UCRL-TR-203128

LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Time series analysis of Adaptive Optics wave-front sensor telemetry data

L. A. Poyneer, D. Palmer

March 24, 2004

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Time series analysis of Adaptive Optics wave-front sensor telemetry data

Lisa A. Poyneer and Dave Palmer

Abstract

Time series analysis techniques are applied to wave-front sensor telemetry data from the Lick Adaptive Optics System. For 28 fully-illuminated subapertures, telemetry data of 4096 consecutive slope estimates for each subaperture are available. The primary problem is performance comparison of alternative wave-front sensing algorithms. Using direct comparison of data in open loop and closed-loop trials, we analyze algorithm performance in terms of gain, noise and residual power. We also explore the benefits of multi-input Wiener filtering and analyze the open-loop and closed-loop spatial correlations of the sensor measurements.

1. Introduction

Wave-front sensing is an important area of adaptive optics. Current methods are susceptible to non-uniform behavior as system conditions change, leading to reduced system performance. We have developed an algorithm that potentially has much more uniform behavior, which will lead to improved system performance. Over the past several months we have been doing short experiments with it in a top-flight Adaptive Optics system. In this project we apply time series analysis to improve our evaluation of experimental data, with the goal of being able to clearly show whether or not our new algorithm will improve performance. This project has shown that though such a definitive statement is not yet possible, there is real evidence that our new algorithm has performance benefits in specific situations. The time series techniques also provide valuable guidance in designing future experiments.

2. Adaptive Optics

In many applications of optical systems, the observed field in the pupil plane has a non-uniform phase component. This deviation of the phase of the field from uniform is called a phase aberration. In imaging systems this aberration will degrade the quality of the images. In the case of a large astronomical telescope, random fluctuations in the atmosphere lead to significant distortion. These time-varying distortions can be corrected using an Adaptive Optics (AO) system, which is a real-time control system composed of optical, mechanical and computational parts. The rest of this section and the following section provide a brief overview of the system components and the relevant problems of interest for this project. Adaptive optics is also applicable to problems in vision science, laser propagation and communication. For a high-level overview, consult this web site.[?] For an in-depth treatment of the astronomical case, consult these books.[?,?]

Our tests were conducted at Lick Observatory on Mt. Hamilton, near San Jose, CA. The 120-inch (3-meter) Shane telescope has an AO system which was developed by LLNL and is jointly maintained with the observatory. In the Lick AO system the phase aberration is measured using a wave-front sensor. The sensor divides the pupil of the telescope into smaller sections (subapertures)

by a small array of square lenslets. Each lenslet forms an spot image of the reference source on a small region of a CCD. At Lick each subaperture forms an image onto a 4 by 4 pixel region. As the phase aberration of the field in the subaperture changes, the spot moves. This motion is directly related to the spatial slope of the phase across the subaperture. By estimating this movement of the spot in both directions on the CCD, the slope of the wavefront in the subaperture can be estimated. These slopes are then sent to a reconstruction algorithm, which generates the phase across the entire pupil based on the local slopes. This reconstructed phase is then applied to a flexible optic (a deformable mirror) which conjugates the phase aberration of the field. This process is done in closed loop and at rates of up to 500 Hz.

3. Wave-front sensing algorithms

It is the process of wave-front sensing that is of particular interest to our research. The wave-front sensing algorithm estimates the wave-front slope from only a few noisy pixels of information on the CCD. High-quality slope estimation is the key first step for a well-performing AO system. The algorithm should be unbiased and have as low noise as possible. Furthermore, it should be robust to changes in observation conditions. The standard WFS algorithm for astronomical AO (the Quadcell method described in detail below) has very low noise, but variable performance as observation conditions change.[?] We have developed an algorithm that is robust to changing conditions within a certain range on the detector. These experiments have been conducted to evaluate the performance of this algorithm in comparison to current methods.

First, we need to discuss the various sources which are used in the AO system. The reference source for the WFS is called the guide star. In the usual case a star is used, which is called the Natural Guide Star (NGS). If a target other than a star is imaged (e.g. a galaxy) a suitable NGS may not be available close to the target. In this case AO will not work unless an artificial star can be created. At Lick this artificial star is made by propagating a sodium (Na) laser into the atmosphere. A layer of positive Na ions exists about 90 km up in the earth's atmosphere. At the layer the laser excites the ions and they emit light at 589 nm. If the laser is well aimed, this region of light emission is small and forms an artificial star, called the Laser Guide Star (LGS). The size of the NGS on the CCD is quite small, but can vary with atmospheric conditions. The size of the LGS on the CCD is much larger and also can vary significantly with conditions. The density and thickness of the sodium layer has an effect, but more importantly the atmospheric conditions (termed the "seeing") can drastically change the spot size of the LGS.

Three different algorithms are examined in this project. All are implemented in the Lick AO control system. The current algorithm used at Lick is called "Quadcell". This is the standard WFS configuration and algorithm used in astronomical AO. It uses the center 2 by 2 pixel region. The x-shift of the spot is simply the difference in counts between the two pixels on the right side and the two on the left, divided by the total number of counts in all four pixels. The quadcell is a special case of a more general algorithm called "Centroiding" which computes the center of mass of the spot image on the CCD and uses that to determine the shift. Provided the spot is small, this provides the lowest noise. However, as the spot size changes, the gain of the estimate changes. Because the spot is so much bigger in LGS mode, the "Binned" algorithm is used. In this case it uses the quadcell formula, but with pixels twice as big. Each large pixel is actually the sum of the

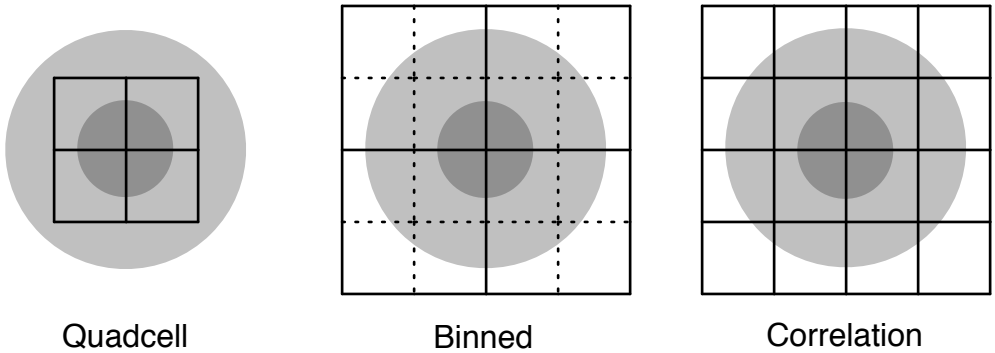

Fig. 1. WFS pixel configurations for the three algorithms, for the 4 by 4 pixel area for each subaperture spot. The effective pixels are outlined in solid lines. The gray circles represent small NGS spots and large LGS spots.

four real pixels in the region. This method will capture signal from a large spot, but at a penalty of higher noise. The new algorithm we are studying is called "Correlation". It is based on a more general WFS algorithm which estimates the shift by correlation with a reference. In this special case it is a compact ratio of sums of pixels in the 4 by 4 pixel region. Its main advantages are that it is robust to changing spot size and does not suffer from significant noise due to using the extra pixels. The pixel configurations and sizes for these three methods are shown in Figure **??**.

