

# UC Santa Cruz

## UC Santa Cruz Previously Published Works

### Title

Laboratory demonstration of optimal identification and control of tip-tilt systems

### Permalink

<https://escholarship.org/uc/item/827319f9>

### ISBN

978-1-5106-5351-1

### Authors

Sengupta, Aditya R  
Gerard, Benjamin L  
Dillon, Daren  
[et al.](#)

### Publication Date

2022-08-30

### DOI

10.1117/12.2630765

Peer reviewed

# Laboratory Demonstration of Optimal Identification and Control of Tip-Tilt Systems

Aditya R. Sengupta<sup>a,b</sup>, Benjamin L. Gerard<sup>a</sup>, Daren Dillon<sup>a</sup>, Maaïke van Kooten<sup>a</sup>, Donald Gavel<sup>a</sup>, and Rebecca Jensen-Clem<sup>a</sup>

<sup>a</sup>Department of Astronomy and Astrophysics, University of California, Santa Cruz, United States

<sup>b</sup>Department of Applied Mathematics and Theoretical Physics, University of Cambridge, United Kingdom

## ABSTRACT

We present the results of testing optimal linear-quadratic-Gaussian (LQG) control for tip and tilt Zernike wavefront modes on the SEAL (Santa Cruz Extreme AO Lab) testbed. The controller employs a physics model conditioned by the expected tip/tilt power spectrum and vibration peaks. The model builds on similar implementations, such as that of the Gemini Planet Imager, by considering the effects of loop delays and the response of the control hardware. Tests are being performed on SEAL using the Fast Atmospheric Self-coherent camera Technique (FAST), and being executed using a custom Python library to align optics, generate interaction matrices, and perform real-time control by combining controllers with simulated disturbance signals to be corrected. We have carried out open-loop data collection, characterizing the natural bench dynamics, and have shown a reduction in RMS wavefront error due to integrator control and LQG control.

**Keywords:** tip-tilt, linear quadratic Gaussian, system identification, laboratory, experiment

## 1. INTRODUCTION

Adaptive optics (AO) systems work to correct wavefront errors in images captured by optical and near infrared telescopes, via a control loop that measures deviations with a wavefront sensor (WFS) and corrects them by applying commands to a deformable mirror (DM) in the loop. AO has been essential to recent developments in astronomy, from ground-based direct imaging and spectroscopy of exoplanets, to precision astrometry.<sup>1</sup> Within AO systems, correcting the wavefront tip and tilt modes at the plane of the DM, which corresponds to  $x - y$  position deviations for the downstream focal plane point spread function, is critical in enabling high-quality AO performance, and can provide substantial benefits if optimized.

A major example of the importance of tip-tilt correction is the Keck Planet Imager and Characterizer (KPIC),<sup>2</sup> which aims to carry out direct imaging and spectroscopy of young exoplanets. It aims to couple light from the planet into an optical fiber so that the light can be transported to an extremely stable spectrograph, producing a high-resolution spectrum. This requires that the point-spread-function (PSF) of the planet be centered on the fiber tip. Any movement in tip-tilt will move the PSF off the fiber, resulting in a reduced amount of planet light being transmitted through the fiber and reducing the resolution of spectra. This process is especially sensitive to high-frequency motion in tip-tilt.<sup>3</sup> A control algorithm for tip-tilt that minimizes motion while damping high-frequency disturbance components is therefore critical for these scientific applications.

Optimal control solutions for tip-tilt generally involve using a recent history of noisy WFS measurements to predict and control for future states. In particular, the case of linear-quadratic-Gaussian (LQG) control<sup>4-7</sup> has been extensively studied. Assuming stationarity of turbulence statistics, linear dynamics perturbed by independent and identically distributed (i.i.d.) Gaussian process noise, and measurements perturbed by i.i.d. Gaussian measurement noise, the LQG controller provides the optimal (minimum-variance) sequence of control inputs. Another solution is predictive wavefront control,<sup>8</sup> which involves using a recent history of noisy WFS

---

Further author information: (Send correspondence to A.S.)

A.S.: E-mail: adityars@ucsc.edu, Telephone: 1 510 697 6549

measurements to predict and control for future states without a particular model for the driving process. LQG has the advantage of being computationally lightweight, involving only two relatively small matrix multiplications per time step.

In this paper, we present the results of experiments run on the UCSC SEAL (Santa Cruz Extreme AO Laboratory) testbed implementing LQG control for tip-tilt. We make use of the FAST (Fast Atmospheric Self-coherent camera Technique) method,<sup>9</sup> which is a focal plane WFS designed for exoplanet imaging science that enables tip/tilt measurement and correction with the coronagraphic science image (see appendix A for a review of FAST), and demonstrate via timing experiments that the software and hardware setup used is robust with arbitrarily computationally intensive control laws, and with a hardware-limited loop rate of 100 Hz. We discuss the mathematical foundation of LQG control adapted to this case, including system identification of disturbances and controller dynamics, and demonstrate improved control residuals as a result of this identification.

In section 2, we describe the mathematics behind LQG control for tip-tilt and present simulation results. In section 3, we describe the software and hardware setup, calibration steps, and characterization of the control loop timings for an arbitrary control law. In section 4, we present and discuss the results of bench testing, and discuss future work.

## 2. METHODS

### 2.1 Linear-quadratic-Gaussian control

Optimal control refers to a strategy for applying control inputs to a dynamic system such that its state minimizes a cost function. In the case of adaptive optics, we consider linear discrete-time dynamic systems with states  $\mathbf{x}_k$ , measurements  $\mathbf{y}_k$ , and with the process and measurement perturbed by i.i.d. Gaussian noise  $\mathbf{w}_k \sim \mathcal{N}(0, W)$ ,  $\mathbf{v}_k \sim \mathcal{N}(0, V)$  respectively, where  $W, V$  are known covariance matrices.

$$\begin{aligned}\mathbf{x}_{k+1} &= A\mathbf{x}_k + B\mathbf{u}_k + \mathbf{w}_k \\ \mathbf{y}_k &= C\mathbf{x}_k + D\mathbf{u}_k + \mathbf{v}_k\end{aligned}\tag{1}$$

LQG control depends on solving for the optimal observer (estimating the state given the noisy measurements) and controller (computing inputs based on the estimated state). These optimal expressions are derived in Appendix B. Both the observer and controller problems come down to solving a certain quadratic matrix equation, known as the discrete algebraic Riccati equation (DARE). We numerically solve the DARE with a series of methods, since it can be ill-conditioned with certain models and having multiple methods makes the implementation more reliable. We first use the solver in the Slycot<sup>10</sup> Python wrapper to the SLICOT numerical control library,<sup>11</sup> then the `scipy.linalg.solve_discrete_are` method,<sup>12</sup> and finally numerical iteration, in which we repeatedly set  $P$  equal to the right-hand side of 14, and iterate until we reach a fixed point. We denote the solution to the DARE by  $P = \text{dare}(\dots)$ , where the arguments are functions of the model matrices shown in Table 1.

