

Variable-order adaptive control of a microelectromechanical steering mirror for suppression of laser beam jitter

Néstor O. Pérez Arancibia

Neil Y. Chen

James S. Gibson

Tsu-Chin Tsao

University of California

Mechanical and Aerospace Engineering

Los Angeles, California 90095-1597

E-mail: nestor@seas.ucla.edu

Abstract. We present an adaptive control scheme for laser-beam steering by a two-axis microelectromechanical systems (MEMS) fast steering mirror. Disturbances in the laser beam are rejected by a μ -synthesis feedback controller augmented by the adaptive control loop, which determines control gains that are optimal for the current disturbance acting on the laser beam. The variable-order adaptive controller is based on an adaptive lattice filter that implicitly identifies the disturbance statistics from real-time sensor data. Experimental results demonstrate that the adaptive controller significantly extends the disturbance-rejection bandwidth achieved by the feedback controller alone. The experimental results also illustrate the value of the variable-order capability of the adaptive controller. © 2006 Society of Photo-Optical Instrumentation Engineers. [DOI: 10.1117/1.2363189]

Subject terms: adaptive control; optical jitter; laser applications; engineering; feedback; adaptive filtering; wavefront compensation.

Paper 050425R received Jun. 13, 2005; revised manuscript received Mar. 31, 2006; accepted for publication Apr. 1, 2006; published online Oct. 26, 2006.

1 Introduction

Precise steering of laser beams has a wide range of applications in fields such as adaptive optics, wireless communications, and manufacturing process. The control problem is to position the centroid of a laser beam at a desired location on a target plane some distance from the laser source with minimal beam motion, or jitter, in the presence of disturbances. In applications, the most common jitter sources are turbulence in the atmosphere through which the beam travels and vibration of the optical bench. Turbulence-induced jitter may be rather broadband,¹⁻⁴ while vibration-induced jitter typically is composed of one or more narrow bandwidths produced by vibration modes of the structure supporting the optical system. Also, some beam-steering mirrors have lightly damped elastic modes that produce beam jitter. This is the case with the microelectromechanical systems (MEMS) mirrors used in the experiment presented here. These mirrors, which are used in free-space optical communications systems, have a torsional vibration mode about each steering axis.

Because the disturbance characteristics often change with time, optimal performance of a beam-steering system requires an adaptive control system. Recent research on jitter control has produced adaptive control methods that employ least-mean-square⁵ (LMS) adaptive filtering and recursive least-squares (RLS) adaptive filtering.^{6,7} The trade-off is between a simpler algorithm (hence computational economy) with LMS versus faster convergence and exact minimum-variance steady-state performance with RLS.

This paper employs an RLS lattice filter in the adaptive controller, and introduces a variable-order adaptive control

scheme that exploits the order-recursive structure of the lattice filter. The capability to vary the order of the filter in the adaptive controller is important because optimal gains can be identified faster for lower-order filters, while higher-order filters are required for optimal steady state rejection of broadband disturbance. Thus, low filter orders can be used initially for fast adaptation without undesirable transient responses, and the filter order can be increased incrementally to achieve optimal steady state jitter rejection.

Section 2 describes the experimental hardware and configuration. Section 3 describes the system identification of the mirror dynamics and transfer functions required for control system design. Section 4 describes the design of the control system, which consists of a linear time-invariant (LTI) feedback control loop augmented by the adaptive control loop. Experimental results for two sets of experiments, each with multiple jitter bandwidths, are presented in Sec. 5.

2 Description of the Experiment

The experimental system is shown in Figs. 1–3. The main optical components in the experiment are the laser source, two Texas Instruments MEMS fast steering mirrors (FSMs), and an On-Trak position sensing device. Figure 2 shows the path of the laser beam from the source to the position sensor. After leaving the laser source, the beam reflects off the mirror FSM 1, which serves as the control actuator, then reflects off the mirror FSM 2, which adds disturbance to the beam direction, and finally goes to the sensor. As shown in Fig. 1, a lens between FSM 1 and FSM 2 and another lens between FSM 2 and the sensor focus the beam to maintain small spots on FSM 2 and the sensor.

Each mirror rotates about vertical and horizontal axes, denoted by axis 1 and axis 2, respectively. The outputs of

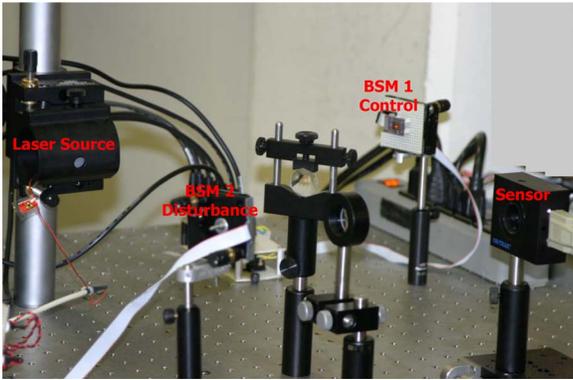


Fig. 1 Photograph of the laser-beam-steering experiment.

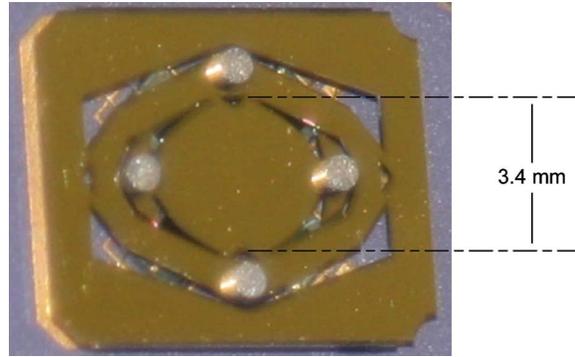


Fig. 3 Texas Instruments TALP1000A MEMS fast steering mirror (3.2- \times 3.6-mm elliptical mirror); axis 1, vertical; axis 2, horizontal.

the sensor are the horizontal and vertical displacements of the centroid of the laser spot on the plane of the On-Track optical sensor. The sensor axes are labeled axis 1 and axis 2, respectively, to correspond to beam deflections produced by the mirror rotations. Thus, in the sensor plane, axis 1 and axis 2 are horizontal and vertical, respectively.

Computer 1 has a Texas Instruments TMS320C6701 digital signal processor (DSP). This DSP runs both feedback and adaptive controllers and sends actuator commands to FSM 1. Computer 2, a PC running the xPC Target real-time operating system, sends disturbance commands to

FSM 2. For the experiments reported here, the sample-and-hold rate is 2000 Hz, which is approximately 15 times the larger of the two natural frequencies of the MEMS mirrors.

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, as indicated in Fig. 2. Note that the only measurements used by the adaptive and feedback controllers are the two signals from the On-Trak sensor. The MEMS steering mirrors used here have internal

