UNIVERSITY OF CALIFORNIA, SAN DIEGO

Harmonic Distortion Correction in Pipelined Analog to Digital Converters

A dissertation submitted in partial satisfaction of the requirements for the degree
Doctor of Philosophy

in

Electrical Engineering (Electronic Circuits and Systems)

by

Andrea Panigada

Committee in charge:

Professor Ian Galton, Chair
Professor James Buckwalter
Professor William S. Hodgkiss
Professor Lawrence E. Larson
Professor Thomas T. Liu

2009

The dissertation of Andrea Panigada is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____

Chair

University of California, San Diego

2009

To Rossella, Giulia and Tommaso.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to thank my advisor Professor Ian Galton. If I hadn't met him, I would have never thought of even starting a PhD program. His guidance, his enthusiasm and his daily advice have determined the success of this project.

I would like to thank all my lab colleagues for their friendship and support. Special thanks to Nevena Rakuljic, for her help in the layout, her technical advice, and her daily moral support.

I want to thank my wife Rossella, for giving me two wonderful children, for making every day of my life worthwhile and for having unconditional faith in me. Even in the toughest times of my graduate career, she managed to maintain the nicest atmosphere in the family, which is where all my energy and resources come from.

Finally I want to thank my parents, for teaching me the value of freedom.

Chapter 1, in full, has been published in the IEEE Transactions on Circuits and Systems I: Regular Papers, volume 53, number 9, pages 1885-1895, September 2006. The dissertation author is the primary investigator and author of this paper. Professor Ian Galton supervised the research which forms the basis for this paper.

Chapter 2, in full, has been submitted for review to the IEEE Journal of Solid-State Circuits. The dissertation author is the primary investigator and author of this paper. Professor Ian Galton supervised the research which forms the basis of this paper.

VITA

1999	Laurea degree in Electronic Engineering, University of Pavia, Italy
2000 – 2002	Design Engineer, STMicroelectronics, Pavia, Italy
2003	Visiting Scholar, University of California, San Diego
2004	Design Engineer, STMicroelectronics, Pavia, Italy
2005 – 2009	Graduate Student Researcher, University of California, San Diego
2009	Master of Science, University of California, San Diego
2009	Doctor of Philosophy, University of California, San Diego

ABSTRACT OF THE DISSERTATION

Harmonic Distortion Correction in Pipelined Analog to Digital Converters

by

Andrea Panigada

Doctor of Philosophy in Electrical Engineering (Electronic Circuits and Systems)

University of California, San Diego, 2009

Professor Ian Galton, Chair

Pipelined analog to digital converters are widely used in telecommunication systems and instrumentation systems, where wide bandwidth analog input signals need to be converted into medium to high resolution digital signals.

A pipelined analog to digital converter is sensitive to distortion introduced by its residue amplifiers, because such distortion leaks into the digital output signal, thus affecting the converter resolution. To reduce distortion, high performance operational

amplifiers are usually required in the first few pipeline stages, but this causes the power consumption, the area occupation and therefore the cost of the converter to increase. An alternative approach is to design low performance operational amplifiers to reduce area and power, and compensate for the distortion they introduce by calibrating the signal in the digital domain.

This dissertation presents a new digital background calibration technique called Harmonic Distortion Correction, which allows the estimation and correction of the distortion introduced by residue amplifiers in pipelined analog to digital converters. Implemented in a prototype pipelined analog to digital converter together with another digital calibration technique known in literature as DAC Noise Cancellation, Harmonic Distortion Correction has been proven to facilitate low-voltage operation and to enable reductions in power consumption relative to comparable conventional state-of-the-art pipelined analog to digital converters.

Chapter 1 provides a mathematical model for the analysis of the distortion introduced by residue amplifiers in pipelined analog to digital converters, outlines the theory behind the Harmonic Distortion Correction algorithm, and presents the behavioral model of an example pipelined analog to digital converter implementing such technique.

Chapter 2 presents a pipelined analog to digital converter integrated circuit prototype implementing Harmonic Distortion Correction and DAC Noise Cancellation, describes the system level and circuit level design issues and solutions, and provides the prototype measurement results.

Chapter 1

Digital Background Correction of Harmonic Distortion in Pipelined ADCs

Abstract—Pipelined ADCs are sensitive to distortion introduced by the residue amplifiers in their first few stages. Unfortunately, residue amplifier distortion tends to be inversely related to power consumption in practice, so the residue amplifiers usually are the dominant consumers of power in high-resolution pipelined ADCs. This paper presents a background calibration technique that digitally measures and cancels ADC error arising from distortion introduced by the residue amplifiers. It allows the use of higher-distortion and, therefore, lower-power residue amplifiers in high-accuracy pipelined ADCs, thereby significantly reducing overall power consumption relative to conventional pipelined ADCs.

I. Introduction

Pipelined ADCs are widely used in applications that require data converters with resolutions in the range of 10 to 16 bits and bandwidths in the range of 15 to 250 MHz [1– 14]. Such applications include cellular telephone base station receivers, 802.11 wireless LAN receivers, and 802.16 wireless metropolitan area network receivers. In general, pipelined ADCs are attractive when the required bandwidth is too high for oversampling delta-sigma ADCs to be efficient and the required resolution is too high for flash ADCs to be efficient.

Unfortunately, the power consumption of high-resolution pipelined ADCs tends to be large, mainly because of the high-performance op-amps required in the first few pipeline stages. Passive sampling can be used to avoid having an op-amp based sample-and-hold in the first stage, which leaves the op-amps in the residue amplifiers of the first few stages as the dominant consumers of power [15, 16]. Each stage in a pipelined ADC performs coarse digitization of its input signal, but the outputs of the stages are combined such that most of the quantization noise cancels to achieve a high-resolution digitized version of the input signal. However, distortion introduced by the residue amplifiers, particularly those in the first few stages, results in imperfect cancellation which reduces the linearity of the pipelined ADC and increases its noise floor. In general high op-amp gain and bandwidth are required to achieve sufficiently low-distortion closed-loop residue amplifier performance. If it were not for this limitation, much lower-performance op-amps could be used in pipelined ADCs to significantly reduce power consumption.

This paper presents a digital background calibration technique, called the *harmonic distortion correction (HDC) technique*, that digitally measures and cancels ADC error arising from distortion introduced by the residue amplifiers. This makes it possible to reduce the power consumption of the op-amps in a given pipelined ADC without sacrificing ADC accuracy. The HDC technique operates in background during normal operation of the pipelined ADC, so it adapts to environmental changes without the need to interrupt normal operation of the ADC. As with other digital calibration techniques, such as presented in [17] and [18], the HDC technique requires a

significant amount of digital signal processing. However, the reduction in op-amp power consumption is expected to far exceed the increase in power consumption from the extra digital logic.

The HDC technique is based on a different principle of operation than the only other techniques known to the authors that cancel harmonic distortion in pipelined ADCs [19, 20]. The benefit of the HDC technique relative to that presented in [19] is that it works for any pipelined ADC input signal, and the benefits relative to that presented in [20] are that it does not have restrictions with respect to dc input signals and it is not sensitive to amplifier offsets .

The paper consists of three main sections. Section II presents an example pipelined ADC architecture and describes the residue amplifier distortion problem. Section III presents the signal processing details underlying the HDC technique. Section IV presents an implementation example of the HDC technique.

II. The Residue Amplifier Distortion Problem

A. An Example Pipelined ADC

A seven-stage example pipelined ADC architecture is shown in Figure 1. Each stage except the last consists of a 9-level flash ADC, a 9-level DAC, and a *residue amplifier* with a gain of 4. The last stage consists of just a 9-level flash ADC. All the flash ADCs and DACs are clocked simultaneously at a sample rate of $f_s = 1/T_s$. The ideal behavior of each flash ADC is to update its digital output each sample time to whichever of the 9 values,-4Δ, -3Δ, ..., 4Δ, is closest to the input voltage at that

sample time, where Δ is the quantization *step-size* of the flash ADC. Therefore, from a signal processing point of view each flash ADC ideally acts as a 9-level uniform quantizer, and the output of the *k*th flash ADC is given by

$$x_k[n] = v_{in\,k}(nT_s) + e_{ADC\,k}[n], \tag{1}$$

where $v_{in\ k}(t)$ is the flash ADC's input signal, and $e_{ADC\ k}[n]$ is the *quantization error* introduced by the flash ADC. The input *no-overload range* of each flash ADC, and, therefore, the *usable input range* of each pipeline stage, is −4.5Δ to 4.5Δ, because the magnitude of the *quantization error* introduced by the flash ADC is bounded by Δ/2 for input voltages within this range and exceeds Δ/2 otherwise. The ideal behavior of each DAC is to convert the format of its input from a digital representation (e.g., bits) to an analog representation (e.g., voltage) without introducing distortion or noise. Therefore, from a signal processing point of view an ideal DAC performs no numerical operation. It follows that in the absence of non-ideal circuit behavior the input to and output of the *k*th residue amplifier at the *n*th sample time are given by

$$v_k(nT_s) = -e_{ADC\,k}[n], \qquad \text{and} \qquad v_{in\ k+1}(nT_s) = 4v_k(nT_s), \tag{2}$$

respectively.

The outputs of the flash ADCs are combined as shown in Figure 1 to form the output of the pipelined ADC, $x_{out}[n]$. The output, $r_k[n]$, of each digital divide-by-four block is called the *digitized residue* of the *k*th stage. As can be seen from Figure 1, $r_k[n] = (r_{k+1}[n] + x_{k+1}[n])/4$ for $k = 1, 2, ..., 5$, so recursive application of (28) and (2) gives

$$r_k[n] = v_k(nT_s) + \frac{1}{4^{7-k}} e_{ADC\,7}[n]. \tag{3}$$

Hence, in the absence of non-ideal circuit behavior the quantization error sequences from all but the last pipeline stage cancel to give

$$x_{out}[n] = v_{in}(nT_s) + \frac{1}{4096} e_{ADC\,7}[n]. \tag{4}$$

Since $e_{ADC\,7}[n]$ is bounded in magnitude by Δ/2 and the first pipeline stage has a usable input range of −4.5Δ to 4.5Δ, this represents slightly more than 15-bit analog-to-digital conversion accuracy.

With ideal circuit behavior, the magnitude of the quantization error from each flash ADC is bounded by Δ/2, so the analog output of each pipeline stage ideally never exceeds 2Δ in magnitude. However, non-ideal circuit behavior such as comparator offset voltages can cause the analog outputs of the pipeline stages to have magnitudes that exceed 2Δ from time to time. To accommodate such *over-range* conditions, the useable input range of the second through seventh pipelined stages is maintained at −4.5Δ to 4.5Δ instead of −2Δ to 2Δ. In this case, the pipelined ADC is said to have an *over-range* margin of ±2.5Δ. The over-range margin greatly relaxes the performance requirements of the flash ADCs in pipelined ADCs [21].

B. Effect of Residue Amplifier Distortion

The effect of residue amplifier distortion can be demonstrated by considering the pipelined ADC of Figure 1 with all ideal components except for the residue amplifier in the first stage. This scenario is shown in Figure 2, wherein a function, *f*, is

used to represent distortion introduced by the first stage's residue amplifier.

The distortion introduced by a practical residue amplifier tends to be well-modeled as a memoryless, weakly non-linear function of the amplifier's input voltage, so it can be approximated accurately by its first N Taylor series coefficients where N typically is small (e.g., $N \leq 5$ is common). Consequently, the distortion function, f, in Figure 2 is given by

$$f(v_1) = \sum_{n=1}^{N} \alpha_n v_1^{\,n} \,. \tag{5}$$

The same argument used above to obtain (4) implies that the output of the pipelined ADC is now

$$x_{out}[n] = x_{out}[n]\big|_{\text{ideal}} + f\left(v_1(nT_s)\right), \tag{6}$$

where $x_{out}[n]|_{\text{ideal}}$ is the ideal output of the pipelined ADC given by (4).

For example, suppose $\alpha_n = 0$ for all n except $n = 1$. This implies that the distortion is just a gain error, i.e., linear distortion. In the absence of other non-ideal circuit behavior, $v_1(nT_s) = e_{ADC\,1}[n]$ and is, therefore, bounded in magnitude by $\Delta/2$, so it follows from (6) that the maximum magnitude of the error from the non-ideal residue amplifier gain is $f(\Delta/2) = |\alpha_1|\Delta/2$. It follows from (4) that the quantization error introduced by the ideal version of the pipelined ADC has a maximum magnitude of $\Delta/8192$. Hence, a gain error of just $\alpha_1 = 1/4096$ is sufficient to cause the resulting pipelined ADC error to be comparable in magnitude to the pipelined ADC's quantization error. More generally, if $\alpha_n = (\Delta/2)^{1-n}/4096 = 2^{n-13}\Delta^{1-n}$, the nth term in (5) gives rise to an error component in the pipelined ADC output with a magnitude comparable

to the pipelined ADC's quantization error.

The 15-bit, 40-MS/s pipelined ADC integrated circuit (IC) presented in [18] provides a convenient circuit-level example of the issues described above. The ADC is based on the architecture shown in Figure 1 modified to include digital background calibration techniques that cancel ADC error arising from DAC capacitor mismatches and interstage gain errors. The ADC achieves over 90 dB of spurious-free dynamic range (SFDR) and 72 dB of peak signal-to-noise-and-distortion ratio (SNDR) over the 20 MHz bandwidth. To achieve sufficiently low distortion for this level of ADC performance, high-power residue amplifiers are used in the design: the op-amps in the residue amplifiers consume approximately 80% of the 400 mW consumed by the entire IC.

Had the sample-rate been higher than 40 MHz, even higher-performance, and, therefore, higher-power, residue amplifiers would have been required to maintain the same SFDR and peak SNDR. For example, circuit simulations indicate that the pipelined ADC's SFDR and peak SNDR drop to 65 dB and 56 dB, respectively if the sample-rate is increased to 100 MHz without improving the performance of the residue amplifiers[1]. Simulation of the residue amplifier stage indicates that this reduction in performance comes from both linear gain error associated with incomplete settling and from third order distortion; the use of differential circuitry causes the even-order terms to be negligible in this example relative to the target specifications of 90 dB

[1] The digital logic in the IC limited the clock rate to 50 MHz, so this observation had to be made via transistor-level simulation. However, circuit simulation results up to 50 MHz are consistent with measured results.

SFDR and 72 dB peak SNDR, and, although higher-order distortion terms are present, they too are negligible in this example. Later in the paper this example is revisited and an implementation of the HDC technique is described that digitally measures and cancels the error introduced by the residue amplifiers to restore the SFDR and peak SNDR to their target values of 90 dB and 72 dB, respectively.

III. Signal Processing Details of the HDC Technique

A. An *m*th-Order Distortion Correction Example

To demonstrate the basic idea underlying the HDC technique, a simplified case is considered first: the residue amplifier in the first stage introduces only *m*th-order distortion, i.e., $f(v_1) = \alpha_m v_1^m$ for some integer, m, and all other components in the pipelined ADC are ideal.

The HDC technique for this example is shown in Figure 3. A set of m uncorrelated, two-level, pseudo-random, digital *calibration sequences*, $t_1[n]$, $t_2[n]$, ..., $t_m[n]$, each of which takes on values of $\pm A$, is zero-mean, and is independent of the pipelined ADC's input signal, are added to the output of the flash ADC. They are converted to analog form along with the output sequence from the flash ADC, so the input to residue amplifier at the nth sample time is

$$v_1(nT_s) = -e_{ADC\,1}[n] - \sum_{k=1}^{m} t_k[n]\,. \tag{7}$$

The amplitude, A, of the calibration sequences is chosen such that the sum of the calibration sequences has a maximum amplitude of approximately $\Delta/4$. Since the sum of

the calibration sequences is amplified along with the quantization error from the flash ADC, this implies that approximately half of the over-range margin is taken up by the calibration sequences, which leaves the other half of the over-range margin for error associated with non-ideal circuit behavior.

The calibration sequences are subjected to the distortion function of the residue amplifier along with the quantization error from the first pipeline stage, and, by reasoning similar to that presented in the previous section to obtain (3),

$$r_1[n] = v_1(nT_s) + \alpha_m v_1^{\,m}(nT_s) + \frac{1}{4096} e_{ADC\,7}[n]. \tag{8}$$

It follows that the pipelined ADC output prior to correction by the HDC technique is

$$y_1[n] = v_{in}(nT_s) + \alpha_m v_1^{\,m}(nT_s) + \frac{1}{4096} e_{ADC\,7}[n]. \tag{9}$$

The purpose of the HDC logic is to estimate $\alpha_m v_1^{\,m}(nT_s)$ with which to cancel the mth-order distortion in $y_1[n]$, i.e., the second term in (9). It does this by correlating $r_1[n]$ against the product of the calibration sequences, $t_1[n]t_2[n]\cdots t_m[n]$. The correlation involves multiplying the digital sequence

$$s_1[n] = r_1[n] + \sum_{k=1}^{m} t_k[n] = -e_{ADC\,1}[n] + \alpha_m v_1^{\,m}(nT_s) + \frac{1}{4096} e_{ADC\,7}[n] \tag{10}$$

by $t_1[n]t_2[n]\cdots t_m[n]$, a two-level sequence that takes on values of $\pm A^m$, and averaging the result. Since the calibration sequences are zero-mean, uncorrelated with each other, and independent of the pipelined ADC's input signal, it follows that $t_1[n]t_2[n]\cdots t_m[n]$ is uncorrelated with all of the terms in (10) except the term $(m!)t_1[n]t_2[n]\cdots t_m[n]\alpha_m$ that occurs in the expansion of $v_1(nT_s)$, as given by (7), raised

to the mth power. Consequently, the average of $s_1[n]$ times $t_1[n]t_2[n]\cdots t_m[n]$ over n is $(m!)A^{2m}\alpha_m$. The HDC logic multiplies the output of the averager by $K_m = A^{-2m}/(m!)$ to obtain γ_m which is an estimate of α_m. It then multiplies γ_m by $r_1{}^m[n]$ to obtain the estimate of $\alpha_m v_1{}^m(nT_s)$.

To the extent that the calibration sequences have the above-mentioned statistical properties, γ_m converges exactly to α_m as the number of samples averaged by the HDC logic increases; the more samples in the average, the better the estimate of α_m. This convergence occurs regardless of the pipelined ADC's input signal, so the HDC technique performs background calibration, i.e., it functions during normal operation of the pipelined ADC. After an initial *convergence time* during which the averager obtains a sufficiently accurate estimate of α_m that the pipelined ADC's accuracy is limited by non-ideal circuit behavior other than mth-order residue amplifier distortion, the pipelined ADC operates at its full accuracy, and the HDC technique continues to track slow variations in α_m that may occur because of temperature changes or as the device ages.

Although the estimate of α_m has an accuracy that depends only upon the number of samples averaged by the HDC logic, the accuracy of the estimate of $\alpha_m v_1{}^m(nT_s)$ is limited by the presence of unwanted higher-order terms that occur in $r_1{}^m[n]$. For example, it follows from (8) that if $m = 3$ and the small last term of (8) is neglected, then

$$\alpha_3 r_1{}^3[n] \cong \alpha_3 v_1{}^3(nT_s) + \underbrace{3\alpha_3{}^2 v_1{}^5(nT_s) + 3\alpha_3{}^3 v_1{}^7(nT_s) + \alpha_3{}^4 v_1{}^9(nT_s)}_{\text{unwanted terms}}. \qquad (11)$$

Fortunately, as demonstrated in the next section the unwanted terms in (11) tend to be small in practice in which case they can be neglected.

For the special case of $m = 1$, the HDC technique as shown in Figure 3 reduces to the gain error correction (GEC) technique presented in [18] and [22]. Hence, the HDC technique can be viewed as an extension of the GEC technique.

