KAON 1427: Keck Adaptive Optics Facility: Real Time Controller Upgrade

Jason C. Y. Chin[1], S. Cetre[1], P. Wizinowich[1], S. Ragland[1], S. Lilley[1], E. Wetherell[1], A. Surendran[1], C. Correia[1], E. Marin[1], R. Biasi[2], C. Patauner[2], D. Pescoller[2], K. Glazebrook[3], A. Jameson[3], W. Gauvin[3], F. Rigaut[4], D. Gratadour[4], J. Bernard[4]

Organization: W. M. Keck Observatory[1], Microgate[2], Swinburne University of Technology[3], Australian National University[4]

ABSTRACT

The W. M. Keck Observatory Adaptive Optics (AO) facilities have been operating with a Field Programmable Gate Array (FPGA) based real time controller (RTC) since 2007. The RTC inputs data from various AO wavefront and tip/tilt sensors; and corrects image blurring from atmospheric turbulence via deformable and tip/tilt mirrors. Since its commissioning, the Keck I and Keck II RTCs have been upgraded to support new hardware such as pyramid wavefront and infrared tip-tilt sensors. However, they are reaching the limits of their capabilities in terms of processing bandwidth and the ability to interface with new hardware. Together with the Keck All-sky Precision Adaptive optics (KAPA) project, a higher performance and a more reliable RTC is needed to support next generation capabilities such as laser tomography and sensor fusion. This paper provides an overview of the new RTC system, developed with our contractor/collaborators (Microgate, Swinburne University of Technology and Australian National University), and the initial on-sky performance. The upgrade includes an Interface Module to interface with the wavefront sensors and controlled hardware, and a Graphical Processing Unit (GPU) based computational engine to meet the system's control requirements and to provide a flexible software architecture to allow future algorithms development and capabilities. The system saw first light in 2021 and is being commissioned in 2022 to support single conjugate laser guide star (LGS) AO, along with a more sensitive EMCCD camera. Initial results are provided to demonstrate single NGS & LGS performance, system reliability, and the planned upgrade for four LGS to support laser tomography.

Keywords: Adaptive Optics, Keck, laser, Laser Guide Star, laser tomography, Real Time Controller

1 INTRODUCTION

The Keck AO facilities (Keck I and II) have been operating with a Microgate SRL real time controller (RTC) since 2007 (Johansson et al. [1]). This system is based on a Field Programmable Gate Array (FPGA) architecture that processed incoming Shack-Hartmann wavefront sensor (SHWFS) and a natural guide star (NGS) tip-tilt (TT) sensor data. Corrections for image blurring from atmospheric turbulence are then sent to the deformable mirror and tip-tilt mirrors. Average telemetry data is sent to the AO system to offload focus and laser pointing. The RTC also supported both Keck I and Keck II's single laser guide star (LGS) AO systems by controlling the laser's uplink tip-tilt (Wizinowich et al. [2]). In 2016, a near IR tip-tilt sensor (NIRTT) was added to Keck I with modifications to the RTC (Femenia et al. [3]). In 2020, a SAPHIRA detector based IR pyramid wavefront sensor (PWS) and controller was added to Keck II (Bond et al. [4]). Although operational, these upgrades push the capability limits and reliability of the RTC. Together with the Keck All-sky Precision Adaptive optics (KAPA) project [Wizinowich et al. [5]), a higher performance and more reliable RTC is needed to meet next generation of AO capabilities such as laser tomography and sensor fusion. Microgate was selected among several vendors to design and build the new RTC to meet the demands for interfacing with new sensors and deformable mirrors, and the demanding tomographic processing for four LGS. The Microgate consortium includes their partners Swinburne University of Technology (SUT) and the Australian National University (ANU). Each of the entities bring their expertise to the team: Microgate for their hardware and controls expertise, and familiarity with the Keck system; SUT for their computational processing expertise; and ANU for their AO processing and software expertise. The project includes delivery of three systems for the observatory: Keck I, Keck II, and a spare. A major requirement for the design of the new RTC (NRTC) is to develop a system that is more flexible to changes, allowing future modifications while still meeting the computation performance in terms of speed, processing power, and deterministic latency. Previous architectures such as FPGAs limited the ability to make modifications without in depth knowledge and expertise. The NRTC processing design centers on the use of NVIDIA GPUs Tesla processors. The use of the CUDA library for the GPU allowed flexibility and optimization of the code to meet the latency requirement.

The NRTC project will also integrate a more sensitive SHWFS. The current CCD39 camera was limited by noise as the frequency increased above 1 kHz. The integration of the low noise OCAM2K EMCCD camera will allow the KAPA project to place four LGS onto a single camera to support laser tomography. The EMCCD camera was developed by First Light Imaging (Feautrier et al. [6]). The OCAM2K's speed and low read noise allows the AO system to be used for fainter NGS and LGS.

2 SYSTEM OVERVIEW

The NRTC is composed of three hardware subsystems: the Interface Module (IM) rack located near the AO optical bench hardware on the telescope, and the computational engine (CE) and telemetry server (TRS) located in the computer room off of the telescope (Biasi et al. [7]). Figure 1 shows the layout of the NRTC system on each telescope. The distance between the IM and CE is approximately 80 m, connected by OM4 fibers. These fibers are capable of 40 Gbs throughput although only 10 Gbs are needed for this project. The location of the IM is selected to be in proximity to the AO system to support the distance requirements to the AO hardware such as the deformable mirror (DM), wavefront sensor (WFS), and tip-tilt (TT) control stages. The TRS and the CE consume significant power and reside in a temperature controlled computer room off of the telescope. The primary interface to the NRTC is via a Linux server using the observatory's Experimental Physics and Industrial Control System (EPICS) architecture. The EPICS architecture is a distributed control system allowing controls and data across multiple subsystems using TCPIP networks. Users can send, retrieve, and log data across platforms without knowledge of the low level communication protocols of the other systems. The Linux server is also known as the AO supervisory controller (AOSC) which manages the non-real time controls and sequencing of the AO system.

