

Adaptive optics based on analog parallel stochastic optimization: analysis and experimental demonstration

Mikhail A. Vorontsov and Gary W. Carhart

U.S. Army Research Laboratory, Computational and Informational Sciences Directorate, Adelphi, Maryland 20783

Marc Cohen and Gert Cauwenberghs

Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, Maryland 21218

Received September 23, 1999; accepted February 15, 2000; revised manuscript received March 1, 2000

Wave-front distortion compensation using direct system performance metric optimization is studied both theoretically and experimentally. It is shown how different requirements for wave-front control can be incorporated, and how information from different wave-front sensor types can be fused, within a generalized gradient descent optimization paradigm. In our experiments a very-large-scale integration (VLSI) system implementing a simultaneous perturbation stochastic approximation optimization algorithm was applied for real-time adaptive control of multielement wave-front correctors. The custom-chip controller is used in two adaptive laser beam focusing systems, one with a 127-element liquid-crystal phase modulator and the other with beam steering and 37-control channel micromachined deformable mirrors. The submillisecond response time of the micromachined deformable mirror and the parallel nature of the analog VLSI control architecture provide for high-speed adaptive compensation of dynamical phase aberrations of a laser beam under conditions of strong intensity scintillations. Experimental results demonstrate improvement of laser beam quality at the receiver plane in the spectral band up to 60 Hz. © 2000 Optical Society of America [S0740-3232(00)00208-8]

OCIS codes: 010.0010, 010.1080, 110.0110.

1. INTRODUCTION

There is a long-term standoff in adaptive optics between two wave-front control paradigms. One approach, known as wave-front conjugation (wave-front reversal), is based on the reciprocity principle.^{1,2} The second approach, which is the primary subject of this paper, is wave-front control using direct system performance metric optimization.

According to the reciprocity principle, wave-front distortions accumulated along the wave propagation path can be compensated by conjugating the received-wave complex amplitude $\psi(\mathbf{r}, t) = |\psi(\mathbf{r}, t)|\exp[i\phi(\mathbf{r}, t)]$, where $|\psi(\mathbf{r}, t)|$ and $\phi(\mathbf{r}, t)$ are the wave amplitude and the phase at the receiver plane and $\mathbf{r} = \{x, y\}$ is a vector in the plane orthogonal to the direction z of wave propagation. For adaptive systems with a transmitted wave, wave-front conjugated correction means that $A(\mathbf{r}, t) = \alpha\psi^*(\mathbf{r}, t)$, where $A(\mathbf{r}, t)$ is the outgoing wave and α is a coefficient. Thus the wave-front conjugation condition contains both the phase conjugation $u(\mathbf{r}, t) = -\phi(\mathbf{r}, t)$ and the amplitude modulation $|A(\mathbf{r}, t)| = \alpha|\psi(\mathbf{r}, t)|$. Most optoelectronic wave-front conjugation adaptive systems correct only phase and ignore the second requirement posed by the reciprocity principle: amplitude correction.³ This violation of the reciprocity principle does not impact adaptive system performance if the amplitude modulation $|\psi(\mathbf{r}, t)|$ is small—a condition common for astronomical imaging. Phase-only correction is effective if the propagation distance L in the phase-distorting medium is relatively short, i.e., $L \ll kr_0^2$

(Fresnel or near-field diffraction), and phase distortions having a characteristic spatial scale r_0 do not result in strong intensity scintillation.⁴

Under conditions of strong intensity modulation typical for ground-to-ground imaging, laser communication, laser-beam-forming systems, etc., we face at least two additional problems. First, phase-only correction distinctly violates the reciprocity principle, and it is unclear how effective such phase conjugation correction could be. Second, strong intensity scintillations impose serious problems for wave-front measurement techniques.⁵ The problem is not only reduction of the wave-front sensor signal-to-noise ratio. Under conditions of strong intensity scintillations, there may be wave-front dislocations (branch points) offering a challenging problem for detection and replication with a wave-front corrector (especially with the use of continuous faceplate deformable mirrors).^{6,7}

The irony of the situation is that despite tremendous efforts that have been made to measure and reproduce the wave front, it is actually unclear how the phase-only conjugation strategy impacts actual adaptive system performance (image quality, energy density, etc.) under the conditions of strong intensity scintillations. By ignoring wave-front amplitude control, we have implicitly agreed that perfect (diffraction-limited) correction cannot occur. In this sense, then, it seems that there is not much point in strictly adhering to the phase conjugation rule. Rather than performing explicit phase conjugation by equalizing the residual phase $[\delta(\mathbf{r}, t) = \phi(\mathbf{r}, t)$

+ $u(\mathbf{r}, t) \rightarrow \text{constant}$] through measurement, reconstruction, and conjugation of the wave front, the phase correction strategy can be based on direct optimization of a system performance metric $J[u]$ (image quality, energy density, etc.). This approach, known in the early days of adaptive optics as aperture tagging or image-sharpening techniques^{2,8} and at the present time known as wave-front control based on model-free or gradient descent optimization,⁹ has been rather unfairly neglected for decades. This situation is rapidly changing. Several factors contribute to the renaissance of model-free optimization techniques for adaptive optics:

1. Problems related with use of the wave-front conjugation technique in the presence of strong intensity scintillations, as mentioned above, have stimulated a search for new adaptive system architectures.

2. For a number of new adaptive optics applications such as anisoplanatic and active imaging,^{10,11} laser beam focusing through a phase-distorting medium on an extended object,¹² long-distance ground-to-ground laser communication, etc., wave-front conjugation cannot be used in a straightforward manner for wave-front control, at least not without involving the rather sophisticated guide star technique.¹³

3. The appearance of inexpensive wave-front correctors based on microelectromechanical systems¹⁴ and liquid-crystal (LC) phase modulators¹⁵ has created the expectation that the entire adaptive system can be fast, small, and inexpensive. From this perspective the wave-front sensor and the phase reconstruction hardware are the major obstacles to achieving this goal.

4. Recently developed efficient parallel model-free optimization algorithms¹⁶⁻²⁰ and their hardware implementation in very-large-scale integration (VLSI) systems²¹⁻²³ offer efficient signal processing architectures for adaptive optics applications, stimulating the transition to nonconjugated adaptive optics.²⁴

In this paper we consider both theoretical and experimental aspects of adaptive wave-front phase distortion compensation using direct system performance metric optimization. In Section 2 we present a brief review of model-free optimization techniques in adaptive optics, with an emphasis on gradient descent methods based on stochastic approximation of the true gradient. The method is extended to include different requirements for wave-front control as well as information from different wave-front sensor types. It is shown that both can be incorporated into a more general gradient descent optimization paradigm, giving rise to a variety of new control system architectures. In Section 3 we report on the first experiments with adaptive wave-front correction systems based on simultaneous perturbation stochastic approximation (SPSA) optimization using a VLSI adaptive controller. The VLSI system was used in two different adaptive system configurations: one with a 127-element LC phase modulator and the other with beam steering and 37-channel micromachined deformable mirrors. The submillisecond response time of the beam steering and the micromachined deformable mirrors, along with fast (up to 1500 iterative steps per second) parallel analog control signal computation using the VLSI system, al-

lowed for adaptive compensation of dynamical laser beam phase aberrations under conditions of strong intensity scintillations, as described in Section 4.

2. GRADIENT DESCENT OPTIMIZATION IN ADAPTIVE OPTICS

A. Gradient Descent

Wave-front distortion compensation can be obtained by directly optimizing a measured characteristic J of an optical system's performance (system performance metric, cost functional). Dependent on the type of adaptive system, the performance metric might be intensity radiation at the focus,^{2,12} image sharpness,^{8,24} or scattered field statistical moments.²⁵ Gradient descent optimization of $J[u]$ can be described as a time-dependent evolution of the controlling phase $u(\mathbf{r}, t)$ in the direction opposite to the first variation $\delta_u J$ of the cost functional:

$$\tau \frac{\partial u(\mathbf{r}, t)}{\partial t} = -\delta_u J, \quad (1)$$

where τ is a time constant. If the dependence $J[u]$ is known or if it can be defined from a known system's mathematical model, the variation $\delta_u J$ can be directly calculated. In this case the dynamical process (1) describes the well-known continuous-time gradient descent optimization technique widely used in image processing applications.^{26,27} In adaptive optics both the dependence $J[u]$ and the system model that includes wave-front distortions are unknown (model free), and the variation $\delta_u J$ should be determined based on measured data.

The dimensionality of the control law (1) can be reduced by representing the controlling wave front u in the following form:

$$u(\mathbf{r}, t) = \sum_{j=1}^N u_j(t) S_j(\mathbf{r}), \quad (2)$$

where $u_j(t)$ are control signals that may (or may not) describe voltages applied to wave-front corrector electrodes and $S_j(\mathbf{r})$ are the corresponding response functions. Then, instead of Eq. (1), we have a system of ordinary nonlinear differential equations describing control signal evolution during the adaptation process:

$$\tau_j \frac{du_j(t)}{dt} = -\gamma J'_j(u_1, \dots, u_N), \quad j = 1, \dots, N, \quad (3)$$

where τ_j are time constants and $J'_j = \partial J / \partial u_j$ are gradient projections of the cost function $J(u_1, \dots, u_N)$. Here the update coefficient γ is positive for system metric minimization and negative otherwise. In the vicinity of an extremum point, the system (3) is always stable, as $J(u_1, \dots, u_N)$ is a Lyapunov function. Indeed, from Eq. (3) we have

$$\frac{dJ(t)}{dt} = \sum_{j=1}^N \frac{\partial J}{\partial u_j} \frac{du_j}{dt} = -\gamma \sum_{j=1}^N \left(\frac{\partial J}{\partial u_j} \right)^2 \tau_j^{-1} \leq 0 \quad (\text{for positive } \gamma). \quad (4)$$

The inequality (4) provides for monotonic decrease in J for the case of metric minimization (positive γ) and increase otherwise. The discrete version of Eq. (3) reads as

$$u_j^{(n+1)} = u_j^{(n)} - \gamma J'_j(u_1^{(n)}, \dots, u_N^{(n)}), \quad j = 1, \dots, N. \quad (5)$$

The stability of the dynamical system (3) in the vicinity of extrema does not guarantee fast convergence. Neither is it guaranteed that the system converges to the global rather than a local extremum.

