

Kalman filtering to suppress spurious signals in adaptive optics control

Lisa A. Poyneer^{1,*} and Jean-Pierre Véran²

¹Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore, California 94550, USA

²Herzberg Institute of Astrophysics, 5071 West Saanich Road, Victoria, British Columbia, Canada V9E2E7

*Corresponding author: poyneer1@llnl.gov

Received April 2, 2010; revised July 24, 2010; accepted August 13, 2010;
posted August 25, 2010 (Doc. ID 126447); published September 27, 2010

In many scenarios, an adaptive optics (AO) control system operates in the presence of temporally non-white noise. We use a Kalman filter with a state space formulation that allows suppression of this colored noise, hence improving residual error over the case where the noise is assumed to be white. We demonstrate the effectiveness of this new filter in the case of the estimated Gemini Planet Imager tip-tilt environment, where there are both common-path and non-common-path vibrations. We discuss how this same framework can also be used to suppress spatial aliasing during predictive wavefront control assuming frozen flow in a low-order AO system without a spatially filtered wavefront sensor, and present experimental measurements from Altair that clearly reveal these aliased components. © 2010 Optical Society of America

OCIS codes: 010.1080, 010.1285.

1. INTRODUCTION

Significant advances have been made in applying advanced control systems techniques to adaptive optics (AO) wavefront control. Most of these techniques are founded in the Kalman filter and, more broadly, the linear quadratic Gaussian (LQG) formalisms and are as such model-based. Important advances in this area include the general closed-loop Kalman filtering approach of Le Roux *et al.* [1] and the predictive Fourier control (PFC) of Poyneer *et al.* [2], which uses the Kalman framework to independently predict the Fourier modes of the wavefront under the assumption of frozen flow. Petit *et al.* experimentally demonstrated that the Kalman approach corrects tip-tilt vibration [3]. Recent work has focused on models that account for deformable mirror (DM) dynamics, such as Looze [4] and Correia *et al.* [5].

In this paper we use an expansion of the Kalman filter model to incorporate measurement noise that is not temporally white. This model expansion allows us to address two distinct problems in AO control. First, it allows us to produce a tip-tilt controller capable of advanced control (e.g., vibration rejection) of common-path error, while at the same time rejecting non-common-path error, such as vibrations that are sensed by the WFS but do not affect the science path. In his work, Petit specifically focused on rejecting common-path vibrations. He mentioned the possibility of extending his derivations to filter out non-common-path disturbances [see also [6] (in French)], but did not model them in the general framework of colored noise as we do and did not give any example where non-common-path disturbances were present. In this paper we present the theory and give detailed simulation results for the predicted Gemini Planet Imager (GPI) [7] tip-tilt environment.

The second application is to the PFC algorithm to deal with spatial aliasing in the WFS. We present the theory

and derive the controller structure, and present experimental data showing that aliases can easily be detected in a low-order AO system. We leave to further work a demonstration of the effectiveness of this algorithm via simulation.

2. INCORPORATION OF COLORED NOISE

In the standard formulation of the Kalman filter state space model, the measurement noise is temporally white. This means that any specific instant of measurement noise is temporally uncorrelated with all others. This is a reasonable assumption when considering WFS noise due to standard CCD noise sources (i.e., photon noise and read noise). However, noise due to another source may not be temporally white.

The AO state space model can be modified to deal with colored sources of noise. To do so, we follow the procedure detailed in Candy [8] (Subsection 5.8.2). We now present a summary of these derivations.

The AO state space model is composed of two equations: one that describes the temporal evolution of the state variables (AO phase to be corrected) and one that describes the measurement process with white noise. In a very general form, the first equation is

$$\mathbf{x}[t+1] = \mathbf{A}\mathbf{x}[t] + \mathbf{B}\mathbf{w}[t], \quad (1)$$

where \mathbf{x} are the state variables, \mathbf{A} describes the temporal evolution, $\mathbf{w}[t]$ is temporally white driving noise, and \mathbf{B} describes how that noise is used to generate the state. The second equation is

$$y[t] = \mathbf{C}\mathbf{x}[t] + \mathbf{D}\mathbf{u}[t] + v[t], \quad (2)$$

where $y[t]$ is the measurement (by the WFS) of the state as described by \mathbf{C} , and $v[t]$ is temporally white measure-

ment noise. The vector $\mathbf{u}[t]$ and matrix \mathbf{D} are used to incorporate the closed-loop DM correction.

Given Eqs. (1) and (2), we simply specify the variables and values and use the algebraic Riccati equation (ARE) and the state equation to solve for the controller. (See, for example, Sections 2 and 3 of Poyneer and Véran [9].)

In the case of colored measurement noise, we begin with the regular AO model state equation that models common path aberrations, now with slightly different notation:

$$\mathbf{x}_c[t+1] = \mathbf{A}_c \mathbf{x}_c[t] + \mathbf{B}_c \mathbf{w}_c[t]. \quad (3)$$

Now the measurement equation, instead of having white noise $v[t]$ has colored noise $z[t]$ instead:

$$y[t] = \mathbf{C}_c \mathbf{x}_c[t] + \mathbf{D} \mathbf{u}[t] + z[t]. \quad (4)$$

In order to use the Kalman equations on this model, we need to convert it to one where the measurement noise is white. We do this by expressing the colored noise $z[t]$ with its own state space model:

$$\mathbf{n}[t+1] = \mathbf{A}_n \mathbf{n}[t] + \mathbf{B}_n \mathbf{r}[t], \quad (5)$$

and

$$z[t] = \mathbf{C}_n \mathbf{n}[t] + v[t]. \quad (6)$$

In this model $v[t]$ remains the temporally white measurement noise of the original model.

To produce the new colored measurement noise state space model (of two equations) we augment Eqs. (3)–(6). The vectors become

$$\mathbf{x}[t] = (\mathbf{x}_c^T[t], \mathbf{n}^T[t])^T, \quad (7)$$

and

$$\mathbf{w}[t] = (\mathbf{w}_c^T[t], \mathbf{r}^T[t])^T. \quad (8)$$

The matrices for Eqs. (1) and (2) are then

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_c & 0 \\ 0 & \mathbf{A}_n \end{pmatrix}, \quad (9)$$

$$\mathbf{B} = \begin{pmatrix} \mathbf{B}_c & 0 \\ 0 & \mathbf{B}_n \end{pmatrix}, \quad (10)$$

and

$$\mathbf{C} = (\mathbf{C}_c \ \mathbf{C}_n). \quad (11)$$

For use further on, we will also now define a modified measurement matrix

$$\mathbf{C}_{\text{mod}} = (\mathbf{C}_c \ \mathbf{0}) \quad (12)$$

that picks out only the common-path phase and ignores the non-common-path phase. We can solve for the Kalman filter, which will now correct the same aberrations as before, and deal with the colored noise in a sensible fashion.

3. VIBRATION REJECTION FOR TIP-TILT

As mentioned earlier, Petit has already described and experimentally demonstrated a tip-tilt controller that rejects common-path vibration. In that formulation, each vi-

bration is modeled with as a second-order, real-valued process characterized by a temporal frequency and a damping coefficient.

A. State Space Model

In our model we will control just one mode of the tip-tilt disturbance (e.g., just tilt). As such we have a single variable that represents the measured phase and the residual error. In many scenarios this is entirely reasonable. However, if analysis of an AO system reveals correlations between tip and tilt disturbances, for best results the state space model should incorporate both. Doing so is a simple extension of the model presented below, so for simplicity and clarity we will consider only one variable in our model and refer to it as tilt.

In our formulation, primarily due to its heritage from PFC, we choose to model the vibration using a complex-valued parameter α . The magnitude of α sets the bandwidth of the vibration and the phase of α sets the positive and negative temporal frequency of the real-valued vibration. This is implemented with the real $\alpha_r = \mathcal{R}\{\alpha\}$ and imaginary parts $\alpha_i = \mathcal{I}\{\alpha\}$ as

$$\begin{pmatrix} a_r[t+1] \\ a_i[t+1] \end{pmatrix} = \begin{pmatrix} \alpha_r & -\alpha_i \\ \alpha_i & \alpha_r \end{pmatrix} \begin{pmatrix} a_r[t] \\ a_i[t] \end{pmatrix} + \begin{pmatrix} w_r[t] \\ w_i[t] \end{pmatrix}. \quad (13)$$

The state variable a_r is used as the vibration signal. We can deal with N_a vibrations and place their variables in a vector $\mathbf{a} = [a_{r,1}, a_{i,1}, \dots, a_{r,N_a}, a_{i,N_a}]^T$.

