

Frequency-Weighted Minimum-Variance Adaptive Control of Laser Beam Jitter

Néstor O. Pérez-Arancibia, *Member, IEEE*, James S. Gibson, and Tsu-Chin Tsao, *Senior Member, IEEE*

Abstract—This paper introduces a frequency-weighting method for adaptive disturbance rejection. The method constrains the high-frequency gain of the prediction filter in a minimum-variance adaptive controller. In an experimental application, the method is used to control a microelectromechanical system fast steering mirror to suppress laser beam jitter. The paper analyzes the effect of sensor noise on the performance of the adaptive control system and demonstrates that sufficient levels of high-frequency noise and/or disturbance combined with control saturation produce a spiking phenomenon in the output error. The experimental results show that the frequency weighting eliminates the spiking.

Index Terms—Frequency weighting, jitter suppression, laser beam control, minimum-variance adaptive control, microelectromechanical systems (MEMS) fast steering mirrors.

I. INTRODUCTION

IN THE emerging fields of laser communications and high-energy laser systems, many applications require precise pointing of laser beams subjected to disturbances including atmospheric turbulence and vibration of optical benches and components. Vibration-induced jitter typically consists of multiple narrow bandwidths produced by vibration modes of the structure supporting the optical system, while turbulence-induced jitter usually has broader bandwidths [1], [2]. Also, as in this paper, the fast steering mirrors used for actuation may have lightly damped elastic modes that create another source of beam jitter. This complex combination of disturbances, which often change with time, necessitates the use of controllers capable of rejecting disturbances with broadband spectra.

The well-known waterbed constraint described by the Bode integral sensitivity theorem [3] makes it impossible to achieve the needed broadband jitter rejection with linear time-invariant (LTI) controllers. Recent literature on control of jitter in laser beams has introduced adaptive control methods that reject jitter over much greater bandwidths than those achieved by LTI feedback control. For laser beams, adaptive controllers based on least-mean-squares (LMS) adaptive filters are presented in [4]–[8], and adaptive controllers based on recursive least-squares (RLS) filters are reported in [9]–[14]. Adaptive control has been used for disturbance rejection in other applica-

tions, including the reduction of the read–write head position error in computer disk drives [15]–[20]. While LMS algorithms are simple and computationally economical, the more complex RLS algorithms achieve faster convergence and exact minimum-variance steady-state performance.

The hardware and geometry of beam steering applications dictate saturation limits on the magnitude of the control commands. Any fast steering mirrors, but especially the micromirrors used here and in various optical communications systems, can be damaged if the control commands are too large. Also, because small angles of steering-mirror deflection can produce large beam displacements at remote receivers and targets, the mirror rotation must be limited so that the laser beam hits all reflecting surfaces and sensors in the optical path. Either type of constraint can be binding in field applications.

A common characteristic of minimum-variance controllers, adaptive or not, is that they amplify low-level high-frequency noise while minimizing the mean-square values of the output errors. This results from the fact that such controllers must have large high-frequency gains to predict broadband disturbances. In some applications, this amplification results in the fast steering mirrors being driven with potentially damaging levels of high-frequency power, even if the amplitude limit on the driving current is not reached. But more serious is a spiking phenomenon, which is illustrated in this paper, in output error produced by the combination of control signal saturation and amplified high-frequency noise.

This paper presents a method for incorporating frequency weighting in the adaptive control loop to constrain the high-frequency gain of the adaptive filter that generates the adaptive control command. The frequency responses of both experimental and theoretically computed steady-state filters demonstrate that the frequency weighting reduces the high-frequency gains. Experimental results show that the frequency weighting indeed reduces the amplification of high-frequency noise and eliminates the spikes in the output error. Furthermore, because of the high sampling rate and high adaptive filter order used here, the experimental results represent a substantial improvement in performance over results achieved previously with the same microelectromechanical systems (MEMS) steering mirror [13].

This paper is organized as follows. Section II describes the experimental setup and explains some practical issues relating to the real-time implementation of the controller. Section III describes modeling and identification of the open-loop and closed-loop plants. Section IV discusses the disturbance spectrum used in the experiments and the disturbance and sensor noise models used in calculating the theoretical steady-state performance of the adaptive control system. Section V describes the

Manuscript received January 24, 2008; revised January 25, 2009. First published April 24, 2009; current version published June 17, 2009. Recommended by Technical Editor N. Jalili. This work was supported by the U.S. Air Force Office of Scientific Research under Grant F49620-02-01-0319 and Grant F49620-03-1-0234.

The authors are with the Mechanical and Aerospace Engineering Department, University of California, Los Angeles, CA 90095-1597 USA (e-mail: nestor@seas.ucla.edu; gibson@ucla.edu; tsaoc@seas.ucla.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMECH.2009.2017532

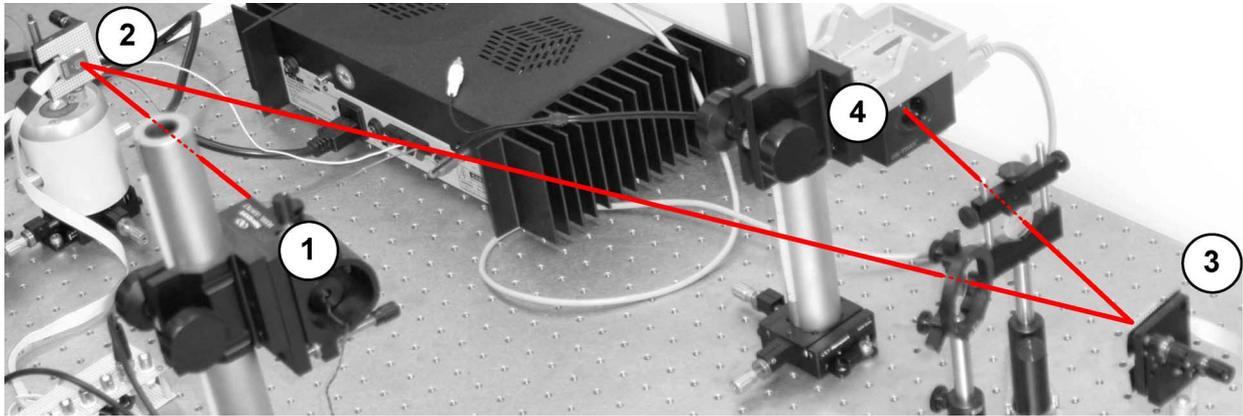


Fig. 1. Photograph of the experiment with optical path shown. Position ①: laser source. Position ②: Fast steering mirror FSM-C (control actuator) mounted on shaker. Position ③: Fast steering mirror FSM-D (disturbance actuator). Position ④: ON-TRAK optical position sensor. FSM-C and FSM-D are TI Incorporated Analog MicroMirrors.

frequency-weighted adaptive control loop and discusses the method for computing the theoretically predicted steady-state performance. The adaptive controller does not require the disturbance and noise models discussed in Section IV, but the theoretical calculations do. Subsequent comparison between theoretical and experimental performance is valuable for evaluating the adaptive control design. Section VI discusses two motivations for the frequency weighting in this paper: the spikes in the output error that sometimes occur without frequency weighting and robustness with respect to high-frequency plant modeling error.

A simulation in Section VI-B demonstrates the fundamental nature of the spikes in the output error that are observed in experimental results in Section VII. The experiments discussed in Section VII demonstrate the capability of the frequency-weighted adaptive controller to suppress broadband jitter in the presence of sensor noise concentrated at high frequencies and in the presence of white sensor noise.

II. DESCRIPTION OF THE EXPERIMENT

Fig. 1 shows the optical bench and the laser path in the experiment. The laser beam leaves the source at position ① in Fig. 1, reflects off the fast steering mirror FSM-C at position ②, then reflects off the fast steering mirror FSM-D at position ③, and finally, reaches the optical position sensor at position ④. Two lenses in the optical path focus the beam on FSM-D and the sensor.

The mirrors FSM-C and FSM-D are identical Texas Instruments (TI) MEMS mirrors used in laser communications for commercial and defense applications. FSM-C is the control actuator, and FSM-D is used to add jitter (i.e., disturbance) of specified bandwidths to the laser beam on both mirror axes. In addition to FSM-D, a second source of jitter is the shaker on which the control actuator FSM-C is mounted. The shaker adds vertical jitter with narrow frequency bands to create the effect of a vibrating platform.

The commanded rotations of the fast steering mirrors are produced by electromagnetic fields with opposing directions. These fields are created by coils with currents commanded by the con-

trol and disturbance computers. The mirrors have a rotation range of $\pm 5^\circ$. The reflecting area of the mirrors is 9 mm^2 . The optoelectronic position sensor at the end of the beam path generates two analog output voltages proportional to the coordinates of the laser beam centroid. In the sensor, quad photodetectors capture the laser intensity and generate current outputs, which are converted to voltage and amplified by an operational amplifier. Further electronic processing of these signals yields two final signals, which are the estimates of the centroid coordinates independent of light intensity.

The output error in the control problem is the pair of sensor measurements, which are the coordinates of the laser beam spot on the sensor. These measurements, in the form of voltages, go to computer 1, which runs both the LTI feedback and the adaptive controllers and sends actuator commands in volts to FSM-C. Computer 2 sends disturbance commands to FSM-D and the shaker. Computers 1 and 2 are PCs running xPC Target, and the sample-and-hold rate for the control system is 5 KHz.

