Inside the LpGBT

Paulo Moreira CERN
Geneva - Switzerland

Ecole IN2P3 de Microélectronique

15 – 19 May 2017

Bénodet - France

Introduction

- Subject:
  - Data transmission systems in the context of HEP.
- What is the problem:
  - Transmission of Physics data from the frontend systems on the detectors to the DAQ systems in the counting room.
  - Transmission of control information from the counting room to the detector systems
- Communication systems [almost] universally process and transmit data in digital form:
  - Thus, [bidirectional] digital transmission between the frontend and the DAQ systems will be considered.
- Practical examples of circuits and systems will be drawn from the GBT and LpGBT projects!

Typical Data Link in HEP – On Detector

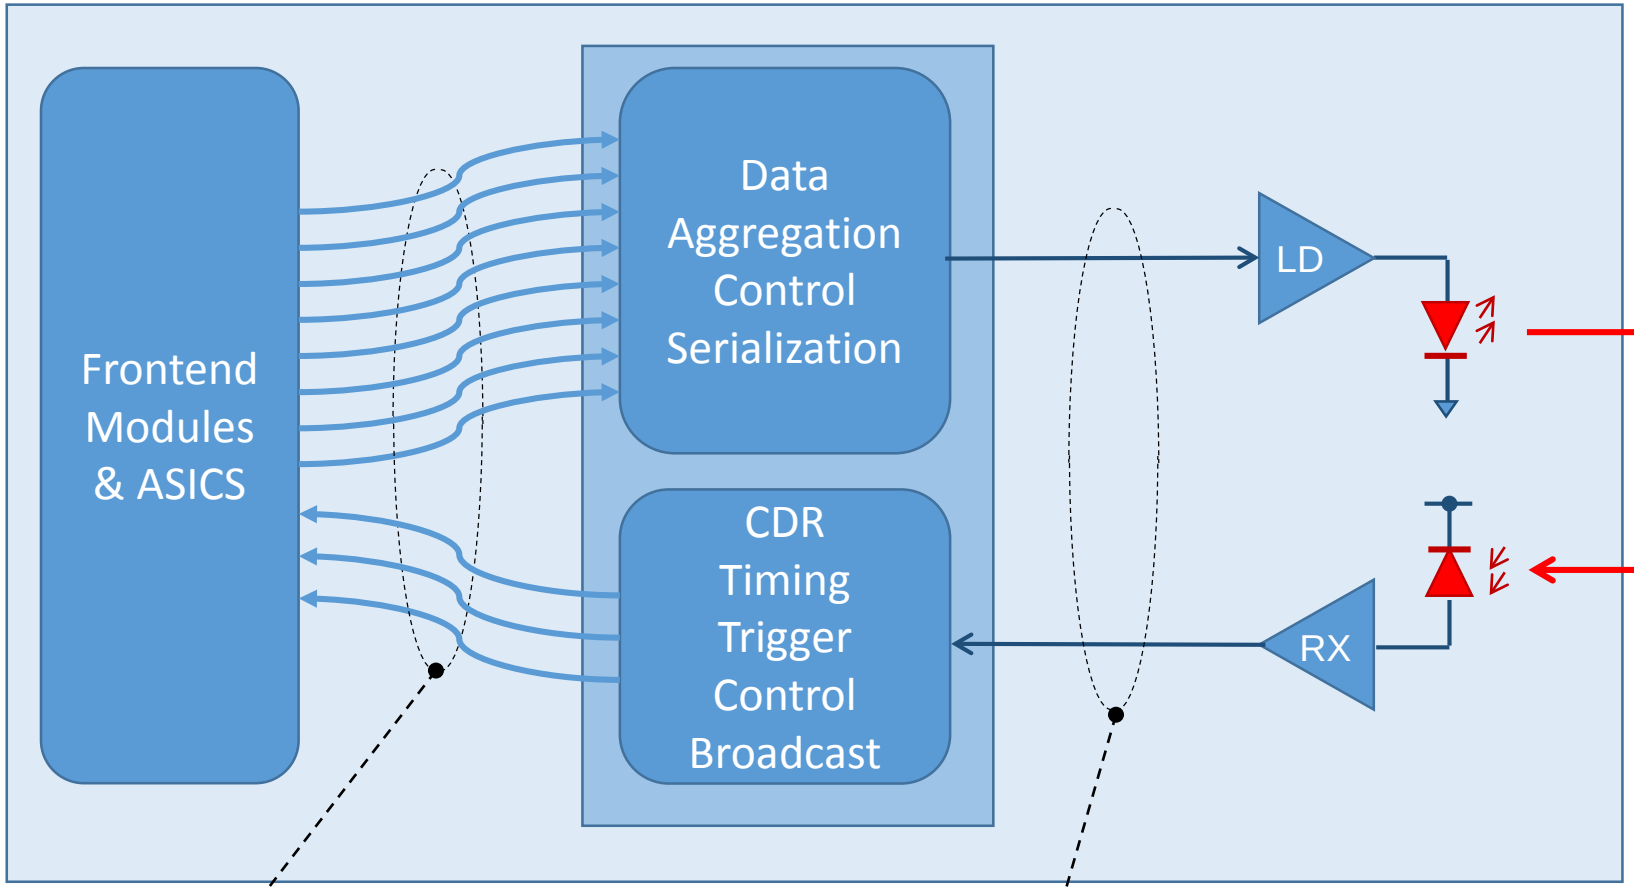

Electrical links [to the frontend electronics]
Data rate: 80 Mb/s to 1.28 Gb/s
(ATLAS pixel is considering 5 Gb/s)
Length: few cm to 1 or 2 meters

Electrical links [to the optoelectronics]
Data rate: 2.56 to 10.24 Gb/s
Length: few cms

Optical links
Data rate: 2.56 to 10.24 Gb/s
Length: 50 to 200 m

Radiation Hardness Required

Typical Data Link in HEP – In the Counting Room

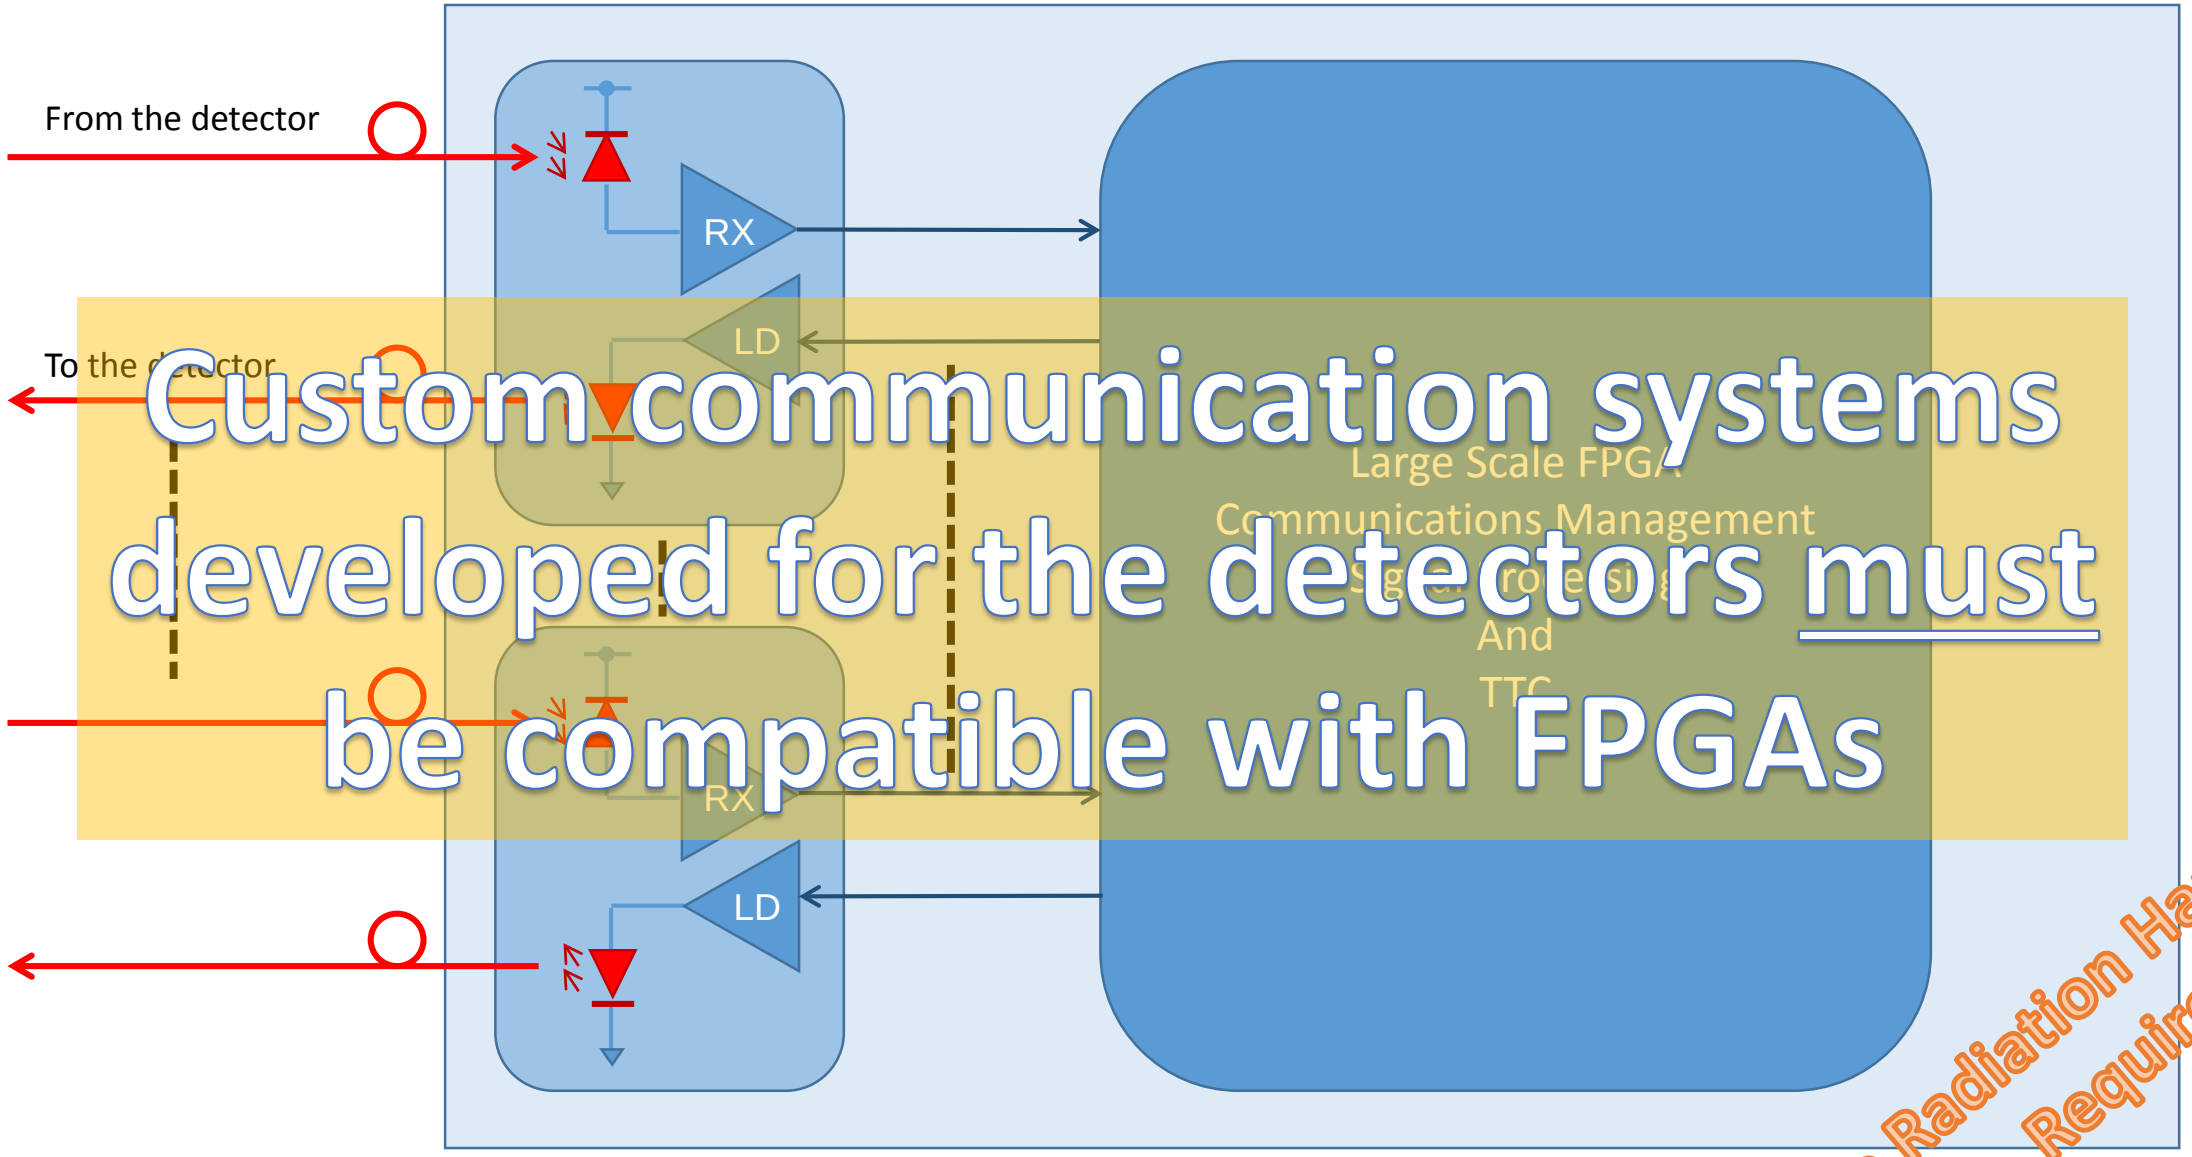

Custom communication systems developed for the detectors must be compatible with FPGAs

No Radiation Hardness Required

HEP Communications Idiosyncrasies

- Typically simpler that long haul communications systems:
  - Localized and short distance systems (< 200 m)
  - Bandwidth is not at a premium!
- However the radiation environment in which they are installed raises several challenges:
  - Total Ionizing Dose (TID) will degrade the performance of the systems to the point that they will eventually fail!
    - ASICs and Optoelectronics must be designed and qualified to guaranty “survival” during the lifetime of the detector!
  - Single Event Upsets (SEU) will disturb momentarily the operation of the systems/circuits!
    - They must be built to mitigate the effects of SEUs and avoid altogether system interrupts
- Besides these two “peculiarities”, HEP systems are based on the same principles as the telecommunication systems:
  - And thus should be compatible with them as much as possible.
  - More specifically, with FPGAs which are the ubiquitous building blocks of HEP DAQ systems!

Outline

- Basic concepts in data transmission systems
  - Building blocks
  - Impairments
  - Optical vs Electrical
- Radiation in electronics and optoelectronics devices
  - Impact of radiation robustness on circuit performance
- The LpGBT – Materializing the concepts
- Connecting the ASICs to the world

Data Transmission Systems
Basic Concepts

Why Digital Communications?

- Digital communication systems are widely used for data acquisition, trigger and control links in HEP:
  - A few exceptions: e.g. the CMS tracker data links
- Advantages:
  - Digital communications can be done virtually error free
    - When appropriate codes are used, transmission errors can be:
      - Detected
      - Corrected
  - They are “natural” for digital systems
  - Easy handling of multiple data sources and destinations
  - Easy re-routing between multiple sources and destination

Optical Link Architecture

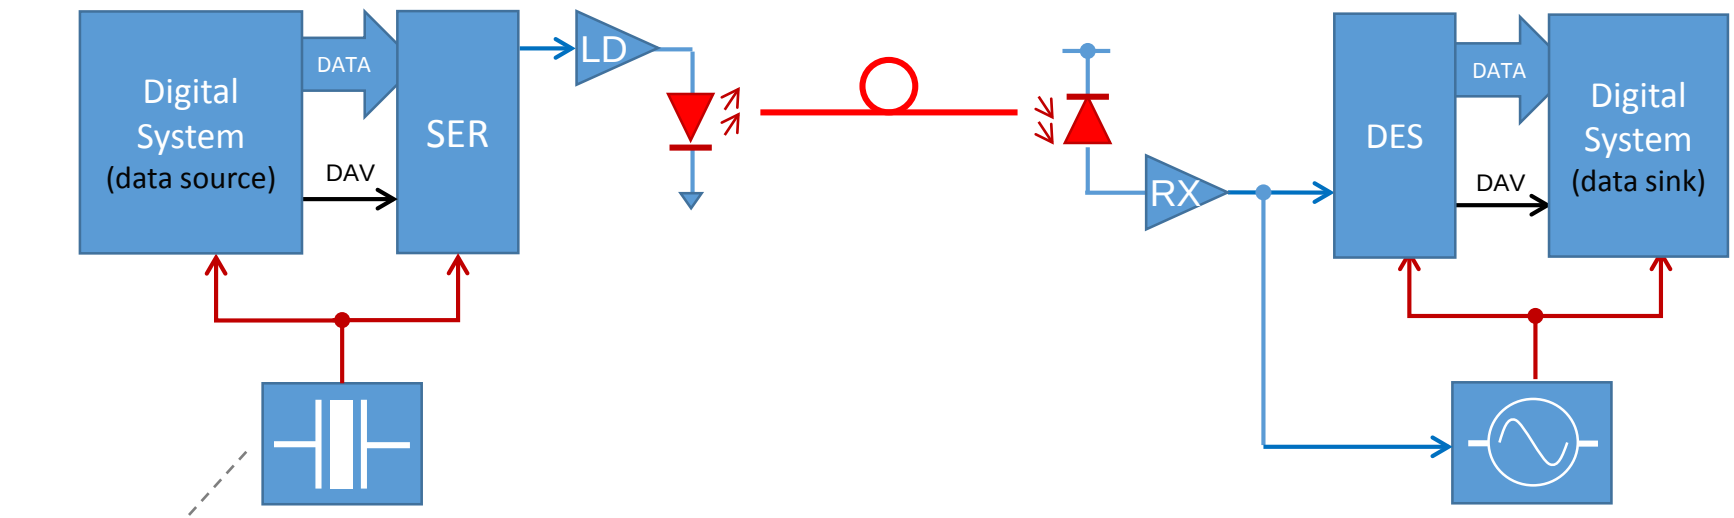

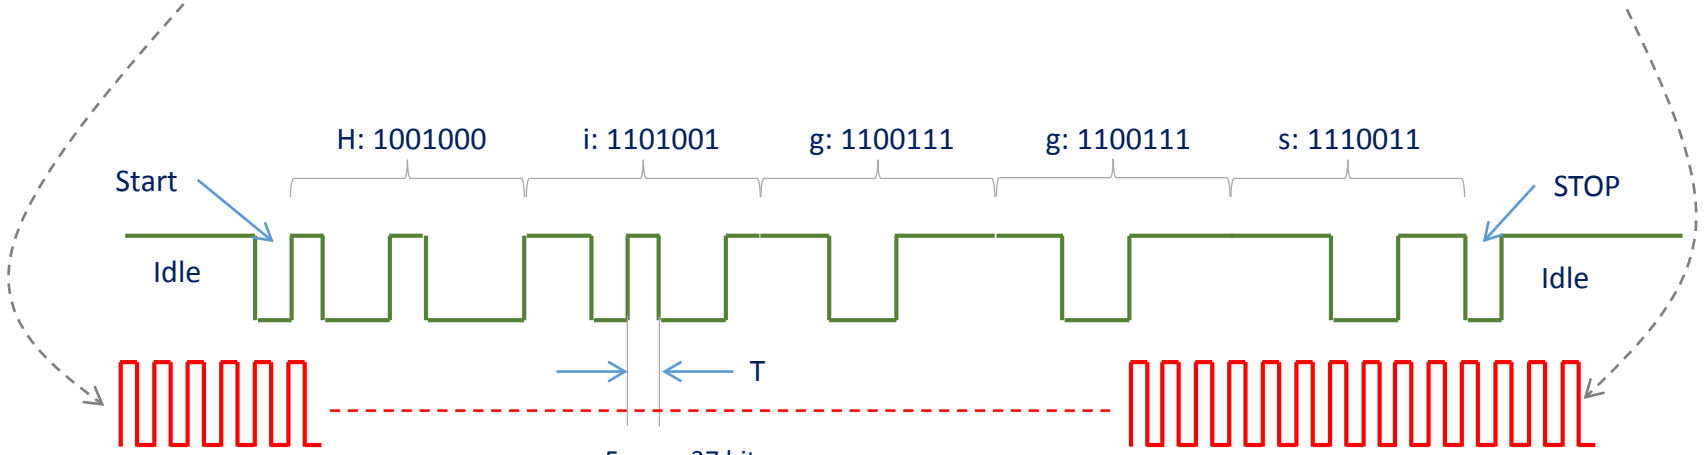

Bit-Rate = 1/T

Frame, 37 bits:

- 35 data bits
- 2 control bits
- "Code efficiency" : 35/37 = 94%

BW – Limited Channel (1/2)

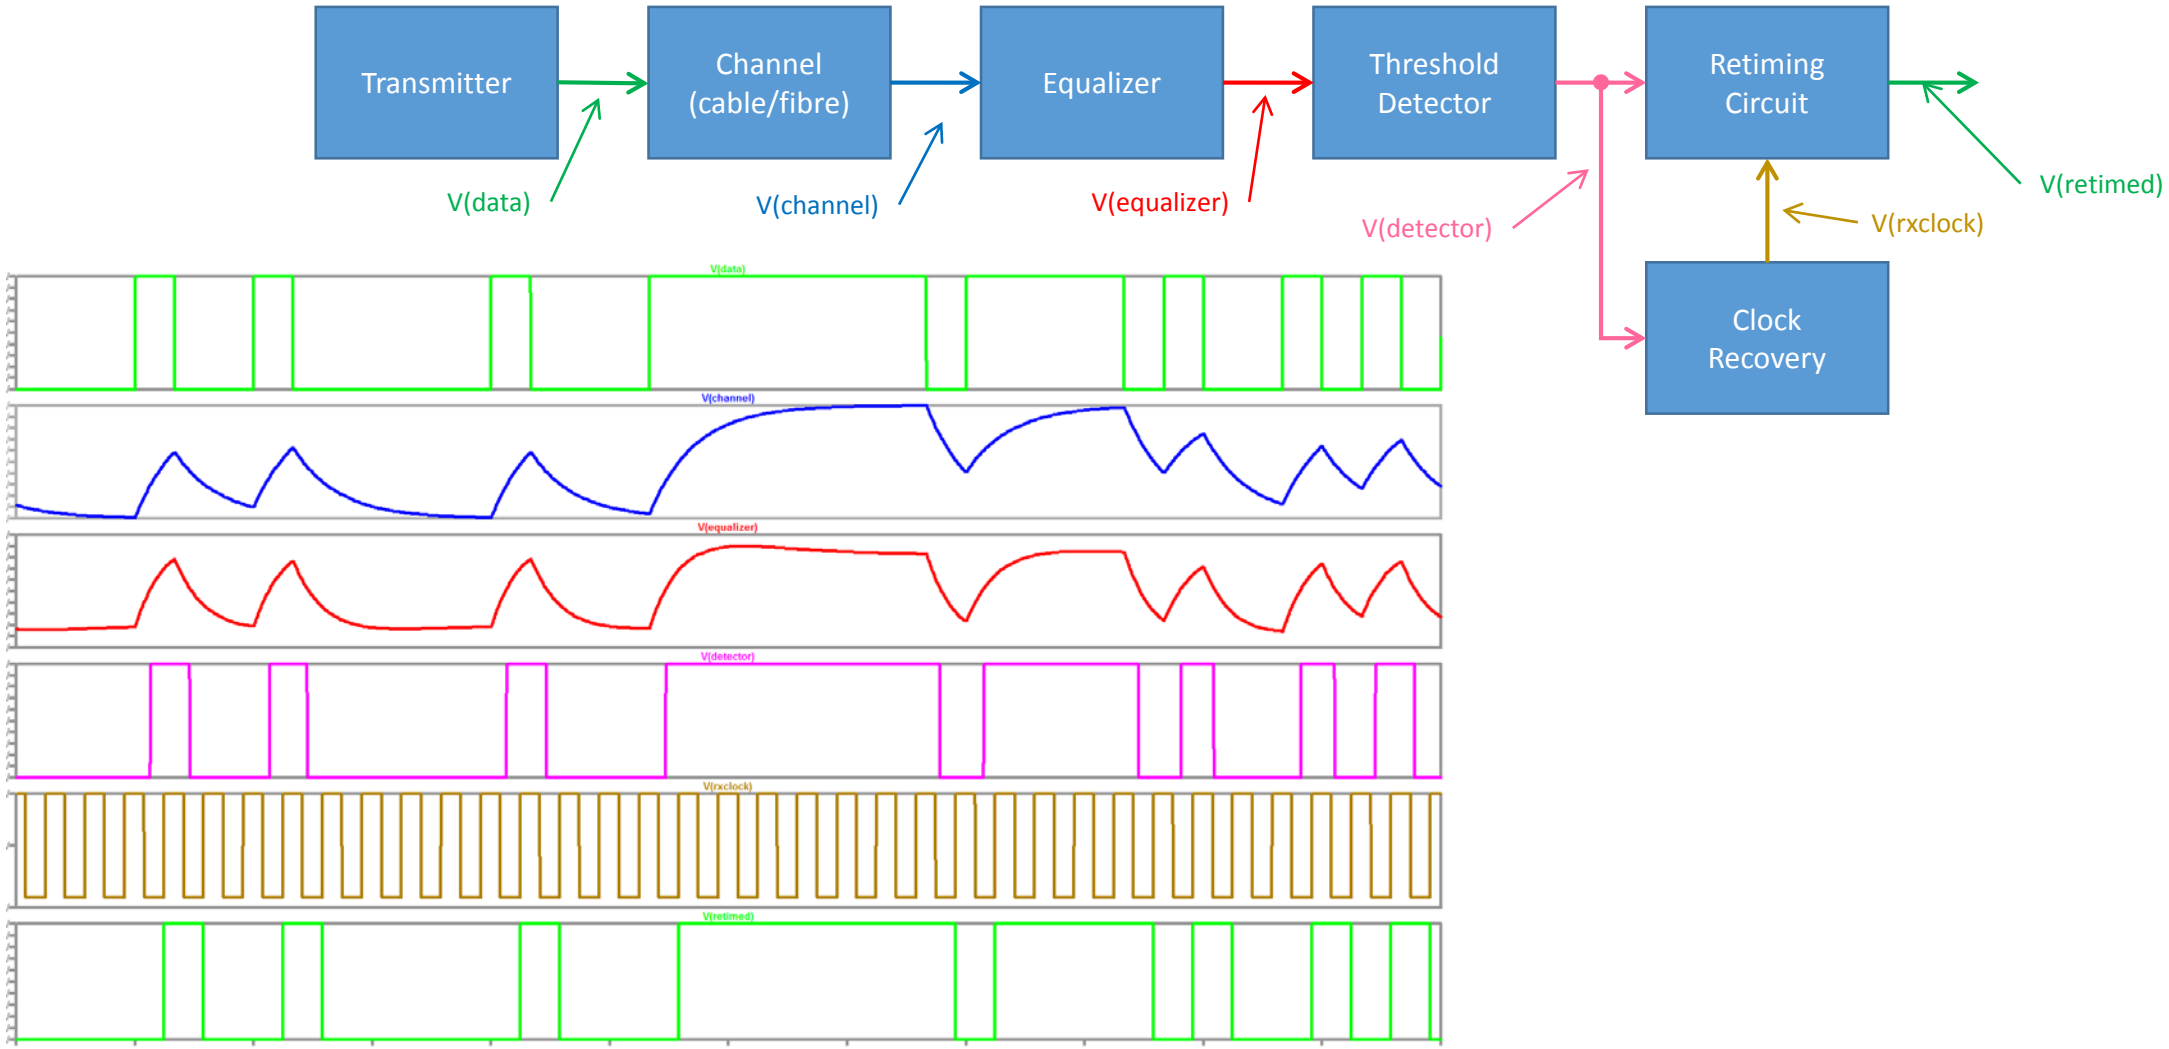

BW – Limited Channel (2/2)

- Due to the limited channel bandwidth the transmitted pulses are broadened in time.
- If the broadening is significant, the signal corresponding to one symbol (bit) will overlap in time the signal of the next symbol!
  - This is called Intersymbol Interference (ISI)
- ISI is seen in an eye-diagram as:
  - Vertical eye-closure:
    - Reduction of the vertical eye-opening
  - Horizontal eye-closure:
    - Random positions of the threshold level crossings (Jitter or Phase Noise)
- Equalization (high-pass filtering in this case) can be used to "restore" the signal high frequency content:
  - Reduces ISI
  - Improves Jitter
  - Can never fully compensate the channel
- Two additional steps are needed to restore the signal to their "original" shape:
  - Threshold detection, restores the symbol levels by comparing a signal with a "pre-defined" level
  - The signal is retimed using the recovered clock, restoring the signal levels with minimum jitter
- Two additional phenomena impair error free data transmission:
  - Attenuation
  - Noise

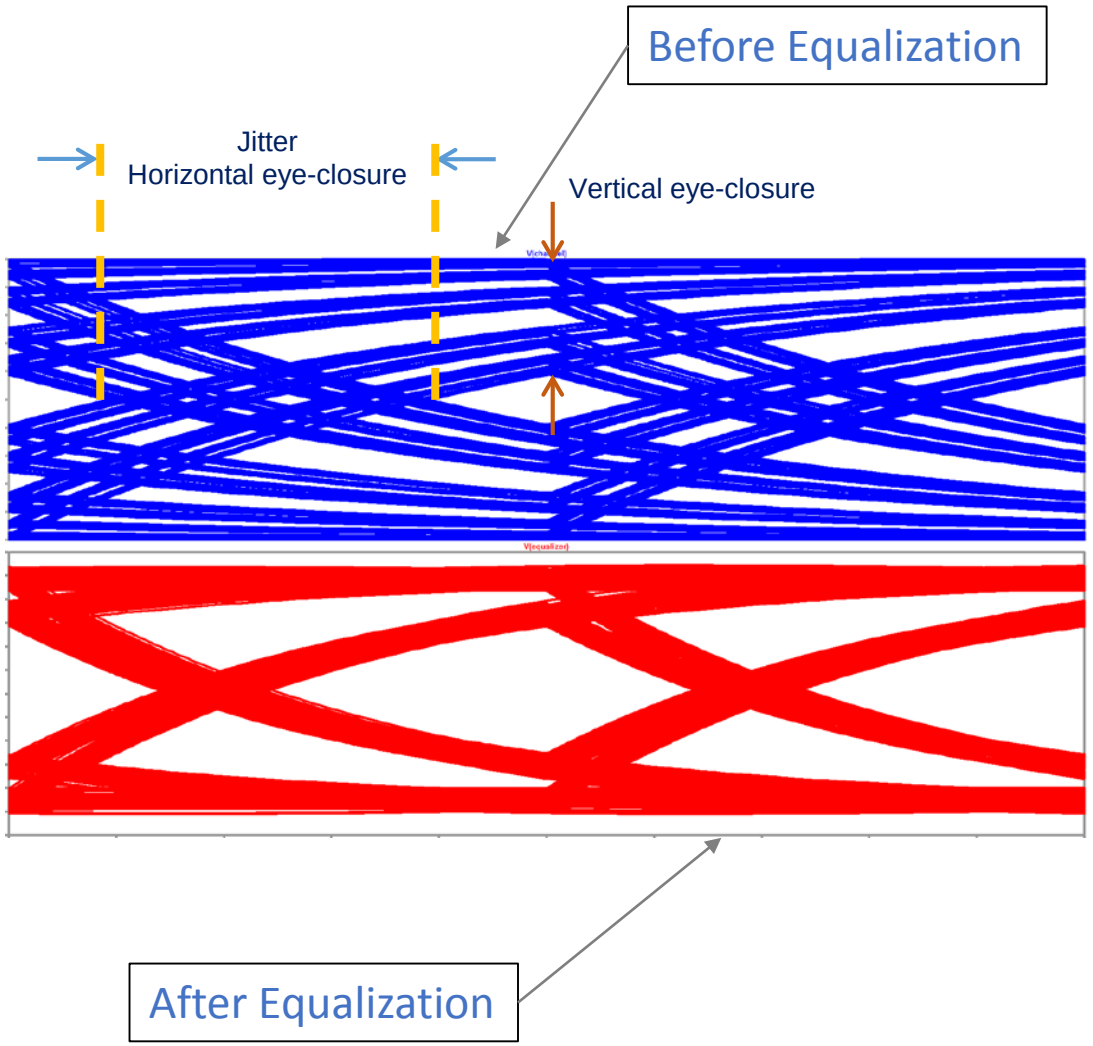

Noise - Amplitude

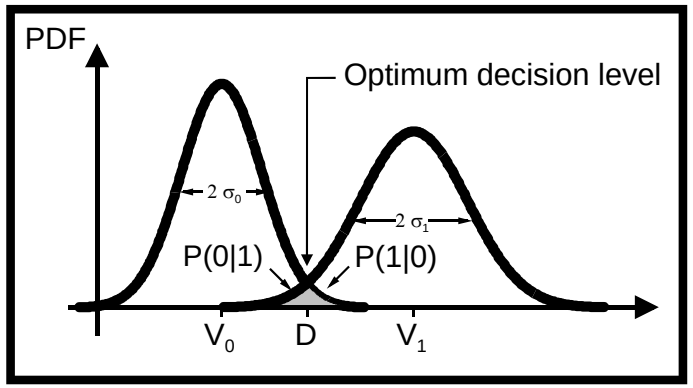

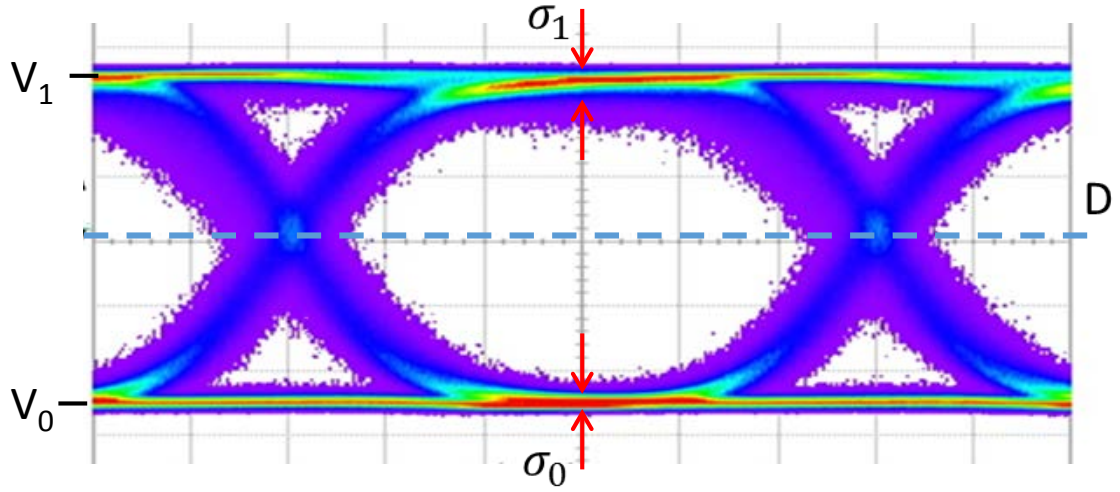

- The average bit error probability:

$$P_e = P(0|1) \cdot P(1) + P(1|0) \cdot P(0)$$

- It is a strong function of the signal to noise ratio:

$$SNR = \frac{V_1 - V_0}{\sigma_n} \quad \text{assuming} \quad \sigma_n = \sigma_1 = \sigma_0$$

- In the limit the measured Bit Error Rate **BER** is equal to P_e
  - SNR = 8.64 (18.7 dB) → BER = 10^{-9}
    - Testing time to "count" 100 errors: 20 s @ 5 Gb/s
  - SNR = 10.84 (20.7 dB) → BER = 10^{-15}
    - Testing time to "count" 100 errors: 231 days @ 5 Gb/s

$$P_e = \frac{1}{2}\, erfc\left(\frac{SNR}{2}\right)$$

Noise - Jitter

- Jitter is phase-noise:
  - Random position of the bit crossings.
- [As seen] Bandwidth limitations cause intersymbol interference that besides reducing the eye-opening also generates jitter
- Amplitude noise will also convert into jitter due to the finite rise and fall times of the signal!
- In summary: bandwidth limitations contribute to eye-closure due to ISI and conversion of amplitude noise into jitter.

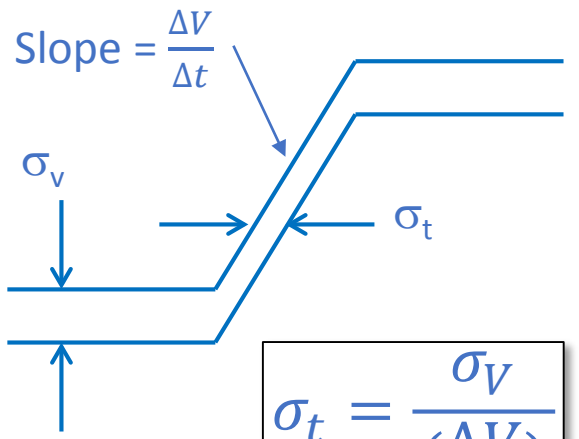

$$\sigma_t = \frac{\sigma_V}{\left(\frac{\Delta V}{\Delta t}\right)}$$

Example:

5 Gb/s PIN – Receiver

V_{dd} = 1.2 V

Sensitivity limit (10^{-9}): SNR = 8.64

$\sigma_V = V_{dd}/SNR$ = 139 mV

BW = 0.7*5 Gb/s = 3.5 GHz

$\Delta V/\Delta t = 2\pi$ x BW x V_{dd} = 26 mV/ps

σ_t = 5.3 ps

GBTIA 5 Gb/s)

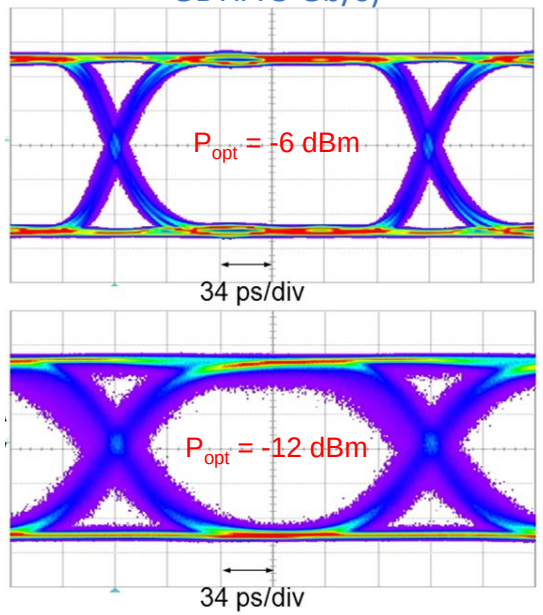

Jitter

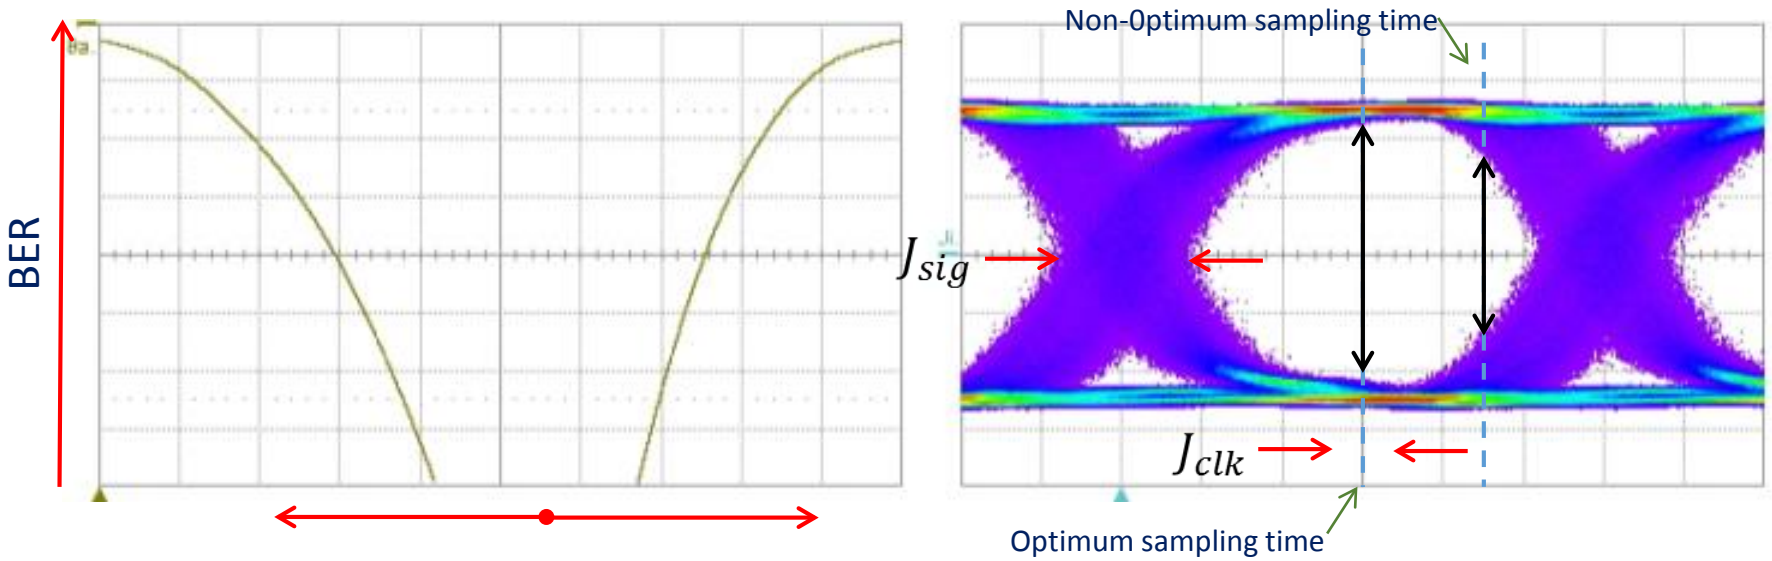

- Deviations from the optimum sampling instant drastically increase the BER because the SNR decreases due to:
  - Signal jitter
  - Signal finite rise and fall times
    - The signal magnitude |V –D| decreases because of the finite raising/fall times of the signal
- Two main causes for non-optimum sampling:
  - Retiming clock static phase error
    - E.g. due to an unbalance in the CDR PLL charge-pump
  - Retiming clock jitter
    - E.g. due to the CDR tracking behaviour or VCO noise

Error Control Coding

- Due to Noise or ISI the received message might differ from the transmitted
  - Some of the transmitted bits will be wrongly detected
- Error control coding introduces extra bits in the transmitted message:
- These allow to:
  - Detect the presence of errors
  - Correct detected errors
- Error control is done at the expense of bandwidth

Even parity: Parity bits are computed such that the number of "1s" in each row and column is even

Transmitted message

Row parity

	[6]	[5]	[4]	[3]	[2]	[1]	[0]	R.P.
H	1	0	0	1	0	0	0	**0**
i	1	1	0	1	0	0	1	**0**
g	1	1	0	0	1	1	1	**1**
g	1	1	0	0	1	1	1	**1**
s	1	1	1	0	0	1	1	**1**
C.P.	**1**	**0**	**1**	**0**	**0**	**1**	**0**	**1**

Column parity

Matrix parity

Received message

	[6]	[5]	[4]	[3]	[2]	[1]	[0]	R.P.
H	1	0	0	1	0	0	0	**0**
i	1	1	0	1	0	0	1	**0**
g	1	1	0	0	1	1	1	**1**
f	1	1	0	0	1	1	0	**1**
s	1	1	1	0	0	1	1	**1**
C.P.	**1**	**0**	**1**	**0**	**0**	**1**	**0**	**1**

Discrepancy between the received row/column parities and the parities computed by the receiver!

Line Coding

- Apart from detecting and correcting errors, coding is also applied to condition the signal to the transmission medium and/or the transmitter/receiver architecture. This is called line coding.
- Most common addressed problems are:
  - DC wander caused by AC coupling ("high-pass" type response):
    - DC blocking between circuits
    - Laser-driver mean power control
    - Offset cancelation circuits in pin-receivers
  - Lack of signal transitions to keep the RX clock locked on the data
- Line codes typically have the following properties:
  - limit the low frequency content in the signal spectrum
  - Guarantee a high density of symbol transitions
- Some popular examples:
  - 8B/10B
  - Scrambling

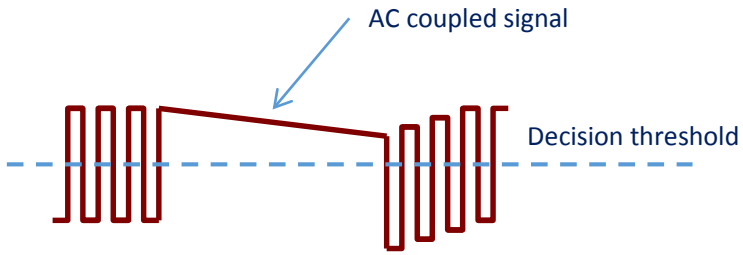

3B/4B

Input	Output		
Disparity	+	0	-
000	1101		0010
001		1001	
010		1010	
011		0011	
100		1100	
101		0101	
110		0110	
111	1011		0100

Why Optical Links in HEP?

- Transmission distance?
  - Telecom: up to transcontinental distances! (long – haul)
  - HEP: at most a few hundreds of meters (short – haul)
- Bandwidth?
  - Telecom: the larger the better!
  - HEP:
    - The larger the better?
    - Does using THz bandwidths per fibre makes sense?
      - Needs data aggregation systems inside the detectors!
      - "Third generation" systems are complex, power hungry [and costly]
    - 5 to 10 Gb/s systems are being engineered for Phase II upgrades
- Especially important for HEP:
  - Small optical fibre cross section (material budget)
  - Immunity to interferences
  - Non EMI generator
  - No contribution to ground loops

Optical vs Electrical

- Optical fibers allow bandwidths and transmission distances that are orders of magnitude that of electrical cables.
- Most systems installed in HEP are "1st Generation" systems, that is they use 850 nm lasers and multimode, graded index, fibers.
- Warning: for such systems the "dispersion limit" limits the transmission distance to less than 200 m at 10 Gb/s
  - Probably [LHC] installed fibers need to be replaced to support 10 Gb/s transmission

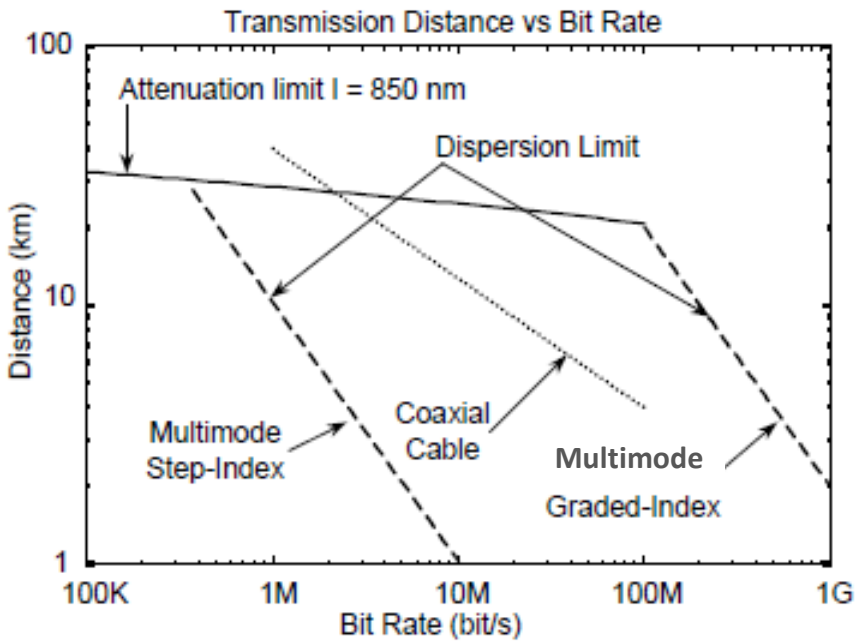

Ecole IN2P3 d'electronique numerique
Aussois, France
23-27 November 1998

Radiation Effects

Radiation (TID) In CMOS Circuits (65 nm)

- TID Induced "I_{on}" Degradation
  - PMOS are more sensitive to TID than NMOS:
    - Minimum size NMOS: -20% @ 200 Mrad
    - Minimum size PMOS: -60% @ 200 Mrad
  - TID induced ΔI_{on} degradation is a function of the channel length L!
    - The longer, the smaller the degradation, but
    - High speed circuits need minimum L transistors!
  - ΔI_{on} degradation is also a function of the channel width W:
    - The wider, the smaller the degradation
  - Enclosed Layout Transistors (ELT) are the least sensitive!
  - Depending on the circuit function, the use of large W might introduce a power penalty!
    - Particularly important for digital circuits where the gate count is high!

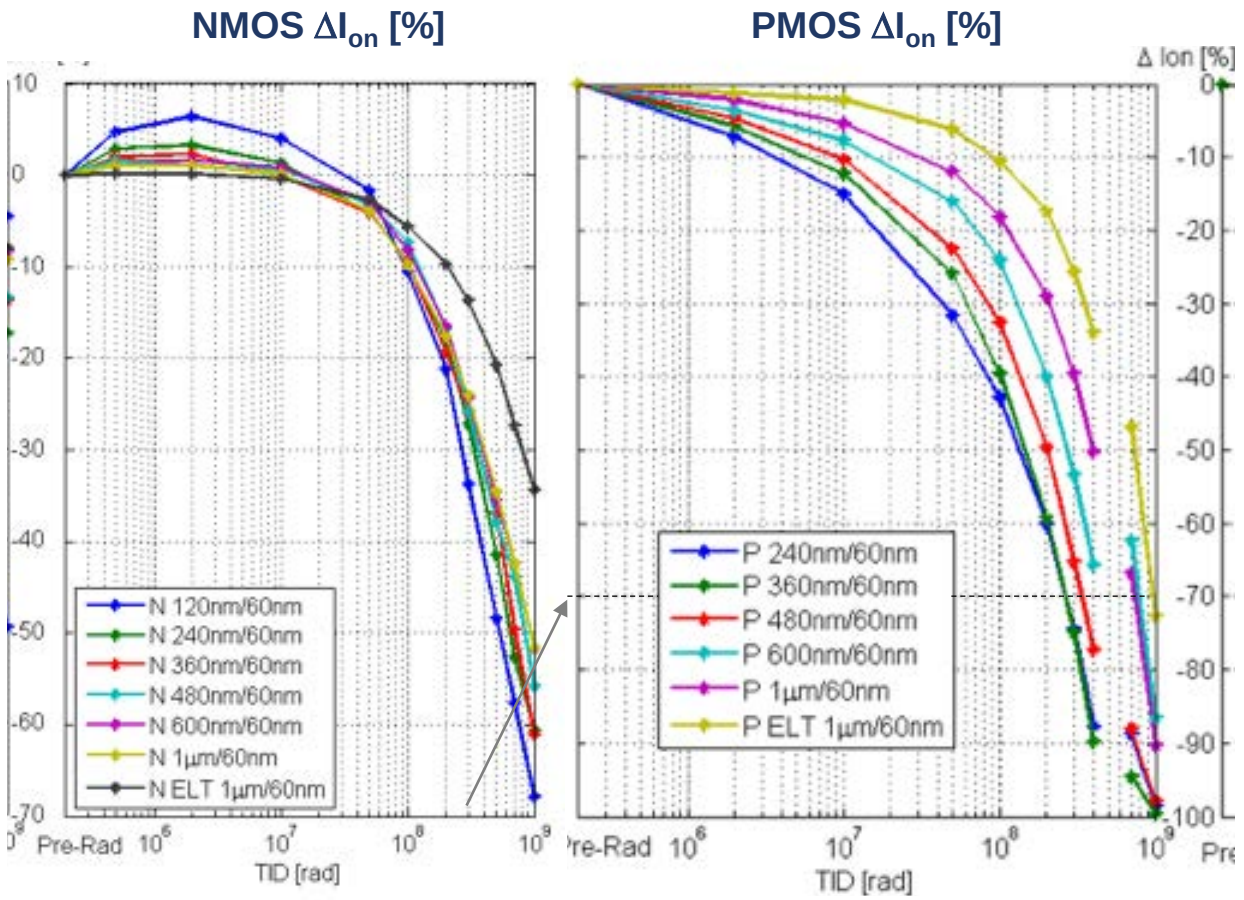

Standard Cell Libraries TID Sensitivity

- Set of ring oscillators irradiated up to 700 Mrad (at room temperature):
  - Annealing not shown in the graph
- As expected, from the tests of individual devices, libraries made out of large devices are radiation harder than the smaller devices counter parts
  - Large devices lead to large power dissipation and larger area circuits!

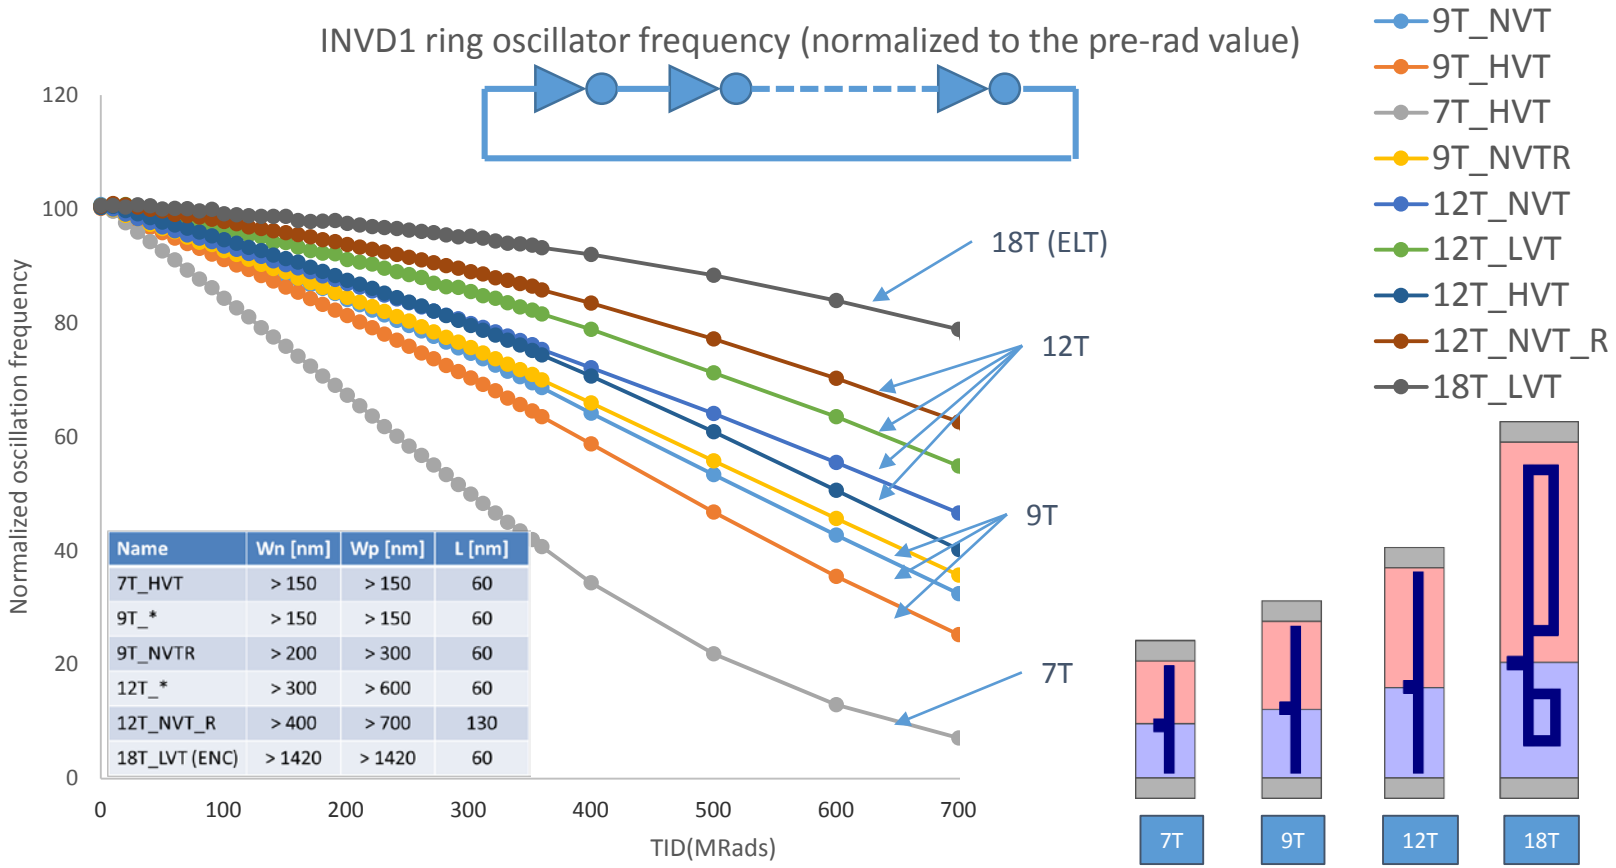

Name	Wn [nm]	Wp [nm]	L [nm]
7T_HVT	> 150	> 150	60
9T_*	> 150	> 150	60
9T_NVTR	> 200	> 300	60
12T_*	> 300	> 600	60
12T_NVT_R	> 400	> 700	130
18T_LVT (ENC)	> 1420	> 1420	60

Single Event Upsets

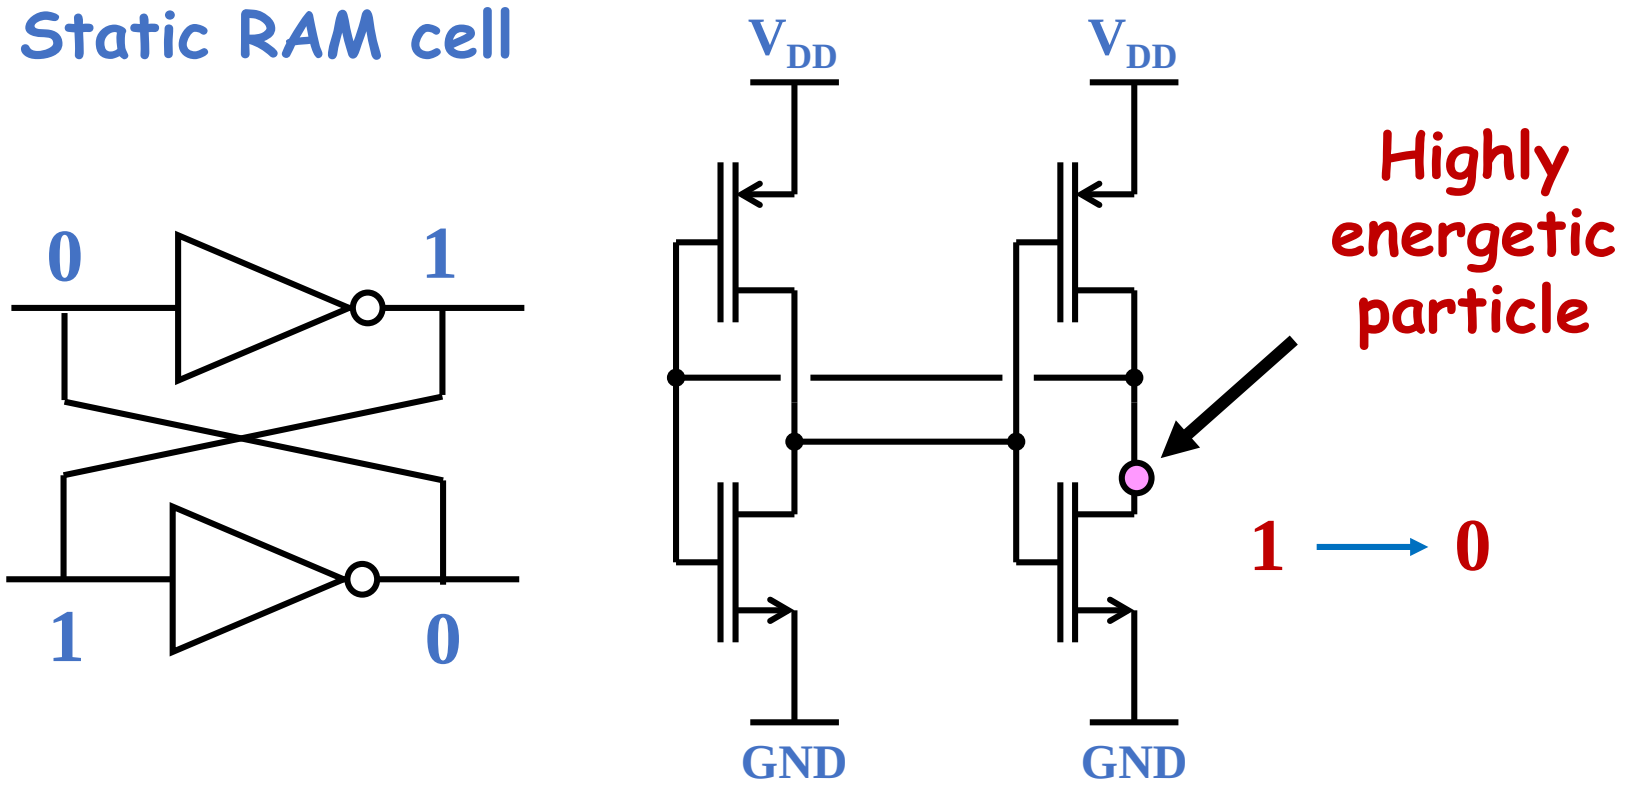

TMR (full triplication)

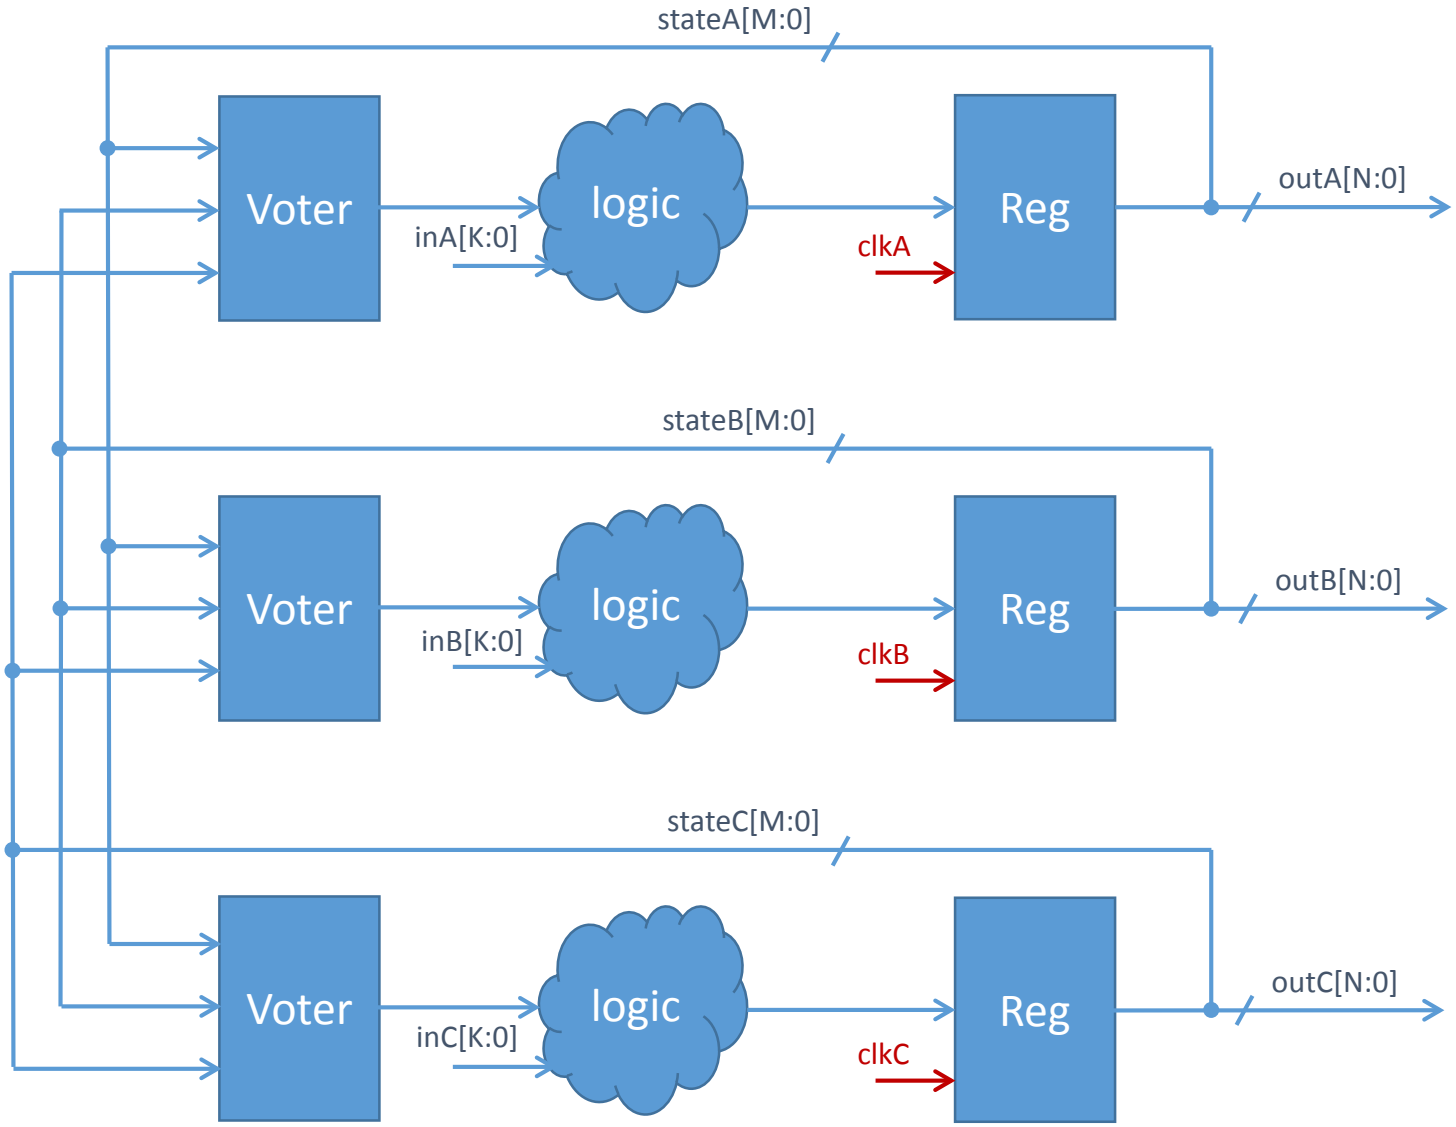

TID Robust Design: Impact on Performance (1/2)

Case study:

- In the LpGBT, for frame alignment, a special frequency pre-scaler is needed:
  - Normally divides by 2 but, upon request, executes a single divide by 3 cycle
  - The pre-scaler waits for the request to be released before it is ready to accept the next one
  - It works at the down-link clock frequency: 2.56 GHz
  - The pre-scaler allows to shift the received down-link frame by 390 ps at a time
- A logic synthesizer was used to synthesise the pre-scaler under two conditions:
  - Relaxed target frequency: 100 MHz
  - Target operation frequency: 2.56 GHz
  - In both cases non-TMR and TMR logic was synthesised

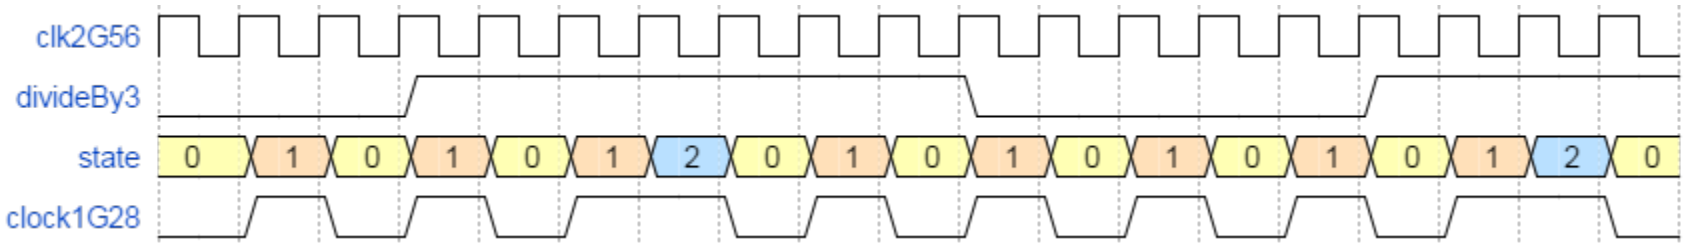

TID Robust Design: Impact on Performance (2/2)

- Gate count:
  - No TMR: 32 Gates
  - TMR (full): 157 Gates [x4.9]
    - Besides triplicating the logic, the registers, the clocks and adding voters, additional buffers are needed to account for the higher loading of the signals!
- As a consequence:
  - TMR can increase the power dissipation by 4.0 to 4.9 times!
  - TMR reduces the operation frequency by a factor between 1.4 to 1.5
- The large device libraries, 12TLVT and 18TLVT (ELT), are the only ones capable of achieving the target operation frequency for this particular circuit with a TMR implementation!
- The 18TLVT library has an edge on radiation tolerance so it was selected for the high speed circuits of the LpGBT!
- The smaller libraries, 7T and 9T have prohibitive degradation with radiation and can only be used in circuits for which the timing margin is very relaxed
  - Notice that, using HVT libraries (or smaller size) does not always pays off in terms of dynamic power if high frequency operation is the target. Larger amounts of buffering and slower rise and fall times ("high short circuit current") might degrade the power consumption figure.

