Bits on speed!

. . . 0010010101 . . .

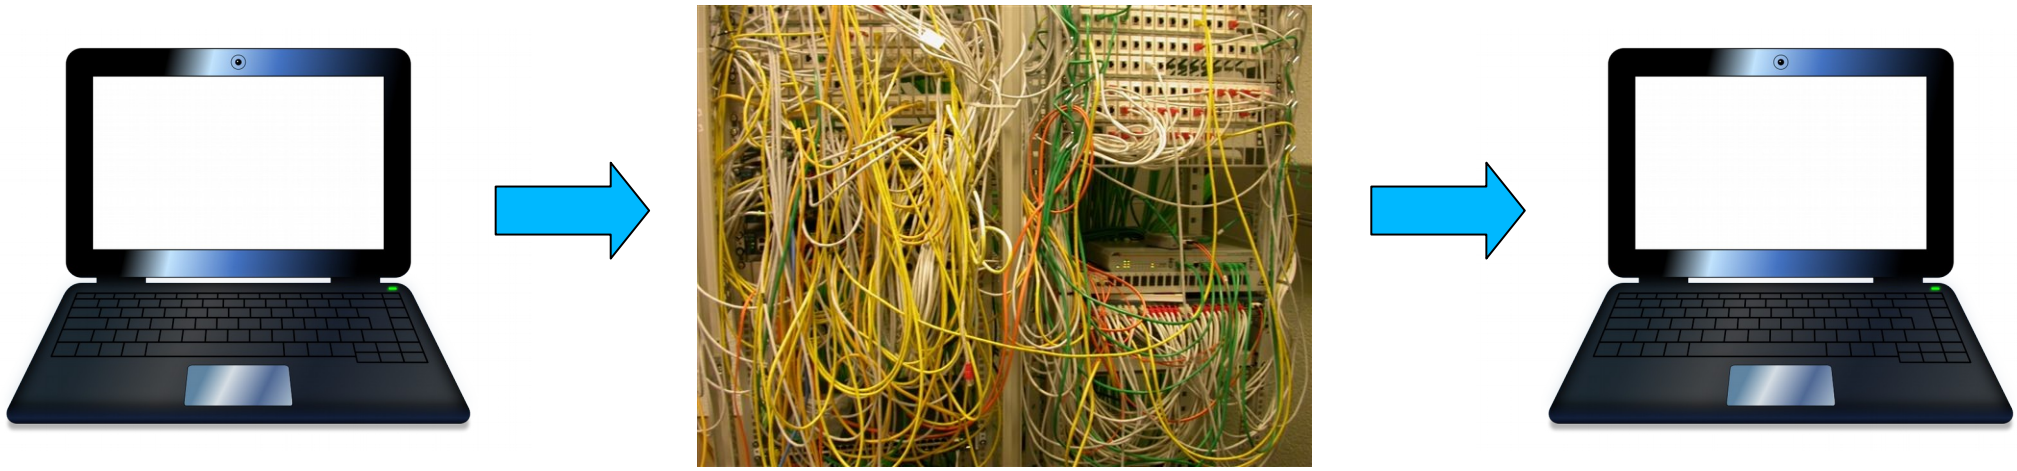

SERDES design in advanced CMOS technologies

miromico

SERDES design in advanced CMOS

- Motivation & Challenges
- Architectural solutions
- Implementation details
- Future trends & remarks

Internet traffic

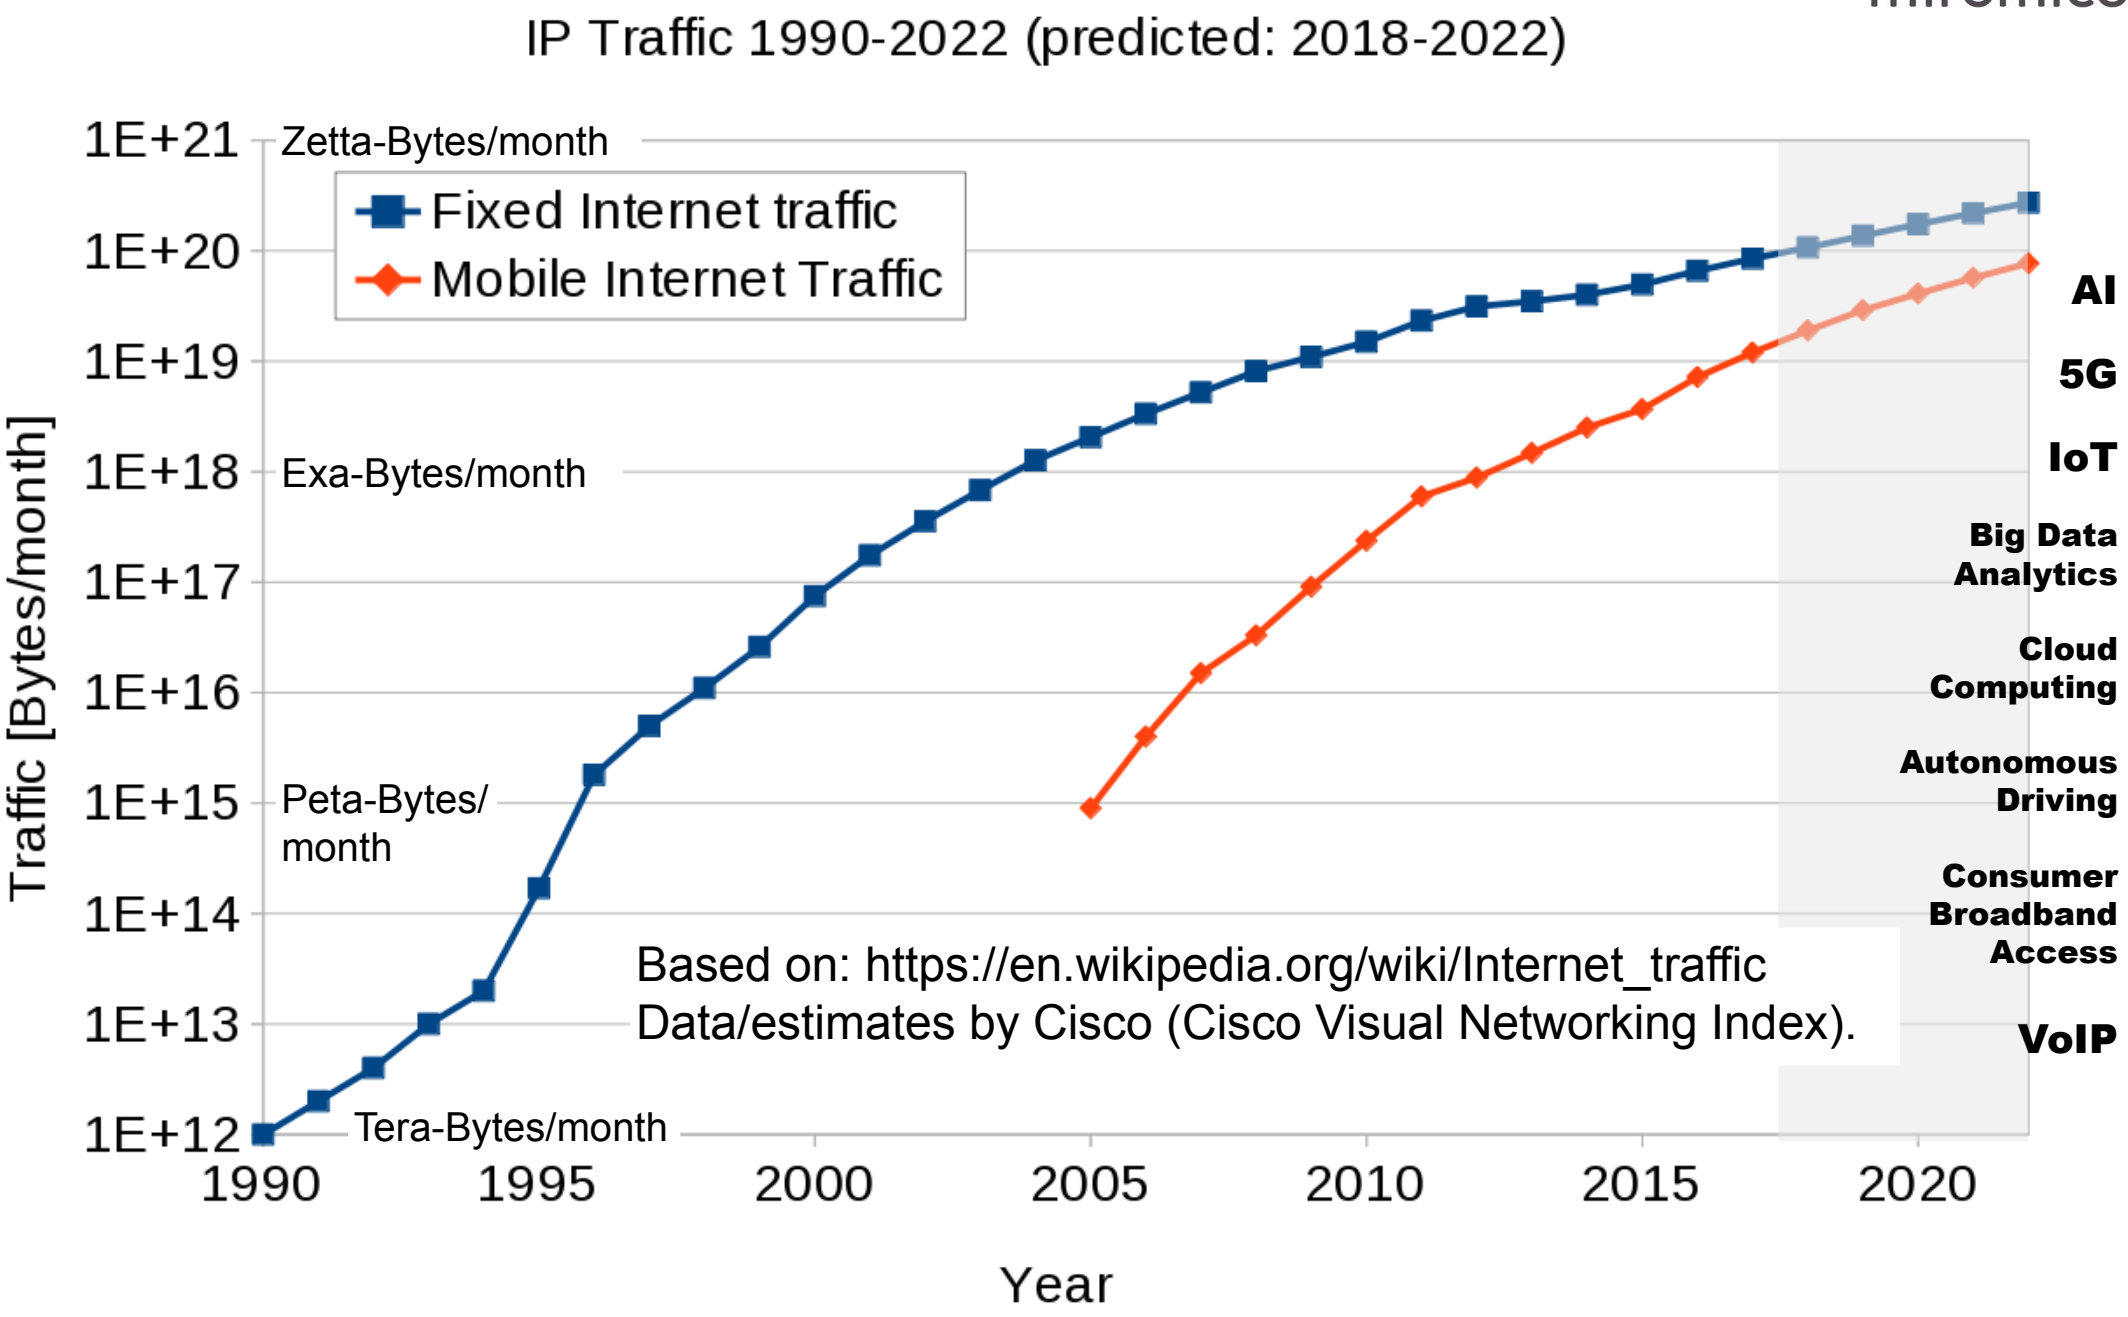

Interconnection speed – buses and networks

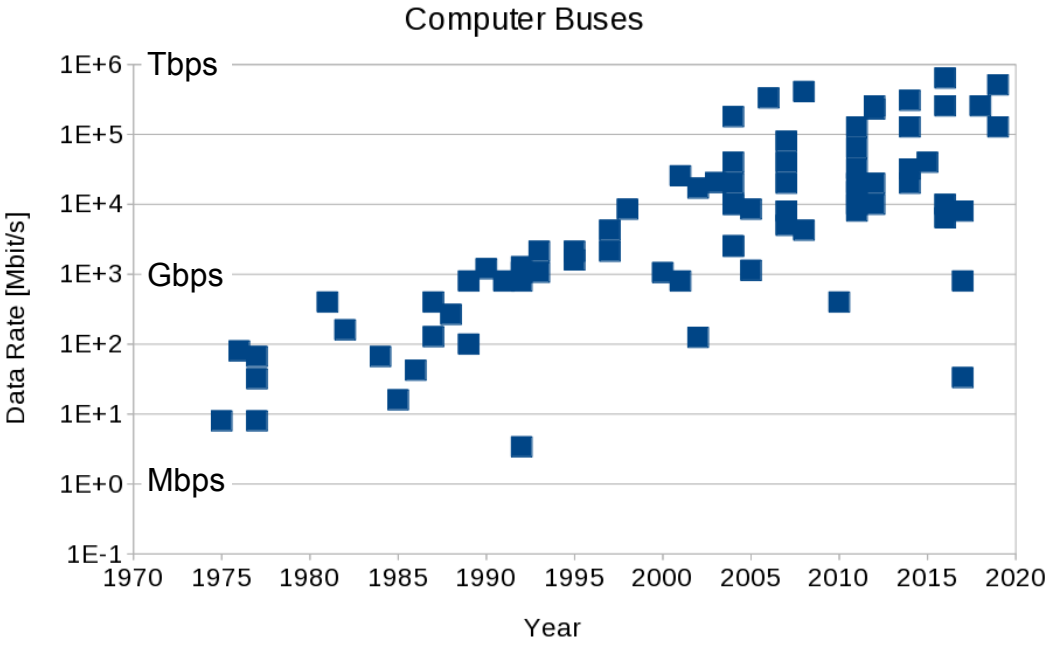

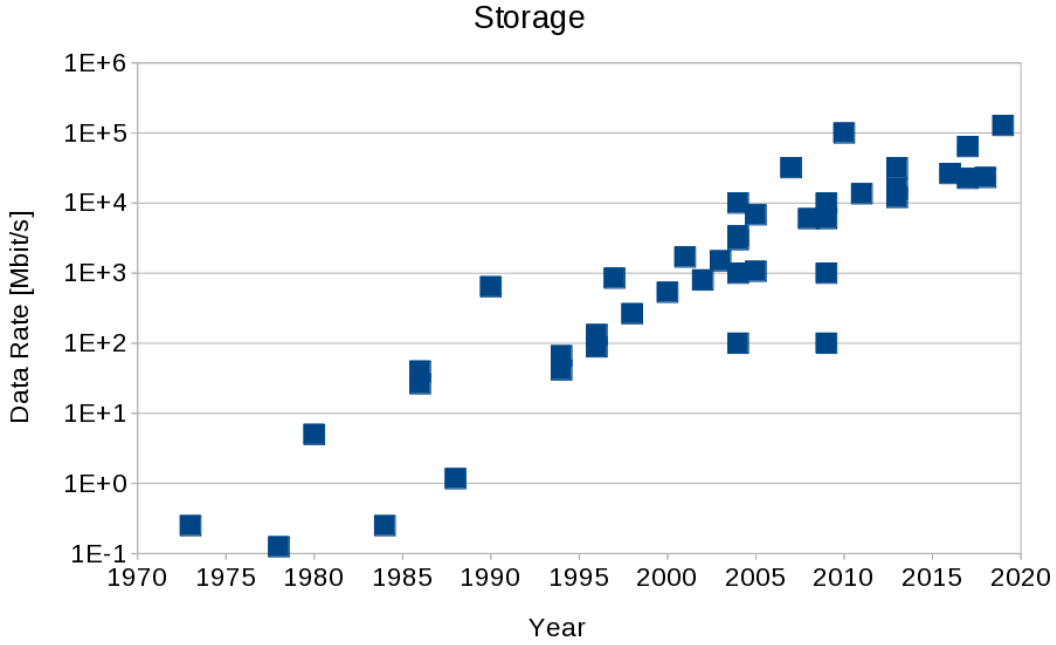

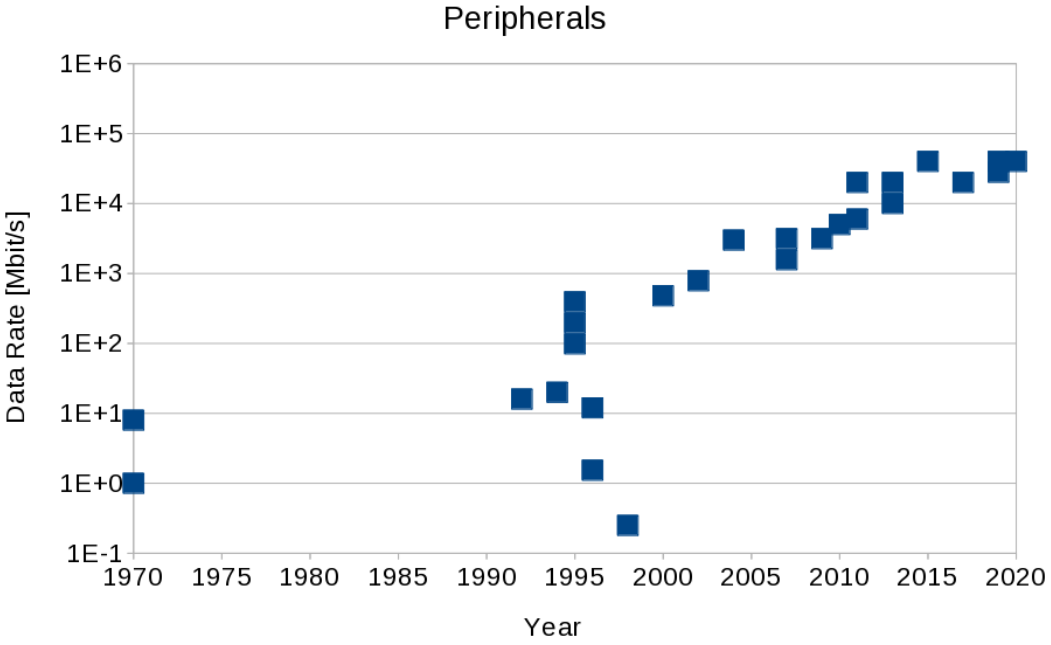

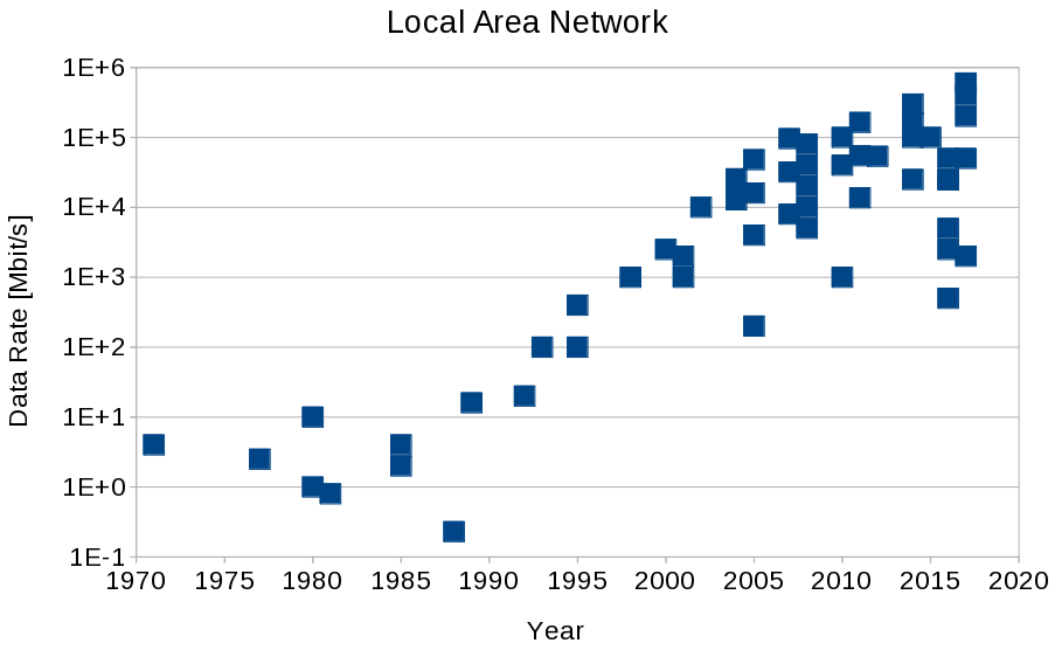

Source: https://en.wikipedia.org/wiki/List_of_device_bit_rates

High Speed Interconnects today

- 99.99% of data processing to-date is silicon transistors
- 90% of long distance signaling is optical
- Networks use different media to overcome their limitations
  - Chip-to-chip, intra-board, backplane, “short” cable interconnects still electrical
  - Move to optical for “longer” distances
- “Longer” becomes shorter at every generation
  - Tighter integration within the package
  - Integrated optics (silicon photonics)

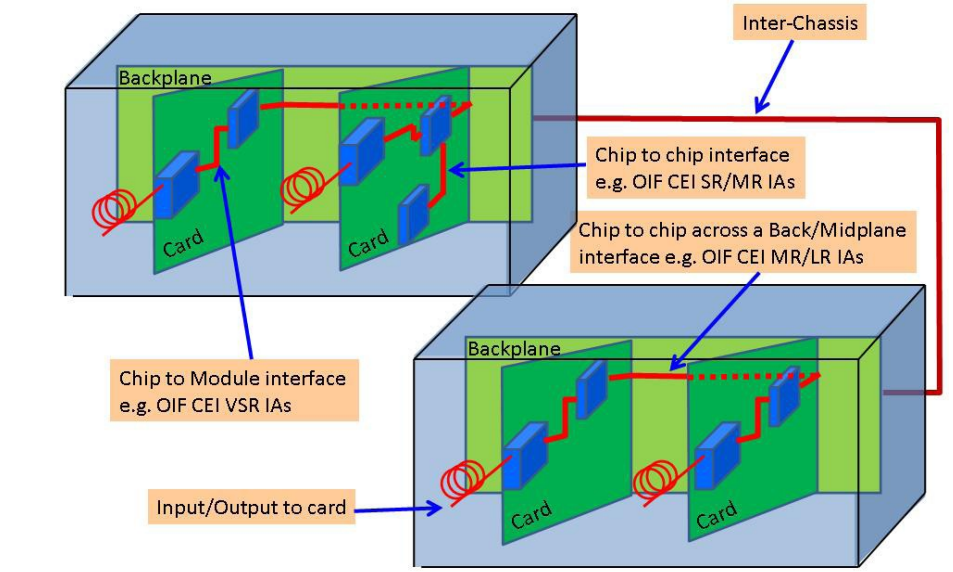

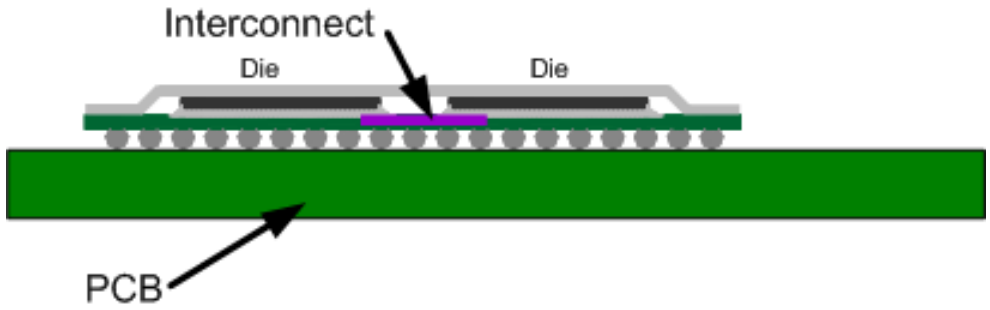

Picture: OIF CEI-56G Application Note

CHIP-TO-CHIP COMMUNICATION SPEED HAS BECOME THE MAJOR BOTTLENECK IN IMPROVING OVERALL PERFORMANCE !

Hierarchical network organization

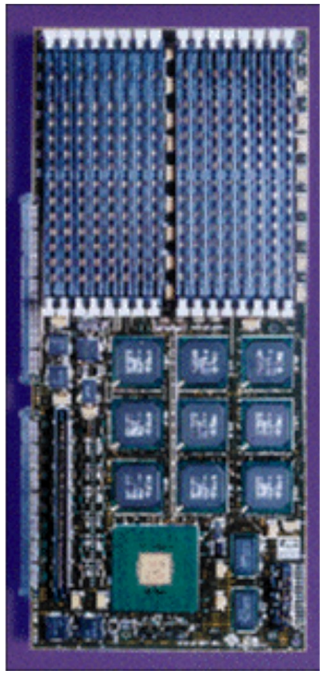

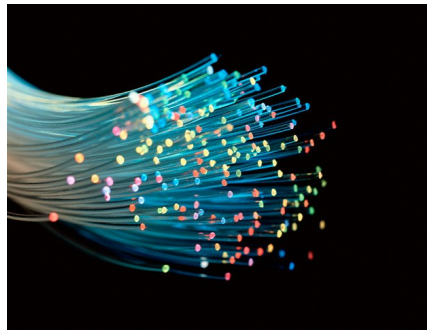

WAN

100m+

LAN
(LR)

<50dB

<100m

Backplane
Copper cable
(MR, LR)

20-35dB

~1m

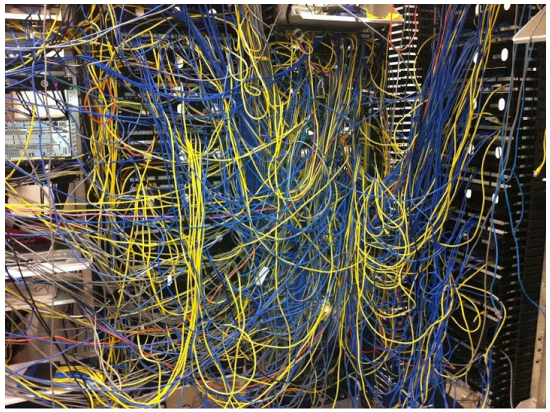

Chip-to-chip
Intra-board
(VSR, SR)

5-20dB

~10cm

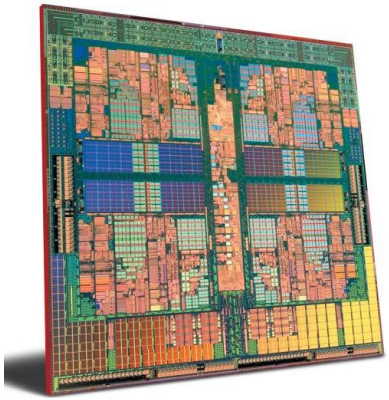

Core-to-Core
Intra-Package
(USR, XSR)

2dB

1-2cm

Max bit density

Max speed
Min power

Distance/Loss

Challenge for the I/O Designers

- *Need to double data-rate every 2-3 years, while satisfying other constraints*
- Same total system cost:
  - Increase complexity in silicon before going to fancy new materials!
- Same total power:
  - Limited by package type, cooling system, etc.
  - …but also by norms & regulations (cannot be too loud!)
  - Power limit capped at ~125W / chip (which is easily reached by a 150-way network-switch in 14nm CMOS, 50% from the I/Os!)

- Same chip size:
  - Doubling data-rate roughly doubles the number of associated gates, but also greatly increases the complexity of required equalization, FEC schemes, etc.
  - A 150-way network-switch in 14nm CMOS is roughly 25mm X 25mm bare die, 30% from the I/Os (cannot be much bigger than that!)
- Same package size:
  - Higher frequency typically means higher problems related to interference, cross-talk, signal integrity and calls for bigger packages
  - A 150-way network-switch in 14nm CMOS has a 62.5mm body-size package (cannot be much bigger than that!)
- Lower Signal-to-Noise Ratio:
  - Higher bandwidth, higher device noise
  - Higher loss in the channel

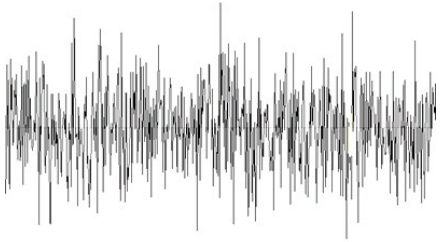

SERDES design in advanced CMOS

- Motivation & Challenges
- **Architectural solutions**
- Implementation details
- Future trends & remarks

The generic Serial-Link

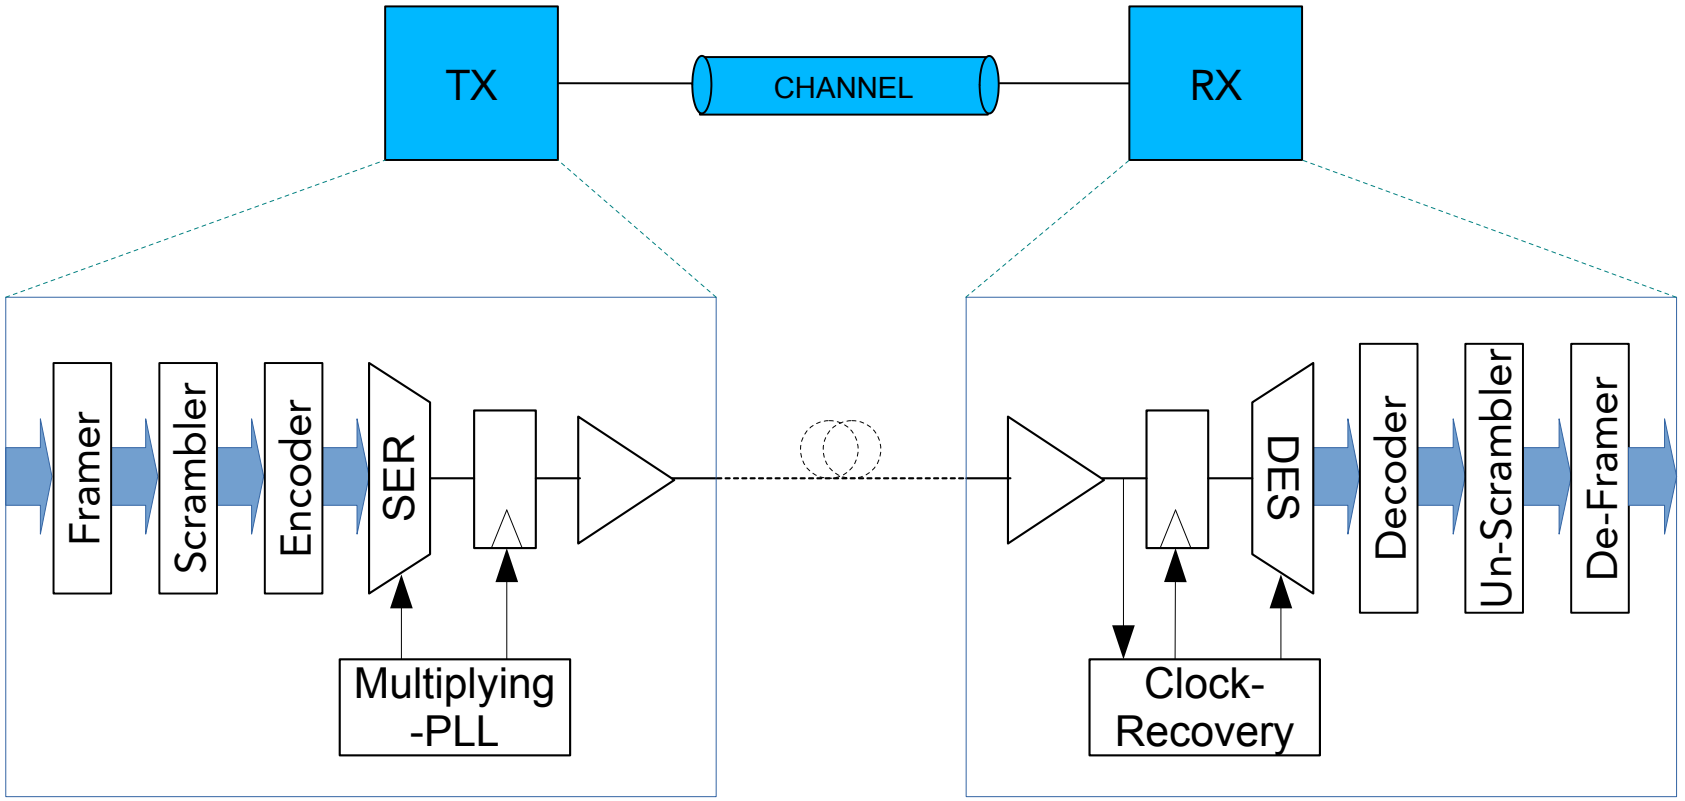

- Why “serial” links? Embed the clock into the data to avoid time skew limit!
  - In “parallel” links the clock is sent separately (eg. source sync., system sync., etc.)
  - However timing-skew limits the speed (<Gbps range) and/or distance (USR, XSR)
  - No time-skew in serial links! BW >> Gbps (but: need encoding and clock-recovery)
  - Serial links have been systematically replacing parallel ones in the past few decades
  - What counts is the datarate-per-pin (...the communication is massively parallel anyway!)