The only physical measurements used by the adaptive and feedback controllers are the two signals from the ON-TRAK sensor, which are essentially noise free. To make the experiments more realistic and challenging, artificial sensor noise is added to the measurements inside computer 1.

III. PLANT MODELS

A. Open-Loop Plant

The open-loop discrete-time plant $P(z)$ is the transfer function that maps the two-channel digital control command to the sampled two-channel output of the optical position sensor. Thus, $P(z)$ is the two-input–two-output digital transfer function for the fast steering mirror FSM-C with a gain determined by the optical position sensor and the laser path length. Output channels 1 and 2 represent horizontal and vertical displacements, respectively, of the beam; input channels 1 and 2 represent commands that drive FSM-C about its vertical and horizontal axes, respectively. Input–output data from open-loop experiments showed negligible coupling between the two channels of the open-loop plant; so henceforth, all discussion and control designs assume

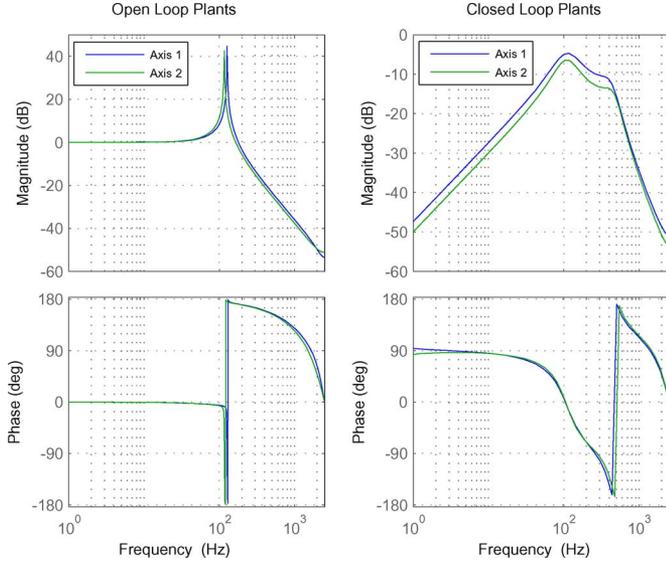


Fig. 2. Bode plots for identified transfer functions \hat{P} (left) and \hat{G} (right).

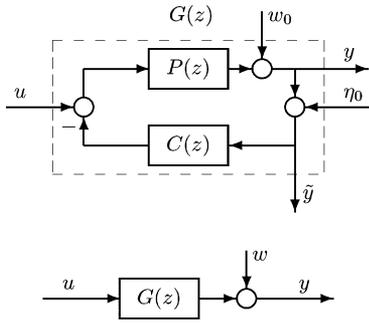


Fig. 3. Top: LTI feedback control system: $P(z)$ = open-loop plant, $C(z)$ = classical LTI feedback controller, y = output error = position of laser spot on optical sensor, u = adaptive control command, w_0 = jitter on laser beam, η_0 = sensor noise, \tilde{y} = noisy measurement fed back to control loops. Bottom: Equivalent system $G(z) = y/u$, w = combined effects of the disturbance w_0 and the sensor noise η_0 with the LTI feedback loop closed.

that $P(z)$ has two uncoupled channels. The subspace system identification algorithm *n4sid* in MATLAB [21] was used to identify an estimate $\hat{P}(z)$ of $P(z)$ from 30 000 samples of input–output data. The uncoupled channels of $\hat{P}(z)$ were identified separately. Fig. 2 shows the Bode plots of $\hat{P}(z)$.

B. Closed-Loop System With LTI Feedback Control

The LTI feedback control loop is shown in Fig. 3. The classical digital controller $C(z)$ consists of an integrator and a notch filter. The integrator gain and notch parameters were tuned to maximize the disturbance-rejection bandwidth. The input u in Fig. 3 is the adaptive control command, and the output y is the position of the laser spot on the optical position sensor. The signal w_0 represents the combined disturbances acting on the system, and the signal η_0 is the sensor noise added to the measured output y to generate the noisy measurement \tilde{y} .

The closed-loop LTI system in the top diagram in Fig. 3, without the output \tilde{y} , can be represented as in the bottom diagram

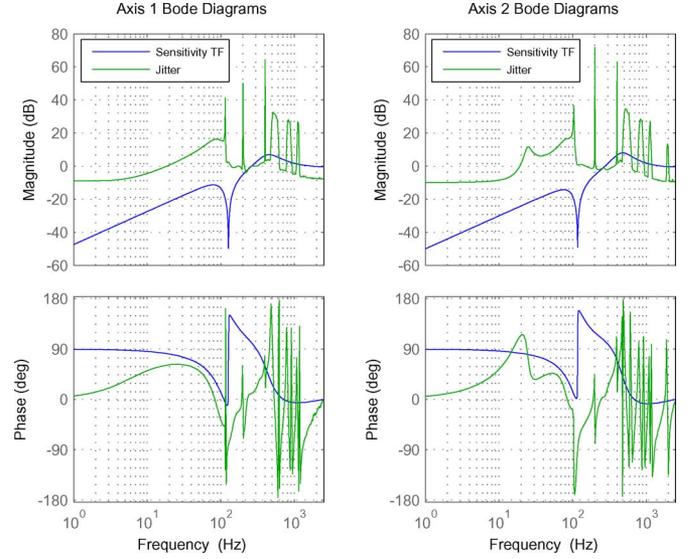


Fig. 4. Bode plots for estimated sensitivity transfer function \hat{S} and identified jitter model W_1 (discussed in Section IV).

in Fig. 3 with

$$G = P(I + PC)^{-1} \quad (1)$$

$$w = (I + PC)^{-1}(w_0 - PC\eta_0). \quad (2)$$

The closed-loop plant dynamics were reidentified by the aforementioned subspace identification method as $\hat{G}(z)$, which has two uncoupled fourth-order channels. Each channel of $\hat{G}(z)$ is stable with one zero outside the unit circle. The Bode plots for $\hat{G}(z)$ are shown in the second column of Fig. 2.

The two-channel output sensitivity transfer function for the closed-loop LTI system is

$$S = \frac{y}{w_0} = (I + PC)^{-1}. \quad (3)$$

An estimate of S , computed as $\hat{S} = (I + \hat{P}\hat{C})^{-1}$, is shown in Fig. 4, which shows that the error-rejection bandwidth is about 200 Hz. Because the gain of the controller $C(z)$ was chosen to maximize the error-rejection bandwidth, the LTI loop amplifies high-frequency disturbance above about 300 Hz.

IV. DISTURBANCE AND IDENTIFIED DISTURBANCE MODELS

The disturbance w_0 is created by the disturbance mirror FSM-D and the shaker. The MEMS mirror FSM-D produces jitter with multiple bandwidths and the shaker produces jitter with two narrow bandwidths, as described in Table I. Passing white noise through filters with the bandwidths in Table I produces the disturbance commands to FSM-D. The sensor noise η_0 , generated by the real-time computer, is the output of a stable linear filter N_0 driven by a white noise sequence ε_η . (For white sensor noise, $N_0 = 1$.) The sequences w_0 and η_0 are independent. Subsequent analysis uses the sequences

$$w_1 = (I + PC)^{-1}w_0 \quad \text{and} \quad \eta_1 = (I + PC)^{-1}\eta_0. \quad (4)$$

TABLE I
JITTER BANDWIDTHS

	Axis 1	Axis 2
	0–100 Hz	0–100 Hz
Bandwidths of jitter commands to FSM-D	500–600 Hz 800–900 Hz 1100–1150 Hz	500–600 Hz 800–900 Hz 1100–1150 Hz 1900–2000 Hz
Frequencies of sinusoidal commands to shaker		200 Hz 400 Hz

Each of the two-input–two-output transfer functions P , C , and \hat{G} has uncoupled channels, and each of the sequences in (4) is modeled as having two uncorrelated channels.

The adaptive controller implicitly identifies certain statistics of the sequence $w + \eta_0$ but otherwise requires no information about any of the sequences in (2) and (4). For theoretical analysis of the performance of the adaptive controller, this paper assumes that all sequences in (4) are stationary with zero mean. Also, the analysis assumes the disturbance models

$$w_0 = W_0 \varepsilon_w \quad \eta_0 = N_0 \varepsilon_\eta \quad w_1 = W_1 \varepsilon_w \quad \eta_1 = N_1 \varepsilon_\eta \quad (5)$$

where W_0 , N_0 , W_1 , and N_1 are finite-dimensional stable LTI filters, and the sequences ε_w and ε_η are independent, stationary, white, and zero-mean, with

$$E\{\varepsilon_w \varepsilon_w^T\} = E\{\varepsilon_\eta \varepsilon_\eta^T\} = I (=2 \times 2 \text{ identity matrix}). \quad (6)$$

It must be emphasized that none of the disturbance models are used by the adaptive controller. The disturbance models W_1 and N_1 are required for computing the theoretical optimal controller and corresponding performance to which the adaptive controller should converge. These theoretical results are compared in Section VII to the experimental performance of the adaptive controller.

The disturbance models W_1 and N_1 were identified from experimental output data taken with the LTI feedback loop closed and with the adaptive control signal $u = 0$. For identification of W_1 , the sensor noise η_0 was zero, and for identification of N_1 , the disturbance w_0 was zero. This identification, performed with the *n4sid* system identification algorithm in MATLAB, produced state-space models in innovations form. Each of the identified models W_1 and N_1 has 64 states for each channel. In principle, W_1 and N_1 can be constructed as $W_1 = (I + PC)^{-1}W_0$ and $N_1 = (I + PC)^{-1}N_0$, but since this construction would be based on estimated models for P and W_0 , the disturbance models W_1 and N_1 were identified directly.