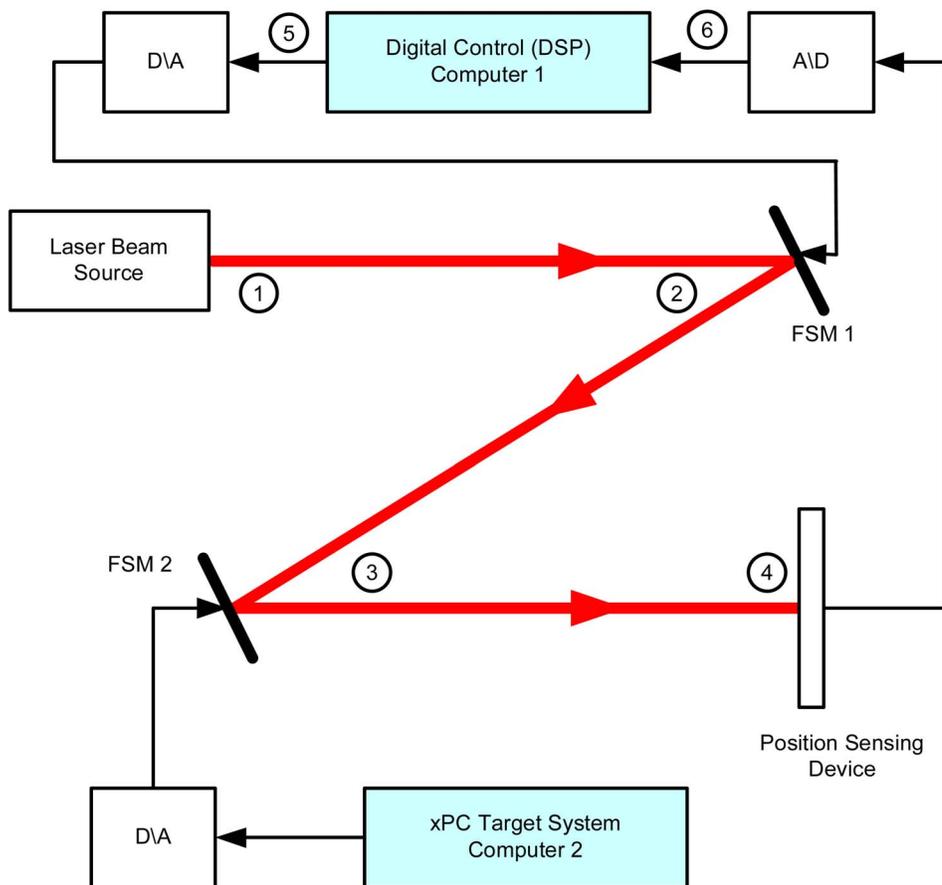


Fig. 2 Diagram of the experiment.

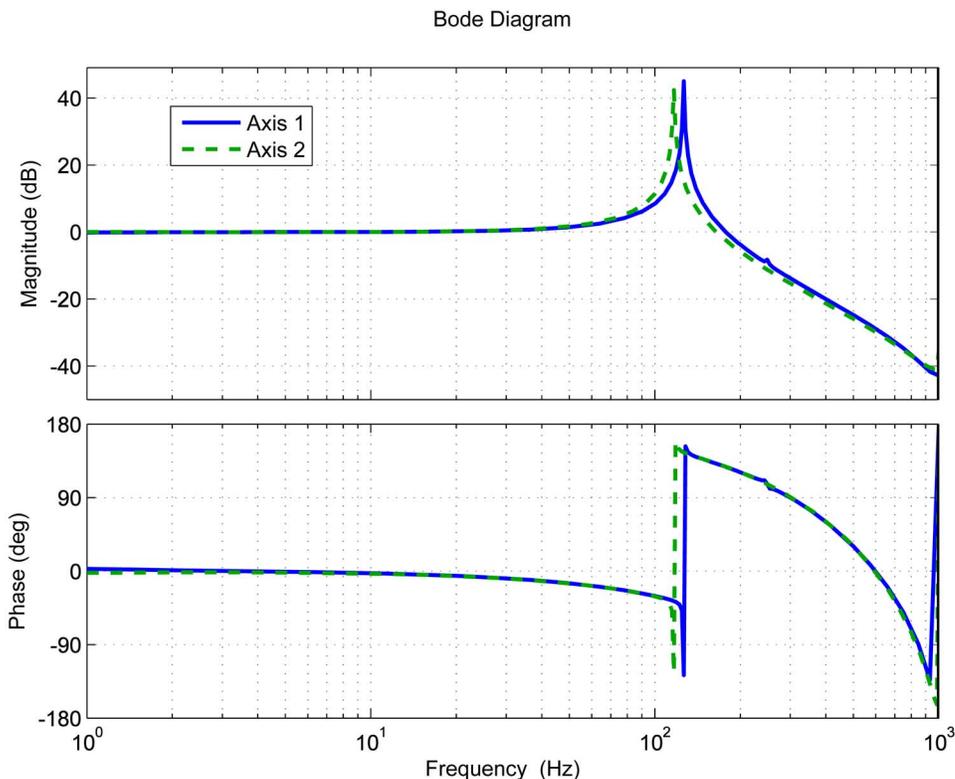


Fig. 4 Bode plots for identified model $\hat{P}(z)$ of open-loop beam-steering mirror FSM 1 (control actuator). Natural frequencies: 126.5 Hz (axis 1) and 119.4 Hz (axis 2).

optical sensors that supply local measurements of the mirror position, but these measurements were not used in the experiments discussed in this paper.

The commanded rotations of the beam-steering mirrors are produced by electromagnetic fields with opposing directions. These fields are created by coils with currents generated by the voltage commands from the control and disturbance computers. The mirrors have a rotation range of ± 5 deg. The reflecting area of the mirrors is 9 mm².

The optoelectronic position sensor at the end of the beam path generates two analog output voltages proportional to the 2-D position of the laser beam centroid. In the sensor, quad photodetectors capture the light intensity distribution and generate currents, which are converted to voltage and amplified by an operational amplifier. Further electronic processing of these voltage signals yields two final signals, which are the estimates of the centroid coordinates independent of light intensity.

3 System Identification

The design of the feedback control system requires an open-loop model of the dynamics of the steering mirror FSM 1, and the adaptive control loop requires an estimate of the transfer function from the adaptive-control commands to the sensor outputs with the feedback loop closed. The open-loop and closed-loop transfer functions are identified by a subspace method^{8,9} using input-output data from two brief experiments in which FSM 1 was driven by white noise. After the first of these experiments, which was open-loop, the feedback controller was designed, and then the

feedback loop was closed for the second system-identification experiment. The discrete-time models were identified for the 2000-Hz sample-and-hold rate. For identification, input-output sequences with 12,000 data points each (i.e., 6 s of data) were generated.

The disturbance actuator FSM 2 has dynamics very similar to those of FSM 1, but the control loops do not require a model of the disturbance actuator. Hence, the system identification uses data generated with FSM 2 fixed.

Experimental results showed negligible coupling between the two channels of each beam steering mirror; i.e., axis 1 commands produced negligible rotation about axis 2 and vice versa. Therefore, an uncoupled pair of single-input, single-output (SISO) transfer functions was identified for the open-loop model of FSM 1. The subspace method identified several higher order mirror modes, but their contribution to the input-output properties of the mirror were deemed insignificant for the purposes of the control. Therefore, a balanced truncation to two states for each mirror axis was chosen for control purposes. The frequency responses of these identified transfer functions are shown in Fig. 4. The true open-loop transfer function from the FSM 1 commands to the sensor outputs (i.e., the open-loop plant) is denoted by $P(z)$, and the identified open-loop plant model is denoted by $\hat{P}(z)$.

As discussed in Sec. 4, the feedback controller did not couple the mirror modes, so a second uncoupled pair of SISO transfer functions was identified with the feedback loop closed and used by the adaptive control loop. The true

