High-Speed Links

Vladimir Stojanovic

(with slides from M. Horowitz, J. Zerbe, K.Yang and W. Ellersick)

Agenda : High Speed Links

- High-Speed Links, What,Where?
- Signaling Faster - Evolution
  - » Circuits
  - » Channel
- System-level design
  - » Channel designer’s view
  - » IC designer’s view
- Demo

What Makes a Link?

- Signaling: sending and receiving the information

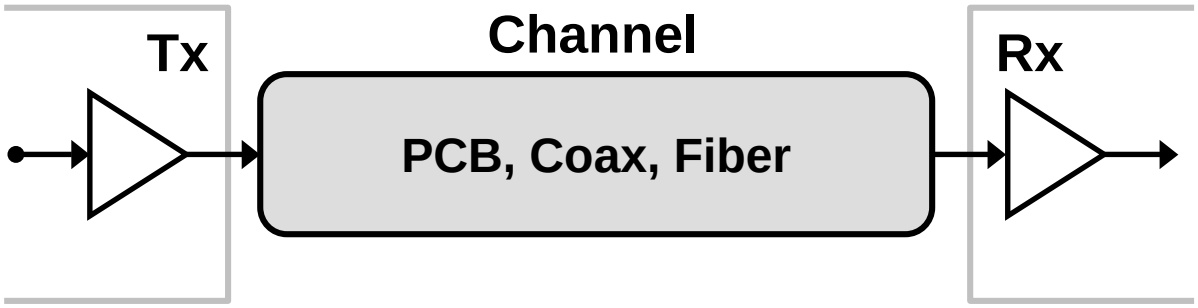

- Clocking: Determining which bit is which

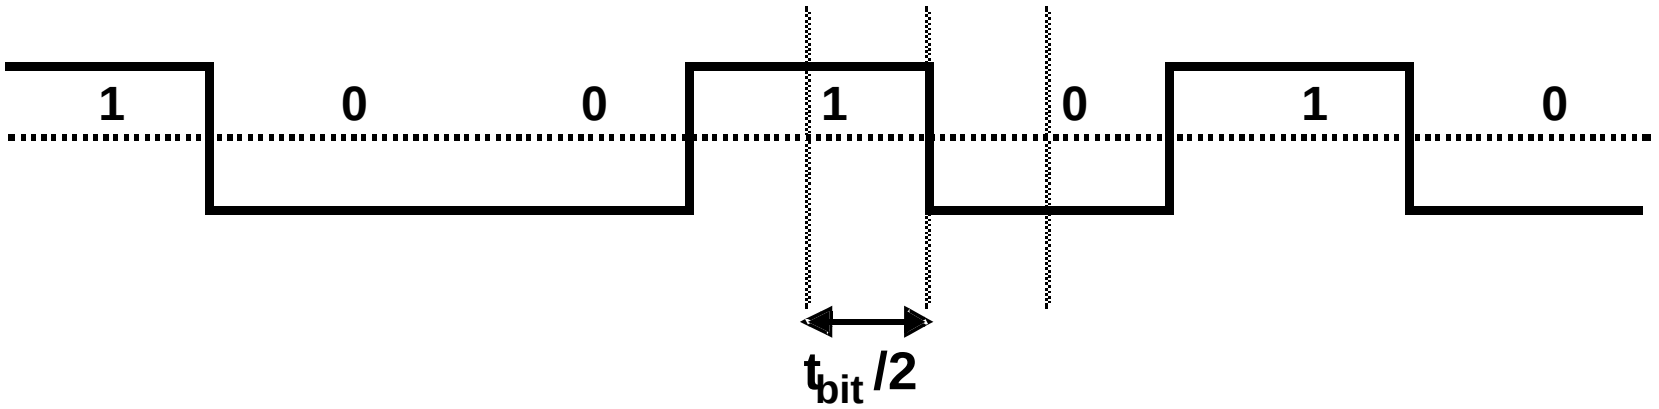

Spanning A Broad Space

- Inverter.....to……..DSL modem
- Metrics
  - » Speed
  - » Latency
  - » Electrical environment
  - » Power & area
  - » Volume

Increasing Chip I/O Bandwidth

- Computers:
  - Main memory:
    - SDRAM100 (100 Mbps) ➲ RDRAM (0.8-1.1 Gbps)
  - Peripherals:
    - PCI (66 Mbps) ➲ Infiniband (2.5 Gbps)
- Networks:
  - Physical Front End:
    - LAN: Fast-Eth (100 Mbps) ➲ Gigabit-Eth (1Gbps)
    - WAN: OC-12 (625 Mbps) ➲ OC-192 (12.5 Gbps)
  - Switch Fabric:
    - 625 Mbps ➲ 2.5 Gbps

Inside the Router

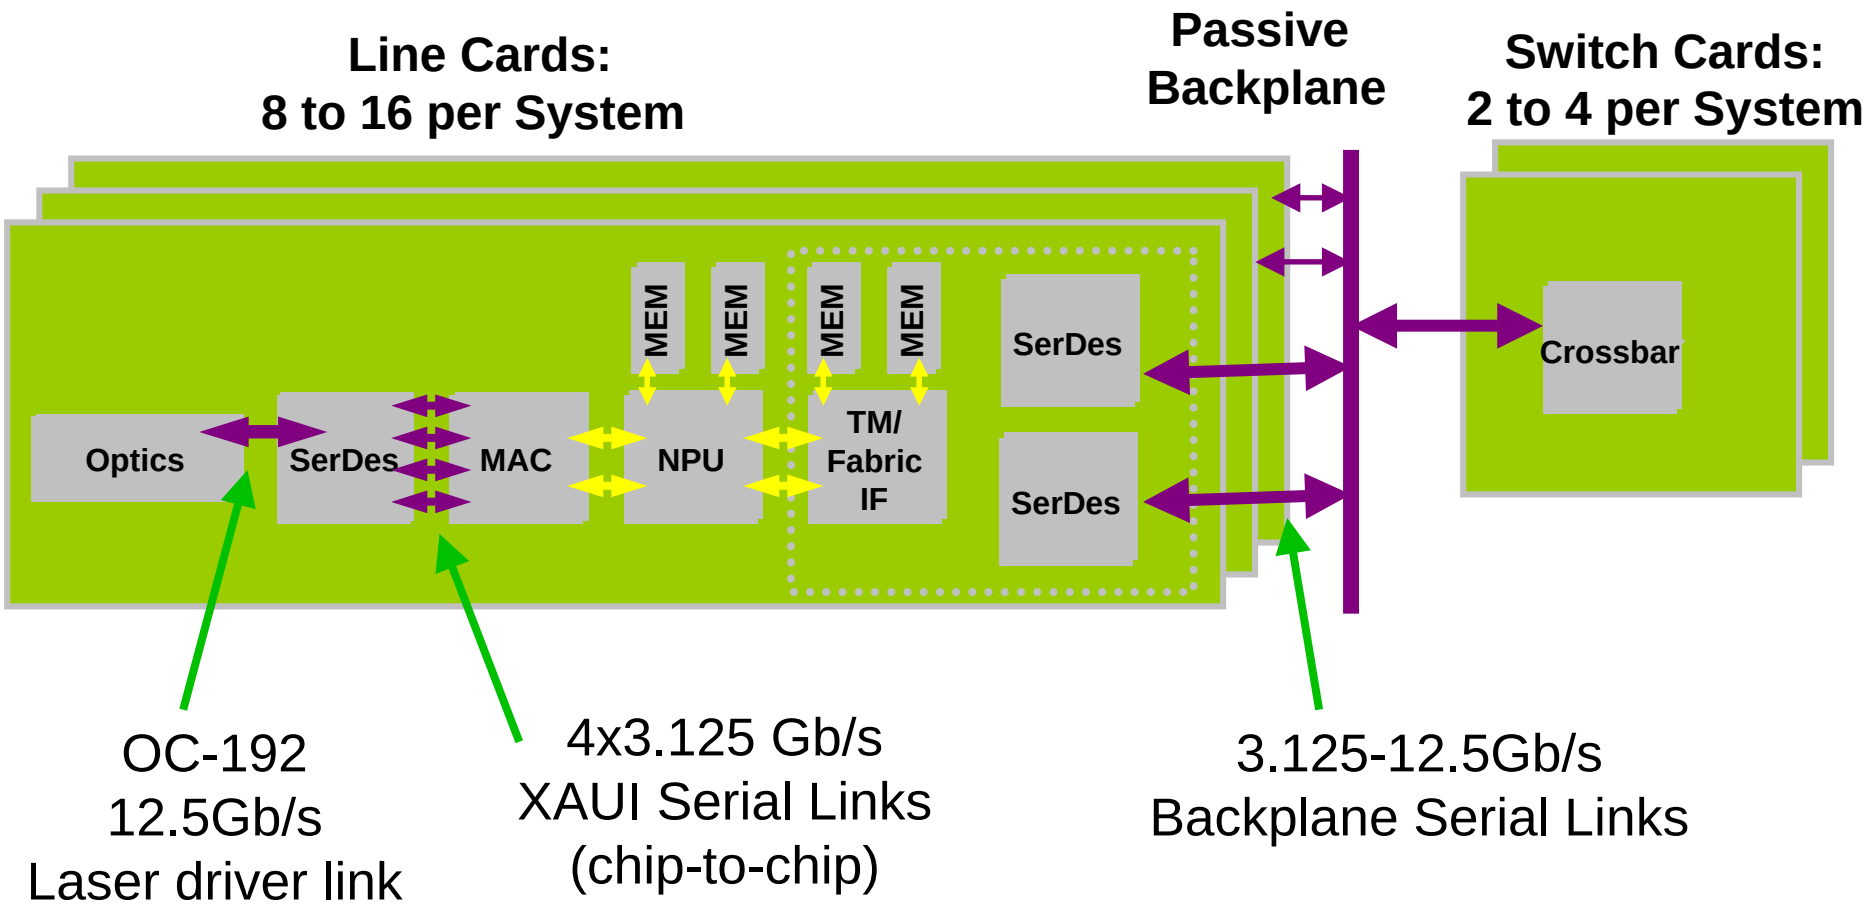

- Regardless of where the links are, there is a constant desire to signal faster and with less power

Serial Link Signaling Over Backplanes - Past

- Channel was not an issue up to 2-3Gb/s

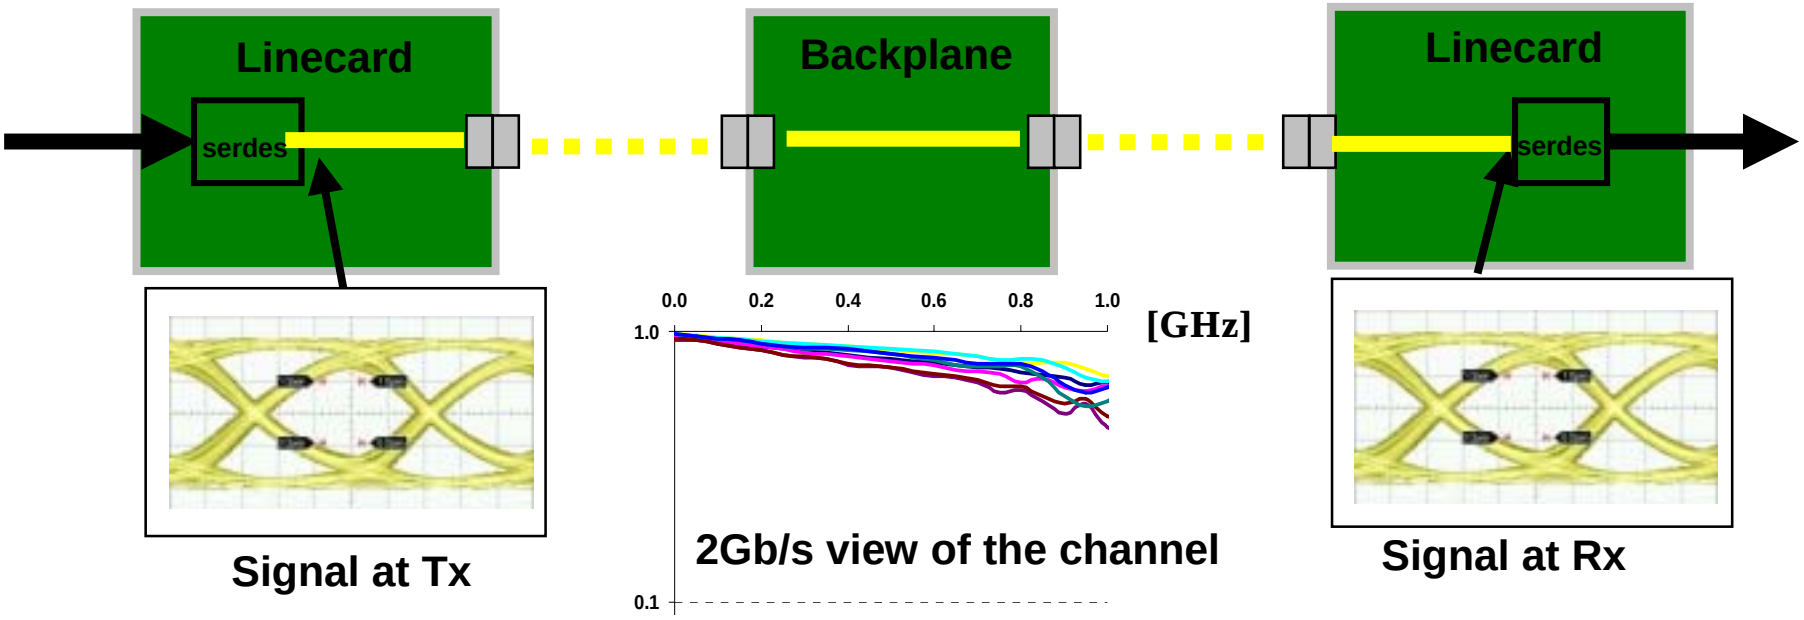

- Designs were limited by transmitter & receiver speed
- Clever circuit design – no communications/SI background needed

Signaling

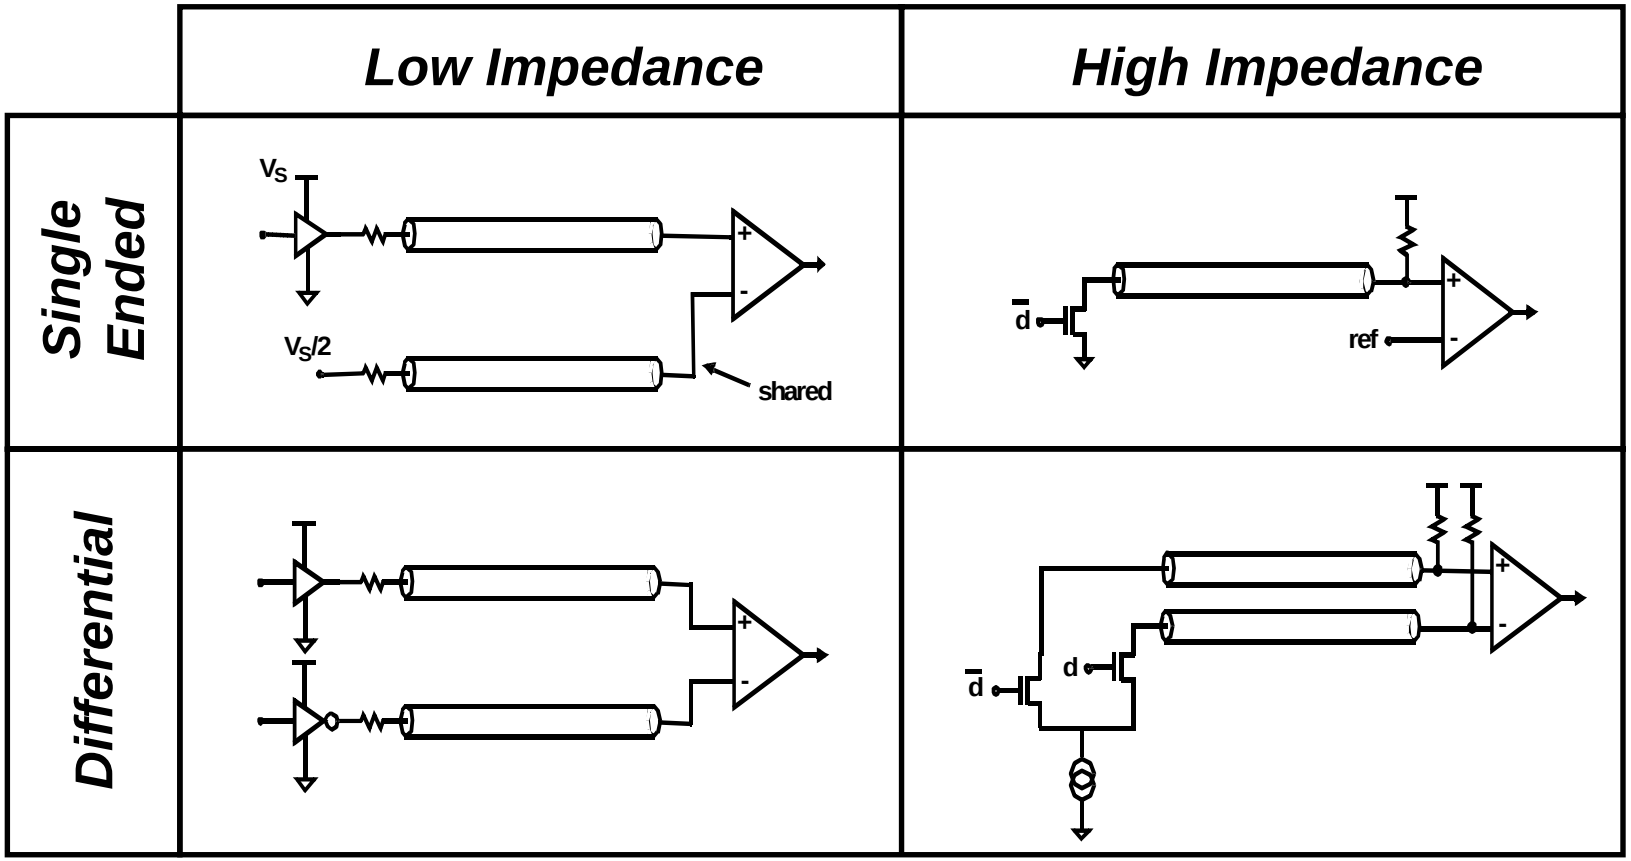
Low Impedance
High Impedance
Single Ended
Differential
V_S
$V_S/2$
shared
ref
$\bar{d}$
d
+
-

Transmitter Design

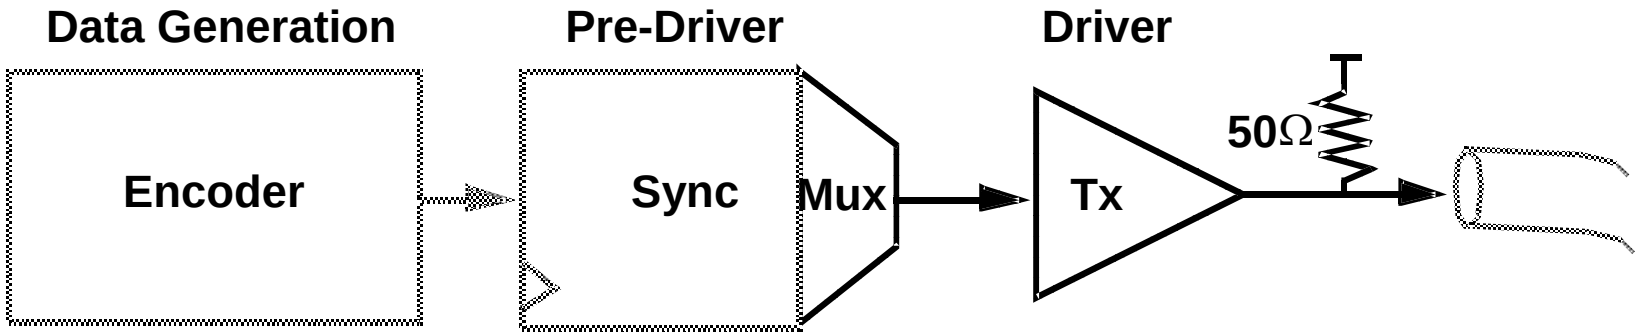

- Critical components: Sync, Mux, Tx
- Design issues:
  - » Slew rate control vs ISI, jitter
  - » Output current and impedance control
- Clock and Driver power dissipation

Output Drivers

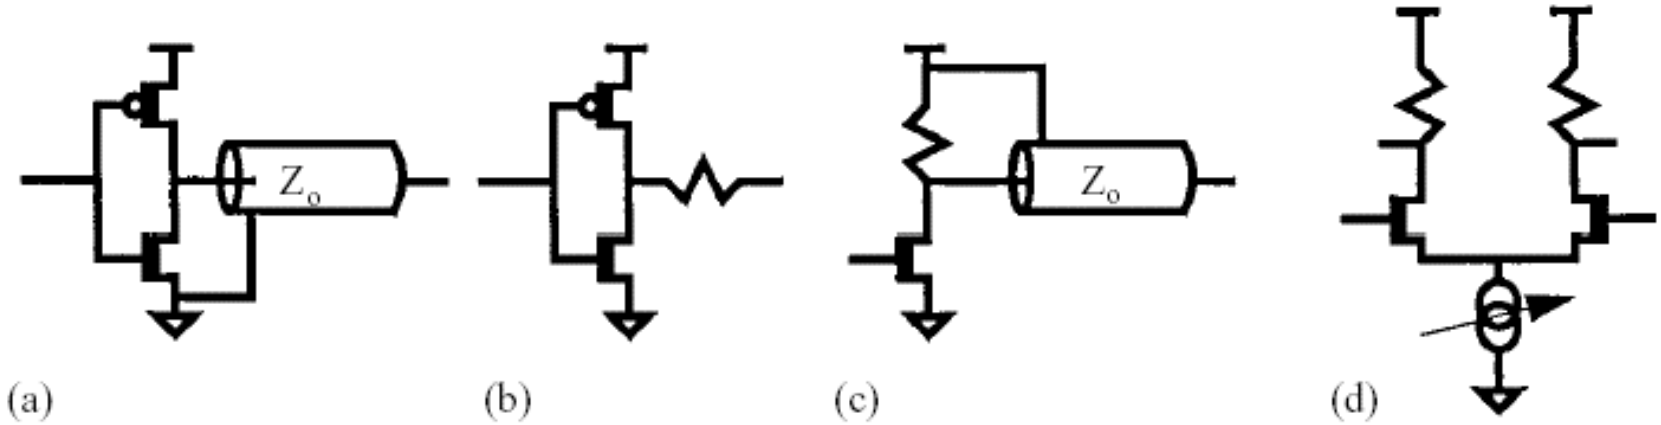

- On-chip clock speed limited to 6-8FO4
- Need to send more bits/clock – multiplex data

Simple Transmitter

- DDR: send a bit per clock edge
- Critical issues:
  - » 50% duty cycle
  - » Tbit > 4-FO4

