Ultra-Low-Power ECG Front-End Design based on Compressed Sensing

Hossein Mamaghanian
School of Engineering
Ecole Polytechnique Fédérale de Lausanne (EPFL)
Email: hossein.mamaghanian@epfl.ch

Pierre Vandergheynst
School of Engineering
Ecole Polytechnique Fédérale de Lausanne (EPFL)
Email: pierre.vandergheynst@epfl.ch

***Abstract*—Ultra-low-power design has been a challenging area for design of the sensor front-ends especially in the area of Wireless Body Sensor Nodes (WBSN), where a limited amount of power budget and hardware resources are available. Since introduction of Compressed Sensing, there has been a challenge to design CS-based low-power readout devices for different applications and among all for biomedical signals. Till now, different proposed realizations of the digital CS prove the suitability of using CS as an efficient low-power compression technique for compressible biomedical signals. However, these works mainly take advantages of only one aspect of the benefits of the CS. In this type of works, CS is usually used as a very low cost and easy to implement compression technique. This means that we should acquire the signal with traditional limitations on the bandwidth (BW) and later compresses it. However, the main power of the CS, which lies on the efficient data acquisition, remains untouched. Building on our previous work [1], where the suitability of the CS is proven for the compression of the ECG signals, and our investigation on ultra-low-power CS-based A2I devices [2] , here in this paper we propose a fully redesigned complete CS-based "Analog-to-information" (A/I) front-end for ECG signals.**

Our results show that proposed hybrid design easily outperforms the traditional implementation of CS with more than 11 times fold reduction in power consumption compared to standard implementation of CS. Moreover our design shows a very promising performance specially in high compression ratio.

I. Introduction

Since early days of the introduction of Compressed Sensing (CS) [3], [4], it has been a focus of the new generation of ultra-low power Analog to information (A2I) readout device design. However, realization of energy-efficient CS-based analog A2I readout devices still face technical problems and the main focus is on using the CS as a low power digital compression technique. One of the main bottlenecks which is preventing the true realization of CS-based readout devices lies on relatively large number of measurements $m = s\log(\frac{n}{s})$, required for decoding (for a given vector of $\mathbf{x} \in \mathbb{R}^n$ with sparsity level s). This bound is even worse in case of the compressible signals like ECG. In a real standard analog realization like Random modulator pre-Integrator (RMPI) [5], [6], this directly translated to the number of required parallel channels. The power consumption of such a realization of CS is far away from being close even to digital implementation of CS or even the traditional Nyquist sampling [7]. The main reason lies in the fact that such a design requires the same number of amplifiers for each channel and it is the most dominant part in power consumption with very large margin. Moreover, having very large number of channels face implementation difficulties for miniaturized WBSNs with existing technologies. Then to have a true analog CS-based readout device, the main focus should be on the minimizing the number of required measurements and as a result less number of channels. This figure could partially be enhanced by more advanced and signal specific recovery algorithms to exploit additional information for underlying signal of interest. This includes all model-based and similar structural sparse recovery techniques [8], [9]. However, the best performance one can ideally achieve is to minimize the required number of measurements to $m = \mathcal{O}(s)$ for exact s-sparse signals.

Here in this paper, we propose a fully redesigned complete CS-based "Analog-to-information" (A/I) front-end for ECG signals. The new design uses an additional ultra-low-power low-resolution channel in parallel to further improve and ease the CS recovery. Such a parallel channel enables us to dramatically reduce the number of required measurements for recovery. This will directly translated to reduction in number of the required parallel channels in a normal CS implementation like RMPI. In this design, CS channel is used as a super resolution path to improve the ultra-low-power, low-resolution samples captured from the parallel path.

