Predictive wavefront control for adaptive optics with arbitrary control loop delays

Lisa Poyneer[1,*] and Jean-Pierre Véran[2]

[1]*Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore, California 94550, USA*
[2]*Herzberg Institute of Astrophysics, 5071 West Saanich Road, Victoria, British Columbia, Canada V9E2E7*
**Corresponding author: poyneer1@llnl.gov*

Received November 13, 2007; revised March 14, 2008; accepted April 7, 2008;
posted April 21, 2008 (Doc. ID 89607); published June 3, 2008

We present a modification of the closed-loop state space model for adaptive optics control that allows delays that are a noninteger multiple of the system frame rate. We derive the new forms of the predictive Fourier control Kalman filters for arbitrary delays and show that they are linear combinations of the whole-frame delay terms. This structure of the controller is independent of the delay. System stability margins and residual error variance both transition gracefully between integer-frame delays. © 2008 Optical Society of America

OCIS codes: 010.1080, 010.1285.

1. INTRODUCTION

Future adaptive optics (AO) systems have ambitious performance goals that will require advances beyond current wavefront control techniques. In particular, an advanced wavefront controller can help reduce the servo-lag (residual atmosphere) error, which scatters light close in near the point-spread-function (PSF) core. This area of research has been quite active recently, though work on applying established signal processing and control theory concepts to AO goes back to Paschall and Anderson's use of linear quadratic Gaussian (LQG) control in 1993 [1].

More recently, Kalman filtering concepts have been used by Gavel and Wiberg [2] and Le Roux *et al.* [3] to formulate optimal controllers that predict the atmosphere. Kulcsár *et al.* [4] have examined the theoretical basis for optimal control; Looze [5] has developed a different discrete-time model to describe the hybrid continuous-time/discrete-time AO system. Experimental work has been done by Petit *et al.* [6] on using Kalman filtering for vibration reduction and by Hinnen *et al.* [7] in using data-driven $\mathcal{H}_2$-optimal control.

In our own recent work we have proposed a computationally feasible and adaptive predictive controller. This method is termed predictive Fourier control (PFC) [8], and it builds upon the closed-loop AO control with a Kalman filtering framework that was developed by Le Roux [3]. In PFC the wavefront is reconstructed in the Fourier basis set [9]. Under the assumption of frozen-flow atmospheric turbulence, the Fourier modes are nearly uncorrelated both spatially and temporally (see Subsection 2.B of [8] for a detailed discussion of this result). This allows each complex-valued Fourier mode to be controlled independently. Closed-loop AO telemetry of the residual phase is recorded and analyzed using a temporal power spectral density (PSD) technique. This allows easy identification of atmospheric layers, which have a highly compact temporal PSD under frozen flow. State space model parameters are directly estimated from closed-loop telemetry and are then used to solve the algebraic Riccati equation (ARE), producing the steady-state Kalman filter, which predicts the atmosphere.

The state space model of Le Roux *et al.*, on which PFC is based, is a discrete-time approximation to the hybrid continuous-time/discrete-time AO control system. This standard model (as described in [10]) is shown in Fig. 1. In this model the continuous-time phase aberration $\phi(t)$ is corrected in closed loop, with discrete measurement noise $v[t]$. The wavefront sensor (WFS) integration and deformable mirror (DM) shaping are done on intervals of length T, producing a frame rate $f_{ao}=1/T$. The total computational latency of the AO system, from the end of the WFS CCD integration to the end of the write to the DM, is given by τ. If $\tau=T$, the AO system can be described with the discrete-time model of Le Roux *et al.* Le Roux's model assumes the signals are sampled at each interval. Kulcsár has shown that when τ is an integer multiple of T, and when the discrete-time phase is the average of the continuous atmospheric phase over that interval of length T, this state space model results in the optimal control algorithm [4]. The difference between sampling and averaging is slight in our case, as the AO frame rate is much higher than the temporal evolution of the continuously varying atmosphere.

If τ is a different integer multiple of T, the state space model of Le Roux can be easily modified to capture that delay. It may not always be the case, however, that τ is an integer multiple of T. For example, in the preliminary design phase of the Gemini Planet Imager (GPI) [11], the system's WFS integration time is $T=500$ μs and the delay is $\tau=737$ μs. Since by its very nature a predictive controller is intimately dependent on the total system delay, it is necessary to determine what the predictive controller will be when the delay τ is a noninteger multiple of T. Furthermore, the models that we use to do Monte Carlo simulations of AO systems are inherently discrete time and therefore cannot simulate fractional delays.

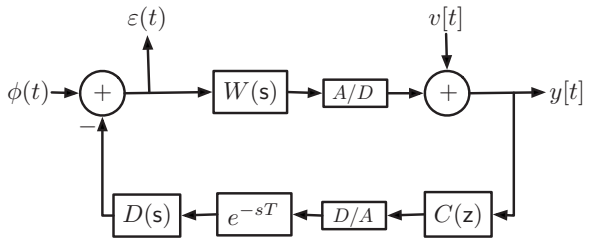

Fig. 1. Block diagram of hybrid continuous-time/discrete-time AO control loop for a single Fourier mode. The phase aberration $\phi(t)$ is corrected in closed loop in the presence of measurement noise $v[t]$. The WFS dynamics are represented by $W(\mathsf{s})$, the DM dynamics by $D(\mathsf{s})$, and the controller delay by $\exp(-\mathsf{s}\tau)$. The discrete-time control law is $C(\mathsf{z})$. A/D and D/A conversion surround the control computer.

In this paper we present a new modified state space model for PFC with arbitrary computational delays. We address the range $0 \leqslant \tau \leqslant 2T$; the treatment is easily extensible to longer delays, if necessary. In addition, we have removed redundant states that were present in the original model, resulting in computational savings. Though we consider only PFC in this paper, this treatment could be applied to the approach of Le Roux or other state space models.

Given this new model, a predictive controller is derived for arbitrary computational delays. Given this state space model, the PFC filter is the optimal temporal filter, since it is the Kalman filter. In this case, where the atmospheric input follows the model, use of a predictive controller provides significantly better performance than an integral controller. As shown here through simulations done in Simulink, variable-delay PFC has graceful behavior between integer time step delays. However, our model makes assumptions about both the control system and the atmosphere. Determination of whether or not this model provides the best answer in a real AO operating environment is left to future work.

2. STATE SPACE MODEL

A. Modeling the Asynchronous Deformable Mirror

The discrete-time state space model mentioned above is valid when the computational delay τ is an integer multiple of T. When the delay is a noninteger, however, the DM takes a new shape asynchronously with the rest of the system. We have chosen to address this by treating the discrete-time signals as if they are constant over an interval of length T and using linear combinations of them. This approach is logical given the DM's actual set-and-hold behavior; it is also used by Hinnen [7] and Looze [5] in their own AO models. We do not distinguish between the DM command that is output by the controller and the actual DM surface. This is because we neglect DM dynamics, as is discussed below.

