

Towards feasible and effective predictive wavefront control for adaptive optics

Lisa A. Poyneer^a and Jean-Pierre Véran^b

^aLawrence Livermore National Lab, 7000 East Ave, Livermore, CA, 94550

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

ABSTRACT

We have recently proposed Predictive Fourier Control, a computationally efficient and adaptive algorithm for predictive wavefront control that assumes frozen flow turbulence. We summarize refinements to the state-space model that allow operation with arbitrary computational delays and reduce the computational cost of solving for new control. We present initial atmospheric characterization using observations with Gemini North's Altair AO system. These observations, taken over 1 year, indicate that frozen flow exists, contains substantial power, and is strongly detected 94% of the time.

Keywords: Adaptive Optics, predictive control, wavefront control

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. There have been several specific proposals in this area recently, including the Strehl-optimal approach of Gavel and Wiberg,¹ the closed-loop Kalman filtering model of Le Roux et al.² (with further experimental work by Petit et al.³) and the data-driven \mathcal{H}_2 -optimal control method of Hinnen et al.⁴

In our own research, we have recently proposed Predictive Fourier Control (PFC).⁵ PFC builds upon the closed-loop Kalman filtering framework of Le Roux. In PFC the wavefront is reconstructed in the Fourier basis set.⁶ Under the assumption of frozen flow atmospheric turbulence, the Fourier modes are nearly uncorrelated both spatially and temporally. This allows each complex-valued Fourier mode to be controlled independently. This independent control makes the Kalman filter computationally tractable. 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.

PFC has several advantages over alternative proposals which make it feasible for use in both current general-purpose AO systems and future specialized high-performance AO systems such as the Gemini Planet Imager (GPI).⁷ First, it is computationally feasible for an AO control computer based on commercial, off-the-shelf technology. This computational efficiency derives from the fact that under frozen flow, the Fourier modes are controlled independently. Other basis sets, such as the actuators or the Zernike polynomials, are temporally and/or spatially correlated, resulting in large matrices.

The second advantage is that the algorithm explicitly measures atmospheric characteristics during closed-loop operation and modifies the predictive control in a block-adaptive fashion.⁸ Predictive control will only improve performance if its model is an accurate (though exactly how accurate remains to be quantified) characterization of atmospheric behavior. The approaches of Gavel & Wiberg and of Le Roux assume that the atmospheric parameters are known a priori or are approximated.

Send correspondence to Lisa Poyneer: poyneer1@llnl.gov, 1 925 423 3360

Despite PFC’s solid theoretical foundation and the positive attributes discussed above, our proposal still has a ways to go before becoming a feasible and effective predictive controller in an operation astronomical AO system. In this paper we discuss two issues in the progress from theory to reality. First, we summarize our recently published work on reformulating the Kalman state space model to deal with arbitrary computational delays. This reformulation not only enables operation with non-integer time-step delays, but also has significant benefit in the computational cost of solving for the new predictive control.

Second, we present initial results from our study of atmospheric conditions as observed with Gemini North’s Altair AO system. Our data set consists of 117 different nights of data, taken during M1 tuning over the course of April 2007 to April 2008. Analysis of this data reveals that frozen flow exists and is strongly detected by Altair 94% of the time. This is highly promising for the implementation of PFC, as that method assumes the frozen flow model.

2. FUNDAMENTAL CONCEPTS

PFC is a method which uses Kalman filtering to provide temporal control of the wavefront. The wavefront is first reconstructed with the technique of Fourier Transform wavefront Reconstruction (FTR).^{6,9} FTR can be viewed in two ways. First, FTR is a filtering technique, in which the slopes are converted to the frequency domain with an efficient Discrete Fourier Transform (e.g. FFTW¹⁰) and then filtered (the perspective taken in [9]). Second, FTR is also a modal reconstruction method, where the wavefront is reconstructed in the Fourier basis set (the perspective taken in [6]).

As mentioned above, the Fourier modes are particularly useful for predictive control because under the assumption of frozen flow turbulence, the Fourier coefficients are both spatially and temporally uncorrelated. This allows each mode to be temporally filtered independently of the others, providing computational efficiency. The Fourier modes are also very useful under frozen flow because the temporal PSD of a Fourier mode has a highly compact and easily modeled form. Each Fourier mode is described by spatial frequencies $f_x = k/(Nd)$ and $f_y = l/(Nd)$, where N is the grid size (48 for GPI) and d is the subaperture/actuator spacing in the pupil (18 cm for GPI). The frequencies variables k and l take the values $-N/2, -(N/2 - 1), \dots, -1, 0, 1, \dots, (N/2 - 2), (N/2 - 1)$.

