LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Experimental results for correlation-based wavefront sensing

L. A. Poyneer, D. W. Palmer, K. N. LaFortune, B. Bauman

July 1, 2005

Disclaimer

Experimental results for correlation-based wavefront sensing

Lisa A. Poyneer, David W. Palmer, Kai N. LaFortune and Brian Bauman

Lawrence Livermore National Lab
7000 East Ave, Livermore, CA, 94551

ABSTRACT

Correlation wave-front sensing can improve Adaptive Optics (AO) system performance in two keys areas. For point-source-based AO systems, Correlation is more accurate, more robust to changing conditions and provides lower noise than a centroiding algorithm. Experimental results from the Lick AO system and the SSHCL laser AO system confirm this. For remote imaging, Correlation enables the use of extended objects for wave-front sensing. Results from short horizontal-path experiments will show algorithm properties and requirements.

Keywords: Wavefront sensing, adaptive optics, laser guide star, high power laser, remote imaging

1. INTRODUCTION

Shack-Hartmann Adaptive Optics (AO) systems use an array of lenslets in the pupil plane to sense the gradient of the wavefront phase across each lenslet area on the pupil.[1] This is done by forming an image of the point source as a 'spot' on the wave-front sensor (WFS) CCD. The spot's shift off null (or off a reference location) determines the average slope of the phase in the lenslet. The location of the spot is normally determined by calculating the centroid (center of mass) of the spot. This method is fast, especially for small numbers of pixels per subaperture spot image, and is in widespread use.

There are, however, many scenarios of interest where a point source reference is not available. These scenarios include solar AO, observation of the Earth from space with lightweight optics and remote imaging over short horizontal paths. In all of these cases, the Shack-Hartmann wavefront sensor would image the scene of interest instead of a point source. Given a pupil-conjugate phase aberration, the subimages would simply be shifted by the same amount as a point source reference. These shifts can be measured and used for wavefront reconstruction. Exactly such an approach is used in solar AO,[2] where correlation of the subimages with a reference image is used for the WFS estimation. Our recent work[3] (termed scene-based wavefront sensing) has developed this method into a form that allows analysis of scene performance based on a single sample image. The periodic correlation of a subimage with a specific reference subimage is taken, then the shift is estimated by parabolic sub-pixel interpolation of the correlation peak location. This method is amenable to calculation of the accuracy and noise propagation of slope estimation. This allows prediction of scene quality and performance under varying illumination conditions. In general, the more high-spatial-frequency content the scene has, the sharper its correlation peak and the lower the slope estimation error. In regards to photon noise, the RMS scene-based WFS measurement error is inversely proportional to the SNR, just as in the case of quadcell centroiding with a point source.[4]

The scene-based WFS correlation algorithm has three specific advantages when used on point sources (which are scenes of a sort) as compared to the centroiding algorithm. First, the propagation of read noise when the subimage has many pixels is substantially lower for correlation than for centroiding. For an $N \times N$ pixel subimage, the read noise error variance for centrioding is proportional to N to the fourth power. Second, correlation is an unbiased estimator is the presence of uniform background. Centroiding is biased by a background, lowering the effective gain of the estimate. Third, correlation is robust to change in spot shape. This is partly due to use of a reference, but partly due to the technique itself. As we shall see in the high-power laser case later on, this effect can be quite significant.

In the remainder of this paper we present experimental results in using the scene-based WFS correlation algorithm. In the first two experiments we present data from working AO systems which have a point-source

---

reference which not quite a true point source. In both of these cases, we present experimental evidence confirming the above assertions about the advantages of correlation over centroiding. In the third experiment we examine the case of scene-based WFS on short horizontal paths for remote imaging. This includes SNR requirements, the ability of an arbitrary scene to correctly measure the phase and the impact of anisoplanatism on scene-based WFS.

Throughout this paper we will examine the temporal power spectral density (PSD) of a signal. The PSD is estimated from experimental data using the periodogram method.[5] This technique allows us to determine the gain of the slope estimate and to analyze the temporal signal and noise characteristics of the estimated slopes.

2. LICK LASER GUIDE STAR ADAPTIVE OPTICS

Many astronomical AO systems which use the Shack-Hartmann have 2×2 pixel subimages of the reference spot and use the quadcell version of centroiding. In the quadcell case, spot sizes changes lead to changes in slope estimate gain due to truncation of the spot on the detector. This has been observed in practice[6] and analyzed with an eye towards compensation.[7] When the WFS is not well aligned or is run off-null, this phenomenon can lead to very poor slope estimates. This effect could be ameliorated by using more pixels on the detector, but the propagation of read noise power to error variance increases by a factor of 20 with the move from 2×2 to 4×4 pixels. This noise increase happens only when centroiding is used as the algorithm. Correlation could be used instead without this substantial read-noise penalty. Not only is the noise propagation lower, but the gain is relatively uniform as the spot size changes. This is of particular interest for the case of the laser guide star. Laser guide star spots are not point sources, and correlation could exploit this fact to give better estimates. We conducted an initial test of Correlation using the AO system at Lick Observatory, where the WFS CCD actually has 4×4 pixels, instead of just 2×2. (In regular operation only the center 2×2 are used, however.) This gives us the flexibility to test alternative algorithms.

2.1. Experimental setup

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.[8] In the Lick AO system the phase aberration is measured with a Shack-Hartmann WFS. Each subaperture forms an image onto a 4×4 pixel region. The slopes are normally calculated with the center 2×2 pixel region with the Quadcell algorithm. 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 sometimes used. In this case it uses the quadcell formula, but with pixels with four times the area. Each large pixel is actually the sum of the four real pixels in the region. This method will capture signal from a large spot, but at a penalty of higher noise. In this special case for Lick correlation is implemented as a compact ratio of sums of pixels in the 4×4 pixel region. The pixel configurations and sizes for these three methods are shown in Fig. 1. The control system runs up to 1 kHz, though our tests were usually at lower rates.

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. For our tests we use a special open-loop comparison mode of the control system. In this mode 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. We also use regular open-loop and closed-loop observations where we get both x- and y-slopes with a single algorithm.

