

Diagnostic Grade Wireless ECG Monitoring

Harinath Garudadri, Yuejie Chi, Steve Baker, Somdeb Majumdar, Pawan K. Baheti, Dan Ballard

Abstract—In remote monitoring of Electrocardiogram (ECG), it is very important to ensure that the diagnostic integrity of signals is not compromised by sensing artifacts and channel errors. It is also important for the sensors to be extremely power efficient to enable wearable form factors and long battery life. We present an application of Compressive Sensing (CS) as an error mitigation scheme at the application layer for wearable, wireless sensors in diagnostic grade remote monitoring of ECG.

In our previous work, we described an approach to mitigate errors due to packet losses by projecting ECG data to a random space and recovering a faithful representation using sparse reconstruction methods. Our contributions in this work are two-fold. First, we present an efficient hardware implementation of random projection at the sensor. Second, we validate the diagnostic integrity of the reconstructed ECG after packet loss mitigation. We validate our approach on MIT and AHA databases comprising more than 250,000 normal and abnormal beats using EC57 protocols adopted by the Food and Drug Administration (FDA).

We show that sensitivity and positive predictivity of a state-of-the-art ECG arrhythmia classifier is essentially invariant under CS based packet loss mitigation for both normal and abnormal beats even at high packet loss rates. In contrast, the performance degrades significantly in the absence of any error mitigation scheme, particularly for abnormal beats such as Ventricular Ectopic Beats (VEB).

I. INTRODUCTION

For remote monitoring of ECG, it is extremely important to maintain the clinical integrity of the signals. Continuous monitoring of ECG is widely used in many clinical settings, including Intensive Care Units (ICUs), post-operative monitoring, emergency care and in ambulatory settings such as Holter monitoring. As interpretation of continuous ECG requires analysis of as many as 10^5 cardiac cycles per patient per day, there has been a need for tools to perform automated labeling and classification of ECG. American National Standard Institute (ANSI) and Association for the Advancement of Medical Instrumentation (AAMI) have established standards for automated tools such as EC57 [1] that are recognized by the FDA in United States. The tools for implementing these protocols, along with ECG databases and annotations by experts for normal and abnormal ECG beats are available in public domain at [2], [3]. In this work, we refer to annotations by experts as ground truth.

H. Garudadri, S. Majumdar and P.K. Baheti are with Qualcomm Research Center, 5775 Morehouse Dr., San Diego, CA, USA 92121. {hgarudad, smajumda, pbaheti} AT qualcomm.com

S.D. Baker and D. Ballard are with Welch Allyn, 8500 SW Creekside Pl, Beaverton, OR, USA, 97008. {Steve.Baker, Dan.Ballard} AT welchallyn.com

Y. Chi is a Ph.D. Candidate, Department of Electrical Engineering, Princeton University, Princeton, NJ 08540. ychi AT princeton.edu

Commercial ECG machines that provide automated labeling and classification benchmark their classification performance against the ground truth. While EC57 does not specify minimum performance requirements for metrics such as Sensitivity (Se – percentage of true events detected) and Positive Predictivity (+P – percentage of detected events that are true), it mandates that such performance metrics with databases in [1] are disclosed. From a clinicians' perspective, a wireless continuous ECG monitoring system should provide diagnostic utility similar to that of a wired system in current standard of care, while enabling non-intrusive form factors for use in free living conditions. This means that there should be no statistically significant degradation in performance due to the wireless link. We observe that 2% or more degradation in classification accuracy from a "wired" baseline performance is clinically significant and about 1% packet loss rate (PLR) can cause significant degradation in performance. In this work, we evaluate the proposed packet loss mitigation approach with MIT and AHA databases using a state-of-the-art commercial ECG arrhythmia classification software and show that the performance does not degrade even at high packet loss rates.

