A 7.1 mW 1 GS/s ADC With 48 dB SNDR at Nyquist Rate

Sedigheh Hashemi and Behzad Razavi, *Fellow, IEEE*

***Abstract*—A two-stage pipelined ADC employs a double-sampling residue amplifier, two interleaved precharged DACs, and a new calibration scheme to correct for residue gain error, offset, and nonlinearity. The coarse and fine stages are implemented as flash ADCs incorporating several techniques to reduce their power, complexity, and kickback noise. Realized in 65 nm CMOS technology and sampling at 1 GHz, the prototype achieves an SNDR of 48 dB at the Nyquist rate and exhibits an FOM of 25 fJ/conversion-step while drawing 7.1 mW from a 1 V supply.**

***Index Terms*—Double-sampling, nonlinearity, offset calibration, pipelined ADCs, precharged DAC.**

I. Introduction

MOST pipelined analog-to-digital converters (ADCs) employ multiple 1.5-bit stages for the sake of speed, modularity, and simplicity of the design. In recent generations, it has been recognized that a multi-bit first stage improves the first residue amplifier's settling speed and relaxes its output swing requirement [1]–[5], but most of these architectures opt for low-resolution stages after the first.

One approach to simplifying a pipelined architecture is to reduce the number of stages to only two, with the hope that the use of a single residue amplifier ultimately translates to lower power consumption. This paper describes a two-stage 1 GHz ADC that incorporates 33 comparators and one open-loop differential amplifier to achieve a resolution of 9 bits with a figure of merit of 25 fJ/conversion-step [6]. We propose a new calibration technique that corrects for various comparator and amplifier imperfections. We also introduce an interleaved precharged digital-to-analog converter (DAC) and a new method of reducing timing mismatches in bootstrapped sampling switches.

Section II provides the background for this work and Section III describes the ADC architecture. Section IV deals with the calibration techniques and Section V with the circuit details. Section VI presents the experimental results.

Manuscript received December 08, 2013; revised February 27, 2014; accepted March 05, 2014. Date of publication April 07, 2014; date of current version July 21, 2014. This paper was approved by Guest Editor Ken Suyama. This work was supported by the DARPA HEALICs program and Realtek Semiconductor.

The authors are with the Electrical Engineering Department, University of California, Los Angeles, CA 90095-1594 USA (e-mail: sedigheh@ee.ucla.edu, razavi@ee.ucla.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSSC.2014.2311812

II. Background

A. General Considerations

The use of two-step ADCs goes back to 1985 [7]. The design in [7] was perhaps the first to employ a resistor ladder as an interstage DAC, but such DACs were deemed slow and gradually replaced by capacitor arrays and morphed into the standard 1.5-bit multiplying DAC (MDAC).

A two-stage pipelined ADC faces three generic issues: 1) the settling speed of the interstage DAC, 2) the fine flash loading seen by the residue amplifier, and 3) the residue amplifier gain error and nonlinearity. In addition, the design must deal with the offset voltages and kickback noise of the comparators in both stages.

The partitioning of the resolution between the stages also merits some remarks. If the first-stage comparators and the residue amplifier sample the input simultaneously (as in a SHA-less front-end), then the resolution of the first stage is limited by a) the sampling timing mismatch between the coarse comparators and the amplifier, b) the offset and thermal noise of the comparators, c) the kickback noise of the comparators, and d) the speed of the interstage DAC.[1] The second stage must deal with the second and third effects as well as the input capacitance of the comparators.

In this work, a resolution of 5 bits is allotted to each stage, with 1 bit of redundancy to relax some of the above issues. Comparator offsets and kickback noise are also corrected to ease the speed–resolution–power trade-offs.

B. Precharged DACs

Resistor-ladder DACs are generally considered slow because their switch resistance, their load capacitance, and the capacitance of the switches themselves give rise to a long time constant—even if the ladder power dissipation is unimportant. For example, a 5-bit ladder driven by a 1-of-n code must drive the load capacitance and the capacitance of 32 switches through the on-resistance, $R_{\rm on}$, of one switch and the Thevenin resistance of the ladder. The issue is particularly acute for common-mode (CM) levels around $V_{\rm DD}/2$, at which even a complementary switch exhibits a high $R_{\rm on}$.

In an ADC environment, it is possible to reduce the settling time of resistor-ladder DACs by precharging the output node

[1]For every doubling of the first stage's resolution, the DAC settling time roughly doubles due to device and routing parasitics. In this design, the settling reaches a significant fraction of the clock cycle for a resolution of above 5 bits.

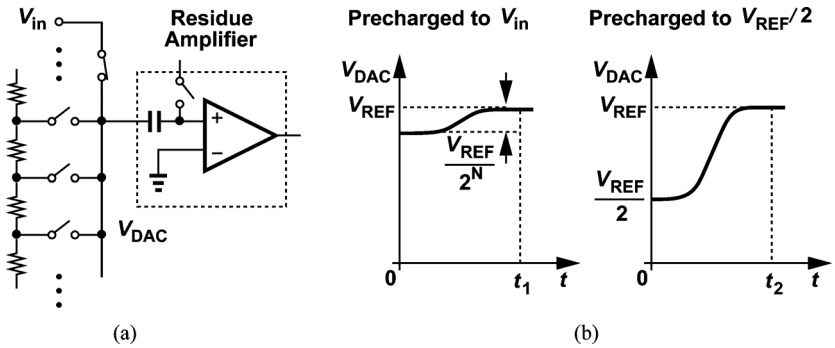

(a) (b)

Fig. 1. (a) Precharged resistor-ladder DAC, and (b) settling performance of a DAC when precharged to input or to $V_{\rm REF}/2$ (precharge occurs before $t = 0$).

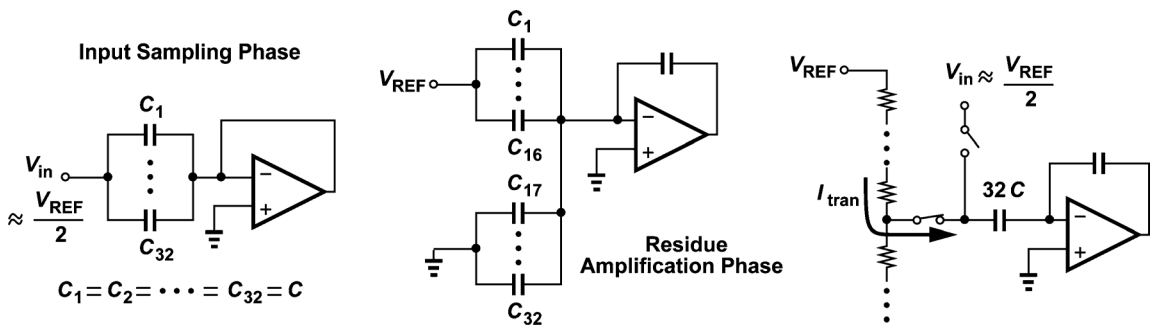

(a) (b)

Fig. 2. Transient currents drawn by (a) capacitor DAC, and (b) precharged DAC.

[4]. Illustrated in Fig. 1(a) for a SHA-less front-end, the idea is to precharge the DAC output node to $V_{\rm in}$ in the sampling mode so that, during conversion, $V_{\rm DAC}$ begins from a voltage close to its final value. This stands in contrast to the ladder DACs used in [7]–[9], which do not utilize precharging. In order to formulate the speed advantage afforded by precharging, we note that the worst-case settling occurs if the sampled voltage differs from the final value by 1 LSB of the DAC. As shown in Fig. 1(b) for an N-bit DAC with a full-scale voltage of $V_{\rm REF}$, the time necessary for the output to reach within $(1-\varepsilon)V_{\rm REF}$, where ε is the relative settling error, is given by

$$t_1 = -\tau \ln(2^N \varepsilon). \tag{1}$$

For example, if $N = 5$ and $\varepsilon = 0.25$ LSB of a 9-bit system, then $t_1 \approx 4.1\tau$. Without precharging to $V_{\rm in}$, on the other hand, one can only choose $V_{\rm REF}/2$ as the best estimate and set $V_{\rm DAC}$ to this value in the sampling mode. The worst-case settling time in this case is equal to

$$t_2 = -\tau \ln(2\varepsilon) \tag{2}$$

which, for the above condition, reaches 6.9τ. Thus, the precharge-to-input operation boosts the DAC speed by about a factor of 1.7. The DAC settling speed can be further relaxed by interleaving, as explained in Section III-C.

