

# Low-Power SAR ADCs: Basic Techniques and Trends

PIETER HARPE <sup>1</sup> (Senior Member, IEEE)

(Invited Paper)

Electrical Engineering Department, Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands

CORRESPONDING AUTHOR: P. HARPE (e-mail: p.j.a.harpe@tue.nl)

This work was supported by the Dutch Research Council (NWO) under Project 16594.


**ABSTRACT** With the advent of small, battery-powered devices, power efficiency has become of paramount importance. For analog-to-digital converters (ADCs), the successive approximation register (SAR) architecture plays a prominent role thanks to its ability to combine power efficiency with a simple architecture, a broad application scope, and technology portability. In this review article, the basic design challenges for low-power SAR ADCs are summarized and several design techniques are illustrated. Furthermore, the limitations of SAR ADCs are outlined and hybrid architecture trends, such as noise-shaping SAR ADCs and pipelined SAR ADCs, are briefly introduced and clarified with examples.

**INDEX TERMS** Analog-to-digital converter (ADC), low power, noise shaping, successive approximation, switched-capacitor circuit.

## I. INTRODUCTION

THE POWER efficiency of electronic devices has become an increasingly important criterion thanks to the advent of small, mobile, or even implantable devices where the available energy is limited. For analog-to-digital converters (ADCs), this has led to the development of architectures, circuit implementations, and algorithms, that together achieve the best power efficiency for a required performance target in terms of resolution and bandwidth. Fig. 1 shows an overview of the energy per conversion versus the resolution [in terms of signal-to-noise-and-distortion ratio (SNDR)] of ADCs, where various architectures are highlighted, using the data published at the ISSCC and VLSI Symposium as of 1997 [1]. The FOMS trend-line indicates designs with an equal power efficiency according to the Schreier-based Figure-of-Merit as defined in [1]. It can be noted that different ADC architectures (such as successive approximation register (SAR) and Delta-Sigma) may achieve a similar FOMS, albeit at different resolutions and, thus, at different power levels. The SAR ADC architecture is particularly efficient at modest resolutions (SNDR from approximately 50–80 dB), but is also holding up well against alternatives at lower resolutions (20–50 dB) and higher resolutions (up to approximately 100 dB).


To further illustrate the broad application scope of SAR ADCs, Fig. 2 visualizes the same data [1], but now in




**FIGURE 1.** Energy per conversion versus SNDR for various ADC architectures. Data from [1].

terms of maximum input frequency versus SNDR. It can be seen that SAR converters (sometimes combined with time-interleaving or noise-shaping) cover almost the entire ADC performance space, with the exception of the very high-resolution corner, which is still exclusive to Delta-Sigma ADCs. Thanks to the broad application scope of SAR converters, their power efficiency, their relatively straightforward architecture, and their portability to newer technology nodes, it has become an often-used ADC architecture for a wide variety of applications, from low-power sensor interfaces to high-speed wireless receivers.

In this review article, partially based on [2], the focus will be on SAR ADCs operating at a relatively low sampling



**FIGURE 2.** Maximum input frequency versus SNDR for various ADC architectures. Data from [1].



**FIGURE 3.** Basic SAR ADC architecture.

frequency (typically in the kHz to tens-of-MHz range), where the main tradeoffs are between power consumption on the one hand, and SNDR (i.e., noise and distortion) on the other hand. For higher sampling frequencies, this tradeoff between power and SNDR is also present, but it will gradually become dominated by dynamic effects, such as settling time, metastability, and jitter. First, the basic design challenges for such low-power SAR ADCs are summarized and several design techniques are illustrated in Section II. Furthermore, the limitations of SAR ADCs are outlined and hybrid architecture trends, such as noise-shaping SAR (NS-SAR) ADCs and pipelined SAR ADCs, are briefly introduced and clarified with examples in Section III. It is important to note that it is by no means possible to give a complete overview within this article and, thus, selected examples are used only for illustration. A brief conclusion is drawn in Section IV.


## II. BASIC SAR ADC TECHNIQUES

### A. BASIC PRINCIPLE OF SAR ADC

A basic SAR ADC (Fig. 3) has four functions: 1) a Sample&Hold block; 2) a comparator; 3) a feedback DAC; and 4) digital logic. After the analog input signal  $V_{in}$  is sampled, the logic performs a binary-search algorithm to determine the  $N$  bits of the output code one by one, based on the decisions made by the comparator. The feedback DAC is consecutively tuned toward  $V_{in}$ . After  $N$  cycles,  $D_{out}$  is an  $N$ -bit digital representation of  $V_{in}$ .

### B. CHARGE-REDISTRIBUTION SAR ADC

While there are different methods of implementing the generic SAR architecture, the most commonly applied solution is the *Charge-Redistribution SAR ADC* (Fig. 4). In this case, the DAC is implemented as an array of binary-scaled capacitors. The array is used to sample  $V_{in}$  first, and during the SAR conversion process, by means of switches, charge redistribution takes place to iteratively determine a new residue voltage  $V_{res}$  which is then quantized by the



**FIGURE 4.** Charge-redistribution SAR ADC architecture.


comparator. Note that all signals in this architecture are in the voltage domain. This architecture is attractive because all components can be made in modern technologies and tend to scale well. Capacitors also tend to have good matching (compared to resistors or transistors), which is beneficial for the linearity. Furthermore, the power consumption of all blocks can be dynamic in this architecture, which enables the inherent scalability of power consumption versus sampling frequency.

