UCRL-JRNL-207495

Optimal modal Fourier transform wave-front control

Lisa A. Poyneer and Jean-Pierre Véran

October 28, 2004

Submitted to the Journal of the Optical Society of America, A

Optimal modal Fourier transform wave-front control

Lisa A. Poyneer (1) and Jean-Pierre Véran (2)

(1): Lawrence Livermore National Lab
7000 East Ave, Livermore, CA, 94551
(2): Herzberg Institute of Astrophysics
5071 West Saanich Road, Victoria, British Columbia, Canada V9E2E7

poyneer1@llnl.gov, Jean-Pierre.Veran@nrc-cnrc.gc.ca

Optimal modal Fourier transform wave-front control combines the speed of Fourier transform reconstruction (FTR) with real-time optimization of modal gains to form a fast, adaptive wave-front control scheme. Our modal basis is the real Fourier basis, which allows direct control of specific regions of the PSF. We formulate FTR as modal control and show how to measure custom filters. Because the Fourier basis is a tight frame, we can use it on a circular aperture for modal control even though it is not an orthonormal basis. The modal coefficients are available during reconstruction, greatly reducing computational overhead for gain optimization. Simulation results show significant improvements in performance in low-SNR situations compared to non-adaptive control. This scheme is computationally efficient enough to be implemented with off-the-shelf technology for a 2.5 kHz, 64×64 Adaptive Optics system.

OCIS codes: 010.1080, 010.1330

1. Introduction

In single-conjugate Adaptive Optics (AO), a guide star is used to make measurements of the wave-front phase. In Shack-Hartmann systems the wave-front sensor (WFS) is an array of lenslets that measures the average gradient of the phase over each subaperture in the pupil plane.[1] These measurements, commonly called slopes, are then sent to a reconstruction algorithm to estimate the phase. The estimate of the measured phase is processed with a control law and then applied to a deformable mirror (DM) to conjugate the phase aberration. The reconstruction step is usually done with a vector-matrix multiply (VMM) and the control law for the commands does not change with time.

Next generation AO systems require more efficient algorithms to deal with higher frame rates and more phase measurements in the pupil. New systems also provide the opportunity

to develop more sophisticated controllers, such as adaptive or predictive methods, which will provide improved correction. There have been significant advances in the last few years in developing these types of new methods.

One major area of advance has been in computational efficiency. For large systems the VMM method is too expensive, as it has computational complexity of $O(n^3)$ to compute a new control matrix and $O(n^2)$ to apply that control matrix, where n is the number of control points. There are now several possible fast algorithms for the reconstruction step. Most focus on faster ways to solve a set of matrix equations. Under certain assumptions the control matrix can become sparse. This allows more efficient application, as has been shown by Ellerbroek[2] with his minumum-variance unbaised (MVU) approach. MacMartin[3] has developed local methods of control, which use iteration or hierarchic methods to ensure proper global results. Shi[4] has tested these sparse methods at Palomar. Gilles and Vogel[5] have applied the Conjugate Gradient technique to the MVU model to produce $O(n \lg n)$ reconstructors. A new approach of Gilles[6] is $O(n)$ for open-loop reconstruction. The method of Fourier Transform wave-front reconstruction (FTR) was initially proposed by Frieschlad.[7,8] This method is not a way to solve a matrix equation, but rather a way to filter the slope measurements to produce a phase estimate. Use of the Fast Fourier Transform (FFT) makes this algorithm $O(n \lg n)$, and other fast Discrete Fourier Transform (DFT) implementations also allow fast reconstruction. Poyneer[9] developed a method of slope management that made the method usable for circular apertures. Poyneer has further developed new filters to improve performance[10] and has experimentally validated[11] FTR in tests at Palomar.

A second area of improvement has been the incorporation of statistical information into the control process. The MVU model incorporates open-loop atmospheric statistics into the model. For closed-loop control, however, information based on past measurements can be used to improve wave-front control. Building on the work of Wallner,[12] Gavel and Wiberg[13] have shown that continuous updating of statistical information of the wave-front can lead to better wave-front reconstruction, using the specific example of frozen-flow turbulence. Modal control is another approach to this problem that does not require any assumption on the nature of the turbulence. Ellerbroek[14] and Gendron[15] developed in parallel the framework of optimal modal control. By controlling the AO system with a set of orthogonal modes, the control problem is simplified such that the overall wave-front correction is maximized by the independent optimization of the control of each mode. This optimization is based on the statistics of the signal and noise inputs for each mode. Optimal modal control was first implemented in ESO's Come-On+ AO system and the required statistical information was derived during a WFS open-loop acquisition, prior to the AO corrected observation.[16] Later in CFHT's PUEO AO system, the statistical information was derived from the analysis of prior closed-loop data, removing the the undesirable overhead of an open-loop acquisition

and enabling the possibility to continuously update the optimization during the AO corrected observation as turbulence conditions change.[17] The same method, with minor variations, was implemented in Gemini's Altair[18] by Véran and in ESO's NAOS AO systems,[19] enabling a significant increase in observing efficiency at the telescope. All the developments mentioned above assume a simple integrator as a control law: only the gain of the integrator is optimized for each mode. However, Dessenne[20] has generalized the concept of optimal modal control to more sophisticated controllers for further improved wave-front correction.

In this paper we present a new method of FTR which is both computationally efficient and incorporates real-time information about the phase aberration to improve correction. This method is called optimal modal Fourier transform wave-front control (OFC); it uses the real Fourier basis as its natural modal set. We will first present the theory of OFC, including discussions of the Fourier basis set, how to make custom modal filters and how to deal with a circular aperture. We will present the optimality criterion for determining modal gains in closed loop from telemetry. End-to-end simulations will demonstrate that OFC significantly improves AO performance as compared to non-adaptive control in the presence of frozen-flow turbulence. In addition, we will show that we can control individual modes, and hence locations in the PSF and that gain optimization compensates for unmeasured system characteristics. Analysis of the computational costs will show that OFC is implementable on currently available hardware.

2. Fourier transform reconstruction as modal control

The method of FTR was originally derived from a discrete model of WFS behavior. FTR uses the DFT to convert slope measurements to the frequency domain, where a filter which inverts the measurement process is applied. The inverse DFT is then applied to convert the estimate back to actuator space. It was recognized[8] that the natural modal basis for this is method consists of the Fourier complex exponentials. The slope and phase signals in the reconstruction problem are always real, which means that their Fourier transforms are always Hermitian. This means that there is a real modal basis consisting of sines and cosines.[21] In this section we define this modal basis and show how the eigenfunction property allows FTR filters to easily be developed for modal control. We then show how this modal basis is valid on an arbitrary aperture through use of the formalism of a Frame and proper slope management.

2.A. The real Fourier basis

Before we present the real Fourier basis, we need to define the signals, variables and transforms that will be used in the discussion. We will be using the DFT on real-valued signals which are $N \times N$ points. For instance, the x-slope signal will be represented as $x[m, n]$. The

indices into these signals will range from 0 to $N-1$. If an index is outside this range, we will assume N-periodicity in both dimensions. The DFT of a signal is defined as

$$X[k,l] = \frac{1}{N} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} x[m,n] \exp\left(\frac{-j2\pi(km+ln)}{N}\right), \tag{1}$$

where we use the engineer's imaginary number $j = \sqrt{-1}$ and the frequency variables k and l. These frequency variables will range from $-N/2+1$ to $N/2$ to reinforce the fact that the DFT of a signal is sampled across a symmetric frequency range.

For now we will ignore the aperture and instead assume that the slope and phase signals are N-periodic. Because these signals are always purely real, the DFTs are always Hermitian. This means we can construct a set of N^2 real modes which form an orthonormal basis (ONB). There are $N/2+2$ cosine modes and $N/2-2$ sine modes.[21] The exact frequency pairs $[k,l]$ that we use are somewhat arbitrary, as long as only one of the two pairs of frequencies $[k,l]$, $[N-k, N-l]$ are chosen due to the Hermitian symmetry. The set of frequency pairs that forms the ONB is called $\mathcal{M}$. Our numbering convention for modes is shown in Fig. 1.

