KAON 791

# Real-time control for Keck Observatory next-generation adaptive optics

Marc Reinig[*a], Donald Gavel[a], Ehsan Ardestani[b], Jose Renau[b]
[a]University of California Observatories, 1156 High Street, Santa Cruz, CA, USA 95064
[b]University of California Santa Cruz, School of Engineering, 1156 High Street, Santa Cruz, CA, USA 95064

## ABSTRACT

The next generation adaptive optics systems for large telescopes will be complex systems far larger, more complex, and with higher performance than any currently installed. This requires adopting new algorithms, technologies, and architectures. The Keck next generation adaptive optics (NGAO) system requires real-time wavefront reconstruction and tomography given input from 7 laser and 3 natural guide stars. Requirements include 2 KHz atmospheric sampling, tomographic atmosphere estimation, and control for 5 deformable mirrors. We take advantage of the algorithms' massive parallelism and realize it on a massive array of FPGAs, GPUs, and multi-core CPUs. This paper presents the current design and analysis of the NGAO system.

**Keywords:** Keck, NGAO, Adaptive Optics, tomography, bandwidth, wavefront sensor, guide star, MOAO

## 1. INTRODUCTION

The next generation adaptive optics systems for large telescopes will be complex systems far larger and more complex with higher performance requirements than any currently installed. This increased complexity, size and performance has led to the need to adopt new algorithms, new technologies, and new architectures to implement them. The Keck Next Generation Adaptive Optics (KNGAO) real-time control (RTC) system is actually 4 independent but coupled AO systems, each with significantly higher performance requirements than the current Keck II AO system. It requires real-time wavefront reconstruction and tomography given input from 4 laser guide stars and 3 natural tip/tilt stars. System bandwidth requirements demand that we sample the atmospheric turbulence 2,000 times a second, produce a tomographic estimate of the atmosphere from them, and use this estimate to control deformable mirrors to correct the effect of the turbulence. To implement the algorithm we take advantage of the massive parallelism inherent in the algorithm and realize it on a massive array of processors, including FPGAs, GPUs, and multi-core CPUs. In this paper, we present the design of the overall real-time control system and describe how it achieves the KNGAO needs.

The next generation AO systems that compute tomography must solve problems at least two orders of magnitude larger than single conjugate systems today and do it at over 1 KHz. Conventional vector matrix multiply approaches are untenable due to the size of the problem and the speed at which it must be calculated. New algorithms must be implemented. Conventional PC based hardware solutions are likewise insufficient to solve the problems and new hardware architectures must be used, typically using massive parallel architectures.

The Keck NGAO system has extreme processing and storage requirements: It must control 10 wavefront sensors (WFSs), 5 deformable mirrors (DMs), 11 tip-tilt (T/T) mirrors, at a 2 KHz frame rate with a computational latency of <600 µsec (from last pixel read from the WFS to last DM command issued). This requires a processing rate of over 700 G-operations per second. It also must be capable of saving telemetry data, streaming at over 3 GB per second from its various components, for extended periods, and storing >50TB per night.

The requirement of modern high-performance AO systems necessitates the use of modern, massively parallel technology. This includes the use of field programmable gate arrays (FPGAs), graphics processing units (GPUs) and multi-core CPUs.

[*] mreinig@ucolick.org; phone 1 (831) 459-4362

GPUs and FPGAs have been proposed[1] and implemented[2 3 4] as solutions for AO for several years. Research has extended from low cost low capability to large and complex high performance systems. The Keck NGAO RTC uses all these technologies to perform tomographic reconstruction and AO compensation of a 40” FoV.

This paper is organized as follows: section 2 presents the issues that led to the requirements of the RTC; section 3 presents the architecture of the RTC; section 4 discusses the decisions on technology that needed to be addressed; and sections 5 and 6 present results and conclusions respectively.

## 2. FUNCTIONS, GOALS, AND LIMITS LEADING TO THE REQUIREMENTS

The Keck NGAO RTC is the real-time control portion of a large complex real-time control system and has multiple modes of operation.

