

# ECEN474: (Analog) VLSI Circuit Design

## Fall 2010

---

### Lecture 26: High-Speed I/O Overview



Sam Palermo

Analog & Mixed-Signal Center

Texas A&M University

# Announcements

---

- Project
  - Preliminary report due Nov 19
- This lecture is not covered in exam 3

# Outline

---

- Introduction
- Electrical I/O Overview
  - Channel characteristics
  - Transmitter & receiver circuits
  - Clocking techniques & circuits
- Future trends & optical I/O
- Conclusion

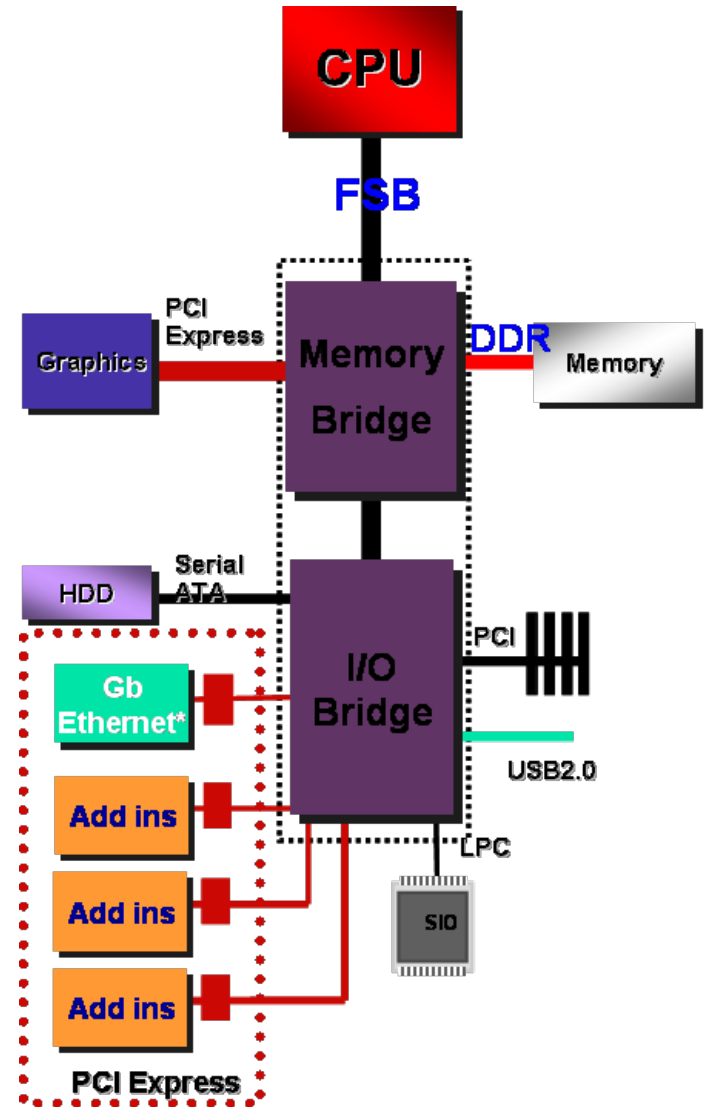
# ECEN 689: Special Topics in High-Speed Links Circuits and Systems

---

- Spring 2011
- <http://www.ece.tamu.edu/~spalermo/ecen689.html>
- Covers system level and circuit design issues relevant to high-speed electrical and optical links
- Channel Properties
  - Modeling, measurements, communication techniques
- Circuits
  - Drivers, receivers, equalizers, clocking
- Project
  - Link system design with statistical BER analysis tool
  - Circuit design of key interface circuits
- Prerequisite: ECEN 474 or my approval


# Desktop Computer I/O Architecture

- Many high-speed I/O interfaces
- Key bandwidth bottleneck points are memory (FSB) and graphics interfaces (PCIe)
- Near-term architectures
  - Integrated memory controller with serial I/O (>5Gb/s) to memory
  - Increasing PCIe from 2.5Gb/s (Gen1) to 8Gb/s (Gen3)
- Other serial I/O systems
  - Multi-processor systems
  - Routers




# Serial Link Applications

- Processor-to-memory
  - RDRAM (1.6Gbps), XDR DRAM (7.2Gbps), XDR2 DRAM (12.8Gbps)
- Processor-to-peripheral
  - PCIe (2.5, 5, 8Gbps), Infiniband (10Gbps), USB3 (4.8Gbps)
- Processor-to-processor
  - Intel QPI (6.4Gbps), AMD Hypertransport (6.4Gbps)
- Storage
  - SATA (6Gbps), Fibre Channel (20Gbps)
- Networks
  - LAN: Ethernet (1, 10Gbps)
  - WAN: SONET (2.5, 10, 40Gbps)
  - Backplane Routers: (2.5 – 12.5Gbps)



# Chip-to-Chip Signaling Trends

Lumped capacitance



**Decade**

**Speeds**


**Transceiver Features**

1980's

>10Mb/s

Inverter out, inverter in

Transmission line




1990's

>100Mb/s

Termination

Source-synchronous clk.

Lossy transmission line




2000's

>1 Gb/s

Pt-to-pt serial streams

Pre-emphasis equalization



Future

>10 Gb/s


Adaptive Equalization,  
Advanced low power clk.  
Alternate channel materials

Slide Courtesy of Frank O'Mahony & Brian Casper, Intel

# Increasing I/O Bandwidth Demand

- Single  $\Rightarrow$  Multi  $\Rightarrow$  Many-Core  $\mu$ Processors
- Tera-scale many-core processors will aggressively drive aggregate I/O rates

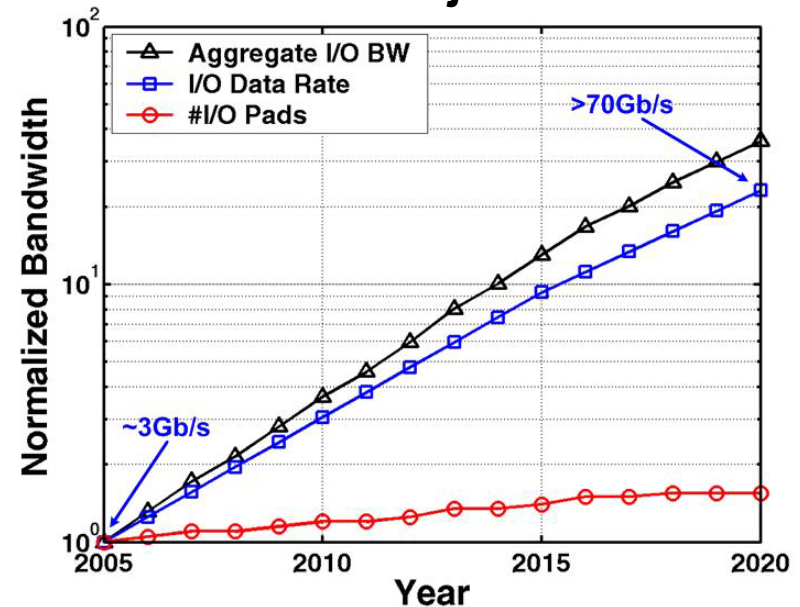
## Intel Teraflop Research Chip



- 80 processor cores
- On-die mesh interconnect network w/  $>2\text{Tb/s}$  aggregate bandwidth
- 100 million transistors
- $275\text{mm}^2$

