Keck Adaptive Optics Note #356

Wavefront reconstruction for the NGWFC

Marcos van Dam

21 September 2005

1 Introduction

There are two wavefront reconstruction philosophies in use with astronomical AO systems. The first is to use a regularized least-squares reconstructor, as used by the current Keck AO system. The second is to use a modal reconstructor, where each mirror mode is optimized separately. The purpose of this investigation is to determine what type of reconstructor the NGWFC should use. The two methods are compared, both qualitatively and numerically through simulations. Based on these findings, we should continue to use the Bayesian reconstructor.

Two other issues are revisited: how to best drive the actuators outside the illuminated pupil and how to best weight the partially illuminated subapertures.

2 Reconstructors

In this section, the two reconstruction philosophies are described.

2.1 Regularized least-squares reconstructors

The standard least-squares reconstructor, R, inverts the poke matrix, H, as follows

$$R = (H^T H)^+ H^T,$$

where $^+$ denotes the pseudo-inverse, which is necessary since $H^T H$ is singular. The Bayesian reconstructor inverts the poke matrix, H, using the covariances of Kolmogorov turbulence, C_φ, and the relative noise in each subaperture, W, as prior information. The covariance matrix of Kolmogorov turbulence, C_φ, is found by calculating the covariance between the points corresponding to each of the 349 actuators using the method of Wallner.[1] A matrix that penalizes local waffle can be used instead of C_φ. There is a noise-to-signal parameter, α, that is adjusted depending on the brightness of the guide star and the strength of the seeing. The reconstructor, R, is given by

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

[1] E.P. Wallner, JOSA **73**, 1771 (1983).

where **1**=[1,1,1,…]^T and η is a small parameter that prevents the reconstructor from producing piston. More descriptions of this regularized least-squares reconstructor can be found in van Dam *et al*.[2]

2.2 Modal reconstructors

Modal reconstructors consider the wavefront to consist of a linear combination of mirror modes. The choice of mirror modes is arbitrary, but the purpose of selecting modes is to enable the control law of each mode to be optimized separately. An excellent discussion of modal control can be found in the description of Altair's reconstruction matrix.[3] A good set of modes should contain "good" modes, where there is much more turbulence than noise, "not so good" modes, where both the rejection and the noise propagation of the mode are reduced, and "bad modes", which contain much more noise than turbulence and are rejected. A good set of modes would appear to be the Karhunen-Loève modes for Kolmogorov turbulence. These modes are orthogonal and statistically independent and the estimates of the power in each mode are statistically independent in open loop. This set of modes was chosen for the analysis that follows. The reconstructor is the given by

$$R = MG((HM)^T HM)^{-1}(HM)^T,$$

where M is the matrix that converts from mode space into actuator space and G is the diagonal matrix that contains the gain for each mode. It can readily be seen that if G is the identity matrix and M is a square matrix, then the reconstructor us the same as the least-squares reconstructor. If M is the identity matrix, then each actuator is a mode. The subaperture weight matrix can also be incorporated as follows:

$$R = MG((HM)^T W^{-1} HM)^{-1}(HM)^T W^{-1}.$$

The advantage of modal reconstructors is that, the residual of each mode can be calculated: $((HM)^T W^{-1} HM)^{-1}(HM)^T W^{-1} s$, where s is the centroid vector, and this information can be used to update the loop gains in matrix *G*.

There is a practical difficulty with estimating the residual in each mode in this way. There are many actuators that, at a given time, are not illuminated and thus have little or no influence on the measured centroids. Hence, the value of the residual at the location of these actuators could be wildly wrong. Unfortunately, because each mode includes every actuator, the estimate for the residual in each mode is wrong. Note that this does not result in an error in the estimate of the wavefront: the sum of all the modes will still produce the right estimate for the wavefront *within* the region of the deformable mirror illuminated by the pupil. But since the wavefront outside the pupil is wrong, the modal coefficient estimates are unreliable. It should be noted that the zonal reconstructor does not suffer from this problem, since the residual wavefront estimates are performed on an actuator by actuator basis and the actuators that lie outside the pupil can be safely ignored.

There is a solution to this problem: the modes can be defined over the illuminated pupil rather than the circular pupil. For a stationary pupil, such as a pupil arising from circular telescope with a central obscuration, this approach can be easily implemented. For a non-circular rotating and nutating pupil, this necessitates that the modal set be continually recreated as the pupil angle changes. This already causes significant problems, since we are trying to optimize the loop gains for each mode. Because the modes themselves keep changing, the previous gain estimates are rendered obsolete even if the atmospheric turbulence does not change. The problem is not insurmountable: there are some games one could play, such as rotating the modes with the pupil. But the complication of the problem is such that modal control would only be worthwhile if it provides a significant benefit.

[2] M.A. van Dam, D. Le Mignant and B.A. Macintosh, "Characterization of adaptive optics at Keck Observatory: part II," Proc SPIE 5490, 174-183 (2004).
[3] J.P. Véran, "Altair's optimizer".

