Digitally Assisted Analog to Digital Converters

Bob Verbruggen

Abstract This chapter discusses how digital assistance can be leveraged in the design of analog to digital converters. Different types of digital assistance are defined, and a few of the possible applications selected for detailed discussion. Finally, an example of an ADC implementation heavily leveraging digital assistance is presented.

1 Introduction

Digital assistance in ADCs can be generally defined as using digital techniques to relax requirements for analog non-idealities. This is not a particularly new concept: as early as 1981 digital correction was used to overcome accuracy limitations in component matching [1]. Since then, technology scaling has massively reduced the power and area cost of digital logic, and the prevalence of digital assistance has increased to match. This has happened to such a degree that in some cases the line between the pure analog architectures and digitally assisted architectures has blurred.

Let us consider for example a conventional 1.5b/stage pipelined converter. Strictly speaking even this architecture is digitally assisted [2]: analog comparator errors are compensated by digitally combining different stage outputs. However, the underlying redundancy technique is so common nowadays that few people would consider this architecture to be overtly "digitally assisted". The scope of this paper will be limited to digital assistance of Nyquist converters that involve clearly defined observation and correction steps for analog non-idealities. A more general overview of digital assistance in data converters is presented in [2].

B. Verbruggen (✉)
Imec, Kapeldreef 75, Leuven, Belgium
e-mail: bob.verbruggen@imec.be

P. Harpe et al. (eds.), *High-Performance AD and DA Converters, IC Design in Scaled Technologies, and Time-Domain Signal Processing*,
DOI: 10.1007/978-3-319-07938-7_2,

Using this definition and the data from [3] the Nyquist converters published at ISSCC from 1997 to 2013 can be divided into two categories, based on whether they use any kind of digital assistance. Figure 1 shows the Walden FoM versus clock speed for these two populations and reveals two interesting trends. First: 40 % of the Nyquist converters published at ISSCC over this time period include digital assistance, which indicates that these techniques are quite common. Second: for clock speeds above 1 MS/s, all of the designs with best Walden FoM leverage digital assistance. Obviously ADC performance cannot be reduced to just the Walden FoM, and many designs achieve excellent performance without any digital assistance, but this trend does lend credence to the idea that digital assistance can be used to lower power consumption in analog to digital converters.

In the remainder of this paper, we will try to show why this is the case. In Sect. 2, we will first briefly define different types of observation and correction. Next, it will be explained how digital assistance can change analog design trade-offs and how this could benefit power consumption in comparators, amplifiers and DACs. In Sect. 4, these ideas are illustrated in a detailed look at the design of a pipelined SAR ADC heavily leveraging digital assistance. Finally, in Sect. 5 conclusions will be drawn.

2 Types of Digital Assistance

A first important classification of digital assistance is based on how analog non-idealities are corrected. In this chapter, digital compensation will refer to a digital mapping of raw output bits to a corrected, final set of output bits, as illustrated in Fig. 2a. On the other hand, digital calibration refers to the use of digital calibration settings to adjust analog circuit parameters and thus directly correct output values of the analog core as shown in Fig. 2b. The relative merits of compensation and calibration depends on the effect to be corrected, and both can be used in conjunction.

Compensation can be done using a simple look-up table to correct for purely static effects [1], or using elaborate digital signal processing to correct for interleaving artefacts, linearity or even self-heating gradients [4]. Digital calibration is applied most commonly to comparator offsets [5], amplifier gains [6] or capacitor values [7], and in all cases involves some digital programmability for said non-idealities.

The choice between compensation and calibration is not a simple one. From the point of view of area and power consumption, calibration is a preferred solution if analog tuning can be implemented with negligible overhead: it avoids the constant digital activity required for compensation logic. However, from the perspective of design-time the flexibility of compensation may be a compelling argument. Some effects, such as DNL degradation, are trivial to calibrate but are extremely difficult to compensate, because they directly affect the quantization of the ADC core. Higher level non-idealities, such as non-linear distortion, which may have many

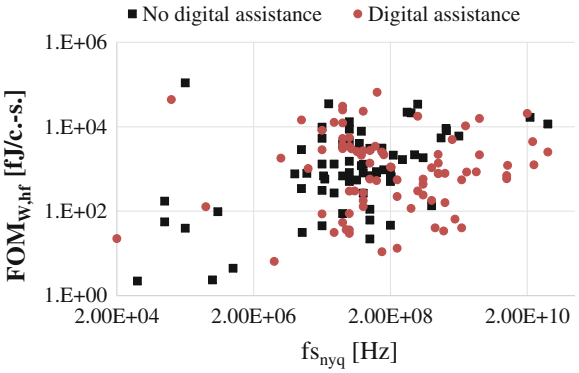

Fig. 1 Nyquist converters published at ISSCC between 1997 and 2013, divided based on use of digital assistance

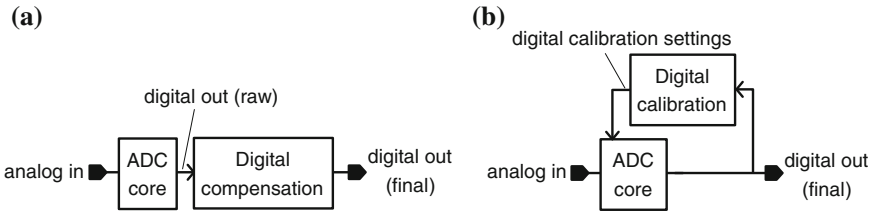

Fig. 2 High-level block diagrams comparing digital compensation (**a**) with digital calibration (**b**)

different analog contributions, are usually difficult to calibrate but comparatively easy to compensate for.

A second important distinction is based on when the ADC core output is observed. In this paper, we will refer to foreground observation if this occurs using a well-known input signal while the ADC operation is interrupted. On the other hand, if observation occurs while the ADC is quantizing an unknown input we will refer to background observation.

The ability to observe non-idealities during continuous ADC operation is a significant advantage to background techniques but this advantage is not without its cost. Since background observation is subject to statistics of the ADC input, it will typically take significantly longer to measure non-idealities with a given confidence level than using foreground techniques. In addition, background observation nearly always requires some assumptions about the ADC input statistics to guarantee convergence. On the other hand, foreground observation is vulnerable to changes of non-idealities after the observation phase, while the ADC is quantizing an unknown input. In many cases, this results in the need to periodically repeat the foreground observation, and thus periods when the ADC is not

quantizing its input. In burst-mode applications, this might be a valid choice, while in others this is unacceptable.

