A Micropower CMOS Algorithmic A/D/A Converter

Gert Cauwenberghs, *Member, IEEE*

Abstract— A low-power and compact VLSI architecture, implementing a bidirectional bit-serial A/D/A (analog-to-digital and digital-to-analog) converter, is presented. Both functions of algorithmic D/A conversion and successive approximation A/D conversion are combined into a single device, converting bits in the order from most to least significant. The MSB-first order guarantees robust implementation, relatively insensitive to component mismatches, offsets and nonlinearities. Also, since the A/D conversion makes use of the intermediate D/A conversion results, matched monotonic characteristics are obtained in both directions of conversion. The final D/A result is available at the end of A/D conversion, and can be used directly in applications calling for analog quantization. More general use of the A/D/A converter allows for bidirectional read/write digital access to local analog information in VLSI. The robust architecture supports dense integration of multiple low-power data conversion units along with digital processors or sensory circuitry in a standard CMOS process. Minimum sizing of active and passive devices in the implementation, to obtain optimal area and energy efficiency, is limited by clock feedthrough and finite gain considerations rather than matching requirements. Experimental results from a prototype VLSI implementation are given. Including control logic, the A/D/A cell measures 216 μm × 315 μm in a 2-μm CMOS process, and achieves 8-b untrimmed monotonicity at 200 μW power consumption for a 20 μs conversion cycle. This corresponds to 4 nJ of energy dissipated per 8-b converted sample.

I. INTRODUCTION

PRESENT trends in integrated data converter design [1], [2] typically target centralized architectures for signal processing, requiring high accuracy and high speed of conversion. At the other end of the spectrum, the data converter cell presented here addresses the need for more compact and energy efficient realizations, for large-scale parallel array processing and integrated applications in hybrid analog-digital technology. Applications range from the integration of in-situ conversion circuits in biomedical sensors and micro-mechanical actuators [3], to large-scale analog quantization [4] as a supporting function in analog VLSI parallel neural computation [5], [6].

Algorithmic data converters [7] offer an efficient implementation with typically less power dissipation and fewer circuit components than other conversion schemes for a given range of conversion bandwidth and resolution. The optimal efficiency is due to the tree-based structure of the conversion process, in which the number of conversion operations scales directly with the number of bits converted.

The author is with the Department of Electrical and Computer Engineering, The Johns Hopkins University, Baltimore, MD 21218 USA.
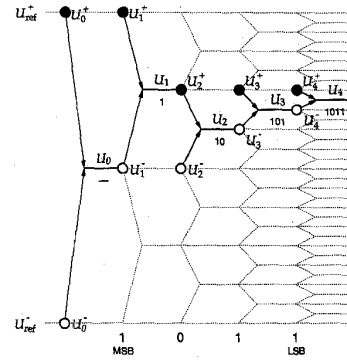
Fig. 1. D/A conversion algorithm.

Algorithmic converters are also more sensitive to device mismatches in the implementation. Nevertheless, advanced compensation schemes allow improved accuracy beyond the intrinsic precision provided by the fabrication process [8].

An MSB-first approach for D/A conversion [9] offers several advantages. First, it supports the same order in bit sequence as is necessary for successive approximation A/D conversion, and hence allows the integration of both A/D and D/A functions in a single bidirectional device, with matched transfer characteristics in both directions of conversion between analog and digital formats. The intermediate D/A conversion results serve directly as successive approximations for A/D conversion, and full D/A conversion for every bit in the successive approximation is no longer necessary. Second, the final result of D/A conversion is available "for free" at the end of conversion. This value corresponds to a quantized version of the analog input, and can be used, for instance, in the dynamic refresh of volatile analog storage [10]–[13]. Finally, the architecture implementing the MSB-first format for algorithmic D/A conversion, as it is reported here, is intrinsically less sensitive to device mismatches than traditional LSB-first approaches. This allows minimum-size implementation with corresponding additional power and area savings.

II. CONVERSION ALGORITHM

Fig. 1 schematically illustrates the MSB-first D/A conversion algorithm, in the case of a 4-b conversion for the input code "1011." The principle and formulation are similar to those presented in [9], with differences pertaining mainly to the implementation architecture.[1] The intermediate states U_i each correspond to the partial conversion of the i most

[1] In particular, the present architecture naturally provides for cancellation of offsets in the active elements, for robust operation under variations in the fabrication process, as described further below.

significant bits of the n-bit input code, and their values are constructed from previous intermediate values conditional on the state of the present bit, leading to the final conversion value in n successive steps. Two reference values, U_{ref}^{+} and U_{ref}^{-}, define the analog conversion range and roughly correspond to the upper and lower extremes of the analog spectrum, respectively. The algorithm employs two analog "registers" (short-term analog memories) U^{+} and U^{-} for storage and later recall of selected intermediate conversion values, and one unit for averaging the two values contained in both registers to construct the next intermediate conversion value. The following sequence of symbolic instructions specifies the algorithm