II. CS-based ECG Front-End

Figure 1 shows the schematic of the proposed ultra-low power front-end for ECG signals. The design includes two parallel paths. The first one is a regular analog CS path that senses the original ECG signal. This is the original CS path and is working at a rate much lower than traditional Nyquist limit. Then, the parallel path is used to assist the CS recovery of our proposed design. This path includes a low-power, low-resolution ADC that in parallel takes samples at Nyquist rate. This channel could give a rough bound of the signal, which later could be used to enhance the performance of the signal reconstruction. As it is shown in the schematic depicted in Figure 1 collected data from both paths are transmitted at a fixed time window. At the receiver side, this additional information from both paths is used to reconstruct the original signal. The power consumption of parallel low-resolution path in this design is extremely low compared to CS path. We will further discuss the power consumption breakdown on

978-3-9815370-4-8/DATE15/©2015 EDAA

the design but, the overall power consumption from this path should be negligible compared to CS path.

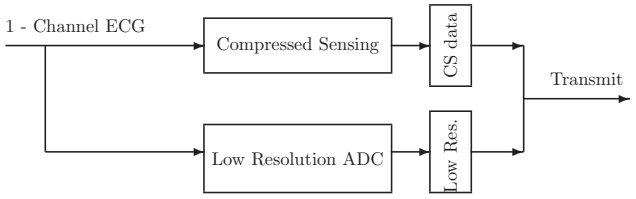

Fig. 1. Schematic diagram of an ultra-low power front-end for ECG signals.

In a fixed size processing window, the output of the original CS channel could be represented as: $\mathbf{y} = \mathbf{\Phi x}$ which are simply linear analog measurement vector from the original ECG signal $\mathbf{x}$.

The output of low-resolution channel is labelled as $\dot{\mathbf{x}}$, which referes to low-resolution quantized version of the original signal $\mathbf{x}$. The output of this channel for an example processing window is shown in Figure 2(a).

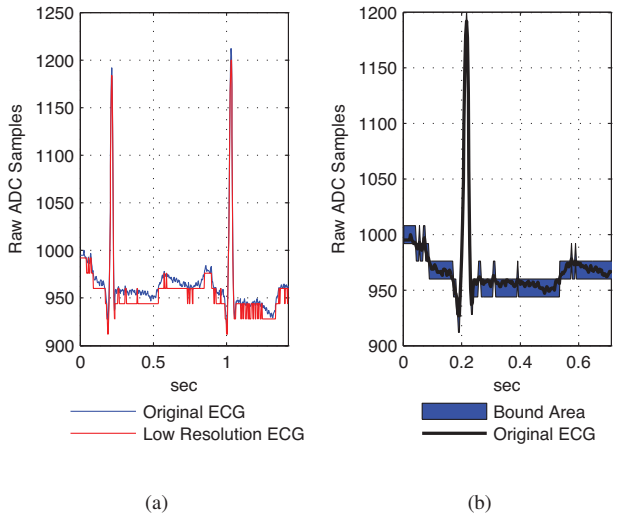

(a) (b)

Fig. 2. a) Output of an example fixed sized window for low-resolution path (7-bit), b) Blue area shows the bounds for the reconstruction results.

This information could be applied in our reconstruction algorithm (Equation 1) as a strong bound on the solution. The modified convex optimization problem now can be written as:

$$\min_{\tilde{\boldsymbol{\alpha}} \in \mathbb{R}^N} ||\tilde{\boldsymbol{\alpha}}||_1 \qquad \text{subject to} \qquad \begin{cases} ||\mathbf{\Phi \Psi} \tilde{\boldsymbol{\alpha}} - \mathbf{y}||_2 \leq \sigma, \\ \dot{\mathbf{x}} \leq \mathbf{\Psi} \tilde{\boldsymbol{\alpha}} \leq \dot{\mathbf{x}} + d \end{cases} \quad (1)$$

where d is the resolution depth step. The new constraint on the solution is really strong since it defines an upper and lower bound for each sample in $\tilde{\boldsymbol{\alpha}}$. This problem is convex and here we use semidefinite-quadratic-linear programming toolbox SDPT3 [10], [11] which is a MATLAB toolbox for solving conic optimization problems.

