Channel Equalization Techniques for High-Speed Electrical Links

Sam Palermo
Analog & Mixed-Signal Center
Texas A&M University

Outline

- Introduction
- Channel characteristics
- Equalizer circuits
- Equalizer adaptation techniques
- Optical interconnects
- Conclusion

Chip-to-Chip Signaling Trends

Lumped capacitance

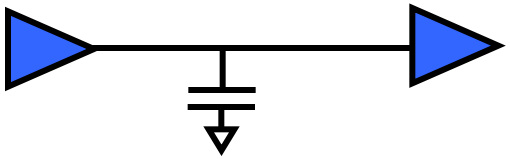

Transmission line

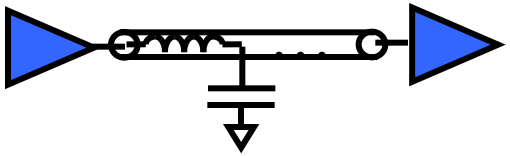

Lossy transmission line

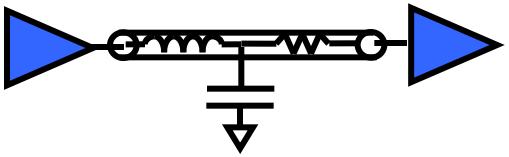

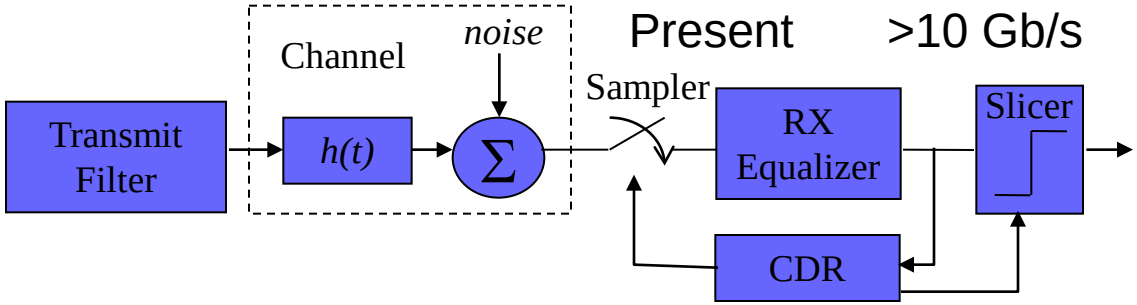

Decade	Speeds	Transceiver Features
1980's	>10Mb/s	Inverter out, inverter in
1990's	>100Mb/s	Termination Source-synchronous clk.
2000's	>1 Gb/s	Pt-to-pt serial streams Pre-emphasis equalization
Present	>10 Gb/s	Adaptive Equalization, Advanced low power clk. Alternate channel materials

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

Increasing Bandwidth Demand

- Single ⇒ Multi ⇒ Many-Core μProcessors
- Tera-scale many-core processors will aggressively drive aggregate inter- and intra-chip bandwidth

Intel Teraflop Research Chip

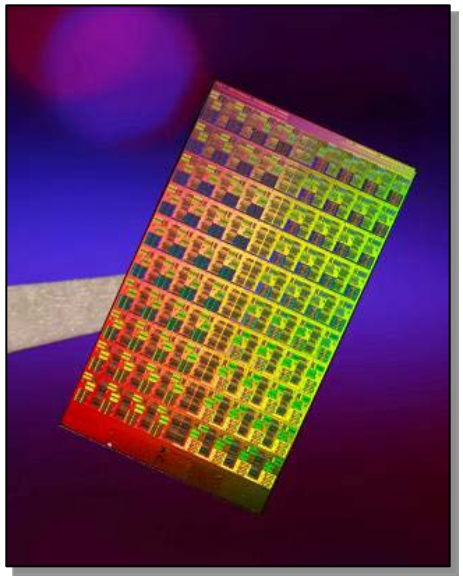

- 80 processor cores
- On-die mesh interconnect network w/ >2Tb/s aggregate bandwidth
- 100 million transistors
- 275mm^2

[Vangal JSSC 2008]

ITRS Projections*

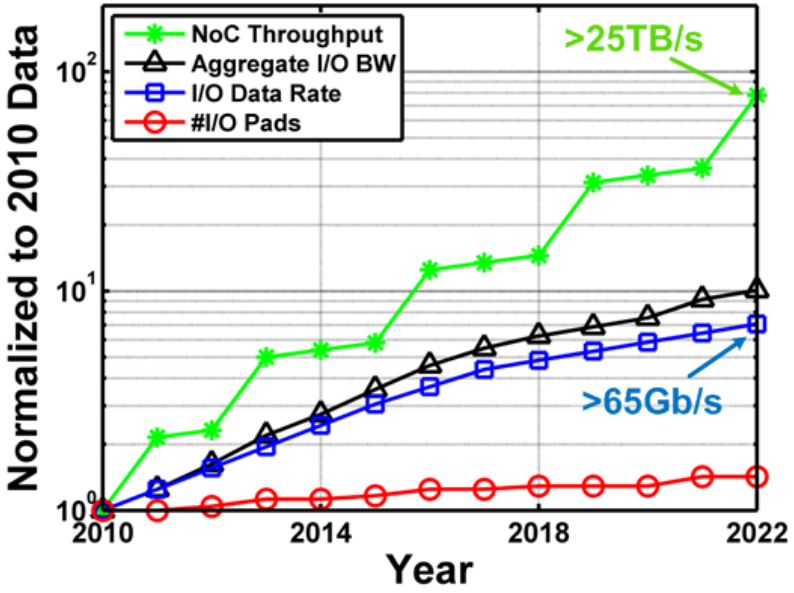

[ITRS 2009]

High-Speed Electrical Link System

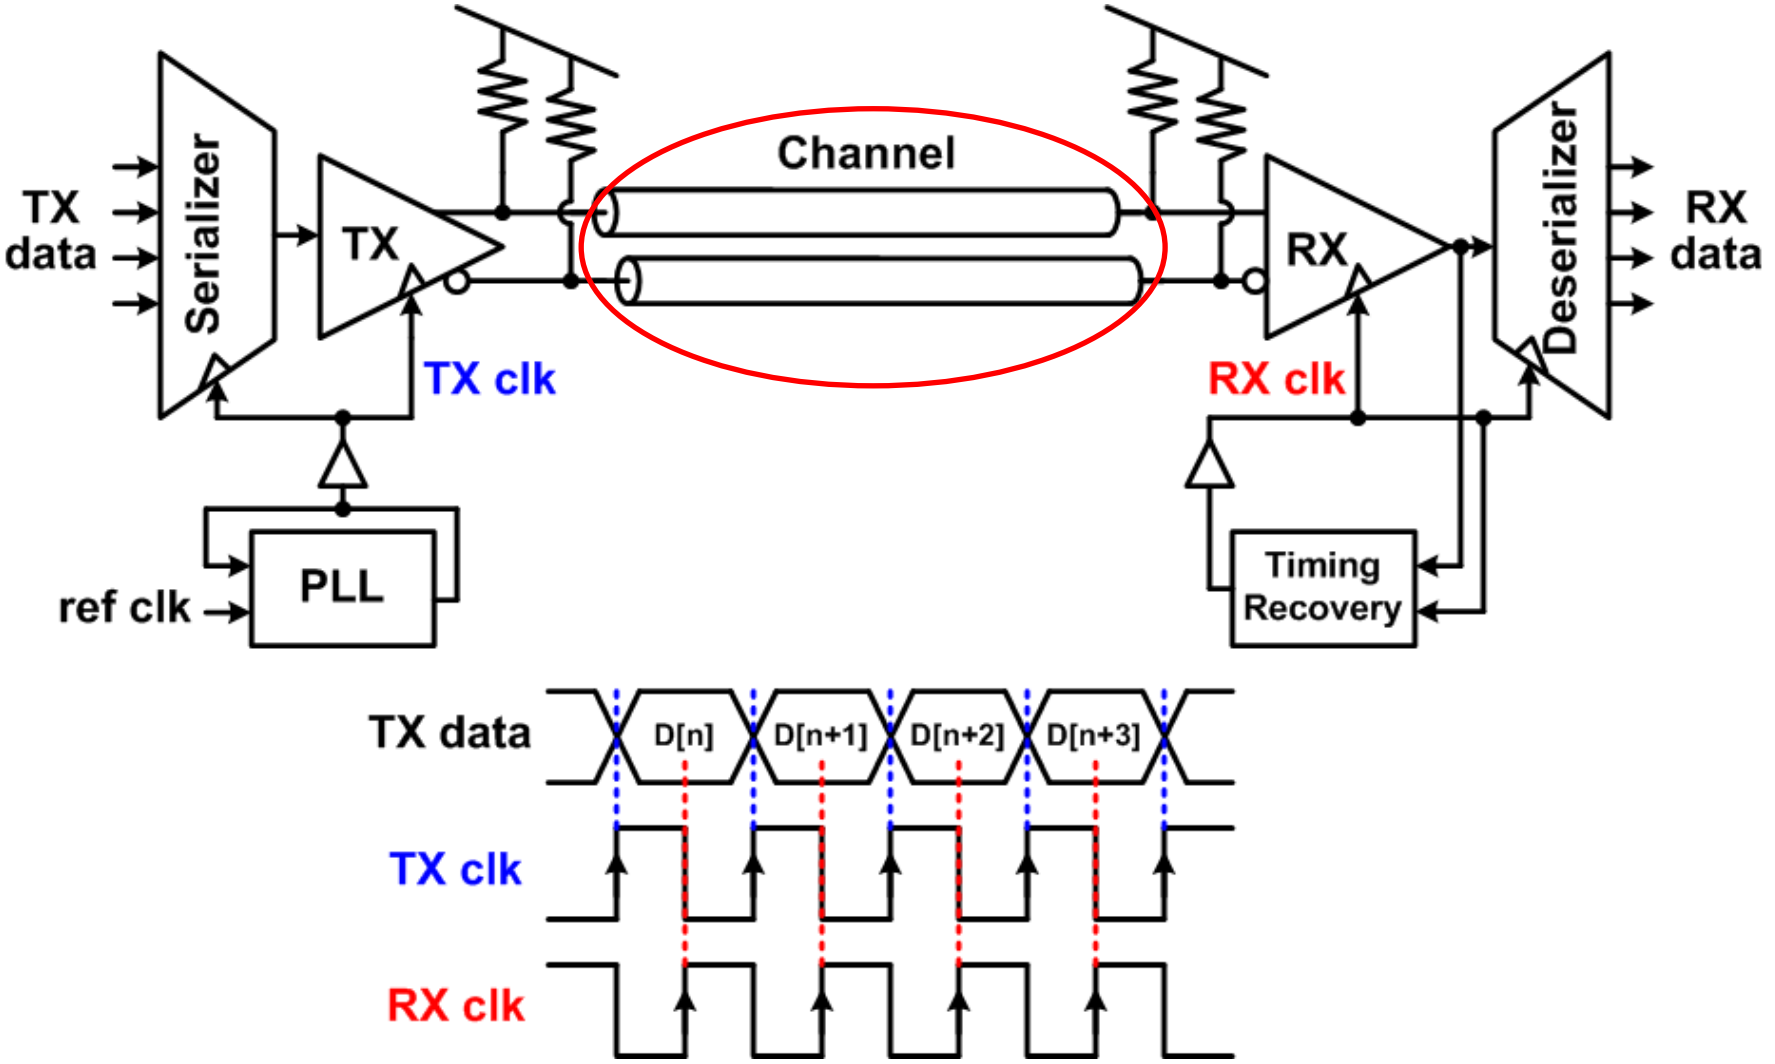

Outline

- Introduction
- Channel characteristics
- Equalizer circuits
- Equalizer adaptation techniques
- Optical interconnects
- Conclusion

Electrical Backplane Channel

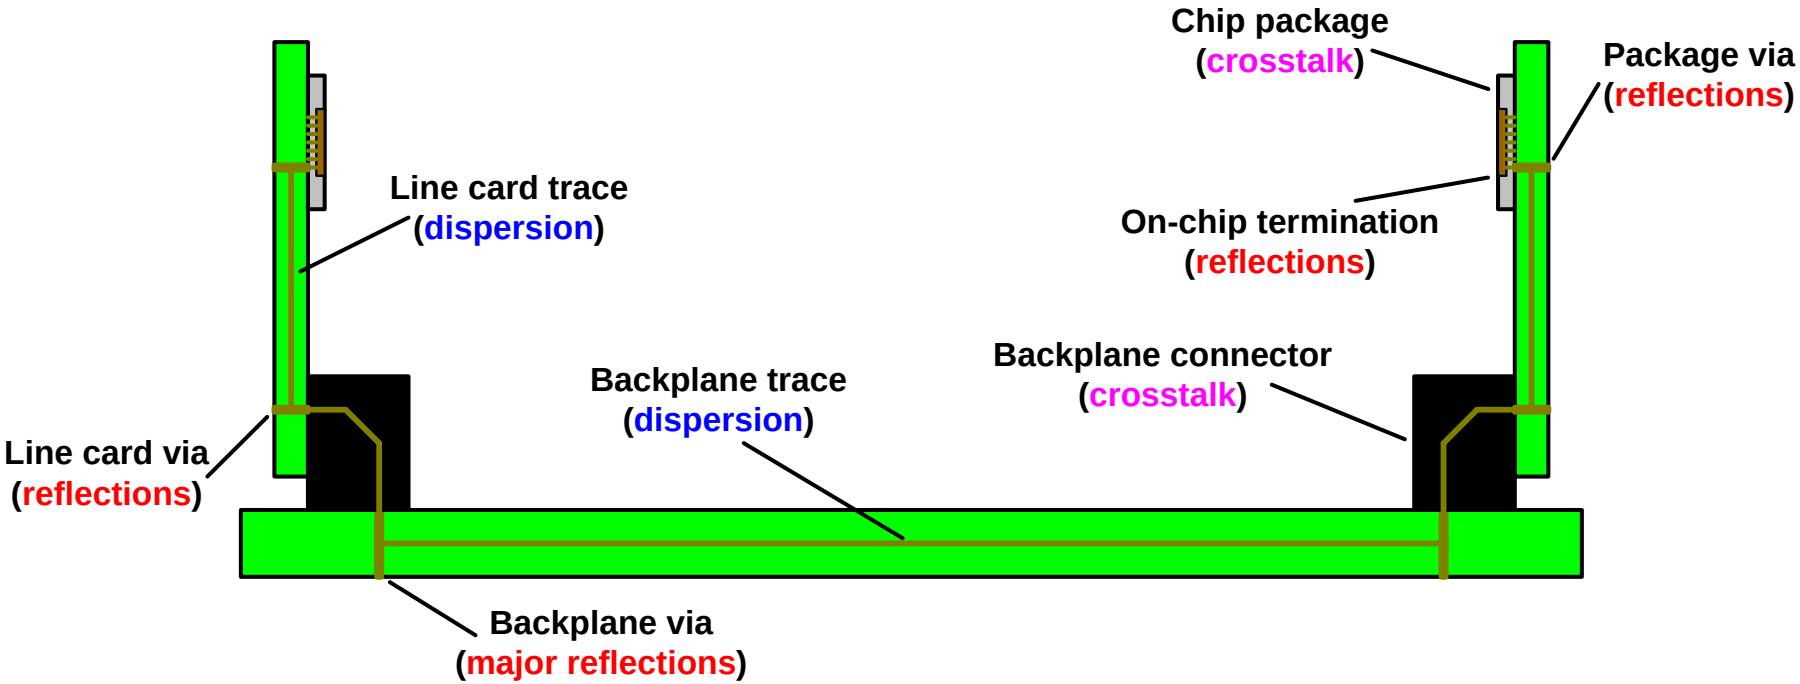

- Frequency dependent loss
  - Dispersion & reflections
- Co-channel interference
  - Far-end (FEXT) & near-end (NEXT) crosstalk

Loss Mechanisms

- Dispersion

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

  - Skin effect, α_R

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

$$\alpha_R = \frac{R_{AC}}{2Z_0} = \frac{\rho L}{\delta_{sd}\pi D 2 Z_0} = \frac{2.61\times10^{-7}}{\pi D 2 Z_0}\sqrt{f}$$

