

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Adaptive control for the IOS adaptive optics system

Jonathan Tesch, Tuan Truong, Lewis C. Roberts Jr., Steve Gibson

Jonathan Tesch, Tuan Truong, Lewis C. Roberts Jr., Steve Gibson, "Adaptive control for the IOS adaptive optics system," Proc. SPIE 11135, Unconventional and Indirect Imaging, Image Reconstruction, and Wavefront Sensing 2019, 1113509 (6 September 2019); doi: 10.1117/12.2527386

SPIE.

Event: SPIE Optical Engineering + Applications, 2019, San Diego, California, United States

Adaptive Control for the IOS Adaptive Optics System

Jonathan Tesch^a, Tuan Truong^a, Lewis C. Roberts Jr.^a, and Steve Gibson^b

^aJet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

^bUniversity of California, Los Angeles, Department of Mechanical and Aerospace Engineering, 420 Westwood Plaza, Los Angeles, CA 90095, USA

ABSTRACT

Future adaptive optics systems will require advanced predictive control algorithms that mitigate loop latency by forecasting disturbances. In many applications, such as laser communications terminals, synthesizing such controllers is challenging due to the non-stationary nature of the disturbance statistics over long periods. We present initial experimental results using a new, multichannel, adaptive control algorithm applied on the Integrated Optical System, an extreme AO system under development for the Laser Communications Relay Demonstration. The adaptive controller implicitly tracks disturbance statistics, and provides broadband wavefront control without the need for open loop downtime. The results illustrate the improved disturbance rejection capabilities of the controller compared to a traditional integrator, even when the exogenous disturbance statistics evolve.

Keywords: Adaptive optics, adaptive control, adaptive filtering, laser communications

1. INTRODUCTION

Next-generation laser communications networks will use medium to large-aperture telescopes as ground terminals. Telescopes in the network will be equipped with adaptive optics instruments (AO) to compensate for atmospheric turbulence and enable injection of the downlink signal into a single-mode fiber. These future AO systems must deliver diffraction limited performance, and provide high levels of disturbance rejection similar to those required for current astronomical AO.¹ Achieving this degree of correction, especially at sites with sub-optimal seeing, necessitates extremely high wavefront sensor frame rates, or advanced control laws that mitigate computational latency by predicting wavefront turbulence. Compared to traditional integrator-based controllers, optimal and predictive algorithms can significantly improve wavefront correction in high-turbulence conditions because they inherently “invert” the incident disturbance power spectrum and plant dynamics.

Current predictive control schemes, such as those in Refs. 2 and 3 among many others, require periods of open loop or pseudo-open loop wavefront measurement, from which a fitting algorithm parameterizes a disturbance model from the observed turbulence statistics. This disturbance model is used, along with a plant model capturing the system latencies, in a linear-quadratic-Gaussian (LQG) framework to synthesize a controller that minimizes a quadratic function of the residual wavefront. In the standard infinite-horizon approach, the resulting controller is linear and time-invariant (LTI), and remains optimal if the turbulence statistics are stationary. However, as LQG

Send correspondence to J. Tesch: E-mail: jonathan.a.tesch@jpl.nasa.gov, Telephone: 1 818 354 1449

controllers have no guaranteed robustness margins,⁴ it can be difficult to predict the degradation of performance as conditions evolve due to changing wind velocities, variable sky conditions, structural vibrations, or telescope pointing. This poses a challenge in AO systems for laser communications, which will require nearly continuous day and night high-performance operation that renders periodic re-optimization of the controller impractical.

This paper presents experimental results using a new type of adaptive controller that continually produces optimal control commands without explicit identification of a disturbance model. In contrast to LTI control laws, this approach employs a recursive least squares (RLS) lattice filter to predict a particular linear combination of future wavefronts, and implicitly tailors the controller error rejection to match the prevailing disturbance power spectra at each time step. The lattice filter is computationally efficient and fully multichannel, allowing multiple spatial modes to be jointly optimized with minimal impact on latency.

The adaptive controller has previously been applied to reduce broadband disturbances in a non-real-time adaptive optics experiment to control multiple spatial modes, and in a real-time beam steering experiment at the UCLA Beam Control Laboratory.^{5,6} This paper presents the first efforts to implement the adaptive controller on a facility-class instrument, the Integrated Optical System (IOS), an “extreme” AO system under JPL development for the Laser Communications Relay Demonstration (LCRD).^{1,7} As a laboratory demonstration, the controller was employed against turbulence generated by an atmospheric turbulence simulator, which injects realistic wavefront disturbances into the IOS optical train. In late 2018 the IOS was installed at the coudé focus of the Optical Communications Telescope Laboratory (OCTL), a 1-meter telescope located at JPL’s Table Mountain Facility. We plan to continue testing against on-sky targets after first light is achieved.

This paper is organized as follows: Section 2 contains a brief overview of the IOS instrument, while Sec. 3 summarizes the adaptive controller in a traditional disturbance rejection framework with a generic plant. Section 4 describes how the adaptive algorithm was implemented in the IOS real-time control software. Section 5 discusses our method for identifying a model of the AO plant. Experimental results pitting the classical integrator and adaptive controller are presented in Sec. 6. Finally, conclusions from the demonstration are discussed in Sec. 7.

2. THE IOS ADAPTIVE OPTICS SYSTEM

The IOS is part of the receive system for Optical Ground Station 1, one of two ground stations in the LCRD optical communications network, and located at the 1-meter OCTL telescope on Table Mountain in Wrightwood, California.¹ Once the geosynchronous LCRD space terminal is launched, the IOS will deliver light from a 1545 nm downlink laser to the receive modem’s single mode fiber, with a required coupling efficiency of 55%. This translates into a Strehl ratio of approximately 70%, which must be met at a 20° elevation angle under median seeing. At Table Mountain, median seeing has been measured at 500 nm to indicate a r_o of 5.2 cm at zenith.