Target: 100 MHz

Library	TMR	Leakage [uW]	Dynamic [uW]	Area [um2]	t_d(critical) [ns]	f_{max} [GHz]
7THVT	YES	0.05	130	770	1836	0.5
9THVT	YES	0.06	116	838	1692	0.6
9T	NO	0.06	20	166	605	1.7
	YES	0.26	99	688	934	1.1
12THVT	YES	0.09	146	963	1174	0.9
12T	YES	0.59	129	842	692	1.4
12TLVT	YES	1.88	124	790	533	1.9
18TLVT (ELT)	NO		45	349	235	4.3
	YES		220	1477	405	2.5

Target: 2.56 GHz

Library	TMR	Leakage [uW]	Dynamic [uW]	Area [um2]	t_d(critical) [ns]	f_{max} [GHz]
7THVT	YES	0.06	1614	896	1305	0.8
9THVT	YES	0.08	1593	1083	1131	0.9
9T	NO	0.07	346	187	447	2.2
	YES	0.37	1392	885	604	1.7
12THVT	YES	0.12	1919	1228	773	1.3
12T	YES	0.78	1727	998	463	2.2
12TLVT	YES	2.55	1688	937	362	2.8
18TLVT (ELT)	NO		544	349	235	4.3
	YES		2189	1441	376	2.7

CML - TID

- CML gates offer the best speed potential in “CMOS” circuits:
  - Only NMOS transistors used;
  - Loads are resistors:
    - Very low parasitic capacitance, compared with C_{db} of the PMOS transistors
  - Speed to a large extent set by the RC time constant in the drain circuit:
    - If I_T (or ΔV) are stabilized the delay suffers little variation with the process parameters!
    - Thus in principle robust against TID
- Main disadvantages:
  - Requires a bias circuit
  - Relatively large power consumption!
    - DC power consumption
  - Relatively large area

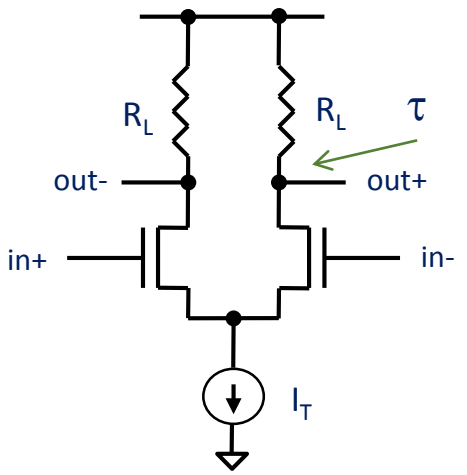

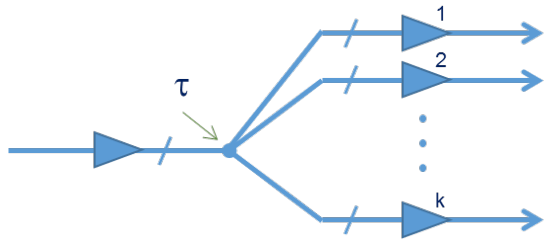

Miller multiplied C_{gd}

Gate resistance

$$\tau = \frac{\Delta V}{I_T}\left[C_{gd} + C_{db} + \left(k + \frac{R_g}{R_L}\right)\left[C_{gs} + (1 + g_m R_L)C_{gd}\right]\right]$$

$$\frac{\Delta V}{I_T} = R_L$$

CML – Buffer Simulations

- Three process corners:
  - SS, TT, FF (actives and passives)
  - Bias current [almost] constant
  - Voltage and Temperature kept constant to assess the effects of process only:
    - This to "simulate" the effects TID < 200 Mrad in the NMOS (SS process)

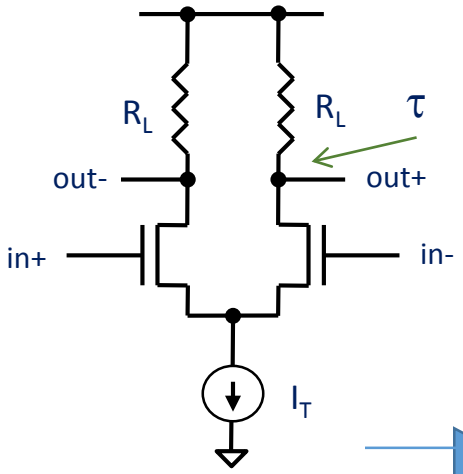

R_L = 1.2 kΩ

W/L = 2.4 µm / 60 nm

I_T = 500 µA

f = 5.12 GHz

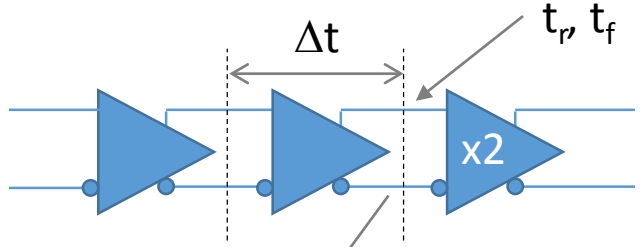

Delay variation very much dependent on the R variations

Power dissipation is high!

	Min	Max	Variation
I_T [mA]	0.49	0.52	6%
R_L [kΩ]	1.04	1.36	27%
Power [mW]	0.59	0.62	5%
ΔV [mV]	532	686	25%
td [ps]	3.6	4.5	23%

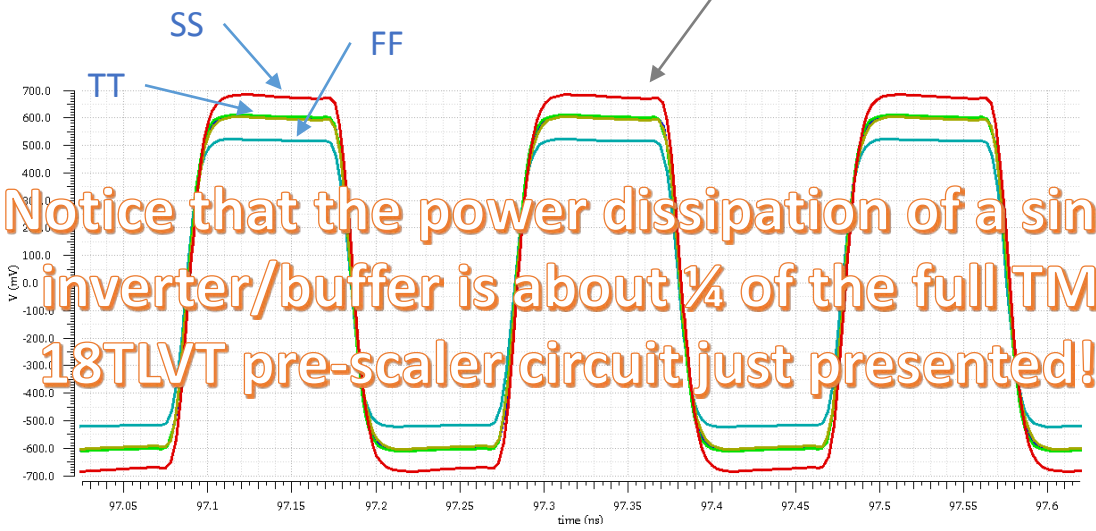

Notice that the power dissipation of a single inverter/buffer is about ¼ of the full TMR 18TLVT pre-scaler circuit just presented!!!

Radiation Damage in Optoelectronics

- Damage mechanism dominated by Displacement Damage (DD) caused by Non-Ionizing Energy Loss (NIEL) from heavy particles (neutral/charged hadrons, energetic leptons).
- **PIN - Diodes:**
  - Increasing of the dark current by several orders of magnitude (from pA to mA)
  - Reduction of the responsivity ($R = I_{ph}/P_{opt}$)
- **VCSELs:**
  - Increase in the threshold current and voltage
  - Decrease of the laser slope-efficiency ($S = P_{opt} / I_{mod}$)
  - VCSELs display higher radiation tolerance than EE diodes

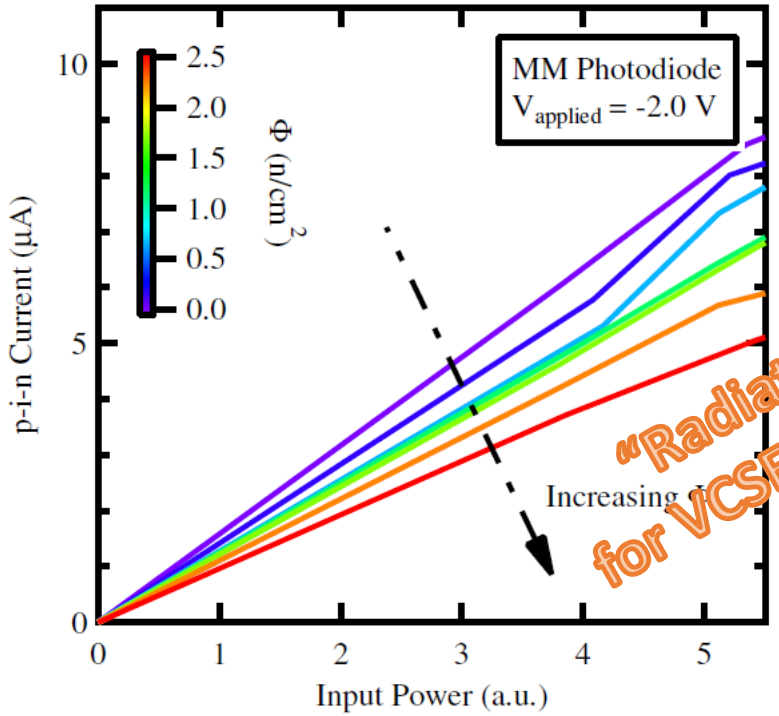

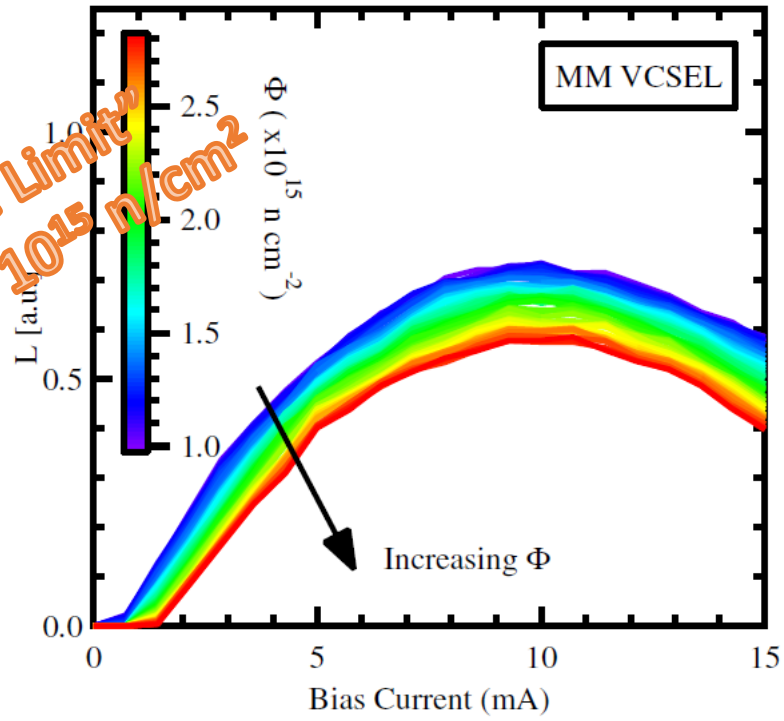

"Radiation Tolerance Limit" for VCSELs and PINs: ~10^{15} n/cm^2

LpGBT ASIC – High Speed SerDes

- Data transceiver with fixed and "deterministic" latency both for up and down links.
  - Clocks and Data
- Down link:
  - 2.56 Gb/s
  - FEC12
- Up link:
  - 5.12 Gb/s or 10.24 Gb/s
  - FEC5 or FEC12
- E-Links:
  - Data rates:
    - 160 /320 / 640 / 1280 Mb/s
  - Count:
    - FEC5
      - Up to 28 @ 160 Mb/s
      - Up to 7 @ 1.28 Gb/s
    - FEC12
      - Up to 24 @ 160 Mb/s
      - Up to 6 @ 1.28 Gb/s

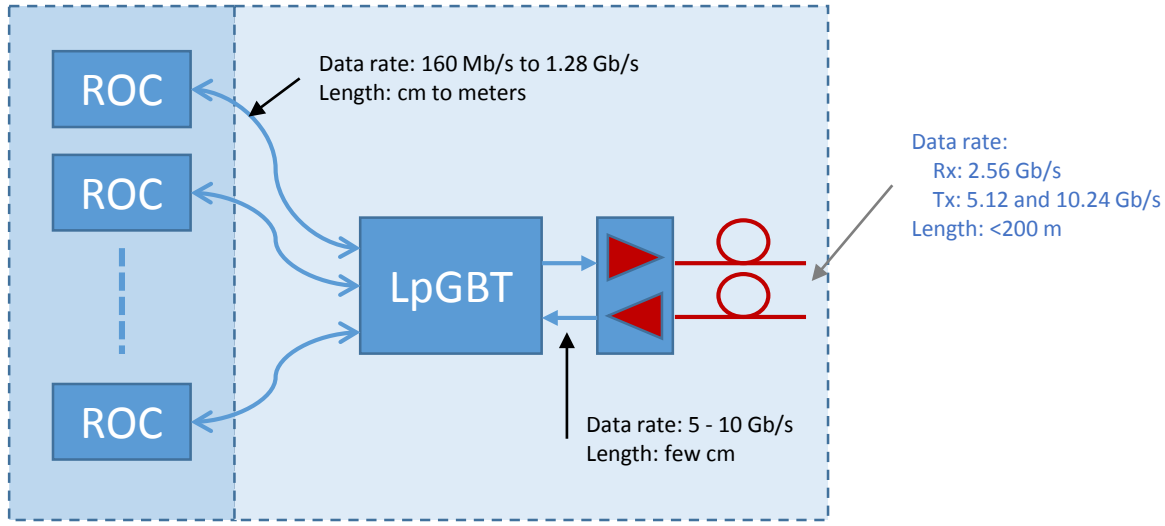

- Power dissipation:
  - Target: ≤ 500 mW @ 5.12 Gb/s
- Small Footprint:
  - Size: 9 mm x 9 mm
  - Fine Pitch: 0.5 mm
  - Pin count: 289 (17 x 17)
- Radiation tolerance:
  - 200 Mrad
  - SEU robust

LpGBT Block Diagram (Simplified)

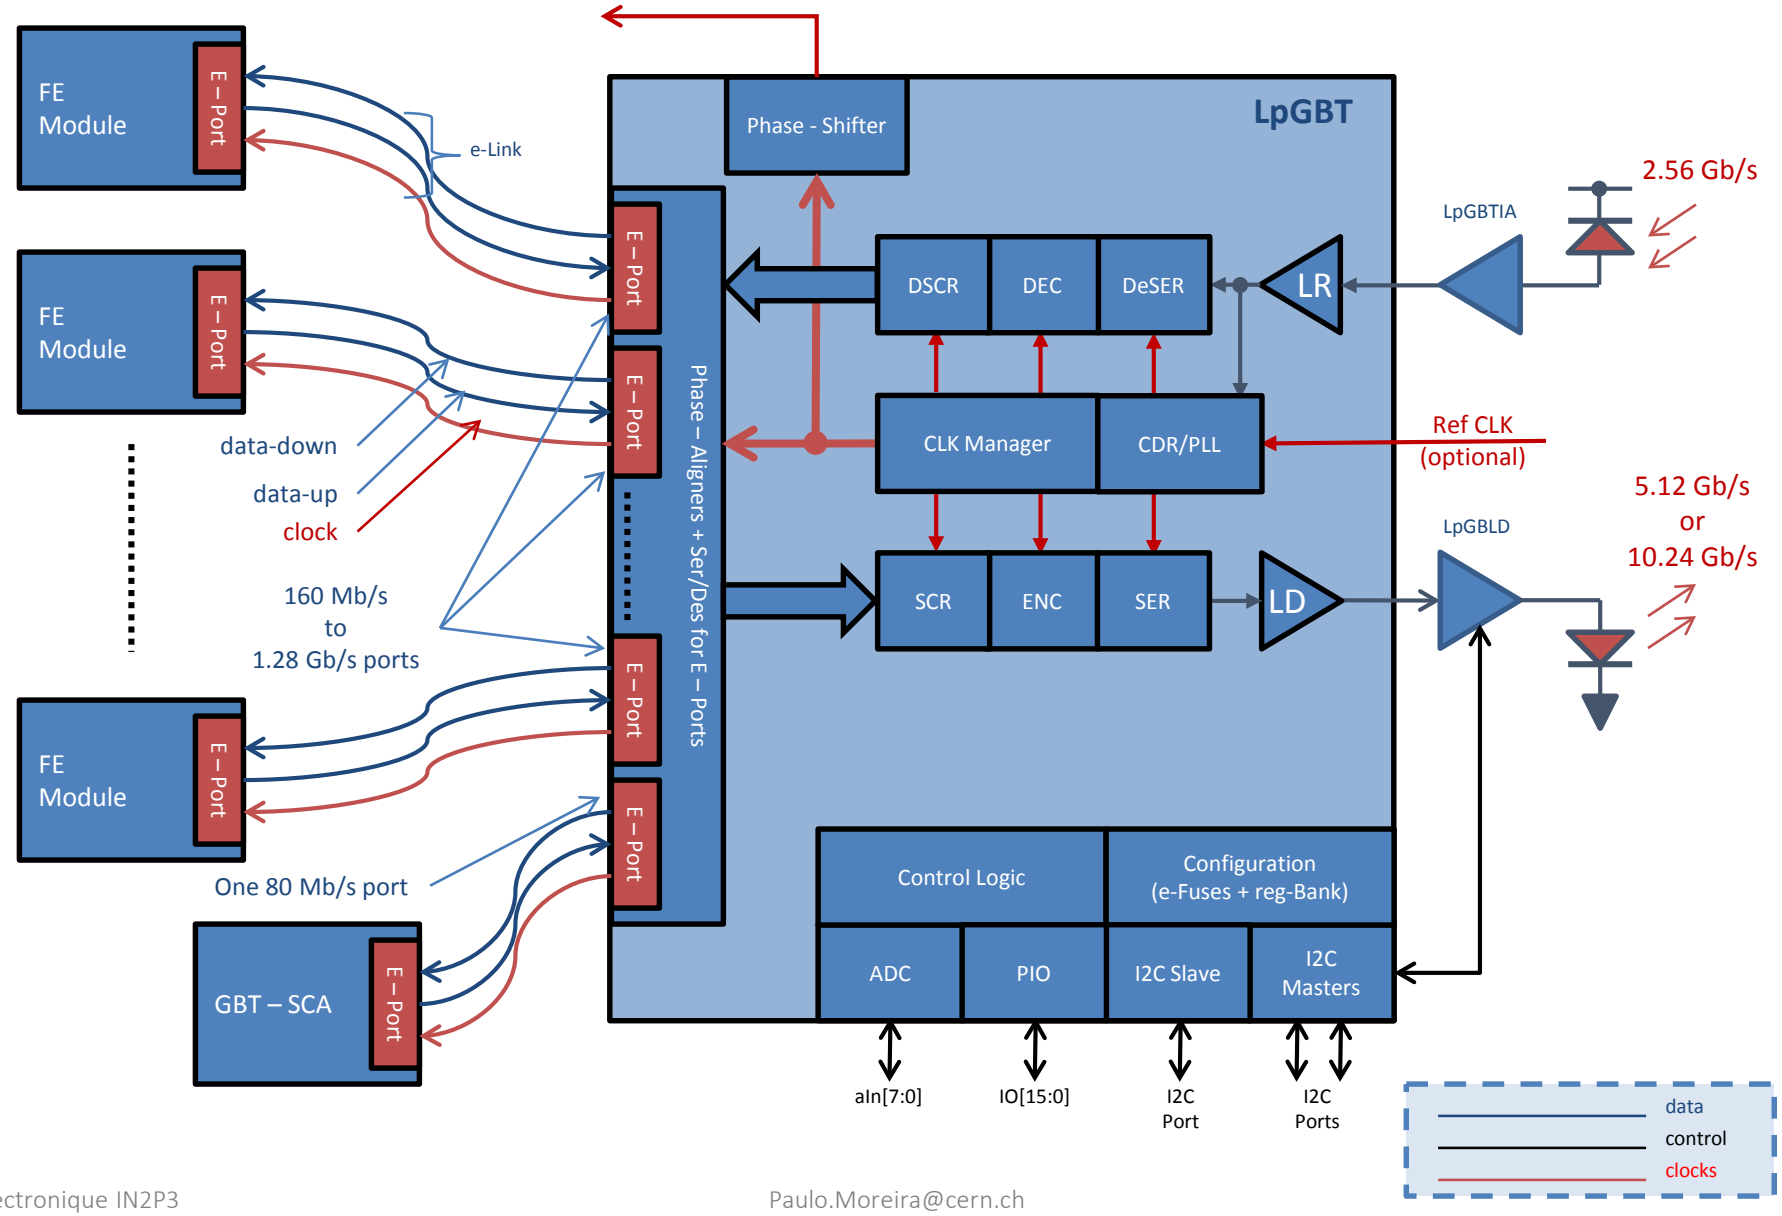

Connecting with the Front-end Modules (ASICs)

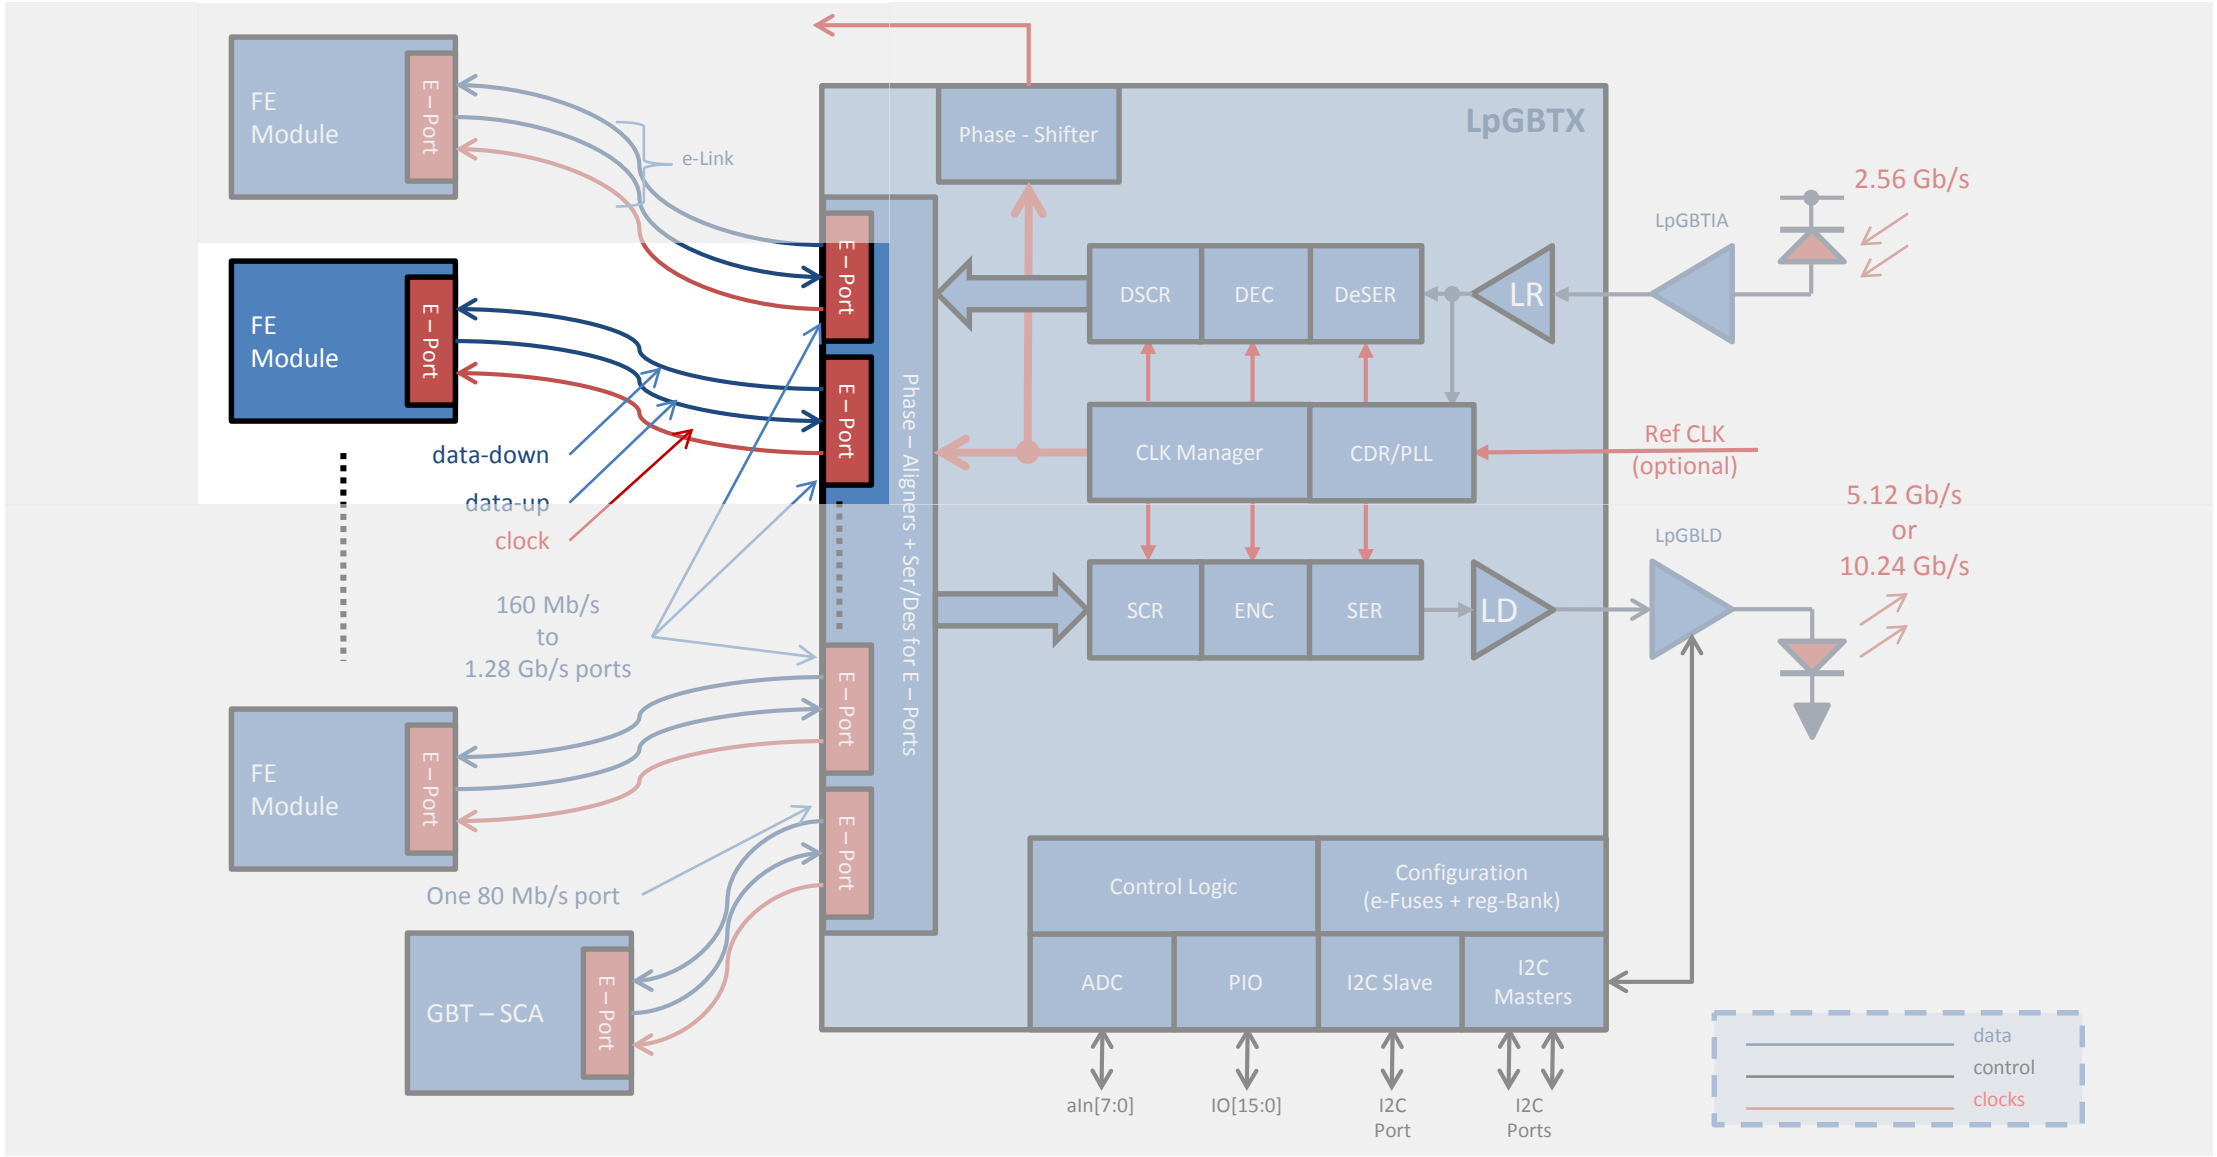

CLPS - eLinks

- Signalling:
  - **CLPS** "CERN Low Power Signalling"
    - [This should] avoid any confusion with LVDS or SLVS
- Main specifications:
  - Link type:
    - Point – to – point
    - Multi – drop transmitter
  - Max data rate:
    - 1.28 Gb/s (NRZ)
  - Max clock frequency:
    - 1.28 GHz
  - Programmable signalling level:
    - 100 mV to 400 mV (single-ended PP amplitude)
    - 200 mV to 800 mV (differential PP amplitude)
  - Common mode voltage:
    - 600 mV (nominal)
  - Load impedance:
    - 100 Ω

Common mode in the middle of the supply (Vdd/2) for best tolerance to ground fluctuations between modules;

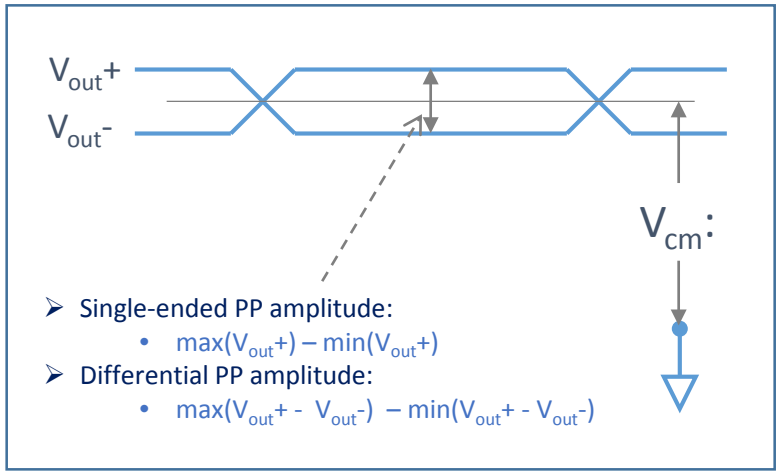

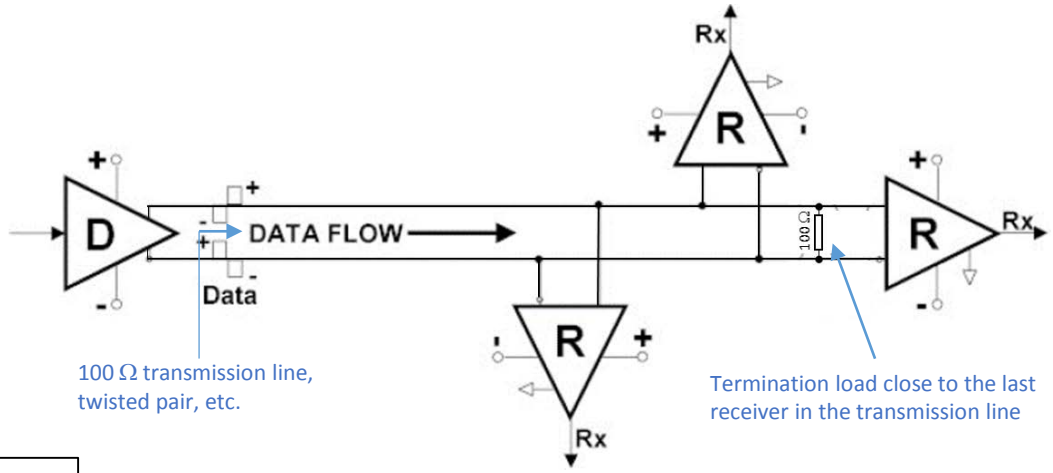

eLink Driver – eTx

- Data rate:
  - Up to 1.28 Gb/s
- Clock frequency:
  - Up to 1.28 GHz
- Driving current:
  - Programmable: 1 to 4 mA in 0.5 mA steps
- Receiving end termination:
  - 100 Ω
- Voltage amplitude in 100 Ω:
  - 100 mV to 400 mV (SE PP amplitude)
  - 200 mV to 800 mV (DIFF PP amplitude)
- Common mode voltage:
  - 600 mV
- Pre-emphasis:
  - Driving current: 1 to 4 mA in 0.5 mA steps
  - Pulse width:
    - Externally timed
    - Self timed: 120 ps to 960 ps in steps of 120 ps
    - Clock timed: T_{bit} / 2

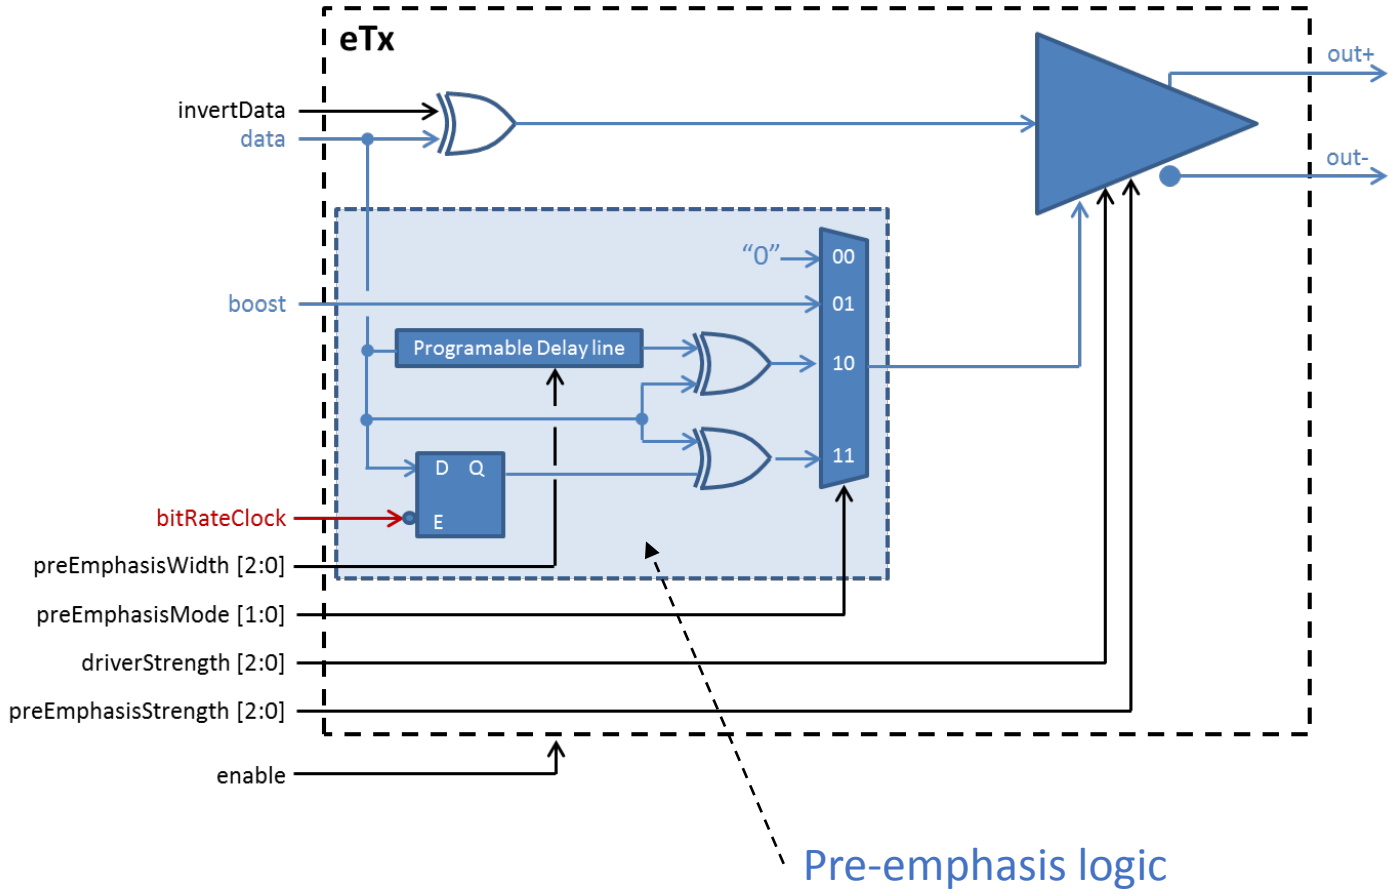

eTx Design

- eTx output stage circuit architecture was driven by radiation tolerance considerations:
  - eTx interfaces with other ASICs so it is important that its performance degrades as little as possible with TID.
- Poly resistors are insensitive to TID
- The circuit relies on having the resistors setting the current and not the transistors

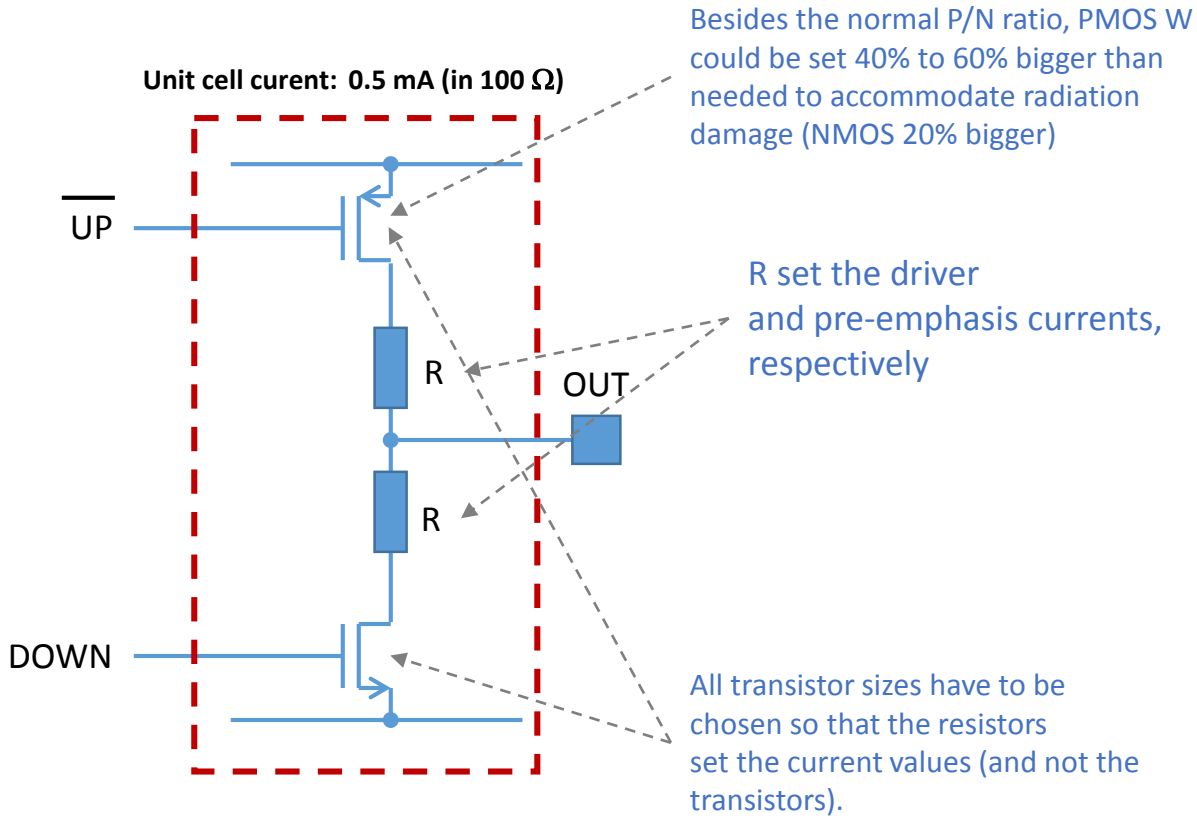

eTx Output Stage is Pseudo Differential

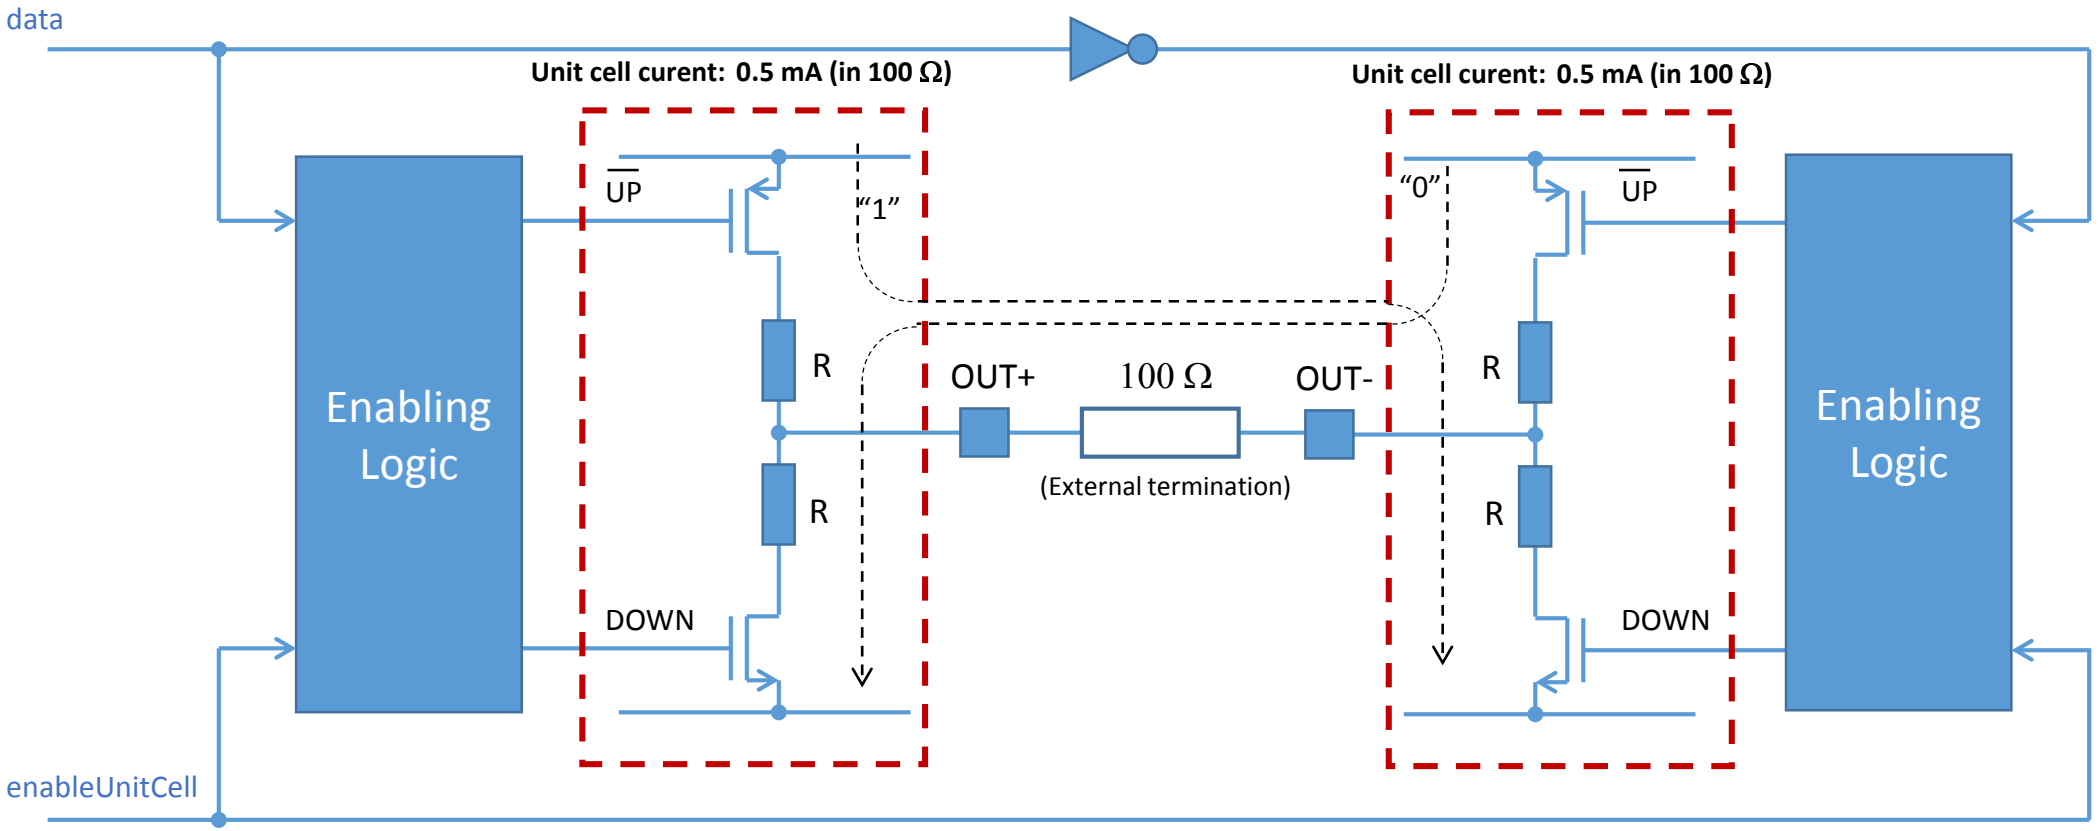

Unit cell can be disabled by keeping: $\overline{UP}$ = "1" and DOWN = "0"

eTx Output Stage is Pseudo Differential

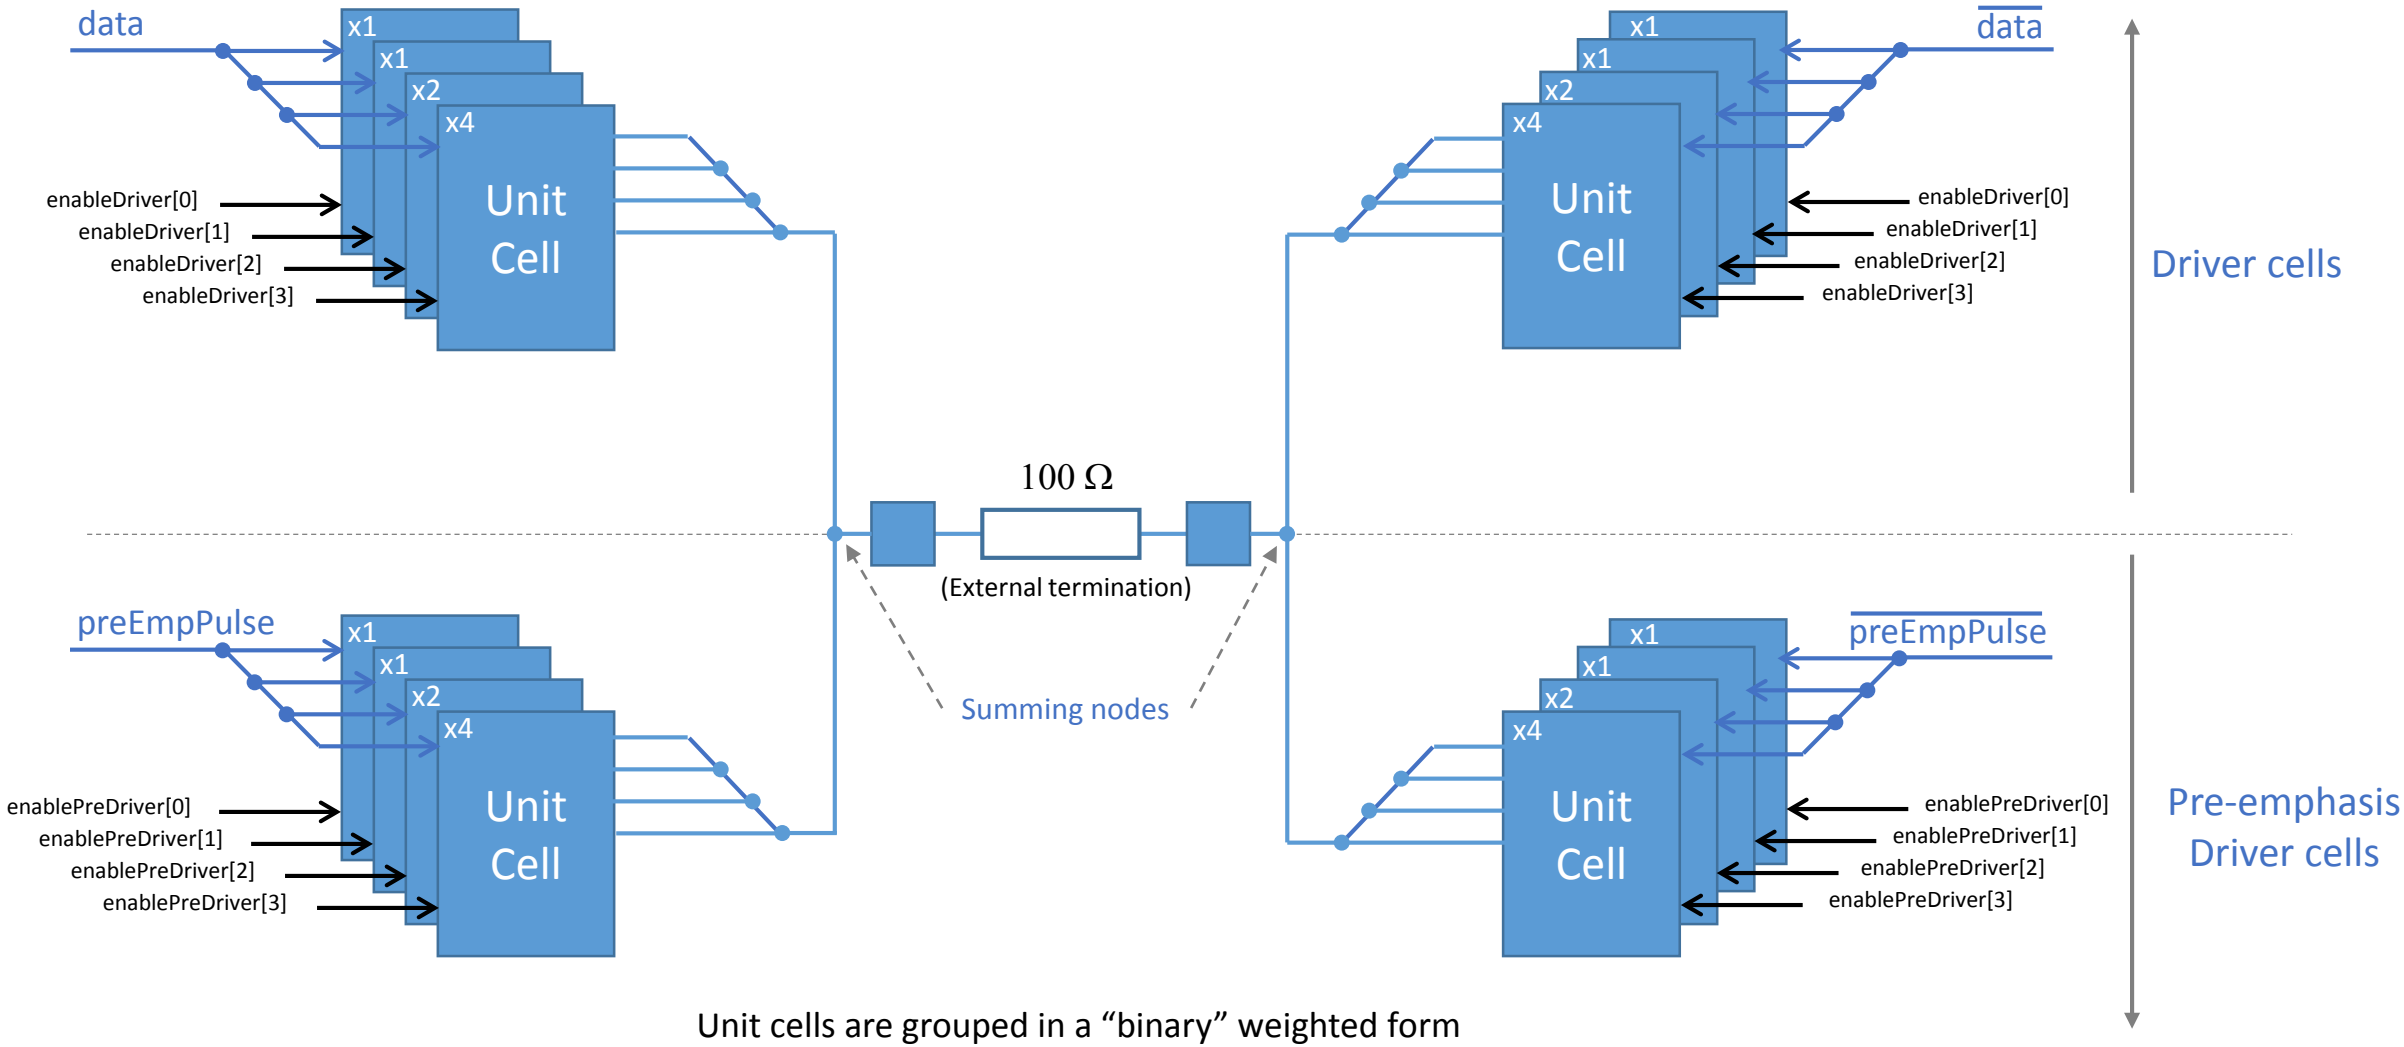

Unit cells are grouped in a “binary” weighted form for current programmability: 1x, 1x, 2x and 4x

eTx – Amplitude control & Power Consumption

Differential output: 500 MHz Clock

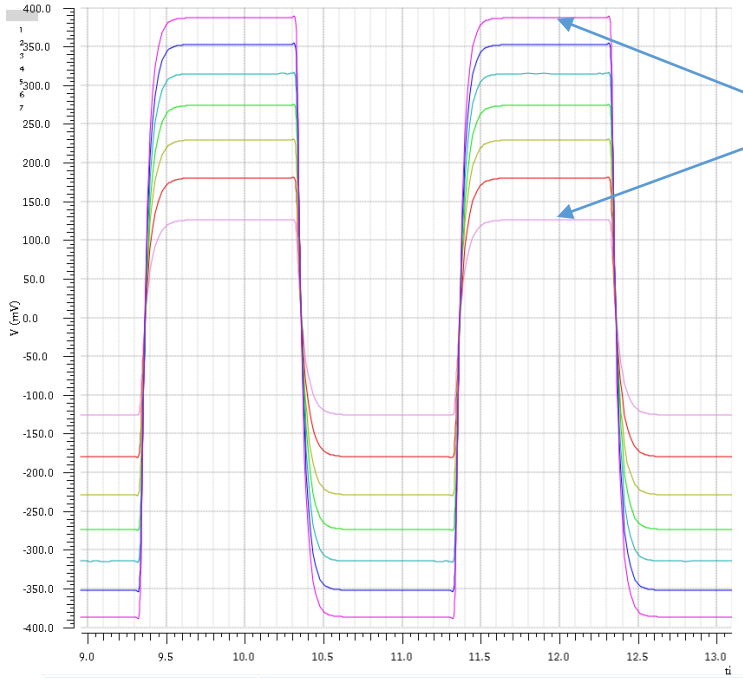

Programmable amplitude

The power consumption depends on the programmed strength (load current) and on the operation frequency. That is also on the dynamic power dissipated in the CMOS circuits that control the output stage!

	Drive strength = 1 (1 mA)		Drive strength = 3 (2 mA)		Drive strength = 7 (4 mA)	
	Ivdd (RMS) [mA]	P [mW]	Ivdd (RMS) [mA]	P [mW]	Ivdd (RMS) [mA]	P [mW]
160 MHz	1.55	1.77	2.51	2.94	3.90	4.41
320 MHz	1.85	2.12	2.88	3.36	4.33	5.10
640 MHz	2.39	2.82	3.54	4.19	5.11	6.07
1.28 GHz	3.48	4.16	4.82	5.76	6.58	7.87

What is Pre-Emphasis (1/4)

- To transmit NRZ data with little ISI the transmission channel must have a bandwidth of at least (rule of thumb):
  - BW = 0.7 × Bit Rate
- For example, for transmission at 5 Gb/s the bandwidth required is:
  - BW = 0.7 × 5 Gb/s = 3.5 GHz
- For illustration purposes lets suppose that:
  - The driver is modelled by an ideal current source
  - And the circuit driven is modelled by a RC network with 3.5 GHz bandwidth:
    - R = 50 Ω
    - C = 0.91 pF
- Simulated eye-diagram:
  - There is very little ISI:
  - The eye is well opened vertically and horizontally
  - The jitter is very low

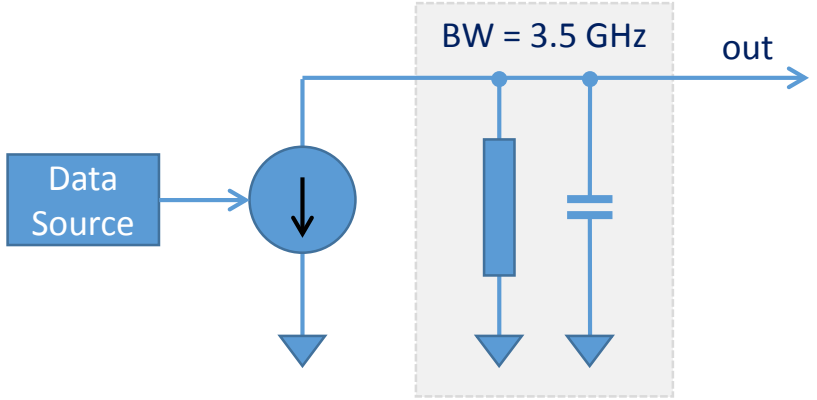

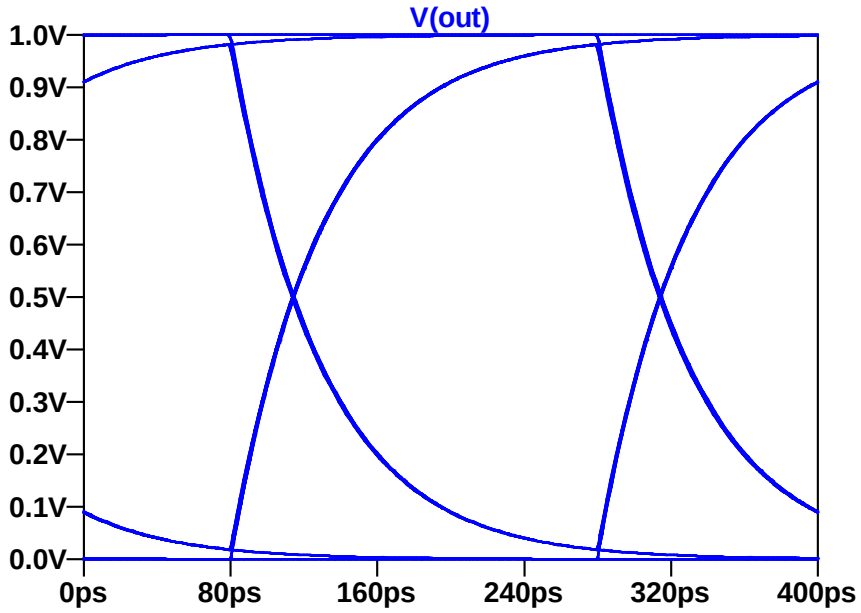

What is Pre-Emphasis (2/4)

- Lets suppose now that the bandwidth of the circuit being driven is four times lower:
  - BW = 3.5 GHz / 4 = 795 MHz
  - R = 50 Ω
  - C = 3.64 pF
- Simulated eye-diagram
  - There are significant amounts of ISI:
  - A “bit” extends for much longer than a bit period
  - The eye-diagram is almost closed vertically and horizontally
  - Jitter is high
- The BER would be “prohibitive” for such a system!

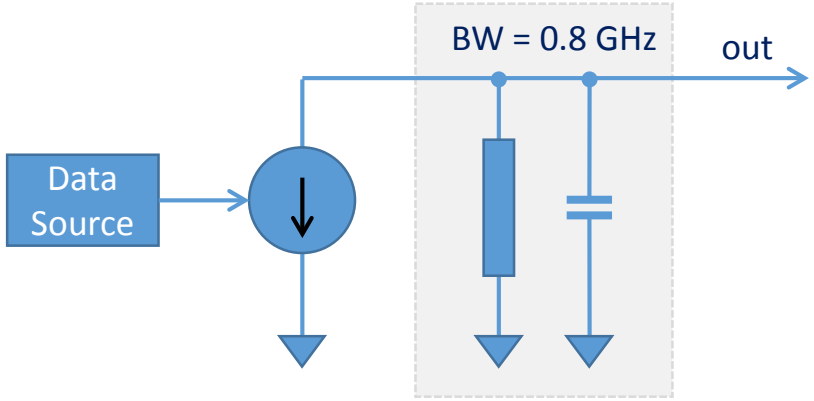

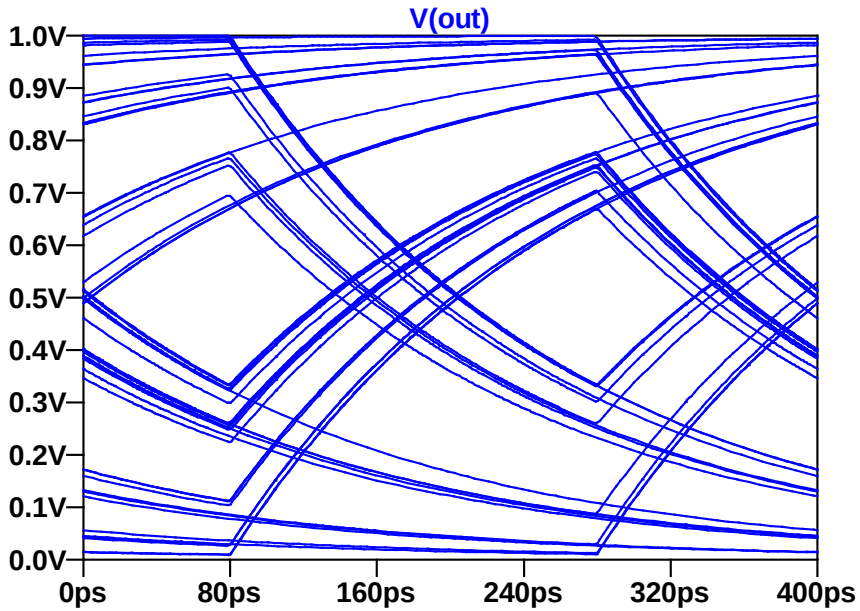

What is Pre-Emphasis (3/4)

- The problem with the low bandwidth is that:
  - For fast successions of “0s”and “1s” the signal has no time to reach the final value before a new “0” or a new “1” is transmitted.
- In an RC type of circuit this can be “easily” overcome:
  - At every “0”-to-“1” or “1”-to-“0” transition drive the circuit to a higher (or lower) voltage than the final one:
    - In our case this is accomplished by using a larger current than the final one
  - Once the voltage reaches the desired amplitude switch to the nominal current.
- In practice, no level crossing detection is made:
  - Immediately after a transition and for a “short” time a current pulse is added to the “nominal” modulation current
  - The pre-emphasis pulse duration is adjusted to open the eye-diagram

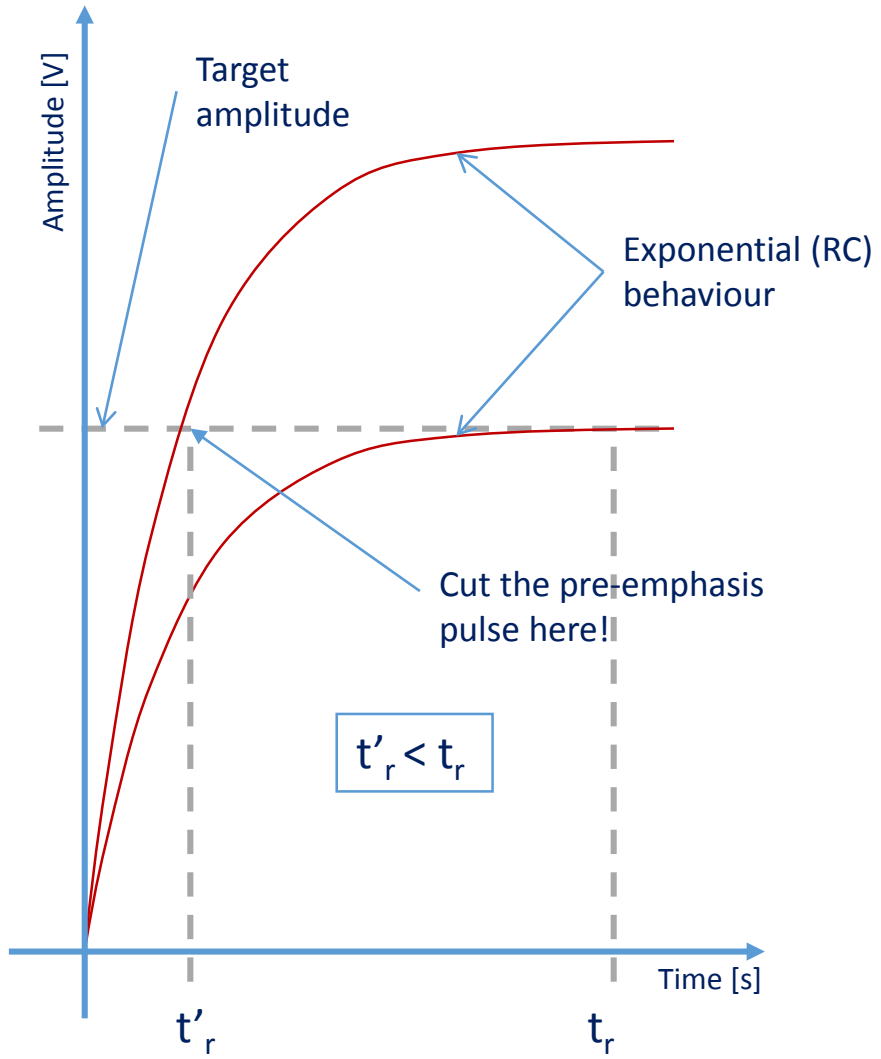

What is Pre-Emphasis (4/4)

Data Rate: 3.5 GHz

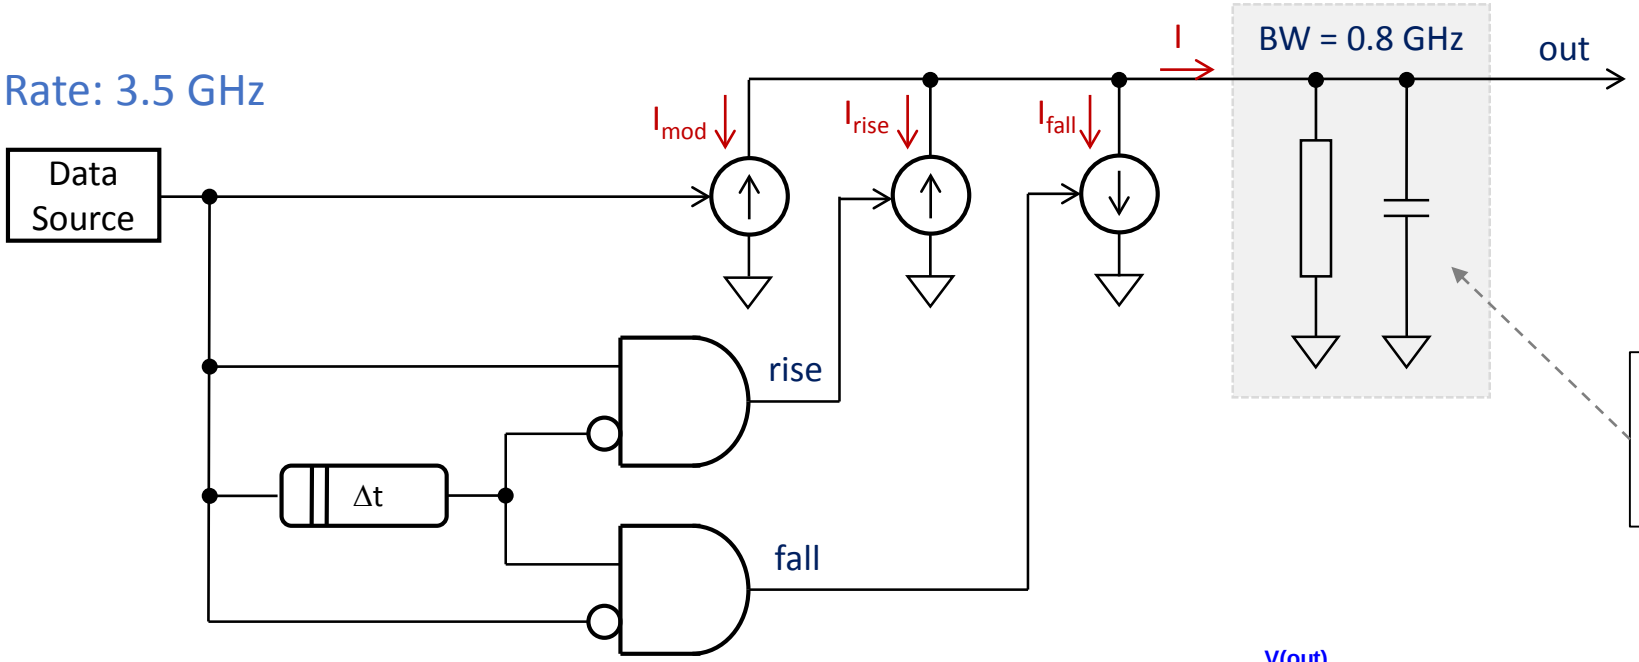

Remember the rule of thumb:
BW = 70% Bit Rate → 2.45 GHz
This network only has 32% of the required bandwidth!

