

Energy-Efficient Long-term Continuous Personal Health Monitoring

Arsalan Mohsen Nia, Mehran Mozaffari-Kermani, *Member, IEEE*, Susmita Sur-Kolay, *Senior Member, IEEE*, Anand Raghunathan, *Fellow, IEEE*, and Niraj K. Jha, *Fellow, IEEE*

[This paper can be cited as: A. M. Nia, S. Sur-Kolay, A. Raghunathan and N. K. Jha, "Energy-efficient long-term continuous personal health monitoring," in IEEE Trans. Multi-Scale Computing Systems \(TMSCS\), vol. 1, no. 2, pp. 85-98, Oct. 2015](#)

[The latest version of this manuscript is available on http://ieeexplore.ieee.org/document/7303932/](http://ieeexplore.ieee.org/document/7303932/)

Abstract—Continuous health monitoring using wireless body area networks of implantable and wearable medical devices (IWMDs) is envisioned as a transformative approach to healthcare. Rapid advances in biomedical sensors, low-power electronics, and wireless communications have brought this vision to the verge of reality. However, key challenges still remain to be addressed. The constrained sizes of IWMDs imply that they are designed with very limited processing, storage, and battery capacities. Therefore, there is a very strong need for efficiency in data collection, analysis, storage, and communication.

In this paper, we first quantify the energy and storage requirements of a continuous personal health monitoring system that uses eight biomedical sensors: (1) heart rate, (2) blood pressure, (3) oxygen saturation, (4) body temperature, (5) blood glucose, (6) accelerometer, (7) electrocardiogram (ECG), and (8) electroencephalogram (EEG). Our analysis suggests that there exists a significant gap between the energy and storage requirements for long-term continuous monitoring and the capabilities of current devices.

To enable energy-efficient continuous health monitoring, we propose schemes for sample aggregation, anomaly-driven transmission, and compressive sensing to reduce the overheads of wirelessly transmitting, storing, and encrypting/authenticating the data. We evaluate these techniques and demonstrate that they result in two to three orders-of-magnitude improvements in energy and storage requirements, and can help realize the potential of long-term continuous health monitoring.

Index Terms—Body area networks, compressive sensing, continuous health monitoring, implantable and wearable medical devices, secure wireless sensor network.

I. INTRODUCTION

RAPID technological advances in biomedical sensing and signal processing, low-power electronics, and wireless networking are transforming and revolutionizing healthcare. Prevention and early detection of disease are increasingly viewed as critical to promoting wellness rather than just treating illness. In particular, continuous long-term health

monitoring, where various physiological signals are captured, analyzed, and stored for future use, is envisioned as key to enabling a proactive and holistic approach to healthcare.

Several trends in computing and communications technology have converged to advance continuous health monitoring from a distant vision to the verge of practical feasibility. Foremost among these is the evolution of implantable and wearable medical devices (IWMDs). Traditionally, medical monitoring systems, such as ECG and EEG monitors, have been used to simply gather raw data, with signal processing and data analysis being performed offline. However, with the continuing performance and energy efficiency improvements in computing, real-time signal processing has become possible. In the last decade, the number and variety of IWMDs have increased significantly, ranging from simple wearable activity and heart-rate monitors to sophisticated implantable sensors. Moreover, advances in low-power wireless communications enable radios to be integrated into even the most energy- and size-constrained devices. This has led to the possibility of composing IWMDs into wireless body area networks (WBANs) [1], [2].

WBANs are opening up new opportunities for continuous health monitoring and proactive healthcare [3]. A typical WBAN for health monitoring consists of (i) implantable and wearable sensors, which are attached to the body or even implanted under the skin to measure vital signs and body signals, e.g., body temperature, heartbeat, blood pressure, etc. and (ii) external devices (which could be smartphones) that act as base stations to collect, store, display, and analyze the data.

Many recent and ongoing research efforts have addressed the design and deployment of WBANs. The CodeBlue project [4] focused on designing wireless sensor networks for medical applications. It included an ad-hoc network to transmit vital health signs to healthcare providers. Otto et al. [5] designed a system architecture to address various challenges posed by the need for reliable communication within the WBAN, and between the WBAN and a medical server. The MobiHealth project [6] offered an end-to-end mobile health platform for healthcare monitoring. Different sensors, attached to a MobiHealth patient, enabled constant monitoring and transmission of vital signals. They considered security, reliability of communication resources, and quality of service (QoS) guarantees.

Notwithstanding advances in IWMDs and WBANs, some key technical challenges need to be addressed in order to enable long-term continuous health monitoring. Due to size constraints and the inconvenience or infeasibility of bat-

Acknowledgments: This work was supported by NSF under Grant no. CNS-1219570.

Arsalan Mohsen Nia is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA (e-mail: arsalan@princeton.edu).

Mehran Mozaffari-Kermani is with the Department of Electrical and Microelectronic Engineering, Rochester Institute of Technology, Rochester, NY 14623, USA (e-mail: m.mozaffari@rit.edu).

Susmita Sur-Kolay is with the Advanced Computing and Microelectronics Unit, Indian Statistical Institute, Kolkata 700108, India (e-mail: ssk@isical.ac.in).

Anand Raghunathan is with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA (e-mail: raghunathan@purdue.edu).

Niraj K. Jha is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA (e-mail: jha@princeton.edu).

tery replacement, IWMDs need to be highly energy-efficient. IWMDs as well as the external devices that aggregate the monitored data have limited storage capacity. Finally, healthcare applications also impose strict requirements for privacy, security, and reliability [2].

This paper aims to address the challenging question of *whether it is feasible to energy- and storage-efficiently provide long-term continuous health monitoring based on state-of-the-art technology*. In this paper:

- We first discuss the traditionally used sense-and-transmit monitoring scheme to establish a baseline for our analyses. We evaluate a system that consists of eight biomedical sensors: (1) heart rate, (2) blood pressure, (3) oxygen saturation, (4) body temperature, (5) blood glucose, (6) accelerometer, (7) ECG, and (8) EEG.
- We present analytical models that can be used to estimate the energy and storage requirements for these biomedical sensors. Our analysis suggests a significant gap between the energy and storage requirements for long-term continuous monitoring and the capabilities of current devices.
- To address the aforementioned gaps in health monitoring, we propose and evaluate three schemes to reduce the overheads of sensing, storing, and wirelessly transmitting the data:
 - 1) First, we explore a simple scheme based on aggregation of samples to amortize the communication protocol overheads and reduce the number of transmissions.
 - 2) Second, we explore anomaly-driven transmission in which the sensors perform on-sensor signal processing to identify time intervals of interest, and only transmit/store data from these intervals.
 - 3) Finally, we explore the concept of compressive sensing (CS) [7], together with a newly developed approach for computation on compressively-sensed data [8], [9], to drastically reduce energy and storage.
- We demonstrate that the proposed schemes can potentially result in two to three orders-of-magnitude reduction in energy and storage requirements, and therefore may be instrumental in enabling continuous long-term health monitoring.
- We compare all proposed schemes and discuss how a continuous long-term health monitoring system should be configured based on patients' needs and physicians' recommendations.

The rest of the paper is organized as follows. Section II describes different components, which form a WBAN and the communication protocols that can be used to connect them together. Section III describes the baseline continuous health monitoring scheme. Section IV presents our analytical models and an analysis of the energy and storage requirements for the baseline WBAN using these models. Section V describes the proposed schemes that include sample aggregation, anomaly-driven sampling, and CS-based computation, and evaluates their energy impact. Section VI evaluates the impact of the proposed schemes on storage requirements. Section VII compares

different schemes and summarizes the medical considerations in configuration and optimization of different sensors. Finally, Section VIII concludes the paper.

II. DIFFERENT COMPONENTS OF A GENERAL-PURPOSE HEALTH MONITORING SYSTEM

In this section, we first describe two fundamental components that form a medical WBAN, namely biomedical sensors and the base station. Second, we discuss the communication protocols, which can be used to connect them together.

A. Health monitoring with networked wireless biomedical sensors

Biomedical sensors have been used for health monitoring for a long time [10]. They sense electrical, thermal, optical, chemical, and other signals to extract information that are indicative of a patient's health condition. Examples of such sensors include oxygen saturation, glucose, blood pressure, heart rate, ECG, EEG, and several forms of imaging.

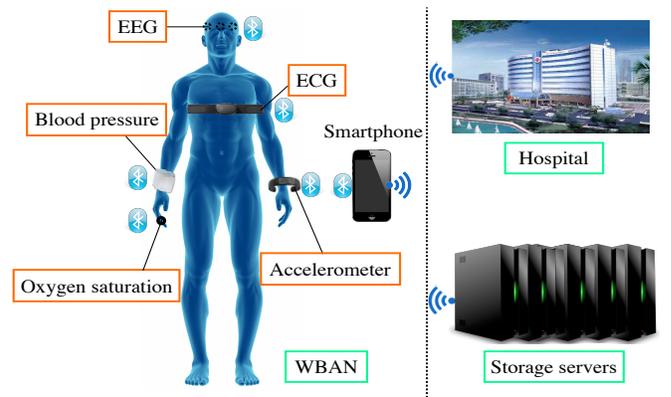


Fig. 1. A personal healthcare system.

