The METIS AO RTC at the beginning of MAIT

Martin Kulas

Max Planck Institute for Astronomy (MPIA)

RTC4AO 6th, Garching, 2023

Outline

- METIS Instrument
- AO RTC Problem Size
- HRTC (Hard Real-Time Core)
  - Hardware
  - Software
  - Performance
  - Numerical Correctness
- SRTC (Soft Real-Time Cluster)

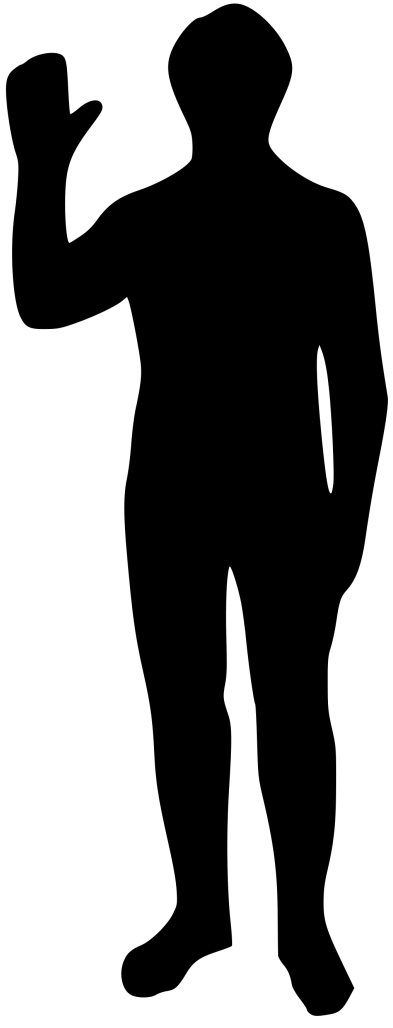

METIS Team

- The METIS instrument is built by an international consortium.
- RTC Team is located at MPIA:
  - Horst Steuer
  - Hugo Coppejans
  - Thomas Bertram
  - Florian Briegel
  - Markus Feldt







METIS Instrument (Mid-infrared ELT Imager and Spectrograph)

