SPARTA Upgrade

Pavel Shchekaturov

AOF GALACSI main interfaces

- sFPDP 2.5 Gb is the main real-time interface for communication
  - DSM
  - 4LGSF - Fast Jitter mirrors (FSM) and Slow Field Steering Mirrors
  - 4 x CCD220 WFSs
  - VIS TT CCD220 WFS
  - SAPHIRA WFS for IRLOS

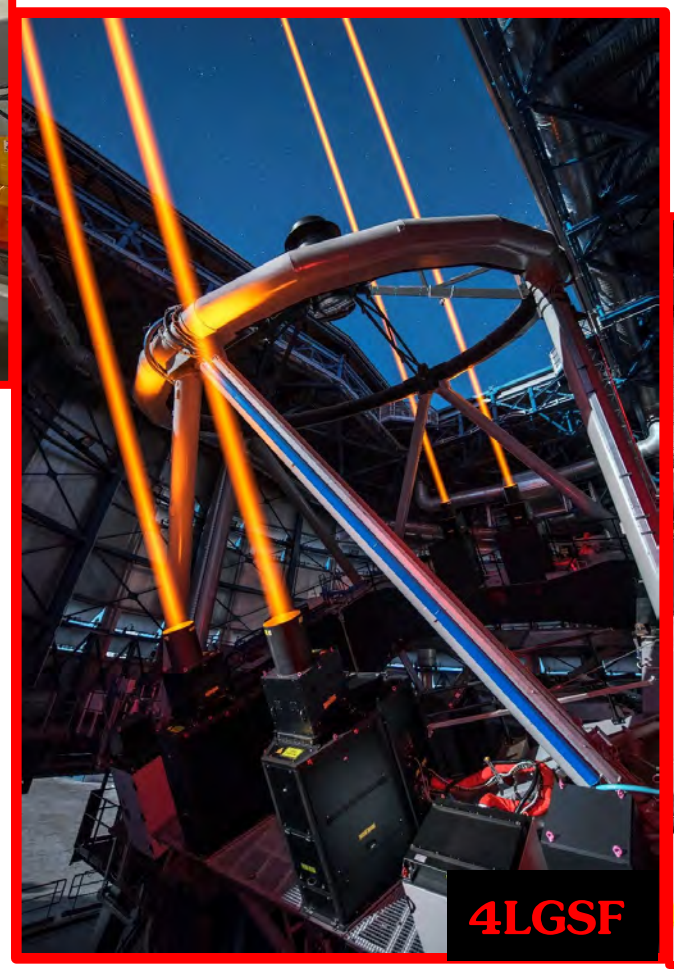

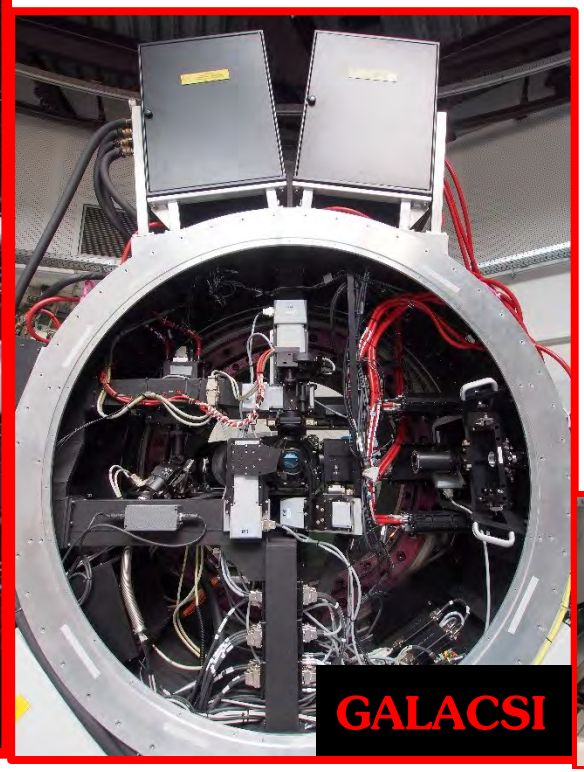

- SPARTA Upgrade does not use Ethernet based interfaces for real-time communications
- GigE, MUDPI and UDP based protocols are/will be supported by the Platform for G+ project

AOF GALACSI modes overview

WFM

NFM

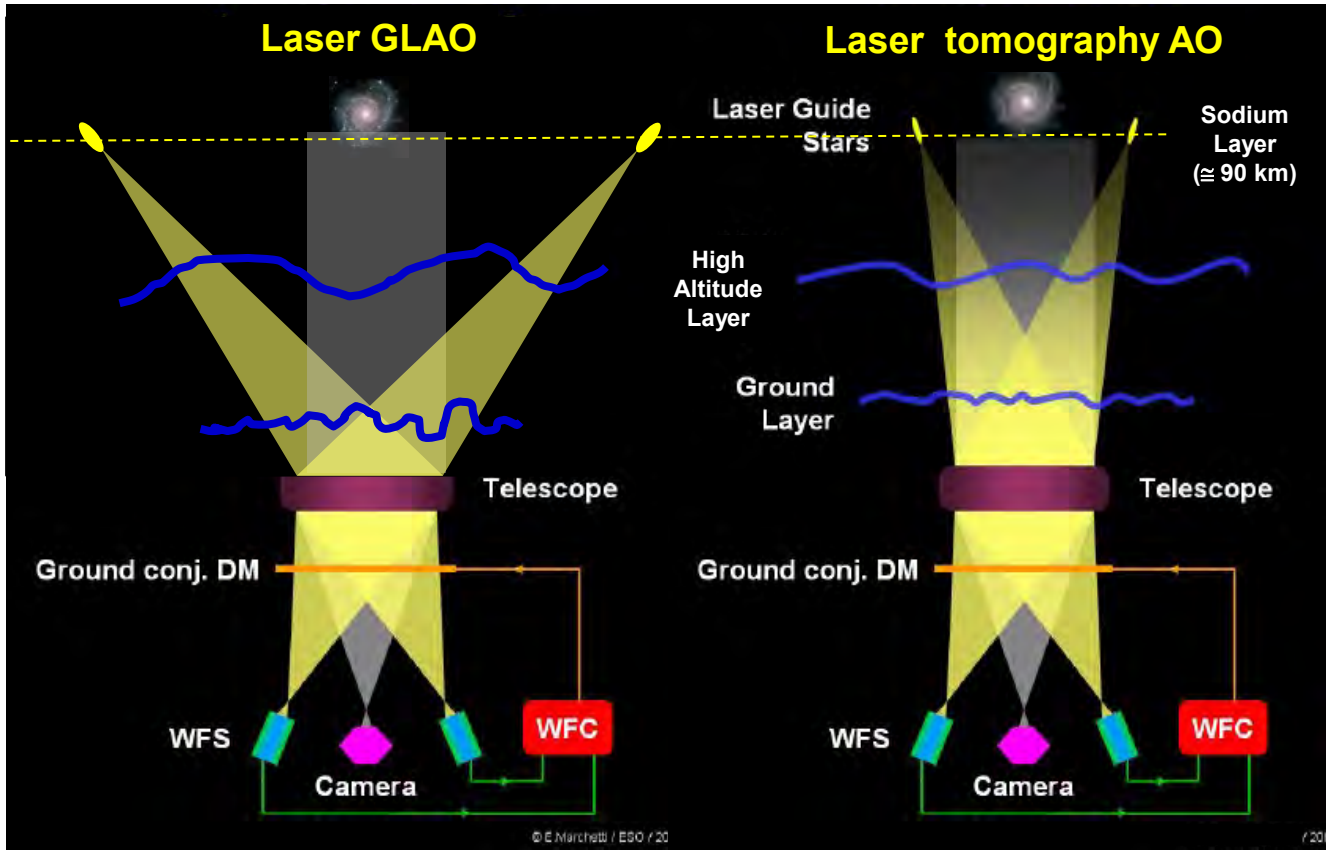

- LGS HO and TT Jitter loops
  - 1 KHz loop rate
  - 4 x LGS SH WFS with 240x240 pixel frame size
  - 40x40 subapertures
  - DSM 1156 actuators
  - 8 Fast Jitter Actuators and 8 Slow FSM actuators
  - 4 x CM with size of 2480x1158
- NFM mode IR loop
  - Frame rate up to 496 Hz (1 Hz, 10 Hz, 100 Hz, 200 Hz, 496 Hz)
  - 256x256 and 128x64 pixels frame sizes
- WFS mode VIS TT loop
  - 200 Hz loop rate
  - 240x240 pixel frame size

Project inception

- January 2019 - SPARTA Red Flag report about Degradation of VPF1 boards lifetime was raised due to the increased failure rate
- SPARTA Upgrade was inspired by the ELT HRTCp prototyping results showing that it is feasible to archive good performance with CPU based HPC HW
- Q1 2019 – Q3 2020 – prototyping phase
- November 2020 – project proposal was submitted
- Some of the design decisions and ideas were reused from ELT HRTCp project

Project motivation

- Replacement of the SPARTA RTC box by single 19 inches rack-mountable HRTC Dell PowerEdge R7525 server with 2 x AMD 7702 Zen2 EPYC Romes CPUs
- Replacement of the sFPDP I/O interface previously implemented in FPGA by 2 x PCIe 2.5 Gb 4 ports sFPDP PCIe FPGA Card
- Implementation of new SW components (SPARTA drivers) previously running on FPGA
- Porting of old SPARTA driver components (PPC based) to new CPU based Platform (SSE, AVX2)

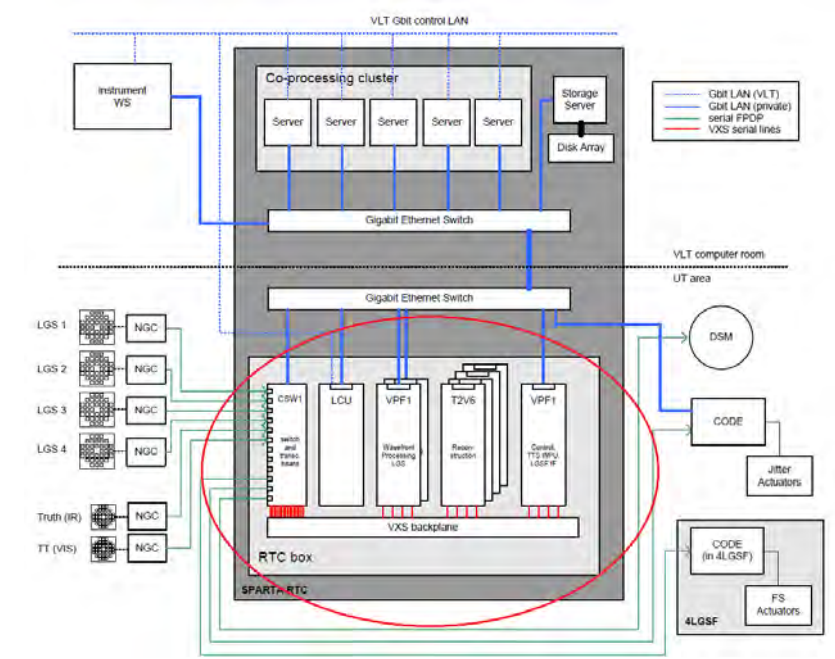

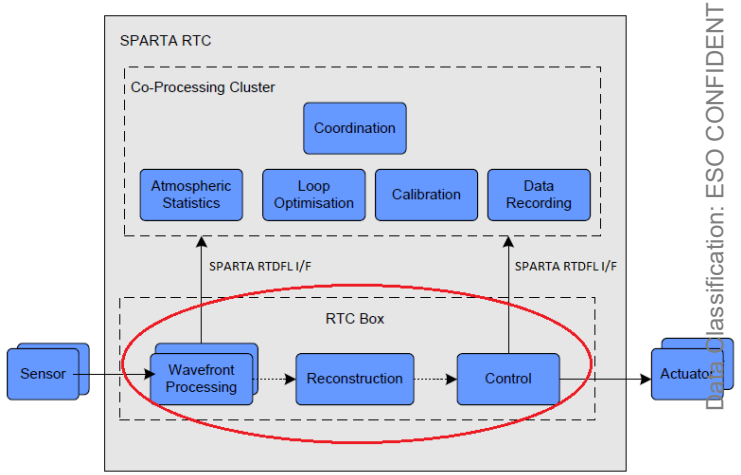

SPARTA Upgrade HRTC HW overview

- HRTC Dell PowerEdge R7525 with 2 x AMD 7702 (2 GHz, 64 Cores / 128 Threads, 256 M, 200 W) Zen 2 CPUs
- 2 x v5051 PCIe 2.5 Gb 4 ports sFPDP cards from NewWaveDV Card based on Xilinx UltraScale+ VU9P FPGA
- 10 Gb Ethernet is used for the telemetry and supervisory network
  - NIC Intel x710 10Gb T-BASE 4 ports for telemetry and supervisory LAN (PXE boot)
  - Cisco Catalist 9300 10Gb T-BASE network switch
  - Supervisory LAN is used to boot HRTC from the GW machine via PXE and have access to HRTC via ssh
- Parallel port PCIe card for latency measurements