The solvability of the DARE results in the *Kalman filter* equations for the observer,

$$P_{\text{obs}} = \text{dare}(A^\top, C^\top, W, V, \mathbf{0})\tag{2a}$$

$$K = P_{\text{obs}}C^\top(CP_{\text{obs}}C^\top + V)^{-1}\tag{2b}$$

$$\tilde{\mathbf{x}}_{k+1} = A\tilde{\mathbf{x}}_k + B\mathbf{u}_k\tag{2c}$$

$$\hat{\mathbf{x}}_{k+1} = \tilde{\mathbf{x}}_{k+1} + K(\mathbf{y}_k - C\tilde{\mathbf{x}}_{k+1}),\tag{2d}$$

and the *linear-quadratic regulator* (LQR) equations for the controller:

$$P_{\text{con}} = \text{dare}(A, B, C^\top C, D^\top D, D^\top C)\tag{3a}$$

$$L = -(D^\top D + B^\top P_{\text{con}}B)^{-1}(D^\top C + B^\top P_{\text{con}}A)\tag{3b}$$

$$\mathbf{u}_k = L\mathbf{x}_k.\tag{3c}$$

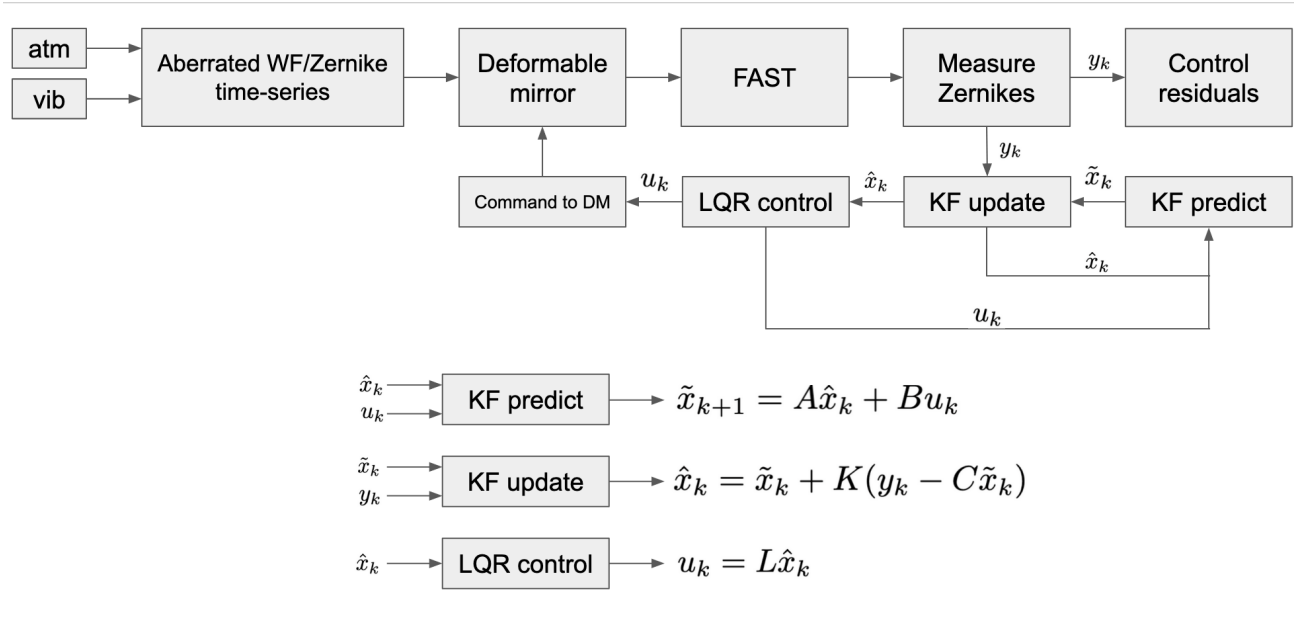


Figure 1. LQG control in an adaptive optics loop. The forward arm consists of the plant (the DM and FAST), and the backward arm contains the linear algebra steps required to implement LQG control in real time, given known values of  $A, B, C, K, L$ .

Matrix	Purpose	Dimension
$A$	State-to-state time evolution	$s \times s$
$B$	Input-to-state time evolution	$s \times p$
$C$	State-to-measurement translation	$m \times s$
$D$	Input-to-measurement translation	$m \times p$
$W$	Process noise covariance	$s \times s$
$V$	Measurement noise covariance	$m \times m$
$Q$	State weighting in cost (not tuned here)	$s \times s$
$R$	Input weighting in cost (not tuned here)	$p \times p$
$S$	State-input correlated cost (not tuned here)	$p \times s$
$P_{\text{obs}}$	Covariance of Kalman filter state estimate	$s \times s$
$P_{\text{con}}$	Inner product kernel for optimal control cost	$s \times s$
$K$	Kalman gain: optimal correction for state estimates	$s \times m$
$L$	LQR gain: optimal function of state to use as input	$p \times s$

Table 1. Summary of matrices in LQG control, with state size  $s$ , measurement size  $m$ , input size  $p$ .

A key benefit of LQG control is that the gains  $K$  and  $L$  can be precomputed, meaning that after the model identification step, the only steps that need to be computed per timestep are the Kalman prediction (2c), the Kalman update step (2d), and the LQR input computation (3c).

Figure 1 depicts an overview of LQG control in an adaptive optics control loop, which shows that control requires only five matrix multiplication operations per timestep. Running LQG control is therefore computationally inexpensive, achieving optimal performance with computations that can be carried out within a fraction of the system's frame delay.

## 2.2 Model building for LQG in an AO context

Table 1 describes all of the matrices involved. To apply LQG control to adaptive optics, it suffices to specify  $A, B, C, D, W, V$  describing the system's time-evolution.

We follow a model choice similar to that in work by Meimon et al.,<sup>5</sup> in which tip-tilt dynamics are the sum of vibration and atmospheric components, as described using autoregressive models of the form

$$x_i = \sum_{k=1}^p a_k x_{i-k} + \epsilon_i, \quad (4)$$

where  $\epsilon_i$  is i.i.d Gaussian noise and  $a_k$  are known coefficients. Here, we restrict models to the AR(2) case, with  $x_i = a_1 x_{i-1} + a_2 x_{i-2} + \epsilon_i$ . Identification consists of analyzing an open-loop timeseries in order to find a combination of autoregressive models that describe its dynamics well.

To identify the system, we first take the power-spectral-density (PSD) of the open-loop timeseries. The PSD is computed using Welch's method with a Hanning window of size  $2^{\lfloor \log_2 \frac{N}{4} \rfloor}$  for a timeseries of length  $N$ , and with the DC component ( $f = 0$ ) removed. We perform a linear fit in log-log space in a frequency range  $\left[\frac{f_s}{120}, \frac{f_s}{3}\right]$ . We mask out any datapoints that are greater than 3 standard deviations, replacing them with a new log-log linear interpolation, to get the modified PSD we use to fit the atmospheric model.

We identify up to  $N_{\text{vib,max}} = 10$  vibrational modes in each of tip and tilt by iteratively finding the largest peak in the PSD, using the *scipy.signal.find\_peaks*<sup>12</sup> function, and again ensuring that it falls within the frequency range  $\left[\frac{f_s}{120}, \frac{f_s}{3}\right]$ . For each one, we make a modified PSD, where all values that are 3 standard deviations above the linear fit are masked out, except those around the identified central frequency. This provides a clear peak that can be well described by an AR(2) model.