Typical Electrical Interconnect characteristics

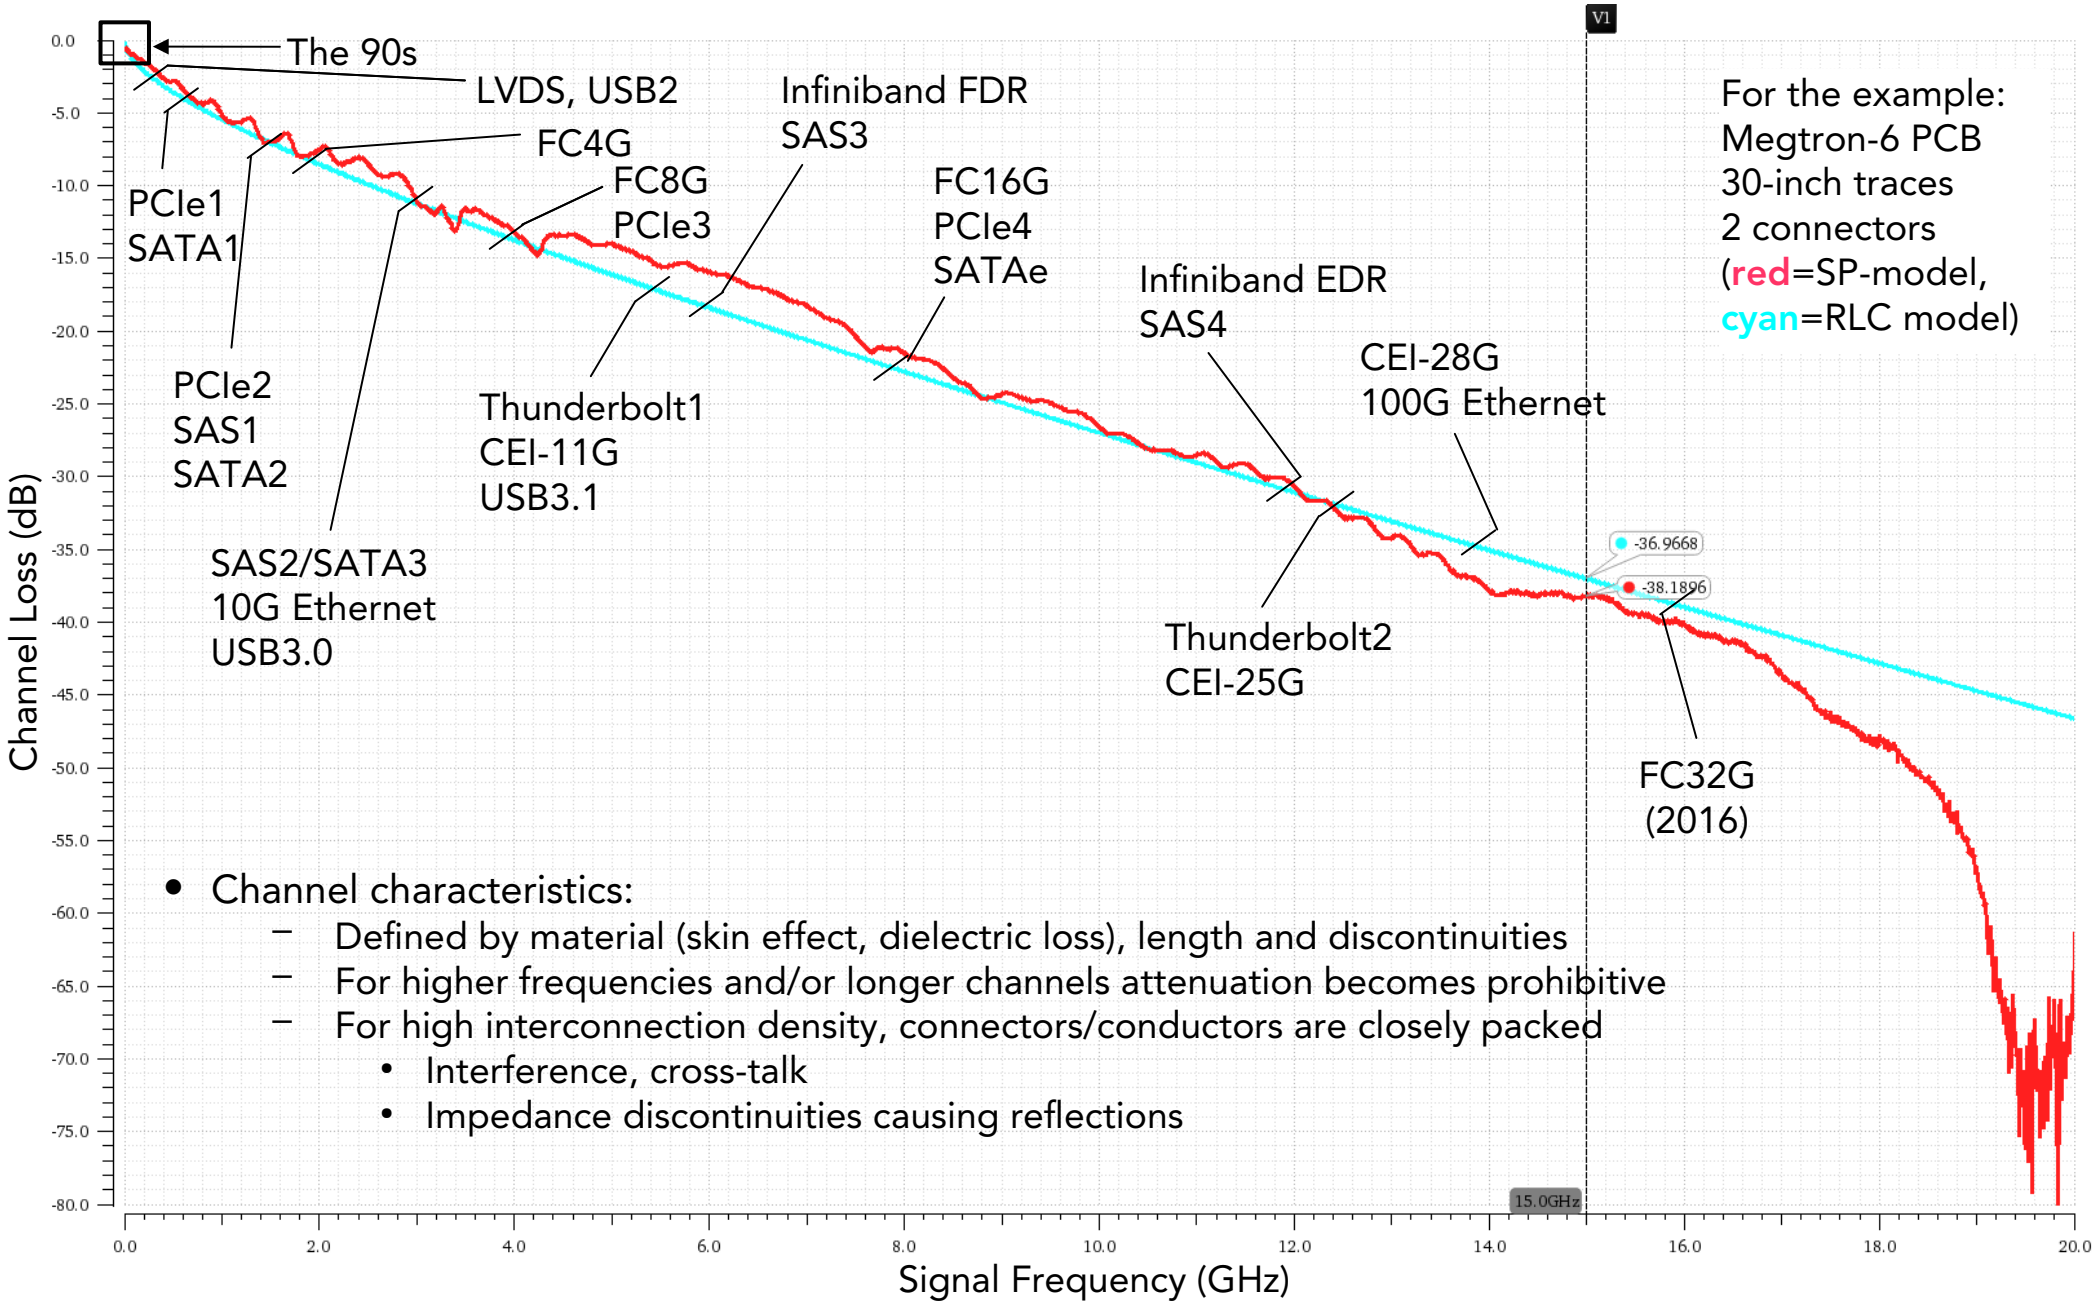

- Channel characteristics:
  - Defined by material (skin effect, dielectric loss), length and discontinuities
  - For higher frequencies and/or longer channels attenuation becomes prohibitive
  - For high interconnection density, connectors/conductors are closely packed
    - Interference, cross-talk
    - Impedance discontinuities causing reflections

Effect of channel loss on the signal?

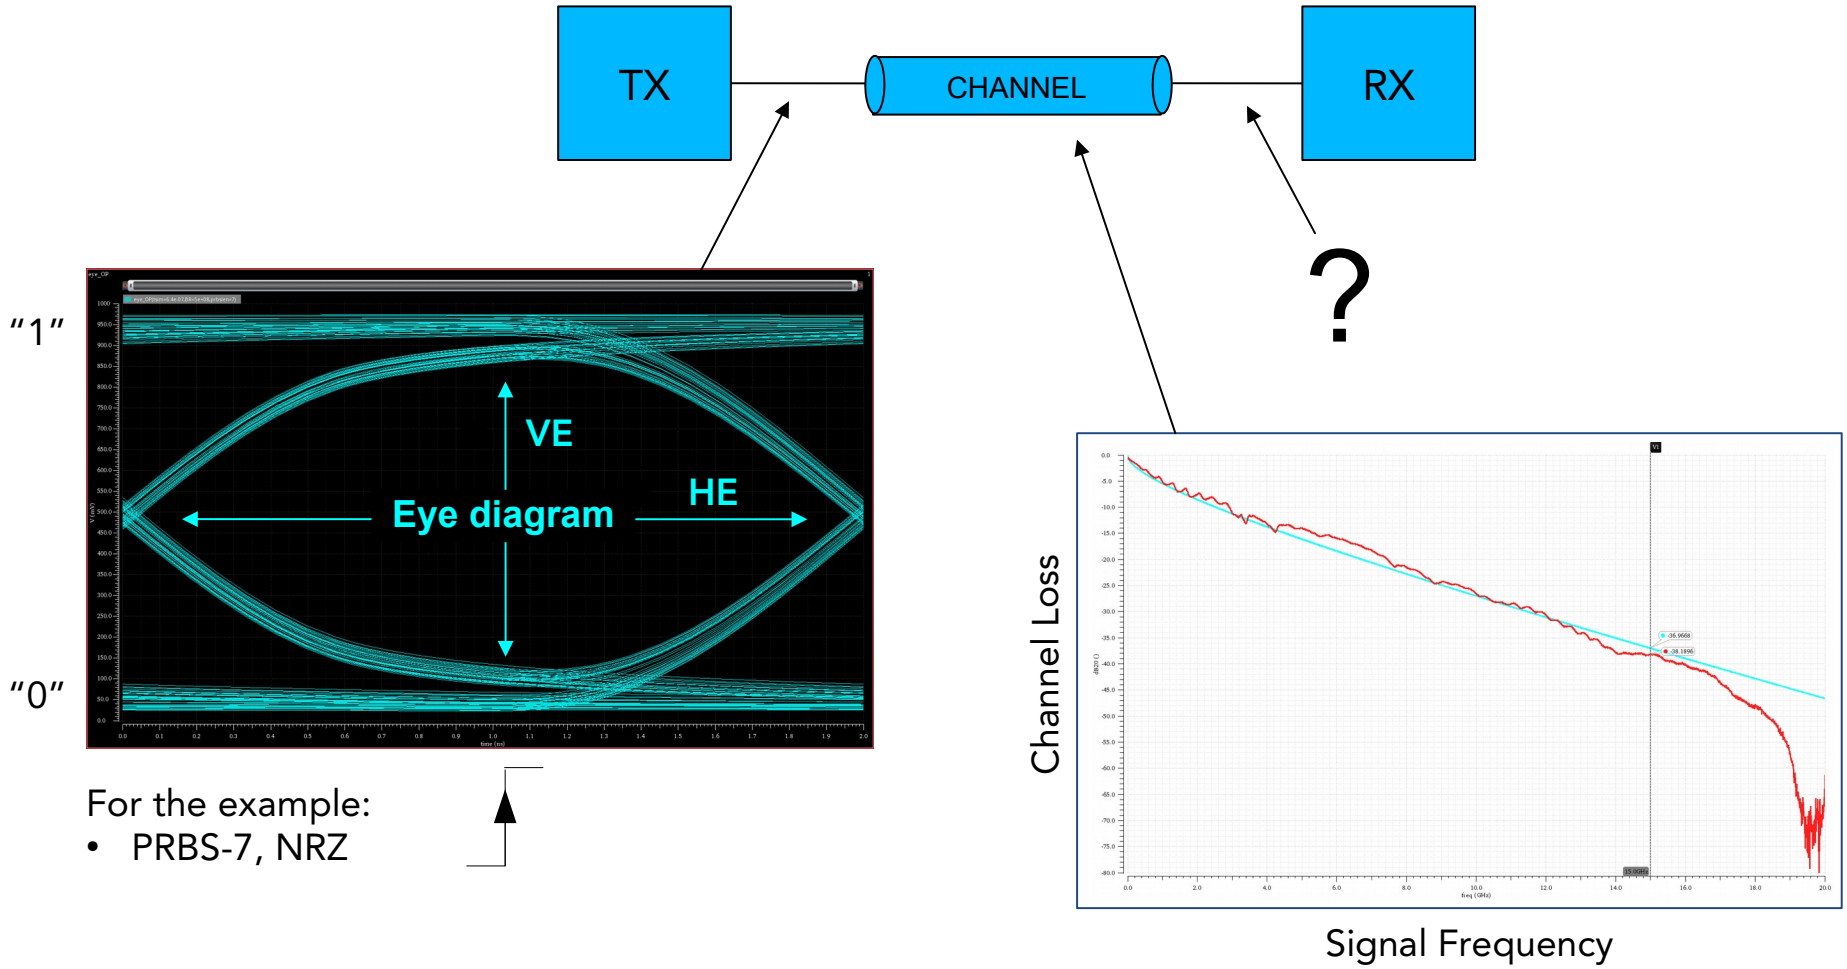

- RX input signal must have sufficient Vertical- and Horizontal-Eye openings:
  - VE must accommodate for slicer sensitivity, offset and noise
  - HE must accommodate for clocking accuracy and signal jitter

Effect of channel loss on the signal?

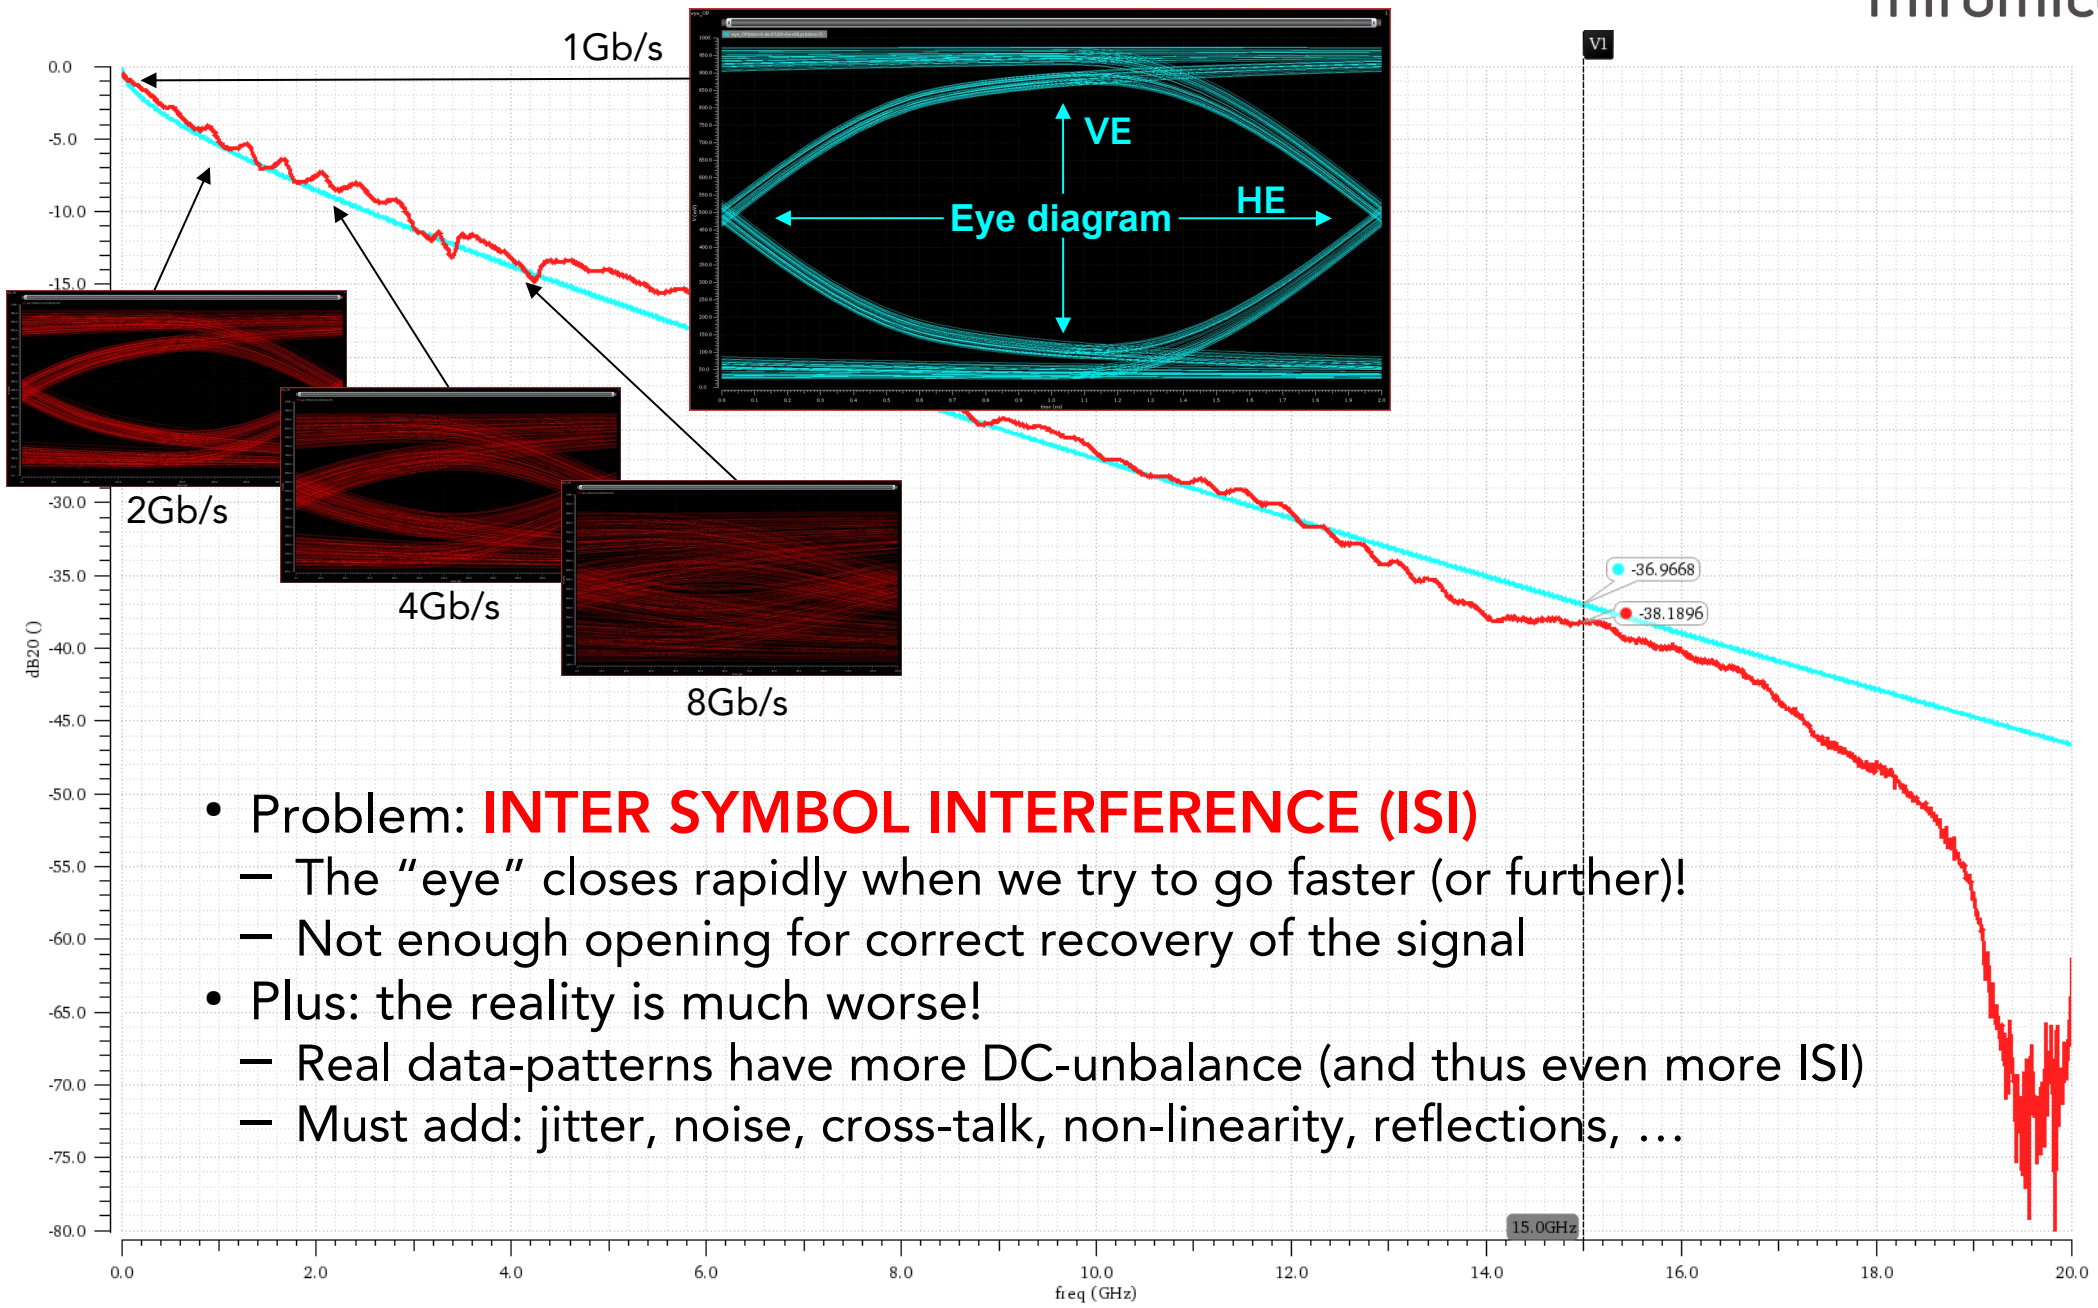

- Problem: **INTER SYMBOL INTERFERENCE (ISI)**
  - The "eye" closes rapidly when we try to go faster (or further)!
  - Not enough opening for correct recovery of the signal
- Plus: the reality is much worse!
  - Real data-patterns have more DC-unbalance (and thus even more ISI)
  - Must add: jitter, noise, cross-talk, non-linearity, reflections, …

Inter Symbol Interference (ISI)

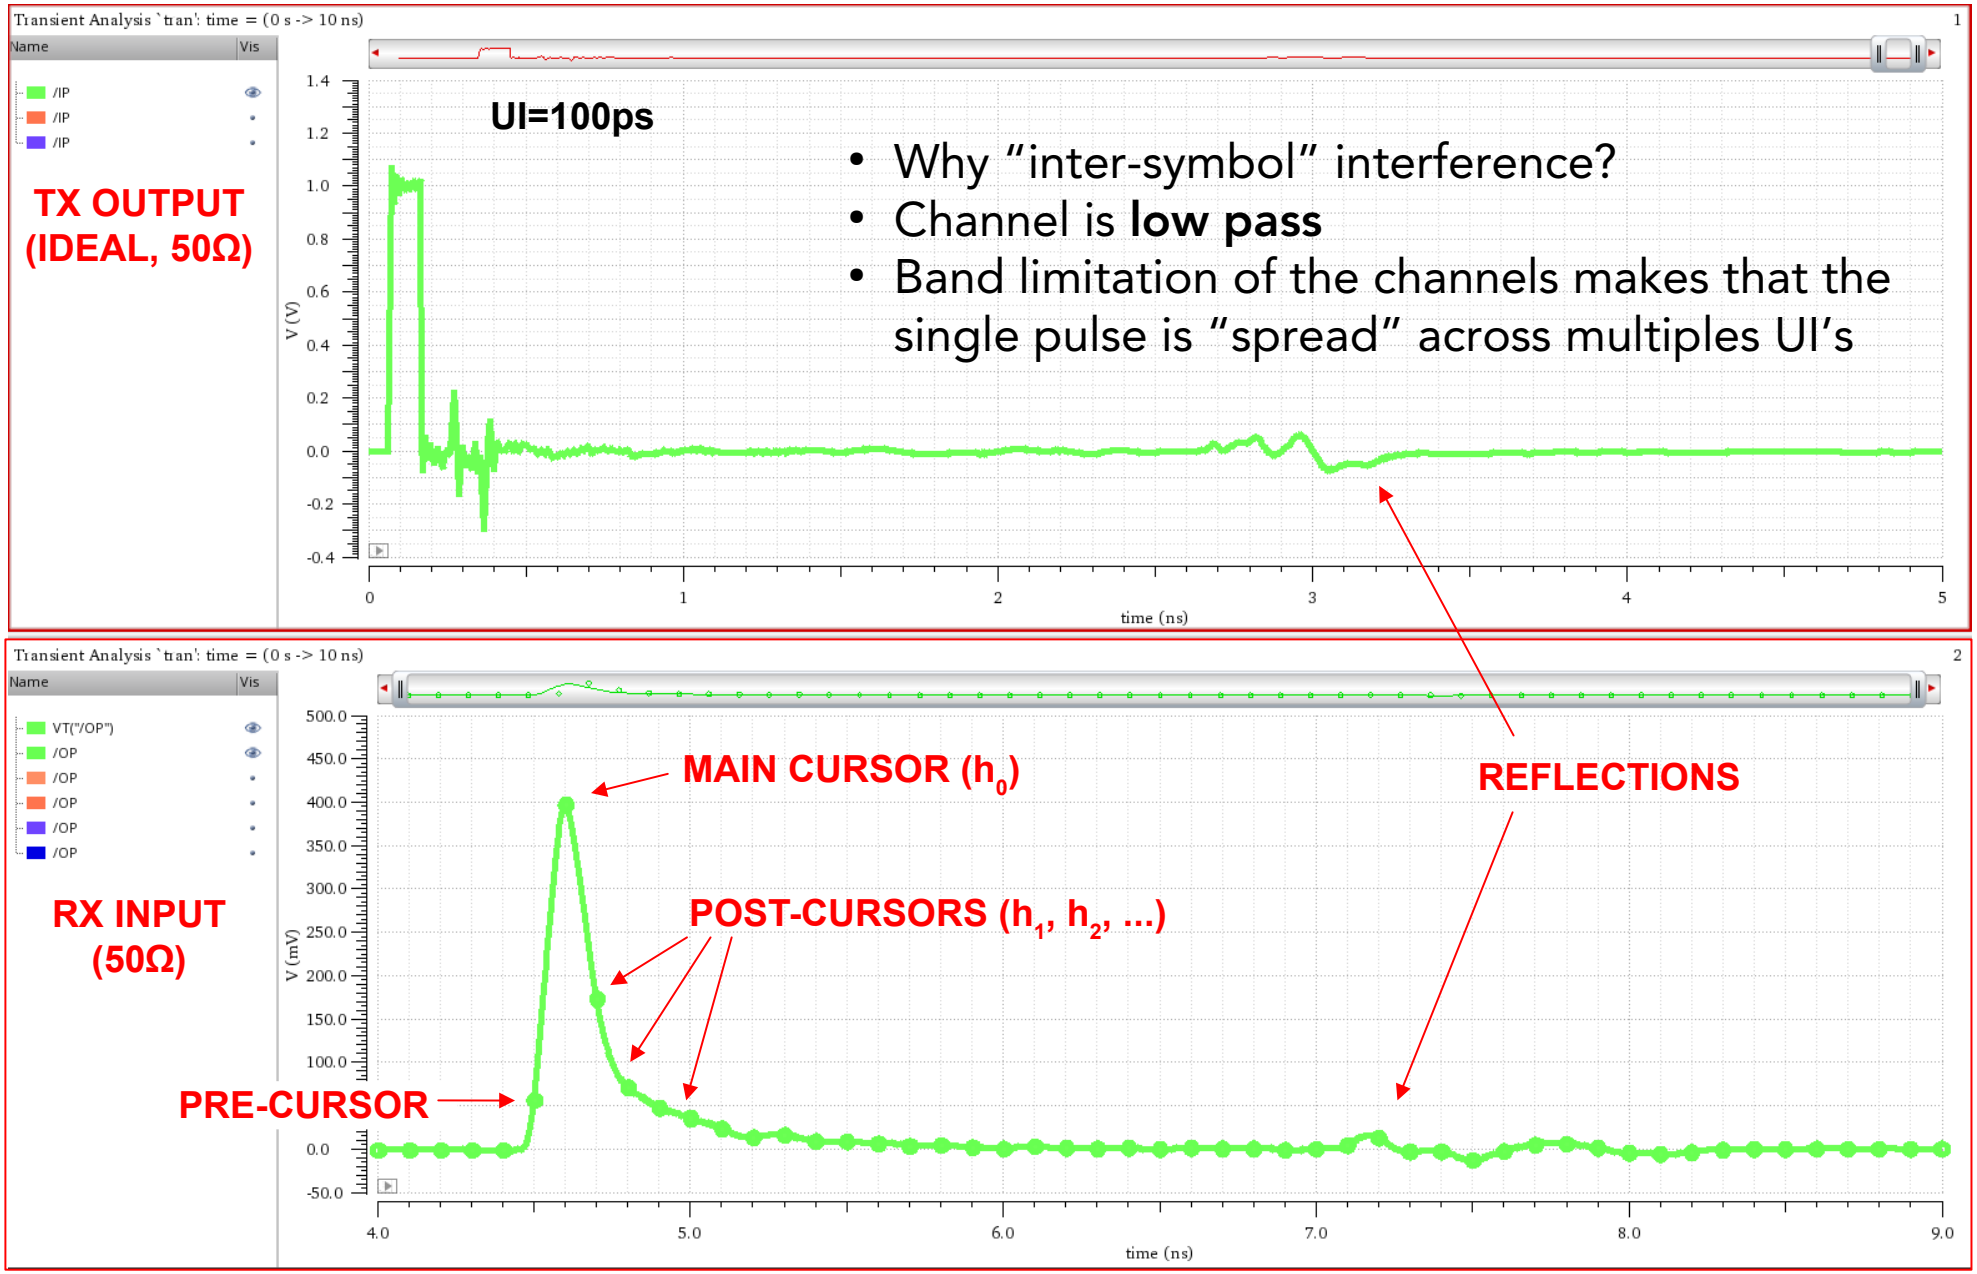

Inter Symbol Interference (ISI)

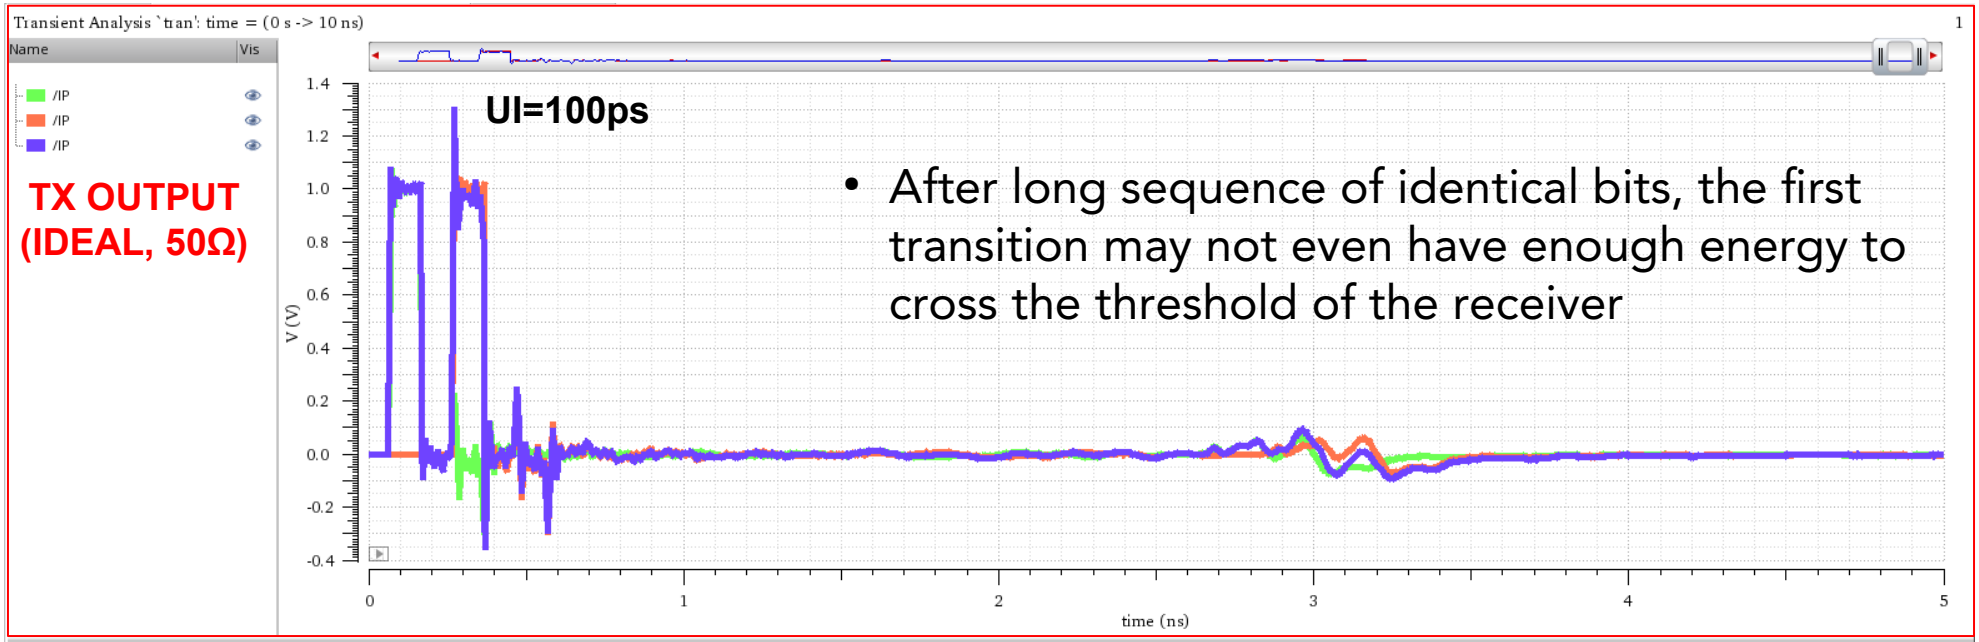

- After long sequence of identical bits, the first transition may not even have enough energy to cross the threshold of the receiver

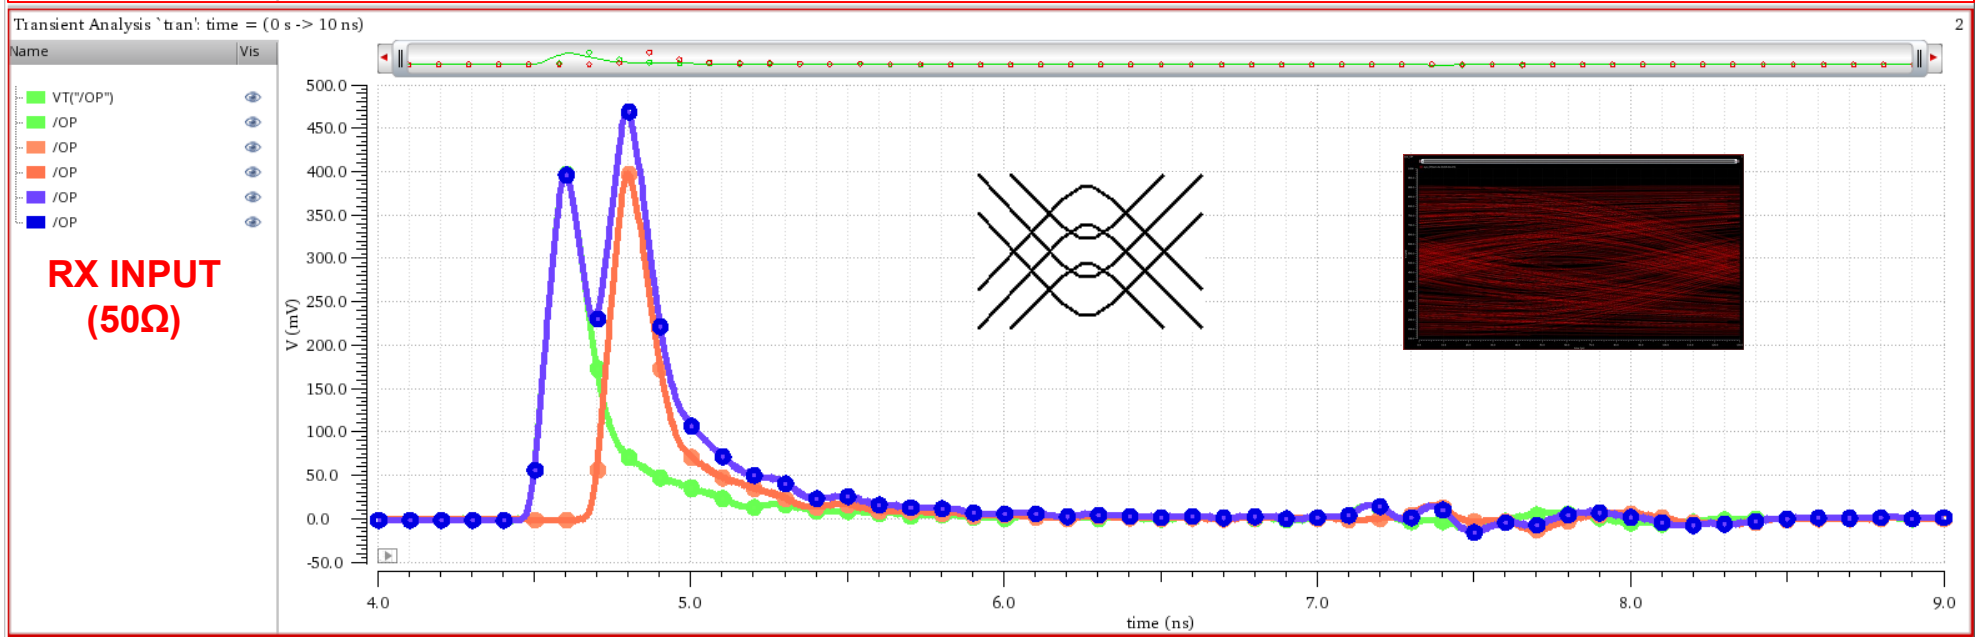

Equalization of the channel (example)

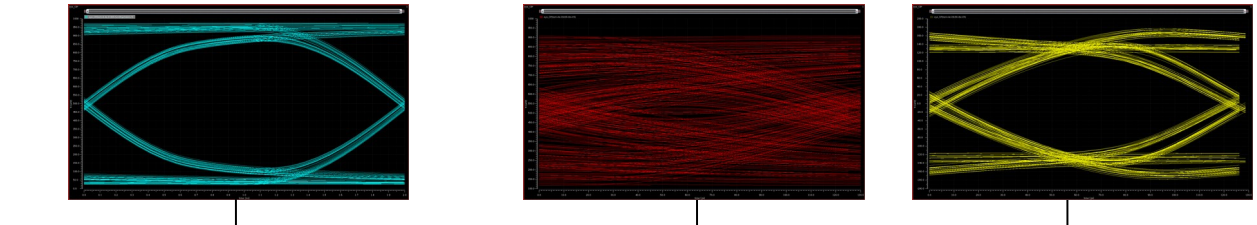

"EQ" TRIES TO REMOVE THE EFFECT OF THE CHANNEL

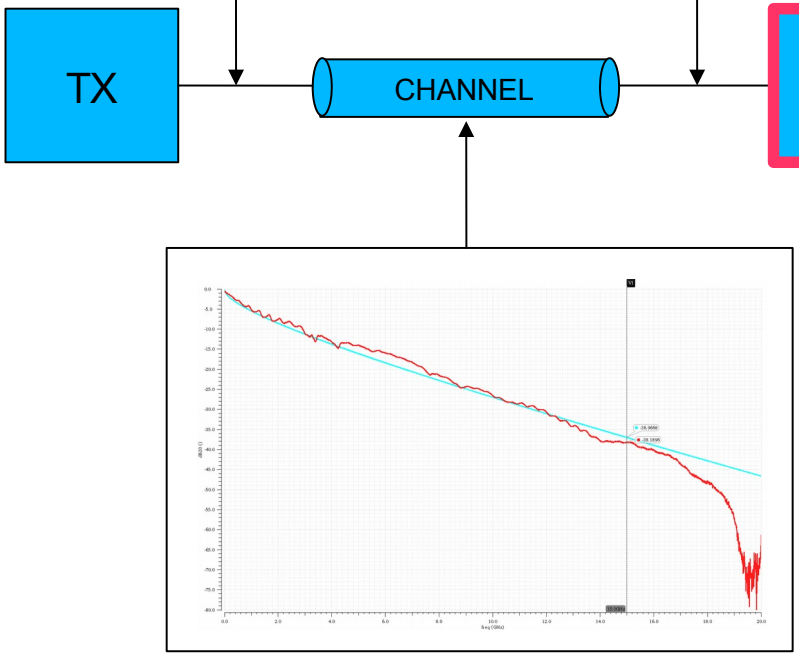

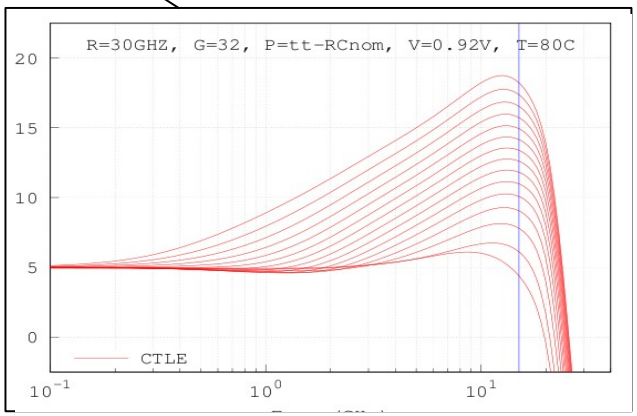

EX.: BY APPROXIMATING THE INVERSE CHANNEL RESPONSE (Linear Equalizer)

Equalizer Classification

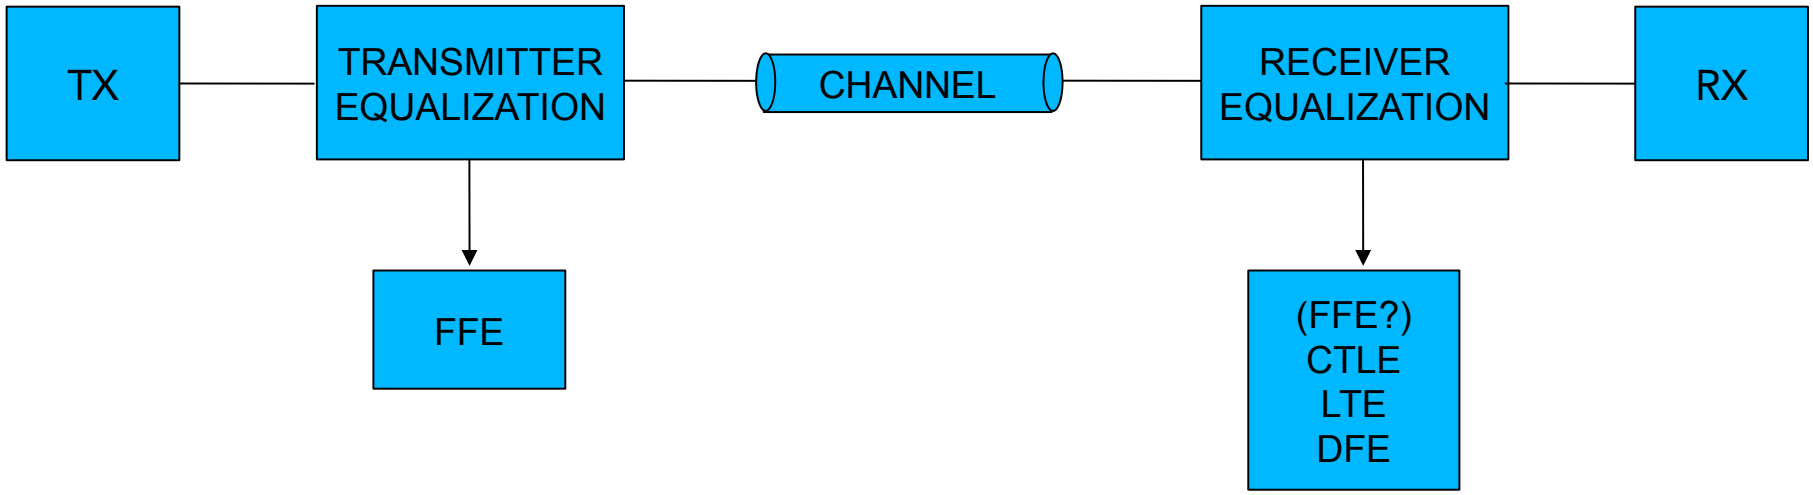

Different kinds:

- Linear vs. Non-linear
- Continuous time vs. Discrete time
- Transmitter vs. Receiver equalization
- Synchronous vs. Asynchronous
- Feed-Forward vs. Feedback-Loop
- Manual vs. Adaptive
- Passive vs. Active
- ...