We have chosen to normalize the modes such that the entire set is an ONB. Four cosine modes exist alone without a sine counterpart. These occur at frequencies $[0,0], [0,N/2], [N/2,0]$ and $[N/2,N/2]$. These cosine modes are

$$\mathcal{C}_{k,l}[m,n] = \frac{1}{N} \cos\left(\frac{2\pi}{N}[km+ln]\right). \tag{2}$$

At the other frequency pairs there are two modes, the sine and the cosine, at the same frequency. The scaling is different to ensure normalization. The cosine is

$$\mathcal{C}_{k,l}[m,n] = \frac{\sqrt{2}}{N} \cos\left(\frac{2\pi}{N}[km+ln]\right), \tag{3}$$

and the sine is

$$\mathcal{S}_{k,l}[m,n] = \frac{\sqrt{2}}{N} \sin\left(\frac{2\pi}{N}[km+ln]\right). \tag{4}$$

These cosines and sines of course exist for $[k,l] \notin \mathcal{M}$. For the cosine modes, we have even symmetry with $\mathcal{C}_{k,l}[m,n] = \mathcal{C}_{N-k,N-l}[m,n]$. For the sine modes we have odd symmetry with $\mathcal{S}_{k,l}[m,n] = -\mathcal{S}_{N-k,N-l}[m,n]$.

We can now use the DFT to calculate the modal coefficients. The DFT analysis equation becomes

$$X[k,l] = D_{k,l} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} x[m,n] \left(\mathcal{C}_{k,l}[m,n] - j\mathcal{S}_{k,l}[m,n]\right), \tag{5}$$

where the scaling the constant $D_{k,l} = 1/\sqrt{2}$ for all values of $[k,l]$ except for $[0,0], [0,N/2], [N/2,0]$ and $[N/2,N/2]$, when $D_{k,l} = 1$. This means that the modal coefficients obtained by the inner products of a signal $x[m,n]$ with all the modes are very easily

extractable from $X[k,l]$. In particular the cosine coefficient is obtained as

$$< x[m,n], \mathcal{C}_{k,l}[m,n] >= \frac{1}{D_{k,l}} \mathrm{Re}\left\{X[k,l]\right\}, \tag{6}$$

and the sine as

$$< x[m,n], \mathcal{S}_{k,l}[m,n] >= \frac{-1}{D_{k,l}} \mathrm{Im}\left\{X[k,l]\right\}, \tag{7}$$

where Re{} takes the real part and Im{} takes the imaginary part.

We now have a real ONB in which the modal coefficients for any signal $x[m,n]$ are obtained with use of the DFT. What makes this formulation particularly useful for wave-front reconstruction is that the Fourier modes are eigenfunctions of linear, shift-invariant (LSI) systems.[22] This means that a complex exponential input at a specific frequency pair $[k,l]$ to an LSI system is output as that exact same complex exponential, scaled by a complex number $H[k,l]$. This function $H[k,l]$ is the frequency response which characterizes that LSI system. The Shack-Hartmann WFS is an LSI system which produces the average gradient over a subregion. The response of the DM is normally modeled as being LSI. Because the Fourier modes are eigenfunctions, an input phase at a specific frequency $[k,l]$ will only ever produce x-slope and y-slope signals at that exact same frequency. We define the response $M_x[k,l]$ as the filter which measures actuators commands placed on the DM with the WFS to produce x-slopes. The y-slope measurement filter is $M_y[k,l]$.

Since the WFS and DM have real-valued impulse responses, the modes have an even more special property. The Hermitian symmetry of the frequency response means that the input of cosine mode $\mathcal{C}_{k,l}[m,n]$ produces an x-slope of $A\mathcal{C}_{k,l}[m,n] + B\mathcal{S}_{k,l}[m,n]$, where $A = \mathrm{Re}\{M_x[k,l]\}$ and $B = -\mathrm{Im}\{M_x[k,l]\}$. The sine mode at that same frequency produces an x-slope of exactly $-B\mathcal{C}_{k,l}[m,n] + A\mathcal{S}_{k,l}[m,n]$. The same holds for the y-slopes, with the cosine mode producing y-slope $C\mathcal{C}_{k,l}[m,n] + D\mathcal{S}_{k,l}[m,n]$ and the sine mode y-slope $-D\mathcal{C}_{k,l}[m,n] + C\mathcal{S}_{k,l}[m,n]$, where $C = \mathrm{Re}\{M_y[k,l]\}$ and $D = -\mathrm{Im}\{M_y[k,l]\}$.

This eigenfunction property results in a simple estimation problem to convert back from x- and y-slope modes to phase modes. In matrix notation, we need to invert the folowing matrix, which converts from the sine and cosine modal coefficients for the phase at a specific frequency $[k,l]$ to the x-slope and y-slope modal coefficients:

$$\begin{bmatrix} A & -B \\ B & A \\ C & -D \\ D & C \end{bmatrix}. \tag{8}$$

The pseudo-inverse which converts from slope modes to phase modes is simply the normalized transpose

$$\frac{1}{A^2 + B^2 + C^2 + D^2} \begin{bmatrix} A & B & C & D \\ -B & A & -D & C \end{bmatrix}. \tag{9}$$

The entries of the pseudoinverse can be directly converted into the reconstruction filter. Recall from above that the projection on to the cosine $\mathcal{C}_{k,l}$ is the scaled real part of $X[k,l]$ and the projection onto the sine $\mathcal{S}_{k,l}$ is the scaled imaginary part of $X[k,l]$. Using the fact that

$$X[k,l] = D_{k,l}\left(< x[m,n], \mathcal{C}_{k,l}[m,n] > -j < x[m,n], \mathcal{S}_{k,l}[m,n] >\right), \tag{10}$$

the pseudoinverse matrix above converts into the DFT of the estimated phase $\hat{P}[k,l]$ as

$$\hat{P}[k,l] = \frac{(A+jB)X[k,l] + (C+jD)Y[k,l]}{|(A+jB)|^2 + |(C+jD)|^2}. \tag{11}$$

Now we step back a bit and look at the initial model. We defined the filters $M_x[k,l]$ and $M_y[k,l]$ which measure the slopes from the actuator commands $P[k,l]$ above. Using the two equations for the forward-process

$$X[k,l] = M_x[k,l]P[k,l], \tag{12}$$

and

$$Y[k,l] = M_y[k,l]P[k,l], \tag{13}$$

the original FTR derivation obtains the best estimate as

$$\hat{P}[k,l] = \frac{M_x^*[k,l]X[k,l] + M_y^*[k,l]Y[k,l]}{|M_x[k,l]|^2 + |M_y[k,l]|^2}, \tag{14}$$

where * denotes complex conjugation. As given above, for a specific frequency pair $[k,l]$, $M_x[k,l] = A - jB$ and $M_y[k,l] = C - jD$. Substituting these expressions into Eq. 14, we obtain exactly Eq. 11, as we should.

2.B. Examples of filters

The modal approach of above allows us to generate customized filters based not just on specific models of WFS and DM behavior, but on simulated or physical AO systems. If we measure the constants A, B, C and D for each frequency pair $[k,l] \in \mathcal{M}$, we can easily form the reconstruction filter.

