Energy-Aware Design of Compressed Sensing Systems for Wireless Sensors Under Performance and Reliability Constraints

Fred Chen, *Student Member, IEEE*, Fabian Lim, Omid Abari, Anantha Chandrakasan, *Fellow, IEEE*, and Vladimir Stojanović, *Member, IEEE*

***Abstract*—This paper describes the system design of a compressed sensing (CS) based source encoding system for data compression in wireless sensor applications. We examine the trade-off between the required transmission energy (compression performance) and desired recovered signal quality in the presence of practical non-idealities such as quantization noise, input signal noise and channel errors. The end-to-end system evaluation framework was designed to analyze CS performance under practical sensor settings. The evaluation shows that CS compression can enable over 10X in transmission energy savings while preserving the recovered signal quality to roughly 8 bits of precision. We further present low complexity error control schemes tailored to CS that further reduce the energy costs by 4X as well as diversity scheme to protect against burst errors. Results on a real electrocardiography (EKG) signal demonstrate 10X in energy reduction and corroborate the system analysis.**

***Index Terms*—Compressed sensing, error correction codes, source coding, wireless sensors, energy efficiency.**

I. Introduction

OVER THE PAST few decades, CMOS scaling and advancements in IC design have progressively reduced the cost of computing, signal processing, and communication. Wireless sensors have been one of the many applications enabled by the miniaturization of CMOS, and now have reemerged as a vital "green" technology to enable efficiency across a variety of industries including building management, agriculture, transportation, and healthcare [1].

Most of the technical obstacles to wider adoption of wireless sensor technology, such as data reliability and node lifetime, can be linked to the stringent energy constraints of each sensor node [2]. As Fig. 1 shows, the energy cost to wirelessly transmit data is typically orders of magnitude greater than any other function that is common within a wireless sensor. Thus, in order to address sensor lifetime it is paramount to minimize the amount of transmitted data. Consequently there have been many efforts across a variety of applications to find low-cost, low-energy data compression schemes that can be implemented locally at the sensor node [13]–[15]. For each scheme, there is a trade-off between data reduction, information integrity, and implementation cost, where the goal is to minimize the total amount of energy required to preserve the desired information.

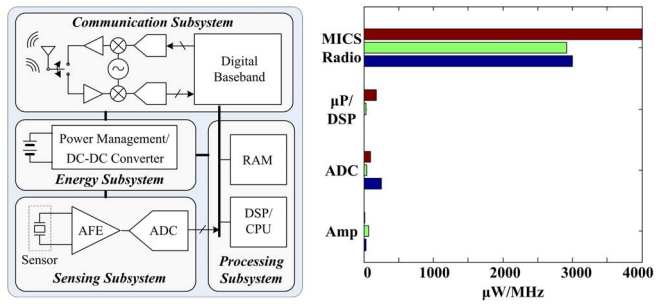

Fig. 1. Common functional blocks within a wireless sensor node and their associated energy costs from *state-of-the-art* publications [3]–[13].

Based on recent research, compressed sensing (CS) shows promise as a potential data compression scheme for wireless sensors. Unlike Shannon sampling theory that explains signal acquisition rates with respect to bandwidth, CS techniques are developed to capture data at approximately the information rate [16]. Furthermore, the compressed samples (or measurements) in CS can be randomly captured, enabling *universal* data compression for sparse signals [17]. In our previous work, we highlighted the applicability of CS techniques for data compression in wireless sensors by demonstrating an energy-efficient hardware realization of the CS data encoder [18], [19].

Having shown that the implementation cost for CS can be minimal, our focus here is to discuss the design of the CS system under a practical, wireless sensor environment. In particular, we explore how impairments from sensor noise, signal quantization and channel noise impact the required transmitted signal energy needed to meet the desired signal fidelity at the receiver. As we will show, CS can not only lower the transmission energy for a given fidelity, but can simultaneously improve the attainable signal fidelity when in the presence of sensor noise. We also show that the degradation in reconstruction performance due to channel errors is gradual for CS, whereas standard lossless source coding algorithms such as Huffman [20] and Lempel-Ziv (LZ) [21] are susceptible to error propagation when decoding corrupted compressed data [22]. Although CS does not require additional layers of coding protocols to insure data integrity, as

Manuscript received October 11, 2011; revised February 21, 2012; accepted March 06, 2012. Date of current version February 21, 2013. This work was supported in part by the MIT CICS, the NSF award ECCS-1128226 and the National Science and Engineering Research Council of Canada (NSERC). This paper was recommended by Associate Editor M. Chakraborty.

The authors are with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: fredchen@alum.mit.edu; flim@mit.edu; abari@mit.edu; anantha@mtl.mit.edu; vlada@mit.edu).

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/TCSI.2012.2215738

is typical with LZ and Huffman [22], channel coding is generally desirable, and we show how CS is specifically amenable to simpler error detection schemes that perform on par with more complex error correction schemes.

The remainder of the paper is organized as follows. Section II first briefly reviews background on CS and sets the terminology used in this paper. In Section III, we describe the framework, parameters, and metrics under which the design space is explored. Section IV quantifies the energy cost of channel errors in a CS system while Section V discusses the impact of different error coding strategies and proposes a couple of coding schemes specific to CS. Section VI, describes the results in the context of a real EKG signal, and finally we offer conclusions in Section VII.

II. Compressed Sensing Background

CS is based on the following key concepts which will be discussed hereafter: signal sparsity, signal reconstruction, and incoherent sampling.

A. Signal Sparsity

CS theory first and foremost assumes that the signal of interest, $\boldsymbol{f}$, has a *sparse* representation in some basis $\boldsymbol{\Psi} = [\psi_1\, \psi_2\, \ldots\, \psi_{\mathrm{L}}]$ where $\boldsymbol{f} = \boldsymbol{\Psi}\boldsymbol{x}$ or equivalently:[1]

$$f = \sum_{i=1}^{L} x_i \psi_i \tag{1}$$

where x_i denotes the ith coefficient of signal vector $\boldsymbol{x}$, and consists of only a few non-zero (or significant) values. For example, a sine wave captured in the time domain requires an infinite number of non-zero samples, whereas it only requires a single non-zero coefficient in the Fourier domain. Fortunately, many sensor signals of interest are compressible and have sparse representations in some signal basis, such as time, frequency, Gabor, wavelets, etc., [23]–[25]. CS techniques are designed specifically for such sparse representations, and are therefore applicable to many sensor systems.

B. Signal Recovery From Incomplete Measurements

When the signal is sparse, CS theory proposes that acquiring only a small number of compressed measurements can capture the necessary information to recover the signal. This framework is shown in Fig. 2, where the N-dimensional input signal, $\boldsymbol{f}$, is directly encoded into an M-dimensional set of measurements $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{f}$ via an $M \times N$ measurement matrix, $\boldsymbol{\Phi}$. When $M < N$ the linear system of equations $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{f}$ is underdetermined, therefore multiple signals, $\boldsymbol{f}$, may produce the same measurements, $\boldsymbol{y}$, making the problem of reconstructing $\boldsymbol{f}$ from $\boldsymbol{y}$ ill-posed. Fortunately, the sparsest solution to $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{f}$ is often unique, and the typically intractable (NP) sparse recovery problem is well approximated by the convex relaxation

$$\min_{\boldsymbol{x}\in\boldsymbol{R}^n} \|\boldsymbol{x}\|_{\boldsymbol{\ell}_1} \text{ subject to } \boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{\Psi}\boldsymbol{x} \tag{2}$$

[1]We refer to $\boldsymbol{\Psi}$ as a basis in this section, but in practice $\boldsymbol{\Psi}$ does not need to be a set of linearly independent vectors. For example, as we will show later, $\boldsymbol{\Psi}$ can be an overcomplete dictionary.

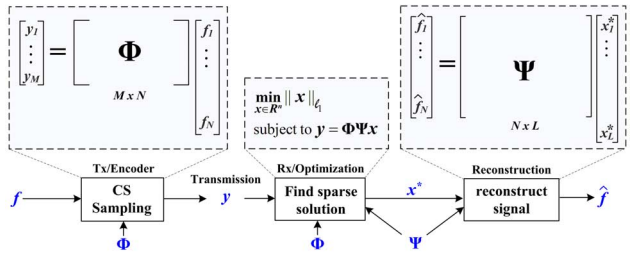

Fig. 2. Compressed sensing framework.

where $\boldsymbol{\Psi}$ is a known $N \times L$ basis matrix and $\boldsymbol{x}$ is the coefficient vector from (1). The recovered signal is then $\hat{\boldsymbol{f}} = \boldsymbol{\Psi}\boldsymbol{x}^*$ where $\boldsymbol{x}^*$ is the optimal solution to (2). The problem of minimizing the ℓ_1-norm in (2) is a linear-programming problem for which efficient techniques have been developed [26]. In the following experiments the signals are reconstructed by solving a modification of (2) using the Lasso-modified LARS algorithm [27]. It should be noted that the signal can be reconstructed using a variety of different algorithms besides what we chose [28]–[33]. The main point is that these reconstruction algorithms have practical implementations for estimating the sparse solution. The practical feasibility of solving (2) (or similar) implies that an N-dimensional signal can in essence be recovered (sometimes perfectly) from a lower number of samples, M, provided that the signal is sparse under $\boldsymbol{\Psi}$. The ratio N/M is essentially the data compression factor (CF) realized by the CS system, and is proportional to the transmission energy saved.

