HPT IP core for high-speed links using Xilinx FPGAs

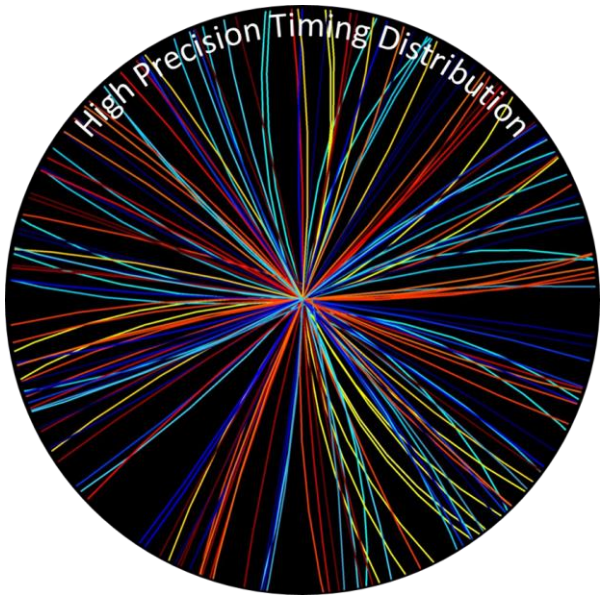

Eduardo Mendes

On behalf of the HPTD team (E. Mendes, S. Baron)

… many thanks to Jan Troska (CERN) and Paolo Novellini (Xilinx)

Motivation

- Investigation triggered by first CMS High Precision meeting
- Interesting to all LHC experiments
- Targeted Ultrascale(+) GTH/GTY transceivers

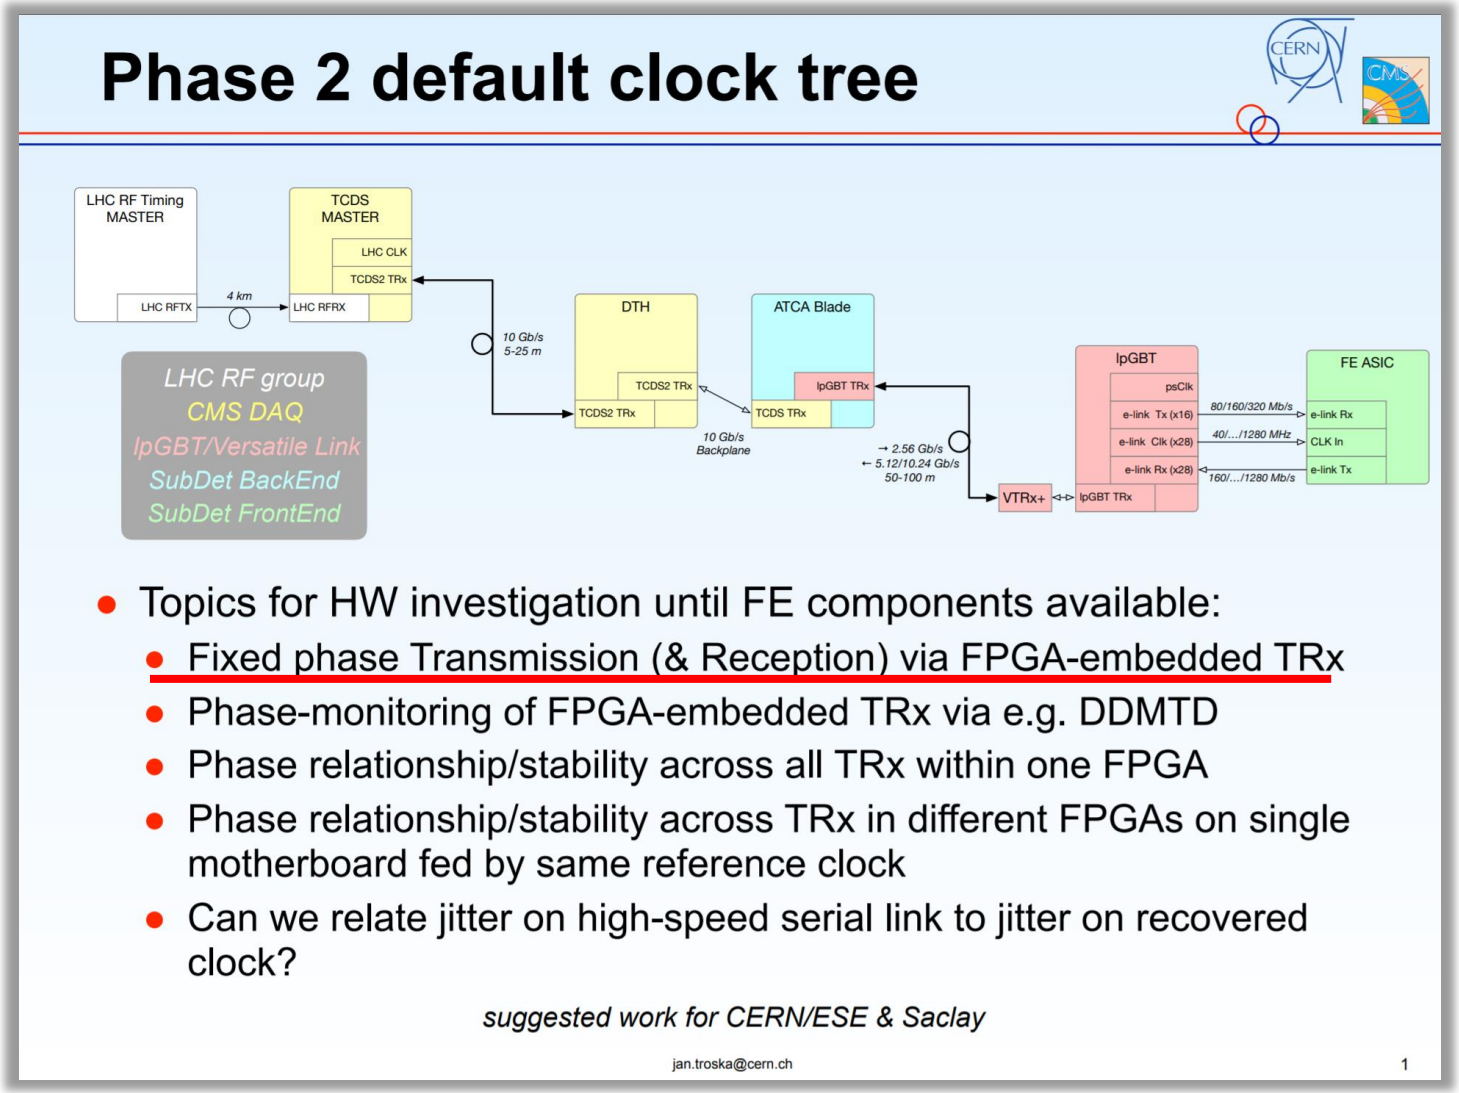

Outline

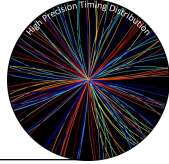

- **Introduction**
- Core architecture
  - Overview
  - Reference design
  - Characterization
- Link cascade
  - Rx recovered clock
  - Application of phase aligner
  - Characterization
- Conclusions

Introduction

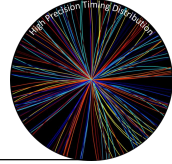

- Traditional implementation of minimal latency variation phase on Xilinx FPGAs is done using the **buffer-bypass** technique
  - GBT-FPGA, TTC-PON, …
- How does it work?

Introduction – Tx buffer bypass

- Phase wr x rd CDC?
  - Divider (÷N) phase

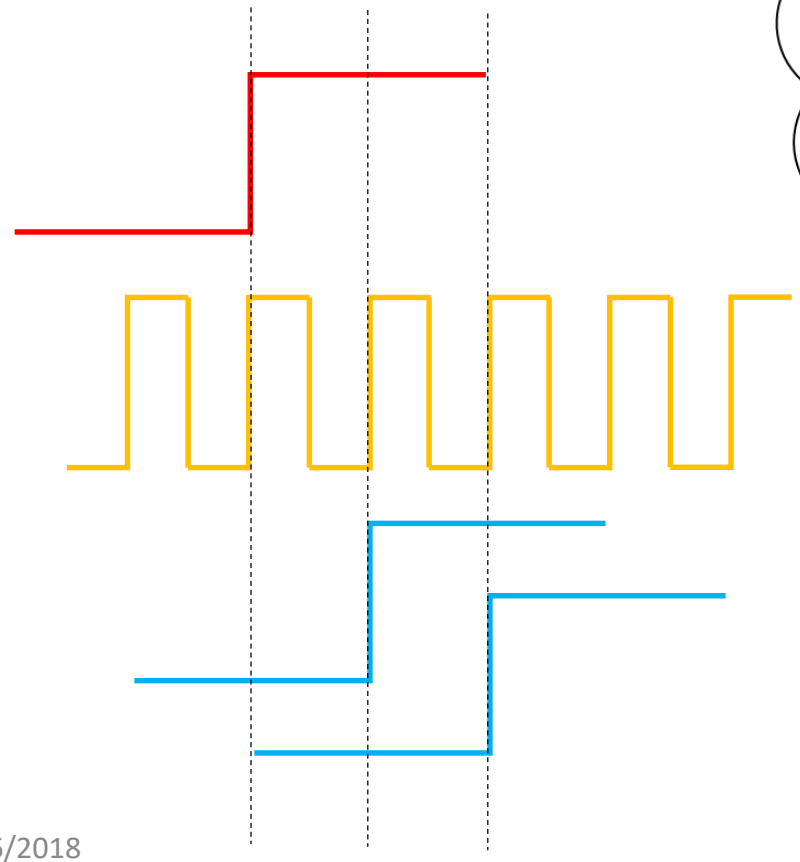

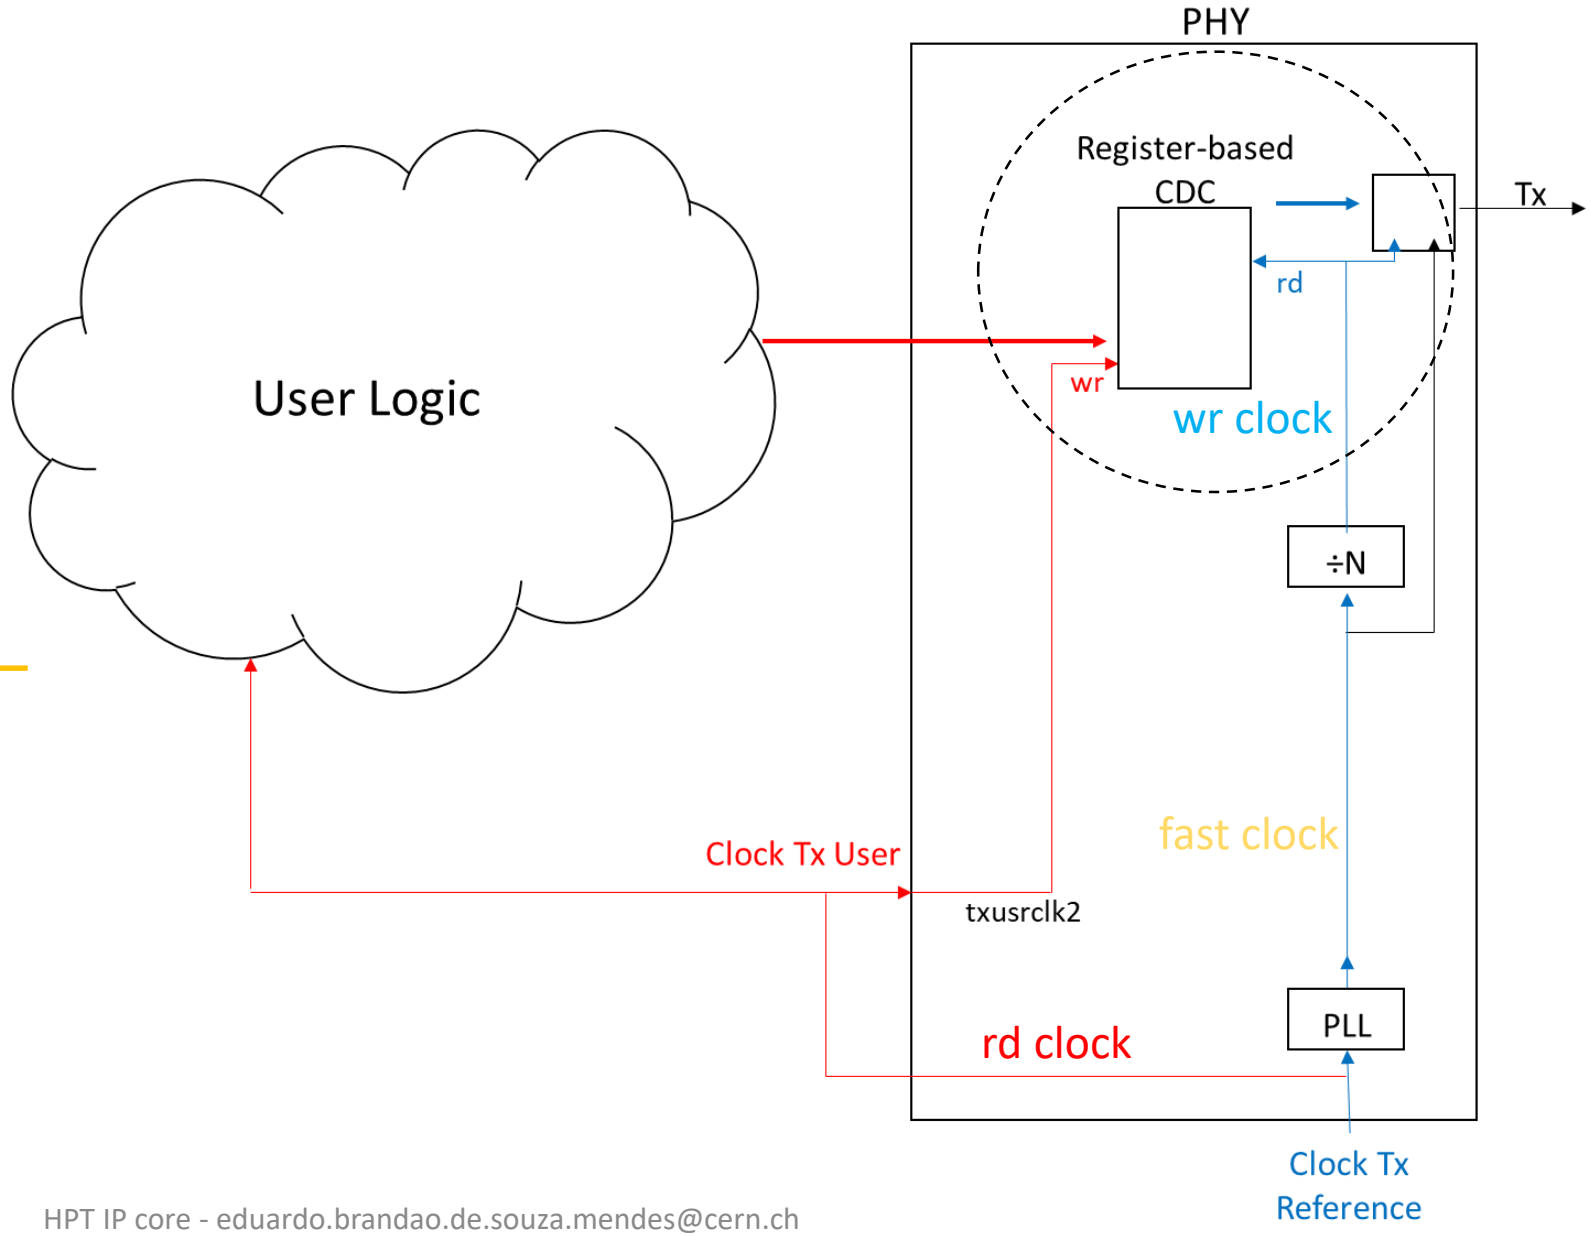

Introduction – Tx buffer bypass

- Phase wr x rd CDC?
  - Divider (÷N) phase
- (1) Tx PI (phase-shift)
  - Initialization
- (2) Delay Aligner (DA)
  - On-the-fly compensation

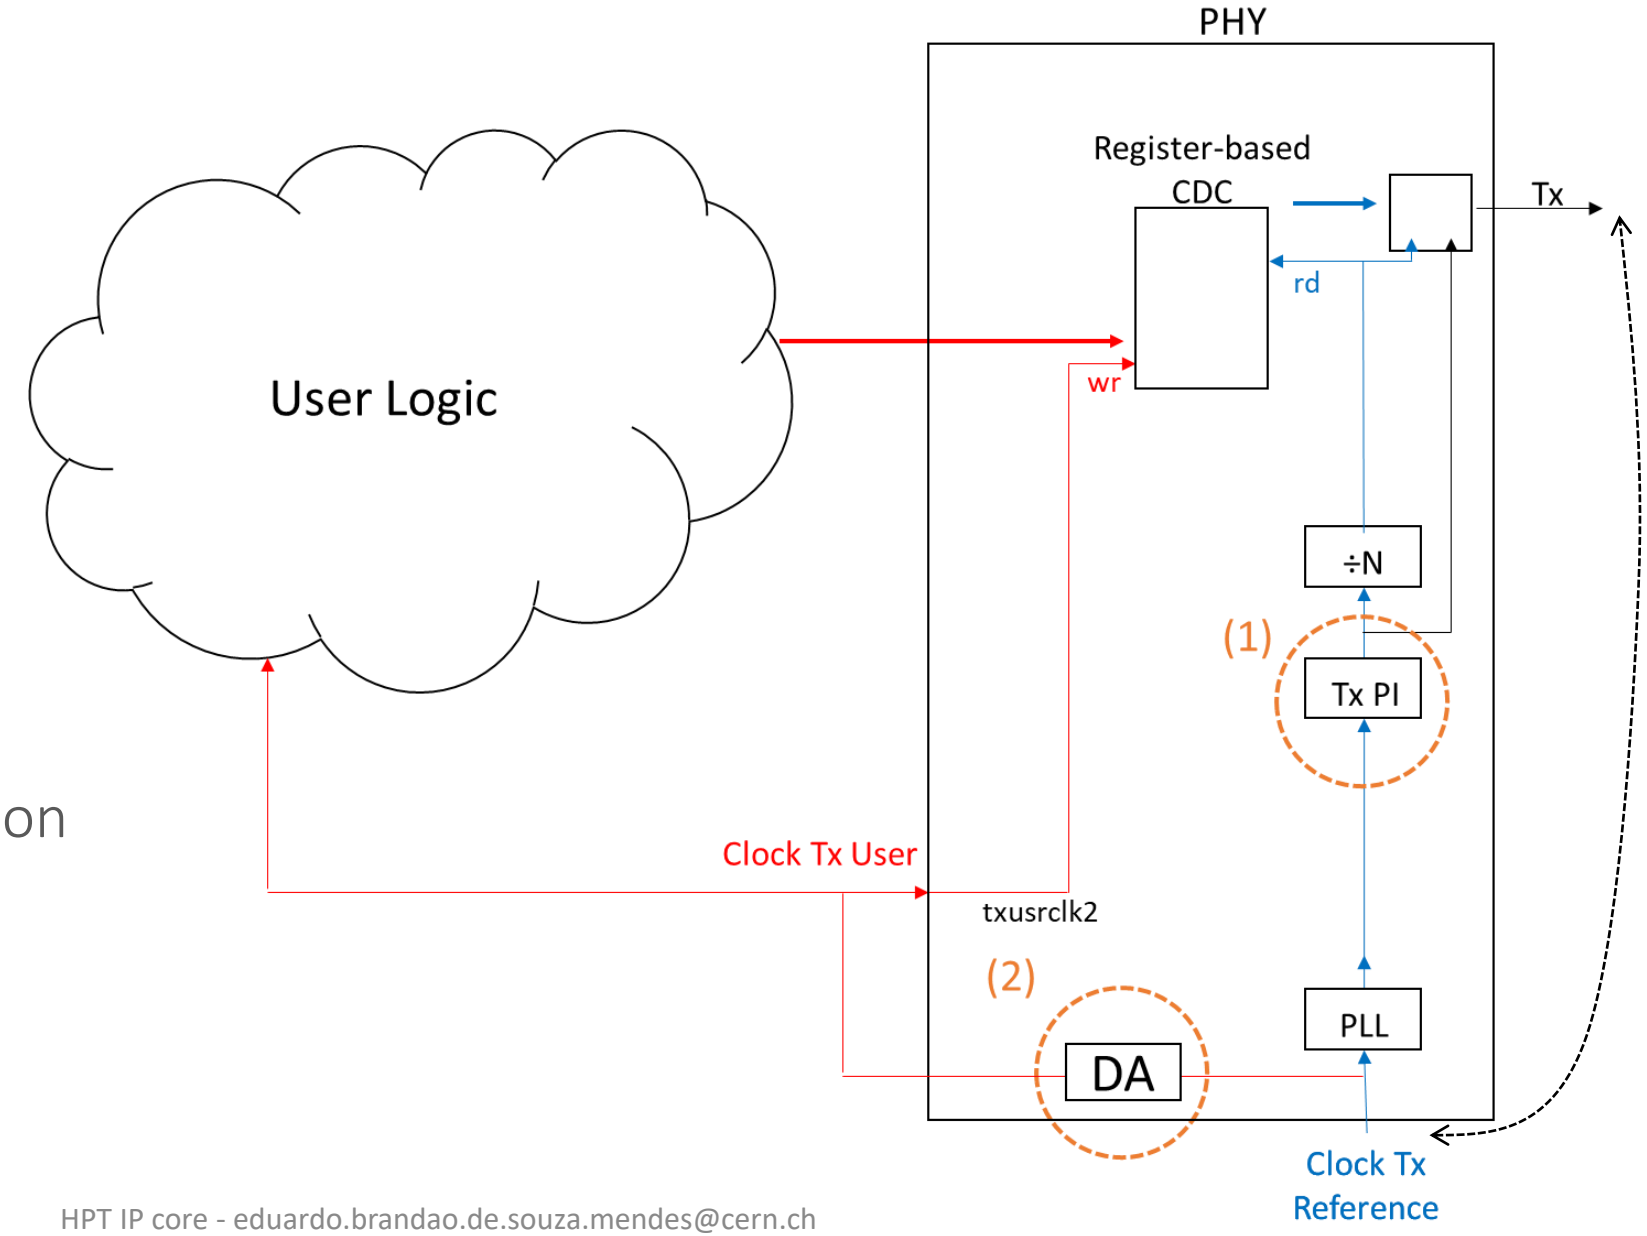

Introduction – deterministic phase requirements?

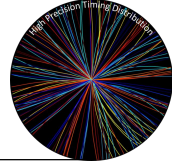

- For HL-LHC
  - Requirement on deterministic phase after start-up is not entirely clear yet
  - Physics phase monitoring?