In this section we will examine three filters. The first is the modified-Hudgin (Mod-Hud) FTR filter which was derived by Poyneer[10] and shown to perform the best of several theoretical filters.[11] This filter assumes that each subaperture measurement is the first difference of the phase at points halfway between actuators bounding the subaperture. These half-pixel shift values of course do not exist, but we can shift the signals with filtering.[22] Using this model, the WFS measurement frequency responses are

$$M_x[k,l] = \exp(j\pi l/N)[\exp(j2\pi k/N) - 1] \tag{15}$$

for the x-slopes and

$$M_y[k,l] = \exp(j\pi k/N)[\exp(j2\pi l/N) - 1] \tag{16}$$

for the y-slopes. Using Eq. 14 above, this produces a filter of

$$\hat{P}[k,l] = \frac{\exp(-j\pi l/N)[\exp(-j2\pi k/N) - 1]X[k,l] + \exp(-j\pi k/N)[\exp(-j2\pi l/N) - 1]Y[k,l]}{4[\sin^2(\pi k/N) + \sin^2(\pi l/N)]}. \tag{17}$$

Recall that this is valid of k, l over the range from $-N/2+1$ to $N/2$. For the piston mode ($[k,l] = [0,0]$) the filter is set to zero to prevent division by zero.

A second way to model the WFS and DM would be to use continuous-domain responses and have the WFS sample the output. This is the approach taken in one dimension by Rigaut;[23] here we present the full two dimensional results. We assume the phase is a continuous signal, scaled along the x- and y-axes such that the subaperture width and sampling interval is 1 unit. In our convention, the WFS slope at point m, n corresponds to the phase region from m to $m+1$ and from n to $n+1$. As mentioned above, the Shack-Hartmann WFS takes the average of the derivative of the phase in a subregion. This produces a response which is a sinc divided by the frequency variable for the derivative's direction. We further assume that the DM is a perfect sinc-interpolator.

Now we can analyze the coefficients A, B, C and D for all modes in the frequency set $\mathcal{M}$. We omit the calculus used to derive the following coefficients. For the x-slopes we have the coefficients A and B as

$$A = \begin{cases} 0, & k = 0 \\ \cos(2\pi k/N) - 1, & k \neq 0, l = 0 \\ (2\pi l/N)^{-1}\{[\cos(2\pi l/N) - 1]\sin(2\pi k/N) + & \\ \qquad [\cos(2\pi k/N) - 1]\sin(2\pi l/N)\} & k \neq 0, l \neq 0 \end{cases}, \tag{18}$$

and

$$B = \begin{cases} 0, & k = 0 \\ -\sin(2\pi k/N), & f_x \neq 0, l = 0 \\ (2\pi l/N)^{-1}\{[\cos(2\pi k/N) - 1]\,[\cos(2\pi l/N) - 1] - & \\ \qquad \sin(2\pi k/N)\sin(2\pi l/N)\} & k \neq 0, l \neq 0 \end{cases}. \tag{19}$$

The coefficients for the y-slopes are very similar, with the frequency variables switched,

$$C = \begin{cases} 0, & l = 0 \\ \cos(2\pi l/N) - 1, & l \neq 0, k = 0 \\ (2\pi f_x)^{-1}\{[\cos(2\pi l/N) - 1]\sin(2\pi k/N) + & \\ \qquad [\cos(2\pi k/N) - 1]\sin(2\pi l/N)\} & l \neq 0, k \neq 0 \end{cases}, \tag{20}$$

and

$$D = \begin{cases} 0, & l = 0 \\ -\sin(2\pi l/N), & l \neq 0, k = 0 \\ (2\pi k/N)^{-1}\left\{[\cos(2\pi k/N) - 1]\,[\cos(2\pi k/N) - 1] - \right. & \\ \qquad \left. \sin(2\pi k/N)\sin(2\pi l/N)\right\} & l \neq 0, k \neq 0 \end{cases}. \tag{21}$$

For the piston mode and the waffle mode ($[k, l] = [N/2, N/2]$) all the coefficients are zero. For the two other cosine-only modes $[N/2, 0]$ and $[0, N/2]$ the coefficients C and D should be automatically set to zero since the sine modes do not actually exist.

We combine these coefficients together to produce a filter (which we shall call Ideal) as specified above by Eq. 11. For piston and waffle the estimated phase is exactly zero,

$$\hat{P}[k, l] = 0. \tag{22}$$

When the y-frequency $l = 0$ but $k \neq 0$ the filter is

$$\hat{P}[k, 0] = \frac{[\exp(-j2\pi k/N) - 1]X[k, l]}{4\sin^2(\pi k/N)}, \tag{23}$$

and in the case of $k = 0$ but $l \neq 0$ the filter is

$$\hat{P}[0, l] = \frac{[\exp(-j2\pi l/N) - 1]Y[k, l]}{4\sin^2(\pi l/N)}, \tag{24}$$

This is exactly the same as the modified-Hudgin FTR filter (Eq. 17) for these frequencies. For all other values of $[k, l] \in \mathcal{M}$ the filter is

$$\hat{P}[k, l] = \frac{j}{8\pi}\frac{[\exp(-j2\pi k/N) - 1][\exp(-j2\pi l/N) - 1]}{\sin^2(\pi k/N)\sin^2(\pi l/N)}\left(\frac{1}{k^{-2} + l^{-2}}\right)\left(\frac{X[k, l]}{l} + \frac{Y[k, l]}{k}\right). \tag{25}$$

This Ideal filter is very similar to the Mod-Hud filter. Only at the highest spatial frequencies do they differ, with the Mod-Hud having lower gain. For $[k, l] \notin \mathcal{M}$ we fill in the filter values using the Hermitian symmetry property.

Our third filter is produced directly from our simulated AO system. This involves placing each mode on the DM, and recording the WFS slopes and measuring the coefficients A, B, C and D empirically. We term this filter Custom. In our simulation we are able to assume a square, N-periodic aperture to do so. In a real system the measurements can be windowed down to a $N/2 \times N/2$ square to directly obtain most modal coefficients. The very lowest frequency modes can be estimated separately. The main impact of using our simulation to measure the modes is that the response of the DM is automatically included. The DM influence function causes high spatial frequencies to be attentuated on the DM. This leads to a filter with higher gains at these frequencies to compensate.

These three filters all have different responses tailored to the specific model which was used. With the exception of invisible modes, each filter will exactly reconstruct for its given model. There is a small difference between Mod-Hud and Ideal. The Custom filter is significantly different, due to the response of the DM. Fig. 2 shows the sum of the magnitude-squared of the x- and y-parts for each filter, which is the PSD of noise propagated from the WFS to the DM in the square aperture case.

2.C. Dealing with arbitrary apertures

With the original FTR filter, a major problem that had to be solved was how to make it reconstruct accurately on a circular aperture. Poyneer developed a method called Extension to extend the slopes outside of the aperture to ensure accurate and low-noise reconstruction. The modal control issue presents a similar problem. The Fourier basis of sines and cosines forms an ONB. Once an aperture is imposed, the sines and cosines are no longer orthogonal to each other and there are more modes than degrees of freedom. Since modal control is based on the modes being independent, how will this work?

It turns out that the Fourier basis has a special property that for arbitrary apertures allows it to be used as if it were an ONB. It is a special type of Frame[24] in which analysis and synthesis can proceed exactly as if it was an ONB. In this subsection we present the basics of frame analysis and prove that for arbitrary apertures, the real Fourier basis forms a special frame that allows both analysis and synthesis.