Most popular acronyms:

- **FFE** → Feed-Forward EQ
- **CTLE** → Continuous-Time Linear EQ
- **LTE** → Long-Tail EQ
- **DFE** → Decision-Feedback EQ

Feed Forward Equalization (FFE)

- It's a **Finite Impulse Response** (FIR) filter:
  - ISI is linear (can be corrected by filtering)
  - Boost bits coming after a transition (HF "pre-emphasis" or LF "de-emphasis")

- Most often seen as de-emphasis **TX-FFE:**
  - Easier to implement in TX (digital domain)
  - Helps for receiver linearity
  - But: Degrades EMI/crosstalk
  - Difficult to be made adaptive
  - Nevertheless widely used whenever the channels response is sufficiently well known in advance (mostly the case)

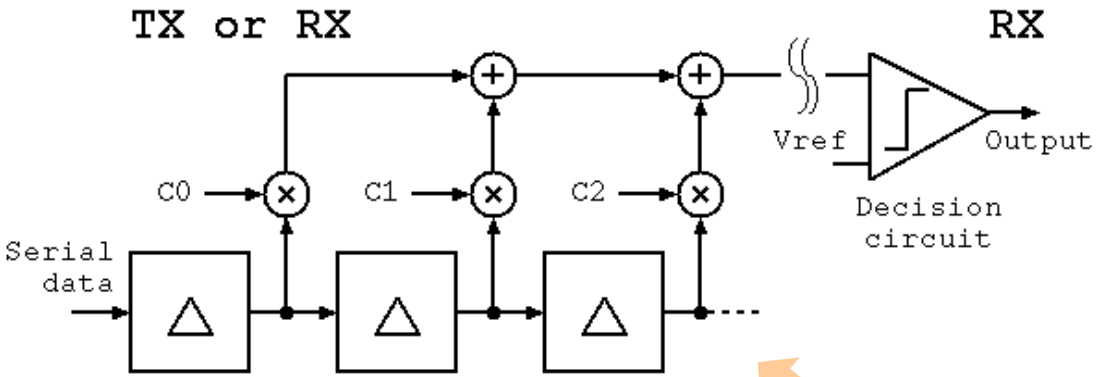

Normally limited to 3-4 taps

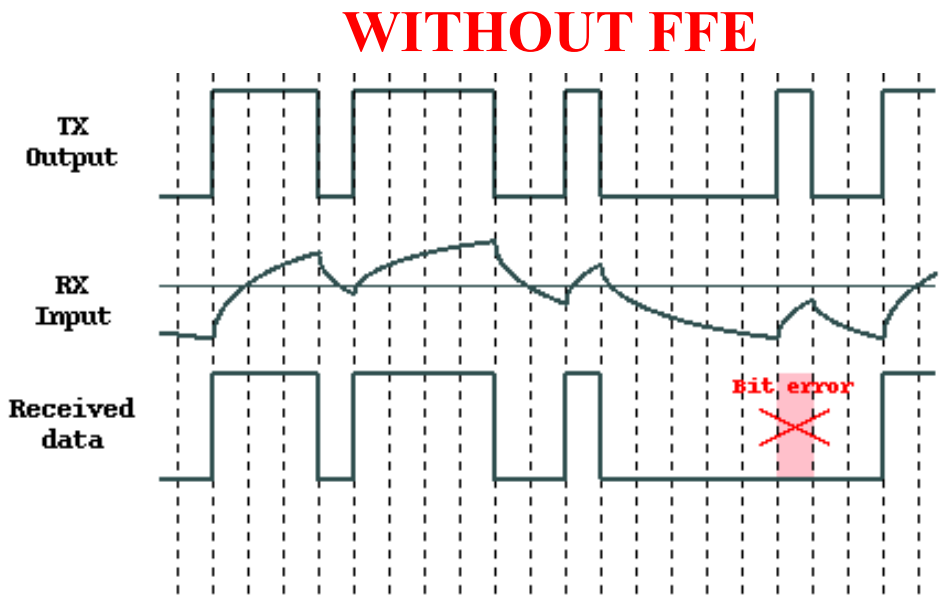

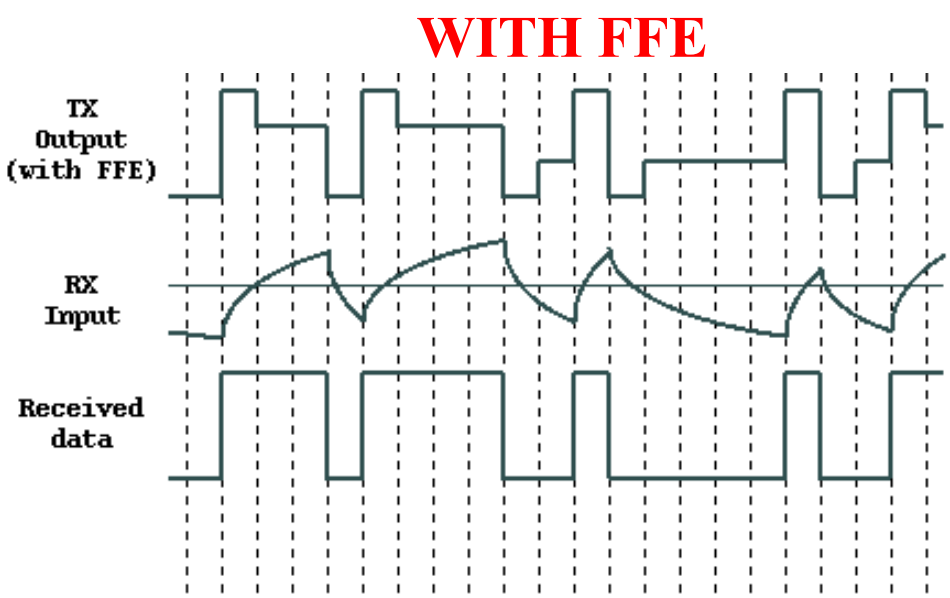

Feed Forward Equalization (FFE)

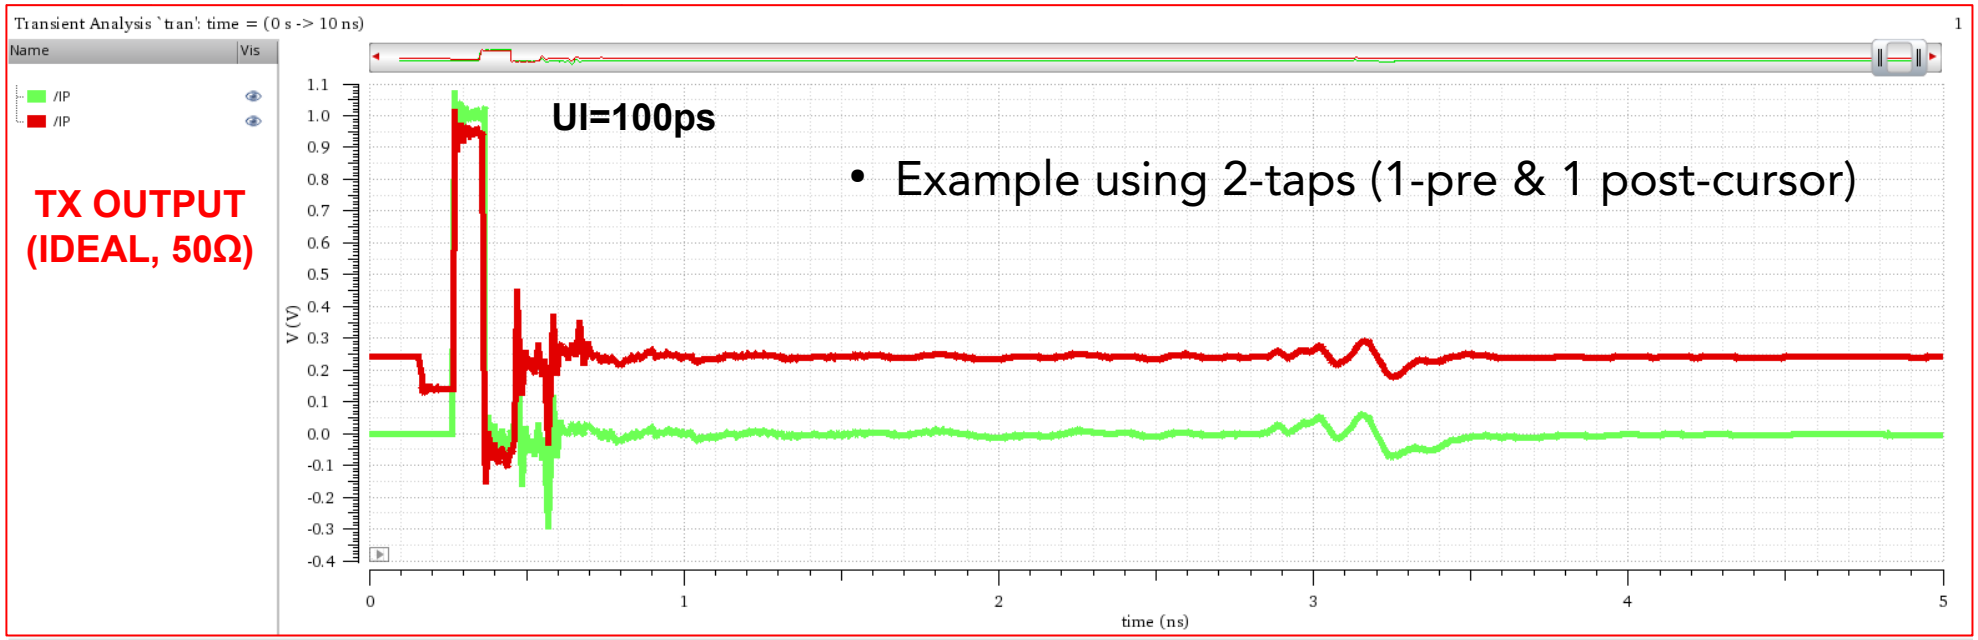

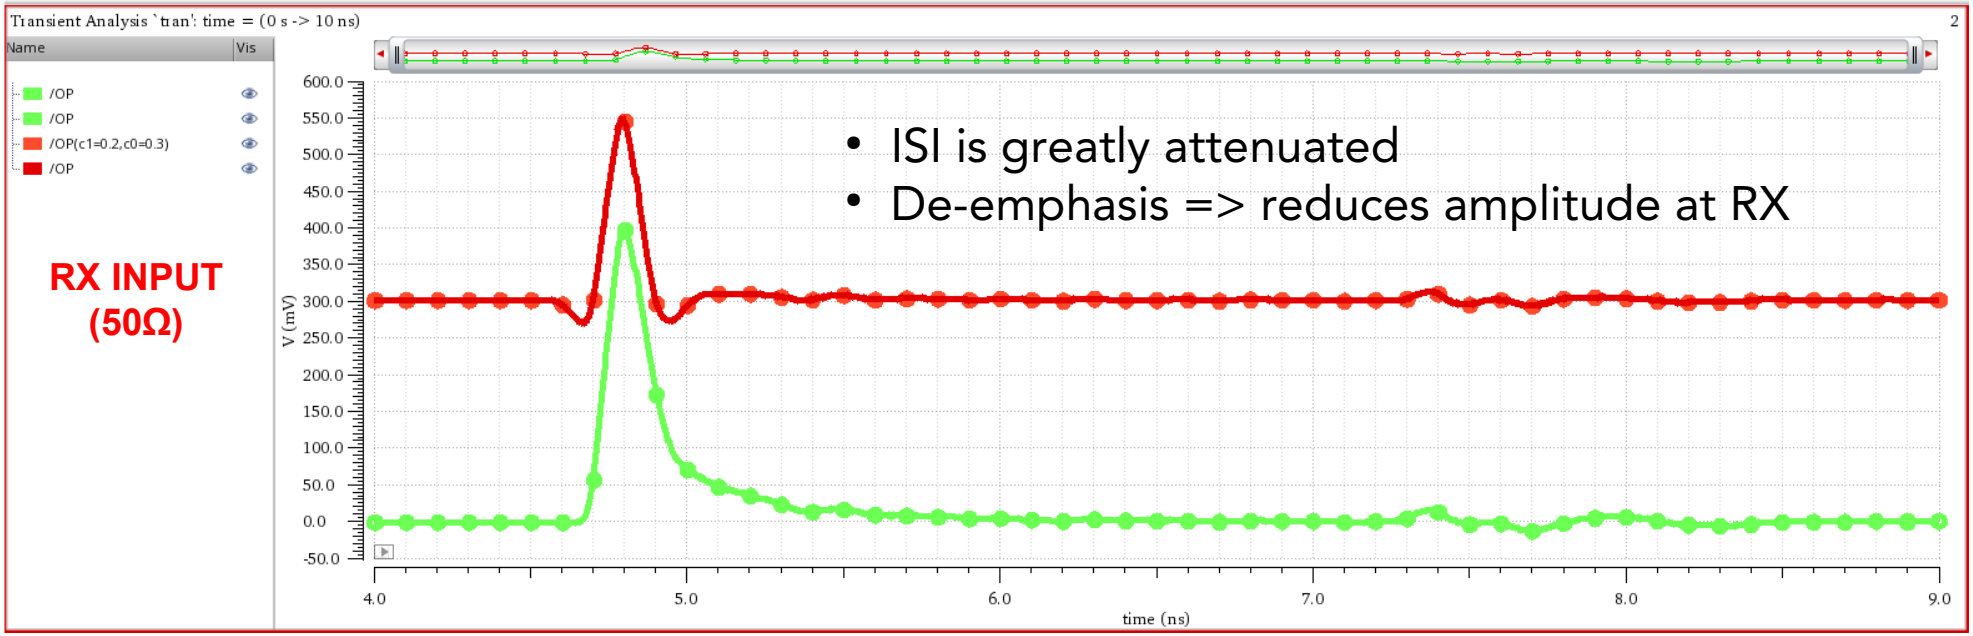

Adding FFE at 8Gbps

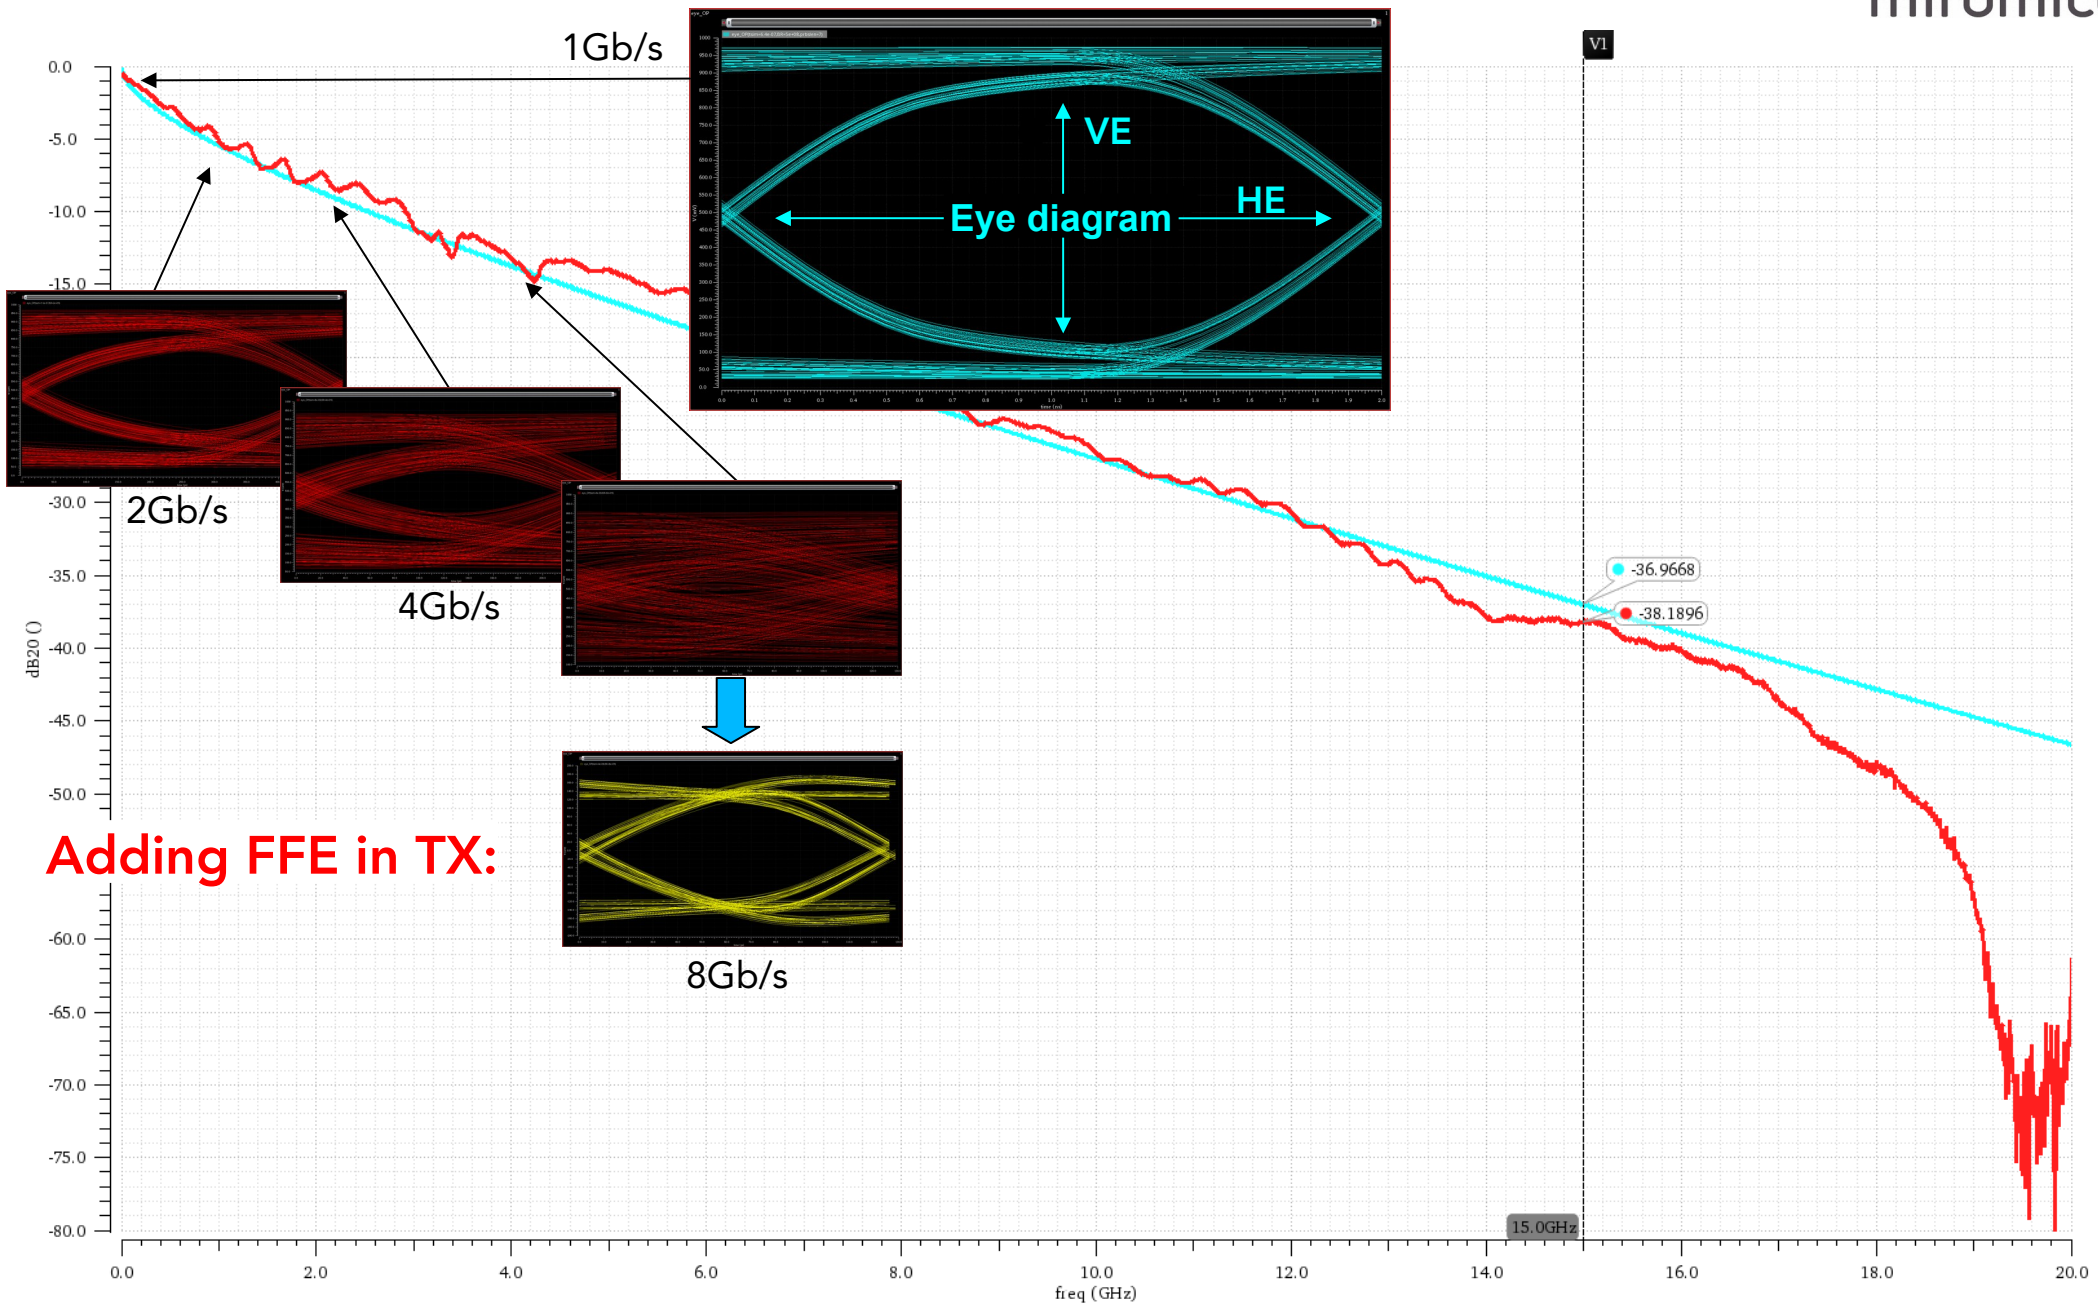

Doubling the rate (16Gbps)

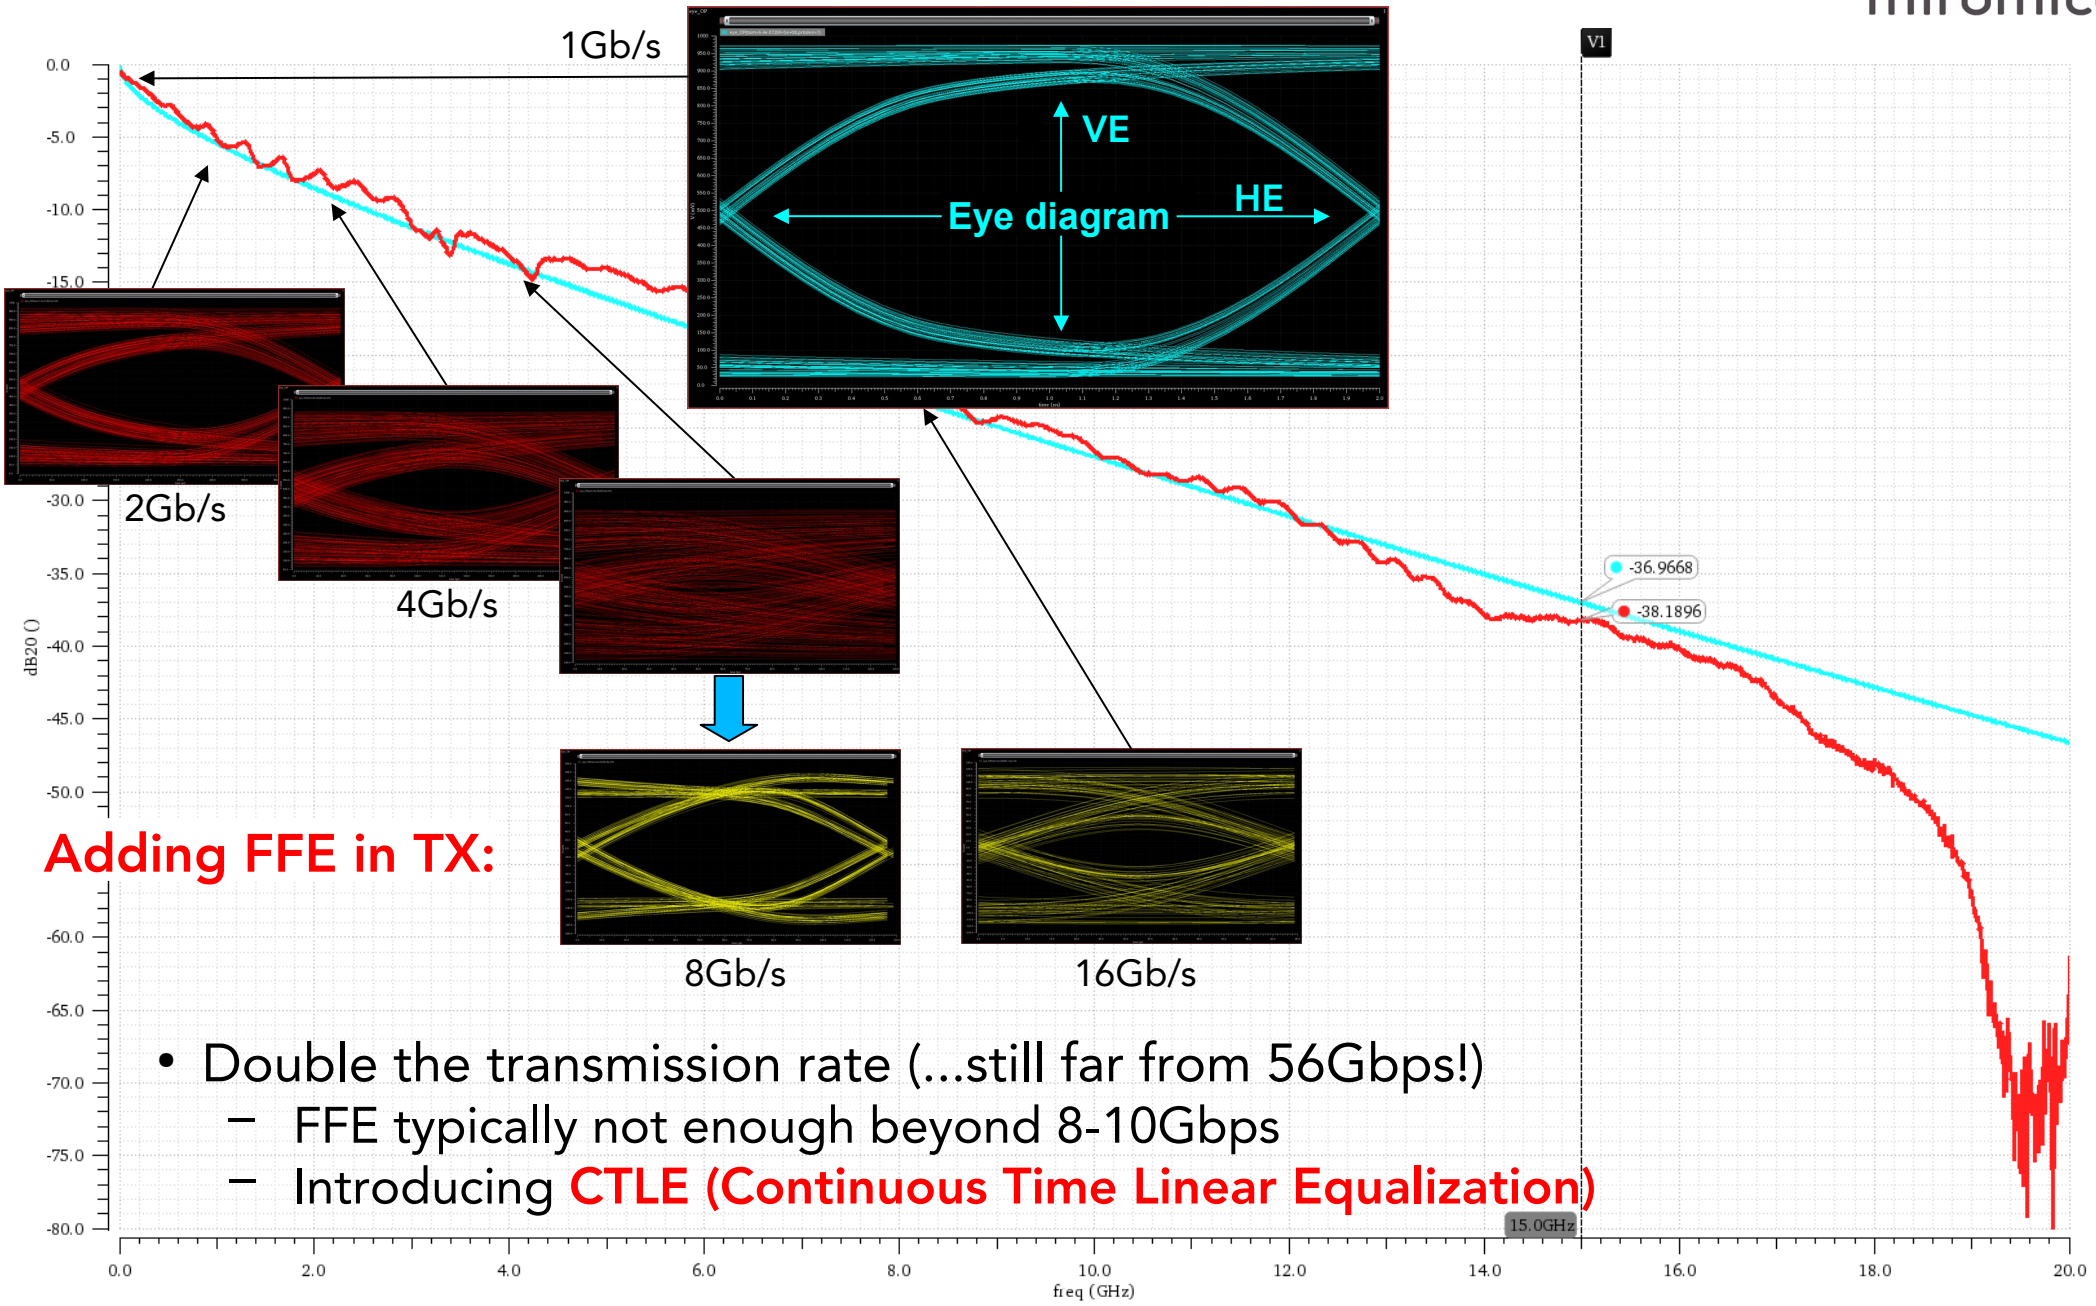

Adding FFE in TX:

- Double the transmission rate (...still far from 56Gbps!)
  - FFE typically not enough beyond 8-10Gbps
  - Introducing **CTLE (Continuous Time Linear Equalization)**

Continuous-Time Linear Equalization (CTLE)

- CTLE counteracts ISI by **boosting high-frequencies** around Fbaud/2
  - Contains at least one zero (peaking amp.), but can have more (LTE, …)
  - Programmable transfer function, for adaptation
  - Unfortunately amplifies noise/interference at the same time
  - Channel equalization typically requires up to 10-12dB peaking (min. across PVT)

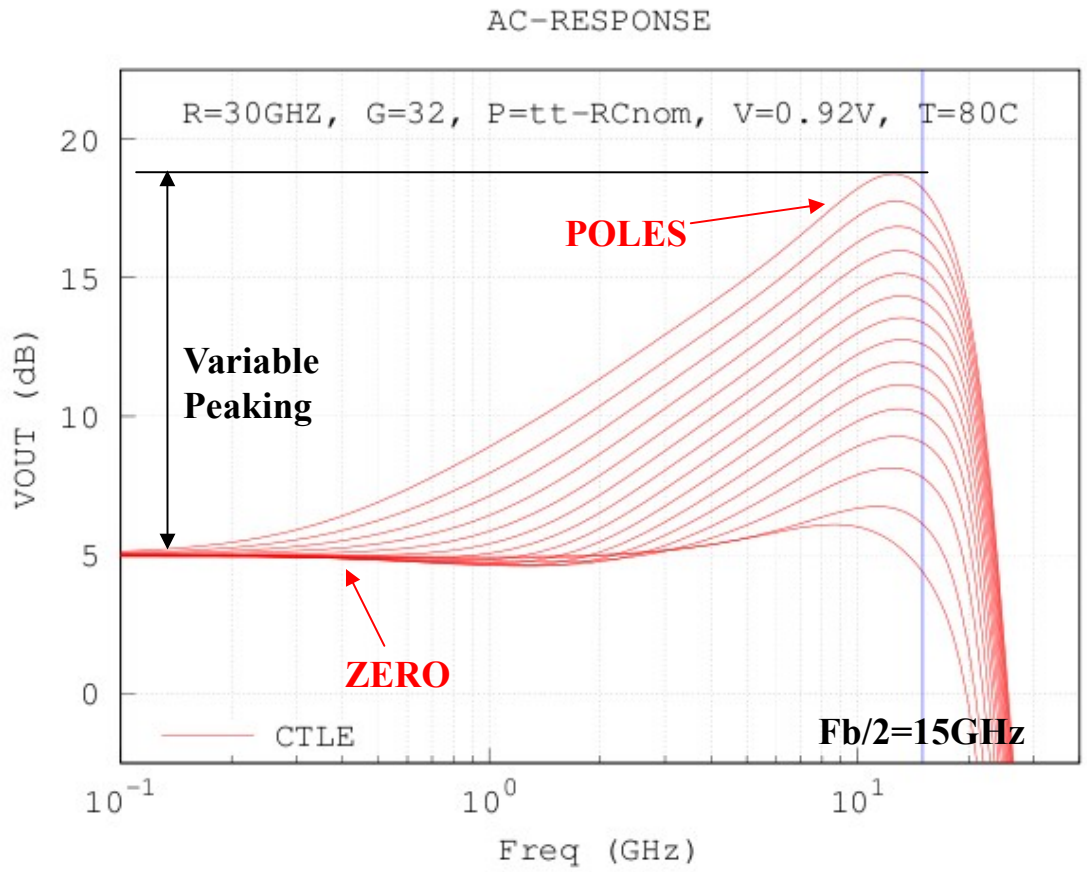

Continuous-Time Linear Equalization (CTLE)

