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

A status update … (a.k.a. teaser slides)

AO4RTC4 workshop, Paris 2016

European Commission

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

AO RTC concept

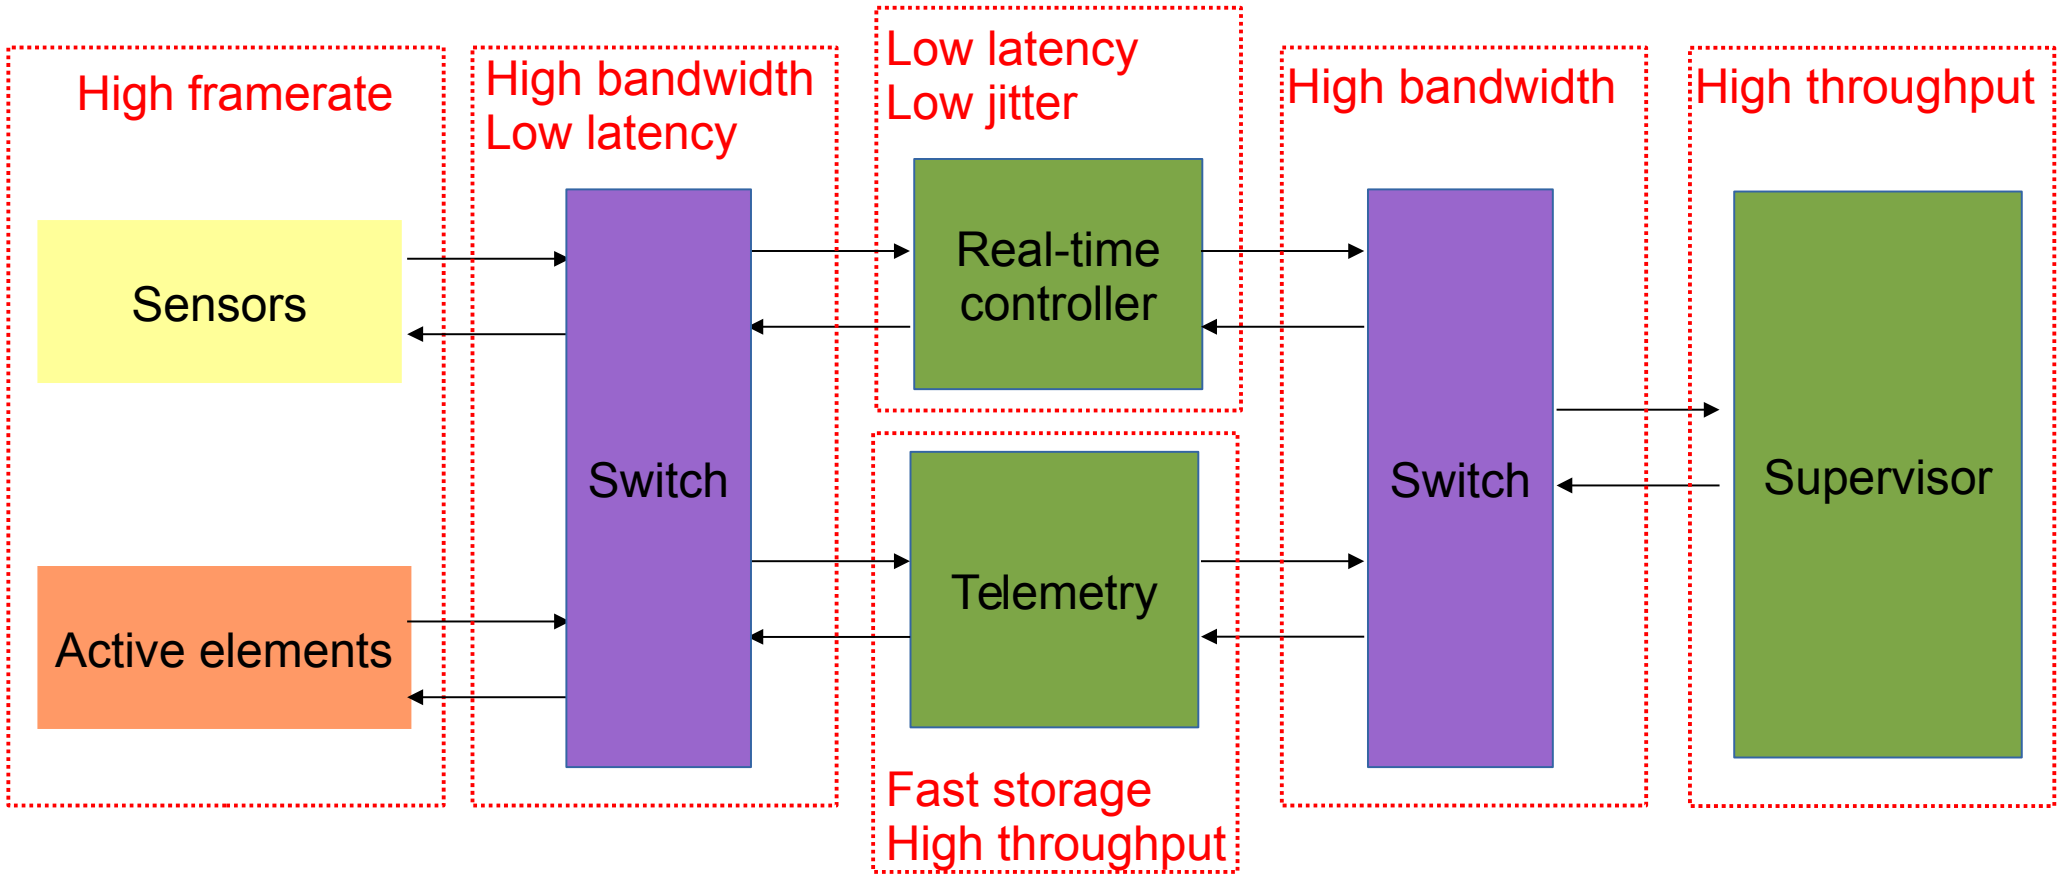

Introduction to Green Flash

- Program objectives: 3 research axes
  - 2 technological developments and 1 validation study
- Real-time HPC using accelerators and smart interconnects
  - Assess the determinism of accelerators performance
  - Develop a smart interconnect strategy to cope for strong data transfer bandwidth constraints
- Energy efficient platform based on FPGA for HPC
  - Prototype a main board, based on FPGA SoC and PCIe Gen3
  - Cluster such boards and assess performance in terms of energy efficiency and determinism
- AO RTC prototyping and performance assessment
  - Assemble a full functionality prototype for a scalable AO RTC targeting the MAORY system
  - Compare off-the-shelf solutions based on accelerators and new FPGA-based concept

What this is about … really

- Find the best trade-off for ELT sized AO systems RTC
  - Comprehensive assessment of existing technologies
  - Development of new custom solutions for comparison
  - Propose new development processes to reduce cost and increase maintainability
- Build a full featured RTC prototype at the largest scale possible
  - Technology down-selection from a number of criteria : performance, cost, compliance to standards, obsolescence, maintainability
  - State of the art system to be assessed in the llab, with a simulator

