Adaptive Optics Reconstruction Utilizing Super-Sampled Deformable Mirror Influence Functions

David Redding, Scott Basinger, Gary Brack, Richard Dekany
Jet Propulsion Laboratory, California Institute of Technology

Ben Oppenheimer
California Institute of Technology

ABSTRACT

Control laws for an adaptive optics system for the Palomar Mountain Hale Telescope are described. These are derived using a linear matrix model of the optics, which gives the Hartman-sensor centroids and the science camera wavefront as functions of deformable mirror (DM) commands and atmospheric phase. The matrices defining this system can be directly measured from the optics, some at finer spatial resolution than the wavefront sensor can resolve. A minimum-wavefront compensator feeds back both the wavefront sensor centroids and previous DM commands, using the atmospheric covariance to smooth the response at the finer spatial scale. The DM command feedback provides direct observability of waffle, piston and other modes unobservable in the wavefront sensor. Compensator gains can be updated using Kalman filtering techniques to track the evolution of the atmospheric covariance matrix.

1.INTRODUCTION

This paper describes adaptive optics control laws for the Palomar Mountain Hale telescope. This system uses a Shack-Hartman sensor observing natural guide stars to drive a 241 active-actuator deformable mirror and a fast steering mirror.[1] First closed-loop results are expected in March 1998.

Shack-Hartman wavefront sensors (WFS) have been used with success in a number of adaptive optics (AO) systems around the world. They provide measurements of the wavefront slope sampled in discrete, contiguous "cells" across the pupil, which can be efficiently processed to estimate the pupil phase error and compute deformable mirror commands. They have a few weaknesses, though.

One problem is that certain modes of the mirror are either unobservable or weakly observable from the WFS. Mirror piston modes, in which every actuator is moved in the same direction and amount; and mirror waffle modes, in which every other actuator is moved in the opposite direction, are 2 examples. This problem can be reduced by using reconstructors that do not amplify unobservable modes, either by explicitly projecting them out of the reconstructor gains, or by using optimal control laws that naturally avoid actuating unobservable modes.[5,6]

The "minimum-wavefront" reconstructor approach used in this paper takes a more robust approach. It uses optimal control gains that do not actuate unobservable DM modes. More importantly, the minimum-wavefront reconstructor feeds back the DM commands, along with the WFS signals, to the estimator. The mirror modes that are not observable in the WFS are strongly observable in the DM commands, and are immediately eliminated. This amounts to using the DM itself as a second wavefront sensor.

A second problem is actuator "pinning," in which some DM actuators are displaced so far as to drive the corresponding WFS centroids out of one cell in into an adjacent cell, or into a spatial filter. In bad seeing conditions, it is possible that large regions of the DM will be locked into a checkerboard pattern through pinning. The minimum-wavefront reconstructor senses this condition through DM feedback.

A third weakness is the limited spatial resolution of the WFS, as compared to the resolution available from interferometry. This problem can be reduced by calibrating the DM influence matrices used to compute the reconstructor gains at a finer spatial scale than the WFS can resolve. The wavefront can then be estimated on a finer grid than the WFS cells. The contribution of the DM to each point is determined through calibrated influence functions, and that of the atmosphere, by interpolation from the WFS measurements,

Part of the SPIE Conference on Adaptive Optical System Technologies • Kona, Hawaii • March 1998
SPIE Vol. 3353 • 0277-786X/98/$10.00

as weighted by the atmospheric covariance matrix. The covariance matrix provides the spatial correlation of the atmospheric wavefront in a statistical form, parameterized by r_0. The influence functions are deterministic, time-invariant, and measurable (though not without error). The DM commands are generated to null the wavefront at the finer spatial scale, producing a response that optimally balances the atmospheric, WFS and DM contributions to the error at that scale.

The reconstructor gains are functions of seeing conditions, as captured by the atmospheric covariance, and of WFS and actuation noise. These factors change with time, are not strictly Kolmogorov, and are otherwise imperfect and unpredictable. The reconstructor errors can be used to improve the atmospheric covariance, however, so that the reconstructor gains will improve over time. This is done using a Kalman filter tracking the AO system outputs. The Kalman filter matches actual response to that predicted by the system linear model, and uses the residuals to update the atmospheric covariance matrix. New reconstructor gains are computed off-line by processing time-sequenced WFS and DM data (the full Kalman filter is too compute-intensive to be used at full scale in real time).

