

Sparse Matrix Wavefront Reconstruction: Simulations and Experiments

Fang Shi^a, Douglas G. MacMartin^b, Mitchell Troy^a, Gary L. Brack^a, Rick S. Burruss^c,
and Richard G. Dekany^b

^aJet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

^bCalifornia Institute of Technology, Pasadena, CA 91125

^cPalomar Observatory, M/C 105-24 PAL, Pasadena, CA 91125

ABSTRACT

Adaptive optics systems with Shack-Hartmann wavefront sensors require reconstruction of the atmospheric phase error from subaperture slope measurements, with every sensor in the array being used in the computation of each actuator command. This fully populated reconstruction matrix can result in a significant computational burden for adaptive optics systems with large numbers of actuators. A method for generating sparse wavefront reconstruction matrices for adaptive optics is proposed. The method exploits the relevance of nearby subaperture slope measurements for control of an individual actuator, and relies upon the limited extent of the influence function for a zonal deformable mirror. Relying only on nearby sensor information can significantly reduce the calculation time for wavefront reconstruction. In addition, a hierarchic controller is proposed to recover some of the global wavefront information. The performance of these sparse wavefront reconstruction matrices was evaluated in simulation, and tested on the Palomar Adaptive Optics System. This paper presents some initial results from the simulations and experiments.

Keywords: Adaptive Optics, Wavefront Sensor, Wavefront Reconstruction, Decentralized Control

1. INTRODUCTION

Many large ground based astronomical telescopes use adaptive optics (AO) systems to provide a real time wavefront compensation for atmospheric turbulence. Adaptive optics uses a reference source to measure the wavefront information and feedback to a deformable mirror (DM) for correction^[1]. The nature of the atmospheric turbulence requires the wavefront to be spatially sampled at sizes on the order of the coherent length of turbulence, which is about 30 cm for near IR, and at speed much faster than the characteristic frequency of the turbulence, which is about tens of Hertz. Current AO systems typically involve hundreds of actuators and system running at speeds of 300 – 1000 Hz. Future large telescopes and the increasing requirements of current AO system will require significantly more actuators and sensors as they will increase proportionally to the square of the telescope's diameter.

Most high order AO systems use a Shack-Hartmann sensor which measures the slope of the wavefront across each subaperture. The correcting wavefront phases at the DM actuator locations are estimated from the measured subaperture slopes. In real time AO operation, the measured wavefront slopes, which are usually represented by centroid positions of the reference source in each subaperture, are multiplied by a wavefront reconstruction matrix to calculate the desired DM actuator positions. This reconstruction matrix is based on a weighted pseudo-inverse of the influence matrix from actuator displacements to sensor measurements, and is in general fully populated. The entire sensor vector is therefore required to give a good estimate of global wavefront error at any point. For an AO system with n subapertures across the entire aperture the number of subapertures and DM actuators are of order n^2 . With 2 slope measurements from each subaperture the number of multiply-accumulation calculations for each wavefront reconstruction is about $2n^4$. Since n increases linearly with the telescope diameter D the computational burden for larger telescopes such as CELT^[2] could be unacceptable.

One approach to reducing computation is based on using spatial Fourier transforms^[3]; the resulting computations are of order $n^2 \log(n)$ rather than n^4 for the full matrix computation. A conjugate gradient approach can be used to solve the problem iteratively retaining only sparse matrix operations, and using pre-conditioning to improve convergence^[4]. More general iterative approaches are briefly discussed in Hardy^[1] and form the basis of the sparse matrix approach of Wild^[5]. Other sparse matrix approaches have also been suggested^[6], however the reconstructors are not in general optimal.

We have studied a sparse matrix method which exploits the relevance of nearby slope measurements for control of an individual actuator, and relies upon the limited extent of the influence function for a zonal deformable mirror. Relying on only the nearby sensor information can significantly reduce the calculation time for wavefront reconstruction. However, with this local approach, performance on low spatial order modes suffers, and a hierarchic approach is needed to recover global modes separately. The performances of these sparse wavefront reconstruction matrices were evaluated in simulations and experiments on the Palomar Adaptive Optics System (PALAO). Detailed simulations of the localized least square sparse matrices and hierarchic controllers have been done by MacMartin and can be found in reference^[7].

This paper describes the local and hierarchic approach for sparse reconstruction matrices, presents experimental evaluations of these sparse matrices obtained at Palomar Observatory, and compares the experimental results with the analytical predictions^[7]. In section 2 we will present a brief introduction of the sparse matrices theory. We then describe our sparse matrices experiment using PALAO and present some of the results which characterize the performance of the sparse matrices in section 3. We will also compare the experimental results with our simulation predictions. Finally, in section 4 we will summarize our experience with these sparse matrices and discuss future development of our sparse matrix technique.

