

Microelectronic Very Large Scale Integrated Systems for Adaptive Optics and Advanced Imaging

Marc Cohen[†], Gert Cauwenberghs[†]
Mikhail Vorontsov[‡] and Gary Carhart[‡]

[†]Johns Hopkins University, Dept. of Electrical and Computer Engineering
3400 North Charles Street, Baltimore, MD 21211

[‡]Army Research Laboratory, Intelligent Optics Laboratory, Adelphi, MD 20783

ABSTRACT

We show that custom adaptive very large scale integrated (VLSI) circuit controllers can directly control commercially available wavefront phase correctors. These controllers achieve real-time wavefront compensation under conditions of strong intensity scintillations. Their control strategy is based on the optimization of a measurable performance index. Optimization is carried out using parallel perturbative stochastic gradient descent. We describe VLSI image-plane processors designed to compute a variety of application specific performance metrics in real-time.

Keywords: adaptive optics, wavefront correction, real-time, VLSI, MEMS mirrors, LC SLM, stochastic optimization

1. INTRODUCTION

Adaptive real-time control of an optical field's wavefront spatial shape is an important task for a variety of applications: astronomical observations from the Earth, ground-to-ground and ground-to-air laser communication and imaging, laser technology, microscopy, lithography and medical imaging. These growing needs for adaptive wavefront control require adaptive systems that are fast (msec. adaptation times), have high resolution (10^6 to 10^8 control channels), small size, low weight and low power. In fact, existing adaptive optics techniques are quite slow (10msec. adaptation time), low resolution ($< 10^3$ control channels), expensive ($> \$1000$ per channel), large and power hungry. This situation is rapidly changing with a new generation of commercially available wavefront compensating hardware: high resolution liquid crystal (LC) spatial phase modulators and micro-electromechanical systems (MEMS). These new devices can potentially provide high speed (sub-msec.) and high resolution wavefront shaping.

To achieve *real-time* wavefront distortion compensation, it is necessary to access all elements of the wavefront controller simultaneously in parallel. For wavefront correctors having many elements ($N > 128$), we anticipate a wiring problem; each control element needs a wire connecting it to its controller. In the future, high resolution wavefront correctors ($N \approx 10^6$ to 10^8) will require controlling electronics to be located locally at each control element. New technologies such as silicon on sapphire and planarized CMOS fabrication will facilitate the fabrication of wavefront corrector and wavefront controller on the same substrate.

What we propose is a flexible strategy to control these wavefront correctors which is applicable to a variety of applications and suitable for real-time implementation. Realizing a model-based control technique in software, or even in hardware, is virtually impossible because models of time-varying atmospheric aberrations are either inadequate or too complex. We suggest a formulation independent of the specifics of the problem, a "model-free" approach which can be realized with massively parallel VLSI hardware. We recently demonstrated that our model-free adaptive VLSI controller directly connected to commercially available wavefront correctors (LC spatial light modulators and MEMS mirrors), achieve compensation of wavefront distortion under conditions of strong intensity scintillations.^{21,20,18}

An additional advantage of the "model-free" approach is that all specific requirements for system performance are "hidden" in the selected optimization metric or cost functional. The metric should be computed based on available wavefront information: image sharpness, laser beam profile, interference pattern, e.t.c. Depending on the type of adaptive optical system, the performance metric might be intensity radiation at the focus,^{13,27} image

Send correspondence to Marc Cohen, marc@bach.ece.jhu.edu

sharpness^{22,17} or scattered field statistical moments.²⁴ The fast computation of system performance metrics is an essential requirement for the success of “model-free” optimization.

In Section 2 we present an overview of model-free optimization applied to adaptive optics and concentrate on a gradient descent based stochastic approximation technique. We describe the VLSI implementation of our parallel perturbative stochastic gradient descent/ascent wavefront controller. In Section 3 we present VLSI image-plane processors for computing performance metrics (laser beam and image quality metrics) compatible with our control strategies. Experimental results from our VLSI adaptive optics controller driving a MEMS micro-mirror wavefront corrector are presented in Section 4. Also shown in Section 4 are some experimental results using our beam quality metric chip to supply feedback in the closed system control loop. The image quality metric chip’s ability to compute a measure of image sharpness is demonstrated through a simple experiment. We summarize and briefly discuss future technologies in Section 5

2. WAVEFRONT CORRECTOR CONTROL

In this section we present our “model-free” approach to direct control of the wavefront corrector. We introduce parallel perturbative stochastic gradient descent in the framework of gradient descent. The resulting algorithm is customized for VLSI implementation. We briefly describe our VLSI Adaptive Optics controller (AdOpt).

2.1. Gradient Descent

Let $\mathbf{u} = u(\mathbf{r}, t)$ be the vector representing the controlled compensating phase with $\mathbf{r} = \{x, y\}$ the vector in the plane orthogonal to the optical axis. Ideally, we wish to minimize the performance metric $J(\mathbf{u})$ by adjusting \mathbf{u} in the direction opposite to the first variation of the cost functional:

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

where τ is a time constant. For piston-type actuators, we can write the controlling phase vector $u(\mathbf{r}, t)$ in terms of a linear combination of each of the contributing elements $u_j(t)$ and the wavefront corrector influence functions $S_j(\mathbf{r})$,

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

Substituting (2) into (1) yields

$$\tau_j \frac{du_j(t)}{dt} = -\beta \frac{\partial J(\mathbf{u})}{\partial u_j}, \quad \text{for } j = 1, \dots, N, \quad (3)$$

with β an update coefficient having positive sign for metric minimization. Since we apply discrete-time control signals to the wavefront controller elements, a discrete time version of (3) is needed,

$$u_j^{(k+1)} = u_j^{(k)} - \gamma \frac{\partial J^{(k)}(\mathbf{u})}{\partial u_j^{(k)}}, \quad (4)$$

with γ the learning-rate coefficient. It can be shown that in the vicinity of an extremum point, the system described in (3) is always stable [ref Gert-NIPS, JOSA00].