4. Experimental setup

The AO control system produces bursts of telemetry for 4096 consecutive time steps. The control rate is variable. For this project we will only analyze the WFS slope telemetry, which is the x- and y-slope estimates for each of forty subapertures. Of these forty subapertures, twelve are on the edges and have noisier behavior, so we will only consider the twenty-eight with full illumination. Two different system modes were used in testing. The first is regular operation. The loop can either be open, where the system observes but does not correct the atmosphere, or closed, where the system is actively correcting the the residual error is observed. The second mode is a special comparison mode which works in open-loop only. In this case the control system computes the slopes using two algorithms on the same CCD data. The y-slopes for each algorithm are output in the telemetry file. This allows direct comparison of two different noisy measurements of the same data.

Using these data, we address two questions. First, direct comparison of the WFS algorithms is done with data from the comparison mode. In this case we estimate noise levels, gains and do optimal estimation on the underlying slope signal that is observed. Doing so helps us evaluate which algorithm is the best option. Second, we use both open-and closed-loop data to analyze the spatial coherence between subaperture signals. We are interested in analyzing how the open-loop phase aberration is auto-correlated across the telescope pupil and how well the closed-loop control corrects this wavefront. Analysis of both power and spatial correlations of the slope series provides information about turbulence strength and spatial coherence. In closed-loop operation the slope series provide information about the residual error in the system. Power spectra and correlations

can provide information about how well the system is performing (and how well each algorithm performs).

5. Time series analysis techniques

A. Model

For the discussion of relevant time series analysis techniques,[?] some notation should first be defined. For the ideal time series, the notation $s(t)$ will be used. The sampled data will be stacked as a column vector $\mathbf{s}$, and indexed with $\mathbf{s}[1]$ to $\mathbf{s}[n]$, where $n = 4096$. Constant parameters will be represented as the gain g and the noise variance σ_w^2. The slopes will be represented by the letters q, b and c for Quadcell, Binned and Correlation. For example, a single observation is modeled as

$$q(t) = s(t) + v(t). \tag{1}$$

Due to the noise characteristics on the CCD, it is reasonable to model the slope measurement noise $v(t)$ as zero-mean, white Gaussian noise. The vector version is simply

$$\mathbf{q} = \mathbf{s} + \mathbf{v}. \tag{2}$$

When we have two measurements of the same signal, the multivariate model is

$$\begin{bmatrix} q(t) \\ c(t) \end{bmatrix} = \begin{bmatrix} g \\ 1 \end{bmatrix} s(t) + \begin{bmatrix} v(t) \\ w(t) \end{bmatrix}, \tag{3}$$

for example when we compare the Quadcell to Correlation. Note that we can only really estimate the ratio of the gains of the measurements, so a single gain g suffices. Because the slope noise arises from the noise on the CCD pixels, both methods have the same underlying noise source and as a rule we cannot assume that the noise series $v(t)$ and $w(t)$ are uncorrelated. Both, however, are still zero-mean, white Gaussian series. Furthermore we will assume that the slope signal is a stationary process.

B. Spectral estimators

The power spectrum of a time series $s(t)$ is denoted by $f_s(\nu)$. To approximate it we will use the periodogram approach over n discrete frequencies ν_k. The discrete fourier transform of a series will be denoted by the capital letter, as

$$S(\nu_k) = \frac{1}{\sqrt{n}} \sum_{t=1}^{n} \mathbf{s}[t] \mathrm{Exp}[-j2\pi\nu_k t]. \tag{4}$$

The periodogram $I_s(\nu_k)$ is $|S(\nu_k)|^2$. The smoothed spectral estimator over a band of L frequencies is given by

$$\hat{f}_s(\nu_k) = \frac{1}{L} \sum_{l=-(L-1)/2}^{(L-1)/2} I_s\left(\nu_k + \frac{l}{n}\right). \tag{5}$$

We can also estimate the smoothed cross-spectrum between two series

$$\hat{f}_{qc}(\nu_k) = \frac{1}{L} \sum_{l=-(L-1)/2}^{(L-1)/2} Q\left(\nu_k + \frac{l}{n}\right) \overline{C\left(\nu_k + \frac{l}{n}\right)}, \tag{6}$$

and from that obtain the smoothed squared-coherence

$$\hat{\rho}^2_{qc}(\nu_k) = \frac{\left|\hat{f}_{qc}(\nu_k)\right|^2}{\hat{f}_q(\nu_k)\hat{f}_c(\nu_k)}. \tag{7}$$

Note that the smoothing assumes spectral flatness in the region being smoothed. This is not always the case for our signals, especially the open-loop PSD of the slopes. Unfortunately, the data are noisy enough the smoothing is necessary to produce a reasonably-shaped spectrum. In particular, without smoothing the squared-coherence estimates are highly unreliable. For these data sets $L = 21$ qualitatively provided the best results ($n = 4096$).

C. Stochastic signal case

For most of our data the true slope signal $s(t)$ is stochastic. In this case, Wiener filtering is an appropriate technique. First we consider the standard single observation in noise case. Given our model $q(t) = s(t) + v(t)$, we desire the best linear filter to apply to $q(t)$ to generate the estimated true signal $\hat{s}(t)$. In the time domain the ideal filter is given by

$$\hat{s}(t) = \sum_{r=-\infty}^{\infty} a(r)q(t-r), \tag{8}$$

or in the frequency domain by the filter

$$\hat{S}(\nu) = A(\nu)Q(\nu) \tag{9}$$

using the convolution theorem. Switching to finite-length measurements, the best filter is given by

$$A(\nu_k) = \frac{\hat{f}_s(\nu_k)}{\hat{f}_s(\nu_k) + \hat{\sigma}^2_v}. \tag{10}$$

This is the standard Wiener filter as described in Section 3.9. The signal and noise spectra must of course be estimated; this is discussed below. For the comparison case (following the methods of Section 5.4 and 5.5), we can derive a joint Wiener filter for optimal estimation using two observed series. In this case we again use a linear filter for each signal to get the estimated true signal

$$\hat{s}(t) = \sum_{r=-\infty}^{\infty} a_q(r)q(t-r) + \sum_{r=-\infty}^{\infty} a_c(r)c(t-r). \tag{11}$$

We use vector notation and the three equations

$$\mathbf{F}_y(\nu) = \begin{bmatrix} g^2 f_s(\nu) + \sigma^2_v & g f_s(\nu) + \sigma^2_{vw} \\ g f_s(\nu) + \sigma^2_{vw} & f_s(\nu) + \sigma^2_w \end{bmatrix}, \tag{12}$$

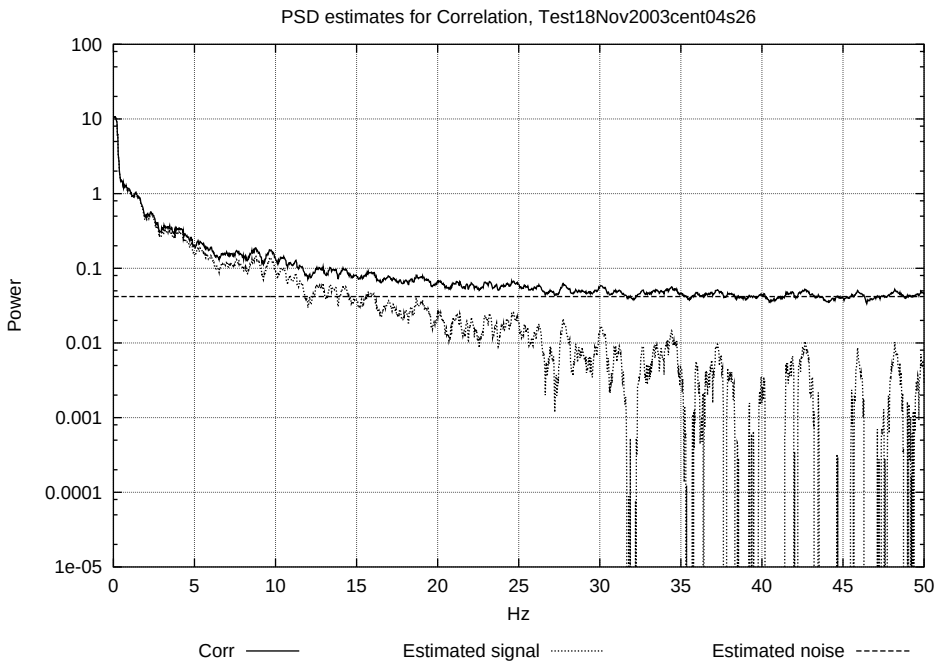

Fig. 2. Example of the PSD fitting. The dark line is the estimated, smoothed PSD for correlation for a specific trial. The straight dashed line is the fit to the noise floor, and the resulting slope signal PSD is also shown.