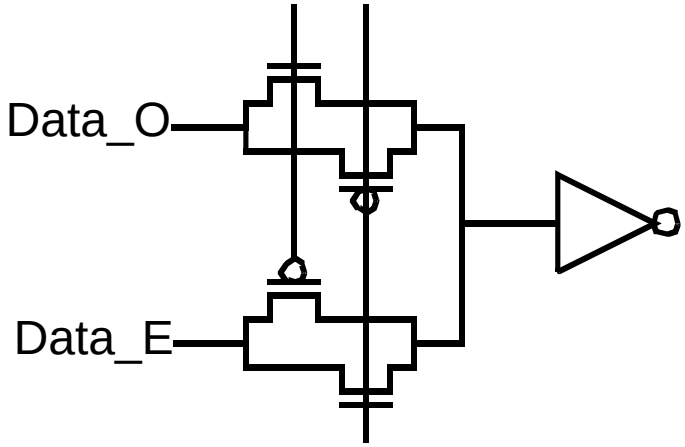

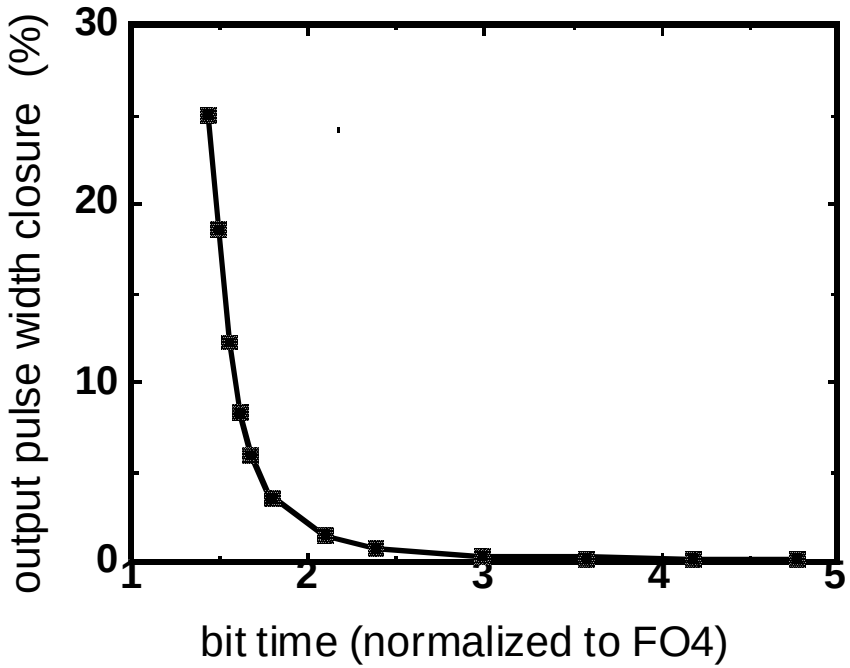

Fastest Transmitter

- Off chip time constant smaller than on chip:
  - Generate current pulse at the output
- Limited only by the output capacitance

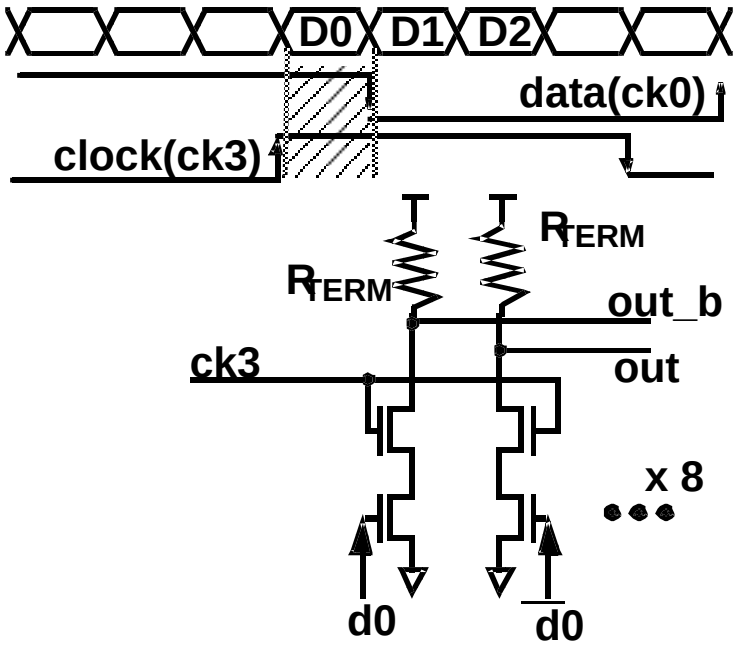

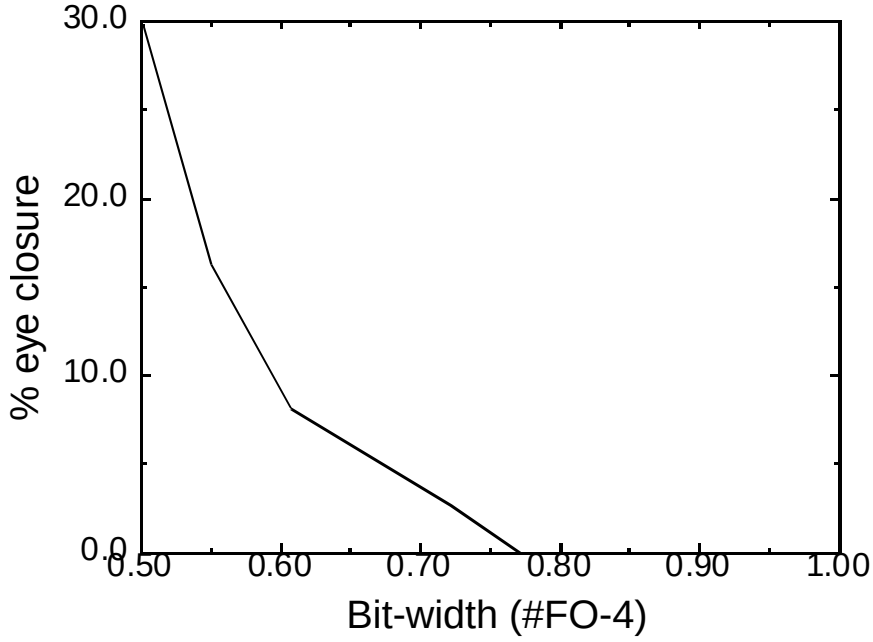

- Limiting time constant 25-Ω*Cpad
- Cpad = 8*Cdriver + Cesd

Simple Receiver

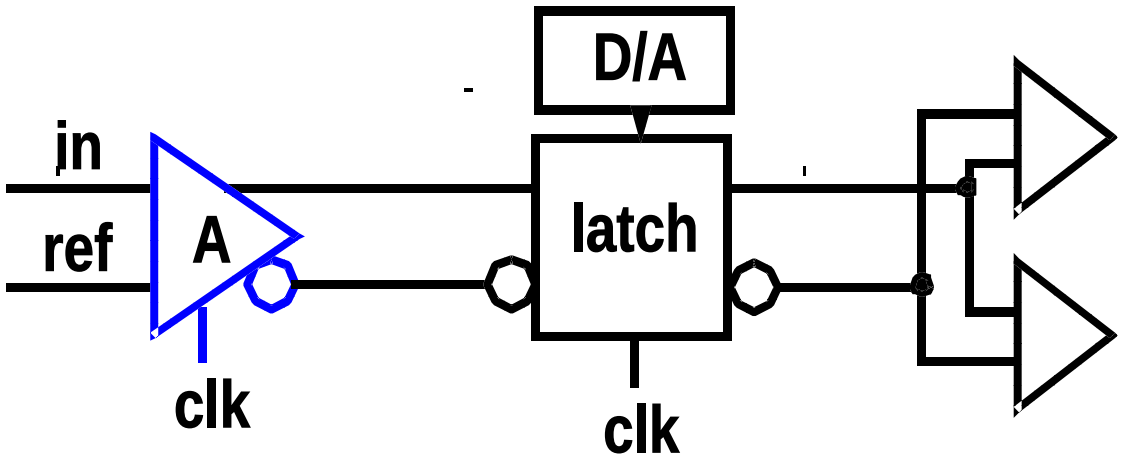

- Regenerative latch has highest gain-bandwidth product of all amplifiers (gain exponential with time – just need to wait long enough)
- Preconditioning stage: filter/integrate rectify CM
- Latch makes decision (4-FO4)
- DAC can be used to compensate offsets

Fastest Receiver

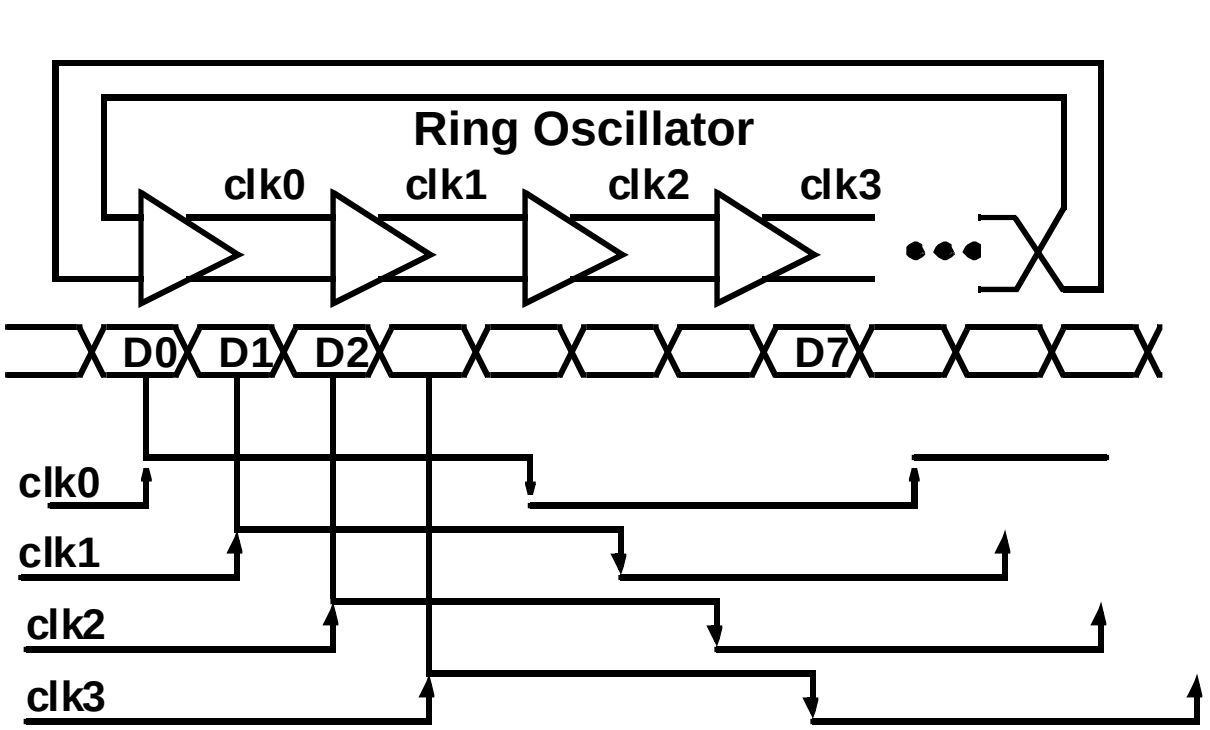

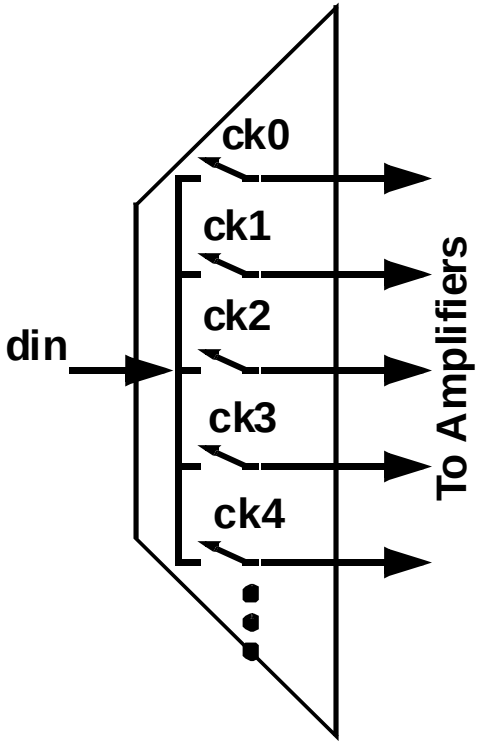

- Use multiple input receivers
  - » Simplest 2, more complex 4-8
  - » Decouples Tbit from latch resolution
  - » Leverage high input impedance amplifiers

Serial Link Signaling Over Backplanes

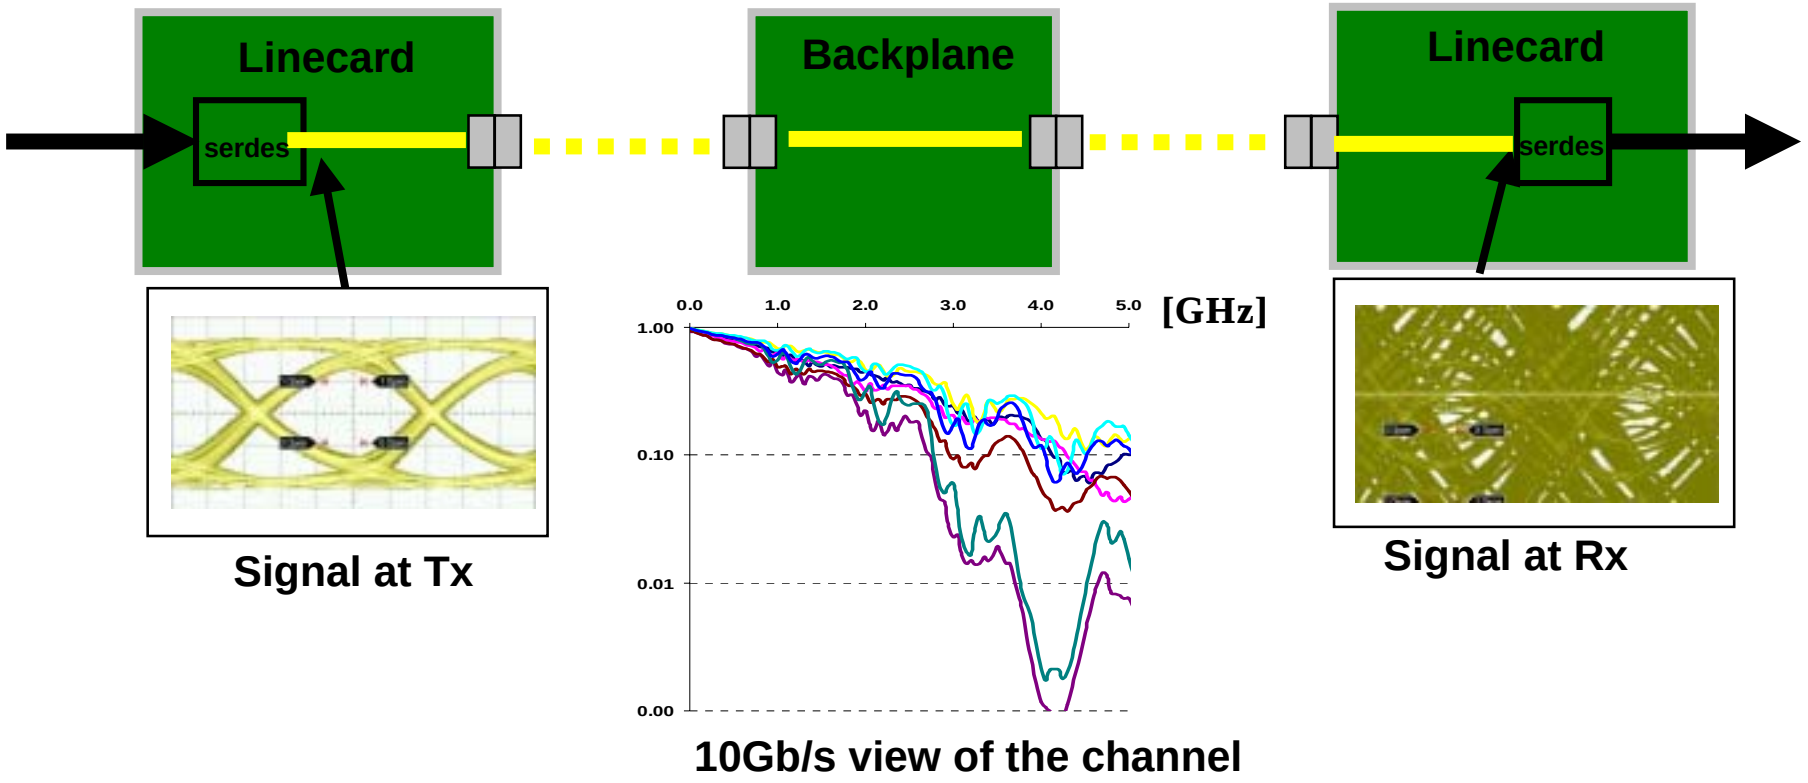

- Now that we've made the fastest Tx & Rx look what happens with the eye
- Need to look more closely into the channel as that seems to be the problem

The Backplane Environment

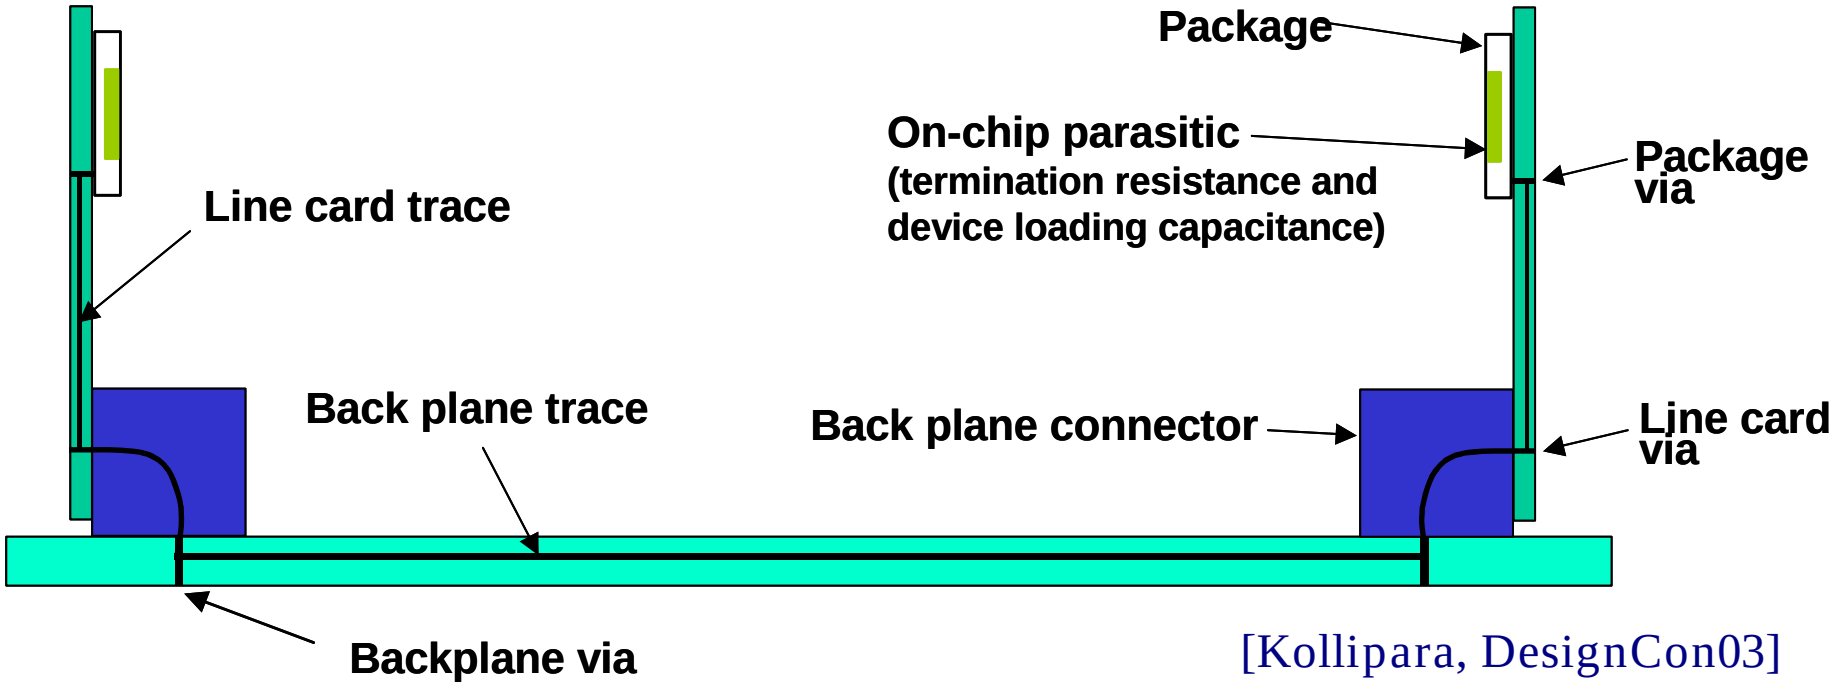

[Kollipara, DesignCon03]

- The problem is there are many sources of Z
  …and thus many possible sources of signal degradation
- Interference
  - » Intersymbol (dispersion, reflections)
  - » Co-channel (crosstalk)

Interference

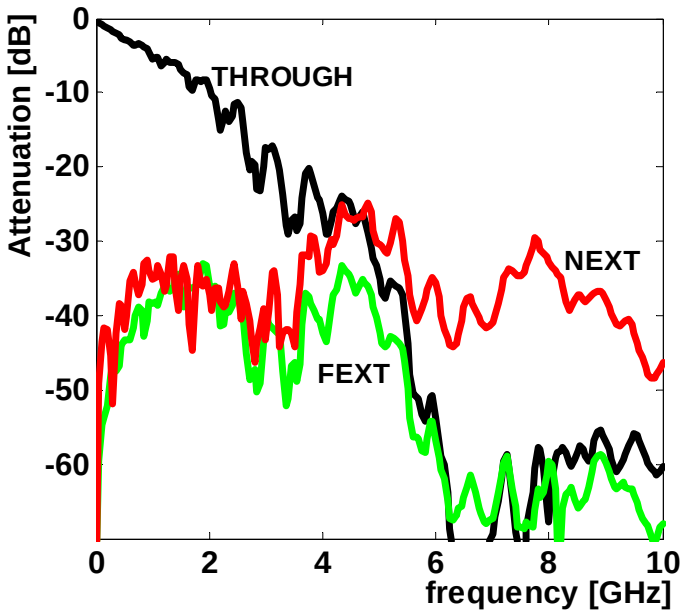

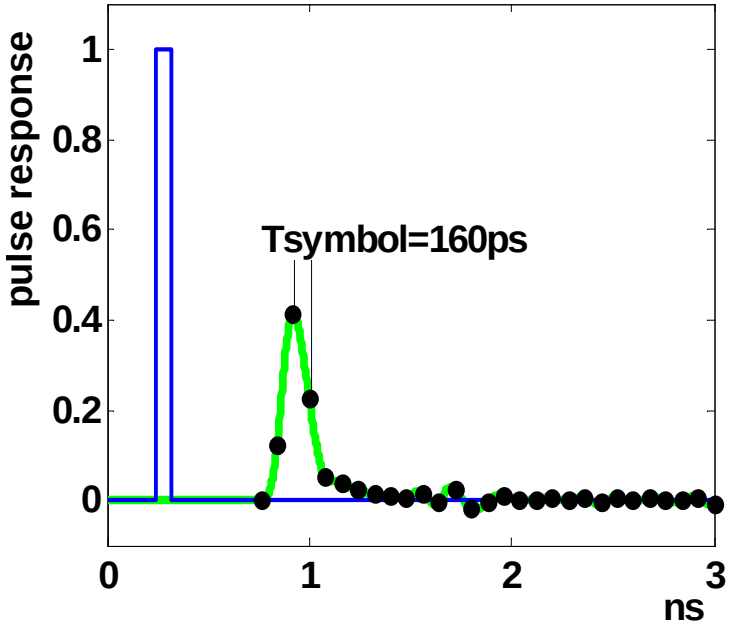

- Inter-symbol interference
  - » Dispersion (skin-effect, dielectric loss) - short latency
  - » Reflections (impedance mismatches – connectors, via stubs, device parasitics, package) – long latency
- Co-channel interference (Far-End & Near-End Crosstalk)

