Spatially filtered wave-front sensor for high-order adaptive optics

Lisa A. Poyneer and Bruce Macintosh

Lawrence Livermore National Laboratory, Livermore, California 94550

Adaptive optics (AO) systems take sampled measurements of the wave-front phase. Because in the general case the spatial-frequency content of the phase aberration is not band limited, aliasing will occur. This aliasing will cause increased residual error and increased scattered light in the point-spread function (PSF). The spatially filtered wave-front sensor (SFWFS) mitigates this phenomenon by using a field stop at a focal plane before the wave-front sensor. This stop acts as a low-pass filter on the phase, significantly reducing the high-spatial-frequency content phase seen by the wave-front sensor at moderate to high Strehl ratios. We study the properties and performance of the SFWFS for open- and closed-loop correction of atmospheric turbulence, segmented-primary-mirror errors, and sensing with broadband light. In closed loop the filter reduces high-spatial-frequency phase power by a factor of 10^3 to 10^8. In a full AO-system simulation, this translates to a reduction by up to 625 times in the residual error power due to aliasing over a specific spatial frequency range. The final PSF (generated with apodization of the pupil) has up to a 100 times reduction in intensity out to $\lambda/2d$. © 2004 Optical Society of America

OCIS codes: 010.1080, 010.7350.

1. INTRODUCTION

In the general case, adaptive optics (AO) systems operate by using sampled measurements of the phase, or a derivative of the phase.[1] The Shack–Hartmann sensor is a common method that samples the phase gradient every subaperture width d in the pupil plane. The reconstructed phase aberration is then applied as a correction, typically in closed loop, with use of wave-front correction devices that have discrete actuators. Because the wave-front sensor (WFS) in the AO system takes sampled measurements of the phase, which is not band limited in the general case, aliasing will occur.

Aliasing is a phenomenon associated with sampling that causes high-frequency signals to be measured as low-frequency signals. An in-depth discussion will not be provided here; please see a text such as Ref. 2 for a thorough treatment. The most important result is that when samples of a signal are taken at intervals of duration T, the signal must be band limited at $1/2T$ (which is half the sampling frequency) for perfect sampling and reconstruction to occur. If the wave front is sampled spatially with period d (the width of a subaperture) in both x and y directions in the pupil plane, the phase must have no spatial-frequency content above $1/2d$. This condition is not met in most cases.

Aliasing will degrade the performance of most AO systems. Because aliasing causes high-frequency phase components to be seen as low-frequency ones, the control system will mistakenly try to correct these phase components. This leads to excess error in the steady state. The characteristics of this aliasing error depend on the power spectrum of the input phase and the particulars of the AO system itself. Analytic calculations have shown[3] that in the infinite-aperture case that corrects atmospheric turbulence, the mean square error due to spatial aliasing is one third of the mean square error due to fitting [i.e., high-frequency phase errors that the deformable mirror (DM) cannot correct].

In a more specialized scenario, aliasing is detrimental in the astronomical high-contrast imaging case, i.e., the search for faint objects close to bright stars. As discussed in Section 2 and in recent work by Sivaramakrishnan *et al.*[4] and Perrin *et al.*,[5] at high Strehl ratios when diffraction has been controlled, the point-spread function (PSF) outside the core approaches the power spectral density (PSD) of the phase. With perfect knowledge of the phase, a deformable mirror should be able to sharply attenuate the components of the phase that correspond to spatial frequencies up to $1/2d$, removing scattered light at radii less than $\lambda/2d$ and producing a PSF with a (square) dark region. This dark region is termed the "basin" of the PSF. Real AO systems with pupil-plane sensors never produce this effect, however, owing to aliasing that fills in the basin. Aliasing of wave-front errors due to the atmosphere or semistatic errors due to the telescope are particularly damaging, since such errors result in slowly evolving speckle patterns[6,7] that can mask interesting faint companions such as extrasolar planets. Removing the aliased power can result in vastly greater sensitivity to companions.

A simple solution to prevent aliasing is to low-pass filter the input signal before the samples are taken. This low-pass filter reduces the power of the signal above the frequency cutoff imposed by the sampling and reduces the amount of aliasing. We propose to implement a nonideal low-pass filter on the phase by using a field stop at a focal plane before the wave-front sensor. This spatial filter will reduce the power in the wave front at high spatial

frequencies and correspondingly reduce the error in the wave-front correction due to aliasing. This configuration will be termed the spatially filtered WFS (SFWFS).

In this paper we examine in detail the design and performance of the SFWFS. The filter, in reasonable situations, can reduce the power of the phase at spatial frequencies above the cutoff by factors of 10^3 to 10^8. Integrated into a closed-loop, noise-free AO simulation, this filtering translates to a reduction in power of the phase residual at low spatial frequencies by a factor of up to 625 times, just by preventing aliasing.

2. PRINCIPLES OF THE SPATIAL FILTER

The basic design for the spatial filter is a hard-edged, square field stop in the focal plane. Its width is λ/d. Given a 10-m telescope with a WFS using 800 nm light and $d = 15.6\,\text{cm}\,(D/d = 64)$, the field stop of the SFWFS is of size 1.05 arcsec. Using subapertures of size 56 cm gives a field stop of size 0.29 arcsec. The filter is square, not circular, because the expected WFS and DM actuator configuration is on a square grid. On the basis of sampling theory, the system should be limited to spatial frequency $1/2d$ to prevent aliasing. Note that a different configuration of WFS and DM (e.g., a hexagonal grid) would have a different controllable region and hence a differently shaped filter. This filter will act as a nonideal low-pass filter on the phase. Low-spatial-frequency (LSF) content below $1/2d$ will be passed, and high-spatial frequency (HSF) content above $1/2d$ will be attenuated. A similar concept for the Pyramid sensor has been proposed[8]; however, it has a field stop (0.1 arcsec) small enough that it rejects controllable spatial frequencies, and no detailed performance analysis was given.

A. Relating the Phase to the Field

The spatial filter exploits the relationship between the frequency content (or power spectrum) of the phase aberration and the field intensity and PSF. This relationship has been recognized for some time,[9] and recent work was been mathematically explicit.[4,5] The latter two references expand the pupil function as a Taylor series and approximate the PSF by several derived terms. This treatment provides insight into both the design of the spatially filtered wave-front sensor and the analysis of its performance. We will briefly summarize the mathematical treatment and then apply the results to the current scenario.