A set of vectors $\{\mathbf{m}_k\} \subset \mathcal{H}$ for a set of indices $k \in P$ is a frame of the space $\mathcal{H}$ if for all signals $\mathbf{x} \in \mathcal{H}$,

$$\exists A, B > 0 \text{ s.t. } A||\mathbf{x}||^2 \leq \sum_{k \in P} | < \mathbf{x}, \mathbf{m}_k > |^2 \leq B||\mathbf{x}||^2, \tag{26}$$

where A, B are the frame bounds. A special type of frame is called the tight frame. This occurs when $A = B \geq 1$. Then the frame bound represents the level of redundancy in the frame. Analysis with a frame is easy; the modal coefficients for a frame are calculated exactly as with an ONB using the inner product. Synthesis with a frame depends on the actual modes in the frame. Let the matrix $\mathbf{A}$ be the matrix which converts from the space of the signal $\mathcal{H}$ to the space of the frame coefficients $L^2(P)$. This matrix has the mode vectors as its rows. This matrix defines the dual modes of the frame as

$$\tilde{\mathbf{m}}_k \equiv \left(\mathbf{A}^T\mathbf{A}\right)^{-1} \mathbf{m}_k. \tag{27}$$

Given these dual modes, any signal $\mathbf{x} \in \mathcal{H}$ can be analyzed and synthesized as

$$\mathbf{x} = \sum_{k \in P} < \mathbf{x}, \mathbf{m}_k > \tilde{\mathbf{m}}_k. \tag{28}$$

This allows a frame to be used for analysis and synthesis, even if the frame does not form an ONB. The dual modes take care of the linear dependences. Examination of Eqs. 26, 27

shows that if $\mathbf{A}^T\mathbf{A} = \mathbf{I}$ then the frame bounds both equal 1 and the dual modes are the same as the frame modes.

First let us look at the full DFT case. We use vector notation and assume some rational scheme to index the 2D signals into the vectors. The matrix $\mathbf{A}_f$ has as its rows each of the real Fourier modes $\mathbf{m}_k$. As with any ONB matrix, the transpose is the inverse, so we obtain $\mathbf{A}_f^T\mathbf{A}_f = \mathbf{I}$. As we would expect for an ONB, the frame is tight with bound 1 and the dual modes are the frame modes.

Now we examine the DFT basis with an aperture imposed. This has all of the same modes, but each mode has all the points outside the aperture removed. Its matrix is $\mathbf{A}_w$. In this case we no longer have orthogonality, and $\mathbf{A}_w\mathbf{A}_w^T \neq \mathbf{I}$. However, $\mathbf{A}_w^T\mathbf{A}_w = \mathbf{I}$ for this special case. The proof of this is straight-forward. We know that for the full square grid, the ONB gives us $\mathbf{A}_f^T\mathbf{A}_f = \mathbf{I}$ and $\mathbf{A}_f\mathbf{A}_f^T = \mathbf{I}$. The matrix $\mathbf{A}_w$ is the same as $\mathbf{A}_f$, except it has some of its columns removed, because these columns represent points outside of the aperture. The matrix $\mathbf{A}_w^T$ is the same as $\mathbf{A}_f^T$, except it has some of its rows removed. These means that the matrix product $\mathbf{A}_w^T\mathbf{A}_w$ is the same as $\mathbf{A}_f^T\mathbf{A}_f$, except that it is missing entries corresponding to those removed rows and columns.

Because the FTR frame is tight with bounds 1, we can analyze and synthesize in the same manner as with an ONB. For an arbitrary set of modes, we would use the inner product to directly calculate the modal coefficients. In the windowed Fourier case, we would use the inner product over only the points in the aperture. However, we can actually use the DFT to calculate the modal coefficients in the windowed case, presuming we window the signal of interest. As shown above (see Eqs. 6,7) we can use the DFT to calculate the modal coefficients for the full real Fourier basis. If we window the signal, making all points outside the aperture zero, the DFT will calculate the modal coefficients for the windowed frame. This DFT may be fast. Better yet, if we have a way to reconstruct the phase such that it is very close to zero outside the aperture, we have the modal coefficients available during the reconstruction step at no extra computational cost. This is a significant advantage over modal schemes which use a single VMM to calculate the actuator commands. In such an approach, the modal coefficients must be calculated with the inner product in addition to the costs of the reconstruction VMM. In the case of Altair, this nearly doubles the computational burden at each time step.[25] In the FTR case, we need do no extra computation to have the modal coefficients available during reconstruction.

We can do this with a new method of slope management. The previously preferred method of slope management extended the slopes outside of the aperture.[9] We have developed a new method which works with the any of the Mod-Huf, Ideal or Custom filters. This new method is called edge correction. The only slopes which are modified are the two x-slopes to the immediate left and right of the aperture in each row and the two y-slopes to the immediate

top and bottom of the aperture in each column. For a specific row, these two slopes are both set to minus one-half of the sum of all other slopes in that row. This ensures that the overall slope across an entire row is zero, and that the area outside the aperture is as flat as possible. Any central obscuration is ignored in the edge correction process. The method of edge correction provides accurate modal coefficient estimates (since the phase is flat outside the aperture) and accurate phase reconstruction. The noise propagation is also low. We have verified this in two approaches. First, we have examined the reconstruction and coefficient estimation properties for each Fourier mode on a given aperture. Both are highly accurate. Second, this method provided performance in end-to-end simulations comparable to the old Extension method.

3. Gain estimation and application

We have chosen the optimal gain estimation method of Dessenne as the way to incorporate information about phase aberration random process into the reconstructor. This method assumes that independent modes are controlled individually. It determines the optimal control law based on telemetry. In particular we have chosen the frequency-domain approach described by Véran for gain estimation instead of a more general time-domain approach as detailed by Dessenne. In this section we will briefly summarize the optimization criterion in the frequency domain, and discuss how it is applied to the FTR modes.

3.A. Optimality criterion

The AO control system is modeled as a discrete controller with only pure delays and the response of the control law. A block diagram of this model is given in Fig. 3. The open-loop response of the controller is given by $H(z)$. The input modal signal to be controlled is given as m. The error signal after closed-loop control is ϵ. The noise of the modal measurement is given by n, which is added after a pure delay. The measured signal s is the error plus noise.

In this model, the closed-loop transfer functions from the input mode and noise to the measurements s are given as

$$\frac{S_{cl}(z)}{M(z)} = \frac{z^{-1}}{1 + z^{-1}H(z)}, \tag{29}$$

and

$$\frac{S_{cl}(z)}{N(z)} = \frac{1}{1 + z^{-1}H(z)}. \tag{30}$$

If the system observes in open-loop, the measurements are

$$\frac{S_{ol}(z)}{M(z)} = z^{-1}, \tag{31}$$

and

$$\frac{S_{ol}(z)}{N(z)} = 1. \tag{32}$$

The transfer function between closed-loop measurements and open-loop measurements is

$$\frac{S_{ol}(z)}{S_{cl}(z)} = 1 + z^{-1}H(z). \tag{33}$$

We want to minimize the squared value of the residual error ϵ, which is just the integral over the PSD of that random process. From the block diagram we see that $S(z) = z^{-1}E(z)+N(z)$. Since the noise at any instant is independent with any past errors, a minimization of $< s^2 >$ will produce a minimization of $< \epsilon^2 >$.

Using the notation that the PSD of the random process s is $S(\omega)$, we therefore want to minimize the cost function

$$\mathcal{J} = \int S(\omega)\, d\omega. \tag{34}$$

The integration is over the entire frequency domain, but we omit the limits of integration for clarity of notation. If we knew exactly the input mode signal and noise PSDs, we would find the control law $H(z)$ which minimizes

$$\mathcal{J} = \int \left| \frac{1}{1 + \exp(-j\omega)H(\omega)} \right|^2 [M(\omega) + N(\omega)]\ d\omega. \tag{35}$$

We don't know $M(\omega)$ and $N(\omega)$, of course, but we can estimate them from closed-loop measurements if we know the control law $H_0(z)$ that was used when those measurements were taken. Using Eq. 33 we obtain

$$\hat{M}(\omega) + \hat{N}(\omega) = |1 + \exp(-j\omega)H_0(\omega)|^2\, \hat{S}(\omega), \tag{36}$$

where $\hat{S}(\omega)$ is our best estimate of the PSD of the closed-loop measurements. This means that our optimization problem is exactly

$$\operatorname{argmin} H(z) \left\{ \int \left| \frac{1}{1 + \exp(-j\omega)H(\omega)} \right|^2 |1 + \exp(-j\omega)H_0(\omega)|^2\, \hat{S}(\omega)\, d\omega \right\}. \tag{37}$$

So our optimization methodology is as follows. We take a series of closed-loop measurements with a fixed, known control law $H_0(z)$. Then we estimate $\hat{S}(\omega)$ from these measurements. In general we will used an unbiased, averaged periodogram technique[22] to obtain $\hat{S}(\omega)$ from a series of consecutive intervals of telemetry data. We invert the control process with the known $H_0(z)$ to obtain the estimated signal and noise PSDs. We then minimize over possible control laws $H(z)$ to find the new optimal control law.

