31st Annual International Conference of the IEEE EMBS
Minneapolis, Minnesota, USA, September 2-6, 2009

Packet Loss Mitigation for Biomedical Signals in Healthcare Telemetry

Harinath Garudadri and Pawan K. Baheti

Abstract— **In this work, we propose an effective application layer solution for packet loss mitigation in the context of Body Sensor Networks (BSN) and healthcare telemetry. Packet losses occur due to many reasons including excessive path loss, interference from other wireless systems, handoffs, congestion, system loading, etc. A call for action is in order, as packet losses can have extremely adverse impact on many healthcare applications relying on BAN and WAN technologies. Our approach for packet loss mitigation is based on Compressed Sensing (CS), an emerging signal processing concept, wherein significantly fewer sensor measurements than that suggested by Shannon/Nyquist sampling theorem can be used to recover signals with arbitrarily fine resolution. We present simulation results demonstrating graceful degradation of performance with increasing packet loss rate. We also compare the proposed approach with retransmissions. The CS based packet loss mitigation approach was found to maintain up to 99% beat-detection accuracy at packet loss rates of 20%, with a constant latency of less than 2.5 seconds.**

I. INTRODUCTION

An important aspect for biomedical telemetry in healthcare applications is to provide an end-to-end reliable communication link, while minimizing sensor power and communication latency. The Harvard CodeBlue project [1] reports that up to 50% packet loss rates were observed in a recent trial involving a multi-hop network. In [2], the authors investigated interference of 802.11 traffic presented to ZigBee nodes in BSN. Their experiments showed $33\%-56\%$ packet loss rate, depending upon the network setup. It is possible to improve the packet loss performance with Quality of Service (QoS) aware networks. One effort to provide QoS in 802.x lower layers [3] argues that traffic pattern based scheduling does not work well for telemedicine applications. They study a dual-channel approach where one channel is reserved for emergency alert messages. In their simulations, they observe packet loss rates of 5 to 25%. Another study found that packet loss rate increases with network congestion, and proposed the use of DiffServ for QoS provisioning [4]. Some authors have also explored the use of Forward Error Correction (FEC) coding in BAN scenarios [5], [6]. Using FEC schemes the authors in [5] observed a residual 3.4% packet loss rate for 2 second latency. In [6], the authors investigated ECG over a GPRS link and were able to achieve 0.1 to 0.3% loss rate for a 3 second latency. It is evident that there is a need for more reliable link, with lower delay and lower power consumption.

Our approach for packet loss mitigation is based on the compressed sensing (CS) paradigm. CS is an emerging signal processing concept, wherein significantly fewer sensor measurements than that suggested by Shannon/Nyquist sampling theorem can be used to recover signals with arbitrarily fine resolution [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 basis representation 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 $M \cdot \log N$ samples are available at the receiver, albeit with some additional computational complexity at the receiver. In the context of BAN, this is desirable as we shift computational complexity to nodes with flexible power budgets and increase life of the sensor. In [10], [11], we described an application of CS for reducing the sensor power for pulse oximeters. An interesting aspect of CS is that it lends itself conveniently for packet loss mitigation, as receivers are equipped to reconstruct signals in some sparse domain. The receiver operation is one and the same whether measurements are omitted at the sender or dropped by the communication channel. In this work, we exploit these capabilities of a CS based receiver for packet loss mitigation in healthcare applications.

The remainder of the paper is organized as follows. In Section 2, we review the CS operations for acquisition at sensor and reconstruction at receiver. In Section 3, we describe the CS based approach for packet loss concealment (CS-PLC). In Section 4, we present experimental results based on ECG signals from the MIMIC database for the proposed scheme and comparisons with traditional approaches such as retransmissions. We present conclusions in Section 5.