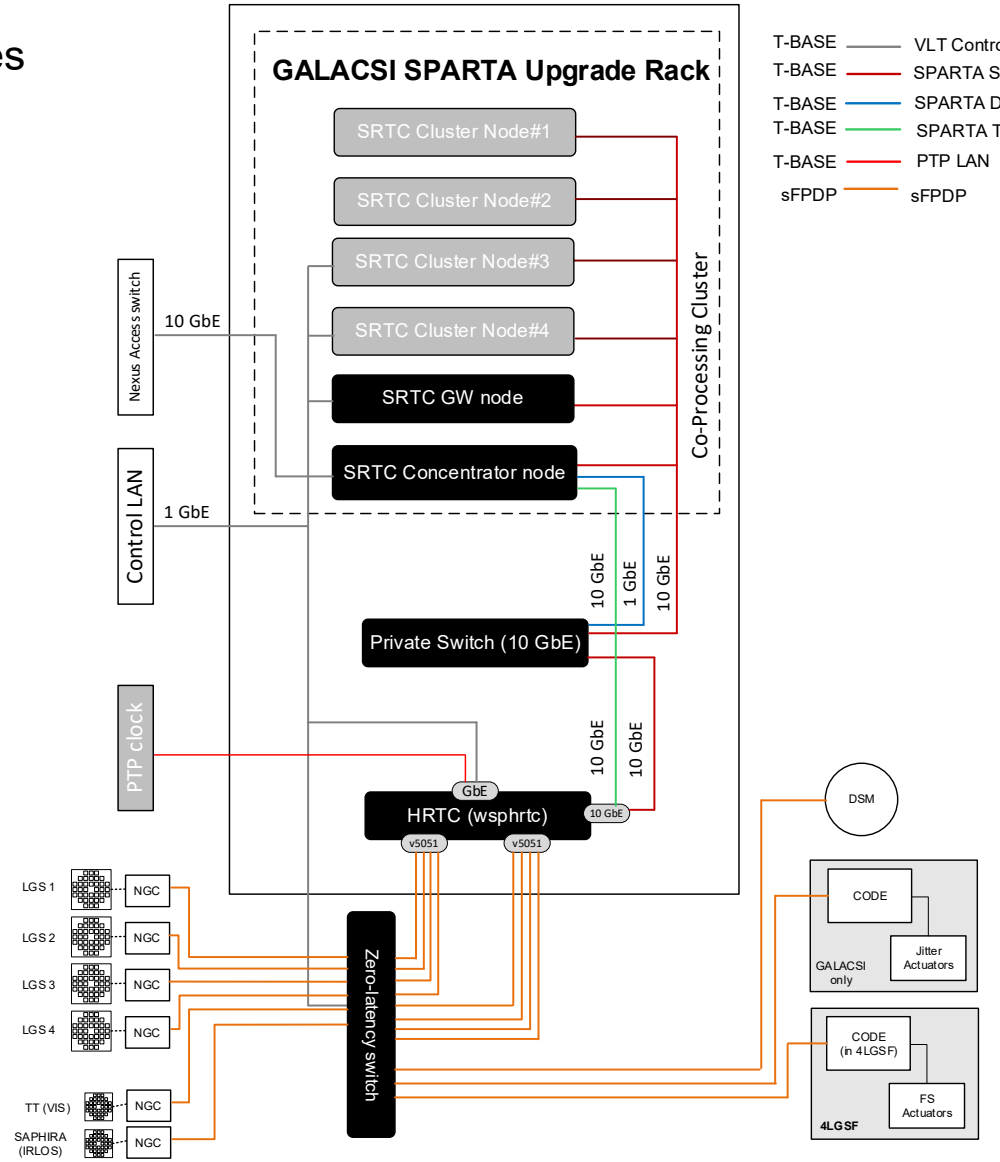

SPARTA Upgrade HRTC main SW concepts

- Keep the interface to the SRTC cluster unmodified as much as possible (SPARTA driver I/F to the supervisors, SPARTA RTDFL I/F for the telemetry distribution and disturbance injection)
- Custom initrtd image without systemd and standard 3.10.0-1127 kernel version from VLTSW2018 (CentOS 7.4)
- HRTC SW is built on the GW machine under the same VLTSW version as SRTC (VLTSW2018), but under different account (different environment setup and GCC 7.3.1 compiler toolchain)
- HRTC part contains CPU bound non-blocking spinning threads
- Custom-made Shared Memory interface is used for non-blocking communication between the real-time threads

SPARTA Upgrade sFPDP I/O interface

- NewWaveDV V5051 PCIe 2.5 Gb 4 ports sFPDP PCIe FPGA Card based on Xilinx UltraScale+ VU9P
- Low level polling non-blocking API, parsing and creating the sFPDP packets by ourself
- The card is configured to use 1 DMA for reception and 1 DMA engine for transmission (most performant configuration)
- Reception delay of pixel block of 5760 bytes on 1 port is 3.2 us
- Transmission delay on block size of 2768 bytes is ~6 us

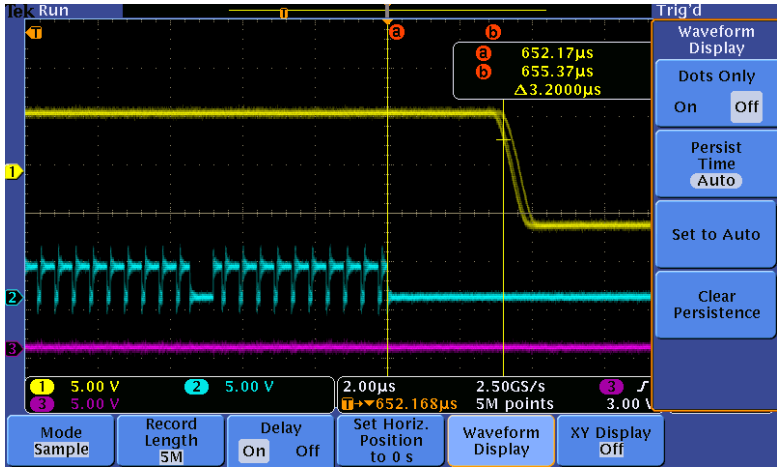

Blue – pixel block of 5760 bytes reception in FPGA by old SPARTA RTC box
Yellow – pixel block of 5760 bytes reception by v5051 on 1 port (RS232 port trigger)

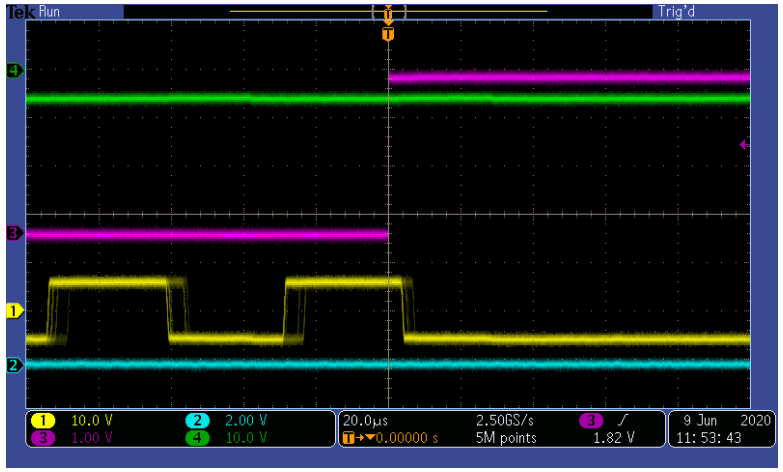

Pink – CCD220 AO NGC WFS_FRAME_SENT signal
Yellow – pixel block of 5760 bytes reception by v5051 on 4 ports (RS232 port trigger)

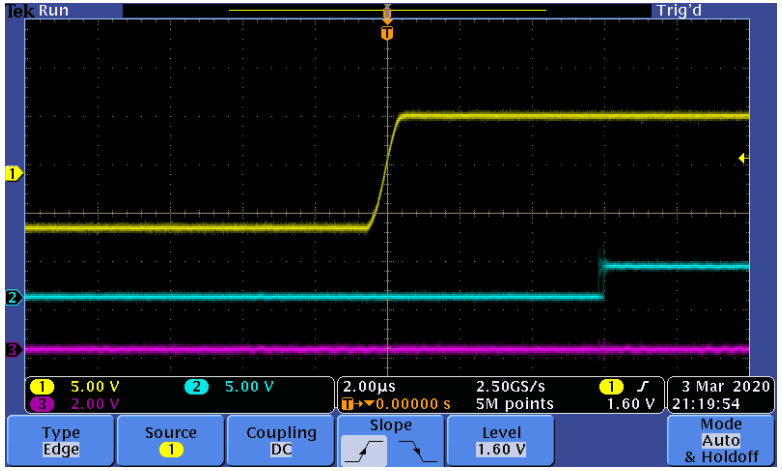

Blue – 2768 bytes reception by FPGA in old SPARTA
Yellow – v5051 transmission start (RS232 port trigger)

AOF GALACSI HRTC control scheme

- 4 x LGS pipelines (4 x WPU and 4 x MVM)
  - LGS/HO reconstruction in the zonal space
  - Shared LGS/HO IIR filter coefficients for all HO modes
  - Loop is frozen if clipping of the command is detected or negative ECHO is received from DSM (last valid command is repeated)
- Jitter loop
- LO VIS/IR loop