In our case we have chosen a integrator-gain control law

$$H(z) = \frac{g}{1 - cz^{-1}}, \tag{38}$$

where the integrator weight c is fixed at something very close to 1 and the gain g is adjustable. This makes the minimization problem of above over a single variable. In this case it is easy to determine the domain over which is control law is stable and develop a closed-form equation for the minimization.

3.B. Using this for FTR modes

For FTR we will apply the gain at the filtering stage during reconstruction. For speed of implementation, we will incorporate this gain into the reconstruction filter, which means that the sine and cosine modes at a specific frequency will have the same gain. This is reasonable, given the fact the the two modes are highly correlated in time and should have similar signal and noise profiles.

We treat both sine and cosine modes with the same weight and solve the following single-variable minimization, where the sine coefficient PSD is $\hat{S}_S(\omega)$ and the cosine coefficient PSD is $\hat{S}_C(\omega)$,

$$\operatorname{argmin} H(z) \left\{ \int \left| \frac{1}{1+\exp(-j\omega)H(\omega)} \right|^2 |1+\exp(-j\omega)H_0(\omega)|^2 \left[\hat{S}_S(\omega) + \hat{S}_C(\omega) \right] d\omega \right\}. \quad (39)$$

Since we know the form of the control law in advance we can reduce this to a simple root finding problem, using the derivative of the cost function. Since we bound our search domain by the stability condition, there has to be a minimum when we are estimating a single parameter. Note that we do this optimization for $N^2/2+2$ filter coefficients. The remaining coefficients of the filter are constructed using the Hermitian symmetry property.

4. Closed-loop simulation results

To test the effectiveness of OFC, we incorporated it into our end-to-end simulation of an AO system. This code is a Fourier optics-based simulation of a high-order AO system for an 8-m telescope. The WFS module determines CCD pixel values using integration over the appropriate regions of the subaperture PSF. This is configured for operation as a quadcell with pixel size $2\lambda/d$, where d is the subaperture diameter. The total number of received photons in a subaperture is based on a combination of standard photometric calibration to Vega and an assumed transmission of the telescope and AO system and CCD efficiency of 0.468. We assumed 8 electrons read noise on the CCD. For a given star magnitude I, the SNR of the WFS measurement can be calculated using standard methods.[26]

The DM is modeled using influence functions and linear superposition. The specific influence function used is that of the Altair DM.[27] The WFS leg has a spatial filter[28] to prevent spatial aliasing, which would introduce error into the modal coefficient estimates. The phase input is a very long phase screen which is translated (without looping) at a nominal speed of 10 m/s across the pupil in a diagonal direction. The phase screen is generated using the spectral factor method[29] and is calibrated so that the r_0 value agrees with the known structure function. For these simulations we have used r_0 of 18 cm (at 500 nm) and a system rate of 2.5 kHz. For the gain optimization we used $t = 128$ samples per periodogram and $a = 8$ individual periodgrams to determine the average, which gave very good performance. Both

tip and tilt and higher order corrections were done with a single-step delay integrator-gain controller, as specified above. We have two main modes of simulation. In the first, several thousand time steps are used to follow the gain optimization through many iterations on a single phase screen. This is used to study convergence and stability of the optimal gains through time. In the second mode, specific filters (e.g. the steady state optimized gains for a specific configuration) are run until the system reaches steady state on each of large set of random phase screens, which enables us to estimate performance in the general case. In these simulations we can examine different portions of the error by removing WFS noise, for example. For all constant-gain results in this section, the gain was set to 0.6 for all modes.

In our results below we use two metrics. For the high-contrast imaging case we are most concerned with the residual wave-front mean square error (MSE) in our controllable spatial-frequency domain out to maximum frequency $1/(2d)$. We calculate the MSE in the band at every time step and take the averaged, modified periodogram PSD estimate over intervals.[28] Both are calculated at the science wavelength of 1.6 microns. Of more interest to astronomical end-users is the metric of Strehl ratio, which we calculate with the standard Marechal approximation.

4.A. Performance improvement with optimal gains

Gain optimization involves a trade-off between errors due to signal and errors due to noise. At high SNRs, higher gains produce less temporal bandwidth error and more measurement error (from WFS noise propagated to the DM through the wave-front reconstruction process and the control law) than in the nominal constant gain case. At low SNRs, the gains are lowered, to produce more bandwidth error, but significantly less measurement error. Fig. 4 shows an example of this for the $N = 48$ case, where there are 44 subapertures across the pupil, resulting in d of 18.2 cm. As described above, we can examine in simulation the individual error components for specific filters. All results given are the median from a set of random phase screen inputs. For the nominal constant gain filter, the bandwidth error does not change with SNR. The measurement noise follows the expected inverse-square law with the SNR. When gain optimization is used, this profile is changed. Using the steady-state optimal filters for several cases, we determined that both bandwidth and measurement errors vary with the SNR, but produce a lower total error than in the constant gain case.

This results in substantial performance improvements at low SNRs. As a particular example of this, consider the results from the $I = 8$ case for three different subapertures sizes: d of 28.6 cm ($N = 32$), d of 18.2 cm ($N = 48$), and d of 13.3 cm ($N = 64$). For $N = 32$ the WFS SNR is 4.89. With constant gain, the MSE is 0.054 and the Strehl is 0.86. Use of the steady-state optimal filter results in a MSE of 0.025 and a Strehl of 0.89. Gain optimization halves the controllable MSE and increases the already high Strehl by 3%. If we reduce the

subaperture size to the $N = 48$ configuration, the WFS SNR is 2.16. The constant gain produces a MSE of .224 and a Strehl of 0.75. With the optimized gains, we obtain a MSE of 0.074 and a Strehl of 0.87. In this $N = 48$ case, use of gain optimization decreases the controllable residual MSE by a factor of three, and increases the Strehl by 12%, to the same level as was obtained in the $N = 32$ system with no optimization. The reasons why the constant-gain $N = 48$ case has lower Strehl than the constant-gain $N = 32$ case are that the WFS SNR is lower and that the noise propagation is higher due to increased system size. However, decreasing d by increasing N results in a larger region in the PSF which is controllable by the AO system, which is desirable in high contrast imaging.

If we further decrease the subaperture size, we see the same type of improvement. For $N = 64$ the WFS SNR is 1.19. Without gain optimization the MSE is 0.612 and the Strehl is 0.50. In the gain optimized case, the MSE is 0.210 and the Strehl is 0.75. Gain optimization decreases the controllable residual MSE by a nearly a factor of three, and the Strehl is increased by 25%, such that performance is comparable with the constant-gain $N = 48$ case.

4.B. Isolation of modes

Though our Fourier modes are not orthogonal on a circular aperture, we can control them as if they were. As a demonstration of this, we took the averaged optimal gains over several iterations at *at* steps from a frozen flow simulation for the $d = 28.6, I = 8$ case. We then constructed three filters - (1) uniform gain of 0.6, (2) optimal gains, (3) optimal gains at $4 \leq k \leq 12$ (and the pairs by Hermitian symmetry) with 0.6 gain elsewhere. We then ran the simulation for 1024 timesteps on the same phase screen and compared the PSDs of the residual phases for the three cases. As shown in Fig. 5, the results of using the filter in case (3) are nearly identical to a combination of the case (2) results for that subset of modes and the case (1) results elsewhere. This is a clear demonstration of the relationship between spatial frequency content and PSF location.[30,31] Though in the general case we will optimize all modes, this implies that if we are computationally limited to optimizing only a fraction of all modes, we can choose the locations in the PSF where the improvement will appear.

