The COSMIC RTC platform

Damien Gratadour

Paris 26 Oct. 2018

AO real-time control

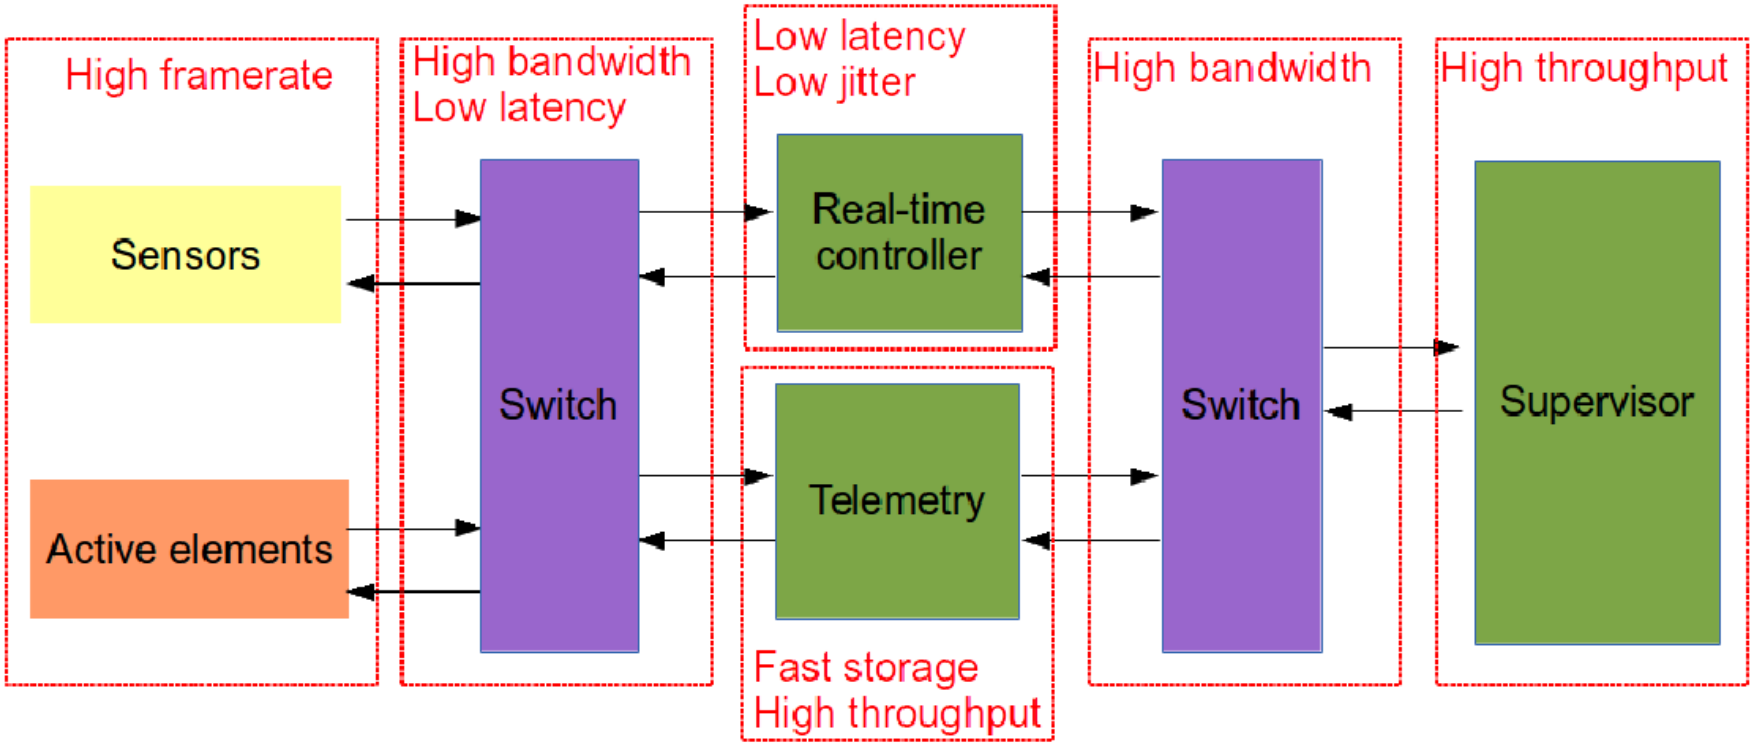

PSL UNIVERSITÉ PARIS
Australian National University
l'Observatoire de Paris

AO RTC

- **From WFS to DM**
  - **Acquire streaming data in real-time** from sensors with low latency
  - **Achieve deterministic computing throughput** on a real-time pipeline including arithmetics and linear algebra
  - **Send commands** to actuators
  - **Supervise** the real-time pipeline with optimisation tasks
  - **Broadcast** telemetry data to clients
  - **Store** raw and processed data
  - Provide AO performance estimation

PSL UNIVERSITÉ PARIS | Australian National University | l'Observatoire de Paris

AO RTC

- **From WFS to DM: performance indicators**
  - End-to-end latency
  - Performance stability (jitter)
  - Data transfer rate for real-time pipeline
  - Compute throughput for real-time pipeline
  - Data transfer rate for telemetry
  - Compute throughput for supervisor / optimizer
  - Data storage capacity