B. Gradient Approximation

In adaptive optics the true gradient components J'_j of the cost function are unknown and should be estimated based on available (measured) information. An estimate (approximation) of the gradient components $\{\tilde{J}'_j\}$ can be obtained by using the aperture-tagging technique. Two methods of aperture tagging were known in the early days of adaptive optics: time division (sequential) and frequency division (multidither technique).^{2,28}

In aperture tagging based on sequential perturbations, gradient estimation of a cost function is achieved by sequentially applying small perturbations δu_j to the control signals and measuring the corresponding changes in the system performance metric:

$$\begin{aligned} \delta J_j &= J(u_1, \dots, u_j + \delta u_j, \dots, u_N) \\ &- J(u_1, \dots, u_j, \dots, u_N), \quad j = 1, \dots, N. \end{aligned}$$

The measured metric perturbation is used for finite-difference gradient estimation: $\tilde{J}'_j = \delta J_j / \delta u_j$. For the sequential perturbation method, the signal-to-noise ratio and hence the accuracy of wave-front control are independent of the number of control channels N , but the time required for gradient estimation is proportional to N .²⁹

The multidithering method allows gradient estimation in parallel. Small perturbations $\delta u_j = \alpha \sin(\omega_j t)$ in the form of harmonic signals with different dithering frequencies ω_j (dithered carriers) and small modulation amplitudes α are applied simultaneously (in parallel) to all control channels. For small modulation amplitudes, the perturbed system performance metric δJ reads as

$$\begin{aligned} \delta J &= \alpha \sum_{j=1}^N \frac{\partial J}{\partial u_j} \sin(\omega_j t) \\ &+ \frac{\alpha^2}{2} \sum_{j,k}^N \frac{\partial^2 J}{\partial u_j \partial u_k} \sin(\omega_j t) \sin(\omega_k t) + \dots \quad (6) \end{aligned}$$

To determine (estimate) gradient components, dithered carriers are demodulated by synchronous detectors and passed through low-pass filters.²⁸ The time required for gradient estimation in the multidithering method does not depend on the number of control channels—a great improvement when compared with the sequential perturbation method.

The multidithering technique has two major drawbacks. To provide efficient signal demodulation, the dithered carriers should be separated by approximately twice the control system bandwidth ω_0 . For an adaptive system with segmented (piston-type) mirrors, the maximum dithering frequency is $\omega_{\max} = [10 + 1.6(N$

$- 1)]\omega_0$.³⁰ As the number of control channels increases, this puts a severe requirement on the deformable mirror bandwidth. For a system with 100 control channels and a desired control system bandwidth of 300 Hz, the required phase corrector bandwidth is approximately 50 kHz. The other problem with the multidithering technique is that the signal-to-noise ratio decreases proportionally with the number of control channels,³¹ thus requiring a high photon-flux density from the reference source.

The mentioned drawbacks of both sequential and multidithering methods for gradient estimation explain why direct system performance metric optimization was almost completely abandoned in favor of phase-conjugation-type systems.

C. Parallel Stochastic Perturbative Gradient Descent

With the recent appearance of the parallel stochastic perturbative gradient descent [also called simultaneous perturbation stochastic approximation (SPSA)] technique,¹⁷⁻¹⁹ in the realm of more general stochastic approximation optimization methods,¹⁶ the current status of wave-front control methods in adaptive optics is bound to change.^{24,32,33}

The parallel stochastic perturbative gradient descent algorithm differs from gradient descent (3) and (5) in the way in which the gradient estimate is obtained. In this method small stochastic perturbations δu_j are applied in parallel (simultaneously) to all control channels. The measured value of the cost function perturbation

$$\begin{aligned} \delta J &= J(u_1 + \delta u_1, \dots, u_j + \delta u_j, \dots, u_N + \delta u_N) \\ &- J(u_1, \dots, u_j, \dots, u_N) \quad (7) \end{aligned}$$

is used directly for gradient estimation: $\tilde{J}'_j = \delta J \delta u_j$. From Eq. (7), with the use of a Taylor expansion,

$$\tilde{J}'_j = \delta J \delta u_j = \frac{\partial J}{\partial u_j} (\delta u_j)^2 + \sum_{k \neq j}^N \frac{\partial J}{\partial u_k} \delta u_k \delta u_j + \dots \quad (8)$$

This expression actually is quite general and applies also to the sequential perturbation and multidithering gradient approximation methods. To yield a fair approximation of the true gradient $\tilde{J}'_j \approx c_j (\partial J / \partial u_j)$, where $c_j = (\delta u_j)^2 = \text{constant}$, the amplitude of the second term in Eq. (8) should be small. The difference between the aperture-tagging and parallel stochastic perturbative methods is the way in which the second term in Eq. (8) is terminated. In the sequential perturbation technique, the second term is zero as a result of the use of sequential perturbations, for which $\delta u_j \neq 0$ only when $\delta u_k = 0$ for all $k \neq j$. In the multidithering technique, the second term contains only harmonic components at the sum and difference dithering frequencies $\omega_j \pm \omega_k$, which for a properly chosen range of dithering frequencies are filtered out by the synchronous detectors.

In the parallel stochastic perturbative technique, a different idea is used. If the chosen perturbations δu_j are random and statistically independent, the second term in Eq. (8) reduces to zero in expected value, and the gradient estimate $\tilde{J}'_j = \delta J \delta u_j$ is consistent and unbiased.³⁴ Thus the approximation $\tilde{J}'_j \equiv \delta J \delta u_j \approx (\partial J / \partial u_j) (\delta u_j)^2$ can be

used directly as the gradient estimate in gradient descent, either in continuous [Eq. (3)] or discrete [Eq. (5)] form. As a matter of fact, one can easily show, by using arguments similar to relation (4), that gradient descent (5) using the parallel stochastic perturbative gradient estimate $\tilde{J}'_j = \delta J \delta u_j$ produces updates that are guaranteed to decrease (for positive γ) or increase (for negative γ) the metric J at every iteration step, provided that the perturbations and γ are suitably small in amplitude.¹⁹

Mathematically accurate convergence properties are obtained when using the parallel stochastic gradient descent algorithm with uncorrelated random “coin flip” perturbations having identical amplitudes $|\delta u_j| = \pi$ and a Bernoulli probability distribution $\delta u_j = \pm \pi$, where $\text{Pr}(\delta u_j = +\pi) = 0.5$ for all j and iteration numbers.^{17,18} In this case $\tilde{J}'_j = \delta J^{(m)} \pi_j^{(m)}$, and the control voltage update rule (5) reads as

$$u_j^{(m+1)} = u_j^{(m)} - \gamma \delta J^{(m)} \pi_j^{(m)}, \quad j = 1, \dots, N. \quad (9)$$

Note that non-Bernoulli perturbations are also allowed in the SPSA, but one must be careful that the mathematical conditions in Ref. 17 are satisfied. To further improve accuracy of gradient estimation, we can use a two-sided perturbation, performing two measurements of the cost function perturbations δJ^+ and δJ^- corresponding to sequentially applied differential perturbations $+\delta u_j/2$ and $-\delta u_j/2$. It follows that $\delta J = \delta J^+ - \delta J^-$ and $\tilde{J}'_j = \delta J \delta u_j$, which produces a more accurate gradient estimate than Eq. (8).¹⁹

Gradient estimation in the parallel stochastic perturbative method combines the best features of the sequential perturbation and multidithering methods without their drawbacks. Indeed, the signal-to-noise ratio in the parallel stochastic perturbative gradient descent is the same as that in the sequential perturbation technique, that is, N^2 times better than in the multidithering technique. This is because for fixed aperture size, the required photon-flux density is independent of the number of control channels N for both sequential and parallel stochastic perturbation methods and is proportional to N^2 for the multidithering technique. On the other hand, gradient estimation is performed as fast as in the multidithering technique but is a factor N times faster than the sequential perturbation technique requires to obtain all components of the gradient. Perhaps an even more significant advantage of the parallel stochastic method is that it lends itself naturally to efficient hardware implementation using analog VLSI technology.^{21–23}

D. Convergence

A potential source of concern for the parallel stochastic perturbative technique would be reduced accuracy of gradient estimation when compared with that of both the sequential perturbation and ideal (noiseless) multidithering methods. The nonzero second (or third) and higher-order terms in Eq. (8) introduce error into the true gradient estimation. This error impacts the convergence speed of gradient descent, which clearly is slower than pure gradient descent based on the true gradient estimation (assuming that precise gradient estimates were continuously available). However, by taking into account the factor N

time loss for gradient estimation using the sequential perturbation method, the parallel stochastic perturbative gradient descent algorithm provides a significantly faster convergence rate than does the sequential perturbation method. Formal analysis of the convergence rate shows that the parallel stochastic method takes a factor $N^{1/2}$ fewer iterations (perturbed measurements of J) than the sequential perturbation method (and thus a factor $N^{1/2}$ more than pure gradient descent) to reach the same level of performance—owing to the mutually uncorrelated statistics of the parallel perturbations.¹⁹ The convergence speed can be further increased quite significantly by using correlated perturbations that match the statistical properties of the phase distortions³³ or an adaptive SPSA algorithm.²⁰

The convergence speed of the parallel stochastic perturbative gradient descent optimization technique under conditions of dynamical phase distortions and strong intensity scintillations is an important issue that we address in Subsection 2.E.