To fit an autoregressive model, we use the PSDs rather than the timeseries so that we can exploit spectral separation between components, as well as separating out measurement noise, which tends to dominate above  $\frac{f_s}{3}$ . We use the bounded Levenberg-Marquardt minimization method as implemented in the Python *lmfit*<sup>13</sup> package to fit the log of a modified PSD to an autoregressive power-law model of order 2.<sup>14</sup>

$$P(f) = \frac{\sigma^2/f_s}{|1 + a_1 \exp(-i2\pi f/f_s) + a_2 \exp(-i4\pi f/f_s)|^2} \quad (5)$$

$$\ln P(f) = 2 \ln \sigma - \ln f_s - 2 \log |1 + a_1 \exp(-i2\pi f/f_s) + a_2 \exp(-i4\pi f/f_s)|. \quad (6)$$

For the atmospheric trend, we fit an AR(2) model as above (fitting for  $a_1, a_2, \sigma$ ), which has been noted by Meimon et al.<sup>5</sup> as providing an adequate description of atmospheric turbulence; however, this order could be increased if necessary to capture more complex dynamics that are correlated in time. For each vibration peak, we fit the same model but for  $f, k, \sigma$ , where  $a_1, a_2$  are related to  $f, k$  by discretization of a damped exponential:

$$\begin{aligned} a_1 &= 2e^{-k2\pi(f/f_s)} \cos\left(2\pi\sqrt{1 - k^2}(f/f_s)\right) \\ a_2 &= -e^{-4\pi k(f/f_s)} \end{aligned} \quad (7)$$

$f$  is provided as an initial value from the peak fit, rather than a fixed parameter, since its true value may be different than the peak-fitted value by up to  $\Delta f = \frac{f_s}{2^{\lfloor \log_2 \frac{N}{4} \rfloor + 1}}$  due to the window size and the Nyquist sampling theorem. It is bounded to within  $\pm \Delta f$  of its initial value.

Finally, we identify the measurement noise by averaging the portion of the PSD in the frequency range  $\left[\frac{f_s}{3}, \frac{f_s}{2}\right]$  (the left endpoint of which may be adjusted: the selected portion of the PSD should be relatively flat, so that there are no dynamics being included in the estimate) and taking the root-mean-square estimate,

$$\sigma_v = \sqrt{\langle f_s P_{f \geq f_s/3} \rangle}, \quad (8)$$

which accurately represents measurement noise as it recovers the average Fourier transform, which sends a Gaussian in time to one in frequency.

Having found the autoregressive coefficients and measurement noise, we build the system matrices. An AR(2) model with process noise  $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$  can be described in state space by

$$\mathbf{x}[k] = [x_k \quad x_{k-1}]^\top, \quad A = \begin{bmatrix} a_1 & a_2 \\ 1 & 0 \end{bmatrix}, \quad B = [0 \quad 0]^\top, \quad C = [1 \quad 0], \quad W = \begin{bmatrix} \sigma^2 & 0 \\ 0 & 0 \end{bmatrix} \quad (9)$$

and to combine models, we stack their  $A$  and  $W$  matrices in block-diagonal form, while concatenating their  $B$  matrices vertically and their  $C$  matrices horizontally. This is combined with  $D = [1]$  and  $V = [\sigma_v^2]$  independently of the models being used. In total, a system matrix for both modes that models atmospheric turbulence and one vibrational mode has the representation (where all blank entries are implicitly zeros),

$$\begin{aligned} \mathbf{x}_k &= \begin{bmatrix} x_k^{tip,atm} & x_{k-1}^{tip,atm} & x_k^{tip,vib} & x_{k-1}^{tip,vib} & x_k^{tilt,atm} & x_{k-1}^{tilt,atm} & x_k^{tilt,vib} & x_{k-1}^{tilt,vib} \end{bmatrix} \\ A &= \begin{bmatrix} a_1^{tip,atm} & a_2^{tip,atm} & & & & & & \\ 1 & 0 & & & & & & \\ & & a_1^{tip,vib} & a_2^{tip,vib} & & & & \\ & & 1 & 0 & & & & \\ & & & & a_1^{tilt,atm} & a_2^{tilt,atm} & & \\ & & & & 1 & 0 & & \\ & & & & & & a_1^{tilt,vib} & a_2^{tilt,vib} \\ & & & & & & 1 & 0 \end{bmatrix} \\ B &= [\mathbf{0}]_{8 \times 2} \\ C &= \begin{bmatrix} 1 & 0 & 1 & 0 & & & & \\ & & & 1 & 0 & 1 & 0 & \end{bmatrix} \\ D &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \\ W &= \text{diag}([\sigma_{tip,atm}^2 \quad 0 \quad \sigma_{tip,vib}^2 \quad 0 \quad \sigma_{tilt,atm}^2 \quad 0 \quad \sigma_{tilt,vib}^2 \quad 0]) \\ V &= \begin{bmatrix} \sigma_{v,tip}^2 & 0 \\ 0 & \sigma_{v,tilt}^2 \end{bmatrix}. \end{aligned}$$

### 2.3 Delay transform

Adaptive optics control loops typically have a frame delay, which is characterized in the model by changing the update rule to

$$\mathbf{x}[k+1] = A\mathbf{x}[k] + Bu[k-d], \quad (10)$$

for a delay of  $d \geq 0$  frames. In the state-space formulation, this can be done by augmenting the state as follows. Let  $\mathbf{x}' = [\mathbf{x} \quad \mathbf{u}_{k-1} \quad \dots \quad \mathbf{u}_{k-d}]^\top$  be the new state, with dimension  $n + pd$ . The corresponding system matrices are updated as follows:

$$A' = \begin{bmatrix} A_{n \times n} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & B_{n \times p} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & I_{p \times p} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & I_{p \times p} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & I_{p \times p} & \mathbf{0} \end{bmatrix}_{(s+pd) \times (s+pd)}$$

## Simulated LQG control results




Figure 2. Simulated LQG control with and without delay. On the open-loop data (blue), control is simulated using three schemes: the integrator (orange), LQG without the delay terms (green), and LQG with the delay terms (red). We see LQG with the delay terms performs better than the others.

$$\begin{aligned}
 B' &= [\mathbf{0}_{n \times p} \quad I_{p \times p} \quad \mathbf{0} \quad \dots \quad \mathbf{0} \quad \mathbf{0}]_{(n+pd) \times p}^\top \\
 C' &= [C \quad \mathbf{0} \quad \mathbf{0} \quad \dots \quad \mathbf{0} \quad D]_{m \times (s+pd)} \\
 D' &= [0]_{m \times p} \\
 W' &= \begin{bmatrix} W & \mathbf{0} \\ \mathbf{0} & \mathbf{0}_{pd \times pd} \end{bmatrix} \\
 V' &= V
 \end{aligned}$$

### 2.4 Method validation through simulations

We validate LQG control and the effects of introducing delays first on a toy problem, then on data that are representative of real tip-tilts. For both of these simulations, we generate sets of process noise  $\mathbf{w}[k]$  and measurement noise  $\mathbf{v}[k]$ , and iterate  $\mathbf{x}[k+1] = A\mathbf{x}[k] + B\mathbf{u}[k] + w[k]$  and  $\mathbf{y}[k] = C\mathbf{x}[k] + D\mathbf{u}[k] + \mathbf{v}[k]$  for various choices of control law  $\mathbf{u}_k(\mathbf{y}_1, \dots, \mathbf{y}_k)$ . Direct comparisons are possible due to the same  $\mathbf{w}_k, \mathbf{v}_k$  being used.

For an identity model (i.e. the state remains constant except for process noise) with a frame delay, which is described by the state-space system  $A = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}, B = [0 \quad 1]^\top, C = [1 \quad 0], D = [0], W = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, V = [1]$ , we simulate four control laws: open-loop, where  $\mathbf{u}_k = \mathbf{0}$ ; integrator control, where  $\mathbf{u}_k = -g\mathbf{y}_k$  given some constant gain  $g$ , in this case 0.1; LQG control for the equivalent system without a frame delay; and LQG control for this system.