Precharged resistor-ladder DACs also offer an interesting advantage over capacitor DACs: the precharge-to-input operation considerably reduces the transient currents drawn from the reference. To explore this point, we consider a 5-bit front-end and examine the total charge drawn from the reference as the ADC goes from the sampling mode to the conversion mode. As shown for a capacitor DAC in Fig. 2(a), with $V_{\rm in} \approx V_{\rm REF}/2$, 16 of the unit capacitors switch to $V_{\rm REF}$ according to the sub-ADC's decision, thereby drawing a total charge of $8{\rm CV}_{\rm REF}$ from the reference. The precharged DAC, on the other hand, requires a worst-case voltage change of $V_{\rm REF}/2^5$ across 32C [Fig. 2(b)], thus pulling a charge equal to ${\rm CV}_{\rm REF}$ from $V_{\rm REF}$. It follows that an N-bit precharged resistor-ladder DAC reduces the reference transient noise by a factor of 2^{N-2} compared to a capacitor DAC. This improvement factor applies to fully differential topologies as well.

Fig. 3 plots the simulated transient currents drawn from the reference for the two cases. Here, 32C = 200 fF (dictated by kT/C noise), the total ladder resistance is 2 kΩ, the analog input is a full-scale sinusoid at 490 MHz, and the clock frequency is 1 GHz. We observe that the capacitor DAC currents are substantially higher. (The two DACs are designed for equal settling times.)

C. Device Stress in Comparators

SHA-less ADC front-ends must ensure reasonable matching between the voltage sampled by the first flash ADC and that sampled by the input capacitor of the multiplying DAC (MDAC) or the residue amplifier. For this reason, the comparator topology shown in Fig. 4(a) is often utilized [10] as its sampling network can track that of the main sampling capacitor. The switches alternately sample the input and the reference, presenting $(V_{\rm in+} - V_{\rm in-}) - (V_{R1} - V_{R2})$ to the comparator.

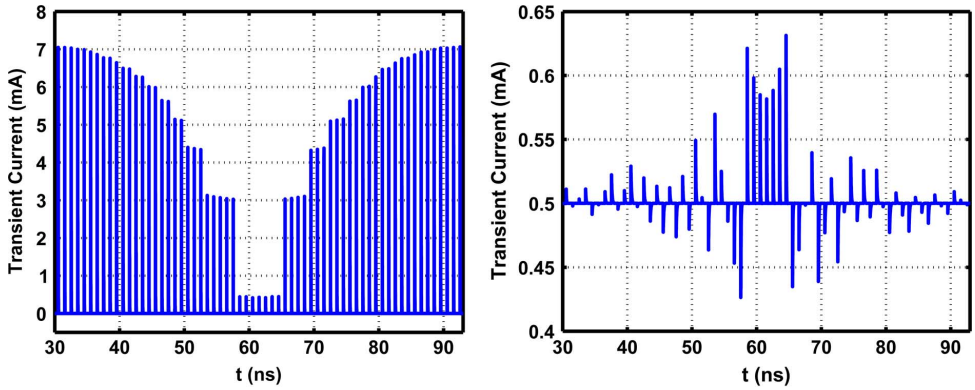

Fig. 3. Simulated transient currents drawn from the reference by (a) a capacitor DAC, and (b) a precharged DAC.

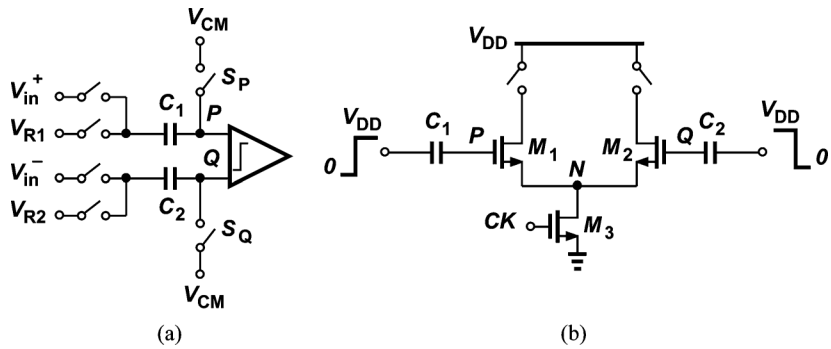

Fig. 4. (a) Sampling network to reduce timing mismatch between coarse ADC and MDAC, and (b) device stress in a StrongArm comparator when operating with rail-to-rail input swings.

This configuration also accommodates rail-to-rail inputs but at the cost of stressing the comparator's input transistors. To illustrate this point, we consider the arrangement depicted in Fig. 4(b), where the input stage of a StrongArm comparator follows C_1 and C_2. Let us assume, as an extreme case, that the differential full scale is equal to $\pm V_{\rm DD}$, i.e., $V_{\rm in}^+$ can be near 0 and $V_{\rm in}^-$ near $V_{\rm DD}$. Also, suppose that V_{R1} and V_{R2} are around $V_{\rm DD}$ and 0, respectively. As nodes P and Q are released from $V_{\rm CM}$ and the input switches connect C_1 and C_2 to V_{R1} and V_{R2}, respectively, V_P jumps from $V_{\rm CM}$ to $V_{\rm CM}+V_{\rm DD}$ and V_Q from $V_{\rm CM}$ to $V_{\rm CM}-V_{\rm DD}$. Consequently, when M_3 turns on and pulls N to zero, M_1 experiences a $V_{\rm GS}$ equal to $V_{\rm CM}+V_{\rm DD}$ and is stressed for about half a clock cycle. We also observe that if the differential pair is *not* clocked, the drain of M_2 remains at $V_{\rm DD}$ while its gate drops to $V_{\rm CM}-V_{\rm DD}$, stressing the transistor. In addition, switch S_P experiences a $V_{\rm DG}$ equal to $V_{\rm CM}+V_{\rm DD}$ when it is off.

The key result here is that the single-ended full scale applied to the above topology must remain below $V_{\rm DD}/2$ so as to avoid device stress. This issue evidently has not been recognized in prior work. The significance of this point becomes clear is Section III-B.

III. ADC Architecture

The proposed ADC exploits several architecture techniques, namely, interleaved precharged DACs, a double-sampling

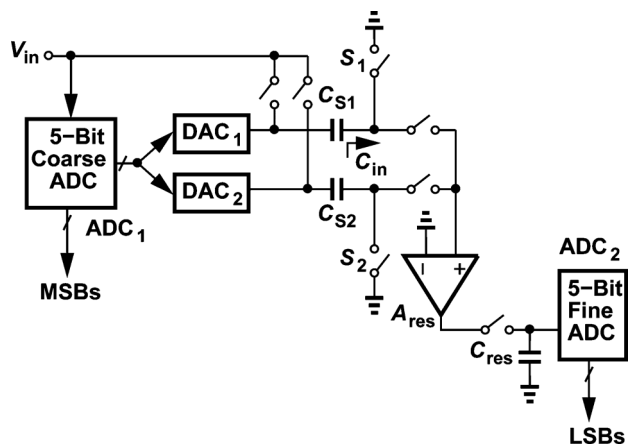

Fig. 5. Conceptual ADC architecture.

residue amplifier, and flash stages with polarity detection. In order to appreciate the role played by each method, we first present the architecture at a functional level and then delve into the details. We remark that, at a sampling rate of 1 GHz with realistic clock transitions and non-overlap times, the ADC must acquire and convert in about 950 ps.