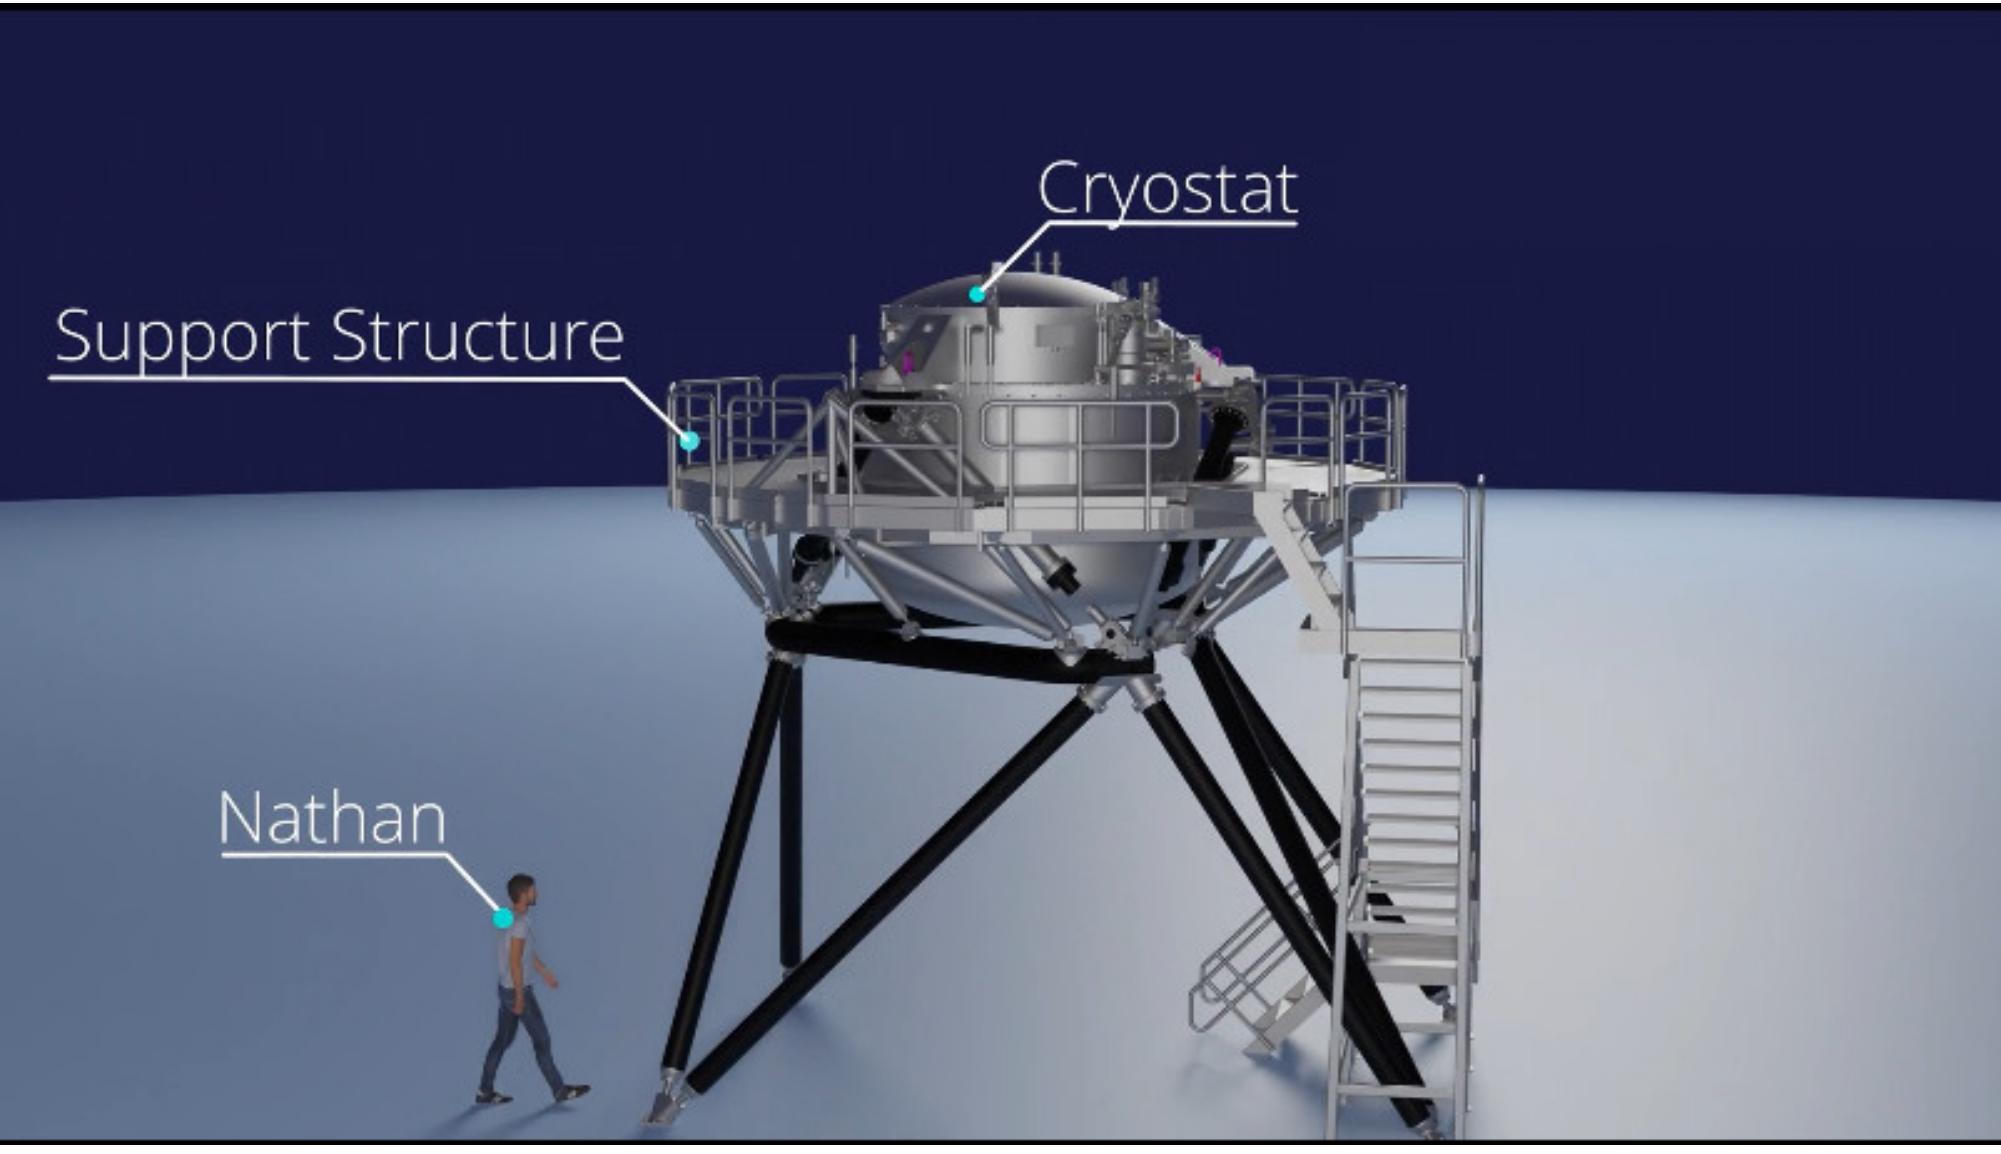

METIS Instrument (as fallback if video fails)

- 1st generation instrument at the ELT (Extremely Large Telescope).
- Science cases:
  - Circumstellar disks and planet formation.
  - Detection and characterisation of extrasolar planets.
  - Star formation and stellar clusters.
  - The physics of galaxies.
  - Evolved stars.
- L, M and N band imaging and spectroscopy.
- High contrast imaging.
- FDR in November 2022.
- First light planned for 2028.

