

Characterization of adaptive optics at Keck Observatory: part II

Marcos A. van Dam,^a David Le Mignant^b and Bruce A. Macintosh^a

^aLawrence Livermore National Laboratory, P. O. Box 808, Livermore, CA 94550

^bW. M. Keck Observatory, 65-1120 Mamalahoa Highway, Kamuela, HI 96743

ABSTRACT

This paper is a continuation of the characterization of adaptive optics (AO) at Keck Observatory (SPIE 5169-01). The bandwidth and measurement noise error terms are often the important sources of wave-front error. Here, we show how the magnitude of these two terms is estimated. First, the Bayesian wave-front reconstructor employed at Keck Observatory is presented and shown to perform better than a conventional SVD reconstructor. A novel technique used to estimate the size of the spot on a Shack-Hartmann wave-front sensor quad cell detector is introduced, along with experimental results using this technique. The spot size is an essential component of the dynamic model of the AO system, which is presented and used to find the bandwidth and noise error terms.

Keywords: Adaptive optics, wavefront sensing, Shack-Hartmann, Keck Observatory

1. INTRODUCTION

Adaptive optics (AO) has found widespread use in astronomical settings to compensate for atmospheric turbulence and telescope aberrations.^{1,2} At the W. M. Keck Observatory, there are identical AO systems on the Keck I and Keck II telescopes.³ In previous work, we described the components and the performance of the Keck AO system.^{4,5} This paper is a continuation of this characterization effort.

In this paper, we focus on the problem of estimating the bandwidth and measurement noise errors. This is important because during an AO observation these are the only two terms that can be optimized to improve the performance. In order to estimate these two terms, a good understanding of the dynamic behavior of the control loop is needed. A good article on how to model an AO system was written by Madec.⁶ Previous papers that model the control loop assume that the controller is modal:⁷⁻⁹ however, many AO systems, including those at Palomar and Lick Observatories, use zonal reconstructors. The Keck AO system uses a Bayesian zonal reconstructor described in Sect. 2. In addition, most previous analyses assume that the centroid of the spot on the Shack-Hartmann wave-front sensor yields direct estimates of the slope of the wave-front.⁷ However, the Shack-Hartmann WFS most commonly used in astronomical AO systems uses a 2×2 array of pixels, known as a quad cell, to measure the centroid. This centroid estimate is inversely proportional to the spot size on the WFS, which must also be estimated. The effect of the wave-front sensing spot size on the control loop is usually ignored. A new method to measure the spot size using slope discrepancy is introduced in Sect. 3. A model of the dynamic behavior of the AO system including turbulence rejection and the noise propagation is developed in Sect. 4 and is used in Sect. 5 to estimate the bandwidth and measurement noise errors.

Further author information: send correspondence to M.v.D, mvandam@keck.hawaii.edu




Figure 1. Images using an SVD (left) and a Bayesian (right) reconstructor. Note the presence of waffle, characterized by four symmetric dots on the corners, when the SVD reconstructor is used.

2. WAVE-FRONT RECONSTRUCTOR

Until 2003, the reconstructor at Keck Observatory was obtained via a matrix inversion of the influence matrix, H , using an SVD algorithm.¹⁰ This has been replaced by a Bayesian reconstructor¹¹ that uses the covariances of Kolmogorov turbulence, C_ϕ , and the relative noise in each subaperture, W , as prior information. There is a noise-to-signal parameter, α , that is adjusted depending on the brightness of the guide star. The reconstructor, R , is given by

$$R = (H^T W^{-1} H + \alpha C_\phi^+ + \eta \mathbf{1}^T \mathbf{1})^{-1} H^T W^{-1}, \quad (1)$$

where $\mathbf{1} = [1, 1, 1, \dots]^T$. Since C_ϕ is the piston-removed phase covariance, obtained using Wallner’s method,¹² the piston eigenmode has a zero eigenvalue. Consequently, C_ϕ is singular and the generalized inverse is needed. As a consequence, C_ϕ^+ does not penalize the piston term at all, and the term $\eta \mathbf{1}^T \mathbf{1}$ is needed to penalize the piston mode and improve the conditioning of the inversion. The value of η is not crucial since any remaining piston term is removed later.