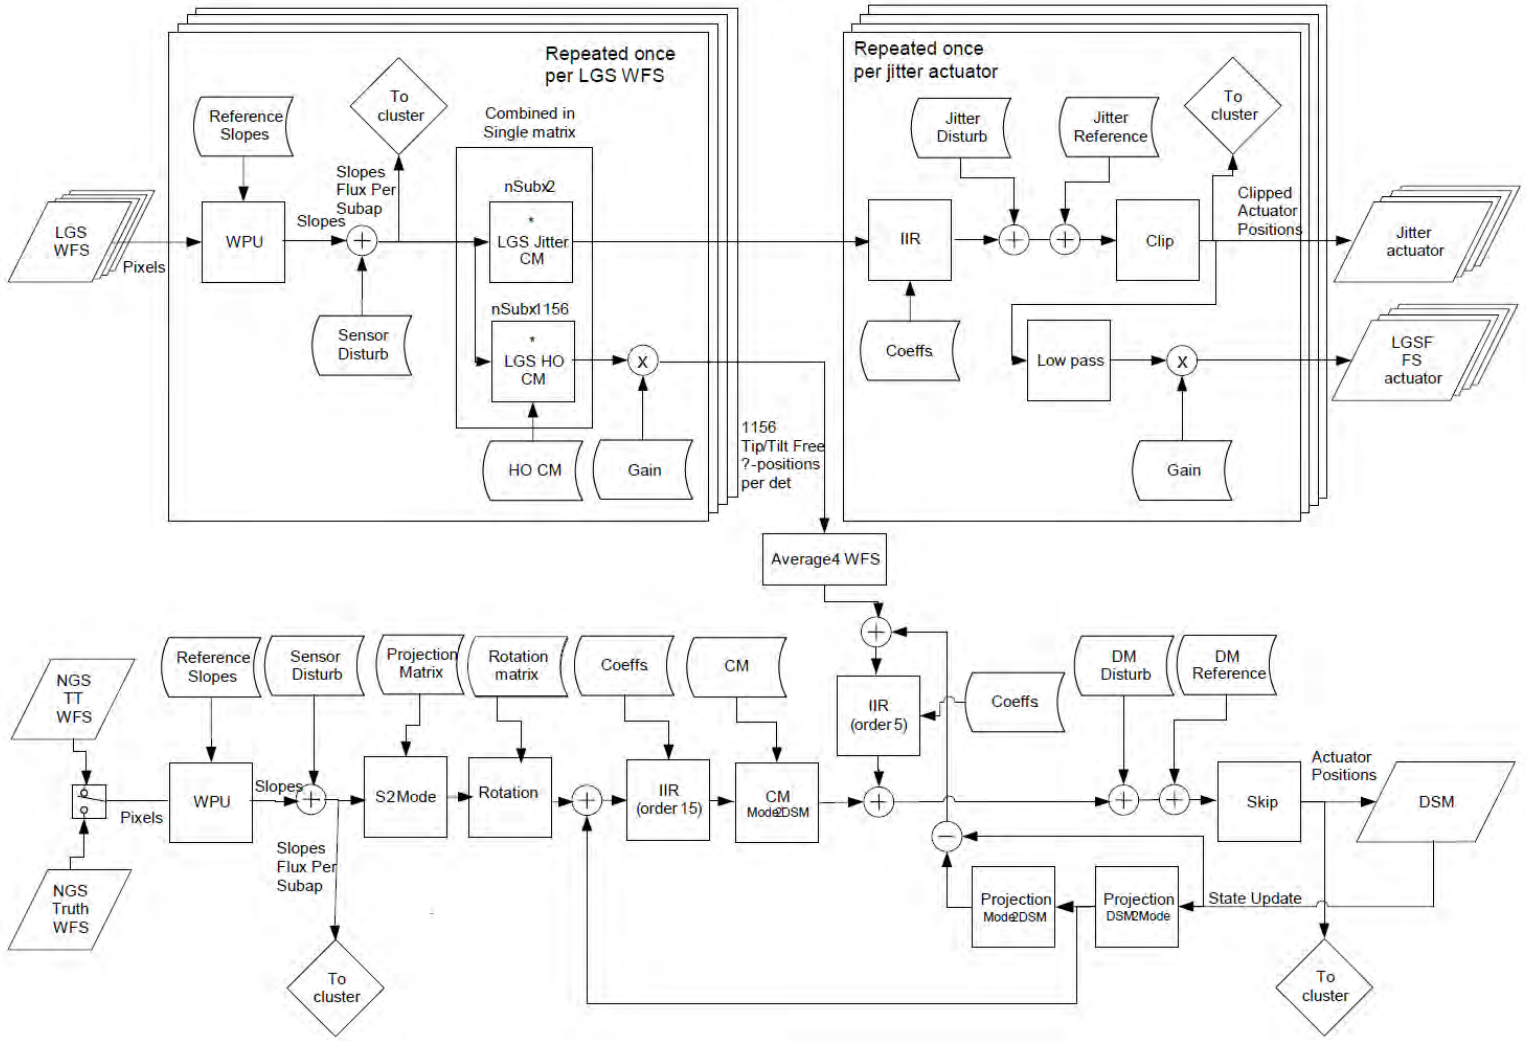

LGS/HO Control Matrixes update optimization

- Reconstruction Supervisor component sends transposed CMs to HRTC
- Network thread writes CM to the double buffer on NUMA node #4 (OS/telemetry NUMA node)
- MVM IO supervisory thread copies small slice of CM on every loop cycle in its idle time into the corresponding worker threads for each MVM pipeline
- When the copy is finished the set the flag in each worker thread to prefetch its own CM slice to bring it into the local cache
- Currently the prefetch phase is just 1 loop cycle. During this loop cycle old CM is used in the computation.

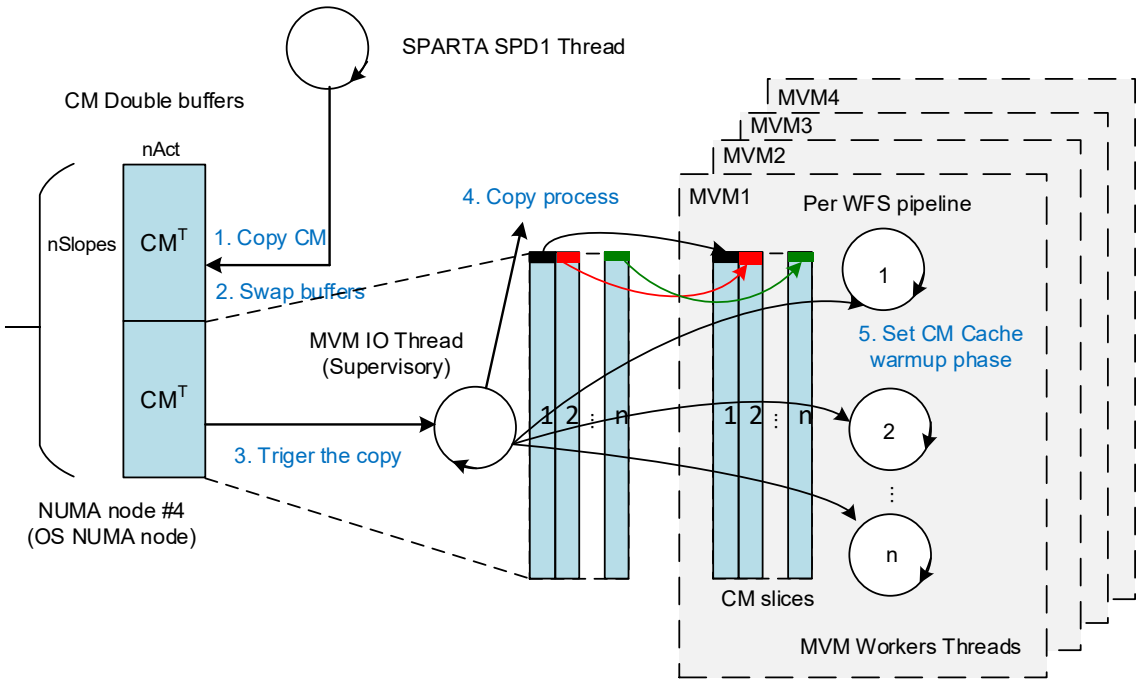

Before CM update optimization

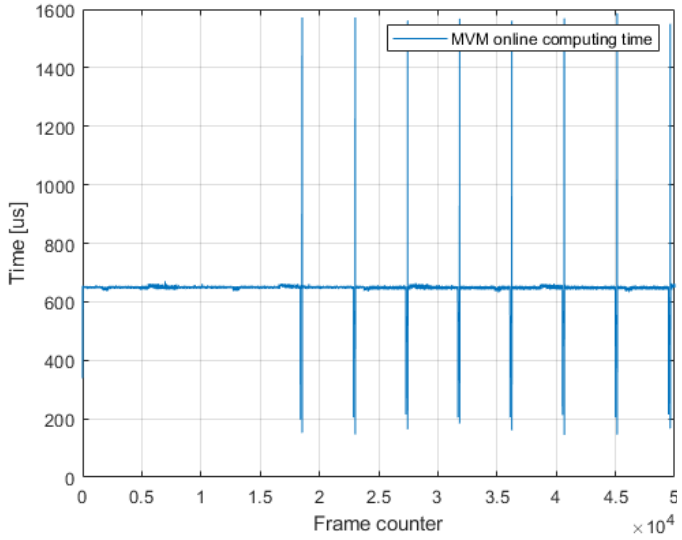

After CM update optimization

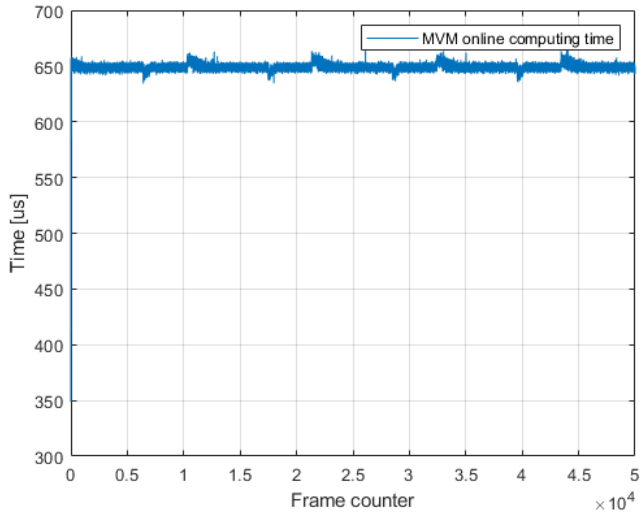

HRTC SW deployment and data flow 1/2

AMD 7702 (64 Cores / 128 Threads, 256 M, 200 W) Zen 2CPUs configured as 4 NUMA nodes per socket

