

Design and Exploration of Low-Power Analog to Information Conversion Based on Compressed Sensing

Hossein Mamaghanian, *Student Member, IEEE*, Nadia Khaled, *Member, IEEE*, David Atienza, *Member, IEEE*, and Pierre Vandergheynst, *Senior Member, IEEE*

Abstract—The long-standing analog-to-digital conversion paradigm based on Shannon/Nyquist sampling has been challenged lately, mostly in situations such as radar and communication signal processing where signal bandwidth is so large that sampling architectures constraints are simply not manageable. Compressed sensing (CS) is a new emerging signal acquisition/compression paradigm that offers a striking alternative to traditional signal acquisition. Interestingly, by merging the sampling and compression steps, CS also removes a large part of the digital architecture and might thus considerably simplify analog-to-information (A2I) conversion devices. This so-called “analog CS,” where compression occurs directly in the analog sensor readout electronics prior to analog-to-digital conversion, could thus be of great importance for applications where bandwidth is moderate, but computationally complex, and power resources are severely constrained. In our previous work (Mamaghanian, 2011), we quantified and validated the potential of digital CS systems for real-time and energy-efficient electrocardiogram compression on resource-constrained sensing platforms. In this paper, we review the state-of-the-art implementations of CS-based signal acquisition systems and perform a complete system-level analysis for each implementation to highlight their strengths and weaknesses regarding implementation complexity, performance and power consumption. Then, we introduce the spread spectrum random modulator pre-integrator (SRMPI), which is a new design and implementation of a CS-based A2I read-out system that uses spread spectrum techniques prior to random modulation in order to produce the low rate set of digital samples. Finally, we experimentally built an SRMPI prototype to compare it with state-of-the-art CS-based signal acquisition systems, focusing on critical system design parameters and constraints, and show that this new proposed architecture offers a compelling alternative, in particular for low power and computationally-constrained embedded systems.

Index Terms—Analog-to-information, compressed sensing (CS), electrocardiogram (ECG) compression, low-power, random modulation pre-integrator (RMPI), spread spectrum, system-level design.

Manuscript received June 01, 2012; revised July 26, 2012; accepted August 23, 2012. Date of publication October 18, 2012; date of current version December 05, 2012. This work was supported in part by the EU under research projects FET-OPEN 225913 (SMALL) and 255931 (UNLocX). This paper was recommended by Guest Editor G. Setti.

H. Mamaghanian, D. Atienza, and P. Vandergheynst are with the School of Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland.

N. Khaled is with the Bioanalytical Science Department, Nestlé Research Center, 1000 Lausanne, Switzerland (e-mail: nadia.khaled@rdls.nestle.com).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JETCAS.2012.2220253

I. INTRODUCTION

SENSING and processing information have traditionally relied on the Shannon sampling theorem, one of the central tenets of digital signal processing. This theorem states that, given a signal of bandwidth Ω , it is sufficient to sample it at twice its the bandwidth (i.e., the Nyquist rate) to ensure faithful representation and reconstruction. However, this traditional analog-to-digital conversion paradigm has been challenged lately. First, there are many situations where Ω is so large that constraints put on sampling architectures are simply unbearable. Systems working in radio-frequency (RF) bands for example, require high sampling rates that severely stress current ADC technologies. Second, even for relatively low signal bandwidths, Nyquist-rate sampling might produce a large amount of redundant digital samples, which causes bottlenecks in systems with limited resource and power storage.

In the absence of extra information, Nyquist rate sampling is essentially optimal for band limited signals, but in many application signals are known to have strong structures that can be exploited with dedicated, often nonlinear, sampling schemes. The recent theory of compressed sensing (CS) [1]–[3], showed that signals that can be represented by sparse superpositions of elementary waveforms can be sampled linearly at rates proportional to their sparsity, which can be dramatically smaller than their Nyquist frequency. This is particularly promising for three main reasons. First, the class of sparse signals is quite broad and covers a wide range of applications, from biosignals, to medical imaging. Second, the linear sampling strategy advocated by CS is universal: the same encoder can be used to sample signals drawn from diverse dictionaries of waveforms. Third, CS is robust: it offers performance guarantees for signals that are only approximately sparse and it is resilient to noise added on the signals or on the collected samples. There is of course a price to pay for these advantages. The decoder in a CS system can be complex, usually it has to solve a convex optimization program to recover the signal, and CS is therefore best suited for applications where constraints are on the encoder. More importantly, the sampling architecture must be redesigned and studied from the ground up; this is the focus of this paper.

There has been a number of interesting initiatives aimed at implementing sub-Nyquist signal acquisition systems. These efforts have resulted in CS-based architectures such as the random demodulator (RD) [5], the modulated-wideband converter (MWC) [6], and others [7]–[10]; see [11] for a complete review.

The present work aims at exploring and designing a low-power (and scalable in implementation complexity) analog front-end for compressive signal acquisition. Our main target applications are embedded systems for ambulatory bio-signal monitoring. In this setting, the concern is not so much high bandwidth, but rather very restricted computational resources and power profiles. We therefore explore simple architectures with a strong emphasis on low-power components. To the best of our knowledge, the present work introduces two main contributions. First, we propose the spread spectrum random modulator pre-integrator (SRMPI), which is a new CS architecture based on the random demodulation techniques. This new design reduces the complexity of state-of-the-art CS-based A2I system implementations while also reducing their power consumption. Second, in order to quantify the potential of CS for low-power sensing and compression, an exhaustive performance and power analysis comparison is performed between different CS-based architectures implementations. In particular, our work carefully quantizes the power figure for our proposed SRMPI hardware implementation based on spread spectrum techniques [12] with the state-of-the-art random modulation pre-integrator (RMPI) [5]. This comparison confirms the potential of our proposed SRMPI architecture as a low power CS-based analog front-end for embedded applications.

Notations: In all the following, lower case letters designate scalar quantities, boldface lower case letters indicate column vectors, and boldface capitals represent matrixes. $\langle \mathbf{a}, \mathbf{b} \rangle$ stands for the inner product of vectors \mathbf{a} and \mathbf{b} . Moreover, m_i and $\mathbf{M}_{i,j}$ are the i th entry of vector \mathbf{m} and the (i,j) th entry of matrix \mathbf{M} , respectively. Finally, $(\cdot)^H$ and $\|\cdot\|_p$ denote the conjugate transpose, and the l_p -norm of a vector, respectively.