PSL UNIVERSITÉ PARIS | Australian National University | l'Observatoire de Paris

COSMIC RTC platform

- **Global system architecture**

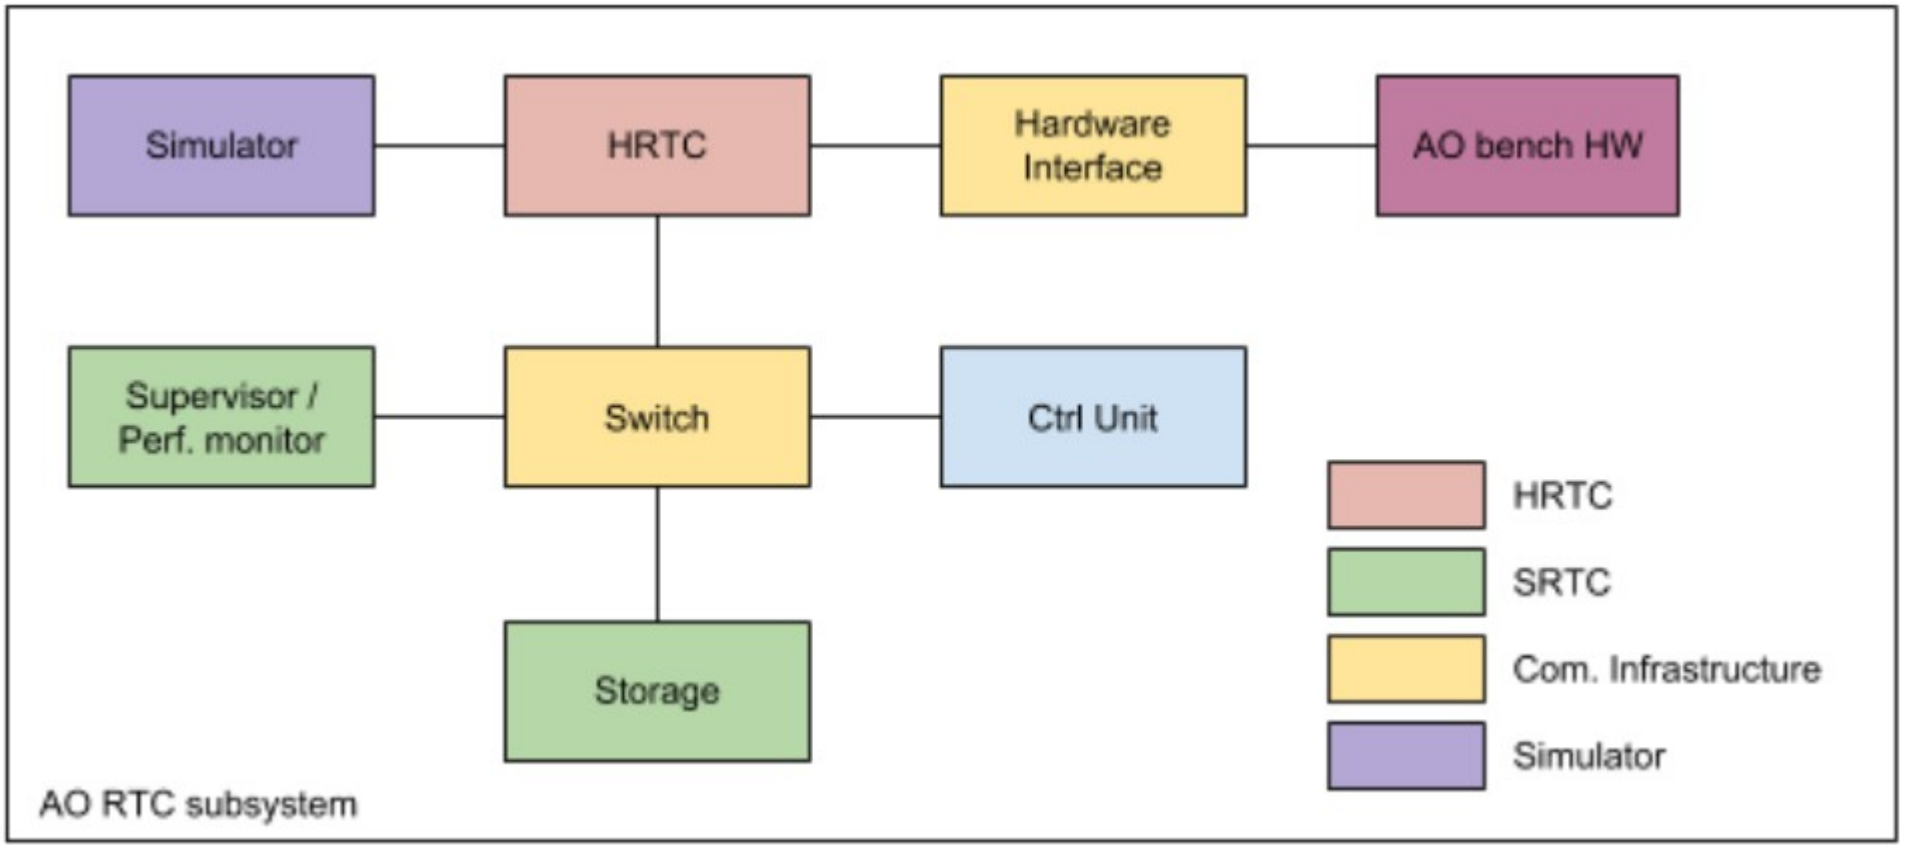

PSL UNIVERSITÉ PARIS
Australian National University
l'Observatoire de Paris

COSMIC RTC platform

- **Global system architecture**

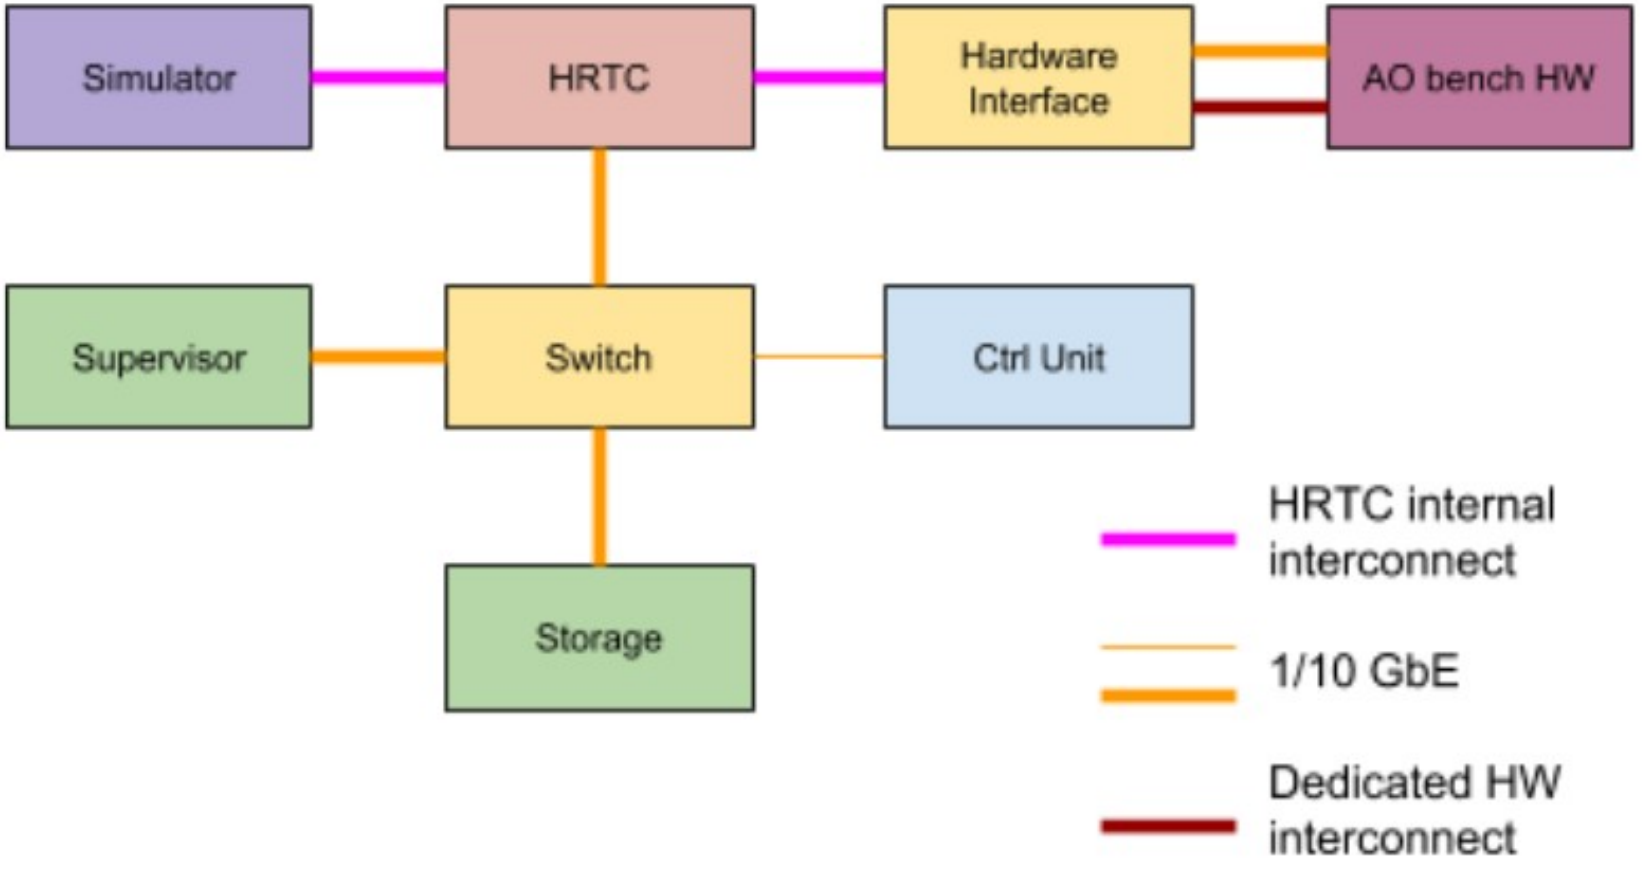

COSMIC RTC platform

- **RTC SW architecture**
  - Vertical and horizontal interactions

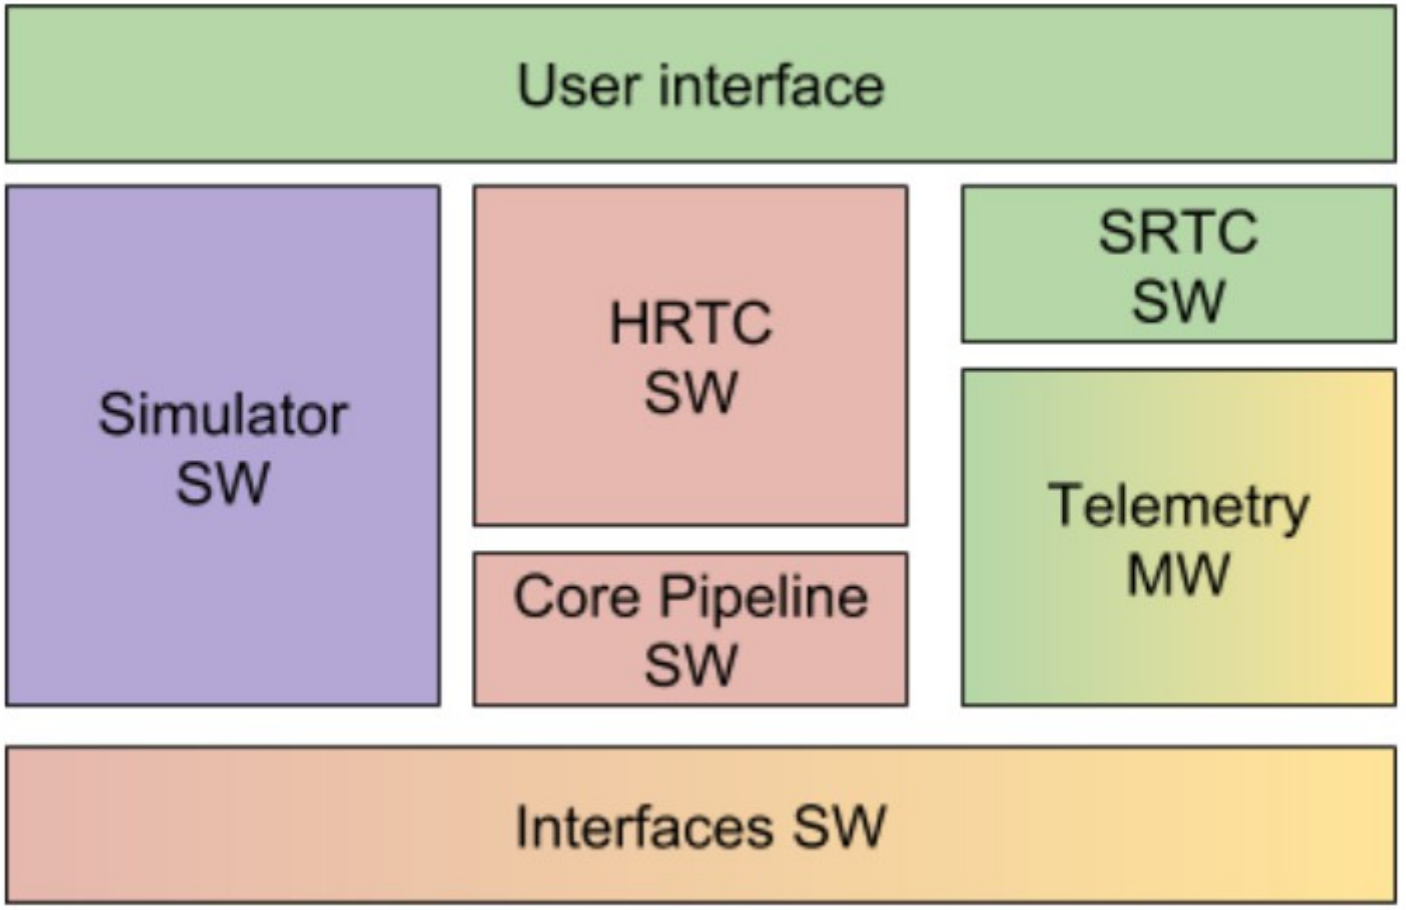

