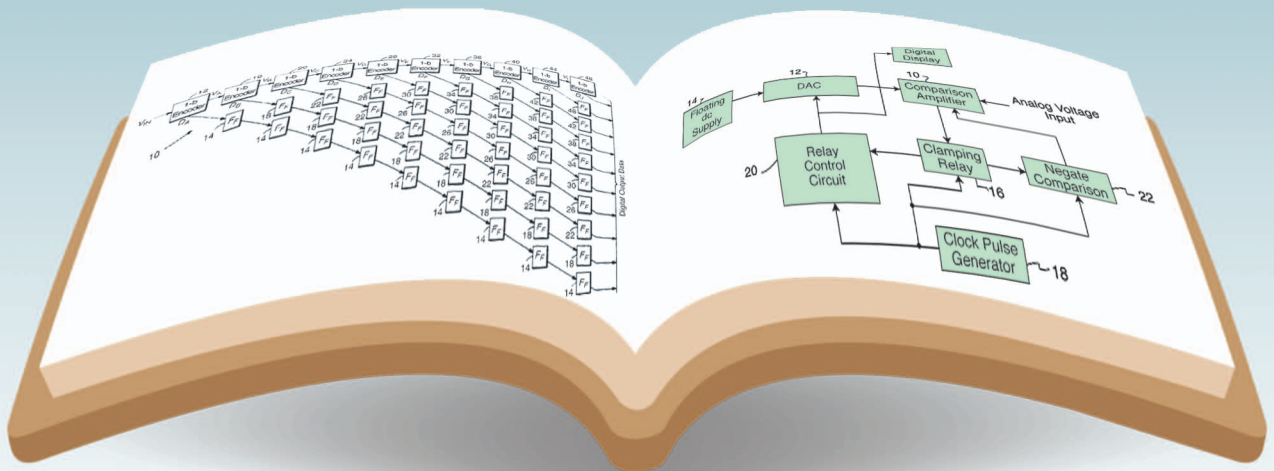



A Tale of Two ADCs



BOOK LICENSED BY GRAPHIC STOCK

Pipelined Versus SAR

//////////
Behzad Razavi

Among numerous Nyquist-rate analog-to-digital converter (ADC) architectures introduced throughout the years, three have found the widest usage: flash, pipelined, and successive-approximation register (SAR) topologies. In this article, we focus on the last two and study their similarities and differences.

Brief History

Pipelined and SAR ADC architectures date back to the mid-1900s. Pipelining in ADCs was evidently first

proposed by Severin of Texas Instruments in a patent filed in 1967 [2]. Figure 1 shows the idea, which remarkably resembles today's 1-bit/stage designs. The SAR topology is even older and can be found in a 1947 paper by Goodall of Bell Laboratories [1] and a patent filed in 1957 by Bell of Consolidated Electrodynamics Corporation [3]. Shown in Figure 2, the latter uses mechanical relays to drive a digital-to-analog converter (DAC), whose output is compared to the analog input by a "comparison amplifier" (a comparator). CMOS pipelined ADCs began to appear in the 1980s [4], [5] and continued to flourish thereafter. However, CMOS SAR designs, having been reported in the 1970s [7], lay virtually dormant until the 2000s, when their potential for low power consumption was recognized. Further historical

Digital Object Identifier 10.1109/MSSC.2015.2442372
 Date of publication: 15 September 2015

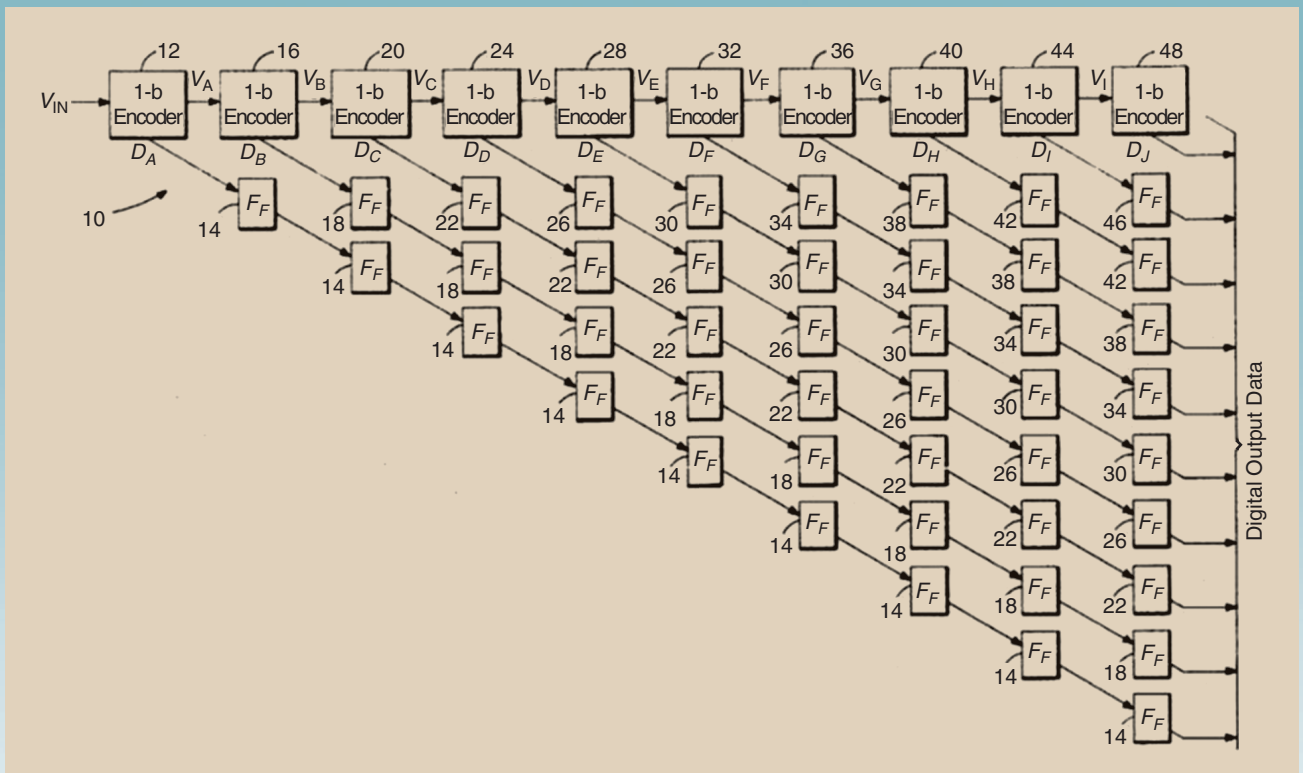


FIGURE 1: A patent showing a 1-bit/stage pipelined ADC.

details can be found elsewhere in this issue [6].