The pupil function is represented by a real aperture function $a(x, y)$ and a real phase function $\phi(x, y)$ in the pupil plane. This complex exponential can be written as a Taylor series such that

$$a(x, y)\exp[\,i\,\phi(x, y)]$$
$$= a(x, y)$$
$$\times \left[1 + i\,\phi(x, y) - \frac{\phi^2(x, y)}{2} - i\,\frac{\phi^3(x, y)}{6} + \ldots \right],$$

$$(1)$$

where i is the imaginary number equal to $(-1)^{1/2}$. This linearization provides direct access to the phase $\phi(x, y)$

and its spectrum in the equations. The field in the focal plane is determined by taking the Fourier transform of the above expression. The field $E(X, Y)$ in the focal plane is

$$E(X, Y) = A(X, Y)*\left[\delta(X, Y) + i\Phi(X, Y) \right.$$
$$- \frac{\Phi(X, Y)*\Phi(X, Y)}{2}$$
$$\left. - i\,\frac{\Phi(X, Y)*\Phi(X, Y)*\Phi(X, Y)}{6} + \ldots \right],$$

$$(2)$$

where $\delta(X, Y)$ is the unit impulse function, $A(X, Y)$ and $\Phi(X, Y)$ are the Fourier transforms of the aperture and the phase functions, respectively, and the $*$ operator denotes convolution.

We will initially consider the monochromatic case for the filter, so the specific wavelength λ of light used will be mentioned only as necessary in the equations. Equation (2) can immediately be applied to the design concept of the spatial filter. The following series of equations are illustrated in Fig. 1. Consider a phase aberration that is a pure cosine with spatial period in the pupil plane l, such that

$$\phi(x, y) = \cos\left(\frac{2\pi y}{l}\right). \qquad (3)$$

Such an aberration is shown in the top row of the figure. The Fourier transform of the phase aberration is simply two impulses at a distance determined by the spatial frequency $1/l$ and the wavelength of light λ,

$$\Phi(X, Y) = 0.5\left[\delta\left(X, Y - \frac{\lambda}{l}\right) + \delta\left(X, Y + \frac{\lambda}{l}\right) \right]. \quad (4)$$

The field in the focal plane is now the unaberrated field due to diffraction $A(X, Y)$ plus a series of weighted replications at distances that are integer multiples of λ/l,

$$E(X, Y)$$
$$= A(X, Y) + 0.5j\left[A\left(X, Y - \frac{\lambda}{l}\right) + A\left(X, Y + \frac{\lambda}{l}\right) \right]$$
$$- 0.125\left[A\left(X, Y - 2\frac{\lambda}{l}\right) + 2A(X, Y) \right.$$
$$\left. + A\left(X, Y + 2\frac{\lambda}{l}\right) \right] + \ldots. \qquad (5)$$

The magnitude of the field in the focal plane is shown in the second row of the figure. A "lump" in the field appears at the characteristic distance, with smaller lumps appearing farther out. The magnitude of these lumps falls off geometrically with distance away. If the spatial frequency $1/l$ is greater than the band limit imposed by sampling of $1/2d$, all of the lumps will be cut off by a field stop with a width of λ/d, effectively suppressing the
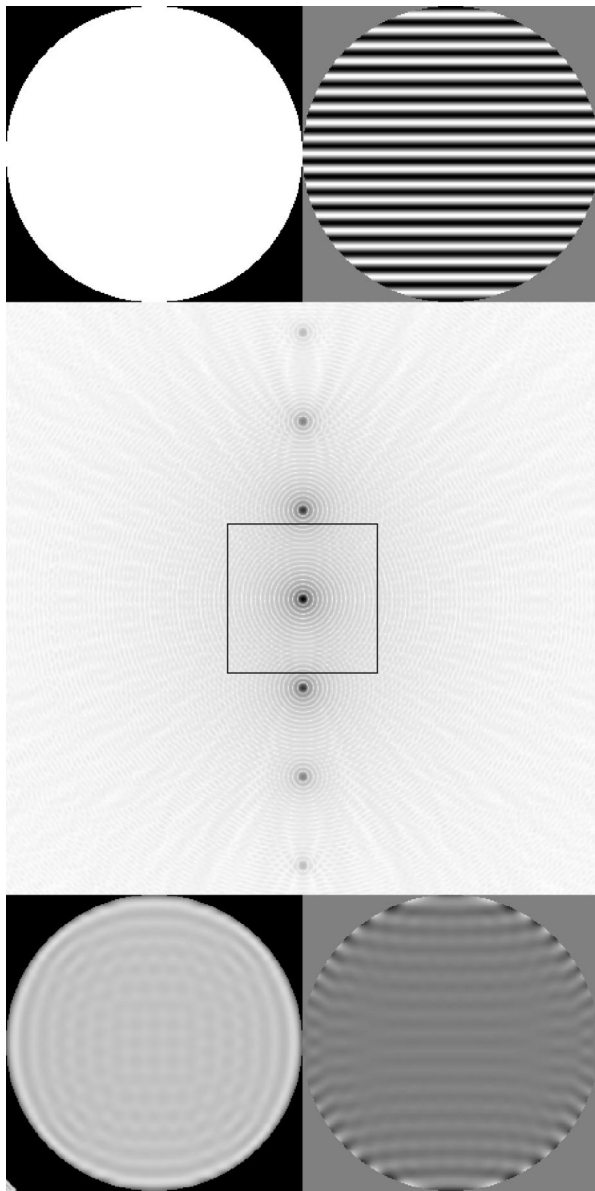
Fig. 1.   Illustration of how the spatial filter works, for a very-large-amplitude $\{\cos[2\pi y(0.6/d)]\}$ phase aberration. Top, magnitude and phase of the field at the pupil plane, linear scale. Middle, magnitude of the field at the focal plane, inverse log scale. Note the large lumps at multiples of the frequency location $0.6/d$ on the sides. The black square is the field-stop edge. Bottom, magnitude and phase in the next pupil plane, same linear scale as top row. The magnitude is slightly reduced, owing to the amount of light rejected at the field stop. The phase is mostly cleaned up, though not entirely.

phase component. If the phase frequency is less than $1/2d$, the first lump will be passed. The lumps farther out (which have much less light) will be cut off, introducing small errors in the passed phase. The final column of the figure shows the filtered magnitude and phase at the next pupil plane. Note the suppression of most of the power of the phase aberration.

The transition zone of the filter from passband to rejection is limited by the diffraction limit of the field $A(X, Y)$. For a circular aperture of diameter D, the FWHM of the Airy disk is at radius $0.5\lambda/D$. This means that the pupil size prevents phase errors at spatial frequencies from

$(1/2d - 1/2D)$ to $1/2d$ from being cleanly passed and errors from $1/2d$ to $(1/2d + 1/2D)$ from being cleanly rejected. For a high-order system this transition region is very small: 64 subapertures across the pupil results in nonideal behavior in the region of $1/2d \pm 0.0156 \times 1/2d$, or $\pm 1.56\%$ of the range.

