

Effects of Non-Uniform Quantization on ECG acquired using Compressed Sensing

Darren Craven, Brian McGinley, Liam Kilmartin, Martin Glavin, Edward Jones
College of Engineering and Informatics
National University of Ireland
Galway, Ireland
{d.craven1,brian.mcginley,liam.kilmartin,martin.glavin,edward.jones}@nuigalway.ie

Abstract—This paper analyzes the effects of quantization on Compressed Sensing (CS) measurements applied to Electrocardiogram (ECG) signals. Two methods of quantization are proposed in this paper: uniform and non-uniform. Reconstruction is performed using a dictionary based on the Mexican Hat wavelet. A distortion-based performance metric Percent Root-mean-squared Difference (PRD) will be monitored at various Compression Ratios (CR) to quantify the impact of quantization. The energy cost of transmission is also evaluated for different levels of quantization and compared, at certain PRD levels. The results demonstrate that non-uniform quantization outperforms the uniform approach and that employing non-uniform quantization improves implementation efficiency for applications with acceptable PRDs above 6.75%. Results show that utilizing non-uniform quantization can increase the CR from 9.8 to 14.1 for a PRD of 30%. Furthermore, this amounts to a 28.91% reduction in wireless transmission per frame from 37.7 μJ to 26.8 μJ considering Bluetooth Low Energy (BLE) as a target wireless communication protocol.

Keywords- ambulatory monitoring; biomedical signal compression; compressed sensing (CS); electrocardiogram (ECG); quantization.

I. INTRODUCTION

The deployment of wireless Body Area Networks (BAN) are transforming modern healthcare services [1]. A BAN typically consists of a biological sensor equipped to capture real-time bioelectric data from a patient. This data is then usually transmitted wirelessly to a smartphone, or computer platform, before transmission to a healthcare provider. These architectures present numerous benefits such as constant patient monitoring, lower healthcare costs and increased patient mobility. The suitability of a BAN hinges on the prerequisite that it can provide an energy efficient implementation. The most significant challenge associated with sustaining a low power implementation is the wireless transmission of the data to the back-end server system accessed by clinicians [2]. Strategies that attempt to minimize the amount of data generated are generally utilized.

This research will focus on the use of a recently introduced lossy compression technique known as Compressed Sensing (CS) [3, 4]. CS provides a low complexity signal acquisition paradigm by simultaneously sampling and compressing. While signal acquisition is highly energy efficient with CS, the

reconstruction of the signal is computationally intensive. In the context of a BAN, reconstruction can be performed outside the wireless environment hence resulting in no significant power implications for the body worn components. In fact, studies have shown CS to be a competitive alternative to state-of-the-art wavelet compression algorithms in terms of overall energy efficiency in a BAN [2]. This is based on the ability of CS to provide a relatively simple acquisition process and the removal of the need for digital compression of the raw sampled signal. Despite this efficiency, CS does not generally provide comparable results to these algorithms in terms of signal reconstruction quality [5].

The power efficiency of CS acquisition is due to the fact that CS acquires fewer measurements than Nyquist rate compression methods. However, to date there has been no investigation of the effects of quantizing these measurements in electrocardiogram (ECG) signal applications. This paper focuses on maximizing the compression gain which can be realized from such quantization. An energy-based distortion metric, namely Percent Root-mean-squared Difference (PRD), has been used to objectively quantify the degradation in signal quality in terms of overall compression gain. This paper additionally exploits the Gaussian distribution nature of CS measurements by implementing a form of non-uniform quantization to increase the Compression Ratio (CR).

II. ECG COMPRESSION

A. Database

The ECG database used in this paper is the MIT-BIH Arrhythmia Database [6]. The database consists of 48 thirty-minute patient records sampled at 360 Hz. The recordings were sampled with 11-bit resolution. For the purposes of this research the records labeled as 100, 107, 115 and 117, which were used in [7], were used and each record was divided into frames of 1024 non-overlapping samples.

B. Performance Metrics

1) *Percent Root-mean-squared Difference (PRD)*: PRD is a measure of the distortion between the original and reconstructed signals.

$$PRD (\%) = \frac{\|X - X'\|}{\|X - \bar{X}\|} \quad (100) \quad (1)$$

where X is the original signal, X' is the reconstructed signal and \bar{X} is the mean of the signal.