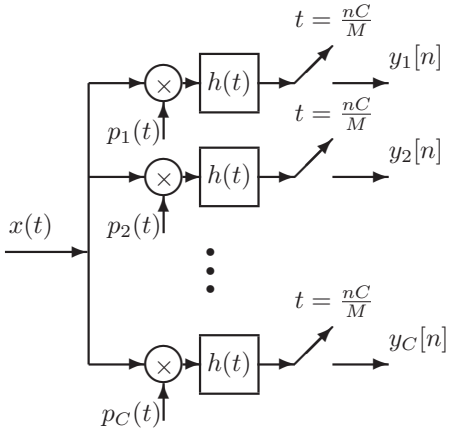

Fig. 3. Block diagram of random modulator pre-Integrator.

III. Methods

A. CS Channel implementation

Figure 3 shows the analog representation of CS, called Random Modulator Pre-Integrator (RMPI). This architecture is a variant of the Random Demodulator (RD) [12] architecture, and is composed of parallel channels of RD. Each block on Figure 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. If number of channels C is equal to number of measurements m, then such an architecture could be used as an exact analog implementation of CS compression. Here, we don't describe the details of the RMPI architecture and we only use the analytical models to calculate the power consumption of such architecture. Interested reader could refer to reference [5], [6], [13], [14].

The proposed design introduces a trade-off between the number of measurements captured from CS high-resolution channel and low-resolution channels. Since the data from low-resolution channel are quantized at a lower resolution, it is clear that they should have very repeated samples. Figure 2(a) shows an example window of the data from this channel. Blue and red curves indicate the original ECG signal at a fixed size window and the corresponding low-resolution samples taken at 7-bit, respectively. This figure shows that as expected, these data are repeated and highly redundant. This means that these data could be efficiently coded and compressed by lossless entropy coding techniques. Apart from the overall compression ratio, implementation complexity and the dictionary size of the encoder, play a very important roles in defining the trade-off between the resolution of parallel channel and required measurements from CS channel.

B. The Entropy Coding of Low-Resolution Channel

Quantized data from low-resolution channel $\dot{\mathbf{x}}$ have a very repetitive nature, since they are sampled at much lower res-

olution. This means that there is a huge redundancy between neighboring samples. Thus, instead of coding directly the quantized samples, we find the difference between consecutive samples in vector $\dot{\mathbf{x}}$ and compress them by lossless entropy coding.

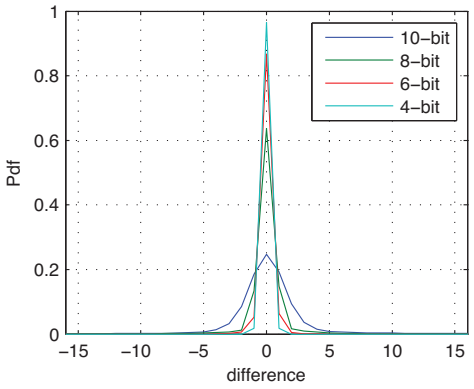

Fig. 4. Pdf of difference between quantized samples from low-resolution channel

Figure 4 shows the probability density function of the difference between quantized samples for different resolution bit depth (i.e., 10, 8, 6 and 4). As expected, the distributions of the difference signal at the output of the low-resolution channel are far from uniform. Consequently, Huffman encoding can be used for further compression. For different ranges of resolution depth, the length of codebook are different. Since this offline-generated codebook should be stored on the node, then the size of required storage is also very important. Figure 5 shows the required memory space in byte for storing the offline-generated codebook for different numbers of bit resolutions.

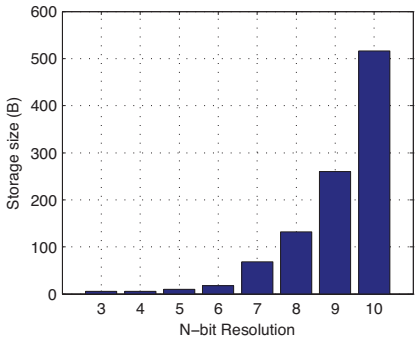

Fig. 5. Storage required in Byte for storing the offline-generated codebook for different values of quantization depth.

