An Ultra Low Energy 12-bit Rate-Resolution Scalable SAR ADC for Wireless Sensor Nodes

Naveen Verma, *Student Member, IEEE*, and Anantha P. Chandrakasan, *Fellow, IEEE*

***Abstract*—A resolution-rate scalable ADC for micro-sensor networks is described. Based on the successive approximation register (SAR) architecture, this ADC has two resolution modes: 12 bit and 8 bit, and its sampling rate is scalable, at a constant figure-of-merit, from 0–100 kS/s and 0–200 kS/s, respectively. At the highest performance point (i.e., 12 bit, 100 kS/s), the entire ADC (including digital, analog, and reference power) consumes 25 μW from a 1-V supply. The ADC's CMRR is enhanced by common-mode independent sampling and passive auto-zero reference generation. The efficiency of the comparator is improved by an analog offset calibrating latch, and the preamplifier settling time is relaxed by self-timing the bit-decisions. Prototyped in a 0.18-μm, 5M2P CMOS process, the ADC, at 12 bit, 100 kS/s, achieves a Nyquist SNDR of 65 dB (10.55 ENOB) and an SFDR of 71 dB. Its INL and DNL are 0.68 LSB and 0.66 LSB, respectively.**

***Index Terms*—ADC, analog-to-digital conversion, circuit noise, CMOS analog integrated circuits, low-power electronics, offset compensation, scaleable, successive approximation register.**

I. Introduction

WIRELESS sensor networks offer a sophisticated platform for environment observation. The vision of a micro-sensor network includes dense, intelligent nodes that are energy-autonomous and that operate and are deployed in an ad hoc manner. Nodes are capable of self-organizing into a collaborative network, and subsequently benefit from spatial diversity through data sharing and multi-hop connectivity [1], [2]. Such networks have broad applications ranging from military surveillance, reconnaissance, and damage assessment to environmental forest fire detection [3] and industrial process monitoring.

The design of sensor node hardware is constrained by several factors. To be energy-autonomous, nodes must be powered entirely by an energy harvesting source. This places demanding, low-energy requirements on the constituent circuits. Ad hoc deployment and operation requires that nodes be fault tolerant and able to adapt to unpredictable environments and network characteristics. Finally, ubiquity places a cost constraint on nodes, reducing their acceptable price per unit to a few cents. Fundamentally, the architecture of an intelligent sensor node consists of an analog-to-digital converter (ADC), a digital signal processor (DSP), and a short range radio. This paper describes the design of an ultra-low-power ADC suitable for sensor nodes [4]. In this context, the ADC has a maximum resolution of 12 bits and a sampling rate of up to 100 kS/s, enabling the conversion of signals with the dynamic range and frequencies expected during environment monitoring. However, since both the performance demands and energy budget are time-varying and unpredictable, the ADC also has a low-power 8-bit mode with a maximum sampling rate of 200 kS/s. Further, the sampling rate can be reduced arbitrarily for linear power savings in both resolution modes. Both oversampling and successive approximation register (SAR) architectures are notable for achieving these specifications at the lowest power levels. As resolutions increase beyond 8 bits, oversampling converters have shown to be the most efficient and have the added advantage of reduced anti-aliasing requirements. In sensor applications, however, events occur sporadically, and the nodes might acquire data only once before having to react. As a result, general Nyquist acquisition is preferred. In this design, the SAR architecture is used, and techniques are developed to efficiently increase the resolution to 12 bits.

Section II describes the global architecture of the ADC and discusses system-level approaches and optimizations for enhancing efficiency and achieving scalability. Section III describes specific circuits that have been used in the implementation of the charge-redistribution DAC and the comparator, which are the two blocks most critical to the speed, power, and precision of the converter. Section IV details the experimental results of the fabricated prototype, and, finally, Section V provides a short conclusion and comparison study for this design.

II. Architecture Design

The main components of a SAR ADC are a sample-and-hold (S/H), a digital-to-analog converter (DAC), a comparator, and a digital state machine (itself called the SAR) [5]. A block diagram of this ADC is shown in Fig. 1. The DAC has been separated into a main-DAC and a sub-DAC, both implemented as passive charge-redistribution capacitor arrays. Additionally, the Clock Manager block and SLEEP signal are used to implement scalability features. The remainder of this section describes the global architecture and associated optimizations affecting the overall conversion process.

A. Optimal Supply Voltage

Lowering the supply voltage is an effective strategy for reducing the power consumption of digital circuits. In the presence of noise, mismatch, finite switch resistance, and distortion, however, the power consumption of analog circuits is liable to increase with reduced supply. Fig. 2 shows the expected energy, with respect to supply voltage, of various components in this

Manuscript received June 14, 2006; revised January 30, 2007. This work was supported by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory, under agreement number F33615-02-2-4005. The IC fabrication was provided by National Semiconductor Corporation.

The authors are with the Microsystems Technology Laboratories, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: nverma@mtl.mit.edu; anantha@mtl.mit.edu).

Digital Object Identifier 10.1109/JSSC.2007.897157

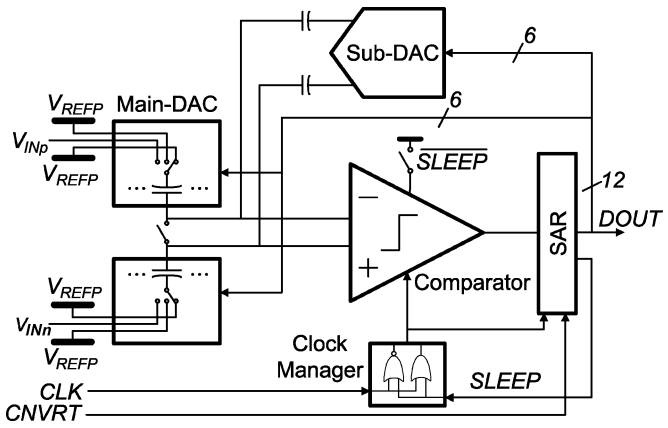

Fig. 1. Implemented ADC block diagram.

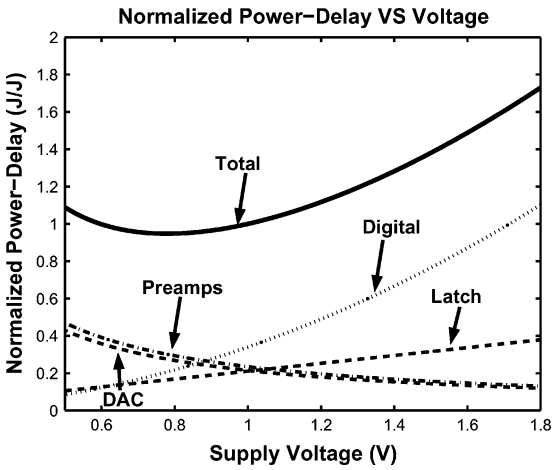

Fig. 2. Normalized ADC energy (power–delay) versus supply voltage.

ADC under the assumption of constant SNR. The relative contributions of the ADC components strongly impact the overall energy profile and are highly implementation dependent. Consequently, for this study, the normalized energies at 1 V are based on actual measurements of the prototype, and these are extrapolated over a range of $V_{\rm DD}$ using fundamental analytical models. Here, the digital circuits follow a quadratic trajectory with respect to supply voltage, while the non-noise-limited analog circuits (i.e., latch) follow a linear trajectory. However, the noise-limited circuits (i.e., preamplifiers and S/H DAC) follow an inverse trajectory.

