

Hardware-Efficient Compressed Sensing Encoder Designs for WBSNs

Jiayi Sheng Chen Yang Martin C. Herbordt
Department of Electrical and Computer Engineering
Boston University, Boston, MA

Abstract—Implanted sensors, as might be used with wireless body sensor networks, must have minimal size and power consumption. In this work we examine digital-based compressed sensing encoders for WBSN-enabled ECG and EEG monitoring, an domain that has received much recent attention. We have two major contributions. The first is using a random Binary Toeplitz matrix rather than Bernoulli. The second is reducing the number of accumulators thereby trading off space for operating frequency. Compared with previous implementations, our new design consumes 1-to-2 orders of magnitude less area and power while still meeting timing constraints and achieving comparable recovery quality.

Keywords—Compressed Sensing; ECG; EEG; Wireless body sensor networks; Toeplitz; hardware-efficient;

I. INTRODUCTION

Wireless body sensor networks (WBSN) has been a hot topic for a number of years. In WBSN-enabled Electrocardiography (ECG) and Electroencephalography (EEG) monitoring systems, the wearable and, especially, the implantable part requires ultra-low-power. Compressed sensing (CS) has been proposed as an effective method that can help meet this constraint [1], [2]. CS is a signal sampling paradigm that enables compression of an input signal with a sampling rate much lower than the Nyquist rate [3]–[5]. It is attractive for WBSN-enabled monitoring systems because the sensor can sample the signal with a reduced number of measurements without losing essential information. Moreover, only this reduced number of measurements needs to be transmitted to the remote telecardiology center where the original signal is reconstructed [2]. A CS encoder is an essential computational part of the sensor embedded in the human body.

We have examined the state-of-the-art digital-based CS encoder designs for WBSN [6]–[8] and found that there is still substantial room for improvement in both power and area reduction. Previous work uses a random Bernoulli matrix as the sensing matrix, but we find that the random binary Toeplitz matrix is often a better candidate. In addition, the ALU part of the previous CS encoders can be optimized further. In particular we find the surprising result that, in this design space, running at a *higher* frequency has substantial

benefit when that is used to reduce chip area and thus leakage current.

In this paper we propose a novel hardware-efficient compressed sensing (CS) encoder architecture that optimizes the power consumption in both aspects, algorithmic and design. The main contributions are as follows:

- 1) We adopt the random Binary Toeplitz matrix as our sensing matrix, which saves a great amount of cost on both sensing and recovery sides. This will be explained further in next section.
- 2) We drastically reduce the area of the ALU in the CS encoder by reducing the number of accumulators. To process an ECG signal, the clock frequency of the CS encoder can be quite low. At this point, leakage power dominates. We find that in this space power consumption benefits substantially from area reduction.
- 3) We further explore and characterize the impact of ALU size on power and correlate design scenario with optimal size.
- 4) ECG and EEG signals have different characteristics and are often captured for different reasons. For EEG signals often only a small fraction of points matter. In our design, we apply a thresholding filter to the input signal to ignore the data points with small amplitude. In this way we further reduce the power consumption of the CS encoder and dramatically reduce the transmission power as well. While thresholding is a well-known method, combining this with the choice of sensing matrix and compute design leads to many trade-offs on both sensing and recovery sides.

To the best of our knowledge, there is no prior work that implements Toeplitz-based CS encoder architecture in hardware. Previous implementations are all Bernoulli-based [6]–[8]. When combined with the other optimizations, our encoder architecture yields 1-to-2 orders of magnitude advantage in both power and area consumption.

The rest of the paper is organized as follows: In the next Section, we review compressed sensing and analyze properties of the binary Toeplitz matrix. In Section III, we describe state-of-the-art Bernoulli-based CS encoder architecture. In Section IV, we present the Toeplitz-based CS encoder architecture. After that come the results where

This work was supported in part by the National Science Foundation through Award #CNS-1405695 and by the Charles Stark Draper Laboratory through the URAD Program Award #SC001-000000834. Email: (jysheng|cyang|herbordt)@bu.edu

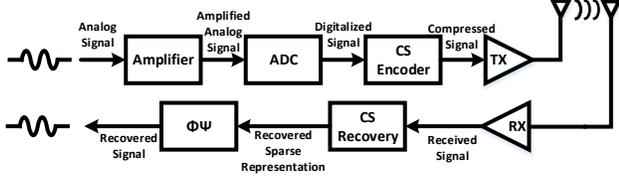


Figure 1. Flow chart of a generic compressed sensing scheme

we first compare the power and area and then verify that the signal from our sensing architecture can be properly recovered. We then describe several sets of experiments and discuss the impact of compression ratio, sparsity, and data set on the quality of the recovered signal.