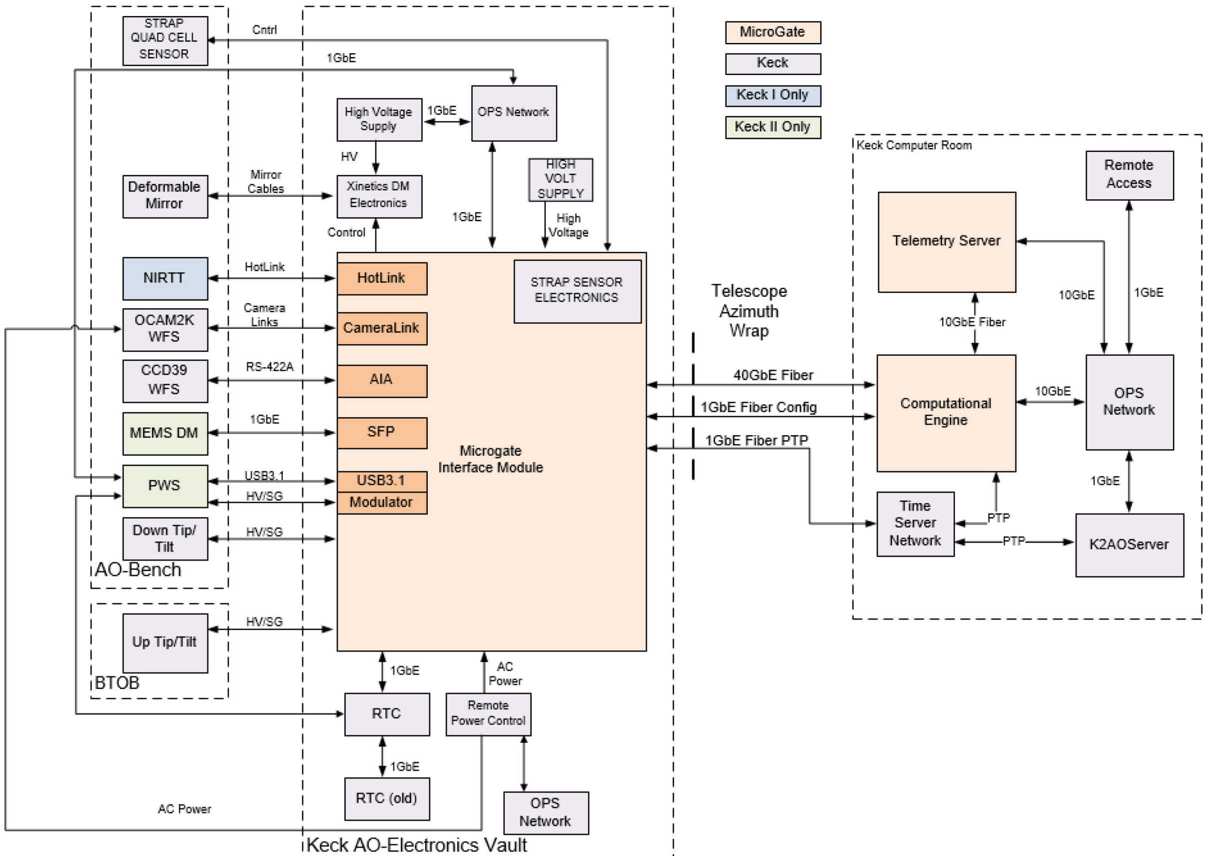

Figure 1: NRTC system architecture and interfaces. The AO system hardware is connected via the Interface Module.

To support logging and diagnostics, each IM downlink data packet from sensors is time tagged. Precision Time Protocol (PTP) format is used to support accuracy and synchronization of the subsystems. PTP is available from the observatory's GPS system and requires PTP enabled network switches to fully utilize its accuracy. The NRTC uses a pseudo version of PTP via a dedicated network specifically for time with minimal traffic. This allows micro-second accuracy as opposed to PTP's nano-second accuracy. All three of the NRTC subsystems, IM, CE, and TRS, along with the AO supervisory controller, use the PTP time from this network.

3 REQUIREMENTS AND DESIGN

The NRTC requirements focuses on its ability to operate across various AO modes. The modes are identified in Table 1 along with the associated hardware. Modes 1 through 4 are natural guide star modes. Modes 5 to 8 are laser guide star modes. Modes 9-11 are laser tomography modes (LTAO) using four laser guide stars. The LTAO modes use a single TOPTICA laser divided into four beams in a 13" square asterism. Mode 11 also uses an asterism but with a wider field (up to 5 arc minutes) requiring four separate OCAM2K cameras along with a higher order deformable mirror to correct the errors. In this mode, the AO bench will require future modifications to allow pickoff of four LGS in the field. In this mode, Microgate must only demonstrate the hardware is capable of interfacing with the additional hardware along with the processing power for the higher data rates and bandwidth performance.

Mode #	Name	SH WFS		PWS	Tip-tilt (+ focus)			DM		Tip-tilt	
		CCD39	OCAM2K		STRAP	NIRTT	PWS	Xinetics	MEMS	TTM	UTT
1	NGS w/ CCD39 WFS	1						1		1	
2	NGS w/ OCAM2K WFS		1					1		1	

3	NGS w/ PWS			1				1		1	
4	NGS w/ PWS & MEMS			1					1	1	
5	LGS w/ SciMeasure & STRAP+TRICK	1			1	1		1		1	1
6	LGS w/ OCAM2K & STRAP+TRICK		1		1	1		1		1	1
7	LGS w/ OCAM2K & PWS		1	1				1		1	1
8	LGS w/ OCAM2K & PWS TT		1				1	1		1	1
9	LTAO w/ STRAP+TRICK		1		1	1		1		1	4
10	LTAO w/ PWS TT		1				1	1		1	4
11	LTAO w/ STRAP+TRICK & new DM		4		1	1		new		1	4

Table 1: Real-time controller AO operational modes.

SH WFS = Shack-Hartmann Wavefront Sensor, PWS = Pyramid Wavefront Sensor, STRAP is a visible tip-tilt sensor, NIRTT = Near Infrared Tip-Tilt, DM = Deformable Mirror, TTM = Tip-Tilt Mirror, UTT = Up-link Tip-Tilt mirror.

In addition to operating with the AO modes, there are two major requirements critical to the NRTC's success. The design of the NRTC shall be flexible to allow for hardware and algorithm modifications. As new data is available such as atmospheric turbulence from MASS-DIMM instruments and Point Spread Function (PSF) reconstruction, the NRTC should be capable of combining these into its algorithms. The software architecture should allow for changes without significantly impacting the performance. The bandwidth requirement is expected to increase from the current CCD39 WFS operating at 1 kHz to the OCAM2K's performance of 2 kHz for non-binning and 3.7 kHz for binning modes. The bandwidth roundtrip processing requirement is set to 500 µsec with a goal of 250 µsec. The data flow from the IM to the CE and commands is shown in Figure 2. The round trip time starts with the sending of the last pixel of camera from the IM and receiving of the commands from the CE. This time does not include the readout time of the camera making the roundtrip time independent of the operating frequency. To meet this time, the data from the various sensors (SHWFS, STRAP, NIRTT, PWS) must be carefully packed by the IM and received by the CE without significant delays and overhead. The uplink commands from the CE are small and insignificant as compared to the downlink data. Microgate developed two custom interface boards between the IM and CE to ensure the data are sent efficiently and directly into the CE's GPU without overhead related to double buffering of data from the CE's CPU. In addition to the CE processing, the downstream and processed data are saved onto a telemetry server. The data is available in real time for processing by external systems such as the PSF server to optimize the AO performance by changing NRTC centroid gains or reconstruction parameters.

