

A CMOS Programmable Analog Memory-Cell Array Using Floating-Gate Circuits

Reid R. Harrison, Julian A. Bragg, *Student Member, IEEE*, Paul Hasler, *Member, IEEE*,
Bradley A. Minch, *Member, IEEE*, and Stephen P. Deweerth, *Member, IEEE*

Abstract—The complexity of analog VLSI systems is often limited by the number of pins on a chip rather than by the die area. Currently, many analog parameters and biases are stored off-chip. Moving parameter storage on-chip could save pins and allow us to create complex programmable analog systems. In this paper, we present a design for an on-chip nonvolatile analog memory cell that can be configured in addressable arrays and programmed easily. We use floating-gate MOS transistors to store charge, and we use the processes of tunneling and hot-electron injection to program values. We have fabricated two versions of this design: one with an nFET injection mechanism and one with a pFET injection mechanism. With these designs, we achieve greater than 13-bit output precision with a 39-dB power-supply rejection ratio and no crosstalk between memory cells.

I. DEFINITION OF AN E-POT

MODERN analog, neuromorphic, or mixed-mode VLSI chips typically have large numbers of inputs and analog parameters, and the number of available pins is often a limiting factor in these systems. Often, these parameters and other biases are stored off-chip as voltages programmed by potentiometers or other sources, one pin per variable (see Fig. 1). Pins are also needed for inputs and outputs. The number of pins available on a chip is limited by the perimeter of the chip, which grows as the square root of die area. If we could move the analog parameters and circuit biases onto the chip, we would save more pins for input, output, and diagnostics (see Fig. 1). Also, storing biases on-chip could significantly reduce circuit board complexity.

If we think of potentiometers as easily modifiable voltage sources, then we would like to build an electronic version of the pot, an *e-pot*, that can be placed on the chip itself. If large numbers of these cells can be placed on one chip, then they will supply more biases than could the few pins required to program them. We have the following requirements for our e-pot.

- 1) *Nonvolatile*: An e-pot must maintain its state indefinitely in the absence of power.

Manuscript received April 2000; revised November 2000. This work was supported by the Center for Neuromorphic Systems Engineering under the National Science Foundation's Engineering Research Center Program. The work of R. R. Harrison was supported by DARPA and NIMH under grants. This paper was recommended by Associate Editor T. S. Lande.

R. R. Harrison is with the Electrical Engineering Department, University of Utah, Salt Lake City, UT 84112 USA (e-mail: harrison@ee.utah.edu).

J. A. Bragg, P. Hasler, and S. P. Deweerth are with the Department of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA (e-mail: phasler@ee.gatech.edu).

B. A. Minch is with the School of Electrical Engineering, Cornell University, Ithaca, NY 14853 USA.

Publisher Item Identifier S 1057-7130(01)02012-2.




Fig. 1. Using electronic potentiometers (e-pots). (a) On typical analog VLSI chips, many of the pins are consumed by bias voltages set with off-chip potentiometers. (b) By storing these voltages on-chip, many pins are freed for I/O.

- 2) *Small*: We would like to get reasonable numbers of these devices on one chip in order to have many “virtual pins.”
- 3) *Few Pins*: Because our goal is to save pins, a large array of e-pots should require few pins for I/O and biases.
- 4) *Wide Voltage Range*: An e-pot should be able to store any voltage from ground to the power-supply rail.
- 5) *“Tweakable”*: We must be able to move an e-pot’s voltage up or down smoothly, just like turning the knob on a potentiometer.
- 6) *Individually Addressable*: We must have control over each e-pot, like an array of real potentiometers (i.e., no global erase necessary).
- 7) *Flexible Supply Voltage*: The e-pot array should function using a wide range of supply voltages and only require additional voltage sources during the programming process.