These general methods for observation and correction provide a huge design-space when implementing an ADC with digital assistance. Proper choices in this design-space often require some knowledge about the system integration, in addition to ADC design know-how.

3 Designing with Digital Assistance

In this section we will discuss a few common applications for digital assistance in recent ADCs. We will consider how digital assistance changes design trade-offs for comparators, amplifiers and DACs. We will not discuss digital assistance for interleaved channels, despite its usefulness and popularity.

3.1 Calibration of Comparator Thresholds

Comparators are at the heart of nearly all ADC architectures, as they are required to perform quantization. Their accuracy is limited by offset and noise, but not all architectures are affected equally by these non-idealities. SAR ADCs, for example are generally quite robust to offset since in a conventional SAR ADC a comparator offset does not introduce any non-linearity. However, they are highly sensitive to comparator noise, as this noise is added directly to the input. Redundancy, which is nearly omnipresent in pipelined ADCs and gaining popularity fast in SAR ADCs, can reduce sensitivity to both offset and noise. In flash-based architectures, however, comparator accuracy is still key, and calibration is an extremely useful technique.

Indeed, the stringent requirements on offset can be all but completely avoided by adding threshold calibration. This is illustrated in Fig. 3 for comparators with a dynamic input pair followed by a latch. These comparators are activated on a falling clock edge and generate a rising slope at the drains of the input pair, depending on the input voltage. This slope is then resolved into digital levels by a latch in different configurations [8, 9]. Threshold calibration can be added by controlling the capacitance to the drain of the input pair [10], by controlling the body terminals of the input pair [11] or by controlling the voltage on the gates of a redundant input pair [12], among others [13, 14]. These calibration schemes have different advantages and disadvantages: using load capacitance is low-noise but slightly degrades speed, using body voltage tracks environmental changes reasonably well but requires a separate N-well and using a redundant input pair is probably the highest speed solution but slightly degrades noise.

Assuming comparator offset is calibrated, the final comparator accuracy is limited by noise. In nearly all cases, for a given power consumption and speed

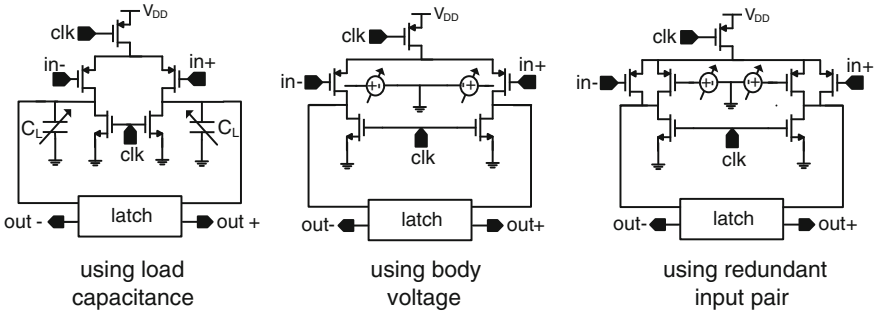

Fig. 3 Popular methods for comparator calibration

comparator noise spread is significantly lower than comparator offset spread, which implies a significant reduction in power consumption. In addition to these power savings, offset calibration typically significantly reduces comparator input capacitance. This is especially useful in flash-based architectures, where many parallel comparators are required, and target speeds are typically high.

Comparator calibration also enables a few architectural changes. One example is the noise-tolerant SAR ADC shown in Fig. 4 [15]. This design uses two offset-calibrated fully dynamic comparators in parallel: one of these is designed for low power and consequently has fairly high noise (HN), whereas the other is optimized for low noise (LN) at a power penalty. In [15] the HN comparator is activated for the first eight SAR cycles, after which the LN comparator resolves two final cycles. By giving the final cycles 1b redundancy with respect to the first eight, the ADC can tolerate a fairly large r.m.s. noise in the HN comparator without an SNR penalty. Since only two low noise comparisons are required instead of nine, comparator power can be significantly reduced. Without offset calibration, this scheme would need far more redundancy to compensate for differences in offset between the two comparators.

Another example is the comparator-controlled SAR ADC proposed in [16] and illustrated for a 4b example in Fig. 5. In this example, the input is tracked on two pseudo-differential 3b DACs and the input of 4 dynamic comparators. At a rising edge of the clock, the tracking switch opens and the first comparator is activated. When this comparator decides, either the positive or negative DAC MSB is discharged and a ready signal is generated. This ready signal is delayed to allow time for DAC settling, and then activates the second comparator. This comparator in turn generates feedback and asynchronously activates the next comparator in line. When all comparators have decided, the outputs of the 4 comparators can simply be latched to obtain the final output code. This arrangement thus implements a conventional SAR algorithm without the need for a controller, but in the absence of comparator calibration, this configuration would be crippled by offset.

Regardless of how comparator offset calibration is implemented and leveraged, a method for controlling the calibration codes is required. The most straightforward choice is foreground calibration: applying the desired threshold at the

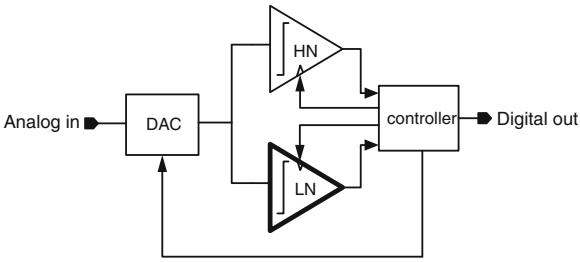

Fig. 4 Simplified block diagram of noise-tolerant SAR ADC

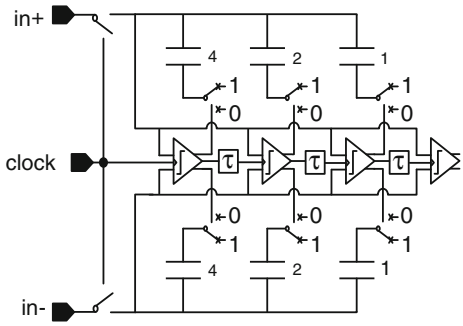

Fig. 5 Simplified block diagram of comparator-controlled SAR

comparator input and changing the calibration code until the comparator outputs as many zeroes as ones [5], or some variation on this scheme [14]. Background schemes also exist but vary wildly depending on architecture: adding a redundant ADC channel and averaging a rotating set of comparator pairs in a flash architecture [17], observing output code density [18, 19] or using redundancy to detect erroneous decisions [18, 19].