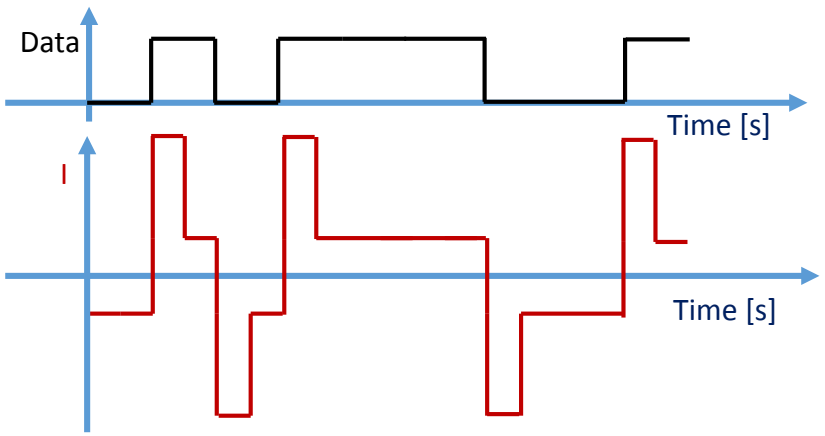

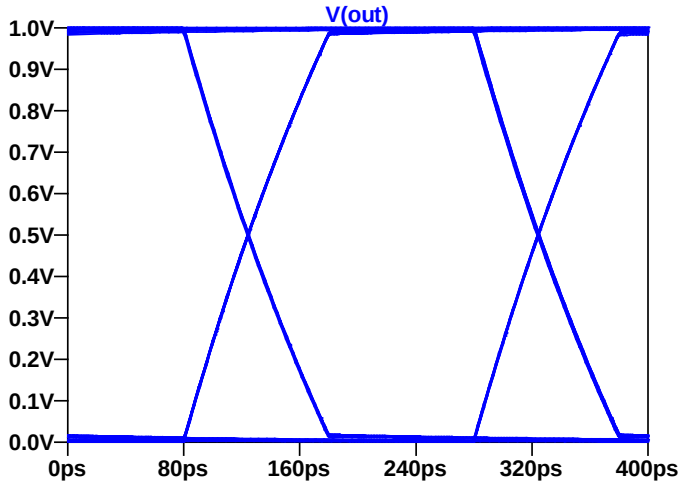

$\Delta t = 0.5 \times T$

I_{mod} = "1"

$I_{rise} = 1.35 \times I_{mod}$

$I_{fall} = 1.35 \times I_{mod}$

$I = I_{mod} + I_{rise} + I_{fall}$

In the Frequency Domain

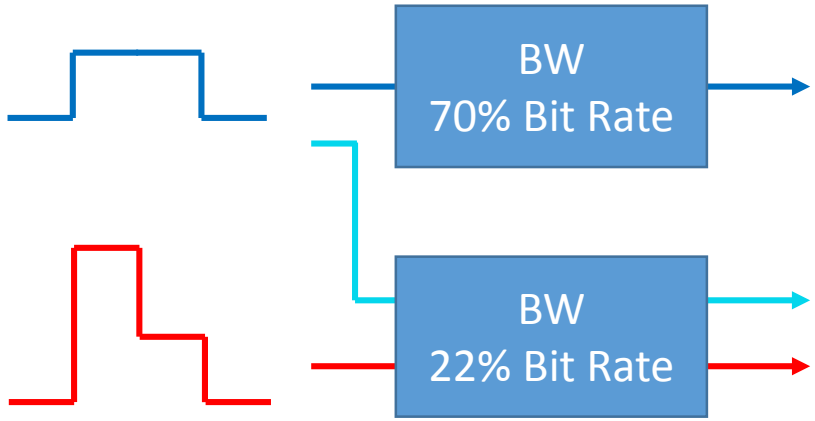

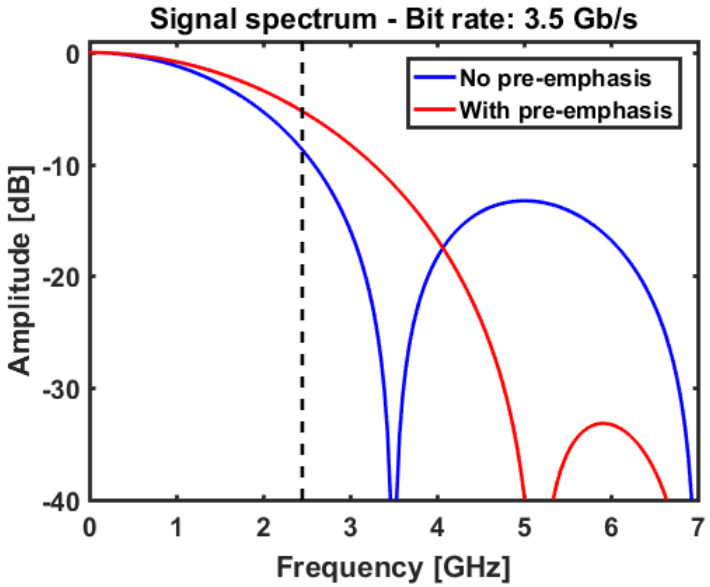

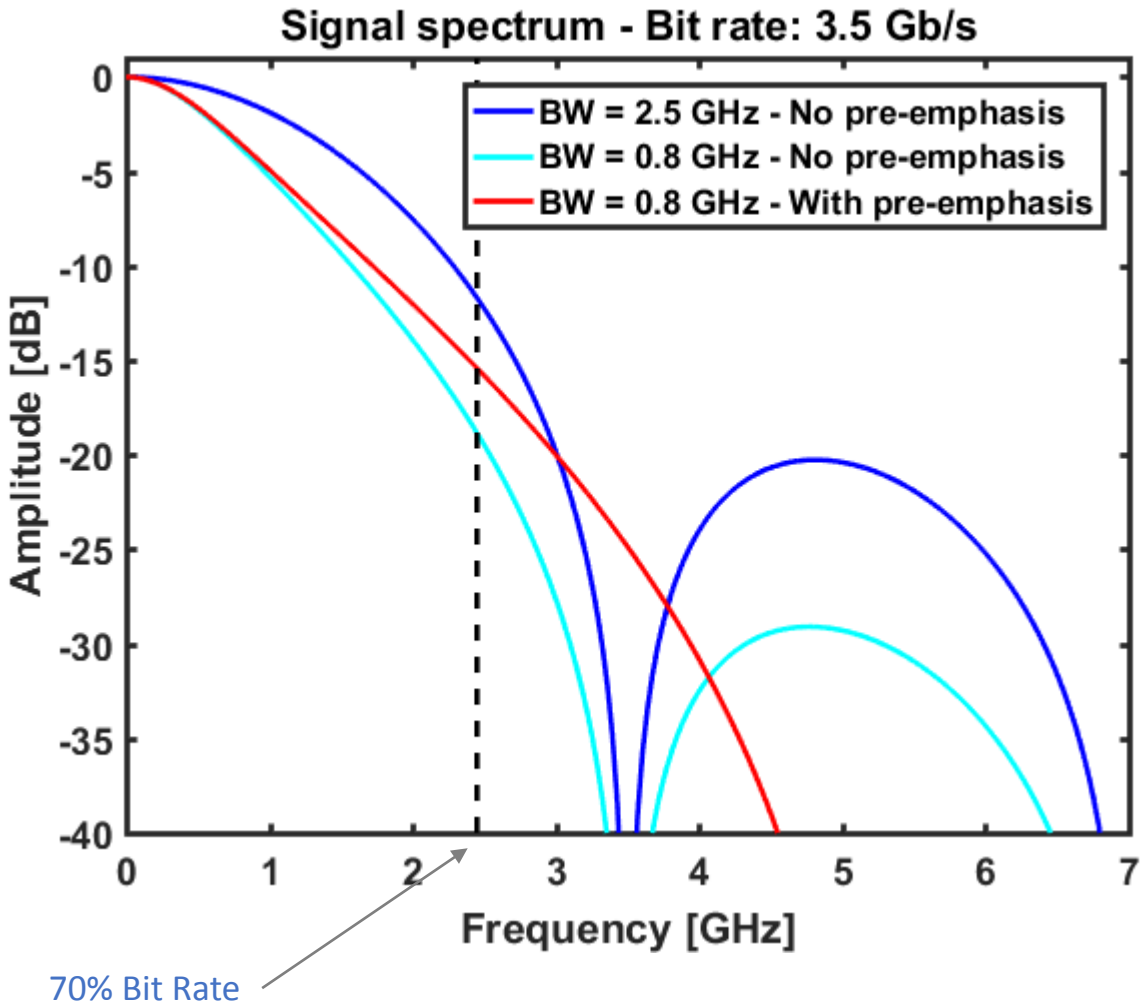

eTx Pre-Emphasis Functions

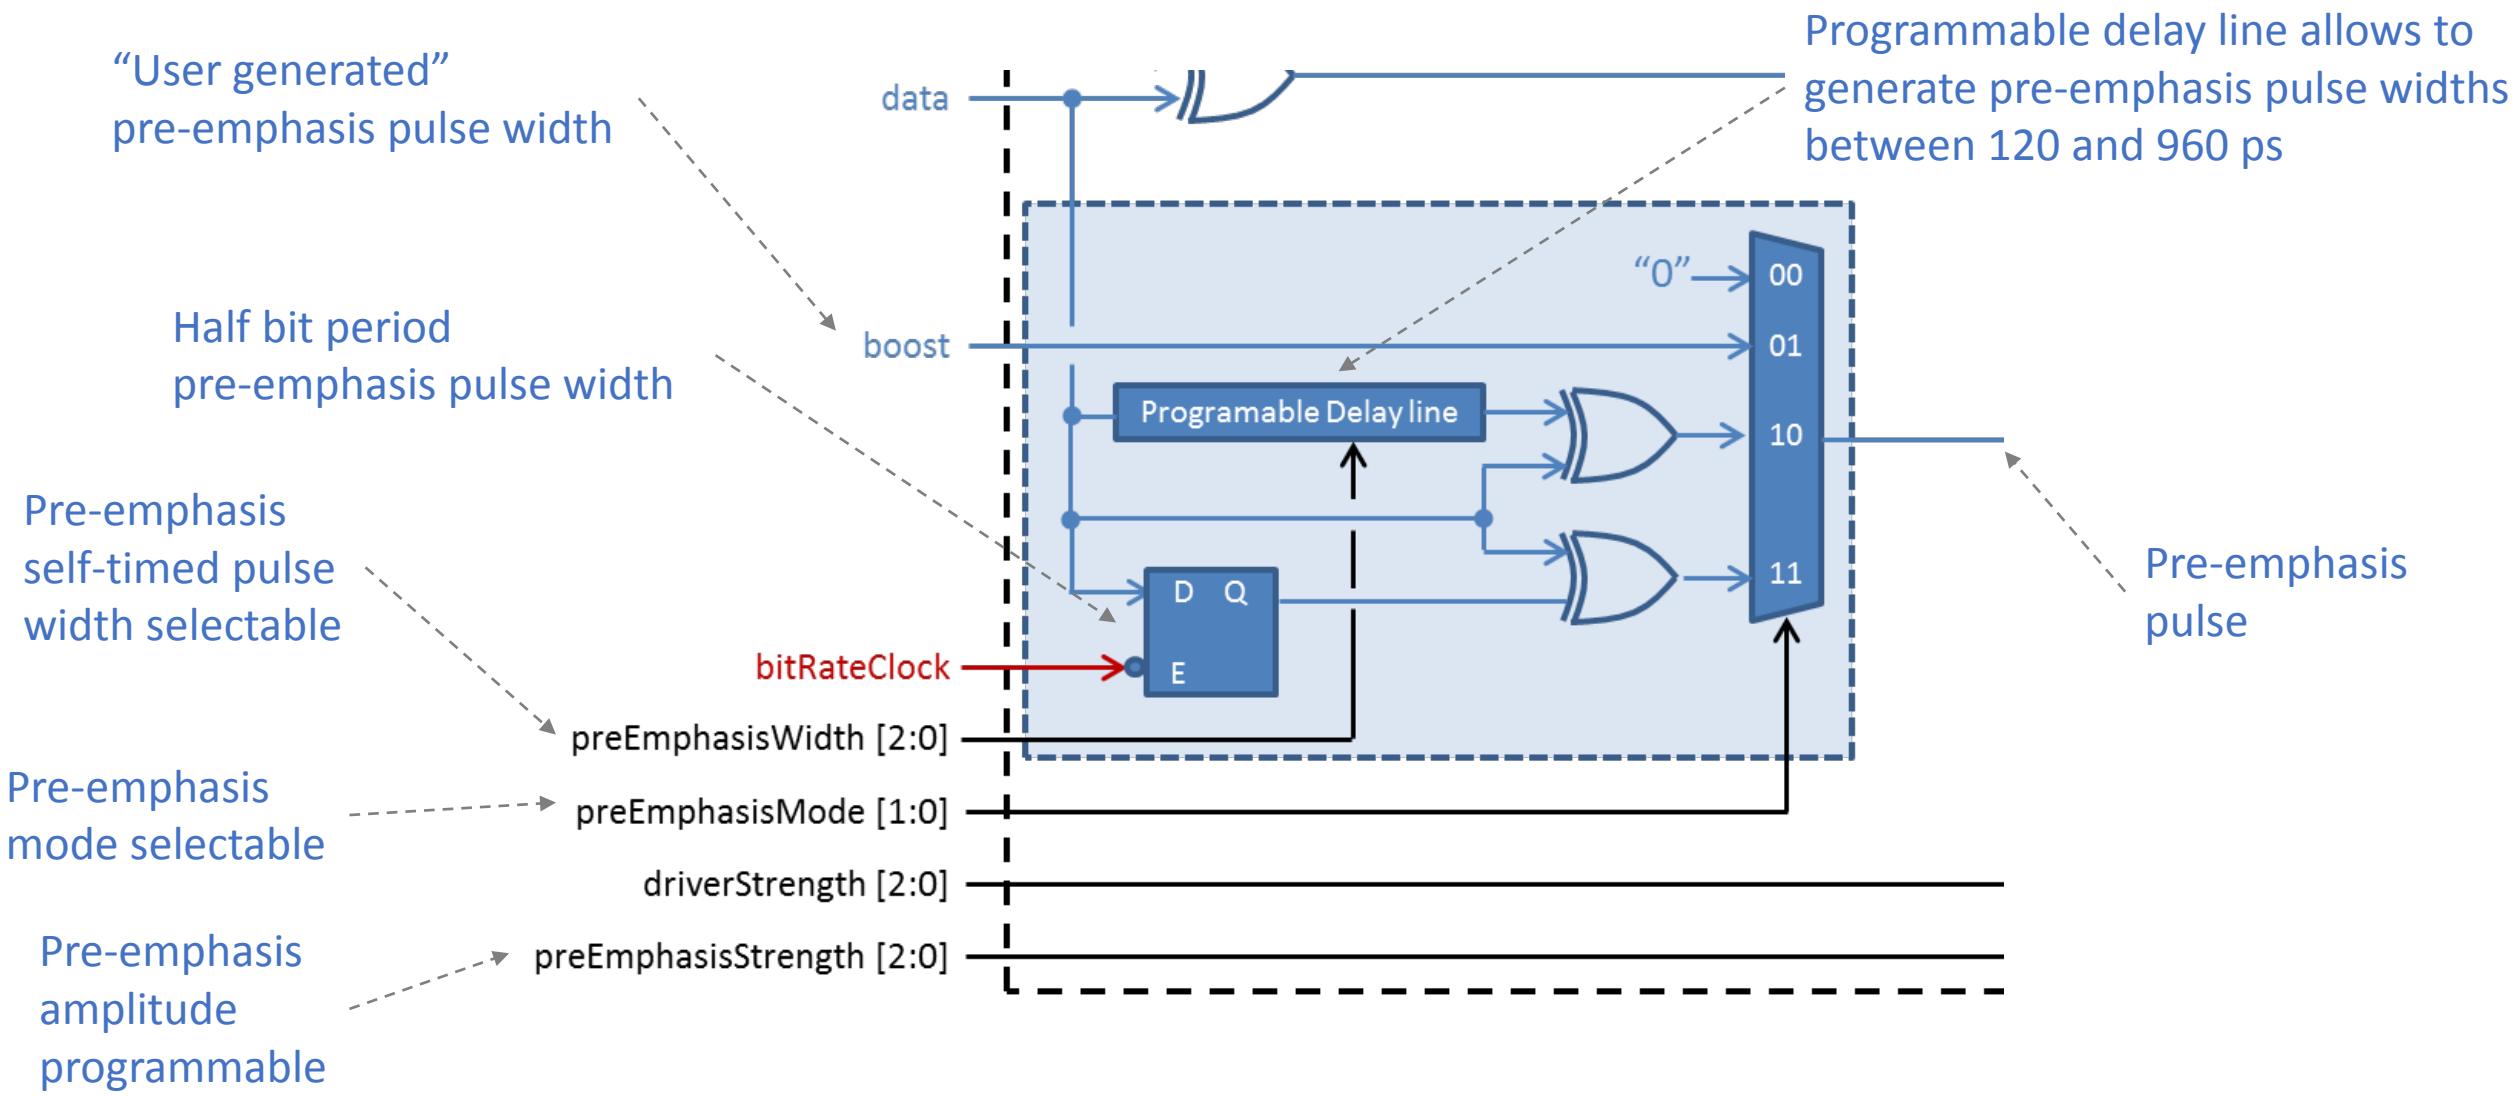

eTx – Pre-Emphasis

500 MHz Clock

Output with low capacitive loading

With Pre-Emphasis
- C_L = 1 pF
- Driver Strength = 2 mA
- Pre-Emphasis Strength = 1 mA
- Pulse width = T/2

Output with high capacitive loading

With Pre-Emphasis
- C_L = 10 pF
- Driver Strength = 2 mA
- Pre-Emphasis Strength = 0, 1, 1,5 and 2 mA
- Pulse width = T/2

Single ended

Differential

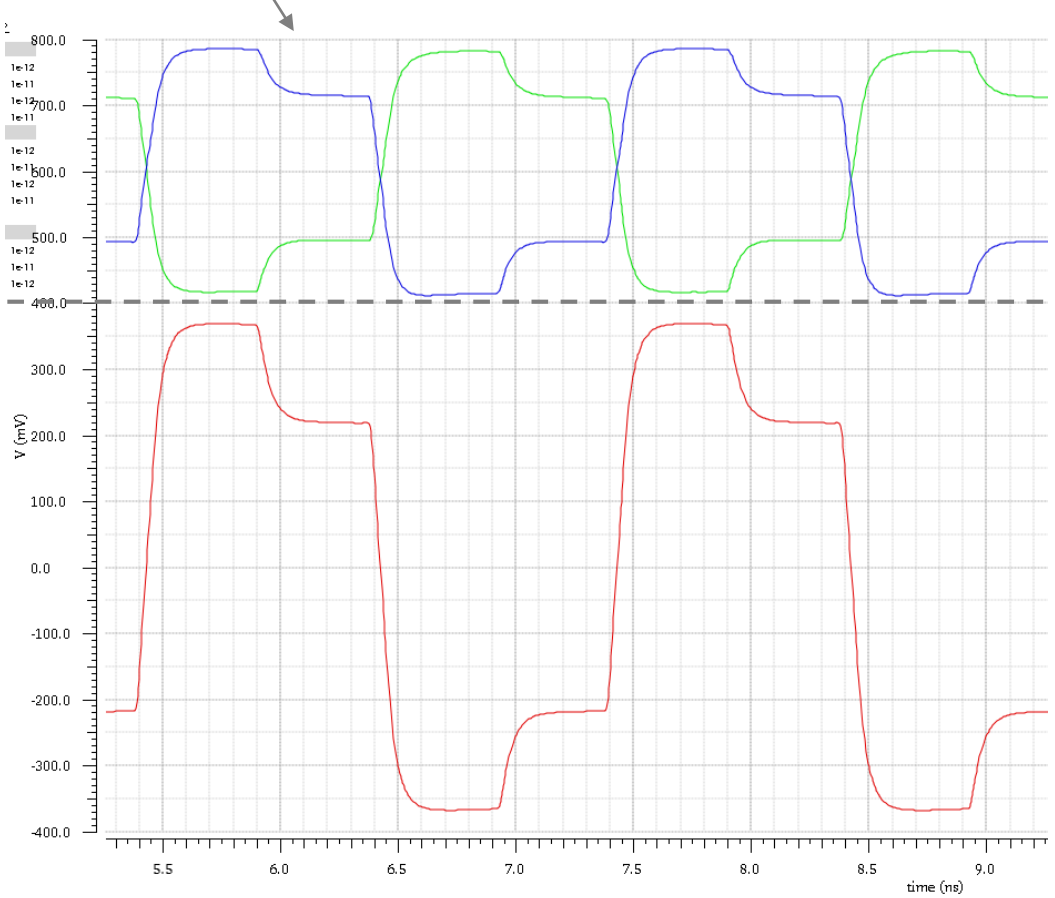

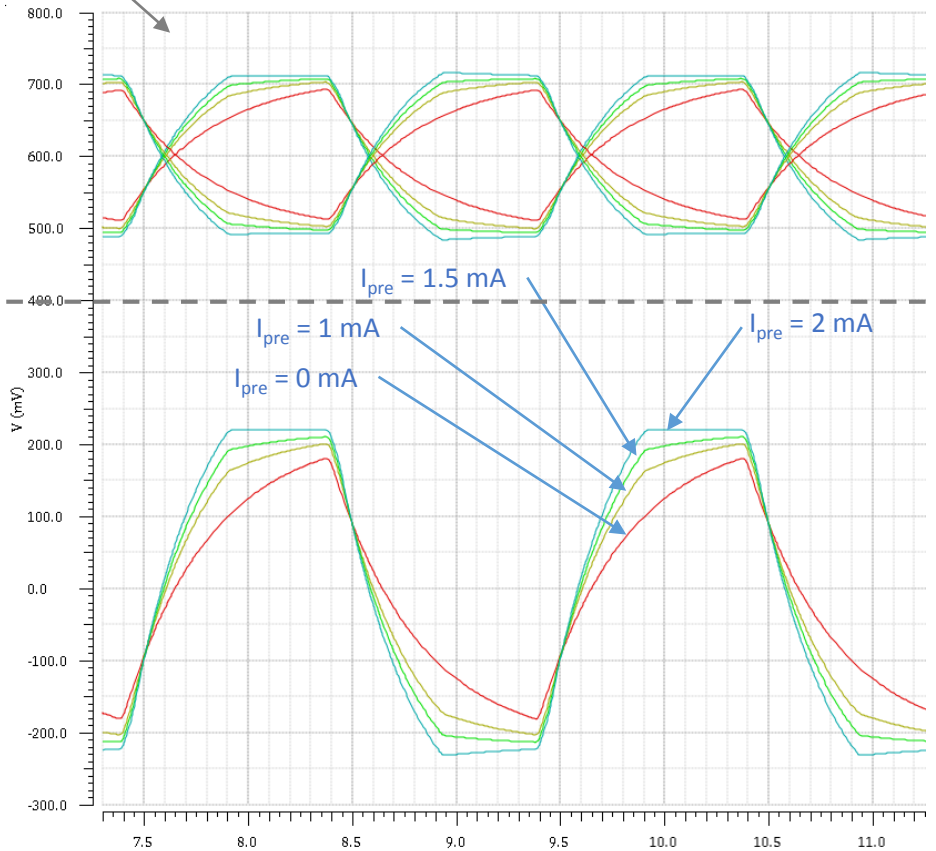

eTx – Pre-Emphasis

Data: PRBS 7
Data rate: 2.56 Gb/s
C_L = 2 x 5 pF

No Pre-Emphasis
- Driver Strength = 2 mA
- Pre-Emphasis Strength = 0 mA
- Pulse width = n.a.

With Pre-Emphasis
- Driver Strength = 2 mA
- Pre-Emphasis Strength = 2 mA
- Pulse width = T/2

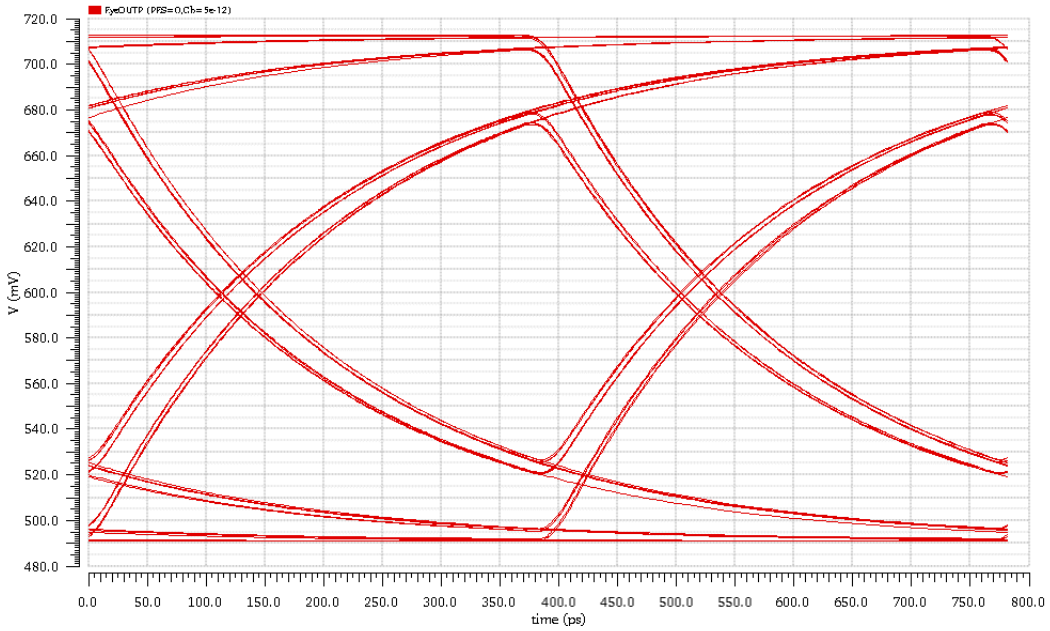

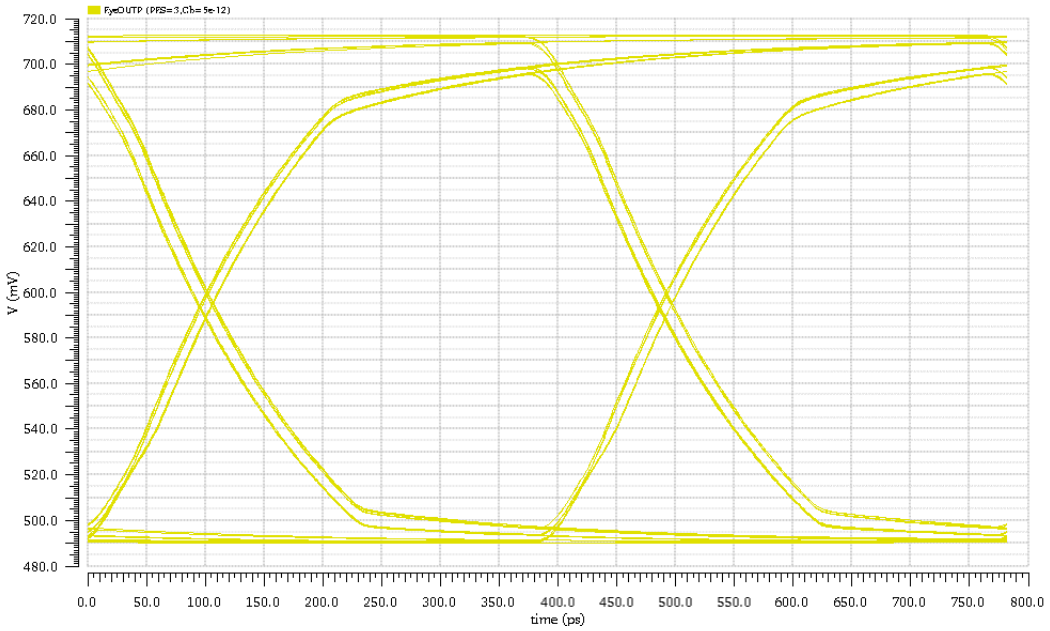

More on Practical [Annoying] Details of Designing for TID Tolerance!

Some Practical Aspects of TID Robustness

- ELT gates use non-minimal device sizes:
  - NMOS: W_n = 1.42 um
  - PMOS: W_p = 3.74 um
  - Basically dictated by the geometry constraints and balancing between NMOS and PMOS current driving
- These “large” devices will introduce noise in V_{dd} and V_{ss} when switching:
  - They drive relatively important loading capacitances (other ELT gates)
  - During the gate signal transition time they represent a low impedance between the supply and ground
- How much noise goes into the supply and ground?

(**ELT** = **E**nclosed **L**ayout **T**ransistor)

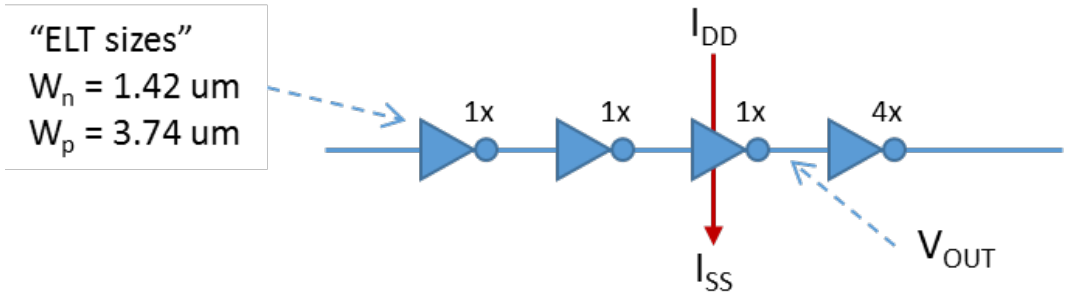

- Test case:
  - Take the “most innocent gate” (minimum size ELT inverter – LVT)
  - Load it with other gates (a reasonable number is 4)
  - Drive it with a 1.28 GHz clock (that is the maximum specified frequency for the eTx)
  - Measure I_{dd} and I_{ss}

Test Case Simulation

- Simulation conditions:
  - C0 (TT, V_{dd} = 1.2V, T = 27C)
  - C4 (FF, V_{dd} = 1.08V, T = 100C)
  - C16 (SS, V_{dd} = 1.08V, T = 100C)
- Main observations:
  - The inverter is [easily] 1.28 GHz capable [all corners]
  - Large peak currents go into the VSS and VDD power rails!
    - These current peaks can reach: 1.3 mA for the fastest PVT conditions!
  - This is for a simple inverter!
    - What are the consequences for a complex circuit or full scale ASIC?

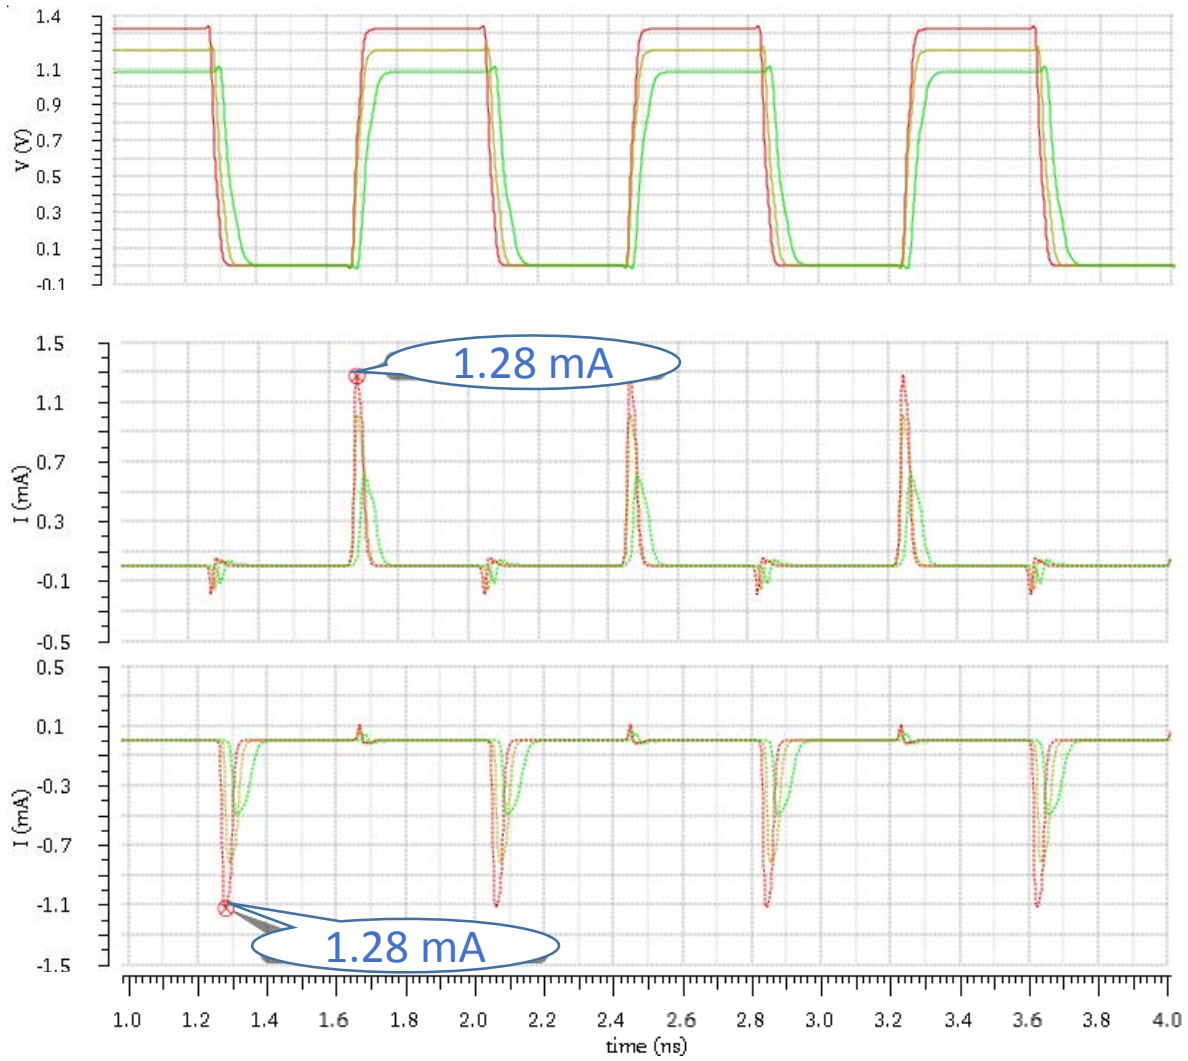

eTx Driver Supply Noise!

- The logic driving the output stage exhibits large current peaks!
  - Up to 14.4 mA
- In the LpGBT there might be:
  - 17 Data Drivers
  - 33 Clock Drivers
  - All working synchronously
- This could represent up to 720 mA of [noise] current being injected in the power supply and ground!
- This noise has to be mitigated!

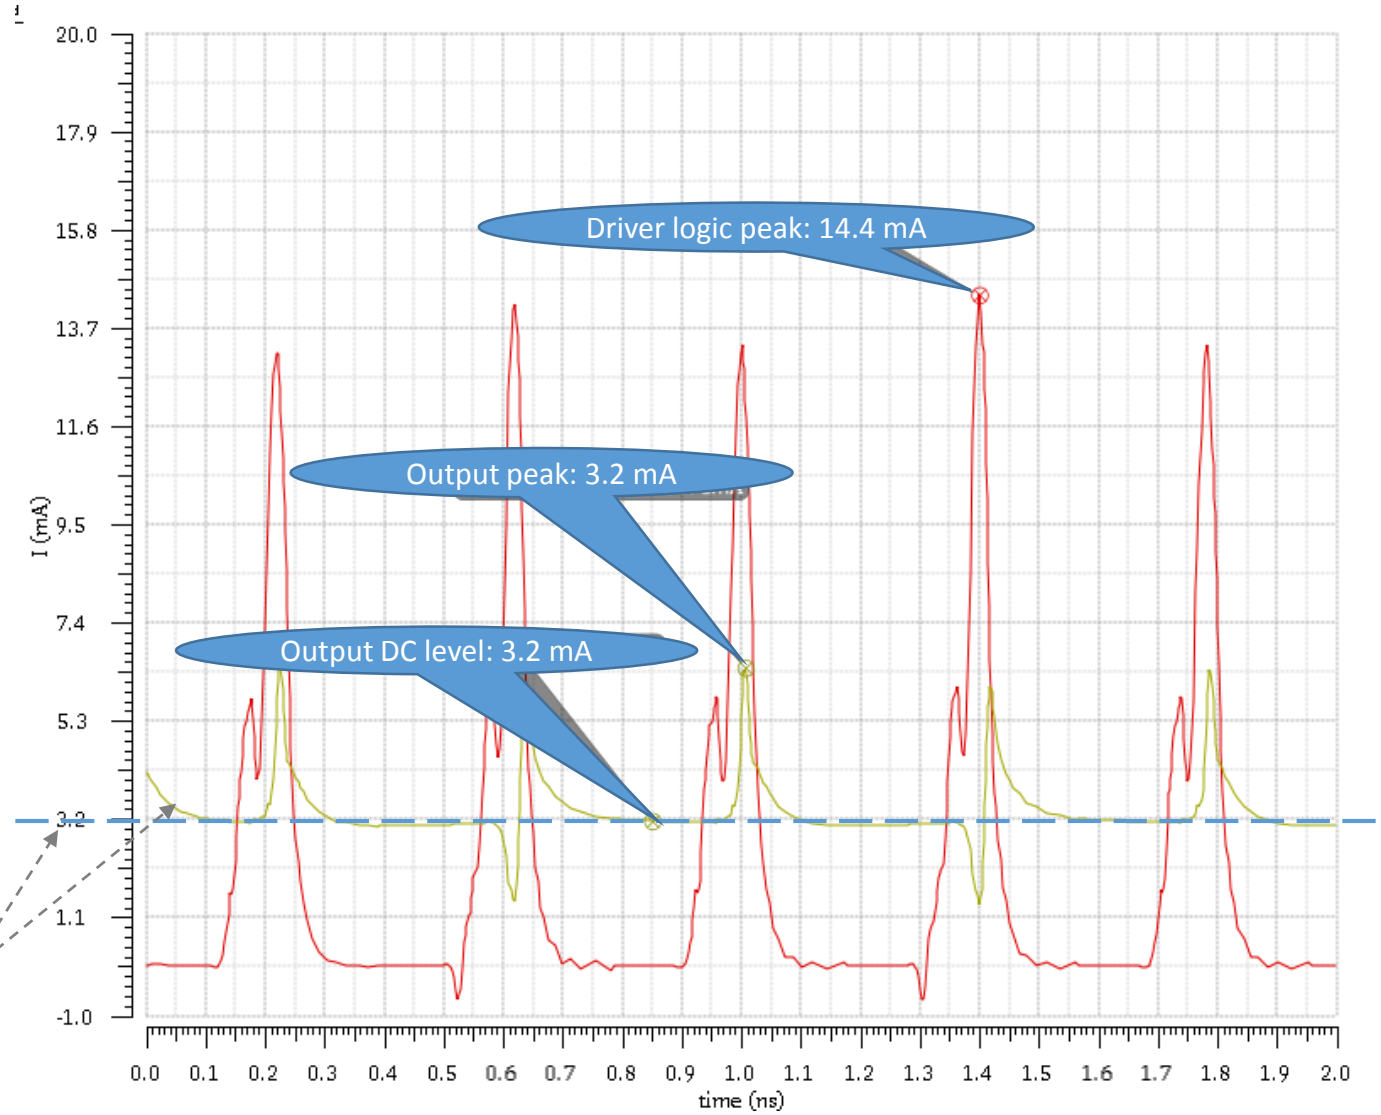

Current in the output stage is “constant” except when the output switches!

eTx Driver Supply Noise!

- Solution:
  - Decrease the peak amplitude by spreading the current pulse in time;
- How
  - RC filter the IO supply and ground locally in the I/O cell
- For this to be effective the I/O cell must be enclosed in a deep well:
  - Otherwise the substrate resistance would “short-circuit” the ground filtering
- Values adopted:
  - R = 10 Ω, C = 10 pF

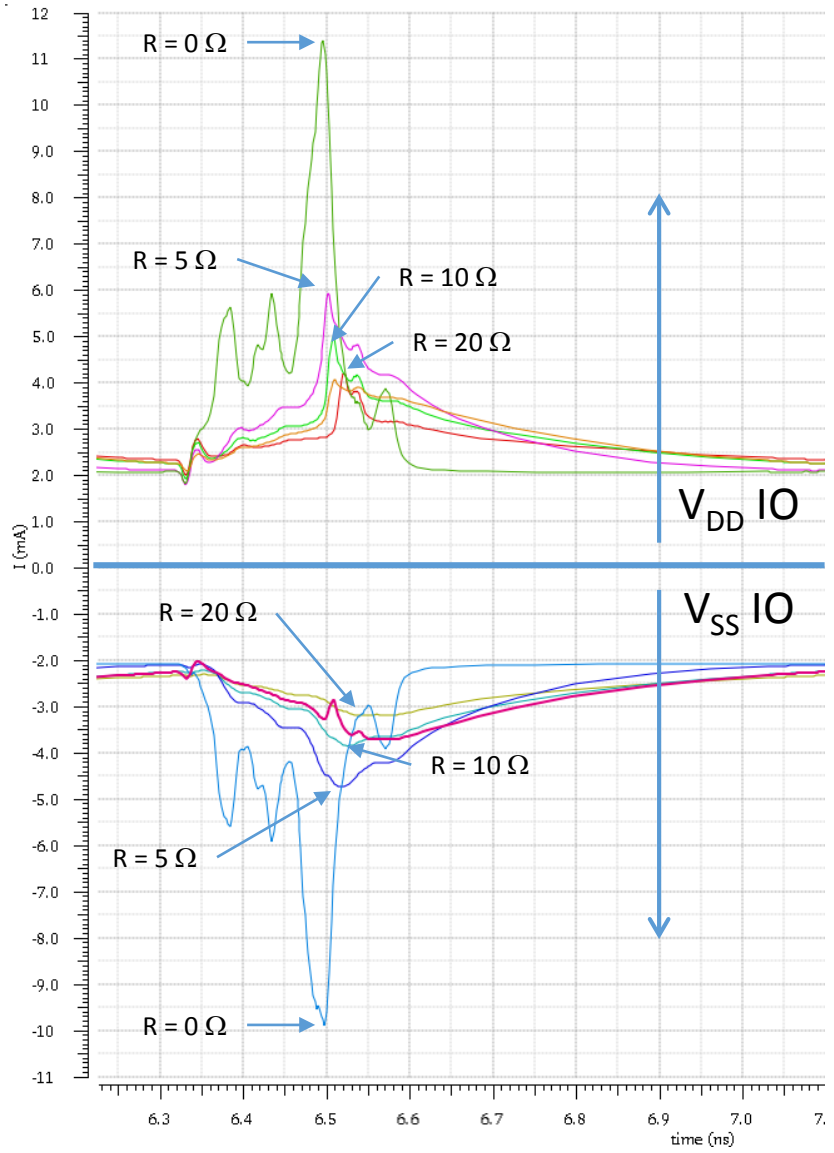

Scheme Adopted for the I/O Power

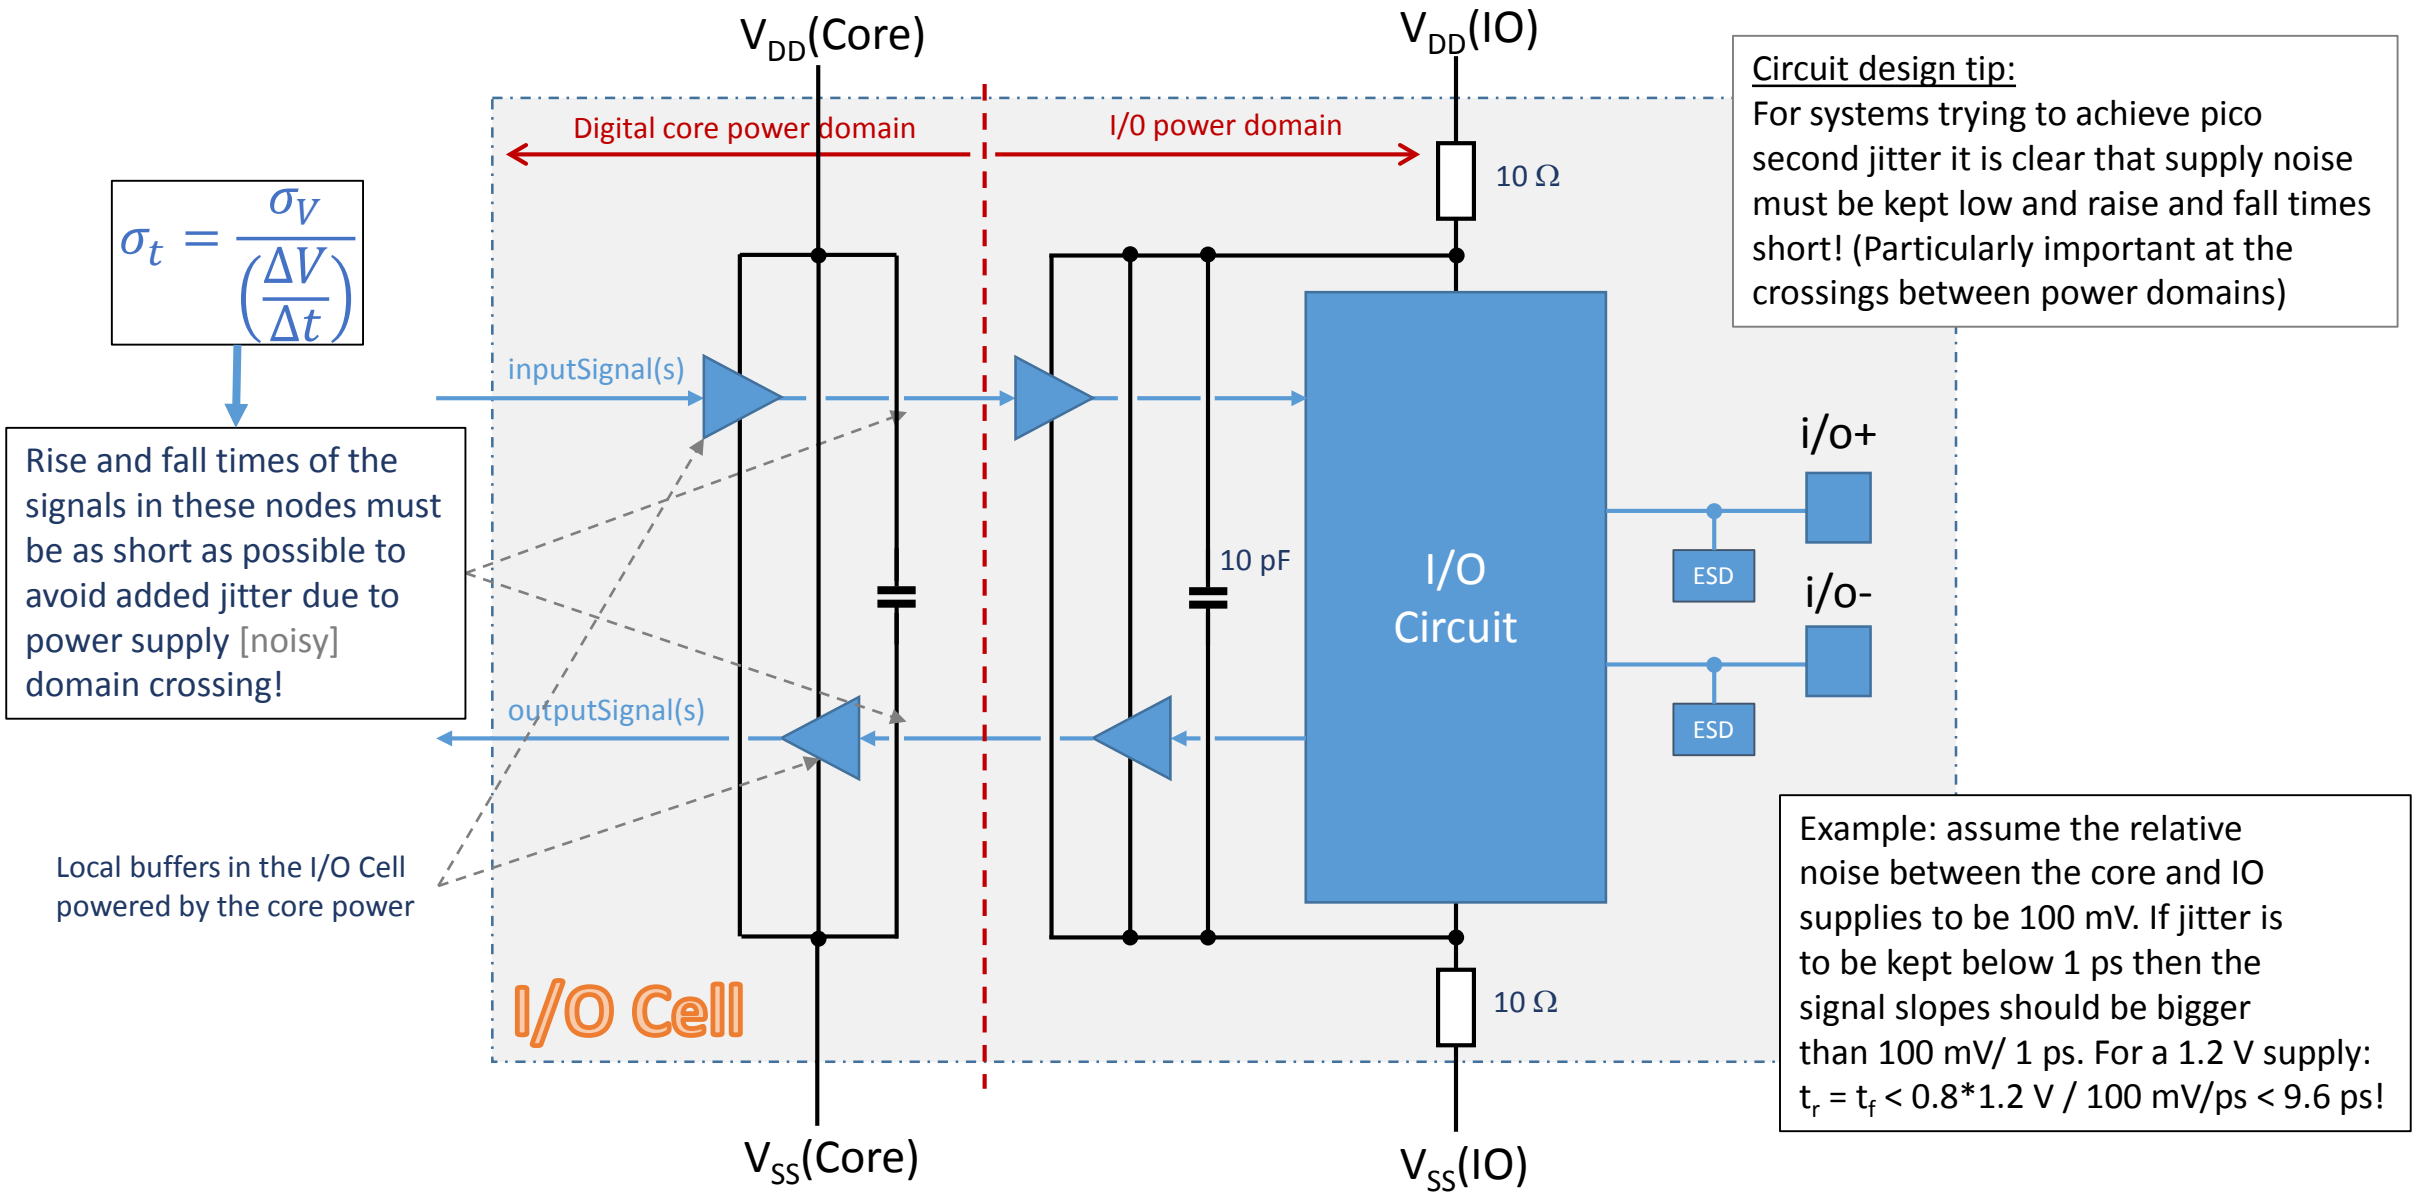

Receiving data from the Front-end Modules (ASICs)

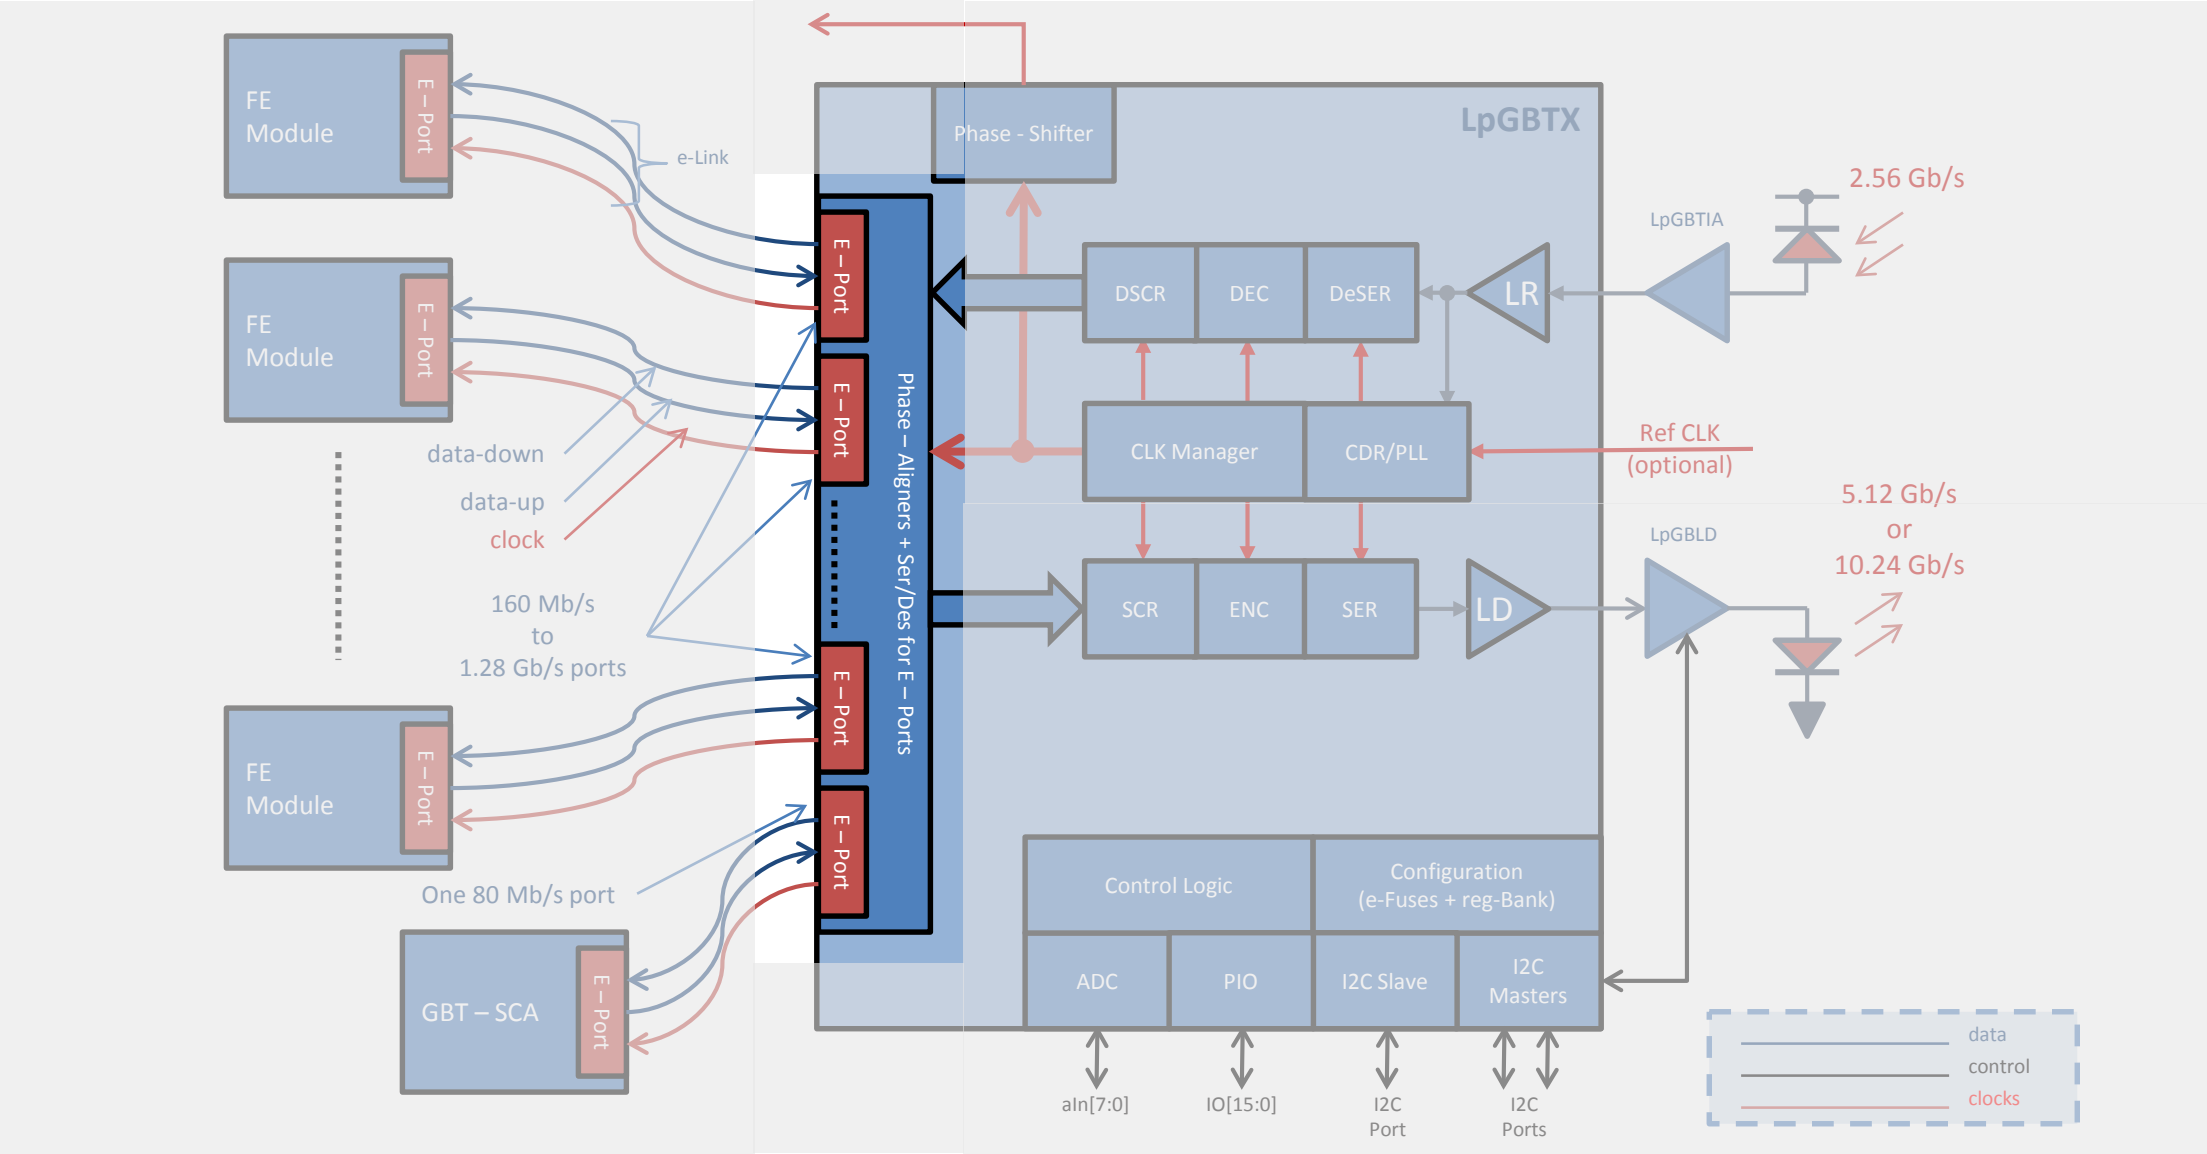

Optimum Sampling

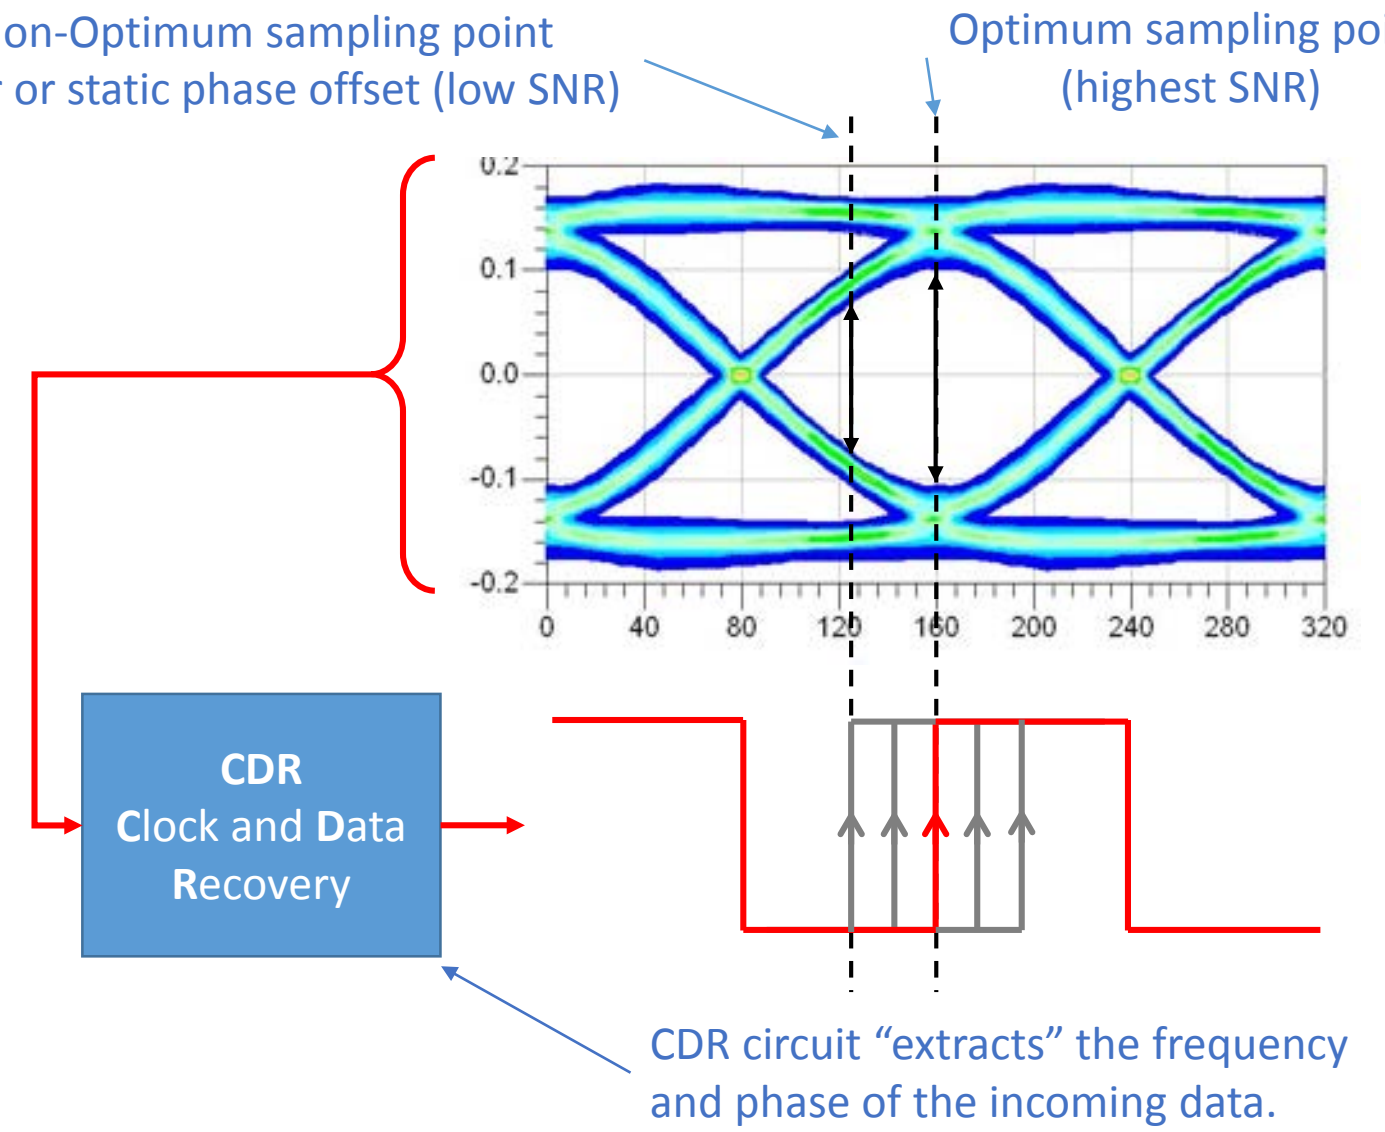

- The LpGBT is the clock source to the front-end modules;
- All the clocks generated by the LpGBT are synchronous with the LHC machine clock;
- Thus the LpGBT “knows” exactly the frequency of the incoming data!
- A CDR circuit is thus not needed for each ePort.
- However, the incoming phase of the data is not known!
- A mechanism is thus needed to sample the incoming data at the optimum sampling point!