Quantization by Binary Search

Analog-to-digital conversion comprises two fundamental functions, namely, sampling and quantization. The latter can be viewed as an operation wherein an “analog estimate” is identified and the digital equivalent of this estimate is created. In a flash ADC, for example, the reference ladder generates the analog estimates, and the comparators identify one as the closest lower value, producing its digital value at the output.

In their basic form, both pipelined and SAR ADCs use binary search to compute analog estimates that successively converge toward the input voltage. Suppose, as shown in Figure 3(a), that an ADC having an input range of 0 to V_{REF} senses an input V_{in} . Binary search begins by nominating $V_{REF}/2$ as the best analog estimate. Next, since $V_{in} > V_{REF}/2$, the search identifies $3V_{REF}/4$ as a better approximation, etc. That is, in each cycle the ADC compares V_{in}

with the most recent analog estimate and directs the search according to the *polarity* of their difference. We call this difference the “residue” and denote it by $V_{in} - \alpha V_{REF}$, where $\alpha = 1/2, 3/4, 5/8$ etc., in the above example. The goal of binary search is to reduce the residue to less than 1 least significant bit (LSB). Figure 3(b) illustrates how the decision result in one cycle leads to comparison with the proper analog estimate in the next cycle.

Pipelined ADCs

Basic Operation

To arrive at the basic pipelined architecture, we first note that the binary search begins with a residue of $V_{in} - V_{REF}/2$, which has the same polarity as $2(V_{in} - V_{REF}/2) = 2V_{in} - V_{REF}$. We can thus form $2V_{in} - V_{REF}$ as the residue and benefit from the $2 \times$ amplification that it provides before going to the next binary search cycle. This method is attractive if the function $f(V_{in}, V_{REF}) = 2V_{in} - V_{REF}$ can be realized efficiently and compactly.

Shown in Figure 4 [8] is a popular implementation known as the “multiplying DAC” (MDAC) stage. In the sampling (acquisition) mode, C_1 and C_2 ($= C_1$) track V_{in} while node X is kept at zero by the unity-gain amplifier. In the amplification mode, C_1 is “flipped” around the operational amplifier (op-amp), and the left plate of C_2 jumps to V_{REF} . The output voltage thus settles to $2V_{in} - V_{REF}$ if the op-amp gain is high and the mismatch between C_1 and C_2 small.

We now ponder operation in the next cycle. If the residue, $2V_{in} - V_{REF}$, is subjected to a subsequent MDAC stage, we have a new residue equal to $f(2V_{in} - V_{REF}, V_{REF}) = 2(2V_{in} - V_{REF}) - V_{REF} = 4V_{in} - 3V_{REF}$, whose polarity allows comparison of V_{in} with $3V_{REF}/4$. But, according to Figure 3(b), this comparison is meaningful only if $V_{in} > V_{REF}/2$; in the case of $V_{in} < V_{REF}/2$, we must compare with $V_{in}/4$. This observation points to the need for a residue equal to $4V_{in} - V_{REF}$ if $V_{in} < V_{REF}$. However, this value cannot be obtained recursively from the function $f(V_{in}, V_{REF}) = 2V_{in} - V_{REF}$.

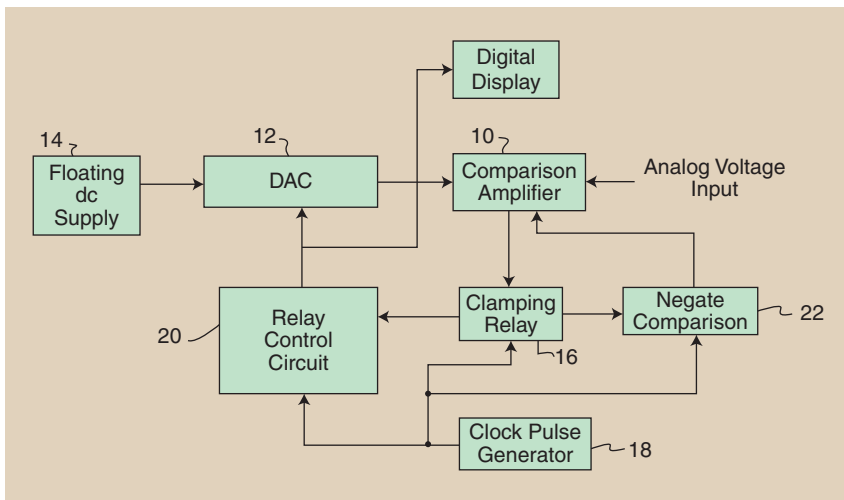


FIGURE 2: A patent showing a SAR ADC.

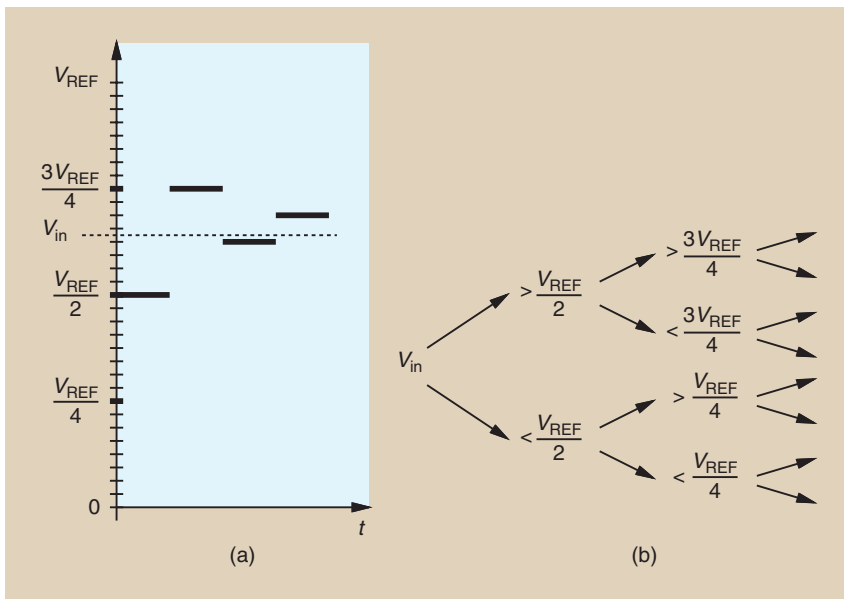


FIGURE 3: (a) A search for analog estimates; (b) decision-directed binary search.