The identified disturbance model W_1 is shown in Fig. 4, where the left and right Bode plots correspond to axis 1 and axis 2, respectively. These Bode plots represent the statistics of the disturbance to which the laser beam is subjected. Fig. 4 compares W_1 to the LTI output-disturbance sensitivity function. These plots show that much of the jitter is far beyond the error-rejection bandwidth of the LTI feedback controller.

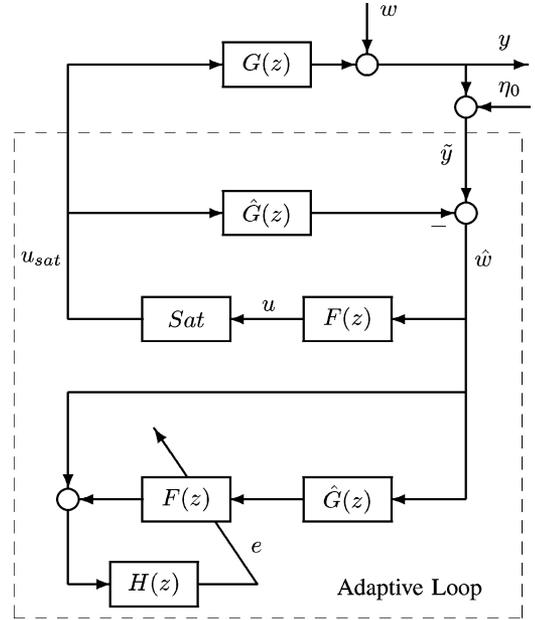


Fig. 5. Block diagram of the adaptive control system.

V. FREQUENCY-WEIGHTED ADAPTIVE CONTROL

A. Adaptive Control Loop

The class of adaptive controllers considered in this paper is illustrated in Fig. 5. Since the two channels of the plant are modeled as uncoupled and each control loop has two independent channels, the discussion in this section is simplified by considering a single channel of the beam steering system. Thus, in this section, all signals are scalar sequences and all transfer functions are single-input–single-output (SISO). The equations and block diagrams are the same for each channel. The transfer functions $G(z)$ and $\hat{G}(z)$ (discussed in Section III) and the frequency-weighting filter $H(z)$ are assumed to be stable.

For each channel, the adaptive controller has two copies of the adaptive filter $F(z)$. The gains are updated in the copy of $F(z)$ in the bottom portion of Fig. 5 to minimize the rms value of the tuning signal e . The other copy of $F(z)$ uses the same gains to generate the adaptive control signal u . The block labeled *Sat* in Fig. 5 saturates the adaptive control command so that $|u_{sat}| \leq 1$ V. The saturation limit is chosen to satisfy two requirements: that the driving current for FSM-C not exceed the specified limit, and that the laser beam not veer off the reflecting surface of FSM-D.

The adaptive filter $F(z)$ has finite-impulse response (FIR), so that it can be written as

$$F(z) = \sum_{l=0}^L f_l z^{-l} \quad (7)$$

with gains f_l and order L . However, the adaptive filter here has the lattice realization as in [22], rather than the realization in (7). The lattice realization has an order-recursive structure discussed in [13], [14], and [22], which allows variable-order adaptive control. The lattice filter generates adaptive control commands of all orders $n \leq L$. During adaptation, lattice-filter

orders $n < L$ are used, with the order increasing to the maximum order L in steady state. For the experiments and simulations in this paper, $L = 80$.

For high orders of F , the RLS lattice filter is more computationally efficient and numerically stable than the classical RLS algorithm, which has been used for low-order filters in various adaptive control applications other than control of laser beams (e.g., [18]–[20]). Lattice-filter-based variable-order adaptive control (without frequency weighting) and the improved transient response that it provides are discussed in detail in [13].

The lattice gains, or reflection coefficients, are updated by the RLS lattice algorithm in [22] to minimize the rms value of the tuning signal e in Fig. 5. Thus, the least squares criterion for the adaptive filter is

$$J_e(F, H) = \text{rms}(e) = \sqrt{E\{e^2\}}. \quad (8)$$

In evaluation of the performance of the adaptive controller, the error signal of primary interest is the laser beam position y in Fig. 5. This signal is available to the control loops only when there is zero sensor noise. However, in the experiments described in this paper, y can be measured (though it is not used by the control loops), and Section VII presents values for the performance index

$$J_y(F, H) = \text{rms}(y) = \sqrt{E\{y^2\}}. \quad (9)$$

B. Theoretical Minimum-Variance Performance

The theoretical results here facilitate analysis of the effect of the frequency-weighting filter H on the performance of the adaptive controller. Comparison of the experimental results in Section VII to the theoretically predicted performance provides an important evaluation of the adaptive controller.

The theoretical analysis here assumes that

$$G(z) = \hat{G}(z) \quad (10)$$

and that the adaptive control command u is not saturated. When (10) holds, the input to the adaptive filter F is

$$\hat{w} = w + \eta_0 = (I + PC)^{-1}(w_0 + \eta_0) = w_1 + \eta_1. \quad (11)$$

According to (4), w_1 depends only on the jitter produced by the disturbance mirror and the shaker, and η_1 depends only on the sensor noise.

The tuning signal is

$$e = H\tilde{y} + H[F\hat{G} - \hat{G}F]\hat{w}. \quad (12)$$

When the gains in F have converged to steady-state values, the transfer functions F and \hat{G} commute, so that the tuning signal is $e = H\tilde{y}$. During adaptation, the time-varying filter F does not commute with \hat{G} exactly, but according to swapping lemmas used in analysis of adaptive control [23]–[25], the term $[F\hat{G} - \hat{G}F]\hat{w}$ in (12) diminishes as the gains in F converge.

For theoretical steady-state analysis, the disturbance models discussed in Section III are used, and the sequences ε_w and ε_η are mutually independent with unit variance in each channel. It follows from Figs. 5 and 6 that the steady-state tuning signal e

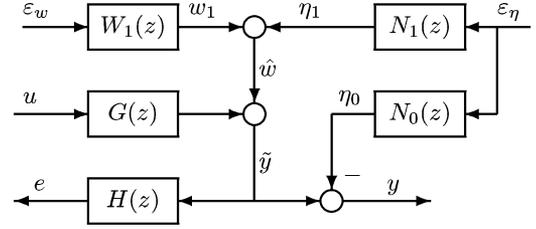


Fig. 6. Theoretical closed-loop system. $G(z)$ = closed-loop plant in Fig. 3, $W_1(z)$ = jitter model, $N_0(z)$ and $N_1(z)$ = sensor noise models, $H(z)$ = weighting filter.

and beam position y satisfy

$$e = (H + F\hat{G}H) \begin{bmatrix} W_1 & N_1 \end{bmatrix} \begin{bmatrix} \varepsilon_w \\ \varepsilon_\eta \end{bmatrix} \quad (13)$$

$$y = \begin{bmatrix} (I + \hat{G}F)W_1 & [(I + \hat{G}F)N_1 - N_0] \end{bmatrix} \begin{bmatrix} \varepsilon_w \\ \varepsilon_\eta \end{bmatrix} \quad (14)$$

and the theoretical steady-state rms values of these signals can be calculated as

$$J_e(F, H) = \left\| (I + F\hat{G}H) \begin{bmatrix} W_1 & N_1 \end{bmatrix} \right\|_2 \quad (15)$$

$$J_y(F, H) = \left\| (I + \hat{G}F)W_1 \quad [(I + \hat{G}F)N_1 - N_0] \right\|_2. \quad (16)$$

In steady state, the adaptive control loop minimizes $J_e(F, H)$ over the set of FIR filters of order L .

To compare results from experiments with different noise statistics, the following normalized $J_y(F, H)$ will be used:

$$\bar{J}_y(F, H) = \frac{J_y(F, H)}{\min_F J_y(F, H=1)(N_0 = N_1 = 0)} \quad (17)$$

where the denominator is the minimum, over FIR filters F of order L , of $J_y(F, H)$ in (16) with $H = 1$ and $N_0 = N_1 = 0$.

For theoretical analysis, with the assumed disturbance models and the condition (10), the minimization of $J(F, H)$ over FIR filters of order L can be formulated as a Wiener–Hopf problem [26], although the formulation is rather complicated for the problem here. The current problem falls into the class of problems discussed in [27], which provides an efficient numerical algorithm for computing the optimal filter F and the minimum output-error variance from plant and disturbance models. This algorithm was used to compute the theoretical steady-state filters F in Section VII of this paper. The theoretical performance indexes $J_y(F, H)$ and $\bar{J}_y(F, H)$ were given by (16) and (17). Application of the algorithm in [27] to the current problem is based on the block diagram in Fig. 6.

VI. MOTIVATION FOR FREQUENCY WEIGHTING

A. High-Frequency Gain of the Adaptive Filter

The adaptive filter F that generates the adaptive control command typically has large high-frequency gains, which amplify high-frequency sensor noise and increase the sensitivity of the closed-loop system to high-frequency modeling error in \hat{G} . The frequency response of F depends on the frequency content of the jitter and noise sequences, as well as on the plant transfer function. Fundamentally, F predicts the jitter, and prediction filters

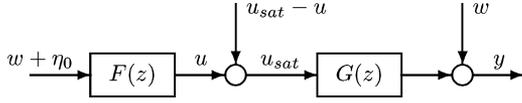


Fig. 7. Block diagram extracted from Fig. 5. The input sequence $u_{sat} - u$ produces the spikes in the output sequence y .

typically have large high-frequency gain. Also, the minimum-variance F is, in a certain sense, an approximate inverse of the plant transfer function G , which rolls off at high frequencies, as shown by the Bode plot of \hat{G} in Fig. 2. The frequency-weighting filter H in Fig. 5 is used to reduce the high-frequency gain of F .