Figure 2 shows the results. As expected, we see that LQG control for the true system provides the best performance, and further that the controlled residuals are spectrally flat.




Figure 3. Simulated LQG control for three-vibration signal

We further test LQG against integrator control with a model containing three vibration peaks, the results of which are shown in Figure 3. As expected, the LQG residuals are flat both in time and frequency, and the empirical rejection transfer function (closed-loop PSD divided by open-loop PSD) shows strong inverse peaks at each of the vibration peaks.

### 3. SETUP AND EXPERIMENTS

#### 3.1 Lab setup

We made use of the FAST<sup>15</sup> (Fast Atmospheric Self-coherent camera Technique) method on the SEAL<sup>16,17</sup> (Santa cruz Extreme Adaptive optics Laboratory) testbed for this work. FAST was implemented using an Andor Zyla 5.5 sCMOS camera and an ALPAO 97 actuator deformable mirror, with the intermediate wavefront sensing steps carried out within the Python codebase described in §3.3.

Instead of generating turbulence on the bench, we simulate disturbances by applying tip and tilt commands in closed loop that are then corrected. This has the advantage of allowing us to define disturbances that precisely match the dynamic model of the LQG controller, or that deviate from it in known ways.

#### 3.2 Calibration and measurement

Before experiments can be run, several calibration steps are executed. These steps ensure the optics are aligned and will produce the desired results in simpler test cases.

The first step is optical flattening, which is described in detail in a separate work by Gerard et al.,<sup>17</sup> but also summarized here. This works by optically flattening the ALPAO DM with a Thorlabs Shack Hartmann Wavefront Sensor (model WFS-20). We use a classical poke matrix, singular value decomposition (SVD)-based pseudo inverse to generate the command matrix while minimizing the propagation of unwanted DM modes such as waffle, and then iteratively closed the loop to optically flatten the ALPAO DM, typically converging at  $\sim 40$  nm rms. We repeat this process  $\sim$ daily in order to mitigate known ALPAO quasi-static drift effects.

After flattening, optical alignment is run. This process consists of a tip/tilt grid search to find the optimal best-flat position of the camera image. In order to find this optimal position, we look at the modulation transfer function (MTF) of the image at each candidate position, possibly after manually steering in order to reach a






Figure 4. Linearity of FAST in each Zernike mode

region where interference effects are clearly visible. The MTF is compared against a pre-computed ideal MTF, which consists of a central lobe surrounded by two side lobes, by computing the contrast ratio of a side lobe to the central lobe. We run this grid search once over a coarse grid and once over a fine grid within the optimal region. Indicators of reaching an optimal position are clearly visible and separated side lobes in the MTF, and a contrast ratio heatmap that peaks roughly in the center of the grid. This optimal position is saved as the best flat, and an averaged image taken at this position is saved for use in measurements.

From an aligned position, an interaction matrix is computed, by applying ideal Zernike shapes and noting the change relative to the best flat in  $L_+$  space. This results in a reference vector, whose outer product with itself is the interaction matrix. The SVD-based pseudo-inverse of this matrix is the command matrix, which we use to measure Zernike coefficients from images. The SVD cutoff to compute the command matrix is optimized to account for changing day-to-day conditions in the lab. The Zernike modes in an image are then measured by computing the change in  $L_+$  space relative to a flat reference, as above, and projecting the result onto the basis of ideal Zernike responses in the command matrix.

We check system linearity by applying linear perturbations in each Zernike mode, and plotting against the actual measured output in each mode. Under ideal conditions, this graph is close to  $y = x$  in the mode being perturbed, i.e. a unit input results in a unit output, and close to  $y = 0$  in other modes, i.e. no cross-talk between modes. This holds until the perturbation is sufficiently off the focal plane mask, significantly decreasing the intensity of interference fringes. If these conditions are not visibly satisfied, the above steps are rerun. Figure 4 shows representative linearity plots on the bench.

### 3.3 Software setup

A SEAL adaptive optics experiment is completely specified using the system optics (the actual hardware or a simulation mode), the duration of an experiment, and the disturbance schedule (specifying the tip and tilt disturbance values to be put on at each timestep). This is paired at runtime with a controller, which is a function

that takes in a system measurement and returns the tip and tilt output to apply. In this way, multiple controllers can be run on the same random process in simulation, and can be run sequentially on the bench.

Before running an experiment, we check hardware alignment by ensuring that no component of the Zernike measurement at the best-flat DM position exceeds a threshold (arbitrarily set at  $10^{-3}$  of the maximum possible actuator input). If this is the case, alignment is automatically rerun and the command matrix is renewed. If this automatic alignment fails after a few iterations, the experiment terminates, and we attempt manual realignment.

After checking alignment, the control loop is run by starting two independent processes, as shown in Figure 5. The first generates a schedule of open-loop disturbances in tip and tilt to apply at each timestep, and applies them to the DM at the appropriate times. The second runs the control loop: in each iteration, it gets an image from the camera, processes it to generate the measured Zernike coefficients, computes the command based on that measurement, and applies it to the DM. The second process also logs the times at which each event happens, the measured coefficient values, and the control commands applied.

### 3.4 Control loop delays

Although Python is not a language typically used for real-time control, we developed a solution, described below, to enable sufficient high-speed performance. This solution combined with the preferred user interface and modular abilities of Python made this approach the best choice for the purposes of our laboratory demonstration.

A key aspect to running real-time control in Python is achieving concurrent execution, so that the delay between a disturbance and its correction is either small enough to be negligible, or a relatively constant multiple of the frame rate so that the controller can compensate for it. This is achieved in our implementation by spin-locking each thread, forcing it to stay in a dummy infinite loop until the system time reaches the time of the previous frame plus  $\Delta t$ . This works in conjunction with Python *multiprocessing.Process* objects, which run separate instances of Python and therefore avoids interdependences that would otherwise interfere with the execution order due to the Global Interpreter Lock.<sup>18</sup>

To account for slight imprecision in the timing controls in Python, a SEAL experiment will drop a frame if the loop computation time ever exceeds the loop rate. For instance, if a loop iteration was set to end at  $t = 0.01$ s, but the iteration takes 0.013s to complete, the experiment will spinlock until  $t = 0.02$ s so that the system matches the underlying discrete-time model. In practice, at most one frame is lost per 1000-frame experiment and most control experiments do not drop any frames in this way, so overall control performance is not significantly affected by the implementation.

Figure 6 shows a schedule, as measured on the bench, for a typical run of this method. We note that each corrective DM action is applied within half a frame of the corresponding camera exposure, and the loop runs at 100 Hz without any events occurring out of sync or drifting over time. This shows us that the experimental setup is robust enough to satisfy the assumptions of LQG control.