2.2. Results: slope estimation gain

Given these three algorithms, the variable gain effect is easy to explore in simulation. This is done by generating spots of varying sizes and using the WFS algorithms to estimate slopes. Fig. 2 shows the response of three different WFS algorithms for a large range of spot sizes in a setup similar to the Lick WFS. For very small spots, none of the algorithms are linear, and all produce the same slope estimate. As the spot size increases, the algorithms approach unity gain. As spot size further increases, the gains of Quadcell and Binned drop off
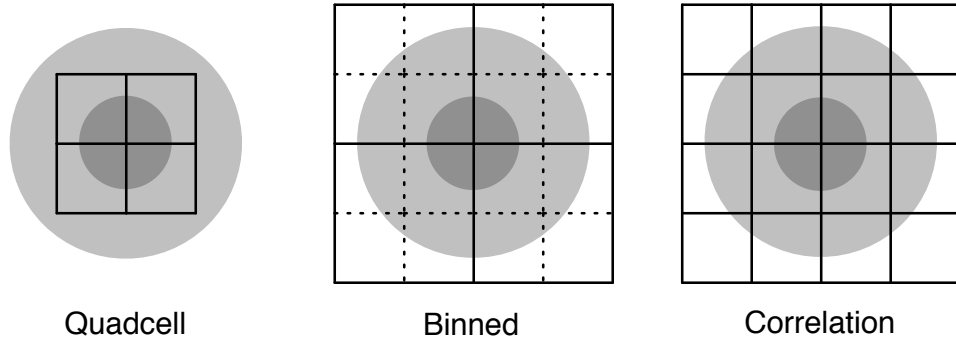
Figure 1. WFS pixel configurations for the three algorithms, for the 4 by 4 pixel area for each subaperture spot. The effective pixels are outlined in solid lines. The gray circles represent small natural guide star spots and large laser guide star spots.
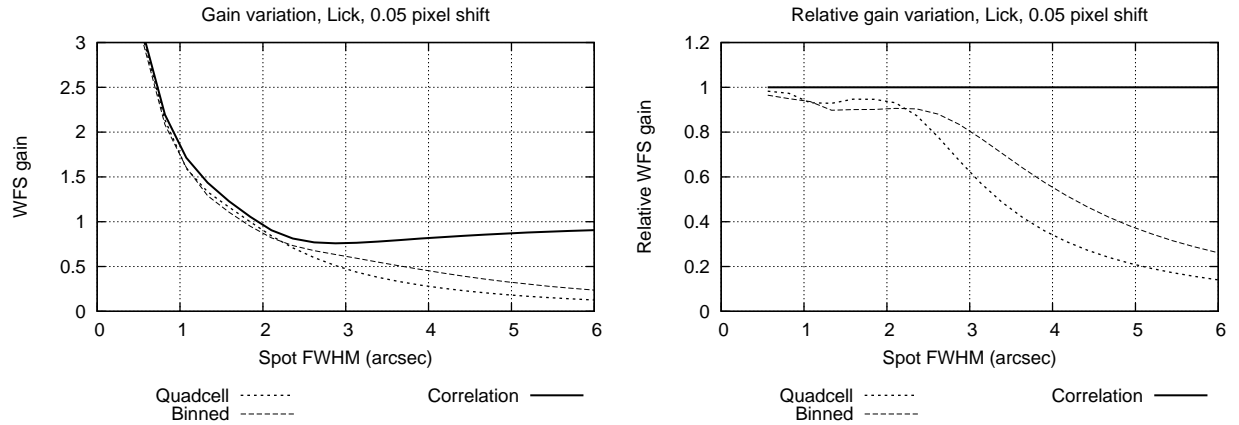
Figure 2. [Left]: Simulation of WFS gain for the Lick case, with 1.9 arcsec pixels and varying spot size. For a fixed shift of 0.05 of a pixel on the WFS, the spot size is changed. As spot size increases, the slope estimate gains for Quadcell and Binned drop uniformly. Correlation stays much more uniform. [Right]: Same data, but plotted as gains relative to Correlation.

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.

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

Nov7WPH	⊡	Nov7RPL	▾	Mar10LGS	△
Nov7WPL	◦	Nov18LGS	▯		
Nov7RPH	▴	Feb5LGS	⊙		

Nov7WPH	▢	Nov7RPL	▽	Mar10LGS	●
Nov7WPL	⊙	Nov18LGS	×		
Nov7RPH	△	Feb5LGS	■		
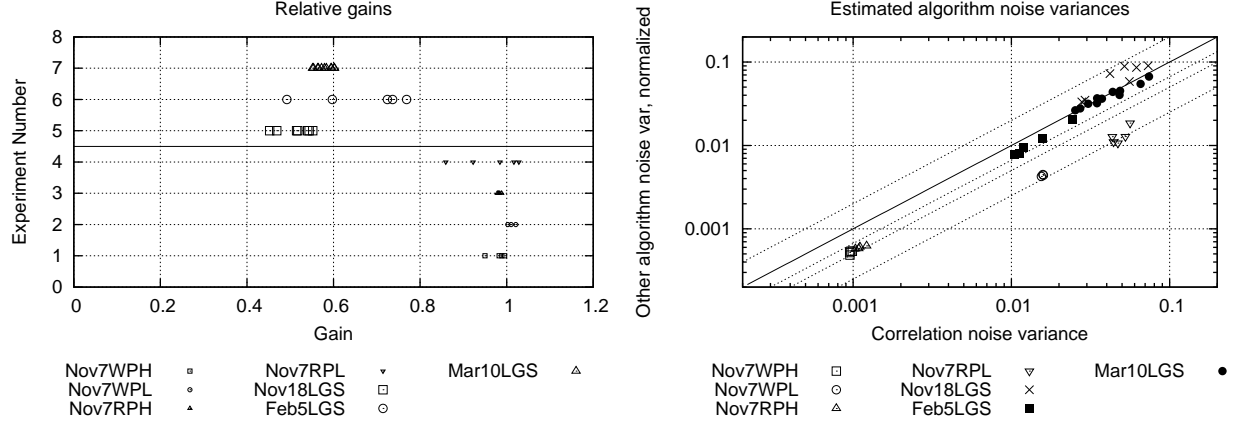
Figure 3. [Left]: For each experiment and trial, the relative gain between Correlation and the other algorithm is shown. For small spots the gain is around 1; for large spots the gain on Quadcell or Binned goes down to as low as 0.5. [Right]: Comparison of measurement noise variances for Correlation to Binned and Quad, for direct comparison trials. For the very small point sources, the noise for Correlation is higher, unsurprisingly.

further, while Correlation remains relatively uniform near a gain of 1. In our experiments we will only have access to the relative gains, which are shown in the right panel of the figure.