COSMIC RTC platform

- **HRTC HW architecture**

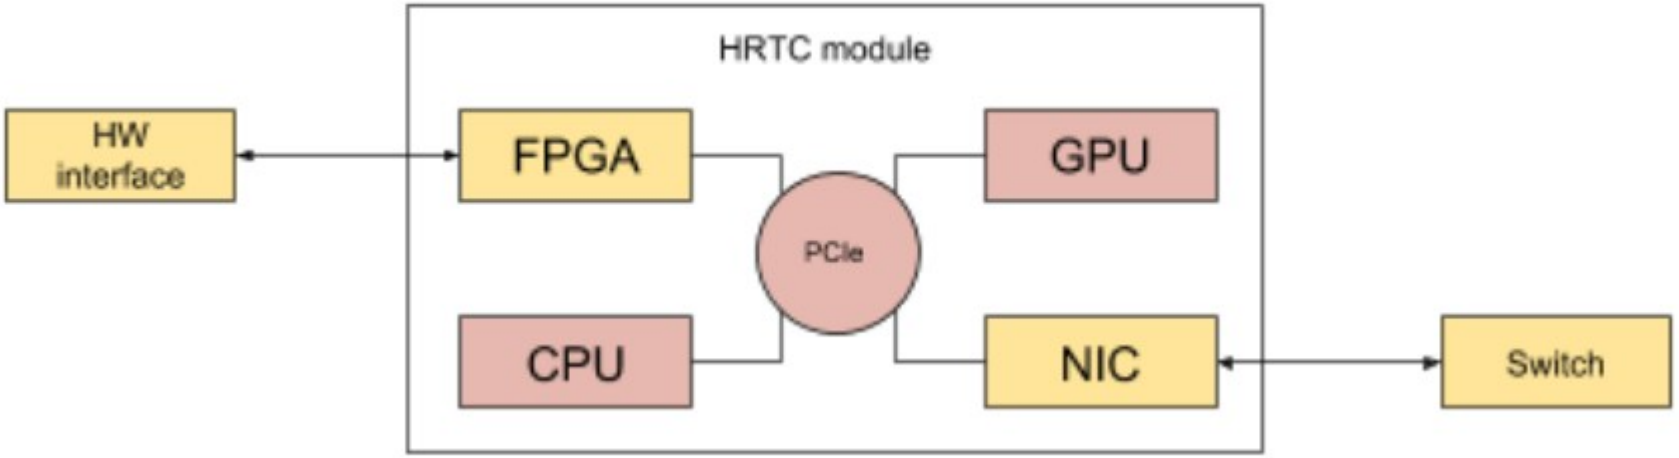

COSMIC RTC platform

- **HRTC HW architecture**

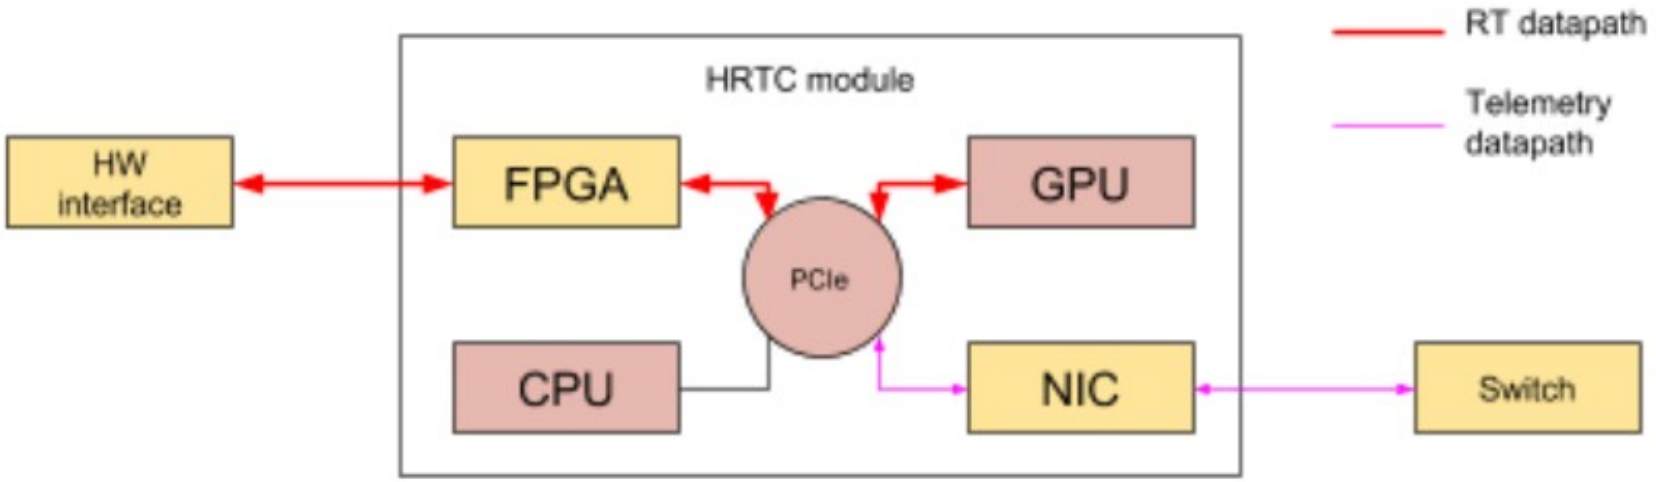

COSMIC RTC platform

- **HRTC SW architecture**

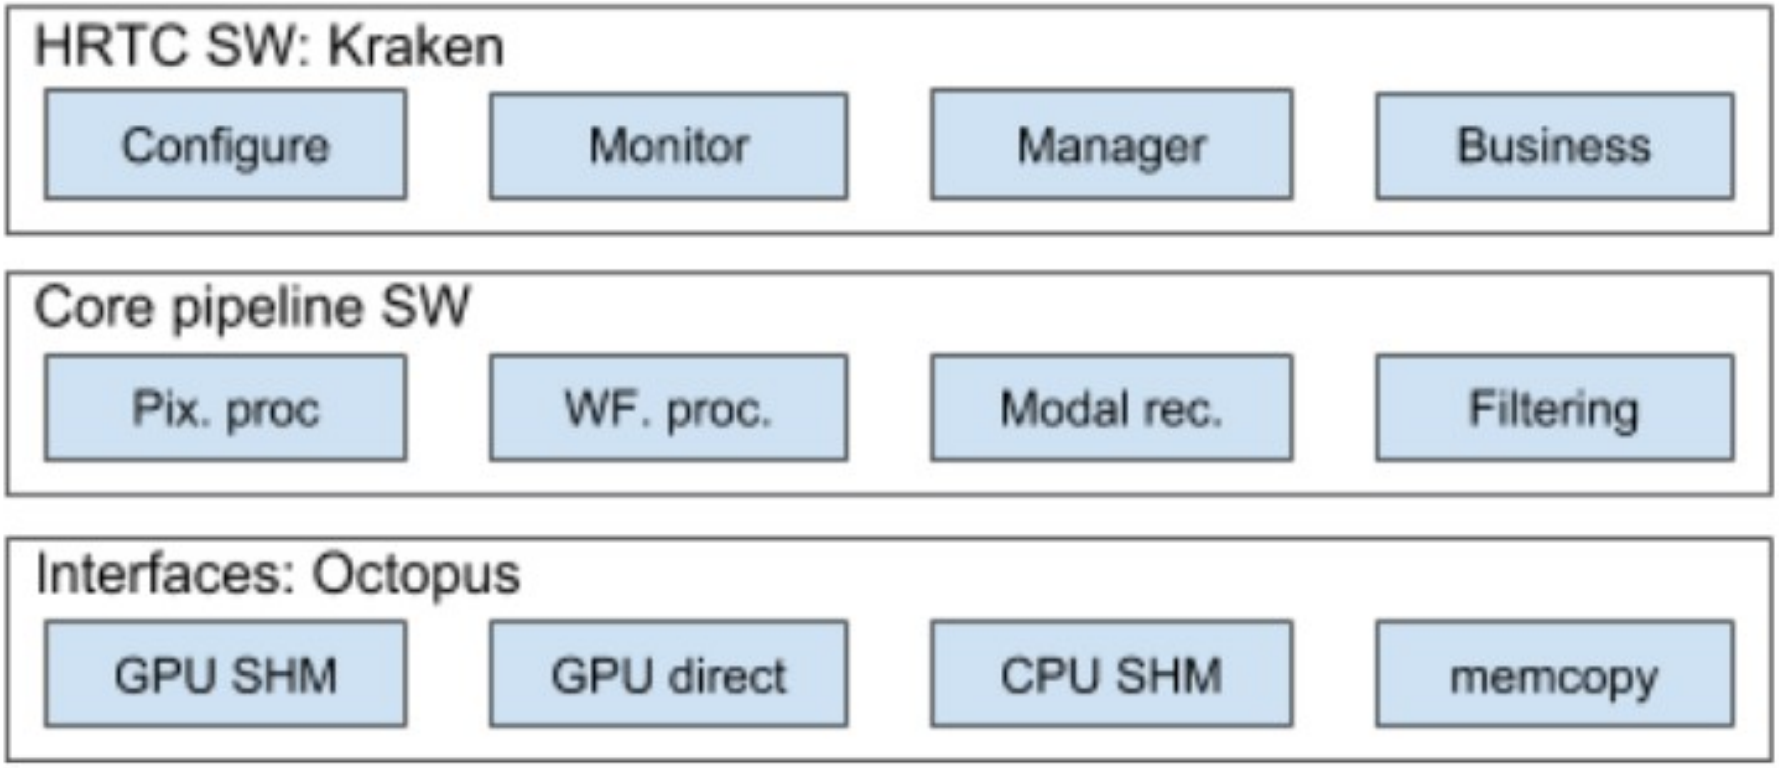

PSL UNIVERSITÉ PARIS | Australian National University | l'Observatoire de Paris

COSMIC RTC platform

- **HRTC SW design**

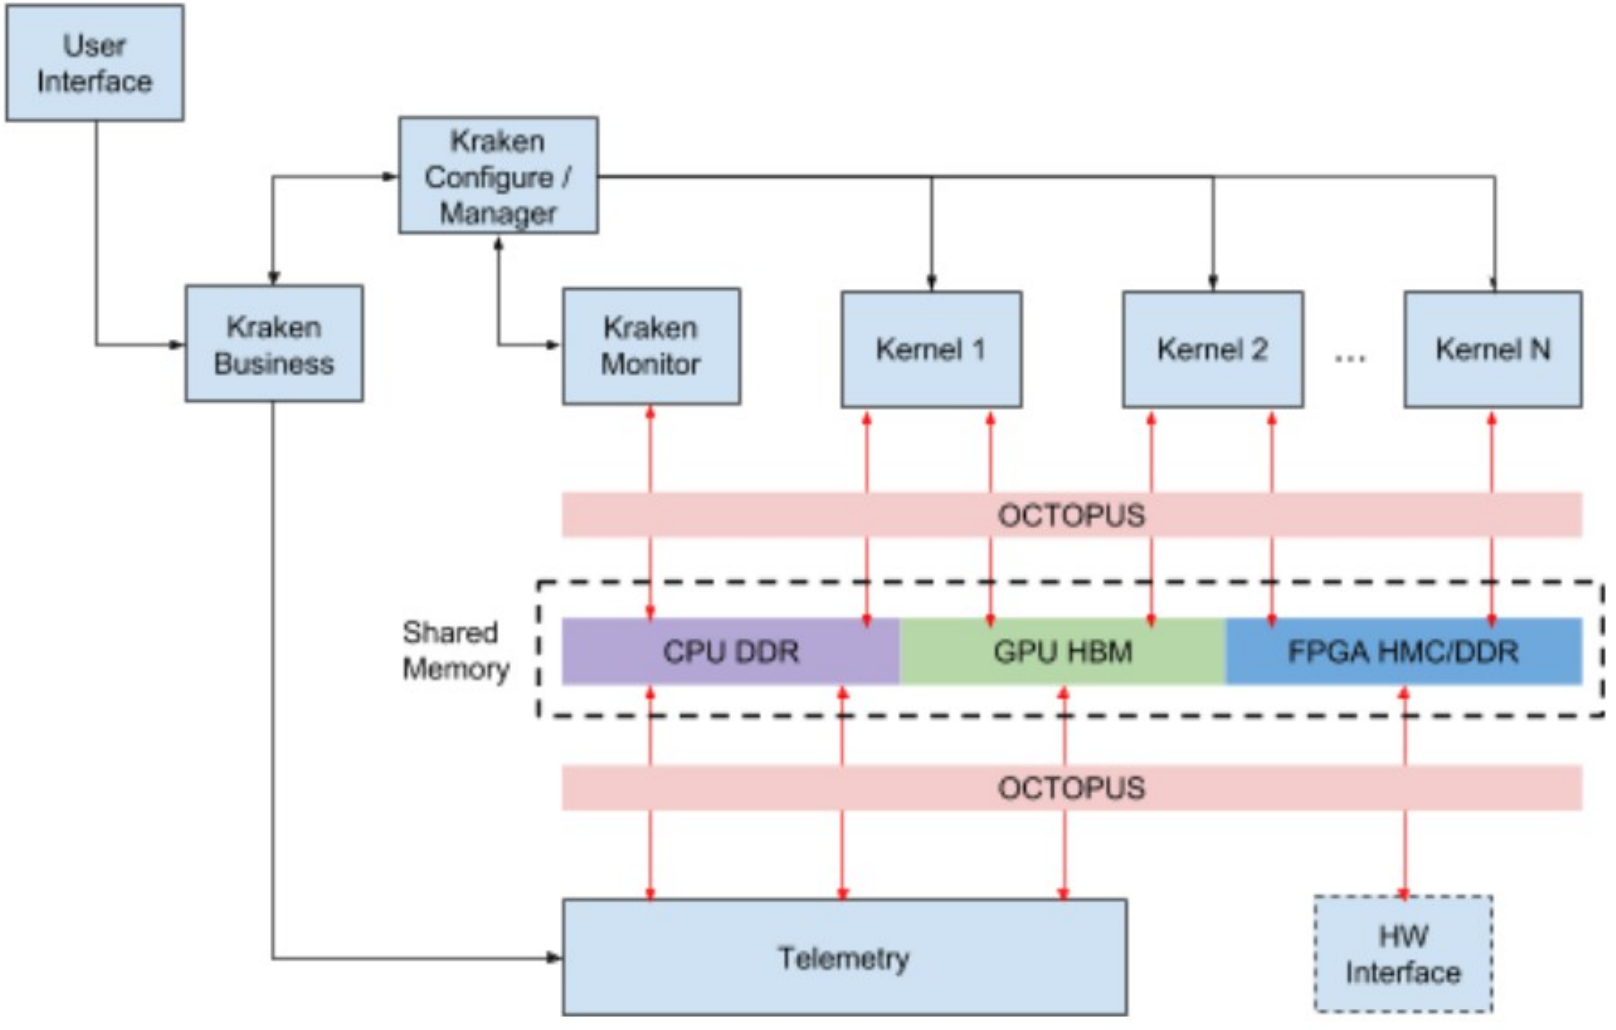

COSMIC RTC platform

- **SRTC SW design**