The discussion in this paper begins by defining the plant to be controlled, in the form of two linear matrix equations. Deformable mirror and other component characteristics are reviewed. Two distinct control approaches are then developed. The first is a "centroid nulling" least-squares controller, which seeks to minimize centroid errors and is implemented using a single matrix multiply. The second is the "minimum-wavefront" compensator, which is a fixed-gain Kalman filter feeding back both the WFS measurements and the DM commands. Its implementation is more complex, as it requires 2 matrix multiplies and a vector add. The excess processing, however, takes place during dead time, in preparation for the next computing cycle, and does not add greatly to the required processor throughput capability.

The compensator gains can be updated off-line using the full Kalman filtering update equations processing the telemetry stream. The updated gains are then be linked back in to the real-time processor.

Use of the optimal gains without the DM feedback is also considered. A differential form of the minimum-wavefront reconstructor, which does not use DM feedback, is compared to the full compensator in time simulations.

Examples are worked through, showing that the minimum-wavefront control is the superior performer. Driving the centroids to null does not necessarily produce minimum wavefront error; conversely, the minimum-wavefront solution often leaves small non-null centroids. Evolution of the reconstructor gains is discussed, and the paper concludes with a review of the implementation of the controller in hardware.

2.AO SYSTEM REPRESENTATION

This section introduces a linear matrix model for Hartman-sensor based adaptive optics systems. The model includes separate equations for the WFS centroids and the exit pupil wavefront. The wavefront equation uses a finer sampling of the pupil then does the centroid equation.

The Palomar AO system uses the Hale 5-meter telescope to feed an off-axis parabola, which reimages the primary mirror onto the deformable mirror (Fig. 1). A second off-axis parabola refocusses the beam, which is then split to send light into both a science imaging camera and the WFS. Within the WFS, a lens creates a second pupil image at a lenslet array. The resulting spot array is focussed on a CCD array, with each spot nominally aligned at the center of a 4-by-4 pixel subarray. A FSM is located between the DM and the telescope to implement tilt corrections. A spatial filter limits the field of view of each centroid to 1 WFS cell.

The Palomar pupil is mapped onto a 16-by-16 cell Shack-Hartman WFS, with the corners of the WFS aligned approximately with the DM actuators. The DM is a 349-actuator Xinetics model, with 241 actuators individually driven. The pupil is illustrated in Fig. 2, as is a 7 x 7 subsample of the pupil illustrating supersampling of the subapertures. In the pupil diagram, the actuators are indicated by black dots; the WFS cells are formed with the actuators at each corner. In the subsample, the actuators are indicated by hollow dots; The smaller solid dots are points where the wavefront is evaluated. These over-sample the pupil compared to the WFS cells, allowing resolution of higher spatial frequency effects than can be seen in the WFS. Spacing of the actuators is 7 mm at the DM; on the PM-scale pupil it is 31.25 cm.

The atmospheric state equation governing the transition of the atmospheric phase from one time step (denoted by a subscript i) to the next is given by:

$$a_{i+1} = a_i + \lambda_i \tag{1}$$

The a_i states are the phases of the wavefront at the entrance pupil of the telescope, as sketched in Fig. 1, at time step i, sampled at the small solid dots indicated in Fig. 2. The mean of a_i is zero, and the covariance at a given instant of time is A_i. A_i can reasonably be generated assuming Kolmogorov turbulence, as a function of r_0. As shown in Fig. 2, the a_i oversample the WFS cells by a factor of m = 4. The dimension of a_i is $(mn+1)^2$ or 625 for the subsample shown. For the Palomar example, we use m=2, n=16, so the dimension of a_i is 1089. The matrix dc/da thus has dimension 241-by-1089.

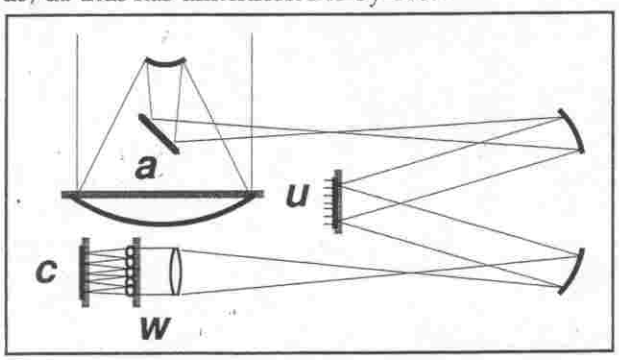

Figure 1. Optics layout for Palomar AO.

The time evolution of the atmosphere is here treated as a noise process, driven by the vector λ_i with covariance Λ_i.