E. Restrictions for Wave-Front Control

Consider yet another important issue in the use of gradient descent optimization for adaptive optics. In most adaptive optics applications, the system performance metric J is not sensitive to a constant wave-front shift by an arbitrary value v , that is, $J[u(\mathbf{r}, t)] = J[u(\mathbf{r}, t) + v]$ for any v . As a result, wave-front control based on a conventional gradient descent algorithm in the form (1), (3), or (5) may cause an undesirable drift of the aperture-averaged mean phase value during the adaptation process. This drift could exceed physical limits on the operating range of the wave-front correction device and result in a decrease of, and eventually in the loss of, system performance metric sensitivity with respect to applied control signal perturbations.

To prevent this drift, one should modify the original (measured) cost function J to include an additional penalty term $J_a[u] = \frac{1}{2}[\bar{u}(t) - u_0]^2$ accounting for the drift of the mean phase:

$$J_1[u] = J[u] + \eta J_a[u], \quad \bar{u}(t) = s^{-1} \int u(\mathbf{r}, t) d^2\mathbf{r}. \quad (10)$$

Here $\bar{u}(t)$ corresponds to the aperture-averaged (mean) phase, u_0 is the desired mean phase value, and s is the aperture area. The weight coefficient η is positive for cost function minimization and negative otherwise. Similar to Eq. (1), the first variation of the cost functional J_1 includes two terms: $\delta_u J_1 = \delta_u J + \eta \delta_u J_a$. Because the dependence $J_a[u]$ is known, the variation $\delta_u J_a$ can be directly calculated: $\delta_u J_a = -[\bar{u}(t) - u_0]$. Correspondingly, the evolution equation (1) originated from the gradient descent optimization procedure transforms to

$$\tau \frac{\partial u(\mathbf{r}, t)}{\partial t} = -\eta [\bar{u}(t) - u_0] - \delta_u J, \quad (11)$$

where $\delta_u J$ again should be determined based on measured data by using, for example, the parallel stochastic perturbative method. In the case of an adaptive system with N control channels as described above, we have

$$\tau_j \frac{du_j(t)}{dt} = -\eta[\bar{u}_N(t) - u_0] - \gamma J'_j(u_1, \dots, u_N), \quad (12)$$

where

$$\bar{u}_N = N^{-1} \sum_{j=1}^N \alpha_j u_j(t), \quad \alpha_j = (s/N)^{-1} \int S_j(\mathbf{r}) d^2\mathbf{r}.$$

For a piston-type wave-front corrector, $\alpha_j = 1$ and $\bar{u}_N = N^{-1} \sum u_j(t)$. The discrete version of Eq. (12) is conveniently implemented in analog VLSI, as described in Section 3, as

$$u_j^{(m+1)} = u_j^{(m)} - \eta[\bar{u}_N^{(m)} - u_0] - \gamma \delta J^{(m)} \pi_j^{(m)}, \quad j = 1, \dots, N. \quad (13)$$

Similarly, the cost functional may include additional penalty terms, for example a term J_g depending on the gradient of the controlling phase $u(\mathbf{r}, t)$:

$$J_1[u] = J[u] + \eta J_a[u] + \chi J_g[u], \quad (14)$$

where

$$J_g[u] = s^{-1} \int |\nabla u(\mathbf{r}, t)|^2 d^2\mathbf{r}.$$

The purpose of the penalty term J_g is to suppress noise and provide smoothness in the controlling wave-front phase $u(\mathbf{r}, t)$. In the case of a piston-type wave-front corrector, the penalty term suppresses the appearance of 2π phase jumps between neighboring phase-shift elements. The weight coefficients η and χ quantify the relative importance attached to the several penalty terms used in the optimization. The additional term J_g results in the following change in the wave-front control dynamics:

$$\tau \frac{\partial u(\mathbf{r}, t)}{\partial t} = d \nabla^2 u(\mathbf{r}, t) - \eta[\bar{u}(t) - u_0] - \gamma \delta J(t) \delta u(\mathbf{r}, t), \quad (15)$$

where $d = 2\chi$ is an equivalent diffusion coefficient. Control law (15) represents a nonlinear diffusion process with an additional stochastic term $\tilde{J}' = \delta J \delta u$ originating from the parallel perturbative gradient estimate. In the case of a piston-type wave-front corrector, the diffusion term describes coupling between neighboring phase-shift elements similar to those analyzed in Ref. 33. Nonlinear dynamics described by a discrete model of diffusion equation (15) could be implemented by using analog VLSI hardware as well. Translinear metal-oxide semiconductor circuits in the weak inversion region of operation implement diffusor-type elements at the pixel level, with a characteristic speed on the order of microseconds, while dissipating just microwatts of power.^{35,36}

F. Wave-Front Sensor Information Fusion

Wave-front control based on system performance metric optimization does not assume or make use of direct information about wave-front phase distortions. On the other hand, when additional wave-front information is available, it can be effectively used in the adaptation process to improve the speed and the accuracy of convergence. The

idea of incorporating different requirements (metrics) for wave-front control under the same umbrella, in the form of a generalized system performance metric as discussed above, can be extended as a tool for fusing information obtained from different adaptive optics sensors. Despite the variety of wave-front sensor types, each wave-front sensor performs a transformation of the wave-front phase distortion $\phi(\mathbf{r}, t)$, or more commonly the phase error $\delta(\mathbf{r}, t) = \phi(\mathbf{r}, t) + u(\mathbf{r}, t)$, into an intensity modulation $I_\delta(\mathbf{r}, t) = F\{\delta(\mathbf{r}, t)\}$ measured by the sensor's photoreceiver. For the sake of simplicity, we assume that the intensity modulation $I_\delta(\mathbf{r}, t)$ is measured in the phase plane or its conjugate plane and that both $I_\delta(\mathbf{r}, t)$ and $\delta(\mathbf{r}, t)$ are geometrically matched. The phase-to-intensity transformation operator F is dependent on the wave-front sensor used. Typically, in adaptive optics, we know what the ideal intensity distribution $I_\delta^0(\mathbf{r})$ should be in the absence of phase distortions. The information from the wave-front sensor can be incorporated into gradient descent control by introducing an additional term J_s into the chosen system performance metric:

$$J_1[u] = J[u] + \mu J_s[u], \quad (16)$$

where

$$J_s[u] = \frac{1}{2} \int [I_\delta(\mathbf{r}, t) - I_\delta^0(\mathbf{r}, t)]^2 d^2\mathbf{r}.$$

It can be shown that parallel stochastic perturbative gradient descent optimization applied to the new metric results in the following wave-front control dynamics:

$$\tau \frac{\partial u(\mathbf{r}, t)}{\partial t} = -\mu [I_\delta(\mathbf{r}, t) - I_\delta^0(\mathbf{r})] \times \delta I_\delta(\mathbf{r}, t) \delta u(\mathbf{r}, t) - \gamma \delta J(t) \delta u(\mathbf{r}, t), \quad (17)$$

where $\delta I_\delta(\mathbf{r}, t)$ is the measured perturbation of the intensity modulation that is due to the wave-front perturbation $\delta u(\mathbf{r}, t)$. The coefficient μ defines preferences in utilizing the information obtained from both the system performance metric and the wave-front sensors. Note that in contrast to the system performance metric perturbation $\delta J(t)$, the wave-front sensor provides a spatially distributed response $\delta I_\delta(\mathbf{r}, t)$. The product $\delta I_\delta(\mathbf{r}, t) \delta u(\mathbf{r}, t)$ in Eq. (17) describes local correlation between the wave-front perturbation and the corresponding intensity change registered by the wave-front sensor. This correlation provides a local feedback between the wave-front sensor and the wave-front corrector. As has been recently demonstrated by using a diffractive-feedback-type adaptive system, the presence of local feedback can significantly accelerate adaptation convergence speed.³⁷

The requirements for wave-front smoothness and the absence of a mean phase drift mentioned above are easily incorporated into a modified stochastic gradient descent rule (17):

$$\begin{aligned} \tau \frac{\partial u(\mathbf{r}, t)}{\partial t} = & d \nabla^2 u(\mathbf{r}, t) - \eta [\bar{u}(t) - u_0] - \mu [I_\delta(\mathbf{r}, t) \\ & - I_\delta^0(\mathbf{r})] \delta I_\delta(\mathbf{r}, t) \delta u(\mathbf{r}, t) - \gamma \delta J(t) \delta u(\mathbf{r}, t), \end{aligned} \quad (18)$$

where d , μ , and η are weight coefficients.

Adaptive wave-front control based on the various forms of gradient descent optimization algorithms studied above gives rise to different control system architectures. How effectively the wave-front control methods perform in actual adaptive optical systems needs to be validated through direct experimental study.