In addition to the biomedical sensors, an important component of a WBAN, as shown in Fig. 1, is the base station or hub, a more capable device that aggregates data from the biomedical sensors, visualizes health data for the patient, performs simple analytics, and communicates the health data to remote health providers or health databases. The base station, which could be a customized device or a commodity mobile device such as a smartphone, contains a more capable processor, data storage, and one or more wide-area network interfaces.

B. Communication protocol

A key consideration in the design of a WBAN is the communication technology (radio and protocol) used to connect the medical sensors with the base station. Energy efficiency, security, and interoperability are some of the key factors that must be considered in this context.

Dementyev et al. analyzed the power consumption characteristics of three popular emerging standards – ANT, ZigBee, and BLE – in a duty-cycled sensor node scenario [11]. They

found that BLE achieves the lowest power consumption, followed by ZigBee and ANT. Most of the power consumption differences can be attributed to the time taken for a sensor to reconnect to the base station after waking up and the efficiency of the sleep mode used between transmissions of successive packets. In addition to low power consumption, BLE provides several other advantages for continuous health monitoring:

- 1) Smartphones have become dominant over other forms of base stations for potential use in the health monitoring system. BLE benefits from the widespread use of Bluetooth technology since BLE can be easily integrated into classical Bluetooth circuitry, and almost all new smartphones support BLE.
- 2) BLE is optimized for use in devices that need to communicate small packets wirelessly.
- 3) BLE is optimized to provide a low-rate (< 270 kb/s) wireless data transfer. As shown later, the maximum transmission rate of all sensors is much less than 270kb/s.
- 4) BLE provides a long transmission range (more than 100 meters) that enhances user convenience.
- 5) Due to the privacy and safety concerns in medical systems, security is a key consideration in WBAN design. BLE supports strong encryption (Advanced Encryption Standard) to provide confidentiality as well as per-packet authentication and integrity.

Thus, in our work, we use BLE for short-range transmissions between medical sensors and the base station.

III. BASELINE CONTINUOUS HEALTH MONITORING SYSTEM

In this section, we first describe our baseline WBAN targeted at long-term continuous health monitoring that consists of eight sensors. Then, we discuss its energy and storage requirements.

A. Baseline WBAN

As mentioned earlier, we use eight biomedical sensors in the WBAN. In the baseline WBAN, each sensor node gathers raw data at a specific sampling frequency related to its application. Then, the node generates a BLE packet using a single sample and sends the raw data to the base station for further analysis. In this scheme, each sensor transmits the sample as soon as it is gathered, and the base station is responsible for processing. In order to implement the WBAN, first, it is required to specify the sampling rate for each sensor. This rate must be chosen in such a way that the requirements of different applications are met. The rates vary significantly from one sensor to another. Moreover, the same sensor may need to have different sampling rates in different applications [12]. We have investigated the range of possible sampling rates for each sensor by reviewing the medical literature published between 1997 and 2014. Next, we provide these ranges for various sensors.

- **Heart rate:** The heart rate is commonly sampled at 6-8 Hz frequency. For example, this sampling rate is currently

used in fetal heart rate monitors [13]. While the typical human heart rate is 65-82 beats per minute (bpm), the rate can sometimes exceed 180 bpm. These considerations suggest a sampling rate of 2-8 Hz [14].

- **Blood pressure:** During a typical ambulatory blood pressure monitoring session, the blood pressure is commonly measured every 15 to 30 minutes over a 24-hour period [15]. In some cases (e.g., occurrence of a hemorrhage), the blood pressure should be sampled at a much higher frequency. For example, Adibuzzaman et al. have investigated the use of a blood pressure waveform sampled at 100 Hz to monitor physiological system variations during a hemorrhage [16].
- **Oxygen saturation:** The sampling rate of continuously-monitored oxygen saturation is suggested to be in the 0.001 Hz to 2.00 Hz range [12], [17], [18]. For example, Evans et al. use measurements at 5-min intervals (sampling rate of 0.003Hz) to monitor critically ill, mechanically ventilated adult patients during intrahospital transport [17].
- **Temperature:** The body temperature normally fluctuates over the day. Continuous monitoring of these small fluctuations is suggested by different researchers for a variety of applications [12], [19]. For example, Simon et al. suggest measurements at 10-min intervals to determine the influence of circadian rhythmicity and sleep on 24-hour leptin variations [19]. However, some applications require a higher sampling rate (e.g., 1 Hz) [12]. Thus, we assume the sampling rate of the body temperature sensor to be in the 0.001 Hz to 1 Hz range.
- **Blood sugar:** Blood sugar measurements every 5 to 15 minutes are used in a variety of medical applications [12], [20]. However, some applications, such as continuous glucose monitoring to detect a sudden rise or drop in the glucose level of diabetics, require a higher sampling rate (~ 100 Hz) [12].
- **Accelerometer:** An accelerometer is widely used for physical activity detection. Its sampling rate typically lies in the 30 Hz to 400 Hz range. However, a lower sampling rate (e.g., down to 2 Hz) might be enough for some applications [12], [21]–[23].
- **ECG:** Determining the frequency content of an ECG signal by investigating its frequency spectrum is usually difficult because it is hard to distinguish between frequency components of signal and noise. Berson et al. record over-sampled ECG signals and then apply different low-pass filters to them [24]. They describe the effect of filtering on amplitude variations, concluding that at least a sampling frequency of 50-100 Hz is necessary to prevent amplitude errors. Moreover, Simon et al. demonstrated that a 1000 Hz sampling rate is enough for the majority of ECG-based applications [25]. We consider ECG sampling rates in the 100-1000 Hz range.
- **EEG:** Traditionally, the range of EEG frequencies that was accepted to be clinically relevant was in/below the gamma band (40-100 Hz). However, filtering of the EEG signal at around 70 Hz and using at least a 200 Hz sampling rate are commonly suggested by medical

literature [26]. Moreover, recent studies have shown that EEG signals may also have physiological relevance in high-frequency bands (e.g., 100-500 Hz) [26], [27]. Based on the above discussion, we consider EEG sampling rates in the 100-1000 Hz range.

Next, we consider the sampling resolution of each sensor, where resolution is defined as the number of bits required for representing a sample. We reviewed several recent publications in the biomedical literature to obtain these resolutions.

- **Heart rate:** An accurate and compact low-power heart rate sensor for home-based health care monitoring is described and implemented in [28]. It shows that a resolution of 10 bits is appropriate for providing an accurate measurement of the heart rate.
- **Blood pressure:** We consider 16 bits of resolution for blood pressure samples, which is commonly used in commercial blood pressure monitoring devices [12].
- **Oxygen saturation:** We consider 8 bits of resolution for oxygen saturation based on the data reported in [29], [30].
- **Temperature:** The body temperature varies within the 35 to 42°C range. An 8-bit resolution is sufficient for body temperature sampling.
- **Blood sugar:** Measurements of blood sugar are based on color reflectance. The meter quantifies the color change and generates a numerical value that represents the concentration of glucose. A 16-bit resolution has been shown to be adequate for blood sugar monitoring devices [31].
- **Accelerometer:** We consider 12-bit resolution, which has been used in a variety of wearable accelerometer applications and commercial devices [12], [21], [32].
- **ECG:** Ultra low-power ECG sensors, which are commonly used in long-term monitoring, support 8 or 12 bits of resolution [33]–[35]. A resolution of 8 bits may result in a small but noticeable quantization error. Researchers have shown that greater than 8 bits of resolution will meet ECG requirements [36]. Therefore, we assume a resolution of 12 bits.
- **EEG:** Several low-power wearable EEG sensors [37], [38] use 10- or 12-bit ADC units. The recording should represent samples down to 0.5 μV and up to plus/minus several millivolts. We consider a 12-bit resolution.

Table I summarizes information on sensors, their resolution and sampling rate, and the maximum wireless data transmission rate.

TABLE I
RESOLUTION, SAMPLING RATE, AND MAXIMUM TRANSMISSION RATE

Sensor	resolution (bits/sample)	Sampling rate (Hz)	Maximum transmission rate (bits/s)
Heart rate	10	2-8	80
Blood pressure	16	0.001-100	1600
Oxygen saturation	8	0.001-2	16
Temperature	8	0.001-1	8
Blood sugar	16	0.001-100	1600
Accelerometer	12	2-400	4800
ECG	12	100-1000	12000
EEG	12	100-1000	12000

B. Energy and storage requirements

Next, we discuss energy and storage requirements for a continuous health monitoring system.