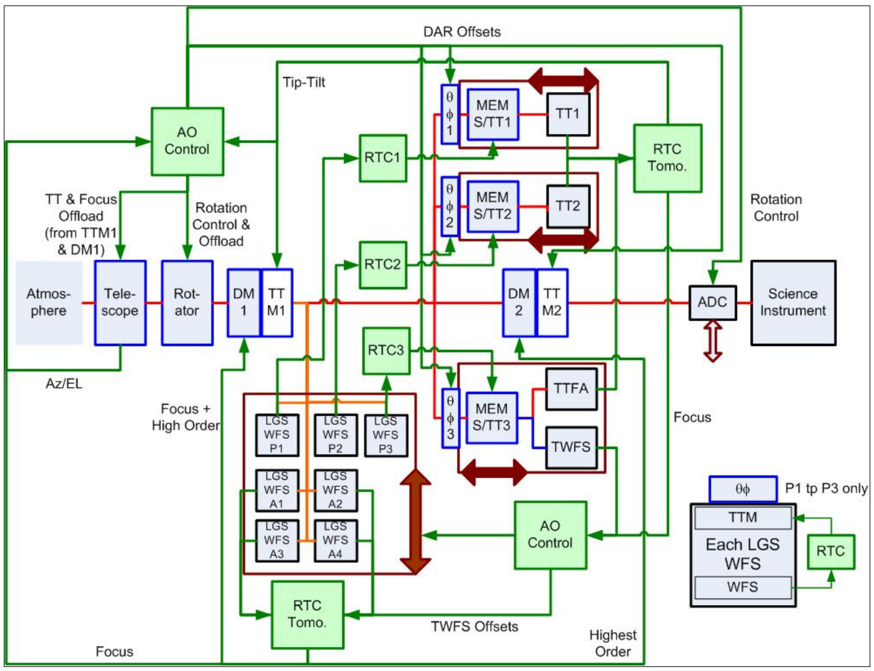


Figure 1 LGS AO control loops for the science instrument[5]

### 2.1 RTC requirements

The RTC’s requirements have been flowed down from a set of parameters determined by the primary science cases for the upgrade[6 7].

Cost is an overriding element in any development project due to budgetary limitations and the following RTC goals were filtered through the lens of budget constraints.

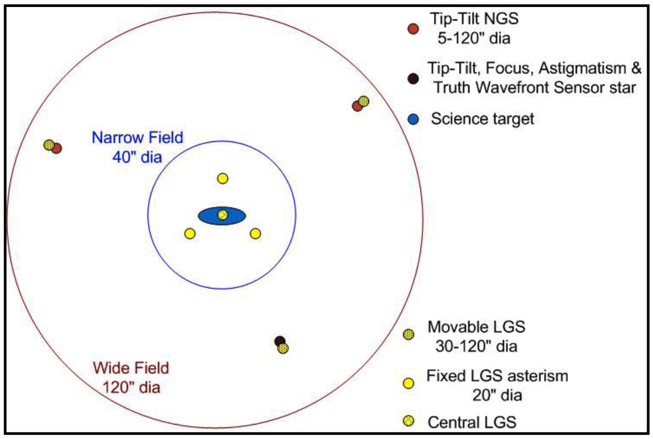


Figure 2 LGS "3+1" asterism for tomography of the central science field, plus three patrolling LGS for image sharpening two TT and one TTFA NGS[8]

**Increase sky coverage**: In order to increase sky coverage, the NGAO needs 3 T/T stars to give a good estimate of the T/T over the 40" science FoV. This required that we be able to sharpen dim IR T/T stars with AO to be able to use them to achieve the desired sky coverage. This in turn created the need for 3 complete and independent patrolling laser guide stars (LGSs) coupled with 3 independent high order (HO) AO systems, one to correct each T/T star. These systems operate open-loop. One of the T/T stars is selected to also provide focus and astigmatism information, which we off load to the non-real time AO Control system for processing.

**Improve Strehls**: the goal was to have Strehls of >80% in the K-band and AO correction in the red (Strehls of 15% - 25% at 750 nm). This led to the decision to use 60 sub apertures across the primary correcting the 40" FoV.

**Meet the temporal error budget**: computational latency and frame rate we set to be 600 µsec and 2 KHz respectively to meet this error budget. The SNR error budget and laser power limitations also played a factor in these values.

**Correct for an optical path delay (OPD) of >6µm:** typical seeing conditions at Keck exceed the stroke capabilities of our HO DMs. Consequently, we must use two DMs for correction: a high-stroke low-order (LO) DM with 400 actuators and high order DMs with 1020 or 4092 actuators depending on their use. The LO DM corrects light for all WFSs and the science object and operates closed loop. All HO DMs operate open-loop.

**HO correction of a FoV of 40"**: to correct over a FoV of 40", we must generate a tomographic estimate of the atmosphere in the field. To do this, 4 LGSs are needed. This 3D estimate of the atmosphere is then used to correct the aberrations in the science direction. The correction in the science direction is calculated by projecting a correcting wavefront along the path to the science object though the tomographic estimate of the atmosphere. The LO and HO modes are separated from this wavefront and the LO are used to control the LO DM. This DM is common to all WFSs and the science object, it is controlled closed loop. The HO modes are sent to the HO science DM correcting the science object only. This DM is controlled open-loop. The tomography engine also has the capability to provide correction in the direction of multiple objects off axis within the 40" tomographic volume above the telescope. However, in the current design of the opto-mechanical system, there are no deployable field units (DFUs) to take advantage of this capability.

**Operate HO DMs open-loop:** This requires that we must be able to place the DM in a desired shape with a single set of commands without the aid of a control loop to null any errors. Unfortunately, the response of the DMs is non-linear. This means we must characterize the DMs and use this knowledge in a DM command processor to transform the desired shape into a set of DM commands that will achieve that shape.

**Control T/T on 11 objects**: since all guide stars must be T/T stabilized, the RTC must control 10 T/T mirrors on the 10 guide stars, plus the T/T stage of the LO DM.

**Adapt to changing conditions and configurations**: the RTC must be able to adapt to changing conditions. These include compensating for different obscurations as the telescope rotates, $C_n^2$, layer heights, long term quasi static aberrations in the telescope or the atmosphere, etc.