METIS SCAO from a Programmer's Perspective

Property	Size
WFS type	P-WFS with four sides
WFS camera frame rate	1 kHz
WFS camera frame size	192 x 192 pixels as int32
Pupil diameter	90 subapertures
Number of subapertures	6,376
Wavefront (mirror) command	4,868 elements

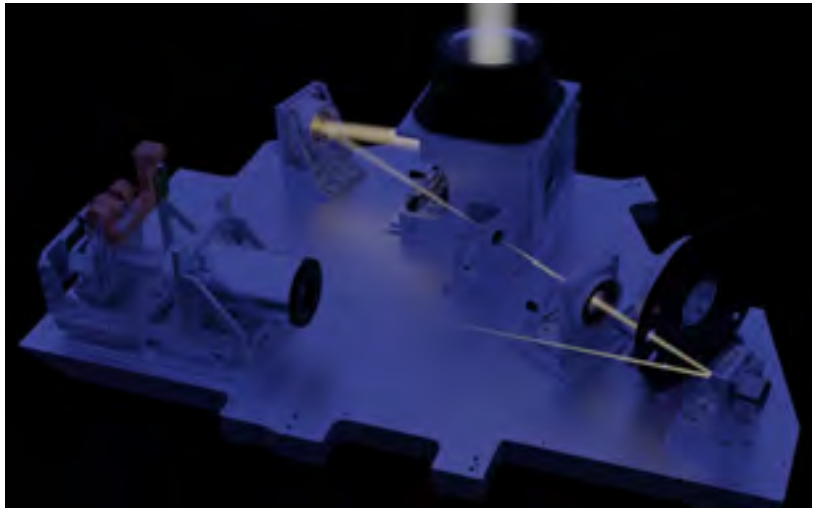

RTC Problem Size

- Main task: Run Wavefront Control Loop (WFC) with hard real-time requirement: RTC computation time limit is 909 µs (elapsed time between first received WFS data by RTC and last transmitted command data by RTC).

WFC task	**[MFLOP/cycle]**	**[MB/cycle]**
Pixel calibration	0.1	1.3
WF signal computation	0.1	0.7
WF reconstruction (MVM)	124.1	**248.3**
Temporal control	< 0.1	0.4
Σ	124.4	250.7

Performance aspect	**Rate**
WFC computational power	137 GFLOP/s
WFC memory throughput	**276 GB/s**
Pixel reception throughput	148 MB/s
CCS command throughput	20 MB/s
Telemetry throughput	286 MB/s

- Challenge: Memory bandwidth.

Solution Strategy

- Interoperability with ELT environment
- Maintainability (Commercial off-the-shelf)
- WFC loop in hard real-time
- WFC loop optimizations in soft real-time
- System context:

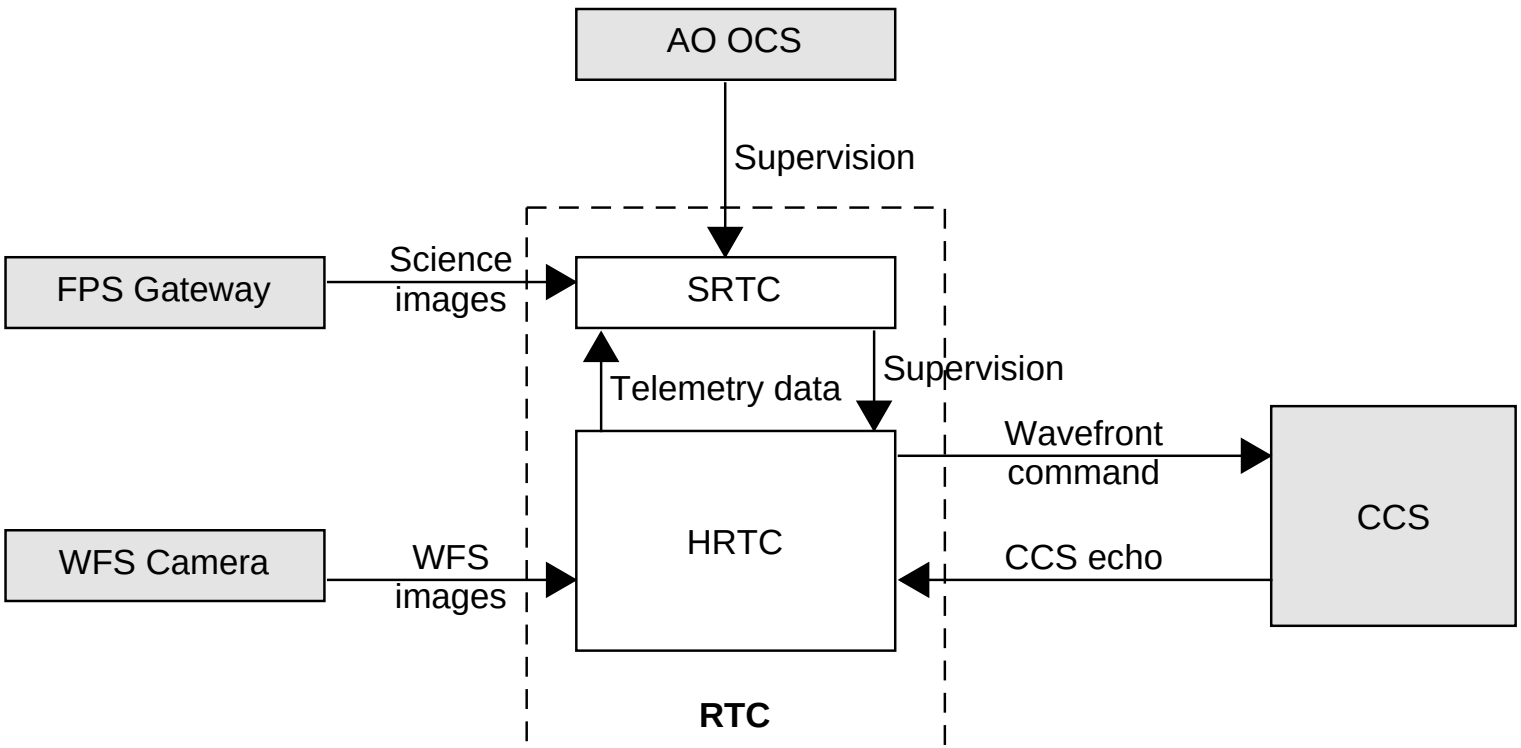

- RTC = HRTC + SRTC

HRTC Prototype: Hardware

Component	Description
GPU server	Asus Barebone ASUS ESC4000A-E10.
CPU	AMD EPYC 7542 with 32 cores @ 3.4 GHz.
Main memory	512 GB DDR4-3200 (data rate: 25.6 GB/s).
GPUs	2x NVIDIA A100 (40 GB HBM2) with PCIe 4.0 x16. Measured internal memory throughput: 1.1 TB/s.
Power supply units	1+1 redundant 1600 W.

