REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
31-07-2008

2. REPORT TYPE
Final

3. DATES COVERED *(From - To)*
15-02-2006 to 31-07-2008

4. TITLE AND SUBTITLE
High Resolution Imaging Testbed Utilizing Sodium Laser Guide Star Adaptive Optics:
The Real Time Wavefront Reconstructor Computer

5a. CONTRACT NUMBER
n/a

5b. GRANT NUMBER
FA9550-06-1-0343

5c. PROGRAM ELEMENT NUMBER
n/a

6. AUTHOR(S)
Dekany, Richard G., Bouchez, Antonin H.

5d. PROJECT NUMBER
n/a

5e. TASK NUMBER
n/a

5f. WORK UNIT NUMBER
n/a

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
California Institute of Technology
1200 E. California Blvd.
Mail Code 201-15
Pasadena, CA 91125

8. PERFORMING ORGANIZATION REPORT NUMBER
n/a

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AF Office of Scientific Research
875 N. Randolph St. Room 3112
Arlington, VA 22203
Dr. Kent Miller, AFOSR/NE
Kent.miller@afosr.af.mil
Technicalreports@afosr.af.mil

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR, DURIP(FY06)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
n/a

12. DISTRIBUTION / AVAILABILITY STATEMENT
Unclassified

13. SUPPLEMENTARY NOTES
None

14. ABSTRACT
We have procured and assembled a real-time wavefront reconstructor computer for the PALM-3000 high-order upgrade to the Palomar adaptive optics system. This innovative computer uses a network of 17 graphics processing units to invert the measurements of a 64 × 64 wavefront sensor and reconstruct the wavefront at 3368 locations across the Hale Telescope pupil, at up to 2 kHz frame rate. A cluster of 10 desktop personal computers house the graphics processors, along with a telemetry recording system and timing hardware. The entire computer cluster is housed in two electronics racks, and connected via a high-speed data switch. We have benchmarked the performance of the system, and demonstrated that it meets the requirements of the PALM-3000 adaptive optics system, with a total latency of just 220 μs per frame. We are currently in the detailed design and integration phases of the optical and software components, and expect to commission the system at Palomar Observatory in early 2010.

15. SUBJECT TERMS
Real-time wavefront reconstructor computer, adaptive optics

16. SECURITY CLASSIFICATION OF:
a. REPORT U
b. ABSTRACT U
c. THIS PAGE U

17. LIMITATION OF ABSTRACT
UU

18. NUMBER OF PAGES
6

19a. NAME OF RESPONSIBLE PERSON
Richard Dekany

19b. TELEPHONE NUMBER *(include area code)*
626-395-6798

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

High Resolution Imaging Testbed Utilizing Sodium Laser Guide Star Adaptive Optics: The Real Time Wavefront Reconstructor Computer

Final Performance Report

Summary

We have procured and assembled a real-time wavefront reconstructor computer for the PALM-3000 high-order upgrade to the Palomar adaptive optics system. This innovative computer uses a network of 17 graphics processing units to invert the measurements of a 64 × 64 wavefront sensor and reconstruct the wavefront at 3368 locations across the Hale Telescope pupil, at up to 2 kHz frame rate. A cluster of 10 desktop personal computers house the graphics processors, along with a telemetry recording system and timing hardware. The entire computer cluster is housed in two electronics racks, and connected via a high-speed data switch. We have benchmarked the performance of the system, and demonstrated that it meets the requirements of the PALM-3000 adaptive optics system, with a total latency of just 220 μs per frame. We are currently in the detailed design and integration phases of the optical and software components, and expect to commission the system at Palomar Observatory in early 2010.

Publications:

Bouchez, A. H., Dekany, R. G., Angione, J. R., Baranec, C., Britton, M. C., Bui, K., Burruss, R. S., Cromer, J. L., Guiwits, S. R., Henning, J. R., Hickey, J., McKenna, D. L., Moore, A. M., Roberts, J.E., Trinh, T. Q., Troy, M., Truong, T. N., Velur, V., "The PALM-3000 high-order adaptive optics system for Palomar Observatory" in Adaptive Optics Systems, edited by Norbert Hubin, Claire E. Max, Peter L. Wizinowich, Proceedings of SPIE Vol. 7015 (SPIE, Bellingham, WA 2008) 70150Z.