Compression ratios for the low-resolution channel are defined as before, and the results are shown in Figure 6. It shows the averaged expected compression ratio for different quantization bit depth. As the resolution increases the compression decreases, since the distribution of the samples is getting closer to the uniform distribution. Finally, because the compressed

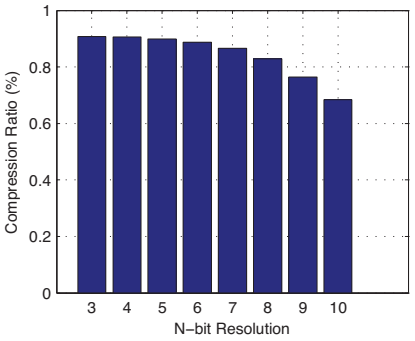

Fig. 6. Average compression ratio for low-resolution path for different bit-resolutions

TABLE I
AVERAGE OVERHEAD FROM LOW-RESOLUTION CHANNEL FOR DIFFERENT BIT RESOLUTIONS

Bit Resolution	10	9	8	7	6	5	4	3
Overhead (D_i %)	26.3	17.6	11.4	7.8	5.6	4.2	3.1	2.3

data from this channel are sent along with compressed sensing data, we calculate the amount of overhead for each resolution. Let us assume that $CR_i, i \in \{3, 4, \ldots, 10\}$ corresponds to the compression ratio obtained by applying Huffman coding and original ECG samples are considered to be at 12-bit resolution. Then the total overhead from low-resolution channel could be described as:

$$D_i = CR_i * {}^{i}/_{12}, \tag{2}$$

where D_i is corresponding percentage of overhead. Table I represents the values of this overhead for corresponding bit resolutions. This number should be added on the final compression ratio obtained from CS channel.

IV. Experimental Results

To validate the performance of the considered compression schemes, we use the MIT-BIH Arrhythmia Database [15], [16] that is the most commonly used database for the comparative study of ECG compression algorithms. This database contains 48 half-hour excerpts of two-channel ambulatory ECG recordings, obtained from 47 subjects studied by the BIH Arrhythmia Laboratory. The recordings were digitized at 360 samples per second per channel with 11-bit resolution over a 10 mV range.

Moreover, to quantify the compression performance while assessing the diagnostic quality of the compressed ECG records, we employ the two most widely used performance metrics, namely *the compression ratio* (CR) and *percentage root-mean-square difference* (PRD). CR is defined as:

$$CR = \frac{b_{orig} - b_{comp}}{b_{orig}} \times 100, \tag{3}$$

where b_{orig} and b_{comp} represent the number of bits required for the original and compressed signals, respectively. The PRD, and associated signal-to-noise ratio (SNR), quantifies

the percent error between the original signal vector $\mathbf{x}$ and the reconstructed $\tilde{\mathbf{x}}$:

$$PRD = \frac{||\mathbf{x} - \tilde{\mathbf{x}}||_2}{||\mathbf{x}||_2} \times 100; \quad SNR = -20\log_{10}(0.01PRD).$$

The trade-off of resolution bit depth for low-resolution channel is considered as 7-bit. This indicates the 7.86% of overhead from low-resolution channel for each processing window. The on-node storage required for storing the offline-generated Huffman codebook is also equal to 68 Bytes.

V. Performance Quality

Figure 7 shows the averaged SNR and PRD values for different compression ratios for both Normal single-lead CS and Hybrid CS, respectively. The figures only show the compression ratios higher than 50%, which is the main region of interest. Results show that, Hybrid CS easily outperforms the normal reconstruction. This superior performance is more outstanding for high compression ratios, where normal CS fails to converge or has very poor reconstruction quality. Moreover, the results also show that "good" reconstruction quality is achieved at 81% of compression compared to 53% for normal CS. Considering the overhead from parallel channel, net compression ratio for Hybrid CS is about 73.14%. Similarly, Figure 8 also shows the box plot for all the records. On each box, the central mark is the median, the edges of the box are the 25^{th} and 75^{th} percentiles, and the whiskers extend to the most extreme data points not considered outliers.