II. CS FRAMEWORK

In this work, we are interested in exploiting the CS paradigm for packet loss mitigation. We focus on ECG signal telemetry here, and note that the methodologies described are applicable to other biomedical signals. Consider a short term segment of an ECG signal, denoted by $N-$dimensional vector $\mathbf{x}$ and f_s is the sampling frequency of $\mathbf{x}$. In order to exploit sparsity we make use of a Gabor basis, consisting of various cosine waves with time support limited by Gaussian window functions at different scales. Let the sparse-domain transform basis be represented by $N \times N$ matrix $\mathbf{W}$. The (i, j) entry of matrix $\mathbf{W}$ is given by,

$$[\mathbf{W}]_{i,j} = \cos\left(\frac{2\pi(i-1)(j-1)}{2N}\right) \times \exp\left(-\frac{(i-1)^2(j-N/2)^2}{wN^2}\right). \quad (1)$$

The authors are with Qualcomm Incorporated, 5775 Morehouse Drive San Diego, CA 92121, USA. Email: hgarudad@qualcomm.com; pbaheti@qualcomm.com

The term w is associated with the width of the Gaussian kernel in the Gabor basis. We normalize each row of the matrix $\mathbf{W}$ such that the corresponding L_2-norm is equal to 1, and we will refer to $\mathbf{W}$ as the sparse-basis. The ECG signal $\mathbf{x}$ is projected on the sparse-basis to generate the corresponding $N-$dimensional representation in transform space and it is given by

$$\mathbf{y} = \mathbf{W}\mathbf{x}. \tag{2}$$

The recently developed CS framework by Candes and Donoho states that if $\mathbf{x}$ is explicitly-sparse with only M non-zero elements in the transform space then $K \geq M \log N/M$ random measurements (i.e. projections of $\mathbf{x}$ on a $K \times N$ random basis) provides sufficient information, with high probability, to enable signal reconstruction with zero error [7], [8]. In real situations the signal is never truly sparse and has some information content throughout the transform-space; however the number of significant components with magnitude greater than ϵ, where $\epsilon \ll \max(\mathbf{y})$, is much smaller than N. The extensions of Candes's result to the case where $\mathbf{x}$ is not explicitly-sparse has been presented in [8]. The CS paradigm still remains valid; however, the reconstruction error does not go to zero. Note that in the CS framework, the measurement matrix is chosen such that it is statistically incoherent with the sparse basis (i.e., independent of signal prior). Typically it is defined as a matrix containing random i.i.d elements (which may be normally distributed as an example).

We now express mathematically the sensing process for $\mathbf{x}$. The first step is to create a measurement matrix $\mathbf{H}$ of dimension $K \times N$, whose elements are independently chosen from the symmetric Bernoulli distribution $\Pr(\mathbf{H}_{i,j} = -1 \text{ or } 1) = \frac{1}{2}$. The next step is to transform the original signal $\mathbf{x}$ using the measurement matrix $\mathbf{H}$. The resulting $K-$dimensional measurement vector $\mathbf{r}$ is given by

$$\mathbf{r} = \mathbf{H}\mathbf{x}. \tag{3}$$

The CS reconstruction framework makes use of the signal prior (i.e., sparsity) and the measurement vector $\mathbf{r}$ to estimate the signal $\mathbf{x}$. In this work we make use of the matching pursuit (MP) algorithm for signal reconstruction from measurement vector $\mathbf{r}$. The MP technique is a greedy algorithm that builds up a signal approximation iteratively by making locally optimal decisions [12], [13]. Each iteration of the MP algorithm costs on the order of $O(KN)$ arithmetic operations. More details about this technique can be found in [13].

III. CS BASED PACKET-LOSS CONCEALMENT

In this section, we address the issue of end-to-end packet losses due excessive path loss, interference from other wireless systems, handoffs, congestion, system loading, etc. We observe that when the signal being transmitted has redundancies, we can treat the packet-losses as lossy compression performed by the channel. We assume that we can identify lost packets at the application layer, e.g. via a Sequence Number field in the headers in lower layers of the protocol stack.

