

Fast Hardware Implementation of Tomography for Multi-guidestar Adaptive Optics

Donald Gavel, Marc Reinig, and Carlos Cabrera
UCO/Lick Observatory, University of California, Santa Cruz, California 95064

ABSTRACT

An adaptive optics system using multiple deformable mirrors and an array of guidestars can correct over a wider field of view than traditional single DM systems and can also eliminate the cone-effect error due to the finite altitude of laser guidestars. In large telescope systems, such as the envisioned 30-meter telescope, or TMT, the extraordinarily large amount of computation needed to implement multi-conjugate adaptive optics at atmospheric turnover rates is prohibitive for ordinary CPUs, even when another ten years of computer development is taken into account. We present here a novel approach, implementing a fast iterative version of the key inverse tomography calculations in an array of parallel computing elements. Our initial laboratory experiments using field-programmable gate arrays (FPGAs) are promising in terms of speed and convergence rates. In this paper we present the theory and results from simulations and experiments.

Keywords: Adaptive Optics, Laser Guide Stars, Tomography Algorithms

1. INTRODUCTION

Several giant astronomical telescopes are now in their concept development stages with apertures of 20 meters or larger. The push toward large aperture telescopes, as well as the fact that astronomers wish to expand on the traditionally narrow adaptive optics field of view, has led to a need for real-time three-dimensional tomographic reconstruction of the atmosphere above the telescope and inside the field of view. Multiple laser guidestar beacons at different field positions can probe this volume. The problem is to efficiently and accurately reconstruct the volume.

We consider two basic types of wide-field adaptive optics correction approaches under consideration for instruments on the UC/Caltech/AURA/ACURA Thirty Meter Telescope (TMT), which is now in the conceptual design stage. Multi-conjugate adaptive optics (MCAO) provides correction over a wide imaging field by correcting the volume of atmosphere in a conjugate space. Wide field imaging cameras and spectrometers can then be placed behind an MCAO system. The current design goal for the NFIRAOS MCAO system on TMT is a 30 arcsecond diameter field. Multi-object adaptive optics (MOAO) uses knowledge of the volumetric distribution of turbulence to make corrections in specific science directions. The MOAO concept is intended for multi-object spectrographs and integral field unit imaging spectrometers where each deployable IFU has a very limited field of view (2 arcseconds on TMT) but a wide field over which these units can be deployed. The goal is a 5 arcminute diameter deployment field. With MCAO, multiple deformable mirrors are placed at conjugate altitudes and make corrections for layers of turbulence. With MOAO, one deformable mirror is dedicated to each science direction, making correction for the integrated turbulence through the atmospheric volume in that direction. In either case, the adaptive optics system must use guidestar information to determine the turbulence throughout the three dimensional volume of atmosphere above the telescope and within the desired field of view. Since the constellation of guidestars provides information in the form of line integrals through this turbulence, the process is analogous to that of computer-aided tomography.

A key challenge of implementing a multi-guidestar adaptive optic system on an astronomical telescope is the extraordinary amount of computation needed to perform this volumetric tomography hundreds of times per second in order to keep up with the changing atmosphere. Early investigation of computational requirements determined that a significant sized supercomputer would be needed to serve a near-IR AO system on a 30 meter telescope. Our reinvestigation of the data processing algorithms and data flow needed for the AO real time data processing and control system has suggested the massively parallel architecture described in this paper.

The hardware to implement the suggested architecture is available today. Custom logic gate arrays (ASICs) could be used, however, particularly flexible option, at least in the initial phases, is the use of field-programmable gate

array (FPGA) logic chips, which are available at low cost and are relatively easy to reprogram. The massively parallel hardware architecture represents a radical departure from the commonly used Von Neuman machine in traditional computers, where data queue up to use a single arithmetic logic unit. Instead, in the massively parallel architecture, large numbers of simple logic units operate on various parts of the data simultaneously. This alternative works very well on algorithms requiring repetitive operations on data with little cross-dependence. Almost all aspects of AO data processing, from wavefront sensor data processing to commanding the deformable mirrors, to the inverse-tomography calculations, fall into this category of algorithm.

A further advantage of our approach is the natural decomposition of the AO processing into clearly established functional sub-units. Some sub-units are associated directly with sensor and actuator hardware, suggesting that they can be separately optimized (and possibly built into) that hardware. The tomography sub-unit carries with it a physical reproduction of the atmospheric volume, holding estimates of the delta optical paths within volume elements. The modularity of the approach allows for ease of maintenance and upgrade, and the system engineering is considerably simplified since each sub unit depends only on design parameters associated with it. We contrast this idea with the “influence matrix” approach to designing AO systems where the AO control problem is treated as a black box with a large number of inputs (DM commands) and outputs (wavefront sensor readings) and the job of the controller is to invert this matrix which depends on all of the AO engineering parameters.

