

Compensating atmospheric turbulence effects at high zenith angles with adaptive optics using advanced phase reconstructors

Grant Soehnel, Michael C. Roggemann, and Glen Archer *

August 29, 2007

Abstract

Atmospheric turbulence degrades the resolution of images of space objects beyond that predicted by diffraction alone. Adaptive optics telescopes have been widely used for compensating these effects, but as users seek to extend the envelopes of operation of adaptive optics telescopes to more demanding conditions, such as daylight operation and operation at low elevation angles, the level of compensation provided will degrade. We have been investigating the use of advanced wave front reconstructors and post detection image reconstruction to overcome the effects of turbulence on imaging systems in these more demanding scenarios. In this paper we show results comparing the optical performance of the exponential reconstructor, the least squares reconstructor, and the stochastic parallel gradient descent algorithm in a closed loop adaptive optics system using a conventional continuous facesheet deformable mirror and a Hartmann sensor. The performance of these reconstructors has been evaluated under a range of source visual magnitudes, and zenith angles up to 67 degrees. We have also simulated satellite images, and applied speckle imaging, multi-frame blind deconvolution algorithms, and deconvolution algorithms that presume the average point spread function is known to compute object estimates.

*All authors are with the Michigan Technological University Department of Electrical and Computer Engineering, 1400 Townsend Dr., Houghton, MI 49931, USA.

1 Introduction

The degradation in performance of imaging systems due to atmospheric turbulence is a well-known problem [1, 2]. The key limiting factor in the resolution of a telescope due to the hardware is diffraction from the finite size of the aperture. Increasing the size of the primary aperture reduces the effects of diffraction. However, corruption of an optical field that has propagated a significant distance through the atmosphere imposes a much greater restriction on resolution than diffraction alone for a large telescope. Varying temperature and pressure in the atmosphere results in a randomly varying index of refraction. The initial effect of this on a propagating optical field is to cause phase errors in the wavefront [2]. As optical fields travel long distances through relatively strong turbulence, the mechanics of wave propagation cause phase errors to produce amplitude errors known as scintillation. Under severe conditions with scintillation, the surface of constant phase known as the wavefront will become badly distorted and discontinuous [3].

An adaptive optics system attempts to compensate for turbulence induced field distortions by reducing phase errors in the field. An example of such a system is shown in figure 1. A deformable mirror makes a phase correction on the incident field, and a wavefront sensor measures a set of local wavefront phase differences on the residual field. There is always a small time delay between sensing and correcting the field [4, 5]. The goal of any adaptive optics phase reconstructor is to start with a vector \vec{s} of wavefront sensor phase difference measurements, and compute a set \vec{p} of relative phases to be used as actuator commands for the deformable mirror.

Traditionally, adaptive optics systems have been used under relatively weak turbulence conditions where scintillation effects are low. The least squares reconstructor has been the most widely used method in these conditions [4, 5]. One basic limitation in keeping turbulence effects minimal is that the zenith angle must be small. As the zenith angle increases, the path through the atmosphere becomes longer and turbulence effects become stronger. Discontinuities known as branch points in the wavefront can result, violating a primary assumption of the least squares reconstructor [6, 7]. There are reconstruction methods that attempt to locate branch points [8, 9, 10], but they are slow and complicated, and their performance in the presence of noise is not well understood. The exponential reconstructor is known to perform well in the presence of phase discontinuities without

the need to locate branch points [7]. Its implementation to an adaptive optics system with a Hartmann wavefront sensor, however, is not straightforward. Here we present a new way of using the exponential reconstructor on Hartmann data.

The new implementation of the exponential reconstructor has resulted in better performance in a closed loop simulation than both the least squares reconstructor and a previously proposed method of implementing the exponential reconstructor on Hartmann data. In addition, the new method of implementing the exponential reconstructor performs better by a larger margin over the other methods as noise levels increase and also as turbulence conditions worsen. This makes the method presented here a good choice for use in severe turbulence conditions and also in low light level conditions.

Advanced phase reconstructors can provide better phase compensation in severe conditions. However, an alternative to wavefront sensing and phase reconstruction is the stochastic parallel gradient descent (SPGD) algorithm [11]. The basis of this algorithm is to generate random perturbations on the deformable mirror shape and then test some measure of performance. New random perturbations are generated and tested until the process converges to an adequate deformable mirror figure. This algorithm is desirable because it does not require any wavefront sensing. The measure of performance is based on the total intensity in the aperture, so it is not as vulnerable to low light levels divided among many subapertures. One concern with this algorithm is its convergence speed. This algorithm has been tested in severe conditions by simulation, and also with a beam compensation experiment.