S. Vangal *et al*, "An 80-Tile Sub-100W TeraFLOPS Processor in 65nm CMOS," *JSSC*, 2008.

## ITRS Projections\*



\*2006 International Technology Roadmap for Semiconductors

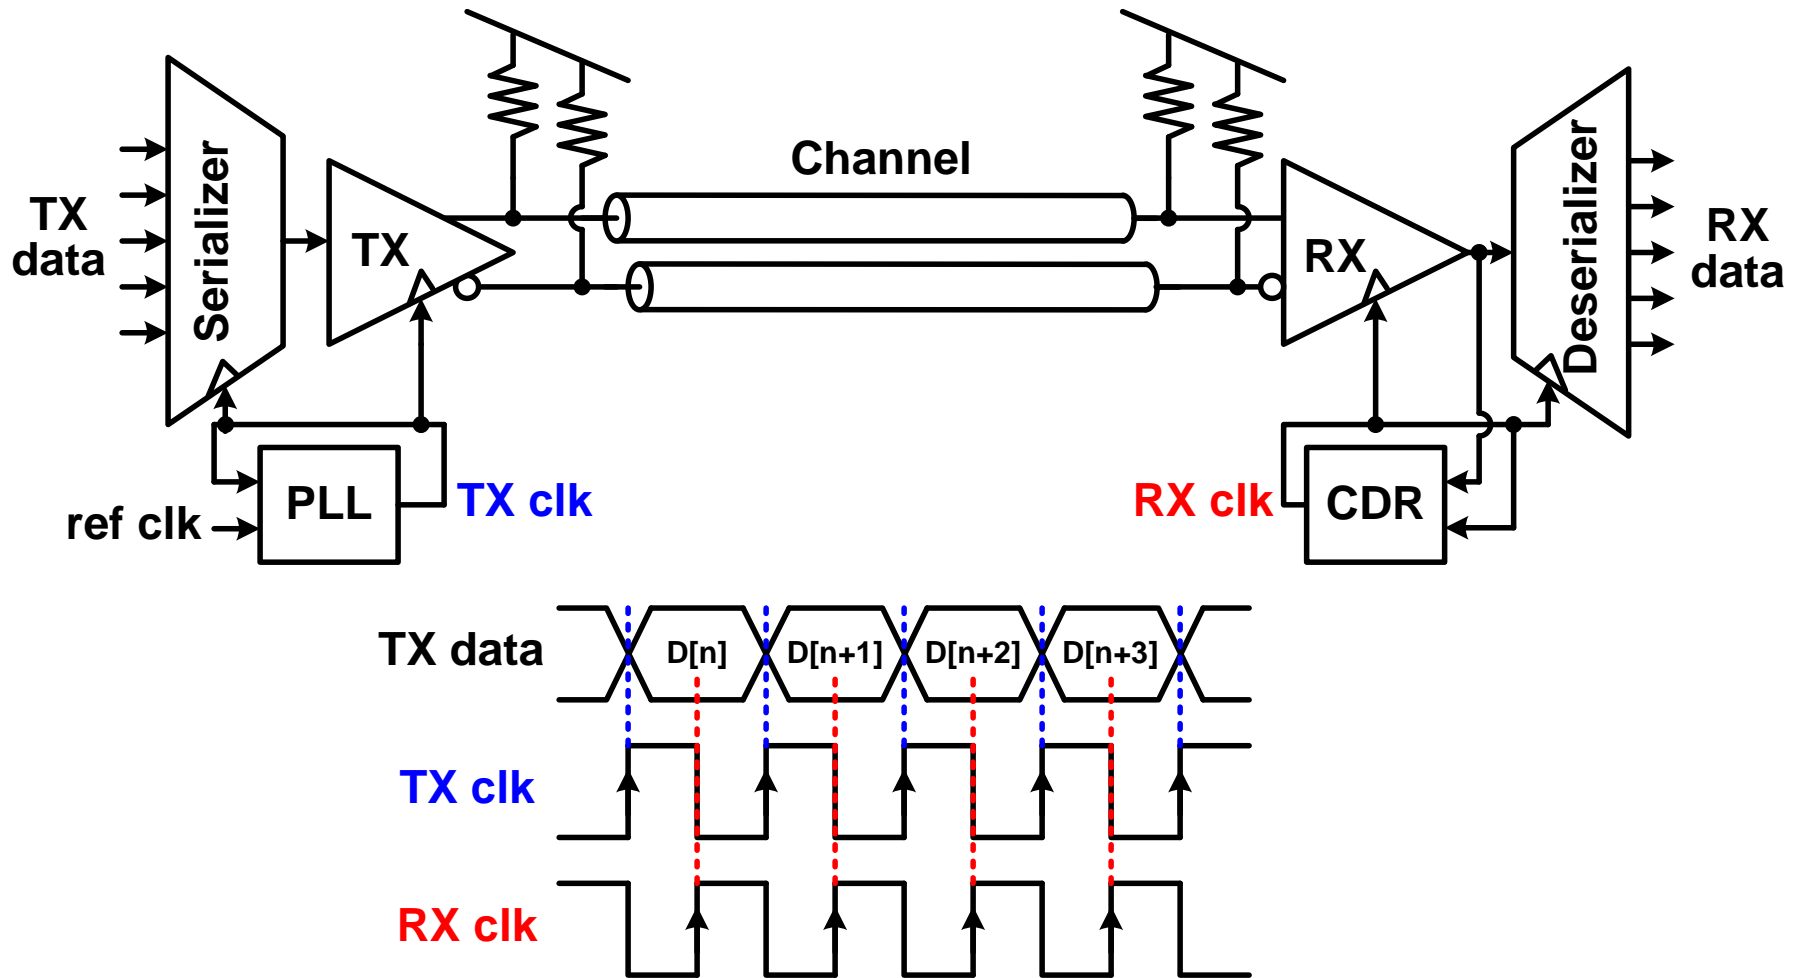


# Outline

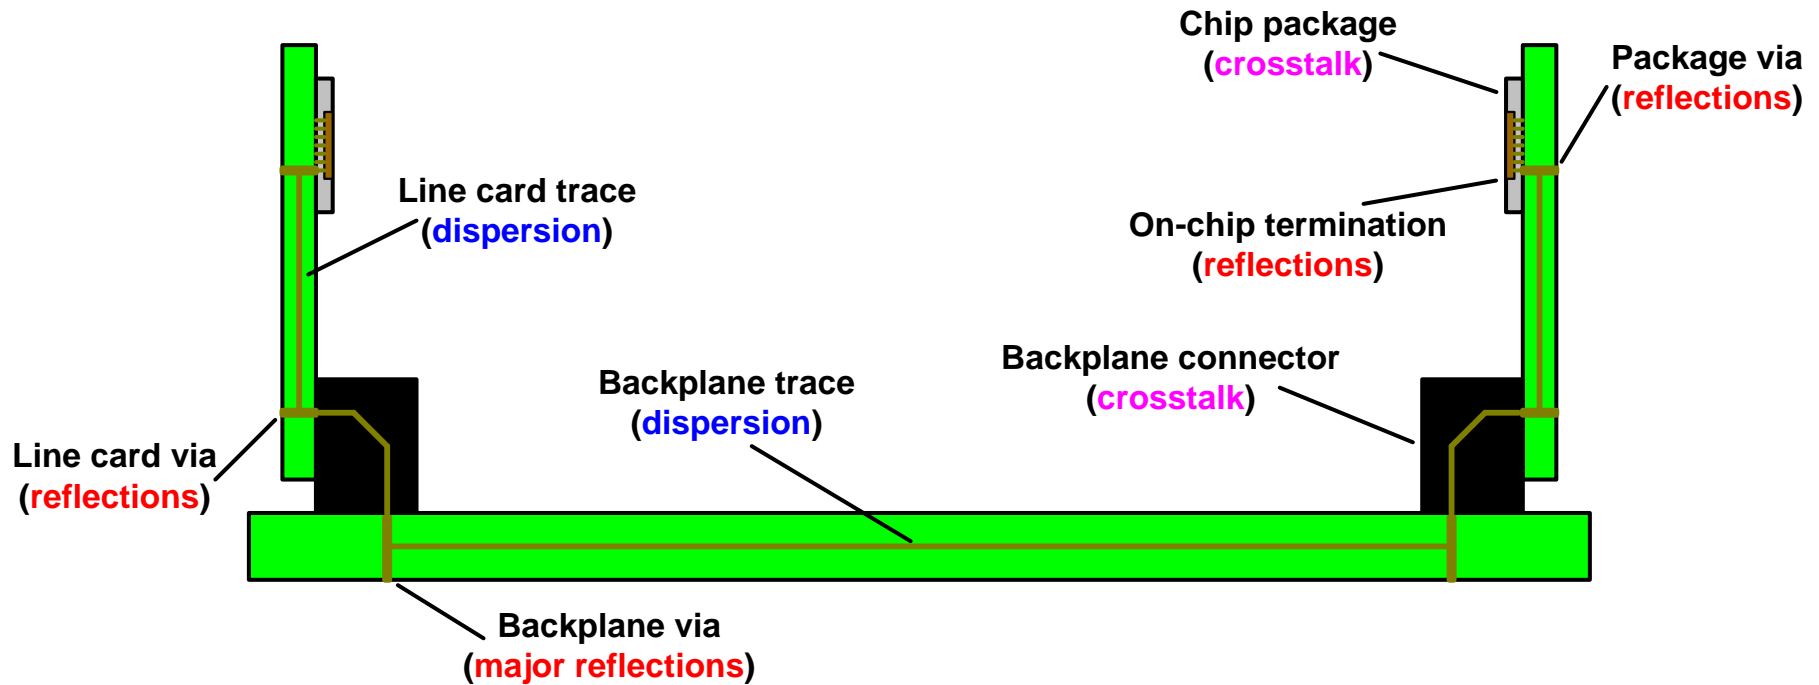
---

- Introduction
- Electrical I/O Overview
  - Channel characteristics
  - Transmitter & receiver circuits
  - Clocking techniques & circuits
- Future trends & optical I/O
- Conclusion

# High-Speed Electrical Link System



# Electrical Backplane Channel



- Frequency dependent loss
  - Dispersion & reflections
- Co-channel interference
  - Far-end (FEXT) & near-end (NEXT) crosstalk

# Loss Mechanisms

- Dispersion

$$\frac{V(x)}{V(0)} = e^{-(\alpha_R + \alpha_D)x}$$


- Skin effect,  $\alpha_R$

$$\text{Skin Depth, } \delta_{sd} = \left( \frac{\rho}{\mu\pi f} \right)^{1/2}$$


