A Practical Architecture for SAR-based ADCs with Embedded Compressed Sensing Capabilities

(Article begins on next page)

20 April 2024

A Practical Architecture for SAR-based ADCs with Embedded Compressed Sensing Capabilities

Carmine Paolino*, Fabio Pareschi*,‡, Mauro Mangia†,‡, Riccardo Rovatti†,‡, Gianluca Setti*,‡

* DET – Politecnico di Torino, corso Duca degli Abruzzi 24, 10129 Torino, Italy.
email: carmine.paolino@studenti.polito.it, {fabio.pareschi, gianluca.setti}@polito.it
† DEI – University of Bologna, viale Risorgimento 2, 40136 Bologna, Italy. email: riccardo.rovatti@unibo.it
‡ ARCES – University of Bologna, via Toffano 2/2, 40125 Bologna, Italy. email: mauro.mangia2@unibo.it

Abstract—In this paper we propose an innovative A/D architecture with the ability to acquire an input signal according to the recently introduced Compressed Sensing (CS) paradigm. The architecture relies on the hardware blocks already found in traditional successive-approximation-register (SAR) A/D converter, requiring only the addition of a limited number of switches. The capacitive array at the core of the circuit is used both by the SAR conversion algorithm and to realize the linear combination of consecutive signal samples, as required by the CS framework. The lack of additional active blocks allows for a remarkable saving in sampling energy with respect to published solutions. The role of some design parameters is investigated and solutions to ease the circuital implementation are analyzed.

I. INTRODUCTION

Compressed Sensing (CS) is a signal processing technique allowing the representation of a broad family of signals with fewer scalars than what the Nyquist-Shannon theorem suggests [1]. The lower acquisition rate makes the implementation of CS suitable when energy and bandwidth are heavily constrained. Interesting applications have emerged in the biomedical field, with the purpose of realizing ultra-low power biosensor nodes [2]–[4].

In this paper we propose an innovative architecture for CS-based sampling that relies exclusively on the capacitive array found in traditional charge-redistribution successive-approximation-register (SAR) A/D converters. Using only the active elements of the original converter, the power consumption is potentially lower than what other CS solutions proposed in the literature have achieved [2], [3], [5], [6]. The capacitive array is used to: i) sample the input at different time steps; ii) hold the samples until the end of the acquisition window; iii) evaluate their linear combination and iv) convert the result into a digital word.

The drawback of the limited hold-time allowed by the (small) capacitive cells is tackled both at the system level (by modifying the matrix representing the CS acquisition process) and at the circuital level (by introducing a circuit to compensate the leakage currents of the pass transistors).

The paper is organized as follows. Sec. II introduces the basic theoretical framework of CS. In Sec. III we identify the main issues of an analog CS implementation and discuss the first, high-level, solution. Sec. IV describes the proposed circuit and the leakage compensator. Finally, we draw the conclusion.

II. RAKENESS-BASED COMPRESSED SENSING

CS relies on the assumption that the signal to be processed is *sparse*. Mathematically, let x_k, $k \in \mathbb{Z}$, be the discrete-time representation of the input signal, and $x \in \mathbb{R}^n$ be a signal window of n consecutive samples x_k. Let also $D \in \mathbb{R}^{n \times n}$ be the sparsity basis such that $x = D\xi$. The input signal is κ-*sparse* if, for any possible x, the vector $\xi \in \mathbb{R}^n$ containing the projection of x on D has at most κ non-null elements, with $\kappa \ll n$.

Under this assumption, all the information contained in x is captured in a *measurements* vector $y \in \mathbb{R}^m$ such that:

$$y = Ax + \nu = AD\xi + \nu, \tag{1}$$

where $A \in \mathbb{R}^{m \times n}$ is the acquisition matrix, and ν accounts for noise and non-idealities in the acquisition process. Since the number of measurements is $m < n$, the acquisition introduces a compression that we can quantify by means of the Compression Ratio $\mathrm{CR} = m/n$.