II. COMPRESSED SENSING AND SPARSE RECOVERY

The main idea behind CS is now quite well known but, for the sake of completeness, we summarize the main concepts. More details can be found in [1]–[3]. Let $\mathbf{x} = \Psi\alpha$ be a real-valued N -dimensional discrete signal vector ($\mathbf{x} \in \mathbb{R}^N$) that is *compressible* in some orthonormal basis $\Psi = [\psi_1 | \psi_2 | \dots | \psi_N]$, where each column is a vector ψ_i , and α represents the N -dimensional coefficient vector. By compressible we mean that the entries of $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_N]$, when sorted in decreasing order of magnitude, decay rapidly to zero; any such a signal is well approximated using a K -term approximation, consisting of the K largest entries of α and setting all other terms to zero, $\mathbf{x} \approx \sum_{k=1}^K \alpha_{r(k)} \psi_{r(k)}$, with $K \ll N$ and $|\alpha_{r(k)}| \geq |\alpha_{r(k')}|$. In essence, compressible signals are well approximated by sparse signals.

Conventionally, one would collect signal samples at the Nyquist rate forming \mathbf{x} and *then* compress it using nonlinear digital compression techniques. CS offers a striking alternative by showing that if \mathbf{x} is compressible, one can recover to a K -term approximation by only collecting roughly $M \approx K$ samples using simple *analog* measurement waveforms, thus sensing/sampling and compressing at the same time. More precisely, we collect $M = O(K \log N/K)$ samples by projecting on sensing waveforms $\{\phi_i\}_{1 \leq i \leq M}$ thus forming the measurement vector $y_i = \phi_i^H \mathbf{x} = \langle \phi_i, \mathbf{x} \rangle, i = 1, \dots, M$. Consequently, the CS linearly compressed data vector $\mathbf{y} \in \mathbb{R}^M$

is described by $\mathbf{y} = \Phi\mathbf{x}$, where Φ denotes the $M \times N$ measurement or sensing matrix with the vectors $\phi_1^H, \dots, \phi_M^H$ as rows. It is important to notice that the sensing matrix Φ does not depend on the signal: CS proposes a simple linear sampling strategy that is only marginally off the optimal but complex best adaptive strategy. To guarantee the robust and efficient recovery of any S -sparse signal, the sensing matrix Φ must obey the key *restricted isometry property* (RIP) [13]

$$(1 - \delta_S) \|\alpha\|_2^2 \leq \|\Phi\mathbf{x}\|_2^2 \leq (1 + \delta_S) \|\alpha\|_2^2 \quad (1)$$

for all S -sparse vectors α . The isometry constant δ_S of matrix Φ must not be too close to one. This property is difficult to verify in practice and it is often replaced by the requirement that the sensing matrix Φ and sparsity basis Ψ must be incoherent [14], [15], i.e., their coherence

$$\mu(\Phi, \Psi) = \sqrt{N} \cdot \max_{1 \leq k, j \leq N} |\langle \phi_k, \psi_j \rangle| \quad (2)$$

is small enough. A universal good choice for the sensing matrix Φ are random matrixes, such as random matrixes with independent identically distributed (i.i.d.) entries formed by sampling: 1) a Gaussian distribution $\mathcal{N}(0, 1/N)$; 2) a symmetric Bernoulli distribution ($P(\Phi_{i,j} = \pm 1/\sqrt{N}) = 1/2$).

If the RIP holds, then accurate reconstruction can be accomplished by solving the following convex optimization problem:

$$\min_{\tilde{\alpha} \in \mathbb{R}^N} \|\tilde{\alpha}\|_1 \quad \text{subject to} \quad \|\Phi\mathbf{x} - \mathbf{y}\|_2 \leq \sigma \quad (3)$$

where σ bounds the amount of noise unavoidably corrupting the data. Many algorithms were introduced to solve this reconstruction problem, including interior-point algorithms [16], [17], gradient projection [18], iterative thresholding [19], and greedy approaches such as orthogonal matching pursuit (OMP) [20], [21]. Our results are based on the basis pursuit denoise algorithm provided in the SPGL1 solver [22].

III. METHODS

The CS sampling and reconstruction steps described in Section II are depicted in Fig. 1. As shown, sampling is essentially a linear projection from an N -dimensional original signal space down to a lower M -dimensional measurement space using a sensing matrix Φ . Since the introduction of CS, several hardware implementations have been proposed in the literature. Some of these implementations are purely digital: CS is used as a simple compression technique that does not require a complex transform. In these cases, CS is applied after Nyquist sampling. More interesting for us is the analog version of CS, where linear measurement \mathbf{y} are collected in the analog domain prior to digitization. Since AD conversion is often the dominant source of power consumption [4], analog implementations of CS can offer interesting low power alternatives since they inherently allow using low rate ADCs.

In this work, our main focus is therefore on the analog implementation of CS. In this section, we first describe a model for sparse signals. We then review the RD architecture, which was the first proposed hardware implementation for CS. We discuss extensions of RD and highlight the main difficulties arising for hardware implementation and limits to their performances.

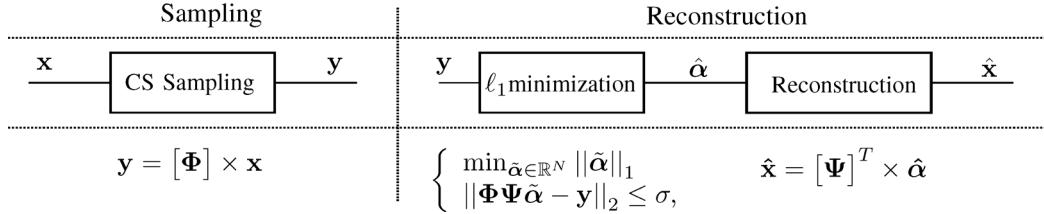


Fig. 1. Block diagram of CS sampling and reconstruction. \mathbf{x} is the original signal being sampled, \mathbf{y} is linear measurements collected using sensing matrix Φ . Ψ is the signal basis under which \mathbf{x} is sparse, and $\boldsymbol{\alpha}$ is the resulting coefficients and $\hat{\boldsymbol{\alpha}}$ the sparse solution found in optimization and $\hat{\mathbf{x}}$ is the reconstructed signal.

A. Signal Model

Suppose our analog signal x has a sparse representation in some basis or dictionary as mentioned before, i.e., the signal can be represented using S parameters per unit of time in some continuous-time basis. More concretely, let the analog signal $x(t)$ be expanded via a basis $\psi_n(t)$, $n \in \{1, 2, \dots, N\}$

$$x(t) = \sum_{n=1}^N \alpha_n \psi_n(t), \quad \text{for } t \in [0, 1]. \quad (4)$$

with $t, \alpha_n \in \mathbb{R}$ and N is analogous to the signal bandwidth. Sparsity means that only a fraction $S \ll N$ of the entries of $\boldsymbol{\alpha}$ are nonzero, i.e., the signal in each time frame is only composed of a few waveforms.