$$\alpha_R = \frac{R_{AC}}{2Z_0} = \frac{\rho L}{\delta_{sd} \pi D 2Z_0} = \frac{2.61 \times 10^{-7}}{\pi D 2Z_0} \sqrt{f}$$

- Dielectric loss,  $\alpha_D$

$$\alpha_D = \frac{\pi \sqrt{\epsilon_r} \tan \delta_D}{c} f$$



Dispersion Loss




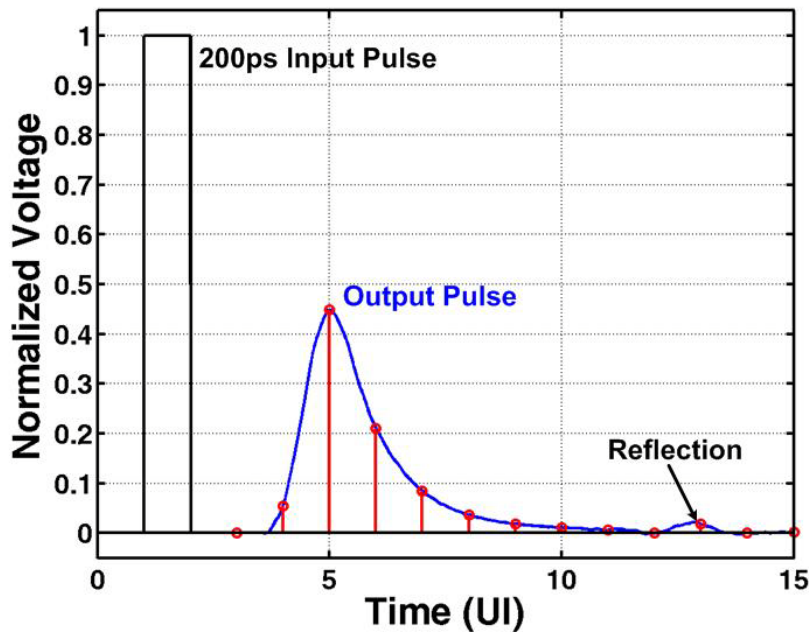

1m 8mil 50Ω stripguide with GETEK dielectric

B. Dally *et al*, "Digital Systems Engineering,"

# Reflections


- Commonly caused by board via stubs and on-chip termination mismatches

$$\frac{V_r}{V_i} = \frac{Z_r - Z_0}{Z_r + Z_0}$$




# Crosstalk


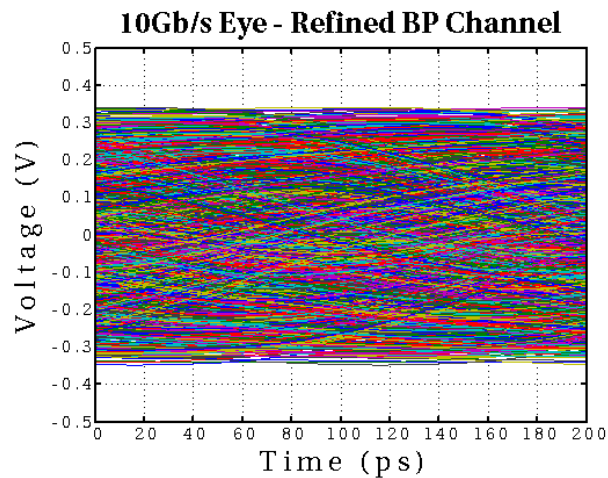



- Occurs mostly in package and board-to-board connectors
- FEXT is attenuated by channel response and has band-pass characteristic
- NEXT directly couples into victim and has high-pass characteristic