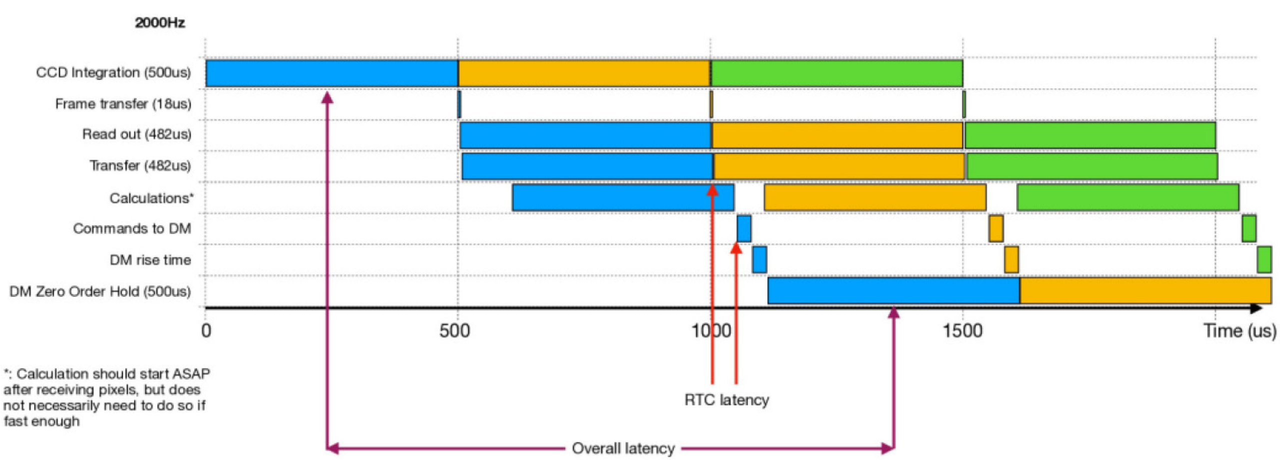

Figure 2: Keck real-time controller data flow timing diagram.

Figure 3 shows the control algorithm for LGS modes 5-8 for a single LGS system in the RTC. This algorithm is similar to the existing operational RTC developed by Microgate for use with the CCD39 SHWFS. Each block is a function within the CE. The color codes show the high (red) and low (green) throughput processing and I/O required by the system. Functions impacted by high and low latencies are highlighted as well.

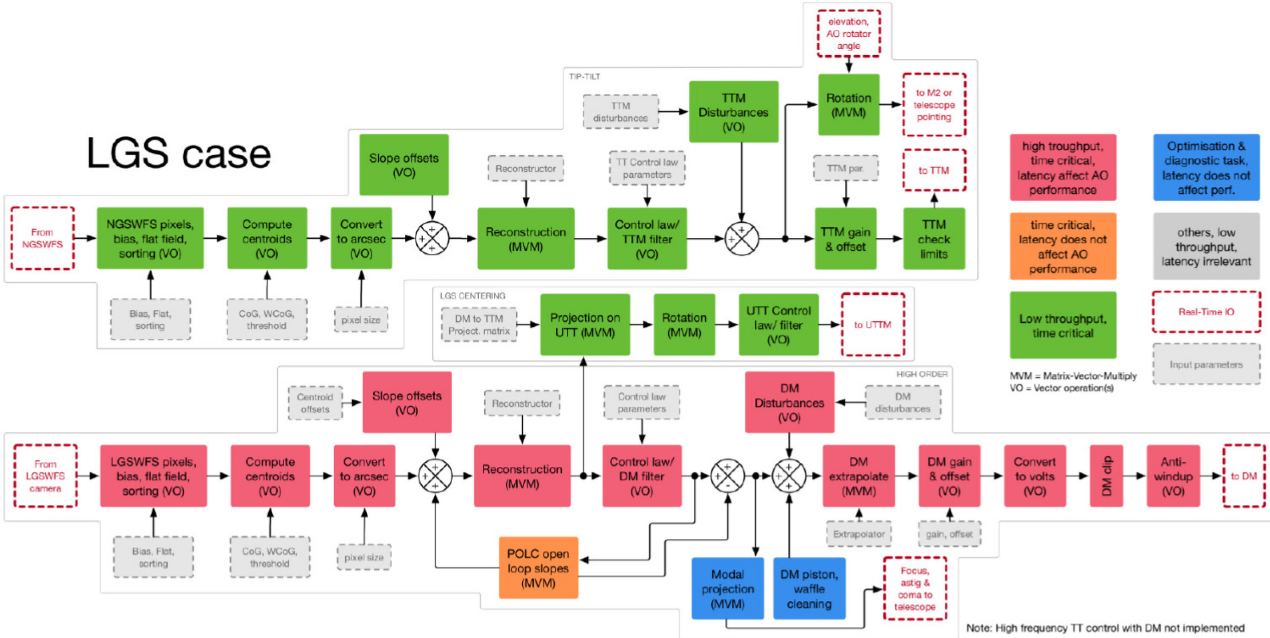

Figure 3: RTC LGS modes 5-8.

For laser tomography modes 9-11 using four LGS, the control algorithm is shown in Figure 4. The differences for the LTAO modes are the four ROIs (regions of interests) instead of a single LGS, the additional processing related to the ROIs, and the pseudo open loop control processing using the DM control feedback. Initially, the NRTC was required to demonstrate its bandwidth capability to support the LTAO modes using one OCAM2K. A new contract has been awarded to Microgate as part of the KAPA project to integrate the infrastructure to support four LGS.

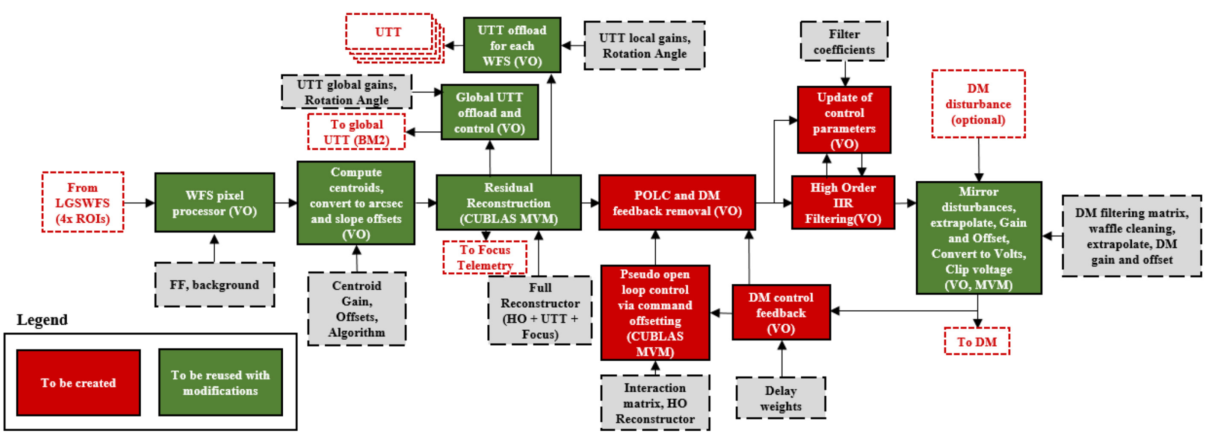

Figure 4: LTAO pseudo open loop control algorithm.

3.1 Interface Module Rack

The IM provides the electronics to communicate with the real-time AO system hardware. Figure 1 shows the hardware interfaces to the various hardware components. The IM rack itself is shown in Figure 5 (left). The IM is divided into two subsystems, the interface module crate (left), and the high voltage control boards (right) for the piezo devices (downlink tip-tilt, uplink tip-tilt, and pyramid sensor modulation). The system can support 15 separate axes of piezo controls and strain gauge sensing. A rendering is shown in Figure 5 (right) of the internal boards for the IM crate which uses a PCIe 8 lanes Gen 3 backplane. By using a PCIe backplane, adding interfaces by extending the PCIe will be easier in the future.