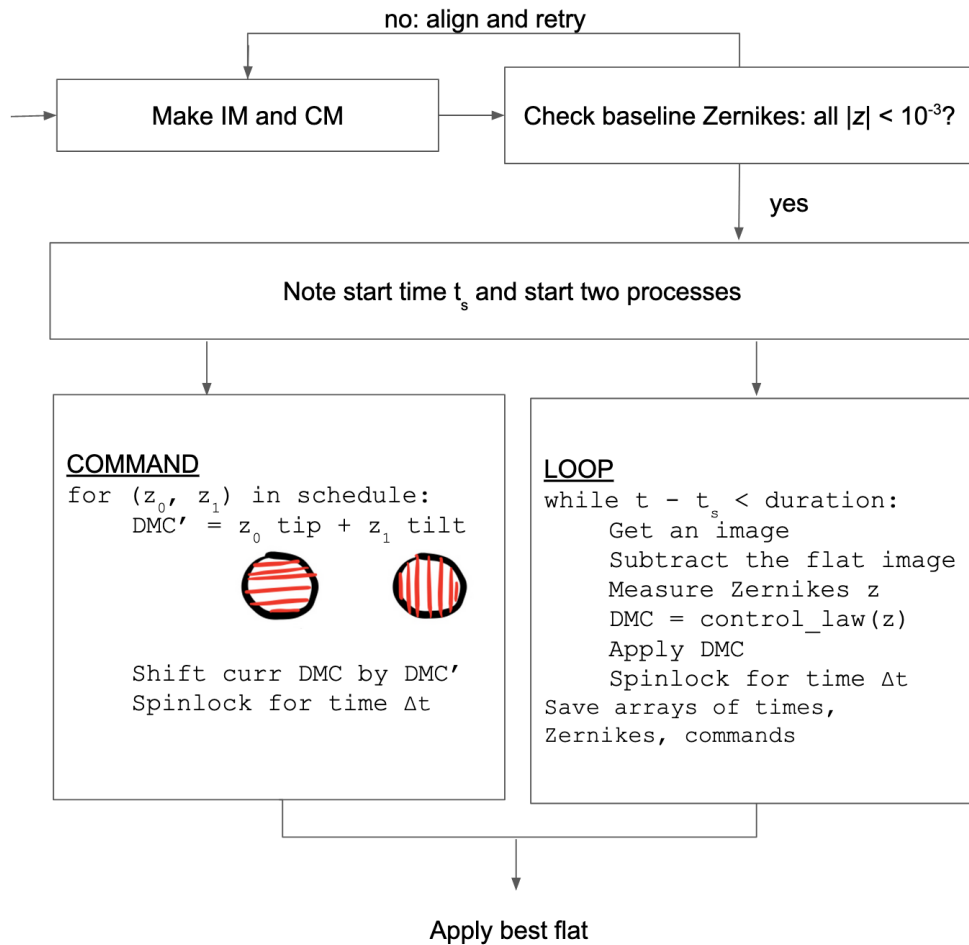
We can further see the effectiveness of this implementation by looking at the delays between consecutive events in the loop. Histograms of these delays for the same typical bench run (lasting 10 seconds/1000 frames) and normalized to 1 frame = 0.01 seconds are shown in Figure 7. We note that the control computation is relatively constant, as we would expect from LQG control (since it consists of the same few matrix multiplications in each iteration), whereas the measurement time is slightly longer but still relatively constant. Since all operations run well within 1 frame, barring a negligible number of dropped frames as described above, we conclude that the bench implementation of closed-loop control is robust.

## 4. RESULTS, DISCUSSION, AND FUTURE WORK

In order to run an LQG control experiment, we first take a 100-second open-loop timeseries. On-sky, this would be equivalently achieved by periodically performing a pseudo-open-loop reconstruction. We use this open-loop time series to build an LQG controller as previously described. For the purposes of control experiments, we specify the AR order of the atmospheric model and the number of vibration peaks to be used in the identification, to ensure that the identified model is able to perfectly match the underlying model used to make the disturbance; on-sky, this tuning would be unnecessary, and we would set these to 2 and 10 respectively.

## ALGORITHM: SEAL EXPERIMENT

**INPUT:** control\_law (function, Zernikes  $\mapsto$  DM command)  
schedule (array of disturbance inputs)  
duration (float, #seconds to run experiment)



**OUTPUT:** times, Zernikes, commands

Figure 5. Computational flow of a SEAL adaptive optics control experiment.




Figure 6. SEAL schedule in closed loop. Each line is the extent of a planned frame, which starts with an exposure (in blue), then a measurement (in orange), then a DM command (in green). For the controller to be in sync, all three dots should appear before the next vertical line.

Delays in the AO loop on SEAL, in #frames




Figure 7. SEAL delay histograms. From left to right, the figure depicts the time from exposure to DM command, exposure to measurement, and measurement to DM command.




Figure 8. SEAL LQG results in the frequency domain

Figure 8 shows the open- and closed-loop PSDs, and the rejection transfer function (their quotient), in both tip and tilt. As expected, it notches out the induced vibration peak, but only in one mode; this may be due to imperfect linearity observed in the other mode, which we discuss further in §4 along with additional possible future work to improve laboratory performance.

From Fig. 4, the FAST linear range is about  $0.1 \mu\text{m PV}$ , converting to  $3.4 \times 10^{-6}$  (ALPAO DM units)<sup>2</sup>. Looking at Fig. 8, this linear range is below amplitudes measured at frequencies  $\lesssim 1$  Hz, which means FAST non-linearities should begin to dominate at these low temporal frequencies, potentially explaining the limited closed-loop gains as we have currently measured. Similarly, the Mode 0 vibration peak centered around 8 Hz, reaches amplitudes above this non-linear regime, whereas the Mode 1 vibration peak shows smaller amplitudes, potentially explaining why the latter mode is corrected but the former is not. A future demonstration where open loop vibration is introduced with the DM at controlled amplitudes within the FAST linear range could help improve performance.

Two additional major ways in which these experiments could be extended would be by increasing the loop rate by identifying and removing computational bottlenecks, and by testing more computationally intensive control algorithms.

The main computational bottleneck in closed-loop control using FAST is the measurement step, which involves a 2D Fourier transform, multiplication in Fourier space by a mask, and an inverse Fourier transform back. This computation takes about 3-5ms on SEAL, and would ideally take  $< 1$ ms. The mask is sufficiently large that an approach using the fast Fourier transform is faster than computing individual components, and the camera dimension is sufficiently large that Fourier methods are faster than spatial-domain convolution. However, since the measurement computation is a static operation over time (consisting of Fourier transforms of the same dimension and a multiplication with the same mask at each iteration), using a lower-level language or offloading some of the computations to a GPU would provide a substantial speedup; in particular, this would make loading the entire image array into Python unnecessary, instead only requiring the much-smaller array of Zernike coefficients to be loaded. Further, the overall operation is linear, suggesting that it could be even further condensed into a single matrix multiplication. A lower-level language would also be well suited to more tightly controlling event timings. In particular, Rust is a good candidate for this task, as its binaries can be loaded as Python libraries and it allows for low-level control and optimization while ensuring memory safety. Porting some or all of the underlying optics and experiment code to Rust would therefore still allow the control algorithms, or other aspects that may require customization, to be written in Python.

This work serves as a proof of concept for testing arbitrarily computationally intensive control algorithms in this setup provided they are delay-tolerant. Due to concurrent processing, the only limiting factor on a control law is the loop rate, and control laws that take longer than the loop rate to compute can still be run, provided they are designed to tolerate this inherent delay. This allows for testing of algorithms like predictive wavefront control<sup>19</sup> with FAST, using prototype models in Python. This lowers the barrier between simulation and bench testing of a control algorithm.

Delay tolerance can be extended further to the controller model. The vibration model used in the LQG controller is a discrete-time realization of a continuous-time process, so by using the underlying infinitesimal generator matrix to time-evolve the underlying model for arbitrary lengths of time, an externally-imposed loop rate could be relaxed. The atmospheric model is based on discrete-time computations, but could be approximated in continuous time via linear interpolation. In the case where the lag is Markovian, which is a reasonable assumption in the vibration LQG case where the problem complexity can increase or decrease noticeably with the changing dimensions due to different numbers of peaks being identified, optimal observers and controllers similar to those used in this work<sup>20,21</sup> can be designed. This would allow the controller to function better in situations where the loop delay is non-integer or difficult to characterize.