A layer of frozen flow is assumed to follow a pure translation across the pupil with velocities v_x and v_y . This frozen flow causes a pure oscillation of a specific Fourier mode. The temporal frequency of the mode is the dot product of the frequency vector with the wave vector, asjusted for the AO frame rate f_{ao} and its inverse the time step $T = 1/f_{ao}$. For PFC, we model each layer for each Fourier mode independently as a first-order auto-regression (termed AR(1)) characterized by the complex number α .

$$a[t] = \alpha a[t - 1] + w[t]. \tag{1}$$

The process is driven by temporally white noise $w[t]$ with a specific variance. 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. 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: $-2\pi T(kv_x + lv_y)/(Nd)$.

When we have many layers of frozen flow, they simply add together. Under this model, the temporal PSD of a complex-valued Fourier mode will have multiple separate peaks, one for each layer at the specified frequencies. Because the modes are complex-valued, the temporal PSD will be asymmetric, so both positive and negative frequencies are possible.

This model is confirmed by our AO simulation. Fig. 1 shows a temporal PSD for a specific Fourier mode when the simulation is run open loop (but with tip-tilt controlled) on a five layer frozen flow atmosphere. The AR(1) model fits this behavior well. Because new phase content does appear as the phase screen translates across the pupil, the peaks have a width, and there is some very low temporal frequency content. This concentrated PSD means that it is easy to identify layers in the atmosphere - the peaks simply need to be located and the corresponding temporal frequencies identified. This is in contrast to other modal basis sets, as is illustrated in the figure on the right side. For the same data, the Zernike modes and actuators values were also analyzed in terms of temporal PSD. These temporal PSDs follow power laws (with different asymptotes) and do not allow easy identification of atmospheric wind parameters.




Figure 1. The temporal PSD of [left] a Fourier mode and [right] either a low-order Zernike coefficient or an actuator. Five-layer frozen flow atmosphere model used in simulation. The five layers each produce distinct peaks in the Fourier mode’s PSD. This allows easy detection. In contrast, the temporal PSDs in other basis sets do not exhibit the same concentration of power or allow easy identification.

3. ORIGINAL AND REFORMULATED MODEL

The original PFC state space model was based on the closed-loop Kalman filter state space model of Le Roux. That model assumed that the computational delay from the end of a WFS integration period to the write of the new DM commands was exactly one time step (equal to a WFS integration interval). For each Fourier mode, a complex-valued state space model was constructed. The state variables were the AR(1) values for each layer (the total number of which is arbitrary), and the phase and DM commands at present and past time steps. In this way, the DM commands were subtracted off the phase signal to produce the actual closed-loop measurement which the AO system uses.

Given the state space model (see [5] for all our equations), the Algebraic Riccati Equation (ARE) is solved in terms of the steady-state error covariance matrix entries. These entries form the coefficients of the temporal filter which predicts the layers of frozen flow atmosphere for that Fourier mode. The specific layer parameters (power levels, temporal frequencies) are estimated from closed-loop AO telemetry. The predictive filter is composed of parallel integral controllers, one for each layer and one for the low-temporal frequency (or “DC”) content. Each layer integrator predicts the residual wavefront measurement two time steps, and predicts the previous DM command by one time step.

Given this structure, the PFC temporal filters are highly dependent on the actual computational delay in the AO system. Though the original model assumed a one time-step delay, in actual implementation this will not always be the case. As an example, at the Critical Design phase the WFS integration time for GPI is 667 μsec and the computation delay is 1000 μsec . This makes the computational delay not 1 time step but 1.5. Since prediction by its very nature depends on the lag between measurement and correction, running GPI with a predictive controller for a 1 time step lag would have potentially poor, or maybe even unstable, performance.

To solve this problem, we have reformulated the state space model. This reformulation is described in detail in our latest paper.¹¹ In this reformulation we do two important things. First, we model the asynchronous DM writes and WFS integrations as a linear combination. If the DM changes shape during the middle of a WFS integration, the state space model does a linear weighting of the DM commands depending on the exact timing. Furthermore, the model assumes that in deriving the best new DM commands, which will be asynchronous, the two intervals where the phase changes are also linearly weighted.