**Collect telemetry and supply diagnostic information**: we need to collect diagnostics and detailed site atmospheric data for system analysis, diagnostics, and PSF reconstruction. This requires a very large, very high-speed disk sub-system.

## 3. ARCHITECTURE

Section 2 presented the performance goals of the system. This section presents the architecture to implement those goals.

### 3.1 Overall System Description

The RTC is a massively parallel pipe-lined system with a frame rate of 2 KHz and latency, from last pixel in to last DM command, of <600μsec. It operates as a slave to the non-real-time AO control system.

It is implemented with GPUs, FPGAs arranged as a systolic array, and multi-core CPUs. It is actually 4 independent but coupled AO systems: 3 that sharpen the tip/tilt/focus/astigmatism stars (TTFA) and one that computes a tomographic estimate of the atmosphere in the science field of view (FoV) above the telescope and corrects for it.

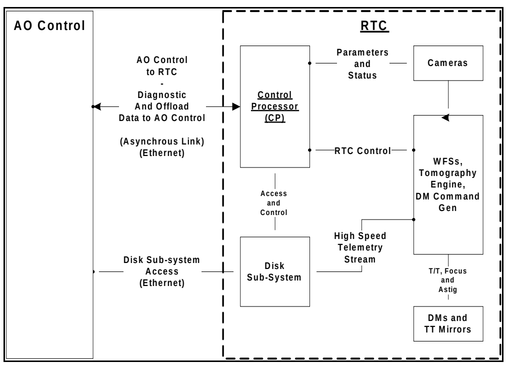


Figure 3 NGAO RTC architecture

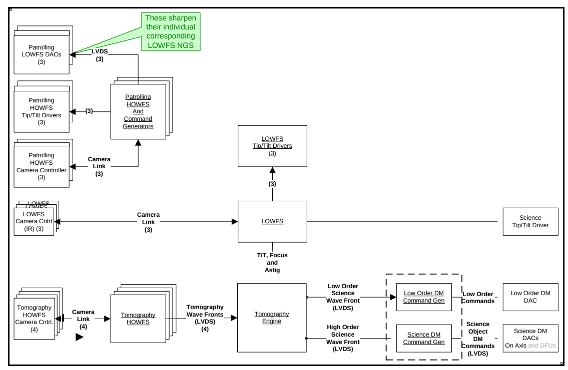


Figure 4 RTC data processing view

Data from ten guide stars (three NGS and seven LGSs) are processed to provide AO for the science object. Each LGS is associated with a separate HOWFS (high order wavefront sensors). Each of these WFSs has significantly higher performance requirements than the current single WFS for the Keck II. All 10 guide stars have individual T/T stabilization. Additionally, the T/T values of the 3 NGSs are calculated by 3 low order wavefront sensors (LOWFS) combined and applied to the overall T/T control through an 11$^{th}$ T/T stage.

Each of the three patrolling HOWFSs is associated with one of the three natural T/T stars. They operate as independent AO systems, sending commands to individual DMs, sharpening the NGS and enabling the use of dimmer stars. The four remaining tomography HOWFSs send their wavefronts to the Tomography Engine. The Tomography Engine reconstructs an estimate of the atmosphere in the 40” FoV. This estimate is used to correct the science object and to provide the LO wavefront information for the LO DM.

### 3.2 Control and telescope system interface

The RTC is a slave system to the non real-time AO Control. The AO Control provides for most human interface (other than low level system diagnostics) and communicates to the RTC through the RTC Control processor using a low level command structure over a TCP/IP socket. The AO Control system provides all control and parameters to the RTC and can request that the RTC send it filtered diagnostic information for status monitoring of the system as well as atmospheric conditions.

Once placed in the AO correcting mode, the RTC runs without intervention. If parameters need to be changed because of a layer height shift, a rotation of the aperture (Keck’s is not circular), a change in the guide star positions relative to the science object, etc. these can all be accomplished without interrupting the operation. The AO Control sends the appropriate new parameters to the RTC. The RTC synchronizes all asynchronous requests from the Control Processor to the start of a frame.

### 3.3 Data flow

Data flows linearly, at high rates, from the HOWFS cameras to their individual reconstructors, see Figure 4. The reconstructed data from the three patrolling HOWFSs are sent to DM command generators that control the HO DMs. These DMs sharpen the T/T stars viewed by the LOWFSs. The T/T data from the LOWFSs is used to create an estimate to correct the T/T in the science field. One of the T/T stars also supplies focus and astigmatism information that is offloaded and processed by the non-real-time telescope system AO control.

The data from the four tomography HOWFS cameras is processed by their individual HOWFSs and the wavefronts are sent to the Tomography Engine. Here, a tomographic estimate of the atmosphere covered by the tomography HOWFSs is created and a correction wavefront for the science object is generated. The low order and high order corrections are split and sent to the low order and high order DM command generators.

Since the DMs in the HOWFSs (patrolling and tomography) operate in open loop, command generators for the WFSs perform a non-linear transform of the desired shape to generate DM commands to achieve that shape. This transform must be extremely accurate since any error will appear directly in the overall error budget for the system.

