

Picosecond-level Clock Distribution System of the LHCb Experiment at CERN

Alberto Perro, Federico Alessio, Paolo Durante, Mitja Vodnik, Sera Conti



25th IEEE Real Time Conference - La Biodola, Elba, Italy - 25th May 2026

Outline


Current Timing Distribution Architecture

Future Requirements and Limitations

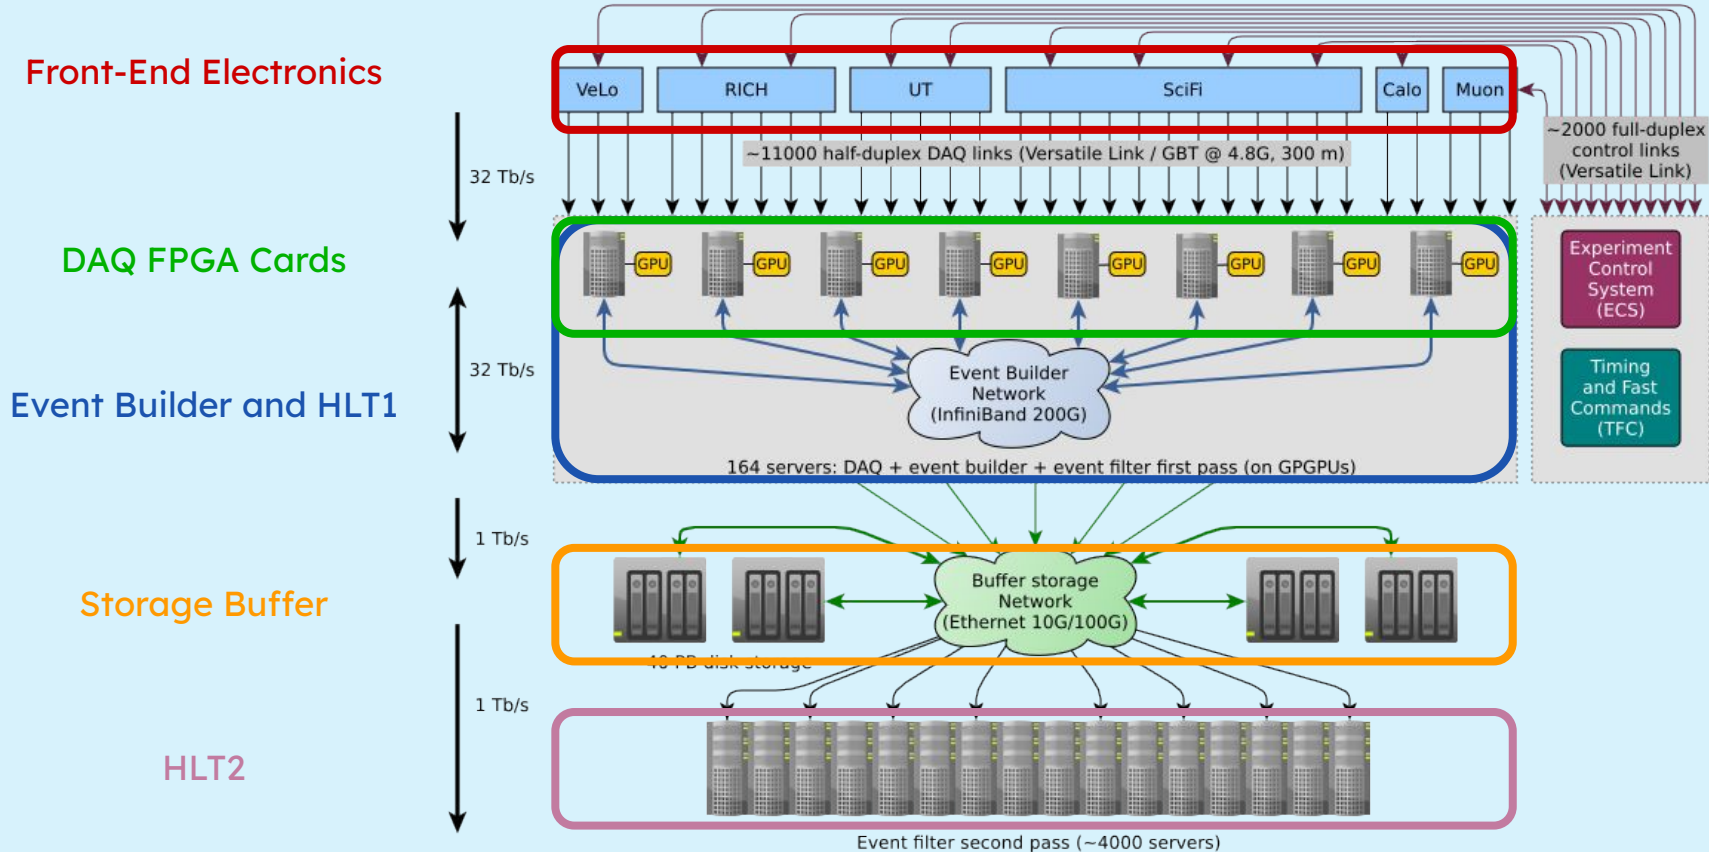
Critical Paths and Temperature Effects

Results and Future Directions

The LHCb Experiment



Online DAQ System Overview



Front-End Electronics

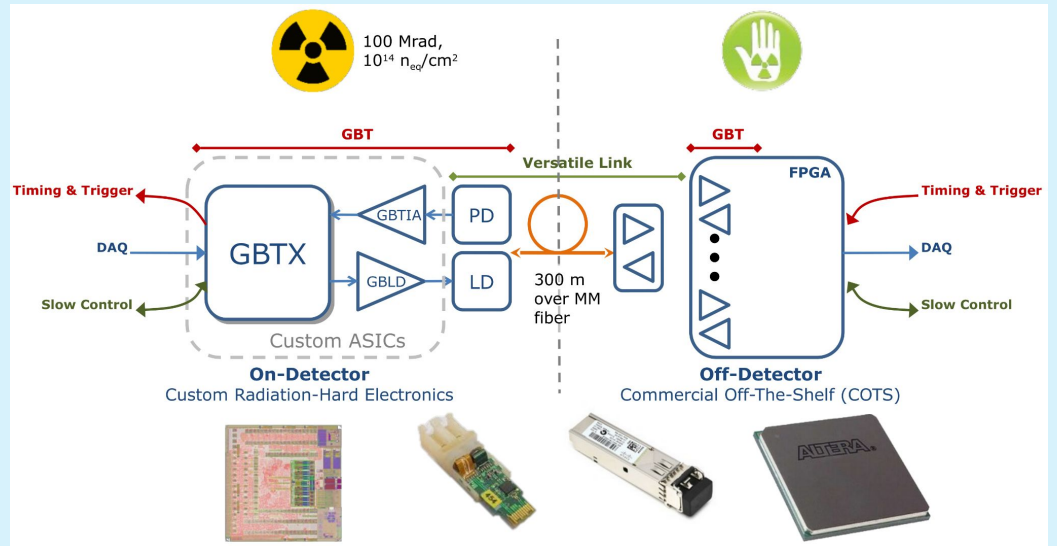
Front-End Electronics (FEE) is the first step in the DAQ chain. Its role is to readout, monitor, and control the sub-detectors.

The DAQ system is synchronous to the LHC bunch clock. It is necessary to **distribute the clock to each FE with deterministic latency**.


This is achieved using CERN's bidirectional radiation-hard **Versatile Link**, composed of the GBTx ASIC and the VTRx transceiver.

The Versatile Link is a multi-gigabit optical link capable of maintaining deterministic latency between FE and Back-End (BE).

Most sub-detectors use Versatile Link also for the data readout.