### C. OVERALL TRADEOFFS

Before looking into the individual components of a SAR ADC, this section reviews some of the system-level choices in terms of optimizing the power efficiency of the ADC. The main tradeoffs in an ADC design are between power consumption, sampling rate (or signal bandwidth), noise, and linearity (as also identified by typical ADC FOMs). Note that there are many other criteria that are also highly relevant (such as chip area, input impedance, supply rejection, etc.), but the focus will be on the former items. For an ADC with dynamic power consumption, the ratio between power consumption and sampling rate will be constant as long as the design is not pushed to either very high speed (which gives a power penalty to force maximum speed) or to very low speed (which gives a power penalty due to leakage). Hence, for low-power SAR ADCs at modest speed, the main tradeoffs will be between power versus noise and linearity.

For the switched-capacitor DAC (Fig. 4), either noise, linearity, or technology limitations might form the bottleneck that ultimately determines its power consumption [3]. In terms of noise, the sampling switch results in a  $kT/C$  noise term, while the switches of the DAC also contribute noise to the DAC output. In terms of linearity, both capacitor mismatch and switch nonidealities cause distortion. To improve the noise and/or linearity of the DAC, typically larger capacitors and better switches are required, both resulting in a higher power consumption. On the other hand, if the target resolution is low, these requirements are so relaxed that the components (capacitors and switches) can be downsized substantially. Ultimately, this downsizing might be constrained by technology limitations (minimum  $W$  and  $L$  for instance) rather than noise and linearity requirements, resulting in an over-design and, hence, power penalty.

The power consumption of the comparator is usually set by noise requirements or technology limitations. Being an active circuit, the comparator adds noise to the system, which



**FIGURE 5.** Leakage in a SAR ADC.


can be counteracted by increasing its power consumption. Like the DAC, technology limitations will play a role at low resolutions. Since the comparator only performs a sign detection, it does usually not add distortion to the system. Note that a comparator offset just resolves into an ADC offset without causing distortion. However, there are some secondary effects that could indirectly still produce some distortion, such as: memory effects, dynamic offset changes due to common-mode variations, and gate leakage.

Finally, the consumption of the logic is only limited by technology, since it does not add noise or distortion as the logic is a digital function.


Overall, for low-resolution SAR ADCs, the technology limits will play a prominent role, and all three blocks of the SAR ADC need to be optimized accordingly. For higher resolutions, the bottleneck will shift toward noise and/or linearity and, thus, the comparator and/or the DAC will dominate overall efficiency. A typical approach to minimize power consumption for low-speed low-resolution applications is to lower the supply voltage, as this reduces both the leakage consumption and the active consumption (assuming it is not limited by noise). A reduced supply further has the benefit that it shifts the transistors in the comparator toward subthreshold, where they achieve higher  $g_m/I_D$ , resulting in better noise efficiency. However, at the same time, this will substantially reduce the maximum speed of operation and may increase the impact from PVT variations. Also, the ON-OFF ratio of switches tends to degrade at lower supplies, making the implementation of the switches more cumbersome. A final challenge at the low-speed operation is leakage, in particular in advanced technology nodes where it tends to be more prominent. First, leakage will increase the power consumption. This can be partially overcome with, e.g., a custom logic design or a specific low-leakage idle-state [4] or power gating [5]. Second, leakage of the sampling switches or comparator gates (Fig. 5) will cause a drift of  $V_{res}$  during the conversion process, usually resulting in distortion.

#### D. COMPARATOR DESIGN

Most comparator architectures can be decomposed in a preamplifier and a latch (Fig. 6), often operated at a CLK signal to make its operation (and power consumption) dynamic. Because of the signal amplification in the preamplifier, the latch will operate faster, its noise requirement will be relaxed and, thus, its power consumption can be



**FIGURE 6.** Typical dynamic comparator architecture composed of preamplifier and latch.



**FIGURE 7.** (a) Dynamic preamplifier and (b) timing diagram [6].

reduced. Since the preamplifier has the most stringent noise requirement, it tends to dominate the comparator's power consumption. On the other hand, the latch is often critical for the operational speed, since it suffers from metastability for small input signals and, thus, will dominate the worst case overall decision time. This is particularly critical for high-frequency converters, while for low-frequency designs, one can often deal with this quite easily by adding sufficient timing margins. Since the focus of this article is on efficiency at low-frequency operation, a couple of preamplifiers aiming for improved power efficiency (i.e., lower power for the same noise requirement) will be described next.