A final note on the expansion of the field is worth making. We can examine the PSF in terms of $A(X, Y)$ and $\Phi(X, Y)$. The PSF is simply the squared magnitude of the field. This produces a matrix of terms, which have been explored in detail.[4,5] For our present purposes, the most important term is $|A(X, Y)*\Phi(X, Y)|^2$, which is labeled the second-order halo term. This is simply the PSD of the phase $\phi(x, y)$ when the phase is windowed by the aperture $a(x, y)$. This should be the dominant contributor to the halo of the PSF, especially when the aperture is apodized. Here we see the direct, though approximate (since many higher-order terms have been neglected), relationship between the PSD and the PSF. This result will be examined in more detail in the following sections.

B.   Performance Characteristics

On the basis of the above analysis, the SFWFS will behave as a nonideal low-pass filter. Except when very close to $1/2d$, a HSF phase aberration produces scattered light that is cleanly and virtually completely rejected. The filter does not fully pass a LSF phase, however, because of the lumps that appear at multiples of the characteristic frequency, as shown clearly in Fig. 1.

This fundamental nonlinearity is due to the way the field is formed and places a limit on SFWFS performance. Consider a phase aberration perfectly band limited to $1/2d$. For low amounts of power, nearly all the scattered light is contained within the $\lambda/2d$ radius of the field stop. As the phase power increases, the scattered light begins to bleed outside this area. By the time the phase aberration is large (i.e., typical of open-loop atmospheric turbulence), the phase has bled out to at least twice the size of the field stop. Figure 2 compares the fields in the focal plane as generated by two distinct components of full $r_0 = 20$ cm atmospheric turbulence: the ideal LSF portion and the ideal HSF portion. The LSF component is dramatically bled over the field stop. The HSF component is almost entirely contained in the region outside the field stop.

This bleeding has a significant practical impact on the effectiveness of the spatial filter. When a pure LSF phase aberration is passed through the stop, the light scattered out beyond $\lambda/2d$ is rejected. When this new field, which equals zero outside $\lambda/2d$, is propagated back to the pupil plane, the phase aberration *gains* HSF content. This new HSF content is necessary to form a field that is zero outside the field stop. In the opposite scenario, when a pure HSF phase aberration is passed through the stop, virtually no HSF phase aberration remains after filter. In short, the SFWFS rejects all the HSF phase components of the true aberration. The large amount of power in the LSF components cause *new* HSF phase aberration to appear after the field stop. In the atmospheric case, this level of HSF power is lower than in the original phase aberration but is still significant. The
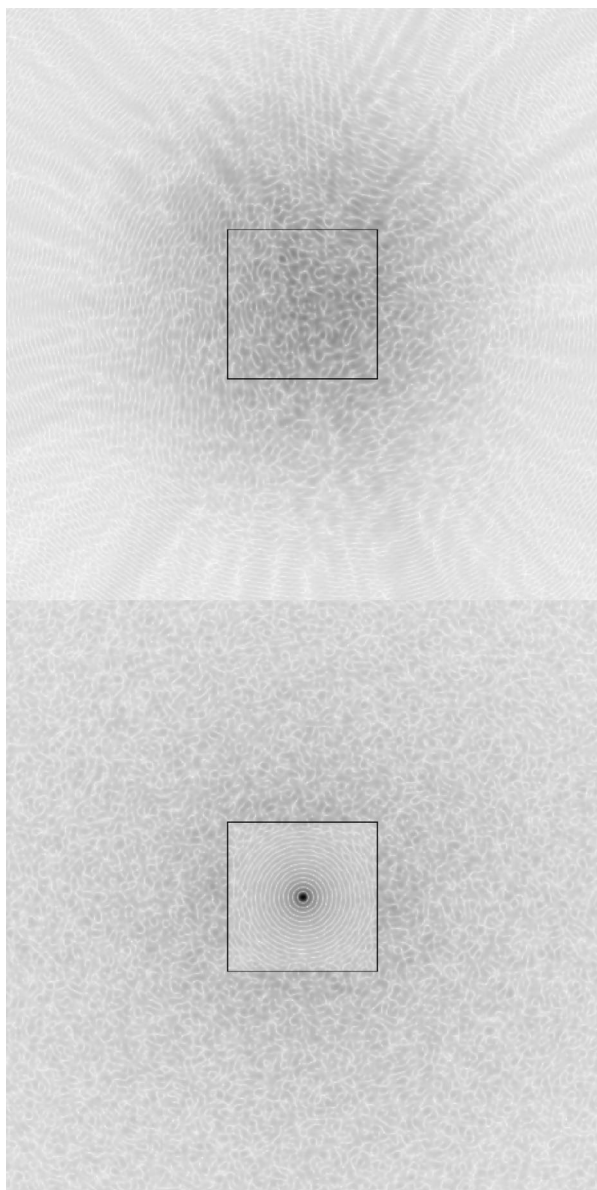
Fig. 2.　Magnitude of the field at the field stop for $r_0 = 20$ cm atmospheric turbulence.　The field out to $2\lambda/d$ from the center is shown, inverse log scale.　The field-stop cutoff edge is shown by the black square.　The atmospheric phase aberration was split into ideal low-pass and ideal high-pass components at the spatial frequency $1/2d$.　Top, field due to the low-frequency portion, which bleeds significantly out beyond the field stop.　Bottom, field due to the high-frequency portion, which scatters light almost exclusively beyond the field stop.　This bleeding leads to reduced spatial filter performance, which is ameliorated by closed-loop AO operation.

best way to reduce this bleeding, and to improve the rejection of HSF phase power by the SFWFS, is to run the AO system in closed loop.　In closed-loop operation, the total LSF power is significantly reduced, mitigating the bleeding phenomenon to a substantial degree.

C.　System Size and Pupil Illumination

The size of the WFS subapertures (the sampling interval) and the wavelength of light set the size of the field stop. For a fixed aperture size, as the number of subapertures

decreases, the field stop gets smaller.　For the spatial filter, this has two important consequences.

First, the size of the atmosphere-degraded spot can be on the order of the field-stop size.　If the spot is significantly larger than the field stop, large amounts of light will be removed from the system.　This results in a lower overall signal-to-noise ratio.　A reasonable case is $r_0 = 20$ cm on an 8-m telescope with WFS at $\lambda = 800$ nm. If the AO system has 32 subapertures ($d = 25$ cm) across the aperture, the SFWFS will pass only 64% of the total light, measured as total magnitude of the pupil function in the pupil plane.　Increasing the system size to 48 subapertures ($d = 16.7$ cm) increases this throughput to 81%.　A 64-across system ($d = 12.5$ cm) has throughput of 92%.　These numbers will improve in closed loop, owing to the suppression of the bleeding.　A typical AO-corrected case for $D = 8$ m would see throughputs of 32 subapertures, 87%; 48 subapertures, 93%; 64 subapertures, 96%.　Note that this throughput affects only the WFS leg of the system.