The asynchronicity and how we model it is illustrated in Fig. 2. At the top is the discrete-time version, where the WFS measurement $y[t]$ is related to the atmospheric phase $\phi[t]$ and the DM commands $d[t]$. In the $\tau = T$ case of Le Roux (which PFC uses), the measurement is delayed by one time step as

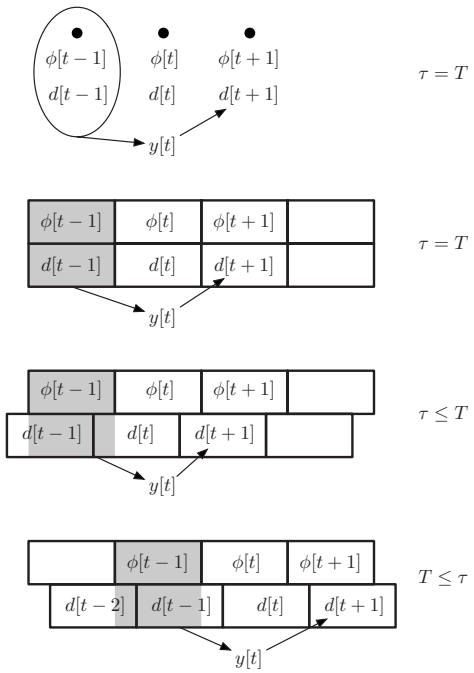

Fig. 2. Comparison of the discrete-time state space model for delay $\tau = T$ (row 1) to the new model, which assumes the signals are constant over an interval of the sampling period (row 2). When $\tau \leqslant T$, the DM signal shifts to the left and becomes asynchronous (row 3). When $T \leqslant \tau$, the DM signal shifts to the right (row 4).

$$y[t] = \phi[t-1] - d[t-1]. \tag{1}$$

Using the Kalman filter result, the new DM command $d[t+1]$ is the best estimate of the phase based on the WFS measurements up to time t as

$$d[t+1] = \hat{\phi}[t+1|t]. \tag{2}$$

To summarize, the measurement $y[t]$ sees the phase $\phi[t-1]$ and uses that information to calculate the DM command, which is applied at next time step $d[t+1]$. This correctly captures the two-frame delay when $\tau = T$ in Fig. 1.

In order to handle asynchronous DM writes, we move from a point-based model to an interval-based model. To do this we assume that all the signals are constant over intervals of length T, instead of existing at discrete points. The new model for $\tau = T$ is shown in the second row; the exact same equations hold. When the computational delay τ is less than T, the DM signals are applied earlier, which causes the DM signals to shift to the left by a fractional amount $\Delta = \tau/T - 1$, as is shown in the third row. The WFS measurement $y[t]$ still sees the phase $\phi[t-1]$, but it now observes two different shapes on the DM. It sees $d[t-1]$ for the first fraction of the interval and $d[t]$ for the second. Written using Δ, the measurement is now modeled as

$$y[t] = \phi[t-1] - \{-\Delta d[t] + (1+\Delta) d[t-1]\}. \tag{3}$$

This simply is a linear combination of the two DM signals, with the weighting dependent on the delay. The second consequence of the asynchronicity is that the new DM command $d[t+1]$ spans two phase instances. Following

the approach above, the new DM command is a linear combination of the phase estimates

$$d[t+1] = -\Delta\hat{\phi}[t|t] + (1+\Delta)\hat{\phi}[t+1|t]. \quad (4)$$

When the computational delay is longer than one step, the DM signal is shifted to the right, as is shown in the fourth row of the figure. Now the measurement is

$$y[t] = \phi[t-1] - \{(1-\Delta)d[t-1] + \Delta d[t-2]\}, \quad (5)$$

and the DM command is

$$d[t+1] = (1-\Delta)\hat{\phi}[t+1|t] + \Delta\hat{\phi}[t+2|t]. \quad (6)$$

In this model we treat all signals as being constant over the intervals of length T, with the DM write asynchronous with the WFS read. This constant behavior is an approximation: The atmosphere phase will in fact vary continuously. For the DM, this assumption ignores any temporal response (i.e., settling time) of the device. How much this approximation reduces the optimality of this model is left for future theoretical work.

These new measurement equations (3) and (5) will be incorporated into the state space model. The new DM command equations (4) and (6) will be used in solving the model for the predictive filter.

B. State Space Model for Fourier Control

Modal Fourier reconstruction analyzes the wavefront in the complex-valued Fourier modal basis set. For a given pupil size D and subaperture size d, there are D/d subapertures across the pupil. The phase is reconstructed on an $N\times N$ grid, where N is usually a few larger than D/d and is chosen for a computationally efficient discrete Fourier transform (DFT), e.g., FFTW [12]. For example, in GPI $D/d=44$ and $N=48$. The Fourier modes, when calculated with the DFT, are indexed by frequency variables k and l and take the values $-N/2, -(N/2-1), \ldots, -1, 0, 1, \ldots, (N/2-2)$, and $(N/2-1)$. This gives each Fourier mode the frequency components, in units of m^{-1}, $f_x = k/(Nd)$, and $f_y = l/(Nd)$.

Each Fourier mode is controlled independently. We can model the process of Shack–Hartmann wavefront sensing and Fourier transform wavefront reconstruction as if the Fourier modal coefficient of the phase were directly sensed. This is possible because we reconstruct and control the modes independently and the fact that the Fourier reconstruction filter can be customized to reconstruct the measured phase with unity gain [9].

The model of the input atmosphere is the same as in the original PFC derivation, where we assume frozen flow of layers of atmospheric phase aberration. When controlling Fourier modes, each turbulent layer becomes a first-order autoregressive process. A complex-valued AR(1) process has the basic form

$$a[t] = \alpha a[t-1] + w[t]. \quad (7)$$

The complex number α has a magnitude of just less than one. The phase of α sets how much the Fourier mode advances in a single time step. This is simply $2\pi T$ times the dot product of the velocity vector of the layer with the frequency vector of that Fourier mode: $-2\pi T(kv_x + lv_y)/(Nd)$.

We assume that the atmosphere is composed of a static layer, where α is a real number just less than one, and L layers of frozen flow. The state variables for these layer components are given by

$$\mathbf{a} = (a_0, a_1, \ldots, a_L), \quad (8)$$

and the autoregression parameters are stored in the matrix

$$\mathbf{R} = \mathrm{Diag}(\alpha_0, \alpha_1, \ldots, \alpha_L). \quad (9)$$

The state space model requires the power levels of each of the driving noises, which set the amount of phase power in each layer. These are given by the covariance matrix

$$\mathbf{P}_w = \mathrm{Diag}(\sigma_{a_0}^2, \sigma_{a_1}^2, \ldots, \sigma_{a_L}^2). \quad (10)$$

We have reduced the state order of the original PFC model to remove redundant states. Doing so produces a state vector that has only the atmospheric layers and the phase ϕ at the previous time step:

$$\mathbf{x}[t] = (\mathbf{a}[t], \phi[t-1])^T. \quad (11)$$

The evolution of this state vector is governed by the equation $\mathbf{x}[t+1] = \mathbf{A}\mathbf{x}[t] + \mathbf{B}\mathbf{w}[t]$. Note here that unlike the original PFC model, the DM signal is not included in the state or its update equations. The state transition matrix is

$$\mathbf{A} = \begin{pmatrix} \mathbf{R} & \mathbf{0}_{1\times(L+1)} \\ \mathbf{1}_{(L+1)\times 1} & 0 \end{pmatrix}. \quad (12)$$

The driving noises are incorporated with

$$\mathbf{B} = \begin{pmatrix} \mathbf{I}_{(L+1)\times(L+1)} \\ \mathbf{0}_{(L+1)\times 1} \end{pmatrix}. \quad (13)$$

This new reduced state model has a measurement equation that incorporates the DM commands as $y[t] = \mathbf{C}\mathbf{x}[t] + \mathbf{D}\mathbf{u}[t] + v[t]$. The measurement vector is always

$$\mathbf{C} = (\mathbf{0}_{1\times(L+1)}, 1), \quad (14)$$

which gives the $\phi[t-1]$ term in Eqs. (3) and (5). The DM shapes necessary in those equations are incorporated with $\mathbf{D}$ and $\mathbf{u}[t]$. The exact values depend on the system delay:

$$\mathbf{D} = \begin{cases} (-1) \\ (\Delta, -1-\Delta) \\ (-1) \\ (\Delta-1, -\Delta) \\ (-1) \end{cases},$$

$$\mathbf{u}[t] = \begin{cases} (d[t]) & \text{if } \Delta = -1, \\ (d[t], d[t-1])^T & \text{if } -1 \le \Delta \le 0, \\ (d[t-1]) & \text{if } \Delta = 0, \\ (d[t-1], d[t-2])^T & \text{if } 0 \le \Delta \le 1, \\ (d[t-2]) & \text{if } \Delta = 1. \end{cases} \quad (15)$$

The final necessary parameter is the variance of the measurement noise $v[t]$, which is $\mathbf{P}_v = (\sigma_v^2)$.