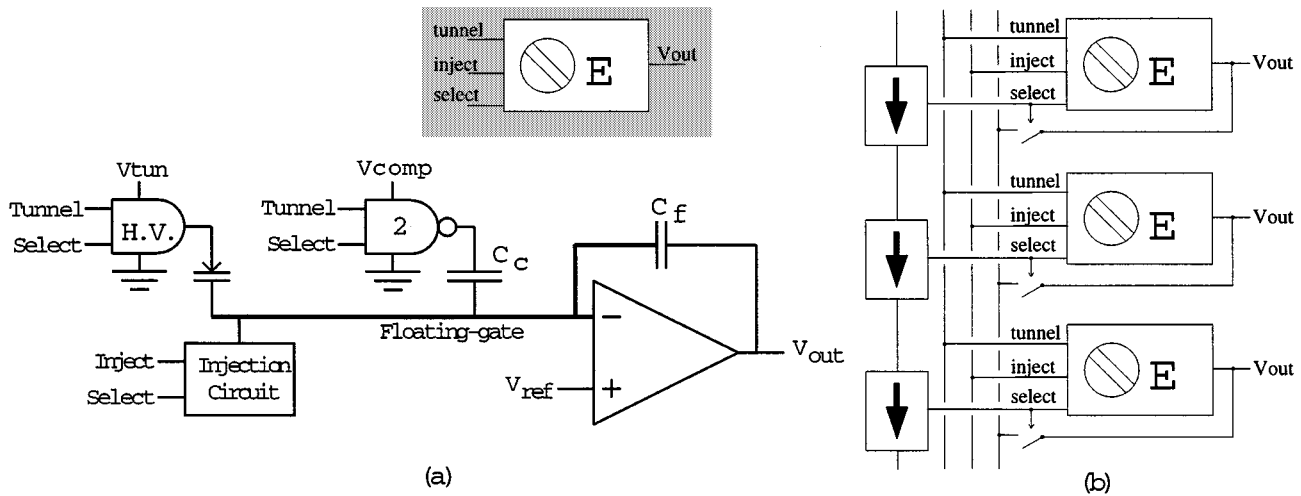


Fig. 2. E-pot circuit schematic. (a) Single e-pot cell. The floating gate is connected to the negative input of an amplifier with feedback capacitor C_f . This pins the floating-gate voltage to V_{ref} and allows V_{out} to be moved from rail to rail by changing the charge of the floating gate Q_{fg} . Charge is removed from the floating gate through tunneling. The tunneling voltage is switched with a high-voltage differential amplifier built with lightly doped-drain nFETs. The capacitive coupling of the tunneling voltage to the floating gate is counterbalanced by switching a lower voltage V_{comp} on a larger capacitor. Charge is added to the floating gate through hot-electron injection. Gate 2 is a pseudo-pMOS NAND gate; we use these symbols for clarity. The high-voltage amplifier is shown in Fig. 5. The amplifier is a ten-transistor nFET-input wide-output-range transconductance amplifier [10]. (b) Array of e-pots with shift register used for addressing. The tunnel, inject, and select lines carry digital signals.

8) *High Programming Precision*: More precision is always desirable.

Recent advances in floating-gate CMOS circuits open up the possibility for building on-chip nonvolatile e-pots [1]–[6]. In this paper, we describe a floating-gate CMOS e-pot array that meets the requirements listed above. We have described preliminary testing of these structures in [2].

An earlier floating-gate memory cell was built by Diorio *et al.* [5]. Our e-pot design has many advantages over this design. By using a high-gain differential amplifier in a feedback configuration, we have explicit control of the floating-gate voltage. In addition, the output amplifier has nFET inputs, preventing any leakage resulting from pFET hot-electron injection. Our memory cells are configured as an addressable array with no crosstalk, and high-voltage switches in each e-pot make a global erase unnecessary. Programming has been further simplified by cancelling the capacitive coupling inherent in tunneling, allowing the output voltage to be monitored during the programming process. Finally, our cells can be designed to use either an nFET-injector or a pFET-injector structure, allowing fabrication in standard CMOS processes.

II. CIRCUIT DESCRIPTION

The circuit schematic is shown in Fig. 2. We designed and fabricated arrays of e-pots in commercially available 2.0- μm BiCMOS and 1.2- μm and 0.8- μm CMOS processes.

We configured the e-pot circuits in a one-dimensional (1-D) array along the edge of each 2.2 mm \times 2.2 mm chip. Addressing circuitry was included in the arrays [see Fig. 2(b)]. In the 2.0- μm process, each e-pot measures 69 μm \times 333 μm , allowing 23 elements to be placed in the array. The 1.2- μm process reduced the element size to 41.4 μm \times 231.6 μm , allowing us to place 39 elements in the array. An e-pot array of arbitrary length consumes 11 pins for control signals and other biases. Once programming is completed, the array needs only three off-chip biases.

In order to make the e-pots nonvolatile, we use floating-gate MOS transistors [1]. A floating gate is a polysilicon node surrounded by SiO_2 , trapping charge on the gate indefinitely. The floating gate is connected to the negative input of a high-gain amplifier with a feedback capacitor C_f . This effectively pins the floating-gate voltage at V_{ref} , the positive input to the amplifier.