$a.$ Initialization ($i = 0$):
 Preset U_0^{+} to U_{ref}^{+} and U_0^{-} to U_{ref}^{-};
 Set U_0 halfway in between U_0^{+} and U_0^{-};
$b.$ Algorithmic iteration ($i - 1$ to i, from $i = 1$ to n):
 If bit i is = "1": set U_i^{+} to U_{i-1}^{+} and U_i^{-} to U_{i-1};
  "0": set U_i^{+} to U_{i-1} and U_i^{-} to U_{i-1}^{-};
 Set U_i halfway in between U_i^{+} and U_i^{-};
$c.$ Termination ($i = n$):
 The final D/A conversion result is given by U_n.

Because this algorithm implements a tree structure for the intermediate values which naturally preserves the order of the input codes, a monotonic conversion is guaranteed even in the case of a bias or nonlinearity in the averaging. Nevertheless, errors in the recall of the stored register values U_i^{+} and U_i^{-} distort the tree structure and affect the monotonicity directly.

The algorithm extends to successive approximation A/D conversion by successively comparing an analog input with the intermediate D/A conversion values, and if the input is larger assigning a "1" to the succeeding bit to be converted, and otherwise assigning a "0." The digital output is then given by the sequence of bits obtained, from MSB to LSB.

III. A/D/A CONVERTER BLOCK DIAGRAM

The block diagram of the analog portion of the bidirectional A/D/A converter, implementing the above algorithm with CMOS-C elements, is shown in Fig. 2. The analog circuit comprises four capacitors C^{+}, C^{-}, $C^{+'}$, and $C^{-'}$, two "bidirectional replication elements" BRE^{+} and BRE^{-}, one latched comparator with additional inverting transconductance output, and a series of switches. Switch SH performs the averaging function by sharing and redistributing charge on C^{+} and C^{-}. $C^{+'}$ and $C^{-'}$ represent the storage registers U^{+} and U^{-}, respectively. Both storage and recall functions are performed by the bidirectional replication elements BRE^{+} and BRE^{-}, which comprise differential transconductance amplifiers OTA^{+} and OTA^{-}, respectively, along with store switches STO^{+} and STO^{-}, and recall switches RCL^{+} and RCL^{-}. The OTA's and switches direct the storage (STO) and recall (RCL) operations by nondestructive replication of voltage from one capacitor terminal to the other, in either direction. To implement offset-free registers, a zero replication error for successive store and recall operation is required, which is satisfied if the input-referred offset voltages for the inverting and noninverting
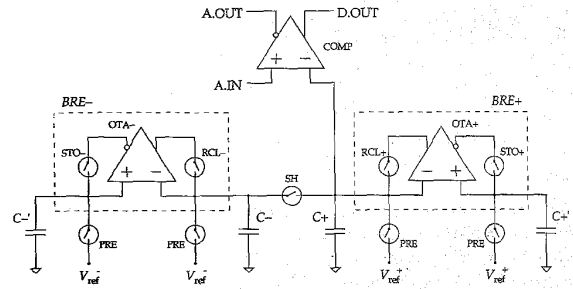

Fig. 2. A/D/A converter block diagram.

outputs of the corresponding OTA are equal. Comparator COMP serves to discriminate successive approximations in A/D conversion. The inverting output A.OUT of COMP is included to buffer the analog capacitive output in D/A conversion mode. By means of negative feedback from A.OUT to the same A.IN analog input to the comparator as used in A/D mode, close matching between A/D and D/A conversion transfer characteristics is achieved.[2]

With this set of functional elements, the above conversion algorithm translates as

a) Initialization:
 Preset C^{-} and $C^{-'}$ to V_{ref}^{-}, and C^{+} and $C^{+'}$
  to V_{ref}^{+} with $V_{\text{ref}}^{-} < V_{\text{ref}}^{+}$ (activate switches PRE).
 Then share the charge on C^{-} and C^{+}
  (activate switch SH).
b) Algorithmic iteration/Successive approximation:
 For n successive steps, iterate:
 If D/A conversion:
  Obtain the next input bit (from MSB to LSB);
 If A/D conversion:
  Compare V^{+} across C^{+} (or V^{-} on C^{-}) with A.IN
   (Evaluate the output D.OUT
   from the comparator)
  set the bit to D.OUT.
 If the bit is "1": Activate BRE^{-} in STO mode
    (activate STO^{-}) and
    activate BRE^{+} in RCL mode
    (activate RCL^{+});
   "0": Activate BRE^{-} in RCL mode
    (activate RCL^{-}) and
    activate BRE^{+} in STO mode
    (activate STO^{+}).
 Then, share the charge on C^{-} and C^{+}
  (activate switch SH).

c) Termination:
 If D/A conversion:
  The converted voltage is represented by either
   V^{-} across C^{-}, or V^{+} across C^{+}.
 If A/D conversion:
  The digital word is given by the sequence
  of the bits from b) in the order processed,
  the MSB first and the LSB last.