A basic dynamic preamplifier [6], together with a sketch of its operation in time, is given in Fig. 7. When CLK is low, there is no power consumption and the capacitances  $C_p$  (often just parasitics) at the output nodes are preset to  $V_{DD}$ . The amplification is initiated on the rising edge of CLK, and the capacitors are gradually discharged by the differential pair current, meanwhile amplifying the differential input signal. Because the amplified signal is only available temporarily [see Fig. 7(b)], it is important that the latch captures the amplifier output in time, but preferably as late as possible, as this maximizes the gain. When CLK is turned low again, the system resets to be ready for the next cycle. Because the energy consumption of a dynamic circuit depends on  $CV_{DD}^2$ , the consumption of this preamplifier approximates  $2C_pV_{DD}^2$ .

In the previous design, the amplification happens only on the falling edge of  $V_{out+,-}$ . In a bidirectional topology [7], both falling and rising edges are used for amplification. As shown in Fig. 8, a pMOS input stage is activated first to start amplification. A level detector at  $(1/2)V_{DD}$  (not shown) will then disable the pMOS stage and enable an nMOS input stage to continue amplification, but now on the falling edge. As a result, compared to the previous preamplifier, a similar overall gain can be achieved, but capacitors  $C_p$  are only charged up to  $(1/2)V_{DD}$ , which saves approximately  $2\times$  the

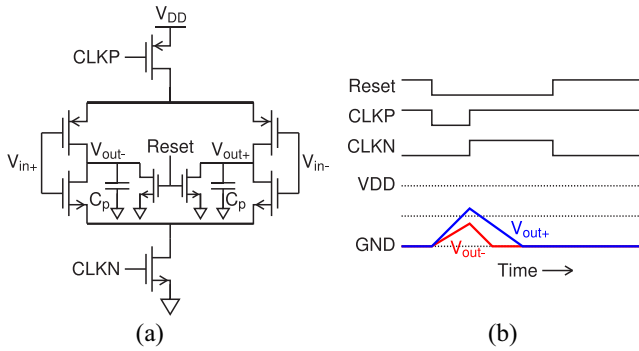


FIGURE 8. (a) Bi-directional dynamic preamplifier and (b) timing diagram [7].




FIGURE 9. (a) Dynamic bias preamplifier and (b) timing diagram [8].

energy, at the cost of some overhead in the level detector and clock controls.

A third option, similar to the first preamplifier but with a series tail capacitor  $C_t$ , is given in Fig. 9 [8]. As  $C_t$  is gradually charged during amplification, it will, at some point, quench the differential pair and, thus, the amplification will stop.  $C_t$  should be designed such that the amplification will stop after the latch has taken over the preamplified signal. Since  $C_p$  is only partially discharged, this will also save power like the bidirectional architecture, while the overhead in this circuit is limited.

### E. DAC DESIGN

The design of the switched-capacitor DAC has two main challenges: 1) the design of switches with sufficient linearity (and bandwidth) and 2) the implementation of a capacitor array that is efficient in power (and area) with sufficient linearity. A particular challenge for the switches is that the ON-OFF resistance ratio tends to degrade in scaled technologies. This is even further exacerbated if the supply is lowered to save power. As a result, techniques are needed to lower the on-resistance and, sometimes, also to increase the off-resistance. To lower the on-resistance, clock boosting [9] or bootstrapping techniques [10] are commonly used. To minimize channel leakage of the sampling switch, for instance, a double switch topology [11] can be used for improved isolation. A further choice is the selection of reference voltages: voltages at/near  $V_{DD}$  and  $V_{SS}$  can be switched relatively easily with a simple pMOS or nMOS transistor,




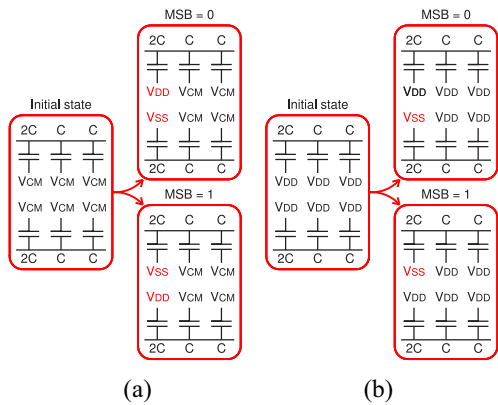
FIGURE 10. (a) Single-ended switched-capacitor DAC and (b) equivalent switch implementation.

respectively. However, references near mid-supply tend to be more cumbersome as they may require boosting techniques. Therefore, a design as in Fig. 10(a), which only uses  $V_{DD}$  and  $V_{SS}$  as references, simplifies the design of the switches: the reference switches are in fact equivalent to digital inverters [Fig. 10(b)], and only the input sampling switches are challenging in terms of signal levels.


The power consumption of the capacitor array can be described as  $\alpha f_s C_s V_{DD}^2$ , where  $\alpha$  is the activity factor (i.e., it depends on how the capacitors are switched exactly, the so-called *switching scheme*),  $f_s$  is the sampling rate,  $C_s$  is the total DAC's capacitance, and  $V_{DD}$  is the supply (acting as the reference). Assuming that the sampling rate is already given, lowering  $V_{DD}$  is one straightforward option (as discussed earlier), but this has a limited range and, at some point, this will limit the SNR, since the signal level also goes down proportional to  $V_{DD}$ . Thus, the most relevant options to improve efficiency in the DAC are the design of the switching scheme and the sizing of the capacitors.