A. Controlling the Array

The e-pot elements are arranged in a 1-D array, with only one of those elements being “active” at a given time. It is important to note that e-pots are still sourcing voltage into the chip when they are not active. The first control signal is a clock that advances the shift register depicted in Fig. 2(b), causing the next e-pot in the array to become active. The V_{out} pin presents the output voltage of the active e-pot, while the sync output presents the logic high signal when, after stepping through the entire array, the shift register rolls over and the first e-pot becomes active again.

The tunnel and inject input signals control charge flow onto and off of the floating node, which in turn controls the output voltage of the amplifier. Tunneling, by removing electrons from the floating gate, increases the floating-gate voltage and reduces the output voltage. Injection, by placing electrons onto the floating gate, reduces the floating-gate voltage and increases the output voltage.

Fig. 3 shows a typical tunnel/inject cycle for an e-pot fabricated in a 2.0- μm BiCMOS process; e-pots fabricated in the 1.2- μm standard CMOS process performed similarly. While programming individual e-pots, we measured no crosstalk to other e-pots due to tunneling or injection in either of the two processes. Fig. 4 illustrates the independent control available over the individual e-pots; in this example, the pots have been programmed in a cosine pattern. We have set e-pots to voltages from 10 mV above ground to the power-supply rail.




Fig. 3. Experimental measurement illustrating e-pot operation. Digital signals control tunneling and hot-electron injection, moving the e-pot output voltage up or down smoothly.

B. Electron Tunneling and Capacitive Compensation

To decrease an e-pot's output voltage, we remove electrons from its floating gate by the process of Fowler–Nordheim tunneling [9]. This process uses a high-voltage source to create an energy barrier thin enough that electrons can tunnel through the gate oxide; typical tunneling voltages used are 35–40 V in 2.0- μm processes and 27–30 V in 1.2- μm processes (based on other measurements, we expect tunneling voltages in the range of 10–12 V in 0.5- μm processes).

We switch high voltages on-chip with a high-voltage differential amplifier built with lightly doped-drain nFETs, which use well diffusion as their drain regions (see Fig. 5). The resulting high-voltage nFETs have breakdown voltages greater than 45 V. In the CMOS process we use, it is not possible to create high-voltage pFETs, so we simply let the pFETs in the diffamp break down. This happens at a V_{ds} of 20 V, so our high-voltage diffamp cannot completely “turn off.” The output varies between around 40V in the “on” mode to around 15 V in the “off” mode. However, the tunneling current depends exponentially on the reciprocal of the oxide voltage, so at $V_{tun} = 15$ V, there is no observable tunneling current. It should be noted that after the e-pots are programmed, the tunneling voltage can be removed and there is no further need for a high-voltage source. By switching the tunneling and injection voltages locally, we achieve individually programmable memory cells, with no need for a global erase.

A major limitation of floating-gate memory cells that are programmed by tunneling results from parasitic capacitances coupling between the tunneling node and the floating gate. Although this capacitance is usually small compared to the total capacitance of the floating gate, the voltage change across this junction is large (typically 25–35 V) when the tunneling voltage is switched on or off. This voltage change couples into the floating gate, causing a large voltage offset in the memory-cell output. Since capacitors connected to floating gates are dc elements, this voltage offset does not decay but




Fig. 4. Output voltages from all 39 e-pots from our 1.2- μm chip, after each element in the array had been programmed to a voltage proportional to the cosine of the e-pot position number. No crosstalk was observed between e-pot elements during programming.

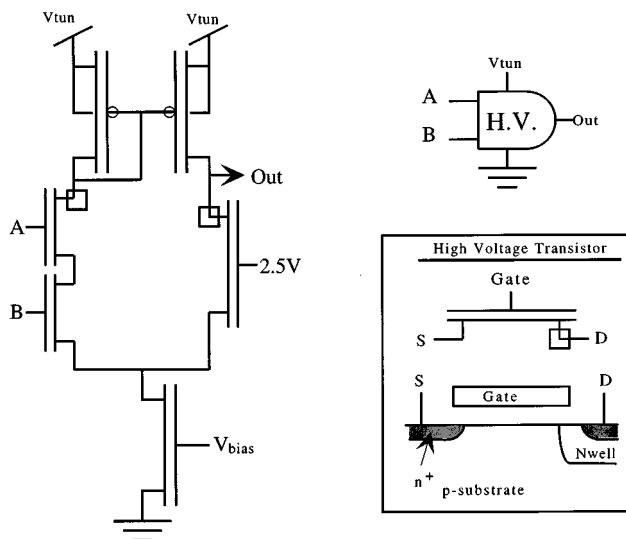


Fig. 5. High-voltage AND gate for connecting and disconnecting high voltages to tunneling junctions. Because the required tunneling voltages are often higher than the breakdown voltage of the source-drain regions to the substrate, we need to use high-voltage transistors. A high-voltage transistor is formed by using a drain region made from a lightly doped n-well. These transistors should be made sufficiently long to eliminate punchthrough effects.

rather is present for the duration of tunneling. This makes for indirect and awkward programming.

