An empirical approach to finding energy efficient ADC architectures

Bengt E. Jonsson

ADMS Design AB, Gåsbackavägen 33, SE-820 60 DELSBO, Sweden
phone: +46 70 622 0005
e-mail: bengt.e.jonsson@admsdesign.com

Abstract—The problem of selecting an optimally efficient ADC architecture at different resolutions is treated using a mainly empirical approach. By analyzing a large amount of measured performance data reported in the literature, the power efficiency of different ADC architectures is investigated. An efficiency hierarchy of ADC architectures is identified, and the low-power enabling features in state-of-the-art designs are summarized. The work shows that there are significant differences between architectures, and also the feasibility of an empirical approach to design optimization.

I. Introduction

The quest for energy efficient solutions is simultaneously driven by multiple and largely independent interests such as the *consumer* craving long battery life for high-performance portable gadgets, communication infrastructure *operators* noting that a significant part of their budget is energy-related, and *governments* implementing legislation and directives for a sustainable society. In addition to that, *specific applications* such as medical implants and wireless sensor nodes put particularly stringent demands on energy-optimized solutions. Noting that *A/D-converters* (ADCs) are typically present in all the systems mentioned above, it is essential to develop low-energy ADC solutions for a broad range of applications and specifications. Several contributions have presented theoretical *power dissipation limits* for A/D-converters [1]-[6], and *empirical performance trends* have been analyzed in [7]-[12]. Some of the theoretical contributions assume a particular architecture such as *pipeline* [3]-[4], while others treat more than one architecture [1], [5]-[6]. The theoretical derivations in [6]-[7], [9]-[10] are compared against relatively large sets of reported measurement data from publications.

The purpose of this work is to illustrate an empirical approach to design optimization and to give an overview of scientific ADCs from an energy viewpoint. Additionally, an *efficiency hierarchy of architectures* will be identified and the *low-power enablers* in the current state-of-the-art summarized. The results and discussions aim to support the selection of energy-optimal ADC architectures based on the performance of prior art. While real-world design projects will have many other design constraints, this paper illustrates a more general search for the overall most efficient architectures at different levels of resolution. A key assumption in this paper is that *empirical design-optimization* (EDO) can be done if a sufficiently large set of design and performance data is available. By systematically analyzing a large number of reported attempts, conclusions can be made regarding optimal architecture and design choices under a particular set of design targets. An optimal data set for EDO contains *all* information about all ADCs ever implemented. As an approximation of the ideal set, the experimental results and measured performance reported in more than 1400 scientific papers published between 1974 and March 2010 have been used in this work. The underlying source data is the same as in [11], and represents a near-exhaustive survey of all scientific ADC implementations ever reported. To the best of the author's knowledge, ADC efficiency has not previously been studied with such large amounts of empirical data.

II. Measures of energy efficiency

Rather than looking for the lowest absolute power dissipation, the present work will focus on the amount of *energy* per converted sample,

$$E = \frac{P}{f_s} \tag{1}$$

By using E, power dissipation is normalized to sampling rate, and E was observed also in [6], [8] and [10]. The definition of E naturally aligns with the concept of charging and discharging ADC-internal capacitors and parasitics at a *sampling rate* of f_s, even if there are architectures where the main power consumption may be from time-continuous or static current. Whereas E normalizes to sampling rate, it does not handle the fact that power dissipation also is a function of nominal resolution N, or the *effective number of bits* (ENOB) calculated from *signal-to-noise-and-distortion ratio* (SNDR) as

$$ENOB = \frac{SNDR_{dB} - 1.76}{6.02} \tag{2}$$

A higher ADC resolution typically requires larger capacitors and more hardware, thus more power dissipation. In order to handle that dependency, a *figure-of-merit* (FOM) that includes N or ENOB could be used. Two figures-of-merit commonly used in the literature are F_1 and F_2:

$$F_1 = \frac{P}{2^{ENOB} f_s} = \frac{E}{2^{ENOB}} \tag{3}$$

$$F_2 = \frac{P}{2^{2ENOB} f_s} = \frac{E}{2^{2ENOB}} \tag{4}$$

As seen in (3) and (4), these FOM simply normalize E to the relative conversion error *amplitude* (F_1) or *power* (F_2). It is however difficult to derive a FOM that is ENOB-independent, and neither F_1 nor F_2 are. The lesser used F_2 is biased towards high-resolution ADCs limited by thermal noise, and it can be shown that F_1 has a sweet spot around the transition between thermal-noise-limited ADCs and other ADCs – currently around ENOB = 9. Thus it will be necessary to simultaneously consider ENOB even when using one of these FOM to represent efficiency. The usefulness of F_1 and F_2 over E is therefore negligible, and efficiency will be viewed as the two-dimensional vector $\{E, ENOB\}$ in this work. A positive side effect is that there's no need to decide which FOM is the most "sound".