B. The HDC Technique for Correction of Multiple Orders of Distortion

By a minor extension of the analysis presented above, it is easy to verify that γ_m converges to α_m even if the residue amplifier's distortion function contains lower-order distortion terms. In other words, even if any of the α_i for $i < m$ are non-negligible in (5), the HDC logic shown in Figure 3 accurately estimates α_m.

However, a complication arises if any of the α_i for $i > m$ are non-negligible. For example, suppose that the HDC technique as shown in Figure 3 is implemented with $m = 3$, but instead of the residue amplifier introducing only third-order distortion, it introduces first-order, third-order, and fifth-order distortion. That is, suppose $f(v_1) = \alpha_1 v_1 + \alpha_3 {v_1}^3 + \alpha_5 {v_1}^5$. In this case (10) becomes

$$s_1[n] = -e_{ADC\,1}[n] + \alpha_1 v_1(nT_s) + \alpha_3 {v_1}^3(nT_s) + \alpha_5 {v_1}^5(nT_s) + \frac{1}{4096} e_{ADC\,7}[n] \qquad (12)$$

with $v_1(nT_s)$ still given by (7). Expanding the fifth-order term in (12) results in several cross-terms that are correlated with the product of the calibration sequences, $t_1[n]t_2[n]t_3[n]$. These terms cause γ_3 to converge to a value that differs from α_3. Specifically, γ_3 now converges to

$$\alpha_3 + \left[30A^2 + 10\left\langle {e_{ADC1}}^2[n]\right\rangle\right]\alpha_5 \qquad (13)$$

as the number of averaged samples increases, where $\langle e_{ADC1}{}^2[n]\rangle$ denotes the average of $e_{ADC1}{}^2[n]$. Therefore, the presence of non-negligible fifth-order residue amplifier distortion prevents the version of the HDC technique shown in Figure 3 from functioning properly because of the unwanted α_5 term in (13).

As another example, consider the same distortion function, but suppose $m = 1$. In this case the HDC logic correlates a single calibration sequence, $t_1[n]$, against $r_1[n]$ to obtain γ_1. It follows from the presentation above that in the absence of third-order and fifth-order distortion, γ_1 would converge to α_1. However, the third-order and fifth-order terms in (12) contain several cross-terms that are correlated with $t_1[n]$. Consequently, γ_1 converges to

$$\alpha_1 + \left[13A^2 + 3\left\langle e_{ADC1}{}^2[n]\right\rangle\right]\alpha_3 + \left[241A^4 + 130A^2\left\langle e_{ADC1}{}^2[n]\right\rangle + 5\left\langle e_{ADC1}{}^4[n]\right\rangle\right]\alpha_5 . \quad (14)$$

From these examples, it is evident that the HDC technique must be modified for cases in which the residue amplifier's distortion function has more than one non-negligible term. The idea is to use N two-level calibration sequences as described above, correlate $r_1[n]$ against $t_1[n]t_2[n]\cdots t_k[n]$ to obtain γ_k for each $k = 1, 2, ..., N$, at which α_k is non-negligible, and estimate the unwanted terms in each γ_k value to obtain an estimate of the corresponding α_k.

For example, suppose again that $f(v_1) = \alpha_1 v_1 + \alpha_3 v_1{}^3 + \alpha_5 v_1{}^5$. In this case, 5 calibration sequences are used, each of which takes on values of $\pm A$ where $A = \Delta/20$. The corresponding HDC logic is shown in Figure 4, where $K_i = A^{-2i}/(i!)$. It calculates γ_1, γ_3, and γ_5 as described above, as well as the averages of $r_1{}^2[n]$ and $r_1{}^4[n]$ which are

denoted as η_2 and η_4, respectively. By the arguments presented above, γ_1 converges to the quantity given by (14), γ_3 converges to the quantity given by (13), γ_5 converges to α_5, and η_2 and η_4 converge to $\langle e_{ADC1}{}^2[n]\rangle$ and $\langle e_{ADC1}{}^4[n]\rangle$, respectively. Therefore, the vector $\boldsymbol{\alpha}' = \mathbf{M}(\eta_2, \eta_4)\,\boldsymbol{\gamma}$ converges to $\boldsymbol{\alpha}$ where

$$\boldsymbol{\alpha} = \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \end{bmatrix}, \quad \boldsymbol{\gamma} = \begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \gamma_3 \end{bmatrix}, \quad \text{and}$$
$$\mathbf{M}(\eta_2, \eta_4) = \begin{bmatrix} 1 & -13A^2 - 3\eta_2 & -241A^4 + 390A^6 + 90A^4\eta_2 + 30\eta_2^2 - 5\eta_4 \\ 0 & 1 & -30A^2 - 10\eta_2 \\ 0 & 0 & 1 \end{bmatrix}. \tag{15}$$

The HDC logic uses the resulting estimated values of α_1, α_3, and α_5 to cancel the corresponding distortion terms in the pipelined ADC's output sequence as shown in Figure 4.

C. Convergence Time

It follows from the presentation above that the γ_k values calculated by the HDC logic can be written as

$$\gamma_k = \frac{1}{k!\,A^k P} \sum_{i=0}^{P-1} s_1[i]\, c[i], \qquad \text{where} \qquad c[n] = \begin{cases} 1 & \text{if } t_1[n]t_2[n]\cdots t_k[n] > 0 \\ -1 & \text{otherwise} \end{cases} \tag{16}$$

and P is the number of samples averaged by the averager blocks. The sign of the product of the calibration sequences, $c[n]$, is a random sequence, so for any finite value of P, γ_k is a random variable.

If the averagers in the HDC logic were ideal, they would evaluate (16) in the limit as $P \to \infty$ in which case γ_k would converge to its ideal value, $\gamma_k|_{\text{ideal}}$. However, P

is finite in any practical averager, so the convergence process is incomplete and this introduces a random estimation error component. The mean squared value of the estimation error, i.e., $E\{(\gamma_k - \gamma_k|_{ideal})^2\}$, can be used to quantify the estimation error. By its definition, $c[n]$ is a white random sequence with zero mean and unity variance. It is independent of the pipelined ADC's input sequence, any term that does not contain one or more of the sequences $t_1[n]$, $t_2[n]$, …, $t_k[n]$ as factors, and any term that contains a calibration sequence other than $t_1[n]$, $t_2[n]$, …, $t_k[n]$ as a factor. With A set to $\Delta/(4m)$ (to provide a specific example), it follows from these properties and (16) that

$$E\left\{\left(\gamma_k - \gamma_k\big|_{ideal}\right)^2\right\} = \frac{1}{P}\left(\frac{4m}{\Delta}\right)^{2k}\left(\frac{1}{k!}\right)^2\left(\frac{1}{P}\sum_{i=0}^{P-1} u_1^2[i]\right), \tag{17}$$

where m is the number of calibration sequences, and $u_1[n]$ is equal to $s_1[n]$ minus the terms that are correlated with $c[n]$. Equation (17) specifies the relationship between the number of samples averaged and the convergence accuracy of the HDC technique. By the design of the pipelined ADC, $|u_1[n]| < \Delta$, so (17), viewed as a function of P, has the form of a bounded sequence divided by P. Hence, as expected this implies that the estimation error goes to zero as $P \to \infty$.

The required convergence time is the minimum value of P for which the HDC logic is able to measure all the α_k values with sufficient accuracy that the error arising from residue amplifier distortion is canceled to the point that the target specifications of the pipelined ADC are met. Equation (17) gives insight into which terms affect the required convergence time. However, a closed-form expression for the required convergence time is not yet known. Hence, as demonstrated in the next section, com-

puter simulations are used to determine the required convergence time on a case-by-case basis.

One insight offered by (17) is that the mean squared estimation error for a given value of P gets worse as k is increased. The number of calibration sequences, m, must at least equal the order of the highest-order distortion term to be measured by the HDC logic, so m is at least as large as k in (17), and the mean squared estimation error is proportional to m^{2k}. Thus, the highest-order distortion term to be measured generally determines the required convergence time. For example, in the HDC technique implementation presented in the next section, the third-order distortion term is the highest term measured by the HDC logic. This term causes the required convergence time to be on the order of 4 billion samples (e.g., 40 seconds worth of samples at a sample-rate of 100 MHz).

D. Overview of Practical Issues

To simplify the presentation the HDC technique has been described up to this point under the unrealistic assumption that the only non-ideal analog component in the pipelined ADC is the residue amplifier in the first pipeline stage. However, as described in the remainder of the paper, the HDC technique is able to function effectively in the presence of realistic circuit non-idealities.

It follows from the analysis presented above that the convergence process works in the presence of any signal that is statistically independent of the calibration sequences. Therefore, circuit noise does not bias the HDC convergence process. This leaves distortion (from components other than the residue amplifier) as the only po-

tential non-ideal circuit behavior that can significantly affect the convergence of the HDC technique. For example, if the DAC in a pipeline stage to which the HDC technique is applied introduces non-negligible, non-linear distortion, the HDC technique will not properly correct for the residue amplifier distortion. Fortunately, with dynamic element matching (DEM) to scramble component mismatches, the DACs in a pipelined ADC can be implemented with extremely high linearity [18, 23]. Moreover, segmentation techniques can be used to create DEM DACs that handle the extra levels required to accommodate the calibration sequences with very little extra hardware complexity or latency [18, 24, 25].

In the examples presented so far, the HDC technique has been applied only to the first pipeline stage, but in general it can be applied simultaneously to as many of the pipeline stages as necessary. As can be deduced from (3) and Figure 1, for each $k = 1, ..., 6$, the combination of pipeline stages k through 7 and the associated digital logic is a pipelined ADC in its own right with a resolution of approximately $2(6 - k) + 3$ bits. Therefore by the reasoning presented above, the HDC technique can be applied simultaneously to any of the first 6 pipeline stages provided calibration sequences are used in each stage that are independent of those used in the other stages. It follows from the architecture of Figure 1 that any distortion introduced by the kth pipeline stage is attenuated by a factor 4^{k-1} referred to the output, so the distortion introduced by all but the first few stages tends to be negligible. Consequently, in practice it is only necessary to apply the HDC technique to the first few stages.

IV. HDC Implementation Example

An example is presented in this section in which the HDC technique is applied to the first three stages of the pipelined ADC shown in Figure 1. The result is shown in Figure 5, Figure 6, and Figure 7: Figure 5 shows a high-level view of the pipelined ADC, Figure 6 shows the pipelined ADC with expanded views of the first pipeline stage and the associated HDC logic, and Figure 7 shows the high-level structure of the DEM DAC used in the first three pipeline stages. The details are described below and computer simulation results are presented to demonstrate the performance of the system.

The residue amplifier distortion for this example is modeled after the behavior observed via transistor-level circuit simulations in the pipelined ADC of [18] for a sample-rate of 100 MHz. Specifically, for each residue amplifier, the non-negligible distortion terms in (5) are $\alpha_1 = -0.0125$, $\alpha_3 = -2^{-6}\,\Delta^{-2}$, $\alpha_5 = -2^{-9}\,\Delta^{-4}$, and $\alpha_7 = -2^{-11}\,\Delta^{-6}$, where $\Delta = 250$ mV is the step-size of the flash ADC. It can be deduced for this case from the results presented in Section II that only the first-order and third-order residue amplifier distortion terms in the first three pipeline stages need be cancelled to achieve 15-bit pipelined ADC accuracy. Therefore, the HDC technique is applied in this example to measure and cancel just these distortion terms.

The details of the first pipeline stage and associated HDC logic are shown in Figure 6. Three pseudo-random $\pm\Delta/16$ calibration sequences are added prior to the DAC, so the sum of the calibration sequences is a four-level sequence that can range from $-3\Delta/16$ to $3\Delta/16$ in steps of $\Delta/8$. The use of three $\pm\Delta/16$ calibration sequences

has two analog circuit implications. The first implication is that the DAC must have a minimum step-size of Δ/8 (instead of Δ as in the fourth through seventh pipeline stages) and enough levels to accommodate the calibration sequences. To avoid exceeding the input range of the DAC, the sum of the calibration sequences are forced to Δ/16 and the HDC estimators for the pipeline stage are disabled when the output of the flash ADC is either at its maximum or minimum value. Therefore, the sum of the calibration sequences and the flash ADC output can take on values of

$$k\Delta + i\Delta/8 + \Delta/16, \quad \text{where } k = -4,-3,...,4 \text{ and } i = \begin{cases} -2,-1,0,1 & \text{if } |k| \le 3 \\ 0 & \text{if } |\mathrm{k}| = 4 \end{cases}, \qquad (18)$$

so the DAC must be able to generate these output levels. The second implication is that the calibration sequences occupy almost half of what would otherwise have been the over-range margin. Specifically, it follows from the discussion in Section II that the over-range margin for each of the first three stages is ±1.75Δ. While this tightens the design constraints on the flash ADC, it is not difficult to handle in practice [18].

As described above, the DAC in each of the first three stages must be capable of generating the output levels specified by (18). This is accomplished by the DAC architecture shown in Figure 7. It consists of a digital DEM encoder block and 15 1-bit DACs. Each one-bit DAC outputs a nominal value of $-q\Delta$ or $q\Delta$ depending upon whether its input bit is 0 or 1, respectively, where q is a *weighting factor* that is fixed for a given 1-bit DAC. There are three 1-bit DACs with q = 1/16, two with q = 1/8, two with q = 1/4, and eight with q = 1/2. With this 1-bit DAC weighting arrangement, for most of the possible input values, $x_{in}[n]$, there are multiple distinct bit vec-

tors, $x_1[n]$, $x_2[n]$, …, $x_{14}[n]$, that give rise to the desired nominal output value. At each sample clock, the DEM encoder pseudo-randomly selects one of these multiple, nominally equivalent vectors.

If all the 1-bit DAC step-sizes were ideal, the pseudo-random selection algorithm in the DEM encoder would have no effect. However, inadvertent component mismatches arise during circuit fabrication which causes the 1-bit step-sizes to deviate from their ideal values. If only one of the possible values of the 1-bit DAC input vector, $x_1[n]$, $x_2[n]$, …, $x_{14}[n]$, were used for each value of $x_{in}[n]$, the step-size errors would cause the overall DAC to introduce harmonic distortion. By pseudo-randomly choosing among the different possible 1-bit DAC input vectors for each input sample, the DEM encoder causes the overall DAC to introduce white noise that is uncorrelated with the other sequences in the pipelined ADC instead of harmonic distortion, and the white noise can be removed in the digital domain by a background calibration technique [18, 23].

From a signal processing point of view the DEM encoder can be viewed as a tree of digital logic blocks called *switching blocks* as shown in Figure 7. Each switching block is labeled $S_{k,r}$ or $S_{k,r}^{seg}$ in the figure, where k and r denote the position of the switching block in the tree. The three switching blocks labeled $S_{k,r}^{seg}$ in the figure are called *segmented* switching blocks because in each case their two outputs affect the input bits to 1-bit DACs with different weighting factors. The ten switching blocks labeled $S_{k,r}$ are called *non-segmented* switching blocks because in each case their two outputs only affect the input bits to 1-bit DACs with equal weighting fac-

tors.

Each switching block operates on a digital input sequence and generates two digital output sequences. The output sequences generated by each segmented switching block, $S_{k,r}^{seg}$, are given by

$$x_{k-1,1}[n] = \frac{x_{k,r}[n] + s_{k,r}[n]}{2}, \qquad \text{and} \qquad x_{1,k+1}[n] = -s_{k,r}[n], \tag{19}$$

where $x_{k,r}[n]$ is the input to the switching block and $s_{k,r}[n]$ is a pseudo-random sequence, called a *switching sequence*. The switching sequence is generated as part of the switching block logic as

$$s_{k,r}[n] = \begin{cases} 0 & \text{if } x_{k,r}[n] \text{ is even} \\ \pm 1 & \text{otherwise (chosen pseudo-randomly)} \end{cases} \tag{20}$$

The output sequences generated by each non-segmented switching block, $S_{k,r}$, are given by

$$x_{k-1,2r}[n] = \frac{x_{k,r}[n] + s_{k,r}[n]}{2}, \qquad \text{and} \qquad x_{k-1,2r-1}[n] = \frac{x_{k,r}[n] - s_{k,r}[n]}{2}, \tag{21}$$

where, as before, $x_{k,r}[n]$ is the input to the switching block and $s_{k,r}[n]$ is a switching sequence given by (20). It can be verified from the results presented in [24, 26], and [27] that the DEM encoder ensures that output level errors in the 1-bit DACs from component mismatches do not cause the overall DAC to introduce harmonic distortion, which is a requirement of the HDC technique.

It follows from (19), (20), and (21) that the data paths through the switching blocks are not clocked, so the DEM encoder could be implemented directly as combinational logic. However, in high-speed pipelined ADCs, latency from the output of

the flash ADC through the DAC in each pipeline stage must be minimized because the larger the latency the less time is available for the residue amplifier following the DAC to settle. In [18] this issue was addressed by implementing the functionality of both the calibration sequence adder and the DEM encoder in parallel as a single layer of digital transmission gates along with some digital logic gates through which latency is not critical. This reduced the latency from the output of the flash ADC through the DEM encoder to that of a single transmission gate. Although the DEM encoder shown in Figure 7 is more complicated than that presented in [18], the same approach has been taken in the computer simulated implementation described below. Since the calibration sequences are known in advance of the flash ADC output data, only the combinational logic component through which latency is not critical is increased in this example relative to the DEM encoder presented in [18].

The practical version of the HDC logic shown in Figure 6 is a direct implementation of ideal version shown in Figure 4, except without fifth-order distortion correction. The primary differences between the practical and ideal versions are that requantization is used to reduce the bit widths of various data buses to reduce digital complexity, and the three averagers are implemented with $P = 2^{32}$ in the practical version. Dithered requantizers are used to perform the requantization as described in [23] to avoid introducing harmonic distortion. Requantization is not necessary, but by reducing data bus widths it greatly reduces the area and power consumption of the HDC logic, yet the quantization noise it introduces adds only slightly to the HDC convergence time. The random dither sequences and calibration sequences in this

example were generated by a single linear feedback shift register of the form described in [28].

At a sample-rate of 100 MHz with $P = 2^{32}$, each HDC block requires approximately 43 seconds to converge. However, the accuracy of each HDC block depends on the accuracies of the HDC blocks in the subsequent stages. Thus, the total convergence time for this example implementation is approximately 2 minutes.

V. Simulation Results and HDC Limitations

The example pipelined ADC with HDC as described above was simulated with various non-ideal circuit effects. The simulated residue amplifier distortion in each stage includes the first through seventh-order distortion terms described above. The 1-bit DAC mismatches were chosen as independent Gaussian random variables; the standard deviations of the 1-bit DACs with step-sizes of Δ, $\Delta/2$, $\Delta/4$ and $\Delta/8$ are 0.30%, 0.42%, 0.60%, 0.85%, of $\Delta = 250$ mV, respectively. The flash ADC threshold errors and residue amplifier offset voltages were chosen as independent Gaussian random variables with standard deviations of 25 mV and 5 mV, respectively. A 10 nV_{rms}^2 white noise signal was added at the input of each residue amplifier to model thermal noise.

Figure 8 (a) shows the power spectral density (PSD) plot[2] of the output of the residue amplifier simulated alone with a 275mV, 6.4 MHz sinusoidal input signal. The amplitude of the input signal is nearly the maximum input that does not overload the next stage of the pipeline. Hence, the output of the residue amplifier consists of

[2] The PSDs were estimated using 16 Hanning windowed periodograms of length 16384.

the 6.4 MHz fundamental tone plus the residue amplifier distortion terms and thermal noise. The plot demonstrates the non-linear behavior of the residue amplifier.