In practice, the implementation of the algorithm is slightly more complex than Eq. (1) indicates. First, the actuators that are not optically conjugate to the corner of a subaperture are “slaved” to the average value of the neighboring unslaved actuators. Second, R is pre-multiplied by a matrix that removes the average x - and y -centroids. Third, the resulting matrix is post-multiplied by a matrix that removes the piston, tip and tilt modes in actuator space.

Equation (1) is the optimal reconstructor in the open-loop case; simulations and extensive sky testing have demonstrated that it also performs very well in closed loop. The reason why it works well is that the inversion of the system matrix is regularized, so that noise in the centroids leads to smaller actuator motions. The new reconstructor has resulted in an increase in the Strehl ratio corresponding to the elimination of 100 nm of wave-front error (in quadrature) for bright guide stars. The most salient difference is the elimination of the four spots in a square pattern in the image, as can be seen in Fig. 1. This is indicative of unsensed waffle on the DM,¹³ which is a problem for Shack-Hartmann wave-front sensors that have the actuators optically conjugate to the corners of the subapertures in what is commonly known as the Fried configuration.¹⁴ The magnitude of the improvement increases as the guide stars get fainter, because α is increased to further suppress the noise at the expense of suppressing some signal.

It is helpful to think of the reconstructor in terms of its eigenvalues and eigenvectors.¹³ Modal reconstructors regularize the reconstruction process by reducing the gain of the modes with low signal-to-noise ratio. Nevertheless, the modes themselves do not change. As a result, the matrix product of the influence matrix with the reconstructor matrix is a diagonal matrix. Regularized least-squares solutions, on the other hand, change the

eigenmodes so that only a few modes contain the bulk of the high spatial frequencies.¹³ As a consequence, there is “cross-talk” between the modes and the analysis of the control loop becomes more difficult.

3. SPOT SIZE ESTIMATION USING SLOPE DISCREPANCY

In this section, we briefly summarize the reasons why an estimate of the size of the spot on the WFS is needed and a novel technique implemented at Keck Observatory. An extensive discussion on the importance of knowing the spot size can be found in the work of Véran and Herriot,¹⁵ while a more thorough presentation of this method is to appear elsewhere.¹⁶ The reader is referred to an excellent paper by Tyler for a more in-depth study of slope discrepancy.¹⁷

The slope discrepancy, \mathbf{s}_\perp , consists of the component of the centroids, \mathbf{s} , that lies in the nullspace of R . The nullspace of R consists of all vectors \mathbf{s} such that $R\mathbf{s} = 0$.¹⁸ The symbol \perp will be used to denote the null space of R .

At Keck Observatory, matrix R has dimensions 349×480 . Depending of the angle of the telescope pupil⁵ on the wave-front sensor, about 50 of the actuators are slaved, leaving about $480 - (349 - 50) = 181$ linearly independent rows. R can be written as a singular value decomposition:

$$R = USV^T, \quad (2)$$

where U is a 349×349 matrix, S is a 349×480 diagonal matrix with the 349 singular values as its diagonal entries and V is a 480×480 matrix. The columns of V , corresponding to the (approximately 181) zero-valued singular values of S , comprise an orthonormal basis, Z , for the null-space of R . It follows that

$$\mathbf{s}_\perp = ZZ^T \mathbf{s}. \quad (3)$$

3.1. Measuring the spot size using slope discrepancy

In this subsection, the process for estimating the spot size using slope discrepancy is described. The relationship between the centroid measurement, \mathbf{s} , the raw centroids, \mathbf{r} and the reference centroids, \mathbf{c} , is:

$$\mathbf{s} = \mathbf{r} - \mathbf{c}. \quad (4)$$

It follows that $\mathbf{s}_\perp = \mathbf{r}_\perp - \mathbf{c}_\perp$. The reference centroids are comprised of aberrations introduced by the lenslet array, the imaging leg and the beamsplitter that splits the light to the two legs.