II. BACKGROUND

A. Compressed Sensing Basics

Traditional sensing and processing rely on the Nyquist-Shannon theorem; this sampling rate, however, can be challenged. In many cases, the signal frequency spectrum is so wide that the sampling rate determined is too large to be satisfied. Even for narrow signal-band application, like WBSN-based ECG and EEG, the Nyquist-rate sampling bring much redundant information since the ECG itself is a sparse signal [2]. Nyquist-rate sampling is expensive for ECG and EEG sensor nodes and impairs the battery lifetime severely. Compressed Sensing has been proposed to address this issue.

The signal flow of compressed sensing is illustrated in Figure 1. After N elements of the analog signal have been sampled, they are sent to the ADC and converted to N digital samples, which can be represented as $X \in R^N$. X is forwarded to the compressed sensing encoder and compressed into M samples. This process can be formulated in Equation 1:

$$Y = \Phi X \quad (1)$$

where $X \in R^N$ is the input signal with N samples, $\Phi \in R^{M \times N}$ is the sensing matrix with M rows and N columns, and $Y \in R^M$ is the compressed signal that has M samples. M is smaller than N .

Compressed sensing theory relies on the prerequisite that the input signal X can be represented by a sparse signal under some basis, which can be formulated in Equation 2.

$$X = \Psi \alpha \quad (2)$$

where $X \in R^N$ is the input signal, $\Psi \in R^{N \times n}$ is the sparsifying basis and $\alpha \in R^n$ is the sparse representation for X . For example, if X is already sparse in the time domain, then Ψ will be just an identity matrix. Or if X is sparse in the frequency domain, then Ψ will be a Fourier transform matrix. Combining Equations 1 and 2, Y can be expressed by Equation 3

$$Y = \Phi \Psi \alpha \quad (3)$$

Besides sparseness, compressed sensing also requires incoherence between the sensing matrix Φ and the sparse representation matrix Ψ in order to minimize the number of measurements (M) needed for the recovery. The coherence of the Φ and Ψ can be evaluated in the following expression [9], [10]:

$$\mu(\Phi, \Psi) = \sqrt{N} \max_{1 \leq k, j \leq N} | \langle \phi_k, \psi_j \rangle | \quad (4)$$

The lower bound on the number of the measurements M can be derived from the value of coherence by the following formula [11]:

$$M \geq C \cdot \mu^2(\Phi \Psi) \cdot S \cdot \log N \quad (5)$$

where C is known to be a small constant around 2 to 2.5 (determined empirically [12]) and N is the number of dimensions of the input signal. S is the sparsity level of the sparse representation, α , which means that the number of non-zero elements of the input signal in the basis Ψ is S . Another constraint that Φ and Ψ must satisfy is the restricted isometry property (RIP) to ensure stable recovery of the S -sparse signal α [13]:

$$(1 - \delta_S) \|\alpha\|_2 \leq \|\Phi \Psi \alpha\|_2 \leq (1 + \delta_S) \|\alpha\|_2 \quad (6)$$

According to [3]–[5], [14], random sensing matrices are very likely to obey the RIP property. Therefore, the sensing matrix can be populated by random variables in various random distributions. Finally Y (with M samples) is transmitted. The recovery algorithm should solve a problem defined in the equation

$$\min_{\alpha \in R^N} \|\alpha\|_{l_1} \text{ subject to } Y = \Phi \Psi \alpha \quad (7)$$

Since M is smaller than the N , the above is an under-determined problem with many possible solutions. It has been proven that, if an l_1 -minimization of α is found, it is very likely to be the closest solution to the original α . After α is recovered, the reconstructed \hat{X} can be acquired by computing $\hat{X} = \Psi \alpha$.

B. Random Binary Toeplitz Matrix

A random sensing matrix is able to meet the RIP requirement so that the signal compressed by it is guaranteed to be properly recovered. Recently, a type of random matrix, the random Toeplitz matrix, has been of interest [1], [15]. An $M \times N$ random Toeplitz is shown here

$$\begin{pmatrix} x_M & x_{M+1} & \cdots & x_{M+N-2} & x_{M+N-1} \\ x_{M-1} & x_M & \cdots & x_{M+N-3} & x_{M+N-2} \\ \vdots & & \ddots & & \vdots \\ x_1 & x_2 & \cdots & x_{N-1} & x_N \end{pmatrix} \quad (8)$$

with entries $\{x_i\}_{i=1}^{M+N-1}$ independent and identically distributed (i.i.d.) random variables that obey a certain distribution.