2. SPARSE WAVEFRONT RECONSTRUCTION MATRICES

For a wavefront sensor which measures the subaperture wavefront slopes s , the influence of the phase from the DM actuators w is defined by the geometric matrix A as

$$s = A w + n \tag{1}$$

where n is the noise in the measurement. In absence of sensor noise, the reconstruction matrix A^+ which is used in the AO system to calculate the desired DM actuator positions from the measured wavefront slopes is the pseudo-inverse of the geometric matrix A ,

$$w = A^+ s = (A^T A)^{-1} A^T s \tag{2}$$

Although the geometric matrix A is sparse due to the localized influence of DM actuators the reconstruction matrix A^+ is in general fully populated. The physical interpretation of this is that the absolute displacement can only be estimated from relative measurements by integrating those measurements to the boundary of the domain; given a slope between two locations, one cannot determine whether a phase point is high or the other is low without considering the entire domain.

Upon a close examination of the reconstruction matrix A^+ one notices that the matrix components for each actuator are dominated by the influence of slope measurements from nearby subapertures and the contributions from other subapertures decline rapidly as the distance between the actuator and subaperture increases. This is because each DM actuator has limited influence over the wavefront phase. A brute force approach to make a sparse A^+ is to simply truncate the matrix so that each actuator uses only the nearby subapertures and the contributions of the slope measurements from other subapertures are discarded. This would significantly limit the number of non-zero components in the matrix. For a regular discrete DM actuator and Shack-Hartmann subaperture layout the truncated reconstruction matrix has a banded structure. This can be coded into the AO reconstructor by a suitable indexing scheme^[8] so that instead of multiplying the entire vector of slopes with the whole row of the matrix only the related slope values and matrix elements are used. In this way the required reconstruction calculation is of order $2kn^2$ instead of $2n^4$. The constant $k = d \times d$ is the number of

subapertures used for each actuators in the sparse matrix, which is only depend the size of local influence region d , and unlike n , it does not increase with the diameter of the telescope D . By reducing the dependence of n to the telescope diameter D from power of 4 to 2, the truncated sparse matrix can significantly reduce the reconstruction calculations for large aperture telescopes. In principle, truncating A^+ by zeroing elements corresponding either to distant sensors or to those with small gain is unfavorable because the direct truncation does not optimize the local information for each actuator. However, the reconstructor gain does decay with distance from the actuator, and thus the performance of truncated reconstruction matrix, as will be shown in the next section, is still reasonable.

The optimal way to utilize the local wavefront slope information around each actuator is to reconstruct the phase for the actuator using the slopes measured from the nearby subapertures using the least square method, an approach we call “localized least square” (LLS) method. Like the truncated sparse matrix, the sparse matrix generated by the LLS method has about $2kn^2$ non-zero components, but the LLS method provides an optimized estimation for each actuator using the slopes from the $d \times d$ subapertures that surround the actuator. The performance degradation of this localized controller relative to the global modes can be predicted by noting that the gain is a spatial high-pass filter with spatial frequency response given by

$$g(k_x, k_y) = 1 - \text{sinc}(k_x d) \text{sinc}(k_y d) \quad (3)$$

Here $k_x = m_x \pi / D$, $k_y = m_y \pi / D$ (m_x, m_y are integers) are the wave-numbers in two orthogonal directions, and thus the gain on the lowest spatial frequency mode is $g_{11} = \text{sinc}^2(\pi d / D)$. Even for modest d/D there is a significant reduction in low frequency gain. While the sparse reconstruction matrix maintains good performance on high spatial wave-number modes, the gain is reduced on roughly the first $(D/d)^2$ low wave-number modes which correspond to the global modes with half-wave scale between d and D . The decrease in gain for the lower spatial wave-number modes has been observed in the experiment.