At the UCO/Lick Observatory Laboratory for Adaptive Optics we implemented a small scale version of back-projection tomography on a commercial FPGA development board in order to gain experience with the algorithm and to understand the scaling to larger systems. The sub-scale prototype runs the inverse tomography algorithm in real-time, but only for a small subsample of the aperture and atmospheric volume space. The algorithm is massively parallel in space so extension to the full size TMT controller only requires replication of the hardware. The inverse tomography algorithm is iterative so we measured the time per iteration with our hardware prototype and then used a computer simulation to determine the iteration rate (number of iterations needed for convergence) for the full scale 30 meter telescope problem. The results of the prototype and computer simulation studies are very encouraging, indicating that a system with on the order of a few hundred FPGA chips (a few 10’s of circuit boards) can serve a 7-9 guidestar AO system with 7800 subapertures per guidestar on the 30 meter telescope at 1kHz frame rate.

2. SYSTEM ARCHITECTURE OVERVIEW

2.1 Data flow through the architecture

The system architecture is depicted in Figure 1. Multiple wavefront sensors, corresponding to each guidestar in the constellation, feed data to a centralized tomography unit. The tomography unit determines an estimate of the delta optical path differences within the atmospheric volume. This information is then in turn used, in MCAO mode, to project best fits onto the finite layers represented by each conjugate deformable mirror, or, in MOAO mode, to project along paths to the multi-object science fields. Finally, since the deformable mirrors commonly have inter-actuator influence functions, a fit must be computed for the actuator commands so that the resulting DM shape best fits the wavefront given the influence functions and other limitations such as limited actuator stroke. The three step process, wavefront measurement, tomography, DM fitting, is evident in the figure. Also evident is the inherent parallelization of operations specific to wavefront sensors or to DMs.

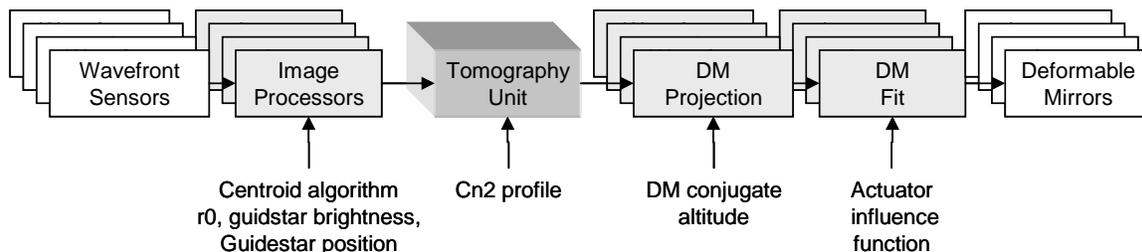


Figure 1. Multi-guidestar AO processing architecture

2.2 Wavefront sensors

Wavefront sensors can be of many different types depending on the type of source and the method of sensing. The most common approach in astronomical adaptive optics today is to use Hartmann sensors to measure the local wavefront slopes on the light from natural star or laser beacons. Another common approach is to measure the local curvature of the wavefront. A possible future technique is to directly measure the phase using a Mach-Zehnder interferometer setup. None of these sensors provide the piston component of the phase. Furthermore laser guidestars are not able to probe tip/tilt component, and, in the case of the 30 meter telescope, sodium beacons may not be able to sense the focus component either, due to the huge uncertainty of the mean altitude of the sodium layer relative to the telescope's depth of focus.

In the proposed control architecture, each wavefront sensor has a wavefront reconstructor associated with it, such that the data presented to the tomography unit is in the form of wavefront phase as a function of position on the aperture. Even sensors dedicated to sensing only low order modes, such as tip/tilt or tip/tilt/focus/astigmatism, will present their data to the tomography unit as wavefronts as a function of position on the aperture. These wavefronts will be lacking the high order components, as the high-order wavefront sensors are lacking the low order wavefront components, but this will not upset the structure of the tomographic algorithm or adversely affect its performance relative to that achievable given the information inherent in the measured data.

The task of finding phase from phase gradients is that of solving an overdetermined set of linear equations

$$\mathbf{s} = \mathbf{G}\mathbf{y} \quad (1)$$

where \mathbf{y} is vector of m wavefront points sampled within the aperture, \mathbf{s} is a vector of $2m$ wavefront slopes, and \mathbf{G} is the gradient operator. The phase from slopes problem is related to that of solving Poisson's equation in two dimensions. The noise-optimal solution, expressed in the Fourier domain is

$$\tilde{\phi}_{est}(\boldsymbol{\kappa}) = \frac{-i\boldsymbol{\kappa} \cdot \tilde{s}(\boldsymbol{\kappa})}{\boldsymbol{\kappa}^2 + \tilde{C}_{nm}(\boldsymbol{\kappa})/\tilde{C}_{\phi\phi}(\boldsymbol{\kappa})} = \frac{-i\boldsymbol{\kappa} \cdot \tilde{s}(\boldsymbol{\kappa})}{\boldsymbol{\kappa}^2} \left(\frac{1}{1 + \alpha(r_0\boldsymbol{\kappa})^{5/3}} \right) \quad (2)$$

where tildes indicate Fourier transforms and $\boldsymbol{\kappa}$ is the spatial frequency. In the noise optimal (Weiner filter) solution shown above, \tilde{C}_{nm} and $\tilde{C}_{\phi\phi}$ represent the power spectrum of the noise and the wavefront aberrations, respectively.