To estimate the gain of the centroiding algorithm in comparison to Correlation, we did several dual-mode test with a variety of sources. In each test there were several individual trials. For each trial, the signal and noise power spectra were calculated, and then used to estimate the gain. 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 seven different experiments. All seven featured Correlation. Two of the seven compared Correlation to Binned, the other five compared Correlation to Quadcell. Table 1 shows the relevant parameters of the experiments. These seven experiments explored different spot sizes and SNRs. For the 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. This set of tests allows us to probe the space addressed in Fig. 2. The probes have very small and small sizes, while the LGS spot is large and of variable size. As shown in the right panel of the figure, as the spot size increases, the relative gain of Binned and Quadcell to Correlation will be reduced.

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 Fig. 3.

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 of experiment. 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 0.58. 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 (see the right panel of Fig. 2).

Now that we have established that there is 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) the inverse-square of the gain. In Fig. 3 [Right] the gain-normalized noise variances for all seven experiments are shown. The noise variance level of each was determined by examination of the temporal PSD and fitting of the noise floor. 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 into the problem. In particular, the relative noise performance is correlated with the gain. As the gain of Binned or Quadcell decreases, that 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.

2.3. Results: closed loop performance

To compare the algorithms in closed loop we could look for a difference in either the power spectra 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 involved between the algorithms. Because
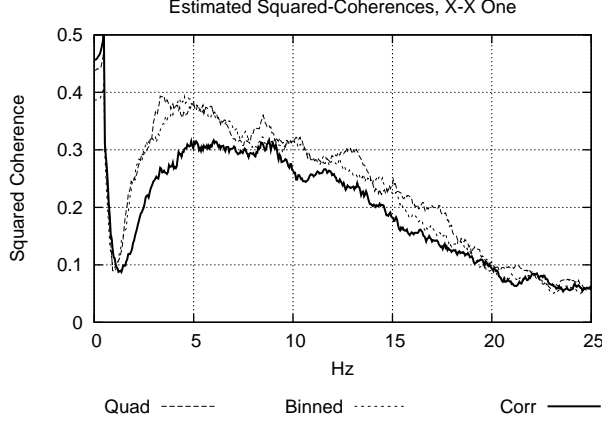
Figure 4. Estimated squared-coherence between x-slopes in neighboring subapertures, for the three algorithms. Correlation has a lower coherence across the band of temporal frequencies, indicating better closed-loop correction.

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. Since we don't know the true gain with high accuracy, we can't use this method conclusively.

In closed-loop the smoothed squared-coherence estimates prove to be very useful. The squared-coherence is a measure of the correlation of the temporal components of two signals.[9] It is calculated in a similar fashion to the periodogram estimate of the temporal PSD. Given two signals $x[n]$ and $y[n]$, with corresponding DFTs $X[k]$ and $Y[k]$, each has a periodogram as $P_x[k] = |X[k]|^2$ and $P_y[k] = |Y[k]|^2$. The smoothed spectral estimates $\tilde{P}_x[k]$ and $\tilde{P}_y[k]$ are the results of applying a box-car average of a specific width to the periodograms. The smoothed squared-coherence is calculated in two parts. First, the cross-spectrum is calculated as $P_{xy}[k] = X[k]Y^*[k]$, where * denotes conjugation. This is then smoothed to produce $\tilde{P}_{xy}[k]$. The smoothed squared coherence is then

$$\tilde{\rho}_{xy}[k] = \frac{|\tilde{P}_{xy}[k]|^2}{\tilde{P}_x[k]\tilde{P}_y[k]}. \tag{1}$$

Because these are normalized by the signal spectra, any gain factor will fall out. Given slope estimates from neighboring subapertures, the squared-coherence is calculated from the time-series. In all three cases of slopes one-subaperture apart (horizontal, vertical and diagonal separations), Correlation had lower squared-coherence than either Quadcell or Binned. See Fig. 4 for a sample of the squared-coherences for the three algorithms in closed loop with the LGS. Lower squared-coherence indicates that neighboring slope measurements are less-correlated and that the atmosphere (which is correlated on these scales) has been better corrected. 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. While the performance of Quadcell or Binned would be improved by a gain compensation, that gain would have to be estimated in some fashion and continuously updated. Correlation provides an algorithm that is robust to changing spot size.

The tests at Lick have shown that in the natural guide star case of small spots, Correlation is a noisier algorithm, because the WFS spots are very small. When the LGS is used, however, the spot size increases dramatically. Both centroiding algorithms suffer from gain variations due to this spot size change, while Correlation does not. This means Correlation provides a uniform and known gain, leading to better closed loop correction.

3. HIGH-POWER LASER ADAPTIVE OPTICS

The Solid State Heat Capacity Laser (SSHCL) project has an intra-cavity AO system to correct for a high power unstable resonator laser.[10] The AO system has some unique features that have posed problems for wave-front sensing. Correlation has specific advantages over a centroiding algorithm in the presence of these problems. First of all, the level of background light as seen on the WFS changes with time. Presence of a background leads to a lower gain on the centroid slope estimate if the background is not removed. In contrast, Correlation gain is insensitive to uniform background. Second, large aberrations in the system lead to spots that have non-diffraction-limited shapes. These variations in spot shape, in particular when spots exhibit smearing or secondary peaks, lead to biased centroid estimates. Correlation with a uniform reference spot ignores these shape changes and successfully finds the true center the the spot. Third, a large number of pixels are used per subaperture on the WFS. The default size is 12×12 pixels. Such a large number of pixels leads to extremely high noise propagation when centroiding is used. This level can be reduced by using a smaller region of pixels (say, 8×8 pixels), but this must be figured out dynamically as the spot location changes. Furthermore, using too small a window leads to spot truncation and a reduction in gain. Correlation allows use of the full region with out the large read noise penalty. For all of these reasons, Correlation should provide better WFS performance in the SSHCL AO system. Our experiments with SSHCL test this prediction.