The second problem that occurs is peculiar to small field stops.　When the phase aberration is large and the field stop is small (approximately $d \geqslant 1.5r_0$) the loss of light results in significant fluctuations in field magnitude across the pupil.　In extreme cases, portions of the pupil can have as little as 20% illumination.　This means that some Shack–Hartmann lenslets will not be able to make correct phase measurements.　The primary reason for this is that the low field amplitude leads to a low number of received photons and hence a lower level of signal in that subaperture.　A secondary effect is that variations in field intensity in a subaperture can lead to spots of abnormal shape, which reduces centroiding accuracy.

3.　SIMULATION DETAILS

A.　Strategy

In even the simplest cases the integrals describing the filtering process do not have closed form.　We have therefore relied on numerical techniques to estimate SFWFS performance.　There are two key issues to be addressed. First, how well does the SFWFS reject HSF phase power? Second, when integrated into an AO system, how does this translate into a reduction in aliasing error?

This first question is addressed by studying the SFWFS in isolation.　Phase aberrations of specific PSDs are filtered, and the phases before and after the SFWFS are compared.　This allows a direct measurement of how effective the filter is.　As shown above, the spatial filter performs poorly when there is substantial power at low spatial frequencies; it is important to verify and characterize this behavior.

The second question is addressed by incorporating the spatial filter into a complete closed-loop end-to-end AO simulation.　This provides an estimate of how phase-rejection performance translates into aliasing-suppression performance.　Furthermore, it allows comparison of phase residuals and PSFs of regular WFS and SFWFS operation.

B.　Numerical Methods

The spatial filter is simulated with Fourier optics.　Given a complex field (which contains the phase aberration) in

the pupil plane of the telescope, the field is Fourier transformed to the focal plane. The field stop is then applied, which makes the field zero outside the edge of the stop. The field is then Fourier transformed back to the pupil plane. Phase unwrapping is conducted, if necessary. For examination of the filter alone, the phase of the field is analyzed. In the end-to-end simulation (see below), the field is then passed along the WFS.

Twice-Nyquist sampling of the PSF is used. The use of 8×8 samples per subaperture in the pupil plane results in a maximum spatial frequency of $4/d$ for the phase. The input phase can be either a deterministic signal or a realization of noise from a specific PSD. In both cases the spatial-frequency content of the phase signals before and after the field stop are analyzed and compared. Doing so correctly can be a fairly subtle endeavor. For a detailed discussion of the spectrum estimation issue, see a standard text such as Ref. 10.

The are two key steps to estimating the filter response. First, the frequency content of the signals must be determined. Second, these PSDs need to be used correctly to estimate the filter response. In all cases we seek to analyze the phase only within a circular aperture (see Fig. 1). Simply Fourier transforming the apertured phase directly leads to artifacts, which include substantial leakage of power from its true frequency to other frequencies. This is conceptually similar to the diffraction of light from the core to large radii in a PSF when the aperture is hard edged. These artifacts are substantially reduced with use of a radially symmetric Blackman window. The formula for this window is, with use of the variable $r = (x^2 + y^2)^{1/2}$,

$$b(r) = 0.42 - 0.5 \cos\left[2\pi\left(\frac{r}{D} + \frac{1}{2}\right)\right]$$
$$+ 0.08 \cos\left[4\pi\left(\frac{r}{D} + \frac{1}{2}\right)\right], \qquad (6)$$

for $|r| < D/2$ and equal to zero elsewhere. This produces a window that smoothly tapers off to zero at the edge of the aperture. This window greatly reduces leakage but at the cost of lowering the resolution of the PSD estimate. Again drawing a parallel with optics, this is similar to the fact that apodization of a pupil leads to both suppressed diffraction and a broader core of the PSF. The output PSDs can have a very large dynamic range, and even given the leakage-suppressing window, some artifacts do still occur. We have carefully determined when leakage obscures the true dynamic range of the PSD and noted this in the text and figures.

In the deterministic case, we have direct access to the Fourier transforms of the input and output signals. The filter response function H is simply the ratio of output to input. The magnitude squared $|H|^2$ of the filter is then determined from the complex response. For the stochastic case we estimate the PSDs of the input and output phase random processes over many realizations with the averaged modified periodogram method. The periodogram is the magnitude squared of the discrete Fourier transform of the random signal. The term "modified" refers to the fact that the signal is windowed, in our case by

multiplication with the Blackman window. The periodograms are averaged over many independent realizations of the noise to produce the averaged modified periodogram estimate of the PSD of the random process. The magnitude squared $|H|^2$ of the filter response is then obtained directly by the ratio of the estimated PSDs of the output and input processes. We confirmed the efficacy of this method by estimating the PSD of the input noise and comparing it with the known PSD used to generate the noise.

After analyzing the SFWFS alone on inputs of different types of noise, we integrate the filter into our end-to-end AO simulation code. AO is applied to static phase errors with no noise to isolate the effects of the AO control system on the output phase of the spatial filter. The baseline design for this is based on a proposed high-order AO system for direct imaging of extrasolar planets.[6] In this case we chose the upper end of the number of subapertures, with $D/d = 62$ across a 10-m mirror. The phase aberration is assumed to be pupil conjugate, and the field at the pupil is created from the aperture function and the phase aberration. This field is sent through the spatial filter code (see above) and then on to the WFS. The WFS is a Shack–Hartmann with 4×4 pixels per subaperture. WFS detector pixel values were created with Fourier optics, and the slope was estimated with a centroid calculation. No Poisson or read noise was added to the pixels on the WFS. The WFS nominally uses 800-nm light. Given the WFS measurements, the phase estimation was done with the Fourier transform reconstructor,[11] and the commands to send to the DM were determined by one-pole control with gain 0.3 and integrator 0.999. The correction was applied to the DM (which has an influence function of a sinc function times a Gaussian) and then to the phase aberration with a single-step delay.

The simulation was run until the error reached a steady state. This residual contains fitting error and aliasing error (when the SFWFS is not used) but does not capture WFS noise error or temporal errors due to phase evolution. Because the SFWFS does not perform as well when there is significant LSF power (see above) a single-step correction of the phase aberration will not fully capture the performance enhancements that it provides. This multistep approach was done solely to reduce open-loop power to closed-loop levels in a self-consistent manner.

After loop convergence the residual error of the system was compared with the input phase to reveal how well the AO system corrected inside the $1/2d$ cutoff. In this case the comparison is between regular WFS operation and the SFWFS. The PSDs of the residual phases and the PSFs generated from these residuals are compared. Diffraction must be suppressed, so the pupil is apodized by using the Blackman window as the aperture in PSF generation. Use of the same function for both windowing and apodization allows easy comparison of the PSFs and PSDs. The science wavelength is 1.6 μm.