Energy consumption can be divided into three categories: sampling, data transmission, and data analysis [39]. Wireless data transmission is usually the major energy-consumer. The available energy in each sensor node is often quite limited. The battery used in the node is typically the largest contributor in terms of both size and weight. Battery lifetime is a very important consideration in biomedical sensors. In particular, battery replacement of implanted sensors may require surgery and, hence, impose cost and health penalties [3]. Thus, biomedical sensors often need to maintain their functionality for months or even years without the need for a battery change. For instance, an implanted pacemaker requires a battery lifetime of at least five years. Furthermore, during communication, biomedical sensors generate heat that may be absorbed in nearby tissue, with possible harmful effects. Hence, the energy consumption should also be minimized from this perspective [3].

Moreover, a WBAN imposes specific storage requirements. Although WBANs facilitate health monitoring and early detection of health problems, physicians usually want access to raw data so that they can independently verify the accuracy of on-sensor processing. Thus, it is important to enable medical personnel to access all or at least important chunks of raw data. However, storing the raw data in the sensor nodes is not feasible for two main reasons. First, IWMD sizes need to be kept small to facilitate patient mobility. Second, adding a large storage to a sensor increases its energy consumption drastically, and as a result, battery lifetime decreases dramatically. Therefore, we may think of storing the data in the base station. However, the base station (e.g., a smartphone) may have its own resource constraints, though much less severe, in terms of storage and battery lifetime. In addition, in order to provide data backup, we may want to periodically send stored data from the base station to storage servers. Therefore, the costs of long-term storage using reliable storage services (e.g., Amazon S3 [40]) should also be considered. Thus, it is important to minimize storage requirements for long-term health monitoring while maintaining adequate information for future reference.

IV. ANALYTICAL MODELS FOR THE EVALUATION OF WBAN'S ENERGY AND STORAGE REQUIREMENTS

In this section, we first describe the analytical models that we use to abstract the essential characteristics of the continuous health monitoring system. Then, we use the model to evaluate the baseline IWMDs.

A. Analytical models

Analytical models can be used to predict system requirements. They are much more efficient than performing simulation. Next, we describe the models used to quantify the energy consumption and storage requirements of the continuous health monitoring system. Table II provides the list of variables used in our models.

TABLE II
VARIABLES, UNIT, AND DESCRIPTION

Variable	Unit	Description
E_{total}	J/day	Total energy consumption of a biomedical sensor
E_s	J/day	Energy consumption of sampling
E_t	J/day	Energy consumption of transmission
E_c	J/day	Energy consumption of computation
E_{ADC}	J/sample	Energy consumption of sampling per sample
f_t	Hz	Transmission frequency
f_s	Hz	Sampling frequency
N	-	Sampling resolution
S	1/day	#samples per day
C	1/day	#transmissions per day
P_{send}	W	Average power consumption in the sending mode
$P_{standby}$	W	Average power consumption in the standby mode
I_{send}	A	Average drained current in the sending mode
$I_{standby}$	A	Average drained current in the standby mode
T_{send}	s	Sending time
$T_{standby}$	s	Standby time
V_{supply}	V	Supply voltage
SR	B/year	Required amount of storage in a year

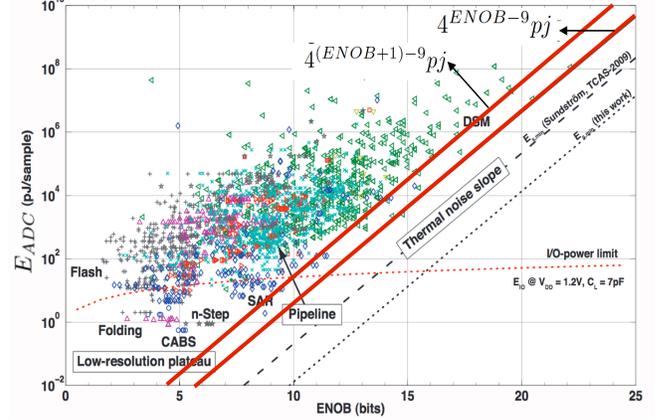


Fig. 2. Scatter plot of the reported E_{ADC} vs. ENOB bits for different ADC architectures: asynchronous (\circ), cyclic (\square), delta-sigma (\triangleleft), flash (+), folding (\triangle), pipeline (\times), successive approximation (\diamond), subranging (\triangleright), n-Slope ($*$), n-Step ($*$), and other (∇) [41].

A.1 Energy consumption

As mentioned earlier in Section III, energy consumption of a sensor has three major components: sampling, transmission, and on-sensor computation. Therefore, we assume that total energy consumption of the sensor (E_{total}) can be written as:

$$E_{total} = E_s + E_t + E_c \quad (1)$$

A.1.1 Sampling energy

Next, we discuss the sampling energy that is consumed by the ADC chip. The total energy consumption of an ADC chip can be divided into: (i) I/O energy, (ii) reference energy, (iii) sample-and-hold energy, (iv) ADC core energy, and (v) input energy [41]. However, separate calculation of these values is difficult. Thus, we use the actual values of the total on-chip ADC energy consumption per sample (E_{ADC}) reported in [41]. It summarizes the experimental results from more than 1400 scientific papers published between 1974 and 2010. Fig. 2 shows the scatter plot of the reported E_{ADC} in each of these papers vs. the effective number of bits (ENOB), where ENOB is defined as:

$$ENOB = \frac{SNR - 1.76}{6.02}, SNR = \frac{P_{signal}}{P_{noise}} \quad (2)$$

ENOB is always less than the resolution for all ADC chips. In particular, for medium-resolution ADCs ($8 \leq N \leq 16$) that are used in biomedical sensors, $ENOB \leq N - 1$ provides a better boundary for the ENOB. For example, Verma et al. presented a low-power 12-bit resolution ADC for WSNs [42]. The ENOB of this ADC is reported to be 10.55 bits.

As shown in Fig. 2, the E_{ADC} of modern medium-resolution ADCs is within the $4^{ENOB-9}pJ$ to $4^{(ENOB+1)-9}pJ$ range. Therefore, the sampling energy consumption per day (E_s) can be upper-bounded as follows:

$$E_s = E_{ADC} * S \quad (3)$$

$$E_{ADC} < 4^{(ENOB+1)-9}pJ \leq 4^{(N-9)}pJ, \quad (4)$$

$$S = f_s \left(\frac{1}{s}\right) * 60\left(\frac{s}{min}\right) * 60\left(\frac{min}{hr}\right) * 24\left(\frac{hr}{day}\right) \quad (5)$$

$$\implies E_s < f_s \left(\frac{1}{s}\right) * 60\left(\frac{s}{min}\right) * 60\left(\frac{min}{hr}\right) * 24\left(\frac{hr}{day}\right) * 4^{(N-9)}pJ \quad (6)$$

Table III shows the upper-bound values of E_s for all the sensors. As discussed later, E_s values for all sensors are negligible in comparison to their total energy consumption. Hence, we can safely assume that $E_{total} \approx E_t + E_c$.

TABLE III
UPPER-BOUND VALUES OF E_s

Sensor	E_s (J/day)
Heart rate	2 e-6
Blood pressure	1 e-1
Oxygen saturation	4 e-8
Temperature	4 e-8
Blood sugar	1 e-1
Accelerometer	2 e-3
ECG	5 e-3
EEG	5 e-3

A.1.2 Transmission energy

In our experiments, we used the Texas Instruments CC2541 Development Kit as the BLE transmission device. To provide a quantitative comparison, we experimentally measured the energy consumption of the transmission chip in a cyclic scenario. In a cyclic transmission, the transmitter takes T_{send} seconds to send the data to the base station and then enters a standby phase for $T_{standby}$ seconds. Hence, the average energy consumption of transmission can be calculated as follows:

$$E_t = (T_{send} * P_{send} + T_{standby} * P_{standby}) * C \quad (7)$$

$$C = f_t \left(\frac{1}{s}\right) * 60\left(\frac{s}{min}\right) * 60\left(\frac{min}{hr}\right) * 24\left(\frac{hr}{day}\right) \quad (8)$$

T_{send} is a fixed value and measured as 2.6 milliseconds for a single packet transmission. $T_{standby}$ depends on the

transmission frequency (f_t):

$$T_{standby} = \frac{1}{f_t} - T_{send} \quad (9)$$

P_{send} and $P_{standby}$ can be obtained by measuring the current drained from the battery with supply voltage V_{supply} . We calculated the average power consumption for a single packet transmission using a standard oscilloscope. P_{send} and $P_{standby}$ were found to be 30.5 mW and 2.5 μ W, respectively, where V_{supply} is set to 2.5 V. In order to measure the power consumption of a single packet transmission, we also considered different packet sizes (varying from 1B to 20B). Our experimental results show that variations in transmission energy of a single packet are negligible when the packet size changes from 1B to 20B. However, since $P_{send} \gg P_{standby}$, a higher transmission rate obviously leads to a higher energy consumption.

