l'Observatoire de Paris — LESIA
Laboratoire d'Études Spatiales et d'Instrumentation en Astrophysique

Durham University

MICROGATE

PLDA

Green Flash

High performance computing for real-time science

Contribution from Observatoire de Paris on final prototypes design Final Design Review, April 6th 2018

European Commission

Project #671662 funded by European Commission under program H2020-EU.1.2.2 coordinated in H2020-FETHPC-2014

Full scale protyping @ OdP

OdP team responsible for full scale prototype based on GPU

- From SCAO to LTAO/MCAO
- Includes HRTC, SRTC and ComIF
- Several simulator strategy: hard real time HW simulator and slower frame-rate end-to-end simulator

Main contributions

- Consistent SW stack architecture across subsystems
- Identify and use optimized libraries in dedicated pipelines
- Optimized SW blocks for HRTC and SRTC
- Assemble HRTC+SRTC subsystems at full scale

SCAO prototype

- General architecture
  - RTC subsystem : no supervisor module

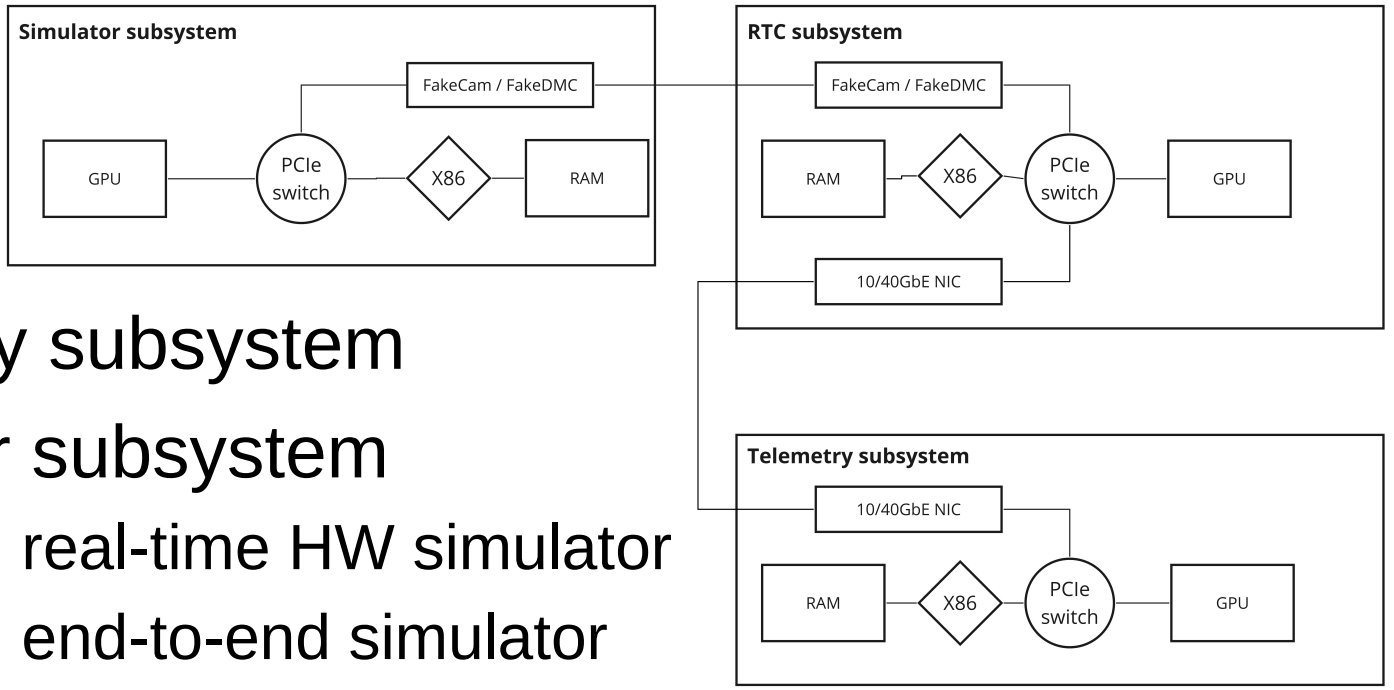

  - Telemetry subsystem
  - Simulator subsystem
    - With real-time HW simulator
    - With end-to-end simulator

SCAO prototype

- Relying on small servers: allow for quick debugging
  - Top: RTC subsystem
    - Hosting a K40 GPU
    - Hosting a smart NIC (Altera Stratix V or Xilinx KUS)
  - Bottom: telemetry subsystem
    - Hosting either a standard NIC or a smart NIC

