A Compressed-Sensing Sensor-on-Chip Incorporating Statistics Collection to Improve Reconstruction Performance

[1]Vahid Behravan, [2]Shuo Li, [1]Neil E. Glover, [1]Chia-Hung Chen, [3]Mohammed Shoaib, [1]Gabor C. Temes and [1,2]Patrick Y. Chiang

[1]Oregon State University, Corvallis, OR, USA

[2]Fudan University, Shanghai, China [3]Microsoft Research, Redmond, Washington, USA

***Abstract* — Reconstructing signals accurately is a critical aspect of compressed sensing. We propose a compressed-sensing sensor-on-chip that compresses and also extracts key statistics of the input signal at sampling time. These statistics can be used at the receiver to significantly improve the accuracy of reconstruction. When compared against a conventional compressed-sensing system, our experimental measured results demonstrate an improvement of as much as 9-18 dB in the signal-to-error (SER) of the reconstructed signal, depending on input data type and compression factor.**

***Index Terms* — compressed sensing, sensor-on-chip, reconstruction error, statistics collection.**

I. Introduction

Power consumption is one among the most critical constraints for sensor networks and future Internet-of-Things devices. The most power-consuming block in a sensor system-on-chip is the radio (transmitter), typically responsible for more than 60-80% of the total sensor's energy usage [1]. Thus, one direct way to reduce the radio's energy consumption is to minimize the amount of transmitted data, such as by using data compression. Compressed sensing is an attractive way of compressing data since it requires no prior knowledge about the input signal, as long as the signal exhibits a property called sparsity. It is broadly applicable because many types of sensor data exhibit sparsity, including biomedical signals [2].

Although compressed sensing is a signal-agnostic method of compression, incorporating additional statistics about the input signal (along with its sparsity) can provide more information that can aid reconstruction algorithms at the receiver. Hence, such statistics collection (SC) can improve the accuracy of signal reconstruction. However, to realize systems based on this principle, sensor nodes will need to transmit additional data (statistics) to the receiver, which may increase transmission-energy costs.

In the literature, various systems have previously implemented the compressed-sensing framework [3],[4]. However, all of them employ blind reconstruction with no SC. They do not incorporate any extra information about the input signal to aid reconstruction, resulting in less-than-optimal SERs at the receiver.

In this paper, we implement for the first time a compressed-sensing framework whose reconstruction is aided by statistics collection. We demonstrate how to fuse sensor data and statistics information together to improve the signal-to-error ratio (SER) of reconstruction by as much as 18 dB. Furthermore, using various low-power design techniques, we show that for an ECG data input, SC can be performed with just 0.4μW/12.6μW (dynamic/static power) at a sampling rate of 96 kHz (an overhead of <1% of the communication energy assuming a 1mW radio).

The paper is organized as follows. In Section II, we provide background on compressed sensing, reconstruction, and the design of the proposed SC approach. We then describe the micro-architectural details of the various sub-blocks of the sensor-on-chip (SoC) in Section III. We present the digital circuit implementation of the SoC in Section IV, measurement results in Section V, and conclusions in Section VI.

II. System Overview

A. Background on Compressed Sensing

In compressed sensing, compression of digitized data is performed using matrix multiplications. In particular, an uncompressed input vector $\boldsymbol{f}$ of size N is multiplied by a measurement matrix $\boldsymbol{\Phi}$ of size M×N, producing a measurement vector $\boldsymbol{y}$ of size M. Since $\boldsymbol{\Phi}$ is typically a matrix of random numbers, $\boldsymbol{y}$ is a vector of random linear projections of $\boldsymbol{f}$ onto $\boldsymbol{\Phi}$. To reconstruct the original signal $\boldsymbol{f}$ from the measured signal $\boldsymbol{y}$, we need to solve a set of equations, where the number of equations is much less than the number of unknown variables. In general, there is no unique solution for these types of equations. However, the theory of compressed sensing [5] shows that if the signal $\boldsymbol{f}$ is sparse in any basis $\boldsymbol{\Psi}$ (*e.g.* time, Fourier, wavelet), then there exist known techniques to reconstruct the original signal with minimum error using convex optimization [5]. In other words, if there is a transformation matrix $\boldsymbol{\Psi}$ such that $\boldsymbol{f} = \boldsymbol{\Psi}\mathbf{x}$ (and therefore $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{f} = \boldsymbol{\Phi}\boldsymbol{\Psi}\mathbf{x}$) and $\mathbf{x}$ is sparse, then reconstruction is possible. Since reconstruction is complex, in typical sensing systems, it is usually performed at an energy-unconstrained receiver, such as in a cellphone or the cloud.

