

Power Optimization of Pipelined ADCs with High-Order Digital Gain Calibration

Mohammad Taherzadeh-Sani and Anas A. Hamoui

Department of Electrical & Computer Engineering, McGill University, Montreal, Canada
 anas.hamoui@mcgill.ca

Abstract— Digital calibration techniques are widely utilized to linearize pipelined A/D converters (ADCs). However, their power dissipation can be prohibitively high, especially when high-order gain calibration is needed. For high-order gain calibration, this paper proposes a design methodology to optimize the data precision (number of bits) within the digital calibration unit. Thus, the power dissipation of the calibration unit can be minimized, without affecting the linearity of the pipelined ADC. A 90-nm FPGA synthesis of a 2nd-order digital gain-calibration unit shows that the proposed optimization methodology results in a 59% reduction in power dissipation.

I. INTRODUCTION

The linearity of a pipelined A/D converter (ADC) is primarily degraded by the linearity errors in its pipeline stages. Consider a typical pipeline stage with digital redundancy [1], as depicted in Fig. 1. In its switched-capacitor circuit implementation, the primary source of linearity error is the gain error (due to opamp nonidealities and analog component matching) in its multiplying digital-to-analog converter (MDAC) [2].

Digital calibration techniques have been successfully demonstrated to mitigate the linearity errors due to analog-circuit nonidealities in pipelined ADCs [3-7]. As described in Section II, these techniques model the gain error in the MDAC of a pipeline stage as a zero-order (constant) error [4,5] or a 2nd-order (nonlinear) error [3,6,7]. This error is then estimated and corrected for in the digital domain. Compared to zero-order gain calibration methods, 2nd-order gain calibration techniques relax the requirements on the dc gain of the opamp in the MDAC, for achieving a given linearity. Hence, these higher-order gain calibration techniques are more suited for low-power analog design, especially in nano-scale digital CMOS technologies. In these scaled technologies, high dc gains for the opamps are difficult to achieve at low power, due to the low supply voltages and the poor intrinsic gains of the MOS transistors.

However, the digital circuit implementation of high-order gain-calibration techniques can require prohibitively-high power dissipation. This is because the previously-reported algorithms for 2nd-order gain calibration [8] are computationally expensive to implement, requiring high-precision multipliers, adders, accumulators, and registers. Accordingly, this paper demonstrates how these algorithms can be simplified in order to render their digital circuit implementation more affordable (in terms of power requirements).

This paper proposes a design methodology for minimizing the power dissipation of the digital calibration unit in pipelined ADCs that incorporate 2nd-order gain calibration. It first identifies the circuit components in the digital implementation of the calibration unit that dissipate a significant amount of power. It then demonstrates that, rather than using the full data precision (the full number of bits) for all signals within these circuit components [8], the data precision of carefully-selected signals can be reduced to minimize the power dissipation of the calibration unit, without affecting the linearity of the ADC.

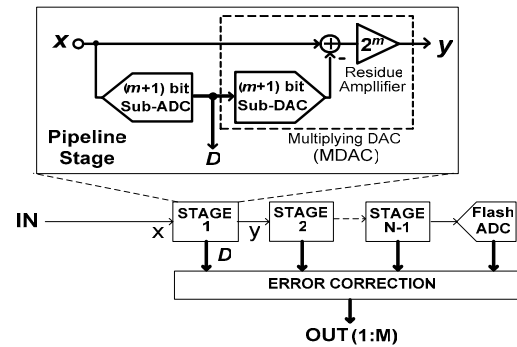


Fig. 1 An M -bit N -stage pipelined ADC. Its 1st pipeline stage has an m -bit effective resolution and a 1-bit redundancy for digital error correction.

To confirm the achievable savings in digital power dissipation using the proposed design methodology, the digital calibration unit of the pipelined ADC must be designed and implemented for various data precisions. Since custom silicon implementation of the calibration-unit design for each data precision is expensive, the calibration-unit designs for various data precisions are synthesized on a state-of-the-art 1.2-V 90-nm FPGA and their power requirements are computed and compared.