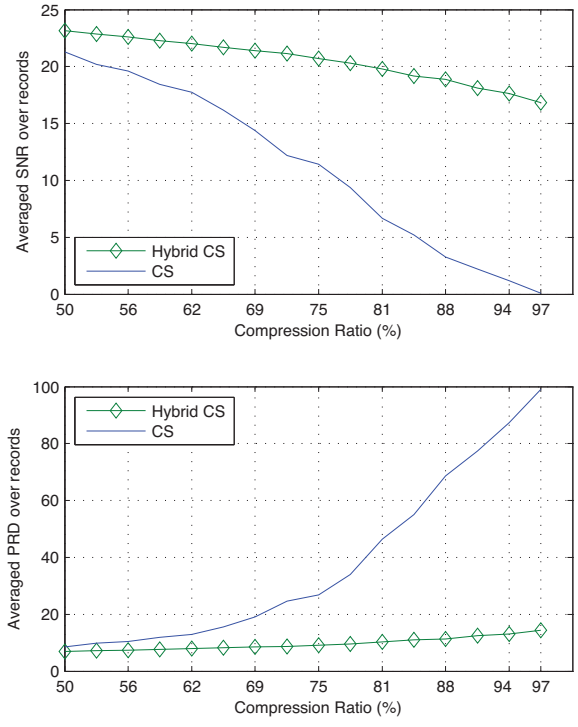

Fig. 7. Averaged top) SNR and bottom) PRD for different compression ratios for both Hybrid and normal CS

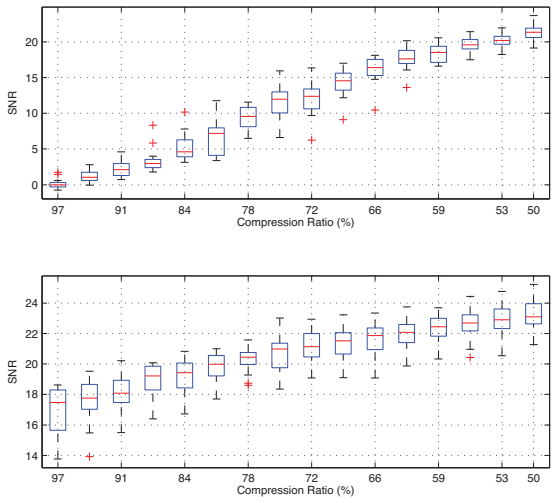

Fig. 8. Box plots for all the records of database for top) normal and bottom) Hybrid CS reconstruction

It is also noteworthy that even with a very high net compression ratio like 85.14% (97% from CS channel minus 7.86% overhead from low-resolution channel), it is still possible to have a reasonable signal to noise ratio higher than 17dB.

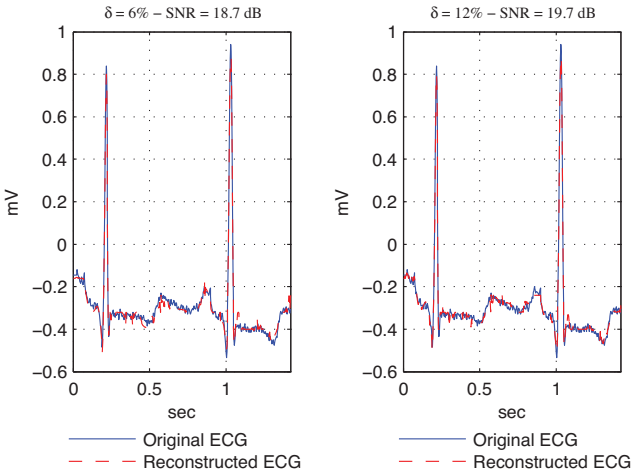

(a) (b)

Fig. 9. Example original ECG window and reconstructed signal over different number of $\delta = \frac{m}{n}$ (ranges: 6%, 12%, 25%)

Additionally, Figure 9 shows a sample window of the reconstructed and original signal for different percentage of $\delta = \frac{m}{n}$. In each plot the corresponding values for δ and SNR are presented in the title. Results show that even with a very high compression ratio of 6%, the output SNR is 18.7 dB, which is very promising.