3.1. Experimental setup

The SSHCL AO system was run in open loop with both a probe laser and the active laser sources for WFS. The raw WFS CCD data was saved from a series of consecutive measurements. For the probe source 100 consecutive frames were used; for the the active laser 20 consecutive shots were used. The probe allows measurement of latent error in the cavity; the active laser measures the heat-induced aberrations during operation. These CCD frames were saved and then were analyzed off-line using the centroiding and correlation algorithms. As in the Lick experiment, the temporal PSDs were calculated and used to determine relative gains and noise levels between the algorithms.

3.2. Results: noise levels

The probe measurement of the cavity allows an analysis of the noise level of the system. The probe measures the turbulence on the optical bench, but this aberration is of a low-enough level such that the effects of noise can be seen in the slope estimates. For the probe trial data of 100 shots, the slopes were estimated as described above. In addition to Correlation on 12×12 pixels, centroiding on 8×8 and 12×12 pixels was done.

After slope estimation, the temporal PSDs were calculated and compared. The comparison of the signal portion reveals the gain (as in the Lick case). Due to background levels, the centroider had a gain of only 0.58 for the 8×8 vase, and a gain of 0.50 for the 12×12 case. This gain is relative to the Correlation measurement, which is assumed to have a gain of 1. After this gain estimation, the temporal PSDs were scaled by the gains such that all algorithms measured the signal equally. Fig. 5 shows these temporal PSDs for the x-slopes. The Correlation method clearly has a lower noise floor. For the 8×8 pixel case, the centroid slope estimate had a gain-normalized noise floor 3 times higher then Correlation's. Increasing the pixel area to 12×12 produced a noise floor about 8 times higher. Though in SSHCL noise is not the dominant consideration, this result is applicable to any Shack-Hartmann AO system that uses many pixels in each subimage on the WFS CCD, as many vision AO systems do.

3.3. Results: accuracy

In SSHCL we are more concerned with the accuracy of a measurement as opposed to the propagation of sensor noise to the estimate. This is for three reasons. First, when the active laser is in use, the SNR on the WFS is quite high. Second, the phase aberrations increase in the laser with each shot, and in a complete run the AO system will have under 100 shots in which to correct these rapidly changing aberrations. In such a case we are more concerned with estimate bias as opposed to estimate variance. The third reason is that special characteristics of the AO system make centroiding susceptible to significant bias. As mentioned above, background levels, field of view truncation and spot shape changes all contribute to estimate bias when centroiding.
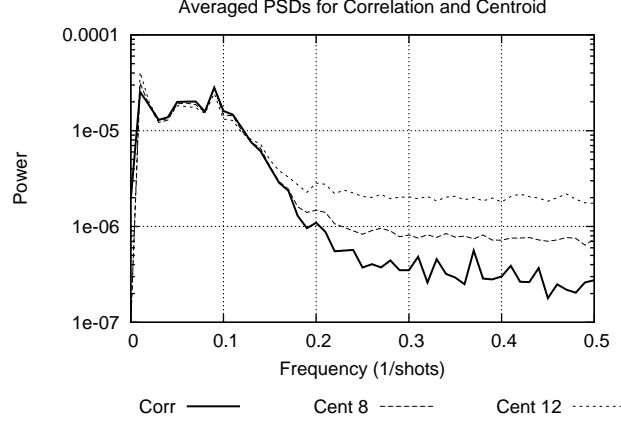
Figure 5. Temporal PSDs of slopes, with the probe source. Centroiding on 8 × 8 (Cent 8) and 12 × 12 (Cent 12) pixel areas are compared to Correlation (Corr). The centroiding noise floor is substantially higher, and increases with the number of pixels used.
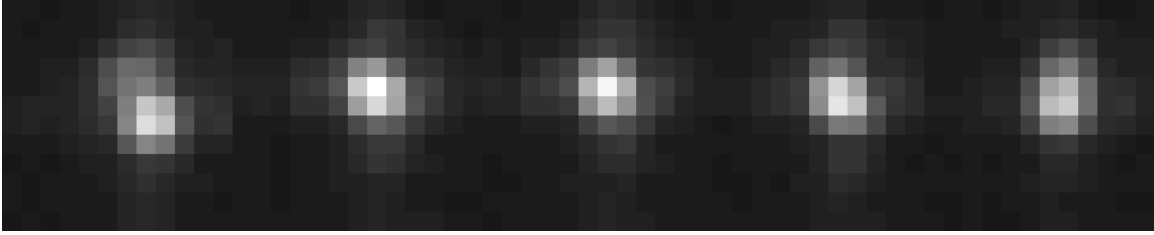
Figure 6. Five different spots from the active laser in SSHCL. In all five cases, the gain-compensated centroid estimate is off by 0.5 pixels in both x- and y-directions.

For the active laser the same type of temporal PSD analysis was completed. This verified both the high level of signal and the high level of aberration. With the active laser, the measured phase aberration had substantially more power than in the probe case and the noise contribution to the slope signals was negligible.