Recovering x from y is an ill-posed problem, solvable by looking for the sparsest vector $\hat{\xi}$ over all possible ξ that satisfy (1). Mathematically, this is equivalent to finding a solution to the optimization problem:

$$\hat{\xi} = \underset{\xi \in \mathbb{R}^n}{\operatorname{argmin}} \|\xi\|_1 \qquad \text{s.t.} \ \|AD\xi - y\|_2 < \varepsilon \tag{2}$$

where $\|\cdot\|_p$ is the standard ℓ_p norm, and ε accounts for the effect of ν. The reconstructed signal can be recovered as $\hat{x} = D\hat{\xi}$.

According to the classic CS theory [1], reconstruction is guaranteed by adopting $m = \mathcal{O}\left(\kappa \log(n/\kappa)\right)$ if the elements of A are instances of independent and identically distributed (i.i.d.) random Gaussian values.

Several techniques have been proposed to improve CS performance with respect to the standard approach. In this paper we employ *rakeness*-based CS [4]. It exploits an additional prior, named *localization*, to maximize, on average, the energy collected by each measurement. With the rakeness-based approach, the rows of A are not generated according to an i.i.d. process, but using a multivariate random process defined by a correlation [4]

$$C_A = \frac{1}{2}\left(\frac{C_x}{\mathrm{tr}(C_x)} - \frac{I_n}{n}\right),$$

where I_n is the $n \times n$ identity matrix, C_x is the correlation profile of the instances of x, evaluated as $\mathbf{E}\left[xx^{\mathrm{T}}\right]$, and $tr(\cdot)$ is the trace operator.

Interestingly, CS is still effective when the acquisition matrix is composed of random antipodal values, i.e. $A \in \{-1, +1\}^{m \times n}$, or random ternary values, i.e. $A \in \{-1, 0, +1\}^{m \times n}$, the latter under the assumption that the

Fig. 1. Example of a 8×16 block diagonal sensing matrix A with $n_b = 4$, $m_b = 2$. White blocks correspond to zeroes.


Fig. 2. Performance of a CS-based signal processing chain in terms of ARSNR as a function of CR for a synthetic ECG signal at approx 60 beats/s, with $f_s = 256$ Hz and $n = 256$. Solid lines refer to an ideal system; dotted lines to measurements corrupted by leakage; dashed lines to the proposed architecture with leakage compensation.

number of zeroes is not too high. We refer to [4] for an overview of methods to generate Gaussian, antipodal or ternary sensing sequences with a prescribed C_A.

III. COMPRESSED SENSING USING BLOCK DIAGONAL SENSING MATRICES

By neglecting the noise term, (1) can be formulated component-wise as

$$y_j = \sum_{k=1}^{n} A_{j,k} x_k, \quad j = 1, \ldots, m \qquad (3)$$

where $A_{j,k}$ is the element of A at the intersection of the j-th row and k-th column.

For the sake of a simpler hardware implementation, it is common to limit $A_{j,k}$ to $\{-1, +1\}$, so that the multiply-and-accumulate operations are reduced to simple signed sums [2], [3], [5]. Still, implementing (3) on an analog circuit is quite challenging, in particular when n is large. In fact, the adder required by (3) and the active circuits therein have to be replicated m times. Since large values of n lead to large values of m (assuming a constant CR), such circuits significantly increase the overall power consumption. A high number of terms in the summation worsens the effects of clock feedthrough and charge injection, dramatically increasing the noise component affecting each measurement. At the same time, MOSFET leakage currents introduce a non-negligible signal degradation if the earliest terms in the summation (i.e., $x_1, x_2, \ldots$) are sampled long before the conversion, evaluated at the n-th time step. Targeting the sensing of biomedical signals, leakage currents are indeed the predominant source of noise.

As a practical example, if we consider an ECG signal sampled at $f_s = 256$ Hz, as in [3], taking full advantage of the sparse representation of the signal would require n in the range 200–500. According to (3), using $n = 256$ implies that the sampled value of x_1 has to be preserved for a time period of almost 1 s, a hold time typically unaffordable even for pF-range hold capacitances.

A workaround for both these issues is the design of A as a block diagonal matrix, each block with size $n_b \times m_b$, as in Fig. 1. This allows the reduction of both the number of physical adders (m_b instead of m) and the hold time required ($n_b - 1$ sampling intervals). Furthermore, being $A_{j,k} \in \{-1, 0, +1\}$, the approach ensures both a simple hardware implementation and a good quality of reconstruction. The effectiveness of the solution is confirmed by the empirical results published in the

literature so far, e.g. [5] and [6]. For an exhaustive discussion on the consequences of employing a block diagonal sensing matrix, we refer to [7]. In [4] a method to generate sensing vectors with random ternary values is proposed and it can be easily applied to the block diagonal case.