In conclusion, comparator threshold calibration can be used to shift the critical accuracy constraint from offset to noise, and thus reduce comparator power consumption and input capacitance. In addition, this calibration enables some architectural changes that are potentially useful. The disadvantage is that most offset calibration methods are sensitive to changes of temperature, common-mode input voltage or supply voltage. This sensitivity thus needs to be handled either through recurring foreground calibration or a background calibration loop for threshold calibrated flash converters to be applicable in a real-world environment.

3.2 Digital Assistance for Amplifiers

Conventional pipeline converters rely on closed-loop high-gain amplifiers to perform residue amplification. Feedback serves to linearize the amplifier and to stabilize the residue gain, but it requires high gain to achieve good accuracy, which is increasingly difficult in deeply scaled CMOS. Alternatives to high gain

amplifiers have been explored: comparator-based switched capacitor circuits [20] or ring amplifiers [21] also offer the benefits of closed-loop settling without some of the drawbacks.

Digital assistance offers a different alternative: in [22] a digital loop tracks gain and distortion of an open-loop single stage amplifier. This implementation uses digital compensation: the digital logic which processes the stage outputs uses adjustable coefficients for linear combination of the stage outputs, as well as adjustable coefficients to correct dominant harmonic distortion components from some of the stage amplifiers.

Digital assistance significantly changes the requirements for residue amplifiers. Indeed, if gain stability and intrinsic linearity are not required, residue amplifier requirements can be reduced to:

- some fairly low amount of gain, for example approximately 12 dB in [6]
- distortion at a level that can be corrected in the digital domain, for example no significant non-third-order distortion components in [22]
- sufficiently low input referred noise
- sufficiently fast amplification.

While the above requirements for speed, gain and linearity can be further relaxed through architectural choices, the input referred noise is a fundamental limitation which currently dictates the lower power limit for amplifiers. To explore this power limit, a number of unconventional amplification paradigms has been proposed: charge-domain pipelines [23, 24], dynamic source follower-based amplifiers [25] or charge-steering amplifiers [6, 26].

While some of these would be applicable at low resolutions without digital assistance, most use digital compensation for gain and distortion. The exception is [5] which uses calibration for gain and avoids the need for distortion correction by limiting input and output range in the architecture choice. Using calibration to reduce distortion is probably impractical: it is not straightforward to implement analog controls guaranteed to fully linearize a circuit.

Background algorithms to observe residue gain and distortion are readily available. For example, by adding pseudo-random dither to the input of the residue amplifier and correlating the ADC back-end output with said dither sequence in the digital domain gain and distortion can be estimated [22]. Methods based on output histogram also exist [27]: these do not require any analog dither injection, but typically do impose more stringent assumptions about the ADC input signal for convergence. Even running at fairly low sample rates these algorithms are certainly fast enough to track most temperature variations, but to obtain an accurate measurement of gain and/or distortion at start-up, a foreground observation is often faster and more practical.

In conclusion, digital assistance can shift the critical requirements for residue amplifiers from gain accuracy and distortion to noise. While these relaxed requirements are certainly beneficial in their own right, they have also enabled exploration of a number of non-conventional amplification paradigms in the quest to find a more beneficial noise-power trade-off.

3.3 Digital Assistance for DACs

SAR and many pipeline architectures rely on matched capacitor or resistor arrays, specifically in feedback DACs. Given their current popularity, we will focus the rest of this discussion on capacitor DACs, but some of the conclusions apply equally to resistive or current DACs.

Assuming a careful common-centroid layout to eliminate systematic mismatch, DAC matching is limited by random effects determined by the well-known Pelgrom model [28]:

$$\sigma\left(\frac{\Delta C}{C}\right) = \frac{A_c}{\sqrt{W \cdot L}} \tag{1}$$

While this formula predicts a fairly straightforward 4× area scaling to improve linearity by a factor 2, in some cases this scaling is anything but straightforward. Indeed, as the size of an array increases, it becomes more sensitive to systematic effects, even while tolerance for such effects decreases. This is especially noticeable in high resolution SAR ADCs: in addition to very accurate matching these also need an extremely large number of units, which further increases area due to unit and routing overhead. In state of the art, there seems to be an intrinsic matching plateau around 10.5b resolution [29]: higher resolution SAR ADCs tend to use techniques to either relax matching requirements [1] or reduce unit count [30].

Both compensation and calibration can be applied to capacitor DACs. Compensation simply involves adjusting digital bit weights to match analog feedback weights, but it requires redundancy to ensure a sufficiently fine quantization grid in the presence of mismatch. In the absence of this redundancy, or to simplify digital logic, calibration [7] is also an option.

A common misconception is that avoiding the matching limit allows sizing of the DAC for kT/C noise, with obvious speed and power consumption benefits. This is true from a certain point of view, but it does assume a fixed capacitance density. In practice the matching constraint of (1) enforces a minimum DAC area, but does not place a constraint on DAC capacitance. Indeed, in vertical MOM capacitors the finger distance can usually be more or less freely chosen. Choosing a greater finger distance thus in theory reduces capacitance without affecting matching. The real benefit of avoiding the matching limits is thus in terms of area, and not in terms of capacitance. Theoretical considerations aside, for high resolution arrays reducing size is often very beneficial as it lowers sensitivity to systematic effects. In addition, the theory indicating that A_c should be the same regardless of capacitance density is not always verified during device modeling. As a result, relying on a theoretical extrapolation of device matching might be somewhat risky. In addition, digital assistance allows switching schemes that are very difficult to design for intrinsic linearity [1].

Background calibration algorithms to detect capacitor mismatch have been extensively applied in pipelined ADCs [27], and could be generalized to cover SAR ADCs. However, foreground calibration might in this case be sufficient, since the primary goal is often to compensate for capacitor mismatch which changes only very slightly as a function of environmental conditions. In many cases, since drift is negligible over the chip lifetime, mismatch is even measured at manufacturing and fused into a ROM [1].