Project timeline

- Good convergence with H2020 ETP4HPC / E-ELT project timeline

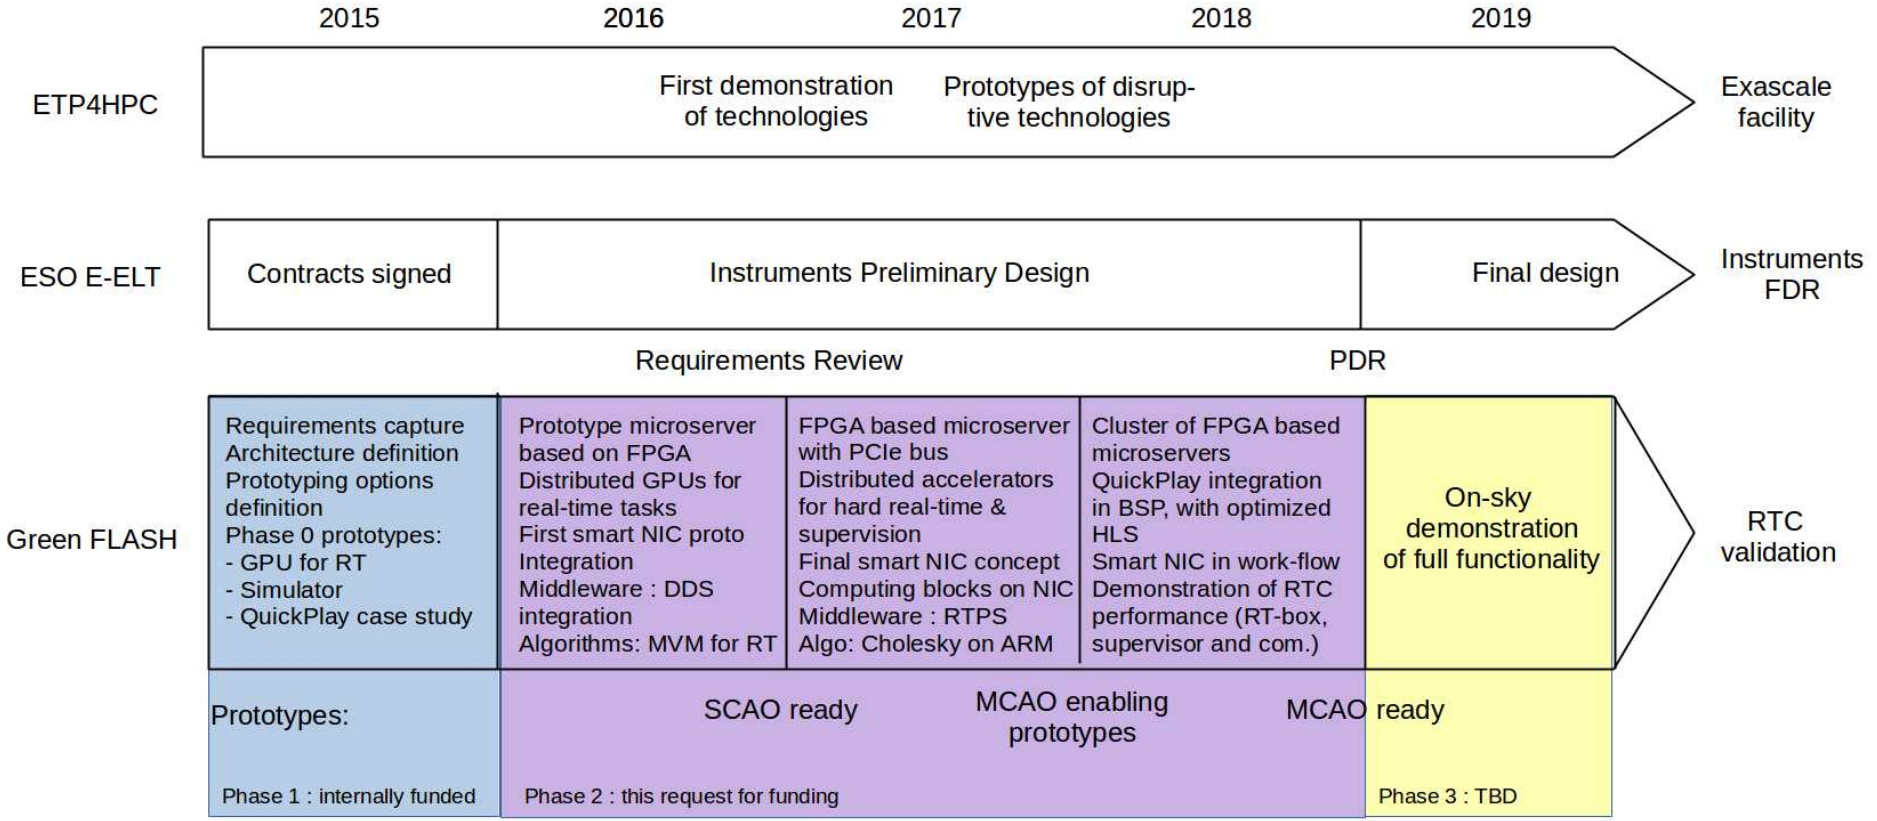

AO RTC concept

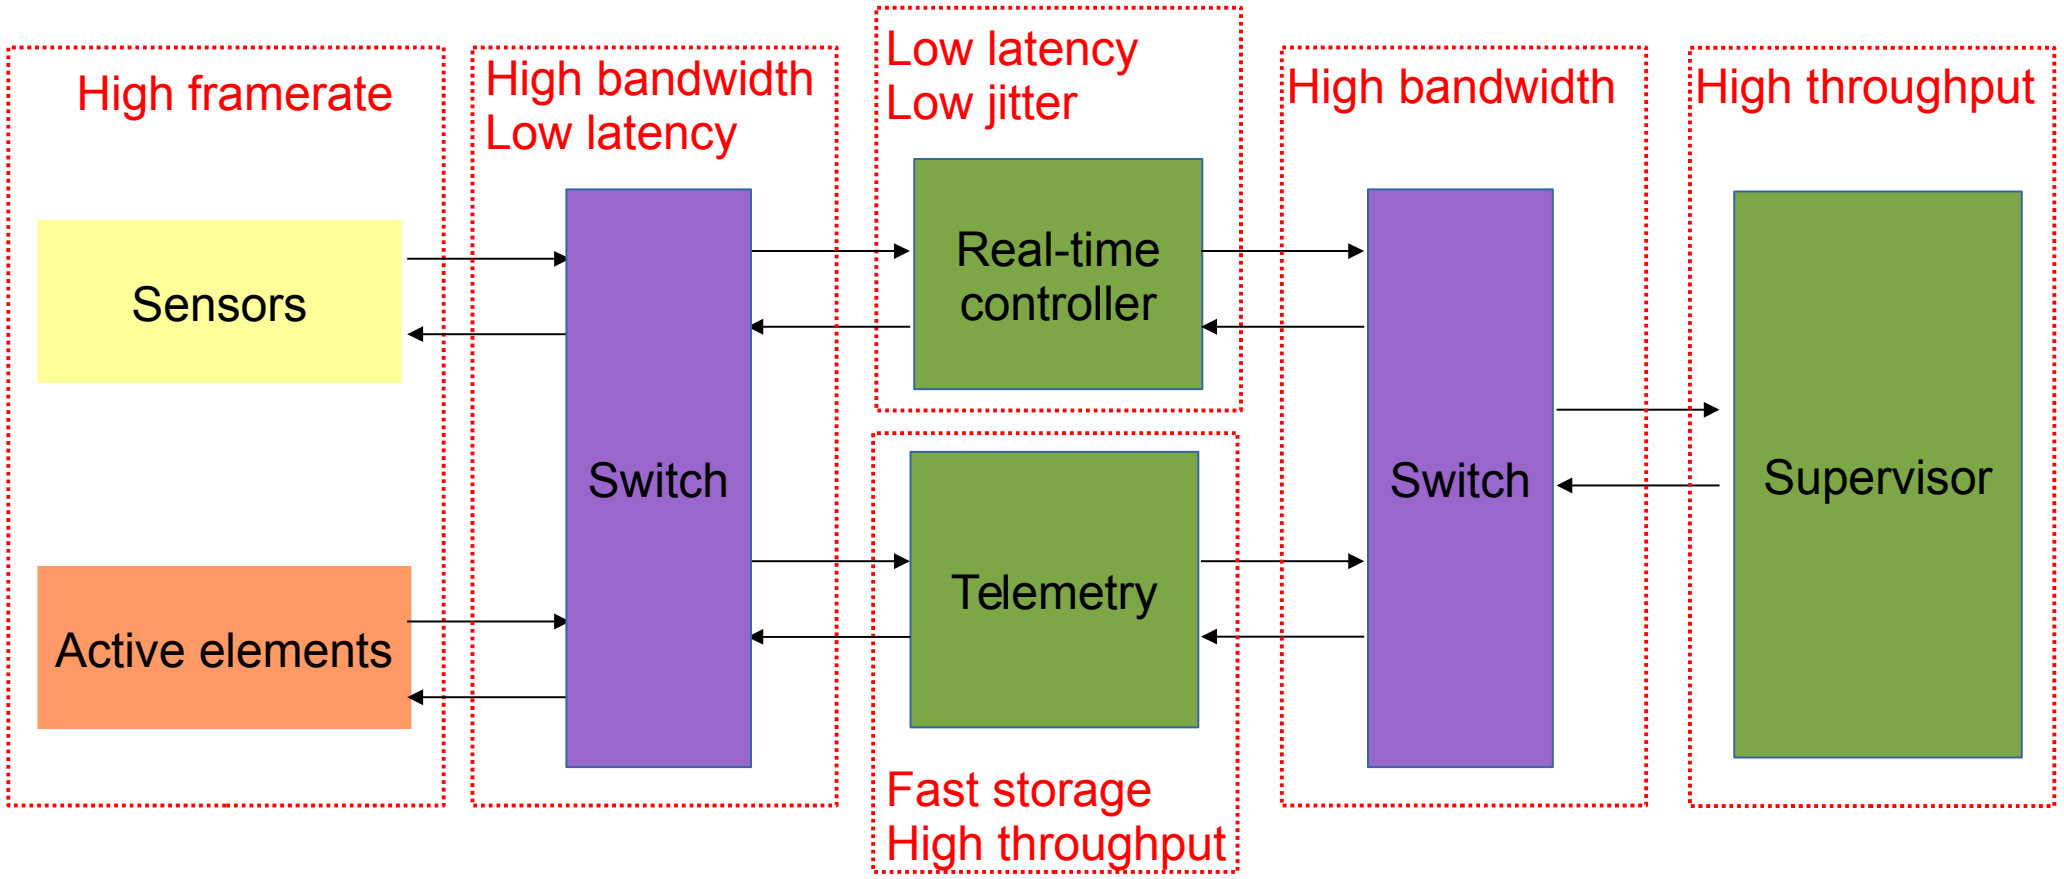

AO RTC concept : RT simulator

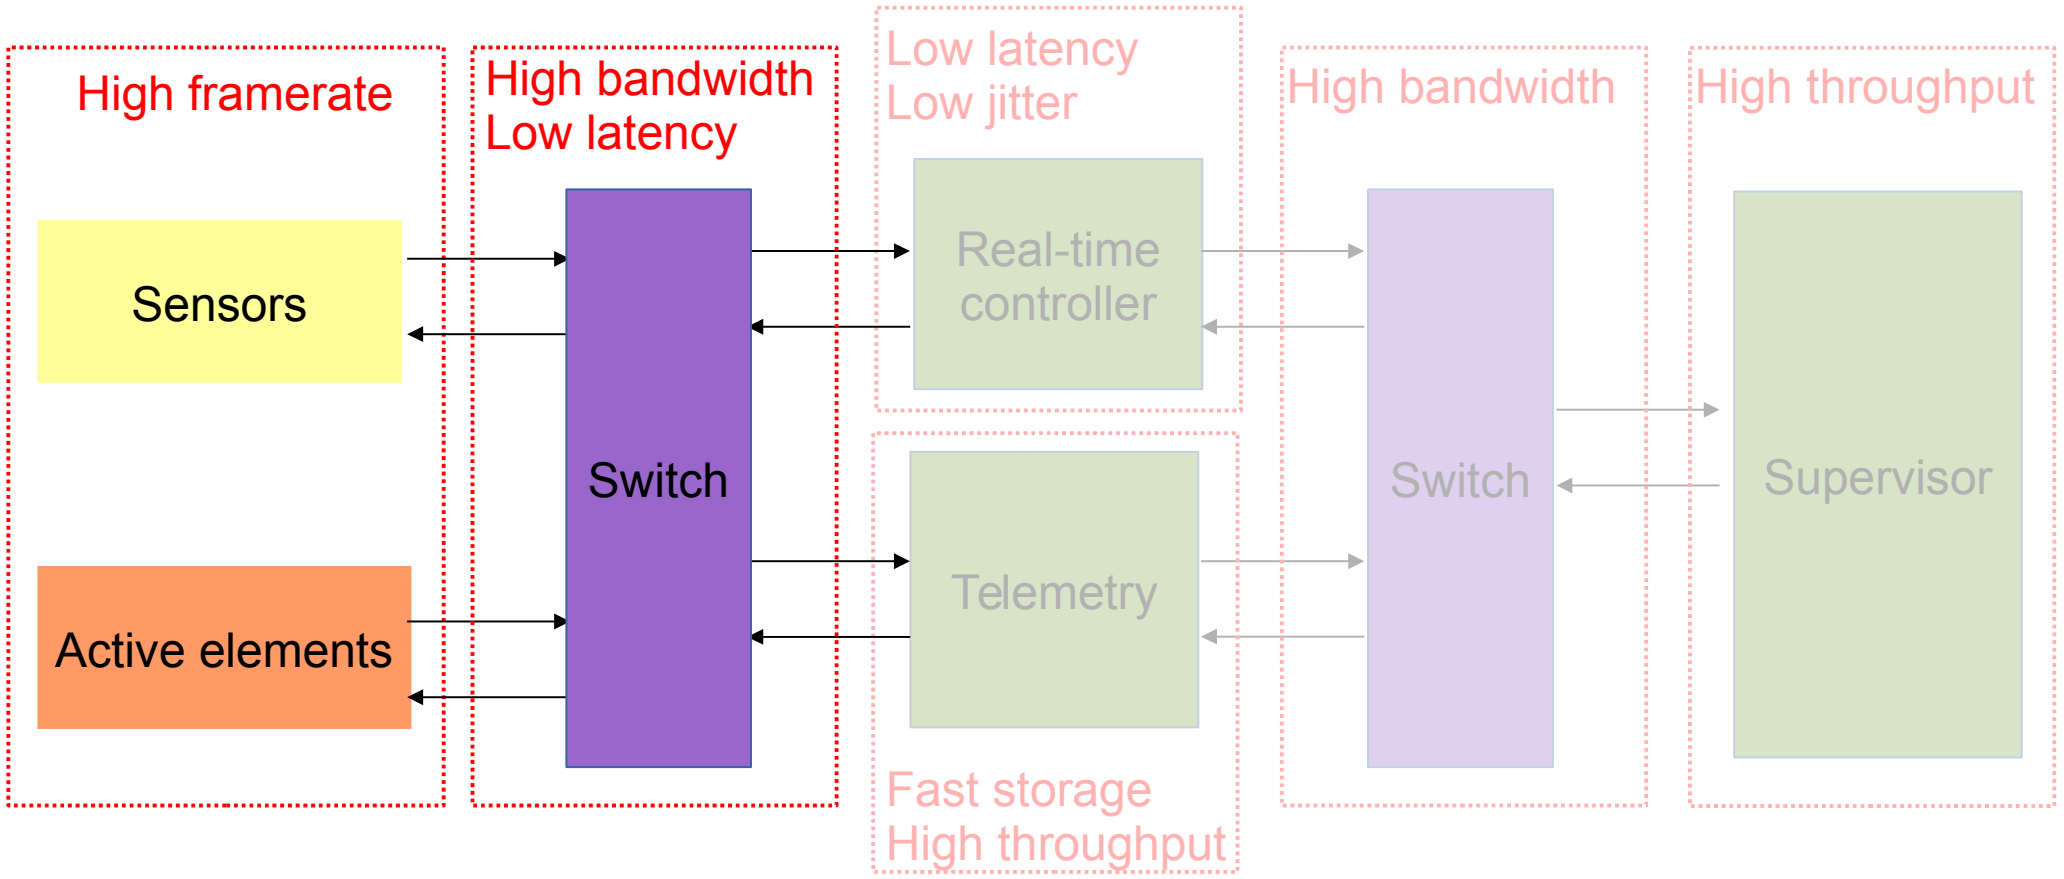

Real-time simulator

- Emulation interface to RTCS

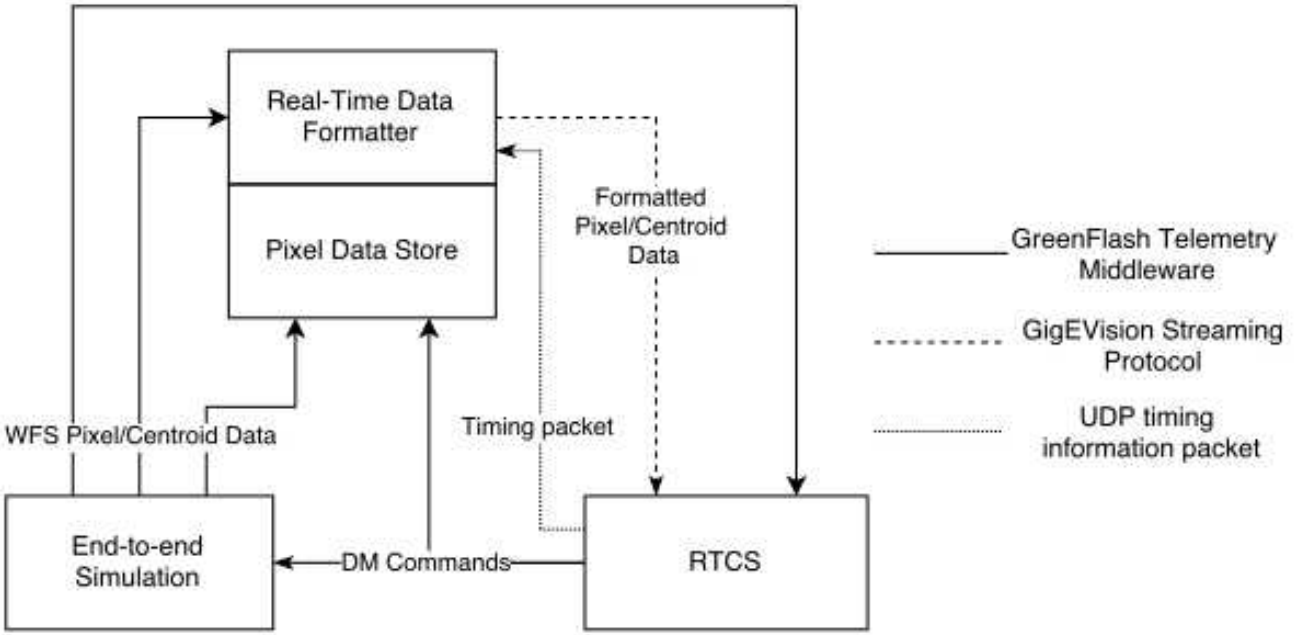

- Concept developed @ Durham, see Andrew Reeves presentation

Real-time simulator

- Several simulation tools :
  - Low level functional test with slow python-based simulation (SOAPy)
  - High-level tests with fast end-to-end simulation (COMPASS)
  - Performance / determinism assessment with HW pixel generator

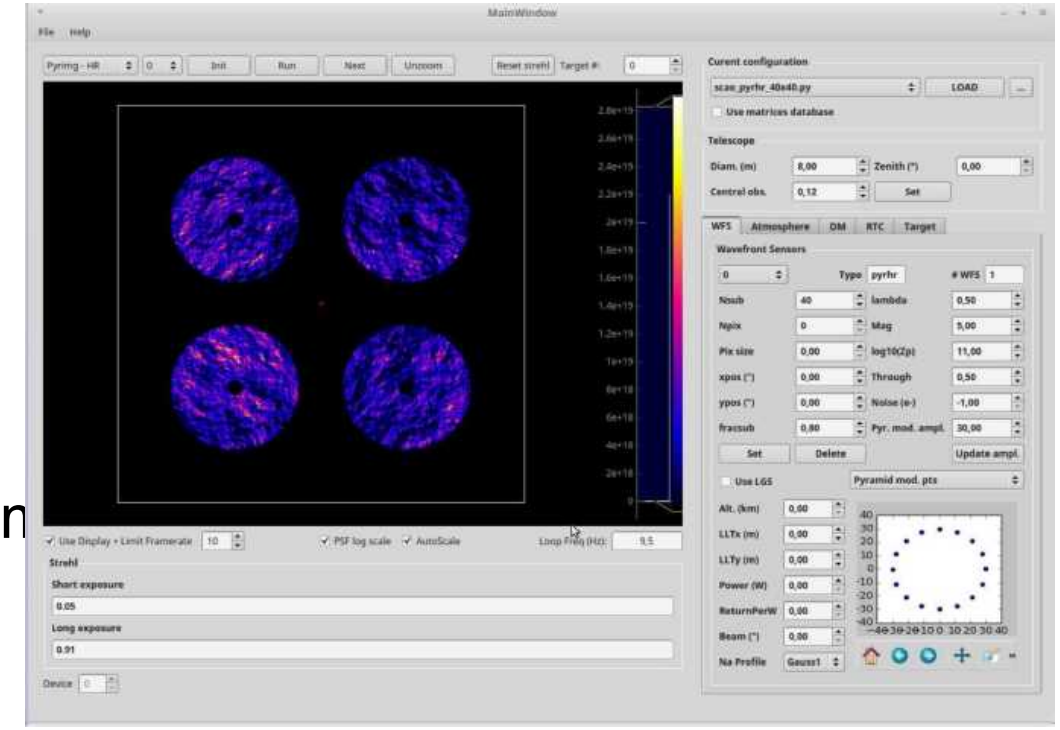

Real-time simulator

- Using COMPASS for E2E should provide a scalable solution over the long term
  - Execution times from F. Ferreira