The data rates are substantial. The total camera data rate is >1.2TB/sec; centroid, wavefront and DM commands add over 0.3TB/sec. Additionally, this and other internal data must be able to be saved to the disk sub-system at full data rate during operation without impacting operation. The disk sub-system is able to record 3GB/sec for over 12 hours (~two days) of operation.

### 3.4 AO control loops and algorithms

As stated, all HO DMs are controlled open loop. The LO DMs are controlled closed loop.

Modal algorithms for reconstruction and tomography (see Figure 5 and Figure 6) have two significant advantages over zonal algorithms: they provide a means to easily both reduce temporal errors by using wind prediction and errors due to systematic vibrations due to wind shake and the like. For details of the algorithms used see [9].

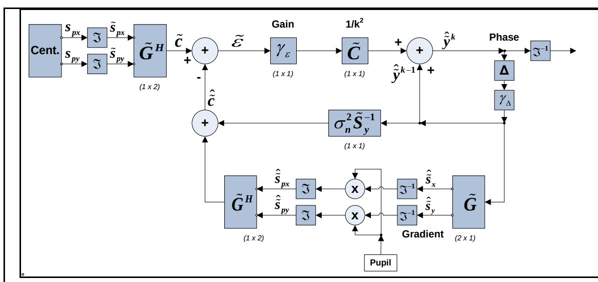


Figure 5 Reconstruction Algorithm[9]

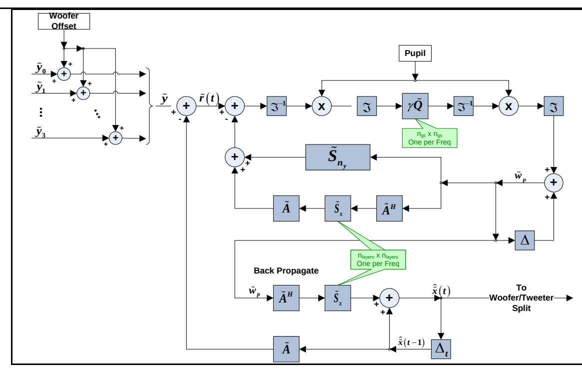


Figure 6 Tomography Algorithm[9]

### 3.5 Wavefront sensors (WFS)

The atmosphere is probed by 10 wavefront sensors (WFS): 3 LOWFS which sense Tip/tilt within a 2' FoV, one of which also senses focus and astigmatism; 3 LGS WFSs to sharpen the 3 T/T stars, allowing the use of dimmer stars; 4 LGS WFSs for tomography to estimate the atmospheric volume over a 40" FoV.

All WFSs operate synchronously at integer multiples of the shortest stare time.

**Tomography HOWFSs (4):** each focuses on an LGS in a constellation. They use a 256x256 camera image, have 4K subapertures (4x4 pixels) and operate at up to a 2 KHz frame rate. The information for the tip/tilt mirror associated with each of these HOWFSs is derived from their centroids. 64x64 subaperture T/T removed wavefronts are sent to the tomography engine.

**Patrolling HOWFSs (3):** each focuses on an LGS placed near an associated NGS, which is used for science object T/T sensing. They use a 128x128 camera image, at up to a 2 KHz frame rate. The T/T removed wavefront from this

HOWFS is sent to a 1K actuator DM to sharpen its associated NGS allowing the system to use dimmer T/T stars, increasing sky coverage. The DM is controlled open-loop.  The information for each tip/tilt mirror associated with each of these HOWGSs is derived from its centroids.

**LOWFS (3):** each LOWFS processes data from an NGS in the 2‘ FoV.  They are used to estimate for T/T, focus, and astigmatism measurements over the science field.  They have ~100x100 camera images that operate at up to a 2 KHz frame rate.  T/T information for the NGS stabilization mirror is derived from this star.  T/T for the science field is estimated by combining the measurements of the 3 LOWFS.

### 3.6 Tomography engine

The Tomography Engine is a systolic array[10].  The massive-parallel computing architecture is based on a systolic arrangement of small processing elements (PEs), all operating in parallel, and with data communication paths to their immediate neighbors. The PEs perform computations locally, and all simultaneously. Data can be transmitted between neighbors in a given direction in one simultaneous shift operation.

With this arrangement of processors, it is possible to map the physical space on to the array directly. At various stages in the algorithm, processors represent:

- 3-D spatial sample points in the atmospheric volume
- 2-D spatial sample points on the aperture associated with each wavefront sensor
- 2-D spatial sample points on the aperture associated with each deformable mirror
- 2-D Fourier domain sample points in each layer of atmosphere
- 2-D Fourier domain sample points on the aperture associated with each wavefront sensor
- 2-D Fourier domain sample points on the aperture associated with each deformable mirror

The algorithm requires very few operation types (+, -, X, shift, load and store).  So the PE can be extremely simple.  A 2D DFT with a systolic array can be performed with circular shifts, first in x, then in y, multiply/accumulating by Fourier constants appropriate to the PEs position and the source of the data shifted in.  The tomography engine performs 5 parallel 88x88 2D DFTs in ~4 µsec.