[2]This assumes, again, that both inverting and noninverting outputs exhibit the same input-referred offset voltage.

Optionally, the offsets between the precharged values on capacitors C^-, $C^{-\prime}$, C^+, and $C^{+\prime}$ in the initialization step a) can be effectively canceled by the additional step of activating both STO^+ and STO^- after a) is completed. Alternative schemes achieving the same or an equivalent result are suggested in [14].

IV. ENERGY CONSIDERATIONS

Thermal noise in analog circuitry defines an absolute limit on the power needed to maintain analog signal processing at a given signal-to-noise margin, or resolution [15]. For the relatively low resolutions considered here, this limit is not practically attainable due to other limitations in analog circuit design, such as switch charge injection noise, parasitic capacitive coupling, device matching, and junction leakage [16]. Ignoring the issue of device matching, both bit-serial and pipelined algorithmic converters dissipate the least energy per conversion since they minimize the number of analog operations performed per converted sample, which is directly proportional to n, the number of bits converted. The issue of device matching and its impact on energy consumption will be addressed in the next section. A rough estimation of the energy consumed per sample is carried out next.

Let α be the duty cycle of the STO and RCL switching operations per bit cycle, let $V_{ref} = V_{ref}^+ - V_{ref}^-$, and let all capacitances of interest equal C. Then, the minimum OTA supply current required to maintain a (slew-rate limited) $V_{ref}/2$ voltage swing at the output driving a load C is given by

$$I_{supply} = \frac{(n+1)\,\beta C V_{ref}}{2\alpha T} \qquad (1)$$

where $\beta > 1$ denotes a factor relating the OTA supply current to the OTA output tail current, and where T is the full conversion time interval including one initialization cycle. Assuming identical bias conditions for the comparator/buffer OTA and the two BRE OTA's, an estimate of the total energy dissipated per conversion cycle T is obtained as

$$E_{sample} \approx \frac{3\beta}{2\alpha}\,(n+1)\,C V_{ref}\,V_{dd}. \qquad (2)$$

This estimate does not include losses due to capacitor precharging and parasitic capacitances, which account for a small fraction of the total energy dissipation. Notice that similar estimates can be obtained for other types of bit-serial or pipelined algorithmic converters. Differences in energy consumption between architectures result from design constraints on the size of C and β.

V. PRACTICAL LIMITATIONS

The lowest achievable limit on energy consumption of any converter architecture relates directly to the design constraints imposed by accuracy considerations. This is especially so in cases where precise matching between components is a concern, and the active and passive devices need to be sized accordingly. The A/D/A architecture presented does not suffer from matching or offset limitations on device sizing, but is still affected by effects of finite gain in the OTA devices, switch charge injection on the storage capacitors, and circuit noise.

Mismatch between C^- and C^+, including parasitics and nonlinear effects (excluding hysteresis), does not affect the monotonicity of conversion. This is due to the convex nature of the charge sharing process, i.e., V is maintained strictly within V^- and V^+. Capacitive mismatch still affects the integral nonlinearity of conversion. Theoretical analysis and simulations, not included here, show that the effect of capacitive mismatch on integral nonlinearity is comparable with that of traditional algorithmic D/A and A/D converters, while the differential nonlinearity is significantly less. For applications where only monotonicity is required and integral nonlinearity is irrelevant, none of the capacitance values are critical. This allows minimum capacitor sizing, limited only by noise and switch charge injection. The latter is partially compensated by dummy switches in a differential arrangement, documented further below.

Offsets in the OTA devices (the two BRE devices inside the D/A, and the comparator/buffer) cancel only to the extent that finite gain effects and switch charge injection effects can be ignored. As stated, offsets in the BRE devices cancel between consecutive STO and RCL operations, provided the input-referred offset voltage of the corresponding OTA is identical for both inverting and noninverting outputs. This is achieved regardless of transistor mismatches, by a symmetrical multiplexed arrangement in a differential OTA circuit, documented further below. Still, finite voltage gain of the OTA and switch charge injection on the capacitors affect a voltage offset between stored and recalled analog register values (a "replication error"), which translates directly into a differential nonlinearity error. The amount of replication error is voltage dependent, and can be reduced by increasing the amplifier gain and the capacitance values, at the expense of power dissipation. The latter defines the practical limits on the energy efficiency of the converter for a given resolution.

VI. CIRCUIT STRUCTURE

While the above architecture for the A/D/A cell can be implemented in various technologies, we describe a compact, process-robust and low-power CMOS implementation below.