The remainder of this paper is organized as follows. Section 2 gives a basic description of the reconstructors studied. Section 3 describes the problem with the geometrical arrangement of the Hartmann wavefront sensor data with respect to implementation of the exponential reconstructor. A previously proposed method of splitting the data into two reconstructor grids is presented, along with the new method which keeps all of the data on a single grid. Section 4 describes the stochastic parallel gradient descent method. Finally, section 5 presents the simulation results.

2 Phase Reconstructors

The goal of a wavefront reconstructor is to use wavefront sensor phase difference measurements to obtain the phase at a set of points. A single sub-

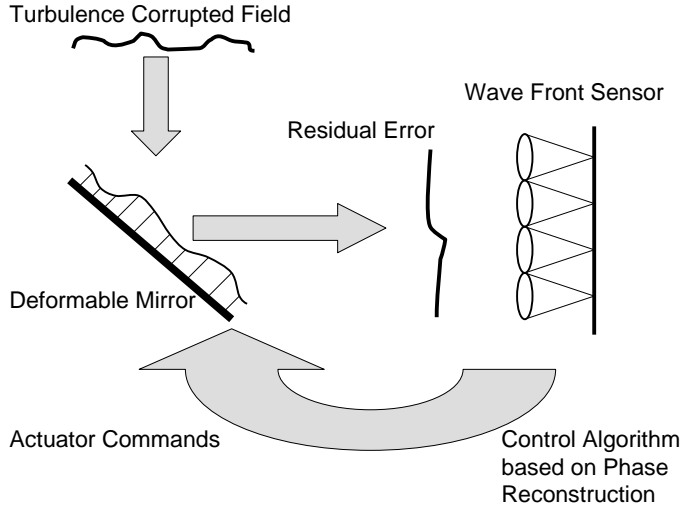


Figure 1: Closed Loop Adaptive Optics System

aperture of a Hartmann wavefront sensor is shown in figure 2. There are two phase difference measurements obtained from each subaperture, one in the x direction, and one in the y direction. For N subapertures, the vector \vec{s} contains all the phase difference measurements $[s_{x1} s_{y1} s_{x2} s_{y2} \dots s_{xN} s_{yN}]$. The actuator locations of the deformable mirror are positioned at each corner of the wavefront sensor subapertures. A vector $\vec{p} = [p_1 p_2 \dots p_M]$ must be obtained where M is the number of actuators.

A least squares reconstructor finds a solution to the problem of mapping \vec{s} to \vec{p} such that a particular measure of the mean square error is minimized. The basic idea behind a least squares reconstructor is to use the actuator and subaperture locations to construct a matrix M such that the slope measurements \vec{s} are related to the estimated phases \vec{p} by

$$\hat{p} = M\vec{s}, \quad (1)$$

where M is an inverse to a gradient matrix G that would solve the problem $\vec{s} = G\vec{p}$. The matrix M can be computed for any arrangement of phase difference measurements and actuator locations.

The exponential reconstructor is based on adding phase differences along a path. This confines it to a geometry in which phase differences link the actuator locations. In addition, the algorithm operates by reducing the grid




Figure 2: Single subaperture of a Hartmann wavefront sensor

size by a factor of 2 repeatedly, and then expanding it by a factor of 2 back to the original size. This requires the grid size for the exponential reconstructor to be $(2^N + 1) \times (2^N + 1)$, where N is a positive integer with phase differences linking those points. This results in the need for $(2^N + 1) \times (2^N)$ x-directed phase difference measurements and $(2^N) \times (2^N + 1)$ y-directed phase difference measurements. An example of this starting grid with $N = 3$ is shown in figure 3. The exponential reconstructor has three parts to its algorithm. In the reduce portion, the grid of phase differences is cut in size repeatedly such that a grid of $(2^N + 1) \times (2^N + 1)$ grid points is reduced to a grid of $(2^{N-1} + 1) \times (2^{N-1} + 1)$. Phase difference measurements in the reduce part of the algorithm are consolidated such that two phase differences along the x or y direction are replaced with one phase difference. This is done until the grid is only of size 2×2 . At this point the least squares solution to the 2×2 grid is calculated. The final step of the algorithm is the build part, in which the 2×2 grid of phase points is expanded such that at each step a $(2^N + 1) \times (2^N + 1)$ size grid becomes a $(2^{N+1} + 1) \times (2^{N+1} + 1)$ grid. This is done until the grid is back to the original size and contains phases at all of the original grid points. A detailed presentation of this algorithm can be found in reference [7].