3. PREDICTIVE FILTER WITH ARBITRARY DELAY

Many AO systems use an integral controller, such as $C(z)=g/(1-0.99z^{-1})$. In the technique of modal gain optimization, which was developed by Gendron and Léna [13], the gain g is optimized. This method, done in closed loop with Fourier modes, forms the GPI baseline control algorithm of optimized-gain Fourier control (OFC) [9]. For PFC we will use AO telemetry to find the $C(z)$, where the form of the predictor is given by the Kalman filter.

Given this state space model, we solve for the temporal filter. Recall that the temporal filter $C(z)$ (as shown in Fig. 1) is the transfer function from the latest measurement $y[t]$ to the DM command $d[t+1]$, which is applied at the next time step (as shown in Fig. 2).

This process begins with the Kalman estimation equation

$$\hat{\mathbf{x}}[t|t]=(\mathbf{I}-\mathbf{K}_s\mathbf{C})\mathbf{A}\hat{\mathbf{x}}[t-1|t-1]+\mathbf{K}_s(y[t]-\mathbf{D}\mathbf{u}[t]). \tag{16}$$

The Kalman gain vector $\mathbf{K}_s$ is a function of the steady-state error covariance matrix $\mathbf{P}_s$,

$$\mathbf{K}_s=\mathbf{P}_s\mathbf{C}^H(\mathbf{C}\mathbf{P}_s\mathbf{C}^H+\mathbf{P}_v)^{-1}. \tag{17}$$

The controller $C(z)$ will be a function of the entries of $\mathbf{P}_s$. This matrix is found by numerically solving the ARE:

$$\mathbf{P}_s=\mathbf{A}\mathbf{P}_s\mathbf{A}^H+\mathbf{B}\mathbf{P}_w\mathbf{B}^H-\mathbf{A}\mathbf{P}_s\mathbf{C}^H(\mathbf{C}\mathbf{P}_s\mathbf{C}^H+\mathbf{P}_v)^{-1}\mathbf{C}\mathbf{P}_s\mathbf{A}^H. \tag{18}$$

The matrix $\mathbf{P}_s$ is indexed by column and then row indices, e.g., $p_{1,0}$ is the second element in the topmost row vector.

A. Details of Solution for Default Case ($\Delta=0$)

The details of the derivation for the $\Delta=0$ case are presented here. To solve the system of equations produced by Eq. (16), we use Z-transform notation with the complex number z, and uppercase letters indicating the Z-transform of time series, e.g., $A_0(z)$ for $\hat{a}_0[t|t]$. Note that state variable $\hat{a}_0[t|t]$ is the estimate of the state a_0 at time t given all the measurements up to and including time t. The control law filters the present measurement $y[t]$ to give the next DM command $d[t+1]$. This produces a transfer function:

$$C(z)=\frac{zD(z)}{Y(z)}. \tag{19}$$

To find this transfer function, we begin with the first $L+1$ scalar equations of the matrix equation (16). These are determined using algebra and the known values of $\mathbf{A}$, $\mathbf{C}$, $\mathbf{D}$, and $\mathbf{P}_s$. For $k=0,\ldots,L$,

$$A_k(z)=\alpha_k z^{-1}A_k(z)+Q^{-1}p_{L+1,k}\left[Y(z)+D(z)z^{-1}-z^{-1}\sum_{l=0}^{L}A_l(z)\right], \tag{20}$$

where the variable $Q=p_{L+1,L+1}+\sigma_v^2$. The $L+1$ equations for the atmospheric states are solved analytically with Gaussian elimination to find expressions for $A_0(z)$ through $A_L(z)$ as a function of $Y(z)$ and $D(z)$.

The final step requires expressing the $D(z)$ in terms of the atmospheric states. By Eq. (4), $d[t+1]=\hat{\phi}[t+1|t]$. The state transition matrix $\mathbf{A}$ shows that

$$\hat{\phi}[t+1|t]=\sum_{k=0}^{L}\alpha_k\hat{a}_k[t|t], \tag{21}$$

which gives us

$$zD(z)=\sum_{k=0}^{L}\alpha_k A_k(z). \tag{22}$$

Solving for the atmospheric states and substituting them into the above equation produces a single equation that relates $zD(z)$ to $Y(z)$. This is then solved to produce the controller

$$C(z)=\left(Q^{-1}\sum_{k=0}^{L}\frac{p_{L+1,k}\alpha_k}{1-\alpha_k z^{-1}}\right)\times\left(1+z^{-1}Q^{-1}\sum_{k=0}^{L}p_{L+1,k}\right)^{-1}. \tag{23}$$

This is the exact same controller as was originally derived for PFC [Eq. (23) of [8]], though now the covariance matrix $\mathbf{P}_s$ is smaller because redundant states have been removed.

B. Solution for $\Delta=-1$

For the (physically unrealistic) case of $\Delta=-1$, the first $L+1$ equations of Eq. (16) are

$$A_k(z)=\alpha_k z^{-1}A_k(z)+Q^{-1}p_{L+1,k}\left(Y(z)+D(z)-z^{-1}\sum_{l=0}^{L}A_l(z)\right). \tag{24}$$

The output of the controller $d[t+1]=\hat{\phi}[t|t]$ expressed in terms of atmospheric states is

$$zD(z)=\sum_{k=0}^{L}A_k(z). \tag{25}$$

The controller is then

$$C(z)=Q^{-1}\sum_{k=0}^{L}\frac{p_{L+1,k}\alpha_k}{1-\alpha_k z^{-1}}. \tag{26}$$

C. Solution for $\Delta=1$

When $\Delta=1$, the first $L+1$ equations of Eq. (16) are

$$A_k(z)=\alpha_k z^{-1}A_k(z)+Q^{-1}p_{L+1,k}\left(Y(z)+D(z)z^{-2}-z^{-1}\sum_{l=0}^{L}A_l(z)\right). \tag{27}$$

The output of the controller $d[t+1]=\hat{\phi}[t+2|t]$ expressed in terms of atmospheric states is

$$zD(z)=\sum_{k=0}^{L}\alpha_k^2 A_k(z). \tag{28}$$

The controller is then

$$C(z) = \left(Q^{-1} \sum_{k=0}^{L} \frac{p_{L+1,k}\alpha_k}{1-\alpha_k z^{-1}} \right) \times \left(1 + z^{-1} Q^{-1} \sum_{k=0}^{L} p_{L+1,k} + z^{-2} Q^{-1} \sum_{k=0}^{L} p_{L+1,k}\alpha_k \right)^{-1}. \quad (29)$$

D. Solution for $-1 \leq \Delta \leq 0$

When the delay τ is not an integer multiple of the WFS integration time T, the derivation is slightly more complex and requires an approximation. For the case $-1 \leq \Delta \leq 0$, the first $L+1$ equations of Eq. (16) are

$$A_k(z) = \alpha_k z^{-1} A_k(z) + Q^{-1} p_{L+1,k} \left(Y(z) + D(z)[-\Delta + (1+\Delta)z^{-1}] - z^{-1} \sum_{l=0}^{L} A_l(z) \right). \quad (30)$$

