An 11-Bit 45 MS/s Pipelined ADC With Rapid Calibration of DAC Errors in a Multibit Pipeline Stage

Imran Ahmed, *Student Member, IEEE*, and David A. Johns, *Fellow, IEEE*

***Abstract*—A technique to rapidly correct for both DAC and gain errors in the multibit first stage of an 11-bit pipelined ADC is presented. Using a dual-ADC based approach the digital background scheme is validated with a proof-of-concept prototype fabricated in a 1.8 V 0.18 μm CMOS process, where the calibration scheme improves the peak INL of the 45 MS/s ADC from 6.4 LSB to 1.1 LSB after calibration. The SNDR/SFDR is improved from 46.9 dB/48.9 dB to 60.1 dB/70 dB after calibration. Calibration is achieved in approximately 10^4 clock cycles.**

***Index Terms*—ADC, analog-to-digital conversion, background, calibration, capacitor mismatch, CMOS, DAC, dual-ADC, missing codes, pipeline, rapid, split-ADC.**

I. Introduction

CAPACITOR mismatch is a major source of missing codes which degrade accuracy and linearity in the output of pipelined ADCs. The relative mismatch between two capacitors is inversely related to the area of the capacitors. Since the accuracy of pipelined ADCs is limited by capacitor mismatch, large sampling capacitors are required in pipelined ADCs to minimize harmonic distortion. As large sampling capacitors require large opamps to achieve a fixed unity gain frequency, relying on increased capacitors sizes to improve linearity can consume a large amount of power. Capacitor mismatch can be reduced by using additional layers to form metal–insulator–metal (MiM) capacitors, however such additional layers are limited to $\sim$0.1% mismatch, and also increase fabrication costs as MiM capacitors are rarely available in process options used in industry which tend to be optimized for digital circuits.

OFDM applications are examples of systems in which the power-linearity tradeoff in pipelined ADCs is very significant. An example output spectrum of an ADC used in OFDM applications is shown in Fig. 1. From Fig. 1, in OFDM applications the bandwidth of each signal bin is relatively narrow, where the noise floor of a single signal bin is given by $P_{\rm NF} f_{\rm BW\text{-}in}/N_{\rm ch}$, where $P_{\rm NF}$ is the power spectral density of the noise, $f_{\rm BW\text{-}in}$ the entire bandwidth of the input signal digitized by the ADC, and $N_{\rm ch}$ the total number of OFDM channels. Thus, in an OFDM system where each channel consumes only a small fraction of the overall bandwidth, a high SNR can be achieved in each signal bin without requiring a very low overall noise floor. Hence, small sampling capacitors with large kT/C noise contributions could ideally be used to minimize power in a pipelined ADC which digitizes the entire OFDM bandwidth. Harmonic distortion in the ADC however leads to interference between signal bins, thus to minimize signal interference the ADC requires low harmonic distortion. Hence, even though OFDM applications typically do not need large sampling capacitors to achieve a low noise floor, large capacitors and as a result large power consumption are required in the pipelined ADCs to maintain a suitably high linearity.

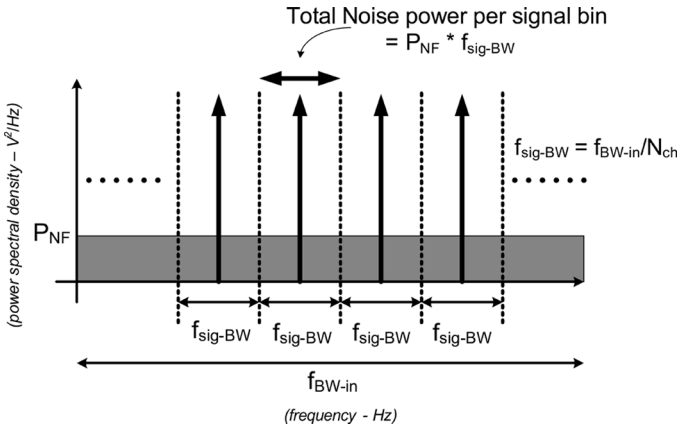

Fig. 1. Example OFDM spectrum.

Low DC opamp gain also results in missing codes in pipelined ADCs. With deeper submicron technology nodes providing less and less intrinsic gain from CMOS devices, an opamp with a large DC gain and reasonable signal swing can only be realized with potentially power-hungry DC gain enhancement techniques such as multi stage opamps (e.g., [1]) and gain boosting [2].

Many digital calibration techniques which measure and compensate the effect of missing codes have been proposed recently to overcome the power–linearity tradeoff in ADCs. As newer deep submicron technologies yield digital circuits with lower power consumption over earlier technologies due to lower supply voltages and smaller device area, digital calibration has become increasingly attractive to realize low power linear ADCs. Digital calibration techniques can be broadly categorized into either foreground (e.g., [3], [4]) or background calibration techniques (e.g., [5]–[16]).

Manuscript received November 28, 2007; revised February 8, 2008. This work was supported by the National Sciences and Engineering Research Council of Canada (NSERC) Collaborative Research Development Grant with Gennum Corporation. Fabrication services were provided by the Canadian Microelectronics Corporation (CMC).

The authors are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, M5S 3G4, Canada (e-mail: imran@eecg.toronto.edu; johns@eecg.toronto.edu).

Digital Object Identifier 10.1109/JSSC.2008.923724

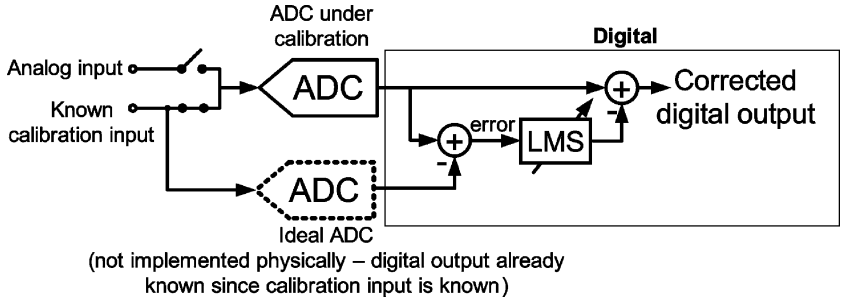

Fig. 2. Principle of foreground calibration.

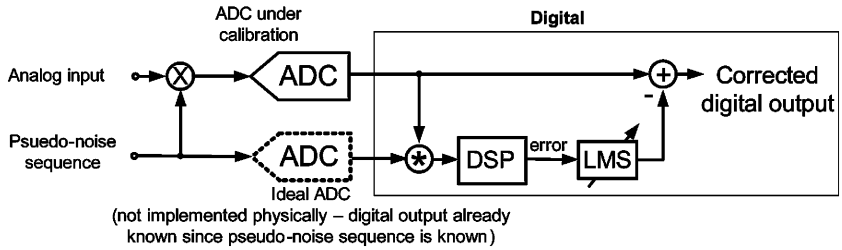

Fig. 3. Principle of background calibration.

As illustrated in Fig. 2, foreground calibration requires the operation of an ADC to be interrupted so that a known input sequence can be applied to the ADC, where by comparing the output of the ADC to the expected ADC output under ideal conditions the impact of missing codes can be quantified and corrected. The advantage of foreground schemes is that calibration can be achieved within a small number of clock cycles as the error signal labeled in Fig. 2 is highly correlated with the error sources causing the missing codes. The disadvantage of foreground calibration is that the ADC is required to be taken offline every time calibration is performed, which in some applications may not be possible.