Over the years, many switching schemes were proposed [12], mostly to improve efficiency, but also to improve linearity or to reduce the number of required capacitors. For illustration, two basic schemes ([13] and [14]) are illustrated for the first switching event (the MSB decision). For the first scheme [Fig. 11(a)], the MSB capacitors (denoted by  $2C$  in the figure) are switched from  $V_{CM}$  to  $V_{DD}$  and  $V_{SS}$  or vice-versa, dependent on the MSB decision. A drawback of this scheme is that it requires three references, which complicates the switch and logic design. However, as advantages, it maintains a steady common-mode level during the conversion process and results in a good power efficiency. The second scheme [Fig. 11(b)] has a pseudo-differential operation. The differential voltage steps made in this design are identical to the first scheme (so functionally the network is equivalent), but this scheme results in common-mode variations during the conversion process and has a lower power efficiency in terms of capacitive charging. However, it requires only two reference voltages and, thus, simplifies the switches and control signals. Therefore, it depends on the ratio of power spent in the capacitor charging process versus in the switch drivers to know which scheme is more efficient in practice.

Besides the switching scheme, the second design challenge is the (layout) implementation of the DAC capacitors. On the one hand, larger capacitors reduce noise and mismatch, but on the other hand, this also increases the power consumption. However, based on back-of-the-envelope estimations [2], [3], it can be noted that low-resolution binary-scaled capacitor




**FIGURE 11.** (a) Merged capacitor switching scheme [13] and (b) monotonic switching scheme [14].




**FIGURE 12.** MIM unit capacitor and array of nine elements.

arrays or often limited by technology limitations (i.e., the smallest capacitance that can be implemented in a practical way) rather than noise or mismatch requirements. To illustrate the technology limitations, a few layout sketches are discussed next. Conventionally (especially in somewhat older mixed-signal technology nodes), MIM capacitors were used to implement the capacitor array. However, while these capacitors are suitable in terms of density and matching if large capacitor values are required, this is less so if the unit capacitor (i.e., the smallest capacitor in the DAC array) has to be minimized. As illustrated in Fig. 12, in such a case the MIM structure tends to be area inefficient due to the relatively large overhead of wiring inside such a unit. Moreover, space is required between the various elements when placed in an array due to DRC rules and required space for interconnect. As a result, this type of capacitor is highly area inefficient if small units are used.

An improvement, commonly used in more advanced technology nodes where the metal pitch is small, is to use MOM capacitors, as sketched in Fig. 13 [4]. While the structure can be drawn in different ways, the key advantages are that metal spacing rules are small, enabling to make small-area and small-value capacitors, and allowing efficient placement in an array. In particular since the DAC output node is shared, the capacitors in fact can be overlaid in an array placement, simplifying the interconnect and reducing area at the same time. To implement a binary-scaled array, the identical units are usually grouped in binary-scaled sets, with a common-centroid approach to compensate potential gradients. By making a custom-designed unit capacitor



**FIGURE 13.** MOM unit capacitor and 3-bit array [4].



**FIGURE 14.** 3-bit delta-length capacitor array [15].


layout, sub-fF units can be designed, but one has to rely on parasitic extractions to determine the capacitor values and possibly on experimental verification to determine matching performance.

As the third example, a delta-length capacitor approach was presented in [15] (Fig. 14). This structure also uses MOM capacitor fingers, but each effective capacitance is determined by the difference of two similar strips of slightly different length. By scaling the length difference in binary steps ( $\Delta, 2\Delta, \dots$ ), binary scaling can be achieved. This binary scaling can be accurate, since the parts of the capacitor strip that do not scale with the length (i.e., the “head” and “tail” of each strip) are exactly compensated by the subtracted capacitor element, such that the effective capacitance is only defined by the length difference. As advantage compared to regular MOM capacitors, this delta-length structure can more easily create small unit elements, and it uses less components in the layout and requires less interconnect, resulting in a reduced chip area.


## F. LOGIC, ALGORITHM, AND LAYOUT DESIGN

In this section, three topics related to the SAR logic are discussed: 1) synchronous versus asynchronous logic; 2) the search algorithm; and 3) impact of the layout on efficiency.

Traditionally, a SAR ADC requires a master clock that is at least  $N \cdot f_s$ , where  $N$  is the resolution and  $f_s$  the sampling rate, such that all internal operations (comparator cycles, storage in register, etc.) can be synchronized to the master clock. At present though, most converters use asynchronous (also called: self-synchronized) logic. In such a case, a master clock is still provided, but it is at a frequency of  $f_s$  and only determines when the converter is tracking the input signal, and when it is converting that signal. The entire conversion process is then done asynchronously with respect to the master clock. To ensure proper timing, e.g., an internal oscillation loop can be used [16], or a delay-line [6]. The main advantage of asynchronous timing is that it reduces the required




**FIGURE 15.** (a) Binary search and (b) search with redundancy.




**FIGURE 16.** (a) Sub-binary radix [21], (b) overlapping binary segments [19], and (c) modified search logic [18]. Note: capacitor values in the figure are for illustration only.

external clock frequency, and that internal timing can be adjusted to the actual circuit operation, typically resulting in a faster conversion. In particular, the analysis in [17] shows that an asynchronous SAR ADC can deal much better with comparator metastability and, thus, reach a faster conversion as compared to a synchronous design. Furthermore, at very low-speed operation, asynchronous timing also reduces the impact of leakage [4].