Lenslet arrays are not perfect: even in the absence of any external aberration, the spots on the CCD are not centered directly below the center of each lenslet but have random x and y displacements, leading to x and y centroids. A large fraction of the power of these centroids (in this case, about $181/480 = 38\%$) lies in the null-space of the reconstructor matrix. In the case where the lenslet array is of extremely high quality, slope discrepancy could be induced by rotating the lenslet array through a very small angle with respect to the CCD, in which case almost all the power of the induced centroids would lie in the null space of the reconstructor.

Define the slope discrepancy component in the reference centroids, \mathbf{c}_\perp as

$$\mathbf{c}_\perp = ZZ^T \mathbf{c}. \quad (5)$$

Let us define a slope discrepancy metric, g , consisting of the least-squares linear fit of \mathbf{s}_\perp to \mathbf{c}_\perp :

$$g = \frac{\mathbf{s}_\perp^T \mathbf{c}_\perp}{\mathbf{c}_\perp^T \mathbf{c}_\perp}. \quad (6)$$

The centroids will change by a factor of ν if the spot size changes, due to turbulence, by a factor of $1/\nu$, where ν lies between 0 and 1. To compensate for this, the reference centroids are scaled by a factor μ , which ideally is equal to ν .

The metric of Eq. (6) is given by:

$$g = \frac{(\nu \mathbf{r}_\perp - \mu \mathbf{c}_\perp)^T \mu \mathbf{c}_\perp}{\mathbf{c}_\perp^T \mathbf{c}_\perp}. \quad (7)$$

Even in the presence of noise and turbulence, the expected value of \mathbf{r}_\perp , $E[\mathbf{r}_\perp]$ is \mathbf{c}_\perp , so

$$\begin{aligned} E[g] &= \frac{(\nu E[\mathbf{r}_\perp] - \mu \mathbf{c}_\perp)^T \mu \mathbf{c}_\perp}{\mathbf{c}_\perp^T \mathbf{c}_\perp} \\ &= \frac{(\nu - \mu) \mathbf{c}_\perp^T \mathbf{c}_\perp}{\mathbf{c}_\perp^T \mathbf{c}_\perp} \\ &= \nu - \mu. \end{aligned} \quad (8)$$

Eq. (8) is the crux of this technique: the slope discrepancy metric is equal to the difference between the true relative scaling of the centroids due to the changing spot size and the scaling of the reference centroids to compensate for this change. If the reference centroid scaling is correct, then $g = 0$.

Since the measurement is made by averaging over a large number of centroid vectors, the noise can be kept low. Atmospheric turbulence will mostly lie in the column space of the reconstructor (since the job of the reconstructor is to reject turbulence!) and in any event, is also averaged. A possible source of systematic error could be high spatial frequency wave-front aberrations aliased onto lower frequency aberrations. These aberrations could stem either atmospheric turbulence¹⁹ or discontinuities in a segmented primary mirror.²⁰ However, the inner product of these centroids with the slope discrepancy in the reference centroids is likely to be small.

3.2. Experimental results of the slope discrepancy

3.2.1. Results with an artificial light source

To verify the validity of this approach under controlled conditions, it was implemented using the artificial light source with no external turbulence. First, the deformable mirror was flattened using an interferometer. Centroids were measured on the wave-front sensor in open loop and were defined to be the reference centroids. Then the tip/tilt and DM loops were closed and a centroids corresponding to a set of 1000 consecutive frames were captured and averaged. This process was repeated with the reference centroids scaled by factors varying from zero to two. The metric of Eq. (6) was calculated and is displayed in Fig. 2, along with the results expected from theory. It can be seen that the slope of the curve is -1 and the metric takes a value of 0 when the reference centroids are scaled correctly, as expected from Eq. (8).

3.2.2. Results on the sky

The experiment of Sect. 3.2.1 was then repeated on the sky and the results are displayed in Fig. 3. The least-squares fit has a slope of -1.04 and an x -intercept of 0.76, indicating that the optimum scaling of the reference centroids is 0.76. A better estimate of the optimum scaling can be obtained by simply averaging the sum of the slope discrepancy metric and the reference centroid scaling. This is equivalent to forcing the gradient to be equal to -1 and gives an optimum scaling of 0.74.

