Karlsruhe Institute of Technology

Institute for Information Processing Technologies

Evaluation of GBT-FPGA for Timing and Fast Control in the CBM experiment

SPONSORED BY THE

Federal Ministry of Education and Research

V. Sidorenko, W.F.J. Müller, W. Zabolotny, I. Fröhlich, D. Emschermann, J. Becker

On behalf of CBM collaboration

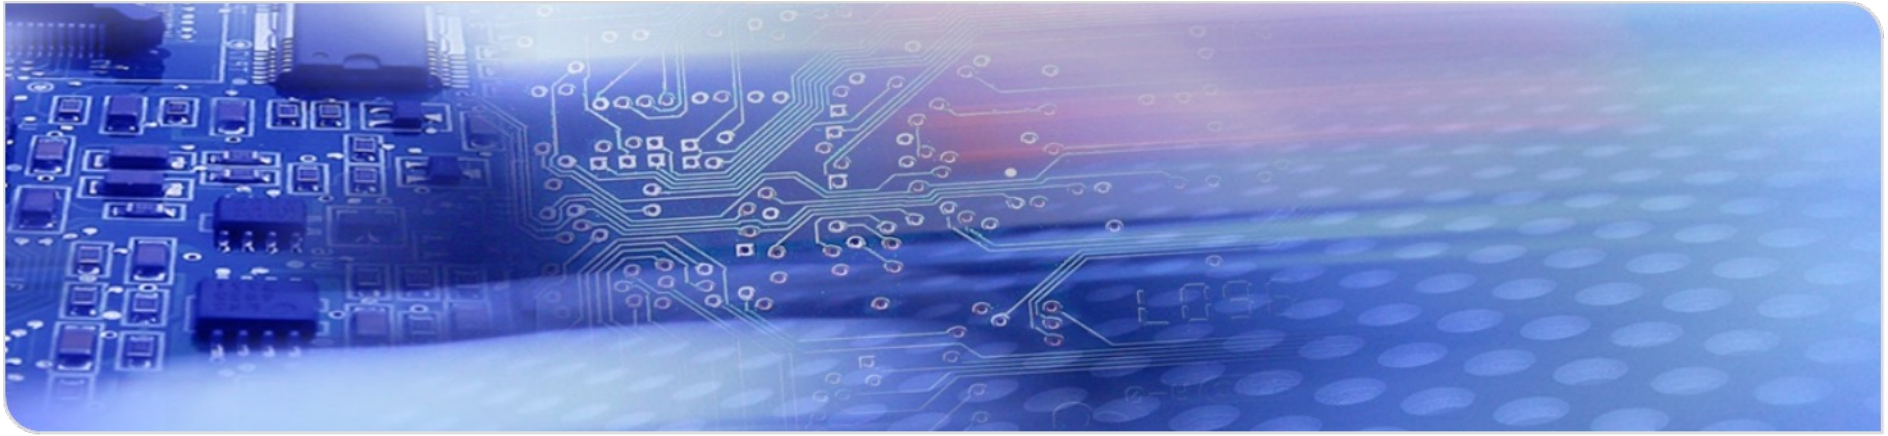

Outline

- Background
- Existing CBM TFC prototype
- Why switch to GBT-FPGA?
- Evaluation
- Summary

Background

CBM experiment

- Peak R_{int} is 10 MHz for Au+Au
- Fast & radiation hard detectors
- 4D tracking (space, time)