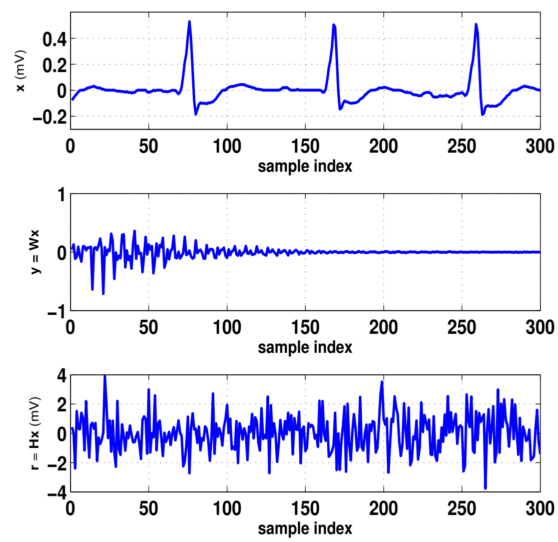

Fig. 1. Sample ECG signal with $N = 300$ (top), Transform of the signal in Gabor space (middle), and Precoded version of the signal (bottom)

At the sender, we perform a precoding operation on sensor data by using the random measurement matrix $\mathbf{H}$. This operation spreads information across multiple packets and enables us to reconstruct the information in lost packets from received packets.

The $K-$dimensional measurement vector $\mathbf{r}$ (i.e., $\mathbf{Hx}$) is referred to as the precoded version of $\mathbf{x}$. We packetize $\mathbf{r}$ into n packets, each with P samples for transmission. Now if the communication channel were to drop some packets, we can make use of the underlying transform-domain sparsity of the signal, and still be able to reconstruct the signal from the correctly received precoded data. Note that the reconstruction fidelity will depend upon the packet loss rate and the signal sparsity structure. Here, we set K equal to N. As the sender requires N samples per transmission, it will introduce a constant latency of $\frac{N}{f_s}$ seconds.

We now consider reconstruction of $\hat{\mathbf{x}}$ at the receiver. Let $\mathbf{H}_c$, a diagonal matrix of dimension $K \times K$ represent the channel and S be the set containing indices of lost packets. The cardinality of the set S represents the number of packets dropped. The elements in the diagonal of $\mathbf{H}_c$ are defined as follows.

$$[\mathbf{H}_c]_{i,i} = \begin{cases} 0 & \text{if } \lceil \frac{i}{P} \rceil \in S \\ 1 & \text{otherwise} \end{cases} \tag{4}$$

The pre-multiplication of $\mathbf{H}_c$ with $\mathbf{r}$ will essentially provide us the precoded data samples that were successfully received. We denote the resulting vector as $\hat{\mathbf{r}}$ and it is given by $\mathbf{H}_c\mathbf{H}\mathbf{x}$. The MP algorithm from [13] is used to generate the signal estimate $\hat{\mathbf{x}}$ from received data $\hat{\mathbf{r}}$. Note that the term $\mathbf{H}$ in the MP algorithm needs to be replaced by $\mathbf{H}_c\mathbf{H}$ in this case. The Gabor basis ($\mathbf{W}$) from eq. (1) is used to enforce sparsity during the CS reconstruction. Figs. 1(a) and 1(b) illustrate

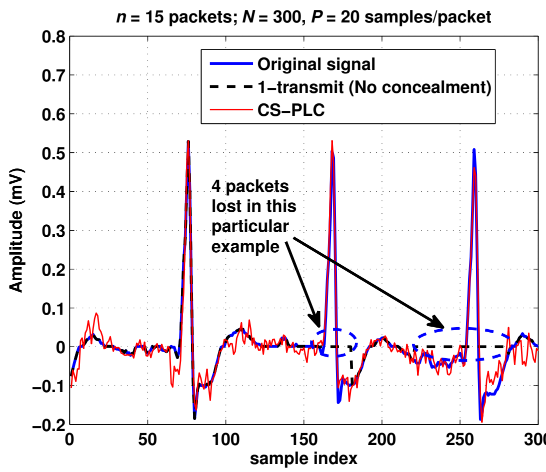

Fig. 2. Signal reconstruction example using CS-PLC and 1-transmit (with no concealment).