The tomography algorithm, see Figure 6, calculates the minimum variance solutions for the estimated volume given the wavefront measurements presented.  These calculations are embarrassingly parallel.  The calculations for any voxel or subaperture (or their Fourier transforms) require little or no information from others and no global storage access. Additionally the local storage requirements are less than 2 KB.  This type of problem maps very well to a systolic array, where each PE communicates only with its immediate neighbors.  This minimizes communication delays between elements.  The PEs are also pipelined so virtually no latency is lost in calculations involving PE to PE communications.

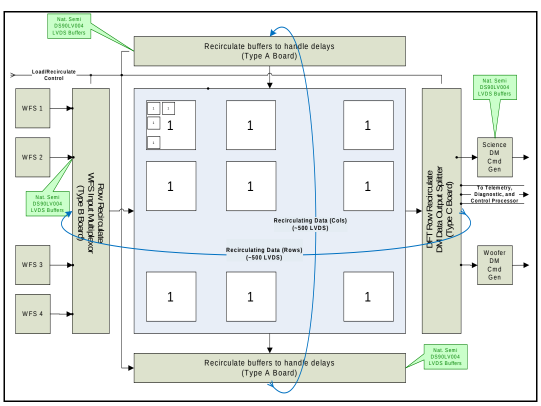


Figure 7  Tomography Engine Systolic Array

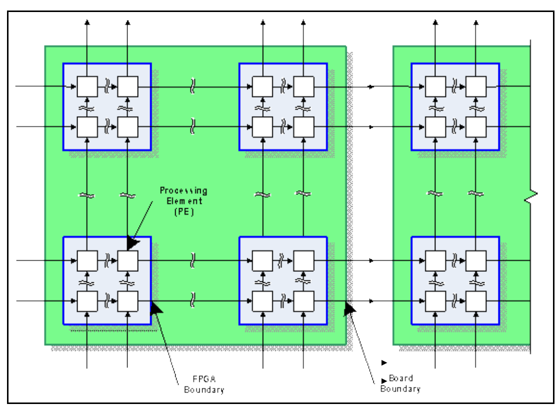


Figure 8  Tomography Engine detail

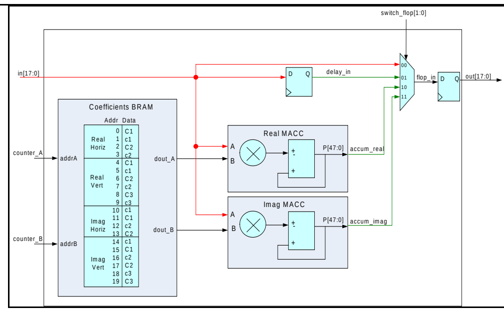


Figure 9 Tomography Engine Processing Element

### 3.7 DMs and DM command generation

There are five DMs controlled by the RTC. After the Tomography Engine has generated the wavefront to be used to correct the science object, the low order modes are extracted and sent to a 400 actuator low order high-stroke DM. This DM corrects the light for all WFS in the 120' FoV and the science object. It operates in closed loop mode. The high order modes are sent to the 4092 actuator high order DM, which corrects only the science object. This DM operates in open loop. The three DMs of the patrolling HOWFSs used to sharpen the NGSs have 32x32 actuators.

All four of the HO DMs operate open-loop and therefore we must have the ability to give them commands that will place the correct shape on their surface. Since there is no feedback on errors, we cannot rely on the closed loop control to correct errors for us. The command generators correct for non-linearity in the DMs response and the effect of the influence of their neighboring actuators and convert the desired shape into a set of actuator commands that will attain that shape directly, without the aid of a closed loop.

### 3.8 Other functions and modes

In addition to the operations required to process the data for reconstruction, tomography, and DM command generation, several other systems and processes are supported.

**Control Processor:** the control processor provides an asynchronous interface between the rest of the RTC and the non-real-time AO control. It provides all internal control of the RTC. This includes processing commands from the AO Control, monitoring all RTC systems, keeping logs of significant events, sending messages to the telescope system AO Control to keep it updated on the RTC's status.

**Parametric control:** all data to compensate for changes in camera response, layer height, meta pupil size for tomography, preconditioning data, selection of different centroid algorithms, reference centroids and offsets for long term or static aberrations, flat, etc. are sent to the RTC by the AO Control during operation to keep its operation in the selected range in real-time without disrupting the science exposure.

**Status and diagnostics:** the RTC is a complex system. It operates as a slave to the AO Control system. As a result, it must keep the AO Control appraised of its current state and any parametric data necessary to monitor the operation of the AO system. This includes power control and environment status (power, temperature, air flow, vibration …), data for PSF reconstruction, atmospheric monitoring, etc.

**Telemetry and the disk sub-system: t**his sub-system is not intended to be used for archival storage. It is intended to be used to capture real time data that may be of use for regenerating PSFs, atmospheric conditions and diagnostics. It is designed to be able to capture >3 GB/sec of streaming telemetry data with no loss during operations. Total data capacity is > 50TB. This is approximately 2 nights of data capturing all sources at full rate. Normal data needs would not be this great, but during development and diagnostic operation, it will be required. To insure that the state of the RTC when data was recorded can be determined, the data is recorded as FITS files, with the complete RTC configuration in the header.

