

Body Sensor Networks: A Holistic Approach from Silicon to Users

Benton H. Calhoun¹, John Lach¹, John Stankovic², David D. Wentzloff³, Kamin Whitehouse²,
 Adam T. Barth¹, Jonathan K. Brown³, Qiang Li², Seunghyun Oh³, Nathan E. Roberts³, Yanqing Zhang¹
¹Dept. of Electrical and Computer Engr. ³Dept. of Electrical Engr. and Computer Sci.
²Dept. of Computer Science University of Michigan
 University of Virginia

Abstract—Body sensor networks (BSN) are emerging cyber-physical systems that promise to improve quality of life through improved healthcare, augmented sensing and actuation for the disabled, independent living for the elderly, and reduced healthcare costs. However, the physical nature of BSNs introduces new challenges. The human body is a highly dynamic physical environment that creates constantly changing demands on sensing, actuation, and quality of service. Movement between indoor and outdoor environments and physical movements constantly change the wireless channel characteristics. These dynamic application contexts can also have a dramatic impact on data and resource prioritization. Thus, BSNs must simultaneously deal with rapid changes to both top-down application requirements and bottom-up resource availability. This is made all the more challenging by the wearable nature of BSN devices, which necessitates a vanishingly small size and, therefore, extremely limited hardware resources and power budget. Current research is being performed to develop new principles and techniques for adaptive operation in highly dynamic physical environments, using miniaturized, energy-constrained devices. This paper describes a holistic cross-layer approach that addresses all aspects of the system, from low-level hardware design to higher-level communication and data fusion algorithms, to top-level applications.

I. INTRODUCTION

Body sensor networks (BSN) are emerging cyber-physical systems that have the potential to revolutionize many aspects of life. Physiological parameters of the human body can be used for real-time medical monitoring, and longitudinal data accumulated from many individuals can help diagnose and treat disease. This information can be used to augment bodily functions through drug delivery, augmented sensory stimulation for the deaf or blind, and support for the movement of prosthetic limbs. The greater context of the body such as social interactions and location can also be sensed and fused with physiological data for improved interpretation and actuation. Overall, BSNs promise to improve quality of life

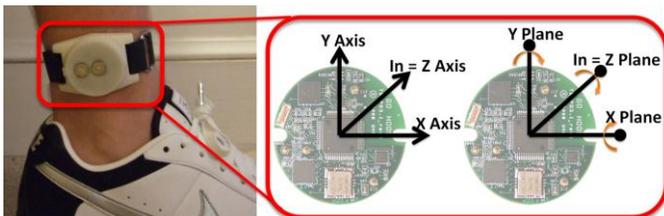


Figure 1. An example of a currently deployed BSN for the sick and elderly is the TEMPO device.

through improved health, augmented sensing and actuation for the disabled, independent living for the elderly, and reduced healthcare costs.

However, the physical nature of BSNs introduces new challenges. The human body is a highly dynamic and unpredictable physical environment that creates constantly changing demands on sensing, actuation, and quality of service (QoS). For example, a BSN application will likely perform different operations when a person walks, sleeps, exercises, or has a medical emergency, and will therefore use different sensors/actuators and will have different requirements on the fidelity, confidence, and latency of data. At the same time, rapid physical movements of the body constantly change the network topology, wireless channel characteristics, and opportunities for energy harvesting. Thus, BSNs must continuously adapt to rapid changes in both top-down application requirements and bottom-up resource availability.

To further complicate matters, the wearable nature of BSN devices necessitates a vanishingly small size, and therefore extremely limited hardware resources and power budget. An example of a currently deployed BSN for the sick and elderly is the TEMPO device shown in Figure 1 [1]. TEMPO is a custom inertial BSN developed at the University of Virginia that provides sensing with six degrees of freedom (three axes of both linear acceleration and rotational rate) and wireless data streaming in the form factor of a wristwatch. With such a device the most common feedback from users is to reduce the form factor and extend the time between battery recharges [2]. BSN devices should ultimately have extremely small volumes of 1 cm³ or less. Battery energy density, however, does not scale well down to these sizes. This precludes the use of many existing solutions that deal with dynamic environments, such as the protocols used for cell phone communication today that overpower channel dynamics with conservative coding and large transmission power. The target size for BSN devices limits the energy budget to the range of 10's to 100's of Joules, 2-3 orders of magnitude less than a cell phone battery.

The main goal of this paper is to present a vision for BSNs that incorporates principles and novel ideas across all layers of the system and that are required to meet CPS challenges. Section II surveys two important aspects of related work, but this paper is not a comprehensive survey of BSNs. Section III then presents new principles and techniques for adaptive

operation in highly dynamic physical environments. This includes a holistic cross-layer approach that simultaneously addresses all aspects of the system, from low-level hardware design to high-level communication and data fusion algorithms. Several open research questions are highlighted in Section IV.

II. RELATED WORK

CPS challenges for a BSN arise from three domains: applications, devices, and wireless communication. Section II.A reviews related work on applications. Device challenges come from the need for small size and low power budgets. Section II.B overviews device hardware related work.