- Implemented a new technique for deterministic latency on Xilinx FPGA which can be potentially useful for HL-LHC
  - Concept suggested by Xilinx engineer (Paolo Novellini - https://indico.cern.ch/event/598467/)

- Briefly explained in last talk
  - First proof-of concept for KU(+) implemented in software

- This talk:
  - In depth overview of technique – light VHDL core developped
  - Characterization for Kintex Ultrascale GTH
  - Link cascade concept

Outline

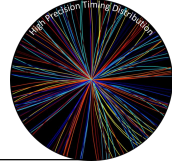

- Introduction
- **Core architecture**
  - Overview
  - Reference design
  - Characterization
- Link cascade
  - Rx recovered clock
  - Application of phase aligner
  - Characterization
- Conclusions

Core architecture: overview

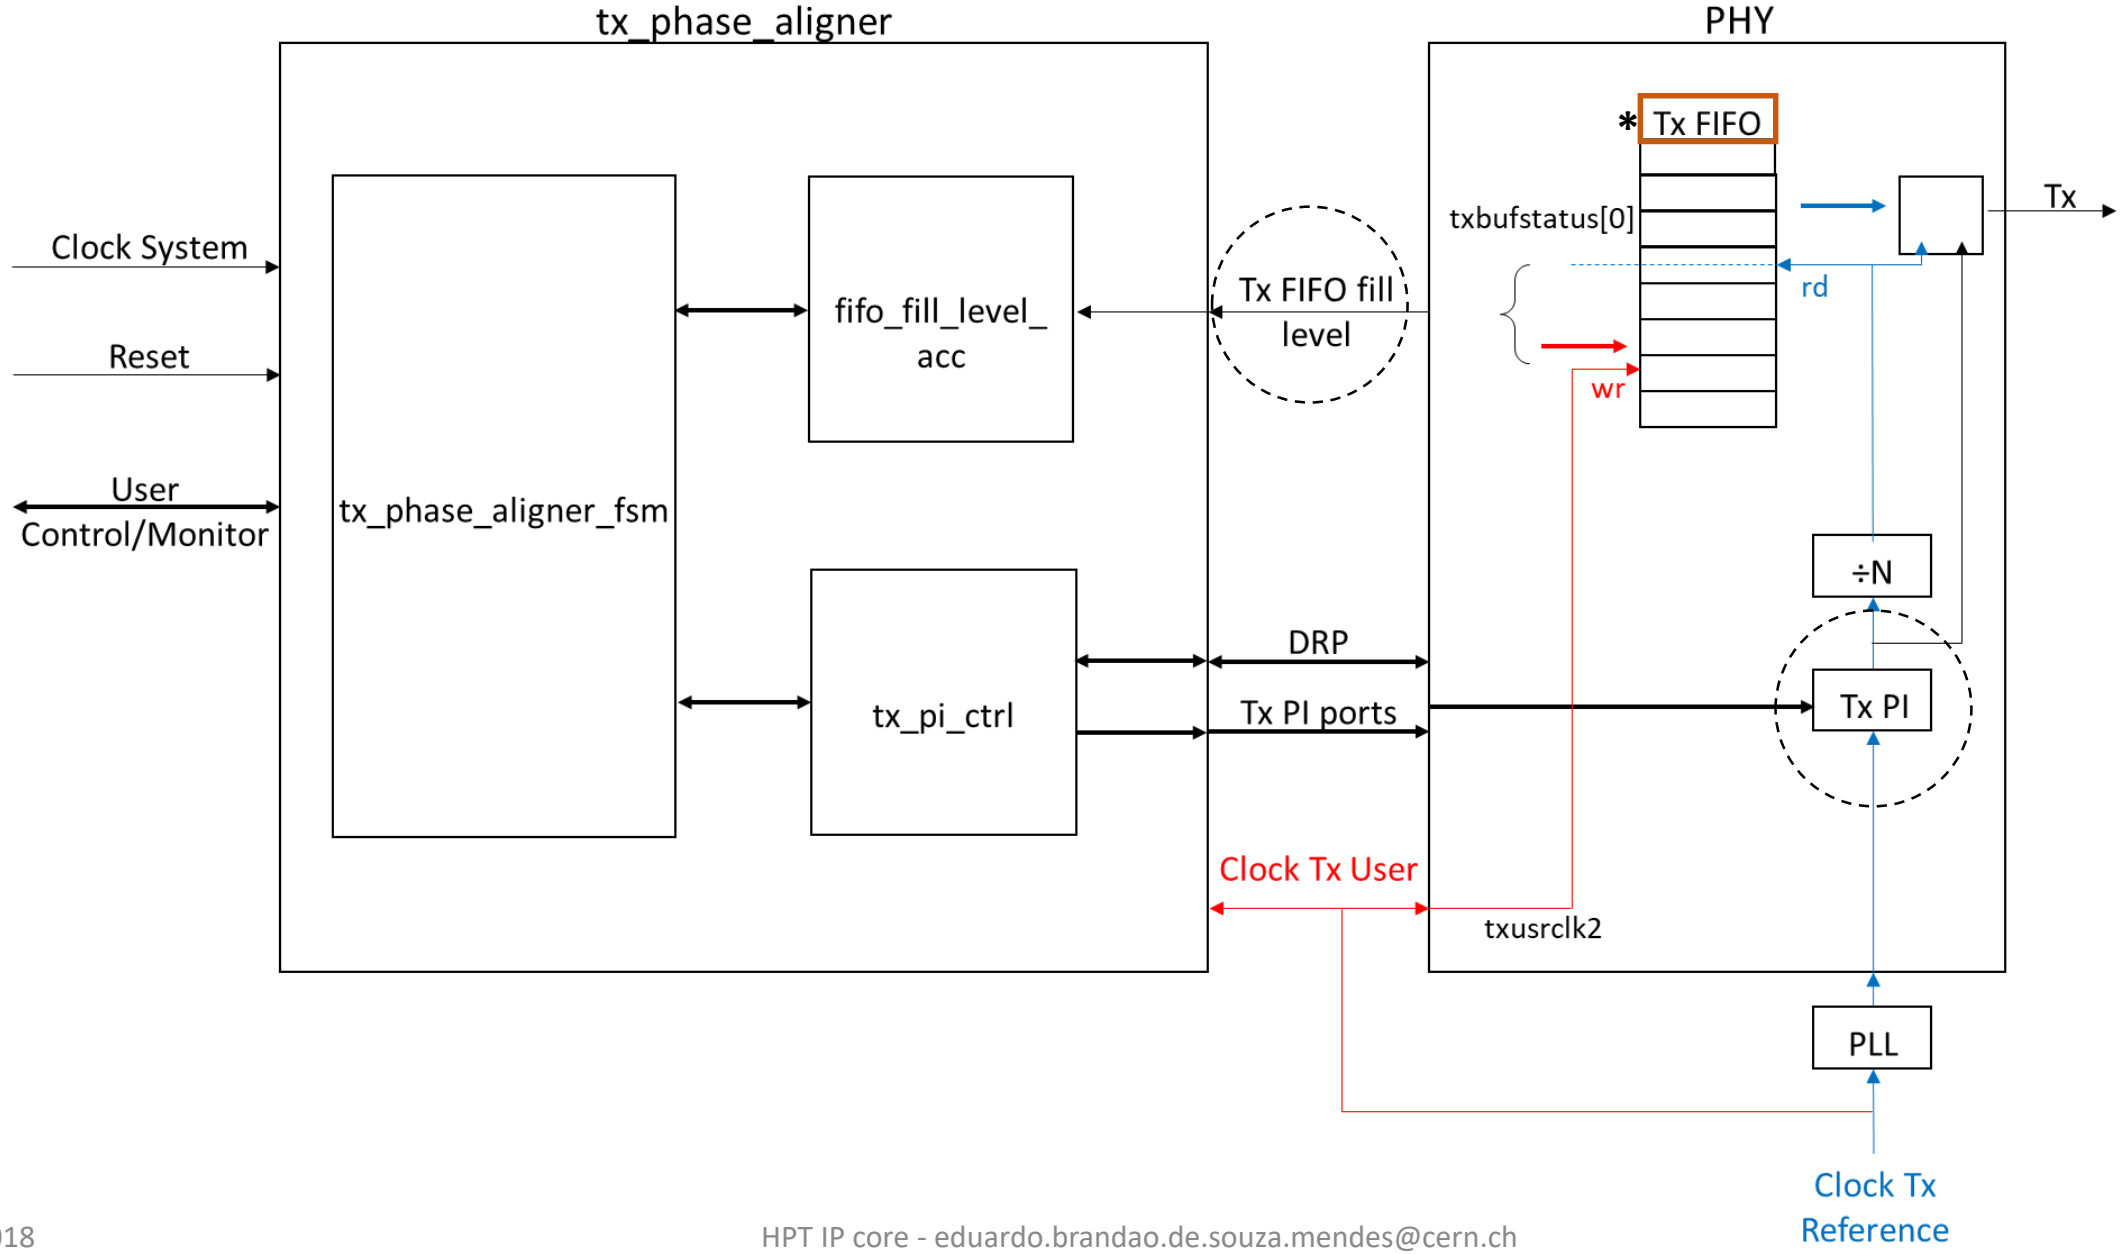

Core architecture: overview

- Flavours
  - Minimal variation (**FINE_ALIGN**)
  - Deterministic latency (**UI_ALIGN**)
- Protocol-independent
  - **It could work on any existing link (simple patch)**
  - requires minor modification on transceiver

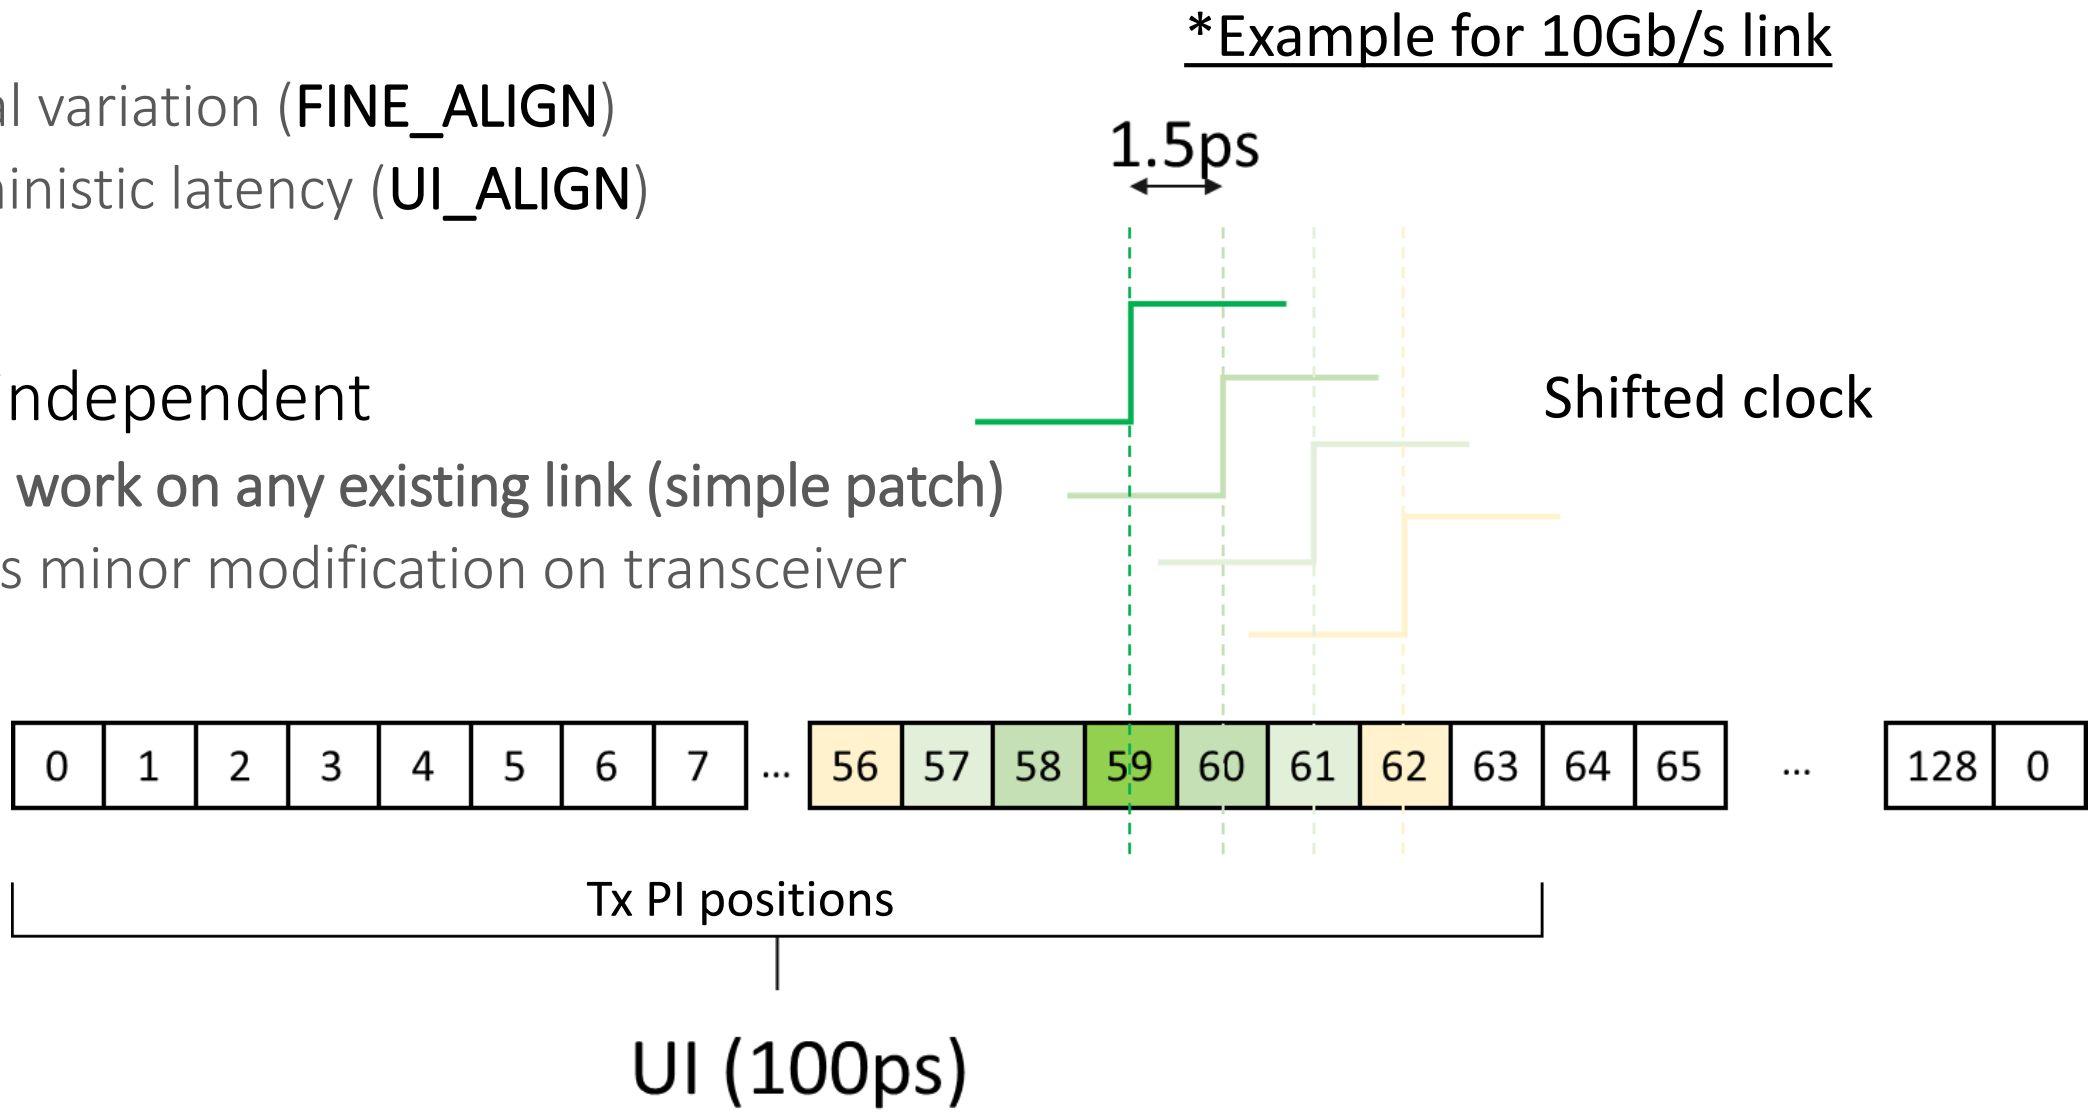

Core architecture: reference design

- Reference design targetting **KCU105** board
  - Kintex Ultrascale FPGA
  - Inspired on transceiver example design
  - https://gitlab.cern.ch/HPTD/tx_phase_aligner (under request - just contact us)
- Contain simulation + hardware
  - Vivado 2016.2
- **Resource usage (core only)**
  - Very light: ~120 CLB LUTs, ~120 CLB REGs
- Latency minor increase (w.r.t. buffer-bypass)
  - For KU-GTH: +2 Tx Word Cycles (320MHz)
- Compatible with LpGBT-FPGA data-rate (10.24Gb/s)

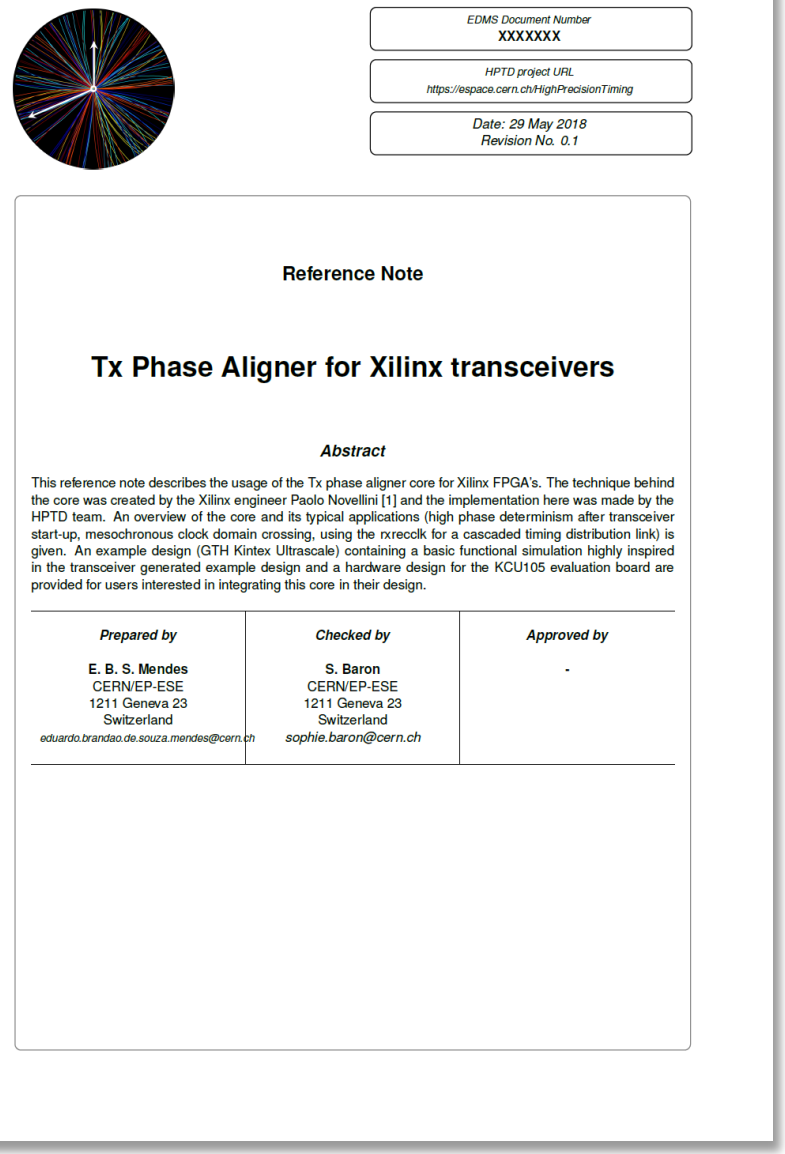

EDMS Document Number
XXXXXXX

HPTD project URL
https://espace.cern.ch/HighPrecisionTiming

Date: 29 May 2018
Revision No. 0.1

Reference Note

Tx Phase Aligner for Xilinx transceivers

Abstract

This reference note describes the usage of the Tx phase aligner core for Xilinx FPGA's. The technique behind the core was created by the Xilinx engineer Paolo Novellini [1] and the implementation here was made by the HPTD team. An overview of the core and its typical applications (high phase determinism after transceiver start-up, mesochronous clock domain crossing, using the rxrecclk for a cascaded timing distribution link) is given. An example design (GTH Kintex Ultrascale) containing a basic functional simulation highly inspired in the transceiver generated example design and a hardware design for the KCU105 evaluation board are provided for users interested in integrating this core in their design.

Prepared by	*Checked by*	*Approved by*
E. B. S. Mendes CERN/EP-ESE 1211 Geneva 23 Switzerland *eduardo.brandao.de.souza.mendes@cern.ch*	**S. Baron** CERN/EP-ESE 1211 Geneva 23 Switzerland *sophie.baron@cern.ch*	-

Characterization: Tx only setup

- FPGA board
  - Tested: KCU116 –GTY (shown in last talk), ZCU102-GTH, KCU105-GTH (shown here)

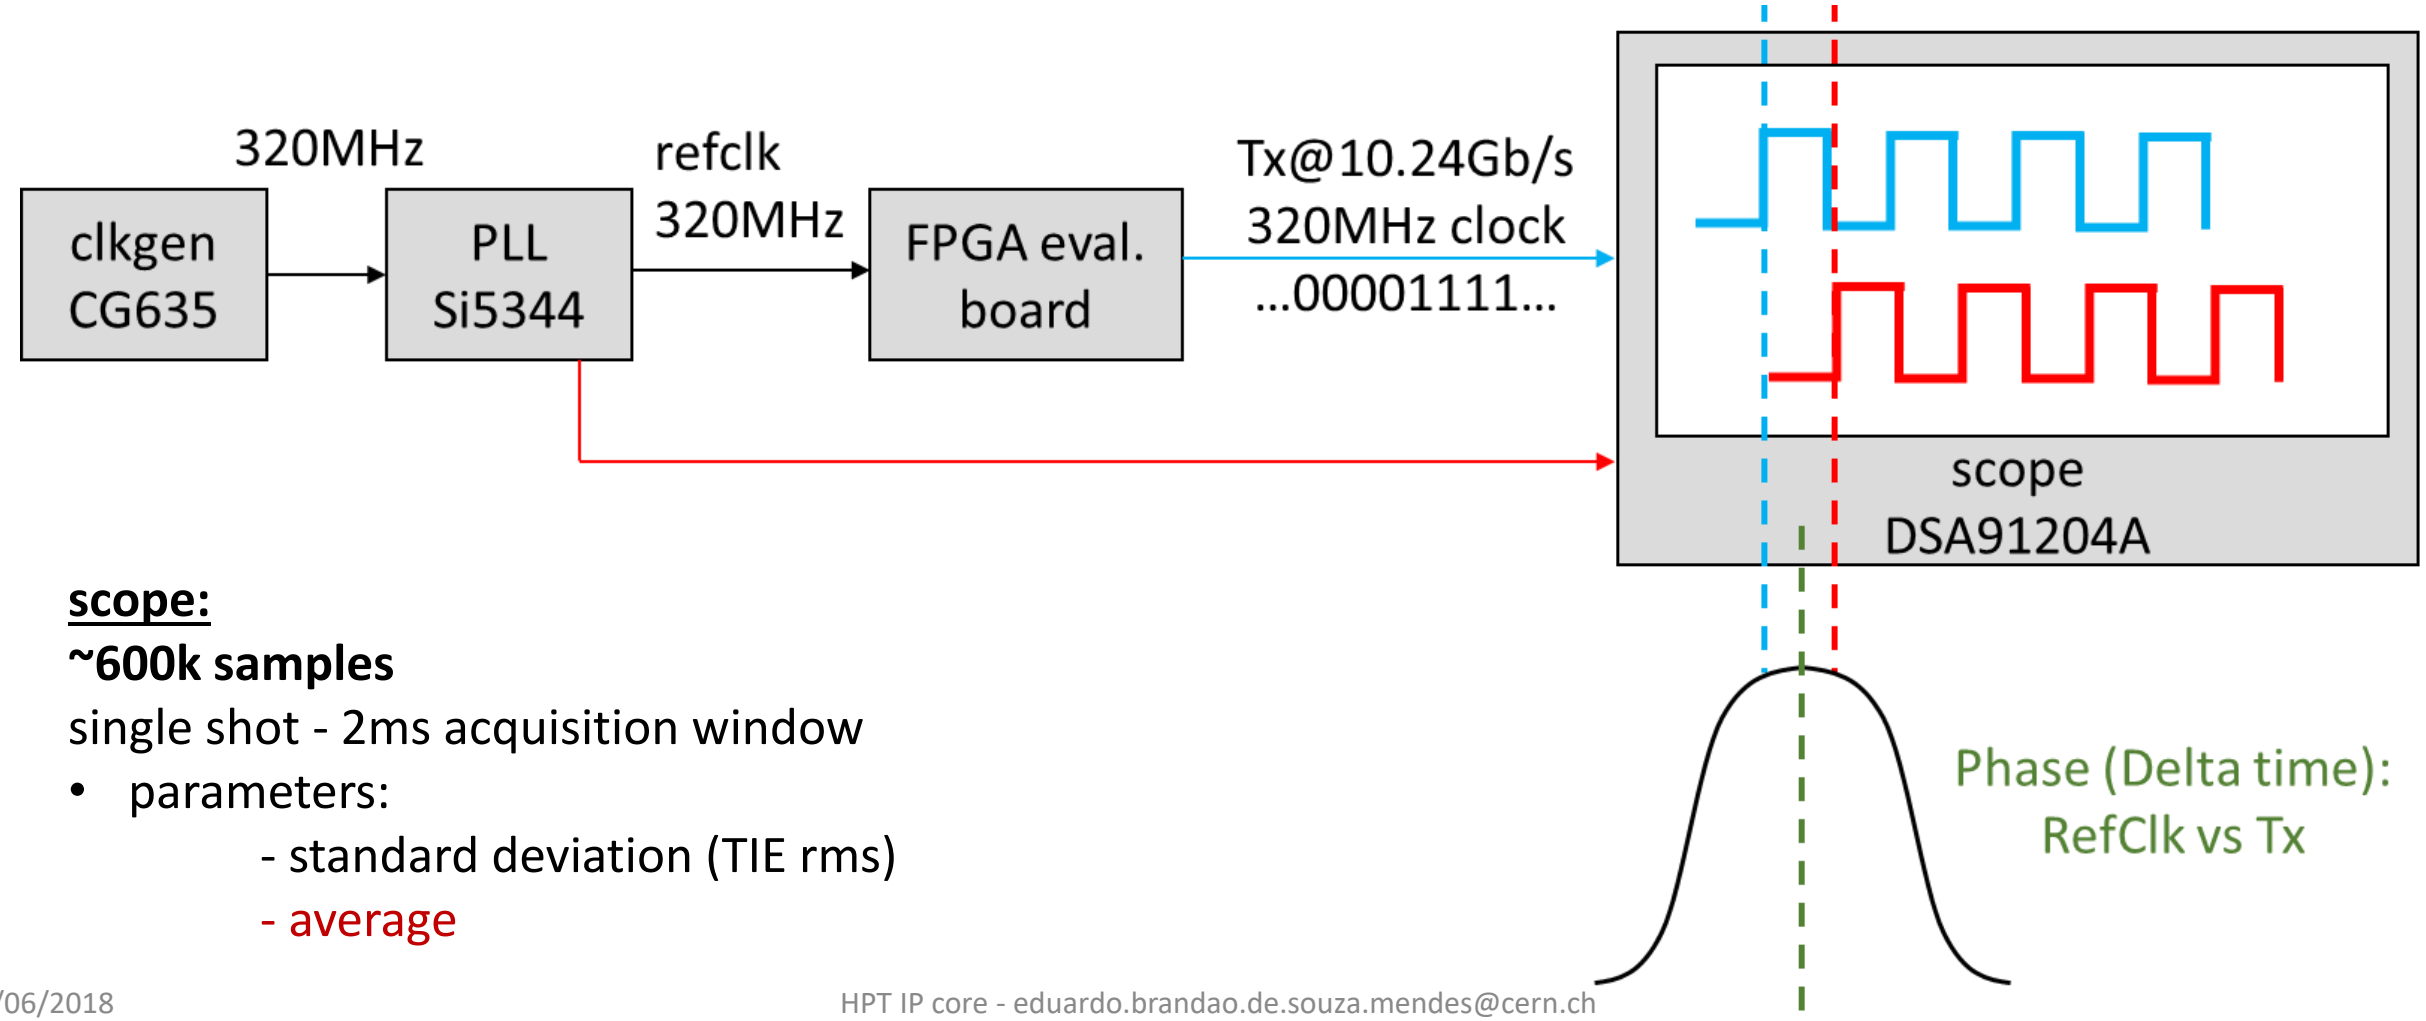

scope:

~600k samples

single shot - 2ms acquisition window

- parameters:
  - standard deviation (TIE rms)
  - average

Characterization: Tx only with buffer-bypass

- Buffer-Bypass: a.k.a. fixed latency (technique used for GBT-FPGA, TTC-PON)
- Average: reset at every acquisition

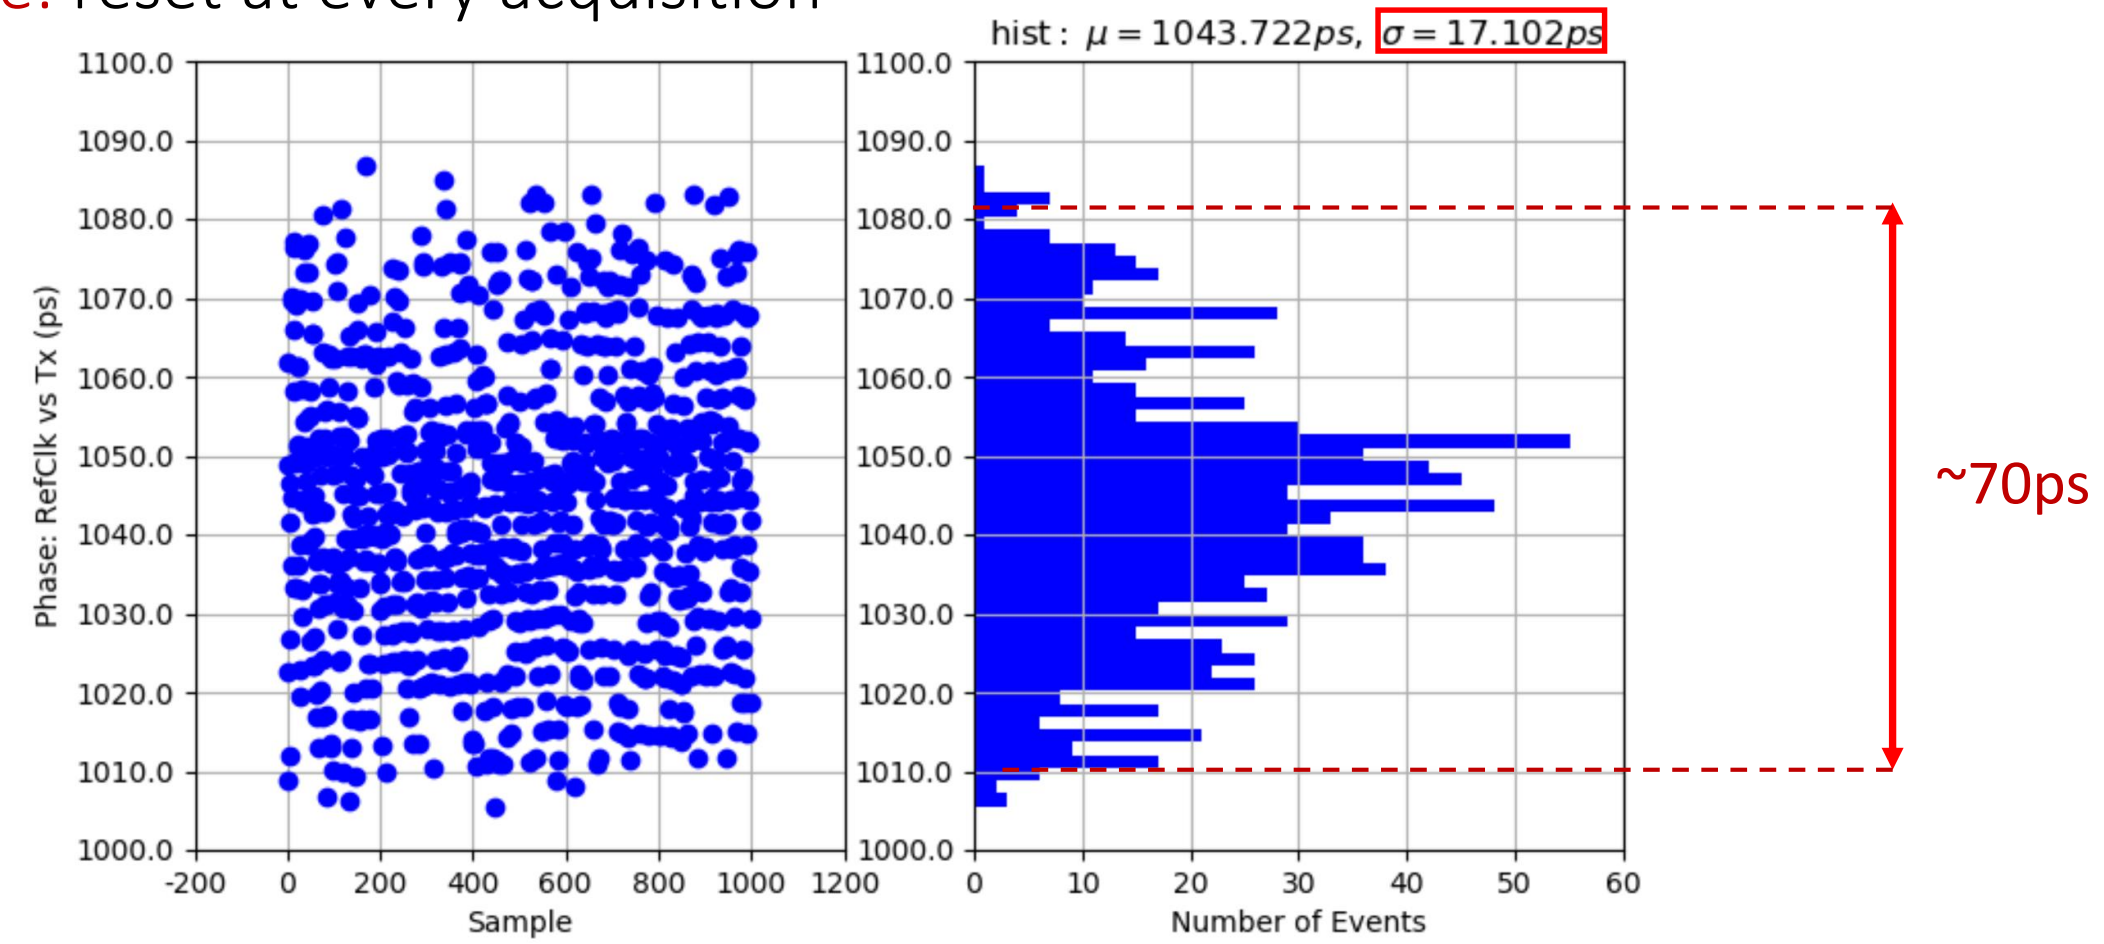

Characterization: Tx only with core presented

- Tx aligner core: **FINE_ALIGN** mode (i.e. minimal latency variation)
- Average: reset at every acquisition

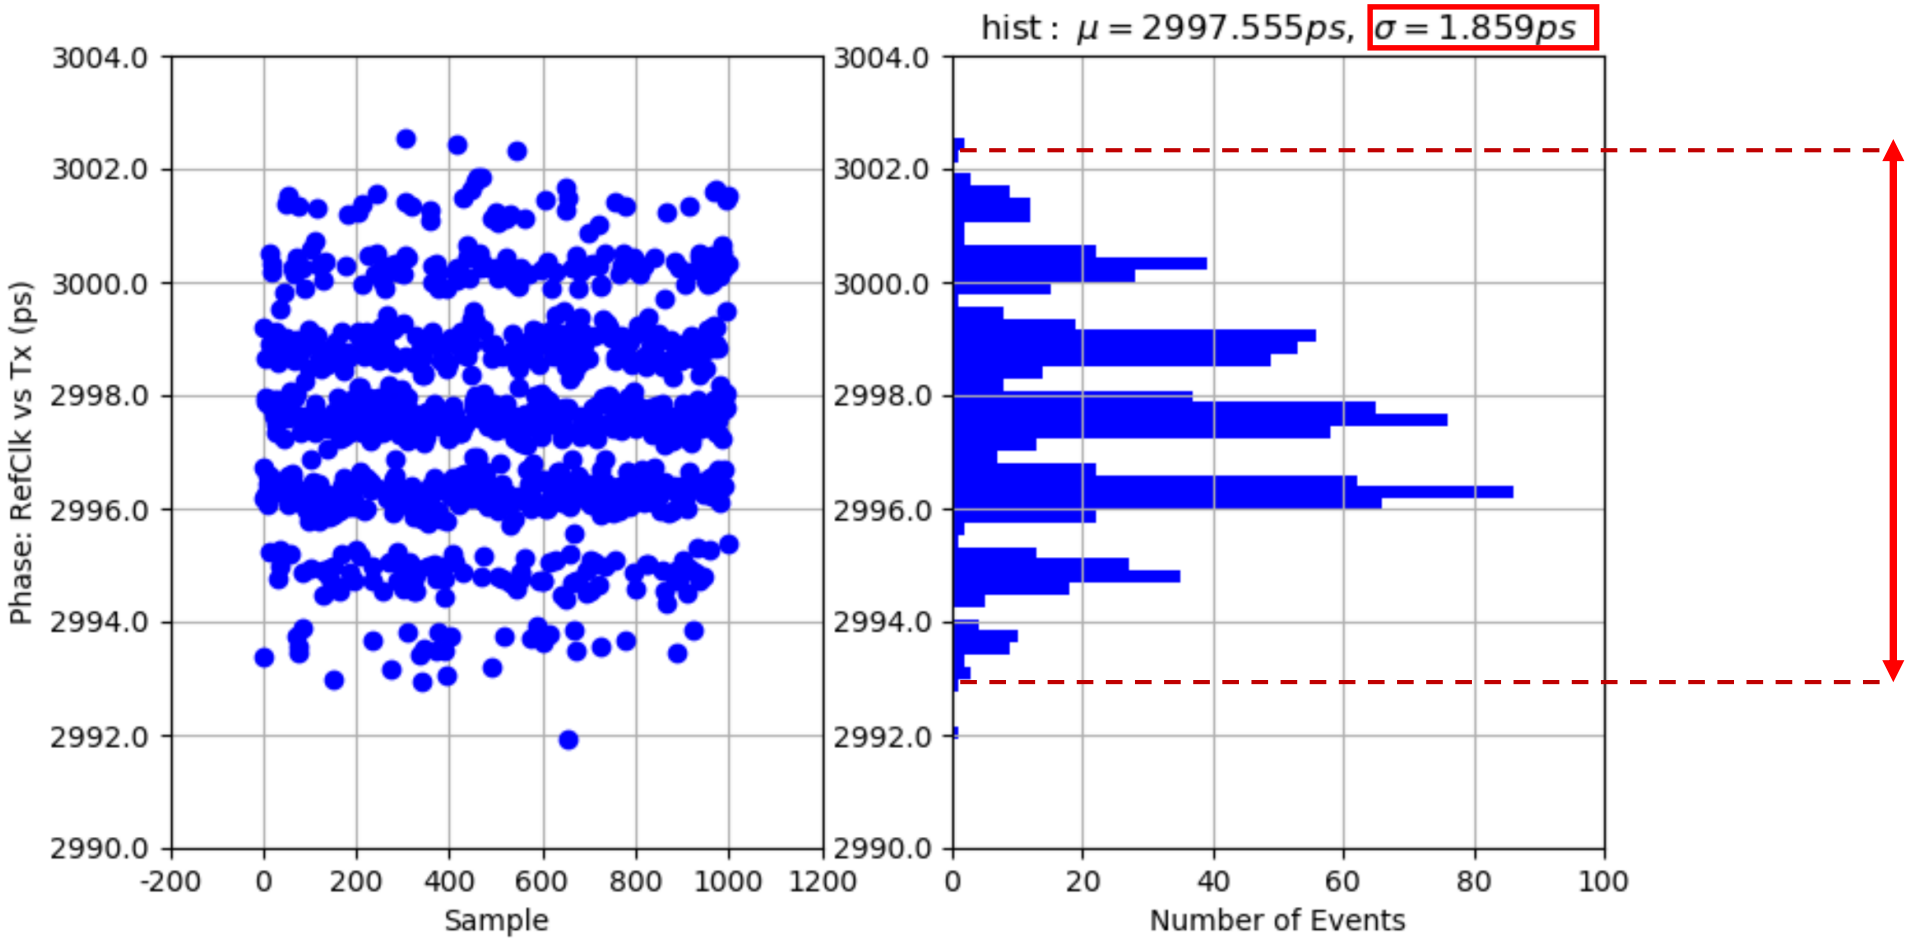

Characterization: Tx only with core presented

- Tx aligner core: **UI_ALIGN** mode (i.e. deterministic latency)
- Average: reset at every acquisition (results similar to no-reset)

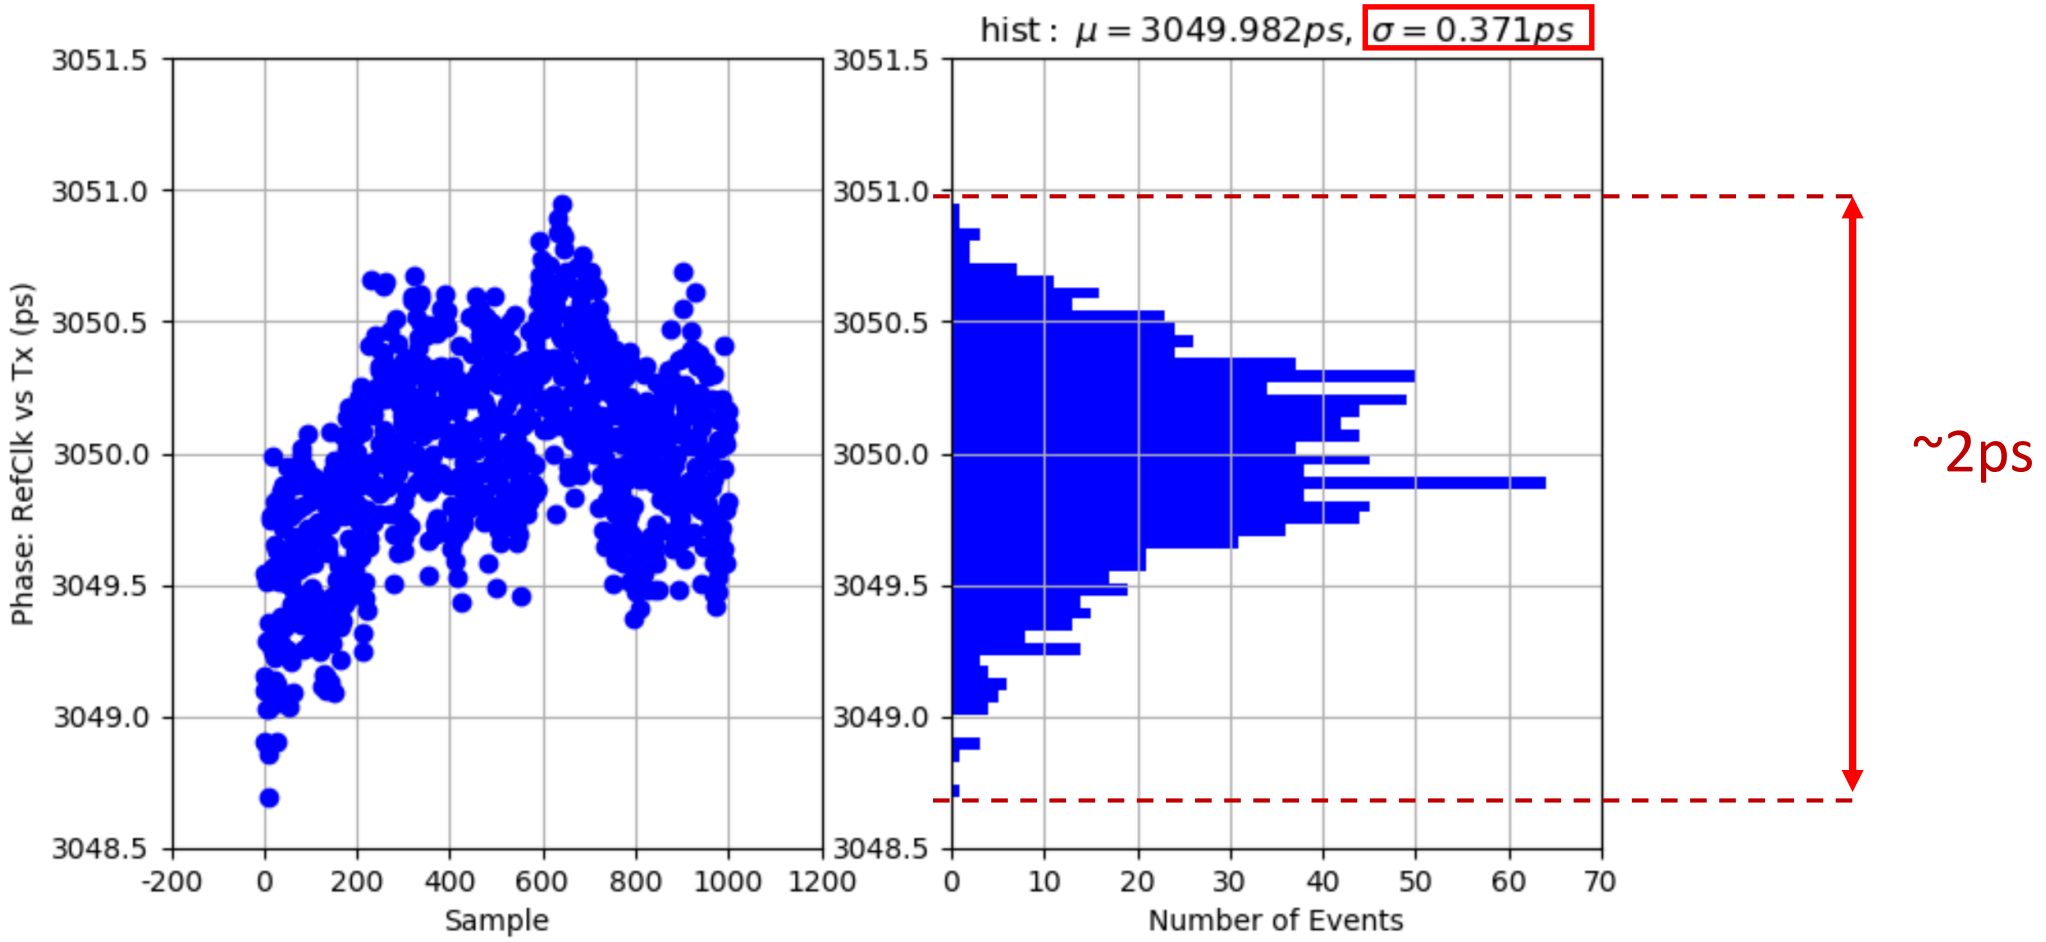

Tx phase aligner: summary

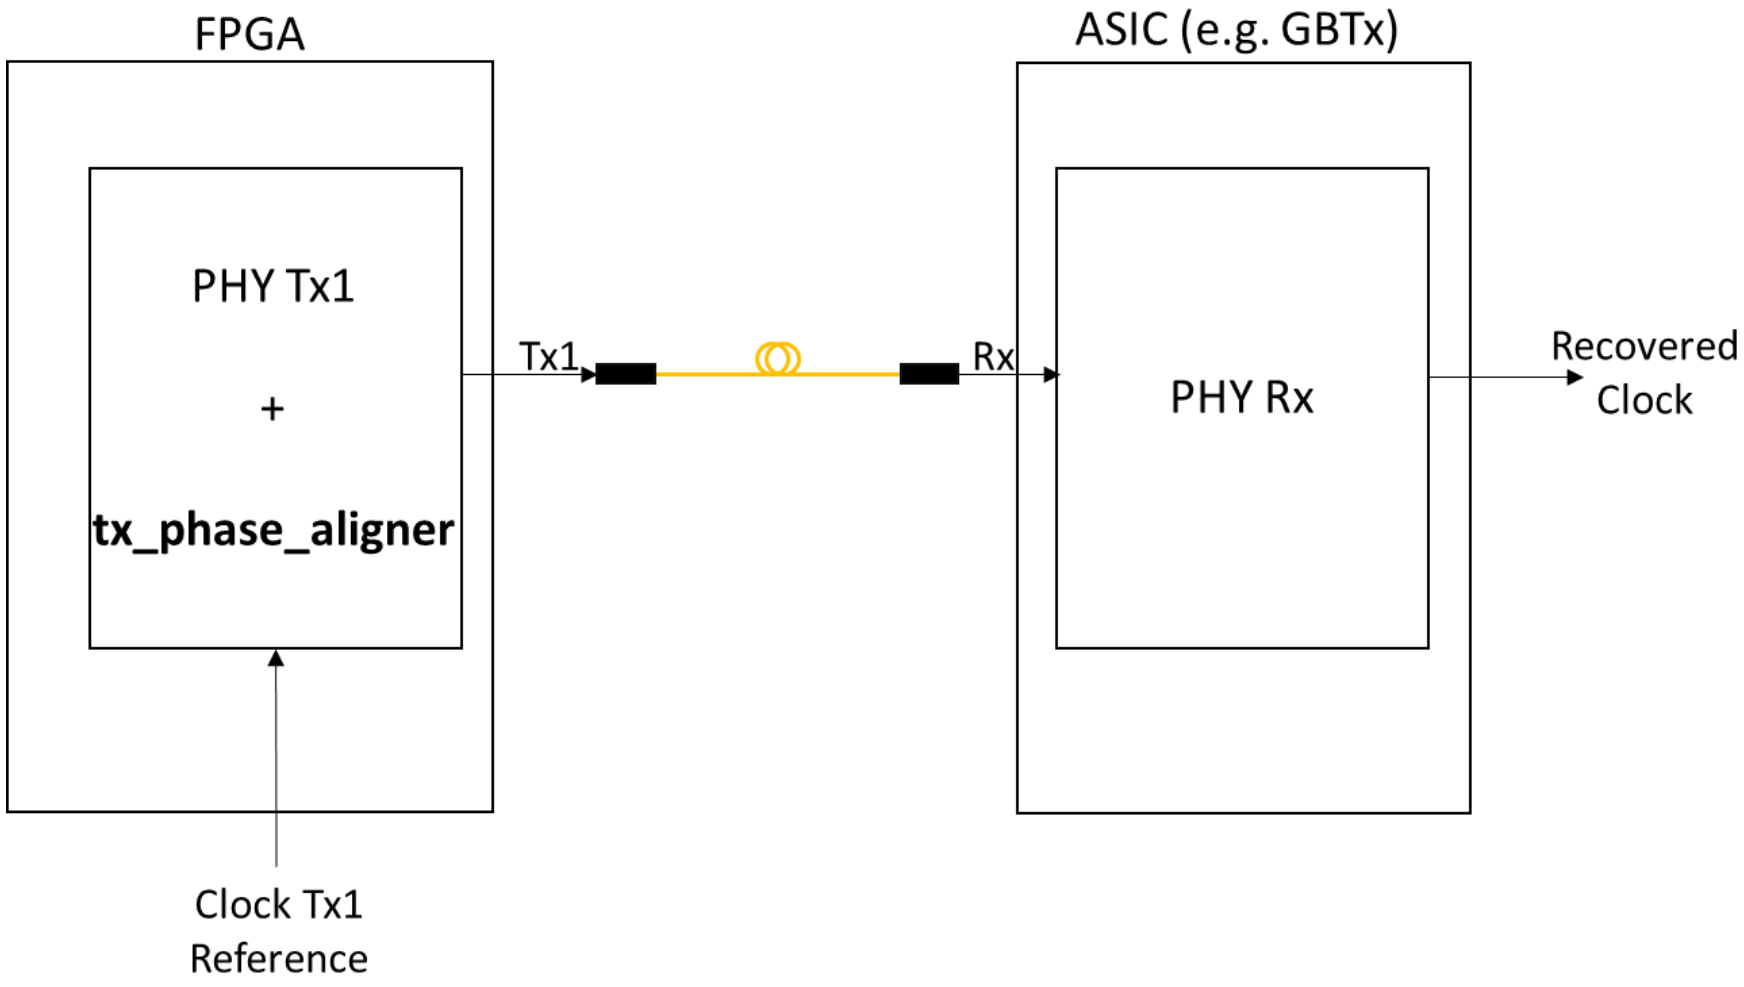

- How do we go to a cascaded link?