Up eLink – Phase Alignment

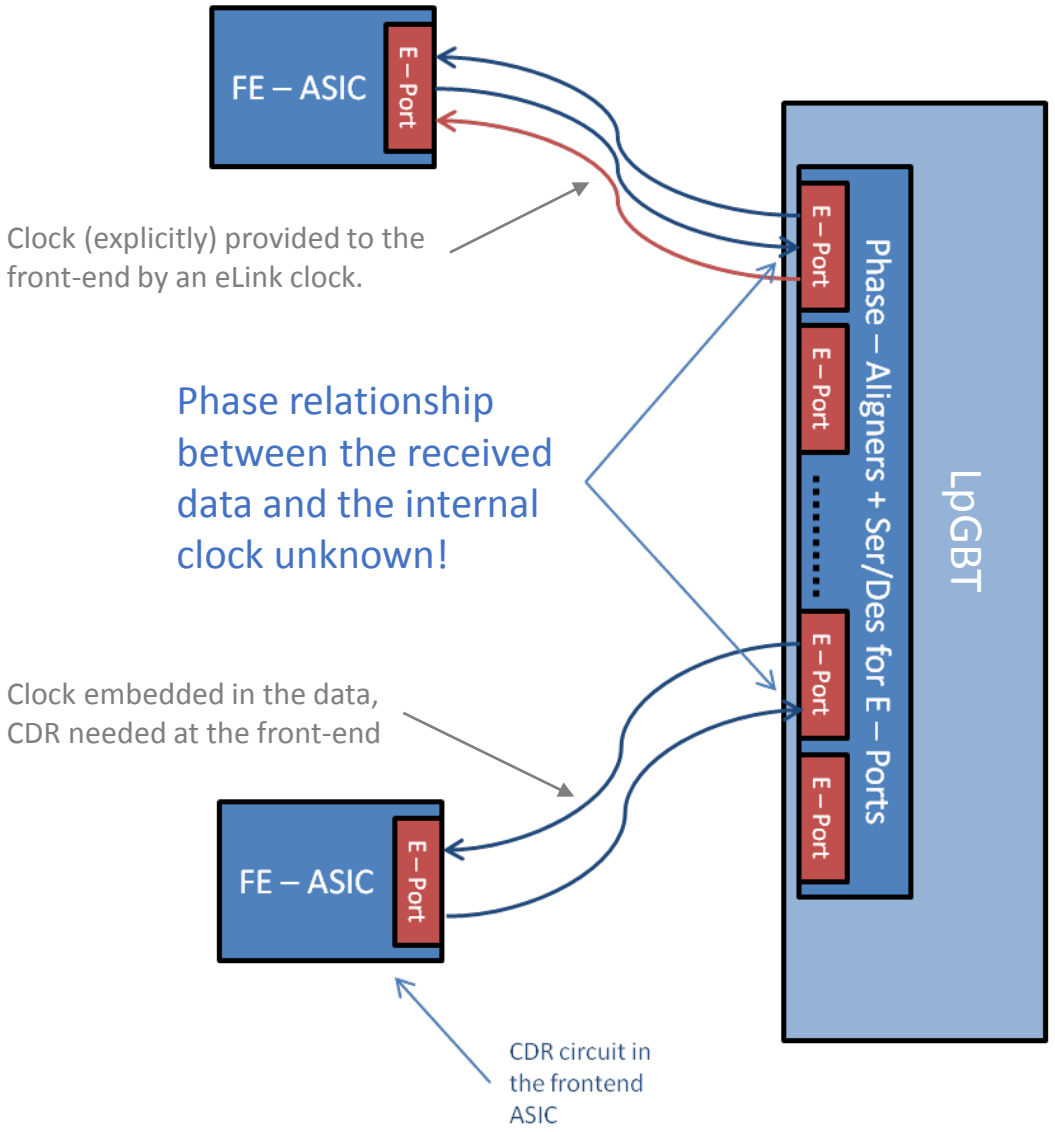

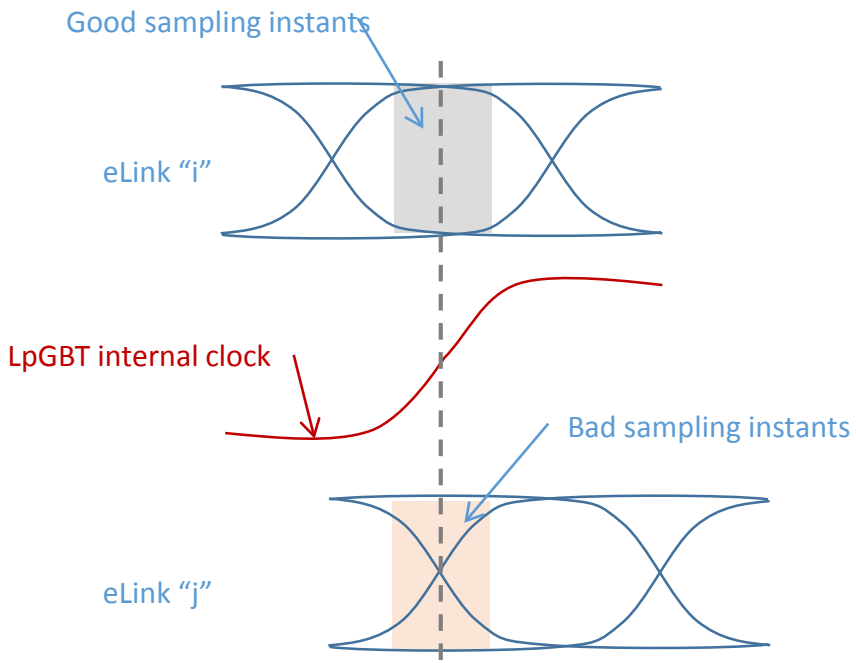

- The phase of the incoming data signals is "unknown" in relation to the internal sampling clock!
- There are up to 28 eLink inputs (potentially) all with random phase offsets
- The solution:
  - "Measure" the phase offset of each eLink input
  - Delay individually each incoming bit stream to phase align it with the internal sampling clock

Phase aligner - Principle of operation

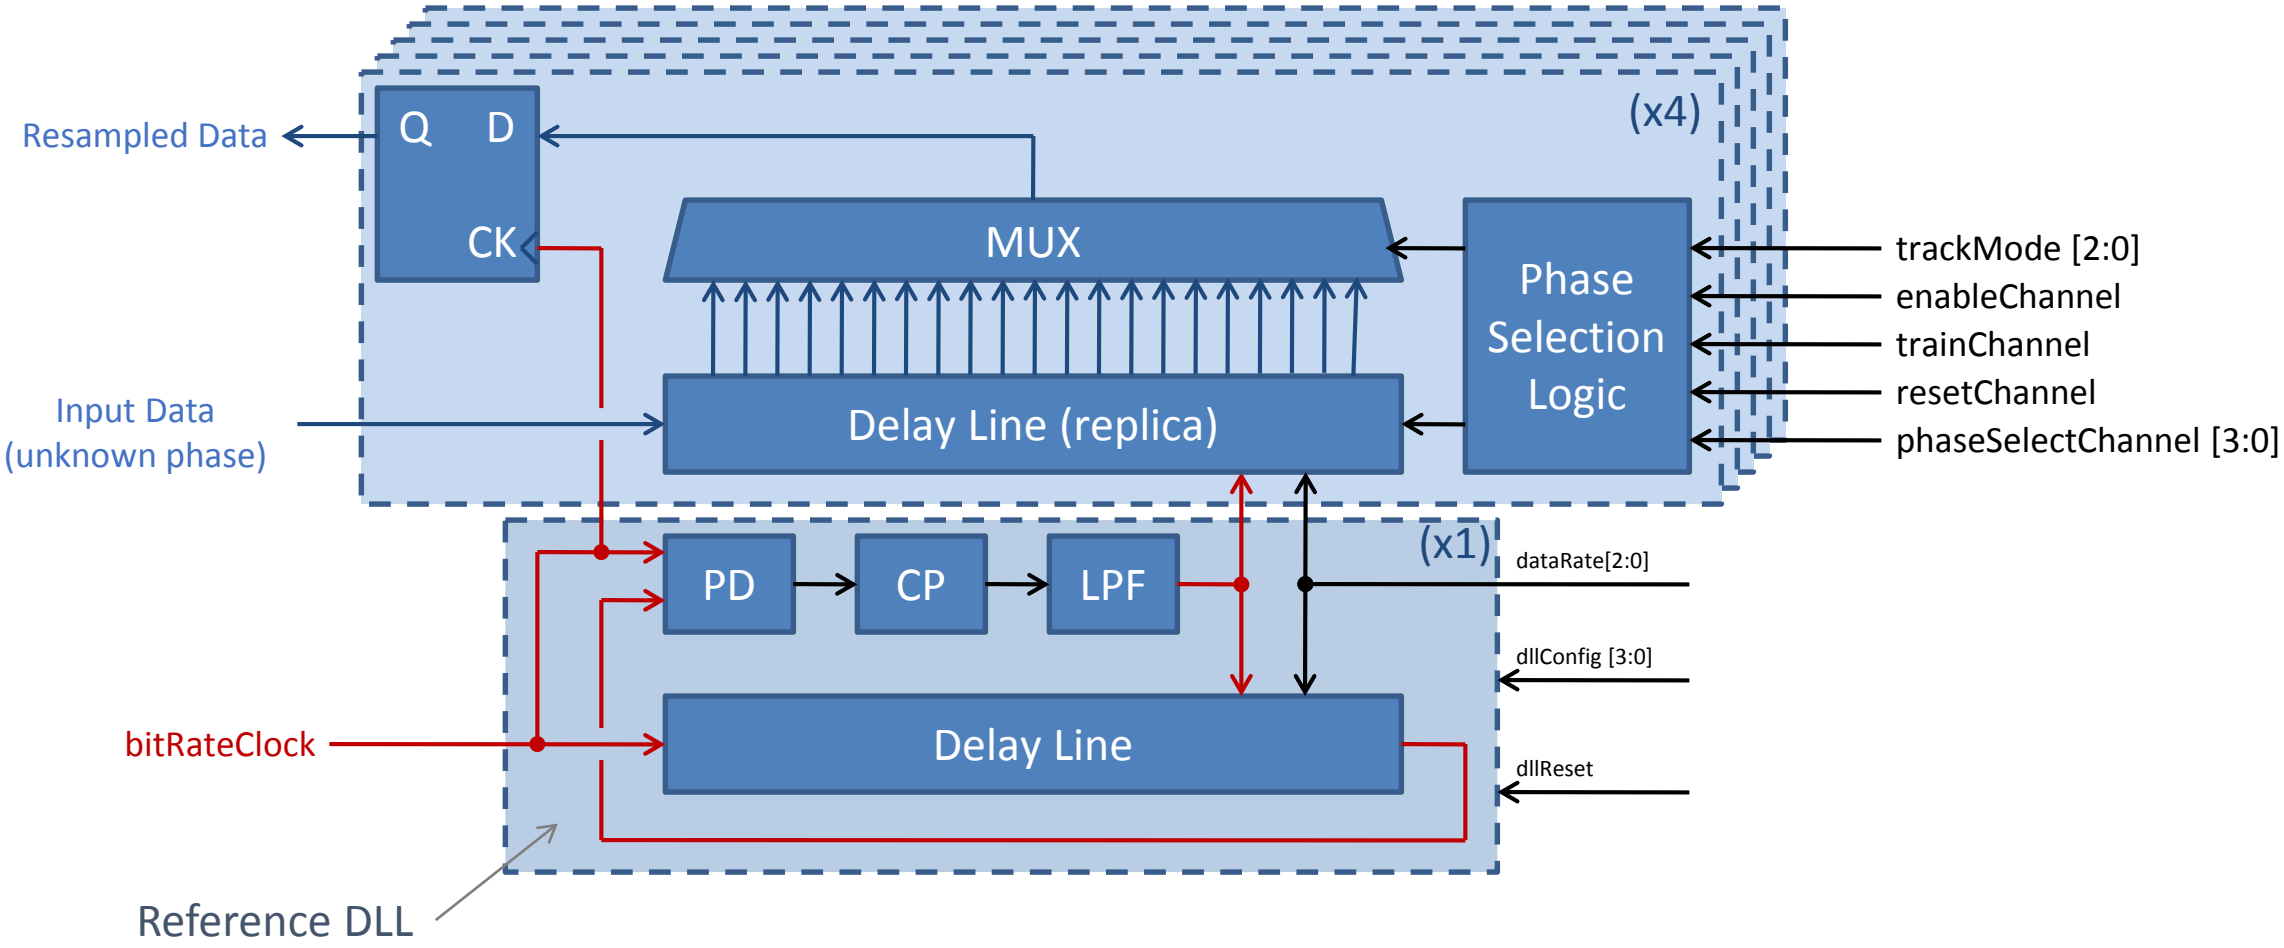

Phase - Selection

1. "Examine" all the phases
2. Detect where the data "edges" are in relation to the clock;
3. Choose the phase that has the edges better centered around the clock
4. Once aligned, the PA can track the data phase wanders that cover virtually a full clock cycle:
   - To allow for this the delay line covers more than one bit period: 1.75 × T_{bit}
   - And, during initialization only phases 4 to 11 are allowed

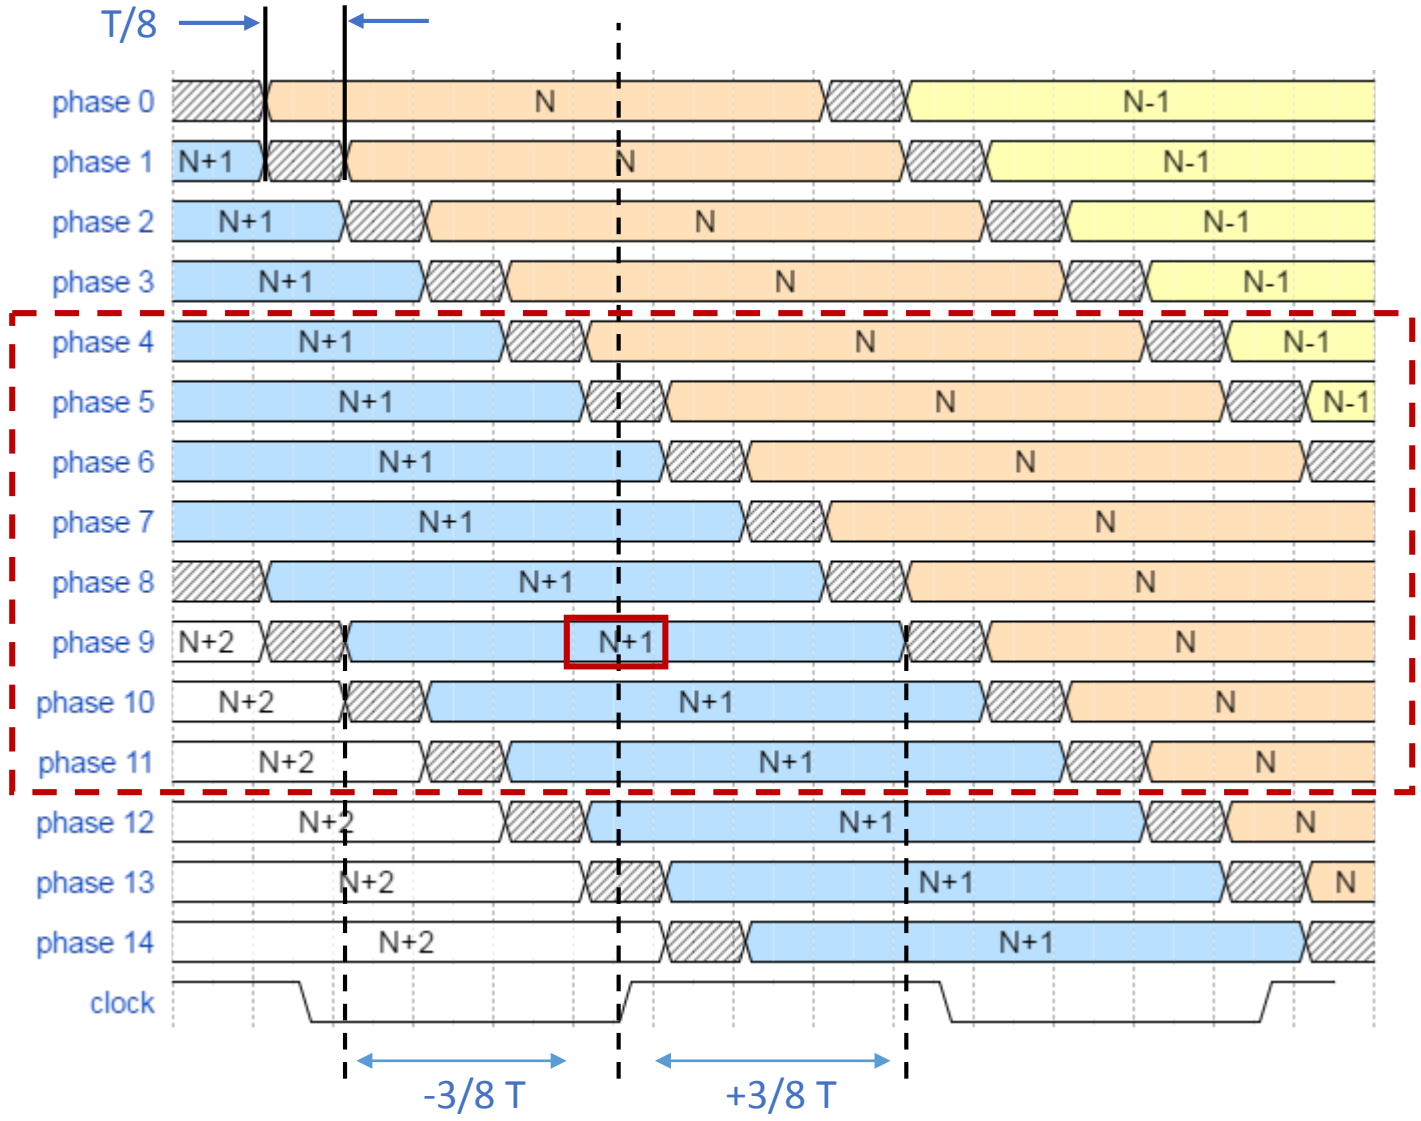

Phase Generation – Principle

- The scheme depends on “accurately” matching the “unit cell delay” with the bit period:
  - $\Delta t = T_{bit}/8$
- For that, the reference DLL contains 16 “unit cells” and it is calibrated to twice the bit period ($f_{bit}/2$):
  - Leading to $\Delta t = T_{bit}/8$
- The exception is the 160 Mb/s case where the DLL is calibrated at the same frequency as the bit rate (160 MHz):
  - Leading to $\Delta t = T_{bit}/16$
  - $\Delta t = T_{bit}/8$ requires two unit cells
  - To cover 1.75 × T_{bit} the length of the delay line is doubled!

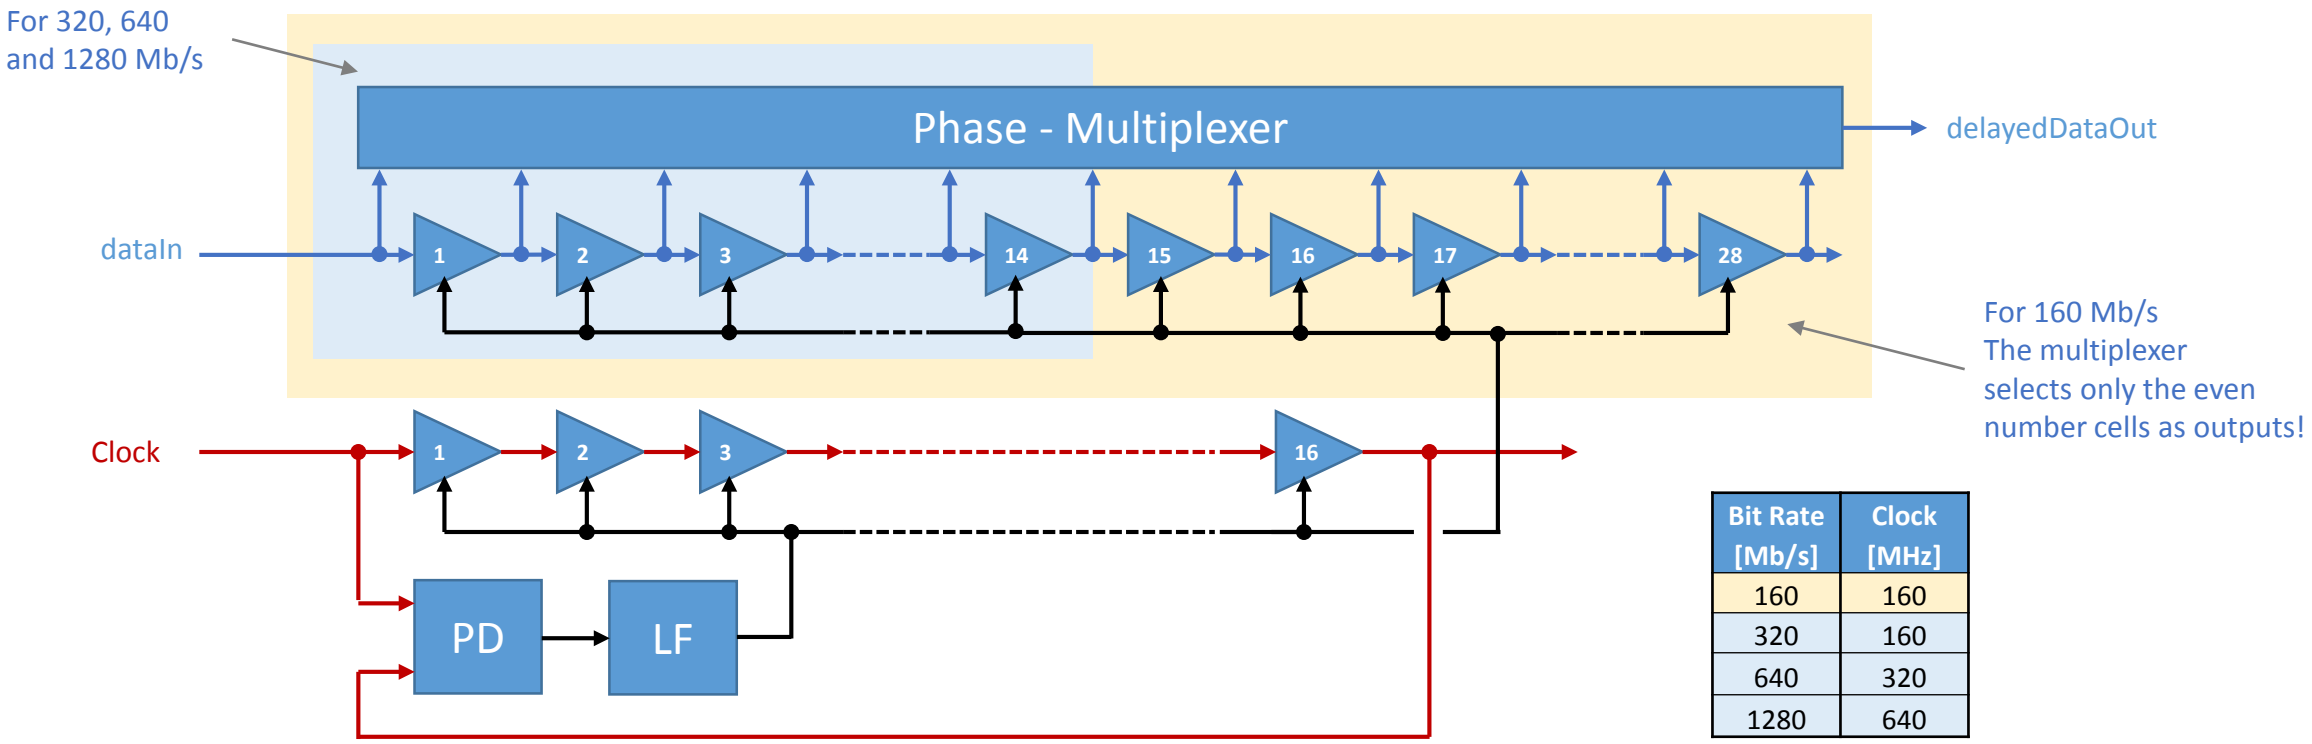

Bit Rate [Mb/s]	Clock [MHz]
160	160
320	160
640	320
1280	640

Unit Cell Delay

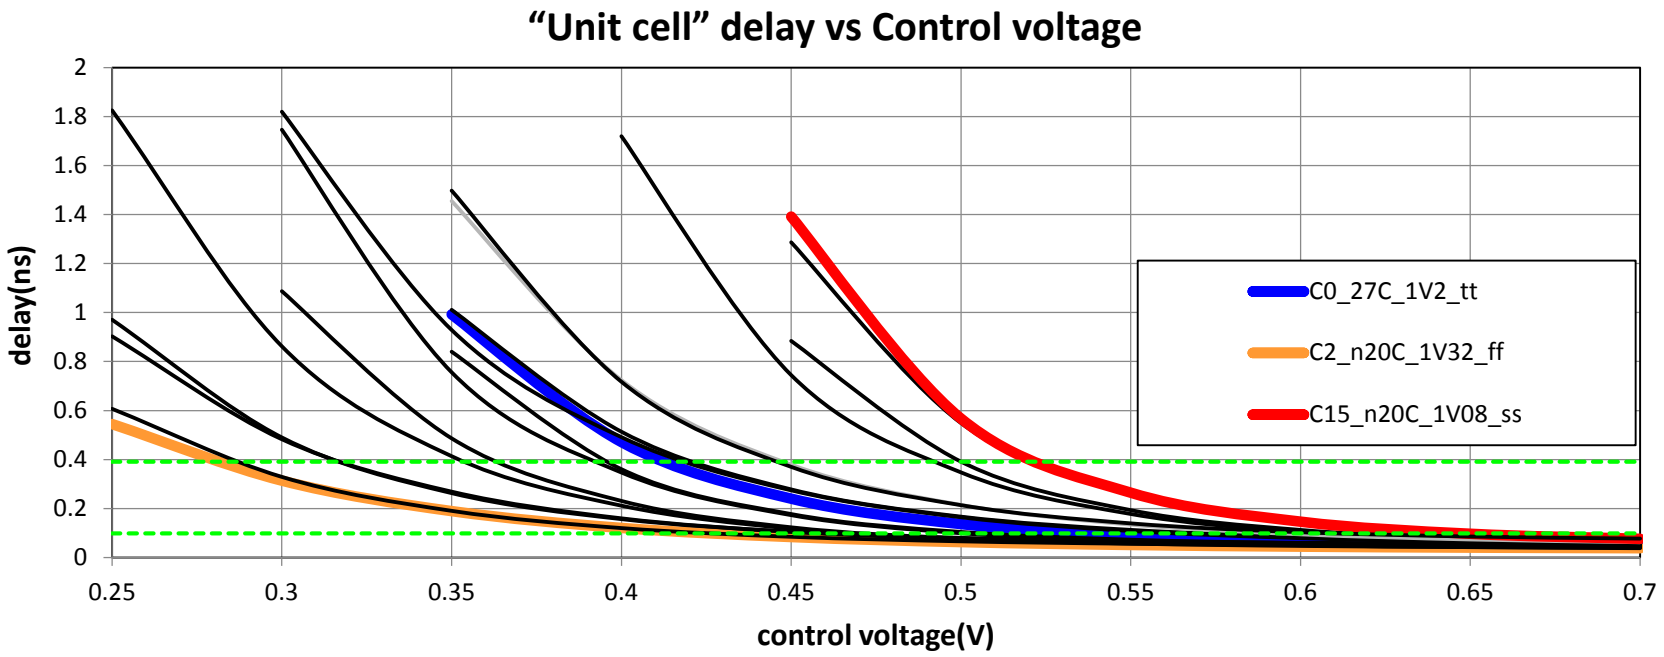

- 160 / 320 Mb/s: Δt = 390.6 ps ($T_{bit}/8$)
- 640 Mb/s: Δt = 195.3 ps
- 1.28 Gb/s: Δt = 97.6 ps ["Just about" for all process corners]

Probably the max data rate this method can be comfortably used in 65 nm CMOS!

Phase Generation – Performance (640 Mb/s)

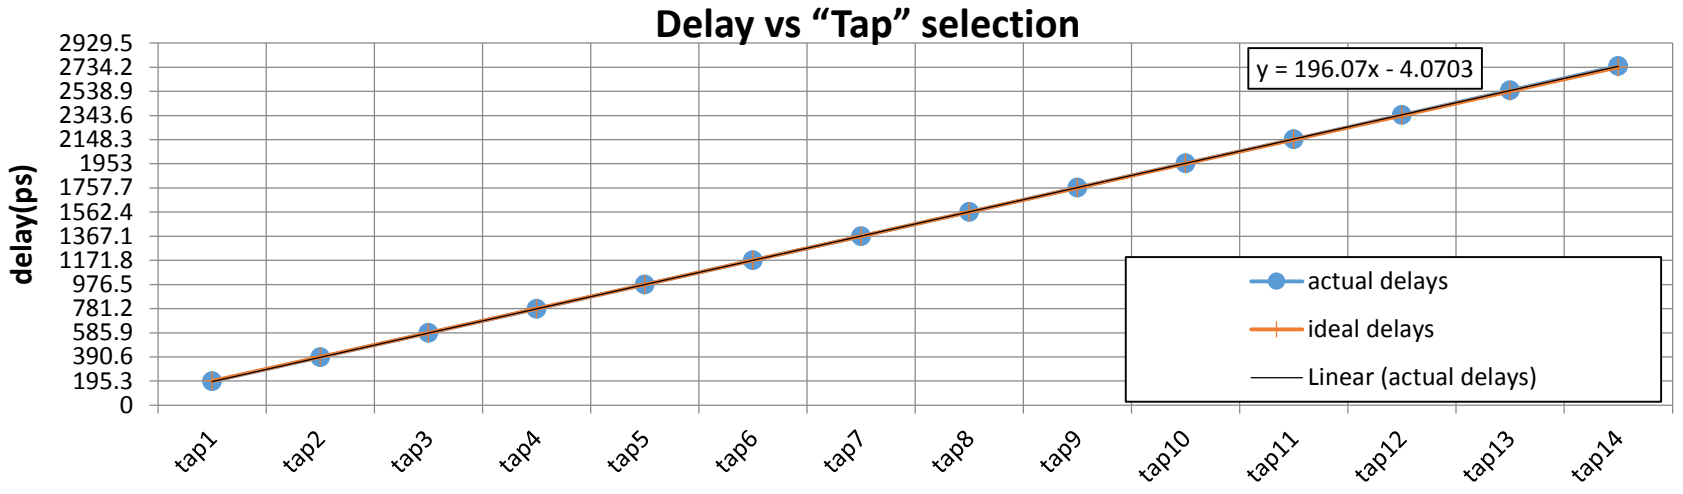

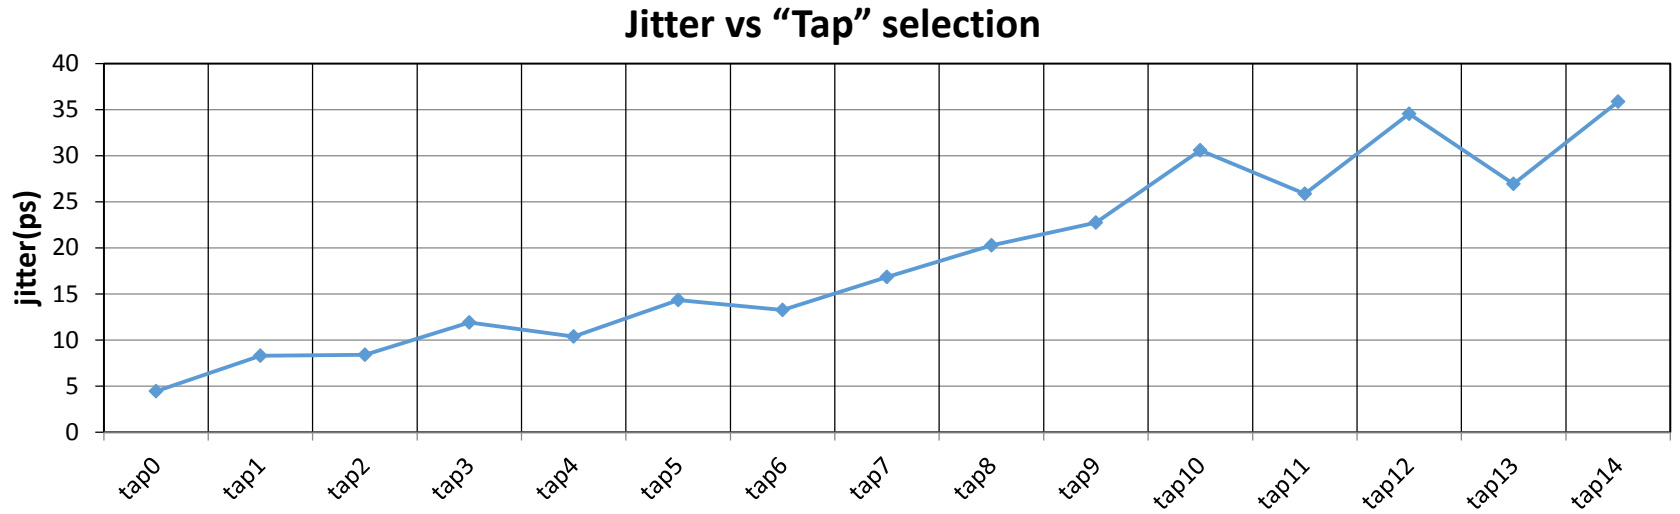

B = 640 Mb/s
T = 1540.8 ps
ΔT = 192.6 ps

J_{max} = 35.7 ps

J_{max}/T = 2%
$J_{max}/\Delta T$= 18%

Edge Detection – Principle

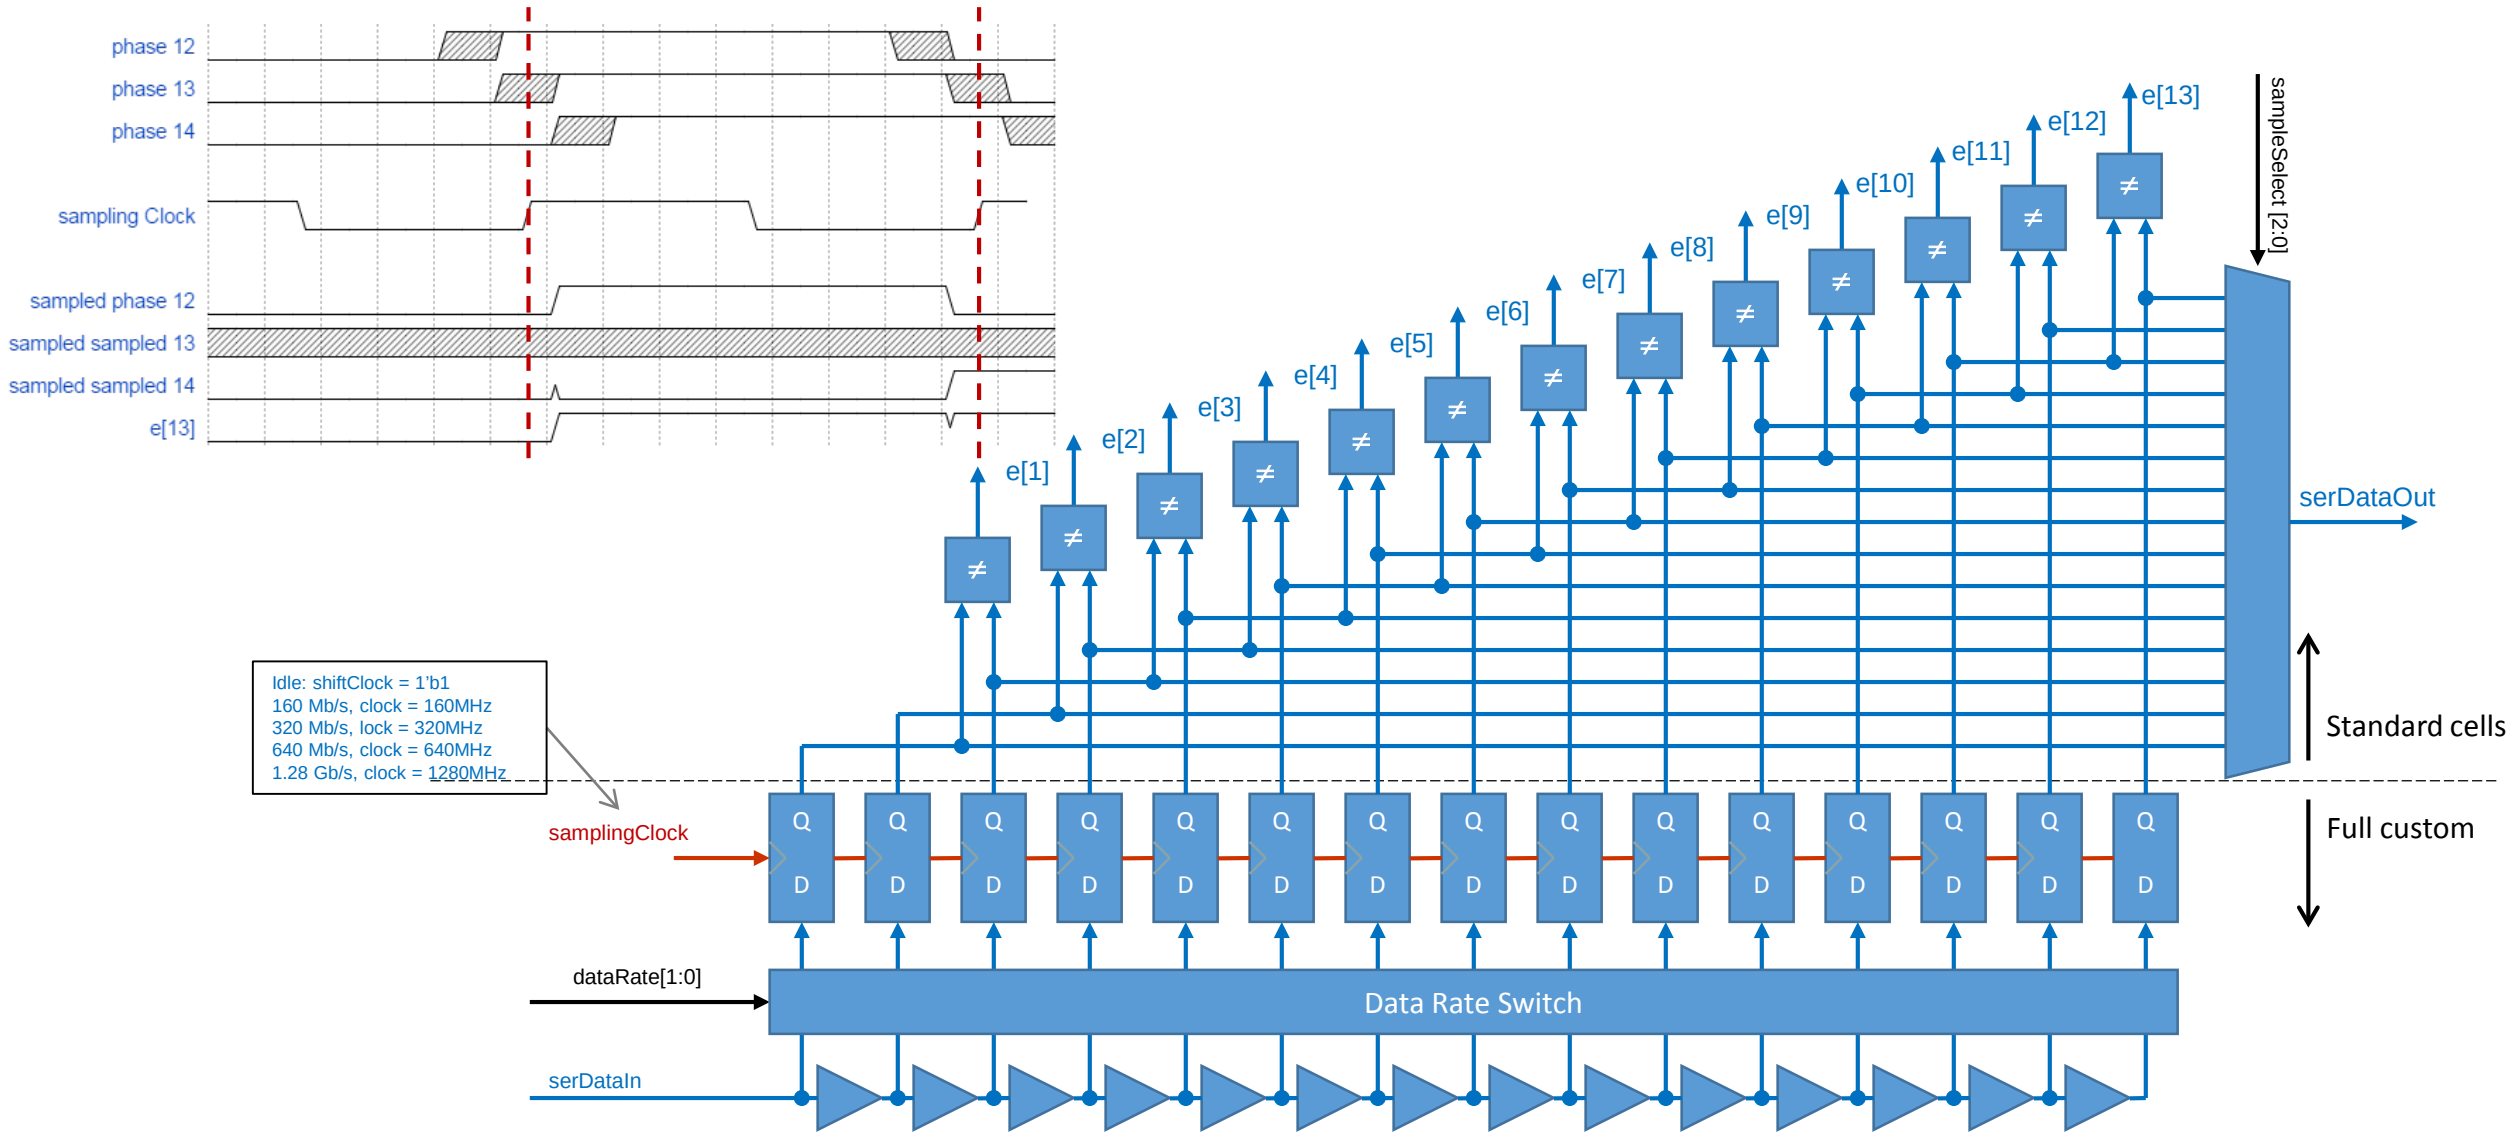

Phase Alignment vs Power Dissipation (1.28 Gb/s)

- Three modes of operation:
  - Automatic phase tracking
  - Training with learned static phase
  - Static phase selection
- Automatic phase tracking requires the lock state machine to constantly monitor the edge transitions and update the phase selection when necessary
- Static phase-selection requires the operator to select the proper phase (the selection is unlikely to change during a full run).
- In this mode, it is possible to reduce the power consumption by:
  1. Disable the delay-line outputs, except the one required
  2. Prevent the signal from propagating further though the delay line.

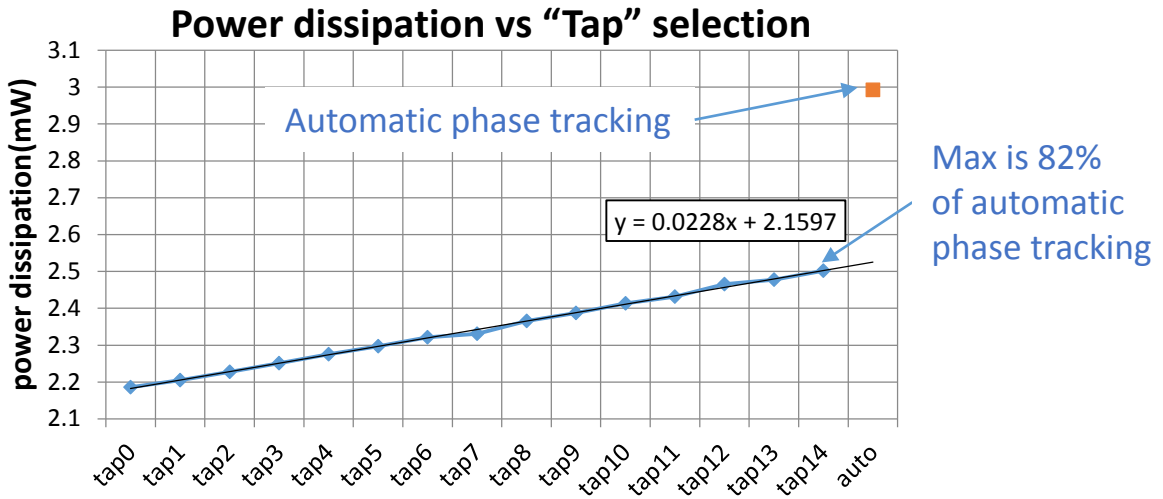

Delay Line power dissipation [mW]			
Mode	Static		Automatic
Corner	1st Tap	Last Tap	
C0	2.7	3.4	4.3
C2	6.7	7.9	9.2
C15	1.8	2.3	3.0

Details That Count!

- In the "Training mode", the circuit switches from "automatic" to "Static phase", disabling the unused outputs and stopping the propagation of the signal further down the delay-line.
- However, the absolute value of the "unit" cell delay is affected up to 8% even if careful buffering is added!
- To avoid this, a dummy gate is added and enabled each time the output gate is disabled.
- This allows to maintain the same loading conditions for all the cells and minimizes the delay mismatch!
  - Less than 0.2% in C0

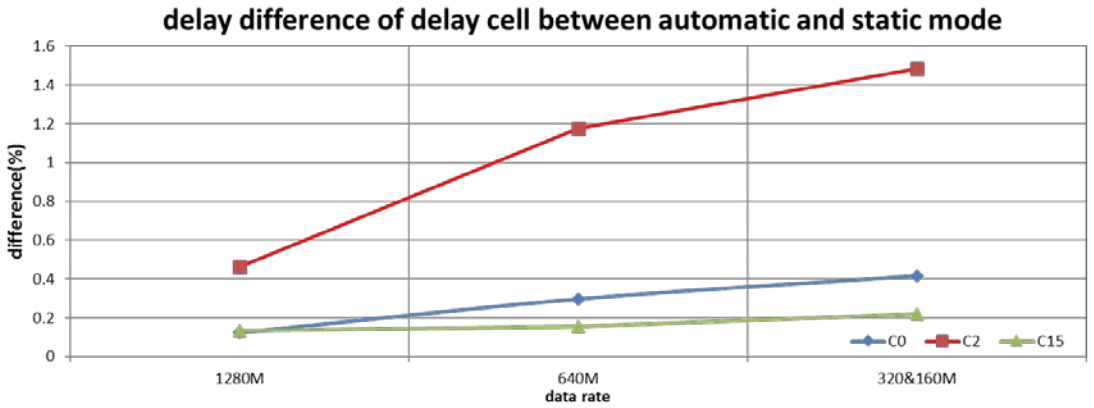

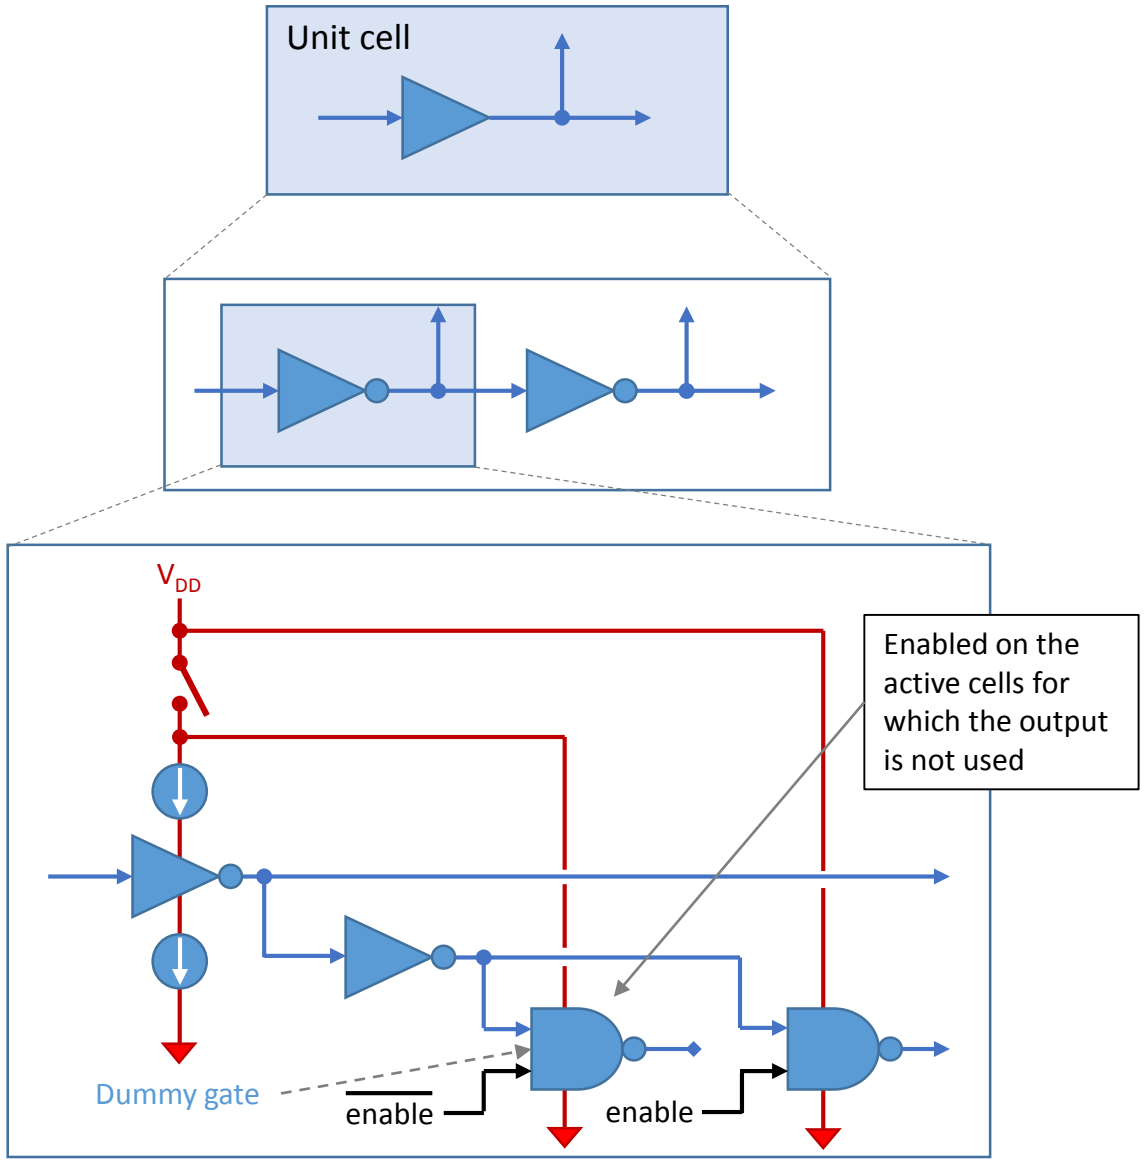

Phase – Aligner Layout

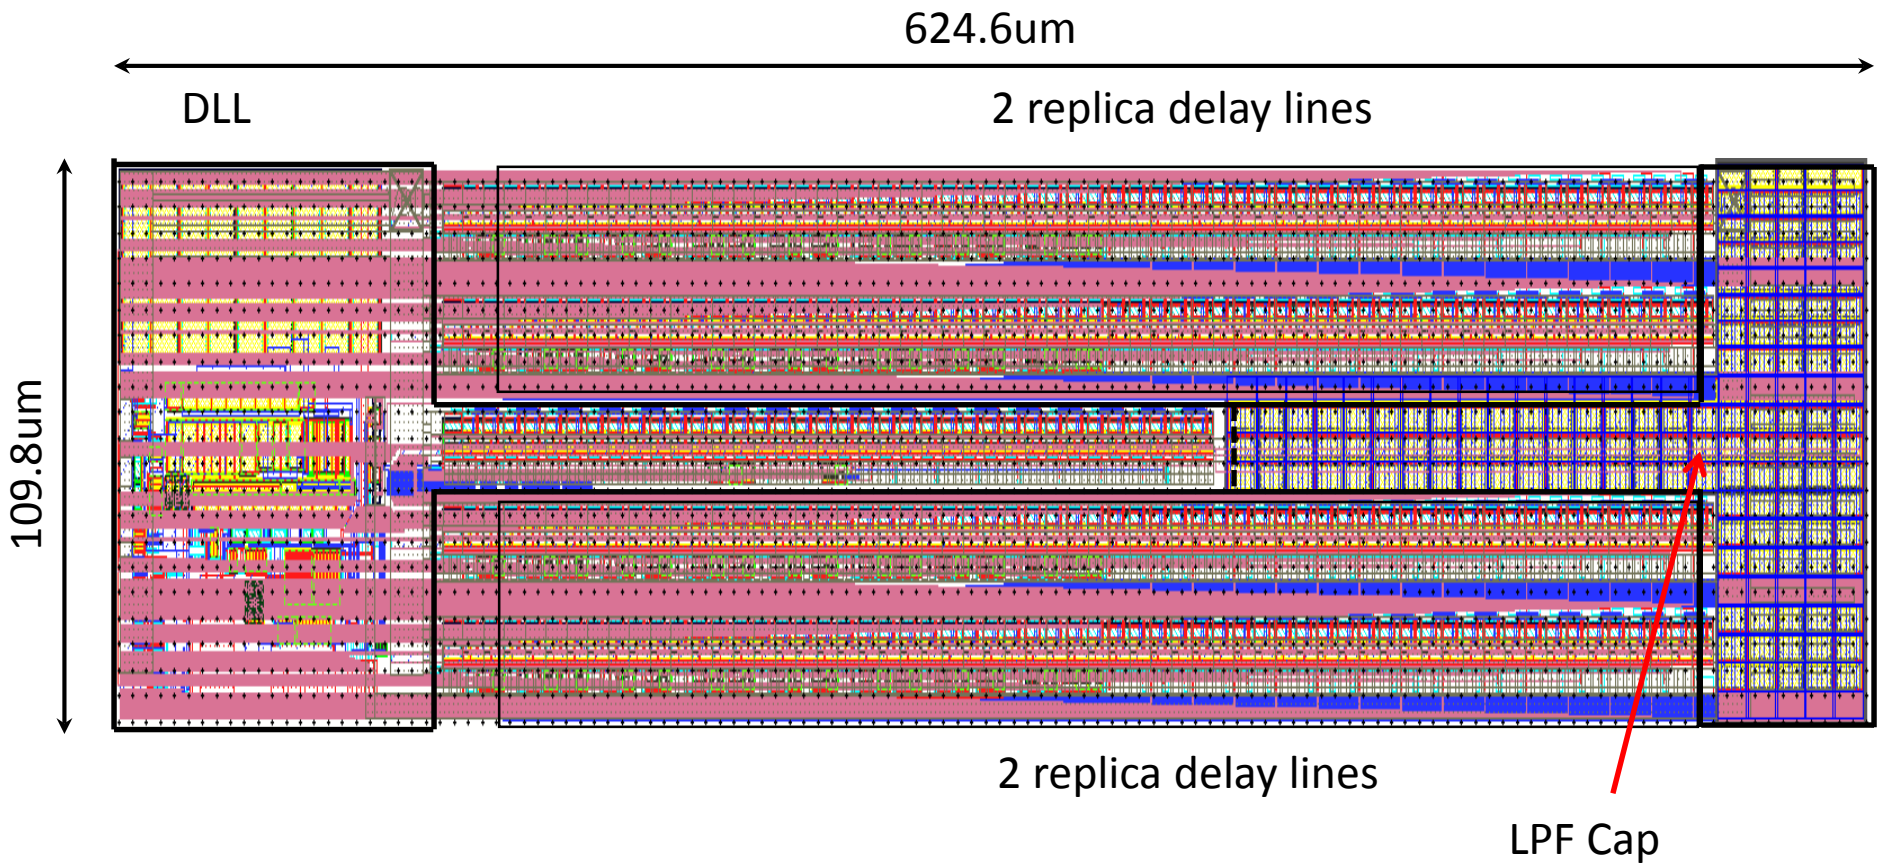

Connecting With the Counting Room

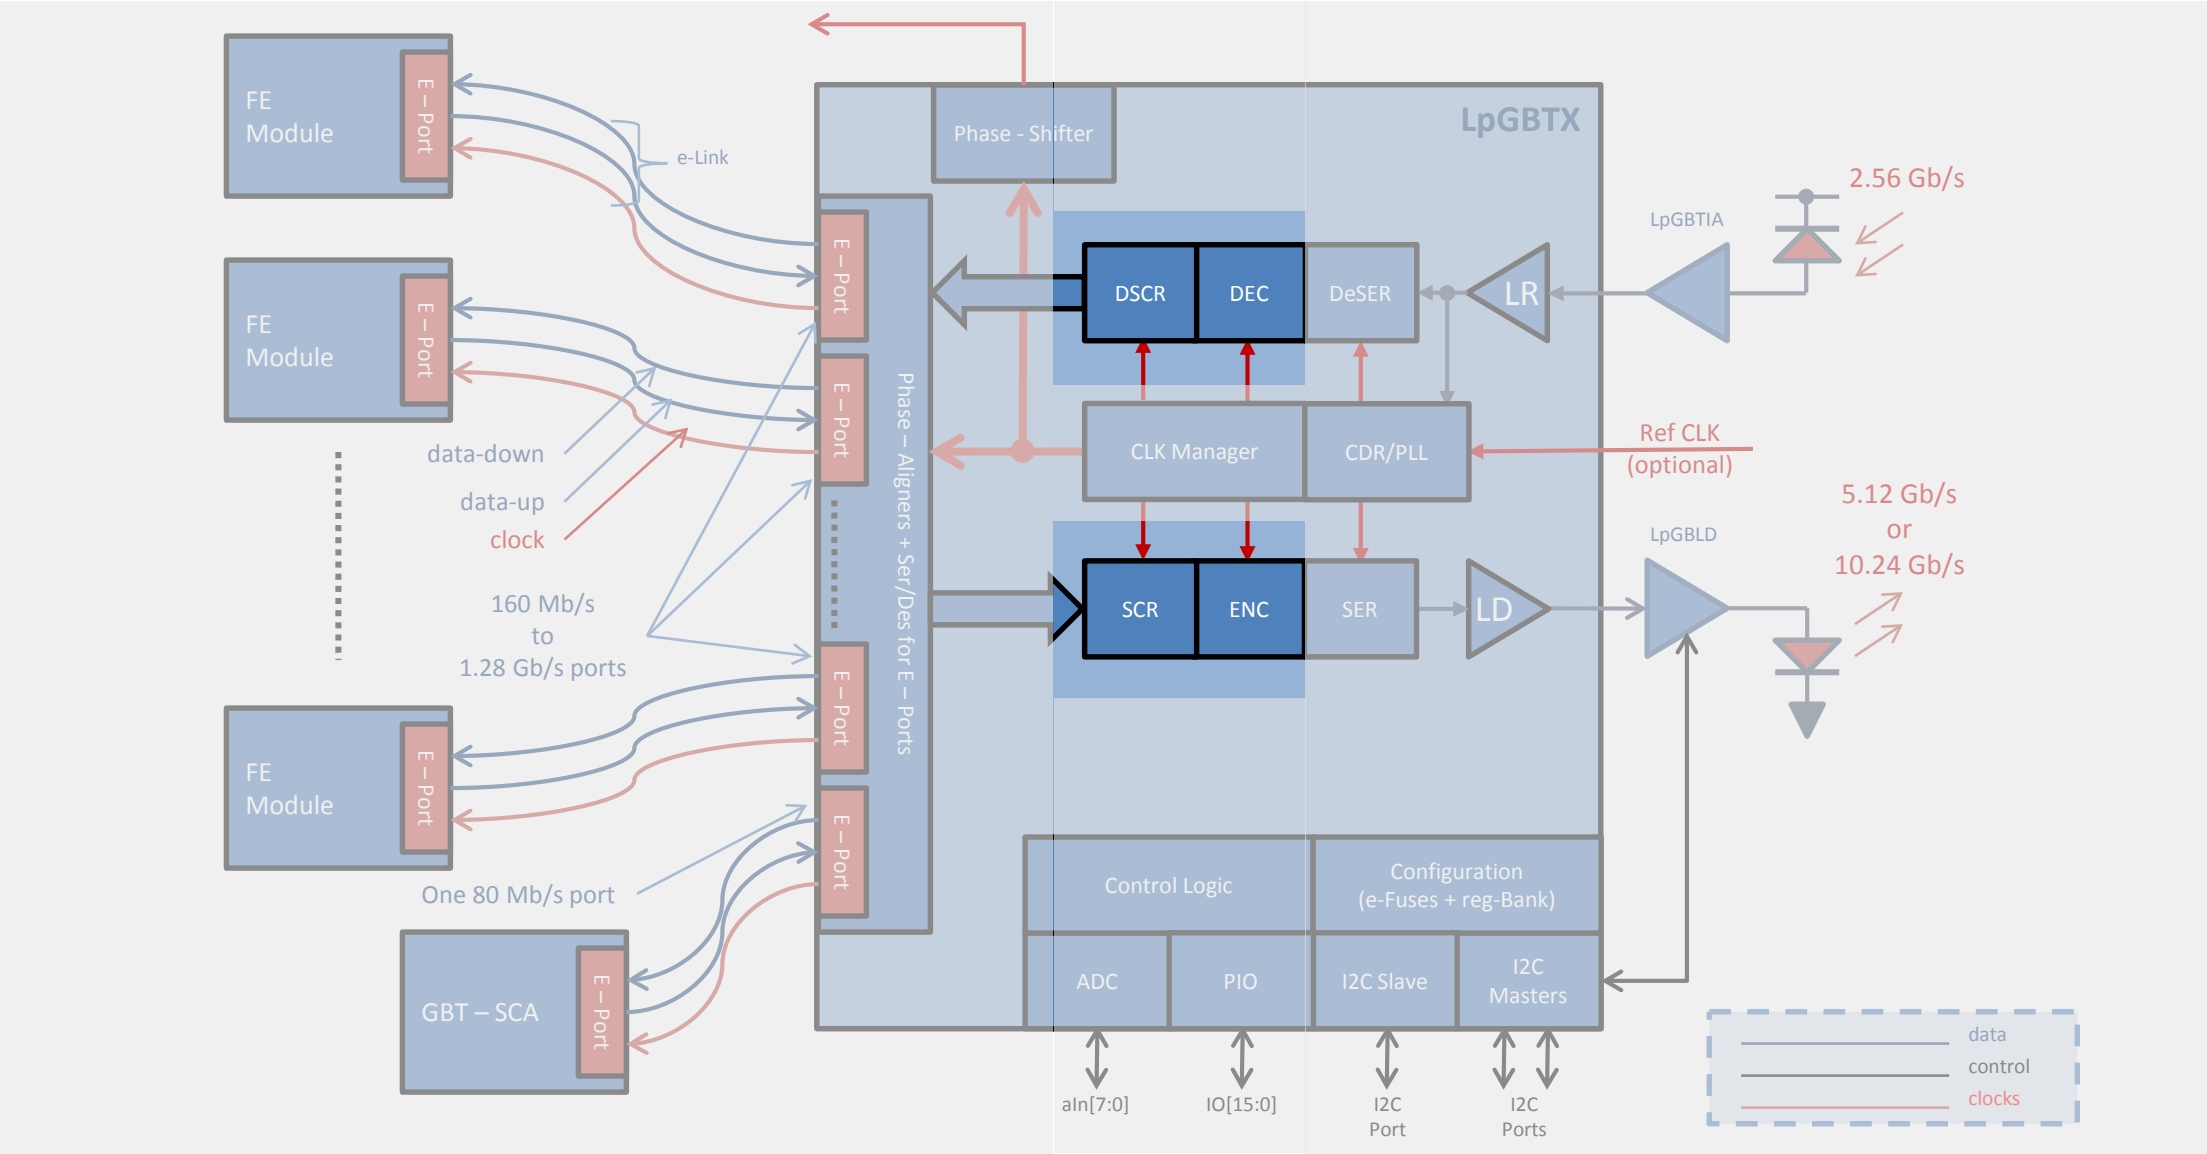

High – Speed Links

- The LpGBT supports the following data rates:
  - Down link: 2.56 Gb/s
  - Up-link: 5.12 / 10.24 Gb/s
- In all cases data is transmitted as a frame composed of:
  - Header
  - The data field
  - A forward error correction field: FEC5 / FEC12
- The data field is scrambled to allow for CDR at no [additional] bandwidth penalty
- Efficiency = # data bits/# frame bits

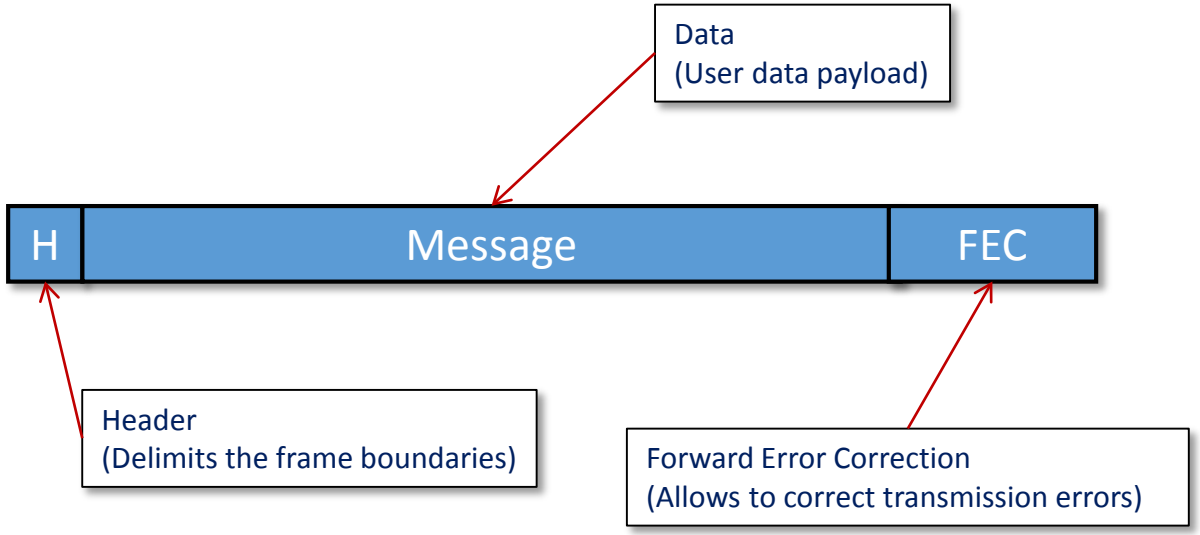

	Down-link	Up-Link			
	2.56 Gb/s	5.12 Gb/s		10.24 Gb/s	
		FEC5	FEC12	FEC5	FEC12
Frame [bits]	64	128		256	
Header [bits]	4	2		2	
Data [bits]	36	116	102	232	204
FEC [bits]	24	10	24	20	48
Correction [bits]	12	5	12	10	24
Efficiency	56%	91%	80%	91%	80%

The order of operations is important

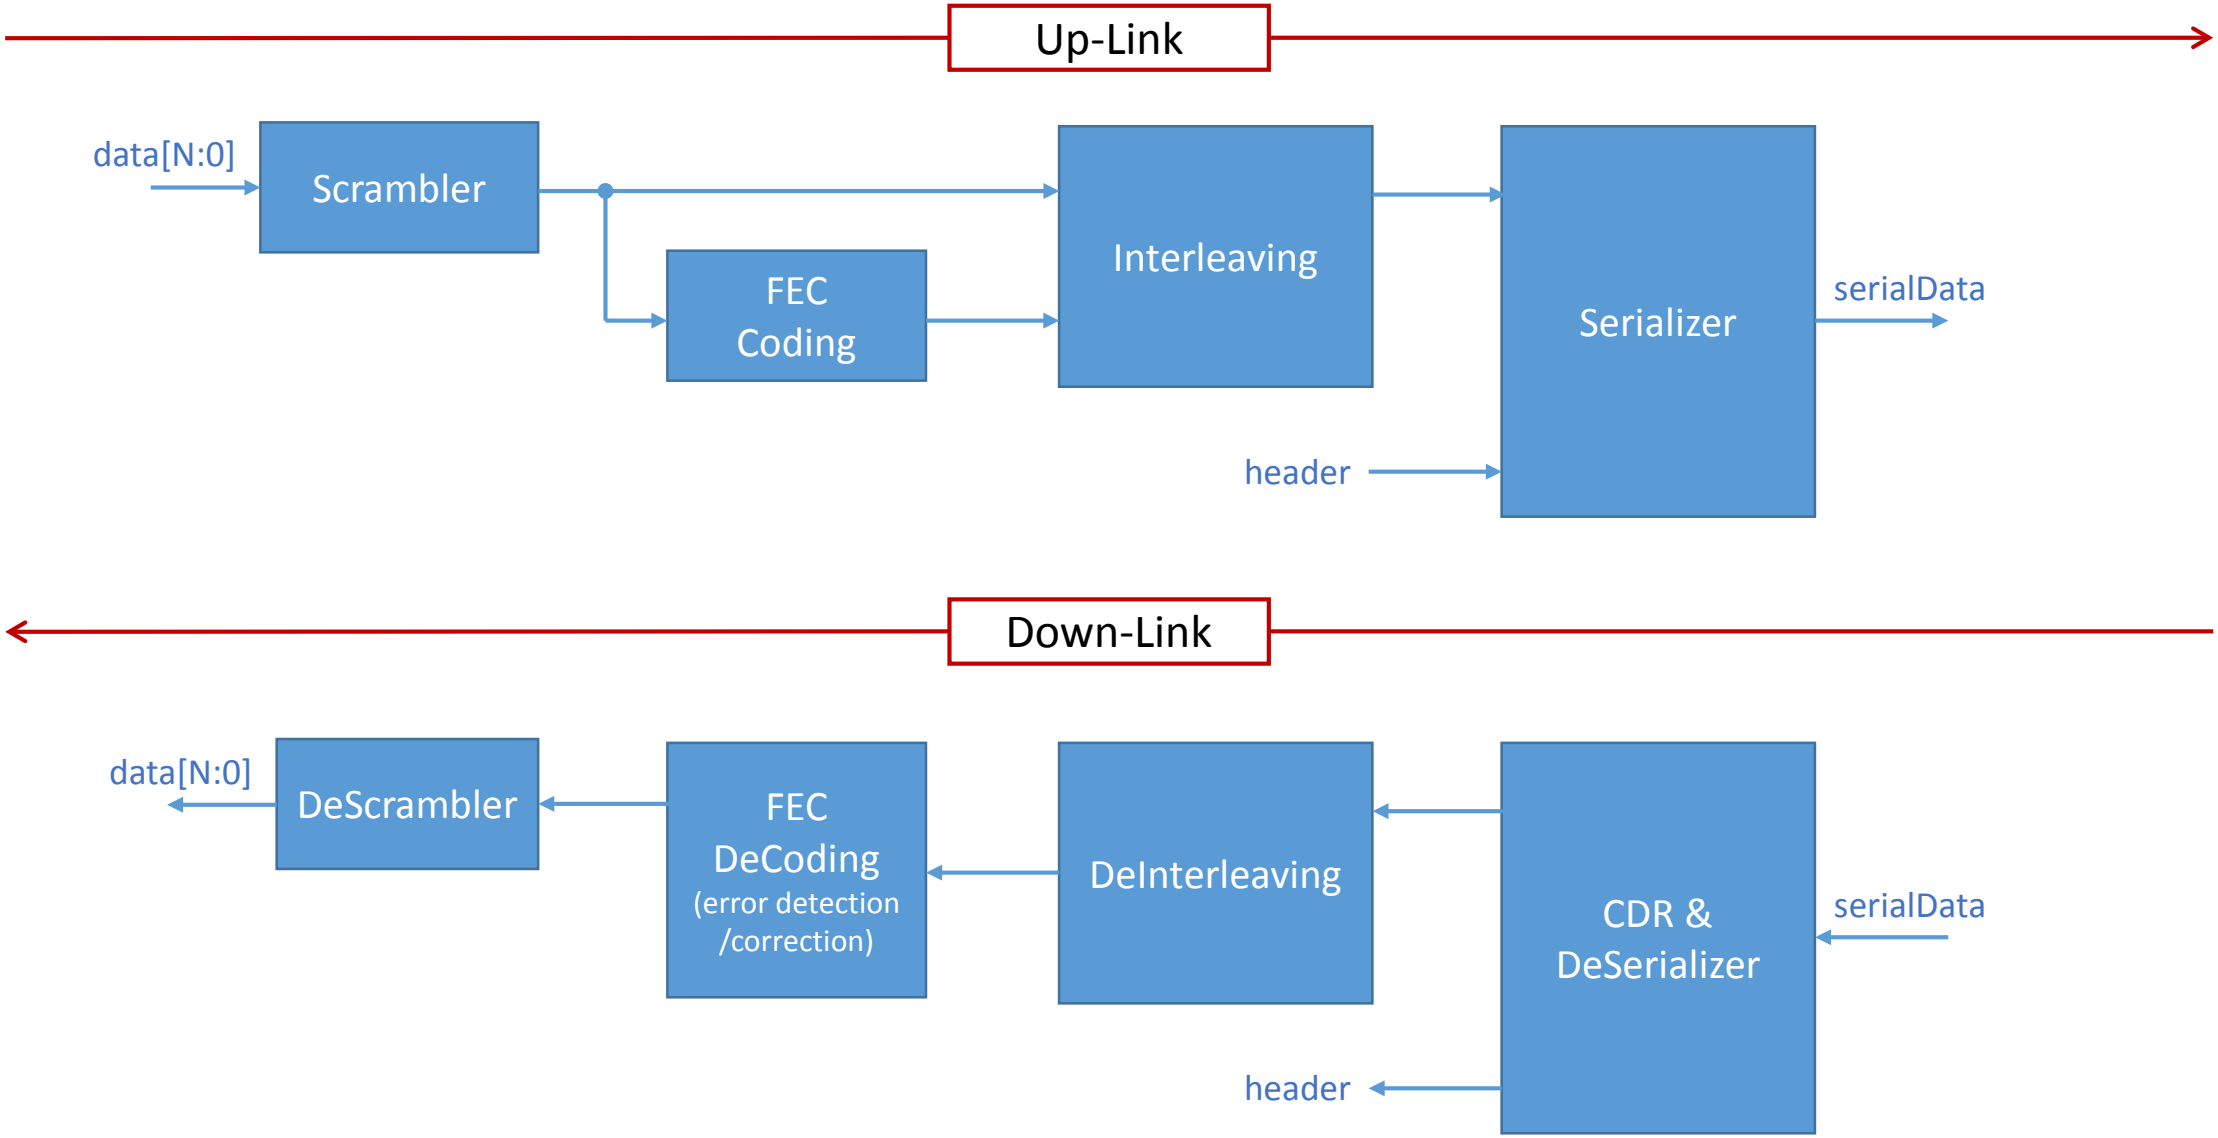

Scrambling

- Converts a bit stream into a “random” sequence of bits
- Objective:
  - Remove the “DC” contents of the data (so called DC balance)
- Needed for:
  - Transmission over an “AC” coupled channel
  - Guaranty enough transitions so that the receiver can recover the clock from the data
- Scrambler/DeScramblers types:
  - Synchronous
  - Self-synchronizing (used in the LpGBT)

Serial scramblers require the shift registers to run at the bit rate!

