Floating-Gate Adaptation for Focal-Plane Online Nonuniformity Correction

Marc Cohen, *Student Member, IEEE,* and Gert Cauwenberghs, *Member, IEEE*

***Abstract*—Stochastic adaptive algorithms are investigated for online correction of spatial nonuniformity in random-access addressable imaging systems. The adaptive architecture is implemented in analog VLSI, integrated with the photosensors on the focal plane. Random sequences of address locations selected with controlled statistics are used to adaptively equalize the intensity distribution at variable spatial scales. Through a logarithm transformation of system variables, adaptive gain correction is achieved through offset correction in the log-domain. This idea is particularly attractive for compact implementation using translinear floating-gate MOS circuits. Furthermore, the same architecture and random addressing provide for oversampled binary encoding of the image with equalized intensity histogram. The techniques apply to a variety of solid-state imagers, such as artificial retinas, active pixel sensors, and IR sensor arrays. Experimental results confirm gain correction and histogram equalization in a 64 $\times$ 64 pixel adaptive array integrated on a 2.2-mm $\times$ 2.25-mm chip in 1.2-μm CMOS technology.**

***Index Terms*—Adaptation, analog VLSI, CMOS imager, equalization, floating gate, focal plane, nonuniformity.**

I. Introduction

Since the seminal work by Carver Mead on neuromorphic floating-gate adaptation in the silicon retina [1], few groups have addressed the problem of online adaptive correction of nonuniformities on the focal plane in solid-state image sensor arrays [2] and neuromorphic vision sensors [3], [4], while most efforts have concentrated on nonadaptive correction using on-chip [5] or off-chip calibrated storage. Gain and offset nonuniformities in the photosensors and active elements on the focal plane contribute "salt-and-pepper" fixed-pattern noise at the received image, which limit the resolution and sensitivity of imaging systems. Flicker noise and other physical sources of fluctuation and mismatch make it a necessity to correct for these effects *online*, which is problematic since the image received is itself unknown. Existing "blind" adaptive algorithms for online correction are complex and the amount of computation required to implement them is generally excessive. Integration on the focal plane would incur a significant increase in active pixel size and a decrease in spatial resolution and fill-factor of the imager along with an increase in power consumption.

Manuscript received April 1, 2000; revised November 1, 2000. This work was supported by NSF Career Award MIP-9702346 and ONR Young Investigator Award N00014-99-1-0612. This paper was recommended by Associate Editor T. S. Lande.

The authors are with the Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD 21218 USA (e-mail: marc@bach.ece.jhu.edu; gert@bach.ece.jhu.edu).

Publisher Item Identifier S 1057-7130(01)02025-0.

In this paper we present a class of *stochastic* adaptive algorithms which integrate general nonuniformity correction with minimal, if not zero, overhead in the number of active components on the focal plane. In particular, we use floating-gate adaptive CMOS technology to implement a two-transistor adaptive gain element for online focal-plane compensation of current gain mismatch. The algorithms make effective use of the statistics of pixel intensity under randomly selected sequences of address locations and avoid the need for extra circuitry to *explicitly* compute spatial averages and locally difference the result. The resulting stochastic algorithms are particularly simple to implement.

The same stochastic algorithms for adaptive nonuniformity correction which take advantage of the spatial statistics of image intensity can be used to perform image intensity equalization and normalization on the focal plane. Equalization is a useful property because it maximizes the available dynamic range and assigns higher sensitivity to more statistically frequent intensities. At the same time, the image is converted into digital form avoiding the need for explicit analog-to-digital conversion.

In Section II we formulate the stochastic algorithms for adaptive nonuniformity correction. We show how a simple logarithm transform on offset correction allows us to use the same algorithms for gain correction. We briefly discuss intensity equalization in Section III as a natural extension of the stochastic rules of Section II. In Section IV we describe and analyze the floating-gate translinear current-mode VLSI implementation of the adaptive pixel. The system architecture including the external circuits used for experimental validation of our VLSI imager is described in Section V. Section VI shows experimental results for gain nonuniformity correction and image intensity equalization. We discuss these results in Section VII and give our conclusions in Section VIII.

II. Adaptive Nonuniformity Correction

Nonuniformity correction can be approached using two strategies: apply a uniform reference image to the static imager and ensure that all pixel outputs are equal [1], or drift natural scenes across the imager where each pixel subtracts its output from its spatially low-pass filtered output to derive an error signal [6]. The former is referred to as *static* nonuniformity correction (SNUC) and the latter as *scene-based* nonuniformity correction (SBNUC). Our imager can accommodate either type of mismatch correction strategy. The SBNUC algorithm has been implemented on the focal plane in CMOS and IR-based imagers [2] and has been successful in reducing offset mismatch.

In this paper, we will concentrate on SNUC to reduce current gain mismatch in a photo-transistor based CMOS imager or silicon retina. We will show how by applying a controllable voltage offset on a floating-gate transistor in each pixel, we achieve an adjustable, adaptive pixel current gain. Our system architecture also allows us to perform SBNUC through controlling the statistics of random address sequences.

First we set up the problem in terms of established online algorithms for offset correction. Then we show how this same algorithm can be extended to gain mismatch reduction through a simple logarithm transformation of system state variables.