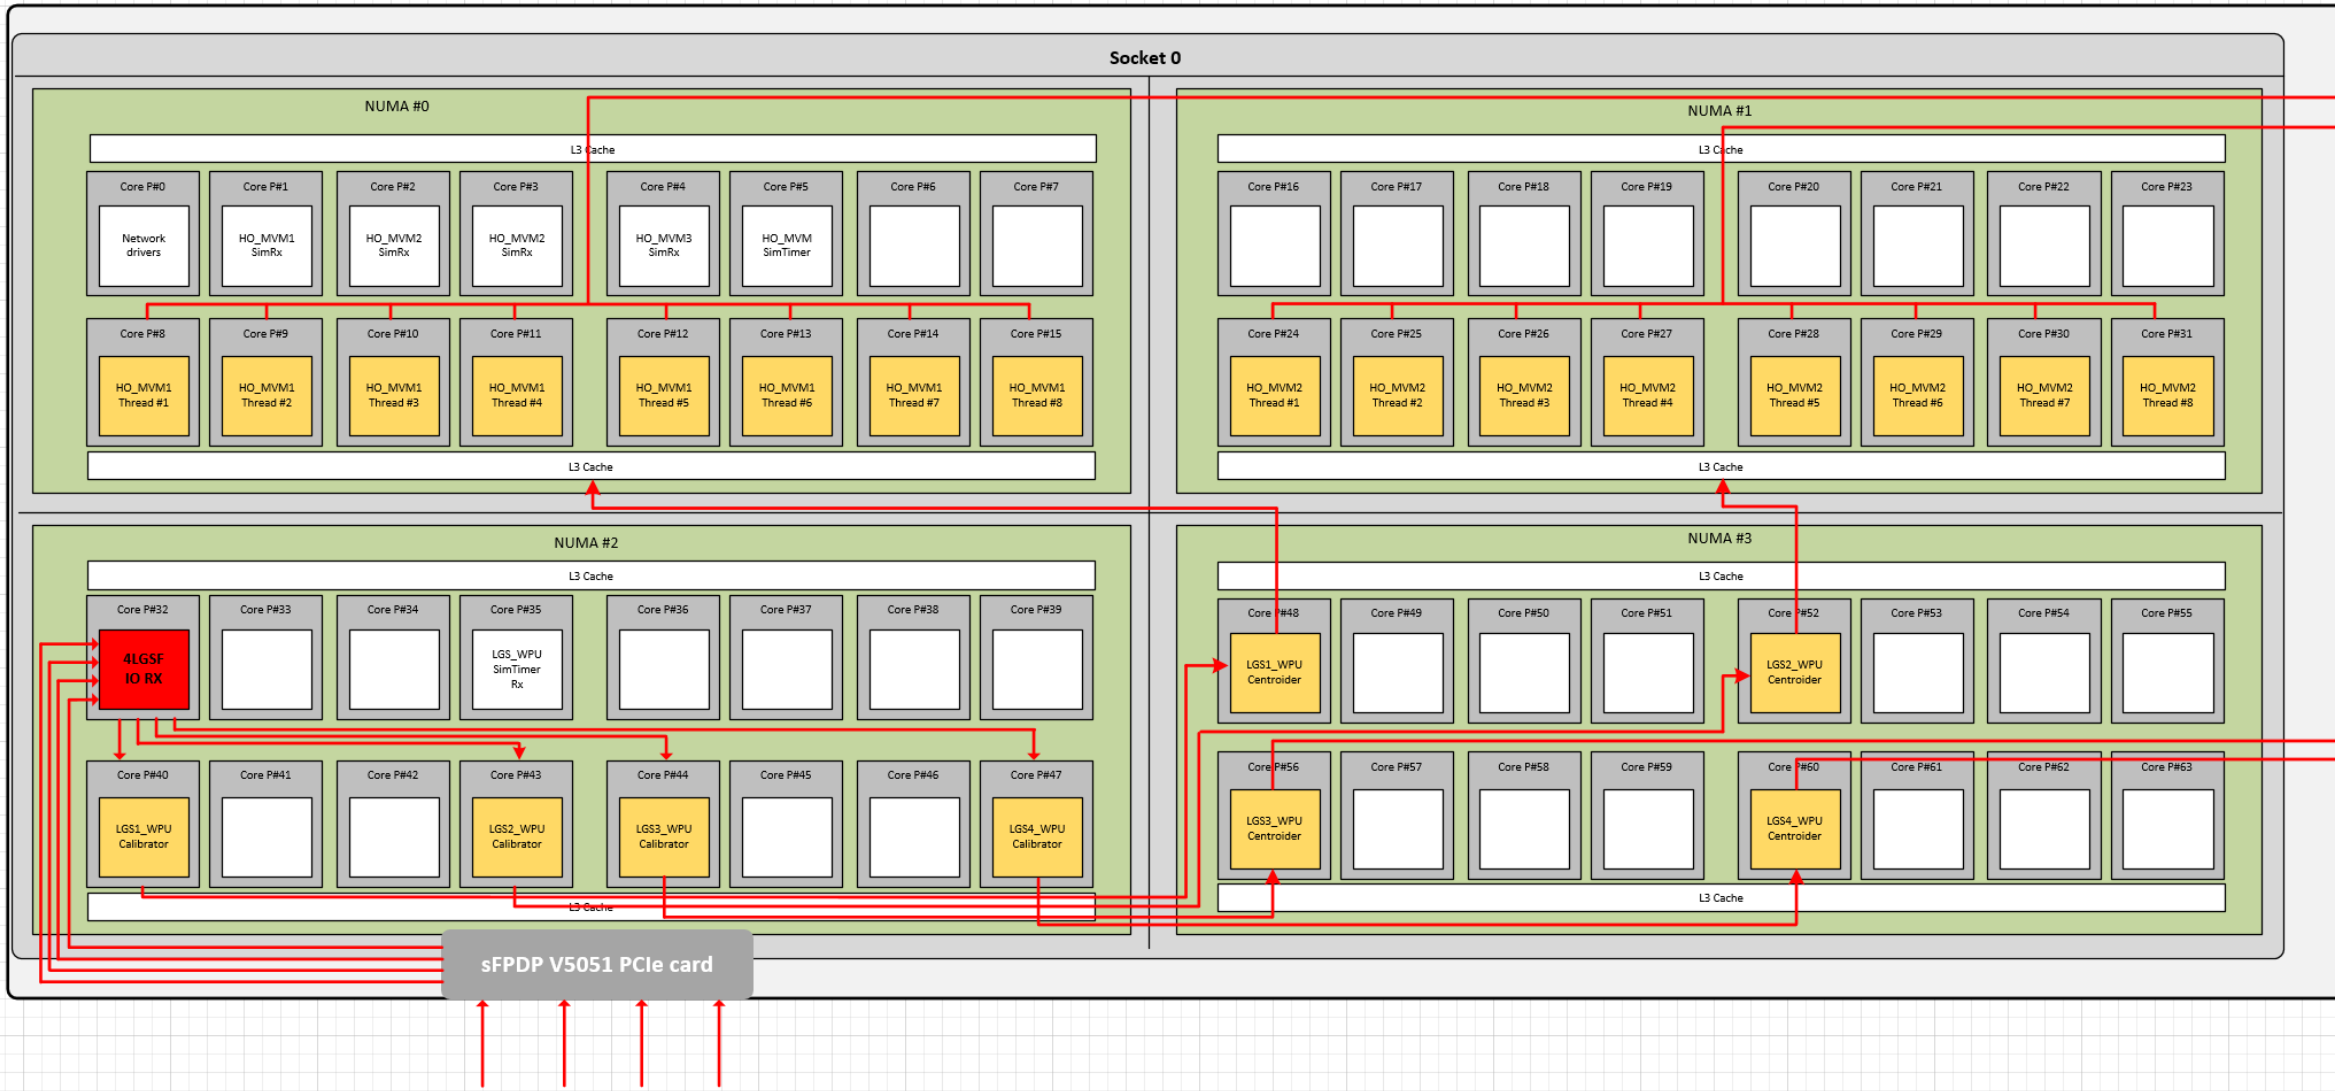

HRTC SW deployment and data flow 2/2

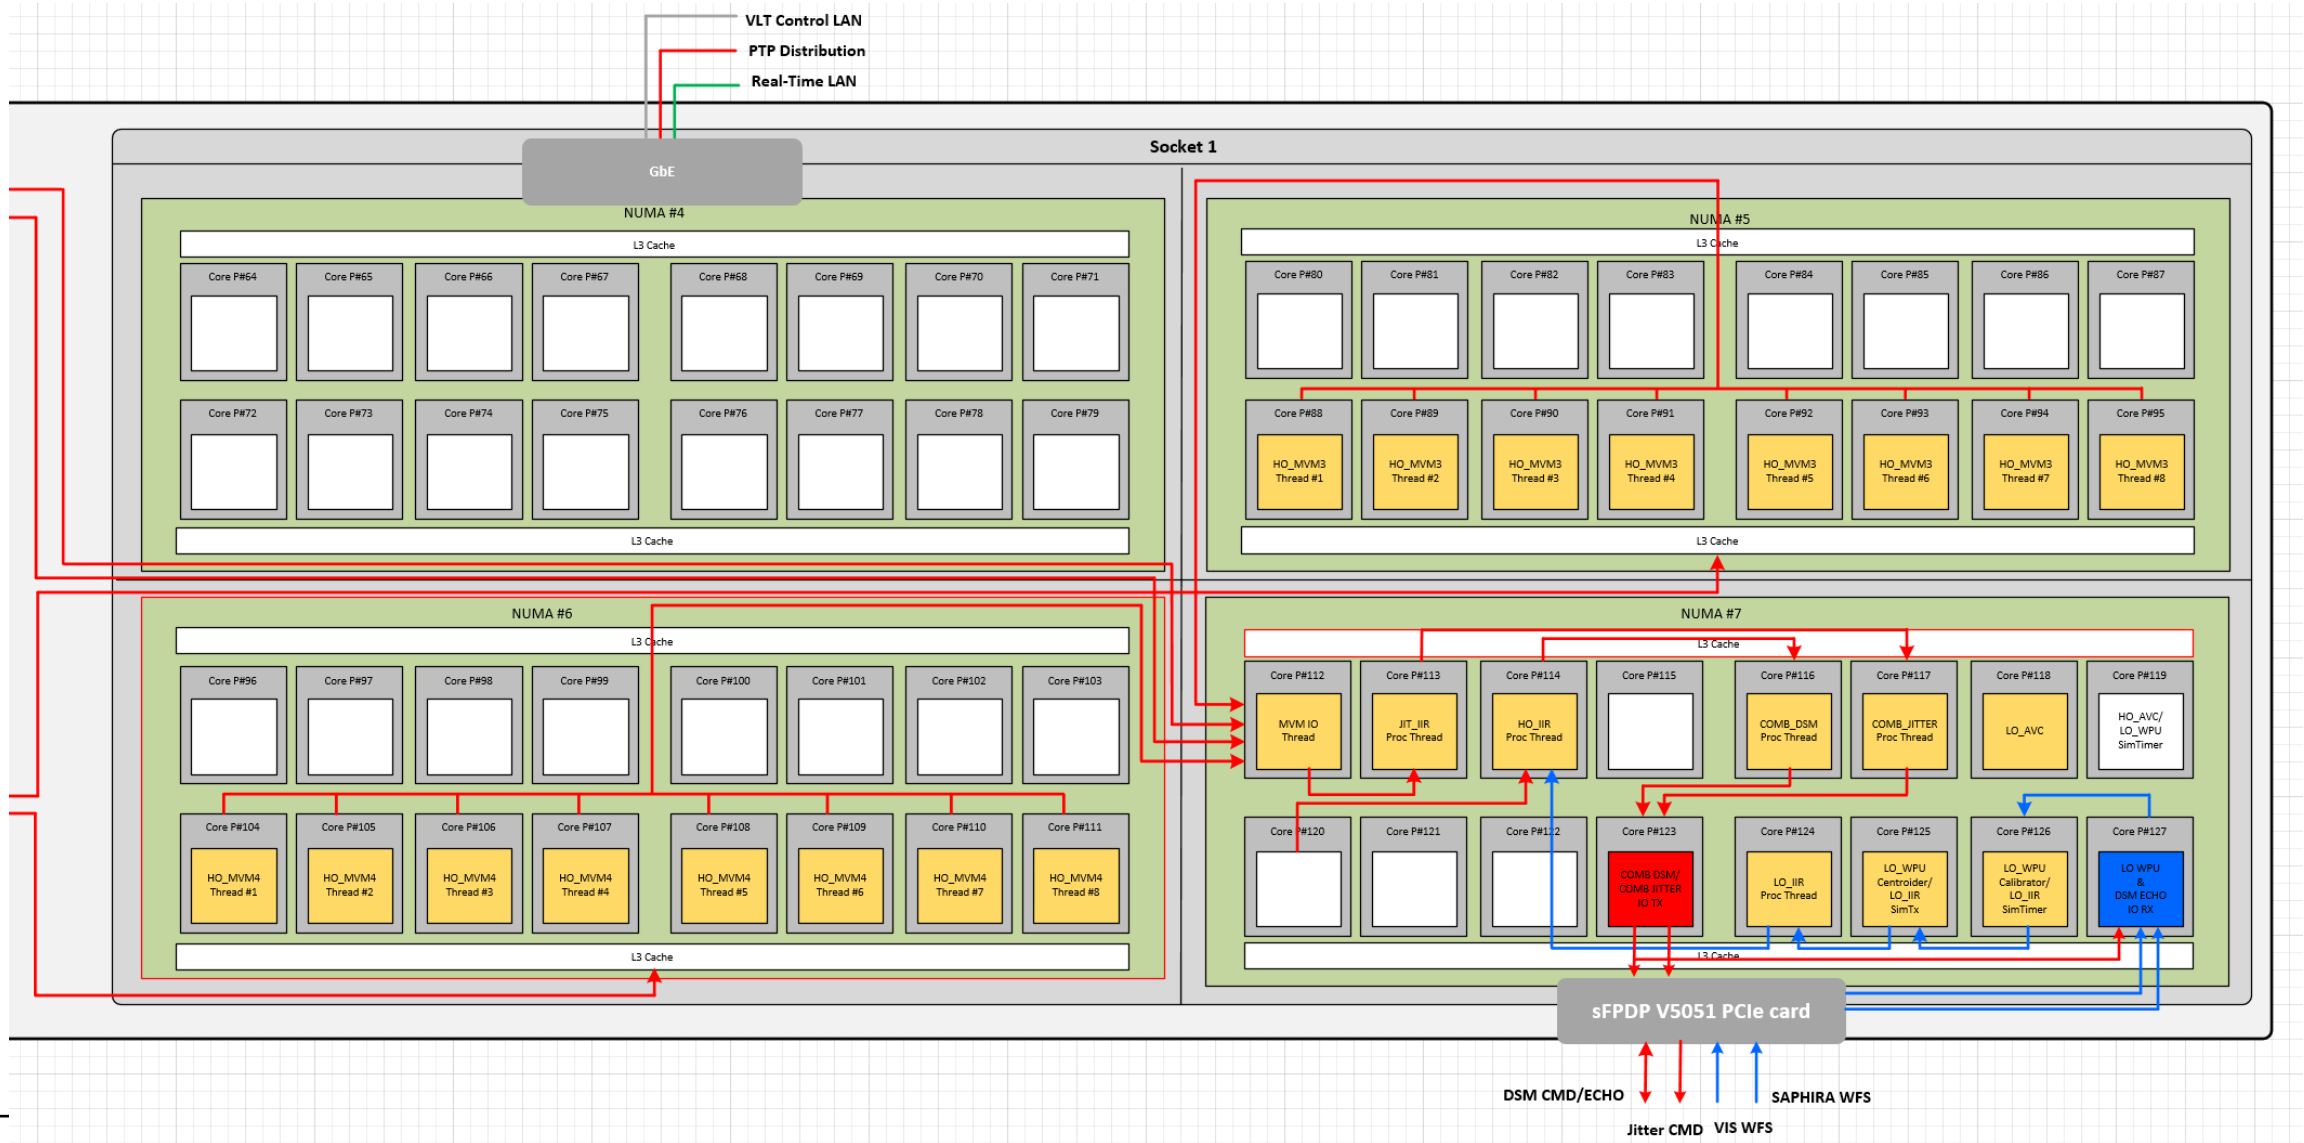

Key optimization HW & SW aspects

- Diskless PXE booted system without systemd
- Isolate of real-time performance critical NUMA node/cores from OS schedule (isolcpus, nohz_full, rcu_nocbs)
- Statically build binary
- Keep the I/O data processing threads local to NUMA node of PCIe cards (sFPDP, NICs and etc.) and avoid putting additional performance critical tasks on the same NUMA node
- Non-blocking spinning threads and polling drivers API
- Atomic based Shared Memory Interface between various computing threads
- Lightweight atomic based signaling mechanism
- Configuration update throttling (one map update per loop cycle, CM update copy in small chunks)