A. Functional Description

Shown in Fig. 5 is a conceptual single-ended diagram of the system. The ADC consists of a coarse 5-bit flash stage $(\rm ADC_1)$,

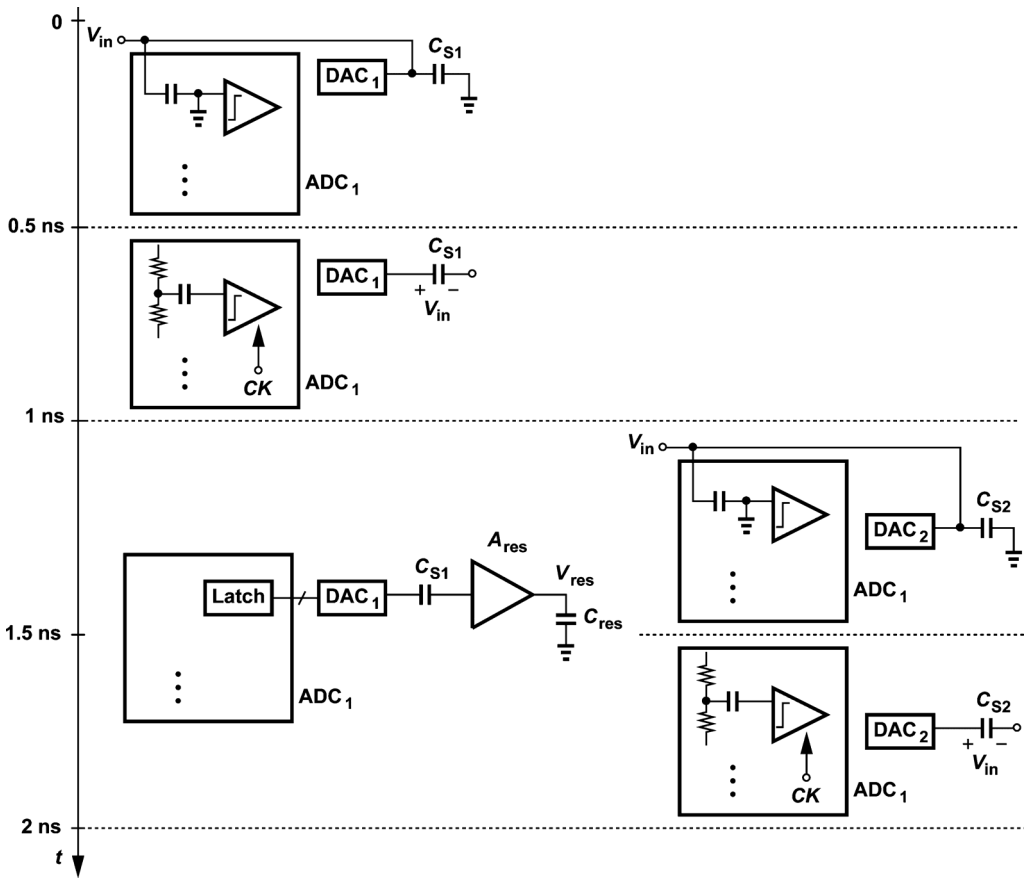

Fig. 6. Detailed operation of ADC blocks from 0 to 2 ns.

two time-interleaved DACs (DAC_1 and DAC_2), an open-loop residue amplifier (A_{res}), and a fine 5-bit flash stage (ADC_2). The pipelining occurs by sampling the output of A_{res} on C_{res}. The two flash stages operate at full clock rate and consume only dynamic power whereas the DACs and the amplifier have about one full period to settle and draw static currents.

The operation of the ADC is illustrated step by step in Fig. 6. In the first 0.5 ns, the comparators in ADC_1 and C_{S1} sample the input. The output of DAC_1 is also precharged. In the second 0.5 ns, ADC_1 is clocked to perform coarse conversion while the input is held on C_{S1} and at the DAC output. In the next clock cycle (from 1 ns to 2 ns), ADC_1's decision (stored in latches that follow the comparators) drives DAC_1 and the residue amplifier. Thus, DAC_1 and A_{res} have an entire clock period to settle.[2] Moreover, in this cycle, ADC_1, DAC_2, and C_{S2} acquire the input for the first 0.5 ns, and ADC_1 makes a decision in the second 0.5 ns. At the end of this period, C_{res} holds V_{res}, ADC_2 begins to convert, and DAC_2 and A_{res} begin to settle.

The proposed architecture allocates half a clock cycle to each of the flash ADCs, allowing compact, low-power implementations, and about one clock cycle to $\mathrm{DAC}_{1,2}$ and A_{res} settling, relaxing their speed–power trade-offs. These points are studied in greater detail below.

[2]Not shown in Fig. 6 for clarity, ADC_2 samples the residue voltage in the second half of each cycle (from 1.5 ns to 2 ns) and converts in the following half cycle to perform pipelining.

B. Coarse ADC

A 5-bit flash stage that must respond in less than 500 ps potentially consumes high power, presents a large input capacitance, and generates a great deal of kickback noise. In this work, these issues are mitigated by a "sliding" architecture and through the use of small transistors within the comparators along with offset cancellation.

The objective of sliding is to employ 16 comparators but resolve 5 bits. This is accomplished by sensing the polarity of the differential input and, accordingly, switching the reference inputs of the comparators to the appropriate half of the reference ladder. Fig. 7 shows the realization of the coarse 5-bit ADC. The circuit employs one comparator as a "polarity detector" (PD) and another 15 as a quantizer. After the sampling of V_{in} is completed, the PD is clocked to determine whether $V_{\mathrm{in}} > 0$ or $V_{\mathrm{in}} < 0$. Accordingly, the reference levels of the quantizer slide to the top or bottom half of the full scale, i.e., to $[V_{\mathrm{REF}}^{+} V_{\mathrm{mid}}]$ or $[V_{\mathrm{mid}} V_{\mathrm{REF}}^{-}]$, respectively. The PD decision and the sliding take about 170 ps, after which the 15 comparators are clocked. The flash decision is then converted to a 1-of-n code and stored in latches whose outputs are gated by DAC_1 Enable and DAC_2 Enable for the purpose of interleaving the two DACs.

The proposed flash architecture offers several advantages over conventional designs. The use of a polarity detector halves

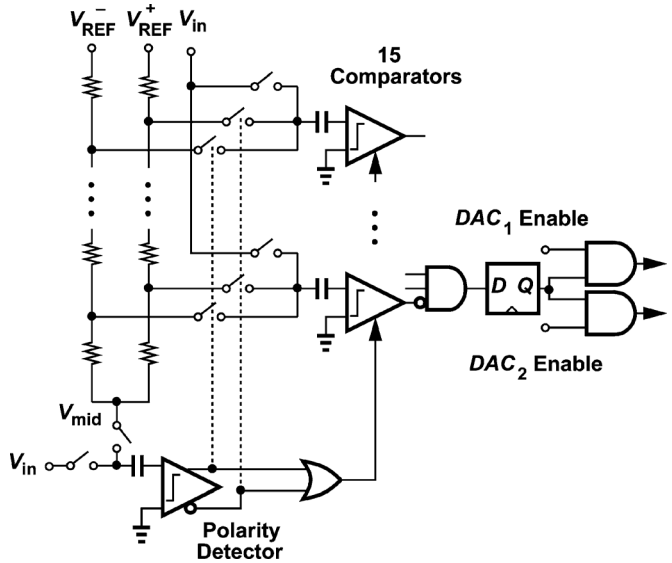

Fig. 7. Realization of the coarse flash stage (ADC_1).

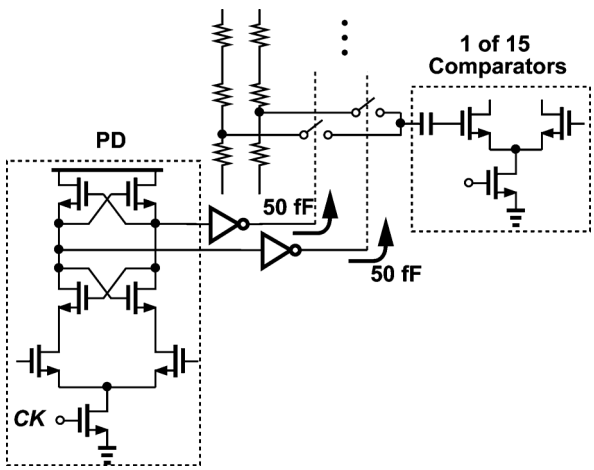

Fig. 8. Polarity detector using StrongArm topology to drive the sliding switches.

the number of comparators and hence their power consumption, input capacitance, and kickback noise. Furthermore, the reference sliding scheme presents only half of the full-scale swing to the capacitors' input devices. For example, if V_{in}^+ in Fig. 4(a) is near 0, the PD and the sliding operation ensure that V_{R1} is less than $V_{\mathrm{REF}}/2$, limiting V_P to less than V_{DD} if $V_{\mathrm{CM}} \approx V_{\mathrm{DD}}/2$. This attribute of our approach stands in contrast to the architecture in [11], which also employs a PD but avoids clocking half of the comparators rather than sliding their references. Both architectures incur a speed penalty due to the polarity detector and its load capacitance.