B. SC-Assisted Reconstruction

Various algorithms are available for reconstruction [6]. Most of them attempt to minimize the l1-norm of $\mathbf{x}$ using an iterative approach. The gradient-projection for sparse reconstruction (GPSR) algorithm is one such algorithm that minimizes the l1-norm and optimization error simultaneously, as shown below [7]:

$$x^* = arg\min_{x} \tau\|\boldsymbol{x}\|_1 + 0.5\|\boldsymbol{y} - \boldsymbol{\Phi x}\|_2^2 \qquad (1)$$

where τ is a non-negative real parameter. In the above formulation, note that τ can also be a column vector with the same number of non-negative entries as $\mathbf{x}$. In this case, the objective function then weighs each element of $\mathbf{x}$ differently such that the following relationship holds true:

$$x^* = arg\min_{x} \tau^T abs(x) + 0.5\|y - \boldsymbol{\Phi} x\|_2^2 \qquad (2)$$

In this paper, our system first extracts some statistics about the sensor's input signal (*i.e.* timing location of the R-peaks in the ECG). Based on these statistics, during reconstruction at the receiver, we choose smaller values for the elements of vector τ that correspond to more important timing locations of the input signal. Using this technique, we can significantly decrease the reconstruction error (results shown in Section V). One drawback of this method is that the input signal $\mathbf{x}$ must be sparse in time. To the best of our knowledge, there are currently no similar algorithms for signals that are sparse in other domains (*e.g.*, wavelet), so we leave this for future work.

III. SENSOR-ON-CHIP ARCHITECTURE

Fig. 1 shows the block diagram of the proposed SC-assisted compressed-sensing SoC. It consists of an analog front-end (AFE), analog-to-digital converter (ADC), and digital-processing module. The digital-processing module comprises of specialized elements for compressed sensing (CS), statistics collection (SC), and packet management (PM) that organizes all of the required data into a packet format and relays them to the serial interface. Next we present implementation details on each of the sub-blocks.

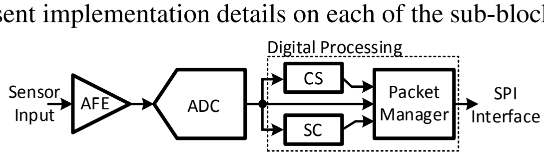

Fig. 1. Block diagram of the proposed SC-assisted compressed-sensing SoC.

A. Analog Front-end

The AFE comprises of a chain of current-feedback instrumentation amplifier (CFIA), a fixed-gain amplifier (FGA) and a programmable-gain amplifier (PGA). It also includes a DC servo-loop (DSL) and a differential output buffer (X1). Fig. 2 shows block diagram of the AFE and its output for an ECG input. The FGA is used to adjust the bandwidth of the amplifier from 250-935 Hz, while the PGA is used to adjust the gain between 40-53.7 dB. The DSL rejects the offset of the input signal and also sets the high pass corner of the amplifier. The schematic of the CFIA is shown in Fig. 3. It includes three transconductance stages, one each for the input signal (V_{in}), feedback input (V_{fb}), and DSL output (V_{os}). The current through the input and feedback differential pairs are equal in order to obtain equivalent gains, while a common-source stage is used to drive the load. The measured AFE performance is listed in TABLE I.