Truong, T. N., Bouchez, A. H., Dekany, R. G., Shelton, J. C., Troy, M., Angione, J. R., Burruss, R. S., Cromer, J. L., Guiwits, S. R., Roberts, J. E., "Real-time wavefront control for the PALM-3000 high order adaptive optics system" in Adaptive Optics Systems, edited by Norbert Hubin, Claire E. Max, Peter L. Wizinowich, Proceedings of SPIE Vol. 7015 (SPIE, Bellingham, WA 2008) 70153I.

Personnel Supported:

This award is limited to equipment purchases for development of the real-time wavefront reconstructor computer. Personnel involved in this project include

Richard Dekany, Caltech, PI
Antonin Bouchez, Caltech, Other Personnel
Christoph Baranec, Caltech, Postdoc
Matthew Britton, Caltech, Other Personnel
Khanh Bui, Caltech, Other Personnel
John Cromer, Caltech, Other Personnel
John Henning, Caltech, Other Personnel
Jeff Hickey, Caltech, Other Personnel
Dan McKenna, Caltech, Other Personnel
Anna Moore, Caltech, Other Personnel
Velur, Viswa, Caltech, Other Personnel
John Angione, JPL, Other Personnel
Rick Burruss, JPL, Other Personnel
Stephen Guiwits, JPL, Other Personnel
Jennifer Roberts, JPL, Other Personnel
J. Chris Shelton, JPL, Other Personnel
T.Q. Trinh, JPL, Other Personnel
Mitchell Troy, JPL, Other Personnel
Tuan Truong, JPL, Other Personnel

Interactions/Transitions:

This project was presented at the international meeting “SPIE Astronomical Telescopes and Instrumentation 2008**,** 23 - 28 June 2008.

There are no patent disclosures to report.

Introduction

PALM-3000 is a high-precision adaptive optics (AO) upgrade to the successful Palomar Observatory Adaptive Optics System (PALMAO) on the 5.1 meter Hale Telescope at Palomar Observatory, currently under development at the California Institute of Technology (Caltech) and the Jet Propulsion Laboratory (JPL). It will use a 3368 active actuator deformable mirror to compensate the atmospheric wavefront on scales as fine as 8 cm at the telescope pupil, using both natural and sodium laser guidestars (NGS and LGS). Such fine correction provides unique new science opportunities, from extreme contrast on bright stars in the infrared to all-sky diffraction-limited imaging and spectroscopy in the visible. A suite of new back-end instruments is being developed to exploit these capabilities, including the SWIFT visible-light integral field spectrograph, Project 1640, a near-infrared coronagraphic integral field spectrograph, and 888Cam, a high frame rate visible light imager.

The real-time wavefront reconstructor computer requirements for PALM-3000 are roughly two orders of magnitude greater than those of the current PALMAO system, which used a single 241 active actuator deformable mirror and 16×16 subaperture Shack-Hartmann wavefront sensor. In contrast, PALM-3000 must control 3719 degrees of freedom (the 3368 actuators in the high-order tweeter mirror, 349 actuators in a low-order woofer mirror, and a tip/tilt mirror) using a 64×64 subaperture wavefront sensor at a minimum framerate of 2 kHz. The PALM-3000 wavefront reconstructor computer was therefore specified to allow full vector matrix multiplication (VMM) of the 8192 simultaneous slope measurements (2 in each subaperture) by a 3719×8192 element reconstruction matrix, with a total latency of less than 250 μs.

Candidate Computer Architectures

Several candidate hardware architectures were evaluated for this challenging computational task, including fixed-point and floating-point digital signal processor boards, and a network of graphics processor units (GPUs). Given the in-house experience with the Texas Instruments (TI) floating point DSP architecture from the PALMAO system, we began the evaluation with two commercially available DSP-based systems.

