Improving the Energy Efficiency of Continuous Health Monitoring Devices

Pramod. R. Bokde
Department of Electronics Engineering
Priyadarshini Bhagwati College of Engineering, Nagpur, India

Abstract— In This research, we evaluate the energy consumption and storage requirement for the baseline scheme using the proposed models for the health monitoring devices. Since each sensor has its own sampling rate and resolution, its energy consumption differs from that of others. They correspond to the minimum and maximum sampling rates, respectively. The minimum/maximum battery lifetimes are reported assuming that each sensor node uses a regular coin cell battery (CR2032). A regular coin cell battery is commonly used in biomedical sensor. Not surprisingly, ECG and EEG sensors are seen to consume the most amount of energy. Thus, these sensors are the main obstacles to providing long-term health monitoring.

Key words: ECG, EEG, Energy Efficiency

I. Introduction

In this paper, we first propose three schemes for signal processing and transmission that can be used in a WBAN. Then, we evaluate and compare these schemes from the energy perspective. We divide the sensors into two different categories based on their transmission rate: low-sample-rate sensors (heart rate, blood pressure, oxygen saturation, temperature, blood sugar, accelerometer) and high-sample-rate sensors (EEG and ECG). Then, we use the following three schemes to reduce the energy consumption of each node.

- We accumulate multiple samples in one packet before transmitting the raw data in order to decrease the number of transmitted packets. The base station is responsible for processing and storage of the raw data. This approach is applicable to both high-sample-rate and low-sample-rate sensors.
- We process the data in high-sample-rate sensors (EEG and ECG) using traditional signal processing methods. Then, we transfer just a fraction of the raw data from the sensor node for storage in the base station based on the result of computation.
- We suggest using CS-based computation in high-sample-rate sensor nodes before data transmission. Again, we just transfer a small fraction of the raw data from the sensor node for storage in the base station based on the result of on-sensor computation.

Although on-sensor computation leads to some extra computational energy consumption, it reduces transmission energy consumption significantly due to the reduction in the amount of data transmitted. This is especially true when the transmission rate of a sensor is very high and important events (e.g., seizure, heart attack) are rare. However, in the case of low-sample-rate sensors, the decrease in transmission energy does not offset the increase in computational energy. Therefore, we do not employ any on-sensor computation for low-sample-rate sensors. We estimated the minimum/maximum energy consumption of each sensor in different scenarios, and based on that, we computed the minimum/maximum battery lifetime.

II. Sample Aggregation

In practice, we do not usually need to transmit the data as fast as we gather them. Thus, we could first accumulate multiple samples (up to 20 B) in one packet and only then transmit the packet. The total number of bits transmitted remains the same. However, the average number of transmitted packets per second is reduced due to the accumulation. The number of samples that can be accumulated in a single packet varies from one device to another based on its resolution. In addition, the data processing algorithm in the base station might have been optimized with a specific number of required samples in mind. Therefore, the number of samples per packet may need to be varied between 1 and the maximum number. For the devices being evaluated, Table 1 shows the maximum number of samples that can be gathered into a single packet.

Sensor	No. of Samples
Heart Rate	16
Blood Pressure	10
Temperature	20
Oxygen Saturation	20
ECG	13
EEG	13
Blood Sugar	10
Accelerometer	13

Table 1 : Maximum number of samples in one packet

In order to calculate the total energy consumption of a sensor, we also need to consider the storage energy required for storing multiple packets before transmission. To store 20 B, which is the maximum number of bytes that can be sent in a single transmission, we consider the energy consumption of a 160-cell buffer. This storage energy remains fixed for the maximum and minimum transmission rates. However, the maximum (minimum) energy consumption is calculated as the energy consumption of transmission using the maximum (minimum) rate plus the energy consumed by the 160-cell buffer. Using the SRAM cell energy reported for the 90 nm technology node], we calculate the minimum and maximum energy consumption of each device, as shown in Table 2. The minimum and maximum battery lifetimes of each sensor are shown in Table 3. Relative to the baseline, this method provides up to 13.58X reduction in maximum energy consumption for low-sample-rate sensors.

The maximum and minimum energy consumptions of high-sample-rate sensors are reduced by 12.98X and 12.83X, respectively.