**RTC sub-system Test Bench:** this system is not shown; however, while it is not attached to the telescope system, the Test Bench is an integral part of the NGAO RTC. It has the ability to simulate inputs and capture output for and sub-system of the RTC. This provides us the ability to perform incoming inspection on any RTC sub-system and diagnose problems while a sub-system is removed from the RTC. It also provides a convenient method to test new algorithms and components without interrupting the normal operation of the RTC itself.

**Wind analysis and compensation:** wind is the primary source of temporal error for which the AO system needs to compensate. Vibration in the telescope system will also degrade the science image and must be detected and compensated for[11]. The systematic error can be seen in a temporal-frequency analysis of the AO compensation residuals.

### 3.9 Physical Architecture

To limit heat dissipation in the dome and minimize potential sources of vibration, the minimum set of equipment is left on the Nasmyth. The rest has been moved to the computer room below the dome floor. Communications between the Nasmyth and the computer room is though high speed fiber, see Figure 10.

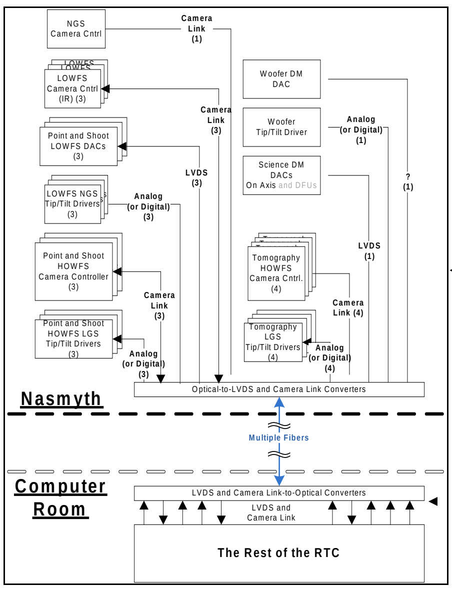


Figure 10 The RTC physical architecture

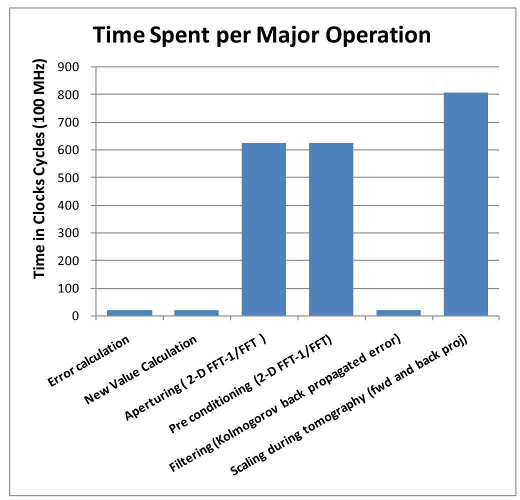


Figure 11 Time Spent in major operations for tomography algorithm

## 4. TECHNOLOGY ISSUES AND DECISIONS

### 4.1 Challenges

The key tech~~nical challenges in designing the RTC~~ were: the massive data processing requirements; the extremely high I/O rates for telemetry, diagnostics and data processing; handling the open-loop command generation; the size and complexity issues of such a large system; and balancing technology choices against the algorithm selections.

Processing the data requires approximately 1 terra operations per second, with latency < 600 µsec, at a 2 KHz frame rate and a data throughput rate of >1-GB/Sec. The computational requirements exceed those of the previous generations by over two orders of magnitude. Computational loads for tomography alone approach terra operations per second. Traditional solutions using CPUs and DPSs are no longer adequate given their capabilities in the low giga-operations per second. Operating with the Fourier mode iterative algorithm required a number of 2D DFT's per iteration. These operations are significantly compute intensive and tend to dominate the computation load, see Figure 11.

The RTC is also a much larger system than the current Keck AO system, with many sub-systems. This increased size leads to more complexity in keeping the various sub-systems synchronized. Intersystem communication can no longer be within a single processor or rack. The increased complexity also increased the amount of diagnostic and telemetry data to transport and manage. The disk sub-system to store the telemetry must be able to store a stream of data that can be over 3GB/sec for extended periods (hours) without pausing the stream or losing data.

In addition to reconstruction and tomography, open-loop operations require additional processing of each wavefront to generate HO DM's commands. This includes a large matrix multiply and processing each actuator values through a 16 million entry nonlinear lookup table (LUT).

The use of any algorithm impacts the hardware architecture decision. The hardware must be balanced to implement the algorithms and the algorithms must be chosen to match the capabilities of existing hardware solutions.

### 4.2 Technology choices

**Multi-Core CPUs:** these systems are parallel, but not massively parallel. They are high-speed, have large memory capacity, but memory access but can be slow when multiple CPUs and cores access it simultaneously. They generally have the highest power requirements of any solution. They require detailed understanding of the multi-core CPU architecture for efficient results. However, it is easiest to find this knowledge in the marketplace.

