LESIA l'Observatoire de Paris | PSL
Laboratoire d'Études Spatiales et d'Instrumentation en Astrophysique
Australian National University

F. Ferreira, A. Sevin, J. Bernard, D. Gratadour

COSMIC RTC platform: architecture, performance and prototyping for MICADO SCAO

WFS & AO Workshop 2020

CONTENT

- COSMIC quick overview
- Software design
- Performance
- Prototyping activities for MICADO

COSMIC RTC PLATFORM

GREENFLASH LEGACY

A common open platform for AO RTC

- Leverage GreenFlash solutions: GPU + FPGA
- Based on heterogeneous architecture to implement main functions
- Increase the readiness level of the CACAO SW stack already used in SCExAO
- Future RTCs based on COSMIC: MICADO, Keck, MAVIS

Partnership & Collaborations

- Academia-industry partnership:
  - Observatoire de Paris
  - Australian National University
  - Swinburne University of Technology
  - Microgate
- International collaborations:
  - HPC centers: KAUST, BSC
  - Subaru Telescope
  - Industry: Nvidia, Thales

SOFTWARE DESIGN

OVERVIEW

- Modularity & interoperability: independent processes, interfaces abstraction
- Scalability: independent processes, unit abstraction
- Configuration & Control: Python based controller & UI

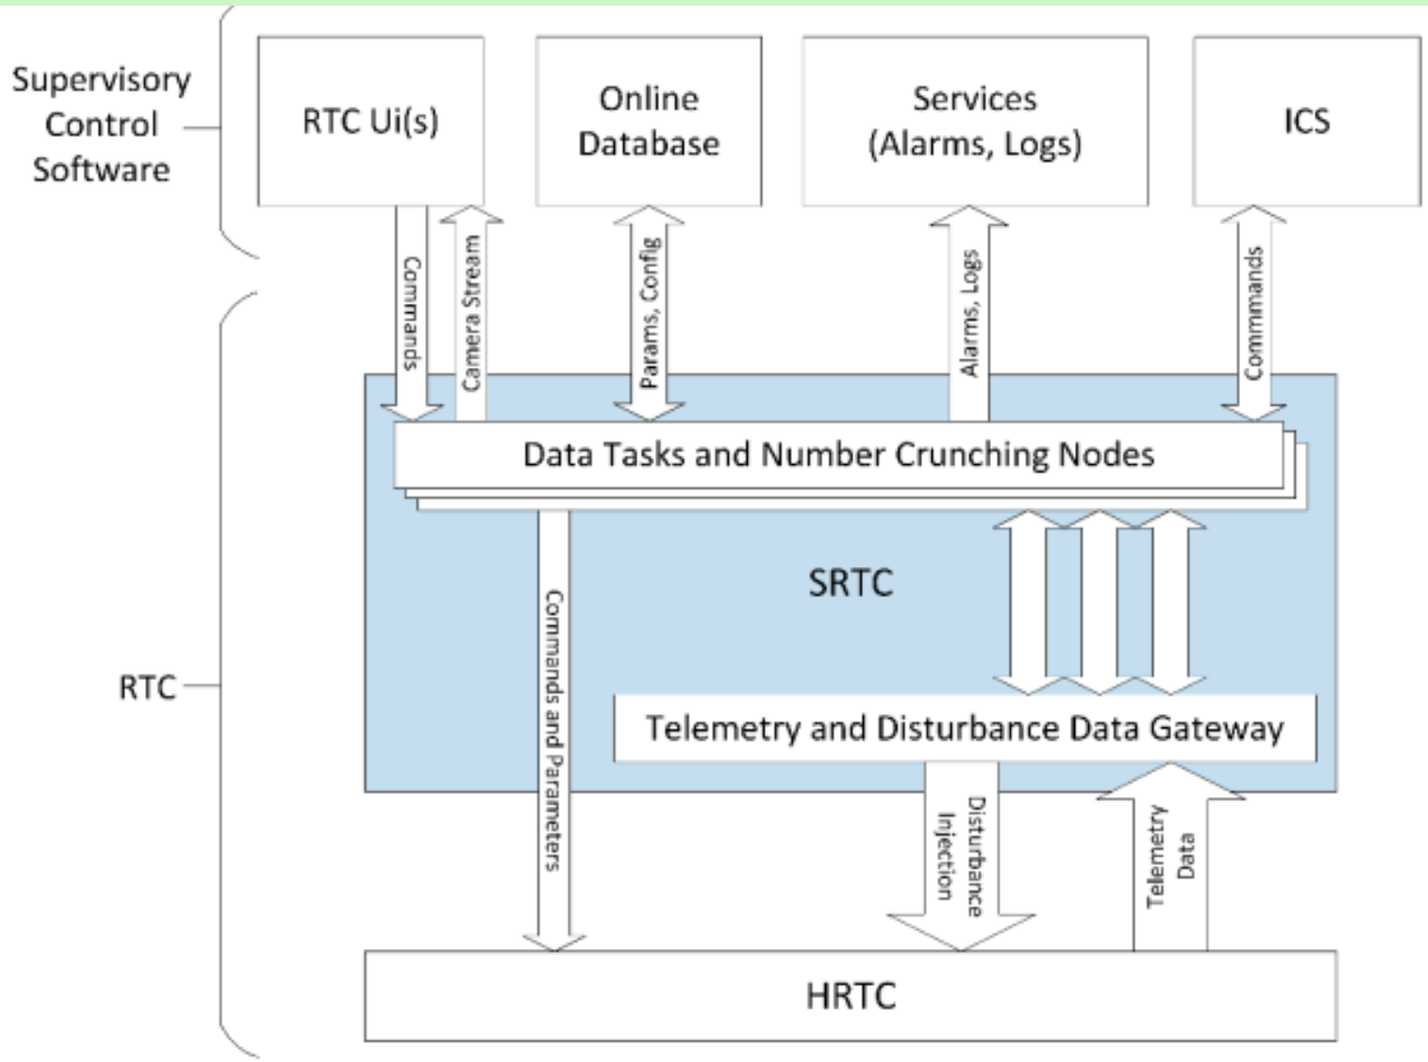

SOFTWARE DESIGN

OVERVIEW

- Modularity & interoperability: independent processes, interfaces abstraction
- Scalability: independent processes, unit abstraction
- Configuration & Control: Python based controller & UI