The paper outline is: Section II reviews the gain-calibration techniques for pipelined ADCs. Section III demonstrates the need for high-order gain calibration. Section IV studies the computational complexity of 2nd-order gain calibration. Section V describes the proposed design methodology for data-precision optimization in digital calibration units with 2nd-order gain calibration. Section VI demonstrates the achievable saving in power dissipation.

II. GAIN CALIBRATION TECHNIQUES

Consider the “capacitor-flip-over” MDAC in Fig. 2, which is widely utilized to realize switched-capacitor pipeline stages with a 1-bit effective resolution (Fig. 1). The MDAC output signal can be expressed as:

$$y = (1 - \delta g)(2x - DV_{ref}) \quad (1)$$

where δg represents the overall gain error due to opamp nonidealities and capacitor mismatches [4]. Here, digit D (i.e., the sub-ADC output of the 1st pipeline stage in Fig. 1) is ± 1 or 0, depending on the input signal x . In this paper, a unity reference voltage ($V_{ref} = 1$) is assumed for simplicity.

The gain error δg in equation (1) is not constant, but rather a nonlinear function of the MDAC (opamp) output signal y . This gain error can be modeled, using a k^{th} -order Taylor-series expansion, as:

$$\delta g = \delta g_0 + \delta g_2 y^2 + \delta g_4 y^4 + \dots + \delta g_k y^k \quad (2)$$

Here, a fully differential MDAC is assumed, resulting in a zero value for the odd-order coefficients in the series expansion of δg .




Fig. 2 Capacitor flip-over MDAC for a 1-bit pipeline stage (Fig. 1).

A. Zero-Order Gain Calibration

To minimize the computational complexity, most calibration techniques model the gain error δg as a constant error [4,5]:

$$\delta g = \delta g_0 \quad (3)$$

Since only errors due to δg_0 in the Taylor-series expansion in equation (2) are then digitally calibrated, such techniques are classified as zero-order gain-calibration techniques.

B. Second-Order Gain Calibration

To account for the unavoidable nonlinear gain errors in the analog circuits, 2nd-order gain calibration techniques have been proposed [1,6,7]. Here, the gain error δg is modeled as a 2nd-order nonlinear function of the MDAC output signal y :

$$\delta g = \delta g_0 + \delta g_2 y^2 \quad (4)$$

Gain errors due to δg_0 and δg_2 are then digitally calibrated.

III. IMPORTANCE OF HIGH-ORDER GAIN CALIBRATION

In a CMOS opamp, the dc gain A varies with the output voltage due to the dependency of the output resistance of a MOS transistor on its drain-source voltage. This nonlinear variation in dc gain A can be modeled as a function of the opamp output voltage y using [9]:

$$A = \begin{cases} A_0 \left(1 - \left(\frac{y}{y_{sat}} \right)^2 \right) & \text{for } y < y_{sat} \\ 0 & \text{for } y \geq y_{sat} \end{cases} \quad (5)$$

where A_0 is the maximum dc gain and y_{sat} is the output saturation voltage of the opamp [9]. Hence, using a Taylor-series expansion, the inverse of the opamp dc gain can be expressed for $y < y_{sat}$ as:

$$\frac{1}{A} = \frac{1}{A_0} \left(1 + \left(\frac{y}{y_{sat}} \right)^2 + \left(\frac{y}{y_{sat}} \right)^4 + \dots \right) \quad (6)$$

Assume that the gain error δg in a switched-capacitor MDAC is only due to the finite dc gain of its opamp. Based on negative feedback theory, the value of δg can then be expressed as [10]:

$$\delta g = \frac{1}{1 + A\beta} \approx \frac{1}{A\beta} \quad (7)$$

where β is the feedback factor of the MDAC during its charge-transfer clock phase [10]. Hence, (6) and (7) result in:

$$\delta g = \frac{1}{A_0\beta} \left(1 + \left(\frac{y}{y_{sat}} \right)^2 + \left(\frac{y}{y_{sat}} \right)^4 + \dots \right) \quad (8)$$

Accordingly, δg increases with decreasing A_0 . Furthermore, the effect of the k^{th} -order error term on the value of δg increases with $(y_{sat})^k$. In scaled digital CMOS technologies, the high-order terms have: a) a significant contribution to the total gain error δg due to the low supply voltage (and, hence, the low saturation voltage y_{sat} for the opamp); and b) the low intrinsic gain of the MOS transistors (and, hence, the low dc gain A_0 of the opamp). Thus, high-order gain calibration techniques are critical for designing high-resolution ADCs in nano-scale CMOS technologies using moderate-gain opamps.