A. Canceling Offset Nonuniformity

Fig. 1(a) schematically demonstrates the offset correction technique. The set of system equations is

$$y = x + o, \qquad z = y + q = x + o + q \tag{1}$$

where

- x input random (sensor) variable;
- y received input with unknown offset o;
- q applied offset correction;
- z corrected output.

For offset compensation, we want

$$o + q \equiv \text{const} \quad \forall \text{ pixels}. \tag{2}$$

A simple (gradient descent) adaptive rule to achieve this is [6]

$$\Delta q = -\alpha(z - z_{\text{ref}}) \tag{3}$$

which adjusts the output z on average toward a reference constructed by expressing variable degrees of smoothness in the image

$$z_{ref} = \begin{cases} \langle z \rangle, & \text{for SNUC} \\ \langle z \rangle_{\text{local}}, & \text{for SBNUC} \end{cases} \tag{4}$$

where the $\langle \cdot \rangle$ symbol represents spatial averaging at global and local scales, respectively, and α denotes the adaptation (or "learning") rate. Circuits implementing a locally differenced diffusive kernel, with adjustable space constant, to perform the computations in (3) are presented in [2]. We introduce a stochastic version of the rule in (3)

$$\Delta q_{\mathbf{r}(k)} = -\alpha \left(z_{\mathbf{r}(k)} - z_{\mathbf{r}(k-1)} \right) \tag{5}$$

where the subscripts $\mathbf{r}(k)$ and $\mathbf{r}(k-1)$ denote pixel addresses at consecutive time steps $(k-1)$ and k, respectively. Taking expectations on both sides of (5), for a particular pixel selected at time k, yields

$$\mathrm{E}\left[\Delta q_{\mathbf{r}(k)}\right] = -\alpha \left(z_{\mathbf{r}(k)} - \mathrm{E}\left[z_{\mathbf{r}(k-1)} \right] \right)$$

which depends on the statistics of the consecutive address selections as determined by the conditional transition probabilities (densities) $p(\mathbf{r}(k-1)|\mathbf{r}(k))$. Therefore, by controlling the statistics $p(\mathbf{r}(k-1)|\mathbf{r}(k))$ through proper choice of the random sequence of addresses $\mathbf{r}$, we can implement, on average, the spatial convolution kernels needed to implement both SNUC and SBNUC in (4). In particular, for a random sequence with $\mathbf{r}(k-1)$ and $\mathbf{r}(k)$ independent [i.e., $p(\mathbf{r}(k-1)|\mathbf{r}(k)) = p(\mathbf{r}(k-1))$]

$$\mathrm{E}\left[z_{\mathbf{r}(k-1)} \right] = \langle z \rangle. \tag{6}$$

Whereas, if $\mathbf{r}(k-1)$ and $\mathbf{r}(k)$ are related by embedding memory in the address sequence (e.g., through inertia, or imposing limits on $\Delta\mathbf{r}$)

$$\mathrm{E}\left[z_{\mathbf{r}(k-1)} \right] = \langle z \rangle_{\text{local}}. \tag{7}$$

Equation (5) is a stochastic online version of SNUC and likewise, (7) implements stochastic SBNUC. Hardware requirements can be further simplified by thresholding the update in (5) into the pilot-rule

$$\Delta q_{\mathbf{r}(k)} = -\alpha \text{ sign } \left(z_{\mathbf{r}(k)} - z_{\mathbf{r}(k-1)} \right) \tag{8}$$

with fixed-size update increments and decrements.

B. Canceling Gain Nonuniformity

The gradient descent formulation [6] also adaptively compensates for gain mismatch, although it does not prevent the gain from becoming negative. Our approach is to relate gain correction, under the positivity constraint imposed by current-domain circuits, to offset correction through a logarithm transformation. This transformation has a physical meaning which can be exploited in the hardware implementation as discussed in the next section. Fig. 1(b) schematically illustrates the concept of gain mismatch correction in relation to Fig. 1(a).

The system is governed by

$$y' = a x', \quad z' = A y' = A a x' \tag{9}$$

which transform into

$$\ln z' = \ln A + \ln a + \ln x' \tag{10}$$

so that for gain nonuniformity correction

$$\ln A + \ln a \equiv \text{const} \quad \forall \text{ pixels}. \tag{11}$$

By identifying correspondending terms (in particular, $\ln A = q$ or $A = e^q$ and $\ln a = o$) in (1) and (10), and because of the monotonicity of the logarithmic map, the learning rule (8) can be rewritten as

$$\Delta q_{\mathbf{r}(k)} = -\alpha \text{ sign } \left(z'_{\mathbf{r}(k)} - z'_{\mathbf{r}(k-1)} \right) \tag{12}$$

which in turn can be expressed as a stochastic online learning rule with relative gain increments

$$\Delta A_{\mathbf{r}(k)} = -\alpha A_{\mathbf{r}(k)} \text{ sign } \left(z'_{\mathbf{r}(k)} - z'_{\mathbf{r}(k-1)} \right). \tag{13}$$