Packet losses occur in wireless networks due to fading, interference, congestion, system loading, etc. Popular choices for radios in Body Area Networks (BAN) such as Bluetooth, Zigbee, etc. operate in the crowded 2.4 GHz band, along with IEEE 802.11. In [4], the authors investigated the interference of 802.11 traffic presented to ZigBee nodes in BAN and found 33% – 56% packet loss rate, depending upon the network setup. Another study [5] based on Zigbee reported packet losses as much as 50% in a clinical trial involving remote ECG monitoring. Note that the ECG signal can be quite sparse in time domain, where important events like QRS-complex occur over a short period of time. Thus, packet losses can result in significant loss of clinically relevant data. Previously, we identified the need for action and proposed an approach to mitigate packet losses [6] based on Compressive Sensing (CS).

CS is an emerging signal processing concept, wherein significantly fewer sensor measurements than that suggested by Nyquist-Shannon sampling theorem can be used to recover sparse signals with high fidelity [7], [8], [9]. The measurements in CS framework are generally defined as inner-products of the signal with random basis functions. CS relies on the assumption that the signal of interest is sparse in some representation basis with only M non zero elements, where $M \ll N$ and N is the signal dimensionality. These signals can be recovered faithfully if an order of $M \cdot \log N/M$ samples are available at the receiver, albeit with

some additional computational complexity at the receiver. We consider lost packets as random sampling by the wireless channel and leverage CS to reconstruct the signal from fewer measurements. We project ECG in to a random space, by pre-coding at the sensor; this reduces sparsity in time and spreads information over an entire frame. Each frame is divided into multiple packets for transmission over the air. The receiver is then able to reliably reconstruct samples in the Nyquist domain, even under packet losses.

It is highly desirable for wearable sensors to last as long as possible, preferably a week with user-friendly form factors, in order to support early discharge from hospitals and home monitoring. In [10], we presented software implementation details and real-time validation of the the proposed approach. Here, we present a power efficient hardware implementation of CS pre-coding for a multi-lead wireless ECG sensor. This enables us to offload computationally demanding operations, including sensing artifacts mitigation, to receivers with better battery budgets.

In Section 2, we review the CS operations for signal measurements at sensor node and reconstruction at receiver node. In Section 3, we describe our method for evaluating the diagnostic integrity of received ECG signals. In Section 4, we present results of EC57 procedures on standard ECG databases for both CS based source compression and CS based channel error resiliency. We present conclusions in Section 5.

II. CS OPERATIONS AT SENSOR AND RECEIVER

In this section we briefly review the CS framework for sensing and reconstruction of sparse signals and present applications in ECG telemetry.

A. Overview of CS based Error Mitigation

Consider a short term segment of a signal $x(n)$, of length N , denoted by an N -dimensional vector \mathbf{x} with f_s as its Nyquist sampling frequency. Let the matrix \mathbf{W} represent basis functions, consisting of $N \times N$ elements. We normalize each row of the matrix \mathbf{W} such that the corresponding L_2 -norm is equal to 1. The transform domain representation of the signal, \mathbf{y} , can be computed as

$$\mathbf{y} = \mathbf{W}\mathbf{x} \quad (1)$$

Sparsity: \mathbf{x} is an M -compressible signal if there are only M significant components for a given transform \mathbf{W} . Let the total energy of the N -component segment be E_N and the total energy contributed by any group of M components be E_M . Then the number of significant components may be obtained by finding the smallest M such that $1 - \frac{E_M}{E_N} \leq \epsilon$, where $\epsilon \ll 1$. We call the ratio $\frac{M}{N}$ as the sparsity ratio and note that a smaller value of this ratio indicates higher compressibility.