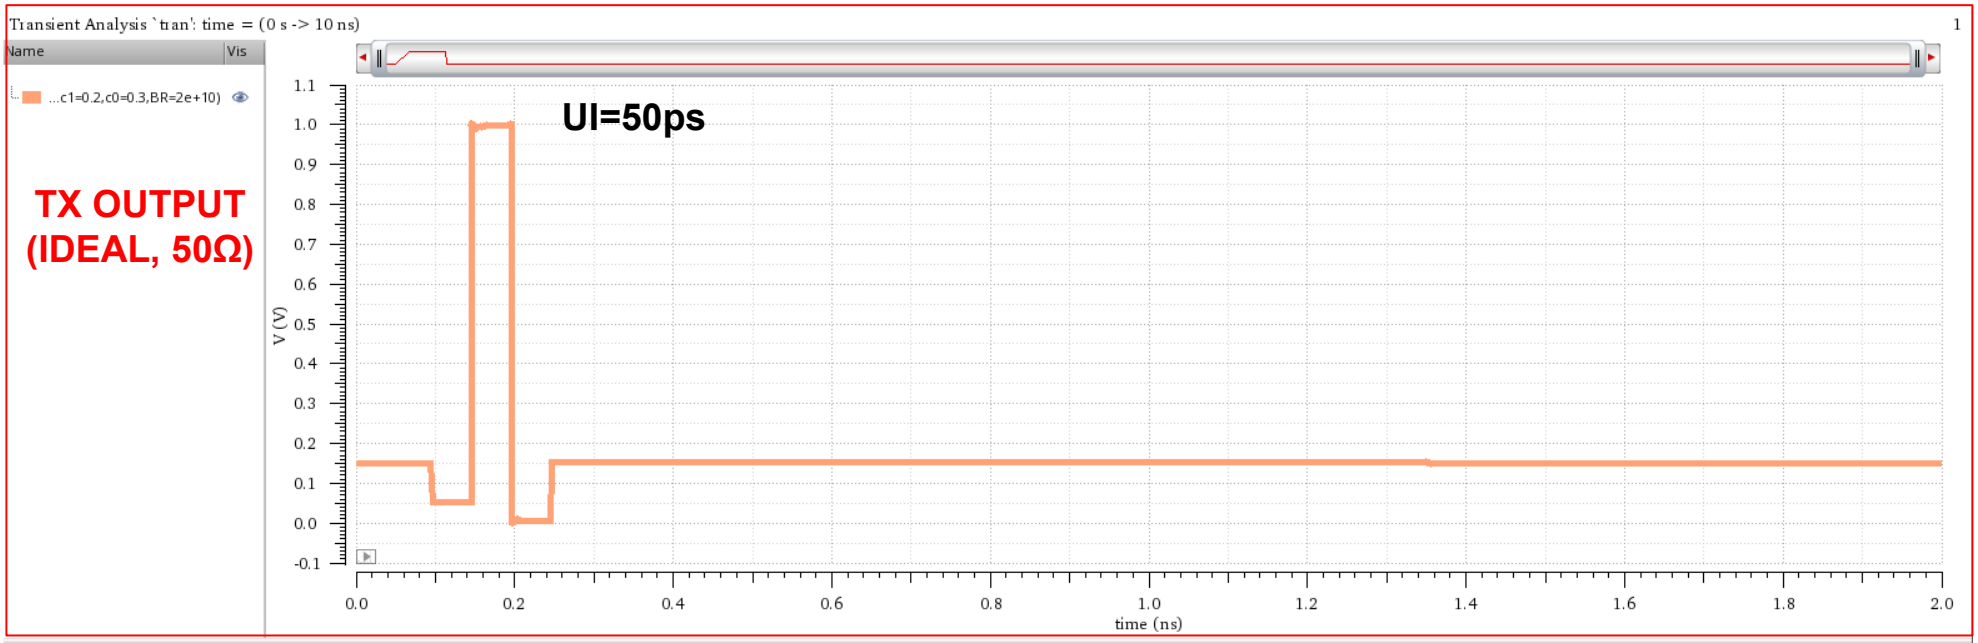

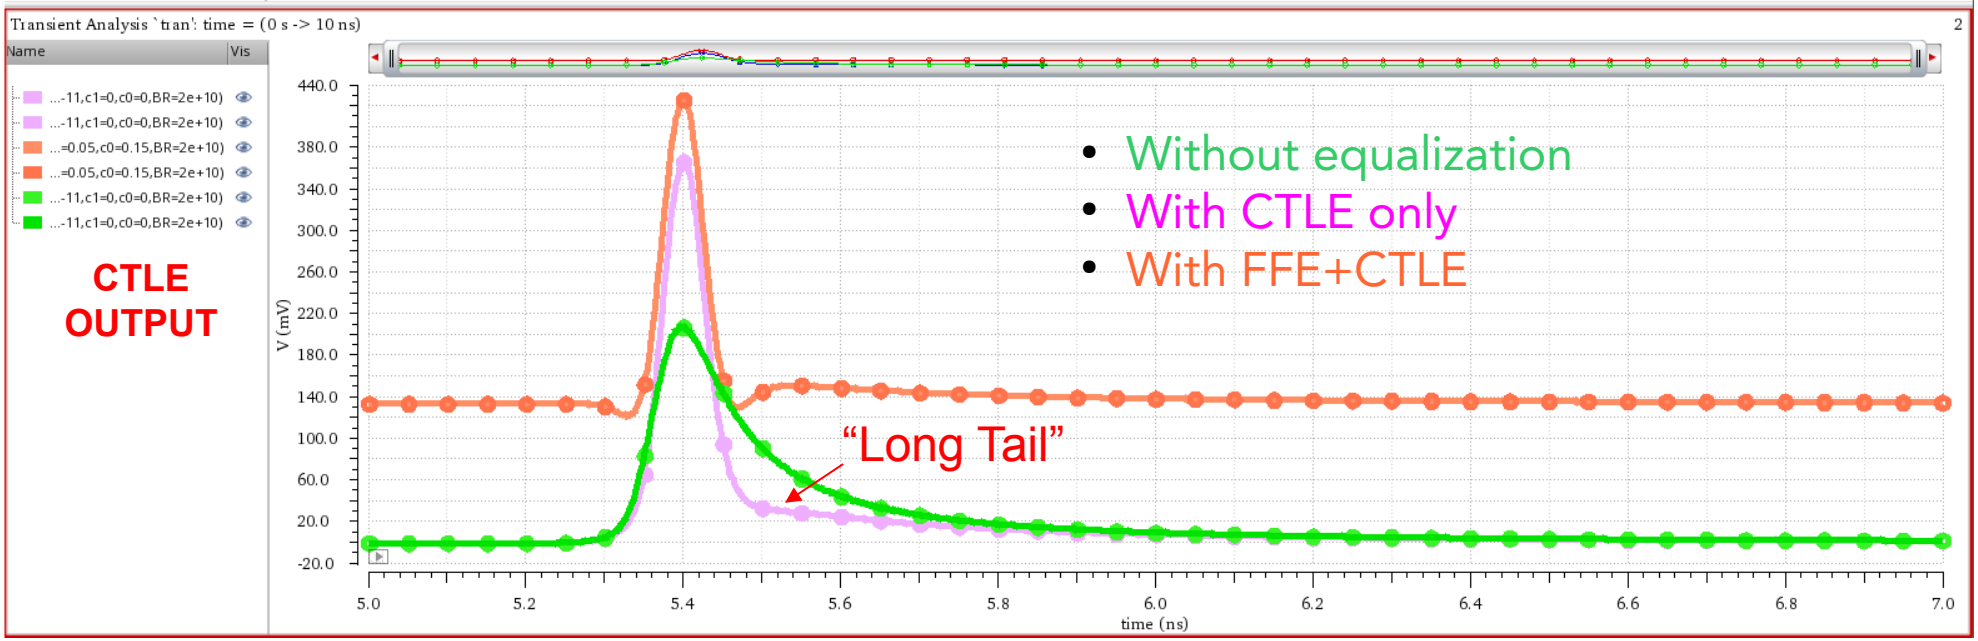

Long-Tail Equalization (LTE)

- LTE is also a Continuous-Time Linear Equalizer
  - Naming underlines that it operates **outside the reach** of other equalizers
  - Most manufacturers add LTE function at 16~20Gbps node (and beyond)
  - Addresses (roughly) low-freq channel attenuation due to skin-effect
  - Need up to 7-10dB emphasis in the mid-band & pole programmability (0.4~2GHz)

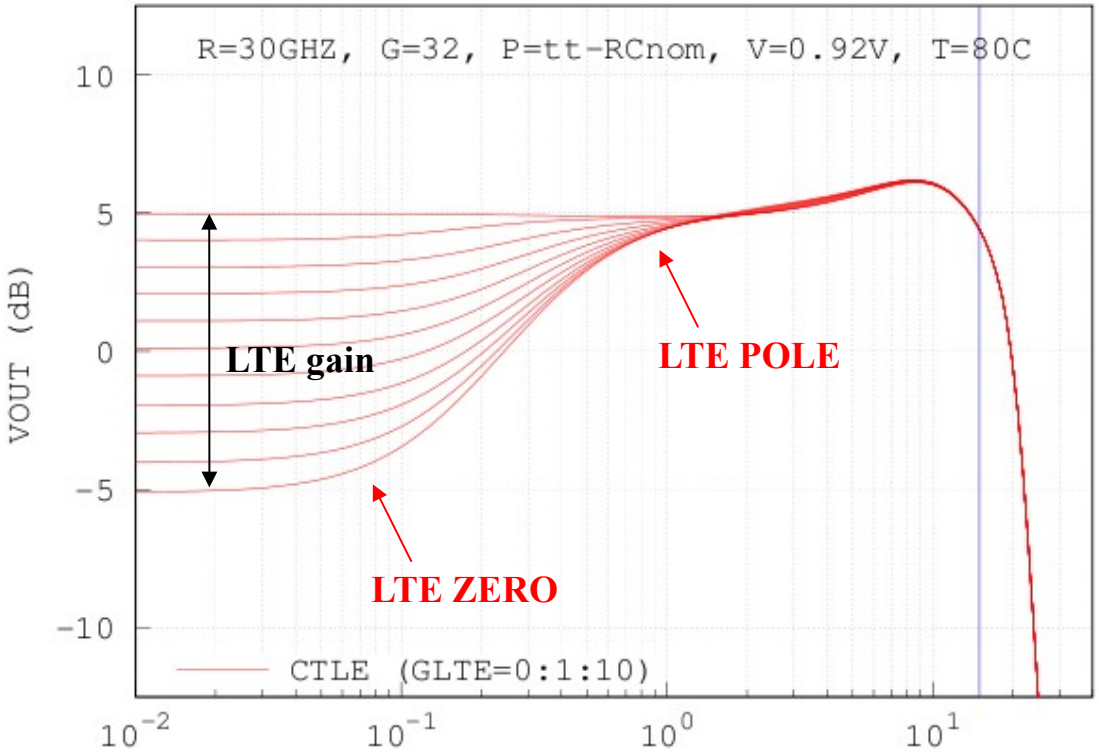

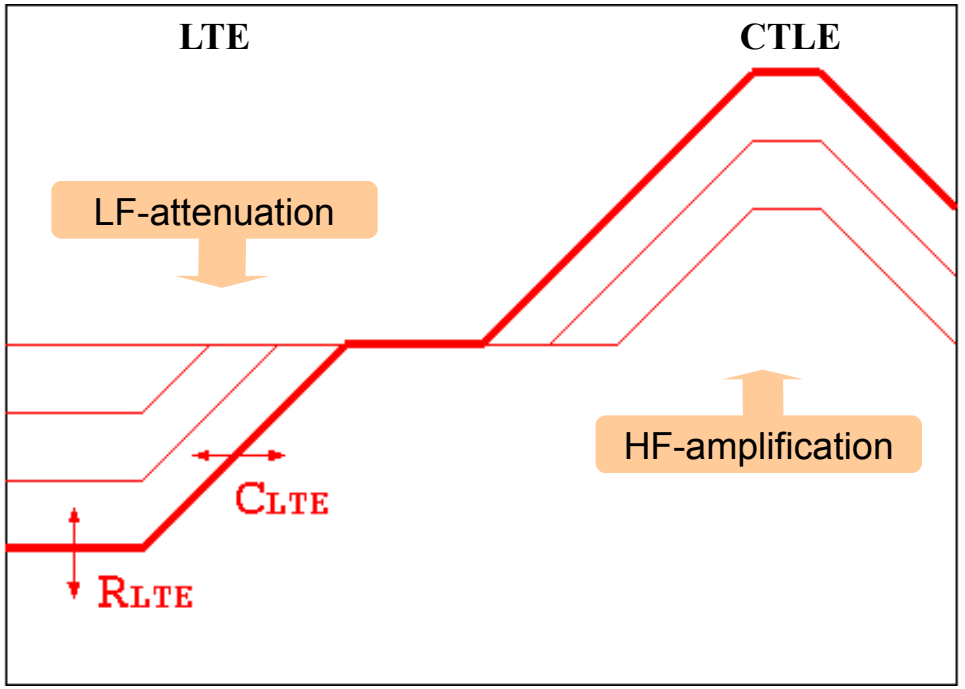

  - Increases EQ capability of the CTLE by combining LF-dump & HF-boost
  - Often requires bulky passive devices (RLC). Tends to become a luxus in modern technologies optimized for digital

Long-Tail Equalization (LTE)

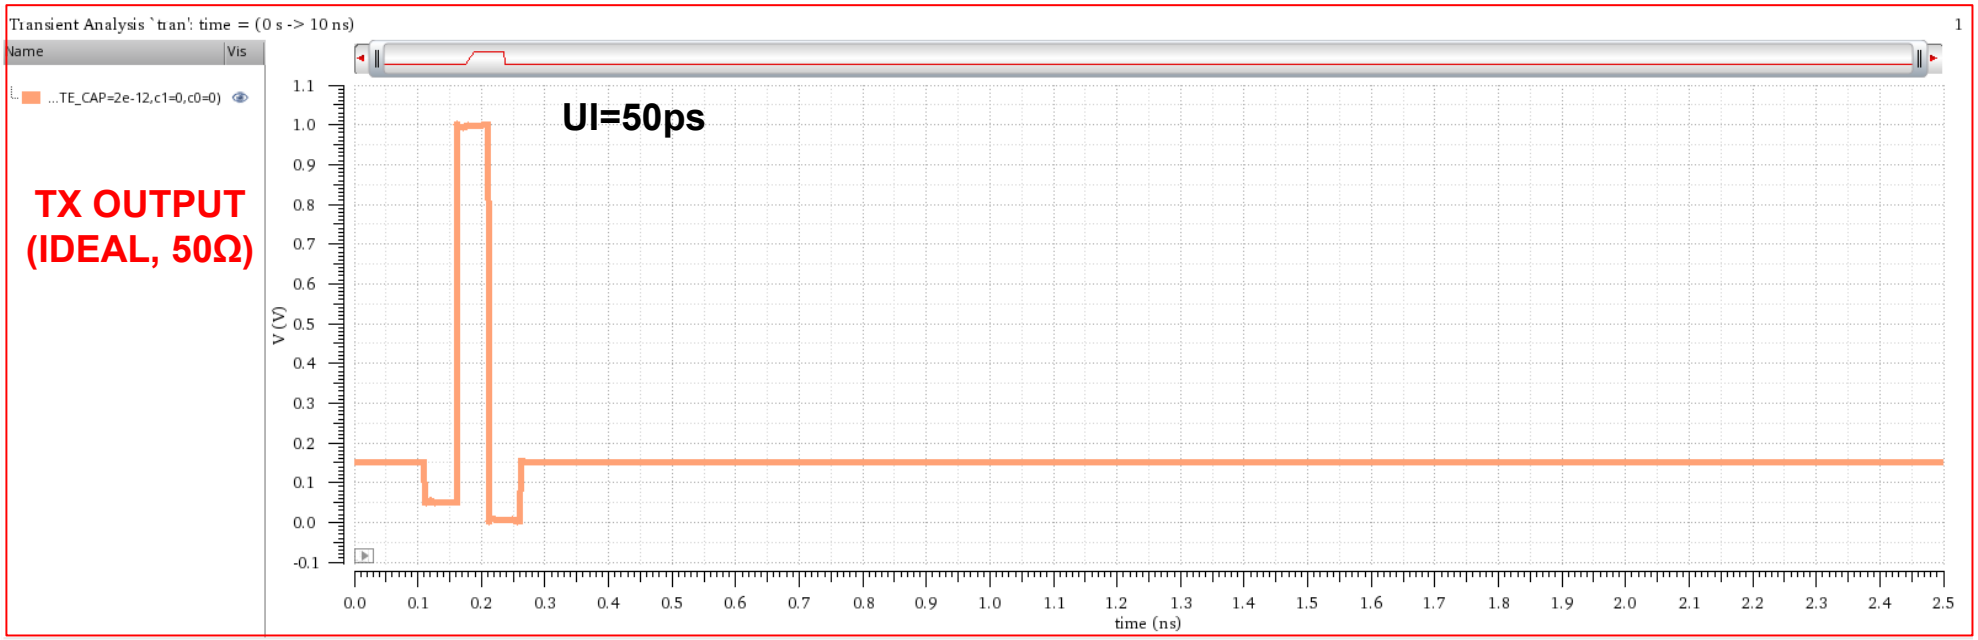

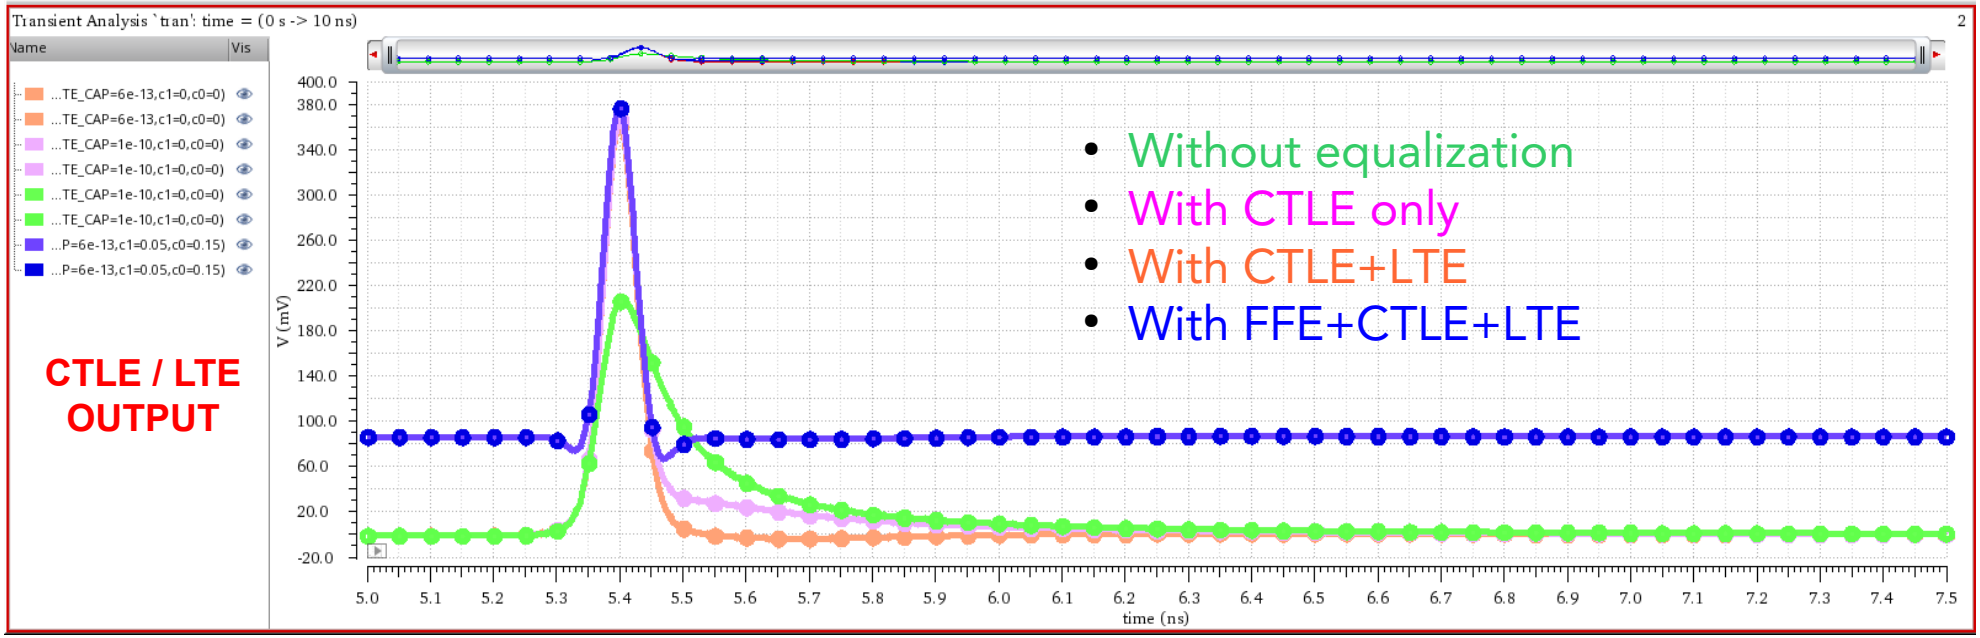

Adding CTLE/LTE at 16Gbps

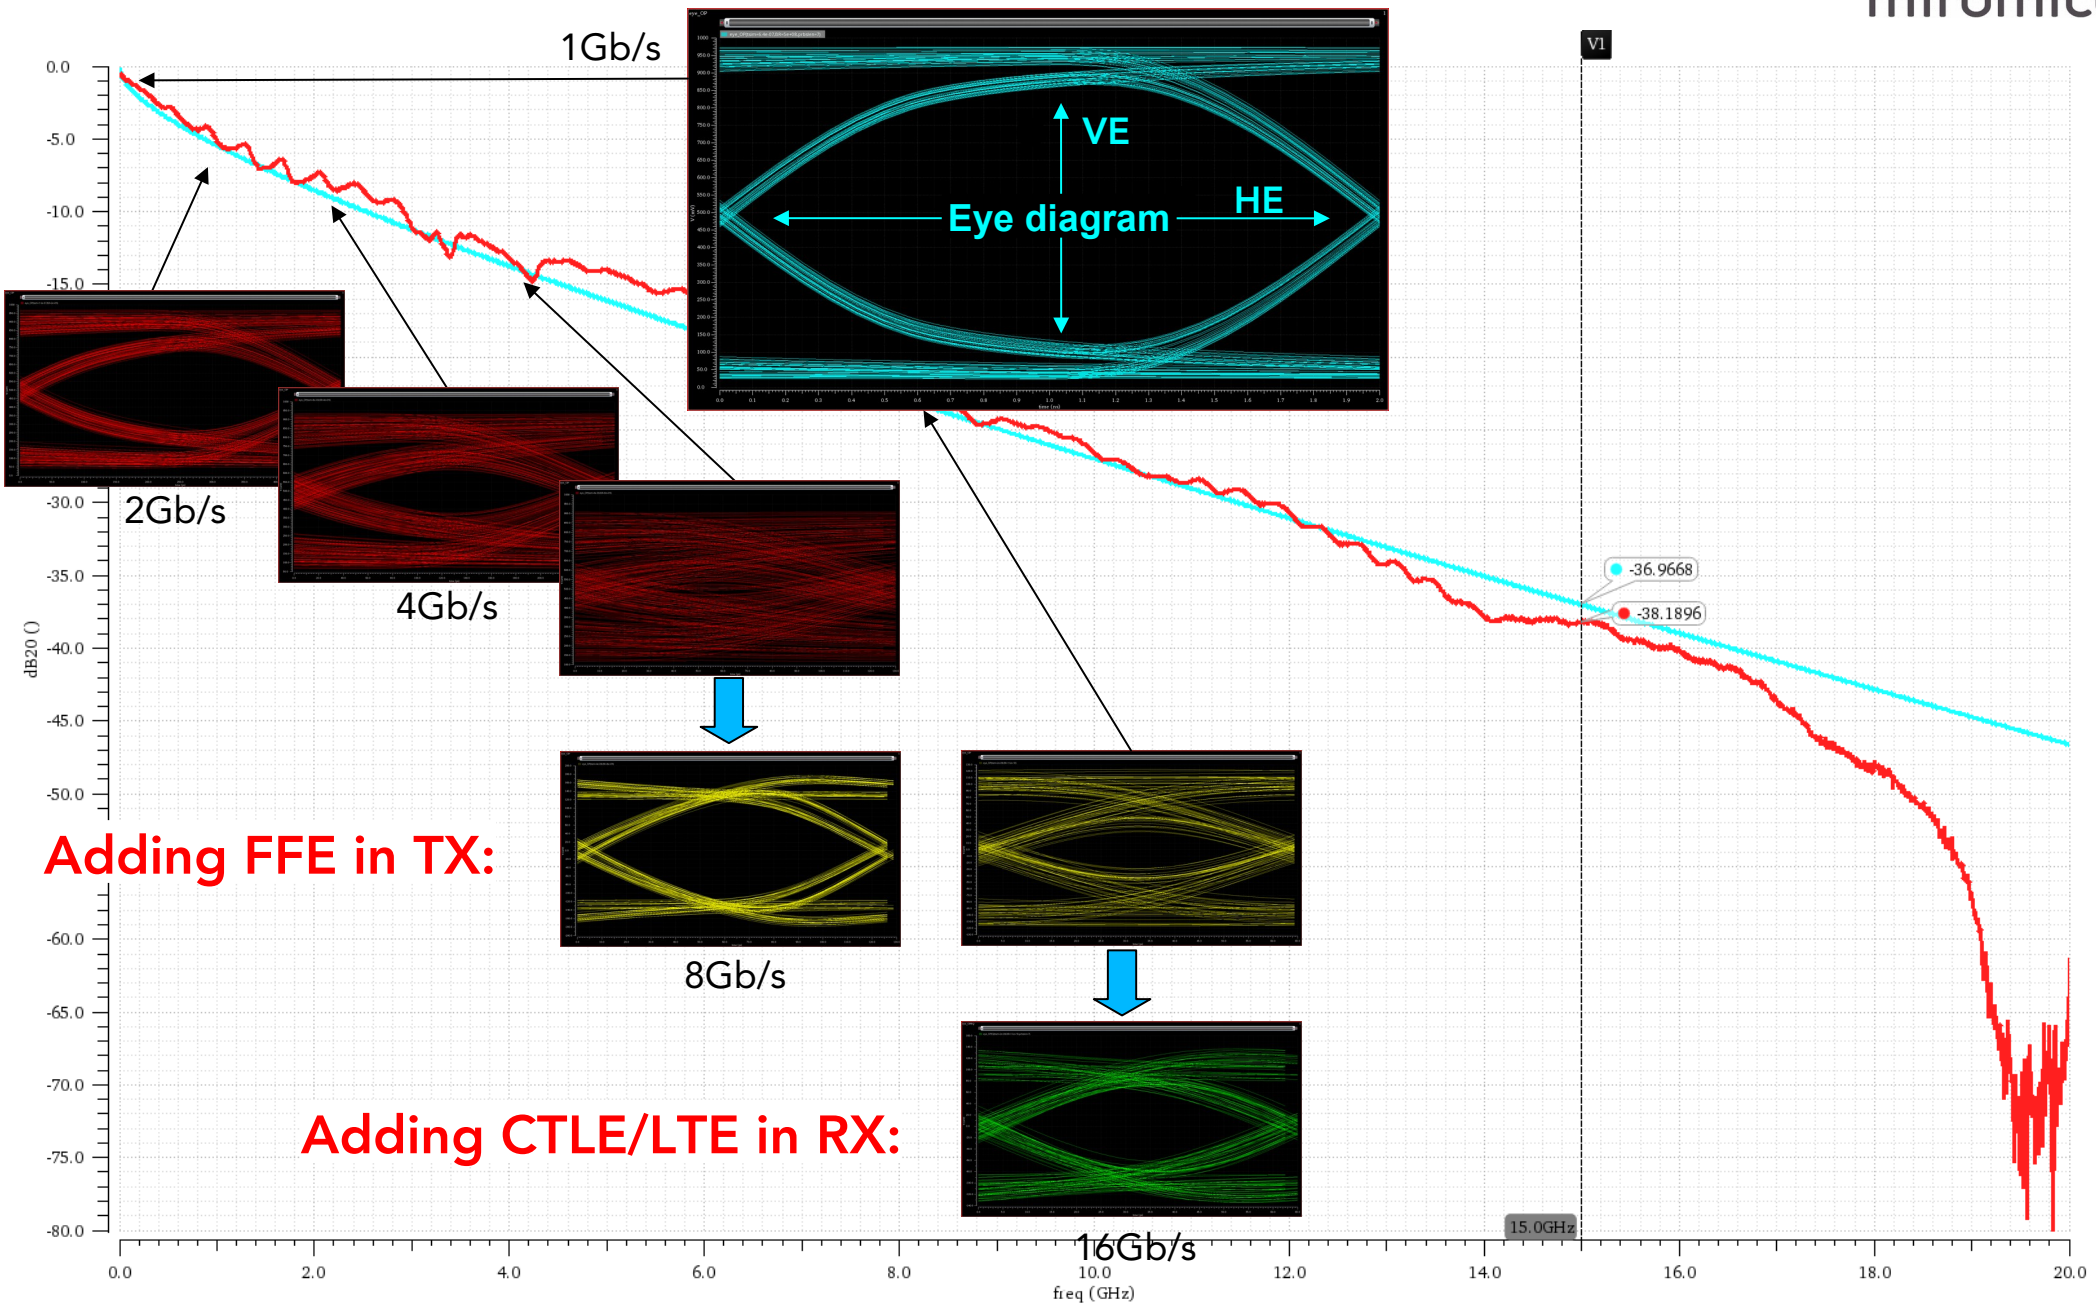

Doubling the rate (32Gbps)

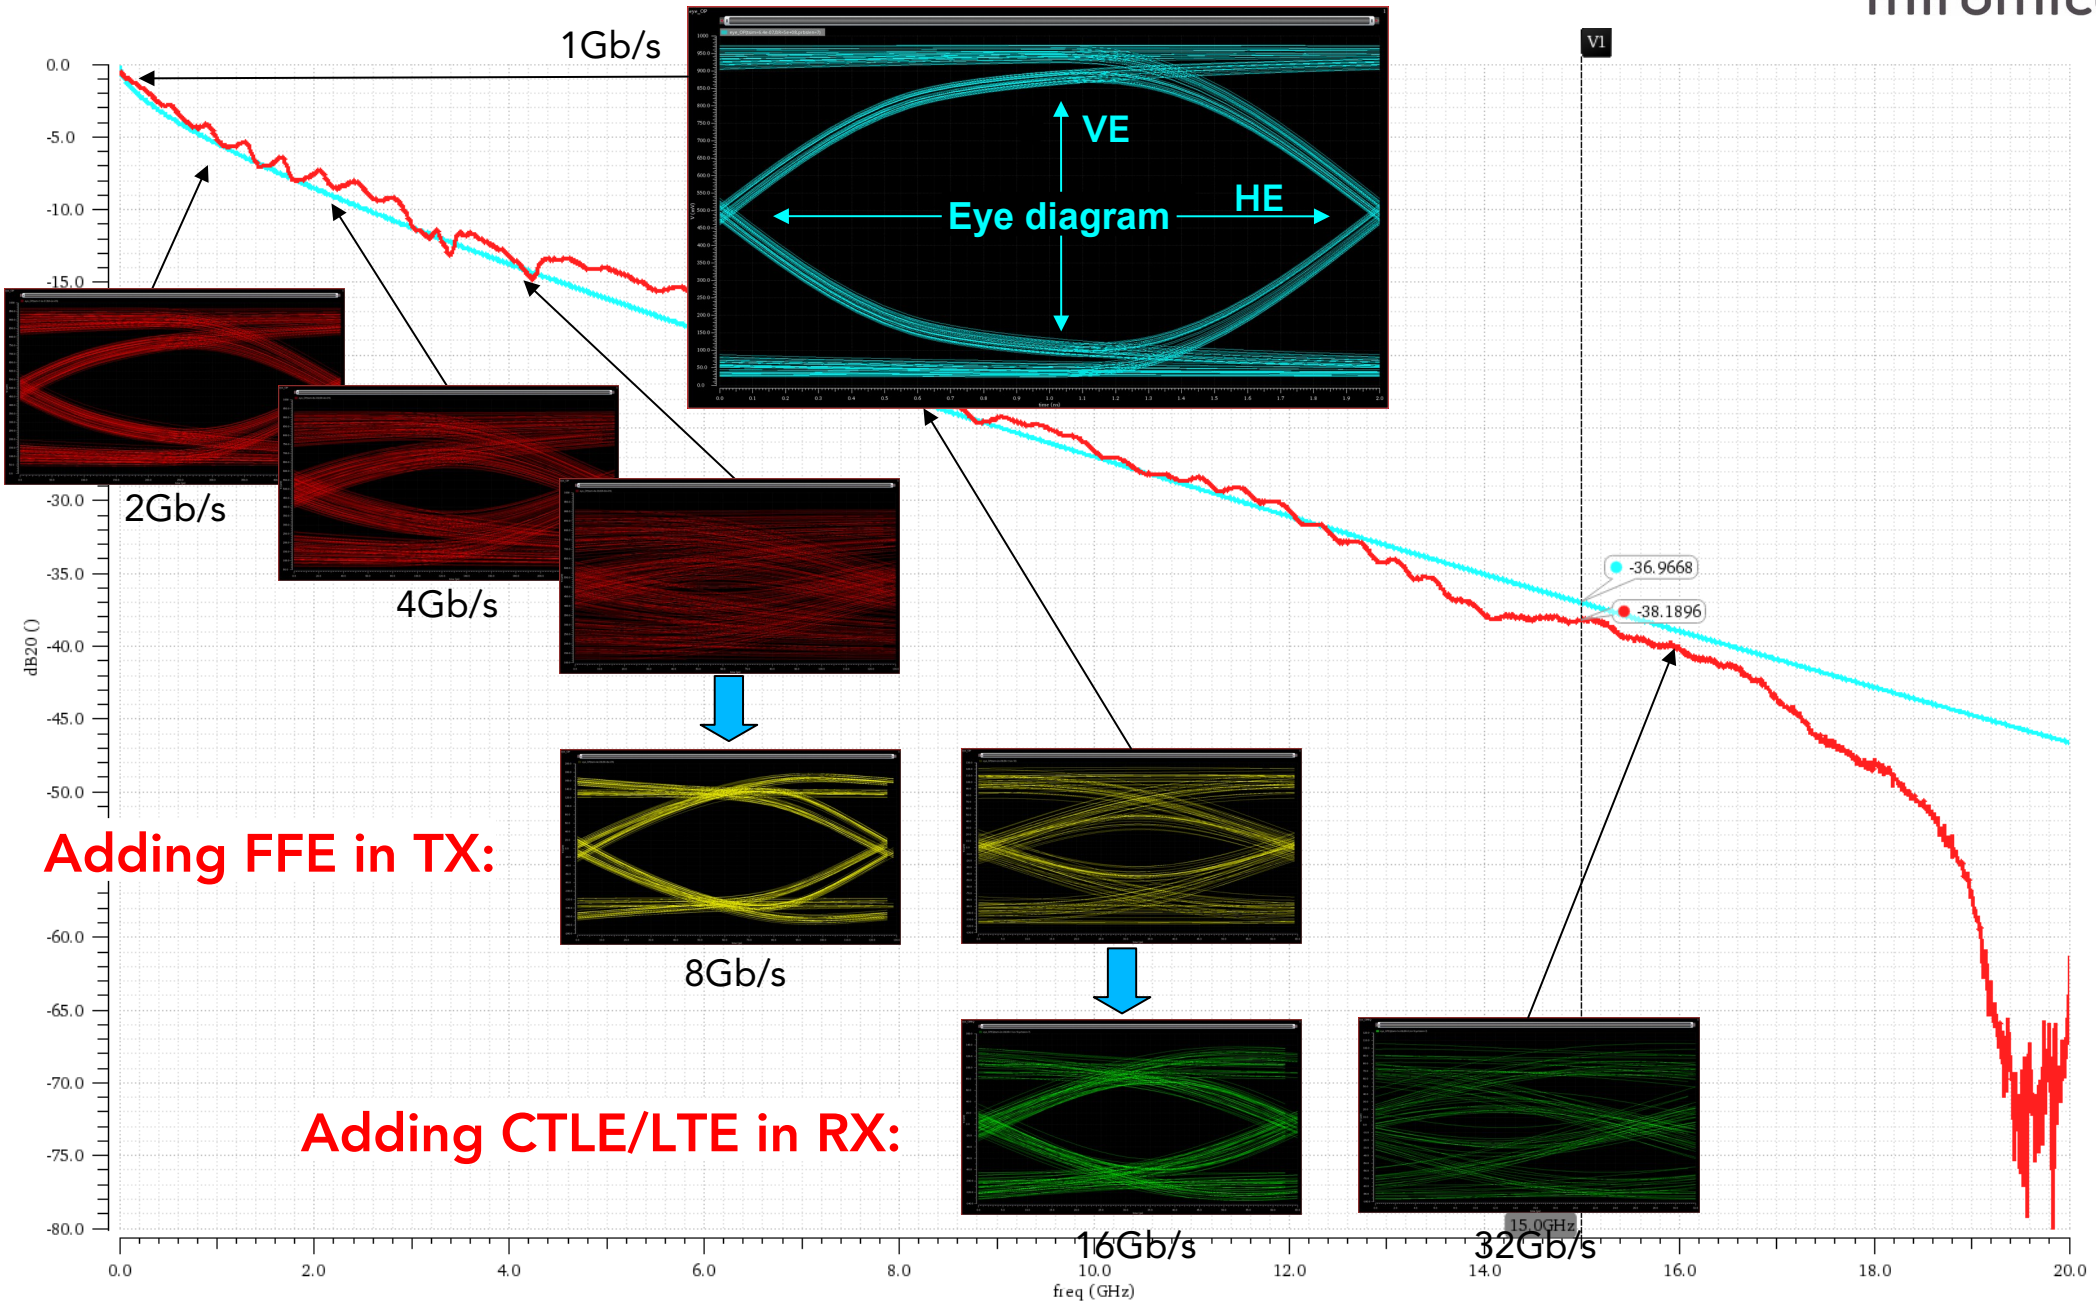

Decision Feedback Equalization (DFE)

- FFE/CTLE have several shortcomings:
  - On 35+dB loss channels simple FFE/CTLE equalization becomes very challenging
  - Amplify signal, noise and interference, without necessarily recovering SNR
  - It is anyway just a rough approximation of the inverse channel-response
  - CTLE/LTE tend to be bulky

- On the other hand: we really only need to **cancel ISI at the sampling point!**
  - If we "know" the previous symbols (the already received ones)
  - And we know how a single bit is distorted going through the channel (ISI)
  - We can calculate the ISI (using a FIR) and remove it from the incoming bit

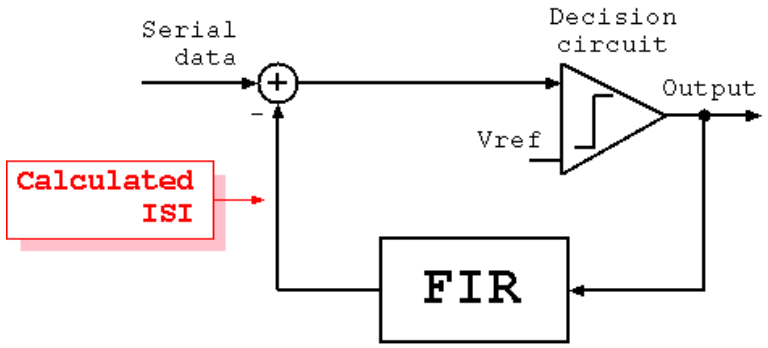

[Austin 1967]

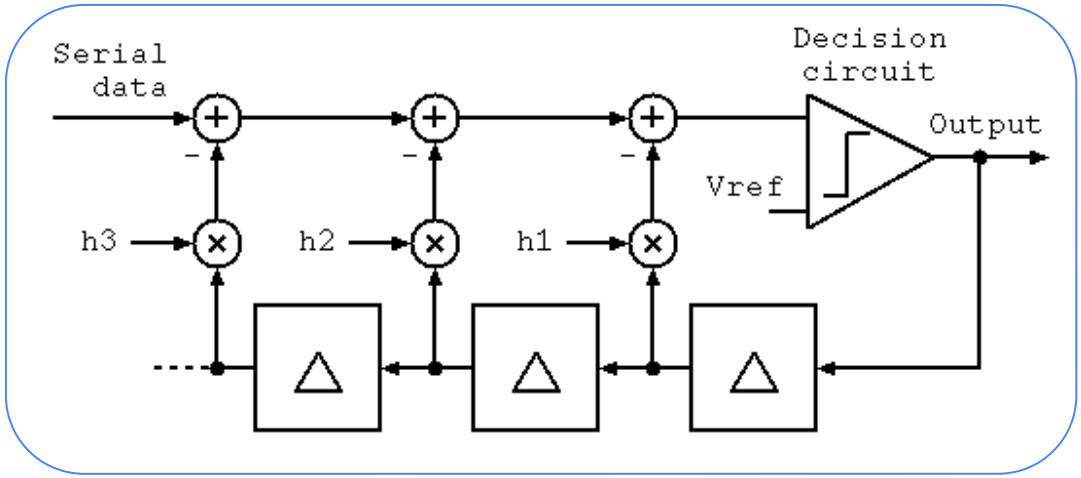

Delay → Scale → Substract

DFE: THE GAME CHANGER!