Sensor	Minimum (J/day)	Maximum (J/day)
Heart Rate	1.50	4.05

Blood Pressure	0.67	38.36
Temperature	0.58	0.95
Oxygen Saturation	0.66	1.26
ECG	52.65	520.24
EEG	52.65	520.24
Blood Sugar	0.62	63.45

Table 2: Minimum and maximum values of total energy consumption while using the sample aggregation Scheme.

Sensor	Minimum (days)	Maximum (days)
Heart Rate	626.26	1650
Blood Pressure	36.95	4216.23
Temperature	2645.20	4125.69
Oxygen Saturation	2030.08	4153.85
ECG	5.15	54.45
EEG	5.15	54.45
Blood Sugar	38.92	4153.85

Table 3: Minimum and maximum battery lifetime of different sensors while using sample aggregation scheme.

III. ANOMALY-DRIVEN TRANSMISSION

We evaluate a process-and-transmit scheme that is more appropriate for high sample- rate sensors (ECG and EEG), which consume significant amounts of energy. If we first process raw data in the sensor nodes themselves and then just transmit some small chunks of data based on the processing results, we can reduce the transmission rate significantly. In this scenario, whenever we detect an abnormal activity, we are required to transmit the raw data corresponding to the abnormal event, in order to facilitate offline evaluation of the data. The computational energy in each sensor node and data transmission rate directly depend on the intended application. We evaluated seizure detection and arrhythmia detection as applications for EEG and ECG sensors, respectively. The traditional computation that we have considered for seizure/arrhythmia detection is as follows. First, we sample the signal at the Nyquist sampling rate. Second, we use a feature extraction algorithm (spectral energy analysis for EEG and Wavelet transform for ECG) to extract the important feature of the signal and build a feature vector. Third, we classify the feature vectors using a binary classifier.

Let us consider an EEG sensor first. We assume signal processing in this sensor is based on a traditional algorithm for seizure detection. The frequency of epileptic seizures varies from person to person. In some cases, seizures may even be separated by years. On the other extreme, seizures might occur every day. Williamson et al. [1] studied 90 patients and reported the mean seizure frequency and mean duration to be 4.7 per month (range: 3 to 9 per month) and 3.8 minutes (range: 1 to 20 minutes), respectively. Based on their result, if the EEG sensor just transmits the small fraction of data corresponding to seizures, the sensor needs to transmit information over a duration of 17.8 minutes per month, on an average. Table 4 shows the average total energy consumption of the EEG sensor when we use the traditional signal processing method described in [2, 3] and only transmit important chunks of data whenever an abnormality is detected.

The minimum (maximum) value corresponds to the minimum (maximum) sampling frequency. In this scheme, the processing module consumes the major part of energy. Relative to the baseline, it provides up to 177X reduction in total energy consumption for the EEG sensor. Table 5 shows the minimum and maximum battery lifetimes of the EEG sensor in this scheme.

Next, we consider ECG sensors, and assume that the signal processing method is the traditional computation method for arrhythmia detection, as discussed in [3]. Unlike seizure, the frequency of occurrence of arrhythmia varies significantly. There are different types of arrhythmia: each may lead to intermittent or consistent symptoms. Therefore, it is difficult to predict the frequency of occurrence for arrhythmia.

Fig. 1 shows the total energy consumption and battery lifetime of the ECG sensor with respect to frequency of occurrence of arrhythmia in a day, respectively. We assume that after detecting an abnormal event, the sensor transmits the information of a standard one-minute ECG strip to the base station.

Sensor	Minimum (J/day)	Maximum (J/day)
EEG	36.27	38.83

Table 4: Average total energy consumption of the EEG sensor for anomaly driven method.

Sensor	Minimum (J/day)	Maximum (J/day)
EEG	69.53	74.44

Table 5: Average battery lifetimes for the EEG sensor for the anomaly –driven method

IV. CS-BASED COMPUTATION AND TRANSMISSION

As the third scheme, we evaluate an approach for computation and data transmission that can reduce the energy consumption of EEG and ECG sensors significantly. As mentioned earlier, since the total energy consumption of EEG and ECG sensors is very high due to their high data transmission rates, if we can process the raw data in these sensors and transmit only small chunks of data upon the occurrence of an abnormal event, the transmission energy may be reduced significantly. However, now the computation energy becomes the major energy bottleneck. Hence, we try to reduce it through CS-based computation. First, we briefly describe CS.