$$\left[\begin{array}{cc} f_{sq}(\nu) & f_{sc}(\nu) \end{array}\right] = \left[\begin{array}{cc} g f_s(\nu) & f_s(\nu) \end{array}\right], \tag{13}$$

and

$$\left[\begin{array}{cc} f_{sq}(\nu) & f_{sc}(\nu) \end{array}\right] = \left[\begin{array}{cc} A_q(\nu) & A_c(\nu) \end{array}\right] \mathbf{F}_y(\nu), \tag{14}$$

to obtain the best filter, which is expressed as

$$\left[\begin{array}{cc} A_q(\nu_k) & A_c(\nu_k) \end{array}\right] = \frac{\hat{f}_s(\nu_k)\left[\begin{array}{cc} (\hat{g}\hat{\sigma}_w^2 - \hat{\sigma}_{vw}^2) & (\hat{\sigma}_v^2 - \hat{g}\hat{\sigma}_{vw}^2) \end{array}\right]}{\hat{f}_s(\nu_k)\left(\hat{g}^2\hat{\sigma}_w^2 + \hat{\sigma}_v^2 - 2\hat{g}\hat{\sigma}_{vw}^2\right) + \hat{\sigma}_v^2\hat{\sigma}_w^2 - (\hat{\sigma}_{vw}^2)^2}. \tag{15}$$

Again, we will need to have good estimates for all the the parameters $\hat{f}_s(\nu_k), \hat{\sigma}_v^2, \hat{\sigma}_w^2, \hat{\sigma}_{vw}^2$ and $\hat{g}$. This is done by using ensemble averages over all subapertures for a given trial. For all 28 measurements, the periodogram was calculated for each. Then the average of these 28 periodograms was smoothed with $L = 21$. Given this smoothed estimate of $\hat{f}_q(\nu_k)$, the white noise level was estimated by fitting a straight line to the noise floor. The resulting signal power spectrum was the difference between the noisy signal spectrum and the noise spectrum, since the noise and signal are uncorrelated. Figure **??** shows an example of this. For the comparison case, the gain ratio $\hat{g}$ between the two signals was estimated by taking the median of the ratio $[\hat{f}_q(\nu_k)/\hat{f}_c(\nu_l)]^{(1/2)}$ in the frequency band of the signal. The correlation of the noises $\hat{\sigma}_{vw}^2$ was estimated by calculating the smoothed squared

coherence $\hat{\rho}^2_{vw}(\nu_k)$ from the two measurements, fitting a flat curve to the noise region and dividing by the estimated noise variances to get

$$\hat{\sigma}^2_{vw} = \left(\hat{\rho}^2_{vw}\hat{\sigma}^2_v\hat{\sigma}^2_w.\right)^{1/2} \tag{16}$$

6. Analysis results - comparison case

There were several comparison experiments, each with several trials. For each trial, the signal and noise power spectra were estimated as described above, and used to estimate the gain $\hat{g}$. So for each experiment we have a set estimated gains and noise variances which allow an analysis of which method has lower noise and how the gain of the slope algorithms vary.

There were two factors that constrained the number of valid trials. First, the WFS can behave in a non-linear fashion if the spot is too far away from the center. So any subaperture signal with an average value of greater than 0.25 was not used. Second, the implementation of correlation in the control system automatically limits the output to a maximum magnitude of 1. This is because the largest output of Quadcell is of magnitude 1. However, correlation can accurately measure larger shifts, and since the gain is higher (as we shall see) this clipping corrupts the true signal. Therefore we have not used any slope signal from correlation where greater than five percent of the sample have a value of magnitude 1, because the random signals will be corrupted if too many samples are truncated.

First we discuss the results of the gain and noise variance estimations. Then we apply Winer filtering techniques to the pairs of measurements.

A. Comparing method gains and noises

There were seven different experiments. All seven featured the correlation algorithm. Two of the seven compared Correlation to Binned, the other five compared Correlation to Quadcell. The Table shows the relevant parameters of the experiments. For each there were 11 trials, though in a few cases very large shifts lead to fewer trials having useable data (see above for restrictions on valid data sets). We had some specific predictions for algorithm performances for these cases. In particular, an important property of the Quadcell algorithm is that as the spot size increases, the gain will go down. For very small spots the WFS is non-linear and both methods should have the same

Exp. name	Source	Size	SNR	Algo.
Nov7WPH	White light point source probe	very small	high	Quad
Nov7WPL	White light point source probe	very small	low	Quad
Nov7RPH	Red light source probe	small	high	Quad
Nov7RPL	Red light source probe	small	low	Quad
Nov18LGS	Laser Guide Star on-sky	large	low	Quad
Feb5LGS	Laser Guide Star on-sky	large	low	Binned
Mar10LGS	Laser Guide Star on-sky	large	low	Binned

Table 1. Seven comparison experiments for correlation. Table shows date, type of source used, source size on WFS CCD, qualitative SNR of experiment and comparison algorithm.

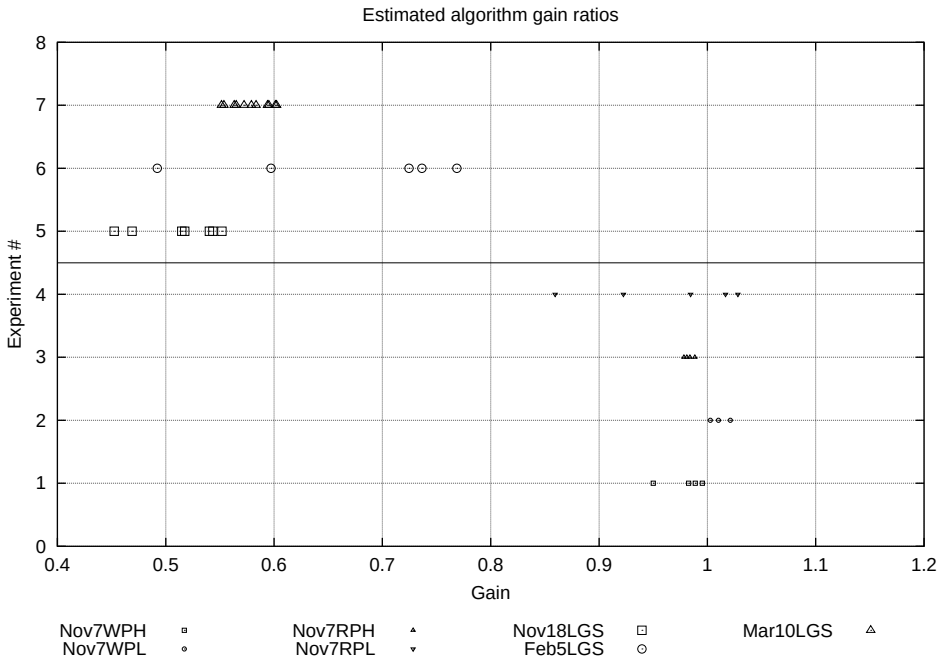

Fig. 3. For each experiment and trial, the relative gain between Correlation and the other algorithm is shown. For small spots the gain is around 1; for large spots the gain on Quadcell or Binned goes down to as low as 0.5.

gain. As spot size increases, the gain of Correlation decreases and becomes constant at 1 for a wide range of well-sampled spots. off, while the gain of Quadcell (and of Binned) continues to decline. These seven experiments explored different spot sizes and SNRs. For white light probe the spot is is approximately 9 times under-sampled according to Nyquist, which is very small compared with the pixel size. The red light probe has a larger size, but is still small compared to the pixels. The LGS is quite large, with the spot bigger than the quadcell area.