As shown in Fig. 3, background calibration continuously measures and removes missing codes, thus has the significant advantage that the ADC is not required to be taken offline to perform calibration. Several different schemes have been proposed recently to implement background calibration, where the vast majority of the schemes use a statistics based approach to realize calibration. In a statistical scheme the input of the pipeline stage under calibration is effectively modulated by a known pseudo-random sequence, where by correlating the digital output of the ADC with the known pseudo-random sequence the impact of missing codes can be determined. To avoid significantly altering the ADC output spectrum the pseudo-noise sequence is typically made very long to avoid correlations with the analog input, as well as small in amplitude so that the injected pseudo-random sequence which appears as an additional white noise source at the output only consumes a small portion of the dynamic range. With statistics based background calibration schemes however, since the digital output of the ADC is highly correlated with the analog input and weakly correlated with the pseudo-random sequence, a large number of clock cycles are required to accurately extract the pseudo-random sequence from the digitized analog input in the ADC output. For example, in [16] $\sim 10^7$ cycles were required to achieve 13-bit linearity, and in [12] $\sim 10^8$ clock cycles were required to achieve $>$14-bit linearity. In [13] it was shown empirically that statistical techniques required on the order of 2^{2N} clock cycles to calibrate gain errors only. For 11-bit linearity approximately 4 million clock cycles are required to only correct gain errors using statistics-based background calibration.

In an industrial environment where ICs are mass produced, ICs are tested for functionality by automated testers. In ADCs which use background-statistical techniques long calibration times can lead to excessive test times, limiting IC production throughput and thus reducing revenue. For example, with 4 million calibration cycles, even with a reasonably high sampling rate of 40 MS/s, 1/10th of a second would be required at minimum to test the ADC. For higher resolution and/or lower speed ADCs the test time can be much higher [13]. In the interest of larger production throughput it is highly desirable to reduce calibration time.

Reducing calibration time has become an active area of research over the past few years. One topology which has proven to be highly effective in reducing calibration times in background schemes is the "dual-ADC" or "split-ADC" approach [9], [13], [14]. As shown in Fig. 4, the split ADC takes a single ADC and splits it into two almost identical ADCs where each ADC has half the area, and half the thermal noise floor (thus half the power) of the overall ADC. Hence, power and area of the split-ADC topology to a first order are not increased over a conventional ADC. Each ADC is identical, except the residue transfer curve of the stage under calibration in one ADC is designed differently than the other. As a result when the ADCs are free of errors both ADCs produce the same output, however when errors are included each ADC produces different outputs. Thus, the error signal used for adaptation/calibration, which is formed by the difference of the two split ADCs, is very weakly correlated with the analog input as it appears as common

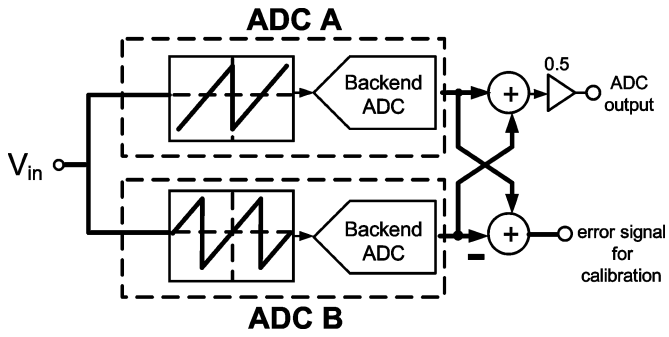

Fig. 4. Example split-ADC topology.

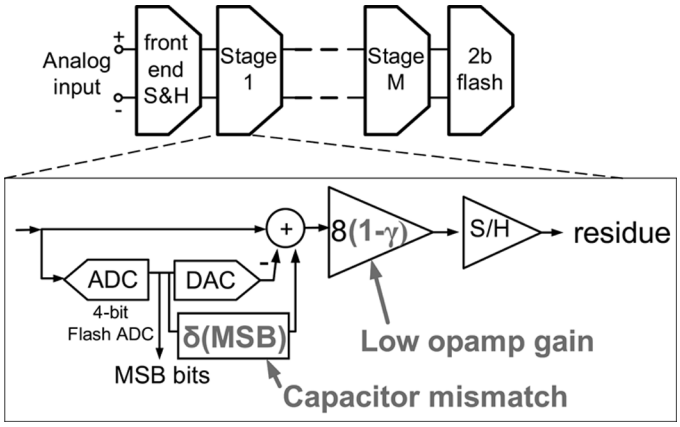

Fig. 5. Pipeline topology, first stage shown in detail including error sources.

mode to the split ADCs and very highly correlated with the error sources in the ADCs [9]. Hence, calibration can be attained in the background within a small number of clock cycles.

Previous publications used the dual ADC technique to calibrate only gain errors in 1.5 b/stage based ADCs [9], [13], [14]. In this work [17] the split-ADC approach is extended to also rapidly correct for DAC errors in multibit pipeline stages using a simple digital realization. A proof-of-concept 11-bit pipelined ADC was fabricated in a 1.8 V 0.18 μm CMOS process where the calibration scheme was used to correct DAC errors in the 4-bit first stage of each split-ADC.

The organization of the discussion in this paper is as follows: Section II reviews gain and DAC errors and their impact in pipelined ADCs. Section III discusses the architecture used to implement rapid calibration of both DAC and gain errors in a multibit stage of a pipelined ADC. Section IV discusses the circuits used to implement the ADC of this work. Section V discusses measured results of a prototype fabricated in a 1.8 V 0.18 μm CMOS process where by using the calibration scheme of this work the INL of an 11-bit ADC at 45 MS/s was improved from $+6.1/-6.4$ LSB to $+1.1/-1$ LSB after calibration. The SNDR/SFDR was improved from 46.9 dB/48.9 dB to 60.1 dB/70 dB after calibration, where calibration was achieved in $\sim 10^4$ clock cycles. Section VI concludes the paper.

II. Review of Gain and DAC Errors in Pipelined ADCs

Fig. 5 illustrates the topology of a 4-bit pipeline stage with 1-bit overlap to allow for sub-ADC errors, and Fig. 6 shows an example circuit implementation of the topology. Fig. 7 illustrates the residue transfer function of the stage when no errors are present.

A. Gain Errors

Due to mismatch between the sampling capacitors $C_0 - C_{15}$ and the feedback capacitor C_f, and also due to low DC gain from the opamp in Fig. 6, the ideal gain of 8x of a 4-bit pipeline stage is modified by $(1-\gamma)$. As shown in Fig. 8 the modified stage gain results in a fixed number of missing codes at every MSB transition (i.e., constant DNL errors or constant jumps in INL at every transition of the bits resolved by the first stage). Prior split-ADC algorithms and the majority of prior calibration schemes are concerned with estimating γ and scaling the backend digital code by $(1-\gamma)^{-1}$ to compensate for the nonideality.

B. DAC Errors

Capacitor mismatch between each of the sampling capacitors C_0–C_{15} in Fig. 6 result in errors in the pipeline stage's DAC which are a function of each MSB bit resolved. As shown in Fig. 9 DAC errors result in each linear segment of the residue transfer curve being shifted up or down by different static random values $\delta(i)$. Hence, DAC errors result in non-constant missing codes at every MSB transition, yielding substantial harmonic distortion. Since unique errors are produced at each MSB transition a gain calibration scheme (which only corrects for the average number of missing codes at each transition) cannot be used to digitally correct the ADC output. To correct DAC errors (e.g., [8], [12], [16]) a separate corrective term for *each* MSB transition is required for calibration, significantly increasing the complexity of the correction scheme over gain-only correction techniques. For example, with a 4-bit pipeline stage including 1-bit of redundancy 15 correction parameters for 16 unique DAC outputs are required to be estimated.

It is noted that in DAC calibration schemes all MSB bits are required to be exercised to measure each unique DAC error. Since an MSB bit can change at most only once every clock cycle, DAC calibration schemes can require $N-1$ times more cycles to converge (where N is the number of unique MSB outputs) than gain calibration schemes. Hence, in an industrial environment it is even more critical for calibration schemes which correct DAC errors to have short calibration times.

Comparing Figs. 8 and 9, it is noted that missing codes produced by gain errors look the same as missing codes produced by DAC errors where the DAC error is constant at every MSB transition. Thus, in a DAC calibration scheme (where the missing codes are corrected as a function of *each* MSB), the gain errors are also corrected in addition to DAC errors. Hence, the calibration scheme proposed in this work which demonstrates rapid calibration of DAC errors in a multibit pipeline stage, is a more general or super-set solution to previously published split-ADC topologies which only address the calibration of gain errors.