HRTC Prototype: Software

- HRTC software is customized for the needs of METIS RTC.
- Computing pipeline contains core functionality:

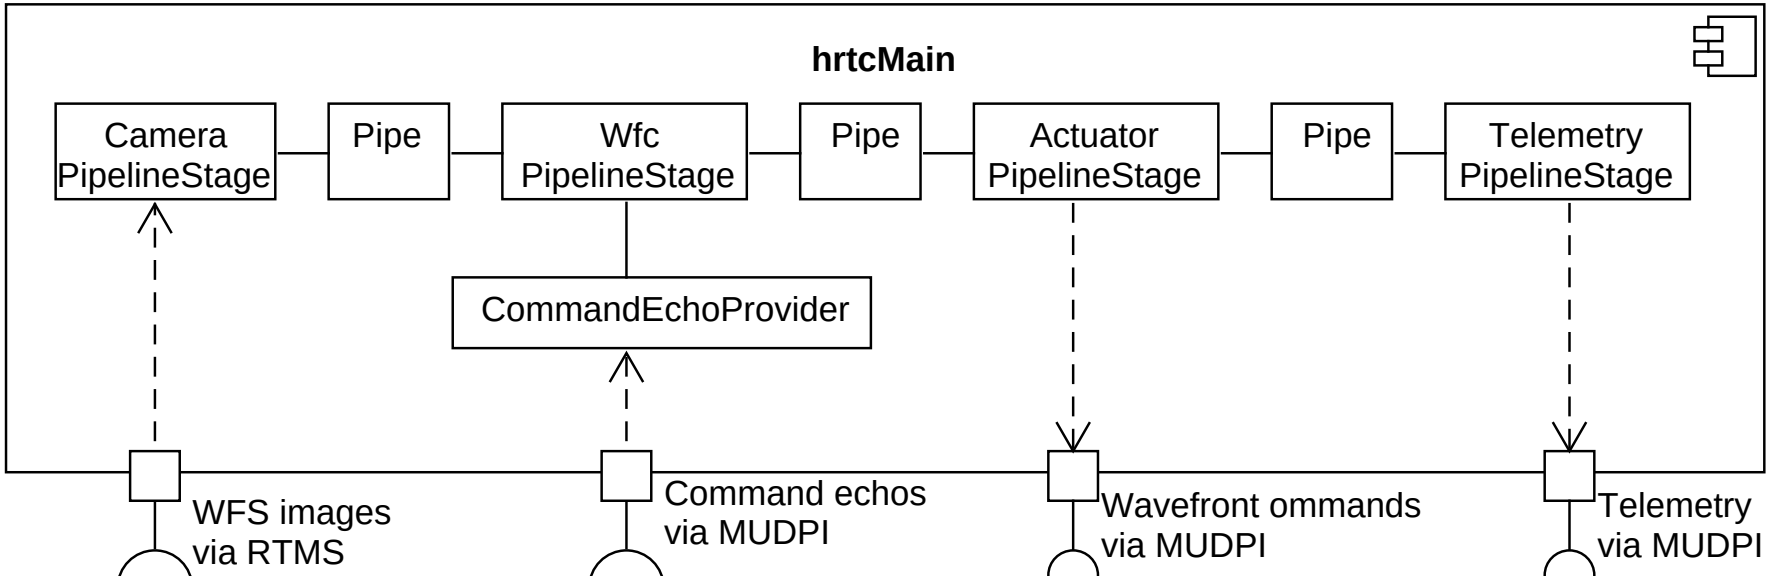

- Some parts of temporal controller are manually optimized with AVX2 instructions.
- Most software dependencies are available in the ELT Development Environment, e.g. GCC, Boost, NVIDIA CUDA compiler and driver. Exception: ZeroC ICE will be replaced in order to be compliant with ESO standards (possibly HTTP).