The output of the controller $d[t+1] = -\Delta\hat{\phi}[t|t] + (1+\Delta)\hat{\phi}[t+1|t]$, expressed in terms of atmospheric states, is

$$zD(z) = \sum_{k=0}^{L} (-\Delta + [1+\Delta]\alpha_k) A_k(z). \quad (31)$$

Solving for the atmospheric states and substituting in produces the controller

$$C(z) = \left(Q^{-1} \sum_{k=0}^{L} \frac{p_{L+1,k}}{1-\alpha_k z^{-1}} [-\Delta + (1+\Delta)\alpha_k] \right) \times \left(1 + z^{-1} Q^{-1} \sum_{k=0}^{L} p_{L+1,k} \times \left[\frac{1 - (-\Delta + [1+\Delta]\alpha_k)(-\Delta + [1+\Delta]z^{-1})}{1-\alpha_k z^{-1}} \right] \right)^{-1}. \quad (32)$$

Note that when $\Delta = -1$, this reduces exactly to Eq. (26), and when $\Delta = 0$, this reduces exactly to Eq. (23). Executing long division on the bracketed ratio inside the final summation produces

$$(1+\Delta)[(1-\Delta) + \Delta\alpha_k] + \frac{(1+\Delta)\Delta z^{-1}(1-\alpha_k)^2}{1-\alpha_k z^{-1}}. \quad (33)$$

If the angles of α_k's are small, $(1-\alpha_k)^2 \approx 0$ and the control law simplifies to

$$C(z) = \left(Q^{-1} \sum_{k=0}^{L} \frac{p_{L+1,k}}{1-\alpha_k z^{-1}} [-\Delta + (1+\Delta)\alpha_k] \right) \times \left(1 + z^{-1} Q^{-1} (1+\Delta) \sum_{k=0}^{L} p_{L+1,k}[(1-\Delta) + \Delta\alpha_k] \right)^{-1}. \quad (34)$$

This assumption that the angles of the α_i's are small is a reasonable one for PFC. Given a 20 m/s wind, the highest temporal frequency seen by PFC is around 80 Hz, or just 8% of Nyquist for GPI's 2 kHz AO. This makes the angle of $\alpha = 0.25$ rad, which qualifies as small. This assumption does not hold in all cases. Kalman filtering, as shown by Petit *et al.* [6], can correct for vibration at half the Nyquist frequency. In that case the angle of α is $\pi/2$ rad. In such a case the above approximation cannot be made.

E. Solution for $0 \leq \Delta \leq 1$

For the case $0 \leq \Delta \leq 1$, the first $L+1$ equations of Eq. (16) are

$$A_k(z) = \alpha_k z^{-1} A_k(z) + Q^{-1} p_{L+1,k} \left(Y(z) + D(z) z^{-1}[(1-\Delta) + \Delta z^{-1}] - z^{-1} \sum_{l=0}^{L} A_l(z) \right). \quad (35)$$

The output of the controller $d[t+1] = (1-\Delta)\hat{\phi}[t+1|t] + \Delta\hat{\phi}[t+2|t]$ expressed in terms of atmospheric states is

$$zD(z) = \sum_{k=0}^{L} ([1-\Delta] + \Delta\alpha_k)\alpha_k A_k(z). \quad (36)$$

Solving for the atmospheric states and substituting in produces the controller

$$C(z) = \left(Q^{-1} \sum_{k=0}^{L} \frac{p_{L+1,k}\alpha_k}{1-\alpha_k z^{-1}} [(1-\Delta) + \Delta\alpha_k] \right) \times \left(1 + z^{-1} Q^{-1} \sum_{k=0}^{L} p_{L+1,k} \times \left[\frac{1 - \alpha_k z^{-1}(1-\Delta+\Delta z^{-1})(1-\Delta+\Delta\alpha_k)}{1-\alpha_k z^{-1}} \right] \right)^{-1}. \quad (37)$$

Note that when $\Delta = 0$, this reduces exactly to Eq. (23), and when $\Delta = 1$, this reduces exactly to Eq. (29). Again, using the small-angle approximation, the fraction inside the final summation can be simplified to produce

$$C(z) = \left(Q^{-1} \sum_{k=0}^{L} \frac{p_{L+1,k}\alpha_k}{1-\alpha_k z^{-1}} [(1-\Delta) + \Delta\alpha_k] \right) \times \left(1 + z^{-1} Q^{-1} \sum_{k=0}^{L} p_{L+1,k} + z^{-2} Q^{-1} \Delta \times \sum_{k=0}^{L} p_{L+1,k}\alpha_k (2 - \Delta + [\Delta - 1]\alpha_k) \right)^{-1}. \quad (38)$$

F. General Form and Implementation

These two forms of the predictive controller, Eqs. (34) and (38), appear to be fairly complex algebraic expressions. However, they share the same simple underlying structure, which is illustrated in the block diagram of Fig. 3. Just as in the original PFC derivation, the predictive filter is made up of two parts: parallel first-order filters that predict each layer and a stabilizing high-pass filter.

First, the residual measurement $y[t]$ is fed through a first-order filter that we term a layer-compensating inte-

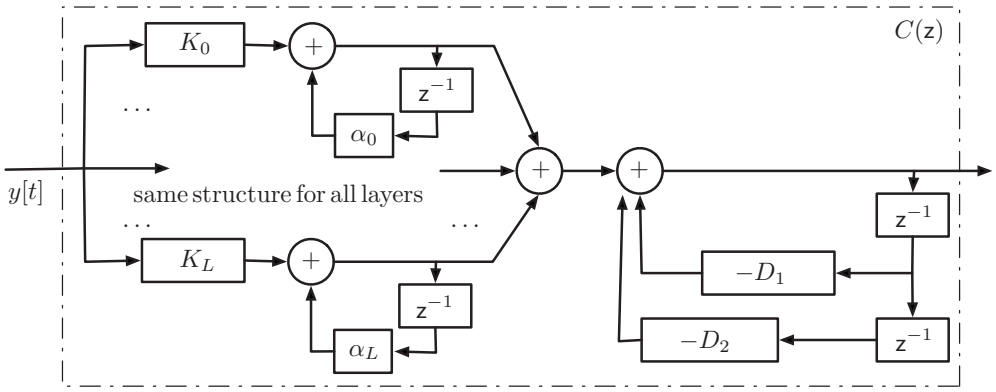

Fig. 3. Flow diagram illustrating the implementation of the predictive filter for $0 \leq \tau \leq 2T$. The coefficients used in the filter depend on Δ and are taken from Eqs. (39)–(41). Note that when $\tau \leq T$, $D_2 = 0$.

grator. The gain on the integral controller is given by K_k. From Eqs. (34) and (38) this gain is

$$K_k = \begin{cases} Q^{-1} p_{L+1,k}[-\Delta + (1+\Delta)\alpha_k] & \text{if } -1 \leq \Delta \leq 0, \\ Q^{-1} p_{L+1,k}\alpha_k[(1-\Delta) + \Delta\alpha_k] & \text{if } 0 \leq \Delta \leq 1. \end{cases} \tag{39}$$

This gain predicts the WFS measurement by the proper amount for that specific layer and provides a scaling based on the relative strengths of the layers.

The high-pass filter (which is a lead filter) is usually either first or second order. The two coefficients D_1 and D_2 are given by

$$D_1 = \begin{cases} Q^{-1}(1+\Delta)\sum_{k=0}^{L} p_{L+1,k}[(1-\Delta) + \Delta\alpha_k] & \text{if } -1 \leq \Delta \leq 0, \\ Q^{-1}\sum_{k=0}^{L} p_{L+1,k} & \text{if } 0 \leq \Delta \leq 1, \end{cases} \tag{40}$$

and

$$D_2 = \begin{cases} 0 & \text{if } -1 \leq \Delta \leq 0, \\ Q^{-1}\Delta\sum_{k=0}^{L} p_{L+1,k}\alpha_k[2 - \Delta + (\Delta - 1)\alpha_k] & \text{if } 0 \leq \Delta \leq 1. \end{cases} \tag{41}$$

This consistency in structure makes the predictor easy to implement in a system with variable computational delays. The structure is the same, while the exact filter coefficients change with delay. Each filter coefficient K_i, D_1, or D_2 can depend on both Δ and the entries of the steady-state error covariance matrix $\mathbf{P}_s$.