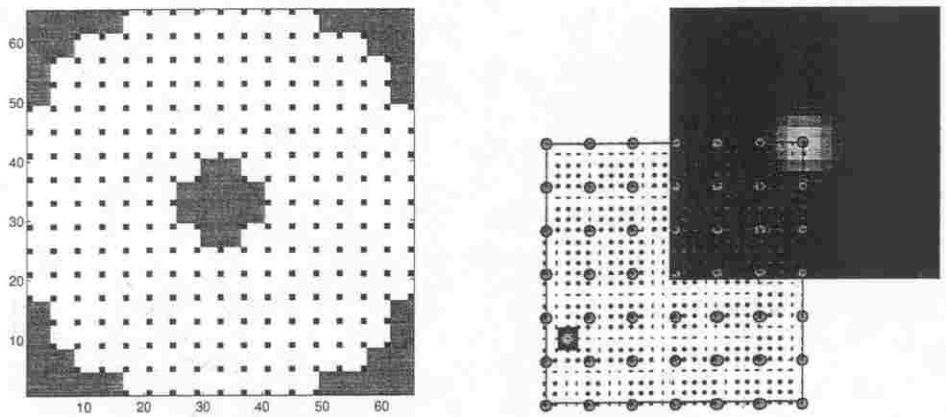

Figure 2. Pupil diagram showing actuator placement and WFS cells; second plot shows DM influence function and 4× oversampled wavefront state locations.

The measurement equation captures the influence of the DM and atmospheric phase errors on the Hartman centroids. At the instant of the centroid measurement, at the start of a new control cycle:

$$c_i = \frac{dc}{da} a_i + \frac{dc}{du}(u_{i-1} + q_{i-1}) + r_i \tag{2}$$

Here c_i is a vector of dimension 512 containing the x- and y-centroids for each of the 256 subapertures, at time t_i. The DM actuator commands that were applied in the previous time step are in the u_{i-1} vector, which is of dimension 241. The q_{i-1} vector contains actuation noise terms, with covariance Q_i. Actuation noise captures gain calibration errors.

The sensitivity of the centroids to DM actuators – the matrix dc/du – can be measured easily by pushing actuators and recording the centroid signals that result, then dividing by the amount pushed. With idealized actuators, the result of pushing a single actuator is to shift the centroid of the four immediately

adjacent Hartman cells directly away from the actuator, at 45° to the grid layout. Actual hardware produces more complication, coupling the response over 36 or more adjacent WFS cells, causing actuator scale factor and alignment problems, etc.

The r_i vector contains 512 centroid measurement error elements, which are driven by detection noise as processed by the centroider two-stage filter. These are assumed to be zero-mean and uncorrelated, with a standard deviation set by sensor and filter characteristics determined chiefly by subaperture SNR.

Equation 3 models the influence of the DM and atmosphere on the WFS and science camera wavefront, represented by a vector w_i. It is the objective of the AO system to drive this to zero. The wavefront equation takes the form:

$$w_i = \frac{dw}{da} a_i + \frac{dw}{du}(u_{i-1} + q_{i-1}) \tag{3}$$

The wavefront vector w_i corresponds to the differential optical pathlengths of rays originating at the entrance pupil at the locations of the elements of a_i, (the small solid dots in Fig. 2) and then traced to the exit pupil located in the WFS or science camera. It has the same dimension of $(mn+1)^2$-by-1. The sensitivity of the wavefront with respect to the atmospheric states, dw/da, is nominally an identity matrix of size $(mn+1)^2$. Deviations occur due to imperfections in the AO and telescope system optics, but these can be calibrated.

The sensitivity of the wavefront to DM actuations, given in matrix dw/du, records the influence function of each actuator at the m-times higher spatial scale of a_i and w_i. It is measured directly, using an interferometer looking back through the AO system at a point source located at the paraxial focus of the telescope. The DM actuators are pushed in turn, the WF recorded and normalized as before.

Influence functions have been measured for the Xinetics 349-actuator DM with 7 mm actuator pitch.[3] For our model we are using an "average" actuator. As shown in Fig. 3, each actuator produces a significant response for a radius of about 2 actuator spacings, pushing up and down by amounts that decrease with increasing distance from the actuator. There is significant curvature to the response that is not resolved at the WFS cell scale, but is captured by the m-times finer sampling provided by w_i.

The validity of this model depends on accurate measurement of the various component matrices. Other effects such as anisoplanatism and scintillation are not modeled. Nonlinearity effects, model range-of-validity considerations (with respect to changes in the optics in particular) and other factors are not addressed here. These effects reduce absolute performance for the approaches considered here, but are not expected to change the relative performance.