When we assume a Kolmogorov spectrum for the wavefront and a white noise distribution for the slope measurements, the second equality results, with $\alpha = (0.027)^{-1} (2\pi)^{11/3} (\sigma_n d_a)^2$ where σ_n is the standard deviation of the noise in slope units, d_a is the size of the subaperture, and r_0 is the Fried parameter characterizing the strength of the Kolmogorov wavefront aberrations.

As implied by equation (2), slope to phase calculations can be performed rapidly in the Fourier domain using the fast Fourier transform. An algorithm developed by Poyneer¹ mitigates problems at the boundary of a finite aperture through an approach that extends the slope measurements artificially beyond the aperture so as to keep the slopes curl-free ($\nabla \times s = i\boldsymbol{\kappa} \times \tilde{s} = 0$) throughout the domain of the fast Fourier transform. Once the Fourier coefficients of s are determined, the scalar equation (2) can be applied at each spatial frequency independently, hence it is amenable to massive parallelization. The trick of performing the fast Fourier transform itself in a massively parallel architecture is addressed in the section below on inverse tomography. The Poyneer extension algorithm is massively parallelizable since each pseudo-measurement point outside the aperture depends only on a few points at the nearest point within the aperture.

We should point out that all of the other practical aspects of pre-processing data from a Hartmann sensor are massively parallelizable. These include reading data from the sensor array, background removal, flat fielding, and centroiding. Since multiple processors can be interfaced to the array along parallel interface lines, the array itself should be designed with many parallel readout amplifiers to minimize transfer time. Since the sensor to computer interface would no longer be the bottleneck, these amplifiers can be optimized for lowest read noise at reasonable pixel rates. We might envision a customized Hartmann sensor CCD chip that includes on it an ASIC chip that performs the preprocessing in parallel and outputs centroid data at a parallel interface port. Carrying this idea further, the entire Poyneer Fourier based reconstructor can be implemented on this chip, and the output would be wavefront

phase (again at a parallel interface port). The chip would take the parameters α and r_0 of equation (2), which depend on seeing conditions and the brightness of the guidestar, as inputs.

2.3 Inverse tomography system

Given data from many wavefront sensors (including the tip/tilt and focus sensors) the job of the tomography system is to obtain a best estimate of the delta optical paths in the volume of turbulence above the telescope and within the science field of view (Figure 2). The tomography problem is represented by the underdetermined set of linear equations

$$\mathbf{y} = \mathbf{A}\mathbf{x} \quad (3)$$

where \mathbf{x} is a vector formed of all the delta-optical path length values within the three dimensional volume and \mathbf{y} is a vector formed of all the wavefront measurement samples. The matrix \mathbf{A} is a special kind of matrix, which we call a forward-propagation matrix; it acts only to form line integrals of \mathbf{x} through the volume. \mathbf{A}^T has a special meaning; it is a back-propagation matrix which would act on a \mathbf{y} vector so as to deposit measured wavefronts at each altitude along lines leading back to the respective guidestars.

A typical sampling for the \mathbf{x} voxels would be every few centimeters in the transverse plane (parallel to the ground), corresponding to the finest sampling of any of the wavefront sensors, and every few kilometers in the vertical direction, corresponding to the altitudes resolvable by the guidestar constellation.

There are several computational techniques for iteratively solving (3). The algorithms we will focus on here can be implemented with massive parallelization. First of all let's examine the properties of the solution to (3). Since the equation is underdetermined (i.e. we assume there are many more volume elements than there are wavefront sensor measurements), exact solutions are non-unique. Amongst the non-unique solutions we can choose one of our liking. For example, the minimum norm solution happens to be the one whose elements have the least average square value. A more preferable solution might be the minimum variance solution, the one constrained to solve (3) but otherwise was most likely given the second order statistical properties of \mathbf{x} a-priori. In any case, all the solutions along the constraint hyperplane $\mathbf{y} = \mathbf{A}\mathbf{x}$ differ from each other by a value that is *invisible to the wavefront sensors*, that is, $\mathbf{A}(\mathbf{x}_1 - \mathbf{x}_2) = 0$. Thus the particular choice of unique solution that meets constraint (3) is purely due to an a-priori preference of the system implementer! Of course, all measurements have noise, that is, $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$ where \mathbf{n} is noise, in which case a minimum variance solution might be pulled even further towards its a-priori preference by coming off the constraint hyperplane a bit, but only by roughly the standard deviation of the noise.