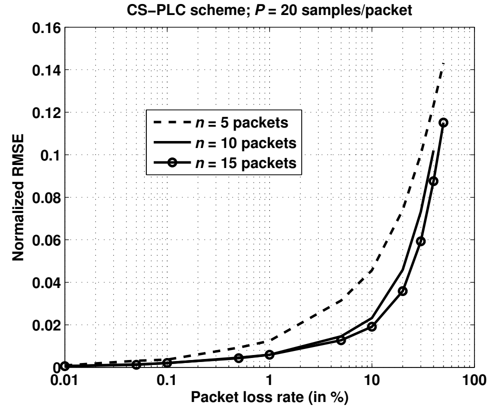

Fig. 3. RMSE comparisons of CS-PLC scheme for $n = 5, 10$ and 15 packets (which correspond to latencies of 0.8, 1.6 and 2.4 seconds respectively).

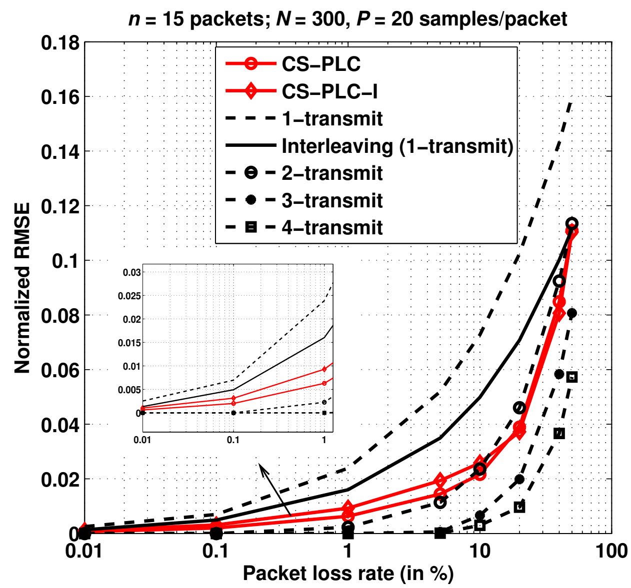

Fig. 4. Comparisons for various PLC schemes in terms of normalized RMSE.

an example with a short segment of ECG signal ($\mathbf{x}$) and the corresponding representation $\mathbf{y}$ (i.e., $\mathbf{Wx}$) in transform space respectively. Fig. 1(c) shows the precoded version of $\mathbf{x}$ given by $\mathbf{r} = \mathbf{Hx}$.

Fig. 2 provides an illustration of the CS-PLC reconstruction described above. The solid curves in blue and red represent the original and reconstructed signals respectively. The dashed curve in black represents $k-$transmit, where $k = 1$ (i.e., single transmission) with zeros substituted for lost packets. In this example, two ECG peaks were lost due to packet losses and the CS-PLC was able to estimate ECG signal peaks with high fidelity.

We also tried an alternative scheme, wherein the precoding is done for every P ECG samples instead of $n{\cdot}P$ samples. This is followed by sample level interleaving of the precoded data across the length of $n{\cdot}P$ samples. We refer to this as CS-PLC-I. Although we precode the original ECG data over shorter durations as compared to CS-PLC, the interleaving step allows to spread the signal information across a longer duration. CS-PLC-I further reduces the complexity at the sender and increases its longevity. We also compare proposed approaches with "retransmissions". Retransmissions has smaller bandwidth penalty compared with FEC, but considerable complexity at the sensor as packets will have to be buffered at the sender. There is also a latency penalty that is proportional to the round trip time (RTT).