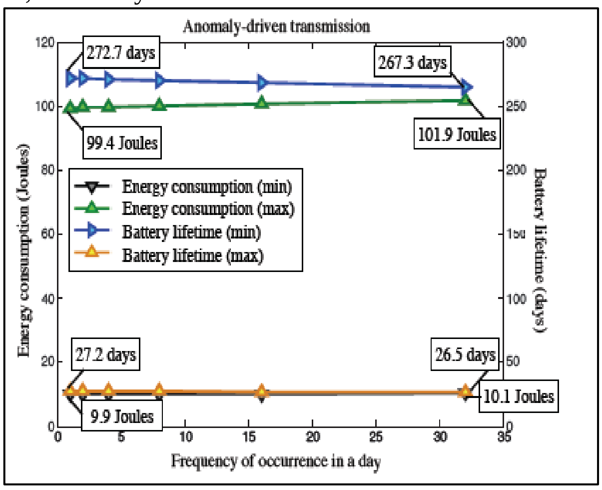

Fig. 1: Energy consumption and battery lifetime of the ECG sensor for the anomaly-driven method with respect to frequency of occurrence of arrhythmia in a day.

CS (also called compressive sampling or sparse sampling) is a signal processing approach for efficiently sampling and reconstructing a signal [4]. The common goal of various signal processing approaches is to reconstruct a signal from a finite number of measurements. Without any prior knowledge or assumptions about the signal, this task is not feasible due to the fact that there is no way to reconstruct an arbitrary signal in an interval in which it is not measured. However, under certain conditions and assumptions, the signal can be reconstructed using a finite number of samples. In the CS approach, a signal can be recovered from far fewer samples than required by Nyquist sampling. Recovering a signal using the CS approach relies on two fundamental principles: sparsity and incoherence.

1) Sparsity: This requires that the signal be sparse in some domain (i.e., the signal's representation in some domain should have many coefficients close to or equal to zero). CS can be used to compress an N-sample signal X that is sparse in a secondary basis. Previous research has shown that ECG and EEG signals are sparse enough in the Wavelet transform space [5] and Gabor space [6], respectively.
2) Incoherence: This indicates that unlike the signal of interest, the sampling/ sensing waveforms have an extremely dense representation in the transformed domain.

The main limitation of the classical CS approach is as follows. Although the signal can be recovered using only a few samples, the traditional signal processing methods are not designed to process the compressed form of the signal. Therefore, the signal needs to be reconstructed before processing by the traditional signal processing methods. Unfortunately, reconstruction of a signal from its compressed representation is an energy-intensive task and cannot be performed on sensors due to their energy constraints. In WBANs, it is often necessary to process the data sampled by the biomedical sensors, e.g., to detect anomalies or compute statistics of interest. In this work, we evaluate a modified version of the classical CS approach that enables ECG and EEG signals to be processed on the sensor without being reconstructed (Fig. 2).

The need for reconstruction can be circumvented by performing signal processing computations directly in the CS domain. Shoaib et al. have developed precisely such a method [2, 3], and demonstrated applications to various biomedical signals. This method reduces the computation energy significantly because much fewer data samples need to be processed. Generally, this method consists of three steps:

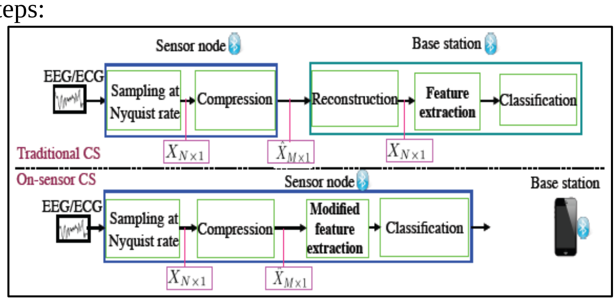

Fig 2: Traditional CS vs. on-sensor CS-based computation.

1) First, we compress the signal of interest using a low-rank random projection matrix. If we can represent the signal (X) as ψ * s, where s is a vector of K-sparse coefficients, a low-rank random matrix Φ can be found to transform X to a set of M samples where O(Klog(N/K)) < M << N. We can then use the following equation for obtaining the compressed samples (denoted by X'): $X'_{Mx1} = \Phi_{MxN} X X_{Nx1}$
2) Second, we generate a feature extraction operation in the CS domain (H') from its equivalent in the Nyquist domain (H) by minimizing the error in the inner product between feature vectors. For any feature extraction method, which can be represented by matrix H, we can derive an equivalent H' matrix in the CS domain [2, 3].
3) Third, we compute $Y' = H' \times X'$ and provide Y' to the classification process.