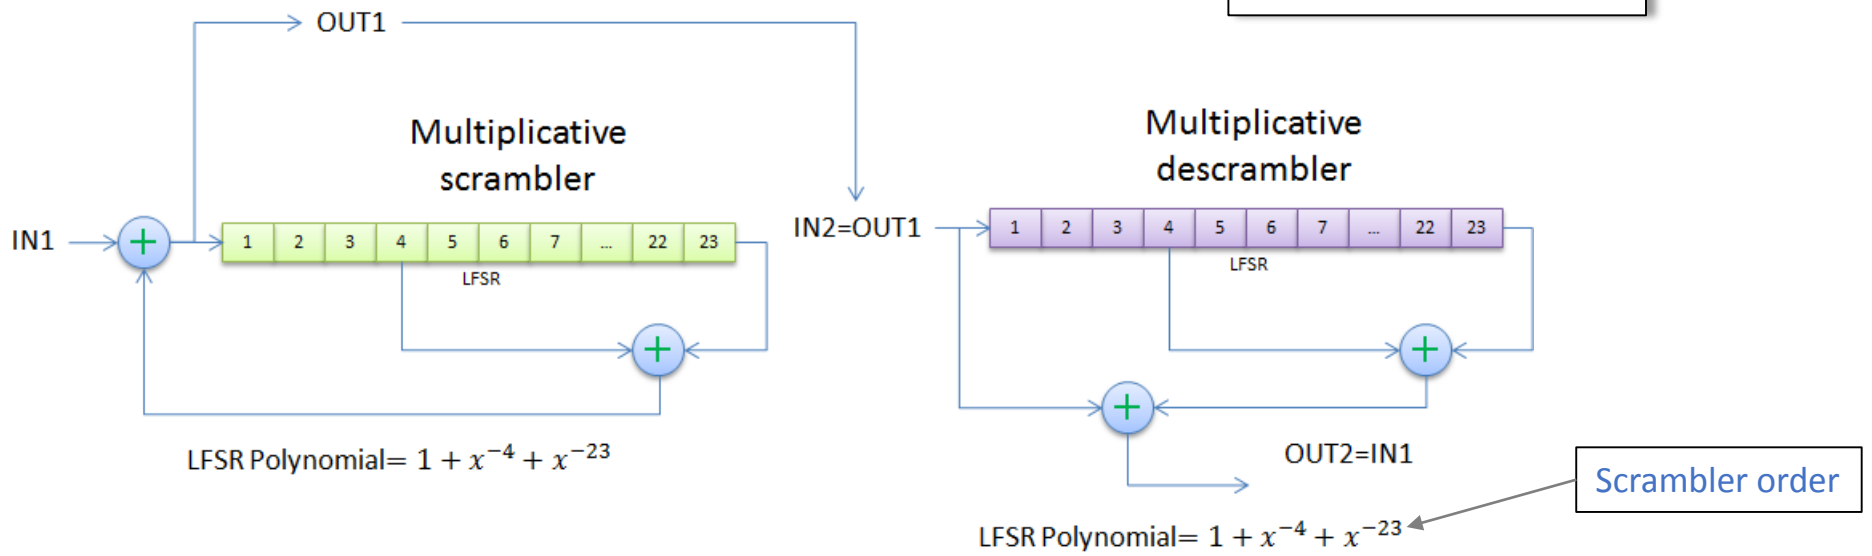

Parallel Implementation

Scrambler

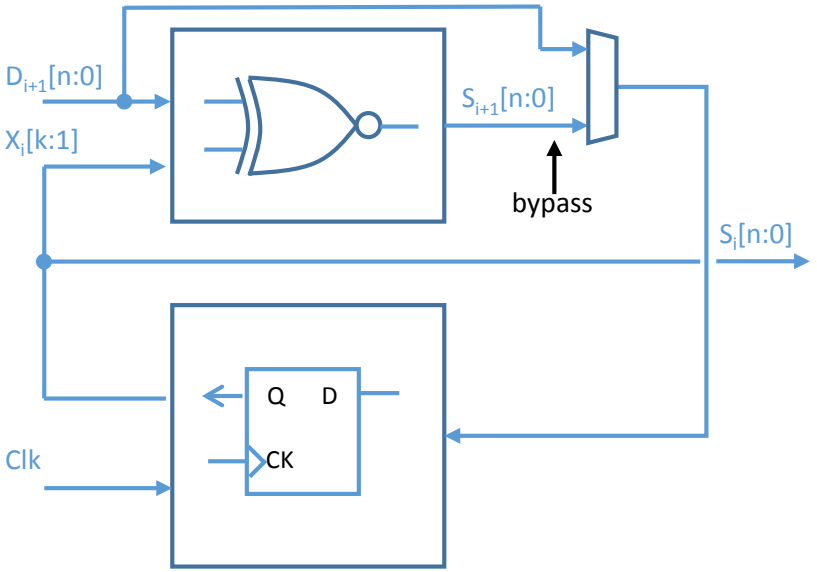

DeScrambler

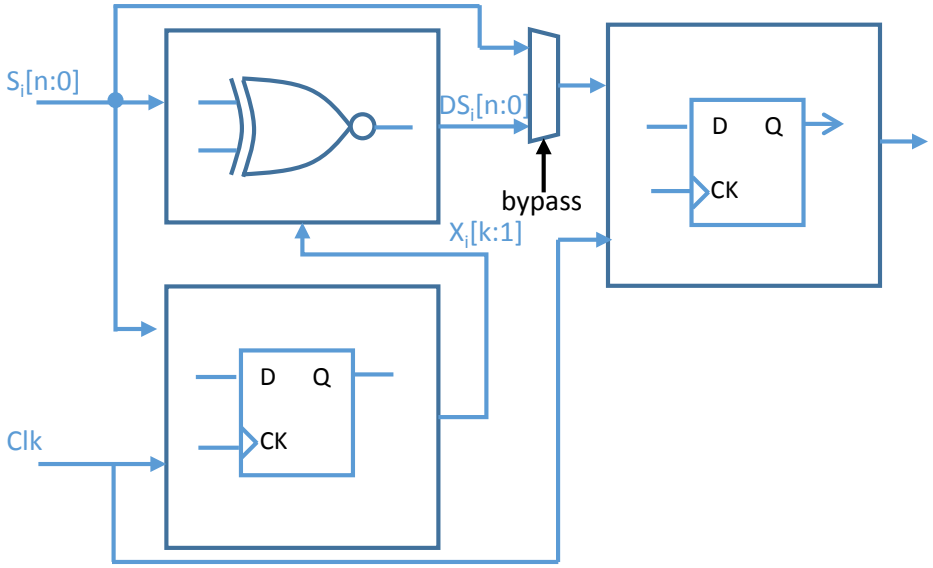

	Down-link 2.56 Gb/s	Up-Link			
		5.12 Gb/s		10.24 Gb/s	
		FEC5	FEC12	FEC5	FEC12
Data [bits]	36	116	102	232	204
Scrambler width [bits]	36	58	51	58	51
Scrambler order	36	58	49	58	49
Number of scramblers	1	2		4	
Recursive equation	eq 1	eq 2	eq 3	eq 2	eq 3

eq 1: $S_i = D_i$ xnor S_{i-25} xnor S_{i-36}
eq 2: $S_i = D_i$ xnor S_{i-39} xnor S_{i-58}
eq 3: $S_i = D_i$ xnor S_{i-40} xnor S_{i-49}

Error Correction Codes

- Due to Noise, Intersymbol Interference or SEUs information might be corrupted as it passes through a channel
- Forward Error Correction (FEC) gives the possibility of correcting errors without asking back the transmitted information
- This is achieved by adding "parity" bits to the transmitted data
- Bandwidth is tradeoff against transmission robustness

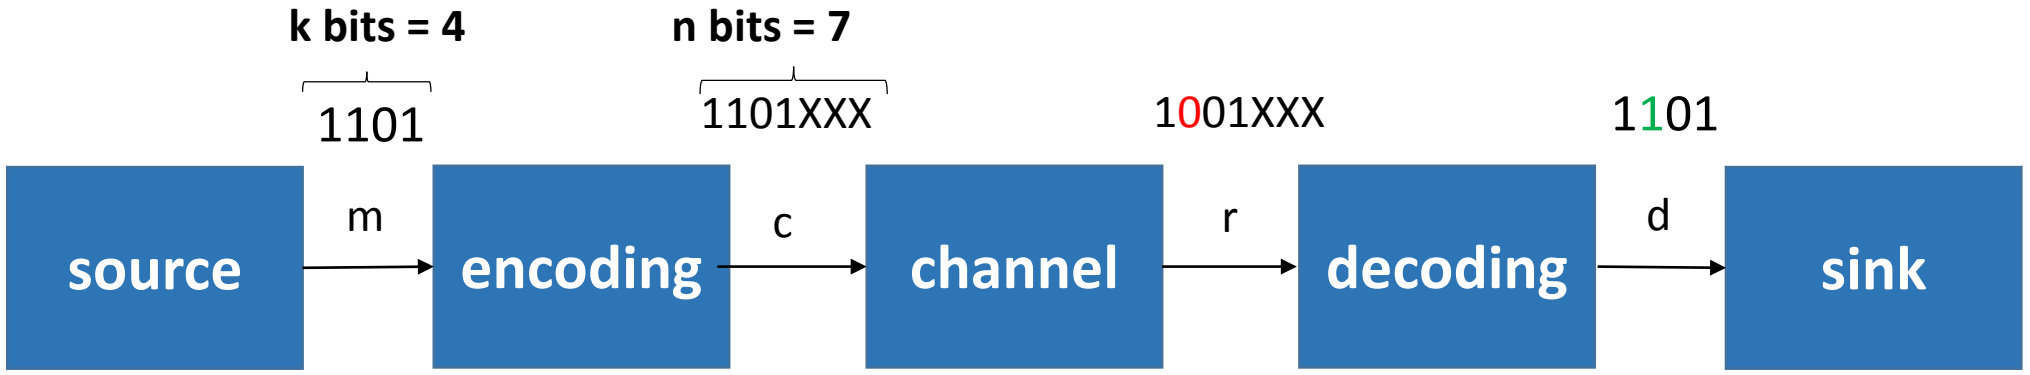

Reed-Solomon Codes

- Non-binary code:
  - Based on symbols of length **m**-bits
- **k**-symbols of message are encoded into a **n**-symbol code word:

$$n = 2^m - 1$$

- The number of parity symbols is:

$$n - k$$

- Allowing to correct up to

$$t = (n - k)/2$$

symbols

- Example:
  - RS(7,5)
    - n = 7 (code word length)
    - k = 5 (symbols to be coded)
  - m = $\log_2(7+1)$ = 3 bits
  - Error correction capability:
    - t = (7-5)/2 = 1 symbol

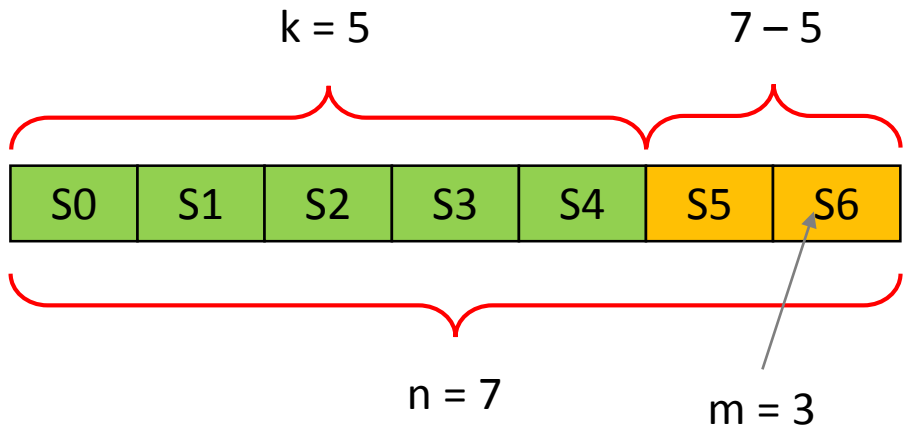

LpGBT Down Link

- Down-link FEC:
  - 4 codes interleaved
  - RS(5,3), m = 3 bits
    - This code is a “shortened” version of RS(7,5)
    - Due to frame size restrictions only 5 out of the 7 symbols are transmitted!
      - “Missing” symbols have to be used at coding time and assume received (known) by the decoder
- This code can correct:
  - Any random error in a frame of 36 bits (not considering parity bits)
  - Any burst error of length 10 bits
  - It can decode burst errors of length 12 bits if it is contained in the boundary of four symbols

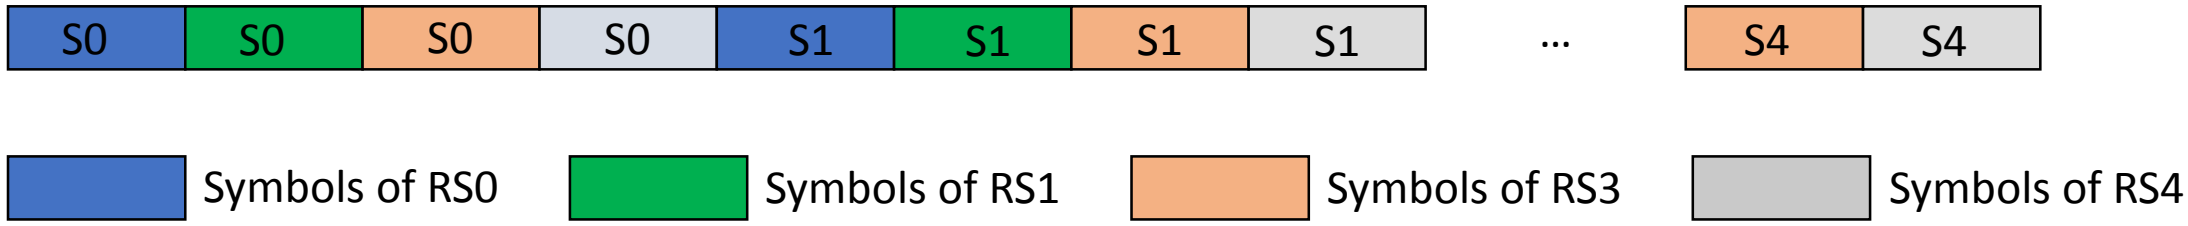

Experimental: LpGBT Down Link (LpGBT-FPGA)

System dominated by random noise

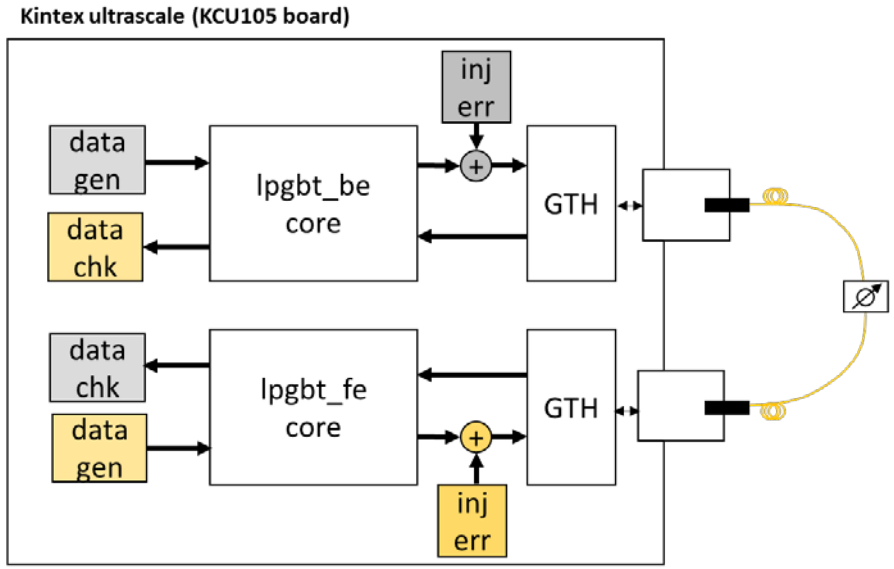

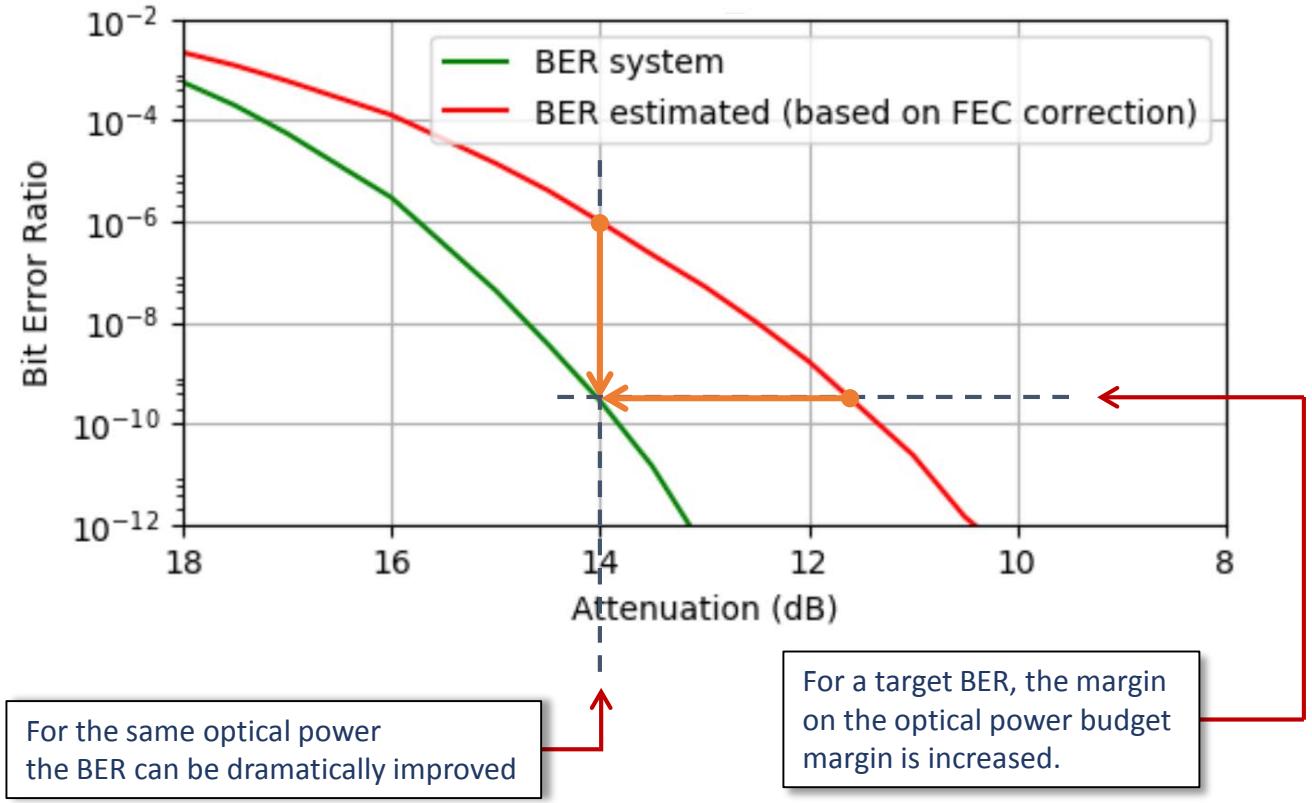

Experimental: GBTX Down - Link

- Radiation levels at 20 cm radius from the beam:
  - 2 ×10^{15} neutrons/cm2
  - 1 × 10^{15} hadrons/cm2
  - 50 Mrad total dose
- High rates of Single Event Upsets (SEU) are expected for SLHC links:
  - Particle "detection" by Photodiodes used in optical receivers
  - SEUs on Receivers,
  - SEUs on Laser-drivers
  - SEUs on SERDES circuits
- Experimental results confirmed that:
  - Error correction is mandatory to achieve error-rates ≤ 10^{-12}
- Upsets lasting for multiple bit periods will occur on PIN detectors!
- Upset lasting for multiple frames can occur in commercial TIAs

Test conditions:
- 70 MeV proton beam
- Flux of 2 × 10^8 p/cm^2/s
- ≈ 2 × the flux in the HL-LHC trackers

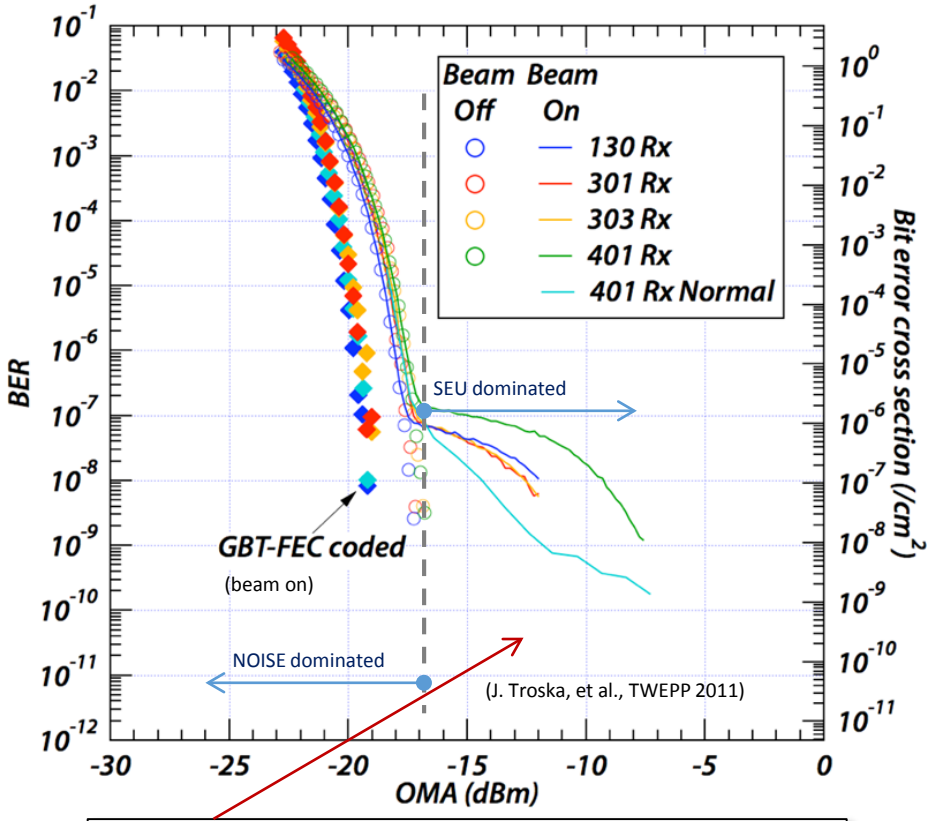

No errors were observed in the FEC-coded data-stream in the SEU dominated region with beam on!

High-Speed Serializer

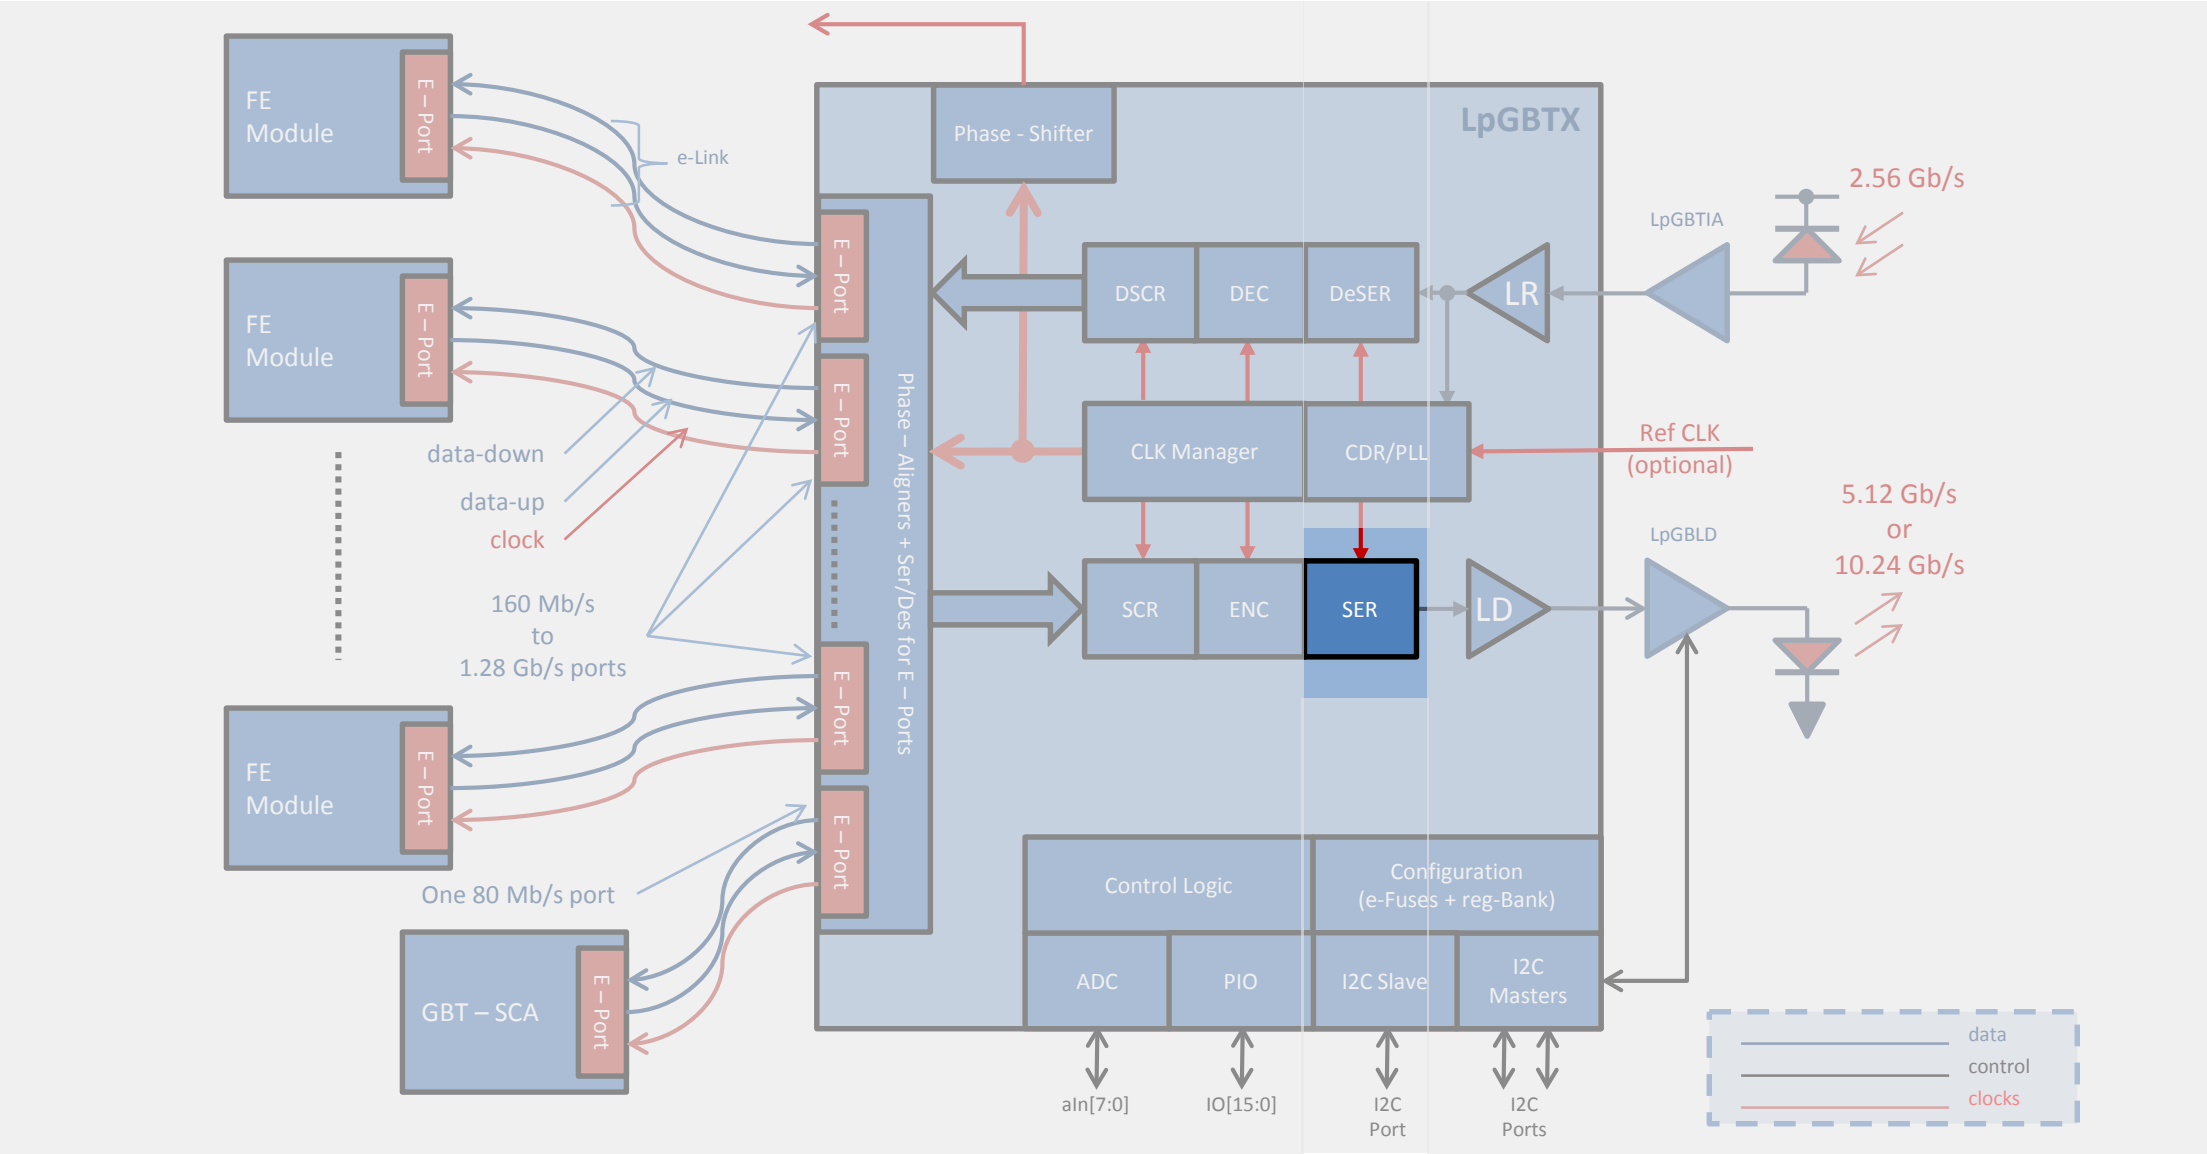

Serializer

- The LpGBT transmitter works at two data rates:
  - 5.12 Gb/s and 10.24 Gb/s
- This means that the Serializer has to convert the parallel data at 40 MHz into a serial bit stream:
  - 128 – bits @ 5.12 Gb/s
  - 256 – bits @ 10.24 Gb/s
- So what is a Serializer?
  - A “glorified” name for a shift register!

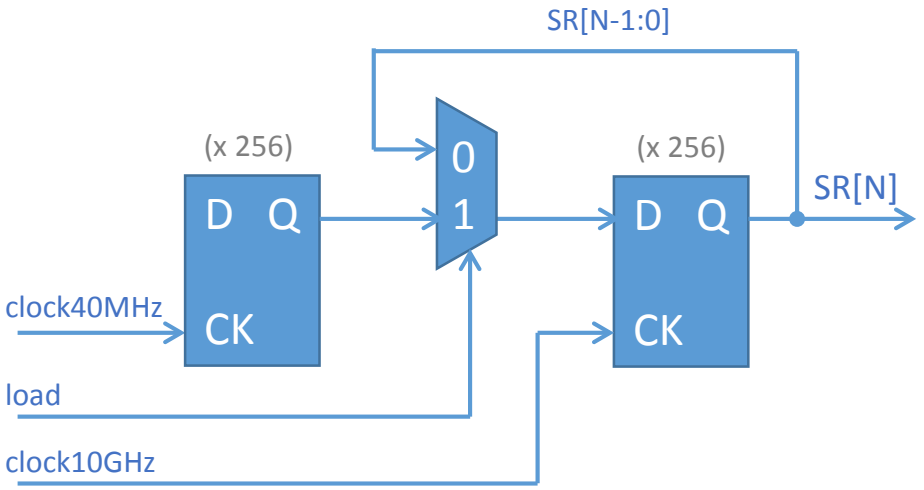

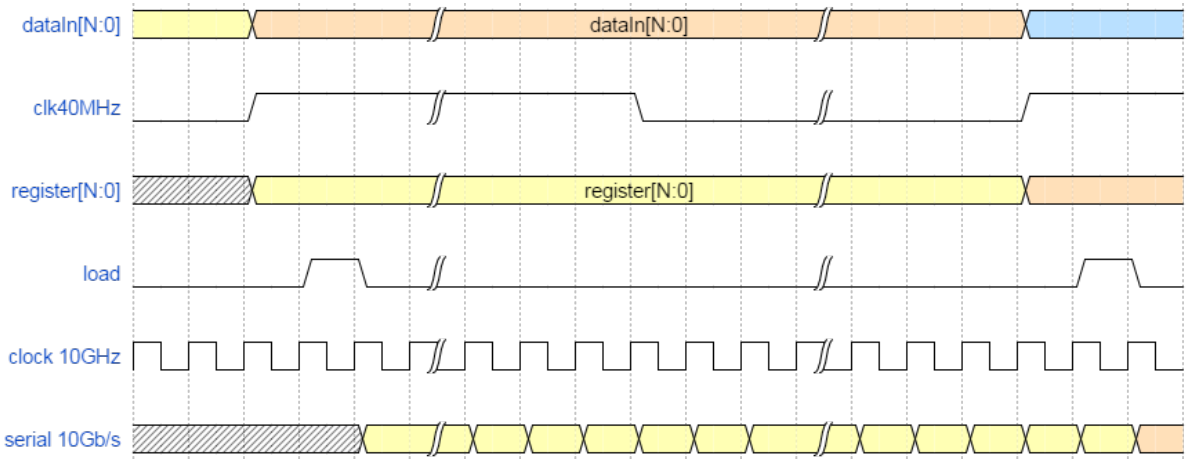

Should a Serializer be Implemented as a Shift Register?

- 10 Gb/s, LpGBT case:
  - 256 registers running at 40 MHz
  - 256 registers running at 10 GHz
  - 256 Multiplexers
    - MUX has to react to 100 ps pulse
    - MUX is reducing the setup time for the shift register
  - Timing between the “load” and the 10 GHz clock is critical!
- FFs:
  - Total of 512 need
  - All the FFs in the shift register need to be 10 GHz capable [not the input register]
- And how about power consumption?
  - Power [in CMOS] consumption is proportional to frequency

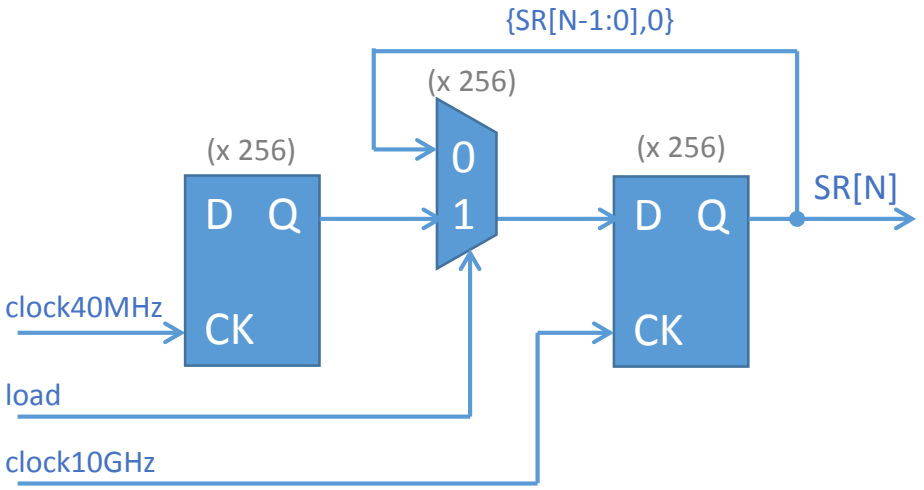

40 MHz

FF power in mW/GHz

$$P = 256 \cdot \frac{10.24\ GHz}{256} \cdot P_0 + 256 \cdot 10.24\ GHz \cdot P_0$$

$$\boxed{P = (256 + 1) \times 10.24\ GHz \times P_0}$$

P(LpGBT) = 257 x 260uW = 66.8 mW

“Multi – Level Serializer”

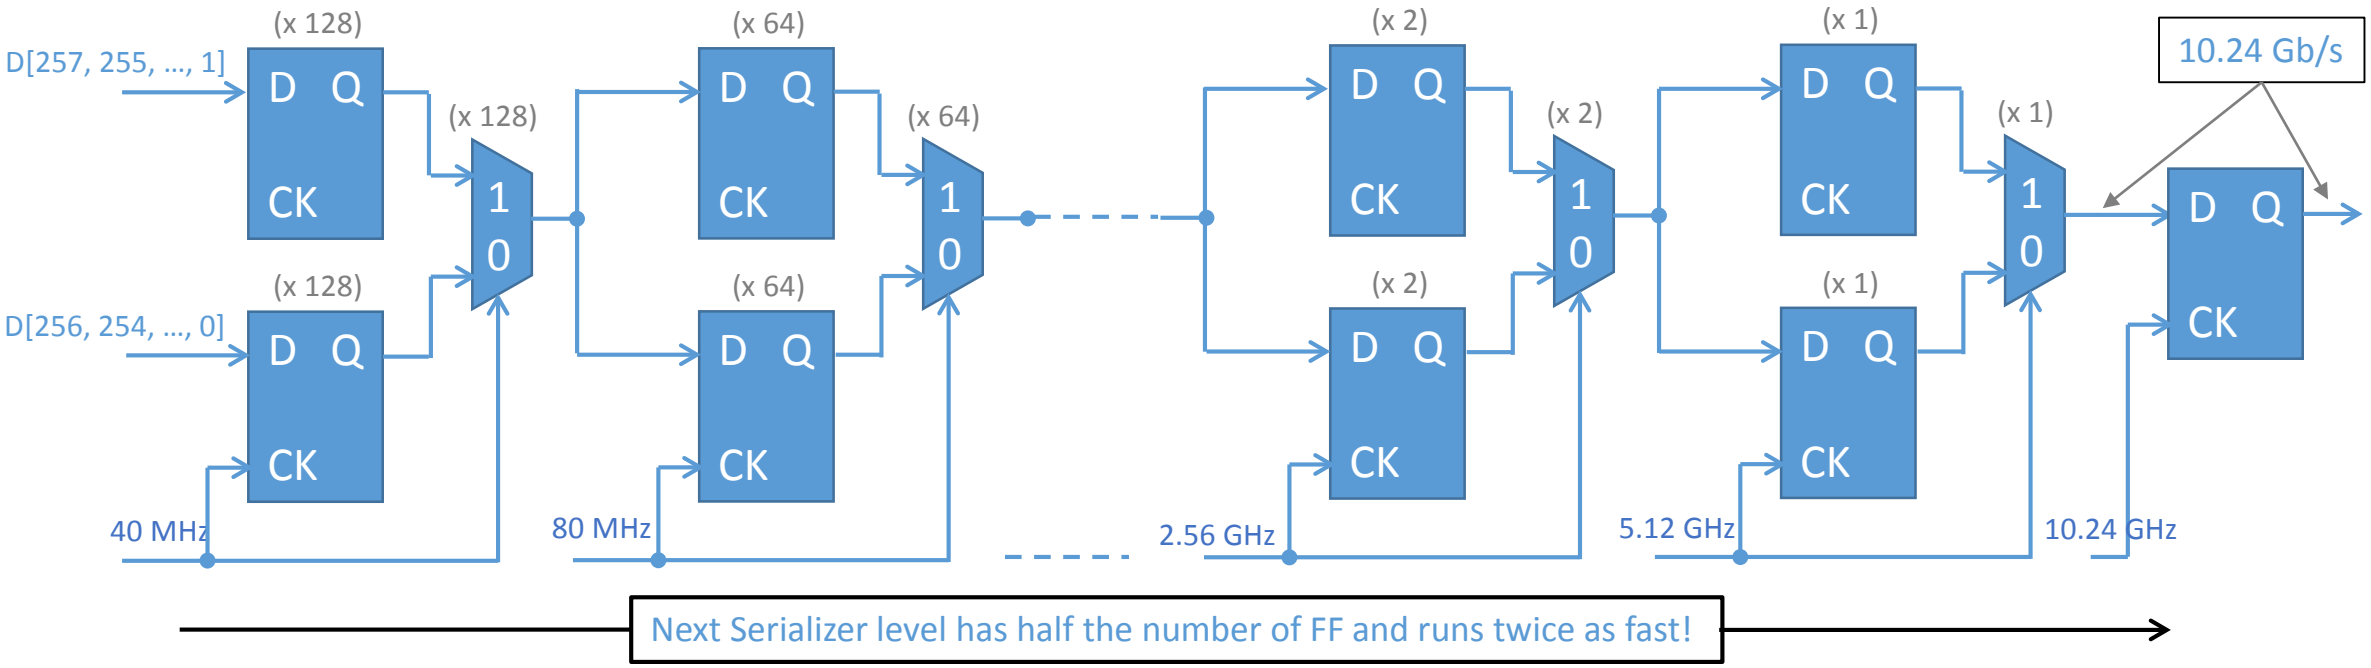

- Next level:
  - Has half the number of FF
  - Runs twice as fast
- The power per Serialization level remains constant:

$$P_L = 10.24\ GHz \times P_0$$

- Ten levels are needed to go from 40 MHz to 10.24 GHz
- The Serializer power is:

$$P = 10 \times 10.24\ GHz \times P_0$$

- Although this architecture requires “twice” as much FFs (1023), it only requires 10/257 = 3.9% of the power consumption!!!

LpGBT Serializer

- Since the clock frequency increases with the serialization level, it is possible to optimize the speed-power at each level:
  - Not possible for the simple shift register
- And a similar power saving goes in the clock tree!
- Also only the last 2 / 3 stages are timing critical
- The LpGBT does not use the last “resampling” stage at 10.24 GHz.
- At the output of the last MUX the signal is already at 10.24 Gb/s
  - Double data rate

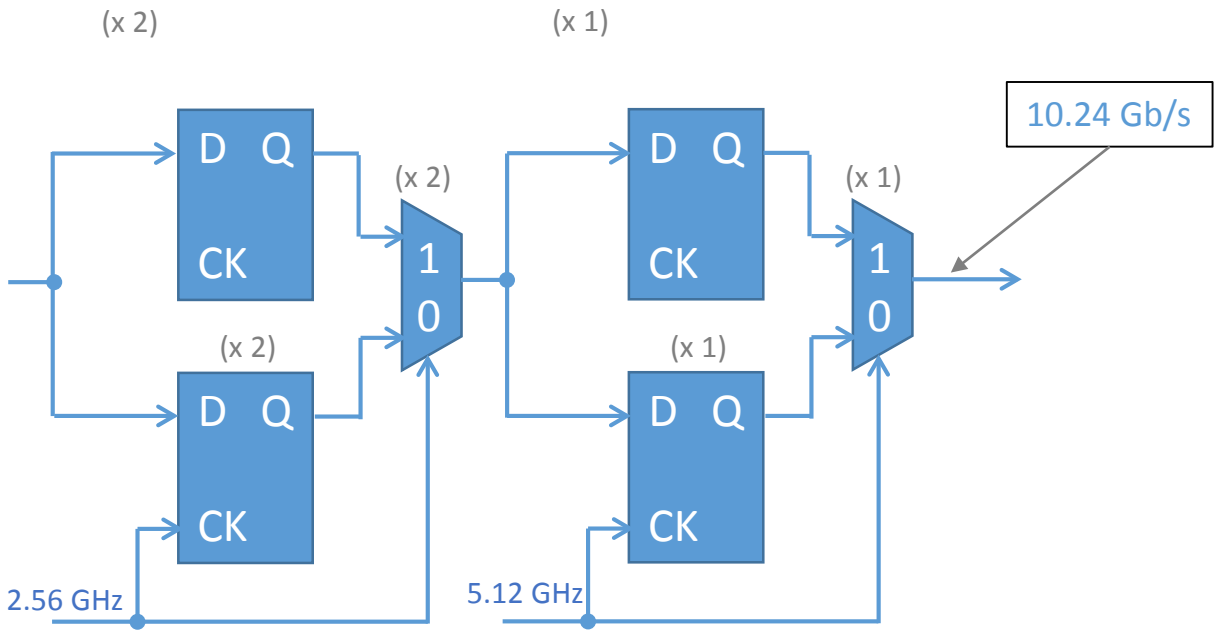

- Low jitter requires thus the last MUX to be fast and the 5.12 GHz clock to have “perfect” duty-cycle

5.12 and 10.24 Gb/s

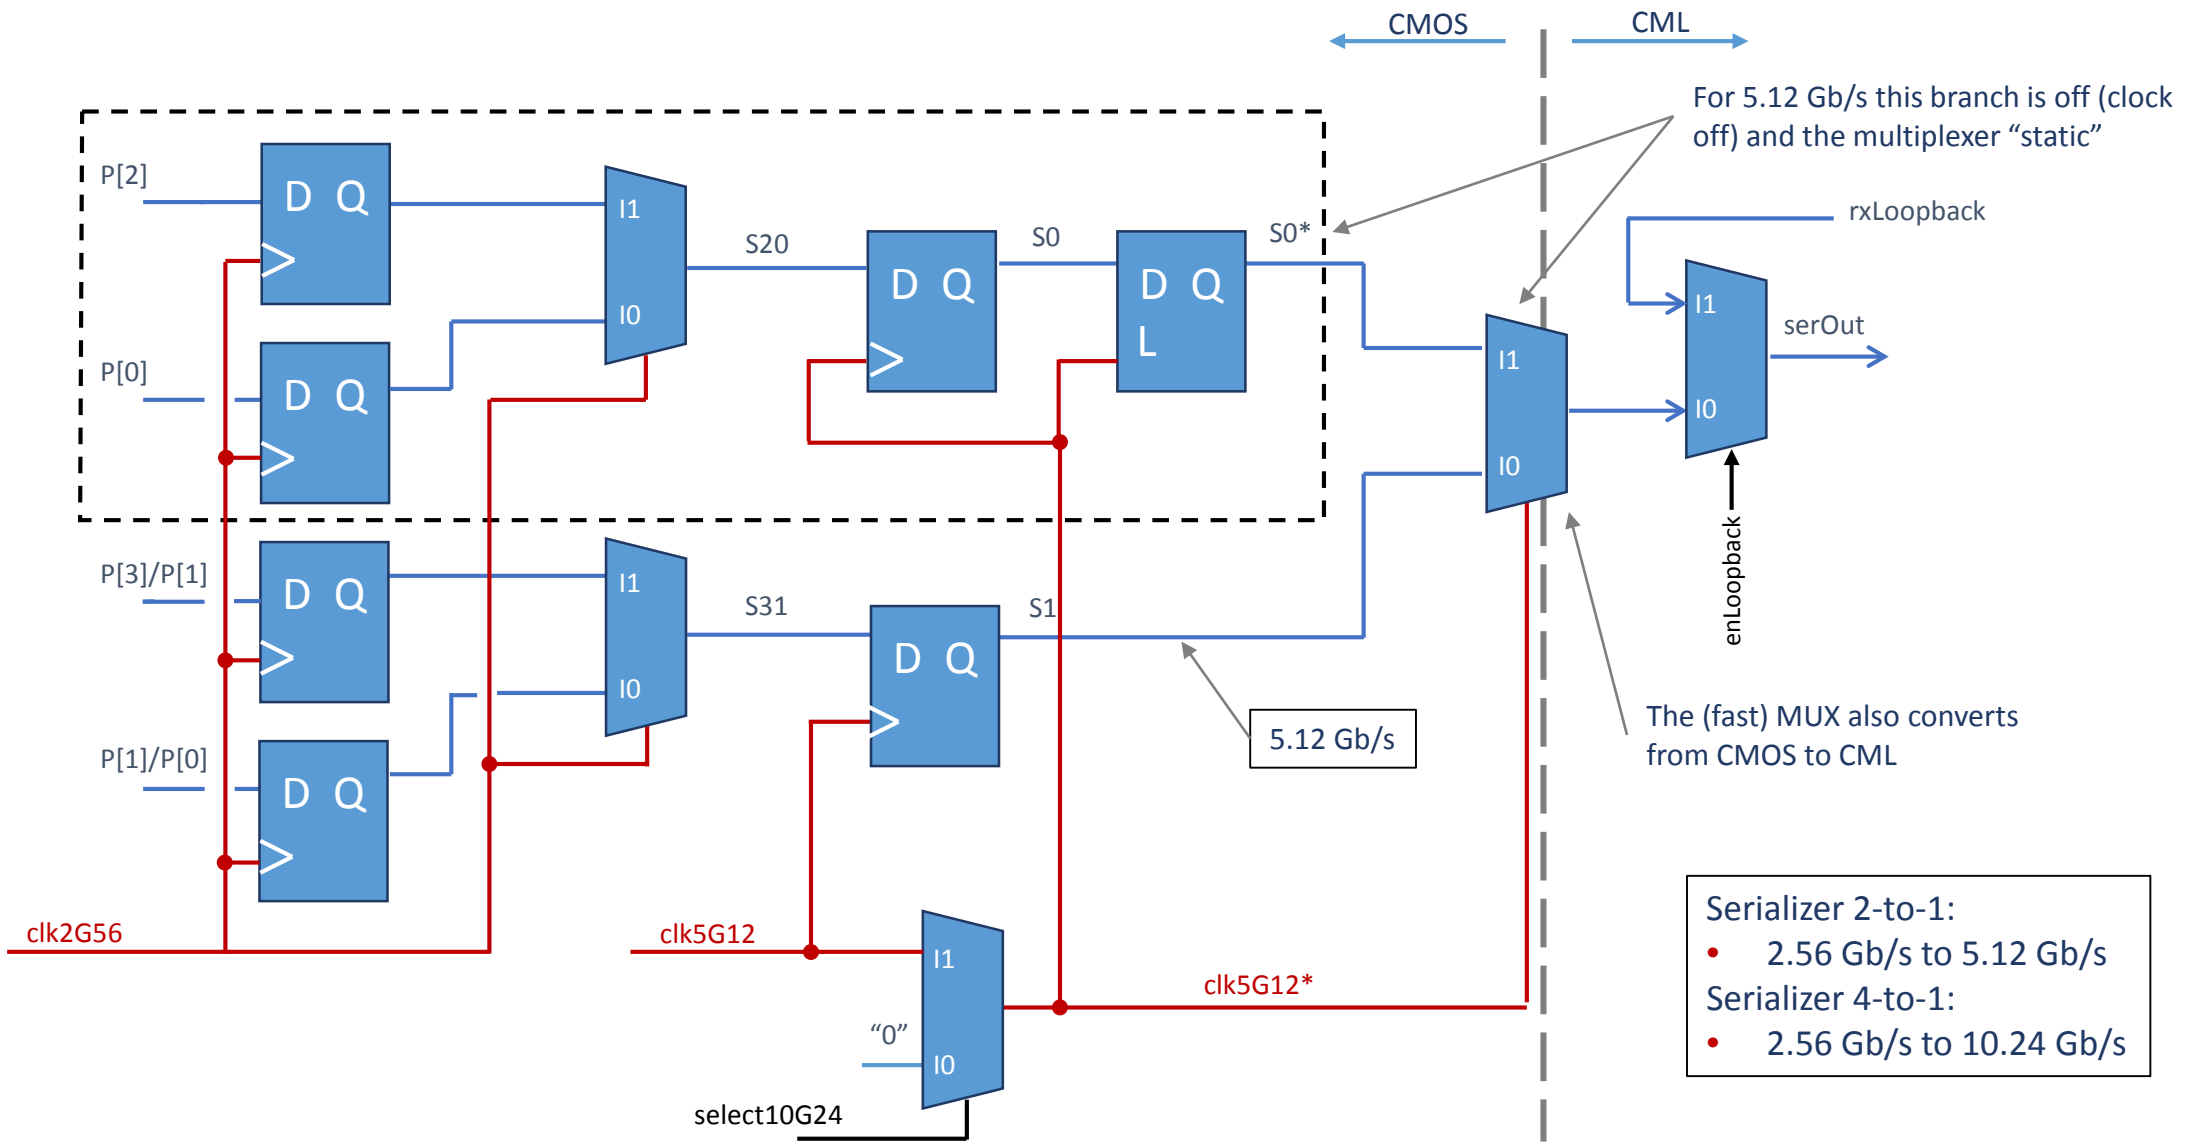

Serializer 2-to-1:
- 2.56 Gb/s to 5.12 Gb/s

Serializer 4-to-1:
- 2.56 Gb/s to 10.24 Gb/s

Fast Multiplexer

- Basically an AOI gate!
- Advantages:
  - Simplicity
  - Speed potential of CML
- Disadvantages:
  - “Perfect” symmetry between select and ~select required for low duty-cycle distortion

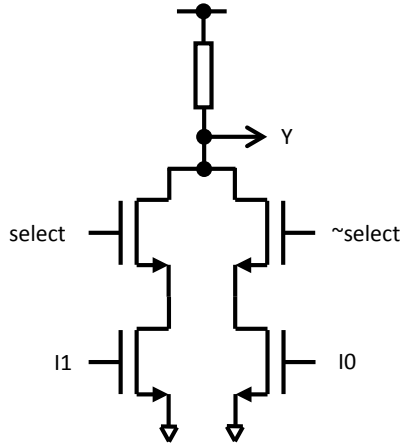

Y = ~(I1 & select + I0 & ~select)

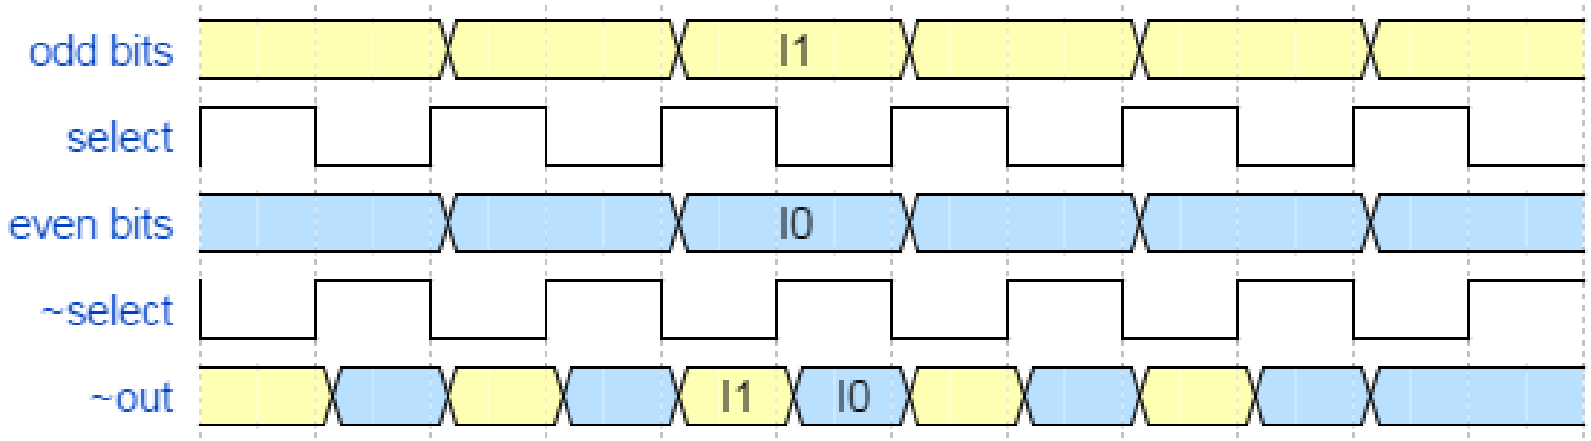

Fast Multiplexer

- Advantage:
  - A “natural” building block for “double data rate”
- Disadvantages:
  - Data dependent jitter
  - Amplitude “badly defined”

Voltages in these nodes depends on the history of I1 and I0

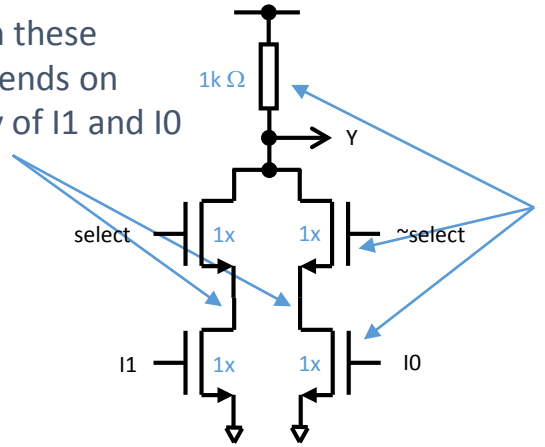

Amplitude depends on the relative “strength” of the resistors and the transistors

“Typical corner”

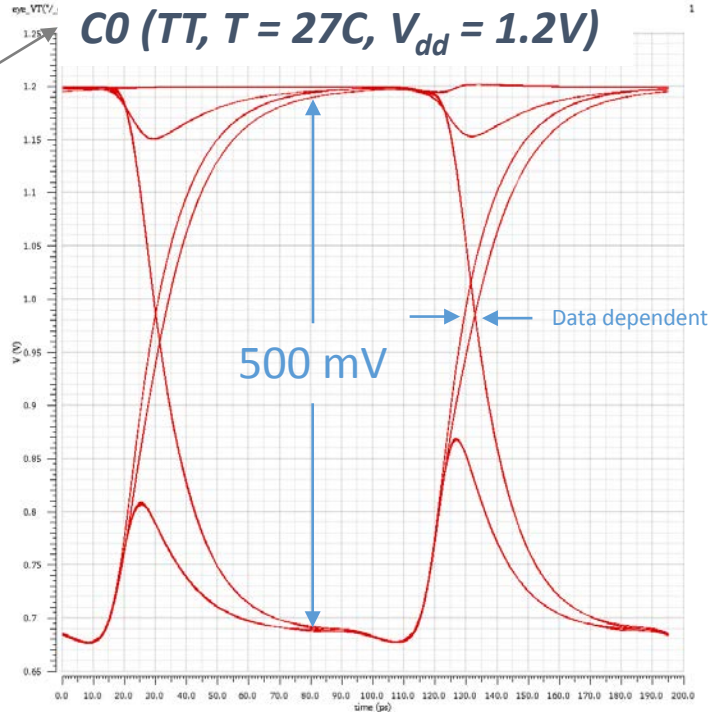

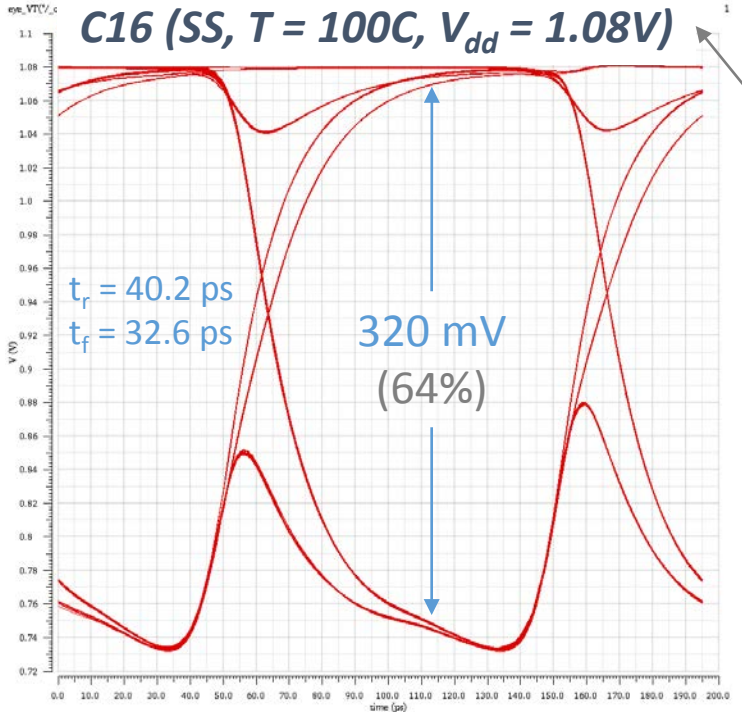

“Worst corner”

Signal Anatomy

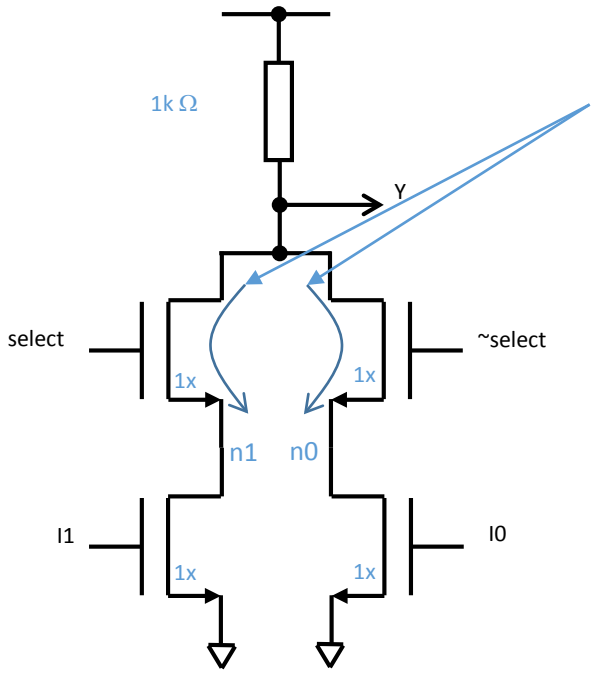

When “select” or “~select” transistors turn on, there is charge sharing between node “Y” and nodes “n0” or “n1”, pulling node “Y” down. The amount of charge shared depends on the history of nodes “n0” and “n1”.

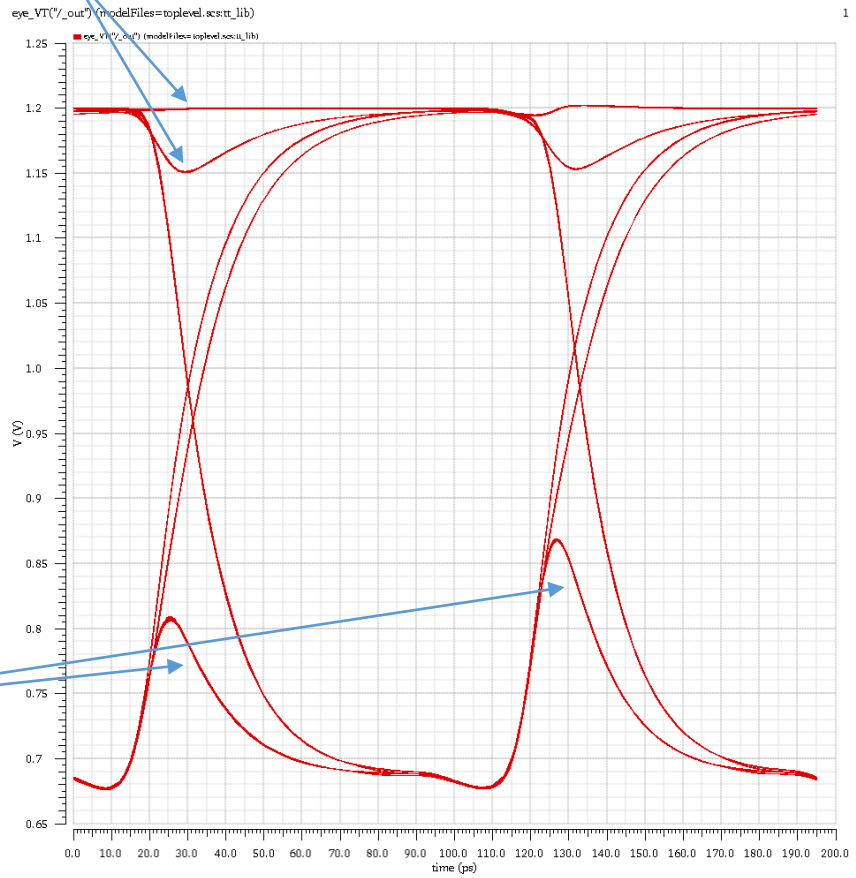

For the short period that both “select” and “~select” transistors are weakly turned on, the resistor pulls node “Y” high!
The height depends on the states of “I0” or “I1” inputs.

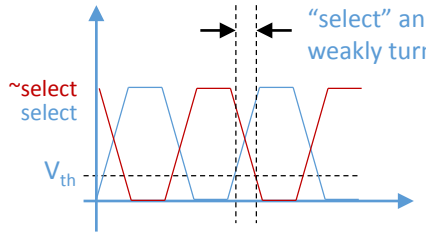

Improving Data Dependent Jitter

History in theses nodes is cleared while the branch is “inactive” reducing the data dependent jitter

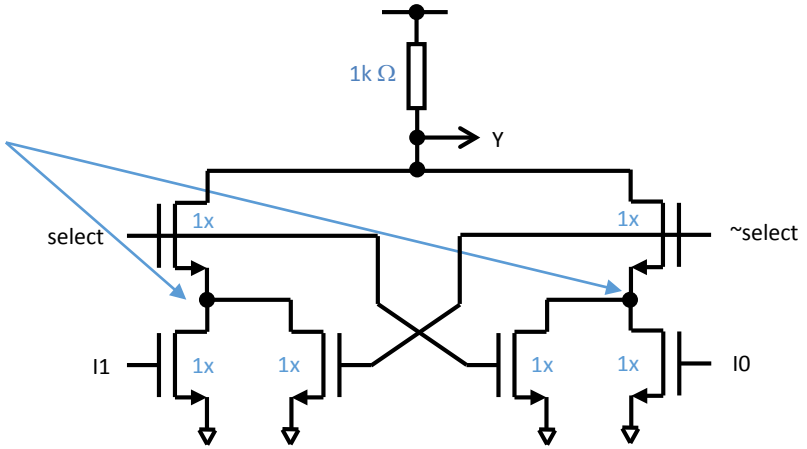

Always the same amount of charge sharing!

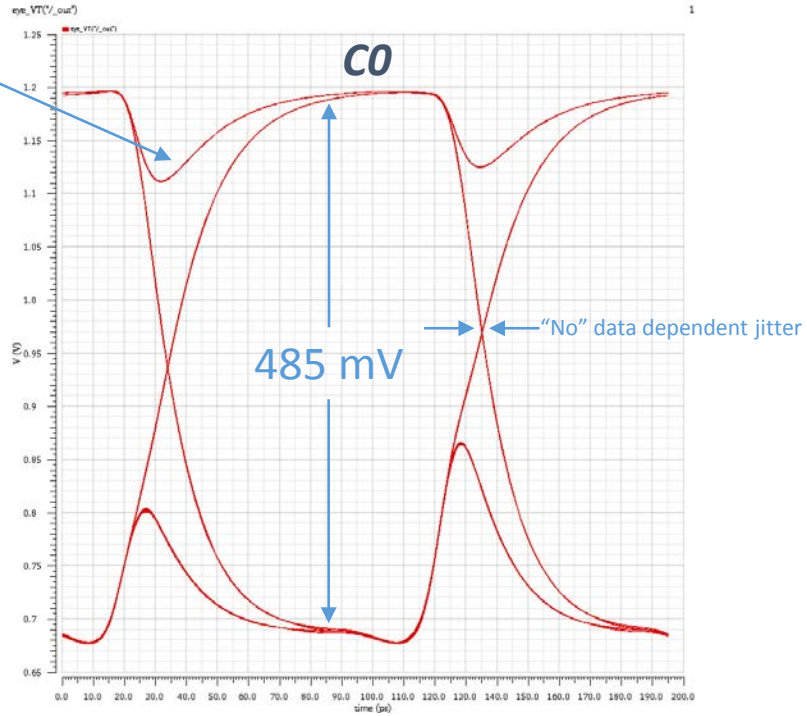

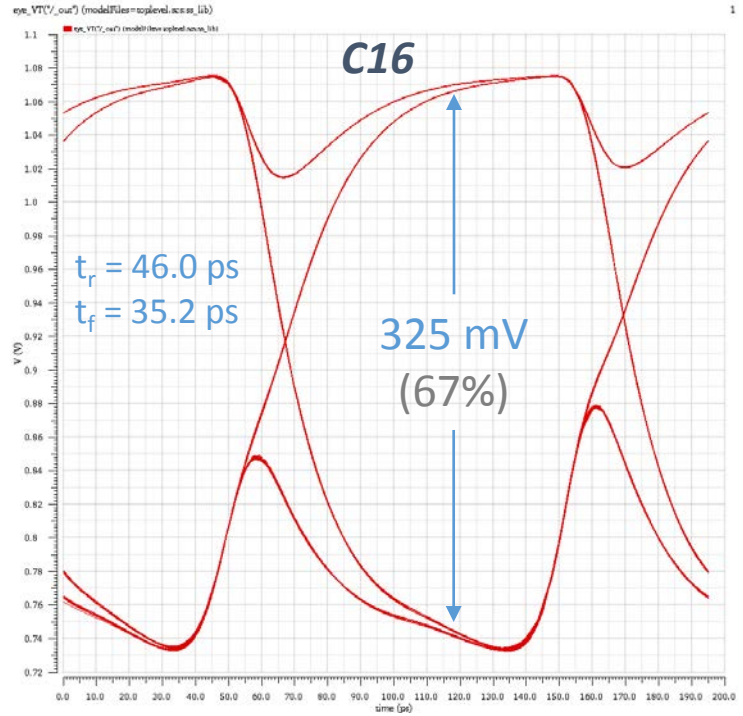

“Amplitude Stabilization”

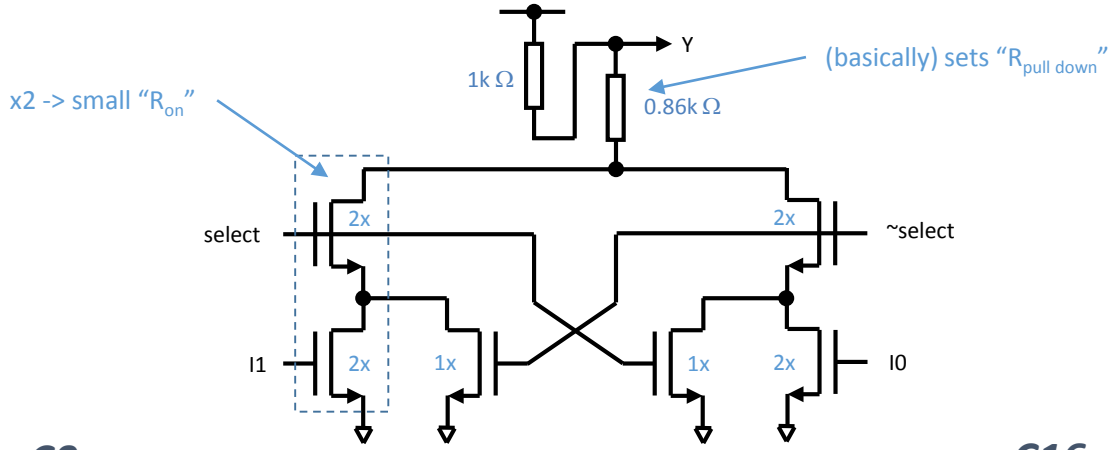

C0

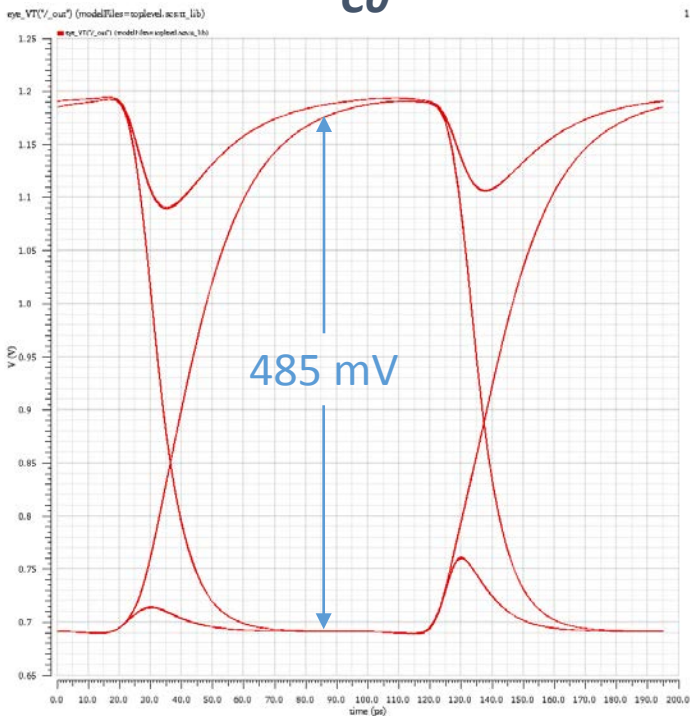

C16

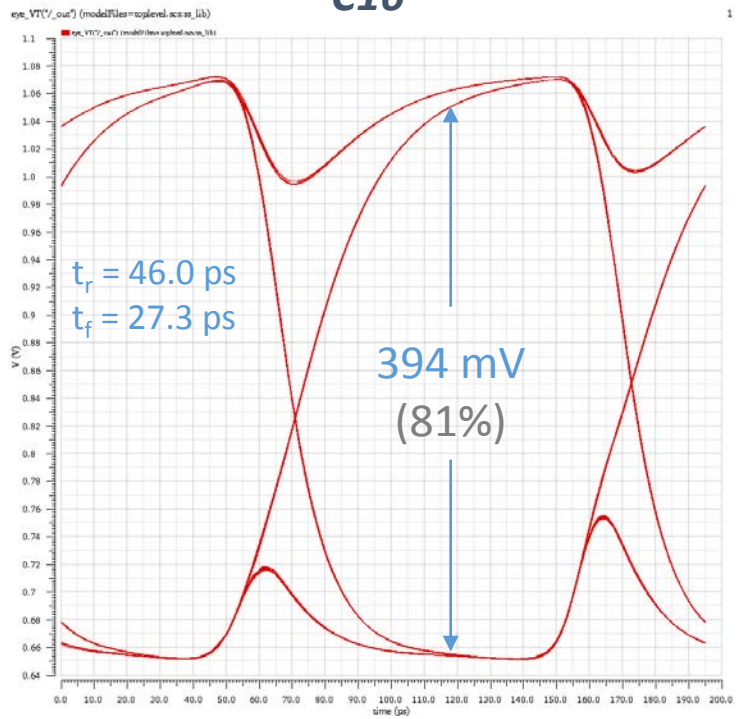

5.12 and 10.24 Gb/s

- For 10.24 Gb/s operation “select” and “~select” alternately select the “I1” and “I0” inputs
- For 5.12 Gb/s operation the “select” input is kept at “0” and the “~select” at “1”. The input “I0” passes thus inverted to the output.
- The circuit is single ended! To interface with the CML logic, two circuits used as “pseudo differential” are needed!