3. ADAPTIVE WAVE-FRONT CONTROL USING ANALOG VERY-LARGE-SCALE INTEGRATION PARALLEL STOCHASTIC OPTIMIZATION: EXPERIMENTAL RESULTS

In this section we describe experimental results obtained in two different adaptive optical systems using a custom-designed mixed-mode analog–digital VLSI controller (called the AdOpt chip), which implements the discrete-time parallel stochastic perturbative gradient descent optimization algorithm (13) with two-sided perturbation and aperture-averaged phase control. Through the experiments the following problems were addressed:

1. Adaptation process convergence speed for adaptive systems having different numbers of control channels and different wave-front corrector types.
2. Impact of local extrema on adaptive system performance.
3. Adaptive correction of dynamical phase distortions.
4. Adaptation under the condition of strong intensity scintillations.

A. Very-Large-Scale Integration Stochastic Gradient Descent Controller

A single AdOpt chip provides for storage, parallel stochastic perturbation, and corresponding gradient descent adaptation of 19 parallel output control signals u_j , $j = 1, \dots, 19$. The chip simultaneously generates the parallel perturbations, performs updates according to the rule (13) from the externally supplied instantaneous system performance metric J , and calculates the mean value $\bar{u}_N = N^{-1} \sum u_j(t)$ used in the update rule. The AdOpt chip design provides for a scalable and expandable control system architecture in which the metric J can be arbitrarily constructed.

In the experiments described below, two adaptive systems using AdOpt chips were implemented: one with 39 (three AdOpt chips) control signals and another with 127 (seven chips) control signals. The VLSI controller consists of the following main components:

1. A pseudorandom vector sequence generator with independent elements $\pi_j^{(m)}$, $j = 1, \dots, 19$, having identical amplitudes and a Bernoulli probability distribution.
2. A circuit actuating the perturbations in parallel onto the control signals.

3. A circuit performing parallel stochastic perturbative gradient descent updates to the control signals in accordance with the control algorithm (13).

4. An analog memory circuit, which stores the control signal values u_j between sequential updates.

5. An optional output driver designated to a specific (LC-based) wave-front corrector.

The design of the AdOpt chip was optimized to control a multielectrode nematic LC phase modulator as well as microelectromechanical system multielectrode mirrors. The VLSI system supplies the output control voltages with an amplitude in the range ± 5 V. The circuit serves parallel updates at adaptation rates up to 200 kHz, well in excess of the bandwidth of typical wave-front correctors currently available. The VLSI system block diagram and the chip micrograph are shown in Fig. 1. A detailed description of the circuits, the operation, and the (measured) technical specifications of the AdOpt VLSI controller are given in Refs. 38 and 39.

B. Control System Architecture

Parameters of the control algorithm (13) implemented on the AdOpt chip were supplied by using a personal computer, interfacing with the chips through a circuit board with integrated microcontroller. These external parameters are perturbation amplitude π , sign and magnitude of the update coefficient γ , and reference value u_0 used for stabilization of the control signal mean values \bar{u}_N . The personal computer also provided timing signals and supplied the sign and the magnitude of the measured cost function perturbation δJ . Computer control of the VLSI system parameters not only provided flexibility for algorithm parameter optimization but also allowed implementation of in-the-loop reinforcement learning for (on-the-fly) parameter value adaptation. An example of this type of control system architecture is given in Ref. 40.

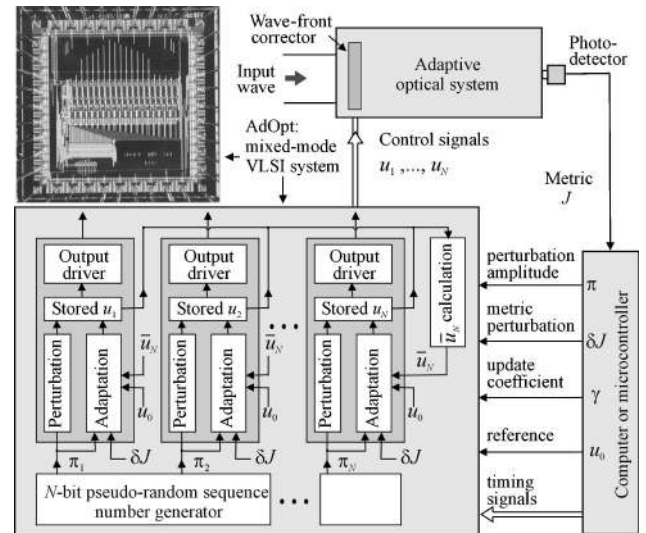


Fig. 1. Wave-front control system architecture with block diagram of the multichannel mixed-mode VLSI stochastic gradient descent controller (AdOpt system) and (top left) micrograph of the 19-channel system, a 2.2-mm \times 2.25-mm chip fabricated by using 1.2- μ m complementary metal-oxide semiconductor technology.

C. Wave-Front Correctors

The wave-front correctors that we used in the experiments were a multielectrode nematic LC phase spatial light modulator (HEX127) from Meadowlark Optics Inc.⁴¹ and a 37-electrode silicon micromachined deformable mirror from OKO Technologies,^{42,43} shown in Fig. 2. The HEX127 phase modulator has 127 hexagonal LC cells. Each cell is 1.15 mm in diameter with 36- μm spacing. The HEX127 aperture size is 12 mm in diameter. The VLSI system (seven AdOpt chips) simultaneously (in parallel) supplied 127 controlling signals to the HEX127 phase modulator. The applied voltages ($0 < |u_j| < 5$ V) provided piston-type wave-front control with individual phase-shift magnitudes covering approximately 2π rad. To prevent mean phase drift, we set the reference voltage value u_0 in the middle of the LC modulator operational range ($u_0 = 2$ V). The adaptation iteration rate, defined as the number of parallel updates in the control signals per second, was limited to 15 iterations per second because of the relatively large (~ 100 -ms) response time of the LC phase modulator. This iteration rate includes both two-sided perturbation and calculation of the control signal mean values \bar{u}_N .

An adaptation iteration rate more than 100 times faster (up to 1600 iterations per second) was achieved by using the same controller (now with three AdOpt chips) with the OKO micromachined deformable mirror. The OKO mirror consists of a chip with a silicon nitride membrane coated with aluminum, forming a mirror attached to the substrate. The membrane shape is electrostatically controlled by voltages applied to the 37 hexagonal control electrodes. The VLSI system outputs were connected to the OKO mirror through high-voltage amplifiers. Each of the 37 amplifiers provided output voltages in the range 0–200 V. The OKO mirror aperture size is 15 mm, of which only the central part (12 mm) is used. The maximum deflection amplitude of the mirror surface is approximately 6 μm for the center electrode and 1.5 μm for the electrodes located at the edge of the mirror. The OKO mirror had a rather strong (~ 4 - μm -magnitude) initial aberration in the form of astigmatism and defocus.

D. Experimental Setups

To characterize the performance of adaptive wave-front correction using the VLSI parallel stochastic optimization technique, we used two simple adaptive focusing systems, shown in Fig. 2. In both systems we used a linearly polarized laser beam from an argon laser ($\lambda = 0.514$ μm) expanded to a diameter of 12 mm. The laser beam power inside a pinhole placed in the focal plane of lens L_1 was measured by a photodetector and used as the system performance (beam quality) metric J . In the system shown in Fig. 2(a), wave-front correction was performed by using the HEX127 phase modulator. The 37-electrode OKO mirror and an additional mirror for wave-front tilt control were used in the system shown in Fig. 2(b). Images of the laser beam in the focal plane of lens L_1 (plane of the pinhole) were registered by a CCD camera (CCD₁ in Fig. 2). To observe the phase pattern formed by the LC phase modulator, we set the input beam polarization at an angle of $\pi/4$ with respect to the LC phase modulator optical axis and orthogonal to the optical axis of the polarizer P_2 .




Fig. 2. Schematic for adaptive laser beam focusing systems used in the experiments: (a) adaptive system with LC multielectrode spatial phase modulator (HEX127 phase SLM) and (b) system with 37-channel micromachined OKO mirror and beam steering mirror (tilt control). Pictured are the geometry of the HEX127 spatial light modulator electrodes, the OKO mirror phase profile with equal voltages (50 V) applied to the mirror's two electrodes (peak value phase deviation of 2π rad), and a photograph of the OKO mirror. Focal lengths corresponding to lenses L_1 and L_2 are 14 in. (35.6 cm).

The lens L_2 imaged the plane of the HEX127 onto the camera CCD₂. The optical axis of the polarizer P_1 was parallel to the LC optical axes, providing registration of only the phase-modulated component of the optical wave inside the pinhole.

E. Compensation of Self-Induced Aberrations

The compensation of random static phase distortions that are introduced by an adaptive system wave-front corrector itself (self-induced aberrations) is a commonly used technique for the evaluation of adaptive optical system performance: operational range, adaptation process stability, and convergence. For the case of self-induced phase distortions, we know that ideal compensation can potentially be achieved. Thus we can directly compare system performance metrics corresponding to undistorted J_{opt} , randomly distorted J_0 , and adaptively compensated $J(m)$ wave front (at the m th iteration). The typical way to create self-induced phase distortions is to initialize the wave-front corrector electrodes with random control voltages. In our case, when the wave-front corrector (HEX127 or OKO mirror) was connected to the VLSI system, this technique could not be applied directly. To create highly distorted wave fronts, we used a different method. Computer control of the sign of the AdOpt chip update coefficient γ allowed us to execute repeating sequences (trials) of cost function J minimization ($\gamma < 0$) and maximization ($\gamma > 0$). Each adaptation stage in-

cluded 512 iteration steps. During the minimization stage, the adaptive system created phase distortions, resulting in a highly distorted laser beam intensity in the plane of the photodetector. During the maximization stage, the adaptive system compensated these self-induced phase distortions so that the laser beam power was concentrated inside the pinhole. Each maximization–minimization trial was repeated 100 times to obtain averaged adaptation evolution curves.