The PD in Fig. 7 must satisfy certain speed and precision requirements. Driving 32 switches and interconnects, i.e., 50 fF of capacitance, the PD employs a StrongArm comparator and one buffer (Fig. 8) and turns on the desired sliding switches in approximately 100 ps. The inputs of the 15 comparators then take only 70 ps to settle to their corresponding references because their sampling capacitors now appear in series with the gates

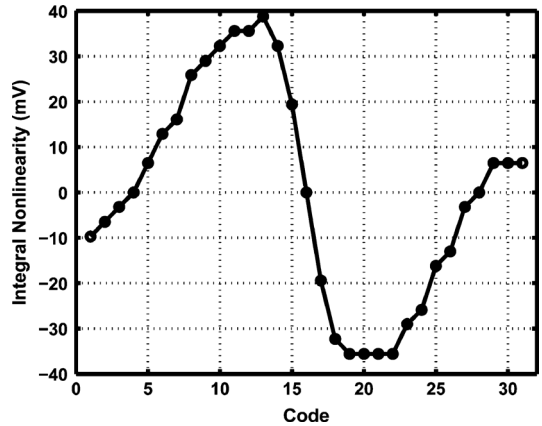

Fig. 9. Kickback-noise-induced INL of the coarse stage.

of input transistors. Thus, the comparators are clocked about 170 ps after the PD.

The coarse stage incorporates the following values: a ladder resistance of 2 kΩ, a StrongArm comparator design exhibiting a raw 3σ offset of 36 mV and consuming 40 μW at 1 GHz, and a sampling capacitor of 18 fF at the input of each comparator. The value of this capacitor is dictated by two factors: the timing mismatch with respect to the main sampling path, and the signal attenuation caused by this capacitor at the input of the coarse comparator.

The offset voltage of coarse comparators must also be managed. With one bit of redundancy, a maximum canceled offset of 4 LSB[3] is budgeted, thus leaving an ample margin of 4 LSB for the timing mismatch and comparator thermal noise.

The choice of the reference ladder resistance is governed by two factors: 1) the time constant associated with the comparators' input network, and 2) the kickback noise. In this design, the former is dominated by the on-resistance of the switches and amounts to 10 ps. The latter, on the other hand, creates a large signal-dependent error because the hard switching operations within the StrongArm comparator produce considerable kickback noise at its input. Fig. 9 plots the simulated kickback-noise-induced integral nonlinearity (INL) of the ladder for a total resistance of 2 kΩ. The error voltage reaches a maximum of 39 mV, severely tightening the error budget provided by 1 bit of redundancy. However, the calibration technique described in Section IV-A reduces the kickback noise along with the comparator offset to less than 4 LSB.

C. Interstage DAC

As mentioned in Section III-A, the ADC employs two time-interleaved precharged DACs. Shown in Fig. 10, the DACs share a single ladder, and their output nodes, V_{DAC1} and V_{DAC2}, are precharged to V_{in} alternately. Upon completion of the sampling mode, V_{in} is stored on the parasitic capacitance, C_{p1} or C_{p2} ($\approx$ 85 fF), providing a close estimate of the final DAC voltage. With a 2-kΩ ladder, the worst-case time constant at these nodes is less than 80 ps because C_{S1} or C_{S2} is in series

[3]Throughout the paper, LSB is the least significant bit of the overall ADC and equals 3.9 mV.

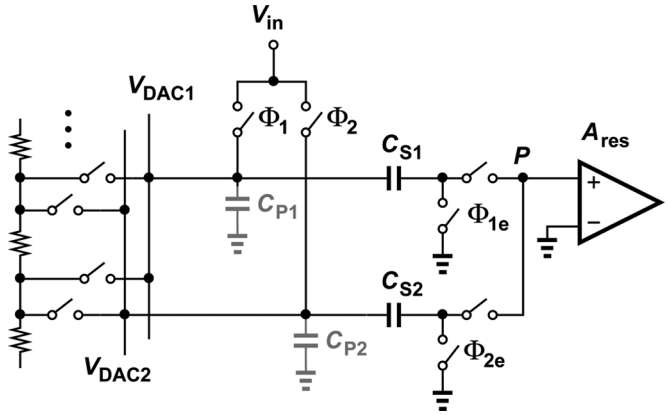

Fig. 10. Two interleaved precharged DACs operating with a double sampling scheme to enhance the speed.

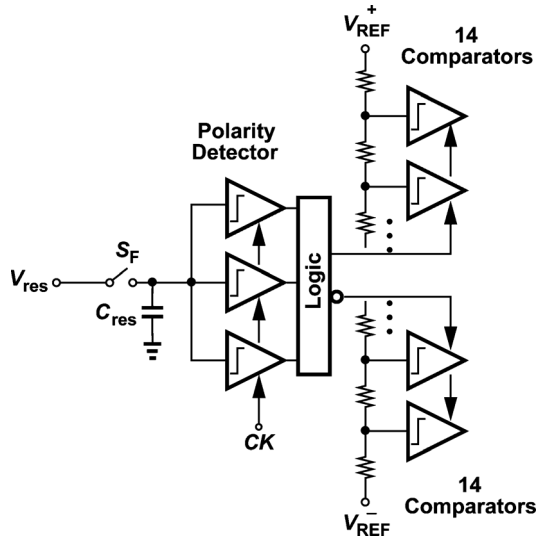

Fig. 11. Realization of the fine flash stage (ADC_2).

with the input capacitance of A_{res}, C_{in}. However, the DAC voltage is attenuated by a factor of 0.75 due to the voltage division between C_{S1} (or C_{S2}) and C_{in}. Since this attenuation also appears in V_{in}, it can be absorbed by A_{res} and calibrated (Section IV). To suppress dynamic errors resulting from the charge stored on node P in Fig. 10, a reset switch (not shown) shorts this node to its common-mode level for about 75 ps every cycle.

The DAC ladder ultimately determines the INL of the overall ADC and is realized as proposed in [12], achieving a linearity of about 11 bits. This ladder has 512 uniformly spaced taps, some of which are used for D/A conversion and some for calibration (Section IV). By virtue of precharging, the DAC output voltage requires about 320 ps to reach within 0.25 LSB of the its final value. The residue amplifier begins amplification simultaneously but has another 600 ps for complete settling.

While it is tempting to utilize the coarse ADC reference ladder for the DACs as well, the residual kickback noise on the former would significantly corrupt the latter. For this reason, two separate ladders are used.

D. Fine ADC

With the residue sampled on C_{res} in Fig. 5, the fine ADC has a half clock cycle for conversion. Depicted in Fig. 11, this stage draws upon the architectures in [11], [13] but with a 2-bit polarity detector so as to relax the noise and offset required of the PD's constituent comparators. The PD is clocked after residue sampling on C_{res} is completed, producing a decision in 150 ps that enables the top or bottom comparator bank. The power consumption and kickback noise are therefore halved.

IV. Calibration Techniques

The push for high speed and low power naturally calls for small devices, requiring additional methods of dealing with their imperfections. A key issue to be borne in mind is the hardware and power overhead that such methods impose. Most significant in this design are the offset and kickback noise of the comparators and the offset, gain error, and nonlinearity of the residue amplifier. This section presents two precision techniques that ultimately afford an FOM of 25 fJ/conversion-step at low frequencies and 34 fJ/conversion-step at the Nyquist rate.

A. Offset Cancellation

The ADC clocks a total of 33 comparators every 1 ns, requiring a response time of about 100 ps to 250 ps. In order to achieve a low power consumption, we can apply "linear" scaling [14] to a reference comparator design. That is, we begin with a design having a power consumption of P, an input offset of V_{OS}, and an input noise of V_n. We then scale down all of the comparator's transistor widths by a factor of α, thus reducing P to P/α, and raising V_{OS} to $\sqrt{\alpha}V_{\mathrm{OS}}$ and V_n to $\sqrt{\alpha}V_n$.[4]

With nearly-minimum-size transistors, the offset may demand a prohibitively large overhead in terms of the number of devices that appear in the signal path. For example, programmable capacitor arrays [15] attached to the internal nodes of the StrongArm comparator both degrade the speed *and* raise the power dissipation ($\propto fCV_{\mathrm{DD}}^2$). Such limitations prescribe an *upper* bound for the raw offset that the design can target. Moreover, foreground calibration (as in our work) is somewhat prone to temperature drifts, reaching diminishing returns if very small residual offsets are desired. This constraint places a *lower* bound on the residual offset that can be achieved.

In this work, the two flash stages employ a cancellation technique that resides entirely outside the comparator. Illustrated in Fig. 12(a), our approach calibrates comparator number j for a decision threshold of V_j as follows: 1) connect one input to V_j, 2) change the other input by means of DAC_j until the comparator output changes, and 3) freeze the DAC_j content. To minimize the effect of comparator noise, this procedure is repeated 10 times and the average value is chosen.

The dedicated DACs in Fig. 12(a) appear formidable, but they are greatly simplified if we recognize that the reference ladder itself can serve as DAC_j, DAC_{j+1}, etc. This approach is feasible

[4]The comparator speed remains unchanged so long as the load capacitance can also be scaled down by a factor of α.