A.1.3 Computation energy

Computation energy varies significantly from one biomedical application to another. In most applications, the computation energy can be divided into feature extraction energy and classification energy. Since a feature extraction function can be converted into matrix form, the feature extraction energy can be estimated as the energy consumption of a matrix multiplication function. The classification energy can be estimated based on the reported values of classification energy per vector for various methods. However, obtaining a general model for computation energy is difficult because of its dependence on the application. In this work, when we consider on-sensor computation energy, we use the values reported in [8], [9].

A.2 Storage requirement

Next, we provide an analytical model for estimating the amount of required storage for one-year storage of raw medical data. When there is no on-sensor computation, this only depends on the sampling frequency (f_s) and sampling resolution (N):

$$SR = f_s \left(\frac{1}{s}\right) * N(bits) * \left(\frac{1B}{8bits}\right) * 31536000\left(\frac{s}{year}\right) \quad (10)$$

However, simple on-sensor computation can significantly decrease the amount of required storage. For example, if the computation method can efficiently detect points of interest from the raw data, we may only need to store those specific points for further analysis. Moreover, on-sensor data compression (e.g., in CS-based applications) can also decrease the number of transmitted bits from the sensor to the base station by compressing the raw data before transmission.

B. Evaluation of the baseline WBAN

Next, we evaluate the energy consumption and storage requirement for the baseline scheme, described in Section III, using the models described above.

B.1 Evaluation of the energy consumption

Since each sensor has its own sampling rate and resolution, its energy consumption differs from that of others. Table IV shows the minimum and maximum amounts of energy consumption for different devices in this baseline scenario. They correspond to the minimum and maximum sampling rates, respectively. Table V shows the battery lifetime of each sensor. 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 sensors. 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.

TABLE IV
MINIMUM AND MAXIMUM VALUES OF TOTAL ENERGY CONSUMPTION

Sensor	Minimum (J/day)	Maximum (J/day)
Heart rate	13.99	55.23
Blood pressure	0.26	686.88
Oxygen saturation	0.26	14.00
Temperature	0.26	7.13
Blood sugar	0.26	686.88
Accelerometer	14.00	2747.52
ECG	686.88	6868.80
EEG	686.88	6868.80

TABLE V
MINIMUM AND MAXIMUM BATTERY LIFETIMES OF DIFFERENT SENSORS

Sensor	Minimum (days)	Maximum (days)
Heart rate	48.8	192.90
Blood pressure	3.93	10125.69
Oxygen saturation	192.86	10125.69
Temperature	378.68	10125.69
Blood sugar	3.93	10125.69
Accelerometer	0.98	192.86
ECG	0.39	3.93
EEG	0.39	3.93

B.2 Evaluation of the storage requirement

Next, we evaluate the baseline system from the storage perspective. We readily realize the baseline transmission scheme requires a significant amount of storage. Table VI shows the minimum and maximum amounts of storage required for long-term health monitoring in this system. The minimum (maximum) value corresponds to the minimum (maximum) sampling frequency. Since EEG and ECG signals require the largest amount of storage, we mainly target these signals for storage reduction.

V. IMPROVING THE ENERGY EFFICIENCY OF CONTINUOUS HEALTH MONITORING

In this section, 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

TABLE VI
MINIMUM AND MAXIMUM STORAGE REQUIRED FOR LONG-TERM STORAGE

Sensor	Minimum (MB/yr)	Maximum (GB/yr)
Heart rate	75.18	0.29
Blood pressure	0.07	5.87
Oxygen saturation	0.03	0.06
Temperature	0.03	0.03
Blood sugar	0.07	5.87
Accelerometer	90.23	17.62
ECG	4511.26	44.06
EEG	4511.26	44.06

(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.

Each scheme is discussed in the following subsections and compared against the baseline scheme. We estimated the minimum/maximum energy consumption of each sensor in different scenarios, and based on that, we computed the minimum/maximum battery lifetime.

A. 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 20B) 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 VII shows the maximum number of samples that can be gathered into a single packet.

TABLE VII
MAXIMUM NUMBER OF SAMPLES IN ONE PACKET

Sensor	#Samples
Heart rate	16
Blood pressure	10
Oxygen saturation	20
Temperature	20
Blood sugar	10
Accelerometer	13
ECG	13
EEG	13

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 20B, 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 90nm technology node in [43], we calculate the minimum and maximum energy consumption of each device, as shown in Table VIII. The minimum and maximum battery lifetimes of each sensor are shown in Table IX. Relative to the baseline, this method provides up to 13.58 \times reduction in maximum energy consumption for low-sample-rate sensors. The maximum and minimum energy consumptions of high-sample-rate sensors are reduced by 12.98 \times and 12.83 \times , respectively.

TABLE VIII
MINIMUM AND MAXIMUM VALUES OF TOTAL ENERGY CONSUMPTION WHILE USING THE SAMPLE AGGREGATION SCHEME

Sensor	Minimum (J/day)	Maximum (J/day)
Heart rate	1.50	4.07
Blood pressure	0.65	69.38
Oxygen saturation	0.65	1.33
Temperature	0.64	0.98
Blood sugar	0.65	69.38
Accelerometer	1.70	212.13
ECG	53.52	529.36
EEG	53.52	529.36

B. Anomaly-driven transmission

Next, 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

TABLE IX
MINIMUM AND MAXIMUM BATTERY LIFETIMES OF DIFFERENT SENSORS
WHILE USING SAMPLE AGGREGATION SCHEME

Sensor	Minimum (days)	Maximum (days)
Heart rate	663.39	1800
Blood pressure	38.92	4153.85
Oxygen saturation	2030.08	4153.85
Temperature	2715.10	4218.75
Blood Sugar	38.92	4153.85
Accelerometer	12.73	1588.24
ECG	5.10	50.45
EEG	5.10	50.45

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 [8], [9], [44]–[46].

Let us consider an EEG sensor first. We assume signal processing in this sensor is based on a traditional algorithm for seizure detection, as described in [8], [9]. The frequency of epileptic seizures varies from person to person. In some cases, seizures may even be separated by years. Williamson et al. [47] 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 X shows the average total energy consumption of the EEG sensor when we use the traditional signal processing method described in [8], [9] 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 177 \times reduction in total energy consumption for the EEG sensor. Table XI 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 [9]. 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. 3 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.

TABLE X
AVERAGE TOTAL ENERGY CONSUMPTION OF THE EEG SENSOR FOR THE
ANOMALY-DRIVEN METHOD

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

TABLE XI
AVERAGE BATTERY LIFETIMES FOR THE EEG SENSOR FOR THE
ANOMALY-DRIVEN METHOD

Sensor	Minimum (days)	Maximum (days)
EEG	69.53	74.44

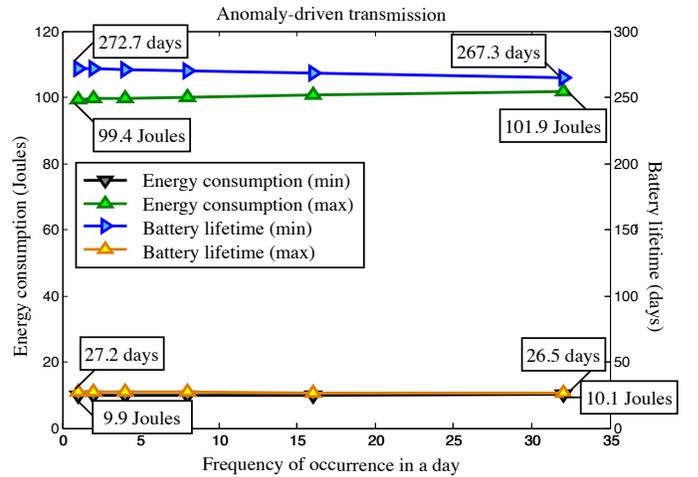


Fig. 3. Energy consumption and battery lifetime of the ECG sensor for the anomaly-driven method with respect to frequency of occurrence of arrhythmia in a day.

C. 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.

CS (also called compressive sampling or sparse sampling) is a signal processing approach for efficiently sampling and reconstructing a signal [7]. 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 Ψ . Previous research has shown that ECG and EEG signals are sparse enough in the Wavelet transform space [48] and Gabor space [49]–[51], 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. 4). 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 [8], [9], 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:

- 1) First, we compress the signal of interest using a low-rank random projection matrix. If we can represent the signal (X) as $\Psi * s$, 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 $O(K \log(N/K)) < M \ll N$. We can then use the following equation for obtaining the compressed samples (denoted by \hat{X}):

$$\hat{X}_{M \times 1} = \Phi_{M \times N} \times X_{N \times 1} \quad (11)$$

- 2) Second, we generate a feature extraction operation in the CS domain (\hat{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 \hat{H} matrix in the CS domain [8], [9].
- 3) Third, we compute $\hat{Y} = \hat{H} \times \hat{X}$ and provide \hat{Y} to the classification process.