2.2. The Model-Free Approach: Stochastic Gradient Approximation

In adaptive optics, true gradients of the cost function $\partial J(\mathbf{u})/\partial u_j$ cannot be accurately modeled (especially not for anisoplanatic distortions), and must be estimated from measurements. In early days of adaptive optics, the method of gradient approximation was known as aperture-tagging.¹³ Two techniques were used; time-division (*sequential*) and frequency-division (*multi-dither*). In the former, small sequential control signal perturbations δu_j are applied and correlated with measured system performance δJ_j to compute finite difference gradient approximations. In the latter, small perturbations of harmonic signals $\delta u_j = \alpha \sin(\omega_j t)$ with different dithering frequencies ω_j and small modulation amplitudes α are applied in parallel to all u_j . To estimate gradient components in the measured metric δJ_j , frequency dithered carriers are first demodulated by synchronous detectors and then low-pass filtered. There are

however disadvantages to both methods. For sequential aperture-tagging, the time required for gradient estimation is proportional to the number of control channels N . For multi-dither aperture-tagging, the time required for gradient estimation is independent of N but the bandwidth requirement for the wavefront corrector grows approximately as $2N\omega_0$, where ω_0 is the control system bandwidth.²³

Among model-free optimization techniques, parallel perturbative stochastic gradient descent,⁴ is perhaps the most promising for adaptive optics applications.^{28,16} This algorithm, also called simultaneous perturbation stochastic approximation (SPSA),²⁶ lends itself well to VLSI implementation.⁷ Sequential¹⁵ and parallel forms^{5,6} have been implemented in VLSI hardware.

In this method, small stochastic perturbations $\delta u_j^{(k)}$ are applied simultaneously in parallel to all control channels $j = 1, \dots, N$. The measured change in performance,

$$\delta J = J(u_1 + \delta u_1, \dots, u_j + \delta u_j, \dots, u_N + \delta u_N) - J(u_1, \dots, u_j, \dots, u_N) \quad (5)$$

is used to estimate the true gradient $\partial J / \partial u_j$. Taylor series expansion of (5) yields

$$\delta J = \sum_{j=1}^N \left(\frac{\partial J}{\partial u_j} \right) \delta u_j + \frac{1}{2} \sum_{j=1}^N \sum_{i=1}^N \left(\frac{\partial^2 J}{\partial u_j \partial u_i} \right) \delta u_j \delta u_i + \dots \quad (6)$$

Multiplying both sides of (6) by δu_l results in

$$\delta J \delta u_l = \frac{\partial J}{\partial u_l} (\delta u_l)^2 + \left[\sum_{j \neq l}^N \frac{\partial J}{\partial u_j} \delta u_j \delta u_l + \sum_j \sum_i \frac{\partial^2 J}{\partial u_j \partial u_i} \delta u_j \delta u_i \delta u_l + \dots \right] \quad (7)$$

The quantity in square parentheses on the right hand side of (7) is a “noise-term”. Assuming perturbations $\{\delta u_l\}$ are random variables, we take expectations on both sides of (7),

$$\langle \delta J \delta u_l \rangle = \frac{\partial J}{\partial u_l} \langle \delta u_l^2 \rangle + \langle \text{noise-term} \rangle \quad (8)$$

Correlating δJ with a particular perturbation δu_l yields an estimate of the true gradient component in direction l . Choosing statistically independent perturbations ($\langle \delta u_j \delta u_i \rangle = \sigma^2 \delta_{ij}$) having zero means ($\langle \delta u_j \rangle = 0$) and symmetric probability densities (about their means) results in the expected value of the noise-term being of order σ^4 . When we replace the true gradient in (4) with the stochastic approximation in (8), the update rule becomes,

$$u_j^{(k+1)} = u_j^{(k)} - \gamma \delta J^{(k)} \delta u_j^{(k)}. \quad (9)$$

Overall, parallel perturbative stochastic gradient descent takes \sqrt{N} fewer iterations than sequential aperture-tagging and \sqrt{N} more than gradient descent.⁴ The signal-to-noise ratio for this technique is the same as that for sequential aperture-tagging, that is N^2 better than the multidither technique.¹⁸

2.3. VLSI Implementation: The Adaptive Optics Controller

We generate pseudo-random Bernoulli distributed perturbations with identical amplitudes $|\delta u_j^{(k)}| = \sigma$ and random signs $\text{sgn}(\sigma) = \pi = \pm 1$ on chip using linear feedback shift-registers. We apply these perturbations differentially to obtain a more accurate and reliable estimate of the true gradient,

$$\delta J^{(k)} = J(u_1^{(k)} + \delta u_1^{(k)}, \dots, u_N^{(k)} + \delta u_N^{(k)}) - J(u_1^{(k)} - \delta u_1^{(k)}, \dots, u_N^{(k)} - \delta u_N^{(k)}). \quad (10)$$

