Integrating lpGBT links into the Common Readout Units (CRU) of the ALICE Experiment

Erno DAVID on behalf of the ALICE collaboration
Wigner Research Centre for Physics, HUN-REN

This work was supported by
OTKA K-135515 and NKFI 2021-4.1.2-NEMZ_KI-2022-00009
grants

3 Oct 2023
TWEPP 2023, Geremeas, Sardinia, Italy

lpGBT Coming in ALICE…

- During the LS2 upgrade a brand new DAQ and trigger system had been developed for ALICE for Run 3 and Run 4
- The upgrading sub-detectors are now connected to the DAQ and Trigger systems with rad-hard **GBT links** through the CRU
- This enables the delivery of timing & control with deterministic latency and taking of data through a single fiber connection
- The GBTx ASIC is not available any more and the new lpGBT supersedes it for new developments or system additions

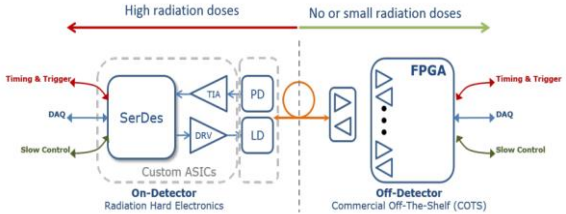

- During the coming LS3 upgrades, the new FE systems (e.g. ITS3, FoCal) will (have to) use lpGBT links to connect to the CRUs
- The lpGBT links have to be integrated into the existing CRU FW while keeping the compatibility with the existing O2, TRIGGER, and DCS systems

Main features of the present GBT links

- 4.8 Gb/s downlink
- 4.8 Gb/s uplink
- Front-end components:
  - GBTx ASIC
  - external slow-control (I2C, SPI, etc.) controller ASIC (SCA)
  - Versatile Link (VL) optical components

VS.

Main features of the new lpGBT with VTRX+

- 2.56 Gb/s downlink,
- 5.12/10.24 Gb/s uplink
- Front-end components:
  - lpGBT ASIC
  - internal slow-control controllers (I2C, SPI, GPIO, ADC, etc.) (and optional external SCA)
  - VL+ optical components

Project Goal

The goal of this *lpGBT in CRU* project is a *feasibility study* with the following consecutive steps:

1. Implementing the *lpGBT-FPGA IP* developed by the CERN EP-ESE team in the CRU *Arria10 GX FPGA* in a single-link **standalone design** and test the functionality at the lowest level (links only with pattern generators / checkers)

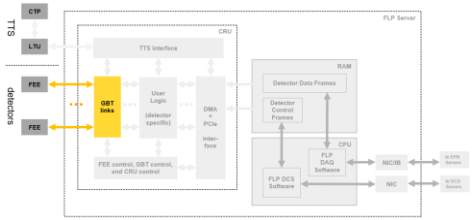

2. **Integrating** the *lpGBT-FPGA IP* in the common CRU FW by removing the GBT interfaces and
   - adding 1 lpGBT interface with full clock & trigger integration
   - adding 12 lpGBT interface (2 TRX banks) with full clock & trigger integration
   - optional: adding 24 lpGBT interface with full clock, trigger, and data integration

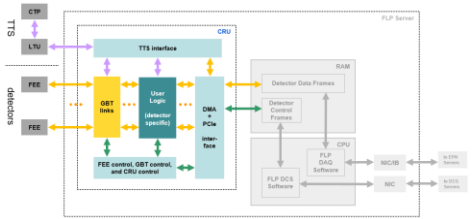

- The aim is not a production firmware, but to study and demonstrate that the replacement of the GBT links with lpGBT links in the CRU is possible while keeping the compatibility with the connected systems (O2, TTS, DCS)

Test Setup (simplified)

The setup is the same at Wigner RCP and
in the CERN CRU-TRIG integration lab

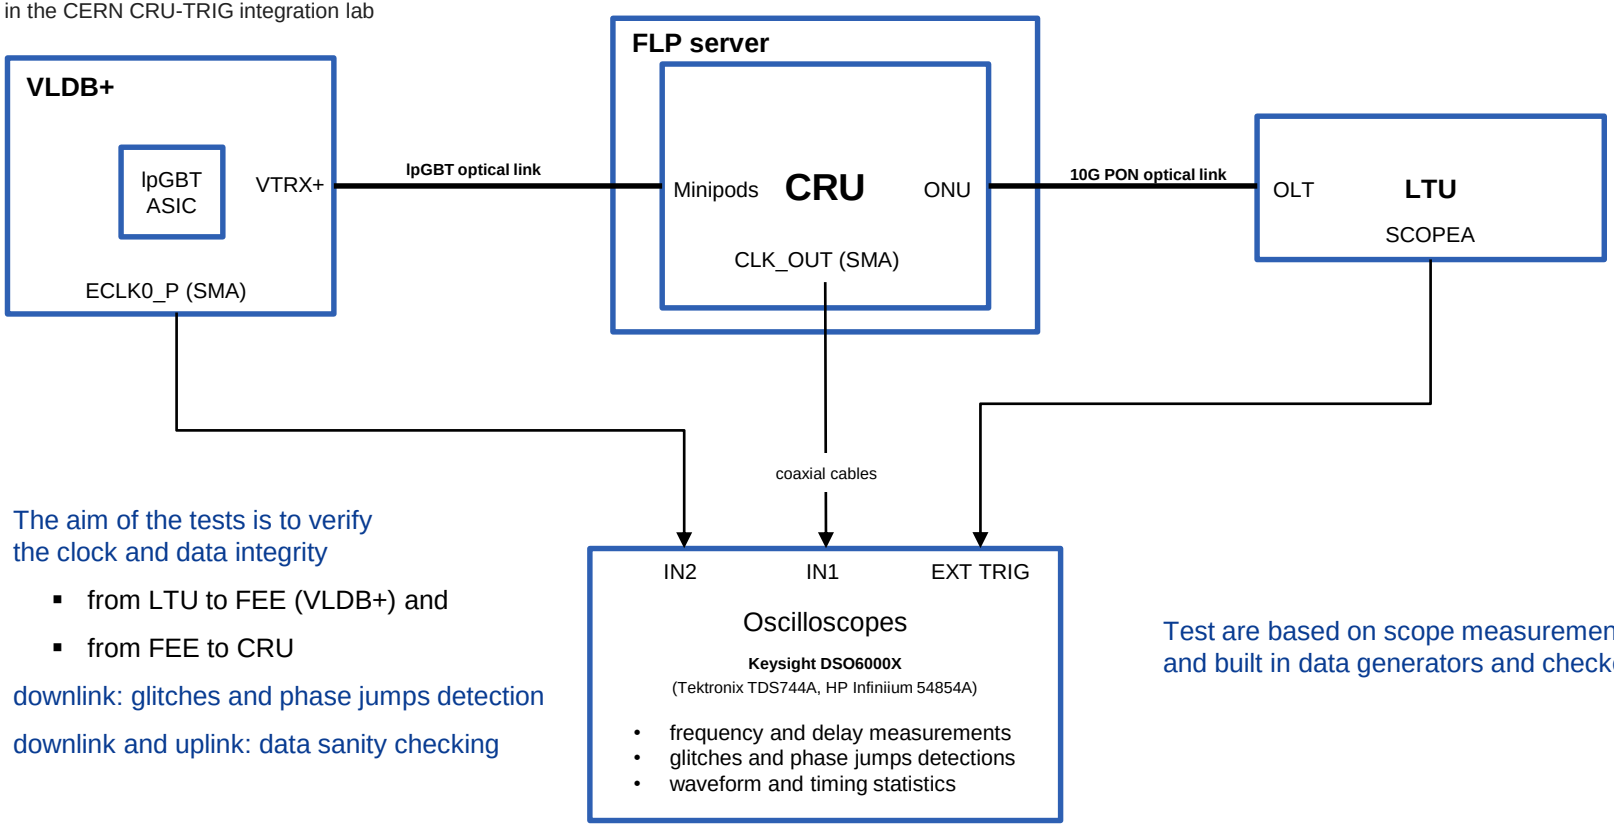

The aim of the tests is to verify
the clock and data integrity

- from LTU to FEE (VLDB+) and
- from FEE to CRU

downlink: glitches and phase jumps detection

downlink and uplink: data sanity checking

Test are based on scope measurements
and built in data generators and checkers

GBT / lpGBT Differences

DOWNLINK

Header, internal and external control channels, data channel, forward error correction

differences:

- 64-bit @ 40 MHz vs 120-bit @ 40 MHz
- 32-bit payload vs 80-bit payload
- 2.56 Gb/s vs 4.8 Gb/s
- TX parallel clock 320 MHz vs 240 MHz