In order to make our memory cells “tweakable,” we compensate for this effect by switching a lower voltage on a larger capacitor in the opposite direction. When we switch the tunneling junction to a high voltage, we simultaneously switch another node from a positive “compensation voltage” to ground. This node is coupled to the floating gate with a capacitance that is many times larger than the tunneling junction capacitance. The compensation voltage V_{comp} is set by an off-chip bias, so the capacitive coupling can be precisely nulled by the user (see Fig. 6). This approach is a practical solution for e-pot biases for a wide range of CMOS processes.




Fig. 6. Illustration of compensating for capacitive coupling during tunneling. Normally, capacitive coupling through the tunneling junction produces large offsets in V_{out} during tunneling. By tuning V_{comp} to the correct value, we can compensate for this effect.

C. Hot-Electron Injection

To increase the e-pot output voltage, we can add electrons to the floating gate by the process of hot-electron injection. There are two methods of performing hot-electron injection: one using a pFET as the injecting device and one using an nFET (see Fig. 7). While any fabrication process can be used to create a pFET injector, nFET-injector structures must be fabricated in a BiCMOS process. Additionally, modern fabrication processes make use of structures such as lightly doped drains and channel spacers that, while making it possible for normal CMOS circuits to function at submicrometer feature sizes, prevent nFET-injector structures from operating properly. These structures have no effect on pFET injector structures, which have been fabricated in processes with feature sizes as small as $0.5 \mu\text{m}$.

To perform pFET injection, we place a large voltage (6 V) across the source and drain of a normal pFET, while holding the gate voltage slightly below that of the source. Holes, after they have diffused across the channel, are exposed to high electric fields at the drain edge of the drain-to-channel depletion region and lose large amounts of potential energy as they drop to the drain voltage. Some of these holes strike Si atoms in the semiconductor lattice, generating an electron-hole pair via hot-hole impact ionization. These electrons travel back into the channel region, gaining energy as they go. When their kinetic energy exceeds the 3.1-eV silicon-silicon-dioxide energy barrier, the “hot” electrons can be injected into the oxide and transported to the floating gate [4].

To perform nFET injection, we use a p-type base implant to create an nFET with a high threshold voltage (around 6 V). This implant is only available in BiCMOS processes, which is why we are restricted to those processes if we wish to use nFET-injector structures. We can operate this device in subthreshold with a gate voltage of 5.5 V and inject electrons to the floating gate by raising the drain voltage to 4 V. Electrons approaching the drain gather energy from the strong electric field, and some gain enough energy to surmount the 3.1-eV energy barrier im-

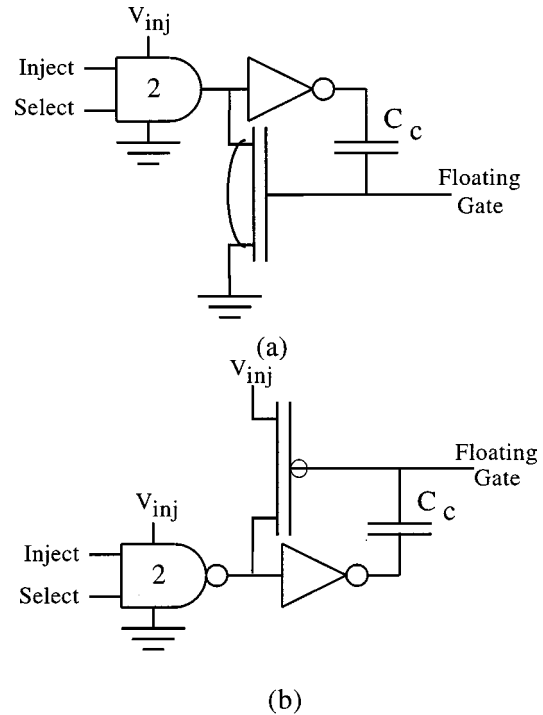


Fig. 7. Two methods for hot-electron injection programming. In both cases, an inverter is used in the compensation circuitry to offset the capacitive coupling into the floating gate through the drain-to-channel capacitor. Since we operate the MOSFETs in subthreshold, the overlap gate-to-drain capacitance is fixed, and therefore the matching in compensation circuitry is limited to gate-capacitor matching. (a) An nFET injection mechanism. (b) A pFET injection mechanism.

posed by the gate oxide. These “hot” electrons are attracted to the floating gate, which is more positive than the drain. Because the floating gate must be around 5.5 V, we use a 6-V power supply for the e-pot and set the floating-gate voltage with the feedback circuit described.