According to the sampling theorem, we should sample signals of the form 4 by sampling at twice the bandwidth. However, these signals have only few degrees-of-freedom. In consequence, it is reasonable to expect that we can acquire them by sampling at roughly the sparsity level S . Stirling's approximation shows that there are about $\exp\{S \log(N/S) + O(S)\}$ ways to select S distinct integers in the range of $\{1, 2, \dots, N\}$. Therefore, it takes $O(S \log(N/S))$ bits to encode the nonzeros elements in the coefficient vector $\boldsymbol{\alpha}$. Ideally this is our bound for the necessary amount of measurements M to solve (3) and reconstruct the S -sparse approximation of the original signal $x(t)$.

B. Random Demodulator

Let us now describe the RD architecture for sampling an analog sparse signal [5], as depicted in Fig. 2. The first stage is a demodulator whose input signal $x(t)$ is multiplied by a continuous time sequence of pseudo-random numbers $P_c(t)$ to obtain a continuous time demodulated signal $z(t)$. Starting with a sequence of pseudo random numbers of $d_n = \{d_0, d_1, \dots, d_{N-1}\}$ that take values ± 1 with equal probability, it is used to create a continuous chipping sequence

$$P_c(t) = d_n, \quad t \in \left[\frac{n}{N}, \frac{(n+1)}{N} \right) \text{ and } n = 0, 1, \dots, N-1. \quad (5)$$

This (ideal) demodulation signal takes values ± 1 over each time frame and switches between the levels randomly at or faster than the Nyquist rate of the input signal x . The final stage is a standard ADC to sample the signal. A low-pass filter is used prior to ADC to prevent aliasing.

The demodulator

$$z(t) = x(t) \cdot P_c(t), \quad t \in [0, 1] \quad (6)$$

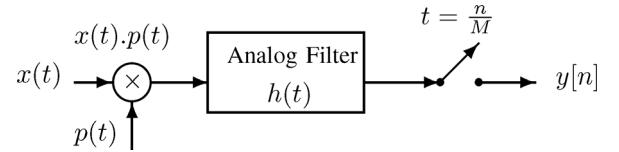


Fig. 2. Block diagram of the RD. The system includes a pseudo-random number generator, an analog anti-aliasing filter and a sampler.

acts by spreading the frequency content of the signal so that it is not destroyed when it is low pass filtered prior to sampling at rate $M \ll N$. In a discrete time form this demodulation action corresponds to the map $x \rightarrow z = \mathbf{D}x$ where

$$\mathbf{D} = \begin{bmatrix} d_0 & & & \\ & d_1 & & \\ & & \ddots & \\ & & & d_{N-1} \end{bmatrix}. \quad (7)$$

The low pass filter is used as a simple accumulator that sums the demodulated signal for duration of $(1/M)$. The filtered signal $y(t)$ is sampled every $(1/M)$ seconds to obtain the measurements vector \mathbf{y} . After each sample, the low pass filter is reset. In summary

$$y[m] = \int_{-\infty}^{\infty} x(\tau) P_c(\tau) h(t - \tau) d\tau \Big|_{t=\frac{m}{M}} \quad m \in \{0, 1, \dots, M-1\}. \quad (8)$$

This approach is called *sample and dump* sampling. Suppose that number of measurements M divides N , then each measurements is the sum of N/M consecutive entries of the demodulated signal. Therefore, we can represent the action of accumulating and sampling as an $M \times N$ matrix \mathbf{H} , whose r th row has N/M consecutive entries starting in column $rN/M + 1$. An example for $M = 2$ and $N = 6$ is

$$\mathbf{H} = \begin{bmatrix} \beta_0 & \beta_1 & \beta_2 & 0 & 0 & 0 \\ 0 & 0 & 0 & \beta_0 & \beta_1 & \beta_2 \end{bmatrix} \quad (9)$$

where β_r are the filter coefficients, which in case of an ideal accumulate-and-dump sampler are all unit entries. Inserting the signal model (4), we can further simplify and write

$$y[m] = \sum_{n=1}^N \alpha_n \int_{\frac{m}{M}}^{\frac{(m+1)}{M}} \psi_n(\tau) P_c(\tau) h\left(\frac{m}{M} - \tau\right) d\tau. \quad (10)$$

It is now clear that we can rewrite this equation in CS matrix form. Concatenating the sampling matrix and basis $\mathbf{V} = \Phi\Psi$,

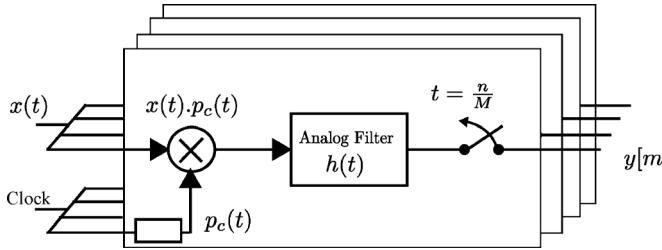


Fig. 3. Block diagram of random demodulator pre-integrator.

where $\Phi = HD$ and $\mathbf{V} \in \mathbb{R}^{M \times N}$, then element $V_{n,m}$ for row m and column n reads

$$V_{n,m} = \int_{\frac{m}{M}}^{\frac{(m+1)}{M}} \psi_n(\tau) P_c(\tau) h\left(\frac{m}{M} - \tau\right) d\tau. \quad (11)$$

As it is clear from (9), the rows of the sensing matrix Φ are coherent and this coherency increases, as we decrease the number of measurements M . This is one of the disadvantages of the RD sampler.

C. Random Demodulation Pre-Integrator

The RMPI is a variant of the RD architecture, which is composed of parallel channels of RD. Each block on Fig. 3 represents a RD channel. The input signal $x(t)$ is fed into channels in parallel and it is demodulated with a unique sequence of pseudo random numbers. Then, as in a normal RD structure, it is low pass filtered and sampled. Here again the continuous time pseudo random signals $P_c(t)$ are chipping signals that take values ± 1 and simply just invert the input signal polarity. The alternating frequency of $P_c(t)$ is at the Nyquist rate and the sequences are independent for each channel. The RMPI structure allows to further reduce the ADC rate by reducing the coherency between rows of the sensing matrix for the same number of measurements compared to the RD architecture. However, RMPI needs more (one for each channel) multiplexers or mixers (for implementing the demodulator section), each of them operating at the Nyquist rate.