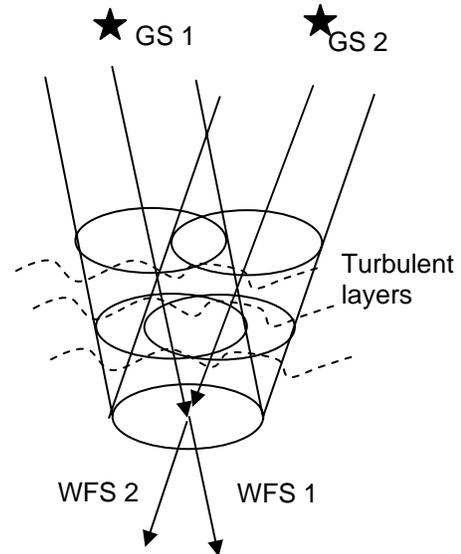


Figure 2. The volume of turbulence above the telescope is probed by guidestar rays.

The invisible modes are well known to adaptive optics engineers. For example spurious waffle patterns associated with rectilinear DM actuator grids are due to the wavefront sensor's extremely low sensitivity to this pattern. Once a big problem in standard minimum norm solutions, waffling was suppressed by imposing an a-priori restraint on the strength of these modes in the DM command vector². In tomography with laser guidestars, additional invisible modes associated with various combinations of unknown tip/tilts have been elaborated upon in the literature³.

Iterative algorithms that solve (3) and provide one of the (user selectable) unique solutions will take the form

$$\begin{aligned}
\mathbf{v}_{k+1} &= \mathbf{v}_k + \Delta \mathbf{v}_k \\
\Delta \mathbf{v}_k &= f(\mathbf{C} \mathbf{e}_k) \\
\mathbf{e}_k &= \mathbf{y} - (\mathbf{A} \mathbf{P} \mathbf{A}^T + \mathbf{N}) \mathbf{v}_k \\
\mathbf{x} &= \mathbf{P} \mathbf{A}^T \mathbf{v}_\infty
\end{aligned} \tag{4}$$

where \mathbf{P} , \mathbf{N} , and \mathbf{C} are positive-definite matrices. In the simplest form of the iteration, $f(\cdot)$ is simply multiplication by a scalar, chosen small enough to keep the iteration stable. In conjugate gradient iteration, the $f(\cdot)$ operation is a first order regression with varying coefficients, the coefficients chosen such that successive steps $\Delta \mathbf{v}_k$ are orthogonal to all previous steps⁴. Both these variations have exponential stability properties which make them attractive for real-time implementation. Given the strong stability properties the system will track time varying input data even when previous partially converged estimates are retained as the starting point when new data arrive. With fixed input, \mathbf{y} , the algorithm (4) converges to $\mathbf{e} \rightarrow 0$ and $\mathbf{x} \rightarrow \mathbf{P} \mathbf{A}^T (\mathbf{A} \mathbf{P} \mathbf{A}^T + \mathbf{N})^{-1} \mathbf{y}$.

If \mathbf{P} , \mathbf{N} , and \mathbf{C} are diagonal matrices, the entire algorithm (4) can be implemented in the spatial domain with massive parallelization over space. The \mathbf{A} operation is simply line integrals through the volume, which when distributed over space and over guidestars, takes only $O(N_l)$ time steps where N_l equals the total number of vertical layers to be summed. \mathbf{A}^T is even easier since the back-propagation can distribute data to all layers and guidestars simultaneously. In the conjugate gradient algorithm the $f(\cdot)$ operation is identical and independent for all the argument vector elements, and so can be distributed over space. However the coefficient determination in the conjugate gradient algorithm requires summing squares over space, an operation that takes $O(\log(m))$ time-steps if hierarchical sub-regions are summed, where m equals the total number of elements in the \mathbf{y} vector (the number of wavefront sensor data points).

We call \mathbf{P} and \mathbf{N} post-conditioners; they determine the nature of the solution (e.g. for a minimum variance solution they are the covariance of the volume delta-optical paths and the measurement noise, respectively). We call \mathbf{C} a pre-conditioner; the choice of \mathbf{C} determines only the convergence rate of the algorithm. Ideally, $\mathbf{C} = (\mathbf{A} \mathbf{P} \mathbf{A}^T + \mathbf{N})^{-1}$, in which case these algorithms converge in one step. Using a preconditioner matrix \mathbf{C} that approximates this inverse (without actually having to invert this gigantic matrix!) generally reduces the number of iterations required.