Consider first the results obtained by using the adaptive system with the HEX127 LC phase modulator.

The normalized dependence $\langle J(m) \rangle$, $m = 1, \dots, 1024$ (averaged adaptation evolution curve), is shown in Fig. 3(a). The adaptive system with 127 control channels was able to “successfully” introduce phase distortions, leading to severe degradation of the beam quality metric (minimization stage), and then correct these self-induced aberrations. The correction dynamical range J_{\max}/J_{\min} [$J_{\max} = \max\langle J(m) \rangle$ and $J_{\min} = \min\langle J(m) \rangle$] ranged between 10 and 15. The normalized standard deviation of the beam quality metric, $\sigma_J = \langle [J(m) - \langle J \rangle]^2 \rangle^{1/2} / \langle J \rangle$, shown in Fig. 3(b), characterizes fluctuations in the beam quality metric level. The standard deviation σ_J was on the order of 15%–20% for beam quality metric maximization and 30%–35% for metric minimization. The results presented in Fig. 3 were obtained when the perturbation amplitude π resulted in the performance metric perturbation δJ having a normalized standard deviation $\sigma_{\delta J}$ approximately 3% for metric maximization and 7% for metric minimization.

Two observations can be made from the analysis of these adaptation trials: (1) The evolution curve in Fig.




Fig. 3. Experimental results of phase distortion compensation in the adaptive system with LC phase modulator: (a) averaged laser beam quality metric and (b) standard deviations for metric J and metric perturbation δJ obtained from averaging over 100 trials of beam quality metric maximization ($m \leq 512$) and minimization ($m > 512$). The time scale is 30 s per 512 iteration steps, indicated by the dot in (a).

3(a) shows the existence of two characteristic phases of the adaptation process: a relatively fast convergence during the first 100–150 iterations, followed by a decrease in the convergence rate. Convergence occurred approximately 1.5 times faster for metric minimization than for metric maximization. (2) We did not observe correction to a uniform wave front. Instead, during each trial, adaptation resulted in a different “optimal” phase pattern. This suggests that there are a number of wave-front phase patterns that correspond to rather high and approximately equal quasi-optimal beam quality metric values (i.e., local maxima and minima of the cost function). Examples of these quasi-optimal phase patterns for performance metric maximization ($m = 512$) and minimization ($m = 1024$), as well as the corresponding focal plane images of laser beam intensity distributions, are shown in Fig. 4.

We expected that the optimal beam quality metric value achieved by the adaptive system (global maximum) should correspond to a uniform phase, which can be obtained by applying the same reference voltage to all HEX127 SLM electrodes. The measured value J_{ref} was compared with the maximum average metric value J_{\max} . Surprisingly, the obtained value J_{\max} was a factor 2.4 times larger than the expected optimum J_{ref} . We also compared the achieved beam quality metric with the quality metric J_0 corresponding to zero voltages applied to all 127 electrodes of the phase modulator. For this case we obtained $J_{\max}/J_0 = 2.05$, i.e., the adaptive system found a “better” wave-front phase than the one that we expected. This result can be explained if we assume the presence of an initial phase aberration in the optical system. The adaptive system not only removed the self-induced aberrations but also (at least partially) compensated the aberrations that are due to the system’s optics.

Quite different adaptation dynamics were observed for the adaptive compensation of self-induced wave-front distortions in the system with micromachined deformable mirror shown in Fig. 2(b). The OKO mirror provides continuous deformation of the wave-front shape. This wave-front phase deformation not only may result in focal plane laser beam intensity spatial redistribution but also produces noticeable shifts in the center of mass of the laser beam. With the same small pinhole (25 μm) as that used in the experiments with the HEX127 phase modulator, we observed highly unstable adaptation dynamics. Increasing the pinhole size resulted in a loss of beam quality metric sensitivity and significantly slowed the adaptation process. Two modifications were made to achieve stable and efficient wave-front correction:

1. Instead of a pinhole, we used a computer-generated gray-scale mask having a relatively small (100- μm -diameter) transparent central part and a linearly decreasing transparency coefficient inside a circle of 1.0-mm diameter. The laser beam power measured behind the mask was used as the beam quality metric.

2. We included additional wave-front x – y tilts control by using a beam steering mirror. The averaged beam quality metric evolution curves $\langle J(m) \rangle / \max\langle J(m) \rangle$ for the maximization–minimization trials discussed above are shown in Fig. 5 for different adaptation scenarios. Ex-




Fig. 4. Phase (left column) and focal plane intensity (right column) patterns obtained in the adaptive system with LC phase modulator during two subsequent maximization–minimization trials. (a), (b) Pattern corresponding to metric maximization ($m = 512$) for the first trial and (c), (d) that for the second trial. (e)–(h) Phase and intensity patterns corresponding to metric minimization: (e), (f) the first trial, and (g), (h) the second trial.

clusive adaptive control of wave-front tilt resulted in fast convergence, but the achieved beam quality metric value was relatively low (curve 1 in Fig. 5) because control of wave-front tilts alone cannot correct for the initial aberrations introduced by the OKO mirror itself. At the same time, wave-front tilt control resulted in a relatively small level of fluctuations in the beam quality metric [see the standard deviation curve $\sigma_J(m)$ in Fig. 5(b)].

On the other hand, exclusive adaptive control of the microelectromechanical system mirror (curve 2) resulted in compensation of both the self-induced and initial phase

distortions and achieved a beam quality metric value more than 2.5 times greater. The drawback of wave-front control using only the micromachined mirror is the low convergence rate and the significant level of beam quality metric fluctuation [see curve 2 in Fig. 5(b)]. In contrast, when both the OKO and tilt control mirrors were used, the adaptation performance improved dramatically (curve 3). The adaptation process typically converged in approximately 50–60 iterations (~ 40 ms), with the lowest fluctuations in the performance metric. Thus incorporation of wave-front tilt control is an important factor for acceleration and stabilization of the adaptation process in adaptive systems with controllable mirrors having a continuously deformable surface.

F. Convergence Rate Issues

Dependence of the convergence rate on the number of control channels N was investigated by using an experimental setup with the LC phase modulator [Fig. 2(a)]. A diaphragm was placed in front of the HEX127 phase modulator to decrease the input beam size in such a way that only the central part of the phase modulator, containing a certain number N of the LC cells, was open. For each aperture size, the input beam power was adjusted to have the same optimal beam quality metric value as that corresponding to an undistorted input wave




Fig. 5. Self-induced phase distortion compensation in the adaptive system with beam steering and micromachined mirrors: evolution curves for (a) averaged beam quality metric and (b) standard deviation. Metric maximization corresponds to $0 < m \leq 512$, and metric minimization corresponds to $m > 512$. The time scale is 0.34 s per 512 iteration steps, indicated by the dot in (a). The evolution curves 1–3 correspond to 1, wave-front tilt control only; 2, OKO mirror control only; and 3, control of both the beam steering and micromachined mirrors. The photographs show focal plane intensity distributions at the end of (top) maximization and (bottom) minimization trials.




Fig. 6. Normalized adaptation evolution curves for beam quality metric maximization–minimization (averaged over 100 trials) in the system with LC phase modulator for different control channel numbers N . The photographs on the right show typical phase patterns for $N = 7, 19$, and 127 .

front. The aperture diameters corresponded to $N = 7, 19, 37, 61, 91$, and 127 (full aperture size) LC cells. For each aperture size, we performed the experiments for adaptive compensation of self-induced phase distortions described in Subsection 3.E. Depending on the number of the LC cells inside the open aperture (number of the control channels N), we obtained averaged evolution curves for beam quality metric maximization and minimization, as shown in Fig. 6. The convergence rate was notably higher for $N = 7$ and 19 than for the fully opened aperture ($N = 127$). The interesting result is that between $N = 37$ and 127 , the convergence rate was approximately constant, apparently independent of the number of control channels N .

The experiments demonstrated that in the case of a piston-type wave-front corrector approximately 80–100 iterations of parallel stochastic gradient descent optimization are required to achieve 80% of the maximum level of the beam quality metric. The adaptation rate was higher for the system with deformable mirror and tilts control (~ 50 – 60 iterations).

Two circumstances should be taken into account when predicting actual adaptive system convergence rates by using self-induced phase distortion compensation experiments. First, in the experiments with self-induced phase distortions, the convergence rate was estimated based on compensation of phase aberrations that were independent from trial to trial—temporally uncorrelated phase aberrations. In some sense, temporally uncorrelated phase distortions as obtained in the alternating maximization–minimization trials constitute the worst possible scenario for an adaptive system—each maximization trial starts from specially prepared “bad” conditions. Under actual conditions an adaptive system compensates continuously changing phase aberrations. The temporal correlation of phase distortions can result in some improvement of the adaptation process convergence rate. Second, “natural” wave-front aberrations have components that cannot be compensated by a chosen wave-front corrector. The presence of these components can still impact (typically negatively) the adaptive process convergence rate.

4. DYNAMICAL WAVE-FRONT DISTORTION CORRECTION

The high adaptation rates (>1500 iterations per second) achieved in the VLSI-based adaptive system with both micromachined and beam steering mirrors allowed us to perform experiments with dynamical phase aberrations originating from an electric heater and fan-induced turbulent flows. The goal of these experiments was to analyze the efficiency of the parallel stochastic perturbative gradient descent technique in the presence of severe intensity scintillations—conditions under which traditional adaptive optics based on wave-front measurements and conjugation do not perform effectively. We considered two adaptation scenarios typical for ground-to-ground laser communication and laser beam focusing systems, here referred to as adaptive receiver and adaptive transmitter systems, respectively.