The total optimum supply voltage occurs at 800 mV. Although none of the active blocks in the SAR ADC have large linearity requirements, biasing circuits and floating analog switches do impose a practical limit on how far the supply voltage can be reduced. Specifically, for biasing, stacked pMOS and nMOS diode structures occur frequently, and their reliability is increased if at least one of the devices is in strong inversion. Similarly, the conductance of analog transmission switches is drastically reduced at low supplies when passing midrail voltages, and their nonlinear nature introduces severe distortion in the input sampling switch even though on-chip charge-pumps are used. Thus, with consideration to the threshold voltages in the target technology, a supply voltage of 1 V was selected, which is very close to the broad minimum in Fig. 2.

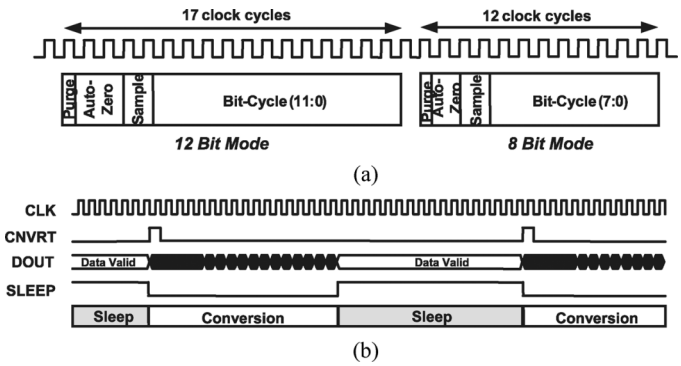

Fig. 3. ADC conversion waveforms showing (a) 12-bit/8-bit conversion plans and (b) power-gating control.

B. Conversion Plan

The conversion plans for the 12-bit and 8-bit modes of this ADC are shown in Fig. 3(a). The conversion starts by purging the DAC capacitors so that they can be used to derive a suitable auto-zeroing reference. Additionally, the auto-zeroing and sampling operations are separated. As a consequence, sampling is delayed with respect to the start of the conversion. This delay is undesirable in some applications. However, as described in Section III, purging, auto-zeroing, and sampling in separate phases improves the common-mode rejection, low-voltage operation, and noise performance of the ADC.

C. Sample Rate Scaling

Assuming a constant supply voltage, the power consumption of both digital circuits and analog circuits (in weak inversion) is directly proportional to their operating speeds. It follows then that the ADC energy-per-conversion does not depend on the conversion rate. In the case of digital circuits, the energy required for a logic transition on a node with capacitance C is simply $CV_{\rm DD}^2$, independent of frequency. Additional forms of energy, namely, direct-path and leakage energy are also present but to a lesser extent. In the case of analog circuits, for example a single-stage amplifier, the energy, or power–delay product, in weak-inversion (which is the most energy-efficient regime) is set by $g_m\tau$. Specifically, g_m is the amplifier transconductance, which is proportional to the bias current, and therefore power, while τ is the output time constant. In weak inversion, this can be expressed as $\alpha I_B R_{\rm OUT} C_{\rm OUT}$, where α is a physical constant. In a sampled system, where analog processing is performed in the discrete-time domain, a reduced sampling rate implies reduced amplifier bandwidth. Consequently, $R_{\rm OUT}$ can be increased. Then, to maintain some required gain (given by $g_m R_{\rm OUT}$), g_m, and therefore I_B, can be decreased by the same factor, yielding no change in the total power–delay.

These results suggest that ADC power scaling with respect to sampling rate can be achieved by increasing the clock period and appropriately adjusting the analog bias currents. This approach has indeed been demonstrated successfully [6], [7]. Alternatively, however, power scalability can also be achieved by performing conversions at a constant, maximum rate, and then clock-gating the digital circuits and power-gating the analog circuits between active conversions when a reduced sampling rate is desired. Power-gating eliminates the analog bias currents

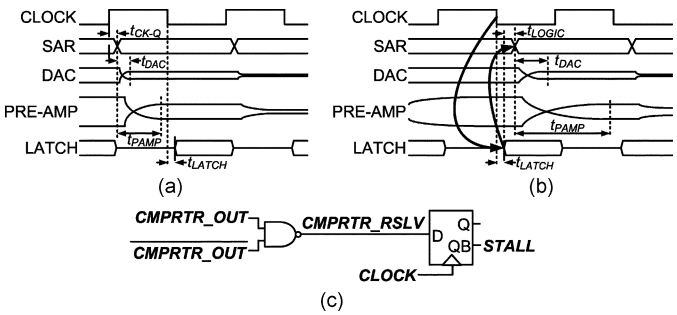

Fig. 4. Waveforms showing (a) standard bit-cycling, (b) self-timed bit-cycling, and (c) circuit to detect excessive latch delay.

in a manner similar to [8]. The former approach suffers from the reliability of altering bias currents and output resistances and the overhead of controlling the adaptation. The latter approach suffers from limitations of long bias-up times between sleep–active transitions. In this SAR ADC, however, the only active block, the comparator, does not require a long bias-up time, so the associated overhead is negligible. Accordingly, as shown in Fig. 3(b), this approach has been adopted. Specifically, the input pulse, CNVRT, initiates a conversion, and, upon completion, the SAR logic asserts a SLEEP signal to clock-gate the digital state machine and power-gate the comparator. The case where the sampling rate has been scaled to half the maximum rate is shown.

D. Self-Timed Bit-Cycling

A straightforward timing scheme for controlling bit-cycling is shown in Fig. 4(a). The DAC and preamplifiers settle during the first half of the clock cycle, and the latch resolves during the second half. In Fig. 4(b), the self-timed scheme used in this design is shown. Here, a latch decision (to either logic value) clocks the SAR logic. This triggers the next bit-cycle phase, and the DAC and preamplifiers immediately begin settling to their subsequent values [9]. Consequently, their settling time can be longer by $(t_{\rm CLK}/2)+t_{\rm CK-Q}-t_{\rm LATCH}-t_{\rm LOGIC}$, where $t_{\rm CLK}$ is the clock period. In this design, the preamplifiers dominate the bit-cycling time (i.e., the logic and latch delays are small), and their settling time can be nearly doubled, considerably reducing their power consumption. It is worth noting that, in SAR conversion, critical bit decisions never occur successively. As a result, a critical bit decision to resolve one of the MSBs will benefit from a short $t_{\rm LATCH}$ since the previous decision will always be fairly relaxed. A critical bit decision to resolve one of the LSBs does not limit the speed of the ADC since the remaining dynamic range is reduced and no recovery time from overdrive is required.

Self-timing does, however, introduce a failure mode since it relies on a latch decision. Very long latch resolution times can stall the bit-cycling process and, therefore, must be detected. The circuit in Fig. 4(c) detects whether the latch has failed to resolve by the rising clock edge. The reset state of both CMPRTR_OUT and $\overline{\rm CMPRTR\_OUT}$ is "1". If both remain high after the allocated half clock-cycle, a bit decision is forced.