4. VERIFICATION OF THE BLEEDING EFFECT

The preceding analysis has shown that when large amounts of power are present at low spatial frequencies,

the filter suppression of HSF power will be poor. To verify this, we conducted simulations of the SFWFS by using random realizations of input phase based on atmospheric turbulence. To generate the phase realizations, we used the spectral factor method for generating von Kármán turbulence.[12] In particular, the shaping filter in this case is proportional to

$$r_0^{-5/6}(f_x^2 + f_y^2)^{-11/12}, \tag{7}$$

where r_0 is the atmospheric coherence length and the filter has radial symmetry. Realizations were produced by filtering white noise through the appropriate shaping filter.

For this trial, $r_0 = 20$ cm (at 500-nm wavelength) was used for a 10-m telescope with 32 subapertures across the pupil. In open loop, the SFWFS reduced HSF power by less than a factor of 10. See Fig. 3, where the solid curve is this filter response.

To test the supposition that simply reducing LSF power will improve filter performance, we reduced LSF power as if AO had been applied. We did this by modifying the atmospheric PSD by a parabolic filter up to $1/2d$ to simulate AO correction.[13] The maximum power reduction due to AO for this case was 1000 times for the lowest spatial frequencies. Such a filter was used to show clearly the effects of reducing the phase power, without any added complications of using a full simulation on an AO system. This allowed isolation of the SFWFS from the effects of the AO component models.

Given this new PSD for closed-loop AO, the filter performance was again examined. The improvement in performance is dramatic. Curves for closed-loop SFWFS performance are shown in Fig. 3 for 15, 20, and 25 cm r_0. In all the AO-corrected cases the performance is drastically better. In closed loop the filter suppresses HSF phase power by a factor of 10^3 to 10^7, with slight variations given r_0. The numerical limit imposed on the estimate by the windowing and dynamic range of the PSD is shown as the "numerical limit" curve in the figure.
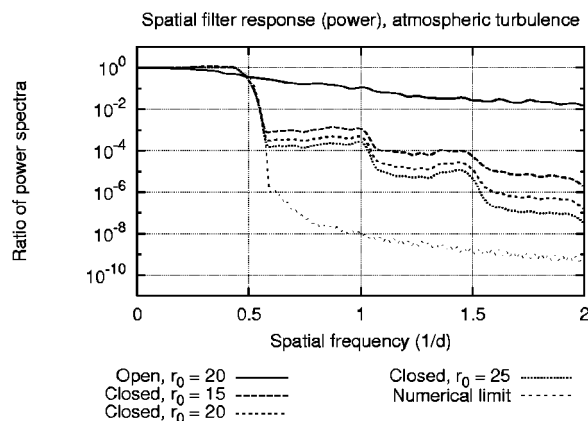

Fig. 3. Response of the spatial filter to atmospheric turbulence, as a ratio of input and output phase power spectra at the spatial filter. For open-loop operation, the filter performs poorly. Closed-loop operation was simulated by application of a parabolic filter to the input phase. For r_0 ranging from 15 cm to 25 cm (at 500 nm), the filtering suppresses HSF phase by a factor of 10^3 to 10^7. As LSF power is reduced, either by AO or by an increase in r_0, the performance of the spatial filter improves.

5. SPATIALLY FILTERED WAVE-FRONT SENSOR INTEGRATED WITH ADAPTIVE OPTICS PERFORMANCE

Numerical simulations, as described above, have established that in open loop, SFWFS performance can be poor. But in closed loop, the SFWFS achieves rejections of HSF phase power by factors of greater than 10^3. In this section we examine whether the SFWFS can still achieve these levels of performance when integrated in a reasonable end-to-end AO simulation. We consider the atmospheric turbulence case discussed above as well as the case of a segmented primary mirror with phase discontinuities.

A. Segmented-Primary-Mirror Phase Aberration
A segmented primary mirror (or "primary") provides an interesting test case for the spatial filter. Edge discontinuities between the segments (due either to phasing or internal segment shape aberrations) have significant high-frequency content. These aberrations are static or quasi-static, perhaps changing slowly as a result of flexure. Because they persist, the aliased phase errors from the segments will fundamentally limit the sensitivity of the PSF. For our simulations we have created a reasonable model based on the actual errors of the segmented primary mirror of the Keck telescopes.[14]

For this model we use the same hexagonal segment geometry as in the Keck. Each segment has internal aberrations, primarily from focus and astigmatism, but also a slight dimpling in the center that has been measured experimentally on real segments.[15,16] The segments also have a small amount of tip and tilt. These phase aberrations produce an average rms surface error of 60 nm on each segment. The segment edges do not line up correctly. On the basis of experimental analysis,[14] the set of 78 edge-height differences between the center of the segment edges in the pupil has an rms value of close to 80 nm, with a maximum difference of magnitude of ~170 nm. This profile for the edge differences is consistent with the simulation results. Accounting for the factor of 2 in the phase error due to reflection off the primary, the rms phase error across the entire mirror surface is 110 nm. This amounts to a Strehl ratio at the science wavelength (1.6 μm) of 0.83 and a Strehl ratio at the sensing wavelength (0.8 μm) of 0.47.