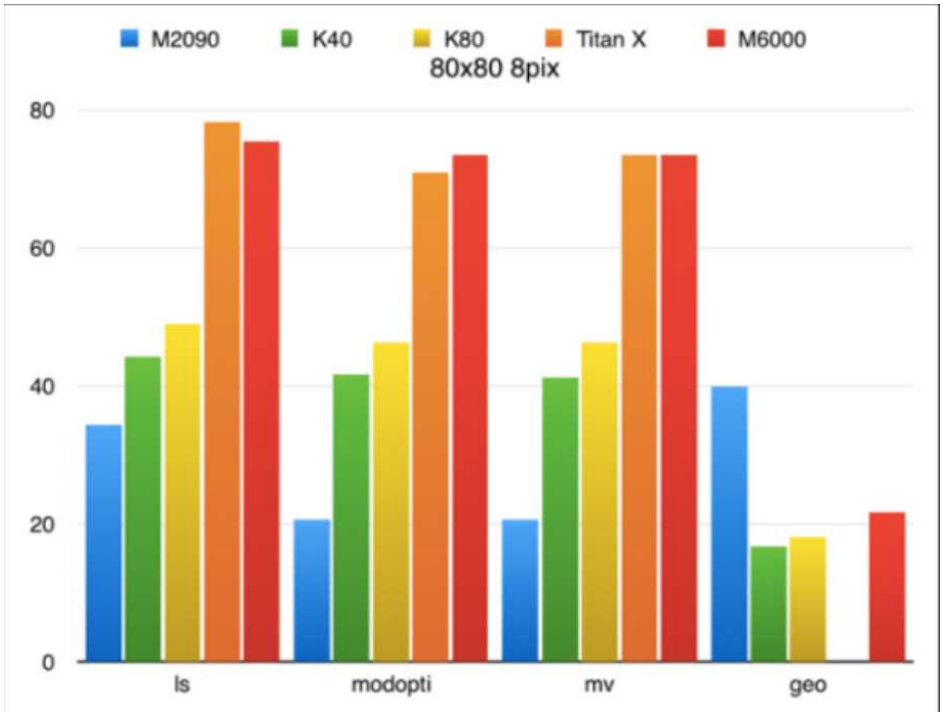

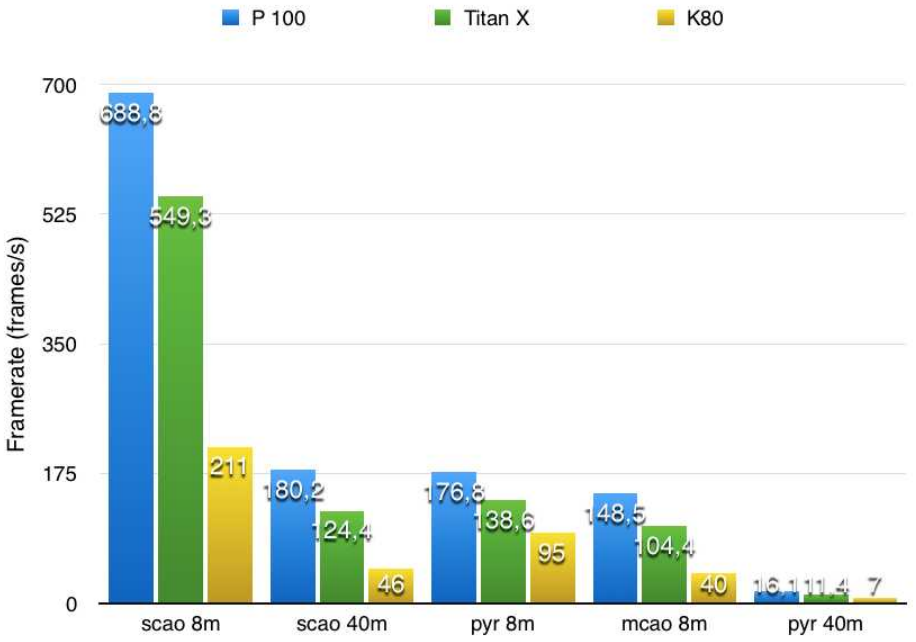

AO RTC concept

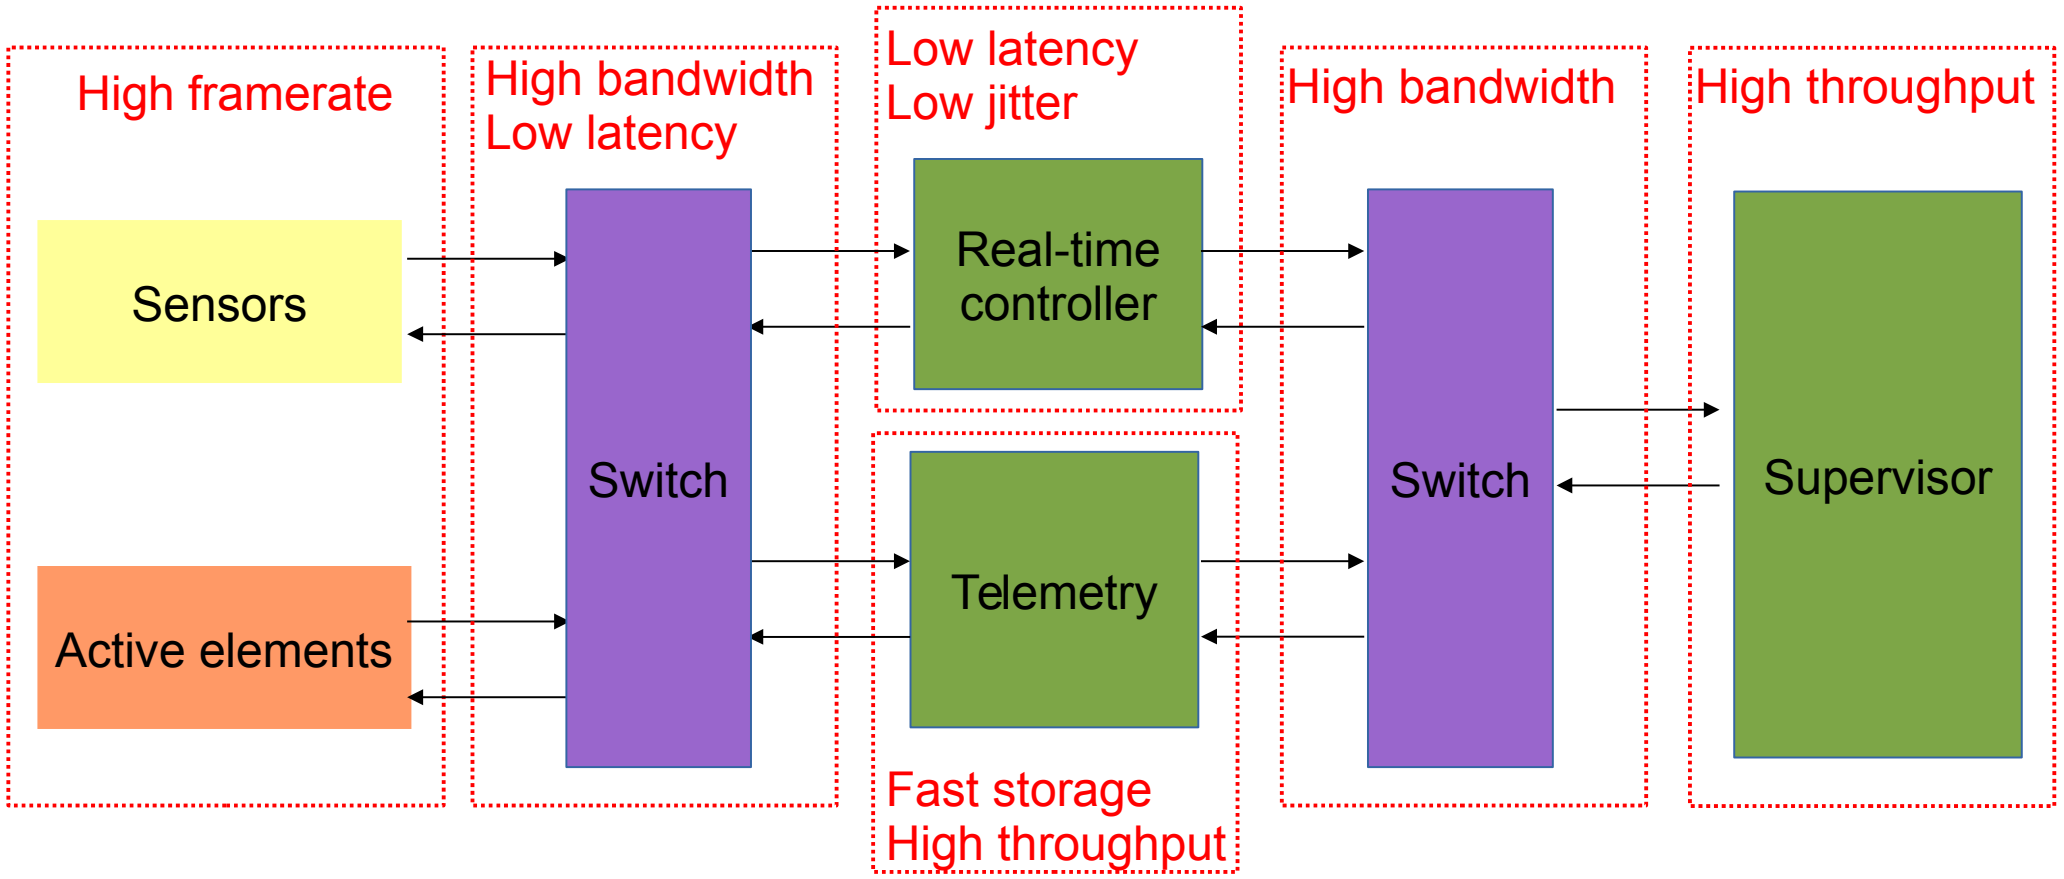

AO RTC concept : data pipeline

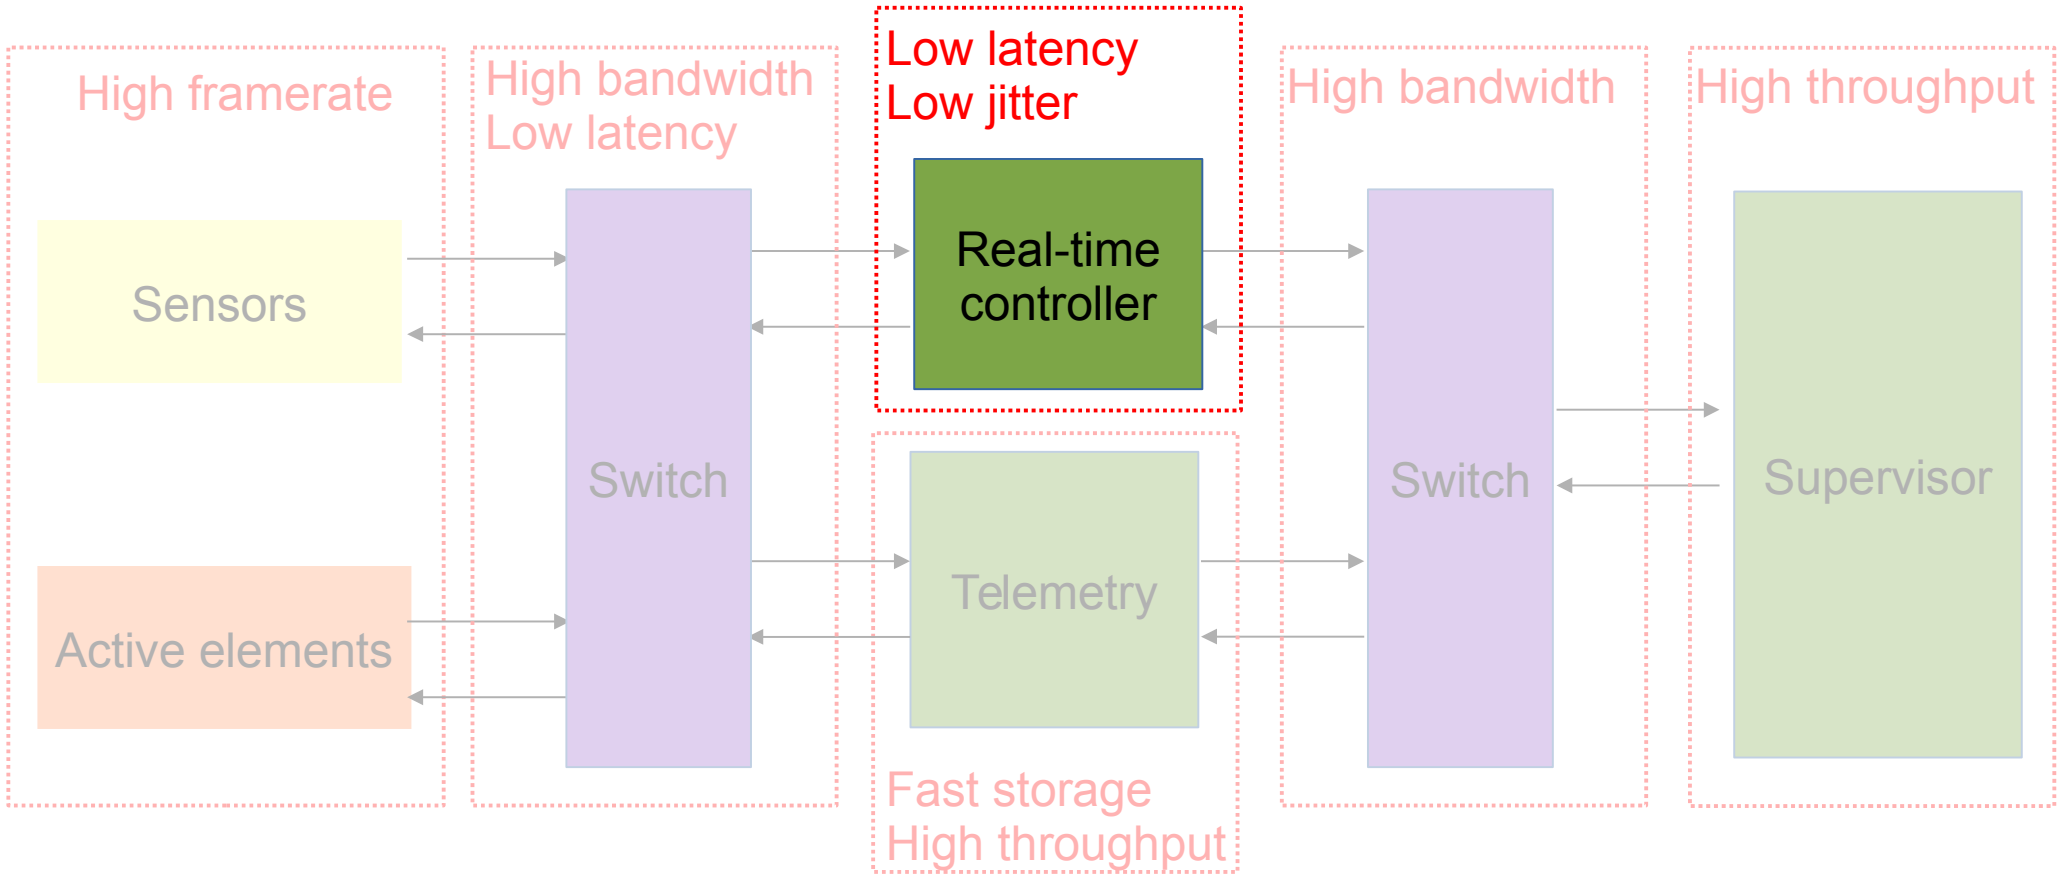

RT data pipeline with accelerators

- Study the achievable perf. on GPUs

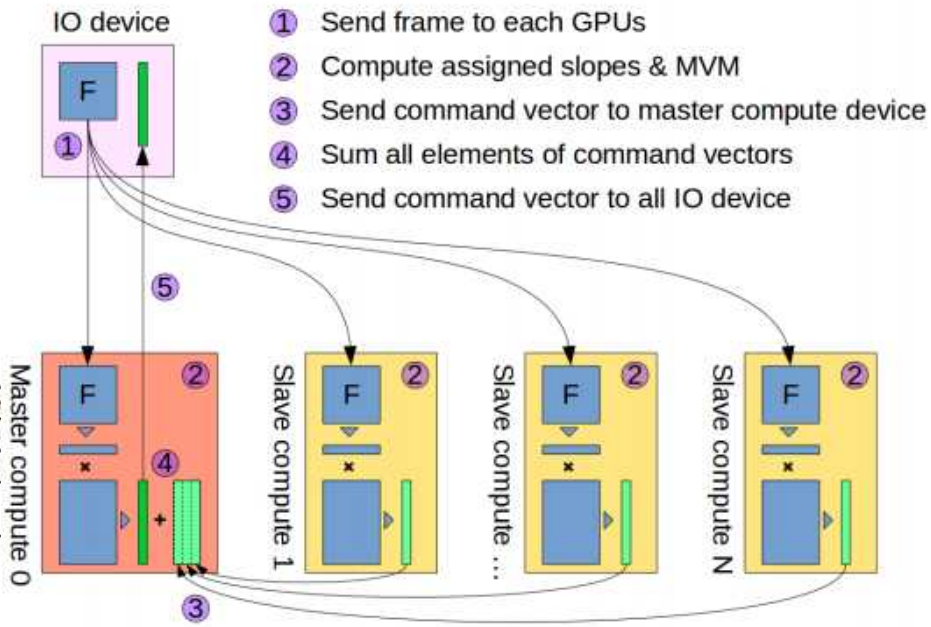

Ctrl	Clock	Average	Min	Max	Jitter	Computation rate	Memory transfe
RTC 1 device.s	device	0.51 ms	0.5 ms	0.52 ms	23 us	2e+02 GFlops	3.9e+02 GB
RTC 2 device.s	device	0.32 ms	0.31 ms	0.33 ms	12 us	3.1e+02 GFlops	6.3e+02 GB