The compression ratio is given by $\alpha = N/M$. 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. 5 shows the sensitivity and FA/h for seizure detection with respect to different compression ratios. A compression ratio α of $8\times$ is seen to maintain sensitivity and FA/h for seizure classification. Moreover, an $8\times$ compression ratio also exhibits similar results for arrhythmia detection [8], [9]. Thus, we assume this ratio for deriving the next set of results.

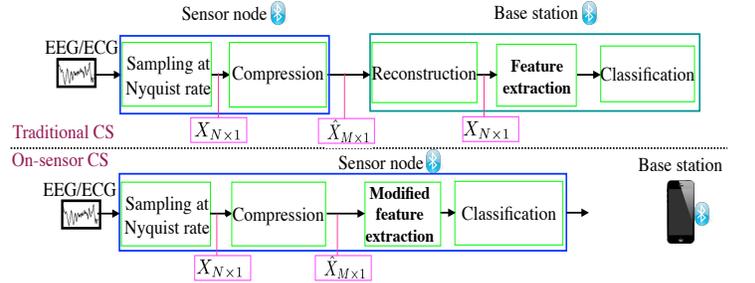


Fig. 4. Traditional CS vs. on-sensor CS-based computation proposed by Shoaib et al. [8], [9].

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 XII) is much less than that of the anomaly-driven signal processing method (Table X). Relative to the baseline, the total energy consumption of the EEG sensor is reduced by up to $724\times$ in this scheme. Table XIII 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. 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.

D. Summary of proposed schemes

Next, we summarize the results.

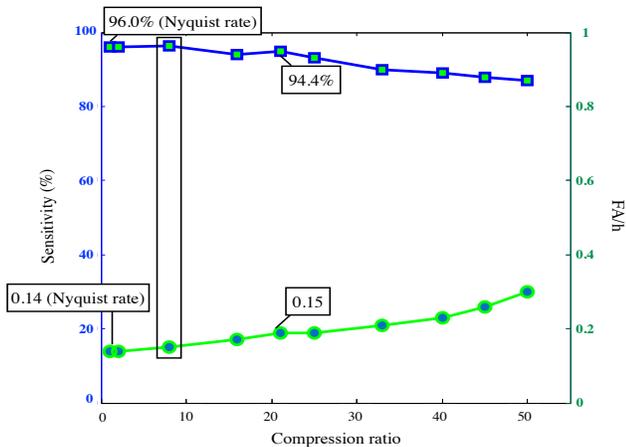


Fig. 5. Sensitivity and FA/h of seizure detection classification with respect to compression ratio. Sensitivity and FA/h CS-based method using $\alpha = 8\times$ are almost equal to the sensitivity and FA/h of the traditional method using Nyquist sampling ($\alpha = 1\times$) [8].

TABLE XII
AVERAGE TOTAL ENERGY CONSUMPTION OF THE EEG SENSOR FOR CS-BASED COMPUTATION

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

Fig. 7 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. 8 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.

VI. STORAGE REQUIREMENTS

We have described three schemes for decreasing the energy consumption of sensors: (i) sample aggregation, (ii) anomaly-driven, and (iii) CS-based computation in the node. The first scheme cannot reduce the amount of required storage because we just accumulate multiple packets in order to reduce the number of transmissions, but we still transmit all the data. However, if we can process the raw data in the sensor nodes and just transmit a chunk of raw data that is essential for future analysis, we would be able to reduce the amount of required storage significantly.

When anomaly-driven or CS-based signal processing is done on the sensor node, the sensor node samples, processes, and then transmits the data based on the result of processing. However, in the case of CS-based computation, the data can be transmitted in compressed form and reconstructed in the base station or server for further processing if needed.

Let us consider EEG sensors first. Based on the results in [47], we assume the mean seizure frequency and mean seizure duration to be 4.7 per month and 3.8 minutes, respectively. Therefore, as mentioned earlier, the EEG sensor needs to transmit information for a duration of 17.8 minutes per month, on an average. Table XIV shows the average amount of storage

TABLE XIII
AVERAGE BATTERY LIFETIMES OF THE EEG SENSOR FOR CS-BASED COMPUTATION

Sensor	Minimum (days)	Maximum (days)
EEG	284.43	389.45

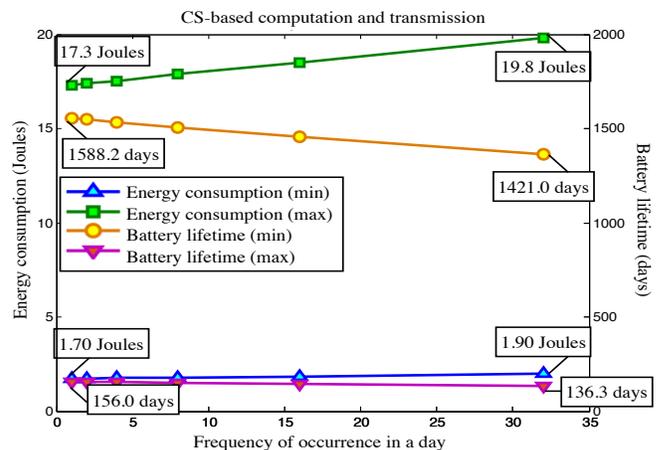


Fig. 6. Energy consumption and battery lifetime of the ECG sensor for the CS-based method with respect to frequency of occurrence of arrhythmia in a day.

required for storing the raw data in the two schemes for seizure detection based on EEG signal analysis. In this table, the minimum (maximum) value corresponds to the minimum (maximum) sampling frequency. The anomaly-driven scheme can be seen to reduce the amount of storage required for storing these signals by $2418\times$. The CS-based scheme provides another $8\times$ reduction on top of this.

As mentioned earlier, unlike seizures, the frequency with which arrhythmia occurs may vary significantly. In order to provide a quantitative analysis for storage requirements in the case of arrhythmia detection, we assume that, after each detection, the sensor transmits the information of a standard one-minute ECG strip to the base station. Fig. 9 shows the amount of required storage for the anomaly-driven and CS-based schemes with respect to the frequency of occurrence. Again, we observe the significant advantage of the CS-based scheme.

TABLE XIV
AVERAGE STORAGE REQUIRED FOR LONG-TERM STORAGE OF PROCESSED DATA

Sensor	Minimum (MB/yr)	Maximum (MB/yr)
EEG (Anomaly)	1.87	18.65
EEG (Compressed)	0.23	2.33

VII. CHOOSING THE APPROPRIATE SCHEME AND HARDWARE PLATFORM

In this section, we first compare the different schemes we presented, and discuss how the appropriate scheme can be chosen for each sensor. Second, we discuss two different types of hardware platforms: application-specific integrated

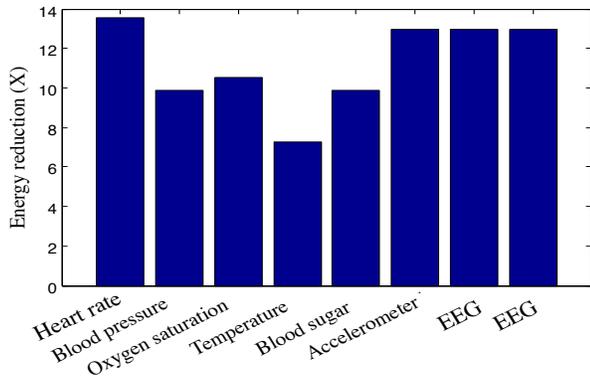


Fig. 7. Energy reduction in each sensor when the sensor accumulates multiple samples in one packet. Raw data are assumed to be gathered at the maximum frequency.

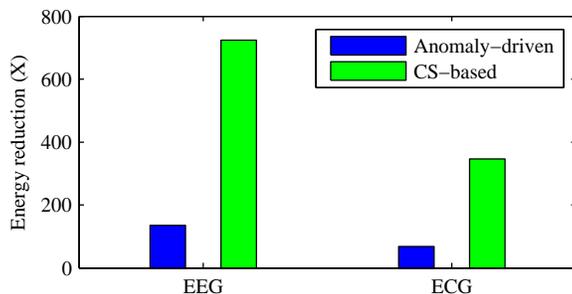


Fig. 8. Energy reduction in EEG and ECG sensors. The number of arrhythmia events in a day is assumed to be 32, and raw data are assumed to be gathered at the maximum frequency.

circuit (ASIC) and general-purpose commercial products. We describe the potential benefits of using ASIC hardware.

A. The appropriate scheme for each sensor

Each scheme has its own advantages and disadvantages. For example, sample aggregation decreases energy consumption at the cost of increased latency. Schemes that use on-sensor computation can significantly increase battery lifetime, however, provide less raw data to the physicians.