However, $\mathbf{P}_s$ does not depend on Δ. This is clear based on an examination of the ARE (18). The ARE depends only on the matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$, none of which depend on Δ. This means that the ARE can be solved independently from knowledge of the exact delay τ. Particular filter forms and the implications of these conclusions will be discussed further in the specific example of a single layer.

G. Exploring the Predictor for One Layer

Let us look at the controller for the case of one layer for five different possible values or ranges of τ. For this case we assume that the atmosphere is composed of a single layer of frozen flow and that no static errors exist. The complex-valued parameter α encodes the phase change with one time step. Multiplication by α is equivalent to predicting by one time step. In the case of a single layer, the coefficients of the predictive filter can be parameterized by a real-valued gain g, which ranges from 0 to 1,

$$g = \frac{p_{1,1}}{p_{1,1} + \sigma_v^2} = Q^{-1} p_{1,1}. \tag{42}$$

This gain g is a function of the signal-to-noise ratio (SNR).

Given g, we can simplify the equations for the filter coefficients. We focus on Eq. (39), which is the gain of the layer integrator,

$$K = \begin{cases} g\alpha & \text{if } \Delta = -1, \\ g\alpha(-\Delta + [1+\Delta]\alpha) & \text{if } -1 \leq \Delta \leq 0, \\ g\alpha^2 & \text{if } \Delta = 0, \\ g\alpha^2([1-\Delta] + \Delta\alpha) & \text{if } 0 \leq \Delta \leq 1, \\ g\alpha^3 & \text{if } \Delta = 1. \end{cases} \tag{43}$$

In the case of integer delays, the coefficient K is a gain (based on the SNR) and a prediction by the appropriate number of time steps. For example, in the $\tau = T$ ($\Delta = 0$) case there is a delay of two whole steps between the WFS integration interval and the interval where the new DM commands are applied. Multiplication by α^2 predicts by two steps. When the delay is a noninteger multiple of T, the coefficient K does a linear approximation with the nearest whole delays. This linear approximation is a good one, particularly when the angle of α is small, which we assume it is.

The layer-integrator coefficient K has two functions—it scales the residual input by a gain g based on the SNR, and then it predicts the measurement by the correct fractional number of steps between WFS and DM. The previous DM command is always predicted by one time step, as would be expected.

In our model for arbitrary delays (see Fig. 2), we assumed that the WFS measurements and DM commands involved a linear combination of the overlapping intervals. The end result is a predictor that deals with noninteger delays by doing the same linear interpolation of the control for whole-integer delays.

H. Implementation

In this development (as in the original PFC derivation) we have presented the predictive controller as a filter from the current measurement to the next DM commands. Representing it as $C(z)$ allows easy implementation and analysis of the temporal frequency responses of the AO system, as will be done in the following sections.

In the original PFC derivation, the state space model had several redundant states. Given this, implementing the Kalman filter as $C(z)$, as opposed to the matrix update equation (16), provided a computational advantage. The controller described by $C(z)$ does not keep track of these redundant states, which provides computational and memory savings. Now that the order of the state space model has been reduced, there is less of an advantage to implementing $C(z)$ directly.

A much more significant benefit to reducing the number of states is the reduction in computational cost for the solving of the ARE. PFC works in a "block adaptive" fashion [14], using telemetry to estimate model parameters. The ARE is then solved, and the predictive filter is updated. In the original proposal [8] the ARE was solved with the doubling algorithm. This algorithm was chosen for both ease of analysis and implementation in the complex-number case. However, the cost to solve the ARE was cubic in the number of states. For many layers, this made the cost of solving the ARE, even when amortized over several hundred time steps, the dominant expense. In the new reduced-state model, there are $L+2$ states, not $L+6$. This significantly reduces the cost; for $L=2$ the doubling algorithm now requires 8 times fewer operations. (See Subsection 7.C of [8] for a detailed discussion of the computational costs.) It is possible that this cost could be reduced even further through either the use of another algorithm, such as Arnold and Laub [15], or an analytic solution.

4. CONTROLLER MODELING AND ANALYSIS

In this study we measure the quality of controller performance using Simulink. This allows implementation of physically realistic AO elements and arbitrary delays. We also use an approximate Laplace transform model of the system to evaluate stability and optimize gains. First, we describe the Simulink model and AO system elements in detail. Then we discuss the Laplace model and how it is used in conjunction with Simulink.

A. Simulink Model

The hybrid continuous-time/discrete-time AO model is given in Fig. 1. This forms the basis for the Simulink model, which is shown in Fig. 4. Our Simulink model is similar to that in [5].

Fourier control deals with complex-valued signals. Because Simulink blocks vary in their ability to handle complex numbers, junctions are used to switch back and forth from a single complex-valued signal and two real-valued signals (real and imaginary parts) as appropriate. The input phase aberration $\phi(t)$ is given by `datainReal` and `datainImag`. After DM compensation, the residual phase error $\varepsilon(t)$ is given by `simout`. This residual error is measured by the WFS, which has its own model given at the bottom of Fig. 4. This implements the WFS integration over an interval of length T. The sampling of the WFS output is done with the $1/\tau s$ block. WFS noise $v[t]$ is added with `RandomNumber`. The lag $\exp(-s\tau)$ is moved up into the WFS block, with no loss of accuracy, to ensure proper continuous-time implementation.

This residual phase estimate is then sent to either an integral controller or the predictive Kalman filter. These are implemented with discrete-time blocks directly from their transfer functions. In particular, the predictive filter

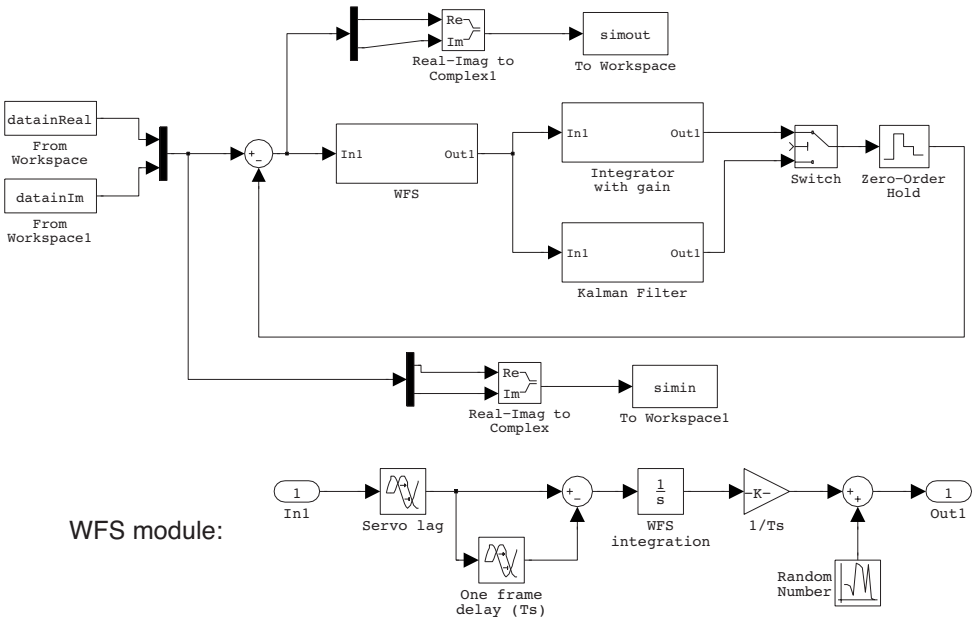

Fig. 4. Simulink model of the hybrid continuous-time/discrete-time AO control loop for a complex-valued Fourier modal coefficient. The WFS module implements the standard transfer function $W(s)$, while the DM module is a Simulink zero-order hold block. The integral controller and Kalman filters are implemented with discrete-time blocks of gains and delays. (The Kalman block is implemented exactly as in Fig. 3.) A switch allows use of either controller.

is implemented exactly as given in Fig. 3. The switch selects which control is used. The signal is then placed on the DM. A standard Simulink `Zero-order hold` block is used to model the DM behavior and convert from discrete to continuous. Any temporal dynamics (e.g., resonances) on the DM are ignored.