The analog portion of the A/D/A cell, Fig. 3, directly implements the diagram of Fig. 2. The OTA's in the BRE's and in the comparator use a standard cascoded current mirror stage, rather than a folded cascode stage, for increased robustness to transistor mismatches at minimum bias current settings.[3] The device sizing in the design provides a dc voltage gain of at least 1000, supporting more than 8 b of resolution.

The configuration of the switches internal to the OTA's and the comparator has been chosen so as to provide close matching between the input-referred offsets of the inverting and noninverting outputs. In particular, the configuration makes use of the fact that for all OTA and COMP devices only one of the inverting and noninverting outputs needs to be active at any given time. The switches connected to D and

---

[3] This precaution is necessary since the current mismatch among identical minimum size transistors in a typical CMOS process could in the extreme case amount to 50% or more in the weak inversion region [17], possibly causing cutoff in a folded cascode mirror unless the bias current is made quite larger than strictly needed.
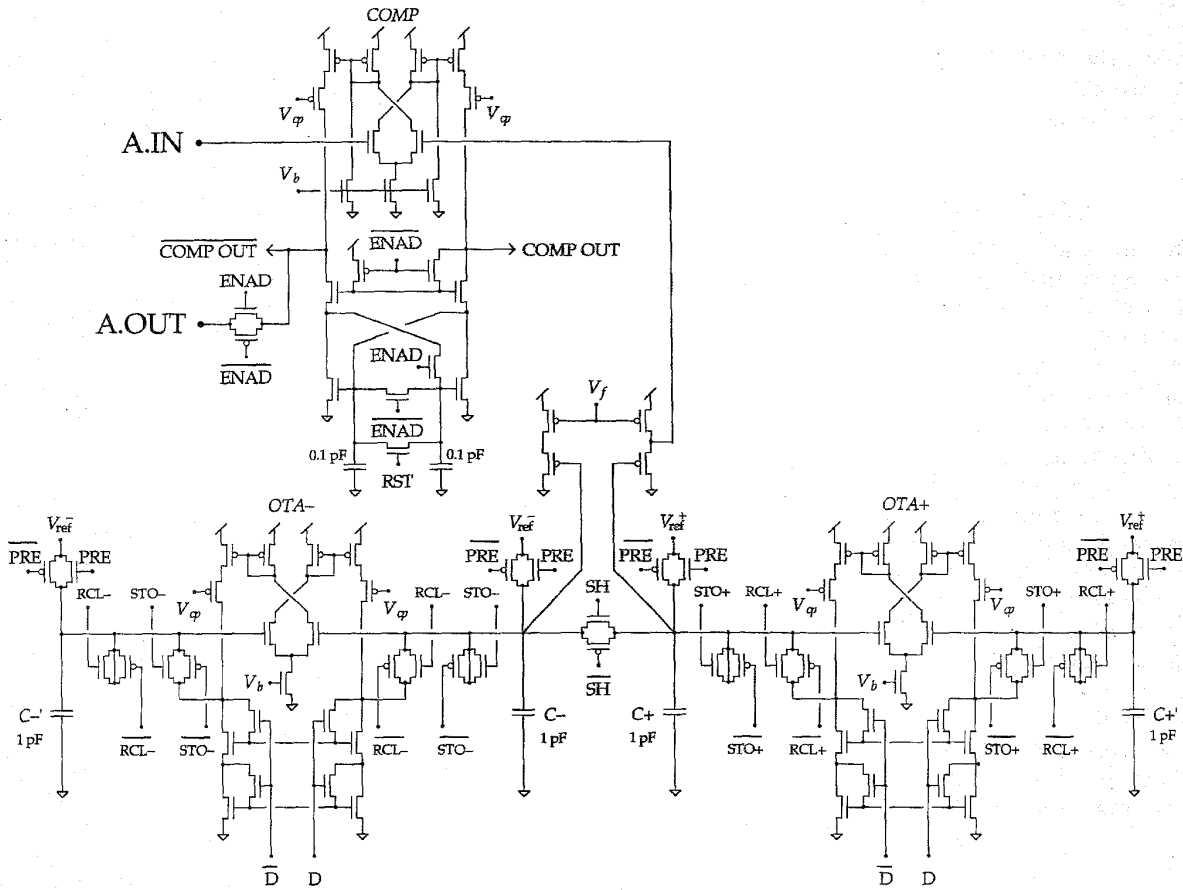
Fig. 3. Detailed schematic of the A/D/A converter cell analog circuitry.

$\overline{D}$ internal to the OTA's control the orientation of the OTA outputs in a single-ended configuration, to enable store and recall operations. Likewise, the switches ENAD and $\overline{\text{ENAD}}$ internal to the COMP device activate either the comparison or analog buffer output (D or A.OUT), to control the operational mode of the converter (A/D or D/A conversion). By using the same input and output stage transistor structures in the OTA and COMP devices for the different configurations, the same input-referred voltage offset applies to all modes of operation. The comparator, when activated in comparison mode (ENAD active), contains a cross-coupled output stage for latched operation. An active-high pulse on RST triggers and fixes the binary output.