Decision Feedback Equalization (DFE)

- DFE does not attempt to invert channel, just remove ISI before taking the decision
  - Does not amplify noise and interference
  - Can cope with an arbitrary channel response (takes care of reflections)
  - Time "depth" is limited by number of feedback taps (trade-off vs power)
  - Error accumulation problem (bursts of errors)
  - Time critical & power hungry (but feasible: demonstrated in CMOS down to UI~30ps)

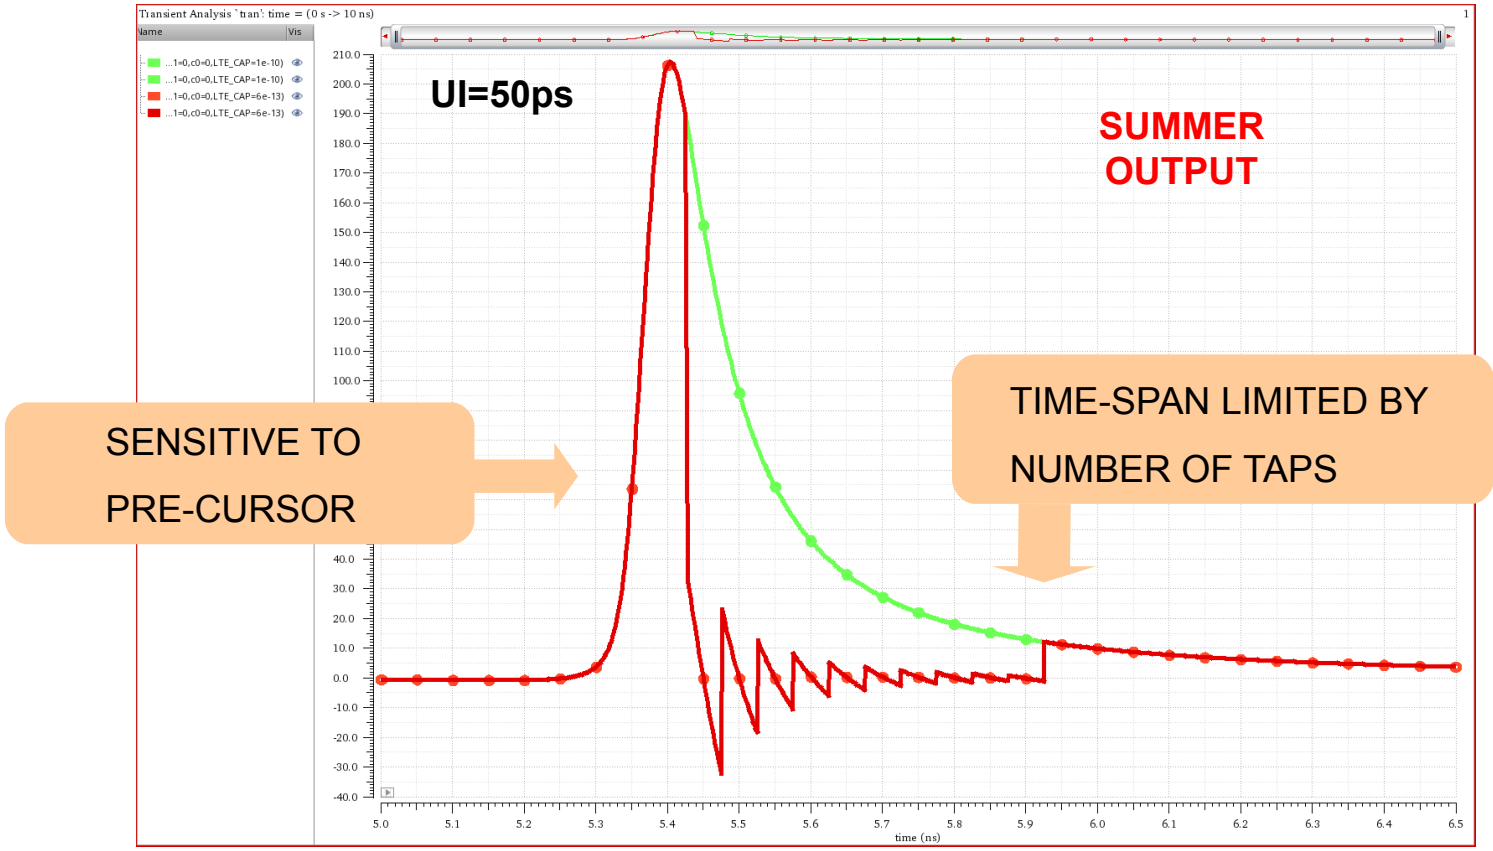

Adding DFE at 32Gbps

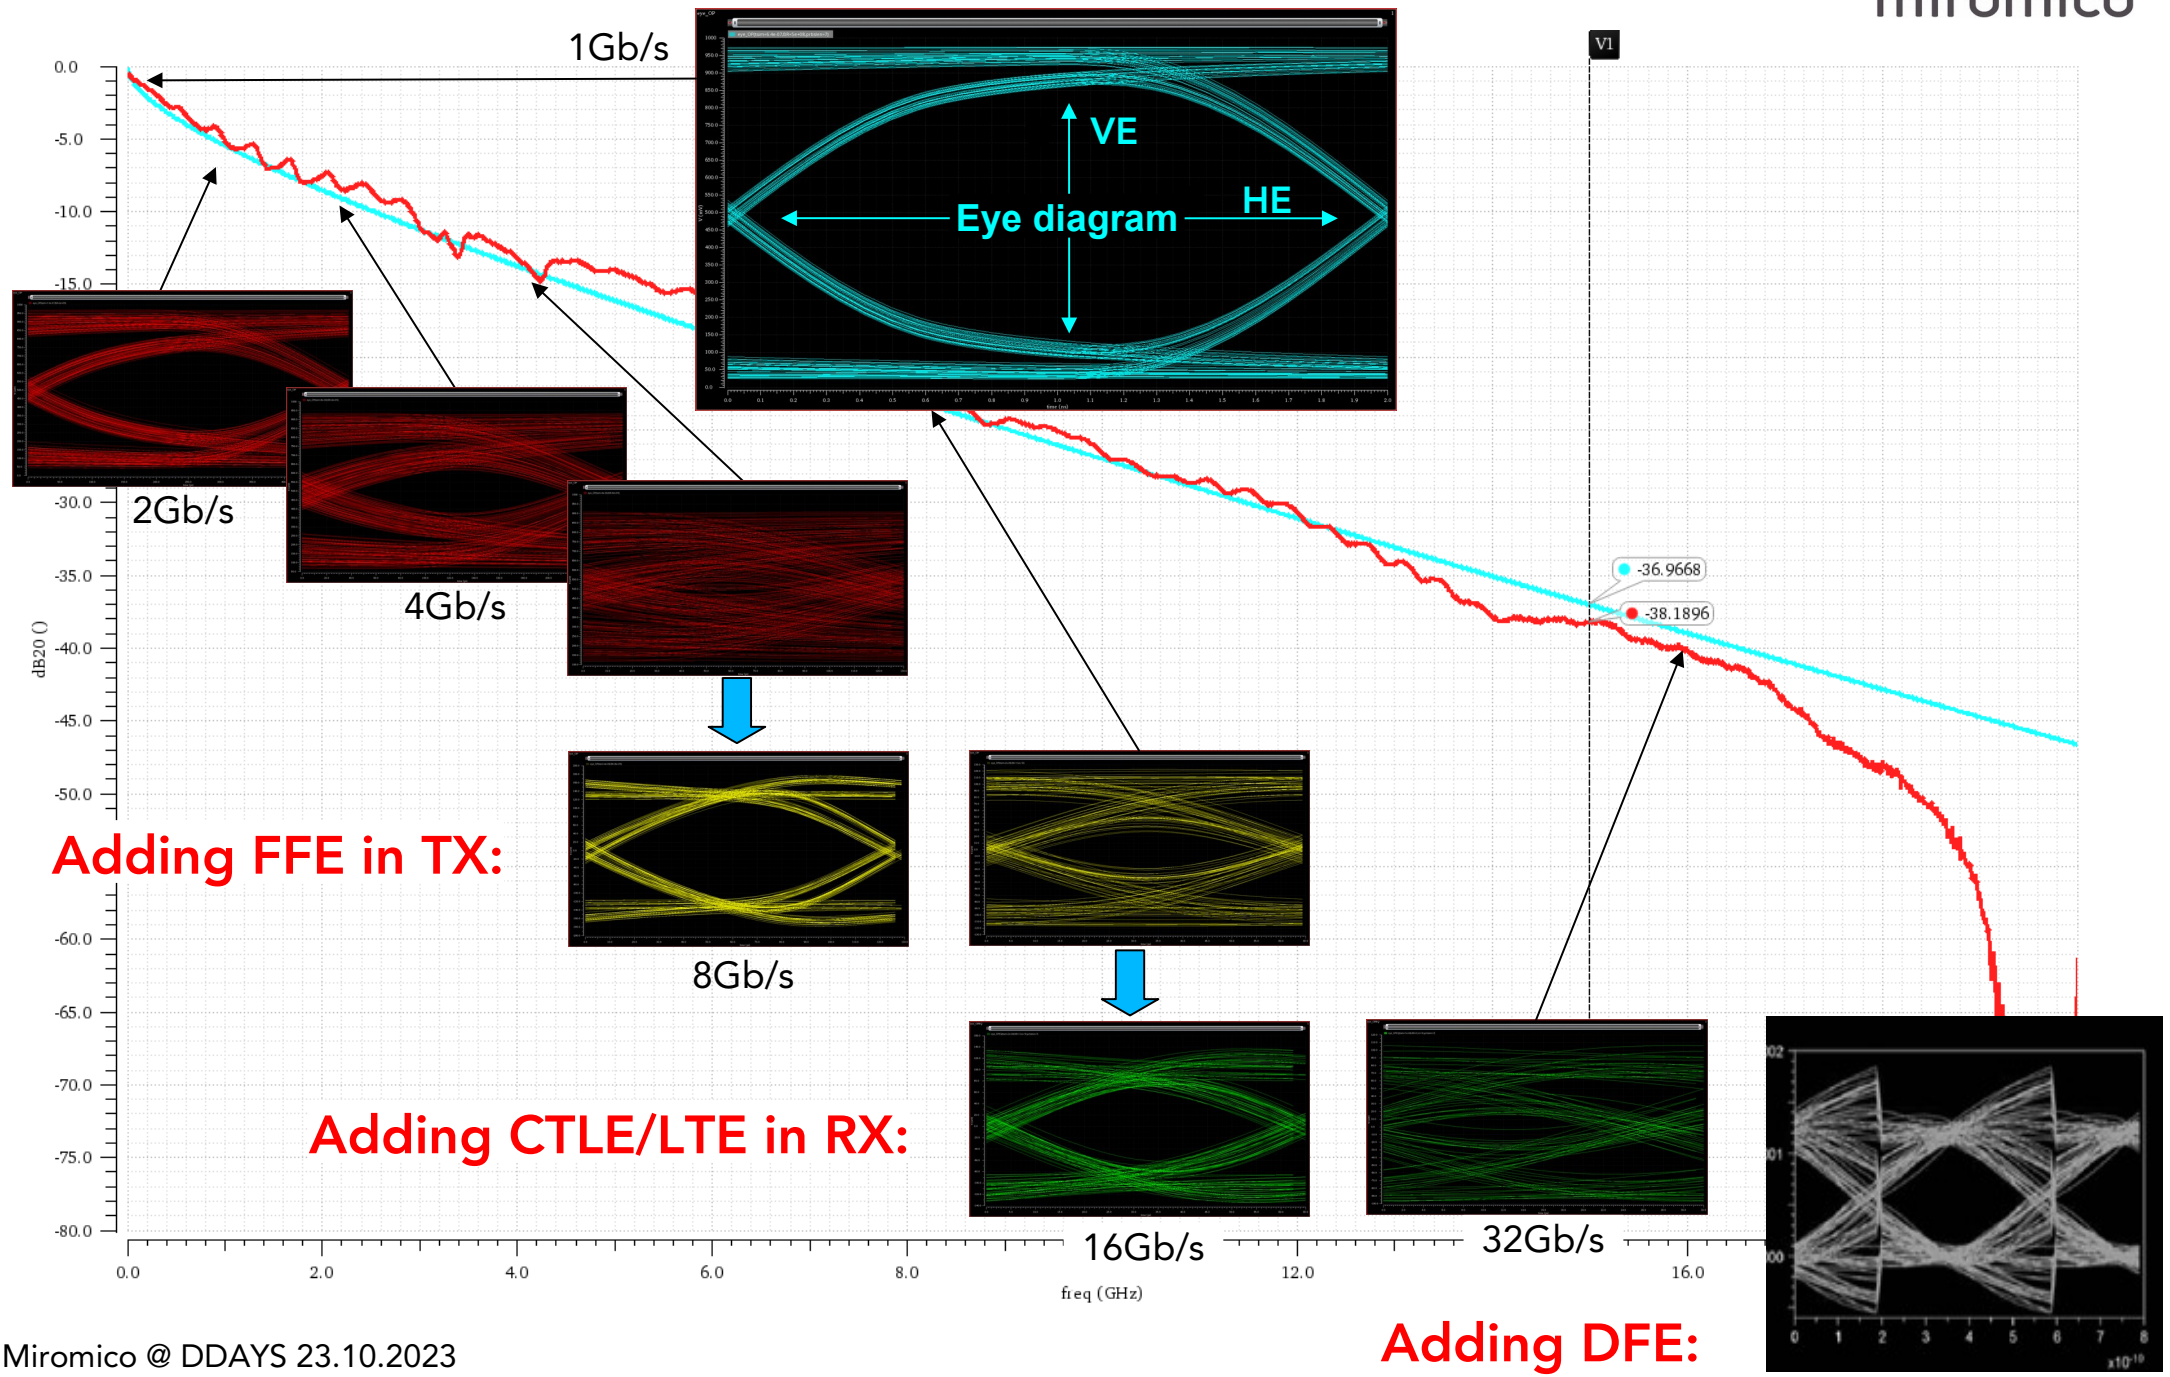

The right mix

	Advantages	Disadvantages
FFE	- Simple	- Non-adaptive - EMI, Xtalk
CTLE / LTE	- Adaptive	- Amplifies noise - Rough correction - Bulky
DFE	- No noise/xtalk amplification - Works on arbitrary channel	- Error accumulation - Complex, timing critical - Time-span limited by power

- FFE/CTLE/LTE usually "just" functional to robust lock of CDR/DFE
- DFE has also limited applicability and tends to be power hungry
- **Combined use of different equalization types to obtain the best results**
- Other kinds of equalization exist but are not mentioned here

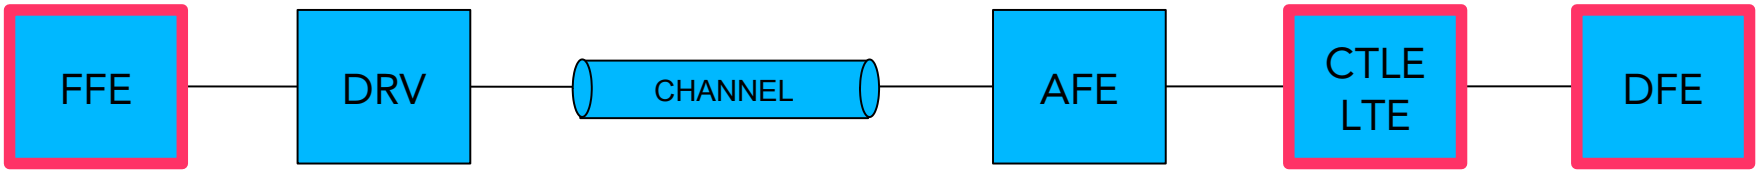

SERDES design in advanced CMOS

- Motivation & Challenges
- Architectural solutions
- **Implementation details**
- Future trends & remarks

Line Driver (DRV)

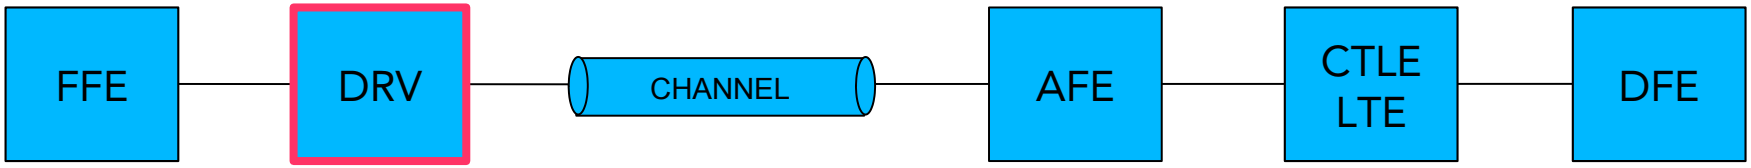

- Line termination
- ESD/LU protection
- Broadband adaptation
- CML vs. SST

Line termination & protection

- Current mode driver (CML)

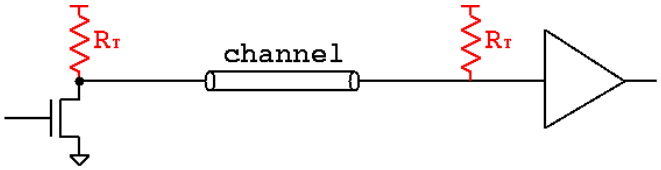

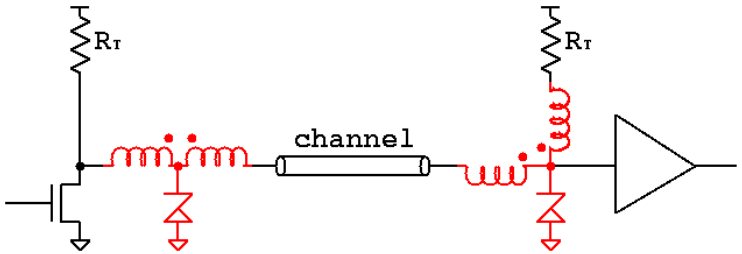

- Voltage mode driver (SST)

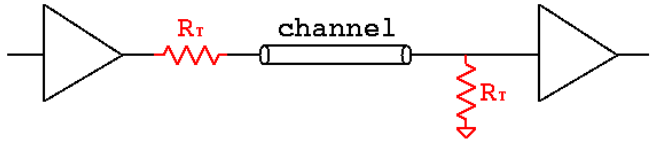

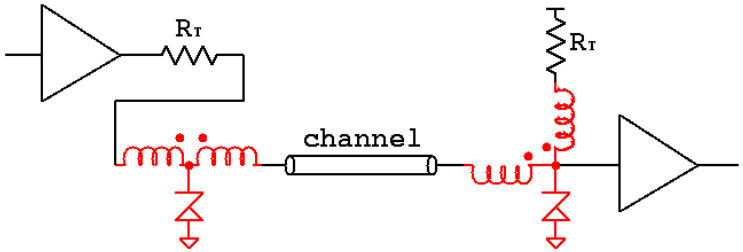

- Line termination (& back-termination) is a MUST for high-speed wireline comm.
  - Less signal but integrity: broadband transfer, avoid multiple reflections
  - Parallel (CML) vs. Series (SST) termination
- Unfortunately ESD/LU protection is also a must for manufacturability
  - Protection diodes represent a significant parasitic capacitance
  - T-coils for better bandwidth & return loss [Galal JSSC 2004]
  - Series termination (SST) offers superior ESD protection

Line Driver

- Current mode driver (CML)

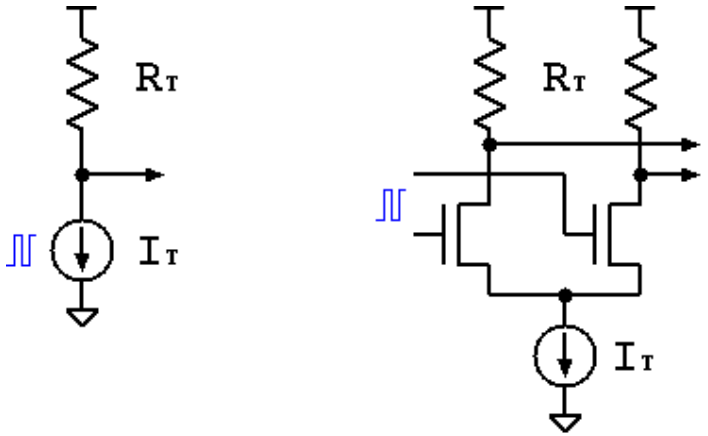

- CML stands for Current Mode Logic
- Unipolar (can use the fastest transistor)
- Purely differential
- Need higher voltage & current for same launch amplitude
- Input drive not directly CMOS (somewhat tricky)
- Asymmetric output drive

- Voltage mode driver (SST)

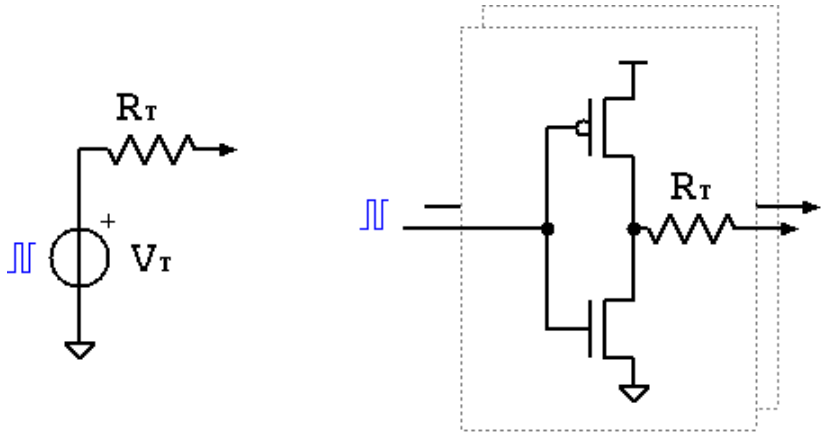

- SST stands for Series Source Terminated
- Switch output between supplies (need almost symmetrical devices). Uses the full supply
- Pseudo-differential
- Input drive is digital, CMOS compatible
- Symmetric output drive
- Switches impedance is part of the termination (more variable)

Line Driver

- Current mode driver (CML)

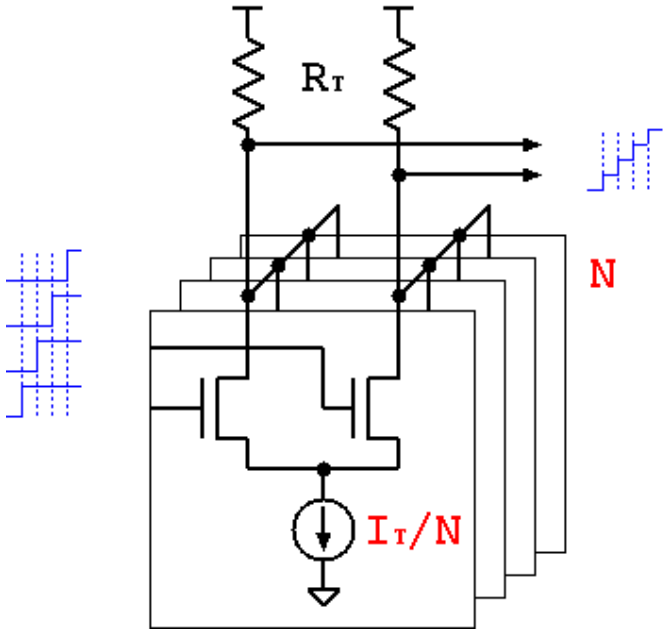

- Tail current sets the amplitude (can use segments to regulate amplitude while keeping constant current density)
- Impedance calibration done with prog. resistors (switches on top)

- Voltage mode driver (SST)

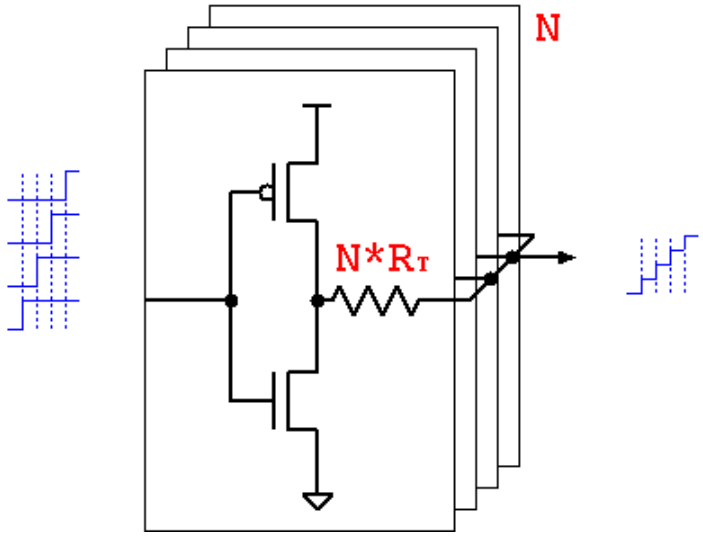

- Amplitude is defined by a ratio of unit elements (up's and down's) → intrinsic good matching. Function of the supply voltage
- Impedance calibration performed with programmable switches

ESD/LU protection

- CML
  - all driver diffusions are directly connected to the output PAD
    - only wire series resistance
    - current limiting is in the device
      - drain extensions
      - silicide mask
      - longer transistors
    - all diffusions are “HOT”
  - Calibration switches have series R
    - LU ring needed depending on R

- SST
  - series resistors limit ESD discharge currents
  - depending in the value of the R may still require LU rings

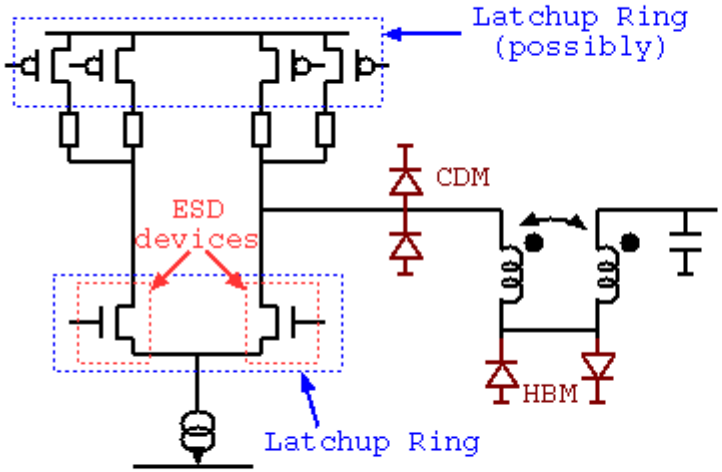

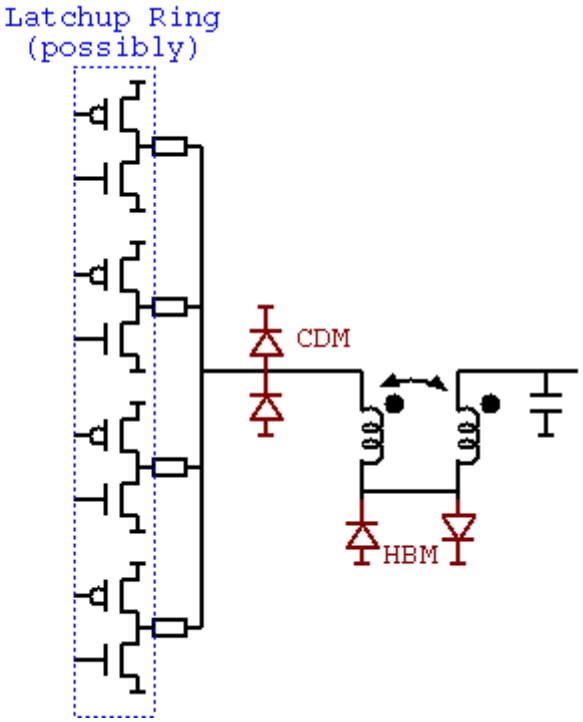

Feed Forward Equalizer (FFE)

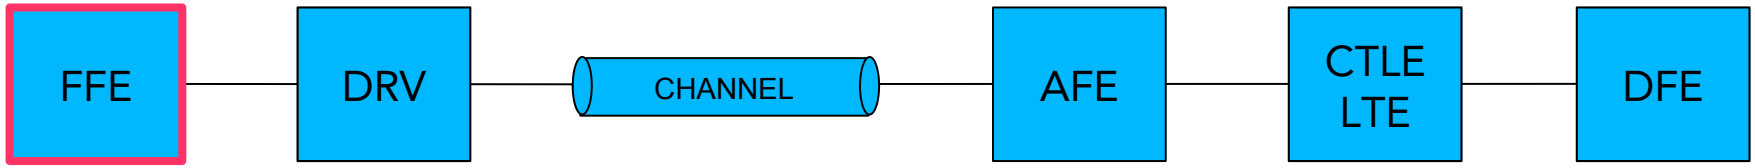

Adding FFE to the Line Driver

- FFE
  - summation of $C_i b_i$ terms
  - envelope is well defined
- Use superposition:
  - split the driver in multiple slices
  - assign slices to coefficients
  - shift the bits from one to the next
- CML
  - easy to implement the summation
- SST
  - naturally segmented
  - the work happens in the digital domain
  - tricky to precalculate the configurations

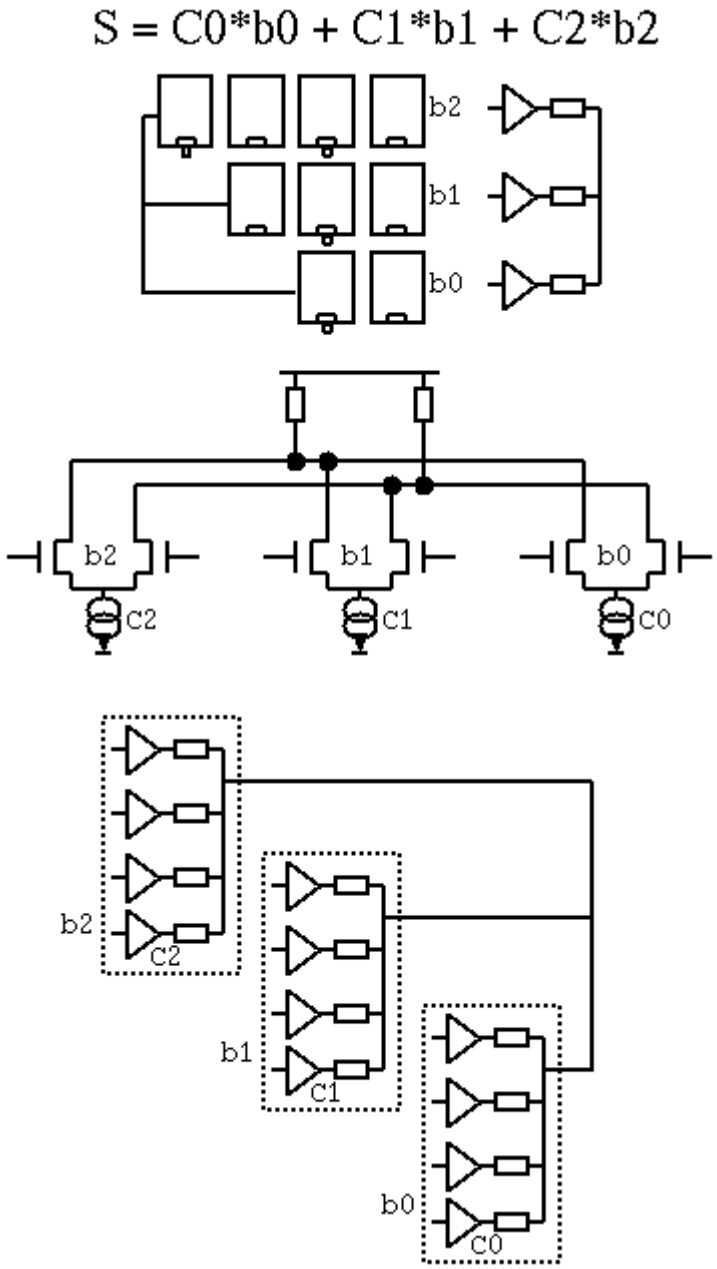

RX → Analog Front End (AFE)

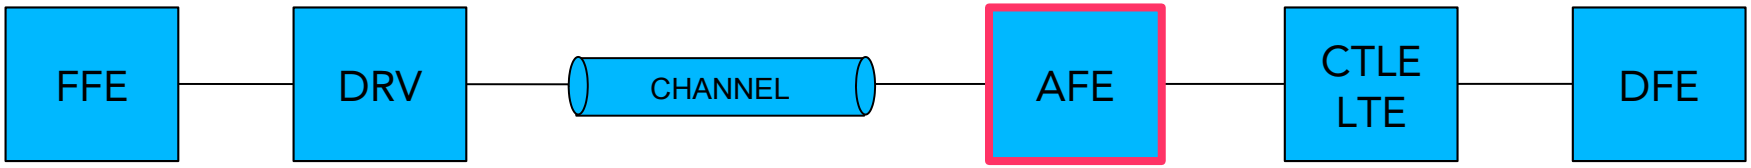

- Termination, ESD
- AC coupling
- AGC (attenuator, VGA)

RX → Analog Front End (AFE)

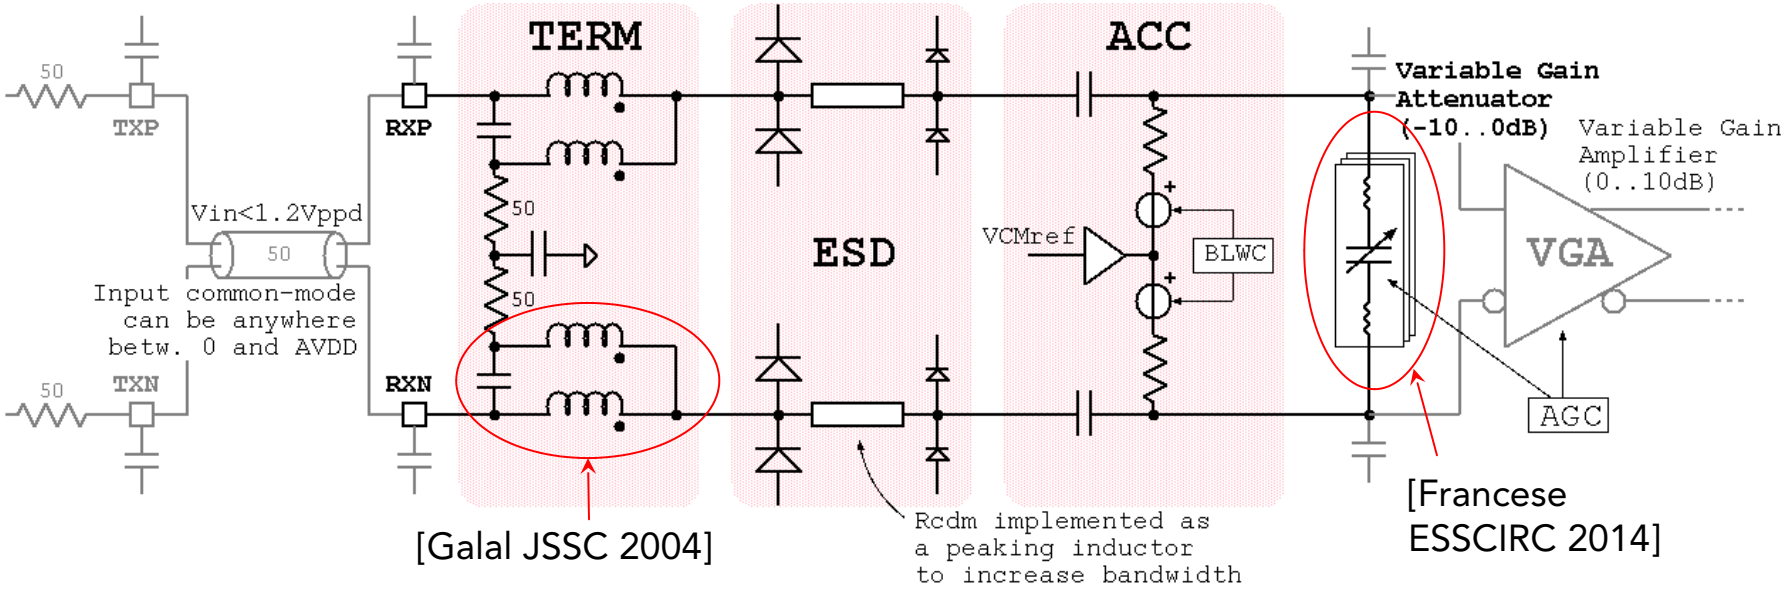

- Termination & ESD protection
  - Broad-band (from DC to baud-half and more)
  - Return Loss (common-mode & differential)
  - Two-stage ESD protection (HBM, CDM)
- On-chip AC coupling
  - AC-coupling: needed to decouple input common-mode of the VGA
  - On-chip: reduces system cost & area (can have >100 lanes / chip!)
  - On-chip: eliminates Z-discontinuity related to discrete cap on PCB
  - But: requires Baseline Wander Compensation (BLWC)
- Passive attenuation (supports AGC)
  - Capacitive divider (...careful not to destroy RL!)

