A 10-b 1-GHz 33-mW CMOS ADC

Bibhu Datta Sahoo, *Member, IEEE*, and Behzad Razavi, *Fellow, IEEE*

***Abstract*—This paper describes a pipelined analog-to-digital converter that resolves 4 b in its first stage and amplifies the residue by a factor of 2, thereby relaxing the opamp linearity, voltage swing, and gain requirements. Calibration in the digital domain removes the effect of capacitor mismatches and corrects for the gain error. Using a one-stage opamp with a gain of 10 and realized in 65-nm CMOS technology, the ADC digitizes a 490-MHz input with a signal-to-(noise+distortion) ratio of 52.4 dB, achieving a figure of merit of 97 fJ/conversion-step.**

***Index Terms*—Analog-to-digital converter (ADC), pipelined ADC, multibit capacitor mismatch calibration, gain error calibration, low-gain opamp, low power.**

I. Introduction

CIRCUIT and architecture innovations continue to improve the performance of pipelined analog-to-digital converters (ADCs) across a wide range of speeds and resolutions. This unabated trend is partially owed to the increasing "raw" speed of MOS devices, with their deteriorating analog properties corrected by novel techniques.

Among various pipeline imperfections, the opamp nonlinearity proves to be particularly problematic. Nonlinearity calibration techniques have been studied extensively [1]–[5], but they typically require multipliers operating at the full rate and hence considerable power consumption in the digital domain. Another issue that arises in foreground calibration of nonlinearity is that the opamp open-loop gain must be high enough to avoid significant error in the correction as the temperature varies. It is therefore preferable to avoid nonlinearity calibration.

This paper describes the design of a 10-b 1-GHz ADC that requires no opamp nonlinearity correction [6]. The ADC incorporates a multibit front end and digitally calibrates capacitor mismatches and gain errors. Realized in 65-nm CMOS technology, the prototype improves the figure-of-merit (FOM) by about a factor of two with respect to the state of the art.

Section II describes the ADC architecture and various design tradeoffs. Section III presents the calibration algorithm. Section IV deals with the critical building blocks, and Section V summarizes the experimental results.

Manuscript received September 05, 2012; revised February 12, 2013; accepted February 28, 2013. Date of publication April 03, 2013; date of current version May 22, 2013. This paper was approved by Associate Editor Lucien Breems. This work was supported by DARPA HEALICS Program, Realtek Semiconductor, and Pullman Lane Productions.

B. D. Sahoo was with the Electrical Engineering Department, University of California, Los Angeles, CA 90095 USA. He is now with Amrita University, Amritapuri, India.

B. Razavi is with the Electrical Engineering Department, University of California, Los Angeles, CA 90095 USA (e-mail: email: 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.2013.2252518

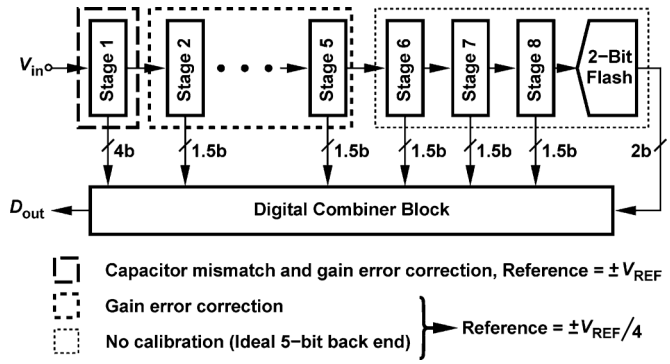

Fig. 1. ADC architecture.

II. ADC Architecture

The design of the opamp used in the first multiplying digital-to-analog converter (MDAC) is greatly relaxed if the first sub-ADC resolves several bits. Traditionally, the residue gain of the MDAC is chosen to be equal to 2^M, where M denotes the sub-ADC resolution, so that the next stage in the pipeline operates with the same full-scale range. While avoiding reference scaling and hence additional resistor dividers, this approach still requires the opamp to accommodate a maximum output swing equal to half of the ADC input range. Consequently, the opamp linearity and settling speed must meet stringent conditions.

A. Architecture Details

In this work, the first MDAC provides a nominal gain of 2 and the first sub-ADC resolves 4 bits. Since the MDAC output swing is confined to approximately ± 150 $\mathrm{mV_p}$, the opamp can be realized as a single stage, offering a more favorable speed-power trade-off than that of a high-gain MDAC and outweighing the power penalty due to the additional reference divider.

Fig. 1 shows the ADC architecture. A 4-b SHA-less front-end is followed by seven 1.5-b stages and one 2-b stage. In the first stage, capacitor mismatches and the gain error are calibrated and in stages 2–5, only the gain error. With ± 0.5 bit of redundancy and an MDAC gain of 2, the reference is scaled down by a factor of 4 for the remaining stages [3]. The foreground calibration begins from stage 5 and proceeds backwards so that the calibration of stage $j (j < 6)$ can assume that the subsequent stages form an "ideal back end." The ADC is designed somewhat conservatively in that the nine stages in fact provide a resolution of 11 b[1]. The performance of the ADC is limited to 10 b by kT/C noise and opamp noise.

The ADC employs circuit and power scaling in the pipeline: the MDAC in stage 2 is scaled down by a factor of 2 and in stage 3 by another factor of 2. Stages 4–8 are identical and are scaled down by 25% with respect to stage 3. With a gain of 2 in the first

[1]For the calibration to be 11-b accurate, an additional bit is required [4], [5], [7], [8], hence the overall resolution of 12 b.

stage, the noise of the second stage must be taken into account. This noise lowers the SNR to 59 dB.

B. Resolution of First Sub-ADC

It is desirable to maximize the first sub-ADC's resolution M, but a number of factors must be considered. In addition to obvious imperfections such as the input capacitance, power consumption, and kickback noise, all of which rise with the resolution, we must also quantify certain error budgets that *fall*. Let us assume 0.5 b of redundancy. First, since the front-end does not contain an explicit sample-and-hold circuit, the timing mismatch Δt between the sub-ADC and the MDAC must remain well below 0.5 LSB of the former. For a maximum analog input frequency of f_0 and a peak amplitude of A_0, we must choose M such that the timing-mismatch-induced voltage error $2\pi f_0 A_0 \Delta t$ is a fraction of 0.5 LSB $= A_0/2^M$

$$2^M < \frac{1}{2\pi f_0 \Delta t}. \tag{1}$$

For example, if $f_0 = 500$ MHz and $\Delta t =$ 10 ps, then (1) gives an upper bound of 5 b for M. While $\Delta t = 10$ ps appears rather pessimistic, it should be borne in mind that a 5-b flash ADC does stretch over several hundred microns, making it difficult to minimize Δt.[2]

The second error margin relates to the 3σ offset $V_{\rm OS}$ of the sub-ADC comparators and its tradeoff with the input capacitance. We optimistically assume that $V_{\rm OS}$ arises from only the threshold mismatch of one differential pair:

$$V_{\rm OS} = \frac{3A_{V_{\rm TH}}}{\sqrt{WL}} \tag{2}$$

where $A_{V_{\rm TH}}$ is in the range of 3 to 4 mV·μm for 65-nm technology. For a full-scale range of 1 V and $M = 4$, a WL of, say, 1 μm^2 yields a $V_{\rm OS}$ of about 1/6th of the LSB and a sub-ADC input capacitance of 15 × 1 μm^2 × (15 fF/μm^2) ≈ 225 fF plus about 50 fF of interconnect capacitance, all comfortable values. For $M = 5$, on the other hand, a WL of 1 μm^2 severely tightens the budget for other errors and presents an input capacitance nearly equal to that of the MDAC. Offset cancellation can ameliorate these issues but was not adopted in this work.

Two other effects further limit the error budget as M increases. The input-referred noise of the comparators merits attention, especially if the design employs small devices along with offset cancellation. Also, as explained in Section III, the capacitor mismatches in the DAC and the MDAC gain error produce some residue overrange, which must be accommodated by the opamp.