**17" Server Channel Response**



# Channel Performance Impact



# Channel Performance Impact





# Outline

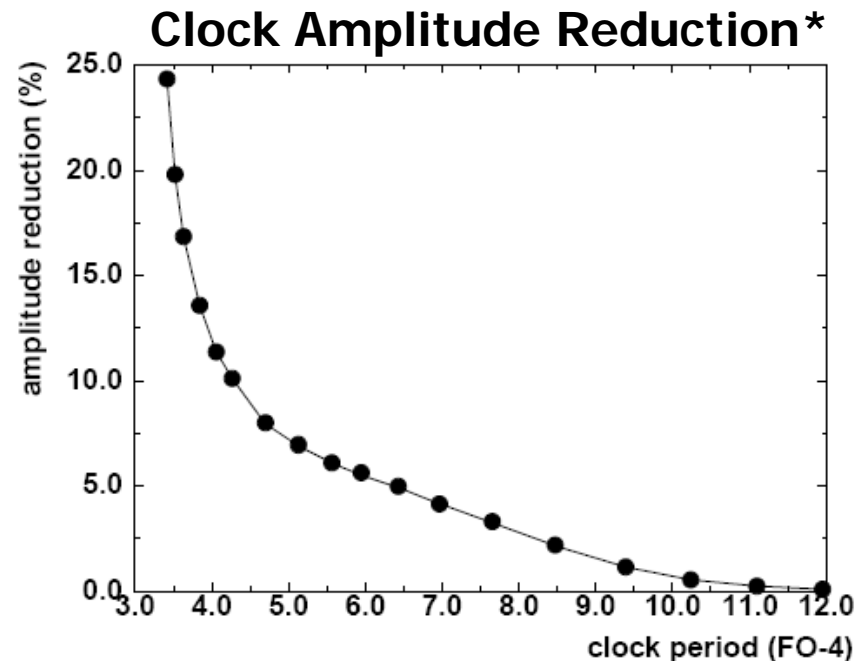

---

- Introduction
- Electrical I/O Overview
  - Channel characteristics
  - Transmitter & receiver circuits
  - Clocking techniques & circuits
- Future trends & optical I/O
- Conclusion

# Link Speed Limitations

- High-speed links can be limited by both the internal electronics and the channel
- Clock generation and distribution is key circuit bandwidth bottleneck
  - Requires data mux/demux to use multiple clock phases
  - Passives and/or CML techniques can extend circuit bandwidth at the expense of area and/or power
- Limited channel bandwidth is typically compensated with equalization circuits

$t_{FO4}$  in 90nm ~ 30ps




\*C.-K. Yang, "Design of High-Speed Serial Links in CMOS," 1998.


# Multiplexing Techniques

- Data mux/demux operation typically employs multiple clock phases
- $\frac{1}{2}$  rate architecture (DDR) is most common
  - Sends a bit on both the rising and falling edge of one differential clock
  - 50% duty cycle is critical
- Higher multiplexing factors with multiple clock phases further increases output data rate relative to on-chip clock frequency
  - Phase spacing/calibration is critical

2:1 Mux




8:1 Multiplexing TX\*




\*C.-K. Yang, "Design of High-Speed Serial Links in CMOS," 1998.

# Current vs Voltage-Mode Driver

- Signal integrity considerations (min. reflections) requires  $50\Omega$  driver output impedance
- To produce an output drive voltage
  - Current-mode drivers use Norton-equivalent parallel termination
    - Easier to control output impedance
  - Voltage-mode drivers use Thevenin-equivalent series termination
    - Potentially  $\frac{1}{2}$  to  $\frac{1}{4}$  the current for a given output swing



**Current-Mode**

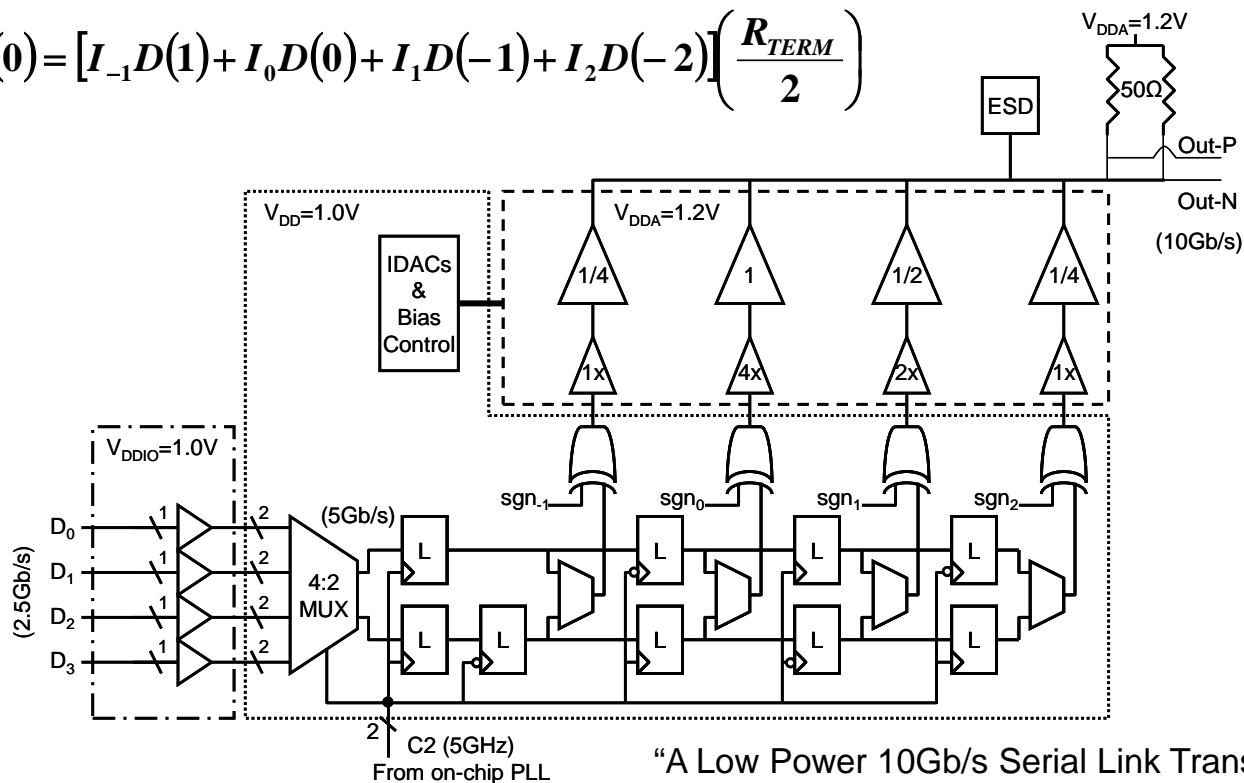


**Voltage-Mode**

# TX FIR Equalization




- TX FIR filter pre-distorts transmitted pulse in order to invert channel distortion at the cost of attenuated transmit signal (de-emphasis)

$$V_{out}(0) = [I_{-1}D(1) + I_0D(0) + I_1D(-1) + I_2D(-2)] \left( \frac{R_{TERM}}{2} \right)$$