In conclusion, DACs can be compensated if redundancy is available, and can be calibrated in general. In theory, the benefit from doing so is in terms of array area, not in terms of power or speed, which should be taken into account when deciding whether to go for intrinsic matching or not, especially at medium resolution. It can be done in background, but in many cases a foreground-only approach is also quite valid, since the non-ideality to be corrected is often quite stable over environmental conditions.

4 Pipelined Dynamic SAR

In this section we will discuss the design of an analog to digital converter for a software defined receiver. Power consumption will be aggressively minimized using digital assistance. The basic architecture is reused from the dynamic pipelined SAR architecture of [6], which scales favorably to the 0.9 V supply 28 nm technology. Comparator offsets and residue amplifier gain are automatically calibrated on-chip, in either a foreground or a background mode. For these non-idealities calibration is preferred over compensation because of the relative ease of implementing analog controls. Channel offset is observed and compensated off-chip, as compensation simply requires a half-rate digital addition. Channel gain mismatch is observed off-chip but calibrated on-chip to avoid fairly power-hungry digital multipliers that would be required for channel gain compensation. The ADC prototype achieves a peak SNDR of 59.8 dB at 410 MS/s for 2.1 mW of power [19].

4.1 Architecture

The implemented ADC consists of two interleaved channels with a full-rate front-end sampling switch as shown in Fig. 6. The channels each consist of a 6b coarse SAR stage, a dynamic residue amplifier, a 7b fine SAR stage and are accompanied by a calibration engine. The calibration detects and corrects comparator offsets and residue amplifier gain either in the foreground or in the background. Since each channel is observed individually, the on-chip calibration engine does not correct interleaving effects such as offset or gain mismatch. Both of these are observed externally, and offset is compensated externally while channel gain is calibrated.

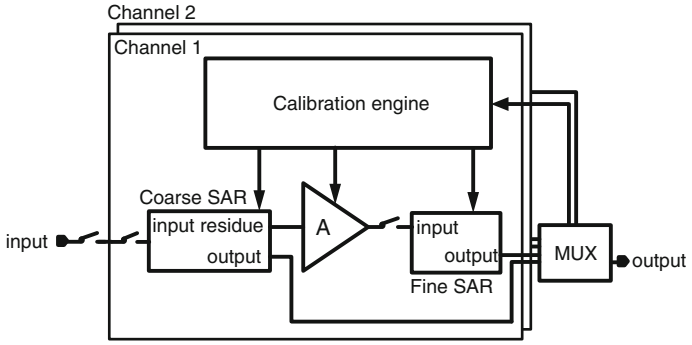

Fig. 6 High-level block diagram of ADC

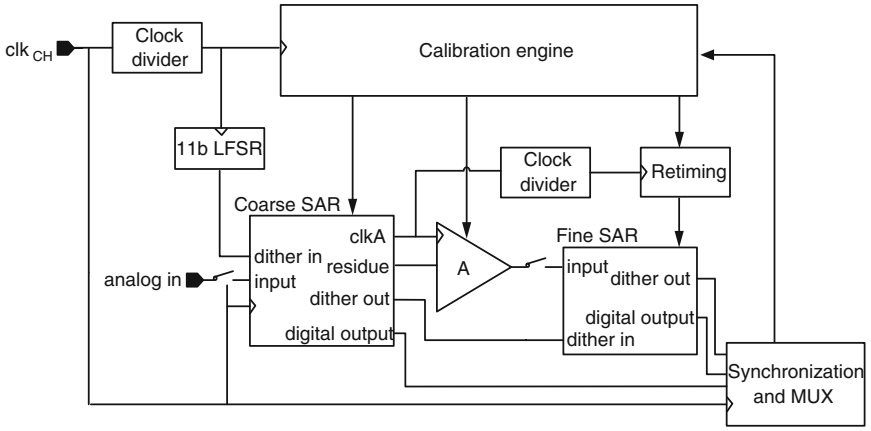

Fig. 7 Detailed block diagram of channel

A detailed view of one of the interleaved channels is shown in Fig. 7. Each channel operates with one synchronous main channel clock used for sampling and the start of the coarse SAR operation. After the coarse SAR has finished its conversion, the residue amplifier is asynchronously activated using clkA and in turn activates the fine SAR. Since the actual calibration engine operates on a single clock, its outputs are retimed using clkA to ensure calibration inputs change only when appropriate. Both the main calibration engine clock and the retiming clock are divided by a programmable factor ranging from 1 to 1,024 in binary scaled steps which allows a trade-off between calibration settling time and digital power consumption. An 11b linear feedback shift register (LFSR) is used to generate a 2,047 length PRBS dither sequence that is synchronized in each SAR stage and finally sent to the calibration engine for use in the residue amplifier gain calibration as will be explained in the section on the calibration algorithms.

4.2 Core ADC Design

The coarse SAR itself is implemented similar to [6, 16] using a comparator controlled step-down DAC. There are three potential benefits to this arrangement. First, the power consumption and delay in the SAR controller is avoided, which admittedly is a fairly minor advantage in deeply scaled CMOS. Second, the above arrangement avoids constraints on common-mode dependence of comparator offsets [31]. Indeed, in the step-down DAC used here the DAC common-mode changes every cycle. This results in a different comparator input common-mode in every cycle, and since in most comparator designs the comparator offset is common-mode dependent, this changing common-mode could introduce comparator errors. While these errors can be mitigated using redundancy or a modified comparator design [31], having multiple comparators ensures that each comparator operates at a specific common-mode, and is thus insensitive to common-mode dependent offset. Finally, the above arrangement trivially allows cycle-specific choices for comparator sizes and DAC settling delays. For example, since MSB settling is typically critical in a SAR ADC, one could easily use a larger delay value for the delay line between the first and second comparator than for the later cycles.

There are also two disadvantages. First, the need to implement multiple comparators rather than a single one: while this an energy-neutral operation using dynamic comparators, it does impose a small area penalty. The more significant disadvantage is the fact that each of the 6 comparators in the coarse SAR have potentially different offsets, which need to be calibrated to ensure accuracy. This can be done in foreground or background, but definitely requires digital logic, with a further area penalty. In practice, the area increase due to the calibration logic, especially combined with the area for multiple comparators, is larger than whatever area is saved by eliminating the need for a controller in the first place, so this is not a beneficial choice from an area perspective. In addition, the digital calibration logic requires some energy, which should be lower than the energy typically required by the SAR controller. For reasonable speed ADCs this is usually the case: whereas a SAR controller is active and thus consumes energy in every ADC cycle, calibration logic only needs to be activated often enough to track environmental changes. Since in most environments, kilocycle/second calibration speeds are sufficient, a comparator-controlled SAR results in energy savings even for an ADC running at a fairly modest 10 MS/s.