The last issue that plays a role in the choice of M relates to the conversion time of the sub-ADC $t_{\rm sub}$. As M increases, so does $t_{\rm sub}$ (for a given power consumption). Consequently, the timing budget allocated to the settling of the MDAC decreases, requiring a greater bias current in the opamp.

Based on the foregoing observations, a resolution of 4 b has been selected for the first sub-ADC in this work. Fig. 2 conceptually shows the first stage in single-ended form, where C_p denotes the opamp input capacitance. By virtue of bootstrapped switches, the front-end sampling speed is fast enough to require only 25% of the clock period, leaving the other 75% for the sub-ADC conversion and the MDAC settling [4]. The thermometer code generated by the sub-ADC switches the bottom plates of C_1-C_{15} to 0 or $V_{\rm REF}$, producing the amplified residue, $V_{\rm res}$, at the MDAC output. Note that $C_1 + \cdots + C_{16}$ is chosen equal to 480 fF, degrading the SNR from its ideal value of 62 dB by about 1 dB due to kT/C noise.

[2]With an acquisition time constant of 30 ps, the MDAC and sub-ADC bandwidth mismatch negligibly affects the performance.

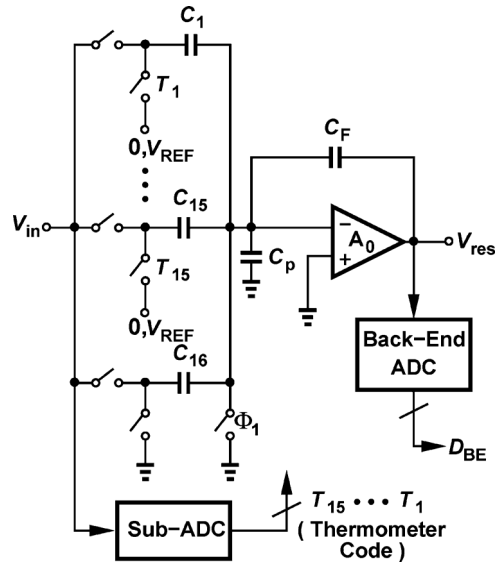

Fig. 2. 4-b MDAC architecture.

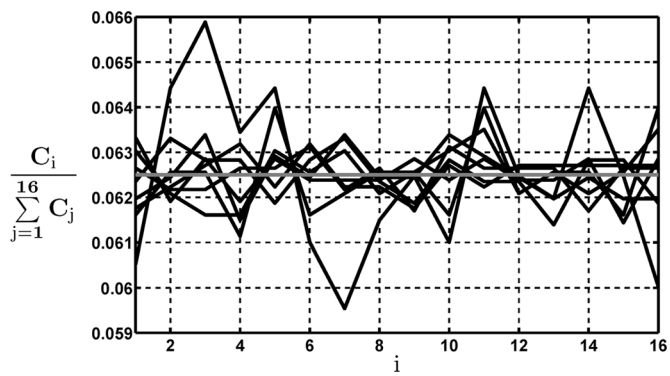

Fig. 3. Measured capacitor mismatch of 8 prototype chips.

III. Calibration Methodology

While relaxing the design of the opamp, a multibit sub-ADC distorts the residue characteristic through the DAC capacitor mismatches. Since $C_1 = \cdots = C_{16} =$ 30 fF in Fig. 2 are chosen according to kT/C noise limitations rather than matching considerations and since C_1 and C_{16} are several hundred microns apart in the layout, some means of mismatch calibration is necessary. Fig. 3 shows as an example the measured deviation and mean values of the first stage capacitors of eight prototype chips, revealing a worst-case error of about 5.6%. (These values are computed from the digitized residue measurements of the ADC prototype described in Section V).

A. Problem Statement

In order to develop an algorithm for the correction of capacitor mismatches, we first formulate their effect on the residue

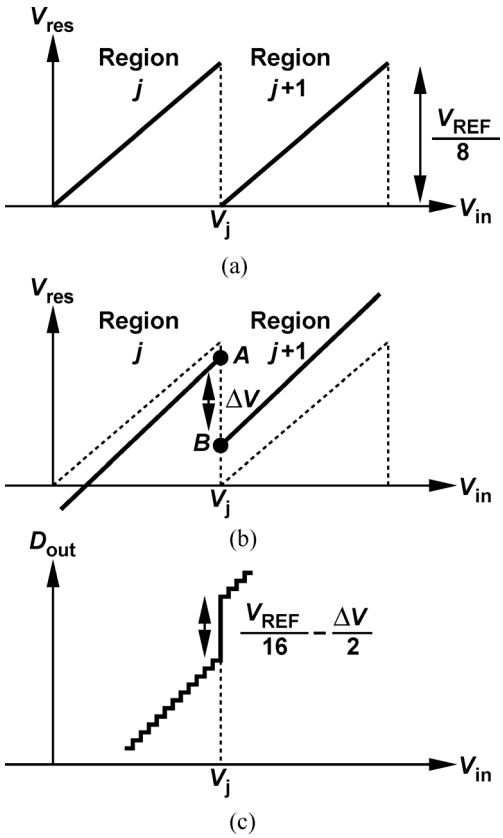

Fig. 4. (a) Ideal residue characteristic. (b) Residue characteristic with capacitor mismatch. (c) ADC input–output characteristic.

characteristic. In the single-ended model of Fig. 2, we ideally have $V_{\rm res} = 2V_{\rm in} - (j/8)V_{\rm REF}$, where j denotes the height of the thermometer code, thus obtaining the characteristic shown in Fig. 4(a). With capacitor mismatches, on the other hand, we write

$$V_{\rm res} = \alpha V_{\rm in} - \beta_j V_{\rm REF} \tag{3}$$

where

$$\alpha = \frac{C_1 + \cdots + C_{16}}{C_{\rm eq}} \tag{4}$$

$$\beta_j = \frac{C_1 + \cdots + C_j}{C_{\rm eq}}. \tag{5}$$

Capacitance $C_{\rm eq}$ is ideally equal to C_F in Fig. 2, but must now include the mismatches, the effect of the finite gain of the opamp, and the input capacitance of the opamp

$$C_{\rm eq} = C_F + \frac{1}{A_o}\left(C_F + C_p + \sum_{m=1}^{16} C_m\right). \tag{6}$$

We make the following observations.

1) The residue gain α is the same in all the 16 regions $j = 0, \ldots, 15$, but unequal to 2 due to the accumulation of capacitor mismatches in the numerator of (4) and other nonidealities in the denominator.
2) The *vertical* shift ΔV, varies from one region to another according to the capacitor mismatches [Fig. 4(b)].

Specifically, in the transition from region j to region $j+1$, the residue exhibits a jump equal to

$$\Delta V = (\beta_{j+1} - \beta_j)V_{\rm REF} \tag{7}$$

$$= \frac{C_{j+1}}{C_{\rm eq}} V_{\rm REF}. \tag{8}$$

(For now, we assume the subsequent stages have infinite resolution so that the digital jump is equal to the analog value ΔV). In the ideal case, $C_{j+1} = C$ and $C_F = 8C$. Assuming now that $C_{j+1} = C + \Delta C$ and $C_{\rm eq} \approx 8C$, we obtain

$$\Delta V = \left(1 + \frac{\Delta C}{C}\right)\frac{V_{\rm REF}}{8} \tag{9}$$

concluding that the overall ADC input/output characteristic incurs an unwanted jump equal to $(\Delta C/C)V_{\rm REF}/8/2$, where the factor of 2 accounts for the voltage gain of the first stage.[3] For this error to be less than, say, 0.25 LSB, we must have

$$\frac{\Delta C}{C} < \frac{1}{256}. \tag{10}$$

Note that the error in each residue transition is given by the mismatch between only one DAC capacitor and C_F. This result agrees with the equation provided by [7].