To overcome this difficulty, we modify the MDAC operation to

$$f(V_{in}, V_{REF}) = 2V_{in} - V_{REF} \quad \text{if } V_{in} > \frac{V_{REF}}{2} \quad (1)$$

$$= 2V_{in} \quad \text{if } V_{in} < \frac{V_{REF}}{2}. \quad (2)$$

Figure 5(a) plots $V_{res} = f(V_{in}, V_{REF})$. As shown in Figure 5(b), the implementation compares V_{in} with $V_{REF}/2$ before residue generation, directing the left plate of C_2 to V_{REF} or 0 accordingly. In this case, the capacitors and the comparator sample the input simultaneously. Since the circuit

does not employ a main sample-and-hold (SHA), it is also called a “SHA-less” front end.

The foregoing principle readily lends itself to pipelining: while one MDAC stage is in the amplification mode, the next can be in the acquisition mode and vice versa [Figure 5(c)]. The concurrent operation of the MDAC stages means that the conversion speed is limited by only the acquisition and amplification times of one stage, typically the first one.

The compact implementation along with pipelining makes this “1-bit/stage” architecture an efficient solution,

especially for resolutions not achievable by flash topologies (approximately 8 bits and above).

This architecture entails five sources of inaccuracy: kT/C and op-amp noise, capacitor mismatch, finite op-amp gain, op-amp nonlinearity, and comparator offset. Of these, noise directly trades with power dissipation, capacitor and op-amp imperfections can be calibrated in the digital domain, and comparator offset is accommodated by a simple modification described below. High-speed designs must also cope with the mismatch between the sampling instants of the capacitors and the comparator in Figure 5(b).

One may surmise that the hardware and power dissipation of the above architecture grow linearly with the resolution; for each additional bit, we can simply add one more MDAC stage to the end of the pipeline. Today’s designs, however, are noise-limited, requiring that, for one more bit, the front-end stage (and possibly the subsequent stages) incorporate proportionally larger capacitors and higher op-amp bias currents. In other words, the hardware and power of the overall pipeline grow approximately *exponentially* with the resolution.

1.5-Bit/Stage Architecture

Let us return to the sources of error. It can be proved that capacitor mismatch and finite op-amp gain alter the slope of the residue plot [Figure 6(a)], whereas comparator offset, V_{OS} , shifts the decision point away from $V_{in} = V_{REF}/2$ [Figure 6(b)]. Both cases may exhibit a residue “over-range,” but the former introduces an error in V_{res} even outside the over-range regions. We deal with the latter case in this section, modifying the architecture so as to tolerate large comparator offsets.

Our subsequent study is more easily followed if we consider a fully-differential system, one wherein the differential input voltage can vary from $-V_{REF}$ to $+V_{REF}$. The residue plot of Figure 5(a) is redrawn

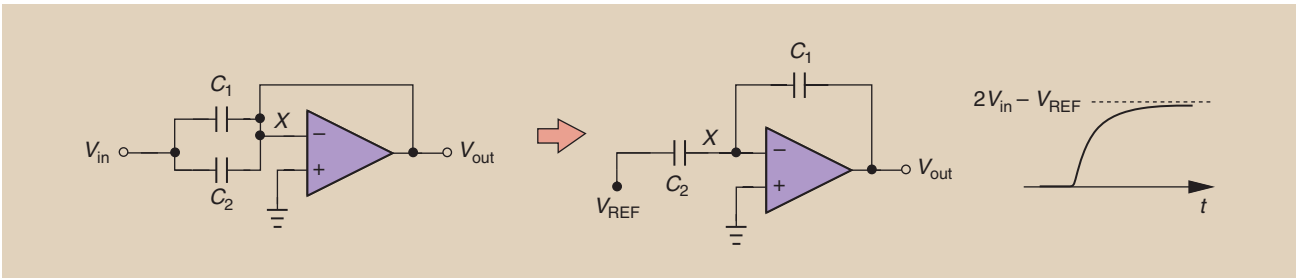


FIGURE 4: The MDAC implementation.