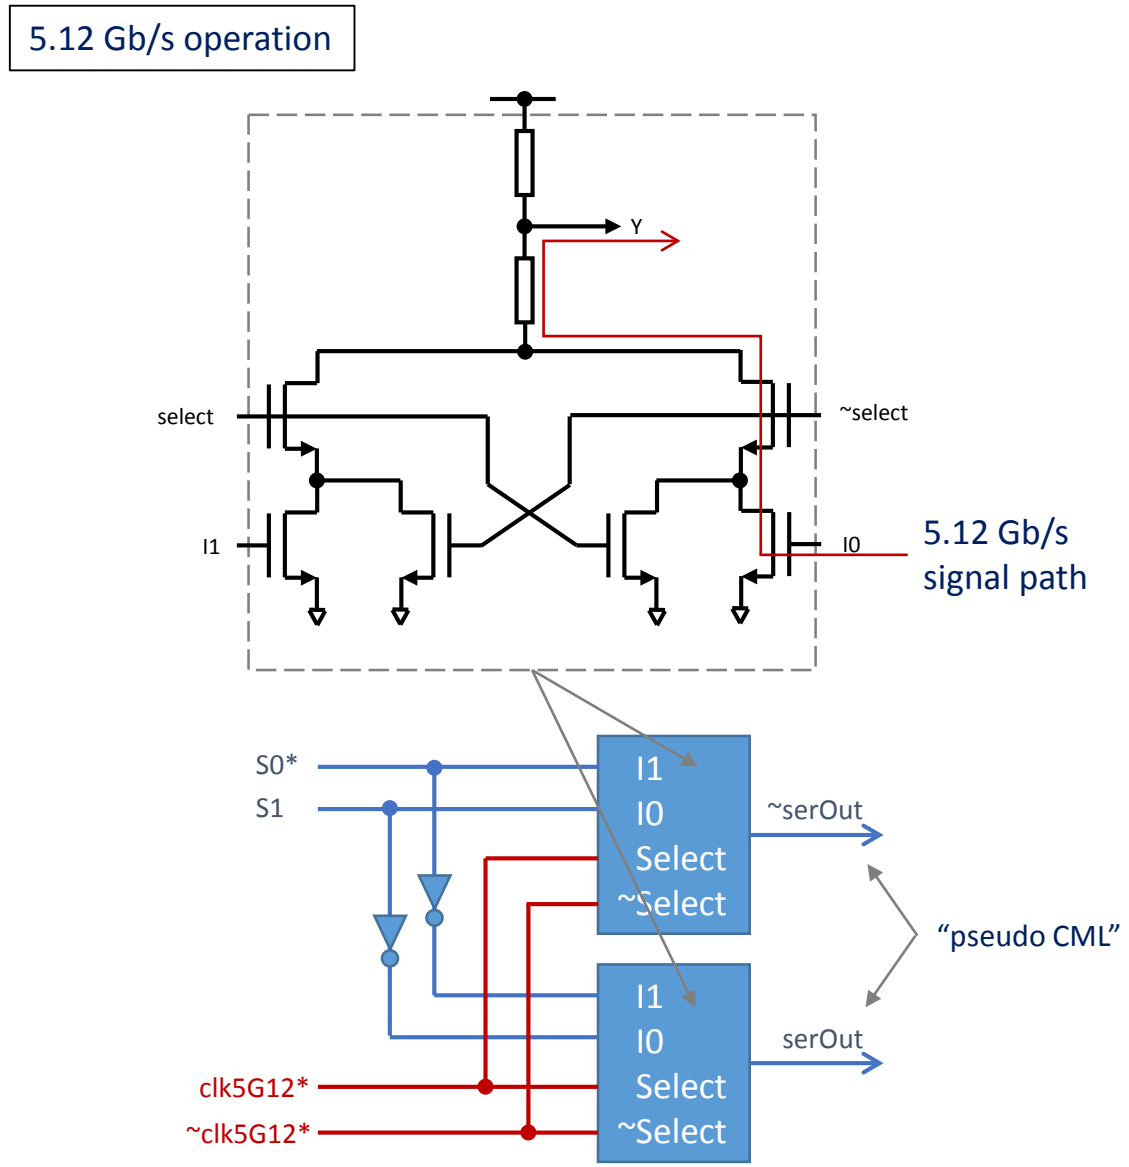

High Speed Line – Driver

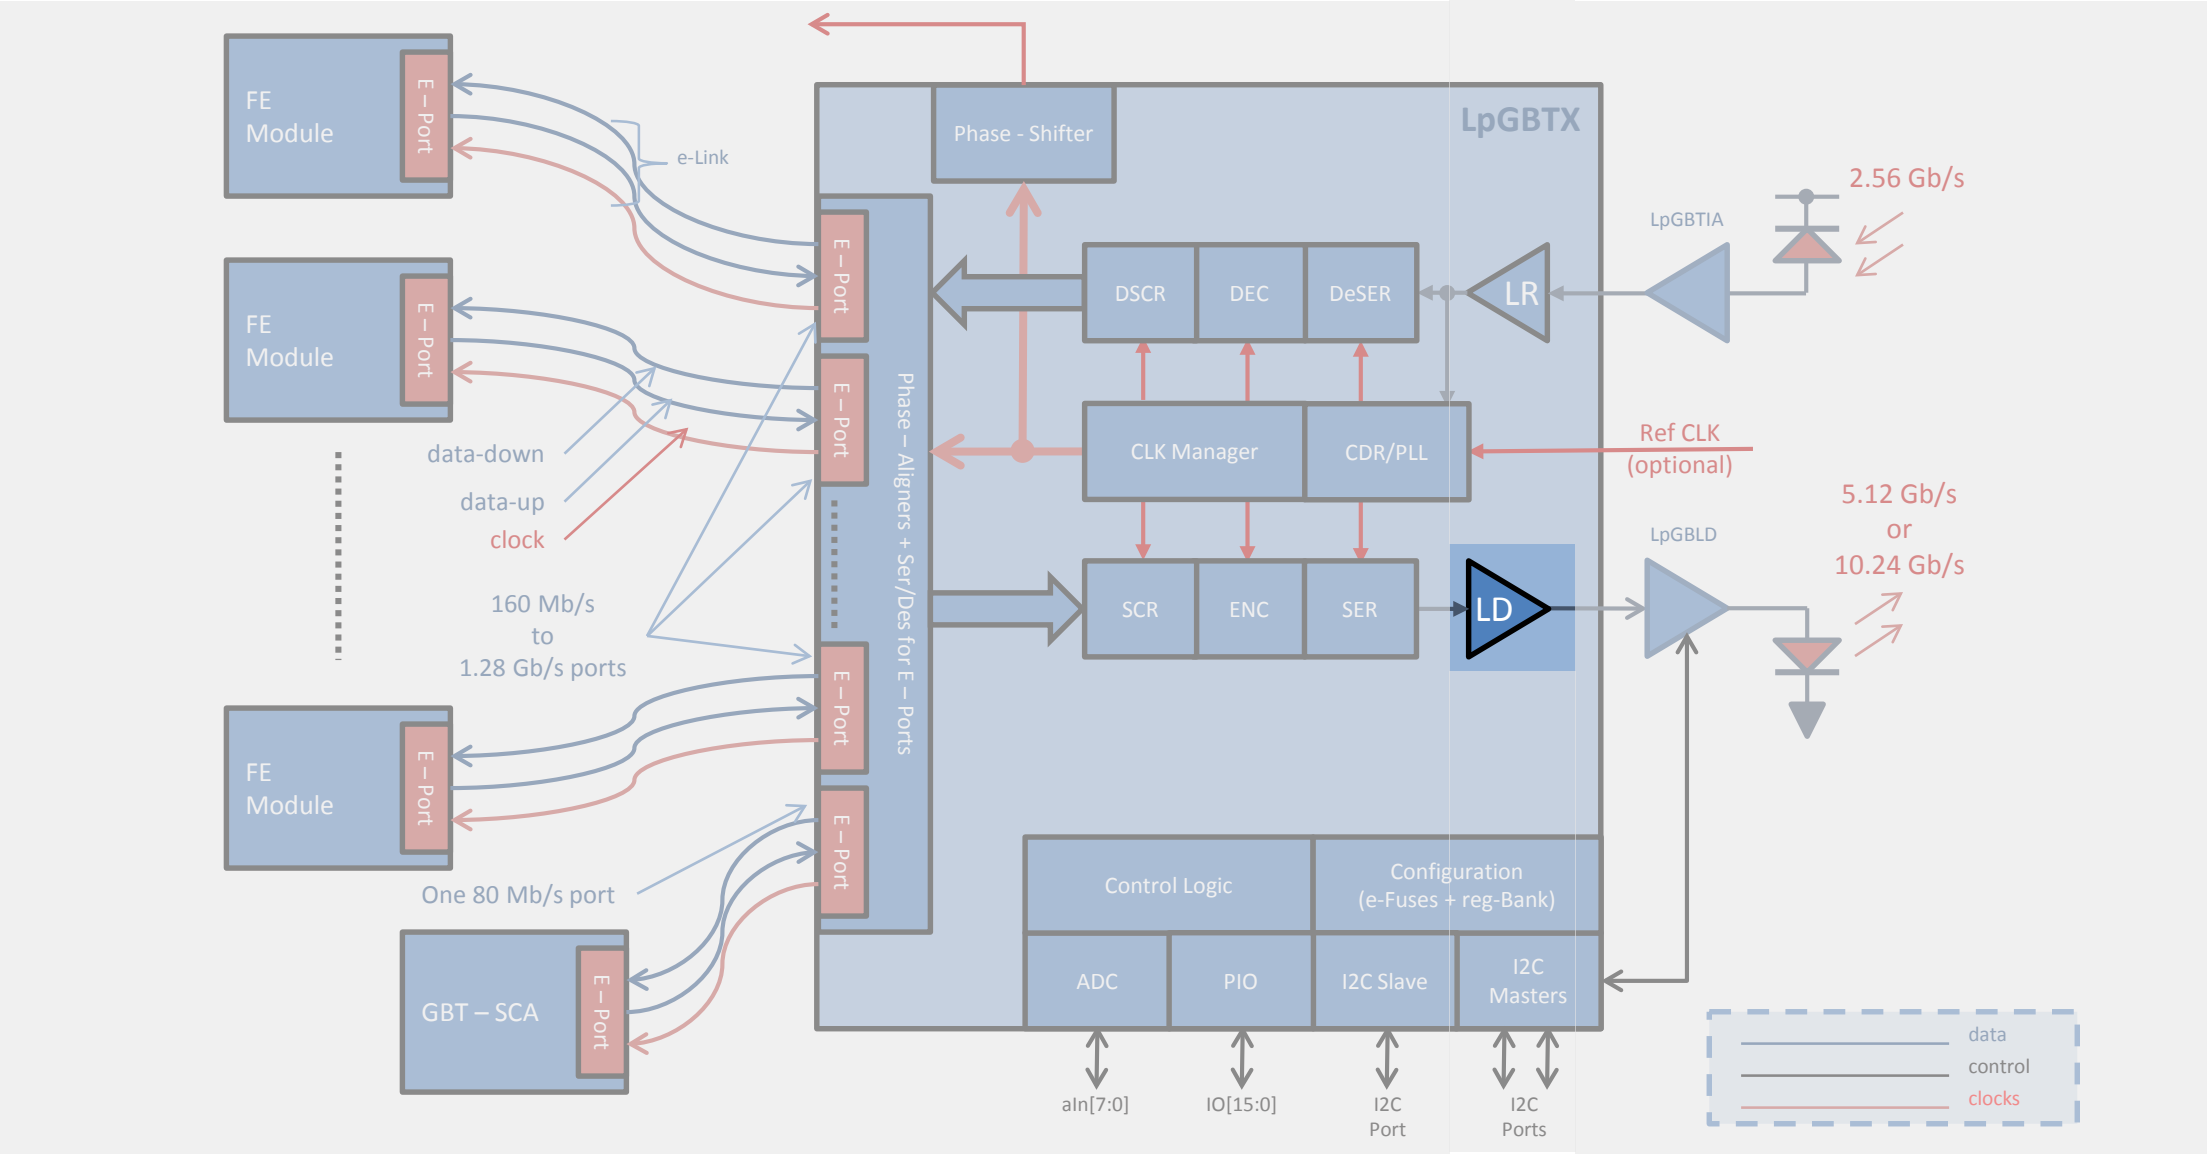

Line – Driver Topology

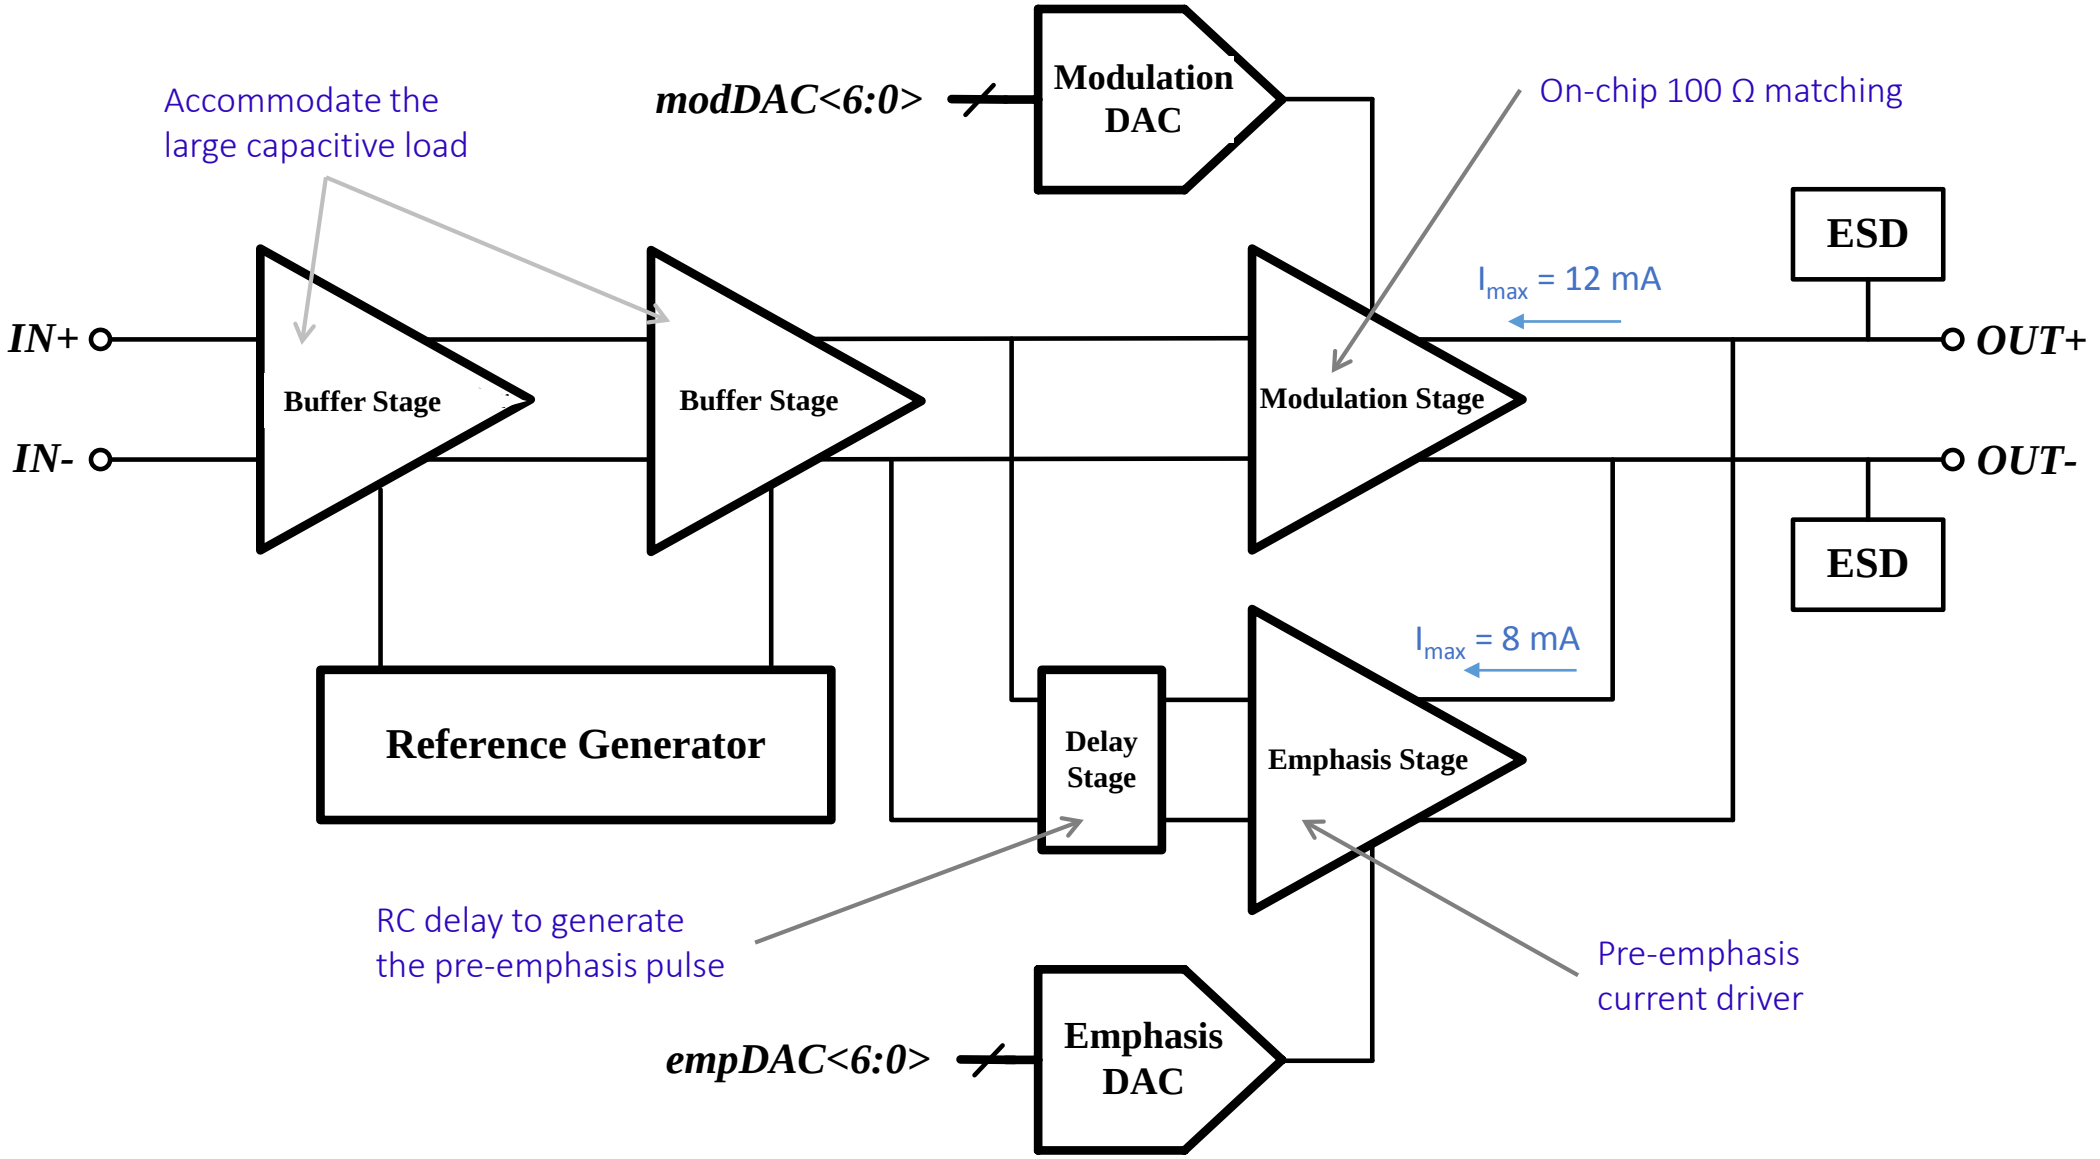

Bandwidth Broadening using Inductive Peaking (1/2)

- Add an inductor in series with the load resistor
- The bandwidth can be extended up to 1.85 times
- For optimum group delay response the BW gain is 1.6 times
- The circuit displays a second order transfer function
  - The frequency response is characterized by the ratio **"m"**

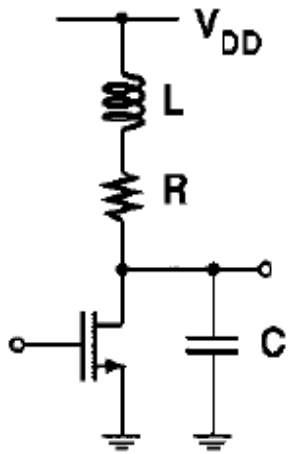

$$L = m.R^2.C$$

Factor m	Normalized f_{3dB}	Response
0	1.00	No shunt peaking
0.32	1.60	Optimum group delay
0.41	1.72	Maximally flat
0.71	1.85	Maximum bandwidth

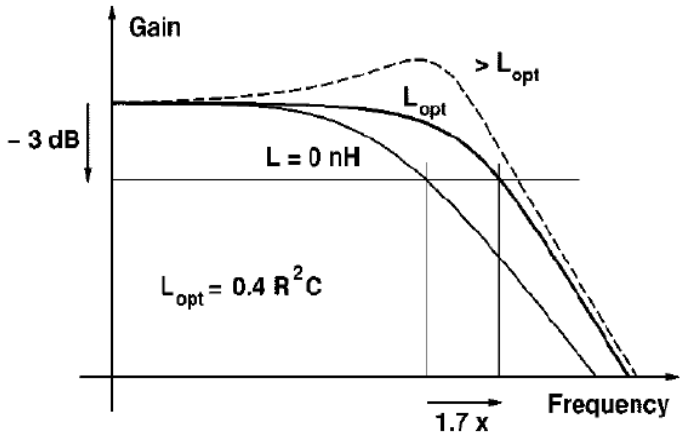

Bandwidth Broadening using Inductive Peaking (2/2)

Power optimization:

- For a CML stage the bandwidth can be enhanced by increasing the tail current:
  - By reducing R_L at fixed ΔV
- Or, for by using inductive peaking for a given current and R_L
- A specified bandwidth can thus be achieved with lower current if inductive peaking is used.

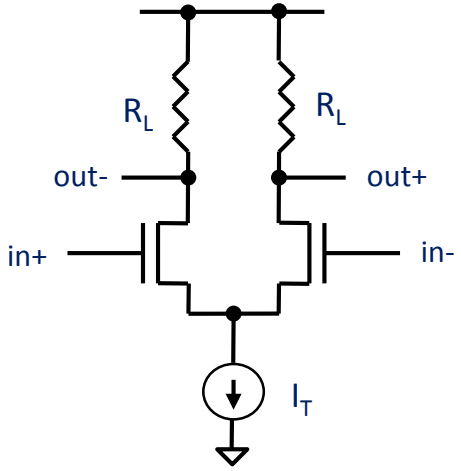

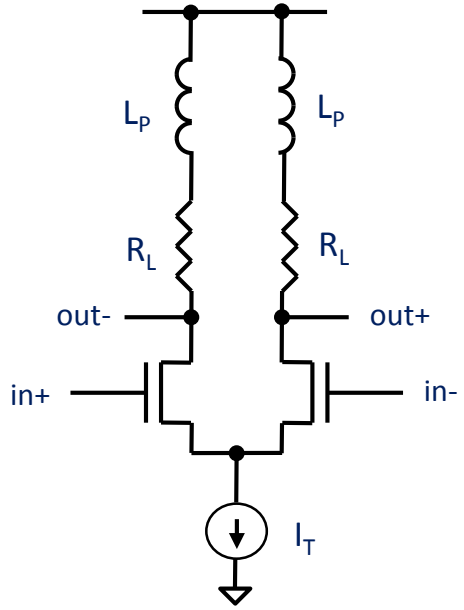

2nd Buffer Stage

R1/R2: 100Ω

M1/M2: 60µm/60nm

L1: 3.9nH

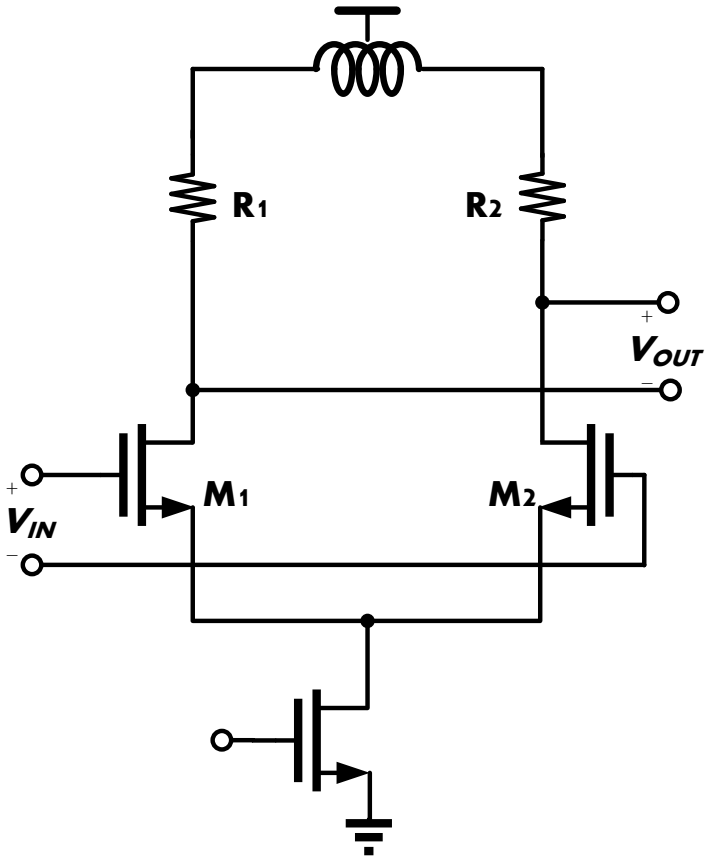

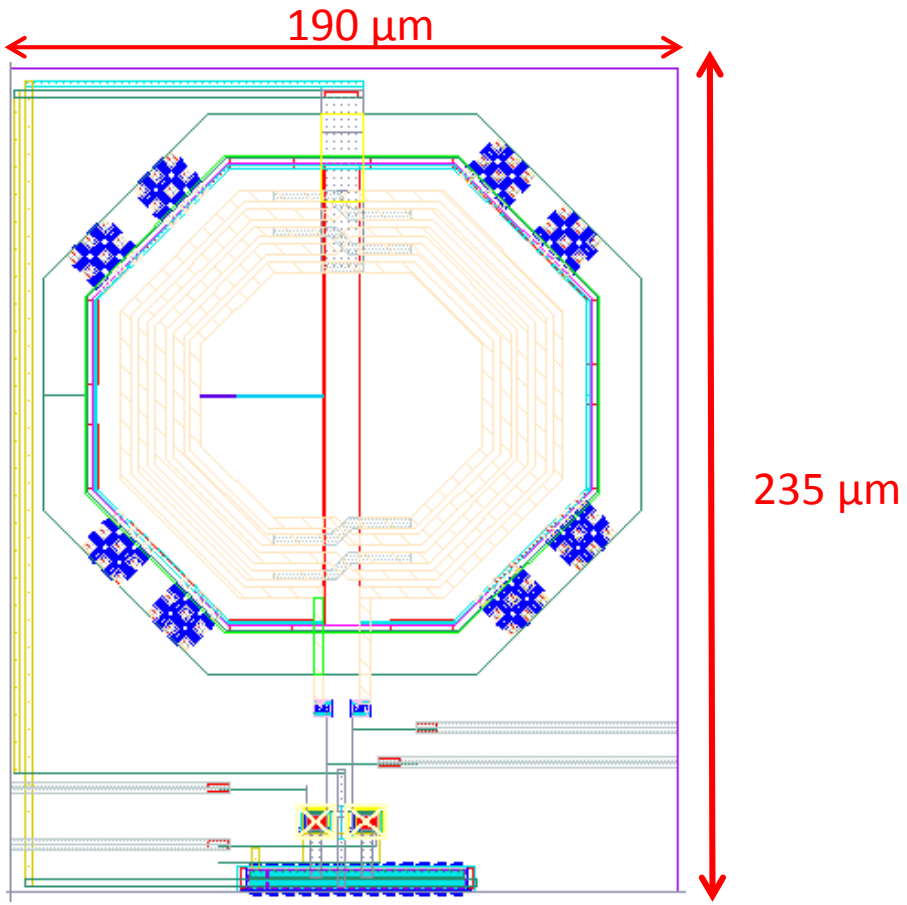

Modulation Stage

R_1/R_2: 50Ω

M_1/M_2: 60µm/60nm

L_1: 1.5nH

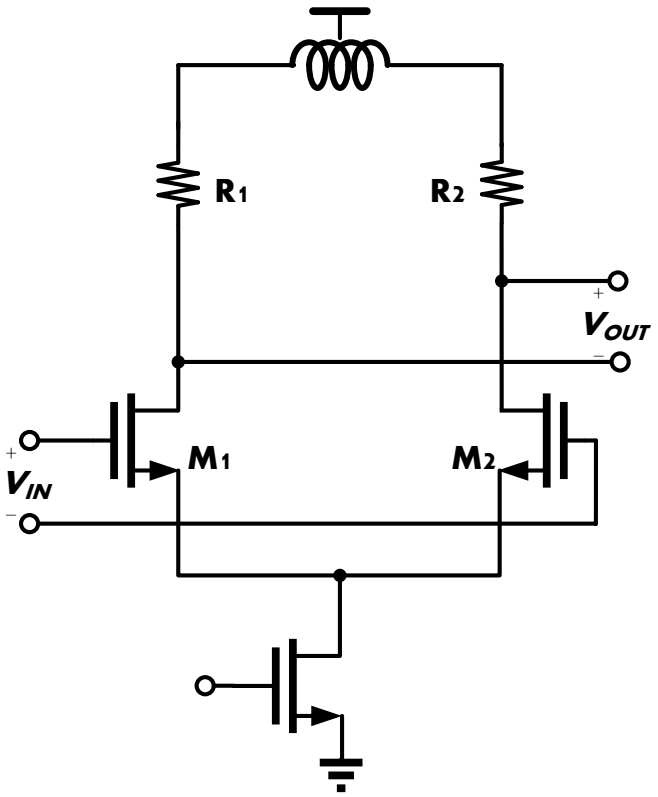

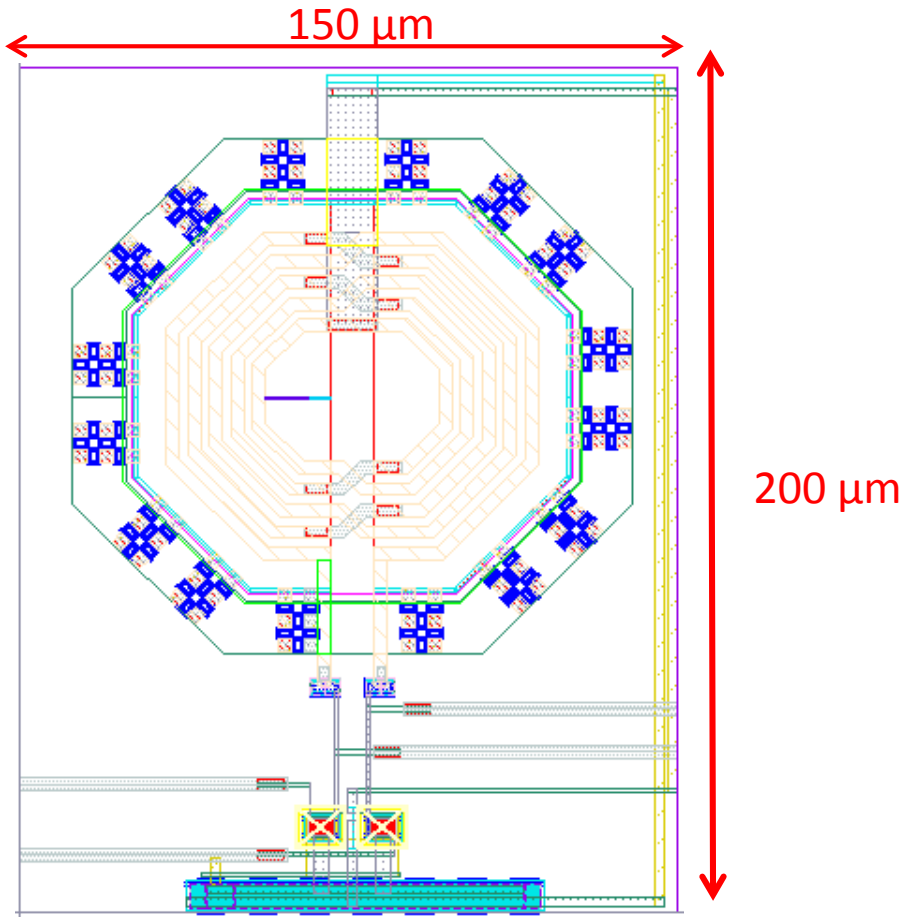

Delay “Stage”

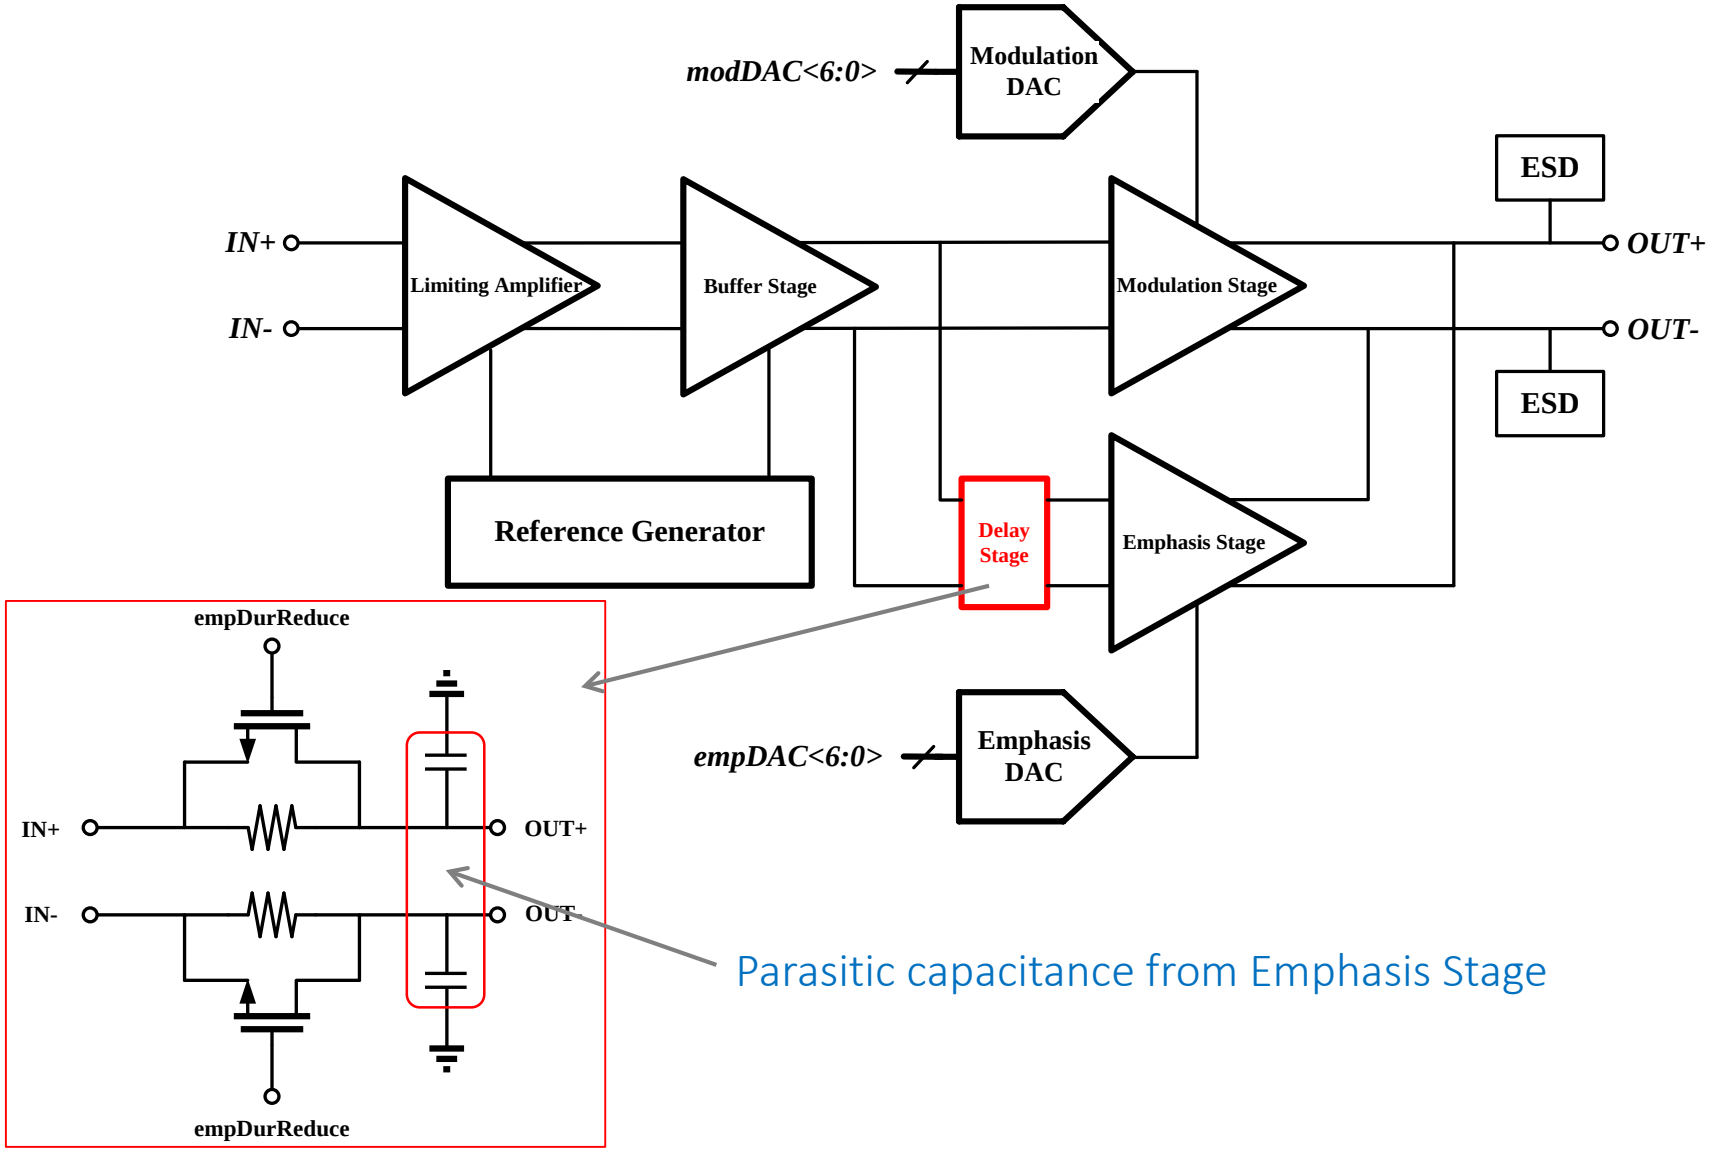

Parasitic capacitance from Emphasis Stage

10 Gb/s Post-layout Simulations

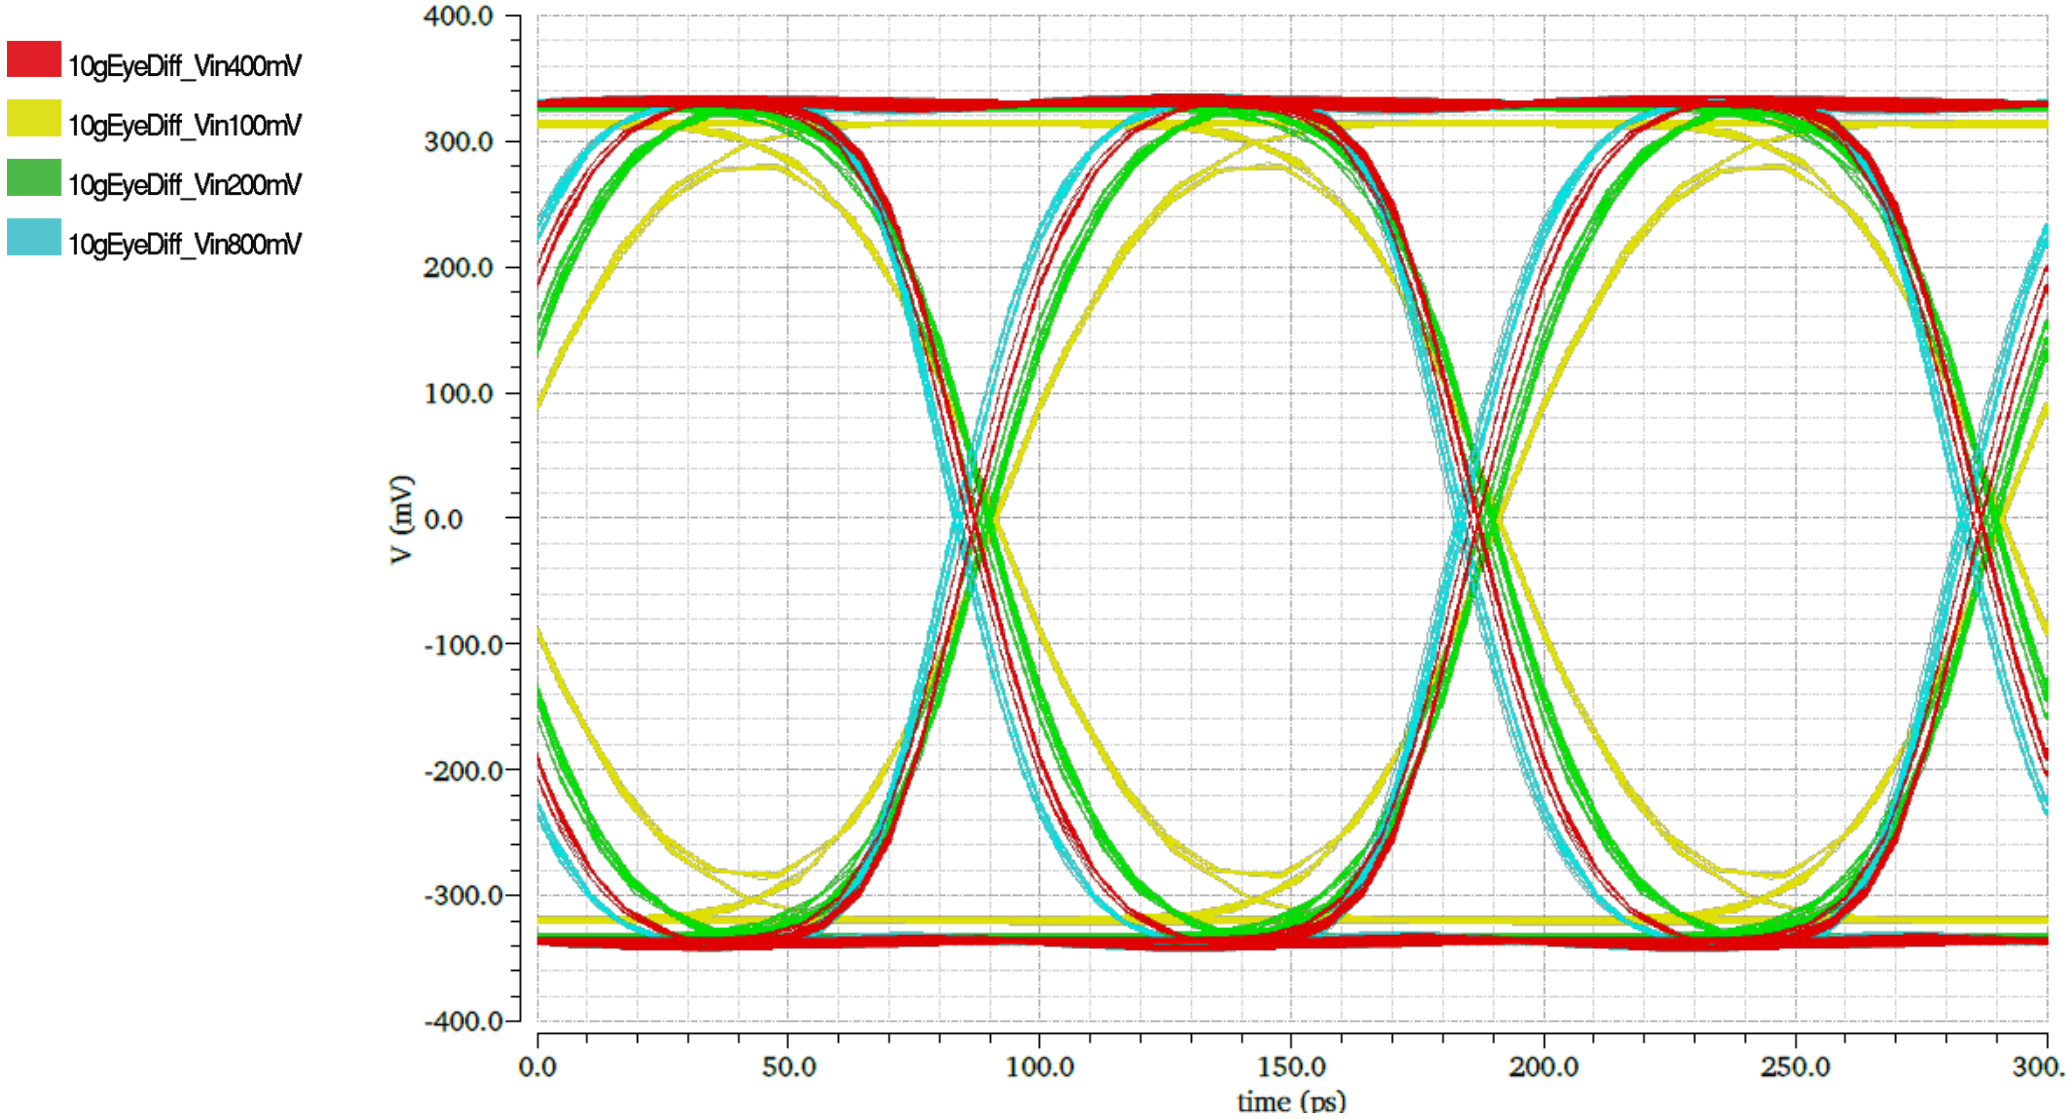

10 Gb/s Pre-Emphasis

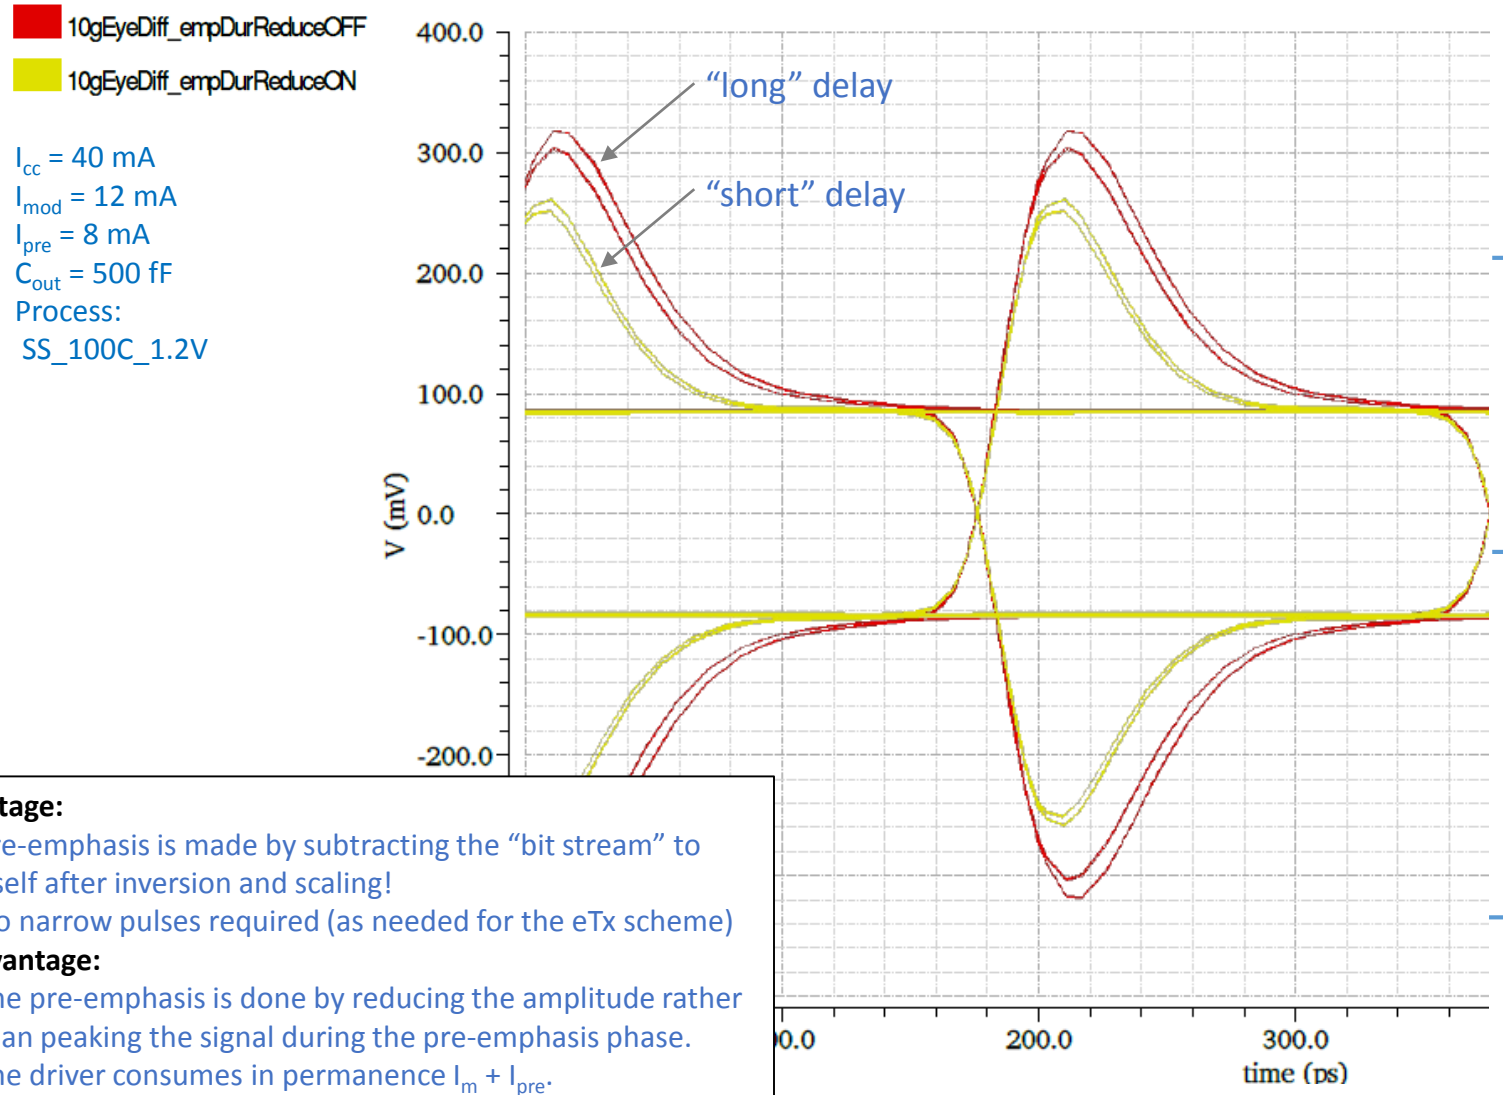

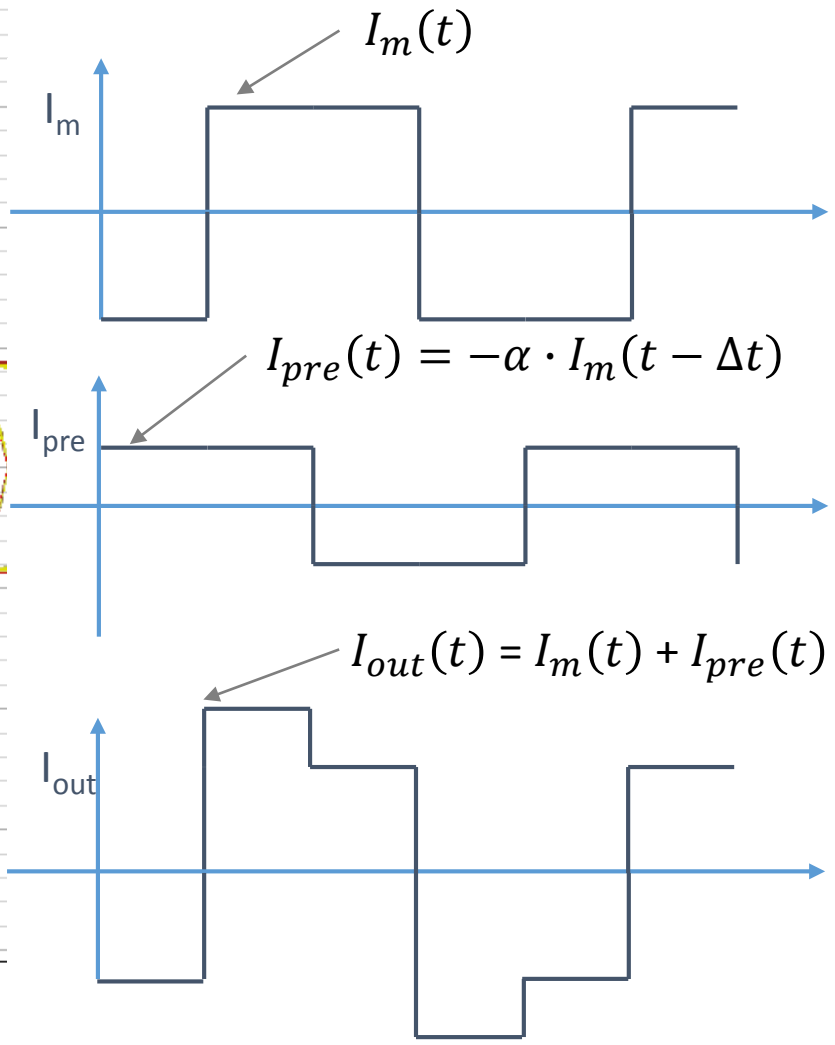

Advantage:

- Pre-emphasis is made by subtracting the "bit stream" to itself after inversion and scaling!
- No narrow pulses required (as needed for the eTx scheme)

Disadvantage:

- The pre-emphasis is done by reducing the amplitude rather than peaking the signal during the pre-emphasis phase.
- The driver consumes in permanence $I_m + I_{pre}$.

High-Speed Line Receiver

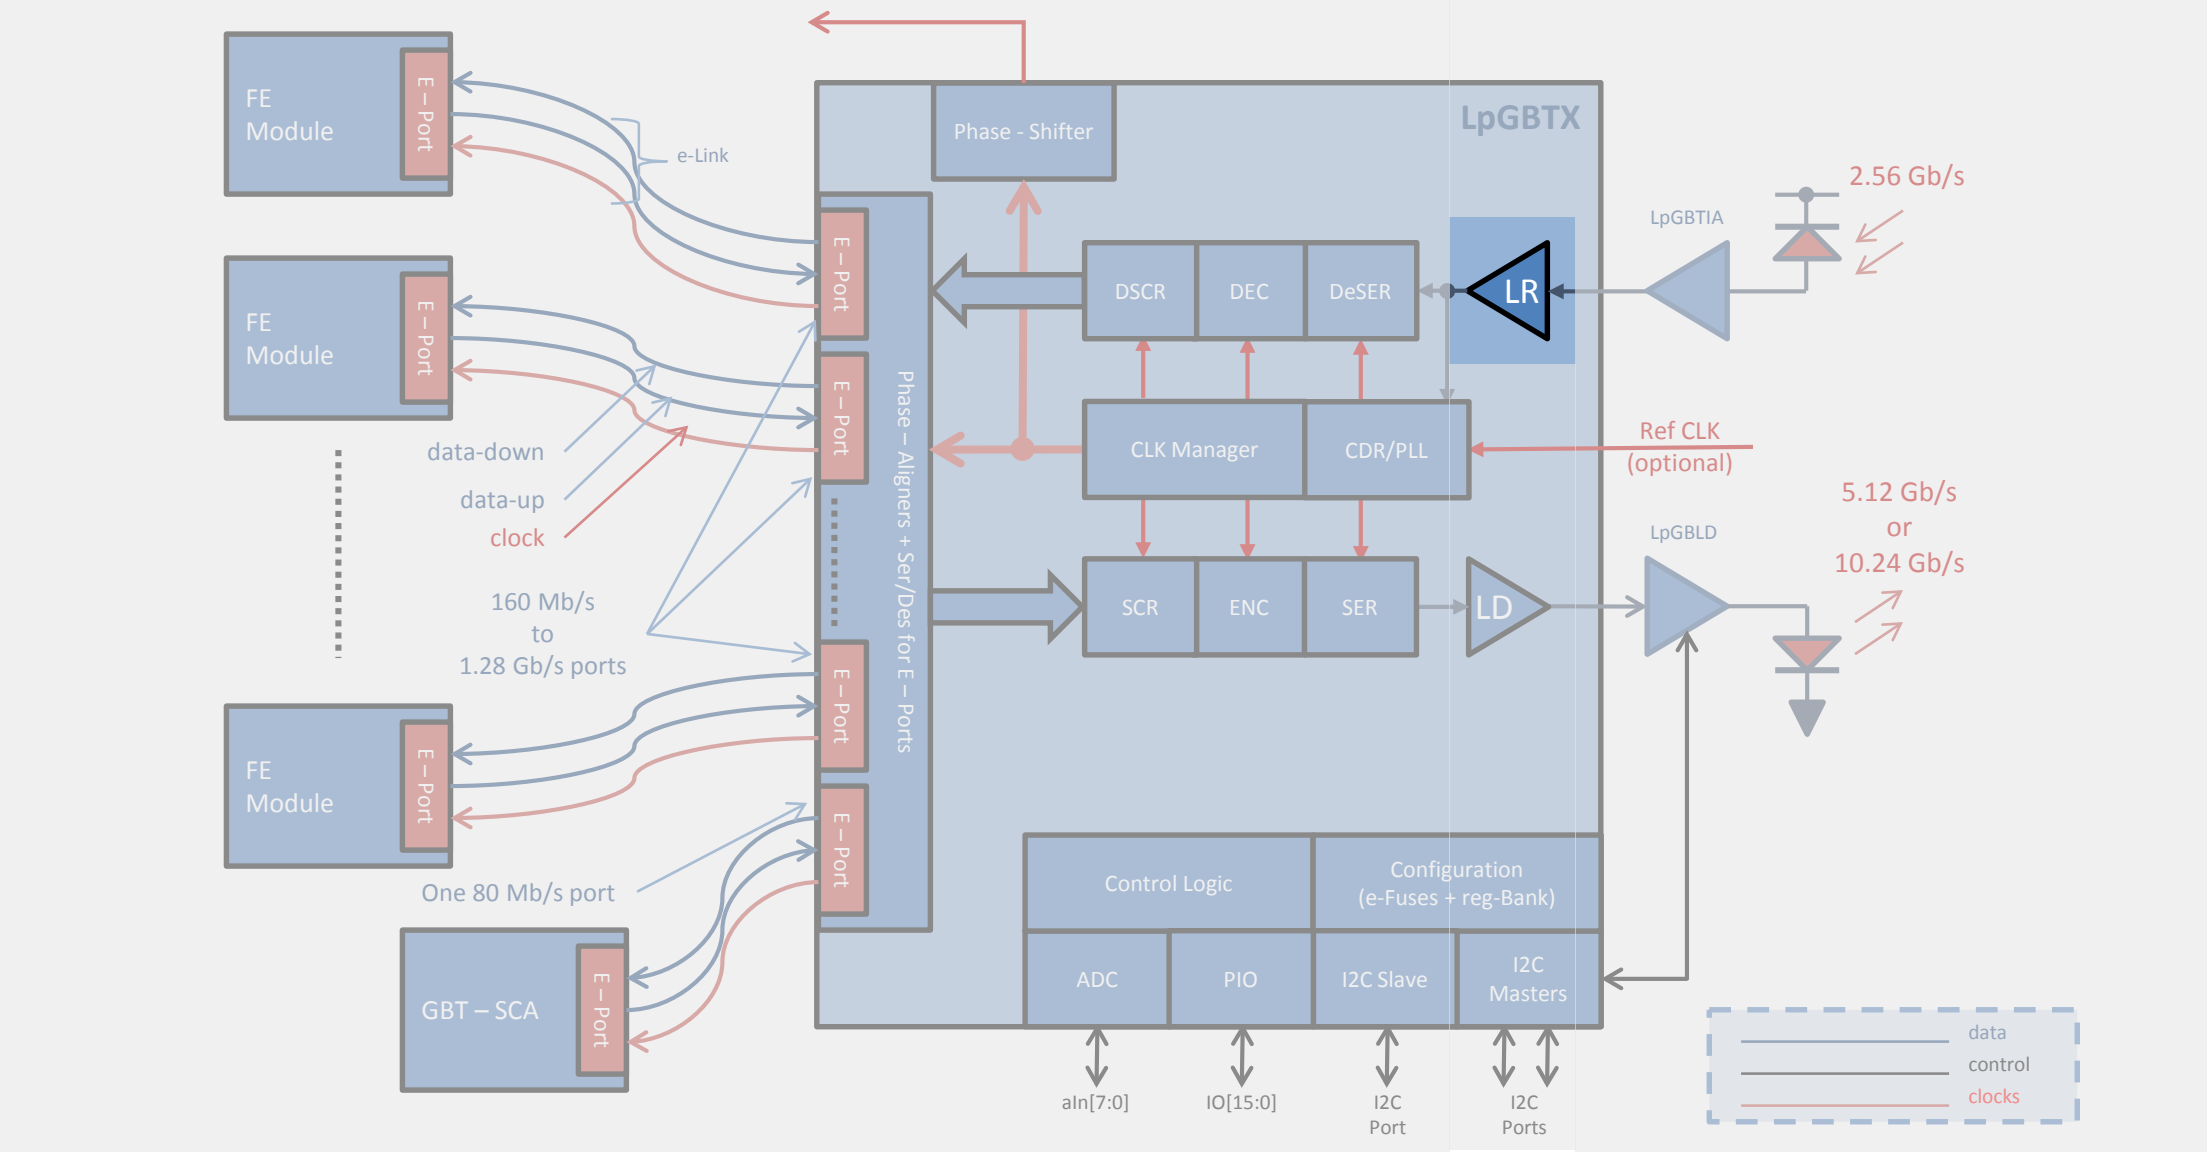

Equalization

5 cm Differential Strip Line in FR4

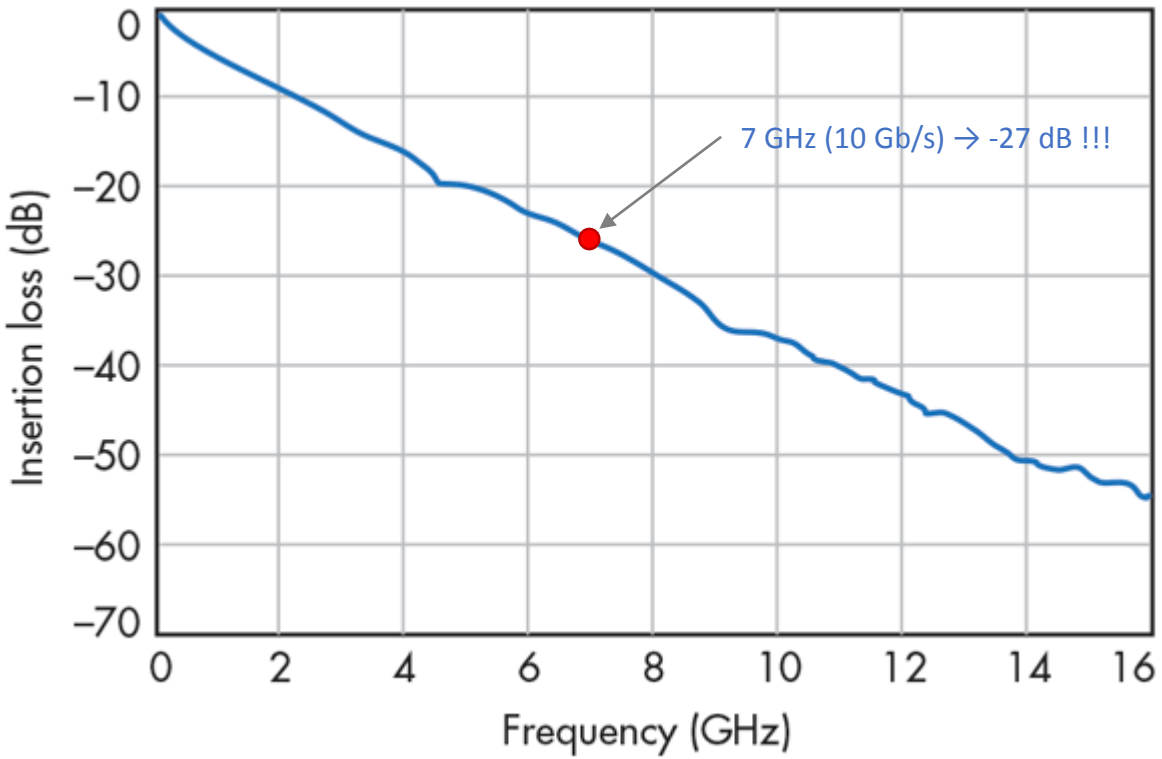

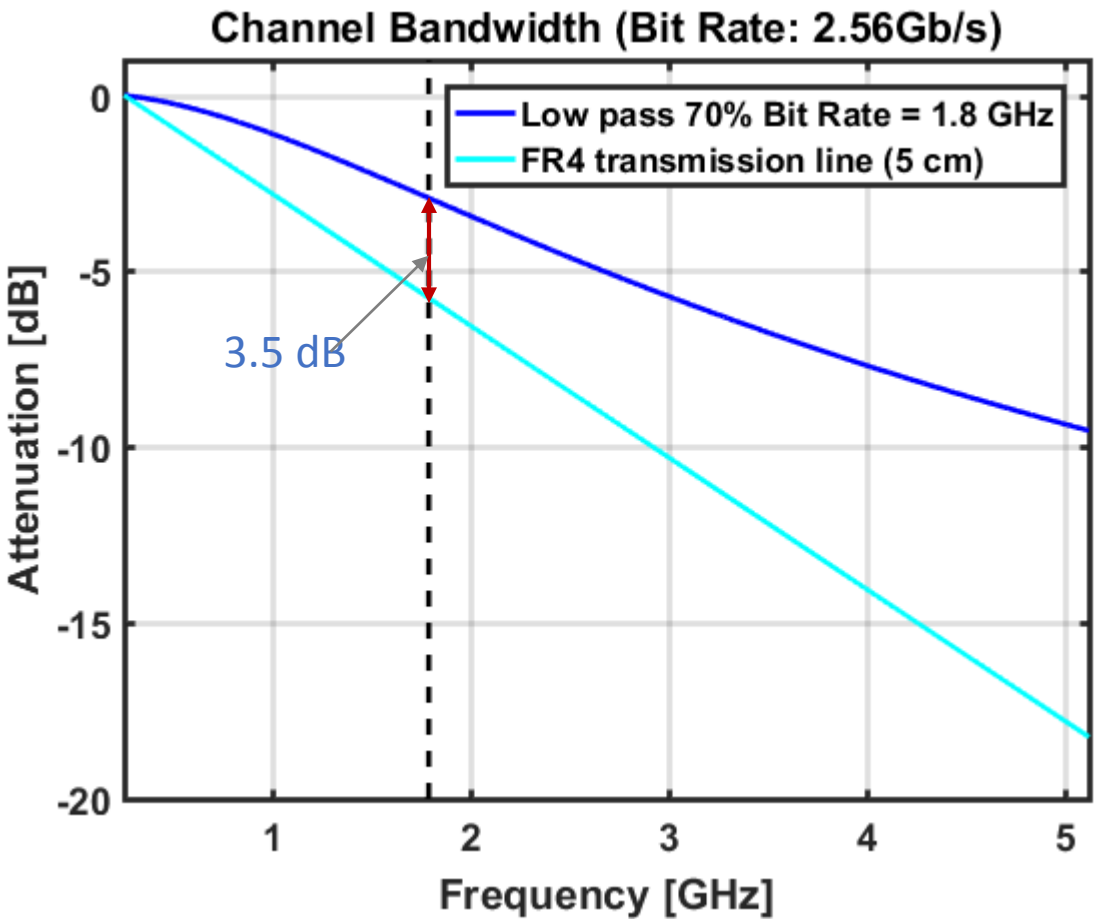

High Pass – Filter Equalization

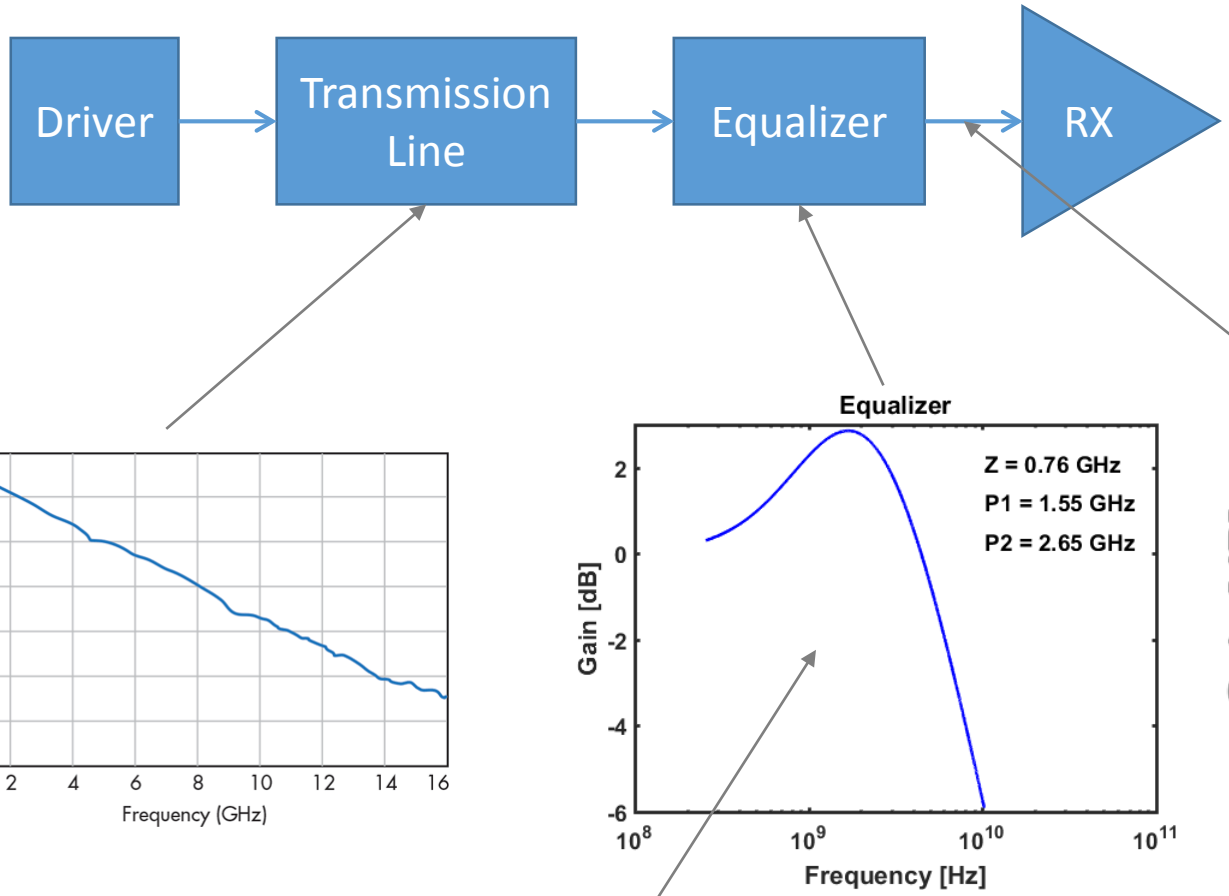

$$H(s) = K\frac{(s - \omega_Z)}{(s - \omega_{P1})(s - \omega_{P2})}$$

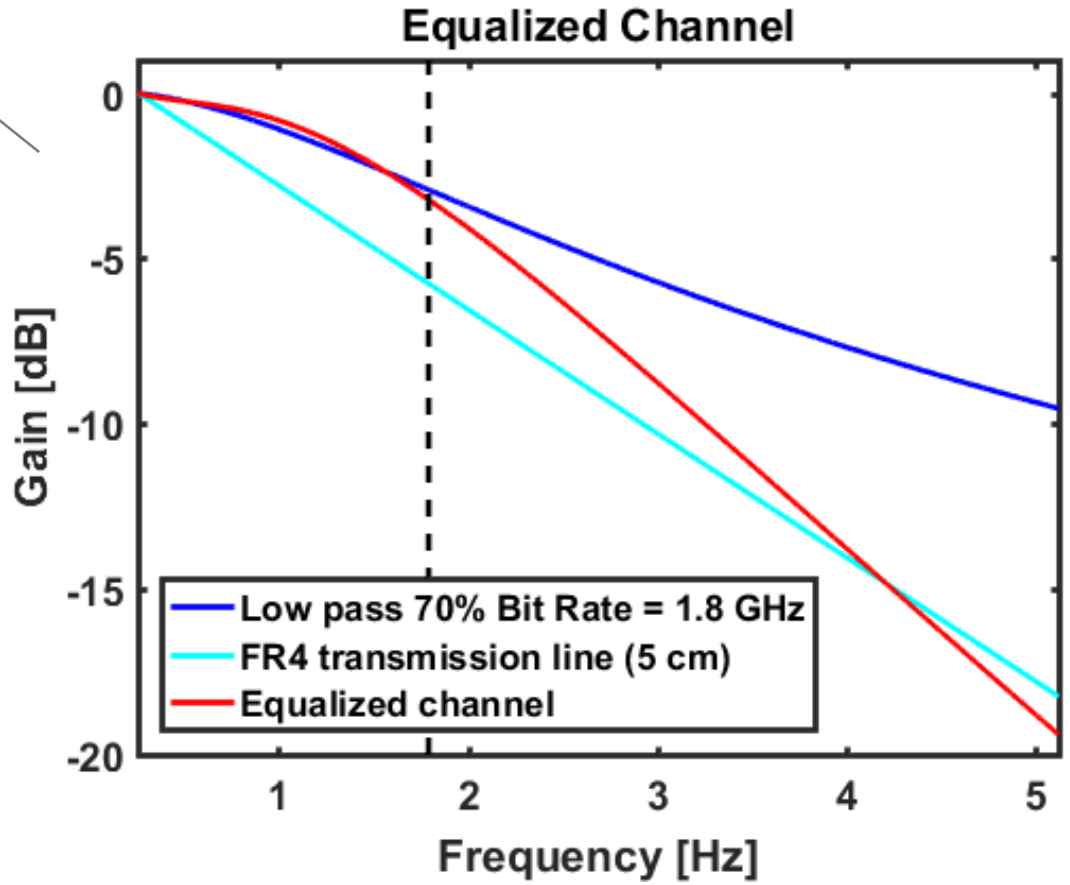

Equalization vs Pre-Emphasis

- Both pre-emphasis and Equalization are attempts to restore the baseband signal spectrum
- Pre-emphasis (done at the transmitter):
  - Tries to generate a wave shape with an "exaggerated" spectral contents at the frequencies that are most attenuated by the channel [typically the high frequencies]
  - No degradation of the SNR
- Equalization (done at the receiver):
  - Amplifies more the high frequency contents of the spectrum than the low frequency. It is an attempt to achieve a combined response of the channel and equalizer that will approach a channel that has no intersymbol interference
  - Degradation of the SNR
- Because of the SNR degradation resulting from equalization, the best approach is to combine pre-emphasis and equalization:
  - Enhance the transmitted high frequency content rather than do all the high frequency peaking at the receiver side

The Equalizer in the LpGBT

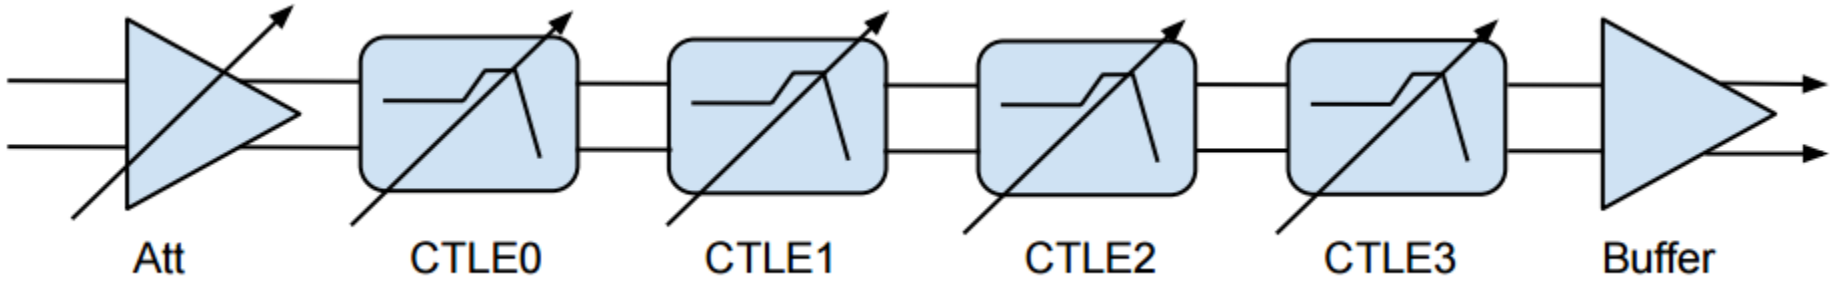

Passive attenuator: allows
to handle signals as high as 1 V
Attenuation: 0, -3.5 and -9.5 dB

Four equalizing stages allow
to flexibly control the shape
of the transfer function.