We model the other tilt components (that are not vibration), which are usually low-temporal-frequency errors due to atmospheric turbulence and windshake, with a state vector \mathbf{g} and transition matrix \mathbf{G} :

$$\mathbf{g}[t+1] = \mathbf{G}\mathbf{g}[t] + \mathbf{B}_g \mathbf{w}_g[t], \quad (14)$$

where the driving white noise is $\mathbf{w}_g[t]$ with noise matrix \mathbf{B}_g . For now we leave the order of this model arbitrary at N_g state variables.

In Poyneer and Véran [9] we presented our reduced-order closed-loop AO state space model that dealt with arbitrary control loop delays. In this treatment we will use the reduced state vector (which produces computational savings) and assume a full time step delay and not utilize the arbitrary-delay features of the modes. Here we expand our model, using the procedure outlined above in Section 2.

In this model the state vector \mathbf{x}_c contains state variables necessary for the time evolution of the common-path phase aberration, in this case the atmosphere; windshake $\mathbf{g}[t]$; vibrations $a_r[t], a_i[t]$; etc. In addition, it has the aberration at the previous point in time $\varphi[t-1]$:

$$\mathbf{x}_c = [\mathbf{g}^T, \mathbf{a}^T, \varphi[t-1]]^T. \quad (15)$$

The state transition matrix is

$$\mathbf{A}_c = \begin{pmatrix} \mathbf{G} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0}^T & \alpha_{r,1} & -\alpha_{i,1} & \cdots & 0 & 0 & 0 \\ \mathbf{0}^T & \alpha_{i,1} & \alpha_{r,1} & \cdots & 0 & 0 & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ \mathbf{0}^T & 0 & 0 & \cdots & \alpha_{r,N_a} & -\alpha_{i,N_a} & 0 \\ \mathbf{0}^T & 0 & 0 & \cdots & \alpha_{i,N_a} & \alpha_{r,N_a} & 0 \\ 1,0,\cdots,0 & 1 & 0 & \cdots & 1 & 0 & 0 \end{pmatrix}. \quad (16)$$

The driving noise vector $\mathbf{w}_c[t]$ has $N_g + 2N_a$ independent Gaussian variables. The noise covariance matrix \mathbf{P}_w is diagonal; the exact variances will be specified when the model is applied. The matrix \mathbf{B}_c controls the feeding of the driving noise into the state. This makes the matrix be $N_g + 2N_a$ columns and $N_g + 2N_a + 1$ rows:

$$\mathbf{B}_c = \begin{pmatrix} \mathbf{B}_g & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{2N_a} \\ 0 & \mathbf{0}^T \end{pmatrix}. \quad (17)$$

Assuming that the computational delay (WFS read, all computer calculations up to application of the mirror commands) is one frame, $\mathbf{u}[t]$ is just the last mirror command $d[t-1]$. This makes the matrix $\mathbf{D} = [-1]$. The measurement matrix is $\mathbf{C}_c = [\mathbf{0}^T, 1]$, which is $N_g + 2N_a + 1$ elements long and all zeroes except the last entry, which picks out $\varphi[t-1]$. Note that this implicitly assumes that (i) pure tilt can be produced by the device used to correct for tilt errors, and (ii) no other effects degrade the WFS measurements. In a real AO system, assumption (i) is very reasonable (corrector is a tip-tilt mirror or a deformable mirror that can produce tilt very accurately), but assumption (ii) may not be correct due to spatial aliasing on the WFS. WFS aliasing could be rejected by the use of a spatial filter in the system, or taken into account in our model, as explained in Section 4.

For the N_b non-common-path vibrations, we model them in the same way as the common-path vibrations, using the variable b and the parameter β instead of a and α . This makes $\mathbf{n}[t] = [b_{r,1}, b_{i,1}, \cdots, b_{r,N_b}, b_{i,N_b}]^T$. The non-common-path state vector $\mathbf{x}_n[t] = [\mathbf{b}[t]^T, \rho[t-1]]^T$, where $\rho[t-1]$ is the last total non-common phase error. The transition matrix is then

$$\mathbf{A}_n = \begin{pmatrix} \beta_{r,1} & -\beta_{i,1} & \cdots & 0 & 0 & 0 \\ \beta_{i,1} & \beta_{r,1} & \cdots & 0 & 0 & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & \beta_{r,N_b} & -\beta_{i,N_b} & 0 \\ 0 & 0 & \cdots & \beta_{i,N_b} & \beta_{r,N_b} & 0 \\ 1 & 0 & \cdots & 1 & 0 & 0 \end{pmatrix}. \quad (18)$$

The driving noise $\mathbf{r}[t]$ for this is $2N_b$ independent Gaussian variables. The noise covariance matrix \mathbf{P}_r is diagonal; the exact variances will be specified when the model is applied. Similarly,

$$\mathbf{B}_n = \begin{pmatrix} \mathbf{I}_{2N_b} \\ \mathbf{0}^T \end{pmatrix}. \quad (19)$$

For the measurement part of the non-common-path state model, the measurement matrix simply picks off $\rho[t-1]$ with $\mathbf{C}_n = [\mathbf{0}^T, 1]$, which is $2N_b + 1$ elements long and all zeroes except the last entry.

At any given time, the command to be sent to the tilt correcting device is $\varphi(t)$ (unity gain feedback), since we are assuming one frame of pure delay in the system. $\varphi[t]$ can be readily computed from $g_1[t]$ and $\mathbf{a}[t]$ in the estimated state vector.

B. Determination of Control

We now have all the matrices necessary to complete our state space model for correcting atmosphere and windshake tilt plus common-path and non-common-path vibrations. All we have to do is determine the signal characteristics (i.e., for a vibration, we need to know the power level to set the driving noise variance, along with the bandwidth and temporal frequency to set α or β .)

Given the state space model, we can apply the Kalman filter equations to update the state vector. Given the updated state vector, we then determine the new control signal for the wavefront correction device. As in our previous work [2,9], we assume that the best mirror command to apply at the next time step is $\hat{\varphi}[t+1|t]$, which is our best estimate of the phase at the next time step based on all measurements up to and including $y[t]$.

We calculate the control as follows. We begin with the new best state estimate that incorporates the current measurements:

$$\hat{\mathbf{x}}[t|t] = (\mathbf{I} - \mathbf{K}_s \mathbf{C}) \mathbf{A} \hat{\mathbf{x}}[t-1|t-1] + \mathbf{K}_s (y[t] - \mathbf{D} \mathbf{u}[t]). \quad (20)$$

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

$$\mathbf{K}_s = \mathbf{P}_s \mathbf{C}^H (\mathbf{C} \mathbf{P}_s \mathbf{C}^H + \mathbf{P}_v)^{-1}. \quad (21)$$

This matrix is found by numerically solving the ARE:

$$\mathbf{P}_s = \mathbf{A} \mathbf{P}_s \mathbf{A}^H + \mathbf{B} \mathbf{P}_w \mathbf{B}^H - \mathbf{A} \mathbf{P}_s \mathbf{C}^H (\mathbf{C} \mathbf{P}_s \mathbf{C}^H + \mathbf{P}_v)^{-1} \mathbf{C} \mathbf{P}_s \mathbf{A}^H. \quad (22)$$

We now obtain $\hat{\varphi}[t+1|t]$ from $\hat{\mathbf{x}}[t|t]$ as

$$\hat{\varphi}[t+1|t] = \mathbf{C}_{\text{mod}} \mathbf{A} \hat{\mathbf{x}}[t|t]. \quad (23)$$

The multiplication $\mathbf{A} \hat{\mathbf{x}}[t|t]$ produces our prediction one step ahead, and then $\mathbf{C}_{\text{mod}} \mathbf{A}$ picks out the state components that sum to the phase $\varphi[t+1|t]$. This is our new command to the wavefront corrector that will be used at time $t+1$, namely, $d[t+1]$.

In the case of tilt we are not computationally limited as in high-order wavefront control, so we implement the controller in matrix form via Eq. (20) and then the final dot product to obtain $\hat{\varphi}[t+1|t]$. We can efficiently implement the matrix multiplications by taking advantage of the fact that \mathbf{A} is sparse.

C. Stability and Relationship to LQG

To fully assess the stability of our closed-loop control, we will consider our state space model and control signal in the context of LQG control. We choose to follow the de-

tailed treatment of Looze [10], section 2.D, though other treatments exist in the adaptive optics literature.