To validate these results, another method was used to compute the centroids. This technique requires dedicated telescope time and is not suitable for use during a science observation. However, it is very accurate and for this reason it is used as a baseline to compare with the method introduced in this paper. The technique consists of closing the loops and taking an image on the science camera using the artificial light source. Then all the x or y reference centroids are offset by small constant amount, ranging between -0.2 and 0.2, and the displacement of the science image, which is proportional to the spot size, is monitored. The procedure is then repeated using a guide star on the sky, and the ratio of the spot size on the artificial light source to the spot size on the star can be calculated.

Figure 4 plots the results with a star and with the artificial light source. The ratio of the gradient of these two curves is 0.82 for x (plotted) and 0.79 for y . Data taken using these two methods to measure the spot size were interleaved to eliminate the possibility of the seeing changing from one test to the next. There is a

Figure 2. Slope discrepancy metric, g , measured on the artificial light source as a function of reference centroid scaling, α .

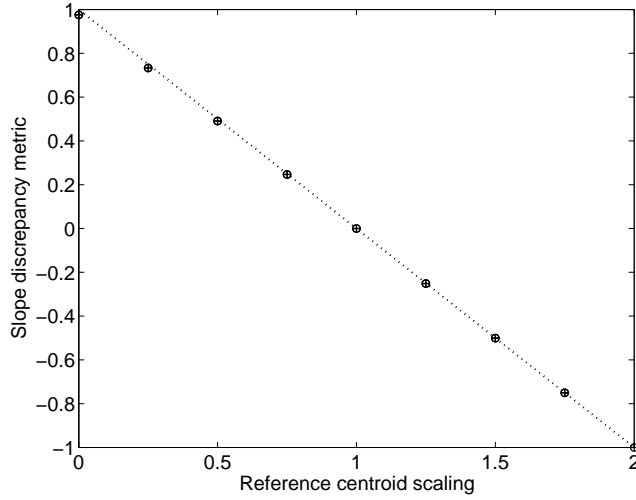



Figure 3. Slope discrepancy metric, g , measuring with a star as a function of reference centroid scaling, α .



slight discrepancy in the results and future work will concentrate on gathering more data and improving the reliability of the slope discrepancy estimates. The discrepancy is likely to consist of systematic errors due to partially illuminated subapertures not being as well-controlled as fully illuminated subapertures.

4. WAVE-FRONT CONTROLLER MODELING

The wave-front controller is described in detail in Johansson *et al.*²¹ and the features of the controller required in the model are simply stated here. A schematic of the TT and DM feedback loops for the Keck AO system is displayed in Figure 5. Both loops have a similar form, but the compensators and the compute delays differ. The wave-front controller can operate at frame rates in the range of 55 to 672 Hz. The upper limit is set by the combination of the camera readout time and the speed at which the computers can multiply the reconstructor matrix by the centroid measurements. There are separate control loops for the TT and the DM. The compute delay times from the time that the CCD is read until the time that the TT and DM voltages are updated, are 1.65 ms and 2.13 ms respectively.




Figure 4. Plot of the mean x -centroid value as a function of displacement of the star (*) and artificial light source (circle)

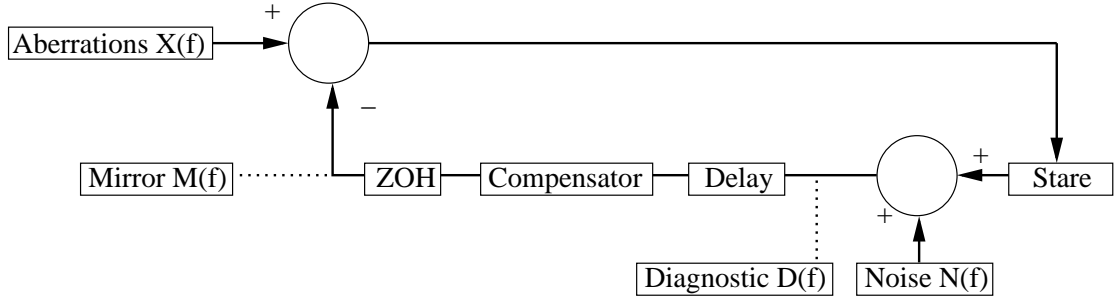


Figure 5. Schematic of the control loop. The diagnostics measure the centroids straight after the addition of the noise and the mirror is driven by the signal coming from the ZOH.