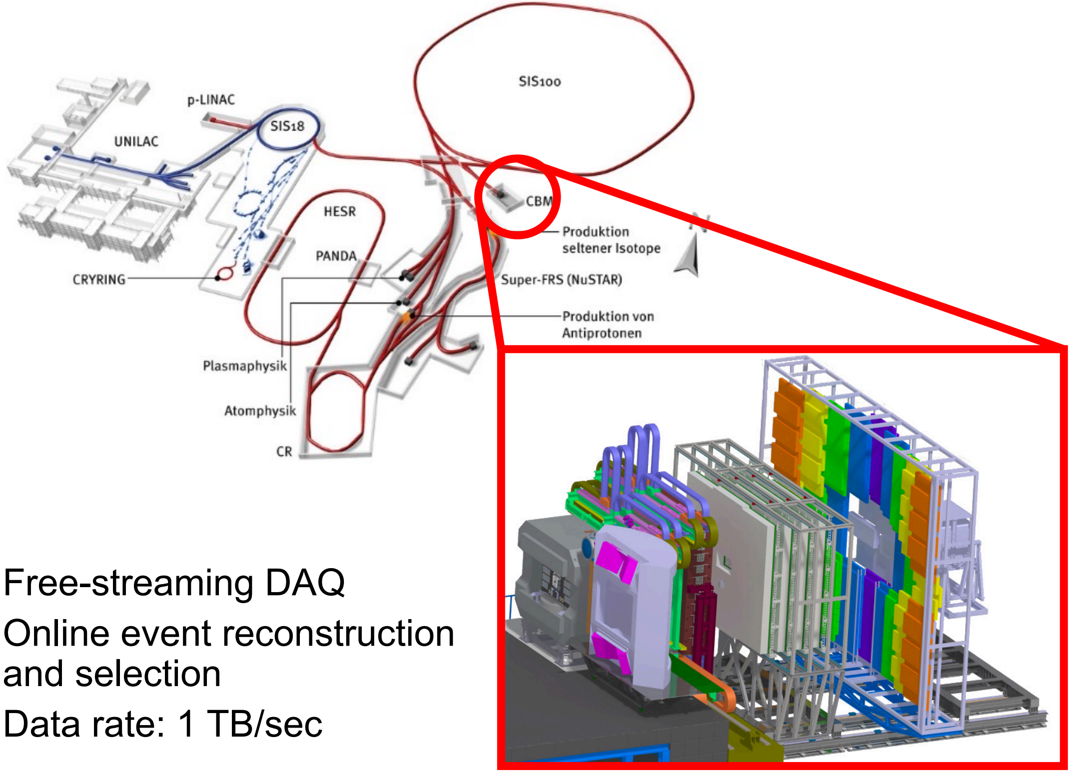

Photo as of May 2022
Source: https://www.gsi.de/forschungbeschleuniger/fair/bau_von_fair/bilder_und_videos

- Free-streaming DAQ
- Online event reconstruction and selection
- Data rate: 1 TB/sec

Timing and Fast Control

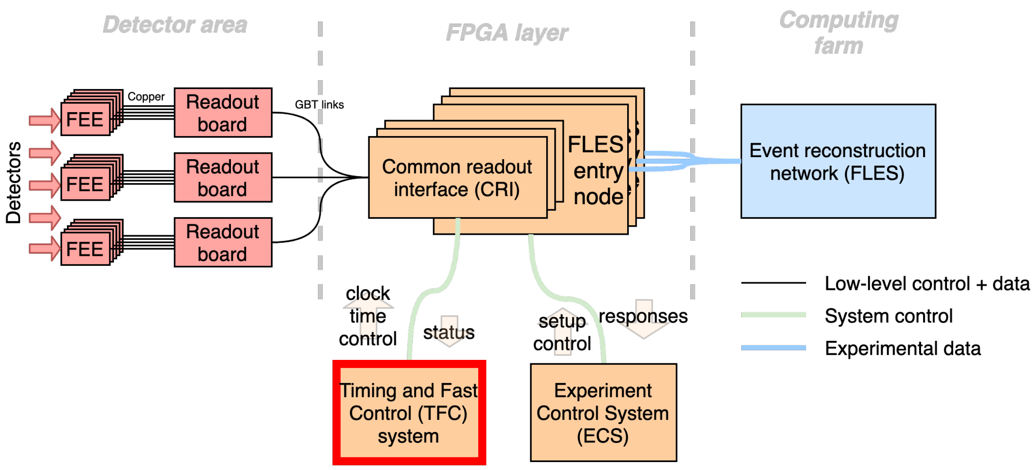

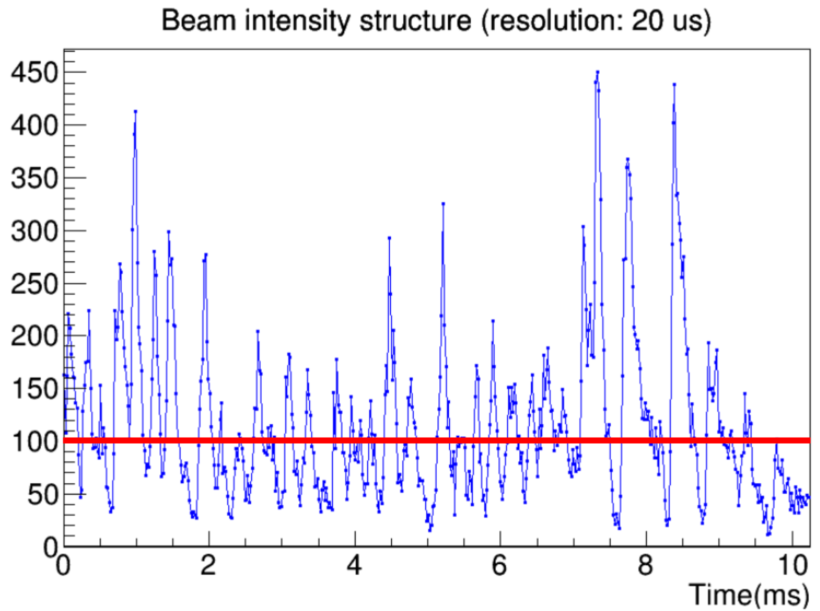

Throttling (study done by X. Gao):

- Versatile fast control network required

Source: X. Gao, D. Emschermann, J. Lehnert, and W. F. J. Müller, “Throttling Studies for the CBM Self-triggered Readout,” presented at the Topical Workshop on Electronics for Particle Physics, Mar. 2020. doi: 10.22323/1.370.0085.