Numerical results on the effect of using a block diagonal sensing matrix are shown in Fig. 2. They are obtained by simulating the CS-based acquisition and subsequent reconstruction of an artificial ECG signal generated according to [8]. Reconstruction is performed by solving (2) with the SPGL$_1$ toolbox[1]. The sparsity basis D is the orthonormal Symlet-6 wavelet basis [9] and the generation of A follows the rakeness approach outlined in [4]. An additive 50 dB Signal-to-Noise Ratio (SNR) perturbation has been introduced on x to model non-idealities of the system. The figure of merit considered here is the Reconstruction SNR (RSNR), defined as:

$$\text{RSNR[dB]} = 20 \log_{10} \left(\frac{\|x\|_2}{\|\hat{x} - x\|_2} \right)$$

The plot shows the average value of the RSNR (ARSNR) observed over 1000 Monte Carlo trials. The solid lines of Fig. 2 are obtained by an ideal system where no measurements degradation is considered. As expected, a higher n_b results in a better quality of reconstruction, with the full matrix providing the best outcome.

Conversely, the dotted lines include a leakage-discharge model developed from actual data on a 180 nm CMOS technology. The model describes the discharge of the hold capacitors for a realistic configuration of the switches, at 300 pA and in the extremely unfavorable condition of $85\,°C$. The total sampling capacitance C_{tot} of the SAR array is kept constant to emulate equal area occupation and conversion power. Each hold capacitance is therefore $C_h = C_{tot}/n_b$. For a fair comparison with respect to the solution proposed in [3], the value of C_{tot} has been set to 16 pF. The effect on reconstruction quality is opposite with respect to the ideal

[1]http://www.cs.ubc.ca/~mpf/spgl1/

Fig. 3. (a) Proposed architecture during redistribution and conversion in digital form (top); acquisition of modulated signal samples (bottom). Curly braces highlight the splitting of the largest capacitances and the grouping of the smallest ones during acquisition. The structure has $C_{tot} = 8C$, $n_b = 4$ and $C_h = 2C$. Greyed-out elements are not used in the considered phase. (b) Leakage compensator, to be placed around every hold capacitor in the bottom part of (a)

case: a lower n_b shortens the acquisition window, reducing the degradation of measurements due to leakage and resulting in a better reconstruction performance. Yet, it is clear that the original information cannot be recovered, being superimposed to a strong noise component, thus requiring the hardware compensation of leakage currents.

IV. CIRCUIT IMPLEMENTATION

The proposed architecture is shown in Fig. 3a. The top half depicts an example of a traditional charge-redistribution SAR ADC [10] composed of a 3-bit weighted array and a 2-bit C-2C array. Implementing the least significant bits (LSBs) by means of a C-2C structure increases the resolution with low area and energy overhead.

In a switched-capacitors SAR the input signal is typically sampled simultaneously on the entire array, holding the top plates at ground. Opening SW_0, the top plates become an isolated node in which the signal-dependent stored charge is preserved. By driving the input switches sequentially, the SAR algorithm ensures that the array voltage approximates the comparator reference with increasingly finer steps. The conversion process thus requires a number of cycles equal to the number of input capacitors, generating one bit at the end of each cycle.

In the proposed architecture, the conversion behavior is unchanged. However, a few modifications have been introduced during the acquisition phase in order to implement the operations described by equation (3), namely, the decomposition of the largest capacitors and an updated timing of the control signals.

A. Managing a Signal Window on the Capacitive Array

The evaluation of each y_j requires, as a first step, modulation of the input signal by a stream of ± 1. This is allowed by SW_{in}, which selects the direct (V_{in}^+) or inverted (V_{in}^-) replica of the

input according to the value $A_{j,k}$. The presence of the inverted replica comes for free in a differential implementation.