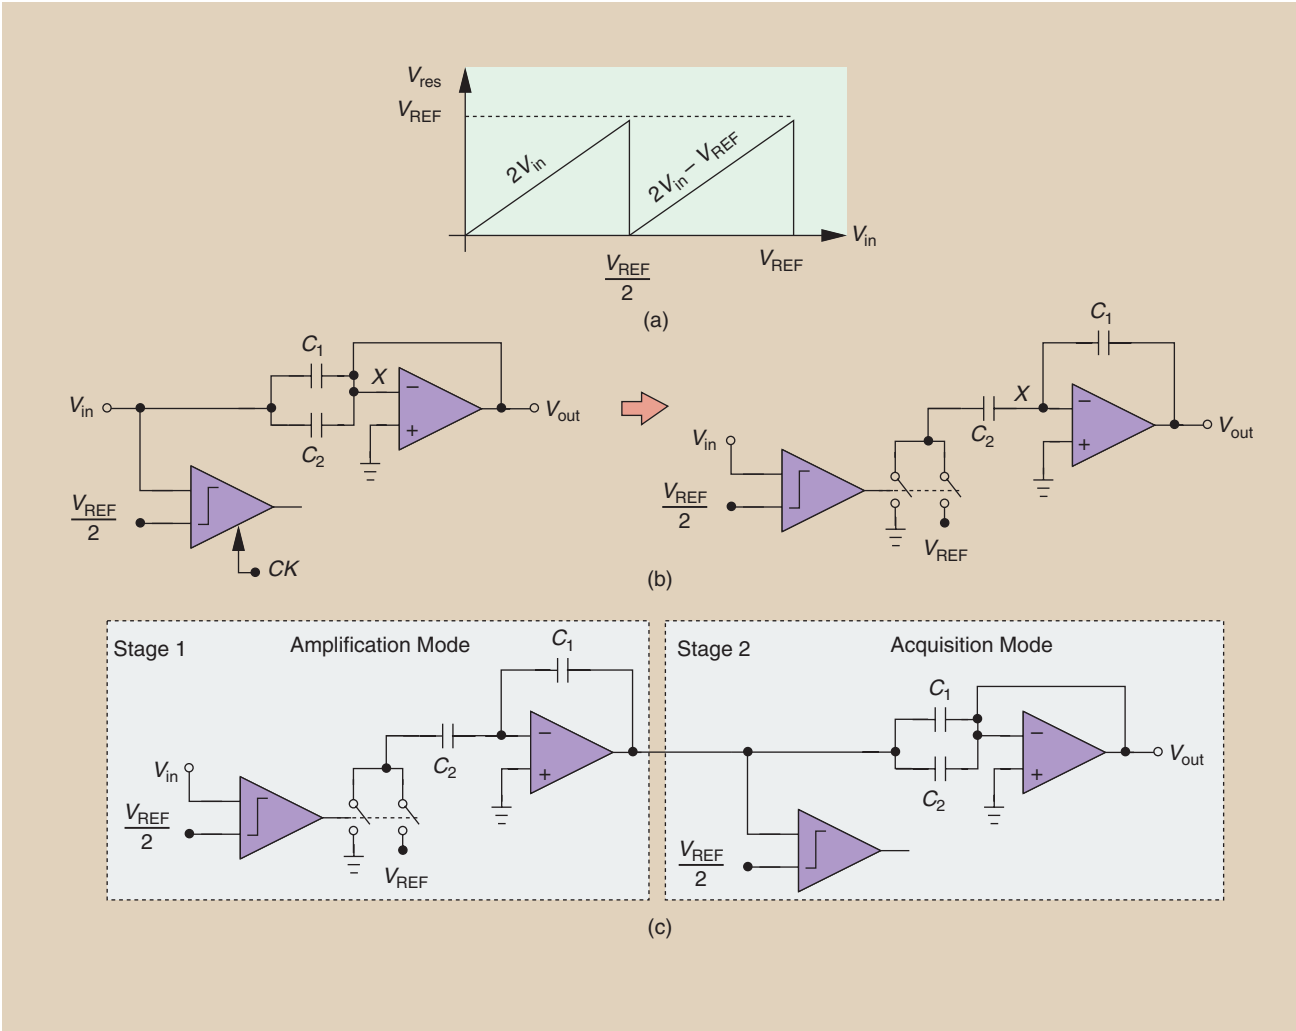


FIGURE 5: (a) A residue plot, (b) use of a comparator to generate proper residue, and (c) two pipelined stages.

in Figure 7(a) for such a system and expressed as

$$\begin{aligned}
 f(V_{in}, \pm V_{REF}) &= 2V_{in} - V_{REF} \quad V_{in} > 0 \quad (3) \\
 &= 2V_{in} + V_{REF} \quad V_{in} < 0. \quad (4)
 \end{aligned}$$

Note that the comparator now compares V_{in} with zero. For example, if $V_{in} = 0$ but the comparator interprets it to be $V_{in} = +2 \text{ mV}$, then the




FIGURE 6: The effect of (a) capacitor mismatch and (b) finite op-amp gain on the residue.




FIGURE 7: (a) The residue plot for a differential system, (b) residue behavior avoiding overrange or underrange, and (c) a 1.5-bit/stage implementation.

decision does not occur at $V_{in} = 0$, and $2V_{in} + V_{REF}$ continues into the positive V_{in} territory, exceeding $+V_{REF}$. Our objective is to confine $2V_{in} + V_{REF}$ to less than $+V_{REF}$. This means that the decision must be made well *before* V_{in} reaches zero. Similarly, if V_{in} approaches zero from positive values, the decision must appear for a sufficiently positive V_{in} so as to avoid $V_{res} < -V_{REF}$.

These thoughts lead us to the residue behavior depicted in Figure 7(b), where the decision at $V_{in} = 0$ is replaced with one at $V_{in} = -V_{REF}/4$ and another at $V_{in} = +V_{REF}/4$ [9]. We now have

$$V_{res} = f(V_{in}, \pm V_{REF})$$

$$= 2V_{in} - V_{REF} \text{ if } V_{in} > +\frac{V_{REF}}{4} \quad (5)$$

$$= 2V_{in} \text{ if } +\frac{V_{REF}}{4} > V_{in} > -\frac{V_{REF}}{4} \quad (6)$$

$$= 2V_{in} + V_{REF} \text{ if } -\frac{V_{REF}}{4} > V_{in}. \quad (7)$$

In essence, we avoid errors around $V_{in} = 0$ by simply amplifying the input by a factor of 2 and not adding or subtracting V_{REF} . Depicted in Figure 7(c), the implementation employs two comparators and is called the “1.5-bit/stage” architecture. This topology can tolerate a comparator offset as high as $\pm V_{REF}/4$. We say half a bit of “redundancy” is used in each stage so as to accommodate the offset. The low power consumption of comparators

very much favors this topology over the 1-bit/stage counterpart.

The resolution of the stages in a pipeline need not be limited to 1.5 bits. Particularly compelling is the use of a flash stage [more than two comparators in Figure 7(c)] in the front end as it eases the gain and output swings required of the op-amp. For example, a 4-bit flash sub-ADC with 1-bit redundancy allows an eightfold decrease in the op-amp’s open-loop gain while consuming negligible power.

Errors resulting from the capacitor mismatch and finite op-amp gain in Figure 7(b) can be removed by digital calibration. The reader is referred to the vast body of work on this topic [10]–[15]. The performance of pipelined ADCs is limited primarily by that of their constituent op-amps, an issue that has become more serious in low-voltage nanometer technologies and created a migration to SAR architectures.

SAR ADCs

Basic Operation

The SAR ADC, too, performs a binary search to find the closest analog estimate of the input. But, in contrast to the concurrent operation of the MDAC stages in a pipelined architecture, the SAR topology employs a single stage and circulates around it over multiple clock cycles while the analog input is frozen. Shown conceptually in Figure 8, the converter incorporates a comparator, some logic, and a DAC in a (negative) feedback loop. The ADC begins by setting the register’s content to 10000, thus producing $V_{DAC} = V_{REF}/2$. The comparator is then clocked to determine whether V_{samp} is less or greater than this value—as previously illustrated in Figure 3(a). This decision loads the proper value into the register, directing V_{DAC} to $V_{REF}/4$ or $3V_{REF}/4$. To resolve N bits, therefore, the loop takes N clock cycles.

The SAR operation can also be viewed as the convergence of a negative-feedback loop: the high gain of the comparator (which acts as a bang-bang subtractor here) forces V_{DAC} to approach V_{samp} —possibly after some “ringing.”

The SAR architecture derives its beauty from three properties: the ability to achieve high resolutions, the absence of op-amps, and the ability to operate with zero static power dissipation. Moreover, the comparator

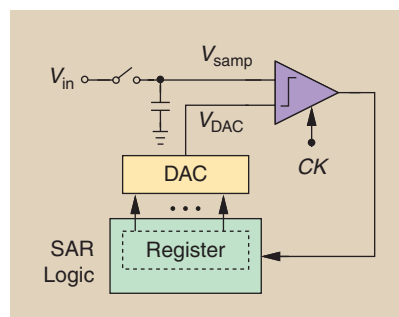


FIGURE 8: Basic SAR architecture.

offset only shifts the overall characteristic rather than distort it. This ADC exemplifies an architecture that is “as digital as it can get.”