III. Intensity Equalization

The constant terms (*const*) both in offset (2) and gain (11) correction are undefined and not regulated during the adaptation. This problem can be circumvented by properly normalizing the acquired image. One particularly attractive approach to normalization is to equalize the image intensity histogram, which in addition to mapping the intensity range to unity also produces a

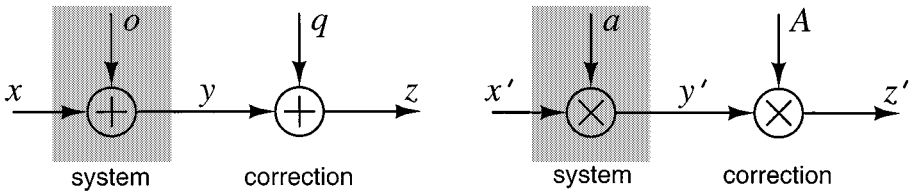

(a) (b)

Fig. 1. (a) Offset correction. (b) Gain correction.

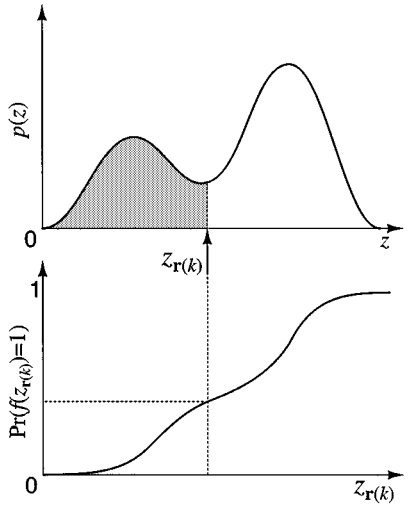

Fig. 2. Input intensity probability density function (top) and corresponding mean-rate transfer function (bottom) for intensity equalization and normalization.

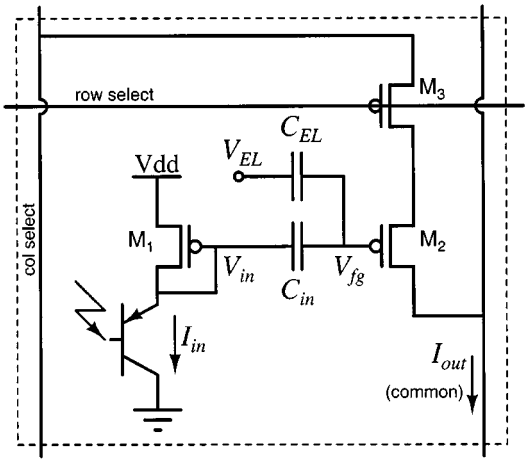

Fig. 3. Floating gate adaptive pixel circuit.

maximum entropy coded output [7]. Incidentally, the same stochastic algorithms (8) and (13) for nonuniformity correction can also be used for histogram equalized image coding.

Pixel intensities are mean-rate encoded in a single-bit oversampled representation akin to delta-sigma modulation [8], although without the need for integration nor any other processing at the pixel level. This could be compared with a popular scheme for neuromorphic multi-chip systems, the address-event communication protocol [9], in which sparse pixel-based events such as spiking action potentials are communicated asynchronously across chips. In the technique described here, addresses are not event-based, but supplied synchronously, with prescribed random spatial statistics.

In particular, we code the image in terms of the bits obtained by comparing $z_{\mathbf{r}(k)}$ and $z_{\mathbf{r}(k-1)}$ as in (8) or $z'_{\mathbf{r}(k)}$ and $z'_{\mathbf{r}(k-1)}$ as in (13). If larger, a "1" symbol is transmitted, else a "0." The selected address is either part of the transmitted code or is generated at the receiver end from the same random seed. Thus, the code is defined as

$$f\left(z_{\mathbf{r}(k)}\right) = \operatorname{sign}\left(z_{\mathbf{r}(k)} - z_{\mathbf{r}(k-1)}\right). \tag{14}$$

By selecting random addresses with controlled spatial statistics as in (4), this code effectively compares a selected pixel's intensity with a base value that is either a global or local average. The probability of "1" is the fraction of pixels in that neighborhood with intensity lower than the present pixel. This is illustrated in Fig. 2, in which the mean-rate pixel activity is given by the cumulative probability density function

$$\Pr\left(f\left(z_{\mathbf{r}(k)}\right) = 1\right) = \int_0^{z_{\mathbf{r}(k)}} p(z)\,dz. \tag{15}$$

This corresponds to intensity equalization and normalization of the image, a desirable feature for maintaining a large dynamic range in image acquisition [10]. As seen in Fig. 2, the coding transfer function assigns higher sensitivity to statistically more frequent intensities. The uniform distribution and maximum entropy encoding obtained by this transformation is a well-known result and appears to take place in biological phototransduction as well [7]. The mechanism of image equalization as achieved here is unique in that it is due to statistical techniques in an oversampled representation, and the statistics of the address sequence can be tailored to control the size of the neighborhood for different spatial scales of intensity normalization.