C. Incoherent Sampling

In addition to signal sparseness, the conditions for sparse recovery also include the *incoherence* between the sensing matrix ($\boldsymbol{\Phi}$) and the signal basis ($\boldsymbol{\Psi}$). The coherence parameter $\mu(\boldsymbol{\Phi}, \boldsymbol{\Psi})$ measures the largest correlation between any row of $\boldsymbol{\Phi}$ and any column of $\boldsymbol{\Psi}$, and is defined as

$$\mu(\boldsymbol{\Phi}, \boldsymbol{\Psi}) = \max_{k,j} |\langle \varphi_k, \psi_j \rangle| \tag{3}$$

where the coherence $\mu(\boldsymbol{\Phi}, \boldsymbol{\Psi})$ ranges between 1 and N (where φ_k and ψ_j are unit vectors). The more incoherent the matrices $\boldsymbol{\Phi}$ and $\boldsymbol{\Psi}$ are, the fewer the number of measurements that are needed to recover the signal. A lower bound on M to recover an S-sparse signal (a signal with S non-zero terms out of N in the basis $\boldsymbol{\Psi}$) was shown to be ([34])

$$M \geq C \cdot \mu^2(\boldsymbol{\Phi}, \boldsymbol{\Psi}) \cdot S \cdot \log N \tag{4}$$

where C is a small known constant (empirically $\sim 2 - 2.5$ [33]) and N is the dimensionality of the signal to be recovered. Thus, when the normalized coherence parameter μ/N is roughly constant, an S-sparse signal can be reconstructed from approximately $S \cdot \log(N)$ measurements. Since S is the *rate of innovation* (or information rate) of the sparse signal, (4) indicates that the required number of samples for successful recovery is proportional to the information content of the signal. Random sensing matrices with independent entries are typically incoherent with most fixed bases as the sample size becomes large [17]. Therefore, CS sampling with random matrices can

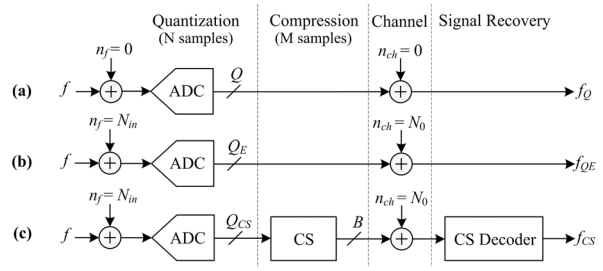

Fig. 3. Evaluation system models for an input signal block of N samples: (a) a baseline system with only quantization error, (b) a system with input noise, quantization noise, and channel noise, and (c) a CS system with input noise, quantization noise, compression error and channel noise.

be viewed as a universal encoder that enables good signal reconstruction without requiring *a priori* knowledge of the signal basis ($\mathbf{\Psi}$).[2] In the context of wireless sensors, this enables a generic infrastructure for data acquisition and compression that is agnostic to the type of signals being acquired, provided that they are sparse.

III. Evaluation Framework

Since most work in CS theory is focused on determining asymptotic performance bounds, a simulation framework must be established to evaluate the performance of a practical CS-based wireless system when the block size (N) is finite. Thus, in this section we describe a set of models and performance metrics for comparison and choose an appropriate signal test set to capture the salient characteristics of the CS algorithms.

A. System Models and Performance Metrics

Fig. 3 shows three system models chosen to highlight the main issues and challenges of designing the system. Fig. 3(a) shows an idealized wireless sensor node infrastructure, which assumes that the input signal to the ADC is noiseless ($n_f = 0$), and that the quantized data is received error-free ($n_{ch} = 0$). The quantization resolution, Q, is chosen to meet the error requirement of the application. Since any practical system will require Q to be finite, we treat this as the baseline performance for comparison with subsequent systems.

Fig. 3(b) shows a more complete model that includes signal noise (n_f) and channel noise (n_{ch}). The last system model shown in Fig. 3(c) includes the effects of compression with CS. Although the CS framework does not necessarily require the input to be quantized, it does require the dynamic range of the input to be preserved. So for Fig. 3(c), Q_{CS} represents the number of bits that correspond to the input dynamic range. To compare the systems shown in Fig. 3, we adopt the *percent root-mean-square difference (PRD)* metric, which is commonly used in quantifying information loss in biomedical signals [37]. This metric is defined as

$$\mathrm{PRD} = 100\sqrt{\frac{\sum_{n=1}^{N}|f[n]-\tilde{f}[n]|^2}{\sum_{n=1}^{N}|f[n]|^2}} \tag{5}$$

where $\tilde{\boldsymbol{f}}$ is the estimate $(f_Q, f_{QE}, f_{\mathrm{CS}})$ of the input signal $\boldsymbol{f}$.

B. Reconstruction Error and Quantization Error in CS

To better understand the performance comparison of CS to more traditional systems, we first discuss the noiseless ($n_{ch} = n_f = 0$) signal recovery performance of CS. CS is generally a lossy compression scheme whose reconstruction performance is dependent on the number of significant features (S) in the signal, the block size (N) of signal samples to compress, the resolution of each signal sample (Q_{CS}), the number of compressed measurements (M), and the resolution of each measurement (B). For any random combination of a measurement matrix, $\mathbf{\Phi}$, and an input signal, $\boldsymbol{f}$, there is some non-zero probability that there will be an ill-conditioned pair of $\boldsymbol{f}$ and $\mathbf{\Phi}$ that results in poor reconstruction error.

To illustrate these dependencies, we constructed random 4-sparse input signals of length $N = 1000$ from an over-complete dictionary of a sample-shifted Gaussian pulses (used as $\mathbf{\Psi}$) to pass through the noiseless CS system. The signals were generated by drawing on a uniform random distribution over $[-1, 1]$ to assign the sign and magnitude, and over [1, 1000] to assign the position of each pulse in the 4-pulse signal. This signal dictionary was chosen as an example to provide a simple signal model to understand, similar to an N × N identity matrix, yet better representing signals acquired in an over-sampled system. Also, as we will demonstrate later, it enables reasonably good reconstruction of real EKG signals.

For each randomly generated signal, the resolution of each compressed measurement (B) is allowed to be as large as needed to accommodate the measurement range, but at a minimum resolution equivalent to the input quantization (Q_{CS}). So for example, if Q_{CS} consists of 4 fractional bits (i.e., $Q_{\mathrm{CS}} = 5$ bits, which includes 1 sign bit), then the resolution of each measurement (B) will also have only 4 fractional bits, but will extend the number of integer bits as needed to handle the additional range from random accumulation of N input samples. To maintain the hardware simplicity described in [19], the measurement matrix, $\mathbf{\Phi}$, is chosen to be a random Bernoulli matrix (± 1 entries) generated from a randomly seeded pseudo-random bit sequence (PRBS) generator.

Fig. 4 shows example PRD distributions (in dB) for 10 000 different input signals for system configurations of $M = 50, 100$ and $Q_{\mathrm{CS}} = 8, 12$. Fig. 4 also shows the PRD corresponding to the net PRD ($\mathrm{PRD}_{\mathrm{net}}$), average PRD ($\mathrm{PRD}_{\mathrm{avg}}$) and the PRD from just quantizing the inputs to $Q = Q_{\mathrm{CS}}$ bits (PRD_Q). The net PRD is defined as

$$\mathrm{PRD}_{\mathrm{net}} = 100\sqrt{\frac{\sum_{k=1}^{K}\sum_{n=1}^{N}|f_k[n]-\tilde{f}_k[n]|^2}{\sum_{k=1}^{K}\sum_{n=1}^{N}|f_k[n]|^2}} \tag{6}$$

[2]It should be noted that while $\mathbf{\Psi}$ does not need to be known at the sensor/encoder, it does need to be known in order to recover the original signal. If $\mathbf{\Psi}$ is not known, a practical approach to determining $\mathbf{\Psi}$ is to train $\mathbf{\Psi}$ based on captured uncompressed data first as described in [19].

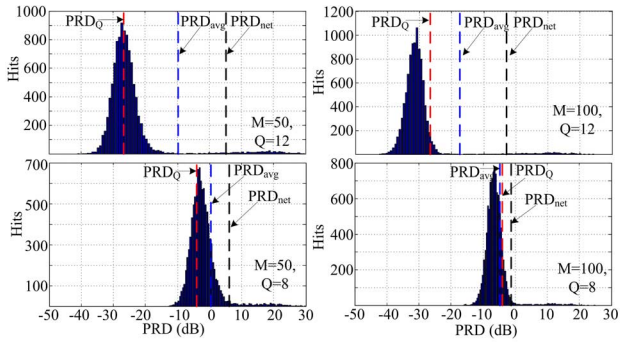

Fig. 4. Distribution of PRD over 10 000 input signals of length 1000 (N) where $M = 50, 100$ and $Q_{\mathrm{CS}} = 8, 12$.

where K is the total number of input signal blocks (10 000 in this example), and is equivalent to the PRD if all K input blocks were considered as a single input signal. The distributions highlight the fact that a small subset of signal blocks can result in errors many orders of magnitude greater than the majority of signal blocks, and that these signals largely dictate the value of $\mathrm{PRD}_{\mathrm{net}}$. The distributions also show that for a 4-sparse signal, the inherent *net* reconstruction error, from taking less than 100 measurements, is greater than the error due to quantization to 8 bits, as $\mathrm{PRD}_{\mathrm{net}}$ is largely independent of signal resolution above 8 bits.

Meanwhile, the average PRD represents what the expected PRD for one block of an S-sparse, N sample signal would be

$$\mathrm{PRD}_{\mathrm{avg}} = \frac{100}{K}\sum_{k=1}^{K}\sqrt{\frac{\sum_{n=1}^{N}|f_k[n]-\tilde{f}_k[n]|^2}{\sum_{n=1}^{N}|f_k[n]|^2}}. \qquad (7)$$

To distinguish the importance of these two measures, $\mathrm{PRD}_{\mathrm{avg}}$ is equivalent to the time-averaged error performance, whereas $\mathrm{PRD}_{\mathrm{net}}$ is the accumulated error performance. These results are consistent with the discrepancy shown between the worst case and average case analysis of CS reconstruction error in [35]. Depending on the application, one metric may be more important than the other.