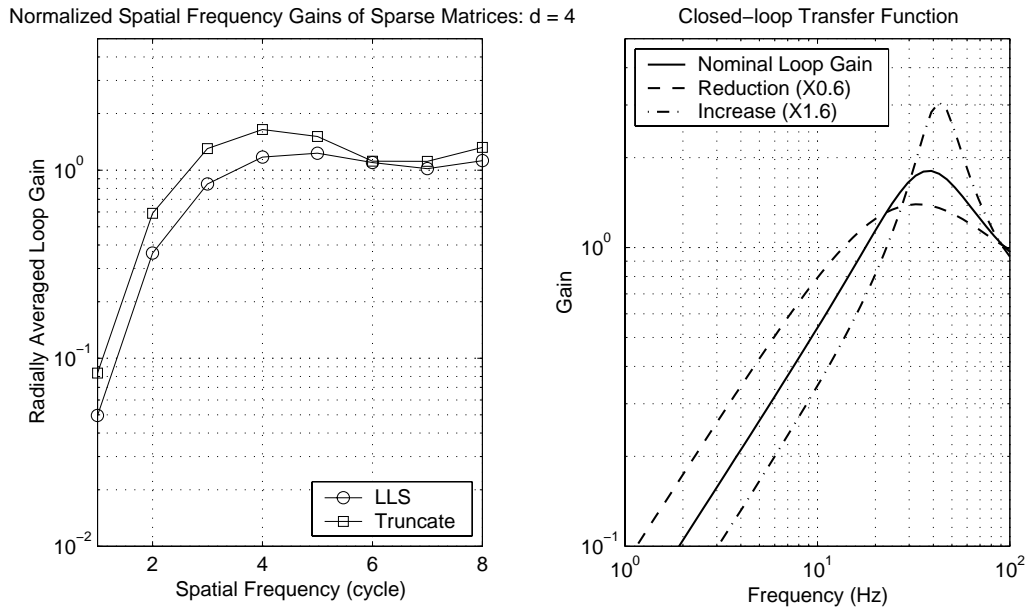


Figure 1. (a). (Left panel) Radially averaged spatial gains for a LLS and a truncated sparse matrix with $d = 4$. The gains are normalized by that of the full matrix. (b). (Right panel) The close loop transfer function for a nominal, a decreased (by a factor of 0.6), and an increased (by a factor 1.6) gain.

The spatial filtering behavior of the sparse reconstructors is shown in Figure 1(a) for both the truncated and LLS reconstructors with local influence size $d = 4$. The gains are normalized by the full least square reconstructor and are plotted as a function of spatial frequency. The LLS reconstructor deviates from the sinc^2 prediction due to averaging

over all modes with a given spatial frequency, however, the general trend can still be observed. The gain of both sparse reconstructors is decreased for low spatial frequencies (global modes), which leads to a drop in the closed loop performance for these modes. There is also an increase in the gain for certain characteristic modes, which is significantly larger for the case of truncated reconstructors. The effect of the gain variations can be seen from the corresponding transfer functions shown in Figure 1(b) which plots a nominal gain, a decreased gain, and an increased gain. This is the expected closed loop decrease (or increase) in the amplitude of disturbances as a function of frequency. With the nominally tuned gains for the PALAO system, then the -3 dB bandwidth is 13-15 Hz, and the peak of the sensitivity is about 1.8 near 40 Hz. A reduction in gain corresponds to a reduction in the closed loop bandwidth, and a corresponding decrease in performance. However, an increased gain, as evident at certain spatial frequency in Figure 1(a), can also be a problem due to the increase in the amplification at frequencies above the bandwidth. If the nominal gain is chosen correctly so as to minimize the overall performance, then any gain greater than unity will degrade performance. Thus, although the truncated reconstructor has higher gains, it should result in lower performance for an optimally tuned system.

A hierarchic approach can be used to recover global mode performance without losing the computational savings of the localized controllers. Performance at low spatial frequencies can be improved by estimating the information that the localized estimator does not by adding extra global reconstructions. This consists of three steps: (1) spatial filtering (e.g. averaging over a region of several subapertures) of the sensor information and condensing into a reduced set of data; (2) estimating global parameters from this condensed data, and (3) distributing these global parameters over the entire DM to combine with the local estimations. Although it is not the only choice, the global parameterization we use is chosen to be geometrically similar to that of the full problem, but with different spatial scale. By maintaining geometric similarity, the same process can readily be repeated over multiple scales. A different control bandwidth could be obtained for different spatial scales (e.g. to compensate for different disturbance covariance) by computing the control as the sum of contributions from the different hierarchic layers. The computation burden for the hierarchic control will increase due the extra layer(s) reconstruction calculation for the global modes. However, the condensed data set limits the number of extra calculations to the order of $2(D/d)^2$. For the two-layers reconstructor demonstrated here, the hierarchic controller scales as $n^{8/3}$. In the limit of $\log_2(n^2)$ layers one might expect the total hierarchic reconstruction calculation of about $n^2 \log_2(n)$, a scaling similar to that of FFT approach^[3].

3. SPARSE MATRICES EXPERIMENTAL RESULTS

The Palomar Adaptive Optics system^[9] has been used to test the sparse matrices described in above section. A facility AO instrument for the Palomar 5-meter Hale Telescope, PALAO has a Shack-Hartmann wavefront sensor with 16 subapertures across the diameter of the telescope, a DM with 241 active actuators, and a fast steering mirror (FSM) for overall tip-tilt correction. Both the subapertures and actuators lie in a square grid with 4 DM actuators at the corners of each subaperture in the standard Fried geometry. An infrared camera system, the Palomar High Angular Resolution Observer (PHARO), provides IR imaging and coronagraph for PALAO. PALAO became a facility instrument in May of 2000 and the system routinely achieves Strehls of 50 – 60% (in K band), thus providing a good testbed for AO technique development.