Figure 8 (b) and (c) show PSD plots of the pipelined ADC with a −1dB relative to full scale 6.4 MHz sinusoidal input signal. Figure 8 (b) shows the case with the HDC technique *disabled*, and Figure 8 (c) shows the case with the HDC technique *enabled*. Comparison of Figure 8 (b) and (c) indicate that the HDC technique improved the simulated SNDR and SFDR by 26 dB and 30 dB, respectively. Numerous other simulations performed by the authors with different input signals, and different random mismatches, ADC thresholds, and DAC mismatches, exhibit similar results.

Before computing the PSD estimates for the simulation results shown in Figure 8 (b) and (c), the components of the final output signal corresponding to DAC mismatches and thermal noise were removed so as not to obscure the effect of the HDC technique. Removal of the components corresponding to DAC mismatches can be achieved in a practical implementation via the DNC technique presented in [18] and [23]. However, the DNC technique is not necessary for the HDC technique to function provided dynamic element matching DACs are used to ensure that error introduced by DAC mismatches does not contain significant harmonic distortion.

One potential limitation of the HDC technique is not demonstrated by the implementation example described above. The version of the correction scheme presented in Figure 4 and Figure 6 is not accurate if the error is too big, i.e. if the α_n coefficients in (5) are too large – this could happen, for instance, if an open loop residue amplifier configuration as in [19] is used instead of a classical closed loop configura-

tion. For example, if the distortion function given by (5) can be written as

$$f\left(v_1(nT_s)\right) = \alpha_1 v_1(nT_s) + \alpha_3 v_1^3(nT_s) \tag{22}$$

and the digitized residue, $r_1[n]$, in Figure 6 is given by

$$r_1[n] \approx (1+\alpha_1)v_1(nT_s) + \alpha_3 v_1^3(nT_s)\,, \tag{23}$$

the correction signal $d_1[n]$ is

$$\begin{aligned} d_1[n] \approx{}& \left(\alpha_1' + \alpha_1'^2\right)v_1(nT_s) + \left(\alpha_3' + 4\alpha_1'\alpha_3' + 3\alpha_1'^2\alpha_3' + \alpha_1'^3\alpha_3'\right)v_1^3(nT_s) \\ &+ \left(3\alpha_3'^2 + 6\alpha_1'\alpha_3'^2 + 3\alpha_1'^2\alpha_3'^2\right)v_1^5(nT_s) + \ldots \end{aligned} \tag{24}$$

Subtracting (24) from the uncorrected output given by (6), using (22), and assuming that $\alpha_n \approx \alpha'_n$, the pipeline output is

$$\begin{aligned} x_{out}[n] \approx{}& x_{out}[n]\big|_{ideal} - \alpha_1^2 v_1(nT_s) - \left(4\alpha_1\alpha_3 + 3\alpha_1^2\alpha_3 + \alpha_1^2\alpha_3\right)v_1^3(nT_s) - \\ &- \left(3\alpha_3^2 + 6\alpha_1\alpha_3^2 + 3\alpha_1^2\alpha_3^2\right)v_1^5(nT_s) - \left(3\alpha_3^3 + 3\alpha_1\alpha_3^3\right)v_1^7(nT_s) - \alpha_3^4 v_1^9(nT_s) \end{aligned} \tag{25}$$

Comparing (25) to (6), it is clear that HDC removes most of the distortion provided the α_n coefficients are sufficiently small. However, in some applications this may not be the case, in which case the remaining unwanted terms in (25) may not be negligible for the given application. In such cases, the modified correction technique shown in Figure 9 can be used. A similar analysis to that presented above indicates that the pipelined ADC output is now

$$x_{out}[n] \approx x_{out}[n]\big|_{ideal} - \frac{3\alpha_3^2}{\left(1+\alpha_1\right)^2}v_1^5(nT_s) - \frac{3\alpha_3^3}{\left(1+\alpha_1\right)^3}v_1^7(nT_s) - \frac{\alpha_3^4}{\left(1+\alpha_1\right)^4}v_1^9(nT_s) \tag{26}$$

Equation (26) shows that linear and third-order distortion has been removed, while the remaining unwanted terms are smaller than or comparable to the respective terms in (25). The price paid for the accuracy improvement is increased complexity. Al-

though both schemes require the same number of multipliers, the extension of the latter scheme to correct for higher-order harmonics would result in a more complex hardware.

Another limitation of the HDC technique has not been highlighted by the example presented in section IV. Had it been necessary to apply HDC to correct fifth order residue amplifier distortion, a problem would have arisen for the chosen pipelined ADC architecture and target specifications. Specifically, the fifth-order distortion term for this case is so small that high-order distortion from the coarse quantization performed by the flash ADCs in each stage becomes significant and distorts the HDC technique's estimate. Equation (12) represents the signal used to estimate the first stage's residue amplifier distortion terms under the assumption that the following stages are either ideal or perfectly corrected. A more accurate expression for $s_1[n]$ is

$$s_1[n] = -e_{ADC1}[n] + \alpha_1 v_1(nT_s) + \alpha_3 v_1^3(nT_s) + \alpha_5 v_1^5(nT_s) + \sum_{k=2}^{7} \lambda_k e_{ADCk}[n] + \ldots, \qquad (27)$$

where the λ_k is the amplitude of uncanceled flash ADC error from the kth stage. Therefore in the absence of perfect cancellation, every flash ADC contributes error in (27). The error is largely quantization noise which tends to be highly correlated with $v_1[n]$ and therefore with the pseudorandom sequences. The smaller the α_n coefficients to be estimated by the HDC technique, the more significantly the imperfectly cancelled flash ADC errors distort the estimated coefficient values. Furthermore, the coarse quantization performed by the flash ADCs is a *hard non-linearity*, so it can not be represented by a small number of Taylor series terms. In conclusion, the HDC

technique, as well as any other scheme (e.g., [20]) that assumes the non-linearity to be estimated is well-modeled by a small number of Taylor series terms fails to work well when the non-linearity to be estimated is very small. In principle, an analog dither signal can be added prior to the flash ADCs to eliminate this problem in cases where very small distortion terms must be measured by the HDC technique [29].

ACKNOWLEDGEMENTS

Chapter 1, in full, has been published in the IEEE Transactions on Circuits and Systems I: Regular Papers, volume 53, number 9, pages 1885-1895, September 2006. The dissertation author is the primary investigator and author of this paper. Professor Ian Galton supervised the research which forms the basis for this paper.

REFERENCES

1. W.T. Colleran, A.A. Abidi, "A 10-b, 75-MHz two-stage pipelined bipolar A/D converter," *IEEE Journal of Solid-State Circuits*, vol. 28, no.12, pp. 1187-1199, December 1993.

2. K. Sone, Y. Nishida, N. Nakadai, "A 10-b 100-Msample/s pipelined subranging BiCMOS ADC," *IEEE Journal of Solid-State Circuits*, vol. 28, no.12, pp. 1180-1186, December 1993.

3. Jipeng Li, Un-Ku Moon, "A 1.8-V 67-mW 10-bit 100-MS/s pipelined ADC using time-shifted CDS technique," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 9, pp. 1468-1476, September 2004.

4. Jong-Bum Park, Sang-Min Yo, Se-Won Kim, Young-Jae Cho, Seung-Hoon Lee, "A 10-b 150-MSample/s 1.8-V 123-mW CMOS A/D converter with 400-MHz input bandwidth," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 8, pp. 1335-1337, August 2004.

5. L. Singer, S. Ho, M. Timko, D.Kelly, "A 12-b 65-Msample/s CMOS ADC with 82-dB SFDR at 120 MHz," in *ISSCC Dig. Tech. Papers*, pp. 38–39, February 2000.

6. R. Jewett, K. Poulton, K.-C. Hsieh, and J. Doernberg, "A 12 b 128 MSample/s ADC with 0.05LSB DNL," in *ISSCC Dig. Tech. Papers*, pp. 138–139, February 1997.

7. C.R. Grace, P.J. Hurst, S.H. Lewis, "A 12-bit 80-MSample/s pipelined ADC with bootstrapped digital calibration," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 5, pp.1038-1046, May 2005.

8. Myung-Jun Choe, Bang-Sup Song, K. Bacrania, "A 13-b 40-MSamples/s CMOS pipelined folding ADC with background offset trimming," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 12, pp.1781-1790, December 2000.

9. P. C. Yu, S. Shehata, A. Joharapurkar, P. Chugh, A. Bugeja, X. Du, S. U. Kwak, Y. Papantonopoulous, and T. Kuyel, "A 14 b 40 MSample/s pipelined ADC with DFCA," in *ISSCC Dig. Tech. Papers*, pp. 136–137, February 2001.

10. K. Nair and R. Harjani, "A 96 dB SFDR 50 MS/s digitally enhanced CMOS pipeline A/D converter," in *ISSCC Dig. Tech. Papers*, pp. 456–457, February 2004.

11. A. Zanchi, F. Tsay, "A 16-bit 65-MS/s 3.3-V pipeline ADC core in SiGe BiCMOS with 78-dB SNR and 180-fs jitter," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 6, pp. 1225-1237, June 2005.

12. Hung-Chih Liu, Zwei-Mei Lee, Jieh-Tsorng Wu, "A 15-b 40-MS/s CMOS pipelined analog-to-digital converter with digital background calibration," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 5, pp. 1047-1056, May 2005.

13. Seung-Tak Ryu, S. Ray, Bang-Sup Song, Gyu-Hyeong Cho, K. Bacrania, "A 14-b linear capacitor self-trimming pipelined ADC," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 11, pp. 2046-2051, November 2004.

14. Y. Chiu, P. Gray, B. Nikolic, "A 1.8 V 14 b 10 MS/s pipelined ADC in 0.18 μm CMOS with 99 dB SFDR," in *ISSCC Dig. Tech. Papers*, pp. 458–459, February 2004.

15. A. M. Abo and P. R. Gray, "A 1.5-V, 10-bit, 14.3-MS/s CMOS pipeline analog-to-digital converter," *IEEE Journal of Solid-State Circuits*, vol. 34, no. 5, pp.599–606, May 1999.

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

17. Yun Chiu, C. W. Tsang, B. Nikolic, P. R. Gray, "Least mean Square Adaptive Digital Background Calibration of Pipelined Analog-to-Digital Converters," *IEEE Transactions on Circuits and Systems I*, vol. 51, no. 1, pp. 38-46, January 2004.

18. E. Siragusa, I. Galton, "A Digitally Enhanced 1.8V 15b 40MS/s CMOS Pipelined ADC," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 12, pp. 2126-2138, December 2004.

19. B. Murmann, B. Boser, "A 12b 75MS/s Pipelined ADC using Open-Loop Residue Amplification," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 12, pp. 2040-2050, December 2003.

20. J. P. Keane, P. J. Hurst, S. H. Lewis, "Background Interstage Gain Calibration Technique for Pipelined ADCs," *IEEE Transactions on Circuits and Systems I*, vol. 52, no. 1, pp. 32-43, January 2005.

21. S. H. Lewis and P. R. Gray, "A pipelined 5-Msample/s 9-bit analog-to-digital converter," *IEEE Journal of Solid State Circuits*, vol. SC-22, pp. 954-961, December 1987.

22. E. J. Siragusa, I Galton, "Gain error correction technique for pipelined analogue-to-digital converters," *IEE Electronics Letters*, vol. 36, no.7, p.617-618, March 30, 2000.

23. I. Galton, "Digital Cancellation of D/A Converter Noise in Pipelined A/D Converters," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 47 no. 3, pp. 185-196, March 2000.

24. A. Fishov, E. Siragusa, J. Welz, E. Fogleman, I. Galton, "Segmented mismatch-shaping D/A conversion," in *Proc. of the IEEE International Symposium on Circuits and Systems*, May 2002.

25. S. Pamarti, L. Jansson, I. Galton, "A wideband 2.4-GHz delta-sigma fractional-N PLL with 1-Mb/s in-loop modulation," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 1, pp. 49 – 62, January 2004.

26. I. Galton, "Spectral shaping of circuit errors in digital-to-analog converters," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 44, no. 10, pp. 808-817, Nov., 1997.

27. J. Welz, I. Galton, "Necessary and sufficient conditions for mismatch shaping in a general class of multibit DACs," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 49, no. 12, pp. 748-759, December 2002.

28. E. Fogleman, I. Galton, W. Huff, H. Jensen, "A 3.3-V Single-Poly CMOS Audio ADC Delta–Sigma Modulator with 98-dB Peak SINAD and 105-dB Peak SFDR," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 3, pp. 297 – 307, March 2000.

29 A. B. Sripad, D. L. Snyder, "A necessary and sufficient condition for quantization errors to be uniform and white," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-25, pp. 442-448, Oct. 1977.

Chapter 2

A 130mW 100MS/s Pipelined ADC with 69dB SNDR Enabled by Digital Harmonic Distortion Correction

Abstract—This paper presents a pipelined ADC with two fully-integrated digital background calibration techniques: harmonic distortion correction (HDC) to compensate for residue amplifier gain error and nonlinearity and DAC noise cancellation (DNC) to compensate for DAC capacitor mismatches. It is the first IC implementation of HDC, and the results demonstrate that HDC and DNC together facilitate low-voltage operation and enable reductions in power dissipation relative to comparable conventional state-of-the-art pipelined ADCs. The pipelined ADC achieves a peak SNR of 70dB and a −1dBFS SFDR of 85dB at a sample-rate of 100MHz. It is implemented in a 90nm CMOS process and consumes 130mW from 1.2V and 1.0V analog and digital power supplies, respectively.

I. Introduction

Pipelined ADCs are advantageous and widely used in applications with signal bandwidths that are too high for oversampling delta-sigma ADCs to be efficient and resolution requirements that are too high for flash ADCs to be efficient. Nevertheless

This work was supported by the National Science Foundation under Award 0515286, by the corporate members of the UCSD Center for Wireless Communications, by the University of California Discovery Program, and by STMicroelectronics.

they are sensitive to distortion introduced by the residue amplifiers in their first few stages, and residue amplifier distortion tends to be inversely related to both power supply voltage and power consumption. Therefore, the residue amplifiers are usually the dominant consumers of power in high-resolution pipelined ADCs, particularly in low supply voltage designs [30- 35].

Recently, digital calibration techniques that measure and cancel pipelined ADC error arising from distortion introduced by the residue amplifiers have been proposed [36- 39]. By relaxing the residue amplifier distortion requirement for a given level of ADC accuracy, they offer the potential to significantly reduce power consumption and supply voltage in high-resolution pipelined ADCs. One such technique, harmonic distortion correction (HDC), is applied to the pipelined ADC described in this paper. It enables the ADC to achieve a peak signal to noise and distortion ratio (SNDR) of 70 dB over its 50 MHz Nyquist band despite the use of residue amplifiers with a DC loop gain of only 23 dB and a unity gain bandwidth of only 200MHz.

The pipelined ADC also incorporates a recently proposed digital calibration technique called DAC noise cancellation (DNC) to compensate for error introduced by DAC capacitor mismatches [40]. Together, HDC and DNC enable the pipelined ADC to achieve state of the art power consumption relative to comparable published ADCs despite a low analog power supply voltage of 1.2V.

The work presented in this paper is the first IC implementation of HDC, so the focus of the paper is to describe the practical implementation issues associated with

HDC. The paper consists of four main sections. Section II provides a brief description of the conventional portion of the implemented pipelined ADC, a model for the residue amplifier distortion, and a brief description of the theory underlying HDC. Section III describes the system-level implementation details and issues associated with HDC, Section IV describes the analog circuit implementation details of the pipelined ADC, and Section V presents measurement results.

II. Background Information

A. The Underlying Pipelined ADC Architecture

A conventional six-stage 14-b pipelined ADC architecture is shown in Figure 10. Each stage except the last consists of a 9-level flash ADC, a 9-level DAC, and a *residue amplifier* with an ideal gain of 4. The last stage is just a 17-level flash ADC. This structure is well known in literature and it is described below using the same notation as in [37].

Each 9-level flash ADC ideally behaves as a uniform quantizer with quantization step of size Δ and input range extending from -4.5Δ to 4.5Δ. In this design, the nominal value of Δ is 170mV. The output of the kth flash ADC at the nth sample time is given by

$$x_k[n] = v_{in,k}(nT_S) + e_{ADC,k}[n], \tag{28}$$

where $v_{in,k}(nT_S)$ is the flash ADC input signal, T_S is the sample interval, and $e_{ADC,k}[n]$ is the error introduced by the flash ADC, i.e., the output minus the input of the flash

ADC with the least significant bit (LSB) of the flash ADC output taken to have a weight of Δ. It is customary to refer to $e_{ADC,k}[n]$ as quantization error although in practice it contains both error arising from quantization as well as error arising from non-ideal circuit behavior such as comparator offset voltages. In the absence of non-ideal circuit behavior it is bounded in magnitude by Δ/2.

The 9-level DAC converts the output of the flash ADC into analog format. The difference between the pipelined ADC input sample and the output of the DAC, called the *residue*, is amplified by the residue amplifier and the result is fed to the next pipeline stage.

It follows from (28) that in the absence of non-ideal circuit behavior, the input to and output of the *k*th residue amplifier at the *n*th sample time are

$$v_k(nT_S) = -e_{ADC,k}[n], \quad \text{and} \quad v_{in,k+1}(nT_S) = 4v_k(nT_S), \tag{29}$$

respectively. In this case the output of the *k*th residue amplifier is bounded between −2Δ and 2Δ, whereas the input range of the subsequent stage to which it is applied extends from −4.5Δ to 4.5Δ. The extra input range is called *over range margin*. Its purpose is to accommodate error from non-ideal circuit behavior such as flash ADC threshold deviations.

As indicated in Figure 11, the output of the divide-by-four block in the *k*th pipeline stage, referred to as the stage's *digitized residue*, can be written as $r_k[n] = (r_{k+1}[n] + x_{k+1}[n])/4$. Therefore, it follows from (28) and (29) that the output of the *k*th stage can be written as:

$$x_{out,k}[n] = x_k[n] + r_k[n] = v_{in,k}(nT_S) + \frac{e_{ADC,k+1}[n]}{4} + r_{k+1}[n] \ , \tag{30}$$

where k = 1, 2, …, 5, and $r_6[n] = 0$. Recursive application of (28) and (29) in (30) gives

$$x_{out,k}[n] = v_{in,k}(nT_S) + \frac{1}{4^{6-k}} e_{ADC,6}[n] \, . \tag{31}$$

This implies that for each k, stages k through 6 together behave as a $(16-2k)$-b ADC. For example, stages 2 through 6 together behave as a 12-b ADC.

B. The Residue Amplifier Distortion Problem

The residue amplifier is usually implemented as an op-amp in a switched capacitor feedback loop [30-35, 39, 41, 46, 49, 52, 53]. When op-amp hard nonlinearities caused by effects like slew rate limiting or clipping are negligible, the residue amplifier tends to be well modeled as a memoryless, weakly nonlinear function of the amplifier's input voltage, as shown in Figure 12, with

$$f(v) = \sum_{i=1}^{N} \alpha_i v^i \ , \tag{32}$$

where α_1 is a linear gain error coefficient and α_i for $i > 1$ are nonlinear distortion coefficients.

Figure 13 shows a simplified representation of the pipelined ADC that includes the effect of op-amp nonlinearity in the first stage. Applying the same reasoning used above to obtain (31), it follows that

$$\begin{aligned} x_{out,1}[n] &= x_{out,1}[n]\big|_{ideal} + f\left(v_1(nT_S)\right) \\ &= v_{in,1}(nT_S) + \frac{1}{4^5} e_{ADC,6}[n] + f\left(v_1(nT_S)\right), \end{aligned} \tag{33}$$

where in the last expression $x_{out,1}[n]|_{ideal}$ has been replaced by (31) with $k = 1$. Therefore, the $f(v_1(nT_S))$ term appears in the output of the pipelined ADC. Furthermore, it follows from (29) that the $f(v_1(nT_S))$ term is a function of the quantization noise from the flash ADC in Stage 1 which is a nonlinear function of the input signal.