4.C. Gain estimation compensates for unknown DM

Three different reconstructors were derived above based on different models of the AO system. The Mod-Hud filter was based on a discrete model and has no information about the WFS or DM behavior. The gain optimization process should be able to compensate for this lack of knowledge. To test this, we ran both the Mod-Hud filter and the Custom filter based on the simulation for 8 iterations of gain optimization in the $d = 28.6, I = 8$ case. The gains were averaged over the last 6 iterations to produce an average optimal gain profile. The ratio of these gains should exactly be the compensation difference between the two filters. This

was confirmed with comparison to the known gain of the Mod-Hud filter, as determined empirically using the end-to-end simulation. Fig. 6 shows these two results, which have only minor differences at the lowest spatial frequencies. This shows that the gain optimization compensates for the DM response. Before optimization the Custom filter provides lower residual error, but after optimization the filters perform equivalently.

4.D. Convergence of gains and stability

It is important that the gain optimization method produce a stable control law which improves performance from the non-optimized case. The stability of the control law is ensured by a setting a limitation on the possible gains, based on a Z-transform analysis of the control law. For a wide range of SNRs from as low as 1 to as high as 58, the gains converged with in a few iterations and remained stable, with slight fluctuations. For very high SNRs, the levels of residual error was very low and consequently fluctuated as the phase aberration changed structure. At medium to low SNRs the level of performance was quite stable, as the measurement noise contribution was significant and uniform.

In our simulation we used a relatively short interval for estimating the PSD. Finding the best length of samples t and number of averages a will depend strongly on atmospheric characteristics, as the sampling should be done on long enough intervals to estimate low-frequency changes in the atmosphere accurately, but fast enough to keep up with any non-stationary changes in the atmosphere, such as significant changes in r_0.

5. Computational Costs

OFC is computationally efficient enough to be implemented with current technology. The detailed computational costs of FTR have already been analyzed,[9] so we will summarize for a specific case here, using the unit of Floating Point Operations (FLOPs) to measure the computational burdens.

At each step in time, the complete FTR process must be completed. Its cost is dominated by two real to complex FFTs on the slope inputs and by one complex to real FFT to obtain the estimated phase. This takes $15N^2 \lg N$ FLOPs, where the computational grid is $N \times N$. The application of the filter using complex multiplication takes 20 FLOPs per grid point, or $20N^2$ FLOPs total. Per timestep, FTR takes approximately $15N^2 \lg N + 20N^2$ FLOPs.

The gain estimation is dominated by the process of generating the periodograms. For a single modal coefficient series of t samples, we first apply a windowing function to provide better estimation. Application of this fixed window takes t FLOPs. We then need to calculate the t-element one-dimensional FFT, at a cost of $2.5t \lg t$ FLOPs. The magnitude-squared of the FFT is then taken, at a cost of $3t$ FLOPs, per periodogram calculation. To conserve memory, the periodogram is added to any previous in the buffer, requiring t FLOPs. This

is a total of of $2.5t \lg t + 5t$ per mode. This is must be done before the next full frame of t samples has been stored, requiring $N^2(5 + 2.5 \lg t)$ FLOPs per timestep.

After the required number a of periodograms has been generated, the optimization process occurs. Since the periodograms have already been averaged, the first step is the conversion to open-loop PSDs. For the $N^2/2 - 2$ frequencies with both cosine and sine modes, the two PSDs are summed at a cost of t FLOPs for each pair. Then these $N^2/2 - 2$ PSDs plus the 4 cosine only PSDs are all converted to open-loop PSD estimates. It will be easy to store the values of the inverse control response $|1 + z^{-1}H_0(z)|^2$ for a large number of gains in memory. Therefore this application takes just t FLOPs for each PSD. This is step takes $(N^2/2 + 2)t$ FLOPs, for a running total of $N^2 t$ FLOPs. The final step is to minimize and find the optimal gain. Again, the values of the closed-loop response should be available in memory for a large number of possible gains. A naive way to implement this would be to do a random search for the minimum, evaluating the total error at k different gains and then searching for the minimum. This takes $(N^2/2+2)2kt$ FLOPs, which is expensive if k is large. A better option would be to do root-finding bisection on the derivative cost function, which would require a much smaller number of evaluations of the cost function. Again the derivative would be stored in memory, making application quick. This would have the same cost, but limit k to a small number, most likely under 10. Ideally this entire process, from averaging to finding the gain, will be completed before t timesteps have passed. This leads to a computational burden of $N^2(1 + k) + 4k$ FLOPs per timestep for this portion.

In summary, the computational costs and implementation of OFC are as follows. At every step FTR computes the phase from the slopes. This takes $15N^2 \lg N + 20N^2$ FLOPs per timestep. From the frequency-domain stage of the reconstruction, the N^2 modal coefficients are extracted into a buffer which stores these coefficients through time. Time is divided into consecutive frames of t samples in length. An optimization cycle proceeds as follows, beginning with the instantiation of a new set of gains. In the first t-frame, no computational occurs and the modal coefficients are stored in the buffer. Then for the next a frames of length t, the periodograms are calculated on the previous frame's coefficients and the periodograms are accumulated. This process requires $N^2(5 + 2.5 \lg t)$ FLOPs per timestep. After a periodograms have been collected, the optimizer switches to processing these periodograms into open-loop PSDs and finding the gains which minimize the error. This process occurs in the the $a+2$-th frame, using data from the first through a-th frames. This requires $N^2(1+k)+4k$ FLOPs per timestep, where k is a small number. The new gains are instantiated at the end, meaning every $t(a + 2)$ timesteps new optimal gains are used, with a lag of $2t$ steps between the last data and the new control.

For a $N = 64$ AO system at 2.5kHz, this process has a maximum computational burden of 1.43 GFLOPs/sec (10^9 FLOPs per second) when $t = 1024$, $a = 8$ and we assume $k = 10$.

As t is reduced the burden also decreases; if $t = 128$ the burden is 1.36 GFLOPs/sec. FTR takes 70% to 75% of the total computation. Given this computational cost, OFC can be implemented for a such a large, fast AO system on currently available off-the-shelf hardware.

6. Conclusions

We have developed the method of optimal modal Fourier transform wave-front control as an efficient and adaptive AO control scheme. End-to-end simulations have shown that optimization of modal gains in closed loop can significantly improve AO performance, reducing controllable residual MSE by as much as three times and increasing Strehl ratios by as much as 25% for our scenario of interest. OFC has a few specific advantages over previous matrix-based implementations of optimal modal control. Because the modal coefficients are available for free during the FTR process, no extra computation is required. Due to the relationship between the Fourier modes and the spatial distribution of light in the PSF, OFC allows the control of specific locations in the PSF. We are currently investigating higher-order control laws and complex gains that may allow a further increase in performance. We plan to implement and validate OFC in the ExAO testbed at UCSC's Lab for Adaptive Optics using a 32×32 MEMS device.

Acknowledgements

We thank Bruce Macintosh for his thoughtful questions and useful input on system-level issues. This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48. The document number is UCRL-JRNL-207495. This work has been supported by the National Science Foundation Science and Technology Center for Adaptive Optics, managed by the University of California at Santa Cruz under cooperative agreement No. AST-9876783.