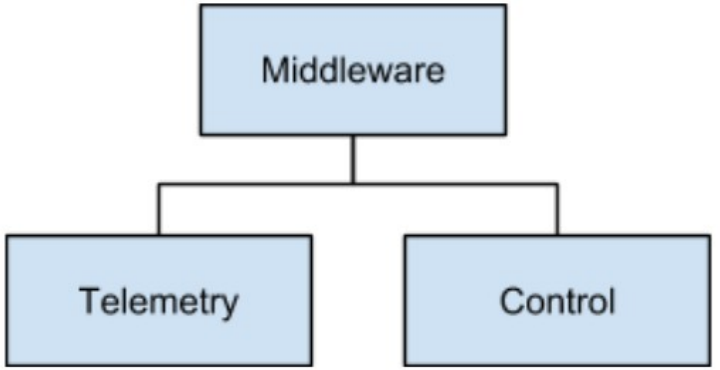

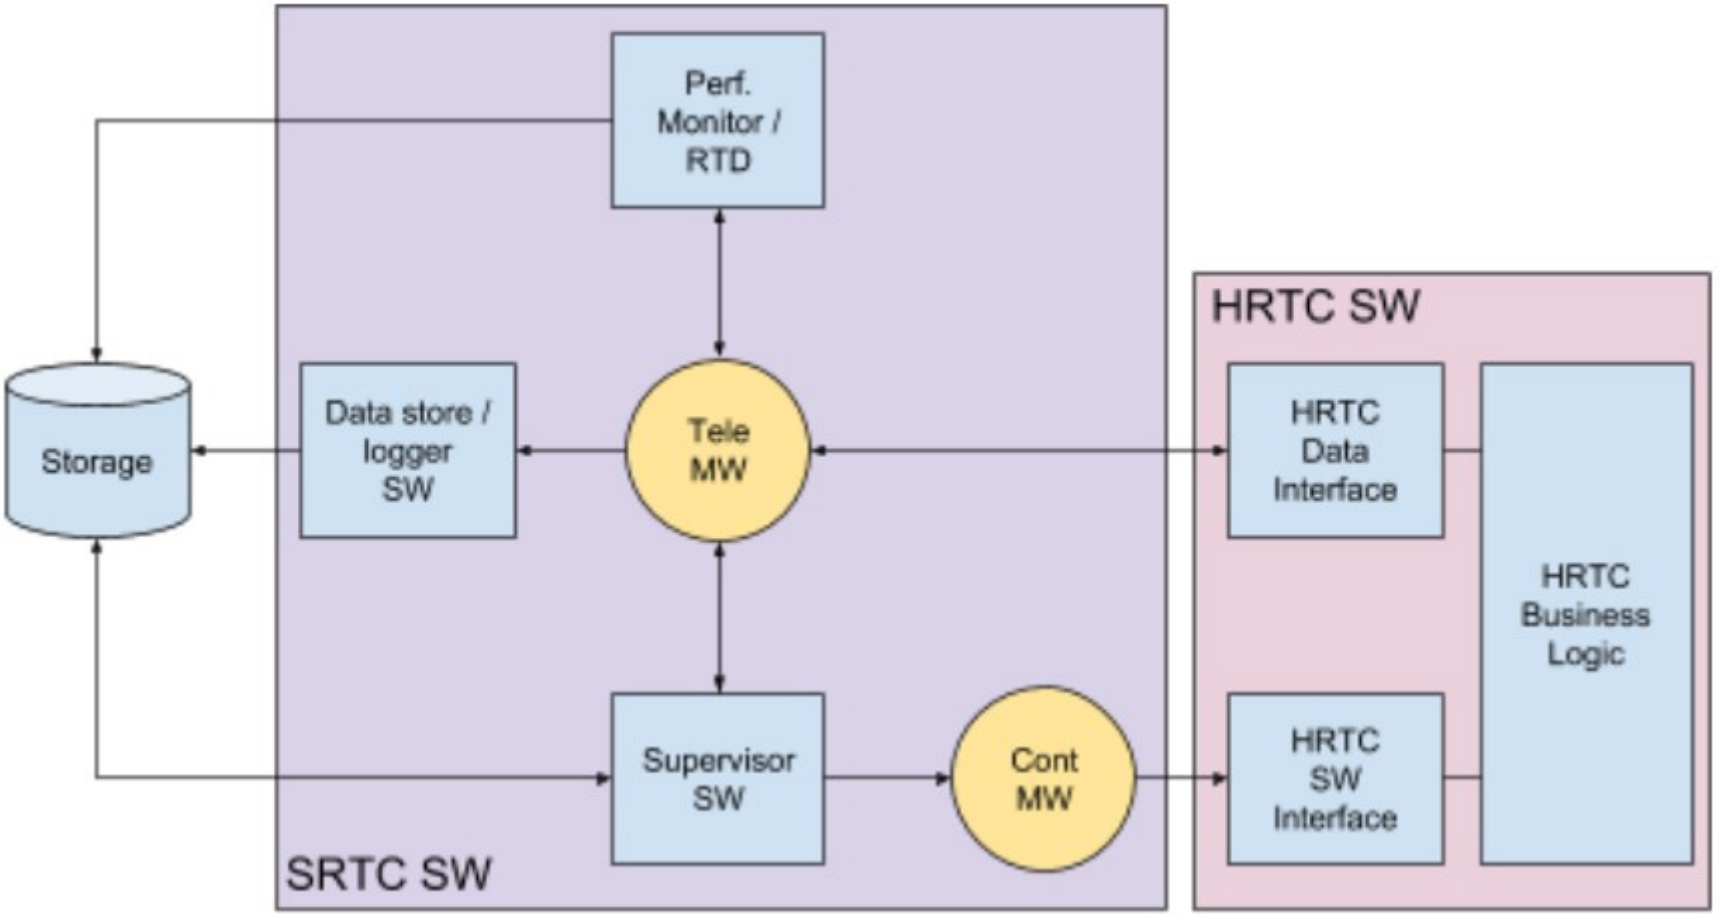

COSMIC RTC platform

- **HRTC / SRTC GPU components**
  - several options could be considered:
    - use high end General Purpose GPUs (e.g. Nvidia Tesla line of products)
    - use high end gaming GPUs (e.g. Nvidia GeForce line of products)
  - HRTC pipeline is dominated by memory bound GDDR5 (gaming) versus HBM2 (GPGPU).
  - Longevity : active cooling (gaming) versus passive cooling (GPGPU)
    - The main problem with getting MTBF values for gaming cards is the really wide range of cooling solutions and clock rates.
  - Ecosystem: GPGPU cards are optimized for cluster usage, including full support for InfiniBand and RDMA, vendor support and cluster management tools

PSL UNIVERSITÉ PARIS | Australian National University | l'Observatoire de Paris

COSMIC components

- **HW interface to provide flexibility**
  - µXComp board from microgate
  - Aria 10 FPGA + HMC + expendable interface