Outline

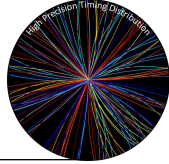

- Introduction
- Core architecture
  - Overview
  - Reference design
  - Characterization
- **Link cascade**
  - Rx recovered clock
  - Application of phase aligner
  - Characterization
- Conclusions

Link cascade: Rx recovered clock

- RXOUTCLK goes through fabric
  - clock which minimal latency variation (w.r.t. header) is achieved for GBT-FPGA, TTC-PON, …
- Ultrascale architecture → RXRECCLK
  - Challenge: UI jumps (w.r.t. header) for different resets (different dividers)

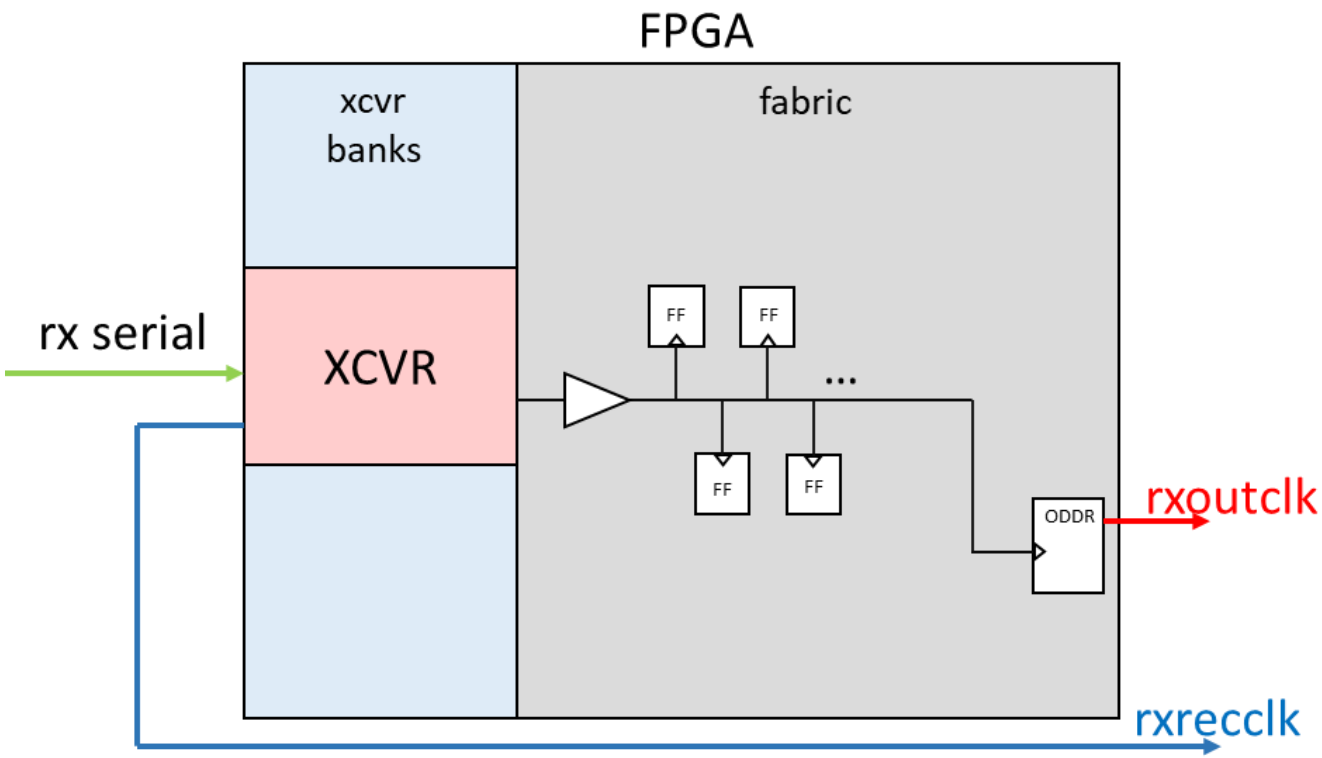

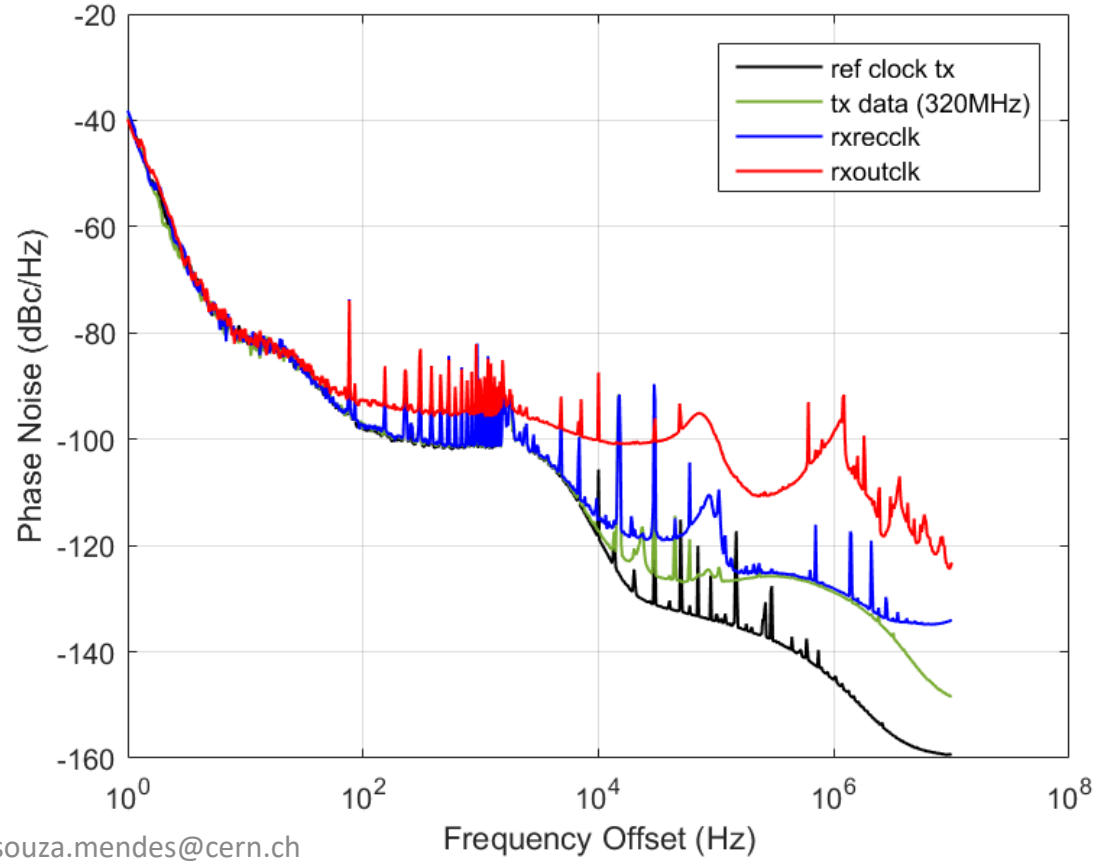

Link cascade: Rx recovered clock (RXUSRCLK2)

- Tx aligner core: **FINE_ALIGN** mode (i.e. minimal latency variation)
- Average: reset at every acquisition
- Observed sometimes ~10ps jumps on GTY-Ultrascale+ (Rx Delay Aligner?)

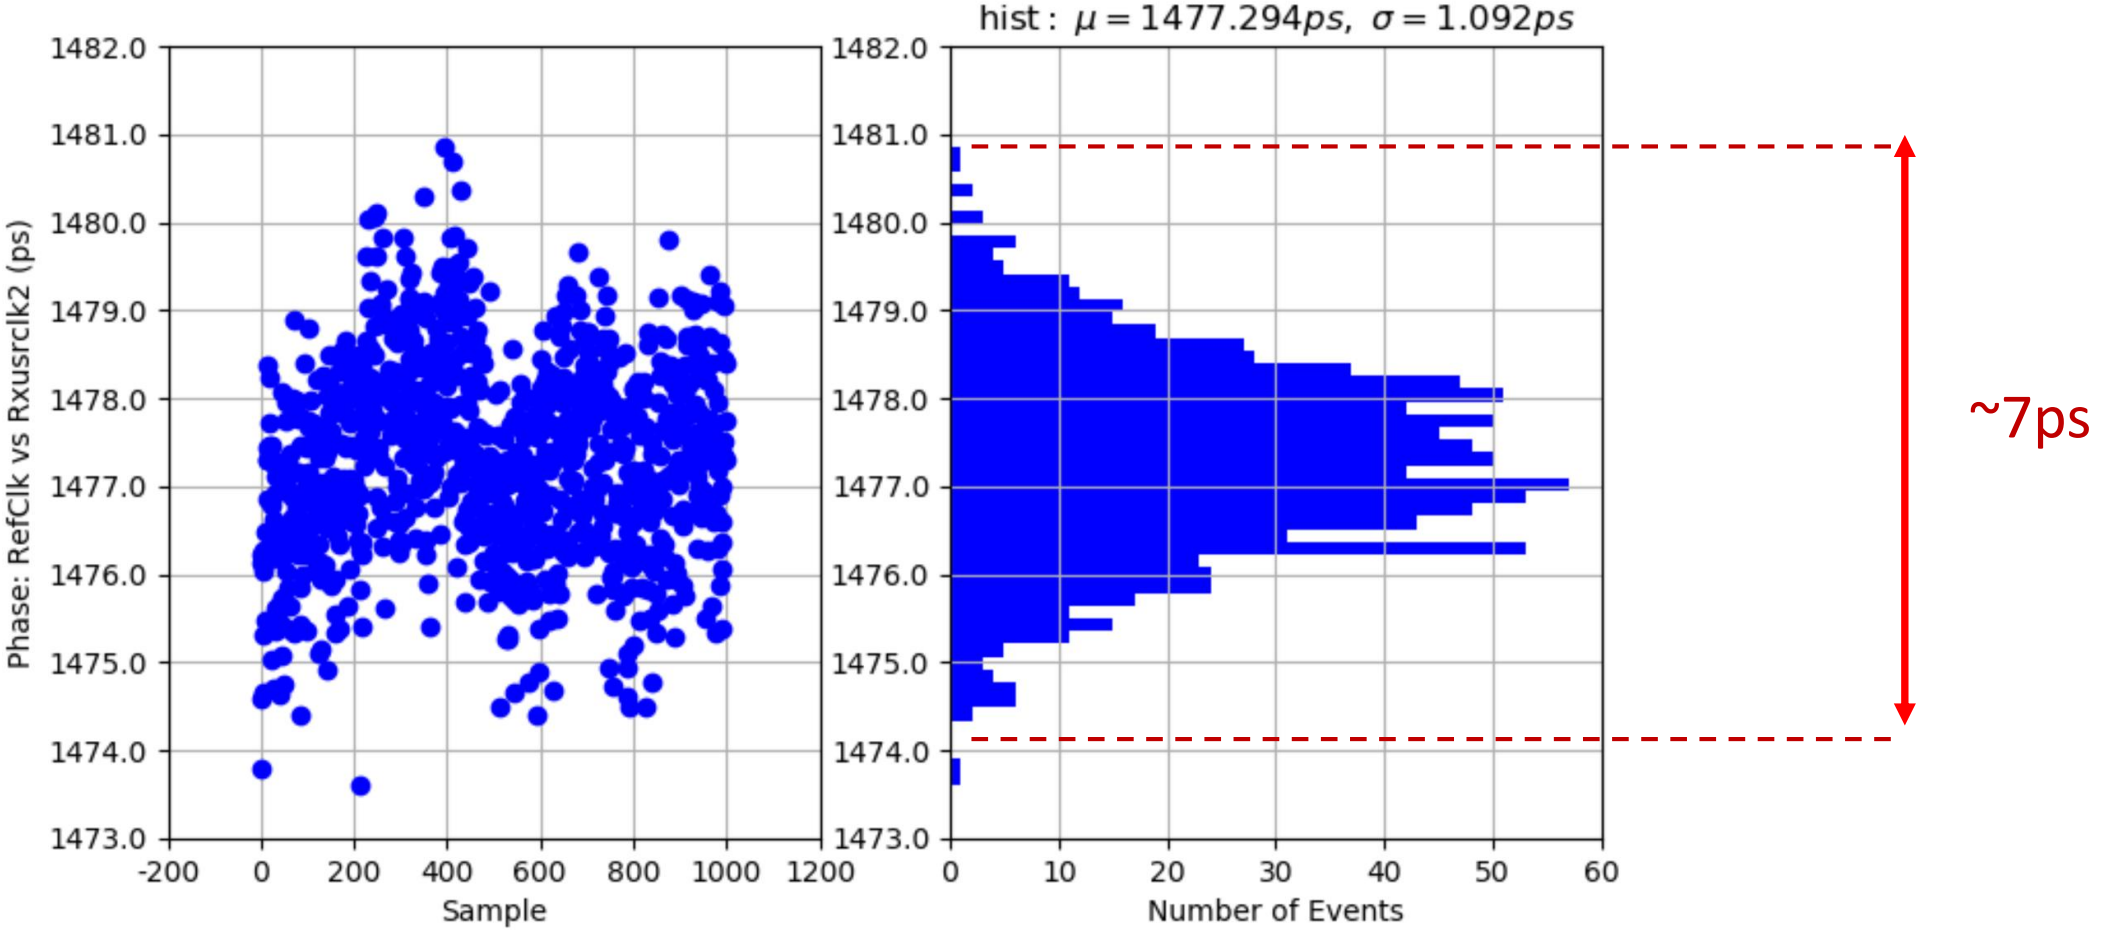

Link cascade: Rx recovered clock (RXRECCLK)

- Tx aligner core: **FINE_ALIGN** mode (i.e. minimal latency variation)
- Average: reset at every acquisition

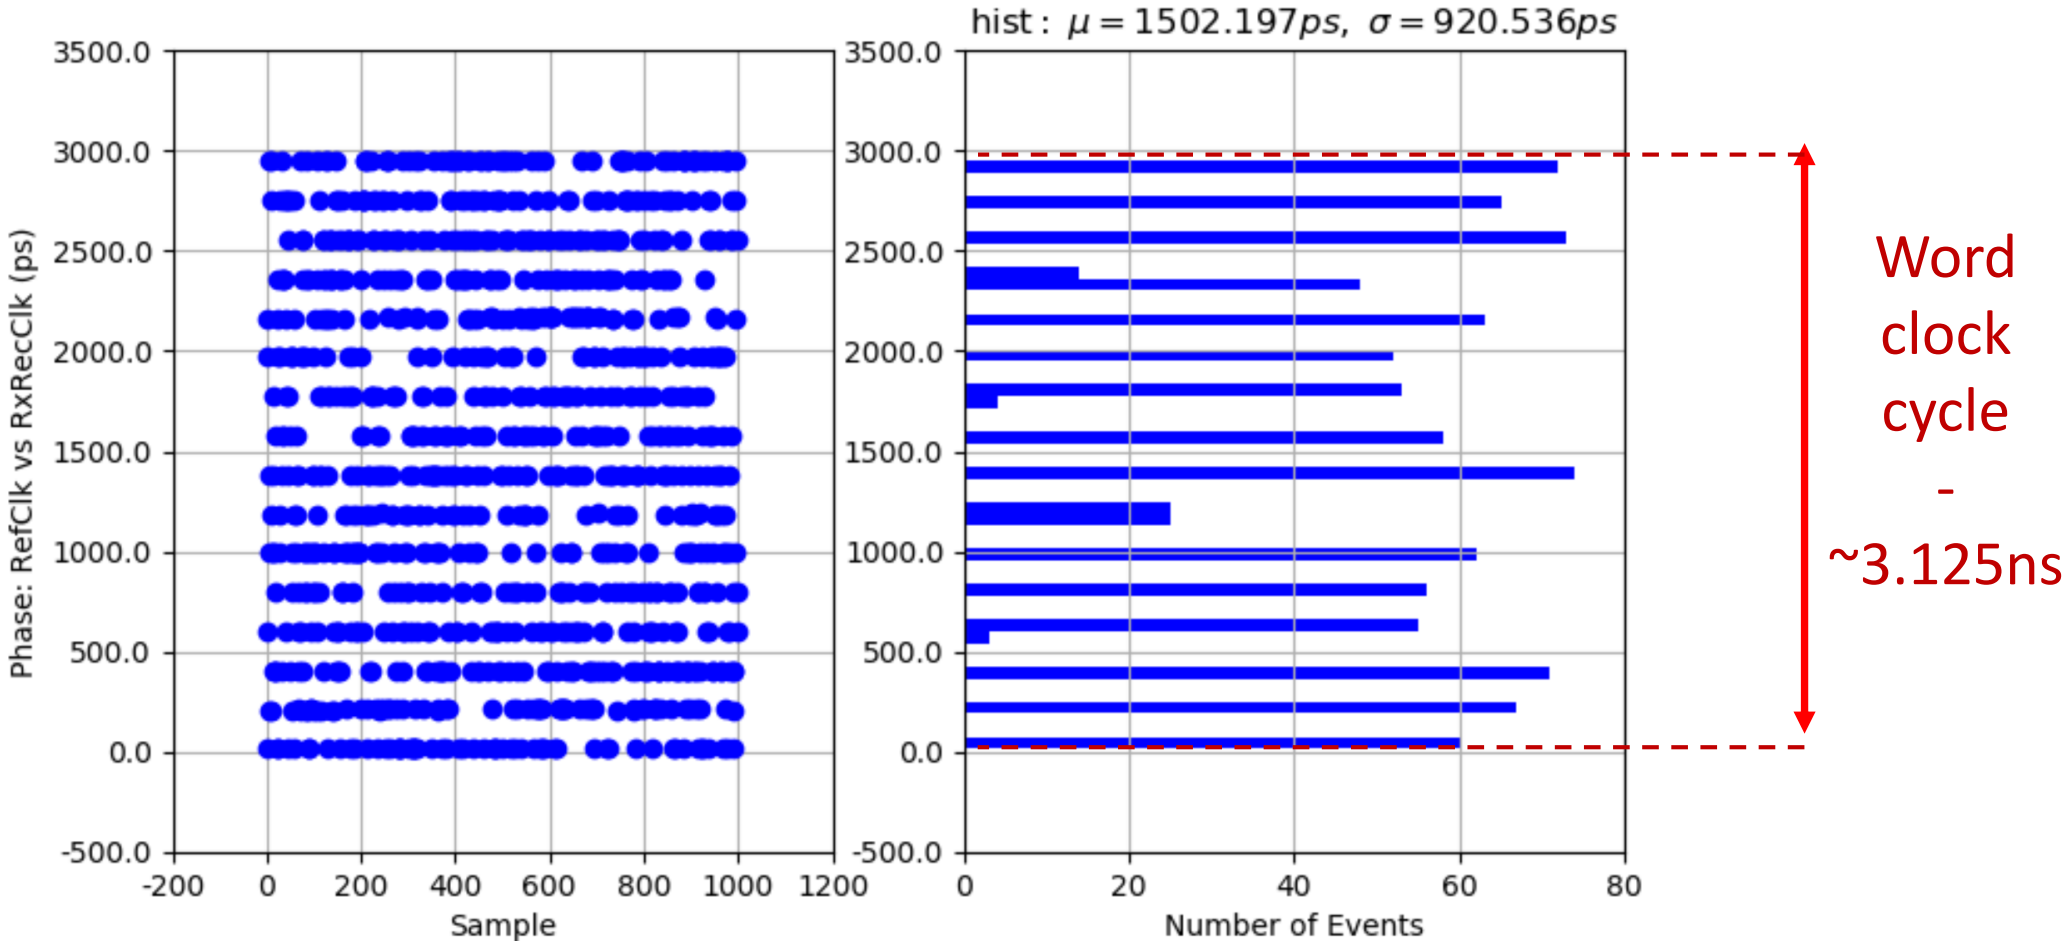

Link cascade: Rx recovered clock (RXRECCLK)

- Tx aligner core: **FINE_ALIGN** mode (i.e. minimal latency variation)
- Average: reset at every acquisition