The second important change is removing redundant states. The original model kept present and past phase values, as well as past DM values. The DM values can be moved out of the state into another vector, and only one phase value (in addition to the AR(1) variables for the layers) need be kept. The current algorithm that is used to solve the ARE has computational cost is cubic in the number of state variables. For a typical number of layers, the new model has half as many states as the old, which reduces the ARE solve cost by a factor of eight.

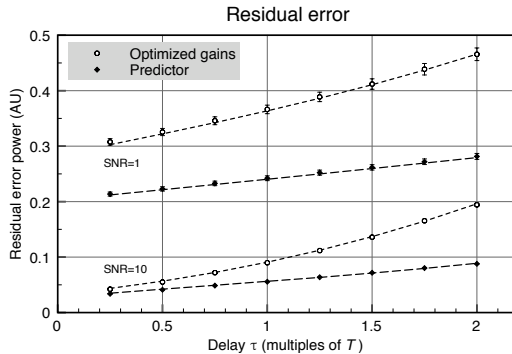


Figure 2. Residual error power (input phase power = 1) for two different SNRs, predictive filter compared to optimized-gain integral controller for a single complex-valued Fourier mode under three-layer frozen flow. The predictor smoothly transitions for arbitrary delays and provides an increasing advantage over the integral controller for longer delays.

The new PFC filter structure is very similar to the original one. It is still composed of parallel layer integrators. Now each integrator predicts the residual wavefront measurement the appropriate (fractional) number of time steps. This is done by a linear weighting of the predictions for integer numbers. The final lead filter which stabilizes PFC also adjusts to the exact computational delay.

This new model and temporal filter assures graceful behavior in terms of both stability and residual error performance as the computational delay is varied. Fig. 2 shows results from both theory and our Matlab Simulink AO model (see [11] for details) for the residual error power when correcting a three-layer atmosphere for a single Fourier mode. The predictor performance is directly compared to an optimized-gain integral controller. As the computational lag increases, so does the benefit of prediction. Because the predictor can selectively notch out layers at higher temporal frequencies, it does not suffer in the same manner as the integral controller when the delay is increased and the bandwidth is reduced.

In our original proposal,⁵ we built upon a solid foundation of Kalman filtering to develop an adaptive, computationally efficient predictive controller. We have further refined¹¹ the algorithm to deal with arbitrary system delays and reduced the cost of solving for the new control. All of our performance tests in simulation, however, have assumed a pure frozen flow atmosphere. Since our simulated phase aberrations conform to our model assumptions, the predictive controller has been highly effective.

Whether or not PFC will be as effective in a real astronomical AO system depends entirely on whether the atmosphere actually exhibits frozen flow. Significant power must be present in the layers for prediction to make a substantial difference. Furthermore, the layer velocities must be stable on the time scales of tens of seconds so that the AO system can measure and adapt to atmospheric conditions before they change. In the next section we present initial results of our astronomical AO telemetry study.

4. ASTRONOMICAL AO TELEMETRY STUDY

To study atmospheric behavior, we have acquired and analyzed AO telemetry data from Gemini North’s Altair AO system. The data are closed-loop DM commands during M1-tuning operations. At the beginning of the night, the AO loop is closed on a bright star and the persistent aberrations on the Altair DM are offloaded to the M1 actuators, so that the static wave-front delivered by the telescope is as flat as possible. We have 117 different 1-minute intervals of data, spread over nights from April 2007 to April 2008.

Though Altair is a relatively low-order AO system (12 subapertures across the 8 m pupil) compared to GPI, and though it uses a matrix-based modal controller instead of Fourier reconstruction, we can still use the AO telemetry to study the atmosphere. First we discuss the data analysis techniques used to obtain estimates of open-loop temporal PSDs of Fourier modes from Altair telemetry. Then we examine these temporal PSDs to look for frozen flow. We do so over both the full 1-minute interval of telemetry, and over 10-second subintervals, as we plan for GPI to update atmospheric parameter estimates and predictive filters every 5 to 10 seconds.

4.1. Data analysis techniques

Though Altair does not use Fourier transform wavefront reconstruction, it is easy to analyze the data as if it did. The Fourier modal coefficients of the phase are simply determined by taking the Discrete Fourier Transform (DFT) of the phase signal embedded in a square grid of size $N \times N$.