HRTC Prototype: Performance Experiment

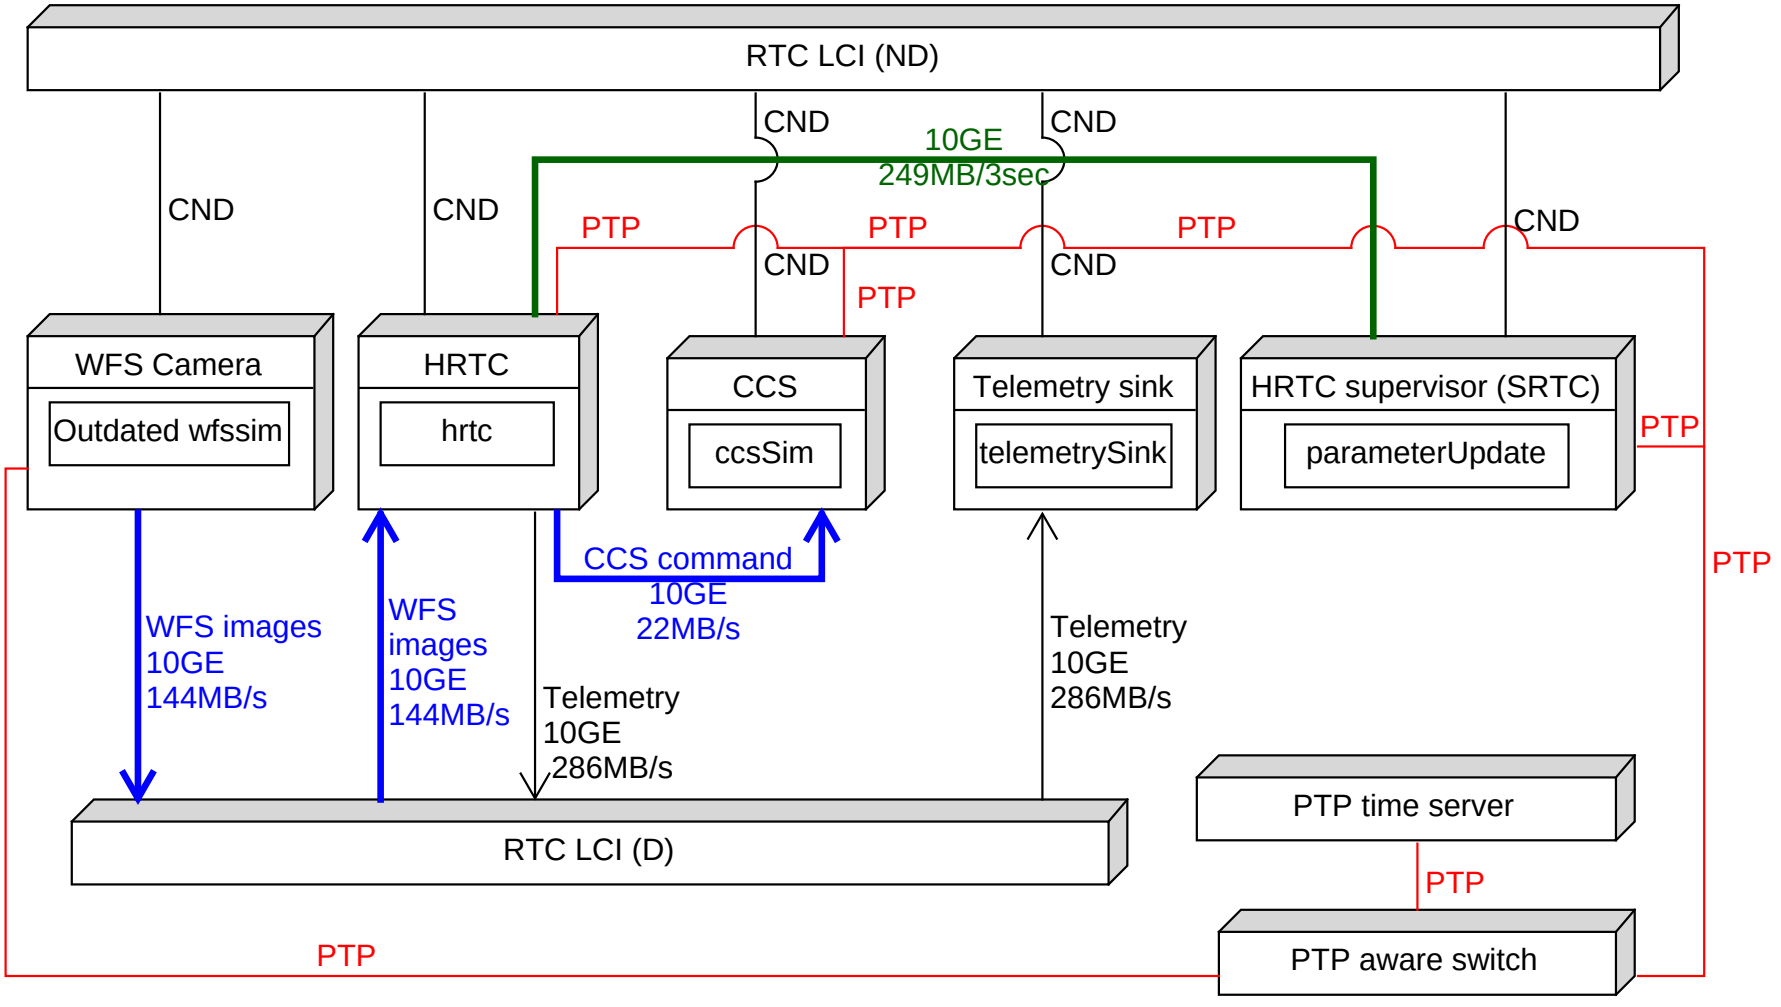

HRTC Prototype: Major Fine-Tuning Actions

- CPU isolations with Linux configuration such as `isolcpus=8-31` (SMT disabled).
- IRQ affinity: All IRQs except the network cards are moved to the first CPU, e.g.
  `echo 1 > /proc/irq/default_smp_affinity` and
  `echo 1 > /proc/irq/$EACH_IRQ/smp_affinity`
- The NVIDIA A100 clock speed has been locked to the maximum operating parameters as defined by the manufacturer, i.e. `nvidia-smi -i 0 -ac 1215,1410` (50 °C and < 120 W; spec: max. 85 °C and 250 W).
- Low-latency network configuration, e.g.
  - IRQs of network cards are moved to isolated CPU cores.
  - Interrupt coalescing is disabled, i.e. `ethtool -C $NIC_NAME adaptive-rx off`
  - Applied profile for low latency network: `tuned-adm profile network-latency`