Equations (4)–(8) suggest that several errors must be calibrated: 1) the mismatches among the DAC capacitors, denoted by ΔC_{j+1} above; 2) the mismatch between C_F and each DAC capacitor; and 3) the residue gain error due to the opamp gain and input capacitance. As explained below, the first two in fact collapse into one. Note that the residue overranges in Fig. 4(b) (measured in practice to be 10 to 15 mV) must also be accommodated by the opamp.

B. Calibration Algorithm

Here, we describe a calibration algorithm for removing the three errors identified above, beginning with the mismatches among the DAC capacitors.

It is possible to deal with the capacitor mismatches in a multibit pipelined stage by dynamic element matching [8]. However, at a sampling rate of 1 GHz, the three-level multiplexing required by this approach constrains the sub-ADC and MDAC timing budget (a total of about 600 ps with nonoverlapping clocks). In addition, [8] injects a pseudorandom analog input to correct for gain error, lowering the overall dynamic range of the ADC by 1 dB. In [9], a foreground capacitor mismatch calibration technique for multibit two-step ADCs has been described that measures the deviation of the ADC input–output characteristic at major carry jumps for a binary-weighted DAC. It then adds or subtracts the deviations in the digital domain. Due to the binary-weighting of the capacitors, the deviation for one region of the residue characteristic is the cumulative sum of those of the lower regions, thus leading to digital truncation errors. In addition, since the technique is not applicable to the MDAC gain error, the work in [9] relies on high-gain opamps.

[3]For a closed-loop MDAC gain of $A_{\rm MDAC}$, we have $C_F = 16C/A_{\rm MDAC}$ and $\Delta V = (1+\Delta C/C)A_{\rm MDAC}V_{\rm REF}/16$. The unwanted input-referred jump is thus equal to $(\Delta C/C)V_{\rm REF}/16$ and hence independent of $A_{\rm MDAC}$.

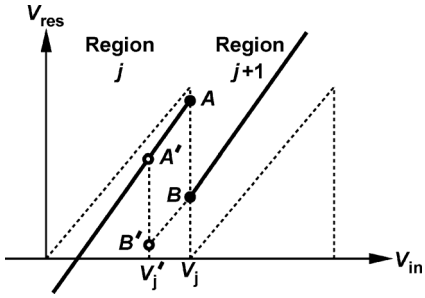

Fig. 5. Effect of ladder nonlinearity on calibration.

The principle behind the calibration algorithm proposed here emerges from (7) and (8): if $(\beta_{j+1}-\beta_j)V_{\rm REF}$ can be measured, so can the value of $C_{j+1}/C_{\rm eq}$. With $C_{j+1}/C_{\rm eq}$ known, α in (4) and β_j in (5) can be reconstructed and hence the first stage input/output characteristic is completely determined. It is interesting to note that the capacitor *mismatches* themselves need not be calculated.

The computation of $(\beta_{j+1}-\beta_j)V_{\rm REF}$ proceeds as follows. We observe that the residue in Fig. 4(b) reaches $\alpha V_j-\beta_j V_{\rm REF}$ as $V_{\rm in}$ approaches V_j (point A) and drops to $\alpha V_j-\beta_{j+1}V_{\rm REF}$ as $V_{\rm in}$ slightly exceeds V_j (point B). We can therefore digitize these two residue values by means of the ADC back end (stages 2–9) and form the difference in the digital domain. This is accomplished by shorting the analog input to V_j (produced by the sub-ADC resistor ladder) and forcing to ZERO or ONE the output of the corresponding sub-ADC comparator, thus exercising points A and B, respectively.

Two questions must now be answered. First, how do the sub-ADC comparator offsets affect the above computation? Since the analog input is directly set equal to V_j and since the comparator makes no decision of its own accord, the comparator offset is immaterial in this measurement. (Of course, as explained in Section II, the comparator offset still affects the error margin in the first stage). Second, since the voltages V_j in Fig. 4(b) are generated by means of a resistor ladder, how does the linearity of the ladder affect the results? Interestingly, the ladder linearity is also immaterial. As illustrated in Fig. 5, even if V_j deviates from its ideal value and moves to V_j', the residue difference between points A' and B' is the same as that between A and B. (However, significant integral nonlinearity in the ladder constrains the error budgets described in Section II).

The foregoing algorithm yields the values of $C_1/C_{\rm eq}$ to $C_{15}/C_{\rm eq}$. It is therefore possible to form the summations $\beta_j=(C_1+\cdots+C_j)/C_{\rm eq}$ in the digital domain (Section IV-C). We must also compute α in (4) so as to obtain the complete first-stage characteristic. Equations (4) and (5) suggest that β_{15} and α differ only by $C_{16}/C_{\rm eq}$. We thus carry out one more calculation: we swap C_{16} and C_1 as shown in Fig. 6, equate $V_{\rm in}$ to V_{15}, and force the corresponding comparator output to ZERO or ONE. It follows that $\alpha=\beta_{15}+C_{16}/C_{\rm eq}$, which can be realized in the digital domain (Section IV-C).

The foregoing computations of $C_j/C_{\rm eq}$ have assumed an infinite resolution for the back end. For $C_j/C_{\rm eq}$ to have an accuracy of N bits, the estimation must proceed with a resolution of $N+1$ bits [9]. Thus, stages 2–9 in fact provide one extra bit for this purpose.

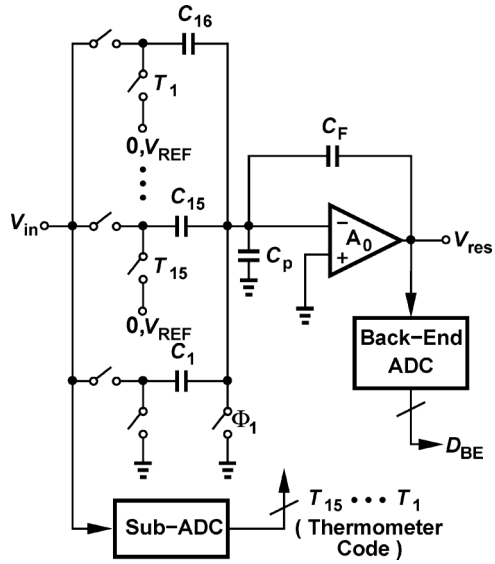

Fig. 6. 4-b MDAC architecture with C_1 and C_{16} swapped.

C. Advantages of Proposed Algorithm

It is worth noting that the simple summations formed above encapsulate not only the DAC capacitor mismatches but also mismatches between each DAC capacitor and C_F in Fig. 2 as well as the effect of the finite gain and input capacitance of the opamp. That is, so long as the linear residue equation $V_{\rm res}=\alpha V_{\rm in}-\beta_j V_{\rm REF}$ applies, the first-stage nonidealities can be readily corrected with no need for high-speed digital multipliers.

Another advantage of knowing the value of α is that it can be adjusted and set close to the nominal value of 2 in the *analog* domain (Section IV-B). This point stands in contrast to conventional pipeline design, where the residue gain must remain safely below 2 so as to avoid residue overrange, thus suppressing the noise of the subsequent stages to a lesser extent and requiring a larger number of stages in the pipeline. The ability to bring the gain close to 2 is particularly attractive in our work as the low open-loop gain of the opamp leads to more than 10% of gain error.

We should also note two disadvantages of the approach in [10] with respect to ours. First, the technique applies only to the mismatch between the input and feedback capacitors in a 1-bit stage. Second, it requires an MDAC gain of less than 2.

The split-ADC method in [11] operates in the background but it imposes three constraints on the design.

1) It requires a dedicated front-end sample-and-hold amplifier and may not lend itself to a SHA-less architecture.
2) To resolve 4 b in the first stage, it actually requires a 5-b sub-ADC. As recognized by, e.g., [4] and [12], a SHA-less front-end saves considerable power and noise.
3) It accumulates truncation errors from one residue region to the next in a manner similar to the work in [9].