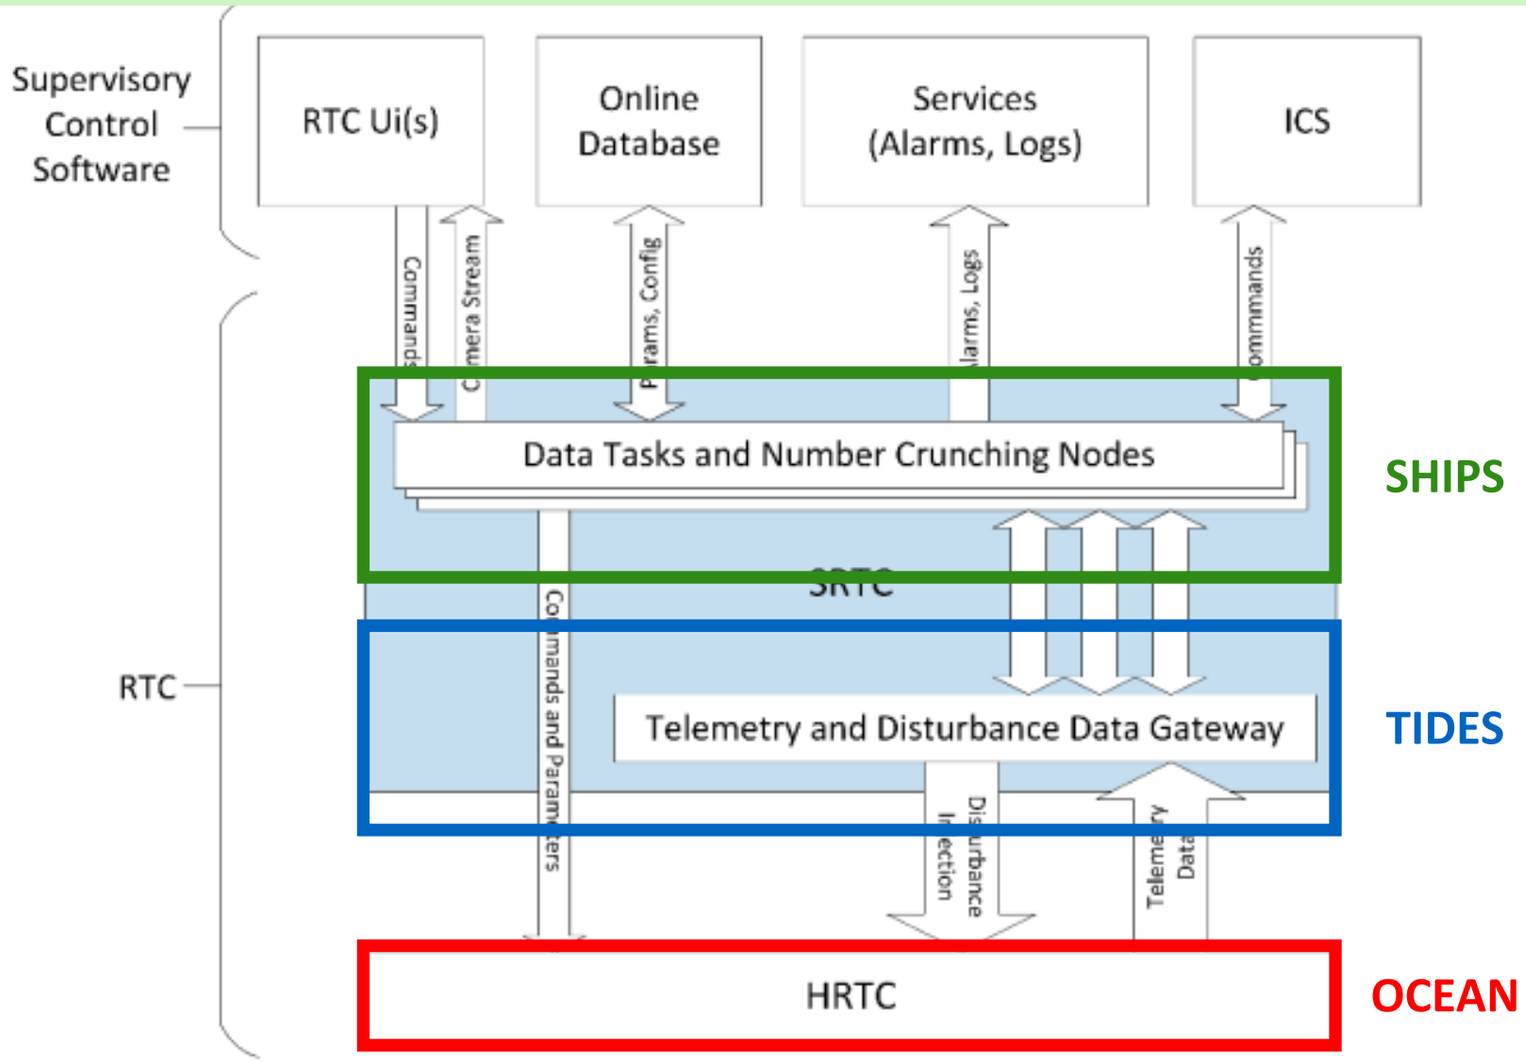

SOFTWARE DESIGN

OCEAN: THE H-RTC

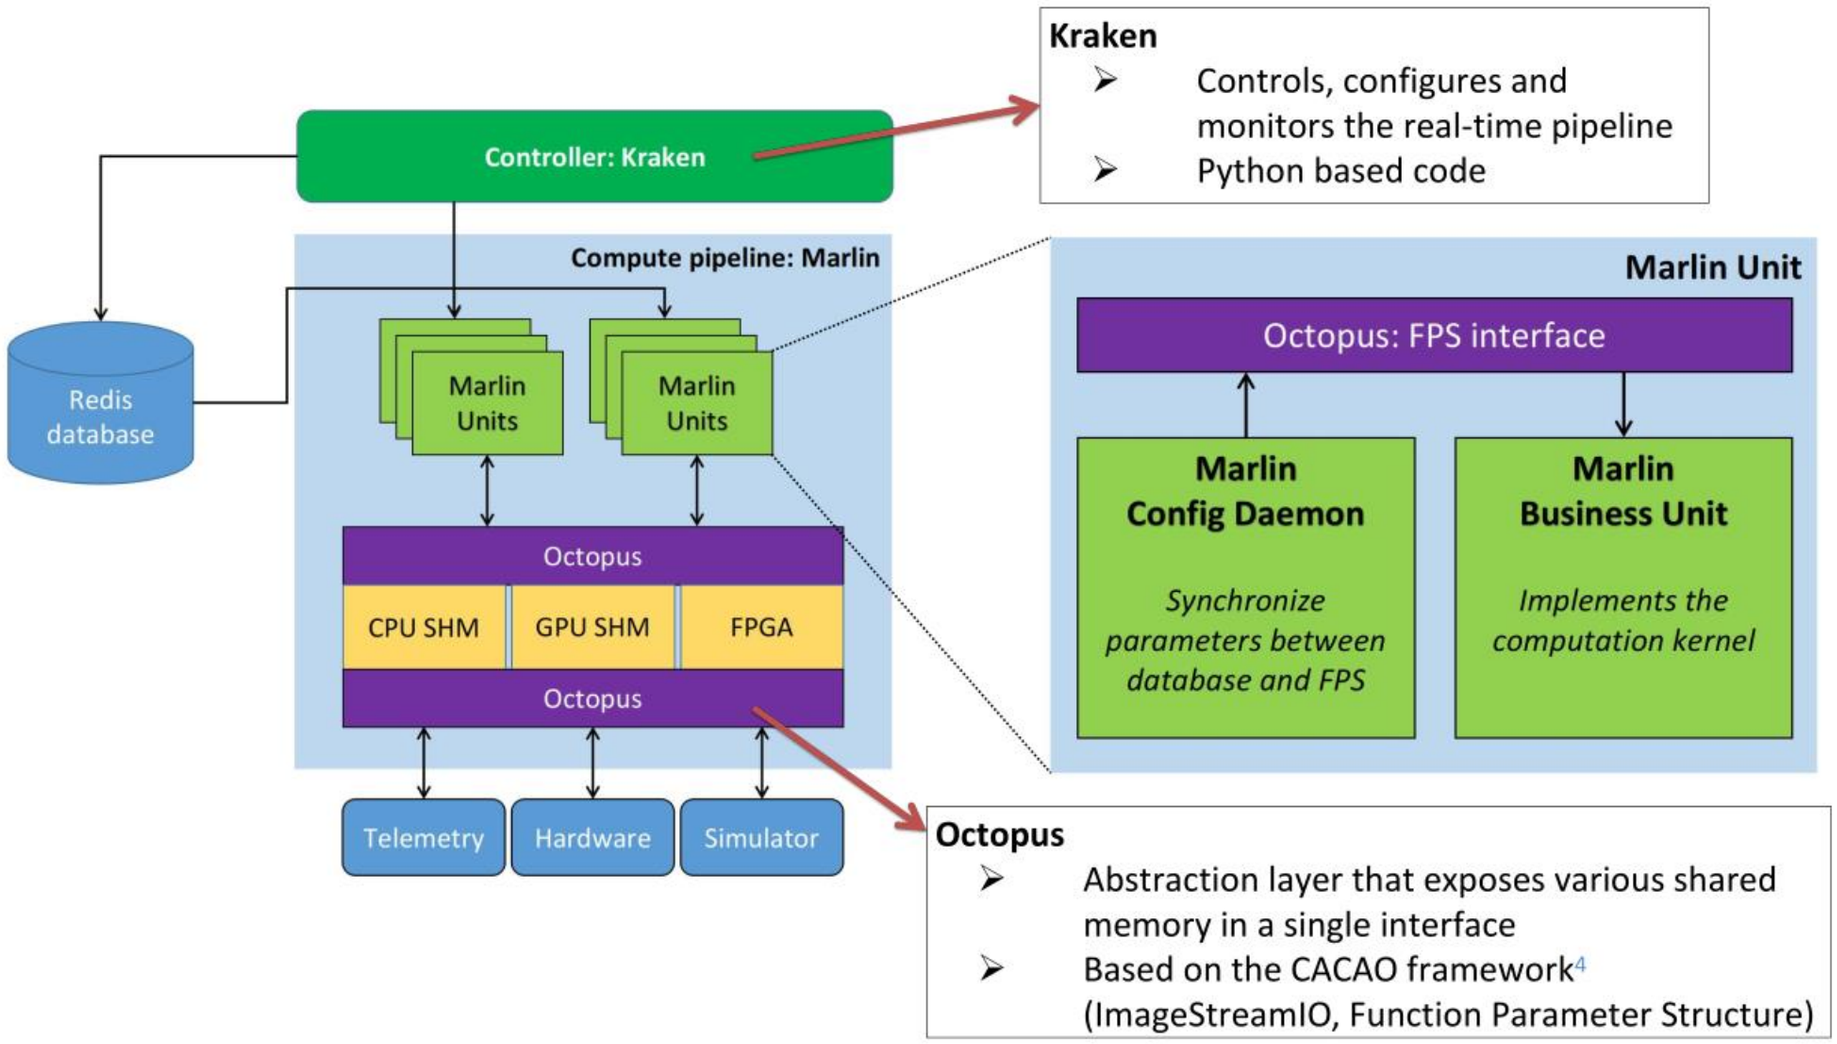

Kraken

- Controls, configures and monitors the real-time pipeline
- Python based code

Octopus

- Abstraction layer that exposes various shared memory in a single interface
- Based on the CACAO framework[4] (ImageStreamIO, Function Parameter Structure)

Marlin

- Defines a container for core real-time components
- Modular and configurable architecture decomposed in business units
- Configuration updated by a daemon that synchronize the parameters between database and FPS (Function Parameter Structure)
- Datas passed by pointer thanks to Octopus double-buffering to ensure low latency

- Marlin: the core real-time software
- Kraken: the controller
- Octopus: the SHM and synchronization manager

SOFTWARE DESIGN

KRAKEN: CONTROL & MONITORING

File Options View Help

Connected

PROCESSES DETAILS SHARE

Name	Status	Control	CPU	Memory	Loop count
CalPixTelemetry (2)	ALIVE	RUNNING	0.1%	188.3 MiB	27158
CoGCentroider (2)	ALIVE	RUNNING	0.1%	1168.1 MiB	27158
ControllerGeneric (2)	ALIVE	RUNNING	0.0%	1486.8 MiB	27158
LoopDataTelemetry (2)	ALIVE	RUNNING	0.0%	188.1 MiB	27158
RawPixTelemetry (2)	ALIVE	RUNNING	0.1%	187.7 MiB	27158
evtcam (2)	ALIVE	RUNNING	1.0%	179.3 MiB	27158
kacou (2)	ALIVE	RUNNING	0.1%	170.3 MiB	27158

RUN STEP PAUSE STOP KILL

SOFTWARE DESIGN

KRAKEN: CONTROL & MONITORING

File Options View Help

Connected

PROCESSES DETAILS SHARE

Name	PID	Status	CPU	Memory	Loop ...	Contr...	Tmux session	Loop ...	Statu...	Creation time (UTC)	CPU affinity	Context sw...	Thre...	Message
MarlinBU_CalPixTelemetry	614105	ALIVE	00	181,784 KiB	30419	0	krakenBusiness_CalPixTelem...	1	0	2020-10-12T07:57:...	16,17	144118	12	4 frame(s) lost, reset loop...
MarlinBU_CoGCentroider	613372	ALIVE	00	1,185,184 ...	30419	0	krakenBusiness_CoGCentroid...	1	0	2020-10-12T07:57:...	9,10	314223	4	4 frame(s) lost, reset loop...
MarlinBU_ControllerGeneric	613548	ALIVE	00	1,511,384 ...	30419	0	krakenBusiness_ControllerGe...	1	0	2020-10-12T07:57:...	10,11	240623	6	4 frame(s) lost, reset loop...
MarlinBU_LoopDataTelemetry	613913	ALIVE	00	181,660 KiB	30419	0	krakenBusiness_LoopDataTel...	1	0	2020-10-12T07:57:...	13,14	192888	12	Slopes and commands ar...
MarlinBU_RawPixTelemetry	614300	ALIVE	00	181,324 KiB	30419	0	krakenBusiness_RawPixTele...	1	0	2020-10-12T07:57:...	16,15	94766	12	4 frame(s) lost, reset loop...
MarlinBU_evtcam	614498	ALIVE	00	172,516 KiB	30419	0	krakenBusiness_evtcam	1	0	2020-10-12T07:57:...	8,5,6,7	35320	6	4 frame(s) lost, reset loop...
MarlinBU_kacou	613730	ALIVE	00	163,324 KiB	30419	0	krakenBusiness_kacou	1	0	2020-10-12T07:57:...	12,13	244328	4	4 frame(s) lost, reset loop...
MarlinConfigDaemon_CalPixT...	614093	ALIVE	00	11,060 KiB	1	1	krakenBusiness_CalPixTelem...	0	0	2020-10-12T07:57:...	0,1,2,3,4,18,19,20,21,22,2...	6148803	3	Running...
MarlinConfigDaemon_CoGCe...	613360	ALIVE	00	10,996 KiB	1	1	krakenBusiness_CoGCentroid...	0	0	2020-10-12T07:57:...	0,1,2,3,4,18,19,20,21,22,2...	6252265	3	Running...
MarlinConfigDaemon_Contro...	613536	ALIVE	00	11,068 KiB	1	1	krakenBusiness_ControllerGe...	0	0	2020-10-12T07:57:...	0,1,2,3,4,18,19,20,21,22,2...	6226754	3	Running...
MarlinConfigDaemon_LoopD...	613901	ALIVE	00	11,048 KiB	1	1	krakenBusiness_LoopDataTel...	0	0	2020-10-12T07:57:...	0,1,2,3,4,18,19,20,21,22,2...	6175046	3	Running...
MarlinConfigDaemon_RawPix...	614288	ALIVE	00	10,928 KiB	1	1	krakenBusiness_RawPixTele...	0	0	2020-10-12T07:57:...	0,1,2,3,4,18,19,20,21,22,2...	6122675	3	Running...
MarlinConfigDaemon_evtcam	614486	ALIVE	00	11,072 KiB	1	1	krakenBusiness_evtcam	0	0	2020-10-12T07:57:...	0,1,2,3,4,18,19,20,21,22,2...	6096047	3	Running...
MarlinConfigDaemon_kacou	613718	ALIVE	00	11,048 KiB	1	1	krakenBusiness_kacou	0	0	2020-10-12T07:57:...	0,1,2,3,4,18,19,20,21,22,2...	6195374	3	Running...