## 5. CONCLUSION

We have demonstrated the effectiveness of LQG control for tip-tilt using FAST, and developed robust testing infrastructure for real-time wavefront control. Our main findings are as follows.

- We present a laboratory demonstration of LQG control with FAST, showing RMS wavefront error improvement compared to an integrator (over all temporal frequencies) and frequency-domain notching of the injected and corrected 9 Hz vibration mode. This is an important milestone for demonstrating the promise of using such LQG technology to suppress vibrations at moderate- to high-speeds with a focal plane WFS.
- In these tests we demonstrate the benefit of including plant dynamics into the LQG model, primarily through the inclusion of a delay term.

These results serve as proof of concept for the use of high-speed focal plane WFSs in control experiments, particularly where the control algorithm has high computational requirements. We have also made use of a pure Python modular codebase, ensuring that varying the control law or features of the control loop for future research is straightforward. The modular structure further ensures that the control loop could be further optimized, possibly increasing the loop rate, without changing the Python interface for defining the control law.

## ACKNOWLEDGMENTS

AS thanks Deana Tanguay and Enrico Ramirez-Ruiz for support throughout the Lamat REU program, and Parth Nobel and Aled Cuda for consultations on implementing concurrent execution.

## APPENDIX A. OVERVIEW OF FAST

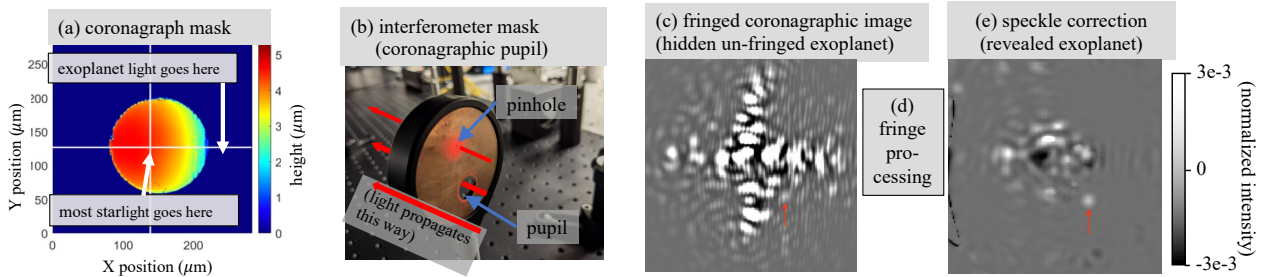


Figure 9. Overview of FAST, adapted from Ref. 17, using images from the Santa Cruz Extreme AO Laboratory (SEAL) high contrast testbed.<sup>16</sup> A custom coronagraph mask (a) is designed to both (i) attenuate starlight while transmitting exoplanet light, and (ii) enhance the stellar fringe visibility using an interferometer mask in the downstream coronagraphic pupil (b). This design enables the detection of an interference pattern on top of the leftover stellar speckles in the coronagraphic image (including for low order modes that are blocked by the coronagraph mask in panel a but still transmitted through the pinhole in panel b), even when only a few photons are detected per pixel (c), enabling coronagraphic science imaging and AO wavefront sensing. The fringed stellar speckles are then processed (d) and further attenuated using both a deformable mirror and via post-processing (e), revealing an un-fringed exoplanet.

Fig. 9 shows an overview of the Fast Atmospheric Self-coherent camera (SCC) Technique (FAST),<sup>15</sup> adapted from Ref. 17. FAST relies on the SCC technique,<sup>22</sup> in which coherent starlight is interfered with itself to enable a measurement and correction of the complex speckle electric field (i.e., phase and amplitude) in the coronagraphic image. Reviewing a mathematical description of the SCC/FAST image, the recorded image on the focal plane detector,  $I$ , produced from noiseless (i.e., without photon/detector/sky background noise) propagation of wavefront error (WFE; i.e., from atmospheric and/or instrumental origin) through to the detector, can be described by

$$I = I_S\{\theta\} + I_P\{\theta\} + I_R\{\theta\} + 2\sqrt{I_S\{\theta\}I_R\{\theta\}} \cos[2\pi(\xi_0/\lambda)\phi\{\theta\}], \quad (11)$$

where  $\theta$  represents the two-dimensional coordinates in the image,  $I_S$  is the un-fringed stellar speckle intensity component,  $I_P$  is the un-fringed planet intensity component,  $I_R$  is the pinhole point-spread-function (PSF) component (i.e., the PSF recorded if the main pupil aperture was blocked), and  $\phi$  is the electric field phase difference between the  $I_S$  and  $I_R$  image components. The spatial scale of the fringe term  $2\sqrt{I_S I_R} \cos[2\pi(\xi_0/\lambda)\phi]$  is set by the Lyot stop's pupil-to-pinhole separation ( $\xi_0$ , projected onto the entrance pupil) and observing



wavelength,  $\lambda$ , enabling sub- $\lambda/D$  spatial scales (between  $\lambda/[0.5\xi_0]$  and  $\lambda/[1.5\xi_0]$ ) to be isolated in the Fourier plane of the image without bias from the un-fringed exoplanet signal  $I_P$ . The phase of a given stellar speckle, which is not normally measureable in a single coronagraphic image (i.e., not encoded in  $I_S$  or  $I_R$ ), is encoded by the relative position of fringes on that speckle, which changes with  $\phi$ . This relative phase can be isolated in the Fourier plane of the image (known as the optical transfer function or OTF, the amplitude of which is known as the modulation transfer function or MTF) via a binary mask multiplication of these “side lobe” spatial frequencies; the complex-valued image plane after Fourier-filtering the MTF sidelobe is commonly referred to as  $I_-$ , representing the fringe amplitude and phase components in the same plane as the detector. This phase measurement is calibrated empirically by applying sine and cosine waves over the full spatial frequency range of the DM (or for the purposes of this paper, analogously with low order Zernike modes). By recording images (i.e., eq. 11) that are then converted to  $I_-$  during this calibration routine to build an AO-based command matrix, DM commands can then be generated from a single image with an unknown speckle pattern pattern via a linear least-squares-based matrix vector multiplication (MVM; ideal for closed-loop AO operations if fringes can be detected on millisecond timescales).