A. BSN Applications

BSNs are used in a growing list of applications including fall detection, gait analysis, monitoring the heart with ECG, pulse oximetry, and detecting Parkinson’s episodes and their severity. Many commercial products are available. For example, Human Recorder Co. LTD and LifeSync are two companies selling wireless ECGs, FaceLake and NatureSpirit are two available pulse oximeter products, and many companies including Wellcore and Philips sell fall detection products. In keeping with the holistic theme of this paper, we focus here on a key cross-layer issue using fall detection as an example.

In BSN fall detection, existing work shows how solutions that expand to include more information from multiple sensors and from higher layers in the system hierarchy (e.g. context) provide better results. As evidence, consider that some fall detectors try to detect the fall event by monitoring thresholds in acceleration. For example, Prado [3] uses a four-axis accelerometer located at the height of the sacrum. Kangas [4] studied acceleration of falls and activities of daily living (ADLs) from the waist, wrist, and head, and showed that measurements from the waist and head were more useful for fall detection. Bourke [5] placed two tri-axial accelerometers at the trunk and thigh and used upper and lower thresholds for both the trunk and thigh. Exceeding any of the four thresholds indicated a fall had occurred. The problem with only using accelerometer thresholds is that other activities such as sitting down quickly or walking vigorously on stairs also generate large vertical acceleration, causing many false positives.

By adding information on body orientation from other sensors, new fall detectors provide a more accurate solution. Noury [6] developed a fall detector BSN consisting of three sensors: a tilt sensor to monitor body orientation, a piezoelectric accelerometer to monitor vertical acceleration, and a vibration sensor to monitor body movements. Noury [7] also developed a sensor with two orthogonally oriented accelerometers and used this system to monitor the inclination and inclination speed to detect falls. Li [8] combined accelerometers (exploiting thresholds), gyroscopes (exploiting postures) and context information (enabling different thresholds to be used in different settings and postures) to further improve fall detection, showing a significant reduction in false positives.

A related application example is using BSNs to identify

individuals at a high risk for falls in order to intervene *before* a fall event occurs. BSNs are therefore being deployed to continuously and non-invasively collect gait and posture data, which can then be analyzed to study the fall risk mechanisms in various populations and ultimately identify high fall risk individuals. For example, TEMPO has been used to classify a “shuffle” gait, which is a prominent cause of falls in the elderly, but the natural variability of gait both within and between individuals complicates the processes preceding classification, such as feature identification, selection, and extraction. Thus, robust and generalized signal and information processing methods are needed to classify inertial gait data both on- and off-node. Using data collected on subjects with a TEMPO node on the right ankle, shuffle gait classification using information-theoretic feature extraction and neural networks yielded nearly 98% accuracy of classifying normal from shuffle gait with as few as two features from a gait cycle and one measurement location (i.e. ankle) as training and test vectors [9]. These principles were also applied to a human subjects study investigating fall risk in end stage renal disease (ESRD) patients on hemodialysis (HD), who experience dramatically higher fall rates than the general population [10]. The study used non-invasive, portable gait, posture, strength, and stability assessment technologies (including TEMPO) to extract mobility parameters known to predict fall risk in the general population both pre- and post-HD for inter-HD periods of two and three days. The results indicated that HD treatment influenced strength and mobility, and inter-dialysis period influenced pre-HD profiles. These results are now being used to identify ESRD patients on HD who are at higher risk for falling and target them with interventions specifically designed for this patient population.

These examples illustrate the numerous application-oriented challenges and opportunities facing BSNs, from coordinated sensing to information extraction from raw sensor data. Solutions will require input from domain experts to ensure that the evolving technologies address the true application requirements. This is true not only in medical applications, but also in those targeting fitness and entertainment. With the proper coordination between BSN technology development and emerging application requirements, BSNs are well positioned to deliver the biofeedback and interactivity necessary for the applications of the future.

B. Hardware for BSNs

In this section, we review existing BSNs built from commercial-off-the-shelf (COTS) parts and custom designed application specific ICs (ASICs) and discuss their limitations and strengths.

COTS BSN nodes generally consist of separately packaged components integrated onto a printed circuit board (PCB). These COTS designs are often based on general wireless sensor network (WSN) motes and share characteristics of general sensor nodes. However, BSN design requirements differ distinctly from WSNs, so BSN nodes largely based on WSNs exhibit inefficiencies, most notably in power consump-

tion. Most COTS platforms include sensors, a front end analog amplifier, a digital filter, a microcontroller, a battery, a reference oscillator, and a radio transceiver. COTS nodes provide solid development platforms that are flexible and easy to build, granting rapid access to prototypes and offering programmability to facilitate changing application requirements, processing algorithms, measurement methods, and communication protocols. For this reason, COTS based platforms are ideal for the development of accurate data acquisition approaches and for clinical data collection. The form factor of COTS based platforms can be small (volume $<1 \text{ cm}^3$, weight $<100 \text{ g}$ [11]) and wearable, using either skin mountable [12] or Velcro chest strap designs [13], although many COTS nodes remain several cm's on a side. Their energy inefficiencies may ultimately limit the full deployment of COTS designs in a wide range of emerging healthcare applications, as it is difficult to prolong the battery lifetime of a COTS platform to much more than 1 day [14][15].