The reconstruction error dependencies are summarized in Fig. 5 where $\mathrm{PRD}_{\mathrm{avg}}$ and $\mathrm{PRD}_{\mathrm{net}}$ are plotted over the M and Q_{CS} design space. Fig. 5 highlights the performance limitations imposed on CS due to the number of measurements and the quantization of the input signal. Below a certain number of measurements ($\sim$50) the average recovered signal error ($\mathrm{PRD}_{\mathrm{avg}}$) is dominated by reconstruction error. However, for larger M, the recovered signal becomes more limited by quantization noise. The same can be said of the net signal error ($\mathrm{PRD}_{\mathrm{net}}$), only the threshold at which this crossover occurs is at a higher value of M. When the error is limited by M, there is a trade-off between compression performance and fidelity of the recovered signal. Even though N is fixed for the results shown in Fig. 5, the impact of changing N can be inferred; in other words, increasing M is similar to decreasing N and vice versa, where the relationship is roughly dictated by (4).

As a point of reference, the original input and reconstructed waveforms corresponding to a PRD of 0.1% and 10% are plotted in Fig. 6; they show that while the quantitative difference is significant between signal recovery for a signal in the bulk of

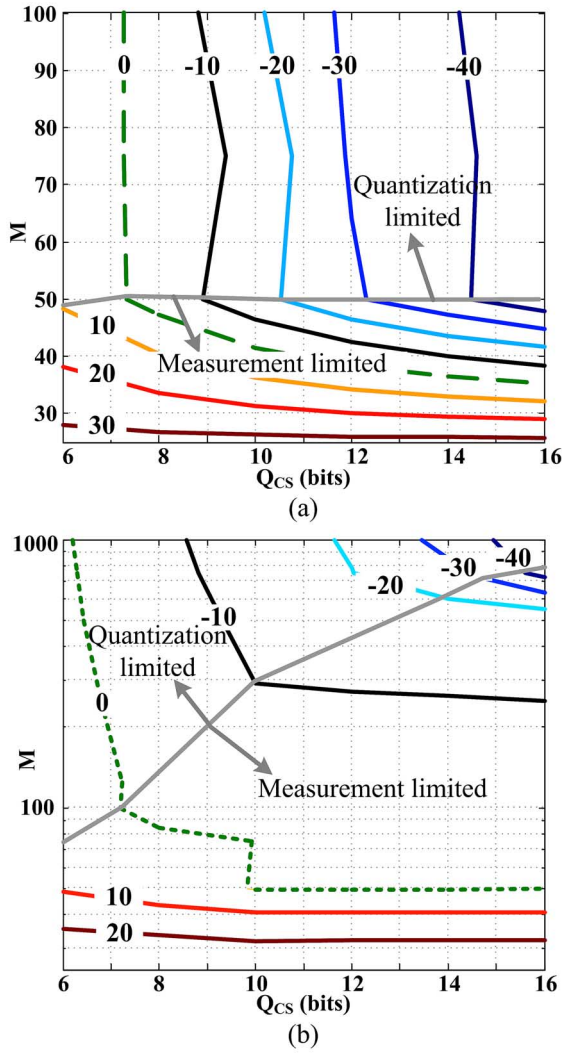

Fig. 5. Contours of the (a) average PRD (in dB) and (b) net PRD (in dB) for 4-sparse signals across the input quantization (Q_{CS}) and CS measurement (M) design space.

the distribution and one in the tail, the perceptual difference is small in either case. In subsequent discussion, we will refer to the $\mathrm{PRD}_{\mathrm{avg}}$ signal as a signal that represents the average reconstructed PRD per block (i.e., the bulk of the distribution). Meanwhile, the $\mathrm{PRD}_{\mathrm{net}}$ signal will refer to the net reconstruction error (i.e., representative of signals in the tail).

C. Compression Performance of CS

Since the purpose of adopting CS theory is for data compression, the compression performance of CS is discussed and compared against the Huffman [20] and LZ [21] source encoding algorithms. The comparisons will be limited to these lossless compression schemes since their recovered signal error is equivalent to just the quantization error. The coding efficiency of each option is measured in bits per sample which is just the number of transmitted bits divided by the number of input samples those bits represent. Fig. 7 plots the resulting compression performance for the $\mathrm{PRD}_{\mathrm{avg}}$ and $\mathrm{PRD}_{\mathrm{net}}$ signals versus their level of quantization. For every case except CS, the signals are each quantized to Q bits before compression. For the CS system, M and Q_{CS} are chosen such that the PRD of the reconstructed

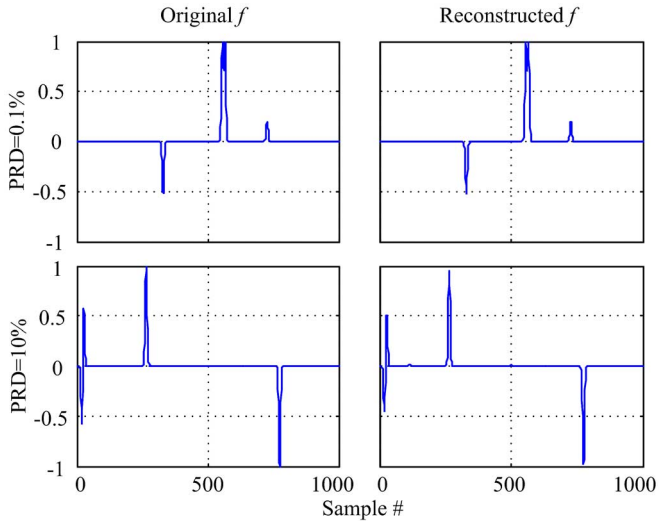

Fig. 6. Original and reconstructed signals for when the PRD ~0.1% (−20 dB) and when the PRD ~10% (20 dB) for $M = 75$, $Q_{\mathrm{CS}} = 10$.

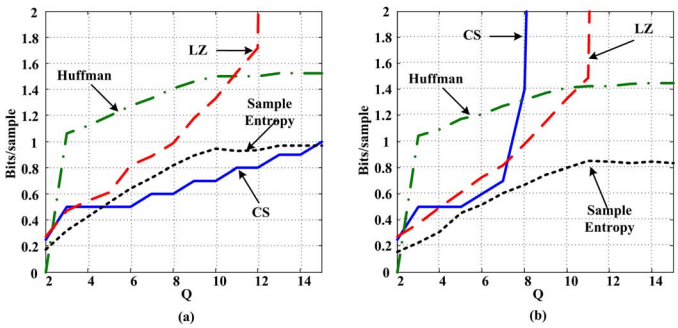

Fig. 7. Number of bits per sample corresponding to the sample entropy, Huffman coding, and LZ coding for the 4-sparse (a) $\mathrm{PRD}_{\mathrm{avg}}$ and (b) $\mathrm{PRD}_{\mathrm{net}}$ signals when quantized to Q bits. For CS, the bits per sample represent $M \cdot B/N$ where M and Q_{CS} are chosen such that the reconstructed PRD matches the output of the baseline system for each target Q value.

signal is equal to or less than the PRD due to quantization to Q bits. The coding efficiency of the CS system is $M \cdot B/N$ where B is about 4 bits larger than Q_{CS}. Additionally, the theoretical entropy per sample is also calculated and plotted where the sample entropy is defined as

$$H(X) = -\sum_{i=1}^{N} p(x_i) \log_2 p(x_i). \tag{8}$$

where p is the probability mass function of X.

As Fig. 7 shows, the coding efficiency of CS roughly tracks the sample entropy and generally outperforms Huffman and LZ when the reconstruction error is not dominant.[3] This should be expected as CS is lossy whereas Huffman and LZ are both lossless algorithms. For $\mathrm{PRD}_{\mathrm{net}}$, the coding efficiency of CS degrades rapidly beyond 8 bits of quantization when the reconstruction error begins to dominate and M must be increased to meet the reconstruction error target.

[3]It should be noted that the results for the Huffman coding algorithm assume that all N samples of the signal have been stored in order to calculate the sample statistics. In both cases, LZ does poorly at higher resolutions since there is less sequence repetition, thus requiring a larger, less efficient code book. Both the Huffman and LZ algorithms would expect to perform somewhat better with longer input sequences. However, such a choice will negatively impact the required hardware [19].

IV. Cost of Transmission Errors

Regardless of whether a source encoding/compression scheme is employed [Fig. 3(c)] or not [Fig. 3(b)], there is an energy cost associated with ensuring a certain quality of recovered signal in the presence of channel noise. In this section, we first quantify the cost of channel errors for both systems when no additional channel coding is applied.

A. Channel Model

To capture the general effect of channel errors, the wireless channel is modeled using an additive white Gaussian noise (AWGN) channel with a noise variance of N_0 and signal energy of E_b. The bit error rate (BER) equals the Gaussian Q-function evaluated at $\sqrt{\mathrm{SNR}}$, where SNR is the signal-to-noise ratio, $\mathrm{SNR} = E_b/N_0$. We choose this for the simplicity of correlating SNR to BER, but our results should generally apply to other channel models as well (e.g., Rayleigh fading) where the final energy results will simply be shifted by their respective error distributions. For both the quantization only and CS systems, the modulation scheme is assumed to be such that the baud rate matches the bit rate (i.e., there is only one bit per symbol). Again, this is chosen for simplicity but the results are directly extendable to alternative modulation schemes. The total number of bit errors per block will depend on how many bits are transmitted per block of N samples. In the simple quantization system this is $N \cdot Q_E$ whereas for the CS system this is $M \cdot B$. Likewise, the total energy cost for each system will also depend on the number of bits transmitted per block and the required SNR per bit. It should be noted that the channel SNR, both in this context and in future discussion, is not rate adjusted so it refers to the energy per transmitted bit as opposed to the energy per transmitted information bit. This convention is chosen in part because for CS, even with no additional channel coding, the delineation between information bits and redundancy bits is not clear.