In order to approach a minimum variance solution, we cannot restrict ourselves to diagonal \mathbf{P} and \mathbf{N} matrices since we need to incorporate the fact that the delta optical paths and the noise (after slope to phase reconstruction) are spatially correlated. Implementing (4) in the Fourier domain gives us this option. If the statistics of \mathbf{x} and \mathbf{n} are spatially invariant then multiplication by covariance matrices \mathbf{P} and \mathbf{N} can be expressed as two dimensional convolutions, or, as implemented in the Fourier domain, componentwise multiplication. Massive parallelization is now distributed over the spatial frequency domain.

The forward-propagation operation $\mathbf{y} = \mathbf{A} \mathbf{x}$ in Fourier space is given by

$$\tilde{\mathbf{y}}_g(\boldsymbol{\kappa}) = \sum_l e^{i\boldsymbol{\kappa} \cdot \boldsymbol{\theta}_g h_l} \tilde{\mathbf{x}}(\boldsymbol{\kappa}, h_l) \tag{5}$$

where $\boldsymbol{\theta}_g$ is the angular location of guidestar g and h_l is the altitude of layer l . The back-propagation operation $\mathbf{w} = \mathbf{A}^T \mathbf{v}$ is given by

$$\tilde{\mathbf{w}}(\boldsymbol{\kappa}, h_l) = \sum_g e^{-i\boldsymbol{\kappa} \cdot \boldsymbol{\theta}_g h_l} \tilde{\mathbf{v}}_g(\boldsymbol{\kappa}) \tag{6}$$

The combined back-propagation, post-conditioning, and forward-propagation operation, $\mathbf{A} \mathbf{P} \mathbf{A}^T$, is written

$$\tilde{y}_g(\boldsymbol{\kappa}) = \sum_{l,g'} C_n^2(h_l) \boldsymbol{\kappa}^{-11/3} e^{i\boldsymbol{\kappa} \cdot (\boldsymbol{\theta}_g - \boldsymbol{\theta}_{g'}) h_l} \tilde{v}_{g'}(\boldsymbol{\kappa}) = \sum_{g'} \tilde{A}_{gg'}(\boldsymbol{\kappa}) \tilde{v}_{g'}(\boldsymbol{\kappa}) \quad (7)$$

where post-conditioning is manifested in the $C_n^2(h_l) \boldsymbol{\kappa}^{-11/3}$ factor, where $C_n^2(h_l)$ is the relative strength of the turbulence at layer l . This operation requires the multiplication by a small N_g by N_g matrix, where N_g is the number of guidestars, which is performed independently at each spatial frequency $\boldsymbol{\kappa}$. The preconditioner matrices, $\tilde{C}(\boldsymbol{\kappa})$, are similarly N_g by N_g matrices which can be easily and accurately precalculated and implemented in parallel across all spatial frequencies. A slightly more complex version of (7) is obtained when the finite altitude laser guidestars form cone beams through the volume⁵. To combine data from laser guidestar cone beams with natural guidestars' plane waves, it is necessary to resample in either the spatial or the frequency domain in the implementation of \mathbf{A} and \mathbf{A}^T . Resampling (using say cubic interpolation) is however also a massively parallelizable operation since the data processing is limited to local clusters of points.

We have described a Fourier domain representation of the iterative algorithm (4), which can be implemented with massive parallelization distributed over the spatial frequency domain and where individual processor units must perform N_g by N_g matrix multiplies for the general, and in particular, the minimum variance solution case.

An important issue to clarify at this point is that a pure Fourier domain approach does not take into account the fact that the wavefront sensor data is valid only on a finite aperture imposed by the telescope. It is not necessary to enforce $\mathbf{y} = 0$ outside the aperture, as a purely Fourier domain method would do, and imposing these extra fictitious constraints tends to generate artifacts in the volumetric estimate and in particular in the estimates of wavefronts in science directions. The remedy is to assure that our algorithm, although taking advantage of the Fourier domain for computational efficiency, actually behaves identical to its spatial domain counterpart. To do this, we transform the residuals, \mathbf{e}_k , back to the spatial domain, impose the aperture (a componentwise multiplication), then convert back to the frequency domain for the remaining steps in the iteration. This is done once each time around the iteration loop so it is important to implement the 2-D fast Fourier transform (FFT) efficiently and in a massively parallel manner. An architecture for doing this using 1-D FFT cores within a FPGA is currently under investigation. Basically, since the 2-D FFT is separable into two 1-D FFTs, each row or column in a 2-D array can be independently transformed. The switch from rows to columns, which in ordinary computer architectures would require transposing the data, is accomplished in the FPGA by placing the FFT cores physically along the diagonal of the 2-D array of data voxels. That way each FFT core is equidistant from its associated vector of data in both dimensions. The FFT within the iteration loop is applied only to data in the two dimensional sensor space, and once for each guidestar.

A summary block diagram of the massively parallelized tomographic engine is shown in Figure 3.