Many COTS nodes employ Bluetooth, Zigbee, or other radios commonly seen in WSN nodes. A survey across different platforms shows that the 868 MHz unregulated band [14][16] and the 2.4 GHz band [15] are popular for COTS platforms, since they provide easy integration with the larger system. The power consumption of these radios makes the COTS node power alarmingly high, often reaching the several 100 mW range [11][17][16] (500 mW for [17]). Use of these radios, while convenient, supports data rates that dramatically exceed the requirement for most BSN applications. The frequency of physiological signals typically range from the 1 Hz to 100's kHz range [11], and data rates in transmission and receiving typically only range up to 100's of kbps [14][17]. A more suitable radio would help to reduce the power of COTS nodes, but convenience seems to trump power as the parameter influencing radio choice in most existing COTS designs.

Non-radio components do not affect the overall COTS power consumption significantly [14]. Instead, the selection of other components such as the analog front end and microcontroller focus on accuracy of measurement, quality of attained signals, noise rejection, and flexibility of programming. For example, [16] employs Microchip's 5 V PIC processor and [17][12] use TI's 3.6 V MSP430. Few developers of COTS nodes discuss attempts to improve the energy or power efficiency in on-node computation or control logic. The analog front end and sensors also employ advanced industry components for their accuracy and noise rejection [16][13]. This design strategy is typical of COTS BSNs nodes, as their main focus to date appears to be attaining clinical quality signal processing to assess the methodology being explored and prototyped on the node.

With the power problem becoming increasingly identified as a major bottleneck, power and energy modeling techniques will aid in power optimization methods. With careful optimization, COTS platforms can achieve self-sustained, autonomous operation with an energy harvester. For example, [18] was able to accomplish a 15 fold power reduction from 7.1 mW to 450 μW by optimizing knobs from the application,

system, and hardware levels, despite having a general COTS architecture such as incorporating a COTS micro-controller and radio. In rare cases, harvesting can provide this much power.

In summary, COTS based designs are characterized by easily programmable components focused on delivering high quality clinical data. Most COTS designs use simple architectures with little to distinguish them from WSN nodes, which reflects designers' focus on utility with low power design as a secondary concern. COTS radios dominate node power and are over provisioned for BSN applications, pointing to the need for using lower power, domain specific radios and employing system level methods to reduce the contribution of radio power to the total. Battery lifetime in COTS designs is still too short ($\sim 1\text{-}3$ days) to allow for wide spread deployment of these nodes across the full spread of BSN applications.

ASIC based BSN chips swing to the other end of the spectrum from COTS designs. These custom nodes are application specific instead of flexible and generic, which follows naturally from the need for excellent efficiency to extend system lifetimes. Custom BSN chip design is still an emerging field, and the number of complete systems in the literature is limited. One reason may be that a complete BSN node requires optimized blocks that each require unique expertise to develop efficiently. An implantable $0.5 \times 1.5 \times 2 \text{ mm}^3$ intraocular pressure sensor in [19], designed to provide continuous feedback for glaucoma treatment, incorporates a solar cell, a MEMS pressure sensor, and a microbattery with a low power SoC. The chip converts the capacitive output of the MEMS sensor to a digital value with a 3.6 V, 7 μW switched capacitor circuit, and an 8 bit 0.4 V, 90 nW, 100 kHz microcontroller stores the data in a 4 kb SRAM. An FSK based transmitter sends one bit 40 mW bursts every 131 μs . With > 10 hours of indoor light a day and measurements less than every 15 minutes, the node can run perpetually from harvested energy.

A glucose sensor on a chip with a wireless transmitter is integrated with a contact lens for diabetes monitoring in [20]. The sensor is inductively powered by a reader held near the eye, communicates with a 2.4 GHz LSK scheme, and consumes less than 3 μW . It utilizes a sub- μW regulator and bandgap reference. The degree of energy autonomy shown in these two designs is ideal for BSNs, and both designs show how block by block optimization and limited flexibility can lead to impressive energy efficiency and miniaturization.

An ECG system in [21] combines a sensor chip housing an analog front end and ADC with controller chip that is integrated onto a flexible band covering most of the chest. The band inductively powers the sensors, which adhere to the body in small disposable bandages underneath the chest band. Each ECG sensor chip [22] consumes an average of 12 μW for the analog front end, clock generation, regulator, and ADC, and the controller chip uses 5.2 mW (for aligned inductors) while powering a sensor and processing the ECG signal [21]. Both the power consumption and design complexity are dominated by the analog circuits for sensing and communication.

