A/D Converter Trends: Power Dissipation, Scaling and Digitally Assisted Architectures

(Invited Paper)

B. Murmann
Stanford University, Stanford, CA

***Abstract*— This paper summarizes recent trends in the area of low-power A/D conversion. Survey data collected over the past eleven years indicates that the power efficiency of ADCs has improved on average by a factor of two every two years. A closer inspection on the impact of technology scaling is presented to explain the observed trend in the context of shrinking supply voltages and increasing device speed. Finally, a discussion on minimalistic and digitally assisted design approaches is used to sketch a route toward further improvements in ADC power efficiency and performance.**

I. Introduction

Analog-to-digital converters (ADCs) are important building blocks in modern electronic systems. In many cases, the efficiency and speed at which analog information can be converted to digital signals profoundly affects system architectures and their performance. Even though modern integrated circuit technology can provide very high conversion rates, the associated power dissipation is often incompatible with application constraints. For instance, the high-speed 6-8-bit ADCs of [1, 2] achieve sampling rates in excess 20 GS/s, at power dissipations of 1.2 W and 10 W, respectively. Operating such blocks in a handheld application is impractical, as they would drain the device's battery within a short amount of time. Consequently, it is not uncommon to architect power constrained applications "bottom-up," by determining the analog/RF front-end and ADC specifications based on the available power budget. A discussion detailing such an approach for the specific example of software-defined radio receivers is presented in [3].

With power dissipation being among the most important concerns in mixed-signal/RF applications, it is important to track trends and understand the relevant trajectories. The purpose of this paper is to review the latest developments in low-power A/D conversion and to provide an outlook on future possibilities.

Following this introduction, Section II provides survey data on ADCs published over the past eleven years. These data show that contrary to common perception, extraordinary progress has been made in improving ADC power efficiency. Among the factors that have influenced this trend are technology scaling, and the increasing use of simplified analog sub-circuits with digital correction. Therefore, Section III takes a closer look at the impact of shrinking feature sizes, while Sections IV and V discuss recent ideas in "minimalistic" and "digitally assisted" ADC architectures.

II. ADC Performance Trends

A. Survey Data and Figure of Merit Considerations

Several surveys on ADC performance are available in literature [3-7]. In this section, we will review recent data from designs presented at the IEEE International Solid-State Circuits Conference (ISSCC) and the VLSI Circuit Symposium. Fig.1. shows a scatter plot of results published at these venues over the past eleven years [8]. Fig. 1(a) plots the energy per Nyquist sample (P/f_s, i.e. power divided by Nyquist sampling rate) against the achieved signal-to-noise-and-distortion ratio (SNDR). This plot purposely avoids dividing the conversion energy by the effective number of quantization steps, as done in the commonly used figure of merit [4]

$$FOM = \frac{P}{f_s \cdot 2^{ENOB}} \tag{1}$$

where

$$ENOB = \frac{SNDR(dB) - 1.76}{6.02} \tag{2}$$

Normalizing by the number of quantization steps assumes that doubling precision would double power, which finds only empirical justification [4]. In fact, if a converter were purely limited by thermal noise, its power would quadruple per added bit (see Section III). However, as this assumption is pessimistic for real designs, it is preferable to avoid a fixed relationship between precision and energy when plotting data across a large range of architectures and resolutions.

Fig. 1(a) indicates that some of the lowest energy ADCs were published at this year's ISSCC. Interestingly, most of these designs target only low to moderate resolution; activity in the high-resolution space appears to be lagging. With respect to (1), included as a straight line for the numerical example of FOM = 100 fJ/conversion-step, it is clearly visible that state-of-the-art high resolution designs (SNDR > 85dB) do not obey the implied 2x increase in power per bit. Furthermore, the most recent low-resolution designs also manage to break away from any linear fit to the overall scatter plot that is based on a slope of 2x per bit.

In addition to an ADC's energy efficiency, the available signal bandwidth is an important parameter. Fig 1(b) plots bandwidth against SNDR for the given data set. In this chart, the bandwidth plotted for Nyquist converters is equal to the input frequency used to obtain the stated SNDR; this

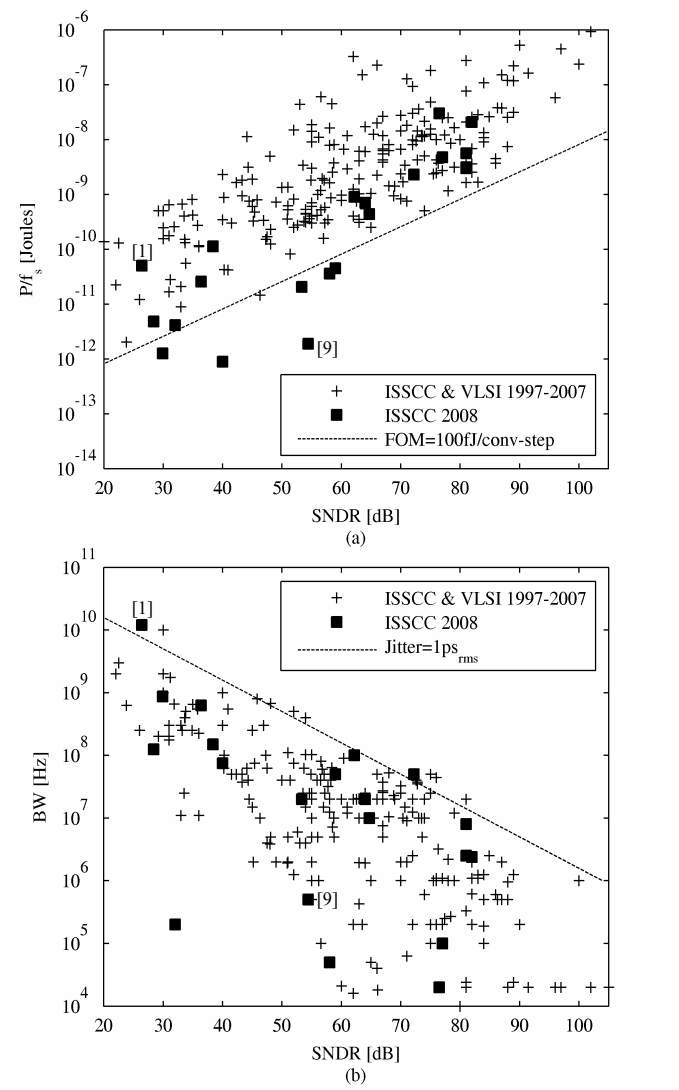

Fig. 1. ADC performance data (ISSCC 1997-2008, VLSI Circuit Symposium 1997-2007). (a) Power efficiency versus SNDR. (b) Conversion bandwidth versus SNDR.