Beginning with a state-space model of closed-loop AO control, the control signal is derived as a solution to the LQG problem via the solution of two AREs. If valid solutions exist for these two AREs—equivalently if two specific matrices are stable—then the closed-loop controller is as well. To use this analysis, we need to recast our model so that the control signal is incorporated into the state [as it is in Eq. (7) of [10]] as opposed to our model where it is in the measurement Eq. (2).

This modeling difference is easy to fix by simply augmenting the state with two additional variables at the head of the vector, and then modifying the existing matrices. We use the following definitions, where the tilde indicates a modified matrix used for the stability analysis given in [10]:

$$\tilde{\mathbf{A}} = \begin{pmatrix} \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} & \mathbf{0} \\ \mathbf{0} & \mathbf{A} \end{pmatrix}, \quad (24)$$

$$\tilde{\mathbf{B}} = [1, \mathbf{0}]^T, \quad (25)$$

$$\mathbf{C}_{\text{aug}} = [0, -1], \quad (26)$$

$$\tilde{\mathbf{C}} = [\mathbf{C}_{\text{aug}}, \mathbf{C}], \quad (27)$$

and

$$\tilde{\mathbf{W}} = \begin{pmatrix} \mathbf{0} \\ \mathbf{B} \end{pmatrix}. \quad (28)$$

Finally, we need to define the weighting criterion $\mathbf{Q}_p = \mathbf{I}$ and calculate \mathbf{Q} as [following Eq. (12) of [10]]

$$\mathbf{Q} = \begin{pmatrix} \mathbf{C}_{\text{aug}}^T \mathbf{Q}_p \mathbf{C}_{\text{aug}} & -\mathbf{C}_{\text{aug}}^T \mathbf{Q}_p \mathbf{C}_{\text{mod}} \\ -\mathbf{C}_{\text{mod}}^T \mathbf{Q}_p \mathbf{C}_{\text{aug}} & \mathbf{C}_{\text{mod}}^T \mathbf{Q}_p \mathbf{C}_{\text{mod}} \end{pmatrix}. \quad (29)$$

To ensure that we measure the residual error relative correctly, we use \mathbf{C}_{mod} , which includes only the common-path phase as measured with \mathbf{C}_c and not the non-common-path as measured with \mathbf{C}_n .

The solution to the LQG problem is given by Eq. (14) of [10]—to minimize the variance of the residual phase $\varphi[t] - d[t]$. Using this analysis, our result is that the LQG control is exactly what we had assumed, namely Eq. (23).

To verify that our model and controller are stable, we now check three matrices. First, as we have done before with Eq. (2), we examine \mathbf{A} . If the eigenvalues of \mathbf{A} have magnitude less than one, our model is stable. For PFC this analysis was trivial: since \mathbf{A} was lower triangular, the eigenvalues are on the diagonal.

For our tilt controller, \mathbf{A} is no longer lower diagonal. However, its structure lends itself to easy analytic calculation: the eigenvalues are 0, α , and α^* for each of the N_a vibrations, and β and β^* for each of the N_b vibrations and the eigenvalues of the matrix \mathbf{G} . By definition the magnitudes of the α 's and β 's are less than one. So as long as \mathbf{G} has eigenvalues with magnitudes less than one, the controller is stable. This is easy to do for reasonable models of atmospheric tilt and wind shake.

To verify that the controller itself is stable, we check the eigenvalues of two other matrices. Following [10], we need to show that $\tilde{\mathbf{A}} - \tilde{\mathbf{B}}\tilde{\mathbf{G}}$ is stable. Since in our case $\tilde{\mathbf{G}} = [0, 0, \mathbf{C}_{\text{mod}}\mathbf{A}\mathbf{A}]$, this is easy to calculate because the eigenvalues are the same as for \mathbf{A} .

Finally, we need to show that $\tilde{\mathbf{A}} - \tilde{\mathbf{L}}\tilde{\mathbf{C}}$ is stable. In our case, $\tilde{\mathbf{L}}$ is obtained from solving the same ARE (with augmentations) as Eq. (22). The eigenvalues of $\tilde{\mathbf{A}} - \tilde{\mathbf{L}}\tilde{\mathbf{C}}$ are the same as those of $\mathbf{A}(\mathbf{I} - \mathbf{K}_s\mathbf{C})$, where the Kalman gains are solved for with Eq. (21) and Eq. (22). So when we solve the ARE to determine the steady-state Kalman gains for the filter, we can just do a quick numerical check to evaluate the eigenvalues to ensure stability.

In summary, our system, following the LQG treatment of Looze, applies the LQG controller and is stable in closed loop as long as our model is stable and the results of solving the ARE for the Kalman gains produce a stable matrix $\mathbf{A}(\mathbf{I} - \mathbf{K}_s\mathbf{C})$.

D. Implementation Issues

We have already described how to calculate the control law in the closed-loop system given a specific state space model. It is necessary to estimate those model parameters. In our previous work we have relied on a block-adaptive implementation where the state space parameters are estimated from closed-loop measurements (see Section 4 of [2]). This was possible because of the computational efficiency of our model identification strategy.

For tilt we plan on doing model fitting based on the same temporal power spectral density (PSD) approach. There are at least four issues that may require future work, if the actual GPI tip-tilt environment warrants it.

First is the identification of common-path versus non-common-path phase for the model. To get any benefit from this type of control, we need to correctly classify the different components. Because the vibrations are due to instrument environment factors (e.g., cryo-coolers), we expect the temporal frequencies not to change, though the power levels probably will. An environmental analysis should be sufficient to identify the frequencies of common-path and non-common-path vibrations that will be used to classify the vibrations.

Second is the question will the tip-tilt controller be updated regularly in a block-adaptive fashion during closed-loop operation, or is it sufficient to do updates irregularly during open-loop time between science targets. This will depend strongly on how much the conditions change, and how computationally expensive the model fit step is. As we already have a method for closed-loop identification using temporal PSD estimation, the fall-back position of open-loop identification can use the same techniques without the closed-to-open conversion of the temporal PSDs of the measurements.

Third, fitting the atmosphere+windshake term model to the data may be a more involved process. In our simulation example below we have reduced the fit to a single variable. However, if we discover that the low-order behavior of the GPI tip-tilt environment is not so simple to fit (e.g., the cutoff varies because of the number of layers in the atmosphere), we will have to develop an efficient model fit strategy.

Finally, the vibration peaks can be identified using the same correlation with a reference peak shape as is used to find layer peaks in PFC. If the vibrations have a significant range of damping coefficients (which will change the peak width), a second step to estimate those may have to be included in the model fit.

In summary, there are several details to be worked out in how best to do the model identification. This will be informed, as will the final model structure, by the actual tip-tilt environment. None of these, however, should pose a significant challenge to implementation.

E. Gemini Planet Imager Tilt Simulation

In a high-contrast AO system such as GPI, it is critical to keep the star extremely well centered on the coronagraph [11]. Tip-tilt requirements are set primarily to preserve GPI's high Strehl ratio for the planetary companions; this produces a residual image motion requirement of 10 mas (goal of 6 mas). Since 8 mas (4.8 in the goal case) are allocated to post-wavefront sensor non-common-path vibrations and other uncontrollable sources, this leaves for our purposes just 6 mas (goal of 3.6 mas). This imposes a very stringent requirement on residual image motion. Because we are controlling a single variable (tilt) we assume below that the performance requirements above apply to the single variable.

Tip-tilt errors have different sources: atmospheric turbulence and telescope windshake, which are usually broadband and originate from the environment outside the observatory and vibrations, which are usually narrowband and originate from the telescope systems and/or the instruments. In this simulation we modeled the atmospheric turbulence using a von Karman spatial power spectrum with a telescope diameter $D=7.906$ m, a turbulence strength of $r_0=14.5$ cm at the wavelength of $0.5 \mu\text{m}$, an outer scale of $L_0=30$ m, and a wind speed of 15 m/s (direction of the wind is assumed to be the same as the direction of the tilt mode). This corresponds to typical (median) conditions at Gemini South. Assuming frozen flow, we derived the temporal power spectrum of the atmospheric tip-tilt from these parameters using the approach described in Conan *et al.* [12].

Since there is no reliable information available for windshake at Gemini South, we used the windshake model proposed for the Thirty Meter Telescope (TMT). This model is based on work from [13], and is detailed in [14]. The temporal power spectrum of the tilt error due to windshake has the same structure as that of the atmospheric tilt, with a plateau at low frequencies and the same $-17/3 \approx -6$ asymptotic power law at high frequency. To be conservative, we took the 85 percentile TMT profile as median windshake profile for Gemini South. This profile has a total of 35 mas RMS of tilt.