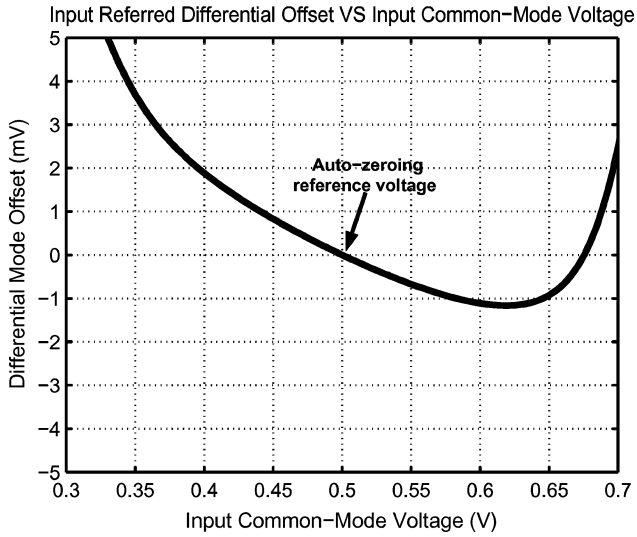

Fig. 5. Differential-mode error with respect input common-mode for and auto-zeroed amplifier.

III. Circuit Blocks

In this section, the block level architectures of the DAC and comparator are described. Since these components are critical with regards to the power consumption, speed, and precision of the entire ADC, much of the design effort has focused on optimizing and improving their performance.

A. DAC Circuits

Single-ended, instead of differential-ended, inputs can be used to greatly ease system complexity. Support of single-ended sampling is provided by means of pseudo-differential sampling, where one of the differential terminals can be set to reference ground. This approach is viable only if the common-mode signal is properly treated. Specifically, the sampled common-mode affects the DAC output common-mode during bit-decisions. For proper offset cancellation in the comparator preamplifiers, however, it is critical that the DAC common-mode be equal to the auto-zeroing reference voltage. This is illustrated in Fig. 5 which shows the differential-mode error at the preamplifier output as its input common-mode varies. Here, 3σ mismatch is applied to the input devices, and the preamplifier is auto-zeroed with an input reference voltage of 500 mV. As shown, no differential-mode error is observed if the input common-mode equals the auto-zeroing reference. However, if the input common-mode deviates by even a few hundred millivolts, mismatch results in a differential current through the preamplifier input devices, causing a large differential-mode voltage error of several millivolts at the output. This limitation is mitigated in two ways: 1) common-mode independent acquisition, and 2) preamplifier auto-zeroing to the voltage of critical SAR decisions. First, as shown in Fig. 6(a), the capacitor arrays are purged of previous charge by shorting their top and bottom plates. Then, as shown in Fig. 6(b), they are switched so that an appropriate auto-zeroing reference can be generated for the comparator. Finally, input sampling is performed.

During sampling, which is shown in Fig. 6(c), the top-plates of the differential capacitor arrays are shorted, and their voltage

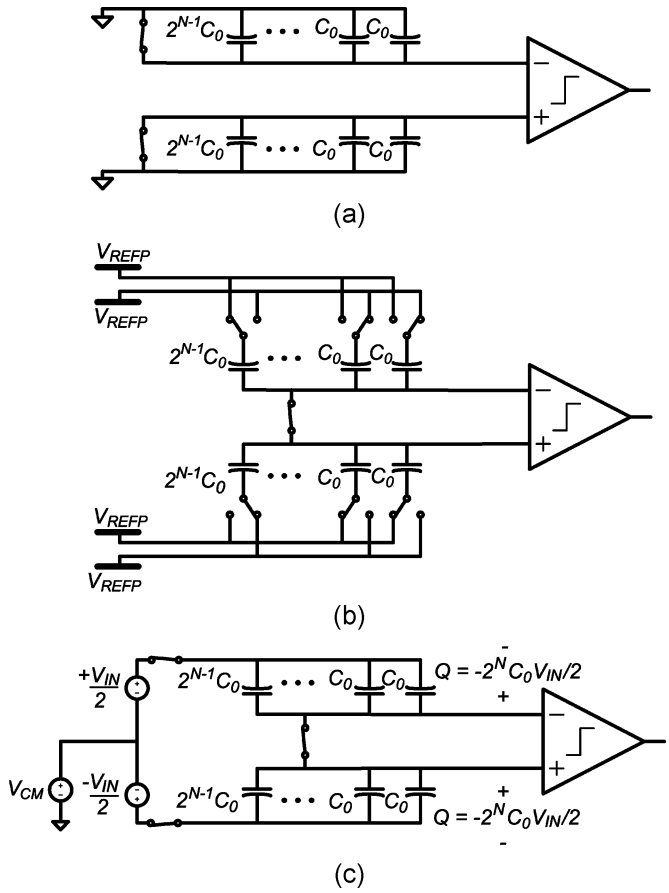

Fig. 6. DAC arrays during (a) purging, (b) auto-zeroing, and (c) sampling.

simply floats to the input common-mode. Since only one switch is required to decouple the positive and negative arrays, the charge injection errors in this network are not subject to switch mismatch. The residual injection error is minimized further by properly sizing the CMOS bottom-plate input switch. Specifically, constituent pMOS and nMOS devices are sized so that the total switch impedance is symmetric about the common-mode voltage expected during differential sampling (i.e., mid-$V_{\rm ref}$). Consequently, the top-plate sampling switch sees a well matched impedance on either side, and its channel charge distributes equally (of course, this only applies to differential sampling since, during pseudo-differential sampling, the common-mode voltage depends on the input signal). An advantage of this sampling network is that, since half the input is sampled on each array, the DAC outputs always remain within the rails [9]. More importantly, however, the top-plate voltage floats to the input common-mode, and, consequently, as desired, only the differential-mode input signal gets sampled. Unfortunately, this does imply that, during pseudo-differential sampling, the top-plate voltage can vary depending upon the input signal. As a result, it is not suitable for auto-zeroing since it can deviate greatly from the common-mode voltage during bit-decisions. Hence, to avoid the preamplifier errors described above, the sampling and auto-zeroing operations must be separated.

Sampling only differential-mode charge guarantees that the DAC output voltage during critical bit decisions will always be centered around midscale. Hence, to avoid differential mode errors in the preamplifiers due to device mismatch, the auto-zeroing voltage should also be at midscale. This voltage must be generated prior to sampling by switching the purged capacitor arrays into the divider configuration shown in Fig. 6(b). Although separating the auto-zeroing and sampling phases increases the conversion time, each can be independently optimized, as described in the following subsections, minimizing the associated overhead. A primary advantage, however, is that since the DAC is used to generate both the auto-zeroing reference as well as the conversion bit-voltages, no floating analog switches are present in the signal path. This eliminates an additional source of charge injection error and enhances the reliability and speed in this low voltage implementation.

B. Comparator Architecture

The comparator is responsible for resolving small inputs near an LSB voltage into full-scale digital values. In this role it has immense gain, speed, and sensitivity requirements. A block diagram of the circuit used is shown in Fig. 7 and consists of two gain paths: the top path is used in 12-bit mode and has three cascaded preamplifiers; the bottom path is used in 8-bit mode and has only one preamplifier. Resolution scaling occurs via two mechanisms. First, bit-cycling is stopped at the desired precision, yielding linear power savings. Additional power savings, of approximately 70% in the preamplifiers, are achieved by selecting the appropriate comparator gain path and disabling the unused path.