  - Dielectric loss , α_D

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

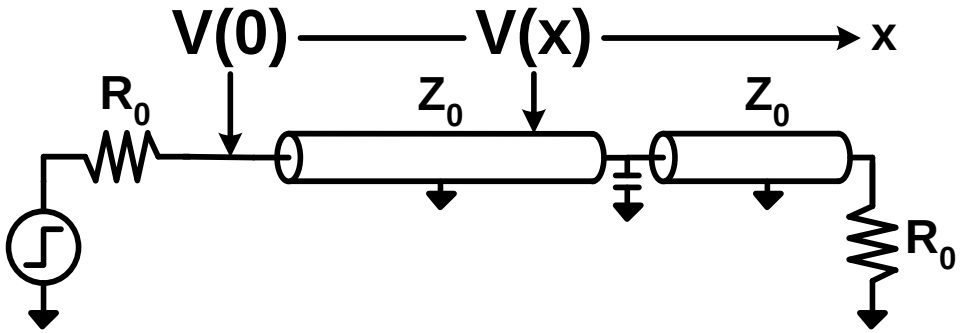

Dispersion Loss

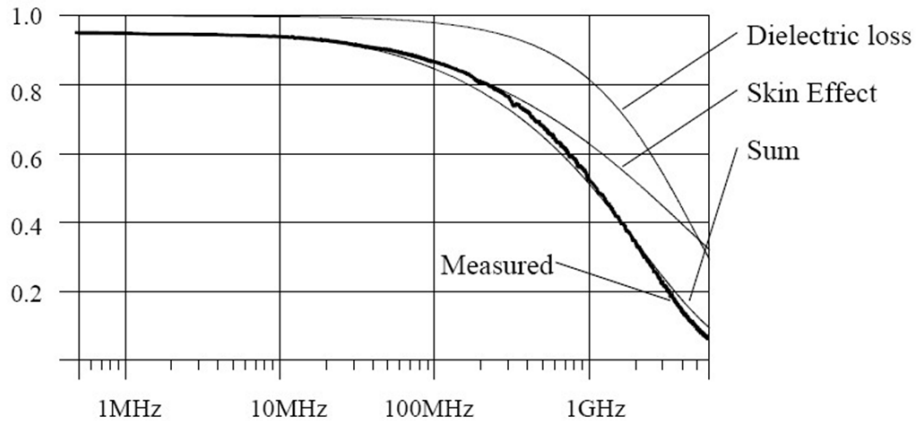

1m 8mil 50Ω stripguide with GETEK dielectric

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

Reflections

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

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

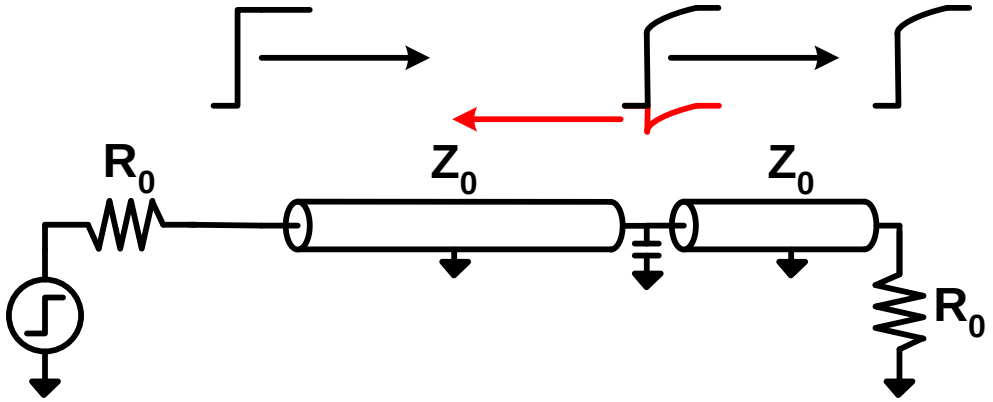

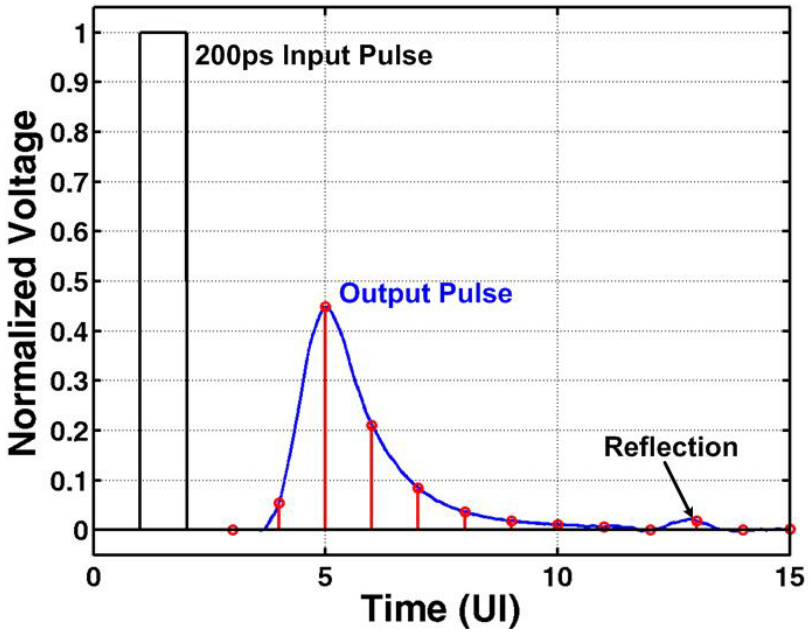

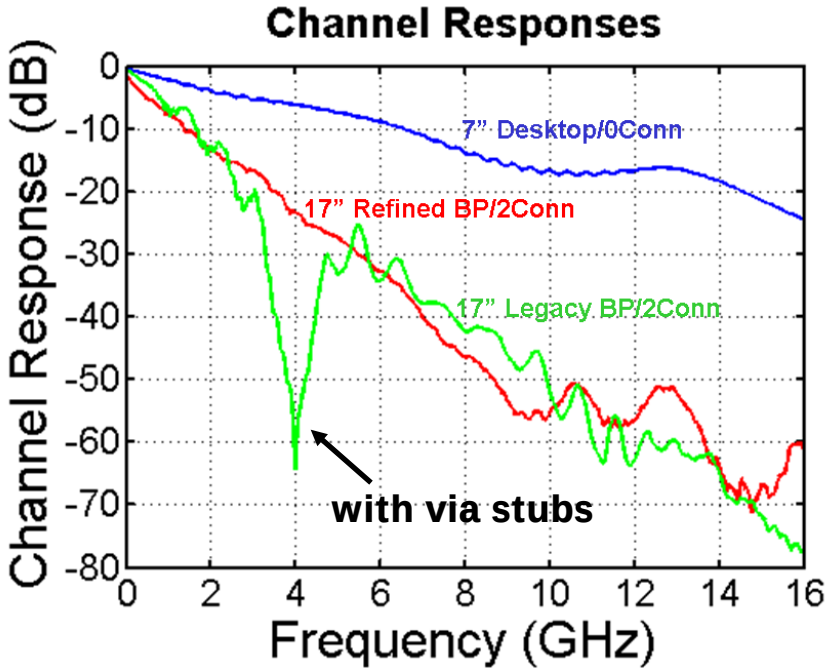

Crosstalk

- 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

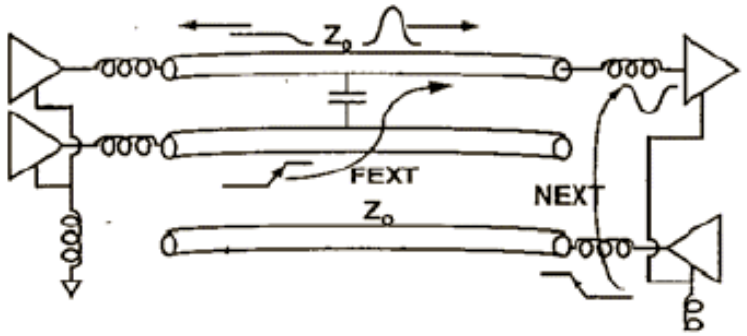

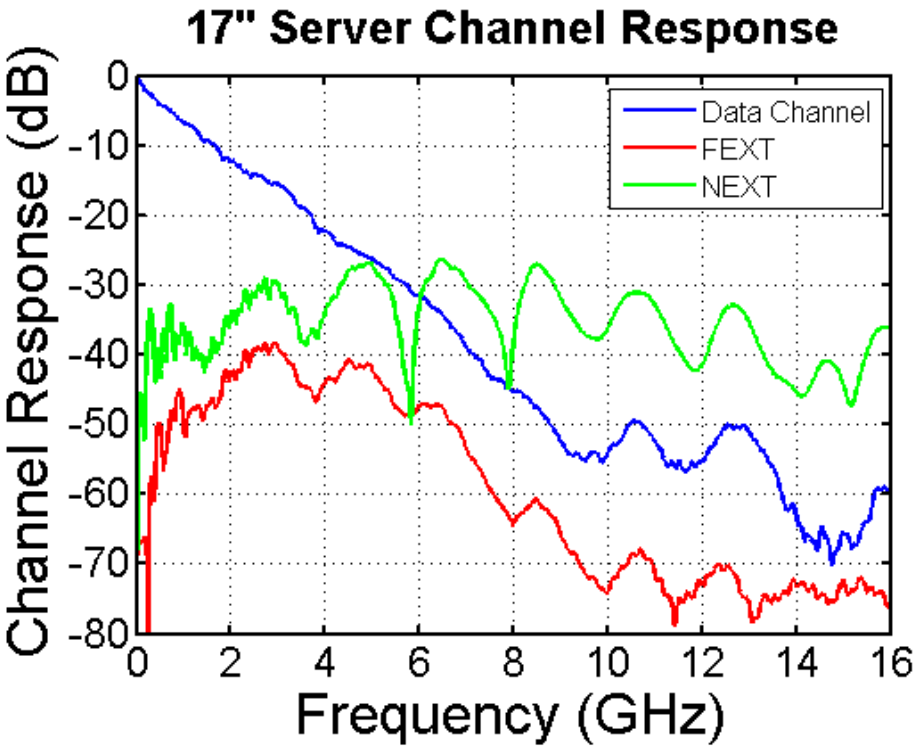

Channel Performance Impact

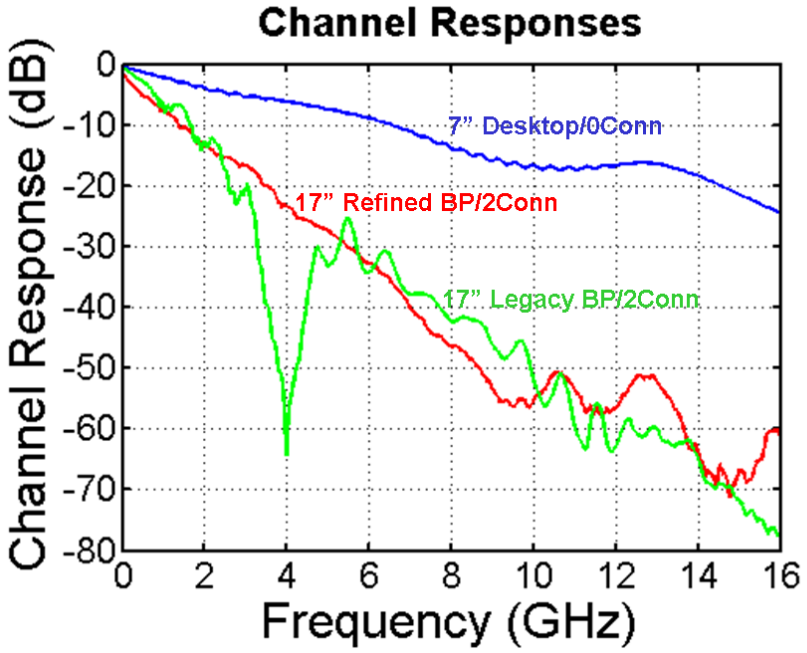

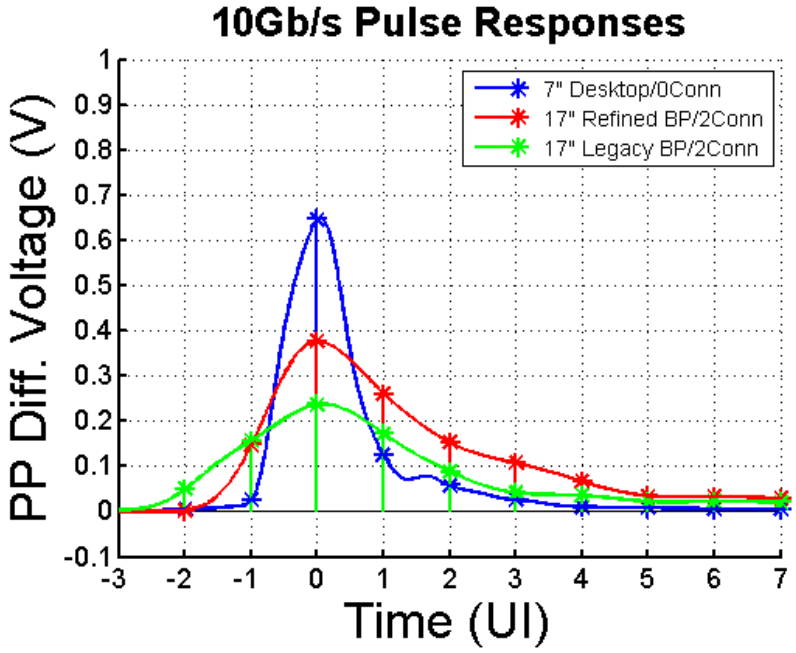

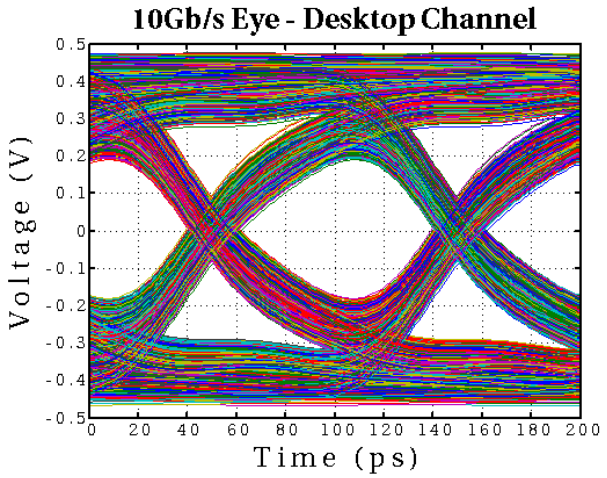

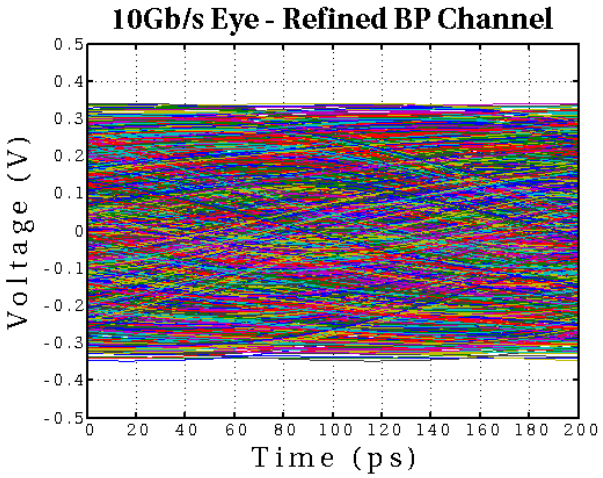

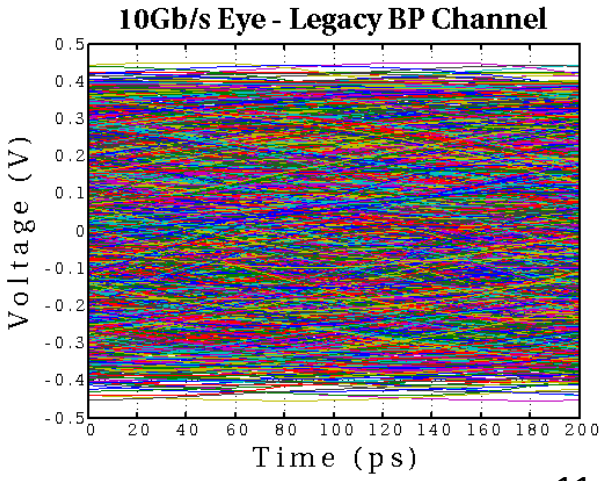

Outline

- Introduction
- Channel characteristics
- Equalizer circuits
- Equalizer adaptation techniques
- Optical interconnects
- Conclusion