The question then is, what type of weighting filter H will lead to an FIR filter F with reduced high-frequency gain? The somewhat counterintuitive answer is a high-pass filter H . The sense in which F inverts G is that F minimizes the 2-norm of $I + F\hat{G}$ weighted as in (15), which shows that a high-pass H penalizes the high-frequency gain of F . The Bode plots in Section VII for both the theoretical and experimental filters F show that high-frequency weighting filters H indeed reduce the high-frequency gains of F . Theory and experiment show that H also reduces $\|F\|_1$ and $\|F\|_\infty$.

Low-pass filters have been used in adaptive control of computer hard drives [18]. However, low-pass filters, whether used as in [18] or as H here or in other ways tried by the authors of this paper, have not reduced the high-frequency gain of F or otherwise improved the performance in the class of applications considered in this paper.

B. Output Spikes Produced by Control Signal Saturation

Experimental results in Section VII show that the combination of control saturation and excessively large high-frequency gain in the filter F produce a spiking phenomenon in the output error when the jitter or the sensor noise has high-frequency components of sufficient amplitude. One might suspect, as the authors first did, that the phenomenon reflects some instability of the adaptive control loop, possibly associated with parameter drift. However, careful analysis of several sets of experimental results indicated that this is not the explanation. The SISO simulation in this section, which uses an LTI filter F in the control system instead of the adaptive filter, produces the same spiking phenomenon and reveals the true nature of the phenomenon.

The block diagram in Fig. 7, which is extracted from Fig. 5 for the case where $\hat{G} = G$, represents the saturated control signal u_{sat} as the sum of the minimum-variance control command u generated by the filter F and the input sequence $u_{sat} - u$. When high-frequency jitter or sensor noise causes the output of the filter F to saturate, the input sequence $u_{sat} - u$ consists of randomly occurring impulses that produce spikes in the output error y .

The simulation represented by Fig. 8 provides the simplest demonstration of the nature of the spiking phenomenon observed in the first set of experimental results in Section VII. In this simulation, the sensor noise is zero (i.e., ε_η , η_0 , and η_1 are zero), so that the simulation results show that sufficient high-frequency power in the jitter alone can produce control saturation and spiking in the output error.

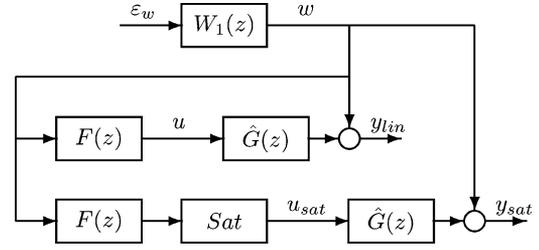


Fig. 8. Block diagram of simulations with LTI filter $F(z)$. Block labeled Sat represents control saturation.

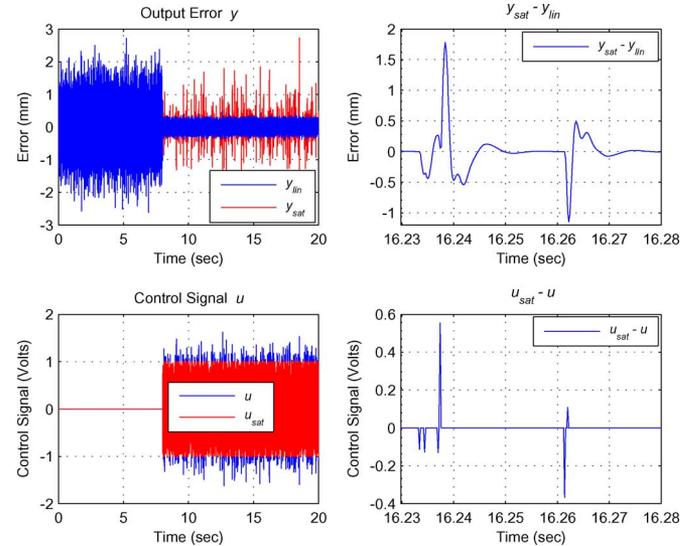


Fig. 9. Results of the simulation in Fig. 8. The control signals u and u_{sat} are zero until $t = 8$ s. Saturation limit for $|u_{sat}| = 1$. Zoomed views on the right show the nature of the spikes in the output error.

The transfer function \hat{G} used in the simulation is the first channel of the \hat{G} discussed in Section III-B, and (10) is assumed. The jitter model W_1 used in the simulation was identified as discussed in Section IV from input–output data from an experiment similar to those discussed in Section VII, except with no sensor noise. This jitter model was scaled by a gain so that, even in the absence of sensor noise, the control signal u saturates. The LTI filter F in the simulation minimizes $J_e(F, H = 1)$, which is equal to $J_y(F, H = 1)$ under the conditions of the simulation. Thus, F minimizes $J_e(F, H)$ as given in (15) with $H = 1$ and $N_1 = N_0 = 0$.

The simulation generates the output error y_{lin} with no control saturation and the output error y_{sat} with control saturation. The plots on the left in Fig. 9 show y_{lin} and y_{sat} and control signals u and u_{sat} . The unsaturated control signal u is optimal for the disturbance in the simulation. However, the large high-frequency gain in F causes u to exceed the saturation limit at random sample times. The plots on the right in Fig. 9 show zoomed views of $y_{sat} - y_{lin}$ and $u_{sat} - u$ for typical spikes in the output error. The input sequence $u_{sat} - u$ is a sequence of impulses at random times, which generates the corresponding sequence $y - y_{sat}$ of impulse responses of the transfer function \hat{G} . These impulse responses are the spikes in y_{sat} .

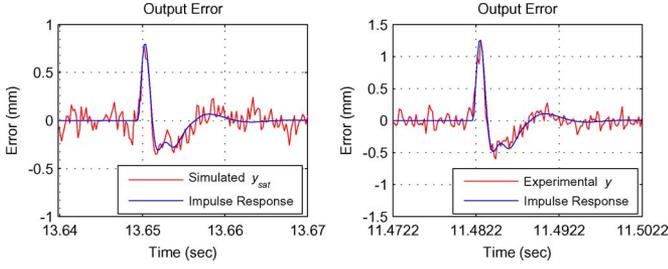


Fig. 10. Left: Single spike in output error y_{sat} from the simulation in Fig. 8 and impulse response of \hat{G} . Right: Single spike in experimental output error y in Fig. 13 and impulse response of \hat{G} .

Occasionally, there is a single isolated impulse producing an isolated spike, as in Fig. 10, where the plot on the left is from the simulation and the plot on the right is from the experimental output error in Fig. 13. Fig. 10 shows that the spikes in the simulation and the experiment are indeed very similar and that they are impulse responses of the plant. It should be emphasized that the filter F in the simulation is LTI whereas the filter F in the experiments is adaptive, although most of the spikes in the experiment are observed in steady state, where the adaptive filter is very similar to the LTI filter in the simulation.

In the experiments and simulations described in this paper, the control signal remains at one saturation limit for very short durations, almost always over one sample period only. However, it seems possible that spikes similar to those here could occur if the control signal remained at one saturation limit for a few consecutive sample times, at least for fast sample-and-hold rates.

While there is extensive literature on control under saturation constraints (e.g., [28]–[35] and many other references), it appears that the spiking phenomenon here has not been seen in previously published examples, where the control bounds are typically binding over longer time intervals. However, certain output transients in [36] and [37], which treat nonadaptive control in a different class of applications, may be related to the impulses observed here.

C. Plant Model Uncertainty

Another important reason for frequency weighting is stability robustness to high-frequency modeling error. The results in Sections V-B and VI-B require (10), but in applications, the modeling error $\Delta_G = G - \hat{G}$ is not zero. This error can affect the stability of the adaptive control system.

Fig. 5, without the bottom part showing how F is identified and without the saturation block, is equivalent to Fig. 11, from which it follows that a sufficient condition for closed-loop bounded-input–bounded-output (BIBO) stability is $\|\Delta_G\|_1 \|F\|_1 < 1$. In the case of LTI F and G , a sufficient condition for asymptotic stability is $\|\Delta_G\|_\infty \|F\|_\infty < 1$. Thus, smaller values of $\|F\|_1$ and $\|F\|_\infty$ increase the stability robustness of the closed-loop system.

VII. EXPERIMENTAL RESULTS

This section presents detailed results for two experiments designed for studying the effect of sensor noise on the performance

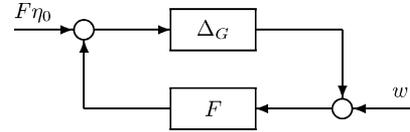


Fig. 11. Closed-loop system with modeling error.

of the adaptive controller. Because the signals from the optical position sensor used in these experiments have negligible sensor noise, the two-channel sensor noise η_0 was added in software so that the statistics of η_0 are known precisely. This is ideal for studying the effects of the noise.

Results for two types of sensor noise are presented: high-frequency noise between 1.5 and 2.5 KHz (the Nyquist frequency) and white noise. The jitter and sensor noise were generated as described in Section IV. In all cases with frequency weighting, the filter H was a high-pass linear-phase FIR filter of order 16. The filter H used for the experimental results reported here had the cutoff frequency 700 Hz.