frequency is not necessarily $f_s/2$. The first interesting observation from Fig. 1(b) is that across all resolutions, the parts with the highest bandwidth achieve a performance that is approximately equivalent to an aperture uncertainty of 1 ps_{rms}. The dashed line in Fig. 1(b) represents the performance of an ideal sampler with sinusoidal input and 1 ps_{rms} sampling clock jitter. Clearly, any of the ADC designs at this performance front rely on a significantly better clock, to allow for additional nonidealities that reduce SNDR. Such nonidealities include quantization noise, thermal noise, differential nonlinearity and harmonic distortion. From the data in Fig 1(b), it is also clear that any new design aiming to push the speed-resolution envelope will require a sampling clock with jitter on the order of ~100 fs_{rms} or better.

In order to assess the overall merit of an ADC (power efficiency and bandwidth), it is interesting to compare the locations of its particular design points in plots (a) and (b). For example, [1] achieves a bandwidth close to the best designs, while showing only average power efficiency. The opposite is true for [9]; this part ranks among the lowest energy designs published to date, but achieves only moderate bandwidth. These examples confirm the intuition that pushing a design toward the speed limits of a given technology will sacrifice power efficiency. To date, there exists no single-number figure of merit that captures this tradeoff in a fair and balanced way across all architectures and resolutions. The same holds true for input capacitance. For example, it is possible to improve the SNDR of most ADC architectures by increasing their input capacitance. An ideal figure of merit would take the power needed to drive the converter input into account.

B. Trends in Power Efficiency and Speed

Using the data set discussed above, it is interesting to extract trends over time. Fig. 2(a) is a 3-D representation of the power efficiency data [Fig. 1(a)] with the year of publication included along the y-axis. The resulting slope in time corresponds to an average reduction in power by 2x approximately every 2 years.

A similar 3-D fit could be constructed for bandwidth performance. However, such a fit would not convey interesting information, as the majority of designs published in recent years do not attempt to maximize bandwidth. This contrasts the situation with power efficiency, which is subject to optimization in most modern designs. In order to extract a

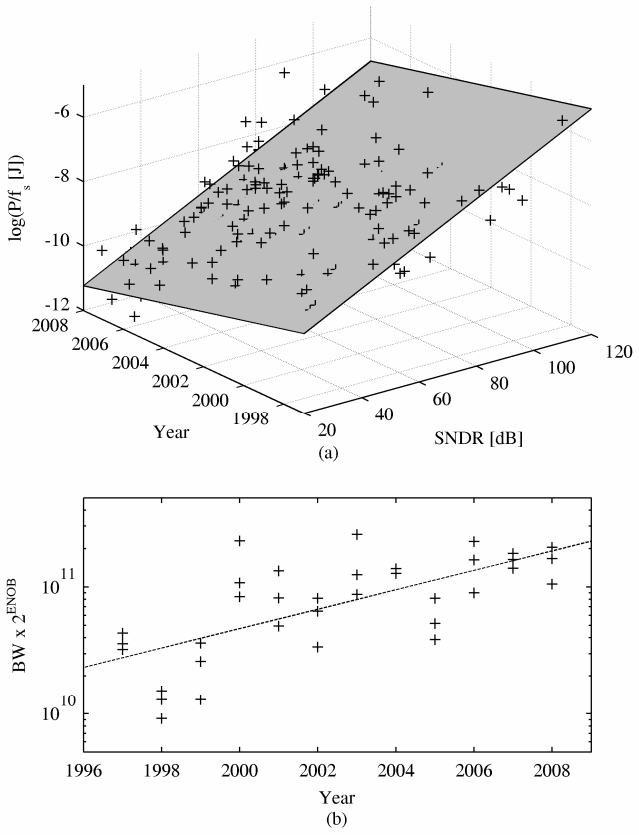

Fig. 2. Trends in ADC performance. (a) 3-D fit to power efficiency. The fit plane has a slope of 0.5x/2 years along the time axis. (b) Fit to speed-resolution product of top 3 designs in each year. The slope of the fit line is 2x/4 years.

trend on achievable bandwidth, Fig. 2(b) scatter-plots the speed-resolution products of the top three designs in each year. This metric is justified by the constant speed-resolution boundary observed from Fig. 1(b). A fit to the data in Fig. 2(b) reveals that performance has doubled every 4 years; a rate that is much lower than the improvement in power efficiency. In addition, as evident from the data points, there is no clear trend as far as the top performance point is concerned; designs of the early 2000's are up to par with some of the works published recently. Consequently, the extracted progress rate of speed-resolution performance should be viewed as a relatively weak and error-prone indicator.

III. Impact of Technology Scaling

As shown above, the power dissipation of A/D converters has halved approximately every 2 years over the past decade. Over the same period, CMOS technologies used to implement the surveyed ADCs have scaled from approximately 0.6 μm down to 65 nm. In this section, we will investigate the role of technology scaling in the observed power efficiency trend. Broader discussions on the impact of scaling are presented in [7, 10, 11].