In order to capture the subsample behavior, the Simulink model operates with an internal sampling of $100f_{ao}$. The atmospheric input is generated at this sampling. This is done by first generating a random time series for each layer following the recursion in Eq. (7). These series are combined into the multilayer atmosphere and then interpolated with cubic splines from sampling at f_{ao} to $100f_{ao}$. This makes the atmospheric input effectively to be continuously varying (as the real atmosphere is), as opposed to the model assumption that it is static over WFS integration intervals.

The Simulink model saves signals at $10f_{ao}$. Oversampling from f_{ao} is necessary to prevent aliasing errors in signal analysis. The baseline AO run is 16384 AO steps, or roughly 8 s at our nominal sampling frequency of 2 kHz. These output signals are analyzed to produce estimates of the residual error variance, which is computed directly in the time domain as the variance of `simout`. This is equivalent to integrating in the frequency domain from $-5f_{ao}$ to $5f_{ao}$. Further increasing the sample rate of the saved signals, and hence the frequency-domain region, did not change the performance estimates. This is because there is negligible power at very high temporal frequencies, given our AO system and input phase aberration.

B. Laplace Transform Model

While the Simulink model provides our performance analysis, an analytic approach is also desirable for two reasons. First, an analytic model allows easy evaluation of system transfer functions and stability margins. Second, the determination of the optimized gain for the integral controller (i.e., OFC) is done in the temporal frequency domain. This requires the transfer functions for the system.

We use a Laplace transform model, based on the approach described by Madec [10]. This model of the hybrid continuous-time/discrete-time AO system is an approximate one. Because of sampling, the system is not linear time invariant (LTI) and hence does not strictly have transfer functions that describe it. However, this approximation is known to be good at frequencies much smaller than the Nyquist frequency, which, considering the typical distribution of the power in the turbulence disturbance, are the frequencies of interest, as long as the system has a sampling frequency greater than a few hundred hertz. We will discuss further below just how much agreement this model has to the Simulink results.

The three continuous elements in Fig. 1 are modeled with Laplace transforms. The WFS integrates over one interval from t to $t+T$, which in Laplace notation is $W(\mathsf{s})=[1-\exp(-\mathsf{s}T)]/(\mathsf{s}T)$. Note that this is exactly the implementation in the Simulink model. The actual WFS is not LTI since it really integrates and samples; this behavior is modeled by the A/D converter after $W(\mathsf{s})$. Likewise, the DM is not LTI; the discrete commands from the controller are placed and held on the DM. The D/A convertor deals with the first part; the set-and-hold behavior has the same transfer function as the WFS, $D(\mathsf{s})=[1-\exp(-\mathsf{s}T)]/(\mathsf{s}T)$. This is different from the Simulink model's DM block, which uses a zero-order hold. As above, we assume that the resonance of the DM is well above the frame rate of the system, which makes the temporal transfer function of the DM equal to 1.

The computational delay is simply $\exp(-\mathsf{s}\tau)$. To include both the continuous and discrete transfer functions, the A/D and D/A conversions are ignored and z must be mapped to s. For restricted domains of integration, the substitution $\mathsf{z}=\exp(\mathsf{s}T)$ is made. For extended domains (e.g., beyond $f_{ao}/2$), the bilinear transform should be used, as done in [5].

Using this method, transfer functions can easily be generated that approximate the temporal responses of the AO system. The error transfer function, which describes the system response from atmospheric input $\phi(t)$ to residual error $\varepsilon(t)$ is

$$H_e(f)=\frac{1}{1+W(\mathsf{s})C(\mathsf{z})\exp(-\mathsf{s}\tau)D(\mathsf{s})}. \tag{44}$$

The noise transfer function from WFS noise $v[t]$ to $\varepsilon(t)$ is

$$H_n(f)=\frac{-C(\mathsf{z})\exp(-\mathsf{s}\tau)D(\mathsf{s})}{1+W(\mathsf{s})C(\mathsf{z})\exp(-\mathsf{s}\tau)D(\mathsf{s})}. \tag{45}$$

These transfer functions can be used to find the margins of the system or to determine the optimal gain for the integral controller.

C. Stability

We will use the Laplace model for two important tasks: evaluating the stability of the controllers and optimizing the gain of the integral controller. These tasks will be successful only if the Laplace model agrees adequately with the actual behavior exhibited by the Simulink model, which is our best implementation of a real AO system and its temporal dynamics. In this section we discuss stability; gain optimization is discussed along with overall performance in the following section.

In order to compare the models, we generate approximate transfer functions from the Simulink model. This is done by sending temporal white noise through the system either from the atmospheric input or the WFS noise input. The PSDs of either input and the residual error are estimated with the averaged, modified periodogram technique [16] from Simulink's saved time series. Following basic principles of stochastic estimation, we estimate the magnitude squared of the effective transfer function by the ratio of these PSDs. For the error transfer function, the atmosphere is temporal white noise and the WFS noise is set to zero. The error transfer function is then estimated as

$$|\hat{H}_e(f)|^2=\frac{\hat{P}_\varepsilon(f)}{\hat{P}_\phi(f)}, \tag{46}$$

where the $\hat{P}(f)$'s are the estimated PSDs. Likewise, this

can be done for the noise transfer function by having zero atmospheric input and white noise on the WFS.

The error transfer functions for an $L=3$ predictive controller with $\tau=1.5T$, as determined by Simulink and the Laplace model, are shown in Fig. 5. Note that there are differences between them, particularly at higher temporal frequencies. This is to be expected, because the DM transfer function used in the Laplace method is only an approximation to the actual behavior of the zero-order hold block in Simulink. When integrated over the frequency domain from $f=-f_{ao}/2$ to $f=f_{ao}/2$, the actual discrepancy is quite small—the integral of the Laplace transfer function is within $<1\%$ of the Simulink result.

Though the AO control system is modeled in Simulink, the transport delay in the WFS poses significant problems when it comes to linearizing and discretizing the model for analysis of margins in Matlab. Given this, and the good agreement of the Laplace transfer functions with the Simulink model, we instead evaluate the stability margins directly by numerical analysis of the Laplace transfer functions.

Just as with the original PFC, the control law $C(z)$ itself is assured to be stable by the model structure of the Kalman filter. As long as the magnitudes of the α_i's are all less than one, the Kalman filter is stable. To analyze the gain and phase margins, we examine $W(s)C(z=\exp(sT))\exp(-s\tau)D(s)$ using standard techniques [17]. The baseline stability requirements are that the phase margin be at least 45 deg, and the gain margin be at least 2. For the controllers with which we deal here, if one of these requirements is met, the other is as well. The predictor, since it has complex-valued coefficients, has slightly asymmetric margins. We always choose the minimum margin. The margins are determined for a single Fourier mode. Because the modes are controlled independently, we assume that if all modes have satisfactory margins, the AO system will as well.

