0.25pJ/b 0.7V 16Gb/s 3-Tap Decision-Feedback Equalizer in 65nm CMOS," ISSCC, Feb. 2014.

16Gb/s Operation

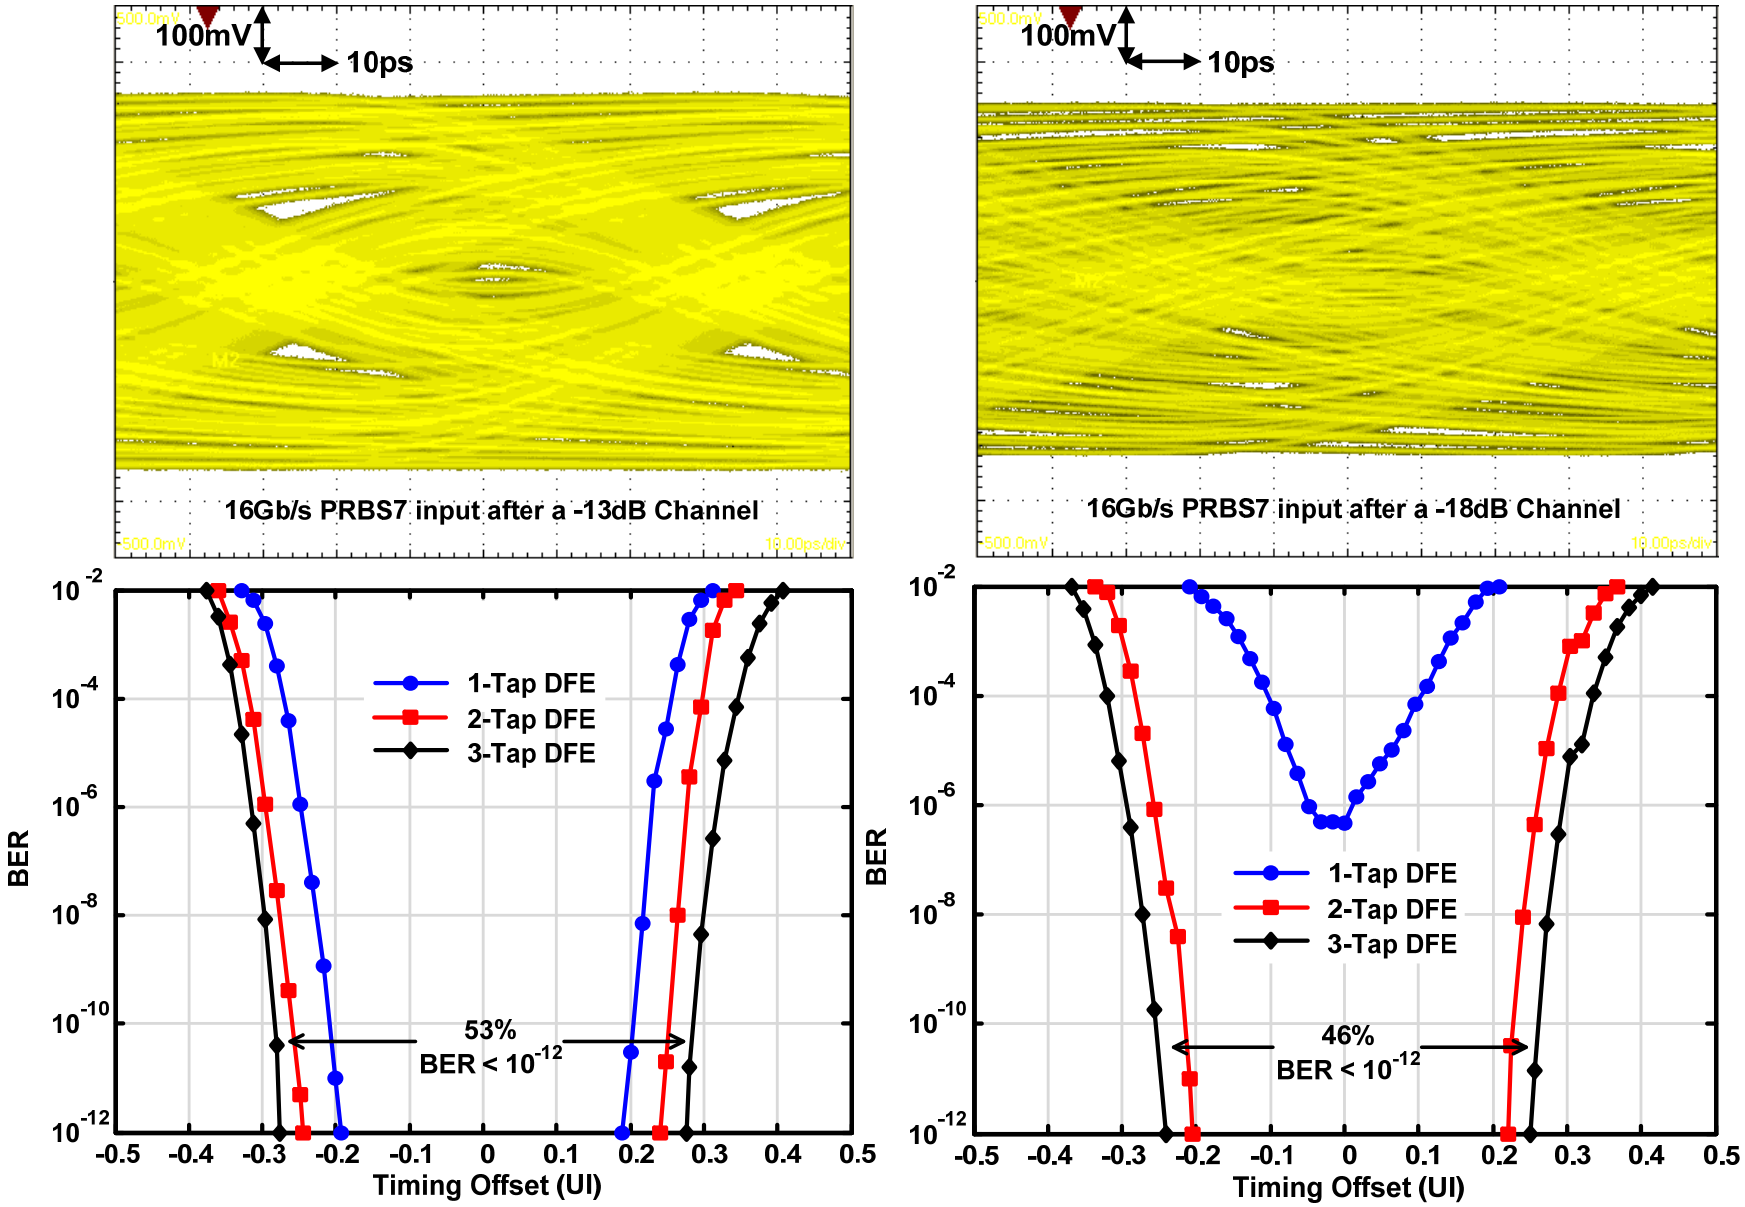

R. Bai, S. Palermo, and P. Chiang, "A 0.25pJ/b 0.7V 16Gb/s 3-Tap Decision-Feedback Equalizer in 65nm CMOS," ISSCC, Feb. 2014.

DFE with Feedback FIR Filter Issues

- DFE critical path timing → speed/power trade-off
- High-loss channels require large number of DFE taps
  - Increases area and power
  - Increases loading → limits speed