Due to a 1 bit redundancy between the coarse and fine stage, the coarse stage is robust to any comparator error up to $\pm 0.5\ \mathrm{LSB}_{\mathrm{coarse}}$. Since comparator offset is calibrated, this redundancy can be used almost completely for errors due to DAC settling or comparator noise. The coarse SAR also adds $\pm 0.5\ \mathrm{LSB}_{\mathrm{coarse}}$ dither to the residue for use in the amplifier gain calibration as will be discussed later. The coarse SAR is designed for intrinsic matching, since the 1 pF required at maximum density is not yet prohibitive in terms of power consumption or speed.

The residue amplifier must amplify the residue so that it can be sampled by the fine SAR converter. Similar to [6] a dynamic amplifier is chosen since it combines

low, purely dynamic power consumption with fast settling. Linearity issues are avoided by using a limited output range and the gain uncertainty is compensated using a calibration. Power supply rejection is not an issue, since the amplifier will be powered from the 0.9 V reference, which must be extremely stable anyway. Common-mode rejection is a potential issue: this approach assumes a fairly stable common-mode, which changes at a speed that can be tracked by background calibration.

The fine SAR architecture is similar to that of the coarse SAR except in the fact that dither is subtracted prior to the first comparator in the fine stage, whereas it is injected after the last comparison in the coarse stage. The top plate of a step-down DAC is connected to the input of 5 high noise, and 2 low noise comparators directly controlling aforementioned DAC. The DAC is implemented using 2.2 fF units, not limited by the relaxed 6b matching requirements. Extra capacitance is added to the DAC top plate to reduce the range of the DAC, thereby reducing the required output range of the residue amplifier. One bit of redundancy between the 5th and 6th cycles of this SAR renders the converter robust to comparator noise errors in the first 5 cycles. This allows the use of a relatively high noise comparator in these cycles whereas the last two comparisons are noise-critical and therefore done in comparators with lower noise and higher power consumption.

As each comparison in both coarse and fine SAR ADCs is done with a different comparator, comparator offset errors potentially affect the converter accuracy. The comparators can be sized for noise if comparator offset is calibrated as will be discussed in the next section. This lowers the power consumption of these comparators to the same level required in a conventional, single-comparator, SAR converter, as offset calibration can be implemented with an extremely small power penalty. All comparators are implemented as in [9] with their offset calibrated using digitally controllable MOS-capacitors.

4.3 Calibration Engine

The calibration engine supports both foreground and background calibration. It is clocked at the start of tracking, which ensures that all its outputs change during the track-time while the coarse SAR and amplifier are reset. The fine SAR calibration inputs are retimed to change at the start of the fine SAR reset, which ensures they also only change while the circuits they control are in reset mode.

4.3.1 Foreground Calibration

Foreground calibration assumes a zero differential input signal at the ADC common-mode. Comparator offsets are calibrated first, followed by amplifier gain calibration. Since all coarse SAR comparators are activated at a different common-mode during normal operation, this common-mode is replicated during calibration.

This is done by shifting the common-mode of the zero differential sampled signal using direct control of the internal DACs before activation of the comparator in question. The output of said comparator is then accumulated during a programmable number of cycles and a binary search drives the comparator threshold to obtain an equal number of zeros and ones at the comparator output.

To calibrate the fine SAR comparator offsets, the coarse SAR DAC is programmed to shift the common-mode amplifier input voltage to the appropriate value. The amplifier is then activated to ensure a zero differential voltage of the appropriate common-mode at the amplifier output. This voltage is then shifted by the fine SAR DAC to the correct common-mode value for each fine SAR comparator. The average comparator output is again driven to 0.5 using a binary threshold search.

In our implementation, all 13 comparators are calibrated sequentially for simplicity, since the direct DAC control allows only one common-mode shift per ADC cycle. Using slightly more complex direct DAC control, multiple common-mode steps per cycle are possible, which could allow calibration of all 13 comparator thresholds to happen concurrently.

In the last step of the foreground calibration the amplifier gain is tuned. The zero differential input voltage is now transformed by the coarse DAC into ± 0.5 LSB_{coarse} amplifier inputs at the correct common-mode for the residue amplifier. The fine SAR stage is then operated normally and its output is accumulated for a programmable number of cycles with positive and negative input each. If the average difference between the two obtained second stage outputs is smaller than 32, the gain of the residue amplifier is increased, if it is larger than 32, the gain is decreased.

4.3.2 Background Calibration

Background calibration is implemented using three different algorithms running concurrently. The first algorithm calibrates the six coarse SAR and first 5 fine SAR comparator offsets using the redundancy available in the ADC [18]. The second algorithm calibrates the final two comparators of the fine SAR based on their average outputs. A final algorithm is used to calibrate the residue gain using dither injection similar to [32].

The first algorithm is illustrated in Fig. 8 for calibration of offsets in the first stage using a one bit redundant second stage, with V_{in1} representing the first stage input and V_{in2} representing second stage input. Since at design-time the ideal residue transfer characteristic is well-known, the nominal upper and lower limits of this characteristic are also known. The implemented algorithm assumes the second stage is ideal, in which case the quantized output residue can only exceed these nominal limits when an error occurs in one of the first stage comparators. Based on the second stage output the sign of the error can be determined, and based on the digital output of the first stage, the specific comparator responsible for the error can be identified.

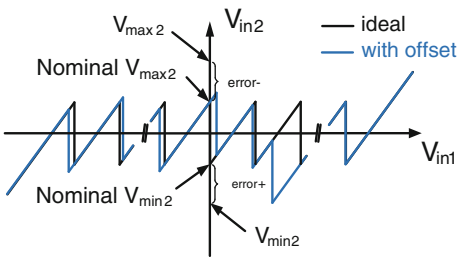

Fig. 8 Illustration of first calibration algorithm

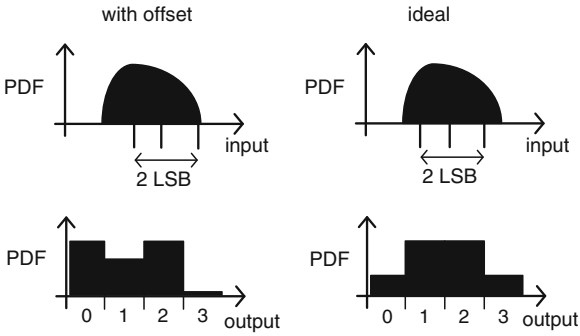

Fig. 9 Illustration of second calibration algorithm

In our implementation the difference between the number of incorrect positive and the number of incorrect negative decisions for each comparator is counted, and when this counter exceeds a programmable limit the comparator offset is adjusted to reduce the probability of the more common type of error. By comparing the relative probabilities of positive and negative errors the impact of comparator noise, incomplete DAC settling and the fact that the fine SAR is not actually ideal are reduced.