In this case, we call \mathbf{W} as the sparsity basis. In the CS paradigm, if one is able to construct a measurement matrix \mathbf{H} of dimension $K \times N$ that is statistically incoherent with the sparse basis \mathbf{W} , then only K measurements given by

$$\mathbf{r} = \mathbf{H}\mathbf{x}. \quad (2)$$




Fig. 1. Sparsity of ECG signals: blue and red plots represent sparsity ratios in frequency and time domains, respectively.

are adequate to estimate \mathbf{y} with a high probability of a small reconstruction error provided [7], [8]

$$K \geq M \log N/M. \quad (3)$$

We observe that the ECG signal has both transient and tonal components and a significant amount of energy is concentrated in very few coefficients in both time and frequency domains. Figure 1 shows the distribution of sparsity for 172 records, comprising a total of 142 hours of ECG data from the MGH database. The x-axis represents all possible values of the sparsity ratio and the y-axis represents the number of segments for a given sparsity ratio. We observed that in the Discrete Cosine Transform (DCT) domain, the average value of the sparsity ratio was 0.08 and for time domain 0.45. It is clear from the graph that ECG signals are significantly sparse in the DCT domain; thus the choice of DCT as \mathbf{W} is justified.

In the CS paradigm, smaller reconstruction errors are obtained with fewer measurements if reconstruction is performed in a sparse domain; however, if the measurement domain is also sparse, then significant information may be lost if $K < N$ samples are measured.

It is clear from Figure 1 that ECG signals are quite sparse in the time domain as well; the mean sparsity across the database being only 0.45. This implies that random sampling directly in time domain, as with a lossy wireless channel, could result in loss of clinically relevant information.

In this work, we are concerned with packet loss mitigation. We consider a lossy channel as performing random sampling of ECG packets over the air and define \mathbf{H}_c the sensing kernel by the channel. Most communication protocols provide a Sequence Number or similar capability to identify the packets were dropped by the channel. We use this information to define a set \mathbf{S} , consisting of the indices of ECG packets lost in the channel. The packet loss rate is given by $E[\#\mathbf{S}]/K$, where $\#\mathbf{S}$ denotes the cardinality of set \mathbf{S} and $E[\cdot]$ represents the expectation operator. Note that the cardinality of set \mathbf{S} is random because of the stochastic nature of the channel. Given \mathbf{S} , we can construct \mathbf{H}_c , of dimension $K \times N$ formed by starting with an $N \times N$ identity matrix and removing rows indexed by the elements of \mathbf{S} .

As random sampling in time domain performed by the channel will result in loss of clinically relevant information, we project the the ECG signal in to a random space, prior to transmission. Let this be the sensing kernel \mathbf{H}_s of dimension $N \times N$, whose elements are independently chosen from the symmetric Bernoulli distribution such that

$$\Pr([\mathbf{H}_s]_{i,j} = 0 \text{ or } 1) = \frac{1}{2}; 0 \leq \{i, j\} < N. \quad (4)$$

Using \mathbf{H}_s from Eq. 4 in Eq. (2) results in N randomly projected measurements denoted by \mathbf{r} . This is sent over the lossy channel as J discrete packets, each packet with P samples of payload. The effective sensing kernel \mathbf{H} to be used in Eq (5) for reconstruction at the receiver is then given by

$$\mathbf{H} = \mathbf{H}_c \mathbf{H}_s. \quad (5)$$

Incoherence: In Compressed Sensing, we require that the measurement kernel \mathbf{H} is incoherent with the representation basis \mathbf{W} . In this context, coherence μ is defined to be the largest element of the product $\sqrt{N} \mathbf{H} \mathbf{W}$ where both \mathbf{H} and \mathbf{W} are $N \times N$ matrices and bounded by $1 \leq \mu \leq \sqrt{N}$. For the Nyquist sampling, \mathbf{H} is an identity matrix, maximally incoherent with \mathbf{W} , with $\mu = 1$. It is well-known that the construction of \mathbf{H} as described in Eq. (4) generates a universal CS encoder that exhibits low coherence with a structured reconstruction basis like the DCT or FFT [11].

Our choice for \mathbf{H}_s described in subsection II-B – a sparse random matrix derived from a subset of a sequence of Bernoulli random variable was driven by sensor power and hardware complexity constraints. μ for this \mathbf{H}_s and DCT basis was 3.88, reasonably close to the lower bound.