A. Adaptive Receiver and Transmitter with Heater/Fan-Induced Turbulent Flow

In the adaptive receiver configuration shown in Fig. 7(a), the electric heater and fan were placed a distance of 3 m in front of the beam steering mirror. The 1500-W “base-board” heater consisted of a 1-m-long heating element oriented parallel to the input laser beam and located 20 cm below the beam. A metal grating with cell size 4 mm was placed on top of the heating element to break the air flow into smaller-sized streams and create refractive-index fluctuations that were small in comparison with the laser beam size (15 mm). The air flow created by the fan contributed to the heated air stream mixture, resulting in a highly turbulent regime. The mutual positions of laser beam, heater, and fan are shown in Fig. 7(a) (bottom right corner).

In the adaptive transmitter configuration shown in Fig. 7(b), two heaters with equivalent power of 1500 W were placed in the middle of the propagation path (25-m length). The adaptive system controlled the wave-front phase of the outgoing beam by using the beam quality data measured at the end of the propagation path.

B. Adaptive Receiver: Heater-Induced Turbulent Flow

The experimental data for the adaptive receiver system were obtained through 100 trials of beam quality metric maximization. Each trial included 500 iteration steps performed during approximately 0.3 s. After each trial the beam quality values $J(m)$, $m = 1, \dots, 500$, were stored. During data storage (~ 0.1 s), the control voltages at both the OKO and steering mirrors were fixed. The results of 100 trials were averaged. Averaged evolution curves $\langle J(m) \rangle$ as well as the normalized standard deviation of the beam quality metric fluctuations, $\sigma_J = \langle [J(m) - \langle J \rangle]^2 \rangle^{1/2} / \langle J \rangle$, are shown in Fig. 8 for the case of turbulent flow created by only the heater [Figs. 8(a) and 8(b)] and by both the heater and the fan [Figs. 8(c) and 8(d)]. Consider first the results presented in Figs. 8(a) and 8(b) (heater only). Without adaptation the level of beam quality metric fluctuations was approximately 80% [see curve 1 in Fig. 8(b)]. Adaptation using only the steering mirror (control of wave-front tilt only) resulted in a 1.2-fold increase of the average beam quality

metric and a 1.25-fold reduction in the metric fluctuation level (σ_J). Adaptive wave-front control using only the OKO mirror led to similar results [see curves 3 in Figs.




Fig. 7. Schematics for adaptive systems used in the experiments with heater/fan-induced turbulence: (a) adaptive receiver and (b) adaptive transmitter configurations. The focal length corresponding to lens L is 70 in. (177.8 cm), and the pinhole size is 2 mm.

8(a) and 8(b)]. The system performance improvement was more substantial and increased 1.7 times for $\langle J \rangle$ and decreased 1.9 times for σ_J when both the micromachined (OKO) and tilt control mirrors were operated at the same time (curves 4). As seen from the initial portion of the adaptation evolution curves, the characteristic convergence rate was on the order of 20–30 iterations if both adaptive mirrors were used and 50–70 iterations for tilt control only. The convergence rate can also be estimated from an analysis of the small dips present in the evolution curves in Figs. 8(a) and 8(c). These dips occur periodically after every 128 iteration steps and resulted from the 10-ms delay used to reset the built-in random sequence generator of the VLSI controller. The dips are caused by these delays in adaptive control and not by drifts in the control voltages, since these were fixed during the reset of the random sequence generator. Because the time delay is short, the phase distortions had relatively little time to drift and the adaptive correction converged within only a few iterations. In the subsequent experiments described below, we activated the random sequence generator reset only once, so these dips did not occur.

Histograms for different adaptation scenarios are shown in Fig. 9. The histogram represents the number of adaptive system states M_{st} as a function of the beam quality metric value J . The histograms in Fig. 9 were calculated by using data from all 100 adaptation trials. Normalized histograms present the probability distributions for the achieved beam quality metric values during the adaptation trials. Another interpretation for the histograms is the distribution of the relative time intervals spent by the adaptive system near a particular metric value J . The data presented by the histograms in Fig. 9 demonstrate significant improvement in system performance through adaptation by using both wave-front cor-

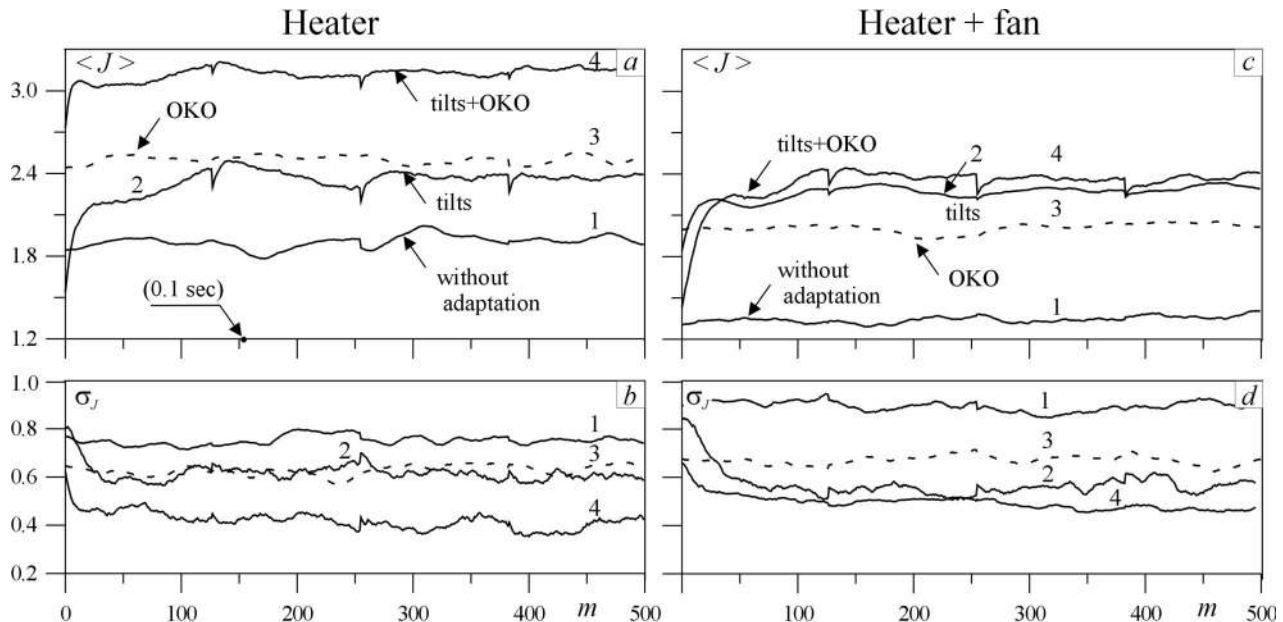


Fig. 8. Experimental results for adaptive receiver system configuration with laser beam propagation through the turbulence created by (a), (b) heater and (c), (d) heater and fan: (a), (c) averaged adaptation evolution curves for beam quality metric and (b), (d) corresponding curves for standard deviation of the beam quality metric fluctuations. The time scale is 0.1 s per 155 iteration steps, indicated by the black dot in (a). The evolution curves 1–4 correspond to 1, disabled feedback control (no adaptation); 2, wave-front tilt control only (tilts); 3, OKO mirror control only (OKO); and 4, control of both the beam steering and micromachined mirrors (tilts+OKO).

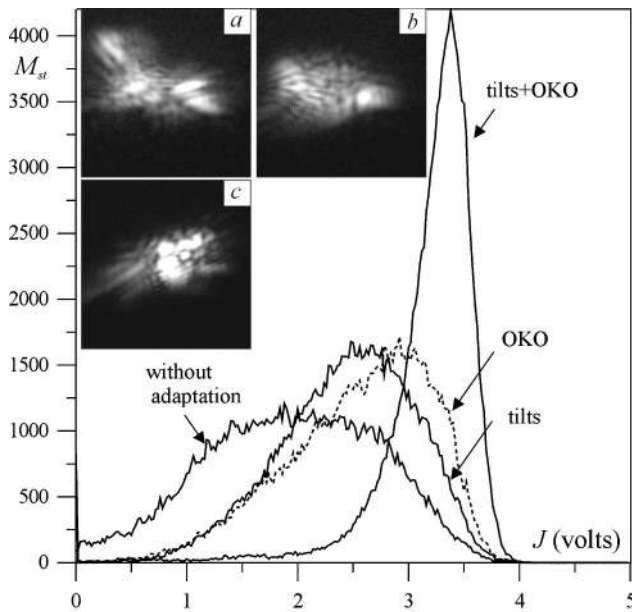


Fig. 9. Histograms for the beam quality metric optimization process in the adaptive receiver with dynamical phase distortions created by an electric heater. Curve labels are the same as those in Fig. 8. The photographs correspond to focal plane intensity distributions in the system (a), (b) without adaptation and (c) with adaptation.

rectors. Typical snapshots of focal plane intensity distributions in the system without adaptation and with adaptation are shown in Figs. 9(a)–9(c). Adaptive wave-front control did not provide complete compensation of phase distortions—the intensity distribution in Fig. 9(c) is still highly aberrated—but did result in a noticeable decrease of the laser beam width and stabilization of position inside the aperture of the focal plane receiver (compare photos in Fig. 9). (In the experiments with dynamical phase distortions, we used a 2-mm pinhole.)