B. PRD vs. Channel Noise

For the raw quantized data stream [Fig. 3(b)], channel errors will degrade the PRD of the received signal. This relationship between channel errors and PRD is plotted in Fig. 8 where PRD_{QE} is plotted against the channel SNR for varying values of Q_E. The signals used to generate the plots are drawn from the sparse set described in Section III. At high channel SNRs, the PRD_{QE} values converge to their respective quantization noise limits. As Fig. 8 shows, to preserve the PRD of the system $(\mathrm{PRD}_{QE} = \mathrm{PRD}_{Q})$ in the presence of any appreciable channel noise, the resolution of the system must increase $(Q_E > Q)$ to counteract the errors due to channel noise.

For CS, a similar exercise can be performed to extract the relationship between the reconstructed signal PRD $(\mathrm{PRD}_{\mathrm{CS}})$ and the channel SNR. However, for CS there are additional degrees of freedom, such as M, N, and B that affect the reconstruction performance. Fig. 9 plots a snapshot of this relationship for both $\mathrm{PRD}_{\mathrm{avg}}$ and $\mathrm{PRD}_{\mathrm{net}}$ signals when $M = 50$, $N = 1000$, and B scales with Q_{CS}. The corresponding distortions from only quantization (PRD_{Q}) are also plotted. As Fig. 9 shows, the signal that is representative of the $\mathrm{PRD}_{\mathrm{net}}$ reconstruction performance essentially has a higher noise floor. As shown in Figs. 4 and

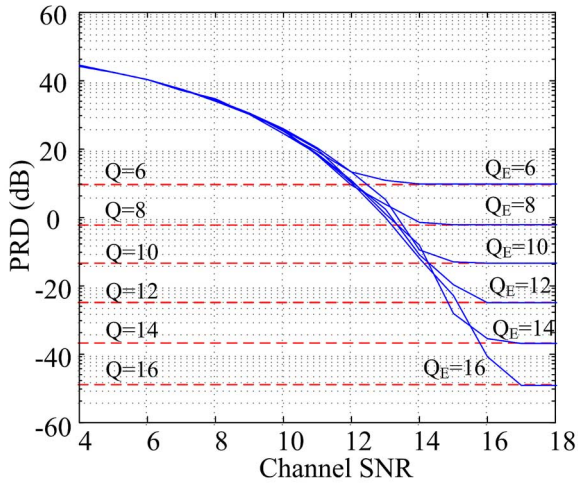

Fig. 8. PRD versus channel SNR for the traditional ADC based system shown in Fig. 3(b).

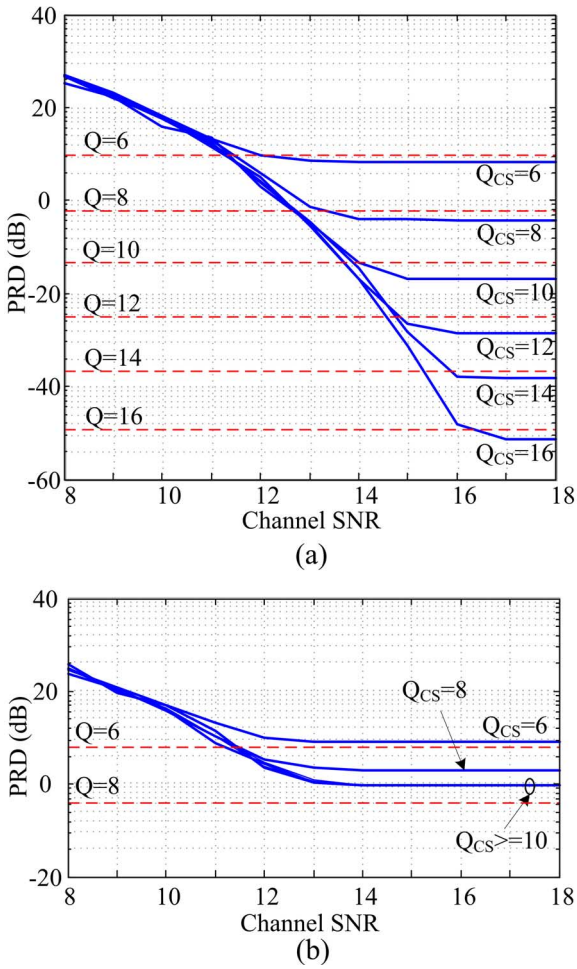

Fig. 9. PRD versus channel SNR for a (a) $\mathrm{PRD}_{\mathrm{avg}}$ and (b) $\mathrm{PRD}_{\mathrm{net}}$ signal in the CS system when $M = 50$ and $N = 1000$.

5(b), to improve the noise floor, the number of measurements (and total energy) must be increased.

C. Energy Cost of Channel Errors

As seen in Figs. 8 and 9, there are many system specifications that can achieve a targeted PRD performance. However,

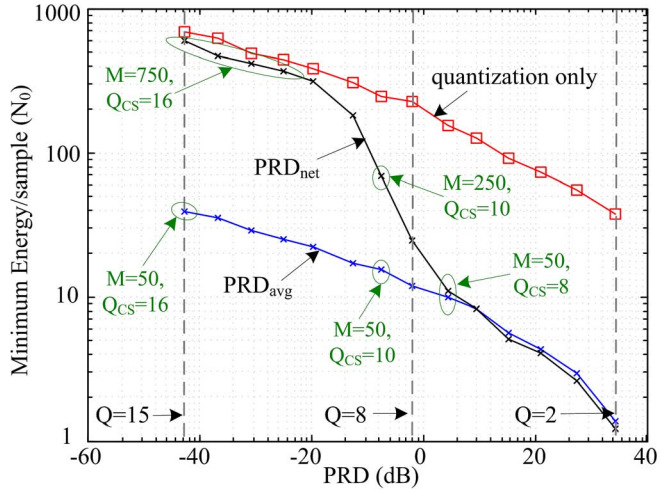

Fig. 10. Minimum energy per sample (in units of channel noise N_0) for each required PRD performance for both the uncompressed ADC system and the CS system for $\mathrm{PRD}_{\mathrm{avg}}$ and $\mathrm{PRD}_{\mathrm{net}}$ 4-sparse signals.

they are not all equally energy-efficient. There is a trade-off between transmitting more bits at a lower SNR and transmitting fewer bits at a higher SNR. For transmitting the raw quantized samples, the energy cost per sample is simply $Q_E \cdot \mathrm{SNR}_{min,E}$, where $\mathrm{SNR}_{min,E}$ is the minimum SNR that enables the system to meet the target PRD. For each CS configuration, the equivalent energy cost per sample is: $M \cdot B \cdot \mathrm{SNR}_{min,CS}/N$.

Fig. 10 plots the minimum energy cost curves over a range of target PRD performances for both systems, where the energy is in units of channel noise (N_0). For reference, several of the minimum energy configurations for the CS system are labeled in addition to the PRD performance associated with error-free quantization to 2, 8, and 15 bits of resolution. Fig. 10 shows the incremental energy cost of additional resolution in both systems and shows the order of magnitude energy savings that CS can provide through data compression. However, to achieve a low PRD ($<1\%$) or high resolution for $\mathrm{PRD}_{\mathrm{net}}$ signals requires an energy cost on par with the uncompressed quantized samples. The reason for this is because more measurements (larger M) are needed to improve the *net* reconstruction error and hence more energy is required. In contrast, if occasional performance degradation is acceptable over brief blocks of time, as with the time average PRD $(\mathrm{PRD}_{\mathrm{avg}})$, then CS can offer order of magnitude energy savings over the entire range of PRD specifications when compared to transmitting the raw quantized samples. However, just as importantly, Fig. 10 shows that channel induced bit errors in the CS measurements do not on average result in catastrophic signal reconstruction as they might with source encoding algorithms such as LZ or Huffman coding.

D. Effect of Signal Noise

So far, the discussion has been limited to noiseless input signals, which suggests that the energy/performance trade-offs described so far are minimum energy scenarios. Any sensor input signals that occur in practice will be somewhat noisy. To capture the effect of signal noise, white Gaussian noise (n_f) is added to the system inputs of Fig. 3(b) and (c). Fig. 11 shows the minimum energy curves when the input noise variance, n_f,

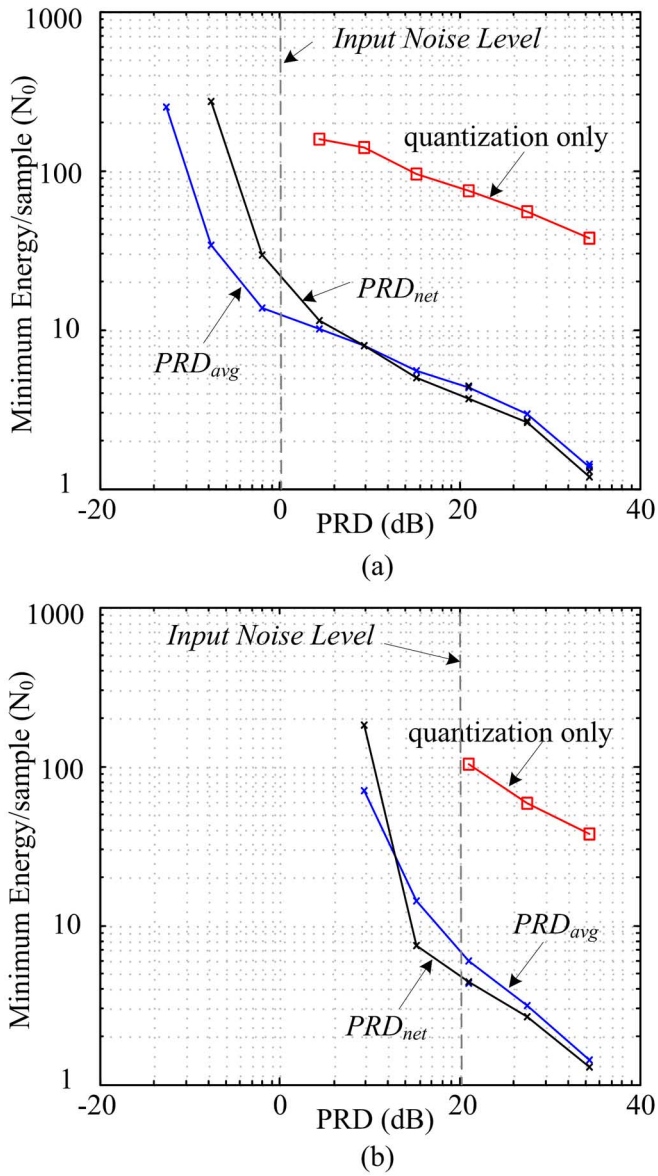

Fig. 11. Minimum energy per sample vs. required PRD performance for both the traditional system and the CS system for $\mathrm{PRD}_{\mathrm{avg}}$ and $\mathrm{PRD}_{\mathrm{net}}$ 4-sparse signals corrupted with noise such that the input PRD equals: (a) 1% and (b) 10%.