Figure 3: Exponential reconstructor grid for $N = 3$

3 Exponential Reconstructor Geometry Implementation

One key challenge in implementing the exponential reconstructor is that it must operate on the geometry shown in figure 3. Phase difference measurements must link actuator points in a square grid. This is referred to as the Hudgin geometry [12]. For a Hartmann wavefront sensor, however, phase difference measurements are positioned across the center of a subaperture while the actuator locations are at the corners of a subaperture. This is called the Fried geometry [13]. Both of these geometries are shown for comparison in Fig. 4. In both drawings a single subaperture is contained inside a square with an actuator at each corner so that there are 4 subapertures shown. A sample g_x and g_y phase difference measurement or value is shown in each drawing.

The solution to adapting Fried geometry data to a Hudgin geometry reconstruction lies in recognizing that wave front sensor slope measurements can be determined with respect to any rotation of the coordinate system. The first step in creating a Hudgin geometry is to obtain phase differences in the diagonal directions of up-right and up-left gradients [7]. This is shown

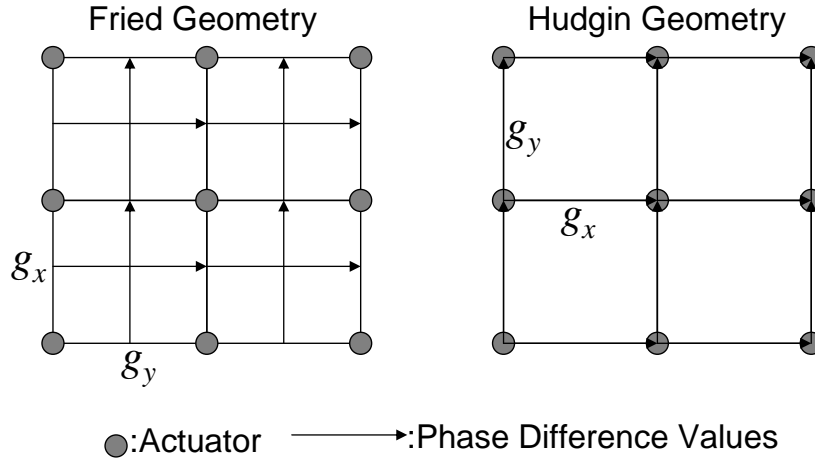


Figure 4: Comparison of the Fried and Hudgin Geometries

in figure 5.

Diagonal phase differences connect actuator locations, so if the coordinates are rotated by 45 degrees the result is identical to the Hudgin geometry. There still remains a problem in that the diagonal phase differences cross each other. The solution proposed by Fried is to split the set of actuator locations and phase difference measurements into two separate grids that can be solved separately [7]. Figure 7 shows this two-grid geometry. Actuators on one of the interlaced grids occupy every other position on the full grid of actuators. There is difficulty, however, in combining the two grids because of waffle error in the reconstruction. The waffle mode in its most basic form is a checkerboard pattern of positive and negative phases across all the actuators. A portion of this pattern containing four subapertures is shown in figure 6. Note that the phase difference measurements due to the waffle mode is shown to be zero. This means that the waffle mode is not detected by the Hartmann wave front sensor. Very little waffle mode would actually exist in a propagated field, but waffle mode can creep into the phase reconstruction in a closed loop adaptive optics system and continue to grow because waffle is unobserved in the wave front sensor. This is a well known problem, especially when converting from the Fried geometry to the Hudgin geometry [14, 15, 16].




Figure 5: Diagonal Phase Difference Geometry




Figure 6: Drawing and Analysis of Waffle Error




Figure 7: Two grid geometry: Open circles and dashed lines are the first grid, colored circles and solid lines are the second grid

Methods for removing waffle error involve either spatial filtering or removing the components of the waffle mode in the frequency domain. When noise is present, as it is in the simulations described here, waffle error is difficult to remove [16]. An example of the reconstructed phase from the two grid geometry is shown in figure 8. It is clear that the waffle mode is present with a varying sign and amplitude across the aperture. A solution to the waffle error was devised such that waffle error doesn't quickly appear in the reconstruction. The first step is to replace each phase difference measurement with two phase differences that each span half the length of the original. These two phase differences will also be half the value of the original as well. This allows a single reconstruction grid to be defined rather than two independent grids. A diagram of this reconstruction grid is shown in figure 9. Using this single grid creates a set of phase points and phase differences that contains phase points at all of the actuator locations along with some phase points that are not at an actuator location. These phase points not at an actuator location in the reconstruction grid are at the center of each subaperture. The reconstruction can then be performed on the single grid, and the phase points that are not at an actuator location can be ignored when matching the reconstructed phases to actuator commands. It was found that this method results in better performance and virtually no waffle error as compared to reconstructing the wave front on two disjoint grids.