Dispersion: Material Loss

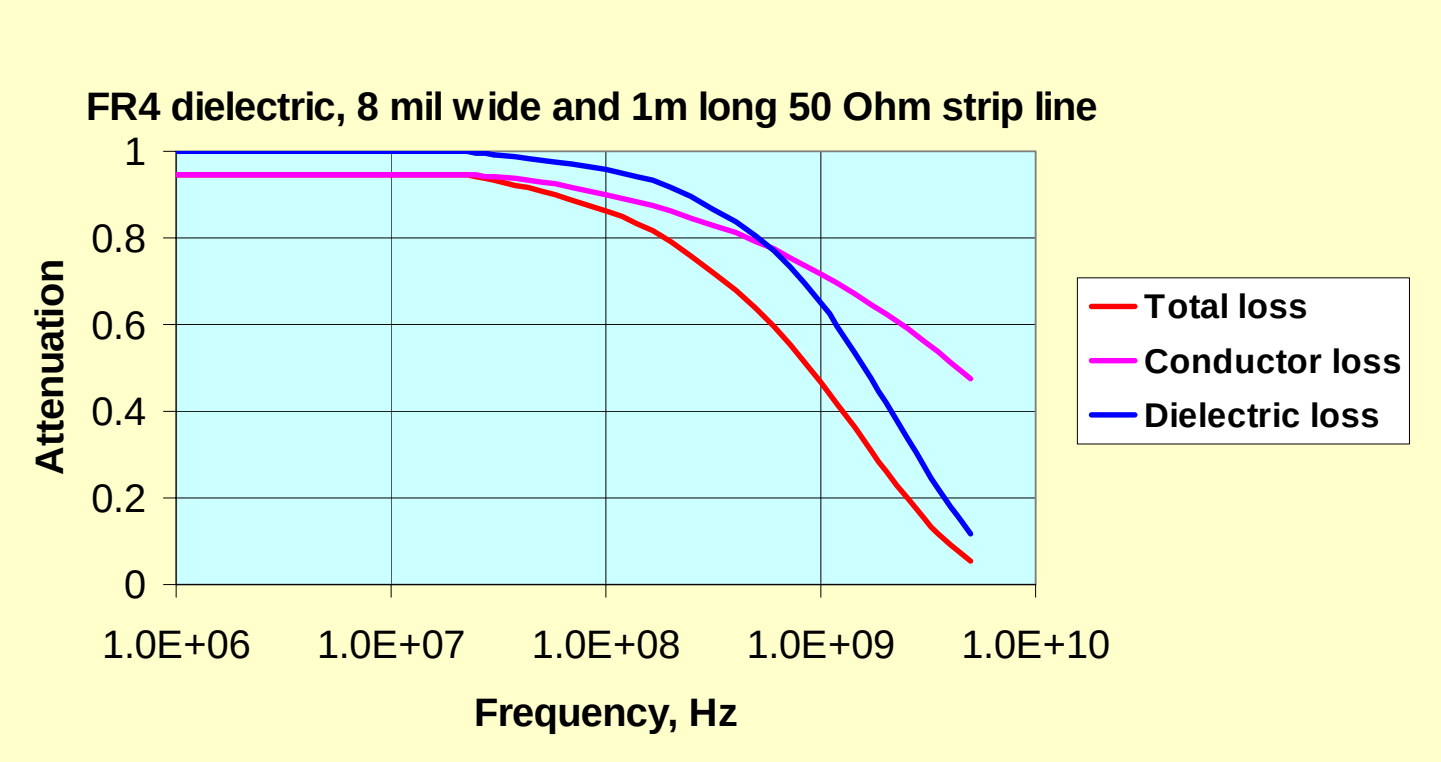

- PCB Loss : skin & dielectric loss
  - » Skin Loss $\propto \sqrt{f}$
  - » Dielectric loss $\propto f$: a bigger issue at higher f

Reflections: Z - Discontinuities

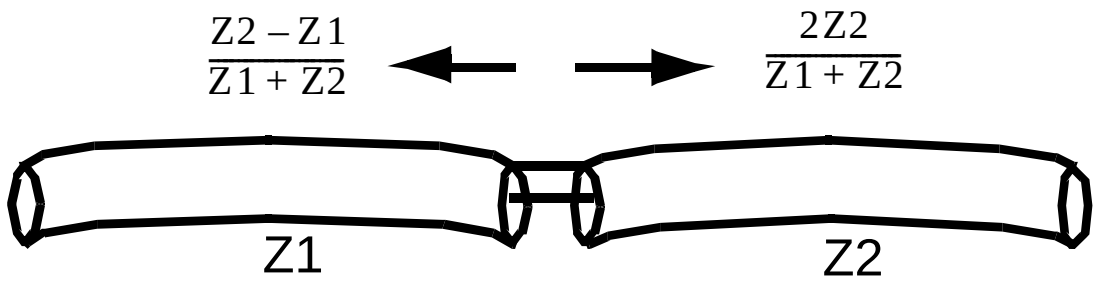

Energy flow into junction = transmitted + reflected energy

- Sources of Reflections : Z - Discontinuities
  - » PCB Z mismatch
  - » Connector Z mismatch
  - » Vias (through) Z mismatch
  - » Device parasitics - effective Z mismatch

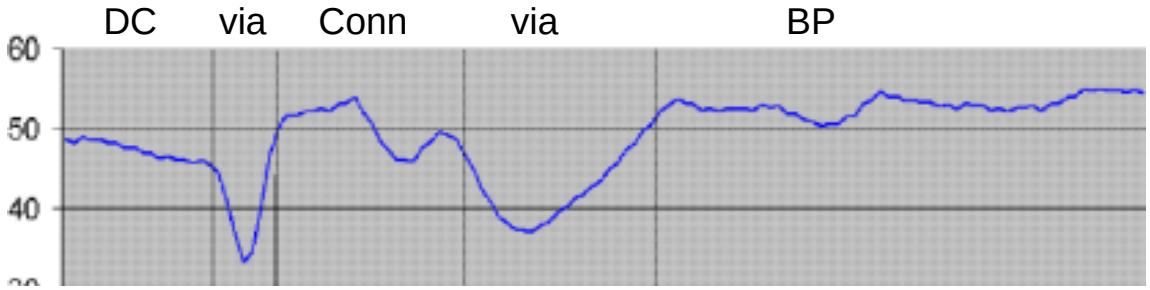

Reflections From Via Stubs

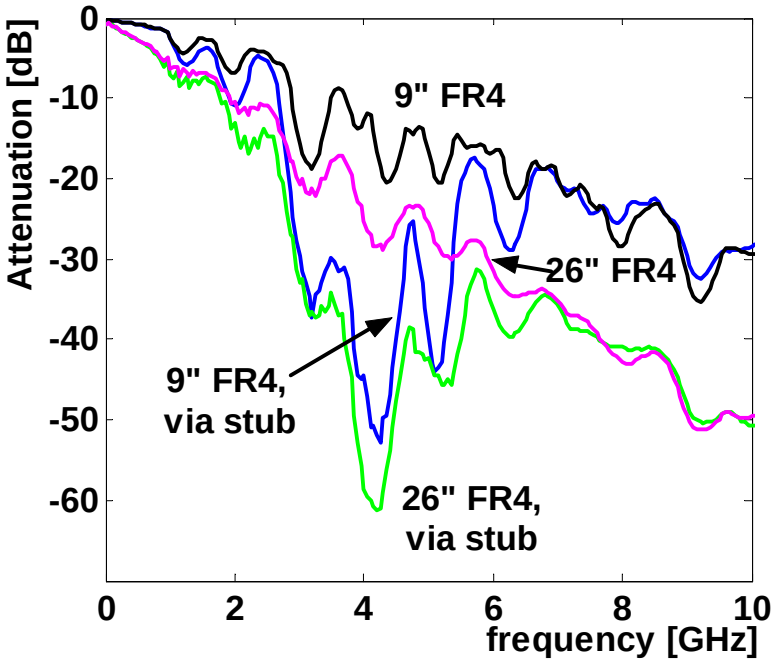

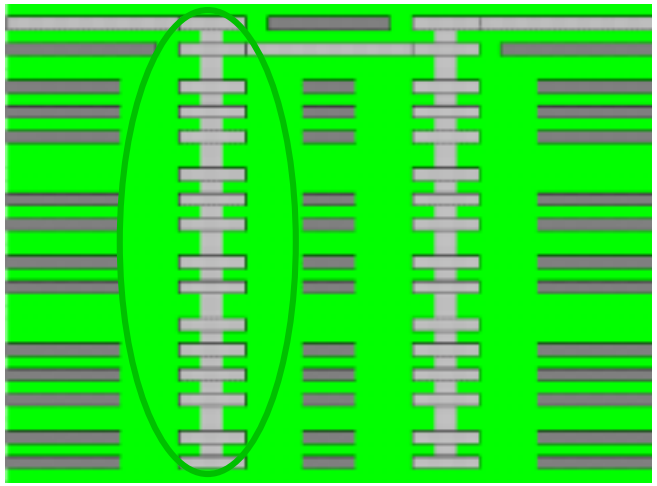

Top layer signaling
results in large via stub

- Additional sources of reflections : stubs
  - » Vias - particularly on thick backplanes
  - » Package plating stubs

Reflections and Crosstalk

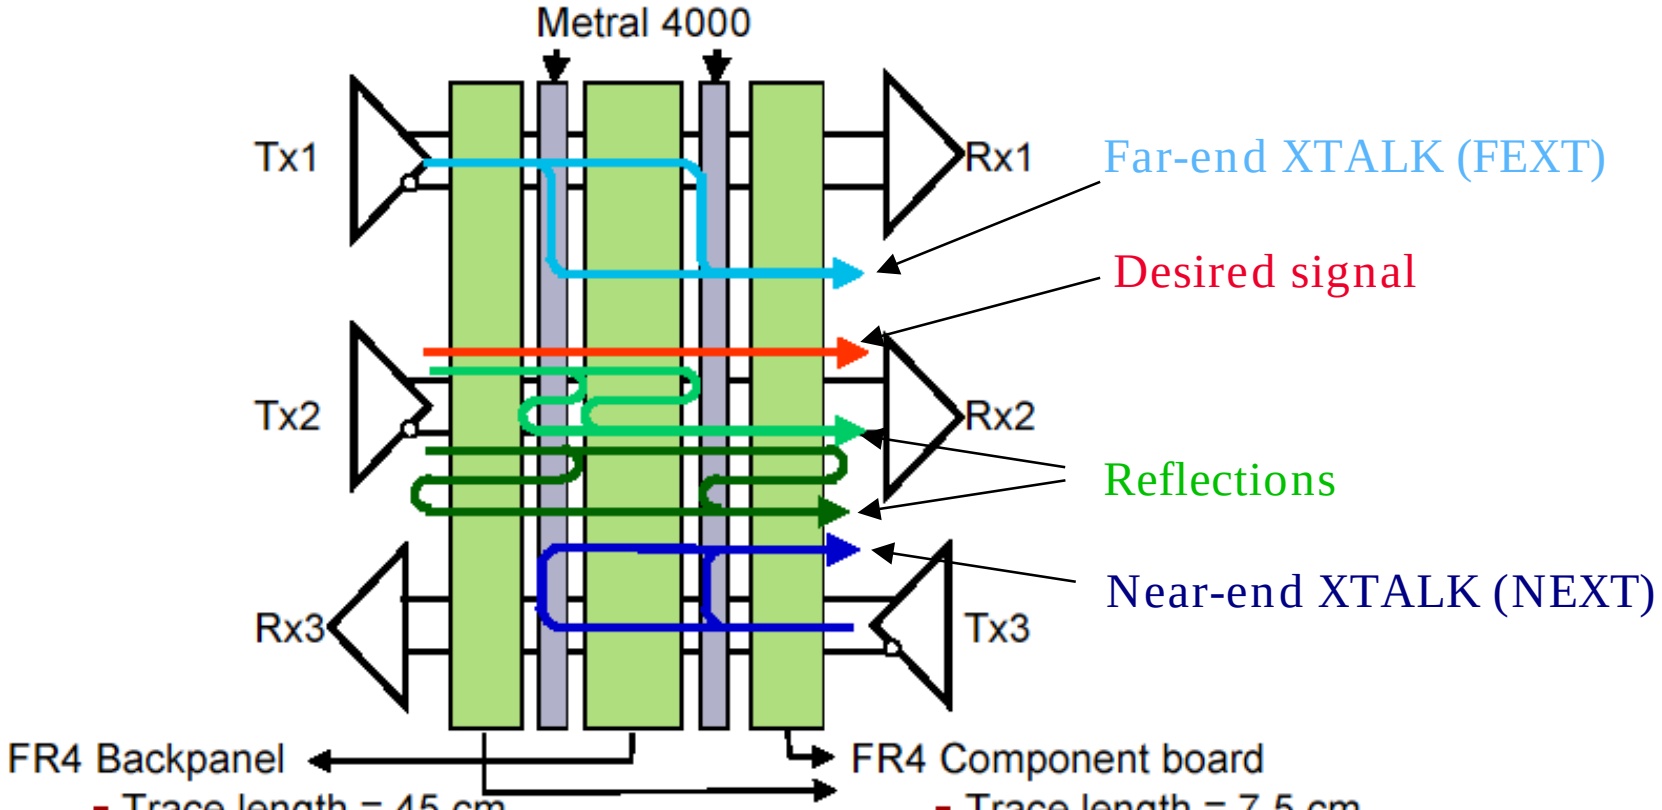

[Sercu, DesignCon03]

Crosstalk

- Many sources
  - » On-chip
  - » Package
  - » PCB traces
  - » Inside connector

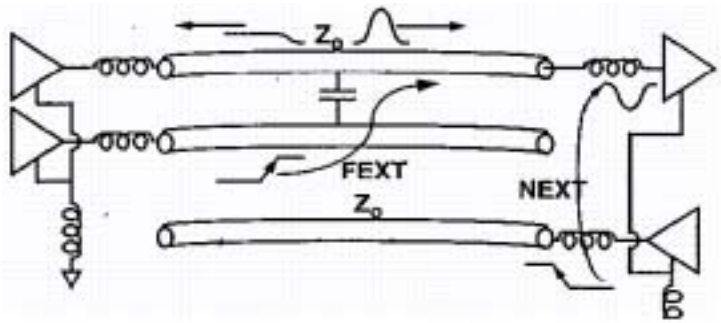

- Differential signaling can help
  - » Minimize xtalk generation & make effects common-mode
- Both NEXT & FEXT
  - » NEXT very destructive if RX and TX pairs are adjacent
    - – Full swing-TX coupling into attenuated RX signal
    - – Effect on SNR is multiplied by signal loss
  - » Simple solution : group RX/TX pairs in connector
  - » NEXT typically 3-6%, FEXT typically 1-3%

A Complex System

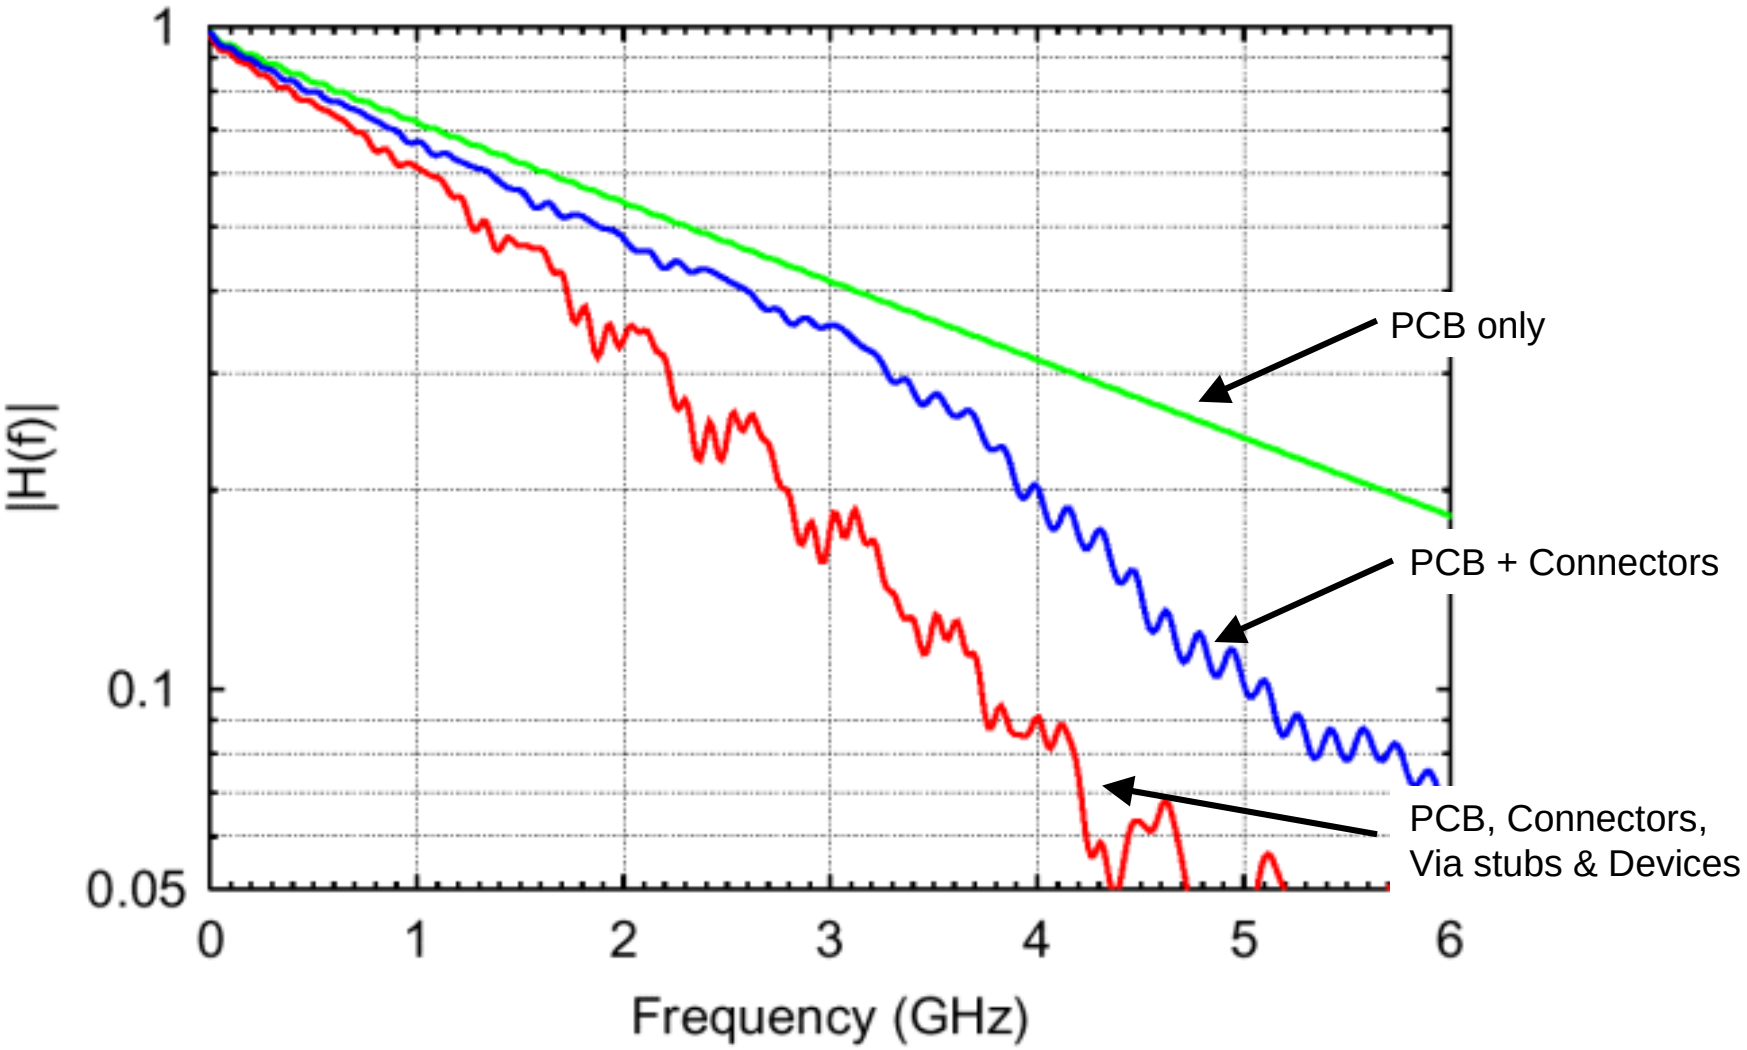

Signaling Faster – System Level Improvements

- Channel designer's view (passive techniques)
  - » Try to make Z-discontinuities go away
  - » Reduce cross-talk (EM isolation)
- IC designer's view (active techniques)
  - » Design circuits that compensate/eliminate interference