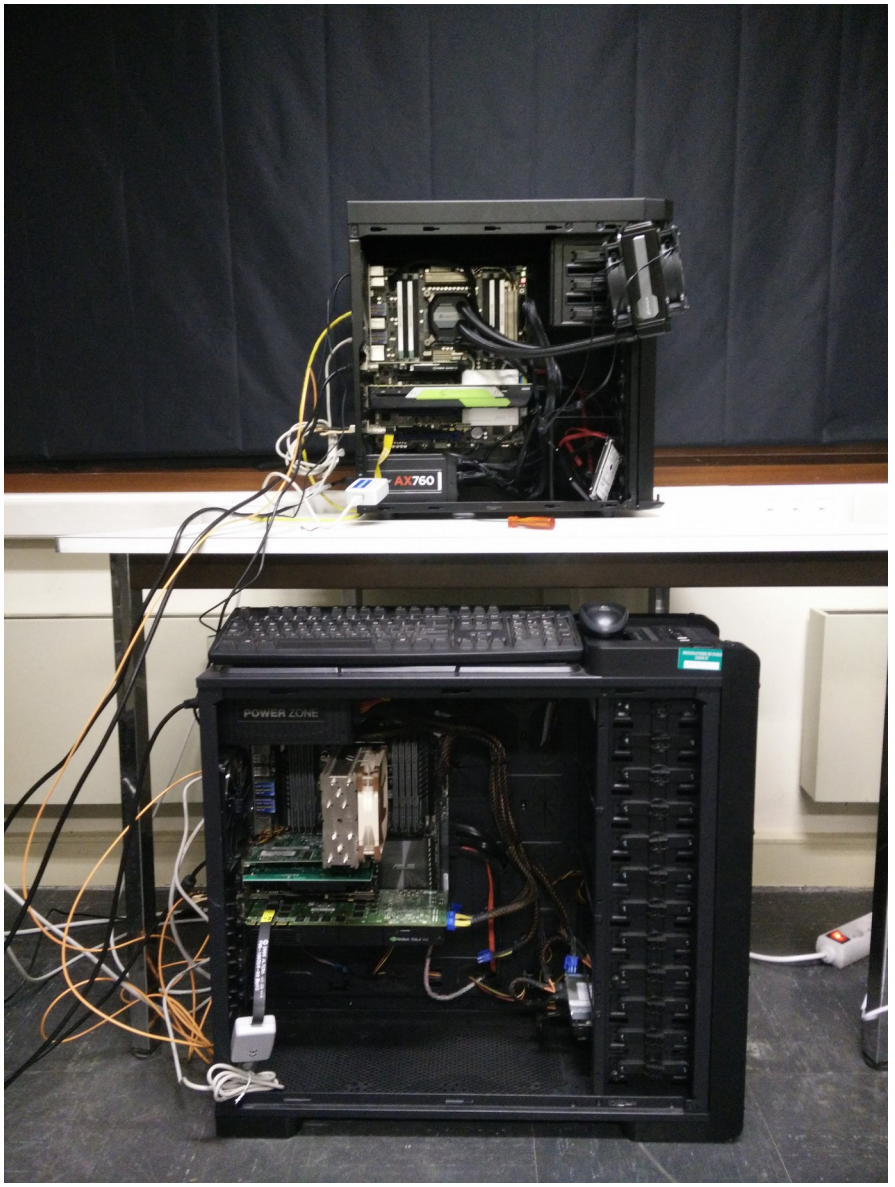

SCAO prototype: datapaths

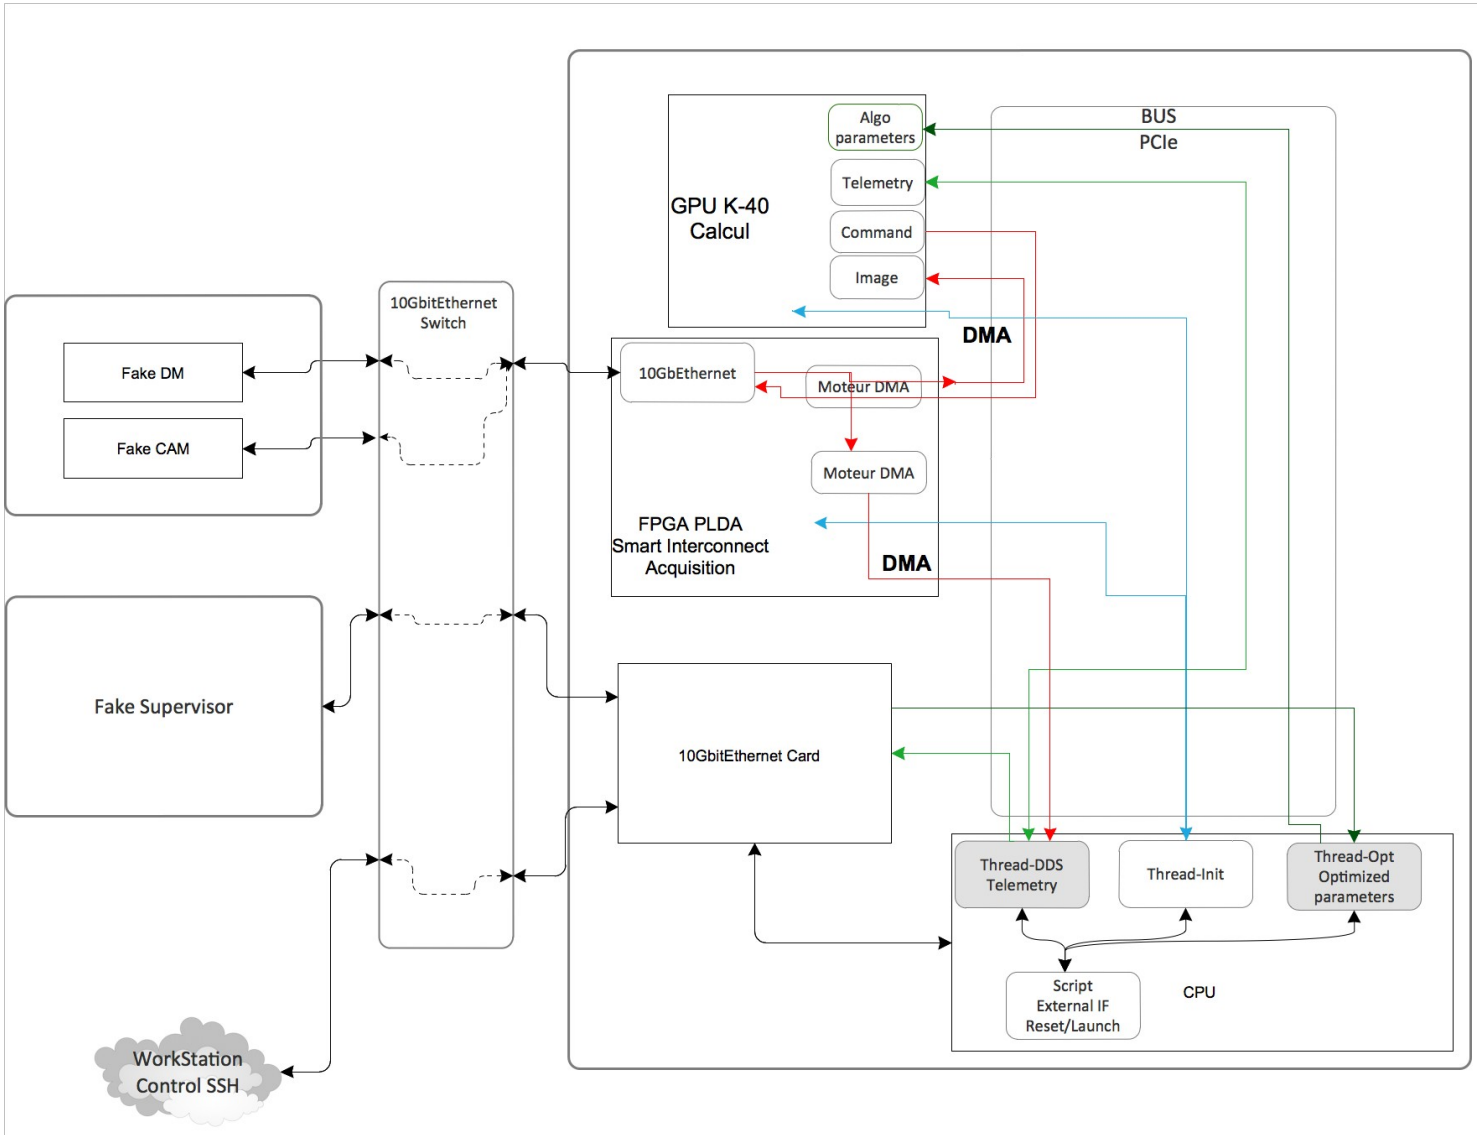

SCAO prototype

- Preliminary performance on an (old) K40 GPU + smart NIC
  - First pixel to last command latency on GPU (blue)
  - First pixel to first command latency on FPGA (orange)