In both methods of injection, switching the injection process on causes a step change in channel voltage, which, by capacitive coupling, causes an offset in the floating-gate voltage. Drain-to-gate overlap capacitance contributes additional coupling to the floating node. To counteract these effects, we tie the switching node to the input of an inverter, the output of which is also capacitively coupled to the floating gate. By adjusting the size of this coupling capacitor prior to fabrication, we can tune the system to reject injection offset for a specific combination of power-supply and injection voltages. The resulting compensation system works well within that narrowly defined voltage range, but as either voltage is changed, the offset voltage becomes significant. In principle, we could compensate for this offset in the same way that we deal with tunneling offset, but to do so, we would need to supply another off-chip bias.

III. CIRCUIT PERFORMANCE

In this section, we discuss performance issues of using these e-pot elements:

- 1) programming accuracy and precision;
- 2) output voltage noise;
- 3) output voltage power-supply rejection;
- 4) long-term drift of the output voltage.

In all cases, the 1.2- μm nFET-injector and 2.0- μm pFET-injector structures showed similar performance.

A. Programming Accuracy and Precision

In order to program voltages precisely, we used a comparator in a feedback loop with the e-pot output voltage, so that the tunneling or injection controls were triggered by the comparator output. The only limitation of this technique is the offset voltage of the comparator, typically less than 5 mV for the particular circuit that we used. In order to measure the accuracy of the memory cell, we programmed the same voltage multiple times and measured the error between the resulting voltage and the reference voltage. To quantify precision, we computed the standard deviation of the resulting voltages.

To measure the accuracy and precision of the e-pot output over a wide range of target voltages, we programmed a single pFET-injector e-pot to produce a series of 79 voltages, uniformly spaced between 0.6 and 4.4 V. At each step, we recorded both the target voltage and the actual output. Performing this procedure 64 times, we found remarkably little deviation from ideal performance across the tested range. Fig. 8 shows the relationship between target voltage and mean output voltage. A straight-line fit to these points shows a gain error of 0.05%.

We observed a 19.3-mV dc offset between the target voltage and the actual output, which can be attributed to imperfect compensation for the offset in the output voltage that takes place during injection. The injection compensation capacitor being used in the pFET-injector circuit was sized to compensate for the switching of an injection voltage roughly equal to the supply voltage for the rest of the chip. These measurements were taken with an injection voltage of 9 V and a supply voltage of 5 V—a much greater difference than originally intended. Further testing has shown that reducing the difference between the two voltages does, in fact, reduce the dc offset error. In future revisions, steps can be taken to reduce this error despite this voltage difference, ranging from an adjustable compensation bias (as is currently being used in the tunneling compensation circuitry) to a capacitor sized to compensate for whatever voltage difference is expected in the final application.

Fig. 9 shows how the observed output error varies as a function of the target voltage once the previously mentioned dc offset has been removed. The vertical bars show the standard deviation of the error over 64 trials, while the points show the minimum and maximum observed errors. There is a systematic increase in the error as the target voltage increases, due to the finite common-mode rejection ratio of the output amplifier. As the output voltage shifts in response to injection switching, the gain applied to that shift is, to some degree, a function of the output voltage.

Fig. 10 shows the distribution of output errors after multiple attempts to source a given target voltage (3.0 V) for e-pots using both nFET-injector and pFET-injector structures. It also shows the distribution of errors when attempting to tunnel to a target. The standard deviations of 450 μV for nFET injection and 470 μV for tunneling correspond to greater than 13-bit precision over the 6-V operating range of the e-pot. The pFET injection e-pot demonstrated a standard deviation of 175 μV over its 5-V operating range, giving an effective precision of 14.8 bit. We




Fig. 8. Accuracy of programming 79 different voltages to the same e-pot element. The response follows a nearly straight line; the gain error is roughly 0.05% across the observed output range.




Fig. 9. Deviation of the e-pot programming voltages from their target voltages. We clearly see the 0.05 gain error in this plot. There is a systematic increase in the error as the target voltage increases, due to offset dependence in our amplifier circuit, responding to the injection switching offset [11]. We show how the observed output error varies as a function of the target voltage. The vertical bars show the standard deviation of the error over 64 trials, while the points show the minimum and maximum observed errors. This deviation comes from noise in both the floating-gate amplifier and the instrumentation circuitry.

made no attempts to measure these data in a low-noise environment, so this number is merely a lower bound on e-pot precision.