[10.5170/CERN-2009-006.342](https://cds.cern.ch/record/105170/files/CERN-2009-006.342)



PCIe40 FPGA Cards

- 1.2M Logic Elements Arria 10 FPGA
- 48 GBT Radiation Hard Optical Links @ 4.8Gbps
- Two SFP+ modules
- Two PCIe Gen3 x8 interfaces

520 cards are used in LHCb, subdivided in three flavors:

- **Readout supervisor (SODIN):** source of real-time control commands and entry point for the LHC bunch clock.
- **Control Card (SOL40):** It interfaces to the slow and fast control of the FEE and to the fast control of the DAQ cards.
- **Readout Card (TELL40):** It acquires data from the FEE, processes it, and delivers it to the host through the PCIe interface.

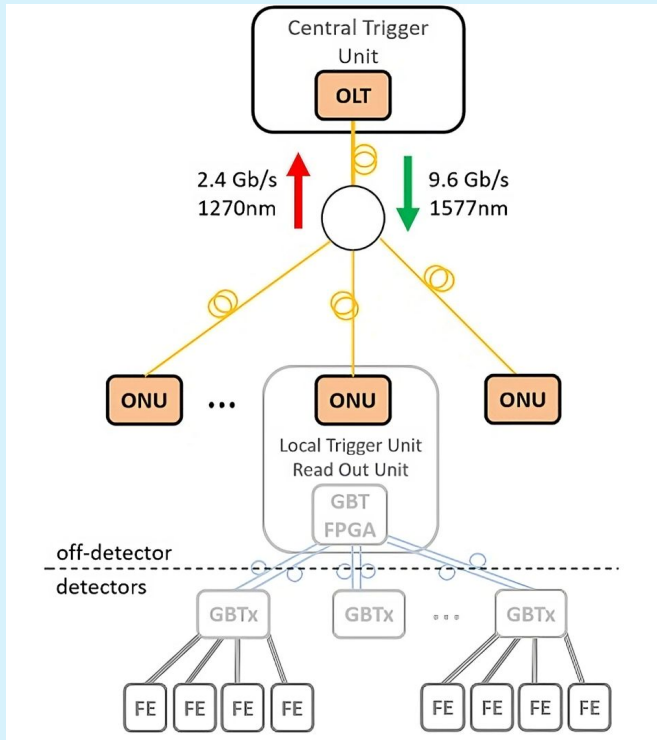
Timing and Fast Control (TFC) Subsystem

The TFC Subsystem is responsible for:

- **Distributing LHC bunch clock** with deterministic latency
- Generating and distributing real-time **commands** with fixed latency
- **Monitoring** all the readout modules in real-time

The TFC backbone is based on the *Timing, Trigger and Control for Passive Optical Networks (TTC-PON)* project.

Multiple Passive Optical Networks (PON) interconnect the PCIe40 FPGA cards.