"A Low Power 10Gb/s Serial Link Transmitter in 90-nm CMOS," A. Rylyakov et al., CSICS 2005

# 6Gb/s TX FIR Equalization Example

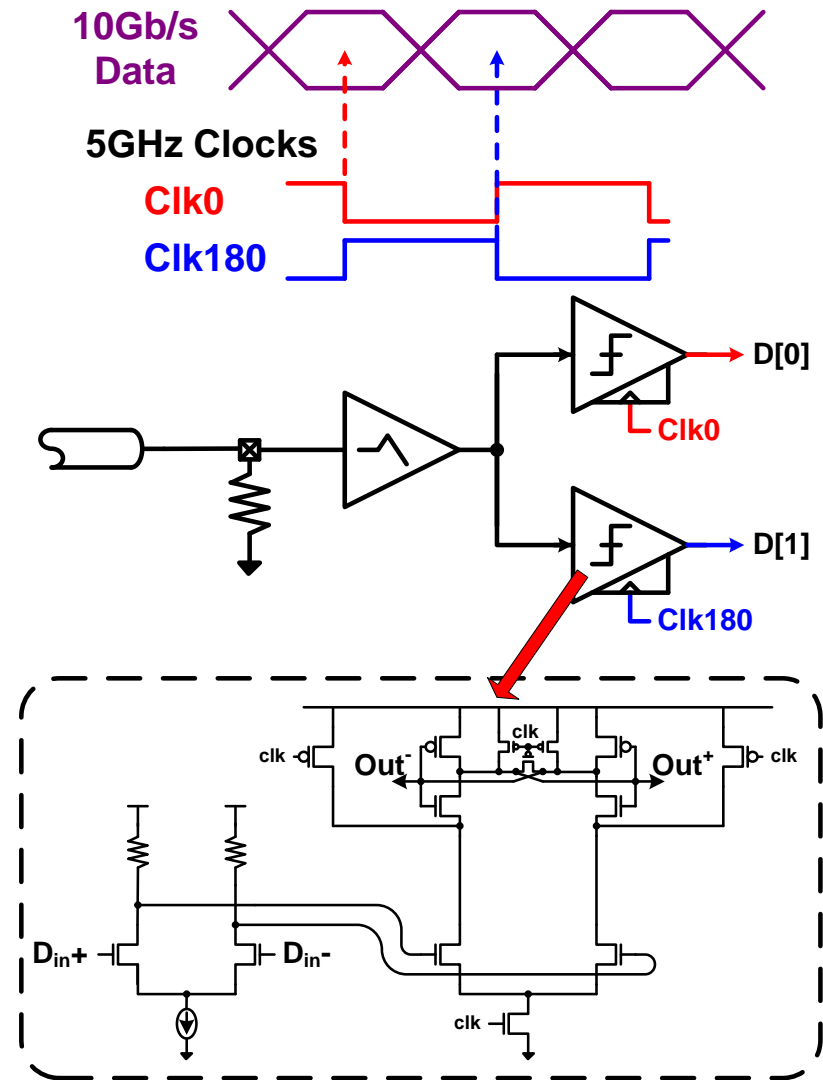


- Pros
  - Simple to implement
  - Can cancel ISI in precursor and beyond filter span
  - Doesn't amplify noise
  - Can achieve 5-6bit resolution
- Cons
  - Attenuates low frequency content due to peak-power limitation
  - Need a "back-channel" to tune filter taps



# Demultiplexing RX

- Input pre-amp followed by comparator segments
  - Pre-amp may implement peaking filtering
  - Comparator typically includes linear-amp & regenerative (positive feedback) latch
- Demultiplexing allows for lower clock frequency relative to data rate and extra regeneration and pre-charge time in comparators



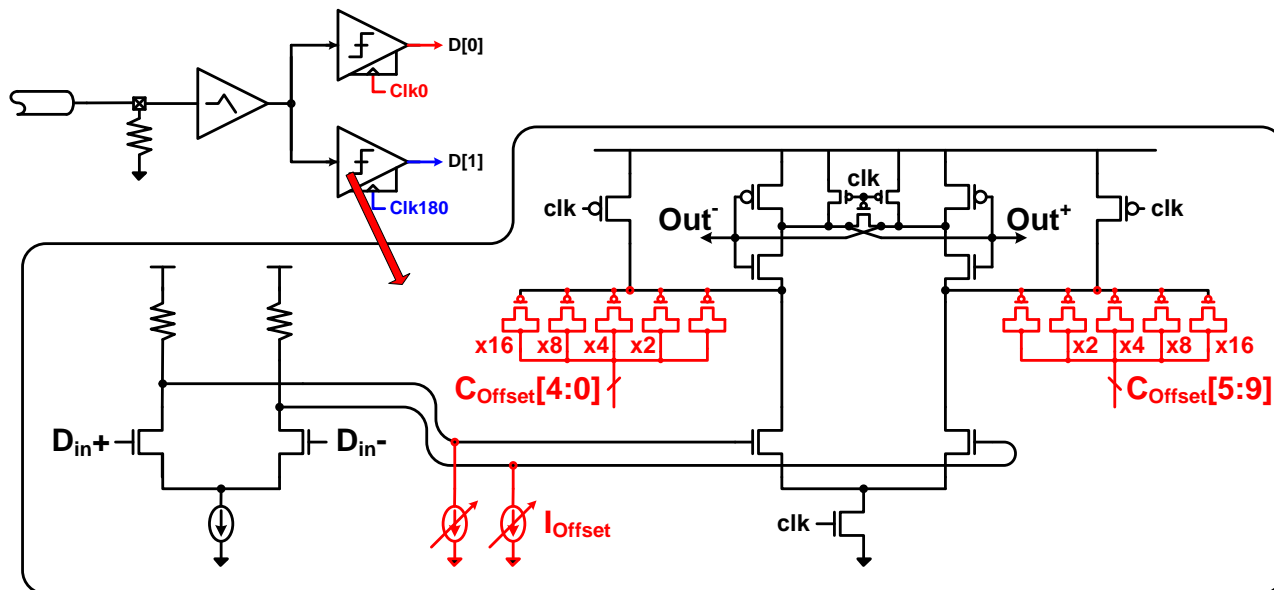
# RX Sensitivity

- RX sensitivity is a function of the input referred noise, offset, and min latch resolution voltage

$$v_S^{pp} = 2v_n^{rms} \sqrt{SNR} + v_{min} + v_{offset}^* \quad \text{Typical Values : } v_n^{rms} = 1mV_{rms}, v_{min} + v_{offset}^* < 2mV$$


$$\text{For BER} = 10^{-12} (\sqrt{SNR} = 7) \Rightarrow v_S^{pp} = 17mV_{pp}$$

- Circuitry is required to reduce input offset from a potentially large uncorrected value ( $>50mV$ ) to near 1mV