III. Energy contributors and limits

The total ADC energy E_{tot} can be split into I/O energy E_{IO}, reference energy E_{ref}, sample-and-hold energy $E_{S/H}$, ADC core energy E_{core}, and the input energy E_{in} delivered by the source. A distinct separation may be difficult to maintain. In ADCs without internal buffer, E_{in} and $E_{S/H}$ could be treated as one, and in ladderless flash ADCs such as [13], there is no E_{ref}. In this work, the reported on-chip E_{tot} will be observed, while theoretical limits for E_{IO} and $E_{S/H}$ (E_{in}) will be used as illustrative reference points.

A. I/O energy

The term "I/O energy" refers to the energy loss in the *digital* input/output circuitry. A minimum E_{IO}-estimate can be derived under the assumption of N parallel CMOS outputs driving a capacitive load C_L. The energy stored in a capacitance C charged to V is

$$E = \frac{CV^2}{2} \tag{5}$$

and the minimum total energy supplied from the source (last inverter) is CV^2 due to the energy loss during charging. Additional energy lost due to tapered drivers or inverter transient current is not included. The average E_{IO} can be estimated from the average bit toggle probability α and output logic swing V_{DD} as

$$E_{IO} = \frac{P_{IO}}{f_s} = N \cdot \frac{\alpha}{2} \cdot C_L V_{DD}^2 \tag{6}$$

In this work $\alpha = 50\%$ (random data) was used. Furthermore, a load capacitance $C_L = 7$ pF has been assumed, with 2 pF attributed to the output pad, and 5 pF to the board interconnect and receiving circuit.

B. Minimum sampling energy

Assume that the ADC input signal is sampled by a capacitor C_s that is reset between each new sample. With a full-scale peak-to-peak input range V_{FS}, the absolute minimum energy E_s required from the source to charge each sample v_m through a resistive switch is

$$E_s = C_s v_m^2 \quad \Rightarrow \quad \overline{E}_s = \frac{C_s V_{FS}^2}{8} \tag{7}$$

If ENOB is noise-limited, and half of the noise power is assigned to the kT/C-noise sampled onto C_s, the full-scale *signal-to-noise ratio* (SNR) over C_s must be

$$SNR_{dB} = 6.02 \cdot ENOB + 1.76 + 3 \tag{8}$$

Thus, the linear SNR is

$$SNR = \frac{\overline{v_{signal}^2}}{v_{noise}^2} = \frac{V_{FS}^2 C_s}{8kT} = 10^{\frac{6.02\,ENOB+1.76+3}{10}} = 3 \times 2^{2\,ENOB} \quad \Rightarrow \quad C_s = \frac{24kT}{V_{FS}^2} 2^{2\,ENOB} \tag{9}$$

which gives an ENOB-defined minimum sampling energy

$$\overline{E}_s = 3kT \cdot 2^{2\,ENOB} \tag{10}$$

The expression in (10) is *the absolute minimum energy required from the source* under the assumptions above and does not include any further loss due to amplifier efficiency, additional circuits such as clock drivers, etc. Assuming more realistic operating conditions, a 16 times larger value of E_s was derived in [6].