Note that the product of \mathbf{H}_c and \mathbf{H}_s is still random and incoherent with \mathbf{W} since \mathbf{H}_s is random by construction and \mathbf{H}_c is obtained by simply choosing a subset of rows from \mathbf{H}_s . Also note that this approach does not discriminate based on the location of the lost measurements and all received packets are of equal importance; therefore, it can handle the bursty errors typically associated with wireless channels.

B. Sensor Side Processing

We consider that case where the sensing matrix \mathbf{H}_s is a square, $N \times N$ matrix that is also invertible and, hence, full-rank. We implemented CS encoding as an online matrix multiplication between \mathbf{H}_s and x . To understand this, consider the matrix multiplication $\mathbf{r} = \mathbf{H}_s \mathbf{x}$ where:

$$\begin{aligned} \mathbf{x} &= \{x_i\}_{i=1:N} \\ \mathbf{r} &= \{r_i\}_{i=1:N} \\ \mathbf{H}_s &= \{h_{ij}\}_{i=1:N, j=1:N} \end{aligned}$$

Then, the elements of the output r are given by:

$$r_k = h_{k1}x_1 + h_{k2}x_2 + \dots + h_{kM}x_M \quad (6)$$

where $k = 1 : N$. We can represent this operation as each x_i contributing the value $h_{ki} * x_i$ to the value of r_k for each $k = 1 : M$. In our memory optimized hardware realization, for each incoming sample x_i , we simply computed its partial




Fig. 2. Fully random (left) and sparse (right) \mathbf{H}_s . Red regions indicate 1 and blue regions indicate 0.

contribution to each of the elements in the output vector y and then discarded it.

The elements of the matrix \mathbf{H}_s were chosen to be part of a random P-N sequence. This was achieved by generating a Linear Feedback Shift Register (LFSR) sequence corresponding to each of the elements h_{ij} and quantizing them to a one-bit value in $\{0, 1\}$. Thus, the matrix multiplication was implemented as additions only, corresponding to the state of the quantized LFSR output. It is also important to keep in mind bit-growth due to the CS encoding operation. Suppose that the incoming x_i are b_1 bits each. If the k^{th} row of the matrix \mathbf{H}_s contains D ones, then the output r_k will require $b_2 = b_1 + \log_2(D)$ bits if full resolution is to be maintained. To reduce the memory required to store and the bandwidth required to transmit r , one approach is to implement a sparse matrix where the number of ones per row is small [12]. In this work, the input \mathbf{x} was 16 bits/sample and the memory budget for the output \mathbf{r} was 20 bits/sample. Therefore, we constructed a matrix with a maximum of 16 ones per row to allow for 4 bits of expansion.

For a real-time implementation, we utilized a double-randomization scheme where each column c_k of the \mathbf{H}_s matrix was a one-bit quantized output of a $\log_2(N)$ LFSR sequence starting at a seed s_k . It is well known that a b -bit LFSR sequence is cyclic with periodicity $2^b - 1$ and is unique for a given starting seed s . Thus, by selecting an LFSR sequence with a maximum length of $N - 1$ and a quantization threshold, we were able to control the density of ones per row. The quantization threshold was selected based on the desired ones-density per row. In order to make the columns statistically independent, the master LFSR sequence that provided the quantized $\{0, 1\}$ values to also point to a starting seed s_k for the k^{th} column. Figure 2 is a graphical representation of fully random and sparse \mathbf{H}_s .

In [13], a similar double-randomization scheme was used to perform CS random projections on biophysical signals such as ECG, followed by under sampling to achieve compression. Their random matrix (\mathbf{H}_s) was of dimension 50×1000 for CS based compression, in contrast with 128×128 for CS based packet loss mitigation addressed here. They demonstrate the efficiency of the scheme with a hardware realization in 90 nm CMOS. The silicon area is $200 \mu\text{m} \times 450 \mu\text{m}$ and consumes $1.9 \mu\text{W}$ [13].