A second popular technique in SAR ADCs is to use redundancy in the search algorithm [18]. As illustrated in Fig. 15(a), a conventional binary-search algorithm divides the input range in two equal parts after the MSB decision (and likewise for the subsequent bit decisions). Unfortunately, if a decision mistake is made due to, e.g., comparator noise or incomplete DAC settling, the signal will go out of range, which results in distortion. With redundancy [Fig. 15(b)], the subranges are adjusted such that they overlap near the decision level. In that case, decision errors (up to the amount of overlap) can be tolerated, as the signal will now remain inside the search space for the subsequent conversions. As advantage, this relaxes the requirements at the decision moment, which can be exploited to increase speed [19] or improve power efficiency [20]. As disadvantage, due to the overlap, more than  $N$  bit-cycles are required to obtain  $N$  bit of information, and some changes to the logic and/or capacitor array are required to implement the nonbinary search. As example, Fig. 16(a) shows a capacitor array with redundancy by using a sub-binary radix. In this way, each step of the algorithm has some overlap, at the cost of relatively complex logic to combine the nonbinary numbers to a final output code. Fig. 16(b) shows an alternative using two binary-scaled segments with redundancy between the two segments. Redundancy is now only available up to a certain point, but the reconstruction logic is easier since most weights are binary scaled. As a third alternative, one could also modify the logic [Fig. 16(c)] to control the capacitors in an arbitrary way to create redundancy. While this simplifies



**FIGURE 17.** Power breakdown of an exemplary SAR ADC [4] based on post-layout simulations.





**FIGURE 18.** Flexible 10-bit SAR ADC [22] with (a) input swapping and (b) stepwise charging.

the capacitor array, it may complicate the logic and capacitor control signals.

While strictly not related to the logic alone, Fig. 17 illustrates the simulated power breakdown of a low-resolution (8 bit) SAR ADC in 90-nm CMOS. First, it can be noted that the logic consumes more power than the DAC or the comparator. That is because, for this low resolution, technology limitations are more critical than noise and linearity requirements. As a result, the logic overhead is high. On top of that, for the same reason, layout parasitics are actually dominating the overall consumption. Most of these losses are due to interconnect between the logic and the switched-capacitor DAC, and inside the logic itself. These results show that for low-resolution ADCs, it is important to minimize logic complexity (e.g., no redundancy, no calibration, and possibly optimized logic cells), and to minimize interconnect (e.g., use a delta-length DAC and compact layout). For higher resolutions the logic and parasitic losses will play a less prominent role as noise and matching requirements in the comparator and DAC will rather quickly dominate the overall consumption.

## G. EXAMPLES FROM LITERATURE

After discussing the basic SAR building blocks, two low-power SAR ADC examples are reviewed here. The flexible 10-bit SAR ADC [22] in Fig. 18 uses various techniques that were described before to save power: a reduced supply, a dynamic-bias comparator, small unit capacitors, a split monotonic switching scheme, and asynchronous timing with a delay line. On top of that, it uses input-swapping to swap the input signal dependent on its polarity, thereby improving driveability. Next to that, stepwise charging of the capacitors is used [6] to reduce the energy that is needed to charge capacitors by using intermediate voltage steps.


**FIGURE 19.** 12-bit SAR-assisted SAR ADC [23].

**FIGURE 20.** Zoom ADC: coarse SAR ADC with fine Delta-Sigma ADC [25].

As a second example, Fig. 19 uses an assisting ADC [23]. Since the main 12-bit SAR ADC requires relatively large capacitors, especially the MSB switching activities would be rather costly. Instead, a small 5-bit assisting ADC is used to take the 5 MSB decisions, the result of which is immediately preloaded into the main ADC. This reduces the number of switching events in the main DAC and, thus, saves substantial power. Redundancy is needed to overcome possible mismatches between the two ADCs. A reduced supply voltage was also used, and a dedicated clock boosting technique and switching scheme were developed to further improve the performance.



### III. SAR ADC LIMITATIONS AND HYBRID EXAMPLES

#### A. SAR ADC LIMITATIONS

Besides the improvement in power efficiency as illustrated in Section II, SAR ADCs have also shown extensive progress toward higher resolution and higher speed. Nonetheless, at some point, it might become beneficial to change the architecture due to some of the inherent disadvantages of the SAR ADC. In terms of resolution: when aiming to improve SNDR, both noise and linearity have to be improved, which does come at a penalty in terms of chip area and power consumption. Since the SAR ADC is mostly a passive structure, it does not benefit from amplification when processing the signal (as opposed to, e.g., a pipelined or Delta-Sigma ADC). In this context, NS-SAR architectures were developed [24] which combine the efficiency of a SAR ADC with noise-shaping properties to extend the resolution. In terms of speed: since a SAR ADC performs all steps sequentially, it ultimately becomes necessary to apply time interleaving or pipelining to further advance the sampling rate. This section briefly reviews NS-SAR ADCs as well as pipelined SAR ADCs, which are two popular architectures to improve resolution and speed, respectively.