4 Stochastic Parallel Gradient Descent: SPGD

The SPGD algorithm for wavefront correction is quite different from conventional wavefront sensing and phase reconstruction techniques. The SPGD approach to deformable mirror control attempts to iteratively optimize some measure of performance by operating on the actuator control signals and some measure of performance. The performance measure, which we will represent with the symbol J , is a function of all the actuator positions. In practice, J is a scalar function of the deformable mirror actuator positions that is determined to be a good measure of performance. In this case encircled energy with an exponential weighting was chosen. The exponential weighting favors the intensity close to the center and only intensity in a chosen radius is integrated. The exact weighting function equation used in the simulation is $\exp(-r/3), r \leq 3$.




Figure 8: Phase Reconstruction with Waffle Mode Present

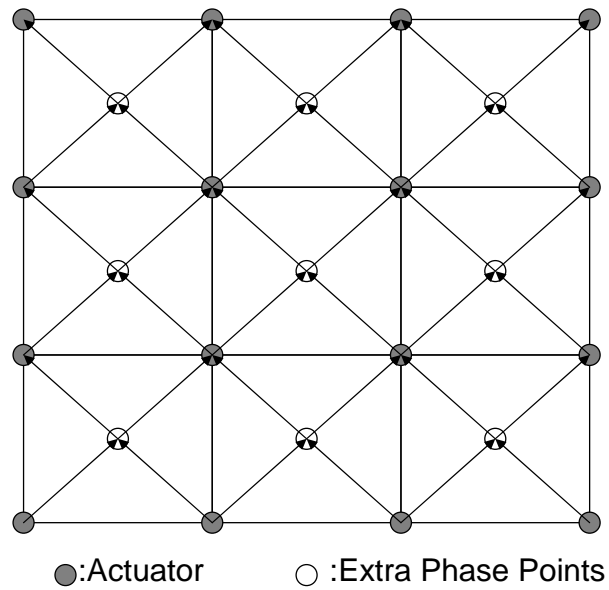


Figure 9: Single grid geometry

The algorithm begins by randomly generating a set of perturbations δu to add and then subtract from the control signals. In this study each actuator perturbation has a magnitude of $+a$ or $-a$ with equal probability. Letting f be the function of the actuator commands that gives the performance metric J , we define

$$J^+ = f(u + \delta u) \quad (2)$$


$$J^- = f(u - \delta u) \quad (3)$$

for the randomly chosen perturbation δu . J^+ and J^- are computed in order to determine what the slope of J is in the direction of δu around the current actuator commands u . The commands u^{n+1} are updated from the current commands u^n by [17]


$$u^{n+1} = u^n + \gamma(J^+ - J^-)\delta u \quad (4)$$

where γ is referred to as the descent strength. γ can be chosen to be a constant through all the iterations, or some choices can be made to vary it. To illustrate considerations for choosing γ and a as they relate to equation 4, consider figure 10. Figure 10 is a plot of the performance metric in a randomly generated direction δu as discussed as the magnitude a in that direction varies around zero in both the positive and negative directions. Two different trials are shown to illustrate different possible curves for the J function as the magnitude along a fixed perturbation direction varies. Note that the J curves shown in the figure are clearly not linear. This means that the calculation of the slope of the J function $\delta J = (J^+ - J^-)$ is only valid for some range about the starting point. The magnitude a of the random perturbations should be chosen so that there is not too much curvature in the J function between $+a$ and $-a$. Also, γ should be chosen to generally keep the magnitude $\gamma(J^+ - J^-)$ multiplied by δu in the update step of equation 4 less than 1. Due to the curvature and nonlinear nature of J , confidence in J is low beyond the magnitude of the perturbation tested. In both simulation and experiment, if the command update magnitude $\gamma(J^+ - J^-)$ is greater than 1 at any iteration, then γ is reduced by .75.