Several modulated samples have to be acquired before conversion, and stored in the array capacitors. Furthermore, in order for the samples to have the same weight on the final result, the storing elements must have uniform value. Therefore the capacitances associated to the most significant bits, which are the largest, are decomposed into smaller elements, each driven by its own set of switches. On the other hand, the entire C-2C sub-array is driven together with some of the smallest scaled capacitors in order to acquire the same sample.

The circuit in the sampling configuration is shown in the bottom half of Fig. 3a. The figure highlights how the capacitance of value $4C$ has been split into two elements, while the entire C-2C sub-array is considered as a whole with the LSB capacitor of the scaled array. The result is that of having $n_b = 4$ hold capacitors of value $C_h = 2C$, each able to store a single modulated sample. Splitting the largest components does not require substantial modification, since ratioed integrated elements are typically realized by several unitary elements connected in parallel. Therefore the only real addition is that of the selection switches to each capacitive sub-element.

Once all the samples in a signal window have been collected, their linear combination is obtained by redistributing the charge stored in the top plates. Using the same notation as in (3), with x_k being a signal sample and $A_{j,k}$ a modulating coefficient, the total charge in the main isolated node of the array in Fig. 3a is

$$Q_j = -C_h \sum_{i=1}^{n_b} A_{j,k} x_k.$$

Opening SW_0 and grounding all bottom plates, the parallel connection of the capacitive elements forces the array voltage to become

$$v_j = \frac{Q_j}{n_b C_h} = -\frac{1}{n_b} \sum_{k=1}^{n_b} A_{j,k} x_k.$$

Apart from the $-1/n_b$ factor, v_j is equivalent to y_j in (3).

From this point on, the array is logically rearranged into its traditional shape (top half of the figure) to finally perform the A/D conversion. This can be achieved by driving simultaneously the sub-elements of a single larger capacitor, so that they behave as the original component, while the LSB switches are now controlled independently.

Concerning the possible decomposition of the C-2C sub-array, since its internal nodes are particularly sensitive to injected noise, it is preferable to leave them untouched. This results in a lower bound for the hold capacitance $C_h \geq C$, which can also be expressed as $C_h = C_{tot}/n_b$.

B. Leakage Compensator

A downside of having $C_h < C_{tot}$ is that the effect of noise, especially leakage, is accentuated. Subthreshold conduction of the pass transistors, as well as the reverse bias currents of their source/drain diffusions, discharge the hold capacitors. Using a larger unitary capacitance, the leakage-induced voltage drop can be reduced at the cost of an increased power consumption and slower operating speed, therefore the need for a leakage-compensation technique.

The architecture reported in [11] and shown in Fig. 3b has been analyzed. Each hold cell is modeled as the parallel connection of the actual hold capacitor C_h, a leakage source I_L and the switches' off-resistance R. The compensator requires a replica of the hold cell, with a scaled capacitance kC_h, $k < 1$, and identical switches so that I_L and R are matched across the cells. Starting from the same initial condition $v_h(0)$, i.e. sampling the same voltage, leakage currents are integrated over time, resulting in a growing voltage difference. Two transconductors inject identical currents, proportional to this difference, in both cells. Intuitively, if the injected currents are larger than the leakage components, the voltage across both cells grows, at a faster rate on the smallest capacitor. Feedback than reduces the compensation current. On the contrary, when the injection is insufficient, both voltages decrease, still faster on the smallest cell. In this case feedback increases the current to inject. Overall, the effect of the feedback loop is to drive the compensation current towards the leakage level, with the only stable condition of both currents being equal.

A more accurate model that considers the effects of R shows that no steady state condition is reached, with the hold voltage described by:

$$
\begin{aligned}
v_h(t) =& v_h(0) \exp\left(-\frac{t}{\tau}\right) u(t) \\
& - RI_L \left[1 - \exp\left(-\frac{t}{\tau}\right)\right] u(t)
\end{aligned}
\tag{4}
$$

with $\tau = G_m R^2 C_h (1-k)$ and $u(t)$ the unitary step function. Voltages vary continuously over time, but the feedback elements can be sized to constrain the variation to an acceptable range. The discharge time constant RC is in fact multiplied by the gain factor $G_m R(1-k)$. In any case, stability of the loop has to be guaranteed. The complete analysis is not included in this paper, but it shows that the capacitive cells behave as a low-pass filter whose pole frequency is controlled by $1/k$. A small k, desirable to minimize the area overhead of the compensator, increases the bandwidth, though bringing the circuit at the onset of instability if the transconductor transfer function singularities are too low in frequency.