To verify the ability of the e-pot structure to provide bias currents, as well as bias voltages, the output of one e-pot element was tied to a current source. Fig. 11 shows the output current as a function of the e-pot output voltage, displaying both the sub-threshold exponential I-V relationship and the quadratic I-V relationship characteristic of above-threshold operation.

B. E-Pot Output Noise

The total root mean square noise present on the e-pot-supplied output voltage was on the order of 1–2 mV. In order to




Fig. 10. (a) Histogram for programming with nFET hot-electron injection—100 trials. (b) Histogram for programming with pFET hot-electron injection—64 trials. (c) Histogram for programming with tunneling—100 trials. In all cases, the target voltage was 3.0 V. All three processes give a precision (V_{dcl} range divided by σ) of greater than 13 bit; pFET injection gives the highest precision, with nearly 15 bit.

determine the noise components present in the e-pot voltage, we measured the frequency spectrum of a pFET-injector e-pot output. The resulting spectrum, shown in Fig. 12, shows conventional $1/f$ noise, which dominates at low frequencies, as well as thermal device noise at higher frequencies. This noise could be reduced by using larger transistors in the output transconductance amplifier, at the cost of increasing the e-pot element size and thereby reducing the number of e-pots that could be placed on a single chip. These noise properties are similar to other floating-gate amplifiers [13], [12]. We showed elsewhere that the tunneling and injection processes do not contribute significantly to the noise levels [13].

We want to investigate how changing the e-pot design will change the amount of output noise. Following [14] and [12], we can model the thermal noise component \hat{i}_o of a subthreshold MOSFETs channel current by

$$\frac{\hat{i}_o^2}{\Delta f} = 2qI. \quad (1)$$

The source of noise from the e-pot comes directly from the noise generated from the wide-output range transconductance amplifier [10]; the e-pot output noise is due to the effect of several nFETs and pFETs. We use the result that the effective current noise of this transconductance amplifier is roughly 5.3 times




Fig. 11. Measurements from using an e-pot for programming a current source. This application shows that we can program through the subthreshold and above-threshold regimes.




Fig. 12. Noise spectrum from an on-chip e-pot cell. We see two types of noise: the conventional $1/f$ noise at low frequencies and the $1/f^{(1/2)}$ thermal noise characteristic of pFETs operating in weak or moderate inversion. We also plot lines showing curve fits to these regions of noise.

that of the bias current transistor [15]. We compute the output voltage noise from the input current noise and the amplifier's effective transconductance. The floating-gate voltage is related to the output voltage through a capacitive voltage divider. Typically, the parasitic capacitance at the floating gate is 20% of the feedback capacitance between the floating gate and output; therefore, the effective transconductance of the input nFETs is roughly 20% less than typical nFET devices. Using this reduced transconductance, we express the signal power of the output-referred voltage noise \hat{V}_{out}^2 as

$$\hat{V}_{\text{out}}^2 = \frac{5.3\hat{i}_o^2}{g_m^2(1 + (\omega\tau_h)^2)} \quad (2)$$

where τ_h is $(C_L + C_f)/g_m$. We can calculate directly the noise per unit bandwidth as

$$\frac{\hat{V}_{\text{out}}^2}{\Delta f} = \frac{10.6qI}{g_m^2(1 + (\omega\tau_h)^2)}. \quad (3)$$

From this expression, we can calculate the total output-noise power as

$$\hat{V}_{\text{out}}^2 = \frac{10.6qI}{g_m^2} \int_0^\infty \frac{1}{1 + (\omega\tau_h)^2} df \quad (4)$$

which evaluates to

$$\hat{V}_{\text{out}}^2 = \frac{10.6qI/g_m}{C_L + C_f}. \quad (5)$$

The lowest total output-noise power occurs at subthreshold current biases. The total output-noise power is inversely proportional to $C_L + C_f$; therefore, a larger capacitor can be used to decrease the total noise.

Now, we would like to calculate the dynamic range of an e-pot. We define dynamic range (DR) as the ratio of the maximum possible linear output swing to the total output-noise power. For a 5-V supply, we can assume conservatively that we have a 4-V linear range (V_{max}). With this definition, and assuming that the e-pot is biased in subthreshold $I/g_m = U_T/\kappa$, we can express the e-pot dynamic range as

$$\text{DR} = \frac{V_{\text{max}}^2}{2\hat{V}_{\text{out}}^2} = \frac{(1.51 \text{ V}^2)\kappa(C_L + C_f)}{qU_T} \quad (6)$$

which is similar to the form for dynamic range for the wide-linear-range amplifier. Using typical C_f capacitors of 2 pF at room temperature ($T = 300\text{K}$), we calculate the dynamic range as approximately equal to 84 dB (14 bit).