These same trends of application focus and extreme optimi-

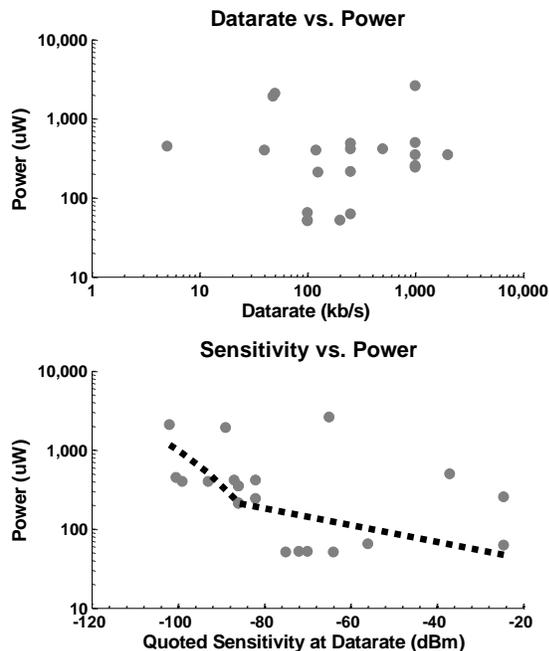


Figure 2. State of the art low power radios and wake up radios. Power is independent of data rate but roughly proportional to sensitivity with a floor of about $50 \mu\text{W}$.

zation emerge in custom components for BSN use. A processor and timer targeting low duty cycle sensing in [23] optimizes sleep power down to below 30 pW and only uses $\sim 300 \text{ fW}$ during active operation at 106 kHz . In another design, a programmable analog front end for biomedical signals in [24] integrates clock generation, tunable filters, and a 12b SAR ADC to produce digitized output samples. The chip consumes 895 nW when acquiring raw ECG data through the front end.

So far, few ASIC nodes leverage system level requirements to reduce power in the radio or analog blocks, instead following the COTS lead by focusing on extracting and communicating raw data. There are a few exceptions. For example, a mixed signal SoC integrating an analog front end and ADC with an 8 bit PIC processor operating in sub-threshold [25] leverages 700 nW processing to reduce the burden on the system radio (not integrated on this chip). The processor uses only $1.5 \text{ pJ/instruction}$ at 280 mV and 450 kHz , and it can extract instantaneous R-R heart rate intervals from a raw ECG sampled at 1 kHz . Transmitting this information instead of raw ECG reduces the wireless data rate by 500X . Also, the processor can successfully maintain accurate computation of heart rate even when it reduces the bias currents in the input amplifier and ADC, causing those analog components to suffer in terms of their block level parameters but permitting the system to function with high fidelity. This allows the full analog front end, ADC, and digital power to drop to only $2.6 \mu\text{W}$ during heart rate extraction and raw ECG acquisition [25]. Similarly, the EEG processing node in [26] includes an analog front end, ADC, and processor for feature extraction. The $\sim 4.3 \mu\text{W}$ power of the chip is dominated by the $3.5 \mu\text{W}$ instrumentation amplifier. Using the chip in combination with a Chip-Con radio shows that on node feature extraction saves 14X system power by reducing the communication load.

Since radios consume so much power when they are active, moving to lower power radio designs is important. Recent low power radios and wake up radios [27], (reported in JSSC, ISSCC, and Symp. on VLSI Circuits) tend to target data rates between 100 and 200 kbps . They typically use the 2.4 GHz ISM or 1.9 GHz PCS bands, use simple modulation schemes such as on-off keying (OOK), report sensitivities from -60 to -80 dBm , and consume between $50 \mu\text{W}$ and $100 \mu\text{W}$. Power for these radios appears to be independent of data rate but proportional to sensitivity. Figure 2 shows that a 10x increase in power leads to roughly a 100x increase in sensitivity. This implies that lower sensitivity radios can offer much lower power, but lower sensitivity may limit communication to periods with better channel characteristics, which is a motivation for the adaptive techniques we describe later in this paper.

In summary, initial forays into custom design for BSN nodes show that extremely energy efficient designs are possible when the hardware is tailored to a specific application and heavily optimized. We observe that opportunities for additional substantial energy savings may be possible by using application requirements at higher levels in the design hierarchy to adjust how the hardware operates. This indicates the potential benefits of supporting ultra low power processing and different power modes than can adjust energy consumption of the hardware as application needs vary.

III. NEW PRINCIPLES AND TECHNIQUES

Current research is underway to develop a combination of software and hardware principles and architectures that span from a new system-on-chip (SOC) based BSN platform to the application layer in order to provide an energy-efficient solution for body sensors that adapts to the CPS challenges of a highly dynamic environment. Building an efficient CPS system for this type of application requires that we analyze issues at every layer of the hierarchy and then iterate to see how the issues and requirements at one layer influence other layers. We have undertaken this exercise, and in this section we propose a hardware platform that we believe contains features important to optimizing the system as a whole. Starting at the hardware level, Section III.A describes the SOC hardware platform, which will have 3 unique architectural aspects identified from our cross-hierarchy analysis to reduce energy usage: (a) an asymmetric radio architecture that uses different bands and modulation for uplink and downlink, (b) low-voltage ICs and hardware accelerators for ultra-low power, adaptive computation, and (c) a low-power receive-only wakeup radio that is highly-tuned to the GSM synchronization symbol, for high accuracy, low-power time synchronization.