FAST’s key innovation compared with previous focal plane wavefront sensing and control is the coronagraphic focal plane mask (FPM; Fig. 9a), which provides a high enough fringe visibility (i.e., fringed vs. un-fringed components) in the final image such that fringes can be detected on millisecond timescales, even when only a few photons are detected per pixel. The classical SCC also generates fringes in the coronagraphic image to enable focal plane wavefront control,<sup>22</sup> but with fringe visibilities about  $10^6$  times lower than Fig. 9c,<sup>15</sup> removing any possibility of AO-based wavefront control on millisecond timescales. Also due to the nature of the FAST FPM, which sends the core of the PSF (i.e., low order modes of the entrance pupil wavefront) into the off-axis Lyot stop pinhole, in turn transmitted through to the SCC image, FAST is sensitive to low order modes of the entrance pupil wavefront<sup>23</sup> while typically the classical SCC is comparatively not sensitive to such modes.<sup>22</sup>

## APPENDIX B. LINEAR-QUADRATIC-GAUSSIAN DERIVATION

Solving the optimal control problem in the infinite-horizon case consists of choosing control inputs  $\mathbf{u}_k$  such that the following cost function is minimized:

$$J(\mathbf{u}_0, \mathbf{u}_1, \dots) = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[ \sum_{k=0}^{n-1} \mathbf{x}_k^\top Q \mathbf{x}_k + \mathbf{u}_k^\top R \mathbf{u}_k + 2\mathbf{u}_k^\top N \mathbf{x}_k \right], \quad (12)$$

where  $Q$  describes the penalty for state deviations,  $R$  describes the penalty for excessive control action, and  $S$  denotes a cross term between these two (and  $\mathbb{E}$  is the expectation operator). We specialize to the case of minimizing the true measurements, i.e. minimizing  $\mathbf{y}_k^\top \mathbf{y}_k$ . We use this formulation, as opposed to the case where  $D = \mathbf{0}$  shown in most sources,<sup>24</sup> because the adaptive optics control problem consists of controlling a measurable property (the tip and tilt positions) of an underlying process, rather than controlling that process itself. The appropriate model for this case has  $B = \mathbf{0}$ , which under the usual formulation would not be controllable. This results in our choices for  $Q, R, S$ :

$$\mathbf{y}_k^\top \mathbf{y}_k = (C\mathbf{x}_k + D\mathbf{u}_k)^\top (C\mathbf{x}_k + D\mathbf{u}_k) = \mathbf{x}_k^\top \underbrace{C^\top C}_Q \mathbf{x}_k + 2\mathbf{u}_k^\top \underbrace{D^\top C}_S \mathbf{x}_k + \mathbf{u}_k^\top \underbrace{D^\top D}_R \mathbf{u}_k. \quad (13)$$

This has the advantage of providing a clear choice for these weighting matrices rather than leaving them as ad hoc parameters; any deviations from optimal performance are necessarily due to model mismatches or similar errors, rather than choosing the weighting matrices suboptimally. In practice, this helps with accurate system identification.

We can perform linear-quadratic-Gaussian control by making use of the separation principle in control theory. Without losing optimality, the problem can be split into two parts:

1. Design of an optimal observer: computing the minimum mean-square-error state estimates  $\hat{\mathbf{x}}_{k+1} \triangleq \mathbb{E}[\mathbf{x}_{k+1} | \hat{\mathbf{x}}_k, \mathbf{y}_k]$ .
2. Design of an optimal noiseless regulator: assuming that our observed states  $\hat{\mathbf{x}}_k$  are the true states  $\mathbf{x}_k$ , computing  $\mathbf{u}_k = \arg \min_{\tilde{\mathbf{u}}_k} J(\tilde{\mathbf{u}}_0, \tilde{\mathbf{u}}_1, \dots)$ .

Both problems can be reduced to the discrete algebraic Riccati equation (DARE), which is a nonlinear matrix equation of the form

$$P = A^\top P A - (A^\top P B + S)(R + B^\top P B)^{-1}(B^\top P A + S^\top) + Q. \quad (14)$$

We denote “ $P$  solves Equation 14” by  $P = \text{dare}(A, B, Q, R, S)$ .

Along the lines of similar LQG derivations,<sup>24</sup> we show how to reduce the observer and controller problems to the DARE, and how to compute the optimal gain, just for the control problem, since the case we use here (with  $D \neq \mathbf{0}$  and with the cost cross term) is relatively uncommon.

Let the cost starting from a state  $\mathbf{x}_0 = \mathbf{z}$  be

$$V(\mathbf{z}) = \min_{\mathbf{u}_0, \mathbf{u}_1, \dots} \sum_{\tau=0}^{\infty} \mathbf{y}_\tau^\top \mathbf{y}_\tau = \min_{u_0, u_1, \dots} \sum_{\tau=0}^{\infty} \mathbf{x}_k^\top C^\top C \mathbf{x}_k + 2\mathbf{u}_k^\top D^\top C \mathbf{x}_k + \mathbf{u}_k^\top D^\top D \mathbf{u}_k. \quad (15)$$

We minimize over one control step by splitting the cost into a one-step term, and the cost from there due to the new state. We pick  $u_0$  equal to the  $w$  minimizing this form of the cost:

$$V(\mathbf{z}) = \min_w (\mathbf{z}^\top C^\top C \mathbf{z} + 2\mathbf{w}^\top D^\top C \mathbf{z} + \mathbf{w}^\top D^\top D \mathbf{w} + V(A\mathbf{z} + B\mathbf{w})) \quad (16)$$

We look for a solution such that  $V$  is quadratic in  $\mathbf{z}$ , i.e. it has the form  $V(\mathbf{z}) = \mathbf{z}^\top P \mathbf{z}$ . Substituting this in to Equation 16, we get

$$\mathbf{z}^\top P \mathbf{z} = \mathbf{z}^\top C^\top C \mathbf{z} + \min_{\mathbf{w}} (2\mathbf{w}^\top D^\top C \mathbf{z} + \mathbf{w}^\top D^\top D \mathbf{w} + \mathbf{z}^\top A^\top P A \mathbf{z} + 2\mathbf{w}^\top B^\top P A \mathbf{z} + \mathbf{w}^\top B^\top P B \mathbf{w}). \quad (17)$$

The  $\mathbf{w}$  minimizing the right-hand side, which we will take as  $\mathbf{u}_0$ , can be found by quadratic minimization:

$$\mathbf{u}_0 = \arg \min_{\mathbf{w}} \mathbf{w}^\top \underbrace{(D^\top C + B^\top P A)}_{L_1} \mathbf{z} + \frac{1}{2} \mathbf{w}^\top \underbrace{(D^\top D + B^\top P B)}_{L_2} \mathbf{w}, \quad (18)$$

which has the standard solution

$$\mathbf{u}_0 = -(L_2)^{-1} L_1 \mathbf{z} = -(D^\top D + B^\top P B)^{-1} (D^\top C + B^\top P A) \mathbf{z}. \quad (19)$$

Substituting this back into 17 and dropping the leading  $\mathbf{z}^\top$  and trailing  $\mathbf{z}$  on each term (the equation has to be true for all  $\mathbf{z}$ , so the corresponding matrix equation is equivalent), we get

$$P = C^\top C + A^\top P A - (C^\top D + A^\top P B)(D^\top D + B^\top P B)^{-1}(D^\top C + B^\top P A), \quad (20)$$

and given  $P$  satisfying this, we compute the optimal control command according to Equation 19.