is chosen to give signal SNRs of 20 dB and 40 dB, which correspond to PRDs of 10% and 1% respectively. As to be expected, when only quantizing, our recovered signal performance is limited by the PRD of the input. The CS system, however, filters some of the input noise during reconstruction and is thus able to achieve both lower energy and better PRD results. So, while the achievable reconstruction performance of CS does not compare favorably with a simple quantizer for noiseless inputs [36], the opposite is true when it comes to more practical noisy sensor inputs.

Another observation that can be drawn from Fig. 11 is that the achievable PRD of the CS system is less limited by input noise than by quantization noise (as shown in Fig. 5). The reason for this is because the input noise is less correlated with the signal (and signal basis) and can be filtered during reconstruction whereas the quantization noise is highly correlated as is typical of oversampled inputs [38]. This result is interesting for its potential hardware implications; for example, an amplifier that drives an ADC could be allowed to have a much higher noise floor compared to the quantization noise of the ADC, which would normally be avoided in non-CS systems.

V. Protection Against Channel Non-Idealities

Thus far we have shown that CS is relatively robust to channel errors unlike typical source coding algorithms (e.g., Huffman or Lempel-Ziv) which are known to be very sensitive to transmission errors and prone to error propagation [22], [39]. Despite the graceful performance degradation in CS, channel coding can still further improve the energy/performance trade-off of the system. Since channel coding algorithms are typically designed without bias towards the message bits, we expect the coding performance of any existing error correction or detection schemes to work equally well with CS measurements as they do with raw quantized samples. However, since single bit errors in the measurement matrix can cause relatively large reconstruction errors, it is not necessarily clear how coding gain translates into energy efficiency. In this section, we propose modifications to some common existing coding strategies for use with CS and examine the energy performance gains from applying these codes.

A. Bit Errors in CS

The performance of channel coding is typically evaluated in terms of BER versus channel SNR, where a lower SNR to achieve the same BER indicates coding gain. However, this type of analysis assumes that each transmitted bit carries the same importance. Ideally, a source code would insure this condition such that the signal error is proportional to the BER of the channel. However, in the case of CS measurements, and even the raw quantized samples, this is not the case. The binary representation of each sample and measurement, while efficient, means that some bits carry more information than others. To illustrate the uneven impact of a single bit error, the average PRD versus bit position error for both CS and the raw samples is plotted in Fig. 12, where higher bit positions indicate the most significant bits (MSB) (highest bit is a sign bit). As Fig. 12 shows, a single bit error in the wrong bit position can result in significant distortion of the reconstructed signal. Given, the uneven distribution of information in the CS measurements, the task is to find the most efficient coding strategy.

B. Error Control

Communication systems and protocols generally revolve around two channel coding strategies or hybrids thereof: error correction and error detection. In error correction protocols, some form of forward error correction (FEC) is applied and sent with the data to detect and correct any errors during transmission, while error detection schemes, such as automatic repeat request (ARQ), rely only on detecting errors at the receiver and communicating back to the sender what corrupted data to resend. In both cases, there is energy overhead to offering some data reliability. Here we examine the performance of these coding strategies as applied to CS.

1) Error Correcting Block Codes: As seen in Fig. 12, errors in the least significant bits (LSB) of the measurement ma-

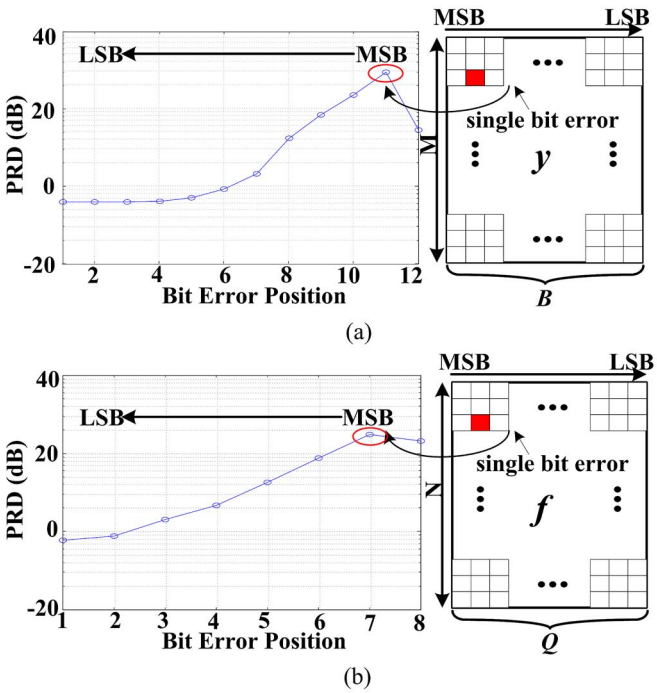

Fig. 12. $\mathrm{PRD}_{\mathrm{avg}}$ versus bit position of a single bit error in the transmitted data for: (a) a CS system and (b) a quantization only system where the input is quantized to 8 bits (i.e., $Q = Q_{\mathrm{CS}} = 8$).

trix have little impact on the reconstruction error. Thus, if the intention is to maximize the performance gain from error correction, it is most efficient to only protect the MSBs in the bit matrix. To assess the effectiveness of this approach, we apply Bose-Chaudhuri-Hocquenghen (BCH) block codes [40] across the M measurements starting with the most significant bits of each measurement. Fig. 13 illustrates the general idea behind the coding scheme as well as showing the reconstructed PRD versus channel SNR for the uncoded and BCH coded CS measurement block when $Q_{\mathrm{CS}} = 12$. Meanwhile, Fig. 14 plots the minimum-energy configurations for the different coding schemes as well as the uncoded case as a function of target PRD. In each of the plots, BCH-T refers to a BCH code with T bits of error correction capability. The actual BCH code that is chosen is determined by the number of bits in the measurement matrix and the desired error correction capability. The most efficient (lowest overhead) BCH(n, k) code is chosen whose message size (k) will cover as many of the message bits $(M \cdot B)$ without requiring filler bits. So for the examples shown in Fig. 13 when $Q_{\mathrm{CS}} = 12$, BCH-5, BCH-10, and BCH-15 correspond to BCH(511, 466), BCH(511, 421), and BCH(511, 376) codes respectively. As Fig. 14 shows, there is roughly a 4X improvement in energy-efficiency over the uncoded data scenario. The saturation of BCH-10 and BCH-15 energy plots indicates that additional error correction capability yields diminishing returns in energy-efficiency.

The above represents a first look on incorporating basic error-correction codes in CS communication schemes. A natural progression of this study is to look at modern iterative coding techniques, e.g., turbo and low density parity check codes [41], [42]. These newer codes are indeed available at comparable block/message lengths as the above BCH codes. On the other

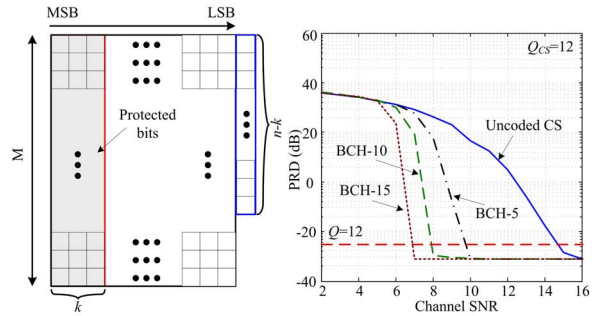

Fig. 13. Example of the proposed BCH based coding scheme and the resulting PRD for a $\mathrm{PRD}_{\mathrm{avg}}$ signal after applying the error correction code to the CS measurements when $M = 75$ and $Q_{\mathrm{CS}} = 12$.