The preamplifier circuit has a nominal gain of 3. However, during auto-zeroing, offset cancellation is performed by storing the observed voltage on the output capacitors. Accordingly, output offset compensation in this manner does not correct the gain compression seen due to device mismatch. However, the preamplifier topology used exhibits minimal degradation, having a gain of over 2.5 in the presence of 3σ device mismatch. Further, it has the benefit of appropriate input and output voltage levels at the selected bias currents. Consequently, output auto-zeroing can be performed by simply purging the auto-zeroing capacitors and shorting them together to high-impedance using the switches shown in Fig. 7.

Since the first preamplifier in the 12 bit cascade is noise limited, a bandwidth limiting output capacitor is used to manage its SNR, giving this stage the longest time-constant and greatest power consumption. However, all the preamplifier bias currents have been minimized, and, due to parasitics, the delays of even the later stages are sizeable. Under these circumstances reducing the gain requirements of the entire cascade saves considerable power. It has been shown that regenerative latches have a superior power–delay product compared to cascaded linear amplifiers [10]. One of the major difficulties, however, is that they can have offsets over 100 mV. Often, in sensor applications, processing is performed on ADC data gathered by multiple different nodes. Consequently, the relative offsets of the ADCs must be small, and offset compensation becomes critical. Accordingly, offset cancelled preamplifiers can be used to obtain a signal swing beyond the latch offset. In this design, however, to maximize the comparator efficiency, an offset compensating latch (described in a later subsection) is used. As a result, even

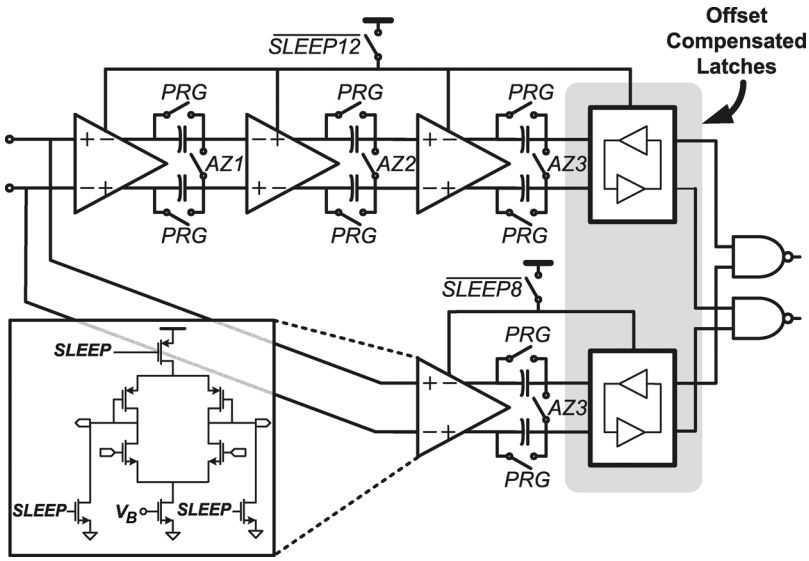

Fig. 7. Comparator block diagram.

in the presence of severe mismatch, the swing requirement at the output of the preamplifiers is just a couple of millivolts.

C. Preamplifier Auto-Zeroing Time Optimization

The passive sampling network has a very fast time constant to minimize the effect of nonlinear sampling switch resistance. Auto-zeroing, however, involves the active preamplifiers and fast operation consumes significant power. Separating the two operations allows the auto-zeroing time to be independently optimized with consideration to noise performance. The preamplifiers are subject to noise during both auto-zeroing and bit decisions, and the relevant networks are shown in Fig. 8(a) and (b). Since the preamplifiers settle just once for auto-zeroing, but 12 times for the bit decisions, it is beneficial to reduce the auto-zeroing noise at the cost of increased settling time.

Analytically, the total noise variance of the first preamplifier is related to the load capacitance during bit-cycling, $C_{\rm L}$, and that during auto-zeroing, $C_{\rm AZ,tot}$, i.e.

$$\overline{v_{\rm no,tot}^2} = \frac{B}{C_{\rm L}} + \frac{B}{C_{\rm AZ,tot}} \tag{1}$$

where B is a noise-setting circuit parameter. Note, this expression is an approximation where the noise from the preamplifier is assumed to dominant above the noise from the auto-zeroing switches. Actual capacitor sizes and circuit parameters validate this assumption. Accordingly, $C_{\rm L}$ can be expressed in terms of $C_{\rm AZ,tot}$:

$$C_{\rm L} = \frac{C_{\rm AZ,tot}}{B' C_{\rm AZ,tot} - 1}. \tag{2}$$

Here, the factor B' is the desired noise variance normalized by B. Now, the total power–delay product of the preamplifier, assuming 12 clock cycles for bit-cycling (to resolve 12 bits) and K clock cycles for auto-zeroing, is given by (3), where K is to be determined. It is assumed here that neither G_M, the amplifier transconductance, nor $R_{\rm out}$, the output impedance, change during the conversion. Note, the sampling energy is not considered since it is set by unrelated time-constants and simply adds a constant to this analysis.

$$\sum_{i=1}^{K+12} G_M \tau_i = K G_M R_{\rm out} C_{\rm AZ,tot} + 12 G_M R_{\rm out} C_{\rm L}. \tag{3}$$

Additionally, in the case that equally complete preamplifier settling is required during auto-zeroing and bit-cycling, the ratio of

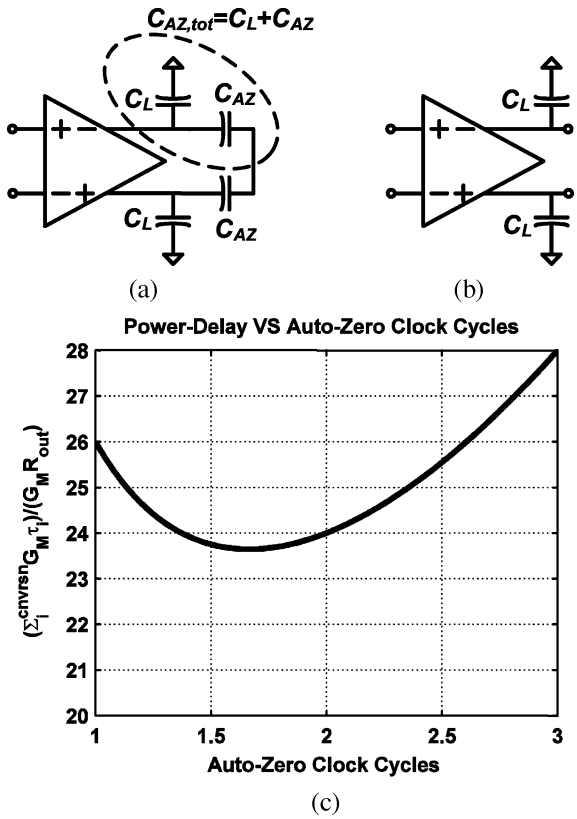

Fig. 8. Preamplifier during (a) auto-zeroing, (b) bit-cycling, and (c) normalized power–delay of auto-zeroed preamplifier.