To maximally exploit this architecture, we identify three cross-layer approaches to enable adaptive, resource-efficient operation (a) creating new networking protocols and data fusion algorithms that exploit the asymmetric, full-duplex radio architecture for adaptive, efficient distributed operations in dynamic wireless environments (III.B), and (b) modeling and predicting the wireless channel based on a combination of

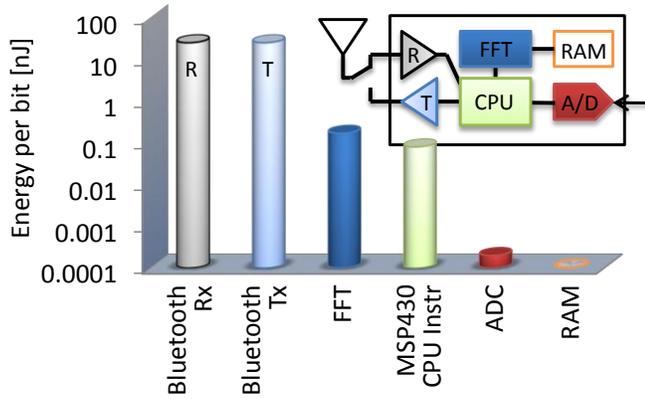


Figure 3. Comparison of energy consumed per bit broken down by task in a wireless sensing node using current state-of-the-art hardware [28][51][52].

low-level cognitive radio techniques and high-level contextual information about the physical world (III.C), and (c) creating an integrated framework for sensing, coding, and signal processing to adaptively balance top-down QoS demands and bottom-up resource and energy availability (III.D).

A. An Ultra-low Power BSN Platform

The extreme energy and form factor constraints of wearable BSN devices necessitate a fundamental departure from the traditional design of wireless embedded sensing devices. Figure 3 illustrates the energy profile of a typical wireless embedded platform today, excluding sensor power: wireless transmission and reception consume the most energy, followed closely by computation. What is required is a new BSN hardware platform to directly address the dominant energy consumers using a custom system on chip (SOC) design that incorporates several key innovations: asymmetric RF communication (to best account for data transfer profiles), adaptive low energy hardware (to adjust to changes in the BSN mode or environment), and synchronization from existing infrastructure (to reduce RF receiver energy during synchronization). This unique architecture is expected to reduce total energy consumption by 1-2 orders of magnitude over existing architectures. The basic SOC architecture is illustrated in Figure 4. A programmable micro-controller provides the heart of the hardware system and interfaces with I/O, memory, and optional hardware accelerators. It also interfaces with an asymmetric radio that uses different bands and modulation for transmission and reception. The receiver can double as a clock harvester, and a digital delay locked loop (DLL) helps the chip

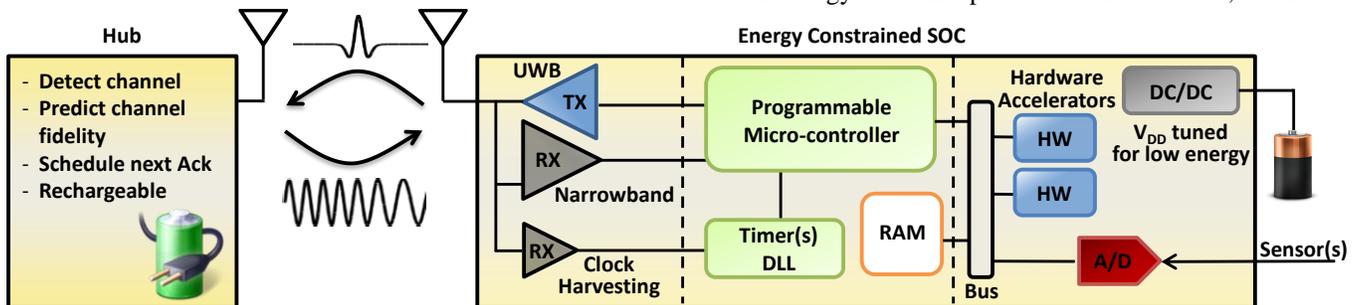


Figure 4. Template for system on chip (SOC) that includes asymmetric RF, programmable ULP processing, and scavenging synchronization signals to reduce energy consumption.

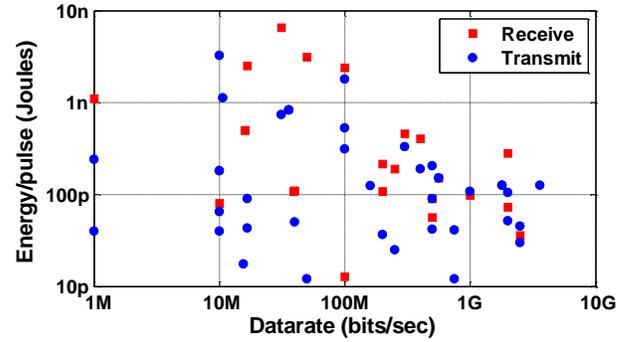


Figure 5. Recently reported UWB transmitters show constant energy per bit over a wide range of data rates.

to remain locked to an external, harvested GSM strobe signal even when that signal becomes temporarily unavailable. On chip voltage regulation allows the operating voltage to be adjusted to meet energy and performance demands. External energy harvesting mechanisms (like thermal gradients) can be added by including a voltage boost regulator. The specific details of the SOC architecture would be tuned for the needs of each BSN application, but the illustration in Figure 4 represents a basic template that captures the major features of an energy efficient hardware platform.