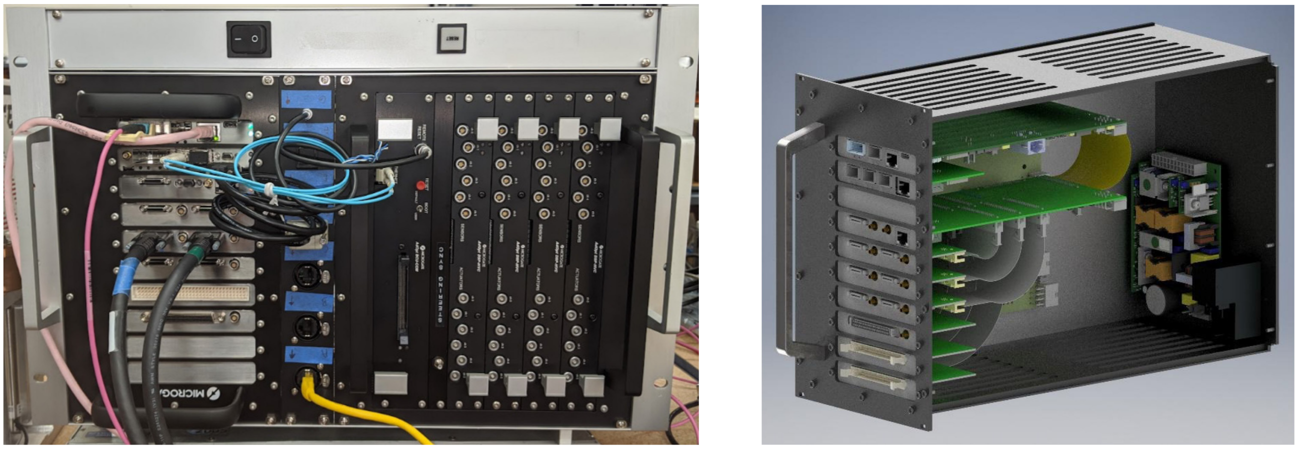

Figure 5: Rack (left) including the interface module and high voltage control boards; and rendering of boards (right) internal to the interface module.

Aside from individual connectors and interfaces to the AO hardware, there are two custom built FPGA boards for data input and output. The IM rack's architecture is shown in Figure 6. The interface board to the CE is known as the µXlink board. This ARM based microprocessor + FPGA based board features standard communications interfaces such as RJ45, USB2.0, USB3, QSFP, SFP+, and a micro D-Sub9 to interface with Keck hardware and the CE. The QSFP interface provides the 40Gb link to the CE to meet future data throughput demands. The RJ45 connections provide standard TCPIP network communications to the Keck supervisory controller and the high voltage control system crate. The D-Sub9 provides the RS485 serial interface to the existing NGS tip/tilt quad cell controller STRAP at over 1000 kHz data rates. The remaining connections support the Boston MEMS DM, PWS, and NIRTT TRICK sensor. Temperature sensors on the board allows the unit to monitor IM crate internal temperatures and issue alarms and shut down the unit if necessary. Status and faults are reported back to the CE and Keck monitoring system. The µXlink also has a FPGA Mezzanine Card Interface to a second FPGA based board controller µXFMC via a ribbon cable.

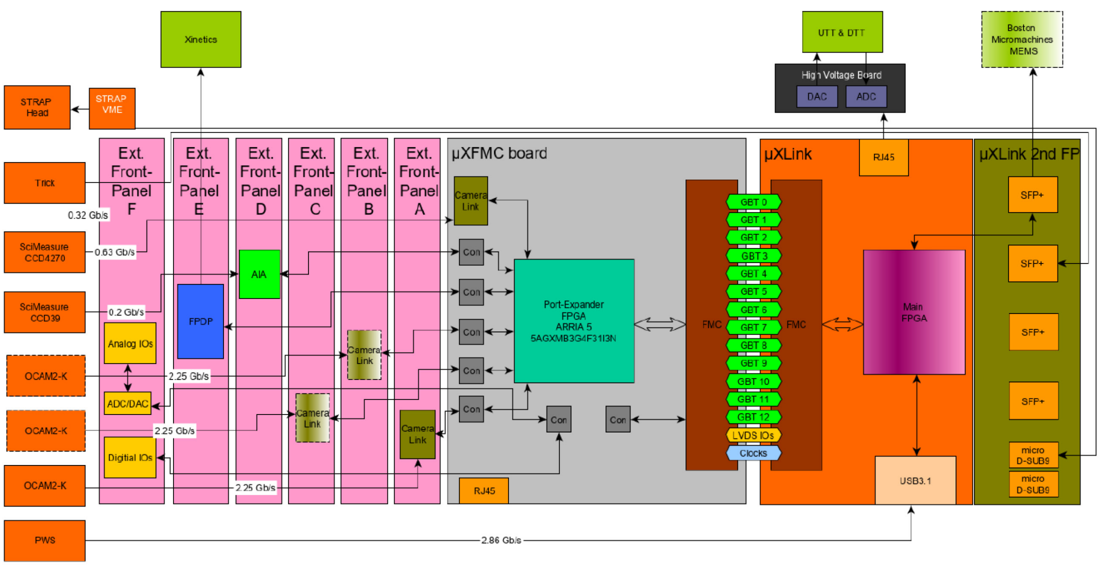

Figure 6: Interface module crate architecture.

Unlike the µXlink board, the µXFMC FPGA board primary function is to provide interfaces to hardware specific communications such as Cameralink. These includes the interfaces to the OCAM2K cameras (up to four in the future), AIA interface to the existing CCD39 SHWFS camera, and outputs to the Xinetics 349-actuator DM. The µXFMC also provides programmable 16 digital and 8 analog I/Os for interfacing with hardware or to provide diagnostic timing. Data between the µXFMC and the CE are sent through the µXlink. The µXFMC must follow strict real time constraints to ensure no data is lost. Firmware modifications to either board can be made or they can be debugged via their JTAG connectors.

3.2 Computation Engine

The CE is an off-the-shelf multi-core Linux server with high performance NVIDIA Tesla V100 GPUs. GPUs that are ideal for the AO algorithm matrix-vector-multiply (MVM) operations. Only one of the two GPUs is needed to operate the NRTC. The second unit is used for non-real time tasks or as a spare. The Supermicro 19-inch rack mounted server has dual Intel Xeon CPUs with eight cores each. To address the challenging environment at 4,000 m, the server is designed with redundancy including a spare GPU, 2.2K power supplies, and 1 TB SSD storage. The CE also includes a µXlink board to support the delivery of high-speed data to/from the IM.

3.2.1 Hardware Architecture

The key to the CE's bandwidth is its ability to parse processing and functionalities between its CPUs and GPU. This is done by a combination of hardware and software. The CE's hardware architecture is provided in Figure 7 along with the theoretical bandwidths between its components. This architecture includes Non-Uniform Memory Access (NUMA) allowing software functions to be isolated on separate CPUs, memory controllers, and PCIe buses. Critical processing is not interrupted by non-scheduled events. One CPU supports the processing of data while the other provides the critical logging and storage of data into the telemetry server. Another key CE design feature is the use of NVidia's GPUDirect Storage functionality. Data from the IM to CE goes directly into the GPU's memory with the use of a µXlink board in the CE. GPUDirect avoids the extra bounce or CPU interactions often required to move data between remote devices and local memory.