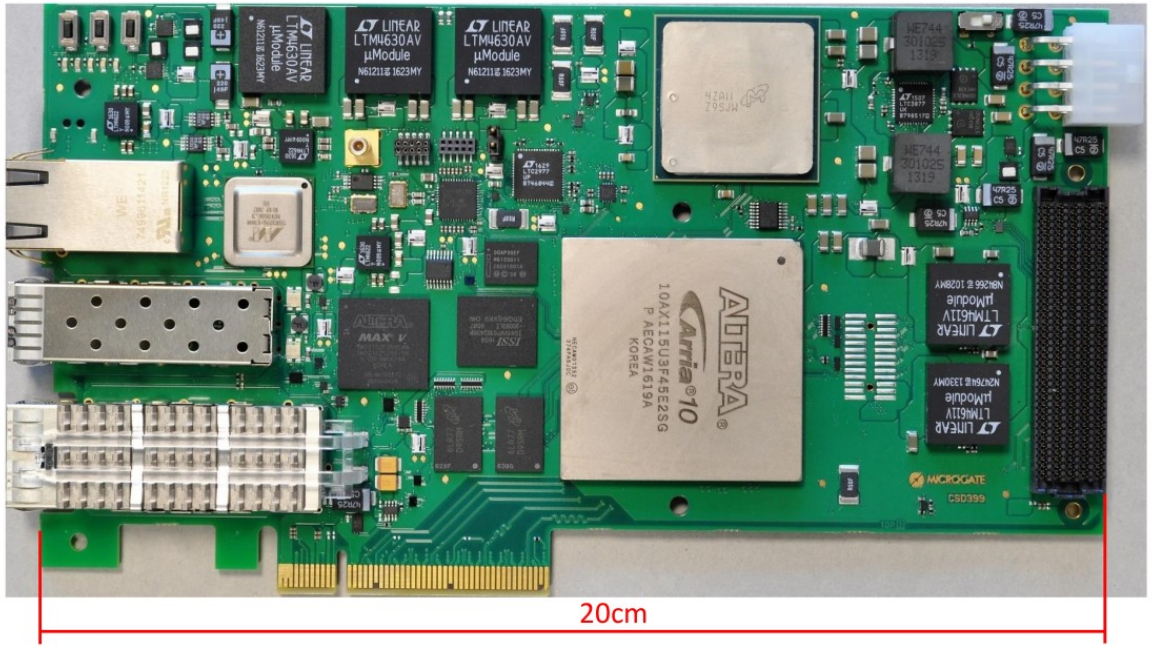

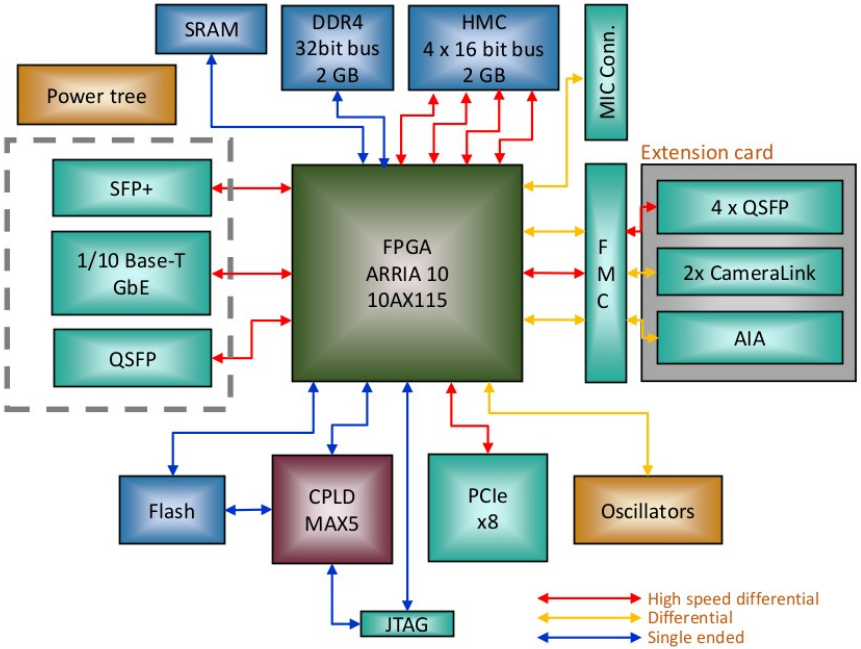

ılian
ıal
sity

l'Observatoire de Paris

COSMIC components

- **Persistent kernels to guarantee time predictable performance**

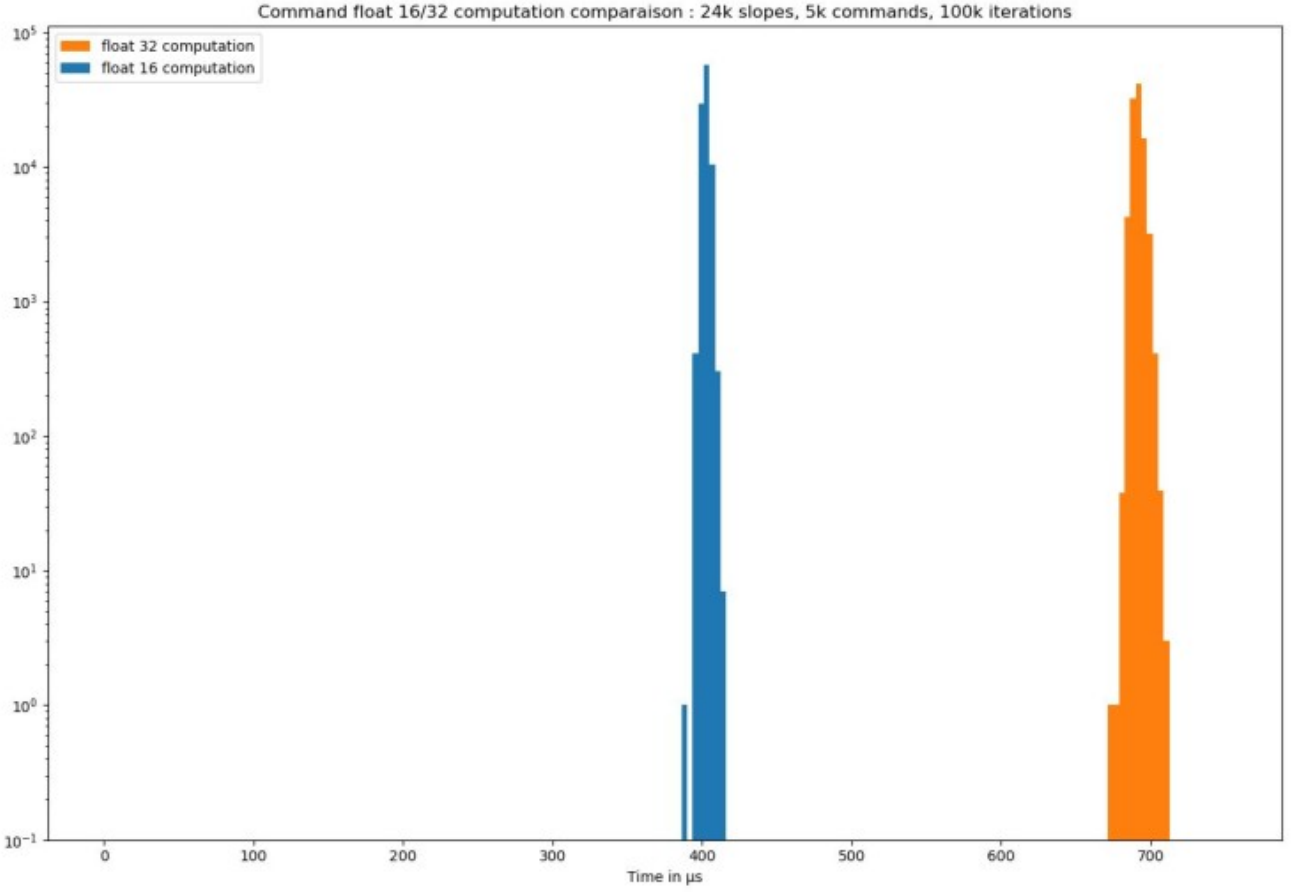

COSMIC components

- **GPUdirect to minimize overall latency budget**

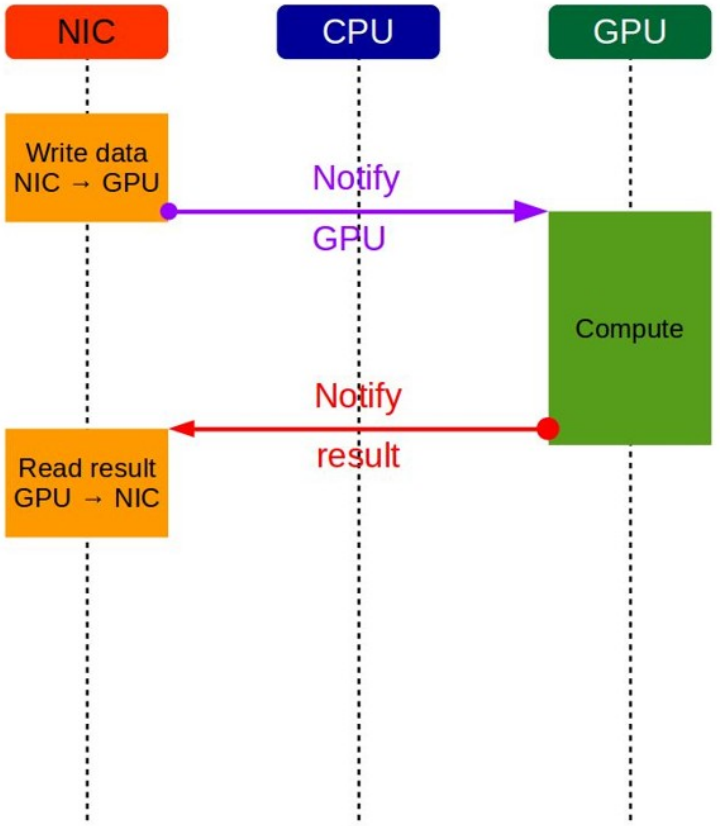

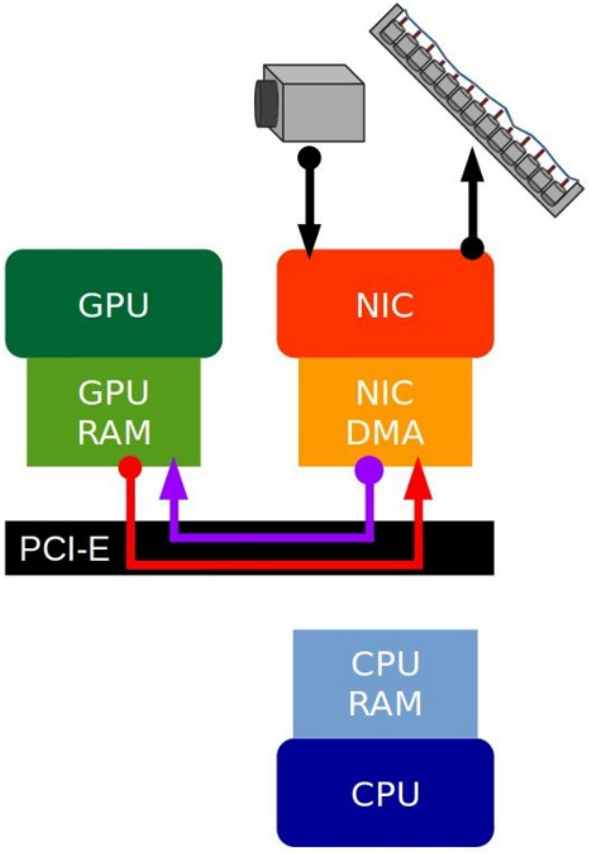

COSMIC components

- **HW and SW simulators to check correctness and performance predictability**

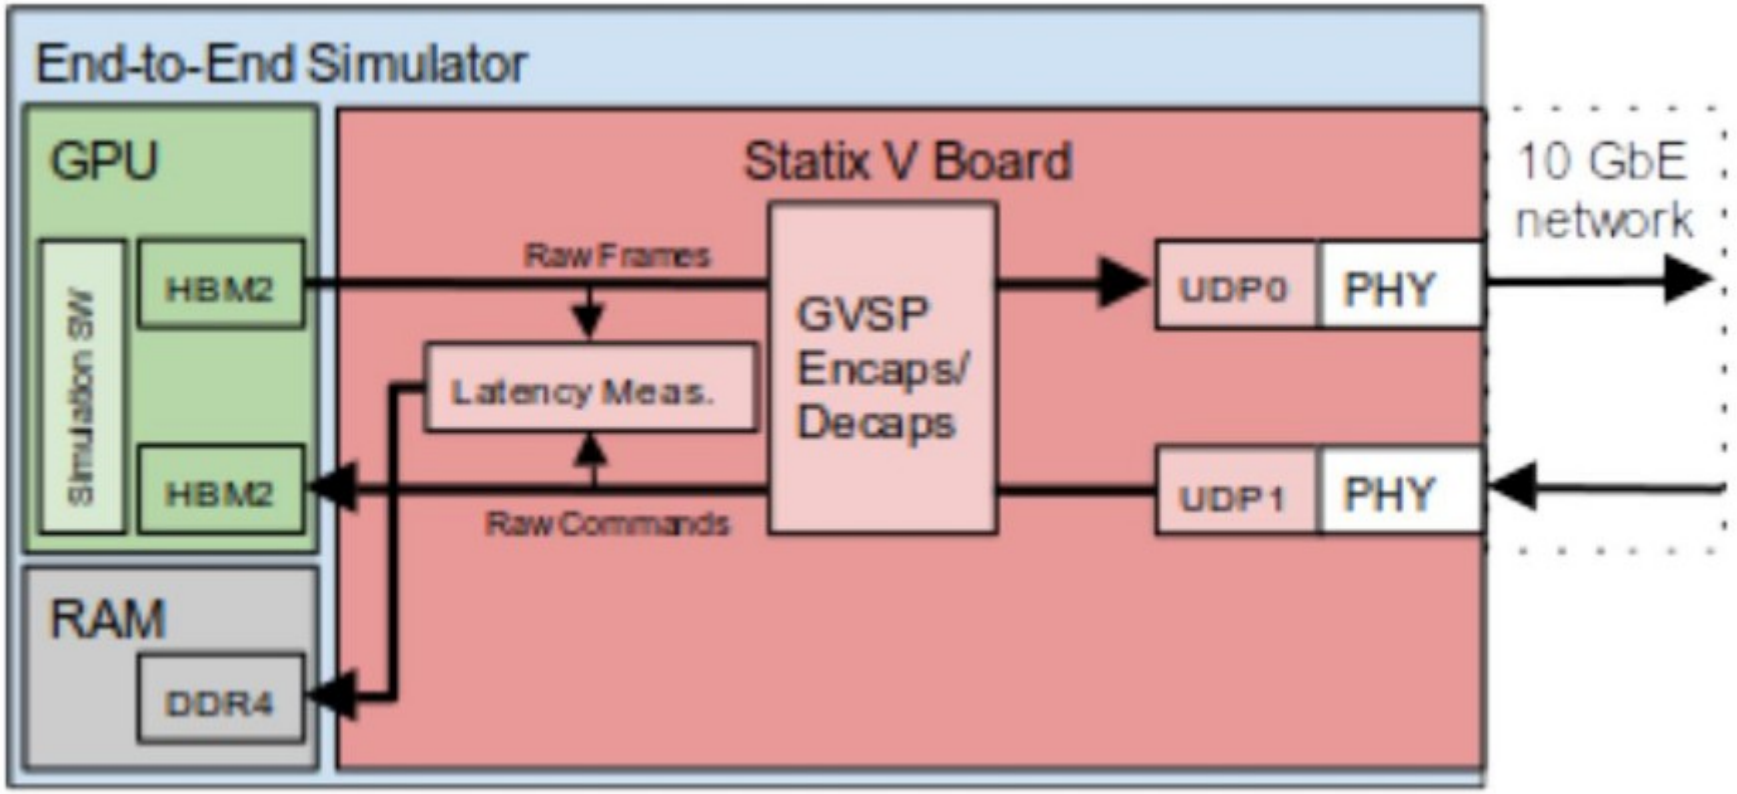

PSL UNIVERSITÉ PARIS | Australian National University | l'Observatoire de Paris

COSMIC prototype

- **Prototype architecture**

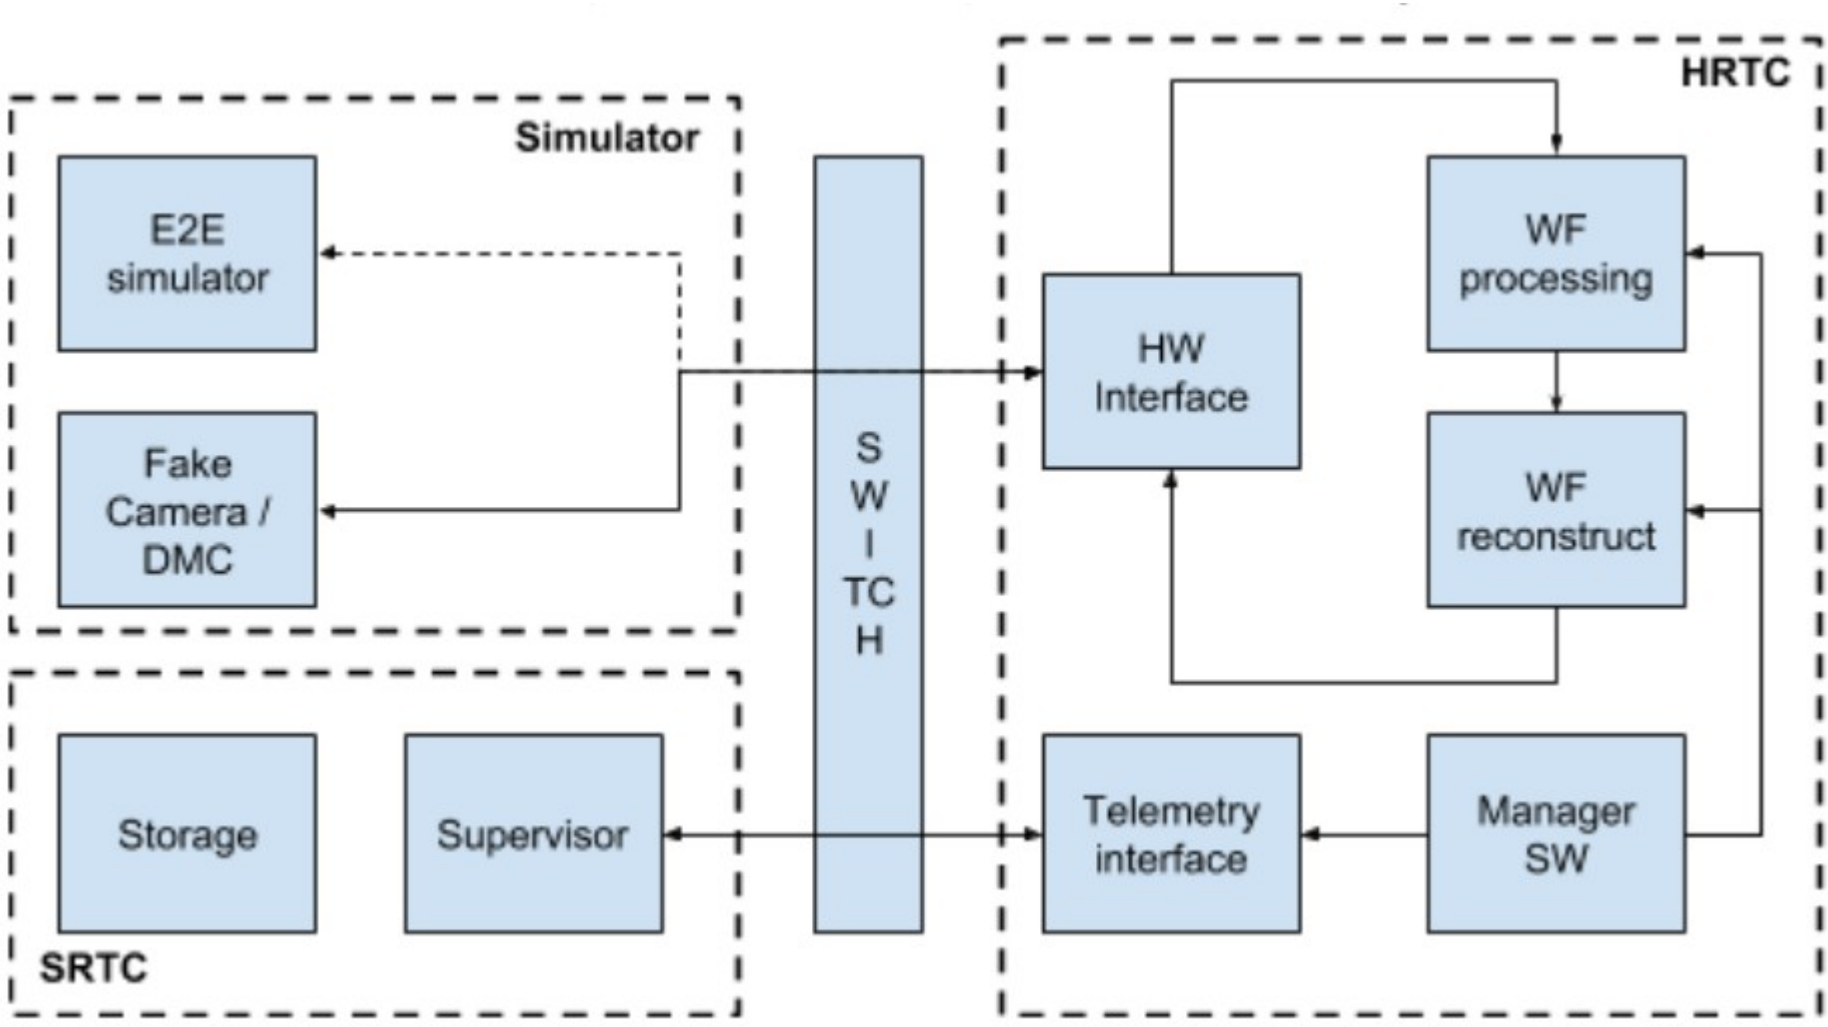

COSMIC prototype

- **Prototype implementation**

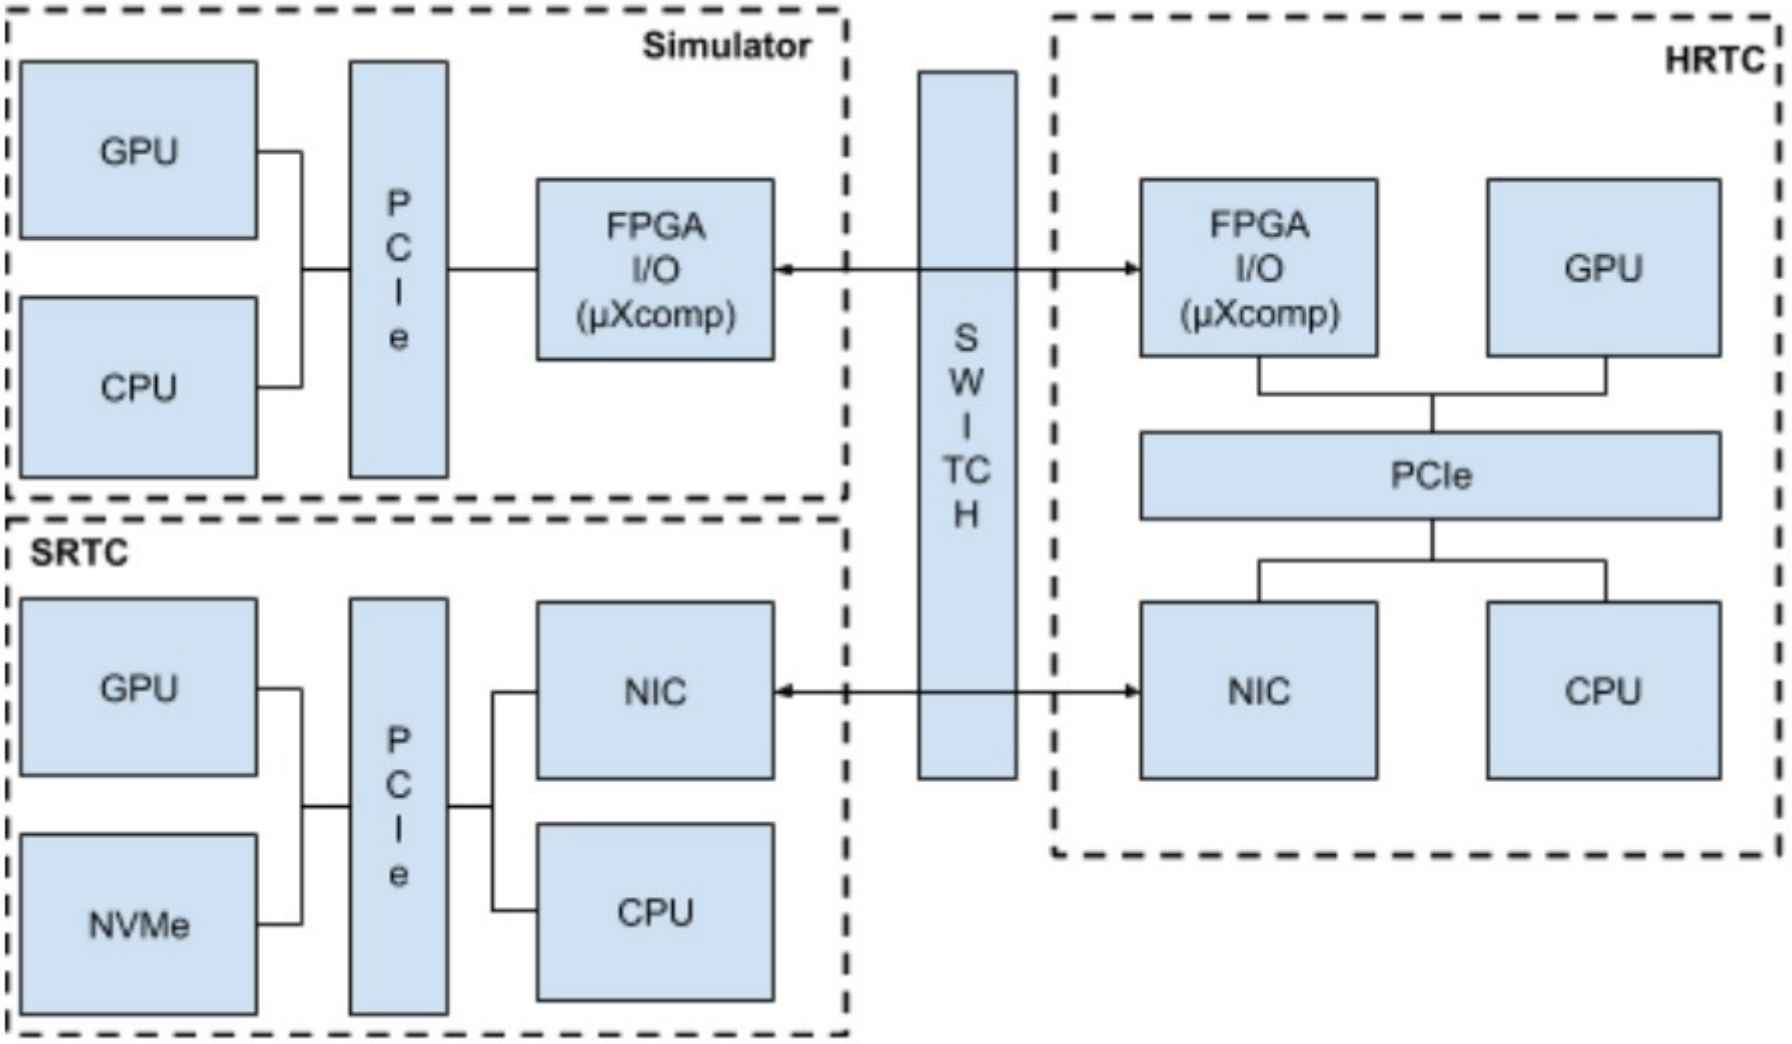

COSMIC partnership

- **Academia-industry partnership**
- **3 academic partners**
  - Observatoire de Paris (France)
  - Australian National University (Canberra)