Since the last two comparators of the fine SAR do not have a redundant comparison available, a second calibration algorithm uses output statistics to calibrate these comparators as illustrated in Fig. 9. Because of the redundancy in our implementation, the PDF at fine SAR comparator input after the first five bits of the fine SAR would ideally be exactly 2 LSBs wide, but because of non-idealities in the preceding residue generation this PDF will be somewhat wider. Assuming that this PDF is more or less symmetric, the optimal placement of the comparator thresholds is such that a symmetric output histogram is obtained. This corresponds uniquely to a 50 % distribution for both comparator outputs. The final two comparator thresholds are thus calibrated by accumulating the comparator output during a programmable number of cycles and adjusting the threshold if the number of positive decisions is larger than 75 % or smaller than 25 %.

The third and final algorithm is used to calibrate the residue amplifier gain as illustrated in Fig. 10, slightly modified from [22]. A pseudo-random signal PN is added to the residue amplifier input using the coarse stage DAC, and an ideally

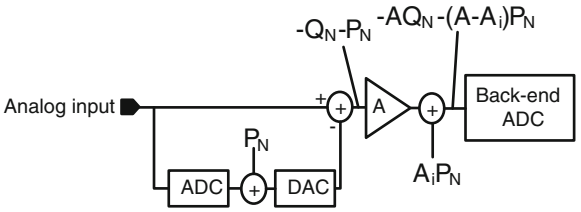

Fig. 10 Illustration of amplifier gain calibration

scaled version of this dither is subtracted at the output of the amplifier by the second stage DAC, before the fine SAR quantization. The fine SAR output is accumulated in one of two registers based on the sign of the dither and when the difference between these two accumulations exceeds a certain value the residue amplifier gain is adjusted. In our implementation we use ± 0.5 LSB_{coarse} and ± 0.5 MSB_{fine} dither amplitude and a 2047 length pseudo-random signal generated using an 11b LFSR. To reduce the complexity only the 4 MSBs of the fine SAR are accumulated. Complexity is further reduced by comparing only the 7 MSBs of the accumulators, instead of calculating the difference between the two accumulator outputs and comparing this difference to a certain threshold. If the 7 MSBs are not the same for both accumulators, gain is adjusted. While in the simplified scheme the gain might be adjusted based on a very small actual difference between the two accumulators, the amplifier gain step is sufficiently small that a few unrequired changes of the amplifier gain do not significantly affect performance.

The calibration engine is described in VHDL and synthesized with a total gate count of 6,080. Block area after place and route is approximately 60 μm by 250 μm for each calibration engine. Area and power consumption could be optimized further by omitting some debugging and testing options.

4.4 Measurements

The ADC prototype has been manufactured in an 1P9 M 28 nm CMOS process with a core chip area of 0.11 mm^2 including the calibration engines but not decoupling (Fig. 11). Individual supply domains with approximately 180 pF decoupling each are used for channel 1 and channel 2; both calibration engines are implemented on a third supply domain with another 120 pF decoupling. The chip area including this decoupling capacitance is 0.25 mm^2, which could be somewhat optimized by using MOS capacitors for decoupling instead of only MOM. Gain and offset mismatch between the two channels are measured by applying non-zero inputs and observing both channel outputs. Gain mismatch is calibrated by changing the top plate DAC capacitance in the coarse SAR, offset is compensated digitally.

INL and DNL performance is measured at 20 MS/s after a 3,000 cycle on-chip foreground calibration and shown in Fig. 11. While the DNL is unremarkable, the INL shows a significant and consistent pattern in both channels and across different dies. This has been identified to be due to unintended front-end layout

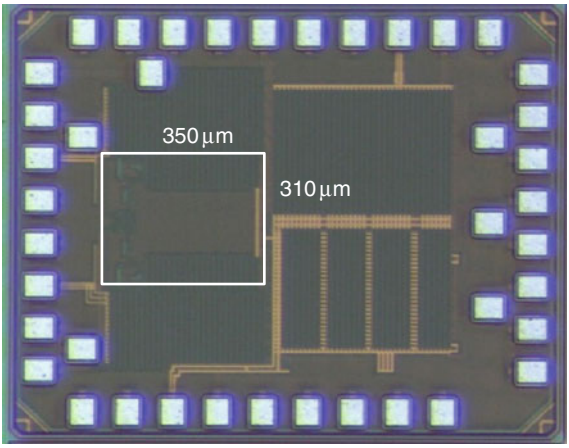

Fig. 11 Chip micrograph with core area indicated

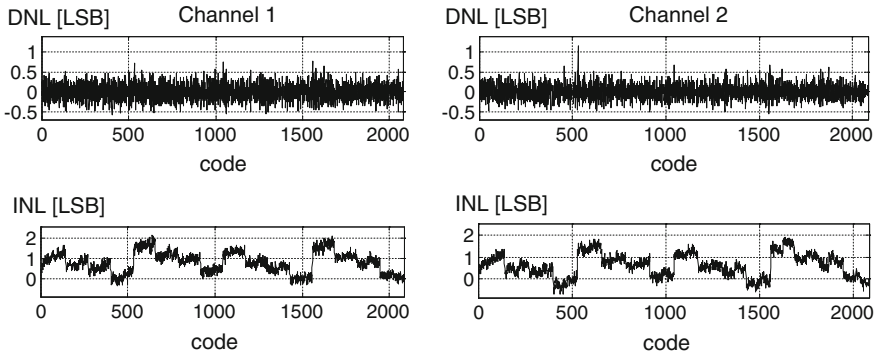

Fig. 12 INL and DNL for channel 1 (*left*) and channel 2 (*right*)

dummies (active and poly) in the first stage DAC. Since redundancy is available, this INL pattern is digitally compensated by adjusting the first stage bit weights for the remaining measurements (Fig. 12).