III. Rapid DAC+gain Error Calibration Scheme

In the calibration scheme of this work two ADCs (ADC A and ADC B) simultaneously process in parallel the same analog

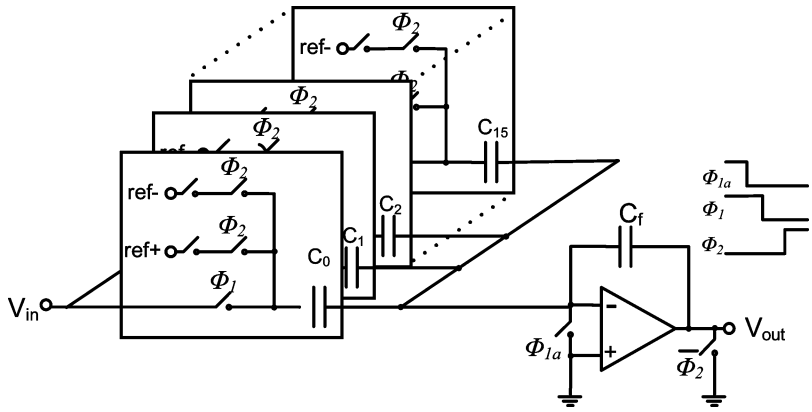

Fig. 6. Example implementation of 4-bit MDAC (shown single-ended, implemented differentially in practice).

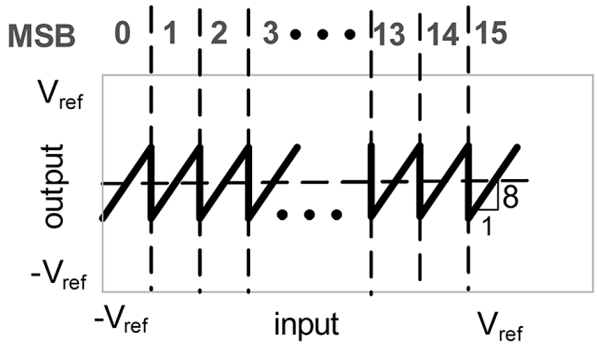

Fig. 7. Ideal residue transfer curve of 4-bit pipeline stage.

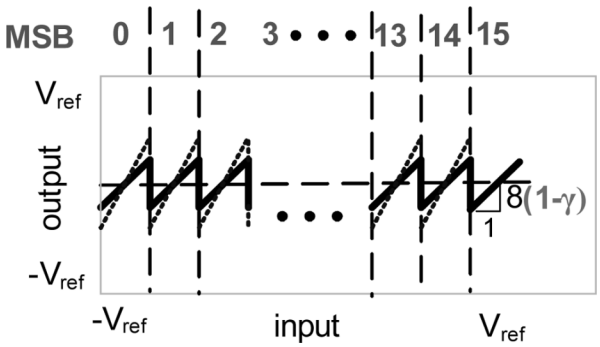

Fig. 8. Residue transfer curve showing impact of gain errors.

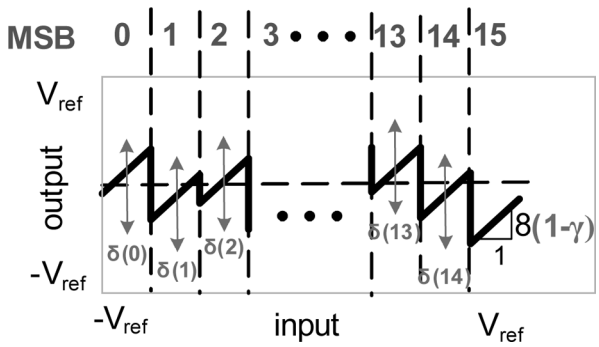

Fig. 9. Residue transfer curve showing impact of DAC and gain errors.

input as shown in Fig. 10. The final ADC output is generated by the average of the two ADC outputs, thus each ADC is designed with half the total capacitance, hence half the power and area of the overall ADC to meet thermal noise requirements [13]. From Fig. 10, ADCs A and B are identical except in each ADC the residue transfer function in the first stage is horizontally offset from the other by approximately 1/2 MSB.

A. Measurement of Missing Codes Due To DAC and Gain Errors

Fig. 11 illustrates the transfer function of key outputs from each split ADC. From Fig. 11 if the analog input to the ADC is such that $\mathrm{MSB_A} = i$, then $\mathrm{MSB_B}$ is either i or $i+1$. The offset between the digital outputs of ADCs A and B for the range of analog inputs where $\mathrm{MSB_A} = i$ and $\mathrm{MSB_B} = i$ is denoted Δ_{i1}, and Δ_{i2} where $\mathrm{MSB_A} = i$ and $\mathrm{MSB_B} = i+1$ respectively, as shown in Fig. 11. In an ideal ADC without DAC or gain errors, the difference between split ADC outputs is constant regardless of the analog input, thus $\Delta_{i1} = \Delta_{i2}$. With ideal ADCs A and B $\Delta_{i1} = \Delta_{i2} = 0$, however, the offset is shown as a constant in Fig. 11 for clarity of illustration.

If DAC and gain errors are included as shown in Fig. 12, each split ADC incurs unique missing codes wherever an MSB changes. As the MSB transitions between each split ADC are staggered, however, ADC A does not incur an error in the first pipeline stage for the same range of analog inputs as ADC B. Thus, the digital output of ADC A can be used as an *ideal reference* to measure the errors of ADC B. The difference between Δ_{i1} and Δ_{i2} precisely gives the error due to missing codes that occurs when $\mathrm{MSB_B}$ changes from i to $i+1$ as shown in Fig. 12. To minimize the effect of zero mean error sources such as thermal noise, Δ_{i1} and Δ_{i2} are averaged before subtraction using a simple first-order IIR filter with transfer function $\mu/[1-(1-\mu)z^{-1}]$. In this work $\mu = 1/64$ was used, where by implementing μ as a power of 2, multiplication by μ can be implemented using simple bit shifts. In a similar manner the unique error due to missing codes at all other MSB transitions can be measured for ADC B.

Errors due to missing codes for ADC A are measured by noting that $\Delta_{i2} - \Delta_{(i+1)1}$ is the error due to missing codes in ADC A when $\mathrm{MSB_A}$ changes from i to $i+1$ as shown in Fig. 12. Hence, the error due to missing codes in ADC A can be determined using already measured values $\bar{\Delta}_{i2}$ and $\bar{\Delta}_{(i+1)1}$. Errors due to missing codes at all other MSB transitions in ADC A are measured using an identical extension as done for ADC B.

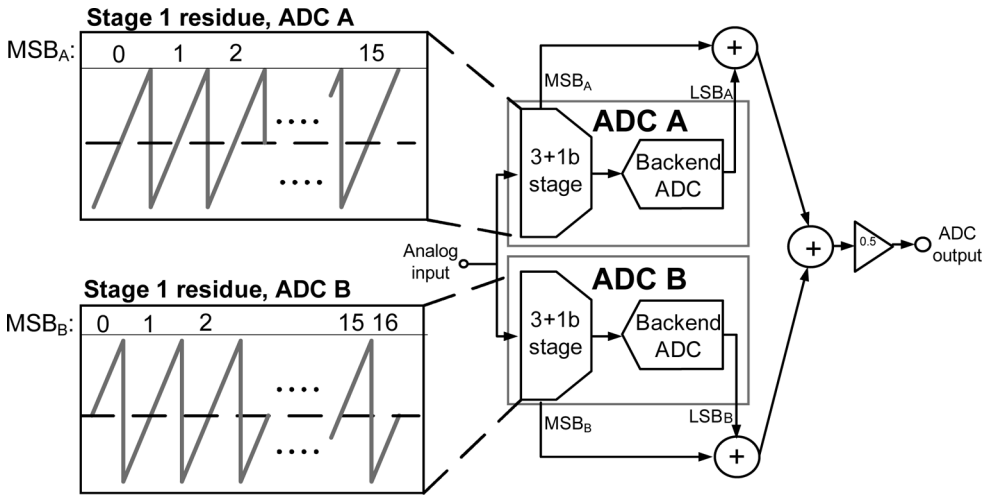

Fig. 10. Dual ADC topology of this work.