Suppose C is the number of channels and $M_c = \{M_0, M_1, \dots, M_{C-1}\}$ is the number of measurements for channel c . We can again rewrite the CS matrix using (13). Now, the element $V_{n,m}^c$ for row m and column n and channel c of the equivalent CS matrix is

$$V_{n,m}^c = \int_{\frac{m}{M_c}}^{\frac{(m+1)}{M_c}} \psi_n(\tau) P_c(\tau) h_c\left(\frac{m}{M} - \tau\right) d\tau \quad (12)$$

where we have $\mathbf{V}^c \in \mathbb{R}^{M_c \times N}$, and the CS matrix for the whole system is

$$\mathbf{V} = \begin{pmatrix} \mathbf{V}^0 \in \mathbb{R}^{M_0 \times N} \\ \mathbf{V}^1 \in \mathbb{R}^{M_1 \times N} \\ \vdots \\ \mathbf{V}^{C-1} \in \mathbb{R}^{M_{C-1} \times N} \end{pmatrix}. \quad (13)$$

In a recent work [23] comparing power consumption profiles for both digital and analog CS, the power efficiency of the RMPI architectures is questioned. The authors claim that their results show that a digital CS implementation can ultimately

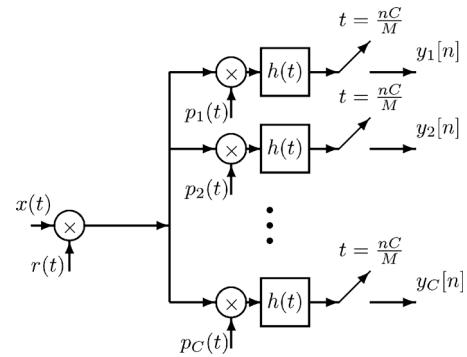


Fig. 4. Block diagram of SRMPI.

outperform analog implementations in terms of overall energy efficiency for data compression in wireless sensors due to the power consumption of amplifiers. However, the RMPI architecture used in their work is a special case where the number of parallel channels is equal to the number of measurements. Thus, an array of M (equal to the number of needed measurements) amplifiers is needed. Moreover, the use of a higher working frequency for the mixers (which should be more than or at least equal to Nyquist rate) makes the analog implementation more costly in terms of power consumption. Hence, the reduction of the sampling rate of the ADC in RMPI or RD structures is done at the expense of using a large number of high frequency mixers and amplifiers (i.e., equal to the number of channels). Although from the circuit viewpoint building the mixer blocks is easy and far simpler than a high rate ADC, there is a significant overhead in terms of overall power consumption. Thus, the application of RMPI seems limited to cases where the signal bandwidth puts unbearable constraints to state-of-the-art ADCs. However, there is a whole realm of applications where signals have a moderate bandwidth, but traditional ADC designs are still too power hungry. In that context, analog CS can offer an interesting alternative, if more strategies, like the proposed SRMPI architecture for power reduction are used.

IV. SPREAD SPECTRUM RANDOM MODULATION PRE-INTEGRATOR

To overcome these problems of the RMPI architecture, we propose a new architecture called *spread spectrum random modulator pre-integrator* (SRMPI), which includes an initial pre-modulation block as fundamental change with respect to the RMPI architecture. The proposed SRMPI design (like RMPI) is a “universal” encoder, which unlike other architectures (like RM) works with signals that are sparse in any fixed domain. Fig. 4 shows the block diagram of the SRMPI architecture. This new premodulation block modulates the original signal $x(t)$ with a random sequence $R(t)$, similar to what is done in the RD structure. Hence, based on the same principle, this modulator block should operate at a rate least equal to Nyquist range and the information of the signal is thus spread over the whole frequency spectrum. The random-modulated signal is then fed to a regular RMPI structure. Thanks to the premodulation block and the signal sparsity features, it is possible to lower the internal channel modulators’ working frequency in the SRMPI

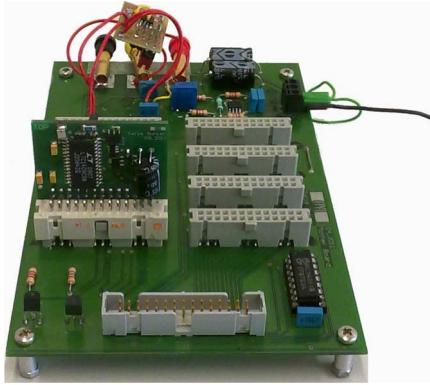


Fig. 5. Analog CS data acquisition board.

design significantly below the Nyquist rate. In fact, the limit to reduce the internal modulators clock depends on the signal sparsity. The equivalent sensing matrix is similar to the RMPI one (12), except that in this case another demodulator should be introduced

$$V_{n,m}^c = \int_{\frac{m}{M_c}}^{\frac{(m+1)}{M_c}} \psi_n(\tau) R(\tau) P'_c(\tau) h_c \left(\frac{m}{M} - \tau \right) d\tau \quad (14)$$

where $R(t)$ and $P'_c(t)$ are the continuous time sequence of random numbers, and where the internal demodulator $P'_c(t)$ can work at rates lower than Nyquist rate. Interested readers can refer to the work by [24] and its references, where the properties of spread spectrum are carefully studied. In this paper, we show the results for the SRMPI architecture (cf. Section VI) having the internal modulators working at half the Nyquist rate.

V. EXPERIMENTAL SETUP

In this section, we present our setup to explore the analog CS design space. The setup is a circuit level implementation of both RMPI and SRMPI as described in Sections III-C and IV. Our hardware implementation consists of a signal acquisition board and a PC with a *Data Acquisition Card* (DAC). The signal acquisition board is a circuit level implementation of an analog CS system with eight channels (Fig. 5). It consists of a main board with eight slots (one per channel) supporting daughter boards implementing each channel architecture. The main board also has a slot for the DAC card. The communication with the PC is done through this card. Each channel includes a modulator at the beginning of the line and a low-pass filter. Each channel has a unique ADC which samples the output of the low-pass filter. The difference between the RMPI and SRMPI structure in our hardware implementation lies on the common modulator on the main input signal line. To test both architectures, this modulator is implemented on the main board and can be bypassed to give the possibility of testing both architectures.

A. Modulator

In both architectures, one of the essential blocks is the random modulator. Since the modulation alternates between values of ± 1 with equal probability, it is simply implemented by changing the polarity of the input signal. In recent works, different circuits have been proposed to implement such modulators [7], [11].

However, the main concern about the implementation of mixers is that the transient time (switching time) should be significantly low compared to the sampling period. The output of the modulation is subsequently integrated (filtered) and any distortion in the switching transient time is treated as additive random noise. Since in our platform the signal bandwidth is in the hertz range, this is not a challenge for the proposed SRMPI implementation. The modulators for both architectures have been implemented using the single pole, double throw (SPDT) CMOS Switch *ADG636* from *Analog Devices*.

In many cases, specially in RF applications with higher sampling frequency, the implementation and calibration of fast mixers with these properties is a crucial part in the design. In these cases, an architecture like SRMPI can drastically ease the problem by reducing the number of high frequency mixers with respect to RMPI without affecting the performance of the overall system as we shall see below.