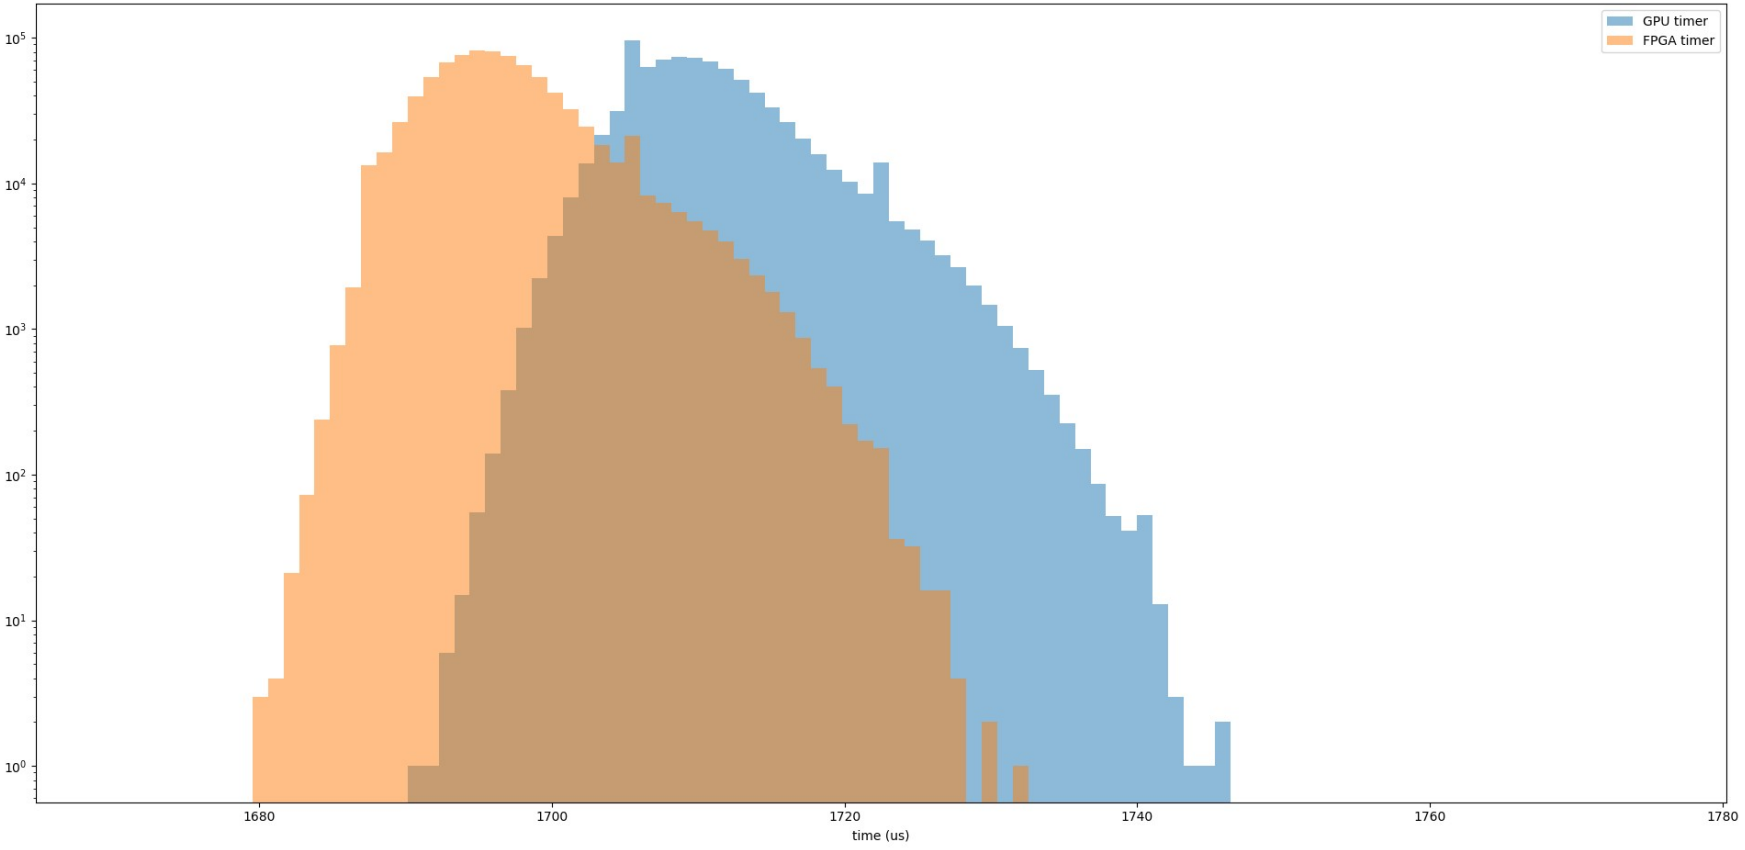

SCAO prototype

- Preliminary performance
  - Time difference between FPGA timings and GPU timings
  - Performance of a single GPU deterministic at the level of 10-20µs

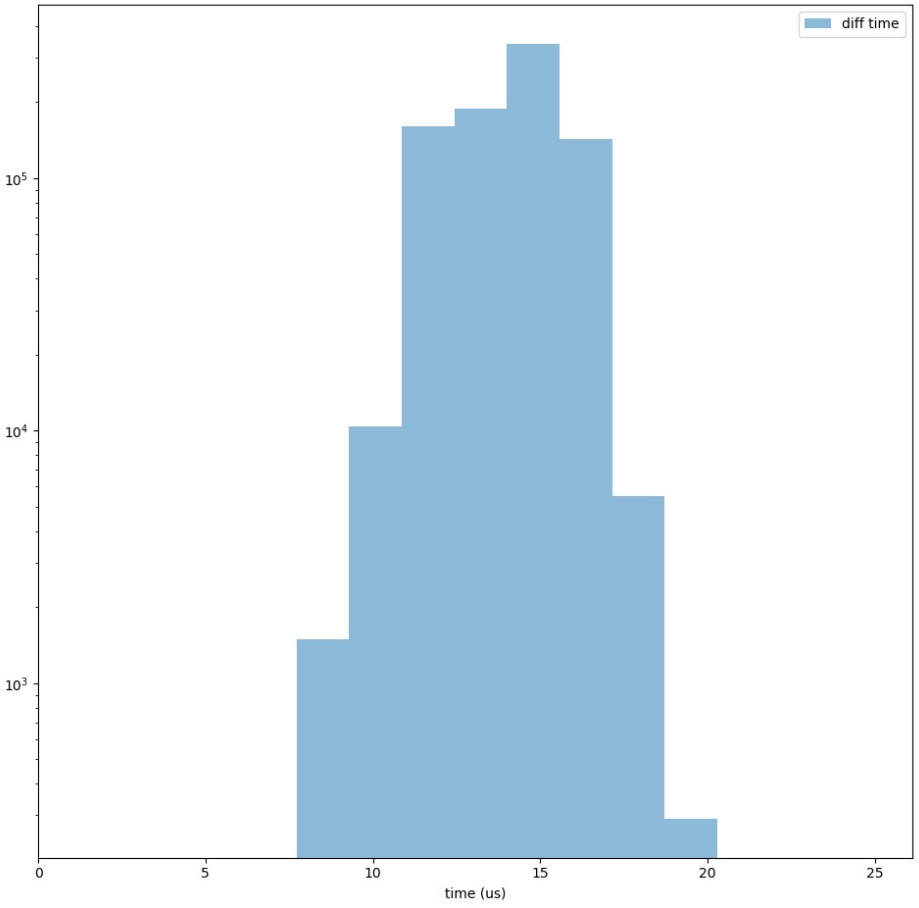

SCAO prototype

- Preliminary performance: transient events
  - Appears to be repeatable
  - Room for investigation