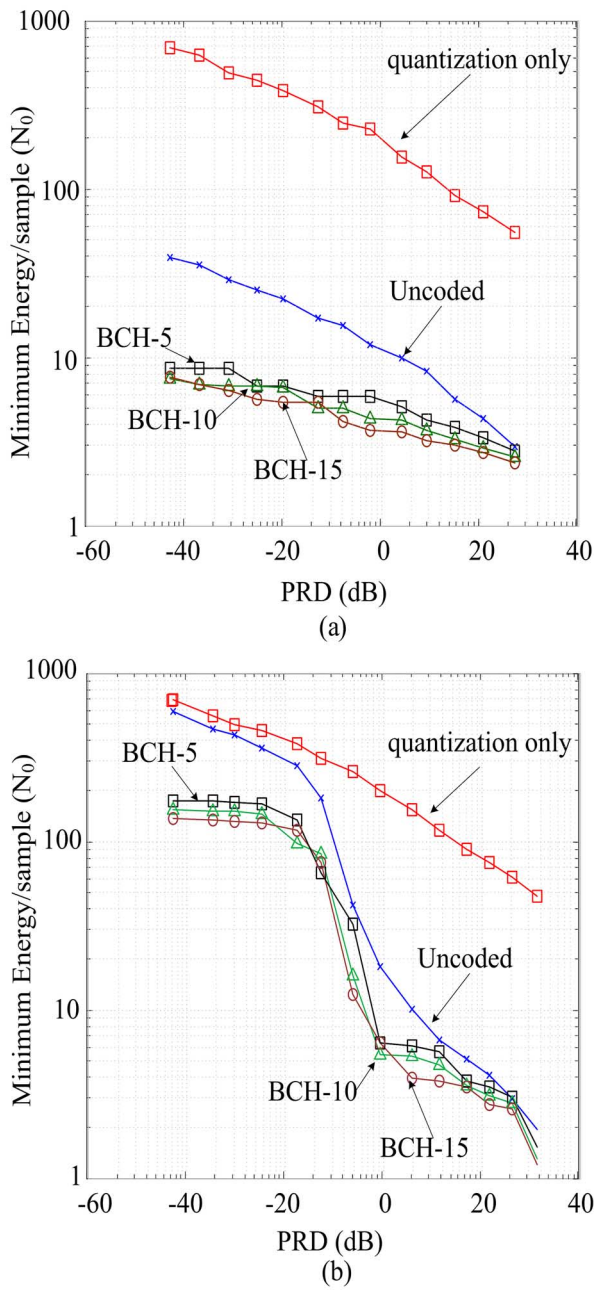

Fig. 14. Minimum energy per sample vs. required PRD performance for coded & uncoded measurements using the BCH coding schemes. Results are shown for signals representative of: (a) $\mathrm{PRD}_{\mathrm{avg}}$ and (b) $\mathrm{PRD}_{\mathrm{net}}$ performance.

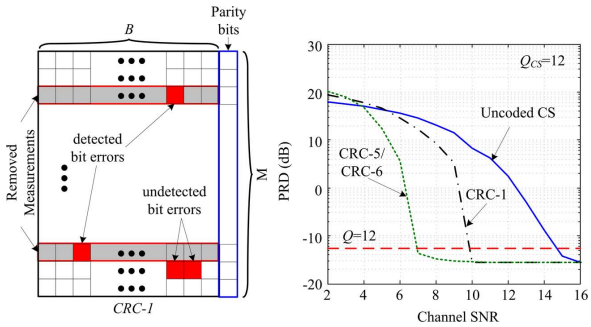

Fig. 15. Example of the proposed CRC-1 based coding scheme and the resulting PRD from applying the codes to the CS measurements when $Q_{\mathrm{CS}} = 12$.

hand, it is well-known that modern codes only perform well (or in other words approach the Shannon-limit) when the block lengths grow very large. As such for modern codes, there exist a fundamental trade-off between the length of the code, and the packetization of the required application, which is beyond the intended scope of this paper. However, in the next subsection we will show that an extremely simple code can provide remarkably similar gains to their more complicated cousins.

2) Cyclic-Redundancy Check (CRC) Coding: CRC codes are popular error detection codes that are easily implementable and provide good error detection [43]. Typical application of error detection schemes (e.g., ARQ) improves data reliability but does not typically improve energy-efficiency since error detection events trigger costly packet retransmissions. However, CS is uniquely suited for error detection schemes since data retransmission may not be required. Recall from Fig. 5 that when we are not measurement limited, decreasing the number of measurements, M, will result in only a soft degradation in reconstruction performance. In the cases where the signal sparsity is lower than expected it is possible that there is no significant performance penalty for small reductions in M. This is attributable to the inherent redundancy of information captured in the compressed measurements.

To take advantage of this characteristic, we propose to use CRC codes to add parity bits *per measurement* (for every B bits) as shown in Fig. 15. At the receiver, when an error is detected, rather than request a retransmission we simply throw those measurements away—i.e., we do not use them for the signal reconstruction. In this scheme, detected channel errors effectively reduce M, which is less destructive than reconstructing the signal with corrupted measurements and more energy-efficient than retransmitting the packet. The results of this coding scheme for CRC-1 (detect odd number of bit errors), CRC-5 (detect ≤ 2 bit errors), and CRC-6 (detect ≤ 3 bit errors) codes are shown in Figs. 15 and 16. Fig. 15 shows the coding gain in regards to PRD for the coded and uncoded cases, while Fig. 16 shows the resulting energy gains for $\mathrm{PRD}_{\mathrm{avg}}$ and $\mathrm{PRD}_{\mathrm{net}}$ signals. From Fig. 16, we see about 4X in energy savings over the uncoded measurements with the proposed error detection scheme. Like the BCH codes, there is a diminishing return in regards to energy-efficiency for higher order error detection. More importantly, we see that error detection schemes, which are less complex and require lower implementation cost, can be just as effective and efficient as error correction schemes.

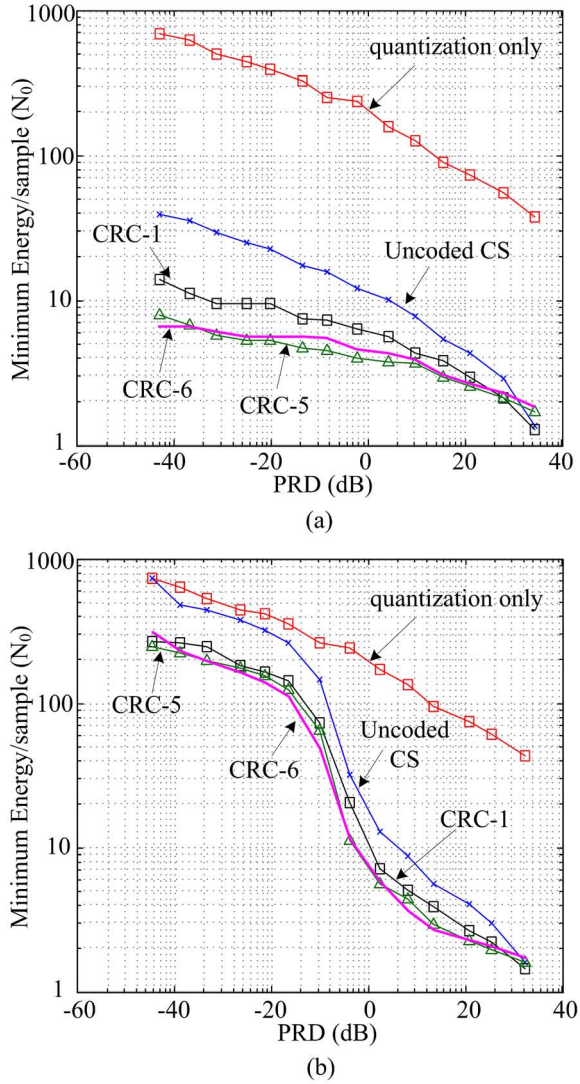

Fig. 16. Minimum energy per sample vs. required PRD performance for coded and uncoded measurements using the CRC error detection codes. Results are shown for signals representative of: (a) $\mathrm{PRD}_{\mathrm{avg}}$ and (b) $\mathrm{PRD}_{\mathrm{net}}$ performance.

C. Diversity Schemes for Burst/Packet Loss

The error correction and error detection schemes discussed previously are only appropriate for random bit errors within a block of measurements. To protect against more catastrophic error events such as longer burst errors and/or packet loss, we borrow ideas from *multiple description coding* [44] and utilize a simple diversity scheme to protect against such error events. Fig. 17 depicts the general idea of the proposed diversity scheme. Here we show a stream of measurement packets, $\boldsymbol{Y}_i$, obtained from data packets, $\boldsymbol{X}_i$, compressed via random sampling matrices $\boldsymbol{\Phi}_{\mathrm{i}}$. In the absence of the proposed protocol, a catostrophically corrupted 3rd packet (Fig. 17) would result in the irrecoverable loss of data $\boldsymbol{X}_3$. On the other hand consider the "2X diversity" protocol shown in Fig. 17. Here, measurements are combined before transmission and the total number of bits transmitted is essentially equivalent to when each measurement is transmitted separately. When there is no packet loss, the individual measurements can be recovered by adding or subtracting two successive packets. In the example shown,

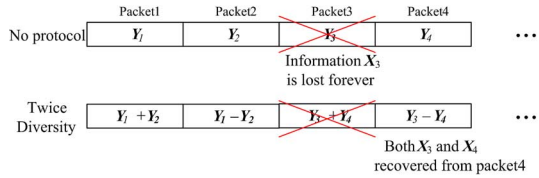

Fig. 17. Diversity protocol for protection against burst errors and dropped packets.

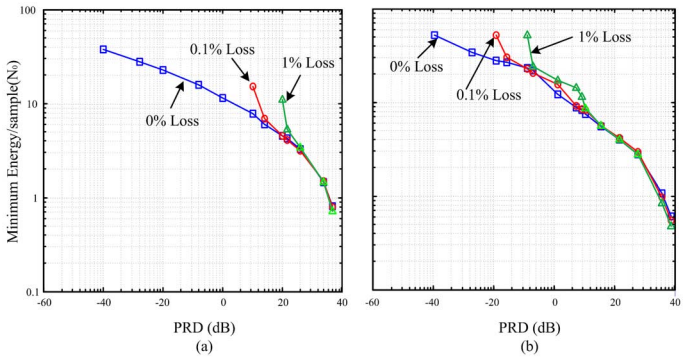

Fig. 18. Minimum energy per sample vs. required PRD performance for 0%, 0.1%, and 1% packet loss for a $\mathrm{PRD}_{\mathrm{avg}}$ signal with: (a) no diversity protocol, and (b) with the diversity protocol.

the loss of the 3rd packet means only the loss of the combination $\boldsymbol{Y}_3 + \boldsymbol{Y}_4$; both data packets $\boldsymbol{X}_3$ and $\boldsymbol{X}_4$ can be jointly recovered from the other received combination, $\boldsymbol{Y}_3$-$\boldsymbol{Y}_4$, The additional complexity of such a scheme requires roughly 2X the compression hardware described in [19]. The performance of the proposed scheme is shown in Fig. 18 for packet loss probabilities of 0.1% and 1%. When a packet is dropped and there is no diversity protocol in place [Fig. 18(a)], the recovered signal quality incurs a big penalty (i.e., PRD is poor). However, when the proposed diversity protocol is added, we can see that we get over 10X in recovered signal performance for roughly 2X the cost. In this example, so long as two consecuitve packets are not dropped, then the signal can be adequately recovered. Additional diversity can be added using this scheme at the cost of reconstruction error and additional hardware.