RX → AFE → AC coupling (VCM & BLWC)

- Baseline Wander (BLW)
  - RC size limited by technology/area
  - Signal can be extremely broadband (plus: have to support legacy systems)
  - Spec: max(BLW) vs. transmission rate
  - Test: use OIF short-stress patterns

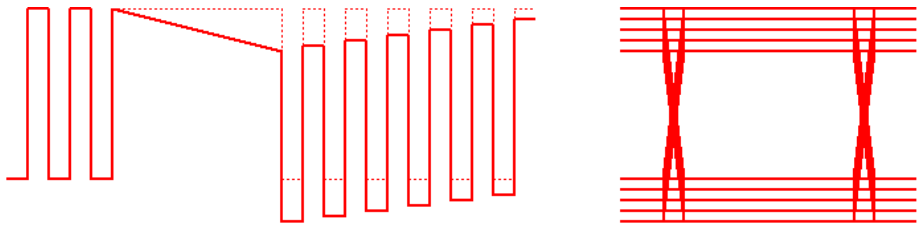

- BLW "compensation"
  - Can compensate BLW or avoid it altogether (IBM solution)
  - Level-shift signal directly from the pads (in a controlled loop)
  - Avoid discharge path by using floating voltage generators

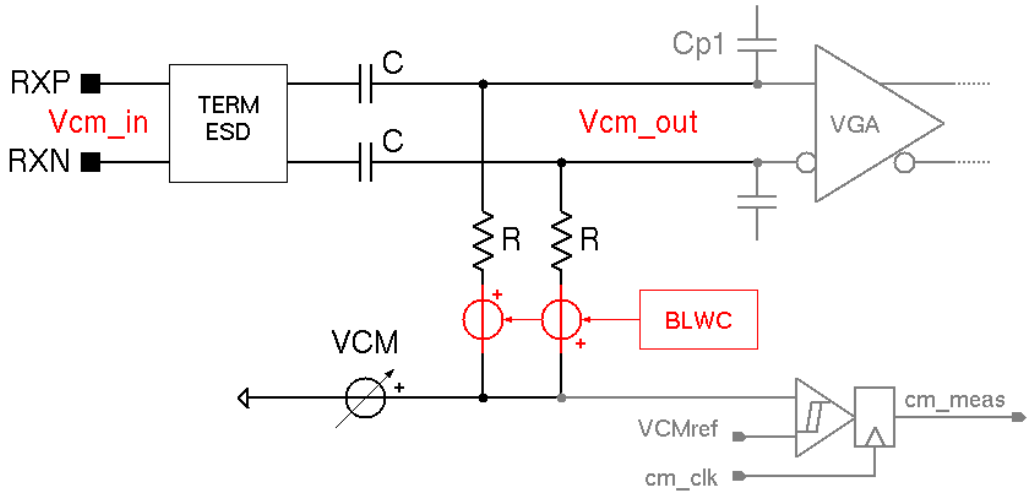

RX → AFE → AC coupling (VCM & BLWC)

- Baseline Wander (BLW)
  - RC size limited by technology/area
  - Signal can be extremely broadband (plus: have to support legacy systems)
  - Spec: max(BLW) vs. transmission rate
  - Test: use OIF short-stress patterns

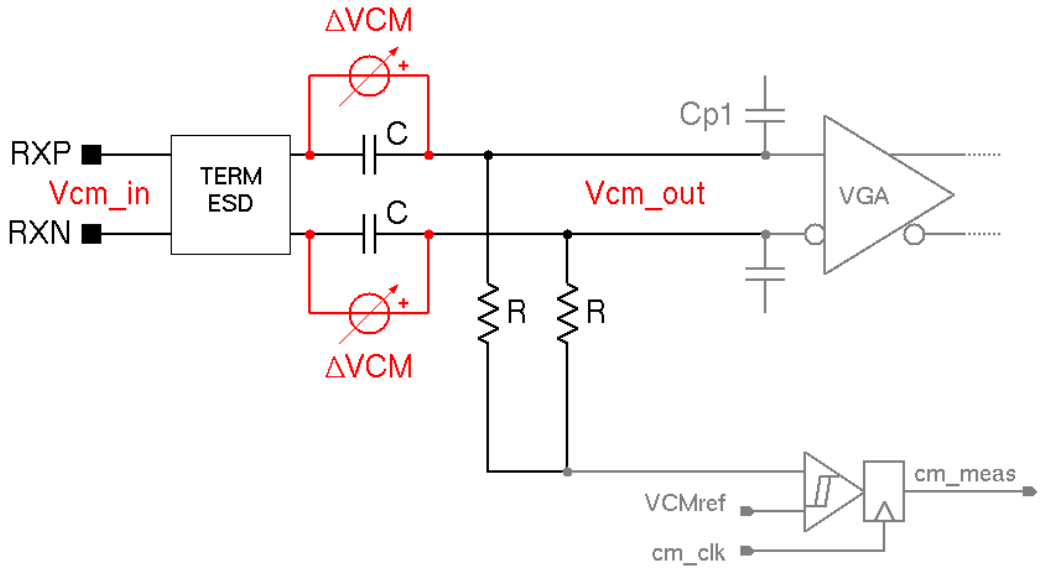

- BLW "compensation"
  - Can compensate BLW or avoid it altogether (IBM solution)
  - Level-shift signal directly from the pads (in a controlled loop)
  - Avoid discharge path by using floating voltage generators

[Gangasani JSSC 11/2014]

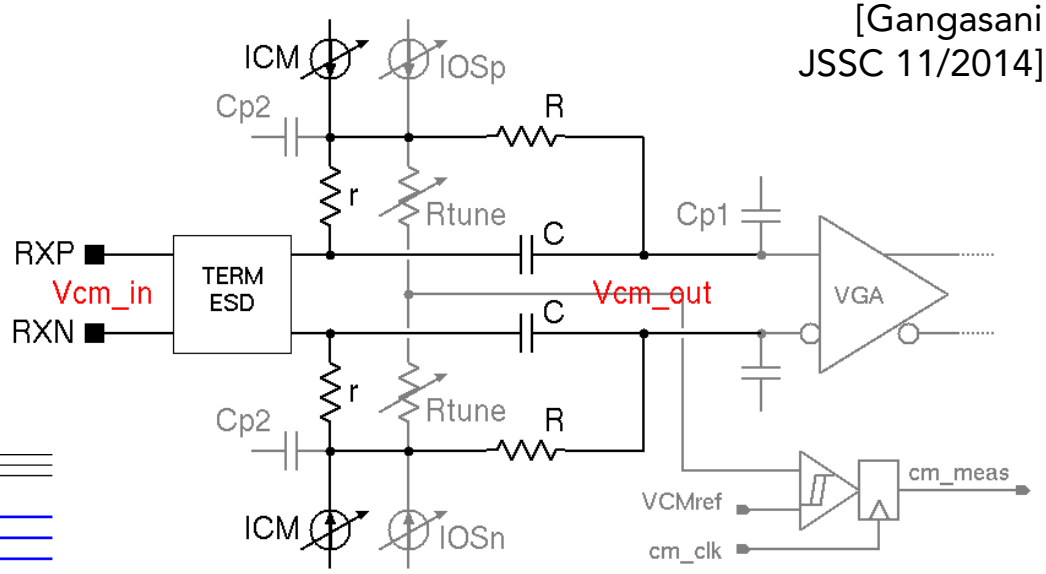

Offset correction

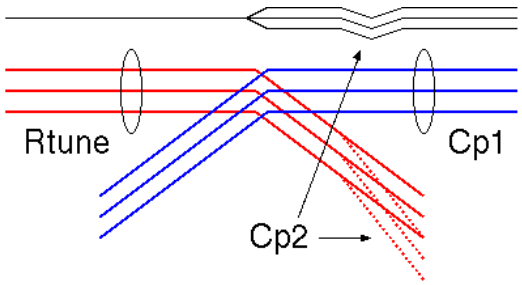

RX → AFE → Variable Gain Amplifier (VGA)

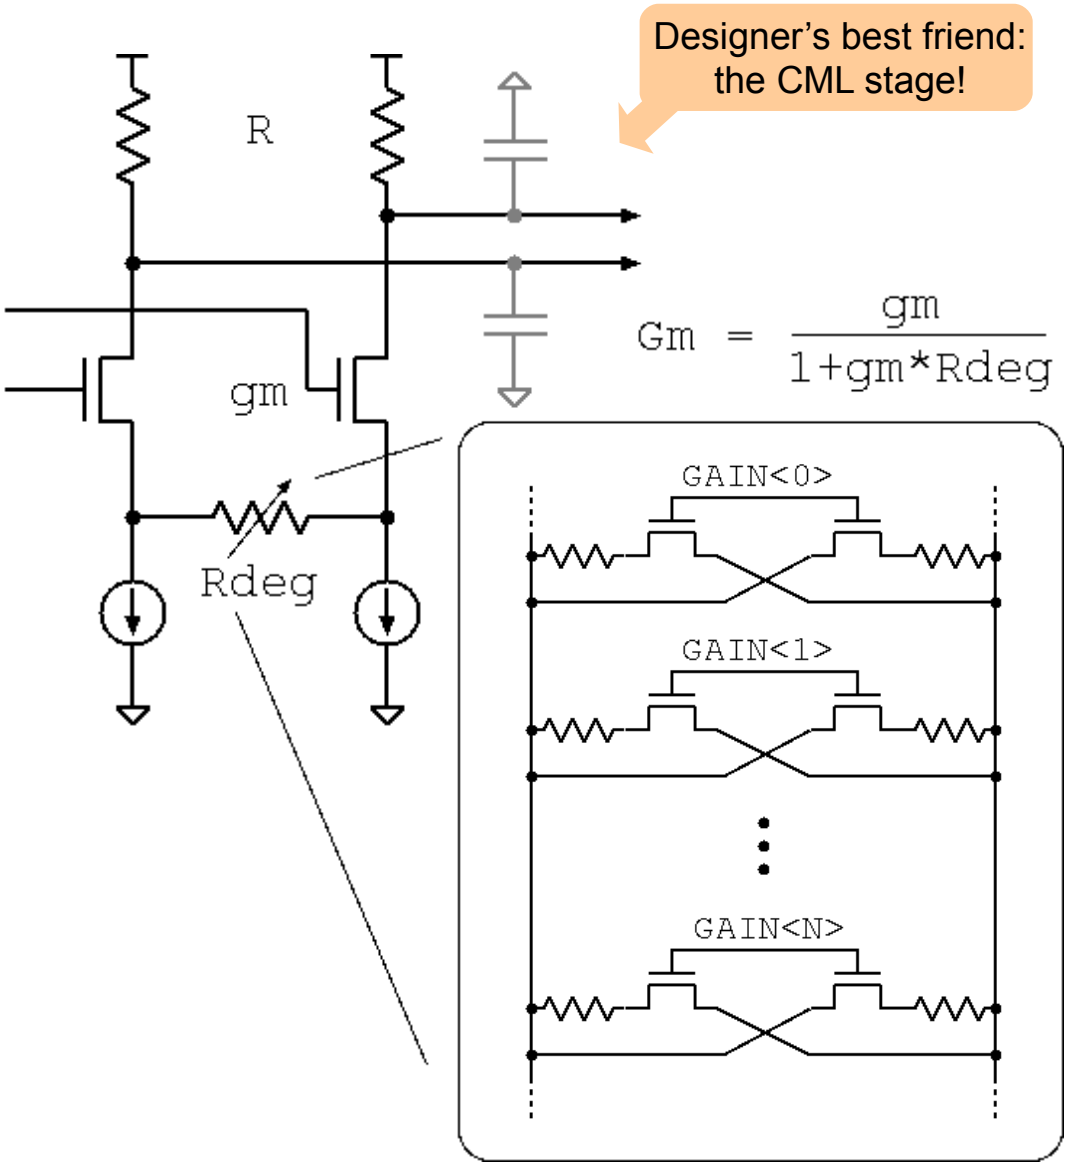

- Function: Variable gain
  - Supports AGC
  - Typical range: 10dB (in steps of 1dB)
  - Difficult to do more due to: UWP, source impedance
  - Typical gain: as much as you can get
  - Difficult to do more than ~6dB (with single CML stage across PVT corners)
- Band broadening techniques
  - Shunt/series peaking, T-coils, etc.
  - Active inductors
  - Anti-miller caps
  - Degeneration by-pass (not for VGA!)
  - Negative caps
  - ...be careful with ringing behaviour!
- Unwanted Peaking (UWP)
  - Unavoidable as much as parasitic cap!
  - Problem if system assumes orthogonal controls for gain/peaking
  - May try to reduce it...

RX → AFE → Variable Gain Amplifier (VGA)

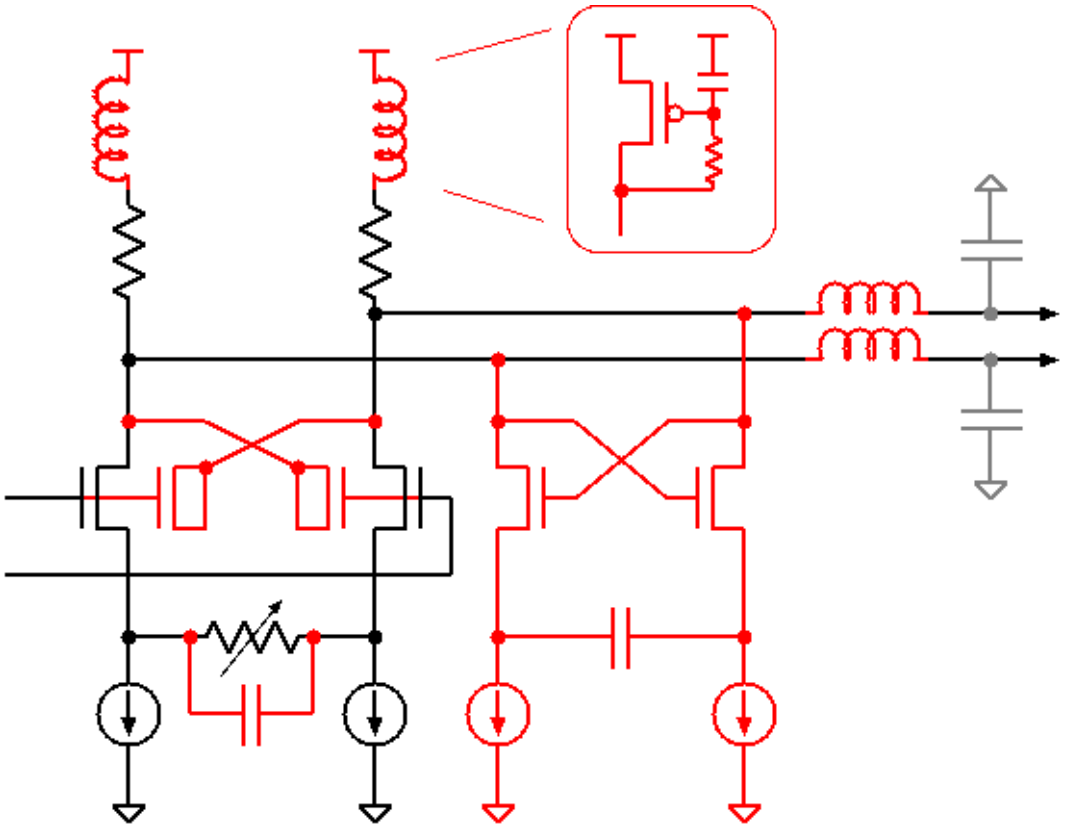

- Function: Variable gain
  - Supports AGC
  - Typical range: 10dB (in steps of 1dB)
  - Difficult to do more due to UWP, source impedance
  - Typical gain: as much as you can get
  - Difficult to do more than ~6dB (with single CML stage across PVT corners)
- Band broadening techniques
  - Shunt/series peaking, T-coils, etc.
  - Active inductors
  - Anti-miller caps
  - Degeneration by-pass (not for VGA!)
  - Negative caps
  - ...be careful with ringing behaviour!
- Unwanted Peaking (UWP)
  - Unavoidable as much as parasitic cap!
  - Problem if system assumes orthogonal controls for gain/peaking
  - May try to reduce it...

RX → AFE → Variable Gain Amplifier (VGA)

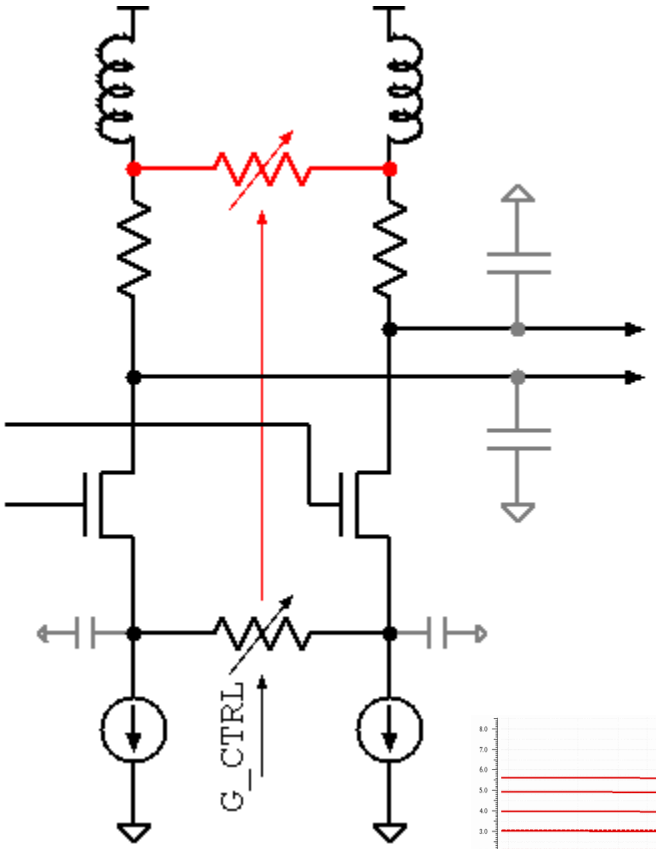

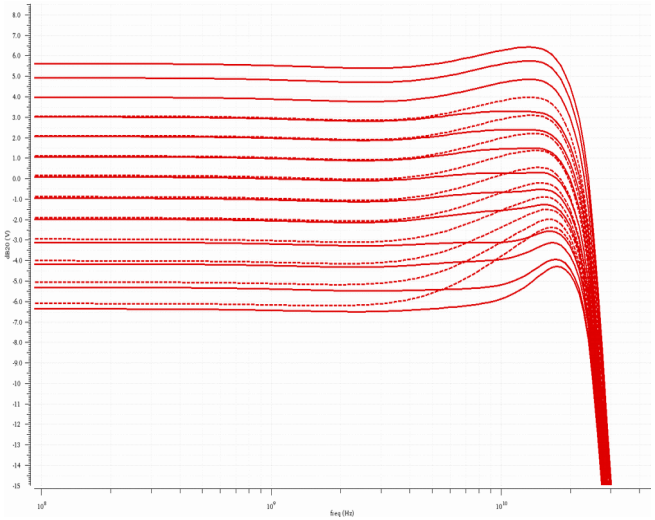

- Function: Variable gain
  - Supports AGC
  - Typical range: 10dB (in steps of 1dB)
  - Difficult to do more due to UWP, source impedance
  - Typical gain: as much as you can get
  - Difficult to do more than ~6dB (with single CML stage across PVT corners)
- Band broadening techniques
  - Shunt/series peaking, T-coils, etc.
  - Active inductors
  - Anti-miller caps
  - Degeneration by-pass (not for VGA!)
  - Negative caps
  - ...be careful with ringing behaviour!
- Unwanted Peaking (UWP)
  - Unavoidable as much as parasitic cap!
  - Problem if system assumes orthogonal controls for gain/peaking
  - May try to reduce it...

RX → Cont. Time Linear Equalizer (CTLE/LTE)

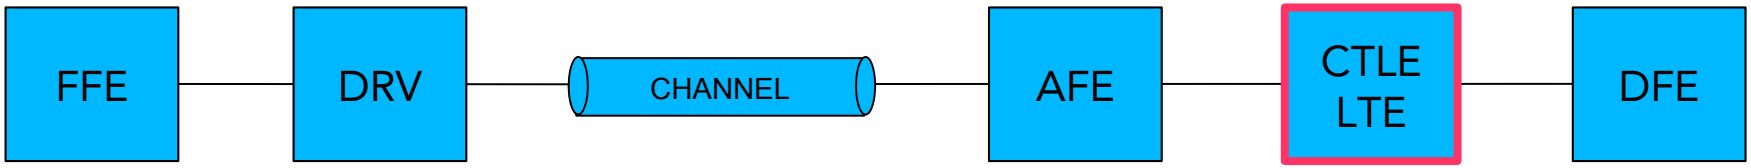

- Passive/Active
- Voltage/Current mode
- Active feedback
- Split path

RX → Cont. Time Linear Equalizer (CTLE/LTE)

	Active		
Passive	Gm-degeneration	Current-mode	Active-feedback
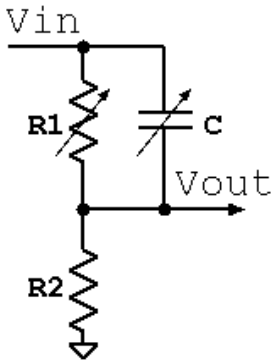	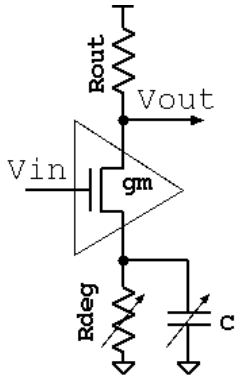	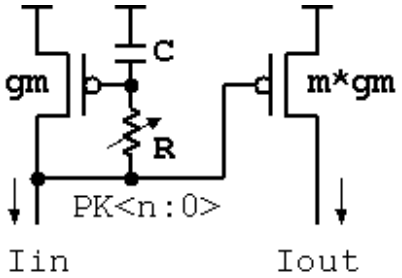	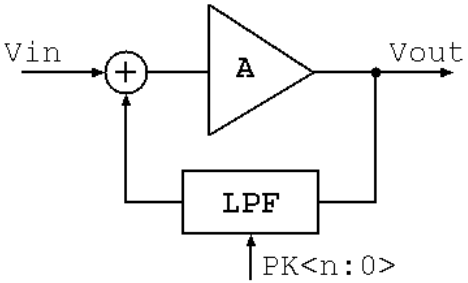
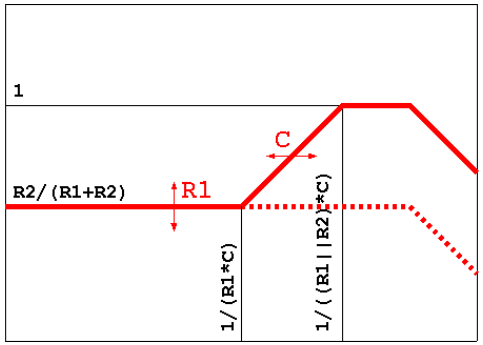	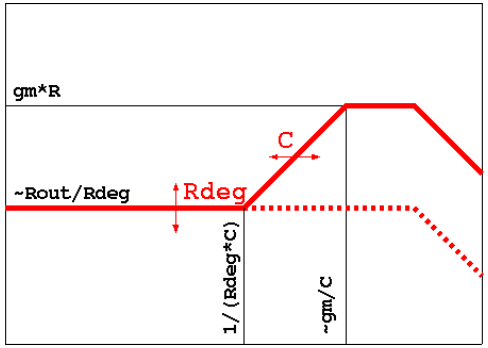	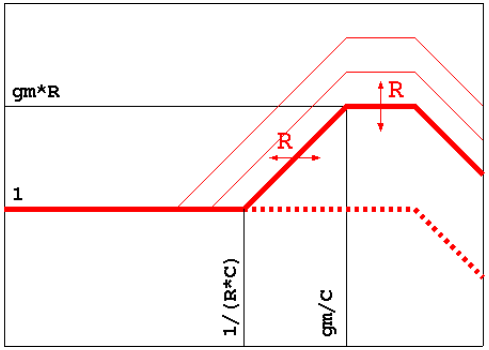	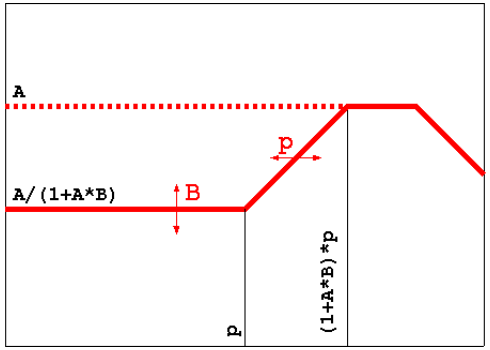

Most important requiremets:

- High-frequency boost
- Minimal low-frequency loss
- Good linearity
- Well-behaved phase response (limited usage of resonant peaking)
- Tunability of the boost
- Low input capacitance (ultimately determining S11)

RX → Cont. Time Linear Equalizer (CTLE/LTE)

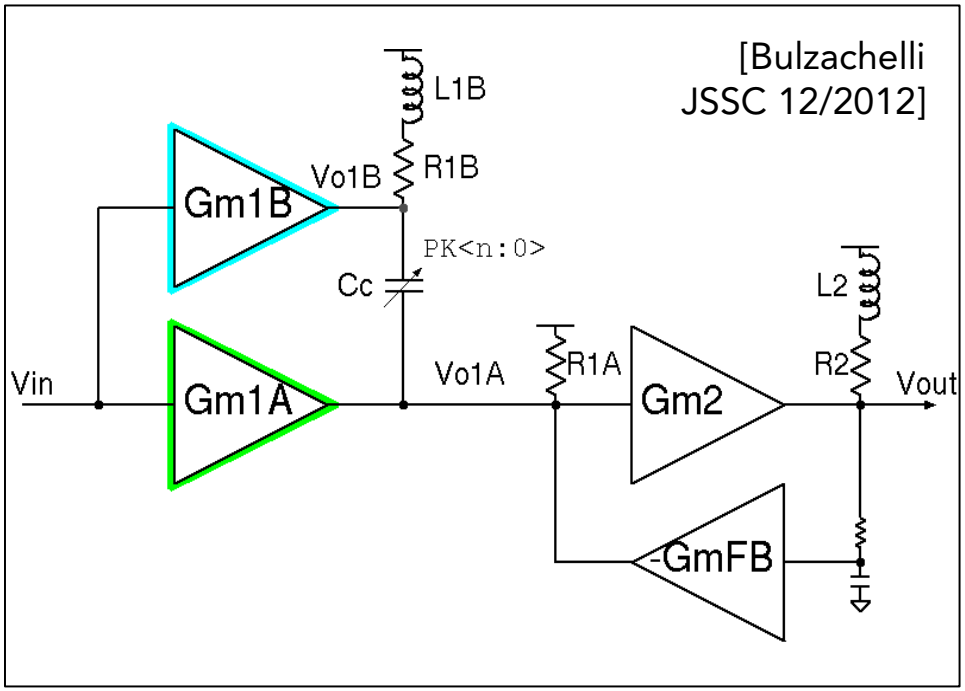

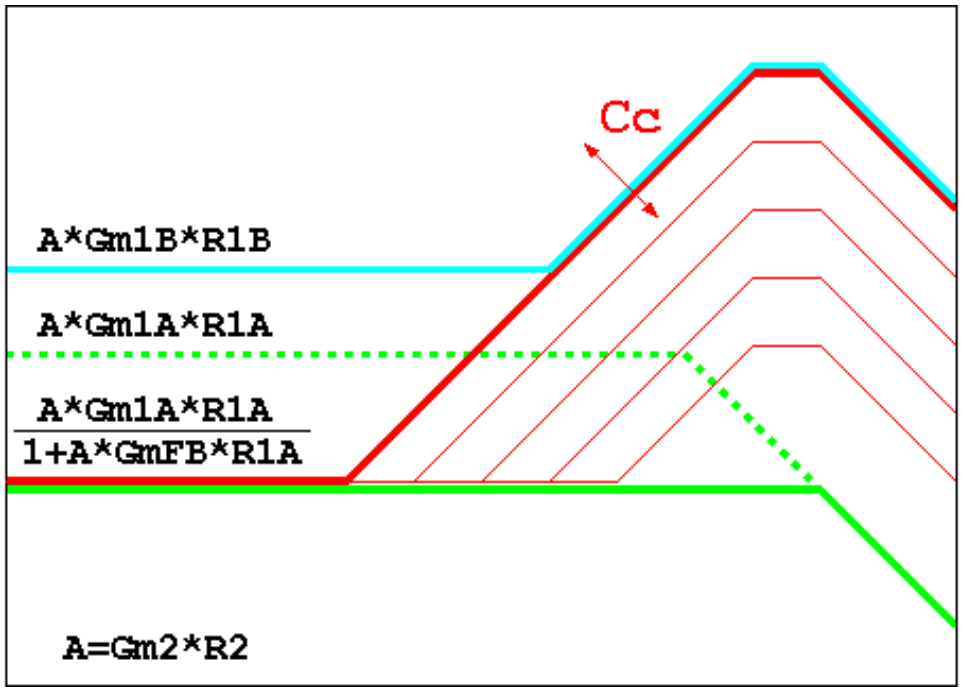

- Peaking amplifier (= active, linear)
  - Split path + AC-coupling
  - Inductive peaking
  - Active feedback

$$\frac{Vo2}{Vin}=\left(\frac{GA\cdot G2}{1+GFB\cdot G2}\right)\cdot\frac{1+s\cdot Cc\cdot R1B\cdot\left(1+\frac{Gm1B}{Gm1A}\right)}{1+s\cdot Cc\cdot R1B\cdot\left(1+\frac{R1A/R1B}{1+GFB\cdot G2}\right)}$$

$$\delta P=\frac{Pmax}{Pmin}=\frac{G\infty(Cc=max)}{G\infty(Cc=0)}=\frac{\left(1+\frac{Gm1B}{Gm1A}\right)}{\left(1+\frac{R1A/R1B}{1+GFB\cdot G2}\right)}$$

RX → Cont. Time Linear Equalizer (CTLE/LTE)

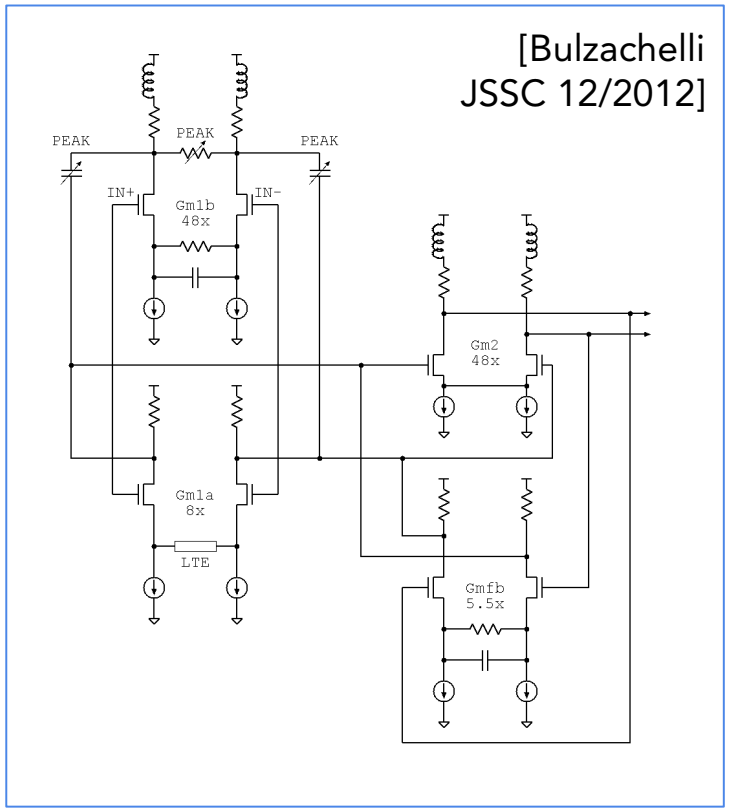

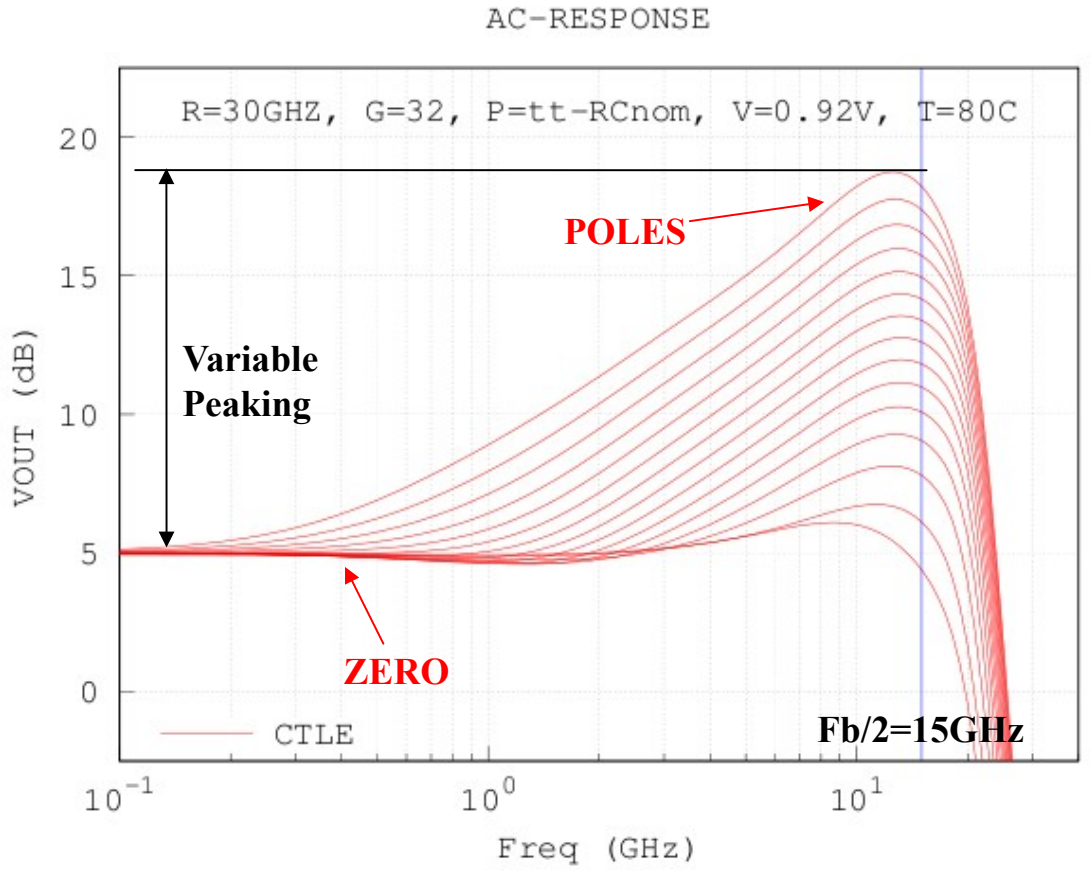

RX → Cont. Time Linear Equalizer (CTLE/LTE)