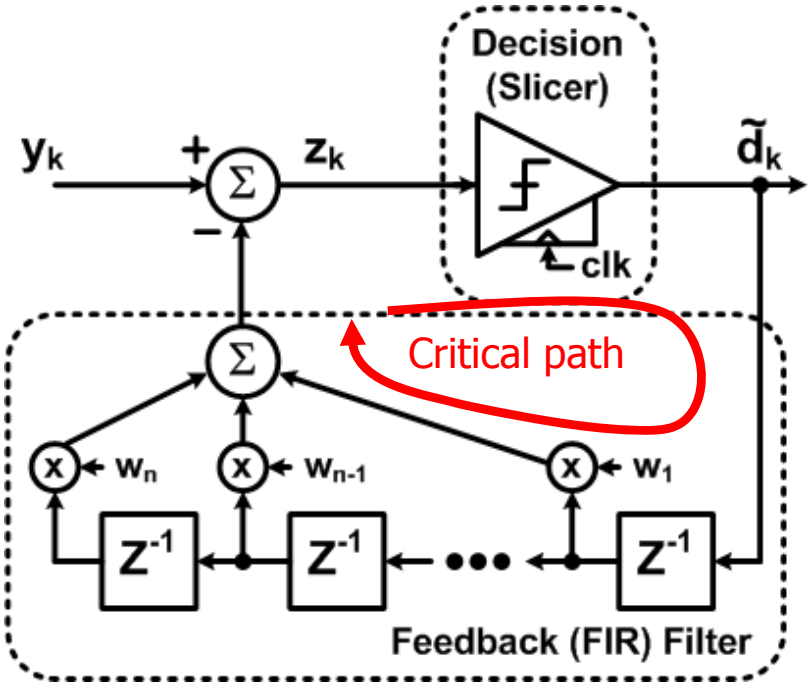

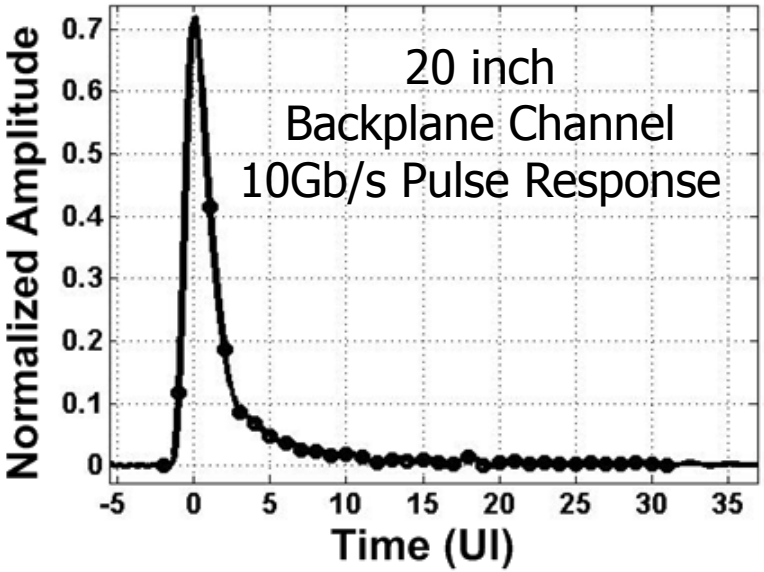

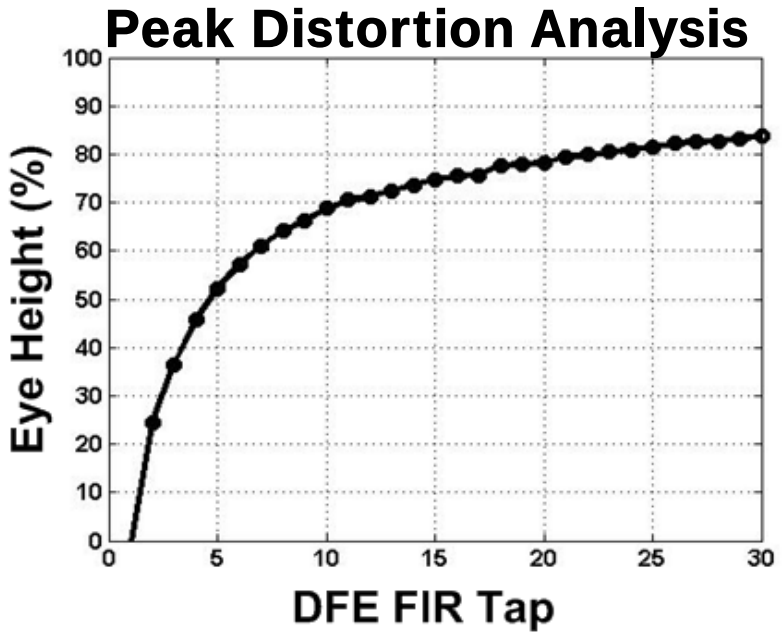

DFE with Feedback IIR Filter

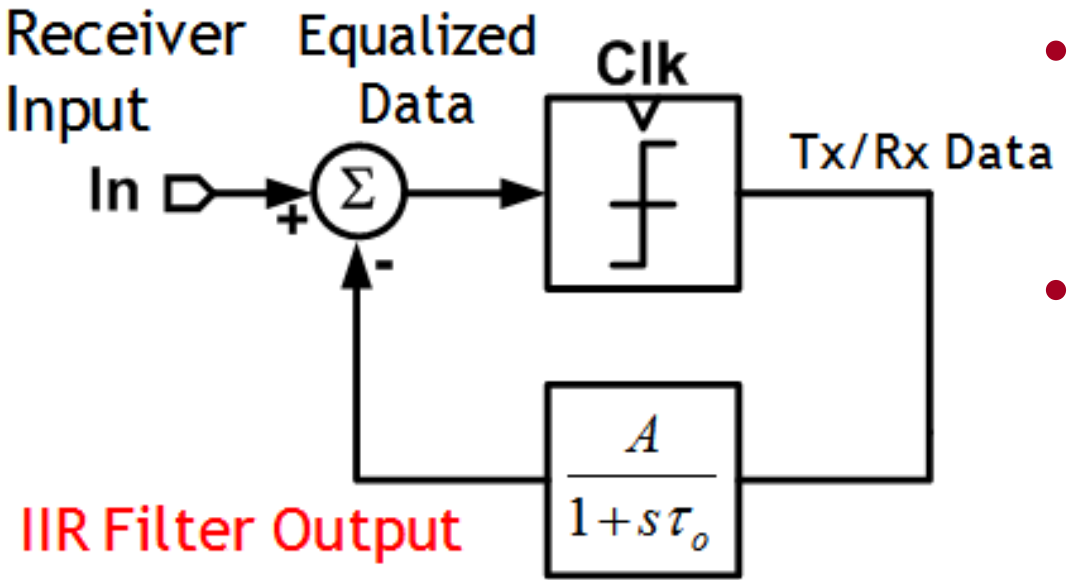

- IIR feedback filter provides efficient long-tail ISI cancellation
- Typical backplane channel well approximated with 2 IIR taps