As described above, for each trial in each experiment, the following parameters were estimated: gain of the other algorithm relative to Correlation and the noise variances of both algorithms. First we examine the relative gains of the algorithms. This was determined by scaling the signal powers for both measurements such that they were equal. The gain results for each trial are shown in Figure **??**. For the very small white light probe and small red light probe, the gains are close to one. This is as predicted; when the spots are very small both algorithms have the same gain. The spread of the results, particularly in the low-SNR red light probe case, could be attributable to the fact that there is very little signal to be measured. Because we are using an internal probe, the only aberration which can be measured in the system is the small amount of air flow on the optical bench. Analysis of the PSDs of the measurements shows a small amount of very low-frequency signal content. In the LGS case the gains of Quadcell and Binned are significantly lower than Correlation. The three sets of data have a significant variation in gain, most likely caused by large differences in the seeing on the nights experiment. Rough measurements of seeing taken during the experiments

indicate that the seeing was good on Feb. 5, but poor on Mar. 10. No measurement was available for Nov. 18. For the Nov18LGS trial the Quadcell algorithm has a gain of around 0.52. For the Feb5LGS trial Binned has a median gain of 0.72, though there is a quite a spread of results. For the Mar10LGS trial Binned has gain around 0.68. It was expected that in the LGS case Quadcell and Binned would have lower gain, since the LGS spots on the WFS CCD are bigger than the pixel size. A gain as low as 0.5 is reasonable for the WFS if the spots are big enough. Some variation in spot size and gain was also expected, though with no quantitative bounds on range.

These data have shown that for a variety of source sizes the relative gains of Binned and Quadcell change. With each trials there was a successive stepping of the steering mirror of the system, which added a DC component to the offsets. (The analysis of these data does not employ time series techniques per se, so the details will not be included in this report.) This allowed us to determine that the gain of the Correlation algorithm was relatively uniform across the three LGS trials where the Quadcell and Binned gains varied significantly. Unfortunately we have no way with this data set to calibrate these trials to calibrate the LGS trials to the probe trials to determine the true gains of the algorithms, as opposed to the relative gains. Designing an experiment along with appropriate analysis techniques is in our experimental plan.

Now that we have established that there in a difference is gain between the methods, we need to examine the noise of the algorithms to determine which makes a higher-SNR measurement of the slope signal. Because of the different gains we must normalize to the same amount of signal power by scaling the variance of Quadcell or Binned (as fitted to the slope measurement PSD) must by $1/\hat{g}^2$. In Figure **??** the gain-normalized noise variances for all seven experiments are shown. This figure plots the variance of the correlation noise versus the normalized variance of the other algorithm. Points on this curve that lie in the lower right half of the plot below the solid line (which is equality) are for trials where Correlation was noisier. Points above and to the left of the line are for trials where Correlation was less noisy.

For the white probe case at high SNR, Correlation had about 1.85 time the noise power of quadcell. At low SNR the performance is worse with Correlation having 3.6 times the noise power. Results for the red source are similar. In the high SNR case Correlation has 1.86 times the power, for the low SNR case 3.97 times the noise power. This result is to be expected. In both experiments the spots are significantly smaller than the pixels. The Correlation algorithm uses 12 extra pixels which essentially have no signal in them, so it is no surprise that the measurement is much noisier.

The LGS trials have different behavior because the spot is much bigger. As shown above there are substantial changes in gain across the three LGS trials, most likely due to changes in the level of atmospheric turbulence. In addition the noise variance for the trials varies by nearly an order of magnitude. This is another indication of variable seeing. In the Nov18LGS experiment Correlation outperformed Quadcell, having only 0.73 times the noise power, a significant performance improvement. Both methods were quite noisy, however. For the Feb5LGS trial, Binned was compared to Correlation. In this case Binned had lower noise variance; Correlation had 1.28 times the noise power. Both noise variances were significantly lower than the previous LGS experiment. In the Mar10LGS experiment, Binned and Correlation were nearly equivalent, with Correlation having 1.03 times the noise power. The noise variances were back up to higher levels.

This small set of data is not enough to make conclusive judgments, but provides some insights

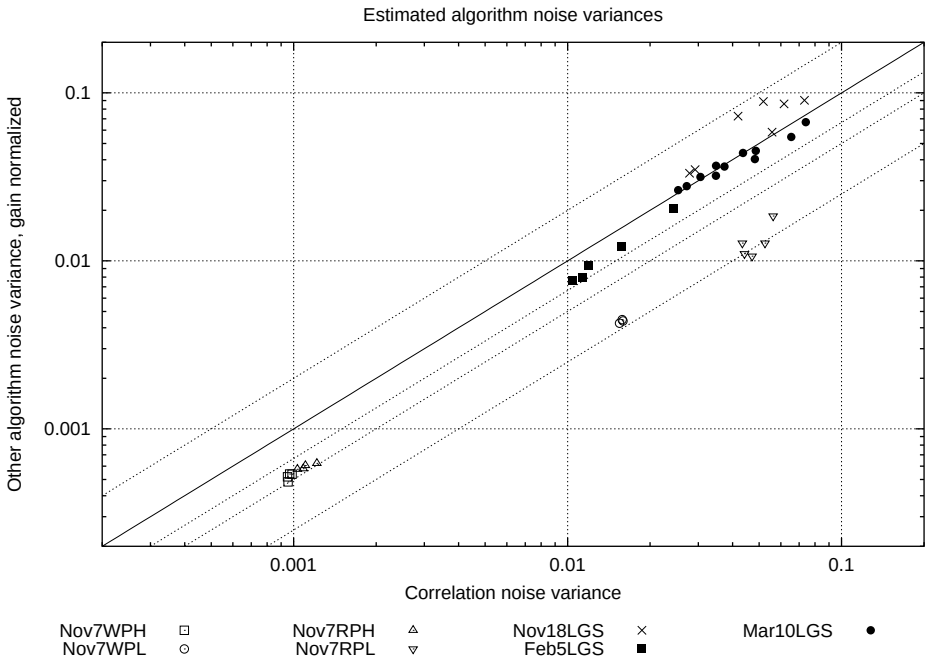

Fig. 4. Comparison of measurement noise variances for Correlation to Binned and Quad, for direct comparison trials.

into the problem. In particular, the noise performance of is correlated with the relative gain. As shown in Figure **??**, the relative gain of Binned or Quadcell decreases, the algorithm becomes more noisy. This holds for both the various trials in the three experiments and all three experiments. This is consistent with the hypothesis of changing spot size. As the spot size becomes bigger, the gain goes down, and more signal is scattered out to the outside pixels. This leads to a noisier Quadcell or Binned algorithm, while Correlation maintains performance independent of spot size.

There are several issues that should be addressed in a comprehensive performance analysis. Trials on the NGS reference should be conducted to compare algorithm performance. Experiments should be regularly conducted over a period of weeks to establish how much spot size (and hence performance) varies over time for LGS and NGS. A way to calibrate the absolute gains of the measurements on LGS and NGS sources would be valuable.