As such, the first major stage of the telemetry analysis pipeline is to convert the telemetry from the specific format provided by the AO system into a time series of Fourier modal coefficients. First, the telemetry vector is mapped into a square grid, producing a two dimensional mapping of the pupil. This is padded into an appropriately sized grid as if for FTR. For Altair data, $N = 16$. Any uncontrolled phase points, such as those from slaved or extrapolated actuators, are discarded and set to zero. Based on provided system information with the telemetry, the data units are converted to nm of phase error. If it has not already been done, piston, tip and tilt are removed from the phase. Finally, the Fourier modal coefficients are calculated, with a normalization constant to preserve total power. For a phase signal $\phi[m, n]$ and the aperture $a[m, n]$, the modal coefficients are

$$\Phi[k, l] = \left(\frac{1}{\sum_{m=0}^{N-1} \sum_{n=0}^{N-1} a[m, n]} \right)^{1/2} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} \phi[m, n] a[m, n] \exp(-j2\pi[mk + nl]/N) \quad (2)$$

These steps produce a properly normalized data cube of Fourier modal coefficients for all time steps in the observation. At any time step, the total power of the Fourier modes (just the sum of their magnitude-squares) is equal to the MSE of the phase conjugated by the DM in the pupil.

The second major stage of the telemetry analysis is evaluating the temporal PSDs of the Fourier modes and correctly compensating for the effects of the control loop and DM. Each temporal PSD estimate is done on segments of length S . The system frame rate f_{ao} sets the maximum estimated temporal frequency of $f_{ao}/2$. The length of the segment S sets the sampling in temporal frequency at f_{ao}/S , and hence the minimum frequency greater than zero which is estimated. For Altair $f_{ao} = 1000$ Hz and $S = 2048$. The temporal PSD will be estimated at multiples of 0.488 Hz. Many segments of length S samples are used to estimate the temporal PSD. This sets the minimum length of time over which a temporal PSD estimate can be calculated. For Altair this is around 10 seconds. This allows the telemetry from the experiment to be examined at either 10-second or 1-minute intervals to characterize the atmosphere.

Given the interval, the first step in PSD estimation is to remove the mean value of the Fourier mode from the time series. This is a standard technique (see [12]) that allows more accurate PSD estimation. This mean-removal subtracts the estimated static error from the signal over the interval. For each Fourier mode, the temporal PSD is then estimated using the averaged modified periodogram technique (again, see [12] for a thorough discussion). The Hanning window $h[t]$ is used on half-overlapped intervals for best SNR. In particular, the estimated temporal PSD is

$$\hat{P}(\omega) = \frac{1}{N_{\text{intervals}}} \sum_{i=0}^{N_{\text{intervals}}} \hat{P}_i(\omega), \quad (3)$$

where

$$\hat{P}_i(\omega) = \frac{1}{S} \frac{1}{\sum_{t=0}^{S-1} h[t]^2} \left| \sum_{t=0}^{S-1} \Phi_i[k, l, t] h[t] \exp(-j\omega t) \right|^2, \quad (4)$$

and $\Phi_i[k, l, t]$ is the i -th interval. The discrete-time frequency ω is converted to the continuous-time frequency f with the equation $\omega = 2\pi f/f_{ao}$. The temporal PSD is normalized by the power in the window such that the sum of the temporal PSD equals the variance of that modal coefficient. At this point the sum of the entire data cube is the time-averaged MSE in the aperture, excluding the static error component.

If the telemetry is pre-compensated for the influence function of the DM, this compensation must be removed. This can only be done with knowledge of the influence function. For Altair a numerical model was used. Generating an influence function with appropriate sampling and grid size allows conversion to a frequency response with just a DFT. For each mode, the temporal PSD is scaled by the magnitude-squared of the transfer function for that mode.

Finally, the response of the control system must be compensated for. Our treatment is based on the standard continuous/discrete-time hybrid system model (see Madec’s chapter in [13]). The goal is to estimate the continuous-time temporal PSD of the atmosphere (given by ϕ) free of influence of the controller. The WFS transfer function is $H_{\text{wfs}}(s) = [1 - \exp(-sT)]/(sT)$, where $T = 1/f_{\text{ao}}$. The DM has the same sample and hold $H_{\text{wfs}}(s) = H_{\text{dm}}(s)$. The computational delay is $H_{\text{delay}}(s) = \exp(-s\tau)$, where τ is the controller delay in seconds. The specific control law $C(z)$ depends on the AO system and is provided with the telemetry. For use with the discrete-time transfer function and discrete-time sampled PSDs, $z = \exp(-sT)$ and $s = j\omega/f_{\text{ao}}$. For Altair, the DM commands $d[t]$ are used; they have temporal PSD

$$\hat{P}(\omega) = \left| \frac{C(z)}{1 + H_{\text{wfs}}(s)H_{\text{dm}}(s)H_{\text{delay}}(s)C(z)} \right|^2 [P_\phi(\omega)|H_{\text{wfs}}(s)|^2 + P_v(\omega)] \quad (5)$$

The transfer function in front is divided out to obtain an open-loop estimate from the DM commands.