D. Back-End Calibration

With 4 b resolved in the first stage, the precision requirements of the subsequent 1.5-b stages are greatly relaxed. Nonetheless, capacitor mismatches and low opamp gain still demand gain

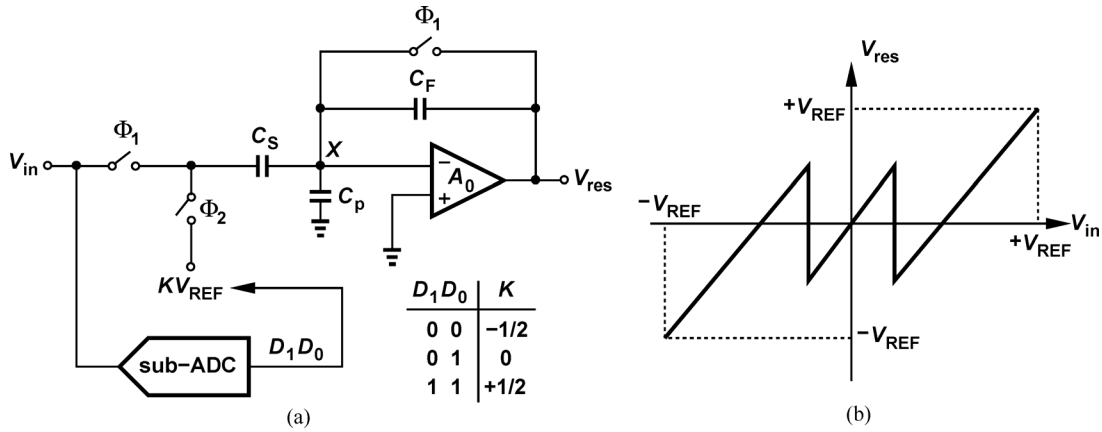

Fig. 7. (a) 1.5-b stage nonflip-around MDAC and (b) its residue characteristic.

calibration in stages 2–5. Shown in Fig. 7(a), the 1.5-b topology used in these stages satisfies the following residue equation:

$$V_{\rm res} = \alpha(V_{\rm in} - KV_{\rm REF}) \tag{11}$$

where $\alpha = C_S/C_{\rm eq}, C_{\rm eq} = C_F + (C_F + C_p + C_S)/A_o$, and $K = 0, -1$, or $+1$.

In order to develop the back-end ADC calibration technique, let us rewrite (11) to obtain $V_{\rm in}$ in terms of the residue as

$$V_{\rm in} = KV_{\rm REF} + \frac{1}{\alpha}V_{\rm res} \tag{12}$$

where $1/\alpha$ is unknown and must be computed. We denote the digital equivalent of $V_{\rm in}$ by $D_{\rm in}$ and recognize that $KV_{\rm REF}$ translates to $K(= 0, \pm 1)$ in the digital domain. Since K is in fact the same as the digital output of the 1.5-b sub-ADC $D_{\rm sub}$, we have

$$D_{\rm in} = D_{\rm sub} + \frac{1}{\alpha}D_{\rm res} \tag{13}$$

where $D_{\rm res}$ is the digitized residue. Providing the desired digital value $D_{\rm in}$ in terms of measurable quantities, this equation simply means that the digital value corresponding to a given analog input is equal to the sub-ADC digital output plus the "input-referred" residue (because the residue is amplified by a factor of α before it is digitized).

We now invoke the calibration technique in [5] to remove the gain error. As illustrated in Fig. 8, we apply three different input levels to the stage under calibration: a small voltage ΔV ($\approx$10 LSB), $V_1(= V_{\rm REF}/4)$ (with the comparator output forced to ZERO), and $V_1 + \Delta V$ (with the comparator output forced to ONE). In each case, the sub-ADC and the remaining stages produce outputs that can be combined according to (13) as

$$V_{\rm in} = \Delta V \rightarrow D_a = D_{{\rm sub},a} + \frac{1}{\alpha}D_{{\rm res},a} \tag{14}$$

$$V_{\rm in} = V_1 \rightarrow D_b = D_{{\rm sub},b} + \frac{1}{\alpha}D_{{\rm res},b} \tag{15}$$

$$V_{\rm in} = V_1 + \Delta V \rightarrow D_c = D_{{\rm sub},c} + \frac{1}{\alpha}D_{{\rm res},c}. \tag{16}$$

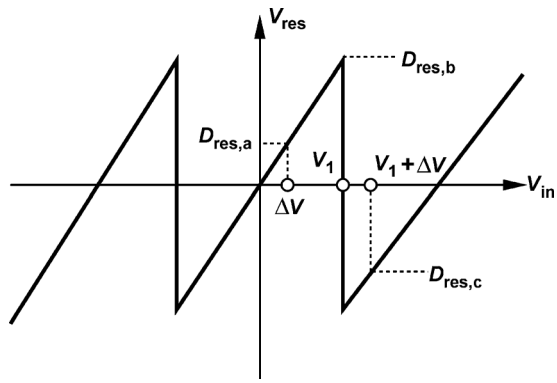

Fig. 8. Perturbation-based calibration for 1.5-b stage.

Note that $D_{{\rm sub},a} = D_{{\rm sub},b}$ because $0 < V_{\rm in} \leq V_{\rm REF}/4$ yields the same sub-ADC output. The key point here is that the difference between D_c and D_b is equal to D_a and hence

$$\frac{1}{\alpha} = \frac{D_{{\rm sub},c}}{D_{{\rm res},a} - D_{{\rm res},b} - D_{{\rm res},c}}. \tag{17}$$

With $1/\alpha$ known, the gain error can be corrected in the digital domain (Section IV-C).

IV. Building Blocks

A. Opamp

The use of a multibit front-end along with a residue gain of 2 allows a simple, efficient opamp design. Shown in Fig. 9, the single-stage opamp is realized with a tail current of 4.8 mA so as to settle in 400 ps while driving the MDAC capacitors and the input of the second stage, a total equivalent capacitance of 250 fF. With $(W/L)_{1,2} = 80\ \mu$m/90 nm and $(W/L)_{3,4} = 96\ \mu$m/200 nm, the circuit provides an open-loop gain of 20 dB and a single-ended output swing of 150 $\rm mV_{pp}$. In order to minimize timing and bandwidth mismatches, the MDAC and the sub-ADC incorporate identical input sampling topologies [4], [8].

Fig. 10 plots the simulated MDAC output spectrum at a sampling rate of 1 GHz and with a sinusoidal input at 490

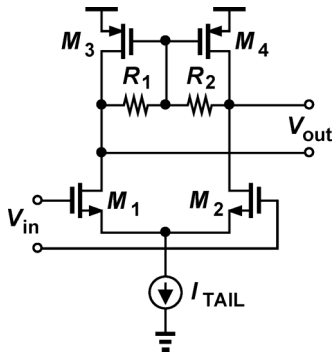

Fig. 9. Opamp circuit.

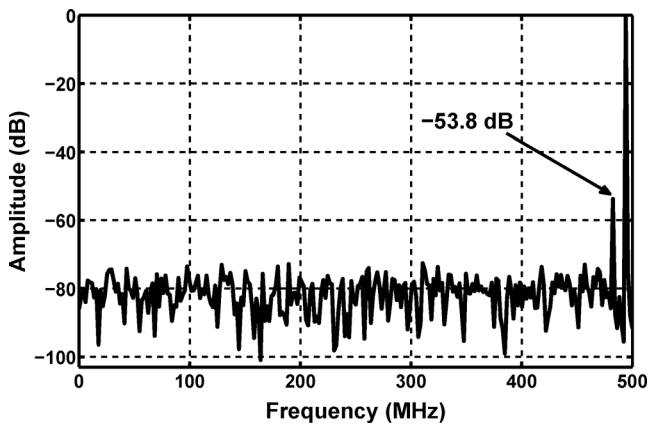

Fig. 10. Simulated MDAC output spectrum.

MHz. It is observed that the aliased third harmonic is about 54 dB below the fundamental for a differential output swing of 2×150 mVpp. Since only 7 b must be hereafter resolved, the circuit exhibits sufficient linearity. The opamps in stages 2 and 3 draw 2.4 and 1.2 mA, respectively, and those in stages 4–8, 800 μA each.