# RX Equalization #1: RX FIR




- Pros

- With sufficient dynamic range, can amplify high frequency content (rather than attenuate low frequencies)
- Can cancel ISI in pre-cursor and beyond filter span
- Filter tap coefficients can be adaptively tuned without any back-channel

- Cons



- Amplifies noise/crosstalk
- Implementation of analog delays
- Tap precision

Eye-Pattern Diagrams at 1Gb/s on CAT5e\*


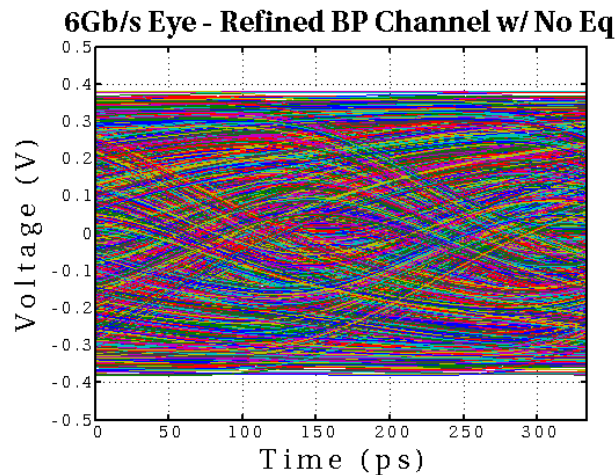


\*D. Hernandez-Garduno and J. Silva-Martinez, "A CMOS 1Gb/s 5-Tap Transversal Equalizer based on 3<sup>rd</sup>-Order Delay Cells," ISSCC, 2007.


# RX Equalization #2: RX CTLE



- Pros
  - Provides gain and equalization with low power and area overhead
  - Can cancel both pre-cursor and long-tail ISI
- Cons
  - Generally limited to 1<sup>st</sup> order compensation
  - Amplifies noise/crosstalk
  - PVT sensitivity
  - Can be hard to tune



# RX Equalization #3: RX DFE




- Pros

- No noise and crosstalk amplification
- Filter tap coefficients can be adaptively tuned without any back-channel


- Cons

- Cannot cancel precursor ISI
- Critical feedback timing path
- Timing of ISI subtraction complicates CDR phase detection


Refined BP Channel 6Gb/s Pulse Responses



6Gb/s Eye - Refined BP Channel w/ No Eq



6Gb/s Eye - Refined BP Channel w/ RX DFE Eq





# Outline

---

- Introduction
- Electrical I/O Overview
  - Channel characteristics
  - Transmitter & receiver circuits
  - **Clocking techniques & circuits**
- Future trends & optical I/O
- Conclusion

# Clocking Architecture #1

## Source Synchronous Clocking



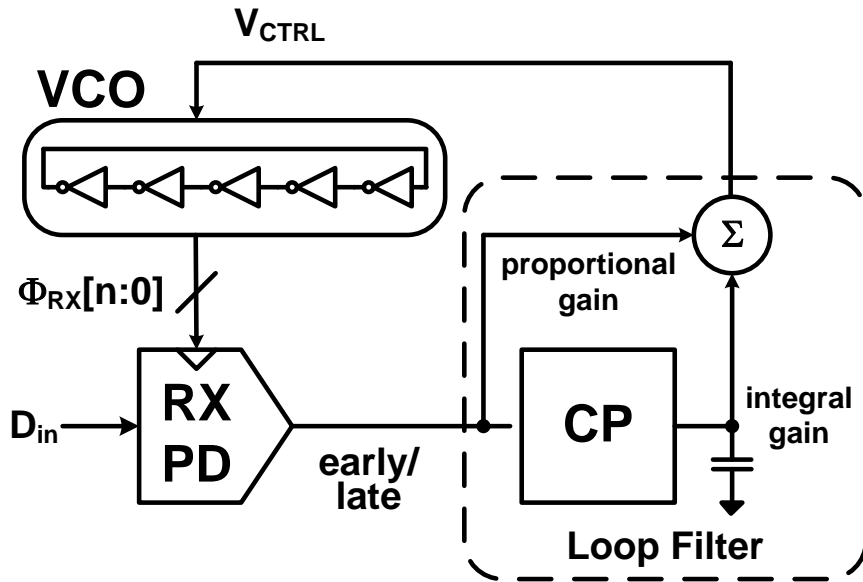
\*S. Sidiropoulos, "High Performance Inter-Chip Signalling," 1998.

- Common high-speed reference clock is forwarded from TX chip to RX chip
- "Coherent" clocking allows high frequency jitter tracking
  - Jitter frequency lower than delay difference (typically less than 10bits) can be tracked
  - Allows power down of phase detection circuitry
    - Only periodic acquisition vs continuous tracking
- Requires one extra clock channel
- Need good clock receive amplifier as the forwarded clock can get attenuated by the low pass channel
- Low pass channel causes jitter amplification