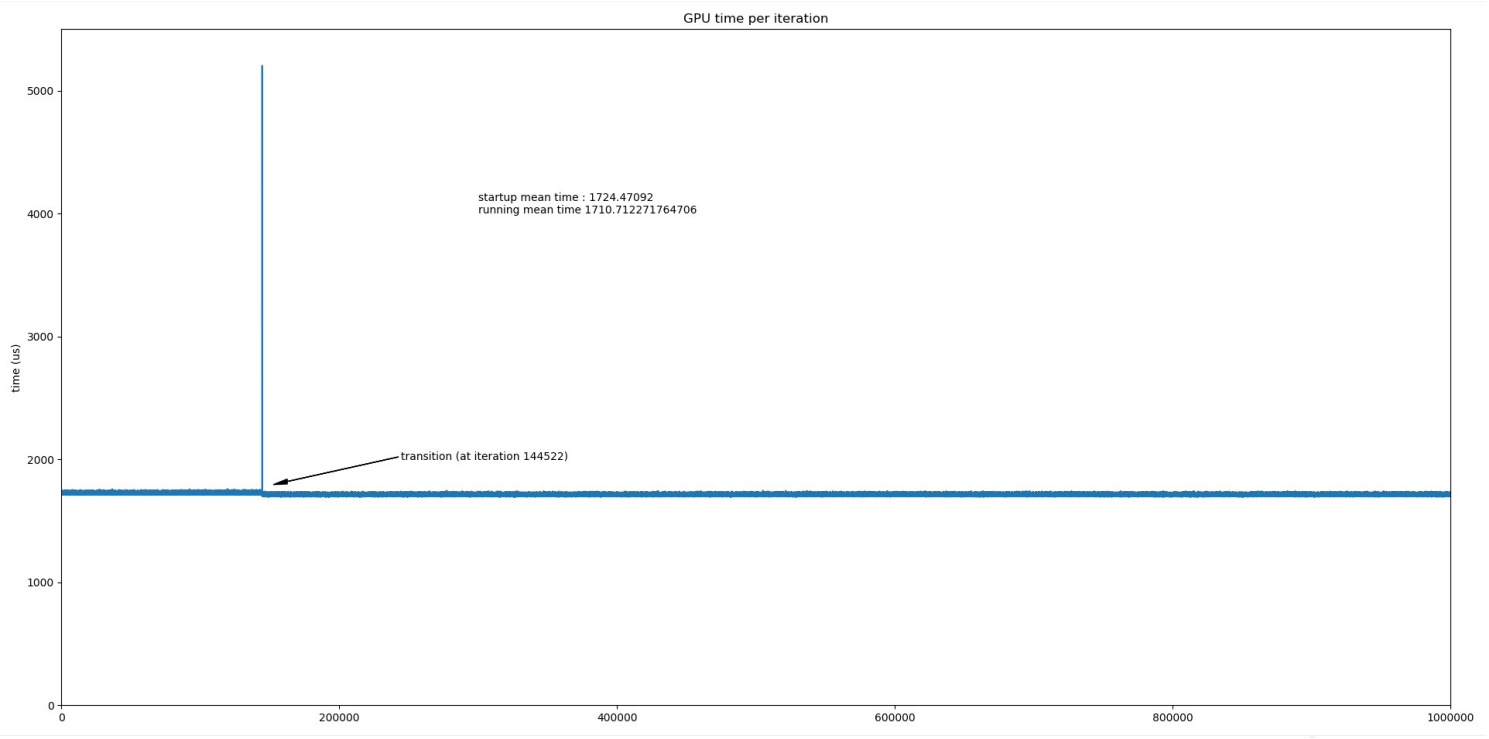

LTAO/MCAO prototype

- Several concepts
  - Full HRTC + SRTC with HW sim.