RUN STEP PAUSE STOP KILL

SOFTWARE DESIGN

KRAKEN: CONTROL & MONITORING

cosmic@poseidon:~/cosmic/ocean/kraken

```
pindexSelected = 0    1
[PID 613177   SCAN TID 613177]  36 cpus   14 processes tracked    Display Mode 5
[h] Help   [F2] CTRL   [F3] Resources   [F4] Triggering   [F5] Timing   [F6] htop (F10 to exit)   [F7] iotop (q to exit)   [F8] atop (q to exit)
Display frequ = 32 Hz  [1174] fscan= 0.99 Hz (  1.00 Hz  1.42 % busy )
Source Code:  line 0 (function )


   STATUS    PID   process name

 ACTIVE     0614498 MarlinBU_evtcam           ON  60 ..84  ITERlim 0/    0/   0  EXEClim 0/    0/   0  ITER    >10ms [   >10ms -    >10ms ] EXEC    3.7us [   3.3us -    4.6us ]  busy =   0.02 %
 ACTIVE     0614486 MarlinConfigDaemon_evtcamOFF
 ACTIVE     0614300 MarlinBU_RawPixTelemetry  ON  60 ..84  ITERlim 0/    0/   0  EXEClim 0/    0/   0  ITER    >10ms [   >10ms -    >10ms ] EXEC  149.0us [  99.4us -  220.8us ]  busy =   0.74 %
 ACTIVE     0614288 MarlinConfigDaemon_RawPixOFF
 ACTIVE     0614105 MarlinBU_CalPixTelemetry  ON  60 ..84  ITERlim 0/    0/   0  EXEClim 0/    0/   0  ITER    >10ms [   >10ms -    >10ms ] EXEC  189.2us [ 141.1us -  322.9us ]  busy =   0.95 %
 ACTIVE     0614093 MarlinConfigDaemon_CalPixOFF
 ACTIVE     0613913 MarlinBU_LoopDataTelemetr ON  60 ..84  ITERlim 0/    0/   0  EXEClim 0/    0/   0  ITER    >10ms [   >10ms -    >10ms ] EXEC  141.3us [ 124.2us -  146.8us ]  busy =   0.71 %
 ACTIVE     0613901 MarlinConfigDaemon_LoopDaOFF
 ACTIVE     0613730 MarlinBU_kacou            ON  60 ..84  ITERlim 0/    0/   0  EXEClim 0/    0/   0  ITER    >10ms [   >10ms -    >10ms ] EXEC  241.0us [ 237.9us -  245.0us ]  busy =   1.20 %
 ACTIVE     0613718 MarlinConfigDaemon_kacou OFF
 ACTIVE     0613548 MarlinBU_ControllerGeneri ON  60 ..84  ITERlim 0/    0/   0  EXEClim 0/    0/   0  ITER    >10ms [   >10ms -    >10ms ] EXEC  348.6us [ 336.8us -  359.0us ]  busy =   1.74 %
 ACTIVE     0613536 MarlinConfigDaemon_ControOFF
 ACTIVE     0613372 MarlinBU_CoGCentroider    ON  60 ..84  ITERlim 0/    0/   0  EXEClim 0/    0/   0  ITER    >10ms [   >10ms -    >10ms ] EXEC   27.8us [  26.6us -   30.2us ]  busy =   0.14 %
 ACTIVE     0613360 MarlinConfigDaemon_CoGCenOFF

Loop time = 0.00031031 s  ( max rate = 3222.58 Hz)
```

```
----------------------------------------------------------------------------[PID 624071] STREAM MONITOR: PRESS (x) TO STOP, (h) FOR HELP----------------------------------------------------------------------------
[h] Help   [F2] sem values   [F3] write PIDs   [F4] read PIDs   [F5] process traces   [F6] access
PIDmax = 4194304    Update frequ = 32 Hz  fscan=19.96 Hz ( 20.00 Hz  0.20 % busy )
  24 streams    Currently displaying    0-  23   Selected 0  ID = 21  inode = 8141888
    inode  name                                type                    cnt0    creaPID  ownPID    frequ       #sem     Semaphore values ....
  8141888 cosmic_CalWFSframeInterface          FLT [400x400x  1]        58469   613037        0  59.88 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8141965 cosmic_commandsInterface             FLT [3229x  1x  1]       58468   613037        0  59.88 Hz  10 sems    1  0  1  1  1  1  1  1  1  1
  8141848 cosmic_WFSframeInterface            UI16 [400x400x  1]        58469   613037        0  59.88 Hz  10 sems    1  0  1  1  1  1  1  1  1  1
  8141876 cosmic_slopesInterface               FLT [22093x  1x  1]       58469   613037        0  59.88 Hz  10 sems    0  1  1  1  1  1  1  1  1  1
  7266161 cosmic_MarlinBU_ControllerGeneric_Ti DBL [  2x  1x1000000]     58459   613548        0  59.88 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8142025 cosmic_endOfPipelineInterface        FLT [  1x  1x  1]         58459   613037        0  59.88 Hz  10 sems    1  0  0  0  1  1  1  1  1  1
  5892722 cosmic_MarlinBU_CoGCentroider_Time   DBL [  2x  1x1000000]     58460   613372        0  59.88 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  7344925 cosmic_MarlinBU_kacou_Time           DBL [  2x  1x1000000]     58459   613730        0  59.88 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  7212029 cosmic_MarlinBU_RawPixTelemetry_Time DBL [  2x  1x1000000]     58459   614300        0  59.88 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8001627 cosmic_MarlinBU_evtcam_Time          DBL [  2x  1x1000000]     58460   614498        0  59.88 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  7645760 cosmic_MarlinBU_LoopDataTelemetry_Ti DBL [  2x  1x1000000]     58459   613913        0  59.88 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  5838663 cosmic_MarlinBU_CalPixTelemetry_Time DBL [  2x  1x1000000]     58459   614105        0  59.88 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8141977 cosmic_cmatInterface                 FLT [3228x22092x  2]           1   613037        0   0.00 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8141989 cosmic_subcmatInterface_1            FLT [3228x11046x  2]           1   613037        0   0.00 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8141948 cosmic_lutPixInterface               I32 [160000x  1x  1]           1   613037        0   0.00 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8141912 cosmic_FlatInterface                 FLT [400x400x  1]              1   613037        0   0.00 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8141936 cosmic_validpixInterface             I32 [  2x22092x  1]            1   613037        0   0.00 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8142037 cosmic_biasInterface                 FLT [3228x  1x  1]             1   613037        0   0.00 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8142008 cosmic_cmatInterfaceCPU              FLT [3228x22092x  2]           1   613037        0   0.00 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8142049 cosmic_signInterface                 FLT [3228x  1x  1]             1   613037        0   0.00 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8142093 cosmic_lutPixInterfaceAcqui          I32 [160000x  1x  1]           1   613037        0   0.00 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8142061 cosmic_lutInterface                  FLT [3228x  1x  1]             1   613037        0   0.00 Hz  10 sems    1  1  1  1  1  1  1  1  1  1
  8141900 cosmic_DarkInterface                 FLT [400x400x  1]              0   613037        0   0.00 Hz  10 sems    0  0  0  0  0  0  0  0  0  0
  8141924 cosmic_RefSlopesInterface            FLT [22092x  1x  1]            0   613037        0   0.00 Hz  10 sems    0  0  0  0  0  0  0  0  0  0
```

```
[cosmic_mo0:bash*                                                                                                                    "cosmic@poseidon:~" 10:16 12-Oct-20
```

SOFTWARE DESIGN

TIDES: DATA STREAMING & CONTROL

Streamer

- Abstraction layer defines data format and architecture (publish/subscribe)
- Available implementation based on RTI DDS with no-copy mechanism
- Foreseen implementation based on raw UDP / MUDPI

Data format

- Frame: *WFS raw or calibrated image*
- SuperFrame: *synchronized slopes and DM command vectors*
- MegaFrame: *synchronized Frame and SuperFrame*
- All formats embed frame counter and timestamp

Controller

- Abstraction layer defines architecture (client/server) and request format
- Request aims to access kraken methods
- Available implementation based on ØMQ