In our sparse matrices experiment all three types of sparse matrices discussed in section 2, including truncation of a full least square matrix, localized least square, and a hierarchic reconstructor, were tested. We tested both truncated and LLS sparse matrices with influence regions varying from $d = 2$ (i.e. each actuator is driven by 2×2 subapertures), which is the sparsest, to $d = 12$. The truncation sparse matrices were generated from a full least square reconstruction matrix. We tested one hierarchic matrix which was constructed with a $d = 4$ LLS method with an additional global mode reconstructor which senses low order modes on 4×4 points over the entire aperture. During tests each sparse matrix as well as the full least square reconstruction matrix was sequentially loaded into the PALAO's wavefront reconstructor and then the AO loop was closed. Wavefront sensor data was recorded for about 60 seconds and three PHARO K band images were taken. The recorded wavefront sensor data includes status of DM and FSM loops, subaperture centroids, subaperture fluxes, FSM positions and residuals, and DM actuator positions and residuals. The PALAO system update rate is 500 Hz and the telemetry recorder can sustain a recording rate of 100 Hz, i.e. every 5th frame of data. The sparse matrices have been tested both on bright ($M_V < 6 \text{ mag.}$) and faint stars ($M_V > 8 \text{ mag.}$) as well as under different

atmospheric seeing conditions. In the experiment, although many elements of the sparse matrices are zeros, the matrices were still maintained in the full matrix format, loaded into the wavefront reconstructor as a regular matrix, and the reconstructor still performed the full matrix calculation. This did not take the advantage of computation savings of the sparse matrices and the number of reconstruction calculation remained the same. However, experiments in this fashion allow us to directly compare the performance of the sparse matrices. Because the experiments have tested the performance of the sparse matrices under real atmospheric turbulence the recorded real time wavefront data have provided us with detailed information on wavefront error's spatial and temporal properties for different reconstruction matrices. Analysis of data has shown that the performance of sparse matrices behaved similarly for both bright and faint stars, though the AO system tends to have larger residual wavefront errors due to the lower signal-to-noise ratio (SNR). To illustrate the difference between the sparse matrices and to distinguish the characteristics of the sparse matrices we use the result from a bright star (SAO64701, a $M_V = 5.9$ K0 star) in this paper.

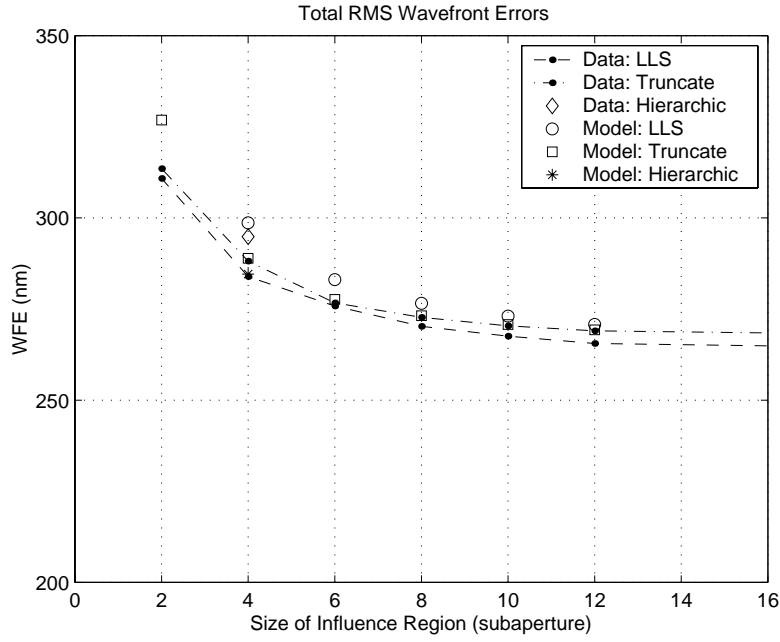


Figure 2. Experimental and model predicted rms residual wavefront errors versus matrix sparseness. The wavefront error is in unit of nm and the subaperture influence size d is in unit of *subaperture*. The WFE uncertainty estimated from the fluctuation of Strehls is about $20 nm$.

Figure 2 shows the rms residual wavefront errors for the AO closed-loop performance using different sparse matrices. The residual wavefront error is calculated from the recorded wavefront sensor data and the PHARO images. For each sparse matrix the closed-loop residual wavefront σ comes from two sources: (1) the AO system error σ_{sys} which is the sum of errors from the servo time delay, residual tip-tilt, DM fitting, DM registration, non-common path calibration, and wavefront sensor measurement error^[9]. This part of the error is estimated using the Strehl ratio from the PHARO images taken when the PALAO is closed loop with the full matrix; and (2) the error caused by the sparse matrix σ_{sparse} which is calculated from the recorded wavefront sensor slopes when the AO loop was closed by a sparse matrix as,

$$\sigma_{sparse}^2 = 4 \cdot \sum_j \frac{(a_j^{Full} - a_j^{Sparse})^2}{n_{act}} \quad (4)$$

where n_{act} is the number of actuator, and the factor of 4 is added in Eq. 4 to account for the DM reflection. The a_j^{Full} and a_j^{Sparse} are the reconstructed DM actuator positions in microns from the recorded closed-loop slopes \mathbf{s}_{sparse} using full matrix \mathbf{A}^{Full+} and sparse matrices $\mathbf{A}^{Sparse+}$ correspondingly, i.e.