A well-known challenge in designing ADCs using modern processes is the diminishing voltage headroom. Since device scaling requires a reduction in supply voltage (V_{DD}), the noise in the analog signals must be reduced proportionally to maintain the desired signal-to-noise ratio. Since noise trades with power dissipation, this suggests to first order that power efficiency should worsen, and not improve, for ADCs in modern technologies. One way to overcome supply voltage limitations is to utilize thick-oxide I/O devices, which are available in most standard CMOS processes. However, using those devices often reduces speed. Closer inspection of the survey data considered in this paper reveals that most published state-of-the-art designs do not rely on thick oxide devices, and rather cope with supply voltages around 1 V.

To investigate further, it is worthwhile to examine the underlying equations that capture the trade-off between supply voltage and power dissipation via thermal noise constraints. In most analog sub-circuits used to build ADCs, noise is inversely proportional to capacitance

$$N \propto \frac{kT}{C} \tag{3}$$

where k is Boltzmann's constant and T stands for absolute temperature. For the specific case of a transconductance amplifier that operates linearly, we can write

$$f_s \propto \frac{g_m}{C} \tag{4}$$

Further assuming that the signal power is proportional to $(\alpha \cdot V_{DD})^2$ and that the circuit's power dissipation is V_{DD} multiplied by the transistor drain current, I_D, we find

$$\frac{P}{f_s} \propto \frac{1}{\alpha^2}\frac{1}{V_{DD}} \cdot \frac{1}{\left(\frac{g_m}{I_D}\right)} kT \cdot SNR \tag{5}$$

The variable g_m/I_D in (5), is related to the "gate overdrive" of the transistor that implements the transconductance. Assuming MOS square law, $g_m/I_D = 2/(V_{GS}-V_t)$ and in weak inversion $g_m/I_D = 1/(n \cdot kT/q)$, with $n \cong 1.5$. Considering the fractional swing (α) and transistor bias point (g_m/I_D) as constant, it is clear from the above expression that power efficiency in noise-limited transconductors should deteriorate at low V_{DD}. In addition, we see that (5) indicates a very steep tradeoff between SNR and energy; increasing the SNR by 6 dB requires a 4x increase in P/f_s.

Since both of these results do not correlate well with the observations of Section II, it is instructive to examine the assumptions that lead to (5). The first assumption is that the circuit is purely limited by thermal noise. This assumption clearly holds for ADCs with very high resolution, but typically few, if any, low resolution converters are impaired by thermal noise.

To get a feel for a typical SNDR value at which today's converters become "purely" limited by noise, it is helpful to plot the data of Fig. 1(a) normalized to a 4x power increase per bit [12]. Fig. 3 shows such a plot in which the P/f_s values have been divided by

$$\left(\frac{P}{f_s}\right)_{min} = 4 \cdot kT \cdot SNR \tag{6}$$

while assuming SNR ≅ SNDR. The pre-factor of 4 in this expression follows from the power dissipated by an ideal class-B amplifier that drives the capacitance C with a rail-to-rail tone at $f_s/2$ [13]. Therefore, (6) represents a fundamental bound on the energy required to process a charge sample at a given SNR.

The primary observation from Fig. 3 is that the normalized data exhibits a visible "corner" beyond which $(P/f_s)/(P/f_s)_{min}$ approaches a constant value. This corner, approximately located at 75 dB, is an estimate for the SNDR at which a typical state-of-the-art design becomes truly limited by thermal noise. Since ADCs with lower SNDR do not achieve

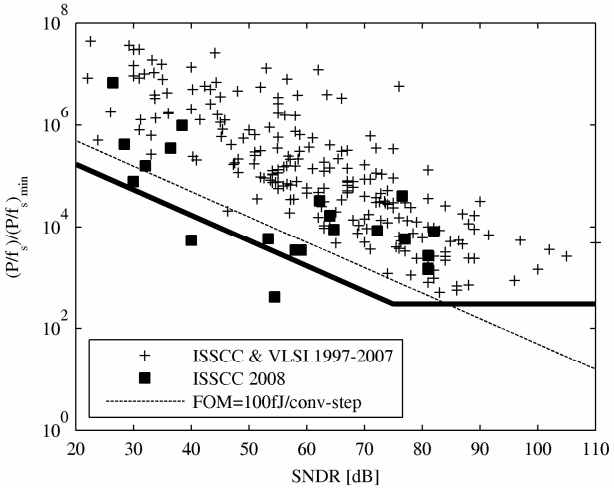

Fig. 3. Data of Fig. 1(a) normalized by $(P/f_s)_{min}$ as given in (6). This illustration suggests the existence of an "SNDR corner." Only ADCs with SNDR > 75dB appear to be primarily limited by thermal noise.

the same noise-limited power efficiency, it can be argued that these designs are at least partially limited by the underlying technology. This implies that over time, technology scaling may have helped improve their power efficiency as opposed to the worsening predicted by (5).

To investigate further, we partitioned the data of Fig. 1(a) into two distinct sets, i.e. high resolution (SNDR > 75 dB) and low-to-moderate resolution (SNDR ≤ 75 dB). We then applied a 3-D fit similar to that shown in Fig. 2(a) to each set and extracted the progress rates over time. For the set with SNDR > 75 dB it was found that P/f_s has halved only every 5.4 years, while for SNDR ≤ 75 dB, P/f_s halves every 1.6 years. The difference in these progress rates confirms the above speculation. For high-resolution designs, (5) applies and scaling technology over time, associated with lower supply voltages, cannot help improve power efficiency. As observed in [7], this has led to a general trend toward lower resolution designs. Since it is very difficult to attain high SNDR at low supply voltages, most applications are steered away from relying on high-resolution ADCs in current fine-line processes. This is qualitatively confirmed in Fig. 4, which highlights the P/f_s data points of ADCs built in CMOS at 90 nm (and $V_{DD} \cong 1$ V) and below.