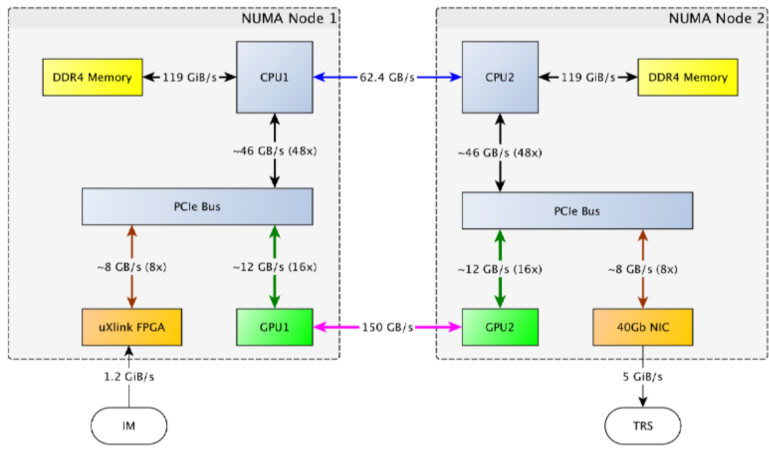

Figure 7: Compute engine hardware architecture.

3.2.2 Software Architecture

The CE's software architecture is shown in Figure 8. The main processing is divided into the Hard Real Time Controls (HRTC), Soft Real Time Controls (SRTC), COMPASS simulator, and the Main Control Module (MCM). The HRTC receives input frames from the IM sensors and performs the GPU processing and sends the resulting commands back to the IM. The SRTC supports the controls and sequencing of the processing, along with sending of raw and processed data to the telemetry server. MCM hosts the interfaces between the Keck supervisory controller and the NRTC. Commands and controls such as sending of the reconstructor, and monitoring of the IM, CE, and the TRS are done through the MCM. The MCM is simplified by using standard interface libraries and protocols which are compiled by both sender and receiver. The asynchronous messaging library ZeroMQ is used to transfer data between the NRTC and the Keck controller. Data is also serialized using the Protocol Buffers (Protobuf) to efficiently transfer data of various sizes and structures. Protobuf is widely used throughout Google to exchange data of various sizes and structures. To support testing and health checks, the CE

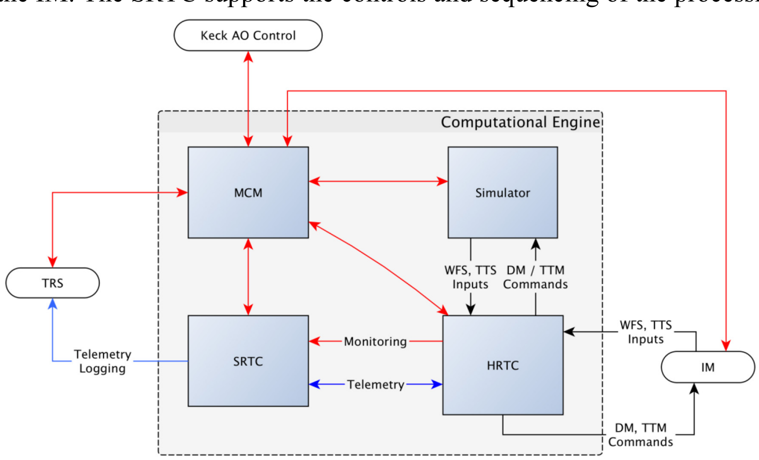

Figure 8: CE's software architecture.

provides an additional simulation mode to test algorithms within the CE or from the IM. Simulation data allows self-testing of the system to verify the health and functionality and is used as part of the Factory Acceptance Testing (FAT) process.

3.2.3 HRTC Software Design and Control

To enable the CE to operate with low jitter and high bandwidth for its HRTC, the design (Ferreira et al. [8]) (Figure 9) and decomposition of the functions must be carefully controlled to maximize data flow and processing efficiency. Tuned pipelines are highly efficient and can be difficult to maintain or modify. Pipelines allowing flexibility and ease of modification are often less efficient and negatively impacted by jitter. The Kraken suite of software is used to provide a flexible implementation of the HRTC while maintaining its bandwidth and jitter requirements. Kraken is a Python based control environment to sequence, manage, and monitor threads for the pipeline. Together with Kraken, two C++ libraries are used to support the real time kernels execution (Marlin) and data interfaces (Octopus). The Marlin abstraction layers provide generic containers for which the real time algorithms are executed. By dividing the processes into individual Marlin business units (BUs), they can be run in parallel as long as they are not highly dependent on information from each other. Each parallel BU can run independently as frame data is provided by Octopus. Computation among the BUs can overlap as partial data arrives from the IM, thus reducing the overall latency. The use of the Octopus abstraction layer controls the shared memory between the CPU, GPU, and µXlink to ensure the efficient transfer of data to the BUs.

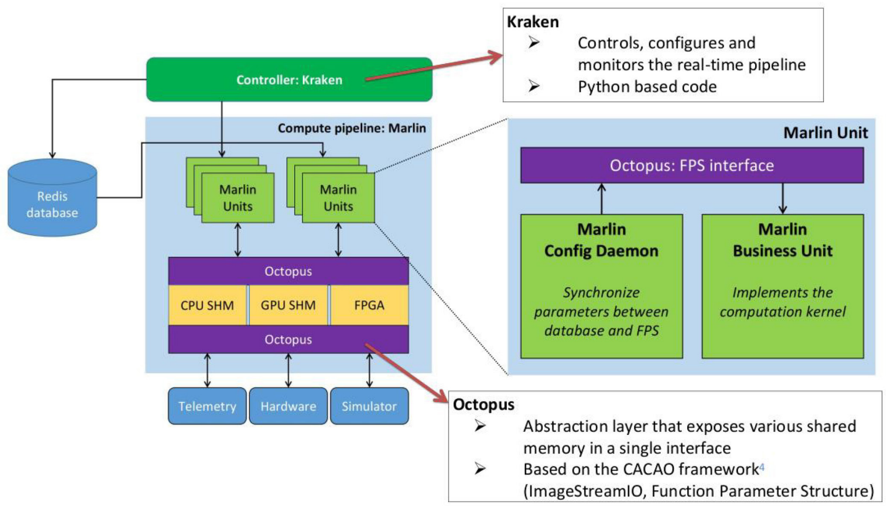

Figure 9: HRTC software design.

3.3 Telemetry Server and Data Management

The telemetry server is a Supermicro Intel Xeon server with a storage system. The storage system is a Seagate Exos X2u12 model with 12 disks bays with each disk capable of storing 12 TB. Although the use of SSD is preferred, the size and cost of required storage resulted in the use of magnetic disks. Seagate's helium filled Exos X12 disks were selected due to its ability to operate 24/7 at a 4500 m environment. To reduce this overall risk, a test was conducted to operate this disk at the summit environment prior to its arrival. The test read and wrote from the disk 24/7 for three months without issue. The 12 drive 144TB capacity is operated in a RAID-6 format giving a usable capacity of 120TB or 12 nights of data storage with an overhead factor of 1.3.

Similar to the MCM, the TRS uses the Protobuf scheme to serialize data and de-serial data from the CE. Post de-serialization, the data is sent to a PostgreSQL server for storage. The PostgreSQL storage format is the same as the original Keck RTC, allowing the Legacy TRS access tools to obtain data from the new TRS. To maximize the use of the TRS, it is expected to respond to multiple clients (up to 10) for data request and to respond within 10 ms with data.