C. Heater and Fan-Induced Turbulence: Fading Effect

The air flow created by the fan significantly increased the beam intensity scintillation and hence the metric fluctuation level (compare the curves in Fig. 8). The standard deviation σ_J increased by up to 100%. The histogram for the system without adaptation displays strong intensity fading characterized by the appearance of a large number of zero metric values (5309 zeros), where the beam exits the pinhole area entirely. Turbulence-induced signal fading is an important factor limiting the performance of laser communication systems. Adaptive wave-front control using the OKO mirror improved the average beam quality metric value by a factor 1.5 but was rather unstable and only slightly decreased signal fading (4406 zeros). Using only wave-front tilt control resulted in significant stability improvement and decreased fading effect (1293 zeros). Quality metric fading was almost completely eliminated (42 zeros) when adaptation was performed by using both beam steering and micromachined mirrors (see the histograms in Fig. 10). Laser beam intensity snapshots averaged over 300 realizations for the system without adaptation and with adaptation are shown in Figs. 10(a) and 10(b), respectively. The beam intensity snapshots were taken in the plane corre-

sponding to the position of lens L in Fig. 7(a). The strong asymmetry of the laser beam in the horizontal direction seen in Fig. 10(a) corresponded to the direction of the air flow created by the fan. Adaptive wave-front control resulted in compensation of the beam asymmetry, as seen in Fig. 10(b).

D. Adaptive Transmitter: Adaptation Spectral Band

In the adaptive transmitter experiments [Fig. 7(b)], each adaptation trial included 4096 iterations that corresponded to a time interval of 2.7 s. Just as in the adaptive receiver experiment, 100 trials were averaged. Two heaters placed in the middle of a 25-m-long path resulted in strong beam quality degradation. The observed beam quality fluctuation level was on the order of 70%–80%. We did not use a fan in this experiment, so that we could emphasize beam quality degradation that is mostly due to small-scale refractive-index fluctuations. The experimental results are shown in Fig. 11. Adaptation resulted in a 1.4-fold increase of the averaged beam quality metric value and a 1.5-fold reduction in standard deviation for beam quality metric fluctuations. The histograms in Fig. 11 show an interesting feature of the system dynamics. Wave-front control using only the OKO mirror demonstrated a relatively high level of the averaged beam quality metric, but at the same time adaptation resulted in an increase of the metric fluctuation level: from $\sigma_J = 0.75$ (without adaptation) to $\sigma_J = 0.9$ with OKO mirror control. This fluctuation level increase and corresponding widening of the histogram indicate instability in the adaptation process as characterized by an occasional loss of control. As seen from the data presented in Fig. 11, add-

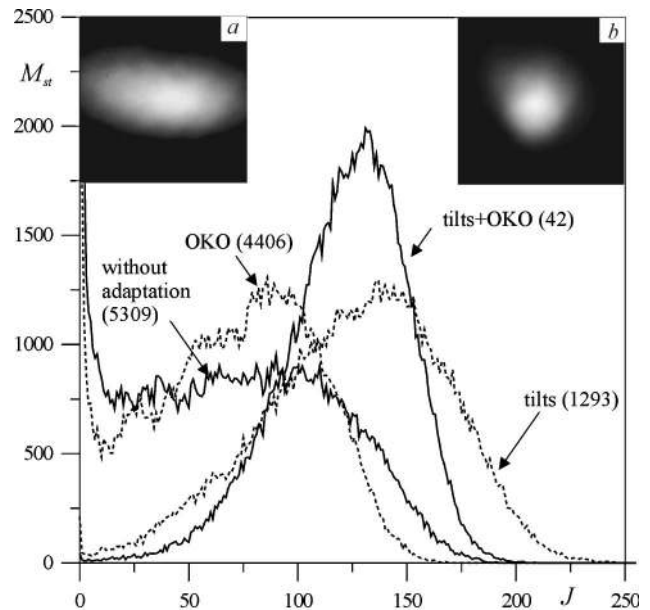


Fig. 10. Histograms for beam quality metric optimization in an adaptive-receiver-type system with dynamical phase distortions created by both heater and fan. The numbers in parentheses correspond to the number of the adaptive system states M_{st} with zero beam quality metric value. The photographs correspond to averaged focal plane intensity distributions in the system (a) without adaptation and (b) with adaptation.

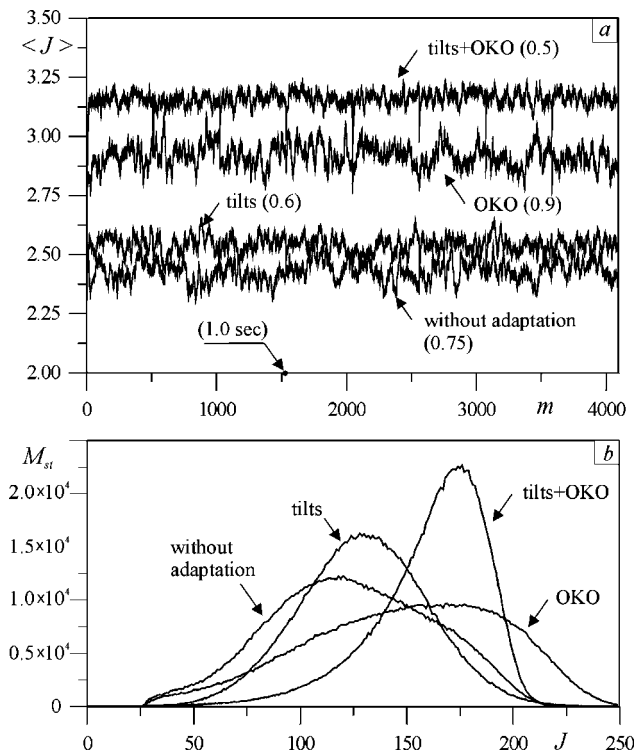


Fig. 11. Experimental results for adaptive transmitter configuration with laser beam propagation through the turbulence created by two heaters: (a) averaged adaptation evolution curves for beam quality metric and (b) corresponding histograms. The numbers in parentheses correspond to the standard deviation σ_J for beam quality metric fluctuations.

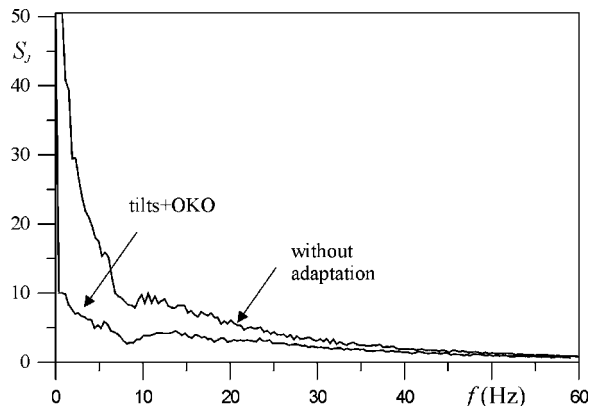


Fig. 12. Beam quality metric temporal spectral energy density averaged over 300 adaptation trials for the adaptive receiver system with feedback control off (without adaptation) and with control using both the beam steering and micromachined mirrors (tilts+OKO).

ing wave-front tilt control significantly improved system performance: large $\langle J \rangle$, smaller σ_J , and narrower histogram.

The beam quality metric temporal spectral energy density S_f is shown in Fig. 12. The dc components in Fig. 12 are clipped to emphasize the changes in the temporal spectral density occurring over the entire frequency range. As seen from the presented data, turbulence created by the heaters resulted in beam quality metric fluctuations in the spectral range of approximately 50–60 Hz, with the most significant fluctuation components occur-

ring near 20 Hz. Adaptive correction using both the steering and OKO mirrors increased the dc component by a factor 1.9 and considerably narrowed the spectrum in the range up to 60 Hz.

5. CONCLUSION

Our results show that adaptive wave-front control using parallel stochastic optimization (simultaneous perturbation stochastic approximation) provides an attractive and practical alternative to traditional adaptive optics techniques based on wave-front measurements and conjugation. The general form of the parallel stochastic optimization technique gives flexibility in designing new control algorithms. As shown theoretically, this flexibility can be used to incorporate various requirements for wave-front phase correction (such as wave-front smoothness and the absence of mean phase drift) and to fuse any information that might be available from different wave-front sensor types. We experimentally demonstrated adaptive optical systems with real-time analog very large-scale integration wave-front control that included a liquid-crystal phase modulator or micromachined deformable and beam steering mirrors that successfully operated under conditions of strong turbulence with large intensity scintillations. Because of the general nature of the stochastic optimization algorithm, the adaptive optics techniques demonstrated here can be applied to a wide variety of optical systems, such as laser communication, beam focusing systems, and imaging systems, that operate in the presence of atmospheric turbulence.

ACKNOWLEDGMENTS

The authors thank James Spall for helpful discussions on stochastic optimization methods, Fred Davidson for discussions about the possibilities to apply VLSI-based adaptive optics techniques to optical communications, and Jennifer Ricklin for technical and editorial comments. This work was performed at the U.S. Army Research Laboratory's Intelligent Optics Laboratory under the Director's Research Initiative program. The AdOpt chip was designed at Johns Hopkins University under U.S. Army Research Laboratory contract DAAH04-96-C-0086 and fabricated through the MOSIS foundry service.

Address correspondence to Mikhail A. Vorontsov, Army Research Laboratory, AMSRL-CI-C, Adelphi, Maryland 20783, or by phone, 301-394-0214; fax, 301-394-0225; or e-mail, vorontsov@iol.arl.mil.