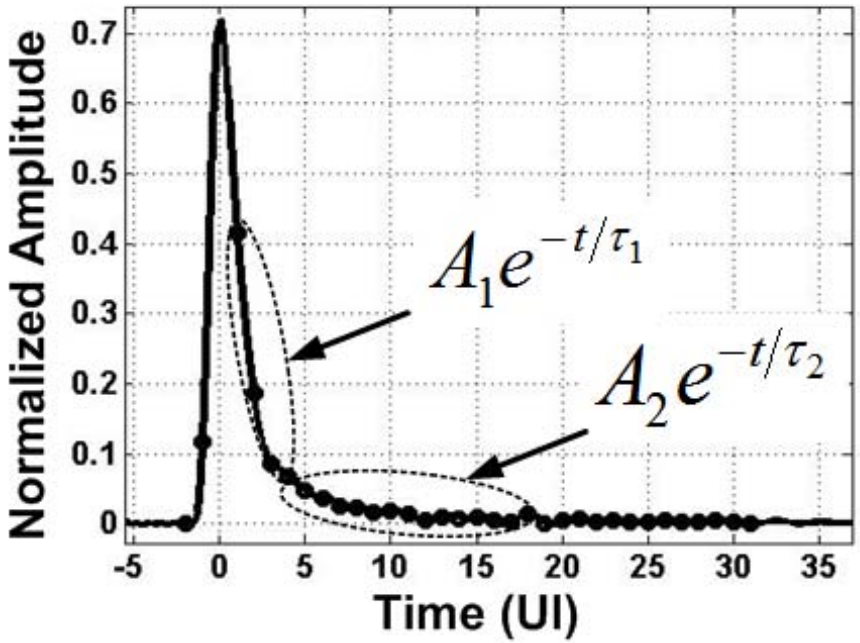

Peak Distortion Analysis

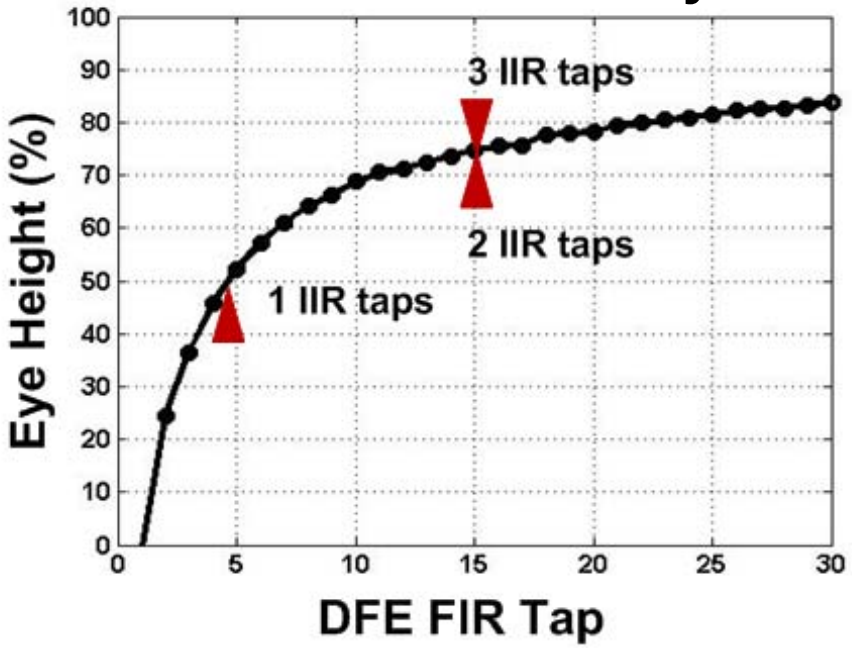

10Gb/s 2-IIR-Tap DFE w/ 35dB Loss Compensation

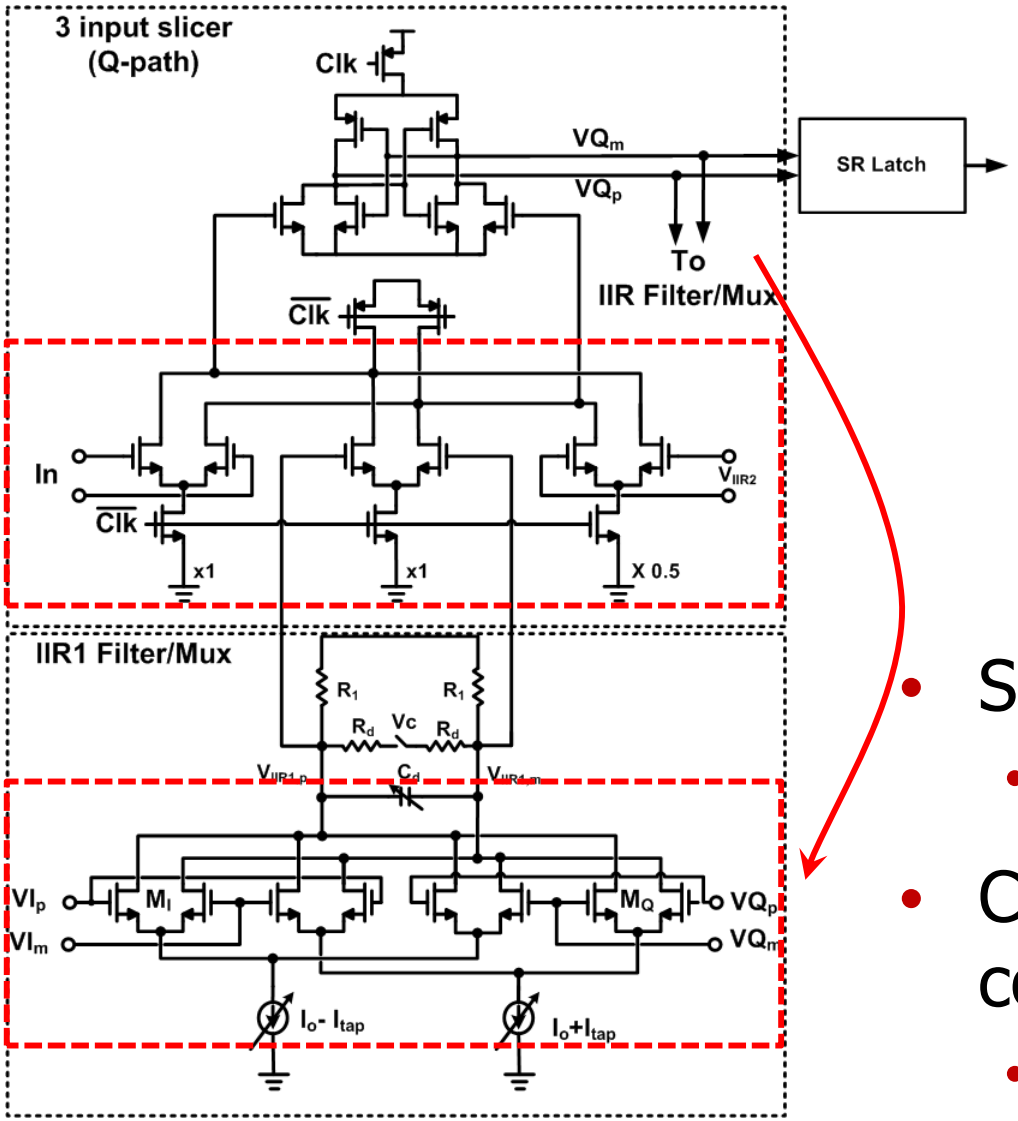

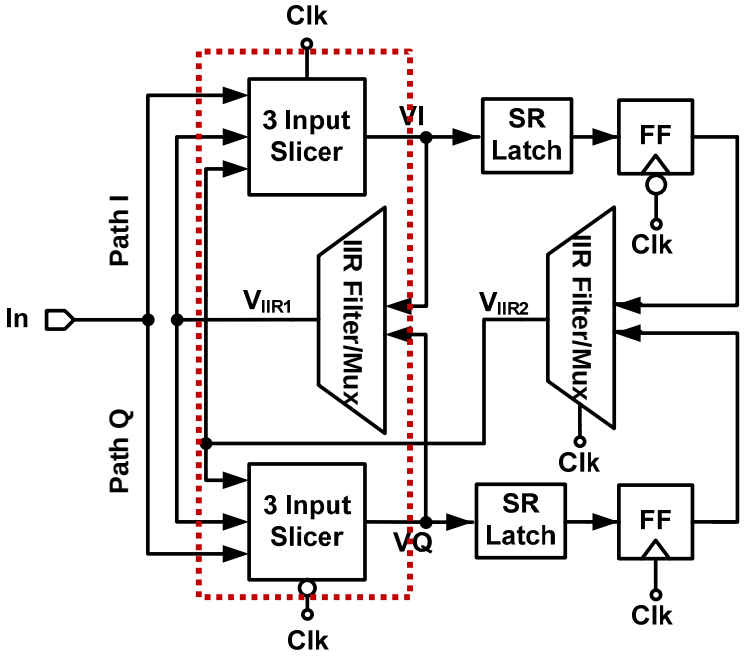

- Summation/slicing merged
  - Three-input double-tail comparator
- Comparator output directly connected to the IIR1 Mux
  - Lowers critical path delay

O. El-Hadidy and S. Palermo, "A 10 Gb/s 2-IIR-Tap DFE Receiver with 35 dB Loss Compensation in 65-nm CMOS," IEEE Symposium on VLSI Circuits, June 2013.