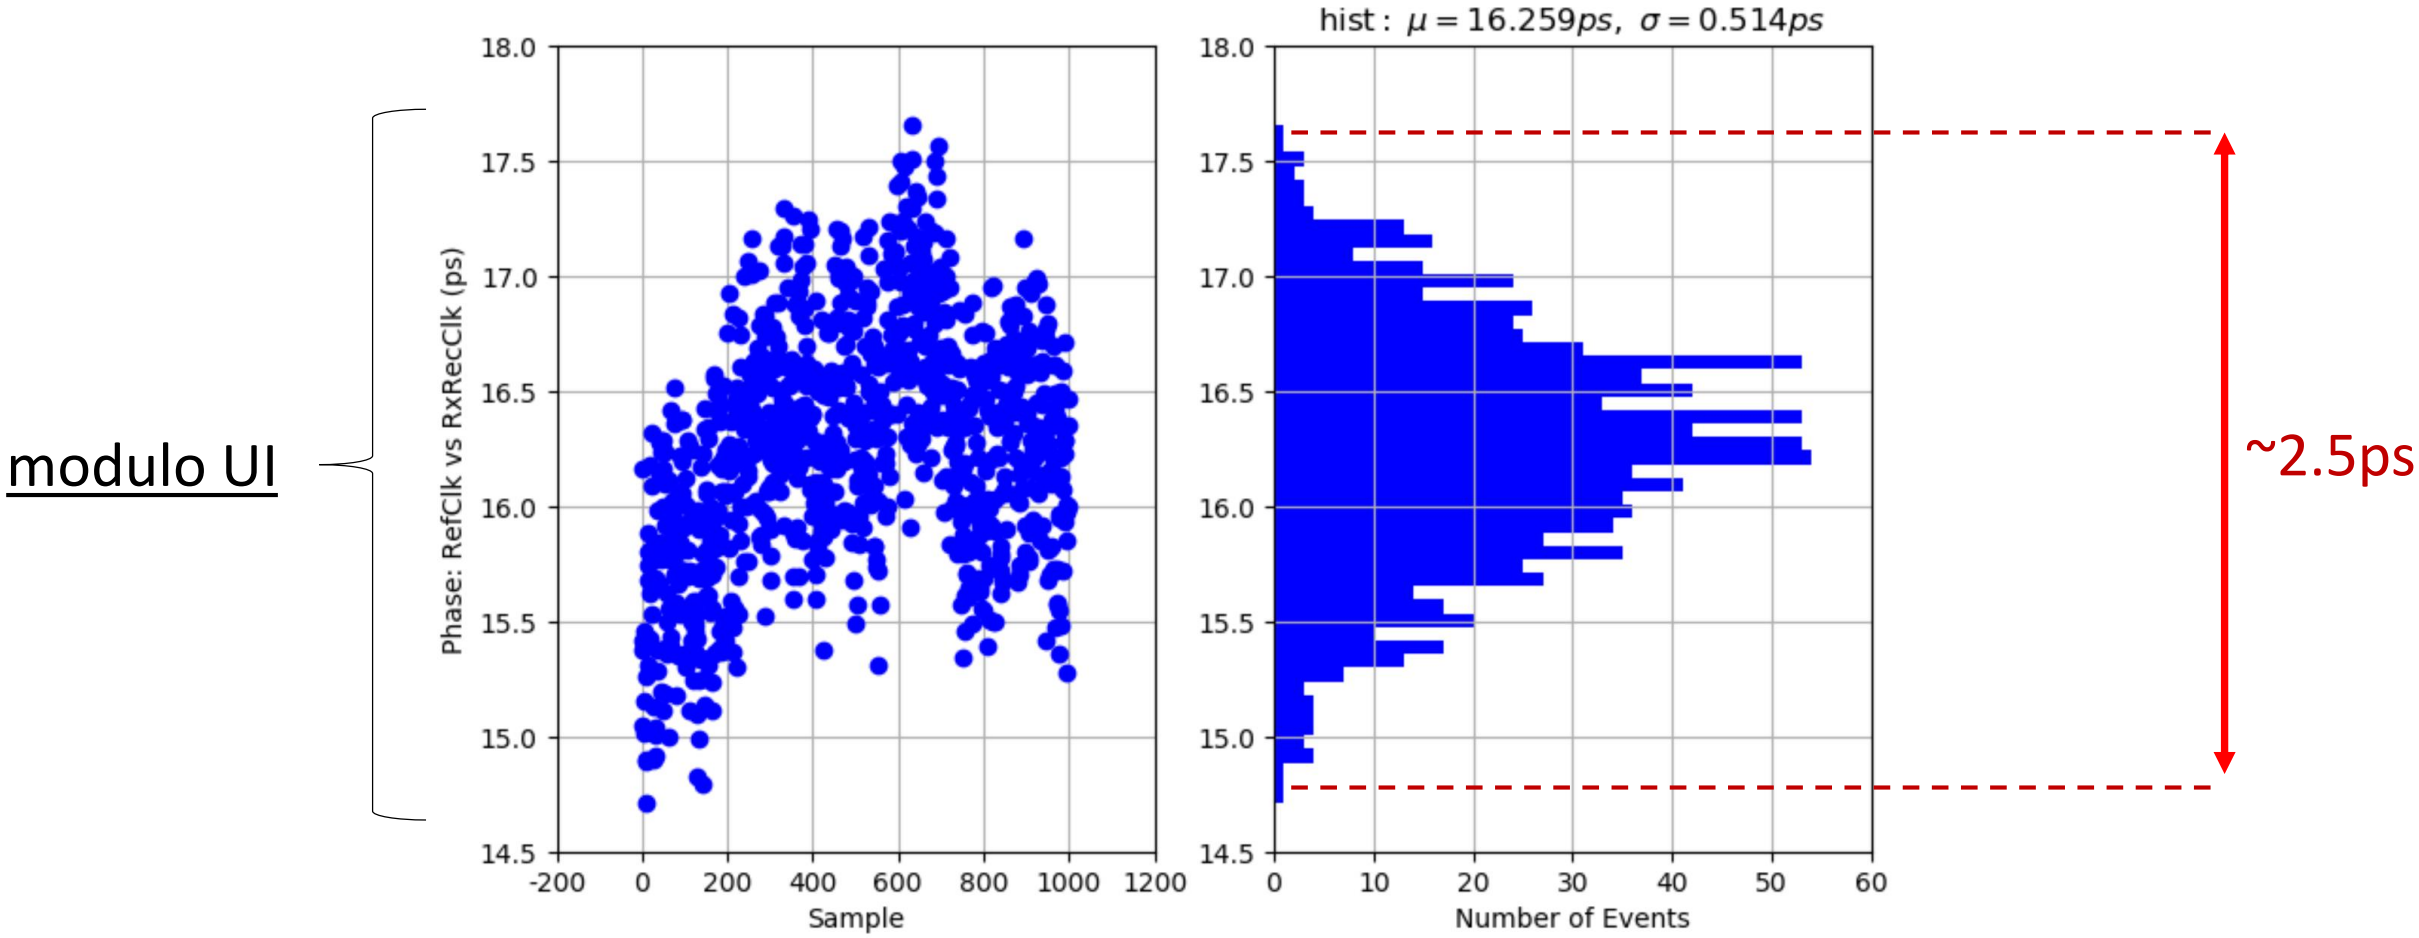

Link cascade: **«traditional»**

- External PLL shift for proper clock domain crossing (CDC)

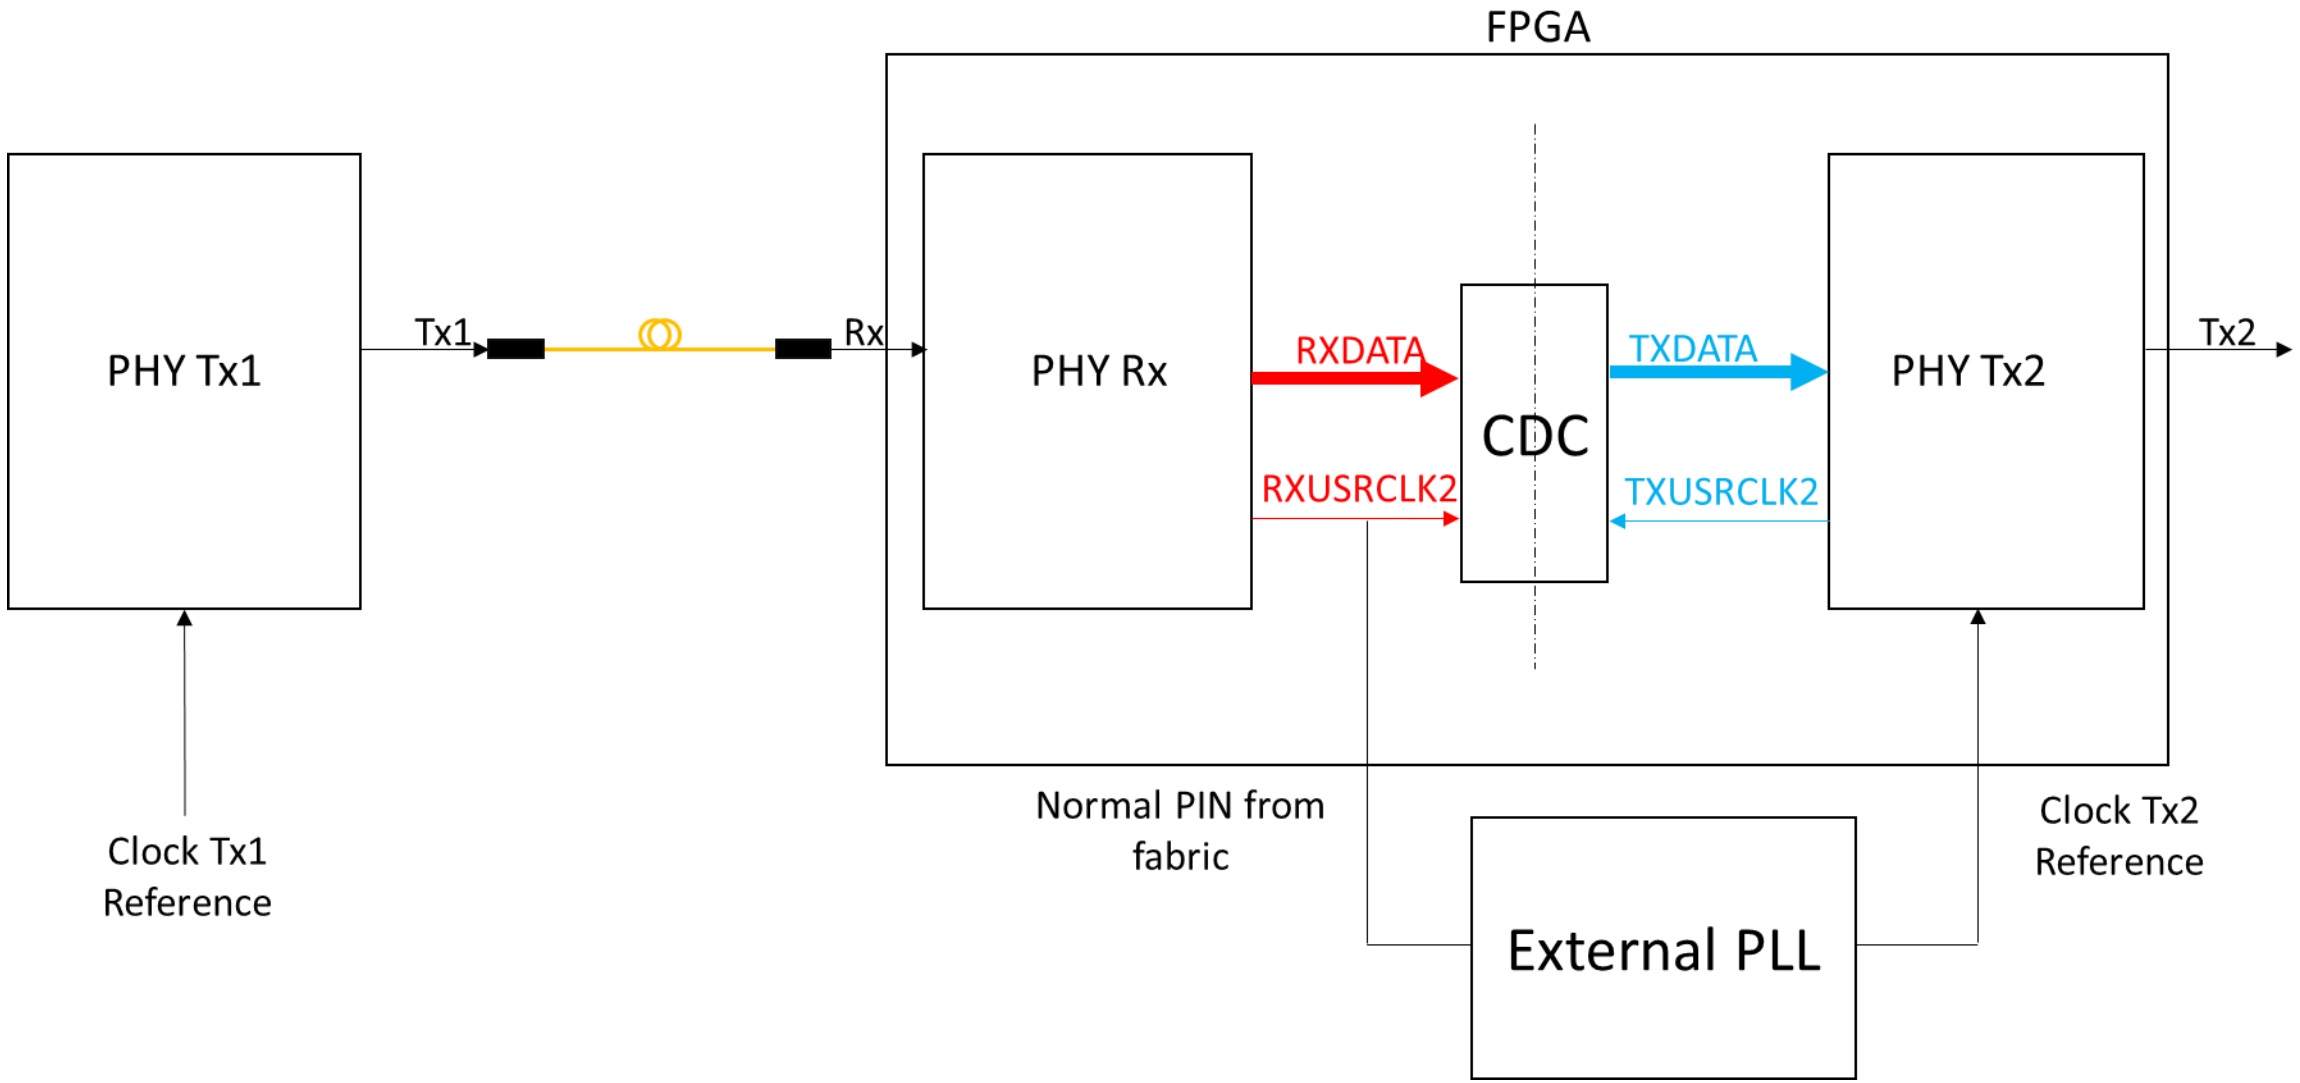

Link cascade: **tx_phase_aligner (local)**

- CDC managed by phase aligner core inside PHY

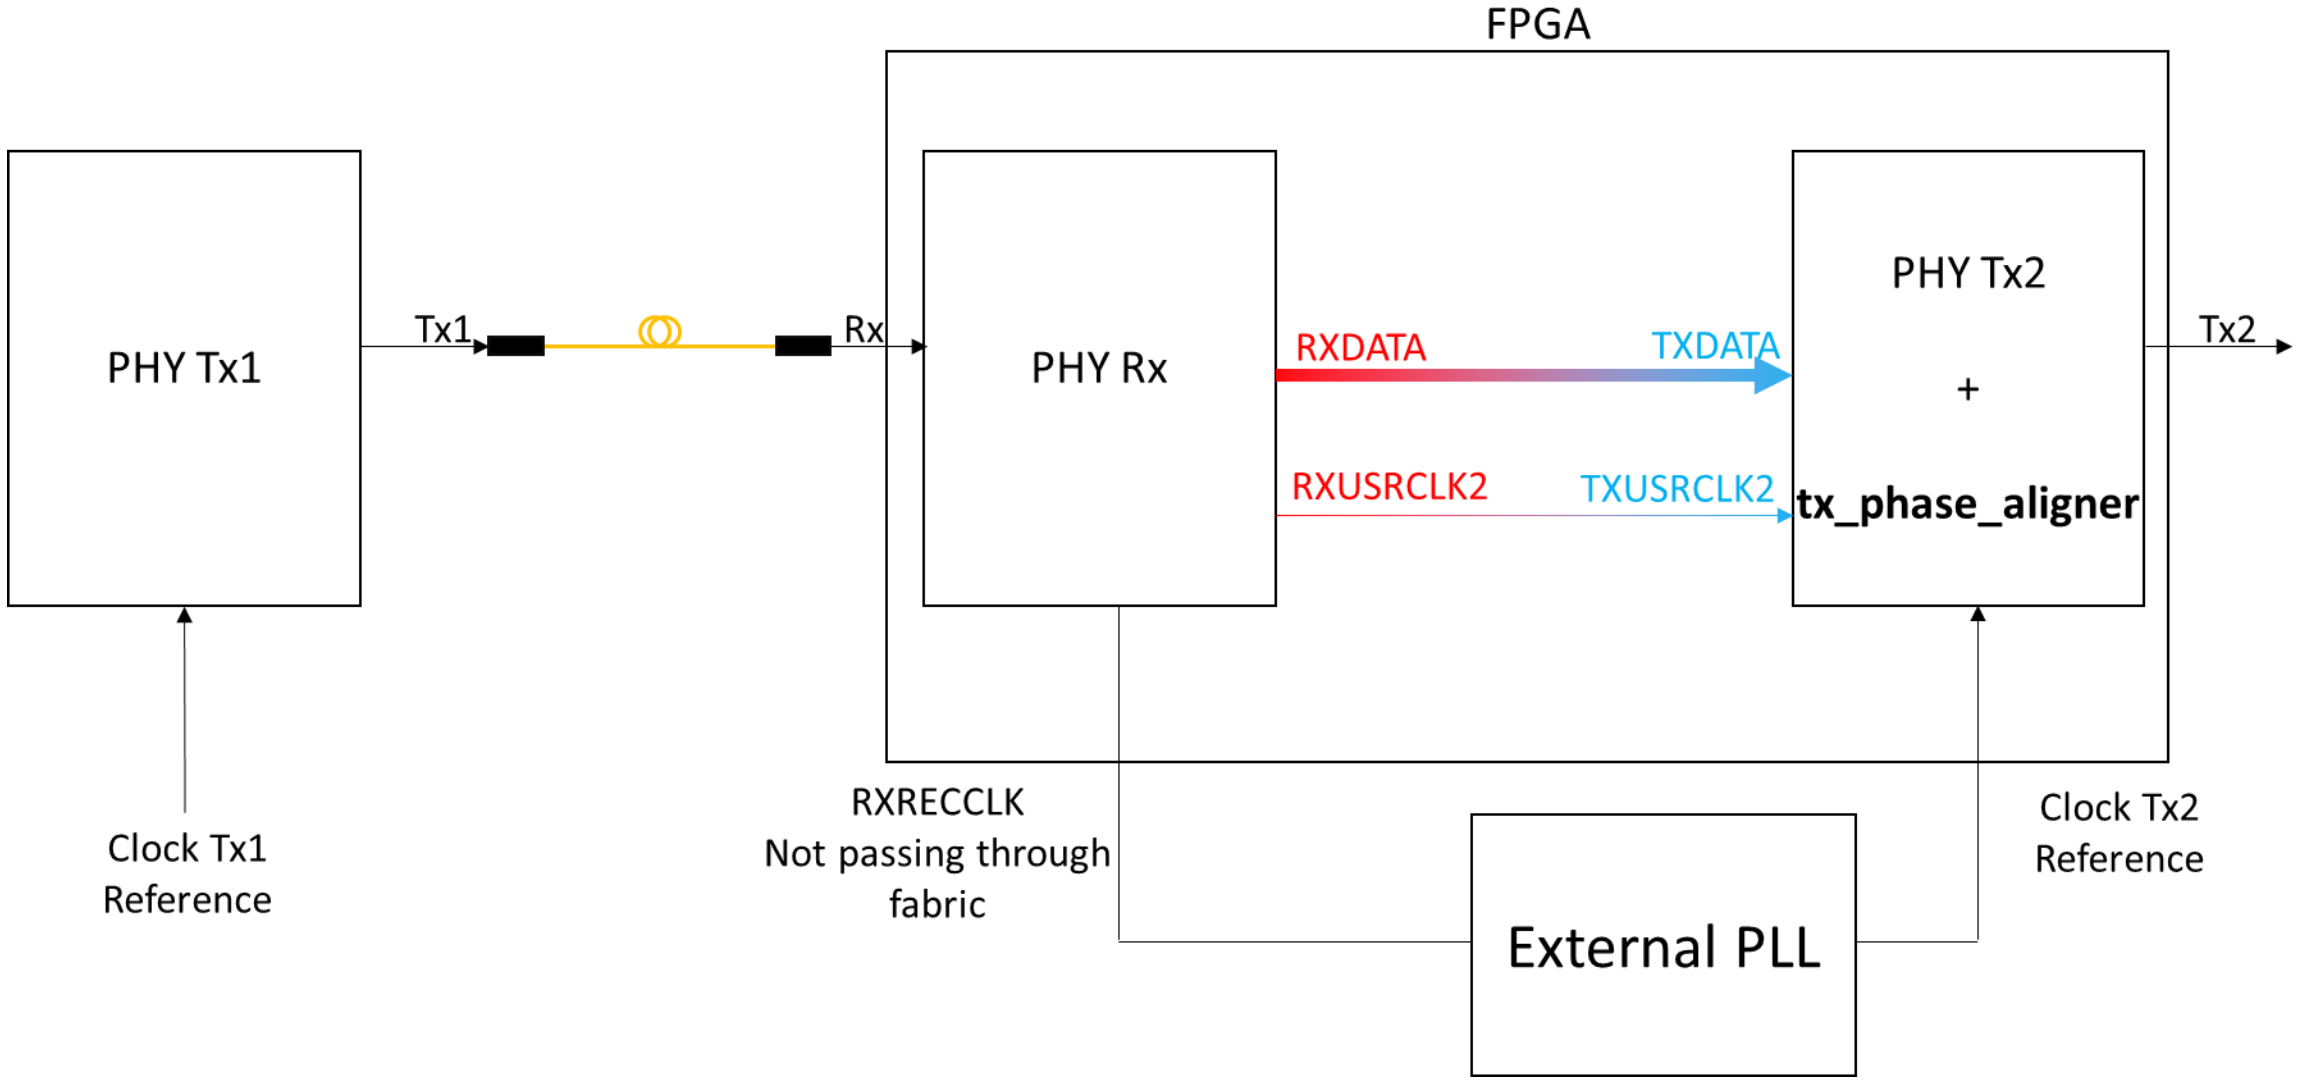

Link cascade: **tx_phase_aligner (system?)**

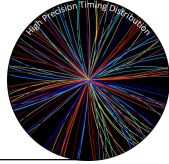

- CDC managed by phase aligner core inside PHY (and high-phase determinism)