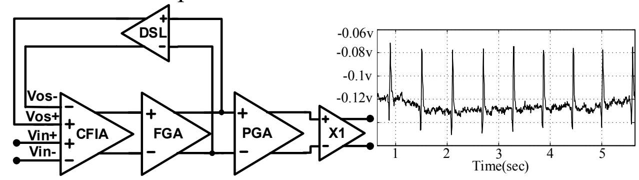

Fig. 2. Block diagram of the AFE.

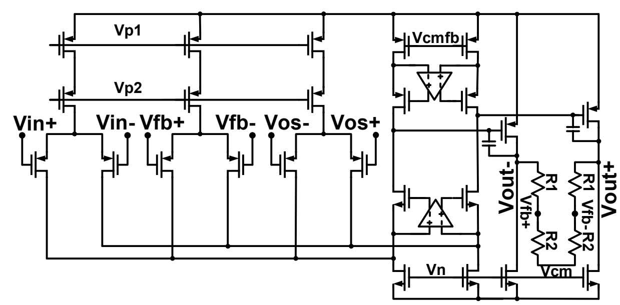

Fig. 3. Schematic of the CFIA.

B. Analog-to-Digital Converter

For high-resolution data conversion, higher-order ΔΣ ADCs are more prone to instability, and exhibit a limited non-overloaded input range. Therefore, we integrate a two-step incremental ADC (IADC), which was previously presented in [8]. The two-step operation achieves nearly third-order noise shaping performance while using only second-order IADC circuitry, improving the energy efficiency significantly. The measured performance shows that at 500S/s, the ADC achieves a 100 dB dynamic range and 91 dB SNDR for a 250 Hz input signal. The device consumes only 10.7 µW of power with an active area of only 0.2 mm^2, which is the smallest among similarly performing ADCs.

C. Digital-Processing Module

On-chip digital processing includes three main blocks. First, the compressed-sensing (CS) block implements compression, where the number of measurements (M) is fixed to 50 in this design. To change the compression factor (CF), the user needs to adjust the length of the input signal (N). Second, the SC block extracts the maximum, threshold point, and for an ECG input, the location and width of the R-peaks. Third, the PM organizes all of the required data into a packet and relays them out of the chip through the SPI port. In the next section, we present more details on each of these blocks.

IV. Digital System Implementation

A. Compressed-Sensing Block

This block implements the multiplication of a set of N input samples (coming from the ADC) with a random M×N matrix $\mathbf{\Phi}$ with randomly selected 0 and 1 values. To guarantee that the columns of $\mathbf{\Phi}$ are independent, we combine the outputs of two linear-feedback shift registers (LFSR) [3]. The compressed-sensing block thus generates M output samples for every N input samples. Each of M output samples are inner products of one row of $\mathbf{\Phi}$ and the N input samples. Since the elements of $\mathbf{\Phi}$ are either 0 or 1, we use a select-and-accumulate (SAC) unit to implement this inner-product computation. The clock signal of the digital blocks has a frequency of 96 kHz and is provided by the ADC. If the corresponding element of $\mathbf{\Phi}$ is 0, the clock is gated and the output of the SAC does not change. If it is 1, then the value of the ADC sample will be added to the output of the SAC. After N ADC samples, the SAC contains the final value and is then stored in a secondary register, resetting the SAC for the next data block. Fig. 4 and Fig. 5 show the compressed-sensing block and LFSR-based random matrix generation block, respectively.

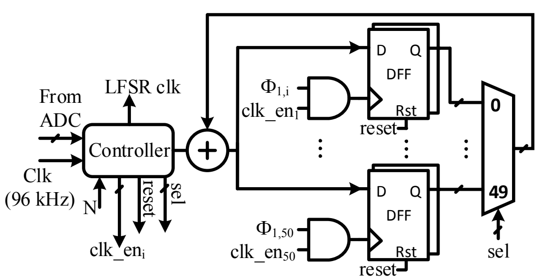

Fig. 4. Block diagram of the compressed-sensing block.

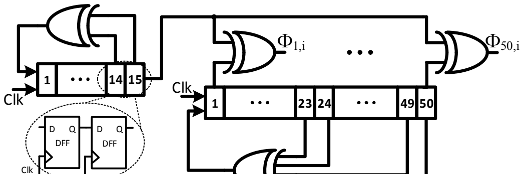

Fig. 5. Block diagram of the module used to generate the elements of the random matrix $\mathbf{\Phi}$.