The conventional way to address this problem is to rely on feedback to suppress the residue amplifier distortion: the higher the gain and bandwidth of the op-amp, the better the suppression. Therefore in a conventional pipelined ADC design, the magnitudes of the α_i coefficients in f are reduced to the point that f has negligible effect on the digital output signal. Unfortunately, this is usually done at the expense of increased power consumption and circuit area.

C. HDC Overview

The alternative approach taken in this design is to use op-amps with larger magnitudes of the α_i coefficients in f, in return for lower op-amp power and area consumption, and then digitally estimate and cancel the resulting nonlinear distortion in the pipelined ADC output via HDC.

As described in Section IV, the op-amp used in this design has an open loop DC gain of 43dB which translates into a residue amplifier DC loop gain of 23dB. Transistor level simulations under typical conditions indicate that the residue amplifier is well-modeled as shown in Figure 12 with $\alpha_1 = -0.06$, $\alpha_2 = 0$, $\alpha_3 = -0.3\text{V}^{-2}$, $\alpha_5 =$

$-2V^{-4}$ and $\alpha_i = 0$ for $i > 5$. Without HDC, the resulting −1dBFS signal-to-noise and distortion ratio (SNDR) for the pipelined ADC would be 43dB, which is 26dB below the target specification.

Extensive circuit simulations run during the design phase of the pipelined ADC IC further indicate that only the distortion introduced by α_1 and α_3 must be corrected to achieve the 70dB SNDR target specification. The magnitudes of α_i for even values of i are negligible because of the differential circuitry used throughout the ADC, and magnitudes of α_i for odd values of $i \geq 5$ are negligible because hard nonlinearities such as slew rate limiting in the op-amps have been avoided. Therefore, HDC is configured in this work to compensate only for residue amplifier distortion associated with α_1 and α_3.

A full description of the theory behind HDC is presented in [37]. The purpose of this section is to provide a brief overview of HDC with enough information to support the subsequent description of its application to the pipelined ADC prototype.

HDC can be applied to each stage of a pipelined ADC to compensate for the distortion introduced by that stage's residue amplifier. In each stage it consists of an *estimation* portion and a *correction* portion. The former estimates the α_i coefficients in (32) for that stage's residue amplifier, and the latter uses the estimates to compensate for the distortion. In the following, the correction portion is described prior to the estimation portion, both in the context of HDC applied to the first pipeline stage.

Correction Portion of HDC

Neglecting the quantization error $e_{ADC,6}[n]$ and assuming ideal behavior of the stages 2-6, $r_1[n]$ can be expressed as

$$r_1[n] \simeq v_1(nT_S) + \underbrace{\alpha_1 \cdot v_1(nT_S) + \alpha_3 \cdot v_1^3(nT_S)}_{f(v_1(nT_S))}. \tag{34}$$

Figure 14 shows the proposed correction method, which implements:

$$\begin{aligned} r_1[n]\big|_{corrected} &= \frac{r_1[n]}{1+\alpha_1} - \frac{\alpha_3}{(1+\alpha_1)^4}(r_1[n])^3 \\ &\simeq v_1(nT_S) + \text{terms with fifth and higher order} \end{aligned} \tag{35}$$

As detailed in [37], the correction process introduces fifth and higher-order distortion terms in $r_1[n]|_{corrected}$, but they can be neglected because the power of the error they introduce in the pipelined ADC output is much lower than the target noise floor.

Estimation Portion of HDC

A digital calibration sequence, $c[n]$, is added to the output of the flash ADC as shown in Figure 15 to enable estimation of the α_1 and α_3 coefficients associated with the residue amplifier in the first pipeline stage. As indicated in the figure, $c[n]$ is converted to analog form by the DAC, so it is subtracted from the input of the residue amplifier. This causes several extra terms related to $c[n]$ to appear as components in the digitized residue. Two of the extra terms are proportional to $\alpha_1 c[n]$ and $\alpha_3(c[n])^3$, and the HDC estimation algorithm uses these terms to estimate the α_1 and α_3 coefficients.

HDC is a background calibration technique so it must estimate the α_1 and α_3 coefficients during normal operation of the ADC. Therefore, $c[n]$ must be such that the terms proportional to $\alpha_1 c[n]$ and $\alpha_3(c[n])^3$ can be measured in the digital residue even in the presence of other, potentially much larger and unknown terms related to the pipelined ADC input signal. Furthermore, it must have a relatively small magnitude so it only occupies a portion of the over range margin of the subsequent pipeline stage.

The simplest known calibration sequence with these properties is a four-level sequence of the form $c[n] = t_1[n]+t_2[n]+t_3[n]$, where the three $t_i[n]$ sequences are 2-level, independent, zero-mean pseudo-random sequences that take on values of $\pm A$ (in this design $A = \Delta/16$). For example, with this calibration sequence the $\alpha_3(c[n])^3$ term in the digitized residue contains the term $6\alpha_3 t_1[n]t_2[n]t_3[n]$. Since $t_1[n]t_2[n]t_3[n]$ is a known, 2-level, zero-mean pseudorandom sequence that takes on values of $\pm A^3$ and is uncorrelated with all the other signal components in the digitized residue, it follows that the average of the product of the digitized residue and $t_1[n]t_2[n]t_3[n]$ converges to $6A^6\alpha_3$ regardless of the input signal to the pipelined ADC.

The HDC algorithm calculates the following correlations

$$\gamma_1 = -\frac{1}{A^2 P}\sum_{n=0}^{P-1} s_1[n]t_1[n], \quad \gamma_3 = -\frac{1}{6A^6 P}\sum_{n=0}^{P-1} s_1[n]t_1[n]t_2[n]t_3[n], \quad \text{and} \quad \eta_2 = \frac{1}{P}\sum_{n=0}^{P-1} s_1^2[n], \tag{36}$$

where $s_1[n] = r_1[n] + c[n]$ and P is the number of samples averaged (e.g., $P = 2^{32}$ was used for most of the measurement results presented in Section V). It can be verified that, provided the residue amplifier is the only significant source of nonlinearity in the

system, these correlations converge to

$$\gamma_1 = \alpha_1 + \left(7A^2 + 3\left\langle e_{ADC,1}^2[n]\right\rangle\right)\alpha_3, \quad \gamma_3 = \alpha_3, \quad \text{and} \quad \eta_2 \simeq \left\langle e_{ADC,1}^2[n]\right\rangle, \tag{37}$$

in the limit as $P \to \infty$ regardless of the input to the pipelined ADC, where $\langle\cdot\rangle$ indicates the infinite time average operation. The HDC algorithm uses these correlation values to calculate the coefficients required by (35) as follows:

$$\begin{aligned} \frac{1}{1+\alpha_1} &= \frac{1}{1+\gamma_1 - \left(7A^2 + 3\eta_2\right)\gamma_3} \\ \frac{\alpha_3}{1+\alpha_1} &= \frac{\gamma_3}{1+\gamma_1 - \left(7A^2 + 3\eta_2\right)\gamma_3} \end{aligned} \tag{38}$$

It follows from (37) that γ_3 is an unbiased estimate of α_3, whereas γ_1 is an estimate of α_1 that is biased by α_3. Therefore, accurate estimation of α_1 requires knowledge of α_3. Unlike HDC, the gain error correction (GEC) technique presented in [41] calculates the equivalent of γ_1 and uses it as an estimate of α_1 directly, so it implicitly assumes that α_3 is negligible. Consequently, highly linear residue amplification is a prerequisite for GEC to function properly.

Simulation results under the typical conditions described in Section II.B indicate that if the HDC correction is performed as indicated in Figure 14 except with α_3 set to zero, then the SNDR and SFDR decrease by 1.7dB and 3dB, respectively. However, if HDC correction is performed as indicated in Figure 14 except with γ_3 set to zero in (38), which is equivalent to the correction performed by GEC, then the SNDR and SFDR drop by about 7dB and 13dB, respectively.

III. Digital Calibration System-Level Details

A. Required Analog Enhancements for HDC

Most of the enhancements required to implement HDC in a pipelined ADC stage are digital. The only exception is that the DAC must be modified as described in this section. To simplify the notation, the description is in the context of HDC applied to first stage.

The calibration signal, $c[n]$, described above takes on values of $-3\Delta/16$, $-\Delta/16$, $\Delta/16$, and $3\Delta/16$ and the output of the flash ADC, $x_1[n]$, takes on values of -4Δ, -3Δ, -2Δ, …, 3Δ, and 4Δ, so a 69-level DAC with step size of $\Delta/8$ is required to represent the signal $x_1[n] + c[n]$. Although a 69-level DAC could be implemented, as explained later it is more convenient to implement a 65-level DAC. Therefore, a 65-level DAC has been implemented in this design. To ensure that $x_1[n] + c[n]$ stays within the 65 level range of the DAC, the addition of $c[n]$ is disabled when $x_1[n] = \pm 4\Delta$, i.e., when the input to the flash ADC has a magnitude between 3.5Δ and 4.5Δ[3].

A 65-level DAC can be implemented by adding the outputs of 64 unity-weighted 1-b DACs. In absence of mismatches among the 1-b DACs, the output of the DAC would be $y[n] = x_1[n] + c[n]$. Unfortunately, mismatches among the 1-b DACs inevitably introduced during fabrication cause the output of the DAC to be

$$y[n] = \alpha_{DAC}\left(x_1[n] + c[n]\right) + \beta_{DAC} + e_{DAC}[n], \tag{39}$$

[3] Disabling the addition of $c[n]$ may also slow down the convergence of the HDC algorithm, because the samples for which $c[n]$ is disabled are not used in the correlation processes described by (36). For example, a sinusoid with a full-scale amplitude of 4.5Δ would cause the addition of $c[n]$ to be disabled 43% of the time. No reduction in convergence time occurs for signals with magnitudes below 3.5Δ.

where α_{DAC} is a constant gain, β_{DAC} is a constant offset, and $e_{DAC}[n]$ is non-constant error referred to as *DAC noise* [40]. If each possible value of $x_1[n] + c[n]$ is mapped to a unique set of input bits to the 64 1-b DACs, then $e_{DAC}[n]$ is a deterministic nonlinear function of the flash ADC output signal, in which case the DAC is a source of nonlinear distortion.

Dynamic Element Matching (DEM) is used in this design to eliminate the DAC as a significant source of nonlinear distortion. The idea behind DEM is that for most values of $x_1[n] + c[n]$, there are multiple ways to set the input bits of the 1-b DACs that would yield $y[n] = x_1[n] + c[n]$ in the absence of mismatches among the 1-b DACs. A *DEM encoder* prior to the 1-b DACs pseudo-randomly selects one of these valid sets of input bits each sample period, in such a way that $e_{DAC}[n]$ has zero mean and is uncorrelated with $x_1[n] + c[n]$.

In a pipeline stage, the propagation delay between the output of the Flash ADC and the output of the DEM encoder must be minimized, because it reduces the time available for the sampling phase or the amplification phase of the stage. Reducing the time for the sampling phase reduces the pipelined ADC's input signal bandwidth, whereas reducing the time for the amplification phase reduces the settling time available for the op-amp. In this implementation, the time allocated for the propagation delay is about 300ps.

One way to achieve this target propagation delay is to use two layers of parallel transmission gates (T-gates). The first layer of $4 \times 64 = 256$ T-gates would compute the sum $x_1[n]$ and $c[n]$, and the second layer of $64^2 = 4096$ T-gates would implement

the DEM encoder. Unfortunately this solution requires 4352 T-gates with an estimated area of 60,000 μm^2 not including the overhead due the routing. Therefore, this strategy was considered impractical.

Instead the segmented DEM DAC shown in Figure 16 has been implemented. The structure has 14 1-b DACs, 8 with weight Δ, 2 with weight $\Delta/2$, 2 with weight $\Delta/4$ and 2 with weight $\Delta/8$. As quantified below, the benefit of this structure is that the DEM encoder can be implemented with much lower circuit area than the DEM encoder mentioned above. However, as explained in [42], a fundamental limitation of segmented DEM DACs is that to achieve the desired $e_{DAC}[n]$ properties they must limit the range of their input sequences to less than the total possible output range of their 1-b DACs. In this case, the input bits of the 1-b DACs shown in Figure 16 could be set to achieve any of 79 output levels, but the DEM encoder is only able to achieve the desired $e_{DAC}[n]$ properties if it limits the range of the DAC to 65 levels. Therefore, the segmented DEM DAC requires 22% more capacitance in its set of 1-b DACs than would be required in a comparable DEM DAC with 64 unity-weighted 1-b DACs. Furthermore, the extra capacitance introduces a 0.5dB increase in KT/C noise. Nevertheless, these drawbacks are considered worthwhile tradeoffs for the reduction in DEM encoder complexity.

The signal processing performed by the DEM encoder is that of the segmented tree-structured DEM encoder shown in Figure 17 with an input sequence given by[4]

[4] The use of notation $c_{k,r}[n]$ for the DEM encoder signals has been chosen to align with [42, 43] ($c_{k,r}[n]$ should not be confused with the calibration sequence $c[n]$).

$$c_{6,1}[n] = 8 \cdot \frac{x_1[n] + \left(c[n] - \Delta/16\right)}{\Delta} + 39\,. \tag{40}$$

The DEM encoder is similar to that presented in [42, 43][5]. It consists of 13 digital *switching blocks*. Switching blocks $S_{6,1}$, $S_{5,1}$, and $S_{4,1}$, are called *segmenting switching blocks* and the remaining switching blocks are called *non-segmenting switching blocks*. Each switching block calculates its two output sequences as a function of its input sequence and one of 13 pseudo-random 1-b sequences, $d_{k,r}[n]$, k = 1, 2, …, 6, and r = 1, 2, …, 6, and 7. The $d_{k,r}[n]$ sequences are designed to well-approximate white random sequences that are independent from each other and $x[n]$, and each take on values of 0 and 1 with equal probability. The switching blocks function as shown in Figure 17 with

$$s_{k,1}[n] = \begin{cases} 0, & \text{if } c_{k,1}[n] = \text{odd}, \\ 1, & \text{if } c_{k,1}[n] = \text{even}, d_{k,1}[n] = 1, \\ -1 & \text{if } c_{k,1}[n] = \text{even}, d_{k,1}[n] = 0. \end{cases} \tag{41}$$

for k = 4, 5, and 6, and

$$s_{k,r}[n] = \begin{cases} 0, & \text{if } c_{k,r}[n] = \text{even}, \\ 1, & \text{if } c_{k,r}[n] = \text{odd}, d_{k,r}[n] = 1, \\ -1 & \text{if } c_{k,r}[n] = \text{odd}, d_{k,r}[n] = 0. \end{cases} \tag{42}$$

for k = 1, 2, and 3, and r = 1, 2, …, 6, and 7.

In this design, the addition of $x_1[n]$ and $c[n]$ and the functionality of the segmented DEM encoder described above are implemented such that all time-critical operations are performed by a layer of 24 parallel T-gates followed by a layer of 64 parallel T-gates. Therefore, the total propagation time is equal to that of 2 cascaded T-

[5] A 69-level version of the DEM DAC could have been implemented by combining the techniques presented in [42] and [44], but it would have been more complex.

gates. The way this is achieved is explained below.

Note from (40) and (41) that the amplitude of $s_{6,1}[n]$ is chosen depending upon the parity of $c_{6,1}[n]$, which in turn depends only on $c[n]$. Therefore the calculations performed by switching block $S_{6,1}$ do not involve the output of the flash ADC, $x_1[n]$, so they can be performed before $x_1[n]$ is available. Similar reasoning applies to switching blocks $S_{5,1}$, $S_{4,1}$, $S_{1,1}$, $S_{1,2}$, and $S_{1,3}$. Therefore the input bits to the bottom 6 1-b DACs in Figure 17 can be computed before the Flash ADC output is available.

By similar reasoning it can be verified that upper output of $S_{4,1}$ is

$$c_{3,1}[n] = x_1[n]/\Delta + 4 + q[n] \tag{43}$$

where $q[n]$ is a function of $c[n]$, $s_{6,1}[n]$, $s_{5,1}[n]$, and $s_{4,1}[n]$, and can take on values of only –1, 0, and 1. A low-latency implementation of (43) is achieved by combining 24 parallel T-gates that can implement the three input-output connection configurations shown in Figure 18. In this design $x_1[n]/\Delta + 4$ can take on values of 0, 1, 2, …, 7, and 8 and is provided by the flash ADC in the form of a thermometer code. Therefore, the mapping shown in Figure 18 results in a thermometer code representation of $x_1[n]/\Delta + 4 + q[n]$. Since $q[n]$ is known before the flash ADC output is ready, the T-gate configuration is selected in advance, thereby minimizing latency.

Figure 19 shows the implementation of the combined $x_1[n] + c[n]$ summer and the segmented DEM encoder. The first layer of 24 T-gates implements (43) as described above, and the second layer of 64 T-gates maps the combined operation of $S_{3,1}$, $S_{2,1}$, $S_{2,2}$, $S_{1,4}$, $S_{1,5}$, $S_{1,6}$, $S_{1,7}$. As demonstrated in [40] and [41] the configuration of such T-gates does not depend on the signal $c_{3,1}[n]$, so it is selected in advance.

Standard logic is used to realize all components for which the timing is not critical. These components include the pseudo-random number generator (PRNG), switching blocks $S_{6,1}$, $S_{5,1}$, $S_{4,1}$, $S_{1,1}$, $S_{1,2}$, and $S_{1,3}$, and all the circuitry that drives the T-gates.

As described above, when $x_1[n] = \pm 4\Delta$ the addition of $c[n]$ must be disabled. To do this with minimal latency, the inputs to the bottom six 1-b DACs in Figure 16 are computed for both cases (addition of $c[n]$ enabled and disabled) and then the correct choice is selected by a multiplexer. For the upper 8 1-b DACs, the T-gates in the first layer provide the disabling function: it can be verified from Figure 18 that when the flash ADC output is full scale (all the thermometer bits are 1's or 0's), $q[n]$ has no effect on the signal.

B. HDC Digital Logic

The details of the HDC block in Figure 15 are shown in Figure 20. The HDC block implements the calculations described in Section II.C with some extra features to reduce complexity.

The signal $s_1[n]$ is requantized to 6 bits prior to the correlators, and its squared value is requantized to 4 bits prior to the η_2 average and dump operation to reduce the size of the HDC logic. *Dithered requantizers* are used to ensure that the resulting quantization noise is zero-mean and uncorrelated with $s_1[n]$ [45], which is sufficient to avoid corrupting the correlations. Although the quantization noise slows down the correlation process slightly, it has been found from simulation and approximate analysis that the increase in convergence time is negligible.

The average and dump blocks shown in Figure 20 compute $A^2\gamma_1$, $6A^6\gamma_3$ and η_2 according to (37). They each average 2^P *valid* samples (i.e. samples for which $c[n]$ was enabled) where P is an integer between 28 and 34 (selectable via serial port control), and then output the average. Each time a new average is produced, the average and dump blocks are reset and start over. When new averages are ready, they are scaled as necessary to obtain γ_1, γ_3 and η_2, and further processed by digital logic that implements (38). Since the estimates are updated only once every 2^P valid samples, all the post processing only needs to produce new values at the same low rate, thus allowing a low-power and low-area realization. Four full speed multipliers and one adder implement (35).

C. Application of HDC and DNC to Multiple Pipeline Stages

In this design, HDC has been applied to the first three stages, as shown in Figure 21. The calibration sequences $c_1[n]$, $c_2[n]$ and $c_3[n]$ are added and three HDC blocks, labeled HDC1, HDC2, and HDC3, are included in the first three stages, respectively. As explained in Section III.D, the coefficient estimation process performed by HDC in a given stage is most accurate when the residue amplifier in that stage is the primary source of distortion. Therefore, the coefficient estimation process is implemented first in Stage 3, then in Stage 2, and then in Stage 1, at which point the cycle is repeated.