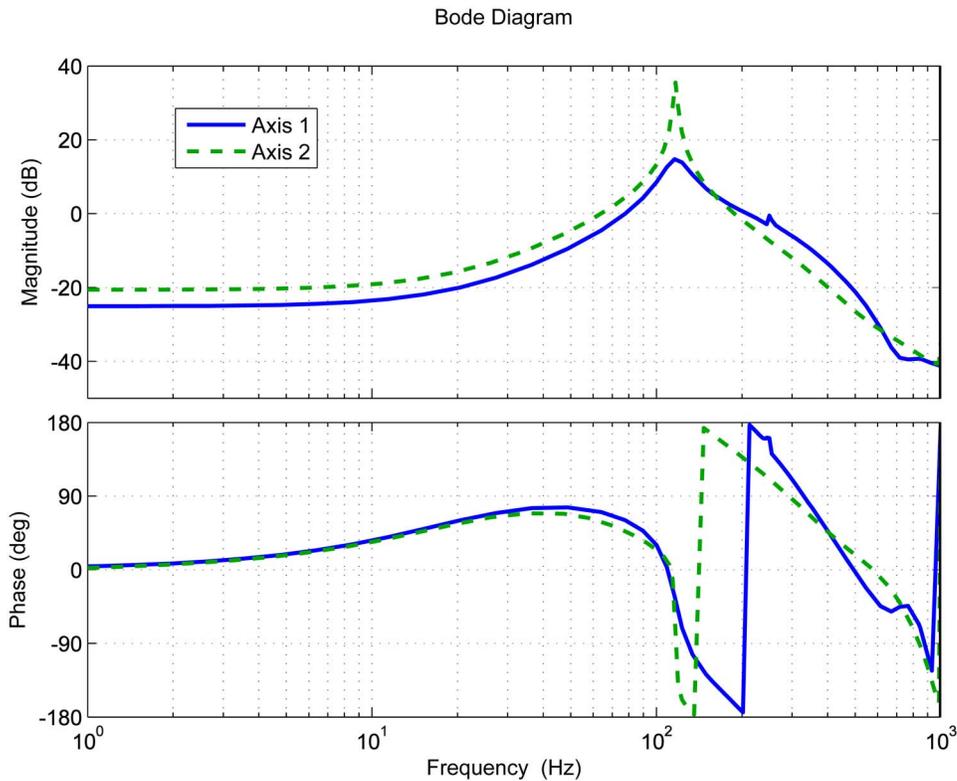


Fig. 5 Bode plots for identified transfer function $\hat{G}(z)$ of beam-steering mirror FSM 1 with LTI feedback loop closed.

closed-loop transfer function and identified closed-loop transfer function are denoted, respectively, by $G(z)$ and $\hat{G}(z)$. The frequency response of $\hat{G}(z)$ is shown in Fig. 5.

4 Control Design

In the control scheme for laser-beam steering presented here, an LTI feedback control loop is augmented by an adaptive control loop. The LTI feedback loop is a μ -synthesis controller designed to achieve two objectives: a disturbance-rejection bandwidth near the maximum achievable with LTI feedback control, and robust stabilization of the beam-steering system. In the adaptive loop, a multi-channel RLS lattice filter implicitly identifies the disturbance statistics in real time. The lattice filter was chosen because of its computational efficiency, numerical stability, and order-recursive structure.

4.1 LTI Feedback Loop

The LTI feedback system is shown in Fig. 6, where $P(z)$ is the open-loop plant and $C(z)$ is a μ -synthesis controller with four states for axis 1 and six states for axis 2. The input u in Fig. 6 is the pair of adaptive control commands, and the output y is the pair of beam displacements measured by the sensor. Four- and six-state μ -synthesis controllers were evaluated for each axis. For axis 1, the six-state controller performed no better than the four-state controller, but for axis 2, the six-state controller provided significantly better jitter rejection.

The discrete-time μ -synthesis method was used to design the feedback controller $C(z)$ to reject the noise w_0 in Fig. 6. This design was based on Fig. 7, where Δ represents the plant uncertainty, and $W_U(z)$ and $W_P(z)$ are the uncertainty and performance weighting filters employed in the design. The μ -synthesis design method uses μ -analysis of robust stability and performance to refine iteratively an H_∞ controller. The D-K iteration in the Matlab μ -Analysis and Synthesis toolbox¹⁰ was used. Also, the guidelines and insights for μ -synthesis design presented in Ref. 11 were followed to maximize the bandwidth of the closed-loop sensitivity function while maintaining robust stability. The uncertainty and performance weighting filters were

$$W_U = \frac{0.8280z - 0.6787}{z - 0.0045}, \tag{1}$$

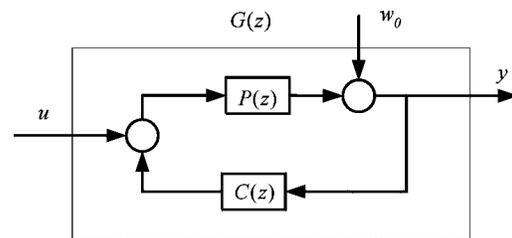


Fig. 6 Block diagram of LTI feedback control system: $P(z)$ =open-loop plant, $C(z)$ = μ -synthesis LTI feedback controller, and $G(z)=y/u$.

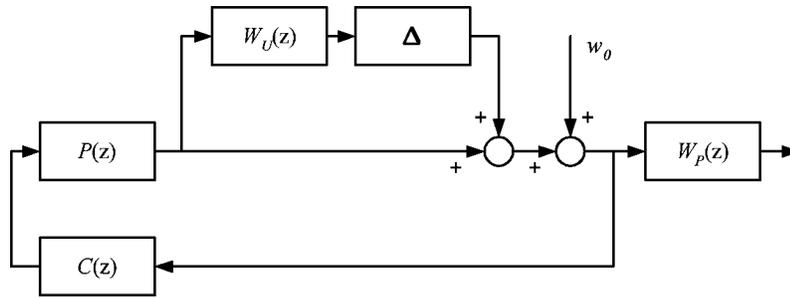


Fig. 7 Block diagram for μ -synthesis design of LTI feedback controller.

$$W_p = \frac{0.0083z^2 + 0.0165z + 0.0083}{z^2 - 1.8373z + 0.8439} \quad (2)$$

Figure 8 shows the two-channel sensitivity function for the modeled beam-steering system with the LTI feedback loop closed. The input for this transfer function is a pair of output disturbances represented by the signal w_0 in Fig. 6, and the output is the pair of measured beam displacements represented by the signal y in Fig. 6. This is the pertinent sensitivity transfer function, since in the experiment the disturbance is added to the beam after it leaves the control actuator. The sensitivity transfer function was computed using the identified open-loop plant model and the μ -synthesis feedback controller, without the adaptive controller.

The feedback controller here was designed to maximize the steady state disturbance-rejection bandwidth, so that

there is some associated amplification of higher frequency disturbance. This trade-off is common in high-performance controllers. The high-frequency amplification can be avoided by accepting less disturbance-rejection below 100 Hz, but a primary purpose of this paper is to demonstrate how the adaptive controller effectively extends the bandwidth of even a high-bandwidth LTI feedback controller.

Although maximizing the bandwidth of the closed-loop sensitivity function was the main objective in designing the LTI feedback controller, a good design should stabilize the lightly damped elastic mode associated with each mirror axis. Figure 5 shows that the feedback controller dampens the natural mode for axis 1 very effectively, reducing the peak by about 30 dB from that in Fig. 4. However, the feedback controller dampens the natural mode for axis 2 less, reducing the peak by about 6 dB from that in Fig. 4.

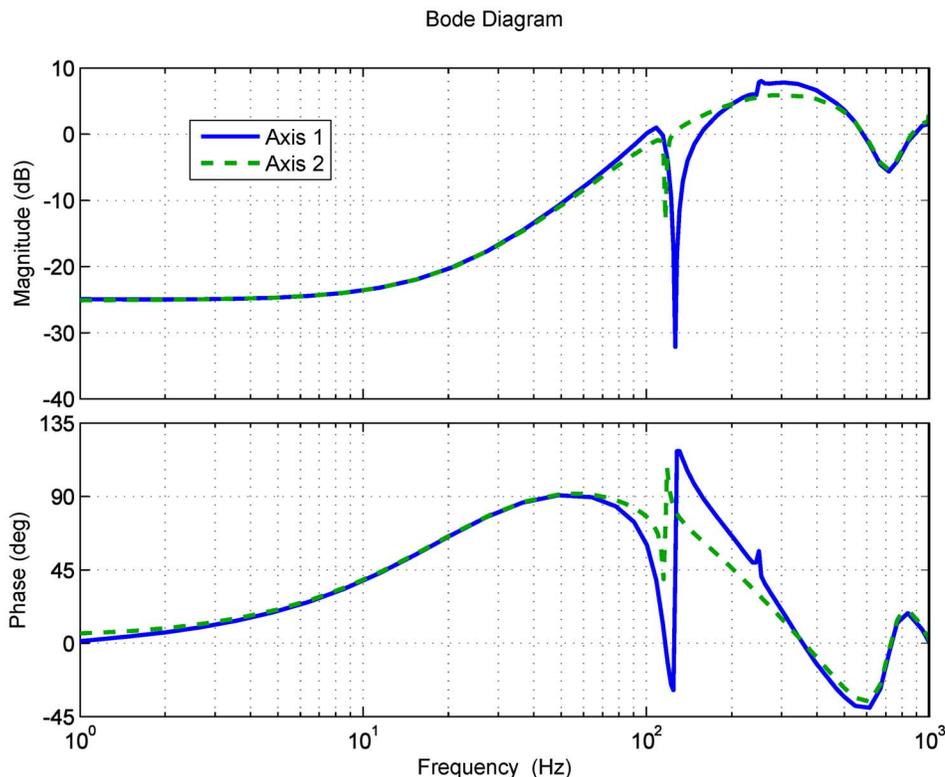


Fig. 8 Bode plots for the sensitivity transfer function $[I - \hat{P}(z)C(z)]^{-1}$.