#### B. SAR ADCS WITH NOISE SHAPING

SAR ADCs and noise-shaping can be combined in different ways. In a zoom ADC [25] (Fig. 20), a SAR ADC coarsely quantizes the input signal, while a Delta-Sigma ADC performs fine quantization within the boundaries set by the SAR ADC (Fig. 21). By co-integrating the DACs and using data weighted averaging (DWA), a high linearity can be achieved as well, and mismatches between coarse and fine conversion are alleviated. As a result, this architecture can achieve a very high resolution, while the Delta-Sigma ADC is relatively relaxed (1-bit quantizer, second-order loop filter).

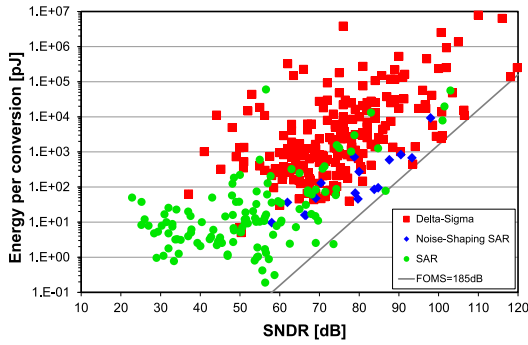

**FIGURE 21.** Example timing diagram of a zoom ADC.

**FIGURE 22.** Conceptual diagram of an NS-SAR ADC.

A second alternative combination is the NS-SAR ADC. Conceptually (Fig. 22), the SAR ADC quantizes the input signal while the residue is integrated in a loop filter and added to the input to perform noise shaping. This architecture is similar to a Delta-Sigma ADC, with a specific architectural implementation. However, several implementation choices are often quite different. For instance, the NS-SAR ADC tends to use a simplified loop filter to keep power efficiency high. In some cases, the loop filter is fully passive (i.e., made with switches and capacitors) [26], while in other cases an active filter is made, typically based on a simplified yet highly efficient amplifier [27]. While the first category benefits from lowest power consumption, the additional  $kT/C$  noise terms either impact the SNR, or otherwise relatively large capacitors might be required. The active filters enable a higher SNR at the cost of increased power consumption. A review of active filters can be found in [27]. Also, since the SAR quantizer can have a reasonably high resolution, both the filter and the oversampling ratio can be relatively modest while still achieving high resolution. While noise shaping improves the SNR of the SAR ADC, it does not improve the linearity which is typically dominated by DAC mismatch. To deal with this, mismatch error shaping (MES) [26], [27], [28] has become a popular technique to shape mismatch errors out of band at the cost of a digital reconstruction function.


For illustration, a recent example of an NS-SAR ADC is shown in Fig. 23 [26]. At the end of the regular SAR conversion, a cluster of four integration capacitors is connected to the DAC output. By charge sharing, this passively integrates the residue on these  $C_{int}$  capacitors. In the next conversion phase, these integration capacitors are now reconnected in



**FIGURE 23.** NS-SAR ADC [26]: (a) residue integration phase and (b) SAR conversion phase.




**FIGURE 24.** Energy per conversion versus SNDR for noise shaping and SAR ADCs. Data from [1].




**FIGURE 25.** Maximum input frequency versus SNDR for noise shaping and SAR ADCs. Data from [1].

series to the DAC, such that the integrated residue signal is amplified and superimposed onto the DAC signal going to the comparator. Overall, this simple implementation with switches and capacitors thus performs residue integration, passive voltage amplification, and addition, to implement noise shaping. Besides that, this design also includes second-order MES to improve linearity, resulting in an overall SNDR of 90 dB for this ADC.


In the past decade, NS-SAR ADCs have become rather popular. To review their present performance and application scope, Figs. 24 and 25 show the energy per conversion and input frequency versus the resolution for SAR ADCs, NS-SAR ADCs, and Delta-Sigma ADCs. As can be seen, the efficiency of NS-SAR ADCs is—at present—in line with (but not better than) SAR ADCs or Delta-Sigma ADCs. However, they nicely form a natural transition between SAR ADCs and Delta-Sigma ADCs in terms of resolution.



**FIGURE 26.** Pipelined SAR ADC with FIA [30].



**FIGURE 27.** Energy per conversion versus SNDR for pipelined (SAR) ADCs. Data from [1].



**FIGURE 28.** Maximum input frequency versus SNDR for pipelined (SAR) ADCs. Data from [1].

### C. PIPELINED SAR ADCs

Pipelined ADCs are known for their ability to increase speed thanks to pipelining of the conversion, but they are also efficient in terms of noise thanks to their interstage gain. Conventionally, such pipelines are made with low-resolution flash stages, but more recently, it became popular to pipeline (two) SAR ADCs [29]. This results in an efficiency improvement for two reasons: first, a SAR stage is relatively efficient, and second, by reducing the number of stages, the number of amplifiers is reduced. As a disadvantage, a medium-resolution SAR stage tends to be slower than a low-resolution flash stage, so the maximum speed may be reduced. As an example, Fig. 26 show a recent example of a pipelined SAR ADC, composed of a 7-bit and a 9-bit SAR ADC [30]. The key to overall efficiency is the implementation of the interstage residue amplifier. While there are many efficient topologies [31], this design uses a so-called floating-inverter amplifier (FIA), which is highly efficient.