B. Segmented Primary Mirror with Adaptive Optics
The same phase aberration was studied for a varying number of subapertures, with sampling adjusted appropriately. In the first step of correction, the SFWFS did poorly at reducing the phase power at HSF. As shown in Fig. 4 by the light dotted curve, the response of the SFWFS to the segmented-primary-mirror phase error is slight. (The response is obtained from the discrete Fourier transform of a single deterministic signal, hence the fluctuations in the curve.) After letting the end-to-end AO simulation converge (see above for more details) the performance of the SFWFS improved dramatically, as expected. For all three cases (30, 50, and 62 subapertures across the 10-m primary) the HSF phase power was sup-
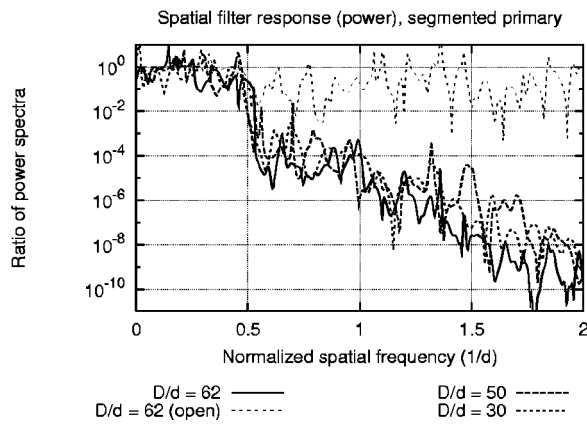
Fig. 4. Response of the spatial filter to segmented-primary phase errors, as a ratio of input and output phase power spectra at the spatial filter. For open-loop operation, the filter performs poorly. End-to-end closed-loop AO operation with no noise (for a varying numbers of subapertures D/d = 30, 50, and 62) results in significantly better spatial filter performance.
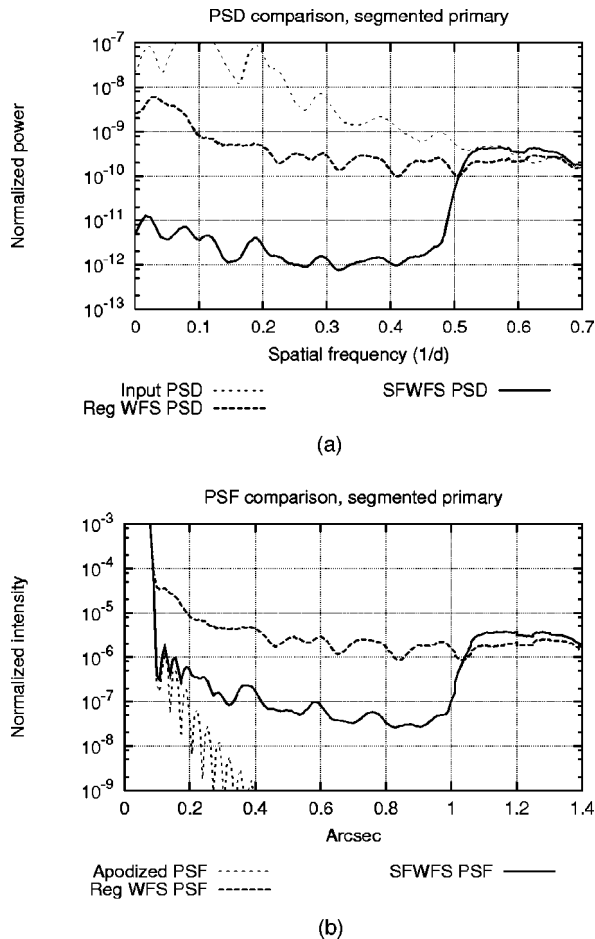
pressed by a factor of 10^4 to more than 10^8. This shows that in this case the AO components do not limit the SF-WFS performance.

Using the end-to-end AO simulation, we produced a residual phase. The PSDs of the two residuals are shown in Fig. 5(a). The prevention of aliasing results in 250 times less residual phase power at LSFs. The residual phase is converted into a PSF by use of apodization to suppress diffraction (see above for details). The PSFs generated by the two residuals are compared directly in Fig. 5(b). The basin is formed in the PSF by the SFWFS, as expected. The intensity in the basin is up to 50 times lower than in the regular WFS case.

C. Atmospheric Turbulence with Adaptive Optics
To confirm that the closed-loop operation on full atmosphere is likely to work, this case was tested in the end-
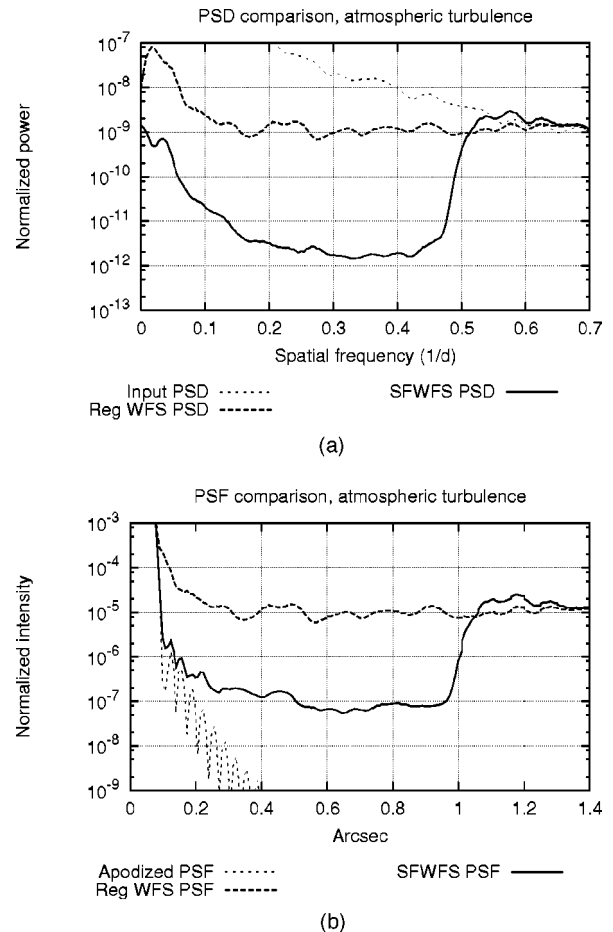

Fig. 5. Comparison of regular WFS with SFWFS after end-to-end AO simulation (D = 10 m, d = 16 cm, D/d = 62). The input-phase aberration is a static segmented-primary-mirror error of 110 nm rms, and there is no WFS noise. (a) PSD comparison. Reduction of aliasing results in 250 times less power in the controllable range in the phase residual. (b) PSF comparison (radial average, diffraction suppressed). This becomes a reduction of PSF intensity in the basin by a factor of 50.

Fig. 6. Comparison of regular WFS and SFWFS after end-to-end AO simulation (D = 10 m, d = 16 cm, D/d = 62). The input-phase aberration is static 20 cm r_0 atmospheric turbulence, and there is no WFS noise. (a) PSD comparison. Reduction of aliasing results in 625 times less power in the controllable range in the phase residual. (b) PSF comparison (radial average, diffraction suppressed). This becomes a reduction of PSF intensity in the basin by a factor of 100. The gain in performance in this case is larger than that in the segmented-primary-mirror case (Fig. 5), because there is more power to prevent in HSF aliasing.

to-end AO simulation. The same system parameters as in the segmented-primary case were used, except that the phase aberration was static atmospheric turbulence with $r_0 = 20$ cm. After AO operation, the residuals both with and without the SFWFS were compared. As shown in Fig. 6(a), use of the SFWFS to reduce aliasing leads to a 625-times reduction in power at LSFs. The radially averaged PSFs, as shown in Fig. 6(b), exhibit the same differences as in the segmented-primary-mirror case. In this case the basin has intensity reduced by a factor of 100 out to the control limit set by $1/2d$.

6. LIMITS ON PERFORMANCE WITH ADAPTIVE OPTICS

The results in the two sections above show how very effective the SFWFS can be at reducing phase power and PSF intensity due to scattered light. The exact performance, however, is dependent on the metric used. Direct examination of the phase PSDs before and after the filter show that reductions up to 10^8 in HSF power can be achieved (see Figs. 3 and 4). Comparisons of the phase residuals after closed-loop AO operation with and without the SFWFS demonstrate up to a 625-times reduction in power at LSFs, which is a significantly lower level of improvement [see Figs. 5(a) and 6(a)]. Finally, the conversion of those phase residuals to PSFs reduces the improvement even more, to a factor of only 50 to 100 times [see Figs. 5(b) and 6(b)].