Choosing an appropriate scheme for each sensor depends on medical considerations such as tolerable latency and patient's condition. Next, we discuss different considerations that should be taken into account by designers, in addition to the battery lifetime and storage requirement.

- 1) Latency: Latency is the time interval between the occurrence of an anomaly and the response that is provided by medical devices, physicians or medical personnel. Tolerable latency depends on the patient's condition.

- *Example 1:* Consider a continuous health monitoring system that is used to monitor a healthy subject who does not have any history of a serious illness. The system can be configured for this subject to provide routine medical check by collecting and sending medical information to physicians or hospitals at

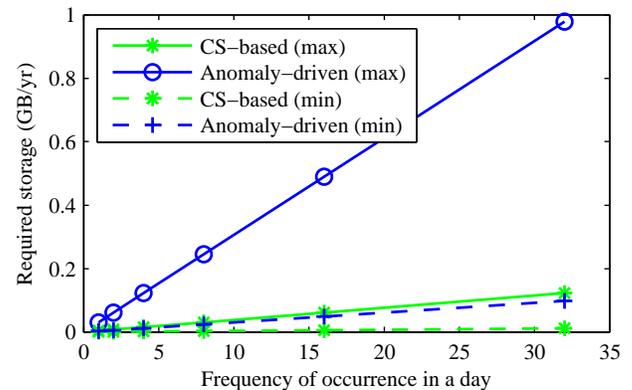


Fig. 9. The amount of storage required for storing important chunks of ECG signals based on the results of computation.

long intervals (e.g., once a day). In fact, latency is not an important factor in this case, and the sensor can be configured to minimize the energy and storage requirements. For example, we can use the CS-based computation method for both EEG and ECG sensors and use the aggregation method for other low-frequency sensors (e.g., temperature) to maximize the battery lifetime of all sensors.

- *Example 2:* Consider a continuous health monitoring system that is used to monitor a subject who has previously been diagnosed with high blood glucose. As a result, any rapid rise in blood glucose should be detected and addressed immediately. In such a scenario, the latency that might be added by using sample aggregation for blood glucose levels may not be acceptable.

Among all the discussed schemes, sample aggregation is the only one that may lead to a non-negligible increase in latency. Therefore, the number of samples that can be aggregated in one packet before transmission can be limited by the latency that can be tolerated.

- 2) Amount of raw medical data transmitted: Physicians may want to examine raw medical data over a specific time period to verify on-sensor computation. The amount of raw information that needs to be transmitted and stored for further analysis varies from one device to another. It also depends on the medical condition of the patient.

Schemes that use on-sensor computation (anomaly-driven transmission and CS-based computation and transmission) only transfer a small portion of raw data containing important information about the occurrence of the anomaly. However, if more medical information is required to be transferred to the base station, the designers should use the other schemes or send more raw data (e.g., over an hour of measurements) after detecting an anomaly.

- 3) Extensibility: This is a design consideration where the implementation takes future modifications of the algorithms into consideration. High extensibility implies that applications of a biomedical sensor can be extended in

TABLE XV
COMPARISON OF DIFFERENT SCHEMES

Scheme	Energy consumption	Required storage	Latency	Amount of raw data transmitted	Extensibility
Baseline	Very high	Very high	Low	All raw data	High
Sample aggregation	Very high	Very high	Varies	All raw data	High
Anomaly-driven	Low	Low	Low	A portion of collected data	Low
CS-based	Very low	Very low	Low	A portion of compressed data	Low

the future with a minimum level of effort. Generally, schemes that rely on on-sensor computation are less extensible in comparison to schemes that transfer raw data to the base station due to the fact that they are designed to minimize the energy consumption and the amount of required storage in certain applications (e.g., arrhythmia detection). Therefore, if a physician wants to change the computation algorithm of the medical device, another device should be designed and used, or at least the device's firmware should be updated each time.

Table XV compares various schemes.

Potentially, different schemes can be used in the health monitoring system for different sensors. Since the sensors are located on different parts of the body, they cannot share on-sensor resources (e.g., the battery). Thus, their battery lifetimes are independent.

We can also use a combination of schemes even in just one sensor. For example, we can combine one of the schemes that uses on-sensor computation (anomaly-driven or CS-based) with the sample aggregation scheme to reduce total energy consumption even more. However, since in anomaly-driven and CS-based schemes, the computation energy is dominant and the transmission energy is only a small fraction of total energy consumption, the addition of the sample aggregation scheme will not provide a significant energy reduction.

B. The hardware platform

An appropriate hardware platform can be chosen from various general-purpose commercial products or else designed as ASIC hardware. General-purpose commercial products enable the designers to implement an algorithm or prototype of a biomedical sensor quickly. However, they are not optimized for the specific application. Anomaly-driven and CS-based schemes use some algorithms to process the raw data on the EEG or ECG sensor nodes before transmission. An ASIC could be designed for these algorithms. In particular, in our computation schemes, the on-sensor computation algorithm uses a support vector machine as a classifier to detect anomalies (arrhythmia and seizure). Specialized processors that enable efficient handling of data structures used by the classifier could reduce computation energy even further. Further energy reduction can be achieved through supply voltage scaling. The total energy is determined primarily by the sum of dynamic (active-switching) energy and the static (leakage) energy. However, reduction in active-switching energy due to supply voltage scaling is opposed by an increase in leakage energy. Therefore, there is an optimal supply voltage at which the circuit attains its minimum energy consumption and still work reliably. This could be addressed in an ASIC. However,

such an ASIC may not be desirable from a cost perspective and does not improve transmission energy.

VIII. CONCLUSION

In this paper, we discussed a secure energy-efficient system for long-term continuous health monitoring. We discussed and evaluated various schemes with the help of eight biomedical sensors that would typically be part of a WBAN. We also evaluated the storage requirements for long-term analysis and storage.

Among the four schemes we evaluated (including the baseline scheme), we showed that the CS-based scheme provides the most computational energy savings (e.g., up to 724× for ECG sensors) because it needs to process much fewer signal samples. For low-sample-rate sensors, we can achieve significant energy savings by simply accumulating the raw data before transmitting them to the base station.

In addition, the CS-based scheme also allows us to reduce the storage requirements significantly. For example, for an EEG sensor based seizure detection application, we achieve total storage savings of up to 19344×. The results indicate that long-term continuous health monitoring is indeed feasible from both energy and storage points of view.

Finally, we compared all proposed schemes and discussed how a continuous long-term health monitoring system should be configured based on patients' needs and physicians' recommendations.

REFERENCES

- [1] S. Ullah, H. Higgins, B. Braem, B. Latre, C. Blondia, I. Moerman, S. Saleem, Z. Rahman, and K. S. Kwak, "A comprehensive survey of wireless body area networks," *J. Medical Systems*, vol. 36, no. 3, pp. 1065–1094, 2012.
- [2] J. Ko, C. Lu, M. B. Srivastava, J. A. Stankovic, A. Terzis, and M. Welsh, "Wireless sensor networks for healthcare," *Proc. IEEE*, vol. 98, no. 11, pp. 1947–1960, 2010.
- [3] B. Latré, B. Braem, I. Moerman, C. Blondia, and P. Demeester, "A survey on wireless body area networks," *Wireless Networks*, vol. 17, no. 1, pp. 1–18, 2011.
- [4] D. Malan, T. Fulford-Jones, M. Welsh, and S. Moulton, "Codeblue: An ad hoc sensor network infrastructure for emergency medical care," in *Proc. Int. Wkshp. Wearable and Implantable Body Sensor Networks*, vol. 5, 2004.
- [5] C. Otto, A. Milenkovic, C. Sanders, and E. Jovanov, "System architecture of a wireless body area sensor network for ubiquitous health monitoring," *J. Mobile Multimedia*, vol. 1, no. 4, pp. 307–326, 2006.
- [6] K. Wac, A. Van Halteren, and D. Konstantas, "QoS-predictions service: Infrastructural support for proactive QoS-and context-aware mobile services (position paper)," in *On the Move to Meaningful Internet Systems*. Springer, 2006, pp. 1924–1933.
- [7] E. J. Candes and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.