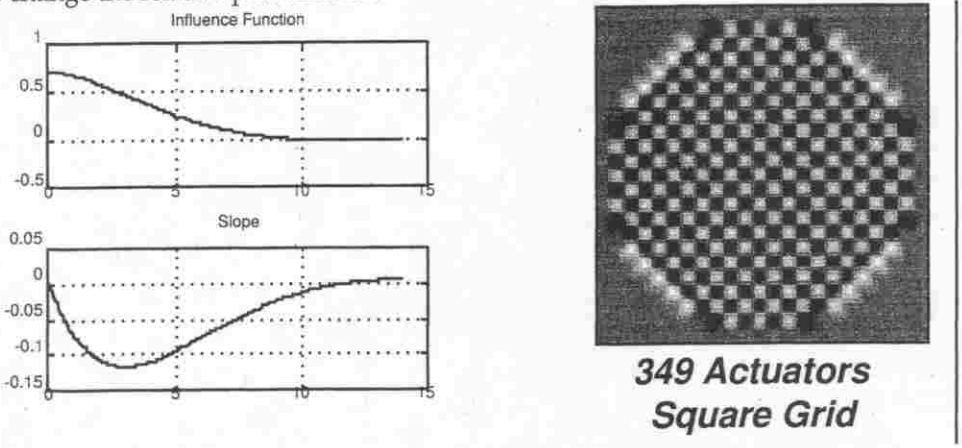

Figure 3. DM influence function; DM waffle mode.

3.AO COMPENSATOR OVERVIEW

The heart of the AO system is its controller. For Palomar, this is implemented in 3 stages, as sketched in Fig. 4. The first stage processes the raw CCD outputs to determine Hartman cell centroids. This "centroider" uses a 2-stage filter, with gains updated according to guide-star and seeing characteristics, to approach quad-cell accuracy in a 4-by-4-pixel format. The centroiding is described in detail in Ref. 2.

The second stage of the controller takes the centroids and determines commands for the deformable mirror (DM) and fast-steering mirror (FSM). This is the AO compensator or "reconstructor," that is the main topic of this paper. In the next sections we describe two basic reconstructor approaches: a centroid-nulling compensator; and the minimum-wavefront optimal reconstructor introduced above.

The final stage is a digital filter, used to temporally smooth inputs to the DM and to improve low-frequency accuracy. The entire AO loop runs at 500 Hz, to achieve a closed-loop bandwidth approaching 60 Hz. Controller gains are updated, following calibration of the optics, and during observation runs, based on wavefront sensor measurements. The new gains are computed at appropriate intervals by "matrix generating software" (MGS), which runs off-line in an engineering workstation-class computer.

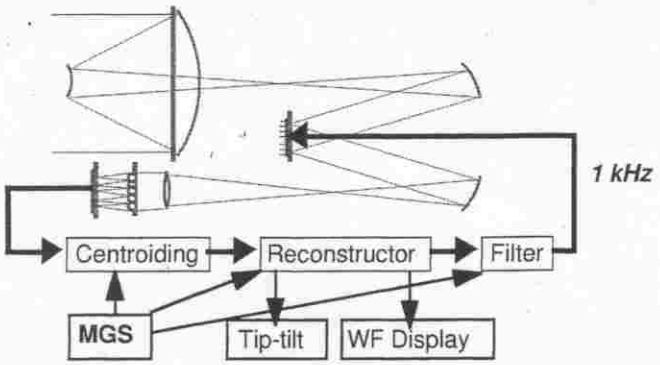

Figure 4. Controller signal flow.

4.CENTROID-NULLING COMPENSATOR

Several systems have been proposed or implemented with compensators that seek to minimize centroid error. For sake of comparison with the minimum-wavefront approach, we formulate a centroid controller here. This discussion is simplified and not intended to duplicate any particular group's approach.

The assumption underlying the centroid controller is that the centroids c_i represent the gradient of w_i in the WFS, and that driving c_i to zero will also flatten the wavefront w_i. This approach is fundamentally limited to WFS resolution limits in its response: knowledge of the influence functions dw/du at the finer scale are not used. No attempt is made to separate the part of w_i due to a_i from that due to u_i. The centroid controller does not explicitly utilize prior knowledge of the wavefront. It does not exploit knowledge of the statistics of the atmosphere or of the measurement and actuation errors.