Figure 21 also shows blocks that implement DNC in the first three pipeline stages. Each DNC block estimates and cancels the DAC noise introduced by the cor-

responding 65-level DAC. The implementation of DNC has been explained in [41], so further description of it is omitted from this paper.

HDC and DNC operate in background during normal operation of the pipelined ADC, so they adapt to environmental changes without interrupting normal ADC operation. However, initial convergence requires 130 seconds in this design (approximately 43 seconds per stage limited by HDC convergence). Although not implemented, simulations indicate that an auto-calibration mode with no changes to HDC or DNC except for zeroing the input signal and using $4c_k[n]$ in place of $c_k[n]$ in each of the first three stages would reduce the initial convergence time to less than a second.

D. Effect of Quantizer Nonlinearity on HDC

The estimation portion of HDC in a given pipeline stage measures distortion in that stage's digitized residue regardless of its source, but the correction portion of HDC is based on the assumption that all of the measured distortion came from that stage's residue amplifier. Therefore, the theory behind HDC implicitly assumes that the residue amplifier is the only significant source of nonlinearity in the pipelined ADC.

Nevertheless, under certain conditions the flash ADCs can be a significant source of nonlinearity in a pipelined ADC. Ideally, only the last stage's flash ADC quantization noise, i.e., $e_{ADC,6}[n]$, appears in the output of the pipelined ADC and the thermal noise level is high enough that it acts like dither and prevents $e_{ADC,6}[n]$ from

introducing significant nonlinear distortion. However, due to non-zero α_1 and α_3, the quantization error from the first 5 stages is not perfectly canceled in practice, and this causes leakage of quantization error to the output. Even in the stages with HDC some leakage of quantization error always occurs, because the estimation process is never perfect.

To simplify the description of the problem, suppose the pipelined ADC has ideal components except that the residue amplifier in the second stage has a non-zero value of α_1 such that a fraction, λ_2, of quantization error, $e_{ADC,2}[n]$, leaks into the digitized residue, $r_1[n]$. In this case it follows from (37) that γ_1 and γ_3 of HDC1 contain terms $\frac{1}{4}\lambda_2\langle e_{ADC,2}[n]t_1[n]\rangle$ and $\frac{1}{4}\lambda_2\langle e_{ADC,2}[n]t_1[n]\ t_2[n]\ t_3[n]\rangle$, respectively. If either of these terms are non-zero, the estimates of the parameters α_1 and α_3 in stage 1 are corrupted. For example, it can be shown that when the pipelined ADC input stays in the range $-\Delta/16$ to $\Delta/16$, $\langle e_{ADC,2}[n]t_1[n]t_2[n]t_3[n]\rangle = A^3\Delta/4$, whereas when the pipelined ADC input stays in the range $\Delta/16$ to $3\Delta/16$, $\langle e_{ADC,2}[n]t_1[n]t_2[n]t_3[n]\rangle = -A^3\Delta/4$. In both cases even a small leakage of $\lambda_2 = 1\times10^{-4}$ causes the estimate of α_3 to have a magnitude of 0.139V^{-2} (recall that $\alpha_3 = 0$ for this hypothetical example). Thus, the estimation error is almost 50% of the typical α_3 value of -0.3V^{-2} achieved by the op-amps used in the pipelined ADC prototype IC.

More generally, for DC and low-amplitude pipelined ADC input signals the correlation between $e_{ADC,2}[n]$ and $t_1[n]t_2[n]t_3[n]$ tends to be nonzero, which corrupts the estimates of α_3. This problem is negligible when the input signal amplitude is

greater than –9dBFS, because in such cases the signal at the input of the second stage's flash ADC is sufficiently busy that $\langle e_{ADC,2}[n]\, t_1[n]\, t_2[n]\, t_3[n] \rangle$ is close to zero.

The problem described above is not unique to HDC. It affects all of the known residue amplifier distortion calibration techniques because they each use a pipelined ADC to measure the distortion from its own residue amplifiers. In each case, the distortion from all sources in the pipelined ADC is measured but the distortion is corrected under the assumption that the residue amplifiers are the only significant distortion sources.

E. An Improved Calibration Sequence

One way to mitigate the problem is to use a calibration sequence consisting of 5 $t_i[n]$ sequences instead of 3 $t_i[n]$ sequences, i.e. $c[n] = t_1[n]+t_2[n]+t_3[n]+t_4[n]+t_5[n]$. The 2 extra sequences add enough dither to decorrelate $e_{ADC,2}[n]$ from $t_1[n]$ and $t_1[n]\, t_2[n]\, t_3[n]$, thus allowing successful convergence even for low-amplitude pipelined ADC input signals.

A drawback of this solution is that the two extra $t_i[n]$ sequences increase the magnitude of the calibration sequence by $\Delta/8$, so additional over range margin is used by the terms in the residue amplifier output associated with the calibration sequence. In this design, the drawback was found to be acceptable because the extra area and power consumption required to correspondingly reduce the flash ADC threshold errors, negligibly increased the overall circuit area and power consumption of the pipelined ADC prototype IC.

Unfortunately, a wiring mistake was made in the pipelined ADC prototype IC which caused the HDC calibration sequence in each HDC-enabled pipeline stage to be the sum of three rather than five $t_i[n]$ sequences. All five sequences are generated on chip for each HDC-enabled pipeline stage, but the wiring mistake prevented two of the five sequences from being used. As described above, the consequence of this mistake is that each stage's estimates of the α_1 and α_3 coefficients tend to be wrong for small-amplitude pipelined ADC input signals.

Normally, each HDC block continually updates its coefficients, but it can optionally freeze the coefficients after convergence. In order to test the pipelined ADC for small-amplitude input signals, the measurement results presented in Section V were obtained by allowing the HDC blocks to converge with a large-amplitude pipelined ADC input signal and then freezing the coefficients via serial port control. Measurements with frozen coefficients for numerous input signals, including small input signals, indicate that full performance is achieved in all cases as expected.

F. A Detection Method for Invalid Correlations

It can be shown that the dithering effect provided by adding the five $t_i[n]$ sequences completely removes any correlation between $e_{ADC,2}[n]$ and the sequences $t_1[n]$ and $t_1[n]t_2[n]t_3[n]$, in the case where $\alpha_1 = 0$. When $\alpha_1 \neq 0$, there are some small ranges of pipelined ADC input signals for which the quantization error $e_{ADC,2}[n]$ is still correlated with $t_1[n]t_2[n]t_3[n]$. For example, suppose $\alpha_1 = -0.06$. Then a pipelined ADC input signal that stays between 0.051Δ and 0.066Δ causes $\langle e_{ADC,2}[n]\ t_1[n]$

$t_2[n]\, t_3[n]\rangle = A^3\Delta/32$, and one that stays between 0.066Δ and 0.081Δ causes $\langle e_{ADC,2}[n]\, t_1[n]\, t_2[n]\, t_3[n]\rangle = -A^3\Delta/32$, both of which would lead to invalid estimates of the distortion parameters.

As seen in the above example, the ranges of inputs for which $\langle e_{ADC,2}[n]\, t_1[n]\, t_2[n]\, t_3[n]\rangle \neq 0$ come in doublets, and each side of the doublet has opposite correlations. Therefore input signals that are sufficiently busy spend roughly equal time in each of the two sides of the doublet, so the effect of this undesired correlation tends to be small. Consequently the existence of these doublets does not seem to affect HDC accuracy significantly.

However, if necessary it is possible to detect invalid correlation results so as to avoid updating the distortion coefficients, until a new successful estimate is available. In the remainder of this section, a method to detect a bad estimate is described[6].

The output $x_2[n]$ of the Flash ADC of stage 2 can be correlated against $t_1[n]t_2[n]t_3[n]$ and the result can be used with the output of the second correlator of HDC1 to generate an estimate of the quantity

$$\left\langle s_1[n]\cdot t_1[n]t_2[n]t_3[n]\right\rangle - \frac{1}{4}\left\langle x_2[n]\cdot t_1[n]t_2[n]t_3[n]\right\rangle . \tag{44}$$

It follows from the reasoning described in Section II that

$$\frac{1}{4}\left\langle x_2[n]\cdot t_1[n]t_2[n]t_3[n]\right\rangle = -6A^6\alpha_3 + \frac{1}{4}\left\langle e_{ADC2}[n]\cdot t_1[n]t_2[n]t_3[n]\right\rangle , \tag{45}$$

and

[6] The authors were not aware of the presence of such doublets until after the tape out, so the detection method described in this section was not implemented in the pipelined ADC prototype IC.

$$\left\langle s_1[n]\cdot t_1[n]t_2[n]t_3[n]\right\rangle = -6A^6\alpha_3 + \frac{\lambda_2}{4}\left\langle e_{ADC2}[n]\cdot t_1[n]t_2[n]t_3[n]\right\rangle, \tag{46}$$

where the last term in (46) is due to the leakage of quantization error. In practice λ_2 << 1 and $|\langle e_{ADC,2}[n]t_1[n]t_2[n]t_3[n]\rangle| < A^3\Delta/2$. Therefore, whenever the magnitude of the estimate of (44) is larger than $\lambda_{2\text{-max}}A^3\Delta/2$, where $\lambda_{2\text{-max}}$ is an upper bound on λ_2, it indicates that $\langle e_{ADC,2}[n]t_1[n]t_2[n]t_3[n]\rangle$ is large enough to cause the HDC estimates to be corrupted. Furthermore, $\lambda_{2\text{-max}}$ need not be a tight upper bound on λ_2. For example, in this system setting $\lambda_{2\text{-max}}$ = 0.01 would work because it would ensure that the estimate of α_3 is precise within 10% assuming $\lambda_2 = 1\times10^{-3}$ and $\alpha_3 = -0.3$, which is sufficient accuracy to achieve the target ADC performance.

IV. Analog Circuit Details

Additional details of the analog and mixed-signal portions of the first three pipeline stages are shown in Figure 22(a). The fourth and fifth stages have a similar structure, except no calibration sequence is added, and therefore a 9-level DEM DAC with a step size of Δ is used in place of the 65-level DEM DAC. Additional details of the last pipeline stage are shown in Figure 22(b).

As shown in Figure 22(a) the continuous time input signal is sampled by two separate passive sampling networks, the outputs of which are connected to the residue amplifier and the flash ADC, respectively [46]. The DAC is realized as a separate circuit connected to the input terminals of the residue amplifier.

A differential switched capacitor unit sampling cell with a simplified timing

diagram are shown in Figure 23. The sampling network of the residue amplifier consists of 8 such unit sampling cells in parallel, whereas each of the 8 comparators of the 9-level flash ADC uses a quarter-size version of a single unit sampling cell. Both the capacitors and switches are scaled proportionally such that both sampling paths have the same nominal time constant. The switches between the top plates of the capacitors and inputs of the op-amp provide isolation from the op-amp during the sampling phase to improve matching between the two sampling networks as described in [41], and ensure that the residue amplifier's α_i coefficients do not depend on the input signal. Bootstrapped switches of the type presented in [41] are used for the continuous-time input sampling switches of both sampling networks to achieve the necessary linearity.

A separate switched capacitor network has been used for the DAC to prevent signal-dependent loading of the voltage references. This benefit comes at the expense of a 3dB increase in KT/C noise and a reduced residue amplifier feedback factor (1/10 versus 1/6) relative to a design in which the DAC and sampling network share the same capacitors.

The sampling network of each 1-b DAC with weight Δ is shown in Figure 24. Scaled versions of the same cell have been implemented for the 1-b DACs with weights $\Delta/2$, $\Delta/4$ and $\Delta/8$ as required to implement the DAC in Figure 17. The input bit b from the DEM encoder to each 1-b DAC controls the Swap Cell during ϕ_{2d}. Bootstrapped switches were not used for the DAC switches.

The residue amplifier is shown in Figure 25. During the amplification phase,

ϕ_{2d}, the feedback capacitor C_f is connected to the op-amp, whereas during the ϕ_{1d} phase C_f is disconnected from the op-amp and discharged. A Miller compensated two-stage op-amp has been used in the residue amplifier for its wide output dynamic range. During ϕ_{1d}, the op-amp is reset by shorting the differential outputs of both stages to reduce memory effects [47]. Switched capacitor common mode feedback circuitry (not shown in Figure 25) controls the common mode voltage of the output stage by adjusting V_{cmfb}.

The simulated DC open-loop gain and unity gain frequency of the op-amp in the first pipeline stage's residue amplifier are 43dB and 1.2GHz, respectively. The corresponding loop gain of the residue amplifier is 23dB at DC and has a unity gain frequency of 200MHz. The current consumption of the op-amp is 4.8mA from a 1.2V supply.

The voltage references V_{refp} and V_{refm} are generated on chip as shown in Figure 26. A set of resistors between the power supply and ground define the desired voltage reference values (nominally set to 950mV and 265mV, respectively), which are buffered by a pseudo-differential voltage follower and decoupled by external capacitors C_{EXT} [48]. The main drawback of this solution is the need for 2 extra pins for external decoupling. However the large external capacitors provide a low impedance over a wide frequency range, and this relaxes the performance requirements of the internal buffers which need only deliver the average current required by the switching load [49]. The total DC current consumption for the reference generation circuitry is 4mA from a 1.2V supply, including the current through the reference ladders.

The same reference voltages V_{refp} and V_{refm} are shared by all the DACs and flash ADCs in the pipelined ADC. The reference ladders that generate the threshold voltages for the flash ADCs are connected between V_{refp} and V_{refm}. In this design, each flash ADC uses a dedicated reference ladder. The common mode voltages used for the switched capacitor circuits are generated as shown in Figure 27.

Latched comparators of the type presented in [50] are used in the flash ADCs. It is a standard latch with preamplifier consisting of 2 resistive loaded differential pairs in series.

The phase generator has been designed with the strategy described in [41]. A dedicated phase is used as a sampling phase of the first stage only (ϕ_1 in Figure 23), such that the sampling instant (falling edge of ϕ_1) happens before any other switching event. This reduces the chance of corrupting the sampling process with disturbances such as from coupling effects. The sampling phase path to the sampling switches has been optimized to reduce jitter via careful layout and the use of a minimum number of inverters. The simulated jitter of the sampling network is 100fs.

To minimize design time the second through fifth pipeline stages are replicas of the first stage with only minor modifications. The only changes are that the residue amplifier, sampling network, and DAC have been scaled by ½ in stages 2 and 3 and by ¼ in stages 4 and 5, and the op-amp has been scaled by ½ in the stages 3 and 4 and by ¼ in the stage 5. More aggressive scaling would have reduced power and area consumption without sacrificing ADC accuracy. Power and area consumption also could have been reduced without sacrificing ADC accuracy if 1.5-b pipeline stages

had been used after the first three stages [51]. However, neither of these design options were taken because they would have increased design time.

The pipelined ADC prototype IC is implemented in 90nm CMOS technology with a deep nWell option, MiM capacitors, and both high and standard threshold voltage transistors. The circuit is partitioned into four power supply domains: (i) *analog*, (ii) *clock drivers & DEM*, (iii) *clock generator*, and (iv) *digital*, each of which is powered by a separate power supply line. The nominal power supply voltages of the four domains are 1.2V, 1.2V, 1.0V, and 1.0V, respectively. To minimize coupling from the substrate, all active components in the analog sections of the IC are in deep n-wells, the digital core and serial port interface (which were laid out with an automated place-and-route tool) are in a single deep n-well, and the capacitors associated with the switched-capacitor portions of the IC were laid out above n-wells. On-chip decoupling capacitance is used to reduce power supply bounce. All pads have ESD protection circuitry. A die photograph is shown in Figure 28. The IC is 2.15mm by 3.35mm and has an active area of $4mm^2$.

The IC is packaged in a 56-pin QFN package with exposed die paddle. All grounds are down-bonded to the exposed paddle. Critical supply pins and the pins that connect the voltage references to external decoupling capacitors are double-bonded to reduce inductance.

V. Measurement Results

Three randomly chosen copies of the pipelined ADC prototype IC have been

tested. Each IC was soldered to one of three identical printed circuit test boards. Measurement results from the three test boards are reported in this section.

Each test board includes input signal conditioning circuitry, a 100MHz low-jitter crystal oscillator and associated clock conditioning circuitry, voltage regulators, and digital circuitry to facilitate acquisition of the output data from and serial port communication with the pipelined ADC prototype IC. The input conditioning circuitry consists of a transformer followed by two passive RC filter stages to convert the single-ended input signal from an SMA connector to differential form and suppress out-of-band noise and distortion. The clock conditioning circuitry uses a transformer to convert the single-ended clock signal from the 100MHz oscillator to differential form. The output swing of the 100MHz oscillator is 3.3V, so high-speed diodes are connected across the secondary terminals of the transformer to limit the amplitude of the differential clock signal to less than 1V. Four voltage regulators provide the four power supplies of the pipelined ADC prototype IC. Five other voltage regulators provide power supplies for the other components on the test board.

Measurements were performed with a variety of single-tone and two-tone pipelined ADC input signals. For each single-tone measurement, the output of a high-quality sinusoidal laboratory signal source was passed through a custom-made passive bandpass filter with a narrow bandwidth centered near the signal frequency to suppress noise and distortion from the signal source, and the output of the bandpass filter was connected to the test board. For each two-tone measurement, the outputs of two identical sinusoidal laboratory signal sources were added and the resulting signal

was passed through a bandpass filter prior to the test board.

The measured performance from the three test boards was found to be nearly identical after calibration by HDC and DNC. Typical measurement results are shown in Figures 20 and 21. Figure 20 shows representative output power spectral density (PSD) plots from the pipelined ADC with a 49.2MHz, 0dBFS single-tone input signal. The grey plot was measured prior to calibration by HDC and DNC, and the black plot was measured after calibration by HDC and DNC. As indicated in the figure, the measured SNDR values of the ADC prior to and after calibration are 42.9dB and 70dB, respectively. Figure 21 shows measured SNDR, SNR, and SFDR values from the pipelined ADC after calibration by HDC and DNC versus frequency and amplitude. The former were measured with −1dBFS single-tone input signals ranging in frequency over the 50MHz Nyquist band and the latter were measured with 19.2MHz single-tone input signals ranging in amplitude from −69dBFS to 0dBFS.

Figure 22 shows plots of typical measured SFDR and SNR values from the pipelined ADC versus the number of values averaged by each of the HDC correlators for a 19.2MHz, −1dBFS single-tone input signal. The data suggests that full accuracy is achieved when 2^{31} or more values are averaged by the HDC correlators. Nevertheless, for all measurements other than those shown in Figure 22 the HDC correlators were configured to average 2^{32} values, which corresponds to approximately 43 seconds of calibration time per stage at a sampling frequency of 100MHz.

For all of the measurements described above, the analog, clock, and digital supply voltages of the pipelined ADC prototype IC were set to their targeted design

values of 1.2V, 1.0V, and 1.0V, respectively, but the power supply for the clock drivers and DEM was set to 1.35V instead of its targeted design value of 1.2V. When this supply is set to its targeted design value of 1.2V, the peak SNDR decreases by approximately 3dB. Although not predicted by simulations, the authors believe that the clock drivers have insufficient strength to achieve full ADC performance at 1.2V.

Table I shows worst-case measurement results for all three test boards under two different power supply voltage scenarios. The two power supply voltage scenarios are denoted V_{DD} *Test Case 1* and V_{DD} *Test Case 2*. In V_{DD} *Test Case 1* all power supplies, except that for the Clock Drivers and DEM, are set to their targeted design values as described above. In V_{DD} *Test Case 2* the analog power supply is reduced to 1.0V, and the digital power supply is reduced to 0.7V. Although the analog circuitry in the pipelined ADC was not designed to work at this reduced power supply voltage, the measured worst-case reduction in SNDR is only 2.2dB because HDC largely compensates for the significant reduction in analog circuit performance.