PERFORMANCE

MICADO SCALE SETUP

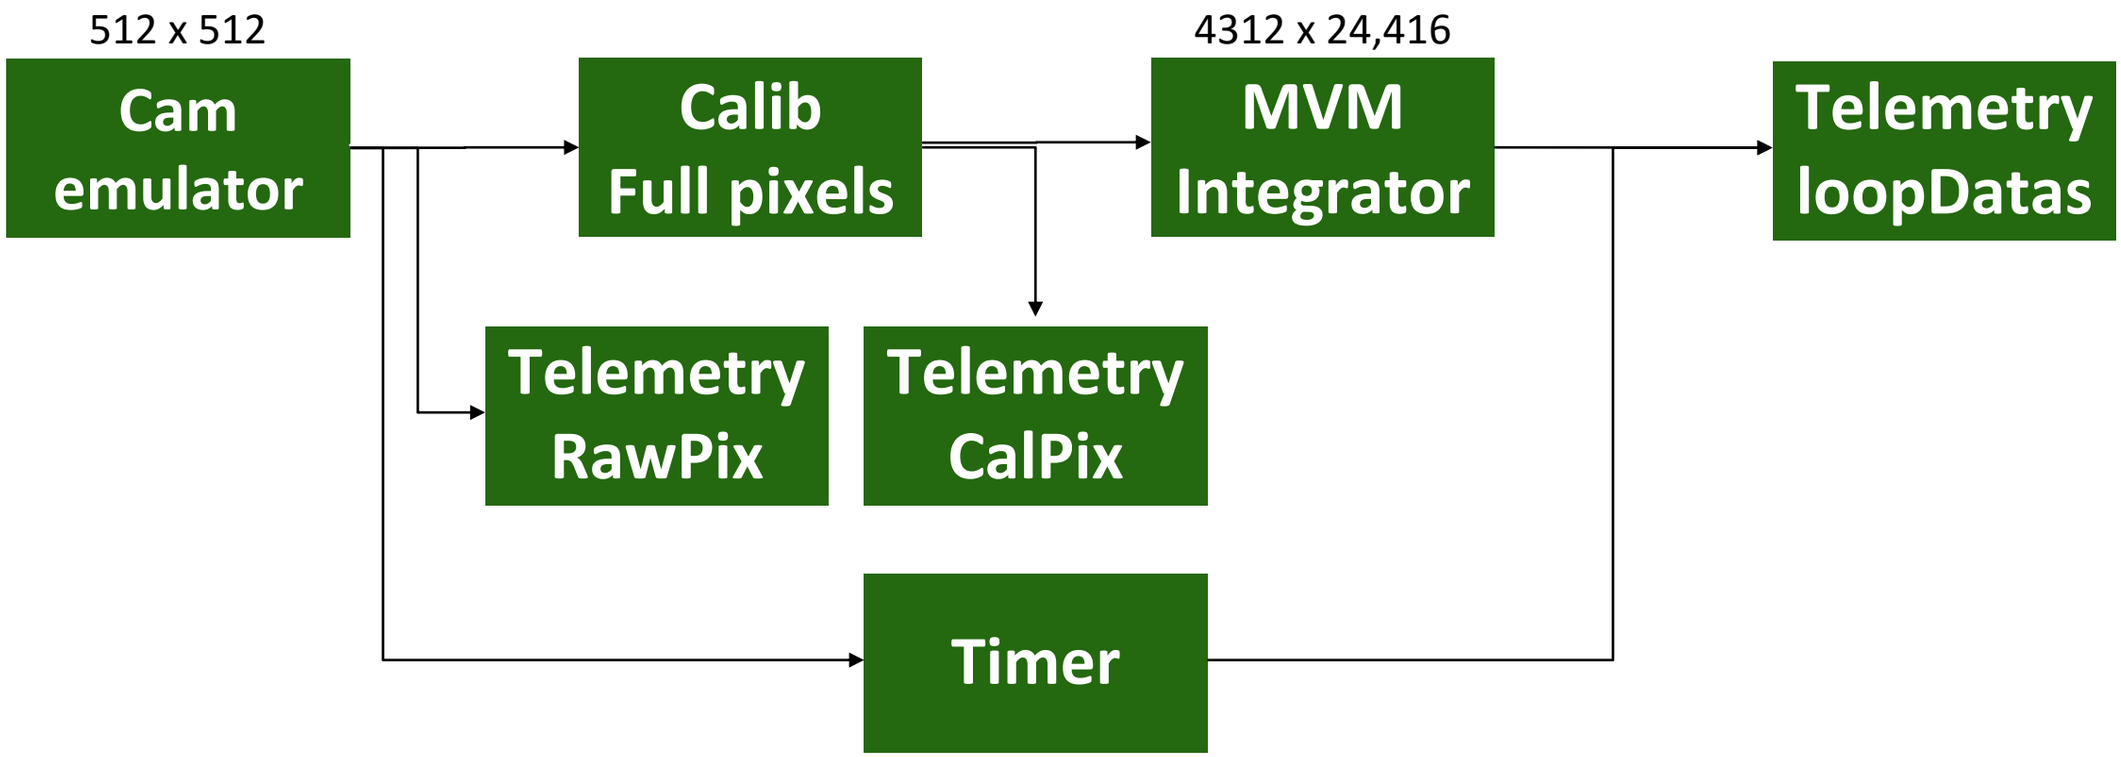

- Timing is done using a dedicated BU and cudaEvents
- Telemetry and data injection is performed during RTC idle
- Injection of large data (command matrix) done by chunks over several iterations
- Core real-time processes are launched with high priority and CUDA MPS
- Launch on DGX-1 V100 over 4 GPUs

PERFORMANCE

MICADO SCALE RESULTS

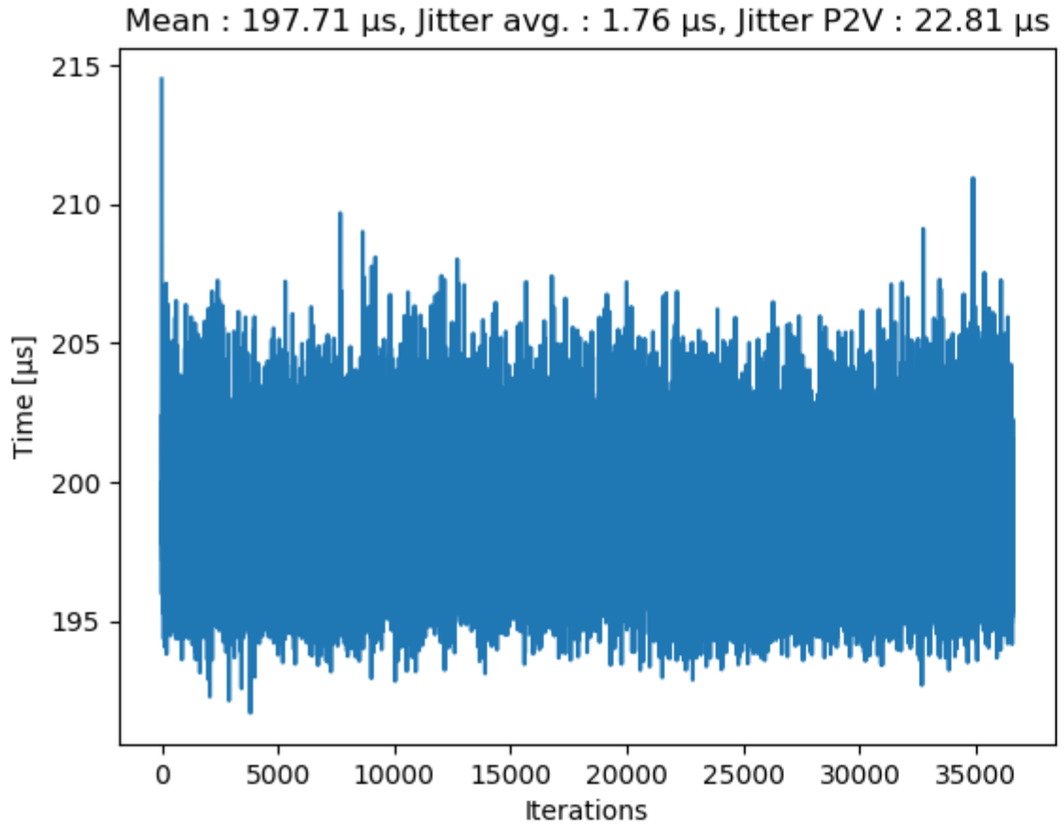

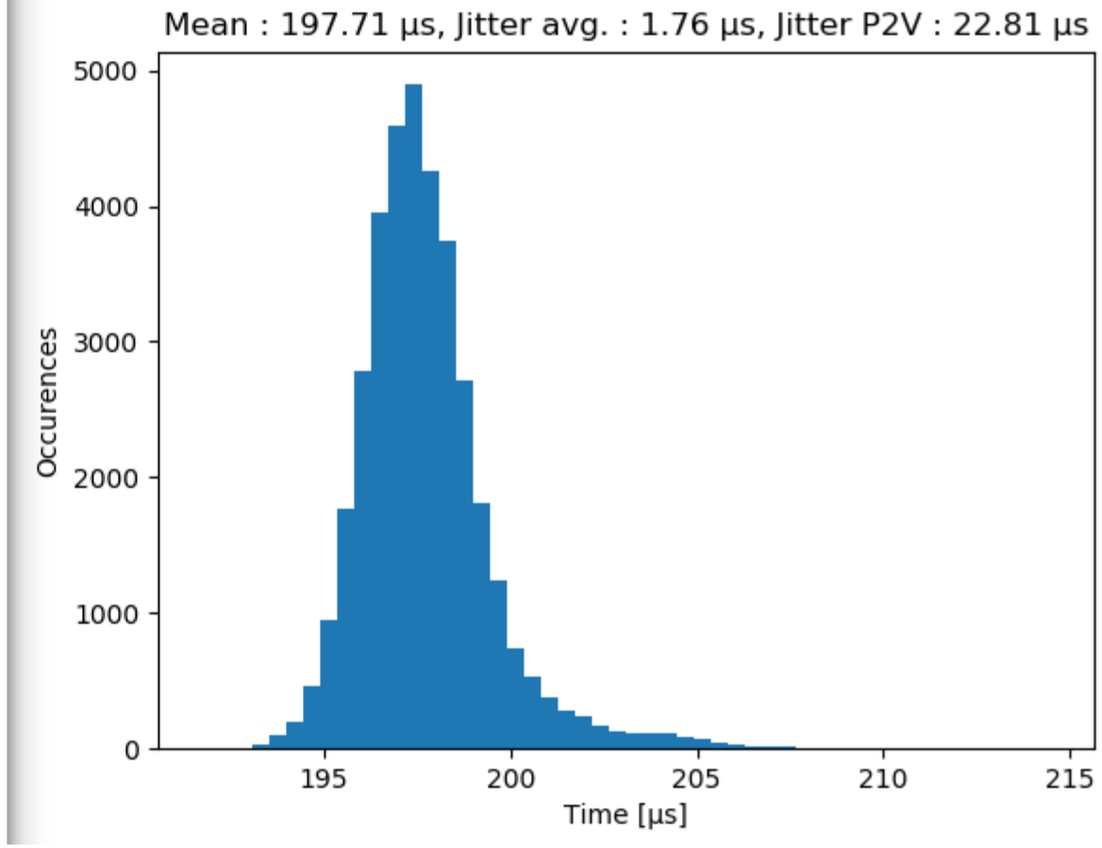

Foreseen performance on new GPU architecture

- Invitation by Nvidia to new DGX-1 A100 proof of concept in august 2020
- Results show up to 40 % performance improvement on COMPASS
- Limited for the RTC because of a CUDA 11 bug (bug report currently under investigation at Nvidia)