IV. Focal-Plane VLSI Implementation

Rather than implementing (13) directly, we make use of the exponential relationship between voltage and current in a (subthreshold) MOS transistor to encode a current *gain* as the exponential of a differential voltage across a *floating gate* capacitor. The increments and decrements Δq in (12) are then naturally implemented by hot-electron injection and tunneling across the floating gate oxide [11].

A. The Pixel

The pixel circuit diagram is shown in Fig. 3. A vertical *pnp* bipolar transistor converts photon energy to emitter current $I_{\mathbf{in}}$

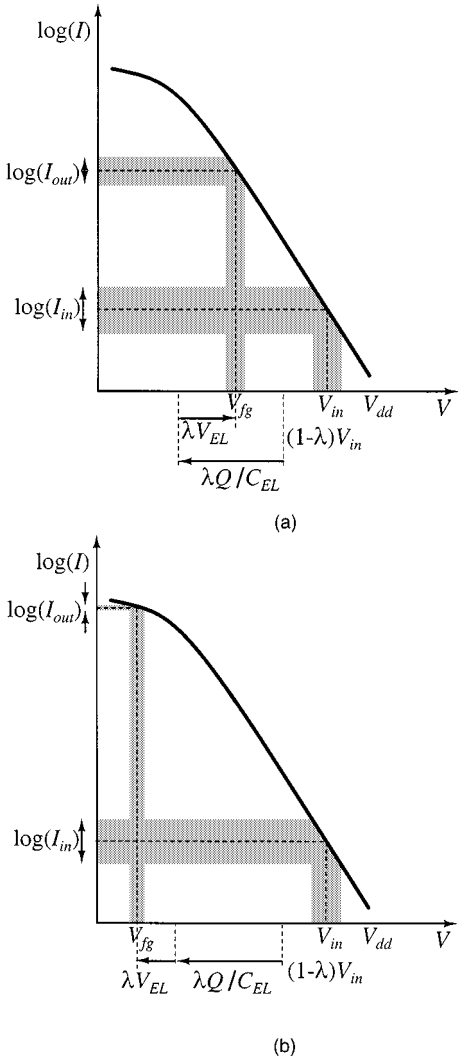

Fig. 4. Pictorial interpretation of the contributions of λ, V_{EL}, and Q/C_{EL} to the pixel current transfer function: (a) for sub-threshold output current $I_{\rm out}$ and (b) for above-threshold output current I_{rmout}.

with current gain β. Transistors M_1 and M_2 form a current mirror with adjustable current gain [12]. M_2 is a floating-gate transistor with two control electrodes; $V_{\rm in}$ and V_{EL} set V_{fg} through capacitive coupling

$$V_{fg} = \lambda\left(V_{EL} + \frac{Q}{C_{EL}}\right) + (1-\lambda)V_{\rm in} \tag{16}$$

with Q the charge injected/tunneled onto the floating gate, $\lambda = C_{EL}/(C_{EL} + C_{\rm in}) \approx 0.3$ and V_{EL} an externally applied global voltage for all pixels. The pixel's output current $I_{\rm out}$ is sourced by transistor M_2 and measured off-chip. Transistor M_3's gate and source provide random access pixel addressing at the periphery, as needed to implement the stochastic kernel. For this pixel design, (16) establishes the following current transfer function in the subthreshold regime:

$$I_{\rm out} = c(I_{\rm in})^{1-\lambda}\exp\left(\frac{-\kappa\lambda Q}{C_{EL}V_T}\right)\exp\left(\frac{-\kappa\lambda V_{EL}}{V_T}\right) \tag{17}$$

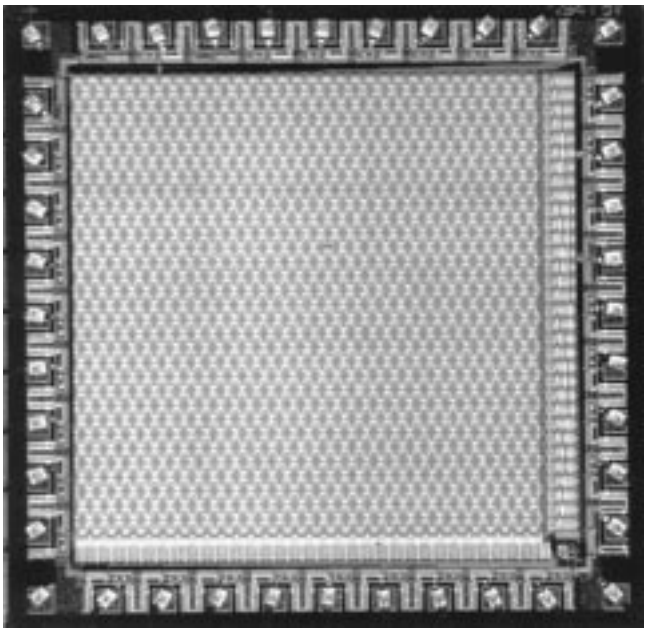

Fig. 5. Micrograph of 64 × 64 pixel adaptive imager chip. Dimensions are 2.2 mm × 2.25 mm in 1.2 μm CMOS technology.