IV. Identifying Efficient Architectures

A first attempt to empirically identify efficient architectures is to look at the energy landscape divided by architecture. A scatter plot of the reported energy/sample in each of the 1400 papers is shown vs. ENOB in Fig. 1. Only data points that report "true ENOB" (SNDR or SNR & THD) have been included, in order to not give unfair advantage to papers reporting SNR-only performance. The separation by architecture reveals significant differences between architectures, which will be discussed below. Consider first the shape of the energy landscape. There are two distinct borders, which will be referred to as the *thermal slope* and the *low-resolution plateau* in this work. For ENOB ≥ 9 the energy per sample seems limited by thermal noise since it quadruples for every bit of effective resolution. The current state-of-the-art almost exactly follows the $E = 2^{2(ENOB-9)}$ **pJ** curve that is included in the plot. At lower resolutions the state-of-the-art levels out to **~1pJ/sample independent of ENOB** down to 2.5-b, where E starts to degrade again. From Fig. 1 it can also be seen that ADCs below 9-b ENOB are comparatively inefficient relative to their $\{ENOB, f_s\}$ performance, and that going below 5-b ENOB does not appear to reduce E at all. Although it is slightly unintuitive, the large body of scientific data suggests that going below 9-b ENOB only gives minor improvements in power efficiency, if any. Hence, *system-level* power optimization should be considered when ADC requirements are lower than 9-b. Rather than trying to optimize the ADC, it can be kept at 8–9-b resolution if that leads to relaxed requirements in other parts of the signal chain.

A. The efficiency hierarchy

A closer inspection of the data in Fig. 1 reveals that the low-resolution energy plateau is defined by a handful of designs that are detached from the main scatter by a factor of 5–10. These designs will be analyzed separately in the following subsection – both because they define state-of-the-art and because several of them are rather special cases of the architectures they represent. The efficiency hierarchy of architectures was therefore identified by analyzing the overall picture – including the main scatter. It was seen that the Δ-∑ *modulator* (DSM), SAR, *flash* and *comparator-based asynchronous binary search* (CABS) architectures are at the top of the efficiency hierarchy. Over a very large number of reported attempts, these architectures have reported a
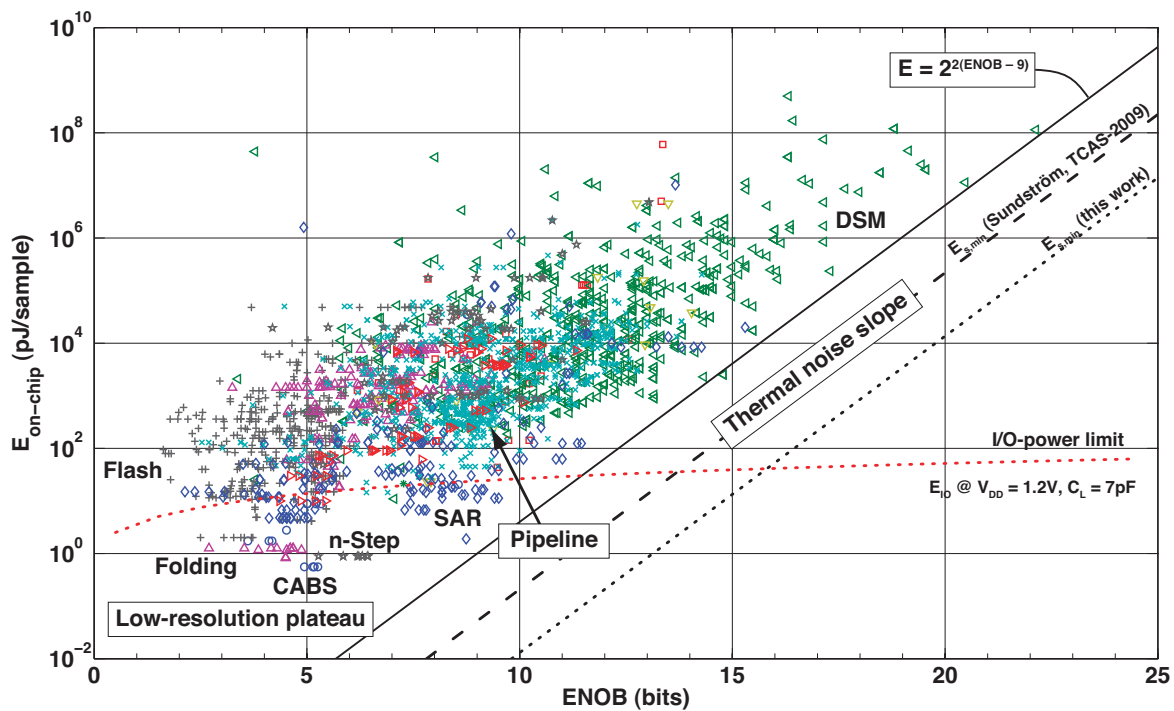
Fig. 1. On-chip energy per sample vs. ENOB for different architectures: Asynchronous (○), Cyclic (□), DSM (◁),
Flash (+), Folding (△), Pipeline (×), SAR (◇), Subranging (▷), n-Slope (✳), n-Step (★) and other (▽).