VI. Power Consumption Model

Within this subsection, we use analytical models for power consumption of our design to evaluate the overall power effi-

ciency of our design. Here we use power model for each block diagram of that design to estimate total power consumption. Power consumption of such a design could be categorized into 3 parts: 1) Analog to digital convertors (ADCs), 2) Accumulator, 3) Amplifiers. We have used the power models for these blocks from [7], where they have quantified analytically power consumption of each block for 90nm technology. The schematic block diagram of the example circuitry for an RMPI implementation of CS is depicted in Figure 10.

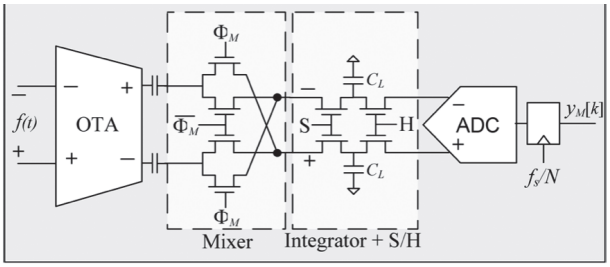

Fig. 10. Schematic block diagram of example circuitry for an RMPI implementation of CS [7]

a) Analog-to-Digital Converter: Power consumptions of the ADCs are directly related to the number of acquired samples, and their resolution bit (B). If the true Nyquist sampling frequency is denoted by f_s, then total number of samples in a second is $\left(\frac{m}{n}\right) \times f_s$, where n is the number of samples in a block and m is the number of required number of channels. Then the resulting power consumptions of array of m ADC are

$$P_{ADC} = \left(\frac{m}{n}\right) \text{FOM} \times 2^B \times f_s \tag{4}$$

where the FOM is figure-of-merit of the ADC that is design specific. The number is about 100 fj/conversion for modern ADCs.

b) Integrator and Sample/Hold: In reference [7], a simplified design is proposed for Operational Transconductance Amplifier (OTA), mixer and integrator, and the equivalent power models are estimated. The corresponding power consumption for the integrator and an ideal Sample/Hold are modeled as:

$$P_{Int} = 2BW_f \times \frac{m \times V_{DD}^2 \times 10\pi \times n \times C_p}{16} \tag{5}$$

where BW_f is the bandwidth of the signal and C_p corresponds to the capacitance for dominant pole of the unloaded OTA $1/2\pi R_0 C_p$.

c) Amplifiers: For the amplifier, the lower bound of the power consumption is typically set by the input referred noise ($v_{ni,rms}$) requirements. For the amplifiers, the figure of merit, which relates the total amplifier current consumption (I_{amp}) to the input referred noise, is the noise efficiency factor (NEF) [17]. The NEF is defined as:

$$\text{NEF} = v_{ni,rms}\sqrt{\frac{2I_{amp}}{\pi.V_T.4kT.BW}} \tag{6}$$

where V_T is thermal voltage (kT/q) and BW is the bandwidth of the amplifier. This number falls in the range between 2 and 3 for many state-of-the-art amplifiers [7]. For the total power consumption of the amplifiers in RMPI architecture we have:

$$P_{amp} = m \times V_{DD}.I_{amp} \tag{7}$$

$$\geq m \times V_{DD}.\frac{\text{NEF}^2}{v_{ni,rms}^2} \times \frac{\pi \times V_T \times 4kT \times BW}{2} \tag{8}$$

and considering the output noise constraint in terms of $v_{ni,rms}^2$, the power model for amplifiers could be simplified as:

$$P_{amp} = 2BW \times 3mn \times 2^{2B_y} \times \frac{G_A^2 \text{NEF}^2}{V_{DD}} \times \frac{\pi(kT)^2}{q} \tag{9}$$

where G_A is the total voltage gain from the input of the amplifier to the input of the ADC. For a detailed descriptions of the power model, we refer the interested readers to reference [7] and reference therin, where the detailed steps are provided.