The centroid control problem is to minimize:

$$J_{cent} = \frac{1}{2} c_i^T c_i \tag{4}$$

subject to the constraint of Eq. 2. The solution is straight-forward:

$$u_i = -\left(\frac{dc}{du}\right)^+ c_i = -G_c c_i \tag{5}$$

Here the gain matrix G_c is of dimension nu-by-nc. The centroid and wavefront residuals following a single control step are:

$$\delta c_i = c_i + \frac{dc}{du}(u_i + q_i) \tag{6}$$

$$\delta w_i = w_i + \frac{dw}{du}(u_i + q_i) \tag{7}$$

5.MINIMUM-WAVEFRONT COMPENSATOR

The minimum-wavefront reconstructor differs from the centroid-nulling controller of the previous section in several respects. It is explicitly designed to minimize the WFS exit pupil wavefront error. It utilizes the finer resolution of the measured dw/du influence functions to extend the resolution of the DM-related component of WFE beyond the WFS resolution limit. It makes use of prior knowledge of the wavefront, as well as knowledge of the statistics of the atmosphere and of the measurement and actuation errors to optimally filter the measurements in applying the control. This increased versatility results in improved performance. It comes at the cost of a more complex implementation in the real-time computer, however.

Conceptually, the wavefront controller operates in two stages at each time step. The first stage explicitly estimates the wavefront a_i at the entrance pupil of the telescope, given the current centroid measurements c_i, the previous DM actuation commands u_{i-1}, and the previous estimate $\bar{a}_{i-1}$ which is now about 1 msec old. The second stage takes the estimate $\bar{a}_i$ and computes the actuator commands u_i. Periodic updates of the covariance of the atmospheric states are performed to keep the gains current.

Actual implementation differs from this picture slightly. There is no explicit evaluation of $\bar{a}_i$ in the on-line processing (though there is in the off-line processing). The computation of the control is split so that part is done in preparation for the next time step, rather than all being done after the receipt of the centroid data. The latter significantly speeds computation of the control, so that the time from receipt of data to actuation of the DM is commensurate with the time for the centroid-nulling controller.

It is convenient to rewrite the measurement equation to lump together the noise terms, namely the measurement noise r and the actuation noise q_i (which represents the uncertainty in the knowledge of the DM state provided by u_i):

$$c_i = \frac{dc}{da}a_i + \frac{dc}{du}u_{i-1} + r_{ci} \tag{8}$$

The combined noise terms r_c have covariance R_c, where:

$$r_{ci} = \frac{dc}{du}q_i + r_i \tag{9}$$

$$R_c = \frac{dc}{du}Q\left(\frac{dc}{du}\right)^T + R \tag{10}$$

The estimation problem is to determine a_i, given centroid measurements c_i corrupted by noise r_i with covariance R, and given known previous actuator commands u_{i-1} corrupted by noise q_{i-1} with covariance Q. The estimated covariance for the atmospheric states a_i is A, which is assumed to be defined initially by the Kolmogorov structure function, and which can updated based on off-line computations.

A reasonable estimate is provided by the minimum-variance estimator, which seeks to balance the uncertainty in the estimate across the sources of that uncertainty. The cost function J_{est} is minimized:

$$J_{est} = \frac{1}{2}\left[\bar{a}_i^T A^{-1}\bar{a}_i + \left(c_i - \frac{dc}{da}\bar{a}_i - \frac{dc}{du}u_{i-1}\right)^T R^{-1}\left(c_i - \frac{dc}{da}\bar{a}_i - \frac{dc}{du}u_{i-1}\right)\right] \tag{11}$$

J_{est} has a minimum at:

$$\delta J_{est} = 0 = \delta\bar{a}_i^T\left[A^{-1}\bar{a}_i + \left(\frac{dc}{da}\right)^T R^{-1}\left(c_i - \frac{dc}{da}\bar{a}_i - \frac{dc}{du}u_{i-1}\right)\right] \tag{12}$$

The nontrivial solution for the estimate is:

$$\bar{a}_i = K\left(c_i - \frac{dc}{du}u_{i-1}\right) \tag{13}$$

The matrix K is the estimator gain matrix:

$$K = E^{-1}\left(\frac{dc}{da}\right)^T R^{-1} \tag{14}$$

Here E is the covariance of the estimate error residuals:

$$E^{-1} = A^{-1} + \left(\frac{dc}{da}\right)^T R^{-1}\frac{dc}{da} \tag{15}$$

The dimension of K is na-by-nc, where na grows with the oversampling factor m.