Despite a resolution of 4 bits preceding the MDAC, the low open-loop gain of the opamp does make the closed-loop gain somewhat sensitive to the temperature. However, the constant-overdrive biasing technique in [13] can be applied to partially cancel the gain variation. Circuit simulations implementing this biasing in the opamp indicate that the SNDR of the overall ADC degrades by 2.8 dB if the temperature rises from 0°C to 75°C and the calibration coefficients for 27°C are used. The proposed ADC therefore requires recalibration if a large temperature change occurs.

B. Programmable-Gain MDAC

Due to the low opamp gain, the MDAC exhibits a gain error as high as 10%. As explained in Section III-C, the ability to compute α for each stage enables us to perform some gain correction in the analog domain. Illustrated in Fig. 11, the idea is to choose the MDAC feedback capacitor, C_F, for a nominal gain *greater* than 2 (about 2.2) and adjust the gain toward a value of 2 by means of the 3-b array $C, 2C$, and $4C$. (In stages 3–8, a 2-b array is used.) In the first stage, $C_F = 120$ fF and $C = 7.5$ fF.

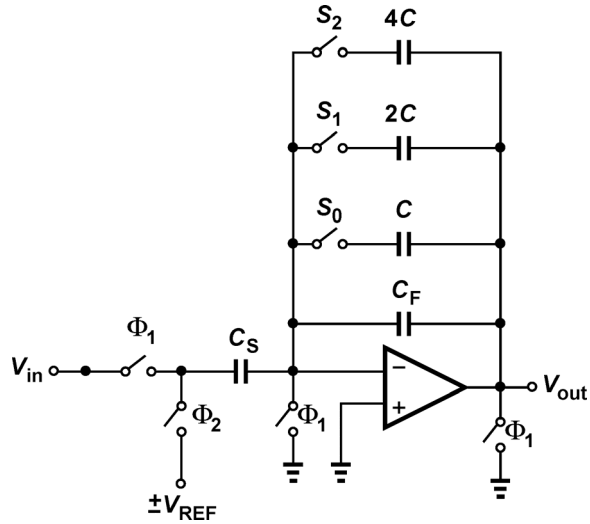

Fig. 11. Programmable feedback capacitors in MDACs.

Since layout parasitics considerably alter C_F and the unit capacitor C, their actual values are chosen after layout extraction. Note that the remaining stages in the pipeline provide enough redundancy to accommodate a residue gain of slightly greater than 2. Thus, a resolution of 3 b for the gain setting of the first MDAC is adequate. For most of the prototype chips that have been tested, only switch S_0 needs to be turned on to provide a gain of about 2.05. For the back-end stages, either S_1 or S_2 is turned on. Of course, the residual gain error is still excessively large for 10-b performance and must be calibrated as explained in Section III.

C. Digital Back End

The experimental prototype reported in this paper does not include the digital back-end. However, detailed design and synthesis of the logic have been carried out to estimate its complexity and power consumption. We remark that the digital back end must perform two distinct functions: 1) computation of various coefficients (α's and β's) during foreground calibration and 2) correction of the digital output of the ADC. The former is frozen after the calibration is finished, drawing negligible power, whereas the latter appears in the high-speed path and must be optimized for minimal consumption.

Let us begin with the first stage of the pipeline. As explained in Section III, the values of α and β_j in (4) and (5) are computed according to the techniques illustrated in Figs. 6 and 4, respectively. This computation proceeds as follows. First, the two residues $\alpha V_j - \beta_j V_{\rm REF}$ (at point A) and $\alpha V_j - \beta_{j+1} V_{\rm REF}$ (at point B) in Fig. 4(b) are digitized, stored, and subtracted, yielding $\beta_{j+1} - \beta_j = C_{j+1}/C_{\rm eq}$. This is repeated for $j = 0$ to 15 and, by swapping C_1 and C_{16} as in Fig. 6, for $j = 16$ as well. Next, the summations $C_1/C_{\rm eq}, (C_1 + C_2)/C_{\rm eq}, (C_1 + C_2 + C_3)/C_{\rm eq}$, etc., are generated, thus forming β_0 to β_{15}, respectively. The value of α is obtained in a similar manner. These operations require approximately 5000 gates, but this logic remains frozen after the foreground calibration is finished. The resulting β_j values are stored in a look-up table, which, as explained below, does appear in the high-speed path.

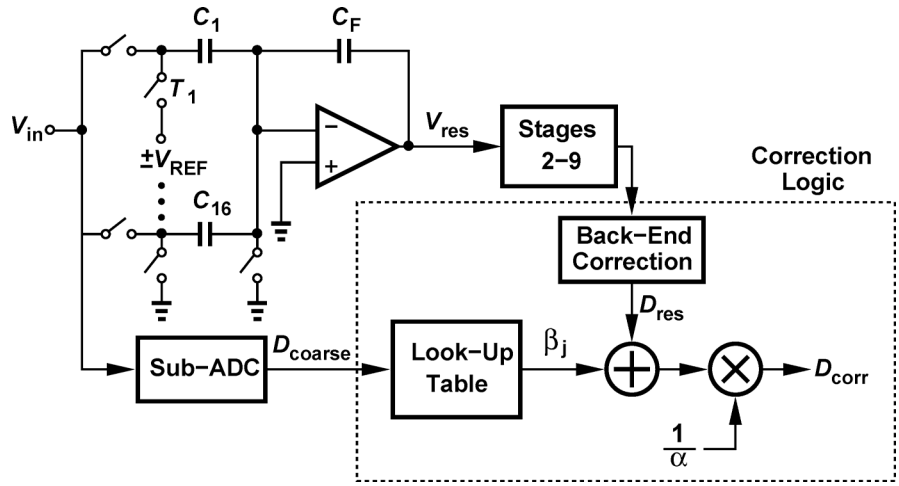

Fig. 12. Conceptual illustration of digital correction at the ADC output.

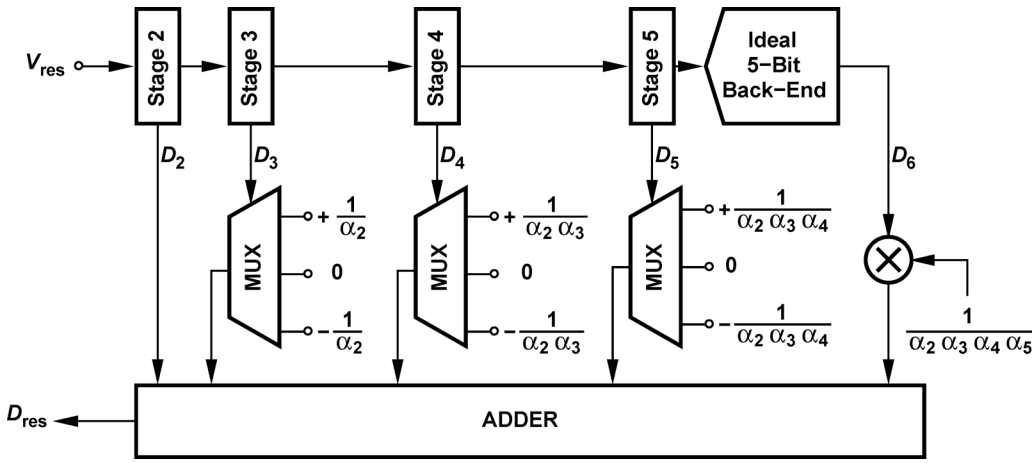

Fig. 13. Back-end correction logic.

We must now decide how these values can be utilized to correct the digital output. Since $V_{\rm res} = \alpha V_{\rm in} - \beta_j V_{\rm REF}$, we have $V_{\rm in} = (V_{\rm res} + \beta_j V_{\rm REF})/\alpha$ and hence in the digital domain

$$D_{\rm corr} = \frac{1}{\alpha}(D_{\rm res} + \beta_j). \tag{18}$$

Here, $D_{\rm corr}$ denotes the overall corrected digital output corresponding to $V_{\rm in}$, and $D_{\rm res}$ the digitized residue produced by stages 2 to 9. The digital back end performing (18) thus emerges as shown in Fig. 12. We make the following observations.