Link with Equalization

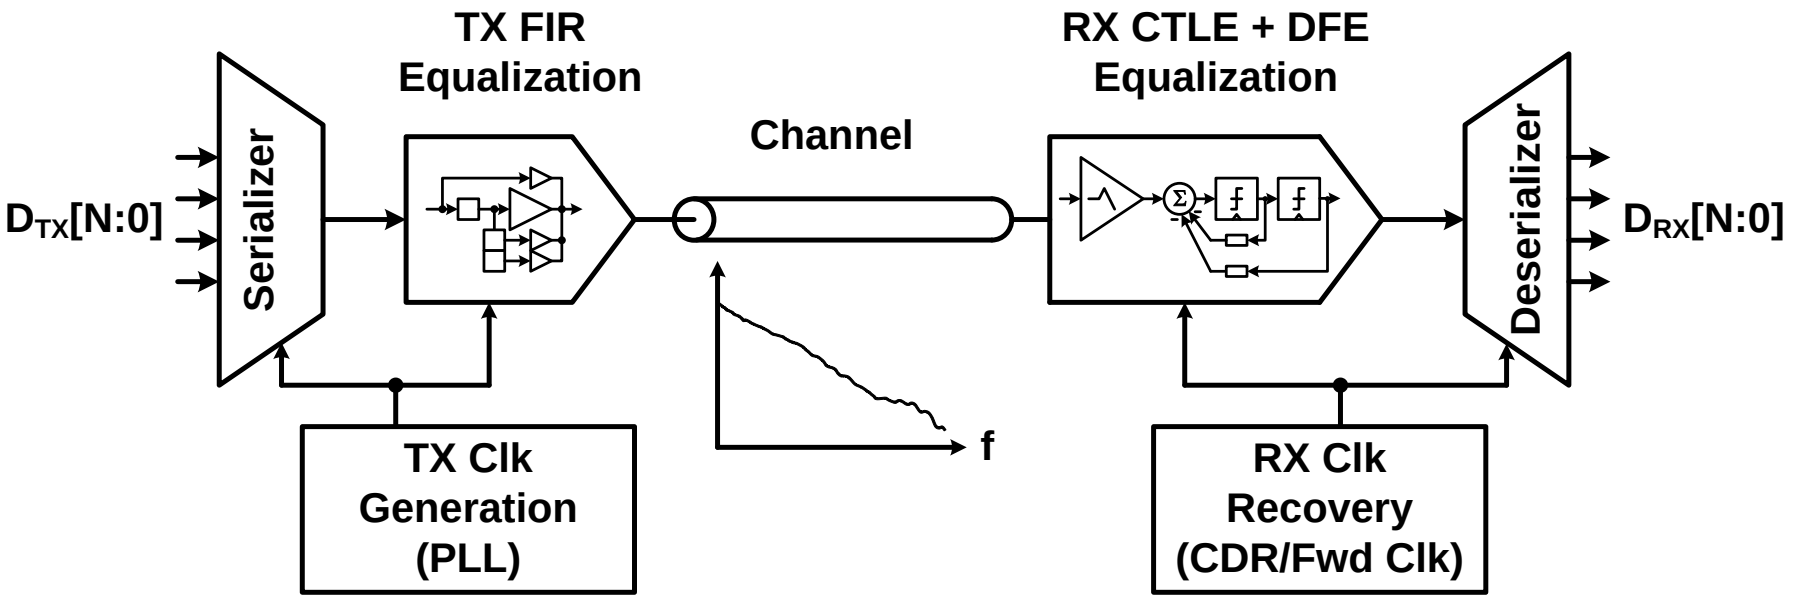

Channel Performance Impact

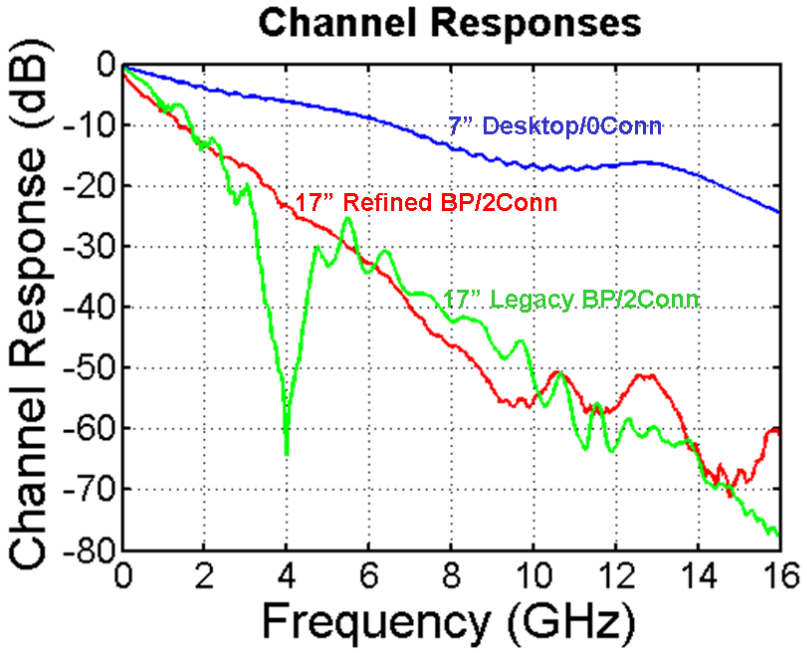

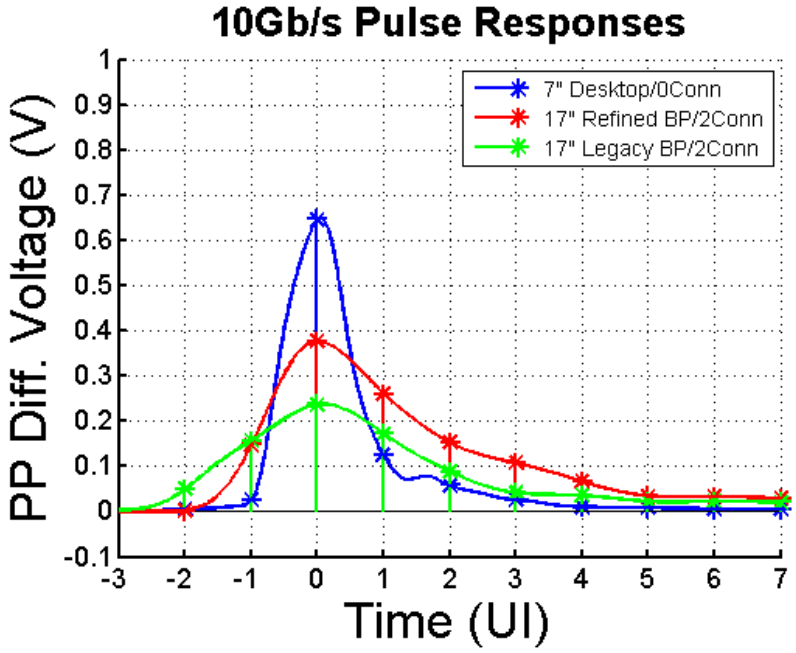

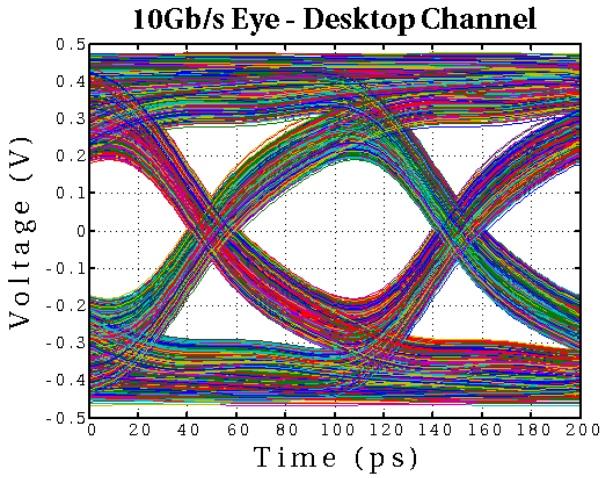

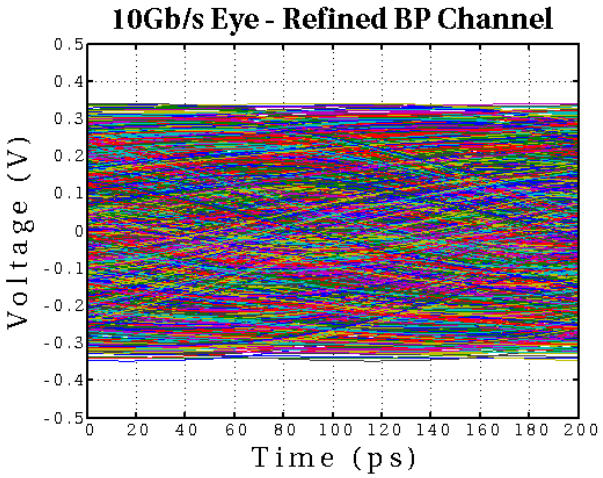

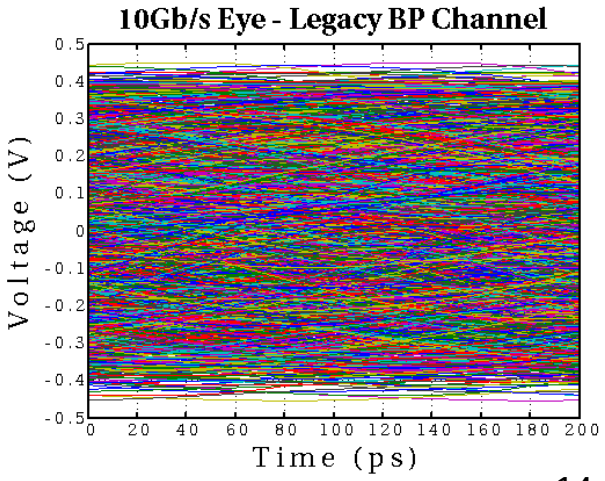

Channel Performance Impact

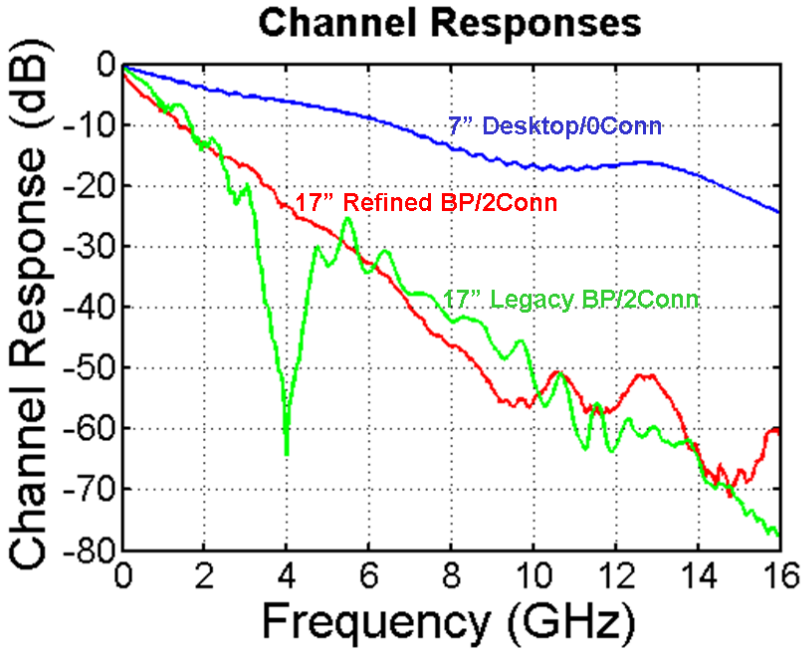

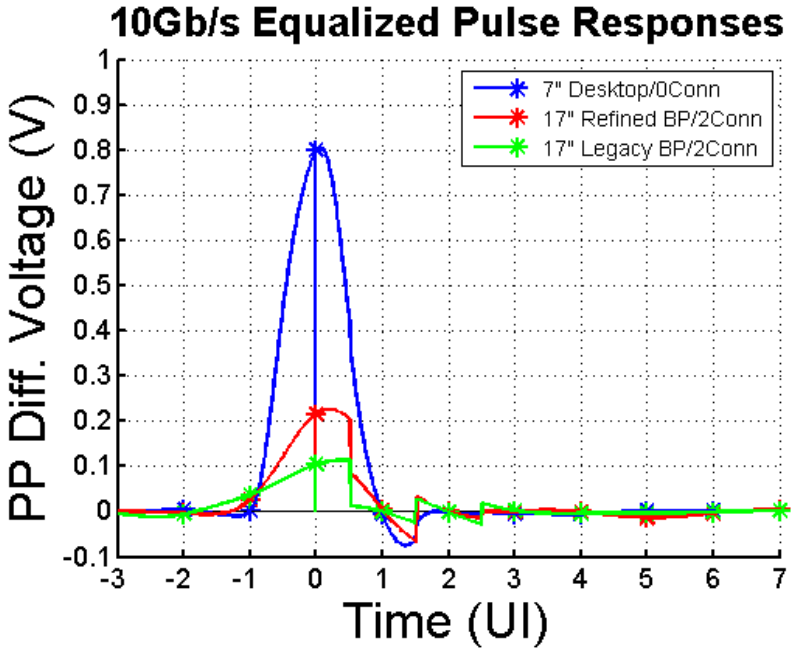

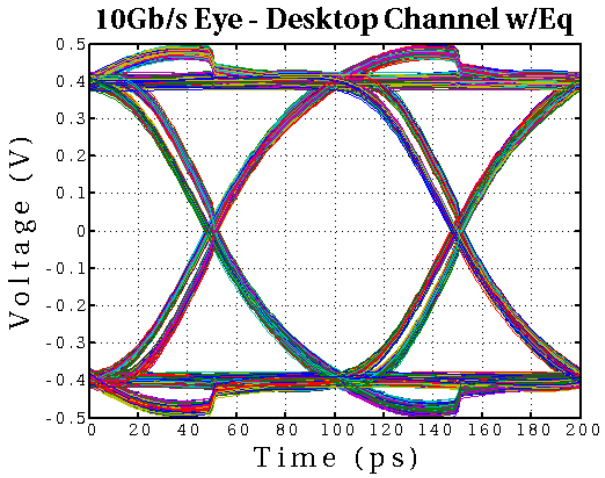

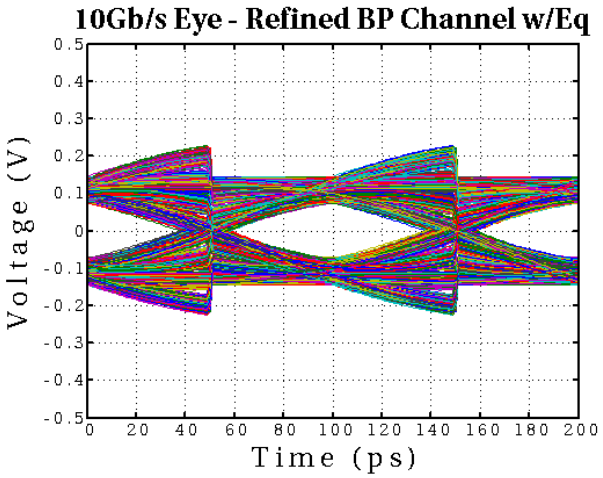

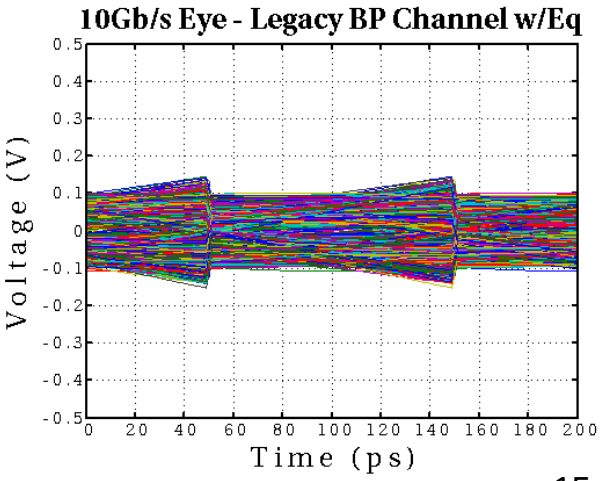

Channel Equalization

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

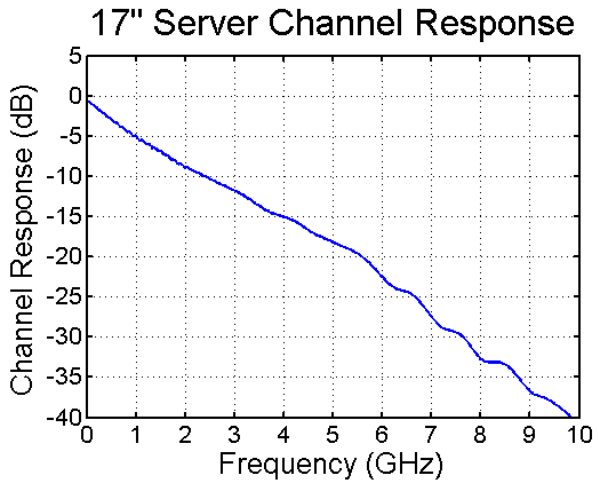

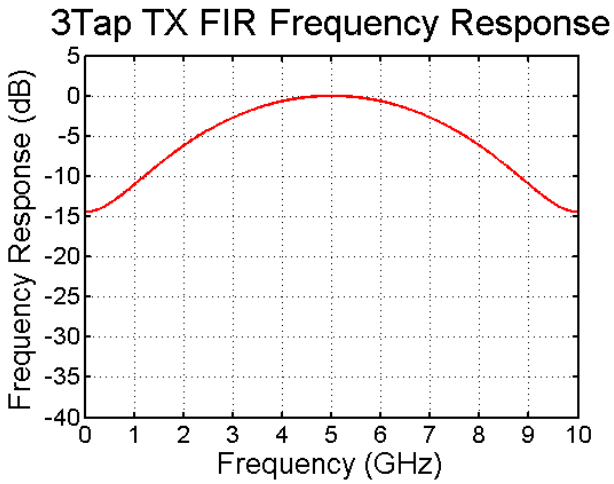

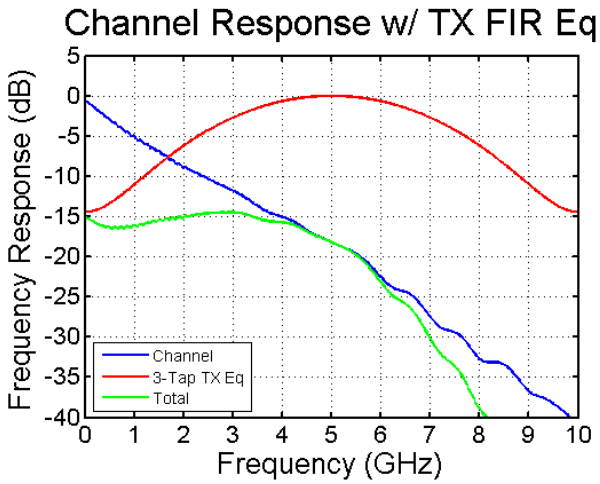

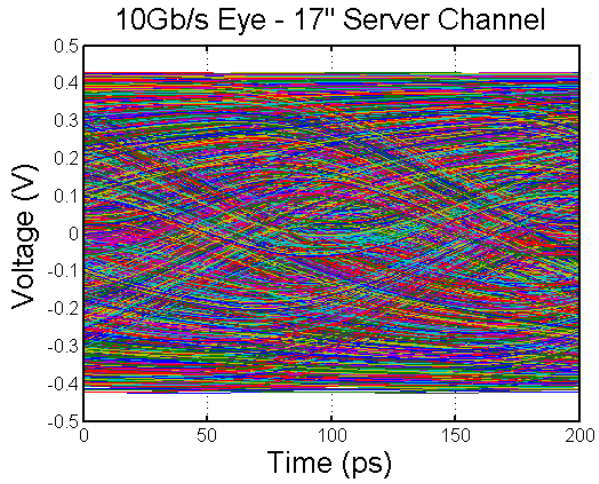

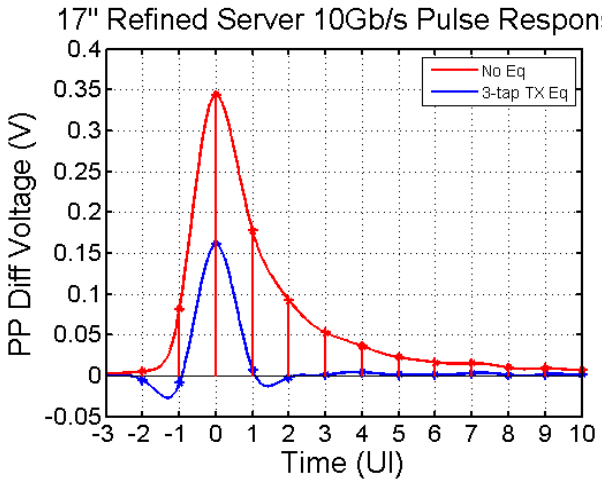

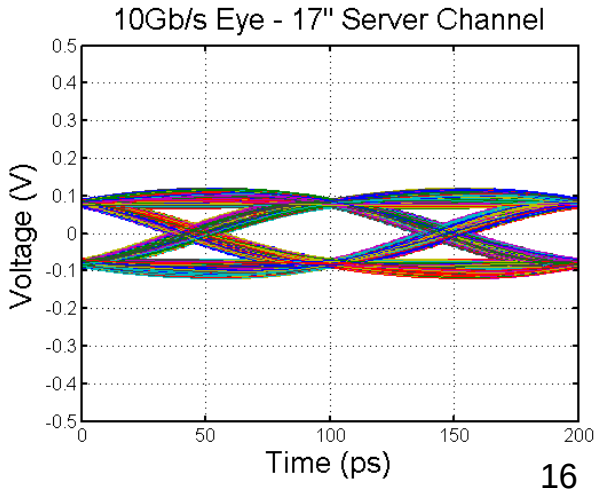

Link with Equalization

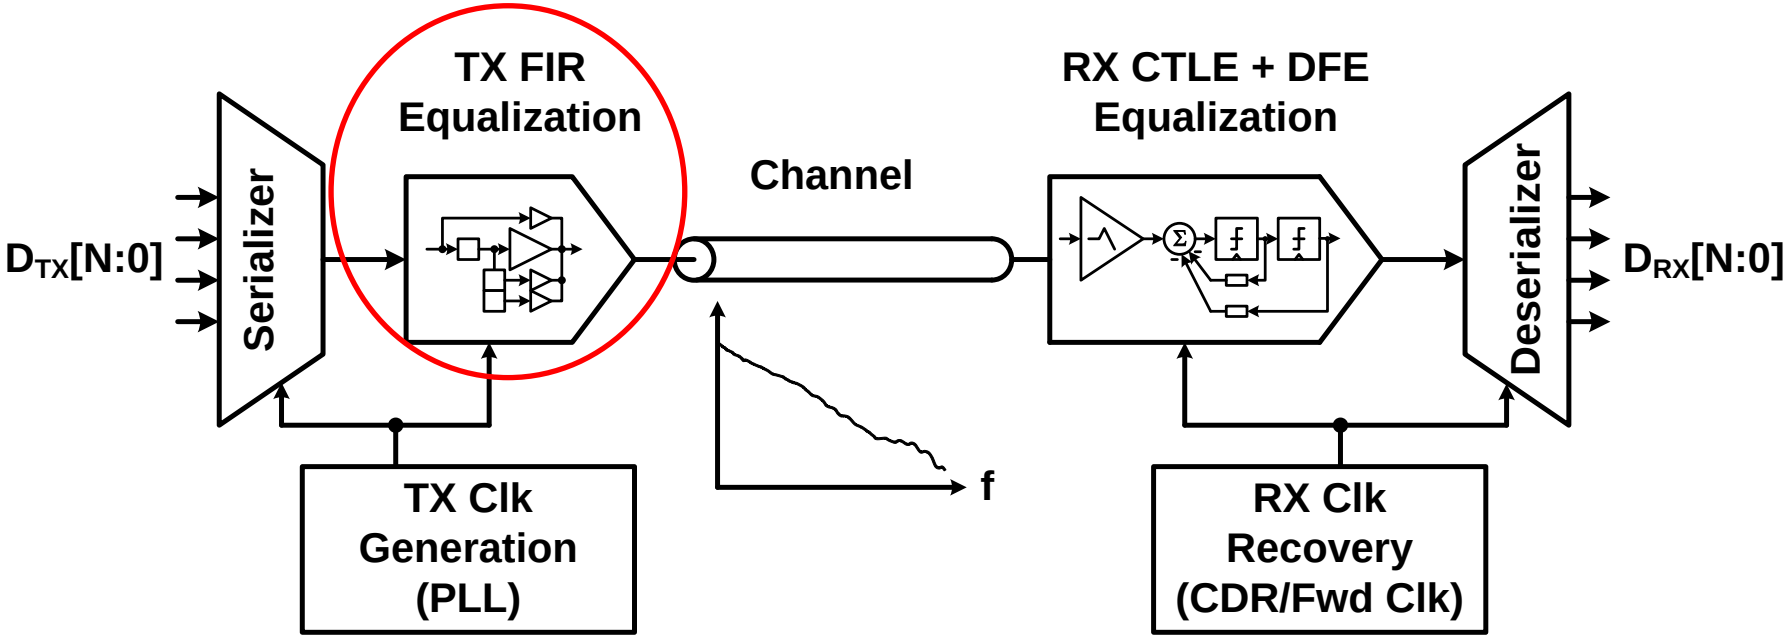

TX FIR Equalization

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

$$V_{out}(0) = \left[I_{-1}D(1) + I_0 D(0) + I_1 D(-1) + I_2 D(-2)\right]\left(\frac{R_{TERM}}{2}\right)$$

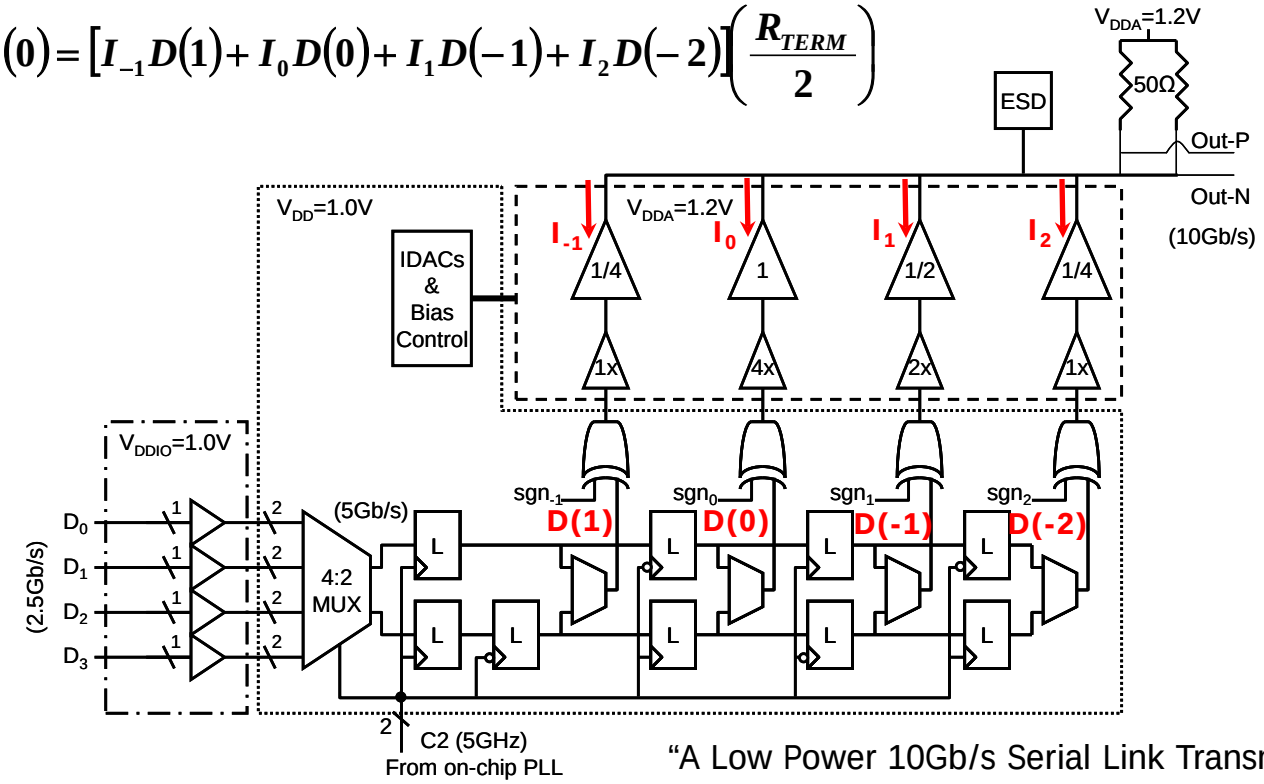

"A Low Power 10Gb/s Serial Link Transmitter in 90-nm CMOS," A. Rylyakov et al., CSICS 2005

6Gb/s TX FIR Equalization Example

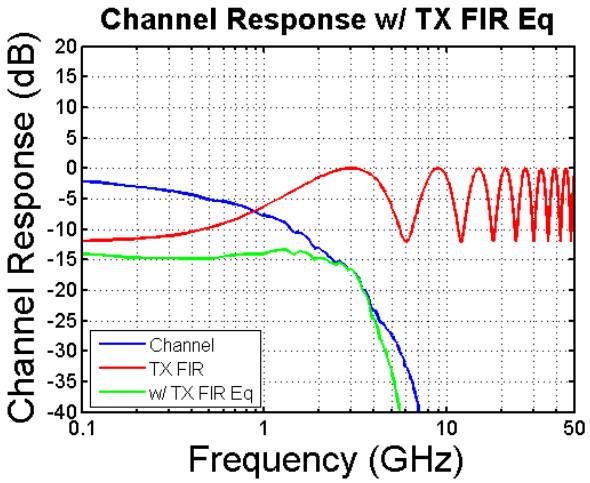

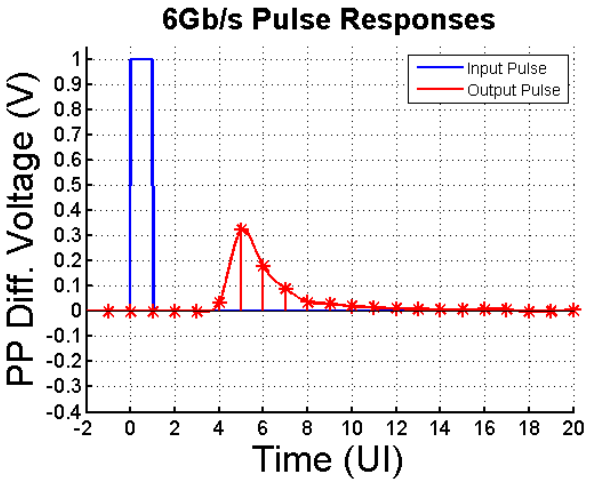

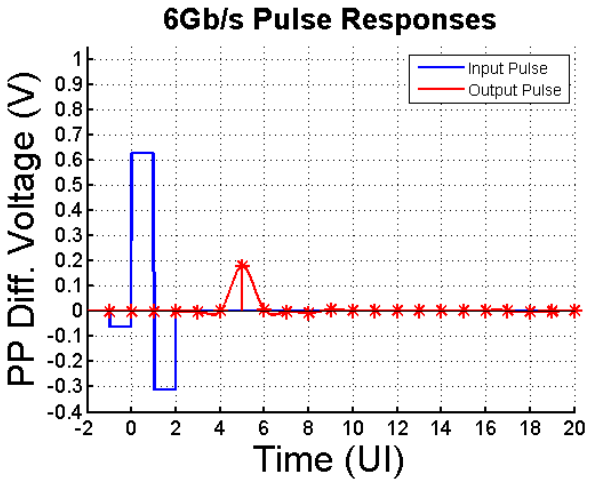

- Pros
  - Simple to implement
  - Can cancel ISI in pre-cursor and beyond filter span
  - Doesn't amplify noise
  - Can achieve 5-6bit resolution
- Cons
  - Attenuates low frequency content due to peak-power limitatior
  - Need a "back-channel" to tune filter taps

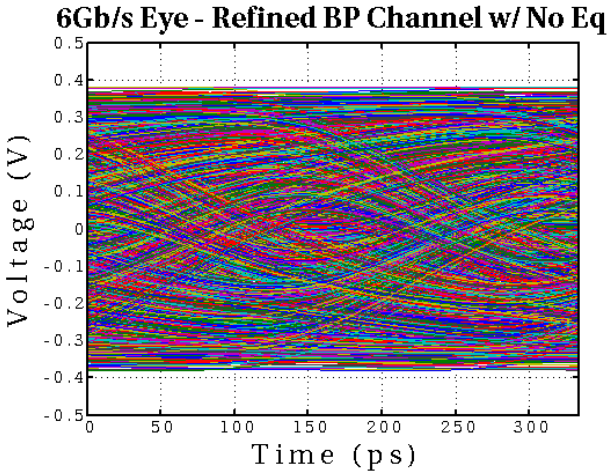

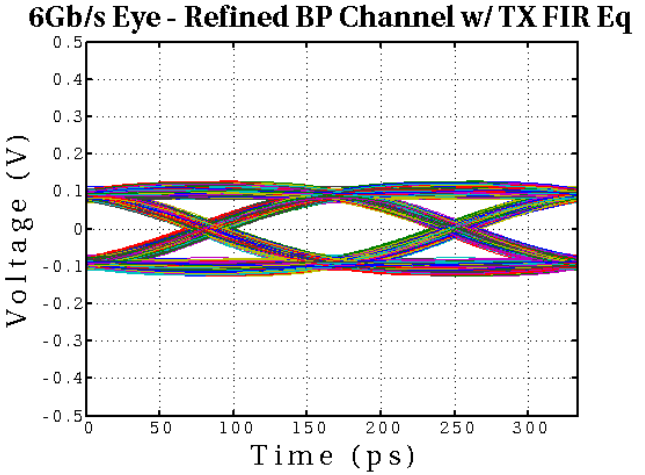

TX FIR Equalization – Time Domain

For 10Gbps : $W(z) = -0.131 + 0.595z^{-1} - 0.274z^{-2}$

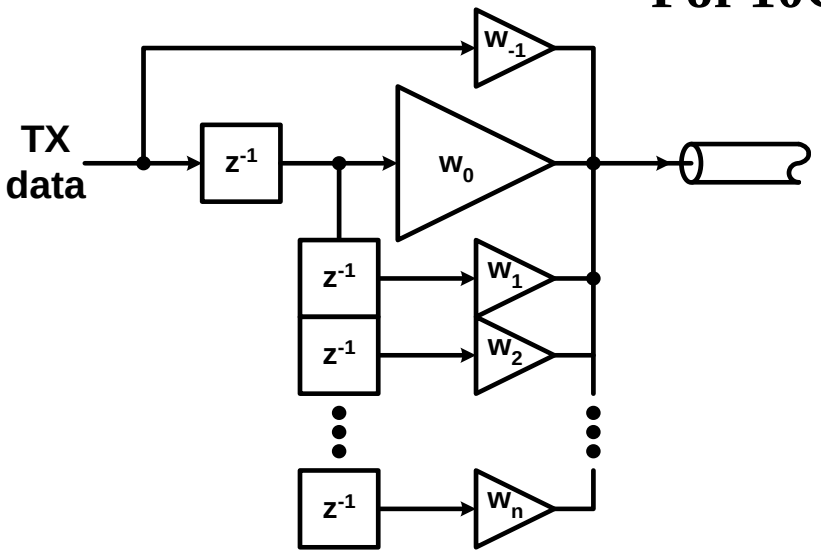

$$\mathbf{W} = \begin{bmatrix} -0.131 & 0.595 & -0.274 \end{bmatrix}$$

Low Frequency Response (Sum Taps)

$$\begin{bmatrix} \ldots & 1 & 1 & 1 & \ldots \end{bmatrix} * \begin{bmatrix} -0.131 & 0.595 & -0.274 \end{bmatrix} = \begin{bmatrix} \ldots & 0.190 & 0.190 & 0.190 & \ldots \end{bmatrix}$$

Nyquist Frequency Response (Sum Taps w/ Alternating Polarity)

$$\begin{bmatrix} \ldots & -1 & 1 & -1 & \ldots \end{bmatrix} * \begin{bmatrix} -0.131 & 0.595 & -0.274 \end{bmatrix} = \begin{bmatrix} \ldots & 1 & -1 & 1 & \ldots \end{bmatrix}$$

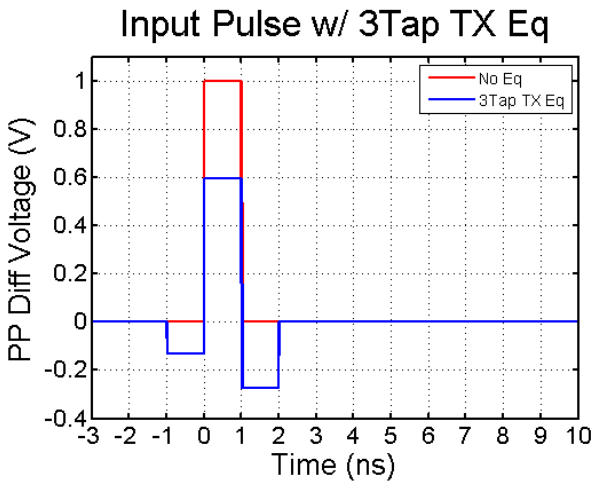

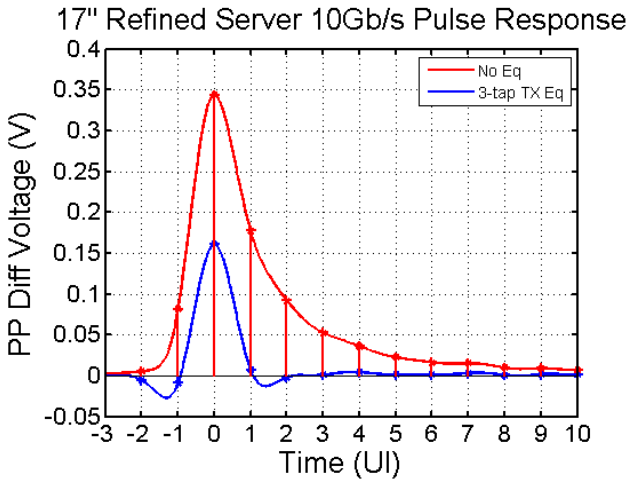

TX FIR Equalization – Freq. Domain

For 10Gbps : $W(z) = -0.131 + 0.595z^{-1} - 0.274z^{-2}$

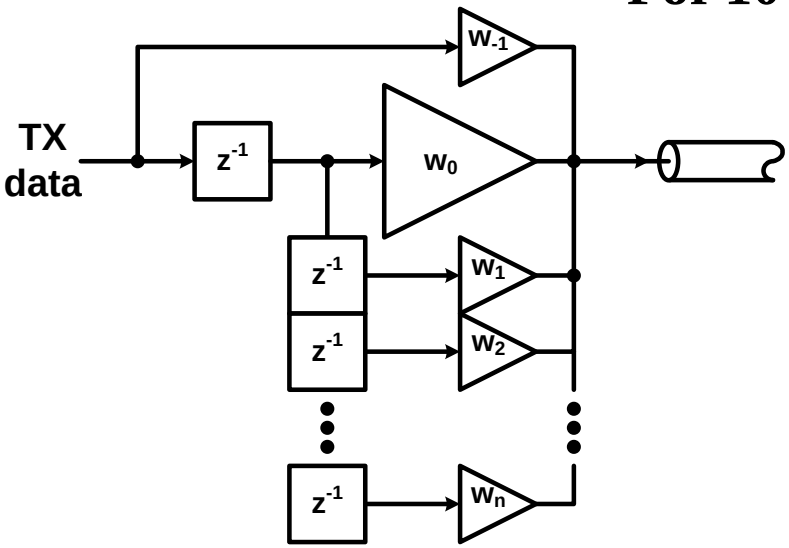

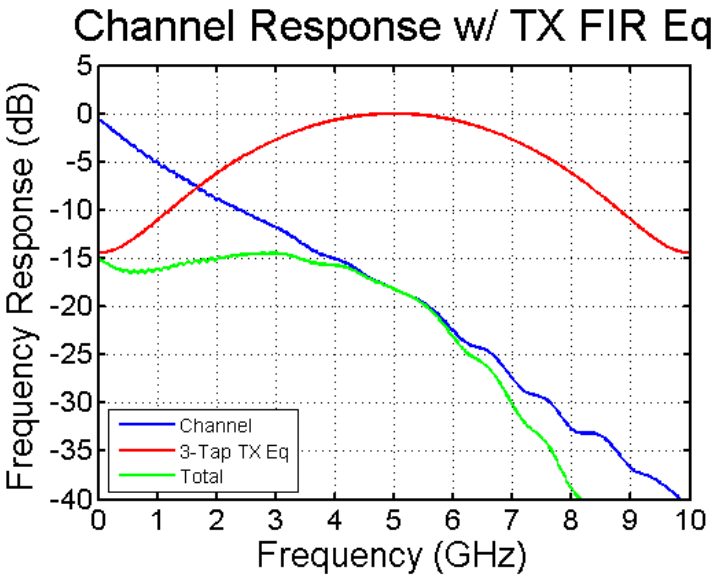

$$W(z) = -0.131 + 0.595z^{-1} - 0.274z^{-2}$$

w/ $z = e^{j2\pi fT_s} = \cos(2\pi fT_s) + j\sin(2\pi fT_s)$

Low Frequency Response $(f = 0)$

$$z = \cos(0) + j\sin(0) = 1 \Rightarrow W(f = 0) = 0.190 \Rightarrow -14.4dB$$

Nyquist Frequency Response $\left(f = \frac{1}{2T_s}\right)$

$$z = \cos(\pi) + j\sin(\pi) = -1 \Rightarrow W\left(f = \frac{1}{2T_s}\right) = -1 \Rightarrow 0dB$$

- Equalizer has 14.4dB of frequency peaking
  - Attenuates DC at -14.4dB and passes Nyquist frequency at 0dB

Note: Ts=Tb=100ps

TX FIR Circuit Architectures

- Direct FIR vs Segmented DAC
- Direct FIR
  - Parallel output drivers for output taps
  - Each parallel driver must be sized to handle its potential maximum current
  - Lower power & complexity
  - Higher output capacitance
- Segmented DAC
  - Minimum sized output transistors to handle peak output current
  - Lowest output capacitance
  - Most power & complexity
    - Need mapping table (RAM)
    - Very flexible in equalization

Direct FIR

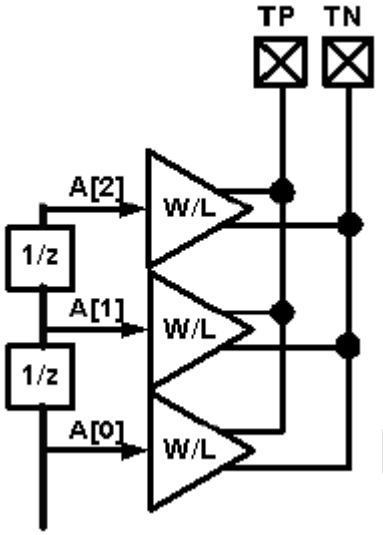

[Zerbe]

Segmented DAC

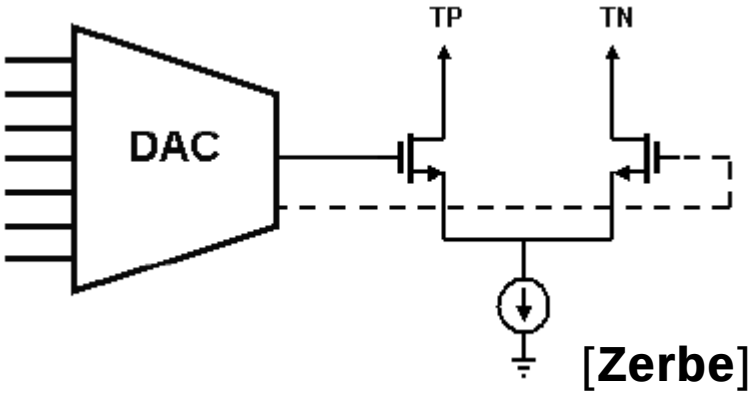

[Zerbe]

Direct FIR Equalization

$$V_{out}(0) = \left[I_{-1}D(1) + I_0 D(0) + I_1 D(-1) + I_2 D(-2)\right]\left(\frac{R_{TERM}}{2}\right)$$