The controller phase of the compensator takes the estimated wavefront and generates commands u_i. The control objective function, to be minimized, is J_{con}:

$$J_{con} = \frac{1}{2} w_i^T w_i \tag{16}$$

The resulting control is computed as:

$$u_i = -\left(\frac{dw}{du}\right)^{+}\frac{dw}{da}\bar{a}_i = -G\bar{a}_i \tag{17}$$

The controller gain matrix G has dimension nu-by-na, where na is increased by the oversampling factor. The total compensator, created by using the estimates $\bar{a}_i$ to drive the control computation, is:

$$u_i = -G_c c_i - G_u u_{i-1} \tag{18}$$

Gain matrices G_c and G_u are:

$$G_c = GK \tag{19}$$

$$G_u = GK\frac{dc}{du} \tag{20}$$

The dimension of G_c is nu-by-nc and that of G_u is nu-by-nu. Neither matrix size is increased by the oversampling factor m. The second product of Eq. 18 can be pre-computed following determination of u_{i-1}, during the time when the processor is normally waiting for the next WFS frame. The computational performance required of the processor is nearly the same as that for the centroid-nulling controller (Eq. 5).

6.KALMAN FILTER FOR UPDATING GAINS

A more elaborate (and computationally intensive) approach to the estimation part of the minimum wavefront compensator is provided by the well-known Kalman filter.[4] This approach extends the minimum-variance estimator just described to update not just the state a_i, but also its covariance A and the estimator gains K. It provides a means to use the measured centroids to improve knowledge of the atmosphere and to tune the response of the compensator accordingly. It is too computationally intensive to implement in the real-time AO loop with current hardware, but may be run off-line as part of the MGS to provide periodic updates to the K matrix. A sequence of centroids and DM actuator commands is recorded and then played back through the filter equations to better estimate the evolving a_i and A. The updated A is then uploaded into the real-time compensator.

The evolution of a_i between updates is given by Eq. 1. For the Kalman filter, the estimate of a_i is now:

$$\bar{a}_i = \bar{a}_{i-1} + K_i\left(c_i - \frac{dc}{da}\bar{a}_i - \frac{dc}{du}u_i\right) \tag{21}$$

The compensator uses gains K_i, producing estimates with error E_i:

$$K_i = E_i\left(\frac{dc}{da}\right)^T R_c^{-1} \tag{22}$$

$$E_i = \left[A_i^{-1} + \left(\frac{dc}{da}\right)^T R_c^{-1}\frac{dc}{da}\right]^{-1} \tag{23}$$

A new value for A is thus:

$$A_{i+1} = E_i + \Lambda \tag{24}$$

Equations 21-24 are evaluated for a recorded sequence of observations made by the real-time system. After some number of iterations, the covariance can be expected to stabilize, so that changes with each iteration are small. This updated covariance matrix is used to update the real-time gains via Eqs. 14 and 15.

Other algorithms for updating the gains could provide for interesting experiments. The spatial scale of the problem could be split. The full Kalman filter can be run for a small number of subapertures, with the correlations generalized to update the high spatial-frequency part of A. Similarly, the full filter could be run at a large spatial scale, by clustering groups of subapertures to form a coarser grid. Better post-processing algorithms exist, such as optimal smoothing, which solves the problem backwards and forwards to provide better estimates of A.[4]

7.DIFFERENTIAL FORM

It is possible to use the optimal gains without the DM feedback – indeed, this is the usual approach.[6] The control for this is simply:

$$\Delta u_i = -G_c c_i \tag{25}$$

The control Δu_i is implemented in a differential sense, in that it is added to the previous control without knowledge of the previous control value. The full minimum-wavefront control of Eq. 18 is an absolute form, feeding back total DM displacement. The differential form has the potential to run away in poorly observed modes; the absolute form has no unobserved DM modes.

More elaborate versions of Eq. 25, which explicitly deal with the uncertainty induced by not knowing the total DM displacement, have been derived.[7] Preliminary results show that they do not differ significantly from the G_c matrix derived for the minimum-wavefront controller.

8.SIMULATION RESULTS

Comparison of the performance of the centroid-nulling and wavefront controllers reveals some interesting differences. Figure 5 shows 3 typical cases simulated at 3 values of r_0. The top third of each frame shows the wavefront w prior to actuating the DM; these random realizations of the atmosphere were generated, using Kolmogorov statistics driven by r_0. WFS read noise and DM actuation noise effects were also included at a low level. The middle third of each frame shows the wavefront compensated by a single full step of the centroid controller (no servo). The bottom third shows the wavefront compensated by the minimum-wavefront control, again in a single full step.