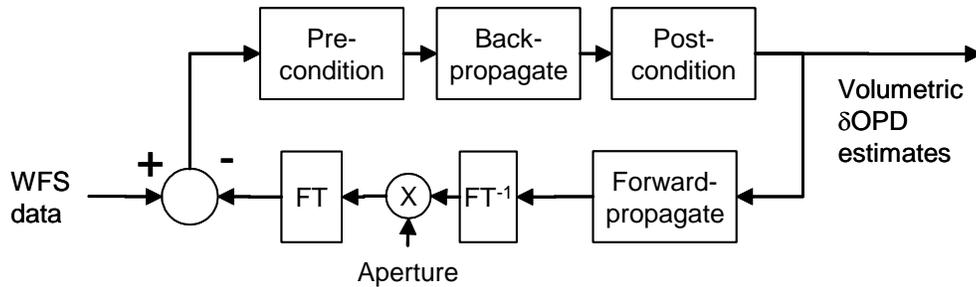


Figure 3. Block diagram of inverse tomography calculation unit.

2.4 Projection to deformable mirrors

Once the tomography solution is found, the volume estimate must be projected onto the deformable mirrors actually placed in the system.

In the MCAO case, an algorithm developed by Tokovinin, LeLouarn, and Sarazin⁶ determines the correction wavefronts that should be placed on a finite set of multi-conjugate deformable mirrors, given delta-optical path

estimates over the volume, so as to minimize the anisoplanatic effects over a given field of view. The algorithm is implemented in the Fourier domain with weighted line integrals through the volume, one per spatial frequency for each deformable mirror. Since we can parallelize over DMs and over the spatial frequency domain, the operation takes only $O(N_l)$ time steps, the same as only one forward projection.

In the MOAO case we must project through the estimated atmosphere along lines leading to the science targets. This is accomplished in the Fourier domain in a manner similar to that for propagating guidestars:

$$\tilde{y}_s(\boldsymbol{\kappa}) = \sum_l e^{i\boldsymbol{\kappa} \cdot \boldsymbol{\theta}_s h_l} \tilde{x}(\boldsymbol{\kappa}, h_l) \quad (8)$$

where $\boldsymbol{\theta}_s$ is the angular location of the science target. This operation, parallelized over science directions (i.e. DMs) and over spatial frequencies will also take only $O(N_l)$ time steps, the same as only one forward projection.

Since the projection to DM operation takes place outside the tomography iterations, the fraction of the overall time spent doing it should be insignificant.

2.5 Deformable mirror fitting

A deformable mirror's actuator commands a_j , combined with the actuator influence function $f(u)$, affect the shape of the DM surface according to

$$\phi(u) = \sum_j f(u - u_j) a_j \quad (9)$$

where u is position on the DM and u_j is the location of actuator j . The least-squares fit solution for actuator commands $\mathbf{a} = [a_j]$ is the solution to the set of linear equations

$$\mathbf{A}\mathbf{a} = \mathbf{b} \quad (10)$$

where

$$A_{ij} = \int w(u) f(u - u_i) f(u - u_j) du \quad \text{and} \quad b_i = \int w(u) f(u - u_i) \phi(u) du \quad (11)$$

and $w(u)$ is the pupil window function (=1 inside the aperture, =0 outside). Note that \mathbf{A} is a positive definite matrix, so we will be able to apply the conjugate gradient algorithm directly. With ϕ given on sample points, as provided by the tomography engine, the integrals above become sums. Calculating the residual in an iterative algorithm becomes

$$\mathbf{e}^{(k)} = \mathbf{b} - \mathbf{A}\mathbf{a}^{(k)} = \left[\int w(u) f(u - u_j) [\phi(u) - \phi^{(k)}(u)] du \right] \quad (12)$$

where

$$\phi^{(k)}(u) = \sum_j f(u - u_j) a_j^{(k)} \quad (13)$$

which are both simple convolutions, and so easy to implement in the Fourier domain. The conjugate gradient algorithm is applied to $\mathbf{e}^{(k)}$ to get the next updated estimate, $\mathbf{a}^{(k+1)}$. Since the bulk of the computations are tied up in the convolutions (12) and (13) we can implement the CG algorithm in the Fourier domain, except that the estimate $\phi^{(k)}$ must be transformed into the spatial domain at each iteration in order to multiply by the aperture then transformed back. The sums of squares needed to form the iteration-varying conjugate gradient coefficients can just as well be calculated in either domain; they are invariant due to Parseval's theorem.

All of the operations are therefore massively parallelizable, with order one clock cycle, plus the time to compute an FFT and inverse FFT, once per iteration. Simulations show promising convergence rates (described in Section 3) therefore we expect the DM fitting step to take a negligible fraction of the total AO processing time.