The compression ratio is given by α = N/M. It denotes the amount of compression obtained by the projection. Because CS leads to a drastic reduction in the number of samples, it has the potential for reducing the energy consumption of various sensors, including biomedical sensors. Direct computation on compressively-sensed data enables classification to be performed on the sensor node with one to two orders of magnitude energy reduction. We exploit this method for long-term continuous health monitoring.

In order to choose a reasonable compression ratio (α), we first need to compare the outcomes of the CS-based method for different compression ratios. Next, we discuss sensitivity (also called recall) and number of false alarms per hour (FA/h) for different compression ratios. Sensitivity represents the true positive rate. It measures the percentage of actual positives that are correctly identified, such as the percentage of seizure conditions that are correctly classified as seizure. FA/h is the number of false positive outcomes in an hour of detection. Such an outcome is an error in classification since a test result indicates the presence of a medical condition that is not actually present.

Fig. 3 shows the sensitivity and FA/h for seizure detection with respect to different compression ratios. A compression ratio α of 8× is seen to maintain sensitivity and FA/h for seizure classification. Moreover, an 8× compression ratio also exhibits similar results for arrhythmia detection [2, 3]. Thus, we assume this ratio for deriving the next set of results.

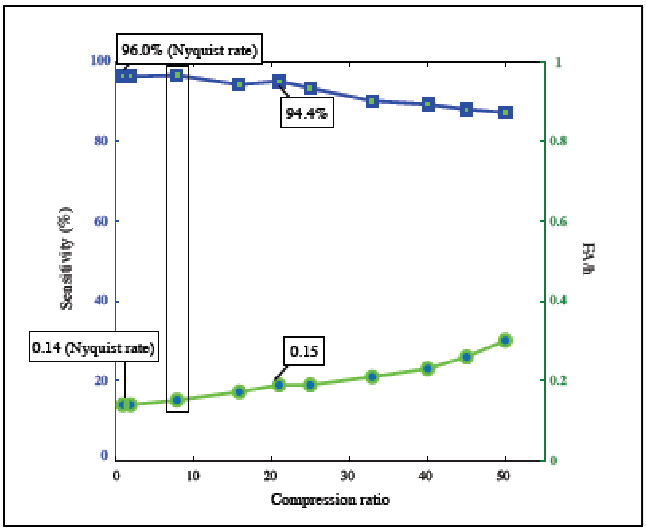

Fig. 3: Sensitivity and FA/h of seizure detection classification with respect to compression ratio.

Next, we examine the EEG sensor in the context of seizure detection. Using the CS-based algorithm for seizure detection, the average value of total energy consumption of the EEG sensor (Table 6) is much less than that of the anomaly-driven signal processing method. Relative to the baseline, the total energy consumption of the EEG sensor is reduced by up to 724× in this scheme. Table 7 shows the battery lifetime of the EEG sensor, which improves by a similar ratio.

Next, we examine an ECG sensor in the context of arrhythmia detection. Fig. 3.6 shows the total energy consumption and battery lifetime of the ECG sensor with respect to the frequency of occurrence of arrhythmia in a day. Similar to the previous scheme, we assumed that after detecting an arrhythmia, the ECG sensor transmits the information of a standard one-minute ECG strip to the base station.

Sensor	Minimum (J/day)	Maximum (J/day)
EEG	6.93	9.50

Table 6: Average total energy consumption of the EEG sensor for CS-based computation.