PROTOTYPING ACTIVITIES

Server hardware configuration

- 18 x Intel Xeon Gold 5220 CPU @ 2.2 GHz
- 2 x Nvidia Quadro RTX 8000 + NVLINK

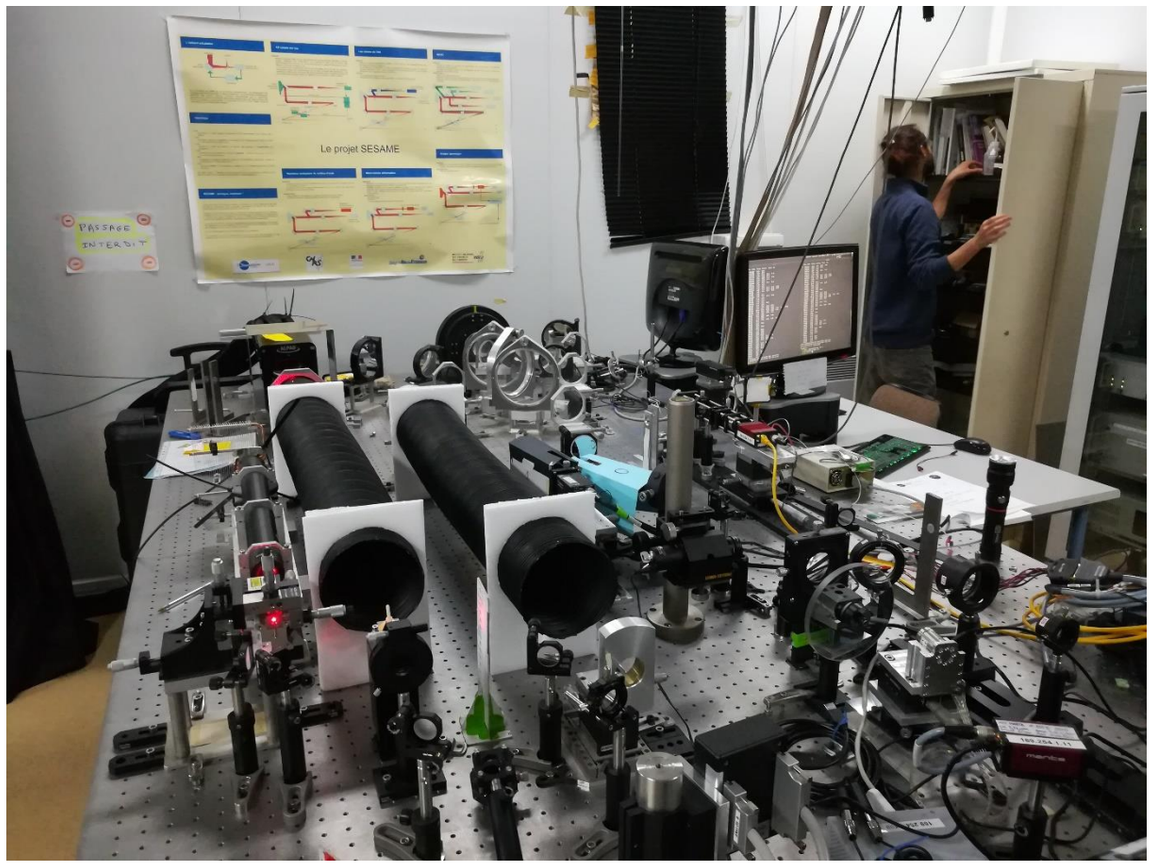

PYRCADO bench

- Pyramid WFS 86x86: 23,148 valid pixels
- Shack-Hartmann WFS 80x80: 12,184 slopes
- ALPAO 4k DM: 3228 actuators
- SLM + COMPASS to produce turbulence

PROTOTYPING ACTIVITIES

FIRST TESTS: LATENCY

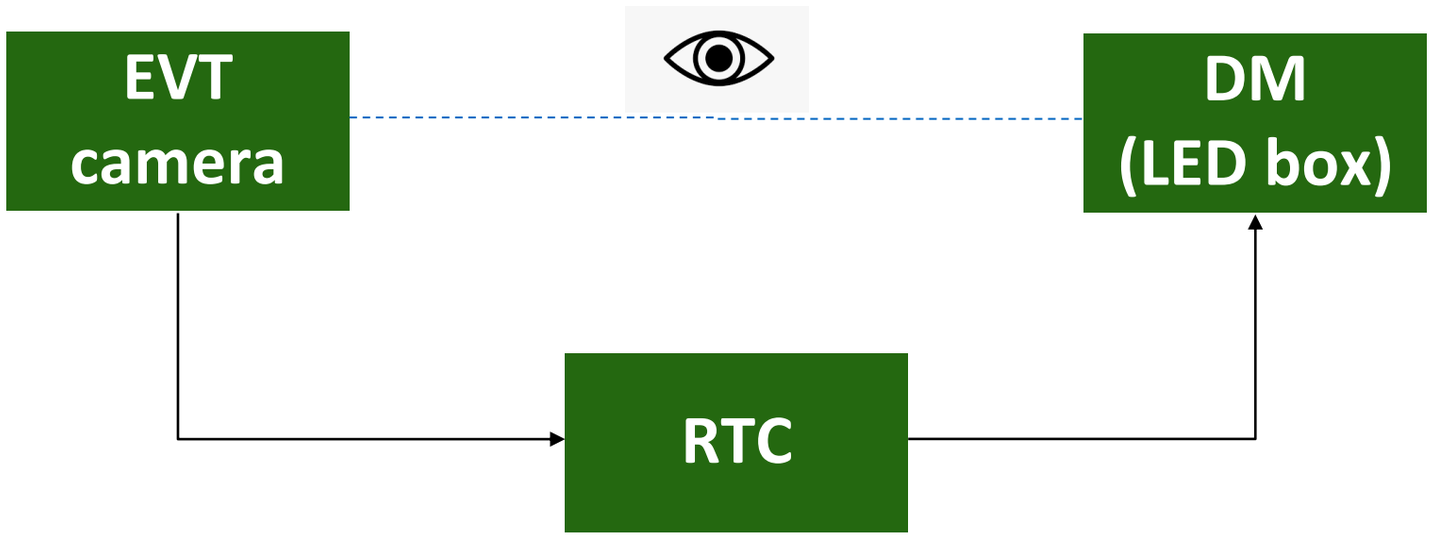

Procedure

- Load a circular buffer of perturbation voltages in the RTC which will light on a LED each 10 iterations
- Retrieve loop data telemetry
- Compare volts and slopes to get the overall latency
- Camera runs at 1 kHz, and RTC is configured for bench scale (24,416 x 3228)

PROTOTYPING ACTIVITIES

FIRST TESTS: LATENCY

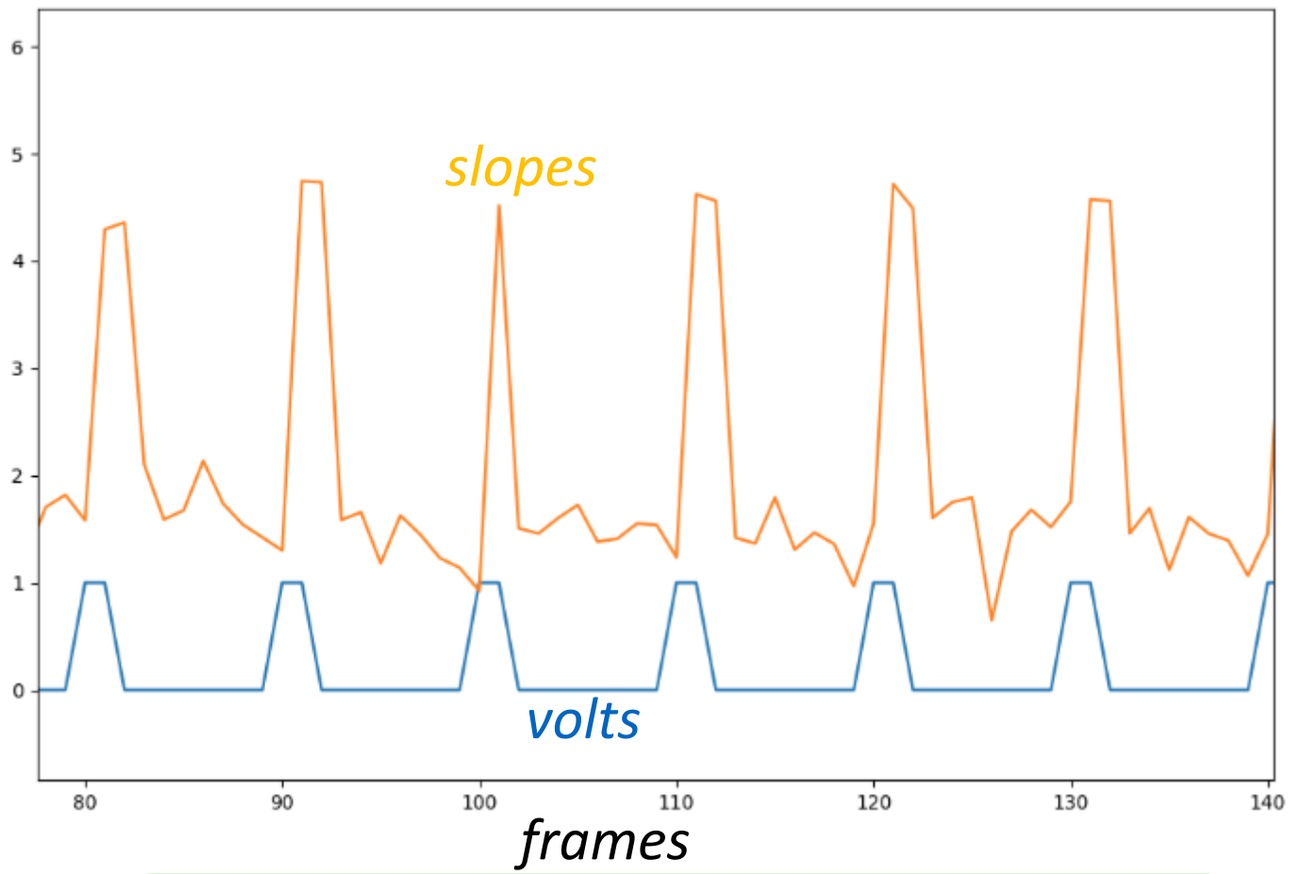

- Slope missing sometimes
- Periodic: each 250 frames

- Camera signals on oscilloscope
- Glitch was due to camera not acquiring for some reason