$$\begin{aligned} \mathbf{a}^{Full} &= \mathbf{A}^+_{Full} \cdot \mathbf{s}_{Sparse} \\ \mathbf{a}^{Sparse} &= \mathbf{A}^+_{Sparse} \cdot \mathbf{s}_{Sparse} \end{aligned} \quad (5)$$

Eq. 4 and 5 calculate the wavefront error caused by the matrix sparseness by calculating phase difference between the reconstructions of a full matrix and the sparse matrix. In Figure 2 the sparse matrix closed-loop wavefront error $\sigma = (\sigma_{sys}^2 + \sigma_{sparse}^2)^{1/2}$ is plotted against the sparseness of the matrices which is indicated by the subaperture influence region size d . The plots show errors from sparse matrices generated in all three approaches as described in section 2. We can see from the plot that the wavefront error increases as the reconstruction matrix become sparser. However, there is little difference between the truncation, LLS, and hierarchic controller. The remaining difference between the truncation and LLS WFE curves are mainly due to the fluctuation of full matrix PHARO image Strehls which have been used to calculate the system error. The WFE fluctuation is estimated form the fluctuation of Strehls to be about 20 nm . The plots show that AO system error is the dominant component error until sparseness is $d \leq 4$. This is also reflected in the Strehl ratio plot shown in Figure 3. Figure 2 also plots the predicted wavefront error from the simulations conducted for various reconstructors in order to compare the predicted performance with the experimental results. The simulations use open loop telemetry data recorded during the experiments with the FSM (tip/tilt) loop closed and the DM loop open. Because the telemetry data are recorded only at 100 Hz, the data are interpolated for simulation. As a result, there is no excitation above the Nyquist frequency of 50 Hz in the simulation and thus, recalling Figure 1(b), the simulation will predict better performance than it should, particularly for case of increases in the loop gain. Thus while we expect the local least squares solutions to outperform the truncated sparse matrices, the simulation predicts the opposite. Other than that, the plots show that the experimental data and simulation prediction agree very well.

Using the calculated residual wavefront errors we can predict the final image's Strehl ratio. The predicted Strehls and the real Strehls from the PHARO images for each sparse matrix are shown in Figure 3. The figure shows that within the fluctuations of data the predicted Strehls in general agrees with the real image.

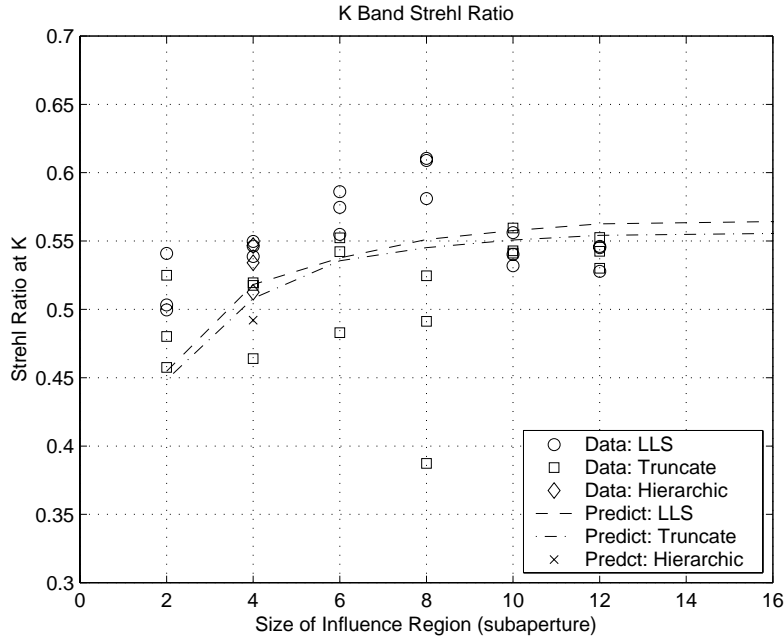


Figure 3. Strehl ratio versus matrix sparseness. Strehl ratio values from both the recorded PHARO images and predictions from the calculated wavefront errors are presented.

Figures 2 and 3 have not shown a difference in performance between the truncation, LLS, and hierarchic control matrices. However, as shown in Figure 4 and 5, more detailed analysis of the recorded wavefront sensor data does show that there are distinguishable differences.