GALACSI SPARTA Upgrade latency measurement setup

- CCD220 pixel stream is multicasted with sFPDP Zero Latency switch to emulate 4 LGS WFSs
- NAOMI I/F sFPDP board is used to receive DSM command (NAOMI command of 1536 bytes instead 1156 as for DSM) and trigger GPIO pin upon the reception
- NAOMI I/F sFPDP board is used to emulate CCD220 for VIS TT WFS in WFM mode
- Jitter CODE LCU is used to receive the Jitter command and trigger GPIO pin when reception is complete

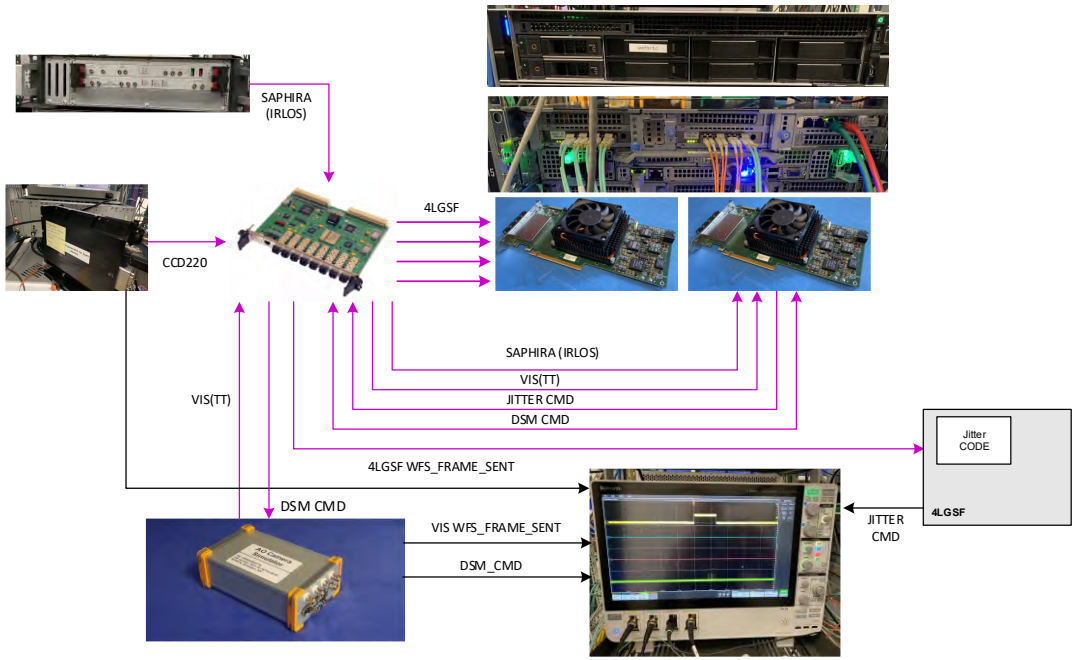

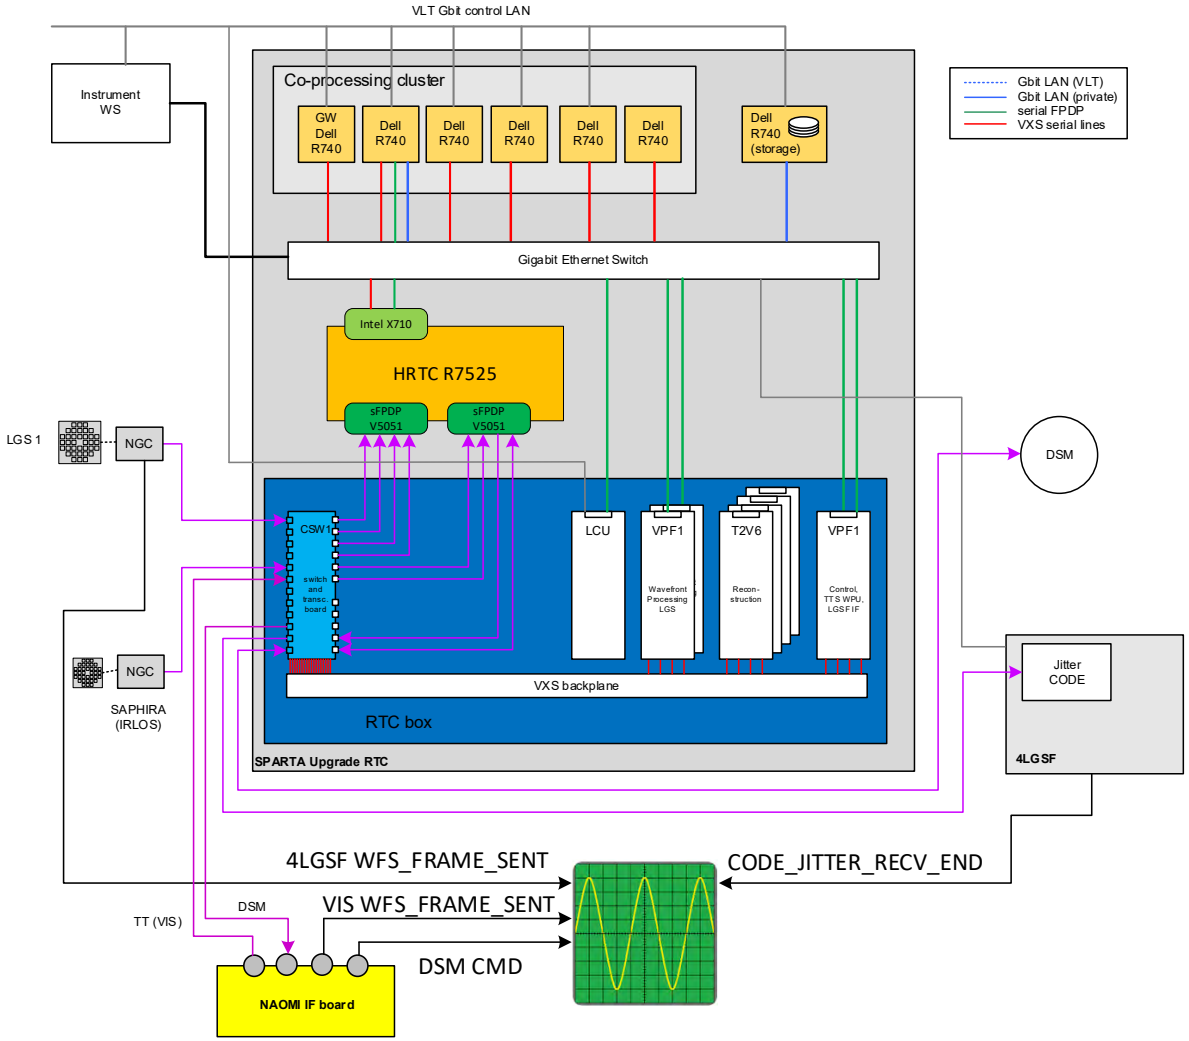

Performance results with the test system in NFM mode

- Continuously running the system for ~48 hours
- LGS, Jitter, LO loops were closed with 0 gain

✓ **LGS HO loop statistics (Meas 3):**

- ✓ Mean: 77.6 us
  Std. Dev.: 2.23 us
  Pk-Pk.: 50.86 us
  Min: 69.6 us
  Max: 120.5 us

✓ **Jitter loop statistics (Meas 4):**

- ✓ Mean: 99.3 us
  Std. Dev.: 2.38 us
  Pk-Pk.: 50.4 us
  Min: 90.5 us
  Max: 141.0 us