[10.22323/1.313.0124](https://doi.org/10.22323/1.313.0124)

TFC Distribution

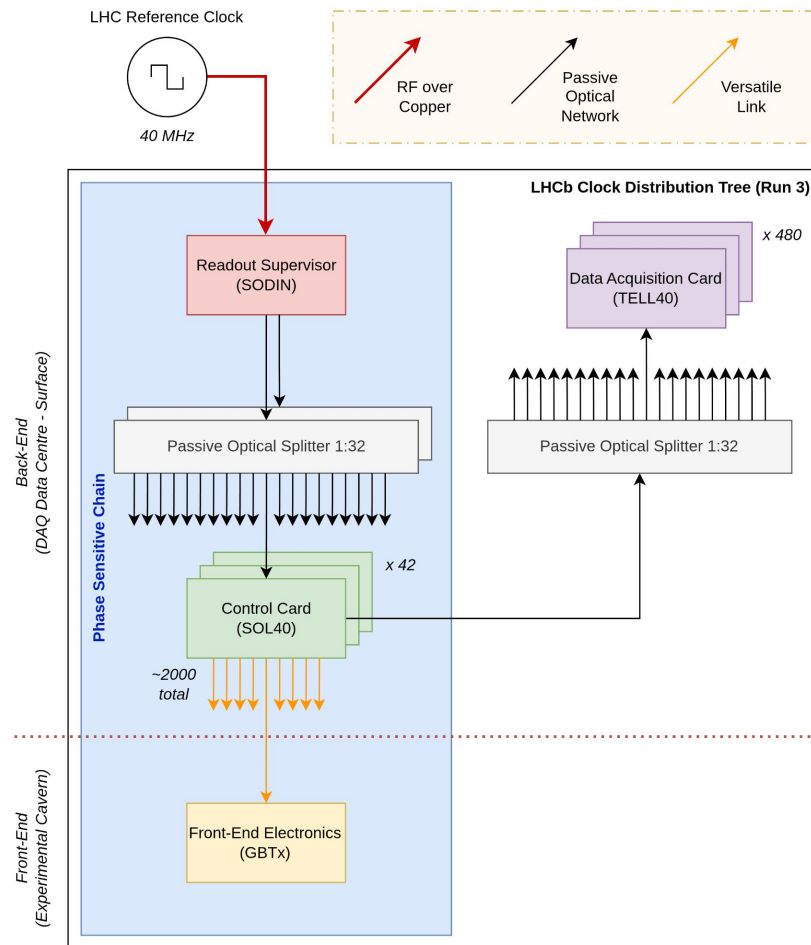
The TFC network uses PCIe40's bidirectional XGS-PON SFP+ modules.

The supervisor (SODIN) is the **entry point of the LHC clock** and **the source of the TFC network**.

Two PON splitters **distribute the generated TFC** commands to the SOL40 Control Cards.

The **SOL40 Cards** relay the TFC information to the FEE using the Versatile Links and to the TELL40 DAQ Cards using a second PON network.

The **phase sensitive part** of the network is the path towards the FE. Once the data is digitized and timestamped, the deterministic latency is no longer required.



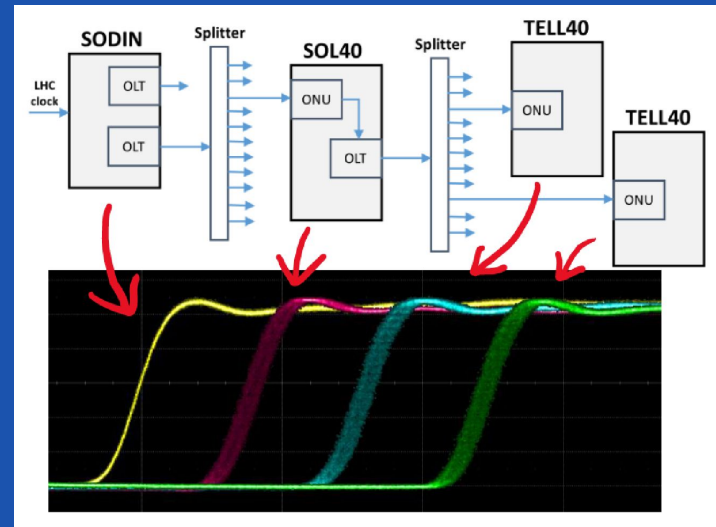
Clock Jitter and Skew Measurements

The current requirements of the TFC are:

- **Clock Skew:** Phase uncertainty of < 500 ps between the LHC bunch clock and each recovered clock
- **Clock Jitter:** Total RMS jitter < 100 ps of each clock
- **Bit Error Rate:** The links' BER must be $< 1.0 \times 10^{-15}$

Measurements done during commissioning report:

- Maximum clock skew of **254 ps** ($\sigma = 76$ ps)
- Random Jitter of **13.5 ps** (considering 76 ps as max)
- BER $< 3.1 \times 10^{-17}$ (95% C.L.)



Future Requirements and Current Limitations

The upcoming LHCb upgrades demand significantly tighter clock distribution requirements across sub-detectors.

To handle increased luminosity, multiple detectors (VeLo, RICH, ECAL) will implement **precision timestamping**:

- Run 4 (2030-2033): RICH requires **< 50 ps** phase uncertainty
- Run 5 (2036-2041): sub-detectors require **< 10 ps** phase uncertainty

Meeting these specifications requires a comprehensive understanding of current system limitations to design the future architecture.

This analysis identifies the constraints of the existing system and proposes solutions to **improve clock quality while reusing current hardware** with minor modifications.