  - Swinburne University of Technology (Melbourne)
- **1 industrial partner**
  - Microgate (Italy)
- Well aligned with PSL-ANU strategic partnership as well as new status of Australia at ESO
- Partners already engaged in Keck RTC upgrade under the responsibility of Microgate

PSL UNIVERSITÉ PARIS | Australian National University | l'Observatoire de Paris

COSMIC collaboration

- **Relying on international collaborations**
- **Collaboration with HPC research centers**
  - Extreme Computing Research Center (ECRC) @KAUST on optimized linear algebra
  - Barcelona Supercomputing Center on programming models for real-time applications
- **Collaboration with Subaru telescope**
  - Increasing the readiness level of the CACAO SW stack
  - Access to the SCExAO instrument to perform on-sky tests of SW / HW solutions
- **Collaboration with industry**
  - Long going collaboration with NVIDIA
  - Ongoing discussions with Thales for leveraging our solutions on other applications
- **Welcoming new collaborations !**

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology
Extreme Computing Research Center

BSC

SUBARU TELESCOPE
National Astronomical Observatory of Japan

NVIDIA

THALES

PSL UNIVERSITÉ PARIS

Australian National University

l'Observatoire de Paris

Conclusion

- **Leveraging Green Flash solutions**
  - Ensure best performance at the lowest cost
  - Modular, upgradable
- **Real-time by design**
  - No need to “tweak” things thanks to a “recipe” of HW and SW components
  - Can be reproduced easily (as soon as you know the recipe)
  - Applicable to various generations of HW (several generations of GPUs and several FPGA boards and vendors)
- **International partnership**
- **Targeting current and future instrumentation**
  - New Keck RTC is based on these concepts
  - Welcome new opportunities !

PSL UNIVERSITÉ PARIS | Australian National University | l’Observatoire de Paris

Thank you

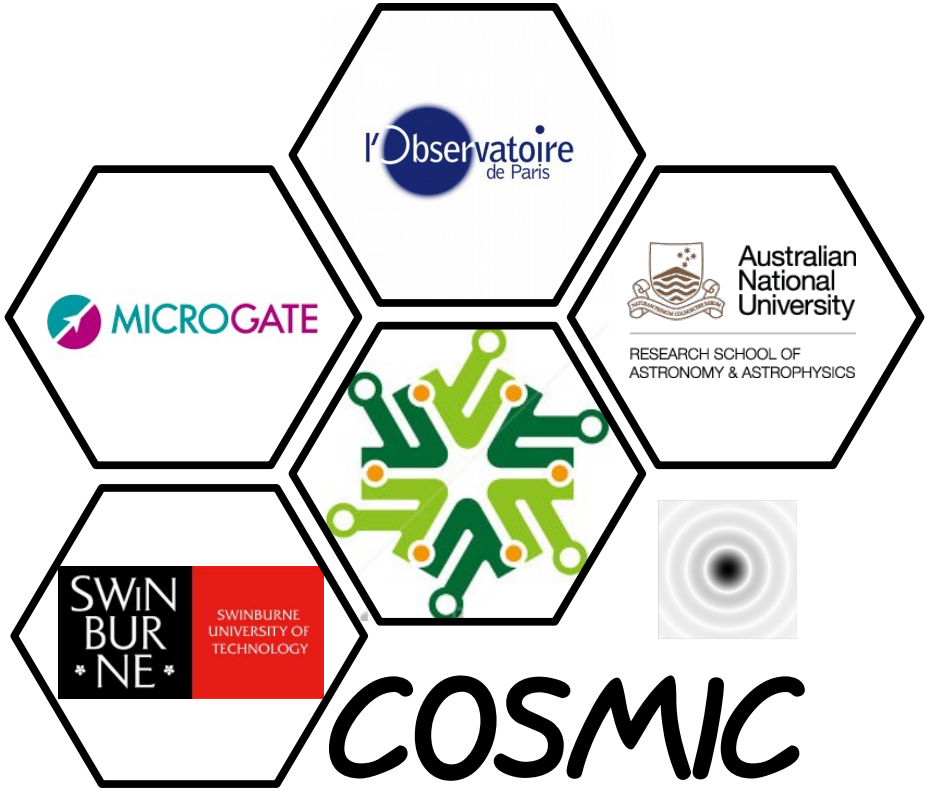

A common open platform for AO RTC