As mentioned in section 2, because the sparse matrices reconstruct the wavefront from local information, correction of the low spatial frequency modes will be degraded. This effect can be seen in the recorded DM residuals, which are the closed-loop DM actuator update values (i.e. centroids multiplied by the reconstructor). Figure 4 shows the spatial power spectral density (PSD) of the DM actuator residuals for the LLS matrices with different sparseness as well as the hierarchic controller. The PSDs shown in the upper panel are radially averaged in each frame and time averaged over 10 second of recorded data. To better see the effect of sparseness of the matrices, the plots in the lower panel are the normalized PSDs which are normalized by the full matrix PSD.

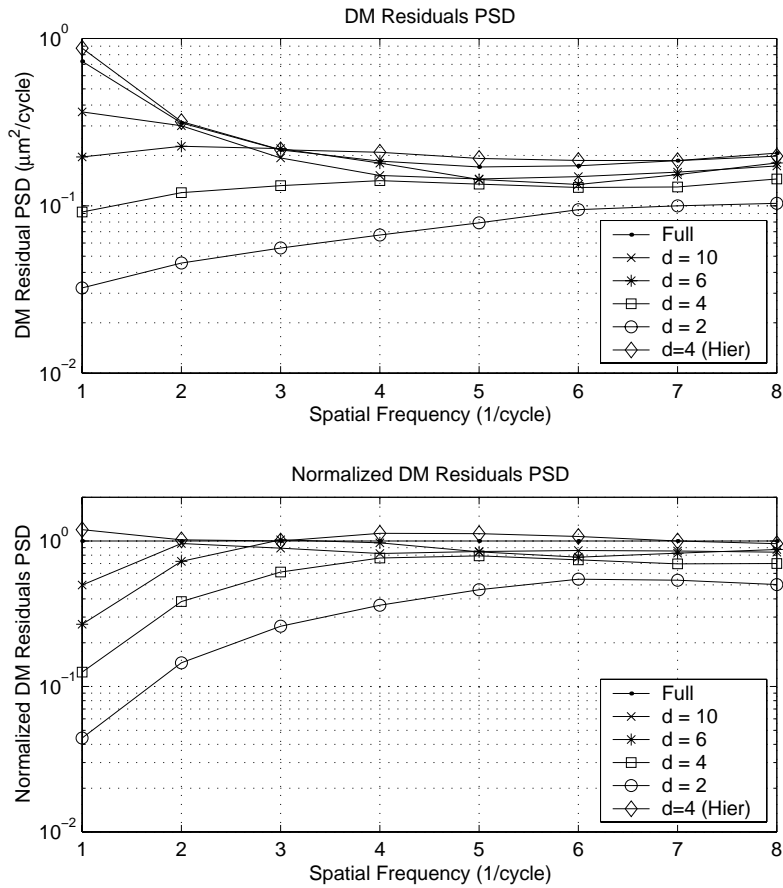


Figure 4. Closed-loop DM residuals PSD. *Upper panel:* DM actuator residuals PSD. The spatial frequency is in unit of cycle across the pupil. *Lower panel:* Normalized PSDs. The normalization is done by the PSD of the full matrix.

Figure 4 shows that as the influence region size d decreases there is less power in the DM residuals at low spatial frequencies, indicating a loss of low spatial order information in the reconstruction. Also, the normalized plots show that for different influence size d the PSD peaks at different spatial frequency. This indicates that the sparse matrix not only loses the gain for lower order modes but also increase the gain at certain characteristic frequency. This phenomenon agrees with our theoretical predictions shown in Figure 1. In contrast the DM actuator residuals from the hierarchic controller do not lose much power for the low order modes. Although it has local influence size of $d = 4$, compared to sparse matrix with the same influence size (the curve with connected squares), the global reconstruction layer in hierarchic controller has recovered most of the gain for the low order modes (the curve with connected diamonds). The DM actuator residuals PSD for the truncation sparse matrices are similar to that shown in Figure 4.