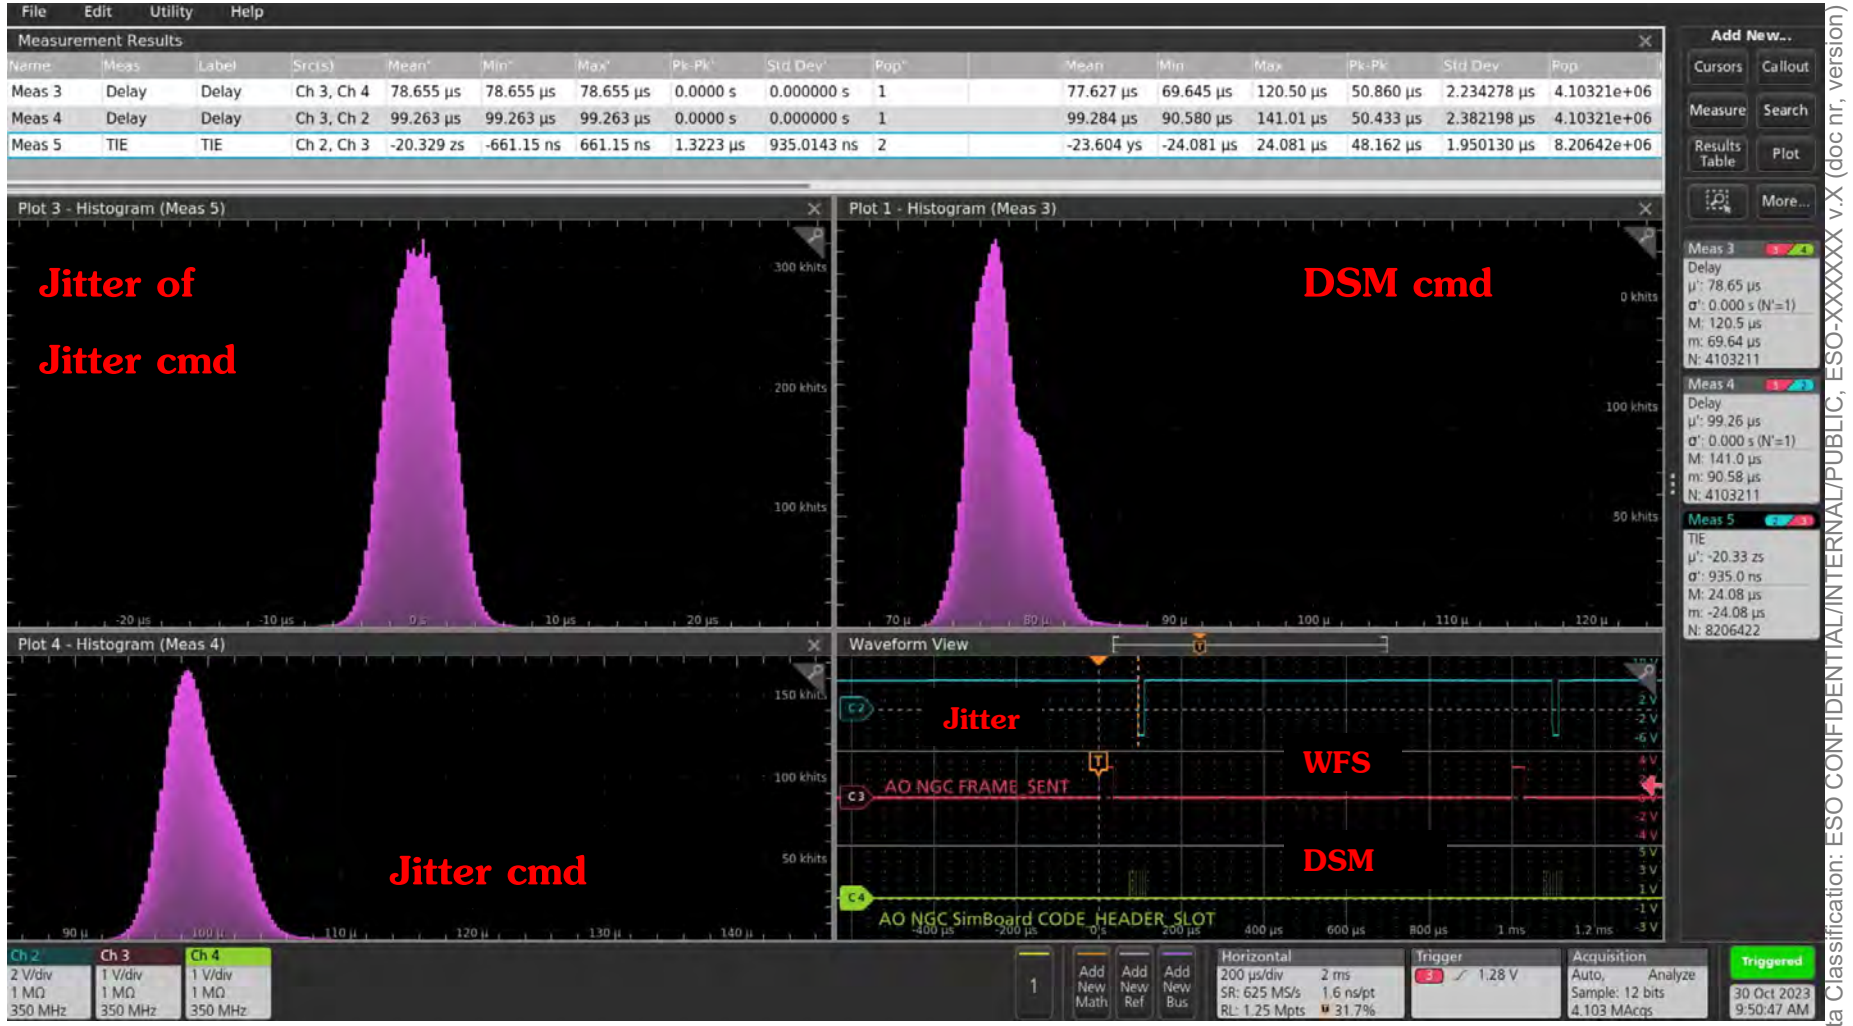

Performance results comparison

Old SPARTA RTC box

- One time measurement
- FPGA system with minimum possible jitter

- **LGS HO loop:**
  - Mean “Latency”: 268 us

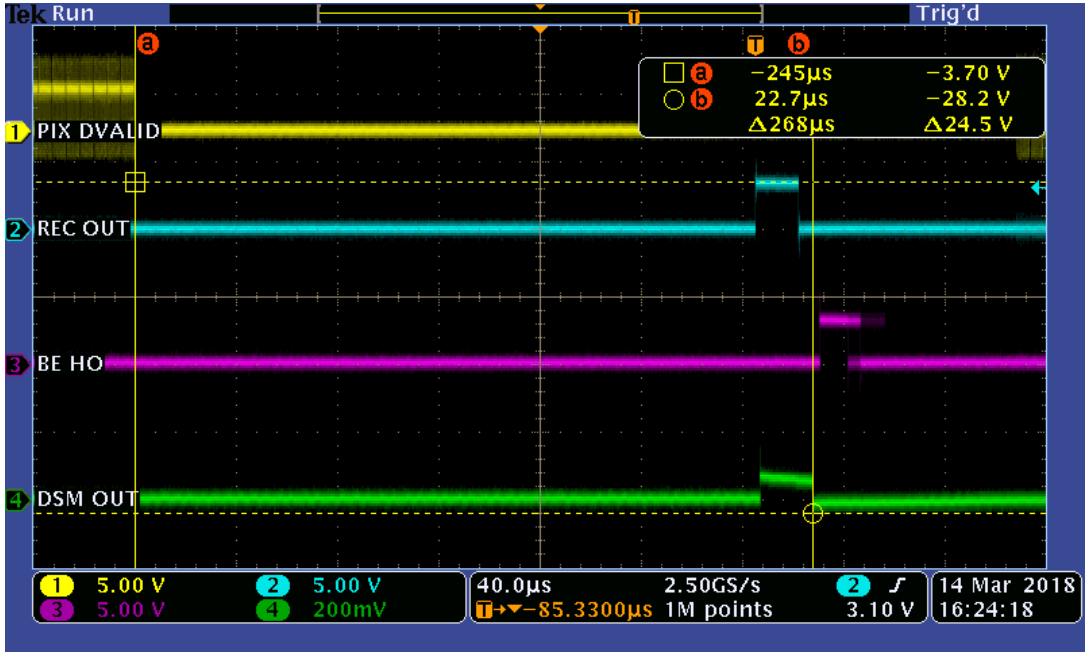

SPARTA Upgrade in WFM mode

- Continuedly running the system for ~12 hours
- LGS, Jitter, LO loops were closed with 0 gain

- ✓ **LGS HO loop statistics:**
  - ✓ Mean: 74.6 us
  - ✓ Std. Dev.: 1.98 us
  - ✓ Pk-Pk.: 19.4 us
  - ✓ Min: 68.3 us
  - ✓ Max: 87.7 us

- ✓ **Jitter loop statistics:**
  - ✓ Mean: 96.1 us
  - ✓ Std. Dev.: 2.17 us
  - ✓ Pk-Pk.: 19.3 us
  - ✓ Min: 88.3 us
  - ✓ Max: 107.7 us

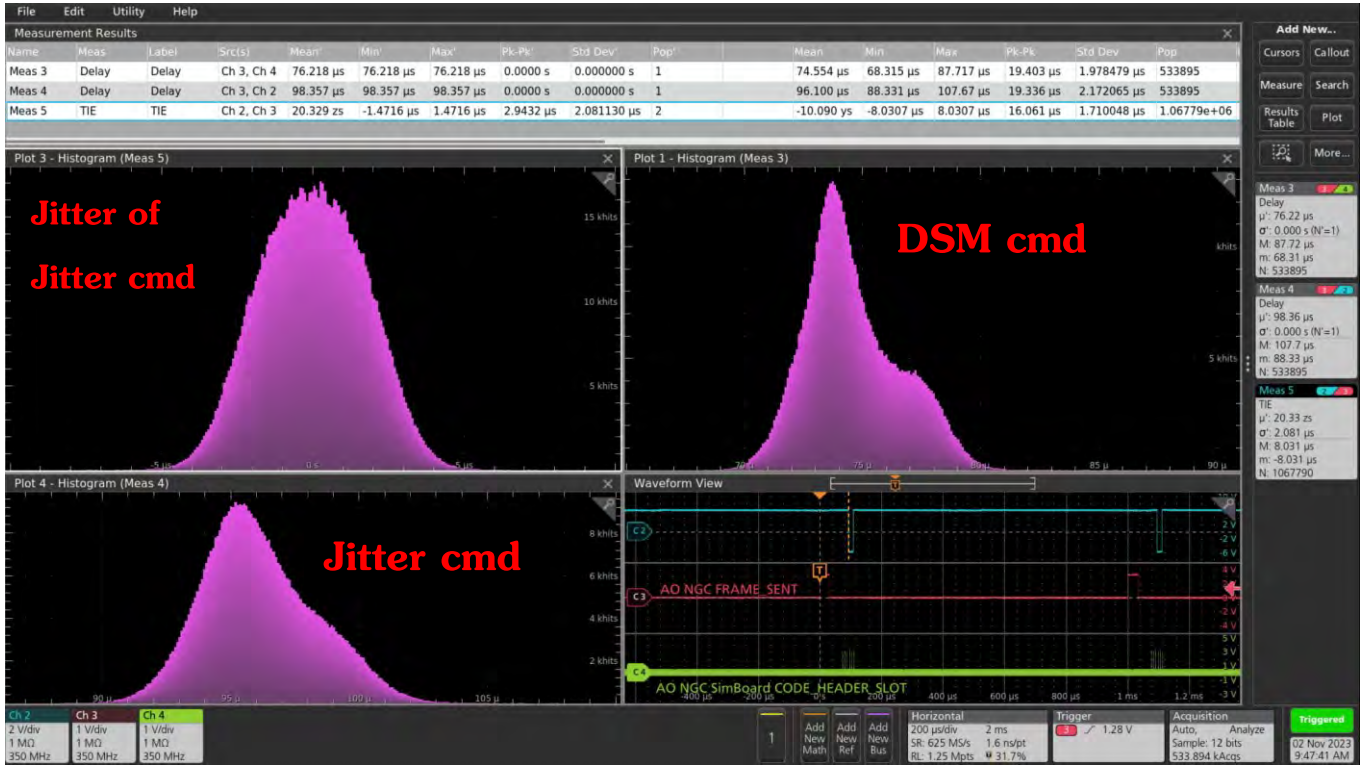

GALACSI SPARTA Upgrade Paranal setup

- Cluster HW is unmodified
- Old SPARTA legacy RTC box and new HRTC are connected such that remote swap of RTCs is possible without HW intervention, thanks to the zero-latency switch