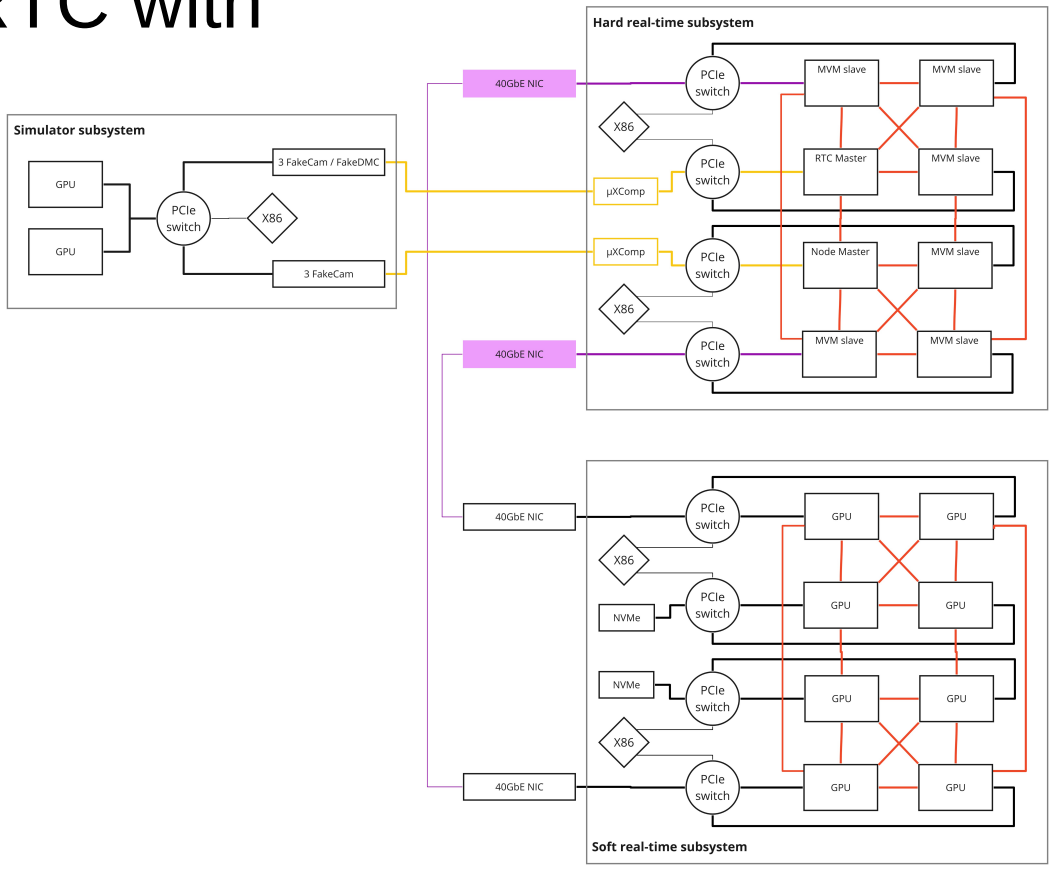

LTAO/MCAO prototype

- Several concepts
  - Full HR HW si
  - HRTC + end-to-end sim. + telemetry subsystem