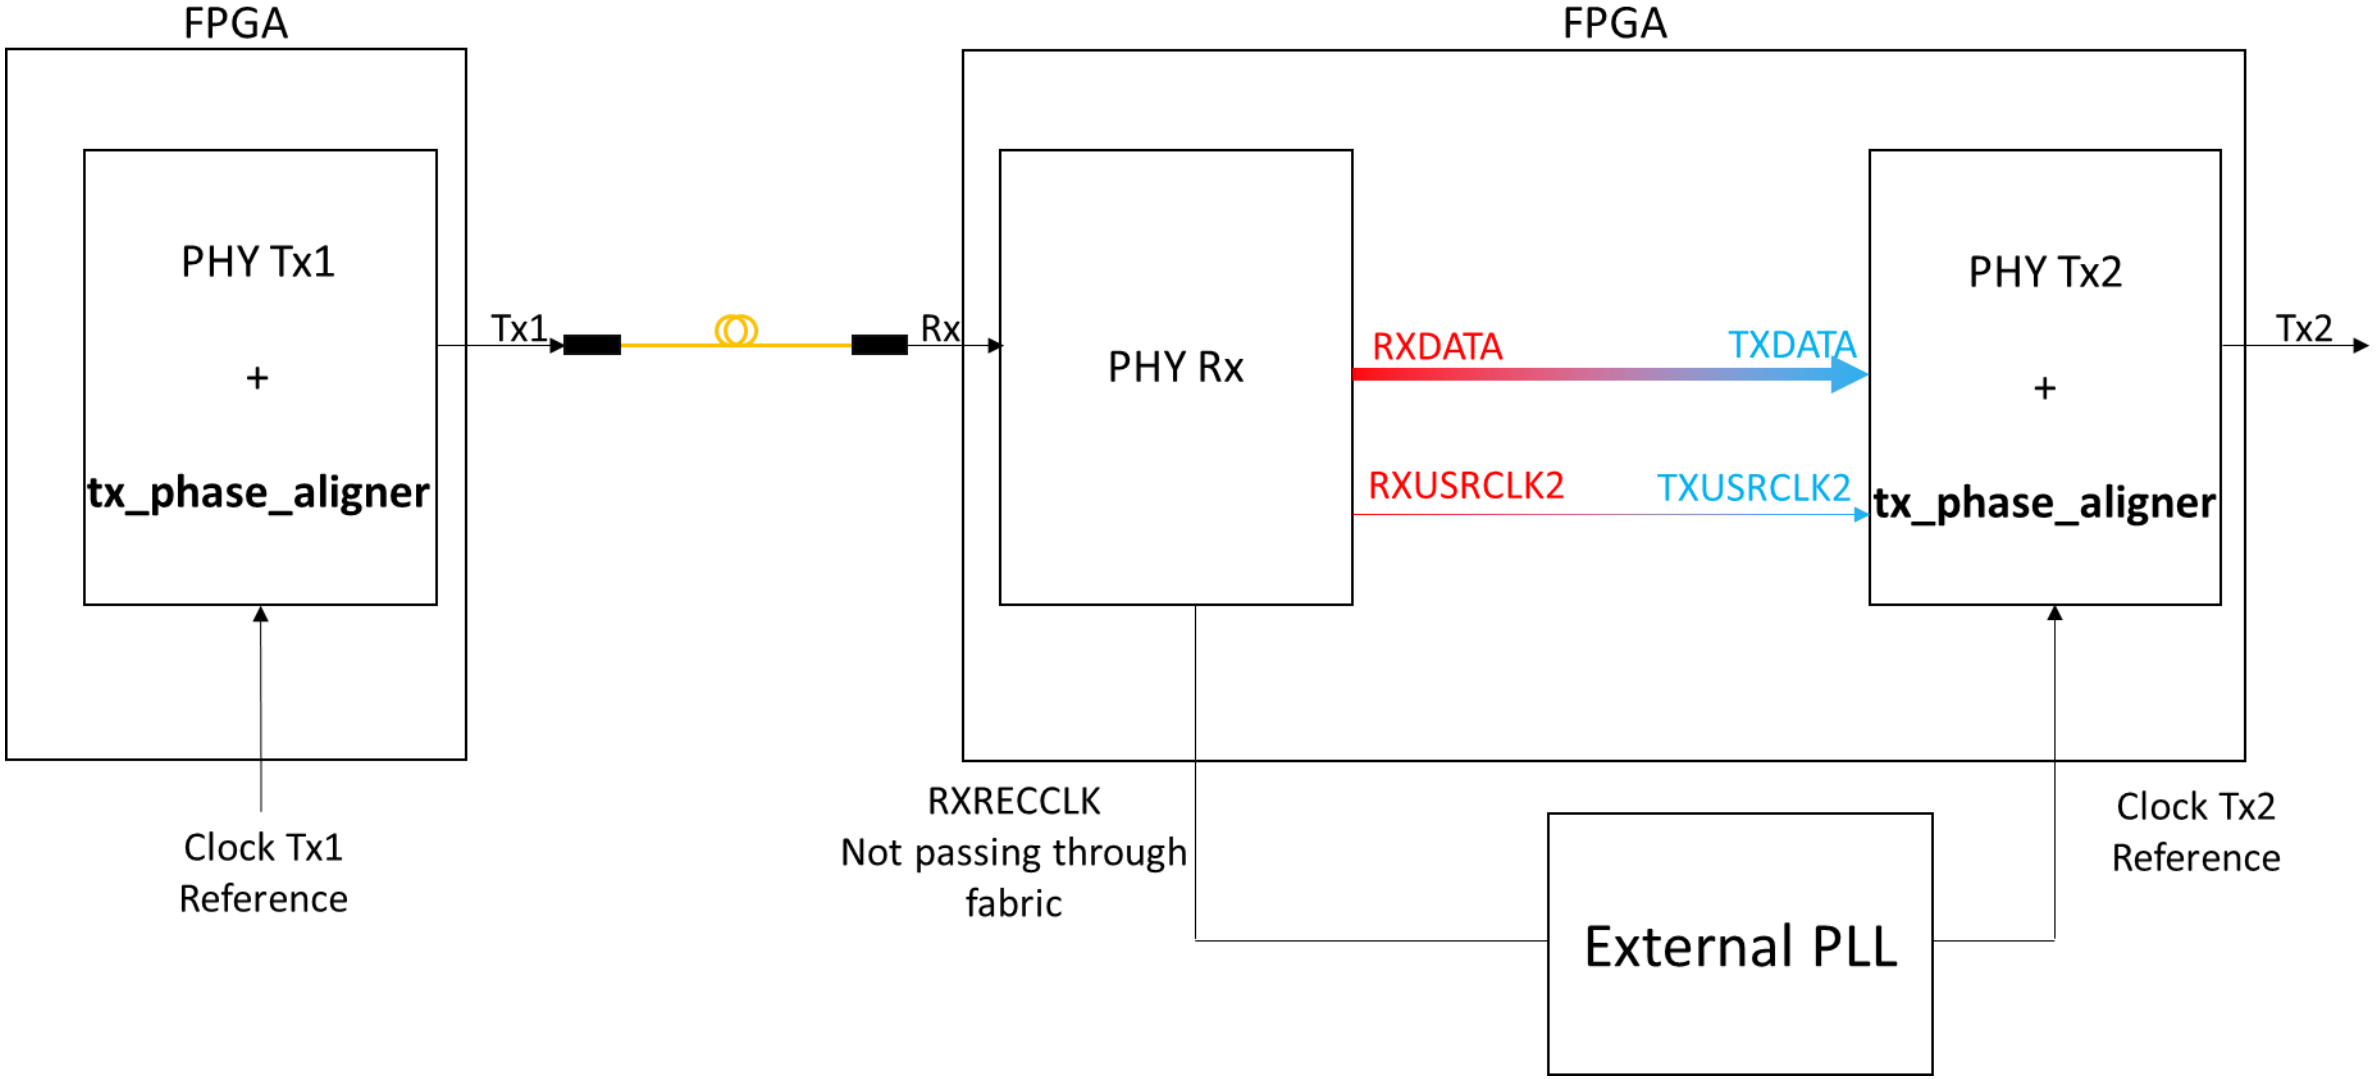

Link cascade: summary

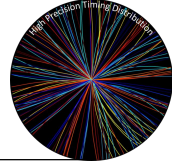

“traditional”

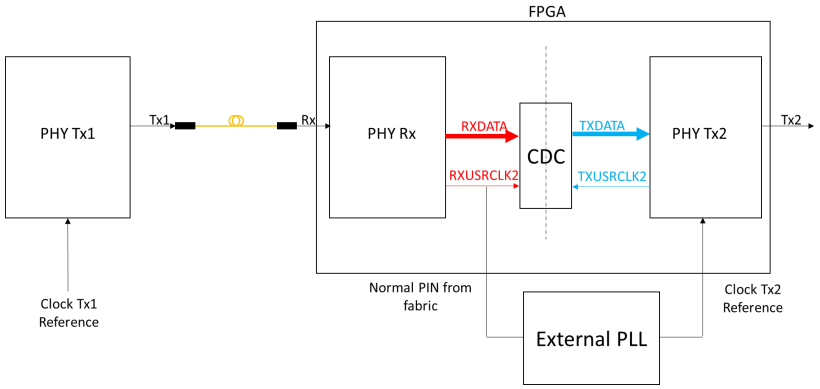

- Clock goes through fabric
- CDC managed by user
- Bigger phase-variations
- Easier Clock to treat

local

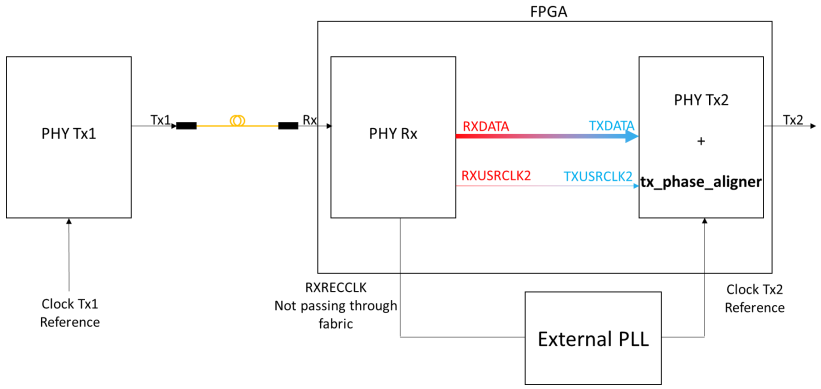

- Clock does not go through fabric
- CDC very simplified
- Phase-determinism depend on who is transmitting to cascaded core
- RXRECCLK will still not be fixed phase

system

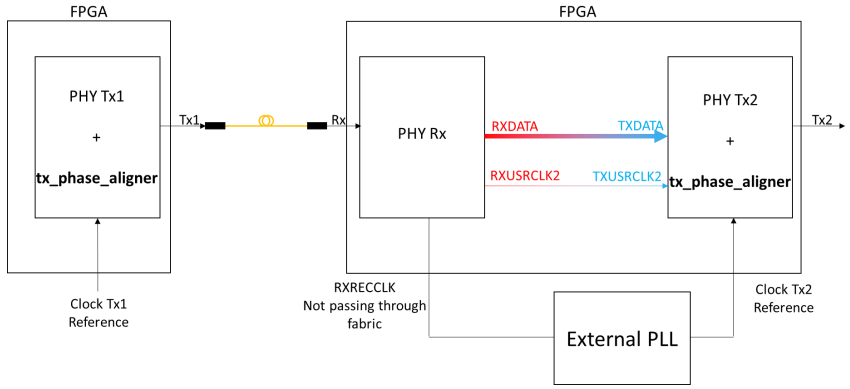

- Clock does not go through fabric
- CDC very simplified
- Higher phase determinism
- RXRECCLK will still not be fixed latency

Characterization: Cascaded setup

- Cascaded test-setup

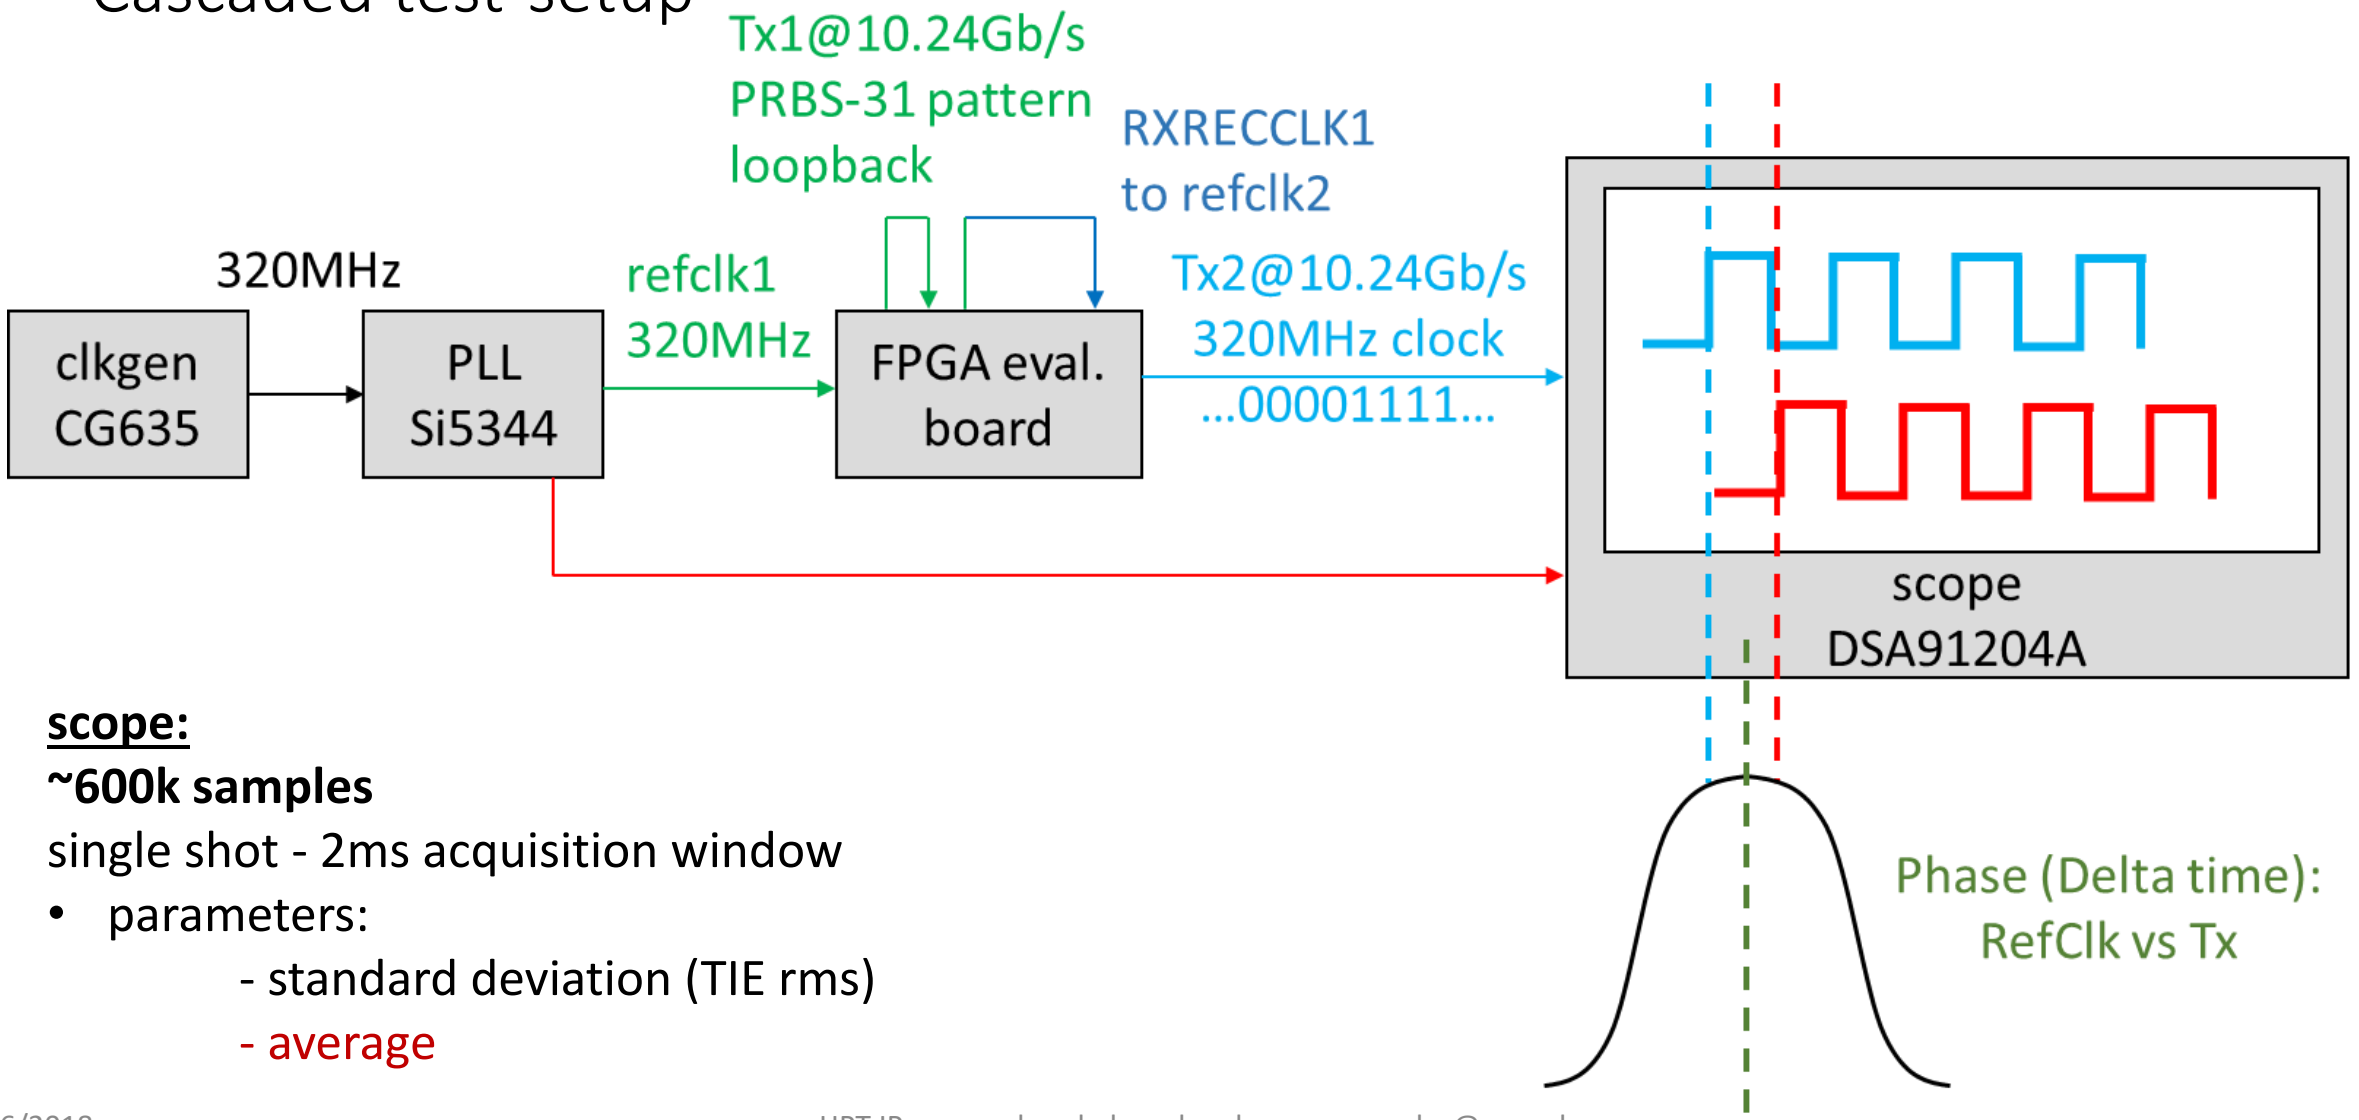

scope:

~600k samples

single shot - 2ms acquisition window

- parameters:
  - standard deviation (TIE rms)
  - average

Characterization: Tx2

- Tx aligner core: **UI_ALIGN** mode (i.e. deterministic latency)
- Average: reset at every acquisition

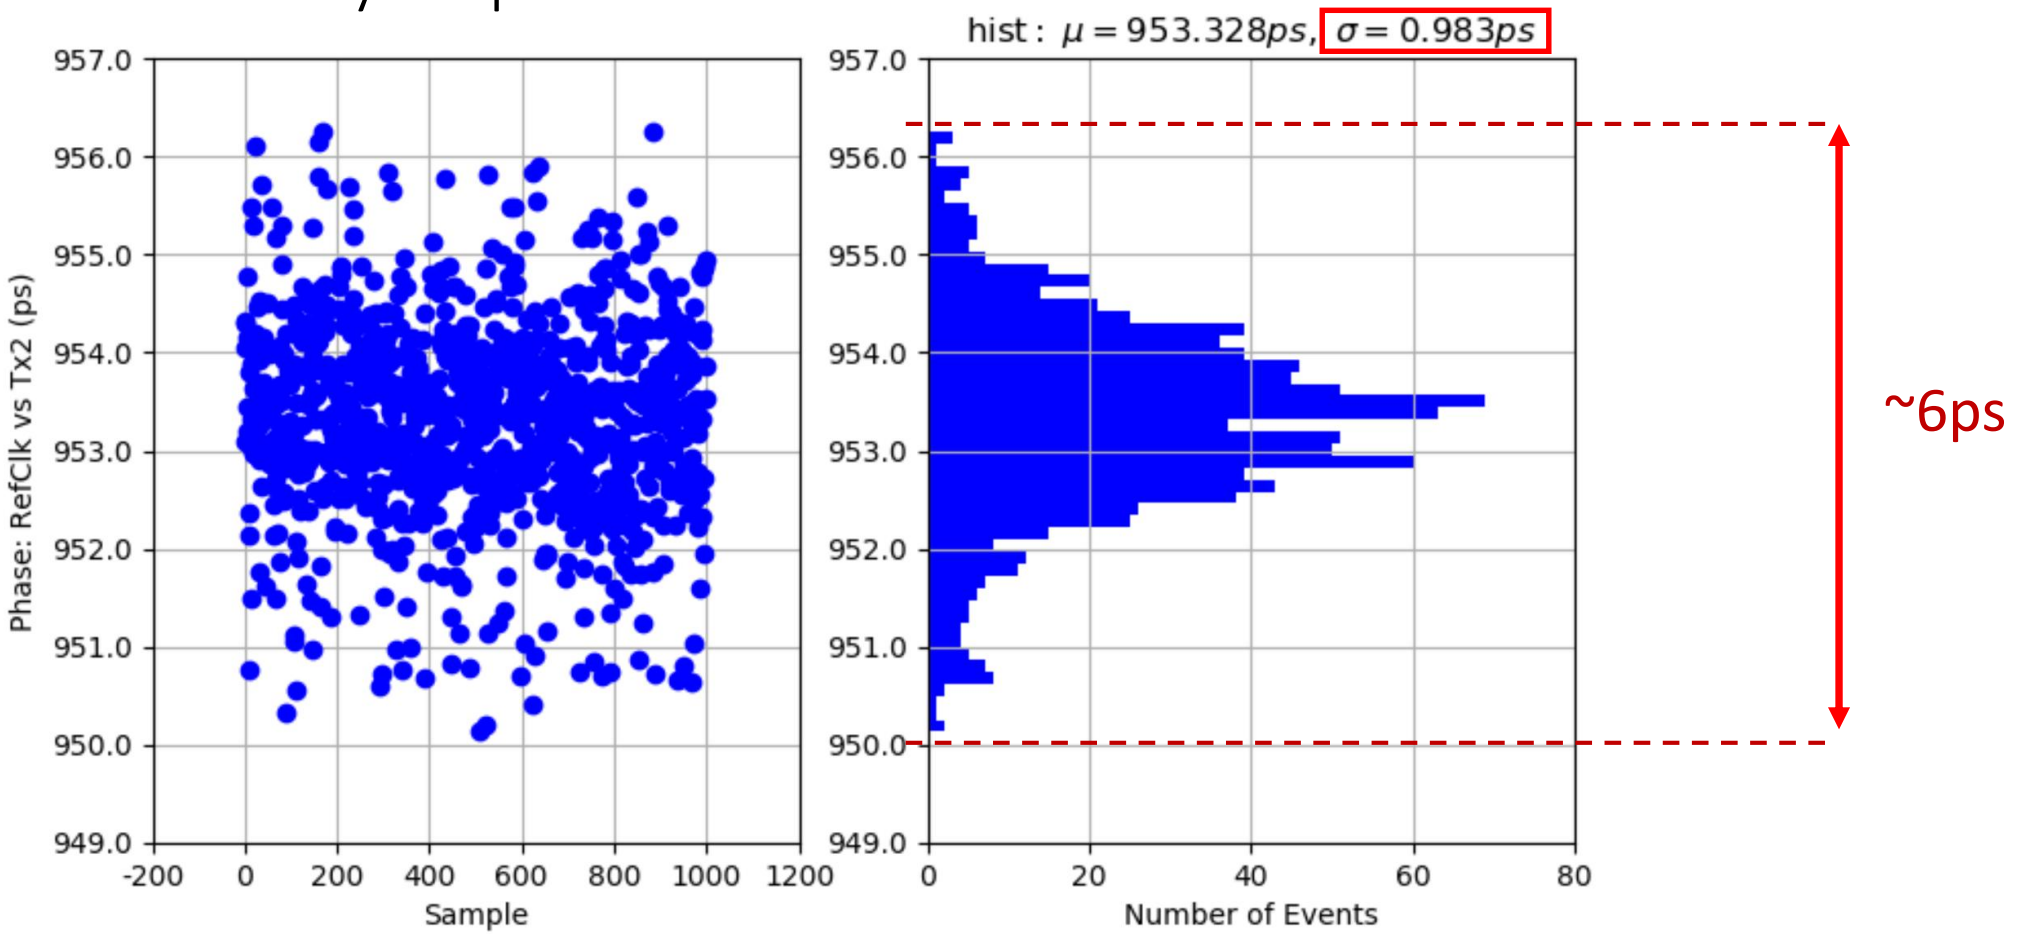

Characterization: Tx2 temperature

- Tx aligner core: **UI_ALIGN** mode (i.e. deterministic latency)