Equalization For Loss : Goal is to Flatten Response

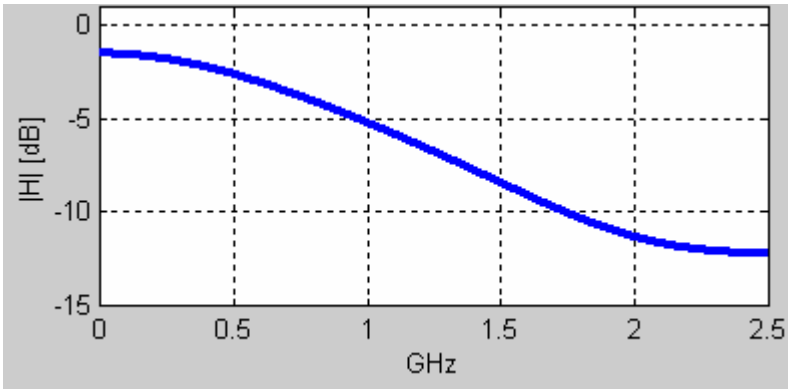

X

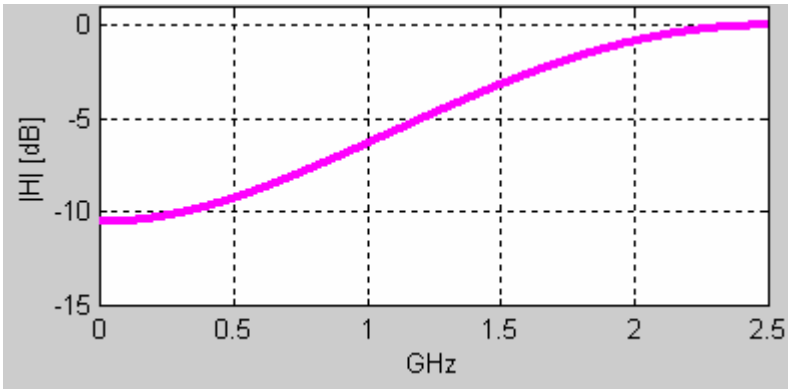

=

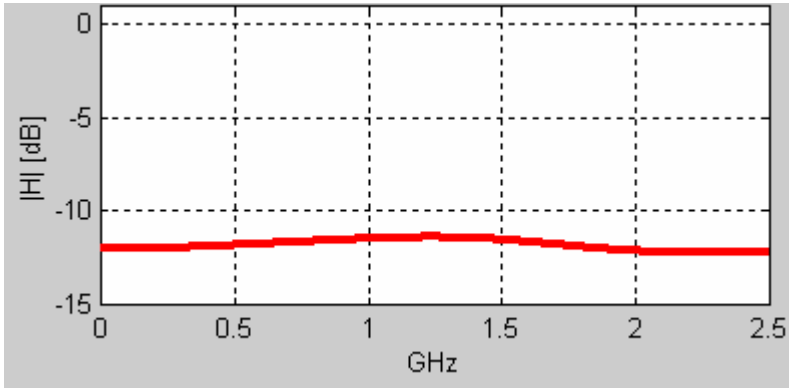

- Channel is band-limited
- Equalization : boost high-frequencies relative to lower frequencies

Receive Linear Equalizer

- Amplifies high-frequencies attenuated by the channel
- Pre-decision
- Digital or Analog FIR filter
- Issues
  - » Amplifies noise
  - » Precision
  - » Tuning delays (if analog)
  - » Setting coefficients
    - – Adaptive algorithms such as LMS

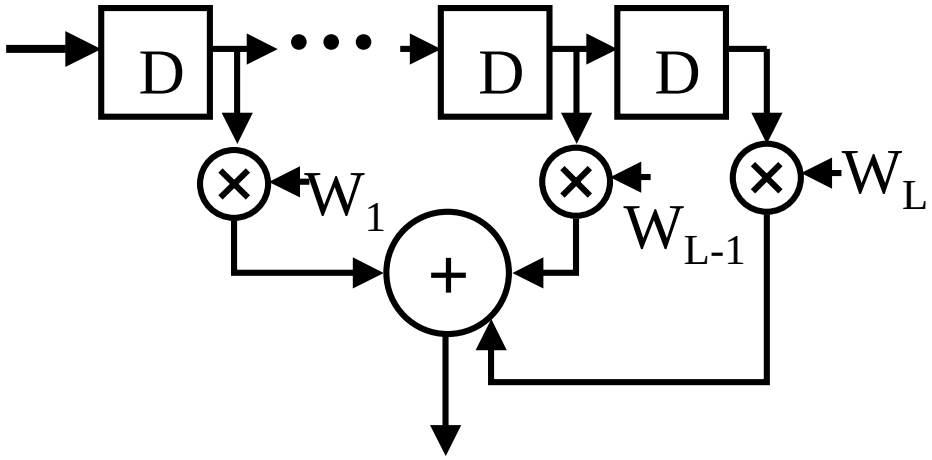

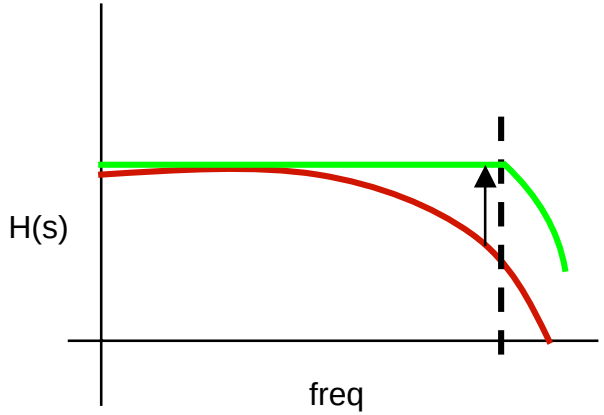

Transmit Linear Equalizer

- Attenuates low-frequencies
  - » Need to be careful about output amplitude : limited output power
    - – If you could make bigger swings, you would
    - – EQ really attenuates low-frequencies to match high frequencies Also FIR filter : D/A converter
- Can get better precision than RX
- Issues
  - » How to set EQ weights?
  - » Doesn't help loss at f

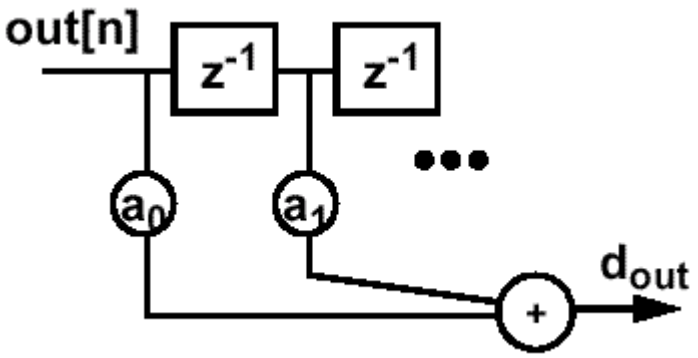

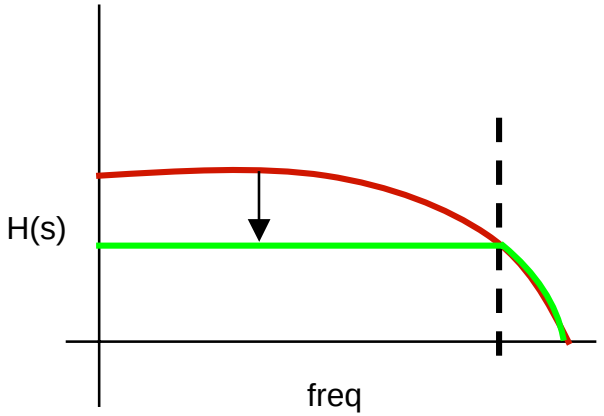

Transmit Linear Equalizer: Single Bit Operation

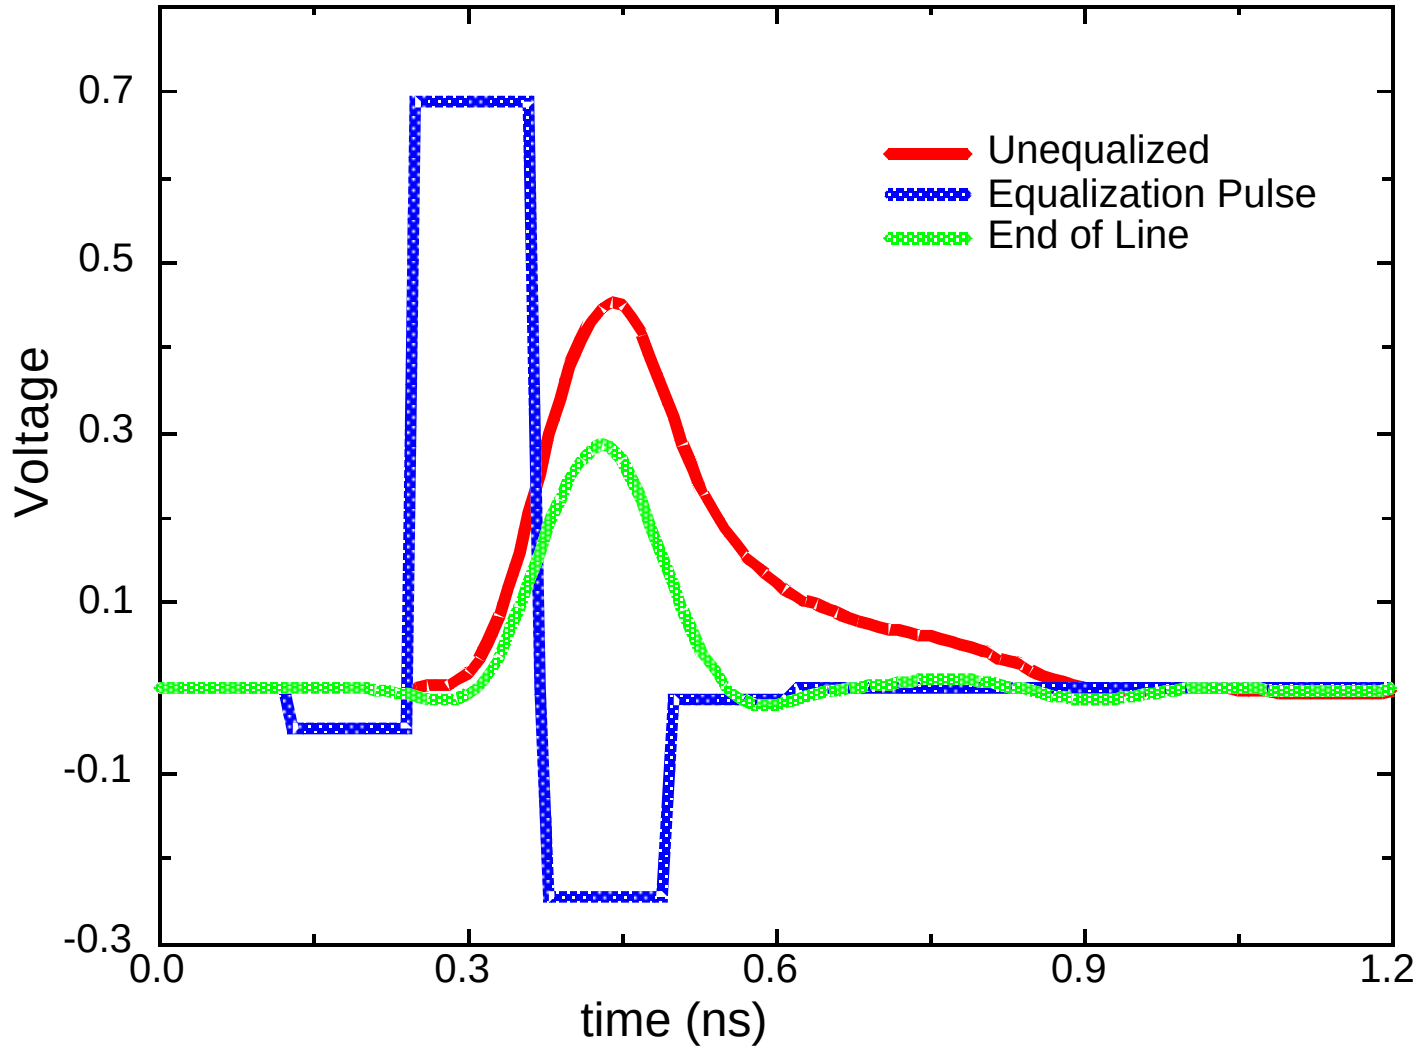

Example : 5Gbps over 26” FR4

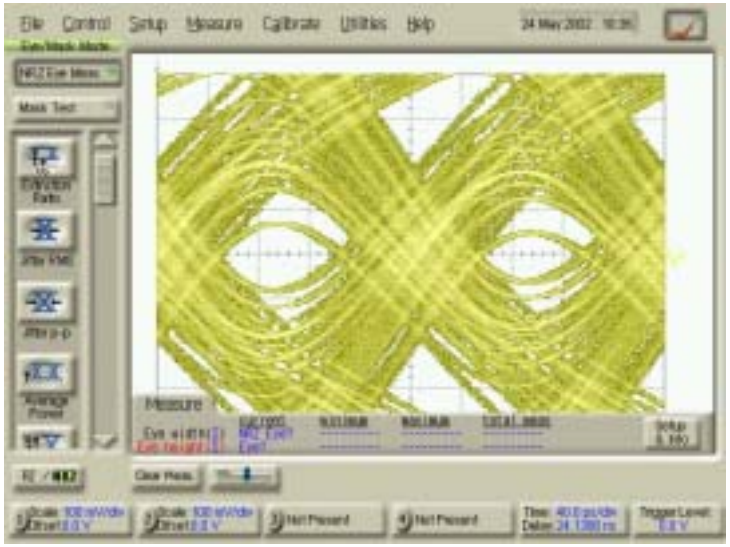

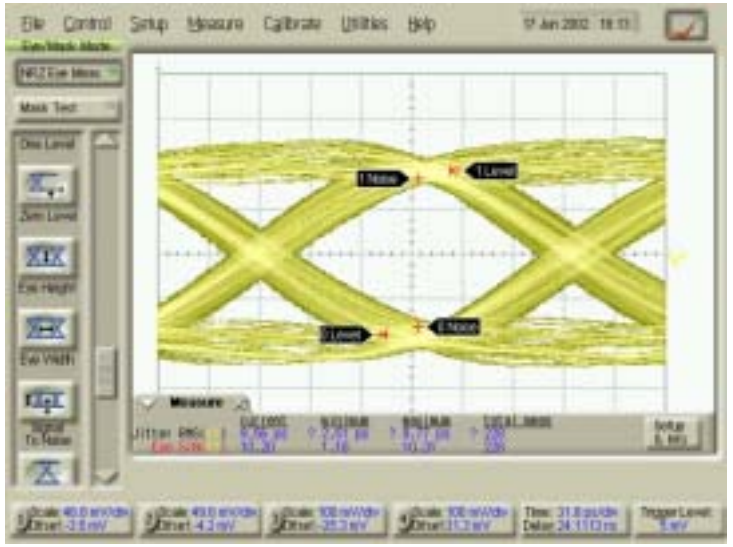

no equalization

with Tx linear equalizer

Decision Feedback Equalization

- Don't invert channel…
  just remove ISI
  - » Know ISI because already received symbols
  - » Doesn't amplify noise
  - » Has error accumulation problem
    - – Less of an issue in links where random noise small
- Requires a feed-forward equalizer for precursor ISI
  - » Reshapes pulse to eliminate precursor

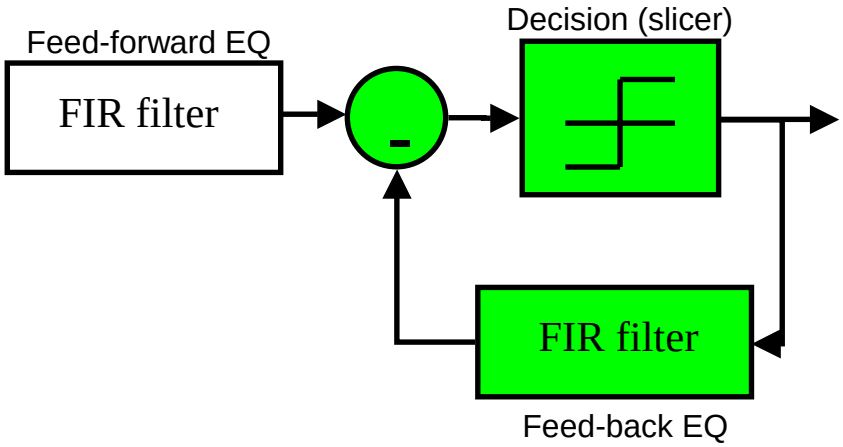

Transmit and Receive Equalization

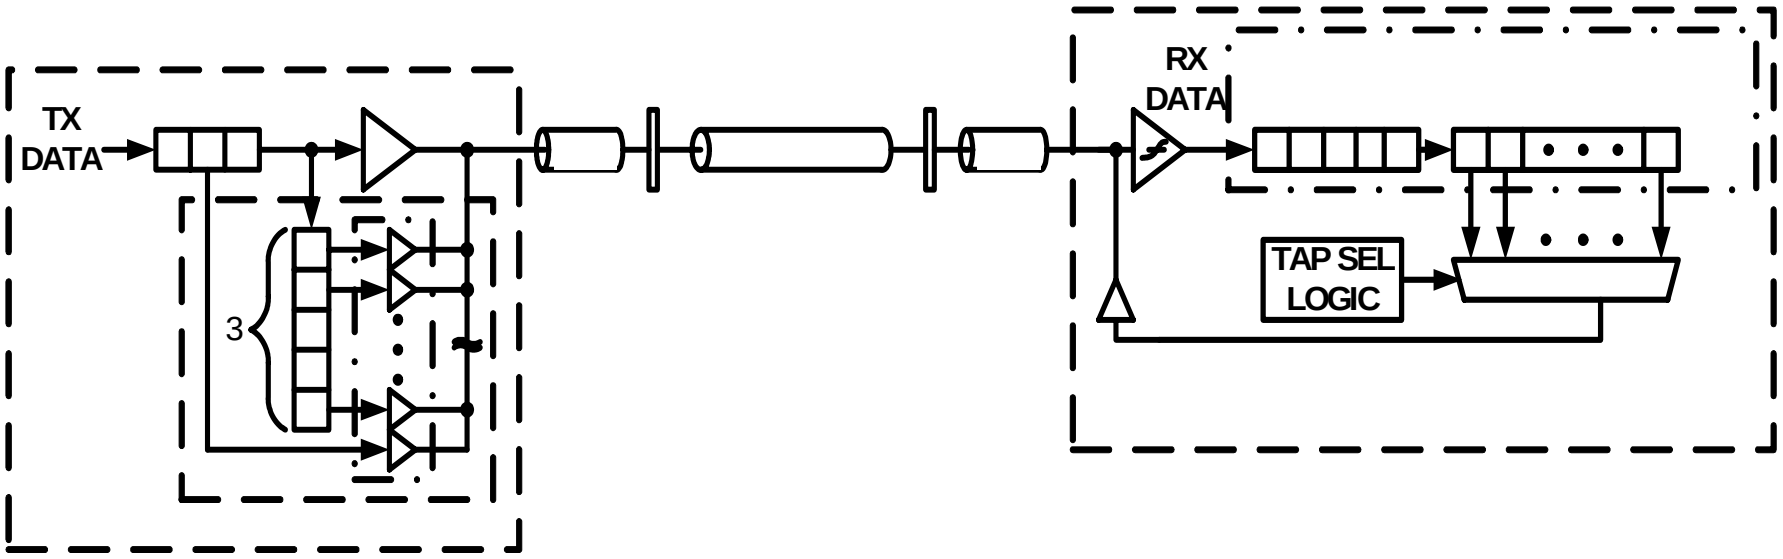

- Transmit and receive equalizers are combined to make a range restricted DFE
  - » Tx equalizer functions as the feed-forward filter
  - » Rx equalizer restricted in performance of loop

Tx & Rx Equalization Ranges

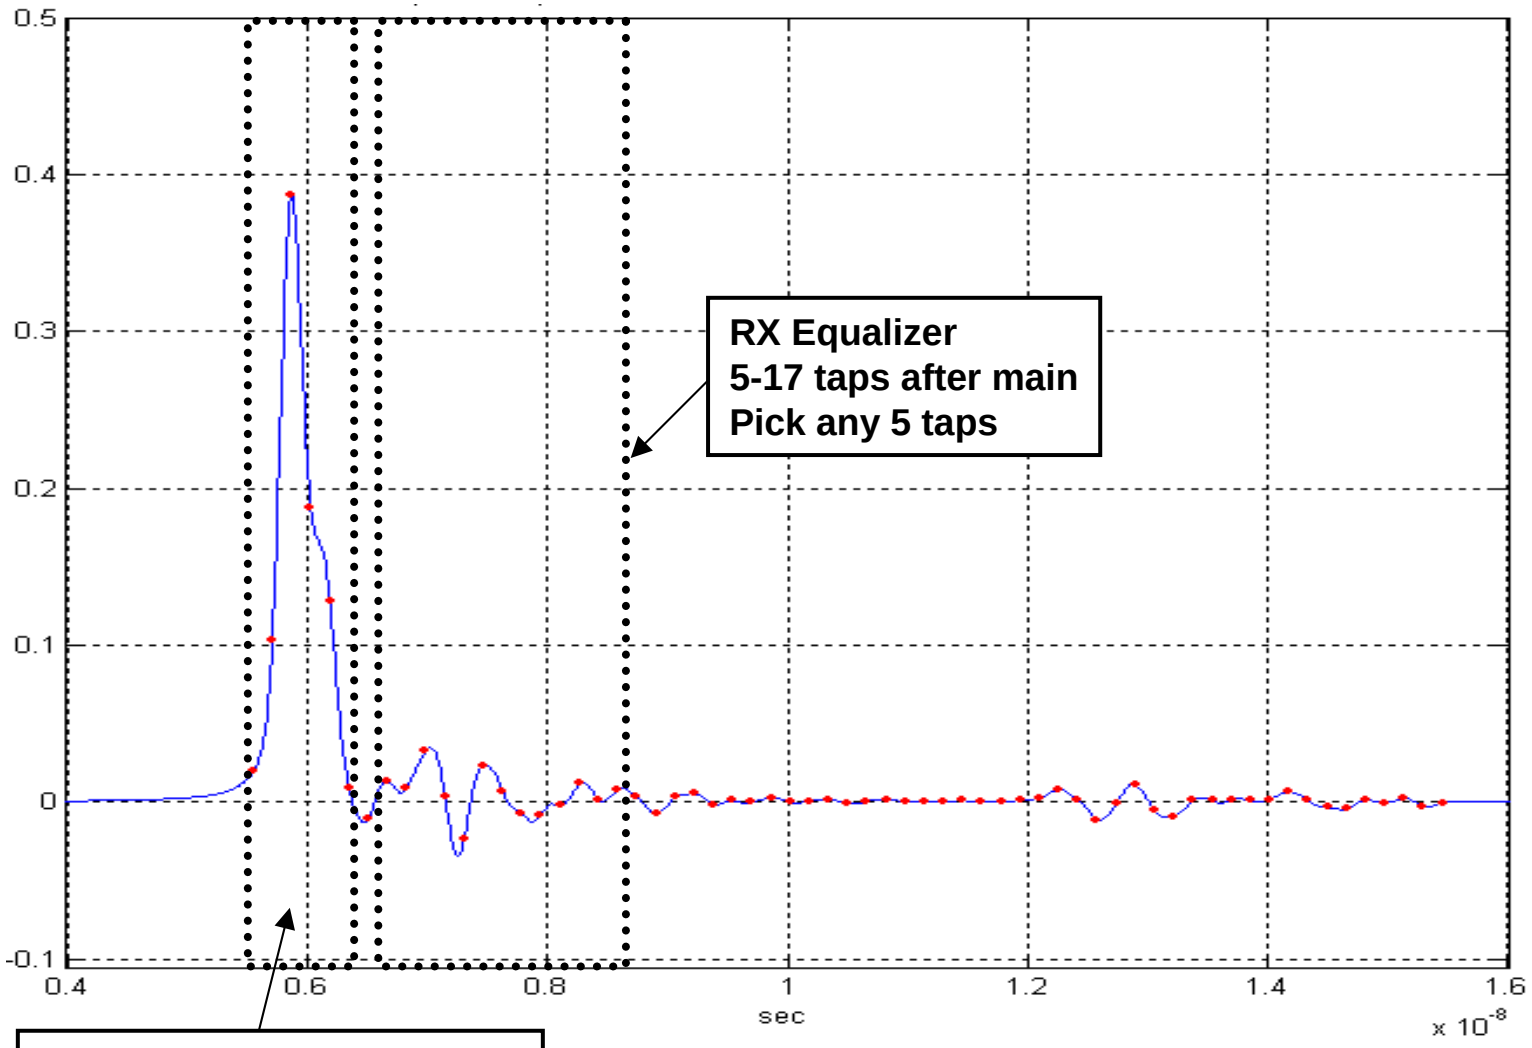

Minimizing Reflections : The Vias

- Minimizing via stubs
  - » Thinner PCBs are better…
    but sometimes impossible
  - » Counter-boring
  - » Blind vias
  - » SMT technology