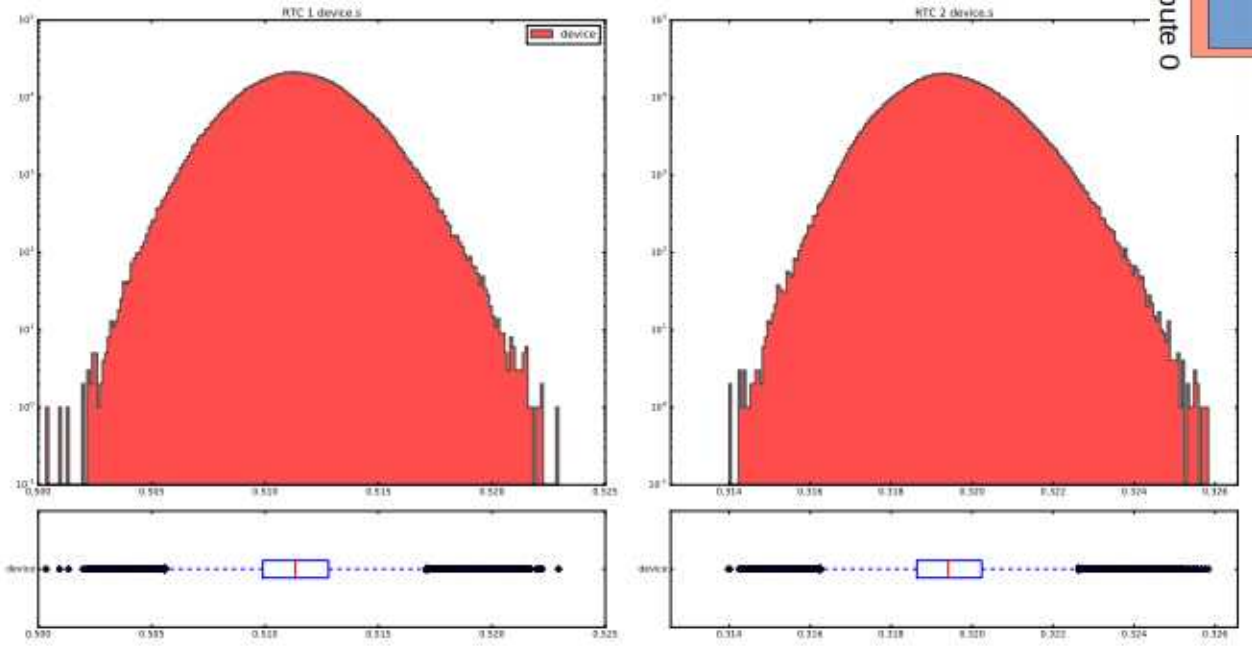

RT data pipeline with accelerators

- Prototype using latest genration GPU cluster

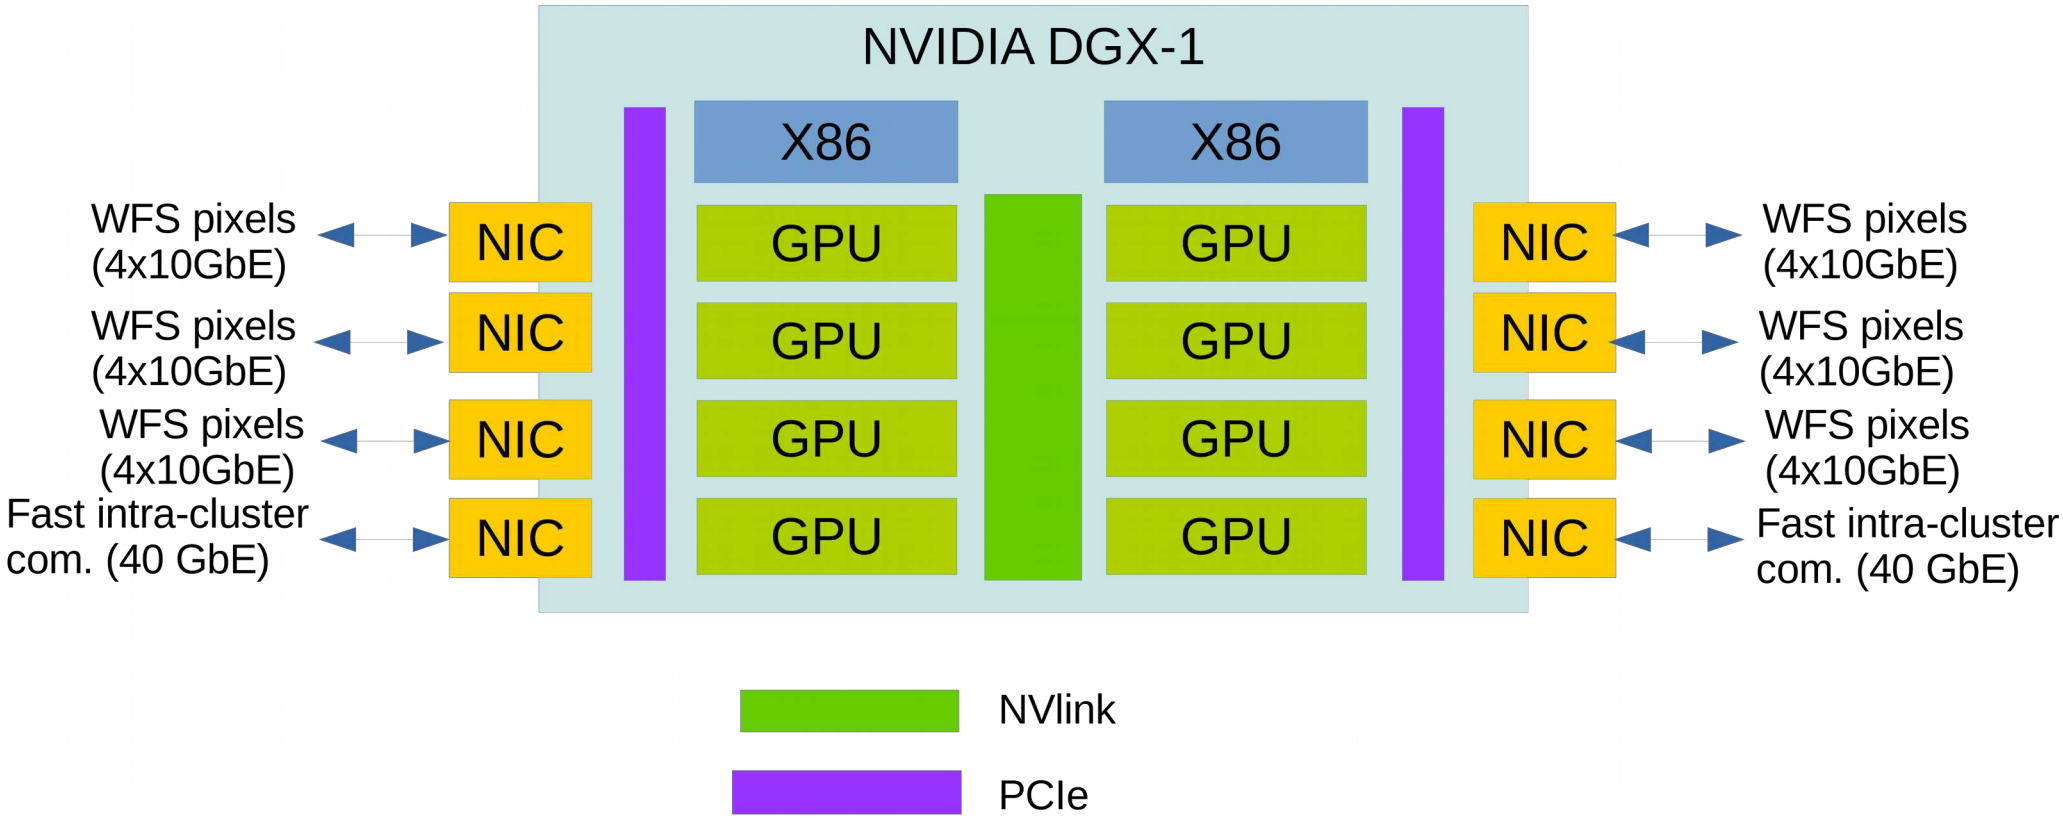

- Concept studied at LESIA, see talks by Julien Bernard & Maxime Lainé

RT data pipeline with accelerators

- Intel Xeon Phi
- First results on Knights Corner (10k meas. vector)

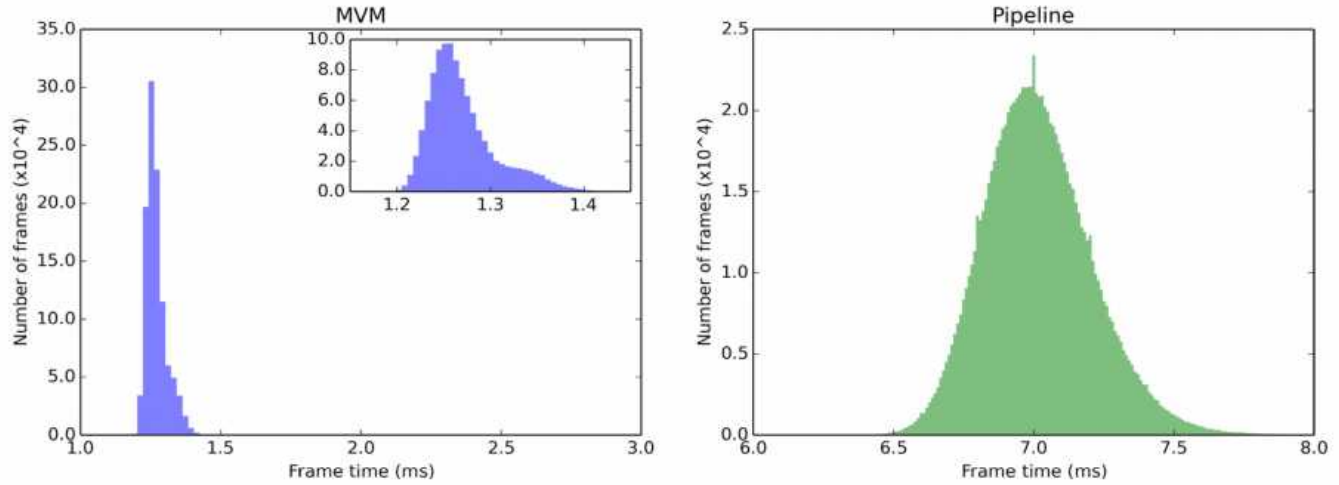

	Knights Corner	Knights Landing predicted	Knights Landing preliminary measured
MVM computation time	1.2-1.3 ms	0.9 ms	0.76ms (0.63us when optimised for thread count)
Memory bandwidth	165 GB/s	250 GB/s	258 GB/s TRIAD, 235GB/s MVM

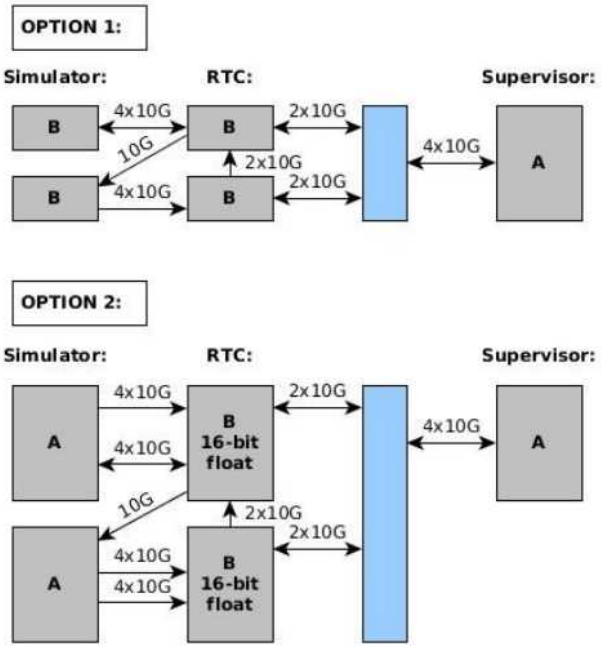

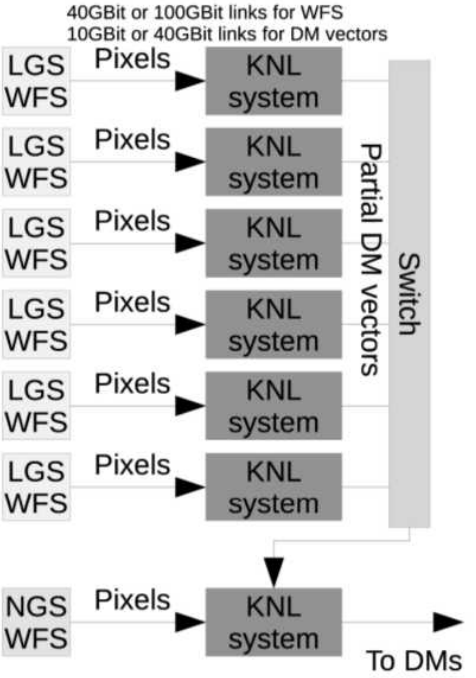

- Concept studied at Durham, see talk by David Jenkins

RT data pipeline with new FPGA microserver

- Exploring new concepts for HPC

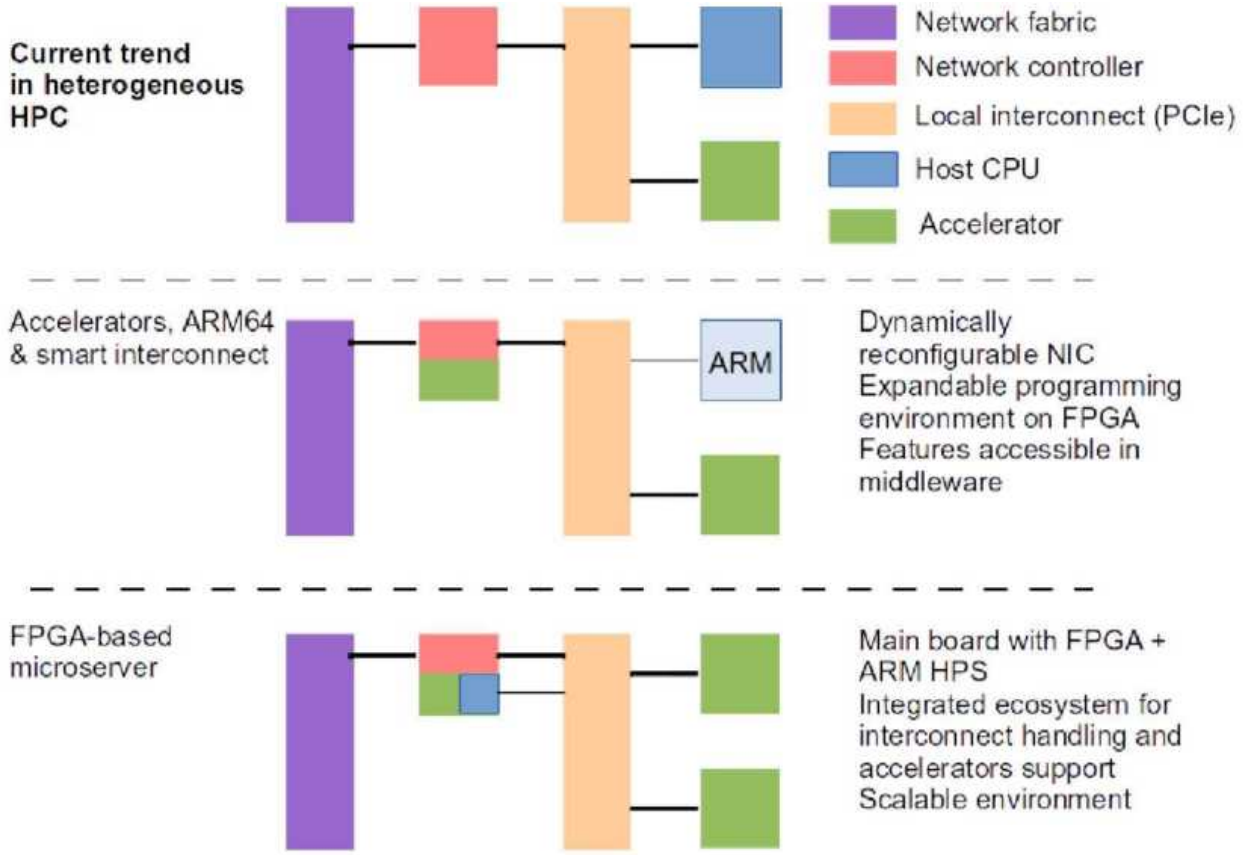

RT data pipeline with new FPGA microserver

- 2 boards concepts : standard (accelerator) / standalone (w/ SOC : microserver)

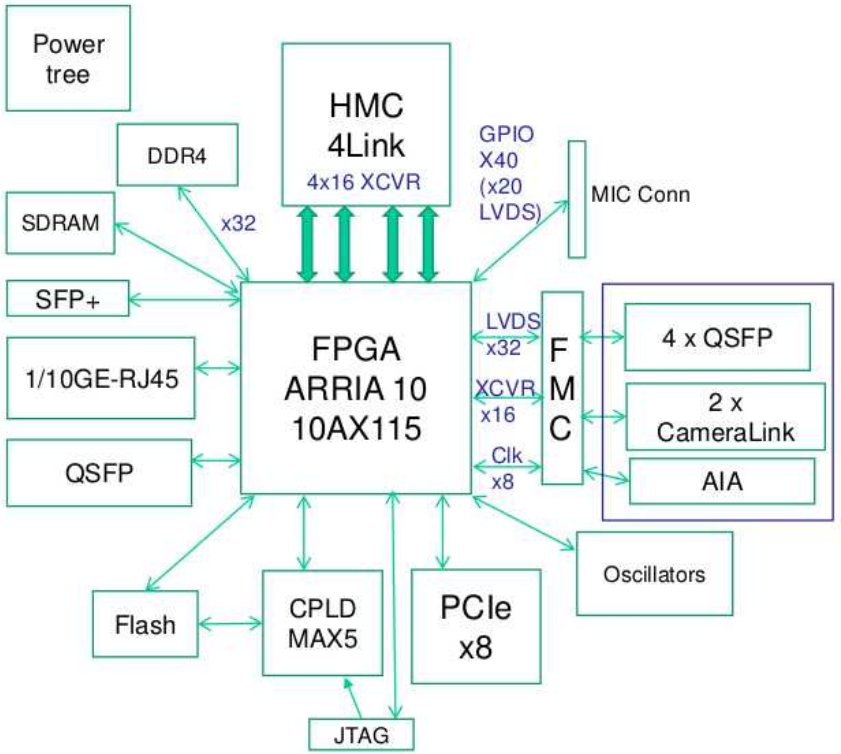

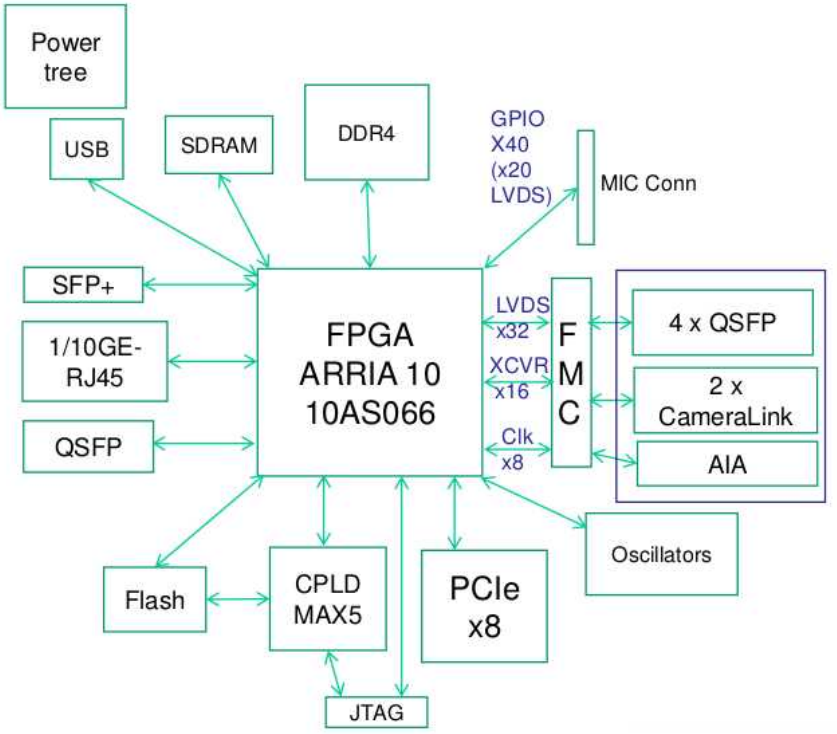

RT data pipeline with new FPGA microserver

- Boards are becoming real @ Microgate !