The Toeplitz matrix has the same value on each diagonal which gives the great advantage that only $M + N - 1$ values need to be stored, rather than $M \times N$. If each entry is random Bernoulli i.i.d., the memory cost to store the whole matrix is reduced still further.¹ In [15], the authors prove that a Toeplitz matrix with i.i.d. entries $\pm 1/\sqrt{k}$ each with probability 1/2 satisfies RIP with high probability.

C. Related Work

Several digital CS encoder hardware designs have been described (see, e.g., [6]–[8], [16], [17]). [16], [17] are both analog-based designs, which are beyond the scope of this paper. Chen, et al. [6] and Suo, et al. [7], [8] both use the random Bernoulli matrix as the sensing matrix. The structure of these two designs is shown in Figure 2.

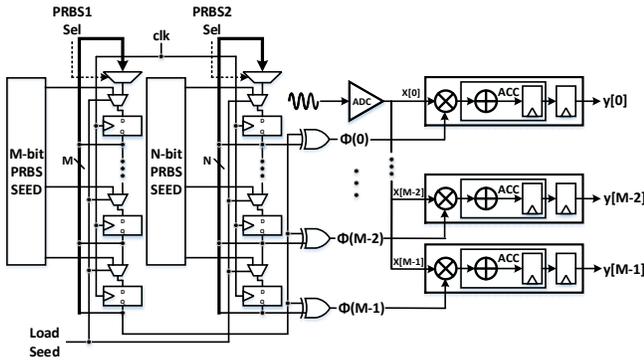


Figure 2. A compressed sensing encoder design based on the random Bernoulli matrix [6]–[8]

We see that this design has two major components. On the left side are two Pseudo-random Binary Sequence (PRBS) generators that generate the random Bernoulli sensing matrix. On the right side is an accumulator array that takes the output of the PRBS generator and multiplies it with the input signal. The accumulator array is basically a matrix-vector multiplication unit that computes the results row by row. Each accumulator unit contains a multiplier and an integrator. The multiplier is simply an XOR gate since each entry of the random Bernoulli matrix is a single bit.

III. IMPLEMENTATION

Our design can be decomposed into two parts, a thresholding filter, which removes the insignificant parts from the original signal and the CS encoder. These are described in the next two Subsections.

¹A second advantage, not used here, is that, while the normal matrix-vector multiplication requires $O(kn)$, the row shifting property of the Toeplitz matrix enables the multiplication between matrix and vector to be done in $O(n \log(n))$ operations by using the FFT. This is because the multiplication between the Toeplitz matrix and a vector is the same as a convolution operation.

A. Thresholding Filter

Full precision EEG data is too detailed for applications like neuron spike sorting. In order to make a correct diagnosis only the information around the neuron spikes needs to be retained. We can therefore apply a thresholding filter to the original signal before it is compressed by the CS encoder. This increases the level of sparsity of the signal improving the compression ratio and energy efficiency (see Equation 5). Thresholding itself is already an effective compression method as it can reduce the original full EEG signal by $5\times$ to $10\times$.

If we define the length of the signal as N and the number of the non-zero elements in the signal as K , then we can evaluate the sparsity level in the equation

$$Sparsity = (1 - K/N) \times (100\%) \quad (9)$$

The thresholding filter sets all the elements smaller than the thresholding value to zero and keeps the other elements unchanged. We can control the sparsity level by adjusting the thresholding value; the higher the thresholding value, the higher sparsity level.

B. Our Compressed Sensing Encoder

Our CS encoder (based on random binary Toeplitz matrix) is illustrated in Figure 3. For an $M \times N$ matrix, the random binary Toeplitz matrix only needs $M + N - 1$ bits. We therefore remove the PRBS generators in the Bernoulli-based implementation and store the sequence with $M + N - 1$ bits instead. When multiplying the Toeplitz matrix and input data vector, our design also takes a different approach. In the Bernoulli-based design, during one clock cycle, all M elements in one column of the sensing matrix are generated simultaneously and then multiplied by one datum from the N -element input vector. In our design only one element per clock cycle is picked out from one row of the sensing matrix to be multiplied by the corresponding element in the input data vector. The temporary product is stored in the accumulator. The accumulator does not output the data until a whole row is computed. Compared with M accumulators in the Bernoulli-based design, our design just needs to instantiate one accumulator. In addition, the multiplier in the accumulator can be implemented with just one XOR gate because the elements in the Toeplitz matrix are also binary values. One drawback in our design is that, dealing with N input data, it takes $M \times N$ clock cycles to finish the computation, while the Bernoulli-based design takes only N cycles. The operating frequency must be adjusted to compensate. We have also parametrized this design to test intermediate points trading off computation for operating frequency.