Asymmetric Wireless Communication

The radio architecture is based on two key observations about BSNs. First, BSNs typically use a star-network topology in which the central hub (e.g. a cell phone) typically has more energy resources when compared to the remote nodes, which may be severely energy-constrained. Second, the requirements on data rate for communication from sensor node to controller (uplink) is higher than that from controller to sensor node (downlink) because the uplink carries the sensed information (data), while the downlink may carry configuration instructions or similar, low-bandwidth information.

The system uses separate RF strategies for uplink and downlink that are optimized for these asymmetric energy budgets and data rates. Narrowband radios designed to communicate over a range that is less than 10 meters typically have equal power consumption in transmit and receive modes [28]. Contrast this to pulsed-UWB radios where the transmitter power is typically 100x lower than the receiver power [29]. Furthermore, the energy/bit of UWB transmitters is typically around 50 pJ/bit, compared to 25 nJ/bit for Bluetooth [28], and UWB energy/bit is independent of the data rate, as shown by

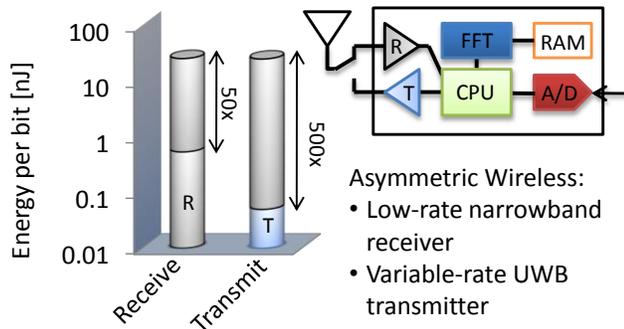


Figure 6. Our proposed asymmetric receive and transmit strategy saves 50x and 500x energy/bit over conventional RF strategies.

the plot of recent UWB transmitter publications in Figure 5. UWB receivers, on the other hand, typically consume 1-to-10 nJ/bit. UWB is, therefore, suitable for transmission from the energy-constrained node, but not necessarily for reception. Low-power narrowband radios can typically achieve lower powers than UWB receivers, and therefore are suitable for reception on the node. The potential energy savings of using a UWB transmitter and a low-power narrowband receiver relative to Bluetooth are illustrated in Figure 6. This allows the node to communicate while using the least amount of resources, at the expense of an acceptable increase in power consumption of the access point. Furthermore, the signal transmitted from the node can be used to measure the channel quality, as discussed in Section III.C.

Adaptive Hardware for Low Energy Computation

Two ultra-low power techniques can be used to reduce the energy consumption required for on-node computation and signal processing: (1) hardware accelerators and (2) low-voltage circuit operation. By reducing the energy of on-node computation, this architecture will expose new energy and latency trade-offs between transmitting raw data and using on-board compression or feature extraction.

A hardware accelerator uses a dedicated circuit to implement a specific function, sacrificing functional flexibility in that component for a 1000x reduction in energy consumption compared to a general-purpose micro-controller [30]. Low-voltage and sub-threshold circuit operation maintains functional flexibility and reduces the energy consumption by over 10x compared to operation at the normal voltage [31]. Figure 7 shows example savings for FFT hardware and CPUs in terms of both energy and CPU cycles. In prior work, successful sub-threshold ICs that implement logic [31], memory [32], and complete micro-controllers [25] as well as systems that adapt energy consumption by moving in and out of sub-threshold [33] were demonstrated. In the future, it is necessary to analyze the processing requirements for different BSNs and develop policies for deciding when to apply low voltage circuit techniques based on the impact for total energy consumption of the system. New research into context sensitive algorithms will determine how much on-node processing makes sense for a given application, and how that information can be combined with simulations of hardware energy con-

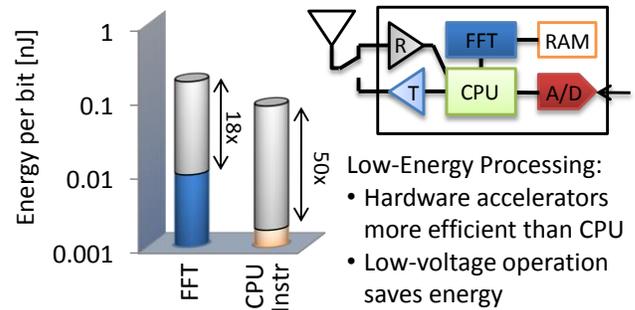


Figure 7. Hardware accelerators and low-voltage digital circuits dramatically reduce the energy relative to existing on-chip computation.

sumption to identify specific sections of the processing that are best implemented using hardware accelerators. For example, control algorithms will likely remain on the micro-controller, but data processing routines (e.g. compression, FFT, etc.) may be best suited for hardware acceleration. In cases where processing energy remains negligible relative to communication, hardware acceleration is unnecessary. In cases where large amounts of processing increase the computation energy substantially, hardware acceleration can be used, and it is also possible to apply low voltage circuit design techniques to minimize energy consumption (e.g. [31]).