SAR ADCs commonly employ capacitor DACs, primarily because the DAC capacitors can also sample the input voltage before the conversion begins. Another advantage of such DACs is their zero static power consumption. Called “charge redistribution” [7], the principle can be illustrated with the aid of Figure 9(a). Here, the DAC consists of binary-weighted capacitors C_1 – C_4 and the dummy unit C_5 . In the acquisition mode, the bottom plates of all capacitors are tied to V_{in} and their top plates to zero.

Next, the bottom plate of C_1 swings to V_{REF} while the others go to ground. The equivalent circuit in Figure 9(b) thus yields $V_X = (8/16)(V_{REF} - V_{in}) + (8/16)(0 - V_{in}) = V_{REF}/2 - V_{in}$. We can see that in this case the voltage at node X is, in fact, the residue (or its negative).

The comparator is subsequently clocked and, according to its decision, one of two actions is taken: (1) If $V_{in} > V_{REF}/2$, then the next residue must assume a value of $3V_{REF}/4 - V_{in}$, which is accomplished by swinging the bottom plate of C_2 to V_{REF} . That is, the next most significant bit (MSB) in the register is raised to ONE. (2) If $V_{in} < V_{REF}/2$, then we must create a new residue equal to $V_{REF}/4 - V_{in}$, by setting the first MSB to ZERO and the next to ONE. Figure 9(c) shows an example of residue settling in this architecture. The conversion continues until the residue falls below 1 LSB. We remark that this architecture resolves 1 bit per cycle.

Viewed as a negative-feedback loop, the charge redistribution topology of Figure 9(a) establishes a virtual ground at node X due to the comparator’s high gain. For this reason, V_X begins at zero and ends near zero, alleviating issues such as the nonlinear junction capacitance of S_1 and the common-mode-dependent offset of the comparator.

In the study of SAR ADCs, we must often monitor the DAC output as a function of time and consider

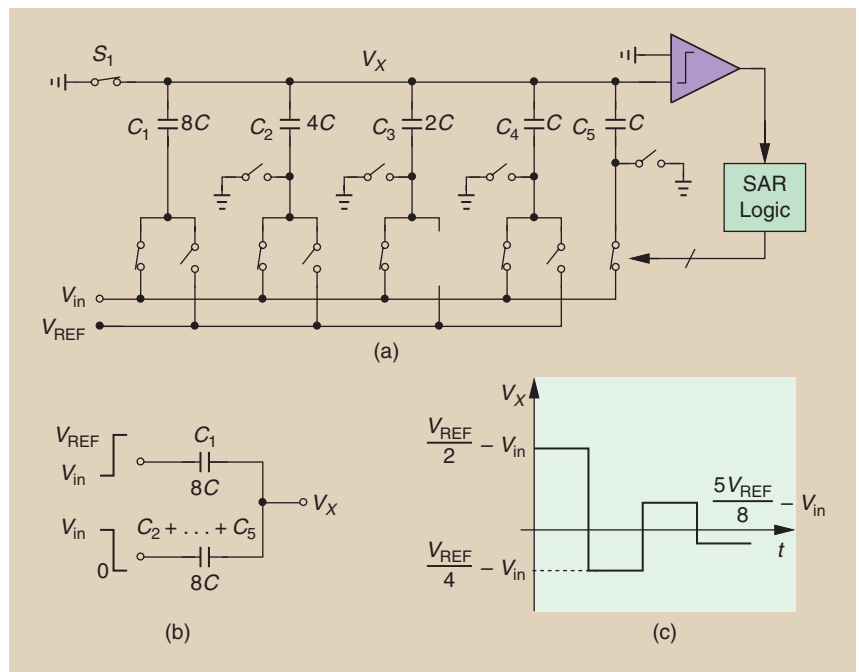


FIGURE 9: (a) A charge redistribution SAR ADC, (b) an operation example, and (c) the virtual ground waveform.

all of its possible trajectories. These trajectories form a “time trellis” as shown in Figure 10, which we will utilize in our time-domain analyses.

SAR Issues

While simple and efficient, SAR ADCs suffer from a number of issues that call for additional circuit and

architecture techniques if a high performance is desired. We describe these issues here and their remedies in the next sections.

First, owing to the multiple clock cycles necessary for each conversion, the SAR architecture is slow. The critical path around the loop in Figure 9(a) entails the comparator

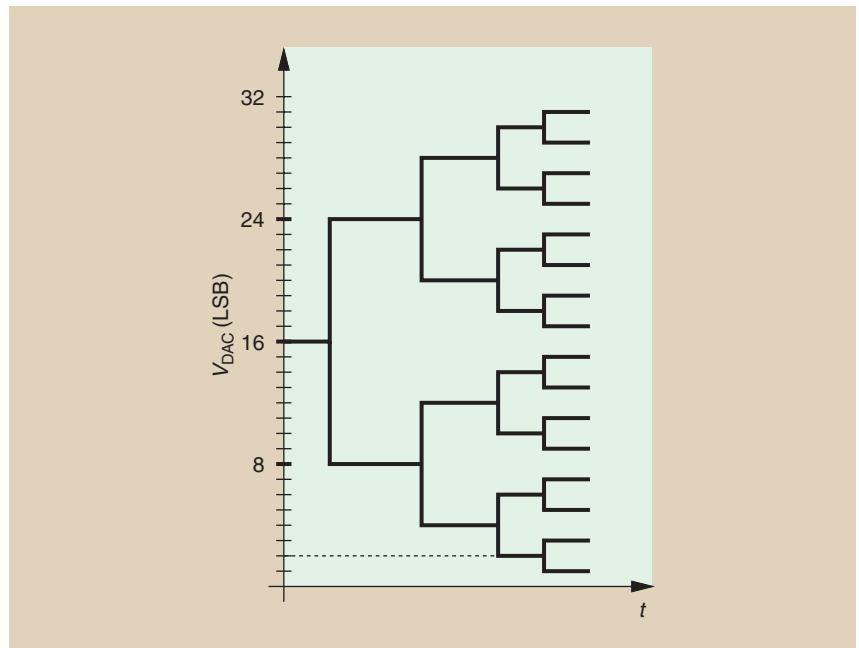


FIGURE 10: A time trellis for studying SAR operation.