The above-discussed situation strongly contrasts the impact of scaling on low-to-moderate resolution designs, as evident from the extracted improvement rate. Quantifying the benefits of scaling on low-to-moderate resolution ADCs from first principles is a complex task, primarily because the involved tradeoffs strongly depend on architecture and design specifics. An analysis that highlights the benefits of scaling in flash and folding ADCs is presented in [14]. In the following paragraphs we will discuss qualitatively the scaling behavior of a moderate resolution pipelined ADC.

Consider the 10-bit, 0.6-µm pipelined ADC described in [15, 16]; this design reflects state-of-the-art in 1996. Close inspection of the design details in [16] reveals that about 30% of the total power in this ADC is dissipated by noise-limited amplifiers. The remaining power is consumed by digital gates, comparators and amplifier stages whose component sizes are set by feature size constraints. To first order, the power dissipation in these latter blocks should scale approximately as $C \cdot V_{DD}^2$, i.e. logic gate energy. Since 1997, we have seen a reduction in process $C \cdot V_{DD}^2$ of approximately 300 times [17]. Yet, this change alone cannot explain the vastly larger improvement factor that 10-bit designs have seen over the past decade; the improvement would be limited to no more than 3.3x in terms of total power.

Clearly, the situation is more complex. First, a circuit that is "limited" by noise may still carry overhead that reduces with scaling. Especially in high-speed designs, amplifier self-loading and parasitic loading at intermediate circuit nodes plays an important role. Technology scaling helps mitigate these capacitances and therefore improves overall efficiency. Unfortunately this effect is hard to quantify.

A more transparent factor is the trade-off between g_m/I_D and the transit frequency (f_T) of the active devices. Switched capacitor circuits based on class-A operational transconductance amplifiers typically require transistors with $f_T > 80\ f_s$. Even for speeds of several tens of MS/s, it was necessary in older technologies to bias transistors far into strong inversion (V_{GS}-V_t > 200 mV) to satisfy this requirement. In more recent technologies, very large transit frequencies are available in moderate- and even weak-inversion. This is further illustrated in Fig. 5(a) which compares typical minimum-length NMOS devices in 180-nm and 90-nm CMOS.

For a fixed sampling frequency, and hence fixed f_T requirement, newer technologies deliver higher g_m/I_D. This tradeoff is plotted directly, without the intermediate variable V_{GS}-V_t, in Fig. 5(b). In order to achieve f_T = 30 GHz, a 180-nm device must be biased such that $g_m/I_D \cong 9$ S/A. In 90-nm technology, f_T = 30 GHz is achieved in weak inversion, at ≅ 18

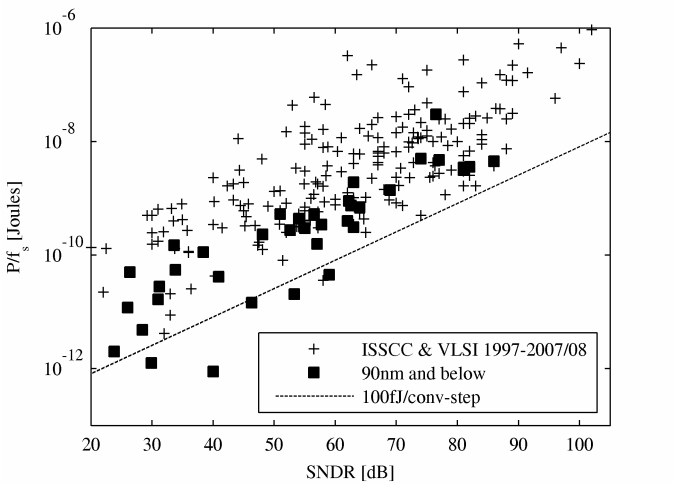

Fig. 4. Power efficiency (P/f_s) for ADCs built in 90-nm CMOS and below.

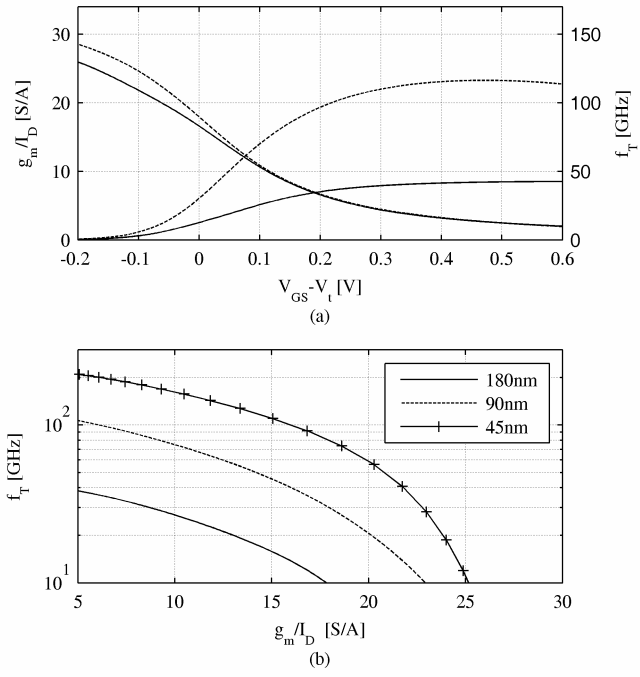

Fig. 5. Tradeoff between g_m/I_D and f_T in modern technologies.

S/A. From equation (5), it is clear that this advantage can counteract the reduction in V_{DD} when going to a newer process. Note that this advantage can only materialize when the sampling speed is kept constant or at least not scaled proportional to the f_T improvement. This was also one of the observations drawn from Fig. 2(b). A converter that pushes the speed envelope using a new technology typically won't simultaneously benefit from scaling in terms of power efficiency.