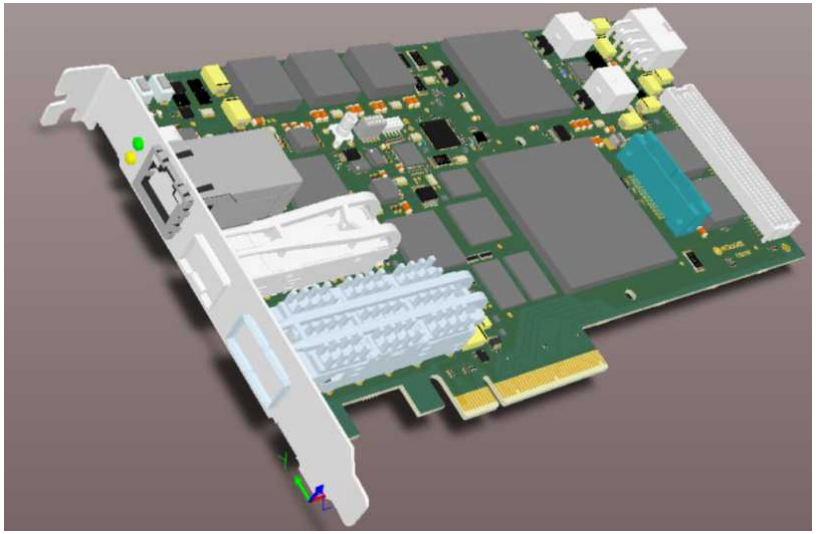

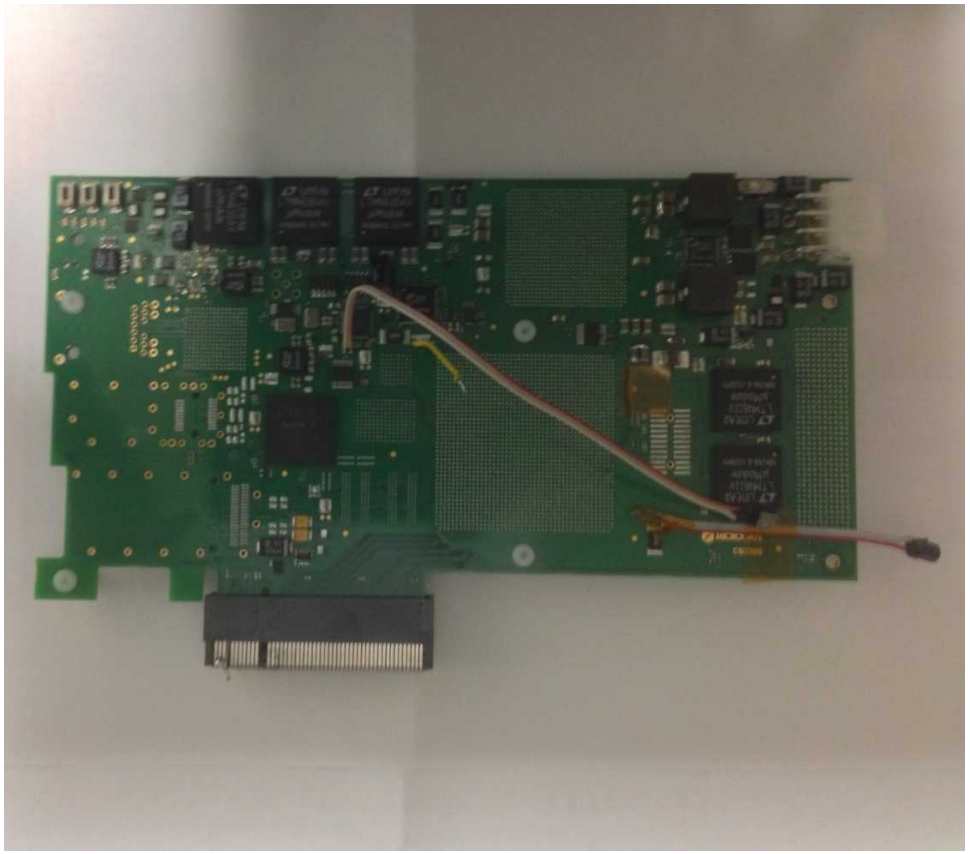

AO RTC concept

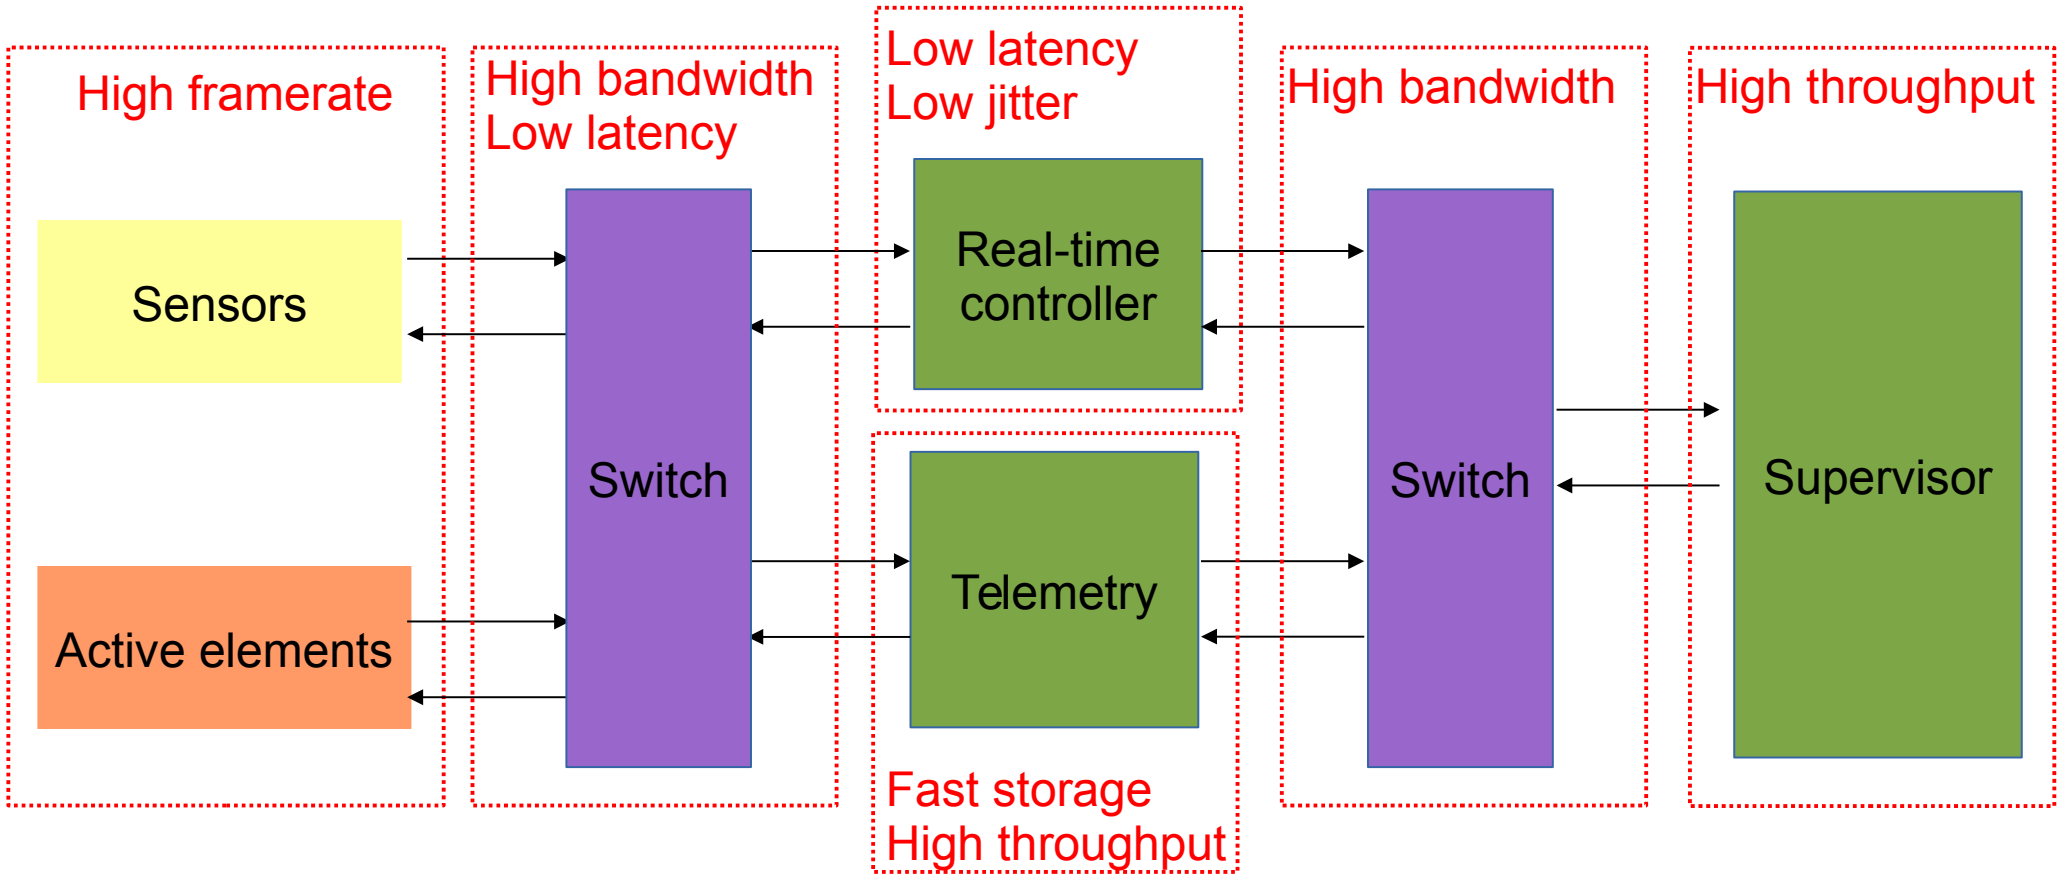

AO RTC concept : smart interconnect

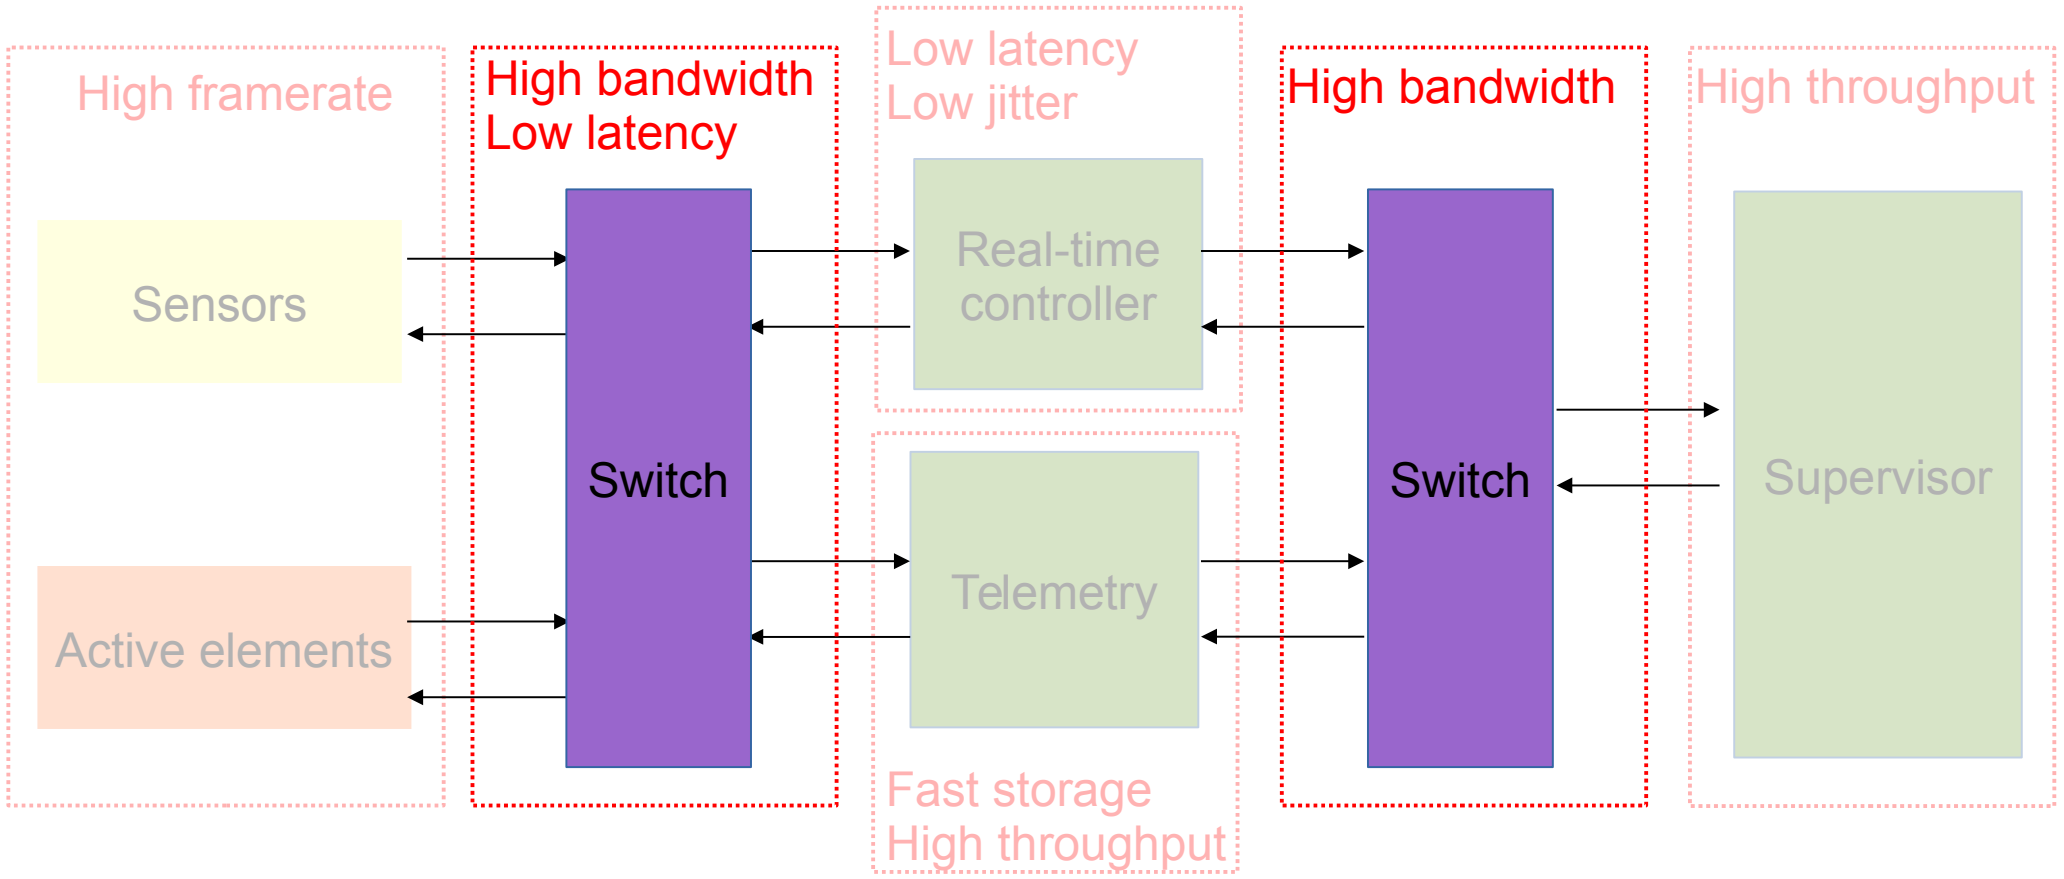

Smart interconnect concept

- Collaboration between PLDA and LESIA
- Proof of concept on a SCAO ready prototype