The performance of present pipelined SAR ADCs is compared against regular pipelined ADCs in Figs. 27 and 28. As expected, the pipelined SAR ADC indeed has a clear advantage in terms of efficiency, while it may have a penalty in the maximum possible frequency for a given resolution.

### IV. CONCLUSION

In the past decades, efficiency, resolution, and speed of SAR ADCs have improved substantially, thanks to improvements



in circuits, algorithms, architectures, and technology. This makes SAR ADCs (and derived hybrid architectures) applicable to a very wide range of applications. While SAR ADCs are still improving further, present research mostly focusses on hybrid architectures, such as the NS-SAR ADC. Extrapolating from Fig. 1, SAR ADCs could still improve substantially at low resolutions ( $<50$  dB), where the designs are presently still limited by technology rather than noise or distortion.

## REFERENCES

- [1] B. Murmann. "ADC performance survey 1997–2021." Jun. 2021. [Online]. Available: <http://www.stanford.edu/~murmann/adcsurvey.html>
- [2] P. Harpe, "High-precision and low-power ADCs," in *Proc. ISSCC Short Course*, 2022.
- [3] P. Harpe, H. Li, and Y. Shen, "Low-power SAR ADCs: Trends, examples and future," in *Proc. ESSCIRC*, 2019, pp. 25–28.
- [4] P. J. A. Harpe et al., "A 26  $\mu$ W 8 bit 10 MS/s asynchronous SAR ADC for low energy radios," *IEEE J. Solid-State Circuits*, vol. 46, no. 7, pp. 1585–1595, Jul. 2011.
- [5] K. Pelzers, H. Xin, E. Cantatore, and P. Harpe, "A 2.18-pJ/conversion, 1656- $\mu$ m<sup>2</sup> temperature sensor with a 0.61-pJ-K<sup>2</sup> FoM and 52-pW stand-by power," *IEEE Solid-State Circuits Lett.*, vol. 3, pp. 82–85, 2020.
- [6] M. van Elzaker, E. van Tuijl, P. Geraedts, D. Schinkel, E. Klumperink, and B. Nauta, "A 1.9 $\mu$ W 4.4fJ/conversion-step 10b 1MS/s charge-redistribution ADC," in *ISSCC Dig. Tech. Papers*, Feb. 2008, pp. 244–245.
- [7] M. Liu, K. Pelzers, R. van Dommele, A. van Roermund, and P. Harpe, "A 106nW 10 b 80 kS/s SAR ADC with duty-cycled reference generation in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 51, no. 10, pp. 2435–2445, Oct. 2016.
- [8] H. S. Bindra, C. E. Lokin, D. Schinkel, A. Annema, and B. Nauta, "A 1.2-V dynamic bias latch-type comparator in 65-nm CMOS with 0.4-mV input noise," *IEEE J. Solid-State Circuits*, vol. 53, no. 7, pp. 1902–1912, Jul. 2018.
- [9] T. B. Cho and P. R. Gray, "A 10 b, 20 Msample/s, 35 mW pipeline a/D converter," *IEEE J. Solid-State Circuits*, vol. 30, no. 3, pp. 166–172, Mar. 1995.
- [10] A. M. Abo and P. R. Gray, "A 1.5-V, 10-bit, 14.3-MS/s CMOS pipeline analog-to-digital converter," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 599–606, May 1999.
- [11] S.-E. Hsieh and C.-C. Hsieh, "A 0.44-fJ/conversion-step 11-bit 600-kS/s SAR ADC with semi-resting DAC," *IEEE J. Solid-State Circuits*, vol. 53, no. 9, pp. 2595–2603, Sep. 2018.
- [12] C.-C. Hsieh, "Nano-watt ultra-low-voltage SAR ADC design," in *Proc. ISSCC Forum 6*, 2018.
- [13] V. Hariprasath, J. Guerber, S.-H. Lee, and U.-K. Moon, "Merged capacitor switching based SAR ADC with highest switching energy-efficiency," *Electron. Lett.*, vol. 46, no. 9, pp. 620–621, Apr. 2010.
- [14] C.-C. Liu, S.-J. Chang, G.-Y. Huang, and Y.-Z. Lin, "A 10-bit 50-MS/s SAR ADC with a monotonic capacitor switching procedure," *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 731–740, Apr. 2010.
- [15] P. Harpe, "A compact 10-b SAR ADC with unit-length capacitors and a passive FIR filter," *IEEE J. Solid-State Circuits*, vol. 54, no. 3, pp. 636–645, Mar. 2019.
- [16] S.-W. M. Chen and R. W. Brodersen, "A 6-bit 600-MS/s 5.3-mW asynchronous ADC in 0.13- $\mu$ m CMOS," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2669–2680, Dec. 2006.