The first system we evaluated was based on fixed-point DSP boards from Lyrtech each featuring four 1-GHz TMS320C6416 DSPs from TI, organized in two clusters of two with the inter-cluster communication bandwidth limited to 528 MB/s by the 64-bit 66MHz PCI bus connecting the clusters. Each DSP is provided with an impressive 8 GMACs of 16-bit peak processing power and 128 MB of dedicated off-chip memory with sustained memory bandwidth of 800 MB/s via a private 64-bit 100-MHz bus. Unfortunately, this memory bandwidth proved inadequate to compensate for the measly 1 MB of on-chip cache that equipped each DSP, given the large size (3368 x 6736) of the reconstruction matrix. As a result, the Lyrtech solution would require a relatively high number of DSPs. This is true for most, if not all, commercially available systems based on current generation of TI DSPs.

The second system we evaluated was based on floating-point DSP boards from Bittware each featuring eight 500-MHz ADSP-TS201 TigerSHARC DSPs from Analog Devices, organized in two clusters of four. Each cluster is provided with 256 MB of shared memory for intra-cluster access via a common 64-bit 83.3-MHz cluster bus at aggregate sustained memory bandwidth, or equivalently intra-cluster communication bandwidth, of 667 MB/s. Like the Lyrtech, the Bittware clusters are also interconnected via a 64-bit 66-MHz PCI bus with identical sustained inter-cluster communication bandwidth of 528 MB/s. Unlike the Lyrtech, each DSP on a Bittware board offers 3 MB of on-chip memory and 3 GFLOPs of 32-bit peak processing power. Based on the performance specifications of the optimized floating point math library from Bittware, we estimated this peak realistically to be between 1 GMACs and 1.3 GMACs, which proved in the final analysis to be the limiting factor in determining the number of DSP boards in the Bittware system. Hence, both DSP-based systems were dropped from consideration due to the high cost, which exceeded not only the budget for this project but also the cost of either GPU-based system by at least 50%.

PALM-3000 Wavefront Reconstructor Computer

The chosen wavefront reconstructor hardware consists of 16 off-the-shelf NVIDIA 8800 Ultra graphics cards which are distributed over 8 dual-core Opteron PCs from HP, each hosting 2 cards, the maximum number of PCI Express (PCIe) x16 cards permitted by PCIe 1.1 Specification. Since our acquisition, PCIe 2.0 has been released with support for up to 4 such cards per PC.

Each NVIDIA 8800 Ultra features 576 GFLOPS on 128 612-MHz single-precision floating-point SIMD processors, arranged in 16 clusters of eight. Each cluster provides 8 K registers and 16 KB of shared on-chip memory organized in 16 banks. Accessing the shared memory is as fast as accessing a register as long as there are no bank conflicts. Also included, but not currently availed by our applications, per cluster is 64 KB of constant memory with a cache working set of 8 KB. The clusters are interconnected via a 384-bit memory bus, providing 103.7 GB/s memory bandwidth to 768 MB of global shared GDDR3 RAM.

It is this winning combination of supercomputing power and unparalleled memory bandwidth that earned NVIDIA GPUs the selection for this project. In particular, it is the ease of programming coupled with unfettered access to the tremendous processing power of the GPU through a simple low-level C programming interface that won NVIDIA over the competing ATI GPU architecture which we also evaluated, though briefly. To aid developers in the software development, NVIDIA provides the CUDA (Compute Unified Device Architecture) environment, which consists of the C cross compiler, debugger, host driver, two mathematical libraries of common usage, FFT and BLAS, and a plethora of code samples, in addition to the GPU API and its runtime.

The configuration is a cluster of eight PCs along with a central node (PC0, with one GPU for low-order wavefront reconstruction in LGS observing mode) and a tenth node hosting a telemetry database are interconnected using a Quadrics QsNet^{II} 16-port switch with extra free ports reserved for future add-on systems. The switch delivers over 900 MB/s of user space to user space bandwidth each direction with latency under 2 μs for a total of 14.4 GB/s of bisectional bandwidth and broadcast capability. This scalable switch combined with the 1-to-N optical splitter that delivers identical complete wavefront sensor frames to each of the N output ports and permits not only future integration but also concurrent executions of additional real-time systems and accelerators.