The IOS has a woofer-tweeter configuration with a fast steering mirror (FSM), a 12×12 low-order DM (LODM), and a 34 actuator diameter high-order DM (HODM). Wavefront sensing is provided via a Shack-Hartmann wavefront sensor (SHWFS) arranged in a Fried geometry with the HODM. A portion of the downlink signal is diverted to a scoring camera designed to be in stable thermal equilibrium with the modem fiber.

An important feature of the IOS is the ability to inject dynamic wavefront disturbances via the atmospheric turbulence simulator (ATS), which contains two computer-controlled rotating phase

plates. Varying the location and velocity of the plates produces turbulence over a wide range of seeing conditions. When in use, a fold mirror projects a laser source through the ATS and into the IOS, mimicking the properties of the eventual downlink signal.

The baseline controller for the IOS is a classical leaky integrator; adaptive control is not required for the instrument to meet its performance requirements. Meeting the Strehl ratio requirements for the system with this design necessitates low-latency computation and WFS rates on the order of 10 kHz. Control at these rates is provided by the real-time control software suite (RTC), which is derived from the architecture employed on the PALM-3000 AO system at Palomar Observatory.⁸ The RTC runs on an eight-chip DSP board, which uses direct memory access (DMA) to retrieve data directly from the WFS frame grabber. Each DSP chip contains eight cores, with each core measuring WFS centroids to compute commands for a subset of DM/FSM actuators. DMA is again used to transfer commands to the DMs and the FSM, bypassing the CPU. This design enables single-digit microsecond latency for a 56×64 pixel WFS frame, and processing rates exceeding 20 kHz. Instrument telemetry is saved to a database stored on a solid-state drive.

For the results in this paper, the IOS was located in a development laboratory at JPL's Pasadena campus. It was shipped to OCTL in October 2018 and is currently undergoing integration and test activities, which will culminate in on-sky testing during the second half of 2019.

3. THE ADAPTIVE CONTROL PROBLEM

We briefly summarize the theoretical aspects of the adaptive controller from the perspective of a generic disturbance rejection problem, before describing the implementation in the IOS. The structure of the adaptive controller is depicted in Fig. 1, and is similar to the version used for control of laser beam jitter in Ref. 6 and for control in a non-real-time AO experiment in Ref. 5. The controller belongs to a class of receding-horizon algorithms, where a sequence of control commands are computed to minimize disturbances over a finite future horizon. At each time step, only the first entry in this sequence is applied. This approach has a rich heritage in process control, and has the advantage of easily accommodating frequency weighting filters and plant nonlinearities, if known.⁹

We assume the AO system (i.e. plant) to be controlled is described by a stable LTI system $G(z)$ represented by the following state-space system:

$$\begin{aligned} x_G(t+1) &= A_G x_G(t) + B_G u(t) \\ e(t) &= C_G x_G(t) + w(t) \end{aligned} \quad (1)$$

where $x_G(t)$ is the state of the plant at time t , $u(t)$ is the adaptive control command, $e(t)$ is the residual error, and $w(t)$ is an output disturbance sequence to be rejected. Typically, e and w are representations of the residual and disturbance wavefronts projected onto convenient spatial modes. The goal of the adaptive controller is to generate the command u that minimizes the variance of the residual error e over a horizon of future time steps. For AO applications, $G(z)$ may contain any system latencies and mirror dynamics, a feature that is critical as WFS frame rates encroach on the physical bandwidth limitations of the active optical elements.

Only $e(t)$ is observable from the WFS output, thus the controller utilizes an internal model of the plant, $\hat{G}(z)$, to provide estimates of $x(t)$ and $w(t)$

$$\begin{aligned} \hat{x}(t+1) &= \hat{A} \hat{x}(t) + \hat{B} u(t) \\ e(t) &= \hat{C} \hat{x}(t) + \hat{w}(t) \end{aligned} \quad (2)$$




Figure 1: Receding horizon adaptive controller connected to plant $G(z)$.

Identification of the system matrices $(\hat{A}, \hat{B}, \hat{C})$ is discussed in Section 5. However, if this identification is accurate and $\hat{G}(z) = G(z)$, then it is evident from the top portion of Fig. 1 that $\hat{w}(t) = w(t)$, and thus is an estimate of the output disturbance sequence.

Next, we define the column vectors $\hat{w}_h(t)$ and $u_h(t)$ as containing the values of $\hat{w}(t)$ and $u(t)$ over a future horizon of length h :

$$\hat{w}_h(t) = \begin{bmatrix} \hat{w}(t+1) \\ \hat{w}(t+2) \\ \vdots \\ \hat{w}(t+h) \end{bmatrix} \quad u_h(t) = \begin{bmatrix} u(t) \\ u(t+1) \\ \vdots \\ u(t+h-1) \end{bmatrix} \quad (3)$$

At each time step, the controller implicitly computes $u_h(t)$ to minimize a quadratic cost function over the future horizon:

$$J_h(\hat{x}(t), \hat{w}_h(t), u_h(t)) = \sum_{k=1}^h [e^T(t+k)Q_1e(t+k) + \hat{x}^T(t+k)Q_2\hat{x}(t+k)] \quad (4)$$

$$+ \sum_{k=0}^{h-1} [u^T(t+k)R_u u(t+k)]$$

$$+ \hat{x}^T(t+h)Q_3\hat{x}(t+h)$$

with weighting matrices $Q_i \geq 0$, and $R_u > 0$. The first summation in (4) is a weighted penalty on the residual error e and plant state \hat{x} over the future horizon. The second summation in (4) is a weighted penalty on the norm of the control command over the same horizon; often, $R_u = r_u I$ for a scalar penalty r_u .