B. Integrator

Another block in both architectures is the integrator, which is used to integrate the output signal after modulation. An integration in the time domain is equivalent to a filter with frequency response $H(S) = 1/S$. But building an ideal integrator is not realizable, and instead a low pass filter used as an integrator. The exact characteristics of this filter must be precisely known at the decoder to implement the full sensing matrix. The only constraint is to have a filter that prevents aliasing before sampling. From an implementation viewpoint, the characteristics of the filter should be stable over time (i.e., time invariant).

In this work, we have used a single-pole filter with custom cutoff frequency $f_c \in [9.52, 497.7]$ Hz with 256 steps. It was implemented via a configurable potentiometer MAX5387NAUD with 256 steps. The transfer function of a single-pole filter is $H(S) = 1/(S + a)$. The pole indicates the time constant of the impulse response $h(t) = e^{-at}$. The ideal integrator corresponds to $a = 0$ and if a is large the impulse response $h(t)$ will decay very rapidly and the filter will quickly forget the past time information. The main concern is to push the pole close to zero.

In both architectures, when a sample is taken, the integrator is reset. A reset control circuit is added to the design to precisely control the timing of this operation. The filter transient time should be small enough compared to the ADC sampling period not to affect the integration. Thus, synchronizing the resets is very important and can introduce more errors in the measurements. In our design, the reset control command along with ADC commands for each channels are sent from the DAC card.

C. ADC and Sampling

A low-power *Linear Technology LTC 1409* ADC is used to sample the output of the filters in each channel. The LTC1409 is a 12-bit ADC with sampling rate of 800 ks/s. To initiate sampling, the DAC card triggers the operation by writing zero on conversion pin of the ADCs. When the sample is taken, it is on hold on the ADC and later DAC the card collects the sampled data of the channels sequentially. The sampling time should be precisely adjusted since any jitter or undesired deviations will introduce more distortion to the measurements. It is a serious challenge in hardware design to reduce the timing deviations in

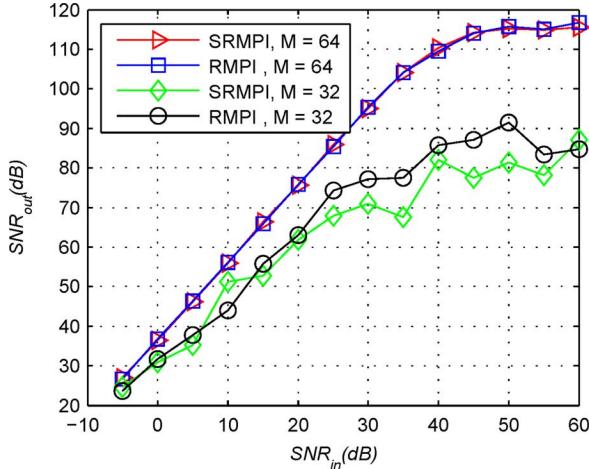


Fig. 6. SNR_{out} versus SNR_{in} for both RMPI and SRMPI architectures with 8 channel for number of measurement $M = 32$ and 64 .

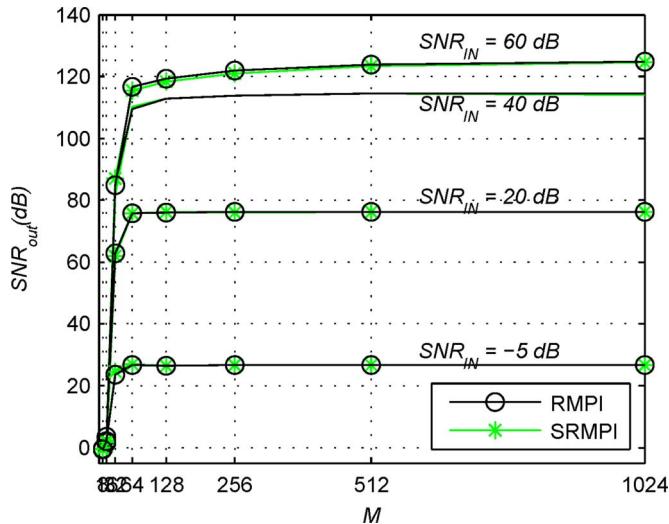


Fig. 7. SNR_{out} versus number of measurements for both RMPI and SRMPI architectures with 8 channel for different input SNR range.

a system. The values of timing deviations should be insignificant compared to the sampling period of the ADC to have valid measurements. In our setup, since we are working at low frequencies the problem is not critical. The data transfer rate for the DAC card is at 80 MHz and far beyond the sampling rate of the ADC. This will ensure that the amount of noise added by sampling is insignificant compared to other parts.

VI. EXPERIMENTAL RESULTS

Fig. 6 compares the output signal-to-noise ratio (SNR), averaged over 100 packets of 2 s of data, for RMPI and SRMPI architectures over a large range of input signal quality. These results were obtained for $M = 32$ and $M = 64$ samples. For SRMPI the internal modulators' frequency is fixed at half the Nyquist rate (512 Hz) and the input signal is composed of $k = 3$ frequency tones. This plot confirms that both architectures could reach the same performance with equal number of measurements and channels. The plot also shows the success of both architecture at recovering the compressed signals even under strong noise corruption. Fig. 7 shows the output signal

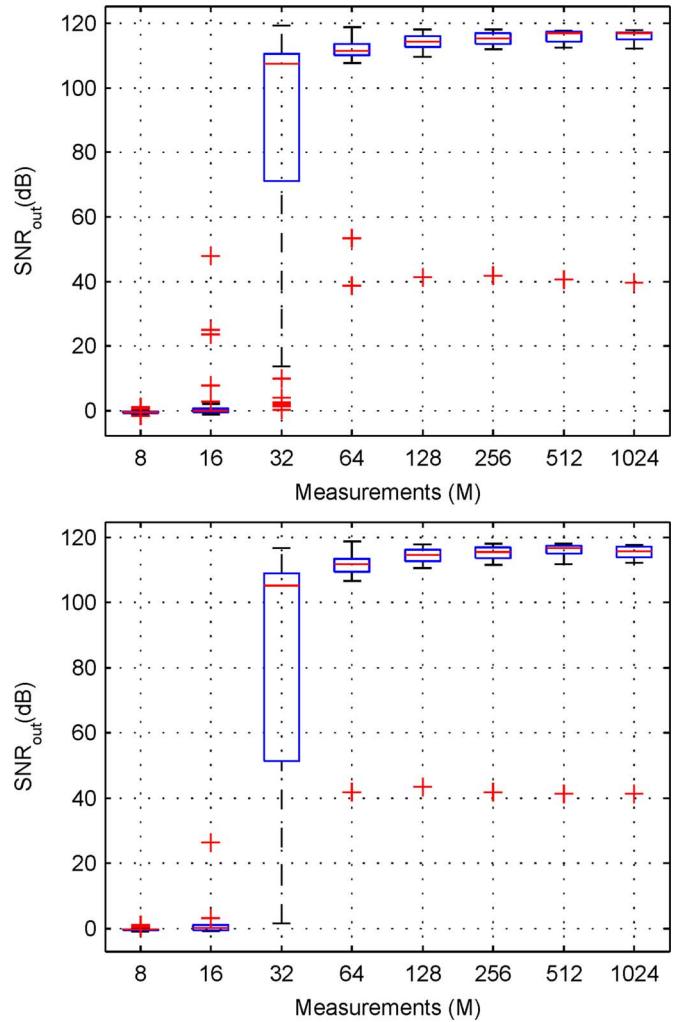


Fig. 8. Box plots for input SNR of 40 dB for RMPI (top) and SRMPI (bottom).