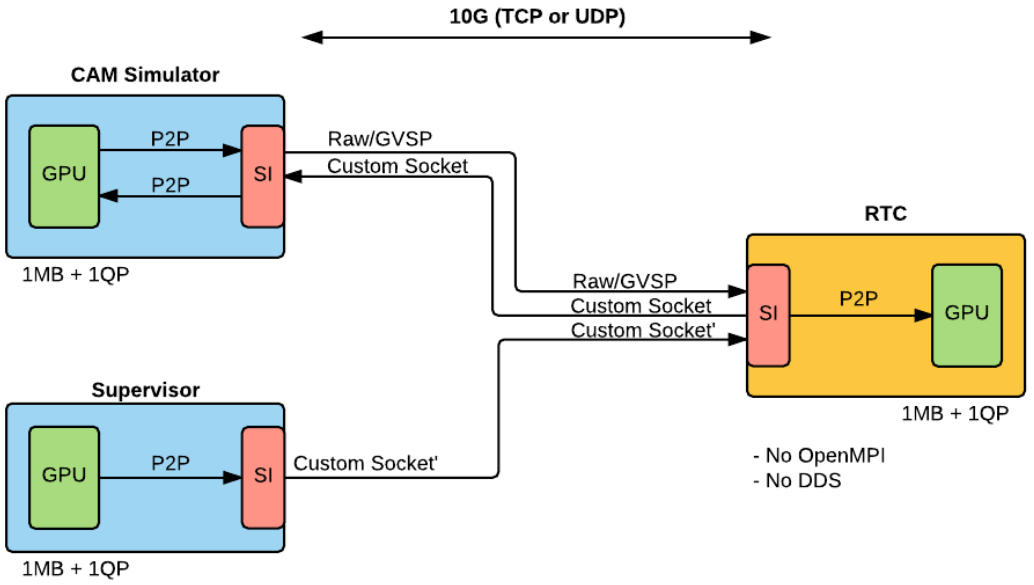

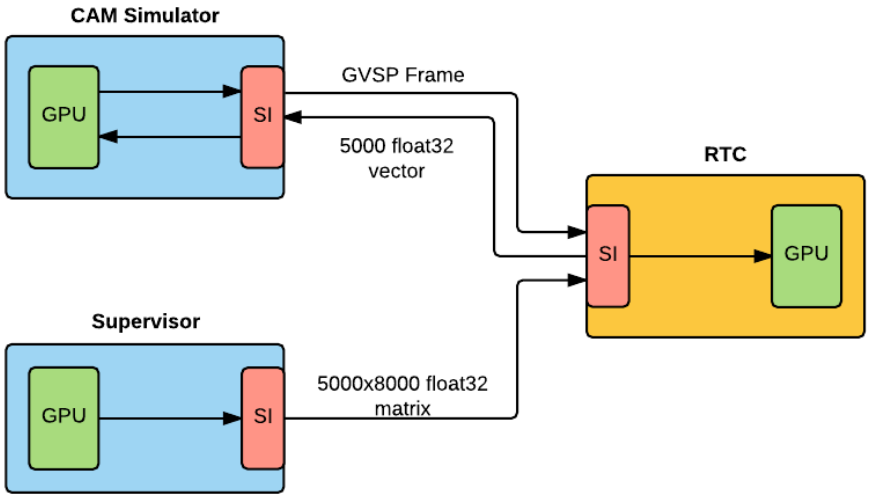

Smart interconnect architecture

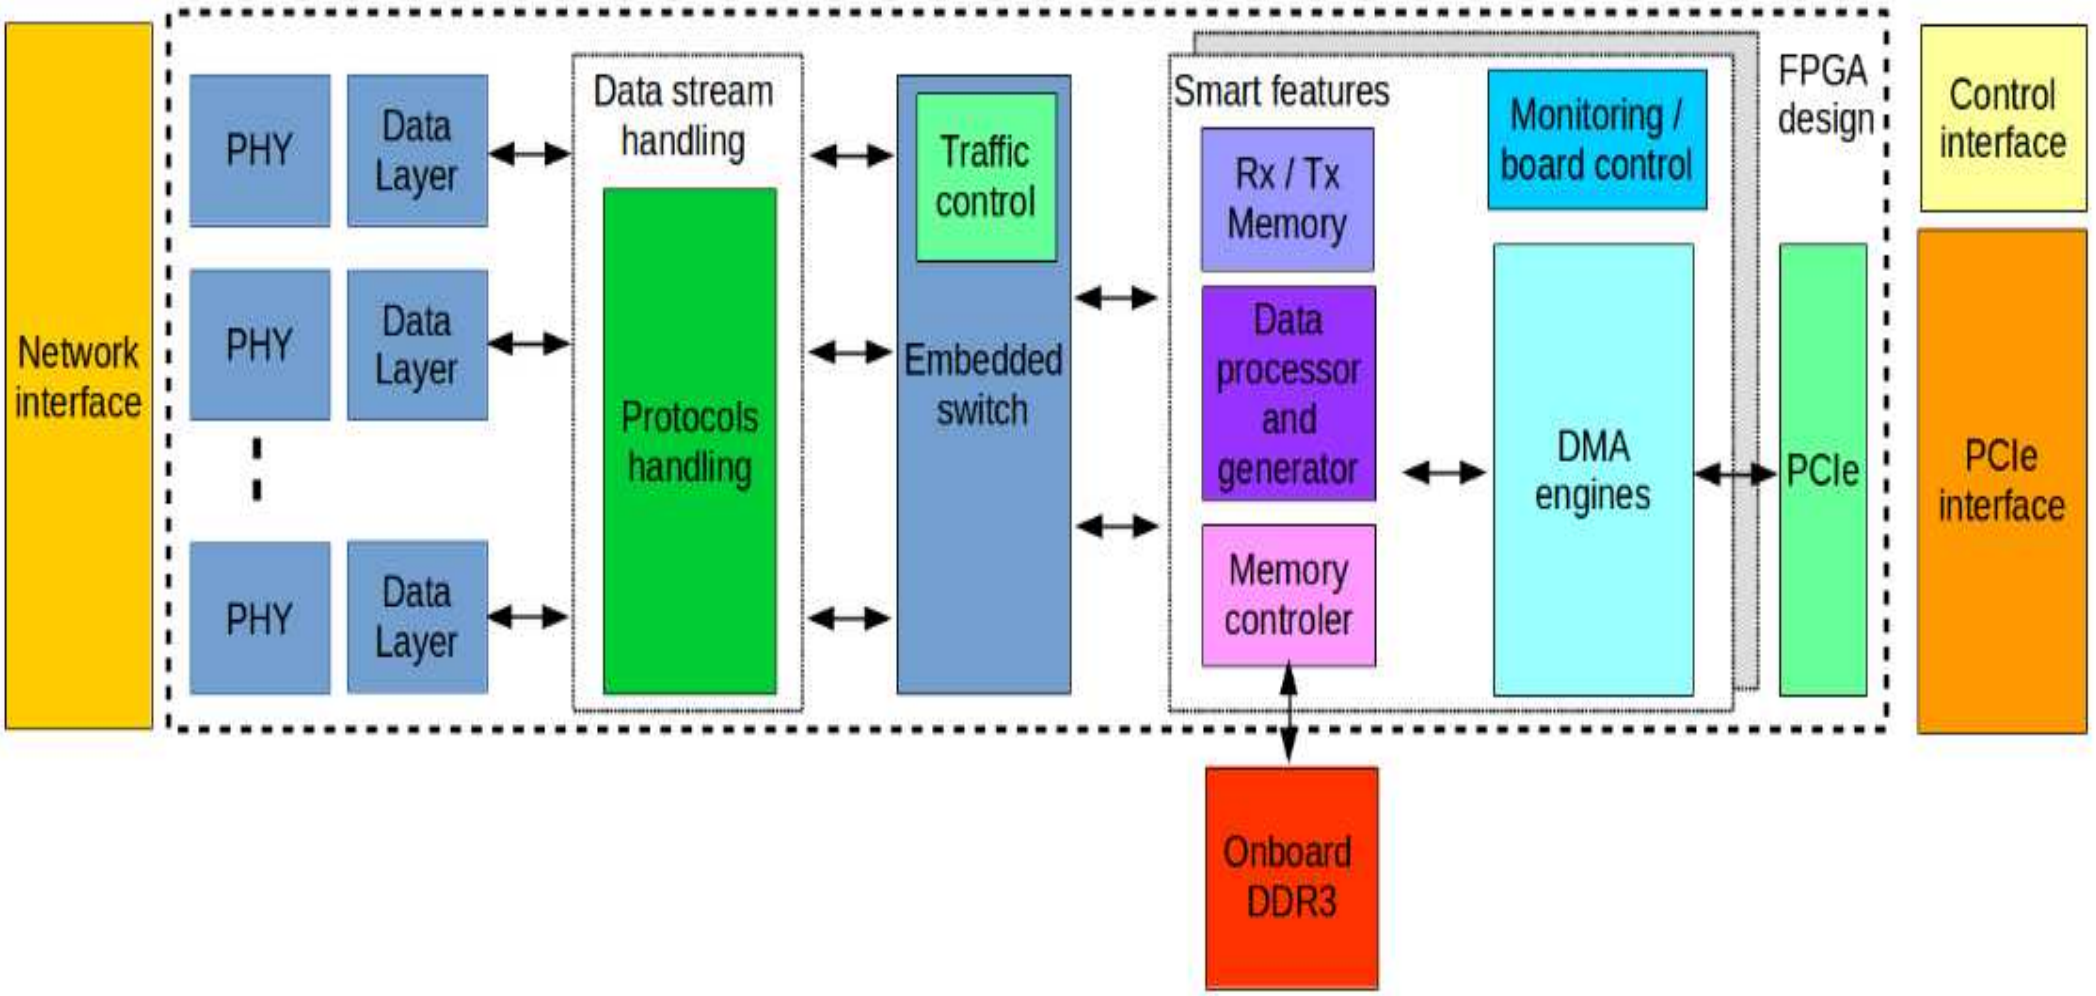

Smart interconnect concept

- Eased devel. process using the QuickPlay tool from PLDA

QuickPlay™

QuickStore™

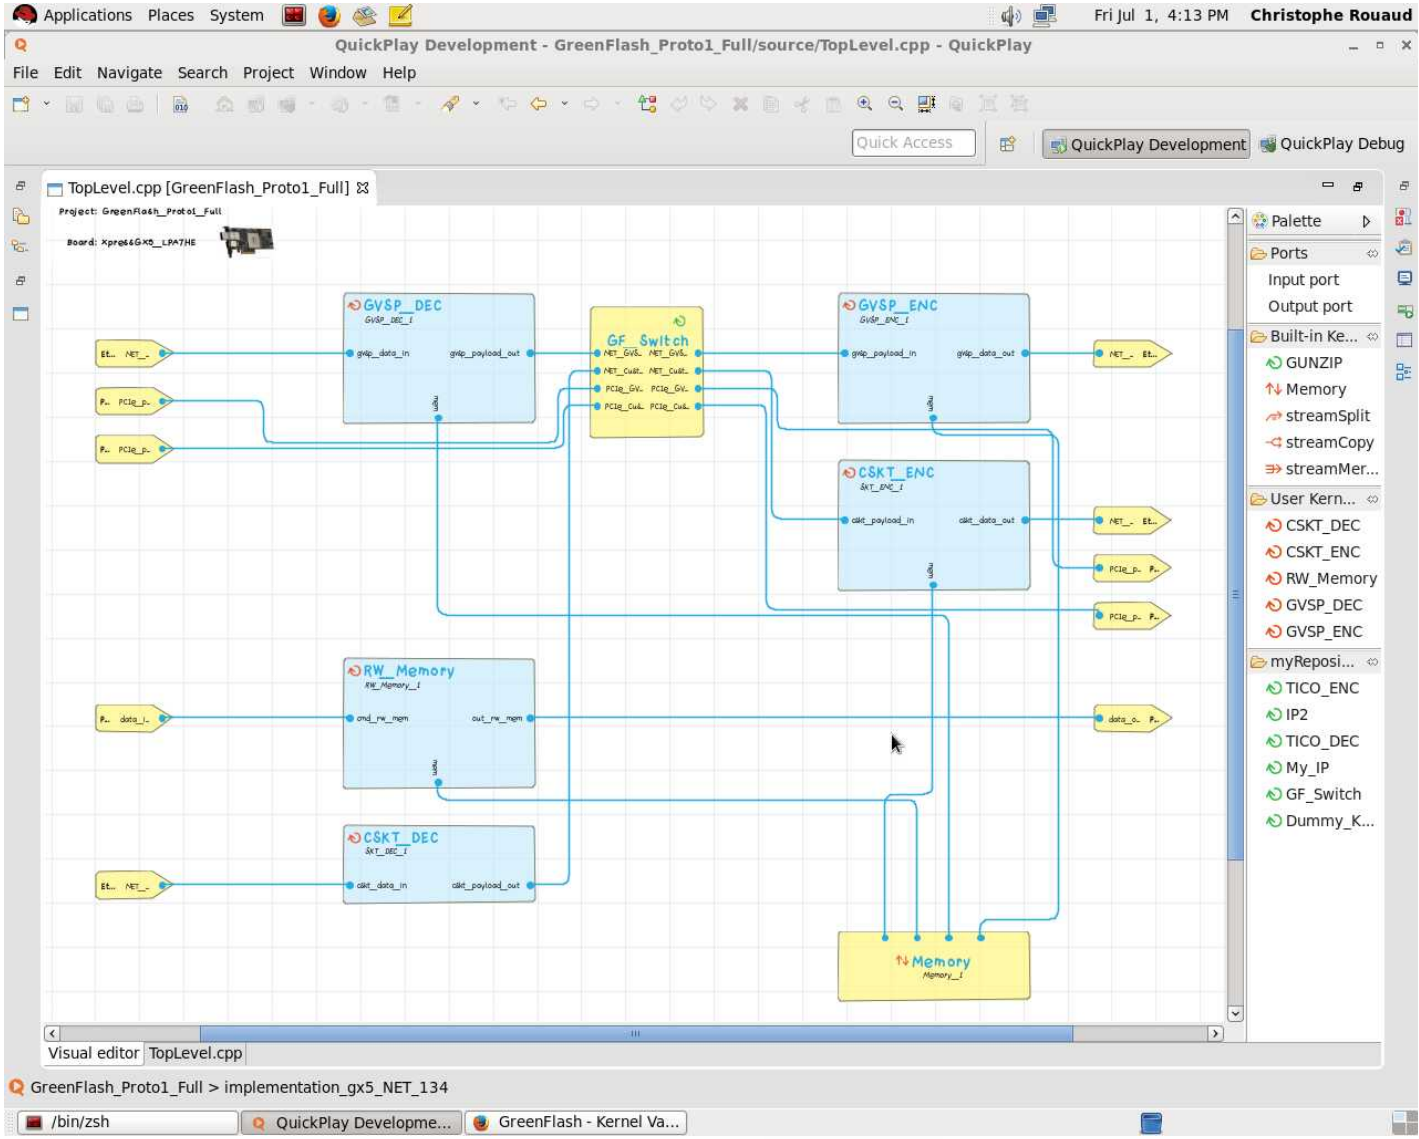

AO RTC concept

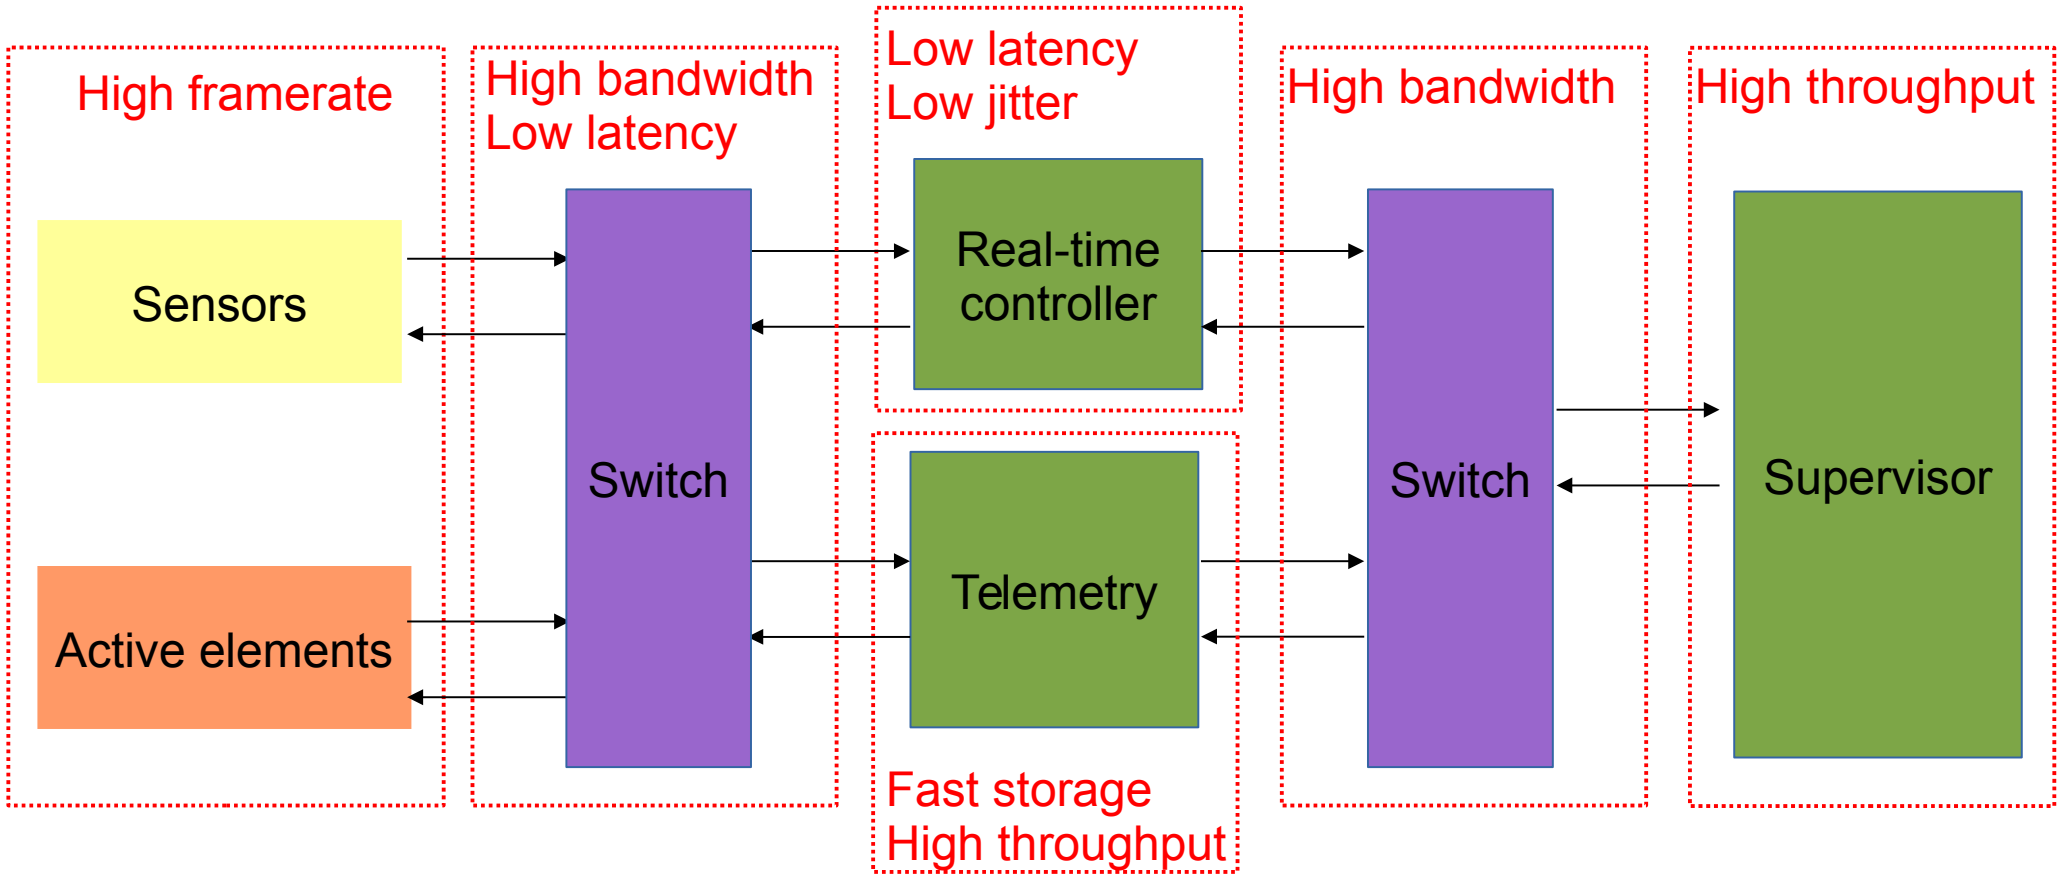