IV. EVALUATION

In this Section, we evaluate our design in two aspects. In the first Subsection, we compare the power and area

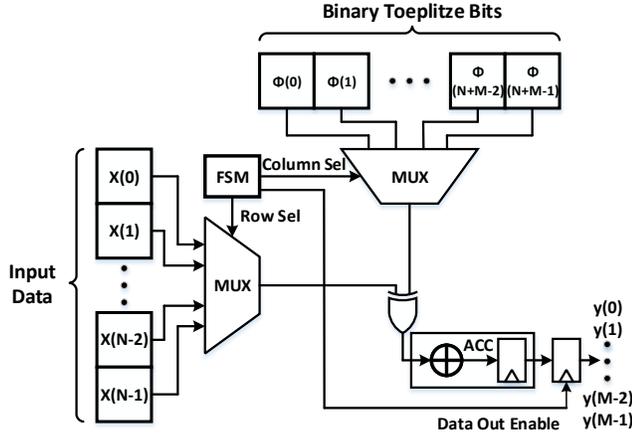


Figure 3. Our hardware-efficient compressed sensing encoder design based on random binary Toeplitz matrix

consumption on two axes: Toeplitz-based versus Bernoulli-based and ALU size versus operating frequency. In the second Subsection we evaluate the recovered signal qualities of the designs. We also extend these results to an initial analysis of an extension from 1D ECG and EEG signals to 2D images.

A. Area and Power

In this Subsection, we present details of our implementations of a Toeplitz-based design and also a reconstruction of the Bernoulli-based designs in [6]. We also examine the use of various numbers of accumulators. Figure 4 shows a comparison of the best new design (Toeplitz and 1 ALU) with the best previous design (Bernoulli and M ALUs). We also explore the design choices that lie in the between: Table I and Table II show area and power consumption of the Toeplitz- and Bernoulli-based design whose ACC numbers range from one to M . Following previous studies, we fix the word width of the input data at 8 bits and the accumulator at 16. We synthesize and place&route both designs using the Cadence Encounter RTL Compiler [18] using a 45nm technology.

Besides the designs in Table I, we also implemented a Bernoulli-based design that has same size settings as the design in [6] where the length of the signal is 1024 and the compression ratio is 0.05. In [6], the authors use a 90 nm technology, therefore the area of the our Bernoulli-based design is expected to be 1/4 of the area in [6]. In fact, the area of Bernoulli-based design is 12926 μm^2 , which is smaller than expected (see [6]). While this is not validation that we have reproduced their design exactly, it does appear that our reconstruction is at least comparable.

From the Figure 4 we note that the Toeplitz-Single-ALU design consumes at least one order of magnitude less area and power than the Bernoulli- M -ALU design. Moreover, when the matrix size increases, the area and power of

Toeplitz-based remains constant, while that of the Bernoulli-based increases linearly with the matrix size. There are two advantages in our design. The first is that the PRBS generator used to generate the random Bernoulli matrix is not needed; rather the random binary Toeplitz matrix is stored, which costs negligible chip space. The second is that we dilate the number of cycles of the CS encoder. In our design, it takes $M \times N$ cycles to compress a vector with length N into a vector with length M , while in the Bernoulli-based design only N cycles are needed. Therefore, our design only instantiates one accumulator while the Bernoulli-based design instantiates M accumulators to reduce a great amount of area. Of course this also changes the operating frequency necessary to meet the processing deadline. In Figure 4, the Toeplitz-based design runs with 1MHz clock while the Bernoulli-based design runs at 20KHz clock, the same as the frequency in [6]. This only marginally diminishes the advantage of the small number of accumulators.

In Table I shows the area consumption of Toeplitz- and Bernoulli-based designs with various numbers of ACCs. Both of them scale linearly with the number of ACCs. But the Toeplitz-based design is always smaller than the Bernoulli-based design when the number of ACCs and CR are the same. In ECG recovery (see next Section), the Toeplitz-based design requires a higher CR than the Bernoulli-based design to achieve the same recovery quality.

In Table II, each row shows the power consumption of Toeplitz- and Bernoulli-based designs with a different number of ACCs and operating frequency adjusted accordingly. In the Table II, the input data rate is fixed to a typical ECG data rate, 1 KHz [1], [19], [20]. Under this condition, we can find that in both designs the leakage power dominates the total power consumption ($> 90\%$). Since the leakage power is mostly related to chip area, area reduction is the most effective approach to reduce the power consumption, which is the exactly our main idea in the Toeplitz-based optimization.