VI. EKG: An Example CS Application

In this section, we apply the analysis discussed in Section IV on a real electrocardiogram (EKG) signal obtained from the MIT-BIH Arrhythmia database [45]. For simplicity we do not consider error control schemes in this section, but simply keep in mind that the same additional energy gains described in earlier sections can be had.

For this example, we are still comparing the performance of the two system models shown in Fig. 3(b) and (c) only using the recorded EKG signal as the input. The recorded signals are inherently noisy so we do not artificially introduce additional signal noise (n_f). Fig. 19(a) shows a segment of the EKG signal used to conduct the experiment. To reconstruct the signal, we used an overcomplete dictionary of Gaussian pulses similar to the one used for the synthetic signals only with three different pulse widths. Fig. 19(b) shows a sample of the reconstructed signal for when $M = 100$ and $N = 1000$, resulting in a PRD of 0.5%. Additional refinement of the signal basis can be performed to improve the reconstruction error, but the results

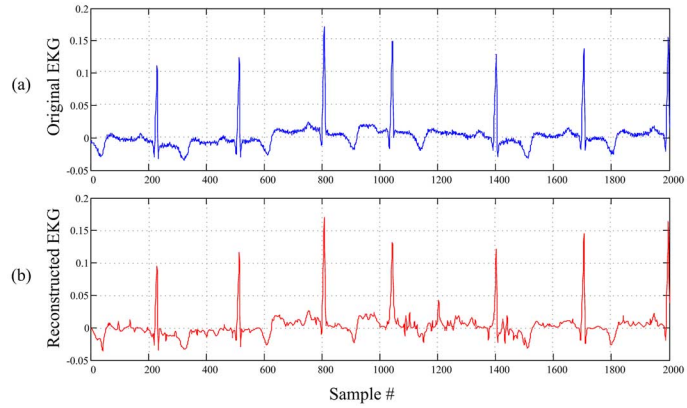

Fig. 19. Original EKG (a) vs. CS reconstruction (b) for $Q = 12$ and $M = 100$.

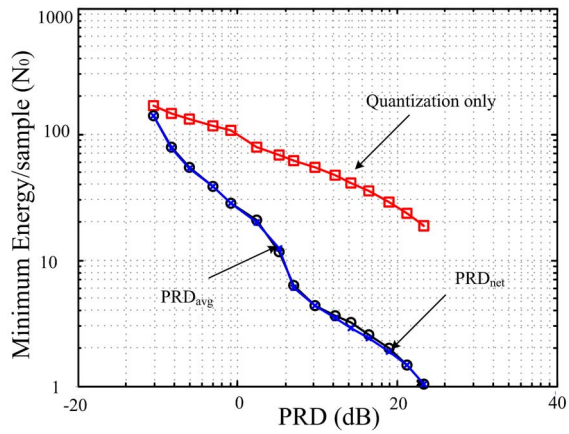

Fig. 20. Minimum energy vs. target PRD for an EKG signal from the MIT-BIH Arrhythmia database.

shown here are on par with the more optimized results described in [46].

The resulting minimum energy vs. target PRD curves (like Figs. 10 and 11) for the EKG signal are plotted in Fig. 20. In general, we see similar results as with the synthetic signals. For target PRDs above 2%, adopting CS provides about 10X in energy gains. For lower PRDs (higher resolution), the performance begins to diminish. In Fig. 20, it can be assumed that the signal is less sparse than the synthetic signals described earlier, so the amount of noise in the signal is large enough such that the $\mathrm{PRD}_{\mathrm{avg}}$ and $\mathrm{PRD}_{\mathrm{net}}$ results are nearly identical (i.e., measurement limited). This is consistent with the noise limited performance predicted by Fig. 11 (where $\mathrm{PRD}_{\mathrm{avg}}$ and $\mathrm{PRD}_{\mathrm{net}}$ converge). It should be noted, however, that our "golden" signal to which we are measuring the reconstruction performance against is a measured signal which contains noise. We typically do not want to reproduce the noise, so if we could fairly determine what the desired signal should be (such that we could compare), we would expect additional performance gains with CS. Despite this limitation, we still see roughly an order of magnitude in energy gains across a range of acceptable PRD values.

VII. Conclusion

In this work we have examined the energy-performance design space for CS in the context of a practical wireless sensor system. We have shown that, even under practical constraints

such as finite resolution and block lengths, CS can be an efficient and robust source encoding/compression algorithm for wireless sensor applications where the signal of interest is sparse. For applications requiring modest resolution performance (<8 bits, PRD$\sim$1%), CS can enable on the order of 10X reduction in transmission energy when compared to raw quantized data. Furthermore, we have shown that CS is robust to channel errors, and is amenable to error control schemes (e.g., CRC error detection) that enable an additional 4X energy reduction and yet have simple hardware realizations. In addition to the system protocols that may be relaxed with greater error resiliency, we have also proposed a diversity scheme for CS that requires no additional transmission costs, provides greater than 10X improvement in recovered signal quality and only requires limited hardware overhead (2X). Finally we show that the design framework and analysis presented is applicable to real world signals, such as EKGs, with a similar order of magnitude reduction in transmission energy costs.

REFERENCES

[1] "OECD, smart sensor networks: Technologies and applications for green growth," Dec. 2009.

[2] I. F. Akyildiz and M. C. Vuran, *Wireless Sensor Networks*. Hoboken, NJ: Wiley, 2010.

[3] N. Verma and A. P. Chandrakasan, "An ultra low energy 12-bit rate-resolution scalable SAR ADC for wireless sensor nodes," *IEEE J. Solid-State Circuits*, vol. 42, no. 6, pp. 1196–1205, Jun. 2007.

[4] A. Agnes, E. Bonizzoni, P. Malcovati, and F. Maloberti, "A 9.4-ENOB 1 V 3.8 μW 100 kS/s SAR ADC with time-domain comparator," in *Proc. Int. Solid-State Circuits Conf.*, 2008, pp. 246–248.

[5] M. Yip and A. P. Chandrakasan, "A resolution-reconfigurable 5-to-10 b 0.4-to-1 V power scalable SAR ADC," in *Proc. Int. Solid-State Circuits Conf.*, 2011, pp. 822–823.

[6] J. L. Bohorquez, J. L. Dawson, and A. P. Chandrakasan, "A 350 μW CMOS MSK transmitter and 400 μW OOK super-regenerative receiver for medical implant communications," in *Proc. Symp. VLSI Circuits*, 2008, pp. 32–33.

[7] S. Rai, J. Holleman, J. Pandey, F. Zhang, and B. Otis, "A 500 μW neural tag with 2 μVrms AFE and frequency-multiplying MICS/ISM FSK transmitter," in *Proc. Int. Solid-State Circuits Conf.*, 2009, pp. 212–214.

[8] B. W. Cook, A. Berny, A. Molnar, S. Lanzisera, and K. S. J. Pister, "Low-power 2.4-GHz transceiver with passive RX front-end and 400-mV supply," *J. Solid-State Circuits*, vol. 41, no. 12, pp. 2757–2766, 2006.

[9] M. W. Phyu, Y. Zheng, B. Zhao, L. Xin, and Y. S. Wang, "A real-time ECG QRS detection ASIC based on wavelet multiscale analysis," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2009, pp. 293–296.

[10] J. Kwong *et al.*, "A 65 nm Sub-Vt microcontroller with integrated SRAM and switched-capacitor DC-DC converter," in *Proc. Int. Solid-State Circuits Conf.*, 2008, pp. 318–319.

[11] R. Yazicioglu, P. Merken, R. Puers, and C. Van Hoof, "A 200 μW eight-channel EEG acquisition ASIC for ambulatory EEG systems," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 3025–3038, 2008.

[12] J. L. Bohorquez, M. Yip, A. P. Chandrakasan, and J. L. Dawson, "A biomedical sensor interface with a sinc filter and interference cancellation," *IEEE J. Solid-State Circuits*, vol. 46, no. 4, pp. 746–756, 2011.

[13] N. Verma, A. Shoeb, J. Bohorquez, J. Dawson, J. Guttag, and A. P. Chandrakasan, "A micro-power EEG acquisition SoC with integrated feature extraction processor for a chronic seizure detection system," *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 804–816, Apr. 2010.

[14] R. R. Harrison *et al.*, "A low-power integrated circuit for a wireless 100-electrode neural recording system," *IEEE J. Solid State Circuits*, vol. 42, no. 1, pp. 123–133, 2007.

[15] A. M. Kamboh, K. G. Oweiss, and A. J. Mason, "Resource constrained VLSI architecture for implantable neural data compression systems," in *Proc. IEEE Int. Symp. Circuits Syst.*, May 2009, pp. 1481–1484.

[16] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[17] E. J. Candes and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.

[18] F. Chen, A. P. Chandrakasan, and V. Stojanović, "A signal-agnostic compressed sensing acquisition system for wireless and implantable sensors," in *Proc. IEEE Custom Integr. Circuits Conf.*, 2010.