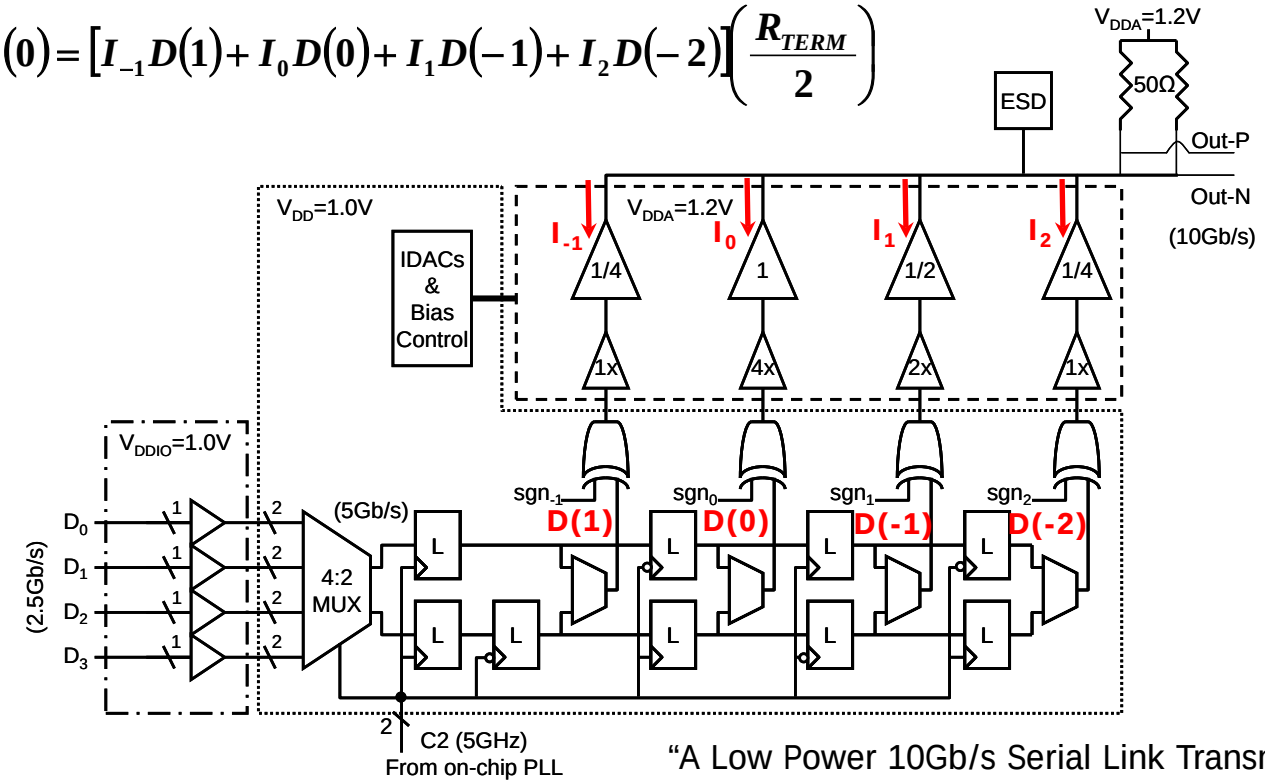

"A Low Power 10Gb/s Serial Link Transmitter in 90-nm CMOS," A. Rylyakov et al., CSICS 2005

Segmented DAC Example

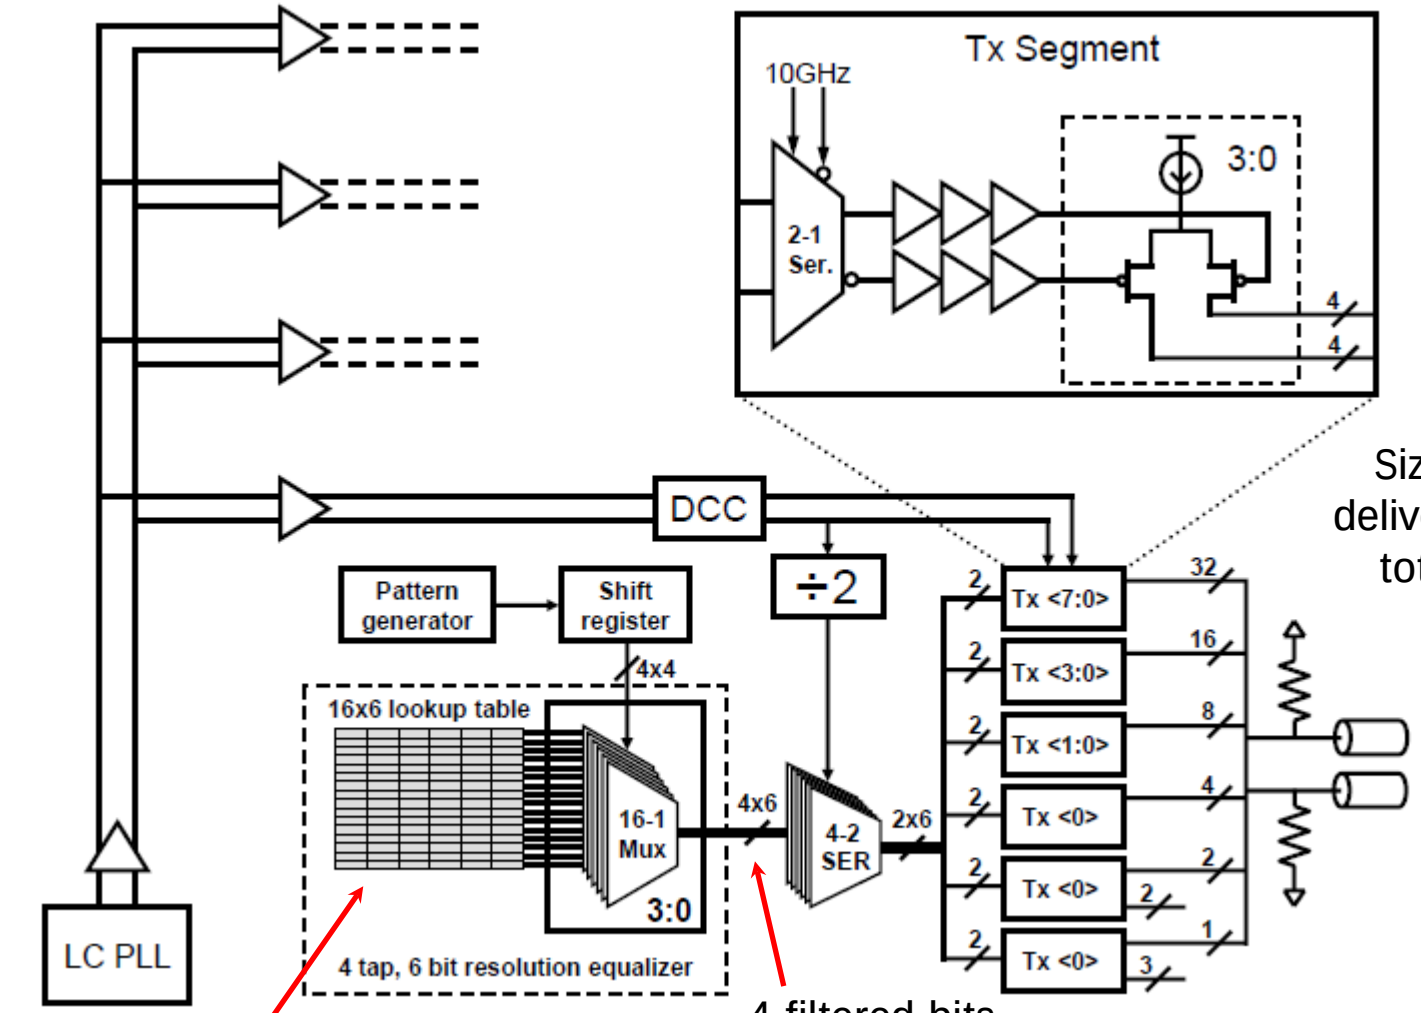

Sized only to deliver maximum total current

Row = 4-bit data pattern
Column = 6-bit weighting

4 filtered bits (parallel) at 6-bit resolution

[**Casper ISSCC 2006**]

Link with Equalization

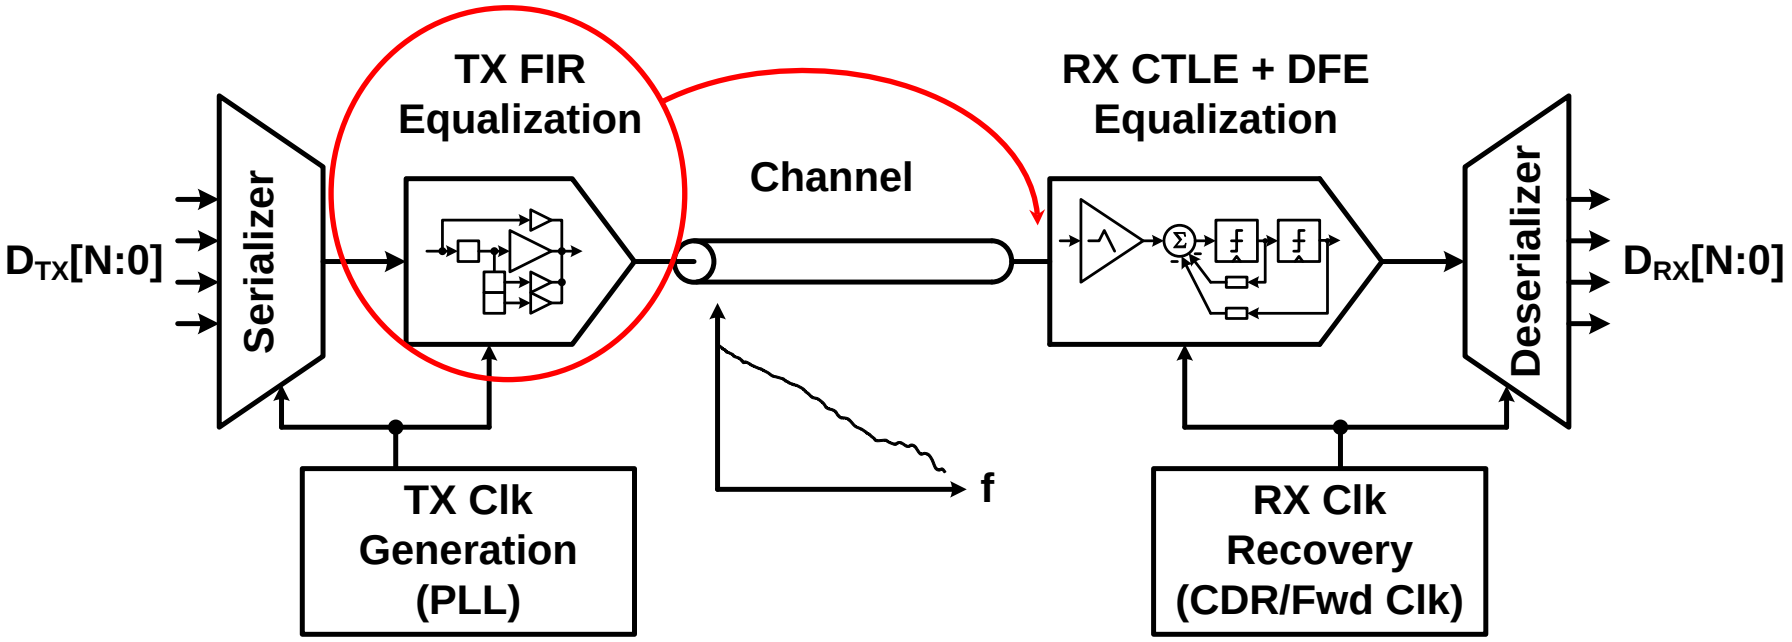

RX FIR Equalization

- Delay analog input signal and multiply by equalization coefficients
- 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

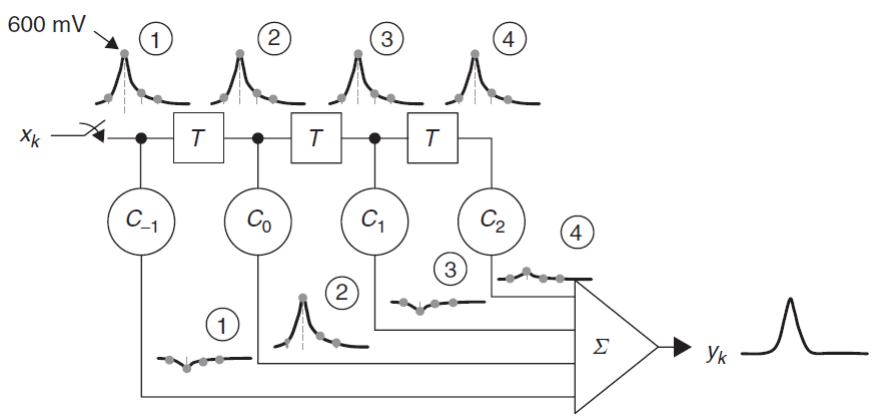

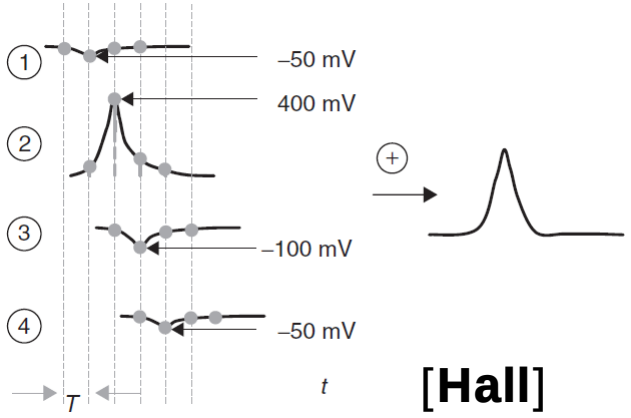

[**Hall**]

RX Equalization Noise Enhancement

- Linear RX equalizers don't discriminate between signal, noise, and cross-talk
  - While signal-to-distortion (ISI) ratio is improved, SNR remains unchanged

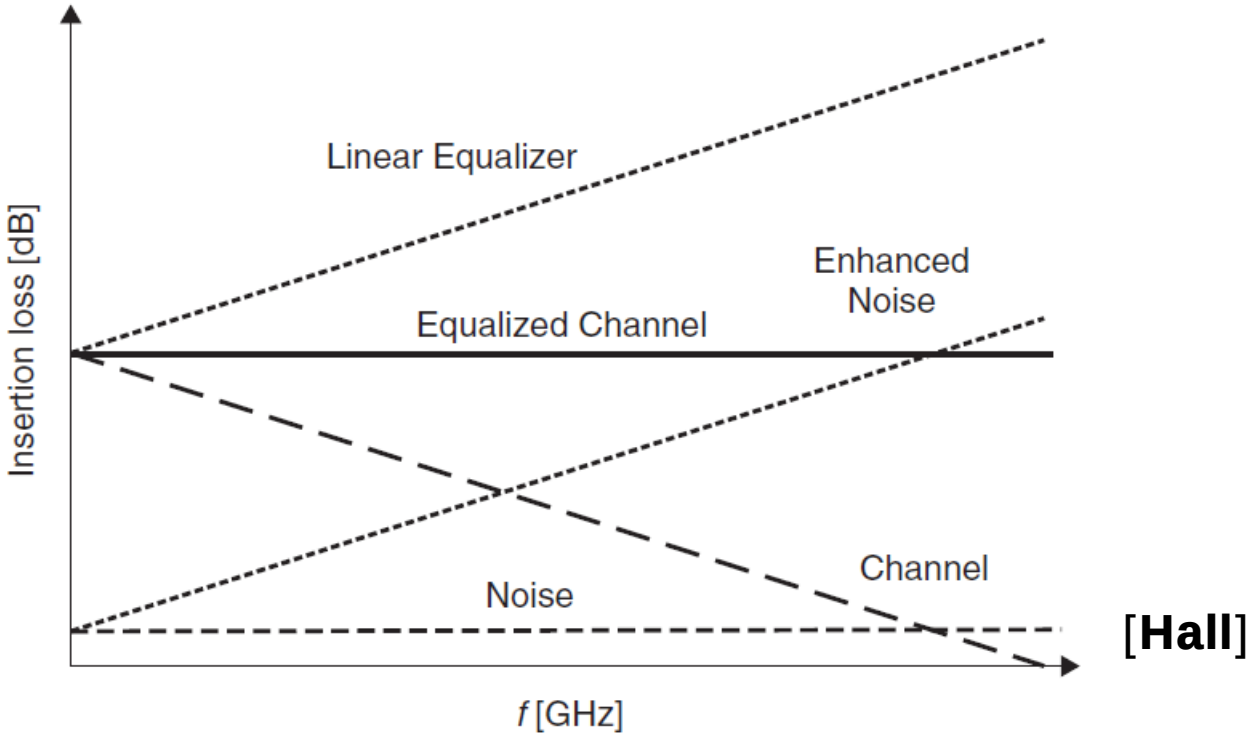

Analog RX FIR Equalization Example

- 5-tap equalizer with tap spacing of $T_b/2$

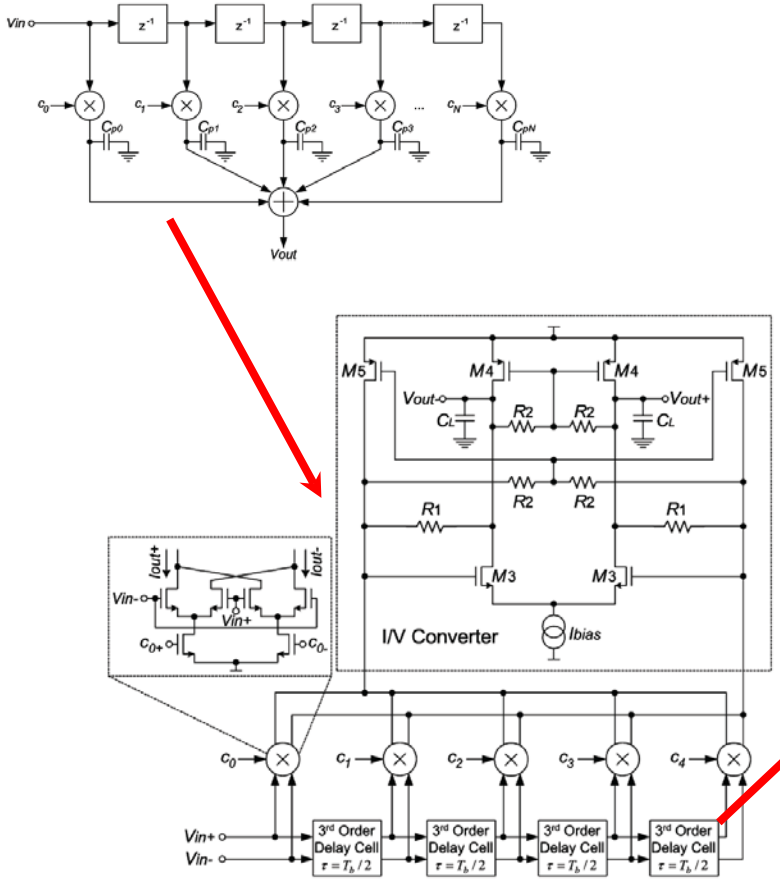

3rd-order delay cell

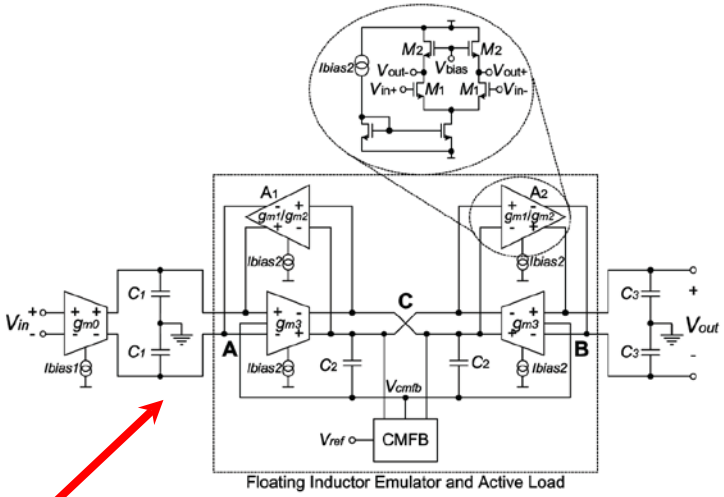

Floating Inductor Emulator and Active Load

1Gb/s experimental results

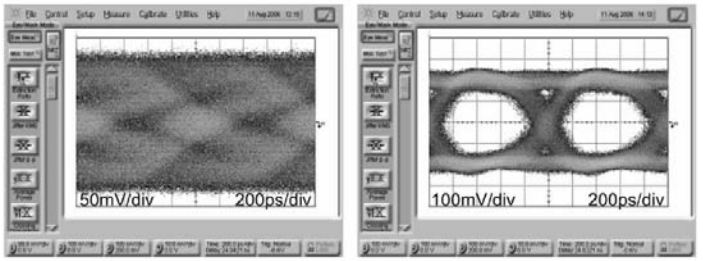

Before Equalizer: 23meters

After Equalizer: 23meters

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

Digital RX FIR Equalization

- Digitize the input signal with high-speed low/medium resolution ADC and perform equalization in digital domain
  - Digital delays, multipliers, adders
  - Limited to ADC resolution
- Power can be high due to very fast ADC

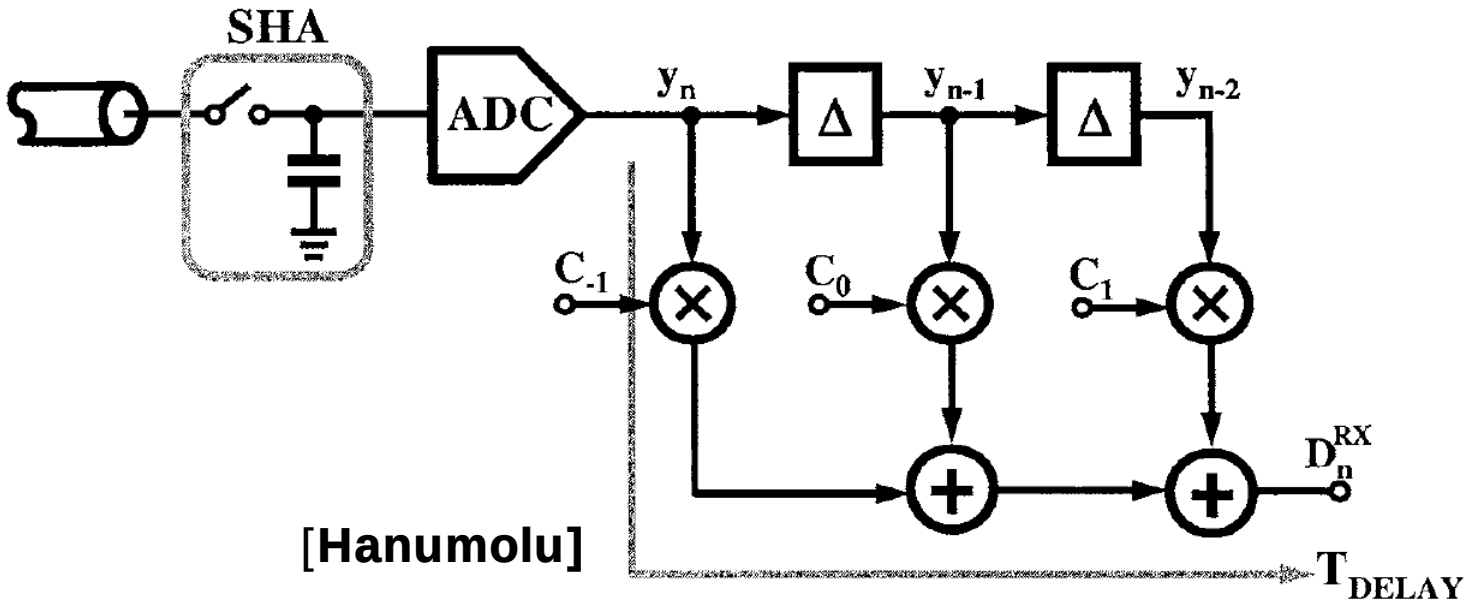

[**Hanumolu**]

Digital RX FIR Equalization Example

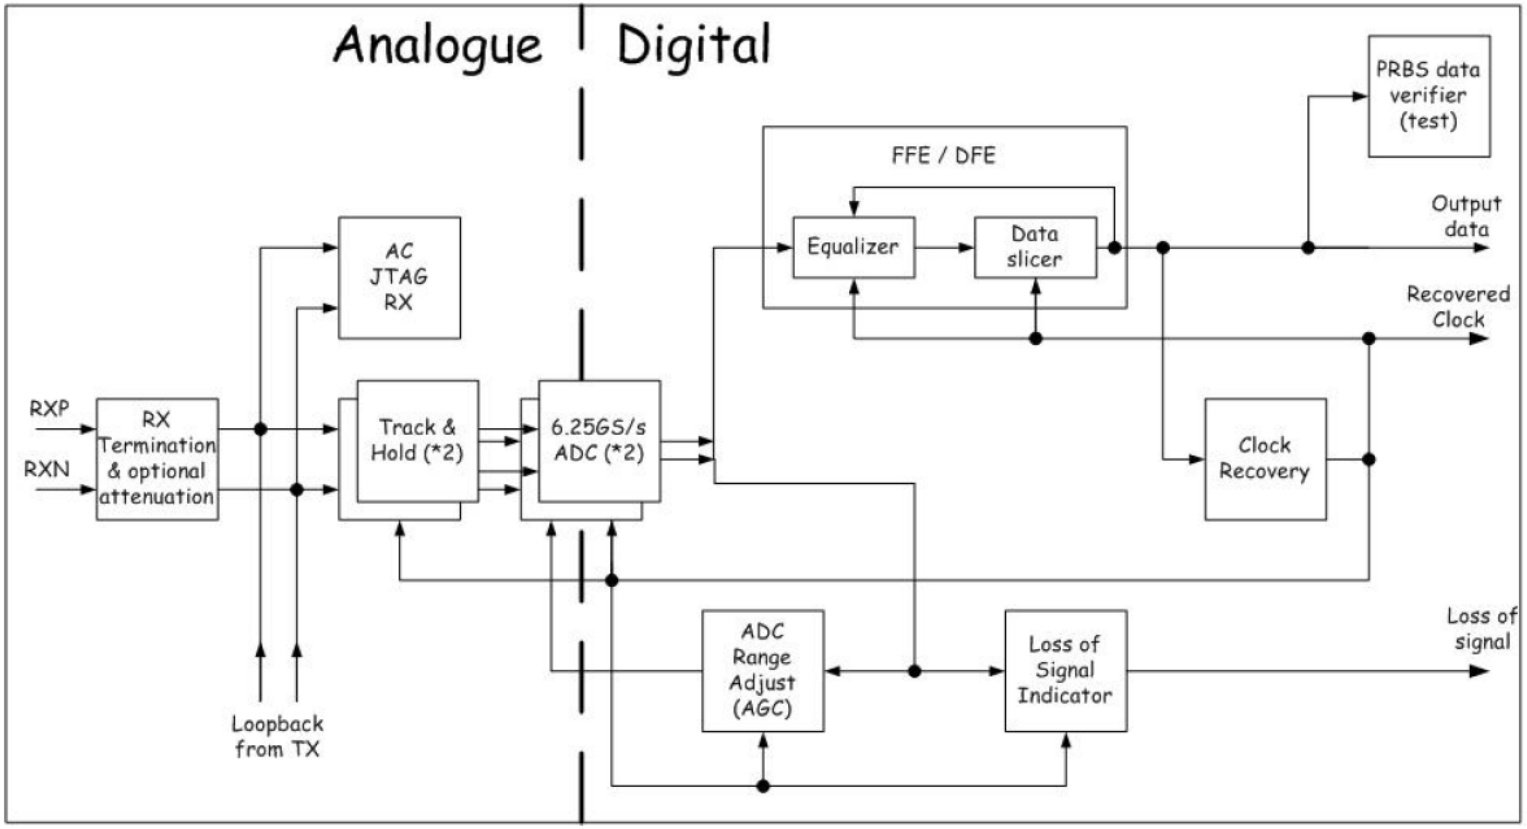

- 12.5GS/s 4.5-bit Flash ADC in 65nm CMOS [**Harwood ISSCC 2007**]
- 2-tap FFE & 5-tap DFE
- XCVR power (inc. TX) = 330mW, Analog = 245mW, Digital = 85mW