3. EXPERIMENT AND SIMULATION STUDY RESULTS

3.1 FPGA implementation

A small laboratory demonstration processor was constructed using a commercial FPGA development kit with a midrange FPGA chip. The processor implements the simple scalar feedback gain form of (4) with $\mathbf{P}=\text{Diag}(C_n^{-2}(h_l))$, $\mathbf{C}=\text{Diag}(1/N_g)$, $f(e_k) = \gamma e_k$, where γ is the loop gain, and with \mathbf{A} and \mathbf{A}^T implemented in the spatial domain. The demonstrator propagates three guidestars through a four layer atmosphere, with wavefronts sampled on four subapertures. A finite state machine, programmed within the gate array, controls the sequence of operations around the iteration loop: subtracting data from estimate to form residuals, pre-conditioning, back projection, and co-adding and post conditioning on the voxel grid.

The objective of the experiment was to determine how much clock time is required for one pass through the algorithm. Expanding the transverse size of the volume or the number of guidestars is accomplished by introducing more processors but will take no additional clock time. A very fast iteration time allows for a large number of iterations. Since the number of iterations needed to converge depends both on the spatial size of the problem and the particulars of the algorithm, we investigated the iteration convergence rate separately with computer simulations described in the next subsection. The FPGA timing tests however established the time per iteration, meeting our goal of <1 microsecond per iteration. This allows for 1000 iterations per millisecond, which is a canonical data measurement cycle time for near infrared science adaptive optics systems. Test results are summarized in Table 1.

Table 1. Timing results from FPGA demonstrator experiments

Element	#Clock cycles	Derived formula	Comment
Load Measured Value	12	3m	Done once
Forward Propagate	27	$N_g(2N_l + 1)$	
Compare	1	1	
Back Propagate	1	1	
Calculate New Estimate	7	$N_g + 4$	
Total (per iteration)	36		720 ns per iteration with a 50MHz clock

For a 7 layer atmosphere with 9 laser guidestars, the time per iteration scales to 3 microseconds. Refinement of the microcode to reduce lag times will enable a 150MHz clock rate, which is well within the range of the present FPGA chip's capability. This will achieve our 1 microsecond per iteration goal. In addition, our initial demonstrator was not parallelized over guidestars (a simplifying design choice for our initial implementation), hence the scaling by N_g in the forward propagate and calculate new estimate steps in Table 1. Parallelization over guidestars can be achieved at the expense of adding more registers but would reduce the iteration time to 420 microsec with a 50 MHz clock.

The FPGA chip in the development package was 20% utilized (2996 of 15360 available logic cells employed). Scaling to a system with 10,000 subapertures (such as for the 30 meter telescope) would require $(1 + N_l)N_{subaps} = 500$ of these chips. Assuming that circuit boards can carry 50 chips, this equates to 10 circuit boards. Parallelizing over guidestars would add another 370 chips.

3.2 Large spatial scale simulation

Computer simulations are necessary to determine number of iterations required to achieve convergence on a full size problem. We simulated a 7800 subaperture system (30 cm subaperture on a circular 30 meter telescope) with five guidestars and 7 layer atmosphere. Convergence is achieved when the line integral through the volume estimate differs in rms from the measured data by no more than the precision of the data. TMT AO error budgets assign on

the order of 40 nm rms for tomographic error in 4 micron rms wavefront conditions (tip/tilt removed). Thus a factor of 100 reduction in residual is required.

We ran the simple scalar feedback gain form of (4) with $\mathbf{P}=\text{Diag}(C_n^2(h_l))$, $\mathbf{C}=\text{Diag}(1/N_g)$, $f(e_k) = \gamma e_k$, where γ is the loop gain, and with \mathbf{A} and \mathbf{A}^T implemented in the spatial domain, the same as in the FPGA demonstrator. However, the computer simulation used 7800 subapertures per guidestar (30 cm subaperture on a circular 30 meter telescope), five guidestars, and a 7 layer atmosphere, making it a full sized TMT problem. This takes about 14 ms per iteration in IDL on a laptop pc. Convergence to 10^{-2} of initial rms was achieved in less than 700 iterations. The relative rms residual after 1000 iterations was 2×10^{-3} .