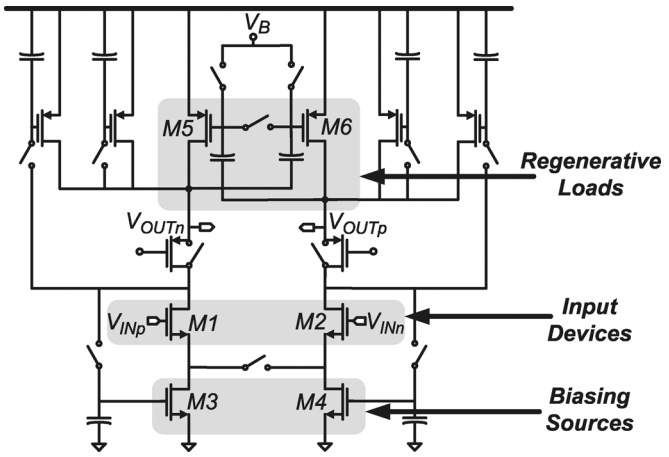

Fig. 9. Simplified offset compensating latch.

$C_{\rm L}$ and $C_{\rm AZ,tot}$ is given by (4), where the settling time during bit-cycling is set to 1 clock cycle.

$$\frac{1}{K} = \frac{\tau_{\rm bit-cycle}}{\tau_{\rm azero}} = \frac{R_{\rm out}C_{\rm L}}{R_{\rm out}C_{\rm AZ,tot}} = \frac{C_{\rm L}}{C_{\rm AZ,tot}}. \tag{4}$$

Finally, substituting (2) and (4) into (3), and normalizing by $G_M R_{\rm out}$, the power–delay can be expressed as in (5):

$$\frac{B'\sum_{i=1}^{K+12}\tau_i}{R_{\rm out}} = K(K+1) + 12\frac{K+1}{K}. \tag{5}$$

This function is plotted in Fig. 8(c), where it is shown that minimum power–delay occurs for $K \approx 1.6$. Based on this result, 1.5 clock cycles have been used to auto-zero the first preamplifier, and the ratio of its load capacitances have been set appropriately.

D. Offset Compensating Latch

A major difficulty with latch compensation for a multi-step SAR converter is that offset compensation requires applying a reference to the input and storing the observed offset. Latches, however, must be reset after every decision making it difficult to preserve the calibration information. The circuit used in this design preserves the calibration biasing by operating over two phases: the auto-zeroing phase occurs once at the start of the conversion, and the reset-resolve phase occurs once during each bit-decision. The circuit used is shown in Fig. 9, and the input devices (M1–M2), regenerative loads (M5–M6), and biasing current sources (M3–M4) are highlighted.

The auto-zeroing phase is shown in Fig. 10, where the gray devices have been deactivated by applying the appropriate rail voltage to their gates. The goal here is to bias M3–M4 such that, when the input voltages are equal, the voltages at the source nodes, $V_{\rm SRC1,2}$, are equal. This will allow the input voltages to be accurately reflected to the source nodes without requiring any special biasing at the gates of the input devices. Initially, $S_{\rm DIFFMD}$, $S_{\rm CSCD}$, and $S_{\rm FB}$ are all closed, and a zero-differential input is applied to M1–M2, as shown in Fig. 10(a). Closing $S_{\rm DIFFMD}$ forces $V_{\rm SRC1,2}$ to be equal, and M3–4/M11–12 get biased under this constraint. Subsequently, $S_{\rm DIFFMD}$ and $S_{\rm CSCD}$ are opened, as shown in Fig. 10(b). Ignoring the feedback connections through $S_{\rm FB}$, opening $S_{\rm CSCD}$ causes $V_{\rm HG1,2}$ to become high-impedance nodes. Opening $S_{\rm DIFFMD}$ causes an incremental change in the branch currents due to device offsets, and these appear at $V_{\rm HG1,2}$ as large incremental voltages. The large voltages get fed back (through $S_{\rm FB}$), rebiasing M3–M4 until the branch currents return to their original values. Accordingly, this causes the $V'_{\rm GS}s$ of M1–M2 to return to their original values, and $V_{\rm SRC1,2}$ are equalized as desired.

During bit-decisions, the latch is in the reset-resolve phase (Fig. 11). Here, $S_{\rm DIFFMD}$ is initially open, and the input voltages to be resolved appear at the source nodes, $V_{\rm SRC1,2}$, through a $V_{\rm GS}$ drop. The branch currents, however, are set only by M3–M4, and regeneration is disabled in M5–M6 by applying a static bias to their gates through $S_{\rm R}$ and $S_{\rm AUX}$, as shown in Fig. 11(a). As a result, voltages generate across the $C'_{\rm R}s$ that, in the presence of all device offsets (including M5–M8), would hold the loads in a metastable state if $S_{\rm R}$ and $S_{\rm AUX}$ were opened. Accordingly, to trigger regeneration, the $S_{\rm R}$ switches are opened and a short time later, $S_{\rm AUX}$ is opened. Delaying $S_{\rm AUX}$ minimizes the charge injection error due to mismatch in the $S_{\rm R}$ switches [11]. Finally, $S_{\rm DIFFMD}$ is closed, as shown in Fig. 11(b). Recall that the input voltages were reflected to the source nodes, $V_{\rm SRC1,2}$, and now, a current path between them is introduced. As a result, the branch currents get perturbed depending on the input voltages, and the latch resolves. M1–M2 introduce a limitation in this structure since the gain from the inputs to their drains can cause a large difference in their $V_{\rm DS}$'s. This causes a difference in the current through their output conductances (i.e., $r'_o s$) which shunts current away from their g_m generators. Consequently, the calibration biasing from the auto-zeroing phase can get disturbed. To minimize this problematic gain, low-impedance diode connected devices, M7–M8, are introduced. During regeneration, the diode connection is broken ($S_{\rm SC}$ switches are opened along with $S_{\rm R}$ switches) to enable strong positive feedback.

E. Offset Compensating Latch—Auto-Zeroing Phase Analysis

Circuit operation during the auto-zeroing phase can also be analyzed through the feedback block diagram shown in Fig. 10(c). Here, the impedance of the pMOS cascode (M9–M12) is assumed to be infinite after $S_{\rm CSCD}$ is opened. In Fig. 10(c), the transresistance, $R_{\rm CSCD}$, is approximately $g_{m1,2}r_{o1,2}r_{o3,4}$, and the resistance at the nodes $V_{\rm SRC1,2}$, is approximately $r_{o3,4}$. The residual incremental voltage at $V_{\rm SRC1,2}$ is given by (6):