In the Table II, the optimal number of ACCs is 1. However this is not always true as when the input data rate increases. In Figure 5, the lines of different colors denote the power consumption of two designs under different input data rates. We find that the optimal number of ACCs shifts from 1 to M as the input data rate increases. This is because the dominating power component transit from leakage power to switching power, which increases linearly with the clock frequency.

B. Reconstruction Quality

In this Subsection, we validate that ECG, EEG, and image signals compressed by a random binary Toeplitz sensing matrix can be reconstructed with the CS recovery algorithm with quality similar to that of the random Bernoulli method. We use Signal-to-Noise-and-Distortion Ratio (SNDR) to

Table I

AREA CONSUMPTION OF TOEPLITZ- AND BERNOULLI-BASED DESIGNS AS A FUNCTION OF NUMBER OF ACCUMULATORS. SIGNAL LENGTH = 1024. COMPRESSION RATIO=0.2 OR 0.4. THE NUMBER LABELED WITH * IS THE TOEPLITZ-BASED DESIGN WITH JUST ONE ACC. THE NUMBER LABELED WITH ** IS THE BERNOULLI-BASED DESIGN WITH M ACCS.

ACC	Matrix Length (N)	Area								
		Gates			Cells			Area (μm^2)		
		Toeplitz	Bernoulli		Toeplitz	Bernoulli		Toeplitz	Bernoulli	
			CR=0.4	CR=0.2		CR=0.4	CR=0.2		CR=0.4	CR=0.2
1	1024	4006*	9715	16658	1665*	3996	7162	3347*	7982	13293
51	1024	23641	28535	35431	10285	12159	21716	18923	22626	28275
102	1024	41513	46437	53054	17719	20080	35824	32935	37107	42338
154	1024	53792	59447	66144	23434	25722	44842	43129	47144	52782
205	1024	72249	76499**	83959	31537	33306**	59018	57510	60935**	67000
256	1024	82390	NA	94200	35489	NA	65484	65969	NA	75172
307	1024	89197	NA	107431	38887	NA	74067	71299	NA	85730
358	1024	108610	NA	120966	46918	NA	83314	86824	NA	96531
410	1024	125405	NA	137731**	54001	NA	92160**	100192	NA	109910**

Table II

POWER CONSUMPTION OF TOEPLITZ- AND BERNOULLI-BASED DESIGNS AS A FUNCTION OF NUMBER OF ACCS. SIGNAL LENGTH = 1024. THE NUMBER LABELED WITH * IS THE TOEPLITZ-BASED DESIGN WITH JUST ONE ACC. THE NUMBER LABELED WITH ** IS THE BERNOULLI-BASED DESIGN WITH M ACCS.

ACC	Power (μW)											
	Internal			Switching			Leakage			Total		
	Toeplitz	Bernoulli		Toeplitz	Bernoulli		Toeplitz	Bernoulli		Toeplitz	Bernoulli	
		CR=0.4	CR=0.2		CR=0.4	CR=0.2		CR=0.4	CR=0.4		CR=0.2	CR=0.4
1	2.66*	13.1	30.3	1.62*	4.28	10.8	31.1*	108	138	35.4*	126	179
102	0.491	0.277	0.609	0.928	0.192	0.447	336	341	400	338	348	401
205	0.446	0.200**	0.463	0.824	0.0154**	0.412	549	528**	617	550	528**	618
307	0.382	NA	0.448	0.596	NA	0.399	659	NA	787	660	NA	787
409	0.4834	NA	0.434**	0.975	NA	0.476**	987	NA	1054**	989	NA	1055**

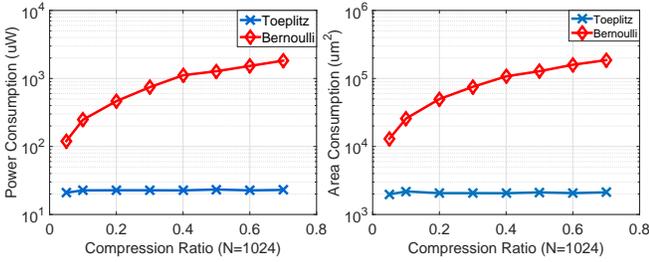


Figure 4. Comparison of power and area for Bernoulli- and Toeplitz-based designs, (The Toeplitz-based design in this figure has just one ACC while the Bernoulli-based design has M ACCs.)

quantify the recovered signal quality:

$$SNDR = 20 \log \frac{\|x\|_2}{\|x - \hat{x}\|_2} \quad (10)$$

where x is the original signal and \hat{x} is the recovered signal.