The measured $\delta J^{(k)}$ is decomposed into its magnitude $|\delta J^{(k)}|$ and $\text{sgn}(\delta J^{(k)})$. As a result we are able to simplify (9) into a form that can be realized compactly in VLSI,

$$u_j^{(k+1)} = u_j^{(k)} - \gamma' \mathbf{XNOR} \left(\text{sgn}(\delta J^{(k)}) \pi^{(k)} \right), \quad (11)$$

where $\gamma' = \gamma / (2\sigma)$ and \mathbf{XNOR} is the Boolean exclusive NOR operation. Figure 1(a) shows a photomicrograph of the AdOpt chip. Each chip controls 19 channels in parallel. Figure 1(b) shows the circuit diagram for each channel which consists of four main parts: an analog memory circuit that maintains each state variable $u_j^{(k)}$, a capacitive divider circuit that differentially perturbs each state variable, a charge-pump circuit that adapts each state variable and an output driver specific to the application. Many chips can be used together in a system that controls many channels. For a complete description of the circuitry, see Edwards *et al.*²⁵

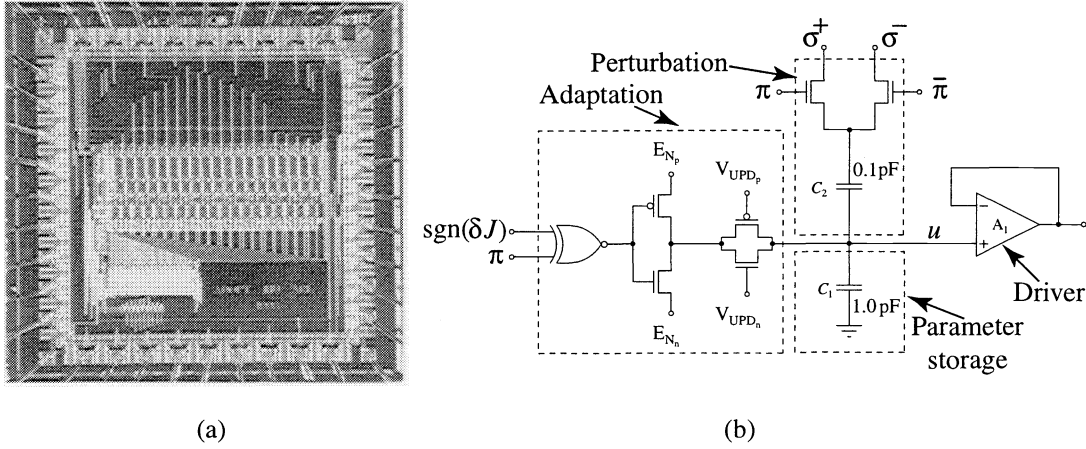


Figure 1. (a) Photomicrograph of the $2.2 \times 2.25\text{mm}^2$ AdOpt chip fabricated through MOSIS in $1.2\mu\text{m}$ CMOS technology and (b) Circuit diagram for a single channel

3. SENSING QUALITY: COMPUTING PERFORMANCE METRICS

For aperture tagging and model-free adaptive optics techniques in general, an appropriate “oracle” of the system’s performance must be specified and computed. Furthermore, the metric $J(\mathbf{u})$ must be defined for the particular application and must be computable in a time much less than the characteristic time of the turbulence.

Several “sharpness” metrics have been suggested.²² For laser beam focusing, metrics involve computing the sum over the image plane of functions of the two-dimensional beam intensity distribution $I(\mathbf{r}, t)$,

$$J_{beam} = \mathcal{F}_r \{I(\mathbf{r}, t)\}. \quad (12)$$

Computing these metrics has proven to be computationally too expensive for real-time applications. More recently, Vorontsov *et.al.*¹⁷ have suggested speckle field metrics based on the Fourier spectrum of $I(\mathbf{r}, t)$. The spectrum can be produced optically in real-time using Fourier optics and measured by a photodetector. The photodetector output voltage requires minimal post processing. The authors have shown that this technique is essentially equivalent to metrics composed of sums of functions of intensity distributions.

Metrics for “image sharpness” and “focus” have been proposed¹⁴ which take the form

$$J_{image} = \int_{\mathbf{r}} \|\nabla I(\mathbf{r}, t)\|^\nu d^2\mathbf{r}, \quad (13)$$

where ν represents some specified norm. Delbrück^{8,9} has built VLSI implementations of (13) with $\nu = 1$ for computing image focus in real-time.

Since we know how to design sophisticated VLSI image-plane processors to compute sums of functions of intensity distributions in real-time,^{19,2,8} we have the freedom to construct metrics designed specifically for a particular application.

3.1. Beam Variance Metric

Our first approach consists of measuring the *variance* of the beam intensity distribution relative to the square of the intensity,

$$\text{BVM} = \frac{N \cdot M}{E^2} \sum_i^N \sum_j^M I_{i,j}^2, \quad (14)$$

where $I_{i,j}$ represents the intensity at the $(i, j)^{th}$ pixel, N and M the number of pixels per row and column respectively and $E = \sum_i^N \sum_j^M I_{i,j}$. This *beam variance metric* (BVM) is suitable for point sources and increases monotonically as the width of the focused beam decreases. We also compute the beam centroid, (X_C, Y_C) .