Fig. 3. ECG Reconstruction example. A segment of the original waveform from record 203 of MIT-BIH is shown in top pane. The pre-coded data along with packet losses (35% in this case) is shown in the middle pane. The reconstructed ECG waveform is shown in the bottom pane. The difference with the original can be seen around 0 amplitude in the bottom pane. The labels at the bottom of top and bottom panes correspond to ground truth and Mortara classifications, respectively.

C. Receiver Side Processing

On the receiver side, many approaches described in literature can be leveraged to reconstruct the Nyquist domain equivalent \hat{x} from the received signal r with missing packets. Examples of such approaches include Gradient-Projection based Sparse Reconstruction (GPSR) [14], Orthogonal Matching Pursuits (OMP) [15], [16], [17]. In our work, we implemented a reconstruction using [17].

Fig. 3 shows a 6-second segment of ECG recording from the record 203 of MIT-BIH database. H_s is a 128×128 matrix and each ECG frame is transmitted in 8 packets. The middle pane shows the data in the random space, with missing segments corresponding to packet losses. It can be seen from the reference (top pane) and test (bottom pane) annotations that there are two locations where an atrial premature beat **A** was mis-categorized as a normal sinus rhythm **N**. All the remaining beats were correctly classified in this segment.

III. VALIDATION FOR DIAGNOSTIC GRADE ECG

As described in the Introduction, ANSI/AAMI specification EC57 [1] provides a framework for validating the performance of automated software tools that classify and label large amounts of ECG waveform data resulting from continuous monitoring. In this section, we present results from applying these “current standard of care” protocols in wired settings to wireless ECG monitoring. The databases we consider are MIT-BIH Arrhythmia Database (48 records of 30 minutes each) and AHA database for ventricular arrhythmia

detectors (80 records of 35 minutes each). Overall, these databases contain a wide variety of cardiac rhythms comprising nearly 250,000 normal and 22,000 abnormal heart beats from multiple subjects. All of the records in the databases come with ground truth of annotations by cardiologists. Commercial ECG machines that provide automated labeling and classification benchmark their performance against this ground truth. The metrics for beat classification performance are Sensitivity (Se – percentage of true events detected) and Positive Predictivity (+P – percentage of detected events that are true), defined as follows [1]:

$$\text{Se} = \text{TP}/(\text{TP}+\text{FN}) \quad (7)$$

$$+\text{P} = \text{TP}/(\text{TP}+\text{FP}) \quad (8)$$

where,

- A correctly detected event is called a true positive (TP)
- An erroneously rejected (missed) event is called a false negative (FN)
- An erroneously detected non-event is called a false positive (FP)
- A correctly rejected non-event is called a true negative (TN).

In this work, we implemented arrhythmia analysis and beat classification using the Mortara algorithm. We re-sampled each record at 500 Hz and scaled to $2.5 \mu\text{V}/\text{LSB}$ to meet the specifications for ECG data as input to the Mortara arrhythmia analysis library. The arrhythmia analysis library was compiled into an executable to read the ECG records and output measurements including heart rate, ST values,

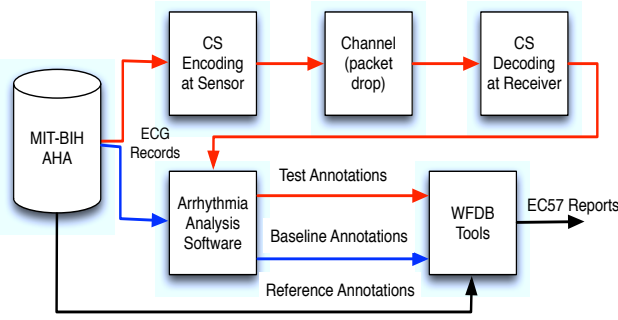


Fig. 4. Experimental setup for ECG validation

QRS amplitudes, etc. along with beat and event classification such as normal sinus rhythm, pre-ventricular, ventricular fibrillation, asystole, bigeminy, pause, etc. The algorithm processes from 1 to 8 leads and can detect QRS complexes as long as at least one lead is valid. The output from Mortara is formatted such that it can be used directly with EC57 tools for comparison with the ground truth annotations by cardiologists provided in the databases. EC-57 requires testing and disclosure of the algorithms' sensitivity and positive predictivity along with Root Mean Square (RMS) heart rate error. The comparison may start after 5 minutes from the beginning of the record. For a beat to be correctly classified, the algorithm must identify the beat with correct classification within 150 ms of the actual event.