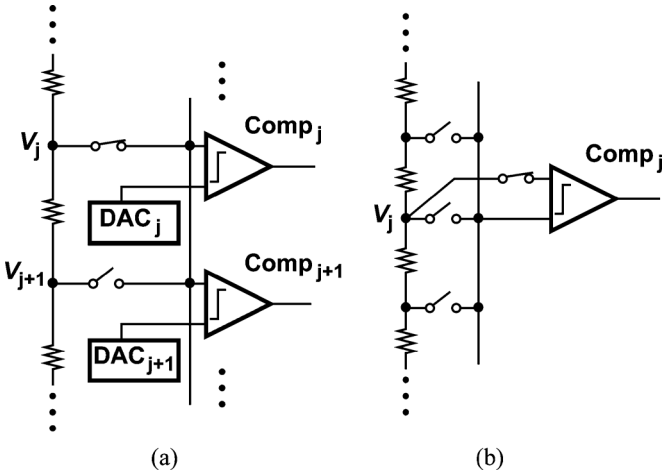

Fig. 12. (a) Comparator offset calibration scheme, and (b) actual implementation.

because each comparator reference voltage must depart from its ideal value by only an amount equal to the raw offset. In other words, as shown in Fig. 12(b), one input of comparator j can slide along a narrow range of the ladder voltages so as to compensate for the offset. With 64 ladder taps and differential implementation, the offset of coarse comparators is reduced to 4 LSB (16 mV). For fine comparators, an additional method is used so as to achieve a finer calibration resolution. The input differential pair of the comparator is decomposed into two halves, and one of the pair's reference voltage is adjusted during calibration. Thus, with a ladder voltage step size of 4 LSB (differential) and a residue gain of 5.6, the offset of fine comparators falls below 0.4 LSB when referred to the input.

It is interesting to note that offset cancellation also removes the error due to kickback noise. This is because a) during offset cancellation, all of the comparators are clocked simultaneously, generating the same INL profile along the resistor ladder as that during actual operation, and b) the value of V_j sampled by the comparator in Fig. 12 is *not* corrupted by this INL profile because it is sampled by the input capacitors of comparators *before* they are clocked and generate kickback noise.

The input-referred noise voltages of the coarse and fine comparators are about 2.3 $\mathrm{mV_{rms}}$. The noise of the latter is divided by 5.6 when referred to the main input.

B. Residue Amplifier Calibration

As explained in Section V-B, the residue amplifier (including C_{S1} and C_{S2} in Fig. 5) exhibits a gain error of 21%, a peak INL of 8 mV, and an input 3σ offset of 9 mV. We propose a calibration method that removes these imperfections by "programming" the fine ADC.

Let us first consider only the gain error. As shown in Fig. 13(a), the residue does not match the full-scale range of the fine ADC, V_{nom}. This issue can be resolved by adjusting the decision thresholds of this stage such that they span the range 0 to V_{max} rather than 0 to V_{nom}.

In order to illustrate the calibration technique, we note that, if a known, precise voltage is applied to A_{res}, then the corresponding fine comparator must make the critical decision, and this can be ensured by adjusting that particular comparator's reference voltage. To account for the capacitive attenuation resulting from $C_{S1,2}$ and C_{in} in Fig. 5, the known voltage must be first sampled by each one of these two capacitors (in a double-sampling manner) and then applied to A_{res}. We thus arrive at the procedure depicted in Fig. 13(b): C_{S1} samples a precise ladder voltage, V_k, the residue amplifier generates $A_{\mathrm{res}}V_k$, and the reference voltage of the critical comparator is adjusted by DAC_k until this comparator trips. As with the coarse ADC, DAC_k is in fact embedded within the fine ADC reference ladder. This procedure is repeated for all of the fine comparators, covering a range of ± 36 LSB around the nominal voltage tap.

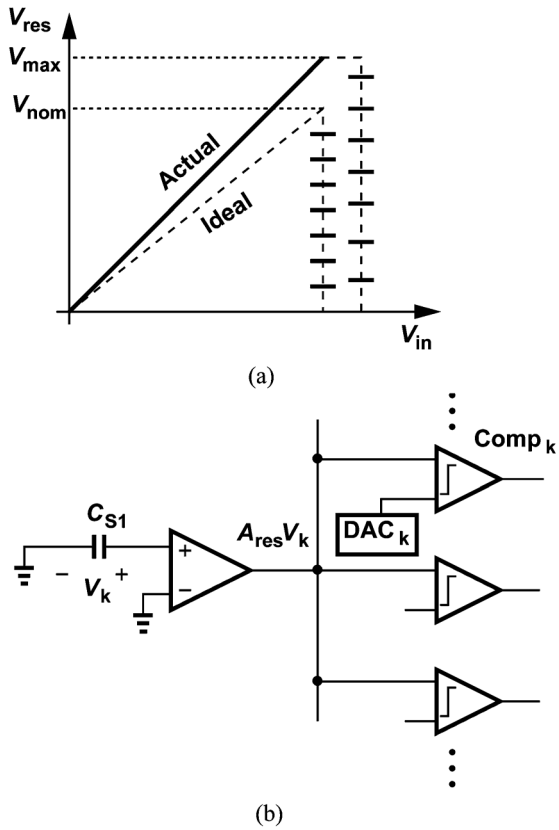

Fig. 13. (a) Residue voltage characteristic in the presence of amplifier gain error, and (b) residue amplifier calibration.

The foregoing description suggests that the proposed calibration technique also corrects the offset and nonlinearity of the residue amplifier and the offset of the fine comparators. This occurs because a) the amplified residue, $A_{\mathrm{res}}V_k$, contains the amplifier's offset and nonlinearity, and b) the decision thresholds of the fine ADC can be adjusted to absorb these imperfections. To see how the calibration linearizes this characteristic, let us consider the ideal situation as shown in Fig. 14(a), where the circles denote the fine ADC decision thresholds. Now, if the residue experiences nonlinearity but the thresholds remain unchanged [Fig. 14(b)], the input voltages at which decisions are made, V_1', V_2', and V_3', are displaced, producing DNL and INL. On the other hand, if we choose the fine thresholds so as

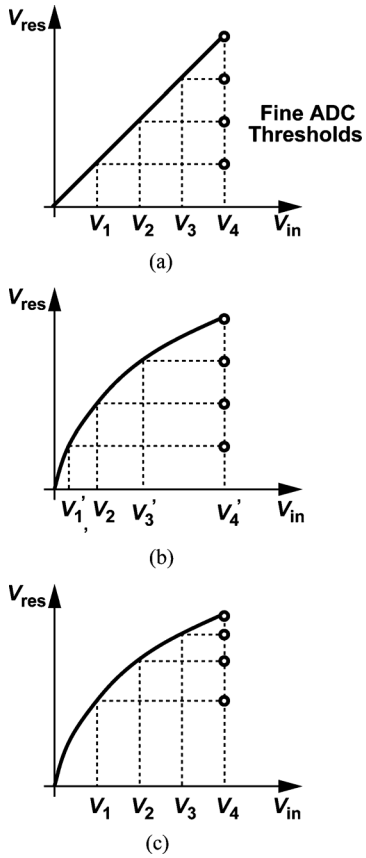

Fig. 14. (a) Ideal residue voltage characteristic, (b) distortion of input-referred threshold levels due to residue nonlinearity, and (c) linearizing residue characteristic by distorting fine ADC decision thresholds.

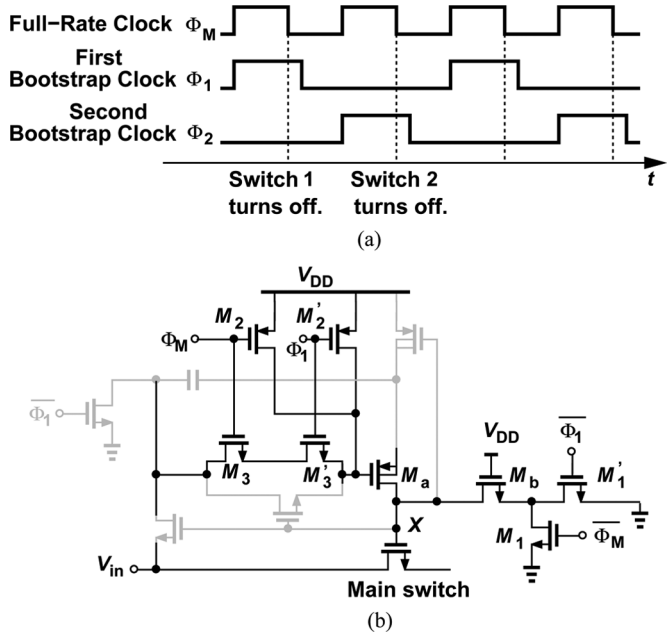

Fig. 15. (a) Falling edges of a full-rate master phase, ϕ_M, used to define sampling instances, and (b) modified bootstrapping circuit incorporating AND function so that ϕ_M overrides ϕ_1 and ϕ_2.

to trigger them at $V_{\rm in} = V_1$, V_2, and V_3 [Fig. 14(c)], then the overall characteristic becomes linear.