Each of the control loops has its own controller. The TT loop is a straight integrator with a variable loop gain, k_{TT} , and a fixed gain scaling of 0.8:

$$y[n] = y[n - 1] + 0.8k_{TT}u[n], \quad (9)$$

where $y[n]$ is the output from and $u[n]$ is the input to the controller at time n . The transfer function can be written as

$$H_{TT}(z) = \frac{0.8k_{TT}z}{z - 1}, \quad (10)$$

where z is the complex Z-transform variable. The DM loop has a double pole compensator with the following temporal response:

$$e[n] = -we[n - 1] + k_{DM}u[n] \quad (11)$$

$$y[n] = ly[n - 1] + e[n], \quad (12)$$

where $e[n]$ is an intermediate quantity. The transfer function is

$$H_{DM}(z) = \frac{k_{DM}z^2}{(z - l)(z + w)}. \quad (13)$$

The compensator, whose function is to increase the bandwidth of the controller,⁶ has its weight, w , set to 0.25. The leak factor, l , set to 0.999 for bright stars and 0.99 otherwise, ensures that invisible modes do not build up in the DM. The loop gain, k_{DM} , is also variable.

The optimum loop gains depend on the turbulence strength and speed, the extent and brightness of the guide star and the frame rate. In routine operation, the loops gains and the frame rate are set using a look-up table that has the median number of ADUs per subaperture per second as its only input. We are in the process of implementing real-time gain optimization using centroid telemetry.

The transfer functions of the remainder of the blocks are as follows:⁶

- Camera stare and the zero-order hold with period T :

$$H_{\text{STARE}}(s) = H_{\text{ZOH}}(s) = \frac{1 - \exp[-sT]}{sT}. \quad (14)$$

- Compute delay with delay time τ_c :

$$H_{\text{DELAY}}(s) = \exp[-s\tau_c]. \quad (15)$$

The compensator block is modeled by substituting $z = \exp[sT]$ into the discrete compensator of Eqs. (10) and (13). In the above equations, $s = i2\pi f$ is the complex frequency variable and f is the frequency. In what follows, all the blocks will be written with f as the argument, since f has a more intuitive meaning than s and is computed directly from the discrete Fourier transform (DFT) of the diagnostics.

5. BANDWIDTH AND NOISE ERRORS

The bandwidth errors are due to the finite sampling rate of the atmospheric turbulence and the delay between the centroid measurements and the DM and TT command updates. The source of the noise errors is the uncertainty in the centroid estimates due to the finite number of photons on the WFS. To calculate these error terms, a good model of the temporal response of the control loop is required.

5.1. Modeling the dynamic performance of the AO system

The transfer function of the entire feedback leg of either the TT or DM loop, $H(f)$, can be written as the product of the transfer function of all the blocks:

$$H(f) = H_{\text{STARE}}(f)H_{\text{DELAY}}(f)H_{\text{TT/DM}}(f)H_{\text{ZOH}}(f). \quad (16)$$

There are two inputs into the control system: the and the turbulence, $X(f)$, and the noise, $N(f)$, which is assumed to have the same power at all temporal frequencies. Likewise, there are two outputs, the mirror position, $M(f)$, and the residual mirror commands obtained from the diagnostics, $D(f)$. The position of the diagnostics in the control loop is just after the addition of the noise, while the mirror position is just after the zero-order hold, as can be seen in Fig. 5. The residual DM commands are obtained by matrix multiplying the reconstruction matrix with the residual centroids. Likewise, the residual TT commands are obtained by taking the average of the x and y centroids. For notational simplicity, we consider the noise to be input before, rather than after, the stare. This assumption has little impact on the transfer functions of the control loops.