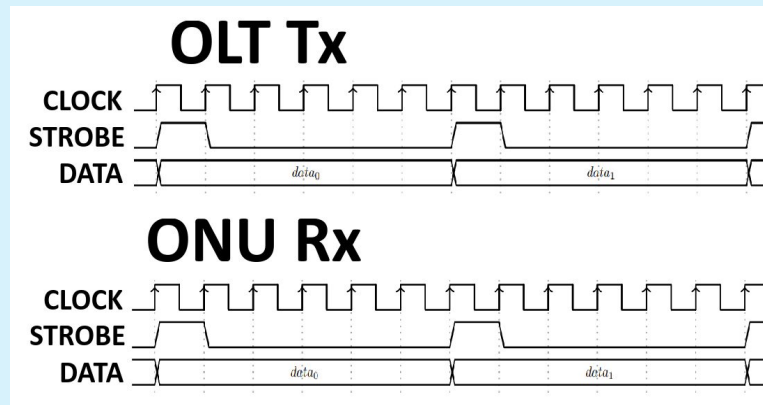
SOL40 Clock Architecture

The SOL40 control card gateway recovers the LHC bunch clock from the TFC stream via Clock and Data Recovery (CDR), which outputs a 240 MHz word clock and a 40 MHz strobe.

FPGA PLLs generate a common 40 MHz system clock from the CDR output. This system clock is used to drive all VL transceiver reference clocks. External PLLs clean the clock signal before distributing it to each transceiver bank.

This architecture introduces **metastability between transceiver channels** and the system clock due to PCB trace length variations, FPGA internal routing delays, and physical pin placement.

The current system **mitigates this by phase shifting** each transceiver bank using both internal and external PLLs to compensate for timing delays.



The source (OLT) synchronizes the data stream using the strobe signal.

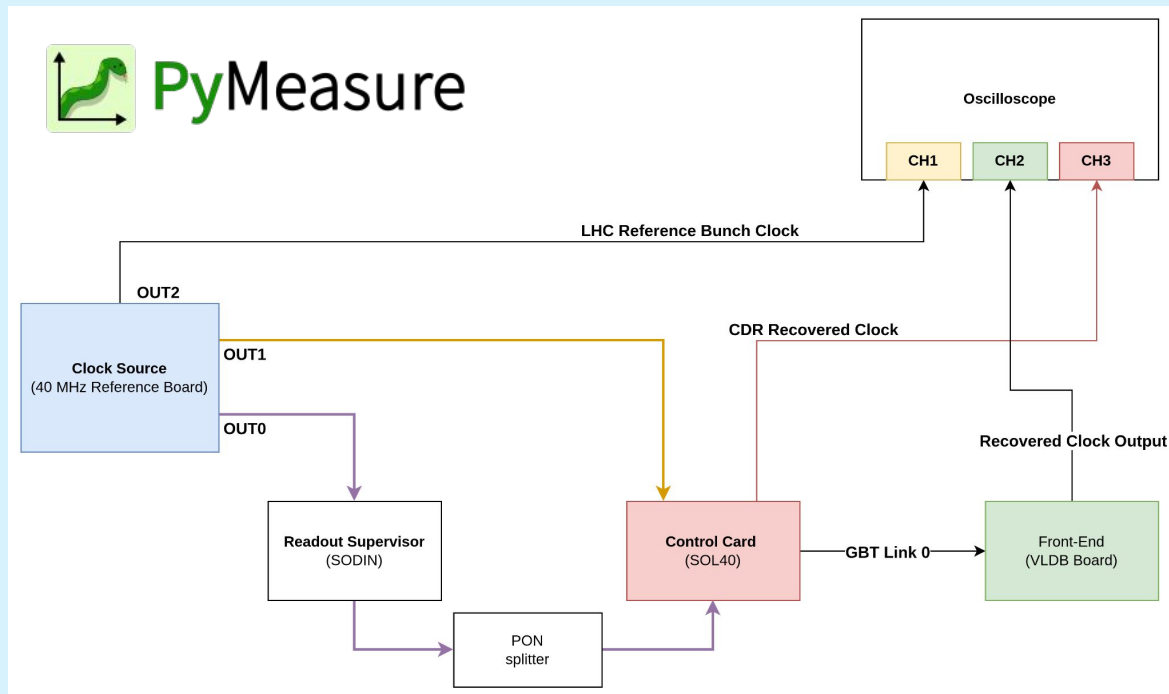
The sink (ONU) recovers the data stream and uses specific codes to recover the strobe signal.

Clock Measurement Testbed

A testbed replicating the clock architecture enables **systematic, automated measurements**.

A 40 MHz reference clock is distributed to both FPGA boards and an oscilloscope. Recovered clocks from the control card and FEE are also connected to the oscilloscope.

Automated measurement software toolkit based on pyMeasure streamlines batch data acquisition.