1) *Reconstruction of EEG signals:* The EEG signals used in this paper are all from the PhysioBank [19]. The reconstruction algorithm is Orthogonal Matching Pursuit (OMP) [12]. The results in Figure 6 show the relationship between SNDR and sparsity (thresholding level) for various compression ratios. Reconstruction of both Toeplitz- and Bernoulli-based compression is demonstrated. Each curve

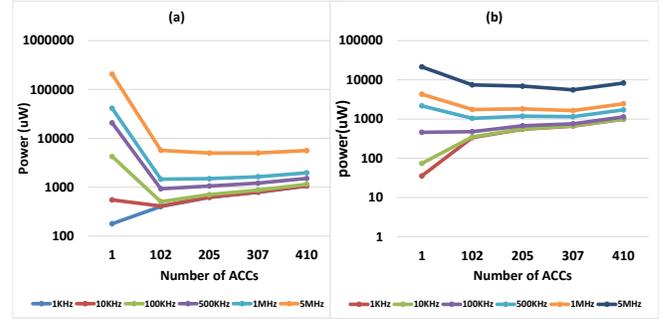


Figure 5. The power consumption for different numbers of ACCs when using different data input rate

corresponds to a different compression ratio. We find that for each compression ratio there is always a sparsity level for which the reconstruction quality is excellent. The higher the compression ratio, the lower the sparsity is required; this means that improving the sparsity level of the signal can reduce the compression ratio required. The compression ratio does not affect the cost of our Toeplitz-based CS encoder, but a low compression ratio reduces the cost of transmission and recovery side processing. For the same compression ratio, the recovery quality of the random Toeplitz matrix

is almost the same as that of the random Bernoulli matrix. In some cases, the Toeplitz is even better than the Bernoulli. Several recovery examples for different compression ratios are shown in the Figure 7. As expected, the higher the sparsity, the better the compression ratio that can be safely used. With CR higher than 0.4, the Toeplitz matrix can compress an elaborate EEG waveform; we therefore have chosen this CR in Tables I and II.

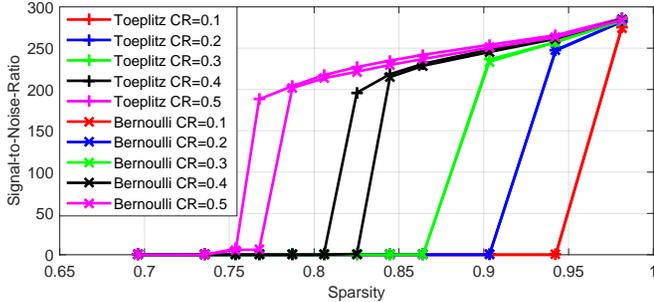


Figure 6. Relationship between SNDR and sparsity for various compression ratios

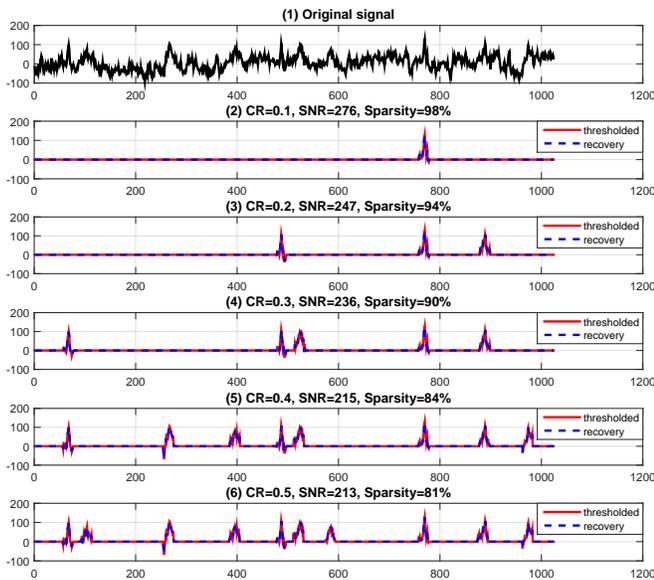


Figure 7. Toeplitz-based EEG recovery examples with different CRs and different sparsity levels (different thresholding level). (1) Original EEG signal (2) CR=0.1, Sparsity=98% (3) CR=0.2 Sparsity=94% (4) CR=0.3 Sparsity=90% (5) CR=0.4 Sparsity=84%. (6) CR=0.5 Sparsity=81%

2) *Reconstruction of ECG signals:* The ECG signals used in this paper are also from the PhysioBank [19]. The reconstruction algorithm is again Orthogonal Matching Pursuit (OMP) [12]. When sparsifying of the ECG signal, it is likely that thresholding-based sparsifying gets rid of too much information. Rather, we use a sparsify dictionary method called K-SVD [21].