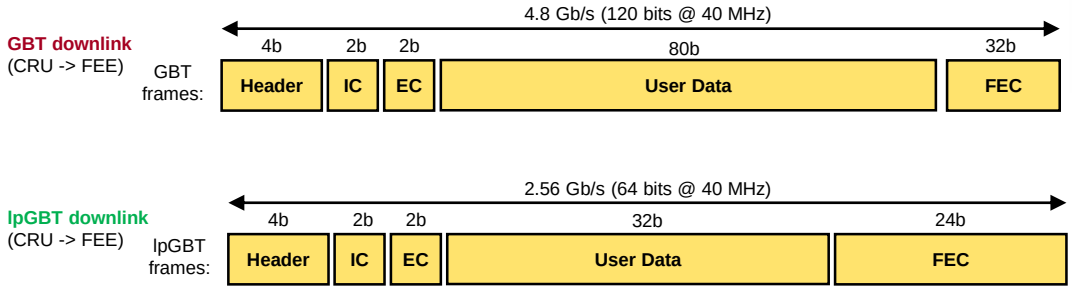

UPLINK

Header, internal and external control channels, data channel, optional forward error correction

differences:

- 128/256-bit @ 40 MHz vs 120-bit @ 40 MHz
- 96/112/192/224-bit payload vs 80/112-bit payload
- 5.12 Gb/s or 10.24 Gb/s vs 4.8 Gb/s
- RX parallel clock 320 MHz vs 240 MHz

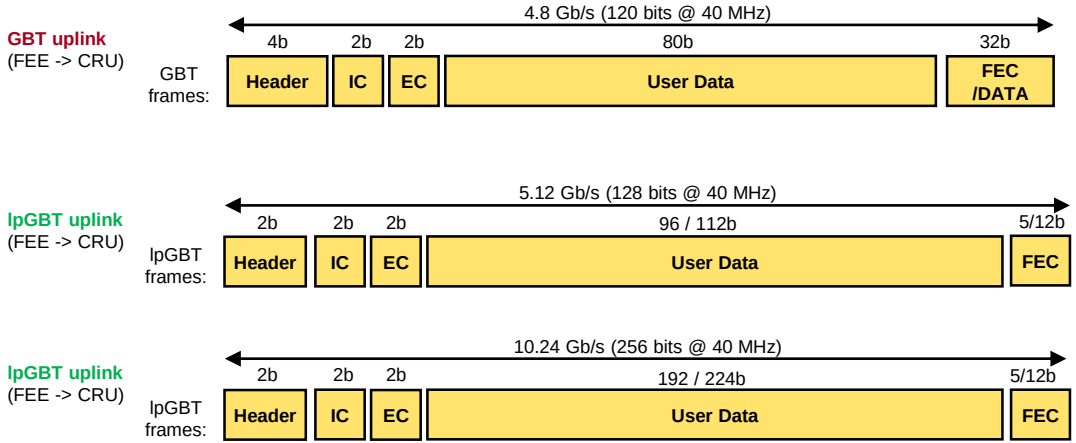

ALICE Common Read-out Unit (CRU)

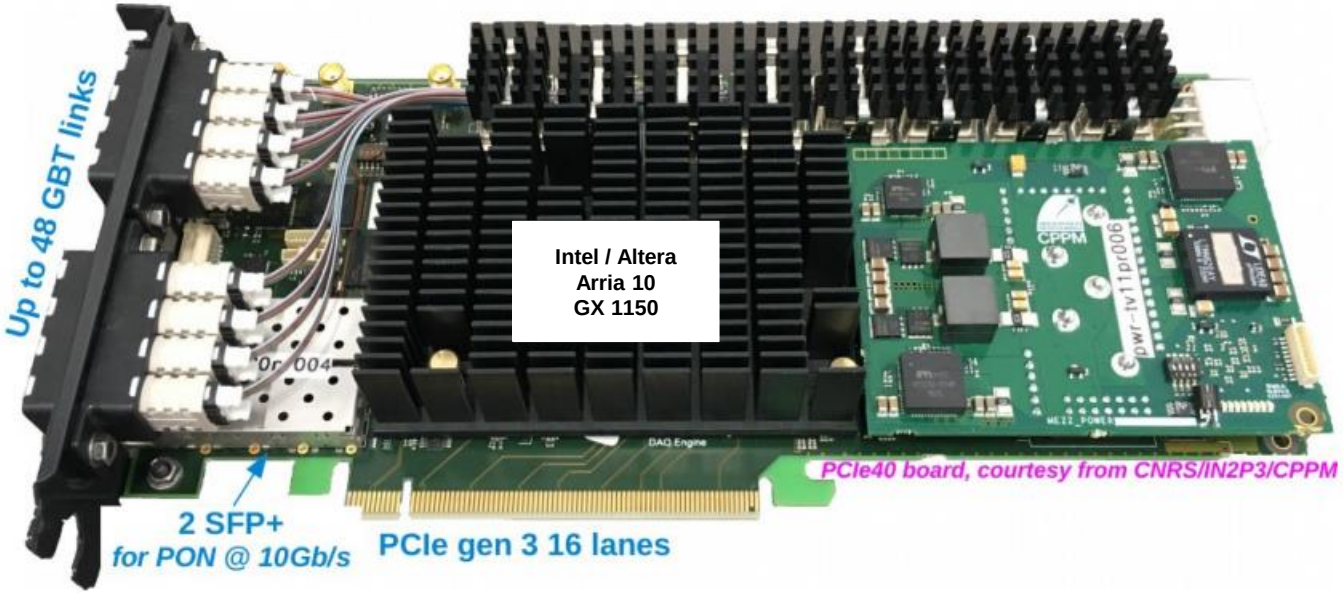

The Common Read-out Units (CRU) are PCIe add-on cards installed in the First Level Processor (FLP) nodes of the ALICE DAQ system. Main tasks of the CRU:

- **Deliver the trigger, timing and read-out control information to the Front-End Electronics**
- **Deliver detector data to the O2 (FLP Servers) with and/or without processing in the CRU FPGA**
- **Transport detector control information between the DCS and the FEE**
- **Take part of the Busy / Drop / Throttle mechanism of the detectors read-out**