Low-power Time Synchronization Harvesting

Time synchronization can allow nodes to reach ultra-low duty cycles while still successfully communicating at pre-determined rendezvous times. However, the reduction of total system energy is limited by clock drift: nodes must periodically resynchronize over the wireless channel, causing energy overhead. The energy overhead of synchronization is amplified in energy-starved BSNs, where nodes often go long periods with their wireless radios off in order to conserve energy. Synchronization energy overhead has been generally recognized as a critical factor in sensor networks, and efforts to reduce the synchronization time for narrowband [27] and UWB [34] radios have been presented. Most notably, a synchronization strategy using a 52 μ W wake-up receiver has been reported, where a wake-up beacon signal generated within the network is used for coarse synchronization of the nodes [27].

An alternative to using wake-up radios to synchronize BSNs is to extract timing signals from one of the ambient wireless signals visible to the BSN. We refer to this as clock harvesting, and the advantage is the high-power wake-up signal is not generated by an access point in the network, but rather is harvested from some other source. For this purpose, we have developed a clock-harvesting receiver (CRX) [35] that synchronizes BSN nodes based on an external GSM signal. The CRX extracts a 21 Hz signal embedded within every broadcast channel of the GSM mobile phone standard.

GSM was selected as a wake-up source because it provides a pervasive and practical signal for use in a BSN. Every GSM cell contains a high-power broadcast channel operating at a fixed frequency. Indoor measurements in a university building show the received power of this channel ranges from -65 dBm

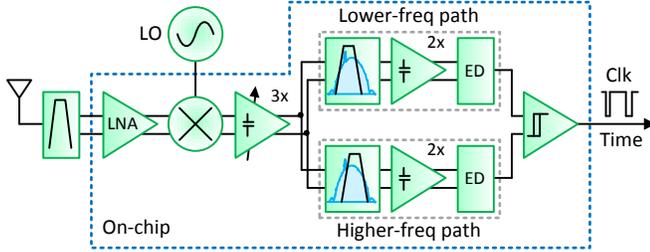


Figure 8. Block diagram of the clock-harvesting receiver that captures a 21Hz reference clock from the GSM broadcast channel.

to -95 dBm. Embedded in every broadcast channel transmission is a tone burst sent periodically for synchronizing cell phones. This signal may be extracted with a low-power receiver, and used to synchronize the reference clocks of BSN nodes. Figure 8 shows a block diagram of CRX to extract this signal. An off-chip SAW filter selects the 1900 MHz PCS band. The input is amplified and down-converted to an IF of 250 kHz, then split into two paths, each with Gm-C bandpass filters that select the tone burst signal. The outputs of the filters are envelop-detected and compared, generating a digital clock output from the GSM signal.

Figure 9 outlines the operation of the receiver in time. The GSM broadcast primarily transmits Gaussian minimum-shift keyed (GMSK) data. Approximately every 46 ms, a pure sinusoidal tone is transmitted. This tone burst lasts for 577 μ s at an offset frequency of 67.7 kHz from the center of the channel. The filter stages in the CRX determine when this frequency burst is present by tuning them to overlapping halves of the broadcast channel. When data is transmitted by a GSM cell, power is spread over the entire channel, passing equal signal levels through each filter. During a tone burst, however, the signals captured along the two filter paths differ. This difference is then detected by the comparator.

A prototype CRX was fabricated in a 0.13 μ m CMOS process and operates from a 1V supply. The LNA can be tuned over a frequency range covering the 1900 MHz band, enabling the selection of any broadcast channel in the US. The measured clock error rate (CER), defined as the number of clock errors to the number of correct clock outputs, is 10^{-3} at an input power level of -87 dBm while the power consumption of the CRX is 126 μ W. Proper operation was verified at input powers up to -5 dBm. The measured jitter at peak sensitivity is 57 μ s, but this reduces to 2 μ s at higher input powers. In sleep mode, the leakage power of the CRX is 81 pW. This CRX was specifically designed with a low sleep mode power to enable a hierarchical synchronization strategy. This allows the CRX to be coarsely duty-cycled by a low-accuracy timer, and fully-powered only momentarily around the 21 Hz GSM clock edges.