B. Signal estimation

These comparison data sets provide a convenient way to explore different types of Wiener filtering their relative performance. For the two signals in each subaperture, the underlying signal can be estimated from the data using four different possible Wiener filters. A case of the input signals is shown in Figure **??**. In this particular trial the gain of Quadcell was estimated to be 0.47. This is clearly visible in an examination of the slope signals as shown in the figure. Once scaled by the gain, the Quadcell signal is quite similar to the Correlation. Using these two measurements the underlying slope signal is estimated.

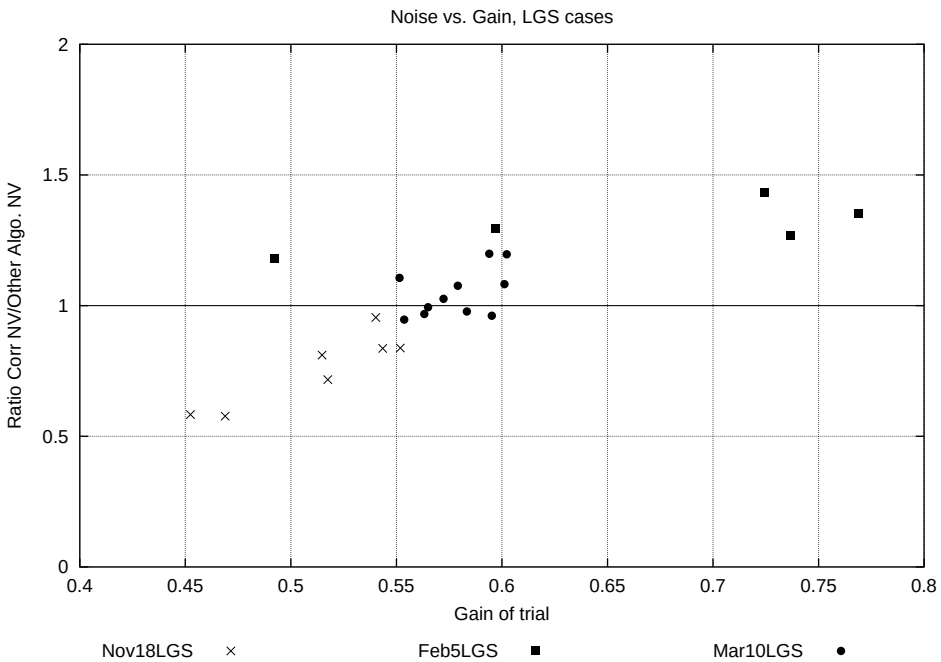

Fig. 5. Ratio of noise power to relative gain. As the relative gain of Quadcell or Binned is reduced (due to increasing spot size) the noise performance of that algorithm relative to Correlation degrades.

Four different Wiener filters were derived for each trial, based on the estimated parameters $g, \hat{\sigma}_v^2, \hat{\sigma}_w^2, \hat{\sigma}_{vw}^2$ and $\hat{f}_s(\nu_k)$. These filters are based on: the signal $\mathbf{q}$ only; the signal $\mathbf{c}$ only; the signals $\mathbf{q}$ and $\mathbf{c}$ assuming the measurement noises were uncorrelated (which they aren't); the signals $\mathbf{q}$ and $\mathbf{c}$ assuming the measurement noises were correlated by the estimated covariance $\hat{\sigma}_{vw}^2$. Using Fourier notation, the filters are as follows. For the Correlation case, the estimate is

$$\hat{S}_c(\nu_k) = \frac{\hat{f}_s(\nu_k)}{\hat{f}_s(\nu_k) + \hat{\sigma}_w^2} C(\nu_k). \tag{17}$$

For Quadcell the estimate must be scaled by the gain of the system, producing

$$\hat{S}_q(\nu_k) = \frac{\hat{g}\hat{f}_s(\nu_k)}{\hat{f}_s(\nu_k) + \hat{\sigma}_v^2} Q(\nu_k). \tag{18}$$

The uncorrelated joint filter is

$$\hat{S}_{qc1}(\nu_k) = \frac{\hat{f}_s(\nu_k)\hat{g}\hat{\sigma}_w^2 Q(\nu_k) + \hat{f}_s(\nu_k)\hat{\sigma}_v^2 C(\nu_k)}{\hat{f}_s(\nu_k)\left(\hat{g}^2\hat{\sigma}_w^2 + \hat{\sigma}_v^2\right) + \hat{\sigma}_v^2\hat{\sigma}_w^2}. \tag{19}$$

The correlated joint filter is

$$\hat{S}_{qc2}(\nu_k) = \frac{\hat{f}_s(\nu_k)\left(\hat{g}\hat{\sigma}_w^2 - \hat{\sigma}_{vw}^2\right) Q(\nu_k) + \hat{f}_s(\nu_k)\left(\hat{\sigma}_v^2 - \hat{g}\hat{\sigma}_{vw}^2\right) C(\nu_k)}{\hat{f}_s(\nu_k)\left(\hat{g}^2\hat{\sigma}_w^2 + \hat{\sigma}_v^2 - 2\hat{g}\hat{\sigma}_{vw}^2\right) + \hat{\sigma}_v^2\hat{\sigma}_w^2 - \left(\hat{\sigma}_{vw}^2\right)^2}. \tag{20}$$

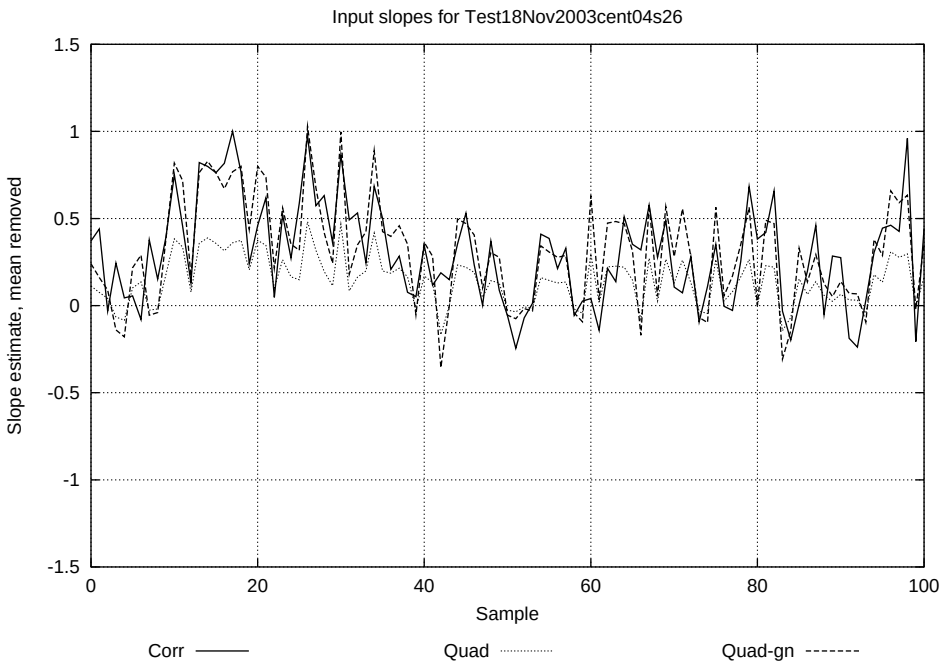

Fig. 6. First 100 samples of Correlation and Quadcell signals from the Nov18LGS trial. Note the scaling by the gain necessary to make Quadcell look similar to Correlation.