REFERENCES AND NOTES

1. J. H. Shapiro, "Reciprocity of the turbulent atmosphere," *J. Opt. Soc. Am.* **61**, 492–495 (1971).
2. J. W. Hardy, "Active optics: a new technology for the control of light," *Proc. IEEE* **66**, 651–697 (1978).
3. In the nonlinear-optics-based or dynamical-holography-based phase conjugation systems, we have true wave-front conjugation (both phase conjugation and amplitude correction); see, for example, B. Y. Zeldovich, N. V. Pilipetsky, and V. V. Shkunov, *Principles of Phase Conjugation*, Vol. 42 of Springer Series in Optical Sciences (Springer-Verlag, Berlin, 1985).

4. C. A. Primmerman, T. R. Price, R. A. Humphreys, B. G. Zollars, H. T. Barclay, and J. Herrmann, "Atmospheric-compensation experiments in strong-scintillation conditions," *Appl. Opt.* **34**, 2081–2088 (1995).
5. B. M. Levine, A. Wirth, H. DaSilva, F. M. Landers, S. Kahalas, T. L. Bruno, P. R. Barbier, D. W. Rush, P. Polak-Dingels, G. L. Burdge, and D. P. Looze, "Active compensation for horizontal line-of-sight turbulence over near-ground paths," in *Propagation and Imaging through the Atmosphere II*, L. R. Bissonnette, ed., Proc. SPIE **3233**, 221–232 (1998).
6. N. B. Baranova, A. V. Mamaev, N. F. Pilipetsky, V. V. Shkunov, and B. Ya. Zel'dovich, "Wave-front dislocations: topological limitations for adaptive systems with phase conjugation," *J. Opt. Soc. Am.* **73**, 525–528 (1983).
7. D. L. Fried, "Branch point problem in adaptive optics," *J. Opt. Soc. Am. A* **15**, 2759–2768 (1998).
8. R. A. Muller and A. Buffington, "Real-time correction of atmospherically degraded telescope images through image sharpening," *J. Opt. Soc. Am.* **64**, 1200–1210 (1974).
9. M. A. Vorontsov, G. W. Carhart, D. V. Pruidze, J. C. Ricklin, and D. G. Voelz, "Adaptive imaging system for phase-distorted extended source/multiple distance objects," *Appl. Opt.* **36**, 3319–3328 (1997).
10. M. C. Roggemann and B. M. Welsh, *Imaging through Turbulence* (CRC Press, Boca Raton, Fla., 1996).
11. G. W. Carhart and M. A. Vorontsov, "Synthetic imaging: non-adaptive anisoplanatic image correction in atmospheric turbulence," *Opt. Lett.* **23**, 745–747 (1998).
12. R. K. Tyson, *Principles of Adaptive Optics* (Academic, Boston, 1991).
13. R. Q. Fugate, "Laser beacon adaptive optics," *Opt. Photon. News* **5**(6), 14–19 (1994).
14. M. C. Roggemann, V. M. Bright, S. R. Hick, and W. D. Cowan, "Use of micro-electromechanical deformable mirrors to control aberrations in optical system," *Opt. Eng.* **36**, 1326–1338 (1997).
15. G. B. Love, J. S. Fender, and S. R. Restaino, "Adaptive wavefront shaping with liquid crystals," *Opt. Photon. News* (October 1995), pp. 16–21.
16. H. J. Kushner and D. S. Clark, *Stochastic Approximation Methods for Constrained and Unconstrained Systems* (Springer-Verlag, New York, 1978).
17. J. C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE Trans. Autom. Control* **37**, 332–341 (1992).
18. J. C. Spall, "A stochastic approximation technique for generating maximum likelihood parameter estimates," in *Proceedings of the American Control Conference* (Institute of Electrical and Electronics Engineers, New York, 1987), pp. 1161–1167.
19. G. Cauwenberghs, "A fast stochastic error-descent algorithm for supervised learning and optimization," in *Advances in Neural Information Processing Systems*, S. J. Hanson, J. D. Cowan, and C. L. Giles, eds. (Morgan Kaufmann, San Mateo, Calif, 1993), Vol. 5, pp. 244–251.
20. J. C. Spall, "Adaptive stochastic approximation by the simultaneous perturbation method," *IEEE Trans. Autom. Control* **45** (to be published); in condensed form in *Proceedings of the IEEE Conference on Decision and Control* (Institute of Electrical and Electronics Engineers, New York, 1998), pp. 3872–3879.
21. B. Flower and M. Jabri, "Summed weight neuron perturbation: an $O(n)$ improvement over weight perturbation," in *Advances in Neural Information Processing Systems*, S. J. Hanson, J. D. Cowan, and C. L. Giles, eds. (Morgan Kaufmann, San Mateo, Calif, 1993), Vol. 5, pp. 212–219.
22. G. Cauwenberghs, "A learning analog neural network chip with continuous-recurrent dynamics," in *Advances in Neural Information Processing Systems*, S. J. Hanson, J. D. Cowan, and C. L. Giles, eds., (Morgan Kaufmann, San Mateo, Calif. 1994), Vol. 6, pp. 858–865.
23. G. Cauwenberghs, "Analog VLSI stochastic perturbative learning architectures," *Int. J. Analog Integr. Circuits Signal Process.* **13**, 195–209 (1997).
24. M. A. Vorontsov, G. W. Carhart, and J. C. Ricklin, "Adaptive phase-distortion correction based on parallel gradient-descent optimization," *Opt. Lett.* **22**, 907–909 (1997).
25. V. I. Polejaev and M. A. Vorontsov, "Adaptive active imaging system based on radiation focusing for extended targets," in *Adaptive Optics and Applications*, R. Tyson and R. Fugate, eds., Proc. SPIE **3126**, 216–220 (1997).
26. B. M. Ter Haar Romey, ed., *Geometry-Driven Diffusion in Computer Vision* (Kluwer Academic, Dordrecht, The Netherlands, 1994).
27. Special issue on partial differential equations and geometry-driven diffusion in image processing and analysis, *IEEE Trans. Image Process.* **7** (1998).
28. T. R. O'Meara, "The multi-dither principle in adaptive optics," *J. Opt. Soc. Am.* **67**, 306–315 (1977).
29. M. A. Vorontsov, G. W. Carhart, D. V. Pruidze, J. C. Ricklin, and D. G. Voelz, "Image quality criteria for an adaptive imaging system based on statistical analysis of the speckle field," *J. Opt. Soc. Am. A* **13**, 1456–1466 (1996).
30. J. E. Pearson and S. Hansen, "Experimental studies of a deformable-mirror adaptive optical system," *J. Opt. Soc. Am.* **67**, 325–333 (1977).
31. S. A. Kokorowski, M. E. Pedinoff, and J. E. Pearson, "Analytical, experimental and computer simulation results on the interactive effects of speckle with multi-dither adaptive optics systems," *J. Opt. Soc. Am.* **67**, 333–345 (1977).
32. G. W. Carhart, J. C. Ricklin, V. P. Sivokon, and M. A. Vorontsov, "Parallel perturbation gradient descent in algorithm for adaptive wavefront correction," in *Adaptive Optics and Applications*, R. Tyson and R. Fugate, eds., Proc. SPIE **3126**, 221–227 (1997).
33. M. A. Vorontsov and V. P. Sivokon, "Stochastic parallel gradient descent technique for high-resolution wavefront phase distortion correction," *J. Opt. Soc. Am. A* **15**, 2745–2758 (1998).
34. This is not necessarily true if the division form $\tilde{J}'_j = \delta J_j / \delta u_j$ is used.
35. G. Andréou and K. A. Boahen, "Translinear circuits in sub-threshold MOS," *Analog Integr. Circuits Signal Process.* **9**, 141–166 (1996).
36. G. Cauwenberghs and M. A. Bayoumi, eds., *Learning on Silicon* (Kluwer Academic, Boston, 1999).
37. M. A. Vorontsov, "High-resolution adaptive phase distortion compensation using a diffractive-feedback system: experimental results," *J. Opt. Soc. Am. A* **16**, 2567–2573 (1999).
38. G. Cauwenberghs, "Analog VLSI recurrent neural network learning a continuous-time trajectory," *IEEE Trans. Neural Netw.* **41**, 827–829 (1994).
39. R. T. Edwards, M. Cohen, G. Cauwenberghs, M. A. Vorontsov, and G. W. Carhart, "Analog VLSI parallel stochastic optimization for adaptive optics," in *Learning on Silicon*, G. Cauwenberghs and M. A. Bayoumi, eds. (Kluwer Academic, Boston, 1999), Chap. 1, pp. 359–382.
40. G. W. Carhart, M. A. Vorontsov, R. T. Edwards, M. Cohen, and G. Cauwenberghs, "Adaptive wavefront control using a VLSI implementation of the parallel perturbation gradient descent algorithm," in *High-Resolution Wavefront Control: Methods, Devices, and Applications*, J. Gonglewski and M. Vorontsov, eds., Proc. SPIE **3760**, 61–66 (1999).
41. *Users Manual*, Meadowlark Optics Inc. HEX127 phase modulator, 1997.
42. G. Vdovin, S. Middelhoek, and P. M. Sarro, "Technology and applications of micromachined silicon adaptive mirrors," *Opt. Eng.* **36**, 1382–1390 (1997).
43. L. Zhu, P. Sun, D. Bartsch, W. R. Freeman, and Y. Fainman, "Adaptive control of a micromachined continuous-membrane deformable mirror for aberration compensation," *Appl. Opt.* **38**, 168–176 (1999).