To explore the nature of the bias, the first comparison was between Correlation and centroiding on 12 × 12 pixels for the 20 shots of the active laser. This temporal PSD comparison revealed that centroiding had a gain of only 0.30 compared to Correlation. After this gain was compensated for (by dividing all the centroid estimates by this gain) the difference between slope estimates was calculated. The RMS difference in slope estimates was 0.45 pixels for all x-slopes and 0.21 pixels for all y-slopes. In this optical setup, 1 pixel of shift corresponds to 560 nm of phase slope per subaperture. For the x-slopes 77% of the slope estimates were off by more than 0.1 of a pixel, while 64% of the y-slopes were off by this amount. This is a substantial amount of bias, which will translate into the reconstructed wavefront. The type of spots which produce these errors are shown in Fig. 6. These five examples are the WFS CCD counts for five different spots in a 12 × 12 pixel area. Each of these five spots produced both x- and y-slope estimates with centroiding that were 0.5 pixel different than the Correlation result. A possible solution to this problem with centroiding is to mask out the background light and secondary peaks. This method does produce slope estimates that are more accurate, but the exact level of thresholding varies with the spot shape and background level. This operation must be figured out dynamically for each spot. When setting the threshold level to a percentage of the peak intensity of each spot, even the best threshold level of 40% to 50% produces estimates where 10% of the x-slopes and 5% of the y-slopes are more than 0.1 pixel in error. Noise levels for thresholding are higher than for Correlation, as well.

Experimental data from the SSHCL laser AO system show how sensitive centroiding is to noise propagation when many pixels are used. Correlation is not sensitive to background levels and spot shape alterations, both of which substantially bias a centroiding algorithm. Correlation provides a robust WFS algorithm for this point-source-like case with many pixels.

4. REMOTE IMAGING WAVEFRONT SENSING

4.1. Experimental setup

Our remote imaging testbed for scene-based WFS consists of a 20-cm Celestron CGE telescope with an open-loop Shack-Hartmann WFS. The setup is flexible to allow the use of different lenslet arrays and a tandem science camera as well as the WFS camera. The telescope has a CGE GoTo mount, which facilitates our testing. The mount is heavy, which reduces wind shake. More importantly, the mount is completely motorized and encoded, which allows for steering via a remote control to target the telescope. The equatorial mount is configured as close as possible to having the two axes of motion be the vertical and horizontal axes. This allows target acquisition without touching the telescope, and makes targeting within a field very easy.

The telescope is mounted onto a longer dovetail than normal, on top of which is mounted an optical breadboard. This provides a stable surface to support the WFS optics and ensures that they are always aligned with the telescope. The optical design of the WFS was customized for our system requirements. The beam from the back of the telescope is brought to a pupil plane. The lenslets are mounted at this location. The mounting allows adjustment of rotation to ensure that the subapertures are aligned with the x- and y-axes of the CCD pixels. Two different lenslet arrays are used: one with 25-mm diameter subapertures in the pupil plane (8 across the aperture), the second with 6.7-mm subapertures (30 across). These two arrays allow flexibility in measuring the phase in the pupil.

Both lenslets produce near-Nyquist sampling of the subaperture PSF on the WFS camera, a DVC 1412AM. The DVC camera is an interline CCD which has data transfer using FireWire. It has adjustable exposure times, frame rates and regions of interest (ROIs) when operated with the DVC View software. This camera was selected because it can achieve high frame rates for small ROIs. The frame rate is dependent on the number of rows read off the CCD. For a full frame, the CCD runs at 10 frames per second (FPS). To capture a near-complete row of the 25-mm subapertures, we use a 96-row ROI at 58.8 FPS. For the 64-row ROI we acquire the center region of one row of 25-mm subapertures or two rows of 6.7-mm subapertures at 71.4 FPS. A single row of 6.7-mm subapertures can be acquired at 90.9 FPS using a 32-row ROI. These high frame rates are necessary to probe the temporal characteristics of the atmosphere. The DVC software supports the acquisition of long series of data at consecutive time steps to produce streams of data. By analyzing a stream, we can probe the temporal behavior of the atmosphere at up to half of the frame rate.

Our experiments in Livermore, CA in September 2004 were mainly conducted over short ($\sim$ 100m) paths which were 1 to 2 m above asphalt in hot weather. As targets we used a LED box which provided a true point source as well as a 'spot' board of white dots on a black background which provided a pseudo point source. For scenes we used several printed targets, including a life-size license plate and a life-size picture of a person's head. We used targets such as cross-hair patterns of varying thicknesses to probe specific problems. These targets were usually setup in direct sunlight. Both the 25-mm and 6.7-mm subapertures were used in the testing.

4.2. Results: SNR

An adequate SNR on the WFS is necessary for reasonable AO operation. This SNR is dependent on the exposure time as well as illumination of the target. Previous radiometric analysis had concluded that adequate light levels would be available for scene-based WFS in remote imaging during daytime. We tested this experimentally to determine for the testbed the actual SNRs. There are two SNRs which we are concerned with. The first is the SNR of the image on the CCD: the ratio of the signal content to the noise content. The second is how this translates to a SNR of slope measurement: the ratio of the true spot motion to the error of that measurement, induced by noise on the image.

The SNR of the image on the CCD can be analyzed through proper calibration. We carried out the CCD calibration following generally accepted methods.[11] This calibration gave us a camera gain of 4.3 electrons per DN and a RMS read noise of 5.1 electrons. From this we can determine the SNR of an image. On a per-pixel basis, the SNR is defined as the ratio of photons (or electrons, it does not matter) received from the signal divided by the standard deviation of the noise, with all sources added in quadrature. For a single pixel assuming only signal and no read noise, the SNR is

$$\frac{S}{N} = \sqrt{p}, \tag{2}$$

where p is the mean number of electrons from the source. If we have Gaussian read noise, this equation for one pixel becomes

$$\frac{S}{N} = \frac{p}{\sqrt{p + \sigma_r^2}}, \tag{3}$$

where σ_r is the RMS read noise. Given our CCD characteristics, the read noise only makes a significant contribution to lowering the SNR when there are fewer than 6 DN from the signal. The per-pixel SNR is over 10 for greater than 28 DN from the source. Given the range of 0 to 4095 on the CCD, this means that we are very rarely read-noise limited. We also took frames in a geometric series of exposure times and verified that the camera response is linear across the full 12-bit (0 to 4095) range of digital number outputs.

Now that we know the per-pixel SNR on the WFS CCD, we need to convert this in to measurement SNR for the slopes. This is highly dependent on the amount of slope signal and the method by which it is measured. First we will consider the error due to photon and read noise on the WFS CCD as a function of type of image used in scene-based WFS. We have derived formulas which describe the scaling laws for error variance for an arbitrary scene, as a function of the photon and read noise profiles.[3] We used this formula to analyze a variety of 16×16 pixel scenes for their error standard deviation in Nyquist-sampled pixels. Given the light levels for an exposure of 0.5 msec in bright sunlight at 100m, the standard deviations were very low.