10Gb/s 2-IIR-Tap DFE w/ 35dB Loss Compensation

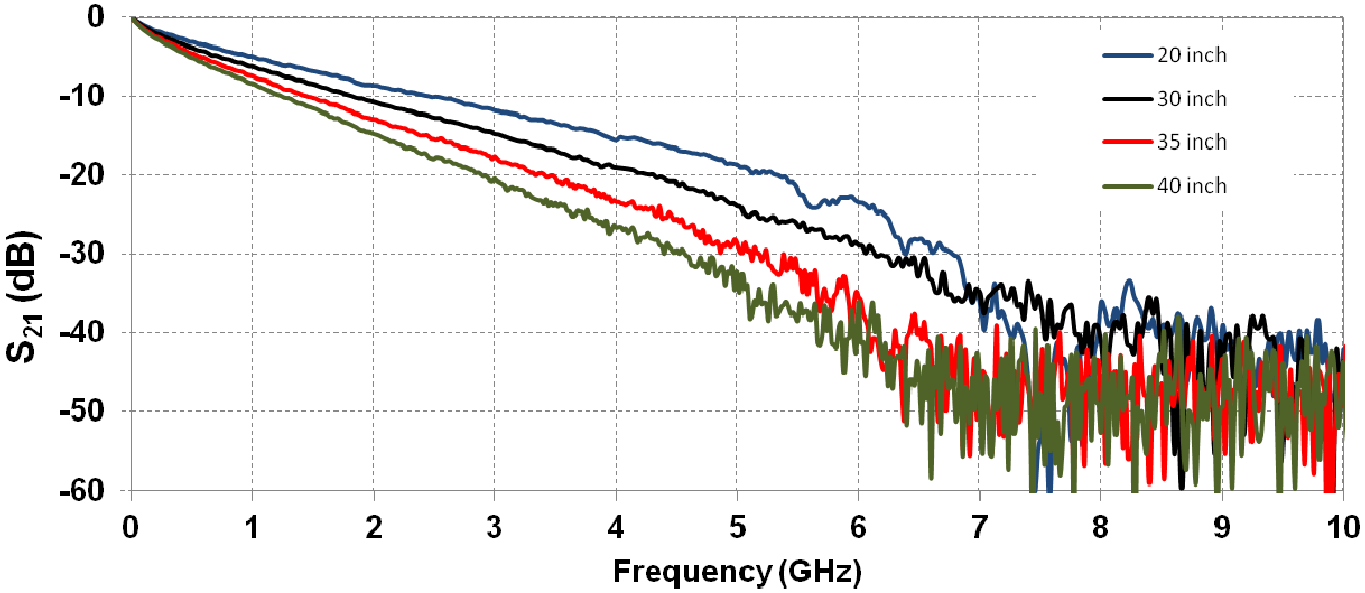

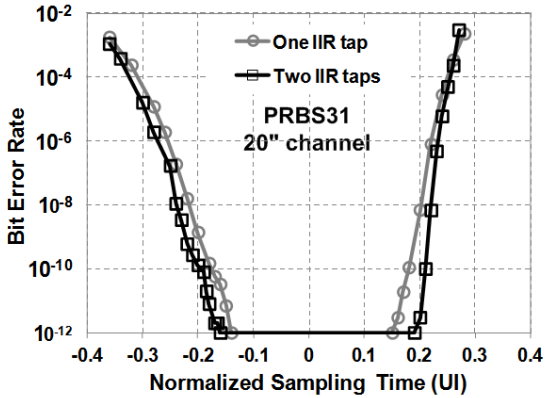

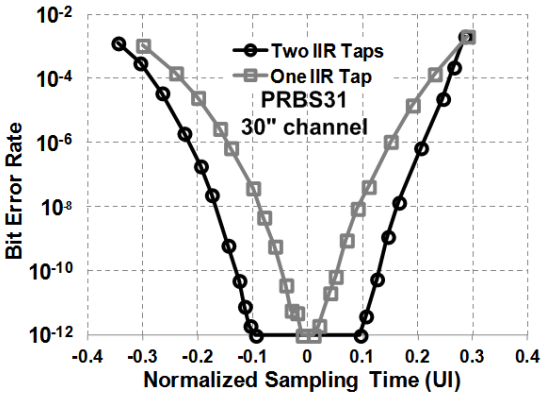

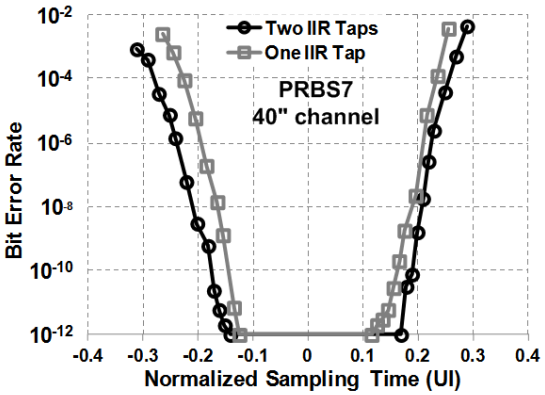

Specification	This Work	[2]	[3]	[4]
Data Rate	10 Gb/s	10 Gb/s	6 Gb/s	3.7 Gb/s
Architecture	2-IIR tap	1-IIR tap	1-IIR tap	3-IIR tap
Channel Loss @ Nyquist	35* dB	27* dB	32.7 dB	8 dB
Eye width BER PRBS Code	31 % 10^{-12} PRBS7	28 % 10^{-9} PRBS7	NA 10^{-12} PRBS7	6 % 10^{-12} PRBS7
Supply (V)	1	1	NA	1.2
Power (mW)	9.9	7	4	17.3
Area (mm^2)	0.0304	0.0175	0.089**	0.0535
Technology	65-nm	65-nm	90-nm	130-nm

* Including ~2 dB of setup loss.
** Include the area of a whole receiver.

O. El-Hadidy and S. Palermo, "A 10 Gb/s 2-IIR-Tap DFE Receiver with 35 dB Loss Compensation ir
on VLSI Circuits, June 2013.