In the end, these four different filters produced quite similar results for the estimated signal for all the trials. As a representative case, the estimated signals from the four filters on the data in Figure **??** are shown in Figure **??**. This behavior is understandable given the input power spectrum and the nature of the Wiener filter. As shown in Figure 2, $\hat{f}_s(\nu_k)$ follows a power law and is substantially larger than the noise in the low frequency band. The two measurements are quite similar at these frequencies, as a visual inspection of Figure 6 shows. Assuming that in this range the noise is essentially zero the joint filter reduces to

$$\hat{S}(\nu_k) = \frac{1}{\hat{g}} Q(\nu_k) + C(\nu_k). \tag{21}$$

So at low spatial frequencies the Joint estimate will be simply the average of the Correlation estimate and the scaled Quadcell estimate. The use of the fact that the noise is correlated makes very little difference. This is also due to the nature of the Wiener filter. Again at low frequencies the noise has no impact and at high frequencies the filter is zero since the measurement is all noise there. Only at mid-range frequencies will the covariance $\hat{\sigma}^2_{vw}$. Since the signal spectrum $\hat{f}_s(\nu_k)$ follows a power law, there is very little power here anyway. Using two measurements instead of one will most likely improve the SNR of the estimate of the true slope signal. Adding more information about correlated measurement noises adds little to the estimate. Because the closed-loop spectrum is quite different, using the correlated measurement noises may improve estimation, but the control system is not set up to provide comparison data in closed loop. Despite al this, Winer filtering is

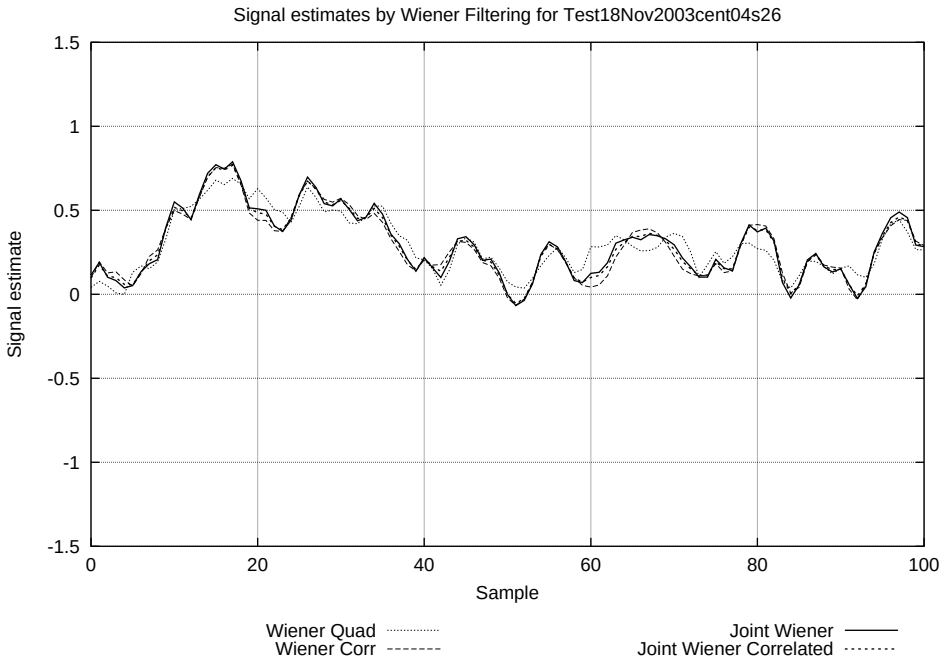

Fig. 7. The four (very similar) wiener filtered signal estimates. Same samples as shown in Figure **??**.

quite useful to obtain optimal estimates of the underlying signal. Wiener filtered estimates are used throughout Section 6.

C. Conclusions

The comparison experiments and time series analysis shows that the Correlation algorithm is noisier for small sources. For the large LGS source the gain of Quadcell or Binned relative to Correlation ranges from 0.5 to 0.8, depending on conditions. For lower gains (larger LGS spots) Correlation has equivalent or better noise performance, but for higher gains (smaller LGS spots) Correlation has worse performance than Binned in the one experiment. Further work is called for to analyze the NGS performance is a similar manner, and to design experiments and data analysis techniques that will enable absolute gain estimation. Special joint-Wiener filters were derived to improve signal estimation using the two measurements, but measurement noise characteristics had small impact due to the power-law nature of the spectrum of the input signal.

7. Analysis results - atmospheric power spectrum case

The WFS telemetry data provide information about the phase aberration which is sensed or corrected by the system. For the open-loop case, we can observe the atmosphere and analyze the spatial correlations of the measured phase slope across the aperture. In the closed-loop case the system is measuring the residual phase that is uncorrected by the system. This is primarily atmosphere that is not fully-corrected due to the delay of the control system, but also contains incorrect compensa-

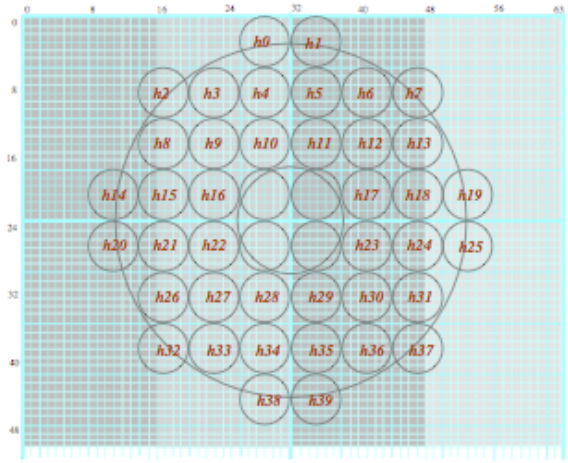

Fig. 8. Lick WFS subapertures on the Lick primary aperture. The lenslets are square, even though they are shown as circular in this figure. We can analyze signal coherences for difference subapertures to examine the spatial coherence of the measured phase slope.

tion due to WFS measurement noise or spatial aliasing of the measurement. For this section of the project we have two experiments conducted on the same night with the LGS reference. We have the open-loop comparison data set Feb5LGS with Binned and Correlation. We also have a set of several closed-loop trials for each of the three algorithms. Ideally we would like to compare open-loop to closed-loop performance, but this data set poses some problems, foremost determining the gains for Quadcell and Binned for the closed-loop trials. Using the techniques outlined above, the Wiener filter for each trial is derived from the estimated signal power spectrum. Using the Wiener filter we estimate the underlying slope signal for each subaperture in each trial. Using these filtered signal estimates, the squared-coherences of pairs of signals are calculated. First we analyze the spatial coherences of the slope signals in open-loop and closed-loop Second, we evaluate the closed-loop performance of all three algorithms with analysis of closed-loop power spectra and coherences.

A. Terminology

The configuration of the subapertures in the Lick WFS is shown in Figure **??**. We choose to estimate the squared-coherence as a function of frequency (see Eqn. 7) for pairs of slope measurements. For a slope in a given direction, we compare it to other slopes of the same direction. Comparing it to the slope of the same direction one subaperture away along the same axis is termed "One". An example of this would be comparing the x-slope in subaperture h3 to that in h4, as shown in Figure **??**. Comparing it to the neighbor perpendicular is called"One-Perp" (x-slopes of h3 and h9). Comparing to the diagonal neighbor is called "One-Diag" (x-slopes of h3 and h10). Results for two subapertures away are labeled similarly.