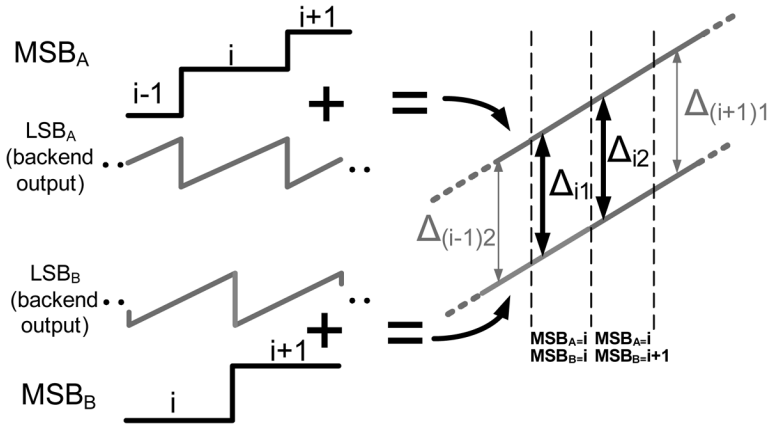

Fig. 11. Transfer curves of first stage (MSB), backend ADC (LSB) and total ADC outputs from each split ADC with no errors.

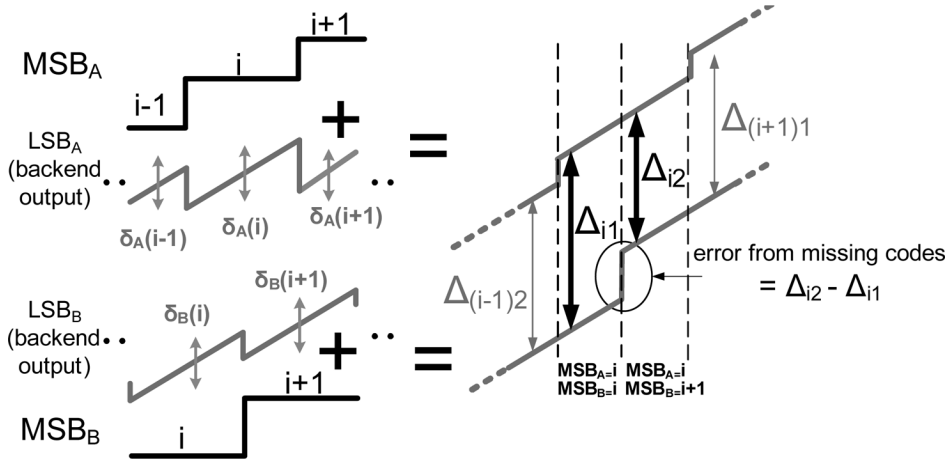

Fig. 12. Transfer curves of key ADC outputs with gain, DAC errors included.

B. Correction of Missing Codes

With the errors from missing codes at each MSB transition measured, each ADC is corrected by shifting each ADC's digital output as a function of the MSB such that overall transfer function of each ADC is free from missing codes due to errors in the first stage as shown in Fig. 13 (the same is done for ADC A).

Rapid calibration is achieved as $\bar{\Delta}_{iB}$ is highly correlated with the number of missing codes; only a small number of clock cycles are required to average out the effects of zero-mean noise. As long as the input is sufficiently busy to generate a sufficient number of estimates of Δ_{i1}, Δ_{i2}, for all i, there is no constraint on the shape of the input signal to the ADC. In contrast statistical techniques use statistical correlations which require many

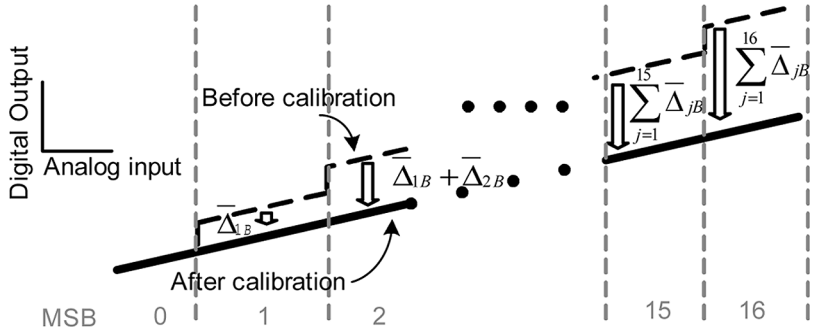

Fig. 13. Illustration of how correction terms for ADC B are derived from estimates of missing codes (correction topology of ADC A is similar).

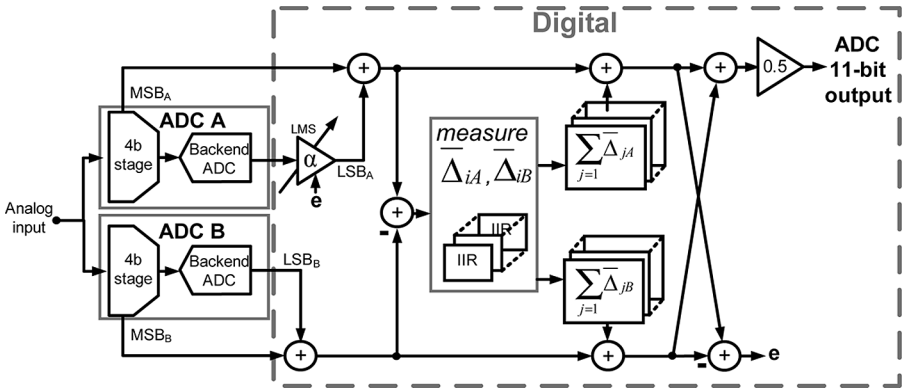

Fig. 14. Full ADC topology of this work.

output samples to extract similar information. The full topology of the ADC including calibration is shown in Fig. 14.

It is noted that the approach of this work is similar to background calibration techniques where a more accurate but slower ADC is used in parallel with the ADC under calibration (e.g., [10]), where by adaptively equalizing the outputs of the fast ADC to the output of the slow ADC calibration can be achieved. In this work since the residue transfer function of one of the split ADCs is offset from the other, ADC A does not suffer an error in the first stage for the same range of inputs as ADC B, thus one ADC can be used as an ideal reference for the other eliminating the need for one of the ADCs to be more accurate than the other. Hence, there is no need to trade higher accuracy with lower sampling rates in the second ADC; both ADCs can operate at the same speed and both ADCs can be used to digitize the analog input. Thus, the power of the additional ADC also goes towards lowering the noise floor in the digital output, unlike [10] where the additional ADC (since it operates more slowly) only aids the correction scheme. Furthermore using the technique outlined in this work both ADCs are calibrated whereas in [10] only one ADC is.

C. Mismatch Between ADCs

Due to random mismatches between the split ADCs, the first stage of ADC A in Fig. 10 will have MSB transitions that are not exactly 1/2 MSB apart, as well as a different input referred offset and different stage gain than ADC B.

The calibration scheme of this work only requires the MSB transitions between ADCs A and B to be nonoverlapping, thus a precise definition of the offset between ADCs is not necessary. As long as the comparators which define the MSB transitions for each split ADC are designed to have an offset within 1/2 MSB no errors are produced by the calibration scheme.

Constant offset between the outputs of ADCs A and B appears as a common-mode shift in both $\bar{\Delta}_{i1}$, and $\bar{\Delta}_{i2}$. Since the number of missing codes at each MSB transition is measured by subtracting $\bar{\Delta}_{i2}$ from $\bar{\Delta}_{i1}$, the common mode is eliminated and thus input-referred offsets of each split-ADC have no impact in the calibration scheme (under the practical assumption that the offsets are not large enough to saturate the output of the pipeline stages).