2.3 Singular value decomposition

A possible choice for the modes is to use the singular value decomposition to select the mode matrix, *M*. The modes here are chosen from ordered from modes containing the least amount of noise to modes containing the greatest amount of noise, irrespective of the statistics of the turbulence. In other words, the modes are ranked in terms of observability. Modes with spatial frequencies equal to half the Nyquist frequency have the highest observability, while modes at the Nyquist frequency, such as waffle, or at very low frequencies, such as piston, have the lowest observability.

The matrix is computed as, in IDL notation:

`SVDC, H,U,V,M` ; M is the mode matrix that converts modes to actuators
`D = H##M` ; D converts modes to centroids
`C =` $(D^T D)^{-1} D^T$; C converts centroids to modes
`R = M##G##C` ; R converts centroids to actuator space

In order for $D^T D$ to be well-conditioned, the modes with the lowest singular values have to be removed. The number of modes and the gains for each mode of the modal gain matrix, G, can be optimized.

3 Simulation results

The simulations were deliberately simplified by omitting known sources of error such as off-null calibrations, telescope vibrations, image sharpening errors, etc. The idea was to have the bandwidth and measurement noise errors dominate, in order to accentuate the effect of the reconstructor. Preliminary results using short simulations showed that when the guide star is bright, the performance of the reconstructors is comparable, with the Bayesian reconstructor slightly better. The reason for this is that the fitting error is the dominant source of error. The parameters used, based on the CCD 39, are displayed in Table 1:

Magnitude	13
Iterations	1000 (+20 to converge)
Frame rate	300 Hz
Read noise	4.4 electrons
Dark current	200 electrons/pixel/second
Pixel size	1.03”
Charge diffusion	0.6 pixels FWHM

Table 1: Parameters used in the simulations

The tip-tilt controller was an integrator with gain 0.45, while the DM controller was a leaky integrator with variable gain.

First, the Bayesian reconstructor was used, starting with a gain of 0.4. The regularization parameter, α, was set to 10 throughout. The optimum gain predicted from the telemetry (residual centroids) was 0.45. In fact, this turned out to be true, as can be seen in Table 2. The same optimization can be done on a mode-by-mode basis, as occurred for the singular modes basis. The modal gain optimization for the Karhunen-Loeve basis did not improve the performance. However, it gave a modest improvement when the simulations were done with a full circular aperture and no central obscuration, which is the space over which these modes are orthogonal. Even then, the results were not as good as for the Bayesian reconstructor. The results imply that the Karhunen-Loeve modes have to be recomputed for every pupil angle for this basis to be useful. In any case, there is no evidence, either in the literature or in these simulations, to suggest that modal reconstructors are better than regularized least-squares reconstructors. Since the Bayesian least-squares reconstructor has been in use at Keck Observatory since 2003 and is the most robust of the methods investigated, it is suggested that we continue to use this reconstructor. It is the most robust method because only one parameter has to be estimated from the telemetry: the optimal loop gain. The modal reconstructors need the gains to be calculated for each mode, which could lead to large errors if the gain finding algorithm does not converge properly.

Reconstructor	Gain	Mode weights	Error (nm)
Bayesian	0.40	Not applicable	277.8
	0.45	Not applicable	275.4
	0.50	Not applicable	277.0
Singular modes	0.45	285 Equal	289.2
	0.45	285 Optimized	279.7
Karhunen-Loeve	0.45	225 Equal	324.0
	0.45	200 Equal	310.8
	0.45	175 Equal	288.5
	0.45	165 Equal	288.8
	0.45	150 Equal	289.3

Table 2: Simulation results showing the wavefront error for each of the three reconstructors

4 PSF estimation

An important concern for astronomers is to know what the PSF is for a particular observation. PSF estimation from control loop data has been successfully implemented at CFHT on PUEO, but there is no currently working PSF estimator for Shack-Hartmann systems. PSF estimation typically uses the covariance matrix for the error in each mode. This presents several problems. For the Bayesian reconstructor, or any other regularized inversion of the system matrix, *H*, there are no modes and hence there is no obvious way to determine to attribute any wavefront error to a location in the PSF. For the modal reconstructors, the modes depend on the pupil angle, and hence change when the illumination of the pupil changes. This means that we need to keep track of not only of the covariance matrix but also the modes themselves, making PSF reconstruction difficult. In addition, even if the modes stayed the same, the pupil doesn't, and the effect of an error in a given mode also depends on the pupil.

5 Conclusion

Simulations show that the Bayesian reconstructor is at least as good as a modal reconstructor with optimized number of modes and modal gains. In addition, the Bayesian reconstructor is computationally cheaper than the other methods. We have a lot of experience with the Bayesian reconstructor and it has proven to be very robust to calibration errors, especially to DM-to-lenslet registration errors. Hence, it is recommended that we continue to use the Bayesian reconstructor.