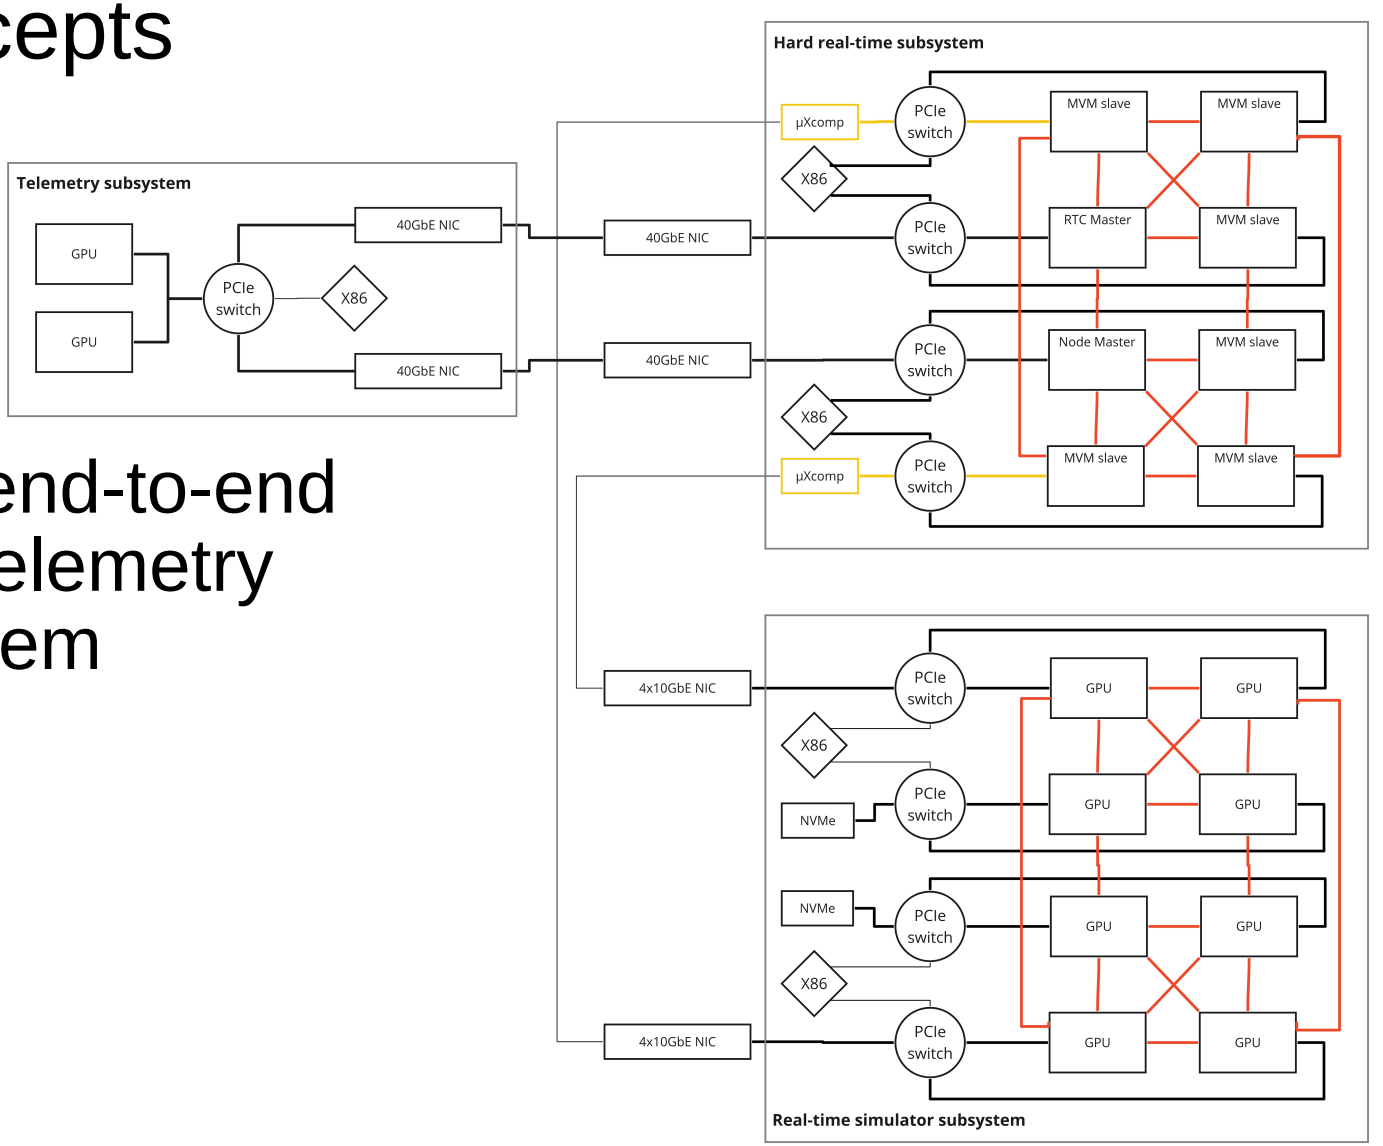

LTAO/MCAO prototype

- Detailed design for the HRTC (data + paths)

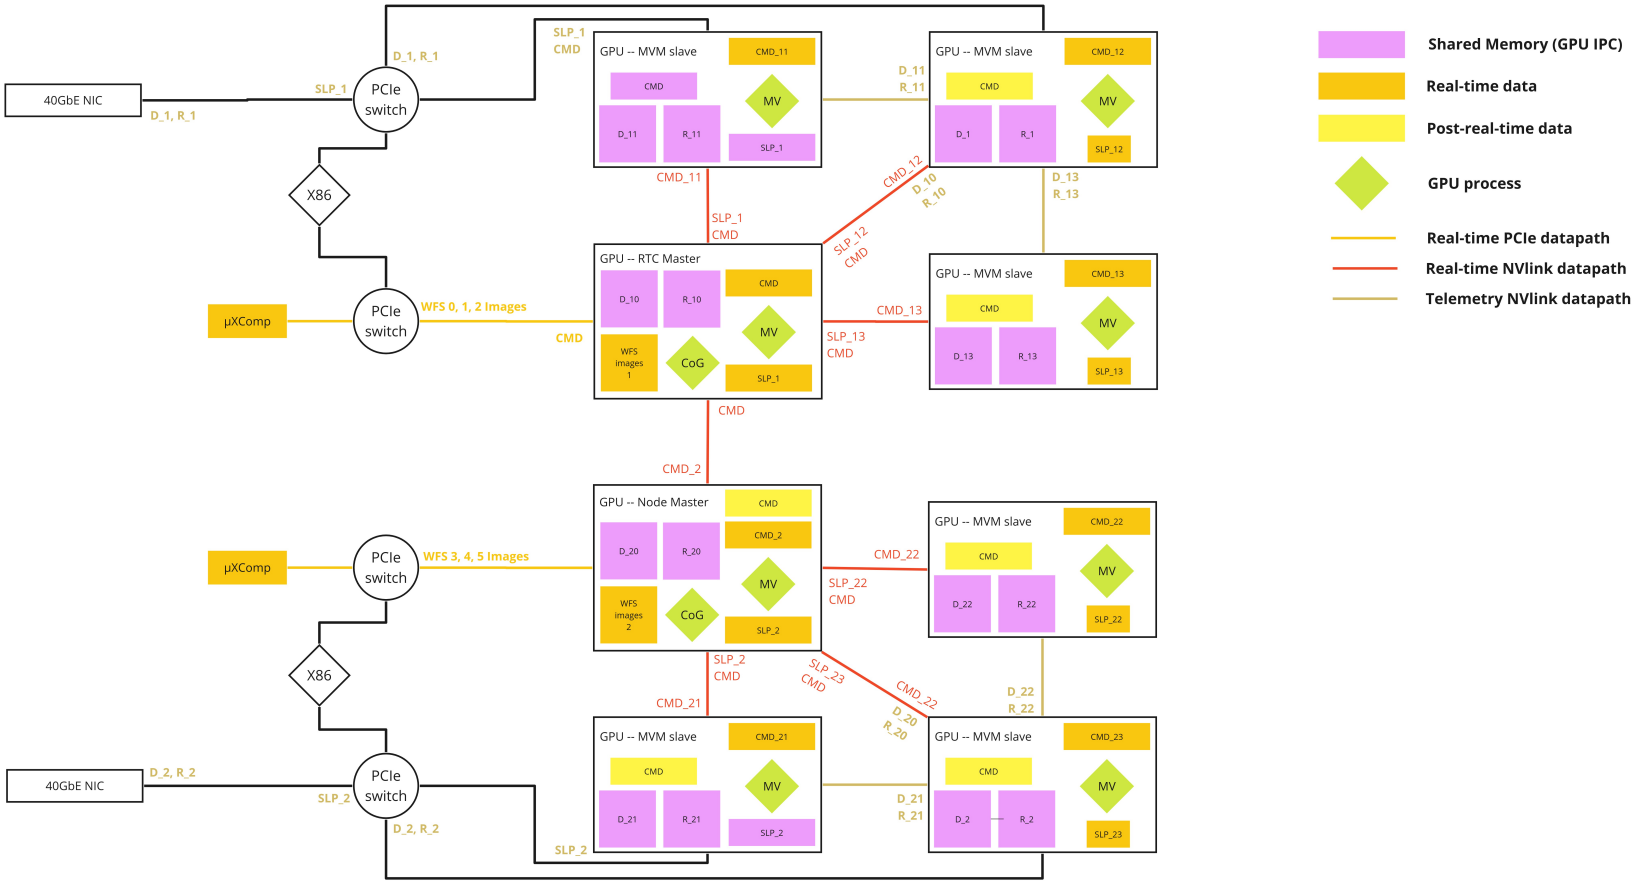

SW stack prototype

- General concept
  - Supports CPU and GPU threads
  - Relies on several levels of Shared Mem.
  - Relies on standard libraries
  - Open Source !