B. Statistics-Collection Block

SC operates by first finding the index and value of the largest element in each set of the N input samples. It also calculates the start index and width of a consecutive number of data samples that are above a user-defined threshold level. To limit the amount of data overhead created by the SC, the SC only extracts the first two locations of where the threshold is passed. The SC also finds the location and width of the first two R-peaks in the N samples. Thus, for every block of N ADC samples, the SC generates ten 16-bit values for the above statistics, resulting in a 20% data overhead. In practice, this overhead can be easily further reduced by using a larger value of M. Fig. 6 shows block diagram of SC block.

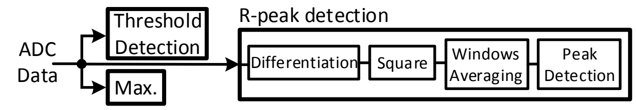

Fig. 6. Block diagram of statistics-collection block.

C. Packet Manager Block

The PM organizes all of the data into a single packet by adding a header and footer, and then relaying the data out of the chip using an SPI interface. SPI is a three-wire serial interface with a clock, serial data, and select signals. Serial bits are transmitted during the clock's rising edges when the select bit is low.

V. Measurement Results

To verify the SoC's measured performance, we generate a random sparse-in-time signal pulse using an external source applied to the SoC input. Next, we collect statistics from the SC module to find the location and width of this pulse. Fig. 7 shows the reconstructed signal with and without using the proposed SC-assisted reconstruction. As can be seen, the reconstruction using SC is almost perfect versus the conventional CS-only version. Note that as discussed in Section II, the current GPSR algorithm currently works only with signals that are sparse-in-time (*i.e.,* $\mathbf{\Psi}$ is a unity matrix), such as UWB signals, neural spikes, and time-sparse ECG waveforms. As part of future work, we plan to use a modification of this algorithm for signals that are sparse in other domains.

Fig. 8 shows the SER for different compression factors with and without using statistics for the same pulse-like input. The SER is defined in Eq. (3) where $\boldsymbol{f}$, $\boldsymbol{f}^*$ are the original and reconstructed signals, respectively:

$$SER = 10\log_{10}\left(\frac{\|f\|_2^2}{\|f - f^*\|_2^2}\right) \tag{3}$$

As shown in Fig. 8, SC-assisted reconstruction improves SER by as much as 18 dB.

Finally, we demonstrate the effect of using statistic collection on the SER for a real ECG signal input, where timing samples that are not located near the R-peak location and below an adjustable threshold value are converted to zero, thereby providing sparsity in the time domain. Measurements of the reconstructed ECG R-peak waveform are shown in Fig. 9. For this real ECG waveform experiment, we observe a 9 dB improvement in the SER. The die photograph of the SoC is shown in Fig.

10, and a comparison with other similar compressed-sensing systems is summarized in TABLE I.

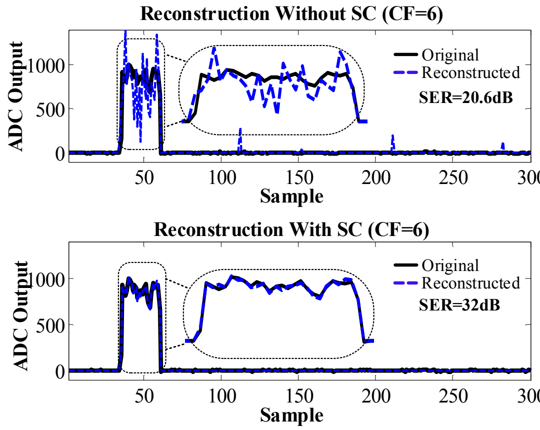

Fig. 7. SC-assisted reconstruction versus conventional reconstruction for a pulse-like waveform.