Fig. 4 depicts the experimental setup used to validate the packet loss mitigation proposed in this study. The black path labeled "Reference Annotations" represents annotations of the ECG waveforms by experts ECG. The blue path labeled "Baseline Annotations" represents the current golden standard in a wireline setting. Each cardiac database was analyzed to provide a baseline performance measure to confirm that introduction of compressive sampling did not affect the system performance and to confirm that the system performance of the test bed used for this paper provides results that match those from prior EC-57 compliance tests.

The red path labeled "Test Annotations" represents wireless case in this study. At the sensor, we used a sparse \mathbf{H}_s of size 128×128 as described in Section II-B. Note that increasing the dimensions of \mathbf{H}_s provides better reconstruction accuracy, at the expense of increase in encoder complexity and additional latency. It is essential from low power perspective that the application layer is optimized for a given radio in BAN. We experimented with 32 packets and 8 packets per frame, corresponding to 4 and 16 ECG samples per packet, respectively. A bursty channel-error model was used to drop packets at loss rates of 0.5%, 1%, 5%, 15%, 25% and 35%. The received data was reconstructed in to Nyquist domain and provided to the Mortara arrhythmia analysis algorithm to generate annotations.

IV. RESULTS

In this section, we evaluate the proposed packet loss mitigation approach with MIT and AHA databases using

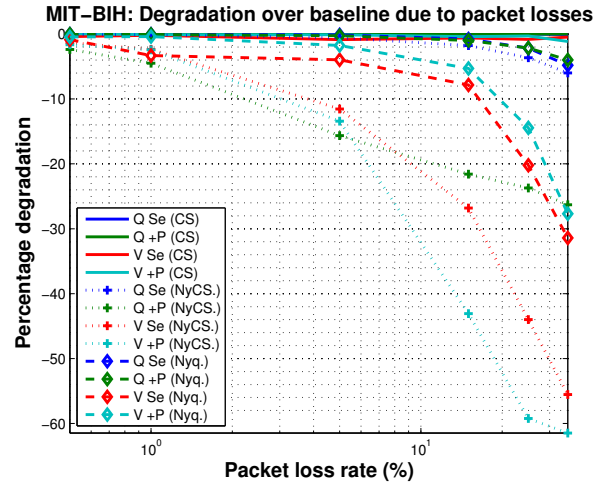


Fig. 5. MIT-BIH Degradation. The legends Q and V correspond to normal QRS sinus rhythms and abnormal VEB rhythms, respectively. The legends Se and +P correspond to Sensitivity and Positive Predictivity, respectively. The solid lines represent the CS based packet loss mitigation approach (CS), the dashed lines represent the case with no random projections at the sensor but with sparse reconstruction at the receiver (NyCS) and the dotted lines represent the Nyquist domain data (NyQ), respectively.

a state-of-the-art commercial ECG arrhythmia classification software and show that the performance does not degrade even at high packet loss rates. We present EC-57 analysis results, specifically Se and +P for normal beats Q (QRS segments) and abnormal beats V (Ventricular Ectopic Beats, VEB; also sometimes called Premature Ventricular Contraction, PVC) in the next section for MIT-BIH and AHA databases.