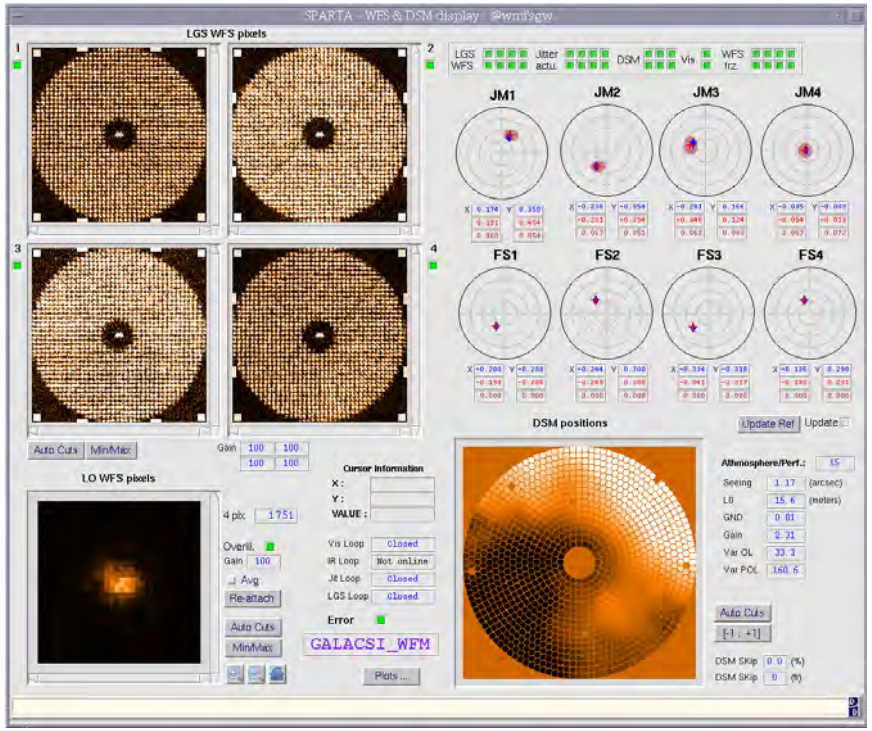

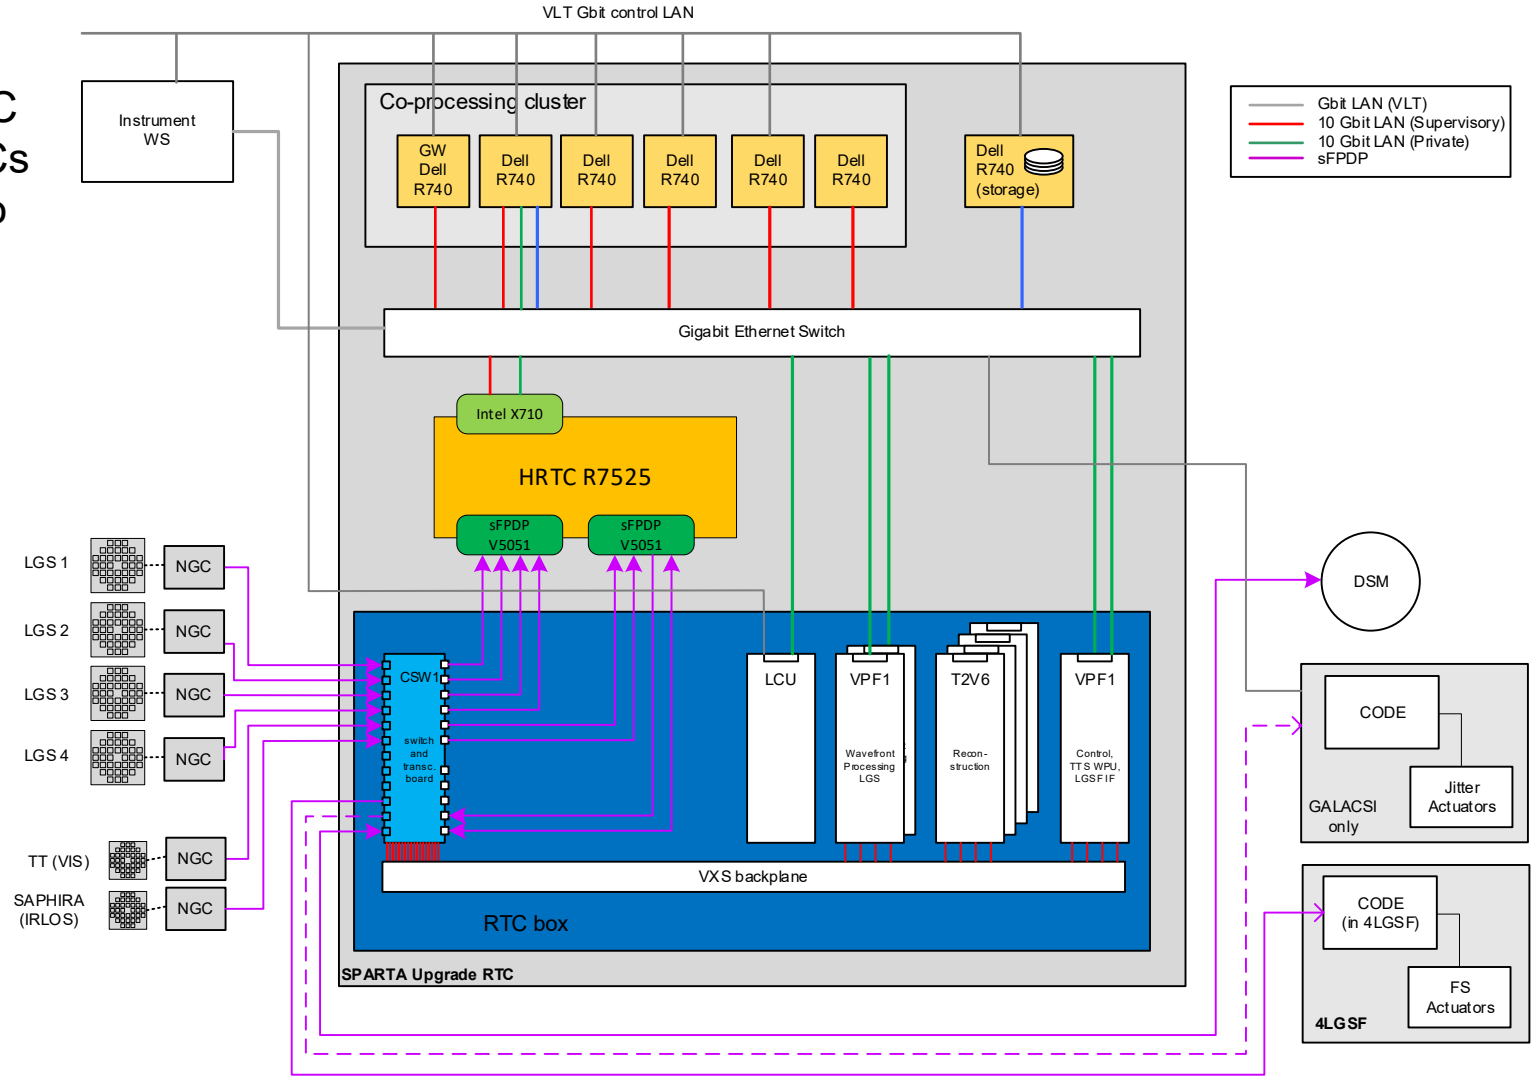

GALACSI SPARTA Upgrade “latency” measurements with timestamping on the real system at Paranal

- **Test preconditions:**
  - Using RDTCS CPU clock
  - Timestamping does not include TX transmission part
  - Reconstruction was executed on dedicated NUMA nodes using 8 cores per MVM
  - Jitter loop is closed on calibration source. LGS and IR loops are closes with 0 gain.
  - IR loop mode: 32x32_SmallScale_500Hz_LowGain

- ✓ **HO loop “latency” statistics:**
  - ✓ Mean: 64.24 us
    Std. Dev.: 1.19 us
    Min: 58.9 us
    Max: 79.13 us

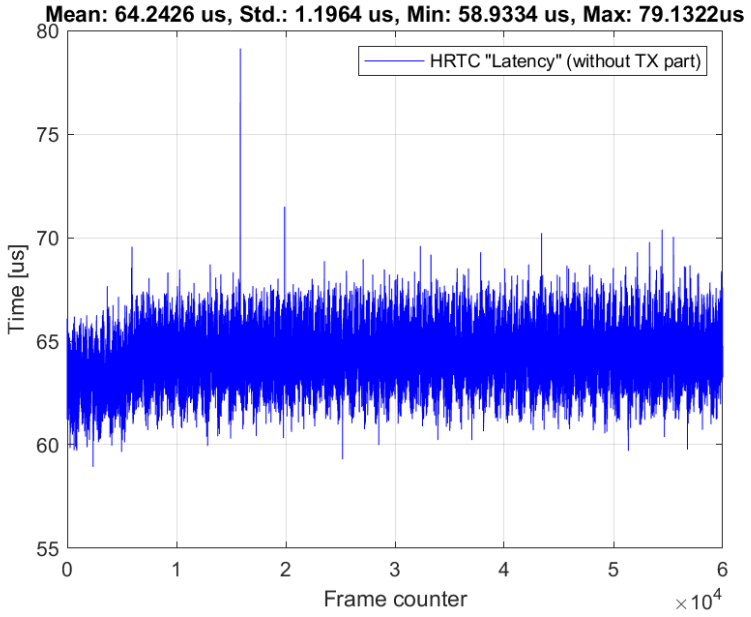

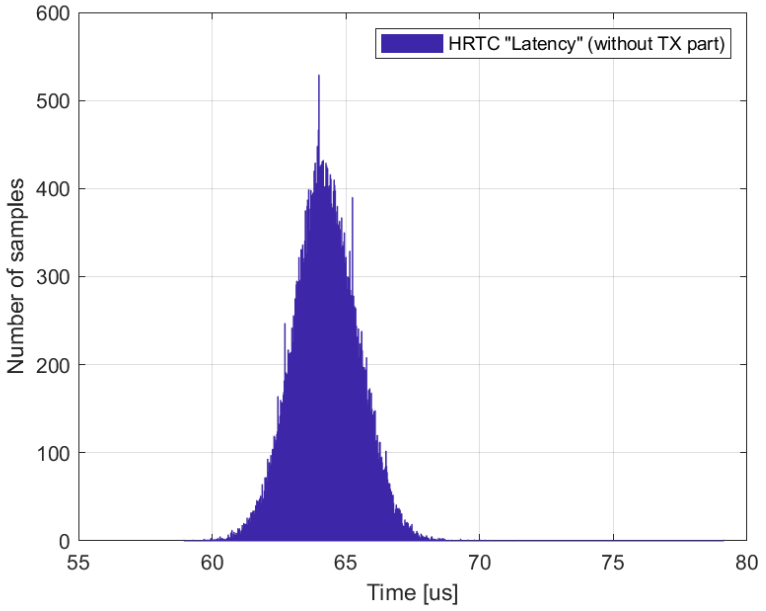

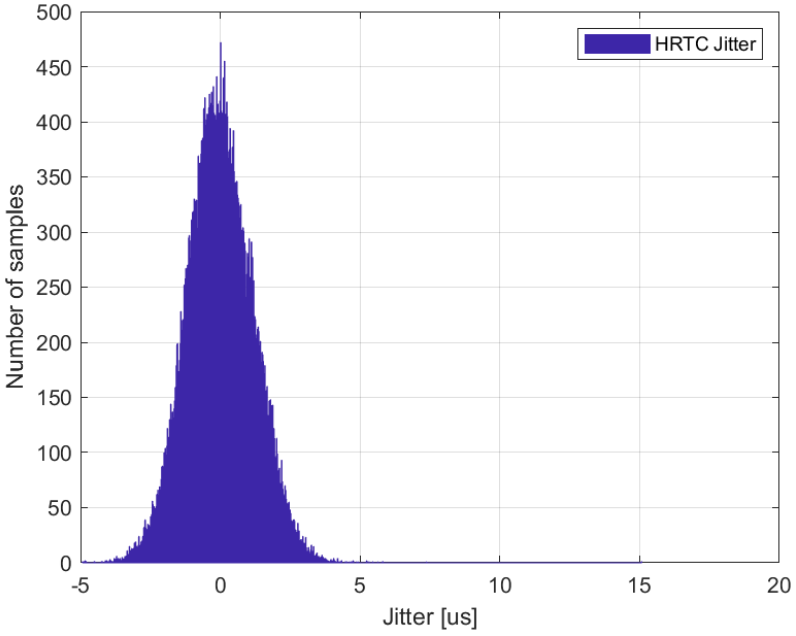

Comparison of Transfer functions measurements 1/2

NFM mode, LGS/HO loop, on sky measurements

❑Old SPARTA RTC box, LGS HO IIR gain -0.3

- ✓ Overall estimated (WFS readout, RTC and DSM) delay **~1.212 ms**
  - ✓ 720 us CCD220 readout time
  - ✓ 120 us DSM delay
  - ✓ **268 us RTC delay**
  - ✓ 104 us unknow?

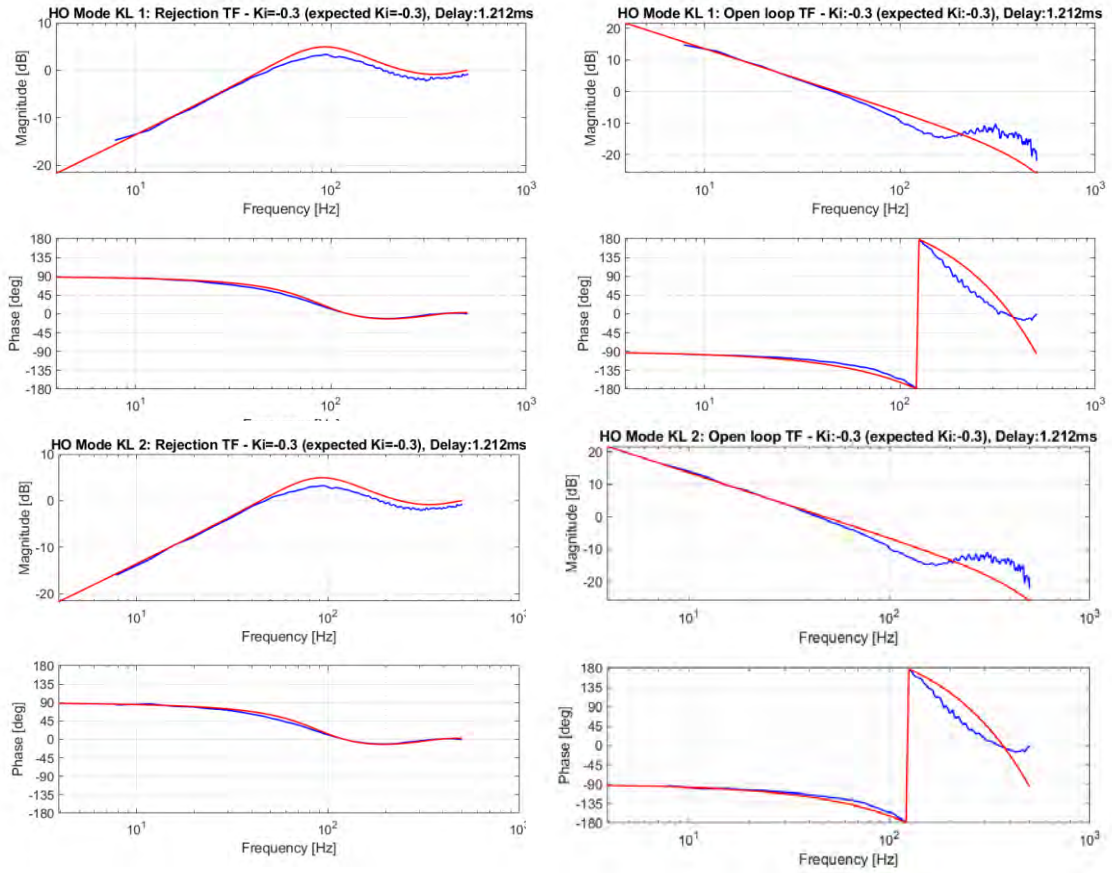

❑SPARTA Upgrade, HO IIR gain -0.4, meas. gain -0.35

- ✓ Overall estimated (WFS readout, RTC and DSM) delay **~1.054 ms**
  - ✓ 720 us CCD220 readout time
  - ✓ 120 us DSM delay
  - ✓ **100 us RTC delay** (including the DSM transmission delay)
  - ✓ 114 us unknow? but it is consistent with old SPARTA RTC box