- [8] 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. 61-1, no. 4, pp. 1105–1118, Apr. 2014.
- [9] 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.
- [10] C. B. Wilson, "Sensors in medicine," *Western J. Medicine*, vol. 171, no. 5-6, p. 322, 1999.
- [11] A. Dementyev, S. Hodges, S. Taylor, and J. Smith, "Power consumption analysis of Bluetooth Low Energy, ZigBee and ANT sensor nodes in a cyclic sleep scenario," in *Proc. IEEE Wireless Symposium*, 2013, pp. 1–4.
- [12] J. Yick, B. Mukherjee, and D. Ghosal, "Wireless sensor network survey," *Computer Networks*, vol. 52, no. 12, pp. 2292–2330, 2008.
- [13] L. D. Durosier, G. Green, I. Batkin, A. J. Seely, M. G. Ross, B. S. Richardson, and M. G. Frasch, "Sampling rate of heart rate variability impacts the ability to detect acidemia in ovine fetuses near-term," *Frontiers in Pediatrics*, vol. 2, p. 38, 2014.
- [14] V. N. Hegde, R. Deekshit, and P. S. Satyanarayana, "Heart rate variability analysis for abnormality detection using time frequency distribution – smoothed pseudo Winger Ville method," *Power (dB)*, vol. 30, no. 20, p. 10, 2013.
- [15] G. Mancina, A. Zanchetti, E. Agebiti-Rosei, G. Benemio, R. De Cesaris, R. Fogari, A. Pessino, C. Porcellati, A. Salvetti, B. Trimarco *et al.*, "Ambulatory blood pressure is superior to clinic blood pressure in predicting treatment-induced regression of left ventricular hypertrophy," *Circulation*, vol. 95, no. 6, pp. 1464–1470, 1997.
- [16] M. Adibuzzaman, G. C. Kramer, L. Galeotti, S. J. Merrill, D. G. Strauss, and C. G. Scully, "The mixing rate of the arterial blood pressure waveform Markov chain is correlated with shock index during hemorrhage in anesthetized swine," in *Proc. IEEE 36th Annual Int. Conf. Engineering in Medicine and Biology Society*, 2014, pp. 3268–3271.
- [17] A. Evans and E. H. Winslow, "Oxygen saturation and hemodynamic response in critically ill, mechanically ventilated adults during intrahospital transport," *American J. Critical Care*, vol. 4, no. 2, pp. 106–111, 1995.
- [18] C. O. F. Kamlin, C. P. F. O'Donnell, P. G. Davis, and C. J. Morley, "Oxygen saturation in healthy infants immediately after birth," *J. Pediatrics*, vol. 148, no. 5, pp. 585–589, 2006.
- [19] C. Simon, C. Gronfier, J. Schlienger, and G. Brandenberger, "Circadian and ultradian variations of leptin in normal man under continuous enteral nutrition: Relationship to sleep and body temperature," *J. Clinical Endocrinology and Metabolism*, vol. 83, no. 6, pp. 1893–1899, 1998.
- [20] H. G. Piper, J. L. Alexander, A. Shukla, F. Pigula, J. M. Costello, P. C. Laussen, T. Jaksic, and M. S. Agus, "Real-time continuous glucose monitoring in pediatric patients during and after cardiac surgery," *J. Pediatrics*, vol. 118, no. 3, pp. 1176–1184, 2006.
- [21] C. Wan-Young, L. Young-Dong, and J. Sang-Joong, "A wireless sensor network compatible wearable u-healthcare monitoring system using integrated ECG, accelerometer and SpO₂," in *Proc. IEEE 30th Annual Int. Conf. Engineering in Medicine and Biology Society*, 2008, pp. 1529–1532.
- [22] Y. Cho, Y. Nam, Y. Choi, and W. Cho, "SmartBuckle: Human activity recognition using a 3-axis accelerometer and a wearable camera," in *Proc. 2nd Int. Wkshp. Systems and Networking Support for Health Care and Assisted Living Environments*, 2008, p. 7.
- [23] K. Sankaran, M. Zhu, X. F. Guo, A. L. Ananda, M. C. Chan, and L. Peh, "Using mobile phone barometer for low-power transportation context detection," in *Proc. ACM Conf. Embedded Network Sensor Systems*, 2014, pp. 191–205.
- [24] A. S. Berson, F. Y. Lau, J. M. Wojick, and H. V. Pipberger, "Distortions in infant electrocardiograms caused by inadequate high-frequency response," *American Heart Journal*, vol. 93, no. 6, pp. 730–734, 1977.
- [25] F. Simon, J. P. Martinez, P. Laguna, B. Van Grinsven, C. Rutten, and R. Houben, "Impact of sampling rate reduction on automatic ECG delineation," in *Proc. IEEE 29th Annual Int. Conf. Engineering in Medicine and Biology Society*, 2007, pp. 2587–2590.
- [26] J. D. Jirsch, E. Urrestarazu, P. LeVan, A. Olivier, F. Dubeau, and J. Gotman, "High-frequency oscillations during human focal seizures," *Brain*, vol. 129, no. 6, pp. 1593–1608, 2006.
- [27] J. Engel Jr., A. Bragin, R. Staba, and I. Mody, "High-frequency oscillations: What is normal and what is not?" *Epilepsia*, vol. 50, no. 4, pp. 598–604, 2009.
- [28] M. Chen, O. Boric-Lubecke, V. Lubecke, and X. Wang, "Analog signal processing for heart rate extraction," in *Proc. IEEE 27th Annual Int. Conf. Engineering in Medicine and Biology Society*, 2005, pp. 6671–6674.
- [29] M. T. Dastjerdi, "An analog VLSI front end for pulse oximetry," Ph.D. dissertation, Massachusetts Institute of Technology, 2006.
- [30] H. Y. Yang and R. Sarpeshkar, "A bio-inspired ultra-energy-efficient analog-to-digital converter for biomedical applications," *IEEE Trans. Circuits and Systems I*, vol. 53, no. 11, pp. 2349–2356, 2006.
- [31] Glucose Meter Fundamentals and Design. [Online]. Available: freescale.com/files/microcontrollers/doc/app_note/AN4364.pdf
- [32] Low-power 12-bit, 3-axis accelerometer. [Online]. Available: http://cache.freescale.com/files/sensors/doc/fact_sheet/MMA8450QFS.pdf
- [33] C. Park, P. H. Chou, Y. Bai, R. Matthews, and A. Hibbs, "An ultra-wearable, wireless, low power ECG monitoring system," in *Proc. IEEE Conf. Biomedical Circuits and Systems*, 2006, pp. 241–244.
- [34] C. J. Deepu, X. Zhang, W. Liew, D. Wong, and Y. Lian, "An ECG-SoC with 535 nW/channel lossless data compression for wearable sensors," in *Proc. IEEE Asian Conf. Solid-State Circuits*, 2013, pp. 145–148.
- [35] L. Yan and H. Yoo, "A low-power portable ECG touch sensor with two dry metal contact electrodes," *J. Semiconductor Technology and Science*, vol. 10, no. 4, pp. 300–308, 2010.
- [36] M. K. Delano, "A long term wearable electrocardiogram (ECG) measurement system," Ph.D. dissertation, Massachusetts Institute of Technology, 2012.
- [37] D. Yates, E. Lopez-Morillo, R. G. Carvajal, J. Ramirez-Angulo, and E. Rodriguez-Villegas, "A low-voltage low-power front-end for wearable EEG systems," in *Proc. IEEE 29th Annual Int. Conf. Engineering in Medicine and Biology Society*, 2007, pp. 5282–5285.
- [38] M. Mollazadeh, K. Murari, H. Schwerdt, X. Wang, N. Thakor, and G. Cauwenberghs, "Wireless multichannel acquisition of neuromodulators," in *Proc. IEEE Conf. Biomedical Circuits and Systems*, 2008, pp. 49–52.
- [39] I. F. Akyildiz, W. Su, Y. Sankarasubramanian, and E. Cayirci, "A survey on sensor networks," *IEEE Communications Magazine*, vol. 40, no. 8, pp. 102–114, 2002.
- [40] C. Doukas, T. Phiakas, and I. Maglogiannis, "Mobile healthcare information management utilizing cloud computing and Android OS," in *Proc. IEEE Int. Conf. Engineering in Medicine and Biology Society*, 2010, pp. 1037–1040.
- [41] B. E. Jonsson, "An empirical approach to finding energy efficient ADC architectures," in *Proc. IEEE ADC Forum*, 2011, pp. 1–6.
- [42] N. Verma and A. P. Chandrakasan, "An ultra low energy 12-bit rate-resolution scalable SAR ADC for wireless sensor nodes," *IEEE J. Solid-State Circuits*, vol. 42, no. 6, pp. 1196–1205, 2007.
- [43] D. Ho, K. Iniewski, S. Kasnavi, A. Ivanov, and S. Natarajan, "Ultra-low power 90nm 6T SRAM cell for wireless sensor network applications," *Networks*, vol. 1, p. 3, 2006.
- [44] Q. Zhao and L. Zhang, "ECG feature extraction and classification using Wavelet transform and support vector machines," in *Proc. IEEE Int. Conf. Neural Networks and Brain*, vol. 2, 2005, pp. 1089–1092.
- [45] K. H. Lee, S.-Y. Kung, and N. Verma, "Low-energy formulations of support vector machine kernel functions for biomedical sensor applications," *J. Signal Processing Systems*, vol. 69, no. 3, pp. 339–349, 2012.
- [46] A. H. Shoeb, "Application of machine learning to epileptic seizure onset detection and treatment," Ph.D. dissertation, Massachusetts Institute of Technology, 2009.
- [47] 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.
- [48] 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.
- [49] M. L. Brown, W. J. Williams, and A. O. Hero, "Non-orthogonal Gabor representation of event-related potentials," in *Proc. IEEE Int. Conf. Engineering in Medicine and Biology Society*, 1993, pp. 314–315.
- [50] P. J. Durka and K. J. Blinowska, "A unified time-frequency parametrization of EEGs," *IEEE Engineering in Medicine and Biology Magazine*, vol. 20, no. 5, pp. 47–53, 2001.
- [51] S. Aveyente, E. M. Bernat, S. M. Malone, and W. G. Iacono, "Analysis of event related potentials using PCA and matching pursuit on the time-frequency plane," in *Proc. IEEE Int. Conf. Engineering in Medicine and Biology Society*, 2006, pp. 2454–2457.