$$v_{\rm src1,2} = \frac{i_{\rm os}}{g_{m3,4}g_{m1,2}r_{o1,2}} \tag{6}$$

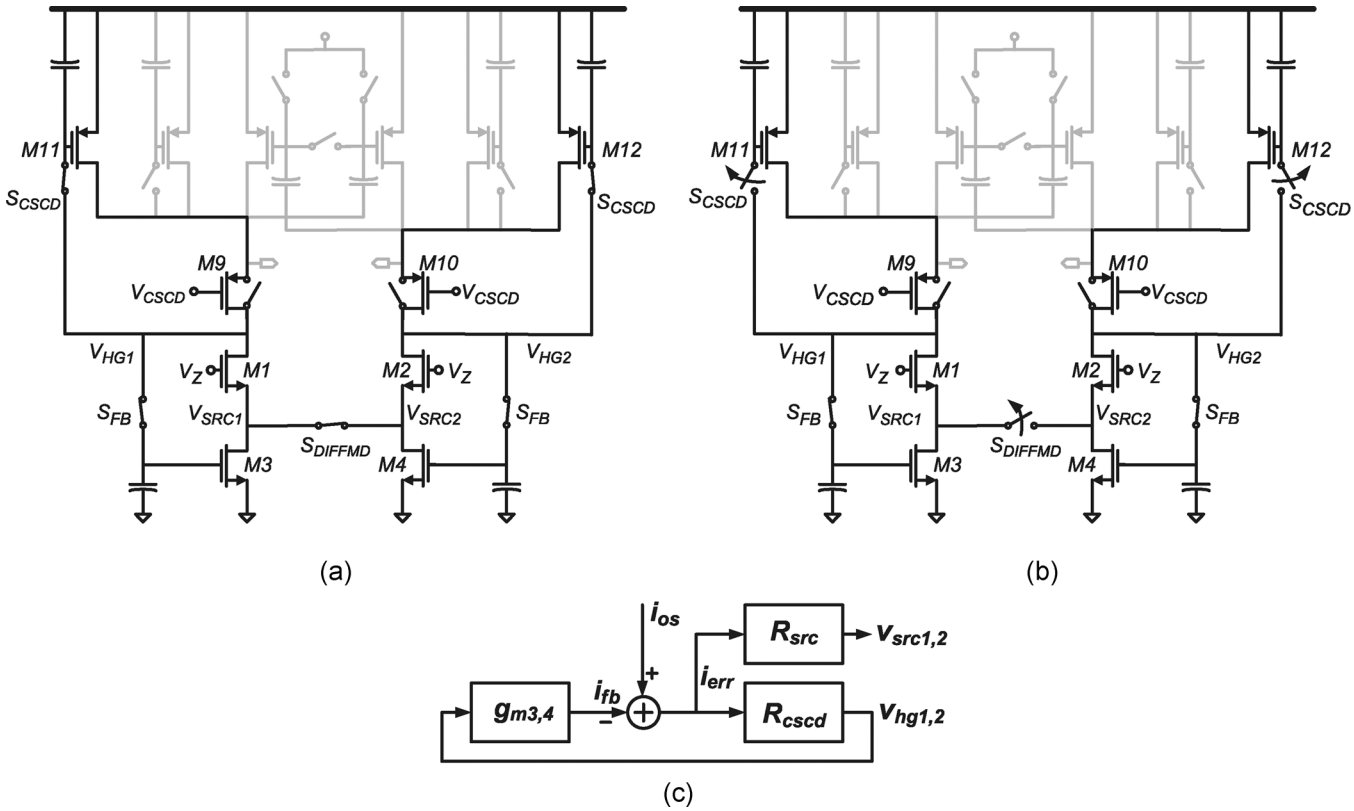

Fig. 10. Offset compensating latch during auto-zero phase: (a) first half, (b) second half, and (c) feedback block diagram.

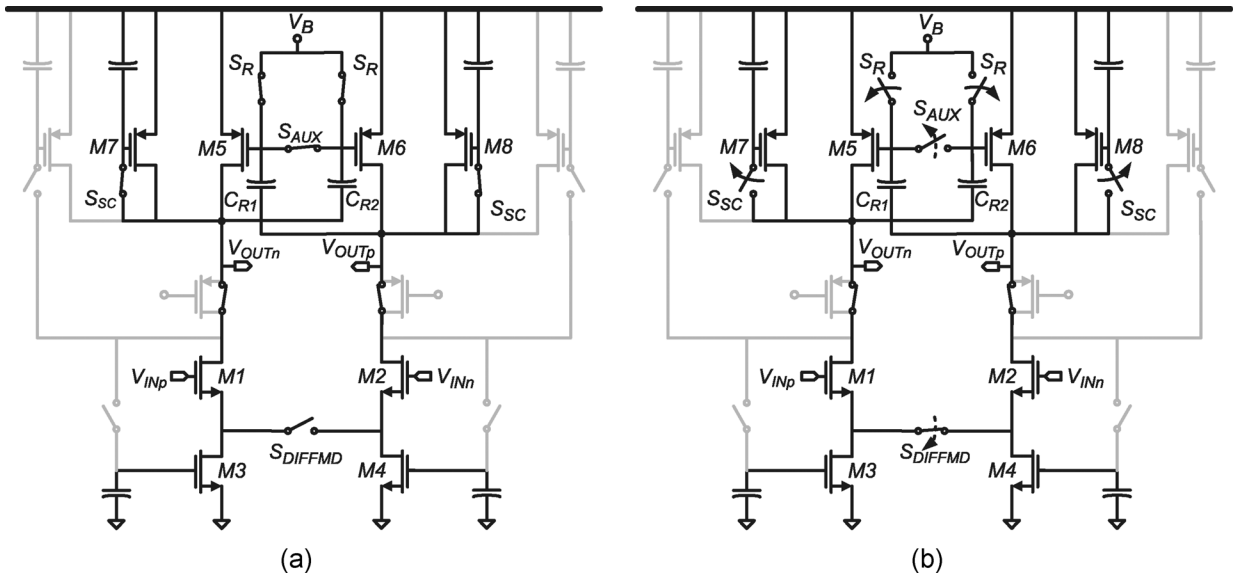

Fig. 11. Offset compensating latch during reset-resolve phase: (a) first half and (b) second half.

where $i_{\rm os}$ is the offset current that flows between the source nodes due to the voltage offset between them. The impedance at the nodes $V_{\rm SRC1,2}$, while $S_{\rm CSCD}$ is closed, relates the offset current to the uncalibrated offset voltage, $v_{\rm os}$. Specifically, $i_{\rm os}$ is approximately given by $g_{m1,2}v_{\rm os}$. Then, the reduction in the offset voltage achieved by the feedback loop is as given by (7):

$$v_{\rm src1,2} = \frac{v_{\rm os}}{g_{m3,4}r_{o1,2}} \tag{7}$$

Specifically, after calibration, the offset is reduced by approximately an intrinsic gain factor.

IV. Testing and Characterization

The low-power ADC was fabricated in a 0.18-μm five-metal two-poly (5M2P) CMOS process. It was packaged in a 0.5 mm pitch TQFP package. A micrograph of the entire ADC is shown in Fig. 12. The active circuits measure 900 μm $\times$ 700 μm.

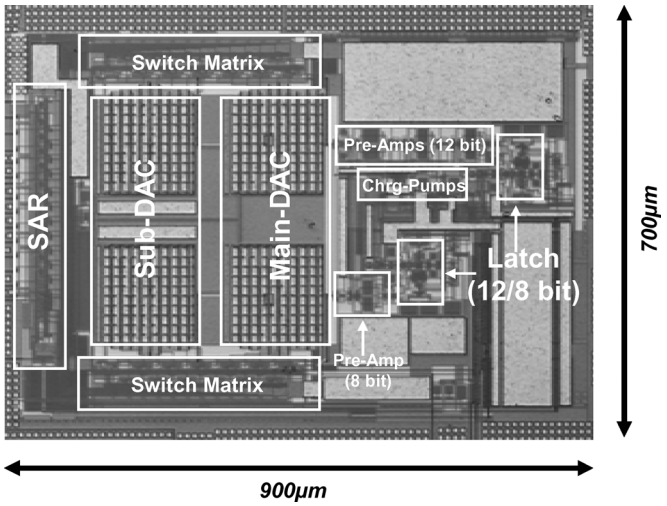

Fig. 12. Micrograph of entire ADC prototype.