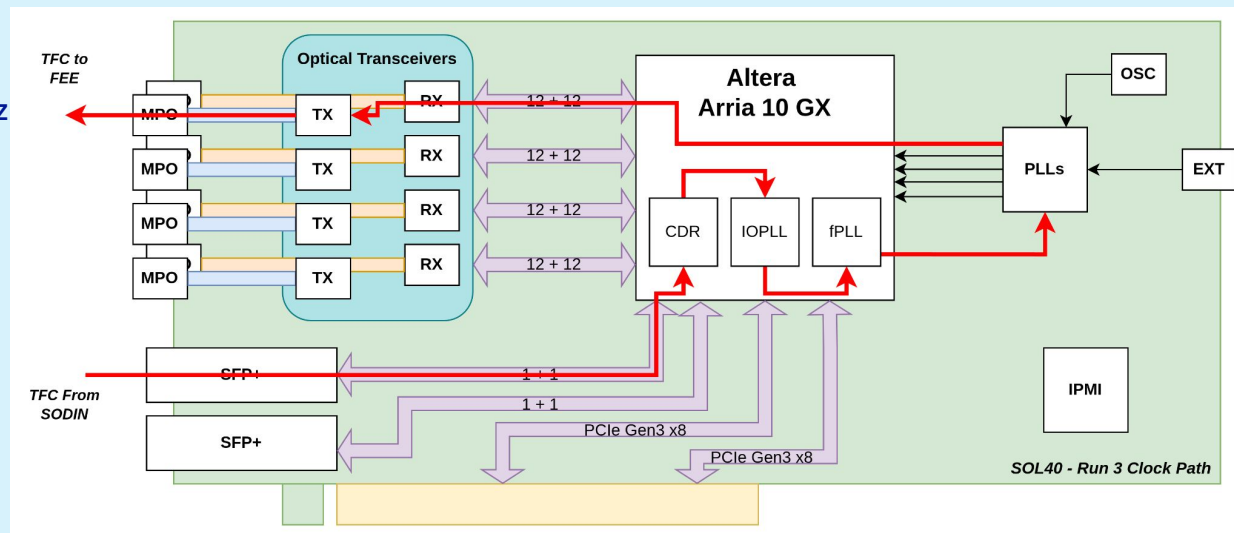
Critical Path Analysis

In the current gateway, the critical path is routed through:

- FPGA IOPLL to recover the 40 MHz from the recovered 240 MHz
- FPGA fPLL to control the phase shift of the global transceiver reference ($\Delta p \sim 100$ ps)
- External PLLs to control the phase shift of each transceiver bank ($\Delta p \sim 72$ ps)

Phase shifting is required to avoid metastability between multiple links.

In the worst case scenario, the theoretical maximum skew introduced by phase shifting is **244 ps**.



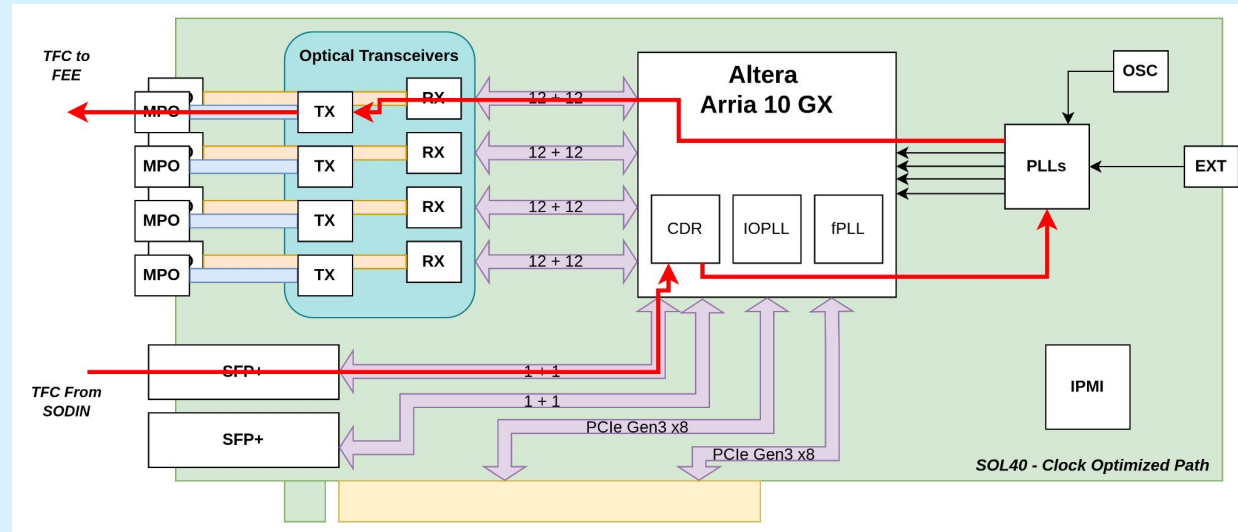
Direct-Path Clock Distribution

Improved the critical clock path by removing the main sources of phase uncertainty:

- **Direct clock** output from transceiver CDR
- **Bypass FPGA PLLs**
- External PLL used in **zero-delay mode**

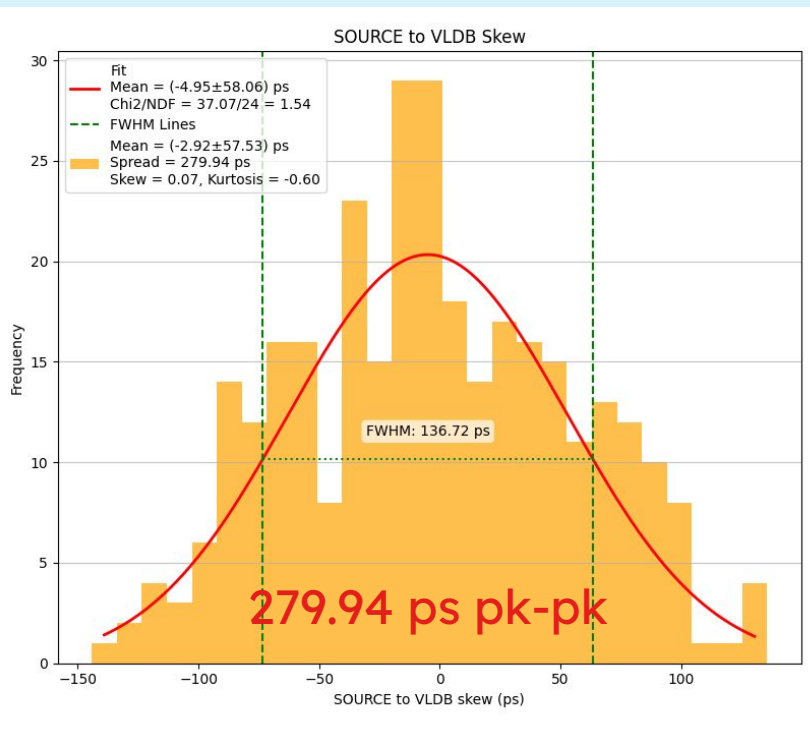
Clock is cleaner, but **no phase shifting** is possible, resulting in metastability.