Arsalan Mohsen Nia received his B.S. degree in Computer Engineering from Sharif University of Technology, Tehran, Iran, in 2012, and M.A. degree in Electrical Engineering from Princeton, NJ, in 2014. He is currently pursuing a Ph.D. degree in Electrical Engineering at Princeton University, NJ. His research interests include wireless sensor networks, Internet of things, computer security, distributed computing, mobile computing, and machine learning.



Mehran Mozaffari-Kermani (M'11) received the B.Sc. degree in electrical and computer engineering from the University of Tehran, Tehran, Iran, in 2005, and the M.E.Sc. and Ph.D. degrees from the Department of Electrical and Computer Engineering, University of Western Ontario, London, ON, Canada, in 2007 and 2011, respectively.

Dr. Mozaffari-Kermani was a recipient of the NSERC Post-Doctoral Research Fellowship in 2011 and the Texas Instruments Faculty Award (Douglas Harvey) in 2014. He currently serves as an Associate

Editor for the IEEE Transactions on Circuits and Systems I and ACM Transactions on Embedded Computing Systems. He also serves as the Guest Editor for the IEEE Transactions on Dependable and Secure Computing for the special issue of Emerging Embedded and Cyber Physical System Security Challenges and Innovations and Guest Editor for the IEEE Transactions on Computational Biology and Bioinformatics for the Special Issue of Emerging Security Trends for Biomedical Computations, Devices, and Infrastructures. He has served as the Lead Guest Editor for the IEEE Transactions on Emerging Topics in Computing for the Special Issue of Emerging Security Trends for Deeply-Embedded Computing Systems in 2014 and 2015.



Anand Raghunathan is a Professor and Chair of VLSI in the School of Electrical and Computer Engineering at Purdue University, where he leads the Integrated Systems Laboratory. His research explores domain-specific architecture, system-on-chip design, embedded systems, and heterogeneous parallel computing. Previously, he was a Senior Research Staff Member at NEC Laboratories America and held the Gopalakrishnan Visiting Chair in the Department of Computer Science and Engineering at the Indian Institute of Technology, Madras. Prof. Raghunathan

has co-authored a book ("High-level Power Analysis and Optimization"), eight book chapters, 21 U.S. patents, and over 200 refereed journal and conference papers. His publications have been recognized with eight best paper awards and four best paper nominations. He received the Patent of the Year Award (recognizing the invention with the highest impact), and two Technology Commercialization Awards from NEC. He was chosen by MIT's Technology Review among the TR35 (top 35 innovators under 35 years, across various disciplines of science and technology) in 2006, for his work on "making mobile secure". Prof. Raghunathan has served on the technical program and organizing committees of several leading conferences and workshops. He has chaired the ACM/IEEE International Symposium on Low Power Electronics and Design, the ACM/IEEE International Conference on Compilers, Architecture, and Synthesis for Embedded Systems, the IEEE VLSI Test Symposium, and the IEEE International Conference on VLSI Design. He has served as Associate Editor of the IEEE Transactions on CAD, IEEE Transactions on VLSI Systems, ACM Transactions on Design Automation of Electronic Systems, IEEE Transactions on Mobile Computing, ACM Transactions on Embedded Computing Systems, IEEE Design & Test of Computers, and the Journal of Low Power Electronics. He was a recipient of the IEEE Meritorious Service Award (2001) and Outstanding Service Award (2004). He is a Fellow of the IEEE, and Golden Core Member of the IEEE Computer Society. Prof. Raghunathan received the B. Tech. degree in Electrical and Electronics Engineering from the Indian Institute of Technology, Madras, and the M.A. and Ph.D. degrees in Electrical Engineering from Princeton University.



Susmita Sur-Kolay (SM05) received the B.Tech. degree in electronics and electrical communication engineering from Indian Institute of Technology, Kharagpur, India, and the Ph.D. degree in Computer Science and Engineering from Jadavpur University, Kolkata, India. She was in the Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA, from 1980 to 1984. She was a Post-Doctoral Fellow in the University of Nebraska-Lincoln, Nebraska-Lincoln, NE, USA, in 1992, a Reader in Jadavpur University from 1993

to 1999, a Visiting Faculty Member with Intel Corporation, Santa Clara, CA, USA, in 2002, and a Visiting Researcher at Princeton University in 2012. She is a Professor in the Advanced Computing and Microelectronics Unit, Indian Statistical Institute, Kolkata. She has co-edited two books, authored a book chapter in the Handbook of Algorithms for VLSI Physical Design Automation, and co-authored about 100 technical articles. Her current research interests include electronic design automation, hardware security, quantum computing, and graph algorithms. Prof. Sur-Kolay was a Distinguished Visitor of the IEEE Computer Society, India. She has been an Associate Editor of the IEEE Transactions on Very Large Scale Integration Systems, and is currently an Associate Editor of ACM Transactions on Embedded Computing Systems. She has served on the technical program committees of several leading conferences, and as the Program Chair of the 2005 International Conference on VLSI Design, the 2007 International Symposium on VLSI Design and Test, and the 2011 IEEE Computer Society Annual Symposium on VLSI. Among other awards, she was a recipient of the President of India Gold Medal from IIT Kharagpur.



Niraj K. Jha (S'85-M'85-SM'93-F'98) received his B.Tech. degree in Electronics and Electrical Communication Engineering from Indian Institute of Technology, Kharagpur, India in 1981, M.S. degree in Electrical Engineering from S.U.N.Y. at Stony Brook, NY in 1982, and Ph.D. degree in Electrical Engineering from University of Illinois at Urbana-Champaign, IL in 1985. He is a Professor of Electrical Engineering at Princeton University. He is a Fellow of IEEE and ACM. He received the Distinguished Alumnus Award from I.I.T., Kharagpur in 2014. He has served as the Editor-in-Chief of IEEE Transactions

on VLSI Systems and an Associate Editor of IEEE Transactions on Circuits and Systems I and II, IEEE Transactions on VLSI Systems, IEEE Transactions on Computer-Aided Design, and Journal of Electronic Testing: Theory and Applications. He is currently serving as an Associate Editor of IEEE Transactions on Computers, Journal of Low Power Electronics, and Journal of Nanotechnology. He has also served as the Program Chairman of the 1992 Workshop on Fault-Tolerant Parallel and Distributed Systems, the 2004 International Conference on Embedded and Ubiquitous Computing, and the 2010 International Conference on VLSI Design. He has served as the Director of the Center for Embedded System-on-a-chip Design funded by New Jersey Commission on Science and Technology. He is the recipient of the AT&T Foundation Award and NEC Preceptorship Award for research excellence, NCR Award for teaching excellence, and Princeton University Graduate Mentoring Award. He has co-authored or co-edited five books titled Testing and Reliable Design of CMOS Circuits (Kluwer, 1990), High-Level Power Analysis and Optimization (Kluwer, 1998), Testing of Digital Systems (Cambridge University Press, 2003), Switching and Finite Automata Theory, 3rd edition (Cambridge University Press, 2009), and Nanoelectronic Circuit Design (Springer, 2010). He has also authored 15 book chapters. He has authored or co-authored more than 400 technical papers. He has coauthored 14 papers, which have won various awards. These include the Best Paper Award at ICCD'93, FTCS'97, ICVLSID'98, DAC'99, PDCS'02, ICVLSID'03, CODES'06, ICCD'09, and CLOUD'10. A paper of his was selected for "The Best of ICCAD: A collection of the best IEEE International Conference on Computer-Aided Design papers of the past 20 years," two papers by IEEE Micro Magazine as one of the top picks from the 2005 and 2007 Computer Architecture conferences, and two others as being among the most influential papers of the last 10 years at IEEE Design Automation and Test in Europe Conference. He has co-authored another six papers that have been nominated for best paper awards. He has received 14 U.S. patents. He has served on the program committees of more than 150 conferences and workshops.