  - » All are costly
    1.1x - 2x

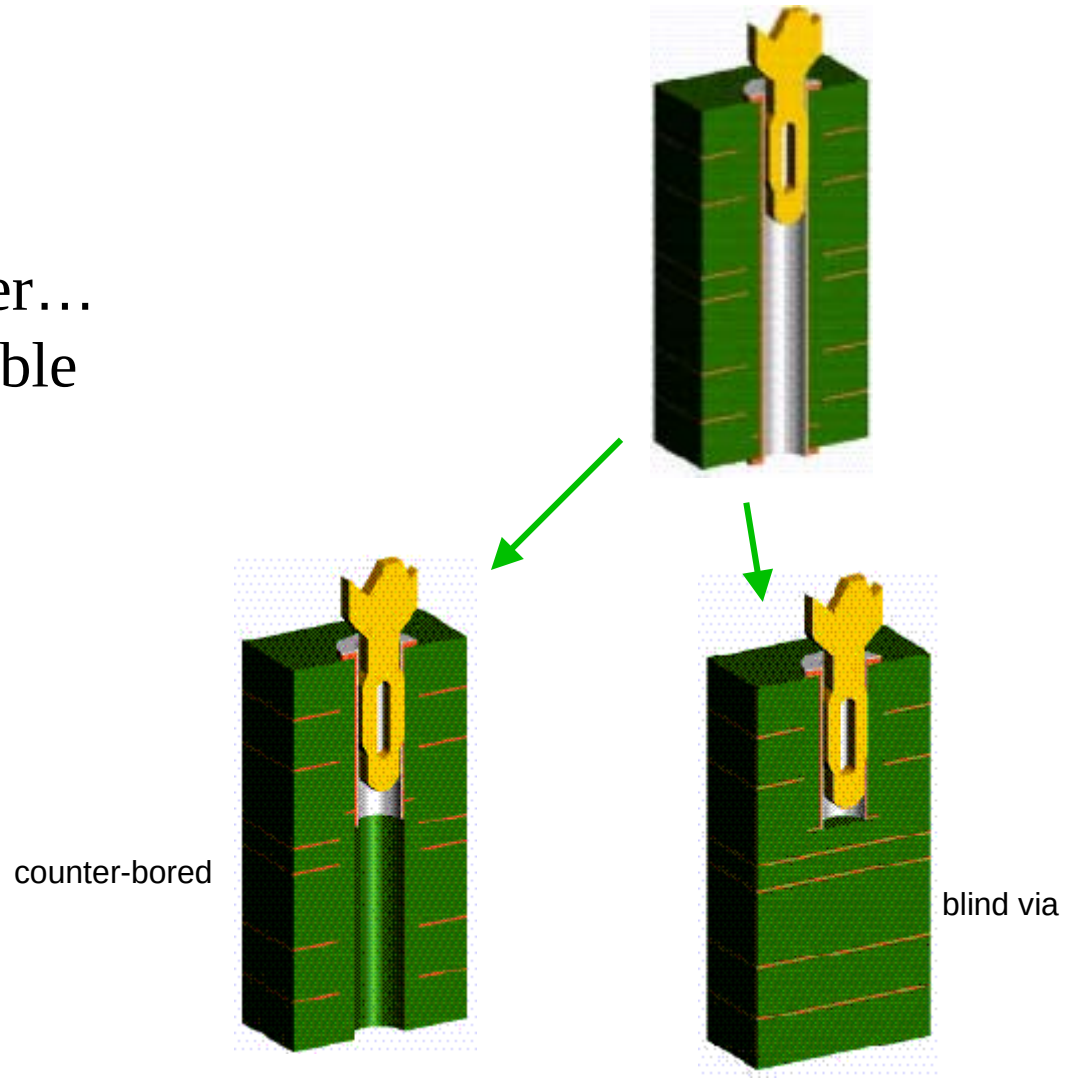

Vias : Effect of Counter-boring

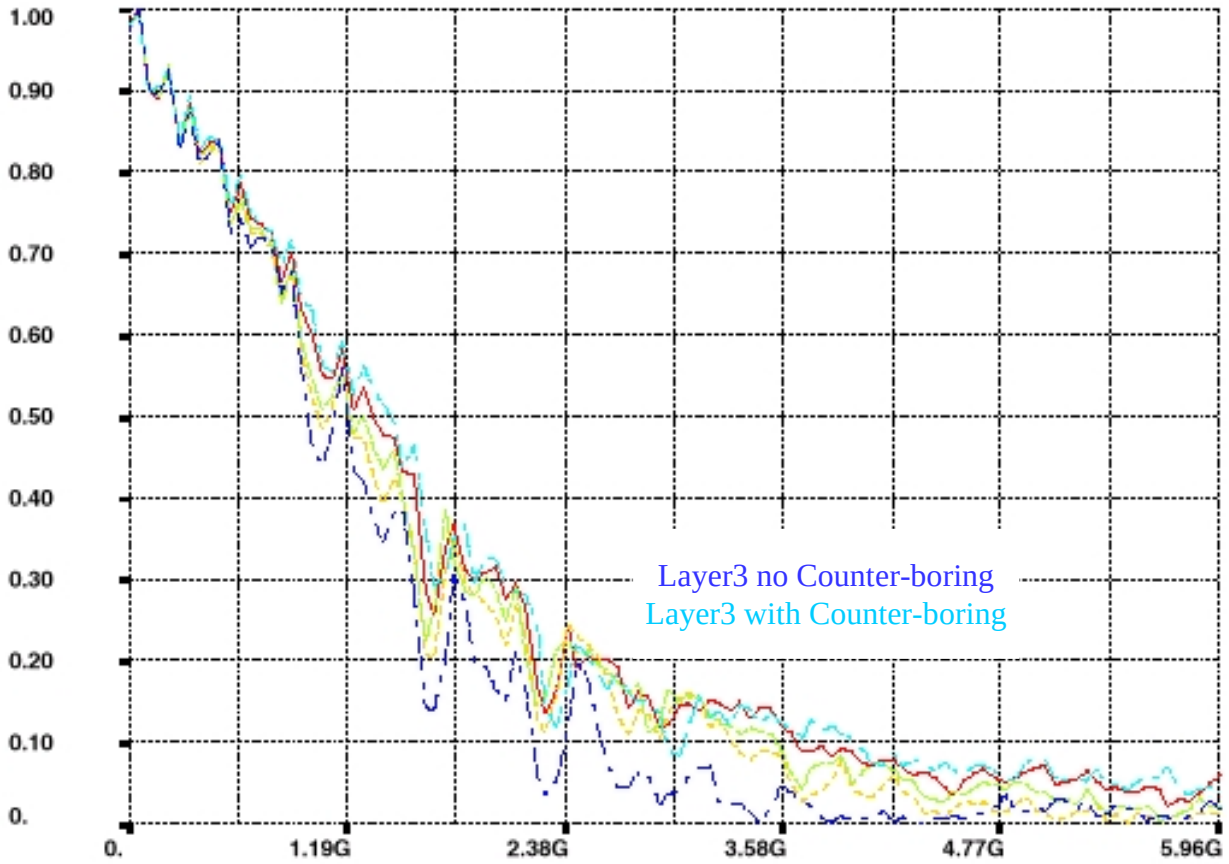

- Counter-boring top layer takes it from highest-loss to lowest-loss & reduces resonance

Minimizing Reflections: Termination Design

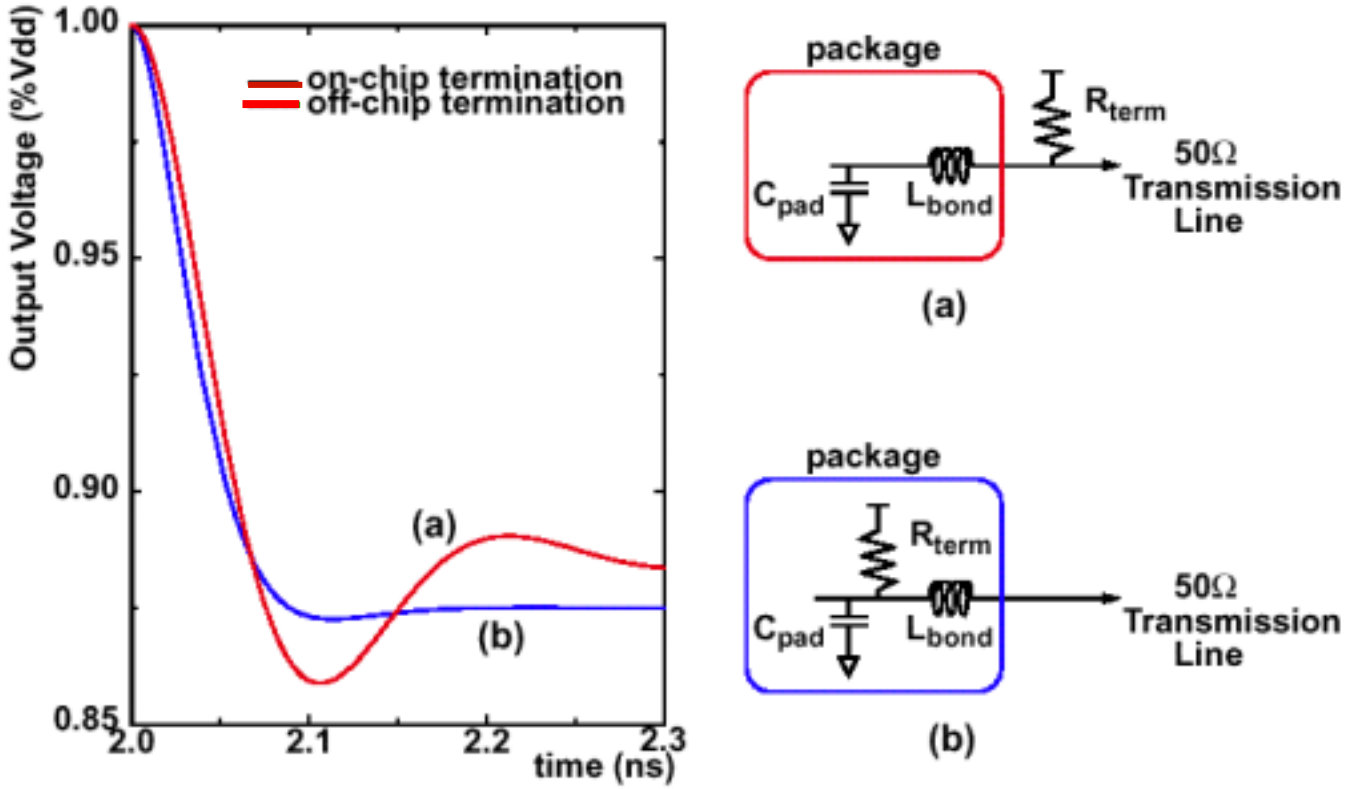

- On-chip termination
  - » Bondwire & pad capacitance part of the channel
    … instead of a stub (which rings)

Minimizing Reflections: FET Terminations

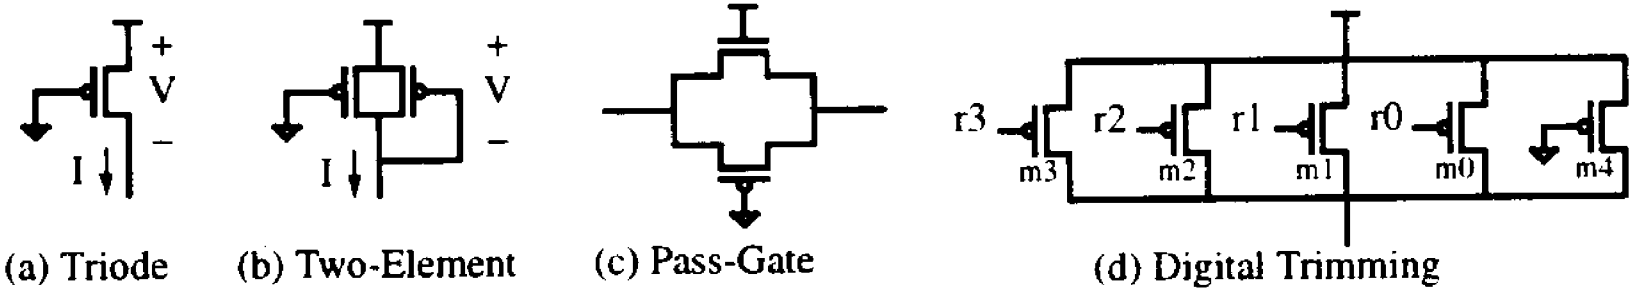

(a) Triode (b) Two-Element (c) Pass-Gate (d) Digital Trimming

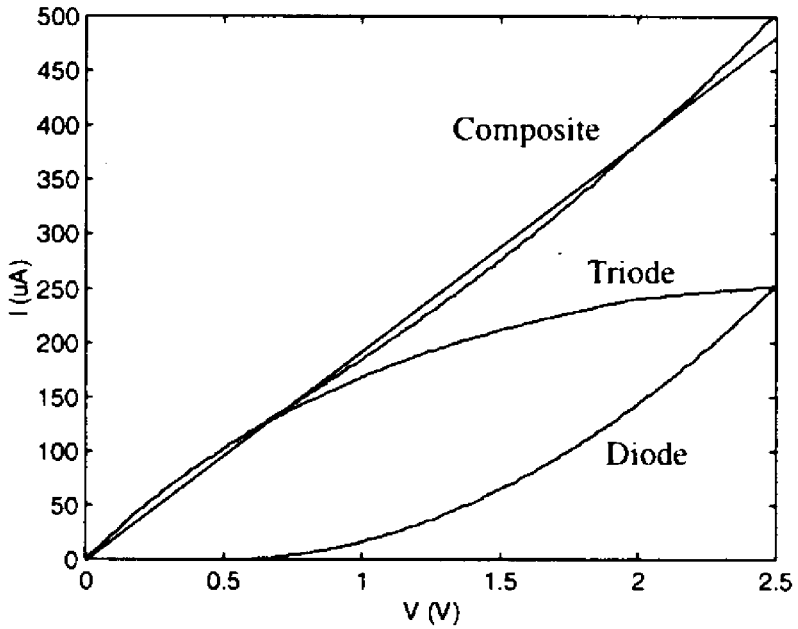

IV-characteristic
of two-element resistor

[Dally]

Alternate Approaches: Multi-Level Signaling

- Binary (NRZ) is 2-PAM
- 2-PAM uses 2-levels to send one bit per symbol
- Signaling rate = 2 x Nyquist

- 4-PAM uses 4-levels to send 2 bits per symbol
- Each level has 2 bit value
- Signaling rate = 4 x Nyquist

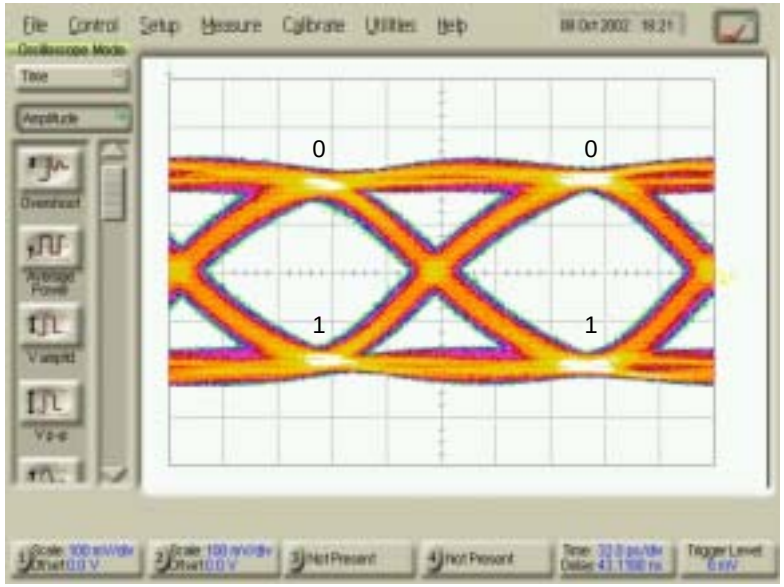

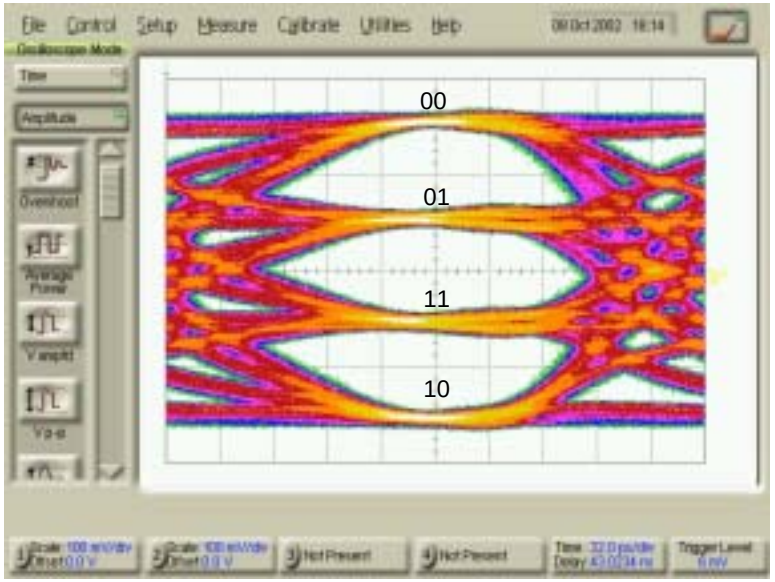

Note : both can be either single-ended or differential

When Does 4-PAM Make Sense?

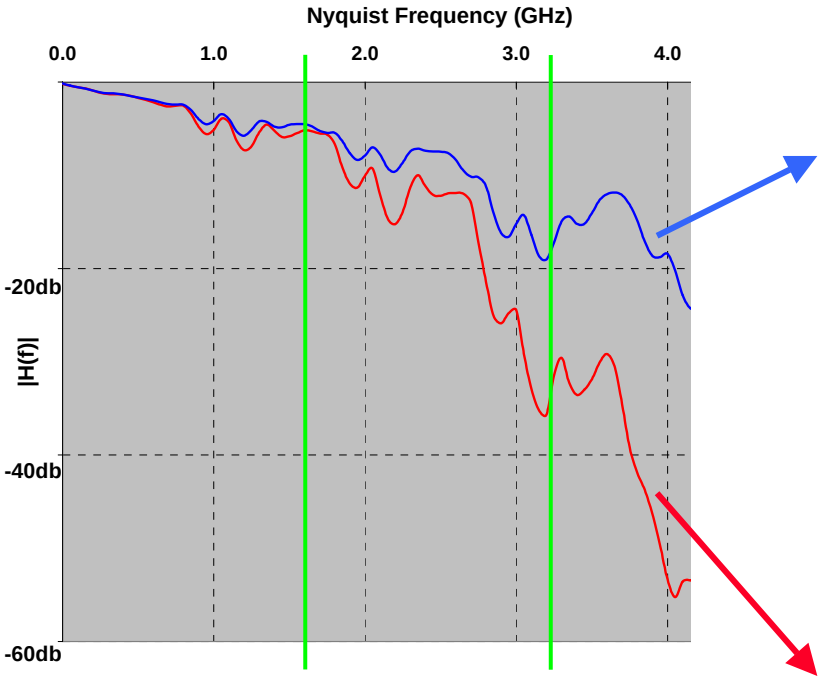

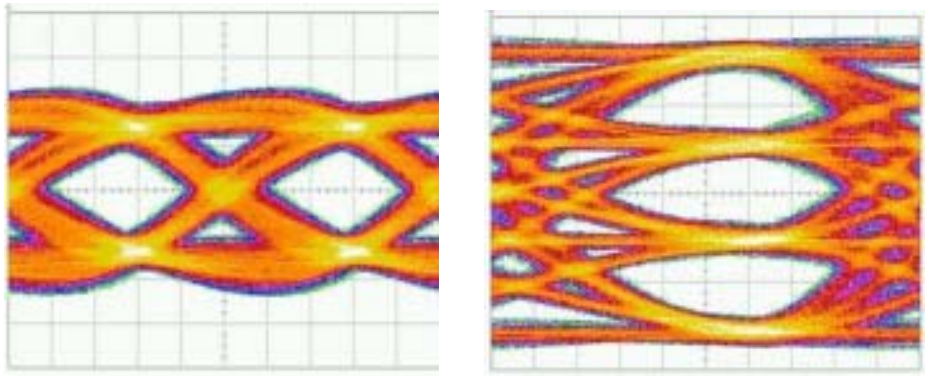

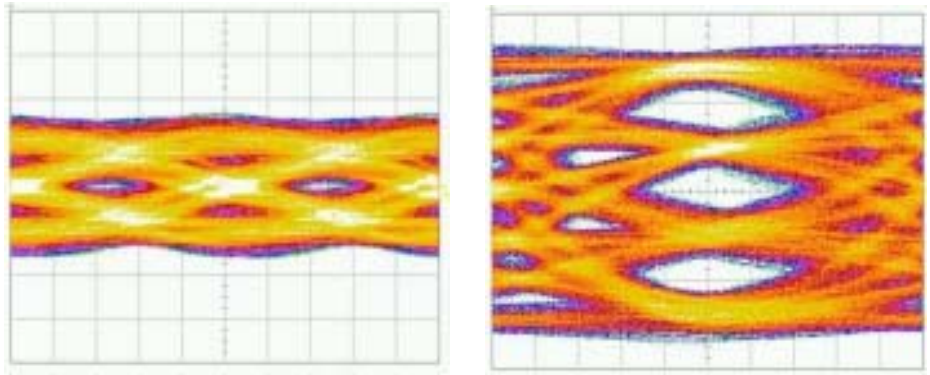

- First order : slope of S21
  - » 3 eyes : 1 eye = 10db
  - » loss > 10db/octave : 4-PAM should be considered