3.1.1. VLSI Implementation: Beam Variance Metric

A photomicrograph of the BVM chip is shown in Figure 2(a). It consists of a 20×20 pixel array surrounded by a ring of dummy pixels to minimize the edge effects.¹ The pixel circuit diagram is shown in Figure 2(b). Each pixel measure $70\mu\text{m}$ on a side. Subthreshold translinear CMOS circuits² in each pixel square the photocurrent while a translinear bipolar circuit¹¹ at the periphery normalizes by the square of the total energy. Row and column decoders at the periphery allow random access for image readout. Centroid circuits at the periphery compute the two-dimensional centroid of the received image.¹⁰

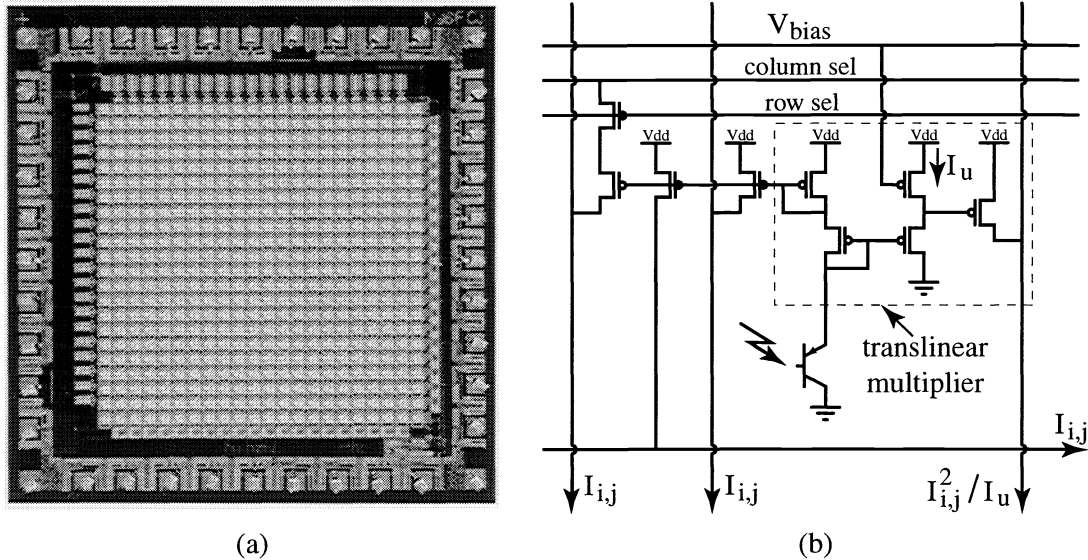


Figure 2. Photomicrographs of $2.2 \times 2.25\text{mm}^2$ BVM chip manufactured through MOSIS in $1.2\mu\text{m}$ CMOS technology: (a) 20×20 pixel BVM chip and (b) pixel circuit diagram.

3.2. Beam Width Metric

Our second approach directly quantifies the average weighted distance or second moment of the intensity distribution of the beam around its centroid,

$$\text{BWM} = X_W^2 + Y_W^2, \quad (15)$$

with

$$\begin{aligned} X_W^2 &= \frac{1}{E} \sum_j (j - X_C)^2 \cdot W_j, & Y_W^2 &= \frac{1}{E} \sum_i (i - Y_C)^2 \cdot W_i \\ X_C &= \frac{1}{E} \sum_j j \cdot W_j, & Y_C &= \frac{1}{E} \sum_i i \cdot W_i \\ W_j &= \sum_i I_{i,j}, & W_i &= \sum_j I_{i,j} \end{aligned} \quad (16)$$

The *beam width metric* (BWM) is a more general metric than the BVM since it is not limited to point sources.

3.2.1. VLSI Implementation: Beam Width Metric

The architecture of the BWM chip is similar to that of the BVM chip. Figure 3(a) shows a photomicrograph of the 40×40 pixel array. The pixel design shown in Figure 3(b) is however different. The BWM does not require squaring of the photocurrent at the pixel level. Its pixel is just $40\mu\text{m}$ on a side. All computations, namely (X_W, Y_W) , (X_C, Y_C) and E , are done at the periphery using subthreshold CMOS translinear circuits.




Figure 3. Photomicrographs of $2.2 \times 2.25 \text{ mm}^2$ BWM chip manufactured through MOSIS in $1.2 \mu\text{m}$ CMOS technology: (a) 40×40 pixel BWM chip and (b) pixel circuit diagram.

3.3. Image Quality Metric

Our approach to quantifying the quality of a received image is to measure image sharpness in terms of the energy at high spatial frequencies. Defocus of an image results in attenuation of the high spatial frequencies. We chose a simple measure of high spatial frequency energy keeping in mind that it must be implementable in VLSI: it is the energy normalized sum over all pixels of the absolute value of pixel photocurrents convolved with a spatial highpass filter,

$$\text{IQM} = \frac{1}{E} \sum_i \sum_j |I_{i,j} * K| \quad (17)$$

$$\text{with } K = \begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix} \quad (18)$$

3.3.1. VLSI Implementation: Image Quality Metric

Figure 4(a) shows a photomicrograph of the IQM chip. It consists of a 22×22 array of pixels with column and row decoders top and right and row and column select shift registers left and bottom. The pixel circuit measuring $120 \mu\text{m}$ on a side is shown in Figure 4(b). Each pixel distributes its current to its four neighbors and receives currents from its four neighbors. The spatial highpass filtered current and its absolute value are computed locally in each pixel. The chip is also a random access imager: each pixel contains a readout transistor which directs its current to the periphery when selected. We have included programmable shift registers to select a particular subset of pixels, so that only their currents contribute to a *local* image quality metric which can be useful for synthetic imaging applications.³

4. EXPERIMENTAL SETUP AND RESULTS

First we demonstrate how the AdOpt chip was used to control a 37 element MEMS mirror and x,y -tilt mirror in a laser beam focusing task under conditions of strong turbulence. We used a simple performance metric: the focused energy through a pinhole. Next we demonstrate that by using the beam quality metric calculated using the BVM chip, we improve the shape of the focused beam while maintaining speed of operation. Lastly we show how the IQM chip produces a valid measure of image quality.