2) *Compression Ratio (CR)*: CR quantifies the reduction in bits required to represent the the original signal divided by the bits needed for the compressed measurements.

$$CR = \frac{(N)(B)}{(M)(q)} \quad (2)$$

where N is the frame size of the original signals (1024), B is the bit resolution of the sampled signal (11 bits with the MIT-BIH Arrhythmia Database), M is the number of measurements sampled and q is the bit resolution of the measurements, post quantization.

3) *Energy Transmission per Frame (E_{TF})*: Higgins *et al.* [8] reported that a BLE transmitter dissipates approximately 33 nJ per transmitted bit over a BLE radio link. This value will be used to quantify the energy implications of quantization in terms of wireless transmission.

$$E_{TF} = (M)(q)(33\text{nJ}) \quad (3)$$

where E_{TF} is the energy transmitted per frame, M is the number of measurements sampled and q is the bit resolution of the measurements.

C. Compressed Sensing

1) Method:

a) *Acquisition*: The acquisition stage of CS is central to its low power operation. Studies have shown that hardware designed to implement CS acquisition can consume less power than an Analog-to-Digital converter (ADC) sampling at the Nyquist rate [9]. The acquisition comprises a relatively simple linear matrix multiplication. The input signal X of N samples is multiplied by a sensing matrix Φ of random entries. Thus the M compressed measurements (Y) are computed as follows.

$$[Y]_{M,1} = [\Phi]_{M,N} [X]_{N,1} \quad (4)$$

b) *Reconstruction*: Successful reconstruction in CS is based on two fundamental concepts; sparsity and incoherence. For typical real-world signals, a sparse vector α exists, which has a small number of non-zero entries (sparse) when represented in the dictionary Ψ . Secondly, the sensing matrix Φ and the dictionary Ψ must be incoherent. Equation (5) describes how X is expressed sparsely in a particular dictionary. Note the sparse dictionary is not necessarily a square matrix and can contain P columns instead of N to create an overcomplete dictionary. The reconstruction process aims to solve for α using convex optimization techniques.

$$[Y]_{M,1} = [\Phi]_{M,N} [\Psi]_{N,P} [\alpha]_{P,1} \quad (5)$$

Solving for α is performed using L1-norm minimization and the Basis Pursuit algorithm [10] is used for this task.

$$\text{Min } \|\alpha\|_1 \text{ subject to } Y = \Phi\Psi\alpha \quad (6)$$

The reconstructed signal X' can then be recovered by multiplying α and Ψ . The dictionary chosen to offer sparse representations of the ECG signals is based on the Mexican Hat wavelet. This wavelet was selected due to the reported performance in [5].

III. QUANTIZATION

Quantization is a simple compression operation that maps a sequence of numbers $Y(N)$ to a smaller set of quantized values $Y_Q(N)$. Quantization occurs as signals are digitized and provides an estimate of the signal by mapping each sample to a codeword based on its voltage. The decoder then reconstructs the signal $Y(N)$ based on this quantized value. This process introduces an error into the signal as the original sample can only be estimated and is not guaranteed to be correctly reconstructed. Generally, the accuracy of this decoding depends on the bit resolution being used.

In terms of CS, quantization of the compressed measurements will be the focus of this study in order to reduce the number of bits being transmitted wirelessly in an ambulatory monitoring architecture. Fig. 1 shows a plot of the Probability Density Function (PDF) curve for the measurement values for patient record 100 resulting from 1000 sensing matrices. In order to try and leverage the Gaussian distribution nature of the measurement values, non-uniform quantization has been used, applying higher resolution in the regions where the measured values occur more commonly. The performance of standard uniform quantization has also been analyzed.




Figure 1. PDF curve of the measurement values for measurement size of $M=512$. In total, the results are averaged from 634,000 (634 frames in record x 1000 iterations) measurement matrices for Record 100.

A. Non-Uniform Quantization

Non-uniform quantization aims to exploit the Gaussian distribution of the PDF curves (Fig. 1). For the purposes of these experiments, two regions have defined: the *high resolution region* and a *low resolution region*. Different step

sizes have been applied in the different resolution regions. A limit value will define the maximum values for the quantizer and to simplify implementation the same limit will be used for both positive and negative values. The boundary between the high and low resolution regions was set at half of this limit value.

Preliminary testing on the percentage of codewords assigned in the high resolution region and suitable limiting values are shown in Table I. The limit values correspond to a particular percentage of area underneath the curve. For example a value of 99.99% corresponds to intervals at the point where 99.99% of the measurements have occurred (in Fig. 1, this interval would use $[-47,47]$ as the limits). Table I analyzes the PRD between the measurement values with uniform quantization and when the compressed measurements have been quantized using different limit values and different resolutions in the high and low resolution regions.