quality for different number of measurements. It is clear that the quality reaches a plateau after roughly $M = 64$ measurements. For an input SNR of -5 dB both architectures can reach a signal quality of 23 dB and even 120 dB for an input SNR of 40 dB with $M = 64$. These figures show the average SNR, but there is a large variance across individual packets. Alternatively, Fig. 8(a) and (b) shows box plots for RMPI and SRMPI, respectively. On each box, the central mark is the median, the edges of the box are the 25th and 75th percentile, and the whiskers extend to the most extreme data points not considered outliers. The red cross markers on the plots correspond to the outliers. Points are drawn as outliers if they are larger than $u + w(u - l)$ or smaller than $l - w(u - l)$, where l and u are the 25th and 75th percentiles (the edges of the boxes), respectively, and w is the maximum whisker length (set to 1.5). The default of 1.5 corresponds to approximately $\pm 2.7\sigma$ and 99.3 coverage if the data are normally distributed. These plots show that with $M = 64$ measurements we can reach values of output SNR higher than 110 dB for both architectures with high probability. These plots also confirm that even outliers never drop lower than the input signal quality (fixed at 40 dB).

We have also defined a success rate, which corresponds to the number of tones that are exactly recovered. Fig. 9 shows

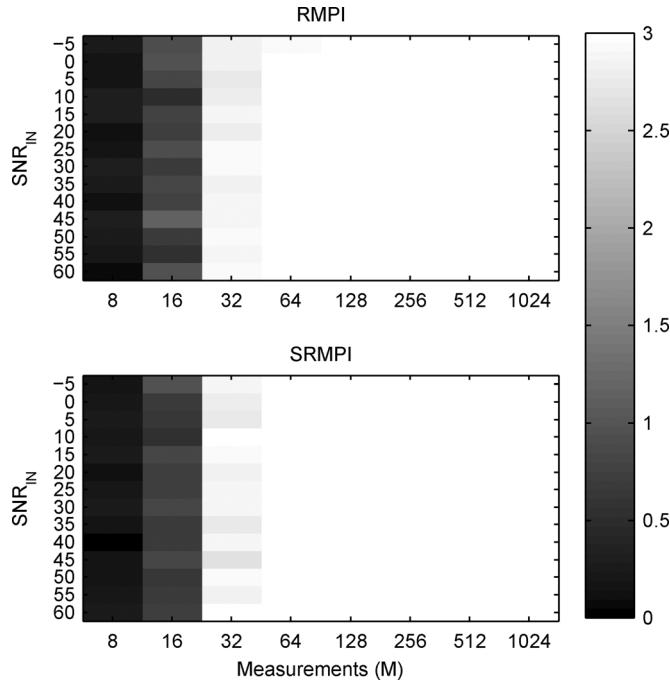


Fig. 9. Averaged successful detected tones of both 8-channel architecture RMPI and SRMPI.

the results for a wide range of input signal quality and number of measurements. Each color indicates the number of frequency tones that were exactly recovered. These results show that full recovery is obtained at all input signal qualities with $M = 64$ measurements and an 8-channel architecture.

VII. POWER ANALYSIS

The previous section described in detail two analog CS architectures, RMP and SRMPI, and carefully characterized both in terms of implementation choices and signal reconstruction metrics (i.e., SNR and success rate). The present section focuses on their power profiles and consumption on the 8-channel platform we developed.

To characterize the power consumption of both architectures, we are interested in the power consumed by two main elements of our implementation: ADCs and mixers. However, the overall power consumption of the physical board includes the consumption by other off-the-shelf components. Nonetheless, we neglect these parts since they can be optimized in a final integrated design (like the path resistors, interfaces with DAC card on PC and so on). This assumption is a simplified way of estimating the overall power consumption of the designed board, but it is valid since the other elements used in both architectures are common and therefore do not change our conclusions. To measure the power, we placed a $330\text{-}\Omega$ resistor in the power path of both elements individually and the voltage drop was measured using an oscilloscope. The corresponding current and power consumption were then calculated in a straightforward way.

Fig. 10(a) shows the ADC consumption during sampling. To find out the exact timing period of the ADC, both conversion start and busy signals were monitored. Both the start and end of the sampling operation are marked on the plot, and the

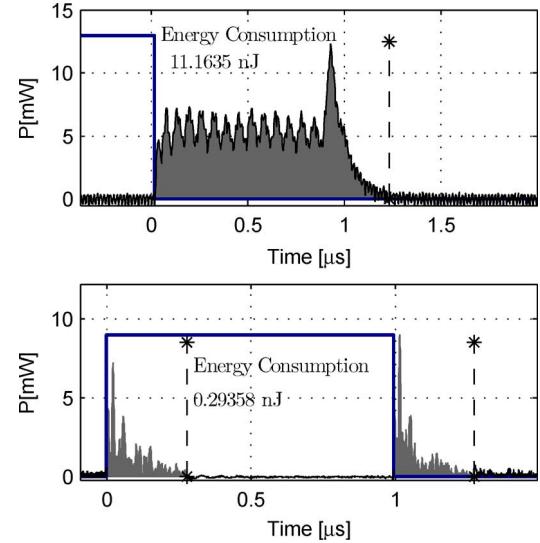


Fig. 10. Power consumption of the ADC (top) during the acquisition time and the mixer (bottom); measured directly on the physical board. The start and end of the periods are indicated on the plots with vertical lines.