For the vibrations, we used three strong vibrations identified by us [15] in our analysis of telemetry data from the Near Infrared Coronagraphic Imager (NICI) WFS at Gemini South. That analysis cannot differentiate common-path vibrations (which affect the science image and should be corrected) from non-common path vibrations (which occur in the WFS path and thus should be ignored). So for the sake of this simulation, we assume that two are common-path vibrations: 4.5 mas RMS at

81 Hz, and 2 mas RMS at 279 Hz, and the other is non-common path: 1.7 mas RMS at 170 Hz.

To model the atmospheric+windshake tilt we considered models \mathbf{G} of different orders. In addition to the simple first-order model, we explored using higher-order models based on Butterworth filters. The flat behavior at low temporal frequencies is captured by the cutoff frequency of the low-pass Butterworth. The higher the order, the sharper the drop-off with temporal frequency. We tried both a second-order Butterworth and third-order. Though the third-order best matched the input temporal PSD, the resulting closed-loop error transfer function had a very low bandwidth that led to excess residual error. As such, we are using the second-order atmospheric model, which has better bandwidth. These different models have various tradeoffs which are very dependent on how well they fit the actual signal inputs. Due to some uncertainty in the exact nature of the GPI tip-tilt environment [15], we will not spend time trying to optimize the model to our assumed conditions; that is a task best left until actual tip-tilt information from GPI is available.

For a second-order model we use two states $\mathbf{g}=[g_1, g_2]$ and the state matrix is

$$\mathbf{G} = \begin{pmatrix} r \exp(j\theta) & 1 \\ 0 & r \exp(-j\theta) \end{pmatrix}. \quad (30)$$

There is only one driving noise that feeds into g_2 , so $\mathbf{B}_g = [[0, 0], [0, 1]]$. It is important to implement this in cascade form, as the closeness of the poles causes numerical errors and instability when multiplied out to direct form.

The first step in the simulation is to generate random time series based on our assumed inputs (e.g., atmospheric tilt as described above). This will be used first to identify the state space model parameters, and then in a closed-loop simulation. A time series is created from the temporal power spectra for each signal using white noise inputs through the spectral factor filtering method. These time series were 32768 time steps long at a sampling rate of 1.5 kHz. Based on our GPI end-to-end simulation as it stood during our Critical Design Phase, the total (tip + tilt) rms measurement noise is 2 mas for a magnitude $I=6$ natural guide star. This provides us with a moderate SNR; the optimized gain for this case is 0.3.

We then estimated the temporal PSD of the open-loop measurements from a single realization of the random time series. We did a model fit procedure to estimate the state space model coefficients. For the vibrations this is the same fit procedure as for the peaks in PFC. For the higher-order atmospheric and windshake models, the fit first selects a cutoff frequency θ and then optimize the r to get a best fit in the mean-square sense.

These parameters (α 's and β for the vibrations and \mathbf{G} matrix, and driving noise powers for everything as well as measurement noise) are then used to populate the full state space. A discrete-time closed-loop simulation on a single variable is then run, with the predictive controller implemented as described in Subsection 3.B, taking advantage of the sparse matrices. This was done for Kalman filters with and without non-common-path errors in their models. For comparison purposes, we also consider an optimized-gain integral controller.

First we will examine the transfer functions of the closed-loop system to illustrate the benefit that the extra modeling provides. Given the model parameters and the ARE solution above, we can determine the transfer functions via numerical evaluation. In Fig. 1, left panel, we show the error transfer functions for the optimized-gain integral controller and of the Kalman filter with the second-order atmospheric+windshake model. There is significantly more attenuation of low temporal frequencies with the second-order atmospheric model. This improved attenuation of the higher-order models can be exploited in high SNR situations. There are two notches in the error transfer function that correspond to the common-path vibrations. The depth of these will depend on the magnitude of the coefficient α used in the state space model and also on the SNR. A magnitude closer to 1 will deepen the notch. In higher-SNR situations these notches will be deeper; in lower-SNR situations they will become more shallow.

Noise transfer functions are shown in Fig. 1, right panel. For the new non-common-path Kalman filter, the deep notch in the noise transfer function now rejects that vibration and prevents it from propagating. There are two other notches in the noise transfer function. These are due to the common-path vibration rejection—each notch in the error transfer function has a matched dip in the noise transfer function, but at a slightly lower magnitude temporal frequency. The Kalman filters in this case all have higher noise propagation than the optimized-gain integral controller, but the second-order atmospheric model shown here is particularly higher.

These transfer functions illustrate how the Kalman filter with a sophisticated model can provide better low-temporal-frequency rejection and correct both common-path and non-common-path vibrations. When applied to the GPI tip-tilt scenario, these controllers enable us to meet our operating goals.

To estimate performance, we ran 32 new and different random time series through the closed-loop simulation. This discrete-time closed-loop simulation assumed a two-step delay from measurement to control, and operated at the same 1.5 kHz at which the inputs were sampled. For this single variable case we assumed perfect sensing of

tilt and perfect correction. We assumed a purely discrete model, where the phase is sampled once per frame; no continuous time approximations were used. (That is, the phase does not change during a time step.) Based on our previous work with both discrete simulations and Simulink [9], this is a reasonable approach for this scenario with full-frame delays.

For each trial, the RMS error on tilt is the square root of the sample variance of the error signal over the entire run (excluding initial convergence of the filter). Over the 32 trials, we calculated the square root of the sum of the squares (RSS) of the each RMS result. The results for this moderate-SNR case are given in Table 1. The optimized gain integrator produced only 0.88 mas RMS residual error due to atmosphere and windshake. However, using the second-order atmospheric model in the Kalman filter provides more rejection (as seen in Fig. 1), with a residual of only 0.022 to 0.024 mas RMS.

As for the vibrations, the integral controller cannot provide selective rejection; it passes through the 4.5 mas RMS common-path vibrations to 5.1 mas RMS residual error. Adding vibration rejection with Kalman filtering effectively removes the vibration, leaving 0.24 mas RMS. The integral controller also passes through the 1.7 mas RMS non-common-path vibration as 1.1 mas RMS error. Using the Kalman filter with non-common-path correction reduces this to 0.15 mas RMS.

The additional modeling terms allow us to improve correction of each different portion of the tilt environment and reduce the total error from 5.4 mas RMS to 2.5 mas RMS. This moves us from just making our requirement to easily meeting and exceeding our goal on the moderate-brightness guide star.

4. PREDICTION OF FROZEN FLOW ATMOSPHERE WITH ALIASING NOISE

The second application of the colored-noise model is to high-order wavefront control. In our previous work on PFC we developed a predictive filter that exploits the fact that each layer of frozen flow causes a concentration of power at a specific temporal frequency for each Fourier




Fig. 1. Error [top] and noise [bottom] transfer functions for the GPI tilt simulation. Use of a higher-order model for atmospheric tilt plus common-path vibrations results in an error transfer function with better low-frequency rejection and notches for the vibrations, as compared to an optimized-gain integrator. Additionally modeling non-common-path vibration produces a noise transfer function that notches out spurious signals. (Shallow secondary notches are a result of the common-path vibration prediction.)

Table 1. Tilt Error (mas RMS), 32 Trials, for GPI Tilt Simulation with Different Controllers^a

Signal	Input	Optimized Gain	Kalman, no NCP	Kalman, NCP
ATM+wind	72.3	0.88	0.022	0.024
CP vibration	4.7	5.1	0.24	0.24
NCP vibration	1.6	1.1	3.0	0.15
WFS noise	2.0	1.0	2.5	2.5
Total residual	—	5.4	3.9	2.5

^aThe Kalman filter uses the second-order ATM+windshake model.

mode. This allows easy identification and correction of multiple layers of turbulence.

However, in all our analysis and simulation we assumed that a spatially filtered wavefront sensor (SFWFS) [16] was used in the AO system. The spatial filter removes aliasing and exposes the underlying temporal error due to incomplete correction of the dynamic atmosphere. It is this error that PFC further reduces with prediction.

The SFWFS works well only in high-order AO systems (see original paper [16] for more discussion of why), though an oversized filter has been shown by Fusco *et al.* [17] to improve performance in a low-order AO test bench. Most present AO systems do not have spatial filters. Rigaut *et al.* [18] derived that for the Shack–Hartmann WFS the aliasing error is one-third the variance of the fitting error. For a curvature WFS, its variance is about equal to the variance of the fitting error [19]. As such, advanced control methods to further reduce temporal error may not have a significant impact on low-order AO systems without a spatial filter. Conversely, methods that could reduce spatial aliasing might be very useful.