CRU Firmware – Main Data Paths

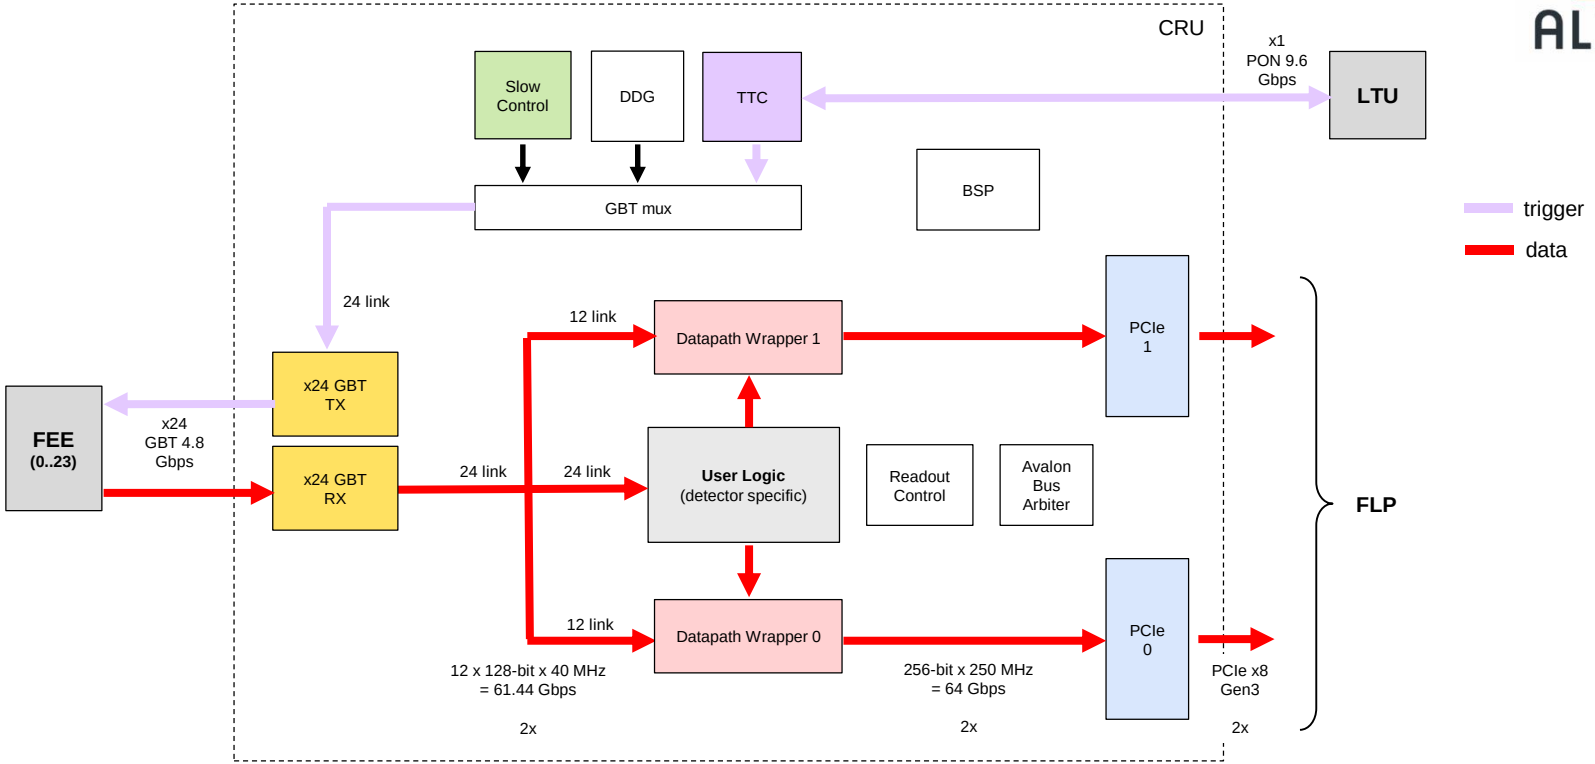

CRU – Clock and Trigger (GBT version)