PAM4 Signaling

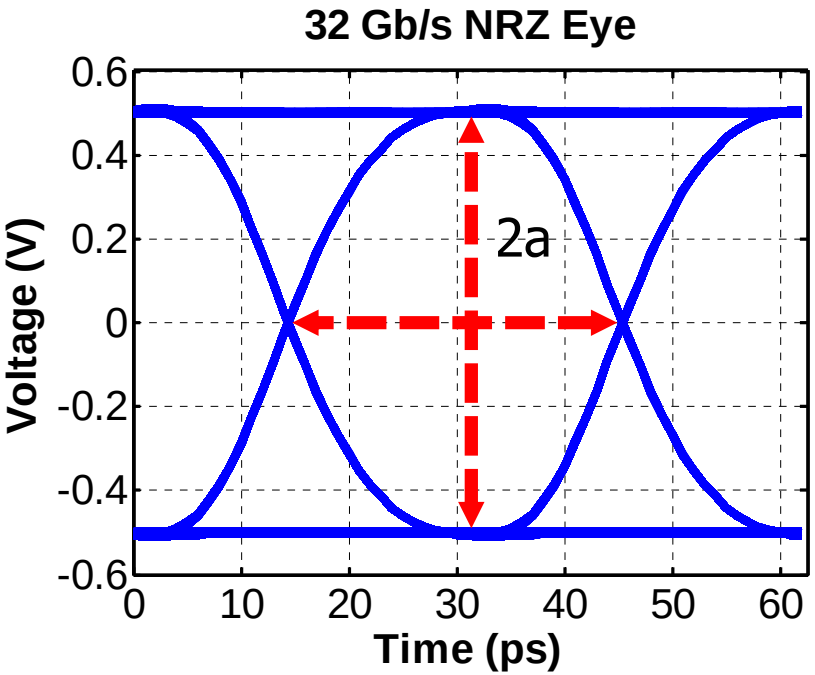

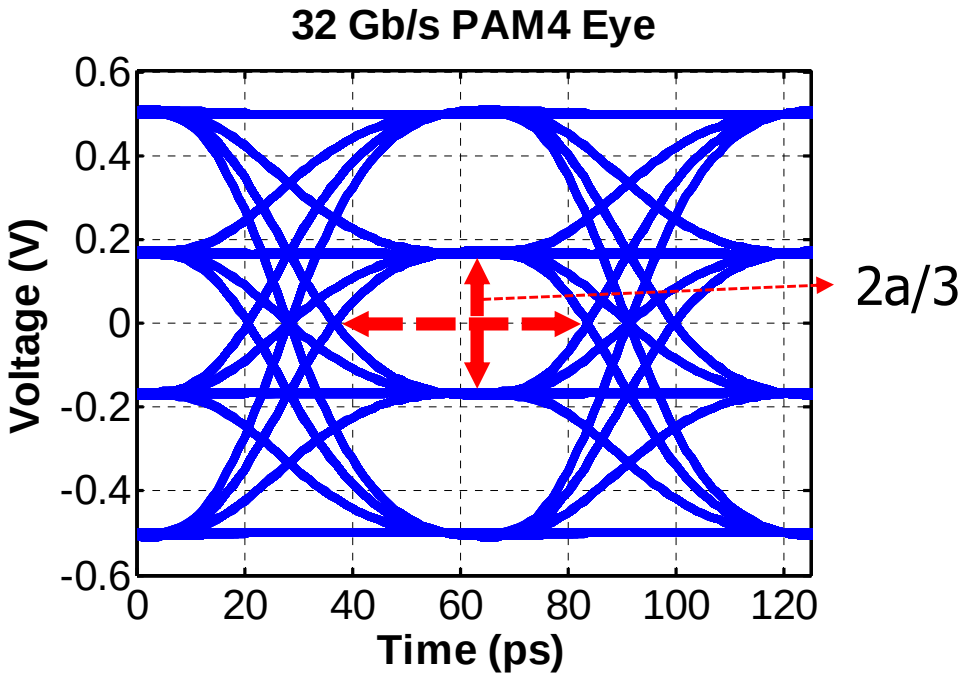

- PAM-4 modulation offers improved spectral efficiency over NRZ
- Main Characteristics:
  - ✓ Lower symbol rate
  - × Lower voltage margin → **Higher sensitivity is required**