With respect to the ADC shown in Fig. 3a, the compensator should be applied across each of the $2C$ cells, therefore, its power consumption has to be extremely low. In [11], compensation of a 0.5 pF capacitor, with $k = 0.2$ and a residual drift of 9.5 mV/sec has been achieved with a current consumption of about 38 nA.

To simulate the effects of compensation, the analytical model in (4) has been applied to the sensing matrix A, obtaining the dashed curves in Fig. 2. They are evaluated considering $G_m = 100\,\mu S$, $R = 1\,G\Omega$ and $k = 0.1$. Having a constant total capacitance of 16 pF, C_h for $n_b = 256$ is too small to allow successful preservation of the samples. As n_b is decreased, performance more closely matches the ideal one. Among the curves shown in the figure, the one for $n_b = 16$ provides the best quality up to $CR = 3$, and even for larger values it is still a suitable candidate for an actual implementation.

V. Conclusion

An innovative switched-capacitor SAR architecture for CS-based acquisitions in the analog domain has been presented. Two techniques to make it practical have been analyzed, namely using a block diagonal sensing matrix and introducing a hardware leakage compensator. Performance gains of these techniques have been validated through algorithmic simulation with real-world parameters, highlighting their importance for the feasibility of the architecture.

References

[1] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. on Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb 2006.

[2] M. Shoaran, M. H. Kamal, C. Pollo, P. Vandergheynst, and A. Schmid, "Compact low-power cortical recording architecture for compressive multichannel data acquisition." *IEEE Trans. Biomed. Circuits and Systems*, vol. 8, no. 6, pp. 857–870, 2014.

[3] F. Pareschi, P. Albertini, G. Frattini, M. Mangia, R. Rovatti, and G. Setti, "Hardware-algorithms co-design and implementation of an analog-to-information converter for biosignals based on compressed sensing," *IEEE Trans. Biomed. Circuits and Systems*, vol. 10, no. 1, pp. 149–162, 2016.

[4] M. Mangia, F. Pareschi, V. Cambareri, R. Rovatti, and G. Setti, "Rakeness-based design of low-complexity compressed sensing," *IEEE Trans. on Circuits and Systems I: Reg. Papers*, vol. 64, no. 5, pp. 1201–1213, 2017.

[5] F. Chen, A. P. Chandrakasan, and V. M. Stojanović, "Design and Analysis of a Hardware-Efficient Compressed Sensing Architecture for Data Compression in Wireless Sensors," *IEEE Journal of Solid-State Circuits*, vol. 47, no. 3, pp. 744–756, Mar. 2012.

[6] D. Bellasi, M. Crescentini, D. Cristaudo, A. Romani, M. Tartagni, and L. Benini, "A broadband multi-mode compressive sensing current sensor soc in 0.16 μ m cmos," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 66, no. 1, pp. 105–118, Jan 2019.

[7] M. Mangia, F. Pareschi, V. Cambareri, R. Rovatti, and G. Setti, *Adapted Compressed Sensing for Effective Hardware Implementations: A Design Flow for Signal-Level Optimization of Compressed Sensing Stages*. Springer International Publishing, 2018.

[8] P. E. McSharry, G. D. Clifford, L. Tarassenko, and L. A. Smith, "A dynamical model for generating synthetic electrocardiogram signals," *IEEE Trans. on Biom. Eng.*, vol. 50, no. 3, pp. 289–294, Mar. 2003.

[9] S. Mallat, *A wavelet tour of signal processing: the sparse way*. Access Online via Elsevier, 2008.

[10] A. Agnes, E. Bonizzoni, P. Malcovati, and F. Maloberti, "A 9.4-enob 1v 3.8μw 100ks/s sar adc with time-domain comparator," in *2008 IEEE International Solid-State Circuits Conference - Digest of Technical Papers*, Feb 2008, pp. 246–610.

[11] L. S. Y. Wong, S. Hossain, and A. Walker, "Leakage current cancellation technique for low power switched-capacitor circuits," in *Proc. of the 2001 Int. Symp. on Low Power Electronics and Design*, 2001, pp. 310–315.