The first difference is between the phase rejection at HSF as seen by the SFWFS and the level of LSF correction after measurement and correction with an AO system. In the aliased AO case, the uncorrected power at HSF (which is aliased in to LSF) is comparable with the LSF residual. In the SFWFS case, the power at HSF is so low that any aliasing that occurs is negligible compared with the LSF residual. The exact level of the LSF residual is due to the AO system itself. Despite the simulation being noise free, it does capture the effects of the Shack–Hartmann WFS spot formation and centroiding, the phase reconstruction, and the DM response. These factors, combined with the control loop, lead to the limit on correction of LSFs. In this simulation the system performance is solely AO limited, not SFWFS limited.

The second discrepancy is between the performance improvement in the residual PSD and in the PSF. In both the atmospheric-turbulence and segmented-primary cases, the SFWFS gives less of an improvement in the PSF than it does in the PSD. This loss of performance is due to the process of PSF formation. In the aliased WFS case, where the residual power at LSF is comparable with that at HSF, the approximation that the PSF halo equals the PSD is a very good one. The second-order halo term $|A(X, Y)*\Phi(X, Y)|^2$ is exactly the PSD since we used the Blackman window for both PSD estimation and pupil apodization. This term dominates the PSF in the halo.

However, as Figs. 5 and 6 show, when the SFWFS is used the PSF basin is not as low as would be expected on the basis of the PSD. The source of the extra light in the basin is a fourth-order term in the PSF expansion: $|A(X, Y)*\Phi(X, Y)*\Phi(X, Y)|^2$. This was determined by analyzing the individual terms in the PSF expansion and

determining the overall contribution of each to the true PSF. This higher-order term becomes visible because of the significant difference in power between the LSF and HSF zones. The convolution of $\Phi(X, Y)$ causes phase power from beyond the controllable region to move in to the controllable region of the basin, where it can be larger than the second-order halo term.

The above may be a special case that contradicts the assertion[4] that the second-order halo term places the ultimate limit on sensitivity of the PSF. The behavior of the fourth-order term is similar in origin to a phenomenon noted by Give'on and Kasdin.[17] We are currently investigating the fundamental limits imposed by the fourth-order term and mitigation strategies. The Blackman window produced answers virtually indistinguishable from the Kaiser window, which is the approximation to the mathematically optimal window.[10] The optimality rests on the assumption of no phase errors. It is possible that the optimal apodization for the AO with SFWFS residual is different.

7. BROADBAND WAVE-FRONT SENSING

The most realistic configuration for the WFS is that it uses a range of wavelengths of light to do wave-front sensing. This poses a challenge for the SFWFS, since the structure of the field is directly dependent on the wavelength of the light forming the field.

Suppose a band of light is used, from shortest wavelength λ_s to longest wavelength λ_l. For simplicity, the spectrum of the AO guide star in the band is assumed to be uniform. A nonuniform spectrum will simply change the degree, but not the location, of the performance degradation. The field stop in the SFWFS is designed for a wavelength λ_f, which is as yet unspecified but is inside the band. The short λ_s light will form a field that is narrower than that of the filter wavelength. The field stop will now let in spatial frequencies up to $(\lambda_f/\lambda_s)/2d$. This unwanted HSF content will be aliased into the wave front but only over a small range of spatial frequencies. Aliasing will occur down to $(2 - \lambda_f/\lambda_s)/2d$. For the long λ_l light, the field stop will be undersized and will have a limit of $(\lambda_f/\lambda_l)/2d$. The top wavelength light will be missing information at the high end of the LSF range.

Both of these properties reduce AO performance in the range of frequencies close to $1/2d$. If the goal is to minimize the zone where phase information is corrupted by either suffering from aliasing or missing information, the best wavelength λ_f for which to design the filter is $2\lambda_s\lambda_l/(\lambda_s + \lambda_l)$. This will give an effective cutoff for the filter of $2\lambda_s/(\lambda_s + \lambda_l)$ times the $1/2d$ desired cutoff.

To demonstrate this performance, we ran complete multiwavelength closed-loop AO simulations on the segmented-primary-mirror case. Five wavelengths spanning the band from 700 to 900 nm were used in the simulation. The field stop was set for the average wavelength of 800 nm. On the basis of the above analysis, we would expect aliasing to affect the phase down to $0.857/2d$; information will be missing down to $0.889/2d$. This translates into a PSF that should degrade in quality at 0.876 arcsec, in contrast to the monochromatic case, where degradation occurs at 1.0 arcsec. As shown in Fig. 7, the
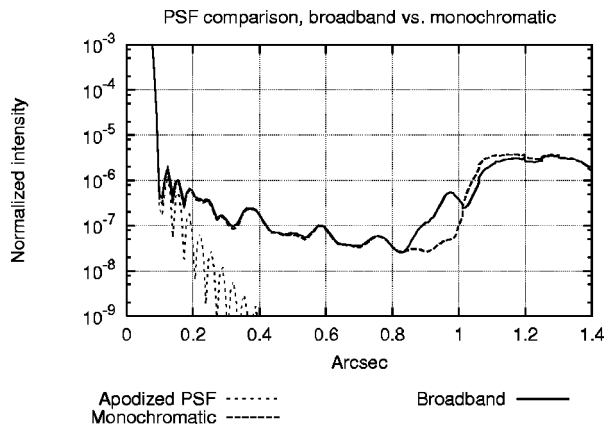
Fig. 7. Comparison of monochromatic (800 nm) and broadband (700–900 nm) SFWFS after end-to-end AO simulation (D = 10 m, d = 16 cm, D/d = 62).   Same segmented-primary-mirror errors as in Fig. 6. PSFs (radial average, diffraction suppressed) are shown.   The only significant difference in PSFs occurs just inside the spatial-frequency cutoff, where the broadband case has increased error owing to both aliasing and incomplete sensing of frequencies near the cutoff.

PSF when sensed with broadband light has this pattern of increased error near the cutoff.

Broadband light with the SFWFS produces a small area of reduced performance that is confined to a region just inside the cutoff frequency $1/2d$.   The size of this area is determined by width of the band and the exact wavelength of light for which the field stop is designed. Care must be taken not to make the band too broad: A band from 500 nm to 1 μm would have, in the best case, an improved area in the PSD and PSF only two thirds as far out from the core as the monochromatic case.

8. MODIFICATIONS TO BASIC FILTER DESIGN

There are several details of implementation of the SF-WFS that could be modified.   In this section we focus on changing the shape of the pupil as seen by the SFWFS and on changing the shape of the edges of the field stop.

Changing the shape of the pupil as seen by the filter is the most obvious choice, as diffractive effects are clearly visible in the field at the stop (see Fig. 1).   However, the bleeding of the field out beyond the cutoff imposed by spatial frequency is not due to diffraction—it is due to the field-formation process.   So apodization of the pupil will not help in this regime.   Where it will help is when the overall phase power is low.   Suppressing diffraction will then reduce leakage from HSF content back inside the field stop and improve suppression.