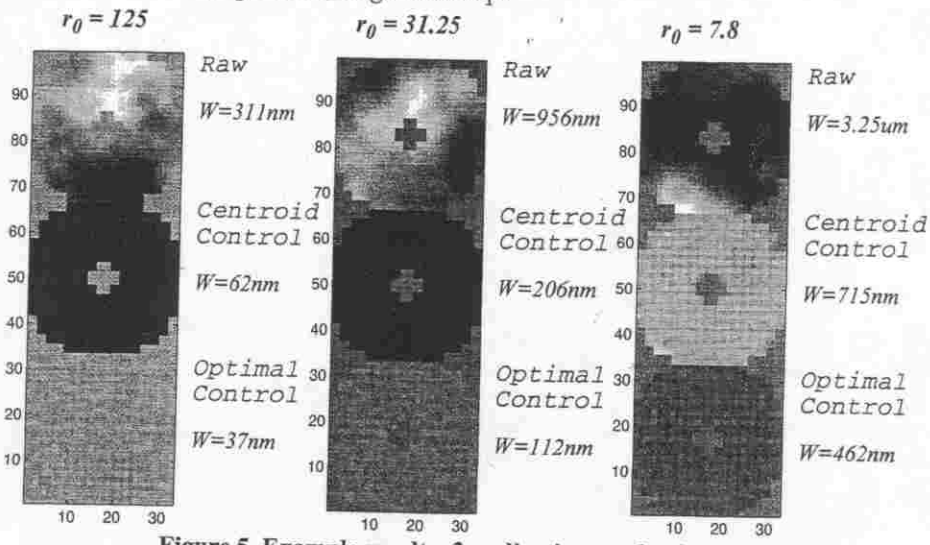

Figure 5. Example results: 3 realizations at 3 values of r_0.

In all cases, the minimum-wavefront control produces lower wavefront error than the centroid-nulling control. The centroid control shows significant piston compared to the minimum-wavefront control; this

piston mode is poorly observed by the WFS. The minimum-wavefront control sees it directly in the DM feedback term of Eq. 18. The residual centroids of the centroid-nulling control (not shown) are significantly smaller than those of the minimum-wavefront control.

Results accumulated using Monte Carlo simulation for the same system show some interesting trends. As seen in Fig. 6, the minimum-wavefront controller outperforms the centroid-nulling controller in WFE, but does so with a significant centroid offset. The optimal control does not occur at a zero-centroid condition. The DM feedback term in Eq. 18 provides an offset point to which the centroids are driven.

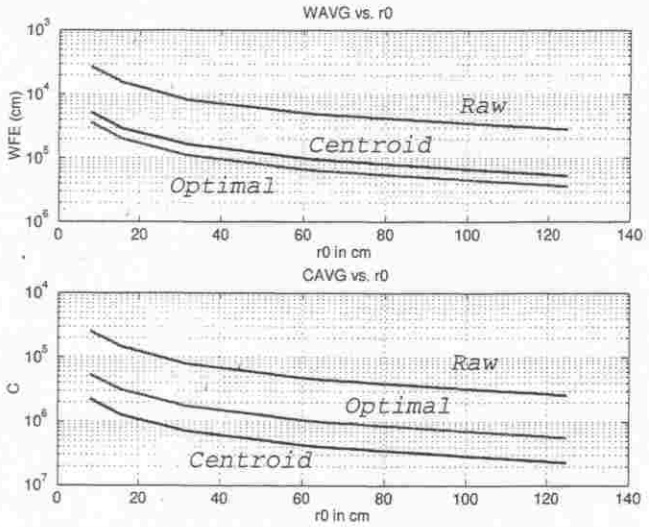

Figure 6. Example results: WF and centroid residuals vs. r_0, for 100 Monte Carlo trials/datapoint.

Figures 7 and 8 compare the time response of the minimum-wavefront controller to that of the differential optimal control described in the preceding section. These results were computed using a simple low-pass filter applied to the to smooth the actuator time response. The atmosphere was calculated at an r_0. of 7.8 cm. It is shifted by half a WFS cell per time step to simulate a moving turbulence layer. Actuator and read noise were included. The raw and compensated wavefronts are shown at 4 times in Fig. 7. The top third of each frame shows the uncontrolled wavefront in the entrance pupil. The middle third of each frame shows the wavefront compensated using the minimum-wavefront control. The bottom third shows the wavefront compensated by the differential optimal controller. As was true for the centroid controller, the differential control is not effective in compensating the poorly observed piston mode.