The formulation of (4) provides a broad design space for tuning the adaptive controller to specific operational scenarios. For example, Q_1 can be chosen to preferentially penalize particular wavefront modes, while R_u can be chosen to constrain the adaptive control commands for specific channels. Temporal frequency-domain weighting functions may also be included in (4) by appending the plant $\hat{G}(z)$. In Ref. 5, for instance, a high-pass filter was incorporated to avoid frequency-dependent nonlinear behavior of the plant that was not captured by the internal model.

One way to solve (4) is to manipulate it to form a standard least-squares problem:

$$J_h(\hat{x}(t), \hat{w}_h(t), u_h(t)) = \left(\tilde{A} \begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix} + \tilde{B}u_h(t) \right)^T \tilde{Q} \left(\tilde{A} \begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix} + \tilde{B}u_h(t) \right) + u_h^T(t)\tilde{R}u_h(t)$$

where the matrices \tilde{A} , \tilde{B} , \tilde{Q} , and \tilde{R} follow from plugging (2) into (4) and simplifying. The optimal control sequence over the future horizon is then

$$u_h(t) = - \underbrace{\left[\tilde{B}^T \tilde{Q} \tilde{B} + \tilde{R} \right]^{-1}}_{\Gamma} \tilde{B}^T \tilde{Q}^T \tilde{A} \begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix} \quad (5)$$

Following the receding horizon control methodology, the adaptive control command $u(t)$ is the first term in the vector $u_h(t)$ given by (5). This command can be written as a linear function of $\hat{x}(t)$ and $\hat{w}_h(t)$

$$u(t) = -K_x \hat{x}(t) - H_w \hat{w}_h(t) \quad (6)$$

with the matrices K_x and H_w extracted from the corresponding rows of Γ . Note that these matrices can be computed offline directly from the plant model; they do not require any knowledge of the disturbance to be rejected.

A key observation from (6) is that the controller does not need to explicitly predict every point in the disturbance over the horizon. Instead, only the specific linear combination $-H_w\hat{w}_h(t)$ is required, along with an estimate of the plant state. This prediction is performed by the top (unshaded) copy of the FIR filter $F(z)$ in Fig. 1, whose input is $\hat{w}(t)$ and output is a prediction of $-H_w\hat{w}_h(t)$. The filter $U(z)$ then forms the optimal control command by combining (2) and (5) into the state-space system

$$\begin{aligned}\hat{x}(t+1) &= (\hat{A} - \hat{B}K_x)\hat{x}(t) - \hat{B}H_w\hat{w}_h(t) \\ u(t) &= -K_x\hat{x}(t) - H_w\hat{w}_h(t)\end{aligned}\quad (7)$$

The role of the adaptive filter, represented by the shaded block in Fig. 1, is to produce the optimal filter gains for the top copy of $F(z)$. This training occurs using past values of $\hat{w}(t)$ to predict a *delayed* version of $-H_w\hat{w}_h(t)$ via the bottom portion of Fig. 1. First, observe that $H_w\hat{w}_h(t)$ can be written as

$$H_w\hat{w}_h(t) = z^h H(z)\hat{w}(t) \quad (8)$$

where $H(z)$ is an FIR filter with matrix gains extracted from H_w . As a result, the output of the filter $H(z)$ in Fig. 1 is $H_w\hat{w}_h(t)$ delayed by h time steps:

$$H(z)\hat{w}(t) = z^{-h} H_w\hat{w}_h(t) \quad (9)$$

At each time step, the RLS filter generates the matrix FIR gains in the shaded version of $F(z)$ to minimize variance of the signal e_r given by

$$e_r(t) = H(z)\hat{w}(t) + F(z)z^{-h}\hat{w}(t) \quad (10)$$

The corresponding FIR gains are then used in the top copy of $F(z)$ to produce a prediction of $-H_w\hat{w}_h(t)$.

In principle, any adaptive algorithm could be used to update the prediction filter gains to minimize (10), including the classical RLS algorithm with $F(z)$ in standard transversal form. The algorithm used here, however, is the more complex lattice filter implementation presented in Ref. 10, which has superior numerical stability due to internal orthogonalization of the input channels. This implementation is also one of the few adaptive update algorithms that can accommodate multiple inputs and outputs. The result is a fully multichannel adaptive controller in the sense that identified gains of $F(z)$ are not constrained to be diagonal – every element in $\hat{w}(t)$ is used to provide a prediction of every element in $H_w\hat{w}_h(t)$. This enables the controller to exploit all spatial correlations present in the disturbance statistics and plant dynamics over the prediction horizon.