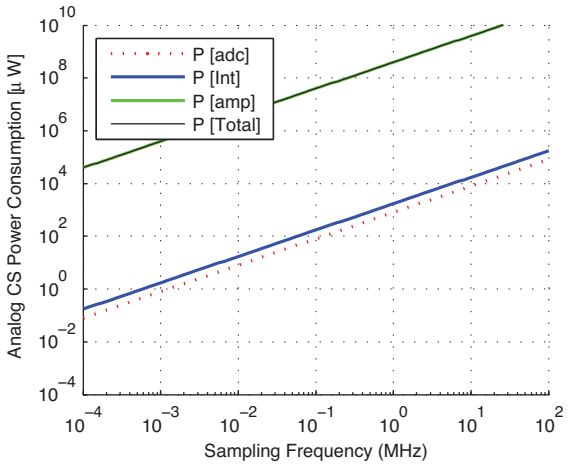

(a) RMPI

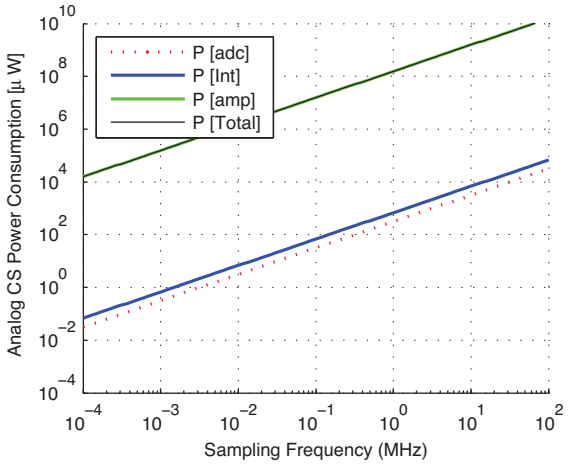

(b) Hybrid CS

Fig. 11. Power consumption breakdown for both RMPI and Hybrid CS architectures.

To compare the power consumption of the proposed design with normal CS implementation (RMPI), for both architectures, we have calculated power consumption for number of required measurements (and then required parallel channels in CS implementation) for fixed $SNR = 20$dB. The required number of measurements equal to 96 and 240 for Hybrid CS and normal CS respectively. The G_A is considered 40dB, which is proper number for a ECG front-end. These numbers may not seem realistic, specially for normal CS implementation. Having this much number of channels (240) may face other technological problem, but for the sake of comparison, it is still valid. Figure 11 shows the results of power consumption breakdown for both architectures. First, it is worth to mention that in such a design with higher number of channels, the dominant part of power consumption-with a large margin- is for amplifier, and the power consumption from the other two parts are just negligible. However, If we want to compare these both architectures, overall power consumption of Hybrid CS is roughly 2.5 times lower than its normal counterpart, which is a huge gain. Since power consumption of the module is directly proportional to the number of measurements.

If we go for much lower number of measurements where still Hybrid CS has acceptable performance quality, this figure gets even better. For only $m = 16$ sample, Hybrid CS reaches the SNR quality of 17dB where for its RMPI counter need $m = 176$ sample. This corresponds to the 11 times reduction in total power consumption.

VII. Conclusion

Even though CS has been very promising to be the solution for the next generation of the ultra-low-power sensing/compression readout devices, but a practical realization still suffers from few important limitations to really validate its low-power consumption claim. Up to now, all analog realizations of the CS have been designed for highly exact sparse signals (few significant base or tone). Some of these limitations lie in the theoretical bound for the number of sufficient measurements for recovery of sparse and compressible signals. This is very critical in case of an analog implementation of the CS, which dictates higher number of parallel channels and as a result many technical problems occurs in implementation. Moreover, such design became highly inefficient considering the overall power consumption. These limitations prevent any practical realization of truly CS-based front-end design especially for signal with medium bandwidth. Here for the first time, a new design is presented for a real analog realization of the CS-based front-end for ECG signals. Our design exploits additional information from an ultra-low-power low-resolution path. This additional information is exploited in recovery algorithm to reduce the number of required measurements. Our results show that this hybrid design easily outperforms the traditional implementation of CS with more than 11X compared to traditional analog implemenattion of CS. Moreover, the performance quality is significantly better in high compression ratios where normal CS even fails to converge.