Figures 5 and 6 present degradation in beat classification as a function of packet loss rate for MIT-BIH and AHA databases, respectively. The solid lines represent the CS based packet loss mitigation approach (CS), the dashed lines represent the case with no random projections at the sensor but with sparse reconstruction at the receiver (NyCS) and the dotted lines represent the Nyquist domain data (NyQ), respectively. Specifically, \mathbf{H}_s is as described in Section II-B for CS, and identity matrix for NyCS and NyQ, respectively; and the sparse signal reconstruction at the receiver is as described in Section II-C for CS, NyCS and "null" for NyQ, respectively. Each ECG frame was packetized into 32 packets in this experiment. The legends Q and V correspond to normal QRS sinus rhythms and abnormal VEB rhythms, respectively. The legends Se and +P correspond to Sensitivity and Positive Predictivity, respectively.

We observe that 2% or more degradation in classification accuracy from a "wired" baseline performance is clinically significant. From Figures 5, 6 it can be seen that performance degrades monotonically for both NyCS and NyQ, compared with CS. This is particularly true for Sensitivity and Positive Predictivity of abnormal rhythms (V), compared with normal sinus rhythms (Q). The Positive Predictivity degradation is also severe, suggesting more false positives, as packet loss rate increases. Without some method of packet loss mitiga-

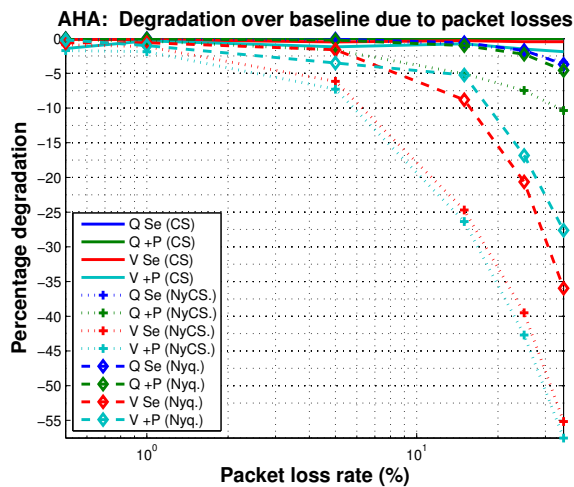


Fig. 6. AHA Degradation. See Figure 5 caption for legend description.

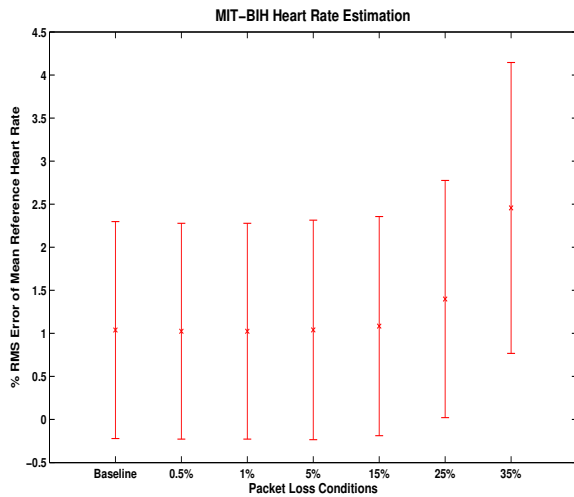


Fig. 7. MIT-BIH: Heart Rate Estimation Error with packet loss conditions.

tion, a 1% packet loss rate can cause clinically significant degradation. While the high packet loss conditions studied here are corner cases, we believe that typical loss rates of around 5% are typical for BAN modems in the crowded 2.4 GHz band. As a reference, packet loss rates of 1 – 3% are commonly used to evaluate voice quality in 3G standards.

Figure 7 shows the performance of heart rate (HR) estimation for the MIT-BIH database, for different packet loss conditions. The mean and standard deviation of Mortara HR estimation error in percentage, against the reference HR from the database for the baseline and various packet loss conditions is shown. There were 8 packets per ECG frame in this experiment. It can be seen that the HR estimation is quite robust in the presence of losses, with the proposed approach.