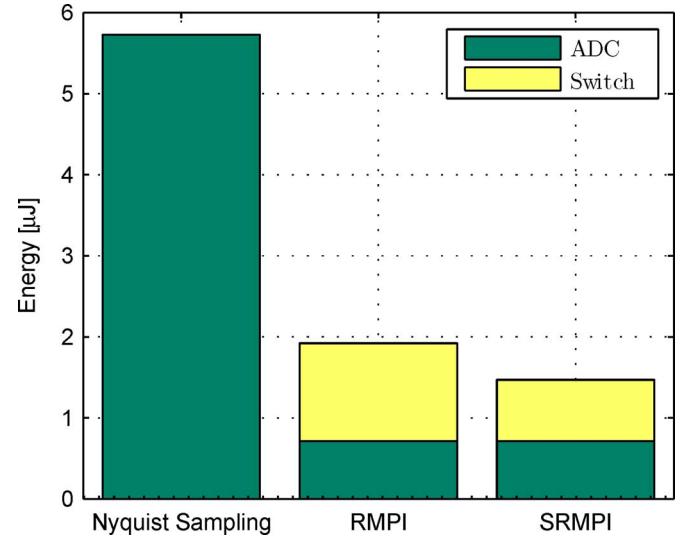


Fig. 11. Energy consumption break down for traditional Nyquist sampling, RMPI and SRMPI architectures for 8 channels.

area below the power curves corresponds to the power consumption of the ADC during one sampling interval. Similarly, Fig. 10(b) shows the power consumption of the mixer during two switching periods. The timing period for a mixer is extracted from the manufacturer's datasheets. Although the power consumption of the mixers are relatively very low compared to the ADC, since they work at a higher rate in both architectures, they consume a significant part of the overall system power consumption.

Therefore, to understand the reduction of power consumption provided by SRMPI over RMPI, we further characterized the energy consumption breakdown of the ADC and mixers in our experimental setup and platform for a period of 1 s with $M = 64$. In particular, Fig. 11 depicts the share of the ADC and mixers in total system power consumption breakdown for the normal Nyquist rate sampling ($F_S = 512$ Hz) and two CS-based

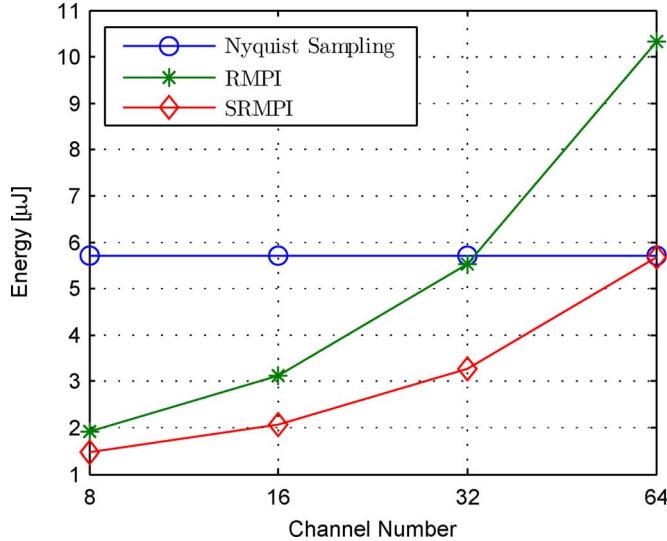


Fig. 12. Energy consumption for traditional Nyquist sampling, RMPI and SRMPI architectures for 8, 16, 32, and 64 channels.

sampling architectures. As shown in Fig. 11, the power consumption of the mixers is significant compared to the ADC. Fig. 11 shows that both architectures can successfully reduce the overall sampling cost. The results show that RMPI can reduce the power consumption by 63% compared to normal Nyquist rate sampling while SRMPI reduces the power consumption by more than 75%. Clearly SRMPI outperforms RMPI by at least 25% in total power consumption. Moreover these results highlight the suitability of analog CS, even at moderate bandwidth, as a low power alternative to traditional sampling.

In addition, the SRMPI architecture shows even higher power saving gains if it is used at high sampling frequency or when a large number of channels is required. In fact in many real world problems like in multi-lead biosignal analysis (e.g., EEG or 12-lead ECG) or in RF applications our preliminary investigations (Fig. 12) show that, for a 32-bit architecture, RMPI only reaches 3% power reduction compared to traditional Nyquist sampling whereas our SRMPI reaches up to 43% energy saving.

VIII. DISCUSSION AND CONCLUSION

The classical analog-to-digital conversion paradigm based on Shannon/Nyquist sampling has been challenged lately, and CS has been proposed as a new emerging signal acquisition/compression paradigm that offers a striking alternative to traditional signal acquisition. For signals having a sparse representation in some dictionary of waveforms, CS can be used to recover those signals from a small number of linear projections, thus enabling efficient sensing, sampling and compression. However, only very few system-wide implementations of CS exist that could serve as reference design and furthermore there is no sensible exploration of the features of each proposed design in the literature to select the best CS-based architecture for different sensing requirements of various types of applications (biomedical sensing, radar systems and communication signal processing, etc.).

In this paper, we have first presented a complete system-level analysis and comparison between state-of-the-art CS-based

signal acquisition systems and then introduced the SRMPI, which is a new design architecture for CS-based analog-to-information (A2I) readout systems. Our experimental results confirm that SRMPI exhibits significantly better overall energy efficiency for a given output quality than any state-of-the-art CS-based signal acquisition systems. Moreover, SRMPI reduces the complexity of the design and final calibration of the system with respect to other CS-based A2I systems. To the best of our knowledge, the results in this work have provided the first system-level power characterization of different analog CS architectures on a real scenario. These results also validate the suitability of using CS-based A2I systems over traditional sampling techniques for highly sparse signals with low-power operation requirements.

ACKNOWLEDGMENT

The authors would like to thank Prof. M. Kayal and S. Sunier for their valuable suggestions and help in the design and implementation of the analog CS platform used in our experimental setup.

REFERENCES

- [1] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [2] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [3] E. Candès and T. Tao, "Near optimal signal recovery from random projections: Universal encoding strategies," *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.
- [4] H. Mamaghani, N. Khaled, D. Atienza, and P. Vandergheynst, "Compressed sensing for real-time energy-efficient ECG compression on wireless body sensor nodes," *IEEE Trans. Biomed. Eng.*, vol. 58, no. 9, pp. 2456–2466, Sep. 2011.
- [5] J. Tropp, J. Laska, M. Duarte, J. Romberg, and R. Baraniuk, "Beyond Nyquist: Efficient sampling of sparse bandlimited signals," *IEEE Trans. Inf. Theory*, vol. 56, no. 1, pp. 520–544, Jan. 2010.
- [6] M. Mishali, Y. Eldar, O. Dounaevsky, and E. Shoshan, "Xampling: Analog to digital at sub-Nyquist rates," *Circuits, Devices Systems, IET*, vol. 5, no. 1, pp. 8–20, Jan. 2011.
- [7] J. Laska, S. Kirolos, M. Duarte, T. Ragheb, R. Baraniuk, and Y. Massoud, "Theory and implementation of an analog-to-information converter using random demodulation," in *IEEE Int. Symp. Circuits Syst.*, May 2007, pp. 1959–1962.
- [8] S. Pfetsch, T. Ragheb, J. Laska, H. Nejati, A. Gilbert, M. Strauss, R. Baraniuk, and Y. Massoud, "On the feasibility of hardware implementation of sub-Nyquist random-sampling based analog-to-information conversion," in *IEEE Int. Symp. Circuits Syst.*, May 2008, pp. 1480–1483.
- [9] S. Kirolos, J. Laska, M. Wakin, M. Duarte, D. Baron, T. Ragheb, Y. Massoud, and R. Baraniuk, "Analog-to-information conversion via random demodulation," in *IEEE Dallas/CAS Workshop Design, Appl., Integrat. Software*, Oct. 2006, pp. 71–74.
- [10] J. Yoo, S. Becker, and A. Emami-Neyestanak, "Design and implementation of a fully integrated compressed-sensing signal acquisition system," in *IEEE Int. Conf. Acoust., Speech Signal Process.*, 2012, pp. 5325–5328.
- [11] S. R. Becker, "Practical compressed sensing: Modern data acquisition and signal processing" Ph.D. dissertation, California Inst. Technol., Pasadena, CA, 2011 [Online]. Available: <http://thesis.library.caltech.edu/6492/>
- [12] F. Naini, R. Gribonval, L. Jacques, and P. Vandergheynst, "Compressive sampling of pulse trains: Spread the spectrum!," in *IEEE Int. Conf. Acoust., Speech Signal Process.*, Apr. 2009, pp. 2877–2880.
- [13] E. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, pp. 1207–1223, 2006.
- [14] D. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2845–2862, Nov. 2001.