4. IMPLEMENTATION IN THE IOS

The structure of the adaptive control implementation in the IOS is shown in Fig. 2. The classical integrator loop for each active optic follows a standard approach. Each WFS image frame is processed to compute subaperture centroids that are relative to offsets established by non-common path calibration with the scoring camera. These centroids are multiplied by a reconstructor matrix

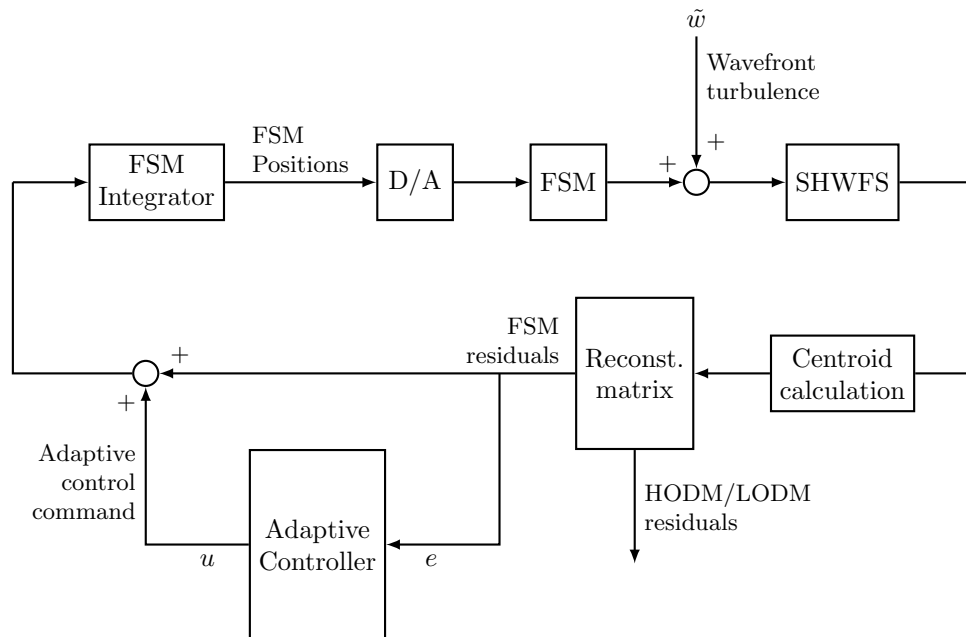


Figure 2: Schematic of the IOS instrument. For clarity, parallel integrator loops for the LODM and HODM, which did not interact with the adaptive controller, are not shown.

to generate wavefront residuals, which are the projections of the wavefront onto the actuator space of the FSM or DMs. If the adaptive controller is not active, each actuator channel is passed through a leaky integrator (i.e. low pass filter) of the form

$$C(z) = -K \frac{z}{z - \alpha} \quad (11)$$

where K is the integrator gain, and α is the leak factor. The resulting actuator “positions” are sent to the driver electronics for each mirror. When this classical loop is closed, $C(z)$ will minimize the wavefront residuals with a fixed disturbance rejection curve characteristic of standard integrator-based controllers, driving the WFS centroids to match the established offsets.

For the results in this paper, we chose to implement the adaptive controller only on the two FSM channels since the real-time latency induced by the RLS lattice filter computation was not initially known. The DM loops were controlled purely by the classical loop, and the reconstructor matrix was constructed to pass all tilt/tip modes to the FSM. Furthermore, while the classical IOS loop is capable of running at WFS rates up to 20 kHz, we chose to perform testing at a 1 kHz WFS rate. However, the two channels were treated as coupled; and hence the implementation of the adaptive controller can be generalized to LODM or HODM channels and higher rates without significant modification.

The adaptive controller interacts with the classical loop by adding the optimal control command to the WFS residuals, and the signals u and e in Fig. 2 correspond to $u(t)$ and $e(t)$ in Fig. 1. This means that the plant $G(z)$ seen by the adaptive controller, which is the transfer function from u to e , includes the closed integrator. It follows that the disturbance $w(t)$ acted upon by the adaptive

controller is the *residual* wavefront error after compensation by the classical loop, i.e.

$$w(t) = S(z)\tilde{w}(t) \quad (12)$$

where $S(z)$ is the disturbance rejection transfer function of the leaky integrator.

Because the adaptive controller assumes a LTI model of the instrument dynamics, including the closed integrator as part of the AO plant can significantly reduce modeling error due to system nonlinearities. First, the integrator helps constrain the WFS centroids to the linear regime of the SHWFS S-curve, and is thus useful in “bootstrapping” the adaptive controller in high turbulence or low light conditions. Second, using the integrator as an inner loop helps mitigate FSM hysteresis, facilitating a linear model of the FSM dynamics.

The “Adaptive Controller” block in Fig. 2 contains all the controller components in Fig. 1 mapping e to u . Because the IOS RTC is optimized for fast matrix multiplication, the LTI blocks in Fig. 1 were implemented using a method developed in Ref. 3, where the state-space form of each system is represented as a matrix multiplication combined with a single-frame delay. For example, for $U(z)$ given by (7), a block matrix, U_{ss} , is defined such that

$$\begin{bmatrix} \hat{x}(t+1) \\ u(t) \end{bmatrix} = \underbrace{\begin{bmatrix} (\hat{A} - \hat{B}K_x) & -\hat{B} \\ -K_x & -I \end{bmatrix}}_{U_{ss}} \begin{bmatrix} \hat{x}(t) \\ H_w \hat{w}(t) \end{bmatrix} \quad (13)$$

At each time step, the elements of $[\hat{x}(t+1) \ u(t)]^T$ corresponding to the state update are “demuxed” from the output vector and sent to a single-step delay to provide $\hat{x}(t)$ at the next time step. Similar matrices were defined for $H(z)$, after it had been formatted in controllable canonical form, and $\hat{G}(z)$. We found this to be an efficient method for implementing arbitrary state-space systems using the existing IOS RTC architecture. The RLS lattice filter algorithm for $F(z)$, which included the h -step delay and produced the prediction of $H_w \hat{w}_h(t)$, was incorporated directly into the RTC software from C code used in Ref. 6. For the two coupled FSM channels, unit testing indicated that these additions added minimal latency to the overall loop.