Dummy switches are added to the capacitors, in conjunction with the STO and RCL switches, for differential compensation of the charge injection due to clock feedthrough. The switch charge injection noise appears as a common-mode input component at the OTA's, and disappears at both sides after de-activation of the selected pair of functional and dummy switches. All storage capacitors are of size 1 pF, to further reduce the effect of charge injection, and to reduce effects of parasitic capacitive coupling. A source follower isolates the capacitor C^+ from the input of the comparator to avoid capacitive coupling to the A.IN analog input. An additional source follower is added to C^- to preserve symmetry for improved matching.

Logic circuitry is integrated along with the analog circuitry in the converter cell, to control operation and to drive the STO and RCL switch waveforms. The control signal ENAD selects the operational mode of the A/D/A converter. A/D conversion mode is achieved when ENAD is in a high state, and D/A conversion mode is established in the low state. ENAD activates the inverting and noninverting outputs of the comparator, and also controls random access to the tri-state D and $\overline{D}$ signal nodes, which represent the bit-serial digital data and its complement, and which can be configured either as digital input (D.IN) or output (D.OUT). The STO and RCL switching waveforms are slew-rate controlled at the switch-off transition to reduce residual clock feedthrough.

The A/D/A converter interfaces with five individual terminals for operation control and I/O: ENAD, A.IN, A.OUT, D and $\overline{D}$. Global connections, common for all instances, include the two reference sources (V_{ref}^+ and V_{ref}^-), and several bias levels and clock waveform signals. The adjustable bias levels are provided to accommodate diverse accuracy, power and speed requirements.

Fig. 4 shows some possible I/O connection configurations of the A/D/A converter, operating as a one-way or bidirec-
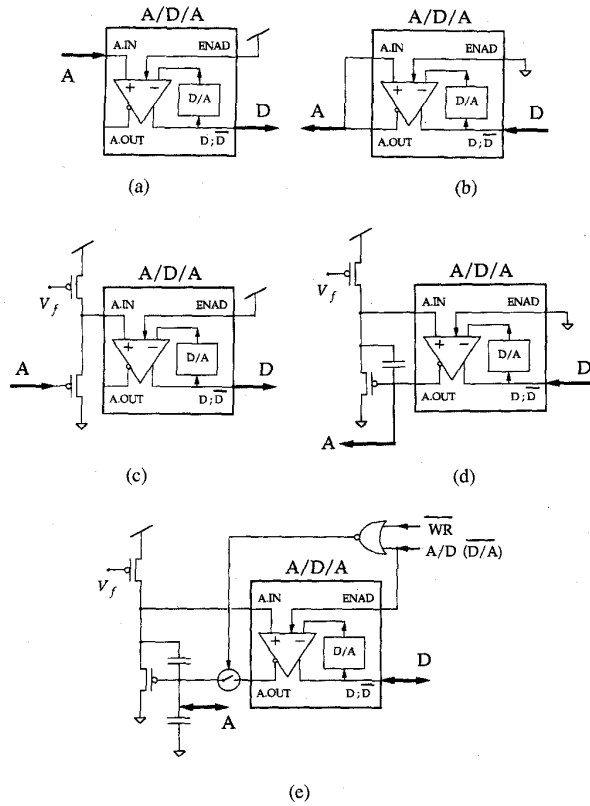
Fig. 4. I/O connection configurations. (a) Simplified A/D configuration; (b) simplified D/A configuration; (c) standard A/D configuration; (d) standard D/A configuration; and (e) bidirectional A/D/A converter implementing analog memory with A and D random access.

Fig. 5. Micrograph of the array of A/D/A converter cells.
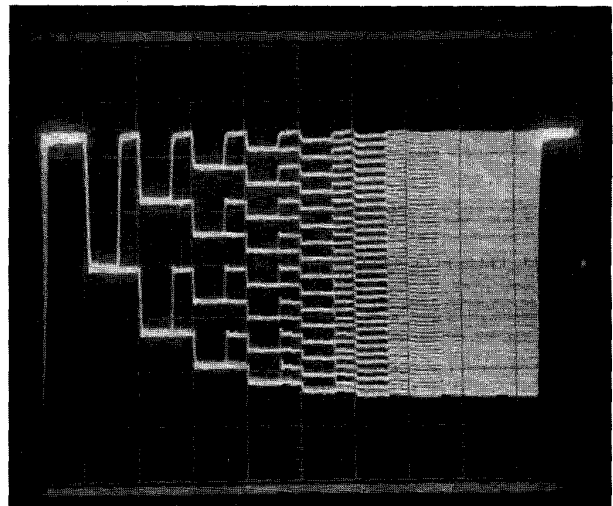
Fig. 6. Recorded algorithmic D/A conversion tree. Horizontal: 10 μs/div; vertical: 0.5 V/div.

tional converter interfacing with a hybrid analog-digital VLSI system. The external source followers at the analog input terminal A.IN in the standard A/D and D/A configurations of Fig. 4(c) and (d) balance the internal source follower stage between D/A and comparator/buffer sections in Fig. 3. In case the internal and external source followers share the same characteristics, their effect on the conversion nonlinearity is eliminated by symmetry. A compensation capacitor across the external source follower may be required for stability in the D/A configuration of Fig. 4(d), depending on the size of parasitic capacitances on the external nodes.