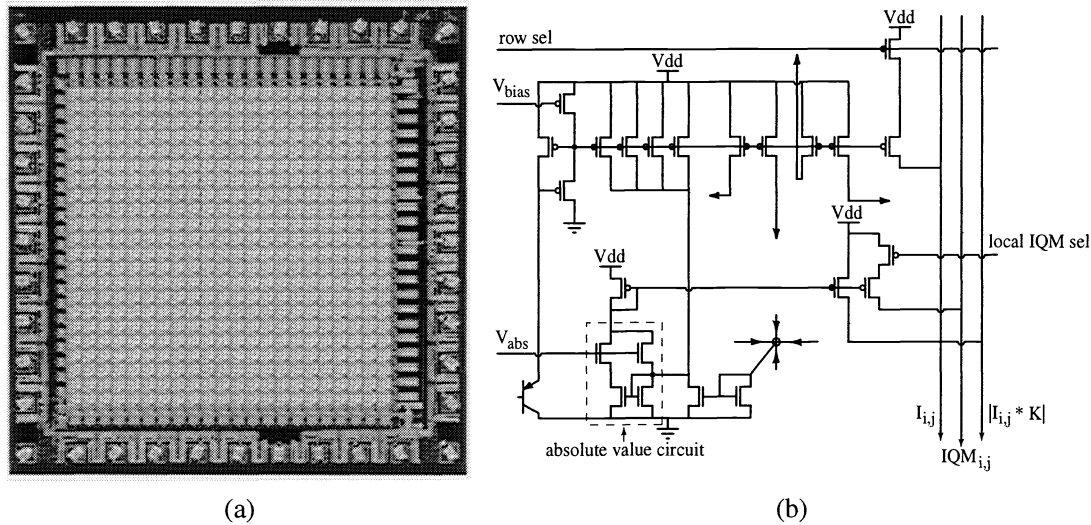


Figure 4. Photomicrographs of $2.2 \times 2.25 \text{ mm}^2$ IQM chip manufactured through MOSIS in $1.2 \mu\text{m}$ CMOS technology: (a) 22×22 pixel IQM chip and (b) pixel circuit diagram.

4.1. Adaptive Optics Controller

We characterized performance of the VLSI systems using an adaptive laser transmitter task. A simplified schematic of our experimental setup is shown in Figure 5. The beam from an Argon laser ($\lambda = 514 \text{ nm}$) was expanded to a diameter of 12 mm, reflected off a 2-degrees-of-freedom x, y -tilt mirror, and then reflected off the 37-element MEMS mirror. The beam passed over a convection heater which introduced turbulence.¹² The now distorted beam passed through a lens which allowed us to image the beam with a CCD camera in the focal plane. The beam-splitter steered the beam through a mask/pinhole (target) also at the focal plane, onto a photodetector. We used the output of the

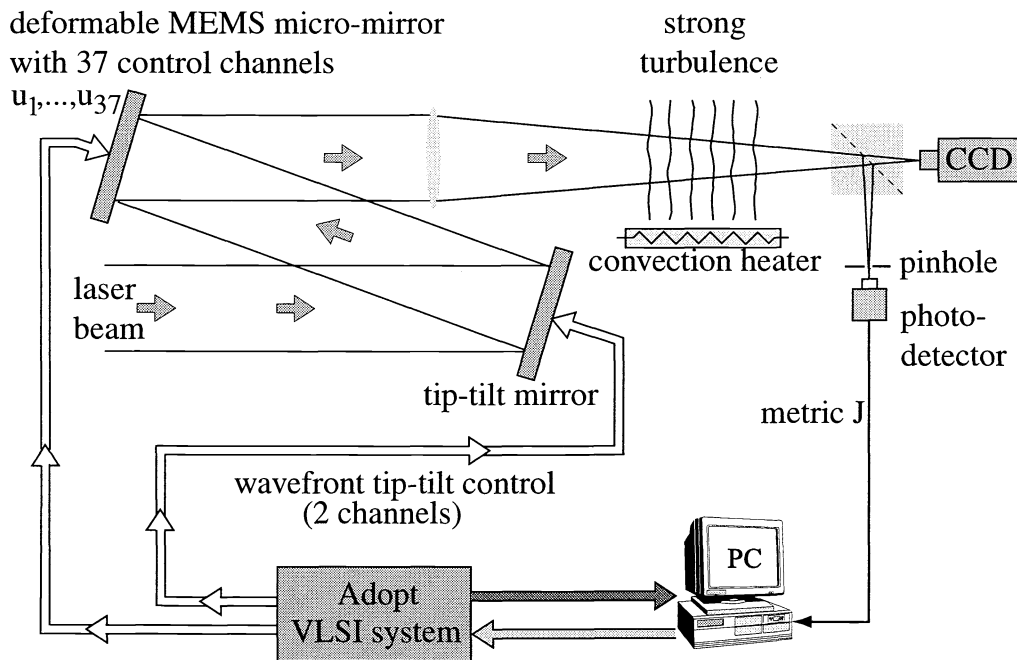


Figure 5. Experimental setup for the real-time laser-beam transmitter system

photodetector (voltage) as a beam-quality metric J . When the beam's wavefront has been optically corrected, J will be large. The photodetector output J was supplied to a PC which computed $|\delta J|$ and $\text{sgn}(\delta J)$ and passed these results on to the AdOpt wavefront controller. Three AdOpt chips were used, two to control the 37 element mirror and a third to control the x,y -tilt mirror.