The transfer function relating the the mirror position to the the turbulence and noise is

$$M(f) = \frac{H(f)}{1 + H(f)}(X(f) + N(f)), \quad (17)$$

while the diagnostics power spectrum is given by

$$D(f) = \frac{1}{1 + H(f)}(X(f) + N(f)). \quad (18)$$

The measurement noise wave-front error, σ_{NOISE} , is due to noise in the centroid measurement propagating to the mirror, $N(f)H(f)/(1 + H(f))$. The square of its value is

$$\sigma_{\text{NOISE}}^2 = \int_{-\infty}^{\infty} \left| \frac{H(f)}{1 + H(f)} \right|^2 |N(f)|^2 df. \quad (19)$$

Since the noise only has power at discrete frequencies, Eq. (19) can be rewritten as

$$\sigma_{\text{NOISE}}^2 = \sum \left| \frac{H(f)}{1 + H(f)} \right|^2 |N(f)|^2, \quad (20)$$

where the summation is over all the sampled discrete frequencies.

The bandwidth error, σ_{BW} , is due to the turbulence that is not compensated by the AO system, $X(f)/(1 + H(f))$. The diagnostics measure this term with an added noise term due to the noise on the centroid measurement propagating through the control loop:

$$D(f) = \frac{X(f) + N(f)}{1 + H(f)}. \quad (21)$$

Assuming that the noise and the bandwidth errors are statistically uncorrelated, the bandwidth squared error is

$$\begin{aligned} \sigma_{\text{BW}}^2 &= \int_{-\infty}^{\infty} \left| \frac{1}{1 + H(f)} \right|^2 |X(f)|^2 df \\ &= \int_{-\infty}^{\infty} \left(|D(f)|^2 - \left| \frac{1}{1 + H(f)} \right|^2 |N(f)|^2 \right) df \\ &\approx \sum \left(|D(f)|^2 - \left| \frac{1}{1 + H(f)} \right|^2 |N(f)|^2 \right). \end{aligned} \quad (22)$$

It follows that knowledge of the transfer function of the AO loop, $H(f)$, and the power spectra of the diagnostics, $|D(f)|^2$, and of the noise, $|N(f)|^2$, can be used to calculate both the bandwidth and the noise error terms.

5.1.1. Calculating the centroid variance

In order to estimate the noise and bandwidth errors, we must first derive the error in the centroid estimates. Let I_1 denote the background-subtracted intensity of one of the pixels in a quad cell. Then its expected value is

$$E[I_1] = pT, \quad (23)$$

and, assuming Gaussian statistics, its variance is

$$\text{Var}(I_1) = (p + d + b)T + \sigma_r^2, \quad (24)$$

where T is the integration time, p , d and b are the photon, dark current and background fluxes in electrons/pixel/second and σ_r is the read noise standard deviation in electrons/pixel/readout. In the analysis that follows, $d = 4470$, $b = 0$ and $\sigma_r^2 = 39.7$. To obtain a simple expression for the variance of the centroid estimates, two simplifying assumptions are made: that there is an equal amount of light in each pixel (*i.e.*, the average centroid is zero) and that the effect of the variance of the denominator in the centroid calculation can be neglected. (what is the effect of this???) Then the variance of the x (or y) centroid estimate, c_x (or c_y), is

$$\begin{aligned} \text{Var}(c_x) &= \frac{\text{Var}(I_1)}{4E[I_1]^2} \\ &= \frac{\text{Var}(I_1)}{4(pT)^2} \\ &= \frac{(p + d + b)T + \sigma_r^2}{4(pT)^2}. \end{aligned} \quad (25)$$

The measured centroid variances agree with Eq. (25). From the variance of the centroids, we can calculate the errors in the TT and DM loops.

5.1.2. Calculating the tip/tilt noise power spectrum

The tip (or tilt) signal sent to the compensator is the average x (or y) centroid value over the 240 active subapertures. Hence the sum of the variance of the tip and tilt estimates in centroid units is $2\text{Var}(c_x)/240$. The next step is to convert this variance into units of WF error. Using the conversion factor for a Gaussian spot between arcseconds to centroid units, m , is²²

$$m = \frac{2 \times 2.355}{\sqrt{2\pi}\text{FWHM}}. \quad (26)$$

Since the FWHM of the spot using the calibration source is 1.25 arc sec, we obtain the result that 1 arcsec is equivalent to 1.50ν centroid units.