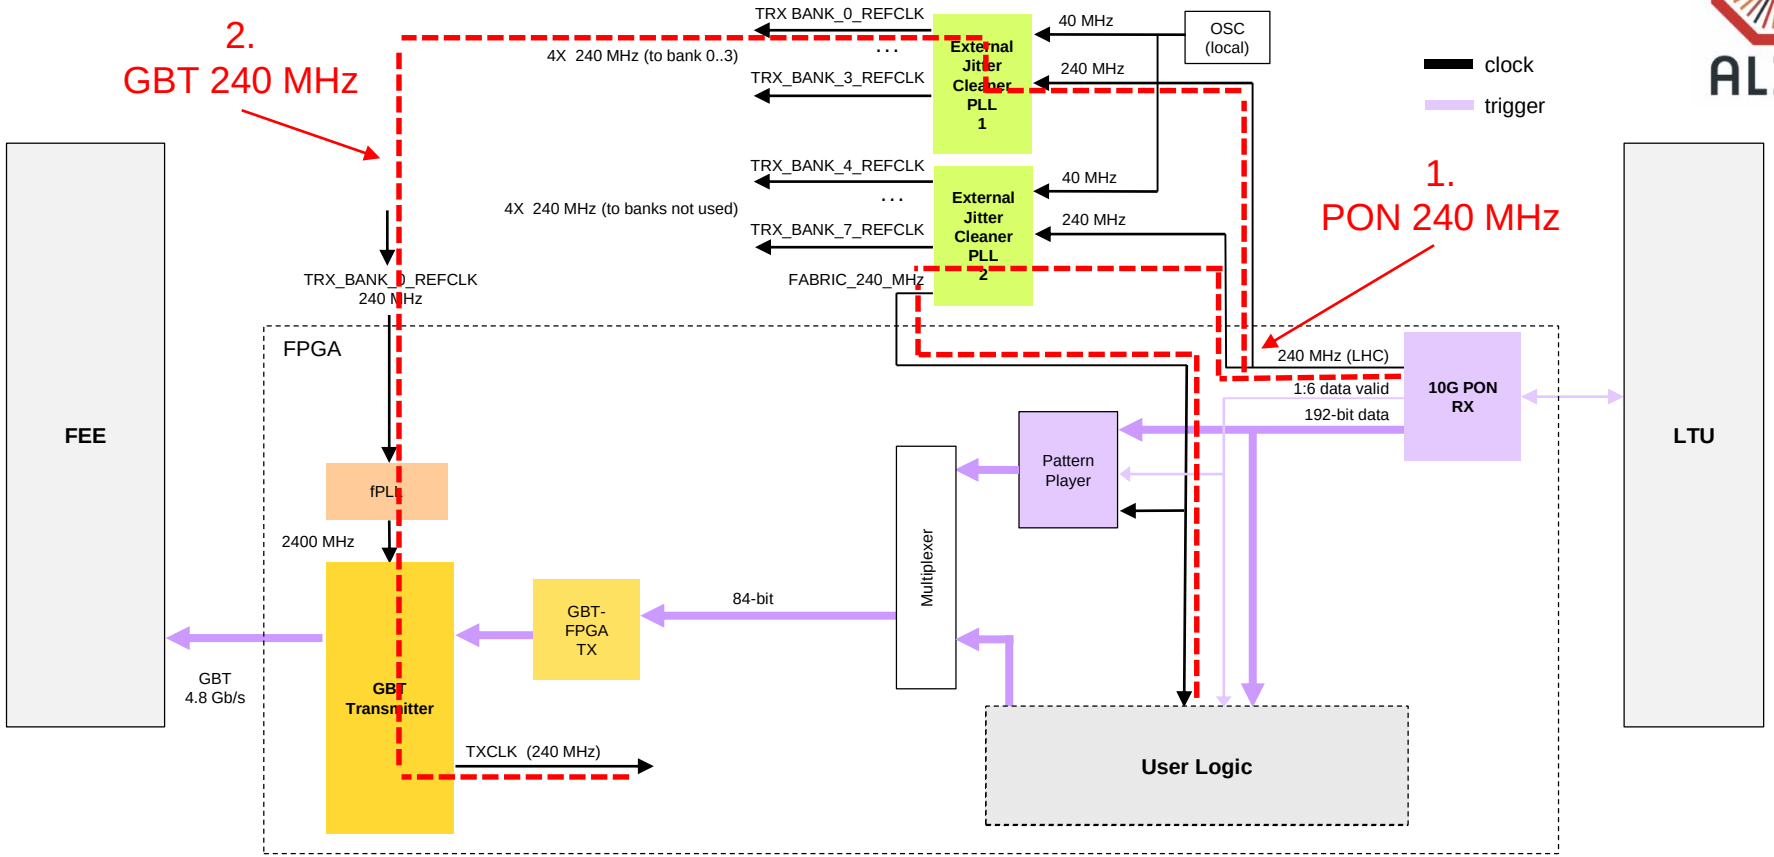

Problem and Concept

GBT CRU

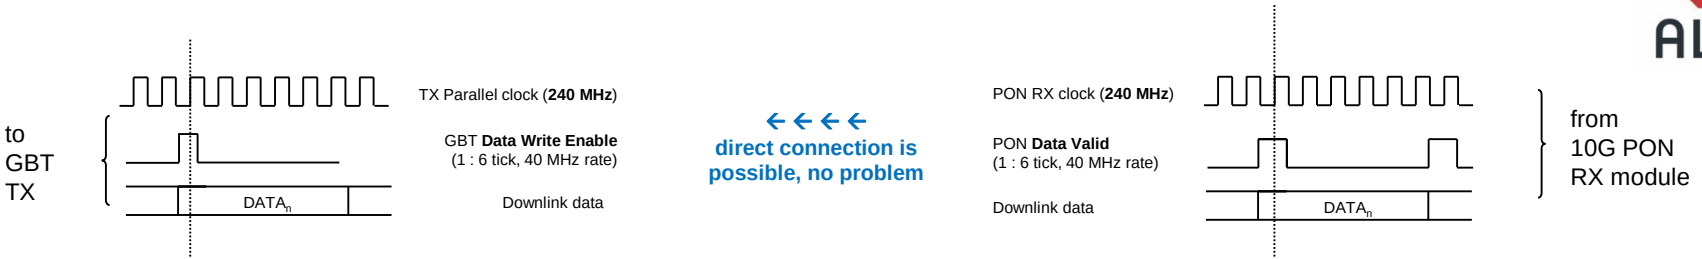

Problem and Concept

GBT CRU

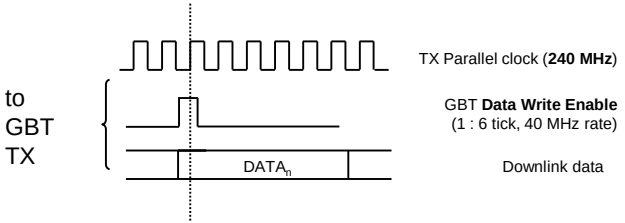

← ← ← ←
direct connection is possible, no problem

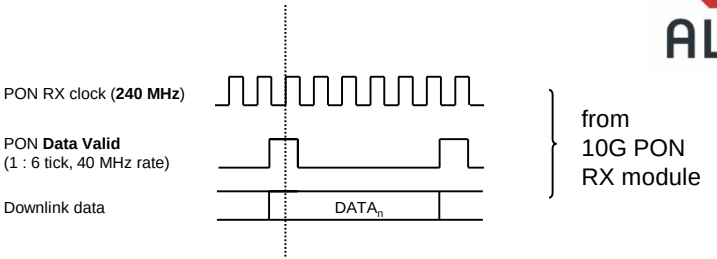

lpGBT CRU

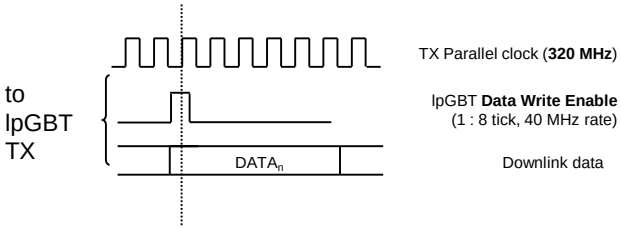

4:3 clock ratio!
?
You have to divide clock
→ phase uncertainty !

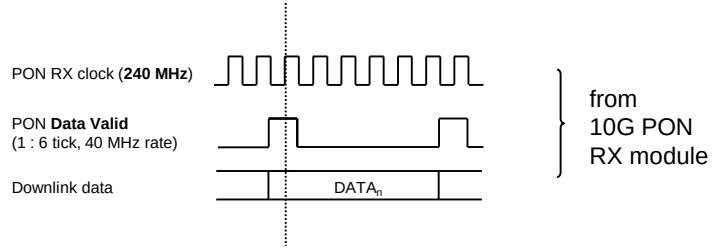

Problem and Concept

GBT CRU

lpGBT CRU

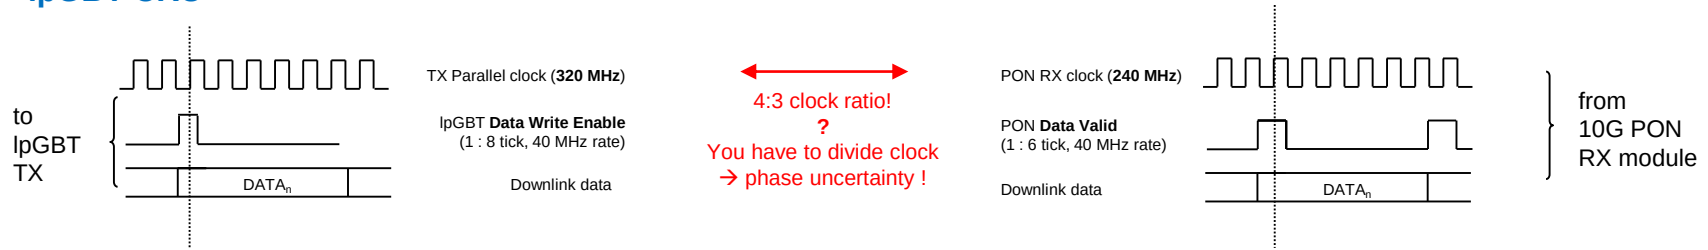

solution (concept)

We have to go down to the common 40 MHz clock domain and align data at both sides to the common 40 MHz clock!

Problem and Concept

GBT CRU

lpGBT CRU

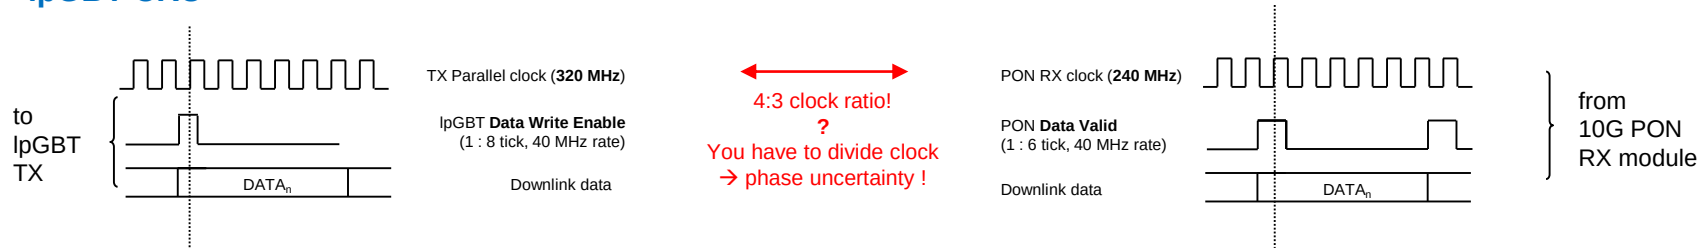

solution (concept)

We have to go down to the common 40 MHz clock domain and align data at both sides to the common 40 MHz clock!

Aligning the 40 MHz Clock to Data Valid

- The PON module recovers the 240 MHz reference clock and a data valid bit with deterministic constant delay to the 40 MHz LHC clock rising edge.
- The IOPLL recovers the 40 MHz dividing the 240 MHz with six. At the output of the IOPLL the two clocks can randomly have six different phase relations.
- If the 40 MHz rising edge is not aligned with the data valid bit then the Control FSM resets the IOPLL.
- Within a reasonable time (practically within a few tries) the data valid bit and the 40 MHz clock will be aligned.

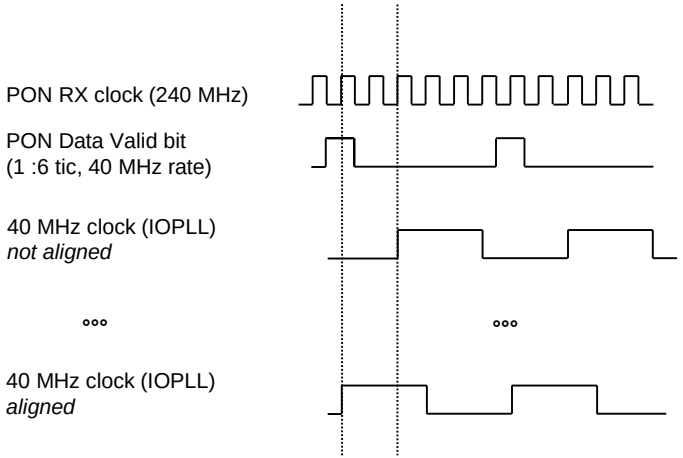

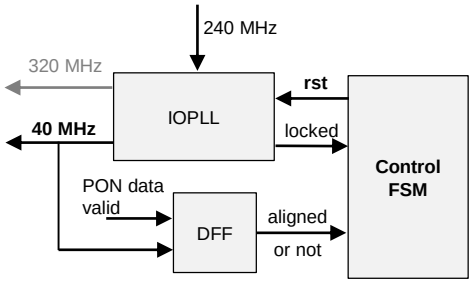

Aligning lpGBT Data Write Enable to the 40 MHz Clock

- The lpGBT data write enable bit is generated in the 320 MHz clock domain with a 40 MHz rate.
- We sample this bit with our 40 MHz clock which is in sync with the 320 MHz clock.
- If it is not aligned with the 40 MHz clock then the Control FSM shifts the write enable bit (bit slip).
- In a maximum of eight steps the write enable bit will be aligned with the 40 MHz clock.