4 FACTORY ACCEPTANCE TESTING AND PERFORMANCE

Due to COVID-19 travel limitations, the project had to find a different route for Factory Acceptance Testing than the traditional route of performing the FAT in person and in different continents. The unknowns of travel restrictions in early 2020 made the schedule and planning difficult. The project decided on a prototype approach to delivery of the initial system. Since a major risk of the project is the ability of the IM to interface with the various hardware, the initial system was delivered for hardware interface verification without the system being fully capable to process data. The hardware interfaces included those within the NRTC among the IM, CE, and TRS. For the success criteria, the focus was on the hardware readiness to communicate with Keck hardware, as well as testing to ensure the hardware is able to operate in the Keck environment (4,000 m). Any subsequent changes needed should be able to be programmed or modified remotely via the network or JTAG interfaces on the µXlink and µXFMC boards. This strategy met with some success in terms of no hardware was sent back to MIC or SUT for changes. The IM was able to communicate with the hardware as cameras and devices were made available to MIC for factory testing. However, significant challenges remained when trying to communicate with devices and controls at the same time. The µXlink and µXFMC ran into difficulties to pack the data orderly without missing any packets and the CE experienced difficulty parsing the data and rebuilding the images, especially for the OCAM2K since data is output from 8 amplifiers simultaneously. Once all the individual devices were successfully demonstrated, the focus shifted to the CE processing data.

To support an automated FAT, a combination of the simulated and preset data is provided to the IM. Python scripts modify user command parameters to generate pipeline outputs in .hdf5 format. These simulated commands include modifying DM, tip-tilt, centroids, and intensities to determine if the system responds correctly. Summary log files and data plots are then generated based on the comparison of the hdf5 outputs with expected results. Output logs are also commented with the NRTC requirements' reference number for compliance verification. The tests allowed demonstration of system compliance and comparison to previous RTC and SHWFS camera data.

4.1 System Bandwidth

The primary performance metric for the RTC is the bandwidth of the processing of the camera image and resultant commands to the DM and tip-tilt stages. The requirement is based on the output of the last pixel from the camera for each frame to the calculated DM and TT commands returning to the IM. The requirement is independent of the frame size and rate since it's based on the final pixel of each frame. The requirement is set at 500 µs with a goal of 250 µs. Because the raw data and frame information are timed tagged, each frame interval is logged along with its jitter. The performance results (Biasi et al. [9]) are presented in Figures 10, Figure 11, Figure 12 for the CCD39 (Modes 1&5), and Figure 13 and 14 for the OCAM2K (Modes 2&6). The red line sets the 500 µs requirement while the green line specifies the 250 µs goal. The blue line represents the output of the last pixel from the camera. Multiple frame rates are used to ensure the processing is independent of frequency. Table 2 summarizes the overall performance for various operating frequencies for both cameras. In these modes, the system complies with the round trip performance goal of 250 µs.

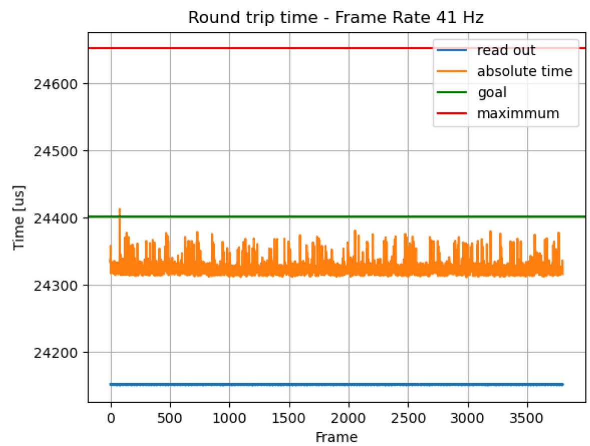

Figure 10: CCD39 performance at 41 Hz.

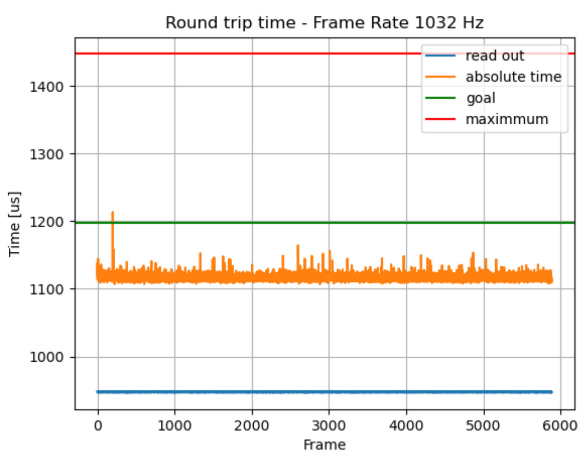

Figure 11: CCD39 performance at 1032 Hz.

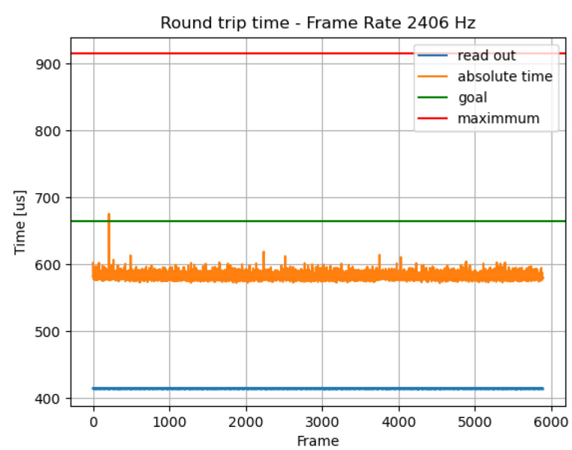

Figure 12: CCD39 performance at 2406 Hz.

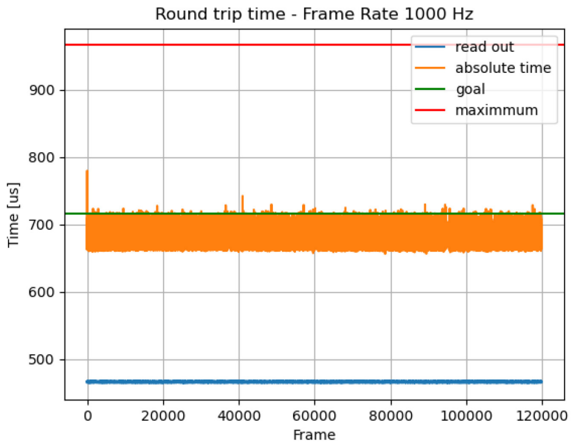

Figure 13: OCAM2K performance at 1000 Hz.

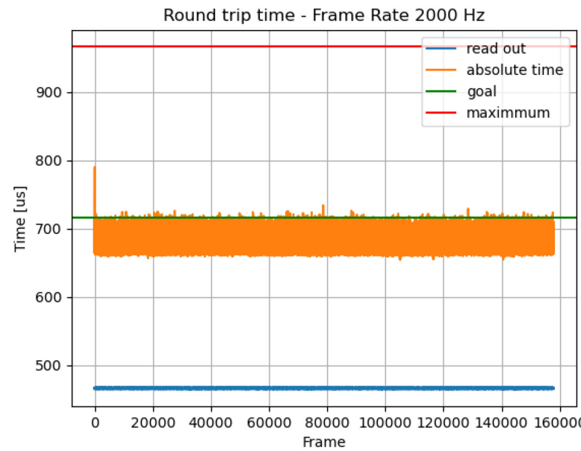

Figure 14: OCAM2K performance at 2000 Hz.