The factor ν is obtained from slope discrepancy, as described in Sect. 3. For the Keck telescope, the RMS WF error due to a tip or a tilt of 1 arc sec is 12.68 μm . Hence, the TT noise power spectrum is

$$|N_{\text{TT}}(f)|^2 = \frac{2}{240} \left(\frac{12.68}{1.50\nu} \right)^2 \text{Var}(c_x). \quad (27)$$

5.1.3. Calculating the deformable mirror noise power spectrum

To calculate the deformable mirror noise PSD, we must convert centroid units into wave-front units. The residual centroids are multiplied by the reconstruction matrix, R , to convert to voltage commands. The voltage commands are then convolved with the actuator influence function, S , to obtain the wave-front induced by the noise. This is achieved using matrix multiplication. Finally, since the size of the WFS spots on the sky increases by a factor of $1/g$ relative to their size on the light source, the wave-front must be divided by ν . Since the noise on the actuator voltages are uncorrelated, we can write

$$|N_{\text{DM}}(f)|^2 = \nu^{-2} |RS|^2 \text{Var}(c_x). \quad (28)$$

There is an implicit assumption in the control loop that the conversion between actuator and centroid space and back to actuator space has no effect on the loop. In other words, it is assumed that the reconstructor is in fact the pseudo-inverse of the system matrix and that $HR = I$. In Sect. 2, it was shown that this does not hold for the Bayesian reconstructor because there is cross-talk between modes. Preliminary simulations show that the main effect of neglecting the reconstructor and system matrices in the analysis is to underestimate the bandwidth and measurement noise errors. In future work, these effects will be included in the model.

5.1.4. Calculating the bandwidth and noise errors

To calculate the residual power spectra, the centroids from the diagnostics must be converted into TT and DM wave-front aberrations in exactly the same way as the noise. Then the power spectrum of the diagnostics is taken using the DFT:

$$|D(f)|^2 = |\text{DFT}[d[n]w[n]]|^2, \quad (29)$$

where $w[n]$ is a normalized windowing function used to avoid spectral leakage due to the non-periodicity of $d[n]$, the residual wave-front as measured by the diagnostics. The window must be scaled to have unit power.

The TT and DM noise errors are calculated by inserting Eqs. (27), (28) and (29) into Eqs. (20) and (22).

6. CONCLUSION

In this paper, the details of a successful wave-front reconstructor, which takes into account the statistics of the atmosphere, was presented. This reconstructor led to improvements in Strehl ratio of over 10% relative to using an SVD reconstructor. A novel technique that uses slope discrepancy to measure the size of the WFS spot on quad cells was presented. Sky tests showed the feasibility of this technique, although more work is needed to improve its accuracy. Finally, rigorous analysis of the AO control loop was undertaken, which can be used to estimate and reduce the bandwidth and measurement noise errors. This analysis is complicated by the fact that there is cross-talk between the different modes of the reconstructor and future work will concentrate on modeling this cross-talk.

ACKNOWLEDGMENTS

This work was performed under the auspices of the US Department of Energy by the University of California, Lawrence Livermore National Laboratory, under contract W-7405-Eng-48. The work has been supported by the National Science Foundation Science and Technology Center for Adaptive Optics, managed by the University of California at Santa Cruz under cooperative agreement No. AST-9876783.