Fig. 3. The 1st pipeline stage (with m -bit effective resolution) of the M -bit pipelined ADC in Fig. 1, with a dither signal PN added at the input of its sub-ADC.

In the following, the 2nd-order gain-calibration algorithm in [6] is summarized to illustrate the computational complexity of 2nd-order gain-calibration techniques.

IV. COMPUTATIONAL COMPLEXITY OF 2ND-ORDER GAIN CALIBRATION

Consider the 1st pipeline stage (with m -bit effective resolution) in the M -bit pipelined ADC of Fig. 1. To calibrate the gain error δg in its MDAC, a known dither signal PN (i.e., a pseudo-random sequence PN) is added at the input of its sub-ADC, as depicted in Fig. 3 [4,6]. Figure 4 illustrates the building blocks of the digital calibration unit which can be utilized for the 2nd-order gain-calibration of the pipeline stage in Fig. 3.

Assume the 1st pipeline stage in Fig. 3 has 1-bit effective resolution ($m=1$). Define:

$$OUT_{cal} = (D + Y_{cal}) / 2 \quad (9)$$

$$\bar{D} = D / 2 - PN / 4 \quad (10)$$

$$Y_{PN} = OUT_{cal} - \bar{D} \quad (11)$$

where Y is the digital representation of the MDAC output y in the 1st pipeline stage (as digitized by the subsequent pipeline stages) and Y_{cal} is the calibrated value of Y . Here, OUT_{cal} is the calibrated M -bit output word OUT of the pipelined ADC (Fig. 1) and D is the digital output of the sub-ADC in its 1st pipeline stage (Fig. 3).

Every clock cycle, the 2nd-order digital calibration unit (Fig. 4) performs the following tasks:

1. In the correction block, the MDAC output Y is corrected based on the estimated values for δg_0 and δg_2 , thereby generating Y_{cal} .
2. Signal Y_{PN} is computed using Y_{cal} , D , \bar{D} , and PN , based on equations (9)-(11).
3. In the estimation block, the values of δg_0 and δg_2 are updated, based on the linearity errors in signal Y_{PN} .

The correction and estimation blocks are described below.

A. Correction Block

The output signal y of the 1st pipeline stage is digitized as Y , by the subsequent pipeline stages (Fig. 1). Thus, Y is available in the digital domain. To digitally correct for the gain error δg , the digitized output signal Y is multiplied by $(1 + \hat{\delta g})$ in the correction block of the digital calibration unit (Fig. 4). Here, $\hat{\delta g}$ is the estimated value of the gain error δg . For 2nd-order gain calibration, the correction algorithm is

$$\hat{\delta g} = \hat{\delta g}_0 + \hat{\delta g}_2 Y^2 \quad (12)$$

$$Y_{cal} = Y \cdot (1 + \hat{\delta g}) \quad (13)$$

where Y_{cal} is the corrected (calibrated) value of Y , while $\hat{\delta g}_0$ and $\hat{\delta g}_2$ are the estimated values for the zero- and 2nd-order gain errors, respectively.

B. Estimation Block

For 2nd-order gain calibration, the values of gain errors δg_0 and δg_2 must be digitally estimated. To illustrate the complexity associated with such 2nd-order gain-error estimation, consider the iterative relations presented in [6] to estimate δg_0 and δg_2 ¹:

$$\hat{\delta g}_0(n+1) = \hat{\delta g}_0(n) + \mu_0 \cdot PN \cdot Y_{PN} \quad (14)$$

$$\hat{\delta g}_2(n+1) = \hat{\delta g}_2(n) + \mu_2 \cdot \left(E[PN \cdot Y_{PN}^3] - 3 \cdot E[PN \cdot Y_{PN}] \cdot E[Y_{PN}^2] \right) \quad (15)$$

Here, n is the iteration index, μ_0 and μ_2 are the update step sizes, and $E[\cdot]$ denotes the average over a large number of samples. It is shown in [6] that iterative relations (14) and (15) converge to the actual values of δg_0 and δg_2 .