Figure 6 gives the phase margins for the case of a high SNR, where the controllers will be most aggressive in pushing the system margins. In fact, for most delays the optimized-gain controller is limited to the maximum gain, which produces a phase margin of 45 deg. The predictor, however, has larger margins. For the noninteger step delays, the phase margins follow a near-linear interpolation among $\tau=0$, $\tau=T$, and $\tau=2T$. This example illustrates one

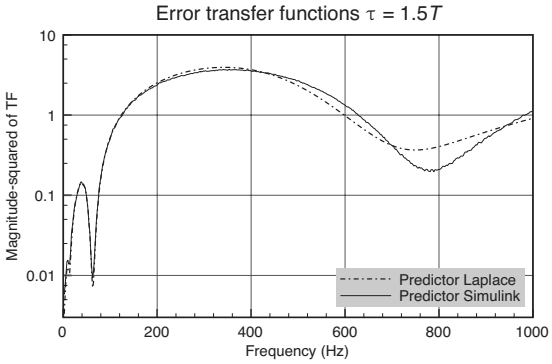

Fig. 5. Error transfer function for a predictive controller with $\tau=1.5T$, based on the Laplace transform model or as determined by running white noise through the Simulink model. There are some small discrepancies due to the modeling of the DM.

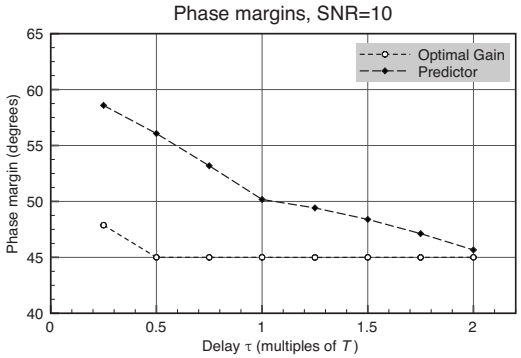

Fig. 6. Phase margins for a high-SNR case for the optimized-gain integrator and the predictor. For most delays τ, the optimal gain is limited to be the maximum gain set by stability. The predictor can be more aggressive and has larger margins for all τ. Note how the margin gracefully transitions between integer time step delays.

advantage of the predictor over the integral controller: At high SNRs the predictor does not run into the same limitation on gain set by stability, which allows it to provide a much better correction. This will be covered in more detail below. For lower SNRs the phase margins are larger for both controllers.

5. RESULTS

A. Methods and Cases Examined

For performance analysis we have picked a specific case that is reasonable given GPI parameters. In this case $L=3$. For $i=0$ to L, $|\alpha_i|=0.99$. The layer temporal frequencies are −40, 13, and 63 Hz. (As noted above, given a 20 m/s wind, a GPI Fourier mode will see layer frequencies up to 80 Hz.) The relative distribution of driving white-noise power is $\sigma^2_{a_0}=3$, $\sigma^2_{a_1}=2$, $\sigma^2_{a_2}=0.5$, and $\sigma^2_{a_3}=1$. Given this distribution, the total power in $\phi[t]$ is normalized to be 1. Then the noise power level is set by fixing the SNR, which is defined as the ratio of the standard deviations σ_ϕ/σ_v. The frame rate is $f_{ao}=2000$ Hz, which makes $T=500$ μs. The controller delay τ can be varied from 125 μs to 2000 μs, in increments of 125 μs, or $T/4$.

The simulations to study the performance of the predictor with variable delays was conducted in Simulink. For each configuration of system delay, 32 different random trials were run, each with a different input atmospheric time series and WFS noise. The residual error power for each trial was estimated by calculating the variance of the residual 8 s time series, which is sampled at $10f_{ao}$. The mean and standard deviation of the error power are calculated from this set of 32 different trials. Recall that the input atmosphere has a power equal to 1, and the WFS noise power is set by the SNR.

The predictive controller was determined by using the above model parameters and numerically solving the ARE (as described in detail in [8]). For this test the predictor was compared to the baseline GPI strategy of OFC [9].

B. Gain Optimization

The gain optimization technique used in OFC is temporal-PSD based. It requires knowledge of the transfer functions [see Eqs. (44) and (45)]. If those transfer functions are not accurate, the optimized gain will be the

wrong gain. First, we describe the gain optimization technique that was used here for the integral controller, and then we evaluate the quality of this gain optimization.

Gain optimization was done with the known PSDs of the wavefront, as opposed to based on data from the simulation. In the OFC proposal, the power on the measurements (given by $y[t]$ in Fig. 1) is minimized, and a Z-transform description of the system is used. For this new arbitrary control loop delay formulation, the Laplace transfer functions will be used. A second discrepancy from the proposal is that the assumption that allowed the measurements $y[t]$ to be minimized as a proxy to the residual error $\varepsilon(t)$ [18] appears to be no longer valid. This assumption was based on the fact that the new WFS noise seen by the measurements is uncorrelated with the residual error. This can also be viewed as the assumption that $\int P_y(f) - P_\varepsilon(f) \mathrm{d}f$ is not a function of the gain g. In the fractional delay case this does not hold for our model. This has implications for actual implementation of OFC with arbitrary delays, as it may require that the signal and noise PSDs be separated out of the joint estimate.

Given this, we directly minimize the residual error through the known power spectra of the input and noise and the transfer functions. This minimization of $P_\phi(f)|H_e(f)|^2 + P_v(f)|H_n(f)|^2$, where the transfer functions were given in Eqs. (44) and (45), produces the gain g. For any given τ, a maximum gain limit is determined such that the phase margin is 45 deg at this gain. For example, when $\tau = 1.5T$, the maximum gain is 0.406. Because we find little discrepancy between the Laplace model and the Simulink model, the gains as optimized by Laplace produce the best performance in Simulink. This was verified by running the Simulink simulation with a variety of different gains on the integral controller. The gain optimization with the Laplace model produced the best performance in Simulink. A specific example of this, for $\tau = 1.5T$ and SNR=1, is shown in Fig. 7. The mean and standard deviation of the Simulink results are given as data points and error bars. When the error bars are not visible, they are less than the size of the data point marker. The theoretical predictions based on the Laplace analysis are given with the curves. The gain as determined by optimization with the Laplace model indeed produces the best performance across a wide range of controller gains. Figure 7 also shows the predictor performance for a variety of Q^{-1}'s; different controllers were generated by manipulating the WFS noise variance input to the ARE solve. The state space model given in Subsection 2.A correctly produces the best performance for a family of related predictors.

For our case we have found that we can use a Laplace transform model of the AO system to optimize an integral controller and, as discussed above, evaluate stability margins. This result disagrees with the recent work of Looze [5], who found that his specific Laplace model did not correctly pick the optimal gains for his Simulink model. There are many differences between our treatment and that of Looze, not limited to the Simulink model, the atmospheric input, the Laplace model, and the control law $C(z)$. It is beyond the scope of this paper to provide a detailed consideration of which factors contribute to the difference in the results. What is sufficient for the purposes of this paper is that our Simulink performance results, as detailed below, are not biased by using an integral controller with a nonoptimal gain.

C. Performance

System performance was determined for a variety of SNRs and delays. The residual error powers and standard deviations for low and high SNRs are shown in Fig. 8. Again, the Laplace results are given with the curves, and the Simulink results are given as data points and error bars. There is very good agreement between the Laplace theory and the Simulink results, which should be expected since there are only small differences between the models.

This formulation of PFC for arbitrary delays allows the predictor to behave in a graceful fashion: The residual error essentially follows a linear interpolation of the performance at whole time step delays. A second important result is that the predictor improves its advantage over the integrator as delays increase. At SNR=10, the predictor has 20% less residual error at the shortest delay and 55%

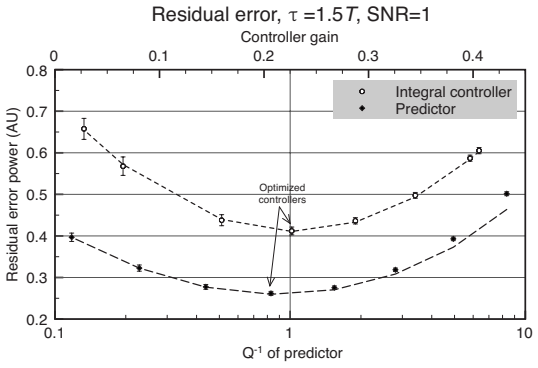

Fig. 7. Residual error performance of the integral controller and the predictive filter for a variety of gains. The optimized gain (as determined by Laplace analysis) results in the best performance in Simulink. The predictor was manipulated by feeding an incorrect WFS noise level into the ARE solve, producing a range of predictors. Again, the predictor produced by the model results in the best performance in Simulink.