REFERENCES

1. M. C. Roggemann and B. Welsh, *Imaging through turbulence*, CRC Press, Boca Raton, Fla. (1996).
2. J. W. Hardy, *Adaptive Optics for Astronomical Telescopes*, Oxford University Press, New York (1998).
3. P. L. Wizinowich, D. S. Acton, C. Shelton, P. Stomski, J. Gathright, K. Ho, W. Lupton, K. Tsubota, O. Lai, C. Max, J. Brase, J. An, K. Avicola, S. Olivier, D. Gavel, B. Macintosh, A. Ghez, J. Larkin, "First light adaptive optics images from the Keck II telescope: a new era of high angular resolution imagery," *PASP* **112**, 315-319 (2000).
4. M. A. van Dam and B. A. Macintosh, "Characterization of adaptive optics at Keck Observatory," *Proc. SPIE* **5169**, 1-10 (2003).
5. M. A. van Dam, D. Le Mignant and B. A. Macintosh, "Performance of the Keck Observatory adaptive optics system," *submitted to Applied Optics* (2004).
6. P.-Y. Madec, "Control techniques," in *Adaptive Optics in Astronomy*, pp 131-54, ed. F. Roddier, Cambridge University Press, Cambridge (1999).
7. E. Gendron and P. Léna, "Astronomical adaptive optics I. Modal control optimization," *Astron. Astrophys.* **291**, 337-347 (1994).
8. C. Dessenne, P.-Y. Madec and G. Rousset, "Optimization of a predictive controller for closed-loop adaptive optics," *Appl. Opt.* **37**, 4623-4633 (1998).
9. C. Dessenne, P.-Y. Madec and G. Rousset, "Sky implementation of modal predictive control in adaptive optics," *Optics Letters* **24**, 339-341 (1999).
10. P. J. Stomski, Jr. and J. C. Shelton, "Compensating for pupil rotation in the W. M. Keck Observatory adaptive optics system," in *Adaptive Optical Systems Technology*, P. L. Wizinowich, ed., *Proc. SPIE* **4007**, 608-619 (2000).
11. N. F. Law and R. G. Lane, "Wavefront estimation at low light levels," *Optics Comms* **126**, 19-24 (1996).
12. E. P. Wallner, "Optimal wave-front correction using slope measurements," *J. Opt. Soc. Am.* **73**, 1771-1776 (1983).
13. D. Gavel, "Suppressing anomalous localized waffle behavior in least squares wavefront reconstructors," *Proc. SPIE* **4839**, 972-980 (2003).
14. D. L. Fried, "Least-squares fitting a wave-front distortion estimate to an array of phase-difference measurements," *J. Opt. Soc. Am.* **67**, 370-375 (1977).
15. J-P. Véran and G. Herriot, "Centroiding gain compensation in Shack-Hartmann adaptive optics systems with natural or laser guide star," *J. Opt. Soc. Am A* **17**, 1430-39 (2000).
16. M. A. van Dam, "Measuring the Shack-Hartmann spot size using slope discrepancy," *in preparation* (2004).

17. G. A. Tyler, "Reconstruction and assessment of the least-squares and slope discrepancy components of the phase," *J. Opt. Soc. Am. A* **17**, 1828-1839 (2000).
18. Gilbert Strang, *Linear algebra and its applications*, 3rd edition, Harcourt Brace, Florida, p. 68 (1988).
19. F. J. Rigaut, J.-P. Véran and O. Lai, "Analytical model for Shack-Hartmann-based adaptive optics systems", in *Adaptive Optical System Technologies*, D. Bonaccini and R. K. Tyson, eds., Proc. SPIE **3353**, 1038-1048 (1998).
20. G. Chanan, C. Ohara and M. Troy, "Phasing the mirror segments of the Keck telescopes II: the narrow band phasing algorithm," *Applied Optics* **39**, 4706-4714 (2000).
21. E. M. Johansson, D. S. Acton, J. R. An, K. Avicola, B. V. Beeman, J. M. Brase, C. J. Carrano, J. Gathright, D. T. Gavel, R. L. Hurd, O. Lai, W. Lupton, B. A. Macintosh, C. E. Max, S. S. Olivier, J. C. Shelton, P. J. Stomski, K. Tsubota, K. E. Waltjen, J. A. Watson and P. L. Wizinowich, "Initial performance of the Keck AO wavefront controller system," in *Adaptive Optical Systems Technology*, P. L. Wizinowich, ed., Proc. SPIE Vol. 4007, 600-607 (2000).
22. M. Schoeck, D. Le Mignant, G. Chanan, P. L. Wizinowich and M. A. van Dam, "Atmospheric characterization with the Keck adaptive optics system I: open-loop data," *Applied Optics* **42**, 3705-3720 (2003).