A 32 Gb/s 0.55 mW/Gbps PAM4 1-FIR 2-IIR Tap DFE Receiver in 65-nm CMOS

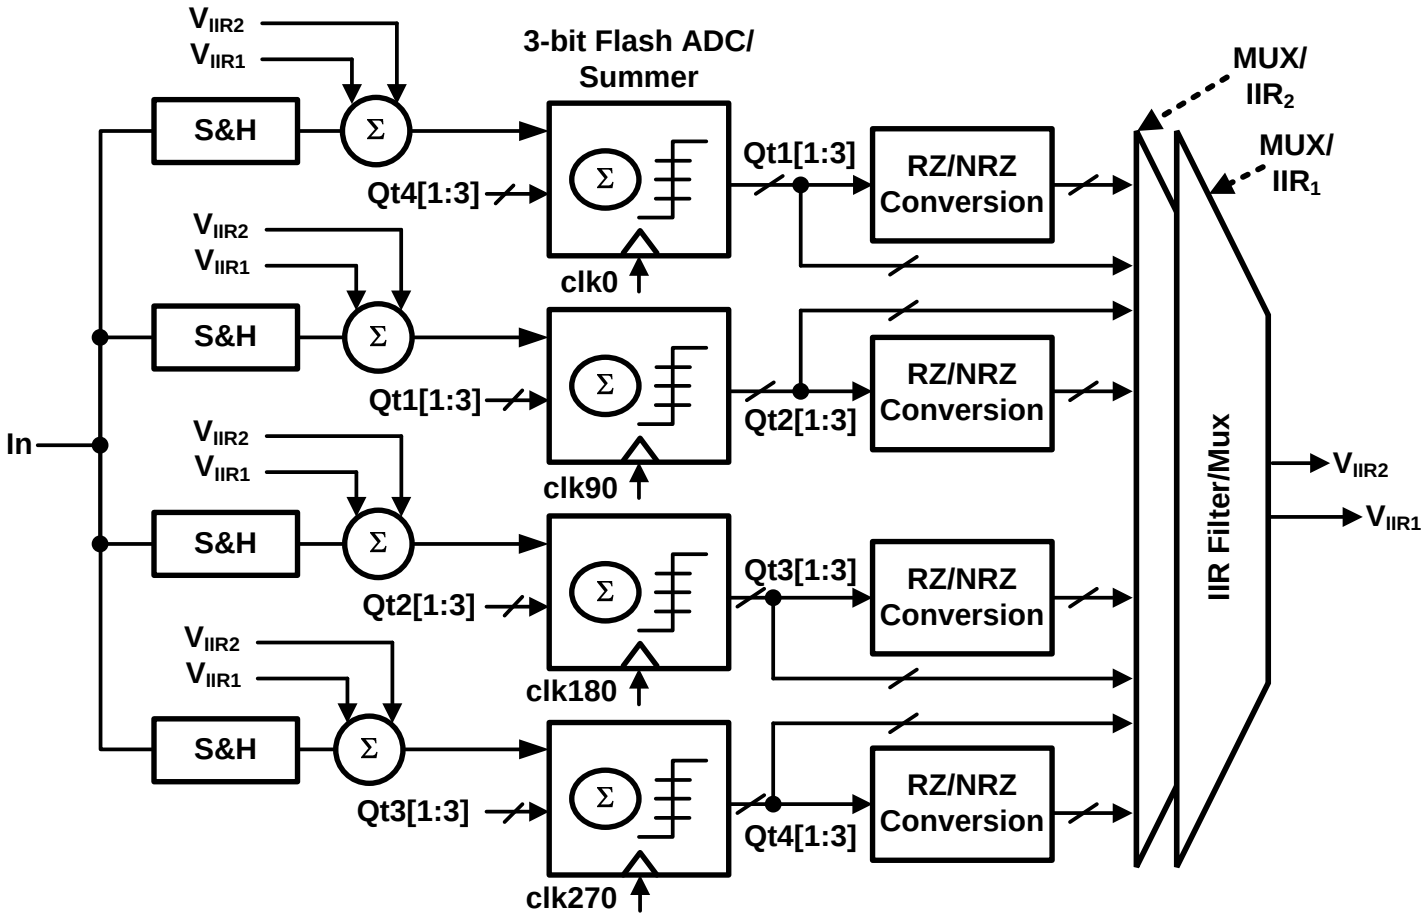

- PAM4 DFE employs 1-FIR tap for 1st post-cursor multi-level ISI cancellation and 2-IIR taps for long tail ISI cancellation
- Multi-level ISI cancellation is achieved with thermometer feedback to tap DACs

O. El-Hadidy, A. Roshan-Zamir, H.-W. Yang, and S. Palermo, "A 32 Gb/s 0.55 mW/Gbps PAM4 1-FIR 2-IIR Tap DFE Receiver in 65-nm CMOS," IEEE Symposium on VLSI Circuits, June 2015.

Dynamic Regenerative Comparator

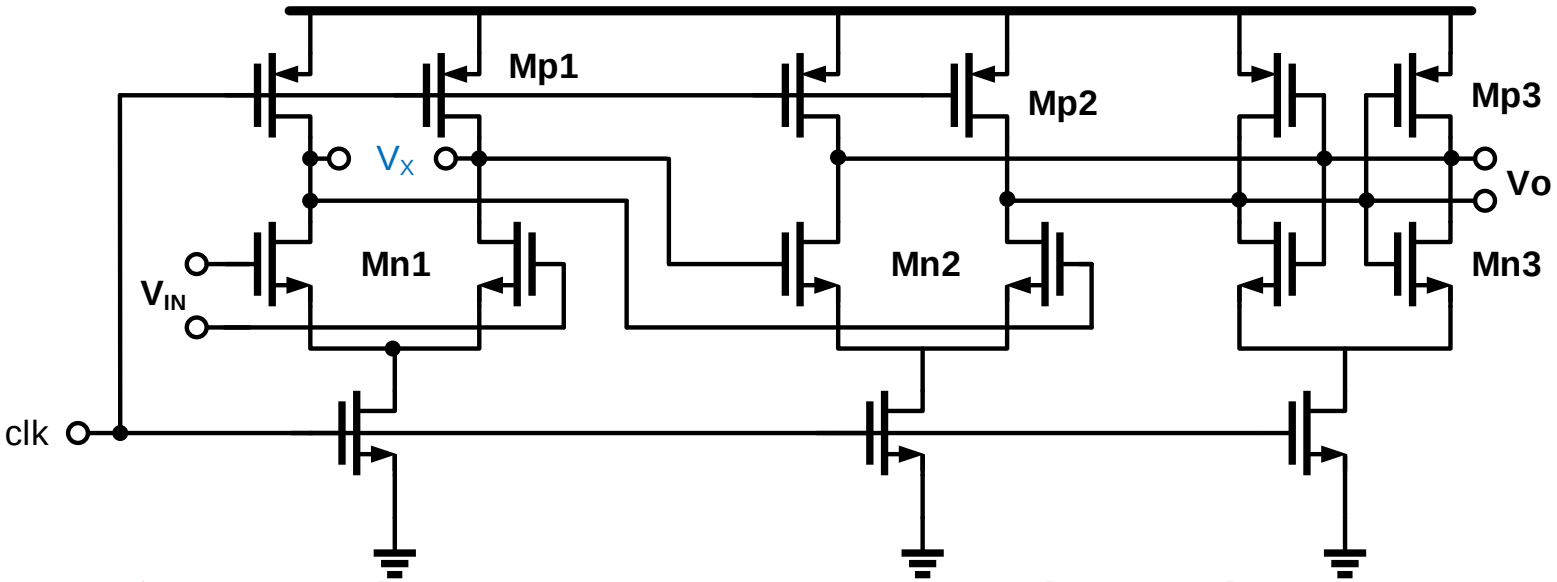

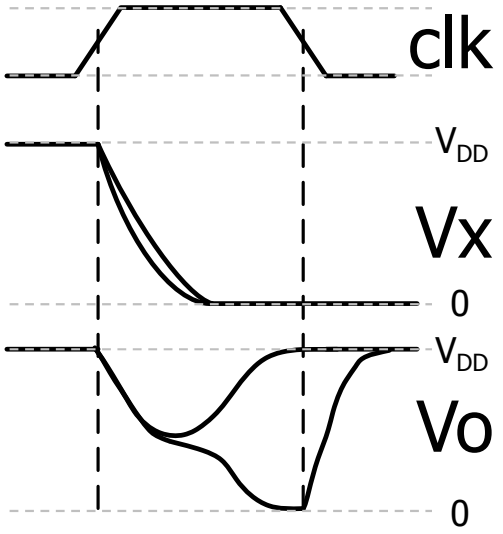

- ✓ Second stage regeneration through small Mn3, Mp3 in parallel with second stage
  - ✓ Full swing output
  - ✓ Smaller delay (versus regenerative comparator)
- ✓ Second stage regeneration current is controlled through NMOS transistor
- ✓ Only requires one clock phase

O. El-Hadidy, A. Roshan-Zamir, H.-W. Yang, and S. Palermo, "A 32 Gb/s 0.55 mW/Gbps PAM4 1-FIR 2-IIR Tap DFE Receiver in 65-nm CMOS," IEEE Symposium on VLSI Circuits, June 2015.