AO RTC concept : supervisor

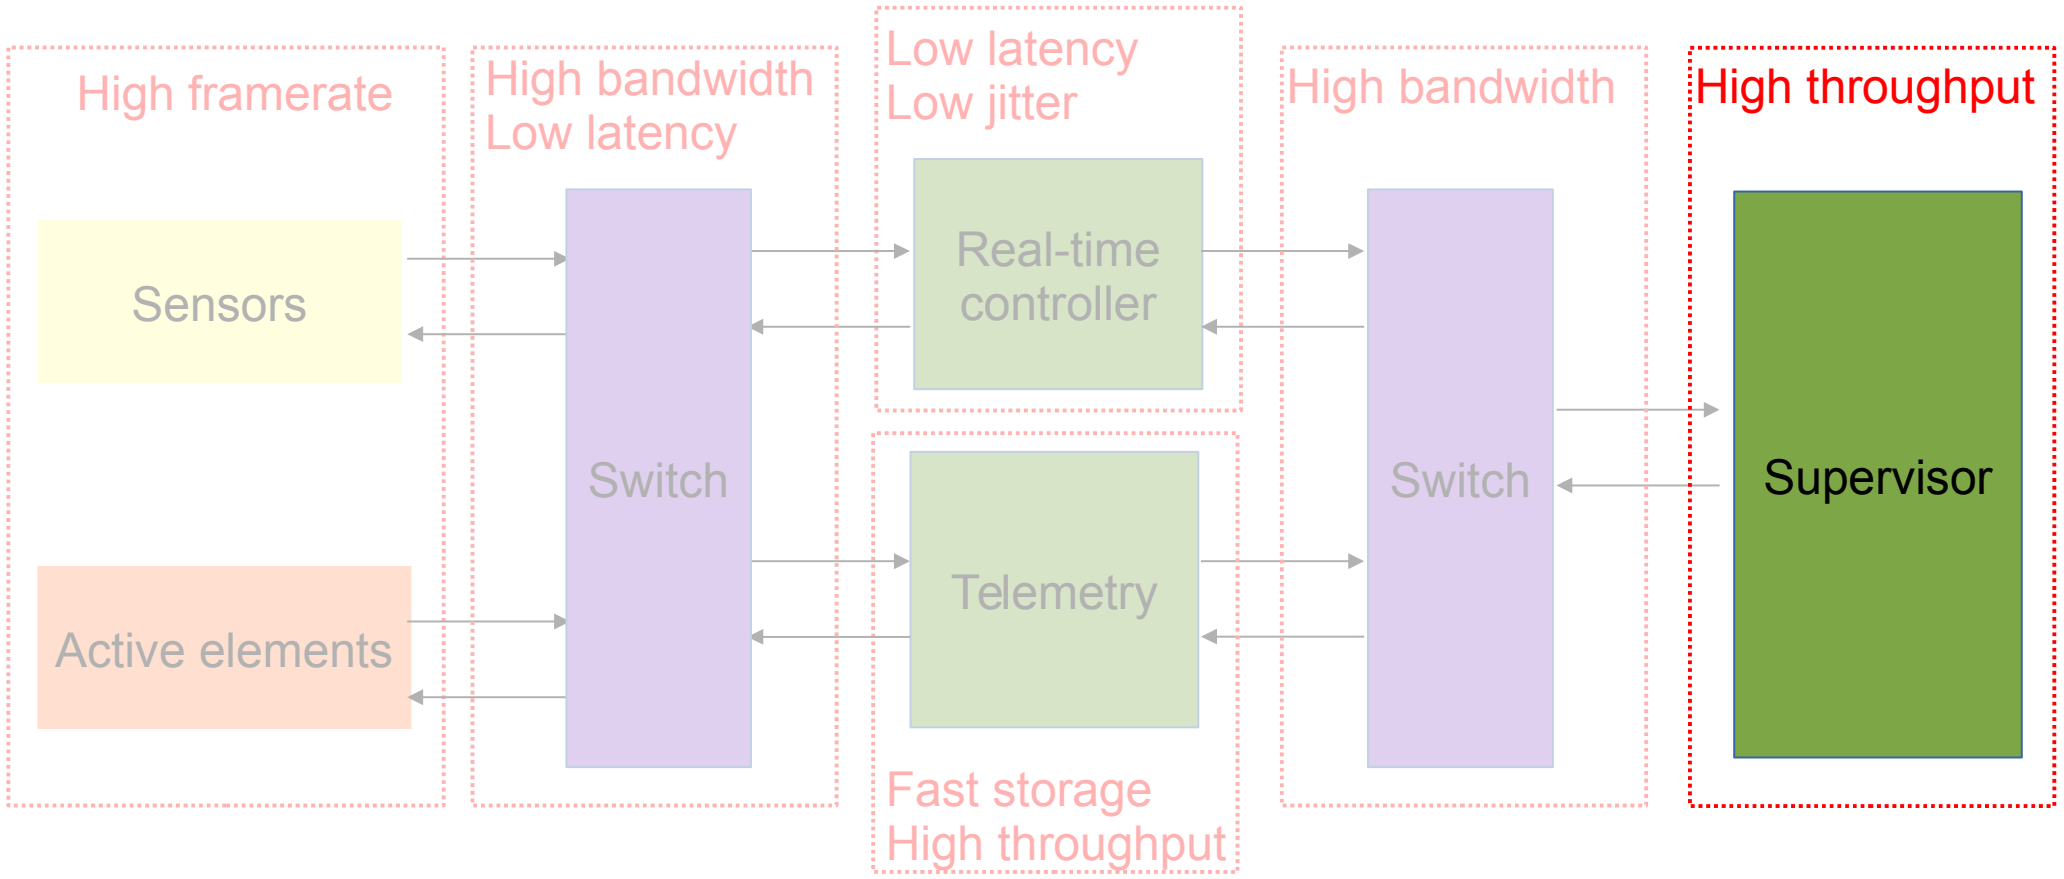

Supervision for tomographic AO

- New, optimized supervision pipeline

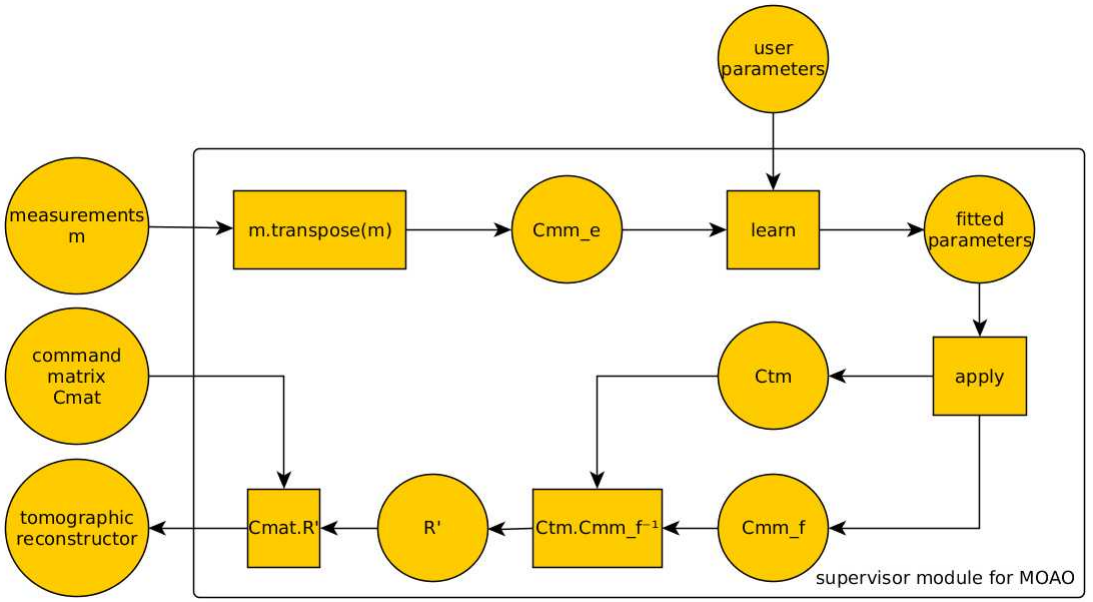

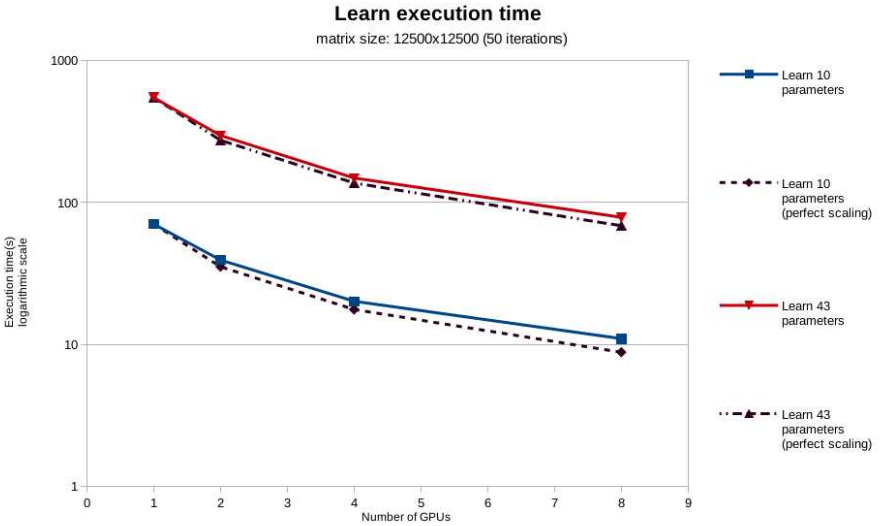

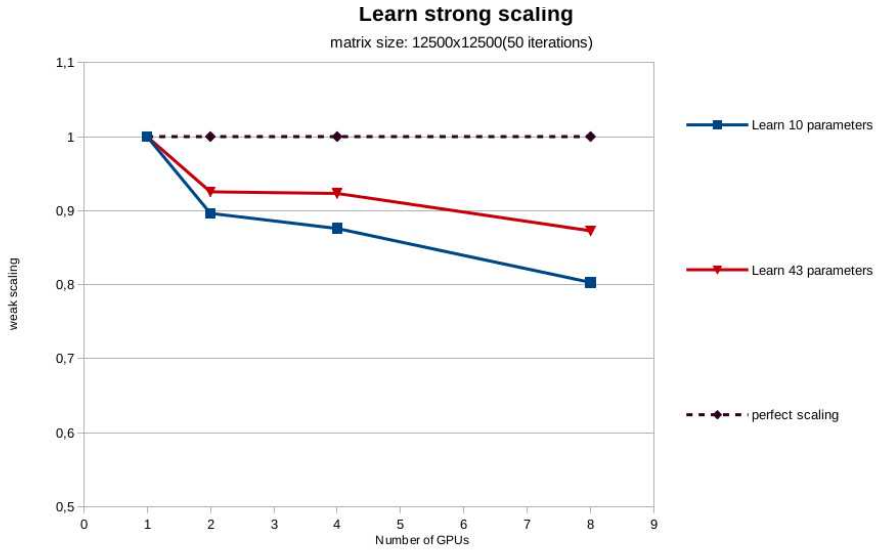

- See talk by Nicolas Doucet
- Collaboration LESIA + KAUST (KSA)