**GPUs:** systems using GPUs to solve highly parallel problems generally have lower power than multi-core CPUs. They are massively parallel, extremely high-speed, and available off the shelf with good development tools. They have very high compute rates, but also have a high penalty for memory accesses. Efficient programming of these devices requires a detailed understanding of the GPU architecture for efficient results. This knowledge is harder to find in the marketplace than that for multi core CPUs.

GPUs do well when they can be loaded with a problem which is large enough to keep their processors busy enough to mask their high memory penalty. Unfortunately, the size of current AO problems is at the lower bounds of the GPU efficiency.

**FPGAs:** systems built with FPGAs generally have the lowest power. The can be configured into massively parallel, high-speed systems. However, their native mode of arithmetic is fixed point and they have limited memory compared to GPUs and multi-core CPUs. However, in our algorithms little memory is needed for the calculations on any particular element and virtually no global memory accesses are needed. FPGAs require considerably more programming than the other implementation possibilities and require detailed understanding of the underlying FPGA architecture to program the hardware for efficient results. The availability of this knowledge is harder to find in the market place.

### 4.3 Technology decisions

The following are the technology decisions that were made based on the specific requirements of the NGAO RTC.

**Fixed point vs. floating point:** We determined that fixed point would provide adequate computational accuracy using scaled integer arithmetic[12]. This takes into account the minimum accuracy of data, 47-bit accuracy of accumulator, and the longest sequence of operations in a single iteration.

**Systolic array:** In bus based systems found in PC architectures, the bus bandwidth must be shared between the data input to be processed, the processed data out, the telemetry and diagnostic data out, instruction fetches and processing parameter I/O. In the systolic array, these paths are all distinct and can run in parallel without interference. Since we had extremely large data transfer requirements for both processing and telemetry, we determined that the systolic array provided enough advantage to select it for tomography.

**Size of problem:** using vector matrix multiplies or the iterative Fourier algorithm for reconstruction, a GPU implementation for 64x64 subapertures could not fit within a single GPU and still meet the timing requirements of <200 µsec. Dividing the problem between multiple GPUs led to increased I/O times for communications between the GPUs. This led to the preferred implementation of a systolic array based on FPGAs. For the algorithms and timing of the Fourier transforms, see Figure 5, Figure 6, and Figure 11.

The tomography WFS algorithm is executed in the tomography engine due to the speed required. However, for the patrolling HOWFS (32x32) the problem is more tractable and we are still considering GPU's as a solution.

## 5. RESULTS

### 5.1 NGAO RTC system specifications

The RTC is a large, complex, massively-parallel, real-time system, see Figure 12. Its specifications push the limits of current technology.

It has a total latency of <600 µsecs and it operates at a frame rate of 2 KHz.

All sub-systems have both hardware and software diagnostics as well as monitoring built in to assist in identifying problems before they occur and in resolving them when they occur.

The RTC operates in both tomographic mode and a NGS mode, which uses no laser guide stars.

Temporarily filtered diagnostic data is offloaded constantly for monitoring by the telescope system AO Control.

The multi-core processors that are used use a real-time Linux kernel for all real-time functions.

The majority of the AO RTC and its heat emission is located away from the dome successfully minimize heat and vibration issues. Each rack on the Nasmyth is equipped with a vibration sensor, which monitors the racks vibration. Each sensor is in turn monitored by the control processor, which will send warnings to the AO Control if the vibrations exceed pre set limits.

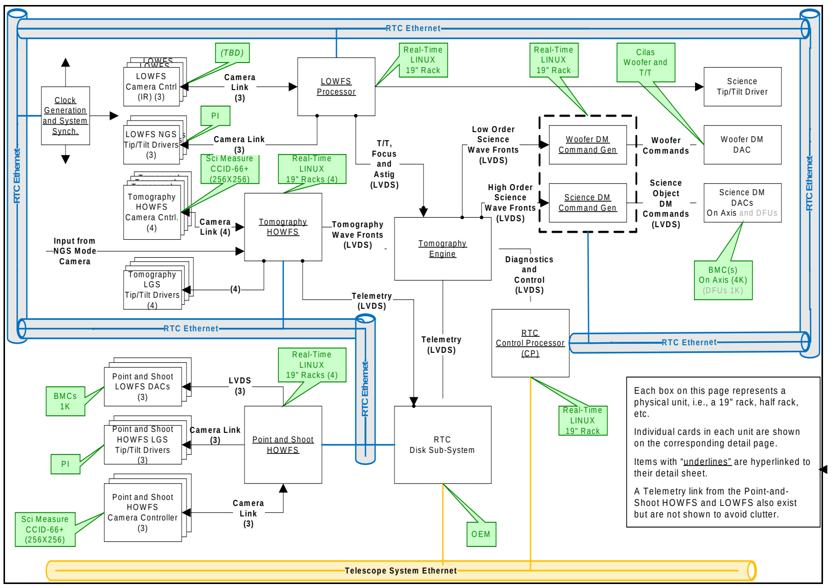


Figure 12 Detailed view of the RTC architecture

The Tomography Engine is implemented as a fixed-point, massively-parallel, systolic array using Xilinx Virtex®-6 FPGAs. Redundancy in both board and FPGA resources allows for easy remote reconfiguration to mitigate the effect of failures during operation.