GP 65nm CMOS Prototype & Measurement Results

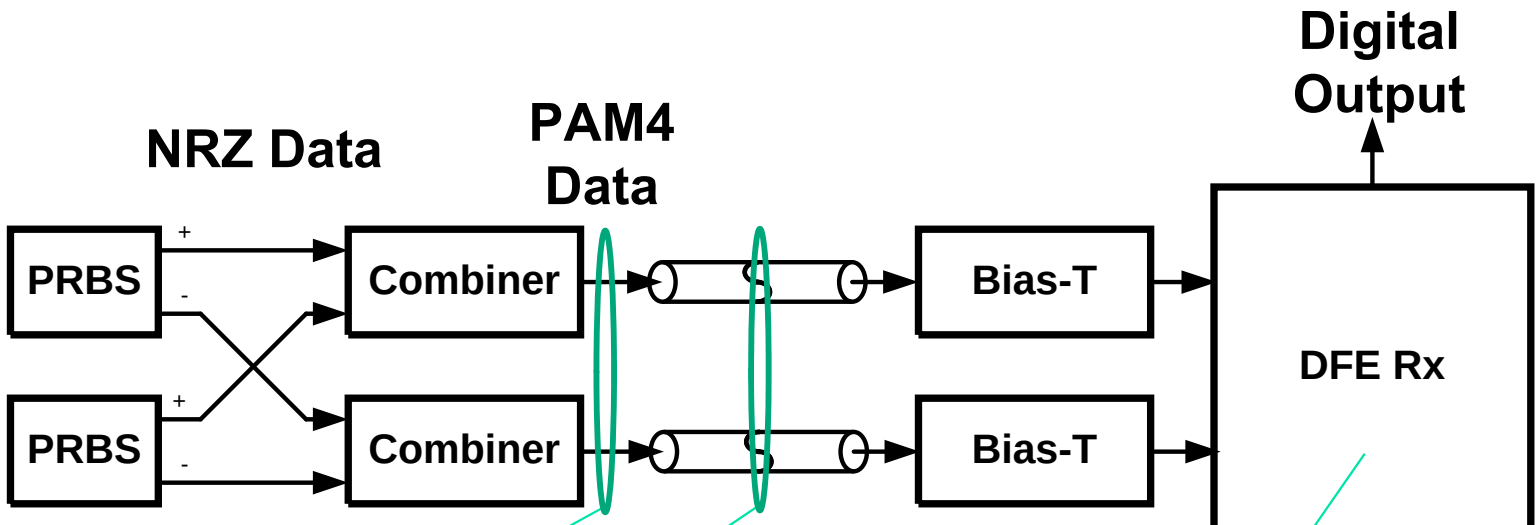

600 mV_{ppd}
25Gb/s PAM4 Data

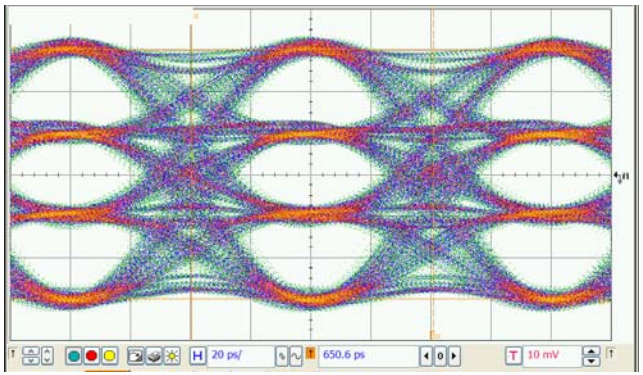

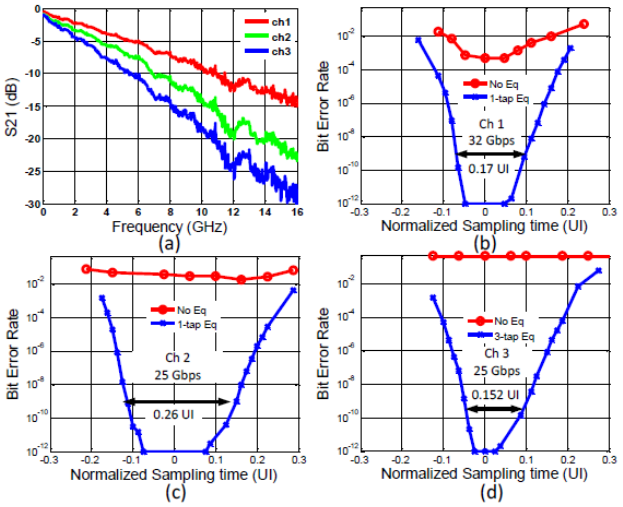

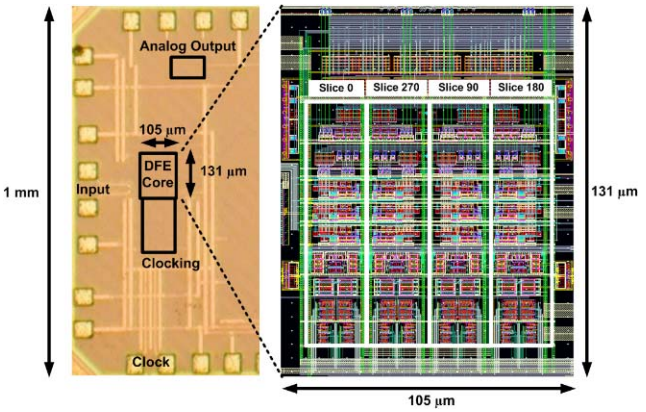

- At 32Gb/s consumes 17.7mW or 0.55mW/Gbps

O. El-Hadidy, A. Roshan-Zamir, H.-W. Yang, and S. Palermo, "A 32 Gb/s 0.55 mW/Gbps PAM4 1-FIR 2-IIR Tap DFE Receiver in 65-nm CMOS," IEEE Symposium on VLSI Circuits, June 2015.

Conclusion

- I/O transceivers need to achieve near 1pJ/b at 10+ Gb/s to support future systems
- Low-voltage operation with parallelism can achieve significant power savings
- Source synchronous architectures reduce clocking complexity
- Circuitry which supports fast power-state transitioning can reduce system average power
- Low-voltage equalizers are necessary to support channel loss for data rates >10Gb/s

Acknowledgements

- Many of the projects discussed today were collaborative works with Prof. Patrick Chiang's group at Oregon St

- Funding support from SRC and TI