The external source follower can be omitted for simplicity, as in Fig 4(a) and (b), though at the expense of a loss in dynamic range and (integral) linearity, due to the nonlinear voltage drop across the internal source follower. This voltage drop applies a nonlinear transformation to the internal D/A conversion results at the input of the comparator/buffer OTA.

An analog memory with bidirectional analog and digital random read/write access is shown in Fig. 4(e), combining the A/D and D/A configurations of Fig. 4(c) and (d) through an analog switch, a storage capacitor, and some simple logic. Digital data is written on the analog storage node by activating WR following D/A conversion. Long-term memory operation is achieved by repetitively performing A/D conversions followed by WR operations in D/A mode, each time refreshing the analog value on the storage node toward its nearest discrete
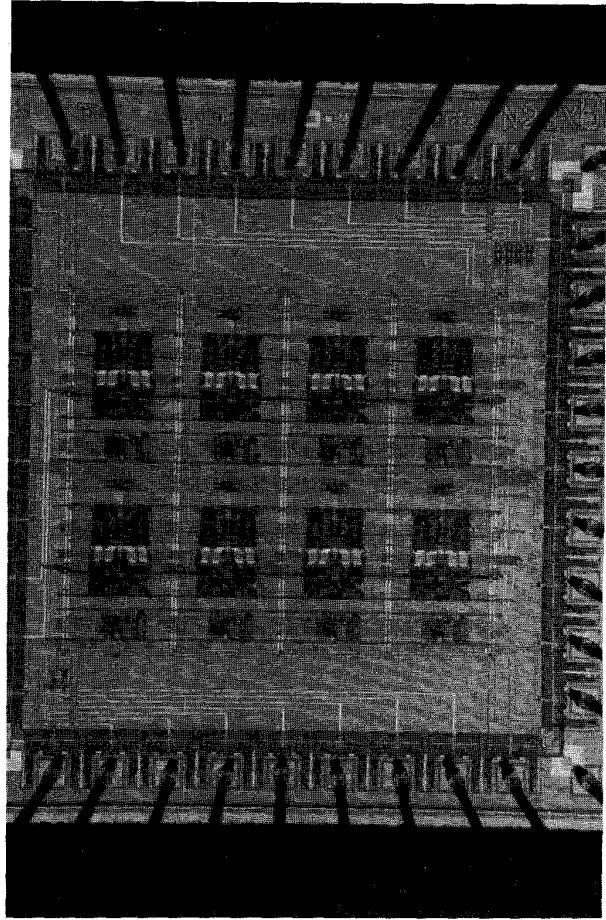
value which is represented by the final approximation D/A value of the successive approximation A/D conversion. An improved scheme with a high degree of fault-tolerance is pre-
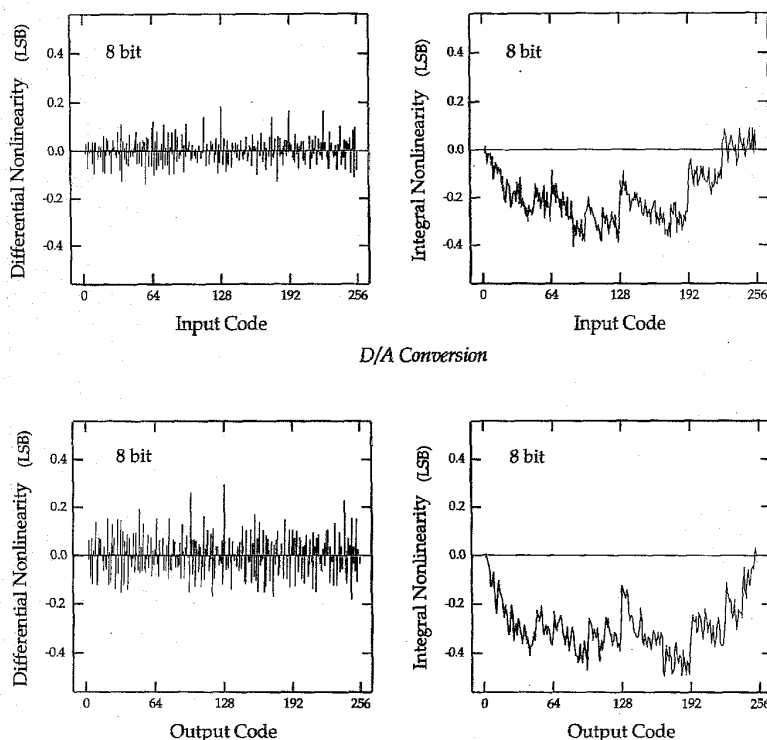
Fig. 7.   Differential and integral nonlinearity of D/A and A/D conversion by the A/D/A converter, at 2 ms conversion cycle.

sented in [13]. The bidirectional configuration of Fig. 4(e) is especially attractive in a multiplexed arrangement of an array of analog memory cells, time-sharing one common A/D/A converter. Examples of applications include synaptic storage in analog VLSI neural networks and focal-plane array image processing for motion detection or nonuniformity correction.