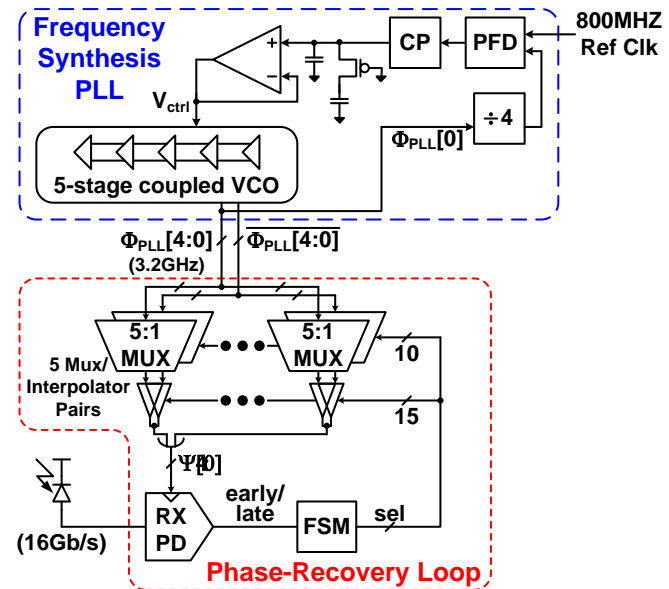
# Clocking Architecture #2

## Embedded Clocking (CDR)

PLL-based CDR

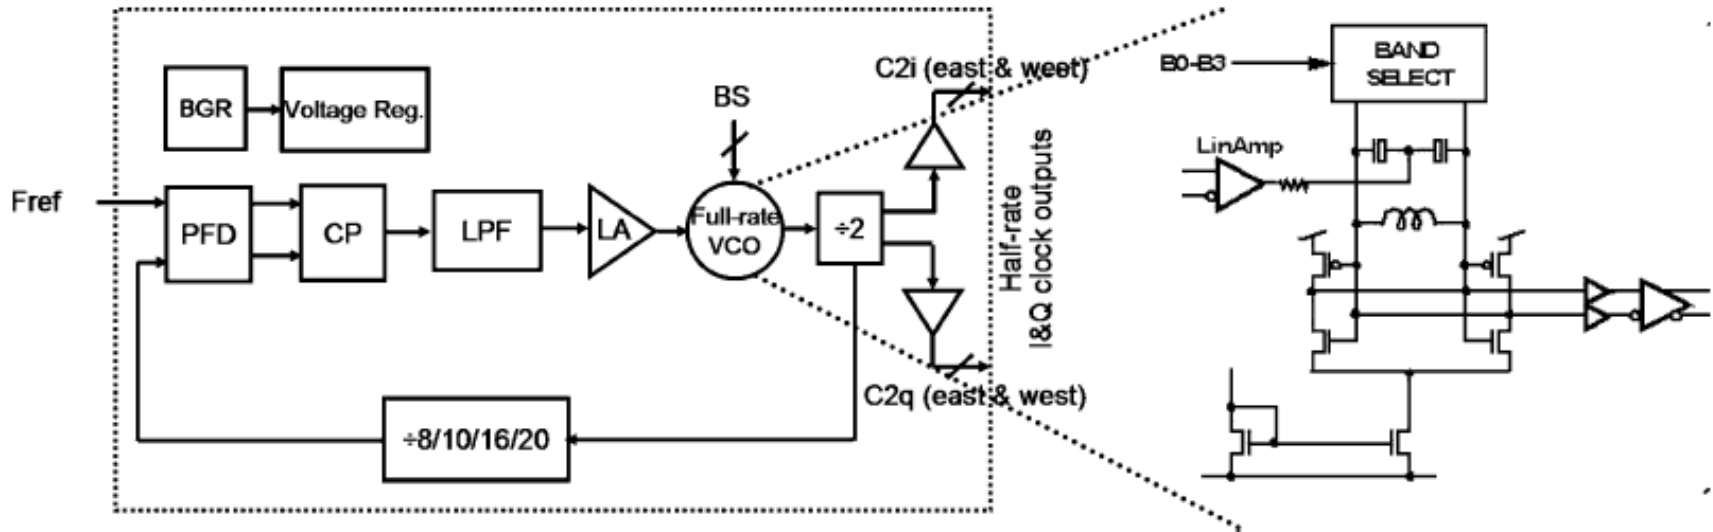


Dual-Loop CDR



- Clock frequency and optimum phase position are extracted from incoming data stream
- Phase detection continuously running
- Jitter tracking limited by CDR bandwidth
  - With technology scaling we can make CDRs with higher bandwidths and the jitter tracking advantages of source synchronous systems is diminished
- CDR can be implemented as a stand-alone PLL or as a “dual-loop” architecture with a PLL or DLL and phase interpolators (PI)


# Phase-Locked Loop (PLL)



\*J. Bulzacchelli *et al*, "A 10Gb/s 5Tap DFE/4Tap FFE Transceiver in 90nm CMOS Technology," JSSC, 2006.

- Used for frequency synthesis at TX and embedded-clocked RX
- Second/third order loop
  - Charge pump & integrating loop filter produces voltage to control VCO frequency
  - Output phase is integration of VCO frequency
  - Zero required in loop filter for stability
- Low-noise VCO (or high BW PLL) required to minimize jitter accumulation

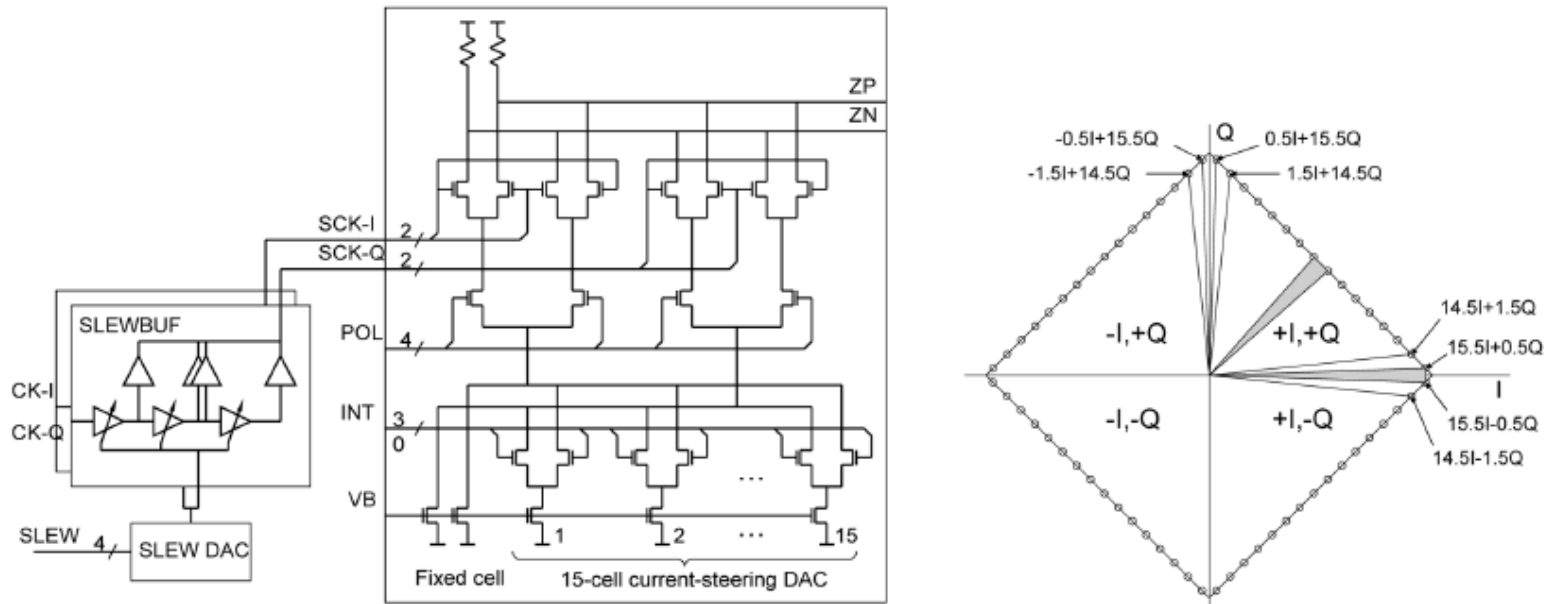
# Delay-Locked Loop (DLL)