HRTC Prototype: RTC Computation Time Performance

#Samples: 300,000
(in 5 minutes)
Median: 382 µs
Mean: 381 µs
Max: 761 µs
Std. dev.: 6 µs
99.99%: 493 µs

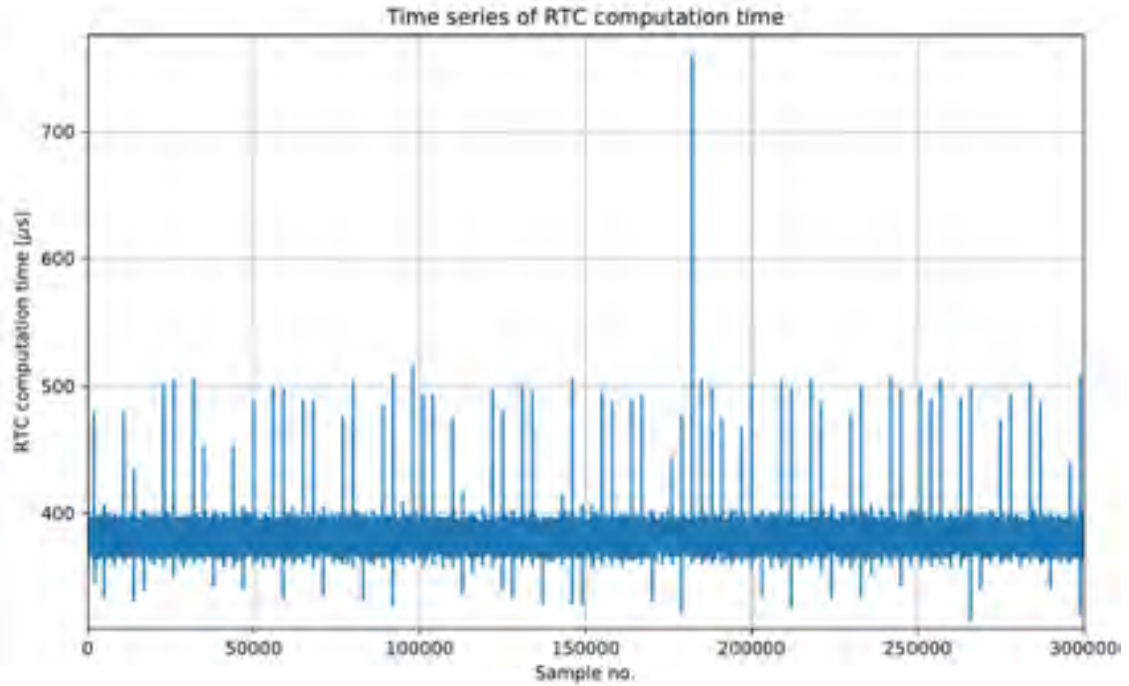

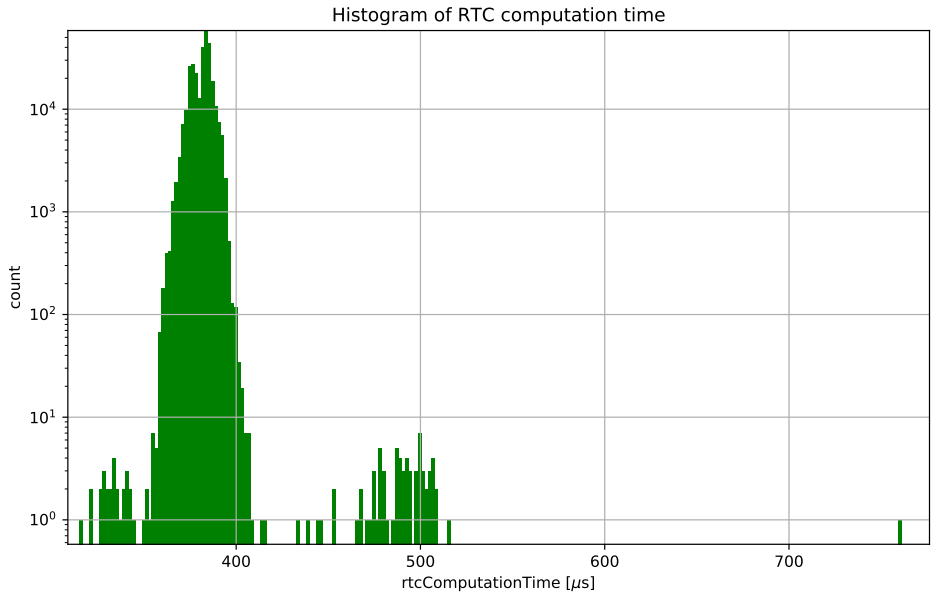

HRTC Prototype: Numerical Correctness Test

- Replaced RTC of COMPASS with METIS HRTC via network.