Fig. 2. Efficiency hierarchy of ADC architectures according to scientifically reported measurement data.

significantly lower energy/sample than the next level in the hierarchy. SAR appears to be the architecture that is most energy efficient over a broad range of resolutions, both along the thermal-noise slope up to 15-b, and along the low-resolution plateau as far down as 2-b. CABS and flash ADCs are only efficient on the plateau, and DSM have state-of-the-art efficiency only along the thermal slope for ENOB ≥ 13-b. ΔΣ ADCs have no competition above 15-b ENOB. The best *pipeline* and *cyclic* ADCs are approximately one order of magnitude less energy-efficient than SAR/DSM along the thermal slope, and thus define the second level in the hierarchy. All other major architectures require at least 5–10 times more energy/sample than pipeline and cyclic. On the plateau, the second level is more crowded. *Subranging*, *folding*, *pipeline*, *cyclic*, *n-step*, and even DSM implementations [14] have all shown equivalent results within 5–10 times the energy of SAR. The energy efficiency hierarchy is shown in Fig. 2. Only the first two levels are shown.

B. State-of-the-art on the low-resolution plateau

The seven ADCs that define the low-resolution energy plateau are listed in order of efficiency in Table I. By reviewing the low-power enablers for each implementation, it becomes evident that common features include a *dynamic comparator* – usually with *trip-point calibration* – that is used in a *ladderless* architecture when applicable. A somewhat less obvious result is that three of the top four implementations used a 1-b MSB stage, effectively detecting the sign before the magnitude is converted by a multi-bit LSB-stage. *Asynchronous-binary-search* architectures are dominant – either stand-alone, or as the LSB-stage. Other architectures are *flash* and SAR. The two designs labeled "*folding*" can also be seen as a sign-detector followed by a flash [16] or pipelined-ABS [17] LSB-stage, but technically they implement a folding architecture.

Ref	Architecture	E [pJ]	@ ENOB$_{max}$	Low-power enablers
[15]	CABS	0.56	5.3	CABS architecture; Ladderless; Dynamic comparator; Trip-point calibration
[16]	Folding	0.85	4.5	1-b MSB stage; Ladderless Flash LSB-stage; Dynamic comparator; Trip-point calibration
		0.82	4.5	
		1.26	4.7	
[15]	Two-step	0.89	6.4	1-b MSB stage; Ladderless CABS LSB-stage; Dynamic comparator; Trip-point calibration
[17]	Folding	1.2	4.9	1-b MSB stage; Ladderless PABS LSB-stage; Dynamic comparator/amplifier; Trip-point calibration
[18]	CABS	1.7	4.2	CABS architecture
[19]	SAR	1.9	8.7	Multi-step capacitor charging; Dynamic comparator; C-DAC
[13]	Flash	2.0	3.7	Ladderless; Dynamic comparator; Trip-point calibration

Table I. State-of-the-art designs on the low-resolution plateau – in order of E.

Ref	Architecture	x_E	@ ENOB	Low-power enablers
[20]	CT-DSM	253	14.8	FF-architecture; Large FSR$_{in}$; Novel excess-delay compensation => low-power op-amps
[21]	DT-DSM	434	22.1	Multiple FB-loops architecture; Architectural optimization => op-amps can slew High OSR spreads noise leak => low-gain op-amps
			20.5	
[22]	CT-DSM	436	14.3	FF-architecture; "Assisted op-amp"
[23]	DT-DSM	663	15.5	FF-architecture; single-comparator, tracking M-bit quantizer
[24]	DT-DSM	728	17.3	FF-architecture; Class-AB op-amps
[19]	SAR	832	8.7	Multi-step capacitor charging
[25]	DT-DSM	861	13.2	FF-architecture; Inverter-based
[26]	SAR	988	15.3	–

Table II. State-of-the-art designs on the thermal slope – in order of x_E.