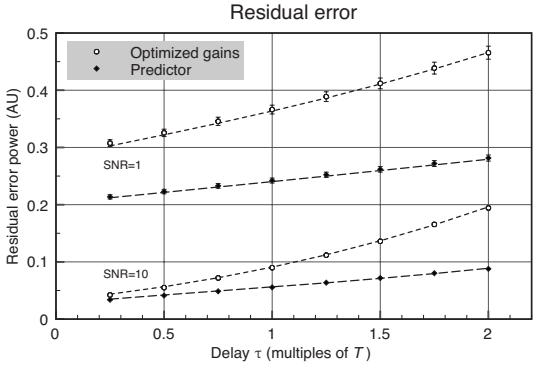

Fig. 8. Residual error for the predictor and the optimized-gain integrator. Input atmospheric power of 1; WFS noise set by SNRs of 1 and 10. Simulink results are the mean (data points) and standard deviation (error bars) of 32 different 8 s Simulink runs. Where error bars are not visible, the standard deviation was smaller than the data point marker size. Theoretical results from the Laplace model are given as dashed curves. The predictor performs gracefully for arbitrary delays, providing more of an advantage over the integral controller for longer delays.

less error at the longest delay. The predictor here has the same performance at $\tau=2T$ as the optimized integrator at $\tau=T$, which means that for this specific input the predictor enables an extra delay of 500 μs. At SNR=1, the predictor starts out with 30% less error, which is increased to 40% less error at the longest delay. At SNR=1, the predictor at all delays had less error than the integral controller at the shortest delay for this input atmosphere.

Some of the predictor's advantage comes from the fact that it can specifically notch out the higher temporal frequency layer components. Each layer that is in the model results in a notch in the error transfer function (as can be seen in Fig. 5), leading to improved rejection of that component. The depth of the notch is partly a function of the power in the layer. In this manner, the predictor can selectively attenuate the input while still limiting noise amplification. In contrast, the gain-optimized integrator must adjust its rejection simultaneously at all temporal frequencies. If the atmosphere does not have as much power in the layers as in this example, the predictor's advantage here will decline.

The predictor is also advantaged in that it has an adjustable lead filter, the coefficients of which are given in Eqs. (40) and (41). As the delay increases, a lead filter appears and then becomes higher order, providing the best adjustment for each delay. In contrast, the integral controller form is fixed for all delays. Which of these provides the most advantage depends on the layer strength. Following this result, an adjustable lead filter could be incorporated with the integral controller to provide better performance with variable system delays.

6. CONCLUSION

We have removed redundant states from our discrete-time state space model of a hybrid continuous-time/discrete-time AO control system and modified it to deal with arbitrary control loop delays. This modification is used to generate a new general formula for PFC. The fundamental structure of the predictive control law is independent of the exact system delay, resulting in ease of implementation. Furthermore, the fundamental coefficients (such as the SNR-based gain for each layer) are independent of the delay. Only the amount of prediction and the specific structure of the stabilizing lead filter depend on the delay τ.

The new PFC controller is essentially a linear interpolation of the controllers at integer time step delays. Consequently, system stability and performance transition gracefully between these whole step delays. A Simulink model of the hybrid system was used to analyze performance. As the computational delay τ increases, PFC provides more of an advantage over a gain-optimized integral controller. This is because it can selectively correct specific temporal frequency bands and adjust its lead filter to the arbitrary delay.

ACKNOWLEDGMENTS

The authors thank the reviewers of this paper for their detailed commentary. We want to particularly recognize the astute observation that the model had redundant states. Our consequent reformulation has resulted in a lower-order model that is simpler and provides significant computational savings. Thanks go to Bruce Macintosh for his insightful commentary during our research for this paper. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory in part under contract W-7405-Eng-48 and in part under contract DE-AC52-07NA27344. The document number is UCRL-JRNL-236046. 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. R. N. Paschall and D. J. Anderson, "Linear quadratic Gaussian control of a deformable mirror adaptive optics system with time-delayed measurements," Appl. Opt. **32**, 6347–6358 (1993).
2. D. T. Gavel and D. Wiberg, "Towards Strehl-optimizing adaptive optics controllers," Proc. SPIE **4839**, 890–901 (2002).
3. B. Le Roux, J.-M. Conan, C. Kulcsar, H.-F. Raynaud, L. M. Mugnier, and T. Fusco, "Optimal control law for classical and multiconjugate adaptive optics," J. Opt. Soc. Am. A **21**, 1261–1276 (2004).
4. C. Kulcsár, H.-F. Raynaud, C. Petit, J.-M. Conan, and P. V. de Lesegno, "Optimal control, observers and integrators in adaptive optics," Opt. Express **14**, 7464–7476 (2006).
5. D. Looze, "Discrete-time model of an adaptive optics system," J. Opt. Soc. Am. A **24**, 2850–2863 (2007).
6. C. Petit, J.-M. Conan, C. Kulcsar, H.-F. Raynaud, T. Fusco, J. Montri, and D. Rabaud, "First laboratory demonstration of closed-loop Kalman based optimal control for vibration filtering and simplified MCAO," Proc. SPIE **6272**, 62721T (2006).
7. K. Hinnen, M. Verhagen, and N. Doelman, "Exploiting the spatiotemporal correlation in adaptive optics using data-driven H_2-optimal control," J. Opt. Soc. Am. A **24**, 1714–1725 (2007).
8. L. A. Poyneer, B. A. Macintosh, and J.-P. Véran, "Fourier transform wavefront control with adaptive prediction of the atmosphere," J. Opt. Soc. Am. A **24**, 2645–2660 (2007).
9. L. A. Poyneer and J.-P. Véran, "Optimal modal Fourier transform wave-front control," J. Opt. Soc. Am. A **22**, 1515–1526 (2005).
10. P.-Y. Madec, "Control techniques," in *Adaptive Optics in Astronomy*, F. Roddier, ed. (Cambridge U. Press, 1999), pp. 131–154.
11. B. A. Macintosh, "Gemini planet imager preliminary design document volume 2: instrument design," Tech. Rep. UCRL-TR-230900 (Lawrence Livermore National Laboratory, 2007).
12. M. Frigo and S. G. Johnson, "The design and implementation of FFTW3," Proc. IEEE **93**, 216–231 (2005).
13. E. Gendron and P. Léna, "Astronomical adaptive optics I. Modal control optimization," Astron. Astrophys. **291**, 337–347 (1994).
14. G. Clark, S. Parker, and S. Mitra, "A unified approach to time- and frequency-domain realization of FIR adaptive digital filters," IEEE Trans. Acoust., Speech, Signal Process. **31**, 1073–1083 (1983).
15. W. Arnold and A. Laub, "Generalized eigenproblem algorithms and software for algebraic Riccati equations," Proc. IEEE **72**, 1746–1754 (1984).
16. A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-Time Signal Processing* (Prentice Hall, 1999).
17. A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals and Systems*, 2nd ed. (Prentice Hall, 1997).
18. C. Dessenne, P.-Y. Madec, and G. Rousset, "Optimization of a predictive controller for closed-loop adaptive optics," Appl. Opt. **37**, 4623–4633 (1998).