IV. Results

In this section, we present quantitative results of the proposed approach for various packet loss conditions. The performance of the PLC schemes are first evaluated in terms of normalized RMSE which is defined as $\frac{\sqrt{E[\|\mathbf{x}-\hat{\mathbf{x}}\|^2]}}{\max\{|\mathbf{x}|\}}$. The term $E[\cdot]$ denotes the expectation operator with Monte Carlo averaging over various realizations of ECG signals ($\mathbf{x}$) and different channel realizations ($\mathbf{H}_c$). All the comparisons presented in this paper are averaged over 20000 monte-carlo channel realizations. The ECG signals in this study are taken from the MIMIC database [14] discussed earlier and the sampling rate is $f_s = 125$Hz. Fig. 3 compares the normalized RMSE performance of CS-PLC for different values of n. The value of P is set to 20 samples (per packet) for all the comparisons in the remainder of the paper. The values of n considered in this comparison 5, 10 and 20 correspond to latencies of 0.8, 1.6 and 2.4 seconds respectively. We observe from fig. 3 that normalized RMSE increases with increasing packet loss rate. This is expected because higher packet loss rates imply reduced amount of reliable data available at the receiver for estimating $\hat{\mathbf{x}}$. Also note that normalized RMSE performance improves with increasing number of packets (n). With higher n we can enforce sparsity over a longer ECG signal duration and thus improve the reconstruction fidelity.

We now compare CS-PLC, CS-PLC-I and retransmissions

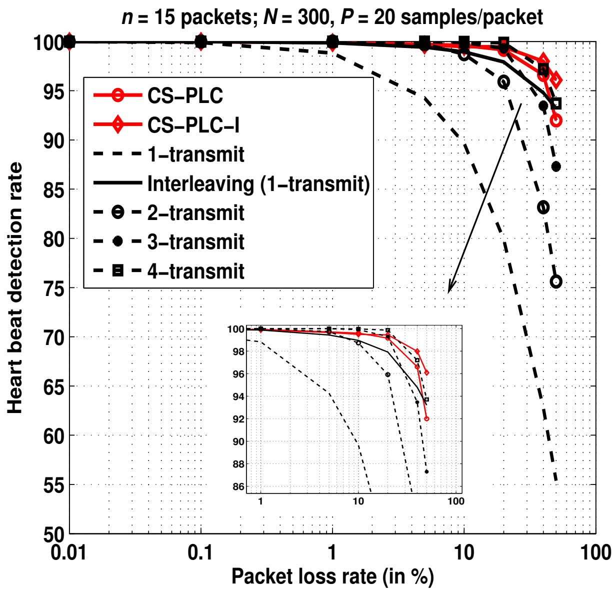

Fig. 5. Heart beat detection performance comparisons for various PLC schemes.

approach in fig. 4. The red curves with "circle" and "square" markers represent the CS-PLC and CS-PLC-I schemes respectively. The value of n is set to 15 packets for the comparisons. Note that the reconstruction fidelity obtained using CS-PLC-I is comparable to the CS-PLC method. The dashed curve in fig. 4 represents the $1-$transmit scheme. Note that at moderate packet loss rate of 1%, 1-transmit method performs 5 times worse, in terms of RMSE, compared with CS-PLC scheme. Next, we employ sample level interleaving (as explained before) in the 1-transmit scheme and quantify the associated RMSE improvement. We observe that with interleaving 1-transmit method performs only 3 times worse as compared to CS-PLC for packet loss rate of 1%. The dashed curves with different markers depict RMSE performance for different values of k in the $k-$transmit approach with no interleaving. Note that with two and three retransmissions we do achieve significant improvement in reconstruction RMSE, however this comes at the cost of increased transmission bandwidth, end-to-end system latency and higher protocol complexity at the sensor.