- Typically used to generate multiple clock phases in RX
- First order loop guarantees stability
- Delay line doesn't accumulate jitter like a VCO
- Difficult to use for frequency synthesis



# Phase Interpolator (PI)

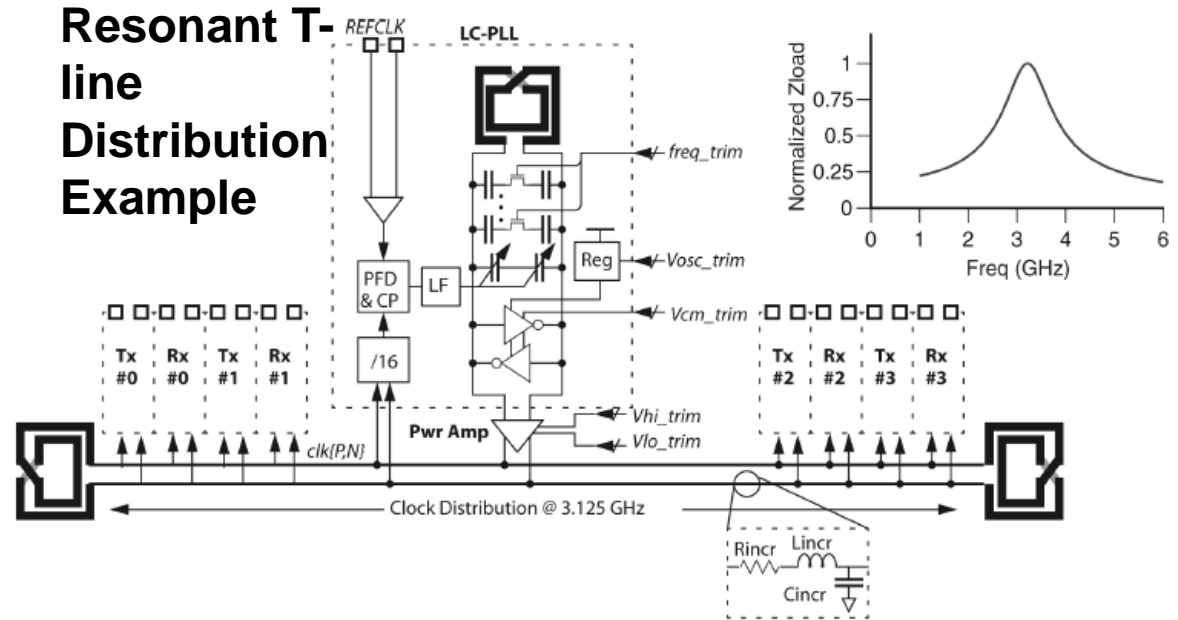


\*J. Bulzacchelli *et al*, "A 10Gb/s 5Tap DFE/4Tap FFE Transceiver in 90nm CMOS Technology," JSSC, 2006.

- Interpolators mix between two clock phases to produce the fine resolution clock phases used by the RX samplers
- Critical to limit bandwidth of PI mixing node for good linearity
  - Hard to design over wide frequency range without bandwidth adjustment and/or input slew-rate control

# Clock Distribution

- Careful clock distribution is required in multi-channel I/O systems
- Different distribution architectures trade-off jitter, power, area, and complexity



\*J. Poulton *et al*, "A 14mW 6.25Gb/s Transceiver in 90nm CMOS," JSSC, 2007.

Architecture	Jitter	Power	Area	Complexity
Inverter	Moderate	Moderate	Low	Low
CML	Good	High	Moderate	Moderate
T-line	Good	Low	Low	Moderate
Resonant T-line	Excellent	Low	High	High


# Outline

---

- Introduction
- Electrical I/O Overview
  - Channel characteristics
  - Transmitter & receiver circuits
  - Clocking techniques & circuits
- Future trends & optical I/O
- Conclusion

# It's about the Energy Efficiency, ...

- Energy efficiency is paramount
  - Emphasis shifting away from maximizing Gb/s to minimizing mW/Gb/s or pJ/bit
- Current commercial high-speed links are  $\sim 10$  mW/Gb/s
- Research caliber links can achieve 1-3 mW/Gb/s at 5-10 Gb/s
  - Emphasis on adaptive voltage scaling, digital calibration techniques, refining electrical channel
- Need to achieve sub-1 mW/Gb/s at data rates  $\sim 10$  Gb/s
- Future systems are projected at even higher data rates (20+ Gb/s)
  - Can we still do electrical?





# Other Trends

---

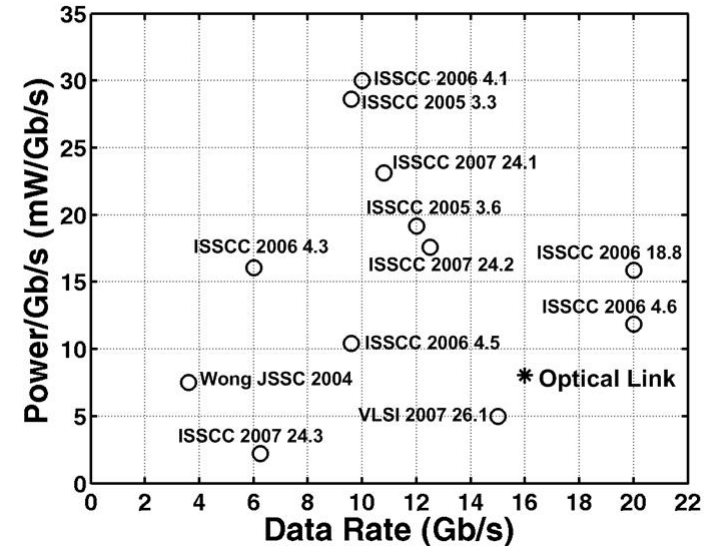
- Can we do better than simple NRZ modulation?
  - Multi-level (4/8-PAM)
  - Multi-tone
  - Duo-binary
- Active crosstalk cancellation
  - Package constraints require high density and high data rate
- ADC-based RX front-ends
  - Get to digital ASAP
  - Allows improved SNR front-ends, but probably doesn't save power

# Chip-to-Chip Optical Interconnects

negligible frequency dependent loss



- Optical interconnects remove many channel limitations
  - Reduced complexity and power consumption
  - Potential for high information density with wavelength-division multiplexing (WDM)



\*S. Palermo *et al*, "A 90nm CMOS 16Gb/s Transceiver for Optical Interconnects," JSSC, 2008.

# Conclusion

---

- High-speed I/O systems offer challenges in both circuit and communication system design
  - High-speed TX/RX, low jitter clocking, and efficient equalizer circuits
- Key issue with scaling high-speed I/O is meeting the energy efficiency targets required by future systems ( $\rightarrow 1\text{mW/Gb/s}$ )
  - Requires circuit improvements and constant electrical channel refinement
  - Optical I/O is a major candidate in this space

# Interested In Research In This Area?

---

- Graduate Students
  - Take the 689 class
- Undergraduate Students
  - Opportunities exist for undergraduate research credits (491)