A last and perhaps even more significant factor to consider is the accrual of design experience, improved optimization, exploitation of process options and refinement of circuit techniques over many generations of technology. For instance, A/D converters in 5-V technologies used to be relatively wasteful in terms of headroom utilization [α-term in (5)]. Newer designs are typically optimized to accommodate signal swings as large as 1 $V_{pp,diff}$ at V_{DD} = 1 V. In addition, we are beginning to see designs that efficiently exploit technology options. For instance, the 10-bit pipelined ADC of [18] uses thin-oxide high-performance analog (HPA) devices to achieve high DC gain using simple, power-efficient telescopic transconductance amplifiers.

Additional directions in the context of design techniques that have recently gained in importance include "minimalistic" and "digitally assisted" approaches. The trend toward such solutions may be explained by the fact that evolutionary grown designs have come very close to practical power limits, imposed by their circuit topologies and associated fundamental constraints. The ideas summarized in the following two sections outline promising directions in this area of research.

IV. Minimalistic Design

Power dissipation in the analog portion of ADCs is strongly coupled to the complexity of the constituent sub-circuits. The goal of minimalistic design is to improve power efficiency and potentially increase speed by utilizing simplified analog sub-circuits.

In architectures that previously relied on op-amp based signal processing, there exists a clear trend toward simplified amplifier structures. Examples include inverter-based sigma-delta modulators [19, 20] and various approaches emphasizing op-amp-less implementation of pipelined ADCs [21-25]. Especially in switched capacitor circuits, eliminating class-A op-amps can dramatically improve power efficiency. This is for two reasons. First, operational amplifiers typically contribute more noise than simple gain stages, as for example resistively loaded open-loop amplifiers. Secondly, the charge transfer in class-A op-amp circuitry is inherently inefficient; the amplifier draws a constant current, while delivering on average only a small fraction of this current to the load. In [26], it was found that the efficiency of a class-A op-amp in a switched capacitor circuit is inversely proportional to the number of settling time constants. For the typical case of settling for approximately 10 or more time constants, the overall efficiency, i.e. charge drawn from the supply versus charge delivered to the load, is only a few percent.

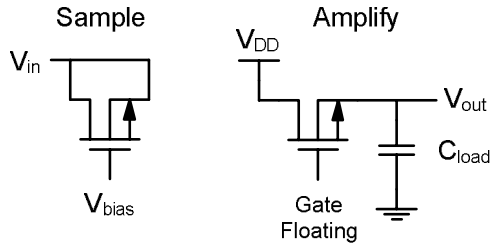

Fig. 6. Dynamic amplifier concept used in [23].

As discussed further in [27], this inherent inefficiency of op-amps contributes to the power overhead relative to fundamental limits. Consider for instance the horizontal asymptote of Fig. 3, located at approximately 300 times the minimum possible P/f_s. The factor of 300 can be explained for op-amp based circuits as follows. First, the noise is typically given by $\beta \cdot kT/C$, where β can range from 5-10, depending on implementation details. Second, charge transfer using class-A circuits, as explained above brings a penalty of approximately 20x. Third, op-amp circuits usually do not swing rail-to-rail as assumed in (6); this can contribute another factor of two. Finally, adding further power contributors beyond one dominant op-amp circuit easily explains a penalty factor greater than 200…400.

A promising remedy to this problem is to utilize circuits that process charge more efficiently and at the same time contribute less thermal noise. A well-know example of an architecture that achieves very high efficiency is the charge-based successive approximation register (SAR) converter [1, 9, 28]. Such converters have seen a renaissance in recent years, primarily because the architecture is well-suited for leveraging the raw transistor speed of new technologies, while being insensitive to certain scaling implications, such as reduced intrinsic gain (g_m/g_{ds}). A problem with SAR architectures is that they cannot deliver the best possible performance when considering absolute speed, resolution and input capacitance simultaneously. This is one reason why relatively inefficient architectures, such as op-amp based pipelined ADCs are still being used and investigated.

In order to make pipelined architectures as power efficient as competing SAR approaches, various ideas are being explored in research. Fig. 6 shows a new single-transistor residue amplification concept that was utilized in the low-power pipelined converter of [23]. In the sampling phase, the transistor is configured as a MOS capacitor in accumulation. During amplification, the drain is switched to V_{DD} and the source drives the discharged capacitive load. At this time, the gate is left floating and the transistor acts as a source follower. V_{out} rises until V_{gs} nears the threshold voltage of the device. Incremental input voltage amplification occurs because of charge conservation. During sampling, signal dependent charge is stored on C_{ox} and gate-source/drain overlap parasitics (C_{ol}). At the end of the amplification phase, the

charge on $C_{gs} = C_{ox}+C_{ol}$ is constant (due to $V_{gs} = V_t$) and all signal dependent charge now appears across $C_{gd} = C_{ol}$. To first order, the voltage gain is set by the ratio $(C_{ox}+2C_{ol})/C_{ol}$. In essence, this scheme mimics charge-redistribution around an operational amplifier, while providing significantly lower noise and highly efficient charge transfer from the supply to the capacitive load.

A general concern with most minimalistic design approaches is that they tend to sacrifice robustness, e.g., in terms of power supply rejection, common mode rejection and temperature stability. It remains to be seen if these issues can be handled efficiently in practice. Improving supply rejection, for instance, could be achieved using voltage regulators. This is custom practice in other areas of mixed-signal design, as for example PLLs [29]. Especially when the power of the ADC's critical core circuitry is lowered by orders of magnitude, implementing supply regulation should be a manageable task.

A second issue with minimalistic designs is the achievable resolution and linearity. Op-amp circuits with large loop gain help linearize transfer functions; this feature is often removed when migrating to simplified circuits. For instance, the amplifier scheme of Fig. 6 is linear only to approximately 9-bit resolution. In cases where simplicity sacrifices precision, it is attractive to consider digital means for recovering conversion accuracy. Digitally assisted architectures are therefore the topic of the next section.