This metastability is mitigated by shifting the system clock instead, affecting only coarse skew in FPGA fabric.

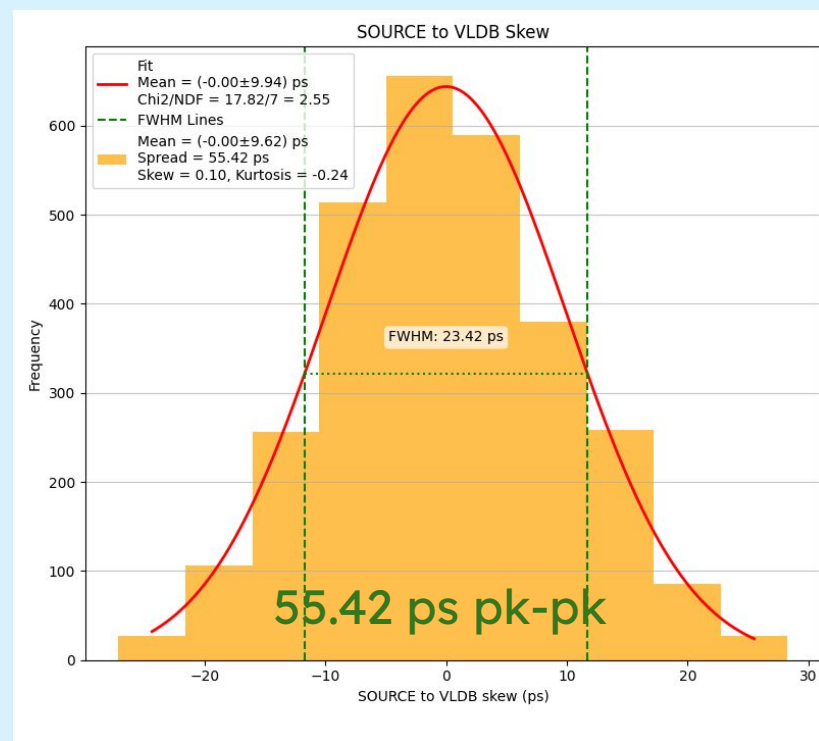


Source to Front End Skew Results

Original Clock Path



Optimized Clock Path



Thermal Effects on Phase Stability


The testbed includes multiple temperature sensors to **monitor thermal influence** during long measurement campaigns.

Temperature variations induce a **linear phase drift**:

- Original gateway: **~17 ps/K**
- Optimized gateway: **~10 ps/K**

Under operating conditions, this thermal effect can cause **> 50 ps** phase deviation over long (>5 hours) timescales.

These effects must be compensated to reach the high-precision timing requirements.

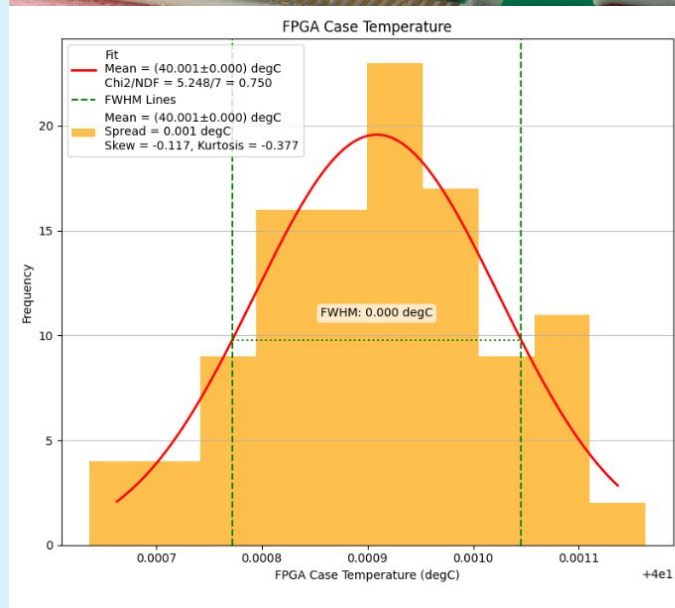



Current gateway temperature-induced phase drift

Active Temperature Compensation

Fine temperature control is achieved using a **thermoelectric cooler** with a high-precision temperature probe mounted directly on the FPGA heat spreader.

The PID control loop maintains **long-term stability of ~1 mK**, effectively eliminating thermal effects on the FPGA fabric.