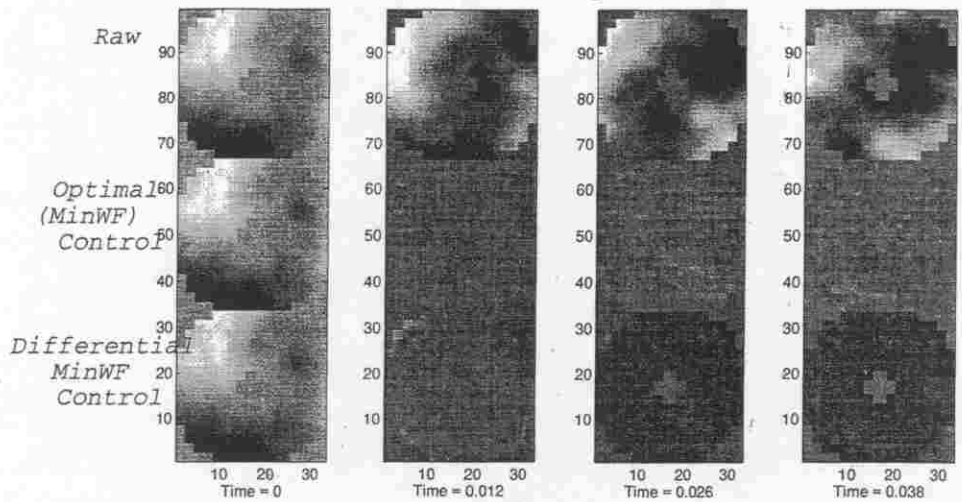

Figure 7. Time sequence.

Performance time histories are plotted in Fig. 8. The absolute form of the minimum-wavefront controller performs significantly better than the differential, though with a finite centroid offset. The difference in the performance is the DM feedback term, which determines the centroid offsets, and which provides the improvement in WF performance.

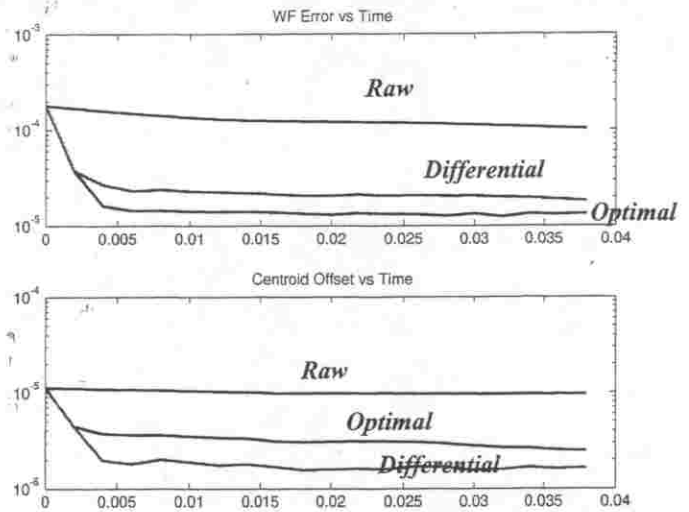

Figure 8. Performance during time response example.

9.CONCLUSION

Optimal reconstructors have been used successfully in the past. The minimum-wavefront compensator presented here differs from previous controllers by including DM actuator displacements as well as WFS signals in the feedback loop. The DM terms provide direct observability of DM modes that are not observable in the WFS signals, eliminating susceptibility to waffle and other problem modes. The DM influence matrices are calibrated at greater resolution than the WFS provides, allowing reconstruction on a finer grid. The atmospheric effects are interpolated onto this grid using the atmospheric covariance matrix. This matrix can be updated during a run using Kalman filtering techniques in an off-line computer, allowing near real-time gain updating. The minimum-wavefront controller is being implemented on the Palomar Mountain Hale 5-meter telescope.

10.ACKNOWLEDGEMENT

This paper benefitted from useful discussions with Mark Milman and Mike Shao. This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with NASA.

11.REFERENCES

1. R. Dekany, "The Palomar Adaptive Optics System," *OSA Technical Digest*, pp. 40-42 (1996).
2. S. Shaklan, "Two-Stage Centroiding Algorithm for Palomar Adaptive Optics," JPL Internal Memorandum, January 23, 1995.
3. B. Oppenheimer, D. Palmer, R. Dekany, A. Sivaramakrishnan, M. Ealey and T. Price, "Investigating a Xinetics Deformable Mirror," SPIE 3126 (1997).
4. A.E. Bryson and Y.C. Ho, *Applied Optimal Control*, Hemisphere Press, 1975.
5. W. Wild, E.J. Kibblewhite and R. Vuilleumier, "Sparse Matrix Wave-front Estimators for Adaptive Optics Systems for Ground-Based Telescopes," Optics Letters, Vol. 20, No. 9, May 1995.
6. W.Wild, "Innovative Wavefront Estimators for Zonal Adaptive Optics," SPIE 3126, pp. 278-287 (1997).
7. D. Redding, "Adaptive Optics Reconstructor Design and Analysis," JPL Internal Memorandum, April 1994.