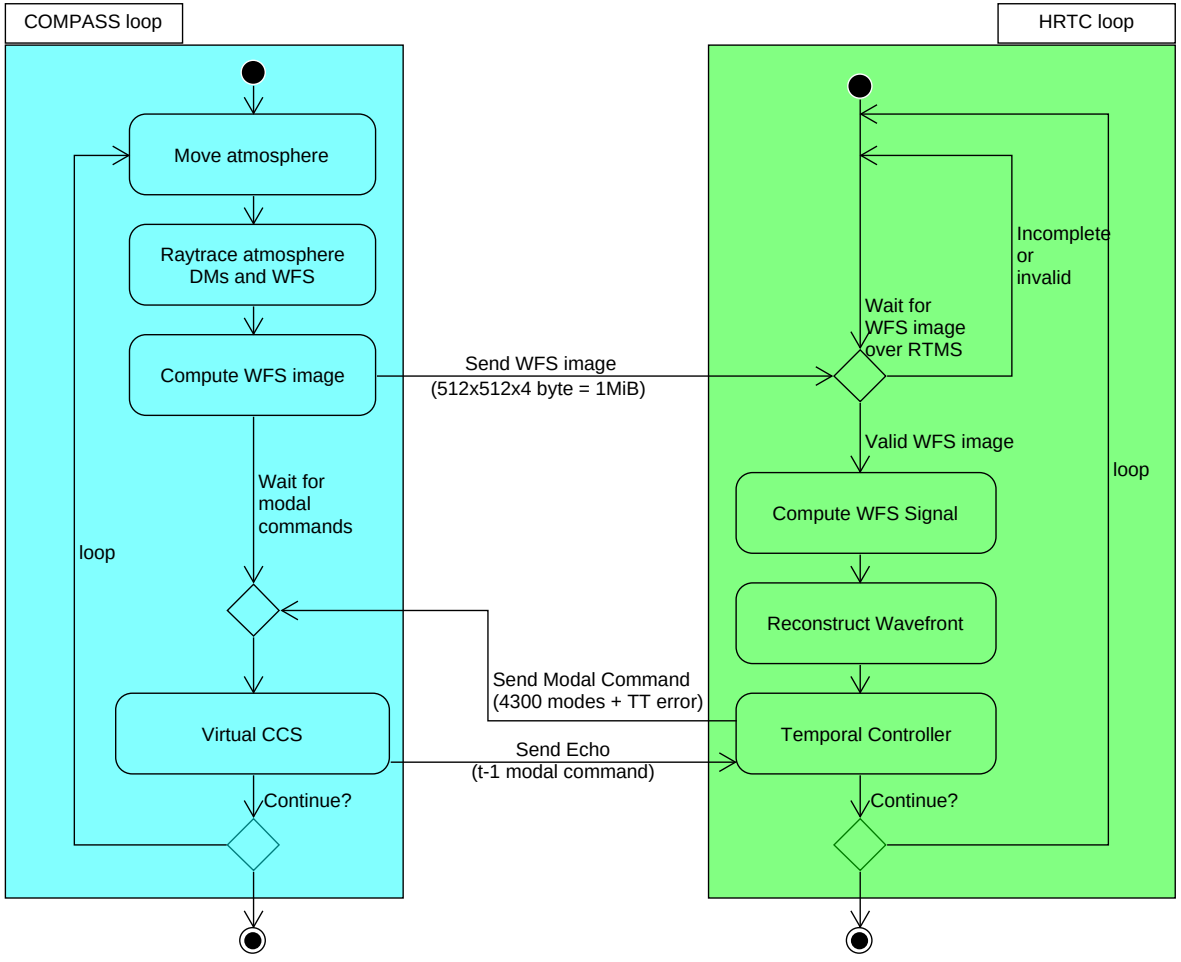

- Result: Strehl ratio stays the same with HRTC in the loop.

SRTC Functionality

- The SRTC supervises and optimizes the HRTC operation.
- It will host many components for different tasks such as estimating parameters and optimization parameters:
  - Command Matrix (CM) Optimizer computes a new command matrix depending of updated configuration parameter such as valid subaperture map or pupil rotation angle. Its operation requires roughly 19 TFLOPS and about 27 GB of memory (FP32).
  - Pupil Position Monitor identifies the lateral pupil position by applying a matched filter algorithm.
- Telemetry requirements are not complete clear, but in worst case, the METIS RTC must be capable to store more than 500 MB/s produced by Telemetry Recorders and Metadata Collector during operation.
- Implementation with ESO RTC Toolkit has started (Pupil Position Monitor).

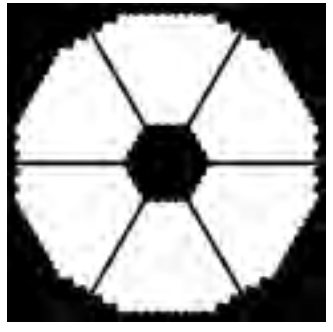

SRTC Hardware Investigation at FDR

- Distributed system built with ESO IT standard servers.