To account for an overall gain mismatch between the two ADCs, a least-mean-square (LMS) adaptive gain term (α in Fig. 14) is also included [13] which scales the backend code of ADC A so as to keep the outputs of ADCs A and B parallel. As ADC B provides an ideal reference for ADC A, the error signal used for the LMS adaptation (which is formed by the difference of the two ADC outputs) is highly correlated with the gain error between ADCs, thus steady state convergence of α occurs within a relatively short time interval. To ensure stability, the adaptation of α is designed to converge more slowly than the loop which measures and corrects the gain and DAC errors of the first pipeline stage.

Ultimately, the accuracy of estimation in errors due to missing codes in the first stage is limited by missing codes and distortion

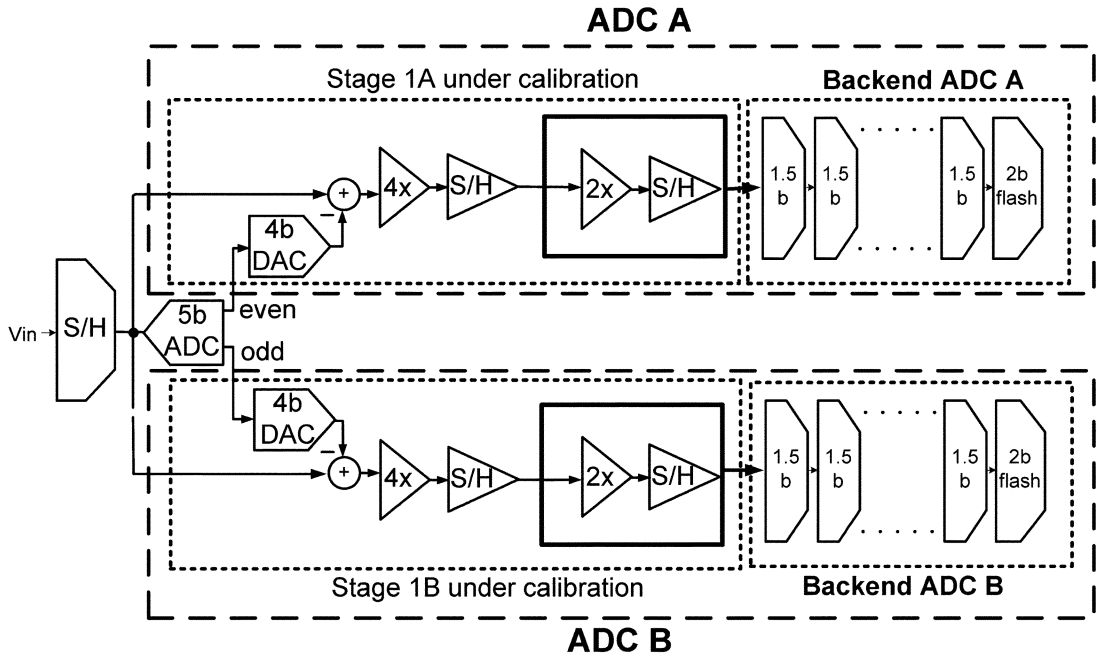

Fig. 15. Analog portion of ADC topology in detail.

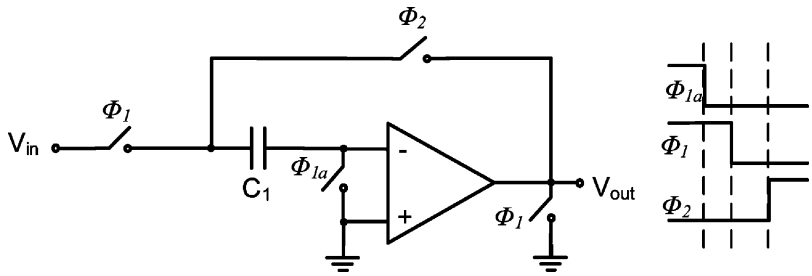

Fig. 16. Sample and hold topology (implemented fully differentially in this work).

in the backend, which although not addressed in this work can be minimized by also calibrating the backend stages.

IV. Circuit Implementation

Fig. 15 shows the architecture of the analog portion of the 11-bit pipeline ADC of this work. Scaling was used in the first three pipeline stages to reduce power consumption [18]. An additional three bits are added to the backend in each split-ADC to improve the accuracy in error estimation. As the last stages in a pipelined ADC consume only a small fraction of the total power, adding extra stages to reduce quantization noise has a minimal impact on ADC power.

A. Front-End Sample-and-Hold

To ensure both ADCs operate on the same analog input, a conventional flip-around front-end sample-and-hold [19] was used and is shown in Fig. 16. The flip-around topology has the advantage of a feedback factor near unity hence is power efficient. The sample-and-hold was designed such that the entire ADC had approximately an 11-bit input referred thermal noise floor relative to a 1.3 V peak-to-peak input sinusoid.

B. 5-Bit Flash ADC

The MSB transitions for each split ADC were generated by a single 5-bit Flash ADC where the even numbered outputs were used for ADC A and odd numbered outputs for ADC B. The 5-bit Flash ADC was implemented using an array of comparators where a single comparator is shown in Fig. 17. Each comparator used a preamp before the slicer to reduce kickback noise from the slicer on to the reference voltages which were implemented with a resistor string. A switched capacitor based topology [20] was used to implement the threshold of the comparator as well as reduce offset from the preamplifier.

From Fig. 15 it is noted that one of the overheads of a split-ADC topology is an increase in resolution of the sub-ADC of the stage under calibration by 1-bit (to ensure the two ADCs do not have overlapping MSB transitions). However since the power of the sub-ADC is only a small fraction of overall ADC power, increasing the resolution of the sub-ADC by 1-bit results in only a small overall increase in analog power. Furthermore for low resolutions in the sub-ADC, as is the case in this work, the power of the sub-ADC is not limited by thermal noise—thus increased resolution can be obtained with only a marginal increase in power so as to sufficiently minimize comparator offsets.

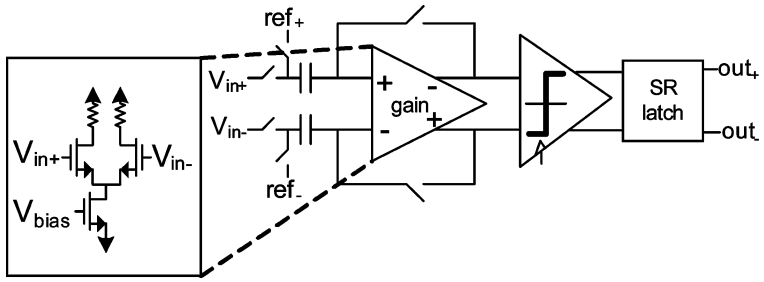

Fig. 17. Comparator topology used in 5-bit Flash sub-ADC.

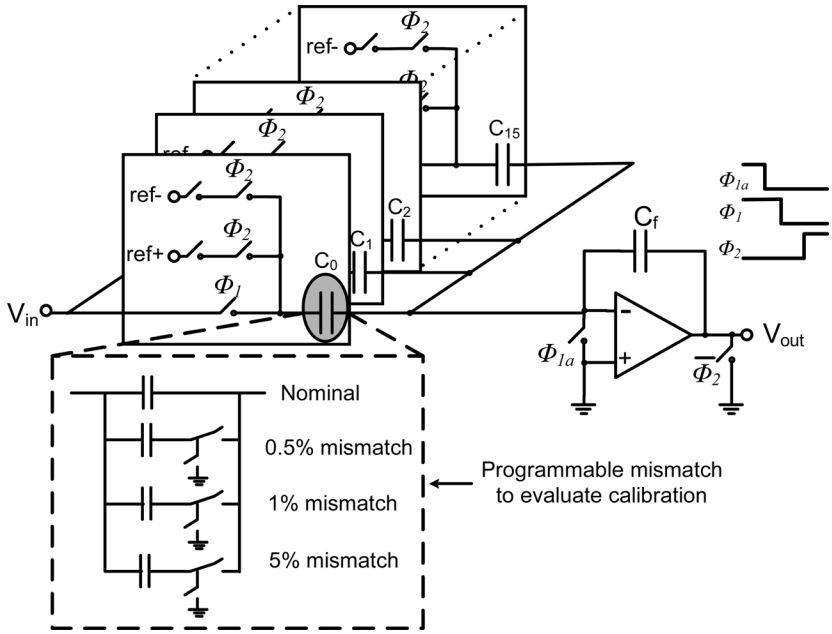

Fig. 18. Topology of first stage MDAC (implemented fully differentially in this work).