Sensor	Minimum (J/day)	Maximum (J/day)
EEG	284.43	389.45

Table 7: Average battery lifetimes for the EEG sensor for CS-based c

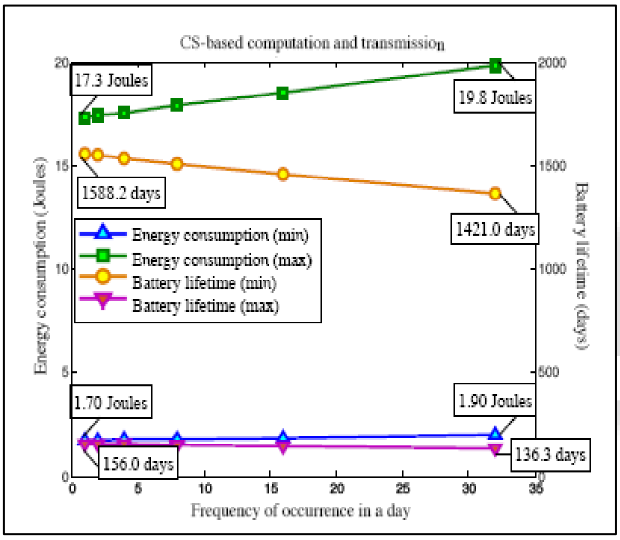

Fig. 4: Energy consumption and battery lifetime of the ECG sensor for the CS-based method with respect to frequency of occurrence of arrhythmia in a day.

V. SUMMARY OF PROPOSED SCHEMES

We summarize the results. Fig. 5 shows the energy reduction in each sensor for the sample aggregation scheme. The energy reduction is an order of magnitude relative to the baseline.

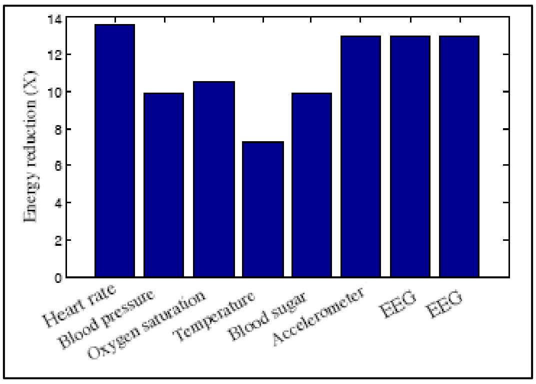

Fig. 5 : Energy reduction in each sensor when the sensor accumulates multiple samples in one packet.

Fig. 6 shows the energy reduction in EEG and ECG sensors when the maximum sampling frequency is employed. The CS-based approach can be seen to result in two to three orders of magnitude energy reduction relative to the baseline.

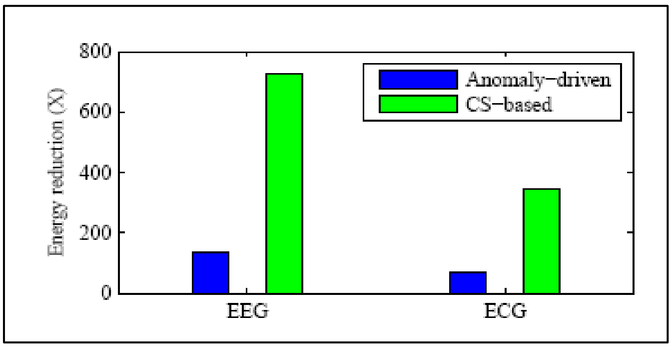

Fig. 6: Energy reduction in EEG and ECG sensors.

REFERENCES

[1] P. D. Williamson, D. D. Spencer, S. S. Spencer, R. A. Novelly, and R. H. Mattson, “Complex partial seizures of frontal lobe origin”, Annals of Neurology, vol. 18, no. 4, pp. 497-504, 1985.

[2] M. Shoaib, K. H. Lee, N. K. Jha, and N. Verma, “A 0.6-107uW energy-scalable processor for seizure detection with compressively-sensed EEG”, IEEE Trans. Circuits and Systems I, vol. 61-I, no. 4, pp. 1105-1118, Apr. 2014.

[3] M. Shoaib, N. K. Jha, and N. Verma, “Signal processing with direct computations on compressively-sensed data”, IEEE Trans. VLSI Systems, vol. 23, no. 1, pp. 30-43, Jan. 2015.

[4] E. J. Candes and M. B. Wakin, “An introduction to compressive sampling”, IEEE Signal Processing Magazine, vol. 25, no. 2, pp. 21-30, 2008.

[5] L. F. Polania, R. E. Carrillo, M. Blanco-Velasco, and K. E. Barner, “Compressed sensing based method for ECG compression”, in Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing, 2011, pp. 761-764.

[6] M. L. Brown, W. J. Williams, and A. O. Hero, “Non-orthogonal Gabor representation of event-related otentials”, in Proc. IEEE Int. Conf. Engineering in Medicine and Biology Society, 1993, pp. 314-315.

[7] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu, “They can hear your heartbeats: Non-invasive security for implantable medical devices”, ACM SIGCOMM Computer Communication Review, vol. 41, no. 4, pp. 2-13, 2011.