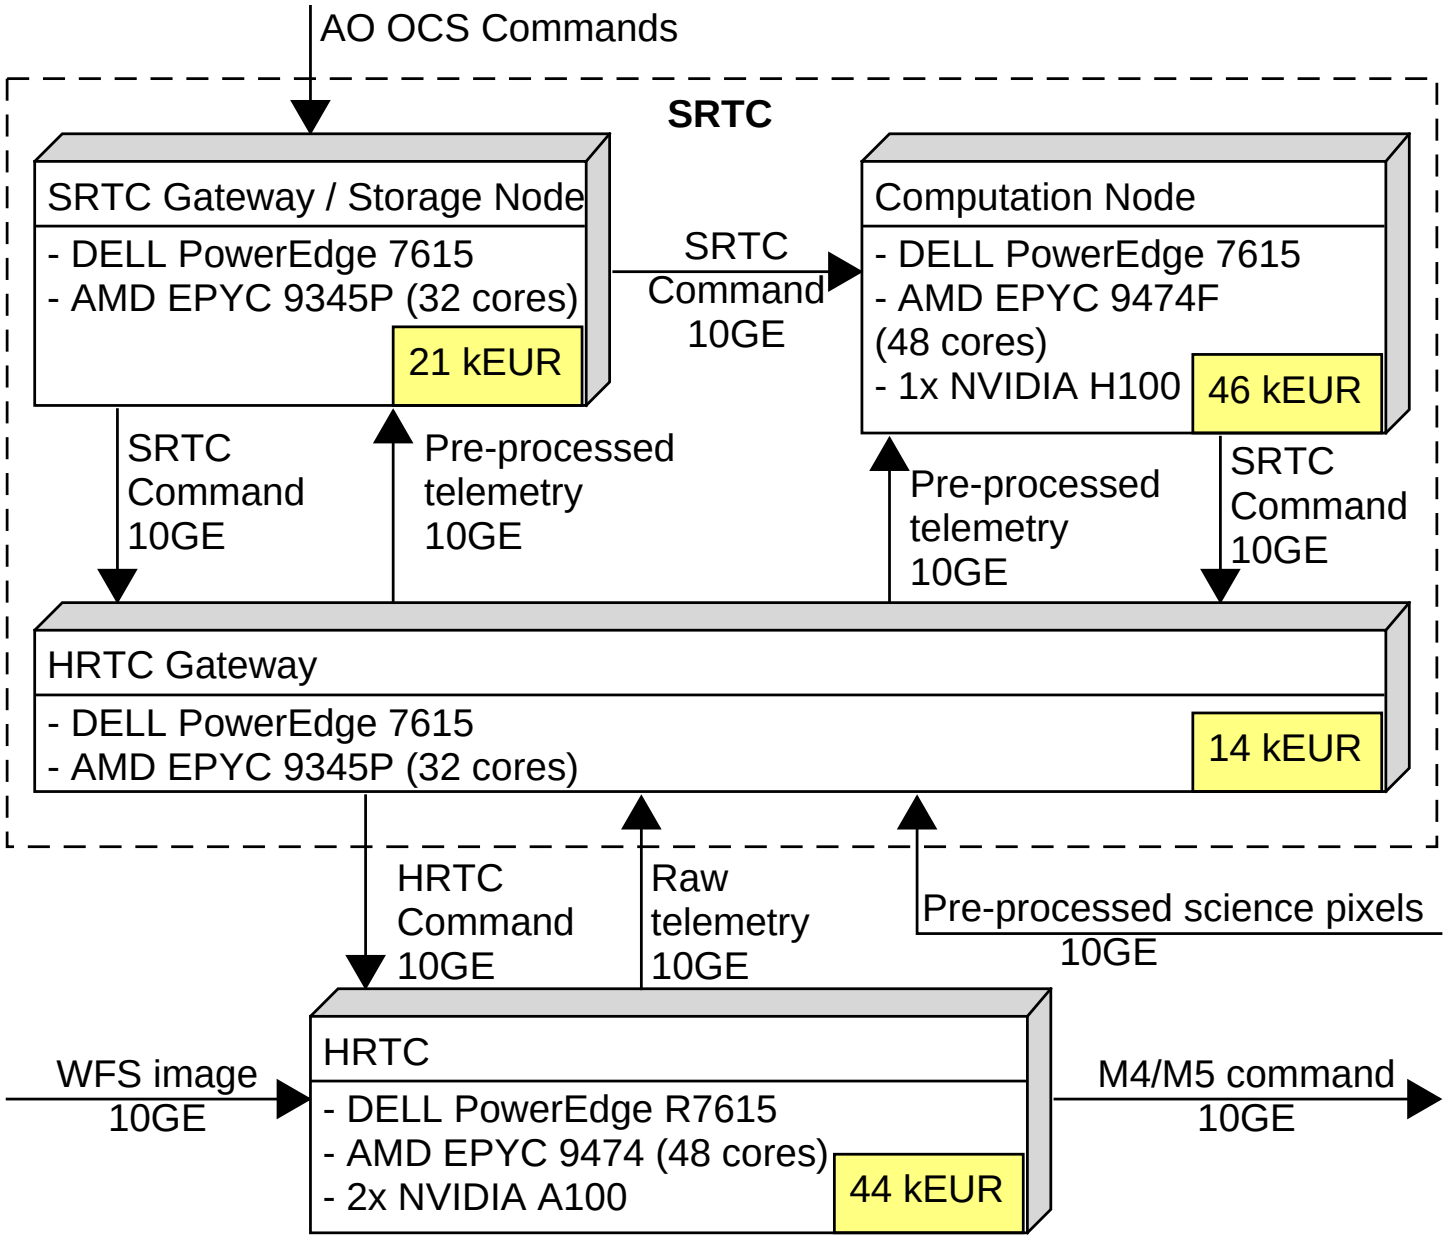

std::exit(0)

- METIS SCAO has passed its FDR in autumn 2022.
- HRTC is capable to meet the performance requirements of the WFC loop in a realistic setup including repeated updates of the command matrix.
- Using a modified version of COMPASS as a framework for the simulation of the telescope, we have demonstrated that we can close the AO loop using the HRTC prototype.
- Further development of the HRTC is expected to run in parallel with the SRTC.
- The SRTC functionality and its hardware demands have been identified and first prototyping activities have started.