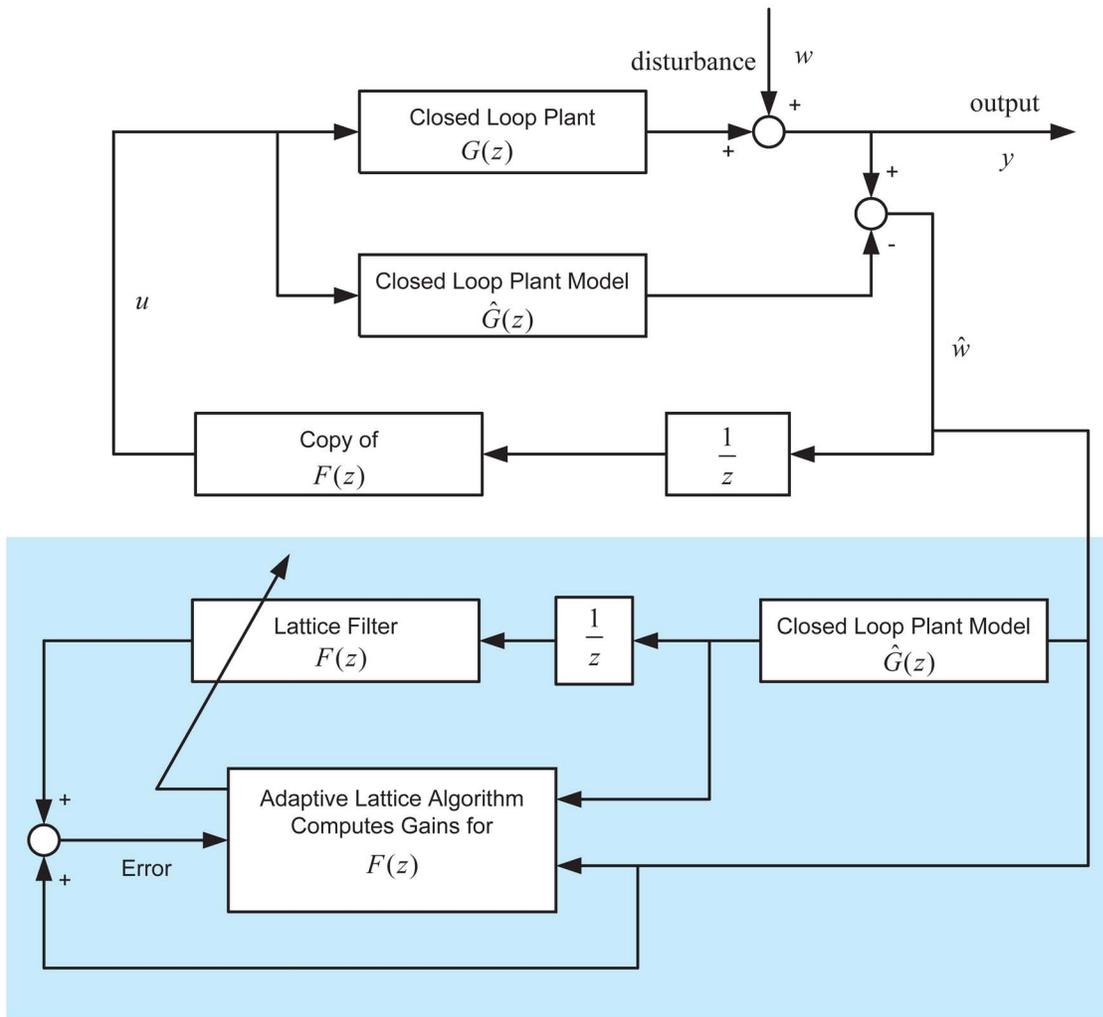


Fig. 9 Block diagram of adaptive control system.

The possible reasons for the poorer performance of LTI feedback loop for axis 2 include poor identification of the axis 2 transfer function, nonlinear actuator dynamics, and poor convergence of the D-K iteration in the μ -synthesis design process. In this experiment, it is possible to provide more damping to the axis 2 mode with different control designs, even a relatively simple PID (proportional integral derivative) feedback, but only at the expense of a lower disturbance-rejection bandwidth. However, the current feedback controller enables us to illustrate the performance of the adaptive loop when one axis is stabilized very well by the LTI feedback loop but one axis is not stabilized well. The experimental results show that the variable-order adaptive controller handles axis 2 well.

4.2 Adaptive Control Loop

In typical beam-steering applications, including adaptive optics and optical wireless communications, the dynamic models of the beam-steering mirrors either are known or can be determined by a one-time identification, as in Sec. 3. The disturbance characteristics, however, depend on the atmospheric conditions in the optical path and on the excited vibration modes of the structure on which the optical sys-

tem is mounted, so that the disturbance characteristics commonly vary during operation of the beam-steering system. Therefore, the adaptive control algorithm presented in this paper assumes known LTI plant dynamics but unknown disturbance dynamics. The adaptive controller requires an estimate $\hat{G}(z)$ of the closed-loop transfer function $G(z)$ in Fig. 6. The RLS lattice filter in the adaptive control loop tracks the statistics of the disturbance and identifies gains to minimize the root mean square (rms) value of the beam displacement.

The adaptive control scheme used here is similar in structure to the adaptive control schemes used in Ref. 6 for experimental adaptive control of a different type of beam steering mirror with much lower bandwidth than the MEMS mirror here, and in Refs. 12–14 for adaptive optics simulations where many sensor and control channels were used but with lower filter orders than used here. The main control-scheme innovation in this paper is the variable-order nature of the adaptive controller, which provides important improvements in transient response during adaptation, as experimental results here illustrate.

Figure 9 shows the structure of the adaptive control

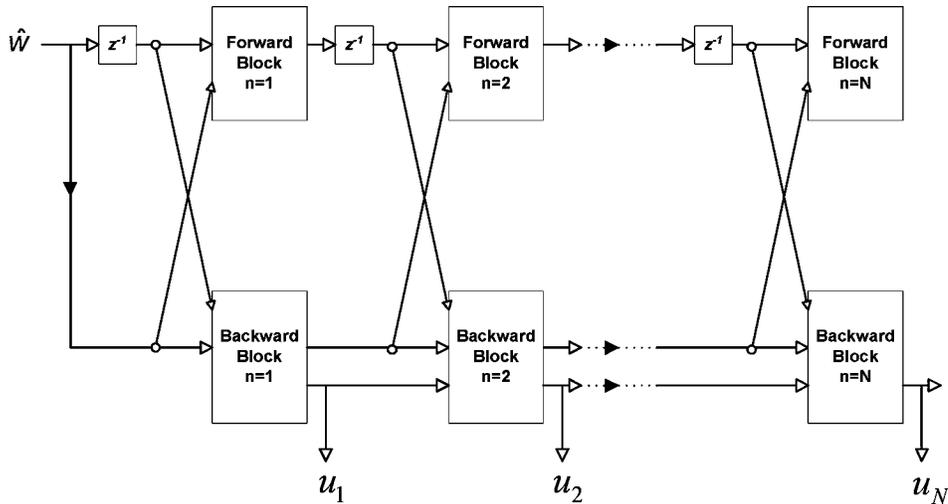


Fig. 10 FIR lattice filter generates adaptive control commands u_n for all filter orders $n \leq N$.

loop. The adaptive finite-impulse-response (FIR) filter $F(z)$ is the main component of the adaptive controller. As shown in the figure, the adaptive controller uses two copies of the FIR filter. The optimal filter gains are estimated in the bottom part of the block diagram in Fig. 9, and these gains are used by the FIR filter in the top part of Fig. 9 to generate the adaptive control signal u .