An experimental setup has been constructed to test the SPGD algorithm. The layout of this experiment is shown in figure 11. The phase distortions are provided by a spatial light modulator (SLM) from Boulder Nonlinear Systems and their Matlab software development kit. A phase screen is generated in Matlab and then written to the SLM. A columnated laser beam is reflected



(a)



(b)

Figure 10: Performance metric J as the magnitude along a particular direction of phase perturbations as δu varies. Two examples are shown.




Figure 11: Block diagram of an SPGD simulation

from this SLM. The second SLM is used to correct the beam which was distorted by the first SLM. Two lenses are placed between the SLM's as shown in order to ensure that the field falling on the correcting SLM matches the field falling on the distorting SLM. The SLM's have 512×512 independent control channels which is much larger than the desired number of variables the SPGD algorithm will operate on. Instead of using the full resolution of the SLM, 32×32 images are written to the correcting SLM. The software automatically scales the smaller images to fit the SLM. This allows for 1024 independent control channels which is similar to the number used for the simulations. A CCD camera is placed in the focal plane of the outgoing beam. The camera image is then used to compute the performance metric for the algorithm.

5 Simulation Results and Discussion

Four methods of wavefront reconstruction were simulated for comparison: (1) the least squares reconstructor, (2) the exponential reconstructor working on the new single grid, (3) the exponential reconstructor working on the two grid geometry, and (4) a combination of the least squares and exponential reconstructor suggested in Ref. [18] known as the slope discrepancy method. This method involves performing a least squares reconstruction, differentiating the resulting phase, and using the exponential reconstructor on the residual phase differences obtained. The reconstructors were tested for fields propagating through the atmosphere, arriving at the top of the atmosphere at zenith angles of 0, 30, 60, and 67 degrees. The amount of light available to the wavefront sensor helps to determine how much noise is present in the measurements. Equation 5 was used to calculate the standard deviation of the noise [2]. The parameter η represents the imperfections of the detector and is set to 1.35 for these simulations. The parameter \bar{K}_W is the number of photons falling in each subaperture during the measurement integration time. Photon counts were obtained from calculations using a $1ms$ integration time on a source of visual magnitude 2, 4, and 6. The transmission loss at each zenith angle was also accounted for. Trials were done with no noise and increasing noise based on the source visual magnitude of 2, 4, and 6.

$$\sigma_n = \begin{cases} \frac{0.86\pi\eta}{(\bar{K}_W)^{1/2}r_0}rad/m & d > r_0 \\ \frac{0.74\pi\eta}{(\bar{K}_W)^{1/2}d}rad/m & d \leq r_0 \end{cases} \quad (5)$$


Plots of the Strehl ratios for all the trials simulated are shown in figures 12 and 13. The best method of wavefront reconstruction with Hartmann wavefront sensor data was found to be the exponential reconstructor operating on the new single grid geometry. This is true at every zenith angle and noise level. There is also a trend that the exponential reconstructor outperforms the least squares reconstructor by a greater margin as the zenith angle increases and as the noise levels increase. This shows that while the exponential reconstructor is always superior, its advantage is even better the more severe the conditions get.

The other two advanced methods of wavefront reconstruction had varied success. The two grid exponential reconstructor tends to only outperform the least squares reconstructor when the zenith angle increases, and the noise level is high. This is because in severe conditions the advantage of


using the exponential reconstructor outweighs the disadvantage introduced by waffle error. The slope discrepancy reconstructor does at least as well as the least squares reconstructor, and tends to do best when the noise level is low. This makes sense because the slope discrepancy reconstructor does two sequential reconstructions based on noisy data. The effect of noise on the least squares reconstruction propagates into the exponential reconstruction. These trends are interesting, but the ultimate result is that the single grid method of using the exponential reconstructor is clearly the best option. It is especially encouraging that the margin of increased performance gets larger the zenith angle and noise increase. This implementation of the exponential reconstructor will expand the range of conditions in which adequate phase correction can be accomplished.