TABLE I AVERAGE PRD OF THE COMPRESSED MEASUREMENTS FOR UNIFORM AND NON-UNIFORM QUANTIZATION METHODS AT 6 BITS RESOLUTION

High Resolution Percentage	Limit Values			
	99%	99.9%	99.99%	99.999%
65%	4.59	2.95	3.41	4.21
70%	4.63	2.91	3.25	3.95
75%	4.95	2.99	3.13	3.72
80%	8.41	3.77	3.17	3.56
Uniform Quantization	4.78	3.52	4.33	5.42

The optimal quantization scenario from Table I in terms of PRD was then used in further testing. The limits coincided with 99.9% of the area underneath each records' PDF curve and the high resolution region has more concentrated intervals at a ratio of 7:3 (70%) to the low resolution region. As a result the non-uniform step sizes for each region were defined as in (7) and (8).

$$\Delta_{hr} \approx \frac{L}{(2^q)(0.7)} \quad (7)$$

$$\Delta_{lr} \approx \frac{L}{(2^q)(0.3)} \quad (8)$$

where Δ_{hr} and Δ_{lr} are the high and low resolution step sizes respectively, L is the limit measurement value, q is the bit resolution, 0.7 and 0.3 are indicative of the ratio between the regions. Note these equations are presented as approximations as $(2^q)(0.7)$ and $(2^q)(0.3)$ are rounded to the nearest integer value if required.

B. Uniform Quantization

Uniform quantization operates by having an equal step size for all intervals inside the defined limits. The step size Δ , is dependent on the bit resolution and the record specific limiting value. The step size is calculated by (9).

$$\Delta = \frac{(2)(L)}{2^q} \quad (9)$$

where Δ is the step size, L is the limit measurement value for a patient record and q is the bit resolution.

IV. EXPERIMENTAL RESULTS

A. PRD

The complete records were compressed and reconstructed using 6, 7, 8 and 9 bit quantization for the non-uniform approach described previously. The same records were also tested without quantization and with 7 bit quantization for the uniform method.




Figure 2. Absolute error shown for two examples of measurements for non-uniform (red) and uniform (black) (a) 6 bit (b) 7 bit quantization.

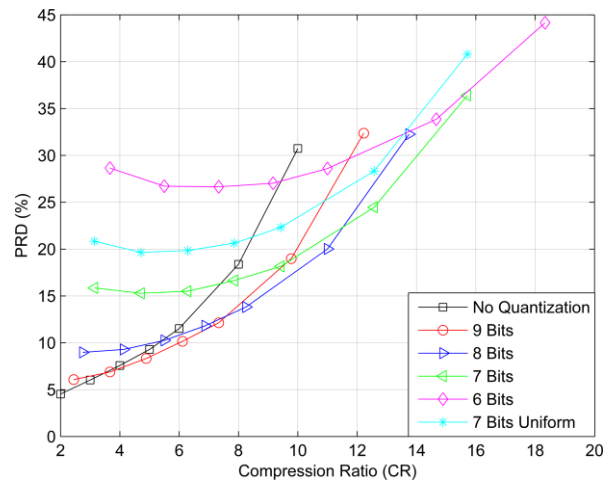


Figure 3. The average PRDs from each record at varying quantization rates. Results using 6,7,8 and 9 bit quantization and with no quantization are shown.

Fig. 2 shows single frame examples of the absolute error with uniform and non-uniform quantization approaches and highlights the benefits of non-uniform quantization. The non-uniform approach has a smaller error for the majority of entries due to the implementation of the *high resolution* region. The occasional error peak from the non-uniform approach is due to

values falling outside the high resolution region where the uniform approach has a smaller step size.

The PRD values are averaged over all the records in Fig. 3. Quantization shows improved compression gains particularly at higher PRD levels. At PRDs below 6.75% there is no benefit accrued for implementing quantization. However, quantization can maximize the compression as acceptable PRD values increase. For this test scenario, quantizing measurements with 7 bits or fewer removes the capacity for achieving PRD values < 9%, even at lower CRs. An increase in CR can be expected at all PRDs above 6.75%. At a PRD of 20%, CR can be increased from 8.3 to 11 and a PRD of 25% allows an extension in CR from 9.1 to 12.7. At a PRD of 30%, each quantization method tested outperforms the no quantization approach.