As mentioned above, the measured performance from the three test boards was found to be nearly identical after calibration by HDC and DNC. However, the α_1 and α_3 coefficients estimated by the HDC blocks (as read from the pipelined ADC prototype IC via the serial port interface) were found to vary significantly from chip to chip. For example, the estimate of α_3 by HDC in the first stage of the pipelined ADC varied by approximately ±30% from chip to chip about a mean of -0.3V^{-2}. Therefore, HDC is at least partly responsible for the observed uniformity of the measurement results.

Table II shows relevant performance data from the pipelined ADC prototype IC along with those from published state-of-the-art ADCs with comparable bandwidths and SNDR values. Two commonly used figures of merit, *FOM*1 and *FOM*2, are included in the table. Although *FOM*1 is more widely referenced in the academic literature than *FOM*2, the latter is often considered more appropriate for ADCs that are SNR-limited because of low supply voltages [52, 53]. As indicated in the table, both figures of merit for the pipelined ADC prototype IC are better (lower) than those for the only other comparable published ADC that operates from a supply voltage below 1.8V, and are better than those for most of the comparable published ADCs that operate from supply voltages at or above 1.8V.

ACKNOWLEDGEMENTS

The authors are grateful to Nevena Rakuljic for technical advice and analog layout assistance, Everest Zuffetti for technical advice, digital synthesis, place and route, Angelo Contini for CAD support, Gerry Taylor and Kevin Wang for technical advice, and Cuong Vu for lab assistance. The authors would like to acknowledge Maurizio Zuffada and Pietro Erratico from STMicroelectronics for support of this research project.

Chapter 2, in full, has been submitted for review to the IEEE Journal of Solid-State Circuits. The dissertation author is the primary investigator and author of this paper. Professor Ian Galton supervised the research which forms the basis of this paper.

REFERENCES

30. S. Devarajan, L. Singer, D. Kelly, S. Decker, A. Kamath, P. Wilkins, "A 16b 125MS/s 385mW 78.7dB SNR CMOS Pipeline ADC,", in *ISSCC Dig. Tech. Papers*, pp. 86-87, February 2009.

31. A.M.A. Ali, C. Dillon, R. Sneed, A.S. Morgan, S. Bardsley, J. Kornblum, Lu Wu, "A 14-bit 125 MS/s IF/RF Sampling Pipelined ADC With 100 dB SFDR and 50 fs Jitter," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 8, pp. 1846-18550, August 2006.

32. Jipeng Li, Un-Ku Moon, "A 1.8-V 67-mW 10-bit 100-MS/s pipelined ADC using time-shifted CDS technique," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 9, pp. 1468-1476, September 2004.

33. Yun-Shiang Shu; Bang-Sup Song, "A 15-bit Linear 20-MS/s Pipelined ADC Digitally Calibrated With Signal-Dependent Dithering," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 2, pp. 342-350, February 2008.

34. Jong-Bum Park, Sang-Min Yo, Se-Won Kim, Young-Jae Cho, Seung-Hoon Lee, "A 10-b 150-MSample/s 1.8-V 123-mW CMOS A/D converter with 400-MHz input bandwidth," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 8, pp. 1335-1337, August 2004.

35. C.R. Grace, P.J. Hurst, S.H. Lewis, "A 12-bit 80-MSample/s pipelined ADC with bootstrapped digital calibration," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 5, pp.1038-1046, May 2005.

36. B. Murmann, B. Boser, "A 12b 75MS/s Pipelined ADC using Open-Loop Residue Amplification," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 12, pp. 2040-2050, December 2003.

37. A. Panigada, I. Galton, "Digital Background Correction of Harmonic Distortion in Pipelined ADCs," *IEEE Transactions on Circuits and Systems - I: Regular Papers*, vol. 53, no. 9, pp. 1885-1895, September 2006.

38. J. P. Keane, P. J. Hurst, S. H. Lewis, "Background Interstage Gain Calibration Technique for Pipelined ADCs," *IEEE Transactions on Circuits and Systems I*, vol. 52, no. 1, pp. 32-43, January 2005.

39. H. Van de Vel, B. A. J.Buter, H. van der Ploeg, M. Vertregt, G. J. G. M Geelen,

E. J. F. Paulus, "A 1.2-V 250-mW 14-b 100-MS/s Digitally Calibrated Pipeline ADC in 90-nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 4, pp. 1047-1056, April 2009.

40. I. Galton, "Digital Cancellation of D/A Converter Noise in Pipelined A/D Converters," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 47 no. 3, pp. 185-196, March 2000.

41. E. Siragusa, I. Galton, "A Digitally Enhanced 1.8V 15b 40MS/s CMOS Pipelined ADC," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 12, pp. 2126-2138, December 2004.

42. K. L. Chan, N. Rakuljic, I. Galton, "Segmented dynamic element matching for high-resolution digital-to-analog conversion," *IEEE Transactions on. Circuits and Systems I: Reg. Papers*, vol. 55, no. 11, pp. 3383-3392, December 2008.

43. K. L. Chan, J. Zhu, I. Galton, "Dynamic Element Matching to Prevent Nonlinear Distortion From Pulse-Shape Mismatches in High-Resolution DACs," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 9, pp. 2067-2078, September 2008.

44. N. Rakuljic, I. Galton, "Tree-Structured DEM DACs with Arbitrary Numbers of Levels," accepted for publication on *IEEE Transactions on Circuits and Systems - I: Regular Papers*.

45. A. B. Sripad, D. L. Snyder, "A necessary and sufficient condition for quantization errors to be uniform and white," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-25, pp. 442-448, Oct. 1977.

46. I. Mehr, L. Singer, "A 55-mW, 10-bit, 40-Msample/s Nyquist-rate CMOS ADC," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 3, pp. 318-325, March 2000.

47. J. P. Keane, P. J. Hurst, S. H. Lewis, "Digital Background Calibration for Memory Effects in Pipelined Analog-to-Digital Converters," *IEEE Transactions on Circuits and Systems - I: Regular Papers*, vol. 53, no. 3, pp. 511-525, March 2006.

48. C. Pinna, A. Mecchia, G. Nicollini, "A CMOS 64MSps 20mA 0.85mm^2 Baseband I/Q Modulator Performing 13 bits over 2MHz Bandwidth", *2004 Symposium on VLSI Circuits Dig. Tech. Papers*., pp. 152-155, June 2004.

49. L. Singer, S. Ho, M. Timko, D. Kelly, “A 12b 65MSample/s CMOS ADC with 82dB SFDR at 120MHz,” in *ISSCC Dig. Tech. Papers*, pp. 38-39, February 2000.

50. A. Bosi, A. Panigada, G. Cesura, R. Castello, “An 80MHz 4× oversampled cascaded ΔΣ-pipelined ADC with 75dB DR and 87dB SFDR,” in *ISSCC Dig. Tech. Papers*, pp. 174-175, February 2005.

51. D. W. Cline, P. R. Gray, “A power optimized 13-b 5-Msamples/s pipelined analog-to-digital converter in 1.2-μm CMOS,” *IEEE Journal of Solid-State Circuits*, vol. 31, pp. 294–303, Mar. 1996.

52. Y. Chiu, P. R. Gray, B. Nikolic, “A 14b 12 MSps CMOS pipeline ADC with over 100 dB SFDR,” *IEEE Journal of Solid-State Circuits*, vol. 39, no. 12, pp. 2139-2151, December 2004.

53. A. Zanchi, F. Tsay, “A 16-bit 65-MS/s 3.3-V pipeline ADC core in SiGe BiCMOS with 78-dB SNR and 180-fs jitter,” *IEEE Journal of Solid-State Circuits*, vol. 40, no. 6, pp. 1225-1237, June 2005.

54. M. Anthony, E. Kohler, J. Kurtze, L. Kushner, G. Sollner, “A Process Scalable Low-Power Charge-Domain 13-bit Pipeline ADC,” *IEEE Symposium on VLSI Circuits*, pp. 222-223, June 2008.

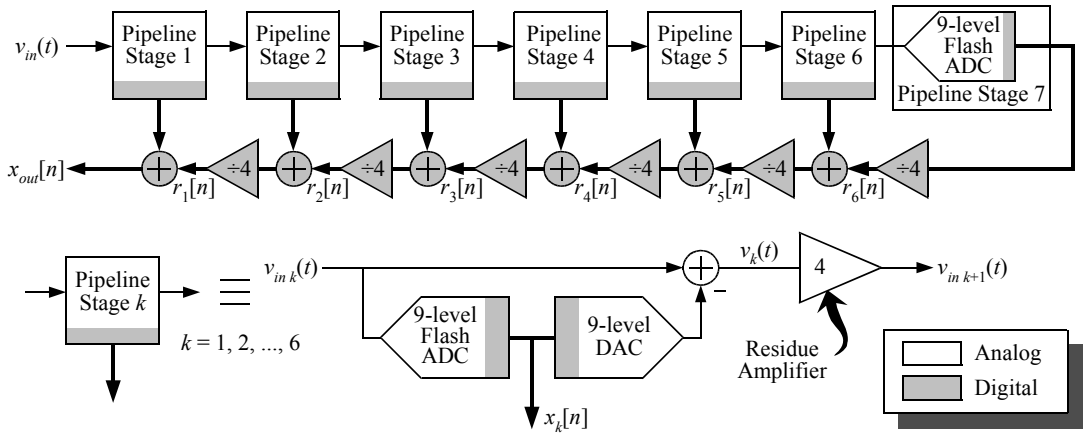

Figure 1 The block diagram of an example 15-bit pipelined ADC.

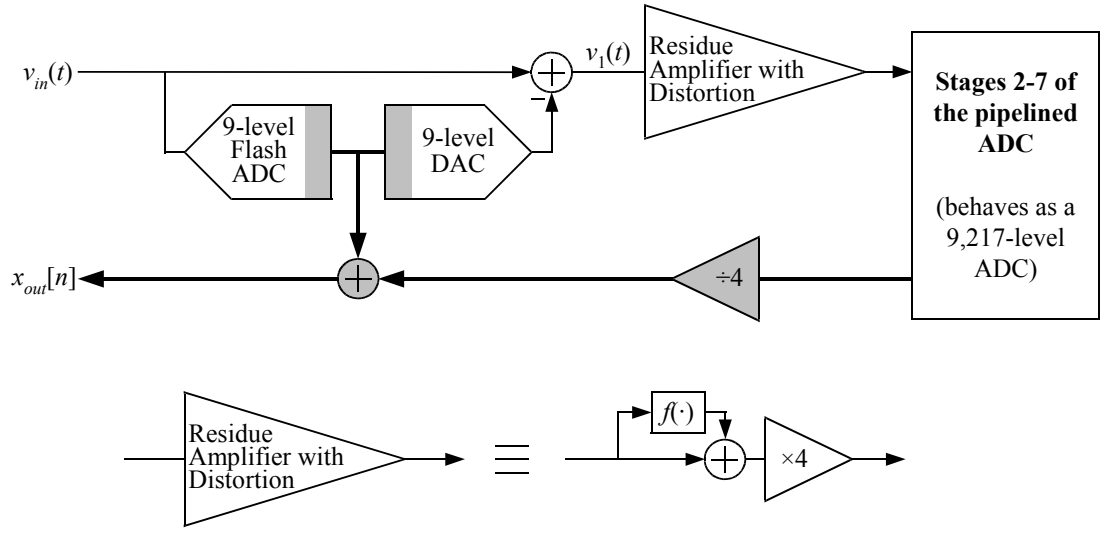

Figure 2 The pipelined ADC of Figure 1 except with a residue amplifier in the first stage that introduces distortion.

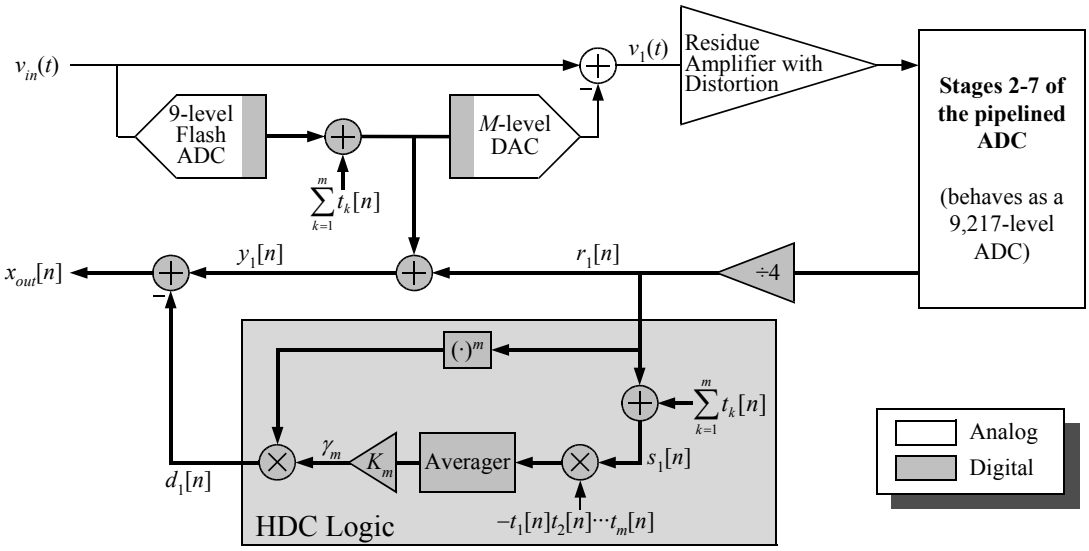

Figure 3 An example of the HDC technique for correction of *m*th-order residue amplifier distortion.

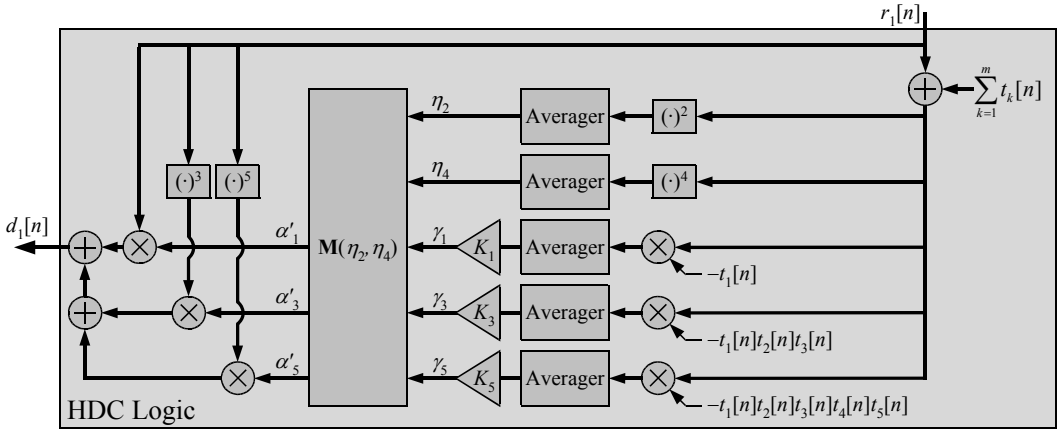

Figure 4 An example of the HDC logic for correction of fist-order, third-order, and fifth-order residue amplifier distortion.

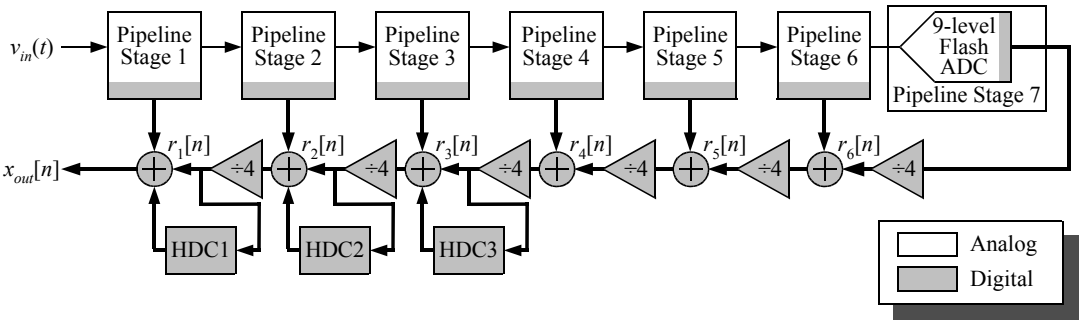

Figure 5 A high-level view of an example pipelined ADC incorporating the HDC technique.

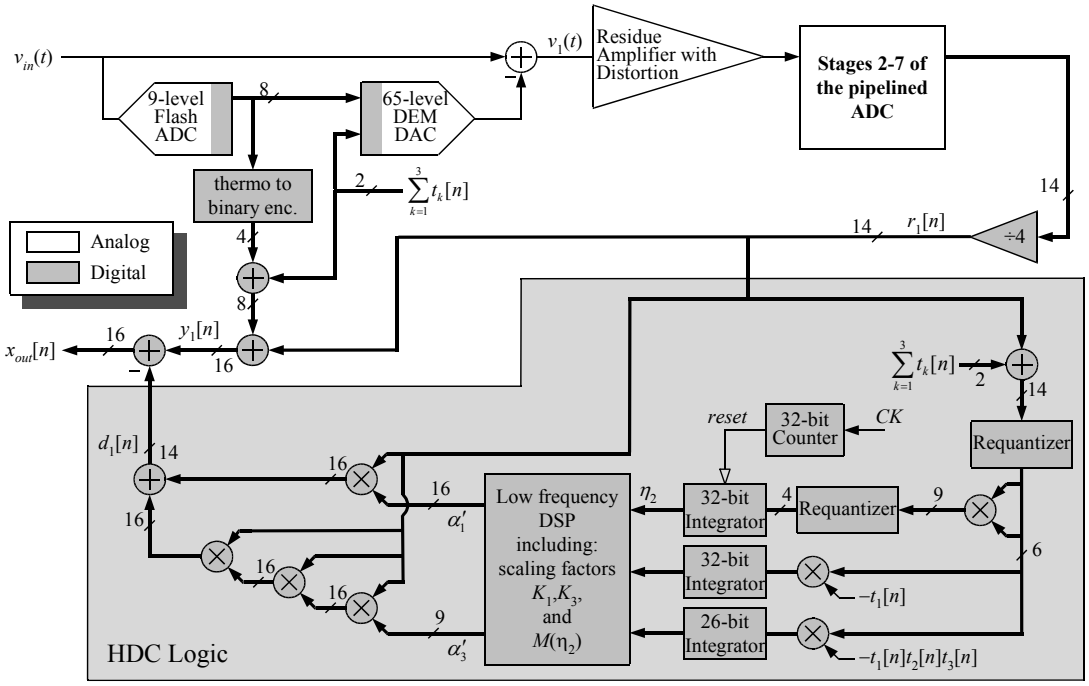

Figure 6 The example pipelined ADC incorporating the HDC technique with expanded views of the first stage and associated HDC logic.