V. Digitally Assisted Architectures

Technology scaling has significantly reduced the energy per operation in CMOS logic circuits. As explained in [30], the typical 0.7x scaling of features along with aggressive reductions in supply voltage have led to a 65% reduction in energy per logic transition for each technology generation. The survey data presented in [17] suggests that a 2-input NAND gate dissipates roughly 1.3 pJ per logic operation in a 0.5-µm CMOS process. The same gate dissipates only 4.5 fJ in a more recent 90-nm process; this amounts to a ~300x improvement in only 10 years. The corresponding progress in ADC energy (based on Section II) amounts to a 32x reduction over 10 years. This means that the relative "cost" of digital computation has reduced roughly by a factor of ten over the past decade.

To get a feel for how much logic can be used to "assist" a converter for the purpose of calibration and error correction, it is interesting to the divide energy per conversion (P/f_s) figures of ADCs by the energy of a single NAND gate. The numbers compiled in Table 1 use data from the fit plane of Fig. 2(a) for 2008, and assume $E_{NAND} = 4.5$ fJ. At low signal fidelity, e.g. SNDR = 30 dB, a single A/D conversion consumes as much energy as toggling approximately 4,700 logic gates. On the other hand, at 90 dB SNDR, more than two million logic gates would need to switch to consume the energy of an A/D conversion at this level of precision.

The consequence of this observation is that in a low-resolution converter, it is unlikely that tens of thousand of gates can be used for digital error correction without exceeding reasonable energy or power limits. A large number of gates may be affordable only if the involved gates operate at a low activity factor or if they can be shared within the system. Conversely, in high resolution ADCs, each analog operation is very energy consuming and even a large amount of digital processing may be accommodated in the overall power budget.

TABLE I

$E_{ADC} = P/f_s$ in today's ADCs [fit data from Fig. 2(b)], relative to logic gate energy ($E_{NAND} = 4.5$ fJ) in 90-nm CMOS.

SNDR [dB]	E_{ADC}	E_{ADC}/E_{NAND}
30	21 nJ	4,700
50	168 nJ	38,000
70	1.35 µJ	300,000
90	10.8 µJ	2,400,000

The following sub-sections provide a non-exhaustive list of opportunities for leveraging digital logic gates in A/D converters.

A. Oversampling

The longest standing example of an architecture that efficiently leverages digital signal processing abilities is the oversampling sigma-delta converter. This architecture uses noise shaping to push the analog quantization error, along with other nonidealities, outside the signal band [12]. Subsequent digital filtering creates a high-fidelity output signal, while the constituent analog sub-circuits require only moderate precision. Even in fairly old technologies, it was reasonable to justify high gate counts in the converter's decimation filter, simply because the analog signal processing energy per sample for typical high-SNDR converters is very large.

A new paradigm that might gain significance in the future is the use of oversampling in traditional Nyquist converters. An example of such an ADC is described in [31]. As we have noted from Fig. 5(b), migrating a converter with a fixed sampling rate to technologies with higher f_T can help improve power efficiency. Ultimately, however, there is diminishing return in this trend due to the weak-inversion "knee" of MOS devices [see Fig. 5(a)]. G_m/I_D no longer improves beyond a certain minimum bias; it therefore makes no sense to target a transistor f_T below a certain value. This, in turn, implies that for optimum power efficiency, one should not operate an ADC below a certain clock rate. Consider for example the f_T versus g_m/I_D plot for 45-nm technology in Fig. 5(b). For $g_m/I_D > 20$ S/A, f_T drops sharply without a significant increase in g_m/I_D. At this point, $f_T \cong 50$ GHz, implying that is possible to build a switched capacitor circuit with $f_{clock} = 50$ GHz/80 = 625 MHz.

To date, there exist only a limited number of applications that require such high sampling rates, and there will clearly remain a number of systems in the future that demand primarily good power efficiency at only moderate speeds. A solution to this situation might be to oversample the input signal by a large factor and to remove out-of band noise (thermal noise, quantization noise, and jitter) using a digital filter. Per octave of oversampling, this increases ADC

resolution by ½ bit. In a situation where a converter is purely limited by noise, this improvement is in line with the fundamental thermal noise tradeoff expressed in (5).

B. Mismatch Correction

Assuming constant gate area (W·L), transistor matching tends to improve in newer technologies. In matching-limited flash ADC architectures, this trend has been exploited in the past to improve the power efficiency by judiciously down-sizing the constituent devices [14]. In order to scale such architectures more aggressively and at the same time address new sources of mismatch in nano-scale technologies, it is desirable to aid the compensation of matching errors through digital means.

In flash-ADCs, there are at least two general trends in this direction. The first is to absorb mismatch errors using a fault-tolerant encoder in conjunction with a comparator bank that contains redundant elements [32]. This approach requires a relatively large number of logic operations per sample. In light of the conclusions from Table 1, such a solution may be most efficient only for technologies below 90 nm.

An alternative approach is to provide redundant comparators that are selectively activated (in a static manner) to minimize the converter's nonlinearity [33]. An extension of this approach is to include digital trimming circuitry in addition to redundant elements [34]. This scheme can yield good power efficiency as it attacks the mismatch problem along two degrees of freedom. An extension of this idea, incorporating redundant channels in a time-interleaved ADC is discussed in [35].

C. Digital Linearization of Amplifiers

As pointed out in Section IV, power-efficient and minimalistic design approaches are typically unsuitable for high-resolution applications, unless appropriate digital error correction schemes are used to enhance conversion linearity. In [20], it was demonstrated that a simple open-loop differential pair used in a pipeline ADC can be digitally linearized to within 12-bit precision. Such a digital correction requires only moderate complexity on the order of a few thousand logic gates. It is foreseeable that the future will bring additional schemes that help improve the nonlinearity of simplified amplifiers; e.g. in the context of sigma-delta modulators.

One key issue in most digital linearization schemes is that the correction coefficients must track changes in operating conditions relatively quickly; preferably with time constants no larger than 1-10 ms. Unfortunately, most of the basic statistics-based algorithms for coefficient adaptation require much longer time constants at high target resolutions [36, 37]. Additional research is needed to extend the recently proposed "split-ADC" [38, 39] and feedforward noise cancellation techniques [40] for use in nonlinear calibration schemes.