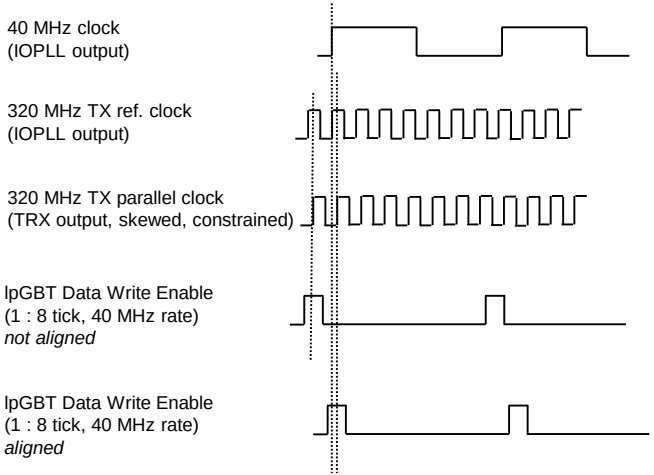

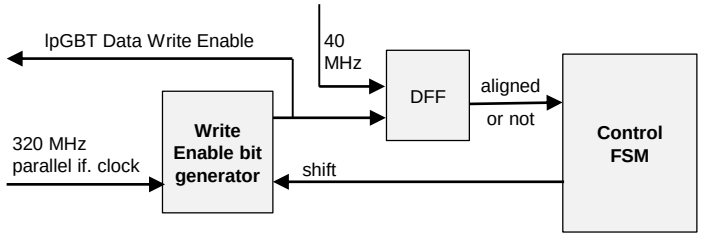

Clock and Trigger Integration / Jitter Cleaning

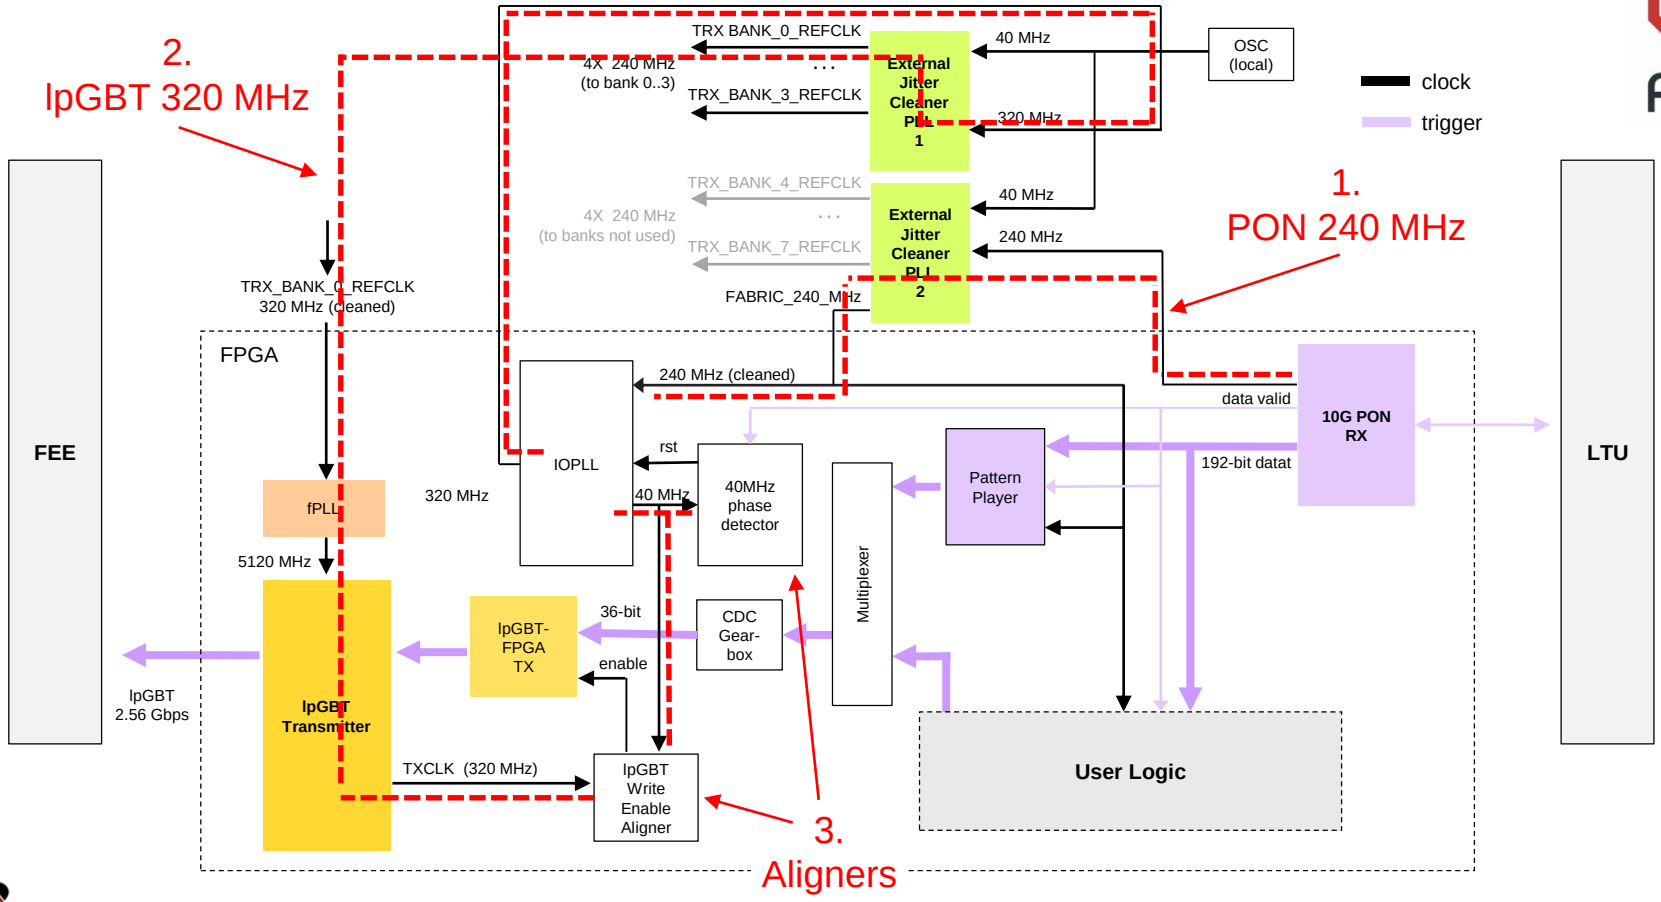

First Test Results (Downlink Clock and Trigger Delay Stability I.)

LHC clock & trigger delivery measurements:
Delay and stability from LTU to FEE (VLDB+)

- **IN1 (yellow): DATA VALID (1:6 tick)**: PON link output (CRU CLK-OUT SMA conn.)
- **IN2 (green): VLDB+ 40 MHz**: lpGBT link output (VLDB+ E0CLK_P SMA conn.)
- **TRIGGER:** external, LTU 40 MHz output (SCOPEA conn.)

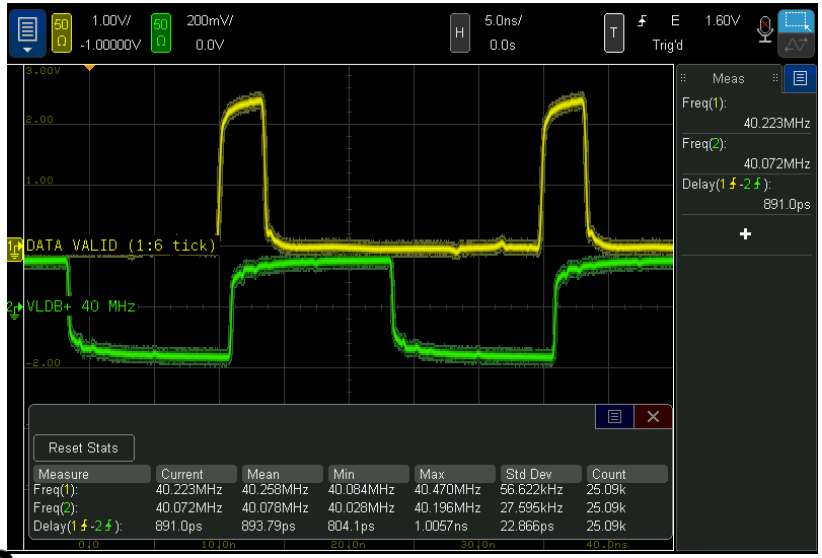

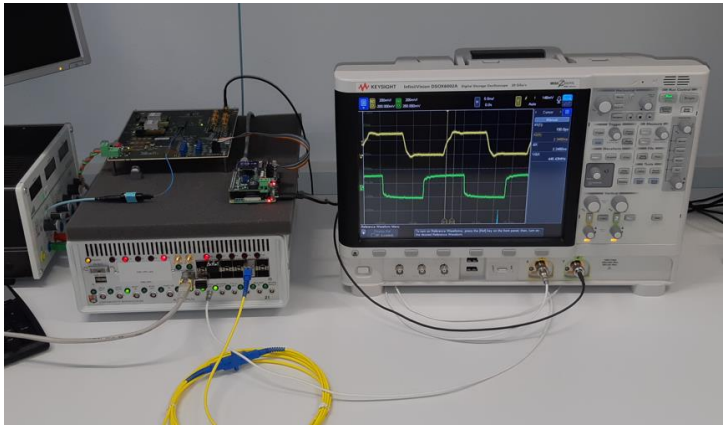

Visualizing the 40 MHz LHC clock, DATA VALID, and lpGBT TX WRITE_ENABLE flags via the clock chain (1 or 12 links being implemented) with infinite persistence for a few hours while challenging the stability with: **optical link disconnections**, **power cycling**, and **PLL forced initializations**

- **DATA_VALID (1 link, 12 links)** → stable with the same delay
  (no glitches, phase jumps observed, continuous 10^{12} clock cycle)
- **TX Write_Enable (1 link, 12 links)** → stable with the same delay
  (no glitches, phase jumps observed, continuous 10^{12} clock cycle)
- **VLDB+ ECLK0 (1 link, 12 link, internal feeding)** → stable with the same delay
  (no glitches, phase jumps observed, continuous 10^{12} clock cycle)

First Test Results (Downlink Clock and Trigger Delay Stability II.)