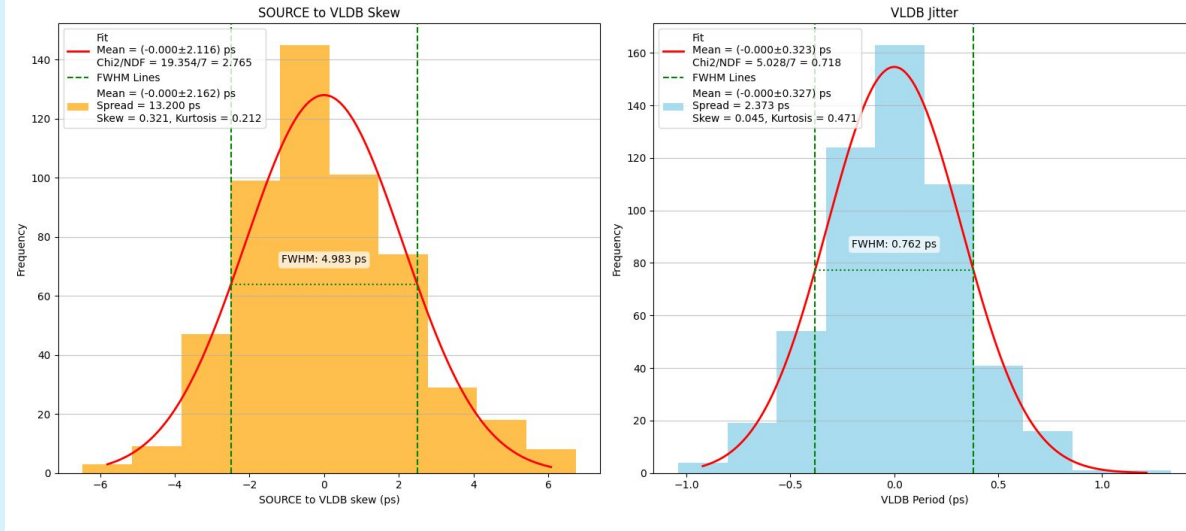


Performance Results

Active temperature compensation further reduces phase uncertainty, achieving **13.2 ps** pk-pk over multiple days, **21x better** than the original configuration.

These results validate the improved stability and robustness of the new clock architecture.

Measurements Recap Histograms



Original	Optimized	Optimized with active temperature control
280 ps	55.4 ps	13.2 ps

Next Steps and Outlook

The automated testbed provides a generic platform for validating future timing components and clock architectures.

The results are promising and suggest the control FPGA boards **may not require replacement** for Run 4. However, the studies presented focus on a single link on a single FPGA under laboratory conditions.

To qualify these results for the LHCb Upgrade, **testing must be extended to larger testbeds** and ultimately the full-scale installation. These studies will become possible during the LHC Long Shutdown beginning in July 2026.

Thanks for your attention

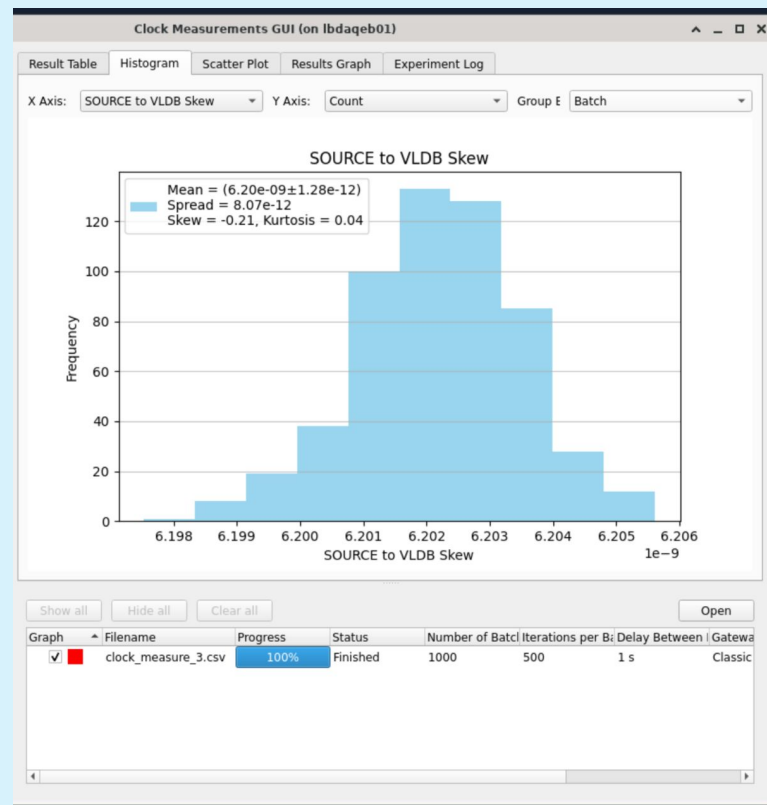
Automated Lab Measurements

Long test campaigns required an automated system to scan all variables and configurations.

A measurement toolkit based on pyMeasure was developed.