It turns out that under frozen flow, aliased components can be picked out and ignored. Below we show experimental evidence that they can be detected, discuss how to model aliasing as a colored noise source, and finally present the control law that predicts the atmosphere while rejecting aliasing.

A. Experimental Observation of Frozen Flow Aliases

In our recent work, with colleagues, we used telemetry from the Keck and Altair AO systems to experimentally validate the frozen flow hypothesis and clearly detect multiple layers of translating turbulence [20]. Here we use some of that data obtained with Altair to demonstrate the detection of aliases. For all the details of the observation nights, experimental methods, etc., please consult [20].

To understand how aliases can be detected, we must first review how frozen flow turbulence is manifested in the Fourier modes of the wavefront aberration. We begin by considering the wavefront phase sampled on the DM actuator or WFS grid spacing d (which are assumed to be the same). For example, GPI's subaperture size in the pupil is approximately $d=18$ cm. The Fourier modes of the wavefront are obtained either in the reconstruction process (i.e., with Fourier transform reconstruction [21] from WFS measurements) or as an after-the-fact discrete Fourier transform (DFT) on the saved telemetry of the estimated phase. It is this second approach that we used in analyzing Altair telemetry.

The DFT is done on an $N \times N$ grid, where N is usually a few larger than D/d , where D is the pupil diameter in meters. For the Altair data, $d=66$ cm, so $D/d=12$, but we use $N=16$. The Fourier modes are indexed by frequency variables k and l and take the values $-N/2, -(N/2-1), \dots, -1, 0, 1, \dots, (N/2-2), (N/2-1)$. This gives each Fourier mode the frequency components, in units of m^{-1} , $f_x=k/(Nd)$ and $f_y=l/(Nd)$. Note that the fundamental limit on the highest frequency measured by the AO system is set by the subaperture size d ; increasing the grid size N on which we reconstruct will not change this.

Under frozen flow, each layer has a velocity vector given by v_x and v_y , with units of m/s. Pure translation of the layer gives rise to an impulse in the temporal frequency power spectrum of a Fourier modal coefficient at

$$f_t = f_x v_x + f_y v_y = \frac{k v_x + l v_y}{Nd}. \quad (31)$$

When the temporal power spectrum of a Fourier coefficient of the wavefront is obtained, a layer of frozen flow shows itself as a concentrated peak of power at that characteristic frequency f_t .

This equation for temporal frequency holds for all Fourier modes, even those beyond the highest frequency set by the actuator spacing d . When a spatially filtered WFS is used, aliasing is prevented and any phase estimate or DM command will not have spurious content. If there is no spatial filter, aliasing will occur. In an open-loop observation mode, spatial frequencies beyond the Nyquist limit set by the actuator spacing d will alias down and appear in the open-loop phase estimate. In closed-loop control, the AO system will try to correct the aliases, leading to their appearing on the DM commands in telemetry.

A given spatial frequency above the Nyquist frequency aliases down to specific controllable modes. This happens through a replication of the signal at multiples of the sampling frequency. (For an introduction to sampling theory, see Chapter 7 of Oppenheim *et al.* [22]). For a spatial frequency k, l in our controllable spatial frequency domain, it will not only have its true spatial frequency content after sampling, but will also have aliases from spatial frequencies $k+Nn_k, l+Nn_l$, where n_k and n_l range from $-\infty$ to ∞ (excluding $n_k=n_l=0$, of course). The power level of these aliases will be lower than the true signal because of the response of the WFS and because under the Kolmogorov model of atmospheric turbulence, higher spatial frequencies have less power.

As an illustration of this, we examine telemetry data from Altair for the night of April 4, 2008. In this case, using the algorithms described in (20), we identified a single layer of atmospheric turbulence with velocity vector $v_x=9.87$ m/s and $v_y=-11.56$ m/s at 90% likelihood. The mapping of these peaks is shown at the top left of Fig 2. The temporal frequency as a function of spatial frequency is clearly visible.

We consider the controllable Fourier mode $k=2, l=6$, which has a layer frequency $f_t=-4.8$ Hz. The four strongest possible aliases will come from $[n_k, n_l]=[-1, 0], [1, 0], [0, -1]$, and $[0, 1]$. The temporal frequency of the aliased

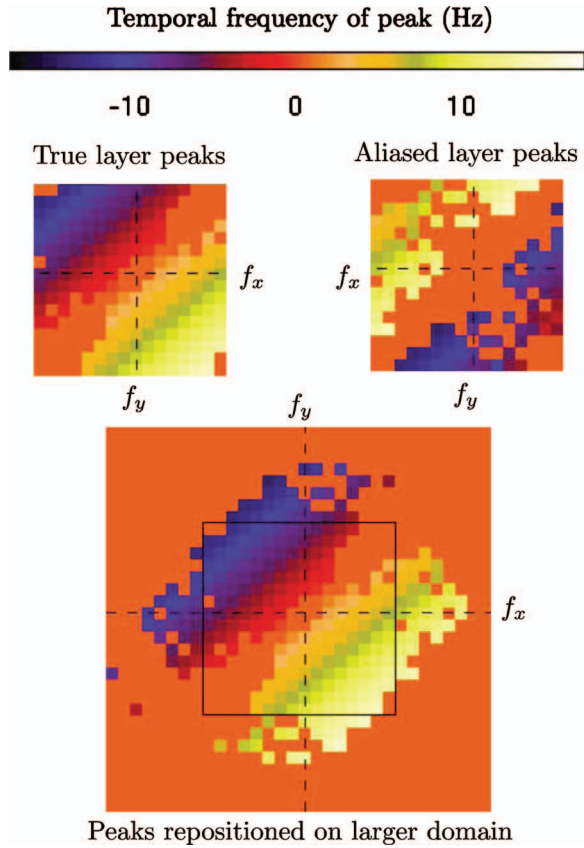


Fig. 2. Maps of found peaks and aliases, Altair April 4, 2008, data set. Each square panel is a f_x, f_y mapping of the spatial frequencies with dashed lines as the axes. The color of each point is the temporal frequency of a peak found in that Fourier mode’s temporal PSD. In “True layer peaks” each peak corresponds to a layer, which is identified with 90% likelihood. In “Aliased layer peaks” each identified peak matches with an alias of that layer. This shows the controllable spatial frequencies that are aliased into. Note the opposite pattern of temporal frequency increase from the true layer peaks. At bottom in “Peaks repositioned,” the aliases are returned to their parent high spatial frequencies, forming a layer map over a larger spatial frequency range. The solid square box outlines the controllable spatial frequencies of the AO system.

component is determined by substituting the frequencies $k + Nn_k, l + Nn_l$ into Eq. (31), resulting in layer frequencies given by

$$f_t = \frac{(k + Nn_k)v_x + (l + Nn_l)v_y}{Nd}. \tag{32}$$

For this specific Fourier mode, the aliases are given in Table 2. Of these four possible aliases, the last, $k=2,$

Table 2. Information on Aliased Components That Might Appear for Controllable Mode $k=2, l=6$ in Altair Data

Alias Offset	Aliased Mode	$(k_a^2 + l_a^2)^{1/2}$	Aliased f_t (Hz).
$n_k=1, n_l=0$	$k_a=18, k_l=6$	19.0	10.2
$n_k=-1, n_l=0$	$k_a=-14, k_l=6$	15.2	-19.6
$n_k=0, n_l=1$	$k_a=2, k_l=22$	22.1	22.1
$n_k=0, n_l=-1$	$k_a=2, k_l=-10$	10.2	12.8

$l=-10,$ should have the most power, as it has the lowest magnitude spatial frequency.

As shown in Fig. 3, the temporal power spectrum clearly shows a secondary peak at 12.8 Hz. Also included in the figure is the temporal PSD of the nearby mode, $k=4, l=7$. This mode has similar behavior, with a true layer frequency of -3.8 Hz and the alias from mode $k=4, l=-9$ at 13.5 Hz. These two modes illustrate a characteristic of aliased frozen flow versus true frozen flow. In true frozen flow, the temporal frequency relationship $f_t = f_x v_x + f_y v_y$ means that as the spatial frequency increases perpendicular to the velocity vector, the magnitude of f_t does as well. For aliases, as the spatial frequency of the controllable mode increases perpendicular to the velocity vector, the magnitude of f_t decreases. In this case, from mode $k=4, l=7$ to mode $k=2, l=6$, the true layer frequency increases in magnitude from -3.8 to -4.8 Hz, but the aliased layer frequency decreases in magnitude from 13.5 to 12.8 Hz. This is clearly seen in the top right panel of Fig. 2. Here the identified peaks that have been classified as aliases of this layer are shown. Note how the magnitude increases toward the origin, whereas the true layer frequencies on the left decrease toward the origin.