The settling behavior of the background calibration is measured by programming the chip in default, center-range settings and applying a full-scale sine wave input. The SNDR calculated in 16 k windows with variable start point is shown in the left plot of Fig. 13 for different divider ratios for the calibration engine clock: the background calibration converges within approximately 50 k calibration cycles per channel. This settling behavior is dominated by the amplifier calibration, which needs at least 2047 cycles per step. Presetting the amplifier gain to a nearly ideal value, settling to steady-state is significantly faster, as shown in the right plot of Fig. 13.

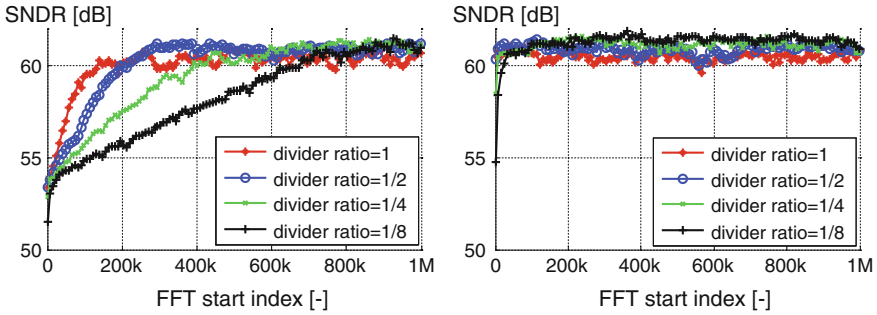

Fig. 13 SNDR versus FFT start index after programming default settings (*left*) and after programming default comparator settings and correct amplifier gain (*right*)

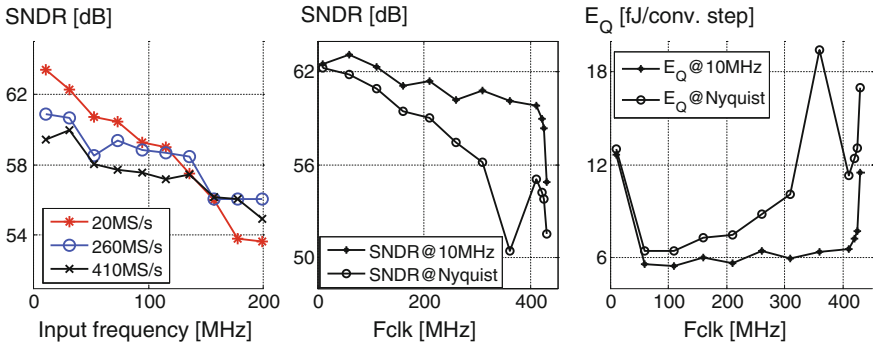

Fig. 14 SNDR versus input frequency (*left*), SNDR versus. clock frequency (*middle*) and energy per conversion step versus clock frequency (*right*)

Dynamic performance is measured with background calibration clocked every 1,024 cycles. The SNDR versus input frequency is shown in the left plot of Fig. 14 for different clock frequencies: at 20 MS/s the ERBW is limited by external clock phase noise, at 260 and 410 MS/s the ERBW exceeds 130 MHz. The SNDR versus clock frequency is shown in the middle plot of Fig. 14 for low and near-Nyquist input frequency. The low frequency SNDR degrades fairly gradually from 62.5 dB at 10 MS/s to 59.8 dB at 410 MS/s. The high frequency behavior degrades gradually at moderate clock frequencies, shows significant degradation around 360 MS/s and then increases somewhat again. The high input frequency degradation at 360 MS/s is due to time skew induced by supply coupling: the clock generation is powered by the channel 1 supply, and at 360 MS/s generation of the sampling clock occurs during DAC reset, which generates some supply bounce. This supply bounce modulates the sampling clock delay and because this only occurs every other sample, a 4.5 ps time skew occurs.

The circuit consumes 100 μA leakage current, 5 pJ per ADC clock cycle and 3.3 pJ per calibration engine clock cycle from a 0.9 V supply. Peak efficiency is 5.5 fJ/conv. step at 110 MS/s and less than 12 fJ at 410 MS/s with a Nyquist input signal as shown in the right plot of Fig. 14. This efficiency is achieved by leveraging digital assistance wherever practical and clocking said assistance at a drastically reduced clock rate, sufficient to track slow environmental changes.

5 Conclusions

It has been shown that digital assistance significantly changes the design tradeoffs in analog to digital converters. Comparator calibration shifts accuracy constraints from offset to noise which significantly reduces comparator power consumption and input capacitance. In addition, comparator calibration allows some alternative architectures that potentially further increase power efficiency. In amplifiers, digital assistance relaxes requirements on linearity and gain precision and shifts the critical requirement to input referred noise. In addition to lowering power consumption in conventional amplifiers, this opens up many alternative amplifier topologies. In DACs, digital assistance can be used to reduce required area for a matching level, which is potentially useful at high resolutions.

Some of these concepts have been illustrated in the design of an interleaved 410 MS/s 11 bit pipelined SAR ADC enabled by digital assistance. A SAR architecture leveraging comparator offset calibration, residue amplifier gain calibration, channel gain calibration and channel offset compensation are assumed during the design phase to relax analog requirements. In addition, due to a design mistake, DAC linearity compensation is used. The measured prototype obtains a peak SNDR of 63.3 dB at 20 MS/s and up to 59.8 dB at 410 MS/s with energy per conversion step below 12 fJ. To date, similar performance has not been achieved with similar energy consumption without digital assistance.

Acknowledgments The author would like to thank Masao Iriguchi, Manuel de la Guia Solaz, Guy Glorieux, Kazuaki Deguchi, Badr Malki, Ewout Martens and Jan Craninckx for their contributions to this paper.