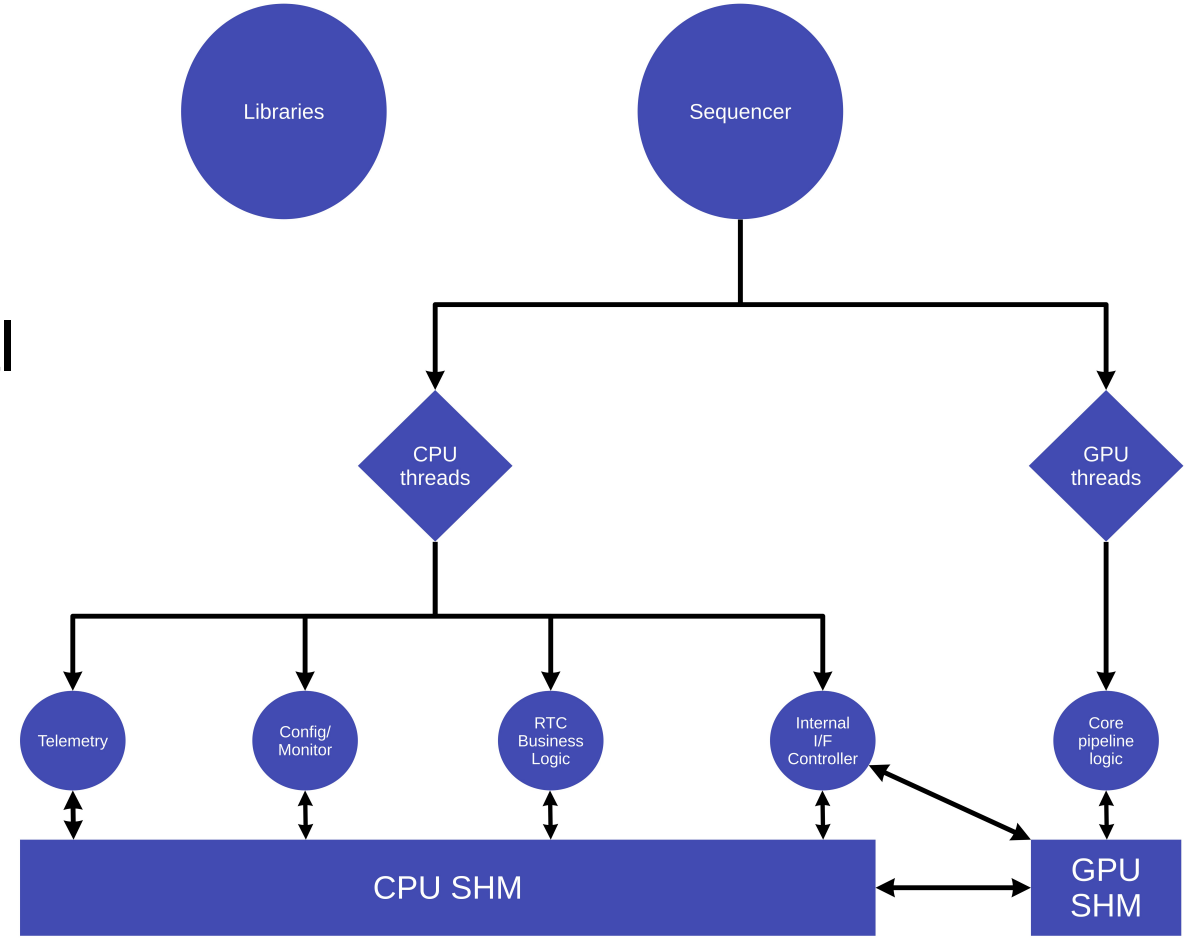

SW stack prototype

- Contributing to the CACAO (Compute And Control for AO) initiative
  - Effort initiated by O. Guyon (Subaru)
  - Based on custom libraries (Interprocess Com., Command Line I/F, numerical tools: MILK), standard libraries (MAGMA, FFTW, GSL) and standard tools (Autotools/Cmake, libtool)
  - Supports GPU offloading (through memcopies)
- Adding an abstraction layer for processes I/F
  - Data I/F and processes synchronization
  - Allow for various HW/SW I/F, ensure modularity
- Converged shared memory model between CPU and GPU processes
- Adding GPU-FPGA DMA libraries
  - Optimized performance when feature is enabled
  - Using the same abstracted I/F
- Adding a python sequencer and PyQt GUIs
- Enabling tight coupling between the real-time core and the COMPASS simulator

CACAO

- Welcoming new contributors
  - Goal is to move forward community support for this SW framework
  - And getting feedback from a wide user community will make it more reliable
  - Open Source model is key to optimize both cost and reliability
- Portability, flexibility, modularity, scalabilty
  - Multi-platform support (including optimized HW approaches)
  - Optimized libraries for predictive control, tomographic AO and more
  - Python I/F, reduce the amount of custom libraries and tools by switching (and adapting) to standards
- Already driving systems at Subaru (SCExAO, AO188) and will soon at Keck and more
- Give it a try ! https://github.com/cacao-org/cacao