C. 4-Bit MDAC

The topology of the multibit MDAC in the first stage of ADC A is shown in Fig. 18. The MDAC for ADC B is identical except slightly modified to account for a 1/2 MSB horizontal offset.

As capacitor mismatch is a process variation its impact can only be properly observed by testing thousands of chips for random variation. To minimize the number of ICs tested each sampling capacitor in the first stage of each split ADC was made individually digitally programmable with 0.5%, 1%, and 5% deviations from nominal values as shown in Fig. 18. Thus, the effectiveness of the calibration scheme of this work could be very quickly evaluated by programming different mismatch profiles in the first stage of each split ADC. The control bits to set the mismatch profile of the first stage in each ADC were set via an on-chip shift-register.

Since the calibration scheme of this work also corrected gain errors, the DC gain of the opamp in Fig. 18 was designed with the minimum DC gain required to adequately suppress distortion from the opamp. The opamp, which had a DC gain of ~50 dB, was implemented using a one-stage p-input folded cascode topology, and is illustrated in Fig. 19. Common mode feedback was implemented using a switched-capacitor-based approach [20].

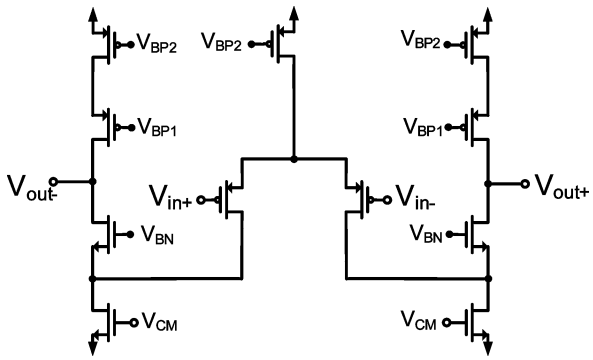

Fig. 19. Opamp used in first stage MDAC.

In a 4-bit pipeline stage including 1-bit of redundancy, the ideal closed loop gain of the MDAC is 8x. It is noted however that the feedback factor β of the MDAC of Fig. 18 is given by

$$\beta = \frac{C_f}{\sum_{i=0}^{15} C_i + C_p + C_f}$$

where C_p is the parasitic capacitance at the input of the opamp. The time constant of the settling time τ of the MDAC given by

$$\tau = \frac{1}{\beta \omega_{ta}}$$

where ω_{ta} is the open-loop unity gain frequency of the opamp. Thus, as the closed loop gain becomes larger C_f becomes smaller and the feedback factor β becomes more sensitive to variations in parasitic capacitance at the input of the opamp. Furthermore, small feedback factors require the opamp to have a large open-loop unity gain frequency to achieve fast closed loop settling. Opamps requiring high unity gain frequencies can approach the limits of their technology, making the design less power efficient. In this design although measurements are presented at 45 MS/s, the design was targeted for a $\sim$2x higher sampling rate. Limitations in the test setup restricted measurements to half the target sampling rate. Thus, in this design to reduce the open-loop opamp unity gain frequency (with the higher sampling rate in mind) as well as sensitivity to parasitic capacitors, the closed loop gain of the MDAC was designed to be 4x rather than 8x. To ensure the backend ADC had a full scale analog input, a pipelined 2x gain stage was inserted between the 4-bit MDAC and backend ADC. The gain error of the 2x stage is effectively lumped into the gain error of the 4-bit MDAC, hence its nonidealities are also calibrated.

It is noted that for ease of implementation in this prototype the calibration scheme requires the input to be sufficiently busy to excite each sampling capacitor to achieve full calibration. However it is conceivable to relax the input signal swing requirement by rotating the sampling capacitors in a known sequence not correlated with the analog input, so that (e.g.) different sampling capacitors can be used to sample a DC input, hence exercise the range of different MSB transitions.

D. Backend Pipelined ADC

The backend ADC was implemented using a cascade of nine standard 1.5-bit/stage pipeline stages followed by a 2-bit Flash ADC as shown in Fig. 15. Gain-boosted [2] folded cascode opamps with a DC gain of 95 dB were used in stages 2 and 3 to minimize missing codes in the backend. The remaining pipeline stages used folded-cascode-based opamps with DC gains of 50 dB.

E. Digital Calibration

To enhance flexibility in the test setup, the digital calibration engine was implemented off chip, where the digital outputs of each split ADC stage were taken off chip and imported into Matlab via a logic analyzer. The digital outputs of the fabricated chip were input to a model of the digital calibration scheme outlined in this work. Since there is no feedback or direct interaction with the analog portions of the ADCs in the calibration scheme of this work, operating on the off-chip digital outputs verifies the calibration scheme without any loss of generality.

Since the missing codes at each MSB transition for both split ADCs are required to be measured, $16 + 15 = 31$ unique missing codes must be measured and estimated. Thus, in a potential on-chip digital implementation each Δ_i is required to be passed through an IIR filter, and the digital output offset by each missing code as shown in Figs. 20 and 21. It is noted that each averaging block in Fig. 21 is only updated once every time its particular MSB bit is excited, i.e., has low activity. Thus, the power of each path in Fig. 21 is fairly low since with a reasonably random input each MSB bit is excited at a small fraction of the overall sampling rate.

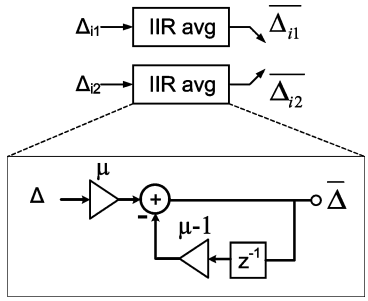

Fig. 20. Detail of IIR filter blocks used in "estimate error" block of Fig. 21.

V. Measured Results

A prototype of the ADC architecture as shown in Fig. 22 was fabricated in a 1.8 V 0.18 μm CMOS process. The area of the analog core was 2.1 mm $\times$ 1.7 mm $=$ 3.57 mm^2. It is noted that the area of the first pipeline stage in each ADC has a large overhead due to the programmability of each sampling capacitor. In a practical design which does not have the additional test circuitry the area of the first pipeline stage could be significantly reduced. As described in Section IV, the ADC was operated at $f_s = 45$ MS/s where the power of the fabricated ADC was 81 mW plus 9.5mW for the reference voltage resistor string used in the 5-bit Flash ADC. It is noted that since the design was optimized for a $\sim$2x higher sampling rate the power could easily be reduced if optimized for 45 MS/s.

Figs. 23 and 24 show the INL and DNL of the ADC before and after calibration, where each DAC element in each split ADC was programmed with a mismatch of either: 0%, 0.5%, 1%, or 5%. Using the calibration scheme of this work, the INL was improved from $+6.1/-6.4$ LSB to $+1.1/-1$LSB after calibration, and DNL from $+1.1/-0.4$ LSB to $+0.45/-0.4$ LSB. It is noted that the residual INL errors after calibration are due primarily to distortion from the backend ADCs, as well as distortion from the front-end sample and hold which sets the best achievable linearity for both ADCs.