Sensor	Frequency (Hz)				Roundtrip Time (µsec)				Readout time (µsec)			
	min	mean	max	jitter	min	mean	max	jitter	min	mean	max	jitter
CCD39	41.35	41.36	41.36	0	159.26	**173.09**	261.07	9.52	24150	24152	24152	0.44
	1031.3	1032.08	1033.33	0.49	159.03	**169.38**	265.84	5.35	946.52	947.65	948.19	0.45
	2402.23	2405.54	2411.9	2.73	157.36	**168.02**	261.78	5.27	413.41	414.32	415.32	0.54
OCAM2K	49.99	49.99	50	0	185.73	**206.51**	315.9	11.5	465.39	466.31	467.53	0.53
	998.4	999.91	1001.26	0.45	189.3	**205.39**	313.76	6.97	464.91	466.31	467.77	0.43
	1994.43	1999.82	2005.88	1.72	188.11	**205.55**	324.01	6.7	464.91	466.31	467.77	0.43

Table 2: New real-time controller system bandwidth performance.

To reach the bandwidth and low jitter performance, multiple optimization strategies were needed. In addition to the use of the µXlink's GPUDirect capability, BUs were synchronized and optimized to minimize interactions among them. CUDA Multi-Process Service (MPS) is used to manage the processes and kernels within the GPUs. Polling is used instead of semaphore for process synchronization. Another key element is to use persistence kernels in the pipeline reducing the overhead, starting and restarting them when needed.

4.2 Auxiliary Performance

Other performance metrics include the response time for the CE's MCM interface and the TRS. In the case of the MCM, the interface is part of the SRTC that does not required fast response. The response time is demonstrated in Figure 15 with a mean time of 0.39 ms. To fully demonstrate the TRS response, a ten-client test was conducted with 1000 requests from the clients. The mean response time of 19 ms is shown in Figure 16. A bimodal mode was seen during this test; but the variances are at the 2 ms level. Both the MCM and TRS response time requirements are at 100 ms.

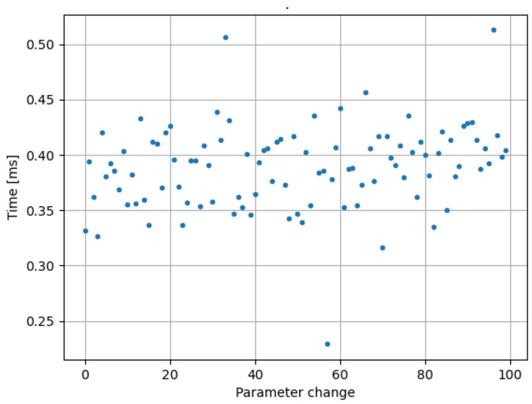

Figure 15: MCM response time.

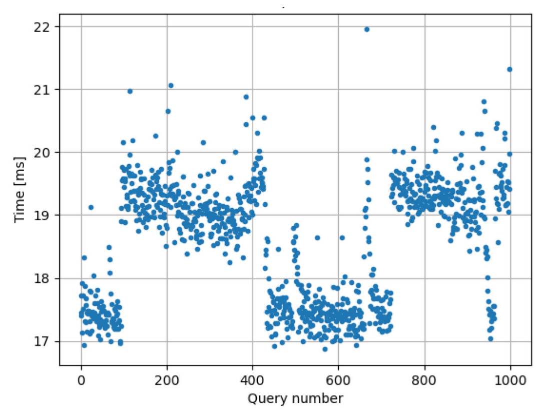

Figure 16: TRS response time from 10 clients.

5 SYSTEM INTEGRATION AND ON-SKY TESTING

The RTC hardware was installed on the Keck I AO system in 2021 and has been undergoing daytime and on-sky testing. With the system fully installed on the Keck I telescope and computer rooms, the NRTC has not experienced any environmental issues such as overheating of the GPUs or failure of the storage disk system. The system was integrated to allow external access for modifications of the software and firmware. Modifications to the µXlinks and µXFMC were required to synchronize the sensor data packets with their proper timing on the IM side and changes were made on the CE side to ensure data was unpacked correctly.

Unrelated to the NRTC, integration of the OCAM2K presented some challenges possibly due to over-illumination of the detector. The issues are shown in Figure 17. The EMCCD gain value was changed from a value of 1 at the left to 600 at the right. A flat field was taken and normalized to the top left quadrant to determine the sensitivity variances across the detector. The gains varied up to a factor of four at the highest gain of 600. Unresponsive or less sensitive columns were also seen between columns For the KAPA system, four LGS will be placed on the detector with each LGS on two quadrants, so sensitivity differences should be minimized.

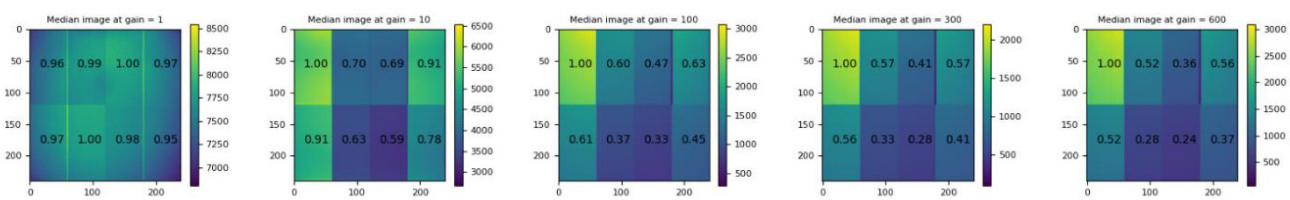

Figure 17: Normalized response for the four OCAM2K readout channels at EMCCD gains of 1, 10, 100, 300 and 600 (from left to right).

Based on discussions with First Light Imaging (FLI), these effects are related to over-illumination of the detector when it experiences higher than expected illumination or as the gain is changed with higher than expected intensities. The system camera does have overprotection in software by sampling the pixels from each quadrant as they arrived over a certain threshold, number of pixels over a certain threshold, and the number of frames this occurs. When the interlock occurs, the EMCCD automatically is set to one to prevent damage. However, the trigger event itself can be detrimental if it occurs often enough. The camera was returned to FLI where the internal voltages were adjusted to reduce the variance to less than a factor of two, and to improve the charge mobility at the gain register. Since two amplifiers shared the same high voltage, further improvements are not possible for individual amplifiers. Future damage will be difficult to correct for as voltages are now at their limits. To further reduce the impact from over illumination, a stricter or more conservative set of

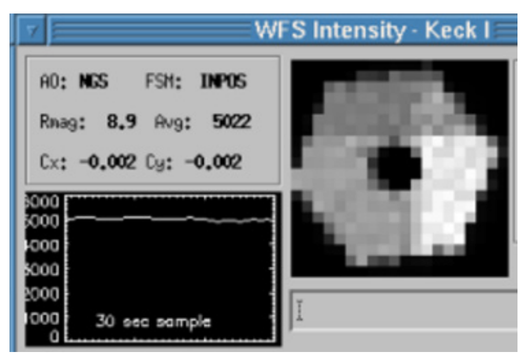

Figure 18: OCAM2K quadrant boundaries seen on the WFS intensity display.

interlocks have been considered to trigger the safety gain of one. Internal sequencing and checks of the camera response are also checked before increasing the EMCCD gain during operations. The boundaries between four readout ports can be seen in Figure 18.