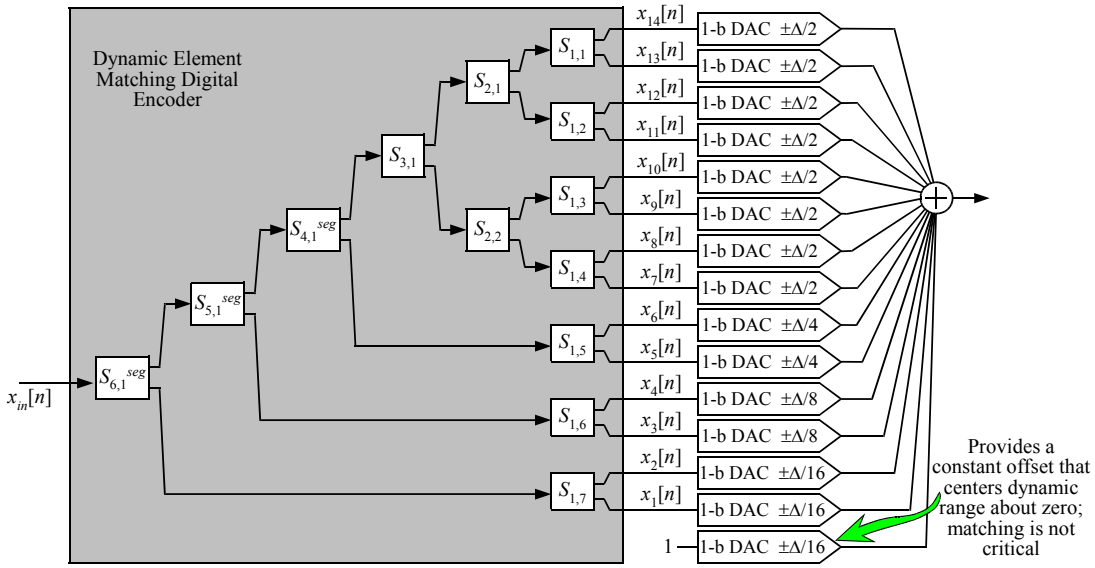

Figure 7 A functional view of the Dynamic Element Matching DAC.

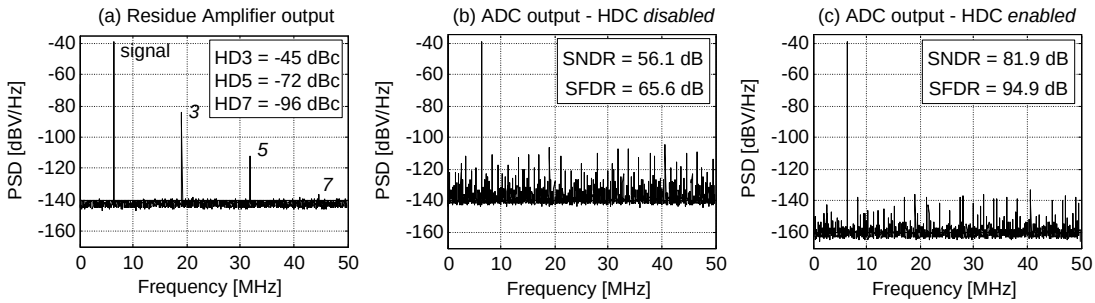

Figure 8 Simulation results for HDC applied to 1st, 2nd and 3rd stages

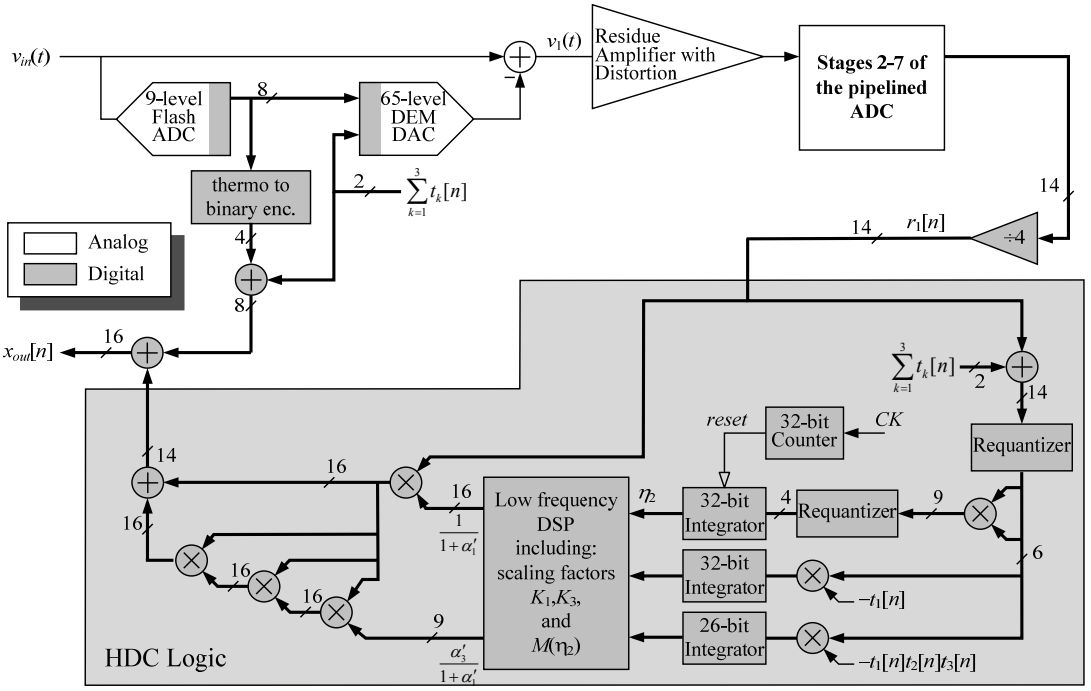

Figure 9 The example pipelined ADC incorporating the HDC technique with improved correction scheme for large distortion coefficients.

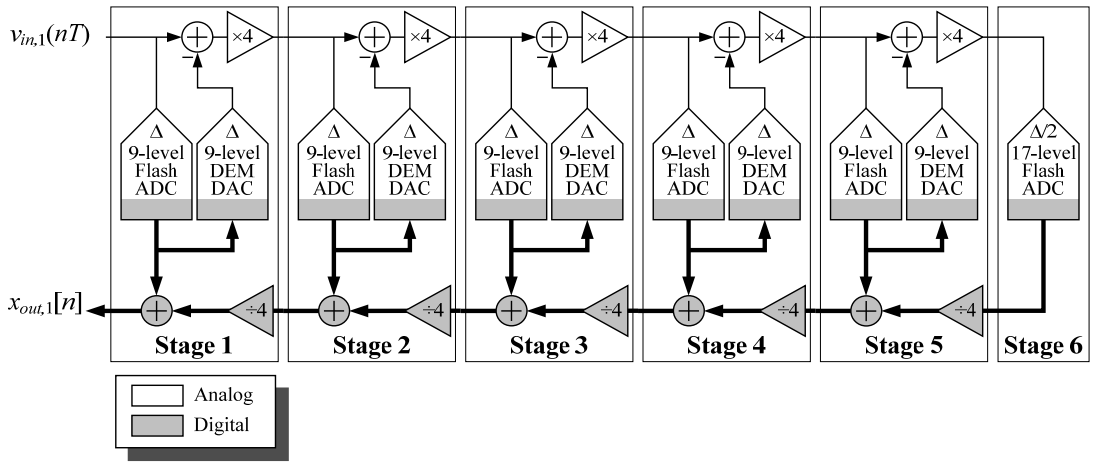

Figure 10 Block diagram of a 14-b pipelined ADC.

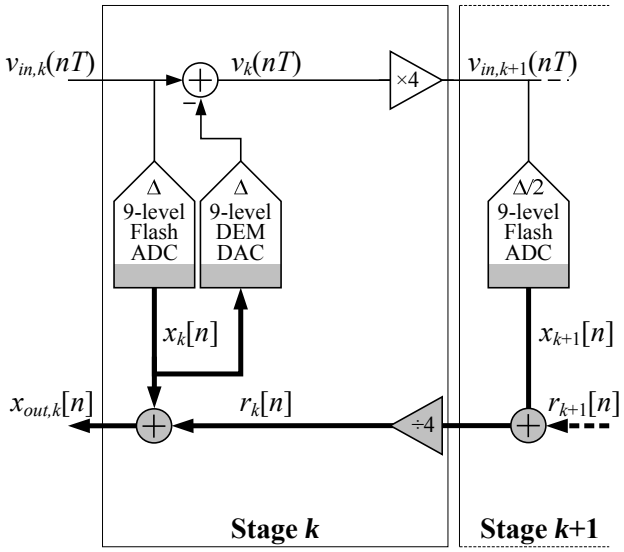

Figure 11 Pipeline Stages *k* and *k*+1.

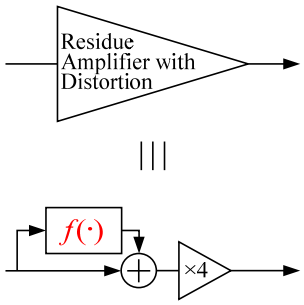

Figure 12 Model of the residue amplifier with distortion.

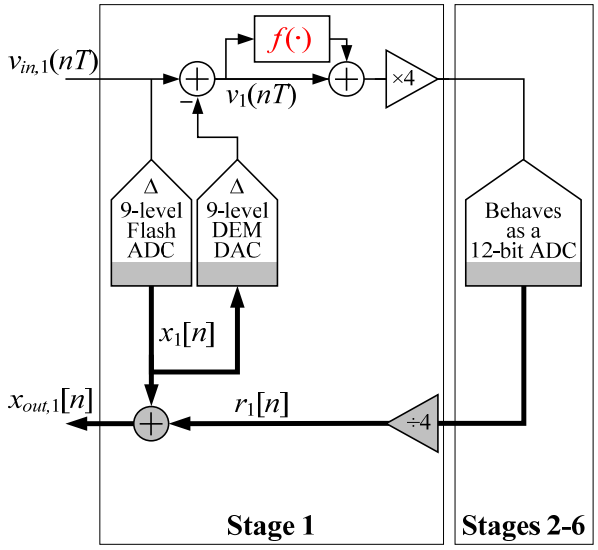

Figure 13 Simplified representation of the 14-b pipelined ADC including distortion from the residue amplifier.

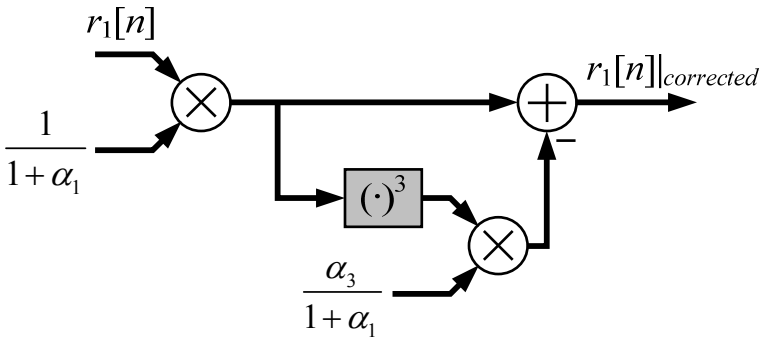

Figure 14 Correction of the distortion in the digitized residue.

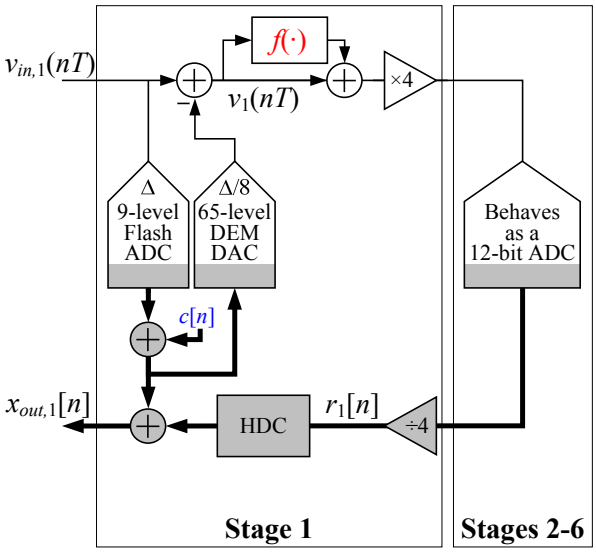

Figure 15 Simplified representation of a 14-b pipelined ADC with HDC applied to the first pipeline stage.

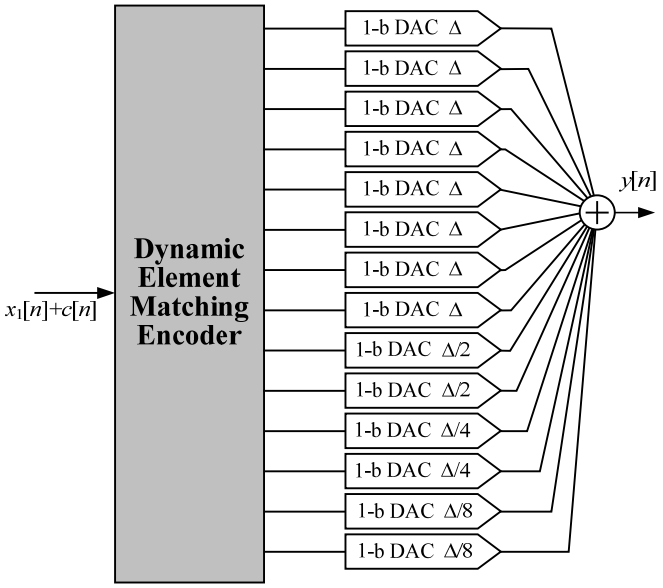

Figure 16 Simplified block diagram of the implemented segmented DEM DAC.

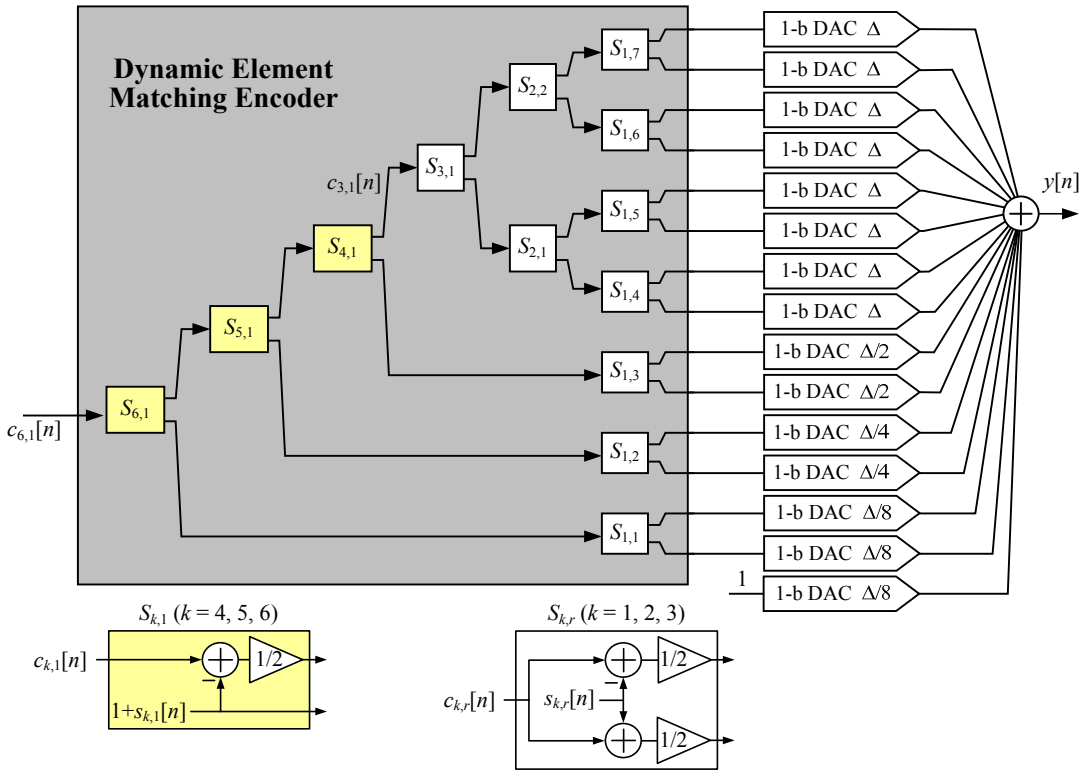

Figure 17 65-level segmented DAC with DEM encoder.

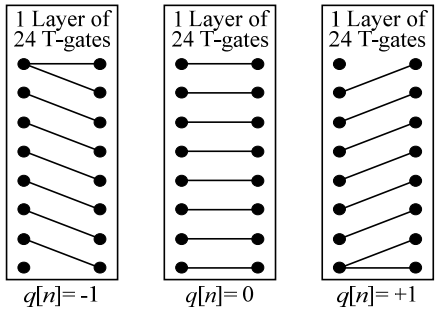

Figure 18 The three configurations of the first layer of T-gates, (only ON switches are shown).

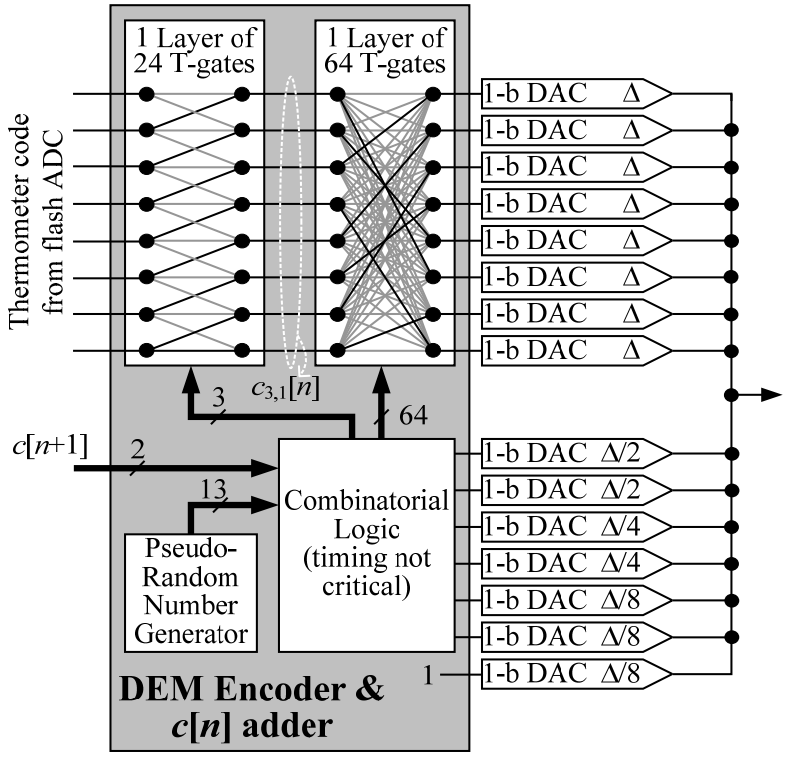

Figure 19 Implementation of the DEM encoder and *c*[*n*] adder.

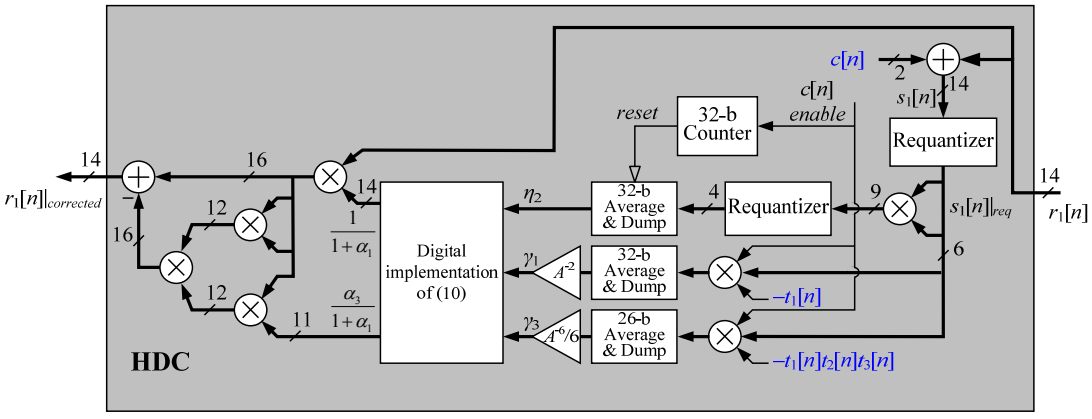

Figure 20 Block diagram of the HDC logic in the first pipeline stage.