A. Effects of the Bandwidth of the Weighting Filter

The Bode plots on the left in Fig. 12 show how the high-pass weighting filter H with different cutoff frequencies theoretically affects the frequency response of the steady-state filter F . The plots on the right in Fig. 12 show how the theoretical values of $\|F\|_1$ and $\|F\|_\infty$ and the normalized rms steady-state position error $\bar{J}_y(F, H)$ vary with the lower cutoff frequency of H . The experimental values of $\|F\|_1$, $\|F\|_\infty$, and $\bar{J}_y(F, H)$ are shown in Fig. 12 for the case of no frequency weighting (cutoff frequency = 0) and for the cutoff frequency 700 Hz.

Table II lists experimental steady-state data for the two types of sensor noise and the case of no sensor noise. The signal-to-noise ratio (SNR) in the table is based on (11) and given by

$$\text{SNR} = \frac{\text{rms}(w_1)}{\text{rms}(\eta_1)}. \quad (18)$$

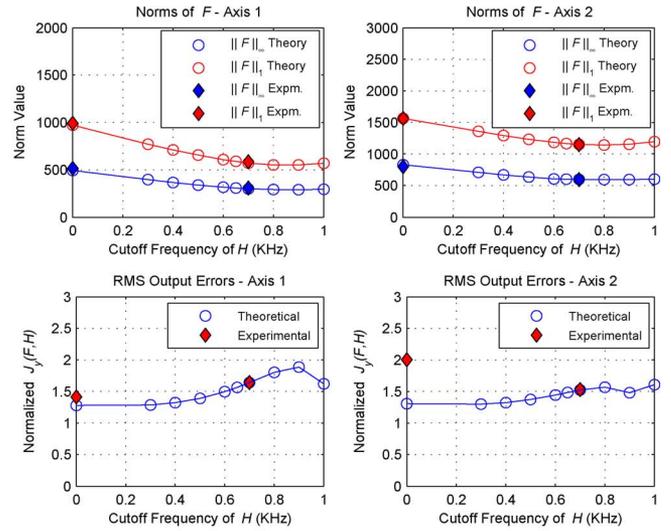
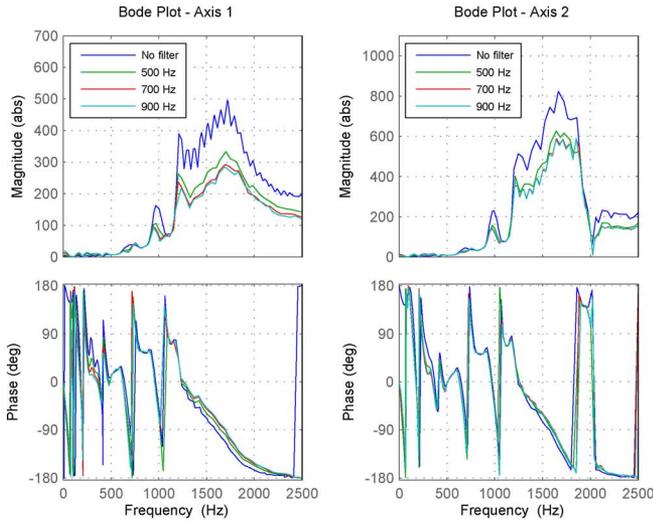
The cutoff frequency 700 Hz was chosen for the weighting filter H because the plots of the theoretical values of $\|F\|_1$, $\|F\|_\infty$, and $\bar{J}_y(F, H)$ indicate that higher cutoff frequencies produce at most negligible further reductions in $\|F\|_1$ and $\|F\|_\infty$ while increasing $\bar{J}_y(F, H)$. The plots in Fig. 12 depend on both the jitter statistics, represented by $W_1(z)$, and the sensor-noise statistics, represented by $N_0(z)$ and $N_1(z)$. In most applications, these statistics are not known, so it is noteworthy that the plots in Fig. 12 indicate that a cutoff frequency between 600 and 700 Hz is a good choice in terms of the tradeoff between reducing the norms of F and increasing $\bar{J}_y(F, H)$, for either type of sensor noise. The adaptive controller requires no information about disturbance and sensor-noise statistics, but ballpark *a priori* estimates can be useful in choosing the cutoff frequency of H .

For the theoretical results in Fig. 12, the theoretically optimal steady-state filter for each noise spectrum and bandwidth of H was computed with the algorithm in [27] for the optimal control problem defined in Section V-B. The corresponding theoretical steady-state performance was given by (16) and (17). In the

High-frequency Sensor Noise

Bode Plots of Theoretical Optimal F

Norms of F and RMS Output Errors



White Sensor Noise

Bode Plots of Theoretical Optimal F

Norms of F and RMS Output Errors

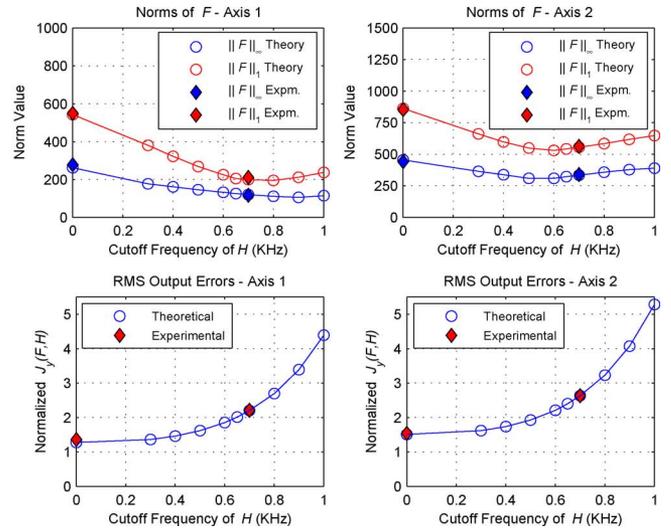
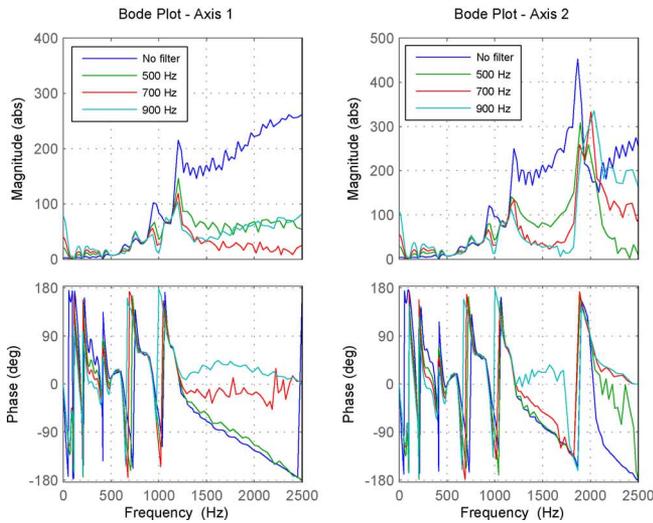
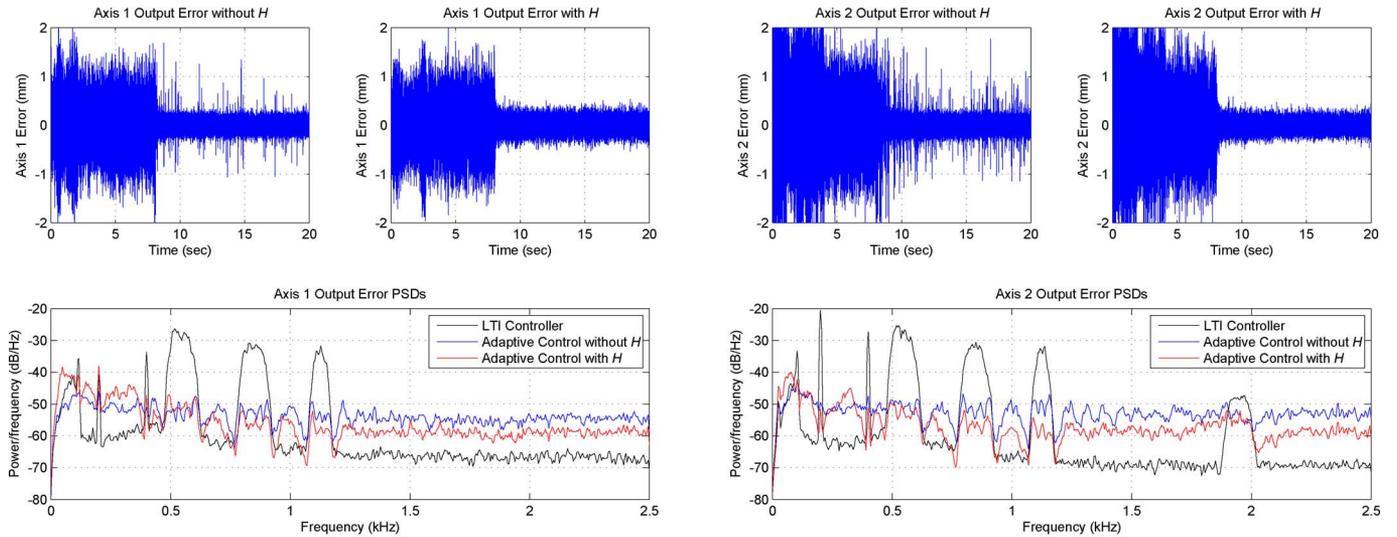


Fig. 12. Plots illustrate effects of H with different cutoff frequencies on the steady-state filter F and steady-state performance of the adaptive controller. Left: Bode plots of theoretical optimal FIR filter F for cutoff frequencies of $H = 500, 700, 900$ Hz. No filter means $H = 1$. Right: Theoretical and experimental norms of F and normalized rms steady-state output error values $\bar{J}_y(F, H)$ [in (17)] versus cutoff frequency of H . Cutoff frequency = 0 means $H = 1$. Experimental results are represented by the symbol \blacklozenge for H cutoff frequencies 0 and 700 Hz.