Alternate Approaches: Simultaneous BiDirectional

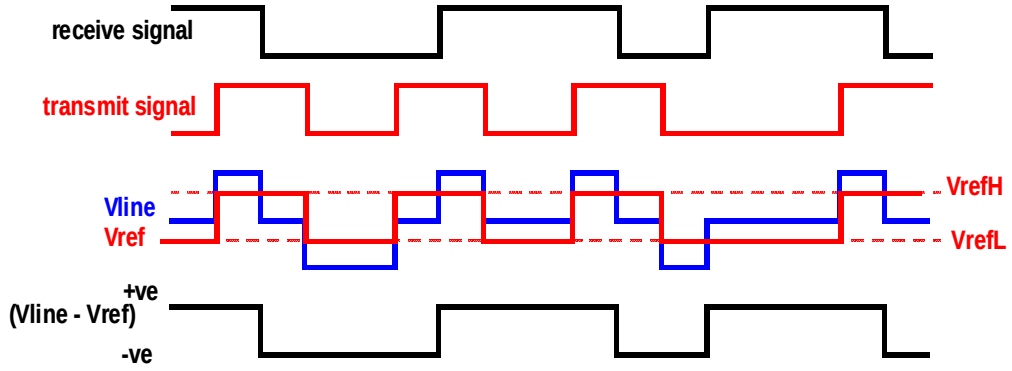

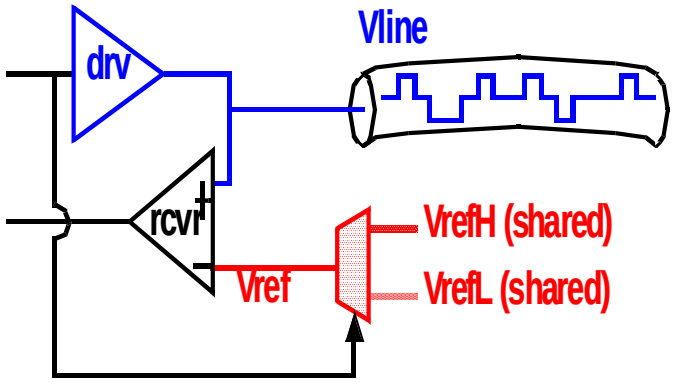

- Two signals at half speed
  - » Makes sense if b/w need equal in both directions
- Issues
  - » Getting ideal timing between TX & RX is tough

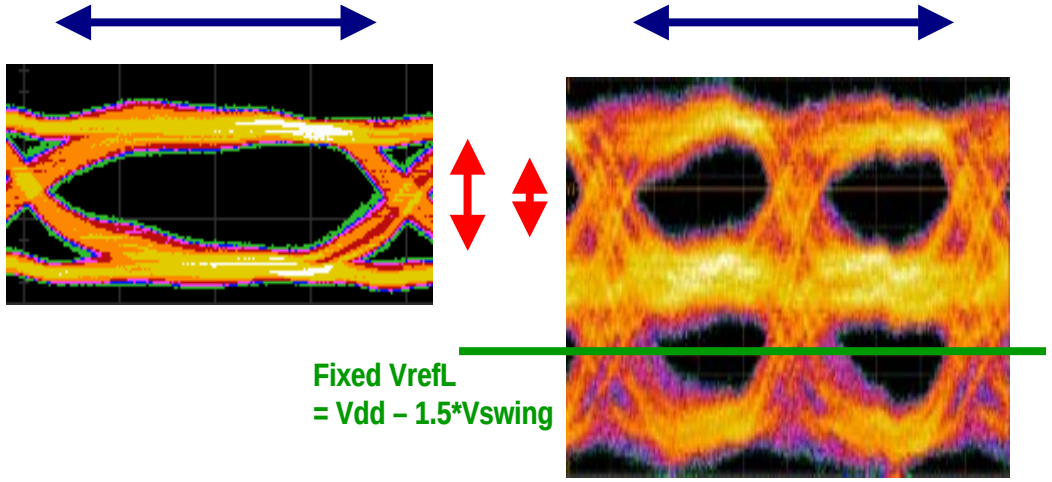

Characterization System

- Multiple
  - » Connectors
  - » Backplane materials
  - » Trace lengths
  - » Layers/via lengths
  - » Via technology
- These slides
  - » 20” Trace length
  - » FR4 non counter-bored
  - » Nelco 6000 2-step counter-bored
  - » Top & bottom layers

- Will show the Rambus 10Gb/s backplane SerDes demo on Friday

An attempt to shift the problem to DSP side

- 8-way DAC (8bit) and ADC (4bit)
- 8GSa/s
- A lot of power (not even including the DSP section)
- DACs and ADCs complex – a lot of parasitic filtering – channel degradation
- Still people are moving in that direction – check out K. Poulton's 20GSa/s 8-bit ADC paper at ISSCC03

Time-Interleaved DACs

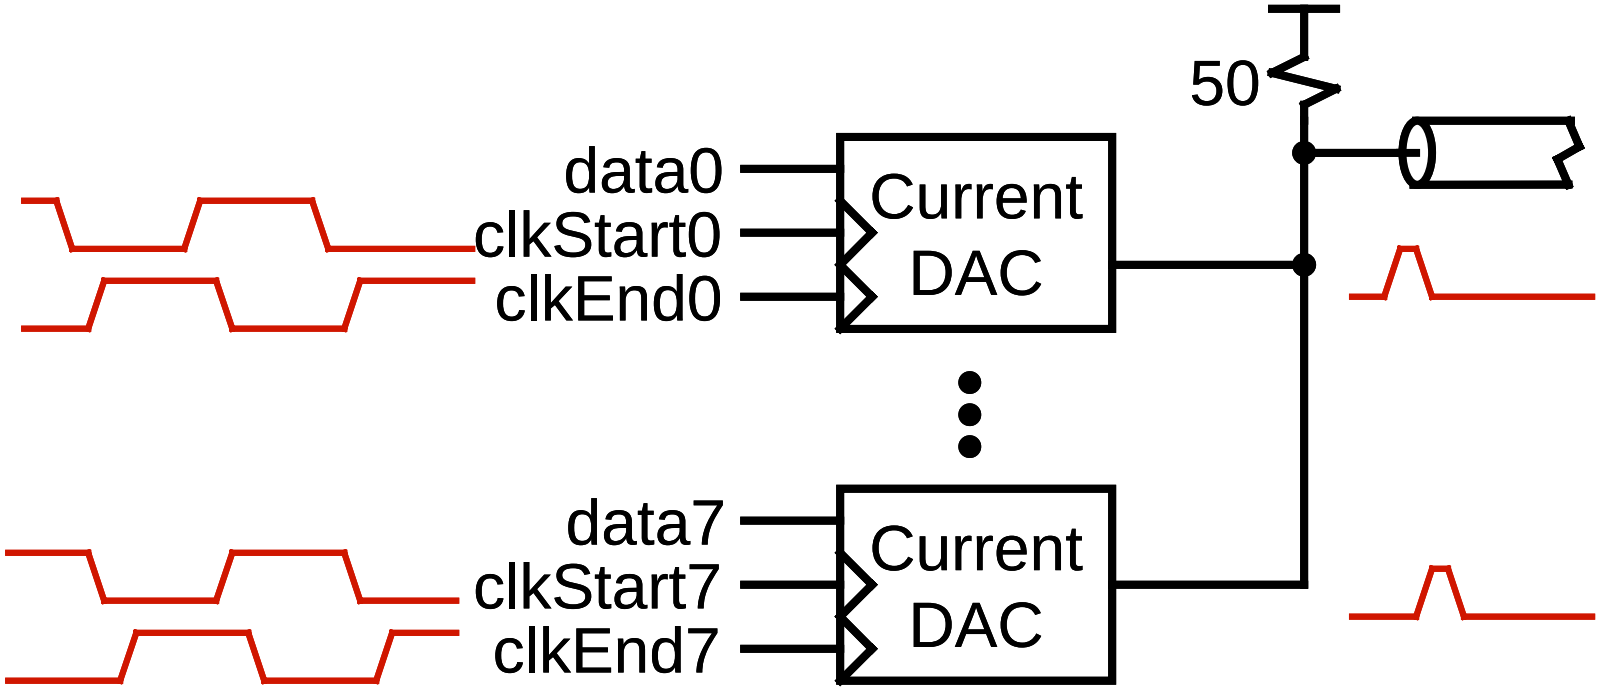

- DACs enabled by overlap of two 1 GHz clocks
  - » Need precise clocks: $3\%_{pp}$ phase noise=>$24\%_{pp}$ symbol
  - » Fast clocks (period of 8 gate delays) limit interleaving
  - » Capacitance of all 8 DACs loads output

DAC Output Circuitry

- RC_{out} = 25Ω * 4.3pF –> 1.5 GHz bandwidth

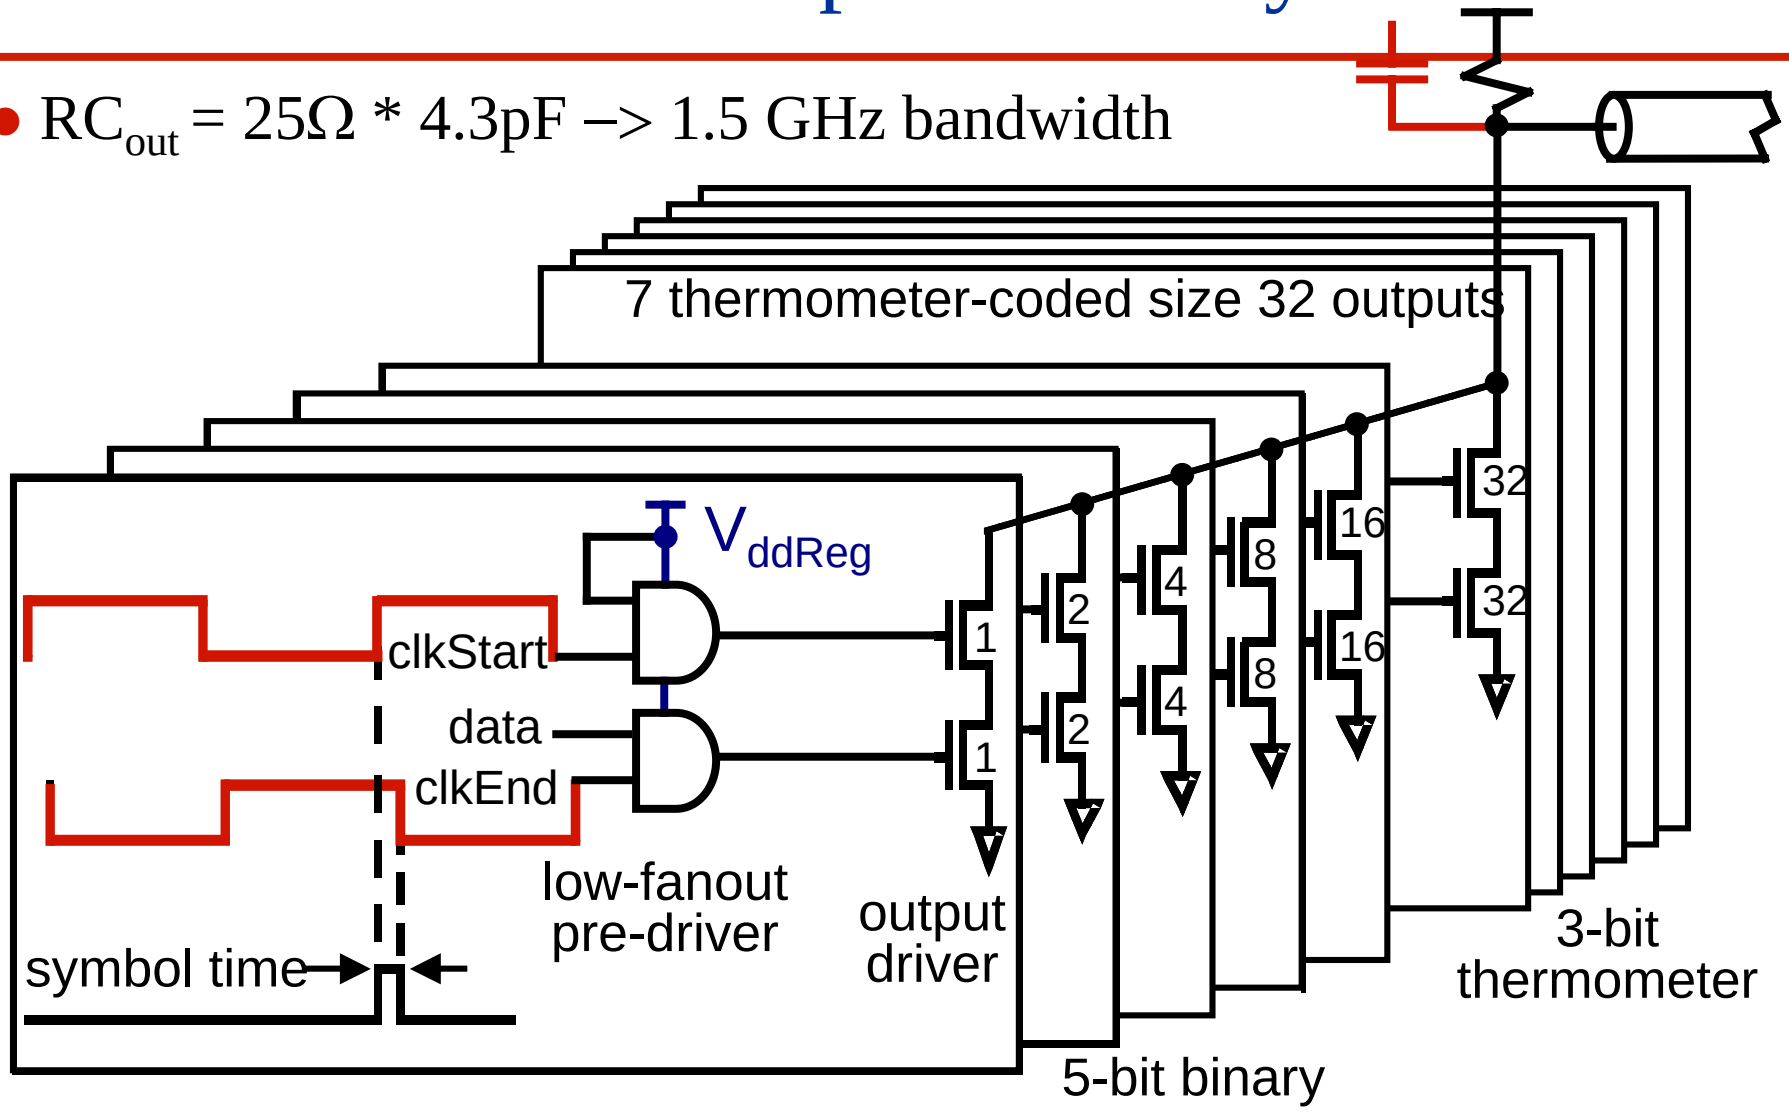

- Predriver V_{ddReg} controls output current

Transceiver Inductors and Clocking

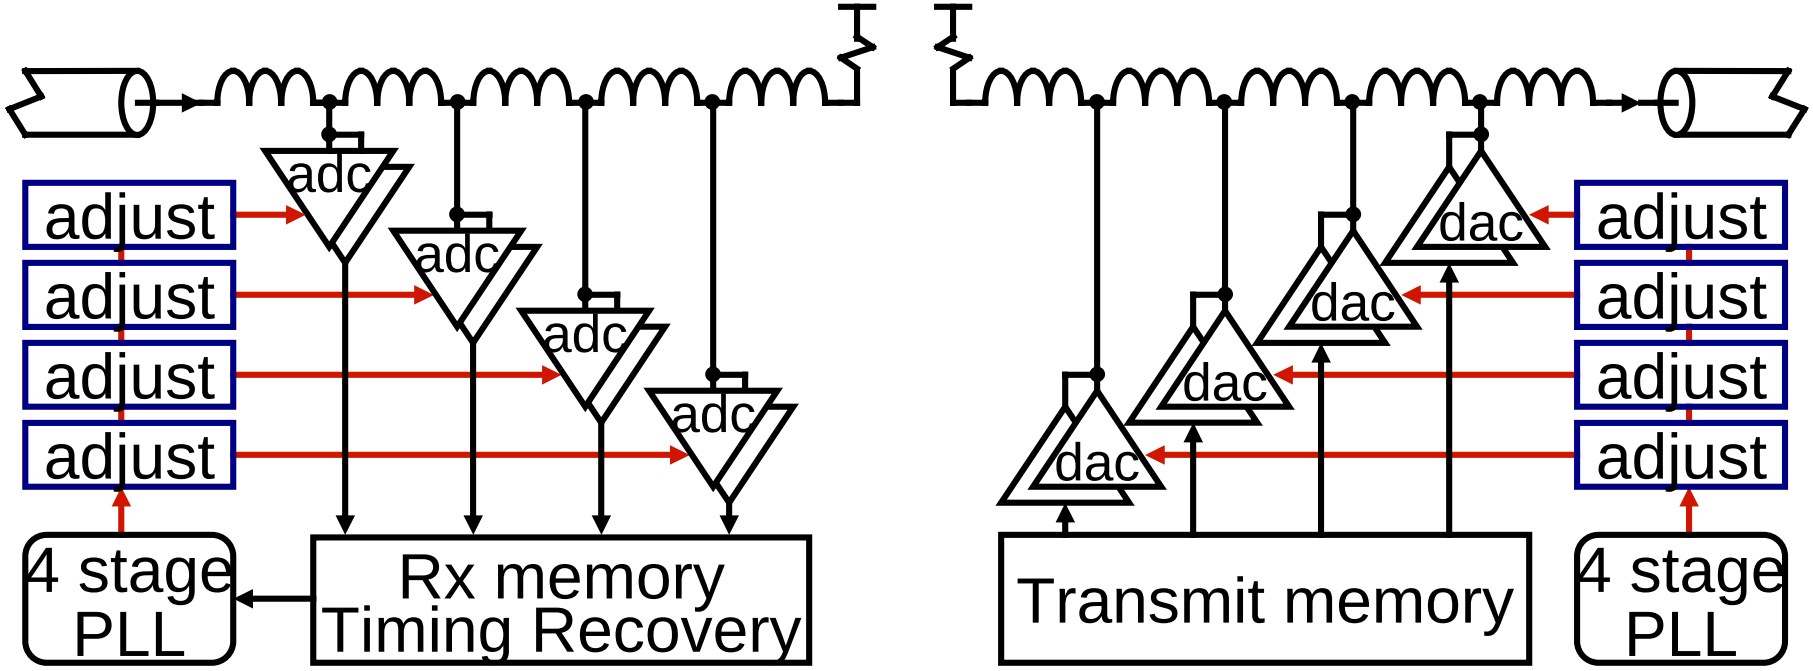

- Phase adjusters correct LC delay, static errors
  - » Adjuster: clock mux, 1/16th-symbol clock interpolator
  - » 8 ADC phase adjusters + 1 for timing recovery
  - » 16 DAC phase adjusters (2 clocks for each DAC)

Next Challenges

- Improving PSR of all circuits in the path
- Integration of many links
  - » Low power, area, portable solutions
- Control of complex architectures
  - » Deal with loss, reflections and crosstalk
- Offset and mismatch
  - » Both voltage and time

- Lots of opportunity for design!