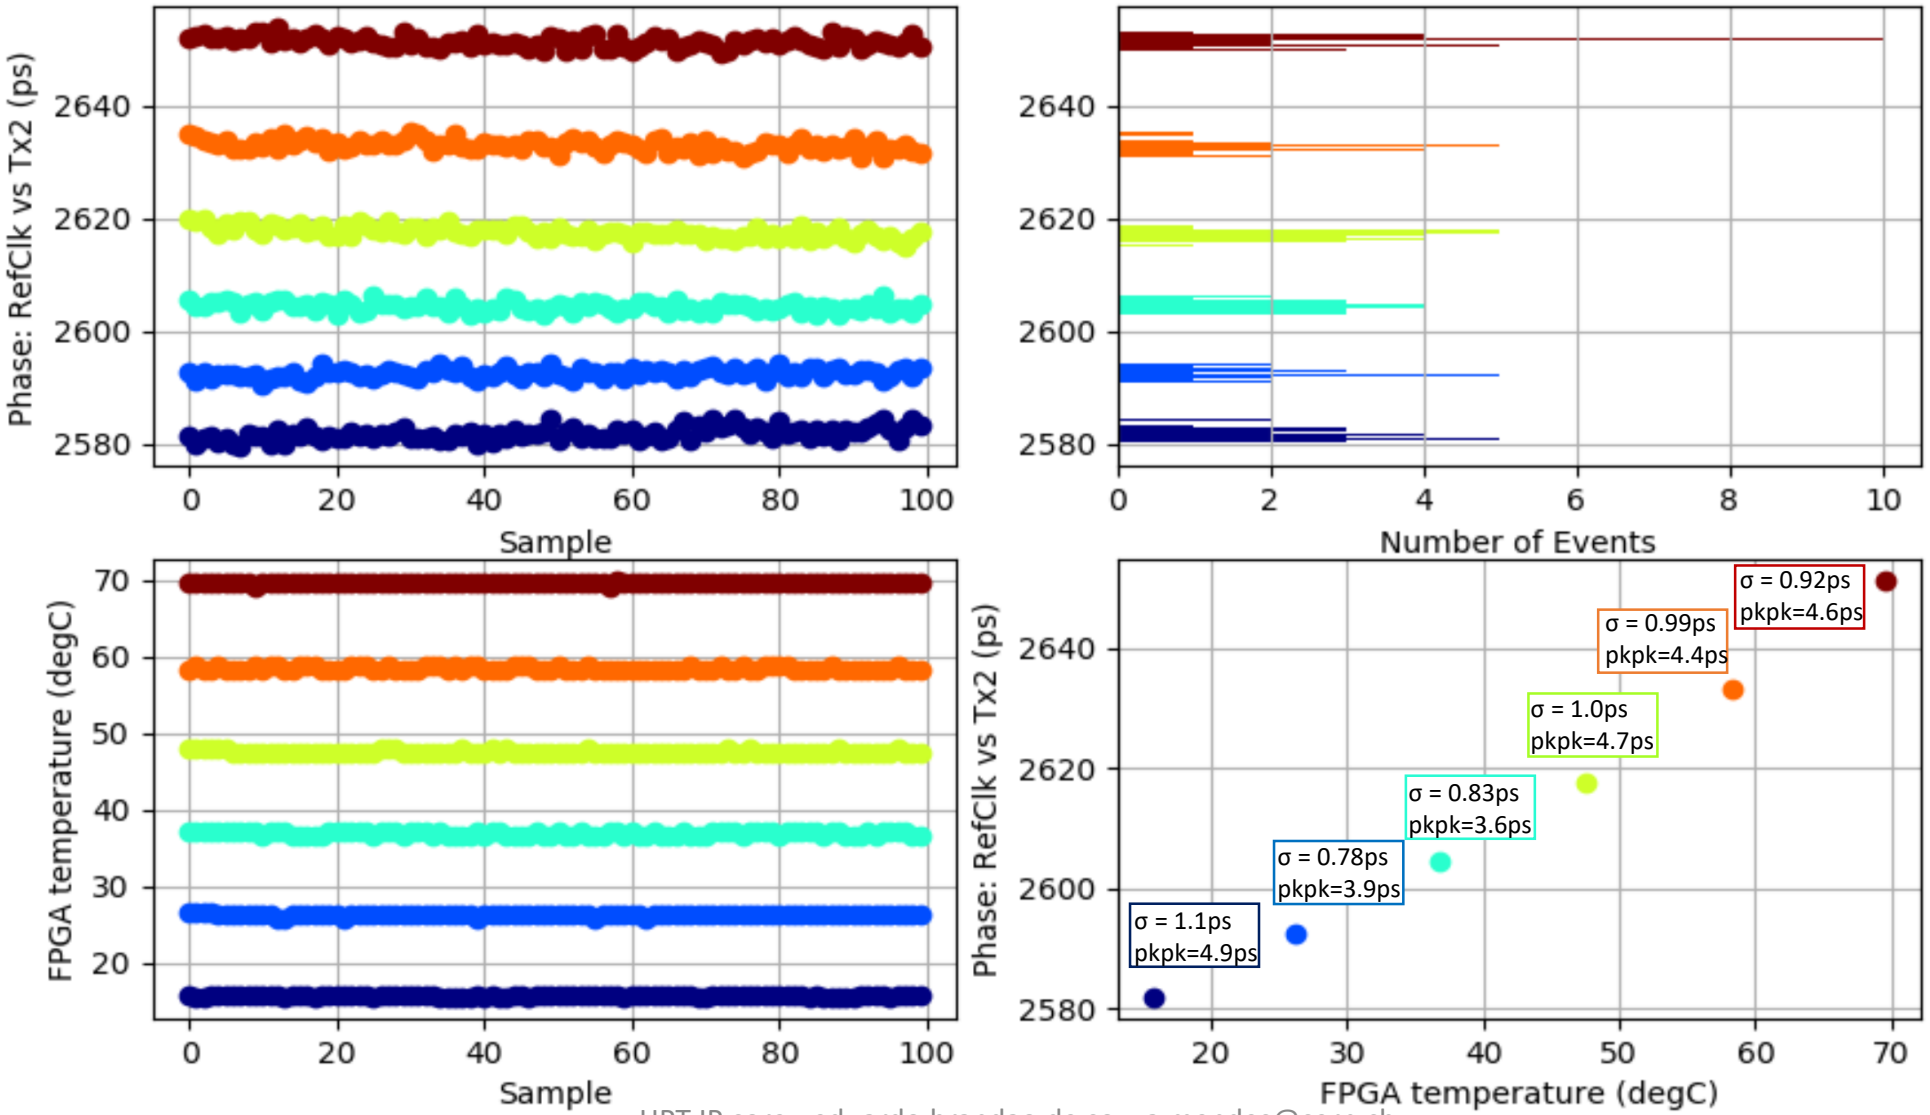

Conclusions

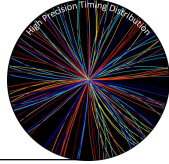

- Implemented a new technique for deterministic phase in Xilinx FPGAs
  - Technique tested with KCU105 (GTH Ultrascale), ZCU102 (GTH Ultrascale+), KCU116 (GTY Ultrascale+)
  - Temperature tests on KCU105
- A very light core was developed in VHDL and is available for interested users
  - https://gitlab.cern.ch/HPTD/tx_phase_aligner (access under request)
  - Compatible with common FPGA cores (TTC-PON, GBT-FPGA, LpGBT-FPGA)
  - If you want to learn more, contact us
- This core can also be useful for cascaded links and for skew tuning

SPARE SLIDES

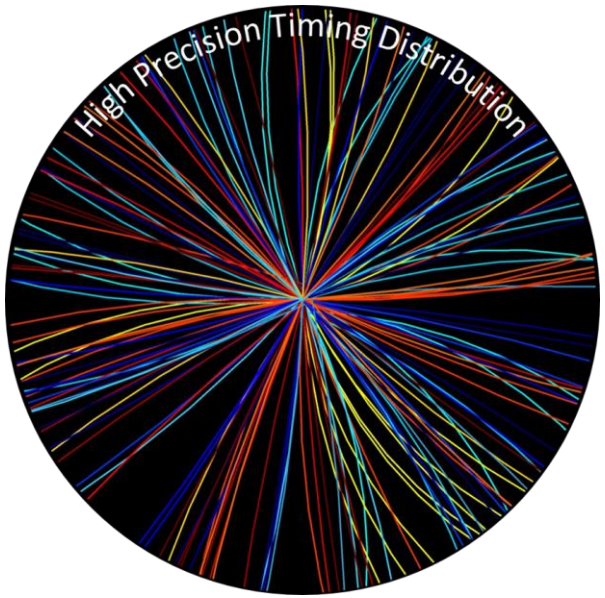

Characterization: Cascaded temperature

- Tx aligner core: **UI_ALIGN** mode (i.e. deterministic latency)

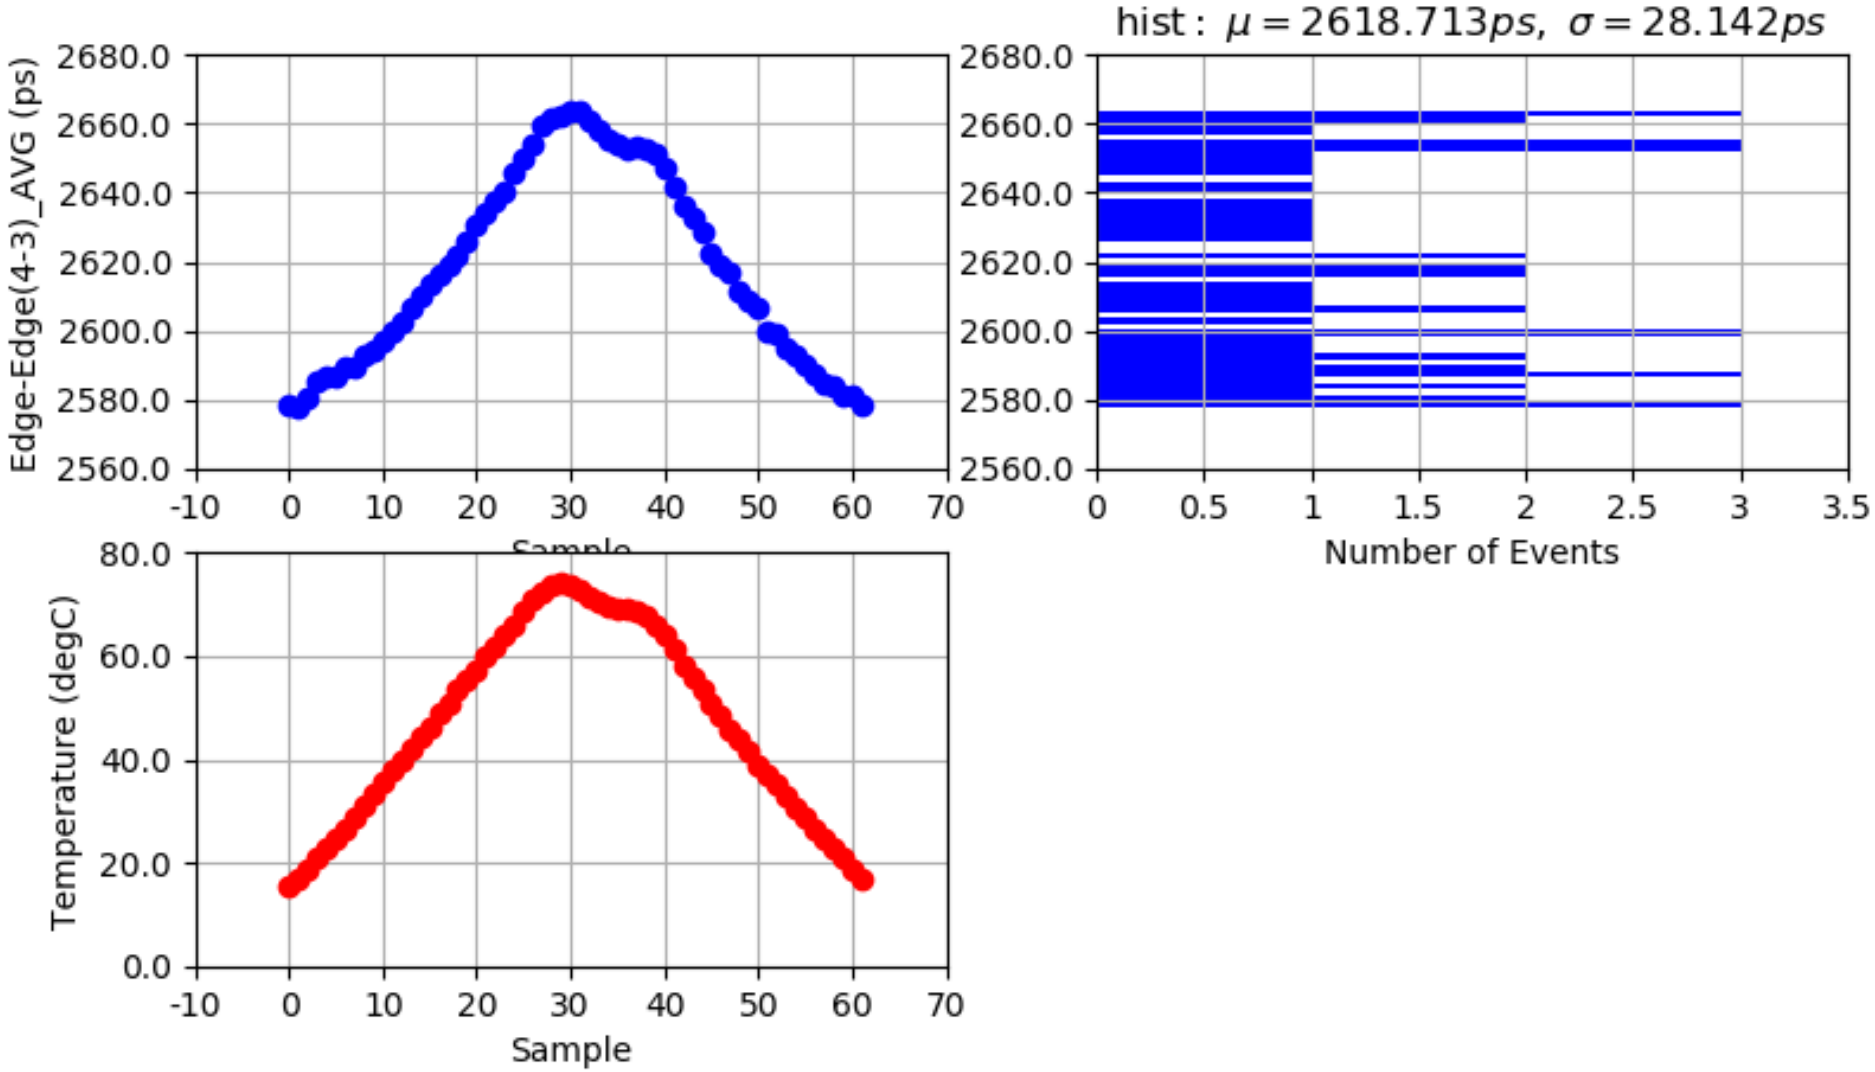

Characterization: Tx only temperature

- Tx aligner core: FINE_**ALIGN** mode (i.e. minimal latency variation)

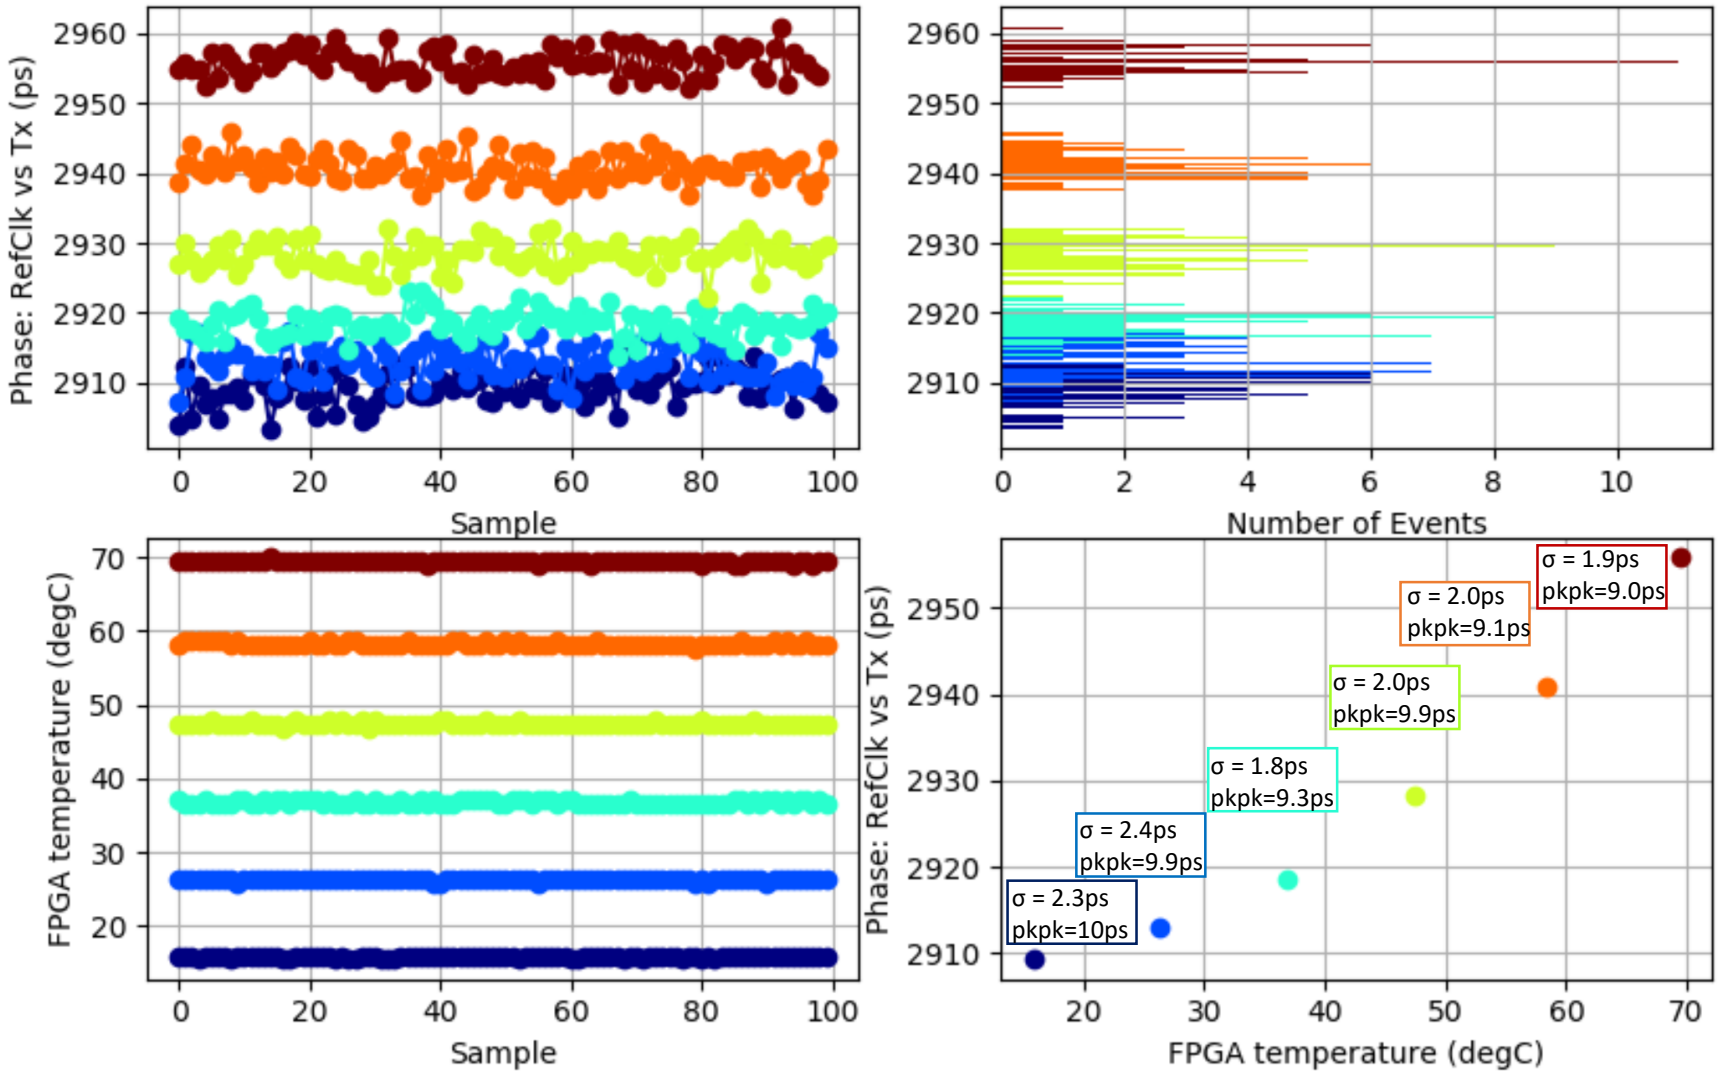

Characterization: Tx only temperature

- Tx aligner core: **UI_ALIGN** mode (i.e. deterministic latency)

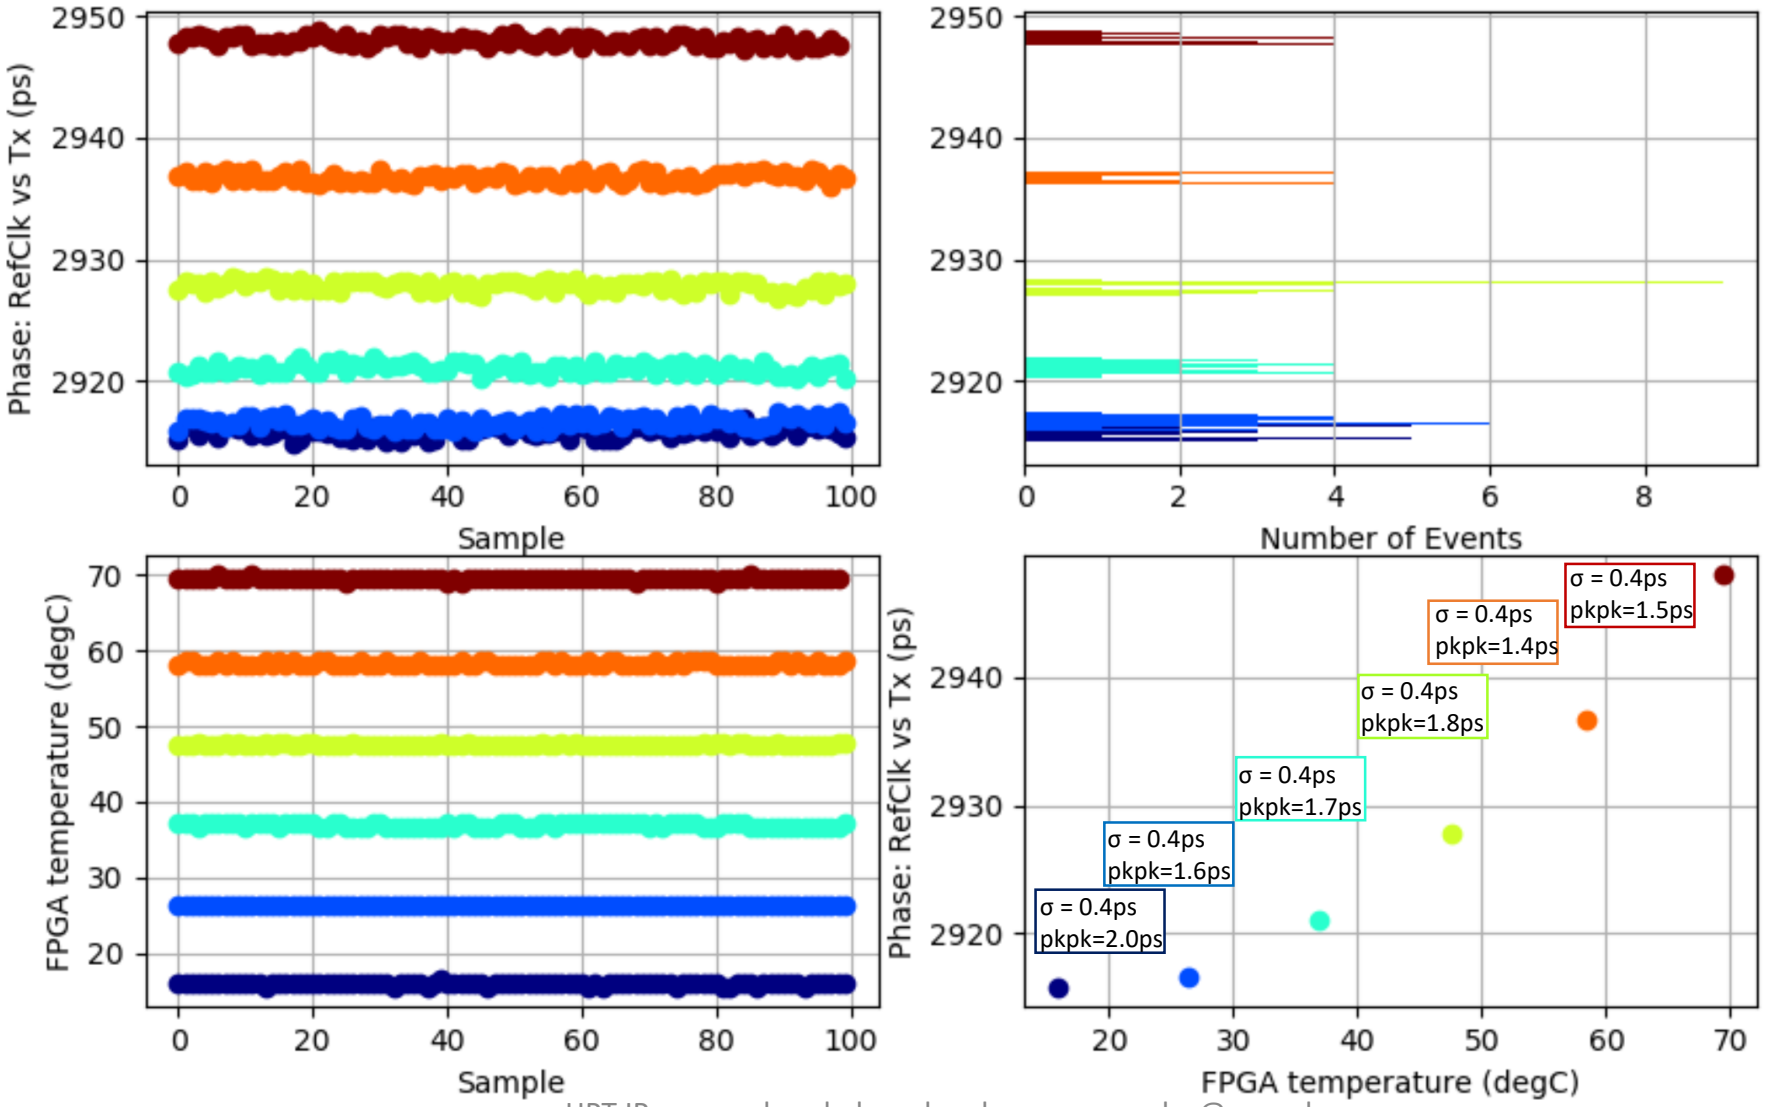

Characterization: Tx only temperature

- Tx aligner core: **UI_ALIGN** mode (i.e. deterministic latency)