After determining the parameters to be used, the SPGD algorithm was simulated for 100 independent realizations of the atmosphere at $\theta_z = 67$. The resulting radially averaged PSFs are shown in figure 14. The Strehl ratios obtained from the simulations are 29.18% for the standard Maui3 turbulence profile and 14.85% for twice the Maui3 profile. With the standard Maui3 profile the performance of the SPGD algorithm is almost always worse than the reconstructors studied. It only performs better at the highest noise levels. The least squares and exponential reconstructors were simulated at twice the Maui3 profile to compare with the SPGD algorithm as well. The resulting Strehl ratios at twice the Maui3 profile are shown in figure 15. Note that the SPGD algorithm does outperform the reconstructors with the single exception of the exponential reconstructor with no noise. The reconstructors are expected to perform very poorly under these conditions as the Fried parameter is $r_0 = 6.61\text{cm}$ for $\theta_z = 67$ at twice the Maui3 profile. The subaperture size for the wavefront sensor is 11.25cm , so the field is clearly undersampled. If a new wavefront sensor and deformable mirror were designed with sampling on the order of $r_0 = 6.61\text{ cm}$, a large increase in performance is likely to be seen. The SPGD algorithm doesn't decrease in performance as rapidly with more severe field corruption because it doesn't depend on a wavefront sensor. Note that doubling the turbulence strength roughly cuts the Strehl ratio in half for the SPGD algorithm. The Strehl ratio is reduced by more than half when using the reconstructors. This suggests a possibility that the SPGD algorithm could be advantageous in using existing hardware to perform corrections in conditions that would be too severe for phase reconstructors.

The experimental setup for testing the SPGD algorithm was used to test




(a)




(b)

Figure 12: Strehl ratios of the reconstructors with: (a) No noise; (b) Magnitude 2 source;



(a)



(b)

Figure 13: Strehl ratios of the reconstructors with: (a) Magnitude 4 source; (b) Magnitude 6 source;

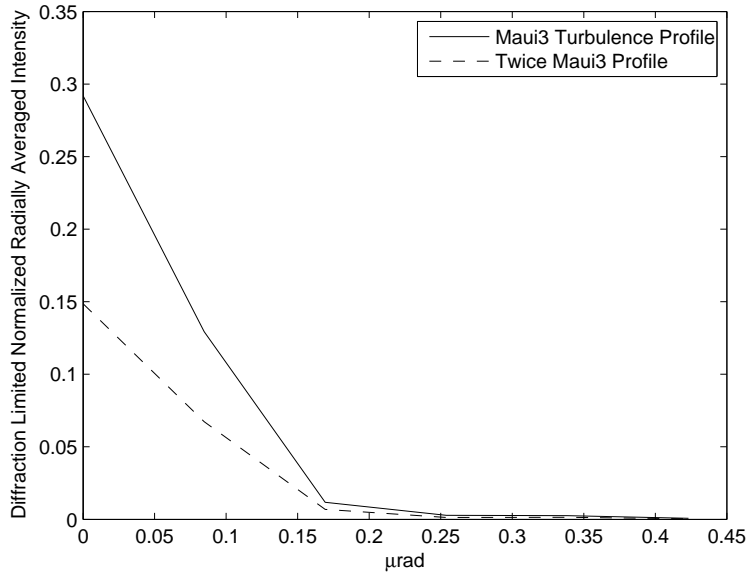


Figure 14: Average PSF over 100 iterations of the SPGD algorithm with $\theta_z = 67$ for the Maui3 turbulence profile and also double the Maui3 profile.

	No Noise	Magnitude 2 source	Magnitude 4 source	Magnitude 6 source
Least Squares	9.2	6.82	9.47	3.32
Exponential	15.24	13.89	12.51	6.24

Figure 15: Strehl ratios obtained for the exponential and least squares reconstructors at $\theta_z = 67$ and twice the Maui3 turbulence profile

the algorithm's convergence with moving phase screens. A 1024×1024 phase screen was generated that has $D/r_0 = 36$. The phase screen written to the distorting SLM is a 512×512 subarray from the larger 1024×1024 phase screen. The distorting SLM is timed to update once per iteration of the algorithm such that the section of the large phase screen written to it progressively moves from left to right. This was done for a range of speeds from 4 iterations per pixel shift to 50 iterations per pixel shift. Linear interpolation was used to update the SLM with a new partially shifted phase screen every iteration of the algorithm. For a single shifting phase screen, the resulting converged spot was averaged for a series of 6 images separated by 100 iterations. This was done well after the algorithm initially converged. This process was then repeated for 30 independently generated phase screens. The final result is a compensated PSF obtained from averaging all 180 images obtained as described.