Existing CBM TFC prototype

CBM TFC prototype now

System requirements:

- Scalability to serve ~250 endpoints
- < 200 ps synchronization accuracy
- < 10 µs fast control response time

Hardware platform:

- BNL-712
- Timing Mezzanine Card

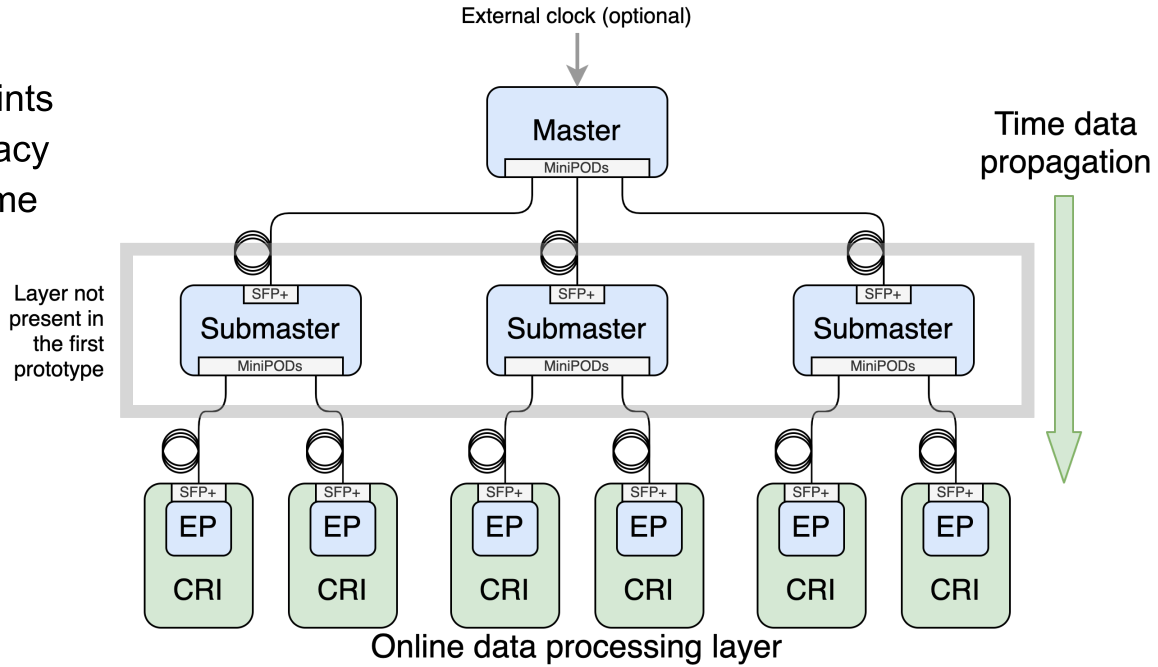

CBM TFC prototype now

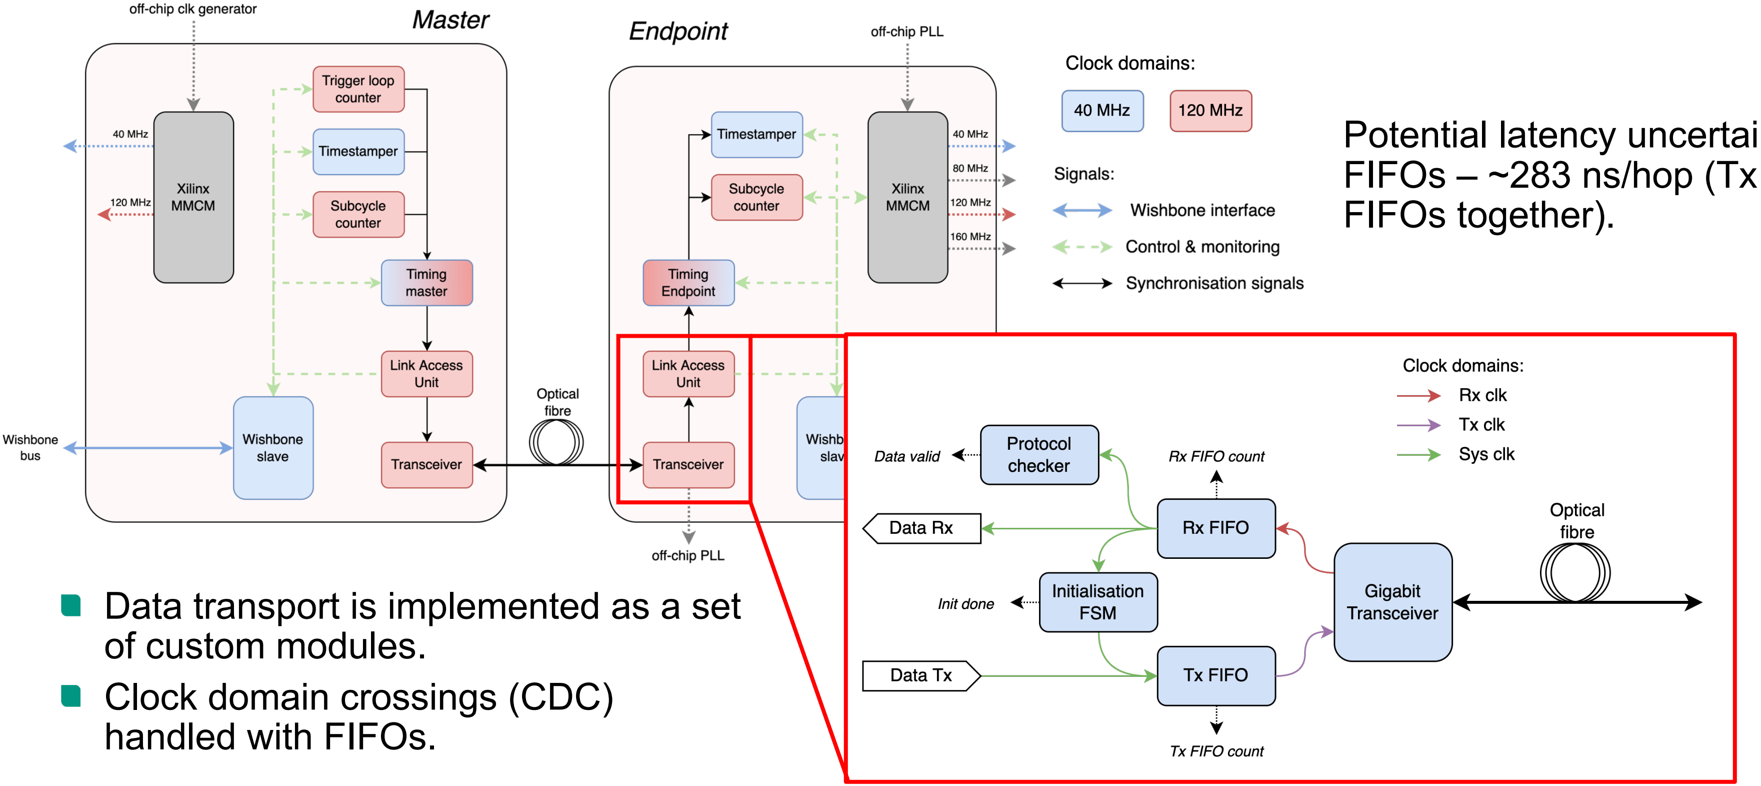

Potential latency uncertainty on FIFOs – ~283 ns/hop (Tx and Rx FIFOs together).