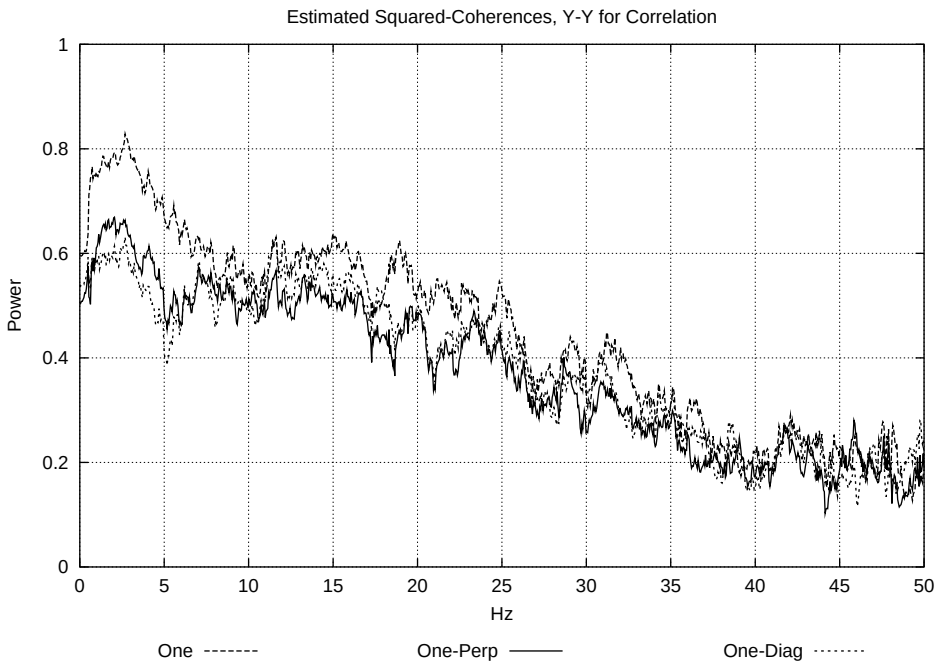

Fig. 9. Open-loop squared-coherences for y-slope signal estimates. Note the high level of temporal correlation for adjacent measurements.

B. Spatial correlations of the phase slope

For both the open-loop and closed-loop data, spectrum estimates and Wiener filtering were conducted as described above. Given the estimated true slope signals for each subaperture, the squared coherences were calculated. The squared-coherences for all valid trials were averaged to get the final estimated squared-coherences. For the open-loop case we can only analyze the y-slopes because the only open loop data we took were in comparison mode during this experiment. A representative plot of the squared-coherences by temporal frequency for the Correlation algorithm is given in Figure **??**. The phase slopes in the "One" case are highly-correlated (squared-coherence of about 0.8) at very low frequencies. The "One-Perp" and "One-Diag" cases are less-correlated across all frequencies, but particularly at very low frequencies. While there is a clear difference in the spatial correlation depending on direction, this result is not necessarily significant by itself. The air mass that causes the phase aberration in the telescope system is translated by the wind. If the wind is blowing in a certain direction, the slope components along that line of travel will be very well-correlated in time, since they are measuring nearly the same signal. These results are averages of trials taken over the period of a few minutes. Several more such experiments should be conducted to see how this wind behavior changes over time.

Wind direction is also a concern in the closed-loop case, as the residual error will be affected by the wind direction and speed. We did the above analysis for all the x-slope and y-slope pairs over five trials for each algorithm and compared the results. The results are shown in Figure **??**. For the

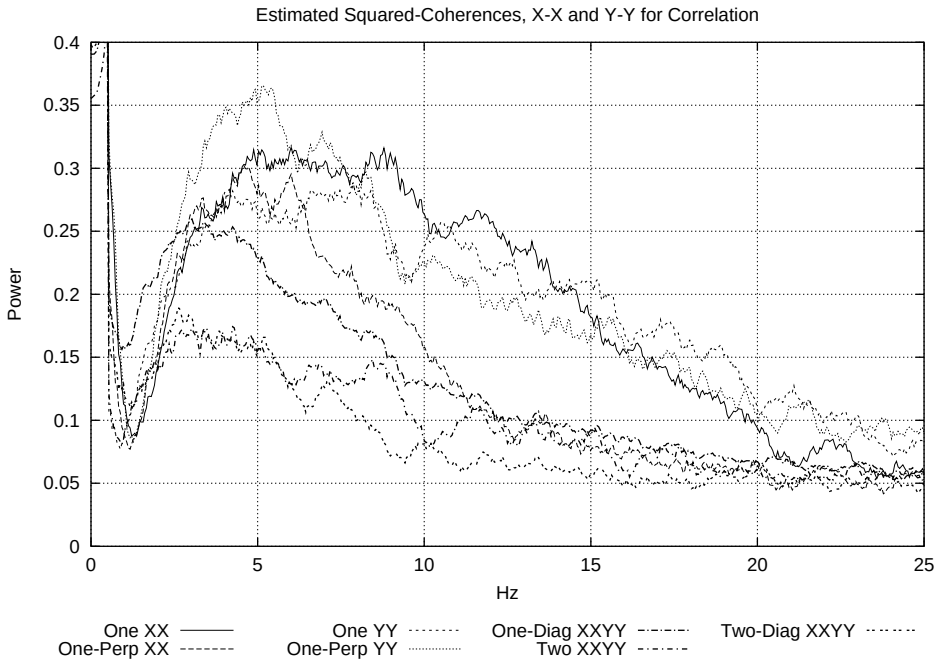

Fig. 10. Closed-loop squared-coherences for x-slope pairs and y-slope pairs. There is significant difference between the x-slope and y-slope results for the "One" and "One Perp" cases. Coherences are much lower than the open-loop case.

"One" and "One-Perp" cases the squared-coherences were significantly different for the x-slopes and for the y-slopes (there was substantially more variation from x- to y-slopes than from trial to trial for a single direction.) Note overall how the coherence of the slopes has gone down significantly. In this case we reach a "noise-floor" in the estimator around 20 to 25 Hz. In the open-loop case this floor wasn't reached until around 40 Hz. The slopes two subapertures away have only a small amount of coherence. (Due to truncation we lacked enough data for the two-subaperture cases for open-loop measurements.)

This analysis would benefit from better experimental design. A series of experiments scattered through time would provide valuable information about the long-term behavior of these spatial correlations, particularly the fact that they may be wind-dependent.

C. Closed-loop comparison of algorithms

To compare the algorithms in closed loop we look for difference in either the power spectrum of the slope signals or the squared-coherence of the slope measurements. Both cases provide information about the residual wave-front which is corrected by the system. The less power in the residual, the better for system performance.

Analysis of the closed-loop power spectrum of the slopes is reasonable, but has a few details that prevent it from being a definitive analysis. By averaging over all the subapertures, a less-noisy estimate of the true power spectrum of the signal is obtained. However, there is a scaling issue

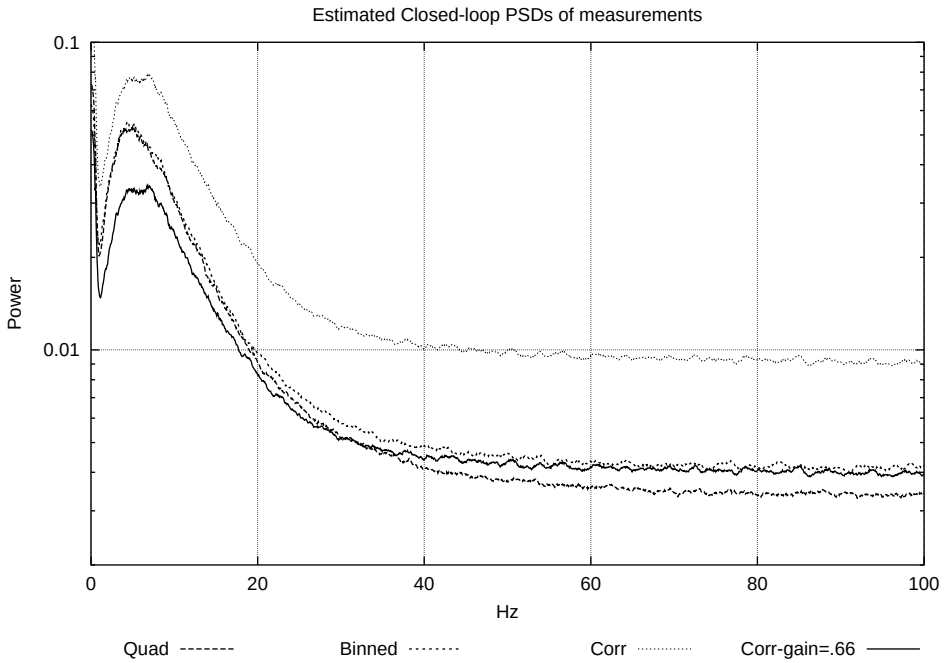

Fig. 11. Estimated closed-loop power spectra of the underlying slope signal (gain-normalized for Correlation) for the three algorithms.