In order to compare performance, the ratio of intensity inside the first zero of the diffraction limited spot to the total intensity was computed. This was chosen because it is a measure relative to the total intensity. Peak intensity and even total intensity were found to vary too much to be used for comparisons. This percentage of encircled energy for each trial is shown in figure 16. For this plot, the range of error is shown representing two standard deviations from either side of the mean. In addition figure 17 shows the radially averaged PSFs obtained from each speed of moving phase screens along with uncompensated and diffraction limited PSFs. Note that the uncompensated PSF is much worse than all the compensated PSFs and also the compensated PSFs approach the width of the diffraction limited PSF. From these plots it can be seen that the tails of the PSF become close to zero and the encircled energy rises above 50 percent at around 20 to 30 iterations per pixel shift of the phase screen. The sample spacing of the phase screens is made to represent sample spacing of 7.5mm in a 3.6m telescope pupil. Typical wind speeds of 7.5m/s would correspond to a pixel shift every millisecond. The results demonstrate that very good compensation could be attained at iteration speeds of $20 - 30\text{kHz}$ in order to get 20 to 30 iterations per pixel shift. Significant compensation can even be attained down to 4 iterations per pixel shift which might correspond to 4kHz .

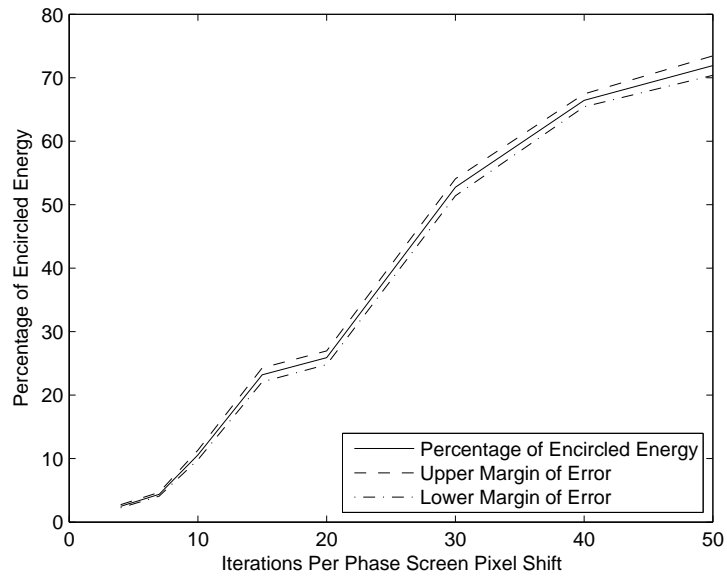


Figure 16: Percent of energy within the first zero of the diffraction limited PSF for different speeds of moving phase screens

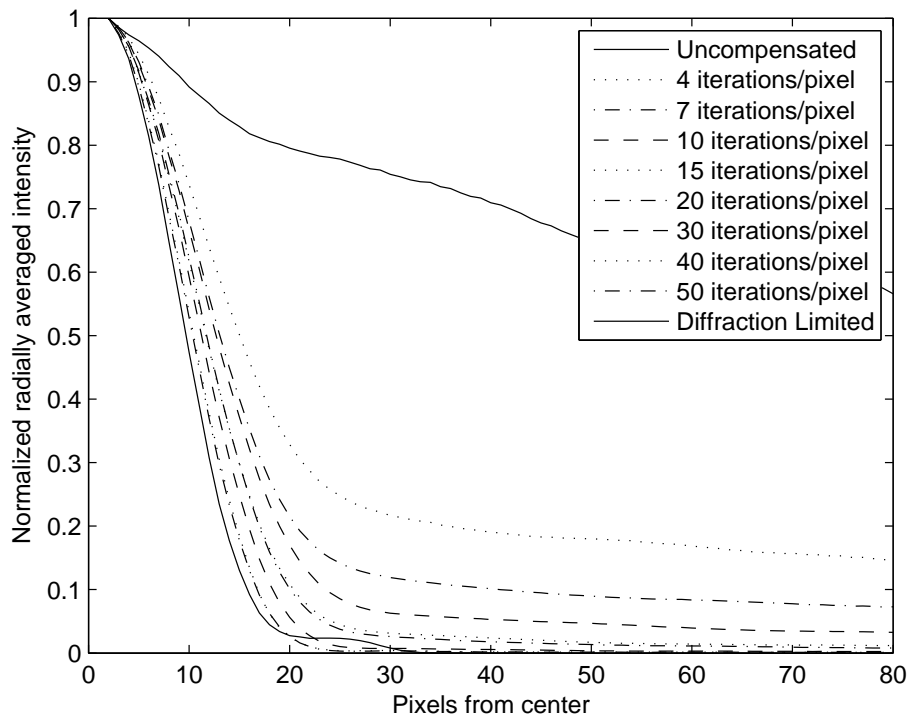


Figure 17: Radially averaged PSF for different speeds of moving phase screens

6 Conclusion

This work has demonstrated a wide variety of phase compensation techniques. As for wavefront reconstruction, the exponential reconstructor is showing a lot of potential for use in severe turbulence and high noise levels. The new method of using it with Hartmann wavefront sensor data is also very feasible. Splitting the measurements in half is a simple solution. It does not gain any resolution, but it keeps the data in a single reconstruction problem and helps the reconstructor perform in closed loop by greatly reducing waffle error. The simulations have shown this implementation has at least modest improvement over all the other methods and in all the conditions tested.