Link with Equalization

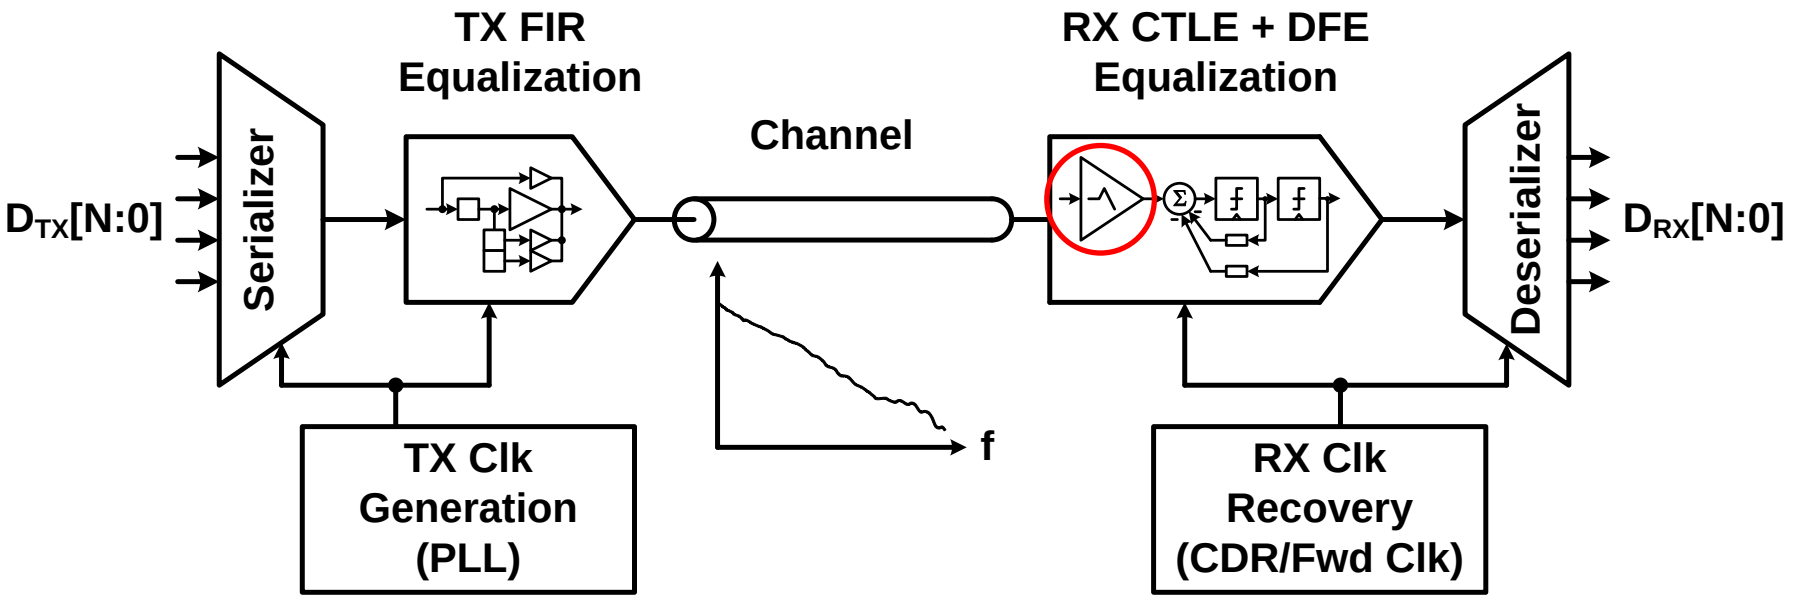

RX CTLE Equalization

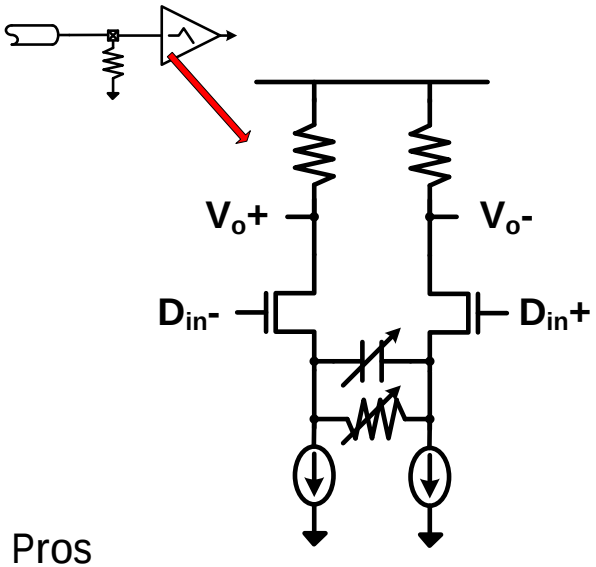

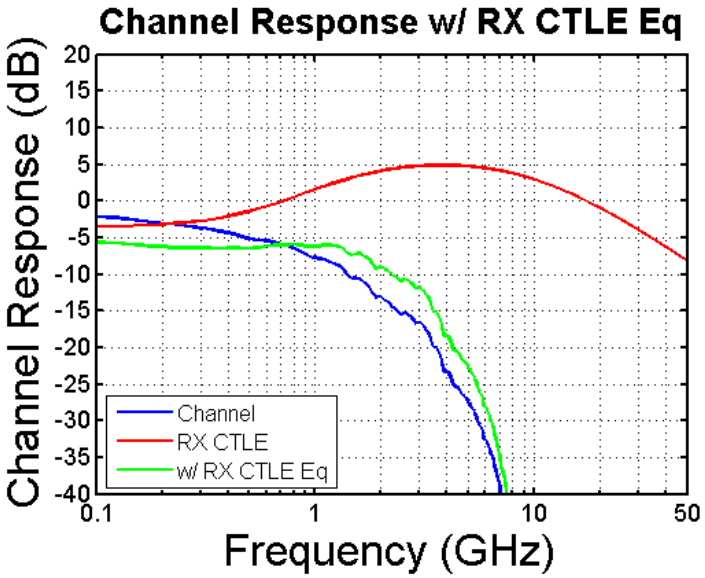

- Pros
  - Provides gain and equalization with low power and area overhead
  - Can cancel both pre-cursor and long-tail ISI
- Cons
  - Generally limited to 1st order compensation
  - Amplifies noise/crosstalk
  - PVT sensitivity
  - Can be hard to tune

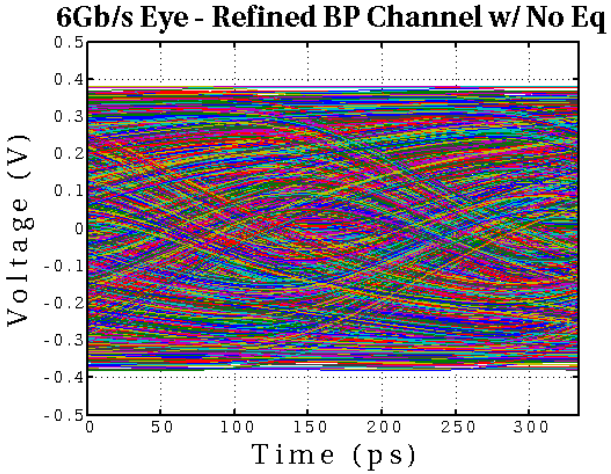

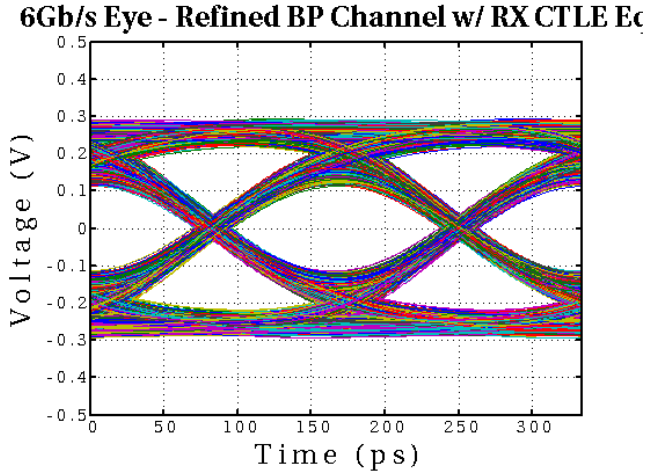

Passive CTLE

- Passive structures offer excellent linearity, but no gain at Nyquist frequency

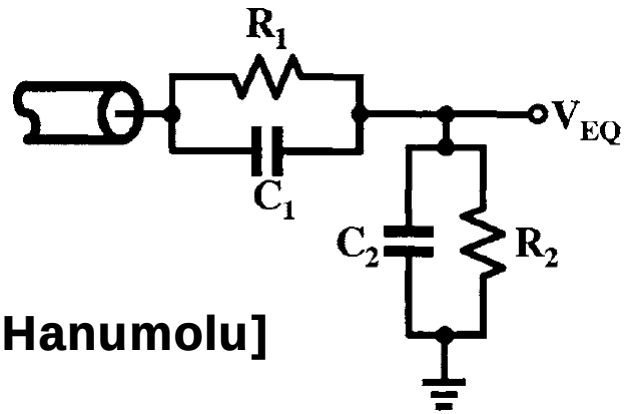

[**Hanumolu**]

$$H(s)=\frac{R_2}{R_1+R_2}\frac{1+R_1C_1s}{1+\frac{R_1R_2}{R_1+R_2}(C_1+C_2)s}$$

$$\omega_z=\frac{1}{R_1C_1},\qquad \omega_p=\frac{1}{\frac{R_1R_2}{R_1+R_2}(C_1+C_2)}$$

$$\textbf{DC gain}=\frac{R_2}{R_1+R_2},\ \textbf{HF gain}=\frac{C_1}{C_1+C_2}$$

$$\textbf{Peaking}=\frac{\textbf{HF gain}}{\textbf{DC gain}}=\frac{\omega_p}{\omega_z}=\frac{R_1+R_2}{R_2}\frac{C_1}{C_1+C_2}$$

Active CTLE

- Input amplifier with RC degeneration can provide frequency peaking with gain at Nyquist frequency
- Potentially limited by gain-bandwidth of amplifier
- Amplifier must be designed for input linear range
  - Often TX eq. provides some low frequency attenuation
- Sensitive to PVT variations and can be hard to tune
- Generally limited to 1st-order compensation

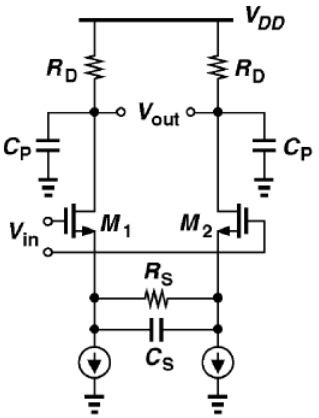

[Gondi JSSC 2007]

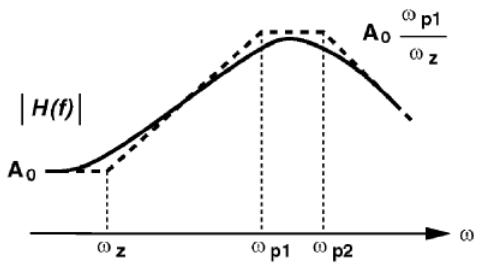

$$H(s) = \frac{g_m}{C_p} \frac{s + \frac{1}{R_S C_S}}{\left(s + \frac{1 + g_m R_S/2}{R_S C_S}\right)\left(s + \frac{1}{R_D C_p}\right)}$$

$$\omega_z = \frac{1}{R_S C_S}, \quad \omega_{p1} = \frac{1 + g_m R_S/2}{R_S C_S}, \quad \omega_{p2} = \frac{1}{R_D C_p}$$

$$\textbf{DC gain} = \frac{g_m R_D}{1 + g_m R_S/2}, \quad \textbf{Ideal peak gain} = g_m R_D$$

$$\textbf{Ideal Peaking} = \frac{\textbf{Ideal peak gain}}{\textbf{DC gain}} = \frac{\omega_{p1}}{\omega_z} = 1 + g_m R_S/2$$

Active CTLE Tuning

- Tune degeneration resistor and capacitor to adjust zero frequency and 1st pole which sets peaking and DC gain

- Increasing C_S moves zero and 1st pole to a lower frequency w/o impacting (ideal) peaking

- Increasing R_S moves zero to lower frequency and increases peaking (lowers DC gain)
  - Minimal impact on 1st pole

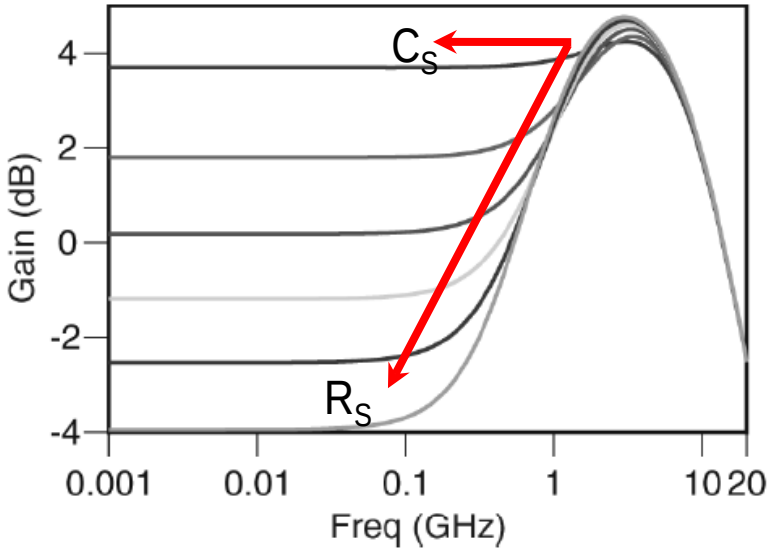

$$\omega_z = \frac{1}{R_S C_S}, \quad \omega_{p1} = \frac{1 + g_m R_S/2}{R_S C_S}$$

Link with Equalization

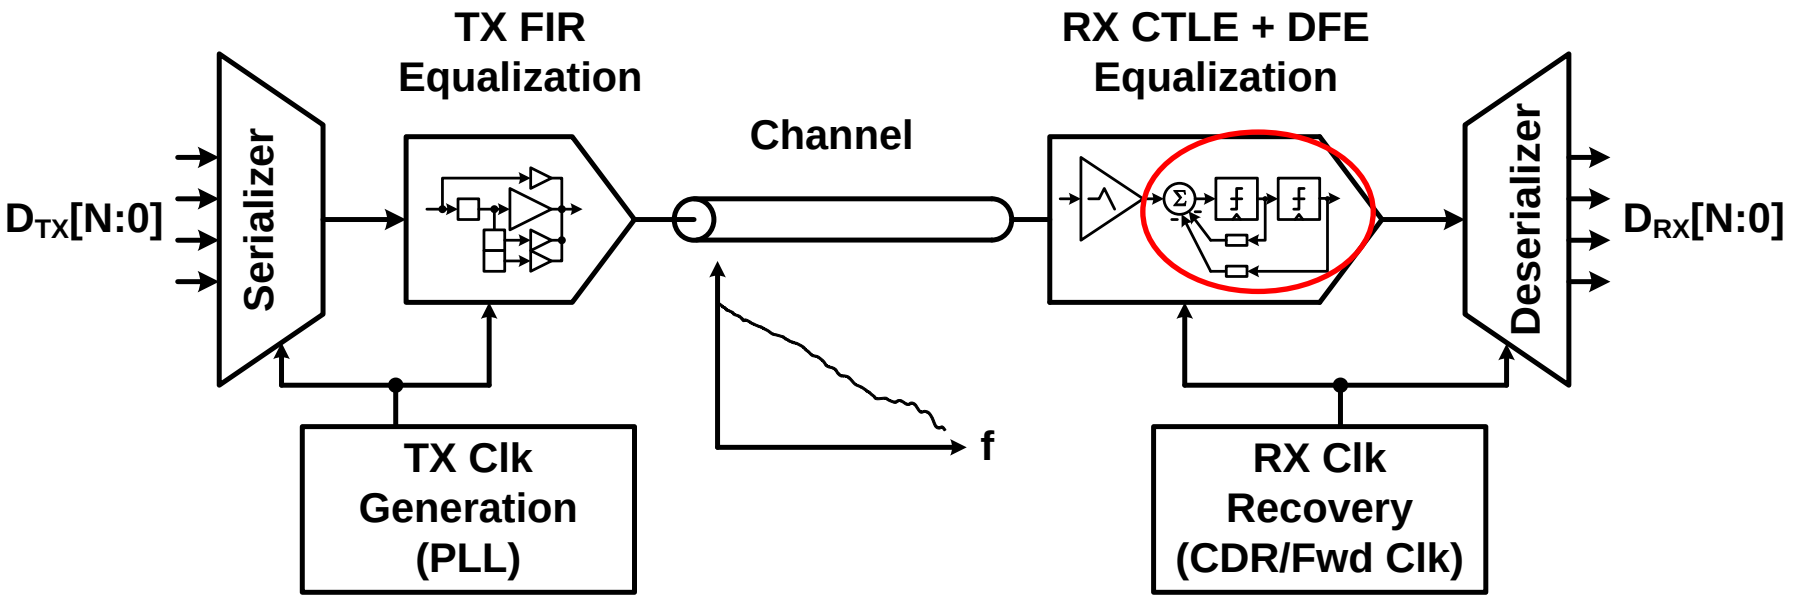

RX Decision Feedback Equalization (DFE)

- DFE is a **non-linear** equalizer
- Slicer makes a **symbol decision**, i.e. quantizes input
- ISI is then directly subtracted from the incoming signal via a feedback FIR filter

$$z_k = y_k - w_1 \tilde{d}_{k-1} \cdots - w_{n-1} \tilde{d}_{k-(n-1)} - w_n \tilde{d}_{k-n}$$

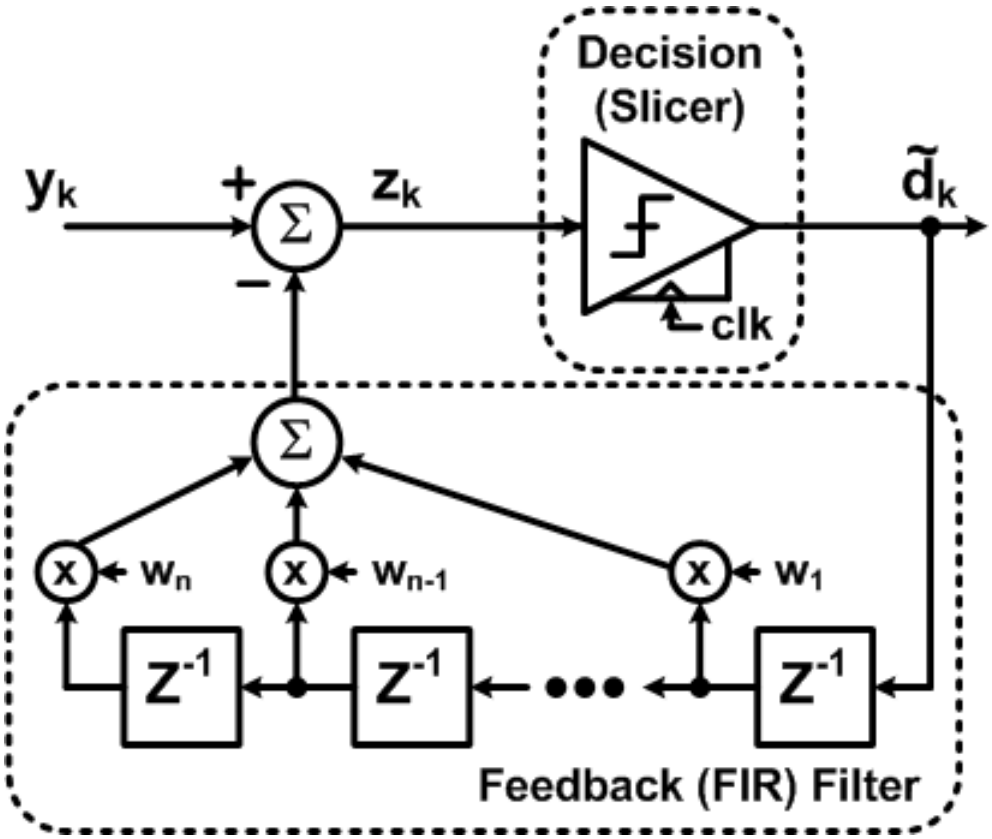

RX Decision Feedback Equalization (DFE)

- Pros
  - Can boost high frequency content without noise and crosstalk amplification
  - Filter tap coefficients can be adaptively tuned without any back-channel
- Cons
  - Cannot cancel pre-cursor ISI
  - Chance for error propagation
    - Low in practical links (BER=10^{-12})
  - Critical feedback timing path
  - Timing of ISI subtraction complicates CDR phase detection

$$z_k = y_k - w_1 \tilde{d}_{k-1} \cdots - w_{n-1} \tilde{d}_{k-(n-1)} - w_n \tilde{d}_{k-n}$$

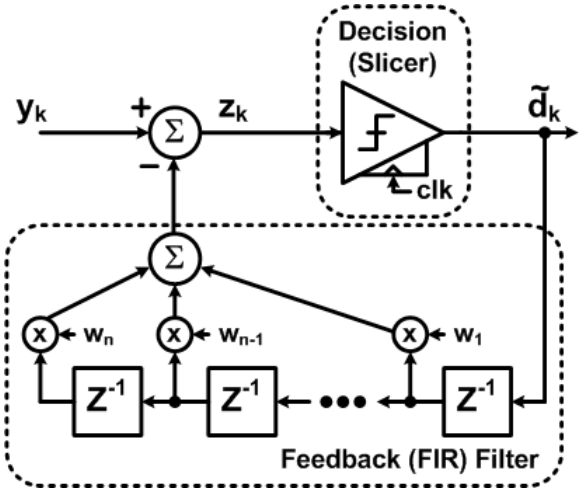

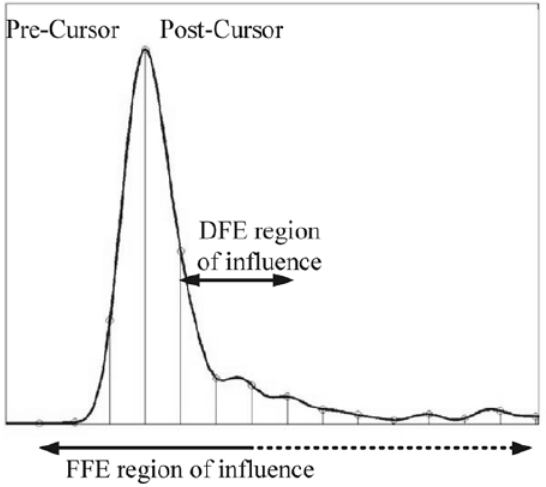

[**Payne**]

DFE Example

- If only DFE equalization, DFE tap coefficients should equal the unequalized channel pulse response values [a_1 a_2 ... a_n]
- With other equalization, DFE tap coefficients should equal the pre-DFE pulse response values