For the results presented in this paper, the two channels of the adaptive controller were uncoupled, although the adaptive lattice filter permits the use of multiple sensor channels, such as multi-axis beam positions and accelerometer measurements, for generating the command for each control channel. A comparison of a variety of experimental results for the jitter-control system here showed that coupling the two channels in the adaptive controller produced no improvement in steady-state performance but the FIR gains converged faster to optimal values for the uncoupled case because this case involves fewer FIR gains.

The disturbance signal w in Fig. 9 is related to the disturbance signal w_0 in Fig. 6 by

$$w = [I - P(z)C(z)]^{-1}w_0. \tag{3}$$

The true sensitivity transfer function $[I - P(z)C(z)]^{-1}$ is approximated closely by the transfer function in Fig. 8.

The lattice structure of the FIR filter that generates the adaptive control commands is illustrated in Fig. 10. The lattice realization of an FIR filter of order N consists of N identical stages cascaded, as in Fig. 10. The details of the algorithms represented by the blocks in Fig. 10 and the RLS lattice algorithm that updates the gains are beyond the scope of this paper. These algorithms are reparameterized versions of algorithms in Ref. 15. The current parameterization of the lattice algorithms is optimized for indefinite real-time operation. The current lattice filter maintains the channel orthogonalization in Ref. 15, which is essential to numerical stability in multichannel applications, and the unwindowed characteristic of the lattice filter in Ref. 15, which is essential to rapid convergence.

As indicated in Fig. 10, each stage of the lattice filter generates an adaptive control command. For $n \geq 1$, the out-

put u_n from the n 'th stage is the optimal control command if an FIR filter of order n is used in the adaptive control loop. For hardware implementation, a maximum filter order N is selected. In each real-time sampling interval, the lattice filter generates the adaptive control commands for all filter orders from 1 to N , and the control algorithm can select which command to use.

Note that, because of the order-recursive structure of the lattice filter, computing the control commands for all orders from 1 to N requires no more computation than computing the control command for the order N alone. Lattice filters are the only RLS algorithms with this property. Therefore, lattice filters are uniquely suited to variable-order adaptive control.

The lattice realization of the FIR filter is quite different from the most common FIR realization, which is

$$u_n = \sum_{k=1}^n B_{n,k}z^{1-k}\hat{w}. \tag{4}$$

While this realization is simpler than the lattice realization, it is less desirable for estimation of optimal gains in adaptive filtering and control. In particular, the optimal gains $B_{n,k}$ for an n 'th-order FIR filter of the form in Eq. (4) are not the first n gains among the optimal gains $B_{N,k}$ for an N 'th-order FIR filter ($N \geq n$). However, the optimal gains (or reflection coefficients) for the first n stages of the lattice filter are the same for all filter orders $N \geq n$.

The capability to vary the order of the filter in the adaptive controller is important because optimal gains can be identified faster for lower-order filters while higher-order filters are required for optimal steady state rejection of broadband disturbance. When the adaptive control loop is first closed or when it is adapting to changing disturbance statistics, lower order control commands should be used initially. The order of the control commands can be increased incrementally as the gains for the higher-order filter stages converge. This procedure, as demonstrated by the

experimental results in Sec. 5, eliminates large transient responses produced by initially incorrect gains in high-order filters.

5 Experimental Results

In the experiments described here, the sample-and-hold rate for control and filtering was 2000 Hz. Two different but partially correlated disturbance commands d_1 and d_2 were sent to the two axes of the disturbance actuator FSM 2. These two tilt command sequences had the form

$$\begin{bmatrix} d_1 \\ d_2 \end{bmatrix} = \begin{bmatrix} 4 & 1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}, \quad (5)$$

where the sequences v_1 and v_2 were obtained by passing independent white noise sequences through bandpass Butterworth filters. The 2-D jitter signal w_0 in Figs. 6 and 7 is the response of the disturbance actuator to the command sequences d_1 and d_2 .

Figures 11–14 compare the error of the beam position on the sensor for experiments with and without the adaptive loop closed, and for experiments with variable-order and fixed-order adaptive control. These comparisons were facilitated by the capability not only to send the same jitter sequences to the disturbance actuator repeatedly but to select the exact point in the disturbance sequence at which to start the adaptive controller. This capability was achieved by sending a pulse from computer 2 to computer 1 at the beginning of each experiment to synchronize the two clocks, although this procedure did produce a nonrepeatable and undetermined offset of less than one sampling interval between the clocks of computer 1 and computer 2.

In the experiments where the adaptive controller was used, the RLS lattice filter began running after 1 s, but no adaptive control commands were sent to the control actuator FSM 1 until 50 time steps later. Thus, the adaptive filter had 50 initial training steps to obtain initial estimates of the FIR gains before the adaptive control loop was closed at 1.025 s. The forgetting factor for RLS estimation was 0.99999. Note that, when the lattice filter began running at 1 s, it had no initial information about the statistics of the jitter or estimates of the FIR gains.

For variable-order adaptive control, Fig. 14 shows how the FIR order used for the control commands changed with time. The initial FIR order was $n=4$, and the order was incremented by 4 at the end of 50-step intervals until it reached the maximum FIR order $N=16$ at 1 s plus 200 steps, or 1.1 s. Thus, the FIR order used for control reached its maximum at 0.1 s after the lattice filter began running with no initial information about the disturbance statistics. For these and other similar experiments, the performance of the adaptive loop was evaluated with several maximum lattice-filter orders. The order 16 yielded better steady-state performance than lower orders, but orders higher than 16 yielded no significant further improvement.

Figures 11 and 12 show two typical sets of experimental results. The time series plotted compare the laser beam position error at the sensor produced by the LTI feedback loop alone with the error produced by the combined LTI feedback loop and variable-order adaptive controller. The PSD plots show the frequency content of the steady-state output errors. The open-loop errors represented by the black

curves in the bottom plots in Fig. 11 were measured with neither control loop closed; hence the open-loop errors are the jitter added to the beam by the disturbance actuator. Table 1 gives the rms values of the steady state error signals in the experiments in Fig. 11. In these experiments, the steady state amplitudes of the mirror displacements were approximately 0.2 deg.

The jitter bandwidths noted in the caption for Fig. 11 refer to the jitter command sequences d_1 and d_2 sent by computer 2 to the disturbance actuator FSM 2. The open-loop PSD plots represent the output w_0 of FSM 2. The transfer function of FSM 2 is very similar to that of the control actuator FSM 1, which is shown in Fig. 4. While FSM 1 is controlled by the LTI and adaptive controllers, no control loop is closed on FSM 2, which is driven only by the jitter command sequences. Hence, even though the jitter commands sent to FSM 2 in the experiments represented by the plots on the right in Fig. 11 had very little power in the vicinity of 110 to 130 Hz, the lightly damped elastic modes of the gimbal for FSM 2 produced the peaks in this range shown in the PSDs on the right in Fig. 11.

The time series in Fig. 11 show rapid convergence to optimal steady-state performance, which this adaptive control algorithm has produced consistently in experiments. In the plots on the right in Fig. 11, the transient response between 1 and 1.25 s is significantly better for axis 1 than for axis 2 due to the lightly damped axis 2 natural mode; however, even for axis 2, the adaptive controller achieves near optimal jitter rejection after less than 0.3 s, or 600 samples. In all cases, the variable-order adaptive controller produces large reductions in position error within 0.1 s.