VII. EXPERIMENTAL RESULTS

An array of eight of the described A/D/A converters has been fabricated on a MOSIS "Tiny" chip in a 2 μm double-poly CMOS process. The chip micrograph of the array of converter cells is shown in Fig. 5. The common supply lines run across horizontally, to enable abutting of the cell with neighboring cells for the compact construction of linear arrays. (The cells have not been compactly placed in adjacent positions on the test chip, since room was available). A scope-plot displaying the binary tree of intermediate D/A conversion voltages captured from the chip is given in Fig. 6. The devices were tested at 8-b resolution for a conversion cycle ranging from 20 μs to 2 ms. For each given conversion speed, the bias levels were adjusted to obtain a globally optimal A/D and D/A conversion performance (in terms of differential linearity). Subsequently, the bias currents were further decreased to the point of perceptible performance degradation. Fig. 7 shows typical plots of A/D and D/A conversion nonlinearities at 2 ms conversion cycle and 8-b length. Whereas a uniform profile in differential nonlinearity was observed under equal biasing and timing conditions for all converters, including those from different chips, differences in the profile of the

integral nonlinearity between different cells were recorded to be more significant. This effect is attributed to the statistical variations of capacitance ratios and comparator offsets across the chip, which by virtue of the conversion algorithm do not affect the monotonicity, but which distort the integral conversion characteristic. The latter should not be considered a disadvantage per se; many applications of data conversion only require monotonicity.

The characteristics of the A/D/A converter and the performance results at different cycle speeds are summarized in Tables I and II. Again, the conversion speed is regulated by optimally adjusting the bias levels of the OTA's and the slew-rate control. As expected, speed can be traded for energy efficiency, but with better accuracy at lower speeds and higher energy efficiency at higher speeds. Optimum efficiency is achieved at 20 μs conversion cycle, at bias settings which correspond to the onset of the strong inversion MOS operation region. The experimental energy dissipation of 4 nJ per sample agrees reasonably with the calculated value of 1.6 nJ from the theoretical prediction (2), with $\alpha = 0.25$, $\beta = 2$, $n = 8$, $C = 1$ pF, $V_{ref} = 3$ V, and $V_{dd} = 5$ V. Unaccounted dissipation in (2) could come from digital circuitry, clock slew-rate control, and settling requirements of the OTA's. From (2), a nearly 10-fold decrease in power dissipation can be expected for the same circuit in 1.2 μm CMOS technology operated at 3 V power supply.

Interestingly, at 20 μs cycle speed, the A/D/A converter performs the equivalent of 25 Msamples/sec at 100 mW of power. While a linear array of A/D/A cells in a parallel multi-

TABLE I
A/D/A CELL CHARACTERISTICS

Technology	2 μm p-well double-poly CMOS
Cell Size	216 μm X 315 μm
Transistor Count	114
Supply	+5 V
Analog Conversion Range	1 V — 4 V
Resolution	8 bit
Cycle Time	≥ 20 μsec

TABLE II
A/D/A CELL PERFORMANCE AT VARIOUS CONVERSION SPEEDS

Conversion Cycle Speed	20 μsec	200 μsec	2 msec
D/A Linearity			
DNL	0.5 LSB	0.4 LSB	0.2 LSB
INL	< 2 LSB	< 2 LSB	< 2 LSB
A/D Linearity			
DNL	0.6 LSB	0.5 LSB	0.3 LSB
INL	< 2 LSB	< 2 LSB	< 2 LSB
Supply Current	40 μA	25 μA	7.5 μA
Power Dissipation	200 μW	125 μW	38 μW

rate arrangement could in principle achieve such bandwidth and efficiency at virtually no loss in precision [18], the present converter has been intended for micro-size instrumentation tasks where the typical scale of power dissipation and circuit size of some of the traditional approaches fail. Unlike dedicated high-speed converters, the converter is able to operate in environments where constraints on power and size would otherwise be too stringent, such as for biomedical implants and micro-mechanical actuators, and for focal-plane array image processing. The present design compares favorably, in size and energy efficiency, with other converters in the same category, e.g., [19], while offering the advantage of bidirectionality in the conversion.