C. Digital Circuit Complexity

To estimate the values of δg_0 and δg_2 in a 2nd-order gain-calibration technique, the iterative relations in (14) and (15) must be digitally implemented. To simplify the digital-circuit implementation of the multiplication by μ_0 and μ_2 in these relations, the values of μ_0 and μ_2 are typically selected as, respectively, 2^{-k_0} and 2^{-k_2} , where k_0 and k_2 are integer constants. Thus, these multiplications can be performed by k_0 and k_2 shifts in the multiplicand. Furthermore, since the pseudo-random signal PN is either +1 or -1, the multiplication by PN in equations (14) and (15) can also be performed by changing the sign of the multiplicand. Moreover, the averaging operation $E[\cdot]$ can be performed using an accumulator.

However, digital multipliers are still required to compute Y_{PN}^2 , Y_{PN}^3 , and $E[Y_{PN}^2] \cdot E[PN \cdot Y_{PN}^3]$ in iterative relation (15) for the 2nd-order error $\hat{\delta g}_2$. These digital multiplications are expensive to implement, in terms of digital power requirements.

In the following, a design methodology is proposed to significantly reduce the complexity of the multiplications and accumulations in iterative relations (14) and (15), thereby minimizing the power dissipation required in the digital circuit implementation of 2nd-order gain calibration algorithms.

V. DESIGN METHODOLOGY FOR POWER OPTIMIZATION

A. Data Precision of Y_{PN}

In the estimation block of a 2nd-order digital-calibration unit (Fig. 4), the complexity of the multipliers and accumulators (used to implement iterative relations (14) and (15)) mostly depends on the data precision (number of bits) of digital signal Y_{PN} . In an M -bit pipelined ADC (Fig. 1), the full precision of Y_{PN} is $(M-m)$ bits, assuming an m -bit effective resolution in the 1st pipeline stage. Reducing the data precision of Y_{PN} reduces the computational complexity and, hence, the required power dissipation for the digital circuits.

Reducing the precision of Y_{PN} by truncating some of its least significant bits (LSBs) is equivalent to adding some quantization noise to Y_{PN} . Assuming this quantization noise to be uncorrelated with signal Y_{PN} and pseudo-random (dither) signal PN , it will be filtered out by the accumulators used to implement iterative relations (14) and (15) [4]. Therefore, a truncated Y_{PN} (i.e., a Y_{PN} with some of its LSBs removed) can be utilized in the estimation block instead of Y_{PN} .

However, there is a weak correlation between Y_{PN} and the

¹ Our goal here is to only illustrate the digital complexity of 2nd-order gain-error estimation, in order to later demonstrate (in Section V) how this digital complexity and, hence, the required power dissipation can be significantly reduced. The reader is referred to [6] for detailed derivations of relations (14) and (15).




Fig. 4. Digital calibration unit for the 2nd-order gain calibration of the pipeline stage in Fig. 3.

quantization noise associated with a truncated Y_{PN} [4]. This affects the convergence of equations (14) and (15). Depending on the amount of quantization noise and its correlation with the input signal, the accuracy of the gain calibration and, hence, the total resolution of the ADC can be reduced. This limits the maximum number of LSBs which can be truncated from Y_{PN} , without affecting the accuracy of the gain calibration. Due to the complex behavior of the quantization noise in the calibration unit and its correlation with the input signal, no mathematical expressions were derived for the optimal data precision. Rather, in Section VI, the optimal data precision (number of bits) for Y_{PN} will be determined through system-level simulations.

The power-saving technique proposed above for the estimation block in 2nd-order gain calibration can also be utilized in zero-order gain calibration, since both zero-order [4,5] and 2nd-order [6] calibration methods essentially use iterative relation (14) to estimate the gain error in their respective estimation blocks.