One of the main potential applications for analog implementation of CS is in High frequency (HF) applications where the sampling frequency is so large. For such a huge sampling rates the equivalent number of bits (ENOB) on a real ADC is very poor. For example, for flash-based ADCs, which are the fastest ADCs, this number is not more than 8-bits for 1GHz sampling frequency. This is the main motivation for designing the new generation of A2I readout devices, like RMPI. Our design has the potential to be used in such a configuration as a super resolution path.

References

[1] H. Mamaghanian, N. Khaled, D. Atienza, and P. Vandergheynst, "Compressed sensing for real-time energy-efficient ecg compression on wireless body sensor nodes," *Biomedical Engineering, IEEE Transactions on*, vol. 58, no. 9, pp. 2456 –2466, sept. 2011.

[2] ——, "Design and exploration of low-power analog to information conversion based on compressed sensing," *Emerging and Selected Topics in Circuits and Systems, IEEE Journal on*, vol. 2, no. 3, pp. 493–501, 2012.

[3] D. Donoho, "Compressed sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289 –1306, april 2006.

[4] E. Candes and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Problems*, vol. 23, no. 3, pp. 969–985, 2007.

[5] J. Tropp, J. Laska, M. Duarte, J. Romberg, and R. Baraniuk, "Beyond nyquist: Efficient sampling of sparse bandlimited signals," *Information Theory, IEEE Transactions on*, vol. 56, no. 1, pp. 520 –544, jan. 2010.

[6] 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 *Circuits and Systems, 2007. ISCAS 2007. IEEE International Symposium on*, may 2007, pp. 1959 –1962.

[7] F. Chen, A. Chandrakasan, and V. Stojanovic, "Design and analysis of a hardware-efficient compressed sensing architecture for data compression in wireless sensors," *Solid-State Circuits, IEEE Journal of*, vol. 47, no. 3, pp. 744 –756, march 2012.

[8] H. Mamaghanian, N. Khaled, D. Atienza, and P. Vandergheynst, "Structured sparsity models for compressively sensed electrocardiogram signals: A comparative study," in *Biomedical Circuits and Systems Conference (BioCAS), 2011 IEEE*, Nov 2011, pp. 125–128.

[9] R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde, "Model-based compressive sensing," *Information Theory, IEEE Transactions on*, vol. 56, no. 4, pp. 1982 –2001, april 2010.

[10] R. H. Tütüncü, K. C. Toh, and M. J. Todd, "Solving semidefinite-quadratic-linear programs using sdpt3," *Mathematical programming*, vol. 95, no. 2, pp. 189–217, 2003.

[11] K.-C. Toh, M. J. Todd, and R. H. Tütüncü, "Sdpt3a matlab software package for semidefinite programming, version 1.3," *Optimization methods and software*, vol. 11, no. 1-4, pp. 545–581, 1999.

[12] 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 *Design, Applications, Integration and Software, 2006 IEEE Dallas/CAS Workshop on*, oct. 2006, pp. 71 –74.

[13] 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 *Circuits and Systems, 2008. ISCAS 2008. IEEE International Symposium on*, may 2008, pp. 1480 –1483.

[14] S. R. Becker, "Practical compressed sensing : modern data acquisition and signal processing," Ph.D. dissertation, California Institute of Technology, 2011. [Online]. Available: http://thesis.library.caltech.edu/6492/

[15] G. Moody and R. Mark, "The impact of the mit-bih arrhythmia database," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 20, no. 3, pp. 45 –50, may-june 2001.

[16] Physionet, "MIT-BIH arrhythmia database." [Online]. Available: http://www.physionet.org/physiobank/database/mitdb/

[17] M. S. Steyaert and W. M. Sansen, "A micropower low-noise monolithic instrumentation amplifier for medical purposes," *Solid-State Circuits, IEEE Journal of*, vol. 22, no. 6, pp. 1163–1168, 1987.