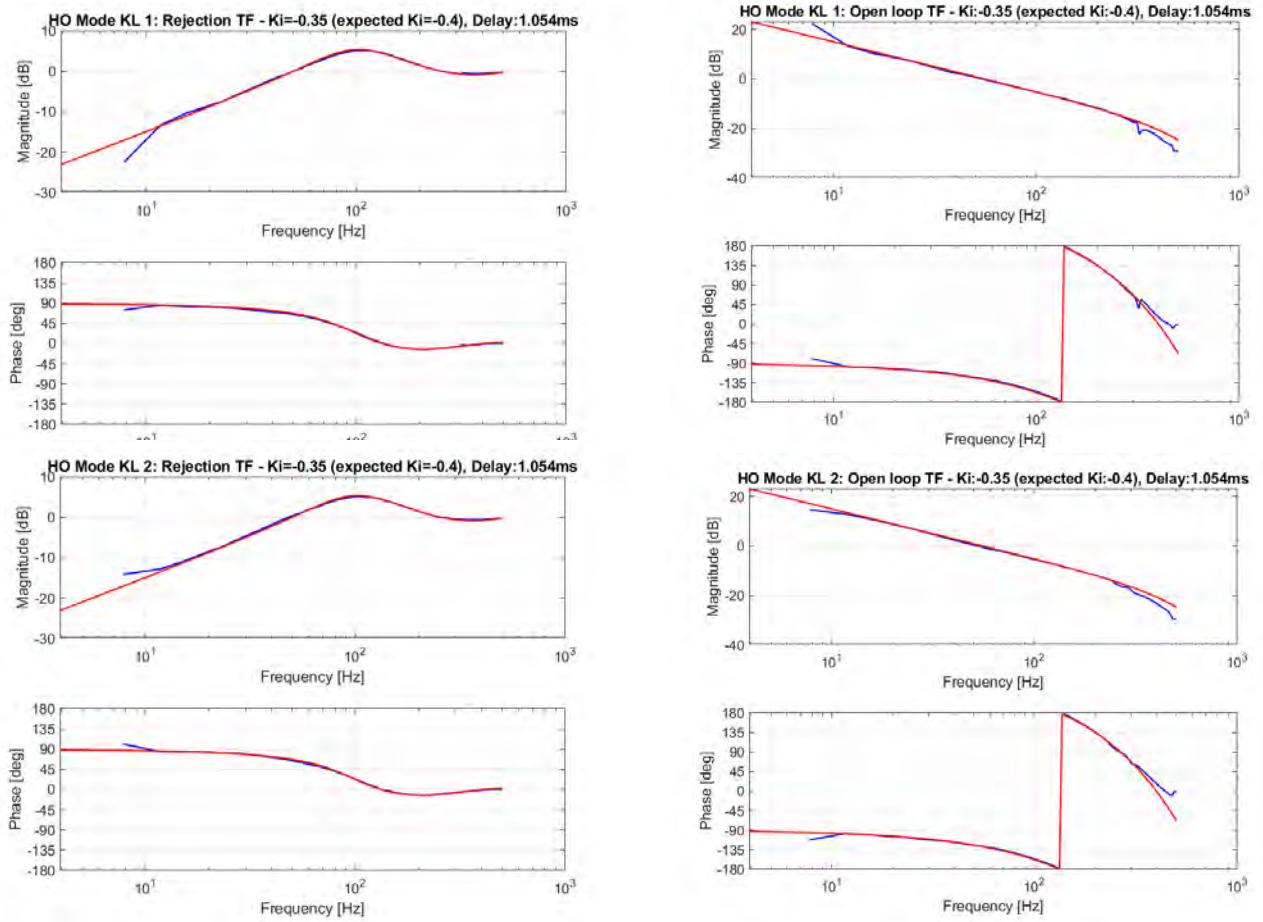

Comparison of Transfer functions measurements 2/2

NFM mode, IRLOS loop, 496Hz, Ki=-0.4

❑Old SPARTA RTC box

- ❑Loop Delay 2.9 ms
- ❑Cut-off frequency: 14Hz

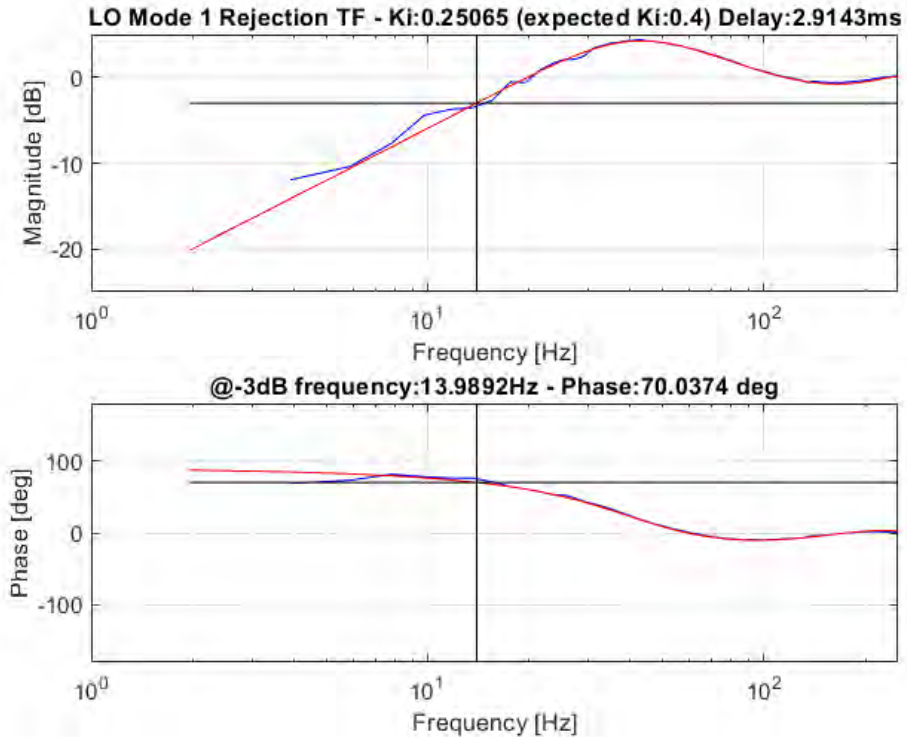

❑SPARTA Upgrade

- ❑Loop Delay 1.4 ms
- ❑Cut-off frequency: 17.4Hz

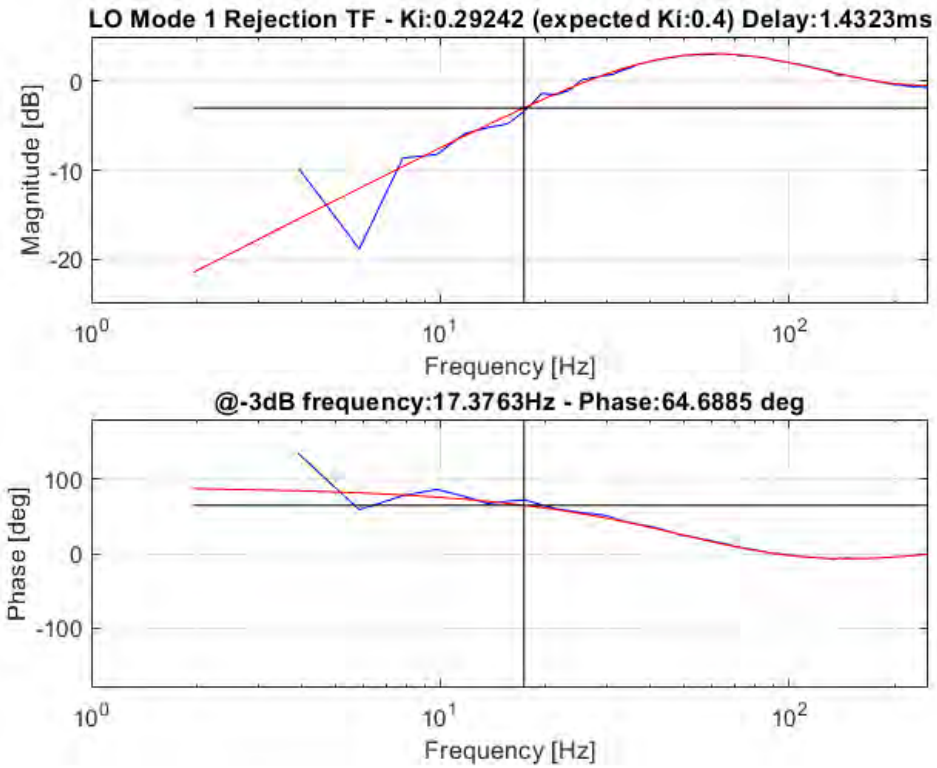

SPARTA Upgrade Project timeline

- January 2019 - SPARTA Red Flag report about Degradation of VPF1 boards lifetime was raised due to the increased failure rate
- Q1 2019 – Q3 2020 – Prototyping phase
- November 2020 – SPARTA Upgrade Project proposal was submitted
- January 2021 – July 2023 – GALACSI SPARTA Upgrade Development phase
- July 2023 - GALACSI SPARTA Upgrade HW installation at Paranal, first on sky test
- October 2023 - GALACSI SPARTA Upgrade commissioning
- **December 2023 - GALACSI SPARTA Upgrade post commissioning AI (current phase)**
- February 2024- GRAAL SPARTA Upgrade HW installation at Paranal
- Q3-Q4 2024 – GRAAL SPARTA Upgrade commissioning
- 2025 – ERIS SPARTA Upgrade HW installation and commissioning
- 2025 – SPHERE SPARTA Upgrade ?

GALACSI SPARTA Upgrade remaining issues & tasks

- Issue with the Jitter loop stability when the spots are not centered. Currently the issue is mitigated by forcing the spiral search of the lasers at every preset.
  - Recalibration of Jitter CM and finetune of Jitter control loop gains for the SPARTA upgrade
  - Implementation of the Anti-Windup filter
- Anti-Vibration Control currently is not working. It was not enabled in operation therefore it is not critical, but to be solved and possibly enabled after SPARTA upgrade
- PTP currently is not available but will be upgraded at Paranal March next year 2024, but it is not critical for GALACSI
- Latency monitoring feature implementation during the operation for statistics in the long term
- Configure the LGS/HO MVMs on 15 cores (core 0 is not usable) to improve the latency. The tuning of the CMs update will be required.

Thank you!

Pavel Shchekaturov
pshchek@eso.org

@ESOAstronomy
@esoastronomy
@ESO
european-southern-observatory
@ESOobservatory