- [15] J. A. Tropp, "Just relax: Convex programming methods for identifying sparse signals in noise," *IEEE Trans. Inf. Theory*, vol. 52, no. 3, pp. 1030–1051, Mar. 2006.
- [16] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1999.
- [17] E. van den Berg and M. P. Friedlander, "Probing the Pareto Frontier for basis pursuit solutions," *SIAM J. Sci. Comput.*, vol. 31, no. 2, pp. 890–912, 2008.
- [18] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE J. Sel. Topics Signal Process.*, vol. 1, no. 4, pp. 586–597, 2007.
- [19] I. Daubechies, M. Defrise, and C. D. Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.*, vol. 57, pp. 1413–1457, 2004.
- [20] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. Asilomar Conf. Signals, Syst. Comput.*, Nov. 1993, pp. 40–44.
- [21] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.
- [22] E. V. Berg and M. P. Friedlander, SPGL1: A solver for large-scale sparse reconstruction Jun. 2007 [Online]. Available: <http://www.cs.ubc.ca/labs/scl/spgl1>
- [23] F. Chen, A. Chandrakasan, and V. Stojanovic, "Design and analysis of a hardware-efficient compressed sensing architecture for data compression in wireless sensors," *IEEE J. Solid-State Circ.*, vol. 47, no. 3, pp. 744–756, Mar. 2012.
- [24] G. Puy, P. Vandergheynst, R. Gribonval, and Y. Wiaux, "Universal and efficient compressed sensing by spread spectrum and application to realistic Fourier imaging techniques," *EURASIP J. Adv. Signal Process.*, vol. 2012, no. 6, 2012.



Hossein Mamaghani (S'11) received the M.Sc. degree in electrical engineering from Sharif University of Technology (SUT), Tehran, Iran, in 2009. Currently, he is pursuing the Ph.D. degree in electrical engineering at École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland. His Ph.D. topic mainly focuses on compressed sensing and universal energy efficient compression scheme for biosignals and he is working jointly with Signal Processing Laboratory (LTS2) and Embedded Systems Laboratory (ESL) at EPFL.

His main research interests include statistical signal processing, compressed sensing, and biomedical signal processing.



Nadia Khaled (M'05) received the M.Sc. degree in electrical engineering from ENSEEIHT, Toulouse, France, in 2000, and the Ph.D. degree in electrical engineering from the Katholieke Universiteit Leuven, Leuven, Belgium, in 2005.

From 2000 to 2005, she was with the wireless research group of the Interuniversity Microelectronics Center (IMEC), Leuven. She then joined ETH Zurich, Switzerland, as a postdoctoral researcher for one year. From 2006 to 2011, she was with EPFL Lausanne, Switzerland, as a scientist. She also held

the position of Visiting Assistant Professor at the University of Carlos III Madrid (UC3M) from 2006 to 2009. As of April 2012, she is a R&D Specialist and Project Manager at the Nestlé Research Center, Lausanne, Switzerland.



David Atienza (M'06) received the M.Sc. and Ph.D. degrees in computer science and engineering from Complutense University of Madrid (UCM), Madrid, Spain, and Inter-University Micro-Electronics Center (IMEC), Heverlee, Belgium, in 2001 and 2005, respectively.

Currently, he is Professor and Director of the Embedded Systems Laboratory (ESL) at EPFL, Switzerland, and Adjunct Professor at the Computer Architecture and Automation Department, the University of Madrid. His research interests focus on system-level design methodologies for low-power embedded systems and high performance systems-on-chip (SoC), including new thermal-aware design for 2-D and 3-D multi-processor SoCs, design methods and architectures for wireless body sensor networks, dynamic memory management and memory hierarchy optimizations, as well as novel architectures for logic and network-on-chip (NoC) interconnects. In these fields, he is co-author of more than 180 publications in peer-reviewed international journals and conferences, several book chapters, and four U.S. patents. He is Associate Editor of Elsevier *Integration*.

Dr. Atienza has received the ACM Outstanding New Faculty Award from the Special Interest Group on Design Automation (SIGDA) in 2012, a Best Paper Award at the VLSI-SoC 2009 conference, and four Best Paper Award Nominations at the HPCS 2012, WEHA-HPCS 2010, ICCAD 2006 and DAC 2004 conferences. He is Associate Editor of IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF CIRCUITS AND SYSTEMS. He is a member of the Executive Committee of the IEEE Council on EDA since 2008, and GOLD member of the Board of Governors of the IEEE Circuits and Systems Society (CASS) since 2010.



Pierre Vandergheynst (SM'08) received the M.S. degree in physics and the Ph.D. degree in mathematical physics from the Université Catholique de Louvain, Louvain-la-Neuve, Belgium, in 1995 and 1998, respectively.

From 1998 to 2001, he was a Postdoctoral Researcher with the Signal Processing Laboratory, Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland. He was Assistant Professor at EPFL (2002–2007), where he is now an Associate Professor. His research focuses on harmonic analysis, sparse approximations and mathematical image processing with applications to higher dimensional, complex data processing. He was co-Editor-in-Chief of *Signal Processing* (2002–2006). He is the author or co-author of more than 50 journal papers, one monograph, and several book chapters. He is strongly involved in technology transfer: he co-founded two start ups and holds numerous patents.

Dr. Vandergheynst is Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING (2007–present). He has been on the Technical Committee of various conferences and was Co-General Chairman of the EUSIPCO 2008 conference. He's a laureate of the Apple ARTS award and of the 2010–2011 De Boelpaepe prize from the Royal Academy of Sciences of Belgium.