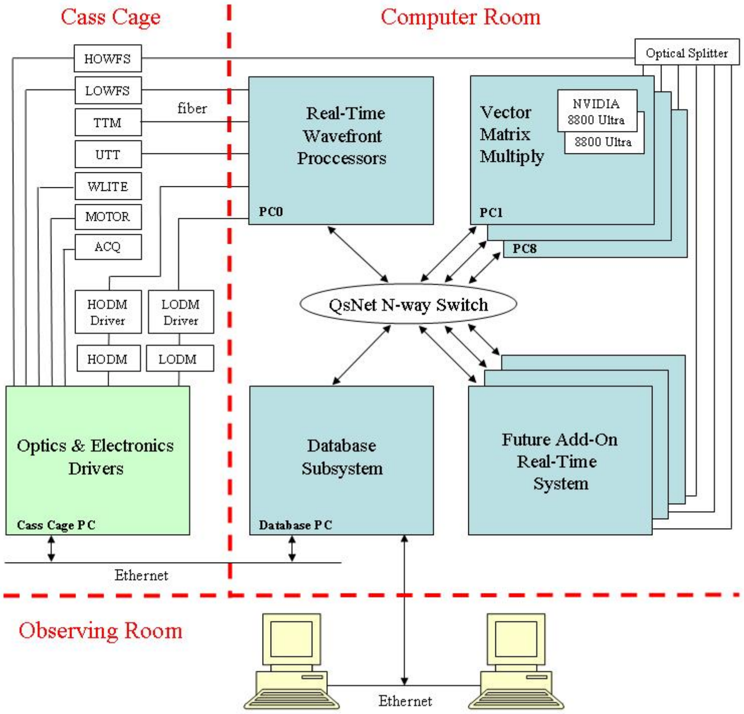

Figure 1: PALM-3000 hardware architecture schematic

Figure 1 illustrates the chosen PALM-3000 real-time computer architecture, and interconnections with other subsystems. Using this architecture, compute-intensive operations are off-loaded and accelerated by low-cost off-the-shelf GPU-based graphics cards hosted in a PC cluster interconnected with one or more ultra low-latency, high-bandwidth switches. High bandwidth data, such as pixels received by the eight PCs from high-order and low-order wavefront sensors (HOWFS and LOWFS), and latency-sensitive commands, such as those issued by PC0 to high-order and low-order deformable mirrors (HODM and LODM), tip/tilt mirror (TTM) and uplink tip/tilt mirror (UTT), are transferred using fiber optics. Low-bandwidth data, such as acquisition camera (ACQ) pixels and hardware status, are sent from the Cass Cage PC to the Database PC via a dedicated 1Gbit Ethernet. Latency-tolerant commands, such as those issued by the Cass Cage PC to motors and white light, including ACQ and DM configurations, are sent via direct connections.

Performance

Figure 2 illustrates the wavefront reconstruction timeline at 2 kHz frame rate. All times shown were measured on the actual purchased hardware, except the DM driver latency which is based on manufacturer's specifications. As soon as the 500 μs integration is complete, the charge accumulated on the detector is shifted to a masked region and a new integration begun. Pixel readout then lasts 500 μs, during which the next integration proceeds. Once half of the detector has been read out, a first batch of pixel data is sent to the VMM computers.

Using 16 parallel GPU processors, computing centroids requires 5 μs while the partial matrix multiplication requires another 196 μs. The results are returned to the central node (PC0) and the second half of the frame processed identically. Summing of the 32 resulting matrices and calculation of the DM commands requires an additional 16 μs and 5 μs, respectively. Thus the total latency from end of frame readout to DM commands is 222 μs, leaving 28 μs of headroom prior to the arrival of the first batch of pixels from the next integration. The total latency from integration midpoint to the settling of the central actuator of the DM is 779 μs.