4.2. Temporal PSD structure and power levels

As discussed above, PFC uses an AR(1) model for each layer of atmosphere. When many layers are present, they simply add together. PFC assumes that an arbitrary number of layers exist, along with a static error term. The static error is represented by a virtual layer in which α is a real number just less than one. To describe an arbitrary atmosphere of L layers, $L + 1$ AR(1) processes are described by their auto-regression parameters $\alpha_0, \alpha_1, \dots, \alpha_L$ and the variances of the driving noises $\sigma_{\alpha_0}^2, \sigma_{\alpha_1}^2, \dots, \sigma_{\alpha_L}^2$. The temporal PSD is simply

$$P_\phi(\omega) = \sum_{i=0}^L \frac{\sigma_{\alpha_i}^2}{|1 - \alpha_i \exp(-j\omega)|^2}, \quad (6)$$

because the layers are uncorrelated with each other. For our Altair analysis, we assume that $\alpha_0 = 0.995$, to provide excellent low-temporal frequency correction. For the layers, we use $|\alpha| = 0.99$.

This temporal PSD consists of a series of peaks. After estimation of $\hat{P}_\phi(\omega)$ from data (as described above) the PSD must be search for these peaks. This stage is essentially finding the model parameters which produce a best-fit to the PSD. Since our initial proposal (see the detailed description in [9]) we have refined the peak-finding process is two ways. First, we have added in an additional slowly-varying term to capture a not-unusual form of atmospheric power seen in the PSDs. This is an additional real-valued α term which has an auto-regression of just 0.8. This produces a broad, flat atmospheric term which extends past 100 Hz. This is seen in Altair data on some occasions. Secondly, we have refined the search-and-project steps. The modified algorithm is as follows:

1. Set range of PSD to search (usually -100 to 100 Hz; will be higher for GPI)
2. Project $\hat{P}_\phi(\omega)$ onto the $\alpha_0 = 0.995$ PSD shape to estimate σ_0^2 .
3. Subtract DC-fit PSD term from data; set to zero any data points with negative values.
4. Project $\hat{P}_\phi(\omega)$ onto the $\alpha_0 = 0.8$ PSD shape, but only in the $|50|$ to $|100|$ Hz ranges. Use this to estimate the power for this mild auto-regression.
5. Subtract this newly fit PSD term from data; set to zero any data points with negative values.
6. Loop through the remaining steps until no significant peaks are found, or the maximum number of peaks (4) is found.
7. Correlate the modified $\hat{P}_\phi(\omega)$ with a preset shape. This correlation can be done in the Fourier domain. The preset shape is the magnitude-squared of the Fourier transform of the Hanning window. This sets the minimum feature size for detecting a temporal frequency component.
8. Locate the maximum of the correlation function using parabolic interpolation on the three points centered on the maximum value. If the maximum value is greater than the significance threshold, proceed. Else, end the search. This is the estimate of the temporal frequency of the peak and is used to generate the complex valued α .




Figure 3. PSD fit for a mid-range Fourier mode. [Left] Data-based PSD (black) and complete fit (red). [Right] The DC (blue), Wing (green) and layer (red) terms of the total fit.

9. Construct the PSD shape (see Eq. 6) for this α_i
10. Estimate the power level σ_i^2 via projection with the PSD shape.
11. Subtract layer-fit PSD term from data; set to zero any data points with negative values.