The PSDs show that, as predicted by Fig. 8, the LTI feedback loop significantly reduces the jitter below about 70 Hz but amplifies jitter above about 150 Hz. The PSDs also show that the adaptive loop yields significant jitter reduction above 70 Hz, thereby extending the bandwidth of the feedback loop. This extended bandwidth accounts for the significant reduction in the amplitudes of the output errors achieved by the adaptive controller, as shown in the plots of the time series.

Another noteworthy point in the PSDs in Fig. 11 is that, while both the feedback loop and the adaptive loop amplify jitter above 150 Hz, the high-frequency amplification often is greater for the adaptive loop—but only where there is very little jitter to begin with, so that the high-frequency jitter is still small with the adaptive loop. This point can be seen also from the zoomed time series in Fig. 12.

The explicit objective of the RLS lattice filter is to minimize the rms values of the output error, and this is accomplished. As is well known, filters and controllers that minimize the rms values of error signals tend to whiten residual errors, so that optimum filters and controllers typically accept some amplification in frequency bands where the disturbance is very low to be able to achieve large reductions in the dominant disturbance power. This is true for minimum variance LTI feedback controllers as well as for high-performance adaptive controllers like the one used here. As the time series show, the jitter in bandwidths with low power is not amplified enough to make it significant.

The adaptation of the lattice filter to the jitter statistics is illustrated particularly well by the PSDs in the 350 - to 360-Hz frequency range. In the plots on the left in Fig.

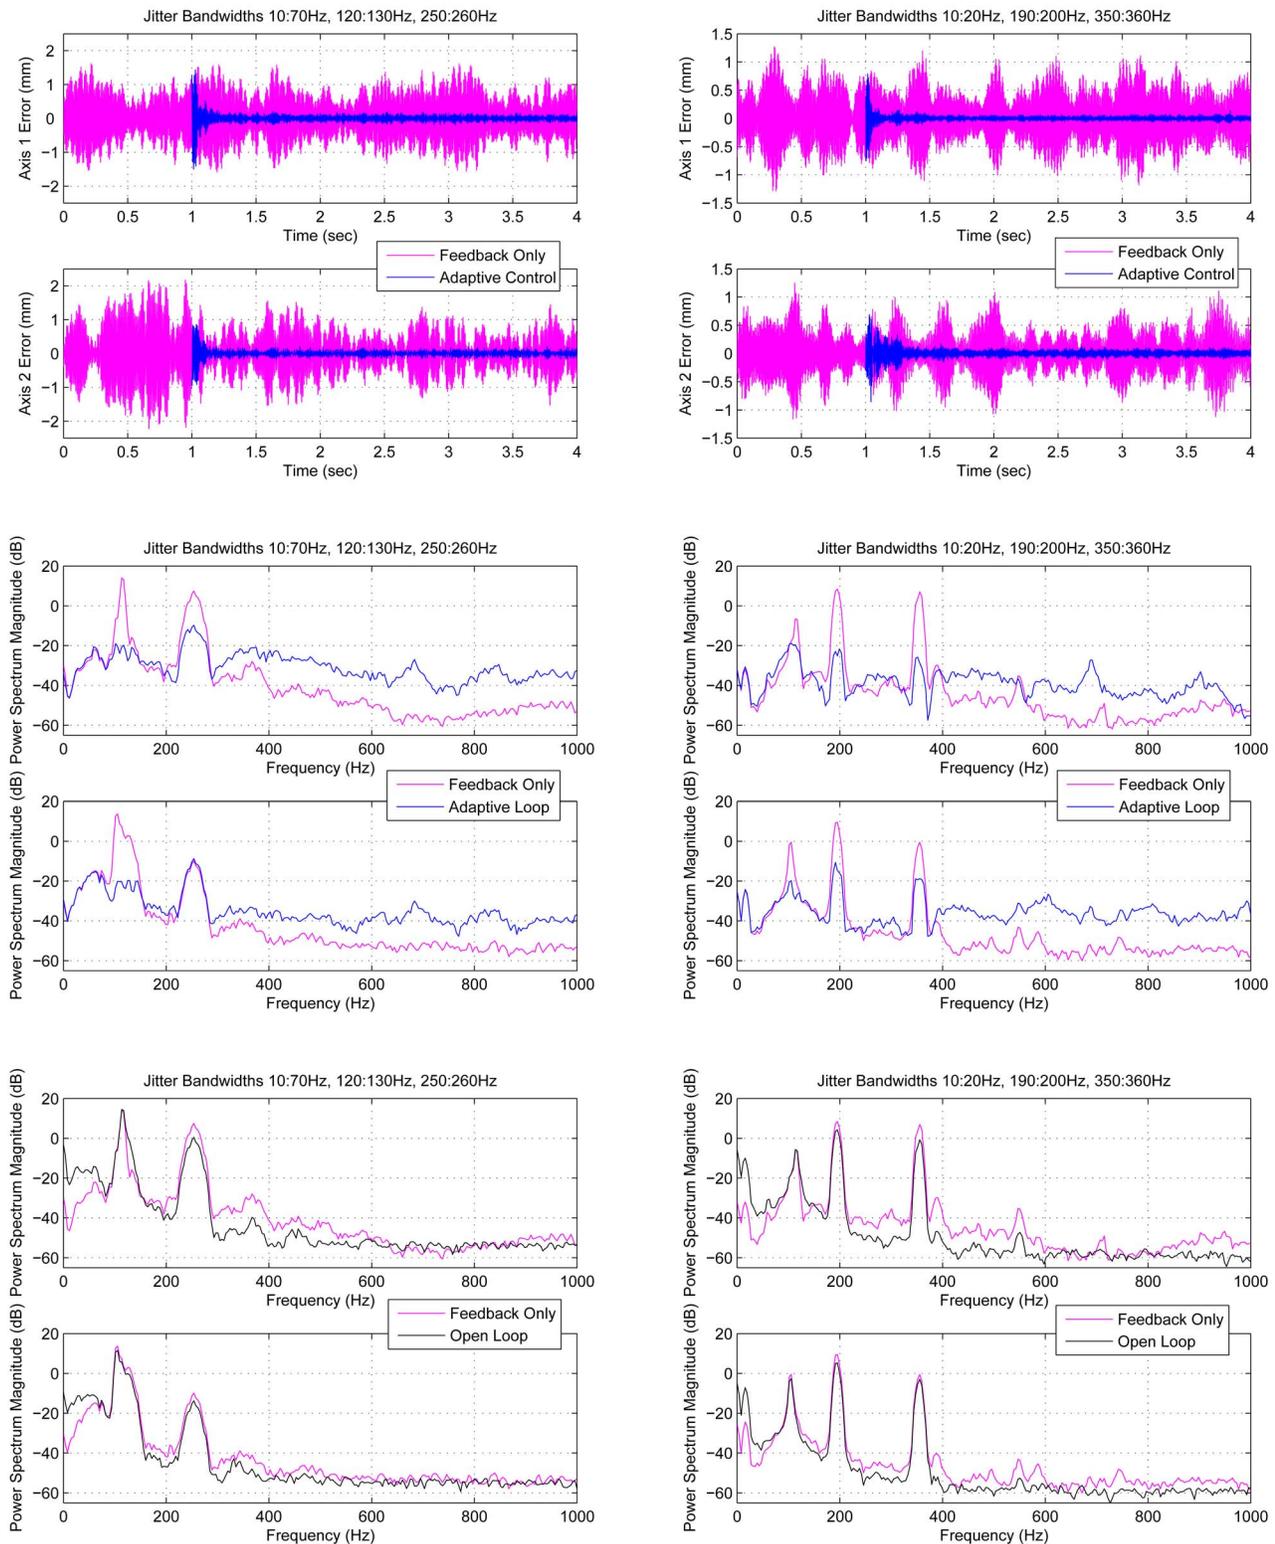


Fig. 11 Left: jitter bandwidths 10:70, 120:130, and 250:260 Hz; right: jitter bandwidths 10:20, 190:200, and 350:360 Hz. Top plots: time series compare output errors produced by LTI feedback only and by variable-order adaptive control combined with LTI feedback. Lattice filter starts running at 1 s; adaptive control loop is closed at 1.025 s. Middle plots: power spectral densities (PSDs) compare output errors produced by LTI feedback only to those produced by variable-order adaptive control combined with LTI feedback; bottom plots: PSDs compare open-loop output errors to those produced by LTI feedback only. (PSDs computed for last 2.5 s = 5000 samples).