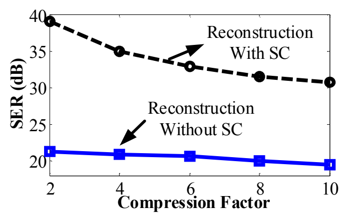

Fig. 8. SER vs. compression factor with and without using signal statistics.

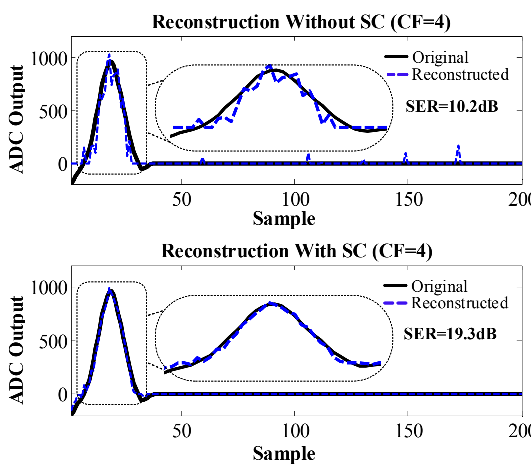

Fig. 9. Impact of SC on reconstruction for an ECG signal.

VI. Conclusion

In this paper we presented a compressed-sensing SoC for compressing biomedical sensor data that incorporates statistics collection (SC). We showed that using some statistics from the input signal can dramatically improve the performance of reconstruction algorithms for signals that are sparse in the time domain. As part of future work, we are modifying the GPSR algorithm such that it can be used for signals that are sparse in non-time domain bases.

Acknowledgement

This work was partially funded by Center for Design of Analog-Digital Integrated Circuits (CDADIC), and the NSF CAREER program.

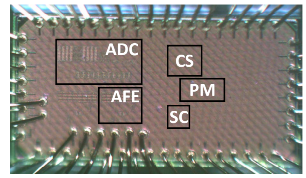

Fig. 10. Die photograph of the proposed SoC.

TABLE I: Performance summary and comparison with previous compressed-sensing SoCs

	Previous CS Sensor-on-Chip [3]	Previous CS Sensor-on-Chip [4]	This Work
AFE	No	No	Power: 10uW Gain: 40/53.7 dB BW:250/476/721 Hz Input Impedance: 2.5G CMRR: 73 dB @50Hz
ADC	SAR 5-bits ENOB < 1.1µW	SAR 6.5-bits ENOB	Incremental 15-bits ENOB 10.7 µW
Compression	2.5µW @ 100 kHz	28nW @ 2kHz	10.7 µW Static 0.7 µW Dynamic @ 96kHz
Statistics Collection	No	No	12.6 µW Static 0.4 µW Dynamic @ 96kHz
Technology	90 nm CMOS	0.13 µm CMOS	65nm CMOS

References

[1] F. Chen, "Energy-efficient Wireless Sensors: Fewer Bits, More MEMS," PhD Thesis, MIT, Sep. 2011.

[2] V. Behravan, et al., "Rate-Adaptive Compressed-Sensing and Sparsity Variance of Biomedical Signals," Body Sensor Networks 2015 (to appear).

[3] F. Chen, et al., "A signal-agnostic compressed sensing acquisition system for wireless and implantable sensors," Custom Integrated CCkts. Conf., Sep. 2010.

[4] Gangopadhyay, et al., "Compressed Sensing Analog Front-End for Bio-Sensor Applications," IEEE J. Solid-State Circuits, vol.49, no.2, pp.426,438, Feb. 2014.

[5] E. Candes, et al., "Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?," IEEE Trans. Information Theory, Dec. 2006

[6] http://statweb.stanford.edu/~tibs/lasso.html

[7] M. Figueiredo, R. Nowak and S. Wright, "Gradient Projection for Sparse Reconstruction: Application to Compressed Sensing and Other Inverse Problems, " IEEE J. Sel. Topics in Signal Processing, Dec. 2007.

[8] C.H. Chen, et al., " A 11 µW 250 Hz BW two-step incremental ADC with 100 dB DR and 91 dB SNDR for integrated sensor interfaces," Custom Integrated CCkts. Conf., Sep. 2014.