1) Each sub-ADC output code must invoke a certain β_j; thus, the β_j look-up table is driven by the sub-ADC.
2) According to (18) $D_{\rm res}$ and β_j must be summed and scaled down by a factor of α, requiring both a high-speed multiplier and a divider. The divider is avoided by applying the Newton-Raphson [14] method to generate $1/\alpha$ (Appendix A). The $1/\alpha$ necessary for 1.5-b stages in (13) is generated in a similar manner.

We must now determine how $D_{\rm res}$ in Fig. 12 is produced. The high-speed back-end correction logic is illustrated in Fig. 13. Note that D_2 represents the residue of the *first* MDAC as digitized by Stage 2, etc. Similarly, α_2 denotes the gain of the residue in Stage 2, etc. The objective is to digitize the first MDAC residue $V_{\rm res}$ and generate

$$\begin{aligned} D_{\rm res} =& D_2 + D_3 \cdot \frac{1}{\alpha_2} + D_4 \cdot \frac{1}{\alpha_2\alpha_3} \\ &+ D_5 \cdot \frac{1}{\alpha_2\alpha_3\alpha_4} + D_6 \cdot \frac{1}{\alpha_2\alpha_3\alpha_4\alpha_5}. \end{aligned} \tag{19}$$

As explained in Section III and above, the values of $1/\alpha_j$ and their cross products are computed during foreground calibration and stored in registers. Fortunately, since 1.5-b stages generate only -1, 0, or $+1$, the multiplication operations in (19) can be replaced with multiplexing. However, one multiplier is still necessary for D_6 and $1/(\alpha_2\alpha_3\alpha_4\alpha_5)$, and another one for the last multiplication by $1/\alpha$ in Fig. 12.

In summary, the digital logic consists of: 1) the low-speed calibration logic, comprising about 5000 gates and 2) the high-speed correction logic, containing an LUT of size 15×16 bits and about 12 000 gates.

Operating at 1 GHz, the correction logic can potentially draw substantial power. This logic was synthesized in 65-nm CMOS technology with various degrees of parallelism and voltage scaling to explore the tradeoff between complexity and power dissipation. Illustrated in Fig. 14 are the results of this effort

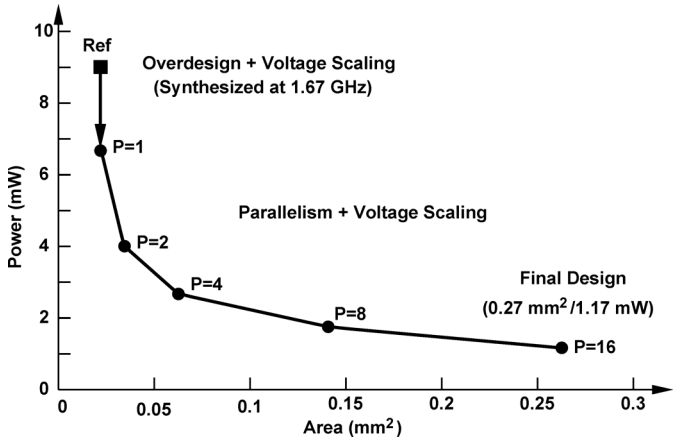

Fig. 14. Power–area tradeoff of digital combiner.

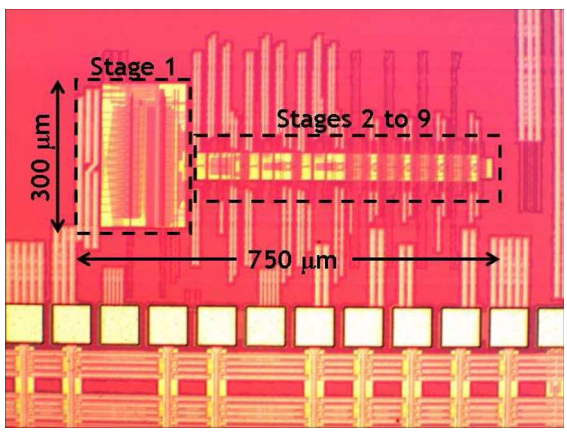

Fig. 15. Chip micrograph.

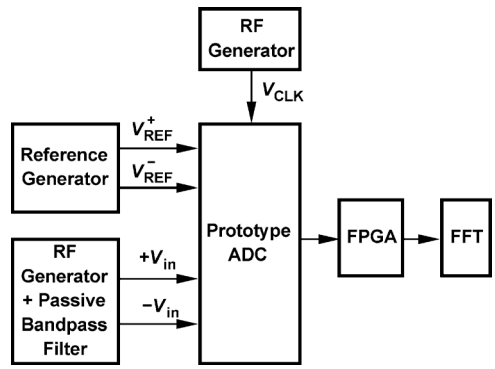

Fig. 16. Block diagram of the test setup.

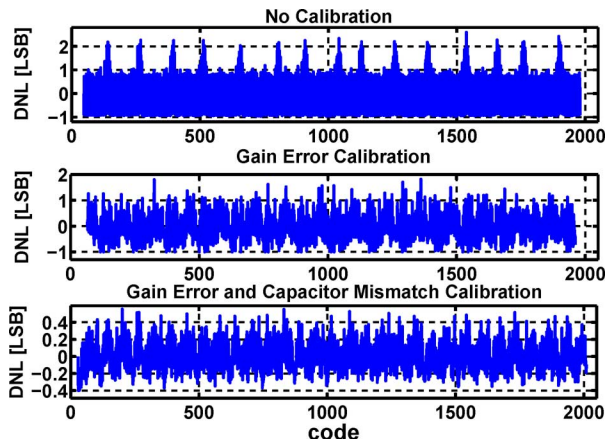

Fig. 17. Measured DNL for 11-b resolution at f_{clk} of 1 GHz.

for a conservative clock frequency of 1.7 GHz. Plotted here is the power consumed by the logic as the number of parallel branches goes from 1 to 16. It is observed that the power varies from 9 mW with no parallelism to about 1.2 mW with 16 parallel branches.[4]

V. Experimental Results

The prototype ADC has been fabricated in digital 65-nm CMOS technology. The die is shown in Fig. 15 and has an active area of 750 μm $\times$ 300 μm. Operating from a 1.2-V supply, the ADC has been tested with the chip directly mounted on a printed circuit board.

The references are provided externally for greater test flexibility. In such a case, we have two options.

1) Add a large on-chip capacitance between $V_{\rm REF}^{+}$ and $V_{\rm REF}^{-}$ so as to absorb the transient currents drawn by the capacitors in the pipeline. This approach minimizes reference disturbance if the bypass capacitance is several orders of magnitude greater than the sum of the ADC capacitors that switch onto the reference on each clock edge. If not sufficiently large, the bypass capacitance leads to a long settling time, degrading the ADC's dynamic performance.

2) Add no on-chip capacitance and let the reference lines remain completely agile. In this case, the references experience considerable ringing but settle fast. Simulations suggest that the latter approach creates a smaller dynamic error in this design and is therefore the preferable choice. The ringing is in fact partially damped by virtue of the two on-chip resistor ladders used in the first and subsequent stages, which introduce a net resistance of 100 Ω between $V_{\rm REF}^{+}$ and $V_{\rm REF}^{-}$.

Fig. 16 shows the ADC test setup. The analog sinusoidal input is derived from an RF signal generator, a balun, and a discrete passive filter to remove its harmonics. The clock is also obtained from an RF generator and converted to non-overlapping phases on the chip. The ADC output is downsampled by a factor of 16 and applied to an FPGA, which contains both the calibration and correction logic and adjusts the residue gain of the stages in the pipeline through an on-chip serial bus (Fig. 11).

Figs. 17 and 18 plot the measured differential nonlinearity (DNL) and integral nonlinearity (INL), respectively, with no correction, with only gain error correction, and with both gain error and capacitor mismatch correction. As mentioned in Section II, the prototype in fact provides a resolution of 11 b. For 10-b performance, the peak DNL and INL values in Figs. 17 and 18 should be divided by two, yielding values of 0.25 and 1 LSB, respectively.