- Data transport is implemented as a set of custom modules.
- Clock domain crossings (CDC) handled with FIFOs.

Why switch to GBT-FPGA?

Why switch to GBT-FPGA?

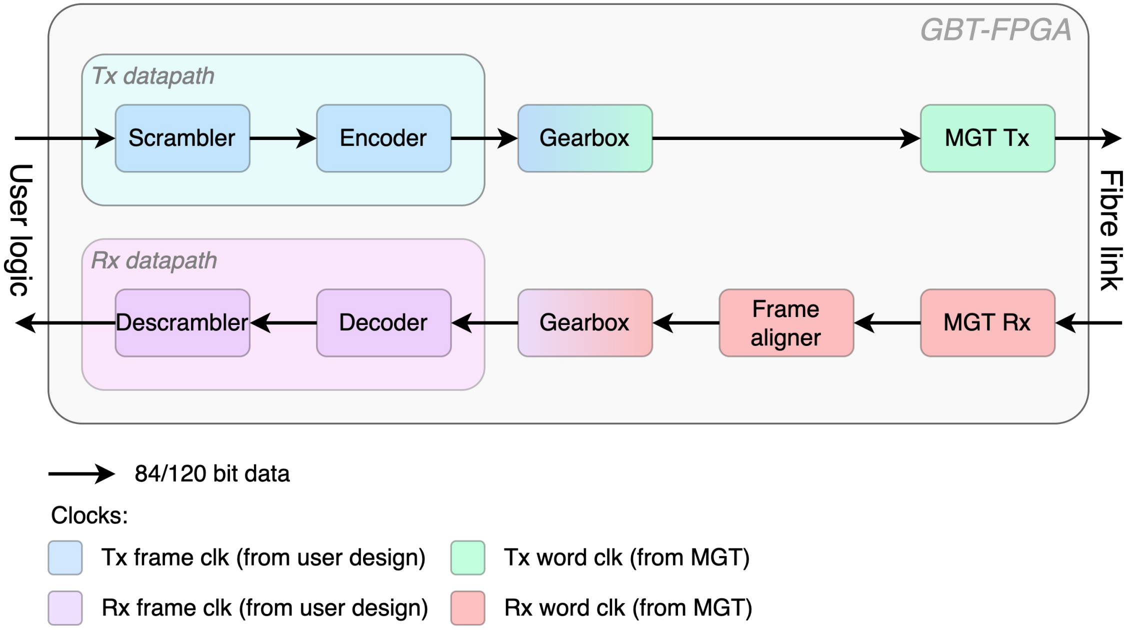

GBT-FPGA core highlights:

- Originally developed at CERN
- Maintained in CBM by Marek Guminski (WUT, Warsaw)
- Latency-optimized flavor with deterministic register-based CDCs
- Forward error correction

Why switch to GBT-FPGA?