In the bottom panel the aliases are returned to the parent spatial frequencies outside the controllable range. Here we can see that aliases from spatial frequencies along the k and l axes are preferentially observed, just as would be predicted based on the response of the Shack-Hartmann WFS (see, for example, Fig. 2 of Jolissaint *et al.* [23].)

The fact that we can easily identify peaks caused by aliasing means that if we can incorporate them into the state space model and produce a Kalman filter that ignores them, we have a way to reduce the aliasing error in the case of frozen flow without the use of the spatial filter. We discuss the state space model and resulting predictive filter next, and then address the issue of actually finding these peaks and knowing that they are from aliasing.

B. Original State Space Model

In PFC, we consider each complex-valued Fourier mode of the phase aberration independently. As discussed above,

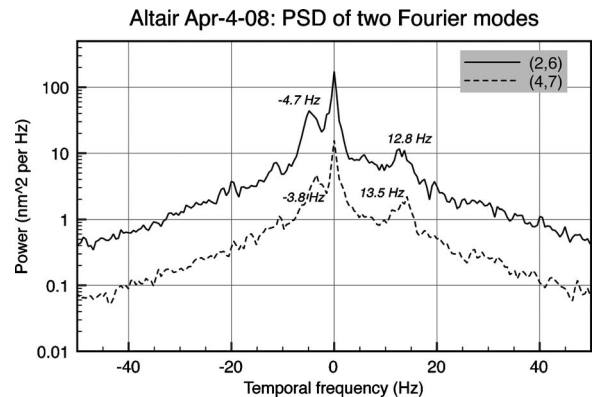


Fig. 3. Temporal power spectrum of Fourier mode 2,6 from a 1-minute-long Altair AO closed-loop run April 4, 2008. The single layer of wind identified for this observation causes the strong peak at -4.8 Hz. The secondary peak at 12.8 Hz is caused by an alias from Fourier mode $k=2, l=-10$. Nearby mode 4,7 has similar behavior and is also shown.

under frozen flow atmospheric turbulence, each layer produces a specific temporal frequency in a Fourier mode. This is modeled by an auto-regressive order (1) process.

$$a[t] = \alpha a[t-1] + w[t]. \quad (33)$$

The complex number α has magnitude just less than one. The phase of α sets how much the Fourier mode advances in a single time step of length T seconds. 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 N_a layers of frozen flow. The state variables for these layer components are given by

$$\mathbf{a}[t] = (\alpha_0[t], \alpha_1[t], \dots, \alpha_{N_a}[t]) \quad (34)$$

and the auto-regression parameters are stored in the matrix

$$\mathbf{A}_\alpha = \text{Diag}(\alpha_0, \alpha_1, \dots, \alpha_{N_a}). \quad (35)$$

The power levels of the driving noises are given by the covariance matrix

$$\mathbf{P}_w = \text{Diag}(\sigma_{\alpha_0}^2, \sigma_{\alpha_1}^2, \dots, \sigma_{\alpha_{N_a}}^2). \quad (36)$$

The reduced state, which has only the atmospheric layers and the phase φ at the previous time step, is

$$\mathbf{x}_c[t] = (\mathbf{a}[t], \varphi[t-1])^T. \quad (37)$$

The common-path portion of the state transition matrix is

$$\mathbf{A}_c = \begin{pmatrix} \mathbf{A}_\alpha & \mathbf{0} \\ \mathbf{1}^T & 0 \end{pmatrix}, \quad (38)$$

where the row vector $\mathbf{1}^T$ causes the phase aberration $\varphi[t]$ to be the sum of all of the layer terms. The driving noises are incorporated with

$$\mathbf{B}_c = \begin{pmatrix} \mathbf{I} \\ \mathbf{0}^T \end{pmatrix}. \quad (39)$$

The measurement vector is

$$\mathbf{C}_c = (\mathbf{0}^T, 1), \quad (40)$$

where the \mathbf{C}_c has $N_a + 2$ elements, all of which are zero except for the last one, which picks off the phase $\varphi[t-1]$ as the measurement. As for tilt, we assume the total system delay is one time step, making $\mathbf{D} = (-1)$ and $\mathbf{u}[t] = (d[t-1])$, where $d[t-1]$ is the DM command applied during the WFS measurement. This model is complete for executing predictive control in the case where the WFS is spatially filtered.

C. State Space Model with Aliases

The aliases, since they are due to frozen flow, follow the same complex-valued AR(1) model as the layer components. Just as with the true phase, the aliases are measured after a delay, so we use the variable $\rho[t]$ to represent the sum of the aliased phase. Given that N_b aliases are identified, we assemble their the state vector

$$\mathbf{n}[t] = (b_1[t], \dots, b_{N_b}[t], \rho[t-1]). \quad (41)$$

The auto-regression parameters are stored in the matrix

$$\mathbf{A}_\beta = \text{Diag}(\beta_1, \dots, \beta_{N_b}). \quad (42)$$

The non-common-path portion of the state transition matrix is

$$\mathbf{A}_n = \begin{pmatrix} \mathbf{A}_\beta & \mathbf{0} \\ \mathbf{1}^T & 0 \end{pmatrix}. \quad (43)$$

The driving noises are incorporated with

$$\mathbf{B}_n = \begin{pmatrix} \mathbf{I} \\ \mathbf{0}^T \end{pmatrix}. \quad (44)$$

The power levels of the driving noises are given by the covariance matrix

$$\mathbf{P}_r = \text{Diag}(\sigma_{b_1}^2, \dots, \sigma_{b_{N_b}}^2). \quad (45)$$

The WFS measurement just picks off $\rho[t-1]$ from the state with $\mathbf{C}_n = (\mathbf{0}^T, 1)$.

D. Determination of Control

Now that we have the full model, we need to generate the controller. First we examine the controller in matrix form. Since only the specific entries of the state space model matrices are different, but not the structure, we use the exact same equations as for the tilt controller in Subsection 3.B.

The only significant difference is that the variable that we are correcting is not a direct signal (e.g., tilt) but a single Fourier coefficient of the wavefront. This coefficient is obtained in closed loop during the reconstruction process, and the controller is applied inside the Fourier domain. Since we have an $N \times N$ grid, we have $N^2/2$ complex-valued Fourier modes to control independently.

So for each Fourier modal coefficient, we begin with our last state estimate $\hat{\mathbf{x}}[t-1|t-1]$. We use the new measurement $y[t]$ to calculate the new state estimate $\hat{\mathbf{x}}[t|t]$. Then we estimate $\hat{\varphi}[t+1|t]$ using Eq. (23). This produces for this Fourier mode the desired phase signal to be compensated. These estimates for all Fourier modes are arranged properly in an $N \times N$ signal, and the inverse DFT is calculated to obtain the spatial phase. This signal is then used with an appropriate fitting algorithm (e.g., influence-function precompensation and voltage-phase conversion [24]) to determine control voltages for the deformable mirror.

Though we could implement the controller with a series of sparse matrix multiplications for each mode, we can analytically determine a more concise form that provides both computational savings and insight into the structure of the filter. Our control law uses the measurement $y[t]$ to determine the phase we want to compensate with the mirror $\hat{\varphi}[t+1|t]$. As such, we can write a transfer function for that controller.

Note that an AO system typically uses an integral controller $C(z) = g/(1 - 0.99z^{-1})$. As we have detailed this kind of derivation before in previous work [2,9], we will omit here the details and simply present the filter results.

First we review the form of the original PFC controller. For a given Fourier mode, the control law is

$$C(z) = \left(\sum_{k=0}^{N_a} \frac{K_k \alpha_k}{1 - \alpha_k z^{-1}} \right) (1 + z^{-1} D_1)^{-1}, \quad (46)$$

where the constants are

$$K_k = Q^{-1} p_{N_a+1,k}, \quad (47)$$

$$D_1 = Q^{-1} \sum_{k=0}^{N_a} p_{N_a+1,k}, \quad (48)$$

and

$$Q = p_{N_a+1} + \sigma_v^2. \quad (49)$$

The p are specific entries in \mathbf{P}_s , which are solved for numerically. For proper implementation, it is important to note that the $p_{k,l}$ are indexed beginning at 0, not 1.