5. IDENTIFICATION OF A PLANT MODEL

Accurate characterization of the AO plant is essential for achieving robust performance with the adaptive controller. Many instruments employing predictive control operate at frame rates much slower than the physical bandwidth of the DMs used for compensation. As such, it is common to largely ignore mirror dynamics and only capture system latencies and the WFS integration time as a series of delays.¹¹ However, because the IOS RTC can operate at very high WFS frame rates, and the closed integrator is included in the plant, the system dynamics cannot be adequately modeled in the manner. Instead, we employed a data-driven subspace identification technique to compute a multichannel LTI model directly, rather than rely on assumed dynamics for the components of Fig. 2.

Traditional subspace methods use white noise input sequences to characterize the matrices of a multichannel state-space model.¹² However, reliably injecting such sequences is not currently possible in the IOS RTC. Instead, we applied step inputs (adaptive control command “pokes”) to the signal u in Fig. 2, with the integrator loop closed at an operational gain, and simultaneously recorded the residual e . This process was performed with the ATS off, and data was averaged over

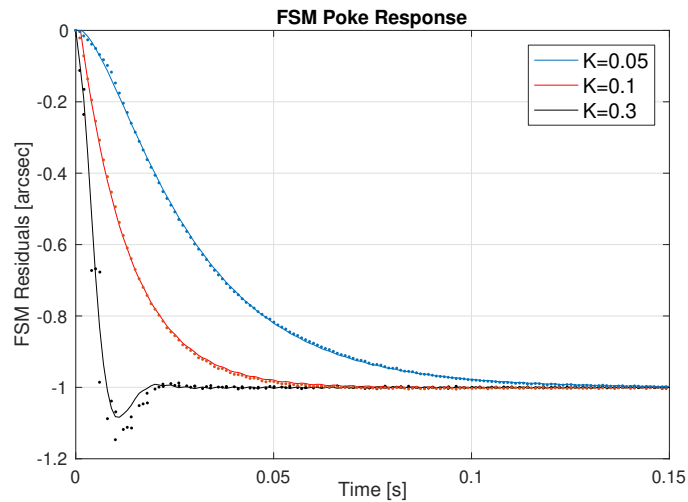


Figure 3: Diagonal elements of the step response of the identified transfer matrix $\hat{G}(z)$, indicated by the solid lines, for several values of the integrator gain K . The residual poke data $\tilde{g}(t)$ used to identify each system is indicated by the dots.

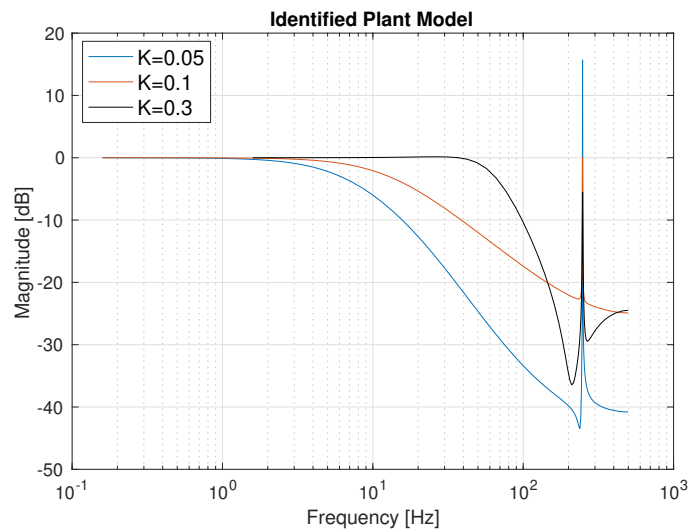


Figure 4: Magnitude response of the diagonal terms of the identified plant model $\hat{G}(z)$.

multiple trials such that $w \approx 0$. Each time series was synchronized to the time of the applied poke so accurate phase data could be captured. The resulting residual time series was the step response of $G(z)$, i.e. the impulse response of $\tilde{G}(z) = zG(z)$.

We then applied the Ho-Kalman realization algorithm¹³ to construct a balanced (observable and controllable) state-space model for $\tilde{G}(z)$, as follows. Let $\{\tilde{g}(0), \tilde{g}(1), \dots, \tilde{g}(n)\}$ be the measured impulse response sequence of $\tilde{G}(z)$, formed by averaging multiple instances of the residual time series $e(t)$. Suppose the state-space matrices of $\tilde{G}(z)$ are $(\tilde{A}, \tilde{B}, \tilde{C}, \tilde{D})$. Then the impulse response

terms (i.e. the Markov parameters) are given by

$$\tilde{g}(t) = \begin{cases} \tilde{D}, & t = 0 \\ \tilde{C}\tilde{A}^{t-1}\tilde{B}, & t > 0 \end{cases} \quad (14)$$

where $t = 0$ corresponds to the time of the applied poke. These terms can be organized into a Hankel matrix that is the product of the observability and controllability matrices of $\tilde{G}(z)$

$$H_o = \begin{bmatrix} \tilde{g}(1) & \tilde{g}(2) & \tilde{g}(3) & \cdots & \tilde{g}(q) \\ \tilde{g}(2) & \tilde{g}(3) & \tilde{g}(4) & \cdots & \tilde{g}(q+1) \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \tilde{g}(q) & \tilde{g}(q+1) & \tilde{g}(q+2) & \cdots & \tilde{g}(2q-1) \end{bmatrix} \quad (15)$$

$$= \begin{bmatrix} CB & CAB & CA^2B & \cdots & CA^{q-1}B \\ CAB & CA^2B & CA^3B & \cdots & CA^qB \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ CA^{q-1}B & CA^qB & CA^{q+1}B & \cdots & CA^{2q-2}B \end{bmatrix} \quad (16)$$

$$= \underbrace{\begin{bmatrix} C \\ CA \\ \vdots \\ CA^{q-1} \end{bmatrix}}_{\mathcal{O}(q)} \underbrace{\begin{bmatrix} B & AB & \cdots & A^{q-1}B \end{bmatrix}}_{\mathcal{P}(q)} \quad (17)$$

where q is the order of the state-space model determined by the rank of H_o . Let the singular value decomposition of H_o be given by $H_o = U\Sigma V^T$. Estimates of the matrices $\mathcal{O}(q)$ and $\mathcal{P}(q)$ are then, up to a similarity transformation, given by

$$\hat{\mathcal{O}}(q) = U\Sigma^{1/2}, \quad \hat{\mathcal{P}}(q) = \Sigma^{1/2}V^T \quad (18)$$