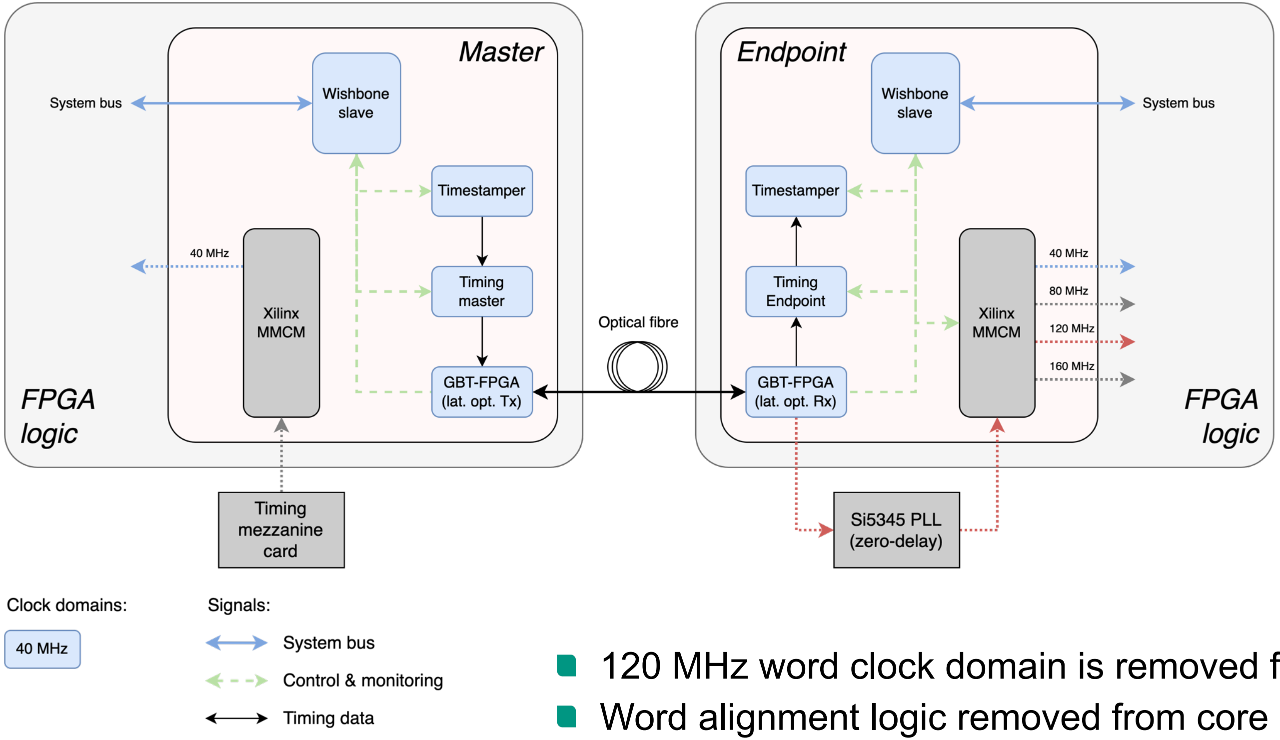

- 120 MHz word clock domain is removed from the TFC layer
- Word alignment logic removed from core design