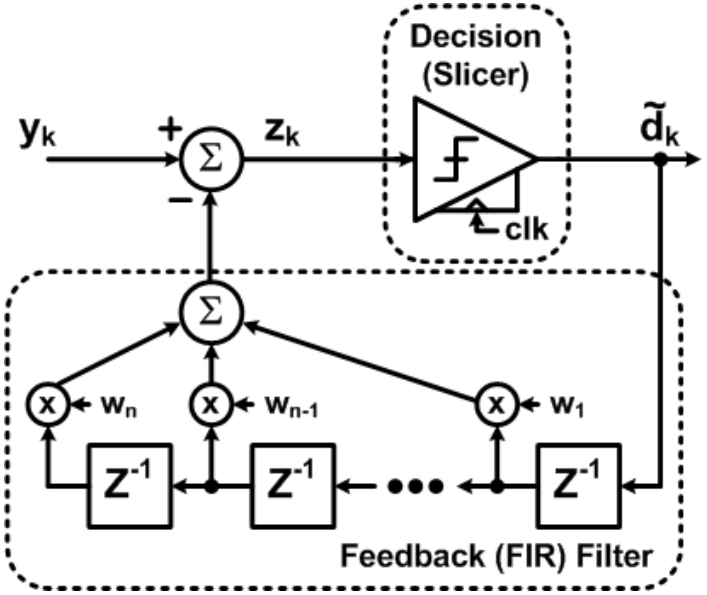

$[w_1\ w_2] = [a_1\ a_2]$

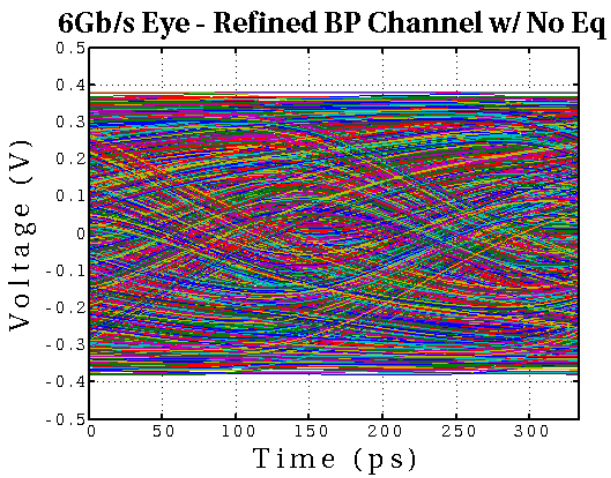

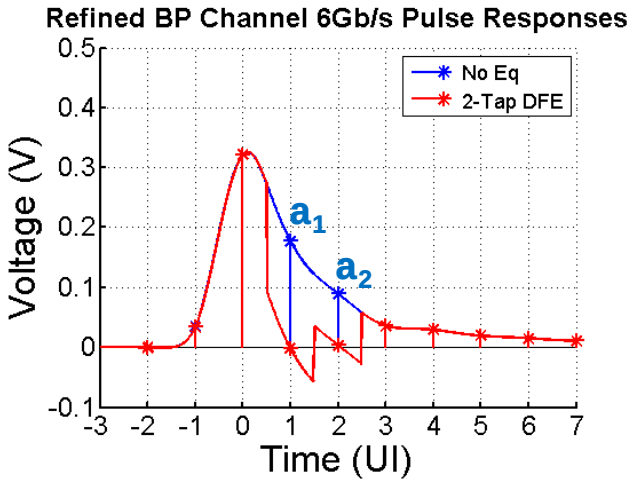

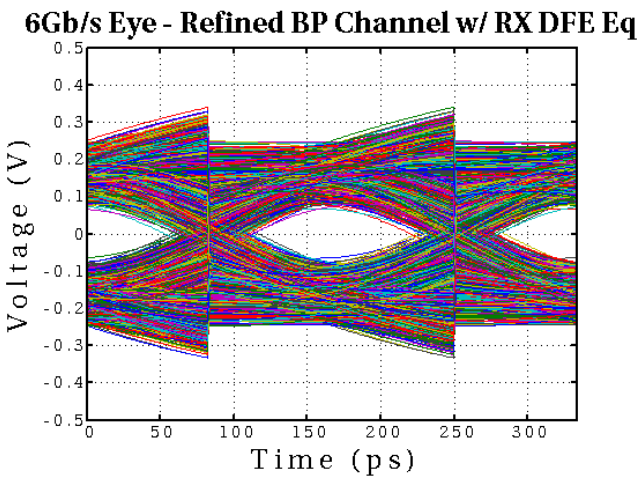

Direct Feedback DFE Example (TI)

- 6.25Gb/s 4-tap DFE
  - ½ rate architecture
  - Adaptive tap algorithm
  - Closes timing on 1st tap in **½ UI** for convergence of both adaptive equalization tap values and CDR

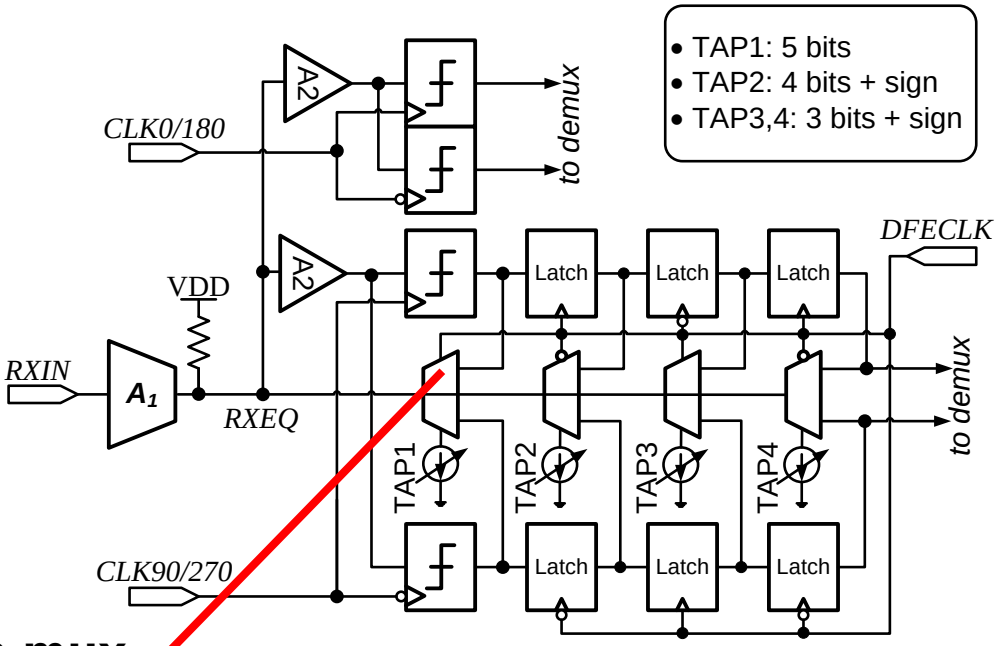

Feedback tap mux

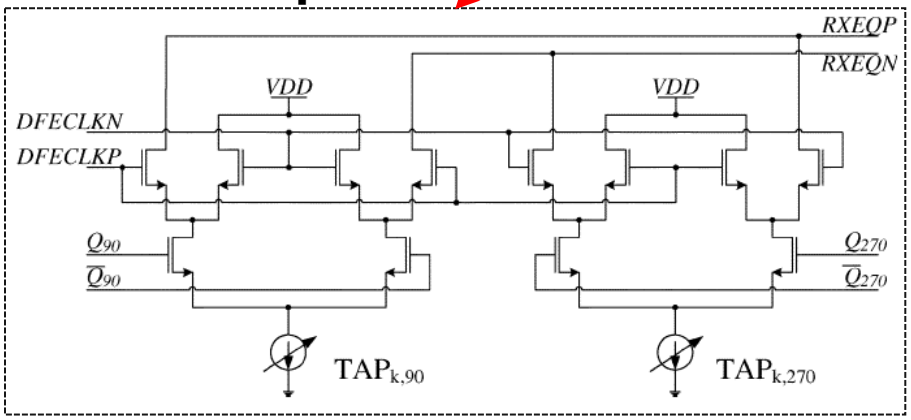

R. Payne *et al,* "A 6.25-Gb/s Binary Transceiver in 0.13-um CMOS for Serial Data Transmission Across High Loss Legacy Backplane Channels," *JSSC*, vol. 40, no. 12, Dec. 2005, pp. 2646-2657

Direct Feedback DFE Critical Path

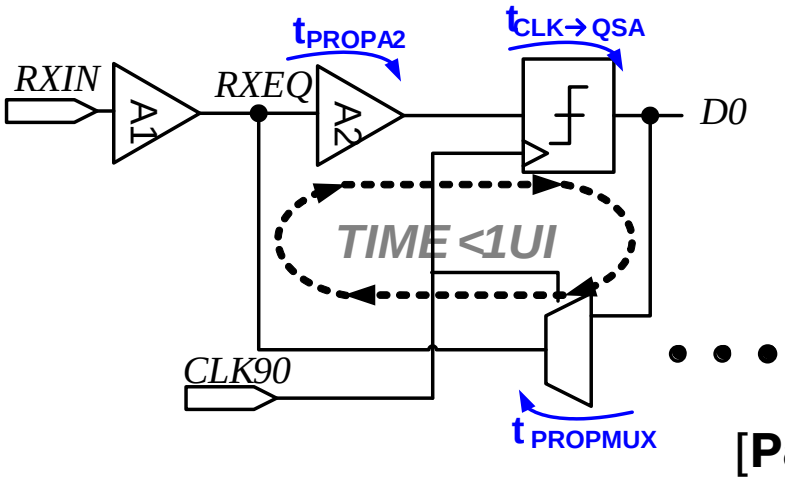

[**Payne**]

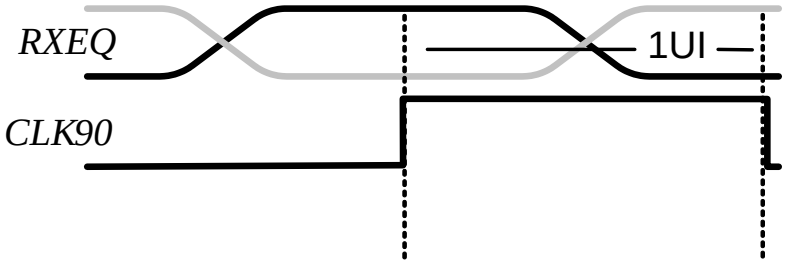

$$t_{CLK\rightarrow QSA} + t_{PROPMUX} + t_{PROPA2} \leq 1UI$$

- Must resolve data and feedback in 1 bit period
  - TI design actually does this in ½UI for CDR

DFE Loop Unrolling

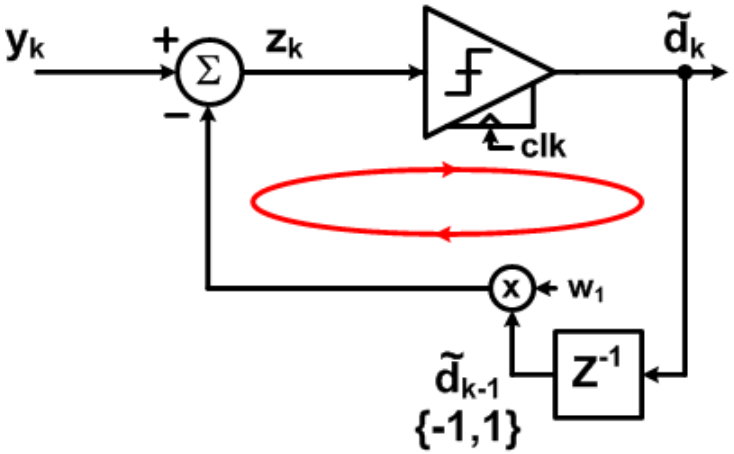

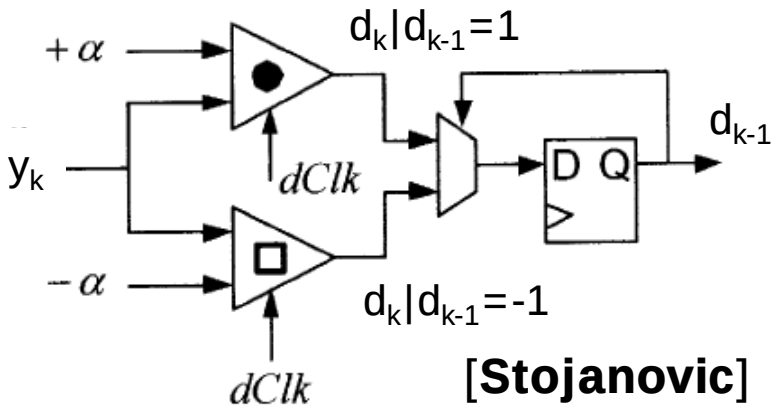

[**Stojanovic**]

- Instead of feeding back and subtracting ISI in 1UI
- Unroll loop and pre-compute 2 possibilities (1-tap DFE) with adjustable slicer threshold
- With increasing tap number, comparator number grows as $2^{\#\text{taps}}$

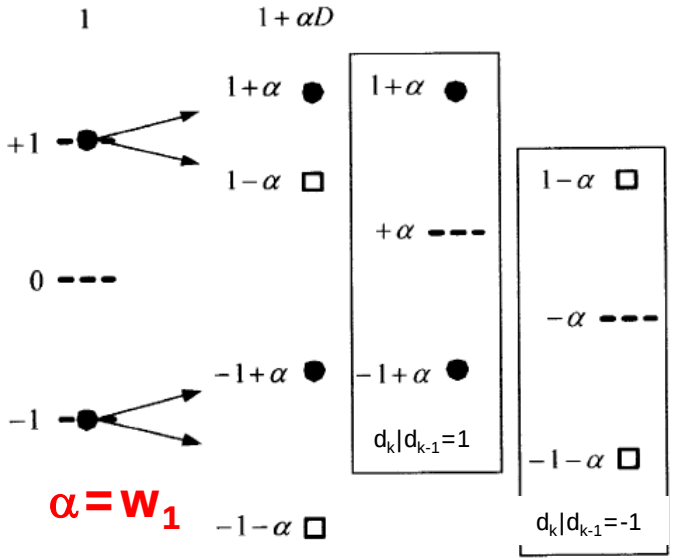

$$\tilde{d}_k = \begin{cases} \operatorname{sgn}(y_k - w_1) \text{ "if" } \tilde{d}_{k-1} = 1 \\ \operatorname{sgn}(y_k + w_1) \text{ "if" } \tilde{d}_{k-1} = -1 \end{cases}$$

DFE Resistive-Load Summer

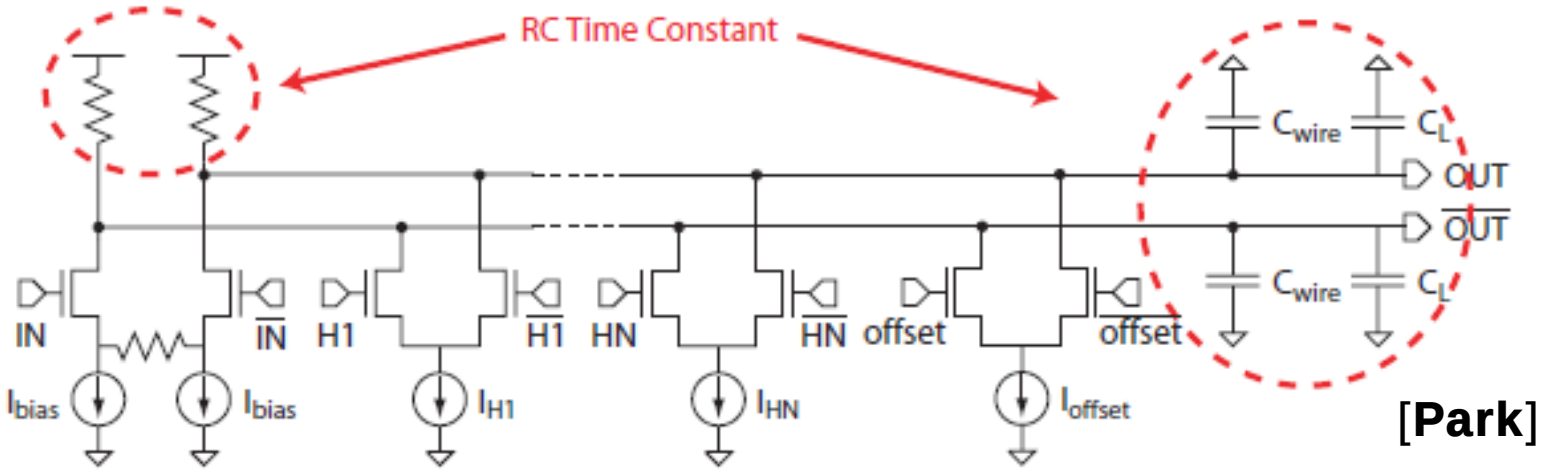

Summer Swing $= IR, \quad \tau = RC$

- Summer performance is critical for DFE operation
- Summer must settle within a certain level of accuracy (>95%) for ISI cancellation
- Trade-off between summer output swing and settling time
- Can result in large bias currents for input and taps

DFE Integrating Summer

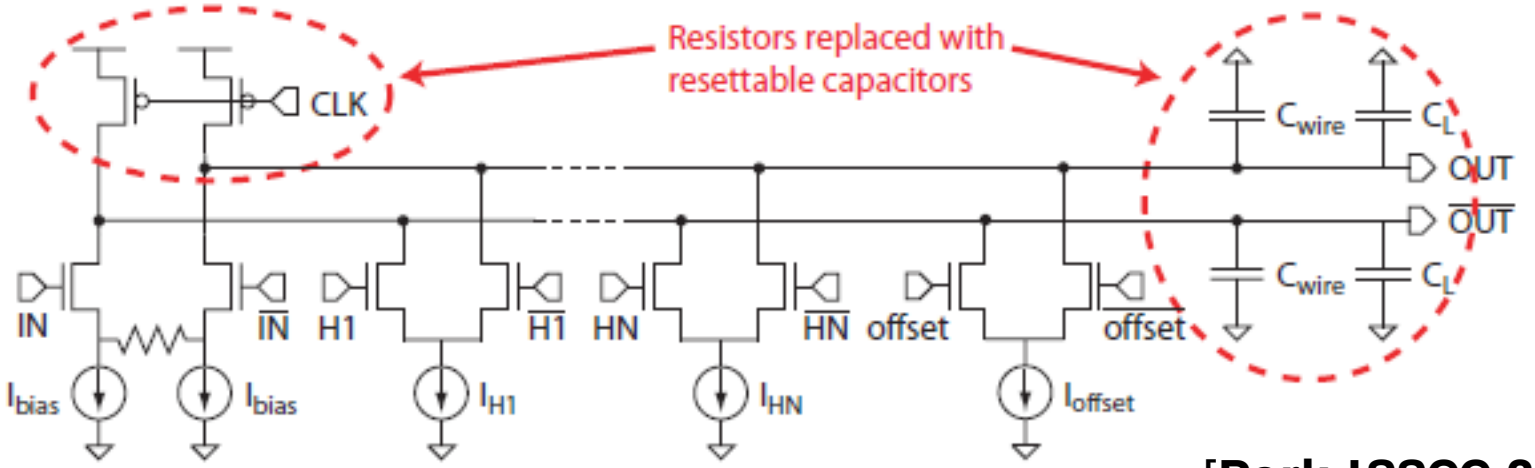

[**Park ISSCC 2007**]

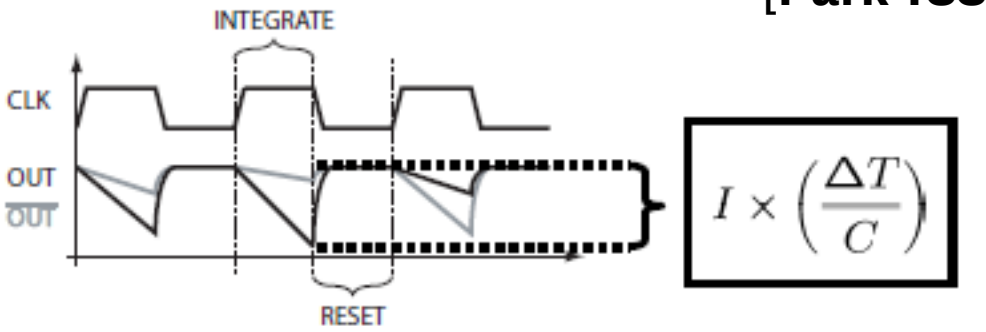

- Integrating current onto load capacitances eliminates RC settling time
- Since $\Delta T/C > R$, bias current can be reduced for a given output swing
  - Typically a 3x bias current reduction

Digital RX FIR & DFE Equalization Example

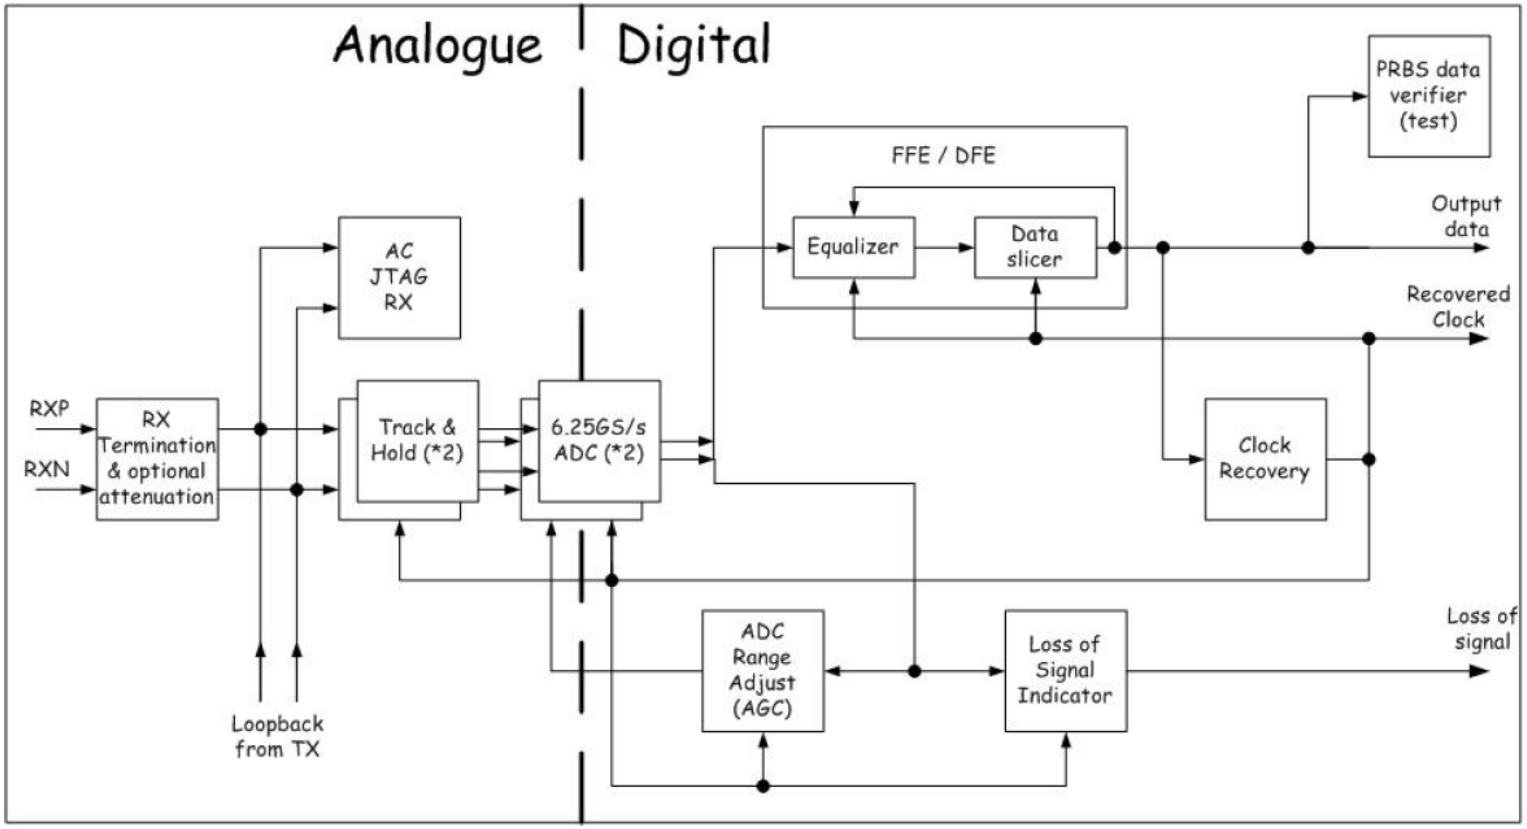

- 12.5GS/s 4.5-bit Flash ADC in 65nm CMOS [**Harwood ISSCC 2007**]
- 2-tap FFE & 5-tap DFE
- XCVR power (inc. TX) = 330mW, Analog = 245mW, Digital = 85mW

DFE with Feedback FIR Filter

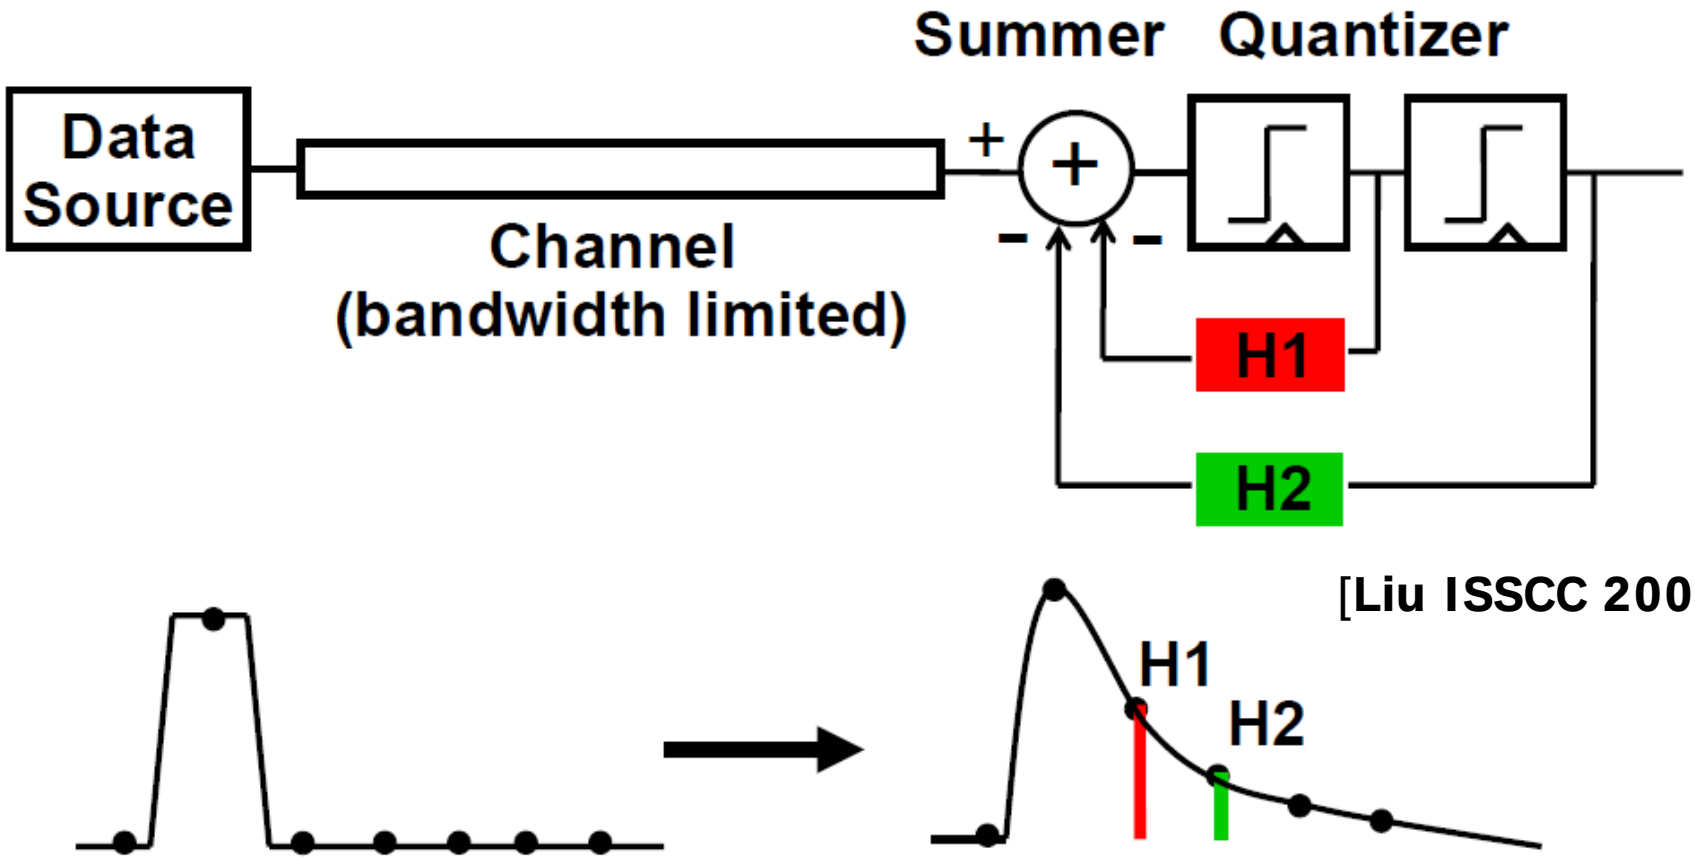

- DFE with 2-tap FIR filter in feedback will only cancel ISI of the first two post-cursors