It follows that the \tilde{B} matrix of $\tilde{G}(z)$ is given by the first n_u columns of $\hat{\mathcal{P}}(q)$, where n_u is the dimension of u , and \tilde{C} by the first q rows of $\hat{\mathcal{O}}(q)$. An estimate of the matrix \tilde{A} can be found by populating a second Hankel matrix with additional samples of the sequence $\tilde{g}(t)$:

$$H_1 = \begin{bmatrix} \tilde{g}(2) & \tilde{g}(3) & \tilde{g}(4) & \cdots & \tilde{g}(q+1) \\ \tilde{g}(3) & \tilde{g}(4) & \tilde{g}(5) & \cdots & \tilde{g}(q+2) \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \tilde{g}(q+1) & \tilde{g}(q+2) & \tilde{g}(q+3) & \cdots & \tilde{g}(2q) \end{bmatrix} \quad (19)$$

$$= \begin{bmatrix} CAB & CA^2B & CA^3B & \cdots & CA^qB \\ CA^2B & CA^3B & CA^4B & \cdots & CA^{q+1}B \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ CA^qB & CA^{q+1}B & CA^{q+2}B & \cdots & CA^{2q-1}B \end{bmatrix} = \mathcal{O}(q)\tilde{A}\mathcal{P}(q) \quad (20)$$

Using (18) and solving for \tilde{A} we have

$$\tilde{A} = \Sigma^{-1/2}U^T H_1 V \Sigma^{-1/2} \quad (21)$$

With $\tilde{G}(z)$ identified, the plant model $\hat{G}(z)$ in (2) is given by $\hat{G}(z) = z^{-1}\tilde{G}(z)$. Note that this method potentially amplifies high frequency noise in the residual signal e , thus care must be taken to ensure sufficient trials are used in computing the impulse response sequence $\tilde{g}(t)$. Alternatively, smoothing filters can be employed to reduce the influence of ambient disturbances.

The dotted lines in Fig. 3 show one channel (in this case, the FSM x -direction) of the time-averaged, residual sequences $\tilde{g}(t)$ in response to a step input applied to the same channel in $u(t)$. Due to inherent channel orthogonality in the reconstructor matrix, the response for the un-poked channel was essentially zero for the same input. Since the leaky integrator controller given by (11) is included in the plant, these responses vary depending on the value of the integrator gain K . Each time series was passed through the Ho-Kalman algorithm described above to produce a balanced realization of the plant model for each integrator gain, the step responses of which are shown as the solid lines in Fig. 3. A similar procedure produced estimates of the plant in the y -direction, yielding a diagonal transfer matrix $\hat{G}(z)$.

For each integrator gain, 3 seconds of 1 kHz data was used in the estimate. The rank of the matrix H_o was chosen to yield a realizations of $\hat{G}(z)$ with 5 internal states, which appeared adequate to capture the gross features of the time series data without overfitting measurement noise. Figure 4 shows the magnitude response of the estimated plant models. At the time the results in this paper were recorded, the 250 Hz spike was included in the plant response, however subsequent analysis suggests this feature was due to unfiltered ambient disturbances and not FSM dynamics. Future versions of the realization algorithm will investigate filtering out this peak from the measured time series data.

6. ADAPTIVE CONTROL RESULTS

Table 1: ATS turbulence settings used in the adaptive control demonstration.

Parameter	Weak value	Strong value
Wind speed	5 m/s	13 m/s
r_o	72.2 mm	72.2 mm
f_g	90.82 Hz	130.4 Hz

With the IOS plant identified for a variety of integrator gains, testing commenced with the ATS active using the settings listed in Table 1. Testing for the results here was performed relatively early in the IOS integration process, and a number of IOS hardware features were lacking at that preliminary stage. First, the ATS did not produce tilt/tip wavefronts with spectra that matched the applied r_o and Greenwood frequency (f_g) values, thus the applied wind speeds were quite high to produce disturbances with appreciable statistics. At the time of the test, the scoring camera had not yet been aligned using non-common-path calibration, and the SHWFS utilized nominal centroid offsets generated by manual optimization. Finally, the HODM was replaced with a flat surrogate mirror. We plan on addressing these issues in future versions of this test, and do not believe they impact the validity of the adaptive control results.

Before applying the adaptive controller, the classical loop was closed with increasing values of K to determine the integrator gain that best minimized rms tilt/tip residuals with the ATS active. This sweep yielded an optimal value of approximately $K = 0.3$, which balanced low-frequency turbulence suppression with amplification of disturbances above the classical bandwidth of 10 Hz.