References

1. M. C. Roggemann and B. Welsh, *Imaging Through Turbulence* (CRC Press, New York, 1996).
2. B. L. Ellerbroek, "Efficient computation of minimum-variance wave-front reconstructors with sparse matrix techniques," J. Opt. Soc. Am. A **19**, 1803–1816 (2002).
3. D. G. MacMartin, "Local, hierarchic and iterative reconstructors for adaptive optics," J. Opt. Soc. Am. A **20**, 1084–1093 (2003).
4. F. Shi, D. G. MacMartin, M. Troy, G. L. Brack, R. S. Burruss, and R. G. Dekany, "Sparse-matrix wave-front reconstruction: simulations and experiments," in *Adaptive*

Optical System Technologies II, P. L. Wizinowich and D. Bonaccini, eds., Proc. SPIE **4839**, pp. 981–988 (2002).

5. L. Gilles, C. R. Vogel, and B. L. Ellerbroek, "Multigrid preconditioned conjugate-gradient method for large-scale wave-front reconstruction," J. Opt. Soc. Am. A **19**, 1817–1822 (2002).
6. L. Gilles, "Order-N sparse minimum-variance open-loop reconstructor for extreme adaptive optics," Opt. Lett. **28**, 1927–1929 (2003).
7. K. Freischlad and C. L. Koliopoulos, "Wavefront reconstruction from noisy slope or difference data using the discrete Fourier transform," in *Adaptive Optics*, J. E. Ludman, ed., Proc. SPIE **551**, pp. 74–80 (1985).
8. K. Freischlad and C. L. Koliopoulos, "Modal estimation of a wave front from difference measurements using the discrete Fourier transform," J. Opt. Soc. Am. A **3**, 1852–1861 (1986).
9. L. A. Poyneer, D. T. Gavel, and J. M. Brase, "Fast wave-front reconstruction in large adaptive optics systems with use of the Fourier transform," J. Opt. Soc. Am. A **19**, 2100–2111 (2002).
10. L. A. Poyneer, "Advanced techniques for Fourier transform wave-front reconstruction," in *Adaptive Optical System Technologies II*, P. L. Wizinowich and D. Bonaccini, eds., Proc. SPIE **4839**, pp. 1023–1033 (2002).
11. L. A. Poyneer, M. Troy, B. Macintosh, and D. Gavel, "Experimental validation of Fourier transform wave-front reconstruction at the Palomar Observatory," Opt. Lett. **28**, 798–800 (2003).
12. E. P. Wallner, "Optimal wave-front correction using slope measurements," J. Opt. Soc. Am. A **73**, 1771–1776 (1983).
13. D. T. Gavel and D. Wiberg, "Towards Strehl-optimizing adaptive optics controllers," in *Adaptive Optical System Technologies II*, P. L. Wizinowich and D. Bonaccini, eds., Proc. SPIE **4839**, pp. 890–901 (2002).
14. B. Ellerbroek, C. V. Loan, N. Pitsianis, and R. Plemmons, "Optimizing closed-loop adaptive-optics performance with use of multiple control bandwidths," J. Opt. Soc. Am. A **11**, 2871–2886 (1994).
15. E. Gendron and P. Léna, "Astronomical adaptive optics I. Modal control optimization," Astron. Astrophys. **291**, 337–347 (1994).
16. E. Gendron and P. Léna, "Astronomical adaptive optics II. Experimental results of an optimized modal control," Astron. Astrophys. Suppl. Ser. **111**, 153–167 (1995).
17. F. Rigaut, D. Salmon, R. Arsenault, J. Thomas, O. Lai, D. Rouan, J.-P. Véran, P. Gigan, D. Crampton, J. F. er, J. Stilburn, C. Boyer, and P. Jagourel, "Performance of the Canada-France-Hawaii Telescope adaptive optics bonnette," Publ. Astr. Soc. Pac. **110**,

152–164 (1998).

18. J.-P. Véran, "Altair's Optimiser," Tech. rep., Herzberg Institute of Astrophysics, Victoria, BC, Canada (1998).
19. G. Rousset, F. Lacombe, P. Puget, N. N. Hubin, E. Gendron, T. Fusco, R. Arsenault, J. Charton, P. Feautrier, P. Gigan, P. Y. Kern, A. Lagrange, P. Madec, D. Mouillet, D. Rabaud, P. Rabou, E. Stadler, and G. Zins, "NAOS, the first AO system of the VLT: on-sky performance," in *Adaptive Optical System Technologies II*, P. L. Wizinowich and D. Bonaccini, eds., Proc. SPIE **4839**, pp. 140–149 (2003).
20. C. Dessenne, P.-Y. Madec, and G. Rousset, "Optimization of a predictive controller for closed-loop adaptive optics," Appl. Opt. **37**, 4623–4633 (1998).
21. W. L. Briggs and V. E. Henson, *The DFT: An Owner's Manual for the Discrete Fourier Transform* (SIAM, Philadelphia, 1995).
22. A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-time Signal Processing* (Prentice Hall, New Jersey, 1999).
23. F. J. Rigaut, J.-P. Véran, and O. Lai, "Analytical model for Shack-Hartmann-based adaptive optics systems," in *Adaptive Optical System Technologies*, D. Bonaccini and R. K. Tyson, eds., Proc. SPIE **3353**, pp. 1038–1048 (Control, 1998).
24. S. Mallat, *A Wavelet Tour of Signal Processing* (Academic Press, New York, 1999).
25. M. J. Smith and J.-P. Véran, "Implementation of the Altair optimization processes," in *Adaptive Optics Systems and Technology II*, P. L. Wizinowich and D. Bonaccini, eds., Proc. SPIE **4839**, pp. 964–971 (2002).
26. S. B. Howell, *Handbook of CCD Astronomy* (Cambridge University Press, Cambridge, 2000).
27. L. Jolissaint, "Personal communication," (Herzberg Institute of Astrophysics, Victoria, BC, Canada, 2004).
28. L. A. Poyneer and B. Macintosh, "Spatially filtered wave-front sensor for high-order adaptive optics," J. Opt. Soc. Am. A **21**, 810–819 (2004).
29. 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 **1237**, pp. 372–383 (1994).
30. A. Sivaramakrishnan, J. P. Lloyd, P. Hodge, and B. A. Macintosh, "Speckle decorrelation and dynamic range in speckle noise-limited imaging," Ap. J. **581**, L59–62 (2002).
31. M. D. Perrin, A. Sivaramakrishnan, R. B. Makidon, B. R. Oppenheimer, and J. R. Graham, "The structure of high Strehl ratio point-spread functions," Ap. J. **596**, 702–712 (2003).

List of Figure Captions

Fig. 1. Frequency grid showing numbering and location of cosine and sine modes for $N = 8$. The four dark squares at $[0, 0], [0, N/2].[N/2, 0]$ and $[N/2, N/2]$ are only cosines. The light grey squares have both a sine and a cosine mode. Modes in white are matched by Hermitian symmetry with a numbered mode.

Fig. 2. Square aperture noise propagation PSDs for each of the three filters. The Mod-Hud and Ideal are very similar, while the Custom filter has higher gains at high spatial frequencies to compensate for the DM response.

Fig. 3. Block diagram of control system model. The input mode coefficient m is compensated via feedback on the error signal ϵ, which is measured in the presence of noise n. The WFS is approximated as a pure delay and the response $H(z)$ is the control law we wish to optimize.

Fig. 4. Measurement and temporal bandwidth MSEs as a function of WFS SNR for the $N = 48$ case. Constant gains lead to uniform bandwidth error and inverse-square measurement error. Gain optimization allows for balancing out these two components to lower overall MSE.

Fig. 5. Fourier modes (and hence locations in the PSF) can be controlled independently even on a circular aperture. Slice along x-frequency axis of long-exposure PSDs of the residual phase error from the end-to-end simulation, with effective k value of gain filter used as index. "Constant gain" is for uniform gain of 0.6 on all modes. "Optimal gain" is for optimal gains on all modes. "Mixed gain" has optimal gains for $4 \leq k \leq 12$ and constant 0.6 gain elsewhere. The Mixed curve follows the Constant PSD up to $k = 3$, where it transitions to the Optimal PSD, and then again at $k = 13$ is begins to increase. This demonstrates that specific modes can be controlled effectively despite their not being orthogonal.