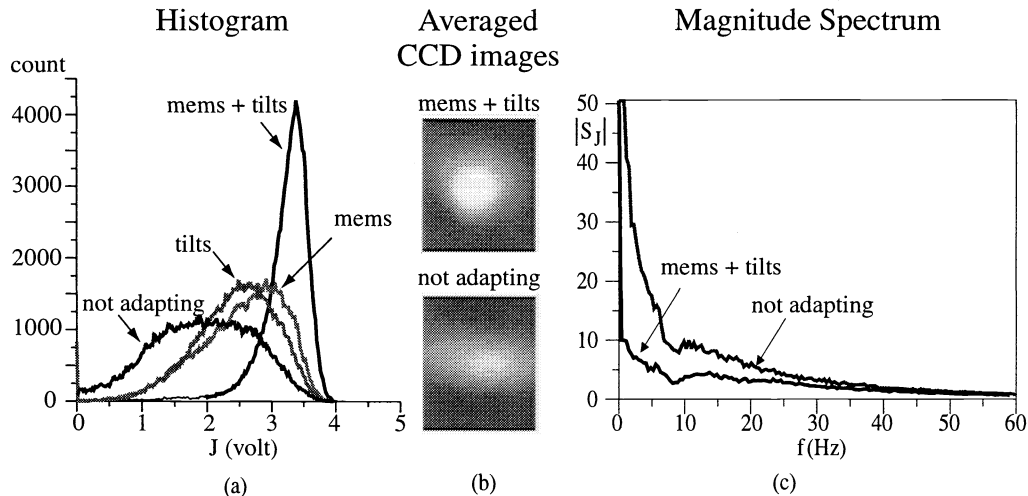


Figure 6. Experimental results for laser beam focusing under conditions of severe intensity scintillations

The results of adaptive laser-beam transmitter in the presence of turbulence are shown in Figure 6. We investigated four different experimental conditions: *not adapting* refers to the system's performance without the AdOpt system (i.e. open-loop), *tilts* refers to the AdOpt system controlling only the x,y -tilt mirror, *mems* refers to the AdOpt system controlling only the 37-element MEMS mirror and *mems + tilts* refers to control of both the x,y -tilt mirror and MEMS mirror. In each experimental condition a trial consisted of starting the system off from a "bad" state (i.e. low performance metric J state) and instructing the system to maximize J . We repeated this 500 times per experimental condition, collecting 500 data points (350msec. of data) at a frame rate of approximately 1.4KHz per iteration.

In Figure 6(a) we plot a histogram of J for all trials from the four experimental conditions. The *not adapting* case shows a broad spread in J with a mean just less than 2. Adding control of *tilts* or *mems* alone reduces the spread in J due to turbulence and improves its mean value $\langle J \rangle$. With both *mems + tilts*, the variance in J is significantly reduced and $\langle J \rangle$ is almost twice that of the reference *non adapting* case. The averaged CCD images corresponding to the *non adapting* and the *mems + tilts* trials are shown in Figure 6(b). Finally, in Figure 6(c) we plot the magnitude spectra of the performance metric $|S_J|$ for the *not adapting* and *mems + tilts* experiments. In both cases, most of the signal energy is below 20Hz, however in the *mems + tilt* case, there is a strong peak of energy at 0Hz (the graph was truncated at 50dB). These magnitude spectra also demonstrate the ability of the system to significantly reduce the effects of strong atmospheric turbulence on the laser focusing task.

4.2. Performance Metric Chips

The purpose of this experiment was to compare the performance of the system and the focused beam shape using the BVM chip to that using the pinhole metric. Figure 7 shows our experimental setup. In this experiment we were not introducing turbulence nor were we attempting to control the tip-tilt mirror (although the BVM also calculates the beam centroids X_C, Y_C). The AdOpt system controlled the 37 element MEMS mirror. Both pinhole and BVM metrics were supplied to the PC.

We performed 100 cycles of metric maximization-to-minimization using first the pinhole and then the BVM chip to supply the metric J . Figure 8 shows the results. We plot the normalized mean performance metric for both *pinhole* and *beam quality* (as measured using the BVM chip) metrics against iteration number for metric maximization followed by metric minimization. The *pinhole* metric produces a higher dynamic range than the *beam quality* metric. The images to the left and right of this plot explain why. To the left we plot the imaged beam on the surface of the BVM chip for metric maximum and to the right for metric minimum. The position and relative size of the pinhole




Figure 7. Experimental setup for comparing BVM to pinhole metric.

are marked for comparison on the chip image for the *pinhole* case. During metric maximization, the *beam quality* metric produces a compact circular beam (high output voltage) while during metric minimization, it produces a diffuse beam (low output voltage). The *pinhole* metric produces a less compact, elliptically shaped beam during maximization and for minimization simply steers the beam out of the pinhole (output voltage close to zero). For both metrics, maximization takes about 1 second while minimization takes about 0.5 seconds.




Figure 8. Experimental results comparing performance and shape of focused beam for *pinhole* and *beam quality* metrics.