Evaluation

Evaluation setup

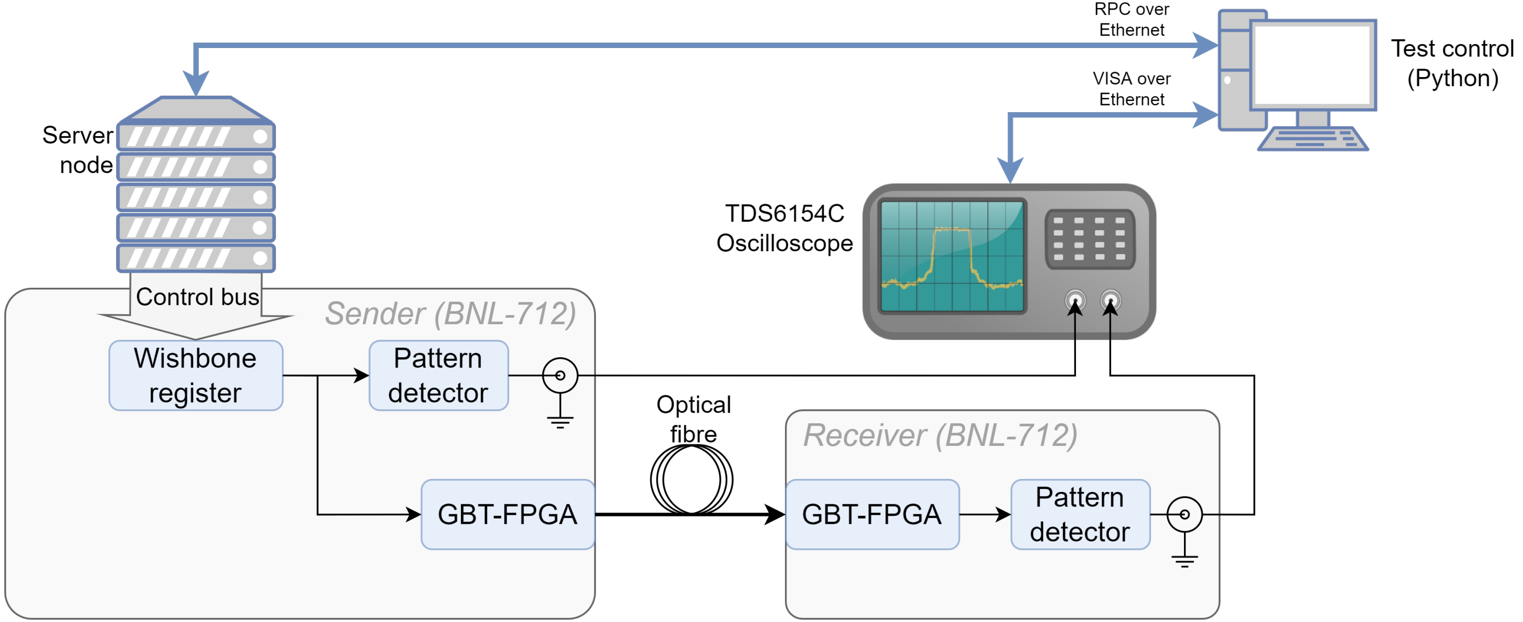

- Measurement speed could be better:
  - collecting 1k samples takes ~70 min
  - ~11 hours for 10 runs
- Fully automated setup

Evaluation procedure

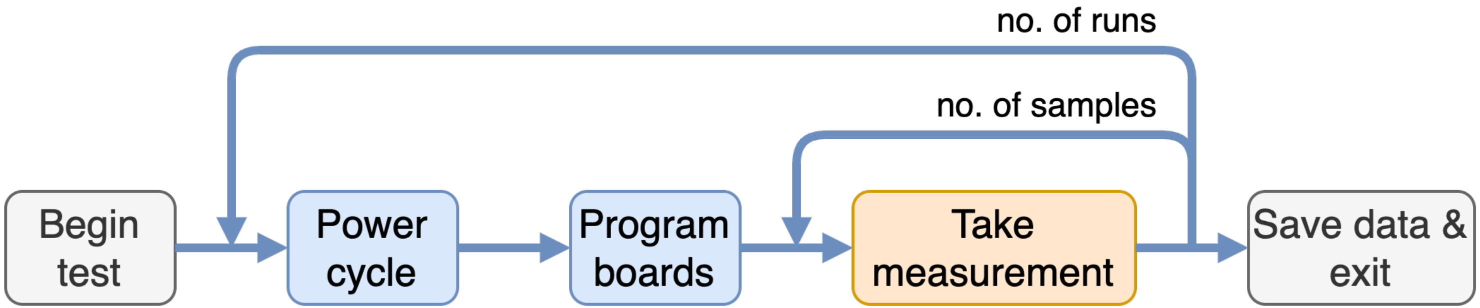

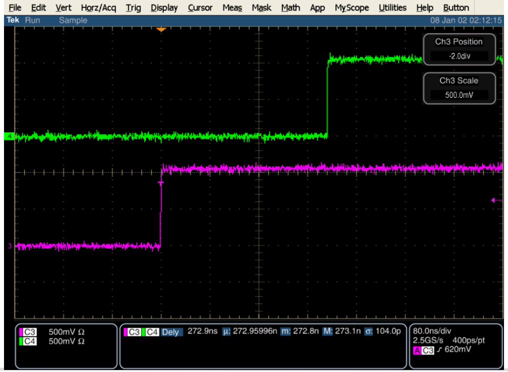

An example of a measurement (sample)

- *Sample* – one instance of latency measurement on the link.
- *Run* – a set of consecutive samples.

Evaluation conditions

- 10 runs, 1000 samples each.
- Direct optical connection from Master to Endpoint.
- Link implementations under test:
  - Existing prototype (no latency optimizations, PLL zero-delay off),
  - Latency-optimized GBT-FPGA (PLL zero-delay on).