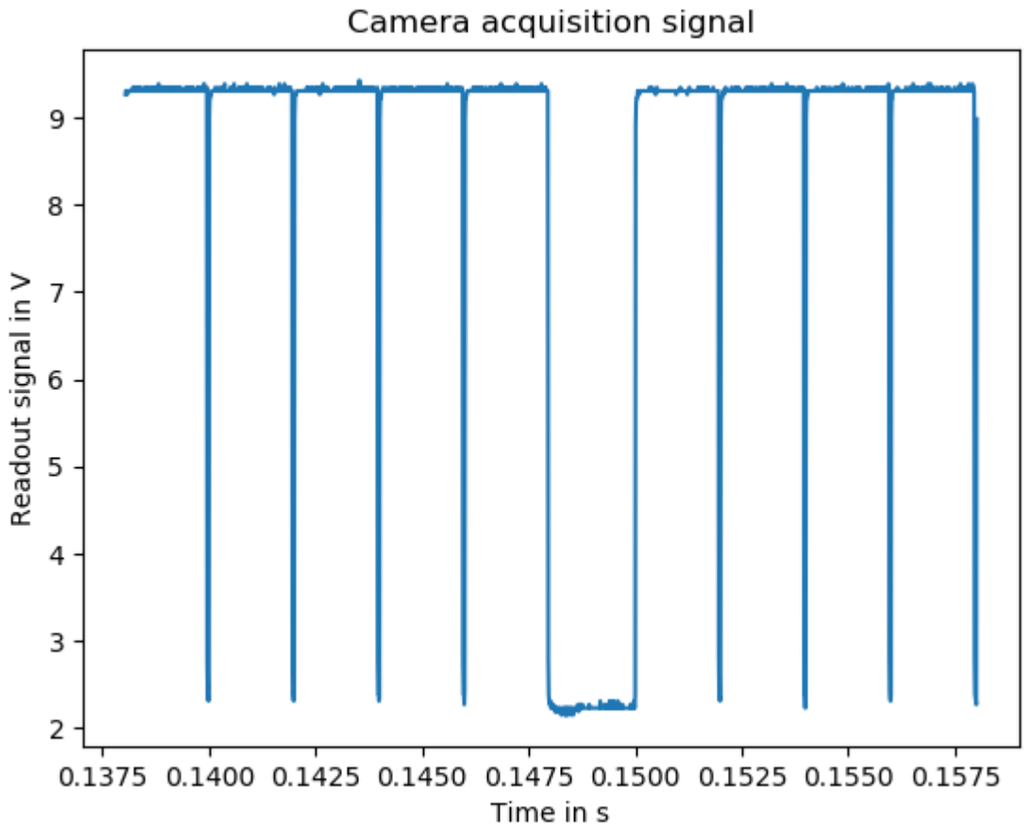

PROTOTYPING ACTIVITIES

WORK IN PROGRESS

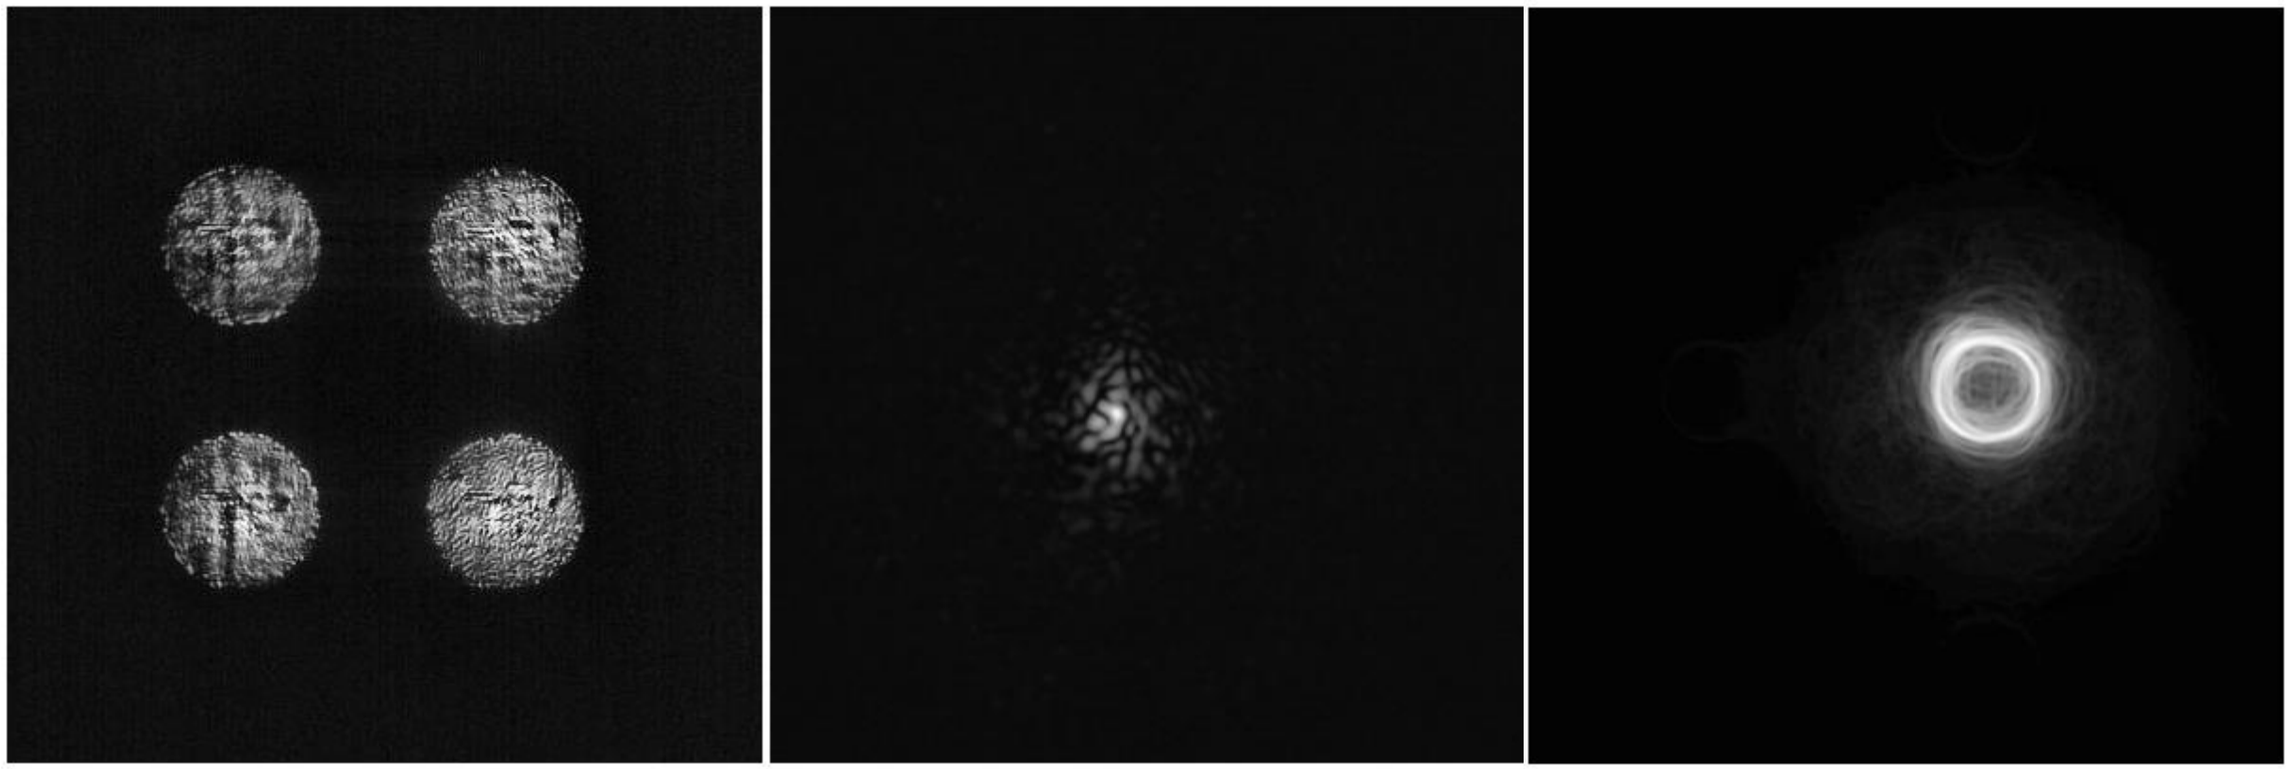

- RTC prototype currently plugged to the PYRCADO bench (PWFS and DM)
- Developments and debugs on-going on in-house “SRTC” software (ADOPT)
- H-RTC located on the prototype, SRTC on workstation: communication over Tides

CONCLUSION

Cosmic SW architecture

- Based on high level of abstraction to reach high modularity and scalability
- Optimized core real-time architecture
- Easy to configure, easy to implement
- Great performance obtained on several scales (MICADO, Keck, MAVIS…)

Prototyping

- RTC prototype already built and running
- First tests on bench confirm high and stable performance in terms of timing
- Work in progress: interaction matrix + closing the loop this week

SOFTWARE DESIGN

OCTOPUS: SHM MANAGER

Common SW for all SHM

- Abstraction layers for shared memory handling
- CPU and GPU SHM implementations based on CACAO software
- FPGA support: work in progress at Swinburne University with Microgate support

Features

- Various type support: uint, int, float, double, half, etc…
- Semaphore based CPU synchronization
- Busy wait synchronization directly on GPU SHM

SOFTWARE DESIGN

MARLIN: CORE REAL-TIME

Abstractions

- Abstraction layers to define BU SHM links and execution
- Interfaces : *FPS*, *inputs, outputs, aux, indices*
- Container to handle multiple BU as a single process
- Execution: *update(), wait(), process(), notify()*
- Execution modes available through SHM control value : RUN, PAUSE, STEP, STOP

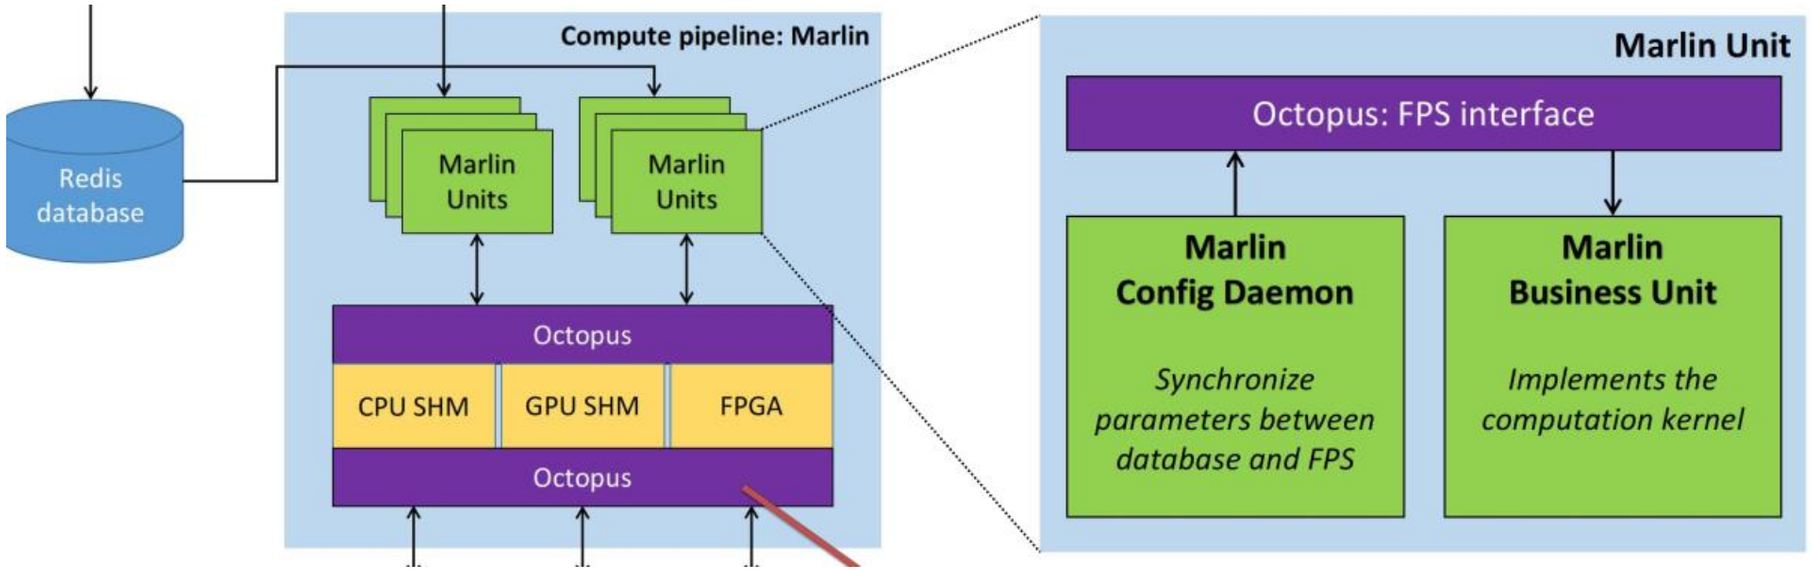

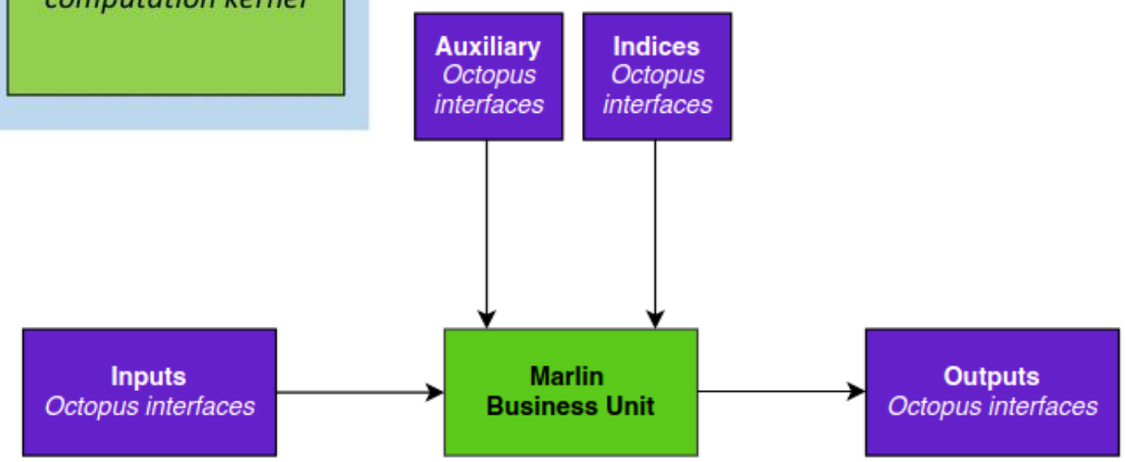

SOFTWARE DESIGN

KRAKEN: CONTROL & MONITORING

Configuration

- Parfiles to define business unit to instantiate
- SHM creation
- Redis database management