The arrangement in Fig. 13(b) inherently calibrates for the offset and nonlinearity errors as well. The value of V_k is chosen equal to 1 LSB (provided by the ladder), the corresponding comparator output is monitored, and $\rm DAC_{\it k}$ is incremented until that comparator's output flips. Next, $\rm V_{\it k}$ is raised to 2 LSB and the procedure is repeated for comparator $\rm k+1$. The calibration continues for $\rm V_{\it k}$ values up to 31 LSB.

The 31 tap voltages necessary for calibration are provided by the resistor ladder used in the interleaved DACs. The precision of the ladder ultimately limits the calibration accuracy. The switches tied to these taps are relatively small as the calibration is performed at low speeds.

The foreground calibration technique proposed here assumes acceptably small temperature drifts in the residue amplifier and comparator imperfections. Simulations indicate that if the amplifier and fine ADC are calibrated at 27°C and the temperature rises to 75°C or falls to 0°C, the maximum INL remains below 5.7 mV (1.45 LSB) and −3.5 mV (0.9 LSB), respectively. In addition, supply variations of ±5% raise the maximum INL by ±0.6 LSB.

V. Building Blocks

The proposed ADC consists of 47 comparators (33 of which are clocked in each cycle), two resistor ladders, sampling capacitors and bootstrapped switches, and a residue amplifier. In this section, we describe the details of the bootstrapping circuit and the amplifier.

A. Bootstrapping Circuit

The double-sampling technique operating on C_{S1} and C_{S2} in Fig. 5 relies on a precise 50% duty cycle for the clock so as to ensure uniform sampling of the input signal. Using (36) in [16], we estimate that a timing mismatch of 0.6 ps between consecutive samples degrades the signal-to-noise ratio by 1 dB at the Nyquist rate. If operating with both the rising and falling edges of a half-rate clock, a double-sampling circuit becomes sensitive to the duty-cycle distortion. It is possible to utilize only the falling (or rising) edge if the clock runs at the full rate, and half-rate "predictive" pulses alternately route this edge to the sampling switches (S_1 and S_2 in Fig. 5) [17]. Here, we propose a method that is applicable to bootstrapped switches.

Our objective is to turn off the main sampling switch on the falling edge of the full-rate clock while the bootstrapping circuit operates at half rate. Fig. 15(a) depicts the waveforms fulfilling this goal: in addition to the devices controlled by the half-rate clocks, ϕ_1 and ϕ_2, we employ other transistors that turn off the main switches on each falling edge of ϕ_M. This is accomplished by implementing an AND function within the bootstrapping network so that ϕ_M can override ϕ_1 and ϕ_2 on its falling edges.

Fig. 15(b) shows the bootstrapping circuit [18] with our overriding devices added. The main switch would be ordinarily turned off by M_1' but is now disabled by M_1 according to ϕ_M. Also, since during the turn-off process, M_a must disconnect the bootstrap capacitor from node X, we insert transistors M_2 and M_3 in parallel and series with the original devices, M_2' and M_3', respectively, so as to turn off M_a by ϕ_M. The timing

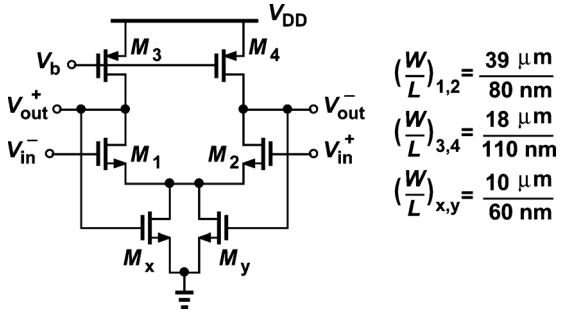

Fig. 16. Residue amplifier topology.

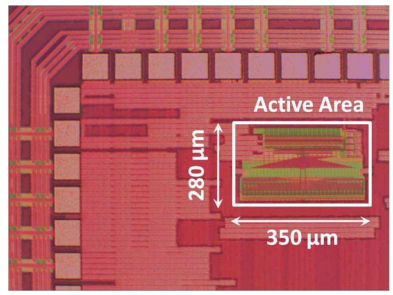

Fig. 18. ADC die photograph.

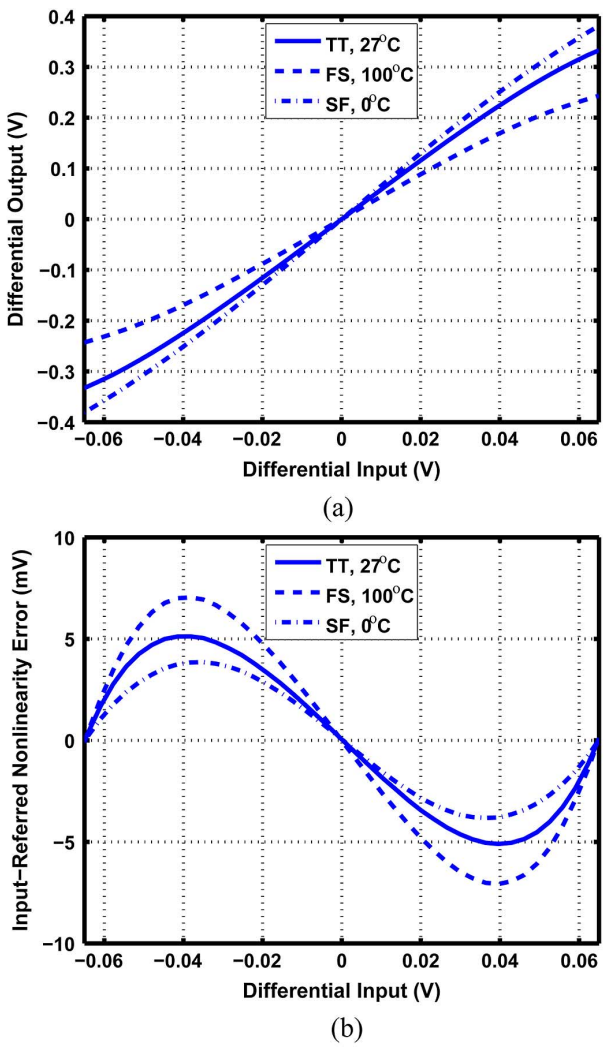

Fig. 17. (a) Residue voltage range, and (b) its nonlinearity (referred to the input).

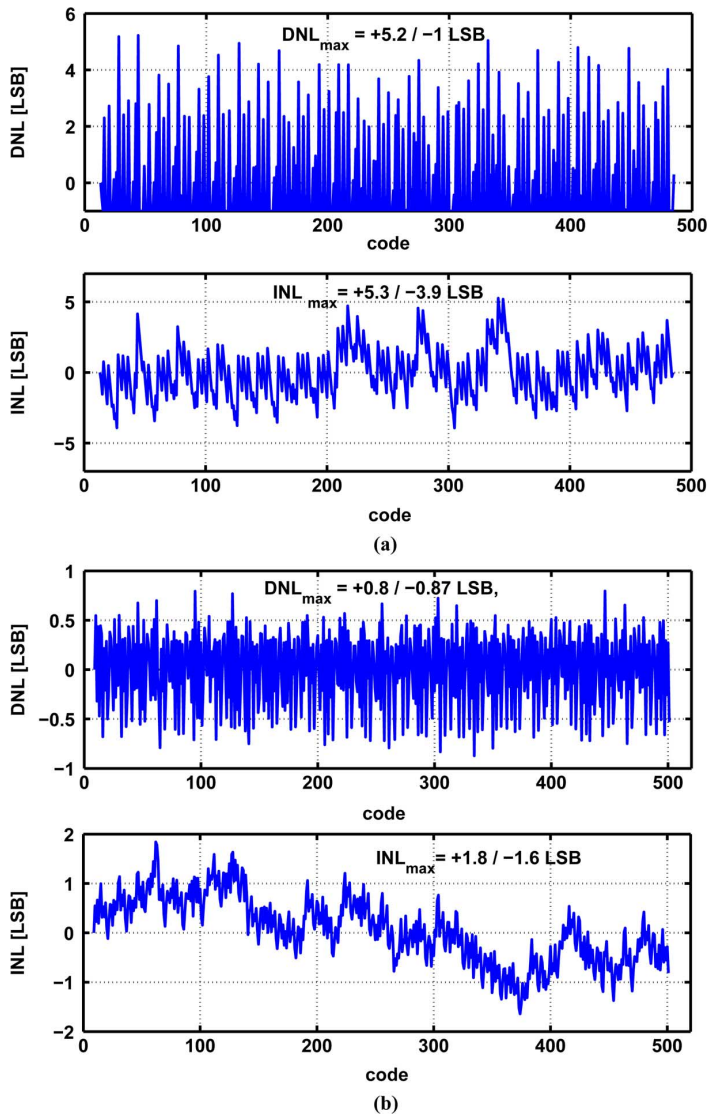

Fig. 19. Measured DNL and INL (a) before, and (b) after calibration.

mismatch now arises only from M_1, M_b, and the main switch. The experimental results reveal a mismatch-induced spur level of -61 dB, suggesting a residual timing error of 0.55 ps.