- [17] A. Waters, J. Muhlestein, and U.-K. Moon, "Analysis of metastability errors in asynchronous SAR ADCs," in *Proc. ICECS*, Dec. 2015, pp. 547–550.
- [18] F. Kuttner, "A 1.2V 10b 20MSample/s non-binary successive approximation ADC in 0.13 $\mu$ m CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2002, pp. 176–177.
- [19] C.-C. Liu et al., "A 10b 100MS/s 1.13mW SAR ADC with binary-scaled error compensation," in *ISSCC Dig. Tech. Papers*, Feb. 2010, pp. 386–387.
- [20] V. Giannini, P. Nuzzo, V. Chironi, A. Baschiroto, G. Van der Plas, and J. Craninckx, "An 820 $\mu$ W 9b 40MS/s noise-tolerant dynamic-SAR ADC in 90nm digital CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2008, pp. 238–239.
- [21] K.-H. Chang and C.-C. Hsieh, "A 12b 10MS/s 18.9fJ/conversion-step sub-radix-2 SAR ADC," in *Proc. VLSI-DAT*, 2016, pp. 1–4.
- [22] H. S. Bindra, A.-J. Annema, S. M. Louwmsa, and B. Nauta, "A 0.2–8 MS/s 10b flexible SAR ADC achieving 0.35–2.5 fJ/Conv-step and using self-quenched dynamic bias comparator," in *Proc. VLSI Symp.*, 2019, pp. C74–C75.
- [23] Y.-S. Hu, K.-Y. Lin, and H.-S. Chen, "A 510nW 12-bit 200kS/s SAR-assisted SAR ADC using a re-switching technique," in *Proc. VLSI Symp.*, 2017, pp. C238–C239.
- [24] J. Fredenburg and M. Flynn, "A 90MS/s 11MHz bandwidth 62dB SNDR noise-shaping SAR ADC," in *ISSCC Dig. Tech. Papers*, Feb. 2012, pp. 468–469.
- [25] S. Karmakar, B. Gönen, F. Sebastiano, R. van Veldhoven, and K. A. A. Makinwa, "A 280  $\mu$ W dynamic zoom ADC with 120 dB DR and 118 dB SNDR in 1 kHz BW," *IEEE J. Solid-State Circuits*, vol. 53, no. 12, pp. 3497–3507, Dec. 2018.
- [26] J. Liu, X. Wang, Z. Gao, M. Zhan, X. Tang, and N. Sun, "9.3 A 40kHz-BW 90dB-SNDR noise-shaping SAR with 4x passive gain and 2<sup>nd</sup>-order mismatch error shaping," in *ISSCC Dig. Tech. Papers*, Feb. 2020, pp. 158–159.
- [27] H. Li, Y. Shen, H. Xin, E. Cantatore, and P. Harpe, "A 7.3- $\mu$ W 13-ENOB 98-dB SFDR noise-shaping SAR ADC with duty-cycled amplifier and mismatch error shaping," *IEEE J. Solid-State Circuits*, vol. 57, no. 7, pp. 2078–2089, Jul. 2022.
- [28] Y.-S. Shu, L.-T. Kuo, and T.-Y. Lo, "An oversampling SAR ADC with DAC mismatch error shaping achieving 105 dB SFDR and 101 dB SNDR over 1 kHz BW in 55 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 51, no. 12, pp. 2928–2940, Dec. 2016.
- [29] C. P. Hurrell, C. Lyden, D. Laing, D. Hummerston, and M. Vickery, "An 18b 12.5MHz ADC with 93dB SNR," in *ISSCC Dig. Tech. Papers*, Feb. 2010, pp. 378–379.
- [30] X. Tang, X. Yang, J. Liu, W. Shi, D. Z. Pan, and N. Sun, "27.4 A 0.4-to-40MS/s 75.7dB-SNDR fully dynamic event-driven pipelined ADC with 3-stage cascoded floating inverter amplifier," in *ISSCC Dig. Tech. Papers*, Feb. 2021, pp. 376–377.
- [31] K. Bult, M. S. Akter, and R. Sehgal, "High-efficiency residue amplifiers," in *Low-Power Analog Techniques, Sensors for Mobile Devices, and Energy Efficient Amplifiers*. Cham, Switzerland: Springer, 2019.



**PIETER HARPE** (Senior Member, IEEE) received the M.Sc. and Ph.D. degrees from the Eindhoven University of Technology, Eindhoven, The Netherlands, in 2004 and 2010, respectively. He started as a Researcher with Holst Centre/imec, Eindhoven, in 2008. Since then, he has been working on ultralow-power wireless transceivers, with a main focus on ADC research and design. He joined the Eindhoven University of Technology in April 2011, where he is currently an Associate Professor of Low-Power

Mixed-Signal Circuits.

Dr. Harpe is a recipient of the ISSCC 2015 Distinguished Technical Paper Award. He is a Co-Organizer of the yearly workshop on Advances in Analog Circuit Design and an Analog Subcommittee Chair for the ESSCIRC conference. He also served as an ISSCC ITPC Member and an IEEE SCS Distinguished Lecturer.