C. State-of-the-art on the thermal slope

The eight implementations that currently define the thermal slope are listed in order of efficiency in Table II. Since E is a function of ENOB, the efficiency is described by the *energy ratio*

$$x_E = \frac{E}{E_s} \tag{11}$$

where E_s is the theoretical minimum energy according to (10), and E is the actual on-chip energy dissipation reported. The converters in Table II are those that have an $x_E \leq 1000$. It means that their *total* on-chip energy/sample is only 250–1000 times above the (unrealistic) theoretical minimum for just the sampling. Using the more realistic, and 16 times larger value for E_s derived in [6], the implementations in Table II are only 16–62 times above the minimum *sampling* energy according to [6] – theoretical *conversion* energy not included. Looking at the low-power enablers for each design yields that minimizing the power dissipated in op-amps is a key to efficient DSM implementations. Five out of six DSM use a *feed-forward* filter topology to reduce the integrator output swing, thus relaxing slew-rate requirements. The state-of-the-art circuits were realized using various low-power amplifier solutions, including *class-AB*, *assisted op-amp*, and *inverters*. Regarding the two SAR converters, it appears that the 16-b design in [26] relies simply on good design and the efficiency of the architecture itself, while the addition of multi-step charging allows the SAR in [19] to extend the thermal slope boundary down to 9-b ENOB. Note that the same 10-b SAR defines state-of-the-art *both* on the plateau and the thermal slope, and thus represents a highly efficient ADC implementation.

V. Conclusion

ADC power efficiency was empirically investigated using a near-exhaustive set of scientifically reported performance. It was found that SAR, DSM, flash, and asynchronous-binary-search architectures are the most energy efficient. Low energy consumption requires architectural as well as circuit-level optimization. Efficient DSM ADCs tend to use feed-forward architectures, while the SAR architecture seems to be inherently suitable for low-energy. At lower resolutions, efficiency is achieved by avoiding static power dissipation. Ladderless architectures using dynamic comparators yield low power. Introducing trip-point calibration allows comparator devices to be scaled below mismatch requirements, thus reducing dynamic power. While the results of the investigation are largely as expected, the work also summarize the efficiency of nearly all scientific ADCs ever reported, and gives a detailed review of the current state-of-the-art. The paper also illustrates the feasibility of finding optimal architectures using a mostly empirical approach. Future work may show how to apply empirical design optimization once the architecture is selected.