C. Power-Supply Rejection

Fig. 13 shows how the e-pot output voltage changes in response to variation in the power-supply voltage. As the supply voltage decreases, the output voltage increases linearly, with a power-supply rejection ratio (PSRR) of 39 dB. This is another advantage the e-pot structure has over a conventional bank of potentiometers, each of which has an average PSRR of 6 dB. Since the output transconductance amplifier is the only e-pot component with a direct connection to the supply voltage, we believe that these variations are due entirely to this amplifier. By optimizing device lengths, and thereby increasing the output amplifier's open-loop gain, the PSRR could be further increased, at the cost of increasing the size of each e-pot element.

D. Long-Term Drift of the Output Voltage

Once a voltage has been programmed, it is important that it remain stable for long periods of time. To test the e-pot's stability, we monitored the output voltage of a single e-pot over a 70-h period. The output voltage increased almost 20 mV in the first 36–40 h, after which it settled to a stable value. This initial increase appears to be due to electron detrapping, as electrons that had become trapped in the gate oxide during injection slowly worked their way out of the oxide and into the floating




Fig. 13. Deviations from ideal output voltage due to changes in the power-supply voltage. All e-pot bias voltages, which are ground referenced, were fixed for these measurements.

gate itself. After 40 h, the population of trapped electrons has been significantly reduced, and the output voltage stabilizes.

IV. CONCLUSIONS

We have built and tested two arrays of nonvolatile analog VLSI memory cells capable of greater than 13-bit precision. The analog values are stored as charge on a floating gate, which is modified through Fowler–Nordheim tunneling and hot-electron injection using either an nFET-injector or pFET-injector structure. The cells are individually addressable and are capable of sourcing precisely controlled voltages for long periods of time with very little noise or drift and a high degree of power supply noise rejection.

We have demonstrated the ability of these structures to provide bias voltages and bias currents. We intend to use them to allow the construction of complex analog and mixed-mode circuits without sacrificing an excessive number of pins for setting circuit biases. They will also greatly simplify board layout, as a large number of potentiometers can be replaced by a simple digital control system. These arrays have been laid out in a standard frame and fabricated in commercially available 2.0- μm BiCMOS and 1.2- μm and 0.8- μm CMOS processes.

REFERENCES

- [1] P. Hasler, B. A. Minch, and C. Diorio, "Adaptive circuits using pFET floating-gate devices," in *Proc. Advanced Research VLSI*, Atlanta, GA, 1999, pp. 215–231.
- [2] R. Harrison, P. Hasler, and B. A. Minch, "Floating-gate CMOS analog memory cell array," in *Proc. Int. Symp. Circuits and Systems*, Monterey, CA, 1998, see also <http://users.ece.gatech.edu/phasler>.
- [3] P. Hasler, C. Diorio, B. A. Minch, and C. Mead, "Single transistor learning synapses with long term storage," in *Proc. Int. Symp. Circuits and Systems*, vol. 3, Seattle, WA, 1995, pp. 1660–1663.
- [4] P. Hasler, B. A. Minch, C. Diorio, and C. Mead, "An autozeroing amplifier using pFET hot-electron injection," in *Proc. Int. Symp. Circuits and Systems*, vol. 3, Atlanta, GA, 1996, pp. 325–328.
- [5] C. Diorio, S. Mahajan, P. Hasler, B. A. Minch, and C. Mead, "A high resolution nonvolatile analog memory cell," in *Proc. Int. Symp. Circuits and Systems*, vol. 3, Seattle, WA, 1995, pp. 2233–2236.