At the end of this process, a set of layer α 's and power levels are obtained for each Fourier mode. As an example of this, see Fig. 3. On the left side of the figure, the data-based PSD estimate is shown, along with the composite PSD fit. On the right the three parts of the fit are shown: the $\alpha = 0.995$ DC term (blue), which is strongly peaked at 0 Hz; the broad, low $\alpha = 0.8$ Wing term (green), which extends out past 100 Hz, and the layers (red). Normalized to the noise variance, the layer frequencies and variances are as follows: For the DC layer with $\alpha = 0.995$, the variance is 1.15. When converted to the PSD, this is 36.1% of the total atmospheric power. The weak auto-regression layer with $\alpha = 0.8$ was variance 27.0, which translates to 23.5% of the total atmospheric power. (Though the variance is high, the total power in the PSD is lower because the PSD is rather flat.) Four peaks were found. Together, they comprise 40.4% of the total atmospheric power in the fit. All have $|\alpha| = 0.99$. The layer frequencies (Hz) and variances are, in order of discovery: 3.59 Hz and 0.97; -13.9 Hz and 0.96; -17.5 Hz and 0.38; and -10.4 Hz and 0.25. The first peak is alone on the right, while the three other peaks form a larger peak on the left.

As shown above, the PSD fit process provides a way to estimate how much total power in is the DC-term and the layers. The DC-term represents the slowly varying error which the integral controller is designed to correct. In Fig. 4 we show the amount of power in the peaks (assumed to be in the layers) for both all spatial frequencies, and the mid-range spatial frequencies $k = 5, l = 5$ and $k = -5, l = 5$. For all spatial frequencies, the median power in the peaks is 30%. Most of the observations have between 20% and 60%. For the specific mid-range Fourier modes, the median is 43%, with a much more flat and broad distribution from 10% to 80%. The amount of power in the total phase is less than in a specific mid-range Fourier mode for a combination of two reasons. First, higher spatial frequencies have higher layer frequencies. Low-frequency Fourier mode layer components are very close to 0 Hz, and as such our calculation of power in the peaks will be very low. Second, the Kolmogorov spatial power spectrum places more power in low frequencies, so the integral over all Fourier modes is biased toward the lower percentages of the low-order modes.

These estimates are not completely accurate because aliasing does exist in Altair. No attempt is made to remove this spatial aliasing from the PSDs. For example, a spatial frequency beyond the Nyquist limit of Altair will alias into a controllable mode. When frozen flow exists, this will cause a peak to appear in the temporal PSD, but this peak is not a true atmospheric aberration. Though we can in some cases identify aliases, we have not developed a robust way to remove all aliased power from the atmospheric estimates. We do not think the inaccuracy will be severe, however. This is because slowly-varying errors in the uncontrollable spatial frequencies will also alias in to the DC-term, offsetting any bias.




Figure 4. Histograms of the percentage of total dynamic atmospheric power which is found in the peaks (assumed to be from frozen flow). [Left]: total power in all controllable Fourier modes; [Right]: power for mid-range Fourier modes $k = 5, l = 5$ and $k = -5, l = 5$.

Altair’s actuator spacing is 3.6 times wider than GPI’s. Given this, GPI will control much higher spatial frequencies than Altair. We expect that for GPI the amount of power in the layers will be higher, because the many more higher spatial frequency modes will be present and these modes should have a higher percentage of their power in the layers.

4.3. Confirming layer existence

In the above analysis, we have assumed that all asymmetric peaks in the temporal PSDs are caused by frozen flow. It is, however, possible that some other source may have caused the peaks. For example, a specific (but symmetric) peak pair might be caused by vibration. The PFC algorithm will identify this and correct it, regardless of its original source.

The underlying concept to identify frozen flow is simple: if a layer of frozen flow exists with velocity vector $\langle v_x, v_y \rangle$ (m/s) exists, the temporal frequency of the peak it causes for Fourier mode with frequency vector $\langle f_x, f_y \rangle$ (m^{-1}) is just the dot product of the velocity vector with the frequency vector: $v_x f_x + v_y f_y$. This means that when the Fourier modes are viewed on a f_x, f_y grid, the temporal frequencies caused by that layer produce a plane, with piston ($f_x = f_y = 0$) always having a peak at 0 Hz.

Because each Fourier mode has its peaks found independently, the PSD-fit process leaves us with a datacube of temporal frequencies. Some temporal frequencies may correspond to peaks which don’t really exist, due to noise in the temporal PSD estimate. Others may be due to vibration, aliasing (as discussed above) or other unknown sources.