The SPGD algorithm is very simple to implement, and it has shown some usefulness. It is not likely to outperform a well sampled wavefront sensor and phase reconstruction system. However, severe turbulence and low light levels can get very taxing on a wavefront reconstructor. The only sensing for the SPGD algorithm is a single measurement of the entire aperture in order to obtain a performance measure. The aperture does not need to be divided into many subapertures for wavefront sensing. This is very helpful in low light levels when each wavefront sensor subaperture would only get a fraction of the little light available. To add to the problem of having many subapertures, as turbulence conditions worsen, more subapertures are needed to adequately sample the field. The SPGD simulations have shown performance comparable to wavefront reconstructors in the most severe noisy conditions. Also, the experiment has shown convergence of the algorithm in changing turbulence. The main consideration for SPGD is whether the convergence is fast enough to overcome the rate of changing turbulence conditions. The results show that improvement can be gained with changing turbulence, and that improvement can be quite good if the algorithm is able to run fast enough.

References

- [1] V. I. Tatarskii. *Wave Propagation in a Turbulent Medium*. Dover Publications, Inc., New York, 1961.
- [2] M. C. Roggemann and B. M. Welsh. *Imaging Through Turbulence*. CRC Press, Boca Raton, FL, 1996.
- [3] David L. Fried and Jeffrey L. Vaughn. Branch cuts in the phase function. *Applied Optics*, 31:2865–2882, 1992.
- [4] Robert K. Tyson. *Principles of Adaptive Optics, Second Edition*. Academic Press, Inc., Boston, MA, 1998.
- [5] John W. Hardy. *Adaptive Optics for Astronomical Telescopes*. Oxford University Press, New York, 1998.
- [6] D. L. Fried. Branch point problem in adaptive optics. *J. Opt. Soc. Am. A*, 15:2759–2768, 1998.
- [7] D. L. Fried. Adaptive optics wave function reconstruction and phase unwrapping when branch points are present. *Opt. Commun.*, 200:43–72, 2001.
- [8] D. C. Ghiglia and Mark D. Pritt. *Two-Dimensional Phase Unwrapping*. John Wiley and Sons, Inc., New York, 1998.
- [9] M. C. Roggemann and A. C. Koivunen. Branch point reconstruction in laser beam projection through turbulence with finite degree of freedom phase only wave front correction. *J. Opt. Soc. Am. A*, 17:53–62, 2000.
- [10] M. C. Roggemann and A. C. Koivunen. Wave front sensing and deformable mirror control in strong scintillation. *J. Opt. Soc. Am. A*, 17:911–919, 2000.
- [11] M. A. Vorontsov, G. W. Carhart, M. Cohen, and G. Cauwenberghs. Adaptive optics based on analog parallel stochastic optimization: analysis and experimental demonstration. *J. Opt. Soc. Am.*, 17:1440–1453, 2000.
- [12] R. H. Hudgin. Wave-front reconstruction for compensated imaging. *J. Opt. Soc. Am.*, 67:375–378, 1976.

- [13] D. L. Fried. Least-square fitting a wave-front distortion estimate to an array of phase-difference measurements. *J. Opt. Soc. Am.*, 67:370–375, 1976.
- [14] Sivaramakrishnan, Makidon, Perrin, Roberts, Oppenheimer, Soummer, and Graham. An analysis of fundamental waffle mode in early aeos adaptive optics images. *The Astronomical Society of the Pacific*, 117:831–846, 2005.
- [15] Donald T. Gavel. Suppressing anomalous localized waffle behavior in least square wavefront reconstructors. *Society of Photo-Optical Instrumentation Engineers*, 2002.
- [16] Lisa A. Payneer, Donald T. Gavel, and James M. Brase. Fast wavefront reconstruction in large adaptive optics systems using the fourier transform. *Astronomical Telescopes and Instrumentation*, 2002.
- [17] Piotr Piatrou and Michael Roggemann. Beaconless stochastic parallel gradient descent laser beam control. numerical experiments. *Optical Society of America*, 2006.
- [18] 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.