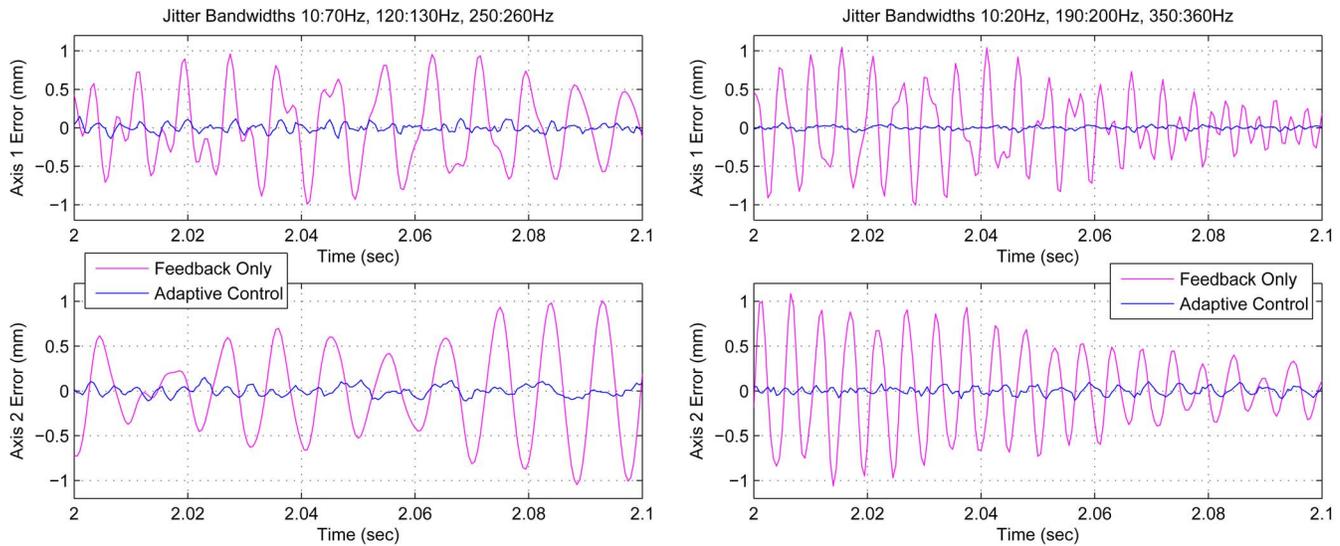


Fig. 12 Zoomed views of error sequences in Fig. 11. Left: jitter bandwidths 10:70, 120:130, and 250:260 Hz; right: jitter bandwidths 10:20, 190:200, and 350:360 Hz.

11, the adaptive control loop amplifies the negligible jitter in this range, but only to a level consistent with minimizing the net rms value of the output error over all frequencies. On the other hand, in the plots on the right in Fig. 11, a large fraction of the jitter power lies in the 350- to 360-Hz range, so the adaptive loop produces the large reductions of output-error power shown for both axes in this frequency range.

The value of the variable-order capability of the adaptive controller is illustrated by Fig. 13, where the curves for LTI feedback control only (first 1 s) and variable-order adaptive control are the same as in Fig. 11. The plots in Fig. 13 show also the output errors produced by the adaptive

controller with fixed order 16, which is the final order of the variable-order adaptive controller. The fixed-order adaptive controller produced large output errors immediately after the adaptive control loop was closed because the lattice filter did not yet have enough data to identify near-optimal gains for the 16th-order control law. The effect is severe in both channels, but more severe for axis 2 because the feedback loop provided so little damping for this axis.

6 Conclusions

We presented a new method for adaptive control of jitter in laser beams. The method has been demonstrated by results

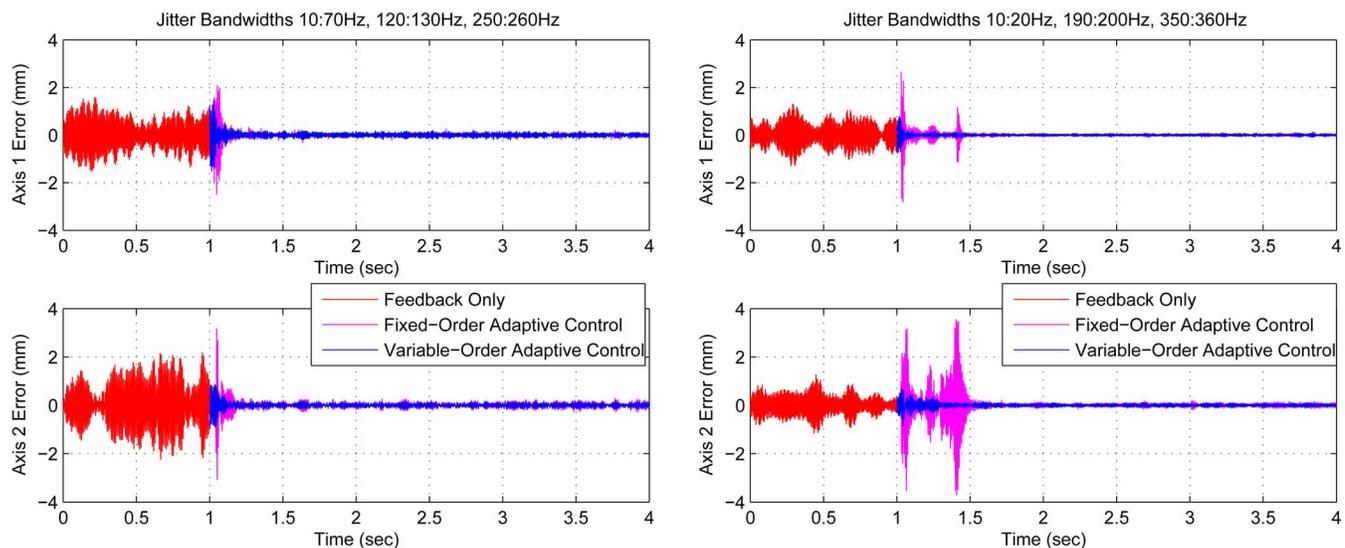


Fig. 13 Comparison of variable-order and fixed-order adaptive control. Left: jitter bandwidths 10:70, 120:130, and 250:260 Hz; right: jitter bandwidths 10:20, 190:200, and 350:360 Hz. Maximum lattice-filter order=16. The lattice filter starts running at 1 s; the adaptive control loop is closed at 1.025 s. The curves for LTI feedback control only and variable-order adaptive control are the same as in Fig. 11.

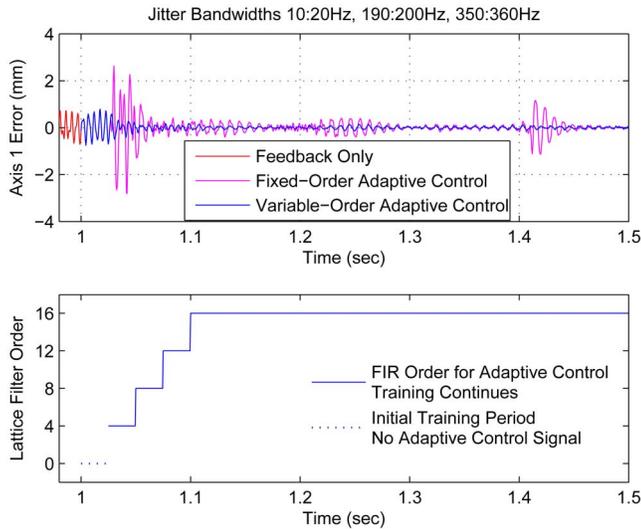


Fig. 14 Top: zoomed view of the time series in the top right plot in Fig. 13; bottom: order of the FIR lattice filter that generated the variable-order adaptive control signal. At the beginning of the initial training period, the lattice filter had no initial information about the statistics of the jitter or estimates of the FIR gains. The adaptive control loop was closed with FIR order $n=4$ after the 50-step ($=0.025$ s) initial training period.

from an experiment employing two-axis MEMS tilt mirrors. Laser beam jitter is rejected by a μ -synthesis feedback controller augmented by the lattice-filter-based adaptive control loop. The rms level of the output error is minimized by the adaptive control loop, which implicitly identifies the disturbance statistics from real-time sensor data and determines control gains that are optimal for the current disturbance acting on the laser beam. Experimental results demonstrate that the adaptive controller significantly extends the disturbance rejection bandwidth achieved by the feedback controller alone. The adaptive lattice filter can perform high-order, multichannel RLS computation in real time at high sampling rates, and the RLS algorithm yields

Table 1 The rms values (in millimeters) of output errors for the last 5000 samples.