Measured S21's: FR4 no C-Bore

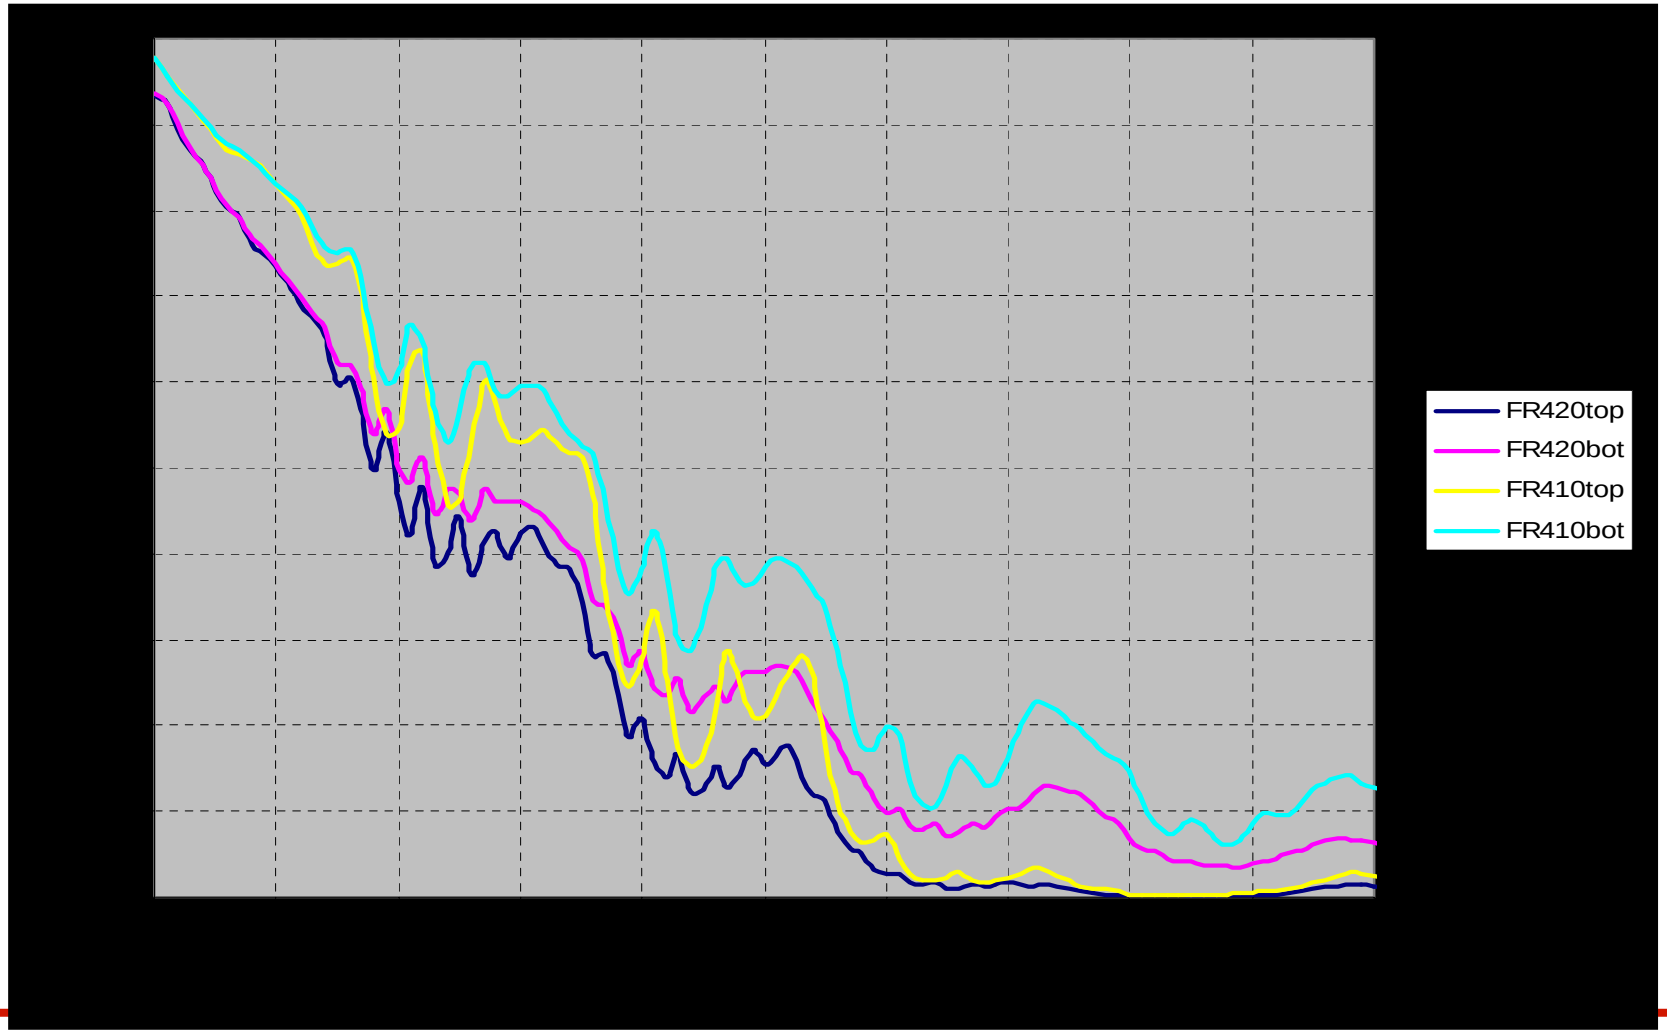

26” FR4 Bot 3.125Gbps, 2P w/EQ

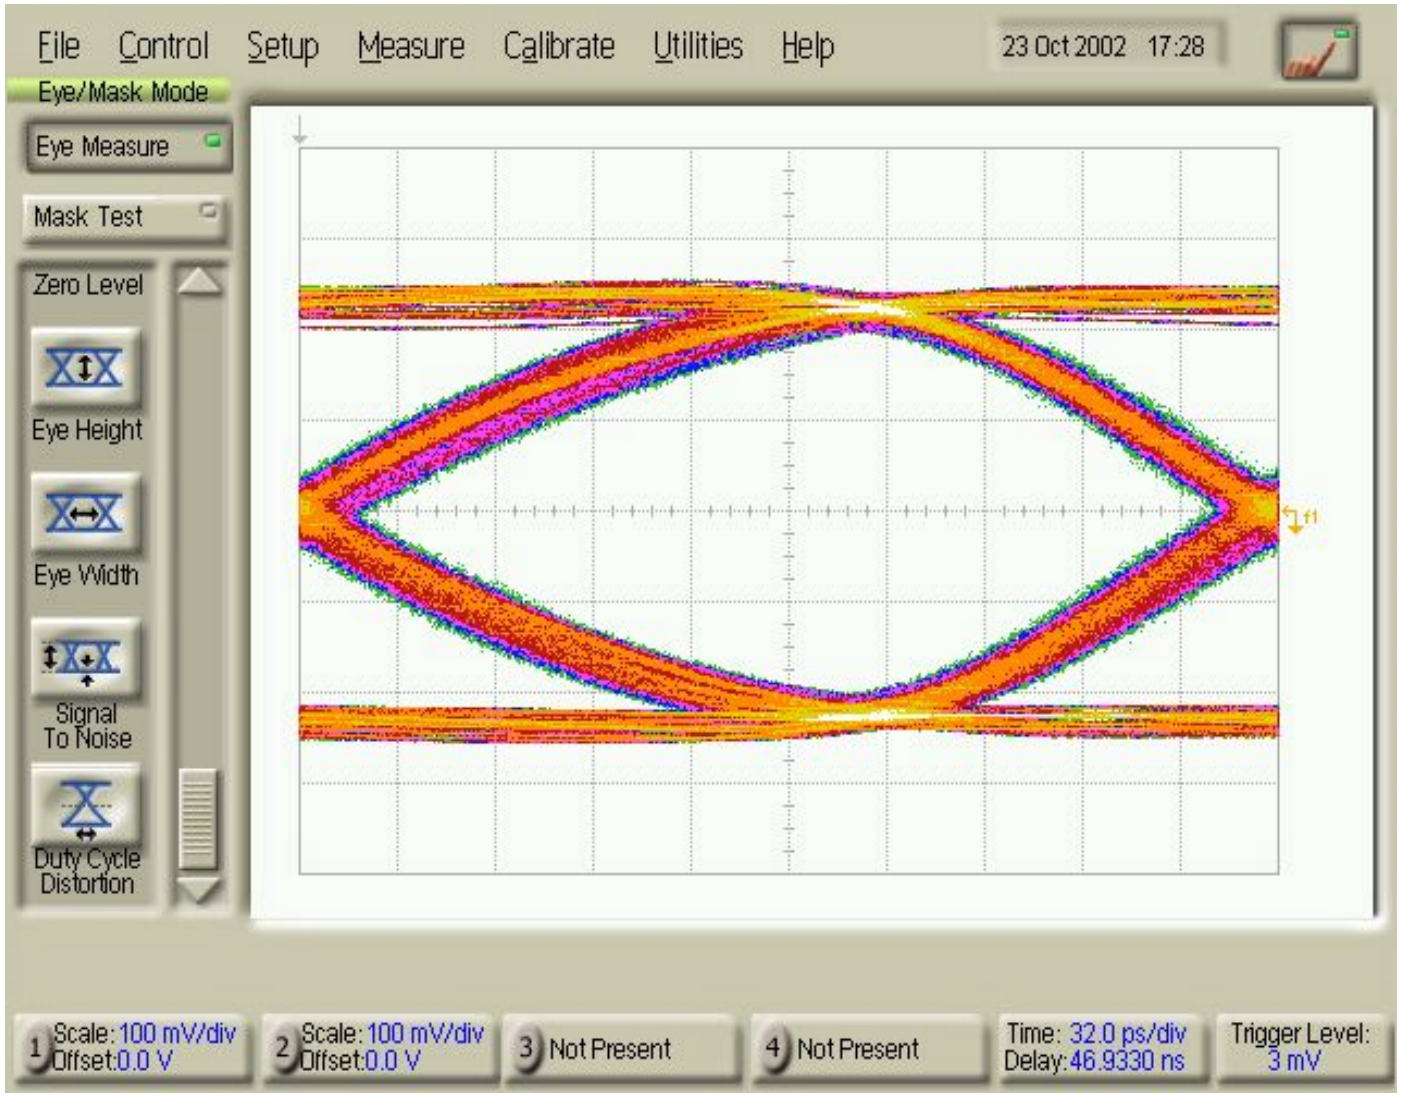

26” FR4 Bot 6.4Gbps, 2P w/EQ

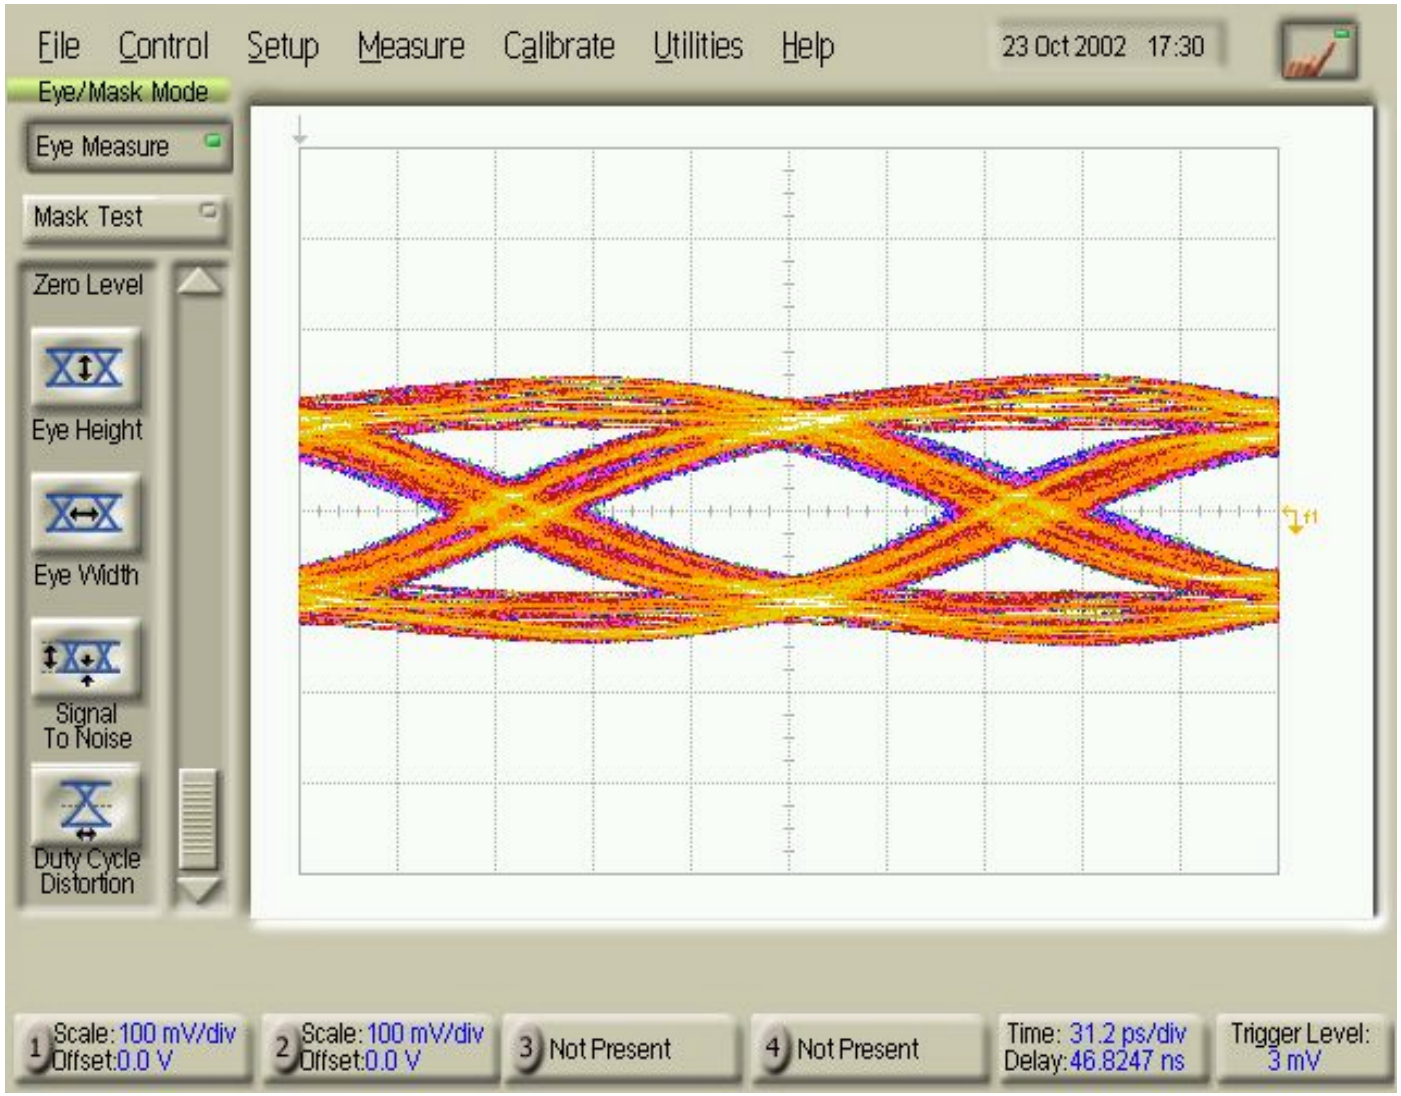

26” FR4 Top 6.4Gbps, 2P w/EQ

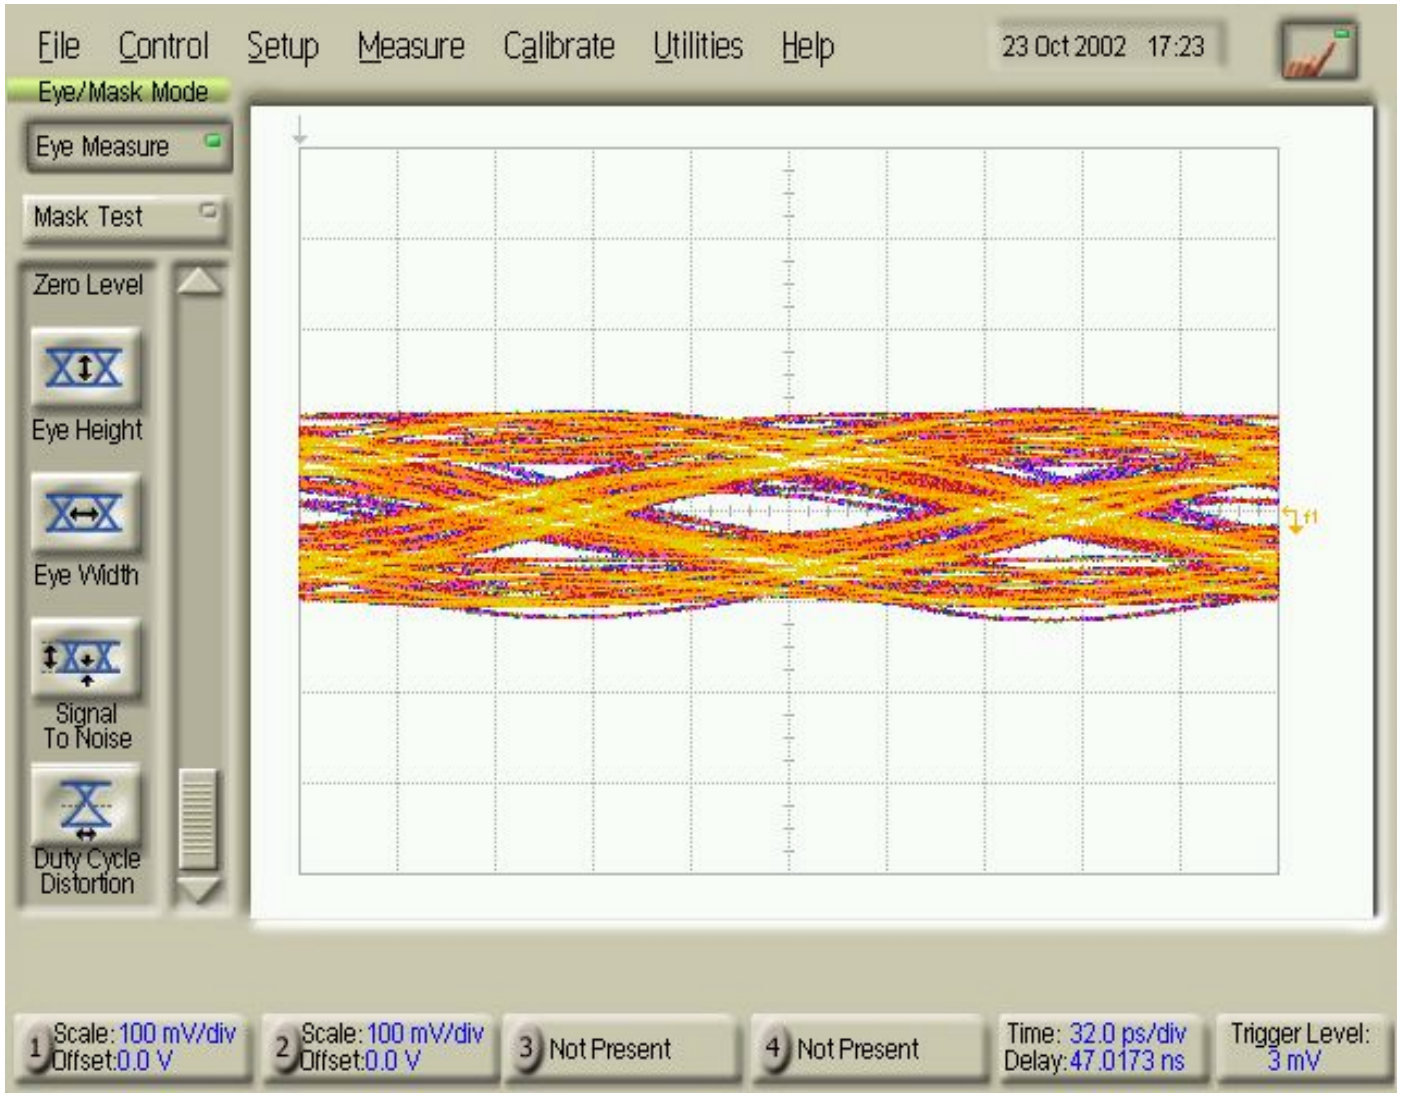

26” FR4 Top 6.4Gbps, 4P w/EQ

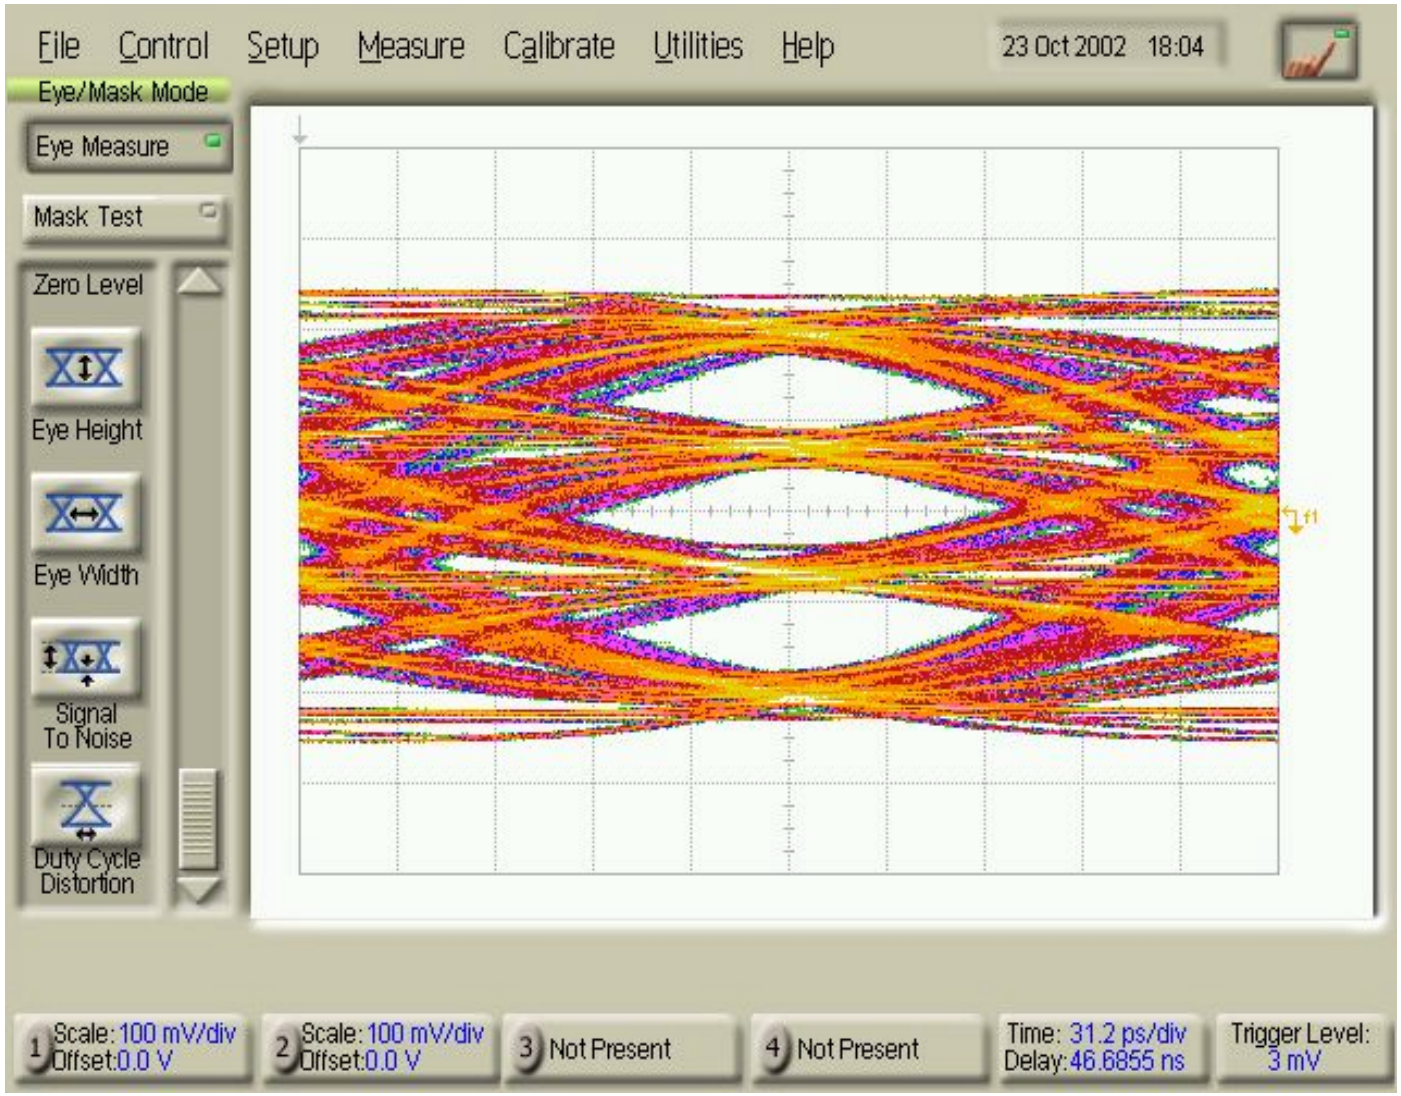

Measured S21's : N6k C-Bore

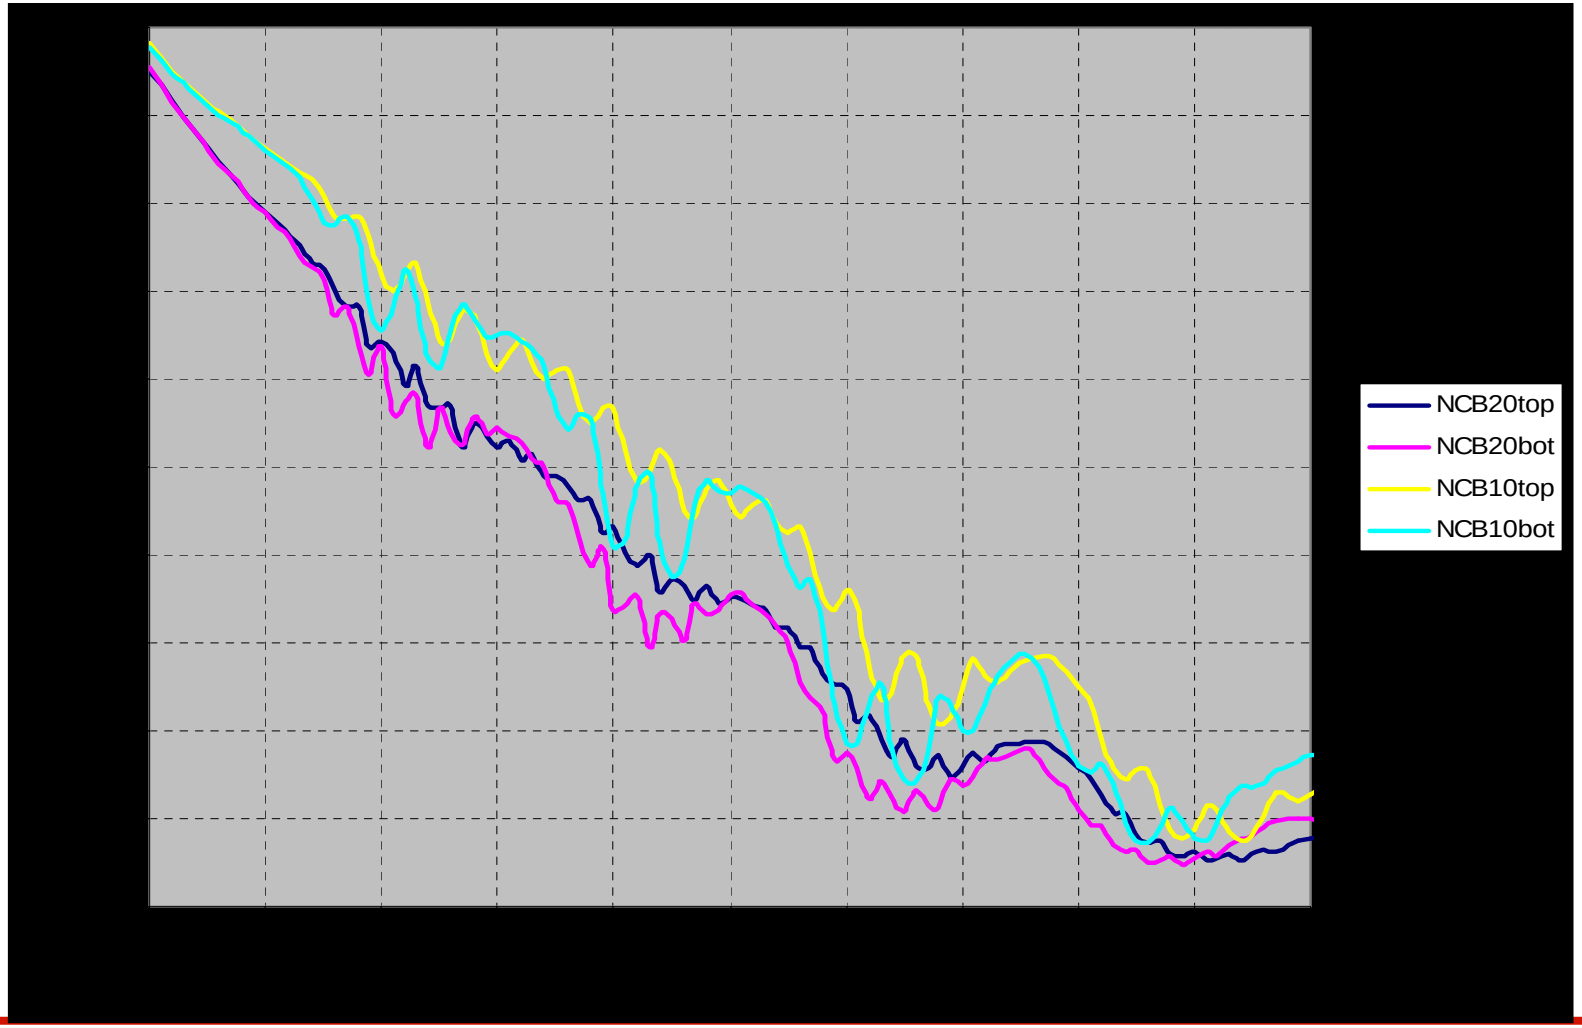