It features:

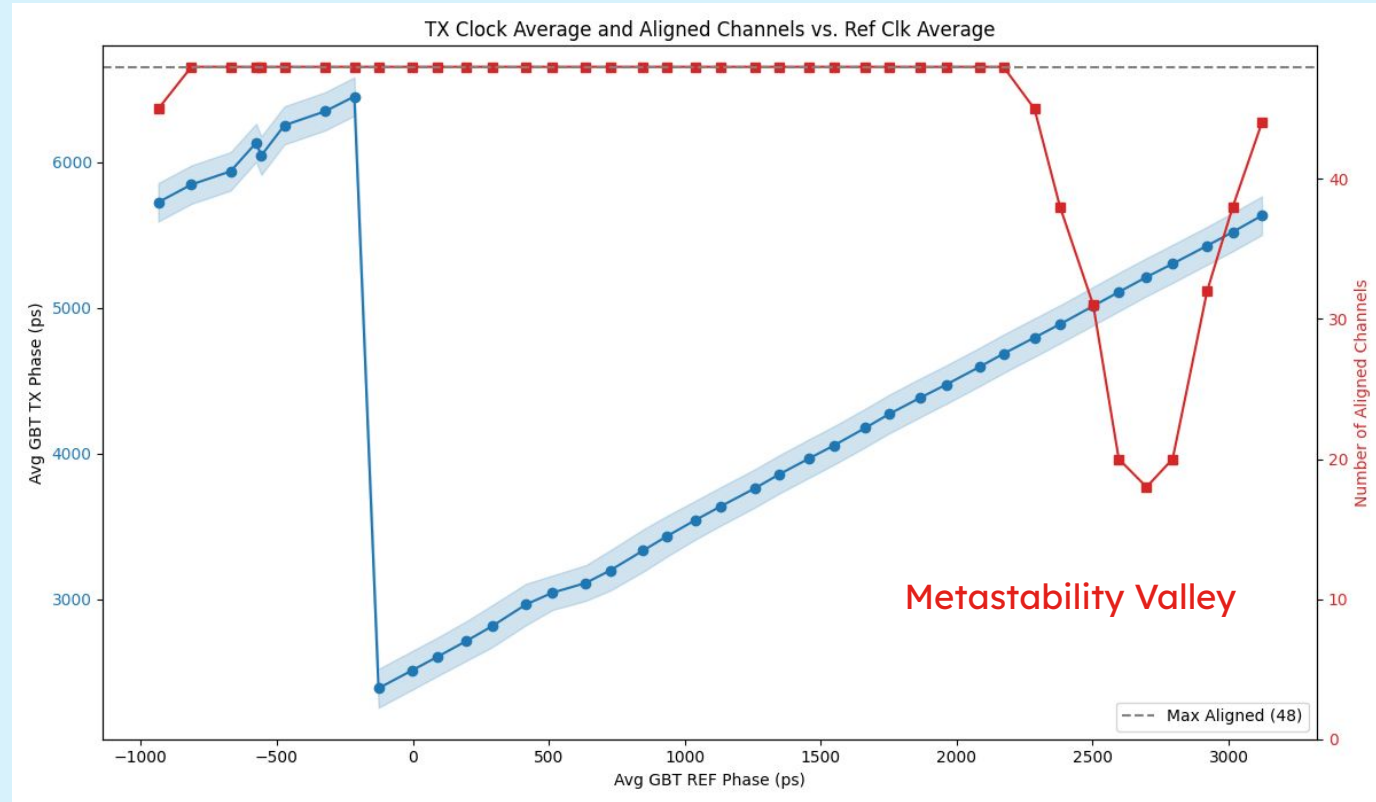
- Automatic Job Batching
- Test Configuration from GUI
- Real-Time Plotting
- Test Procedure Logging
- Full integration of instrumentation (scope, FPGA, temperature, ...)



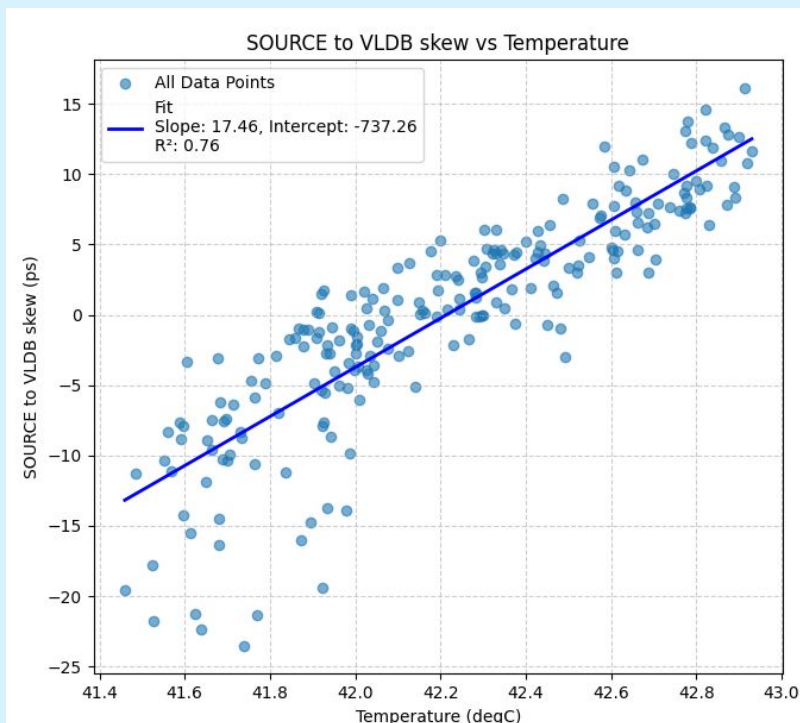
Metastability between Channels

Reference Clock is shifted by the FPGA PLLs.


There is a *metastability valley* where most transmitter cores are not capable of maintaining phase alignment.



Temperature Dependencies



Original Clock Path



Optimized Clock Path