Two options for the clock integration (external/internal) has been tested for delay stability and with simple jitter measurements. (1 link only.)

lpGBT TX_REF: externally distributed to the lpGBT transceiver banks through one of the CRU's two on-board jitter cleaner PLL. (The other filters the incoming 240 MHz PON RX_CLK.)

lpGBT TX_REF: internally distributed to the lpGBT transceiver banks ("Only" the incoming 240 MHz PON RX_CLK is cleaned with one of the CRU's two on-board jitter cleaner PLLs.)

No instabilities or significant difference in jitter has been observed. Further jitter analysis may follow.

Uplink Testing

Link ID	GBT Mode Tx/Rx	Loopback	GBT MUX	Datapath mode	Datapath status	RX freq (MHz)	TX freq (MHz)	Status	Optical power (uW)	System ID	FEE ID
0	GBT/GBT	None	TTC:CTP	Streaming	Enabled	259.55	320.63	DOWN	0.0	0x3	0x0
1	GBT/GBT	None	TTC:CTP	Streaming	Enabled	320.44	320.63	DOWN	0.0	0x3	0x0
2	GBT/GBT	None	TTC:CTP	Streaming	Enabled	321.17	320.63	DOWN	0.0	0x3	0x0
3	GBT/GBT	None	TTC:CTP	Streaming	Enabled	320.02	320.63	DOWN	0.0	0x3	0x0
4	GBT/GBT	None	TTC:CTP	Streaming	Enabled	317.92	320.63	DOWN	0.0	0x3	0x0
5	GBT/GBT	None	TTC:CTP	Streaming	Enabled	319.63	320.63	DOWN	0.0	0x3	0x0
6	GBT/GBT	None	TTC:CTP	Streaming	Enabled	320.63	320.63	UP	900.0	0x0	0x0
7	GBT/GBT	None	TTC:CTP	Streaming	Enabled	319.88	320.63	DOWN	0.0	0x0	0x0
8	GBT/GBT	None	TTC:CTP	Streaming	Enabled	298.17	320.63	DOWN	0.0	0x0	0x0
9	GBT/GBT	None	TTC:CTP	Streaming	Enabled	254.86	320.63	DOWN	0.0	0x0	0x0
10	GBT/GBT	None	TTC:CTP	Streaming	Enabled	257.52	320.63	DOWN	0.0	0x0	0x0
11	GBT/GBT	None	TTC:CTP	Streaming	Enabled	249.85	320.63	DOWN	0.0	0x0	0x0

Test results:

- The x12 link lpGBT module is recognized by the O2 software (link status, TX/RX frequency, etc.)
- The lpGBT ASIC "Loopback Downlink Group Data Source" mode was used with built in checkers in the lpGBT module
- The links were tested one-by-one with a single VLDB+ for 1 hour / link (10^{14} bits) → no errors
- Integration with the Datapath wrapper is ongoing

Summary and Further Work

Summary:

- x12 link lpGBT module (with 10.24 Gbps / FEC12 uplink mode) implemented inside the CRU-FW
- No visible phase jumps in the VLDB+ 40 MHz clock (10^{12} clock cycle)
- Stable data loopback through the lpGBT ASIC ("Loopback Downlink Group Data Source" mode) (10^{14} bits per link)

Further work:

- Connect the lpGBT module with the DMA and Slow Control modules
- Improve stability between rebuilt: floor planning and design lock
- Characterize skew between links
- Clock and trigger analysis
- Test with detector specific user logic at full 10.24 Gbps speed

Thank You for Your Attention!