Simulations with the cases of both full atmospheric turbulence and AO-corrected atmosphere confirm the above finding.   A slight smoothed-edge apodization of the pupil was used.   For any reasonable r_0, pupil apodization makes no difference in the open-loop atmosphere case. In the AO-corrected case, improvement in rejection was seen for large r_0, corresponding to 40 nm or less of total rms phase error.   Apodization appears to be helpful only in cases where overall phase aberration is already quite small.

A second option to reduce diffractive effects is to apodize the field stop itself.   The only significant effect of smoothing the edge transition, however, is to make the filter transition from passband to rejection less sharp.   In no scenario did apodization of the field stop improve performance.

9. CONCLUSIONS

Detailed analysis with the help of Fourier optics simulations has shown that the spatially filtered WFS can substantially improve the performance of an AO system over a specific range of spatial frequencies, resulting in an improved PSF with higher sensitivity.   The PSF will have a dark basin in the controllable region that has been cleaned out as a result of the removal of aliasing.   The field stop very effectively cuts off phase aberrations beyond $1/2d$ in frequency.   Significant low-spatial-frequency power, however, can reduce filter performance. Owing to the nature of field formation, power from low-spatial-frequency phase aberrations bleeds out beyond the field stop and is converted into high-spatial-frequency power.   This behavior places a limit on the effectiveness of the spatial filter when there is a significant amount of low-spatial-frequency power, i.e., in open loop.   Closed-loop performance, however, solves this problem and produces substantial high-spatial-frequency power rejections.   For both full atmospheric turbulence and a segmented primary mirror, the SFWFS can achieve suppression of phase power past $1/2d$ by 10^3 to 10^8.   When integrated into a full noise-free end-to-end AO simulation, this translates up to a 625-times reduction in phase power in the area affected by aliasing.   With the use of apodization, improvements of up to a 100-times reduction in intensity can be obtained in the basin of the PSF.   The broadband WFS case has been analyzed.   It has been confirmed through simulation that the broadband of light will reduce the practical cutoff frequency of the spatial filter, leading to smaller regions where full performance benefits are obtained.

On large telescopes, small AO systems (e.g., those with fewer than 16 subapertures across the pupil) may benefit only slightly from use of the SFWFS.   The proposed 4096-actuator eXtreme Adaptive Optics Planet Imager,[6] however, will implement a spatial filter.   In a future paper, the authors will explore the effects of this approach on high-contrast imaging sensitivity.

The corresponding author, Lisa A. Poyneer, can be reached at poyneer1@llnl.gov.

REFERENCES

1. J. W. Hardy, *Adaptive Optics for Astronomical Telescopes* (Oxford U. Press, Oxford, UK, 1998).
2. A. V. Oppenheim and A. S. Willsky, *Signals and Systems*, 2nd ed. (Prentice-Hall, Upper Saddle River, N.J., 1997).
3. F. Rigaut, J.-P. Veran, and O. Lai, "An analytical model for Shack–Hartmann-based adaptive optics systems," in *Adaptive Optical System Technologies 3*, D. Bonaccini and R. K. Tyson, eds., Proc. SPIE **3353**, 1038–1048 (1998).
4. A. Sivaramakrishnan, J. P. Lloyd, P. E. Hodge, and B. A. Macintosh, "Speckle decorrelation and dynamic range in speckle noise-limited imaging," Astrophys. J. Lett. **581**, L59–L62 (2002).
5. M. D. Perrin, A. Sivaramakrishnan, R. B. Makidon, B. R. Oppenheimer, and J. R. Graham, "The structure of high Strehl ratio point-spread functions," Astrophys. J. **596**, 702–712 (2003).
6. B. Macintosh, J. Graham, L. Poyneer, G. Sommargren, J. Wilhelmsen, D. Gavel, S. Jones, P. Kalas, J. P. Lloyd, R. Makidon, S. Olivier, D. Palmer, J. Patience, M. Perrin, S. Severson, A. Sheinis, A. Sivaramakrishnan, M. Troy, and J. K. Wallace, "Extreme adaptive optics planet imager: XAOPI," *Techniques and Instrumentation for Detection of Exoplanets*, D. R. Coulter, ed., Proc. SPIE **5170**, 272–282 (2003).
7. B. Macintosh, S. S. Olivier, B. J. Bauman, J. M. Brase, E. Carr, C. J. Carrano, D. T. Gavel, C. E. Max, and J. Patience, "Practical high-order adaptive optics systems for extrasolar planet searches," in *Adaptive Optics Systems and Technology II*, R. K. Tyson, D. Bonaccini, and M. C. Roggemann, eds., Proc. SPIE **4494**, 60–68 (2002).
8. M. Feldt, T. F. Henning, S. Hippler, R. Weiss, M. Turatto, R. Neuhauser, A. P. Hatzes, H. M. Schmid, R. Waters, E. Puga, and J. Costa, "Can we really go for direct exo-planet detection from the ground?" in *High-contrast Imaging for Exoplanet Detection*, A. B. Schultz and R. G. Lyon, eds., Proc. SPIE **4860**, 149–160 (2003).
9. C. B. Hogge, R. R. Butts, and M. Birlakoff, "Characteristics of phase-aberrated nondiffraction-limited laser beams," Appl. Opt. **13**, 1065–1070 (1974).
10. A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing* (Prentice-Hall, Upper Saddle River, N.J., 1989).
11. L. A. Poyneer, M. Troy, B. Macintosh, and D. T. Gavel, "Experimental validation of Fourier transform wave-front reconstruction at Palomar Observatory," Opt. Lett. **28**, 798–800 (2003).
12. E. M. Johansson and D. T. Gavel, "Simulation of stellar speckle imaging," in *Amplitude and Intensity Spatial Interferometry II*, J. B. Breckinridge, ed., Proc. SPIE **2200**, 372–383 (1994).
13. A. Sivaramakrishnan, C. D. Koresko, R. B. Makidon, T. Berkefeld, and M. J. Kuchner, "Ground-based coronagraphy with high-order adaptive optics," Astrophys. J. **552**, 397–408 (2001).
14. G. Chanan, C. O'Hara, and M. Troy, "Phasing the mirror segments of the Keck telescopes II: the narrow-band phasing algorithm," Appl. Opt. **39**, 4706–4714 (2000).
15. M. Troy, personal communication (mtroy@jpl.nasa.gov, Jet Propulsion Lab, Pasadena, California, 2003).
16. C. Carrano, personal communication (cjc@llnl.gov, Lawrence Livermore National Laboratory, Livermore, California, 2003).
17. A. Give'on and J. Kasdin, personal communication (agiveon@princeton.edu, Princeton University, Princeton, New Jersey, 2003).