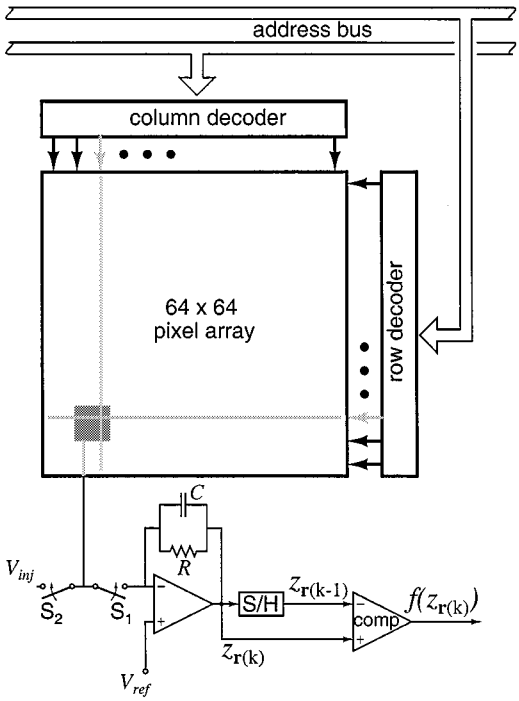

Fig. 6. Chip architecture and system setup for gain mismatch correction and intensity histogram equalization.

where

c	$= (I_0 W/L\exp(V_{dd}/V_T))^{\lambda}$;
I_0	subthreshold leakage current;
W and L	width and length of transistors M_1 and M_2;
V_{dd}	supply voltage;
V_T	thermal voltage;
κ	subthreshold slope factor (back gate coefficient).

The first exponential factor on the right in (17) corresponds to the adaptive gain correction A, while the second exponential factor represents normalization which is globally controlled

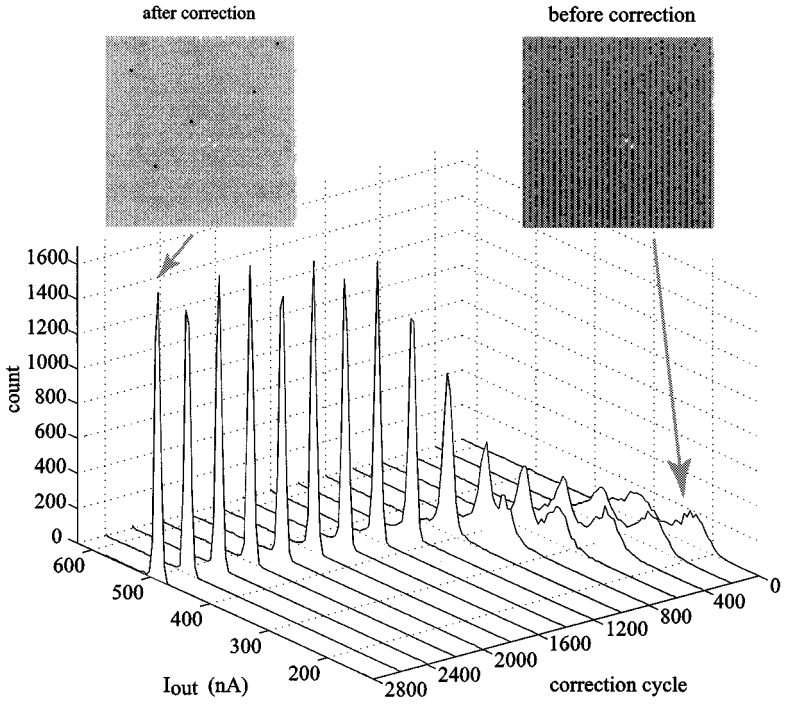

Fig. 7. Time course of gain nonuniformity reduction, as recorded from the adaptive imager chip. Also shown are acquired images before and after gain correction performed under conditions of uniform illumination.

by V_{EL}. By injecting electrons onto (tunneling electrons from) the floating gate [11] we incrementally (decrementally) alter Q, which in turn logarithmically modulates A, and thereby effectively implements the pilot rule (12).

Fig. 4 illustrates the effect of the various contributions to the pixel current transfer function $I_{\rm out}$ versus $I_{\rm in}$, through the floating gate voltage V_{fg} as given by (16). Capacitive division between $C_{\rm in}$ and C_{EL} reduces the voltage swing on the floating gate V_{fg} by a factor $1-\lambda$ relative to the input voltage $V_{\rm in}$. Through the logarithm V-to-I transformation across the MOS transistor in subthreshold, this compresses the dynamic range of intensities in the output image

$$I_{\rm out} = c'(I_{\rm in})^{\gamma} \tag{18}$$

by a factor $\gamma = 1-\lambda$, as shown in Fig. 4(a). Hot electron injection onto the floating gate modulates the charge Q, and thereby corrects the (relative) gain in each pixel individually by correspondingly lowering the floating gate voltage V_{fg}. The electrode voltage V_{EL} allows for a global shift of V_{fg} for all pixels, in either a positive or negative direction as shown in Fig. 4(a) and (b). The effect either way is a global, electronically adjustable scale factor in the gain, which allows for automatic gain control. For lower values of V_{EL}, which bring transistor M_2 above threshold as indicated in Fig. 4(b), a smaller compression factor γ is obtained in the current transfer function (18), although this factor γ then depends on the signal. If the image is subsequently histogram equalized through the oversampled binary encoding (14), the nonlinearity in the transfer function (18) becomes irrelevant.