26” N6k-cb Top 6.4Gbps, 2P

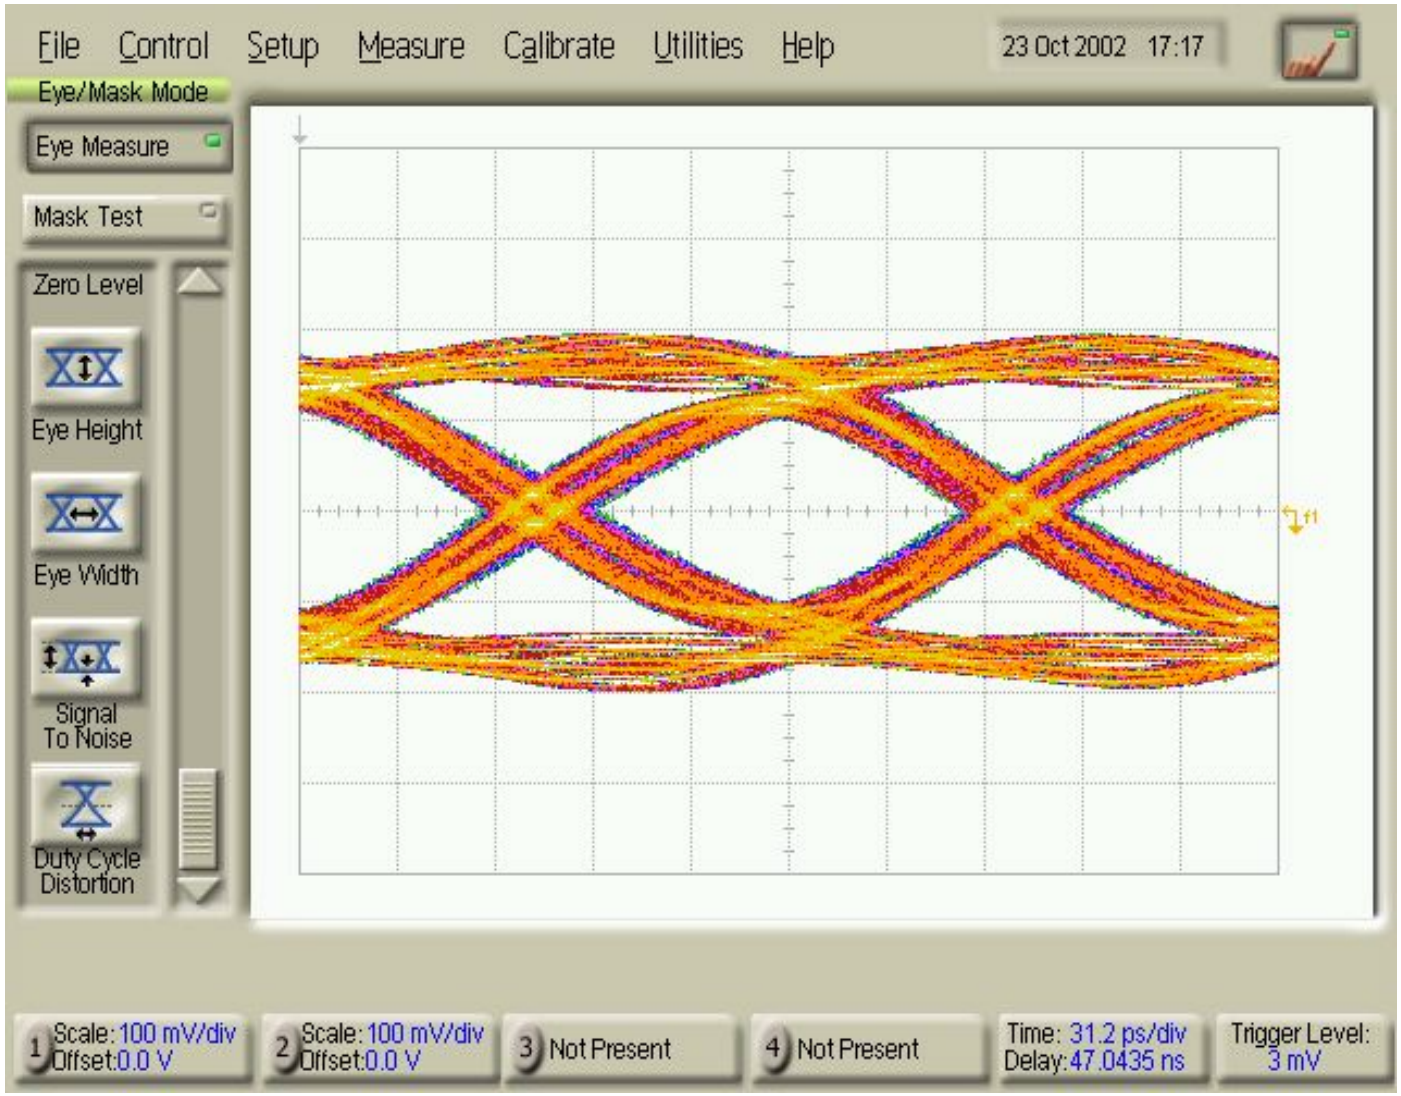

10G Eyes & System Margin Shmoos

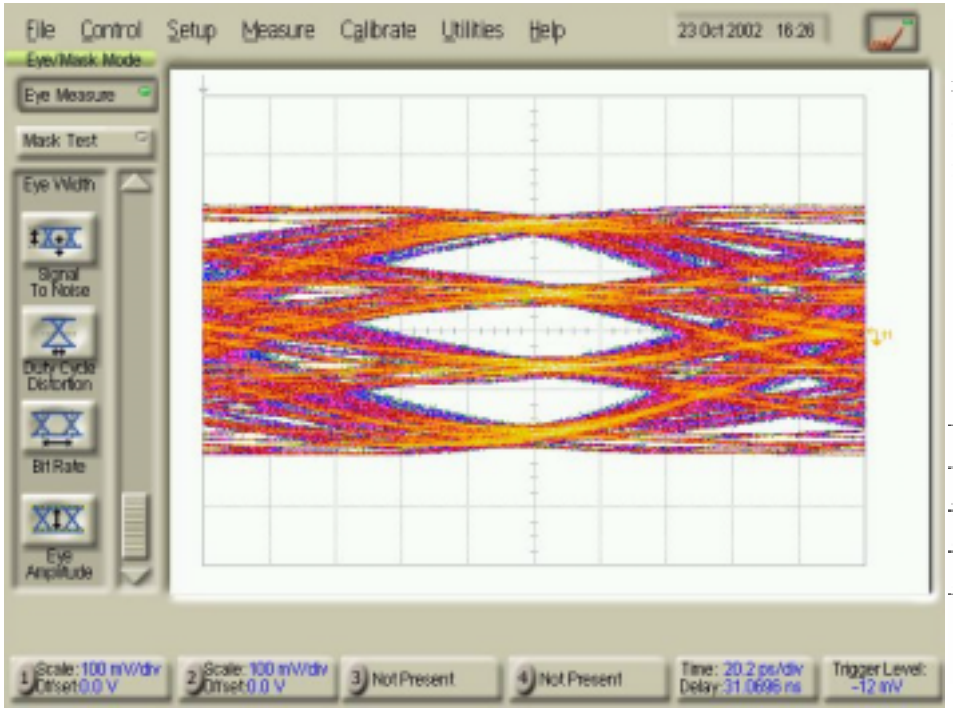

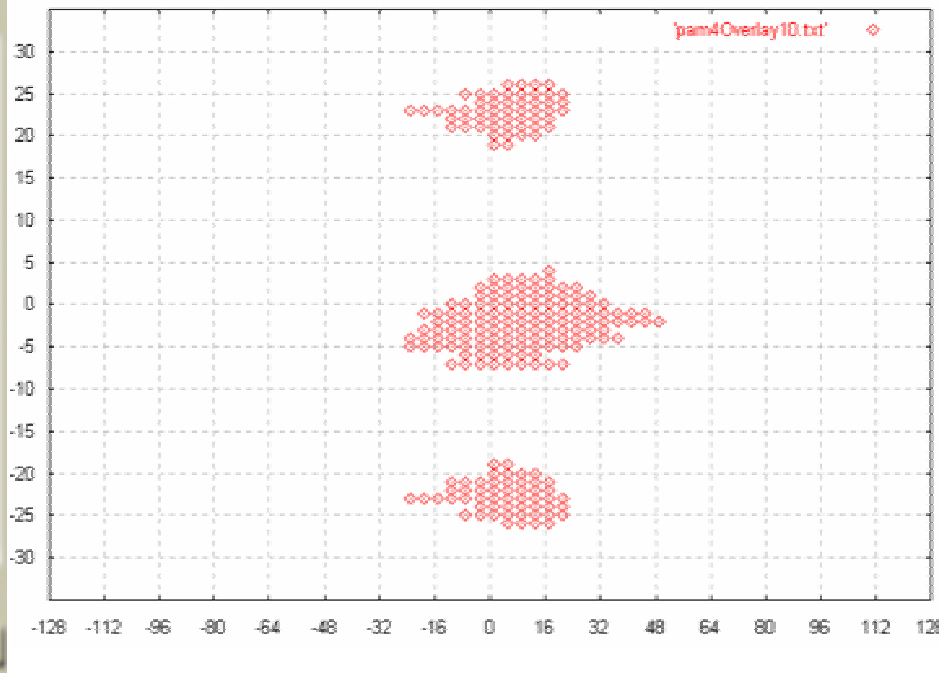

- 3”/20”/3” = 26” Trace + 2 Connectors
- Tested to BER < 10^{-15}

Link Performance vs. Time

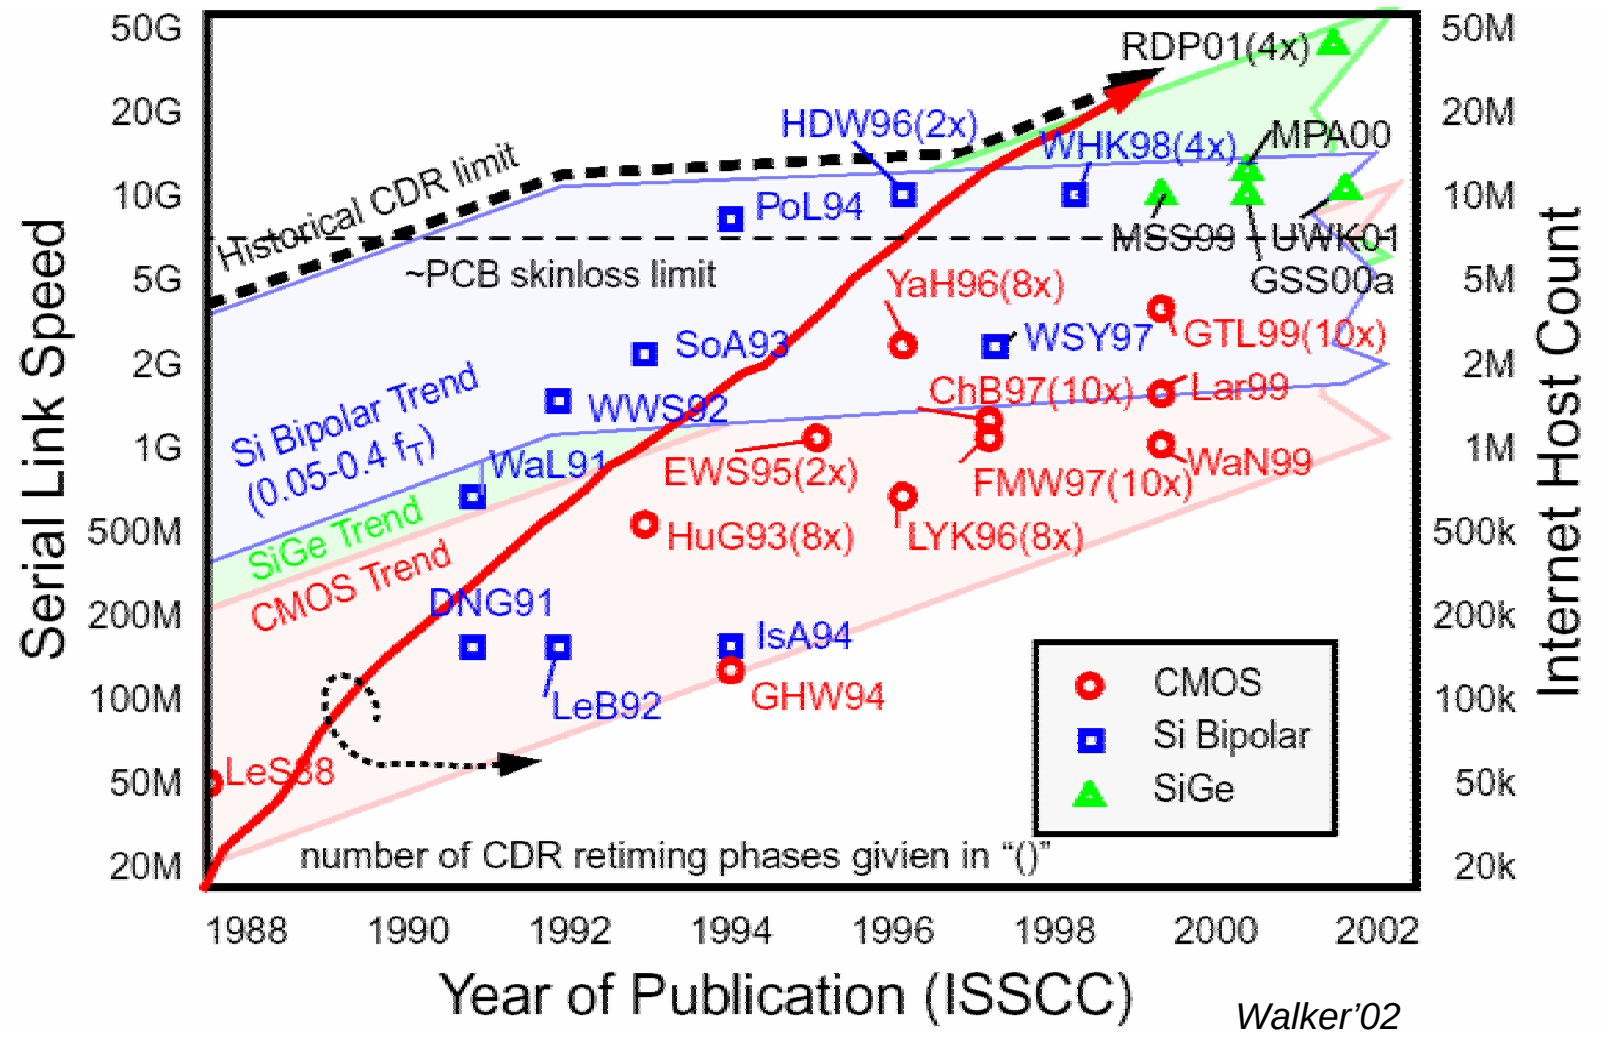

Walker'02

Link Efficiency: Gb/W, Gb/mm^2

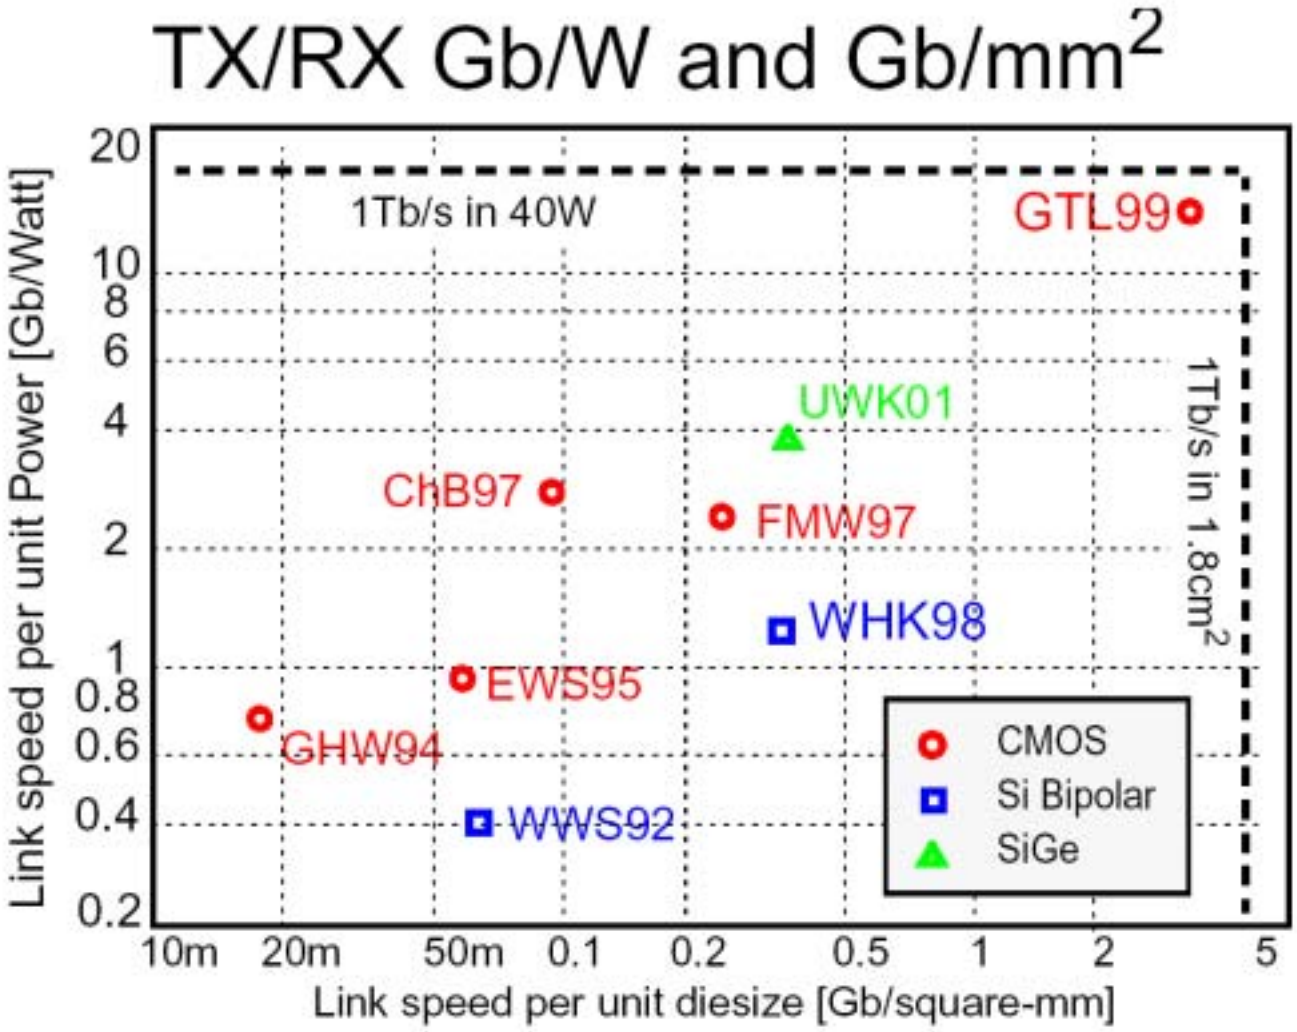

Walker'02

(ISSCC'92-2001)