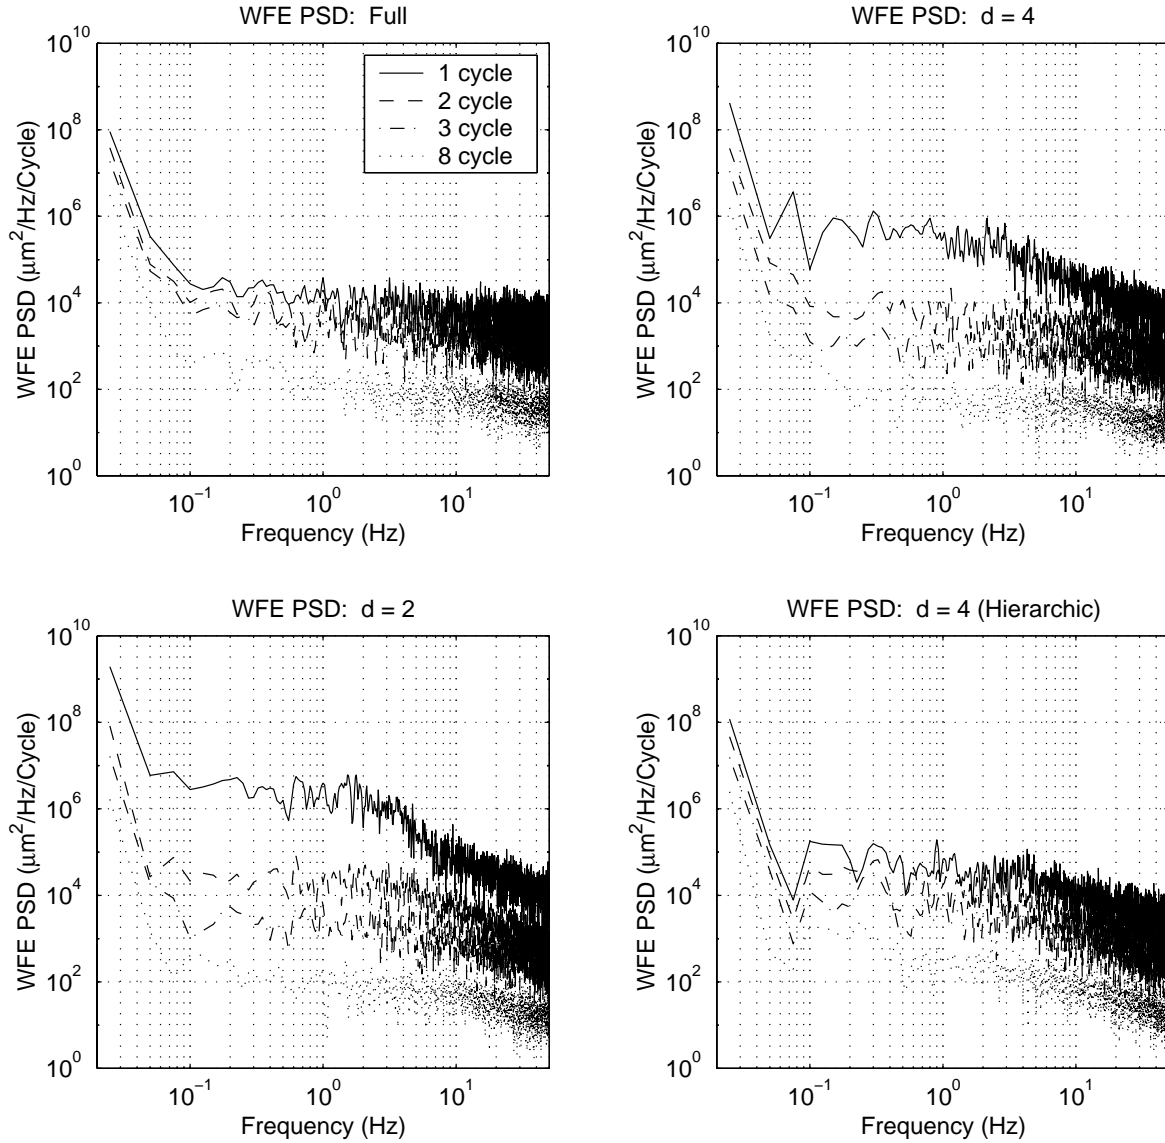


Figure 5. Temporal PSD for the closed-loop wavefront error under different sparse matrices. The temporal PSD is calculated for each spatial frequency and a few selected spatial frequencies are presented and indicated by different line styles. The influence size d of sparse matrix is shown in the title of each panel. To reduce the scattering and make the plots more readable the plots are smoothed for the points with frequency higher than 0.1 Hz.

Besides the PSD of DM residuals we have also looked into the spatial and temporal properties of the closed-loop wavefront errors under the sparse matrices. From the recorded centroid data the wavefront errors are reconstructed using the full matrix. The spatial PSD is then calculated for each frame and radially averaged. Then the temporal PSDs are calculated for each spatial frequency PSD over the 40 seconds of data. Plots in Figure 5 shows the temporal PSD for a few selected spatial frequencies. The temporal frequency ranges from 0.025 to 50 Hz, determined by the data-recording rate of 100 Hz and data length of 40 seconds. Each panel shows results from a sparse matrix with influence region size d indicated in the title of the plot. For a comparison, data from the full matrix is also shown in the upper left panel. As the matrix becomes sparse the low spatial order ($1 - 2$ cycles) wavefront error power increases, especially for the case of $d = 4$ and $d = 2$. For the case of $d = 2$, one can even argue that the cutoff frequency for the low spatial mode (1 cycle) has also been pushed lower to less than 10 Hz compared with the full matrix case. Again the recovery of global modes by the

hierarchic controller has maintained the power of low spatial modes, providing a performance close to the full matrix case (lower right panel). The temporal PSDs for the truncation sparse matrices have the similar features shown in Figure 5.