The results in the Figure 8 show the relationship between reconstruction SNDR and compression ratio for Bernoulli-

and Toeplitz-based compression. Each point in this figure is a representation of 25 runs; the variation for each point is displayed by an error bar. The best and worst cases of the 25 runs are also displayed. The results show that under the same compression ratio, the Toeplitz-based compression is not as good as the Bernoulli-based compression when using dictionary learning to do the sparsification. To achieve the same recovery quality, the Toeplitz-based compression needs to maintain almost twice CR as Bernoulli-based compression. However, from the table I and II, we can see that even the CR is doubled, our Toeplitz-based design with one accumulator still has great advantage over previous multi-ACC Bernoulli-based design. One reconstruction example of Toeplitz-based compression is shown in Figure 9

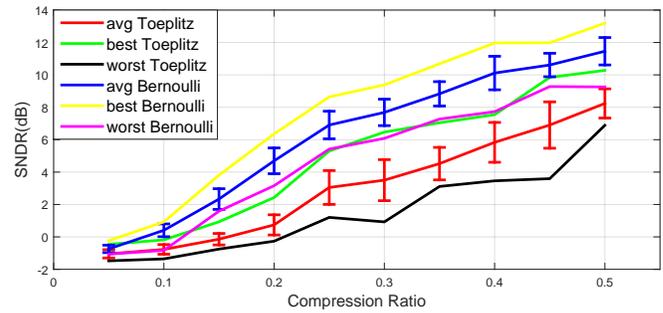


Figure 8. Relationship between SNDR of ECG recovery and compression ratio for random Bernoulli- and Toeplitz-based compression

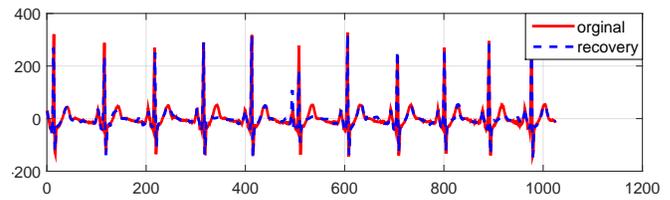


Figure 9. Toeplitz-based EEG recovery example with CR=0.4, SNR=8.77dB

3) *Reconstruction of Images:* When using images, thresholding is again not an appropriate method to sparsify the input data. Instead, we adopt the 2D wavelet transform as the sparse representation matrix. For image input, we use the Treemp2D algorithm [22]. Figure 10 compares the recovery quality of the Toeplitz- and Bernoulli-based compression. The image here is 128×128 . We find that when CR is 0.5, the Toeplitz-based compressed image can be properly reconstructed. When the CR goes to 0.3, the recovered image is blurred. The random Bernoulli matrix performs somewhat better. Tables I and II demonstrate that the power and area of the Bernoulli-based CS encoder scales linearly with the matrix size, while the Toeplitz-based design remains constant. When using the image as input signal, the signal length (N) is much larger than the ECG data, and the sensing matrix size is quadratically related with N . The physical

advantage of the Toeplitz-based design is therefore greater for image inputs. This is likely to allow use of a higher CR to achieve better reconstruction while still using only a fraction of the resources.

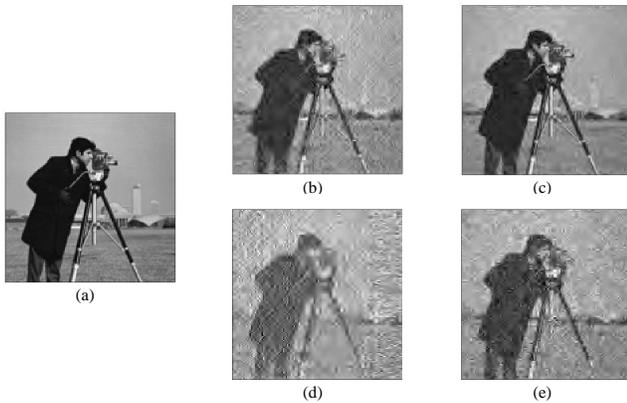


Figure 10. Comparison between Toeplitz- and Bernoulli-based recovered images (a) original image 128×128 (b) Toeplitz-based recovery CR = 0.5 SNDR = 17.11dB (c) Bernoulli-based recovery CR = 0.5 SNDR = 21.04dB (d) Toeplitz-based recovery CR = 0.3 SNDR = 11.50dB (e) Bernoulli-based recovery CR = 0.3 SNDR = 15.18dB