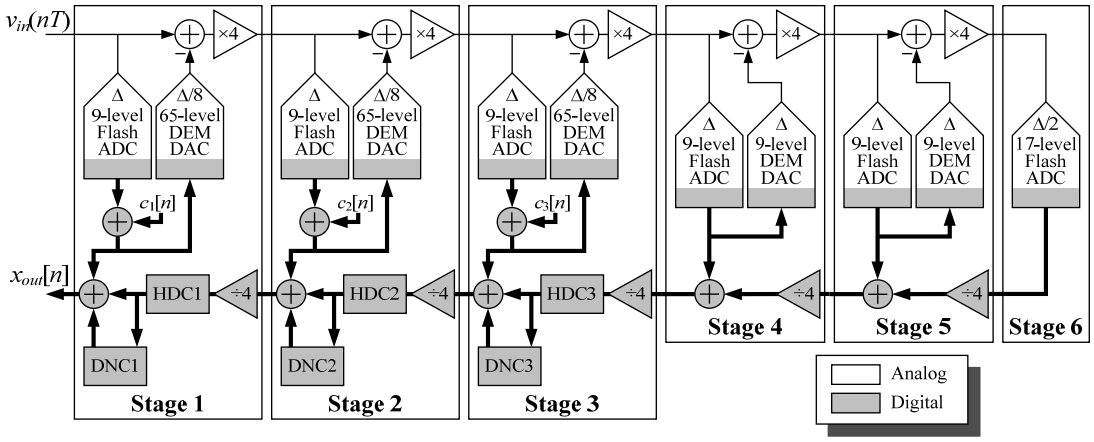

Figure 21 Complete block diagram of the implemented 14-b pipelined ADC.

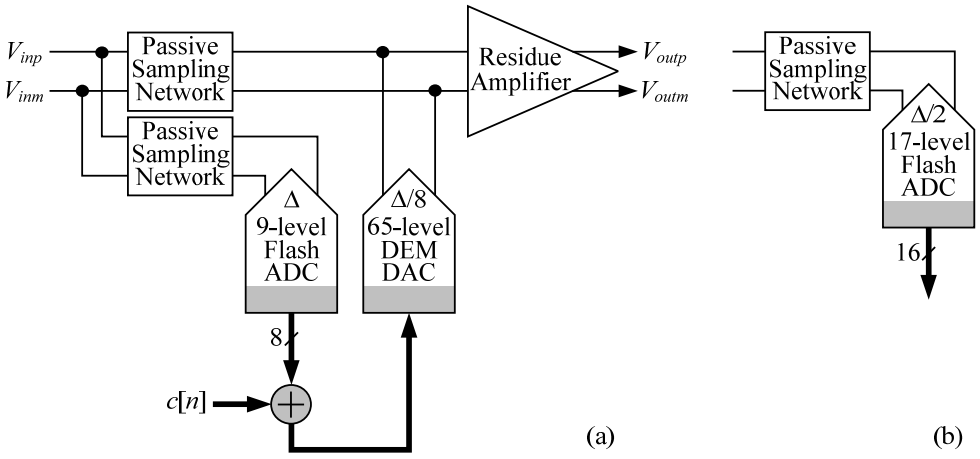

Figure 22 Block diagram of the mixed-signal circuitry in (a) the first three pipeline stages, and (b) the last pipeline stage.

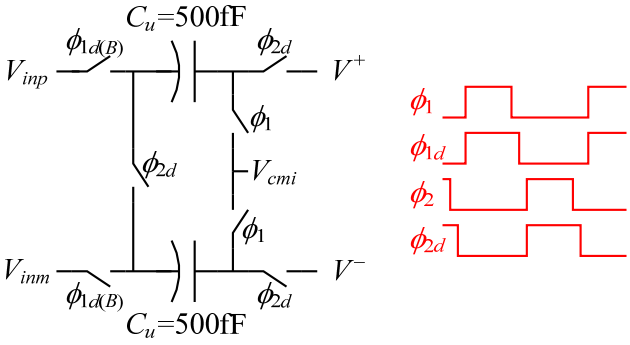

Figure 23 Unit passive sampling network (bootstrap circuit for ϕ_{1d} switches not shown).

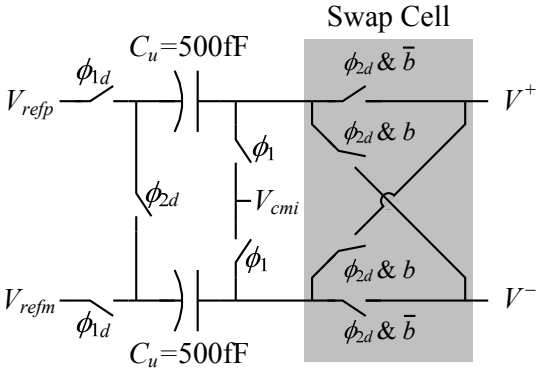

Figure 24 1-b DAC sampling network (weight Δ).

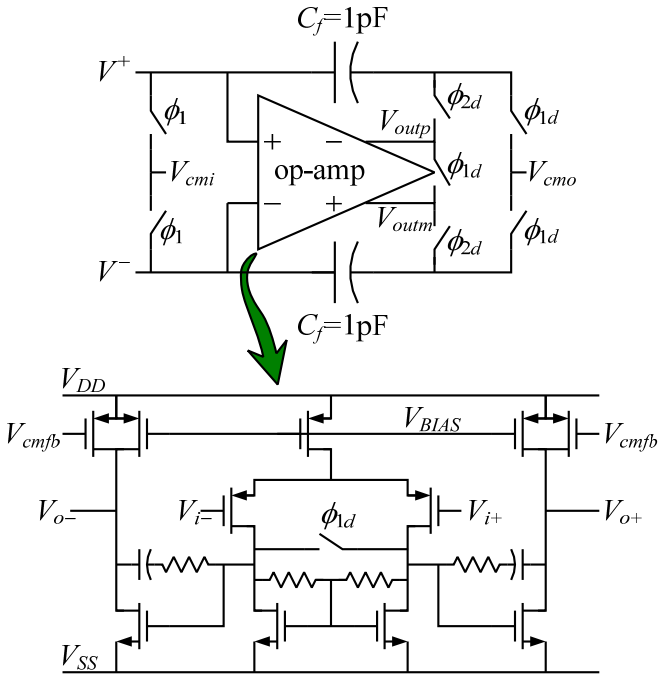

Figure 25 Residue amplifier switched capacitor network and op-amp.

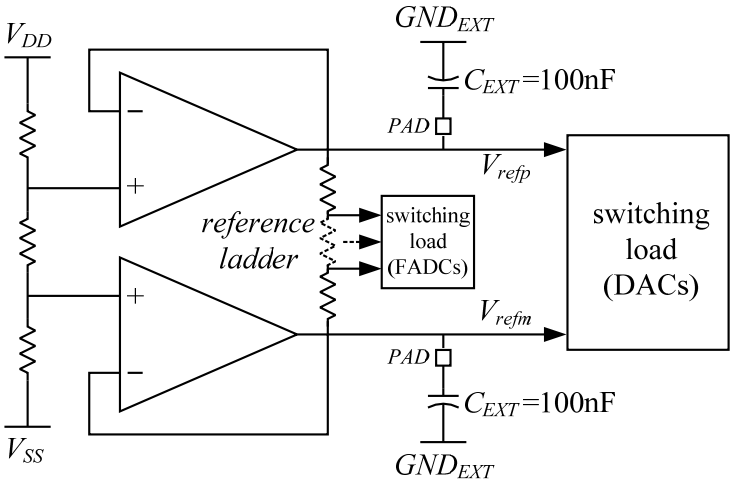

Figure 26 Reference voltage generator.

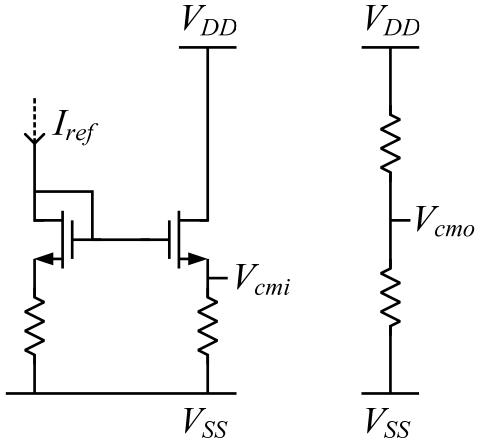

Figure 27 Common mode voltage generators.

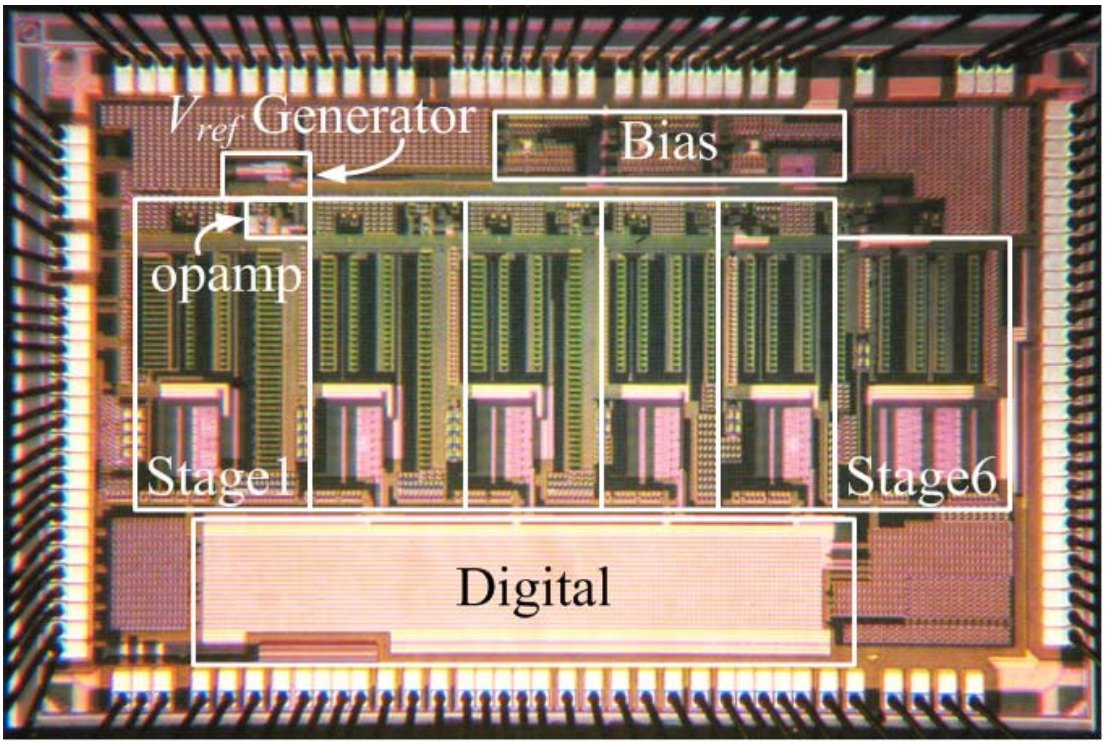

Figure 28 Die photograph.

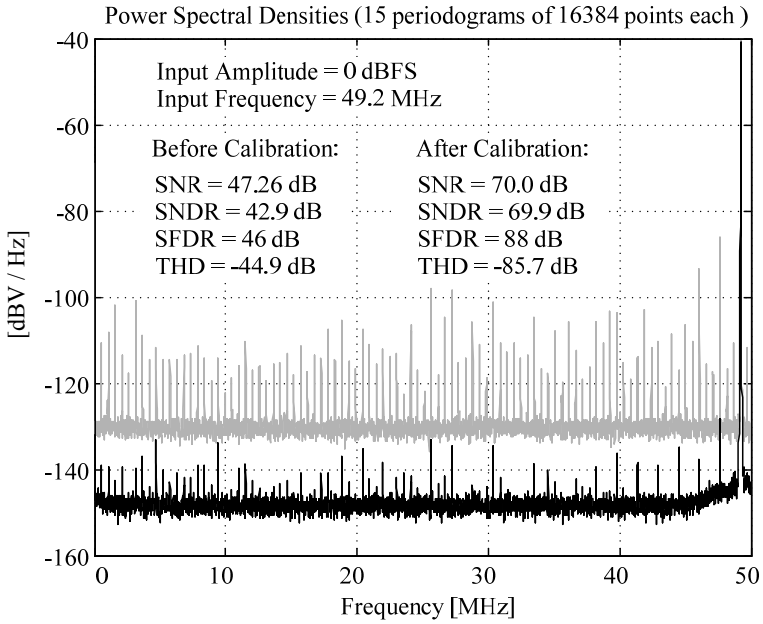

Figure 29 Measured ADC output PSD plots before and after HDC/DNC calibration.

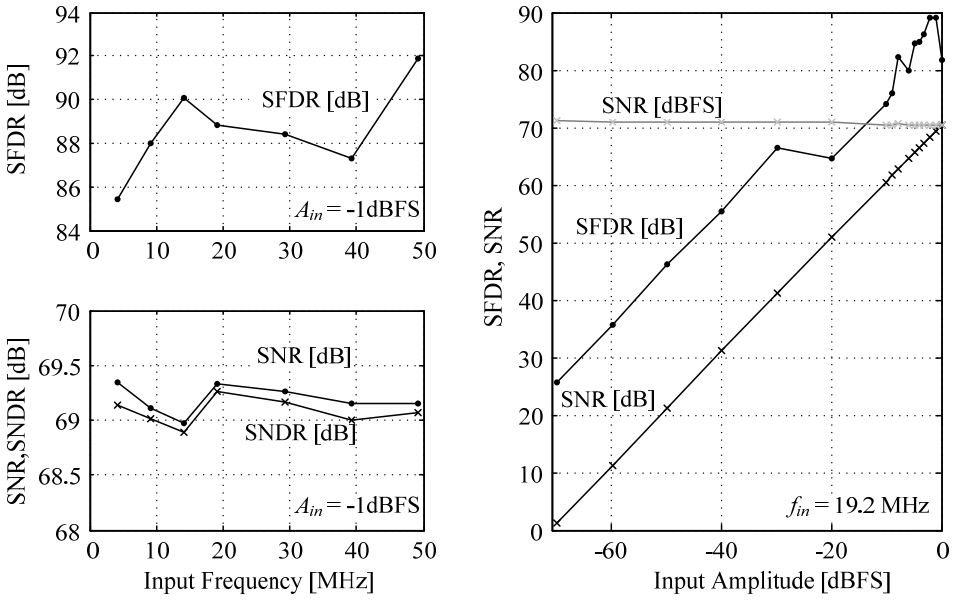

Figure 30 Measured SFDR, SNR, and SNDR versus input frequency and input amplitude.

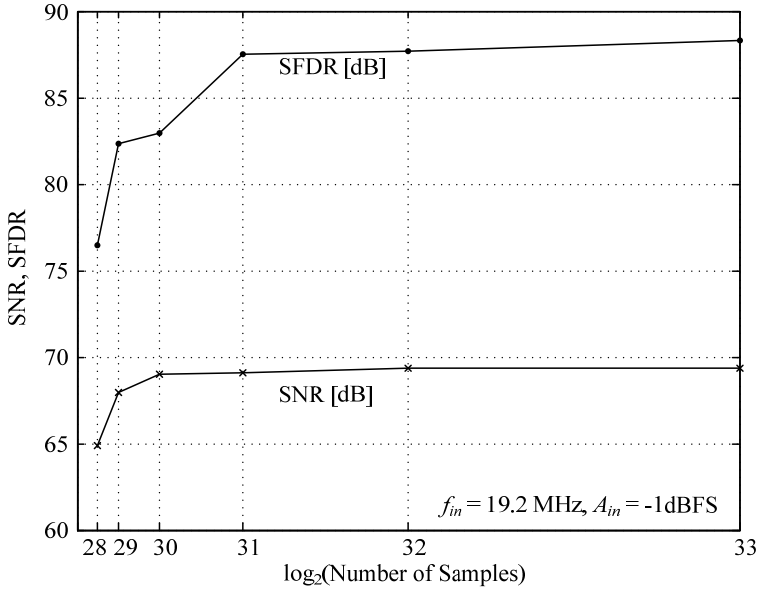

Figure 31 Measured SNR and SFDR versus number of points averaged by HDC.

TABLES

Table I ADC performance summary.

Design Details						
Technology	90 nm CMOS					
Package	56 pin QFN					
Die Size Including Pads and ESD Protection	2.15 mm × 3.35 mm					
Active Area	4 mm^2					
Digital Calibration	on-chip					
Voltage References	on-chip					
Worst Case Measured Results Over Nyquist Band for f_s = 100 MHz						
Power Supplies	**V_{DD} Test Case 1**			**V_{DD} Test Case 2**		
	V_{DD}	Power Diss.		V_{DD}	Power Diss.	
Analog	1.2 V	93 mW	130mW	1.0 V	62 mW	92 mW
Digital	1.0 V†	17 mW		0.7 V†	7 mW	
Clock Generator	1.0 V	1 mW		1.0 V	1 mW	
Clock Drivers & DEM	1.35 V‡	19 mW		1.35 V‡	22 mW	
Input and References						
Input Voltage Range	1.5 $V_{p\text{-}p}$ differential			1.25 $V_{p\text{-}p}$ differential		
Internal V_{refp} / V_{refm}	950 mV / 265 mV			775 mV / 225 mV		
Performance with HDC and DNC On						
Peak *SNR*	70 dB			68.3 dB		
SNDR at −1dBFS	68.8 dB			66.6 dB		
SFDR at −1dBFS	85 dB			75 dB		
2-tone *SFDR* at −1dBFS	86 dB			80 dB		
Maximum *INL*	3.6 LSB			3.8 LSB		
Maximum *DNL*	0.54 LSB			0.39 LSB		
Performance with HDC and DNC Off						
SNDR at −1dBFS	43.3 dB			47.3 dB		
SFDR at −1dBFS	52.3 dB			58 dB		
Performance with HDC on and DNC Off						
SNDR at −1dBFS	64.6 dB			64.3 dB		
SFDR at −1dBFS	85 dB			75 dB		

† The digital circuitry works reliably and full ADC performance is achieved provided this V_{DD} is at least 0.6V.

‡ When this V_{DD} is set to its targeted design value of 1.2V, the peak *SNDR* decreases by approximately 3dB. Although not predicted by simulations, the authors believe that the clock drivers have insufficient strength to achieve full ADC performance at 1.2V.

Table II Comparison to prior work.

Reference or Part Number	f_s (MS/s)	*SNDR* (dBFS)	*SFDR* (dB)	V_{DD} (V)	P_{tot} (mW)	*FOM*1 (pJ/step)	*FOM*2 (pJ·V/step)
[36]	75	68	76	3	314	2.04	6.12
LTC2259	80	73	90	1.8	93	0.32	0.57
AD9233	80	70.5	90	1.8	248	1.13	2.03
ADS6123	80	72.3	89	3.3	318	1.18	3.89
LTC2260	105	73	90	1.8	112	0.29	0.53
AD9233	105	70.5	90	1.8	320	1.11	2.00
ADS6124	105	72.3	84	3.3	374	1.06	3.49
[54]	250	65.9	82	1.8	150	0.37[†]	0.67
[39]	100	70	80	1.2	250	1.00	1.2
This work	100	69.8	85	1.2	130	0.52	0.62
This work	100	67.6	75	1.0	92	0.47	0.47

$$FOM1 = \frac{P_{tot}}{2^{ENOB} f_s} \quad \text{and} \quad FOM2 = FOM1 \times V_{DD} \quad \text{where} \quad ENOB = \frac{SNDR - 1.76\,\text{dB}}{6.02\,\text{dB}}$$

[†] This value is slightly different that that published in [54]. However, the lead author of [54] confirmed to us that the value published in this table is correct.