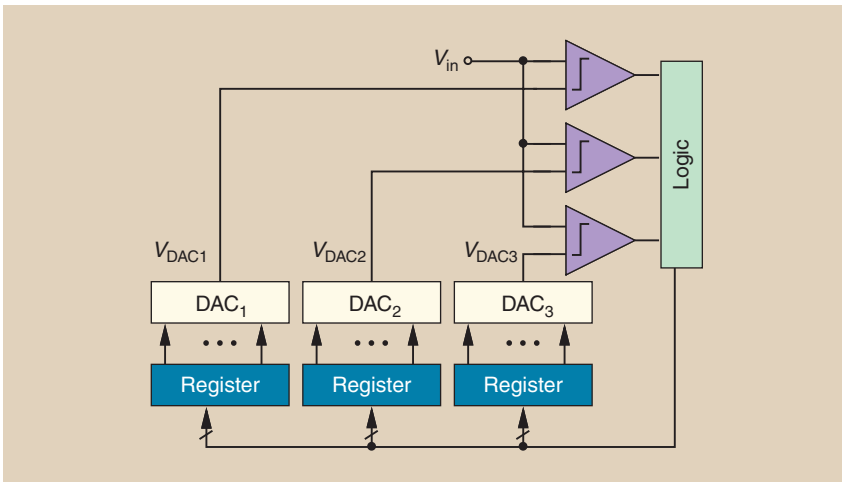


FIGURE 11: Simple 2-bit/cycle SAR ADC.

response time, t_{comp} , the logic delay, t_{logic} , and the settling of the DAC, t_{DAC} . In the 1-bit/cycle architecture studied thus far, one conversion period is approximately equal to $N(t_{comp} + t_{logic} + t_{DAC})$. For a given CMOS technology node, each of these three components has a lower bound regardless of (or weakly dependent on) the power

dissipation, thereby limiting the ADC's speed.

Second, the SAR ADC imposes on its DAC both a high accuracy and a high resolution. For example, a 10-bit design requires 1,024 unit capacitors, or 2,048 if the system is fully differential. By comparison, the pipelined topology of Figure 5(c) incorporates only two capacitors per

stage (with no matching required between the capacitors in one stage and those in the next). Thus, high-resolution SAR ADCs tend to occupy a large area if the DAC unit capacitor size is dictated by matching requirements rather than by kT/C noise.

Third, the lack of residue amplification means that the comparator input noise can limit the performance, as is typical for resolutions of 8 bits and above. A low-noise comparator's response time and power dissipation can far exceed those of a standard design.

Pipelined and SAR architectures also share some issues: the matching of the capacitors must be commensurate with the resolution unless digital calibration is employed, and the circuit generating V_{REF} must exhibit fast settling and low noise.

SAR Speed Improvement

The conversion speed of SAR ADCs can be improved through the use of various techniques [16]–[24]. Three are described here.

The first and most obvious approach is to interleave two or more channels and proportionally raise the conversion rate. The generic interleaving issues apply here as well, including area and input capacitance penalty and the problem of interchannel mismatches. Also

Pipelined and SAR ADC architectures date back to the mid-1900s.

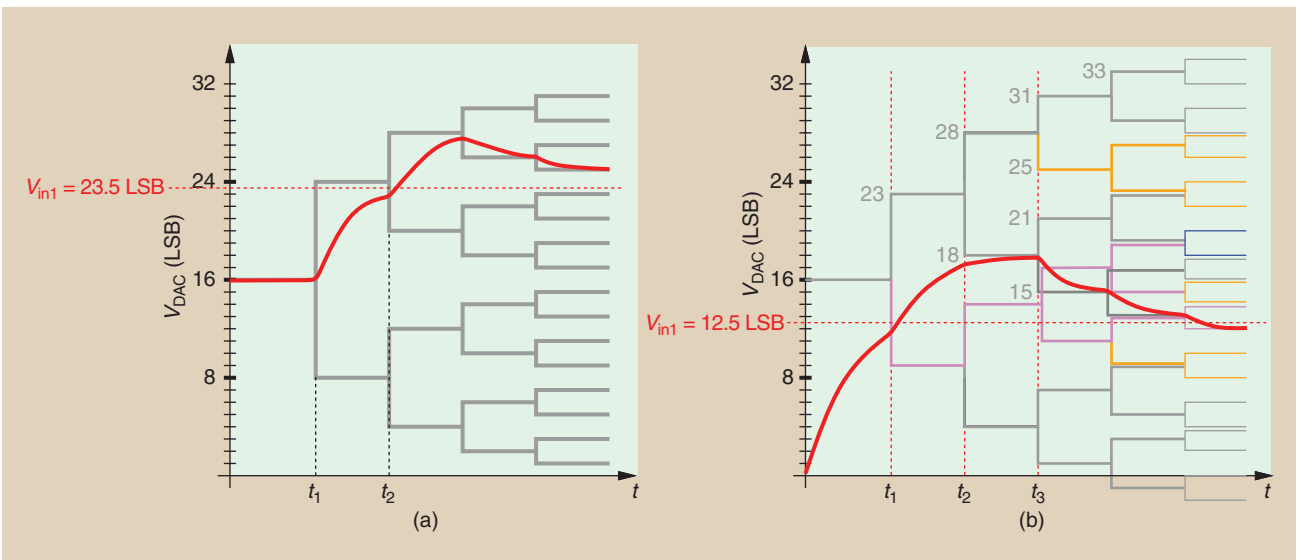


FIGURE 12: (a) SAR response with incomplete DAC settling and (b) use of redundancy to correct the error.