Our method for determining whether a layer of frozen flow exists is a straight-forward but brute-force algorithm. We begin with a domain of possible velocity vectors that are evenly spaced on a square grid. For the Altair data, we searched from -60 to 60 m/s with resolution 0.5 m/s to generate a 241 by 241 grid. Then, for each velocity vector, we calculate the frequency plane $v_x f_x + v_y f_y$. Next, we search over each spatial frequency $\langle f_x, f_y \rangle$ where $|v_x f_x + v_y f_y| > 2$. The peak-finding algorithm cannot find peaks at very low temporal frequencies, so the layer is “visible” only to a subset of spatial frequencies. For each Fourier mode which can “see” that layer, we search the set of found peaks for a peak that is close enough to the desired temporal frequency $v_x f_x + v_y f_y$. We define “close enough” as being within ± 0.75 Hz or 15% of the true value, whichever is less. This means for peak frequencies under 5 Hz, the found peak must be very close to the true value, which reduces false detections.

So for a given layer velocity, all Fourier modes that can “see” that layer are searched to see if they detected a peak that was close enough. The likelihood of a layer is defined as the total number of modes which did find a peak close enough divided by the number of modes which can “see” the layer. This metric is calculated for each velocity vector in the domain, produce a map of the likelihood of there being a layer. This algorithm is robust to false positives. Random peaks, vibrations and aliases show up with likelihoods under 5%. A sample wind map (for the same observation as the PSDs shown in Fig. 3) is given in Fig. 5. This wind map shows that the




Figure 5. Windmap of layer velocity likelihoods. Range -60 to +60 m/s, color scale from 0 (black) to 1 (white). Three layers are clearly identifiable.




Figure 6. Temporal frequency planes for Fourier modes, as induced by the three strongly identified layers. [Left]: known frequencies from identified layers. [Right]: Peaks found. Note that where the frequency is too low (orange regions) the Fourier modes will not identify it.

peaks which were found cluster strongly in three locations in velocity space, indicating three layers. Proceeding in order of likelihood, the three layers have the following parameters. The most strongly detected layer has speed 13.3 m/s and angle 66° , with likelihood 73%. This layer is compact. The next layer is more diffuse; its center has speed 26.4 m/s, angle -6° and likelihood 65%. The weakest layer has speed 8.2 m/s and angle 99° , with likelihood 51%. The likelihoods track back directly into how many Fourier modes identify and correct each layer. As discussed above, the temporal frequencies induced by a layer form a plane. Given a datacube of peak temporal frequencies, we can compose them into matching planes which show where and how many modes found each layer. For the three-layer case given above, Fig. 6 shows the ideal temporal frequency planes for the three layers, and the detected planes. Examination of these frequency planes shows that these wind layers are very clearly found in the data.

In the case shown above it is easy to “eyeball” the wind map and identify the three layers. To automate this process, and to be robust in situations where layers are located close together in velocity space, we have applied the Watershed algorithm (see [14] for a discussion of this well-established algorithm). This algorithm segments the wind map and identifies layers at the maximum inside each region. We have set a threshold for a layer being “strongly detected” by the algorithm at 50% likelihood for the 1-minute observations and 45% for the 10-second observations. (This 10-second PSD estimates, and hence peak temporal frequency estimates, are more noisy, hence the slightly lower threshold. This SNR issue is one which we plan to further explore.) In some cases a layer is clearly identifiable by eye, but its maximum likelihood is under the threshold. In those cases we do not classify that layer as strongly detected, though it may be readily apparent.

The wind map and the corresponding list of strongly detected layers are highly useful. First, we can characterize the the composition of the atmosphere in terms of number of layers, wind velocities, etc. First we examine how often at least one layer is strongly detected. Fig. 7 shows the histograms for the strength of the most likely




Figure 7. Histograms of likelihood of the most likely layer found in each wind map. [Left]: 1-minute observations, [Right]: 10-second observations.

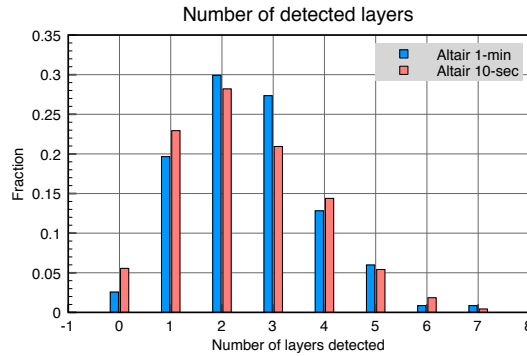


Figure 8. Histograms of the number of strongly detected layers.

layer for all observations. For the 1-minute observations, the median likelihood is 80%. In only 2 cases was a layer not strongly detected, meaning 98% of the time at least one layer was. For the 10-second observations, the histogram has moved slightly towards lower likelihoods. With the 45% likelihood threshold for strong detection, 94% of the time at least one layer was strongly detected. Given this, we see (through the PFC peak-finding process) clear evidence of frozen flow nearly all the time.