[4]The power consumption is determined from the digital tools.

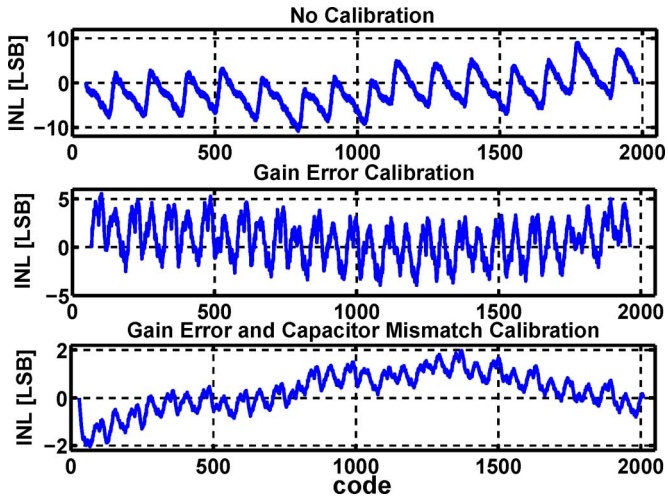

Fig. 18. Measured INL for 11-b resolution at $f_{\rm clk}$ of 1 GHz.

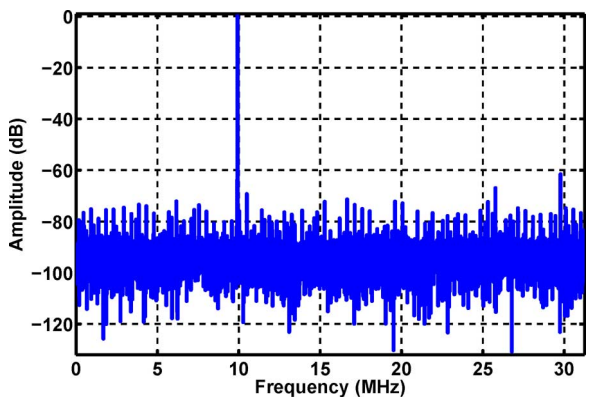

Fig. 19. Pipelined ADC output spectrum for $f_s = 1$ GHz (downsampled by a factor of 16), $V_{\rm in} = 1.2$ V and $f_{\rm in} = 490$ MHz.

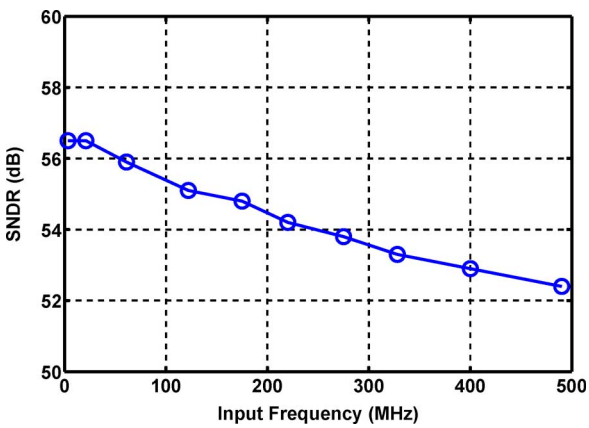

Fig. 20. Measured SNDR versus input frequency with calibration at $f_s = 1$ GHz.

Fig. 19 shows the measured output spectrum for a 490-MHz full-scale analog input sampled at 1 GHz. (Due to downsampling, the fundamental appears at 10 MHz.) The signal-to-(noise+distortion) ratio (SNDR) is equal to 52.4 dB.

Fig. 20 plots the measured SNDR as a function of the analog input frequency at a sampling rate of 1 GHz. (According to simulations, the kT/C and opamp noise limit the SNDR at low input frequencies to 59 dB.)[5]

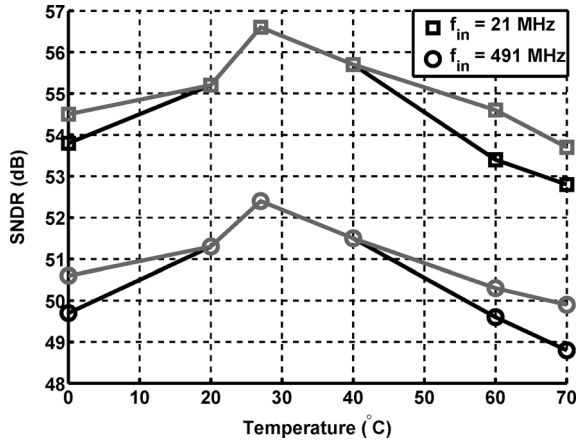

Fig. 21. Measured SNDR versus temperature with calibration performed at 27 °C. Black curves: settings are frozen; gray curves: MDAC feedback capacitors are adjusted below 10 °C and above 50 °C.

TABLE I
PERFORMANCE SUMMARY

Resolution		**11 bits**
Conversion Rate		**1 GHz**
Supply Voltage		**1.2 V**
Input Voltage		**1.2 V_{PP}**
Input Capacitance		**1 pF (Single Ended)**
Power	**Analog**	**15.12 mW**
	Clock	**14.3 mW**
	Reference[7]	**3.6 mW**
	Total	**33 mW**
SNDR		**52.4 dB (f_{in}=490 MHz)**
INL (at 11 bits)		**±2 LSB**
DNL (at 11 bits)		**±0.5 LSB**
FOM		**0.097 pJ/conv**
Technology		**65 nm CMOS 1P9M**
Active ADC Area		**750 μm × 300 μm**

[7]Resistor ladder. Excludes reference generation power.

Fig. 21 shows the measured SNDR of the prototype as a function of temperature while calibration is performed at 27 °C. Two cases are investigated: 1) the circuit's and the calibration settings remain frozen and 2) the programmable feedback capacitor is incremented by one unit as the temperature exceeds 50 °C and decremented by one unit as the temperature falls below 10 °C. The latter case improves the performance and is feasible if the system-on-chip employs a temperature sensor, which is a fairly common situation in today's designs.

Table I summarizes the measured performance of the experimental prototype. It is important to note that the clock network

[5]For a supply variation from 1.2 to 1.4 V, the measured SNDR changes from 56.5 to 56.6 dB at $f_{\rm in} = 21$ MHz and from 52.4 to 52.5 dB at $f_{\rm in} = 491$ MHz.

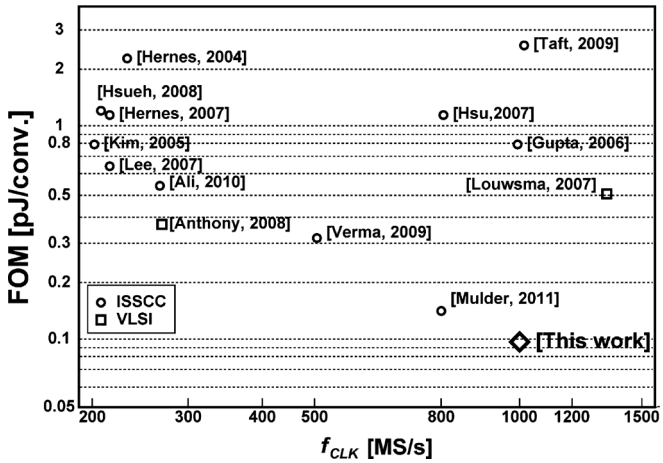

Fig. 22. FOM comparison (reference generation power consumption has been excluded for Mulder, 2011).

(the nonoverlap generator along with the buffers) is designed conservatively and consumes about 43% of the overall power. It is expected that a less conservative design can halve the clock network dissipation.

Fig. 22 compares the figure-of-merit (FOM) of various ADCs having comparable resolution and sampling rates [4], [15]–[26] with the current work. The ADC prototype reported here achieves an FOM of 97 fJ/conversion and hence about a factor of 2 improvement over the 40-nm design in [23].