Restores the signal
to CML levels

Equalizer Stage – RC Degenerated Differential Pair

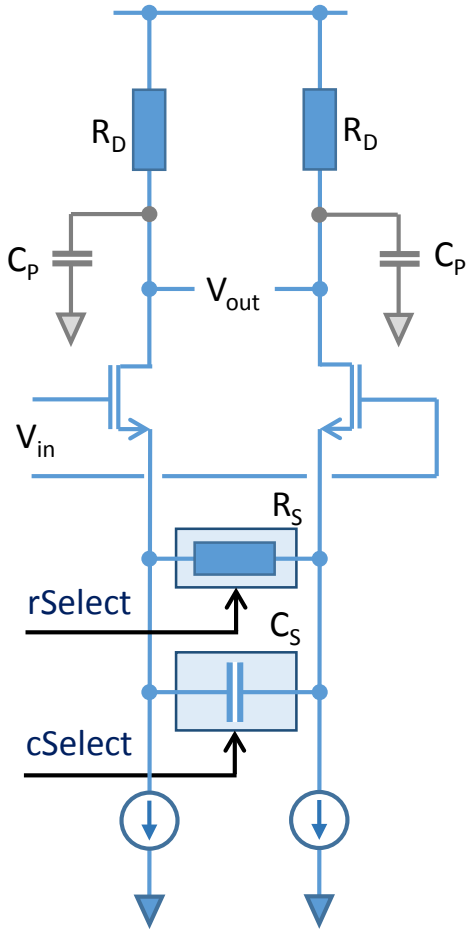

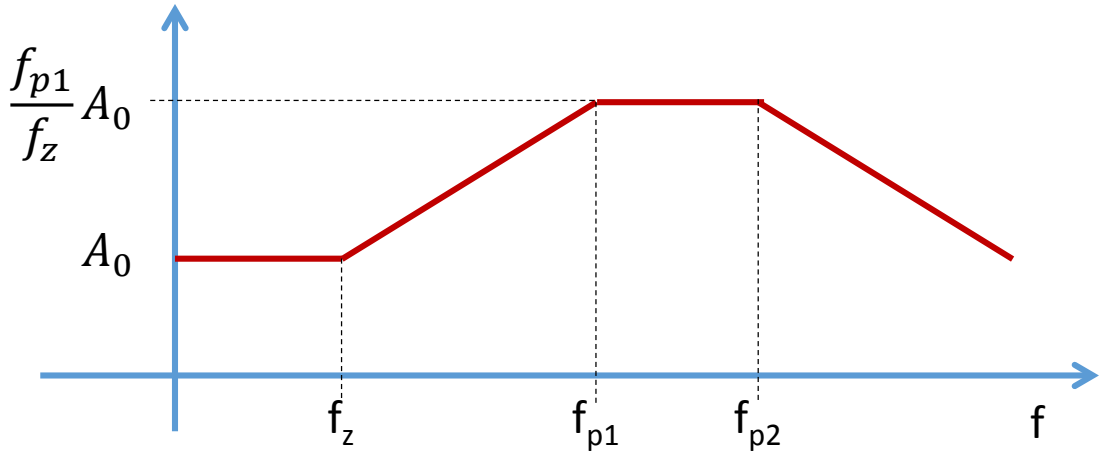

	All Stages	Stage					
Setting		1st and 2nd		3rd		4rd	
	C_s [fF]	R_s [kΩ]	f_z [MHz]	R_s [kΩ]	f_z [MHz]	R_s [kΩ]	f_z [MHz]
1	70	3.0	758	0.6	3789	0.4	5684
2		4.9	464	1.2	1895	1.0	2274
3		7.0	325	2.4	947	1.6	1421
4	140	3.0	379	0.6	1895	0.4	2842
5		4.9	232	1.2	947	1.0	1137
6		7.0	162	2.4	474	1.6	711
7	280	3.0	189	0.6	947	0.4	1421
8		4.9	116	1.2	474	1.0	568
9		7.0	81	2.4	237	1.6	355

$$f_z = \frac{1}{2\pi R_s C_s}$$

$$f_{p1} = \frac{1}{2\pi R_D C_D}$$

$$f_{p2} = \frac{1 + (g_m + g_{mb})\frac{R_s}{2}}{2\pi R_s C_s}$$

$$A_0 = \frac{g_m R_D}{1 + (g_m + g_{mb})\frac{R_s}{2}}$$

First Stage Gain

- Degeneration resistance
  - 3 kΩ, 5 kΩ, 7 kΩ
- Degeneration capacitance
  - 70 fF
- Load cap
  - Cp ≅ 10 fF
    - C_{ds}+ Next stage (C_{gs} + Miller)
- Load resistance
  - 6 kOHm

Rs(kΩ)	fz(GHz)	fp1(GHz)	fp2(GHz)
3	0.76	1.55	2.65
5	0.45	1.25	2.65
7	0.32	1.12	2.65

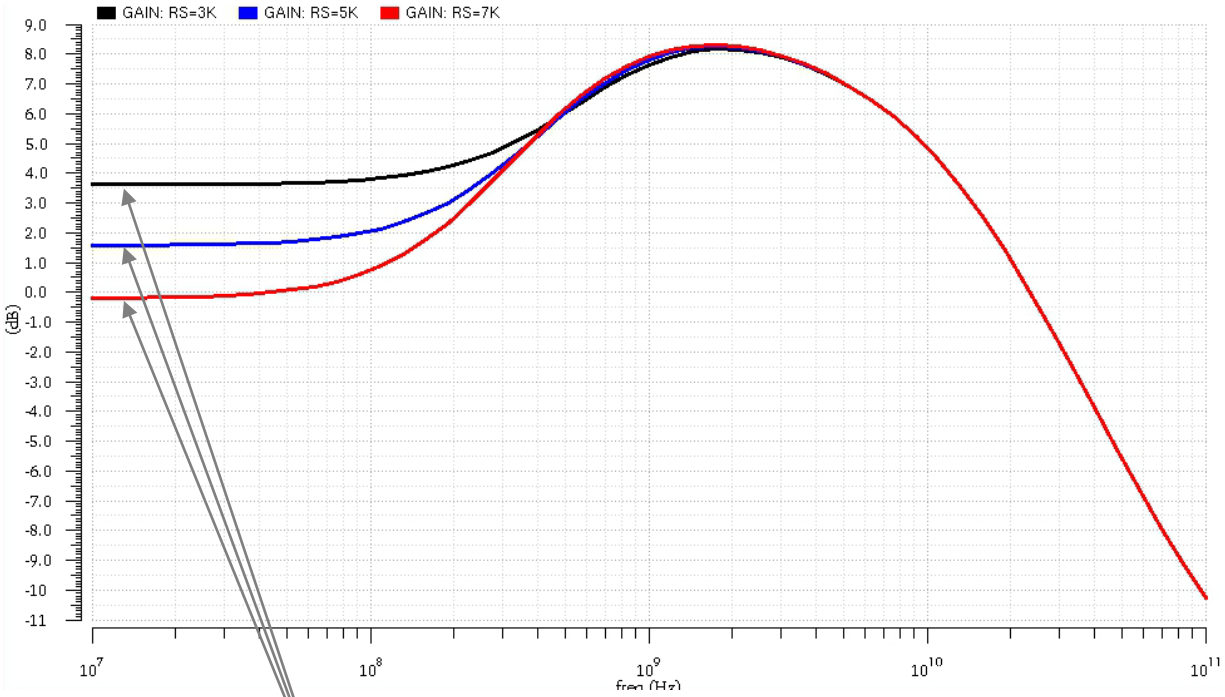

Due to the limited GBW of the technology, peaking is achieved by attenuating the low frequency gain rather than busting the high frequency gain!

75 cm Coaxial Cable

- Simulation conditions:
  - Data rate:
    - 2.56 Gb/s
  - Signal amplitude:
    - 100 mV
  - Process:
    - FF, V_{DD} = 1.08 V, T = 100 °C
- Results (worst case):
  - Power dissipation:
    - 2.34 mW (max)
  - PP jitter:
    - 70 ps (max)
  - SNR:
    - 20 (26 dB) (min)

Worst case Eye diagram (SNR20)

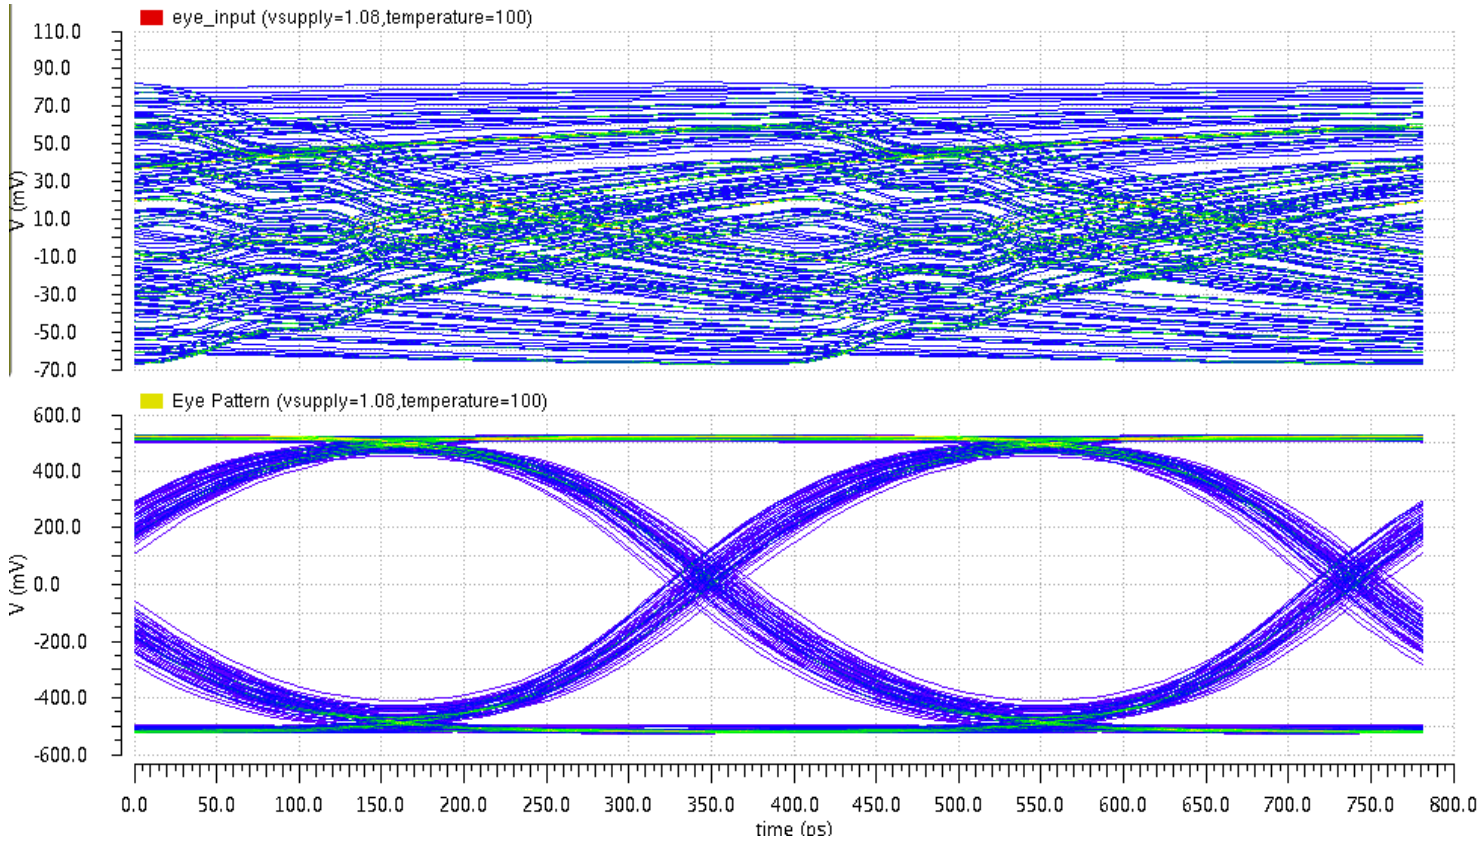

Eye Opening Monitor

- Goal:
  - Monitor the opening of the received data eye diagram and the equalizer's performance
- What is the problem:
  - BER is limited by intersymbol interference (ISI), cross-talk with neighbour channels, reflections due to impedance mismatch
- Solutions:
  - Employ differential signalling, Pre-Emphasis at the transmitter, Continuous Time Linear Equalizer (CTLE) and Decision-Feedback Equalization (DFE) at the receiver
  - The Eye-Opening Monitor (EOM) in the LpGBT will be used to monitor the performance of the down-link and to optimize the equalizer performance (CTLE).

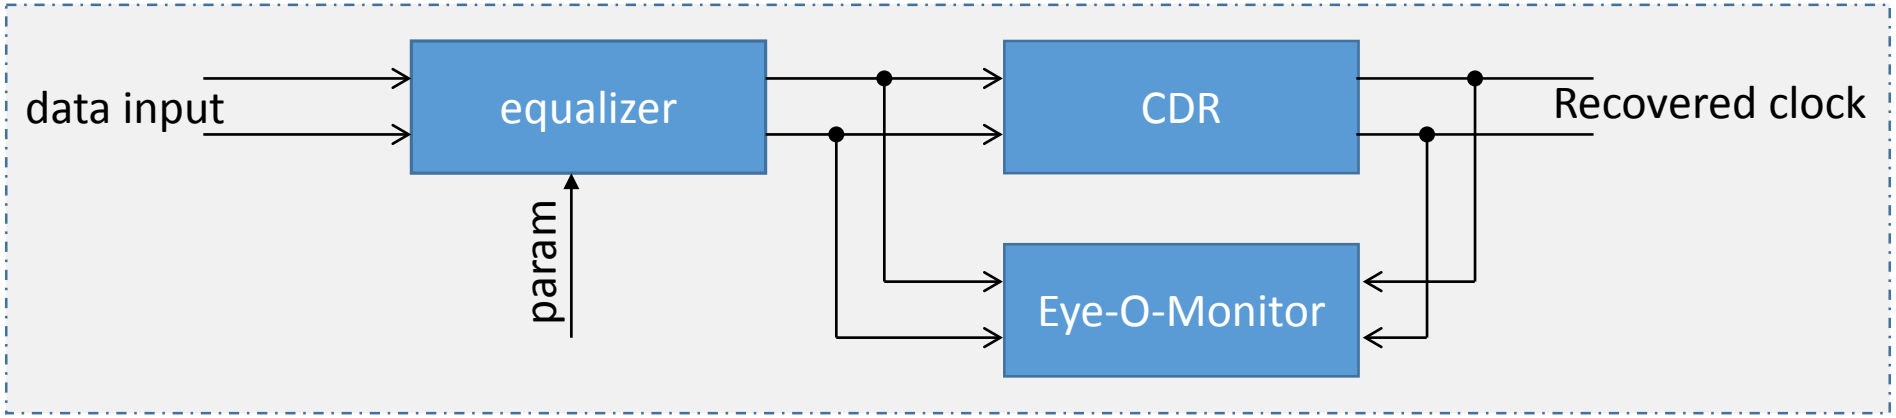

Eye Opening Monitor

- Provides an “eye diagram picture” by using a “signal-scan” approach
- The scan is performed across the time (x-axis) and across the amplitude (y-axis), yielding a “signal density” per point
  - The input signal [data] is compared with a reference voltage “V_{of}”
  - The comparator’s result is sampled by the rising edge of a clock synchronous to the incoming data
  - The sampled result drives a ripple counter to accumulate statistics
  - The counter is enabled for a well defined period.

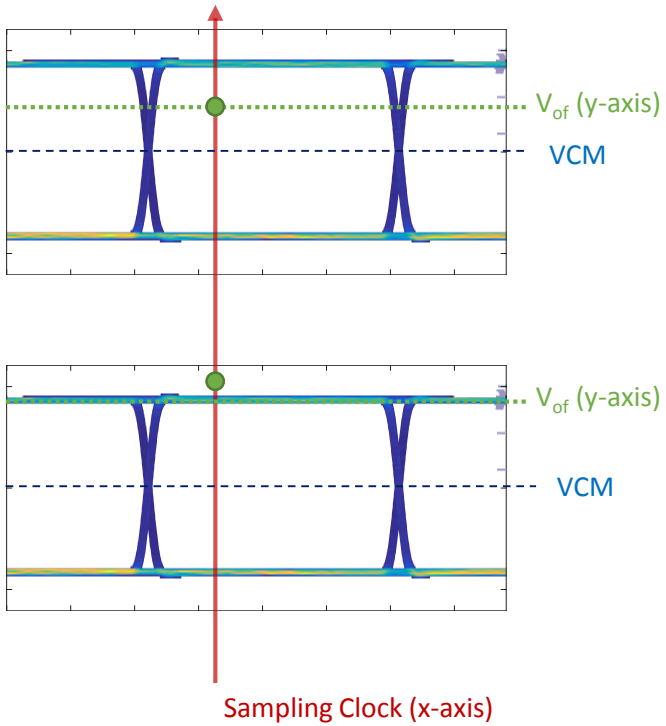

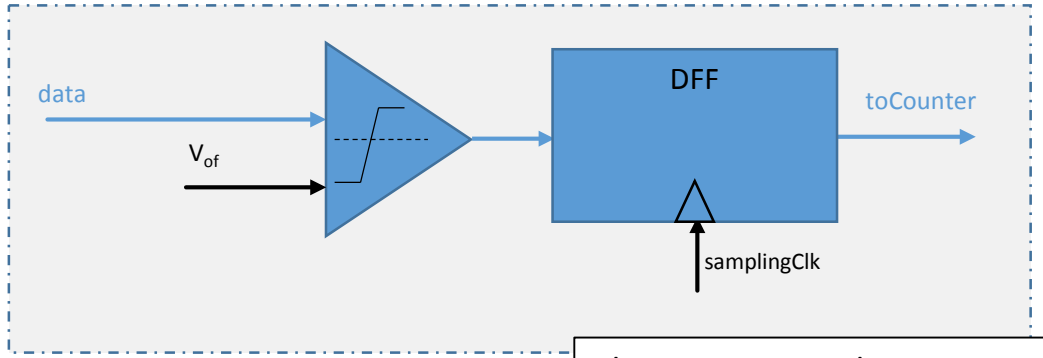

The counter is only incremented when V_{of} is between the *data* amplitude values

Eye Opening Monitor

If V_{of} is within the data amplitude levels the comparator toggles

Signal amplitude ≅ V_{of}
Comparator toggles "statistically"

Uncertainty zone due to noise and ISI

If V_{of} is outside the data amplitude levels the comparator never toggles

V_{of}

VCM

samplingClock

comp
data - vof

1	0	1	0	0/1	0	0/1	0	0	0	0	0
				1	0/1	1	0/1	1	1	1	1

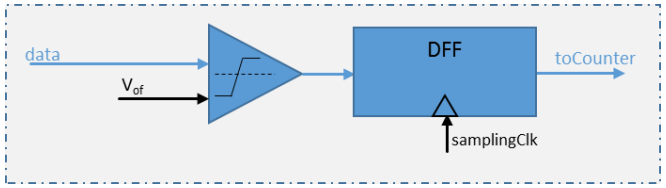

"**toCounter**" toggles when V_{of} is within the input signal amplitude's maximum and minimum at the rising edge of the sampling clock!

EOM in the lpGBT

- Y-axis:
  - 31 points, step = ~20 mV (covers from $V_{DD}/2$ up to V_{DD})
- X-axis:
  - 64 points, step = ~6.1 ps in typical

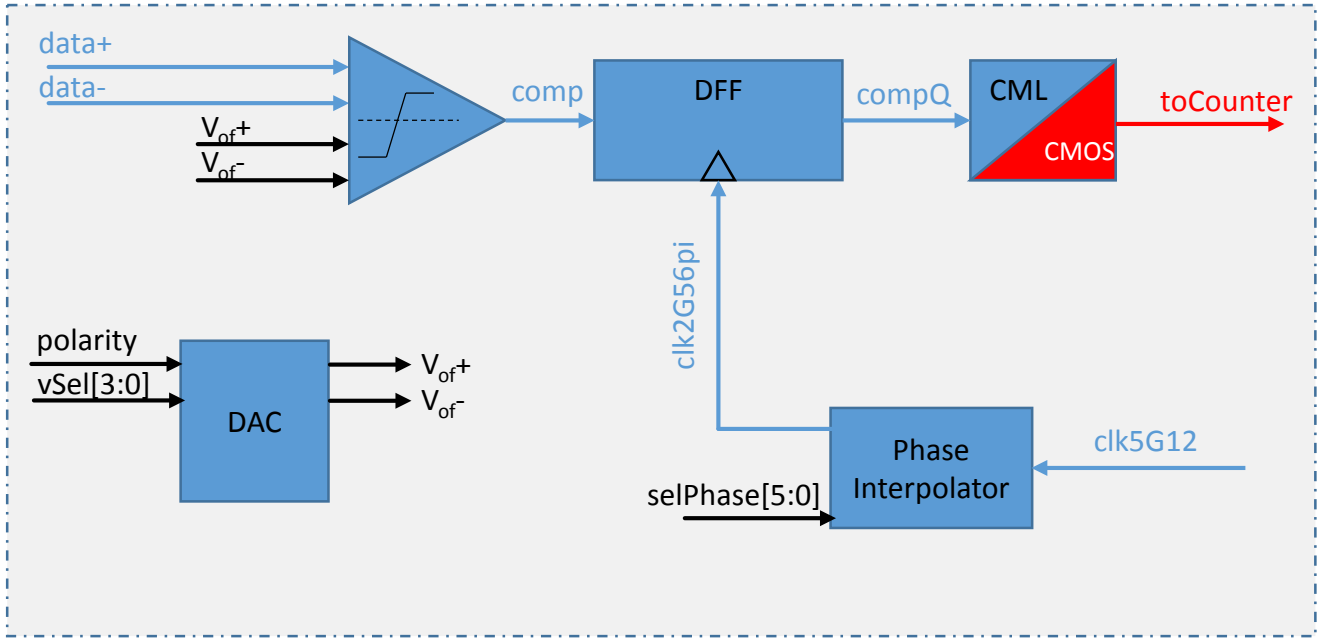

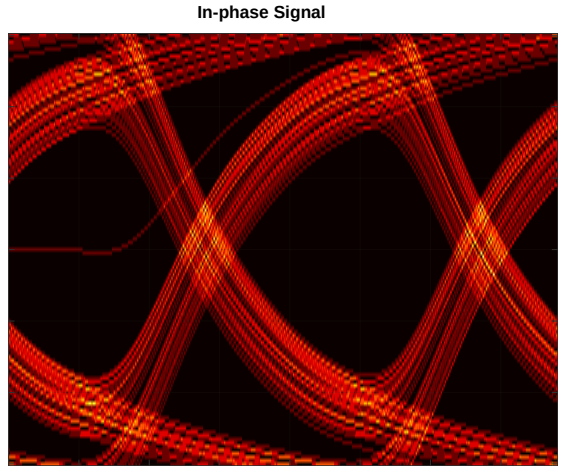

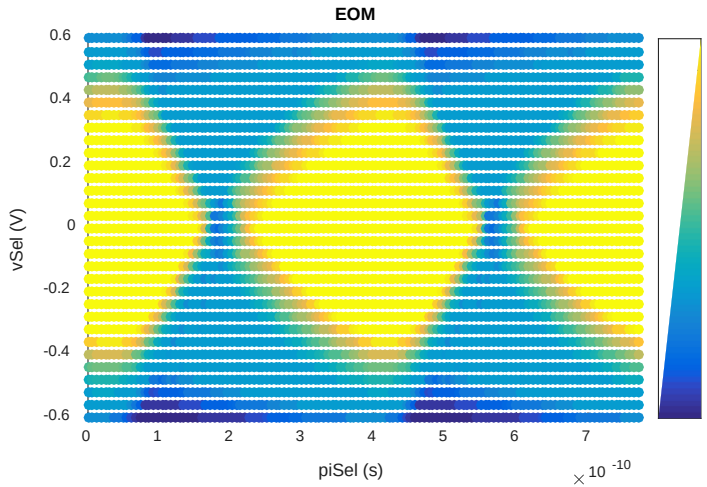

EOM in the lpGBT (y-axis)

- The comparator uses a differential difference amplifier
- This eliminates the dependence on the common-mode of the input signals

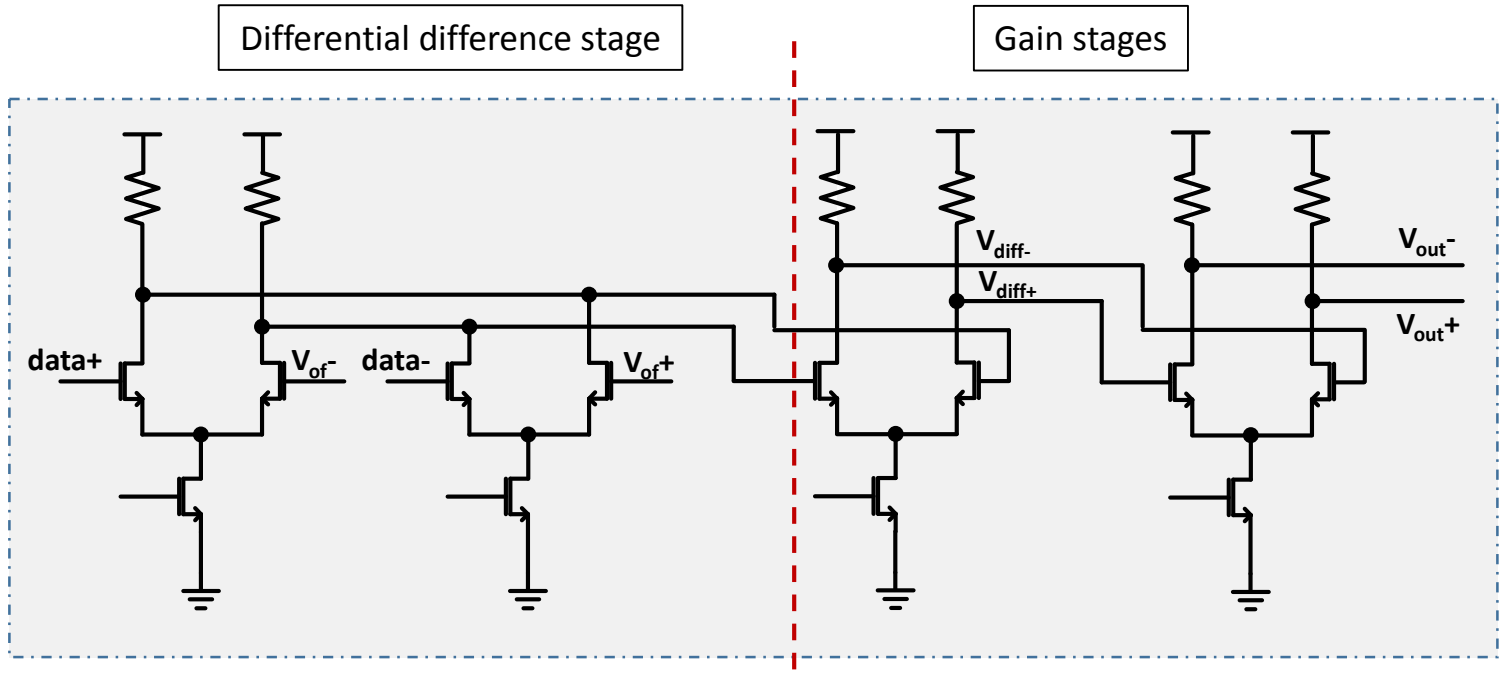

$$V_{diff} = Av[(data+ - V_{of}-) - (data- - V_{of}+)]$$

meaning

$$V_{diff} = Av[(data+ - data-) - (V_{of}+ - V_{of}-)]$$

EOM in the lpGBT (x-axis)

- Phase interpolator (1/2)
  - Receives the VCO clock (5.12 GHz)
  - Generates in-phase "I" and quadrature "Q" signals
  - Uses those for phase interpolation:
    - Full phase rotation with ≅ 6.1 ps resolution (64 possible phases)

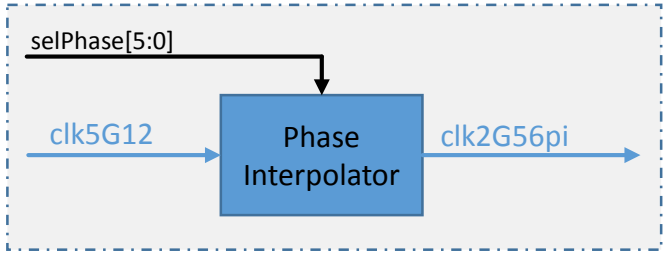

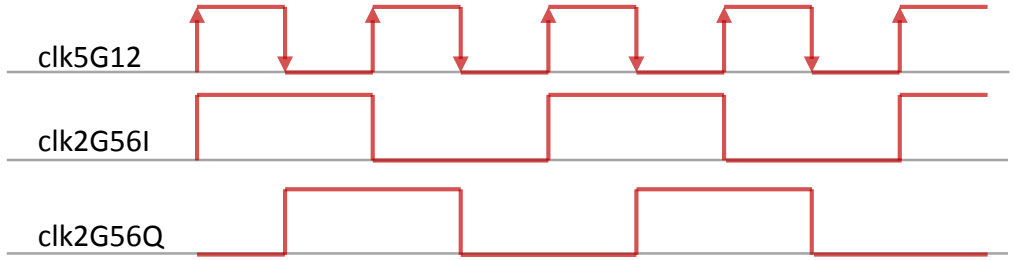

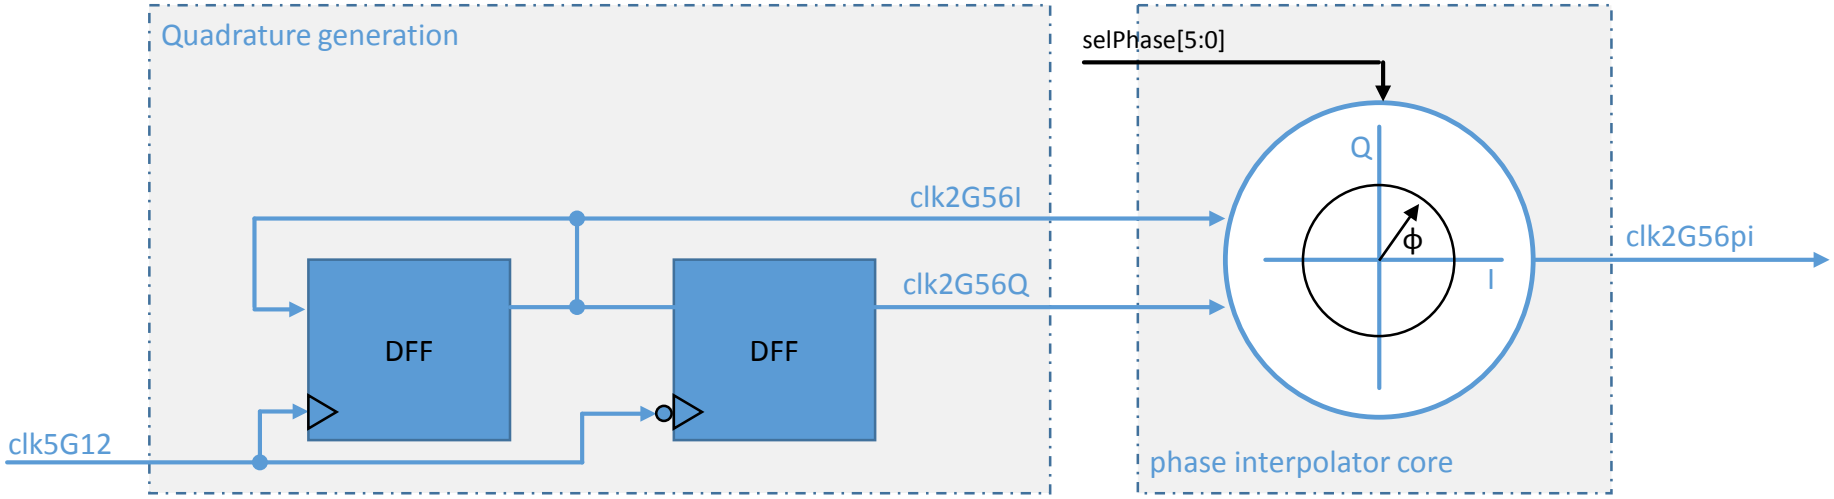

EOM in the lpGBT (x-axis)

- Phase interpolator (2/2)

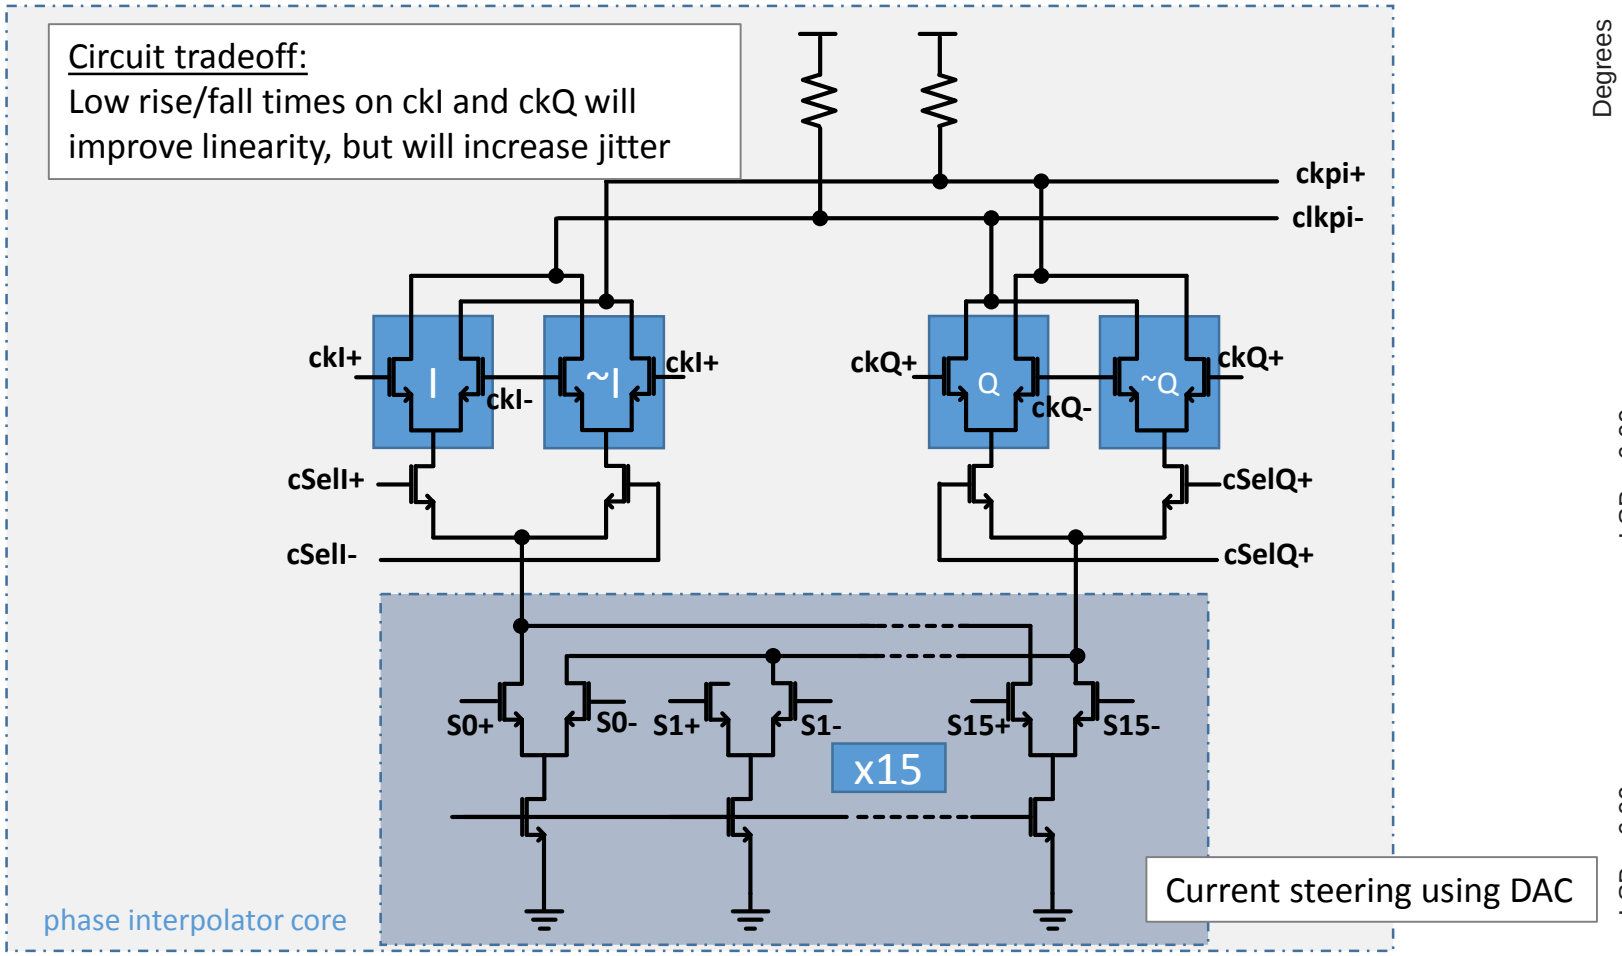

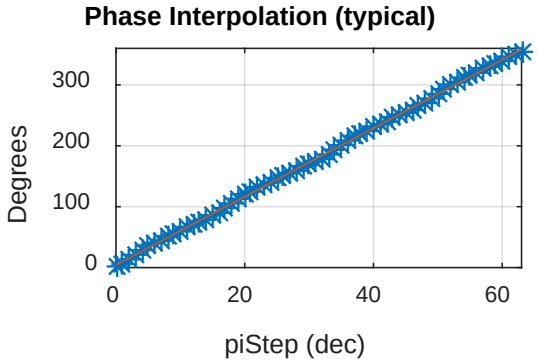

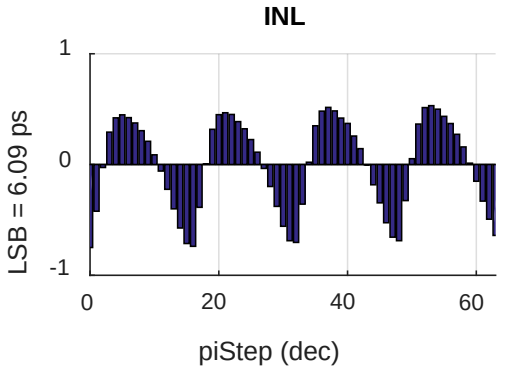

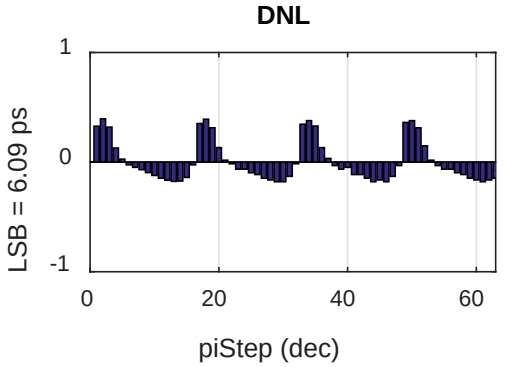

Reference: A 40 Gb/s CDR with Adaptive Decision-Point Control Using Eye-Opening-Monitor Feedback, H. Noguchi et al, ISSCC. 2008

Clock and Data Recovery

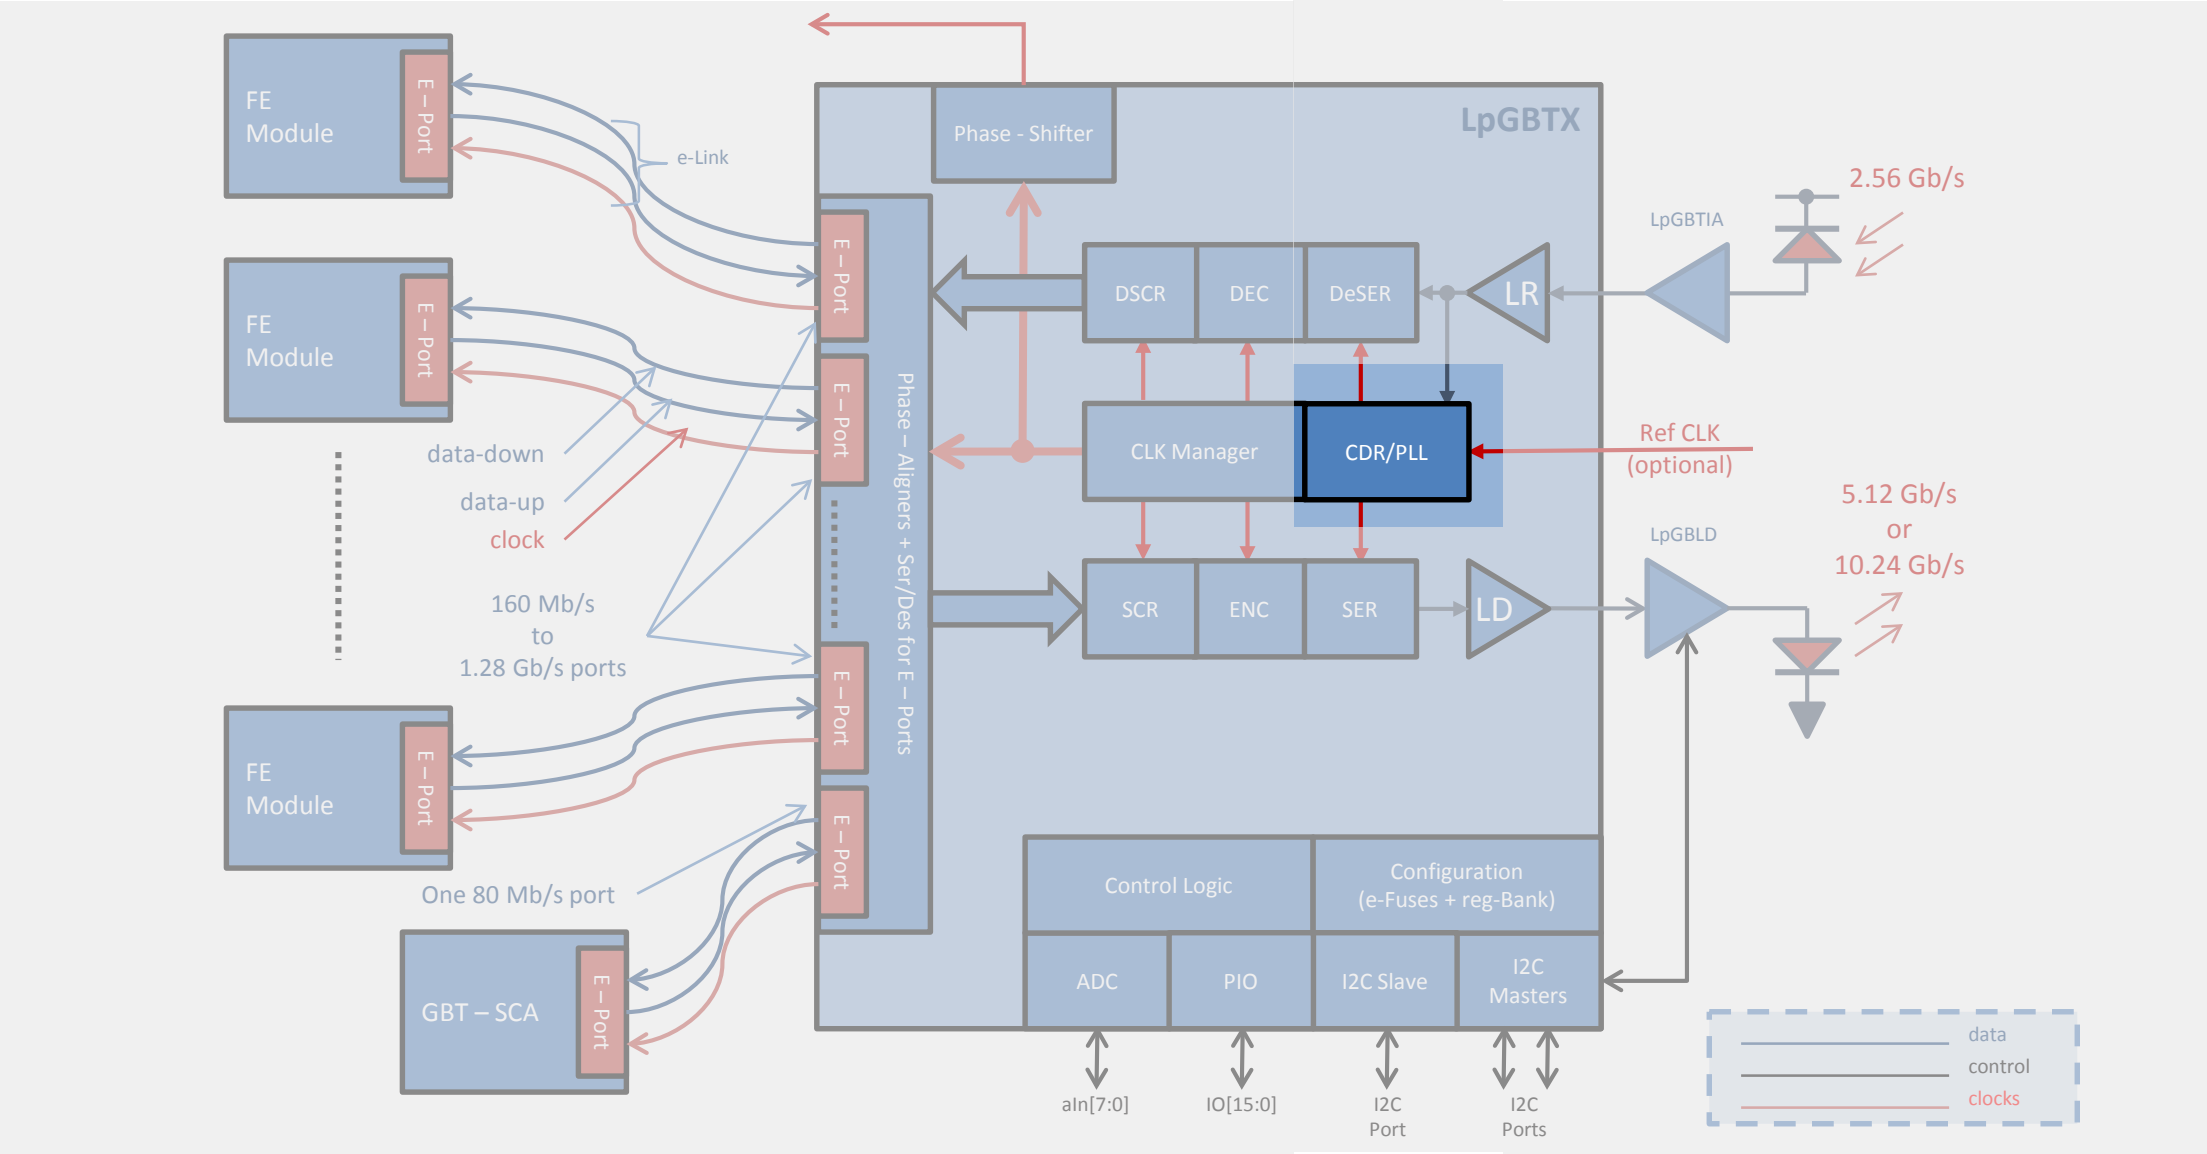

PLL/CDR

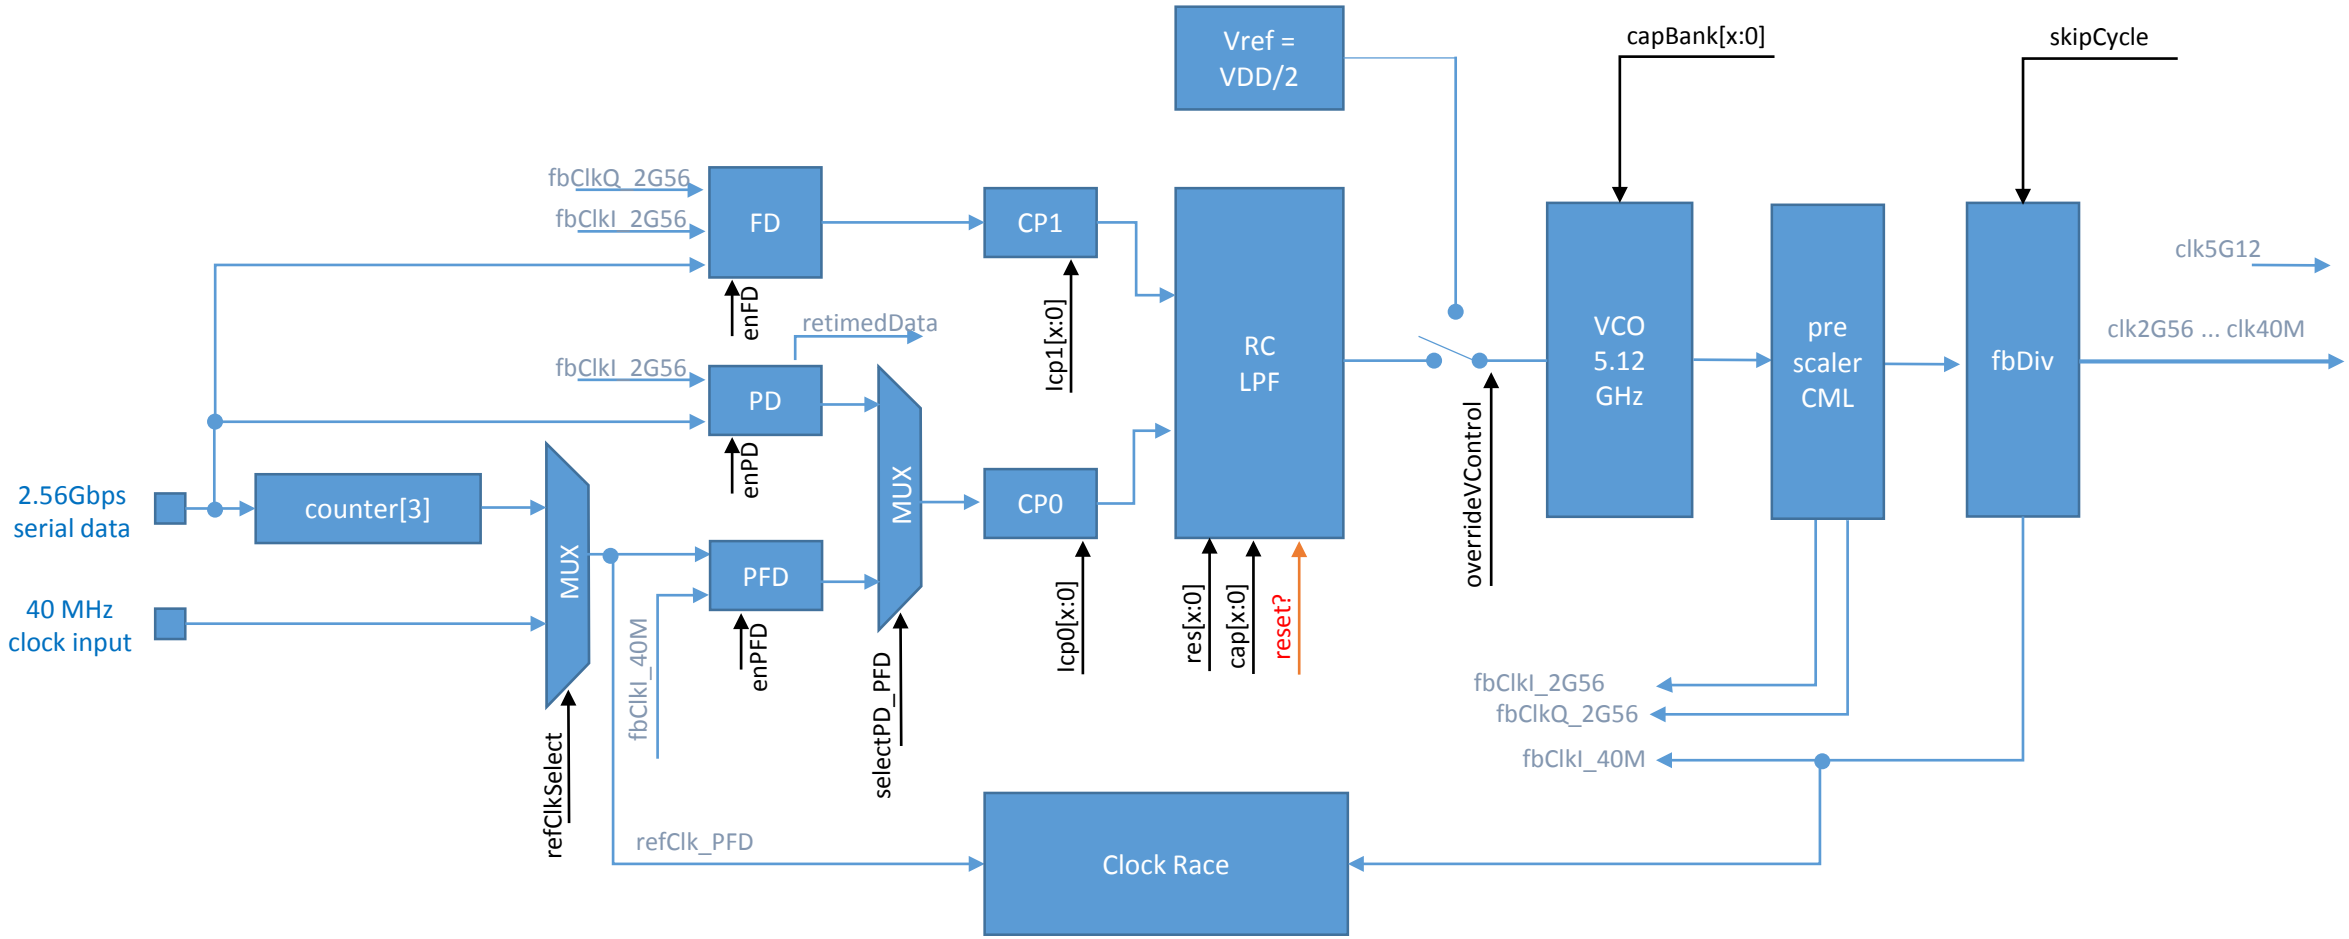
Vref =
VDD/2
capBank[x:0]
skipCycle
fbClkQ_2G56
fbClkI_2G56
FD
enFD
CP1
Icp1[x:0]
RC
LPF
VCO
5.12
GHz
pre
scaler
CML
fbDiv
clk5G12
clk2G56 ... clk40M
retimedData
fbClkI_2G56
PD
enPD
MUX
CP0
Icp0[x:0]
res[x:0]
cap[x:0]
reset?
overrideVControl
2.56Gbps
serial data
counter[3]
MUX
PFD
enPFD
fbClkI_40M
selectPD_PFD
40 MHz
clock input
refClkSelect
refClk_PFD
Clock Race
fbClkI_2G56
fbClkQ_2G56
fbClkI_40M

The LpGBT

- For low jitter the PLL uses an LC oscillator (VCO)
- Advantage:
  - Low phase noise (jitter)
- Disadvantage:
  - Limited tuning range
- Due to PVT it is unlikely the VCO center frequency will be 5.12 GHz (f_{LHC} x 128)
- A calibration of the VCO is required at beginning of operation!
- The process is easy in the PLL mode since a reference clock (f_{LHC}) is available:
  - The VCO control voltage is fixed ($V_{DD}/4$)
  - The VCO clock frequency is compared with the reference.
  - Switched capacitors are switched in or out.
  - Based on the frequency measurement, the best capacitor setting is selected
- But what if the ASIC is used as a receiver and no reference clock is available?

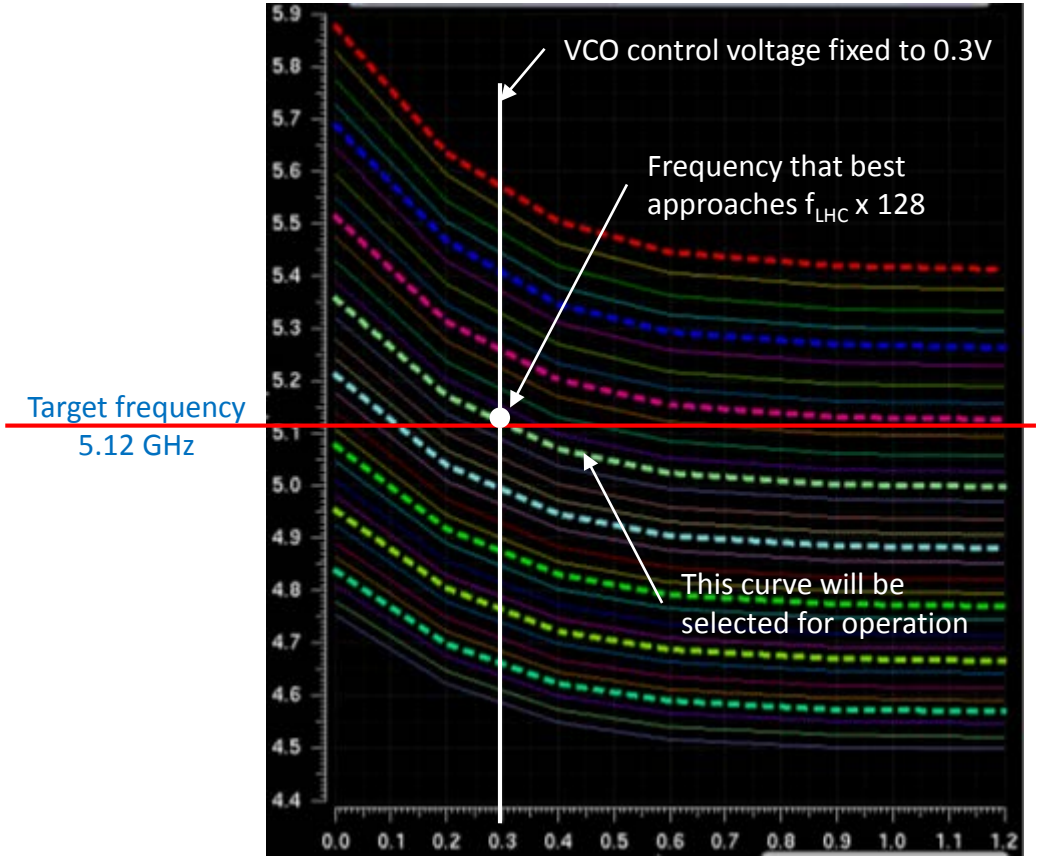

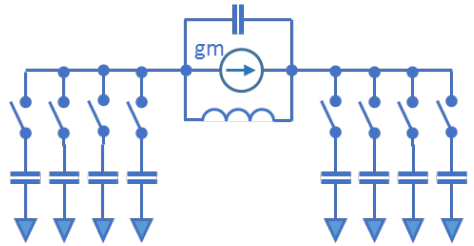

Reference-Less Locking

- When working as a CDR, although a clock reference is not present, the serial data itself can be used to calibrate the VCO centre frequency!

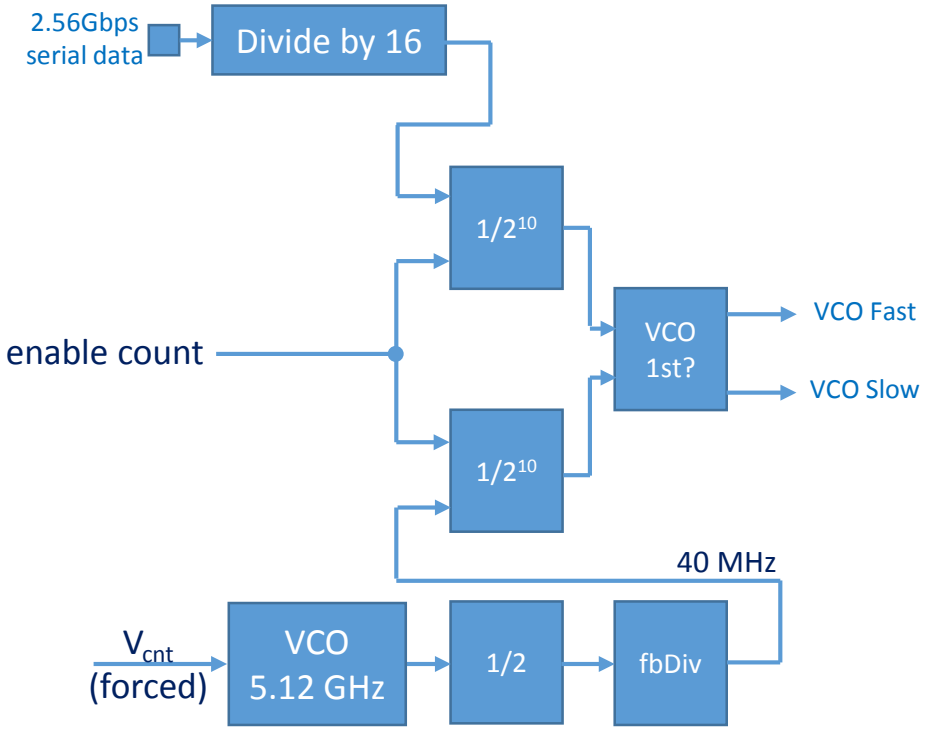

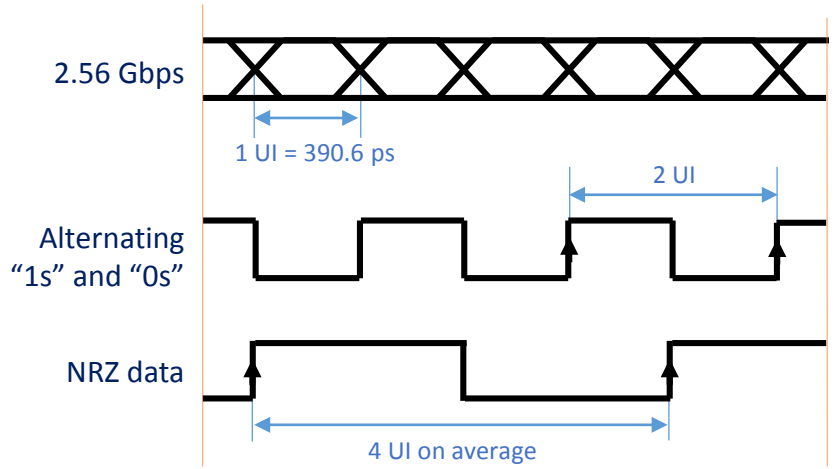

On average, due to NRZ data probability, the distance between two consecutive rising edges is 4 UI

On average, the signal looks like a 2.56GHz/4 = 640 MHz clock!