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

[20] D. A. Huffman, "Minimum-redundancy coding for the discrete noiseless channel," *IEEE Trans. Inf. Theory*, vol. IT-7, no. 1, pp. 27–38, Jan. 1961.

[21] J. Ziv and A. Lempel, "A universal algorithm for sequential data compression," *IEEE Trans. Inf. Theory*, vol. IT-23, no. 3, pp. 337–343, May 1977.

[22] J. A. Storer and J. H. Reif, "Error resilient optimal data compression," *SIAM J. Comput. (SICOMP)*, vol. 26, no. 4, pp. 934–939, 1997.

[23] S. Aviyente, "Compressed sensing framework for EEG compression," in *Proc. IEEE/SP 14th Workshop Stat. Signal Process.*, Aug. 2007, pp. 181–184.

[24] S.-G. Miaou and S.-N. Chao, "Wavelet-based lossy-to-lossless ECG compression in a unified vector quantization framework," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 3, pp. 539–543, 2005.

[25] M. Fira, L. Goras, C. Barabasa, and N. Cleju, "On ECG compressed sensing using specific overcomplete dictionaries," *Adv. Electr. Comput. Eng.*, vol. 10, no. 4, 2011.

[26] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[27] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, "Least angle regression," *Ann. Stat.*, vol. 32, no. 2, pp. 407–499, 2004.

[28] E. J. Candès, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted L_1 minimization," *J. Fourier Anal. Appl.*, vol. 14, no. 5–6, pp. 877–905, Oct. 2008.

[29] A. Zymnis, S. Boyd, and E. Candes, "Compressed sensing with quantized measurements," *IEEE Signal Process. Lett.*, vol. 17, no. 2, pp. 149–152, Feb. 2010.

[30] L. Jacques, D. K. Hammond, and J. M. Fadili, "Dequantizing compressed sensing when oversampling and non-Gaussian constraints combine," *IEEE Trans. Inf. Theory*, vol. 57, no. 1, pp. 559–571, 2011.

[31] U. Kamilov, V. K. Goyal, and S. Rangan, "Optimal quantization for compressive sensing under message passing reconstruction," in *Proc. IEEE Int. Symp. Inf. Theory*, 2011.

[32] G. Peyre, "Best basis compressed sensing," *IEEE Trans. Signal Process.*, vol. 58, no. 5, pp. 2613–2622, May 2010.

[33] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.

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

[35] S. Rangan, A. K. Fletcher, and V. K. Goyal, "Asymptotic analysis of MAP estimation via the replica method and applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 58, no. 3, pp. 1902–1923, Mar. 2012.

[36] V. K. Goyal, A. K. Fletcher, and S. Rangan, "Compressive sampling and lossy compression," *IEEE Signal Process. Mag.*, vol. 25, pp. 48–56, 2008.

[37] M. Fira and L. Goras, "An ECG signals compression method and its validation using NN's," *IEEE Trans. Biomed. Eng.*, vol. 55, no. 4, pp. 1319–1326, Jun. 2008.

[38] B. Widrow and I. Kollar, *Quantization Noise: Roundoff Error in Digital Computation, Signal Processing, Control and Communications*. Cambridge, U.K.: Cambridge Univ. Press, 2008, p. 778.

[39] W. J. Huang and E. J. McCluskey, "Transient errors and rollback recovery in LZ compression," in *Proc. 2000 Pacific Rim Int. Symp. Dependable Comput.*, pp. 128–135.

[40] R. C. Bose and D. K. Ray-Chaudhuri, "On a class of error correcting binary group codes," *Inf. Control*, vol. 3, no. 1, pp. 68–79, 1960.

[41] C. Berrou, A. Glavieux, and P. Thitimajshima, "Near Shannon limit error-correcting coding and decoding: Turbo-codes," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Geneva, Switzerland, May 1993, vol. 2, pp. 1064–1070.

[42] R. G. Gallager, "Low density parity check codes," in *Monograph*. Cambridge, MA: MIT Press, 1963.

[43] W. W. Peterson and D. T. Brown, "Cyclic codes for error detection," *Proc. IRE*, vol. 49, no. 1, pp. 228–235, 1961.

[44] V. K. Goyal, "Multiple description coding: Compression meets the network," *IEEE Signal Process. Mag.*, vol. 18, no. 5, pp. 74–94, Sep. 2001.
[45] A. L. Goldberger *et al.*, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation Electron. Pages*, vol. 101, pp. e215–e220, 2000.
[46] M. Fira, L. Goras, C. Barabasa, and N. Cleju, "On ECG compressed sensing using specific overcomplete dictionaries," *Adv. Electr. Comput. Eng.*, vol. 10, no. 4, 2011.

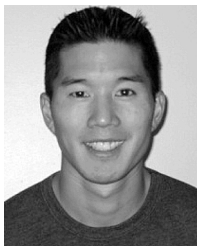

Fred Chen (S'05) received the B.S. degree from the University of Illinois, Urbana-Champaign, in 1997, the M.S. degree from the University of California, Berkeley, in 2000, and the Ph.D. degree from the Massachusetts Institute of Technology, Cambridge, in 2011, all in electrical engineering.

From 2000 to 2005 he was with Rambus Inc., Los Altos, CA, where he worked on the design of high-speed I/O and equalization circuits. He has also previously held a design position at Motorola, Inc., Libertyville, IL. His current research interests include energy-efficient circuits and systems, and circuit design in emerging technologies.

Dr. Chen was a recipient of the 2010 ISSCC Jack Raper Award for Outstanding Technology Directions Paper.

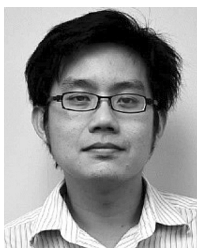

Fabian Lim (S'06–M'10) received the B.Eng. and M.Eng. degrees from the National University of Singapore in 2003 and 2006, respectively, and the Ph.D. degree from the University of Hawaii, Manoa in 2010, all in electrical engineering.

Currently, he is a Postdoctoral Fellow at the Massachusetts Institute of Technology, Cambridge. He has held short-term visiting research positions at Harvard University in 2004 and 2005. From October 2005 to May 2006, he was a Staff Member in the Data Storage Institute in Singapore. From May 2008 to July 2008, he was an intern at Hitachi Global Storage Technologies, San Jose. In March 2009, he was a visitor at the Research Center for Information Security, Japan. His research interests include error-control coding and signal processing, for both communication and storage applications.

Omid Abari (S'08) received the B.Sc. degree in communications engineering with highest honors from Carleton University, Ottawa, ON, Canada, in 2010. He is currently working toward his M.S. degree in electrical engineering and computer science at Massachusetts Institute of Technology, Cambridge, MA.

His research interests include the design of low power, energy-efficient circuits, and systems for wireless communication applications.

Mr. Abari was the recipient of the Merrill Lynch fellowship, the Natural Sciences and Engineering Research Council of Canada (NSERC) Postgraduate scholarship, and Ontario Professional Engineers Foundation for Education Scholarship in 2010.

Anantha P. Chandrakasan received the B.S, M.S., and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994 respectively.

Since September 1994, he has been with the Massachusetts Institute of Technology, Cambridge, where he is currently the Joseph F. and Nancy P. Keithley Professor of Electrical Engineering. His research interests include micro-power digital and mixed-signal integrated circuit design, wireless microsensor system design, portable multimedia devices, energy efficient radios and emerging technologies. He is a coauthor of *Low Power Digital CMOS Design* (Kluwer Academic Publishers, 1995), *Digital Integrated Circuits* (Pearson Prentice-Hall, 2003, 2nd edition), and *Sub-threshold Design for Ultra-Low Power Systems* (Springer 2006). He is also a co-editor of *Low Power CMOS Design* (IEEE Press, 1998), *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000), and *Leakage in Nanometer CMOS Technologies* (Springer, 2005).

Dr. Chandrakasan was a co-recipient of several awards including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, the 1999 DAC Design Contest Award, the 2004 DAC/ISSCC Student Design Contest Award, the 2007 ISSCC Beatrice Winner Award for Editorial Excellence, and the ISSCC Jack Kilby Award for Outstanding Student Paper (2007, 2008, 2009). He received the 2009 Semiconductor Industry Association (SIA) University Researcher Award. He has served as a technical program co-chair for the 1997 International Symposium on Low Power Electronics and Design (ISLPED), VLSI Design'98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Sub-committee Chair for ISSCC 1999–2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, the Technology Directions Sub-committee Chair for ISSCC 2004–2009, and the Conference Chair for ISSCC 2010–2012. He is the Conference Chair for ISSCC 2013. He was an Associate Editor for the IEEE Journal of Solid-State Circuits from 1998 to 2001. He served on SSCS AdCom from 2000 to 2007 and he was the meetings committee chair from 2004 to 2007. He was the Director of the MIT Microsystems Technology Laboratories from 2006 to 2011. Since July 2011, he is the Head of the MIT EECS Department.

Vladimir Stojanovic received his Ph.D. in electrical engineering from Stanford University in 2005, and the Dipl.Ing. degree from the University of Belgrade, Serbia in 1998.

He is the Emanuel E. Landsman Associate Professor of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology, Cambridge. He was also with Rambus Inc., Los Altos, CA, from 2001 through 2004. His research interests include design, modeling and optimization of integrated systems, from CMOS-based VLSI blocks and interfaces to system design with emerging devices like NEM relays and silicon-photonics. He is also interested in design and implementation of energy-efficient electrical and optical networks, and digital communication techniques in high-speed interfaces and high-speed mixed-signal IC design.

Dr. Stojanovic received the 2006 IBM Faculty Partnership Award, and the 2009 NSF CAREER Award as well as the 2008 ICCAD William J. McCalla, 2008 IEEE Transactions on Advanced Packaging, and 2010 ISSCC Jack Raper best paper awards.