D. Digital Correction of Dynamic Errors

Most of the digital correction methods developed in recent years have targeted the compensation of static circuit errors; work on dynamic errors that limit a converter's effective resolution bandwidth has been lagging. In various applications, as for instance sub-sampling base-station receivers, it is desirable to improve the converter's high-frequency linearity beyond raw technology limits [41]. Digital compensation of the relevant frequency dependent nonlinearities in the sampling front-end of ADCs will likely evolve as an attractive area for future research. If digital capabilities in nano-scale technologies continue to improve, dynamic compensation schemes based on relatively complex Volterra series may become feasible [42].

E. System-Synergistic Error Correction Approaches

In the above discussion, ADCs are being viewed as "black boxes" that deliver a set performance without any system level interaction. Given the complexity of today's applications, it is important to realize that there exist opportunities to improve ADC performance by leveraging specific system and signal properties.

For instance, large average power savings are possible in radio receivers when ADC resolution and speed are dynamically adjusted to satisfy the minimum instantaneous performance needs. The design described in [43] demonstrates the efficacy of such an approach.

In the context of digital correction, it is conceivable to leverage known properties of application-specific signals to "equalize" the A/D converter together with the communication channel [44, 45]. For instance, the converter described in [45] uses the system's OFDM pilot tones to extract component mismatch information. In such an approach, existing system hardware, as for instance the FFT block, can be re-used to accommodate ADC calibration.

VI. SUMMARY

This paper has summarized recent trends in the context of low-power A/D conversion. Using survey data from the past eleven years, we have observed that power efficiency in ADCs has improved at an astonishing rate of 2x every 2 years. In part, this progress rate is based on cleverly exploiting the strengths of today's technology. Smaller feature sizes help improve the power dissipation in circuits that are not limited by thermal noise. In circuit elements that are limited by thermal noise, exploiting the high f_T of modern transistors can be of help in mitigating a penalty from low supply voltages.

A promising paradigm is the trend toward minimalistic ADC architectures and digital means of correcting analog circuit errors. Digitally assisted ADCs aim to leverage the low computing energy of modern processes to improve the resolution and robustness of simplified circuits. Future work in this area promises to fuel further progress in optimizing the power efficiency of A/D converters.

Overall, future improvements in ADC power dissipation are likely to come from a combination of aspects that involve improved system embedding and reducing analog sub-circuit complexity and raw precision at the expense of "cheap" digital processing resources.

REFERENCES

[1] P. Schvan, et al, "A 24GS/s 6b ADC in 90nm CMOS," *ISSCC Dig. Techn. Papers*, pp. 544-545, Feb. 2008.

[2] K. Poulton, et al., "A 20-GSample/s 8b ADC with a 1-MByte Memory in 0.18-um CMOS," *ISSCC Dig. Techn. Papers*, pp. 318-319, Feb. 2003.

[3] A. A. Abidi, "The Path to the Software-Defined Radio Receiver," *IEEE J. Solid-State Circuits*, vol. 42, no. 5, pp. 954-966, May 2007.

[4] R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Select. Areas Commun.*, vol. 17, no. 4, pp. 539-550, Apr. 1999.

[5] P. B. Kenington and L. Astier, "Power Consumption of A/D Converters for Software Radio Applications," *IEEE Trans. Vehicular Techn.*, vol. 49, pp. 643-650, Mar. 2000.

[6] K. G. Merkel and A. L. Wilson, "A Survey of High Performance Analog-to-Digital Converters for Defense Space Applications," *Proc. IEEE Aerospace Conf.*, vol. 5, pp. 2415-2427, Mar. 2003.

[7] Y. Chiu, "Scaling of analog-to-digital converters into ultra-deep-submicron CMOS," *Proc. CICC*, pp. 375-382, Sep. 2005.

[8] B. Murmann, "ADC Performance Survey 1997-2008," [Online]. Available: http://www.stanford.edu/~murmann/adcsurvey.html.

[9] M. van Elzakker, et al., "A 1.9µW 4.4fJ/Conversion-step 10b 1MS/s Charge-Redistribution ADC," *ISSCC Dig. Techn. Papers*, pp. 244-245, Feb. 2008.

[10] A.-J. Annema, et al., "Analog Circuits in Ultra-Deep-Submicron CMOS," IEEE J. Solid-State Circuits, vol. 40, no. 1, pp. 132-143, Jan. 2005.

[11] K. Bult, "The Effect of Technology Scaling on Power Dissipation in Analog Circuits," in Analog Circuit Design, M. Steyaert, A. H. M. Roermund, and J. H. v. Huijsing, (eds.), Springer, 2006.

[12] R. Schreier and G. C. Temes, *Understanding Delta-Sigma Data Converters*: IEEE Press, 2005.

[13] E. A. Vittoz, "Future of analog in the VLSI environment," *Proc. IEEE ISCAS*, pp.1372-1375, May 1990.

[14] P. C. S Scholtens, et al., "Systematic power reduction and performance analysis of mismatch limited ADC designs," *Proc. ISLPED*, pp. 78-83, Aug. 2005.

[15] T. B. Cho, et al., "A power-optimized CMOS baseband channel filter and ADC for cordless applications," *Dig. VLSI Circuits Symposium*, pp.64-65, Jun. 1996.

[16] G. Chien, "High-Speed, Low-Power, Low-Voltage Pipelined Analog-to-Digital Converter, MS Thesis, University of California, Berkeley, 1996.

[17] B. Murmann, "Digitally Assisted Analog Circuits – A Motivational Overview," *ISSCC Special-Topic Evening Session* (SE1.1), Feb. 2007.

[18] M. Boulemnakher, et al., "A 1.2V 4.5mW 10b 100MS/s Pipeline ADC in a 65nm CMOS," *ISSCC Dig. Techn. Papers*, pp. 250-251, Feb. 2008.