B. Residue Amplifier

With 5 bits resolved in the first stage and the use of double sampling, the residue amplifier, $A_{\rm res}$, needs to generate only a moderate output swing (± 350 mV) and settle in about 600 ps. These relaxed requirements are indeed essential here because the total interconnect and MOS capacitance seen by $A_{\rm res}$ reaches 200 fF,[5] demanding a high power otherwise.

The low gain of one-stage amplifiers makes negative feedback marginally useful in our work. For this reason, we employ the open-loop topology shown in Fig. 16. The circuit provides a nominal voltage gain of 7.5 with an input-referred noise of

[5]Note that the unselected comparator bank in the fine ADC of Fig. 11 still presents significant capacitance to $A_{\rm res}$.

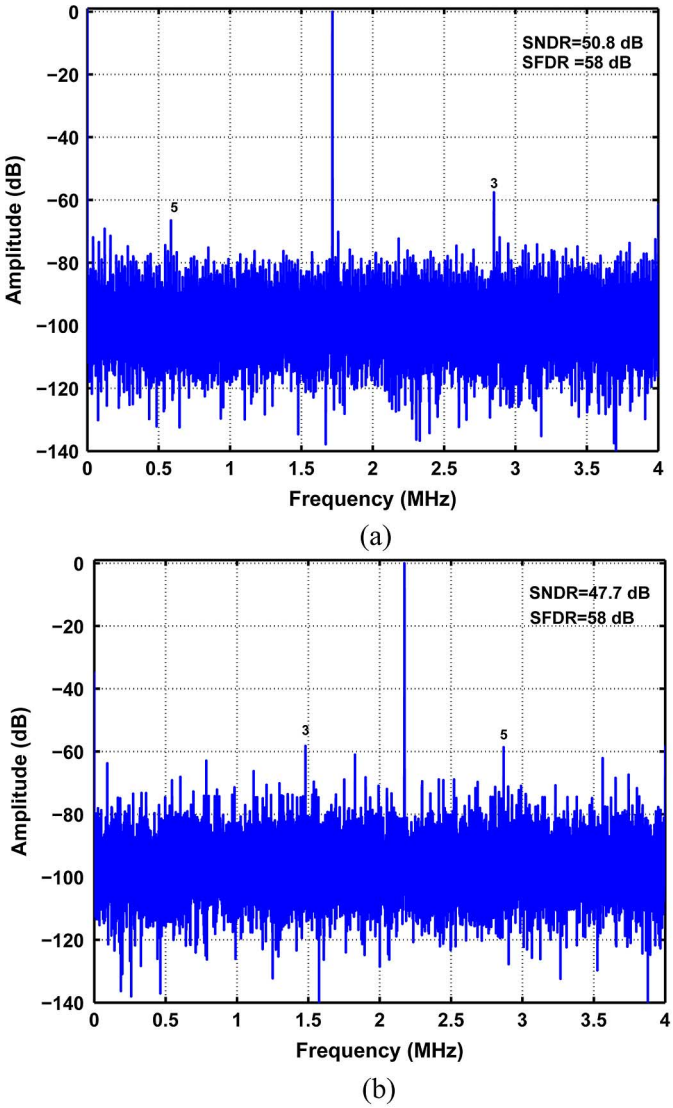

Fig. 20. ADC output spectrum at a sampling rate of 1 GHz and an input frequency of (a) 1.7 MHz and (b) 490 MHz (downsampled by a factor of 125). (Components labeled 3 and 5 are the aliased third and fifth harmonics.)

110 μV_{rms} while drawing a supply current of 2.5 mA. (Due to the attenuation resulting from $C_{S1,2}$ and C_{in} in Fig. 5, the overall residue gain drops to 5.6.) Operating in the triode region, transistors M_x and M_y serve as a simple common-mode feedback circuit with no tail current source.

Fig. 17 plots the simulated input-output characteristic and nonlinearity profile of the amplifier for the nominal case (TT, 27°C) and two worst cases (FS, 100°C and SF 0°C). The former demands that the fine ADC comparators provide a correction range of ±75 mV around the nominal characteristic. We recognize that this much correction also suffices for nonlinearity calibration.

VI. Experimental Results

The prototype ADC has been designed and fabricated in standard 65 nm CMOS technology in an active area of 350 μm × 280 μm. Fig. 18 shows the die photograph.

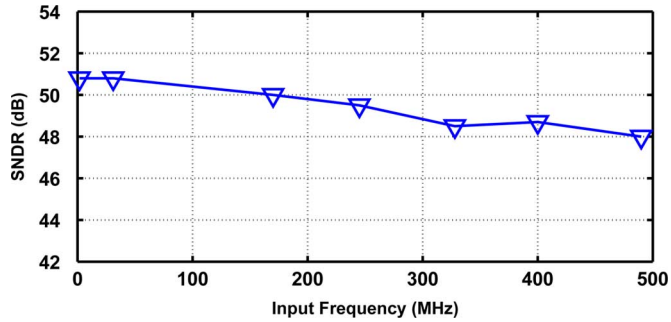

Fig. 21. Measured SNDR versus input frequency at a sampling rate of 1 GHz.

TABLE I
ADC Power Break Down Obtained by Simulation

Coarse Flash Stage	**0.9 mW**
Amplifier	**2.5 mW**
Fine Flash Stage	**1.2 mW**
Reference Ladderes	**1.0 mW**
Clock Generation	**1.4 mW**
Total	**7.1 mW**

Controlled by on-chip registers and a serial bus, the calibration of the two stages proceeds by sending the outputs of individual comparators off chip, interpreting their values, and adjusting the on-chip registers accordingly. For example, the comparator noise averaging mentioned in Section IV-A is performed off chip. The ADC has been directly mounted on a printed-circuit board and tested at a sampling rate of 1 GHz with a 1 V supply. The digital output is downsampled by a factor of 125.

Fig. 19 plots the measured DNL and INL before and after calibration. The peak DNL drops from +5.2/−1 LSB to +0.8/−0.87 LSB and the peak INL from +5.3/−3.9 LSB to +1.8/−1.6 LSB.

Fig. 20 shows the measured output spectrum for full-scale analog inputs at 1.7 MHz and 490 MHz. The ADC achieves a signal-to-(noise+distortion) ratio (SNDR) of 51 dB at low frequencies and 48 dB at the Nyquist, with a spurious-free dynamic range (SFDR) of 58 dB in both cases.

Plotted in Fig. 21 is the measured SNDR as a function of the input frequency at a sampling rate of 1 GHz. Table I shows the ADC's power breakdown and Table II compares the performance to recent art in the resolution range of 8 to 10 bits.

VII. Conclusion

The use of a double-sampling front-end, interleaved precharged DACs, and a new comparator-based calibration scheme allows two-stage pipelined ADCs to operate at high speeds with low power consumption. The calibration technique corrects for comparator offset and kickback noise, and amplifier offset, gain error, and nonlinearity. A 60 nm

TABLE II
ADC PERFORMANCE SUMMARY AND COMPARISON WITH PRIOR ART

	Sahoo VLSI 2012	Hashemi CICC 2012	Hong ISSCC 2013	Lien VLSI 2012	This Work
Resolution (bit)	10	10	9	8	9
Input Cap. (pF)	NA	0.7	0.7	NA	0.5
SNDR (dB) Low Freq./ Nyquist	56 / 52	57 / 53	54 / 51	45 / 43	51 / 48
f_S (MHz)	1000	1000	900	750	1000
Power (mW)	33	36	10.8	4.5	7.1
FoM (fJ/CS) @ Low Freq.	75	70	30	41	25
FoM (fJ/CS) @ Nyquist	97	100	40	50	34
Technology (nm)	65	65	45	28	65
Supply (V)	1.2	1.2	1.2	1.0	1.0
Active area (mm^2)	0.225	0.175	0.038	0.004	0.1

CMOS 1 GHz prototype exhibits an SNDR of 48 dB at the Nyquist rate while consuming 7.1 mW.