V. VLSI System Architecture

An array of 64×64 adaptive pixels is integrated with x and y random-access addressing decoders onto a 2.2-mm $\times$ 2.25-mm chip in 1.2-μm CMOS technology. The micrograph of the prototype fabricated through MOSIS is shown in Fig. 5.

Fig. 6 illustrates the architecture of the chip and the setup used to experimentally validate the concept of reducing the gain mismatch between pixels on the prototype adaptive array. We uniformly illuminate the imager and randomly select a column and row address $\mathbf{r}(k)$. With switch S_1 closed and S_2 open, we measure $I_{\rm out(\mathbf{r}(k))}$ using a transimpedance amplifier to generate a voltage $z_{\mathbf{r}(k)}$. If $f(z_{\mathbf{r}(k)}) = 0$, we open S_1 and momentarily close S_2. The drain of transistor M_2 is pulsed down to $V_{\rm inj} \approx (V_{dd}-8\,{\rm V})$ and a small packet of negative charge is injected onto the floating gate. If $f(z_{\mathbf{r}(k)}) = 1$, we do not alter the gain of the selected pixel and continue by randomly selecting a new pixel. As such we implement a *one-sided* version of the stochastic pilot rule of (12)

$$\Delta q_{\mathbf{r}(k)} = \begin{cases} \alpha, & f\left(z_{\mathbf{r}(k)}\right) < 0; \\ 0, & \text{otherwise.} \end{cases}$$

Because adaptation is active in only one direction, the average level $\langle z \rangle$ drifts in that direction over time. We can use the coupling electrode voltage V_{EL} to compensate for this drift.

After gain nonuniformity correction, the imager can be used to acquire static natural images. Also, using random addresses with prescribed statistics, the output bit from the comparator $f(z_{\mathbf{r}(k)})$ can be accumulated in bins whose addresses are defined by $\mathbf{r}(k)$. The resulting histogram then represents the intensity equalized acquired image.

VI. Experimental Results

The 64×64 phototransistor-based imager was uniformly illuminated using a white light source. We scanned the pixel array before any gain mismatch correction, and again after every 200 cycles of correction, until we judged correction to be completed after 2800 cycles. Each of the 4096 pixels was selected in random sequence every cycle. Fig. 7 shows the evolution of the histograms built from I_{out} recorded from each pixel on the focal-plane versus adaptation cycle number. Also shown are the scanned images from the chip before and after gain mismatch correction.

We measured and plotted the standard deviation of I_{out} ($\sigma_{I_{\text{out}}}$) normalized to the mean $\langle I_{\text{out}} \rangle$, versus $\langle I_{\text{out}} \rangle$, before and after gain mismatch correction. The five different $\langle I_{\text{out}} \rangle$ correspond to five different levels of illumination which we label 1, 2, 3, 4, and 5. Adaptation was done at illumination level with label 5. Fig. 8 plots these experimental results.

We projected a black and white 35-mm slide onto the imager after gain mismatch correction and scanned the array. The slide contained a light-grey character "R" against a dark-grey background (both bitmapped). The resulting image, as scanned from the imager chip, is shown in Fig. 9.

We also projected a 35-mm grayscale (pixelated 64×64) slide "eye" image, shown in Fig. 10(a), onto the imager. The acquired image is shown in Fig. 10(b), and the histogram equalized image obtained from 256-times oversampled binary coding of the chip output is shown in Fig. 10(c).

VII. Discussion

Injecting a negative packet of charge onto the floating gate of transistor M_2 lowers its gate voltage and therefore increases its output current. Consequently, correction is in one direction only, increasing the current gain. Since the efficiency of charge injection depends exponentially on the magnitude of drain-source current through the device [11], pixels having higher I_{out} will inject more each time their drains are pulled to V_{inj}. This positive feedback mechanism can be kept in check either by driving the common drain with a current rather than voltage source, or by appropriately setting V_{inj}, keeping S_2 closed for a fixed time interval ($\approx 100\ \mu$s), and having hysteresis in the comparator which computes $f(z_{\mathbf{r}(k)})$. We choose the latter option for simplicity of the test setup.

The scanned image before correction in Fig. 7 shows strong vertical striations in I_{out}. After the gain mismatch correction procedure, these striations are no longer visible as evidenced by the post-correction image. We see five dark pixels (low I_{out}) in this image. These pixels are "stuck" off and therefore experience negligible injection when they are selected. Ideally, after correction we should expect to see an impulse in the histogram, all pixels having the same I_{out} when uniformly illuminated. In reality we see a single narrow peak in the histogram due to injection efficiency being proportional to current and due to hysteresis in the comparator.