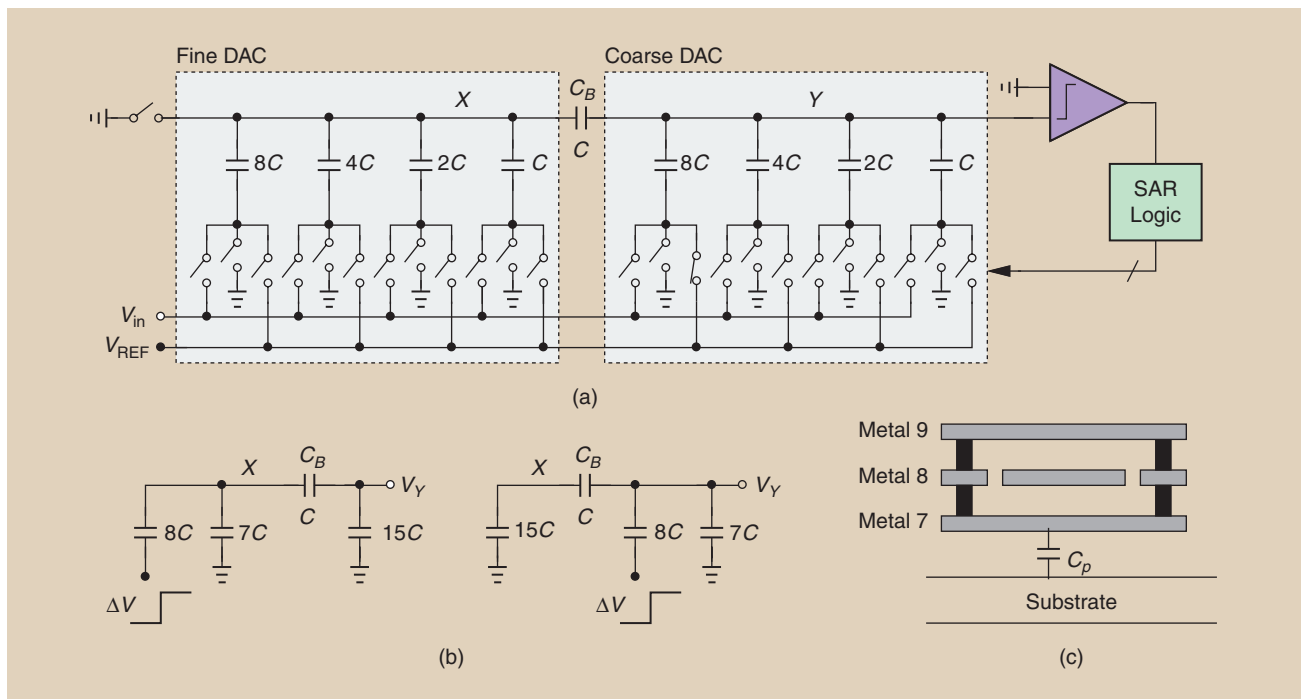


FIGURE 13: (a) A compact DAC implementation using a bridge capacitor, (b) the response to a bottom-plate swing of ΔV , and (c) the realization of C_B for a minimal top-plate parasitic.

problematic is the undesirable coupling of the channels through the reference voltage if it is shared among them.

The second approach is to resolve more than 1 bit per cycle [16]–[18]. For example, three comparators can compare the input with three analog estimates, providing 2 bits of information. But how are the analog estimates created in consecutive cycles? As illustrated in Figure 11, each estimate must be generated by a DAC. In the first conversion cycle, we set $V_{DAC1} = 3V_{REF}/4$, $V_{DAC2} = V_{REF}/2$, and $V_{DAC3} = V_{REF}/4$. If, for example, the comparators decide that $3V_{REF}/4 < V_{in} < V_{REF}$, then the next values of the DAC outputs will be set to $(15/16)V_{REF}$, $(14/16)V_{REF}$, and $(13/16)V_{REF}$, respectively. We observe the considerable growth in the complexity, area, and input capacitance (in the case of the charge redistribution architecture). Furthermore, the comparator random offsets become problematic in this case. It is possible to incorporate a single resistor ladder to serve all three DACs [17], but the switching array necessary to route the ladder taps to the three DAC outputs becomes complex and slow.

Recent work has demonstrated up to 3 bits of resolution per cycle [18].

The third approach relates to the DAC speed and incorporates redundancy [22] to allow incomplete DAC settling in each conversion cycle. To understand this concept, we consider the 5-bit trellis depicted in Figure 12(a), where V_{DAC} begins at $V_{REF}/2 = 16$ LSB but, after the

the design includes no margin (redundancy) for the correction of this error.

Let us ponder how redundancy can be added to accommodate the DAC's incomplete settling. We can make two observations. (1) The comparator decision at $t = t_2$ incorrectly and irreversibly sets 1 bit in the register; the loop has no ability to reset that bit and the resulting

While one MDAC stage is in the amplification mode, the next can be in the acquisition mode and vice versa.

comparator's first decision (at $t = t_1$), does not settle to $3V_{REF}/4 = 24$ LSB in the allotted clock cycle. As a result, the next decision (at $t = t_2$) incorrectly infers that $V_{DAC} < V_{in}$, directing the DAC toward $7V_{REF}/8 = 28$ LSB. Now that V_{DAC} exceeds V_{in} , the loop acts to reduce it in subsequent cycles but, due to the initial error, the final value (25 LSB) remains away from the ideal value (23 LSB). This is because

error in the DAC output persists. (2) Any correction after $t = t_2$ requires extra clock cycles. The redundancy therefore entails both enough over-range "compensation" for the DAC and additional clock cycles. Since the example of Figure 12(a) incorrectly adds 4 LSB to the DAC output, the subsequent cycles must provide at least 4 LSB of correction [23] plus the initial error between 24 LSB and V_{in} .

High-speed designs must also cope with the mismatch between the sampling instants of the capacitors and the comparator.

As an example of SAR redundancy, consider the trellis shown in Figure 12(b) [23], where, for simplicity, we assume the change in V_{DAC} from 0 to 16 LSB is slow. Here, the trellis makes smaller jumps than in binary search but also contains a greater number of conversion cycles. If $V_{in} = 12.5 \text{ LSB}$ and the DAC output does not reach the nominal value of 16 LSB, then at $t = t_2$, an error of $23 \text{ LSB} - 16 \text{ LSB} = 7 \text{ LSB}$ is incurred. The loop—with its high gain—then counteracts this effect by stepping V_{DAC} down by $5 \text{ LSB} + 3 \text{ LSB} + 2 \text{ LSB} + 1 \text{ LSB} + 1 \text{ LSB} = 12 \text{ LSB}$. That is, the loop can correct errors as large as $12 \text{ LSB} - 7 \text{ LSB} = 5 \text{ LSB}$ between 16 LSB and V_{in} . This ability comes at the cost of one extra cycle. The non-binary increments $23 \text{ LSB} - 16 \text{ LSB} = 7 \text{ LSB}$ etc. need no change in the DAC design and can be realized by a look-up table placed after the register [21].

DAC Complexity Reduction

The number of unit capacitors in the DAC can be reduced by various circuit techniques [19], [20], [24]. A common method entails the use of a “bridge” capacitor to scale down the contribution of one section of the DAC to the output, thus creating finer steps. Illustrated in Figure 13(a) for an 8-bit system [24], the idea is to attenuate the effect of the fine DAC output swing by means of C_B . For example, as shown in Figure 13(b), in response to a bottom-plate swing of ΔV , the 8C capacitor in the fine DAC produces $V_Y = (8/255)\Delta V$ whereas the 8C in the coarse DAC yields $V_Y = (16 \times 8/255)\Delta V$. The array therefore achieves 8 bits of resolution using only 31 unit capacitors.