We can also predict how short an exposure would be possible without excess error. Ignoring read noise, we do the same analysis for fewer photons. Extrapolating to an exposure of 0.0625 msec, most of the sample scenes have RMS error of about 0.05 pixels. Some of the lower contrast images have RMS error of between 0.1 and 0.2 pixels. Even for very low light levels we have scenes with low RMS error. At this level the brightest portions of the image are receiving up to 60 DNs of signal per pixel.

These results give us a rough lower limit on the amount of light necessary to conduct useful scene-based WFS. If in bright illumination we can use exposures as short as one-sixteenth of a msec, we could use 1 msec exposures of a target with 16 times less illumination. This gives us a broad range of possible targets. These illumination results hold for a resolved target (an extended scene) at an arbitrary distance. They also hold for arbitrary size subapertures. As long as Nyquist-sampling is maintained on the WFS CCD, the size of the subaperture does not affect the amount of light received in the field of view. Further work will be done to characterize these predictions in actual low-illumination settings (e.g. observations on cloudy days and near dusk).

These results of course depend on the RMS spot motion. The next section will discuss the expected spot motion from theory and the experimental results. For now we will say that for horizontal paths the RMS spot motion is on the order of 1 Nyquist-pixel. This means that measurement SNRs are over 20, even for the very short exposures, and up to 50 for the 0.5 msec exposures mentioned above.

A second way to analyze the slope measurement SNR is to examine the temporal PSDs of the measurements. The PSD estimation technique will also be discussed in detail in the next section. However, the relevant results for SNR analysis are easily summarized. Even the shortest exposures for which we have experimental data (0.1 msec), the temporal PSDs show no noise floor, indicating very high SNRs of slope measurement.

4.3. Results: Measuring Phase

Even if exposure times and frame rates are adequate, an AO system using scene-based WFS will not work unless there is a scene available which will correctly provide the slopes. This depends both on scene content and on the optical design of the system, namely the field of view and subaperture size relative to observation conditions.

To assess this, we conducted many trials in different conditions. The LED box and the spot board were used as references which we assumed would provide accurate measurement of the atmosphere. Observations of these point-source targets were interspersed with observations of our various scenes.

The simplest calculation we can do from the slope estimates is to calculate the total power of the signal and hence estimate the atmospheric coherence length r_0. For a square subaperture of diameter d, the standard deviation of the one-axis spot motion in Nyquist-sampled pixels is[12]

$$\sigma = 2\sqrt{0.18}\left(\frac{d}{r_0}\right)^{5/6}. \tag{4}$$
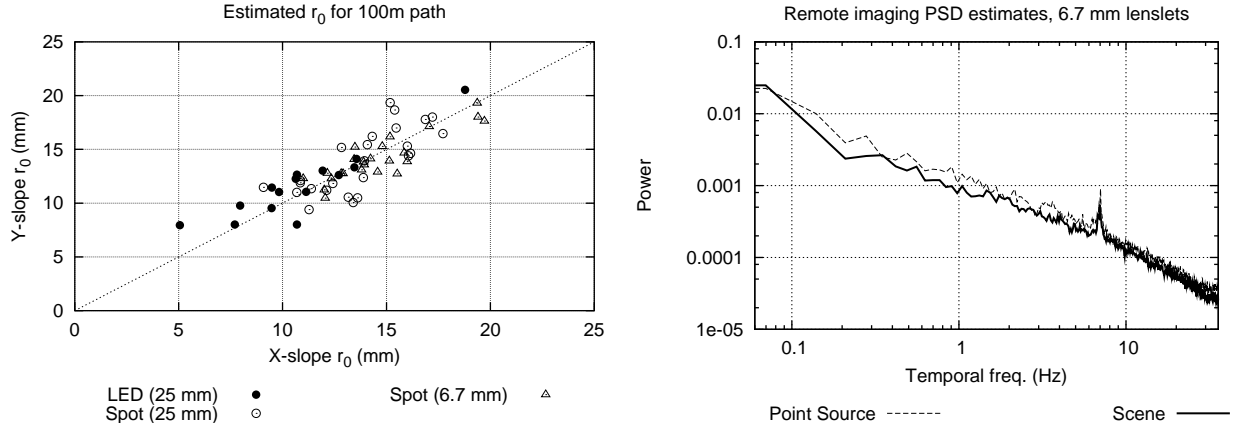
Figure 7. [Left]: Estimated r_0 for 64 different trials with the LED and spot targets over a 100m path. The estimate based on the y-slope is plotted vs. the x-slope estimate. Though there is substantial variation, most values cluster between 10 and 15 mm. [Right]: Temporal PSD estimates for slopes, point source compared to a good scene. The scene measures the same temporal behavior and power as the point source.

Since we know the pixel size and the subaperture diameter d for our lenslet configurations, we can estimate r_0. We have done so for all the point-source data. Fig. 7 [Left] shows a scatter plot of the estimated r_0's for 64 different temporal streams of data using both the LED and the spot board, and with 25-mm and 6.7-mm subapertures. All data were taken over the same path of roughly 100m directly above asphalt in the warm afternoon. The data were taken over three different afternoons in September. This figure shows the significant variation of r_0, though most of the estimates cluster between 10 and 18 mm.

To determine whether or not the scenes are estimating the slopes correctly, we need a more detailed analysis than just the r_0 estimate. This is done with the temporal PSD technique. In this case since we do not have simultaneous measurement of the same aberration with different algorithms (as we did in the Lick and SSHCL cases) we have no exact 'truth'. We can, however, compare measurements that were taken nearly contemporaneously and average over several trials. Fig. 7 [Right] shows two such temporal PSDs, one for a point source reference and the other for a scene, over the same path, but over different instants in time. Both the point source and the scene slopes exhibit the same temporal structure and the same vibration spike (of unknown origin).