It achieves >800 giga-operations per second. Each PE executes 6700 operations per frame for tomography and 4500 operations per frame for reconstruction. There are 88x88 PEs organized in 5 layers giving sustained peak rates of >1 tera-operation per second. A subset of the tomography engine will be prototyped in the detailed design phase of this project.

The disk sub-system is capable of storing multiple nights of data (>50 TB) @ >3 GB/sec for post processing and analysis: atmospheric studies, PSF reconstruction, etc.

| **Item** | **Keck II** | **NGAO** | **Comment** |
|---|---|---|---|
| WFS (HOWFS + LOWFS) | 2 | 10 | NGAO has 7 HOWFSs: 4 for tomography,,3 to sharpen the NGSs, and 3 LOWFS for TTFA |
| AO compensation down to: | 1.2 µm | 700 nm | |
| DMs controlled | 1 (400 element) | 5 | NGAO ( 3-1020, 1-4092, 1-400) |
| T/T mirrors controlled | 1 | 11 | 10 + 1 stage on woofer |
| Tomography | No | Yes | 5-layers, 88x88 subapertures per layer |
| Frame rate | 2KHz | 2 KHz | |
| Compute latency | <500 µsec | <600 µsec | |
| Fourier Transform Speed | N/A | 4 µsec | 5 simultaneous 88x88 2D DFTs |
| Power dissipated | <<20 KW | <20 KW | |
| Wind compensation | Not capable | Capable | Not currently planned for use |
| Vibration compensation | Capable | Capable | Not currently planned for use |
| Corrected FoV | | 40” | |
| MOAO capability | Not capable | Capable | RTC is capable, but, there are no deployable IFUs |

Table 1 Summary of the changes in the Keck RTC for NGAO

## 6. CONCLUSIONS

Computational loads required by the next generation AO systems are extremely large and require the use of massive parallelism. The use of alternatives to multi-core CPU implementations such as GPUs, FPGAs, and novel architectures such as systolic arrays must be used. The decision of what technology to use in different places in these complex systems is in itself difficult and no simple or comprehensive answer exists.

Communication, I/O, and data storage requirements in these systems are also increased to extremely high levels. The need to meet these simultaneously with the needs of the basic data processing plays a significant part in the overall architecture and implementation decisions.

Further, the implementation decisions for today’s RTC systems are closely coupled with the algorithm choices and vice versa. They cannot be made independently. We have chosen FPGAs in a systolic array for the computationally intensive tomography operations. This decision is but one point in a complex space of possible implementations. We suspect that the curve of optimum solutions is rather broad, but one does not often have the luxury of investigating all options to pick the “absolute” best. At the extreme performance levels required, tradeoffs and concessions must be made.

Because compensation and analysis of wind and systematic vibration become significant elements in the error budget for large telescope’s AO systems, we have implemented a system in which it is easy to integrate compensation for these effects at a later time without having to modify the RTC.

## 7. 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.

This material is based upon work supported AURA through the National Science Foundation under Scientific Program Order No. 5 as issued for support of the Telescope Systems Instrumentation Program (TSIP), in accordance with Proposal No. AST-0335461 submitted by AURA.

## REFERENCES


[1] Gavel, D., Reinig, M., “Fast hardware implementation of tomography for multi-guidestar adaptive optics,” Proc. SPIE Vol. 5903, (2005)

[2] Marichal-Hernández, José G., Rodríguez-Ramos, José M. and Rosa, Fernando, “Modal Fourier wavefront reconstruction using graphics processing units,” Journal of Electronic Imaging, vol. 16, issue 2, p. 023005, (2007)

[3] Bouchez, A., Dekany, R., “The PALM-3000 high-order adaptive optics system for Palomar Observatory,” Proc. SPIE 7015 (2008)

[4] Lloyd-Hart, M., Angel, R. “Design of the Adaptive Optics Systems for GMT,” SPIE Vol. 6272, (2006)

[5] Wizinowich P., “Keck Next Generation Adaptive Optics: AO-Level Control Loops,” Keck Adaptive Optics Note 705, February 19, 2010

[6] Gavel D., et al, “Concept for the Keck next-generation adaptive optics system,” SPIE Proc. 7015, (2008)

[7] Wizinowich P., et al, “W.M. Keck Observatory’s Next Generation Adaptive Optics Facility,” Proc. SPIE 7736, (2010)

[8] Wizinowich P., et al, “Next Generation Adaptive Optics: Preliminary Design Manual,” Keck Adaptive Optics Note 768, May 24, 2010

[9] Gavel, D. et al, “Wavefront control algorithms for the Keck next-generation adaptive optics system,” SPIE Proc. 7736, (2010)

[10] Kung, H. T. and Leiserson, C. E, Systolic Arrays (for VLSI). Proceedings of Sparse Matrix Symposiums, pages 256–282, 1978

[11] Brix, M., Naranjo, V., et. al., “Vibration measurements at the Large Binocular Telescope (LBT),” SPIE Vol. 7012 (2008)

[12] Schatzman, James C., “Accuracy of the discrete Fourier transform and the fast Fourier transform,” SIAM J. Sci. Comput. Vol 17, No 5, pp 1150-1166, September, 1996