The number of layers that are strongly detected varies from night to night. Fig. 8 shows the histograms for the number of strongly detected layers. Most observations have 1 to 4 layers. The cases with 5 or more detected layers tend to have a larger region of diffuse, but still detectable, likelihoods, as opposed to a concentrated region. The watershed algorithm finds multiple layers in these diffuse regions.

5. CONCLUSIONS

We have summarized our latest theoretical developments in modify Predictive Fourier Control (PFC) to deal with arbitrary computational delays. This not only enables prediction with non-integer time step delays, but also substantially reduces the cost of solving for new predictive filters.

The main portion of this paper presented initial results from a study of 117 nights of Altair telemetry to look for frozen flow. Using data processing techniques to evaluate Altair circular buffers as if they had been taken with a Fourier reconstruction algorithm, we examine the temporal PSDs of the Fourier modes. This study reveals that a median of 30% of total atmospheric power is due to frozen flow, across all of Altair’s controllable spatial frequencies. Mid-range Altair spatial frequencies have more power, with a median of 43% and a spread of 10% to 80%. Because substantial power exists in frozen flow, the Kalman filter should be effective in improving performance. By examining the temporal frequencies of peaks across all Fourier modes, we can assemble a

likelihood map of wind velocities to identify layers of turbulence. Using this method, between 94% (10-second intervals) and 98% (1-minute intervals) of the observations had at least one layer of strongly detected frozen flow. Most observations had between one and four layers.

This initial study is very promising for PFC. Frozen flow exists nearly all of the time at Mauna Kea and is easily detectable by Altair. Much further work remains, including in generating reliable estimates of how much PFC will improve performance. We will also further refine (if necessary) the model to improve the model fit to the data, and study temporal variability to assess how best to adapt to changing layers in closed loop.

ACKNOWLEDGMENTS

We thank the Gemini Observatory staff for their work in obtaining Altair circular buffers for our analysis. This work 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-CONF-404559. 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. D. T. Gavel and D. Wiberg, "Towards strehl-optimizing adaptive optics controllers," in *Adaptive Optical System Technologies II*, P. L. Wizinowich and D. Bonaccini, eds., *Proc. SPIE* **4839**, pp. 890–901, 2002.
2. 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**, pp. 1261–1276, 2004.
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. Exp.* **16**, pp. 87–97, 2008.
4. K. Hinnen, M. Verhagen, and N. Doelman, "Exploiting the spatiotemporal correlation in adaptive optics using data-driven \mathcal{H}_2 -optimal control," *J. Opt. Soc. Am. A* **24**, pp. 1714–1725, 2007.
5. 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**, pp. 2645–2660, 2007.
6. L. A. Poyneer and J.-P. Véran, "Optimal modal Fourier transform wave-front control," *J. Opt. Soc. Am. A* **22**, pp. 1515–1526, 2005.
7. B. A. Macintosh, "Gemini planet imager preliminary design document volume 2: Instrument design," tech. rep., Lawrence Livermore National Lab submitted to International Gemini Project Office, 2007.
8. G. Clark, S. Parker, and S. Mitra, "A unified approach to time- and frequency-domain realization of fir adaptive digital filters," *IEEE Trans. Acoustics, Speech and Signal Processing* **31**, pp. 1073–1083, 1983.
9. L. A. Poyneer, D. T. Gavel, and J. M. Brase, "Fast wave-front reconstruction in large adaptive optics systems with use of the Fourier transform," *J. Opt. Soc. Am. A* **19**, pp. 2100–2111, 2002.
10. M. Frigo and S. G. Johnson, "The design and implementation of FFTW3," *Proceedings of the IEEE* **93**(2), pp. 216–231, 2005. special issue on "Program Generation, Optimization, and Platform Adaptation".
11. L. A. Poyneer and J.-P. Véran, "Predictive wavefront control for adaptive optics with arbitrary control loop delays," *J. Opt. Soc. Am. A* **25**, pp. 1486–1496, 2008.
12. A. V. Oppenheim, R. W. Schaffer, and J. R. Buck, *Discrete-time Signal Processing*, Prentice Hall, New Jersey, 1999.
13. F. Roddier, *Adaptive Optics in Astronomy*, Cambridge University Press, Cambridge, 1999.
14. R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, Prentice Hall, New Jersey, 2002.