BACK-UP

CRU clocking scheme

Clock Tree PCIe40V2

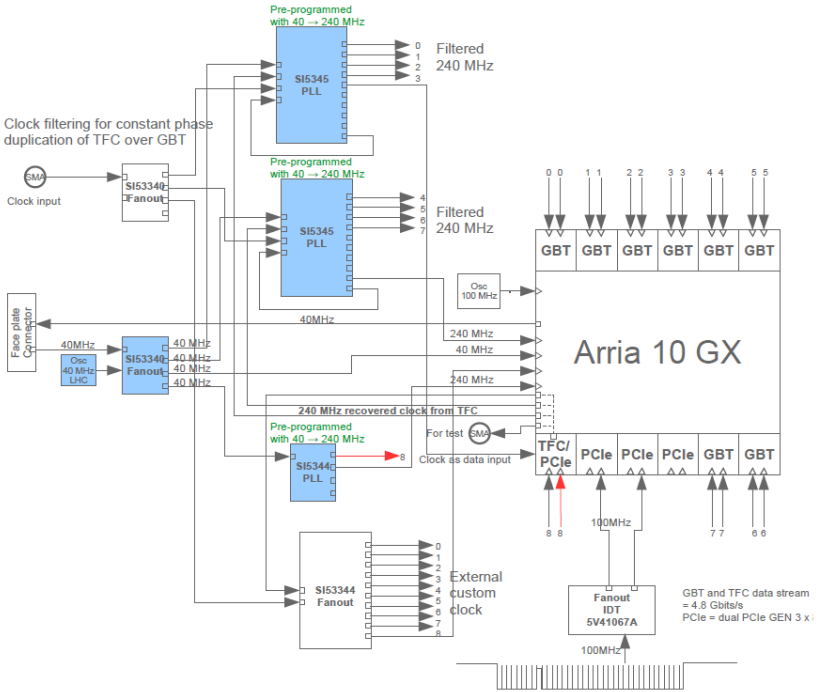

FW clocking

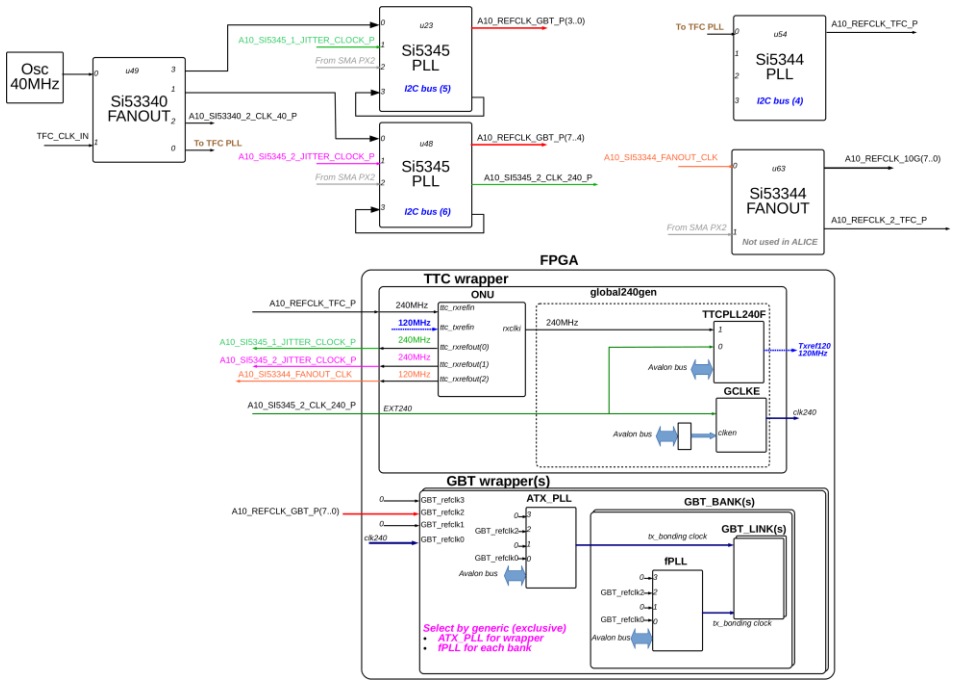

CRU System Connections

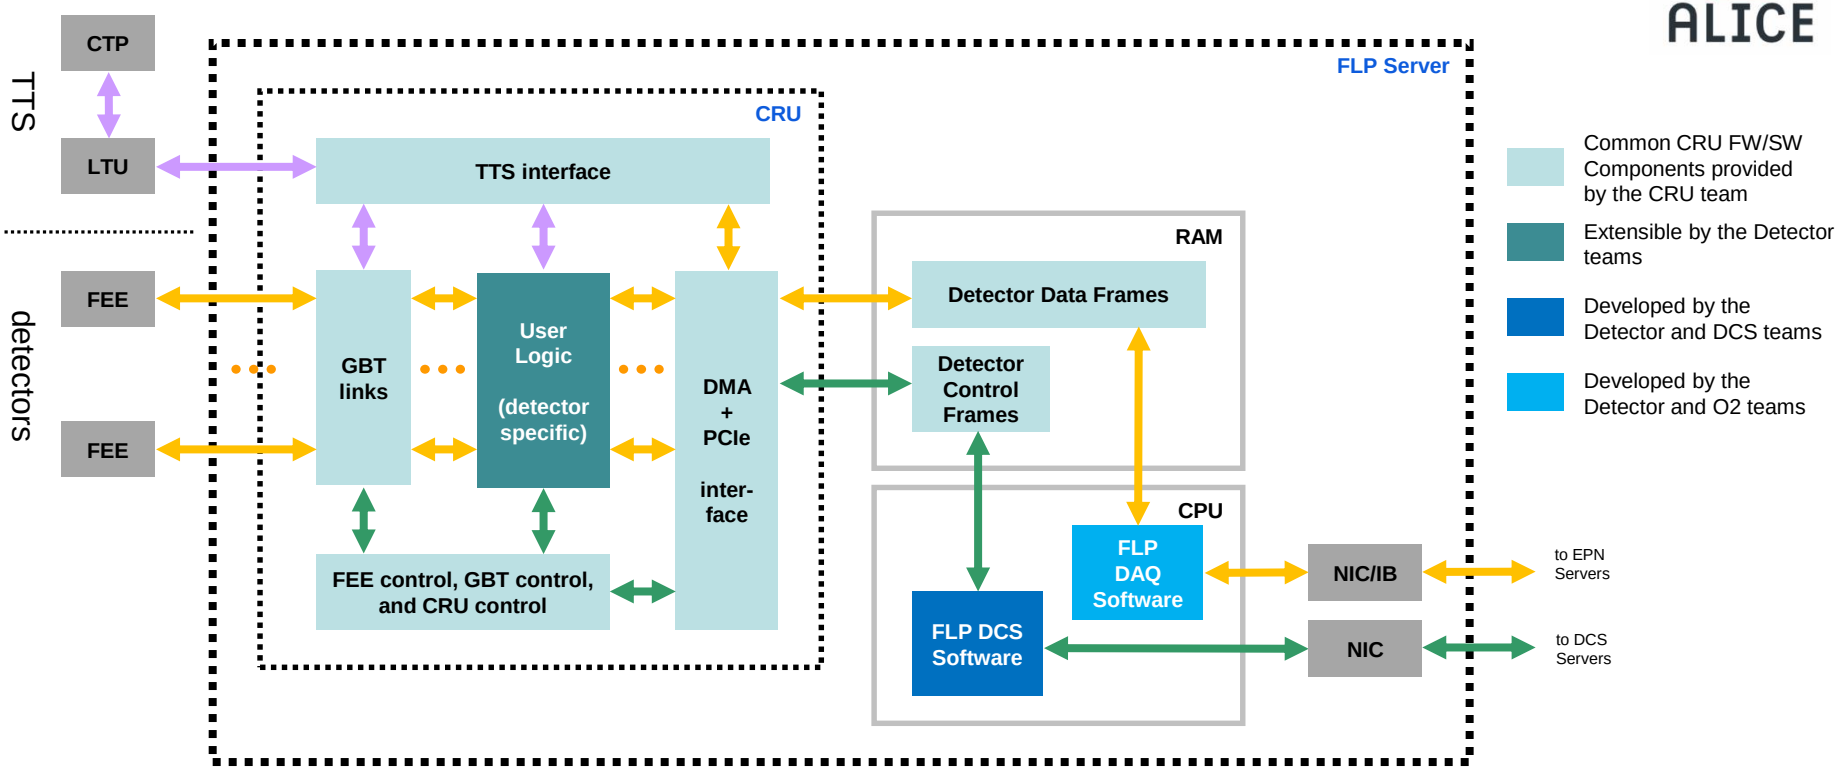

Clock and Trigger Integration / Jitter Cleaning (Option B, alternative)