TABLE I
ADC PERFORMANCE SUMMARY

	8 Bit Mode	12 Bit Mode
Process	0.18µm CMOS, National Semiconductor	
Area	900µm X 700µm	
Voltage Supply	1V	
Input Range	+/- 1V (differential) 1V (single-ended)	
Resolution	8 Bits	12 Bits
Maximum Sampling Rate	200kS/s	100kS/s
Analog Power	7.5µW@100kS/s	11µW@100kS/s
Digital Power	6.2µW@100kS/s	8.8µW@100kS/s
Reference Power	5.3µW@100kS/s	5.3µW@100kS/s
Total Power	19µW @ 100kS/s	25µW @ 100kS/s
INL/DNL	0.68LSB/0.66LSB	0.19LSB/0.16LSB
SNDR (at Nyquist)	49.7dB (f_{in}=100kHz)	65.3dB (f_{in}=50kHz)
SFDR (at Nyquist)	63.2dB (f_{in}=100kHz)	71dB (f_{in}=50kHz)

The remainder of this section describes how the prototype was tested and its measured performance. A summary of the performance is also provided in Table I.

A. Static Linearity

The code density test [12] was conducted using a full-swing, differential sinusoidal input with amplitude of 1 V. To test the 12-bit mode linearity, a sampling rate of 100 kS/s was used, and the frequency of the input signal was 111.381 Hz. Approximately 4 million samples were taken (30 records each with a size of 131 072). The offset of the ADC, determined using the method described in [13], is better than 830 μV among the tested parts. Fig. 13(a) shows the differential nonlinearity (DNL) and integral nonlinearity (INL) with respect to the output code. The maximum DNL is $+0.58$LSB$/-0.66$LSB, while the maximum INL is $+0.68$LSB$/-0.56$LSB. Fig. 13(b) shows the DNL and INL in 8-bit mode. Here, the sampling rate was 200 kS/s, and the input frequency was, once again, 111.381 Hz. The maximum DNL is $+0.16$LSB$/-0.14$LSB, while the maximum INL is $+0.19$LSB$/-0.16$LSB.

B. Dynamic Noise and Linearity

The signal-to-noise-plus-distortion ratio (SNDR) of the ADC was derived using tone testing. Then, from (8), the effective numbers of bits (ENOB) was determined.

$$\text{ENOB} = \frac{\text{SNDR (dB)} - 1.76}{6.02}. \tag{8}$$

Fig. 14(a) shows the ENOB of this ADC with respect to the input frequency. In 12-bit mode, the ADC samples at 100 kS/s and achieves an ENOB of 10.55 bits (65.3 dB SNDR) at its Nyquist rate. In 8-bit mode, the ADC samples at 200 kS/s and achieves an ENOB of 7.96 bits (49.8 dB SNDR) at its Nyquist rate.

A fast Fourier transform (FFT) of the ADC at near-Nyquist operation in 12-bit mode is shown in Fig. 14(b) $(F_S = 100\,\text{kS/s})$ and that in 8-bit mode is shown in Fig. 14(c) $(F_S = 200\,\text{kS/s})$. Odd—order harmonics are clearly visible. This distortion is due to the nonlinearity of the input switch resistance. The SFDR in 12-bit mode is 71 dB.

C. Power Consumption

At the highest 12-bit performance point, corresponding to 100 kS/s, the ADC core (not including I/O) consumes 25 μW from the 1-V supply. At the highest 8-bit performance point, corresponding to 200 kS/s, the ADC core consumes 39 μW from the 1-V supply. In both resolution modes, the power consumption decreases linearly towards zero as the sampling rate is reduced. In 12-bit mode, the power is measured to be approximately 200 nW at 500 S/s.

V. CONCLUSION

A Nyquist-rate ADC, operating from a 1-V supply, has been presented. The SAR architecture, which allows for a mostly passive implementation, was leveraged to achieve micro-power operation. Additionally, the architecture enabled efficient power management, allowing the power consumption to scale linearly as the sampling rate varies between 0 and 100 kS/s. For further power savings, 8-bit or 12-bit operating modes can be selected dynamically.

Improvements in the efficiency of the converter were achieved by employing a variety of techniques. The low supply voltage significantly reduced the overall power consumption of the ADC, while an offset compensating latch minimized the gain requirements of the comparator preamplifiers. Further, self-timing was used to ease their settling time. Finally, the common-mode rejection of the ADC was enhanced by appropriately sampling the input and auto-zeroing the comparator to the voltage of critical SAR decisions.

A widely used figure-of-merit (FOM) normalizes the ADC power consumption to the input bandwidth it can digitize and the dynamic range it achieves:

$$\text{FOM} = \frac{P}{2f_{\text{in}}2^{\text{ENOB}}}. \tag{9}$$

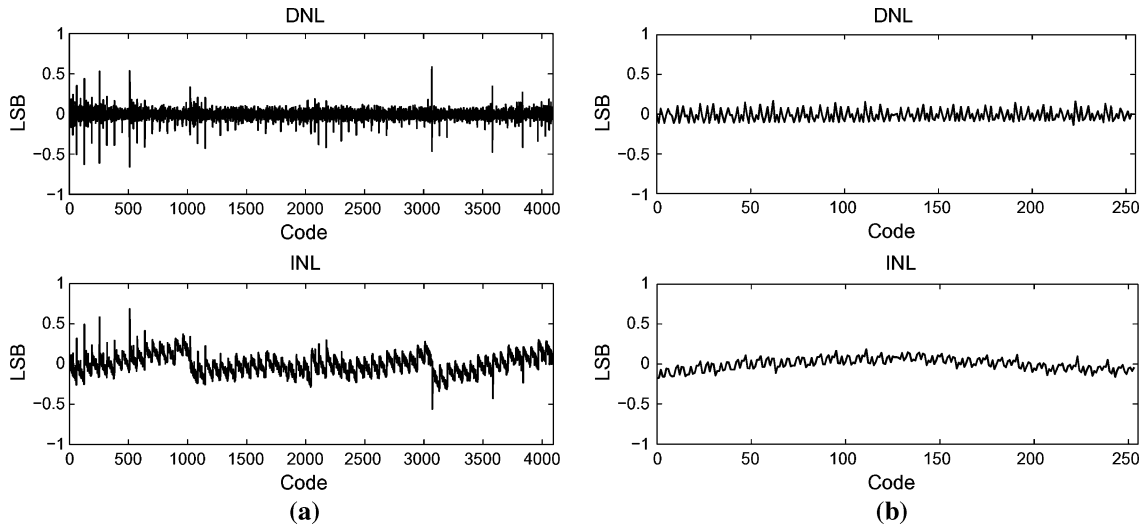

Fig. 13. DNL and INL of ADC in (a) 12 bit mode and (b) 8 bit mode.