Jitter bandwidths 10:70, 120:130, and 250:260 Hz		
	Axis 1 rms	Axis 2 rms
Open loop	0.5147	0.4190
Feedback only	0.5338	0.5274
Adaptive control	0.0591	0.0595
Jitter bandwidths 10:20, 190:200, and 350:360 Hz		
	Axis 1 rms	Axis 2 rms
Open loop	0.2010	0.2098
Feedback only	0.3487	0.3160
Adaptive control	0.0254	0.0374

faster convergence to optimal gains than does the LMS method, which is used more commonly in adaptive disturbance-rejection applications.

An important feature of the adaptive control scheme is the variable-order nature, which exploits the order-recursive structure of the lattice filter. The variable-order adaptive controller exhibits faster RLS adaptation with initial low filter orders without sacrificing the optimal steady-state performance of high-order adaptive filters. This improved adaptation is particularly important in practical applications where large transients cannot be tolerated. In some experiments with the fixed-order adaptive controller, the transients were so large that the laser beam went out of the sensor range. As illustrated by the experimental results reported here, the variable-order adaptive controller eliminates large transient bursts produced by the fixed-order adaptive controller during the initial adaptation.

Acknowledgments

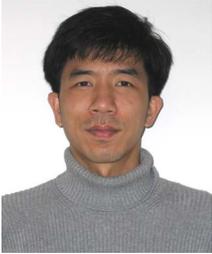
The authors are indebted to Texas Instruments for providing the TALP1000A MEMS steering mirrors used in this research. This work was supported by the U.S. Air Force Office of Scientific Research under Grants F49620-02-01-0319 and F-49620-03-1-0234.

References

1. M. C. Roggemann and B. Welsh, *Imaging through Turbulence*, CRC, New York (1996).
2. R. K. Tyson, *Principles of Adaptive Optics*, Academic Press, New York (1998).
3. B. L. Ellerbroek, "First-order performance evaluation of adaptive optics systems for atmospheric turbulence compensation in extended field-of-view astronomical telescopes," *J. Opt. Soc. Am. A* **11**, 783–805 (1994).
4. R. Q. Fugate, B. L. Ellerbroek, et al., "Two generations of laser guide star adaptive optics experiments at the starfire optical range," *J. Opt. Soc. Am. A* **11**, 310–324 (1994).
5. M. A. McEver, D. G. Cole, and R. L. Clark, "Adaptive feedback control of optical jitter using Q-parameterization," *Opt. Eng.* **43**(4), 904–910 (2004).
6. B.-S. Kim, S. Gibson, and T.-C. Tsao, "Adaptive control of a tilt mirror for laser beam steering," in *Proc. Am. Control Conf.*, pp. 3417–3421, IEEE, Boston, MA (2004).
7. N. O. Pérez Arancibia, S. Gibson, and T.-C. Tsao, "Adaptive control of MEMS mirrors for beam steering," in *Proc. ASME Int. Mechanical Engineering Congress and Exposition (IMECE2004)*, Anaheim, CA, Nov. 2004, IMECE2004–60256, ASME, New York (2004).
8. Y. M. Ho, G. Xu, and T. Kailath, "Fast identification of state-space models via exploitation of displacement structure," *IEEE Trans. Autom. Control* **39**(10), 2004–2017 (1994).
9. P. Van Overschee and B. De Moor, *Subspace Identification for Linear Systems*. Kluwer Academic, Norwell, MA (1996).
10. G. J. Balas, J. C. Doyle, K. Glover, A. Packard, and R. Smith, *μ -Analysis and Synthesis Toolbox*, Mathworks Natick, MA (1995).
11. B.-S. Kim and T.-C. Tsao, "A performance enhancement scheme for robust repetitive control system," *ASME J. Dyn. Syst., Meas., Control* **126**(1), 224–229 (2004).
12. J. S. Gibson, C.-C. Chang, and B. L. Ellerbroek, "Adaptive optics: wavefront correction by use of adaptive filtering and control," *Appl. Opt.* **39**, 2525–2538 (2000).
13. J. S. Gibson, C.-C. Chang, and N. Chen, "Adaptive optics with a new modal decomposition of actuator and sensor spaces," in *Proc. Am. Control Conf.*, pp. 4619–4625, IEEE, Arlington, VA (2001).
14. Y.-T. Liu and S. Gibson, "Adaptive optics with adaptive filtering and control," in *Proc. Am. Control Conf.*, pp. 3176–3179, IEEE, Boston, MA (2004).
15. S.-B. Jiang and J. S. Gibson, "An unwindowed multichannel lattice filter with orthogonal channels," *IEEE Trans. Signal Process.* **43**(12), 2831–2842 (1995).



Néstor O. Pérez Arancibia received his Ingeniero and MEng degrees from the Pontificia Universidad Católica de Chile in 2000 and since 2002 he has been a PhD student with the Mechanical and Aerospace Engineering Department at the University of California, Los Angeles. His current interests include feedback control, adaptive filtering, and mechatronics.



signal processing.

Neil Y. Chen received his BS degree in mechanical engineering in 1993 from National Taiwan University, his MS degree in mechanical engineering in 1995 from National Chiao-Tung University, Taiwan, and his PhD degree in mechanical and aerospace engineering in 2001 from University of California, Los Angeles (UCLA). He has been a postdoctoral scholar at UCLA since 2001. His research interests include system identification, adaptive filtering and control, and



and in 1977 joined the faculty of the University of California, Los

James S. Gibson received his BS degree in aerospace engineering in 1970, his MS degree in engineering mechanics in 1972, and his PhD degree in engineering mechanics in 1975, all from the University of Texas at Austin. He has served on the faculties of the Aerospace Engineering and Engineering Mechanics Department at the University of Texas at Austin and the Engineering Science and Mechanics Department at Virginia Polytechnic Institute and State University,

Angeles (UCLA), where he currently is a professor of mechanical and aerospace engineering. Professor Gibson's research interests include control and identification of dynamical systems, adaptive filtering and signal processing, and infinite-dimensional system theory. His research applications include identification and control of flexible structures, control of laser beams and adaptive optics, modeling, identification, and control of microinertial sensors, control of fluid flow, and active noise control. He has been an associate editor for the *SIAM Journal on Control and Optimization* and for the *IEEE Transactions on Automatic Control*.



Tsu-Chin Tsao received his BS degree in engineering in 1981 from National Taiwan University and his MS degree in 1984 and his PhD degree in 1988 in mechanical engineering from the University of California at Berkeley. He served 11 years on the faculty of the Mechanical and Industrial Engineering Department at the University of Illinois at Urbana-Champaign. In 1999 he joined the faculty of the University of California, Los Angeles (UCLA), where he currently is a professor with the Mechanical and Aerospace Engineering Department. Professor Tsao's research areas are control systems and mechatronics. Recognitions of his research include 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 ASME Dynamic Systems and Control Division in 1997, and the Hugo S. Shuck Best Paper Award from American Automatic Control Council in 2002.