Fig. 25 shows an FFT of the ADC output before and after calibration with a 1.3 Vp-p sinusoid at 2.39 MHz. The SNDR/SFDR of the ADC is improved from 46.9 dB/48.9 dB to 60.1 dB/70 dB. The calibration technique was verified with different programmed mismatches in the first stage as well as different full scale inputs (e.g., sinusoidal, random). In all cases successful background calibration was attained in a short time interval. Fig. 26 illustrates the variation of SNDR and SFDR with input frequency before and after calibration, where the

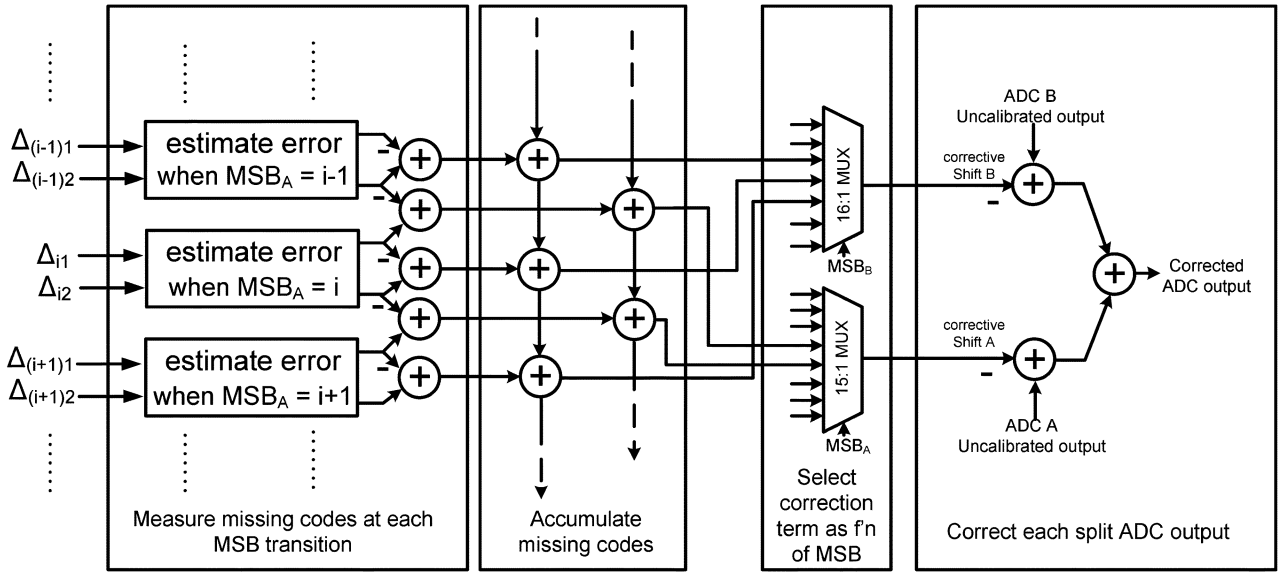

Fig. 21. Digital implementation of calibration.

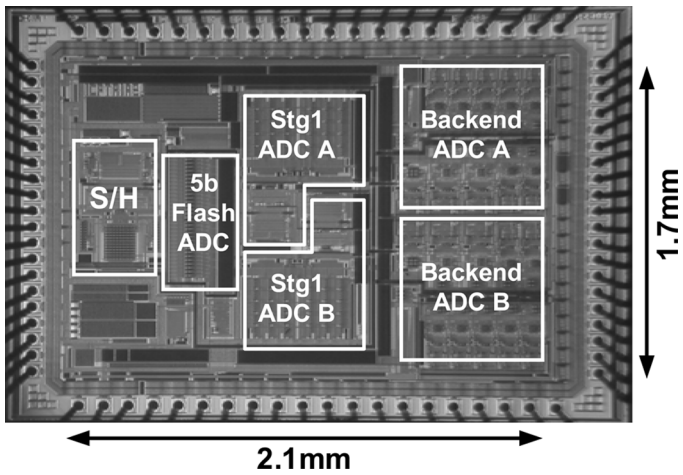

Fig. 22. Micrograph of fabricated IC in 1.8 V 0.18 μm CMOS.

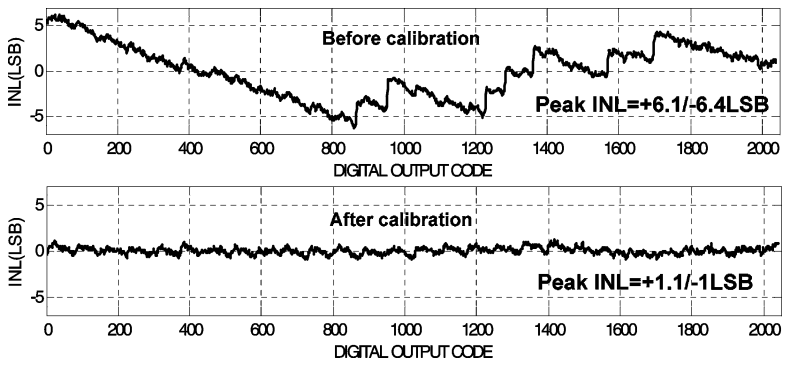

Fig. 23. INL before and after calibration.

fall-off of accuracy for higher input frequencies was attributed to the input switch in the front-end sample and hold.

The improvement of ADC SNDR and SFDR is plotted versus total number of calibration cycles in Fig. 27, where it is shown that steady state is attained within $\sim 1 \times 10^4$ clock cycles or effectively 0.22 ms (with $\mu = 1/64$). From [13] it is noted that if a statistics-based background calibration technique were used $\sim 4 \times 10^6$ clock cycles would be required to correct only the gain

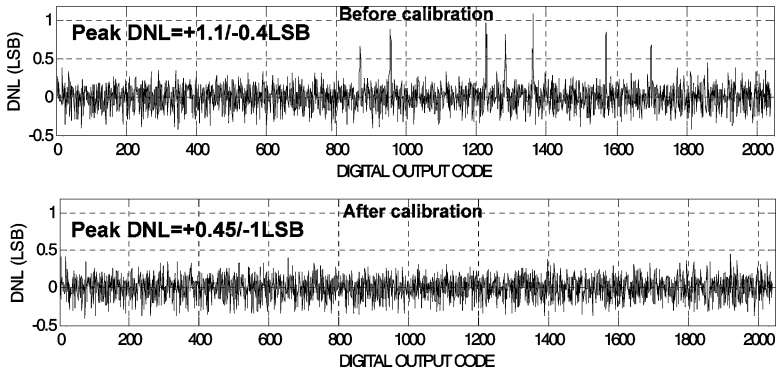

Fig. 24. DNL before and after calibration.

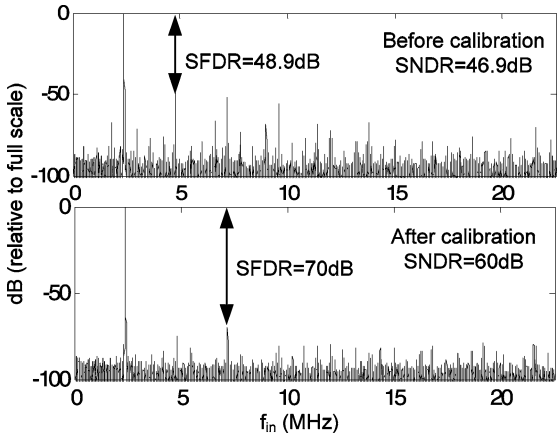

Fig. 25. FFT of ADC output before and after calibration.

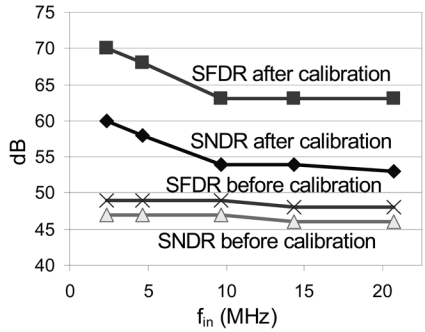

Fig. 26. Variation of ADC SNDR, SFDR with input frequency, before and after calibration.

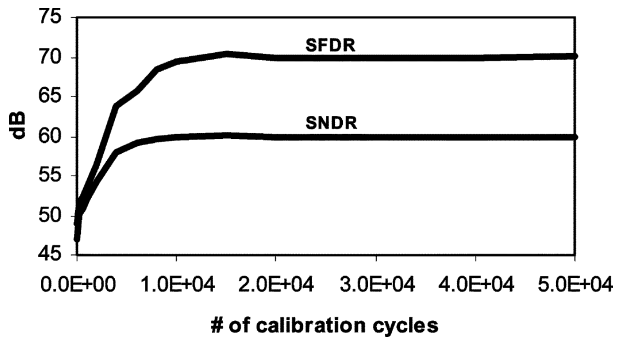

Fig. 27. ADC SNDR, SFDR improvement with number of calibration cycles.