## REFERENCES

- [1] Boehle, A., Ghez, A. M., Schödel, R., Meyer, L., Yelda, S., Albers, S., Martinez, G. D., Becklin, E. E., Do, T., Lu, J. R., Matthews, K., Morris, M. R., Sitarski, B., and Witzel, G., “An Improved Distance and Mass Estimate for Sgr A\* from a Multistar Orbit Analysis,” *The Astrophysical Journal* **830**, 17 (Oct. 2016).
- [2] Delorme, J.-R., Jovanovic, N., Echeverri, D., Mawet, D., Kent Wallace, J., Bartos, R. D., Cetre, S., Wizinowich, P., Ragland, S., Lilley, S., Wetherell, E., Doppmann, G., Wang, J. J., Morris, E. C., Ruffio, J.-B., Martin, E. C., Fitzgerald, M. P., Ruane, G., Schofield, T., Suominen, N., Calvin, B., Wang, E., Magnone, K., Johnson, C., Sohn, J. M., López, R. A., Bond, C. Z., Pezzato, J., Sayson, J. L., Chun, M., and Skemer, A. J., “Keck Planet Imager and Characterizer: a dedicated single-mode fiber injection unit for high-resolution exoplanet spectroscopy,” *Journal of Astronomical Telescopes, Instruments, and Systems* **7**, 035006 (July 2021).
- [3] Hottinger, P., Harris, R. J., Dietrich, P.-I., Blaicher, M., Bechter, A., Crass, J., Glück, M., Pott, J.-U., Bharmal, N. A., Basden, A., Morris, T. J., Koos, C., Sawodny, O., and Quirrenbach, A., “Focal plane tip-tilt sensing for improved single-mode fiber coupling using a 3d-printed microlens-ring,” (2019).
- [4] Le Roux, B., Ragazzoni, R., Arcidiacono, C., Conan, J.-M., Kulcsar, C., and Raynaud, H.-F., “Kalman-filter-based optimal control law for star-oriented and layer-oriented multiconjugate adaptive optics,” in [*Advancements in Adaptive Optics*], Bonaccini Calia, D., Ellerbroek, B. L., and Ragazzoni, R., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **5490**, 1336–1346 (Oct. 2004).
- [5] Meimon, S., Petit, C., Fusco, T., and Kulcsar, C., “Tip-tilt disturbance model identification for Kalman-based control scheme: application to XAO and ELT systems,” *Journal of the Optical Society of America A* **27**, A122 (Sept. 2010).
- [6] Sivo, G., Kulcsár, C., Conan, J.-M., Raynaud, H.-F., Gendron, É., Basden, A., Vidal, F., Morris, T., Meimon, S., Petit, C., Gratadour, D., Martin, O., Hubert, Z., Sevin, A., Perret, D., Chemla, F., Rousset, G., Dipper, N., Talbot, G., Younger, E., Myers, R., Henry, D., Todd, S., Atkinson, D., Dickson, C., and Longmore, A., “First on-sky SCAO validation of full LQG control with vibration mitigation on the CANARY pathfinder,” *Optics Express* **22**, 23565 (Sept. 2014).
- [7] Petit, C., Conan, J. M., Kulcsár, C., Raynaud, H. F., Fusco, T., Montri, J., and Rabaud, D., “First laboratory demonstration of closed-loop Kalman based optimal control for vibration filtering and simplified MCAO,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], Ellerbroek, B. L. and Bonaccini Calia, D., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **6272**, 62721T (June 2006).
- [8] Guyon, O. and Males, J., “Adaptive Optics Predictive Control with Empirical Orthogonal Functions (EOFs),” *arXiv e-prints*, arXiv:1707.00570 (July 2017).
- [9] Gerard, B. L., Marois, C., Galicher, R., and Véran, J.-P., “Fast focal plane wavefront sensing on ground-based telescopes,” in [*Adaptive Optics Systems VI*], Close, L. M., Schreiber, L., and Schmidt, D., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **10703**, 1070351 (July 2018).
- [10] “Slycot: Python wrapper for the subroutine library in systems and control theory (slicot).” <https://github.com/python-control/Slycot> (2020).
- [11] Benner, P., Mehrmann, V., Sima, V., Van Huffel, S., and Varga, A., [*SLICOT—A Subroutine Library in Systems and Control Theory*], 499–539, Birkhäuser Boston, Boston, MA (1999).
- [12] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods* **17**, 261–272 (2020).
- [13] Newville, M., Otten, R., Nelson, A., Ingargiola, A., Stensitzki, T., Allan, D., Fox, A., Carter, F., Michal, Osborn, R., Pustakhod, D., Ineuhaus, Weigand, S., Glenn, Deil, C., Mark, Hansen, A. L. R., Pasquevich, G., Foks, L., Zobrist, N., Frost, O., Beelen, A., Stuermer, azelcer, Hannum, A., Polloreno, A., Nielsen, J. H., Caldwell, S., Almarza, A., and Persaud, A., “lmfit/lmfit-py: 1.0.3,” (Oct. 2021).

- [14] Roberts, S., “Spectral estimation.” [https://www.robots.ox.ac.uk/~sjrob/Teaching/B14\\_SP/b14\\_sp\\_lect4.pdf](https://www.robots.ox.ac.uk/~sjrob/Teaching/B14_SP/b14_sp_lect4.pdf) (2017).
- [15] Gerard, B. L., Marois, C., and Galicher, R., “Fast Coherent Differential Imaging on Ground-based Telescopes Using the Self-coherent Camera,” *The Astronomical Journal* **156**, 106 (Sept. 2018).
- [16] Jensen-Clem, R., Dillon, D., Gerard, B., van Kooten, M., Fowler, J., Kupke, R., Cetre, S., Sanchez, D., Hinz, P., Laguna, C., Doelman, D., and Snik, F., “The Santa Cruz Extreme AO Lab (SEAL): design and first light,” in [*Techniques and Instrumentation for Detection of Exoplanets X*], Shaklan, S. B. and Ruane, G. J., eds., **11823**, 420 – 433, International Society for Optics and Photonics, SPIE (2021).
- [17] Gerard, B. L., Dillon, D., Cetre, S., Jensen-Clem, R., Yuzvinsky, T. D., and Schmidt, H., “First experimental results of the fast atmospheric self-coherent camera technique on the Santa Cruz extreme adaptive optics laboratory testbed: demonstration of high speed focal plane wavefront control of residual atmospheric speckles,” in [*Techniques and Instrumentation for Detection of Exoplanets X*], Shaklan, S. B. and Ruane, G. J., eds., **11823**, 459 – 476, International Society for Optics and Photonics, SPIE (2021).
- [18] “Initialization, finalization, and threads.” <https://docs.python.org/3/c-api/init.html>.
- [19] van Kooten, M. A., Jensen-Clem, R., Cetre, S., Ragland, S., Bond, C. Z., Fowler, J., and Wizinowich, P., “Status of predictive wavefront control on Keck II adaptive optics bench: on-sky coronagraphic results,” in [*Techniques and Instrumentation for Detection of Exoplanets X*], Shaklan, S. B. and Ruane, G. J., eds., **11823**, 446 – 458, International Society for Optics and Photonics, SPIE (2021).
- [20] Feng, X., Loparo, K., Ji, Y., and Chizeck, H., “Stochastic stability properties of jump linear systems,” *IEEE Transactions on Automatic Control* **37**(1), 38–53 (1992).
- [21] Mahmoud, M. and Shi, P., “Robust kalman filtering for continuous time-lag systems with markovian jump parameters,” *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications* **50**(1), 98–105 (2003).
- [22] Baudoz, P., Boccaletti, A., Baudrand, J., and Rouan, D., “The Self-Coherent Camera: a new tool for planet detection,” in [*IAU Colloq. 200: Direct Imaging of Exoplanets: Science & Techniques*], Aime, C. and Vakili, F., eds., 553–558 (Jan. 2006).
- [23] Gerard, B. L., Véran, J.-P., Singh, G., Herriot, G., Lardière, O., and Marois, C., “Fast focal plane wavefront sensing as a second stage adaptive optics wavefront sensor,” in [*Adaptive Optics Systems VII*], Schreiber, L., Schmidt, D., and Vernet, E., eds., **11448**, 483 – 499, International Society for Optics and Photonics, SPIE (2021).
- [24] Boyd, S., “Linear quadratic stochastic control with partial state observation.” <https://stanford.edu/class/ee363/lectures/lqg.pdf> (2008).