On the bottom, we plot the standard deviation divided by the mean $\sigma_J / \langle J \rangle$ for the *pinhole* and *beam quality* metrics against the iteration number. For both metrics, $\sigma_J / \langle J \rangle$ is higher during minimization than maximization. This can be explained by the fact that the signal-to-noise ratio is lower during minimization and also because there

are many more ways to minimize the performance metric than to maximize it. In general, the *beam quality* metric exhibits low $\sigma_J / \langle J \rangle$ because the beam remains on the BVM chip during both minimization and maximization and therefore a metric signal is always present.

We are currently testing the BWM chip using the same experimental setup as that shown in Figure 7.

We also performed preliminary experiments on the Image Quality Metric chip. Figure 9(a) shows our setup. A white light source was used to illuminate a grayscale 35mm slide which was focused down onto the chip surface. The chip was mounted on a translation stage so that it could be moved into and out of the image plane. The 35mm slide consisted of a grayscale checkerboard pattern. We moved the chip 1mm on either side of the image plane in steps of $50\mu\text{m}$. At each position we recorded the IQM computed by the chip and scanned the image captured by the chip. The results are plotted in Figure 9(b). We plot the measured normalized IQM as a function of displacement from the image plane along with ± 3 std. dev. of the measured normalized IQM. We show four images scanned off the chip for various displacements from the image plane.

The results show desired properties of an image quality metric: namely, a high dynamic range of approximately 20; the largest value of IQM is obtained when the image is sharply focused on the chip; and the metric decreases monotonically away from this value for displacements on either side of the image plane.

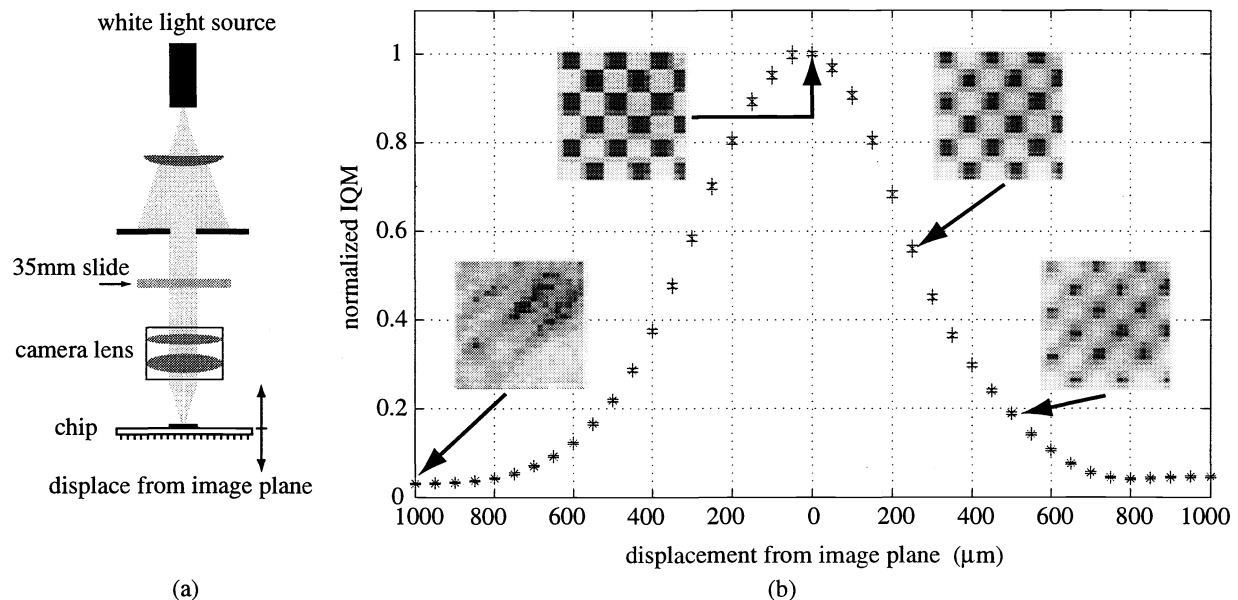


Figure 9. Testing the Image Quality Metric chip. (a) Experimental setup and (b) Experimental Results

5. CONCLUSION

We have described our VLSI “toolbox” for real-time adaptive optics applications. The AdOpt controller chip directly controls wavefront correctors, either MEMS or LC SLM, and our performance metric chips, BVM, BWM and IQM compute various measures of beam/image quality on the focal plane.

Because atmospheric turbulence, especially anisoplanatic turbulence, is too difficult and too computationally expensive to model accurately, we adopted a model-free approach to wavefront control. The technique is based on stochastic gradient descent/ascent performed on an application specific cost functional. We implemented parallel perturbative stochastic gradient descent/ascent in a VLSI chip, the AdOpt controller. We have shown its efficacy in experiments using a laser beam transmitter system and demonstrated significant reduction of wavefront distortion in the presence of severe intensity scintillations.

Since we chose a model-free approach, performance metrics are tailored to the particular application. For laser beam tasks, the shape, energy distribution and centroid of the beam can be computed in real-time at the focal plane. For imaging tasks, a measure of the high spatial frequency content of an image can be similarly computed. We have

demonstrated that our VLSI chips can be used successfully to compute a variety of quality metrics that provide appropriate real-time feedback signals to the AdOpt controller.