Not all scenes are as good as this example, however. In some cases the scene does not have sufficient spatial frequency content to produce good measurements. In previous work[3] we have shown how to predict the error variances of the slope estimate based on scene content. We have done a similar analysis for the gain of the estimate. Even though the scene is actually shifted by the correct amount in the subimage, the correlation estimation technique may return a estimate that is lower by a uniform gain, due to image content. The easiest way to determine this gain on experimental data is to take a sample subimage over the full field of view (which is circular in our optical system) and estimate the gain by simulating image shifts.

In one trial we observed several different scenes with the 6.7-mm lenslets. As a reference, we also observed the spot board. For each target, we calculated the power of the slope measurement (the mean-square error). This MSE was used above in r_0 estimation. For the scenes, we also predicted the gain of the scene. Fig. 8 shows a scatter-plot of the gain of the scene (as predicted from experimental data) and the MSE of the slopes measurements. Among the set of scenes, there is a large variation in the quality of the image in terms of the gain of the slope estimation. Points that lie near the dotted line on the figure are scenes where the total slope power is correct given the estimated r_0 over the trials. While some scenes are clearly poor estimators, this plot shows that given a prediction of the scene gain, a reasonable compensation can be made to measure the true slope. Though not shown, the temporal PSD of scenes with gains of greater than 0.50 have the same structure as that of the point source, as shown in Fig. 7 [Right].

A second major consideration in this remote imaging scenario is anisoplanatism. The phase aberration along the imaging path is distributed, which means that different parts of the imaging target will 'see' different phase
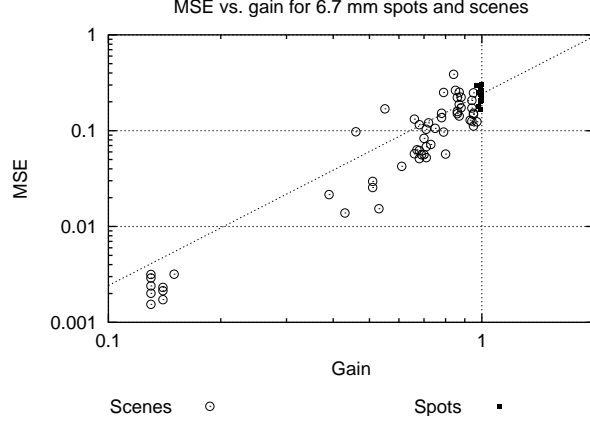
Figure 8. Total slope MSE versus gain of scene used for scene-based WFS, 6.7-mm subapertures. The spot sets the level of power, with the dashed curve showing the MSE for varying WFS gain. X-slope and Y-slope MSEs are shown for three different scenes. Some scenes clearly work better than others.
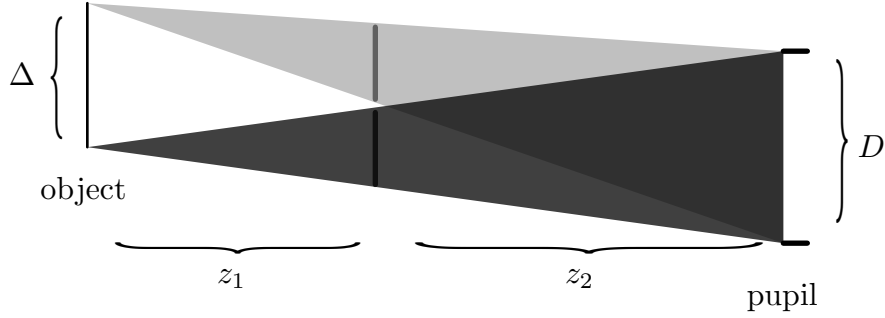

Figure 9. Cartoon rendering of the anisoplanatism issue for subimages. Two points on the object are distance Δ apart. Two cones are formed by tracing rays from these points to the entrance pupil (of diameter D) of the optical system. As the phase aberration moves away from the pupil towards the object, the two points see less phase in common along their paths.

aberrations along the optical path. This impacts not only the full scene which is imaged by the telescope, but the smaller scene which is imaged by each subaperture in the WFS. The field of view of the subimage is restricted on the CCD such that the neighboring subimages do not overlap. In our case the field of view is approximately 1 milliradian for both subaperture sizes. For a target 100m away, this translates to a 10 cm field of view.

If the phase aberration is at the telescope pupil, all parts of the image are affected by it in the same way. As the phase aberration moves away from the pupil towards the object, light from different parts of the image have different aberrations. This is illustrated by the cartoon in Fig. 9. As the phase screen moves away from the pupil towards the object, the phase aberration seen by two parts of the image eventually becomes distinct. This means that instead of the entire subimage being shifted by the same amount (as is the case for a pupil-conjugate phase aberration), different parts of the subimage will be shifted in different directions. This effect is exacerbated as the field of view on the object increases. As the field of view increases, this separation phenomenon occurs for phase aberrations closer to the telescope.

When the phase aberration is very close to the object, it cause local changes in structure which are seen by all subimages. This simply changes the structure of the image from timestep to timestep. Because a new reference subimage is used at every timestep, scene-based WFS is insensitive to phase aberrations very close to the object. Phase aberrations close to the pupil will be well-measured. It is the phase aberrations in the middle

| | Subimage 1 | Subimage 2 | Difference | Simulated Diff. |
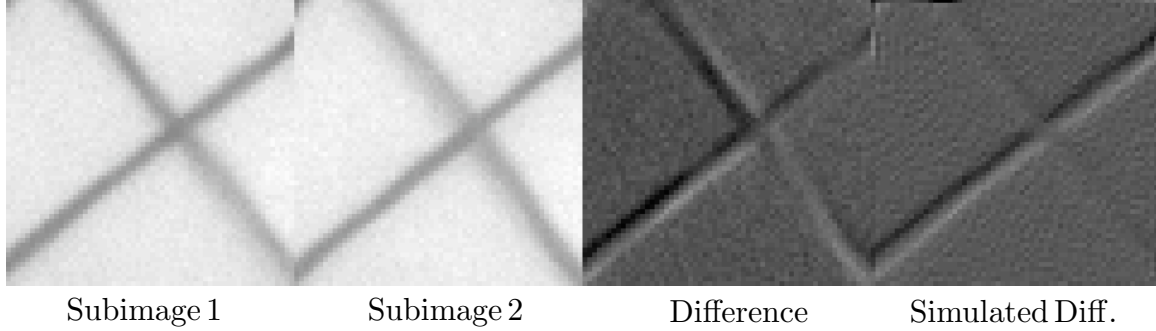
Figure 10. Comparison of field-test subimages from 25-mm case that demonstrate anisoplanatism. On the left are two subimages from subapertures which are 10 cm apart, at the same instant in time. On the right are the actual difference of the two frames, and the simulated difference, based on the scene-based WFS shift estimate. Individual parts of the object have moved, demonstrating anisoplanatism that limits scene-based WFS.

of the path that pose the biggest problem.