If this signal drives a "divide by 16" counter, a 40 MHz clock reference is obtained!

This is a rather imprecise frequency but if further divided the accuracy increases

TID Radiation: Topology Matters (1/2)

- Two PLLs working at 2.5 GHz
- Same circuits except the VCO:
  - LC VCO
  - Ring oscillator VCO
- Same power dissipation
- Same loop dynamics

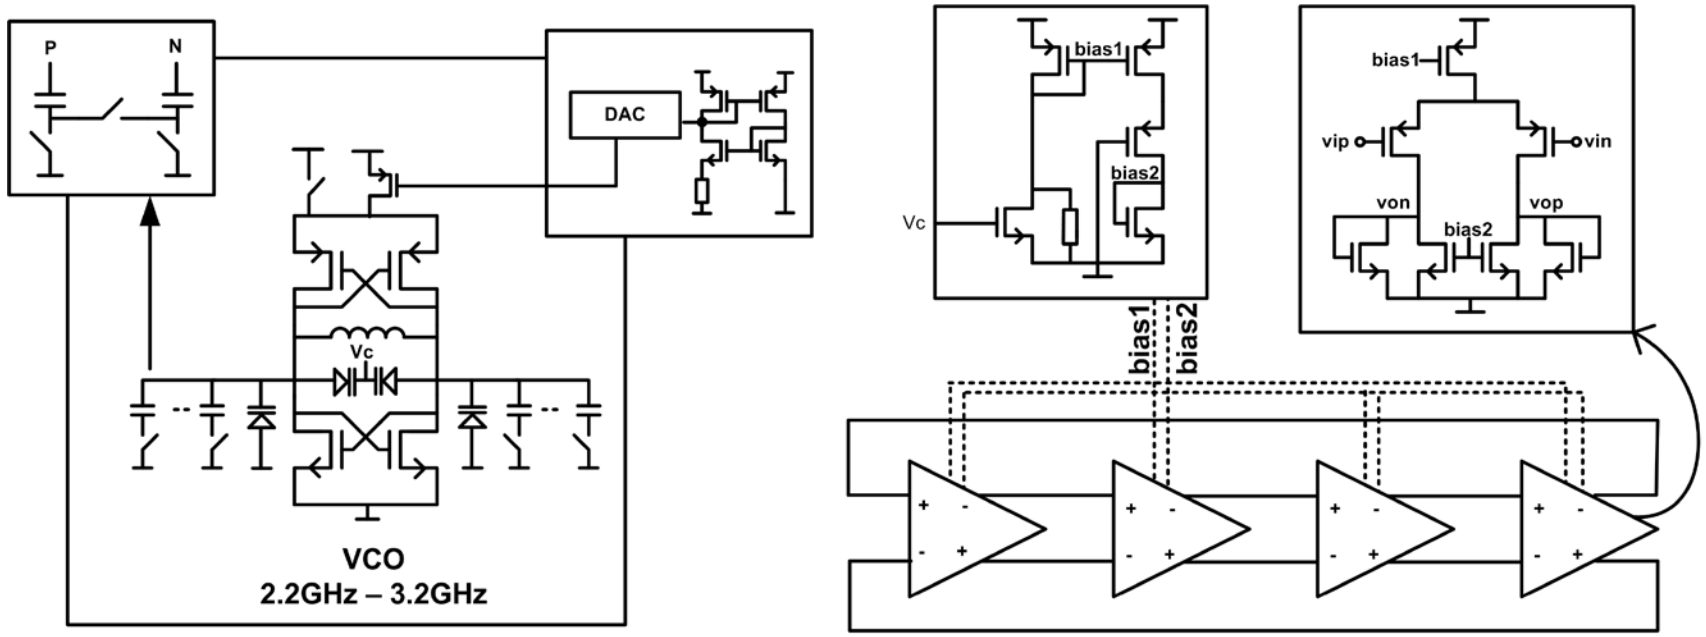

TID Radiation: Topology Matters (2/2)

- Pre/post-rad jitter (rms):
  - LC: 0.3 / 1.0 ps
  - RO: 5.6 / 22 ps
- 600 Mrad + Annealing:
  - LC: $\Delta f < 5\%$ (LC: Passives set the frequency)
  - RO: $\Delta f < 40\%$ (Ring: Actives set the frequency)

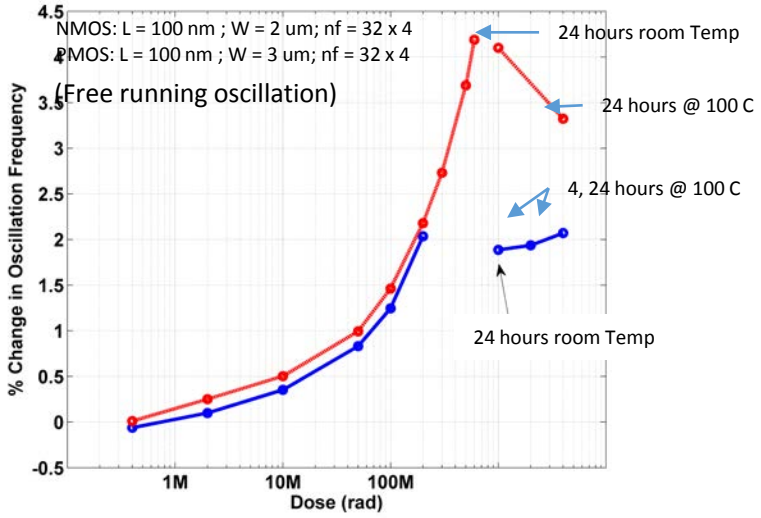

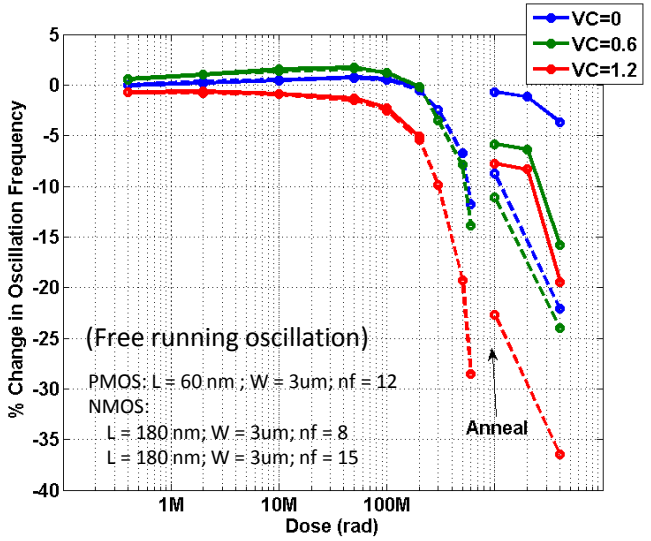

SEU Radiation: Topology Matters – Round 1

- Heavy ion testing:
  - LET: 3.2 to 69.2 MeV.cm^2/mg
- LC oscillator displays a significantly higher sensitivity than the ring oscillator!
  - Contrary to expectations!
- SEU Phase jumps:
  - Type: phase unlock
  - Ring oscillator: both polarities
  - LC: Mainly positive
- Two-Photon Absorption (TPA) laser tests point to the VARACTOR as the main culprit!
  - Total cross section of the LC-oscillator is 4 10^{-5}cm^2 from which 70% is contributed by the varactor area!

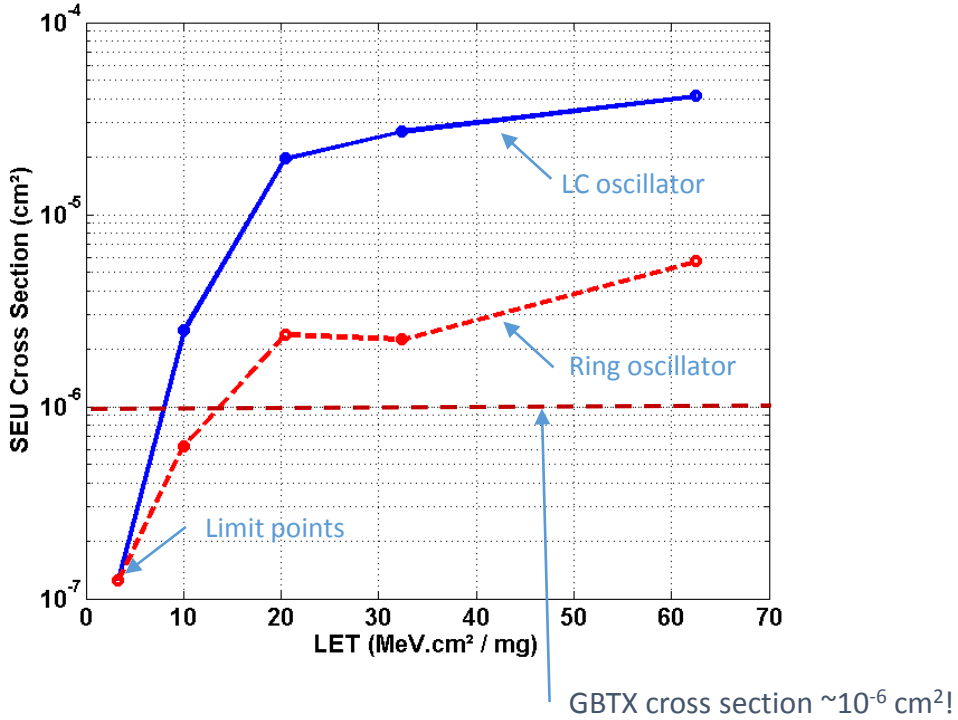

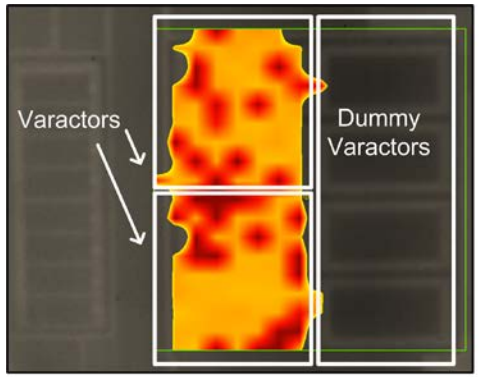

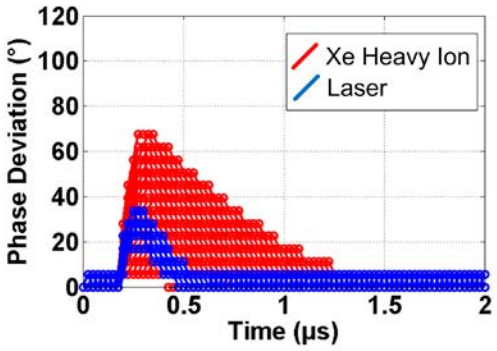

TID Radiation: Topology Matters – Round 2

- A new design prototyped to test the hypothesis:
  - Smaller varactor area
  - Different frequency tuning topology:
    - Grounded vs floating well!

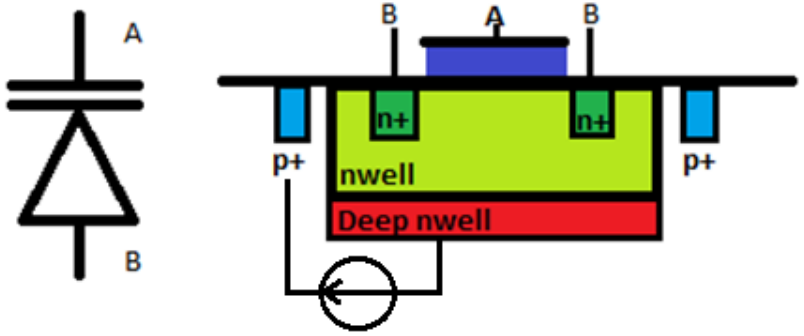

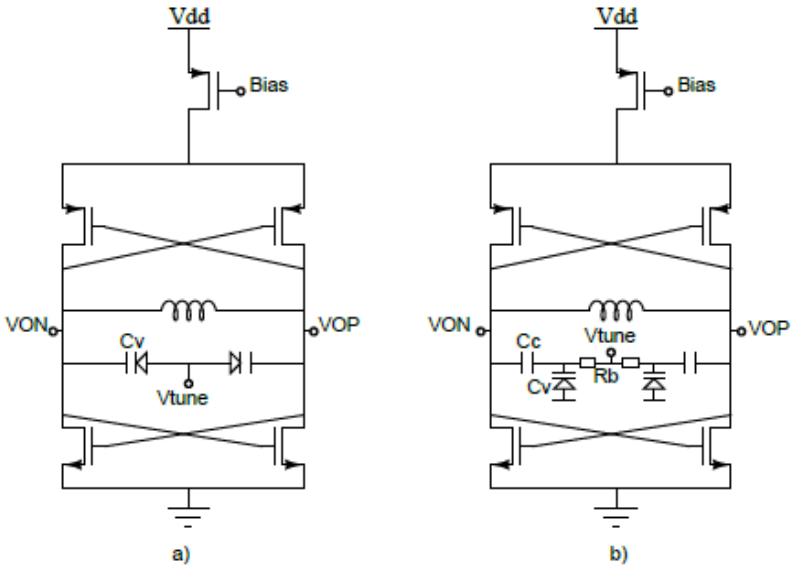

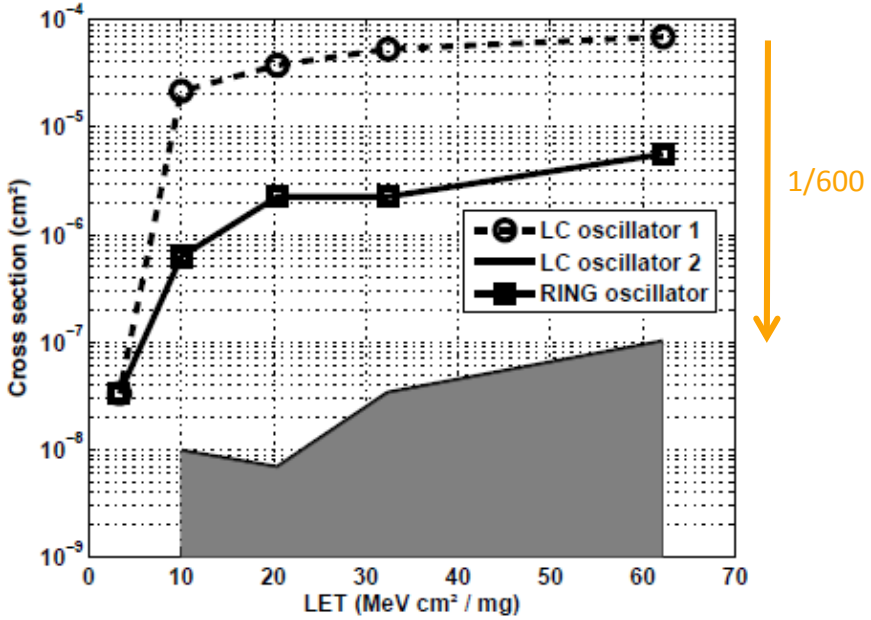

Clock Recovery

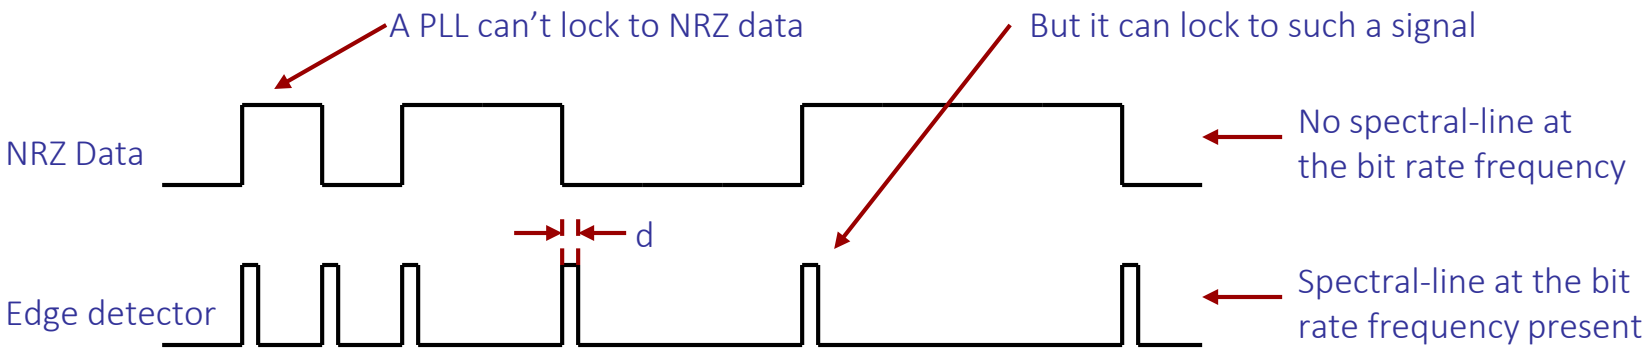

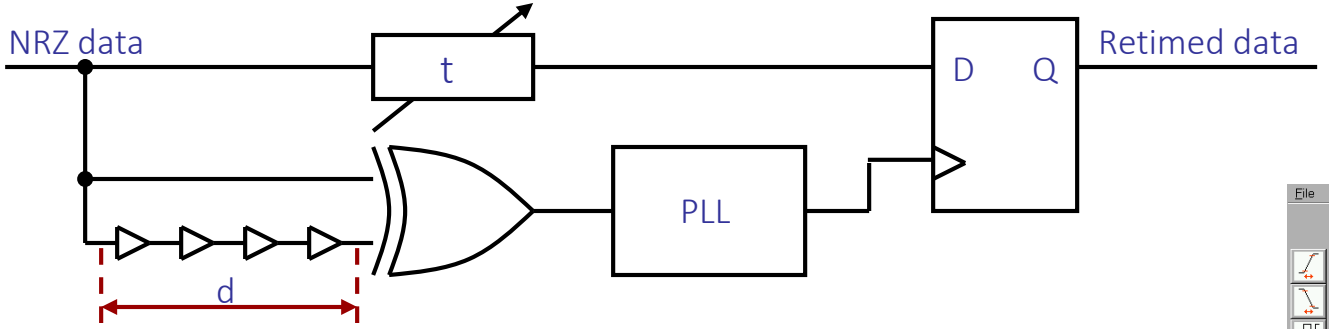

Ideal sampling instant

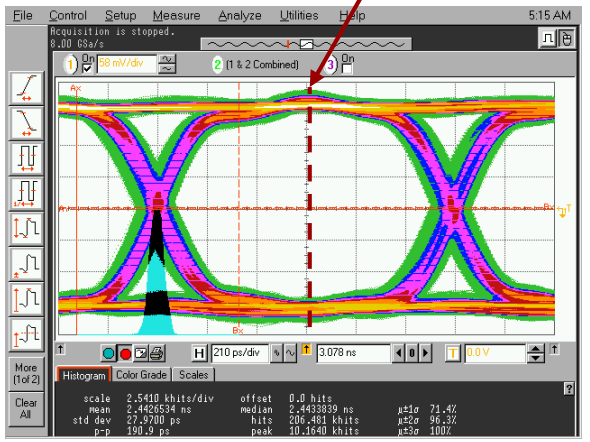

Problem:

Difficult to match the delay of the two paths so that the recovered clock samples the data at the middle of the data eye over all process, temperature and supply variations.

Alexander Phase Detector - Principle

- It is a bang-bang detector:
  - Only early/late information
- To take the Early/Late decision:
  - 1st Look for transitions
  - 2nd Take an Early/Late decision at every transition
- Three samples of the serial data are necessary to find the transition and resolve the phase relationship.
- No data transition present:
  - S1 = S2 = S3
  - S1 ⊕ S3 = 0 and S1 ⊕ S2 = 0
  - Charge-pump: Hold
- Data transition + Early clock:
  - Sample S1 ≠ S3 and S1 = S2
  - S1 ⊕ S3 = 1 and S1 ⊕ S2 = 0
  - Charge-pump: Down = (S1 ⊕ S3) & ~(S1 ⊕ S2)
- Data transition + Late clock:
  - Sample S1 ≠ S3 and S1 ≠ S2
  - S1 ⊕ S3 = 1 and S1 ⊕ S2 = 1
  - Charge-pump: Up= (S1 ⊕ S3) & (S1 ⊕ S2)
- The falling edge of the clock aligns with the data transition instants:
  - In lock, S1 or S3 are thus at the optimum sampling instants

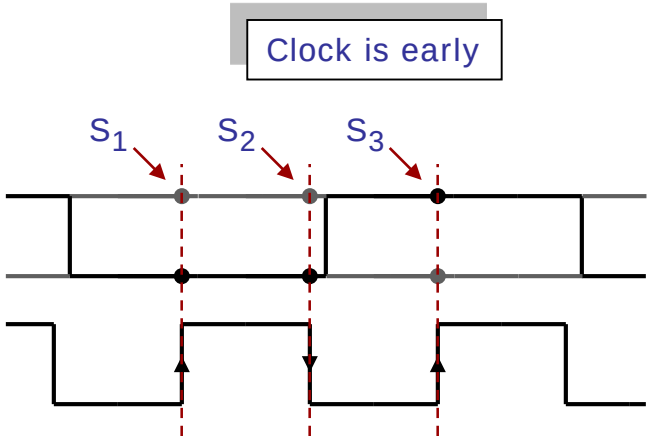

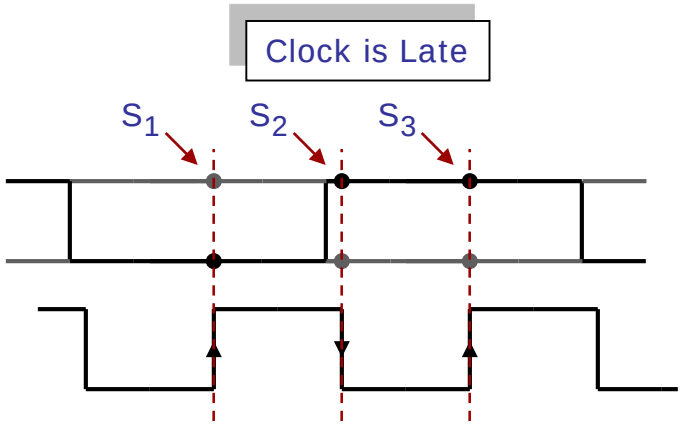

Alexander Phase Detector

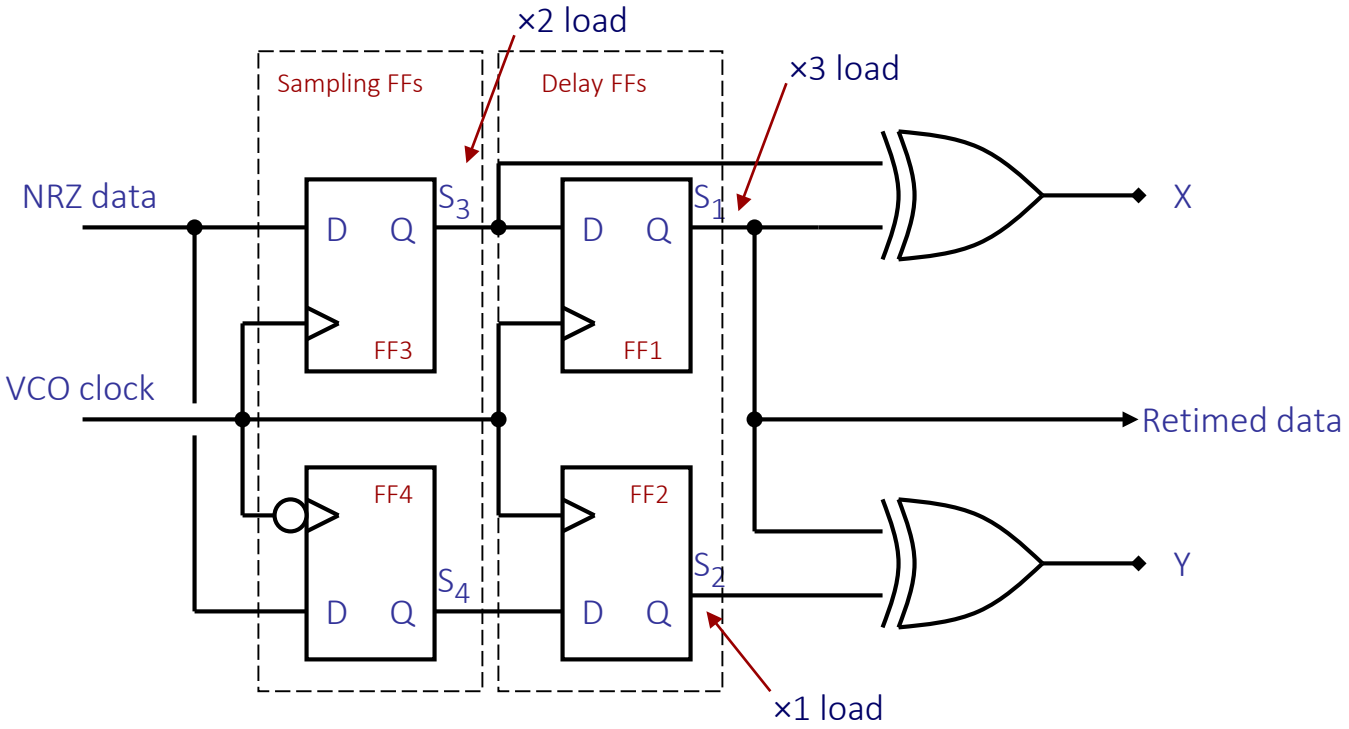

Down = X & ~Y

Up = X & Y

Connectors, PCB & PACKAGE

Reaching the ASIC

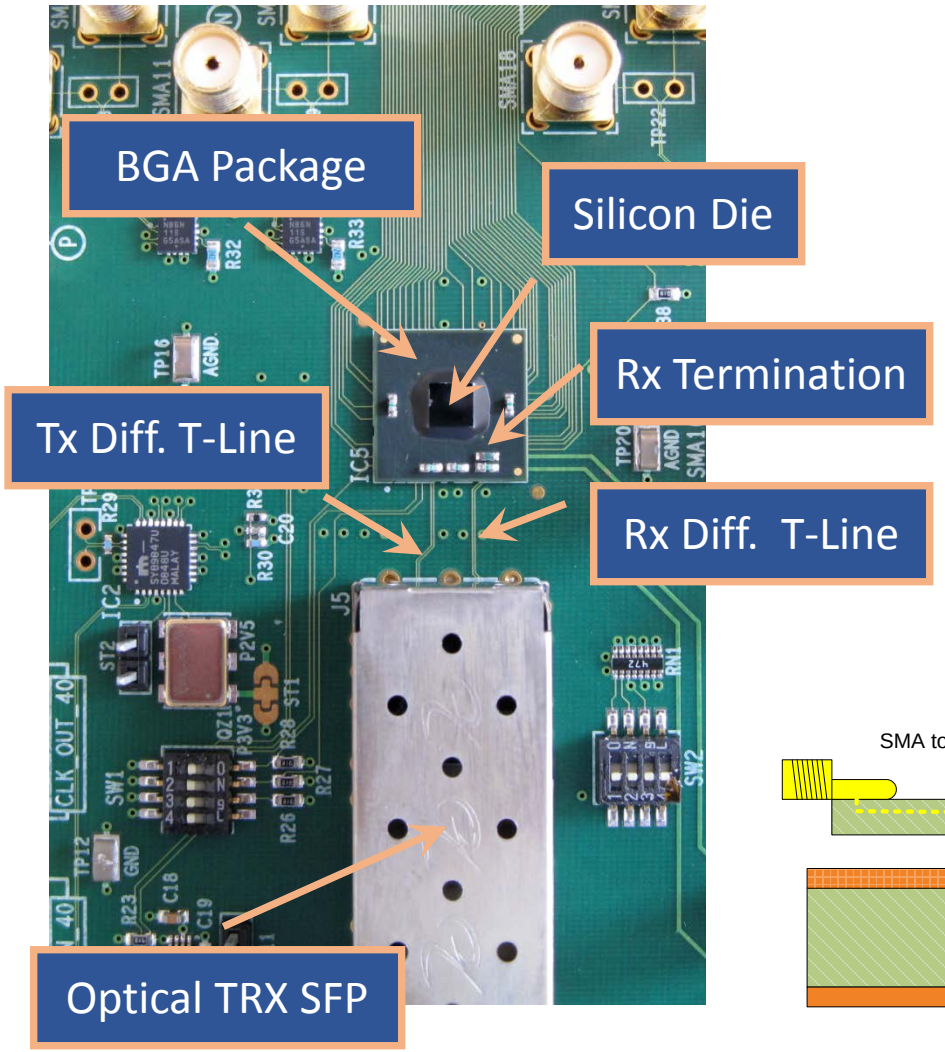

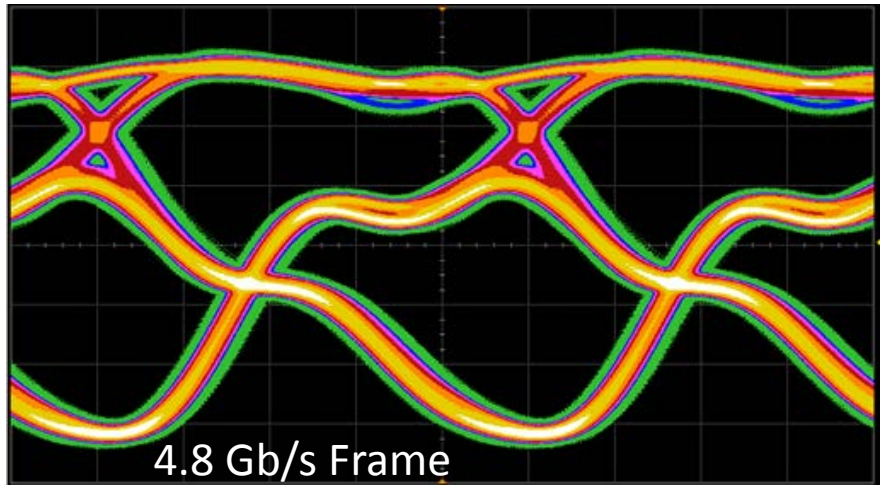

Probe Effect ?
Test board ?
Package ?
Chip Design ?

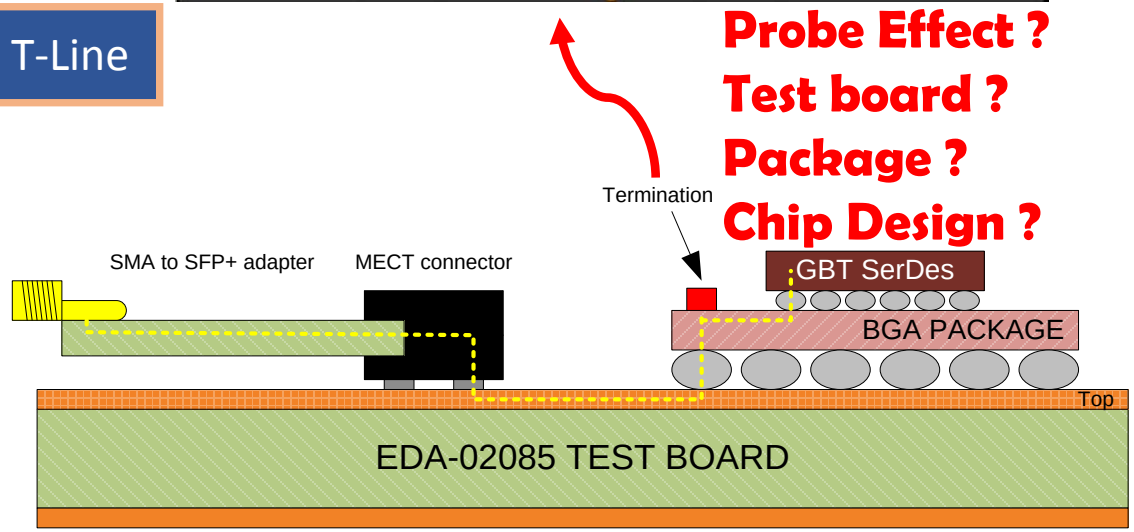

“Build” Models

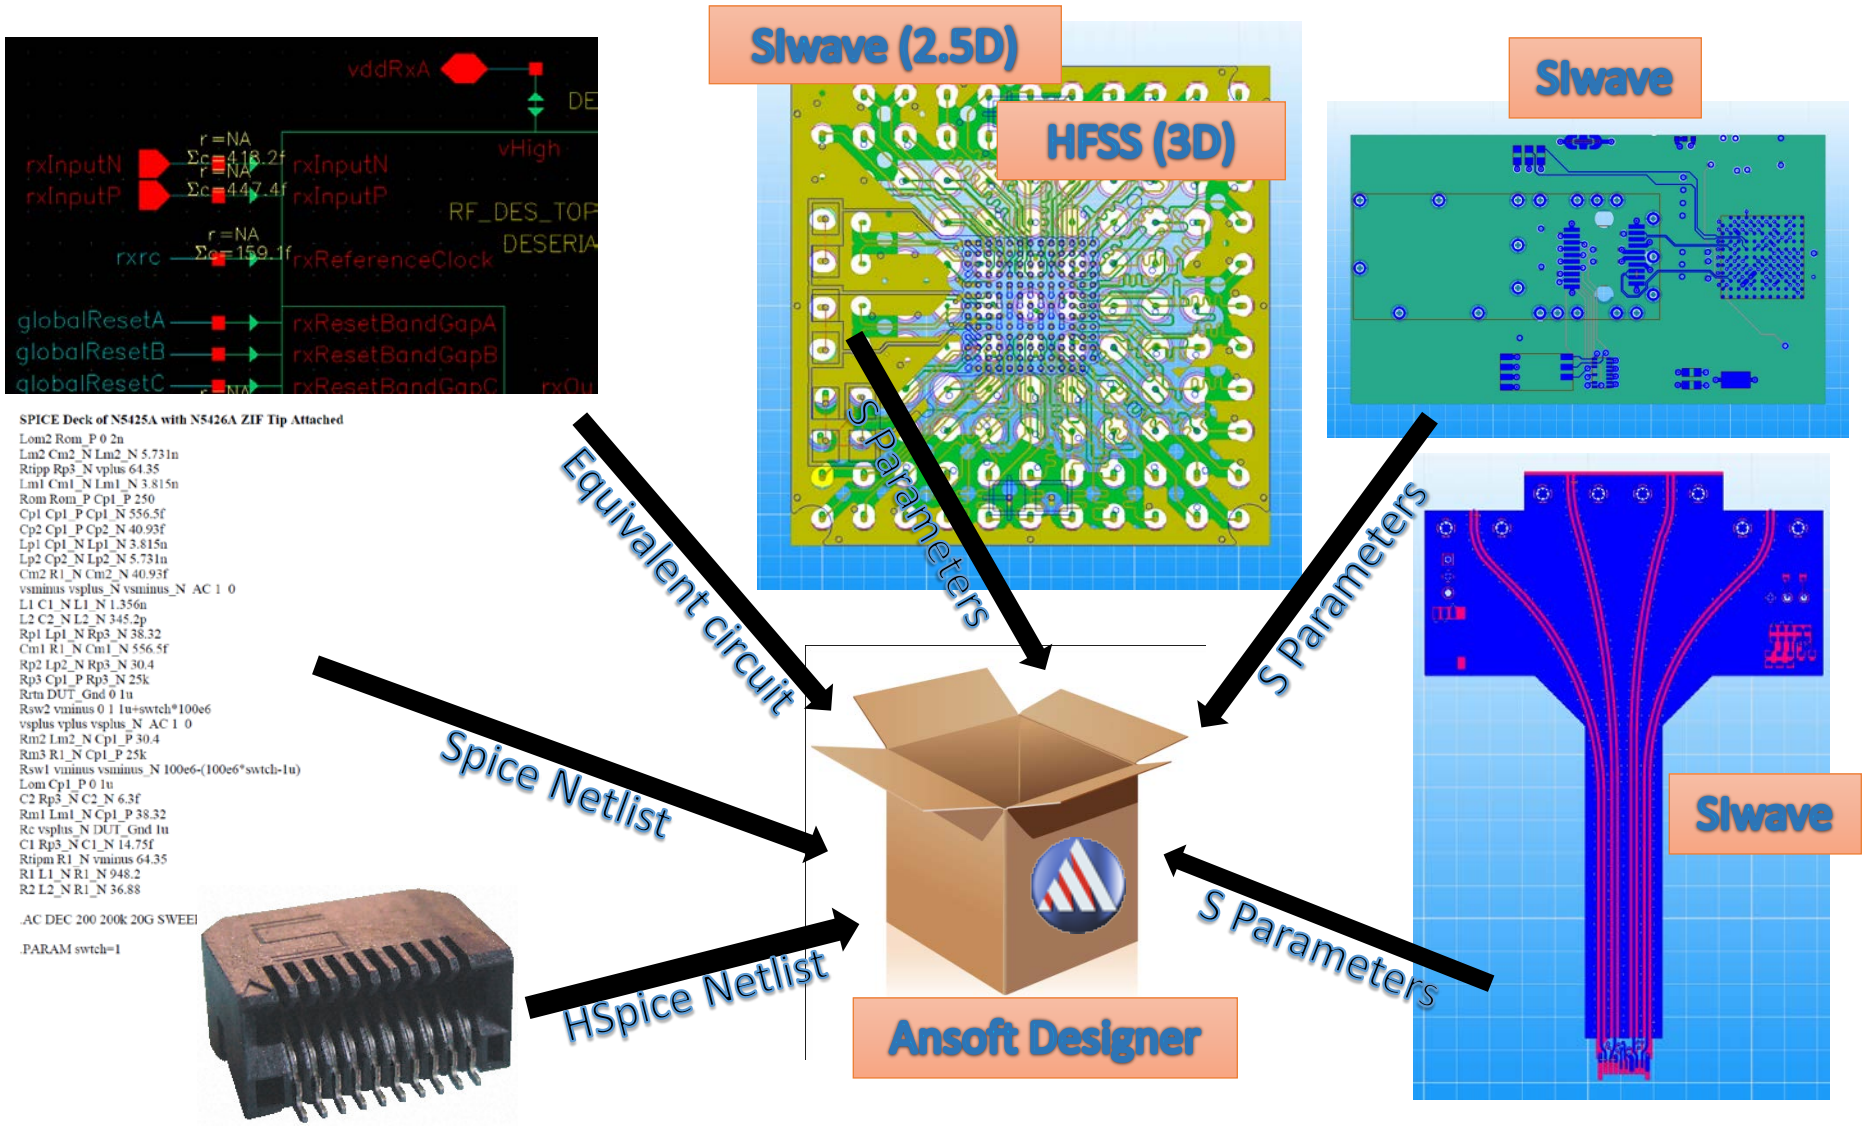

Fully Model the Signal Path

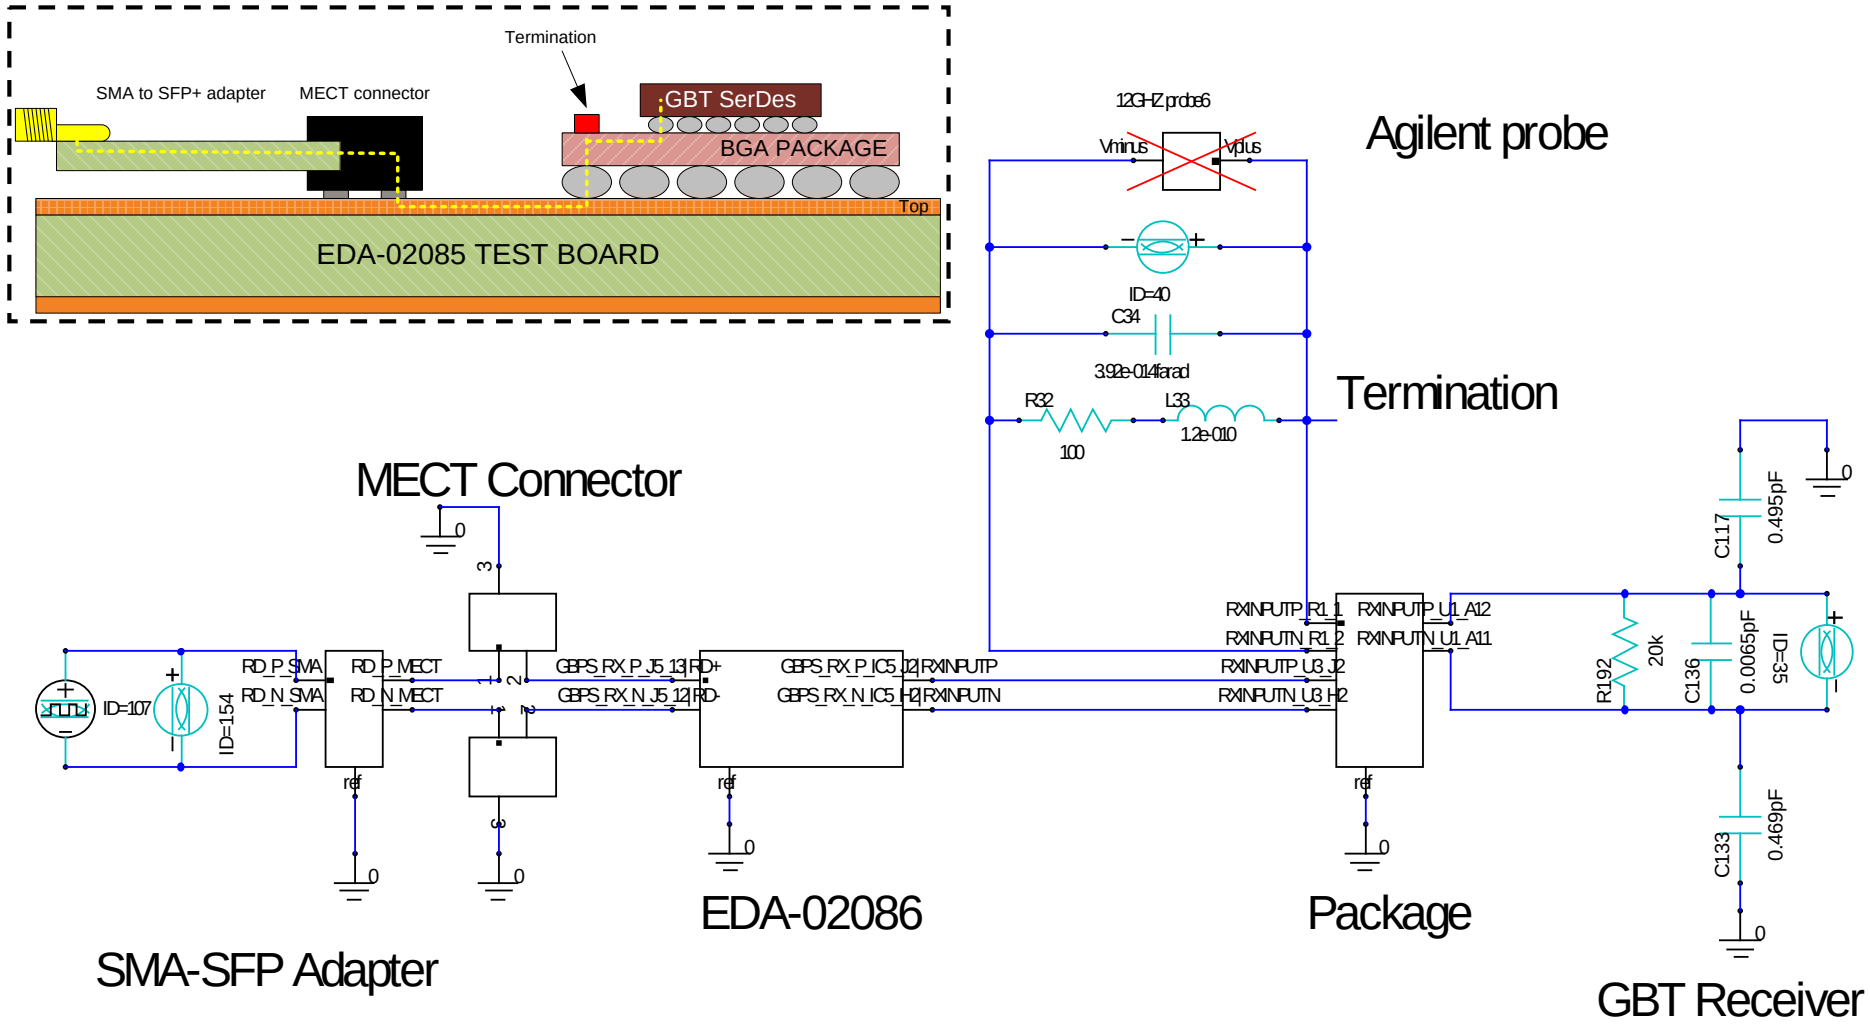

Unpopulated Board (no GBTX)

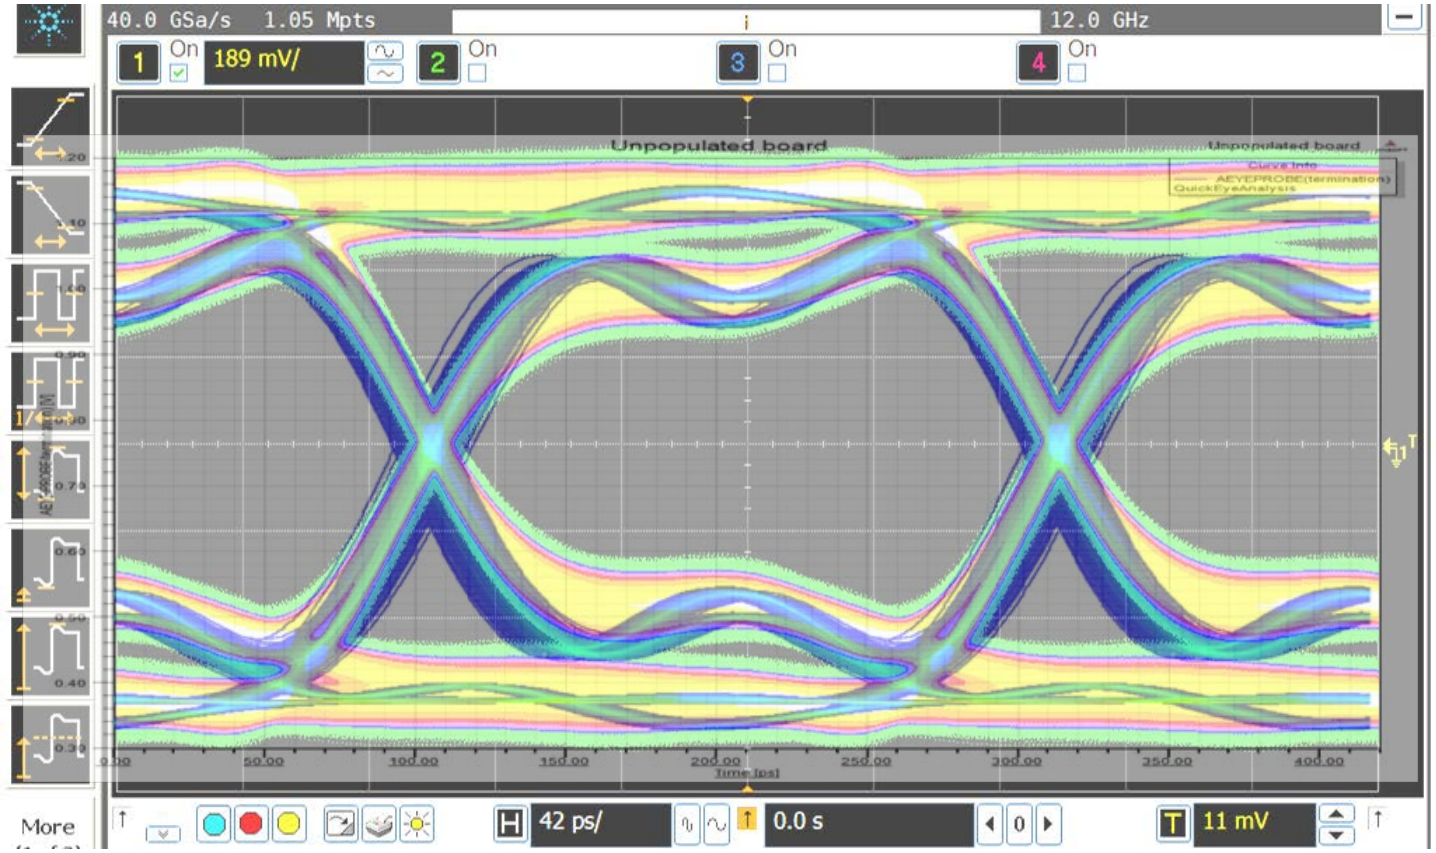

Measurement (yellow) and simulation (blue) with 100 ohms termination at the BGA pads

Fully Populated Board

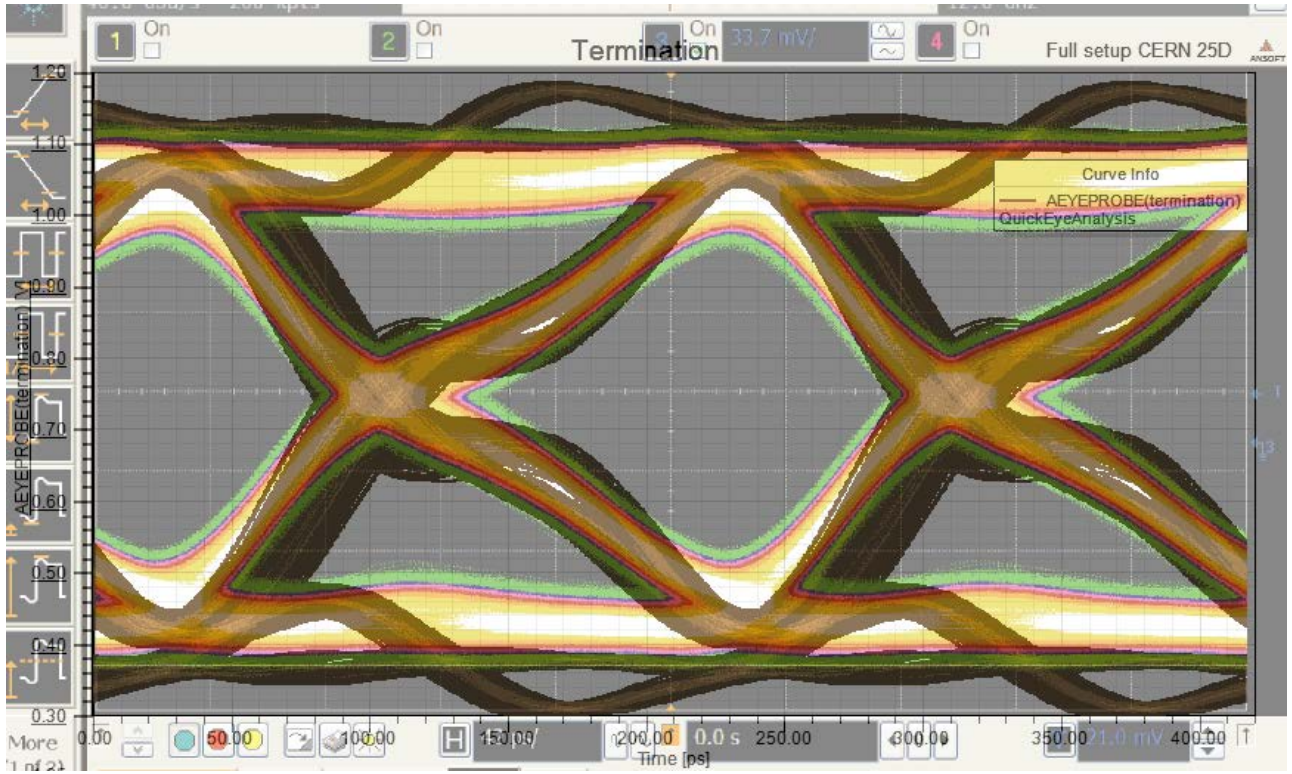

Probe is at the receiver termination , ≈ 3.5 mm from the input buffer.
Measurement (yellow) and Simulation (Brown).

Checking the Models

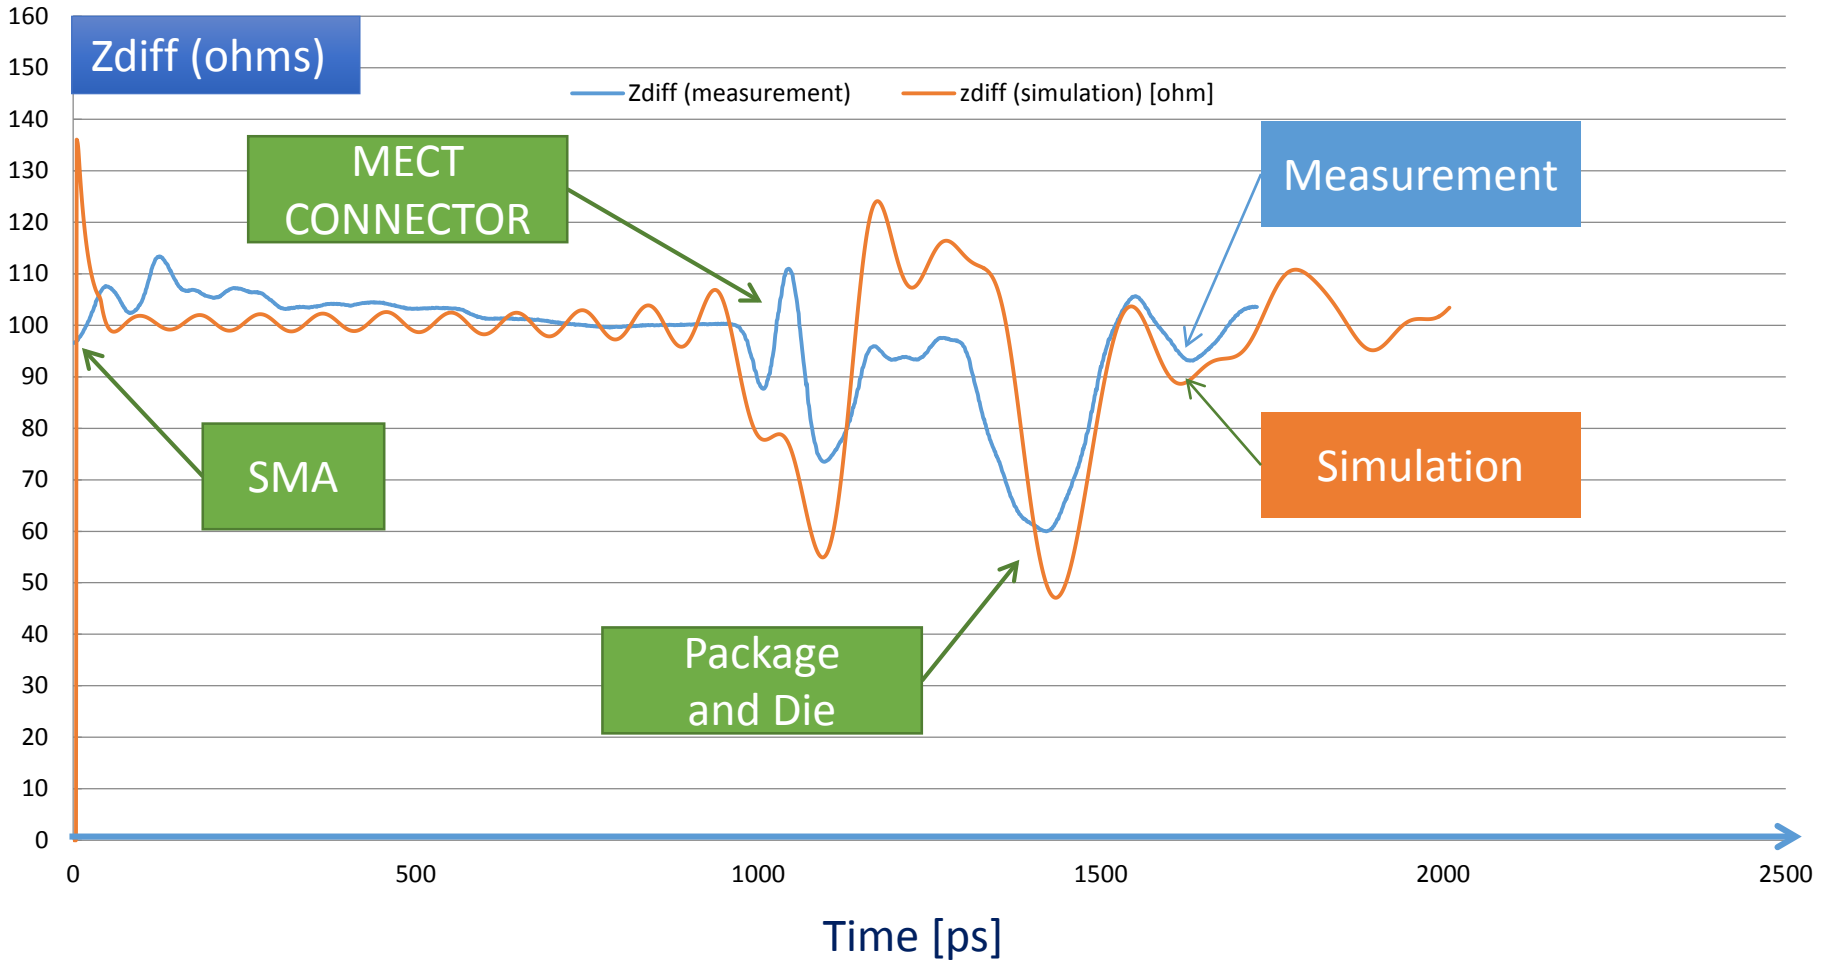

Schrodinger In High Frequency Electronics!?

- Can signals be observed without being disturbed?
- Manufacturers provide equivalent electrical modes for their oscilloscope probes
- Simulate to evaluate how much the system is being disturbed
- In our case the loading effect of the probe was small!

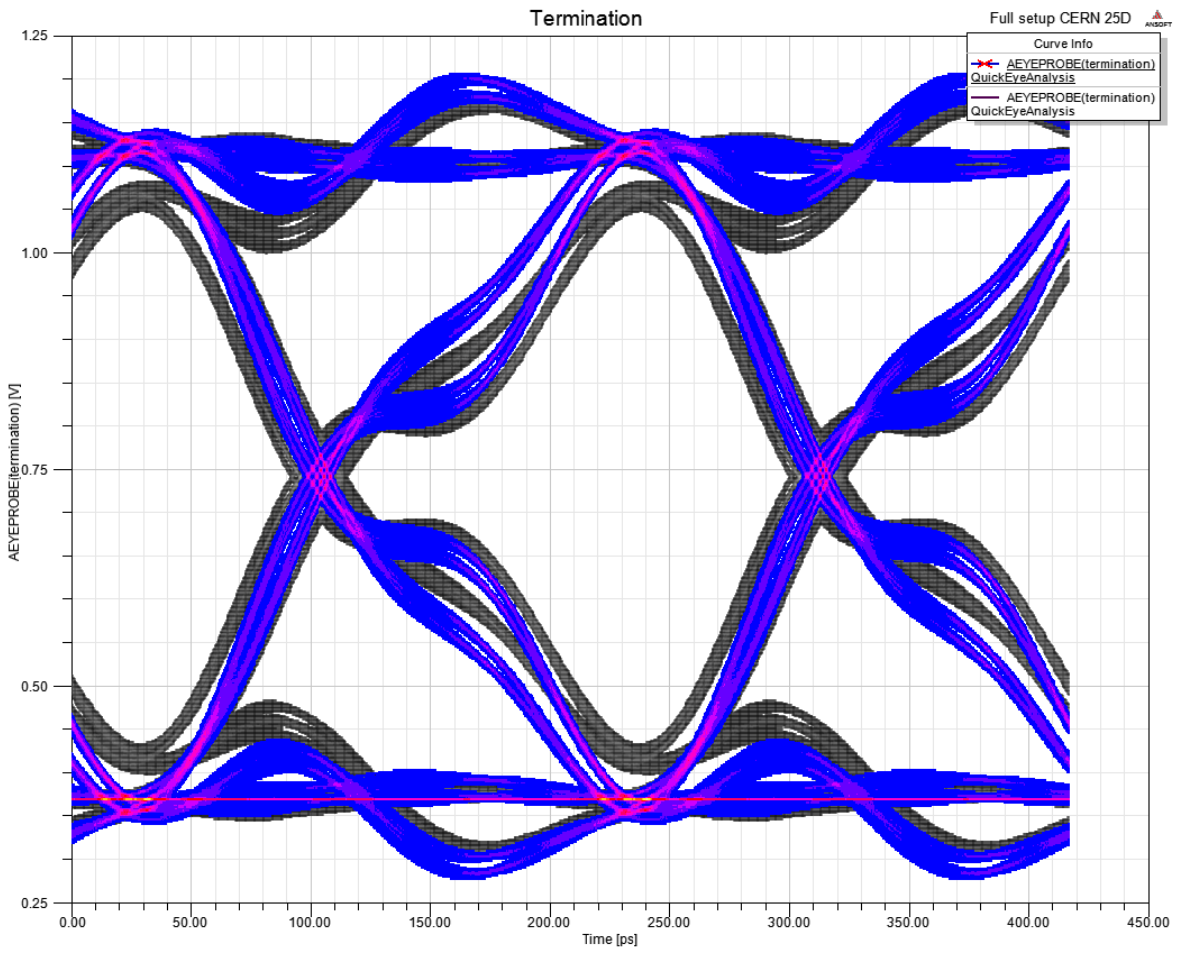

With Probe (gray)

No Probe (blue)

“Virtual Probe” - Simulation

1Gbps

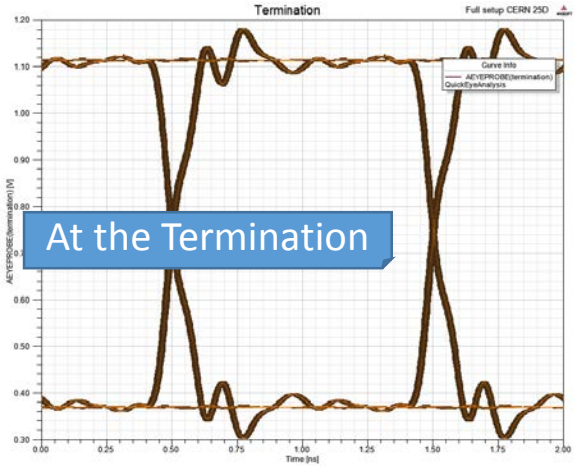

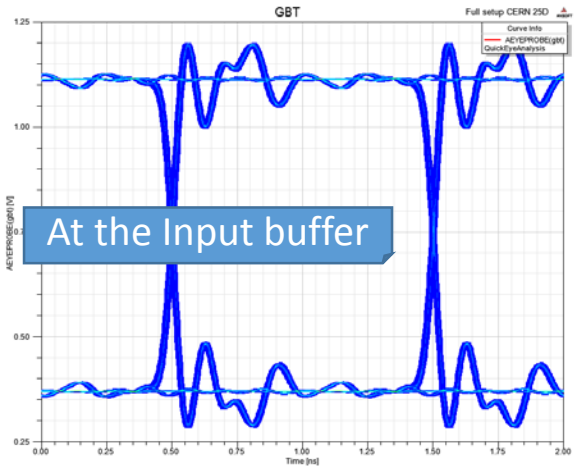

3.2Gbps

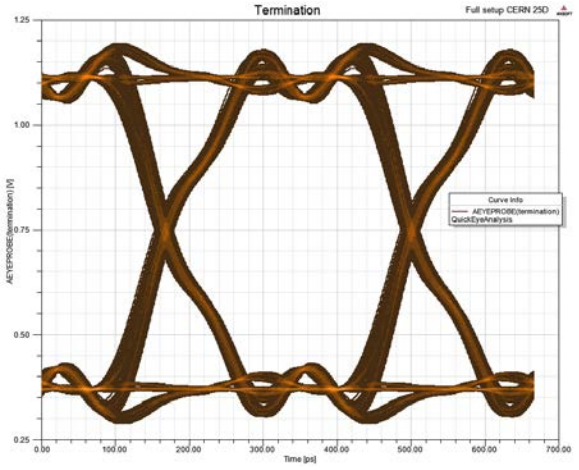

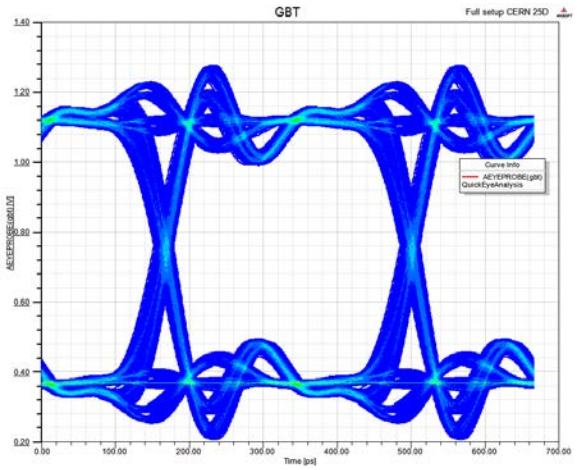

4.8Gbps

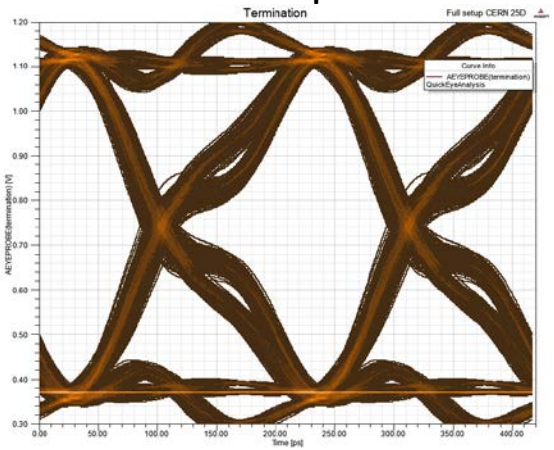

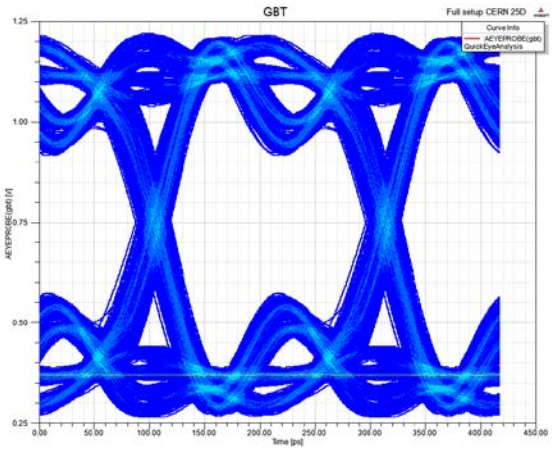

Improving the Package

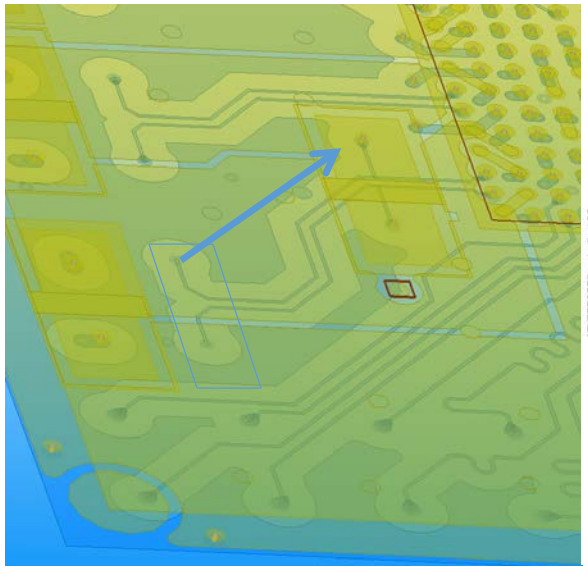

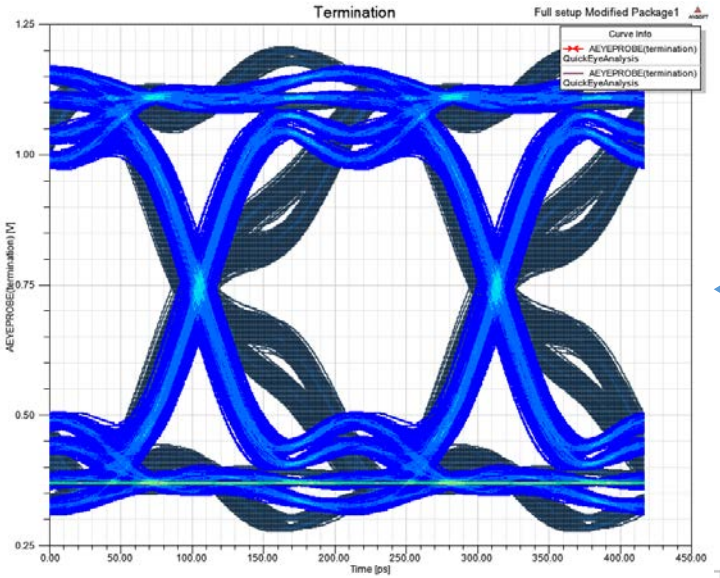

Original design: Dark blue

At the Termination

Move the termination close to the input buffer

At the Input buffer

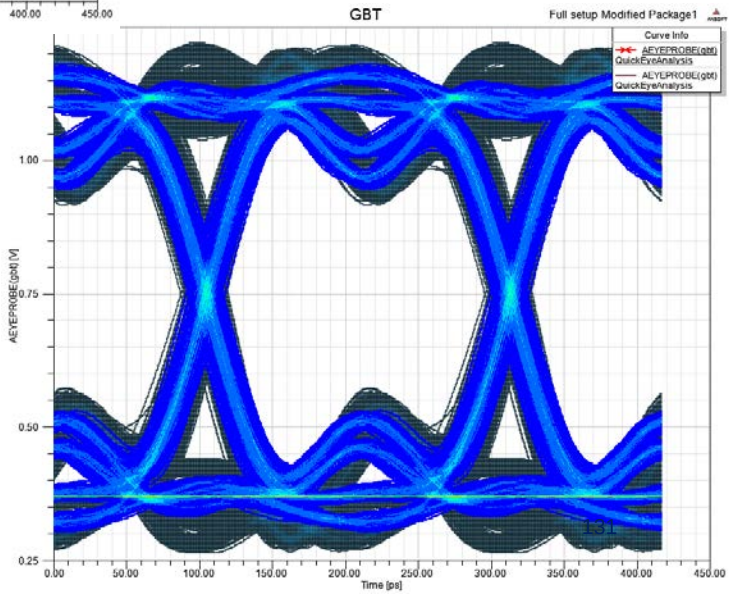

Credits

This talk reflects the work of many people in the GBT and LpGBT projects and it wouldn’t have been possible without their contribution. So I do collectively thank them.

In particular some people have prepared specific simulations, drawings,… for this set of slides. I specially thank their help:

- Bram Faes
- Eduardo Mendes
- Pedro Leitao
- Quan Sun
- Rui Francisco
- Szymon Kulis