VIII. CONCLUSION

A robust architecture for bit-serial bidirectional A/D/A conversion, with corresponding compact circuit implementation, has been presented and experimentally demonstrated. Unique features of the bidirectional cell are the MSB-first order of the serial bits, and the matched conversion characteristics in both (D/A and A/D) directions of conversion. Further, the robustness of the D/A algorithm in the presence of device mismatches supports small device sizing, which in combination with the architectural efficiency yields a more compact design and lower power (for a given bandwidth) than alternative successive approximation approaches. The advantages combine to make the A/D/A quantizer a versatile building block for applications of large-scale analog quantization and parallel array processing in environments under stringent area and power constraints. A simple 8-b A/D/A cell design of size 0.068 mm^2 in 2 μm CMOS, operating at 4 nJ energy consumption per quantized sample, has been demonstrated.

REFERENCES

[1] J. C. Candy and G. C. Temes, Eds., Oversampling Delta-Sigma Data Converters: Theory, Design, and Simulation. New York: IEEE Press, 1992.
[2] P. R. Gray, B. A. Wooley, and R. W. Brodersen, Eds., Analog MOS Integrated Circuits, II. New York: IEEE Press, 1988.
[3] K. D. Wise, "Integrated microinstrumentation systems: Smart peripherals for distributed sensing and control," in ISSCC Tech. Dig., vol. 36, pp. 126–127, 1993.
[4] Y. Horio and S. Nakamura, "Analog memories for VLSI neurocomputing," Artificial Neural Networks: Paradigms, Applications, and Hardware Implementations. New York: IEEE Press, 1992, pp. 344–363.
[5] C. A. Mead, Analog VLSI and Neural Systems. Reading, MA: Addison-Wesley, 1989.
[6] G. Cauwenberghs, "An analog VLSI recurrent neural network learning a continuous-time trajectory," IEEE Trans. Neural Networks, to be published.
[7] R. Gregorian and G. C. Temes, Analog MOS Integrated Circuits for Signal Processing. New York: Wiley, 1986.
[8] C.-C. Shih and P. R. Gray, "Reference refreshing cyclic analog-to-digital and digital-to-analog converters," IEEE J. Solid-State Circuits, vol. SC-21, pp. 544–554, 1986.
[9] K. Watanabe, G. C. Temes, and T. Tagami, "A new algorithm for cyclic and pipeline data conversion," IEEE Trans. Circuits Syst., vol. 37, pp. 249–252, 1990.
[10] L. M. Terman, Y. S. Yee, R. B. Merrill, L. G. Heller, and M. B. Pettigrew, "CCD memory using multilevel storage," IEEE J. Solid-State Circuits, vol. SC-16, pp. 472–478, 1981.
[11] B. Hochet, "Multivalued MOS memory for variable-synapse neural networks," Electron. Lett., vol. 25, no. 10, pp. 669–670, 1989.
[12] E. Vittoz, H. Oguey, M. A. Maher, O. Nys, E. Dijkstra, and M. Chevroulet, "Analog storage of adjustable synaptic weights," in VLSI Design of Neural Networks. Norwell, MA: Kluwer, 1991, pp. 47–63.
[13] G. Cauwenberghs and A. Yariv, "Fault-tolerant dynamic multi-level storage in analog VLSI," IEEE Trans. Circuits Syst. II, vol. 41, pp. 827–829, 1994.
[14] ———, "Method and apparatus for monotonic algorithmic digital-to-analog and analog-to-digital conversion," U.S. Patent 5 258 759, 1993.
[15] E. Vittoz, "Low-power low-voltage limitations and prospects in analog design," in Analog Circuit Design, R. J. van de Plassche et al., Eds. Norwell, MA: Kluwer, 1995, pp. 3–15.
[16] E. Dijkstra, O. Nys, and E. Blumenkrantz, "Low power oversampled A/D converters," in Analog Circuit Design, R. J. van de Plassche et al., Eds. Norwell, MA: Kluwer, 1995, pp. 89–103.
[17] E. Vittoz and J. Fellrath, "CMOS analog integrated circuits based on weak inversion operation," IEEE J. Solid-State Circuits, vol. 12, pp. 224–231, 1977.
[18] A. Petraglia and S. K. Mitra, "High-speed A/D conversion incorporating a QMF bank," IEEE Trans. Instrum. Meas., vol. 41, pp. 427–431, 1992.
[19] V. Valencic, P. Deval, and F. Krummenacher, "8-Bit micropower algorithmic A/D converter," Electron. Lett., vol. 23, no. 18, pp. 932–933, 1987.

Gert Cauwenberghs (S'89–M'92) received the Engineer's degree in applied physics from the Vrije Universiteit Brussel, 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.

He joined the Electrical and Computer Engineering Department of Johns Hopkins University in 1994 as an Assistant Professor. His research covers VLSI systems and algorithms for parallel signal processing and adaptive neural computation.