Fig. 6. Gain optimization compensates for lack of knowledge about the DM in the reconstruction filter. "Empirical" is determined by putting modes through the simulation. "Ratio of optimal gains" is the ratio of the steady-state optimal gains for Mod-Hud and Custom filters on the same frozen-flow input.

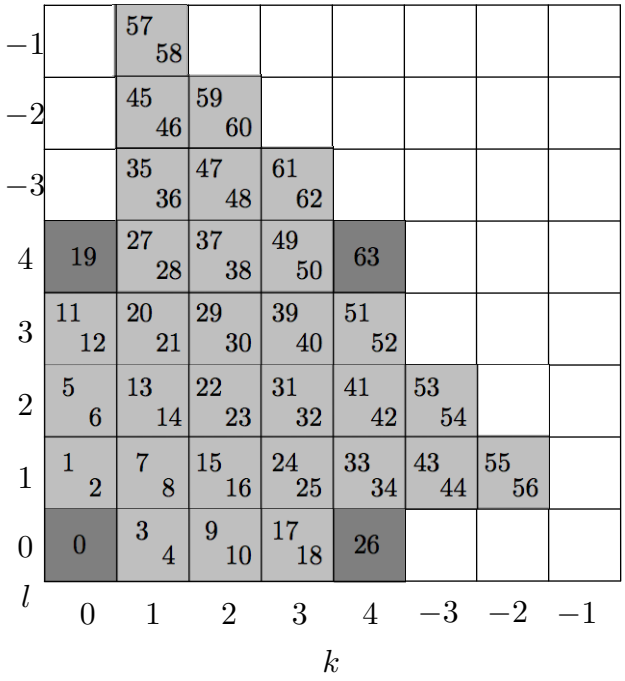
−1
−2
−3
4
3
2
1
0
l
0 1 2 3 4 −3 −2 −1
k
57 58
45 46
59 60
35 36
47 48
61 62
19
27 28
37 38
49 50
63
11 12
20 21
29 30
39 40
51 52
5 6
13 14
22 23
31 32
41 42
53 54
1 2
7 8
15 16
24 25
33 34
43 44
55 56
0
3 4
9 10
17 18
26

Fig. 1.

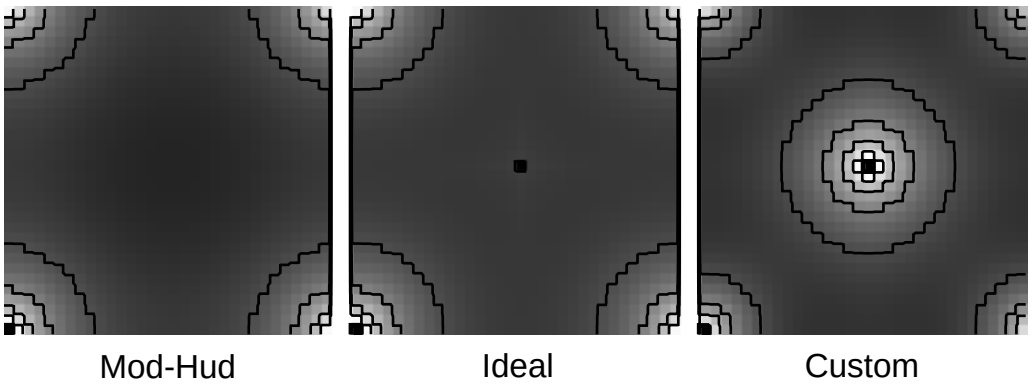

Fig. 2.

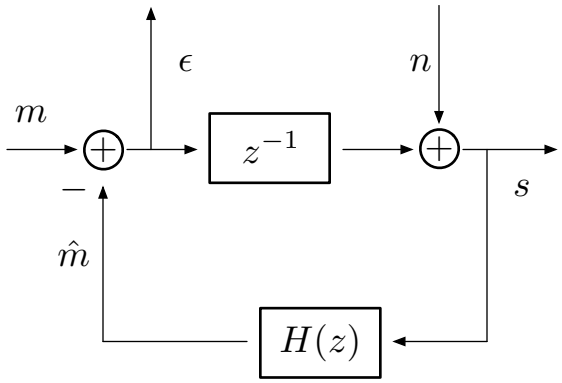

Fig. 3.

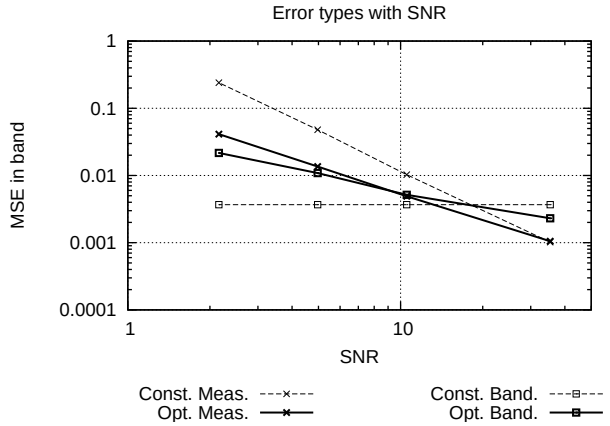
Error types with SNR
MSE in band
1
0.1
0.01
0.001
0.0001
1
10
SNR
Const. Meas.
Opt. Meas.
Const. Band.
Opt. Band.

Fig. 4.

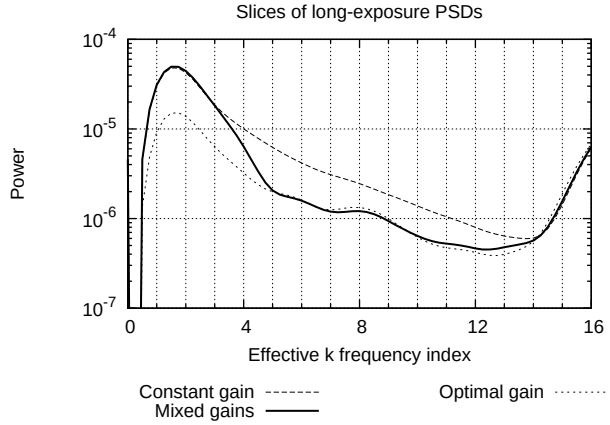
Slices of long-exposure PSDs
Power
10-4
10-5
10-6
10-7
0
4
8
12
16
Effective k frequency index
Constant gain
Mixed gains
Optimal gain

Fig. 5.

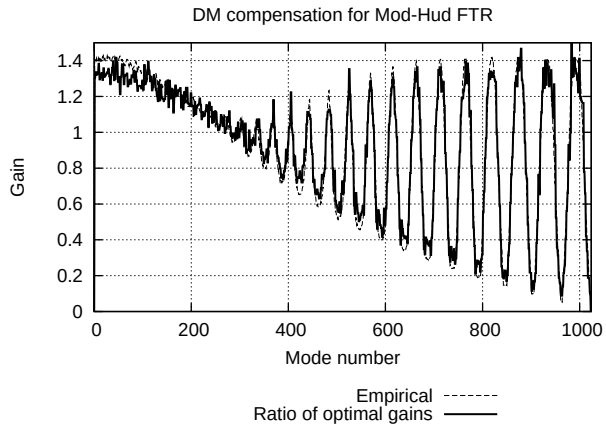
DM compensation for Mod-Hud FTR
Gain
1.4
1.2
1
0.8
0.6
0.4
0.2
0
0
200
400
600
800
1000
Mode number
Empirical
Ratio of optimal gains

Fig. 6.