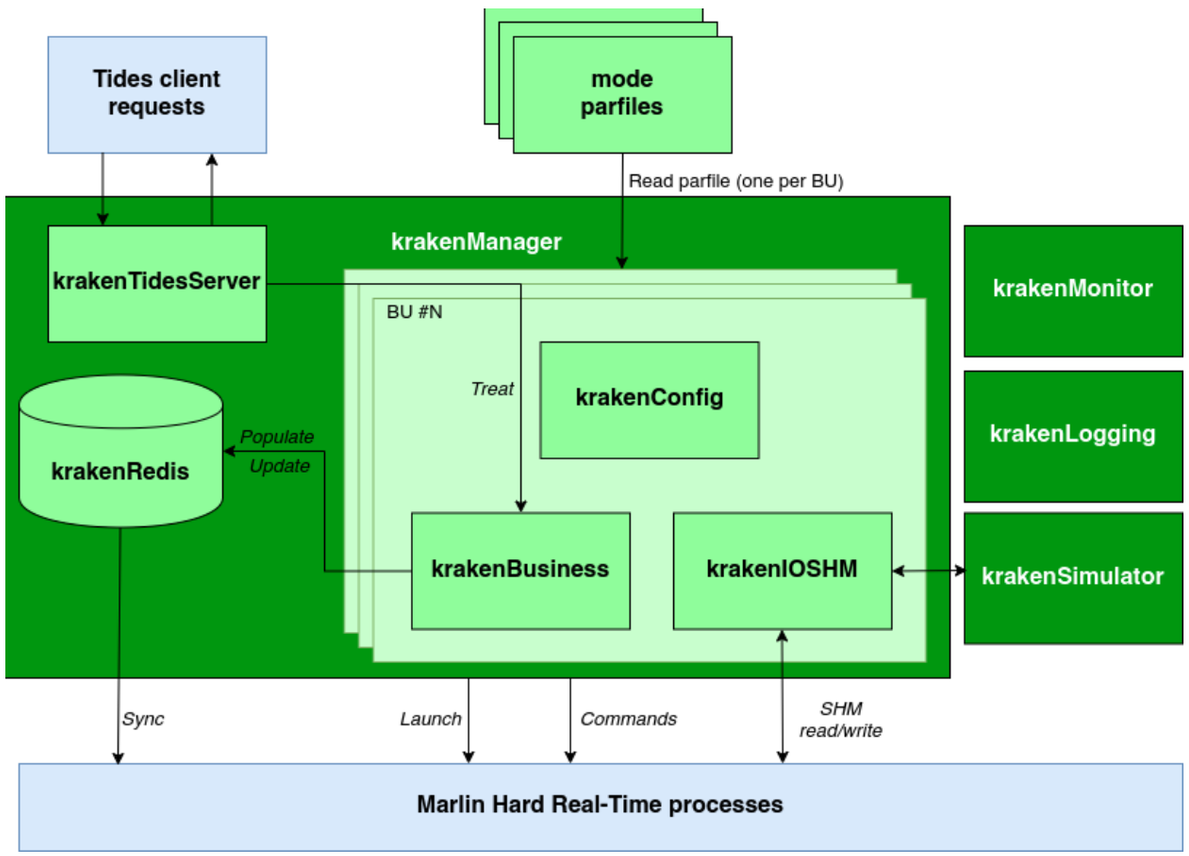

Control

- Launch and control runtime and update of each marlin process
- Tides server to handle remote client request

Monitoring

- GUI to monitor marlin processes health
- Compatible with CACAO processtool CLI: health, timings, fps, shm, etc…

SOFTWARE DESIGN

COMPASS: SIMULATOR

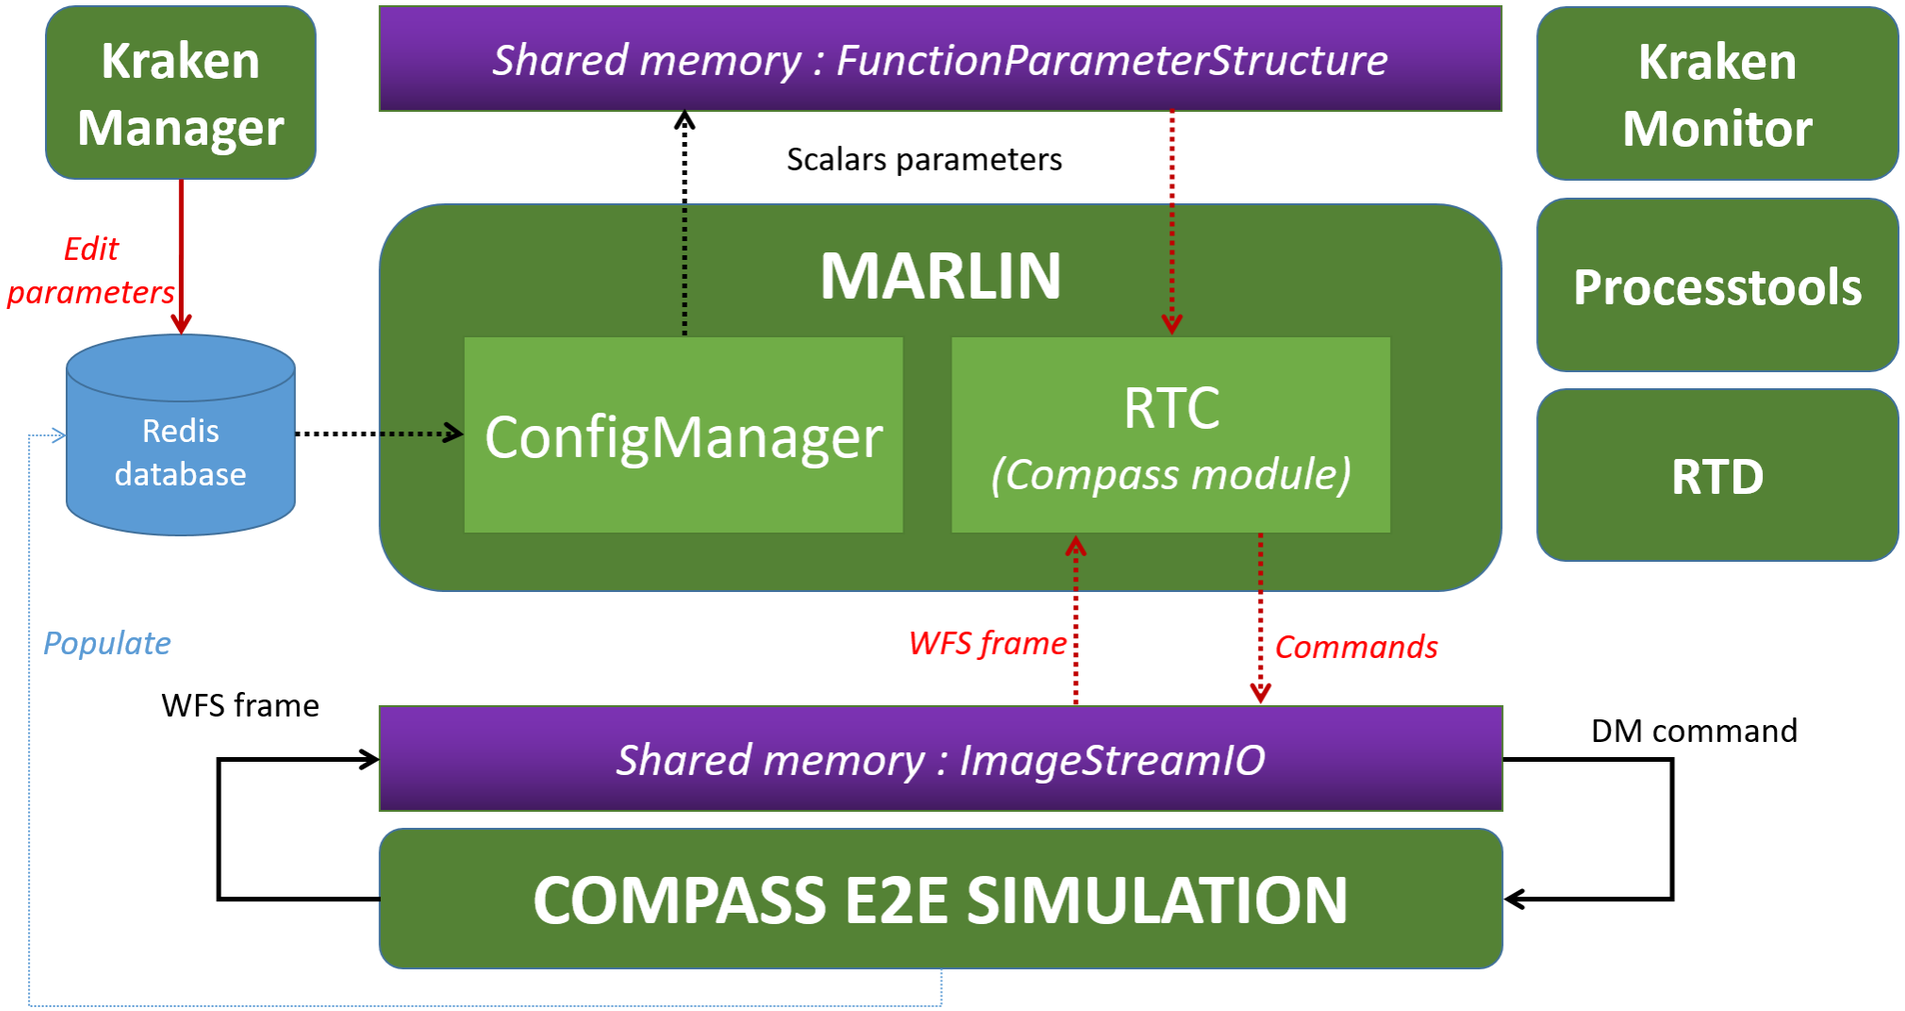

- Validation made with E2E simulation tool COMPASS
- DM outputs direct comparison