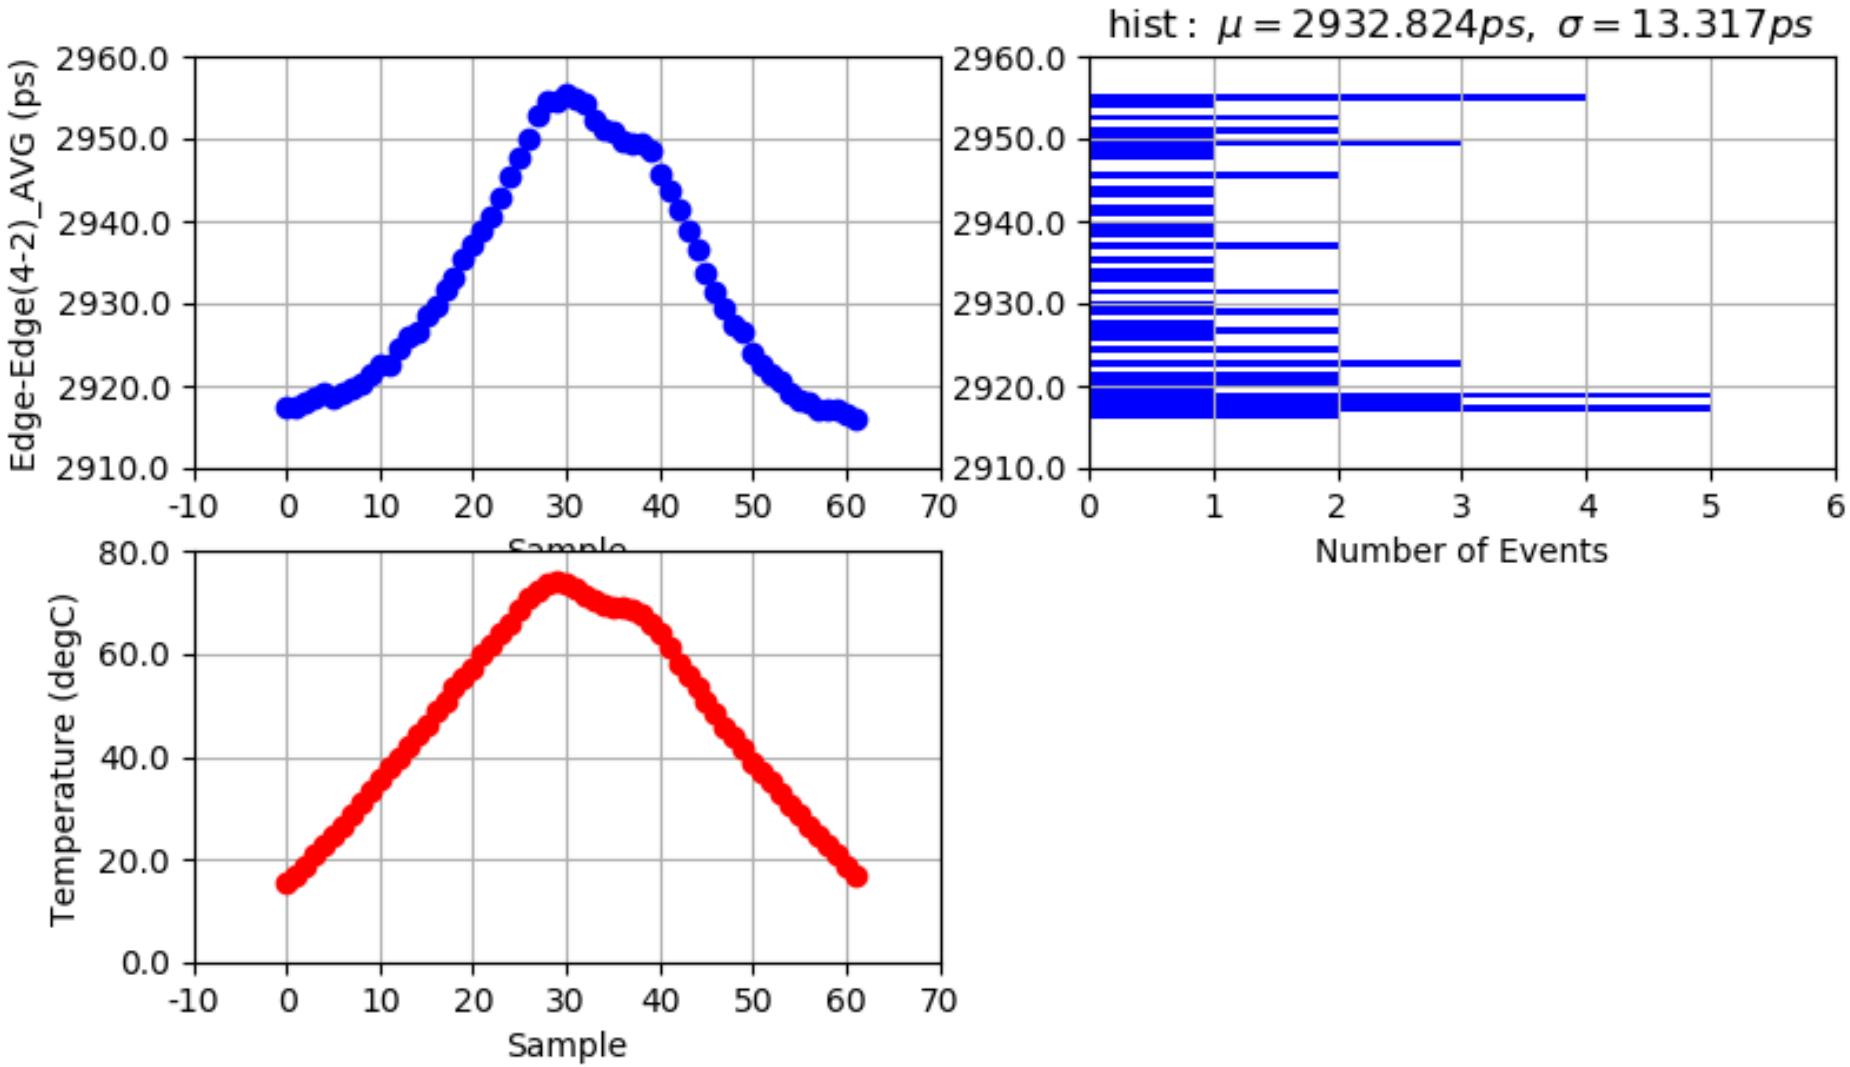

Characterization: Tx only temperature

- Tx aligner core: **UI_ALIGN** mode (i.e. deterministic latency)

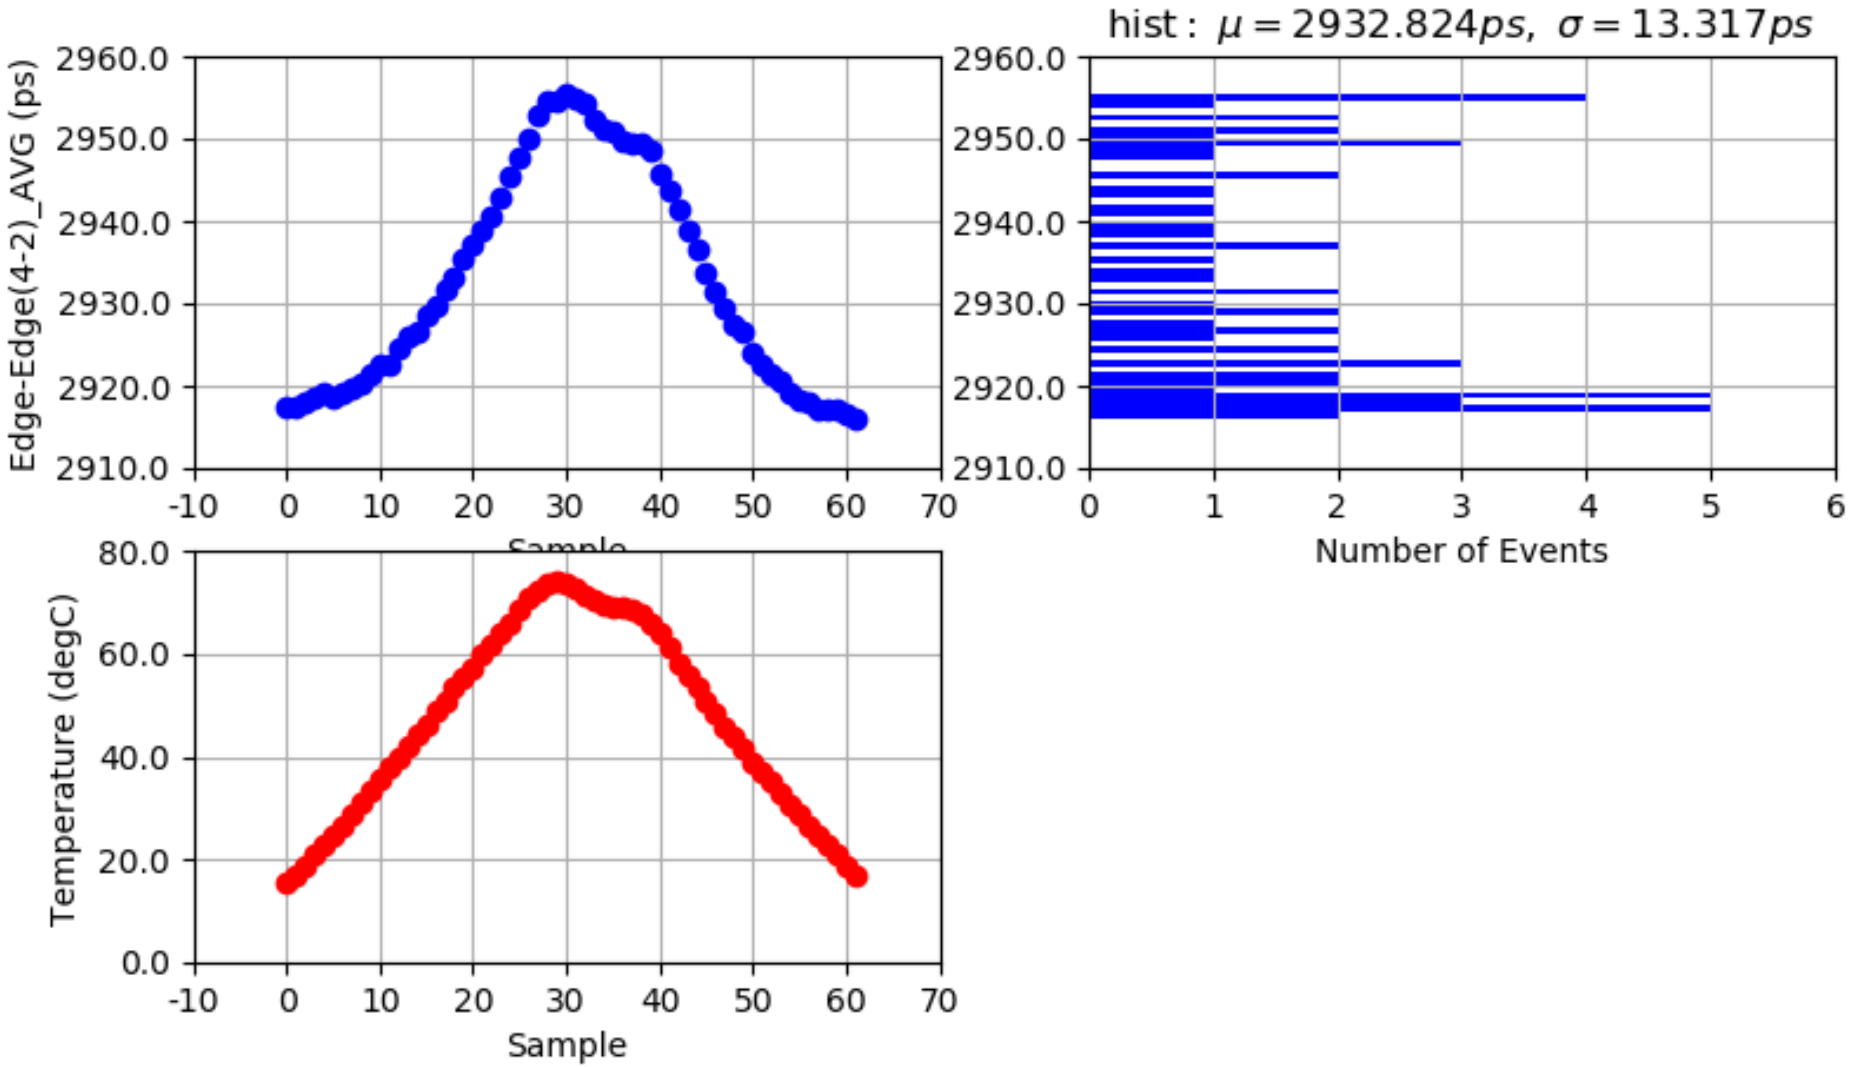

Transmitter path (GTY Ultrascale+)

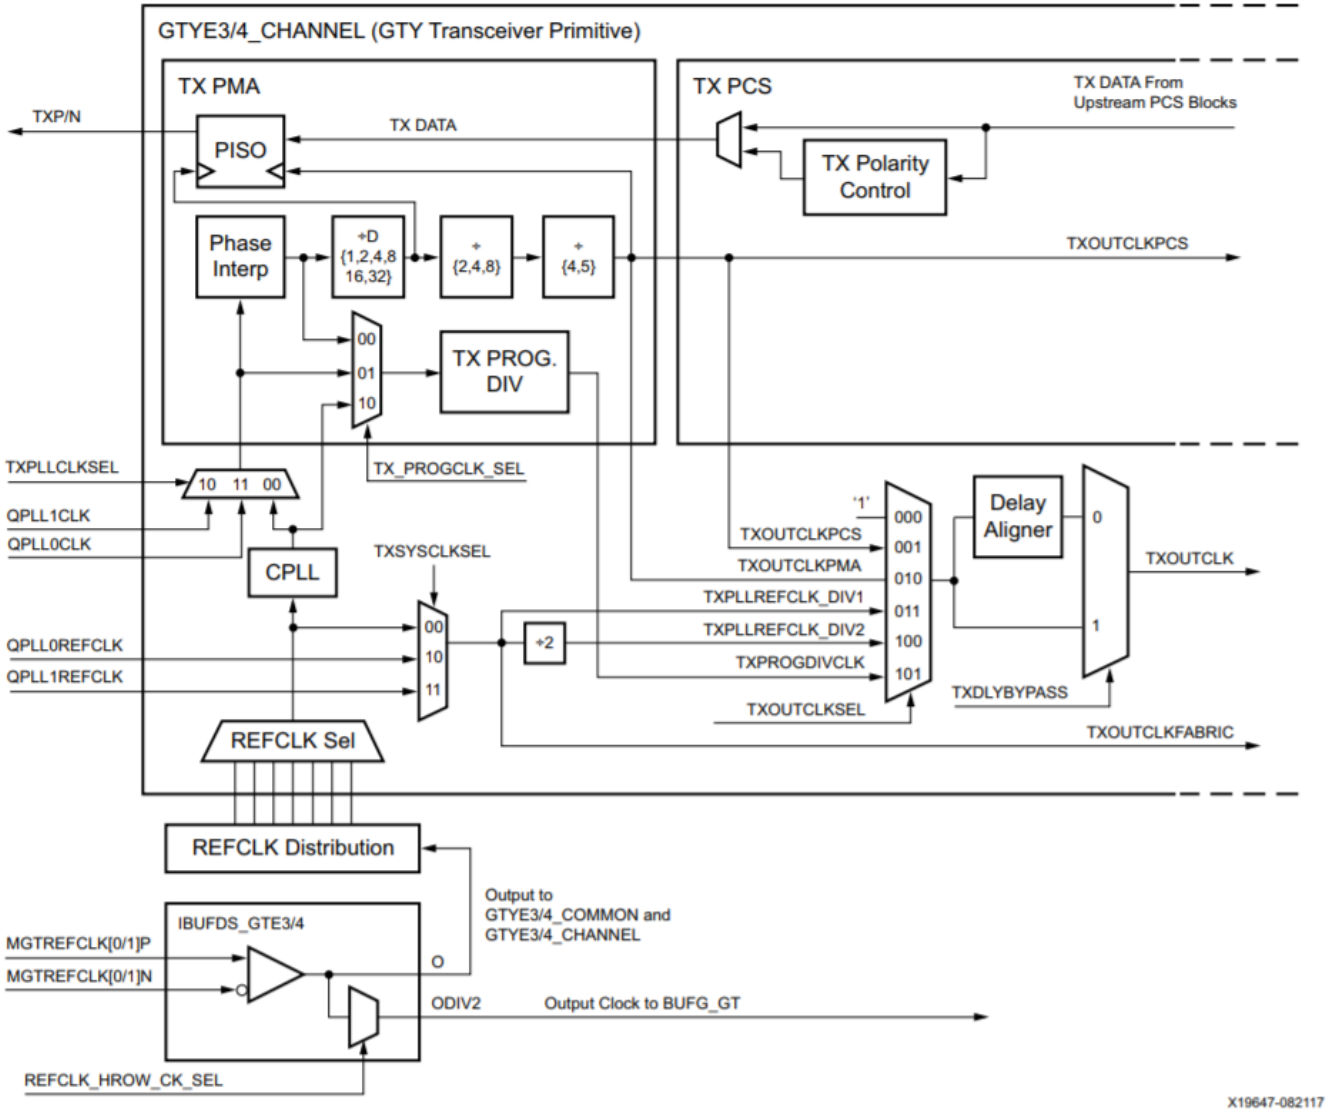

Figure 3-30: **TX Serial and Parallel Clock Divider**

Receiver path (GTY Ultrascale+)

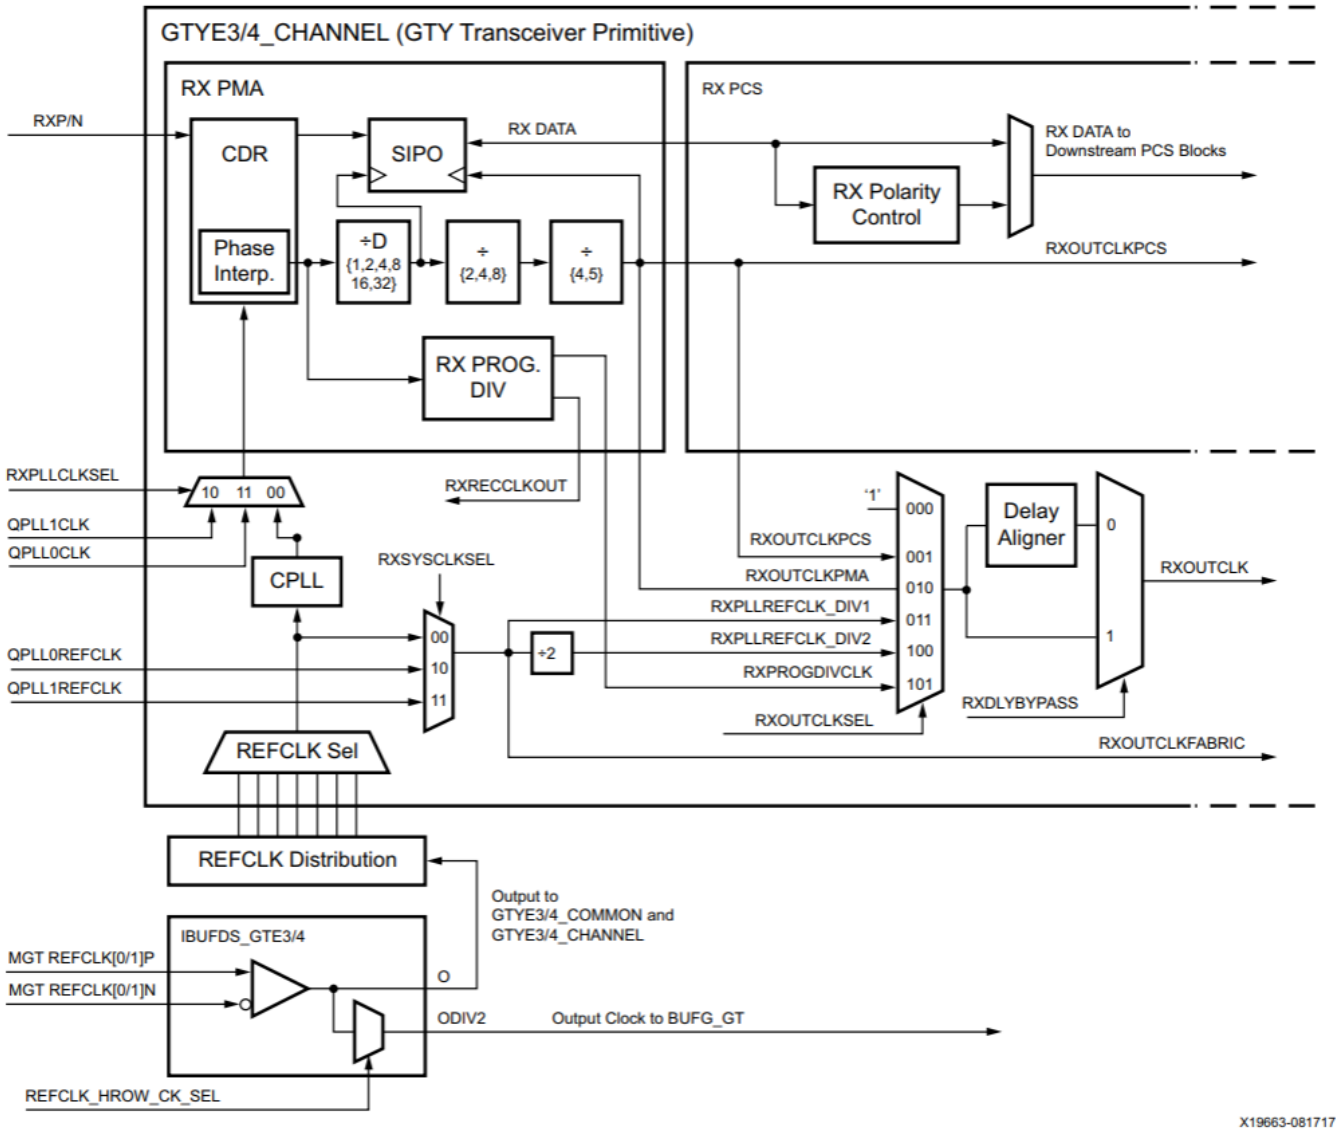

Figure 4-16: **RX Serial and Parallel Clock Divider**