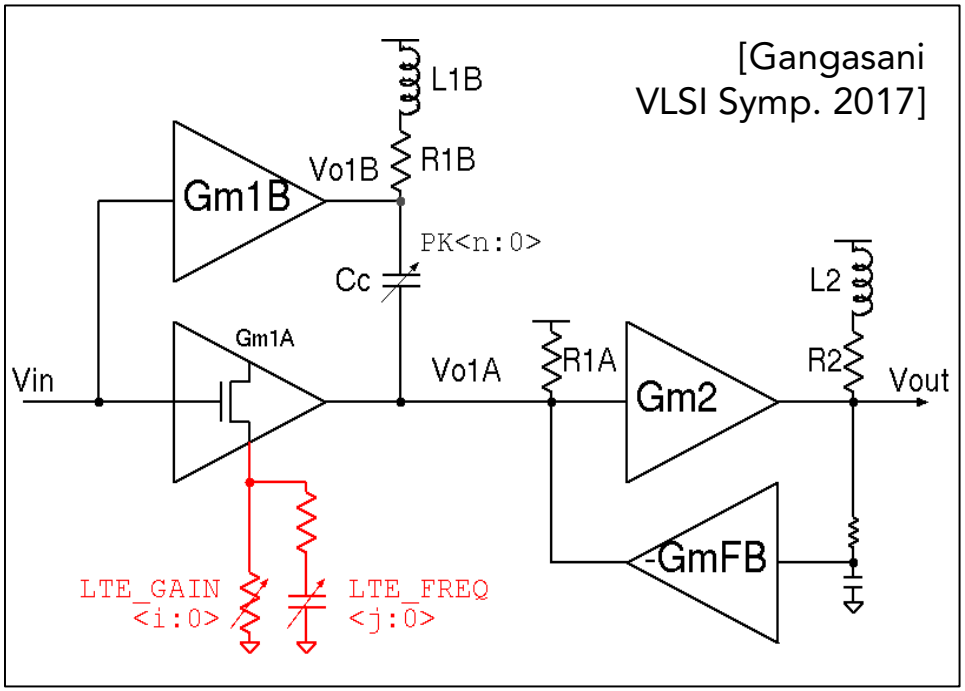

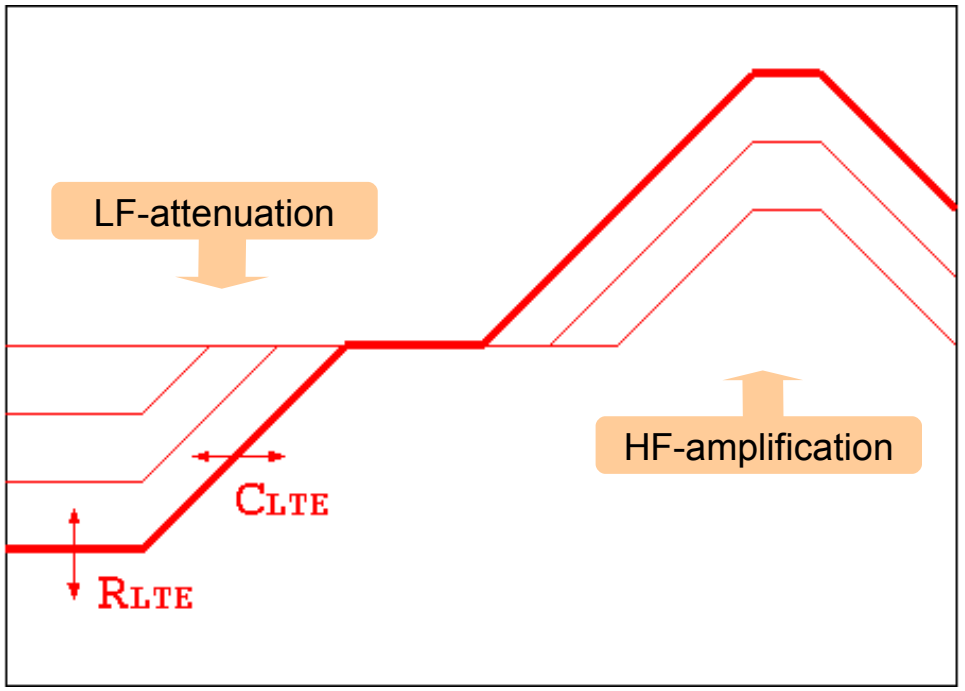

- LTE effectively increases the equalization capability by combining LF-dump & HF-boost (plus: approximate slopes different than 20dB/decade)
- CTLE/LTE: low-frequency shaping and large BW → often require bulky passive devices (RLC), tends to become a luxus in modern technologies optimized for digital

RX → Cont. Time Linear Equalizer (CTLE/LTE)

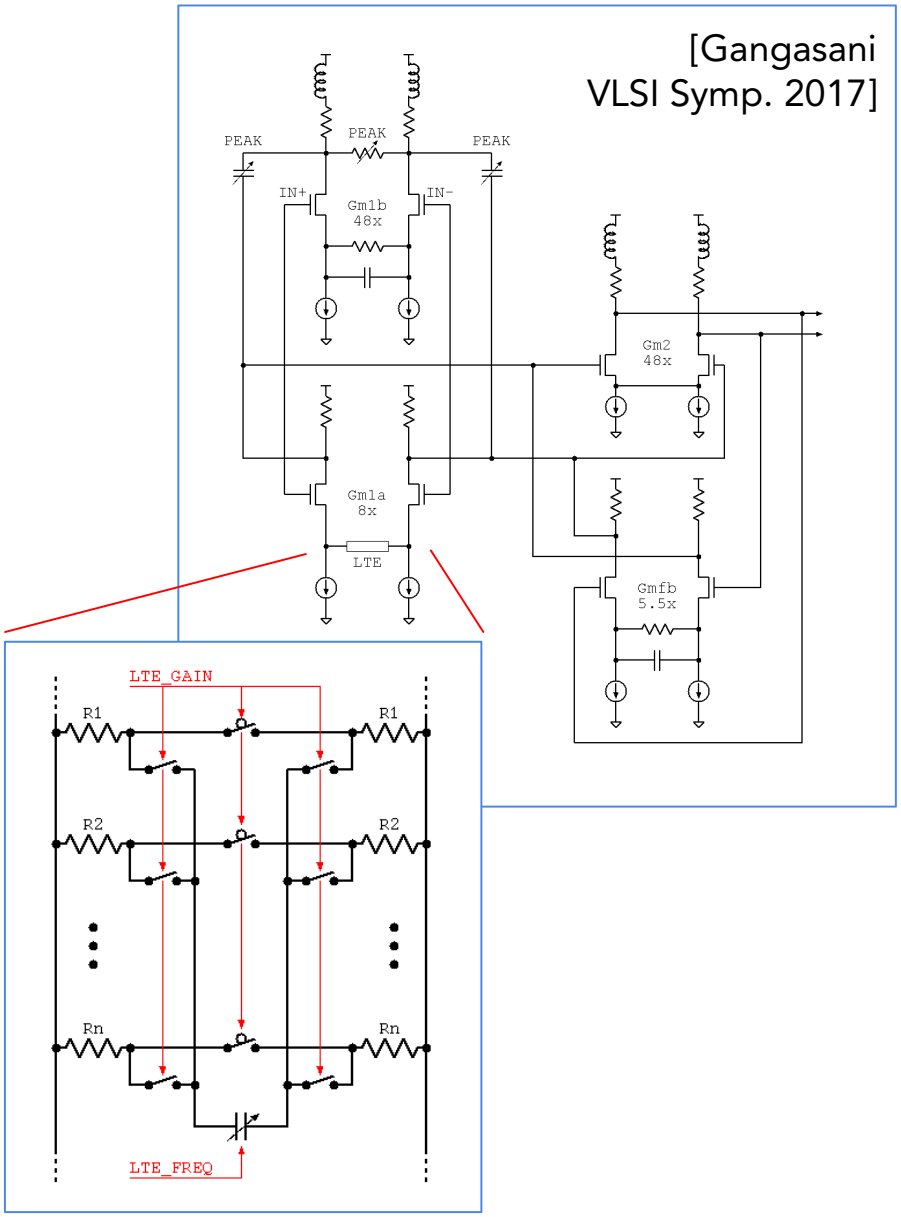

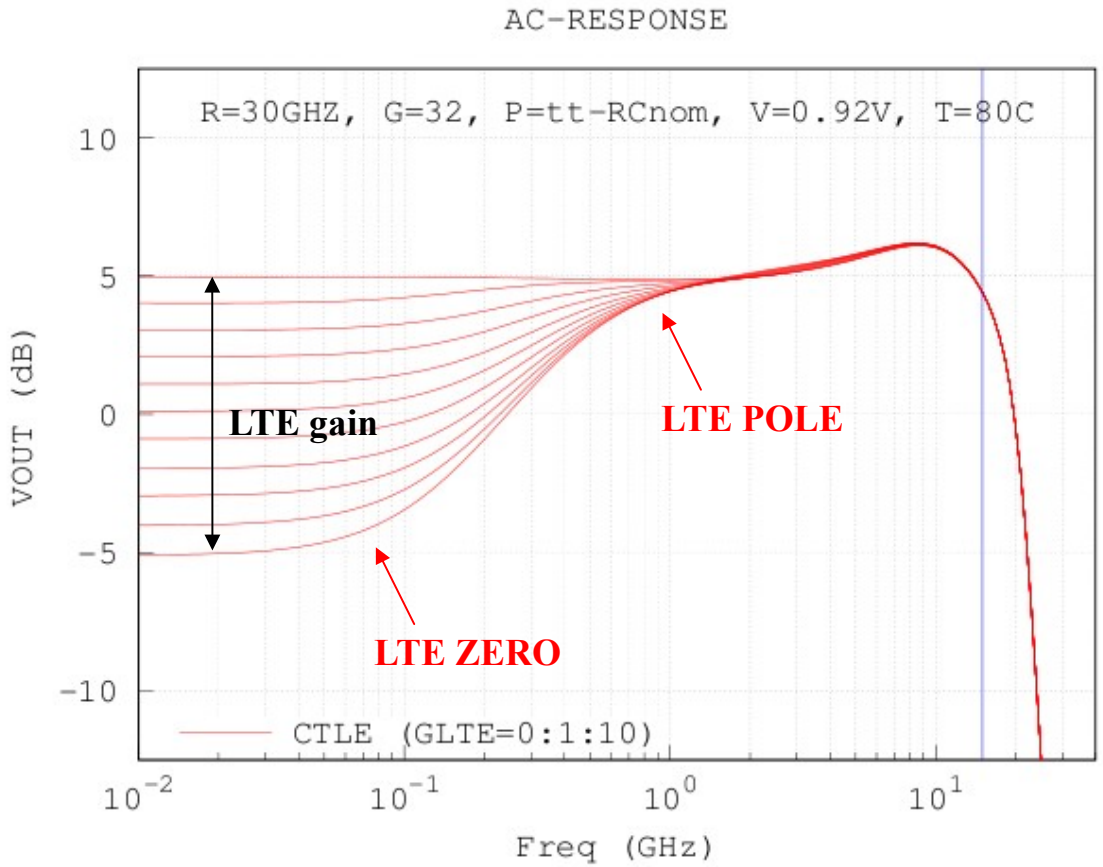

RX → Decision Feedback Equalizer (DFE)

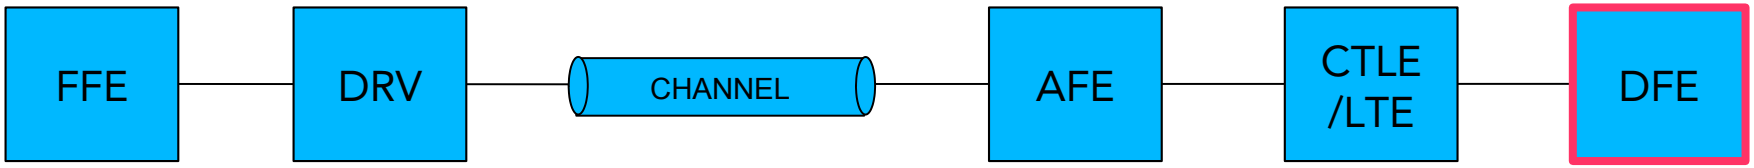

- Speculation
- Sub-rates
- Summer
- Slicer

RX → Decision Feedback Equalizer (DFE)

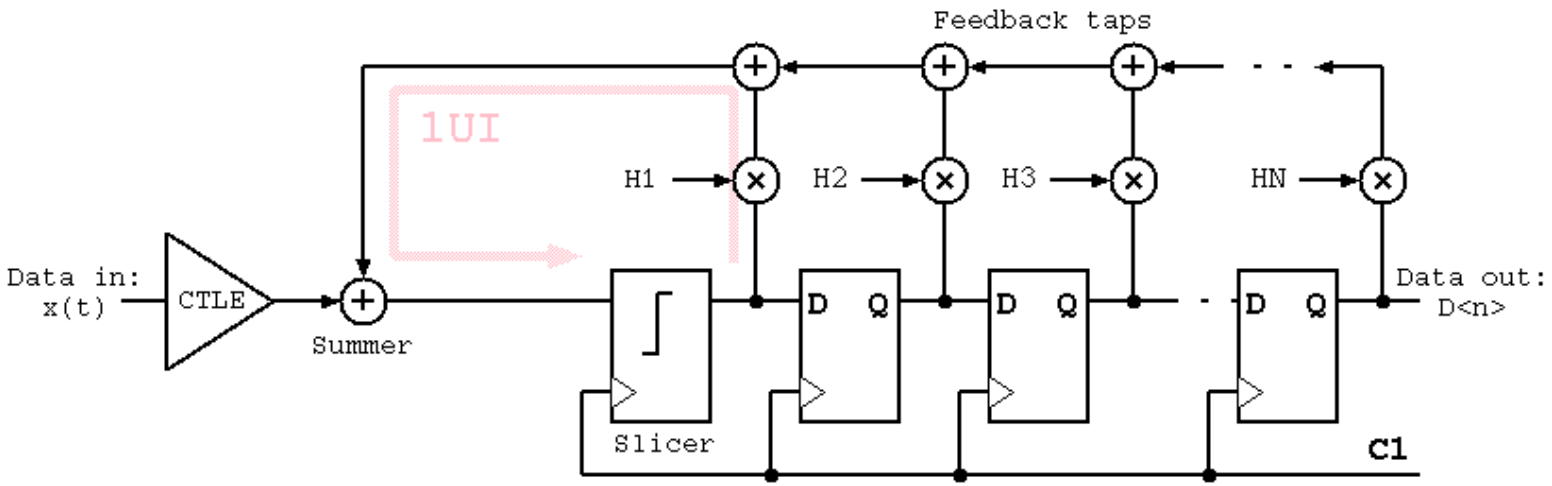

- DFE is powerful but difficult at those speeds! (it's a **feedback loop**)
  - 1-UI for decision, feedback and summation: 31.25ps if running @ 32Gbps
  - Typical slicer decision time for 10mV amplitude is ~30ps (in 14nm CMOS)
  - Typical inverter delay is ~9ps (in 14nm CMOS)
  - Summation settling time with 15GHz BW is ~22ps

However you turn it, this ends up being much longer than 1 UI !

RX → DFE → Unrolling

- We could borrow 1 bit time:
  - Take <u>both</u> decisions and choose when the next bit comes
    - This is **1-level look-ahead** (also called "speculation" or "unrolling")
    - Buys us 1 UI time at the cost of added hardware
    - Sounds fishy? It does work!

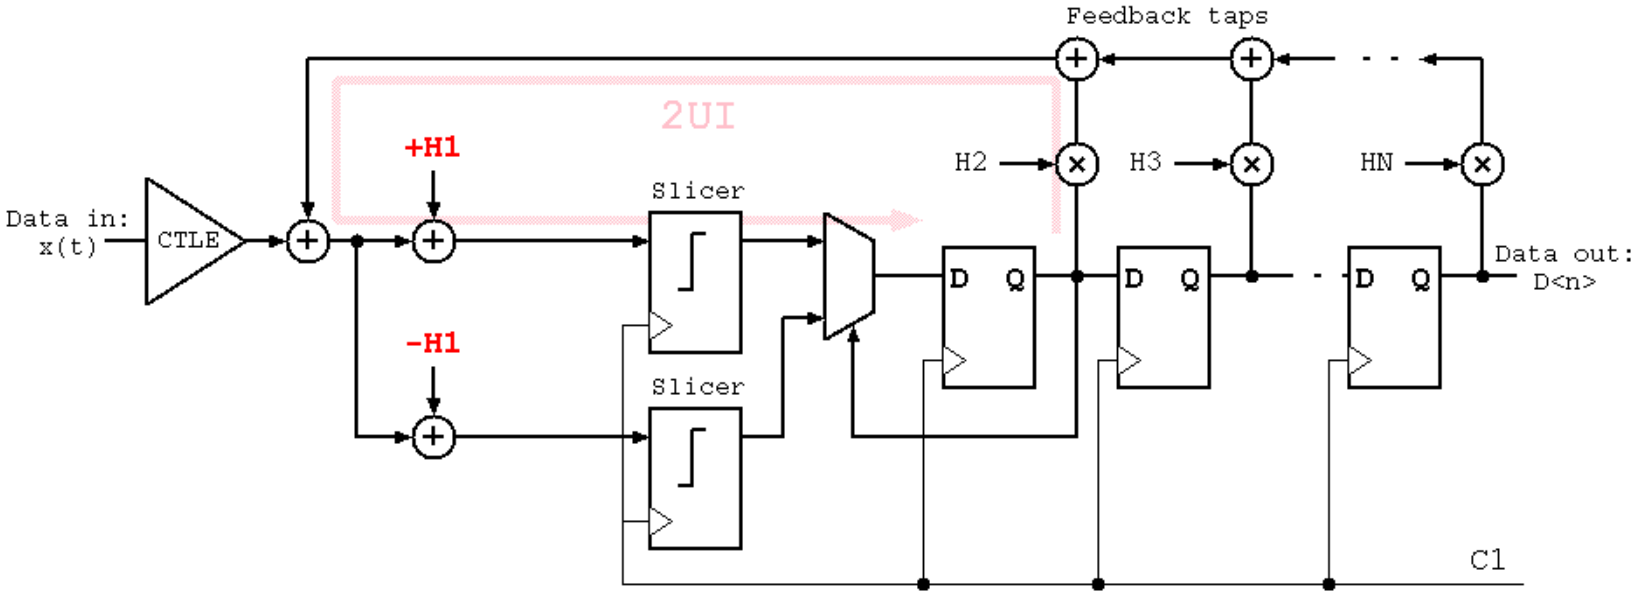

- You can take it further: 2-levels (4 decisions), 3-levels (8-decisions), …
- But: at each unrolling step you are doubling the number of slicers (4x for PAM4)
  - Additional load for the driving stage and for clock distribution
  - Are you also doubling the number of summers? Prefer to shift the slicing thresholds

RX → DFE → Sub-rate clocks

- Half-rate (C2)
  - Further relax the circuit requirements by halving the clock frequency and processing the data on 2 parallel paths (phase shifted by 180⁰)
  - Need to insure duty-cycle accuracy of the C2 clocks
  - Where does the power go? …

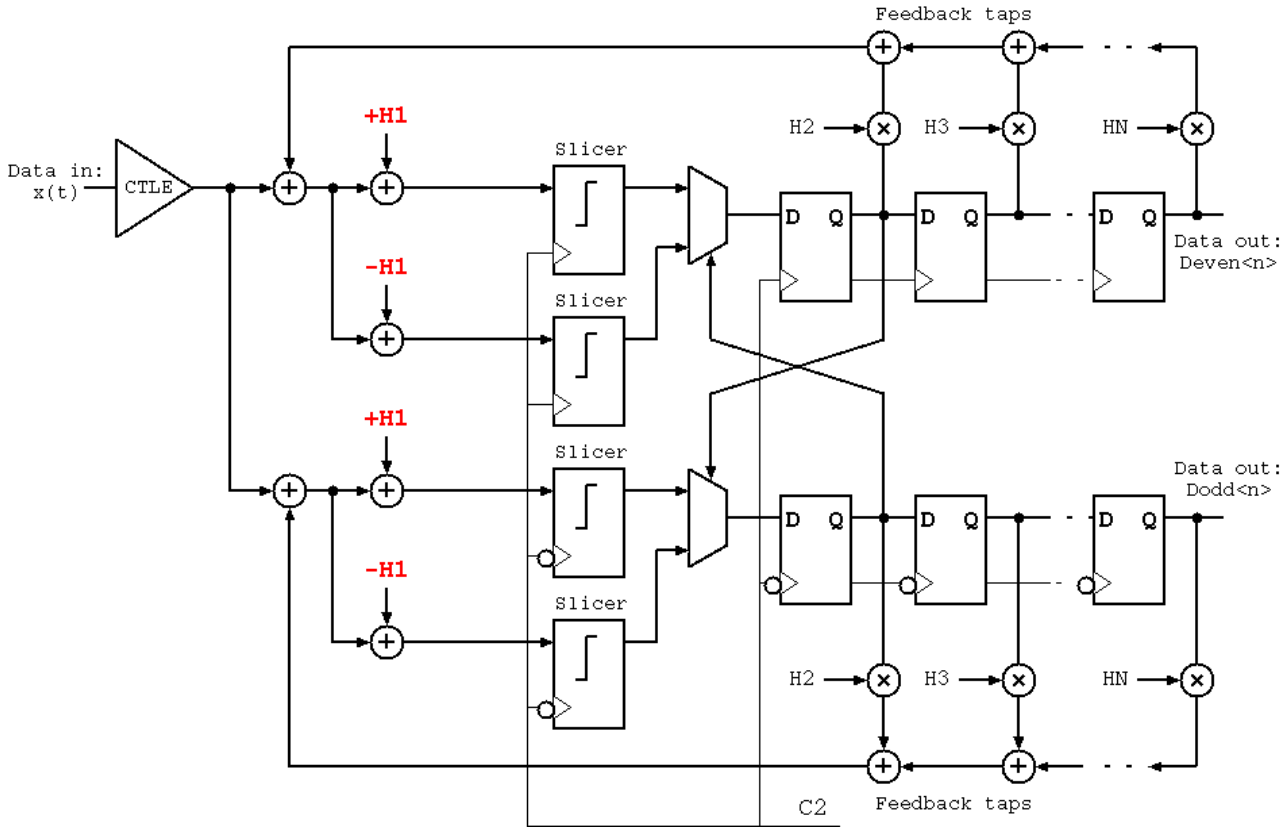

- Quarter-rate (C4)
  - Need to insure quadrature accuracy of C4 clocks, etc.

RX → DFE → Example

- Quarter-rate triple-speculation 15-taps DFE for 32Gbps NRZ

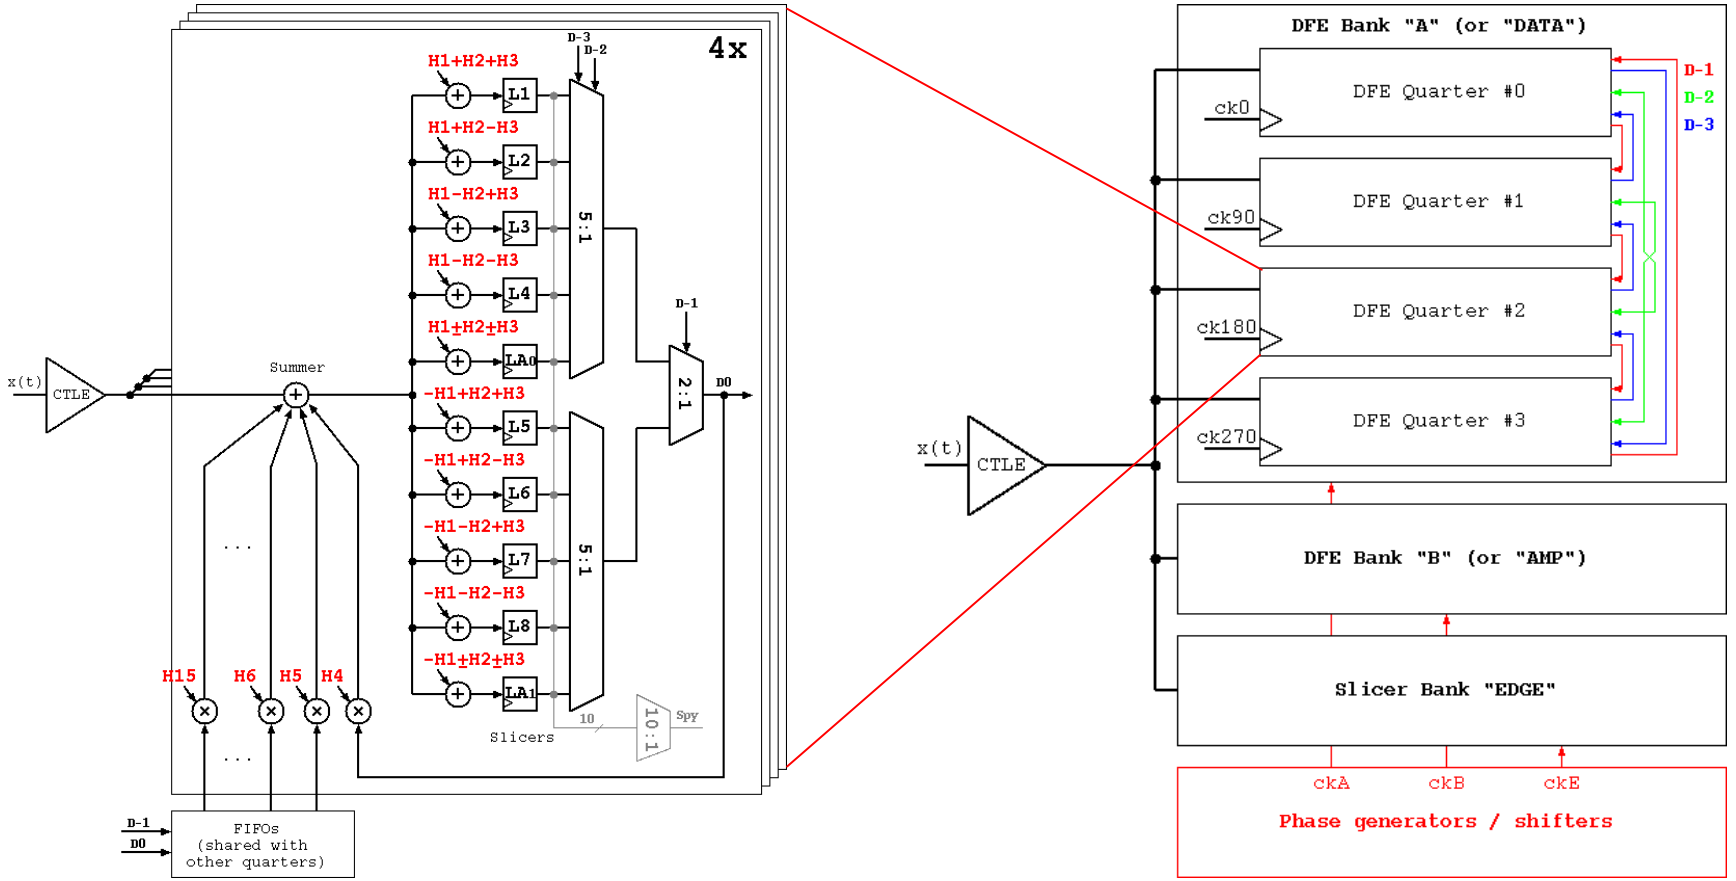

- Must provide measurement channels
  - Latch calibration (offset)
  - DFE taps adaptation
  - Phase detection
  - Test & monitoring

- Redundant banks and/or latches
  - A/B/EDGE
  - DATA/AMP/EDGE
  - Redundant latches (spy) that can be cycled

[Bulzacchelli VLSI Symp. 2017]

RX → DFE → Summer

- Resistively loaded CML summer

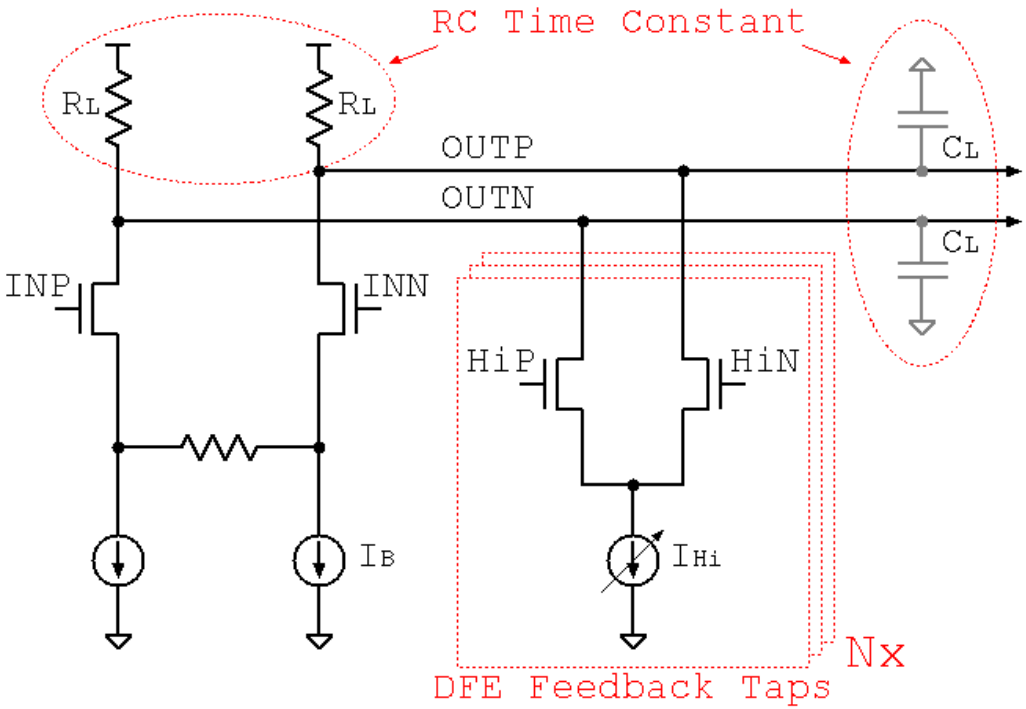

- Current summation implements: signal (linear) + feedback taps
- Tap amplitudes are tuned with the tail currents (IHi) through IDACs
- Required timing sets a limit to max. RL and min. IB (=> min. power)

RX → DFE → Summer

- Current integrating summer

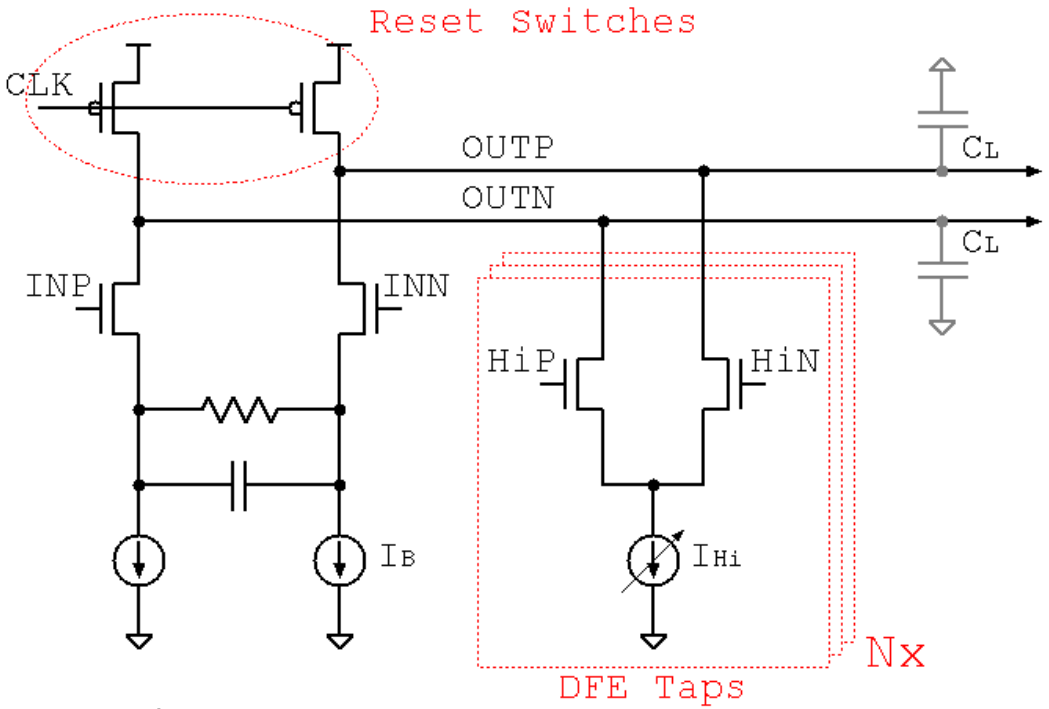

[Park ISSCC 2007]

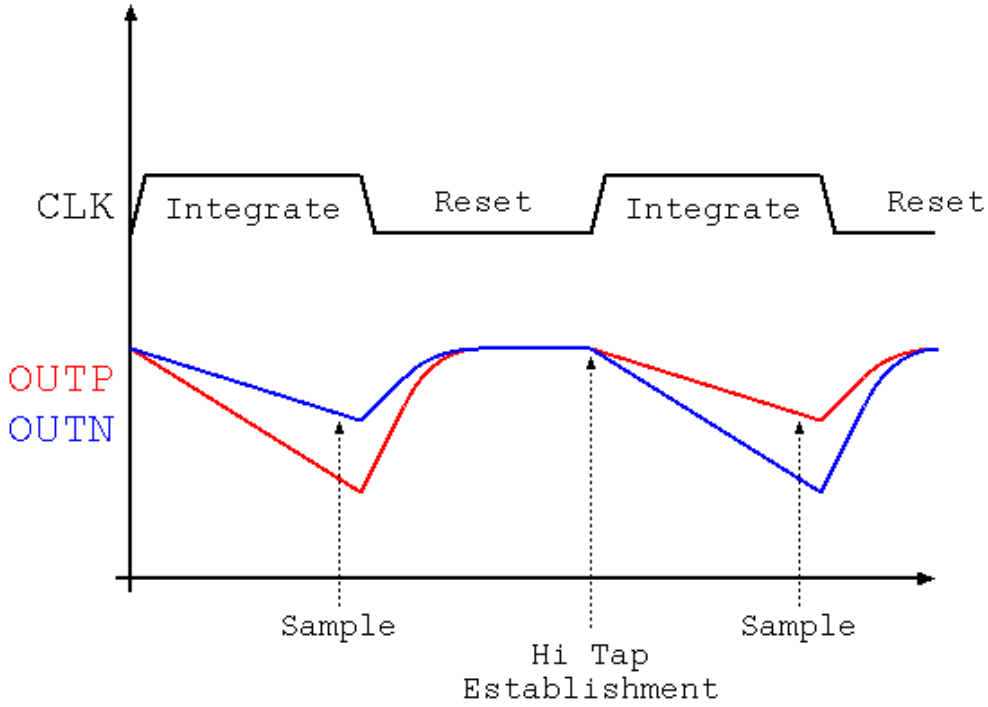

- By replacing resistors with PMOS switches, the IB current can be scaled
- Highest efficiency if only parasitic CL is loading the output, however also highest variability => requires calibration
- Feedback taps (HiP/N) should be established before beginning of integration
- Peaked integrator vs. Sampled integrator? (see next slide)

RX → DFE → Summer

- Peaked integrator vs. Sampled integrator?