PERFORMANCE

MULTI-WFS PROFILE

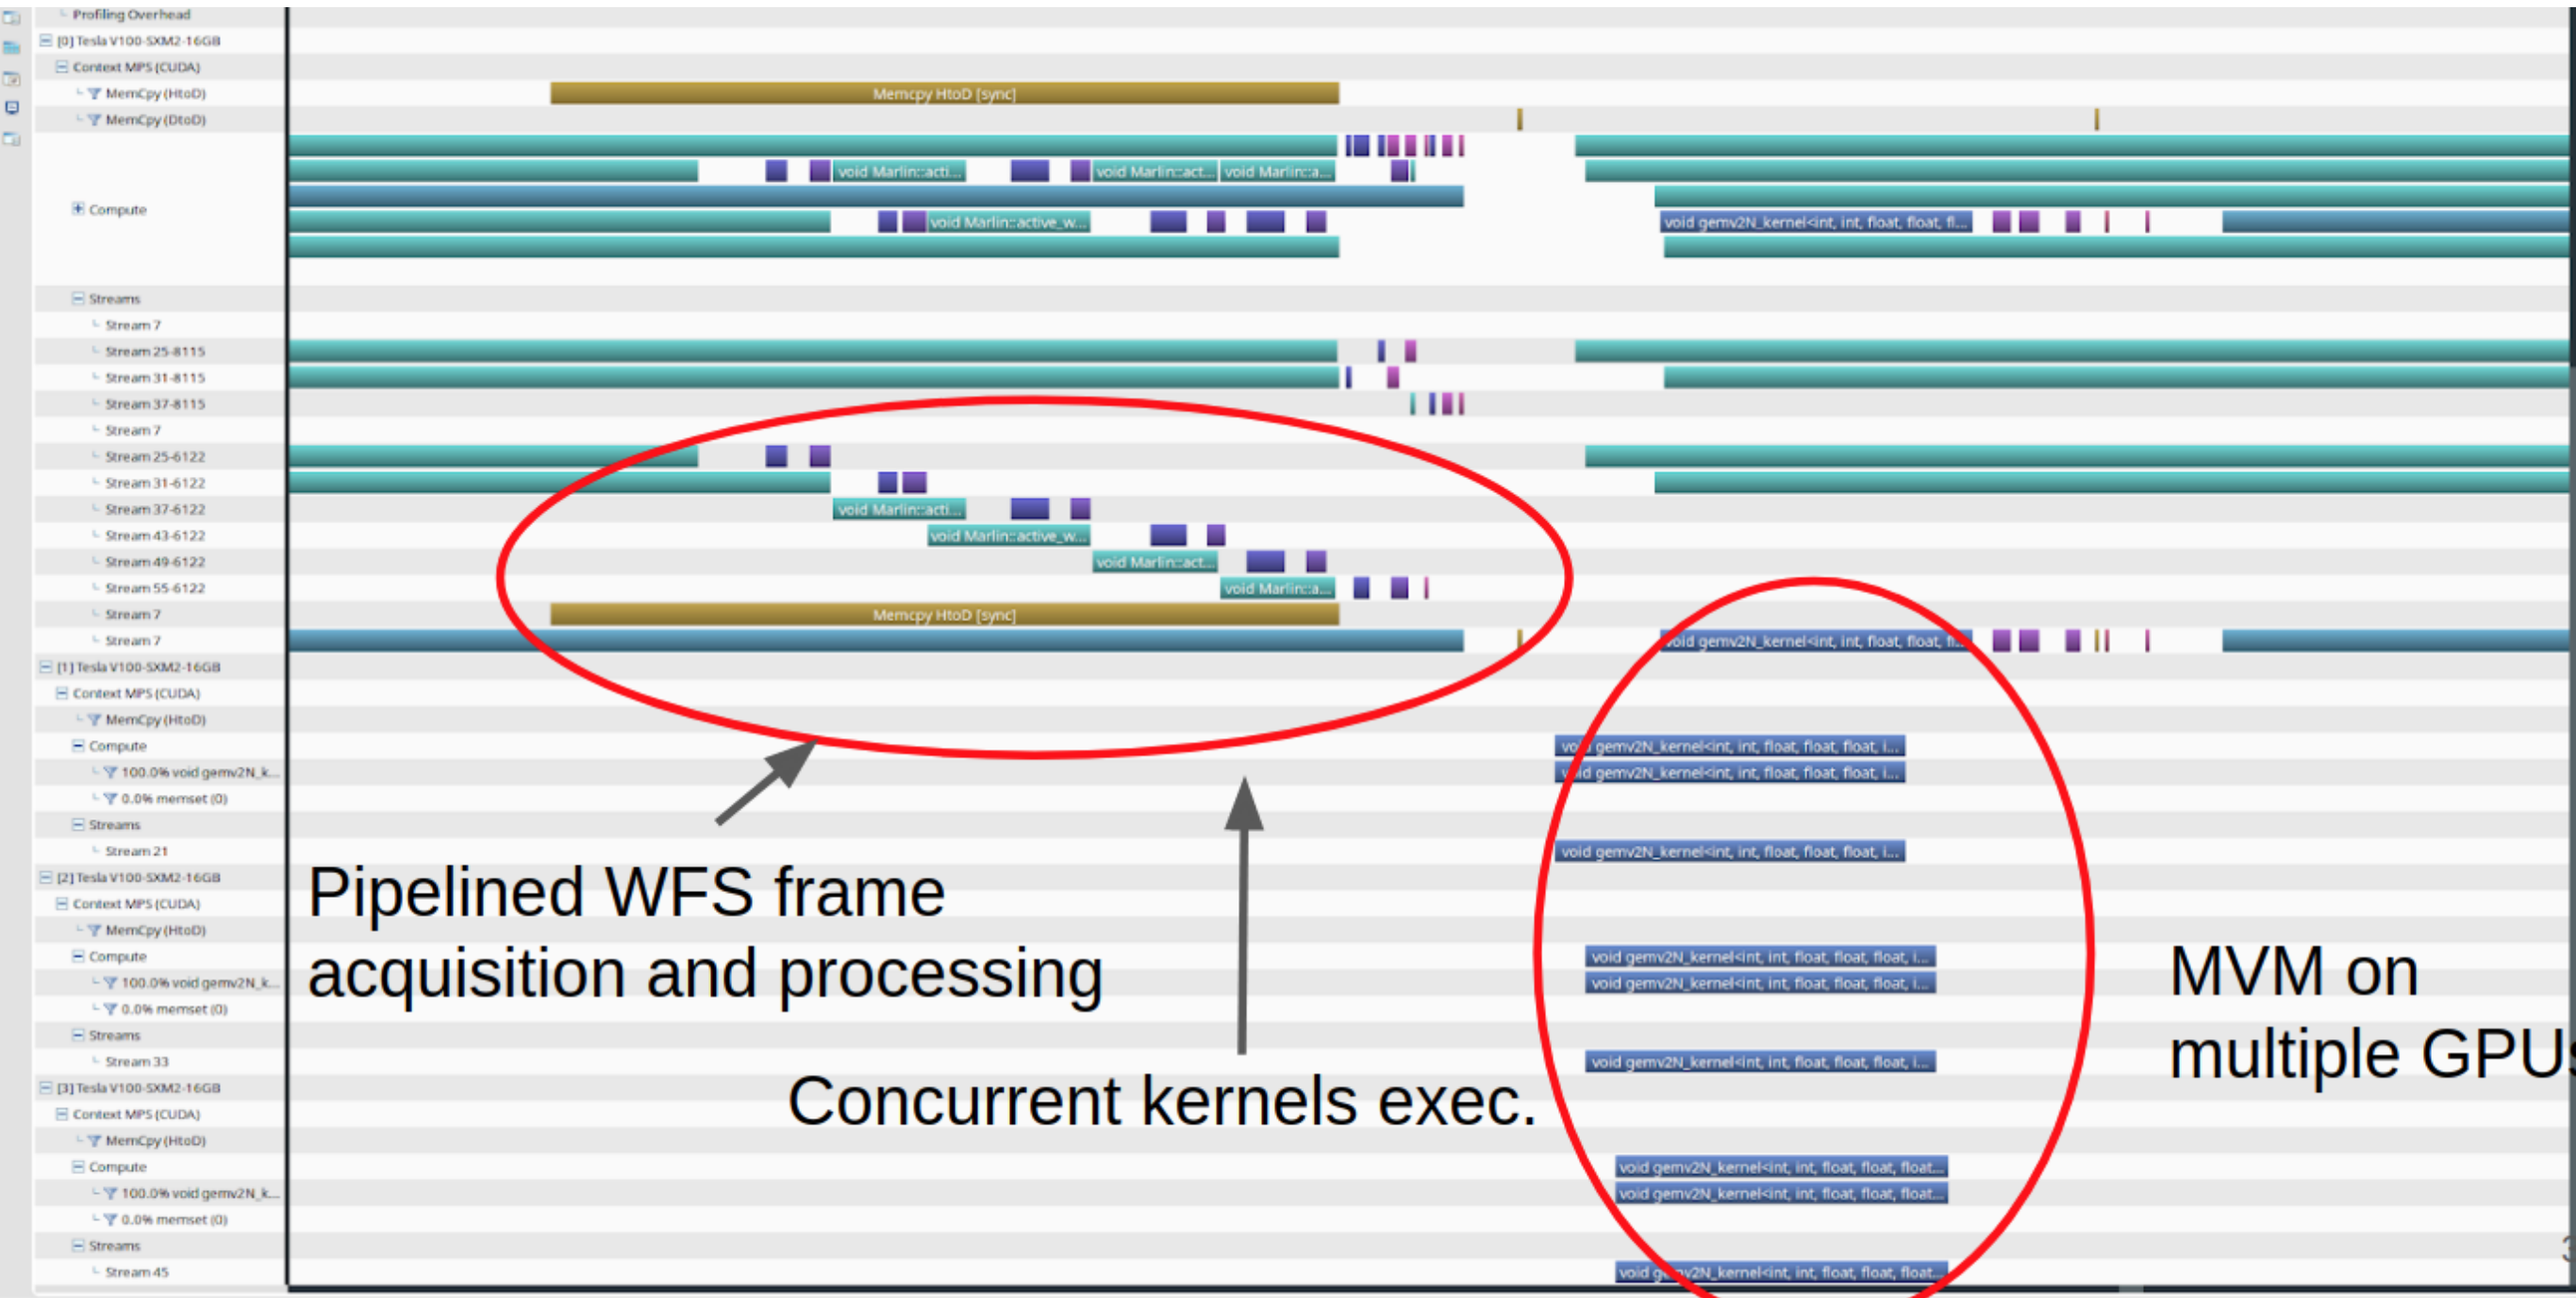

PROTOTYPING ACTIVITIES

FIRST TESTS: LATENCY