Fig. 8 demonstrates that we did in fact reduce gain mismatch and not just $\sigma_{I_{\text{out}}}/\langle I_{\text{out}} \rangle$ as a consequence of increasing $\langle I_{\text{out}} \rangle$ [13]. The pre- and post-correction data lie on two separate curves demonstrating that there is indeed a dramatic reduction in gain mismatch due to adaptation. At low $\langle I_{\text{out}} \rangle$ (i.e., low illumination) there is a reduction in $\sigma_{I_{\text{out}}}/\langle I_{\text{out}} \rangle$ from 70 to

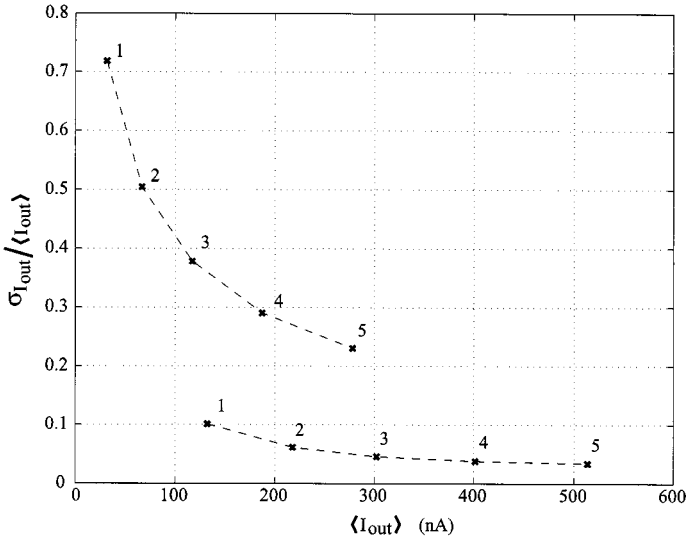

Fig. 8. Experimental pre- and post-corrected $\sigma_{I_{\text{out}}}/\langle I_{\text{out}} \rangle$ versus I_{out} for five different illumination intensities: (*top curve*) before gain correction and (*bottom curve*) after gain correction.

Fig. 9. Example image acquired from the adaptive imager chip after gain mismatch reduction. We projected a light-grey letter "R" against a dark-grey background onto the chip.

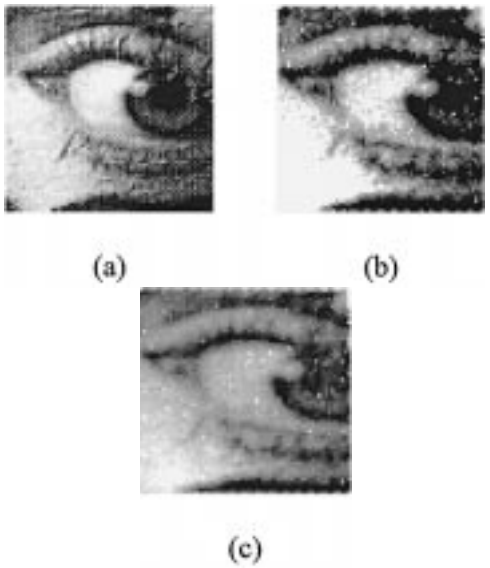

Fig. 10. (a) Original "eye" image; (b) image acquired from the chip; and (c) equalized image obtained from oversampled binary coding of the chip outputs.

10%. At higher $\langle I_{\text{out}} \rangle$ (i.e., high illumination), the reduction is from 24 to 4%.

The scanned image of an "R" after adaptation shown in Fig. 9 gives a clear image mostly free of gradients, striations, and other fixed pattern noise present in the imager before compensation. The remaining "salt and pepper" noise (two pixels in Fig. 9) is an artifact of the inhomogeneous adaptation rates under voltage-controlled hot electron injection in the setup of Fig. 6, which can be alleviated by using a current-controlled setup instead. The "eye" image shown in Fig. 10(c) reveals more (intensity) detail, especially around the iris, in the image after intensity equalization than the acquired image in Fig. 10(b).

VIII. Conclusions

We have introduced a compact pixel design and a strategy for reducing gain mismatch inherent in arrays of phototransistors used in CMOS imagers. We have shown how the learning rule for offset correction can be transformed into the log domain to produce a stable learning rule for online gain mismatch correction. This rule is very naturally implemented by a simple translinear circuit. The pixel incorporates a floating gate transistor which can be incrementally injected with a small packet of negative charge. The injected charge increases the current gain of the pixel in relative terms (i.e., by constant increments on a logarithmic scale).

Experimental results from a custom 64×64 phototransistor-based adaptive pixel CMOS array, fabricated through MOSIS, prove that our pixel design and learning rule were successful for static nonuniformity correction. In addition, we demonstrated intensity histogram equalization and digital coding of the output image, in a binary oversampled representation, by means of the same random-addressing stochastic algorithms and architecture as used for the adaptation.

Acknowledgment

P. Hasler joined the authors with the initial testing of the chip, at the 1998 NSF Telluride Workshop on Neuromorphic Engineering. K. Boahen provided feedback on an early draft of this paper on histogram equalization. The chip was fabricated through the MOSIS service.

References

[1] C. A. Mead, "Adaptive retina," in *Analog VLSI Implementations of Neural Systems*, Mead and Ismail, Eds. Norwell, MA: Kluwer, 1989.

[2] P. O. Pouliquen, A. G. Andreou, C. W. Terrill, and G. Cauwenberghs, "Learning to compensate for sensor variability at the focal plane," in *Proc. Int. Conf. Neural Networks (IJCNN)*, Washington, DC, July 1999.