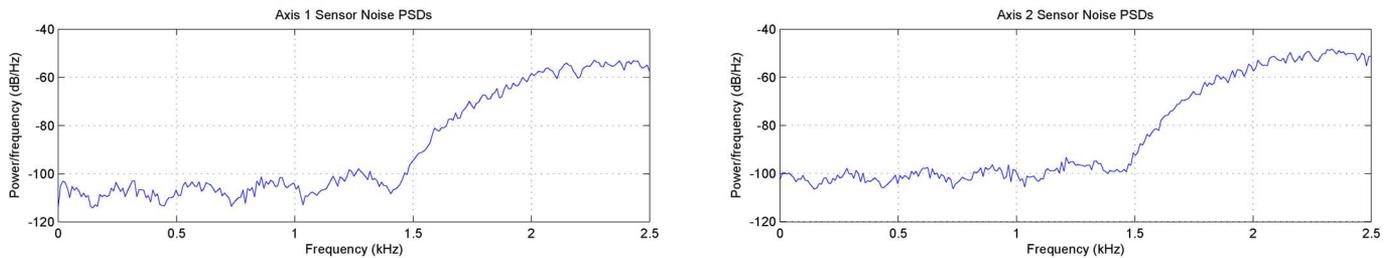
TABLE II
EXPERIMENTAL RESULTS WITH DIFFERENT TYPES OF SENSOR NOISE

	No Sensor Noise		High-frequency Sensor Noise		White Sensor Noise	
	Axis 1	Axis 2	Axis 1	Axis 2	Axis 1	Axis 2
SNR	∞	∞	9.862	9.690	9.843	9.597
$J_y(F, H) = \text{RMS}(y)$, LTI Feedback Control Only	0.379	0.443	0.470	0.561	0.385	0.442
$J_y(F, H) = \text{RMS}(y)$, Adaptive Control without H	0.086	0.078	0.111	0.146	0.099	0.101
$J_y(F, H) = \text{RMS}(y)$, Adaptive Control with H	0.136	0.118	0.129	0.111	0.158	0.172
$\ F\ _\infty$ without H	442.9	839.3	515.5	799.3	277.5	441.8
$\ F\ _\infty$ with H	160.2	457.3	307.3	595.2	116.8	339.4
$\ F\ _1$ without H	866.6	1685.	991.0	1562.	549.7	854.4
$\ F\ _1$ with H	318.7	774.2	583.3	1148.	210.5	562.6

Output-error Time Series and Power Spectral Densities



PSDs for Sensor Noise η_0



Bode Plots of Theoretical and Experimental Steady-state Filter F

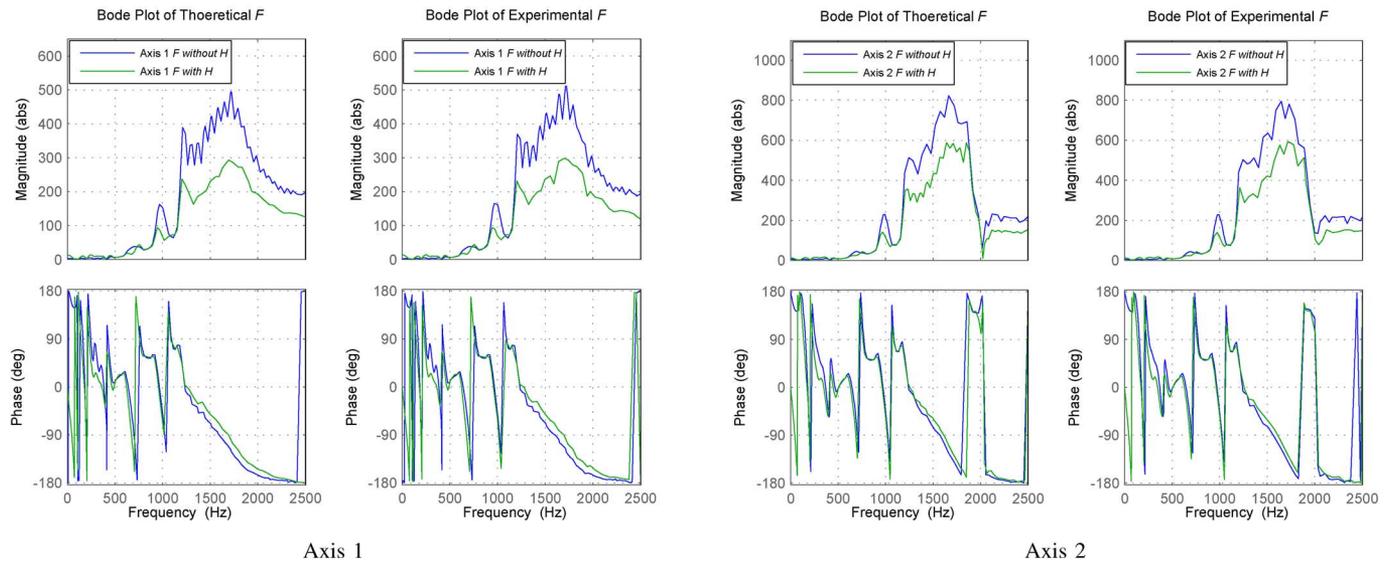


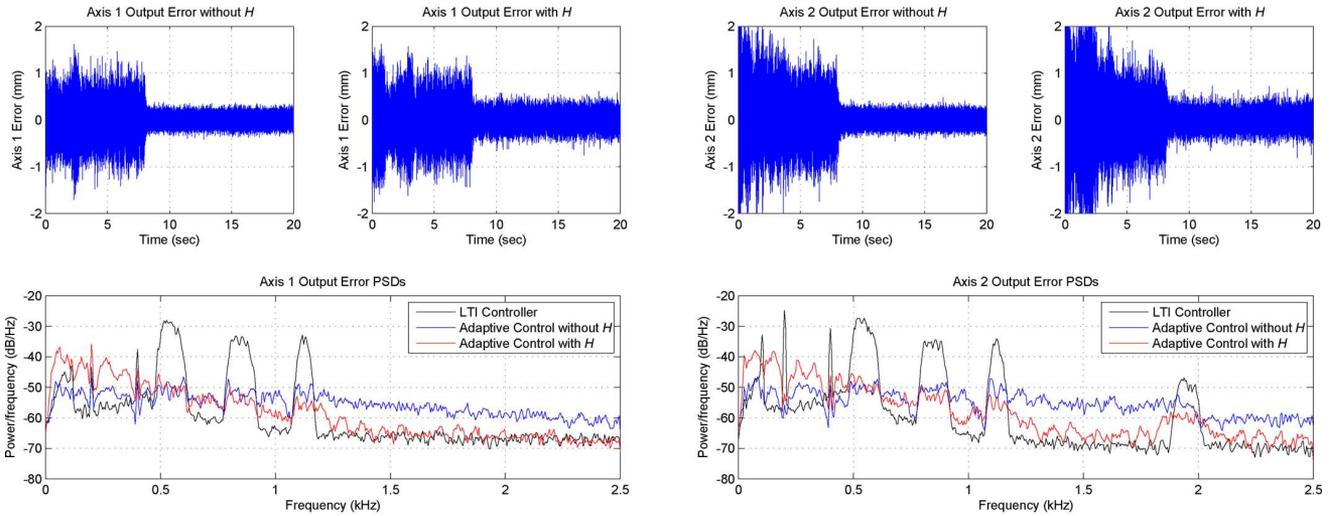
Fig. 13. Experimental results for the case of high-frequency sensor noise. Cutoff frequency of weighting filter $H = 700$ Hz. The LTI feedback control loop described in Section III-B was closed from the beginning of the each experiment, and the adaptive loop was closed at $t = 8$ s.

plots on the right in Fig. 12, the experimental results agree closely with the theoretical results, except for the normalized rms output error in the case with high-frequency sensor noise and no frequency weighting.