Figure 4. Two reconstructed frames of record 115 at the varying quantization rates. Original signal (red) and reconstructed signals (black) are shown for each plot. (a) 9 bits (CR = 12.2, PRD = 28.1%). (b) 8 bits (CR = 13.8, PRD = 34.4%). (c) 7 bits (CR = 15.7, PRD = 39.8%). (d) 6 bits (CR = 18.3, PRD = 58.8%)

Fig. 4 provides a visual evaluation of reconstruction using quantized measurements. An original signal and reconstructed signal using different quantization rates are shown. In general, the QRS complexes in the ECG are preserved even at the lowest quantization rates. The higher PRD values for these rates are due to added noise in the time periods between these complexes. It is clear from Fig. 4 that despite the PRD values reaching approximately 40%, applications requiring QRS feature extraction could expect good performance at these PRD levels, as the QRS complex is well preserved.

B. Energy Consumption

In order to analyze the advantages of implementing this quantization, the energy savings in terms of wireless transmission are explored. The results in Table II are presented considering BLE as a target low-power wireless protocol and the E_{TF} is estimated by assuming a continuous transmission link. The savings from quantization are significant for higher PRDs suggesting that applications requiring low power QRS detection would benefit most from the proposed quantization.

TABLE II ENERGY TRANSMISSION PER FRAME FOR DIFFERENT PRD VALUES

PRD	Non-Uniform Quantization E_{TF} (μ J)	No Quantization E_{TF} (μ J)	Percentage Reduction
10%	62.13	70.41	11.76%
20%	33.82	45.25	25.26%
30%	26.8	37.7	28.91%

V. CONCLUSION

This paper has proposed the use of quantization to maximize the CR in ECG CS applications. A form of non-uniform quantization is compared to a uniform approach. The effects of quantization are shown in terms of PRD and from an energy point of view. The quantization proposed provides significant compression gains with CS for applications where PRD > 6.75% would be acceptable. For example in an application where a PRD of 30% is tolerable, 7 bit quantization can extend the CR from 9.8 to 14.1. From an energy perspective this amounts to a 28.91% reduction in E_{TF} from 37.7 μ J to 26.8 μ J.

ACKNOWLEDGMENT

This work was supported by the Irish Research Council (IRC) and the Higher Education Authority (HEA) under the Program for Research in Third Level Institutions.

REFERENCES

- [1] G. V. Crosby, T. Ghosh, R. Murimi and C. A. Chin, "Wireless Body Area Networks for Healthcare: A Survey," *International Journal of Ad hoc, Sensor & Ubiquitous Computing*, vol. 3, pp. 1-8, 2012.
- [2] H. Mamaghanian, N. Khaled, D. Atienza and P. Vanderghyest, "Compressed Sensing for Real-Time Energy-Efficient ECG Compression on Wireless Body Sensor Nodes," *IEEE Transactions on Biomedical Engineering*, vol. 58, pp. 2456-2466, 2011.
- [3] E. J. Candes, J. Romberg and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, pp. 489-509, 2006.
- [4] E. J. Candes and T. Tao, "Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?," *IEEE Transactions on Information Theory*, vol. 52, pp. 5406-5425, 2006.
- [5] D. Craven, B. McGinley, L. Kilmartin, M. Glavin and E. Jones, "Compressed Sensing for Bioelectric Signals: A Review," *IEEE Journal of Biomedical and Health Informatics*, in press.
- [6] MIT-BIH Arrhythmia Database. [Online]. (2000). Available: <http://www.physionet.org/physiobank/database/mitdb/>.
- [7] L. F. Polania, R. E. Carrillo, M. Blanco-Velasco and K. E. Barner, "Compressed sensing based method for ECG compression," in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 761-764, 2011.
- [8] G. Higgins *et al.*, "The Effects of Lossy Compression on Diagnostically Relevant Seizure Information in EEG Signals," *IEEE Journal of Biomedical and Health Informatics*, vol. 17, pp. 121-127, 2013.
- [9] H. Mamaghanian, N. Khaled, D. Atienza and P. Vanderghyest, "Design and Exploration of Low-Power Analog to Information Conversion Based on Compressed Sensing," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 2, pp. 493-501, 2012.
- [10] S. S. Chen, D. L. Donoho and M. A. Saunders, "Atomic decomposition by Basis Pursuit," *Journal on Scientific Computing*, vol. 20, pp. 33-61, 1999.