V. CONCLUSION

In this paper, we propose a novel hardware-efficient compressed sensing encoder design based on the random binary Toeplitz matrix and use of smaller, higher frequency compute units. We estimate the power and area consumption of our design using the Cadence Encounter RTL Compiler. The power and area consumption of our design shows great advantages over previous CS encoder designs based on the random Bernoulli matrix. We verify that the EEG, ECG and image signals compressed by random binary Toeplitz matrix are both recoverable. For any CR the recovery quality is not as good as Bernoulli-based recovery. But the gap is small—and in the case of EEG negligible—with respect to the benefit to power and area.

REFERENCES

- [1] A. Dixon, E. Allstot, D. Gangopadhyay, and D. Allstot, "Compressed Sensing System Considerations for ECG and EMG Wireless Biosensors," *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 6, no. 2, pp. 156–166, 2012.
- [2] 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, 2011.
- [3] D. L. Donoho, "Compressed Sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [4] E. J. Candes and T. Tao, "Decoding by Linear Programming," *Information Theory, IEEE Transactions on*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [5] —, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *Information Theory, IEEE Transactions on*, vol. 52, no. 2, pp. 489–509, 2006.
- [6] F. Chen, A. Chandrakasan, V. Stojanovic, and B. Rao, "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, 2012.
- [7] J. Zhang, Y. Suo, S. Mitra, S. Chin, S. Hsiao, R. Yazicioglu, T. Tran, and R. Cummings, "An Efficient and Compact Compressed Sensing Microsystem for Implantable Neural Recordings," *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 8, no. 4, pp. 485–496, 2014.
- [8] Y. Suo, J. Zhang, T. Xiong, S. Chin, R. Cummings, and T. Tran, "Energy-Efficient Multi-Mode Compressed Sensing System for Implantable Neural Recordings," *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 8, no. 5, pp. 648–659, 2014.
- [9] J. A. Tropp, "Just Relax: Convex Programming Methods for Identifying Sparse Signals in Noise," *Information Theory, IEEE Transactions on*, vol. 52, no. 3, pp. 1030–1050, 2006.
- [10] D. Donoho and X. Huo, "Uncertainty Principles and Ideal Atomic Decomposition," *Information Theory, IEEE Transactions on*, vol. 47, no. 7, pp. 2845–2862, 2001.
- [11] E. Candes and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Problems*, vol. 23, no. 3, pp. 969–985, 2007.
- [12] J. A. Tropp and A. C. Gilbert, "Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit," *Information Theory, IEEE Transactions on*, vol. 53, no. 12, pp. 4655–4666, 2007.
- [13] E. Candes, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [14] E. Candess, "An introduction to compressive sensing," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 21–30, 2008.
- [15] J. Haupt, W. Bajwa, G. Raz, and R. Nowak, "Toeplitz Compressed Sensing Matrices With Applications to Sparse Channel Estimation," *Information Theory, IEEE Transactions on*, vol. 56, no. 11, pp. 5862–5874, 2010.
- [16] T. Ragheb, J. N. Laska, H. Nejati, S. Kirolos, R. Baraniuk, and Y. Massoud, "A prototype hardware for random demodulation based compressive analog-to-digital conversion," in *Proc. IEEE Midwest Symp. Circuits and Systems*, 2008, pp. 37–40.
- [17] M. Mishali, Y. Eldar, O. Dounaevsky, and E. Shoshan, "Xampling: Analog to digital at Sub-Nyquist rates," *IET Circuits, Devices Syst.*, vol. 5, pp. 8–20, 2011.
- [18] (2013) Encounter Power System Datasheet. Cadence.
- [19] (2015) PhysioBank Archives. National Institute of Biomedical and Imaging and Bioengineering & National Institute of General Medical Sciences. [Online]. Available: <http://www.physionet.org/physiobank/>
- [20] L. Hejfel and E. Roth, "What is the adequate sampling interval of the ECG signal for heart rate variability analysis in the time domain?" *Physiological Measurement*, vol. 25, no. 6, pp. 1405–1411, 2004.
- [21] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *Signal Processing, IEEE Transactions on*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [22] R. G. Baraniuk, V. Cevher, and M. F. Duarte, "Model-Based Compressive Sensing," *Information Theory, IEEE Transactions on*, vol. 56, no. 4, pp. 1982–2001, 2010.