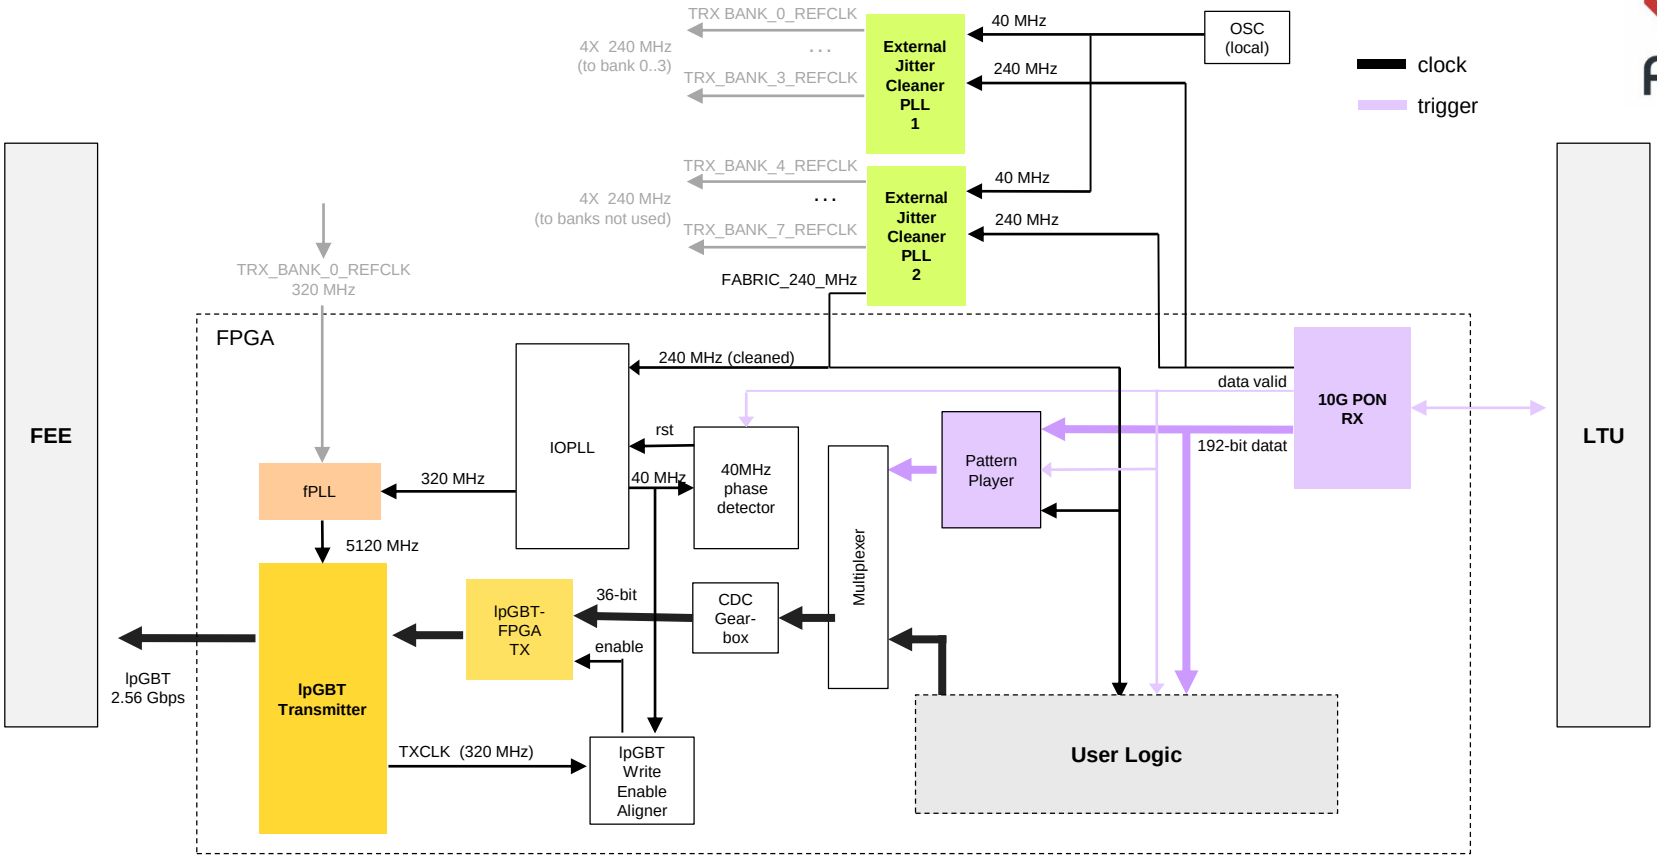

Uplink Integration

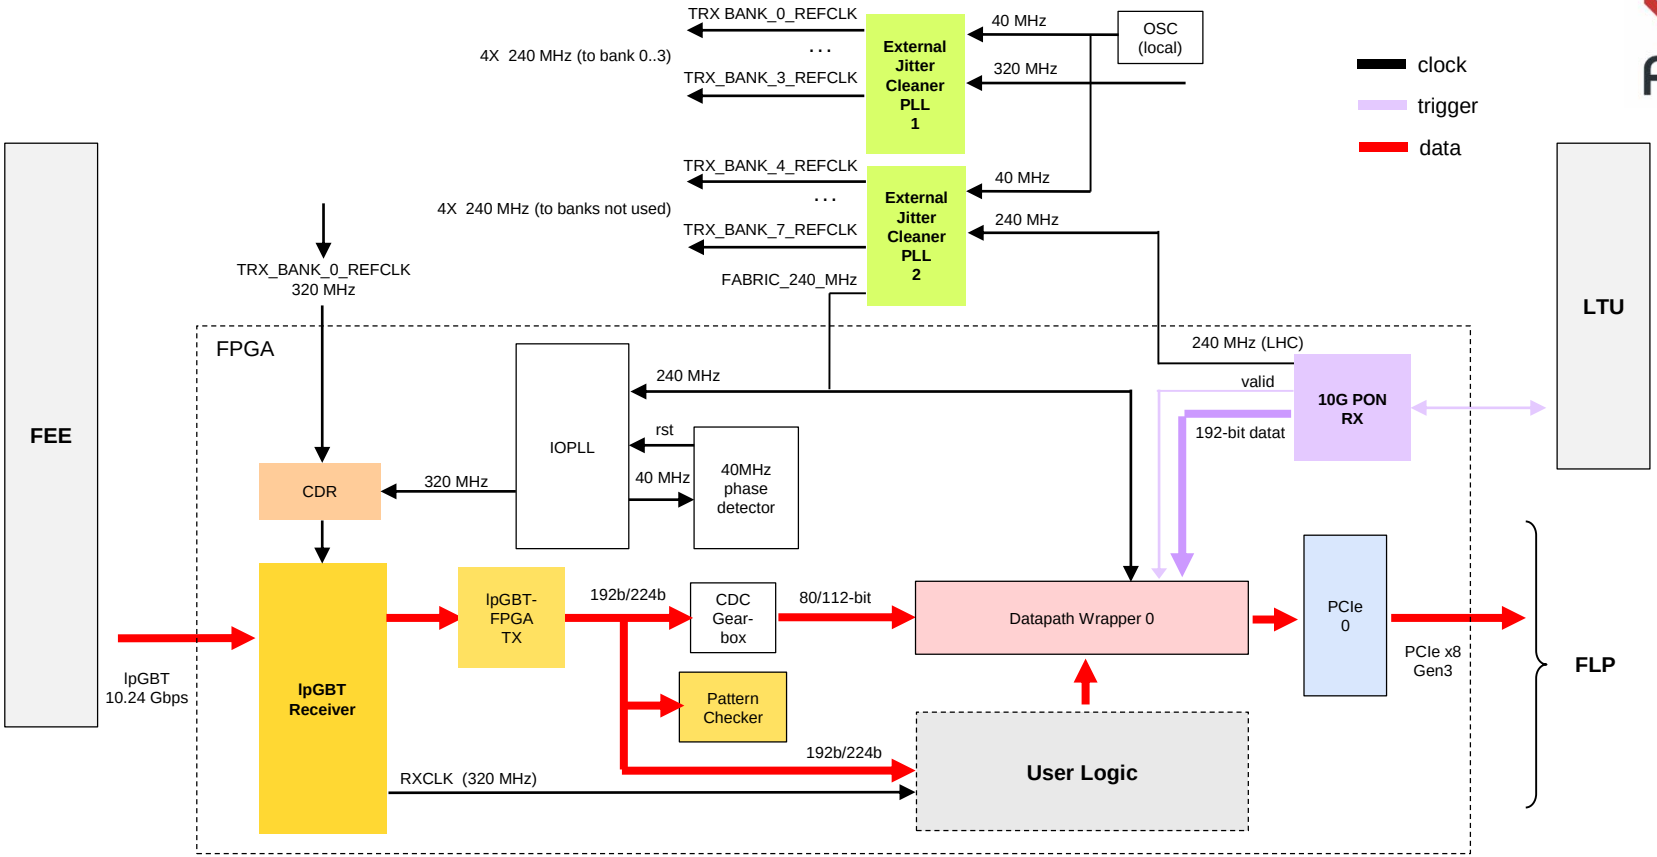

Uplink Testing

For the uplinks testing we have the following options:

- Use the built-in pattern generators in the lpGBT ASIC and check the received payload in the lpGBT module.
- Use the CRU firmware's Datapath Wrapper which is able to record the incoming raw data. Currently it supports normal GBT (80-bit) and wide bus mode (112-bit) mode, so it can check only part of the lpGBT data
- Later: write a dedicated user logic to process the full lpGBT payload and creates the standard ALICE data packets (SAP) and sends them to the PCIe DMA (through the Datapath Wrapper).

For our uplinks testing we have used the built in pattern generators and checkers with x12 lane 10.24 Gbps (FEC12) links. Test results so far:

- We have tested 1 standalone link for a few hours → no errors
- We have tested a 12-link implementation, the links one-by-one with a single VLDB+ for 1 hour /link → no errors
- Integration with the Datapath wrapper is ongoing *(may evolve till the conference)*

Testing the lpGBT Control Channels

- **The lpGBT ASIC supports both EC and IC 2-bit control channels:**
  - EC channel is backward compatible. It can connected to an GBT SCA on the FEE cards
  - IC channel requires upgrade to a newer GBT-SC firmware module provided by CERN Electronics group

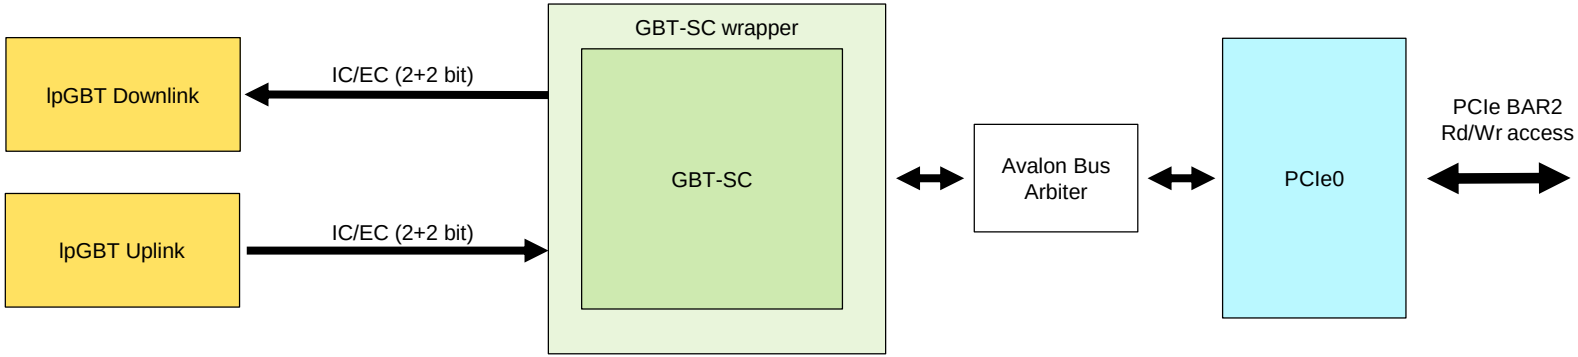

- **For control channel testing we can use the low level IC and EC field tester Python scripts used during the original CRU FW development. We can read and write the ASICs registers with them in the same way.**
- **To do this we have started to update the CRU FW with the new *slow control module (GBT-SC IP)* of the EP_ESE team and modify our software scripts to support the new lpGBT ASIC control registers (IC).**
- **The aim is to test reading out a few registers only as a demonstration.**