The bridge capacitor nonetheless suffers from parasitics to the substrate,

introducing grounded capacitances at X and Y and causing errors in the DAC output. To address this issue, we construct the capacitor as shown in Figure 13(c), where one plate is realized in metal 8 and is caged by metal-7 and metal-8 layers and vias. This geometry ensures negligible electric field lines going from metal 8 to ground and hence a very small parasitic associated with this terminal. The other terminal of the capacitor does exhibit a parasitic, C_p , to ground, but if attached to node Y as in Figure 13(a), it simply gives rise to a DAC gain error, which is tolerable in a SAR environment.

References

- [1] W. M. Goodall, “Telephony by pulse code modulation,” *Bell Syst. Tech. J.*, vol. 26, pp. 395–409, July 1947.
- [2] J. A. Severin, “Technique for high speed analog-to-digital conversion,” U.S. patent 3,599,204, Aug. 10, 1971.
- [3] N. W. Bell, “Analog-to-digital converter,” U.S. patent 3,059,223, Oct. 16, 1962.
- [4] S. Masuda, Y. Kitamura, S. Ohya, and M. Kikuchi, “A CMOS pipeline algorithmic converter,” in *Proc. CICC*, 1984, pp. 559–562.
- [5] S. H. Lewis and P. R. Gray, “A pipelined 5-M sample/s 9-bit analog-to-digital converter,” *IEEE J. Solid-State Circuits*, vol. 22, pp. 954–961, Dec. 1987.
- [6] W. Kester, “A brief history of data conversion,” this issue.
- [7] J. L. McCreary and P. R. Gray, “All-MOS charge redistribution analog-to-digital conversion techniques—Part I,” *IEEE J. Solid-State Circuits*, vol. 10, pp. 371–378, Dec. 1975.
- [8] B. S. Song, M. F. Tompsett, and K. R. Lakshmikummar, “A 12-bit 1-M sample/s capacitor error-averaging pipelined A/D converter,” *IEEE J. Solid-State Circuits*, vol. 23, pp. 1324–1332, Dec. 1988.
- [9] S. H. Lewis, H. Scott Fetterman, George Gross, and T. R. Viswanathan, “A 10-bit 20-M sample/s analog-to-digital converter,” *IEEE J. Solid-State Circuits*, vol. 27, pp. 351–358, Mar. 1992.
- [10] P. Bogner et al., “A 14b 100-MS/s digitally self-calibrated pipelined ADC in 0.13- μm CMOS,” in *ISSCC Dig. Technical Papers*, Feb. 2006, pp. 156–157.
- [11] C. R. Grace, P. J. Hurst, and S. H. Lewis, “A 12-bit 80-M sample/s pipelined ADCs with bootstrapped digital calibration,” *IEEE J. Solid-State Circuits*, vol. 40, pp. 1038–1046, May 2005.

- [12] 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, pp. 2040–2050, Dec. 2003.
- [13] A. Verma and B. Razavi, “A 10-bit 500-MS/s 55-mW CMOS ADC,” *IEEE J. Solid-State Circuits*, vol. 44, pp. 3039–3050, Nov. 2009.
- [14] O. A. Hafiz, X. Wang, P. J. Hurst, and S. H. Lewis, “Immediate calibration of operational amplifier gain error in pipelined ADCs using extended correlated double sampling,” *IEEE J. Solid-State Circuits*, vol. 48, pp. 749–759, Dec. 2012.
- [15] Y. Miyahara, M. Sano, K. Koyama, and K. Suzuki, “A 14b 60 MS/s pipelined ADC adaptively cancelling op amp gain and nonlinearity,” *IEEE J. Solid-State Circuits*, vol. 49, pp. 416–425, Nov. 2013.
- [16] Z. Cao, S. Yan, and Y. Li, “A 32 mW 1.25 GS/s 6b 2bit/step SAR ADC in 0.13- μm CMOS,” *IEEE J. Solid-State Circuits*, vol. 44, pp. 862–873, Mar. 2009.
- [17] H. Wei, C.-H. Chan, U. Chiao, and S.W. Sin, “An 8-b 400-MS/s 2-b-per-cycle SAR ADC with resistive DAC,” *IEEE J. Solid-State Circuits*, vol. 47, pp. 2763–2772, Sept. 2012.
- [18] C. H. Chan, Y. Zhu, S.W. Sin, S.-P. U, and R. P. Martins, “A 5.5mW 6b 5GS/s 4x-interleaved 3b/cycle SAR ADC in 65nm CMOS,” in *ISSCC Dig. Technical Papers*, Feb. 2015, pp. 466–467.
- [19] 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, pp. 731–740, Apr. 2010.
- [20] S. H. Cho, C. K. Lee, J. K. Kwon, and S. T. Ryu, “A 550- μW 10-b 40-MS/s SAR ADC with multistep addition-only digital error correction,” *IEEE J. Solid-State Circuits*, vol. 46, pp. 1881–1892, Aug. 2011.
- [21] F. Kuttner, “A 1.2V 10b 20 M sample/s non-binary successive approximation ADC in 0.13- μm CMOS,” in *ISSCC Dig. Technical Papers*, Feb. 2002, pp. 176–177.
- [22] Z. Boyacıgiller, B. Weir, and P. D. Bradshaw, “An error-correcting 14b 20 μs CMOS A/D converter,” in *ISSCC Dig. Technical Papers*, Feb. 1981, pp. 62–63.
- [23] T. Ogawa, H. Kobayashi, Y. Takahashi, N. Takai, and M. Hotta, “SAR ADC algorithm with redundancy and digital error correction,” *IEICE Trans. Fundam.*, vol. 93, pp. 415–423, Feb. 2010.
- [24] Y. Chen, Z. Zhu, H. Tamura, M. Kibune, T. Hamada, M. Yoshioka, and K. Ishikawa, “Split Capacitor DAC Mismatch Calibration in Successive Approximation ADC,” *Proc. CICC*, pp. 279–282, Sept., 2009.

About the Author

Behzad Razavi (razavi@ee.ucla.edu) is a professor of electrical engineering at the University of California, Los Angeles, where he conducts research in the area of analog and RF integrated circuits. He has served as an IEEE Distinguished Lecturer and is a Fellow of IEEE. He has received four teaching awards and, with his students, seven best IEEE paper awards. His books have been published in seven languages. He received the 2012 Donald Pederson Award in Solid-State Circuits. 