V. CONCLUSIONS

In this work, we addressed the challenges of packet losses in wireless ECG monitoring. We reviewed the ANSI/AAMI

framework for validating the performance of automated software tools that classify and label large amounts of ECG waveform data resulting from continuous monitoring in current standard of care. We observe that 2% or more degradation in classification accuracy from a “wired” baseline performance is clinically significant. We show that, without some explicit packet loss mitigation, packet loss rates as low as 1% can cause clinically significant degradation. We presented a brief overview of CS principles and proposed a packet loss mitigation scheme that consumes very low power at the sensor node. We validated the proposed approach with MIT-BIH and AHA databases and show that sensitivity and positive predictivity of arrhythmia classification is statistically similar to that of “wired” baseline at moderate packet losses observed in typical wireless networks.

REFERENCES

- [1] American National Standard, *Testing and Reporting Performance Results of Cardiac Rhythm and ST Segment Measurement Algorithms*, ANSI/AAMI/ISO EC57:1998/(R)2008, Arlington (VA), 1998.
- [2] *PhysioNet: the research resource for complex physiologic signals*, online at <http://www.physionet.org/>.
- [3] George B. Moody, *Evaluating ECG Analyzers*, available online at <http://mail.physionet.org/physiotools/wag/eval.htm>, Harvard-MIT Division of Health Sciences and Technology, Cambridge, MA, USA.
- [4] J. Hou, B. Chang, D-K Cho, and M. Gerla, “Minimizing 802.11 interference on zigbee medical sensors,” in *BodyNets 2009 conference*, April 2009.
- [5] V. Shnayder, B. Chen, K. Lorincz, T. R. F. Fulford-Hones, and M. Welsch, “Sensor networks for medical care,” in *Harvard University Technical Report TR-08-05*, 2005.
- [6] H. Garudadri and P. K. Baheti, “Packet loss mitigation for biomedical signals in healthcare telemetry,” in *Proc. IEEE EMBS Conf.*, Minneapolis, USA, September 2009.
- [7] D. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, vol. 52, pp. 1289–1306, April 2006.
- [8] J. Romberg, E. Candes and T. Tao, “Stable signal recovery from incomplete and inaccurate measurements,” *Communications on Pure and Applied Mathematics*, vol. 59, pp. 1207–1223, August 2006.
- [9] *IEEE Signal Processing Magazine [Sensing, Sampling, and Compression]*, vol. 25, IEEE Signal Processing Magazine, March 2008.
- [10] H. Garudadri, P.K. Baheti, S. Majumdar, C. Lauer, F. Mass and, J. van de Molengraft, and J. Penders, “Artifacts mitigation in ambulatory ecg telemetry,” in *12th IEEE International Conference on e-Health Networking Applications and Services (Healthcom), 2010. IEEE*, 2009, pp. 338–344.
- [11] E. Candes and J. Romberg, “Sparsity and incoherence in compressive sampling,” *Inverse Problems*, vol. 23, pp. 969–985, June 2007.
- [12] R. Berinde and P. Indyk, *Sparse recovery using sparse random matrices*, MIT-CSAIL Technical Report, 2008.
- [13] F. Chen, A.P. Chandrakasan, and Stojanovic V., “A signal-agnostic compressed sensing acquisition system for wireless and implantable sensors,” in *Custom Integrated Circuits Conference (CICC), 2010 IEEE*, sept. 2010, pp. 1–4.
- [14] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, “Gradient projection for sparse reconstruction: Applications to compressed sensing and other inverse problems,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, pp. 586–598, Dec 2007.
- [15] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, “Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition,” in *Proc. Conf. Rec. 27th Asilomar Conf. Signals, Syst. Comput.*, 1993, pp. 40–44.
- [16] M. F. Duarte, M. A. Davenport, M. B. Wakin, and R. G. Baraniuk, “Sparse signal detection from incoherent projections,” in *Proc. Int. Conf. on Acoustics, Speech, and Signal Proc. (ICASSP)*, May 2006.
- [17] R. Rubinstein, M. Zibulevsky, and M. Elad, *Efficient implementation of the K-SVD algorithm using batch orthogonal matching pursuit*, CS Technion, 2008.