“Smooth” Channel

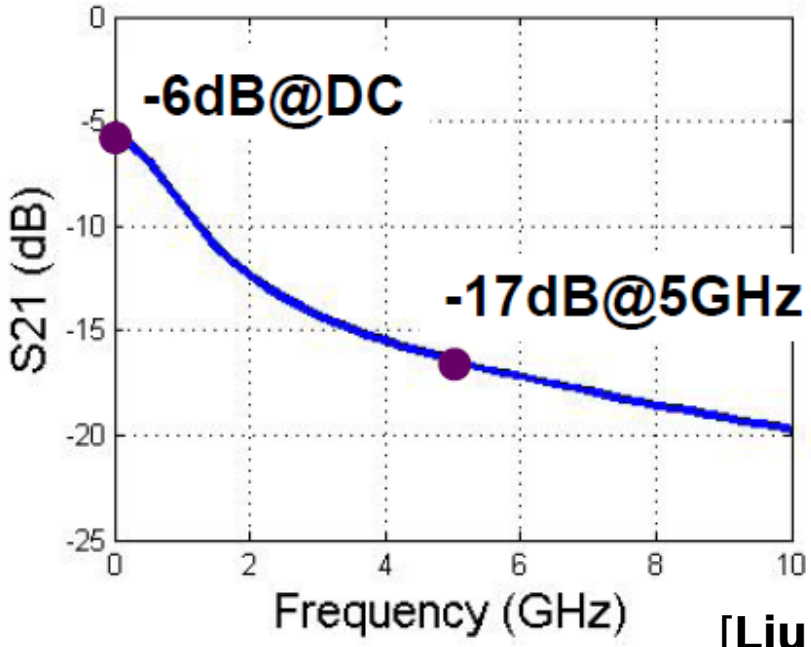

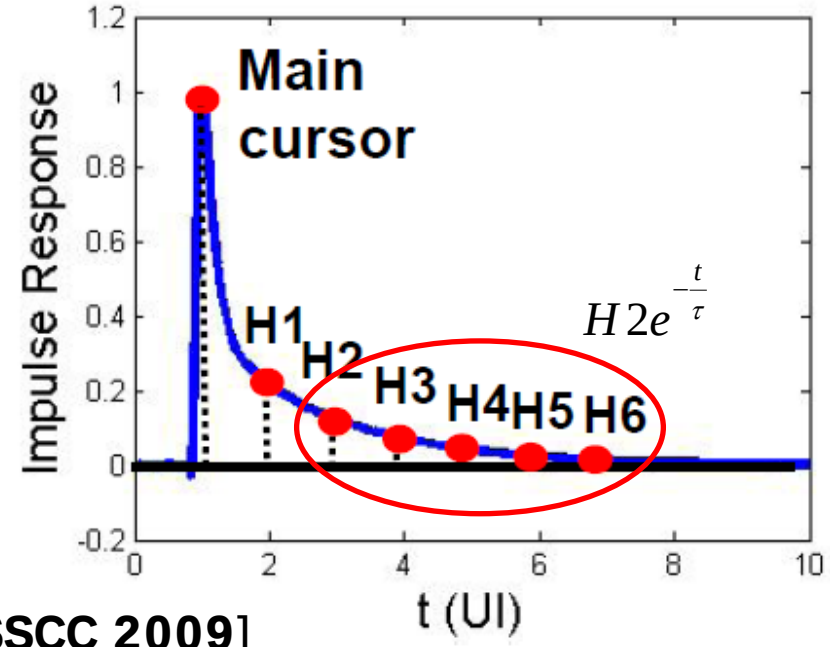

[**Liu ISSCC 2009**]

- A DFE with FIR feedback requires many taps to cancel ISI
- Smooth channel long-tail ISI can be approximated as exponentially decaying
  - Examples include on-chip wires and silicon carrier wires

DFE with IIR Feedback

[**Liu ISSCC 2009**]

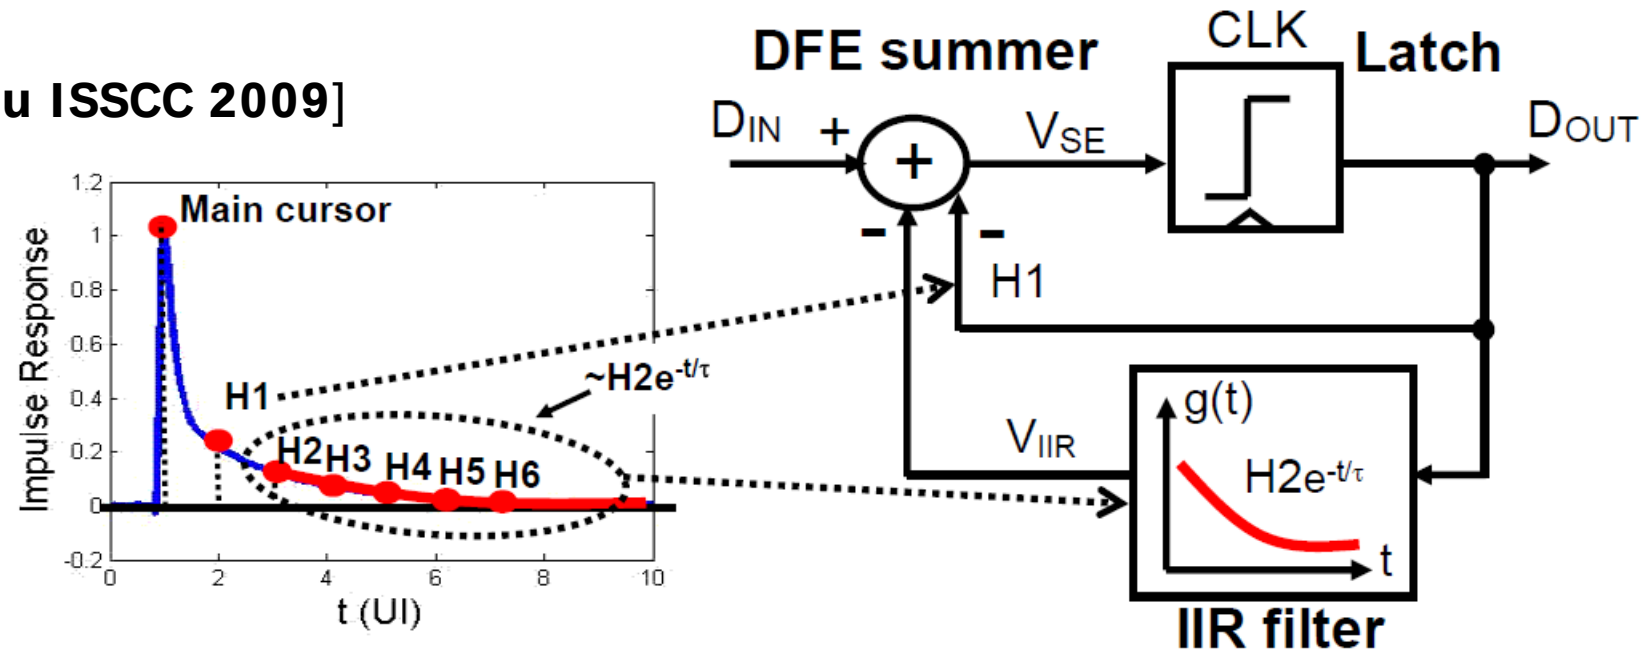

- Large 1st post-cursor H1 is canceled with normal FIR feedback tap
- Smooth long tail ISI from 2nd post-cursor and beyond is canceled with low-pass IIR feedback filter
- Note: channel needs to be smooth (not many reflections) in order for this approach to work well

DFE with IIR Feedback RX Architecture

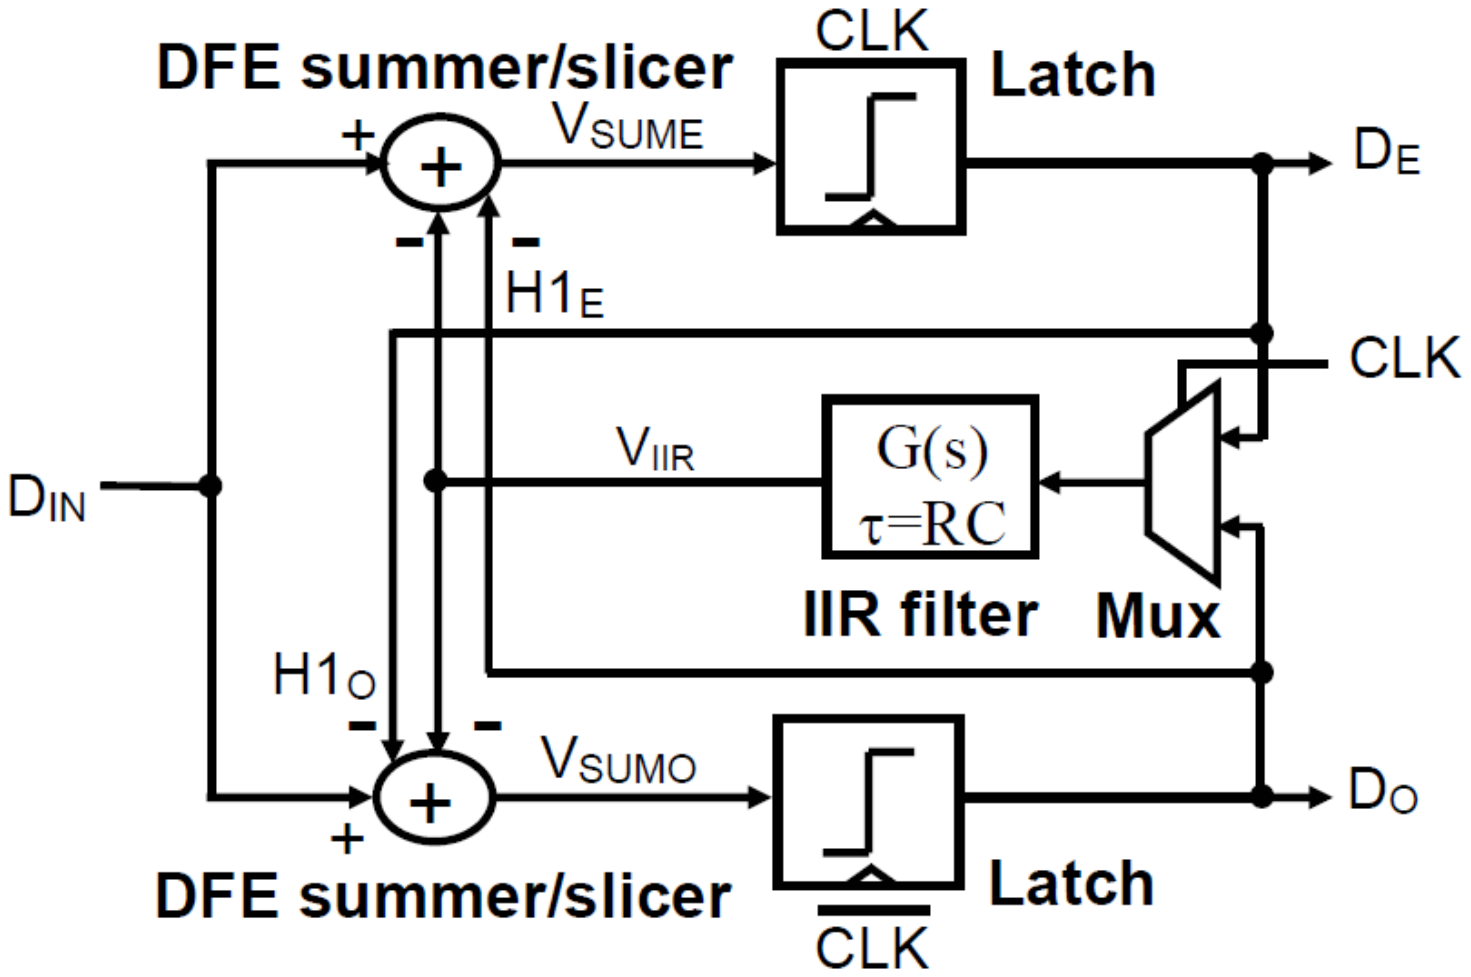

[**Liu ISSCC 2009**]

Merged Summer & Partial Slicer

[**Liu ISSCC 2009**]

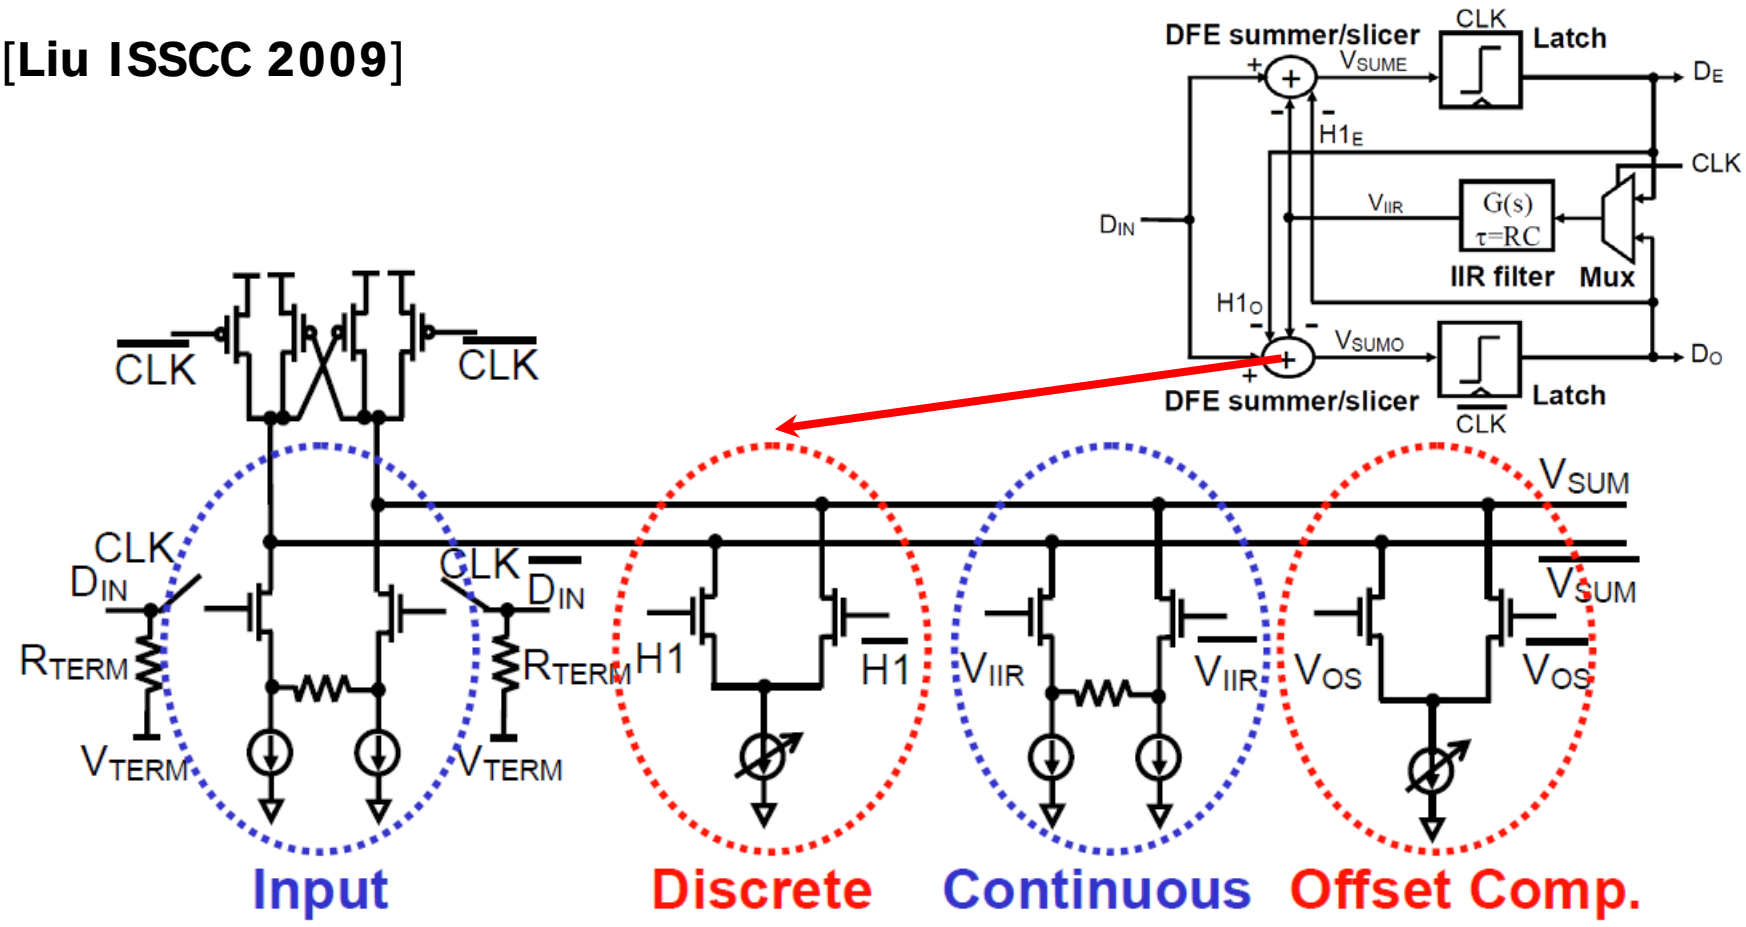

- Integrating summer with regeneration PMOS devices to realize partial slicer operation

Merged Mux & IIR Filter

[Liu ISSCC 2009]

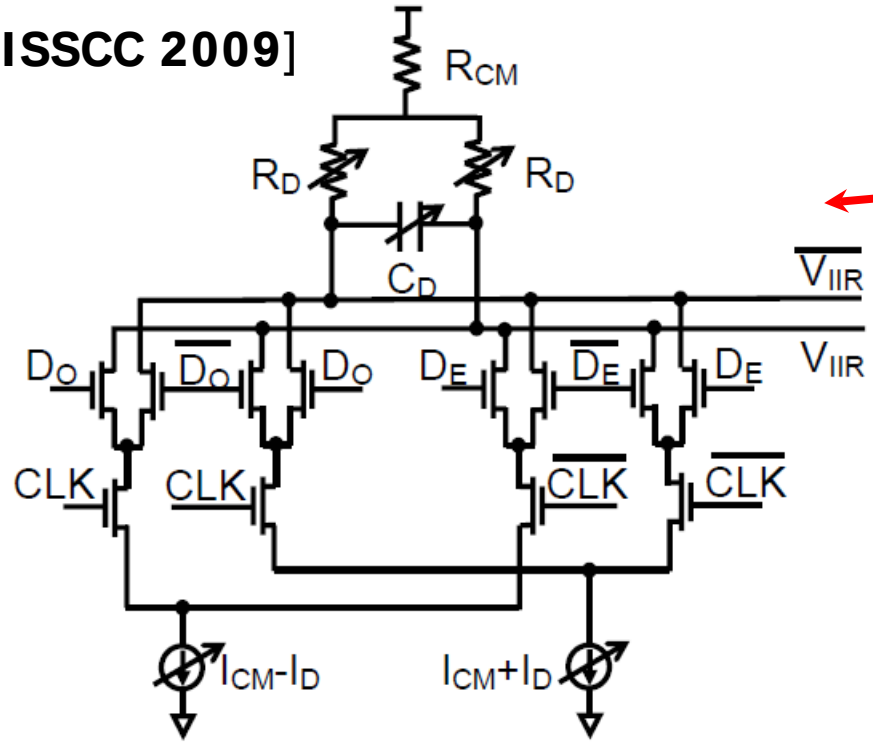

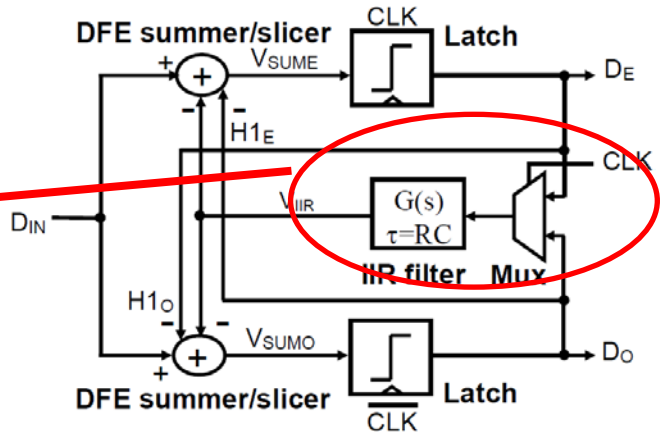

- Low-pass response (time constant) implemented by R_D and C_D
- Amplitude controlled by R_D and I_D
- 2 UI delay implemented through mux to begin cancellation at 2nd post-cursor

Outline

- Introduction
- Channel characteristics
- Equalizer circuits
- Equalizer adaptation techniques
- Optical interconnects
- Conclusion

Setting Equalizer Values

- Simplest approach to setting equalizer values (tap weights, poles, zeros) is to fix them for a specific system
  - Choose optimal values based on lab measurements
  - Sensitive to manufacturing and environment variations
- An adaptive tuning approach allows the optimization of the equalizers for varying channels, environmental conditions, and data rates
- Important issues with adaptive equalization
  - Extracting equalization correction (error) signals
  - Adaptation algorithm and hardware overhead
  - Communicating the correction information to the equalizer circuit

TX FIR Adaptation Error Extraction

- While we are adapting the TX FIR, we need to measure the response at the receiver input
- Equalizer adaptation (error) information is often obtained by comparing the receiver input versus the desired symbol levels, dLev
- This necessitates additional samplers at the receiver with programmable threshold levels

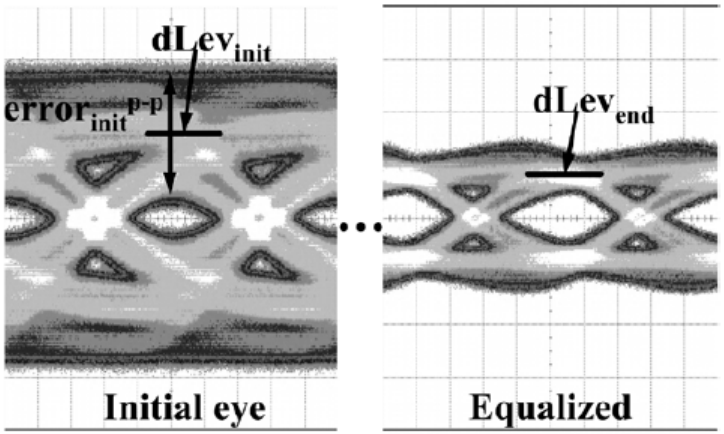

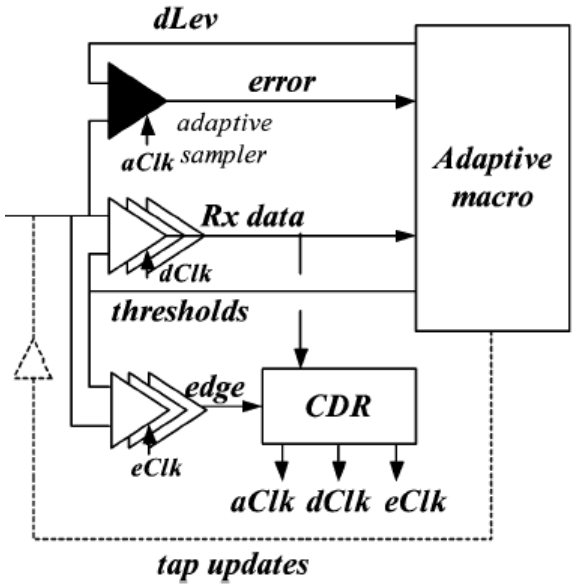

[**Stojanovic JSSC 2005**]