B. Data Precision of Y

In an M -bit pipelined ADC (Fig. 1), the full precision of Y is $(M-m)$ bits, assuming an m -bit effective resolution in its 1st pipeline stage. The value of Y used to compute Y_{cal} in equation (13) and, subsequently, OUT_{cal} in equation (9) should always be implemented at its full precision, since its precision directly impacts the accuracy of the pipelined ADC output.

However, the accuracy of the estimated gain error $\hat{\delta g}$ used in the correction block of a 2nd-order digital calibration unit (Fig. 4) is determined by the accuracy of $\hat{\delta g}_0$, $\hat{\delta g}_2$, and Y^2 , as per equation (12). Since $\hat{\delta g}_0$ and $\hat{\delta g}_2$ are the outputs of two accumulators, their values change moderately around their mean values, owing to the randomness of the accumulated signal. Therefore, having an accurate Y^2 does not significantly improve the accuracy of the estimated $\hat{\delta g}$, since the other components (i.e., $\hat{\delta g}_0$ and $\hat{\delta g}_2$) in equation (12) are not very accurately estimated. Hence, the data precision (number of bits) of Y can also be reduced when computing $\hat{\delta g}$, without affecting the accuracy of the gain calibration. In Section VI, the optimum data precision for Y will be determined through system-level simulation.

The power-saving technique proposed above for the correction block in background digital calibration can also be utilized in foreground digital calibration, since both background [1,4-6] and foreground [7] calibration methods use similar correction blocks.

VI. FPGA SYNTHESIS RESULTS

To estimate the achievable savings in digital power dissipation, the 2nd-order gain-calibration method in [6] is synthesized on an FPGA (EP2C5T144C6, Cyclone II, 90 nm, 1.2 V), with and without following the proposed design methodology for power optimization. To measure the linearity of the pipelined ADC, behavioral simulations of the ADC and its digital calibration unit are performed in SIMULINK [10].




Fig. 5 SNDR of the calibrated pipelined ADC versus :
 (a) no. of bits in Y , with Y_{PN} having its full precision (13 bits)
 (b) no. of bits in Y_{PN} , with Y having its full precision (13 bits)
 (c) no. of bits in Y and Y_{PN}

The pipelined ADC is partitioned into 14 pipeline stages ($M=14$), each with a 1-bit effective resolution ($m=1$). In the SIMULINK behavioral simulations, the input-output transfer function of the MDAC in each pipeline stage is modeled as in equation (1), with its gain error modeled as a 2nd-order error based on equation (4). Only the MDAC in the 1st pipeline stage is assumed to have gain errors. Further, only gain errors due to the nonlinear finite dc gain of its opamp are modeled as in equation (5), assuming a maximum dc gain $A_0=100$ V/V and an output saturation voltage $Y_{sat}=1.414$ V. Hence, considering only a 2nd-order expression for the gain error δg as in equation (4), this results in

$$\delta g = 0.02 + 0.01y^2 \quad (16)$$

for a 1.5-bit MDAC with a feedback factor $\beta=0.5$ (i.e., neglecting any capacitor mismatch or parasitic capacitors). A 0.9-V sinusoidal signal is applied at the ADC input.

The values of step sizes μ_0 and μ_2 in the iterative relations (14) and (15) are set to 2^{-23} and 2^{-17} , respectively. On the FPGA, the averaging in (15) is realized using low-pass filters. Each low-pass filter is implemented using an accumulator:

$$E[x](n+1) = E[x](n) + \mu_e \cdot (x - E[x](n)) \quad (17)$$

where n is the iteration index and $\mu_e=2^{-19}$.

The 2nd-order digital calibration unit (Fig. 4) is synthesized on the FPGA and its *dynamic* power dissipation is estimated using Quartus II. The clock frequency of the FPGA is set at 100 MHz. A random signal Y is applied at the input of the calibration unit (Fig. 4) to estimate its worst-case *dynamic* power dissipation. To estimate its *static* power dissipation, an estimate of the ratio of the calibration-unit area to the total FPGA area is calculated and then multiplied by the total static power dissipation of the FPGA (which is available in its data sheet).