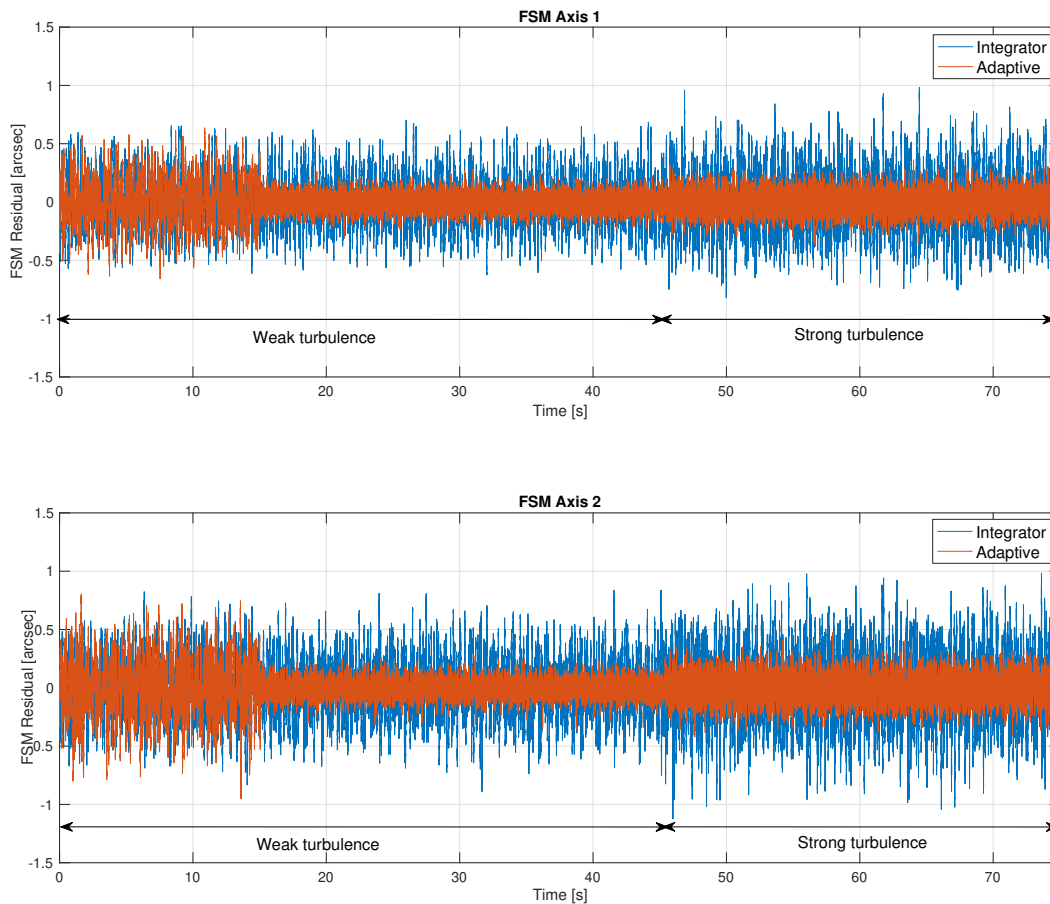


Figure 5: Residual time series data comparing performance with the classical leaky integrator and the adaptive controller for the two FSM channels. The turbulence strength and direction was changed at approximately 45 seconds. The open loop disturbance response is omitted for clarity.

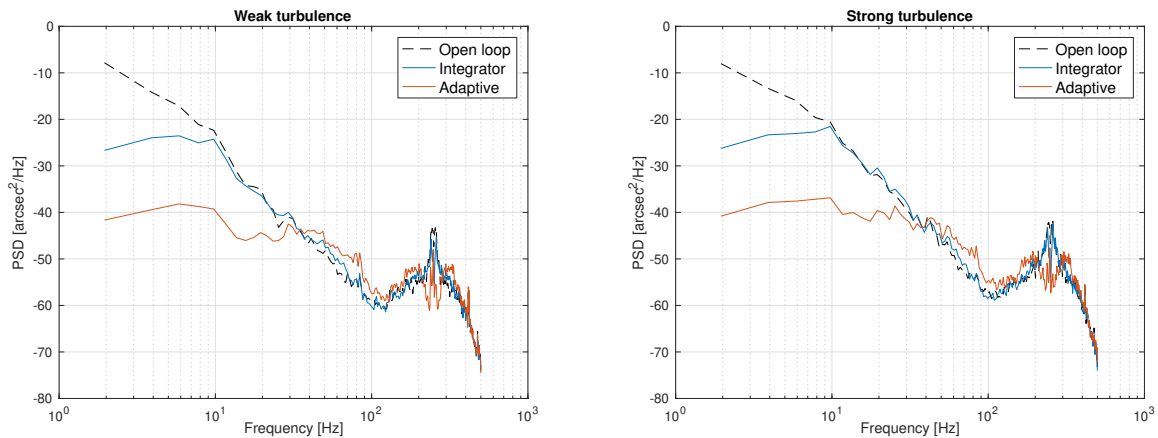


Figure 6: Power spectra for one FSM axis for the weak and strong turbulence regimes.

This value was also chosen for the closed integrator when the adaptive controller was active, thus the plant model corresponding to $K = 0.3$ in Fig. 4 was selected for $\hat{G}(z)$. Note, however, that the adaptive controller does not require the plant to be internally optimal. With the plant identified, the transfer functions $U(z)$ and $H(z)$ in Fig. 1 were computed using (5) and the parameters shown in Table 2.

Both the adaptive and integrator controllers were then closed using an ATS scenario containing weak and strong turbulence regimes, as described in Table 1. After they were activated, no modifications were made either the adaptive or integrator loops for the remainder of the test.

Table 2: Adaptive and integrator controller parameter values used in the demonstration.