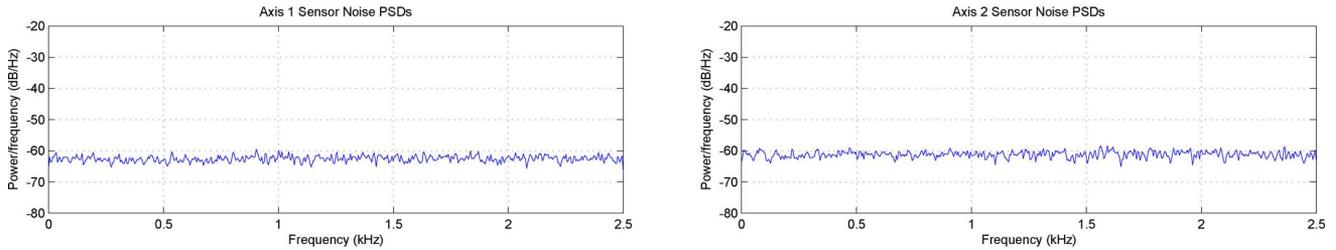
B. High-Frequency Sensor Noise and Control Saturation

Fig. 13 shows experimental results for the case of high-frequency sensor noise. Without frequency weighting, the experimental steady-state output error is larger than predicted

Output-error Time Series and Power Spectral Densities



PSDs for Sensor Noise η_0



Bode Plots of Theoretical and Experimental Steady-state Filter F

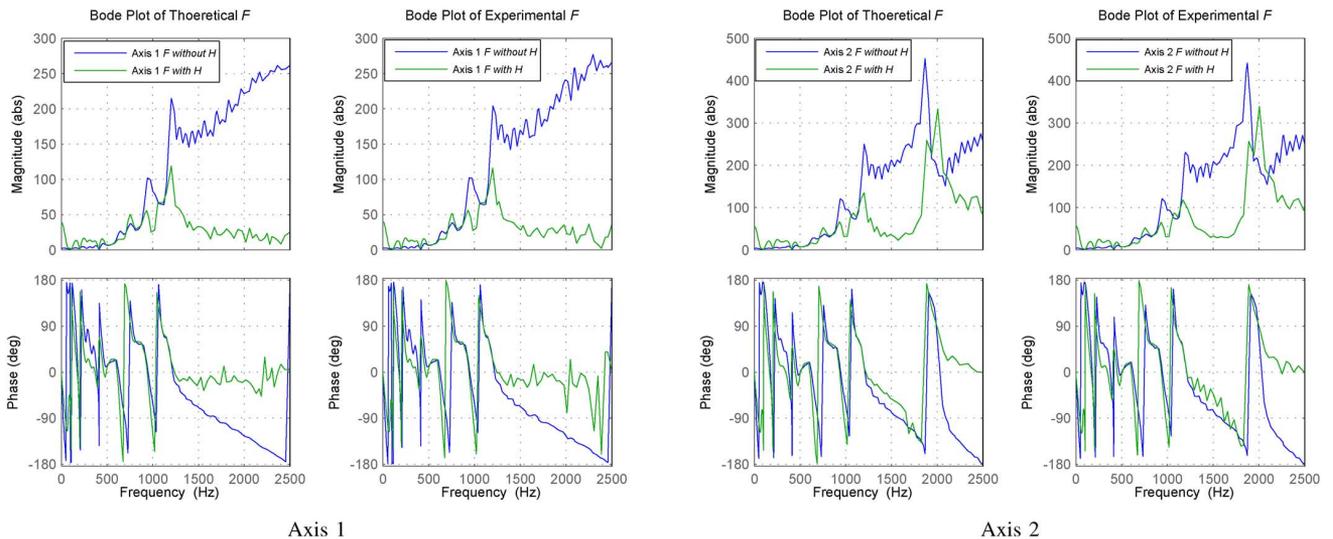


Fig. 14. Experimental results for the case of white sensor noise. Cutoff frequency of weighting filter $H = 700$ Hz. The LTI feedback control loop described in Section III-B was closed from the beginning of the each experiment, and the adaptive loop was closed at $t = 8$ s.

theoretically because control saturation (which is not modeled in the theoretical calculations) produces the spikes shown in the time-series plots of the output errors in Fig. 13. The nature of these spikes is described in Section VI-B. The results in Fig. 13 show that the frequency weighting in the adaptive control loop eliminates the spikes by reducing the high-frequency gains of

the adaptive filter F . The Bode plots in Fig. 13 show close agreement between the theoretical steady-state filters F and those to which the adaptive lattice filter converged in the experiment.

The results in Table II for high-frequency sensor noise show that the adaptive controller without frequency weighting reduces the rms output error by approximately a factor of 4 even while

generating the spikes. However, such spikes in output error are undesirable in any practical control system.

In the absence of control saturation, adaptive control with frequency weighting produces larger values of $\text{rms}(y)$ than adaptive control without frequency weighting. However, when the spikes produced by control saturation are sufficiently severe, the value of $\text{rms}(y)$ can be substantially larger than the linear theory predicts and larger than $\text{rms}(y)$ with frequency weighting. For example, for the high-frequency noise, adaptive control without frequency weighting produces $\text{rms}(y) = 0.146$ for channel 2 whereas adaptive control with frequency weighting produces $\text{rms}(y) = 0.111$. Examination of the time series for channel 2 in Fig. 13 suggests that the rms value of the output error y should be smaller with H .

The power spectral densities (PSDs) of the output errors in Fig. 13 show that, in steady state, the adaptive controller exhibits a common characteristic of minimum-variance controllers: the tendency to whiten residual errors. This means that the adaptive controller reduces the output-error PSDs in bandwidths of large jitter at the minor expense of amplifying low-level jitter and sensor noise. The PSDs show the output errors without frequency weighting to be nearly white, even with the spikes due to control saturation. With frequency weighting, which penalizes high-frequency error more, the adaptive controller amplifies low-level jitter and sensor noise less at high frequencies and more at low frequencies.

The Bode plots in Fig. 13 show that, as predicted theoretically, the weighting filter H reduces the high-frequency gain of the steady-state adaptive filter F . The close agreement between the theoretical F and experimental F is noteworthy.

Another point to note from the PSD plots in Fig. 13 is that the adaptive controller with H reduces the jitter by more than 20 dB in the 1100–1150 Hz band. Fig. 2 shows that the open-loop and LTI closed-loop plants have low gain in this band so that large control power in the 1100–1150 Hz band is required to achieve the jitter rejection shown in Fig. 13.

C. White Sensor Noise

Fig. 14 shows experimental results for the case of white sensor noise. As indicated in Table II, the rms SNRs are almost identical for white and high-pass sensor noise. The jitter statistics are the same for the two cases also. Yet, the spikes do not appear in the output error for the white-noise case because the steady-state filter F has significantly lower high-frequency gain in this case so that the control signal does not saturate.

The high-frequency gains in F are lower with white sensor noise than with no noise, for the same jitter statistics [38]. However, this does not mean that the saturation-induced spikes in the output error never occur with white noise. As the simulation discussed in Section VI-B shows, the spikes can occur even with no sensor noise. The phenomenon occurs when there is sufficient high-frequency power in the control signal to cause control saturation like that in the example in Section VI-B. The high-frequency power in the control signal results from a combination of jitter and sensor-noise statistics and the high-frequency gain of F .

As shown in Table II and the time-series plots in Fig. 14, the frequency-weighted adaptive controller yields larger rms output errors than the adaptive controller without frequency weighting. This is expected because the frequency-weighted adaptive filter minimizes the frequency-weighted rms output error instead of the ordinary rms error.

The Bode plots in Fig. 14 show that, again, the weighting filter H reduces the high-frequency gain of the steady-state adaptive filter F . Again, the theoretical F and experimental F agree closely.

VIII. CONCLUSION

This paper has introduced a frequency-weighting method for minimum-variance adaptive control, and presented experimental results that demonstrate the effectiveness of the method in control of an MEMS mirror for suppression of laser beam jitter. Output errors in both experimental results and a simulation illustrate a spiking phenomenon produced by the minimum-variance controller when high-frequency noise or disturbance causes randomly occurring control signal saturation of short duration. The experimental results show that the frequency weighting presented in the paper eliminates the spiking.

Aside from the frequency weighting method presented here and the spiking phenomenon discussed, the results in this paper represent a significant advance beyond those in [13]. Here, as in [13], lattice-filter-based adaptive control suppressed laser beam jitter in simultaneous multiple bandwidths, but the improved adaptive control loop here rejected jitter at frequencies greater than 1100 Hz, more than three times the highest frequencies at which jitter was controlled in [13] with the same MEMS mirror—which has a natural frequency of vibration at approximately 120 Hz on each axis. This increased frequency range has resulted from the higher sampling rate of 5 KHz and higher adaptive filter order of 80 here, as opposed to the 2 KHz sampling rate and filter order 16 in [13], as well as the frequency weighting in the case of high-frequency sensor noise.

ACKNOWLEDGMENT

The authors are indebted to TI Incorporated for providing the TALP1000A Analog MicroMirrors used in this research.

REFERENCES

- [1] M. C. Roggemann and B. Welsh, *Imaging Through Turbulence*. Boca Raton, FL: CRC Press, 1996.
- [2] R. K. Tyson, *Principles of Adaptive Optics*. San Diego, CA: Academic, 1997.
- [3] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control*. Chichester, U.K.: Wiley, 1996.
- [4] R. M. Glaese, E. H. Anderson, and P. C. Janzen, "Active suppression of acoustically induced jitter for the airborne laser," in *Proc. SPIE, Laser Weapons Technol.*, vol. 4034, Orlando, FL, Apr. 2000, pp. 151–3421.
- [5] M. A. McEver, D. G. Cole, and R. L. Clark, "Experiments in adaptive optical jitter control," in *Proc. SPIE, Smart Struct. Mater. 2003: Model., Signal Process., Control*, vol. 5049, San Diego, CA, Mar., pp. 275–282.
- [6] M. A. McEver, D. G. Cole, and R. L. Clark, "Adaptive feedback control of optical jitter using Q -parameterization," *Opt. Eng.*, vol. 43, no. 4, pp. 904–910, Apr. 2004.
- [7] L. P. Fowler and R. Blankinship, "Experimental adaptive filtering and disturbance feedforward approach for flexible beam train control with single disturbance path," presented at the 2006 DEPS Directed Energy Syst. Symp.: Beam Control Conf., Monterey, CA, Mar. 2006.