On-sky integration started with modes 1 and 2 (see Table 1) using the CCD39 and OCAM2K SHWFS, respectively. The OCAM2K required modified camera fore-optics to match the 200 µm lenslet spacing to four pixels on the detector. The NRTC NGS performance was demonstrated on Jan 14, 2022. The NRTC was able to correct a 10th magnitude star (Figure 19) achieving a FWHM of 53 x 53 mas (left) and 52 x 54 mas (right). Observations were made at a zenith angle of 11°. The images were taken with the OSIRIS imager and the Br-gamma filter. The OCAM2K gain was 10 and a 200 Hz frame rate was used; the DIMM-measured seeing was 0.8". Figure 20 shows NGS AO corrected images at various guide star magnitudes. OCAM2K is in binning mode at a gain of 600 with seeing at 0.45 arcsec. The FWHM performance was 52x53mas with a Rmag of 8.9 operating at 3000 fps (left), 57x57mas with a Rmag of 12 operating at 150 fps (middle), and 68x68mas with a Rmag of 14.8 operating at 150fps (right).

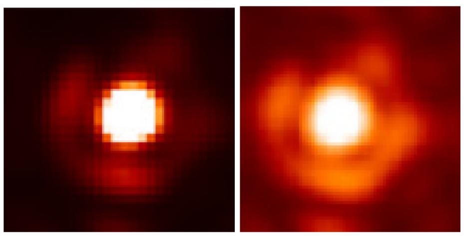

Figure 19: NGS first light AO corrected images.

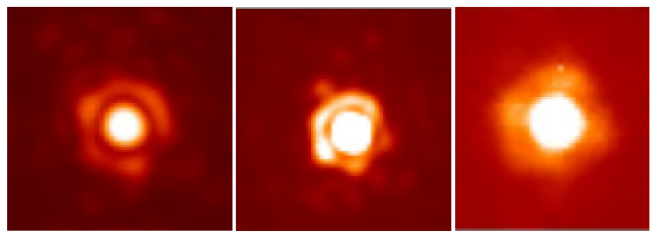

Figure 20: NGS AO corrected images at varying NGS Rmags.

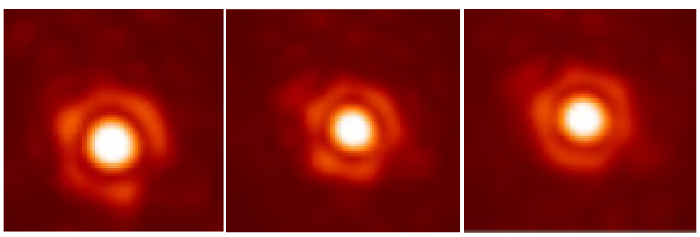

Figure 21: LGS AO corrected images.

LGS performance is demonstrated in Figure 21. The TOPTICA/MPBC operated at 22W with a Rmag=8.4. At this magnitude, the OCAM2K was able to operate at 3kHz with its maximum gain of 600. OSIRIS with the Br Gamma filter was used to capture the images. The FWHM performance was 53x53mas (left), 55x54mas (middle), and 56x53mas (right). With improved calibration of the AO system, we should be able to achieve the diffraction limited FWHM performance of 48mas.

In terms of hardware reliability, the NRTC has not had any spontaneous reboots except early in the project due to improper memory allocations in the CE. There are occasions of the pipelines hanging which we do not fully comprehend; but can easily be restarted once the problem is known. Many of the integration issues are related to the orientation or ±sign accounting of parameters, and proper packing of the data from multiple sensors in the IM and unpacking of the data in the CE. Challenges also included requirements of the TRS to be backwards compatible with database tools developed for the previous RTC. The updated organization of the data in the NRTC made it challenging to decompose the data with the legacy tools. A software middle layer was necessary to operate with the legacy tools.

6 CONCLUSION

A new real-time controller (NRTC) and OCAM2K camera are being integrated with the Keck I AO system in support of the KAPA project. A second, identical NRTC system and camera has been installed with the Keck II AO system and integration will begin as the Keck I system commissioning is completed. Unique aspects of the Keck II system include a SAPHIRA-detector based near-infrared pyramid wavefront sensor and a MEMS deformable mirror. The Microgate Consortium designed and built the NRTC to support existing and future AO capabilities for the Keck AO facilities. The hardware architecture with the use of GPUs allows for future changes and greater processing capacity while the software architecture provides the flexibility needed to test new algorithms and to integrate new sensors and controls. The system itself has greater reliability and has built-in tools to support failure isolation within the system. The NRTC has demonstrated its bandwidth performance in single LGS and NGS modes and is scheduled for an LTAO upgrade using four laser guide stars in 2023. This upgrade will demonstrate one of its main requirements in terms of being flexible to support new algorithms and hardware. Future capabilities planned for the NRTC include demonstrating PSF reconstruction, predictive control, sensor fusion, additional wavefront sensors and a higher order DM.

ACKNOWLEDGEMENTS

The W. M. Keck Observatory is operated as a scientific partnership among the California Institute of Technology, the University of California, and the National Aeronautics and Space Administration. The Observatory was made possible by the generous financial support of the W. M. Keck Foundation. The authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Maunakea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain. The Keck II real-time controller is funded by the NSF Major Research for Instrumentation Program award AST-1727071 (Wizinowich). Funding support for the KAPA system is provided by the National Science Foundation Mid-Scale Innovations Program award AST-1836016 (PI: Wizinowich). We would like to acknowledge the technical contributions of Randall Campbell, James E. Lyke, and Max Service. We also would like to acknowledge the technical team of Florian Ferreira and Arnaud Sevin from Observatoire de Paris. We would like to thank the NRTC science team of Andrea Ghez, Mike Liu and Claire Max for their support.

REFERENCES

[1] Johansson, E. et al., "Upgrading the Keck AO wavefront controllers," SPIE Proc. 7015-121 (2008).
[2] Wizinowich, P. et al., "The W. M. Keck Observatory Laser Guide Star Adaptive Optics System Overview", Publications of the Astronomical Society of the Pacific, 118:297-309 (2006).
[3] Femenia-Castell, B. et al., "Status and new developments with the Keck I near-infrared tip-tilt sensor," SPIE Proc. 9909 (2016).
[4] Bond, C. et al., "Adaptive optics with an infrared pyramid wavefront sensor at Keck," J. Astronomical Telescope, Instruments and Systems, Syst. 6, 039003 (2020).
[5] Wizinowich, P. et al., "Keck All Sky Precision Adaptive Optics Program Overview," SPIE Proc. 12185 (2022).
[6] Feautrier, P. et al., "Characterization of OCam and CCD220: the fastest and most sensitive camera to date for AO wavefront sensing," SPIE Proc. 7736 (2010).
[7] Biasi, R. et al., "MIC-KR-100-01_DER_v3_Keck_RTC_Design_Report," Keck Adaptive Optics Note 1247, (2019).
[8] Ferreira, F. et al., "Hard real-time core software of the AO RTC COSMIC platform: architecture and performance," SPIE Proc. 11448 (2020).
[9] Biasi, R. et al., "MIC-KR-100-12_TRP_v1d4 Real Time Controller Factory Acceptance Test Report," Keck Adaptive Optics Note 1420, (2022).