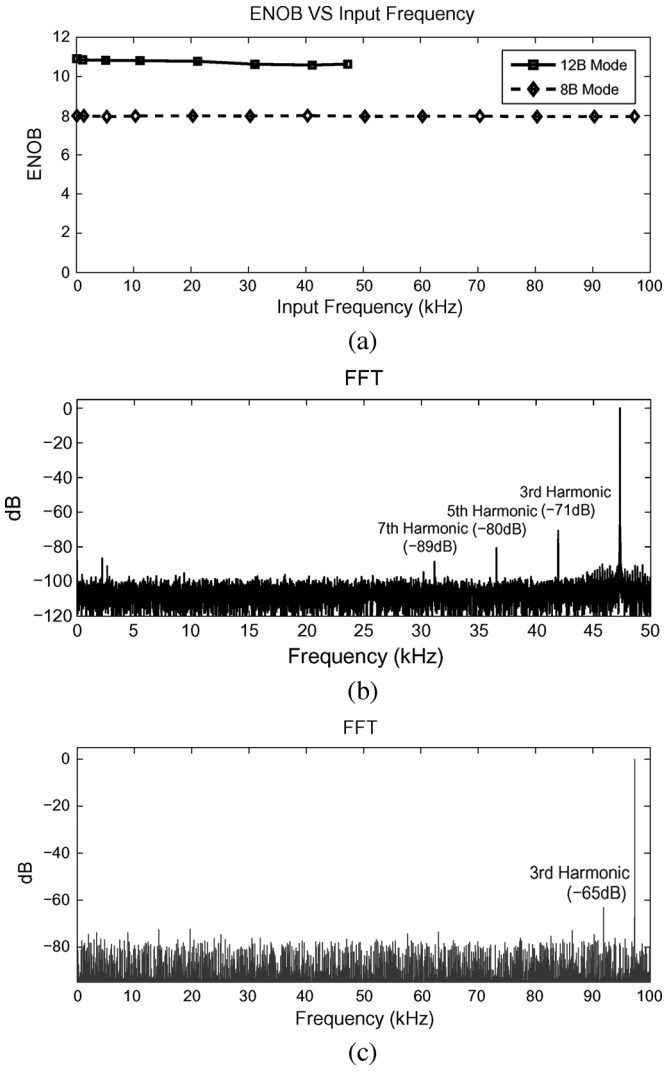

Fig. 14. Dynamic performance showing (a) ENOB versus input frequency for the ADC in 12-bit mode and 8-bit mode, (b) FFT of ADC in 12-bit mode with 47.3 kHz input tone, and (c) FFT of ADC in 8-bit mode with 97kHz input tone.

In its 12-bit mode, the ADC achieves a FOM of 165 fJ per conversion step.

ACKNOWLEDGMENT

The authors would like to thank Prof. H.-S. Lee, B. Ginsburg, and P. Holloway for their support and feedback.

REFERENCES

[1] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar, "Next century challenges: Scalable coordination in sensor networks," in *Proc. ACM MobiCom*, Aug. 1999, pp. 263–270.

[2] R. Min and A. Chandrakasan, "A framework for energy-scalable communication in high-density wireless networks," in *Proc. Int. Symp. Low Power Electronics and Design (ISLPED 2002)*, Aug. 2002, pp. 36–41.

[3] I. F. Akyildiz, Y. S. W. Su, and E. Cayirci, "Wireless sensor networks: A survey," *Comput. Netw.*, vol. 38, no. 4, pp. 393–422, Mar. 2002.

[4] N. Verma and A. Chandrakasan, "A 25 μW 100 kS/s 12 bit ADC for wireless micro-sensor applications," in *IEEE Int. Solid-State Circuits Conf. (ISSCC 2006) Dig. Tech. Papers*, San Francisco, CA, Feb. 2006, pp. 222–223.

[5] J. L. McCreary and P. R. Gray, "All-MOS charge redistribution analog-to-digital conversion techniues—Part I," *IEEE J. Solid-State Circuits*, vol. SC-10, no. 6, pp. 371–379, Dec. 1975.

[6] K. Gulati and H.-S. Lee, "A low-power reconfigurable analog-to-digital converter," *IEEE J. Solid-State Circuits*, vol. 36, no. 12, pp. 1900–1911, Dec. 2001.

[7] G. Geelen, E. Paulus, D. Simanjuntak, H. Pastoor, and R. Verlinden, "A 90 nm CMOS 1.2 V 10 bit power and speed programmable pipelined ADC with 0.5 pJ/conversion-step," in *IEEE Int. Solid-State Circuits Conf. (ISSCC 2006) Dig. Tech. Papers*, San Francisco, CA, Feb. 2006, pp. 214–215.

[8] I. Ahmed and D. Johns, "A 50 MS/s (35 mw) to 1 kS/s (15 μw) power scaleable pipeline 10 bit ADC with minimal bias current variation," in *IEEE Int. Solid-State Circuits Conf. (ISSCC 2005) Dig. Tech. Papers*, San Francisco, CA, Feb. 2005, pp. 280–281.

[9] G. Promitzer, "12-bit low-power fully differential noncalibrating successive approximation ADC with 1 MS/s," *IEEE J. Solid-State Circuits*, vol. 36, no. 7, pp. 1138–1143, Jul. 2001.

[10] J.-T. Wu and B. A. Wooley, "A 100 MHz pipelined CMOS comparator," *IEEE J. Solid-State Circuits*, vol. 23, no. 6, pp. 1379–1385, Dec. 1988.

[11] P. J. Lim and B. A. Wooley, "An 8-bit 200-MHz BiCMOS comparator," *IEEE J. Solid-State Circuits*, vol. 25, no. 1, pp. 192–199, Feb. 1990.

[12] J. Doernberg, H.-S. Lee, and D. A. Hodges, "Full-speed testing of A/D converters," *IEEE J. Solid-State Circuits*, vol. SC-19, no. 6, pp. 820–827, Dec. 1984.

[13] *IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters*, IEEE Std 1241-2000.

Naveen Verma (S'04) received the B.A.Sc. degree in electrical and computer engineering from the University of British Columbia, Vancouver, BC, Canada, in 2003, and the M.S. degree from the Massachusetts Institute of Technology, Cambridge, MA, in 2005. He is currently pursuing the Ph.D. degree at the Massachusetts Institute of Technology.

His research interests include low-power mixed signal circuits in the areas of analog-to-digital converters, SRAMs, and implantable biological systems.

Mr. Verma was the recipient of the Intel Foundation Ph.D. fellowship and the NSERC Postgraduate fellowship.

Anantha P. Chandrakasan (M'95–SM'01–F'04) received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994, respectively.

Since September 1994, he has been with the Massachusetts Institute of Technology, Cambridge, where he is currently the Joseph F. and Nancy P. Keithley Professor of Electrical Engineering. His research interests include low-power digital integrated circuit design, wireless microsensors, ultra-wideband radios, and emerging technologies. He is a coauthor of *Low Power Digital CMOS Design* (Kluwer Academic, 1995) and *Digital Integrated Circuits* (Pearson Prentice-Hall, 2003, 2nd edition). He is also a co-editor of *Low Power CMOS Design* (IEEE Press, 1998), *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000), *Leakage in Nanometer CMOS Technologies* (Springer, 2005), and *Subthreshold Design for Ultra-Low Power Systems* (Springer 2006).

Dr. Chandrakasan has received several awards including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, the 1999 Design Automation Conference Design Contest Award, and the 2004 DAC/ISSCC Student Design Contest Award. He has served as a technical program co-chair for the 1997 International Symposium on Low Power Electronics and Design (ISLPED), VLSI Design '98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Subcommittee Chair for ISSCC 1999–2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions Subcommittee Chair for ISSCC 2004–2007. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He serves on the SSCS AdCom and is the meetings committee chair.