- [6] C. Diorio, P. Hasler, B. A. Minch, and C. Mead, "A complementary pair of four-terminal silicon synapses," *Analog Integr. Circuits Signal Process.*, vol. 13, no. 1/2, pp. 153–166, 1997.
- [7] A. Thomsen and M. A. Brooke, "A floating gate MOSFET with tunneling injector fabricated using a standard double-polysilicon CMOS process," *IEEE Electron Device Lett.*, vol. 12, pp. 111–113, 1991.
- [8] H. V. Tran, T. Blyth, D. Sowards, L. Engh, B. S. Nataraj, T. Dunne, H. Wong, V. Se rin, T. Lam, H. Hazarian, and G. Hu, "A 2.5V 256-level nonvolatile analog storage device using EEPROM technology," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 1996, pp. 270–271.
- [9] M. Lenzlinger and E. H. Snow, "Fowler–Nordheim tunneling into thermally grown SiO₂," *J. Appl. Phys.*, vol. 40, no. 1, pp. 278–283, 1969.
- [10] C. Mead, *Analog VLSI and Neural Systems*. Reading, MA: Addison-Wesley, 1989.
- [11] Y. Tsividis and P. Antognetti, *Design of MOS VLSI Circuits for Telecommunications*. Englewood Cliffs, NJ: Prentice-Hall, 1985, pp. 145–170.
- [12] P. Hasler, "Foundations of learning in analog VLSI," Ph.D. dissertation, California Institute of Technology, Pasadena, CA, 1997.
- [13] P. Hasler, B. A. Minch, and C. Diorio, "An autozeroing floating-gate bandpass filter," in *Proc. IEEE Int. Symp. Circuits and Systems*, Monterey, CA, 1998.
- [14] R. Sarpeshkar, T. Delbruck, and C. A. Mead, "White noise in MOS transistors and resistors," *IEEE Circuits Devices*, pp. 23–29, Nov. 1993.
- [15] R. Sarpeshkar, R. F. Lyon, and C. Mead, "A low-power wide-linear-range transconductance amplifier," *Analog Integr. Circuits Signal Process.*, Nov. 1996.

Reid R. Harrison received the B.S. degree in electrical engineering from the University of Florida, Gainesville, in 1994 and the Ph.D. degree from the California Institute of Technology, Pasadena, in 2000.

After working for a short time at the Jet Propulsion Laboratory and Los Alamos National Laboratory, he joined the Computation and Neural Systems at the California Institute of Technology. He is currently an Assistant Professor of Electrical Engineering at the University of Utah, Salt Lake City. His research interests include silicon implementations of biological sensory systems, low-power analog VLSI circuits, and neuroprosthetics.

Julian A. Bragg (S'00) received the B.S. degree in electrical engineering from The George Washington University, Washington, DC, in 1995. He is currently pursuing the M.D. and Ph.D. degrees in bioengineering from the Georgia Institute of Technology and Emory University School of Medicine, Atlanta.

His research interests include neuroscience, muscle physiology, and biological motor control.

Mr. Bragg is a member of Tau Beta Pi and Eta Kappa Nu.



Paul Hasler (S'87–A'97–M'01) received the B.S.E. and M.S. degrees in electrical engineering from Arizona State University, Tempe, in 1991 and the Ph.D. degree in computation and neural systems from the California Institute of Technology, Pasadena.

He is currently an Assistant Professor in the Department of Electrical and Computer Engineering at the Georgia Institute of Technology. His research interests include low-power electronics; mixed-signal integrated circuits and systems; the use of floating-gate MOS transistors to build adaptive information processing systems and "smart" sensor interfaces; the physics of deep submicrometer devices or floating-gate MOS devices; and analog VLSI models of neurobiological learning and sensory information processing.

Dr. Hasler received the IEEE Electron Devices Society's Paul Rappaport Award in 1996. He is active in the IEEE as a Cochair of the Atlanta section of the IEEE Electron Devices Society, as a Reviewer for IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS and IEEE TRANSACTIONS ON NEURAL NETWORKS, and as Cochair for special sessions in the IEEE International Symposium on Circuits and Systems in both 1998 and 1999.



Bradley A. Minch (S'90–M'91) received the B.S. degree in electrical engineering (with distinction) from Cornell University, Ithaca, NY, in 1991 and the Ph.D. degree in computation and neural systems from the California Institute of Technology, Pasadena, in 1997.

He is currently an Assistant Professor in the School of Electrical and Computer Engineering, Cornell University. His research interests include the analog and digital integrated circuit design, translinear circuits, log-domain filters, and adaptive floating-gate MOS circuits.

Dr. Minch is a member of Tau Beta Pi, Eta Kappa Nu, and Phi Kappa Phi. He received the IEEE Electron Devices Society's Paul Rappaport Award in 1996.

Stephen P. Deweerth (S'85–M'91) received the B.A. degree in mathematics and chemistry from Wartburg College, Waverly, IA, in 1985. He received the M.S. degree in computer science and the Ph.D. degree in computation and neural systems from the California Institute of Technology, Pasadena, in 1987 and 1991, respectively.

Since 1991, he has been with the Georgia Institute of Technology, Atlanta, where he is currently an Associate Professor of electrical and computer engineering. His current research interests are in the area of highly collective computation performed by physical systems. He is primarily involved in the modeling of neurobiological systems that control movement.