Figures 5 and 6 show, respectively, the SNDR at the ADC output and the power dissipation in the digital calibration unit, versus:
 (a) data precision of Y (with Y_{PN} having its full 13-bit precision);
 (b) data precision of Y_{PN} (with Y having its full 13-bit precision);
 (c) data precisions of Y and Y_{PN} .

Based on Fig. 5, the data precisions of Y and Y_{PN} can be reduced from full 13-bit precision to 7 bits, while still achieving an SNDR > 80dB (13 bits). Based on Fig. 6, such data-precision reduction in Y and Y_{PN} (from 13 to 7 bits) saves 59% of the total power dissipation in the calibration unit.

Note that, the reported power-dissipation values are extracted using an automated compilation of the VHDL code for the gain calibration algorithm. With a custom silicon design of the digital calibration unit, this power dissipation can be further reduced [8], making high-order gain calibration techniques even more practical in high-resolution ADCs.




Fig. 6 Power consumption of the digital calibration unit versus:
 (a) no. of bits in Y , with Y_{PN} having its full precision (13 bits)
 (b) no. of bits in Y_{PN} , with Y having its full precision (13 bits)
 (c) no. of bits in Y and Y_{PN} .

VII. CONCLUSION

High-order gain calibration is required to achieve high linearity in pipelined ADCs designed using low-gain opamps in low-voltage digital scaled CMOS technologies. The data precision (number of bits) within the digital calibration unit of the pipelined ADC can be optimized to minimize the power dissipation, without affecting the linearity of the pipelined ADC. The proposed design methodology for power optimization has been demonstrated for 2nd-order gain calibration. Results confirm that this design methodology can significantly reduce the power dissipation of the digital calibration unit in a pipelined ADC that incorporates 2nd-order gain calibration.

VIII. ACKNOWLEDGMENT

This work was supported by the System-on-Chip Research Network (SOCRN), CMC Microsystems, Canada and the Fonds Québécois de la Recherche sur la Nature & les Technologies (FQRNT), Québec, Canada.

IX. REFERENCES

- [1] S. H. Lewis and P. R. Gray, "A pipelined 5 Msample/s 9-b analog-to-digital converter," *IEEE Journal of Solid-State Circuits*, vol. SC-22, pp. 954-961, Dec. 1987.
- [2] M. Taherzadeh-Sani and A. A. Hamoui, "A digital background calibration technique for capacitor mismatch errors in pipelined ADCs," *IEEE Trans. Circuits Syst.-II*, pp. 966-970, Sept. 2006.
- [3] B. Murmann and B.E. Boser, "A 12-bit 75-MS/s pipelined ADC using open-loop residue amplification," *IEEE Journal of Solid-State Circuits*, Vol. 38, no. 12, pp. 2040 – 2050, Dec. 2003.
- [4] J. Li, G. Ahn, D. Chang, and U. Moon, "A 0.9-V 12-mW 5-MSPS algorithmic ADC with 77-dB SFDR," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 4, pp. 960-969, Apr. 2005.
- [5] E. Siragusa and I. Galton, "Gain Error Correction Technique for Pipelined Analog-to-Digital Converters," *IEE Elect. Let.*, pp. 617-618, Mar. 2000.
- [6] J. Keane, P. Hurst, and S. Lewis, "Background interstage gain calibration technique for pipelined ADCs," *IEEE Trans. Circuits Syst. I*, vol. 52, no. 1, pp. 32-43, Jan. 2005.
- [7] C. Grace, P. J. Hurst, and 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] C. Grace, "Digital calibration of interstage gain errors and signal dependent variations in pipelined analog-to-digital converters," Doctoral thesis, Univ. California, Davis, 2004.
- [9] A. A. Hamoui and K. Martin, *Delta-Sigma Data Converters in Low-Voltage CMOS for Broadband Digital Communication*. Dordrecht, Netherlands: Springer, 2008.
- [10] A. A. Hamoui *et al.*, "Behavioral modeling of opamp gain and dynamic effects for power optimization of delta-sigma modulators and pipelined ADCs," in *IEEE International Symposium on Low Power Electronics and Design (ISLPED)*, pp. 330-333, Oct. 2006.