AO RTC concept

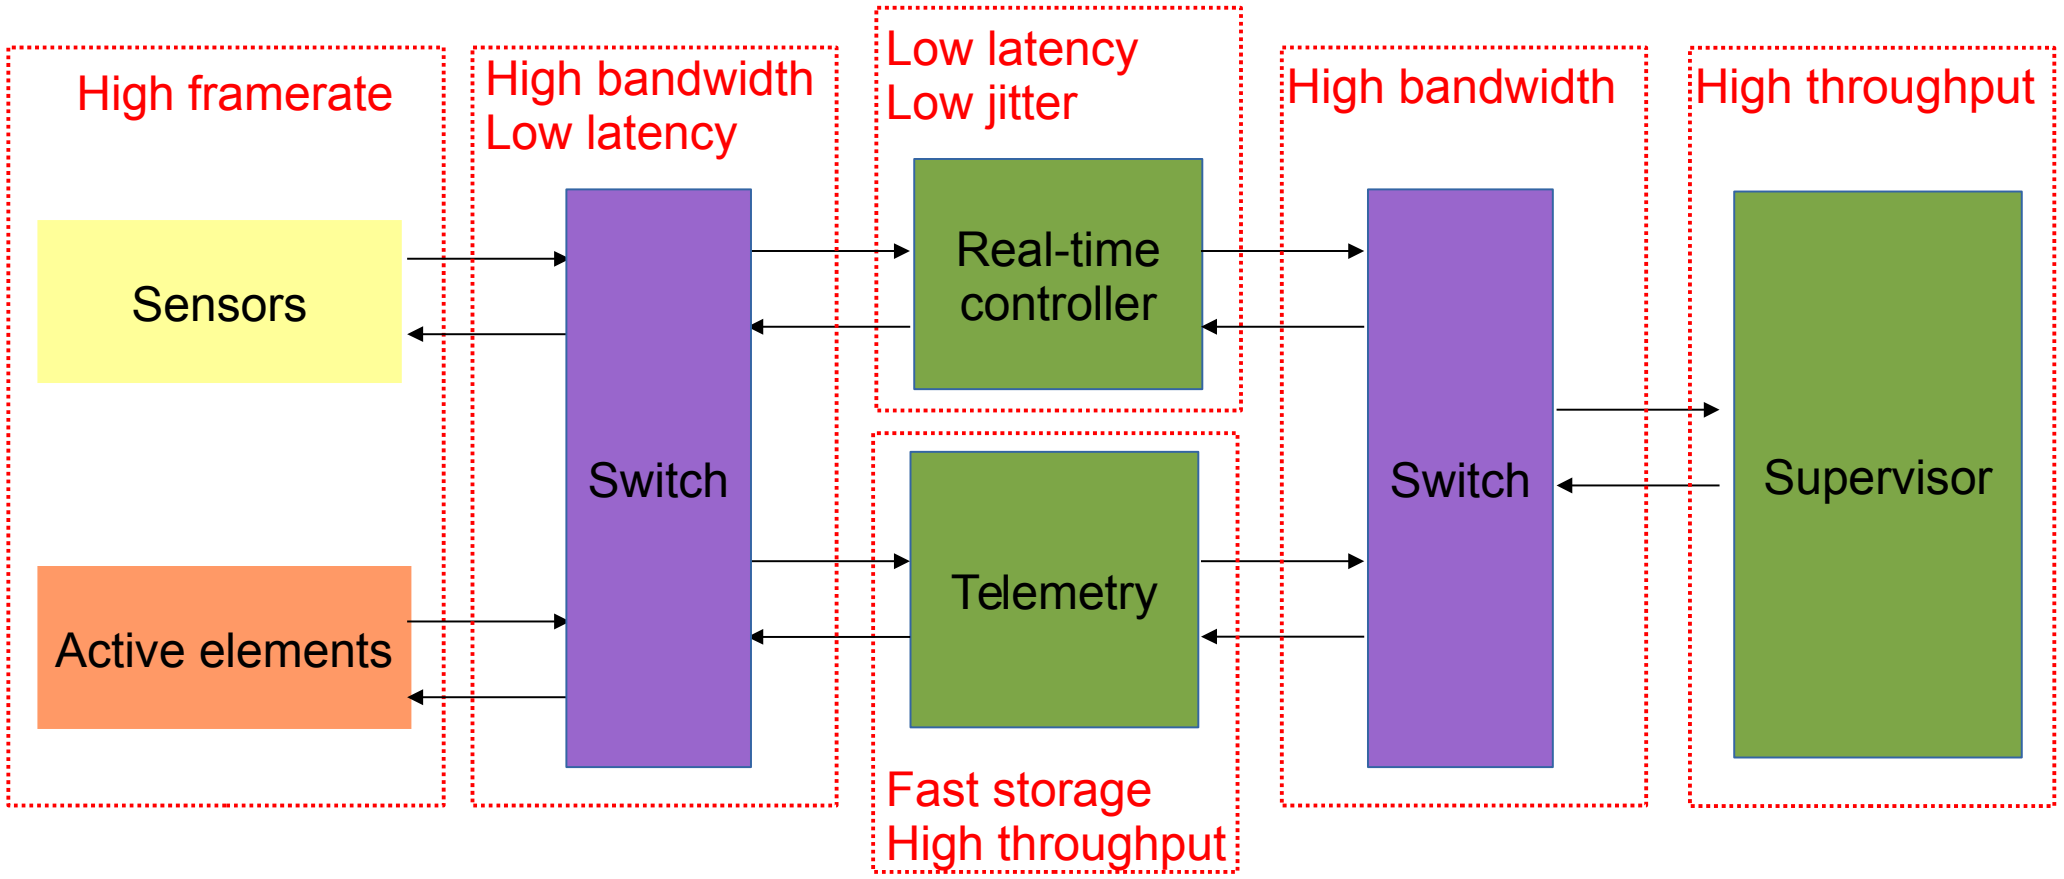

AO RTC concept : SW & MW

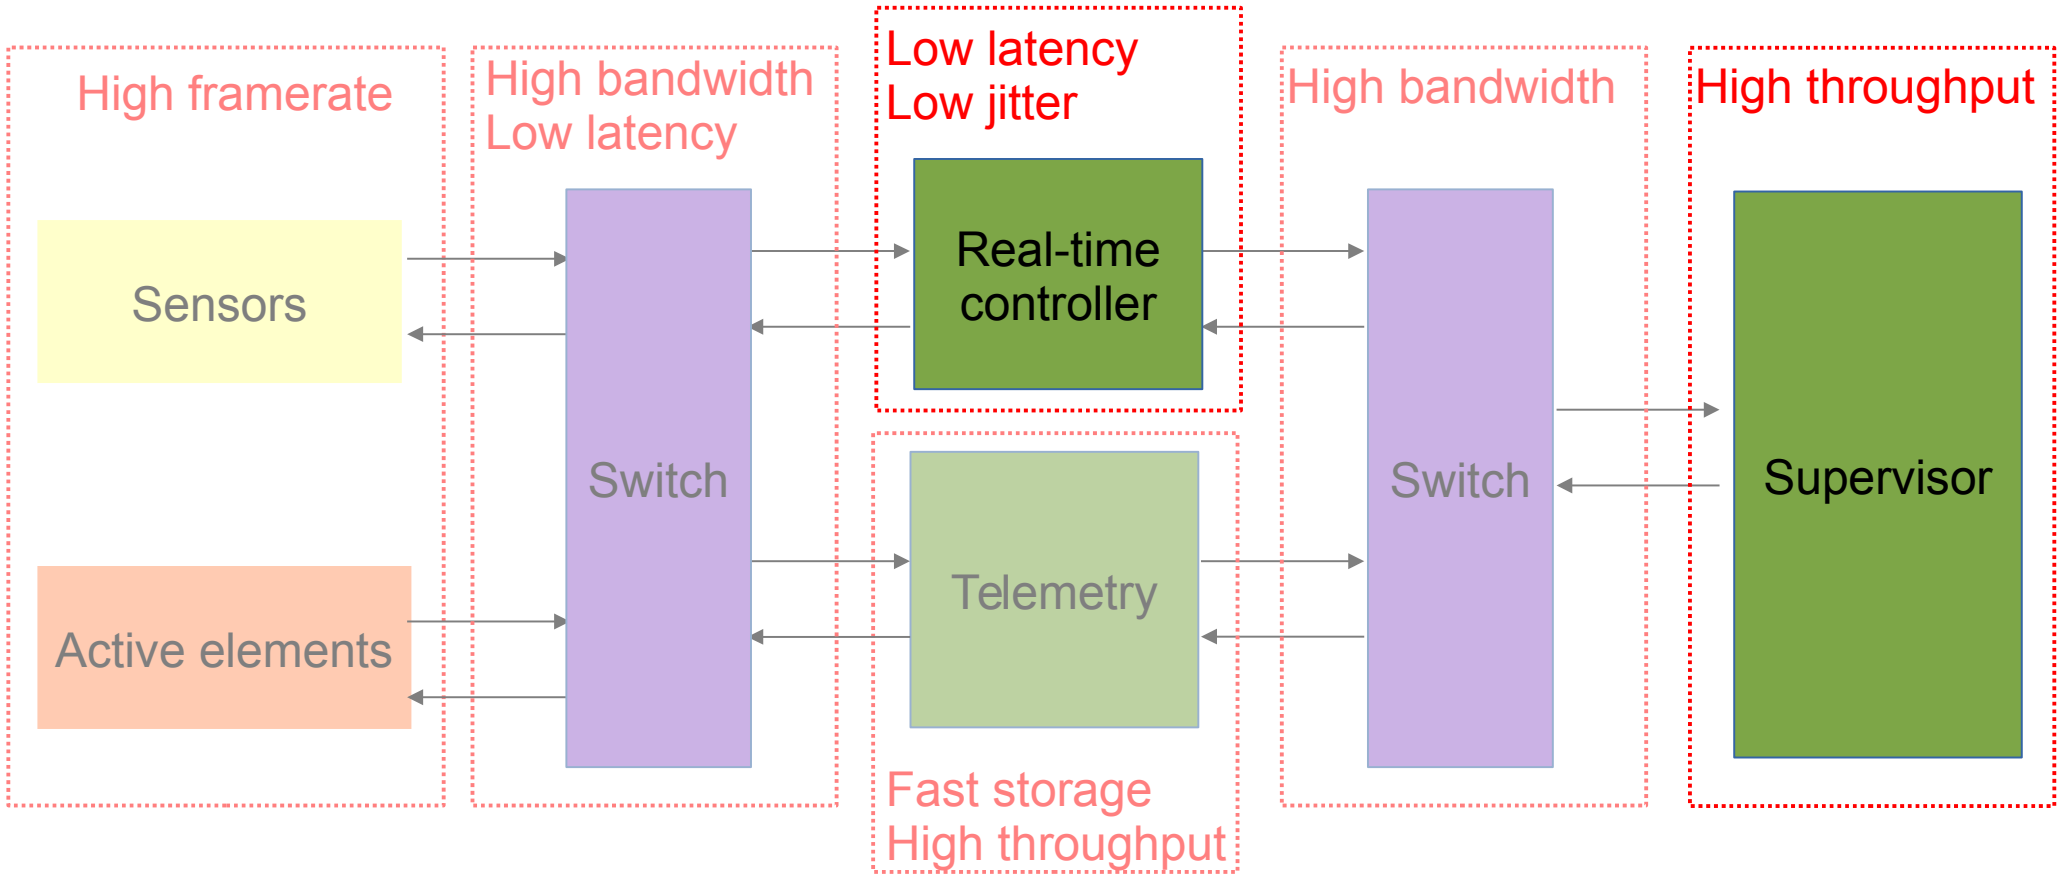

SW / MW stack

- Under development in Paris & Durham
- Run a standard HPC ecosystem on the prototype boards and clusters
- Performance assessed w/r AO application (mainly linear algebra)

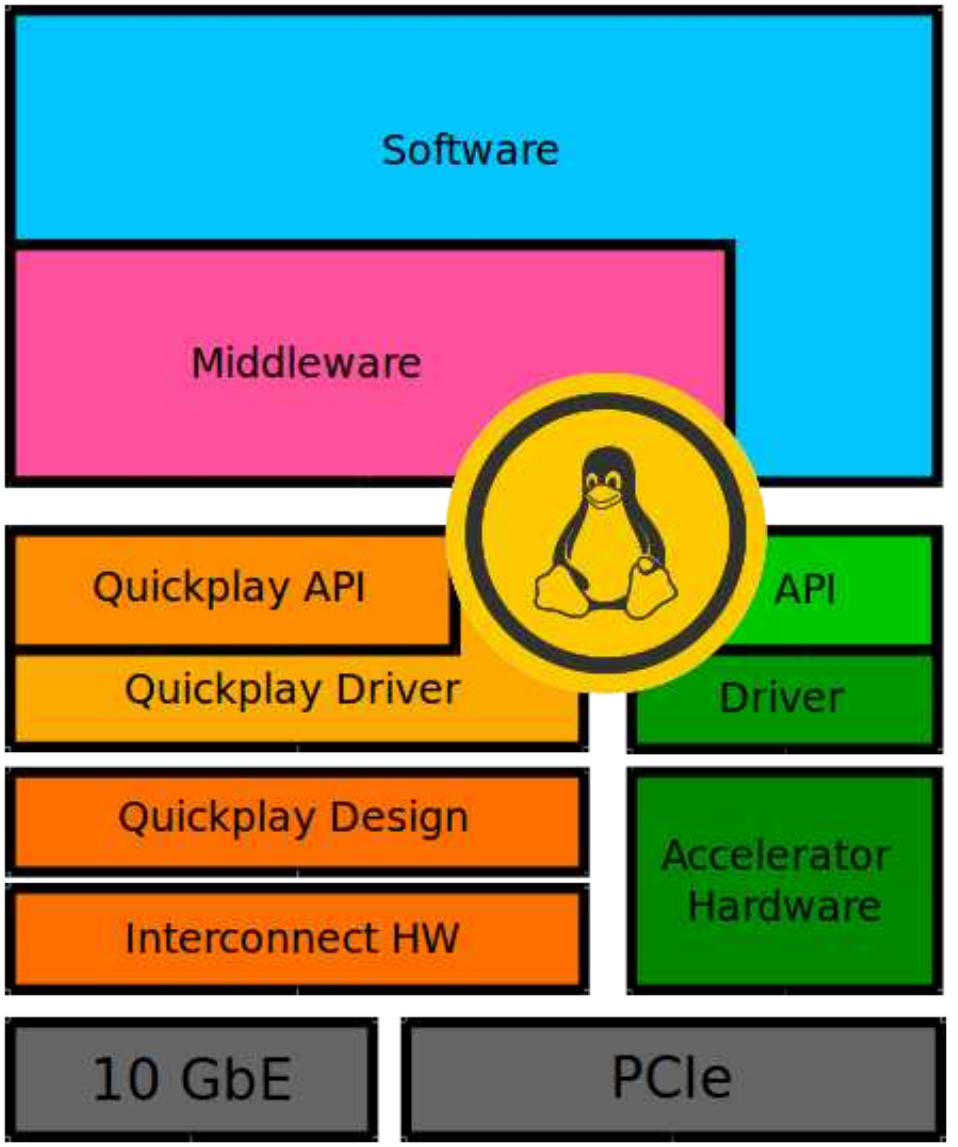

Middleware options

Role	Option 1	Option 2
Data Pipeline	Zero-MQ	MPI
Command/Control	Zero-MQ	ICE
Telemetry	DDS	ICE

- Under study in Durham
- See Eddy Younge talk on wednesda

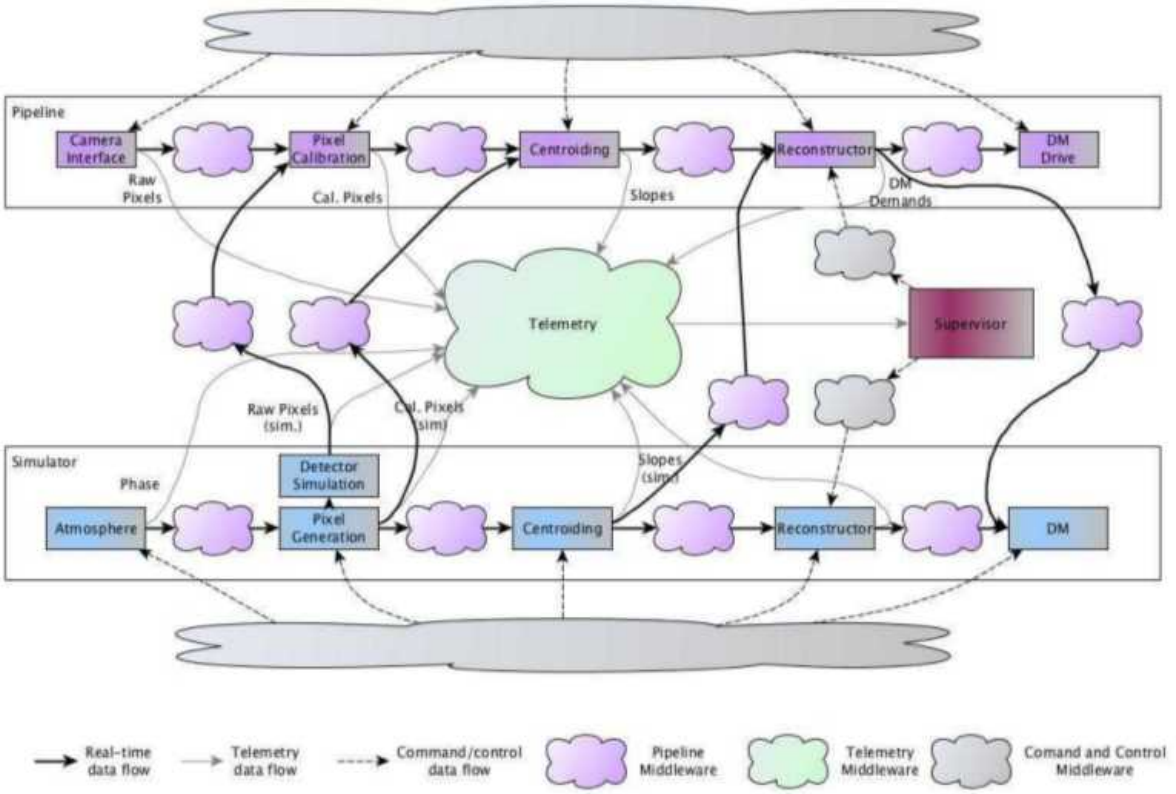

Summary

Project on track

- PDR occurred in Jan. 2016 with feedback from community
- Prototyping activities have started with preliminary results (see per WP presentation)
- Prototyping mid-term review scheduled on Feb. 1rst 2017

Collaborations initiated

- Good feedback from the community on different aspects (HPC + instrumentation)
- Evaluate the convergence and minimize additional effort

More work needed on dissemination

- Populating public website
- Contributions to international conference and publications

Already enhancing the readiness level of commercial solutions

- Contribution to QuickPlay development
- Design of an innovative FPGA board (see Roberto's presentation)