Next, we compare all the PLC schemes presented above with respect to heart beat detection accuracy. This quantity is defined as the rate of correctly identifying the QRS peaks in the ECG signal. The value of 100% indicates perfect beat detection whereas the value of 0% indicates no beat detection. According to the AAMI standards in [15] the beat is correctly detected if it lies within 150 ms of the annotated beat index available beforehand from the database. The details about the monte-carlo simulations are same as in fig. 4. We make several observations from this data. First that heart beat detection rate degrades with increasing packet loss rate as expected. With single transmission (and no interleaving) the detection accuracy obtained at packet loss rate of 50% is equal to 55%. Second the performance improves for the $k-$transmit schemes with increasing k. Third observation is that the proposed CS based PLC schemes perform superior than $3-$transmit method even at very high packet loss rates. Note that at packet loss rate of 50%, CS-PLC-I achieves 96% detection accuracy as opposed to the $3-$transmit method which achieves 87% detection accuracy only.

V. Conclusions

Packet losses can have extremely adverse impact on many healthcare applications relying on BAN and WAN technologies We present CS based packet loss concealment solution from an end-to-end application perspective. Our approach has lower transmission overhead compared with FEC. It is also a simpler scheme compared to retransmissions and has a constant latency. We show that the reconstruction accuracy degrades gracefully as packet loss rate increases. We demonstrated that ECG signals can be recovered with high fideliety, even in the presence high packet loss rate conditions. Our simulations based on the ECG data from MIMIC demonstrate that up to 99% beat-detection accuracy at packet loss rates of 20%, with a constant latency of less than 2.5 seconds.

References

[1] 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.

[2] J. Hou, B. Chang, D-K Cho, and M. Gerla. Minimizing 802.11 interference on zigbee medical sensors. *to appear in BodyNets 2009*.

[3] C. Chigan and V. Oberoi. Providing QoS in ubiquitous telemedicine networks. In *Fourth Annual IEEE International Conference on Pervasive Computing and Communications Workshops*, March 2006.

[4] D. J. Vergados, D. D. Vergados, and I. Maglogiannis. Applying wireless diffserv for QoS provisioning in mobile emergency telemedicine. In *IEEE Global Telecommunications Conference (GLOBECOM)*, pages 1–5, November 2006.

[5] T. Sakamoto, Wei Daming, and K. Hamao. A study of ecg telemonitoring system using mobile internet connections. In *The Fourth International Conference on Computer and Information Technology (CIT)*, pages 593–597, September 2004.

[6] X. Ge, D. Lai, and Z. Fang. A real-time continuous ecg transmitting method through gprs with low power consumption. In *The 2nd International Conference on Bioinformatics and Biomedical Engineering (ICBBE)*, pages 556–559, May 2008.

[7] D. Donoho. Compressed sensing. *IEEE Transactions on Biomedical Engineering*, 52:1289–1306, April 2006.

[8] E. Candes, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59:1207–1223, August 2006.

[9] *IEEE Signal Processing Magazine [Sensing, Sampling, and Compression]*, volume 25. IEEE Signal Processing Magazine, March 2008.

[10] P. K. Baheti and H. Garudadri. Heart rate and blood pressure estimation from compressively sensed photoplethysmograph. *to appear in BodyNets 2009*.

[11] P. K. Baheti and H. Garudadri. An ultra low power pulse oximeter sensor based on compressed sensing. *to appear in BSN 2009*.

[12] 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.

[13] J. A. Tropp, M. B. Wakin, M. F. Duarte, D. Baron, and R. G. Baraniuk. Random filters for compressive sampling and reconstruction. In *Proc. Int. Conf. on Acoustics, Speech, and Signal Proc. (ICASSP)*, May 2006.

[14] G. B. Moody and R. G. Mark. A database to support development and evaluation of intelligent intensive care monitoring. In *Proc. Computers in Cardiology*, pages 657–660, Indianapolis, IN, September 1996.

[15] American National Standard. *Cardiac monitors, heart rate meters, and alarms*. ANSI/AAMI EC13:2002/(R)2007, Arlington (VA), 2007.