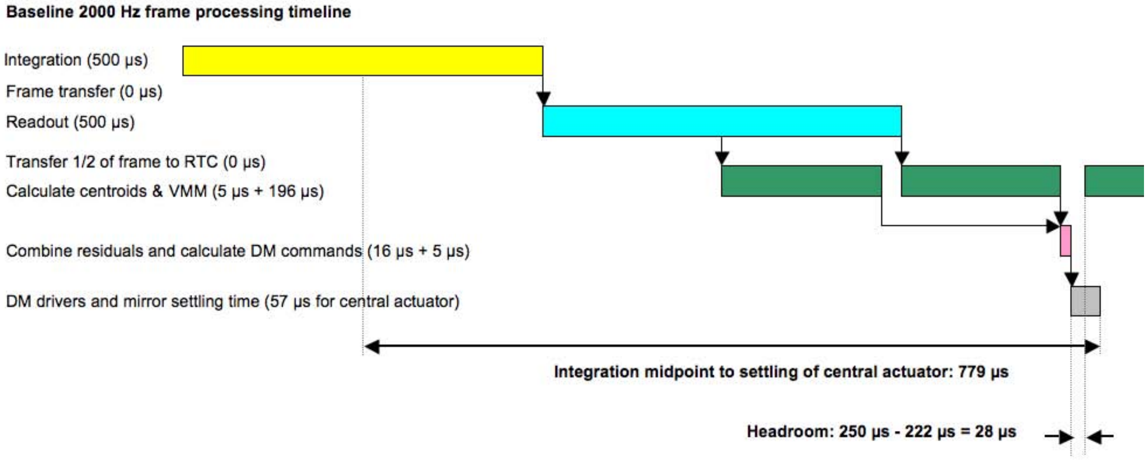

Figure 2: PALM-3000 processing timeline and compute latency at 2 kHz frame rate.

Procurement and Integration

Table 1 lists the procurements made with this award, which allowed the purchase of two identical real-time wavefront reconstructor computers, one to be integrated into the PALM-3000 system at Palomar Observatory and one to be permanently located in the lab at Caltech/JPL for software development and testing during the system's operational lifetime. In addition to the two computer clusters, we purchased one spare of each hardware component (with the exception of the Quadrics switches and RAID units).

We received all components by mid-April 2008, and are now in the process of integrating the computer clusters. Figure 3 shows a recent photograph of the two racks which make up wavefront processor 1, partially populated with 9 PCs (containing 17 GPUs). We expect integration and testing of the wavefront processor computers to be complete by end September 2008.

Description	Count	Cost
Hewlet Packard xw9400 Workstations	21	$92,921
EDT PCI Express digital video input cards	19	$16,792
XFX GeForce 8800 video cards	36	$27,615
Quadrics 16-port switches, network cards, cables	21 + 2	$44,227
SUSE Linux Enterprise Server and Real-time licenses	20 + 2	$4,296
NexSan RAID storage devices	2	$30,732
IRIG Synchronized Timing Cards	5	$7,046
Computer racks	4	$15,455
Grand Total		**$239,084**

Table 1: Real-time wavefront reconstructor procurements

Conclusions

With 3368 active actuators and an interactuator spacing in the pupil of just 8 cm, PALM-3000 will be the first extreme AO system deployed for astronomy. DURIP/AFOSR funding has allowed us to develop an innovative, high-performance wavefront processor computer for this system, based on a networked cluster of off-the-shelf graphics processor units.

Figure 3: *Left:* Rack layout design. *Right:* Partially populated computer racks in our lab.

This architecture is cost-effective and flexible, and provides unmatched ease of development through a simple low-level C programming interface to the GPUs. We have measured a total compute latency of just 222 μs, meeting our requirements for operation at up to 2 kHz. The wavefront reconstructor computer is currently being integrated in our lab, and we look forward to fielding it at Palomar Observatory in early 2010.