TX FIR Adaptation Algorithm

- The sign-sign LMS algorithm is often used to adapt equalization taps due to implementation simplicity

$$w_{n+1}^{k} = w_{n}^{k} + \Delta_{w}\text{sign}(d_{n-k})\text{sign}(e_{n})$$

w = tap coefficients, n = time instant, k = tap index, d_n = received data,

e_n = error with respect to desired data level, $dLev$

- As the desired data level is a function of the transmitter swing and channel loss, the desired data level is not necessarily known and should also be adapted

$$dLev_{n+1} = dLev_{n} - \Delta_{dLev}\text{sign}(e_{n})$$

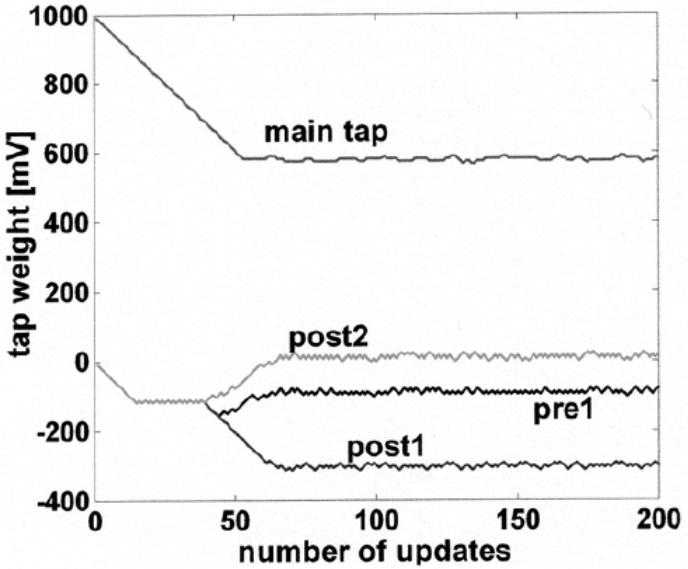

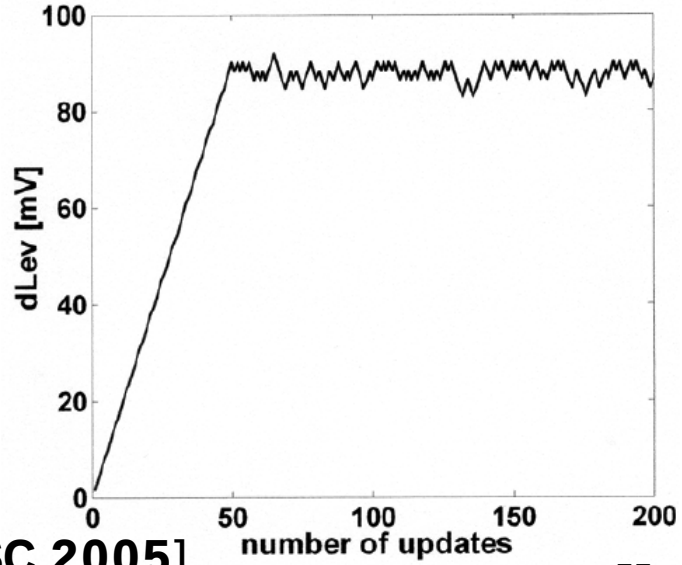

[**Stojanovic JSSC 2005**]

TX FIR Common-Mode Back-Channel

- In order to communicate FIR tap update information back to the TX, a back-channel is necessary
- One option is to use low data rate (~10Mb/s) common-mode signaling from the RX to TX on the same differential channel

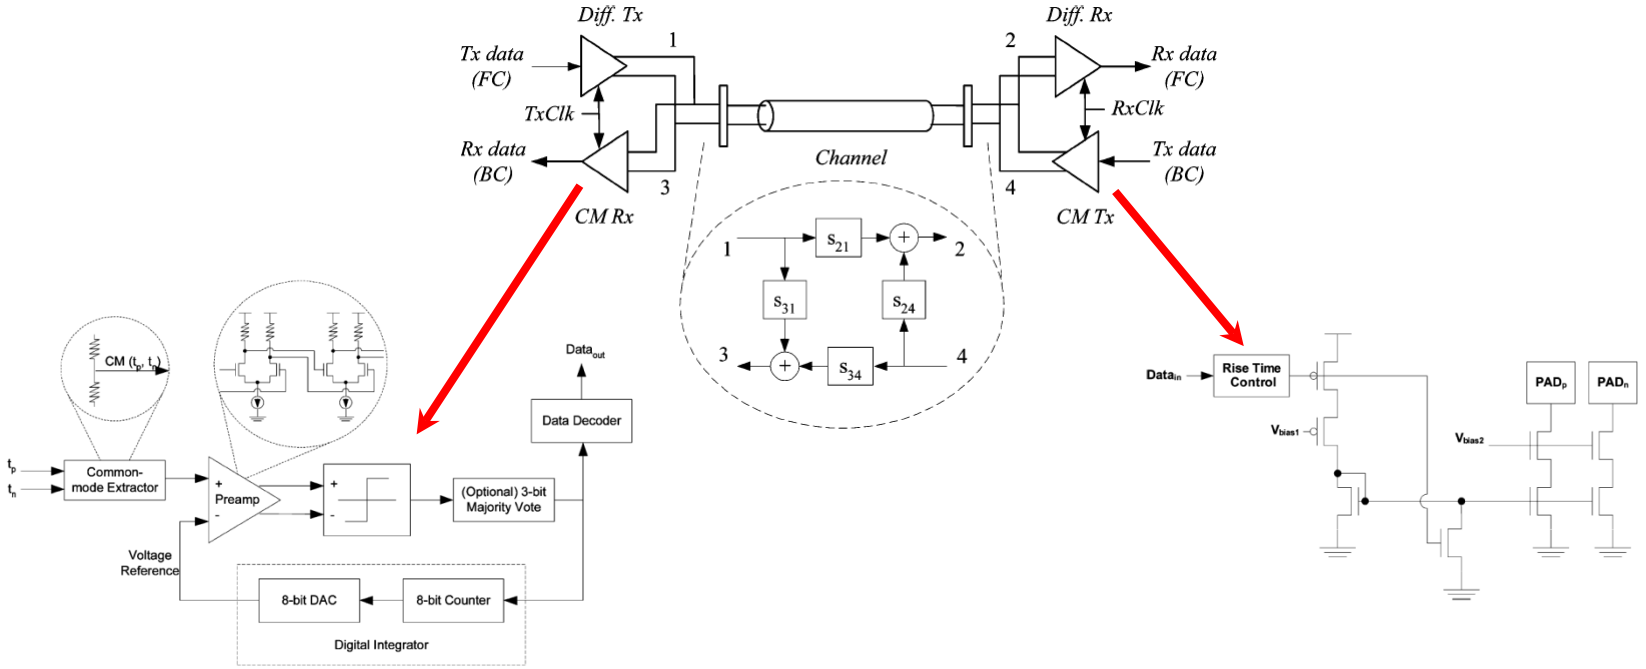

[**Stojanovic JSSC 2005**]

TX FIR Data Encoder Back-Channel

- Another option is to use a high-speed TX channel on the RX side that communicates data back to the TX under adaptation
- Flexibility in data encoding (8B10B/Q) allows low data rate tap adaptation information to be transmitted back without data rate overhead

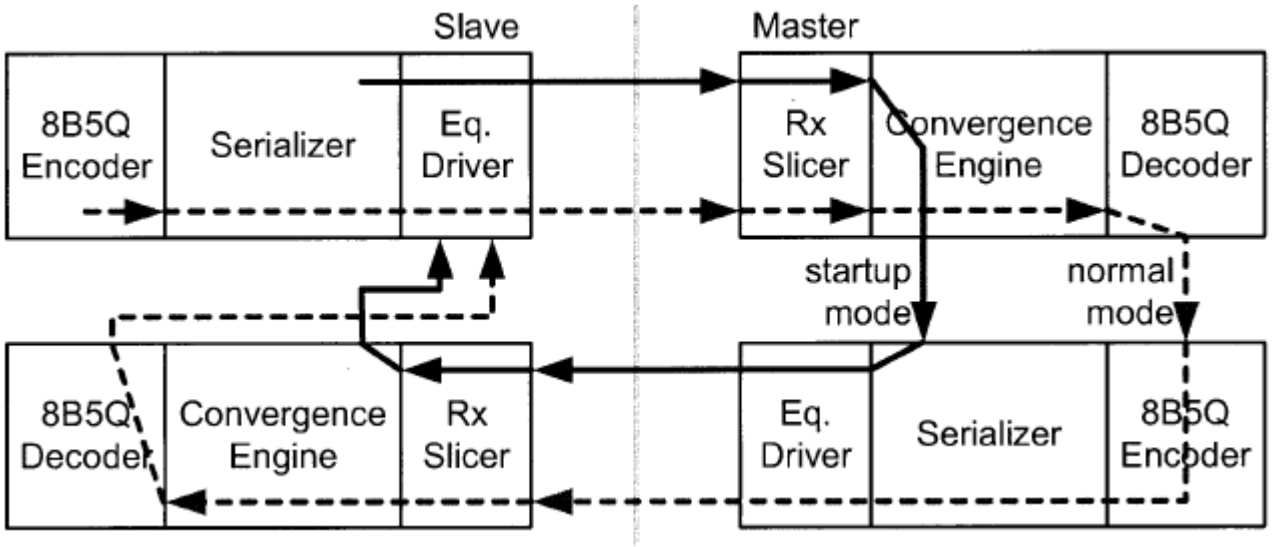

[**Stonick JSSC 2003**]

CTLE Tuning with PSD Measurement

- One approach to CTLE tuning is to compare low-frequency and high-frequency spectrum content of random data
- For ideal random data, there is a predictable ratio between the low-frequency power and high-frequency power
- The error between these power components can be used in a servo loop to tune the CTLE

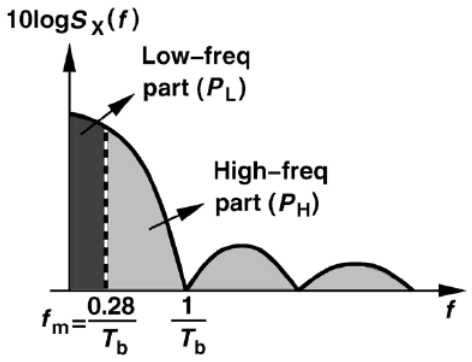

$$s_x(f) = T_b\left[\frac{\sin(\pi f T_b)}{\pi f T_b}\right]^2$$

$$\int_0^{f_m} s_x(f)df = \int_{f_m}^{\infty} s_x(f)df = \frac{1}{4}$$

where $f_m \approx \frac{0.28}{T_b}$

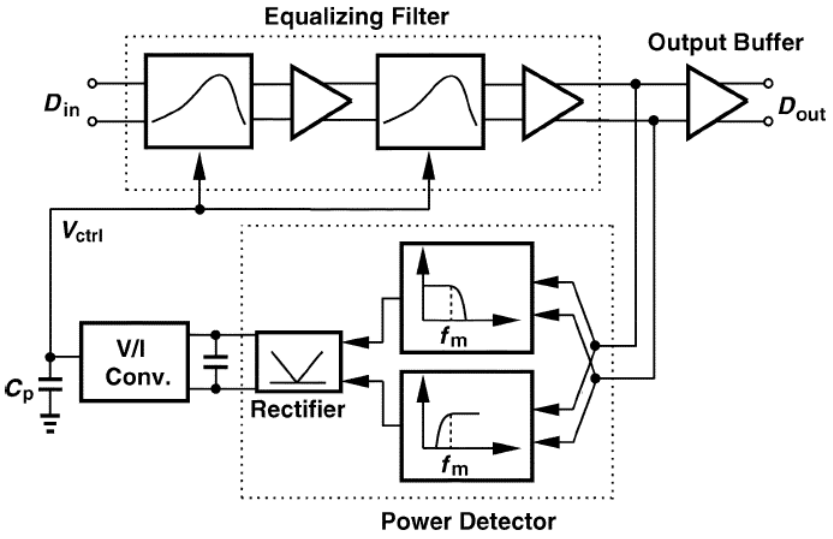

[**Lee JSSC 2006**]

CTLE Tuning w/ Output Amplitude Measurement

- CTLE tuning can also be done by comparing low-frequency and high-frequency average amplitude
- Approximating the equalized data as a sine wave, a predictable ratio exists between the low frequency average and high-frequency average
- Equalizer settings are adjusted until the high frequency peak-to-peak swing matches the low-frequency peak-to-peak swing

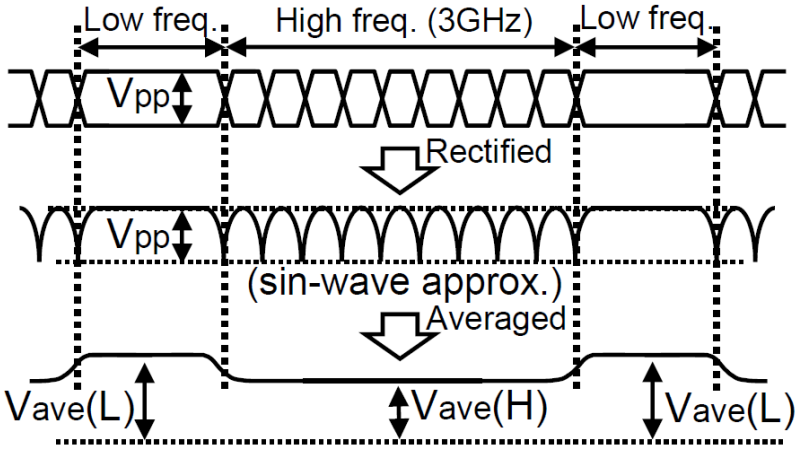

- $V_{ave(L)}=V_{pp}$
- $V_{ave(H)}=(2/\pi)\times V_{pp}$

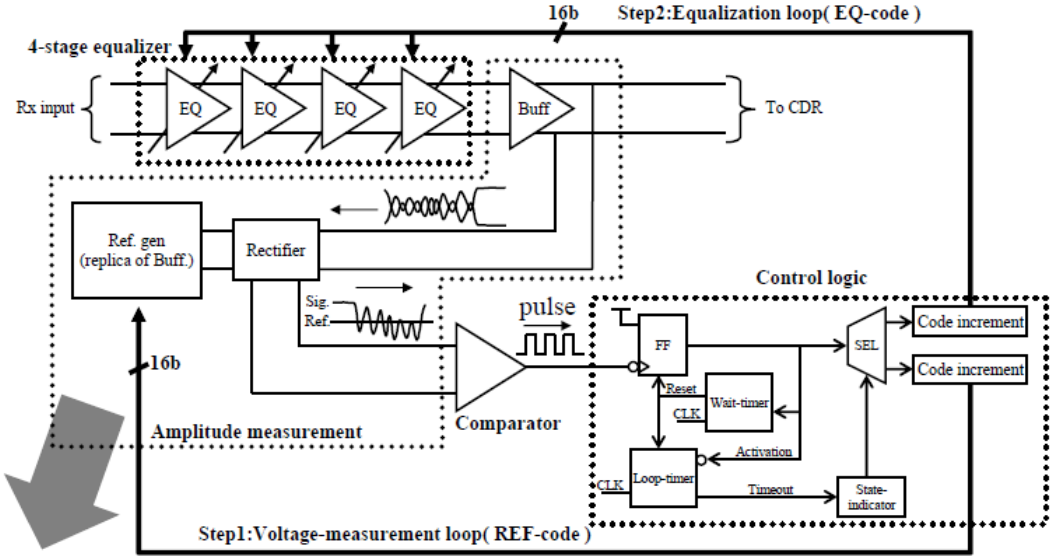

[**Uchiki ISSCC 2008**]

DFE Tuning

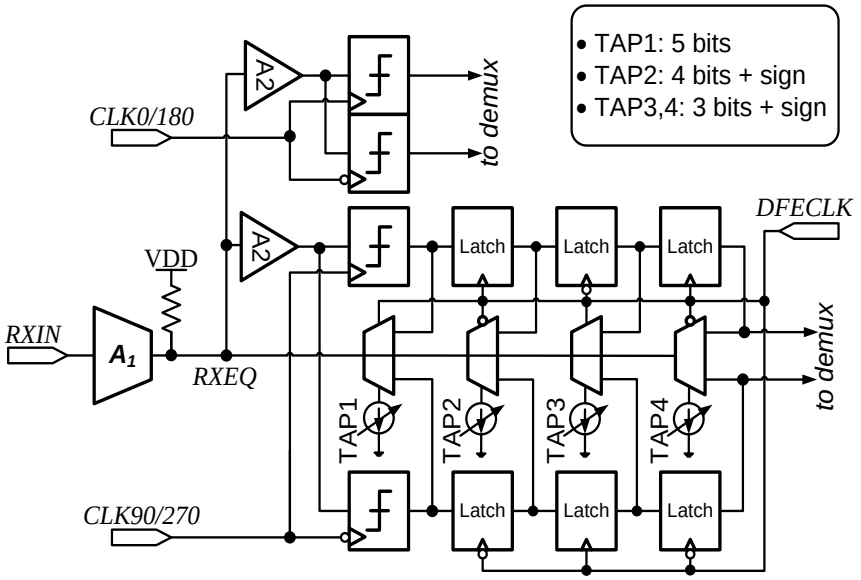

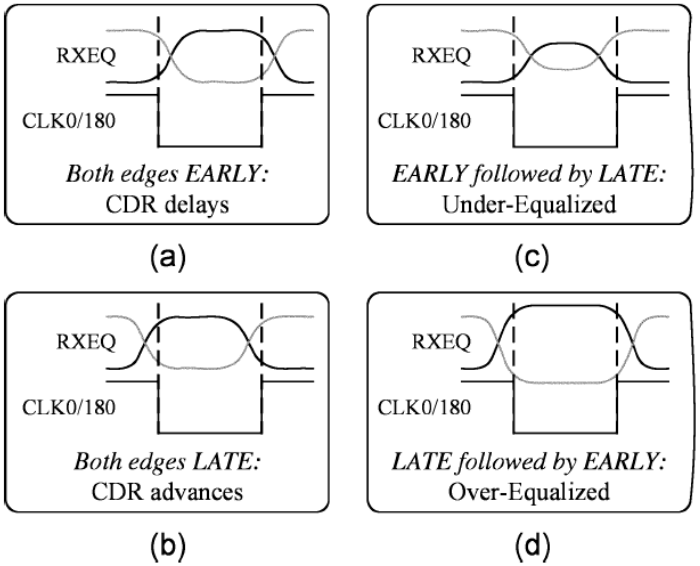

- 2x oversampling the equalized signal at the edges can be used to extract information to adapt a DFE and drive a CDR loop
- Sign-sign LMS algorithm used to adapt DFE tap values

[**Payne JSSC 2005**]

Outline

- Introduction
- Channel characteristics
- Equalizer circuits
- Equalizer adaptation techniques
- Optical interconnects
- Conclusion

High-Speed Optical Link System

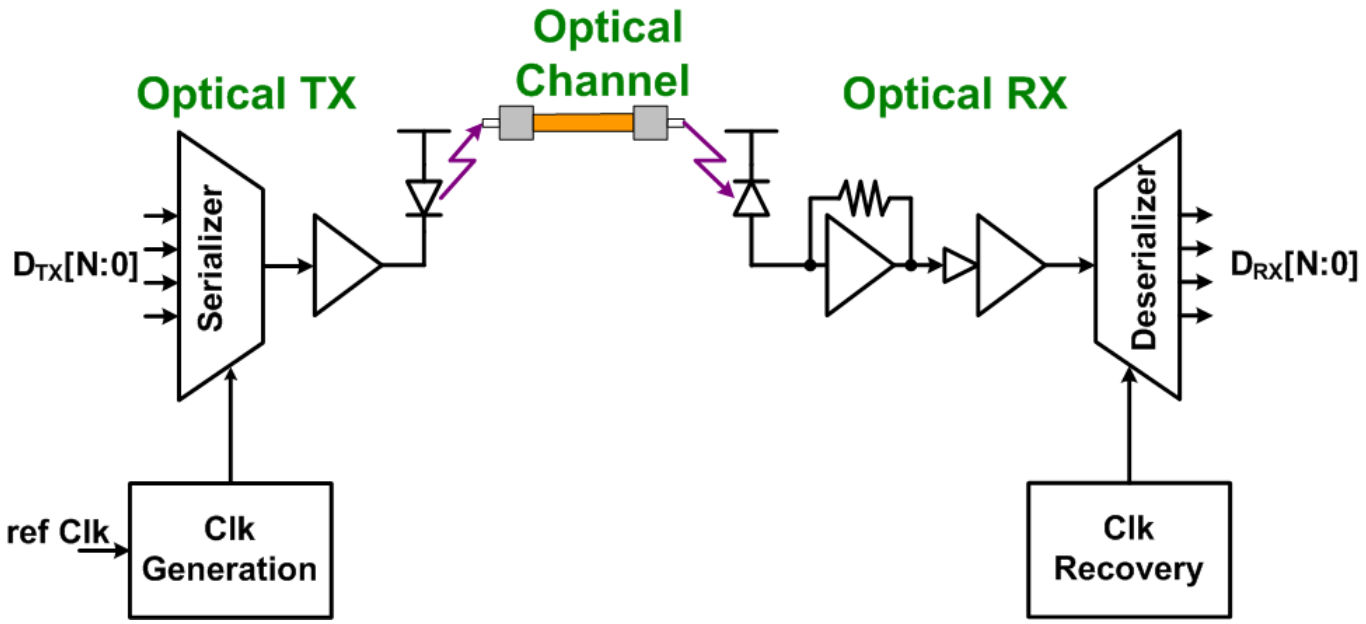

- Optical interconnects remove many channel limitations
  - Reduced complexity and power consumption
  - Potential for high information density with wavelength-division multiplexing (WDM)

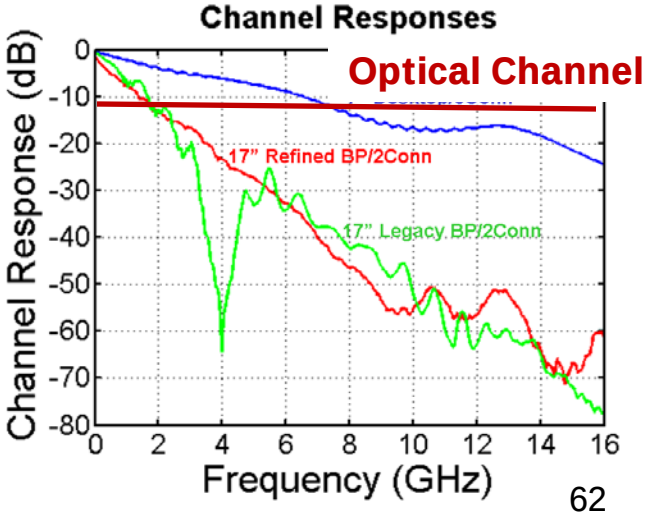

Wavelength-Division Multiplexing

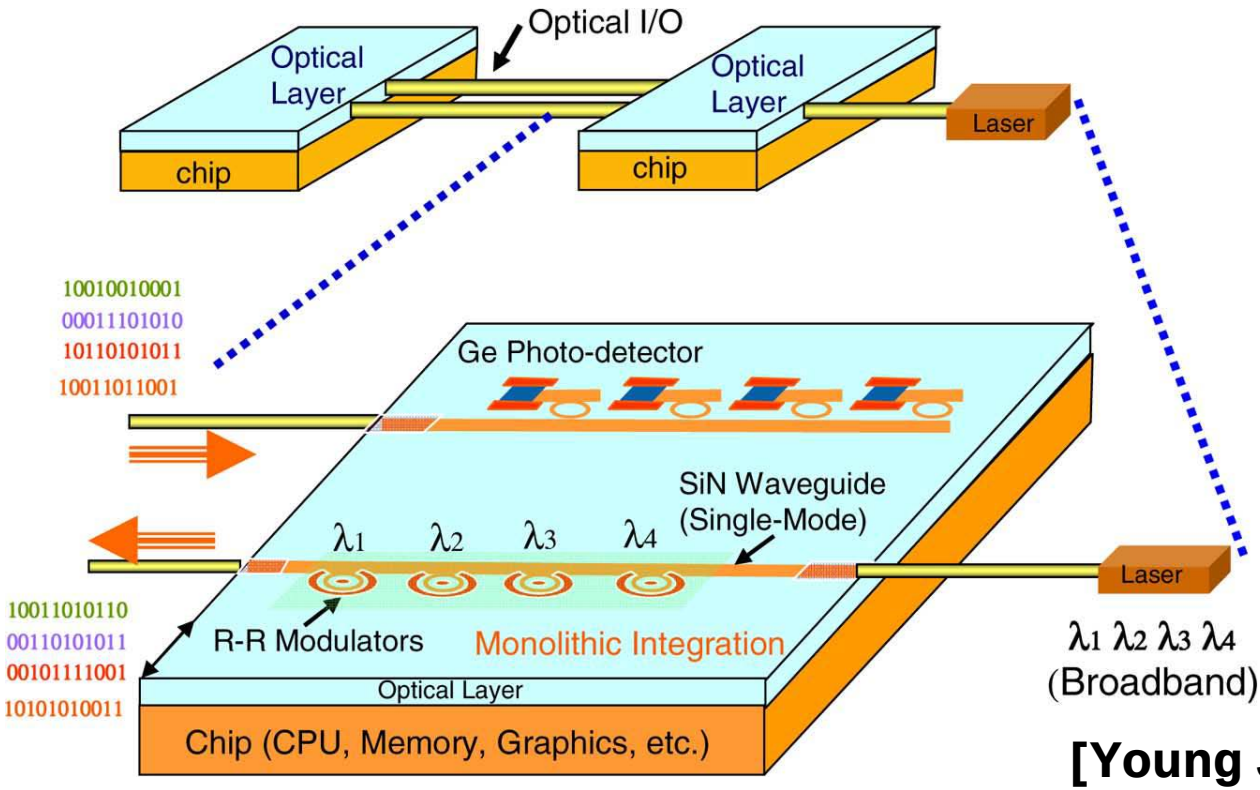

[Young JSSC 2010]

- WDM allows for multiple high-bandwidth (10+Gb/s) signals to be packed onto one optical channel

Conclusion & Future Trends

- Data rates are scaling faster than electrical channel bandwidths, necessitating higher complexity, adaptive, and more efficient equalization circuits
- Nanometer CMOS scaling also provides the potential for more advanced systems to deal with ISI
  - ADC front-ends with complex digital equalization
  - Error-correction coding
- On the horizon are optical interconnect systems which provide the potential for distance-independent bandwidth which scales with the number of wavelengths/channel