- [8] V. A. Skormin, M. A. Tascillo, and T. E. Busch, "Adaptive jitter rejection technique applicable to airborne laser communication systems," *Opt. Eng.*, vol. 34, no. 15, pp. 1263–1268, May 1995.
- [9] N. O. Pérez Arancibia, S. Gibson, and T.-C. Tsao, "Adaptive control of MEMS mirrors for beam steering," presented at the ASME Int. Mech. Eng. Congr. Expo., Anaheim, CA, Nov. 2004, Paper IMECE2004–60256.
- [10] N. O. Pérez Arancibia, N. Chen, S. Gibson, and T.-C. Tsao, "Adaptive control of a MEMS steering mirror for suppression of laser beam jitter," in *Proc. Amer. Control Conf.*, Portland, OR, Jun. 2005, pp. 3586–3591.
- [11] N. O. Pérez Arancibia, N. Chen, S. Gibson, and T.-C. Tsao, "Adaptive control of a MEMS steering mirror for free-space laser communications," in *Proc. SPIE, Free-Space Laser Commun. V*, San Diego, CA, Sep. 2005, vol. 5892, pp. 334–345.
- [12] N. O. Pérez Arancibia, S. Gibson, and T.-C. Tsao, "Saturation and frequency weighting in adaptive control of laser beam jitter," in *Proc. SPIE, Free-Space Laser Commun. VII*, vol. 6709, San Diego, CA, Aug. 2007, pp. 6709–6727.
- [13] N. O. Pérez Arancibia, N. Chen, S. Gibson, and T.-C. Tsao, "Variable-order adaptive control of a microelectromechanical steering mirror for suppression of laser beam jitter," *Opt. Eng.*, vol. 45, no. 10, pp. 104 206-1–104 206-12, Oct. 2006.
- [14] P. K. Orzechowski, N. Y. Chen, J. S. Gibson, and T.-C. Tsao, "Optimal suppression of laser beam jitter by high-order RLS adaptive control," *IEEE Trans. Control Syst. Technol.*, vol. 16, no. 2, pp. 255–267, Mar. 2008.
- [15] A. Sacks, M. Bodson, and P. Khosla, "Experimental results of adaptive periodic disturbance cancellation in a high performance magnetic disk drive," *Trans. ASME, J. Dyn. Syst., Meas., Control*, vol. 118, no. 3, pp. 416–424, Sep. 1996.
- [16] M. Bodson and S. C. Douglas, "Adaptive algorithms for the rejection of sinusoidal disturbances with unknown frequency," *Automatica*, vol. 33, no. 12, pp. 2213–2221, Dec. 1997.
- [17] M. Bodson, A. Sacks, and P. Khosla, "Harmonic generation in adaptive feedforward cancellation schemes," *IEEE Trans. Autom. Control*, vol. 39, no. 9, pp. 1939–1944, Sep. 1994.
- [18] R. Horowitz and B. Li, "Adaptive track-following servos for disk file actuators," *IEEE Trans. Magn.*, vol. 32, no. 3, pp. 1779–1786, May 1996.
- [19] R. Horowitz, B. Li, and J. W. McCormick, "Wiener-filter-based minimum variance self-tuning regulation," *Automatica*, vol. 34, no. 5, pp. 531–544, May 1998.
- [20] K. Krishnamoorthy and T.-C. Tsao, "Adaptive-Q with LQG stabilizing feedback and real time computation for disk drive servo control," in *Proc. Amer. Control Conf.*, Boston, MA, Jun. 2004, pp. 1171–1175.
- [21] P. Van Overschee and B. De Moor, *Subspace Identification for Linear Systems*. Norwell, MA: Kluwer, 1996.
- [22] S.-B. Jiang and J. S. Gibson, "An unwindowed multichannel lattice filter with orthogonal channels," *IEEE Trans. Signal Process.*, vol. 43, no. 12, pp. 2831–2842, Dec. 1995.
- [23] A. Datta, "Robustness of discrete-time adaptive controllers: An input-output approach," *IEEE Trans. Autom. Control*, vol. 38, no. 12, pp. 1852–1857, Dec. 1993.
- [24] P. Ioannou and B. Fidan, *Adaptive Control Tutorial*. Philadelphia, PA: SIAM, 2006.
- [25] M. Krstić, I. Kanellakopoulos, and P. Kokotović, *Nonlinear and Adaptive Control Design*. New York: Wiley, 1995.
- [26] S. Haykin, *Modern Filters*. New York: Macmillan, 1998.
- [27] P. K. Orzechowski, J. S. Gibson, and T.-C. Tsao, "Characterization of optimal FIR gains and minimum-variance performance for adaptive disturbance rejection," in *Proc. Amer. Control Conf.*, New York, Jun. 2007, pp. 1908–1913.
- [28] J. C. Doyle, R. S. Smith, and D. F. Enns, "Control of plants with input saturation nonlinearities," in *Proc. Amer. Control Conf.*, Minneapolis, MN, Jun. 1987, pp. 1034–1039.
- [29] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *Proc. Amer. Control Conf.*, Pittsburgh, PA, Jun. 1989, pp. 1693–1698.
- [30] M. Kanamori and M. Tomizuka, "Model reference adaptive control of linear systems with input saturation," in *Proc. 2004 IEEE Int. Conf. Control Appl.*, Taipei, Taiwan, Sep. 2004, pp. 1318–1323.
- [31] A. Zheng, M. V. Kothare, and M. Morari, "Anti-windup design for internal model control," *Int. J. Control*, vol. 60, no. 5, pp. 1015–1024, Nov. 1994.
- [32] A. Saberi, Z. Lin, A. R. Teel, and P. Daoutatis, "Control of linear systems with saturating actuators," *IEEE Trans. Autom. Control*, vol. 41, no. 3, pp. 368–378, Mar. 1996.
- [33] F. Tyan and D. S. Bernstein, "Dynamic output feedback compensation for linear systems with independent amplitude and rate saturations," *Int. J. Control*, vol. 67, no. 1, pp. 89–116, May 1997.
- [34] N. Kapoor, A. R. Teel, and P. Daoutatis, "An anti-windup design for linear systems with input saturation," *Automatica*, vol. 34, no. 5, pp. 559–574, May 1998.
- [35] Y. Hong and B. Yao, "A globally stable high-performance adaptive robust control algorithm with input saturation for precision motion control of linear motor drive systems," *IEEE/ASME Trans. Mechatronics*, vol. 12, no. 2, pp. 198–207, Apr. 2007.
- [36] T. Nguyen and F. Jabbari, "Disturbance attenuation for systems with input saturation: An LMI approach," *IEEE Trans. Autom. Control*, vol. 44, no. 4, pp. 852–857, Apr. 1999.
- [37] T. Nguyen and F. Jabbari, "Output feedback controllers for disturbance attenuation with actuator amplitude and rate saturation," *Automatica*, vol. 36, no. 9, pp. 1339–1346, Sep. 2000.
- [38] N. O. Pérez Arancibia, "Adaptive control of opto-electro-mechanical systems for broadband disturbance rejection," Ph.D. dissertation, Univ. California, Los Angeles, 2007.



Néstor O. Pérez-Arancibia (S'05–M'08) received the Ingeniero and M.Eng. degrees from the Pontificia Universidad Católica de Chile, Santiago, Chile, in 2000, and the Ph.D. from the Mechanical and Aerospace Engineering Department, University of California, Los Angeles (UCLA), Los Angeles, in 2007.

Since October 2007, he has been a Postdoctoral Scholar with the Laser Beam Control Laboratory and the Mechatronics and Controls Laboratory, Mechanical and Aerospace Engineering Department, UCLA.

His current research interests include feedback control, adaptive filtering, adaptive optics, and mechatronics.



James S. Gibson received the B.S. degree in aerospace engineering, the M.S. degree in engineering mechanics, and the Ph.D. degree in engineering mechanics from the University of Texas, Austin, in 1970, 1972, and 1975, respectively.

He was with the faculties of the Aerospace Engineering and Engineering Mechanics Department, University of Texas, Blacksburg, Austin, and the Engineering Science and Mechanics Department, Virginia Polytechnic Institute and State University, Blacksburg. In 1977, he joined the faculty of the University of California, Los Angeles (UCLA), where he currently is a Professor of mechanical and aerospace engineering. His current research interests include control and identification of dynamical systems and adaptive filtering, with applications to identification and control of flexible structures, control of laser beams and adaptive optics, identification and control of microinertial sensors, and control of fluid flow and noise control. He was an Associate Editor of the *SIAM Journal on Control and Optimization*.

Prof. Gibson was an Associate Editor of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL.



Tsu-Chin Tsao (S'86–M'88–SM'09) received the B.S. degree in engineering from National Taiwan University, Taipei, Taiwan, in 1981, and the M.S. degree and the Ph.D. degree in mechanical engineering from the University of California, Berkeley, in 1984 and 1988, respectively.

He was a member of the faculty of the Mechanical and Industrial Engineering Department, University of Illinois at Urbana-Champaign, for 11 years. In 1999, he joined the faculty of the University of California, Los Angeles (UCLA), where he currently is a Professor in the Mechanical and Aerospace Engineering Department. His current research interests include control systems and mechatronics.

Prof. Tsao has received the *ASME Journal of Dynamic Systems, Measurement, and Control* Best Paper Award for the papers published in the journal in 1994, the Outstanding Young Investigator Award from the ASME Dynamic Systems and Control Division in 1997, and the Hugo S. Shuck Best Paper Award from the American Automatic Control Council in 2002.