Evaluation – existing setup

Latency distribution: 10 runs by 1k samples, boards are power cycled and reprogrammed between runs

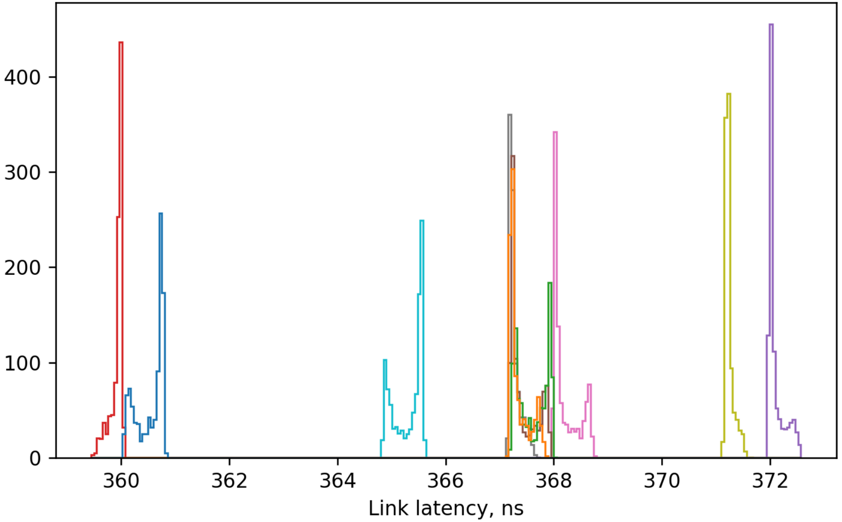

Worst run

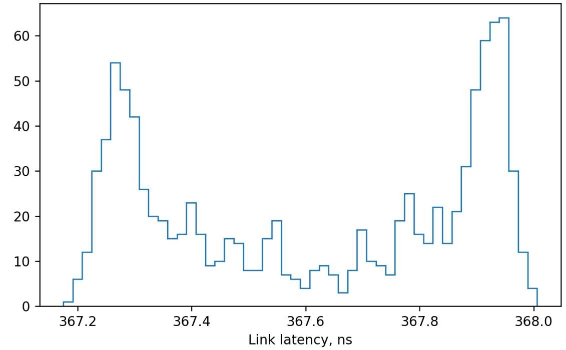

Best run

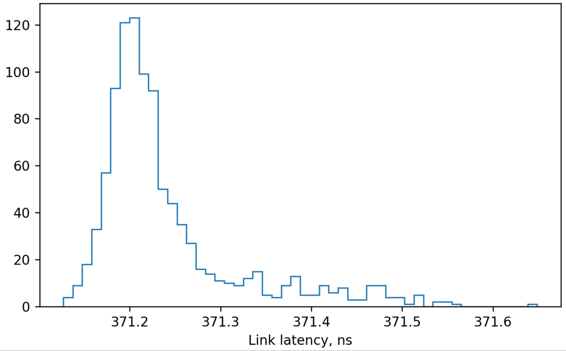

Current TFC prototype (no latency optimizations, PLL zero-delay off):

- Very poor reset-to-reset determinism (> 10 ns variation)
- Latency distributed over up to 1 ns

Evaluation – GBT-FPGA link

Latency distribution: 10 runs by 1k samples, boards are power cycled and reprogrammed between runs