References

[1]   F. Maloberti, F. Francesconi, P. Malcovati, and O. J. A. P. Nys, "Design considerations on low-voltage low-power data converters," *IEEE Trans. Circuits and Systems, pt. I*, vol. 42, no. 11, pp. 853–863, Nov. 1995.
[2]   P. B. Kenington, and L. Astier, "Power consumption of A/D converters for software radio applications". *IEEE Trans. Vehicular Tech.*, vol. 49, no. 2, pp. 643–650, March 2000.
[3]   S. Kawahito, "Low-Power Design of Pipeline A/D Converters," *Proc. of IEEE Custom Integrated Circ. Conf. (CICC)*, San Jose, California, USA, pp. 505–512, Sept., 2006.
[4]   C. Svensson, S. Andersson, and P. Bogner, "On the power consumption of analog to digital converters," *Proc. of NORCHIP*, Linköping, Sweden, pp. 49–52, Nov., 2006.
[5]   B. Bechen, T. v. d. Boom, D. Weiler, and B. J. Hosticka, "Theoretical and practical minimum of the power consumption of 3 ADCs in SC technique," *Proc. of Eur. Conf. Circ. Theory and Design (ECCTD)*, Seville, Spain, pp. 444–447, Aug., 2007.
[6]   T. Sundström, B. Murmann, and C. Svensson, "Power dissipation bounds for high-speed Nyquist analog-to-digital converters," *IEEE Trans. Circuits and Systems, pt. I*, vol. 56, no. 3, pp. 509–518, Mar. 2009.
[7]   R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Selected Areas in Communications*, no. 4, pp. 539–550, Apr. 1999.
[8]   K. G. Merkel, and A. L. Wilson, "A survey of high performance analog-to-digital converters for defense space applications," in *Proc. IEEE Aerospace Conf.*, Big Sky, Montana, Mar. 2003, vol. 5, pp. 2415–2427.
[9]   B. Le, T. W. Rondeau, J. H. Reed, and C. W. Bostian, "Analog-to-digital converters [A review of the past, present, and future]," *IEEE Signal Processing Magazine*, pp. 69–77, Nov. 2005.
[10] B. Murmann, "A/D converter trends: Power dissipation, scaling and digitally assisted architectures," *Proc. of IEEE Custom Integrated Circ. Conf. (CICC)*, San Jose, California, USA, pp. 105–112, Sept., 2008.
[11] B. E. Jonsson, "A survey of A/D-converter performance evolution," *Proc. of IEEE Int. Conf. Electronics Circ. Syst. (ICECS)*, Athens, Greece, pp. 768–771, Dec., 2010.
[12] B. E. Jonsson, "On CMOS scaling and A/D-converter performance," *Proc. of NORCHIP*, Tampere, Finland, pp. 1–4, Nov. 2010.
[13] G. Van der Plas, S. Decoutere, and S. Donnay, "A 0.16pJ/conversion-step 2.5mW 1.25GS/s 4b ADC in a 90nm Digital CMOS process," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 2310–2312, Feb., 2006.
[14] U. Wismar, D. Wisland, and P. Andreani, "A 0.2V 0.44μW 20 kHz analog to digital ∑Δ modulator with 57 fJ/conversion FoM," Proc. of *Eur. Solid-State Circ. Conf. (ESSCIRC)*, Montreux, Switzerland, pp. 187–190, Sept., 2006.
[15] G. Van der Plas and B. Verbruggen, "A 150 MS/s 133 μW 7 bit ADC in 90 nm digital CMOS," *IEEE J. Solid-State Circuits*, Vol. 43, pp. 2631–2640, Dec., 2008.
[16] B. Verbruggen, J. Craninckx, M. Kujik, P. Wambacq, and G. Van der Plas, "A 2.2mW 5b 1.75GS/s folding flash ADC in 90nm digital CMOS," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 252–253, Feb., 2008.
[17] B. Verbruggen, J. Craninckx, M. Kujik, P. Wambacq, and G. Van der Plas, "A 2.6mW 6b 2.2GS/s 4-times interleaved fully dynamic pipelined ADC in 40nm digital CMOS," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 296–297, Feb., 2010.
[18] Y.-Z. Lin, S.-J. Chang, Y.-T. Liu, C.-C. Liu, and G.-Y. Huang, "A 5b 800MS/s 2mW Asynchronous Binary-Search ADC in 65nm CMOS," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 80–81, Feb., 2009.
[19] M. van Elzakker, E. van Tuijl, P. Geraedts, D. Schinkel, E. Klumperink, and B. Nauta, "A 1.9μW 4.4fJ/Conversion-step 10b 1MS/s charge-redistribution ADC," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, pp. 244–245, Feb., 2008.
[20] S. Pavan, N. Krishnapura, R. Pandarinathan, and P. Sankar, "A 90μW 15-bit ΔΣ ADC for digital audio," *Proc. of Eur. Solid-State Circ. Conf. (ESSCIRC)*, Munich, Germany, pp. 198–201, Sept., 2007.
[21] R. Naiknaware, and T. Fiez, "142dB ΔΣ ADC with a 100nV LSB in a 3V CMOS process," *Proc. of IEEE Custom Integrated Circ. Conf. (CICC)*, Orlando, USA, pp. 5–8, May, 2000.
[22] S. Pavan, and P. Sankar, "A 110μW Single Bit Audio Continuous-time Oversampled Converter with 92.5 dB Dynamic Range," *Proc. of Eur. Solid-State Circ. Conf. (ESSCIRC)*, Athens, Greece, pp. 320–323, Sept., 2009.
[23] H. Park, K. Nam, D. K. Su, K. Vleugels, and B. A. Wooley, "A 0.7-V 100-dB 870-μW Digital Audio ΣΔ Modulator," *Symp. VLSI Circ. Digest of Technical Papers*, Honolulu, USA, pp. 178–179, June, 2008.
[24] J.-Y. Wu, Z. Zhang, R. Subramoniam, and F. Maloberti, "A 107.4 dB SNR multi-bit sigma delta ADC with 1-PPM THD at 0.12 dB from full scale input," *IEEE J. Solid-State Circuits*, Vol. 44, pp. 3060–3066, Nov., 2009.
[25] Y. Chae, I. Lee, and G. Han, "A 0.7V 36μW 85dB-DR audio Δ∑ modulator using class-C inverter," *Proc. of IEEE Solid-State Circ. Conf. (ISSCC)*, San Francisco, California, Vol. i, pp. 490–491, Feb., 2008.
[26] K. Y. Leung, K. Leung, and D. R. Holberg, "A dual low power ½ LSB INL 16b/1Msample/s SAR A/D converter with on-chip microcontroller," *Proc. of IEEE Asian Solid-State Circ. Conf. (ASSCC)*, Hangzhou, China, pp. 51–54, Nov., 2006.