Even over a 100m path, the 25-mm lenslet configuration suffered from anisoplanatism in the subimages. Experimental data shows that sometimes the subimages from different subapertures at the same point in time are not merely shifts of each other, but have different structure. An example of this is shown in Fig. 10. The two subapertures are separated by 10 cm in the pupil plane. The difference frame is the actual difference of the two subimages. The simulated difference frame was generated by estimating the subimage shifts using scene-based WFS and then subtracting a sub-pixel shifted copy of subimage 1 from itself. A comparison of the true and simulated differences shows that significant portions of the cross image are shifted locally, as opposed to globally.

Anisoplanatism in the subimage will degrade slope estimation. This effect can be ameliorated somewhat by restricting the the field of view of the subimage. The amount of anisoplanatism is dependent on the observing conditions and how the turbulence is distributed over the optical path. Based on the analysis above, horizontal paths with fully distributed turbulence (as is our experimental case) are most difficult. Paths where turbulence is concentrated at object and pupil (e.g. looking out across a valley to the other side well above the bottom) are much better. Anisoplanatism is also a major factor in designing an AO system which would correct a large field of view in remote imaging. This is inherent to the distributed turbulence and not due to the use of scenes, so we will not discuss this issue further.

5. CONCLUSIONS

We have presented experimental results from three separate tests of the scene-based WFS correlation algorithm. Two of those tests were in systems where the point-source reference is not really a point source and where correlation provides specific benefits. At Lick with the laser guide star, we showed that the gain of the quadcell centroiding algorithm varied substantially on time scales both short and long. This change is gain leads to reduced system performance. In closed-loop, use of Correlation leads to less-coherent slopes and better correction due to this higher and uniform gain. The noise level of Correlation is comparable to centroiding in this case.

For the SSHCL high-power laser AO system, we demonstrated that Correlation provides much lower noise levels and estimates of higher accuracy. When many pixels are used in the WFS, centroiding is susceptible to errors caused by both background light and noise. Correlation is insensitive to these errors and provides lower-noise estimates. The SSHCL system has large aberrations which cause the WFS spots to have deformations from a uniform diffraction-limited spot. These deformities lead to substantial estimation bias when using centroiding. If many pixels are used on the WFS, Correlation is a much more robust algorithm than centroiding.

Our final experiment was using scene-based WFS in a remote imaging application over short horizontal paths. We used a 20-cm telescope for imaging over 100-m paths. In daylight, there was more than adequate light for operation and no noise floor was evident in the slope data, even for exposures under 0.5 msec. Some scenes are better for slope estimation that others, but the gain of the scene can be predicted and normalized for. Good

scenes measure the same temporal PSD for the slopes as the point source references do. Anisoplanatism poses a problem for scene-based WFS on paths with distributed aberrations, though this effect can be mitigated by proper choice of subimage field of view and observing conditions.

ACKNOWLEDGMENTS

REFERENCES

1. M. C. Roggemann and B. Welsh, *Imaging Through Turbulence*, CRC Press, New York, 1996.
2. T. R. Rimmele, "Solar adaptive optics," in *Adaptive Optical Systems Technology*, P. L. Wizinowich, ed., *Proc. SPIE* **4007**, pp. 218–231, 2000.
3. L. A. Poyneer, "Scene-based Shack-Hartmann wave-front sensing: analysis and simulation," *Appl. Op.* **42**, pp. 5807–5815, 2003.
4. G. Tyler and D. Fried, "Image-position error associated with a quadrant detector," *J. Opt. Soc. Am. A* **72**, pp. 804–808, 1982.
5. A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-time Signal Processing*, Prentice Hall, New Jersey, 1999.
6. I. de Pater, S. G. Gibbard, B. A. Macintosh, H. G. Roe, D. T. Gavel, and C. E. Max, "Keck adaptive optics images of uranus and its rings," *Icarus* **160**, pp. 359–374, 2002.
7. J.-P. Véran 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), 2002.
8. J. M. Brase, J. An, K. Avicola, H. D. Bissinger, H. W. Friedman, D. T. Gavel., B. Johnston, C. E. Max, S. S. Olivier, R. W. Presta, D. A. Rapp, J. T. Salmon, K. E. Waltjen, and W. A. Fisher, "Adaptive optics at Lick Observatory: system architecture and operations," in *Adaptive Optics in Astronomy*, M. A. Ealey and F. Merkle, eds., *Proc. SPIE* **2201**, pp. 474–488, 1994.
9. R. H. Shumway and D. S. Stoffer, *Time Series Analysis and its Applications*, Springer, New York, 2000.
10. K. N. LaFortune, R. L. Hurd, E. M. Johansson, C. B. Dane, S. N. Fochs, and J. M. Brase, "Intracavity adaptive correction of a 10-kw solid state heat-capacity laser," in *Laser Resonators and Beam Control VII*, A. V. Kudryashov, ed., *Proc. SPIE* **5333**, pp. 53–61, 2004.
11. S. B. Howell, *Handbook of CCD Astronomy*, Cambridge University Press, Cambridge, 2000.
12. D. L. Fried, "Optical resolution through a randomly inhomogeneous medium for very long and very short exposures," *J. Opt. Soc. Am. A* **56**, pp. 1372–1379, 1966.