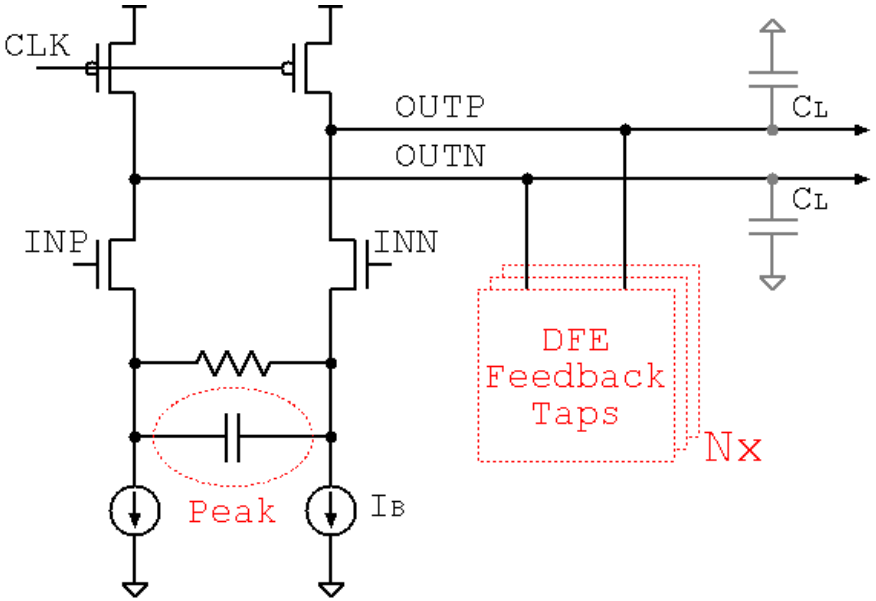

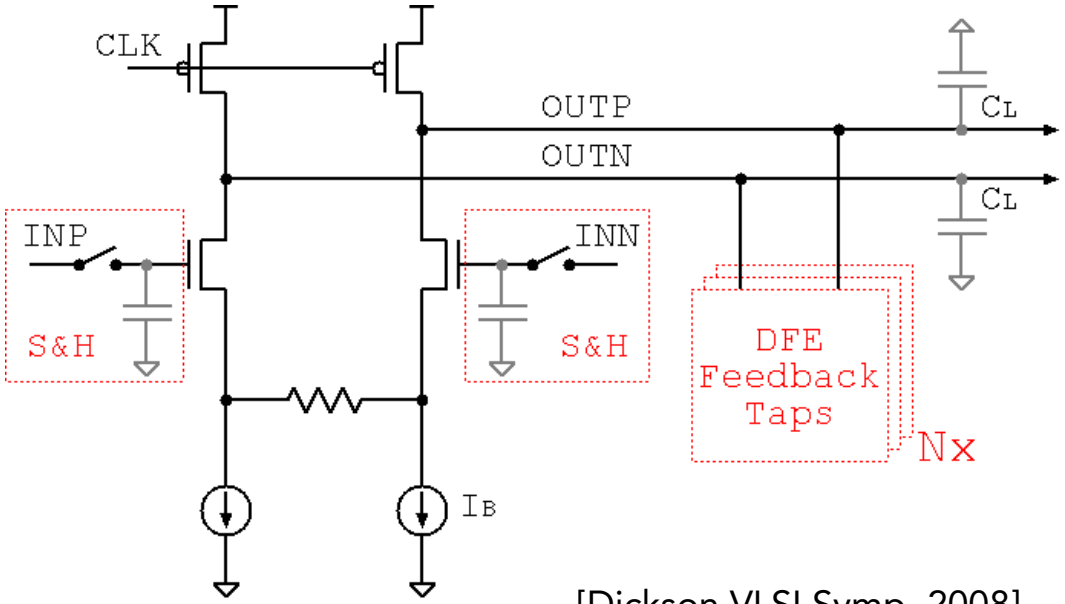

[Dickson VLSI Symp. 2008]

- The integration is an average over the clock cycle ("sinc" loss vs. frequency)
  - Zero transmission (infinite loss) at clock frequency (Fbaud), 3.9dB loss at Fbaud/2, etc.
- Peaked integrator
  - Compensates sinc loss with capacitive peaking
  - Integration time is 1UI (cannot increase gain by increasing integration time)
- Sampled integrator
  - Integrates a sampled (=constant) version of the signal
  - Problems related to sampling noise (kT/C), kick-back, …
  - May allow longer integration times (with lower current) in sub-rate systems (e.g. C4)

RX → DFE → Summer

- Integrating summer with capacitive charge feedback

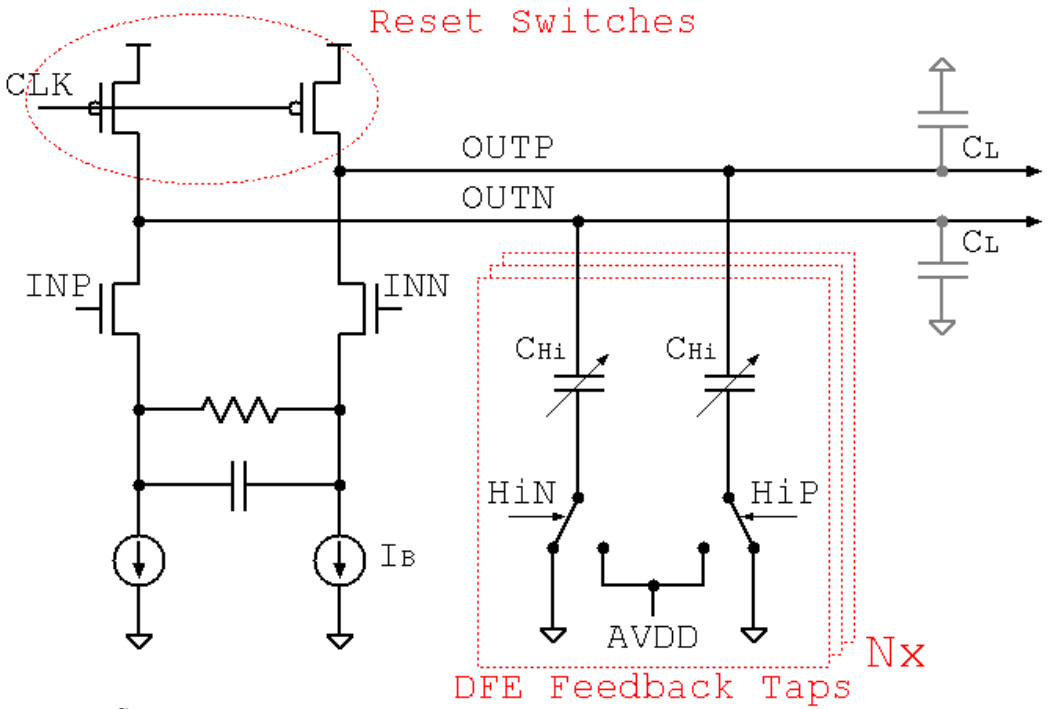

[Toifl VLSI Symp. 2011]

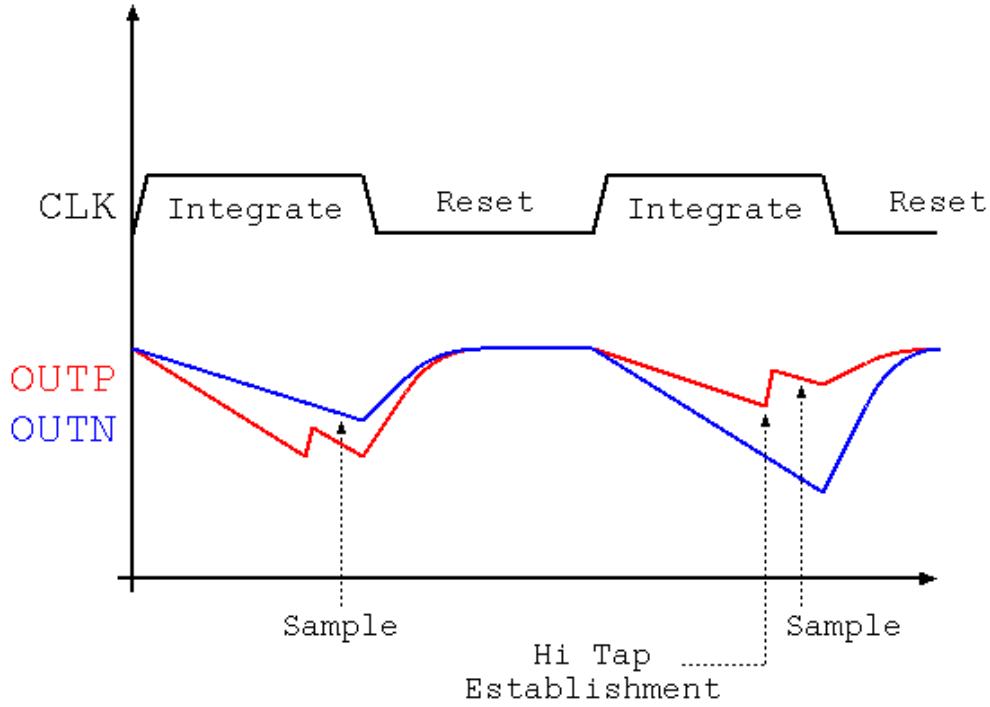

- Charge injection allows to relax the timing
- Feedback tap need not be ready before integration starts
- Tap amplitudes are tuned with the feedback caps (CHi) through CDACs
- Depending on technology: may end up being bulky/parasitic
  - Usually reserve it for the first non-speculated tap (most timing critical)

RX → DFE → Slicer

- Threshold detection (sensing) & Sampling
  - In practice these 2 functions are merged in the clocked-comparator (or "latch")

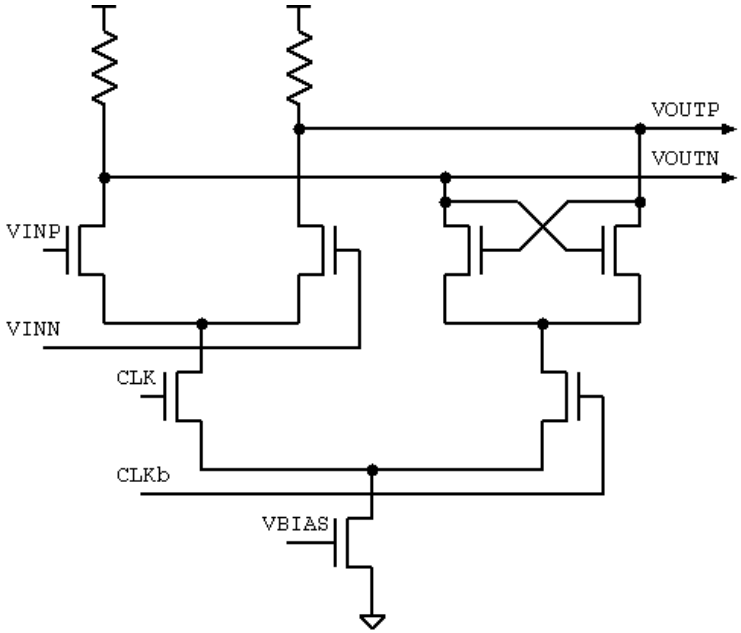

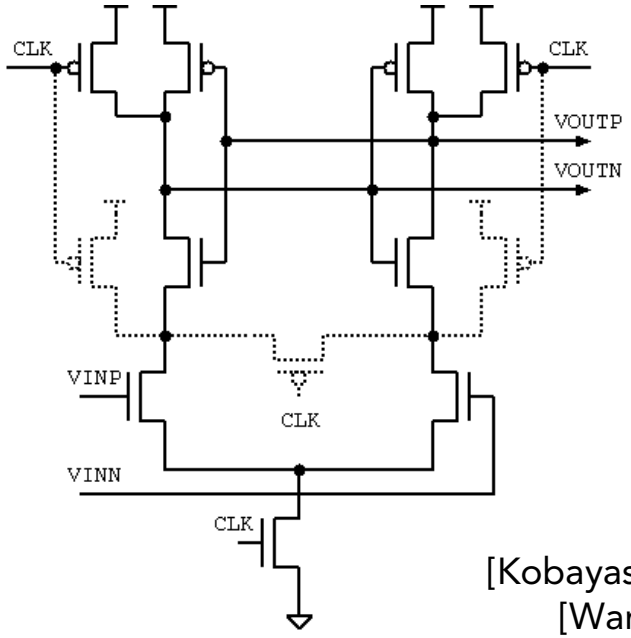

[Kobayashi JSSC 1993]
[Wang JSSC 2000]

- CML latch
  - Trade-off between power and speed
  - Draws static current
  - Reduced swing output levels

- StrongARM latch (=DCVS)
  - Power efficient due to regenerative amplification
  - No static current
  - CMOS output levels
  - Intrinsically faster

Adaptive Equalization

- Why adaptive?
  - In most applications the channel characteristic is not known a-priori
  - Marketing requires covering a large spectrum of applications anyway (same IC)
  - Slowly varying channel characteristic (temperature, humidity, aging, etc.)

- Different approaches
  - Least Mean Square (LMS) and Recursive Least Square (RLS) algorithms
  - Spectrum balancing (CTLE/LTE)
  - DFE taps = negative post-cursors in the pulse response (if possible to measure)
  - Zero-forcing on crossings (maximize eye width)
  - Zero-forcing on data (maximize eye height)

- TX-FFE
  - If EQ is done in the TX, it is difficult/costly to “close-the-loop” for adaptation
  - Mostly open-loop (operator driven) configuration

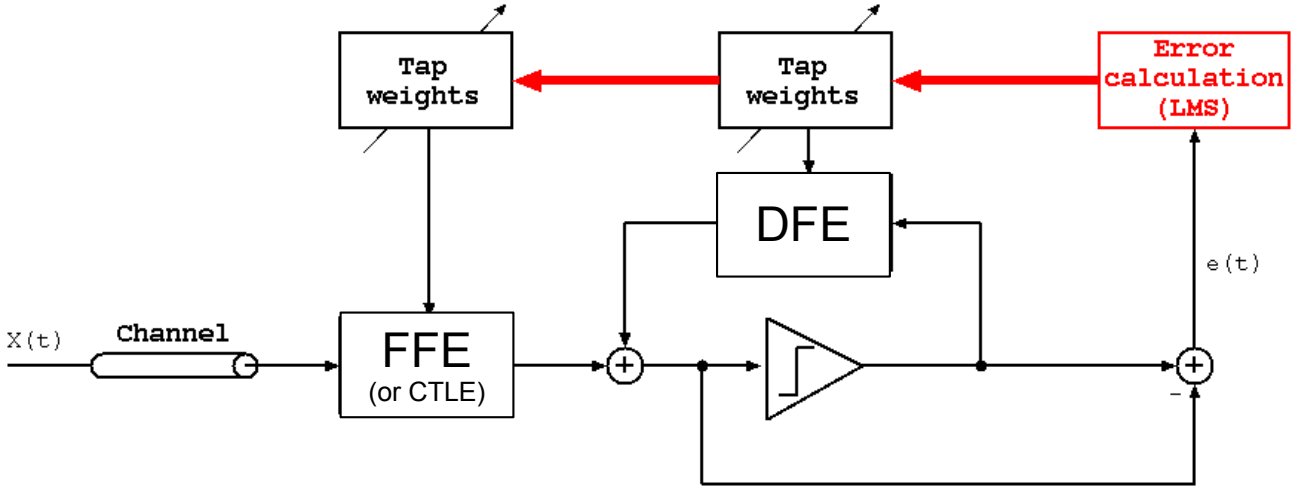

Other system impairments

BER strongly depends on SNR:
BER = ½*erfc(SNR/2)

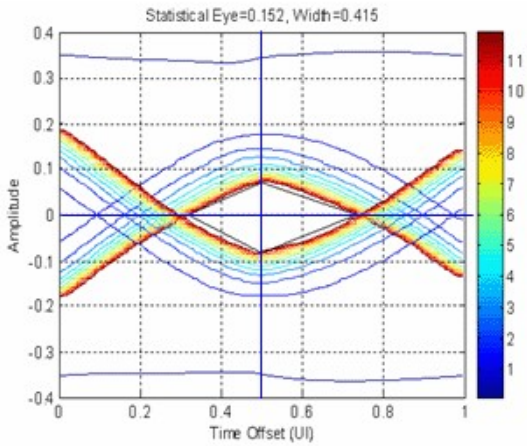

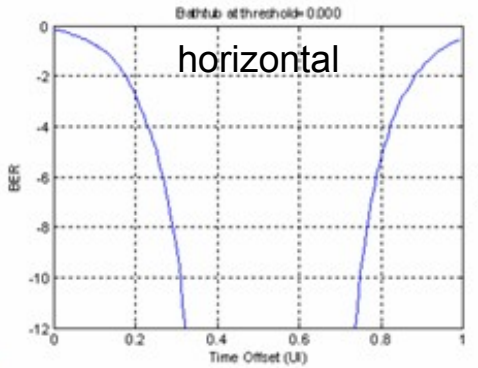

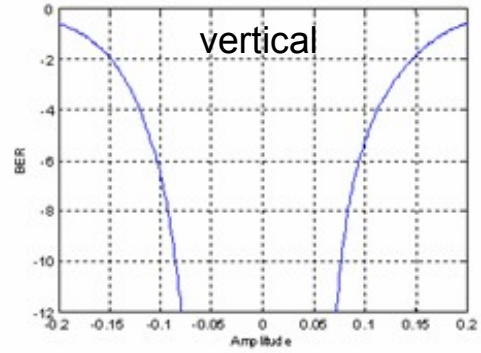

- Sources of errors (reducing SNR, one way or the other)
  - Reflections & Crosstalk
  - Attenuation (insertion loss)
  - Amplitude noise, linear distortion & offset (before slicer)
  - Sampling clock jitter (phase noise) -> reduces SNR, due to non-optimum sampling point
  - Process/Environment variation (PVTs) and device aging

- Avoid/Correct wherever possible
  - Broadband TX/RX termination (for low reflections)
  - Signal conditioning (enhance signal wrt. comparator sensitivity/noise)
  - Offset correction
  - Clock duty-cycle correction (for half-rate and lower)
  - Clock quadrature correction (for quarter-rate and lower)
  - startup / background calibration

SERDES design in advanced CMOS

- Motivation & Challenges
- Architectural solutions
- Implementation details
- Future trends & remarks

Market pushing for 56-112-224Gbps!

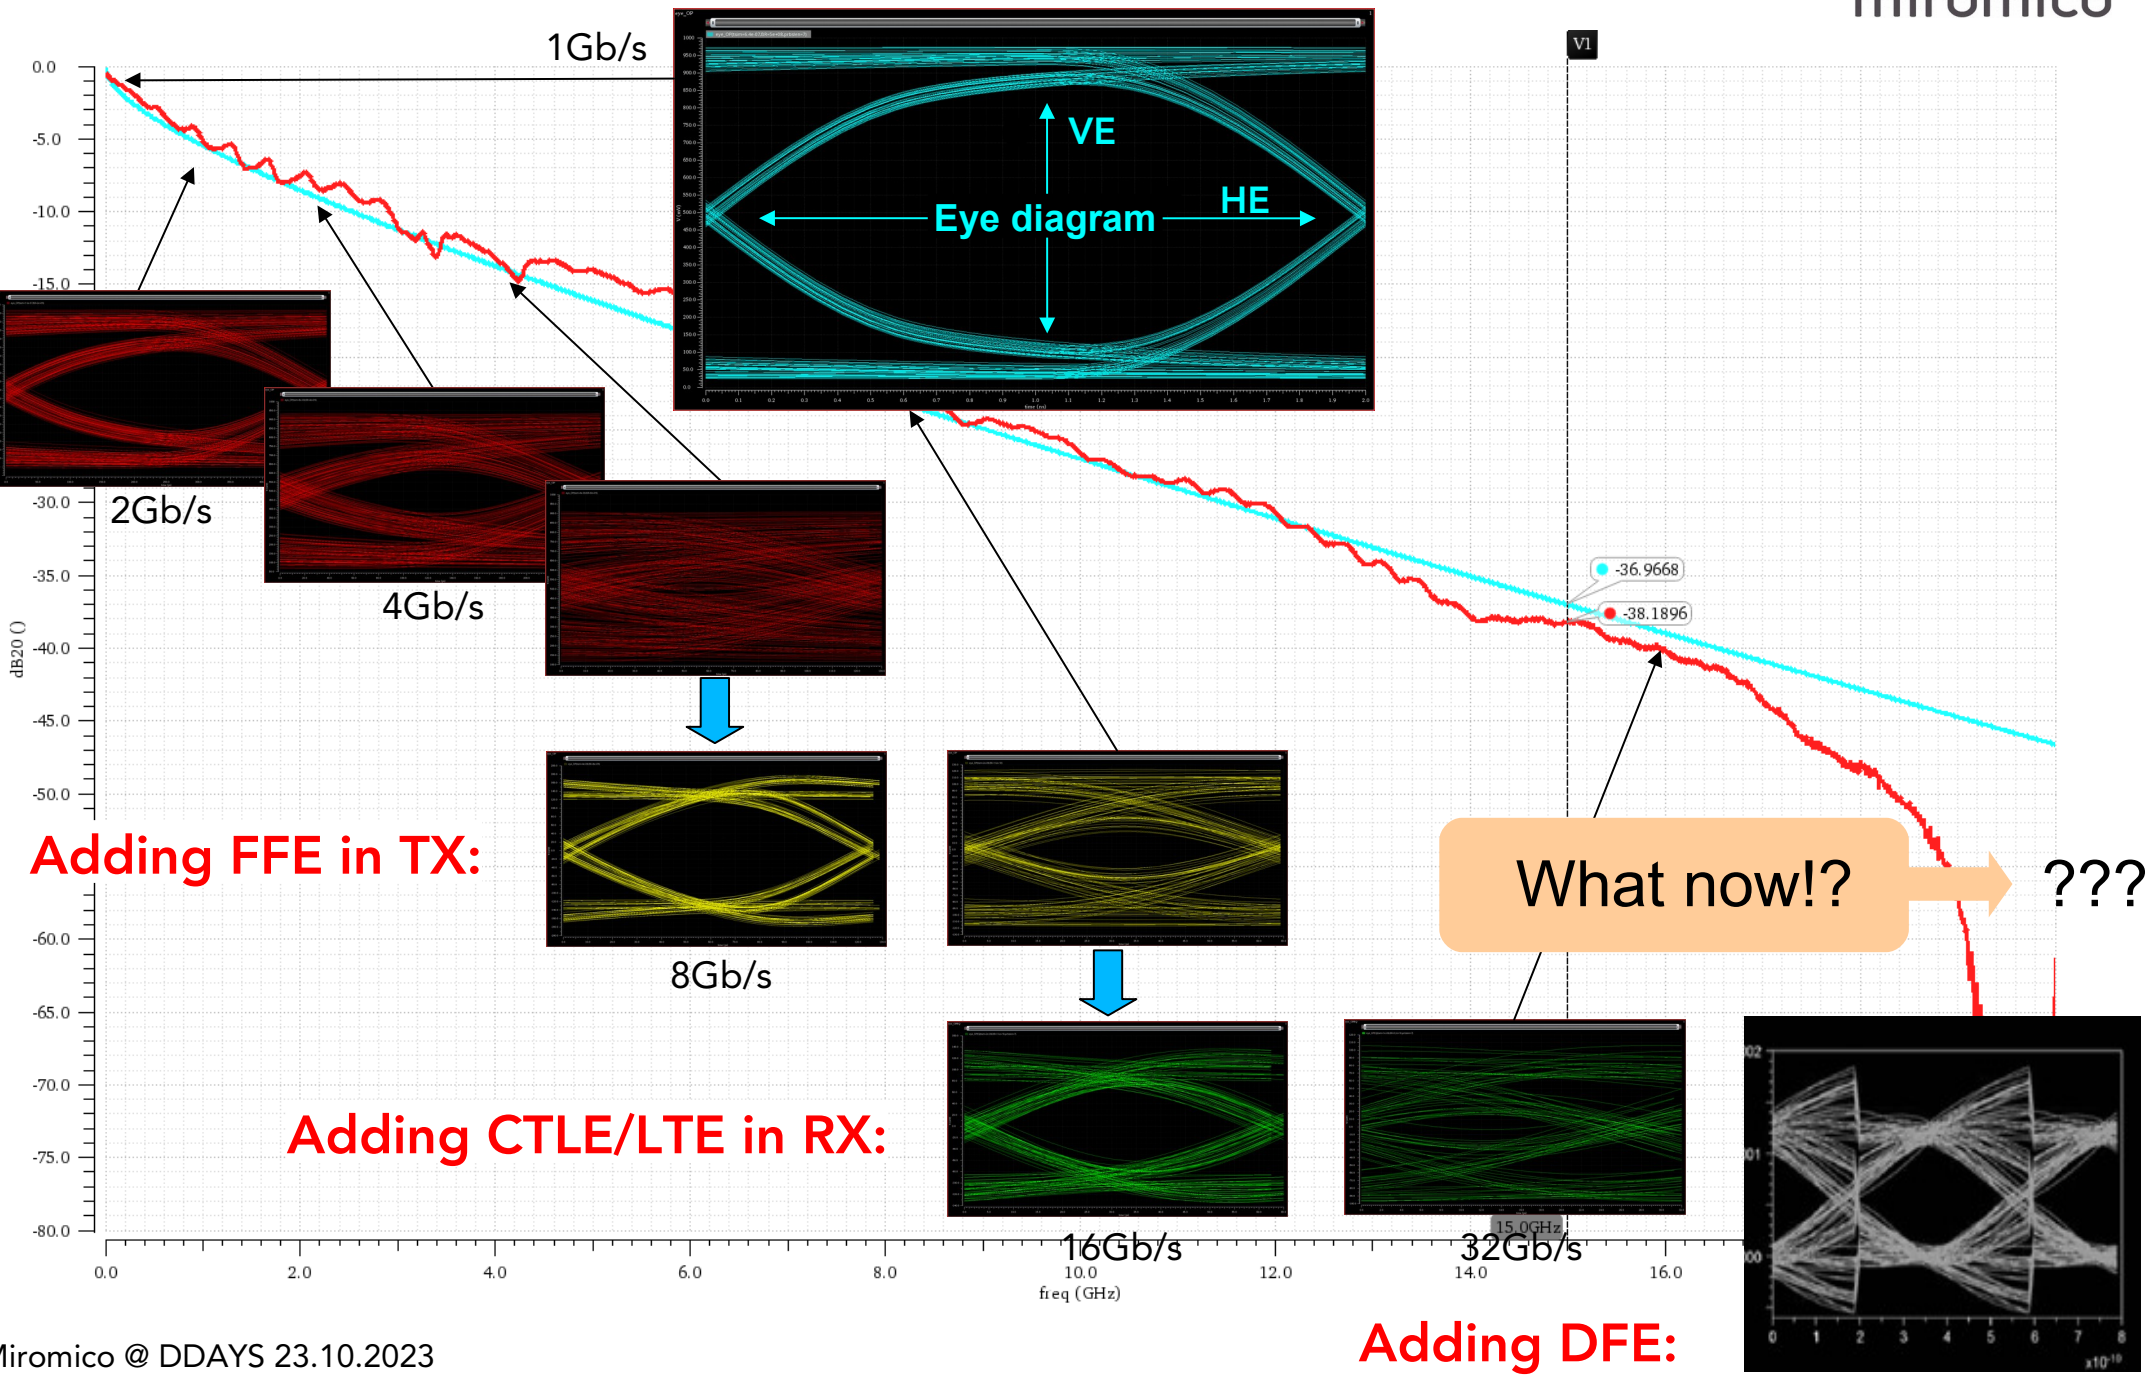

Alternative modulation schemes

Instead of going faster, send more information at the same baud-rate!

- NRZ (aka. PAM-2)
  - Binary levels
  - One transition per bit (higher BW requirements wrt. multi-level schemes)
  - Profits of the full voltage amplitude

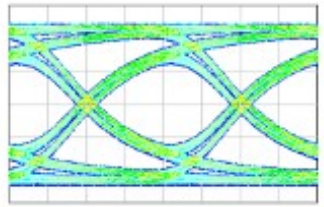

- PAM-4
  - Quaternary levels
  - Requires half of the bandwidth (wrt. NRZ) for the same bit rate
  - ~9.5dB SNR penalty due to more levels in the same voltage amplitude
  - More complex implementation (esp. for Sampling, EQ, CDR, …)

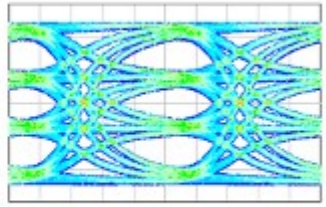

- NRZ vs PAM-4?
  - Forced choice, in case you must comply with a given standard
  - Depends (mainly, but not only) on the channel
    - Compare the channel loss at double the frequency vs. the 9.5dB SNR penalty coming from PAM4 (usually consider ~11dB or higher to account for higher complexity, sensitivity, etc.)
  - Implementation preferences...

Encoding schemes

PAM-4 is usually accompanied by improved error-correction coding

- Forward Error Correction compensates for higher sensitivity of PAM-4
  - FEC codes allow to dramatically reduce required BER (for same channel loss) or allowed channel loss (for same BER). Recover from 10e-5 BER to 10e-15
  - Expensive: large amount of logic, additional latency (though becoming better with new technologies, new schemes, etc.)
- Other encoding schemes also used for
  - Improving DC balance
  - Ensuring minimum transition rate for clock recovery
- Overhead (redundant bits) and code "efficiency"
- Some popular schemes
  - 8/10b, 64/66b, Reed-Solomon(n,m), ...

DAC/ADC based transceivers

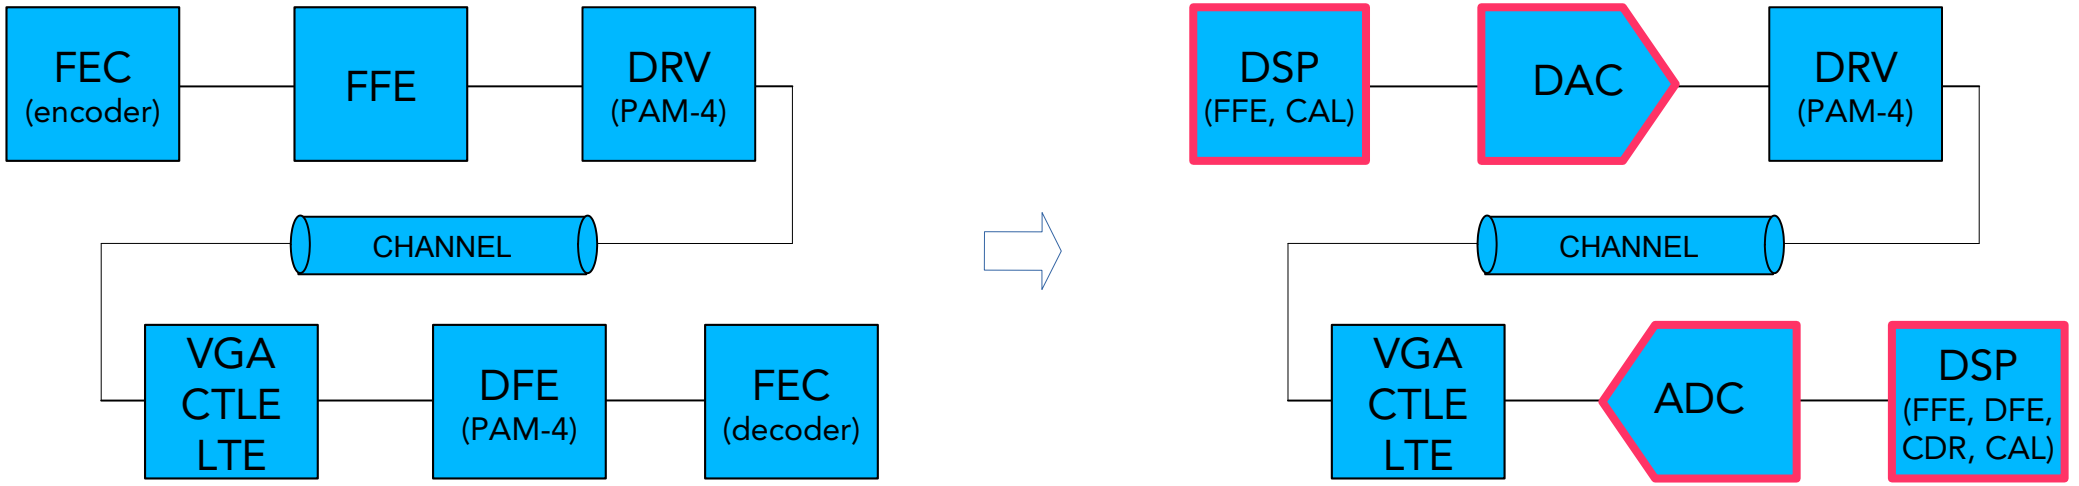

- Slicer-based RX
  - Mature architecture, can be made very efficient
  - DFE-loop timing can drive power very high as the transmission rates keep increasing
  - Complexity and power explode with multi-level modulation (PAM-4)
  - Does not scale well in technologies where Ft & RC bandwidth don't scale correspondingly

- ADC-based RX (& DAC-based TX)
  - Push EQ on the digital side (scalable, portable, versatile, etc.)
  - Intrinsically compatible with multi-level signaling (PAM-4/8/…)
  - Enables more complex EQ in DSP
  - Analog EQ still beneficial to reduce requirements on ADC resolution
  - Analog calibration enables "digital-like" scaling of DAC/ADC front-ends
  - Better power trade-off for modern technologies, higher data-rates, more stringent EQ requirements

Generalized differential signaling

- Research effort to improve bitrate/pin efficiency wrt. differential pairs
  - Single-ended: 1 bit every clock cycle over 1 wire (ratio: 1) → **Most efficient, BUT not differential!**
  - Differential pair: 1 bit every clock cycle over 2 wires (ratio: 0.5)
  - MIPI C-PHY: 2.28 bits every clock cycle over 3 wires (ratio: 0.76)
  - Kandou Chord Signaling (eg. ENRZ, EP3L, CNRZ, ...)
  - ENRZ (= Ensemble-NRZ): 3 bits every clock cycle over 4 wires (ratio: 0.75)

MIPI C-PHY

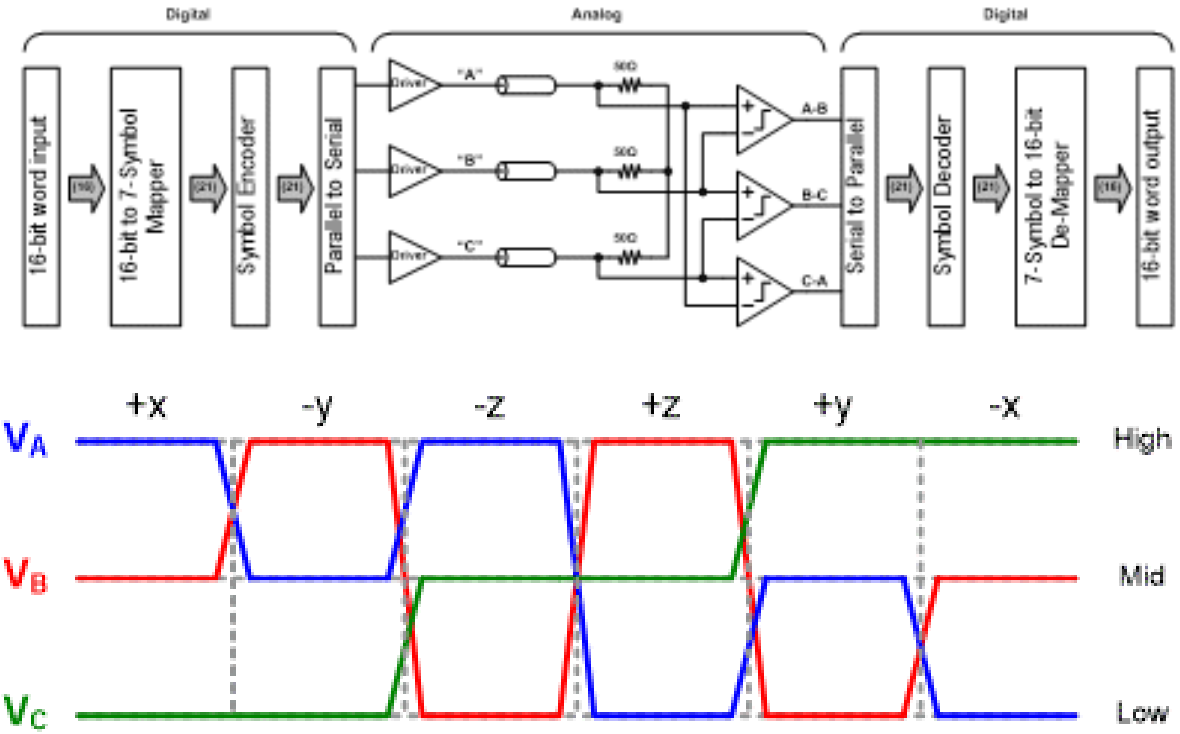

Wire State	Receiver Output		
	A-B	B-C	C-A
+x	1	0	0
-x	0	1	1
+y	0	1	0
-y	1	0	1
+z	0	0	1
-z	1	1	0

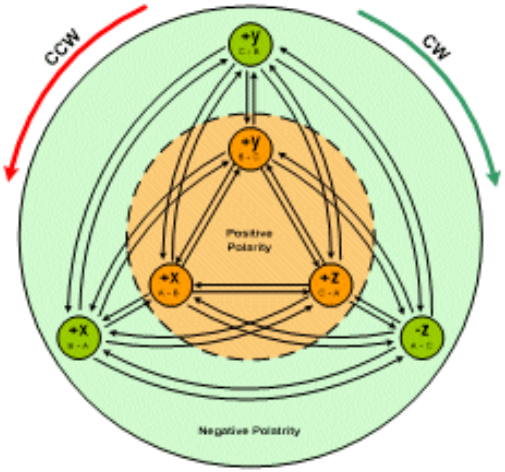

[https://www.design-reuse.com/articles/43501/mipi-d-phy-c-phy-combo-type-architecture.html]

Do you see the light?

- Optical is slowly but steadily replacing copper for long-reach interconnects:
  - Legacy support is keeping us electrical as long as possible
    - Huge investments have been made in existing infrastructures
  - Electrical-to-optical interface is still somehow exotic
    - Integration of modulators, filters and photo-diodes requires large areas
    - Mechanical mounting of the couplers clashes with the heat-sinks or the PCB
    - Difficult to reach the same integration density as the current electrical solutions
  - On-chip purely optical processing isn't a thing yet; we still rely on the good old silicon transistor to do the heavy lifting
- But certainly:
  - Long Reaches (LR) are becoming increasingly difficult on electrical conductors
  - E-O interfaces are moving closer to our chips (inside the package? inside the IC?)
  - Increasing pressure on XSR (and USR) links: chip-to-(optical)module

Transmission over existing copper cables still possible thanks to sophisticated (but inexpensive) silicon solutions, such as coding/equalization schemes and error correction algorithms of ever increasing complexity!

And of course there is more to SERDES design than just EQ...

EQ
FFE
CTLE
DFE
(others?)

TERM/ESD
Line impedance termination
ESD protection
AC-coupling

Continuous adaptation
CTLE peaking regulation
FFE/DFE coefficients

AGC
Programmable amplification
Background gain regulation

SerDes

Serialization
Deserialization

Data encoding
8/10b, 64/66b, FEC, Scrambling

CDR
Clock & Data Recovery
(data must contain traces of the clock, in the form of transitions)

Support functions
Offset cancellation
Clock corrections (DCC, QCC)
Monitoring and measurements
Testability

Logic
Link/Application layer glue logic
Error Detection/Correction

...but each of them would require another talk!

Thanks for your attention!

http://www.miromico.ch