[3] Y. Chiang and J. G. Harris, "An analog integrated circuit for continuous-time gain and offset calibration of sensor arrays," *Analog Integrated Circuits Signal Process.*, vol. 12, no. 13, pp. 231–238, 1997.

[4] A. Pesavento, T. Horiuchi, C. Diorio, and C. Koch, "Adaptation of current signals with floating-gate circuits," in *Proc. 17th Int. Conf. Microelectronics for Neural, Fuzzy and Bio-Inspired Systems*, 1999, pp. 128–134.

[5] A. Aslam-Siddiqi, W. Brockherde, M. Schanz, and B. J. Hosticka, "A 128-pixel CMOS image sensor with integrated analog nonvolatile memory," *IEEE J. Solid-State Circuits*, vol. 33, pp. 1497–1501, Oct. 1998.

[6] D. A. Scribner, K. A. Sarkady, M. R. Kruer, J. T. Caulfield, J. D. Hunt, M. Colbert, and M. Descour, "Adaptive retina-like preprocessing for imaging detector arrays," in *Proc. Int. Conf. Neural Networks (ICNN)*. Boulder, CO: IEEE, Feb. 1993, pp. 1955–1960.

[7] S. B. Laughlin and D. Osorio, "Mechanisms for neural signal enhancement in the blowfly compound eye," *J. Experimental Biology*, vol. 144, pp. 113–146, 1989.

[8] J. Nakamura, B. Pain, T. Nomoto, T. Nakamura, and E. R. Fossum, "On-focal-plane signal processing for current-mode active pixel sensors," *IEEE Trans. Electron. Devices*, vol. 44, pp. 1747–1758, 1997.

[9] K. A. Boahen, "A retinomorphic vision system," *IEEE Micro.*, vol. 16, pp. 30–39, 1996.

[10] Y. Ni, F. Devos, M. Boujrad, and J. H. Guan, "Histogram-equalization-based adaptive image sensor for real-time vision," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1027–1036, 1997.

[11] P. Hasler, B. A. Minch, J. Dugger, and C. Diorio, "Adaptive circuits and synapses using pFET floating-gate devices," in *Learning on Silicon*, G. Cauwenberghs and M. Bayoumi, Eds. Norwell, MA: Kluwer, 1999.

[12] H. Miwa, K. Yang, P. Pouliquen, N. Kumar, and A. Andreou, "Storage enhancement techniques for digital memory based, analog computational engines," in *IEEE Int. Symp. Circuits Systems*, vol. 5, 1994, pp. 45–49.

[13] A. Pavasovic, A. G. Andreou, and C. R. Westgate, "Characterization of subthreshold MOS mismatch in transistors for VLSI systems," in *J. VLSI Signal Process.*, 1994, vol. 8, pp. 75–85.

Marc Cohen (S'95) received the Bachelor of Science degree in electrical engineering from the University of the Witwatersrand, South Africa, in 1978 and a Masters of Science degree in electrical engineering from the same university in 1983. In 1991 he received the Masters of Science degree in biomedical engineering from the Johns Hopkins University, Baltimore, MD. In 1995 he began the Ph.D. program in electrical and computer engineering at the same university.

In 1983 he worked as a Research Engineer for the South African Chamber of Mines where he formulated mathematical models of human thermoregulation in the heat. In 1983 he joined the Research Laboratory of Electronics, Speech Communication Group at the Massachusetts Institute of Technology where he developed instrumentation for measuring the movement of tongue, lips, and jaw during speech production. He was a Fogherty Visiting Scientist at the National Institutes of Health, National Eye Institute from 1991 to 1993, where he developed hardware for extracellular action potential recordings. From 1993 to 1995 he worked as a Principal Analog Design Engineer for Cardiac Pacemakers, Inc.

Gert Cauwenberghs (S'89-M'92) received the Engineer's degree in applied physics from the University of Brussels, Belgium, in 1988, and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology in 1989 and 1994, respectively.

In 1994, he joined Johns Hopkins University where he is now Associate Professor of Electrical and Computer Engineering. From 1998 to 1999 he was on sabbatical as Visiting Professor of Brain and Cognitive Science at the Center for Computational and Biological Learning, Massachusetts Institute of Technology, and at the Center for Adaptive Systems, Boston University. His research interests include VLSI circuits, systems and algorithms for parallel signal processing, adaptive neural computation, and low-power coding and instrumentation.

Dr. Cauwenberghs is Associate Editor of the IEEE Transactions on Circuits and Systems—II: Analog and Digital Signal Processing, and Vice Chair of the IEEE Circuits and Systems Society Technical Committee on Analog Signal Processing. He has organized special sessions at conferences and journal special issues on learning, adaptation and memory, and recently co-edited a book *Learning on Silicon* (Kluwer, 1999). He was Francqui Fellow of the Belgian American Educational Foundation in 1988 and received the National Science Foundation Career Award in 1997, the Office of Naval Research Young Investigator Award in 1999, and the Presidential Early Career Award for Scientists and Engineers (Pecase) in 2000.