VI. Conclusion

The performance of pipelined ADCs can be substantially improved by resolving a larger number of bits in the first stage, but maintaining a low MDAC gain. Such a strategy relaxes the MDAC linearity, gain error, and voltage swing requirements, allowing the use of one-stage opamps. However, the mismatch among the first stage DAC capacitors as well as the MDAC gain error must be removed. This paper proposes a calibration algorithm for these errors and demonstrates its promise by a 10-bit 1-GHz ADC that advances the state-of-the-art FOM by about a factor of 2.

Appendix

The Newton–Raphson method can obviate the need for division in the digital domain. To compute $1/\alpha$, we define a function $f(x) = 1/x - \alpha$ and iteratively seek that value of x which drives $f(x)$ toward zero. The iteration begins with an initial guess x_0 and proceeds according to the following update expression [14]:

$$x_{i+1} = x_i - \frac{f(x_i)}{f'(x_i)} = x_i(2 - \alpha x_i). \tag{20}$$

Since the initial guess for α in this case (≈ 2) is relatively close to the actual value, the technique converges easily.

Acknowledgment

The authors would like to thank Prof. D. Markovic and T. Han for synthesizing and estimating the power of the digital combiner block. They would also like to thank A. Jain and N. A. Jacob for design and FPGA implementation of the digital combiner block and the digital calibration logic, respectively. The authors gratefully acknowledge the TSMC University Shuttle Program for fabrication support.

References

[1] B. Murmann and B. E. Boser, "A 12b 75 MS/s pipelined ADC using open-loop residue amplification," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2040–2050, Dec. 2003.

[2] A. Panigada and I. Galton, "A 130 mW 100 MS/s pipelined ADC with 69 dB SNDR enabled by digital harmonic distortion correction," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3314–3328, Dec. 2009.

[3] H. van de Vel, B. A. J. Butler, H. van der Ploeg, M. Vertregt, G. J. G. M. Geelen, and E. J. F. Paulus, "A 1.2 V 250-mW 14b 100 MS/s digitally calibrated pipeline ADC in 90-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1047–1056, Apr. 2009.

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

[5] B. Sahoo and B. Razavi, "A 12-Bit 200-MHz CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 44, no. 9, pp. 2366–2380, Sept. 2009.

[6] B. Sahoo and B. Razavi, "A 10-bit 1-GHz 33 mW CMOS ADC," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2012, pp. 30–31.

[7] 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.

[8] E. Siragusa and I. Galton, "A digitally enhanced 1.8-V 15-bit 40-MSamples/s CMOS pipelined ADC," *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2126–2138, Dec. 2004.

[9] S. H. Lee and B.-S. Song, "Digital-domain calibration of multistep analog-to-digital converters," *IEEE J. Solid-State Circuits*, vol. 27, no. 12, pp. 1679–1688, Dec. 1992.

[10] A. Karanicolas, H. S. Lee, and K. Bacrania, "A 15-b 1-Msamples/s digitally self-calibrated pipeline ADC," *IEEE J. Solid-State Circuits*, vol. 28, no. 12, pp. 1207–1215, Dec. 1993.

[11] I. Ahmed and D. A. Johns, "An 11-bit 45-MS/s, pipelined ADC with rapid calibration of DAC errors in a multibit pipeline stage," *IEEE J. Solid-State Circuits*, vol. 43, no. 7, pp. 1626–1637, July 2008.

[12] S. Devarajan, L. Singer, D. Kelly, S. Decker, A. Kamath, and P. Wilkins, "A 16-bit, 125-MS/s, 385 mW, 78.7 dB SNR CMOS pipeline ADC," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3305–3313, Dec. 2009.

[13] M. Yoshioka, M. Kudo, T. Mori, and S. Tsukamoto, "A 0.8 V 10b 80 MS/s 6.2 mW pipelined ADC with regulated overdrive voltage biasing," in *ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 452–453.

[14] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C-The Art of Scientific Computing*. Cambridge, U.K.: Cambridge Univ., 1992.

[15] B. Hernes *et al.*, "A 1.2 V 220 MS/s 10b pipeline ADC implemented in 0.13 μm digital CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2004, pp. 256–257.

[16] R. C. Taft *et al.*, "A 1.8 V 1.0 GS/s 10b self-calibrating unified-folding-interpolating ADC with 9.1 ENOB at Nyquist frequency," in *ISSCC Dig. Tech. Papers*, Feb. 2009, pp. 78–79.

[17] K.-W. Hsueh *et al.*, "A 1 V 11b 200 MS/s pipelined ADC with digital background calibration in 65 nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2008, pp. 546–547.

[18] B. Hernes *et al.*, "A 92.5 mW 205 MS/s 10b pipeline IF ADC implemented in 1.2 V/3.3 V 0.13 μm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 462–463.

[19] H.-C. Kim *et al.*, "A 30 mW 8b 200 MS/s pipelined CMOS ADC using a switchedopamp technique," in *ISSCC Dig. Tech. Papers*, Feb. 2005, pp. 284–285.

[20] S.-C. Lee *et al.*, "A 10b 205 MS/s 1.0 mm^2 90 nm CMOS pipeline ADC for flat panel display applications," in *ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 458–459.

[21] A. M. Ali *et al.*, "A 16b 250 MS/s IF-sampling pipelined A/D converter with background calibration," in *ISSCC Dig. Tech. Papers*, Feb. 2010, pp. 292–293.

[22] M. Anthony *et al.*, "A process-scalable low-power charge-domain 13-bit pipeline ADC," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2008, pp. 222–223.

[23] J. Mulder *et al.*, "An 800 MS/s dual-residue pipeline ADC in 40 nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2011, pp. 184–185.

[24] S. M. Louwsma *et al.*, "A 1.35 GS/s, 10b, 175 mW time-interleaved AD converter in 0.13 μm CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, June 2007, pp. 62–63.
[25] S. Gupta *et al.*, "A 1 GS/s 11b time-interleaved ADC in 0.13 μm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2006, pp. 2360–2361.
[26] C.-C. Hsu *et al.*, "An 11-bit 800-MS/s time-interleaved ADC with digital background calibration," in *ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 464–465.

Bibhu Datta Sahoo (M'08) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Kharagpur, India, in 1998, the M.S.E.E. degree from University of Minnesota, Minneapolis, MN, USA, in 2000, and the Ph.D.E.E. degree from the University of California, Los Angeles, CA, USA, in 2009.

From 2000 to 2006, he was with the DSP Microelectronics Group, Broadcom Corporation, Irvine, CA, USA, where he designed analog and digital integrated circuits for signal-processing applications. From December 2008 to February 2010, he was with Maxlinear Inc., Carlsbad, CA, USA, where he was involved in designing integrated circuits for CMOS TV tuners. From March 2010 to November 2010, he was a Post Doctoral Researcher with the University of California, Los Angeles, CA, USA. From December 2010 to December 2011, he was an Assistant Professor with the Department of Electronics and Electrical Communication Engineering, Indian Institute of Technology, Kharagpur, India. Since December 2011, he has been an Associate Professor with the Department of Electronics and Communication Engineering, Amrita University, Amritapuri, India. His research interests include data conversion, signal processing, and analog circuit design.

Prof. Sahoo was the recipient of the 2008 Analog Devices Outstanding Student Designer Award.

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 an Adjunct Professor with Princeton University from 1992 to 1994 and with Stanford University in 1995. He was with AT&T Bell Laboratories and Hewlett-Packard Laboratories until 1996. Since 1996, he has been an Associate Professor and subsequently a Professor of electrical engineering with the University of California, Los Angeles, CA, USA. He 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). He has served as a guest editor and associate editor for the *International Journal of High Speed Electronics*. His current research includes wireless transceivers, frequency synthesizers, phase-locking and clock recovery for high-speed data communications, and data converters.

Prof. Razavi 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 a guest editor and an associate editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS and the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS. He has served as an IEEE Distinguished Lecturer. He was the recipient of 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 corecipient 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.