Parameter	Symbol	Value
SHWFS frame rate	–	1 kHz
FIR filter order	–	10
Prediction horizon	h	10
Command penalty	r_u	0.05
Residual penalty	Q_1	I
State penalty	Q_2	$\hat{C}^T \hat{C}$
Terminal state penalty	Q_3	$\hat{C}^T \hat{C}$
Integrator gain	K	0.3
Integrator leak	α	0.01

FSM residual data for both controllers are shown in Fig. 5. Once the adaptive loop is closed after 15 seconds, it significantly outperforms the integrator during the weak regime, achieving a mean rms error of 0.06 arcsec compared to 0.23 arcsec for the integrator over this period. As the wind speed is increased after 45 seconds, the adaptive controller autonomously tracks the evolving turbulence statistics, and produces a mean rms error of 0.09 arcsec, compared to 0.29 arcsec for the integrator. The adaptive control residuals do increase during this strong turbulence period, however we believe this to be due to artificial saturation limits imposed on the signal u .

Power spectra of the residuals, shown in Fig. 6 for a single FSM axis over the two wind speed

periods, are revealing. In both turbulence regimes the adaptive controller clearly reduces the residual error over a wider range of frequencies. While the integrator disturbance rejection bandwidth is theoretically limited to approximately 10 Hz, the adaptive controller provides significant mitigation of the high frequency peaks around 250 Hz. As with the time domain results, we believe relaxing the saturation limitations on u would allow the adaptive controller to continue to “flatten” the residual power spectrum to the white noise floor.

7. CONCLUSIONS

This demonstration validated several new algorithms and analysis techniques that are applicable to future AO systems. First, the Ho-Kalman realization algorithm was successfully applied to identify a plant model that reflected observed dynamics with the classical integrator loop closed. As future systems possess faster frame rates, mirror dynamics and their accurate characterization will play an important role in developing model-based control algorithms. The software development process utilized a convenient method for implementing the adaptive controller (and other advanced linear control algorithms) in a real-time environment, harnessing the existing matrix multiplication capabilities of the RTC to realize multiple state-space objects with minimal latency impact.

The promising experimental results demonstrate that the adaptive controller can provide significant performance benefits over a standard integrator. An important point is that no explicit open loop disturbance data was used to train the adaptive controller, which successfully tracked statistics both when the loop was initially closed, and when the turbulence statistics were altered. While the results shown here occurred over a relatively limited time period, long-duration tests indicated that performance with the adaptive controller did not degrade over time. It is reasonable to suspect that the adaptive controller would maintain its performance during changing on-sky conditions without significant operator intervention. We expect that employing similar adaptive controllers could significantly increase the performance margin of future systems, enabling lower-cost hardware and adding robustness in countering statistically non-stationary environmental factors.

As noted, the demonstration presented here was performed during the early integration and test phase of the IOS instrument. In future testing, we plan on demonstrating the performance of the adaptive controller against on-sky stellar targets, as well as extending the algorithm to several LODM channels.

ACKNOWLEDGMENTS

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration and funded through the internal Research and Technology Development program.

REFERENCES

- [1] Roberts Jr, L. C., Block, G. L., Fregoso, S. F., Herzog, H., Meeker, S. R., Roberts, J. E., Tesch, J., Truong, T. N., Rodriguez, J. D., and Bechter, A., “First results from the adaptive optics system from LCRD’s Optical Ground Station One,” Advanced Maui Optical and Space Surveillance Technologies Conference (2018).
- [2] Petit, C., Fusco, T., Fedrigo, E., Conan, J. M., Kulcsár, C., and Raynaud, H. F., “Optimisation of the control laws for the SPHERE XAO system,” in [*Proc. SPIE 7015, Adaptive Optics Systems*], **7015** (2008).

- [3] Tesch, J., Truong, T., Burruss, R., and Gibson, S., “On-sky demonstration of optimal control for adaptive optics at Palomar Observatory,” *Optics Letters* **40**(7), 1575–1578 (2015).
- [4] Doyle, J., “Guaranteed margins for LQG regulators,” *IEEE Transactions on Automatic Control* **23**(4), 756–757 (1978).
- [5] Tesch, J., Gibson, S., and Verhaegen, M., “Receding-horizon adaptive control of aero-optical wavefronts,” *Optical Engineering* **52**(7), 071406 (2013).
- [6] Tsuchiya, N., Gibson, S., Tsao, T., and Verhaegen, M., “Receding-horizon adaptive control of laser beam jitter,” *IEEE/ASME Transactions on Mechatronics* **21**(1), 227–237 (2016).
- [7] Roberts Jr, L. C., Block, G., Fregoso, S., Herzog, H., Meeker, S. R., Roberts, J. E., Rodriguez, J., Tesch, J., and Truong, T., “A laser communication adaptive optics system as a testbed for extreme adaptive optics,” in [*Adaptive Optics Systems VI*], **10703** (2018).
- [8] Truong, T., Bouchez, A., Burruss, R., Dekany, R., Guiwits, S., Roberts, J., Shelton, J. C., and Troy, M., “Design and implementation of the PALM-3000 real-time control system,” in [*Proc. SPIE, Adaptive Optics Systems III*], 84472F–84472F, SPIE (2012).
- [9] Kwon, W. and Han, S., [*Receding Horizon Control*], Springer-Verlag London (2005).
- [10] Jiang, S. and Gibson, S., “An unwindowed multichannel lattice filter with orthogonal channels,” *IEEE Transactions on Signal Processing* **43**(12), 2831–2842 (1995).
- [11] Poyneer, L. and Véran, J.-P., “Predictive wavefront control for adaptive optics with arbitrary control loop delays,” *J. Opt. Soc. Am. A* **25**(7), 1486–1496 (2008).
- [12] Katayama, T., [*Subspace Methods for System Identification*], Springer-Verlag London (2005).
- [13] Verhaegen, M. and Verdult, V., [*Filtering and System Identification, a Least Squares Approach*], Cambridge University Press (2007).