We acknowledge that in the future, when very high resolution wavefront correctors become available, we will face a wiring problem: each element of the corrector needs to be connected to its control channel. We expect to address this issue by using new fabrication technologies and architectures to integrate the wavefront corrector with local control circuitry. Silicon on sapphire technology will enable us to build LC SLM modulators with microelectronics on the same substrate. Planarized small feature size CMOS technology will enable us to build MEMS mirrors and microelectronics on the same substrate. Together with the control strategy we have developed, we will be able to build and control wavefront corrector systems having very high resolution with 10^6 to 10^8 state variables.

REFERENCES

1. A. Pavasovic, A.G. Andreou and C.R. Westgate. Characterization of subthreshold mos mismatch in transistors for vlsi systems. *Journal of VLSI Signal Processing*, 8:75–85, 1994.
2. A.G. Andreou and K.A. Boahen. Translinear circuits in subthreshold mos. *Analog Integrated Circuits and Signal Processing*, 9:141–166, 1996.
3. G.W. Carhart and M.A. Vorontsov. Synthetic imaging: nonadaptive anisoplanatic image correction in atmospheric turbulence. *Opt. Lett.*, 23(10):745–747, 1998.
4. G. Cauwenberghs. A fast stochastic error-descent algorithm for supervised learning and optimization. In *Advances in Neural Information Processing Systems*, 5, pages 244–251. Morgan Kaufman, San Mateo, CA, 1993.
5. G. Cauwenberghs. A learning analog neural network chip with continuous-recurrent dynamics. In *Advances in Neural Information Processing Systems*, 6, pages 858–865. Morgan Kaufman, San Mateo, CA, 1994.
6. G. Cauwenberghs. An analog vlsi recurrent neural network learning a continuous-time trajectory. *IEEE Transactions on Neural Networks*, 7(2), March 1996.
7. G. Cauwenberghs. Analog vlsi stochastic perturbative learning architectures. *J. Analog Integrated Circuits and Signal Processing*, 13(1-2):195–209, 1997.
8. T. Delbrück. A chip that focuses an image on itself. In *Analog VLSI Implementation of Neural Systems*, pages 171–188. Kluwer Academic, 1989.
9. T. Delbrück. Silicon retinas for autofocus. In *ISCAS 2000*, Geneva, May 2000.
10. S.P. DeWeerth and C.A. Mead. A two-dimensional visual tracking array. In Jonathan Allen and F. Thomson Leighton, editors, *Advanced Research in VLSI, Proceedings of the Fifth MIT Conference*, pages 259–275. MIT Press, 1988.
11. B. Gilbert. Current-mode circuits from a translinear viewpoint: A tutorial. In F.J. Lidgley C. Toumazou and D.G. Haigh, editors, *Analogue IC Design: the current-mode approach*, pages 11–91. Peter Peregrinus, Ltd., 1990.
12. J.D. Gonglewski and M.A. Vorontsov, editors. *Artificial turbulence for Imaging and wave propagation*. SPIE, 1998.
13. J.W. Hardy. Active optics: a new technology for the control of light. *Proc. IEEE*, 66:651–697, 1978.
14. B. Horn. Focusing. *Massachusetts Institute of Technology. Artificial Intelligence Laboratory*, Artificial Intelligence Memo 160, 1968.
15. M. Jabri and B. Flower. Weight perturbation: An optimal architecture and learning technique for analog vlsi feedforward and recurrent multilayered networks. *IEEE Transactions on Neural Networks*, 3(1):154–157, 1992.
16. 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.
17. 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(7):1456–1466, 1996.
18. M.A. Vorontsov, G.W. Carhart, M. Cohen and G. Cauwenberghs. Adaptive optics based on analog parallel stochastic optimization: analysis and experimental demonstration. *J. Opt. Soc. Am. A*, 17(8), 2000.
19. C.A. Mead. Adaptive retina. In Mead and Ismail, editors, *Analog VLSI Implementations of Neural Systems*. Kluwer Academic, 1989.

20. M.H. Cohen, G. Cauwenberghs, M.A. Vorontsov and G.W. Carhart. Adaptive wavefront correction: a hybrid vlsi/optical system implementing parallel stochastic gradient descent. In *Proc. of SPIE EUROPTO*, Florence, Italy, 1999.
21. M.H. Cohen, G. Cauwenberghs, R.T. Edwards, M.A. Vorontsov and G.W. Carhart. Adopt: Analog vlsi stochastic optimization for adaptive optics. In *Proc. of IJCNN*, Washington DC, 1999.
22. R.A. Muller and A. Buffington. Real-time correction of atmospherically degraded telescope images through image sharpening. *J. Opt. Soc. Am.*, 64(9):1200–1210, 1974.
23. J.E. Pearson and S. Hansen. Experimental studies of a deformable-mirror adaptive optical system. *J. Opt. Soc. Am.*, 67:325–333, 1997.
24. V.I. Polejaev and M.A. Vorontsov. Adaptive active imaging system based on radiation focusing for extended targets. *Proc. SPIE* (3126), 1997.
25. R.T Edwards, M.H. Cohen, G. Cauwenberghs, M.A. Vorontsov, and G.W. Carhart. Analog vlsi parallel stochastic optimization for adaptive optics. In G. Cauwenberghs and Magdy A. Bayoumi, editors, *Learning on Silicon*, pages 359–382. Kluwer Academic, Boston, Dordrecht, London, 1999.
26. J.C. Spall. A stochastic approximation technique for generating maximum likelihood parameter estimates. In *Proceedings of the American Control Conference*, pages 1161–1167, Minneapolis, 1987.
27. R.K. Tyson. *Principles of Adaptive Optics*. Academic Press, Boston, 1991.
28. 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.