[19] Y. Chae et al., "A 0.7V 36µW 85dB-DR Audio ΔΣ Modulator Using Class-C Inverter," *ISSCC Dig. Techn. Papers*, pp. 491-492, Feb. 2008.

[20] R. H. M. van Veldhoven et al., "An Inverter-Based Hybrid ΣΔ Modulator," *ISSCC Dig. Techn. Papers*, pp. 493-494, Feb. 2008.

[21] B. Murmann and B. E. Boser, "A 12-bit 75-MS/s Pipelined ADC using Open-Loop Residue Amplification," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2040-2050, Dec. 2003.

[22] J. K. Fiorenza *et al.*, "Comparator-Based Switched-Capacitor Circuits for Scaled CMOS Technologies," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2658-2668, Dec. 2006.

[23] J. Hu, et al., "A 9.4-bit, 50-MS/s, 1.44-mW Pipelined ADC Using Dynamic Residue Amplification," *Dig. VLSI Circuits Symposium*, Jun. 2008.

[24] M. Anthony, et al., "A Process-Scalable Low-Power Charge-Domain 13-bit Pipeline ADC," *Dig. VLSI Circuits Symposium*, Jun. 2008.

[25] A. Nazemi, et al., "A 10.3GS/s 6bit (5.1 ENOB at Nyquist) Time-Interleaved/Pipelined ADC Using Open-Loop Amplifiers and Digital Calibration in 90nm CMOS, *Dig. VLSI Circuits Symposium*, Jun. 2008.

[26] E. Iroaga and B. Murmann, "A 12b, 75MS/s Pipelined ADC Using Incomplete Settling," *IEEE J. Solid-State Circuits*, vol. 42, no. 4, pp. 748-756, Apr. 2007.

[27] B. Murmann, "Limits on ADC Power Dissipation," in Analog Circuit Design, M. Steyaert, A. H. M. Roermund, and J. H. v. Huijsing, (eds.), Springer, 2006.

[28] D. Draxelmayr, "A 6b 600MHz 10mW ADC array in digital 90nm CMOS," *ISSCC Dig. Techn. Papers*, pp. 264-265, Feb. 2004.

[29] A. Maxim and M. Gheorghe, "A sub-1psrms jitter 1-5GHz 0.13um CMOS PLL using a passive feedforward loop filter with noiseless resistor multiplication," *Dig. RFIC Symposium*, pp. 207–210, Jun. 2005.

[30] S. Borkar, "Design challenges of technology scaling," *IEEE Micro,* vol. 19, pp. 23-29, Apr. 1999.

[31] M. Hesener, et al., "A 14b 40MS/s Redundant SAR ADC with 480MHz Clock in 0.13pm CMOS," *ISSCC Dig. Techn. Papers*, pp. 248-249, Feb. 2007.

[32] C. Paulus, et al. "A 4GS/s 6b flash ADC in 0.13um CMOS," *Dig. VLSI Circuits Symposium*, pp. 420-423, Jun. 2004.

[33] C. Donovan and M. Flynn, "A 'Digital' 6-bit ADC in 0.25-µm CMOS," *IEEE J. Solid State Circuits*, vol. 37, pp. 432-437, Mar. 2002.

[34] S. Park, et al., "A 4-GS/s 5-b Flash ADC in 0.18um CMOS," *IEEE J. Solid-State Circuits*, vol. 42, no. 9, pp. 1865-1872, Sep. 2007.

[35] B. P. Ginsburg and A. P. Chandrakasan, "Highly Interleaved 5b 250MS/s ADC with Redundant Channels in 65nm CMOS," *ISSCC Dig. Techn. Papers*, pp. 240-241, Feb. 2008.

[36] B. Murmann and B. E. Boser, "Digital Domain Measurement and Cancellation of Residue Amplifier Nonlinearity in Pipelined ADCs," *IEEE Trans. on Instrumentation and Measurement*, vol. 56, no. 6, pp. 2504-2514, Dec. 2007.

[37] A. Panigada and I. Galton, "Digital background correction of harmonic distortion in pipelined ADCs," *IEEE Trans. Circuits Syst. I*, vol. 53, no. pp. 1885–1895, Sep. 2006.

[38] J. Li and U.-K. Moon, "Background calibration techniques for multistage pipelined ADCs with digital redundancy," *IEEE Trans. Circuits Syst. II*, vol. 50, no. 9, pp. 531–538, Sep. 2003.

[39] J. McNeill, et al., "'Split ADC' architecture for deterministic digital background calibration of a 16-bit 1-MS/s ADC," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2437–2445, Dec. 2005.

[40] K.-W. Hsueh, et al, "A 1V 11b 200MS/s Pipelined ADC with Digital Background Calibration in 65nm CMOS," *ISSCC Dig. Techn. Papers*, pp. 547-548, Feb. 2008.

[41] P. Nikaeen and B. Murmann, "Digital Correction of Dynamic Track-and-Hold Errors Providing SFDR > 83 dB up to fin = 470 MHz," in preparation.

[42] Y. Chiu, et al., "Least-mean-square adaptive digital background calibration of pipelined analog-to-digital converters," *IEEE Trans. Circuits and Systems I*, vol. 51, pp. 38-46, Jan. 2004.

[43] P. Malla et al., "A 28mW Spectrum-Sensing Reconfigurable 20MHz 72dB-SNR 70dB-SNDR DT ΔΣ ADC for 802.11n/WiMAX Receivers," *ISSCC Dig. Techn. Papers*, pp. 496-497, Feb. 2008.

[44] W. Namgoong, "A channelized digital ultrawideband receiver," *IEEE Trans. Wireless Communications*, pp. 502-510, Mar. 2003.

[45] Y. Oh and B. Murmann, "A Low-Power, 6-bit Time-Interleaved SAR ADC Using OFDM Pilot Tone Calibration," *Proc. CICC*, pp. 193-196, Sep. 2007.