4. CONCLUSIONS AND FUTURE WORK

Real time wavefront reconstruction for future large aperture ground-based telescopes has posed a big challenge on real time computation for AO's wavefront reconstructor. We have proposed an approach to reduce the fully populated reconstruction matrix to a sparse matrix with a banded structure based on the limited influence of the DM actuators. Experiments and simulations have shown that as the matrix sparseness increases the accuracy of the wavefront reconstruction degrades. This has been understood to be caused by the loss of the global mode information. The loss of the lower mode gain in the controller can be predicted. The remedy for this is to add an extra layer of wavefront reconstruction from a much reduced sensor data set to sense only the low order modes. This hierarchic control method has shown to provide much better wavefront reconstruction without much increase of the calculation burden to the reconstructor (scaled as $n^{8/3}$ for the two-layer controller and $\sim n^2 \log_2(n)$ for multi-layer controller). Both simulation and experiment have proved its effectiveness.

The successful generation and implementation of all the sparse matrices, especially the hierarchic controller, have validated our models and simulations. The experimental data have agreed with simulation predictions. The modeling and simulation can provide a crucial tool in the design of sparse matrices for the future AO systems. The experiment on PALAO has been a critical step for us to understand the behaviors of these sparse matrices and will continue to provide the excellent testbed capabilities for the evaluation of new generation sparse matrices.

Sparse matrices tests on PALAO have shown that with the exception of extremely sparse matrices ($d \leq 4$) the dominant limitations of AO performance are still the AO system errors such as calibrations. The decrease of the closed-loop Strehls of images is not significant for the moderately sparsened matrices. Another observation from the experiment is that although in principle the truncated sparse matrices are not optimal compared with the localized least square matrices, the experiment data have shown little difference, which is contradicted to the predictions from model. One speculation is that it was caused by the imperfect gain tuning for the sparse matrices.

The modeling and experimental studies on the sparse matrix technique are still ongoing. We plan to study how to optimize the matrix sparseness taking into account the decreased computation delay together with the seeing condition. The matrix sparseness and the servo time delay are related variables in the system design. A sparser matrix requires less reconstruction calculations which will reduce the servo time delay, but will also introduce more wavefront errors. For any given AO system and atmospheric conditions there should be an optimal sparseness for the system. A realistic model will allow us to find this optimal solution. Another issue is to understand the performance of the sparse matrices on the laser guide star (LGS) AO system. Characteristics of LGS such as non-uniform sensor gain across the pupil may complicate the problem. Other approaches, such as variable influence size may need to be implemented to accommodate it. Finally, so far we have used only the least square method for the wavefront reconstructions in our sparse matrices experiment. When the guide star is faint the least square solution may not be optimal. An optimal sparse estimator may provide a better solution for the faint guide star.

5. ACKNOWLEDGEMENT

This work was performed in part at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Observations at the Palomar Observatory were made as part of a continuing collaborative agreement between Palomar Observatory and the Jet Propulsion Laboratory.

6. REFERENCES

- [1] Hardy, J. W., *Adaptive Optics for Astronomical Telescopes*, Oxford Series on Optical and Image Sciences, 1998.
- [2] Dekany, R., Nelson, J. E., and Bauman, B., "Design considerations for CELT adaptive optics," SPIE **4003**, Munich, Germany (2000).
- [3] Poyneer, L. A., Gavel, D. T., and Brase, J. M., "Fast wavefront reconstruction in large adaptive optics systems using the Fourier transform," *J. Opt. Soc. Am. (A)*, 2002.
- [4] Gilles, L., Vogel, C., and Ellerbroek, B., "A multigrid preconditioned conjugate gradient method for large scale wavefront reconstruction," submitted to *J. Opt. Soc. Am. (A)*, 2002.
- [5] Wild, W. J., Kibblewhite, E. J., and Vuilleumier, R., "Sparse matrix wave-front estimators for adaptive-optics systems for large ground-based telescope," *Optics Letters*, **20**, 955 – 957, 1995.
- [6] Murphy, T. P., Lyon, R. G., Dorband, J. E., and Hollis, J. M., "Sparse matrix approximation method for an active optical control system," *Appl. Opt.*, **40**, 6505 – 6514, 2001.
- [7] MacMartin, D. G., "Local and hierarchic control for adaptive optics," *AIAA Guidance, Navigation and Control Conference*, AIAA 2002-4939, 2002.
- [8] Truong, T., Brack, G. L., Trinh, T., Troy, M., Shi, F., and Dekany, R. G., "Real-time wavefront processors for the next generation of adaptive optics systems: a design and analysis," SPIE paper **4839**-109, Waikoloa, Hawaii, 2002.
- [9] Troy, M., Dekany, R., Brack, G., Oppenheimer, B., Bloemhof, E., Trinh, T., Dekens, F., Shi, F., Hayward, T., and Brandl, B., "Palomar adaptive optics project: status and performance," *Proc. SPIE*, **4007**, 31-40, Munich, Germany, 2000.