ACKNOWLEDGMENT

The authors gratefully acknowledge the TSMC University Shuttle Program for chip fabrication.

REFERENCES

[1] W. Yang, D. Kelly, I. Mehr, M. T. Sayuk, and L. Singer, "A 3-V 340-mW 14-b 75-Msample/s CMOS ADC with 85-dB SFDR at Nyquist input," *IEEE J. Solid-State Circuits*, vol. 36, no. 12, pp. 1931–1936, Dec. 2001.

[2] C. C. Lee and M. P. Flynn, "A SAR-assisted two-stage pipeline ADC," *IEEE J. Solid-State Circuits*, vol. 46, no. 4, pp. 859–869, Apr. 2011.

[3] B. D. Sahoo and B. Razavi, "A 10-b 1-GHz 33-mW CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 48, no. 6, pp. 1442–1452, Jun. 2013.

[4] S. Hashemi and B. Razavi, "A 10-bit 1-GS/s CMOS ADC with FOM = 70 fJ/conversion," in *Proc. 2012 IEEE Custom Integrated Circuits Conf.*, Sep. 2012, pp. 296–297.

[5] S. W. Chiang and B. Razavi, "A 10-bit 800-MHz 19-mW CMOS ADC," in *2013 IEEE Symp. VLSI Circuits Dig.*, Jun. 2013, pp. 100–101.

[6] S. Hashemi and B. Razavi, "A 10-bit 1-GS/s CMOS ADC with FOM = 70 fJ/conversion," in *Proc. 2013 IEEE Custom Integrated Circuits Conf.*, Sep. 2013.

[7] A. G. F. Dingwall and V. Zazzu, "An 8-MHz CMOS subranging 8-bit A/D converter," *IEEE J. Solid-State Circuits*, vol. SC-20, no. 6, pp. 1138–1143, Dec. 1985.

[8] B. P. Brandt and J. Lutsky, "A 75-mW, 10-b, 20-MSPS CMOS subranging ADC with 9.5 effective bits at Nyquist," *IEEE J. Solid-State Circuits*, vol. 34, no. 12, pp. 1788–1795, Dec. 1999.

[9] D. J. Huber, R. J. Chandler, and A. A. Abidi, "A 10 b 160 MS/s 84 mW 1 V subranging ADC in 90 nm CMOS," in *2007 IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 454–457.

[10] I. Mehr and L. Singer, "A 55-mW, 10-bit, 40-Msample/s Nyquist-rate CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 35, no. 3, pp. 318–325, Mar. 2000.

[11] B. Verbruggen, J. Craninckx, M. Kuijk, P. Wambacq, and G. Van der Plas, "A 2.6-mW 6 bit 2.2 GS/s fully dynamic pipeline ADC in 40 nm digital CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 10, pp. 2080–2090, Oct. 2010.

[12] A. Verma and B. Razavi, "A 10-bit 500-MS/s 55-mW CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 44, no. 11, pp. 3039–3050, Nov. 2009.

[13] Y.-H. Chung and J.-T. We, "A 16-mW 8-bit 1-GS/s subranging ADC in 55 nm CMOS," in *2011 IEEE Symp. VLSI Circuits Dig.*, Jun. 2011, pp. 128–129.

[14] S. Ibrahim and B. Razavi, "Low-power CMOS equalizer design for 20-Gb/s systems," *IEEE J. Solid-State Circuits*, vol. 46, no. 6, pp. 1321–1336, Jun. 2011.

[15] G. Van der Plas and B. Verbruggen, "A 150 MS/s 133 μW 7 bit ADC in 90 nm digital CMOS," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2631–2640, Dec. 2008.

[16] B. Razavi, "Design considerations for interleaved ADCs," *IEEE J. Solid-State Circuits*, vol. 48, no. 8, pp. 1806–1817, Aug. 2013.

[17] U.-T. Wang and B. Razavi, "An 8-bit 150-MHz CMOS A/D converter," *IEEE J. Solid-State Circuits*, vol. 35, no. 3, pp. 308–317, Mar. 2000.

[18] M. Dessouky and A. Kaiser, "Vey low-voltage digital-audio $\Delta\Sigma$ modulator with 88-dB dynamic range using local switch bootstrapping," *IEEE J. Solid-State Circuits*, vol. 36, no. 3, pp. 349–355, Mar. 2001.

[19] H.-K. Hong *et al.*, "An 8.6 ENOB 900 MS/s time-interleaved 2 b/cycle SAR ADC with a 1 b/cycle reconfiguration for resolution enhancement," in *2013 IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2013, pp. 470–471.

[20] Y.-C. Lien, "A 4.5-mW 8-b 750-MS/s 2-b/step asynchronous subranged SAR ADC in 28-nm CMOS technology," in *2012 IEEE Symp. VLSI Circuits Dig.*, Jun. 2012, pp. 88–89.

Sedigheh Hashemi received the B.S. and M.S. degrees from the University of Tehran, Tehran, Iran, in 2005 and 2007, respectively, and the Ph.D. degree from the University of California, Los Angeles, CA, USA, in 2012, all in electrical engineering.

She was a postdoctoral scholar in the Communication Circuits Laboratory at the University of California, Los Angeles, in 2013, focusing on the analysis and design of high-speed data converters. She is now with the Analog/RF circuit design group of Qualcomm.

Behzad Razavi (F'03) received the B.S.E.E. degree from Sharif University of Technology, Tehran, Iran, in 1985, and the M.S.E.E. and Ph.D.E.E. degrees from Stanford University, Stanford, CA, USA, in 1988 and 1992, respectively.

He was with AT&T Bell Laboratories and Hewlett-Packard Laboratories until 1996. Since 1996, he has been Associate Professor and subsequently Professor of electrical engineering at the University of California, Los Angeles. His current research includes wireless transceivers, frequency synthesizers, phase-locking and clock recovery for high-speed data communications, and data converters.

Prof. Razavi was an Adjunct Professor at Princeton University from 1992 to 1994, and at Stanford University in 1995. He served on the Technical Program Committees of the International Solid-State Circuits Conference (ISSCC) from 1993 to 2002 and VLSI Circuits Symposium from 1998 to 2002. He has also served as Guest Editor and Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, and the *International Journal of High Speed Electronics*.

Prof. Razavi received the Beatrice Winner Award for Editorial Excellence at the 1994 ISSCC, the best paper award at the 1994 European Solid-State Circuits Conference, the best panel award at the 1995 and 1997 ISSCC, the TRW Innovative Teaching Award in 1997, the best paper award at the IEEE Custom Integrated Circuits Conference in 1998, and the McGraw-Hill First Edition of the Year Award in 2001. He was the co-recipient of both the Jack Kilby Outstanding Student Paper Award and the Beatrice Winner Award for Editorial Excellence at the 2001 ISSCC. He received the Lockheed Martin Excellence in Teaching Award in 2006, the UCLA Faculty Senate Teaching Award in 2007, and the CICC Best Invited Paper Award in 2009 and 2012. He was the co-recipient of the 2012 VLSI Circuits Symposium Best Student Paper Award. He was also recognized as one of the top 10 authors in the 50-year history of ISSCC. Professor Razavi received the IEEE Donald Pederson Award in Solid-State Circuits in 2012.

Prof. Razavi is a Fellow of IEEE, has served as an IEEE Distinguished Lecturer, and is the author of *Principles of Data Conversion System Design* (IEEE Press, 1995), *RF Microelectronics* (Prentice Hall, 1998, 2012) (translated to Chinese, Japanese, and Korean), *Design of Analog CMOS Integrated Circuits* (McGraw-Hill, 2001) (translated to Chinese, Japanese, and Korean), *Design of Integrated Circuits for Optical Communications* (McGraw-Hill, 2003), and *Fundamentals of Microelectronics* (Wiley, 2006) (translated to Korean and Portuguese). He is also the editor of *Monolithic Phase-Locked Loops and Clock Recovery Circuits* (IEEE Press, 1996), and *Phase-Locking in High-Performance Systems* (IEEE Press, 2003).