The original PFC filter has an elegant, easy to understand structure. First, the measured residual $y[t]$ is sent in parallel to integrators for each layer. [This is the summation term in the first set of parentheses in Eq. (46).] These layer integrators predict using α_k and weight based on SNR using K_k . The summed result of the layer integrators is then sent through a high-pass lead filter which ensures stability. [This is the term in the second set of parentheses in Eq. (46).]

The new predictive controller for the anti-alias PFC filter is

$$C(z) = \left(\sum_{k=0}^{N_a} \frac{K_k \alpha_k}{1 - \alpha_k z^{-1}} \right) \left(1 + z^{-1} D_1 + z^{-1} \sum_{k=1}^{N_b} \frac{C_k}{1 - \beta_k z^{-1}} \right)^{-1}, \quad (50)$$

where the constants are

$$K_k = Q^{-1} (p_{N_a+1,k} + p_{N_a+N_b+2,k}), \quad (51)$$

$$D_1 = Q^{-1} \sum_{k=0}^{N_a} (p_{N_a+1,k} + p_{N_a+N_b+2,k}), \quad (52)$$

$$C_k = Q^{-1} (p_{N_a+1+k, N_a+1}^* + p_{N_a+N_b+2, N_a+1+k}), \quad (53)$$

and

$$Q = p_{N_a+N_b+2} + \sigma_v^2. \quad (54)$$

The new filter has very similar structure to that of the original PFC filter. The parallel layer integrator structure remains exactly the same. The exact values of the K_k are now different, as the error covariances involving the aliases are incorporated. [Compare Eq. (47) to Eq. (51).] The high-pass lead filter still preserves the same structure involving D_1 , though with a slightly different value. [Again, compare Eq. (48) to Eq. (52).] Now the second filter stage incorporates alias removal. This is done through the summation term, which is a set of layer integrators for the aliases. Note here that since the integrated residual has already been predicted, no extra prediction

with β_k is necessary in the numerator, as it is with α_k in the first-layer integrators.

This second bank of layer integrators tracks the aliases and subtracts them from the DM command for that Fourier mode. This filter structure is illustrated in Fig. 4. The region with a gray background is the new alias removal portion of the filter.

E. Stability and Relationship to LQG

Because the anti-alias PFC state space model has the exact same form as the tilt controller, we apply the same analysis as in Subsection 3.C. Again, our choice of control signal is the solution to the LQG problem, where we here minimize the variance of the residual error $\varphi[t] - d[t]$ for each Fourier mode. Again, the stability check on the matrix $\tilde{\mathbf{A}} - \tilde{\mathbf{B}}\tilde{\mathbf{G}}$ is passed as long as our model \mathbf{A} is stable. Because \mathbf{A} is lower triangular, it suffices to ensure that the magnitudes of all of the α and β are less than one. As before, $\tilde{\mathbf{A}} - \tilde{\mathbf{L}}\tilde{\mathbf{C}}$ has the same eigenvalues as $\mathbf{A}(\mathbf{I} - \mathbf{K}_s\mathbf{C})$, where the Kalman gains are solved for with Eq. (21) and Eq. (22).

F. Implementation Issues

The significant challenge to implementing this anti-alias PFC filter is classification of the layer peaks in the temporal power spectra. In our original proposal [2] each Fourier mode is dealt with separately. Since the WFS is assumed to have a spatial filter, all peaks found are assumed to be from layers or from other sources (such as vibration) that we want to correct. Each Fourier mode has its own predictive controller generated via the ARE solver independently.

If there is no spatial filter in use, some of those peaks will be caused by aliasing. The problem is, if we consider each mode independently, how can we tell (as in Fig. 3) that for mode $k=2$, $l=6$ the peak at $f_l = -4.8$ Hz is from a real layer of frozen flow, but the peak at 12.8 Hz is from an alias? The only way we can determine if 12.8 Hz is an alias is to know the velocity vector of the layer. The only way to know the velocity vector of a layer is to examine all the Fourier modes and see if the peaks found are consistent with a single (or multiple) layers.

In our recent observation study of atmospheric characteristics in [20], we implemented just such an algorithm. We did this to prove that the peaks that we found in the temporal PSDs of Fourier modes were actually caused by frozen flow. This algorithm is a straightforward, brute-force method that calculates the likelihood of a layer for a range of velocity vectors. For a given velocity vector $[v_x, v_y]$, the layer frequency f_t is calculated for each Fourier mode. The fraction of Fourier modes that found a peak close to (usually within 2 Hz) the desired f_t is the likelihood of a layer of that specific velocity actually existing. A reasonable range of velocities (e.g., from -30 m/s to 30 m/s) for both v_x and v_y is searched, and the significant local maxima of the likelihood function (if any) are selected as true layers.

Once this analysis has been done, we have a list of layers that we will correct. Given the velocity vectors of these layers, it is easy to calculate the possible temporal fre-

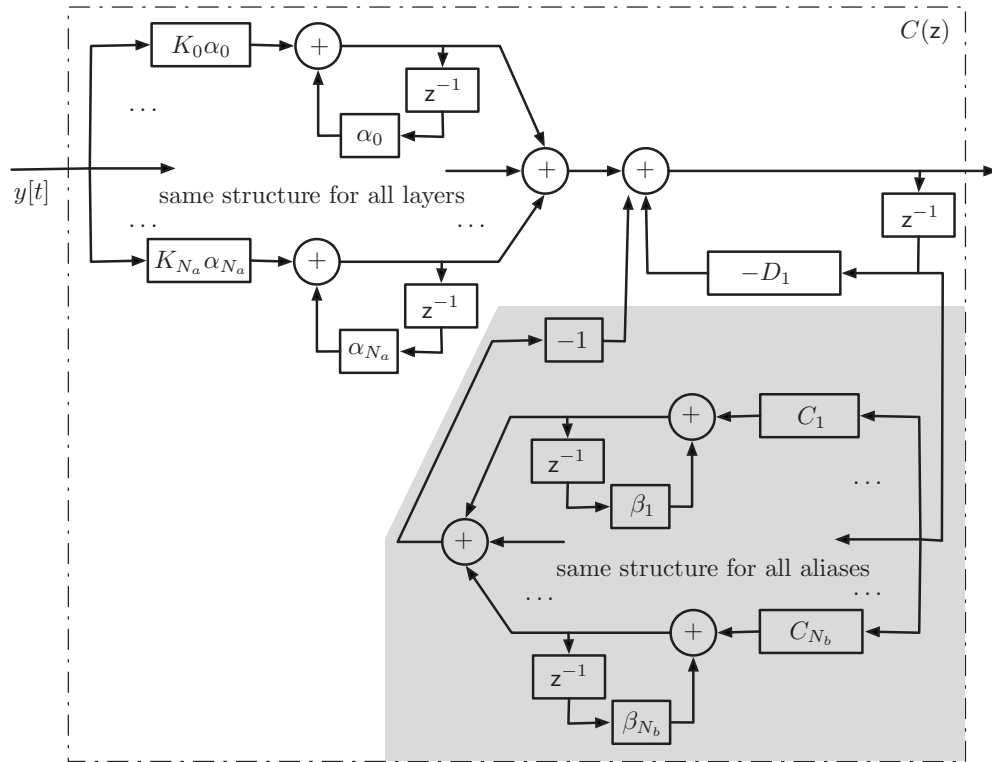


Fig. 4. Block diagram of anti-alias predictive controller for a single Fourier mode. First the layers are predicted in parallel with layer integrators. The result is summed and sent through a second filter. This filter predicts the aliases and subtracts them to produce the DM command. The area shaded in gray is the new portion due to incorporation of alias removal. The regular PFC filter has the exact same structure as the white region.

quencies of the strongest aliases for a constrained range of n_b and n_l and see if any peaks were found in the temporal PSDs.

A second implementation issue is deciding which layers and aliases to correct, and whether or not to correct them in all modes, as opposed to only the modes in which they are identified. As mentioned above, in our experimental data example the layer was identified with 90% likelihood. This means that 90% of the modes that could see that layer actually had a clear peak in the temporal PSD. Some 60% of those modes also had a clear alias peak that was identified. The question then becomes, should we construct a model across all modes for the layer velocity and power, and generate *estimated* state space model parameters, or should we remain fully data-based and use only the peaks and power levels found in each mode? The first method would allow more uniform correction and provide robustness against missed detections. It would also, however, require extra computational cost and introduce an unknown amount of model–data mismatch.