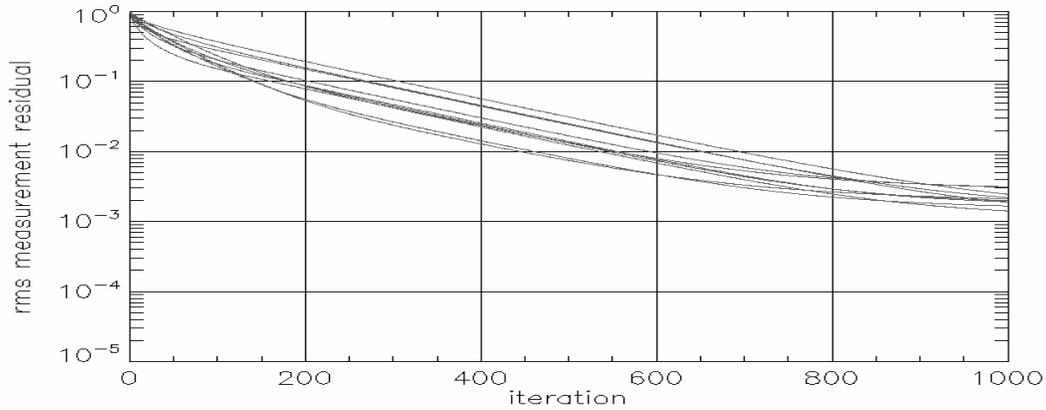


Figure 4. Convergence results from ten independent simulations of the iterative tomography algorithm applied to a 30 meter telescope AO system (7800 subapertures per guidestar, 5 guidestars, 7 layer atmosphere). Algorithm (4) was used with fixed feed back gain iteration and \mathbf{A} and \mathbf{A}^T implemented in the spatial domain. Initial atmospheric realizations were random with a Kolmogorov spatial power spectrum.

3.3 Simulation of the DM fitting procedure

We performed computer simulations of DM fitting procedure, (12)-(13), using non-preconditioned conjugate gradient, and assuming an influence function that is typical of today's deformable mirrors (about 15% interactor influence). A DM with 8797 actuators (7854 inside the aperture and 943 in guard bands) was fit to random Kolmogorov phase screens on a 30 meter aperture. Two tests were performed, one where the desired wavefront was sampled at the same spacing as the actuator grid, and one where the desired wavefront was super sampled by a factor of 3. With equal sampling, the algorithm consistently converged in 3 iterations or less. With super sampling, consistent convergence was attained in less than 20 iterations. Convergence is achieved when fitting error on the sample grid becomes less than the theoretical fitting error as set by the actuator spacing relative to r_0 , or approximately 70 nm rms for the $r_0=20$ cm (at 0.55 micron) cases we chose.

4. CONCLUSIONS

We have presented an alternative control architecture for multi-guidestar adaptive optics tomography based on massive parallelization of the real-time calculations. An FPGA based small-scale prototype implementation was constructed and evaluated in concert with computer simulations of the full size system. Extrapolations from our initial test results indicate that a massively parallel processor system with on the order of a few hundred FPGA chips (a few 10's of circuit boards) can serve a 7-9 guidestar AO system with 7800 subapertures per guidestar on the 30 meter telescope at 1kHz frame rate, even using the simple scalar feedback iteration and running it completely in the spatial domain (i.e. very little preconditioning).

In our ongoing work we are investigating how to implement the FFT (and sums of squares) efficiently within the massively parallel architecture so that a Fourier domain preconditioned conjugate algorithm can be implemented. It is yet to be determined if the extra computations per cycle this entails will be compensated with a corresponding

drop in the number of iterations needed to converge. The FFT appears to be important in the DM fitting procedure as well, but may not be as time-critical there.

ACKNOWLEDGEMENTS

This work was supported in part by the National Science Foundation Science and Technology Center for Adaptive Optics (CfAO), managed by the University of California at Santa Cruz under cooperative agreement No. AST-9876783, and in part by the Gordon and Betty Moore Foundation through its grant to the UCO/Lick Laboratory for Adaptive Optics.

This work incorporates much of what we have learned about MCAO tomography in the past three years through the CfAO sponsored Analysis, Modeling, and Simulation collaborative research efforts. The authors would like to thank in particular Drs Brent Ellerbroek and Luk Gilles of the TMT project office, and Prof. Curt Vogel of U. Montana, Bozeman (all CfAO participants) for the many helpful discussions about iterative algorithms and minimum variance tomography reconstructors during the CfAO Analysis and Modeling workshops.

REFERENCES

1. Poyneer, L., Gavel, D., and Brase, J., *Fast wave-front reconstruction in large adaptive optics systems with use of the Fourier transform*, JOSA-A, **19**, 10, pp2100-2111, 2002.
2. Gavel, D., *Suppressing anomalous localized waffle behavior in least-squares wavefront reconstructors*, SPIE, **4839**, pp. 972-980, 2003.
3. Lloyd-Hart, M., Milton, N., *Fundamental limits on isoplanatic correction with multiconjugate adaptive optics*, JOSA-A, **20**, 10, pp. 1949-1957, 2003.
4. Barrett, R., Berry, M., Chan, T. F., Demmel, J., Donato, J., Dongarra, J., Eijkhout, V., Pozo, R., Romine, C., and Van der Vorst, H., *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, 2nd ed. SIAM, Philadelphia, PA, 1994. http://www.netlib.org/linalg/html_templates/Templates.html.
5. Gavel, D., *Tomography for multiconjugate adaptive optics systems using laser guide stars*, SPIE, **5490**, pp. 1356-1373, 2004.
6. Tokovinin, A., Le Louarn, M., Sarazin, M., *Isoplanatism in a multiconjugate adaptive optics system*, JOSA-A, **17**, 10, pp1819-1827, 2000.