References

1. Z. G. Boyacigiller, B. Weir, P. D. Bradshaw, “An Error-Correcting 14b/20μs CMOS A/D Converter,” IEEE ISSCC 1981
2. B. Murmann, “Digitally assisted data converter design,” IEEE ESSCIRC, September 2013, pp. 24–31
3. B. Murmann, “ADC Performance Survey 1997-2013,” [Online]. Available: http://www.stanford.edu/∼murmann/adcsurvey.html

4. B. Setterberg, et al., "A 14b 2.5GS/s 8-way-interleaved pipelined ADC with background calibration and digital dynamic linearity correction," IEEE ISSCC 2013
5. G. Van der Plas, et al., "A 0.16pJ/Conversion-Step 2.5mW 1.25GS/s 4b ADC in a 90 nm Digital CMOS Process," IEEE ISSCC 2006
6. B. Verbruggen, et al., "A 1.7mW 11b 250MS/s 2× interleaved fully dynamic pipelined SAR ADC in 40 nm digital CMOS," IEEE JSSC, Vol. 47, No. 12, December 2012, pp. 2880–2887
7. A. Shikata, et al., "A 0.5 V 1.1 MS/sec 6.3 fJ/Conversion-Step SAR-ADC With Tri-Level Comparator in 40 nm CMOS," IEEE JSSC, Vol. 47, No. 4, April 2012, pp. 1022–1030
8. T. Kobayashi, et al., "A current-controlled latch sense amplifier and a static power-saving input buffer for low-power architecture," IEEE JSSC, Vol. 28, No. 4, April 1993, pp. 523–527
9. M. Miyahara, et al., "A low-offset latched comparator using zero-static power dynamic offset cancellation technique," IEEE ASSCC 2009
10. M.-J. E. Lee, et al., "Low-power area-efficient high-speed I/O circuit techniques," IEEE JSSC, Vol. 35, No. 11, November 2000, pp. 1591–1599
11. F. H. Gebara, et al., "A body-driven offset cancellation technique in PD-SOI," International Conference on Microelectronics, May 2004, Vol. 2, pp. 567–570
12. B. Verbruggen, et al., "A 2.2 mW 1.75 GS/s 5 Bit Folding Flash ADC in 90 nm Digital CMOS," IEEE JSSC, Vol. 44, No. 3, March 2009, pp. 874–882
13. Y.-S. Shu, "A 6b 3GS/s 11mW Fully Dynamic Flash ADC in 40 nm CMOS with Reduced Number of Comparators," Symposium on VLSI Circuits, June 2012, pp. 26–27
14. V. H.-C. Chen, et al., "An 8.5mW 5GS/s 6b flash ADC with dynamic offset calibration in 32 nm CMOS SOI," Symposium on VLSI Circuits, June 2013, pp. 264–265
15. V. Giannini, et al., "An 820 μW 9b 40MS/s Noise-Tolerant Dynamic-SAR ADC in 90 nm Digital CMOS," IEEE ISSCC 2008
16. G. Van der Plas, et al., "A 150 MS/s 133 μW 7 bit ADC in 90 nm Digital CMOS," IEEE JSSC, Vol. 43, No. 12, December 2008, pp. 2631–2640
17. Y. Nakajima, et al., "A Background Self-Calibrated 6b 2.7 GS/s ADC With Cascade-Calibrated Folding-Interpolating Architecture," IEEE JSSC, Vol. 45, No. 4, April 2010, pp. 707–718
18. Z. Gu, et al., "A novel self-calibrating scheme for video-rate 2-step flash analog-to-digital converter," IEEE International Symposium on Circuits and Systems, May 1992, Vol. 2, pp. 601–604
19. B. Verbruggen, et al., "A 2.1 mW 11b 410 MS/s dynamic pipelined SAR ADC with background calibration in 28 nm digital CMOS," Symposium on VLSI Circuits, June 2013, pp. 268–269
20. J. K Fiorenza, et al., "Comparator-Based Switched-Capacitor Circuits for Scaled CMOS Technologies," IEEE JSSC, Vol. 41, No. 12, December 2006, pp. 2658–2668
21. B. Hershberg, et al., "Ring Amplifiers for Switched Capacitor Circuits," IEEE JSSC, Vol. 47, No. 12, December 2012, pp. 2928–2942
22. B. Murmann, et al., "A 12-bit 75-MS/s pipelined ADC using open-loop residue amplification," IEEE JSSC, Vol. 38, No. 12, December 2003, pp. 2040–2050
23. M. Anthony, et al., "A process-scalable low-power charge-domain 13-bit pipeline ADC," Symposium on VLSI Circuits, June 2008, pp. 222–223
24. N. Dolev, et al., "A 12-bit, 200-MS/s, 11.5-mW pipeline ADC using a pulsed bucket brigade front-end," Symposium on VLSI Circuits, June 2013, pp. 98–99
25. J. Hu, et al., "A 9.4-bit, 50-MS/s, 1.44-mW Pipelined ADC Using Dynamic Source Follower Residue Amplification," IEEE JSSC, Vol. 44, No. 4, April 2009, pp. 1057–1066
26. S.-H. W. Chiang, et al., "A 10-Bit 800-MHz 19-mW CMOS ADC," Symposium on VLSI Circuits, June 2013, pp. 100–101
27. L. Brooks, et al., "Background Calibration of Pipelined ADCs Via Decision Boundary Gap Estimation," IEEE TCAS I, Vol. 55, No. 10, November 2008, pp. 2969–2979
28. M. J. M. Pelgrom, et al., "Matching properties of MOS transistors," IEEE JSSC, Vol. 24, No. 5, October 1989, pp. 1433–1439

29. P. Harpe, et al., “A 2.2/2.7fJ/conversion-step 10/12b 40kS/s SAR ADC with Data-Driven Noise Reduction,” IEEE ISSCC 2013
30. N. Verma, et al., “A 25/spl mu/W 100kS/s 12b ADC for wireless micro-sensor applications,” IEEE ISSCC 2006
31. C.-C. Liu, et al., “A 10-bit 50-MS/s SAR ADC With a Monotonic Capacitor Switching Procedure,” IEEE JSSC, Vol. 45, No. 4, April 2010, pp. 731–740
32. E. J. Siragusa, et al., “Gain error correction technique for pipelined analogue-to-digital converters,” Electronics Letters, vol. 36, no. 7, March 2000, pp. 617–618

http://www.springer.com/978-3-319-07937-0

High-Performance AD and DA Converters, IC Design in Scaled Technologies, and Time-Domain Signal Processing
Advances in Analog Circuit Design 2014
Harpe, P.; Baschirotto, A.; Makinwa, K.A.A. (Eds.)
2015, XII, 418 p. 331 illus., 167 illus. in color., Hardcover
ISBN: 978-3-319-07937-0