A third issue is what to do about aliases from the static layer. Since there is no temporal variation, we cannot separate out the aliased component of the static layer from the controllable component. Aliases from a static error could be calibrated out by measurement (e.g., phase diversity using science images) and appropriate modification of reference centroids.

Finally, we discuss the relationship of our proposed anti-alias filter with the approach of Petit *et al.* [25]. In that LQG approach, specific uncontrollable modes (that have high spatial frequencies above Nyquist) are used in

the estimation given a certain statistical prior. Once these modes are estimated, they are excluded from the control. This allows the excess power at specific controllable frequencies near Nyquist (which already violates the Kolmogorov prior) to be instead estimated as being from an alias.

In contrast, we estimate any aliased component using the frozen flow assumption, not the Kolmogorov spatial power spectrum. To see the difference, consider a pure Fourier mode of phase aberration sensed by the WFS at spatial frequency k, l . If it is a true phase measurement (not an alias) its spatial frequency in the pupil is $f_x = k/(Nd)$ and $f_y = l/(Nd)$. However, in the absence other information, we have no way to tell if this measurement is truly that of a controllable phase, or if it is from an alias (and then which alias). By using a spatial prior, an LQG controller could estimate which portion of the signal is controllable and which is not. In our method, we use the temporal behavior of that Fourier mode to disentangle true phase from alias. First we analyze the temporal PSD and identify any peaks that are due to frozen flow. With knowledge of the true layers, we can then classify those peaks as coming from controllable phase and specific aliases. We then directly control the desired portion, and notch out the aliases. We are not distributing the signal among controllable and uncontrollable modes statistically, but instead directly separating out the components based on their temporal variations.

Because of the complexities of implementation described above, we leave to further work the computationally efficient implementation of the above algorithm and

the determination in simulation of the performance improvement possible with the new alias-suppressing PFC filter.

5. CONCLUSIONS

We have demonstrated how to use a Kalman filtering state space model with colored noise to model AO closed-loop control with temporally non-white noise. Then we modeled and developed filters for two different cases. For the case of tip-tilt control, we modeled non-common-path vibrations and demonstrated through simulation of the anticipated GPI tip-tilt environment that improved modeling of atmospheric+windshake errors, common path vibration, and non-common-path vibration leads to significantly improved correction. Performance is model dependent, but these modeling advances provide us security that once GPI is installed and the actual tip-tilt environment is known, a controller that meets requirements will be implementable.

In the case of frozen flow atmospheric turbulence, we expanded our existing Predictive Fourier Control model to account for aliasing in the WFS. We provided an example of experimental data from Altair where the aliases are clearly identifiable in closed-loop telemetry. We presented the control law which predicts the atmosphere and provides rejection of aliasing. This new filter has the same structure as the original PFC filter, but with the added step of predicting aliases and removing them from the DM commands. A detailed performance study through simulation was left to future work.

ACKNOWLEDGMENTS

We thank the reviewers for their detailed and perceptive commentary, in particular regarding our stability analysis. Thanks to Bruce Macintosh for guidance on system-level specifications and tolerances for GPI. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. The document number is LLNL-JRNL-426892. This work has been supported by the National Science Foundation (NSF) Science and Technology Center for Adaptive Optics managed by the University of California at Santa Cruz under cooperative agreement No. AST-9876783.

REFERENCES

1. B. Le Roux, J.-M. Conan, C. Kulcsár, 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).
2. 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).
3. C. Petit, J. M. Conan, C. Kulcsár, H. F. Raynaud, and T. Fusco, "First laboratory validation of vibration filtering with LQG control law for adaptive optics," *Opt. Express* **16**, 87–97 (2008).
4. D. P. Looze, "Linear-quadratic-Gaussian control for adaptive optics systems using a hybrid model," *J. Opt. Soc. Am. A* **26**, 1–9 (2009).
5. C. Correia, H.-F. Raynaud, C. Kulcsár, and J.-M. Conan, "On the optimal reconstruction and control of adaptive optical systems with mirror dynamics," *J. Opt. Soc. Am. A* **27**, 333–349 (2010).
6. C. Petit, "Étude de la commande optimale en OA et OAMC, validation numérique et expérimentale," Ph.D. thesis (L'Université Paris 13, 2006).
7. B. Macintosh, J. Graham, D. Palmer, R. Doyon, D. Gavel, J. Larkin, B. Oppenheimer, L. Saddlemyer, J. K. Wallace, B. Bauman, J. Evans, D. Erikson, K. Morzinski, D. Phillion, L. Poyneer, A. Sivaramakrishnan, R. Soummer, S. Thibault, and J.-P. Véran, "The Gemini Planet Imager," *Proc. SPIE* **6272**, 62720L (2006).
8. J. Candy, *Model-Based Signal Processing* (Wiley, 2005).
9. L. A. Poyneer and J.-P. Véran, "Predictive wavefront control for adaptive optics with arbitrary control loop delays," *J. Opt. Soc. Am. A* **25**, 1486–1496 (2008).
10. D. P. Looze, "Minimum variance control structure for adaptive optics systems," *J. Opt. Soc. Am. A* **23**, 603–612 (2006).
11. A. Sivaramakrishnan, R. Soummer, L. Pueyo, J. K. Wallace, and M. Shao, "Sensing phase aberrations behind Lyot coronagraphs," *J. Astrophys. Astron.* **688**, 701–708 (2008).
12. J.-M. Conan, G. Rousset, and P.-Y. Madec, "Wave-front temporal spectra in high-resolution imaging through turbulence," *J. Opt. Soc. Am. A* **12**, 1559–1570 (1995).
13. D. G. MacMynowski, K. Vogiatzis, G. Z. Angeli, J. Fitzsimmons, and J. E. Nelson, "Wind loads on ground-based telescopes," *Appl. Opt.* **45**, 7912–7923 (2006).
14. TMT reference document, "Observatory architecture document (OAD)," Tech. Rep. TMT.SEN.DRD.05.002.CCR18, TMT (2007).
15. J.-P. Véran and L. Poyneer, "Evaluation of the T/T conditions at Gemini South using NICI AO telemetry data," in *1st AO4ELT Conference—Adaptive Optics for Extremely Large Telescopes*, Y. Clénet, J.-M. Conan, T. Fusco, and G. Rousset, eds. (EDP Sciences, 2009), 05002.
16. L. A. Poyneer and B. Macintosh, "Spatially filtered wavefront sensor for high-order adaptive optics," *J. Opt. Soc. Am. A* **21**, 810–819 (2004).
17. T. Fusco, C. Petit, G. Rousset, J.-M. Conan, and J.-L. Beuzit, "Closed-loop experimental validation of the spatially filtered Shack-Hartmann concept," *Opt. Lett.* **30**, 1255–1257 (2005).
18. F. J. Rigaut, J.-P. Véran, and O. Lai, "Analytical model for Shack-Hartmann-based adaptive optics systems," *Proc. SPIE* **3353**, 1038–1048 (1998).
19. J.-P. Véran, F. Rigaut, H. Maitre, and D. Rouan, "Estimation of the adaptive optics long-exposure point-spread function using control loop data," *J. Opt. Soc. Am. A* **14**, 3057–3069 (1997).
20. L. A. Poyneer, M. A. van Dam, and J.-P. Véran, "Experimental verification of the frozen flow atmospheric turbulence assumption with use of astronomical adaptive optics telemetry," *J. Opt. Soc. Am.* **26**, 833–846 (2009).
21. L. A. Poyneer, D. T. Gavel, and J. M. Brase, "Fast wavefront reconstruction in large adaptive optics systems with use of the Fourier transform," *J. Opt. Soc. Am. A* **19**, 2100–2111 (2002).
22. A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals and Systems*, 2nd ed. (Prentice Hall, 1997).
23. L. Jolissaint, J.-P. Véran, and R. Conan, "Analytical modeling of adaptive optics: foundations of the phase spatial power spectrum approach," *J. Opt. Soc. Am. A* **23**, 382–394 (2006).
24. L. A. Poyneer, A. Norton, and D. Dillon, "Open-loop shaping of a 4k mems with Fourier-domain pre-compensation," in *Adaptive Optics: Methods, Analysis and Applications*, OSA Technical Digest (CD) (Optical Society of America, 2009), paper JTUC3.
25. C. Petit, J.-M. Conan, C. Kulcsár, and H.-F. Raynaud, "Linear quadratic Gaussian control for adaptive optics and multiconjugate adaptive optics: experimental and numerical analysis," *J. Opt. Soc. Am. A* **26**, 1307–1325 (2009).