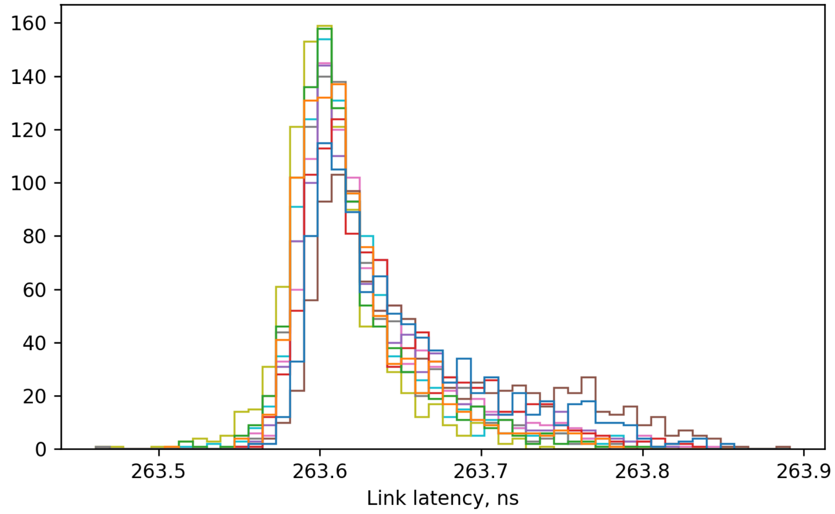

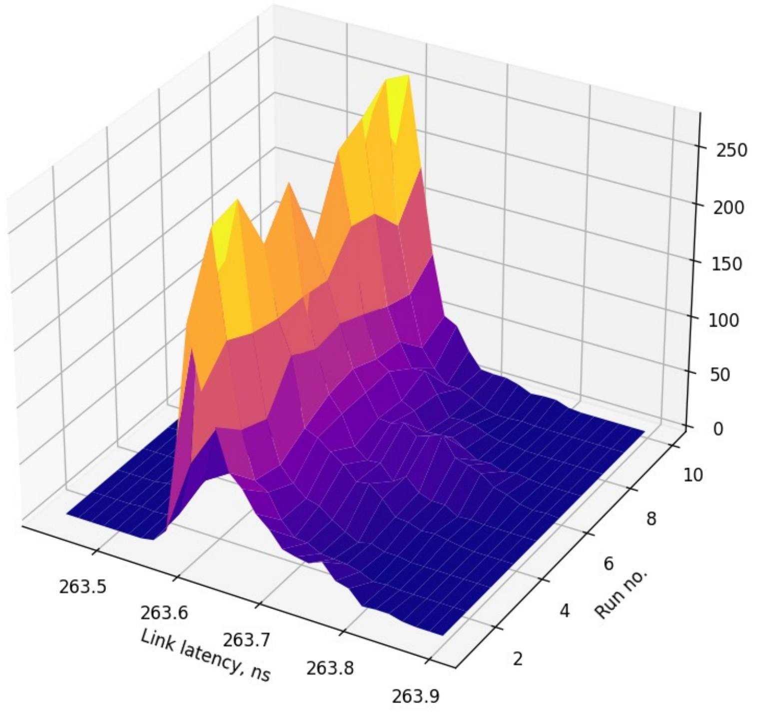

Latency-optimized GBT-FPGA link (PLL zero-delay on):

- Excellent reset-to-reset determinism (< 100 ps variation)
- Improved latency distribution (within 500 ps p-p)

Evaluation – comparison

Latency mode, mean and std values over 10 runs

Existing prototype

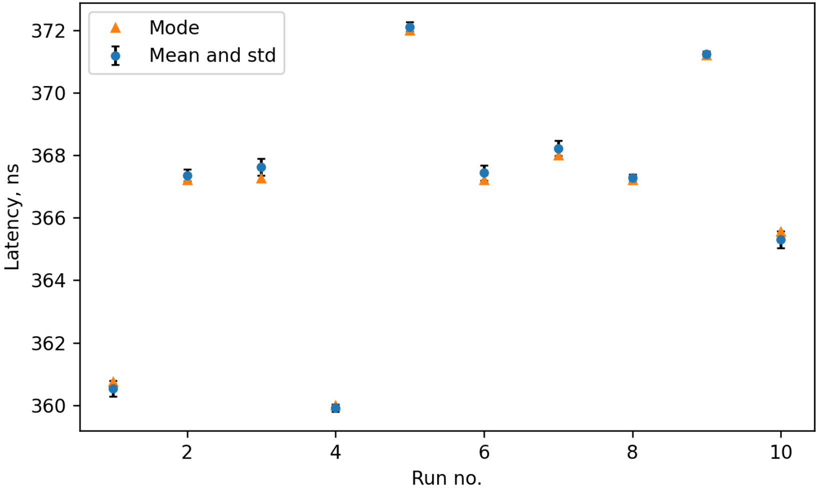

Latency optimized GBT-FPGA

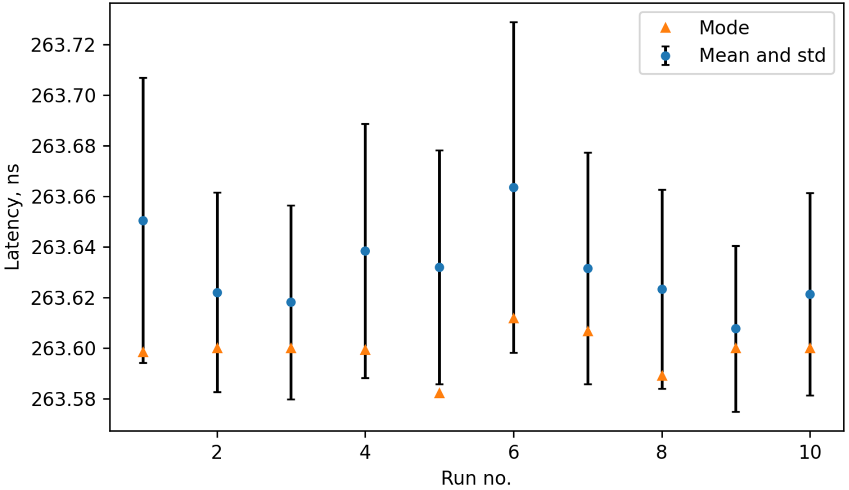

- Absolute latency improved by ~100 ns
- Reset-to-reset determinism improved by a factor of ~200
- Latency distribution within one run improved by a factor of ~2
- In-run variation is much more pronounced within runs than between resets

Summary

- Usage of GBT-FPGA for data transport:
  - Reduces total link latency by ~100 ns
  - Improves in-run determinism by ~2 times reset-to-reset determinism by a factor of ~200
  - Keeps link latency safely within a ~500 ps range
  - Simplifies gateware design by handling low-level transceiver management
- Automated evaluation setup allows for easy link latency characterization

- New challenge: observed in-run link latency variation

Thank you!