involved between the algorithms. Because Quadcell and Binned have been shown to have a lower and variable gain (ranging from 0.5 to 0.8 in open-loop) to properly compare signal power the estimates must be scaled by the correct factor. In this case, because the comparison trials were taken later in the run, we do not have a contemporaneous measurement of the true gain ratio between Correlation and the other algorithms. We chose to scale the estimated power spectrum by the median gain determined by previous experiments. In this case Correlation residual has lower power. This is shown in Figure **??**. The exact scaling is of course unknown, though for the entire reasonable range based on previous data analysis, Correlation residual has lower power. However, it is unknown exactly what the power is. This problem could be ameliorated in future experiments by interleaving comparison tests with closed-loop tests.

In closed-loop the squared-coherence estimates prove to be very useful. Because these are normalized by the signal spectra, any gain factor will fall out. In all three of the one-subaperture distance cases Correlation had lower squared-coherence than either Quadcell or Binned in the signal band. See Figure **??** for an example from the x-slope "One" case. This result is consistent with the fact that Quadcell and Binned have lower gain. If the gain of a slope estimate is less than 1, it effectively lowers the overall gain of the control loop, leading to poorer temporal performance and increased residual error power in closed loop. As shown above, the gain of Quadcell or Binned can vary from 0.5 to 0.8 in open loop, and consequently they should have increased temporal error and residual power in closed loop. This indicates that Correlation is actually providing better correction

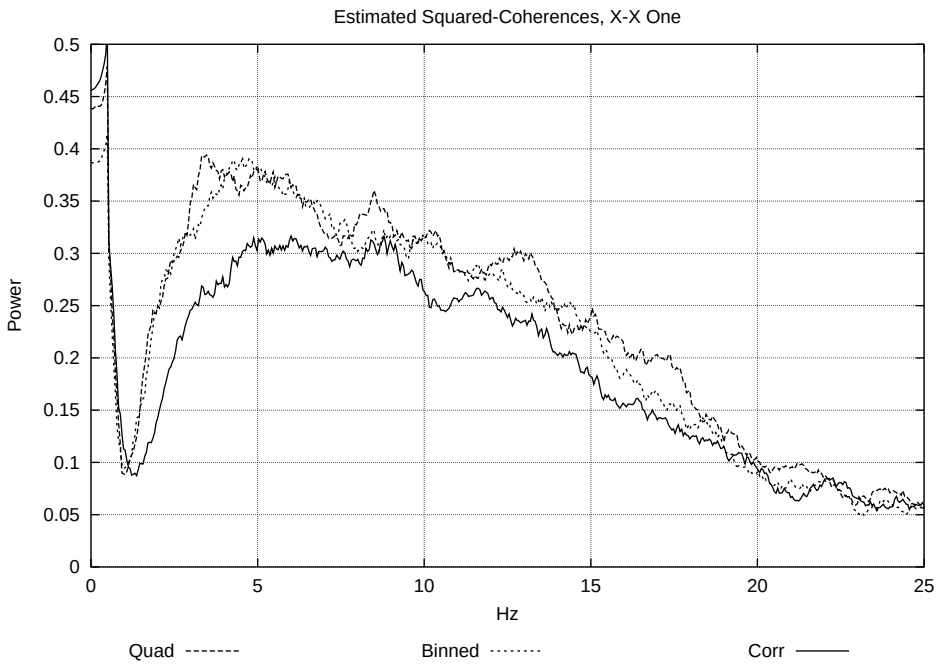

Fig. 12. Estimated squared-coherence for the three algorithms for the x-slope case "one". Note how Correlation has lower squared coherence than the other algorithms across the band.

in closed-loop than either other algorithm.

D. Conclusions

By comparing the estimated slope signals in various subapertures, we were able to determine some useful facts about the system and the algorithms. In open loop the slopes are highly correlated through time with their neighboring measurements. There was a substantial variation in this coherence with direction, which could be caused by wind in the atmosphere. More testing over a larger sample of conditions is necessary to see how this varies through time. In the closed-loop case the slopes are much less correlated, and over a smaller temporal frequency band. Again there is significant variation with direction.

Comparison of the closed-loop trials of each algorithm has shown that Correlation has less residual power and less spatial coherence between slope measurements than Quadcell or Binned. This is most likely due to the gain differences and the resulting temporal errors.

8. Project conclusions and future directions

This project has sought to use time series analysis techniques to compare the performance of Correlation to current WFS algorithms. For the small probe sources, Correlation is a noisier algorithm. For the LGS tests on the sky, Correlation sometimes performs better and sometimes worse in terms of noise. The lower the gain of the Quadcell or Binned (due to larger spots) the better Correlation

does in comparison. The analysis of temporal power spectra and squared-coherences in closed loop showed that the higher gain of correlation resulted in lower residual power and spatial coherence, which should translate into better correction.

One might ask, however, if the gain of Quadcell and Correlation is lower, why not just turn up the gain of the control loop? The problem is that the gain is variable with spot size. It was shown above the the relative gain of Quadcell and Binned varies over a wide range in LGS operation. The optimal gain for Quadcell or Binned will change with time, and would require an adaptive control system to determine, potentially at significant computational expense. Correlation will not have this variation in gain and hence will operate with constant control system parameters. In a standard control system the consistent gain will lead to better performance.

The most important experiment to design next is one which will calibrate the absolute gains of the algorithms. This is necessary to confirm that Correlation has uniform gain in the LGS case. We also must analyze the NGS case in a similar manner to the LGS analysis of this project. We need to determine if gain variations that will affect performance also exist in the NGS case.

There are several ways the analysis of this project could be further deepened with time series techniques. This project has not even touched upon parameter estimation via EM. The analysis presented in this project could be further developed, including deriving error bounds on all the various estimators and coming up with better smoothing techniques to deal with power-law PSDs.

Acknowledgements

This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48. Lisa A. Poyneer can be reached at poyneer1@llnl.gov.

References

1. NSF Center for Adaptive Optics, `http://cfao.ucolick.org`

2. J. W. Hardy, *Adaptive Optics for Astronomical Telescopes,* (Oxford University Press, New York, 1998).

3. F. Roddier, Ed., *Adaptive Optics in Astronomy,* (Cambridge University Press, Cambridge, 1999).

4. J.-P. Veran and G. Herriot, "Centroid gain compensation in Shack-Hartmann adaptive optics systems with natural or laser guide star," J. Opt. Soc. Am. A **17**, 1430–1439 (2000)

5. R. H. Shumway and D. S. Stoffer, *Time Series Analysis and Its Applications,* (Springer-Verlag, New York, 2000).