TABLE I
SUMMARY OF ADC PERFORMANCE

f_s=45MS/s (f_{in}=2.39MHz)		
	Before Calibration	After Calibration
INL (LSB)	+6.4/-6.1	+1.1/-1
DNL (LSB)	+1.1/-0.4	+0.45/-0.4
SFDR	48.9 dB	70 dB
SNDR	46.9 dB	60 dB
Power (Analog core)	81 mW	
Area	3.57 mm^2	
# of calibration cycles	10^4 cycles (0.22ms)	
Input signal swing	1.3V p-p	
Technology	1.8V, 0.18μm CMOS	

error. In this work both gain *and* DAC errors are corrected in less than 1/100th the number of total calibration cycles. Table I summarizes key measurements of the work.

VI. CONCLUSION

A technique to rapidly measure and correct for missing codes introduced by a multibit first stage in a pipelined ADC was presented. Measured results from a prototype in 1.8 V 0.18 μm CMOS show an improvement in INL from $+6.1/-6.4$ LSB to $+1.1/-1$ LSB, and SNDR/SFDR from 46.9 dB/48.9 dB to 60.1 dB/70 dB using the calibration approach of this work. Calibration was achieved in $\sim 10^4$ clock cycles.

REFERENCES

[1] F. You, S. H. K. Embabi, and E. Sanchez-Sinencio, "Multistage amplifier topologies with nested Gm-C compensation," *IEEE J. Solid-State Circuits*, vol. 32, no. 12, pp. 2000–2011, Dec. 1997.

[2] K. Bult and G. J. G. M. Geelen, "A fast settling CMOS opamp for SC circuits with 90-dB DB gain," *IEEE J. Solid-State Circuits*, vol. 25, no. 12, pp. 1379–1384, Dec. 1990.

[3] D. Y. Chang, J. li, and U. Moon, "Radix-based digital calibration techniques for multi-stage recycling pipelined ADCs," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 51, no. 11, pp. 2133–2140, Nov. 2004.

[4] C. R. Grace, P. J. Hurst, and S. H. Lewis, "A 12b 80 MS/s pipelined ADC with bootstrapped digital calibration," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2004, pp. 460–539.

[5] U. Moon and B. S. Song, "Background digital calibration techniques for pipelined ADC's," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 44, pp. 102–109, Feb. 1997.

[6] S.-U. Kwak, B.-S. Song, and K. Bacrania, "A 15-b, 5-Msample/s low-spurious CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 32, no. 12, pp. 1866–1875, Dec. 1997.

[7] O. E. Erdogan, P. J. Hurst, and S. H. Lewis, "A 12-b digital-background-calibrated algorithmic ADC with −90 dB THD," *IEEE J. Solid-State Circuits*, vol. 34, no. 12, pp. 1812–1820, Dec. 1999.
[8] I. Galton, "Digital cancellation of D/A converter noise in pipelined A/D converters," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 47, no. 3, pp. 185–196, Mar. 2000.
[9] J. Li and U. Moon, "Background calibration techniques for multistage pipelined ADCs with digital redundancy," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 50, no. 9, pp. 531–538, Sep. 2003.
[10] Y. Chiu, C. W. Tsang, B. Nikolic, and P. R. Gray, "Least mean square adaptive digital background calibration of pipelined analog-to-digital converters," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 51, no. 1, pp. 38–46, Jan. 2004.
[11] X. Wang, P. J. Hurst, and S. H. Lewis, "A 12-bit 20-Msample/s pipelined analog-to-digital converter with nested digital background calibration," *IEEE J. Solid-State Circuits*, vol. 39, no. 11, pp. 1799–1808, Nov. 2004.
[12] E. Siragusa and I. Galton, "A digitally enhanced 1.8-V 15-bit 40-MSample/s CMOS pipelined ADC," *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2126–2138, Dec. 2004.
[13] J. McNeill, M. Coln, and B. Larivee, "A split-ADC architecture for deterministic digital background calibration of a 16b 1 MS/s ADC," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2005, vol. 1, pp. 276–598.
[14] L. Jipeng, G.-C. Ahn, D.-Y. Chang, and U. Moon, "A 0.9-V 12-mW 5-MSPS algorithmic ADC with 77-dB SFDR," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 960–969, Apr. 2005.
[15] H.-C. Liu, Z.-M. Lee, and J.-T. Wu, "A 15-b 40-MS/s CMOS pipelined analog-to-digital converter with digital background calibration," *IEEE J. Solid-State Circuits*, vol. 40, no. 5, pp. 1047–1056, May 2005.
[16] S. Ray and B.-S. Song, "A 13-b linear, 40-MS/s pipelined ADC with self-configured capacitor matching," *IEEE J. Solid-State Circuits*, vol. 42, no. 3, pp. 463–474, Mar. 2007.
[17] I. Ahmed and D. A. Johns, "An 11-bit 45 MS/s pipelined ADC with rapid calibration of DAC errors in a multi-bit pipeline stage," in *Proc. 33rd European Solid-State Circuits Conf. (ESSCIRC)*, Sep. 2007, pp. 147–150.
[18] P. T. F. Kwok *et al.*, "Power optimization for pipeline analog-to-digital converters," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 36, no. 5, pp. 549–553, May 1999.
[19] W. Yang, D. Kelly, L. 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.
[20] D. Johns and K. Martin, *Analog Integrated Circuit Design*. New York: Wiley, 1997.

Imran Ahmed (S'00) received the B.A.Sc. degree (with honors) from the Division of Engineering Science in 2002 and the M.A.Sc. degree in the Department of Electrical and Computer Engineering in 2004, both at the University of Toronto, Toronto, ON, Canada. Currently he is working towards the Ph.D. degree at the same university.

Since 2000, he has spent a total of 24 months as a student intern at Snowbush Microelectronics, Toronto, working on various leading edge mixed-signal circuits. His research interests include reconfigurable analog design, digital calibration of analog circuits, and very low-power mixed-signal circuits. From August to November 2007, he was an intern at Broadcom Netherlands, in Bunnik. At Broadcom he worked on developing very low-power pipelined ADC architectures.

Mr. Ahmed is a recipient of the Natural Sciences and Engineering Research Council of Canada (NSERC) Canadian Graduate Scholarship (CGS), Post Graduate Scholarship (PGS), Ontario Graduate Scholarship (OGS), and University of Toronto Fellowship. He received first place in the operational category and best overall submission in the 2005 DAC/ISSCC student design competition. His work at ESSCIRC 2007 was awarded the "Young Scientist Award." He is also the recipient of the 2008 Analog Devices Outstanding Student Designer Award.

David A. Johns (S'81–M'89–SM'94–F'01) received the B.A.Sc., M.A.Sc., and Ph.D. degrees from the University of Toronto, Toronto, ON, Canada, in 1980, 1983, and 1989, respectively.

In 1988, he joined the University of Toronto where he is currently a full Professor. He has ongoing research programs in the general area of analog integrated circuits with particular emphasis on digital communications, oversampling, signal processing, PLLs, ADCs, DACs, and adaptive filtering. His research work has resulted in more than 40 publications as well as the 1999 IEEE Darlington Award. He is coauthor of a textbook entitled *Analog Integrated Circuit Design* (Wiley, 1997) and has given numerous industrial short courses. Together with academic experience, he has four years of semiconductor industrial experience during 1980, 1983–85, and 1995, and is co-founder of the microelectronics company Snowbush in Toronto.

Dr. Johns served as an associate editor for IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS (TCAS) PART II from 1993 to 1995 and for TCAS PART I from 1995 to 1997. His homepage is http://www.eecg.toronto.edu/~johns.