B. New Network Protocols and Distributed Data Fusion Algorithms

Numerous wireless protocols and sensor fusion algorithms have been designed for networks of low-power embedded devices, including neighborhood-based in-network processing [36], centralized aggregation [37], mobile agents, and auto-

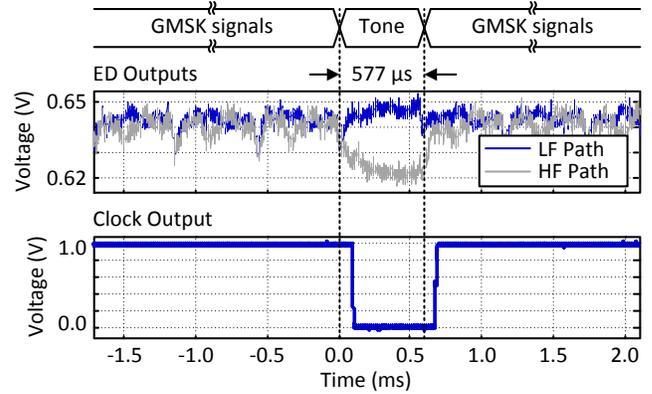


Figure 9. Time response of a frequency burst from a GSM tower and the corresponding clock output of the CRX.

matic program decomposition [38], among others [39]. However, the BSN hardware platform described above has several unique characteristics that preclude these or other existing techniques. First, the asymmetric radio architecture eliminates a basic primitive for wireless communication: the broadcast channel. This will affect protocols at all layers of the network stack that rely on channel sensing, broadcast messages, eavesdropping, and/or direct neighbor-to-neighbor communication. Second, the new platform has a very different resource and energy profile than existing platforms, which will change the design space of efficient BSN applications and protocols. For example, the cost of receiving is an order of magnitude higher than the cost of transmitting, which poses challenges for traditional MAC-style channel arbitration and limits the use of link-layer ACKs for reliable delivery and flow control.

Novel low-power networking protocols and data fusion algorithms are required for BSNs using a holistic approach that percolates the unique characteristics of the BSN platform all the way up and down the application stack. For example, the potential for full-duplex master-slave communication can be used to perform out-of-band coordination between neighboring nodes and to produce an emulated broadcast channel when necessary. These wireless primitives can be used to design a suite of network protocols for network discovery, transmission scheduling, and algorithms for leader election and the efficient calculation of aggregate statistics over a group of nodes can be developed. The end result is a catalogue of cross-layer protocols and algorithms that jointly optimize the costs and needs of sensing, data fusion, communication, and clock synchronization.

Once this catalogue is created, the system will use new techniques to adaptively tune the algorithms and dynamically switch between algorithms and protocols based on the properties of a given application and the current operating conditions. These techniques will use both top-down contextual information as well as bottom-up resource information. For example, changes in heart rate or blood oxygen levels may not be a concern if the user is exercising, but could be a concern if the person is driving or has suddenly fallen. Top-down information that the user is stationary (sleeping or in a vehicle) will lead to the use of algorithms and protocols tuned for low-

power steady state operation, such as time-synchronized sensing and rendezvous-based communication. On the other hand, the context of bodily motion (walking or running) will lead to the use of algorithms tuned for rapid topological changes and high channel dynamics, such as active neighbor discovery, channel monitoring and prediction, and delay-tolerant networking. High-level bodily location, motion, and activities may also indicate the increased availability of energy harvesting and/or the availability of known network resources such as a cell phone or access point. This information can be used to trigger batch operations such as the upload of data logs, or the creation of high-fidelity data for periodic sensor calibration.

New data fusion algorithms must also be used to exploit application logic and top-down sensing requirements to reduce total system energy through strategic data collection. The value of the data from one sensor may depend on the data from another sensor. For example, sensors that detect physical activity for diabetes may not be necessary if blood sugar levels have been detected to be normal. In these cases, the data fusion algorithms can automatically order the sensing and data processing to avoid using unnecessary sensors, increasing communication and data fusion but reducing total energy consumption. Thus, true bi-directional information flow is necessary: sensor data is necessary to detect the situational context, and the context can be used to decide which sensors should be used.

C. Exploiting Periodic or Predictable Wireless Channels

One key aspect of the physical BSN environment is that bodily motion will create bounded dynamics in the wireless channel due to bodily occlusions and the changing, but bounded, proximity of devices attached to the body. This hostile RF environment can be handled by predicting and dynamically adapting to current wireless channel conditions, rather than reacting to them. The cost of supporting dynamic adaptation (e.g. by local processing) will be small relative to the savings provided from reduced wireless communication overhead from, for example, dropped packets or unnecessarily strong coding or transmission power. In order to realize these energy savings, cross-layer approaches to channel prediction and adaptation should incorporate both top-down context about user and bottom-up channel sensing information.

The foundation for a channel prediction scheme can be a set of systematic experiments for BSN channel characterization. Preliminary experiments show high variation in the received signal strength between nodes at different locations on the body (Figure 10), and many other BSN channel modeling experiments corroborate these results [40][41][42]. Two generalized measurement setups for both narrowband (900 MHz and 2.4 GHz) and UWB communication are most prevalent. The first method collects data using a vector network analyzer measuring S21 between a transmitter on the body and a receiver (e.g. [41]). The second method records RSSI or LQI of a signal at a given sampling frequency (e.g. [40]). Path loss models are generated using curve/distribution

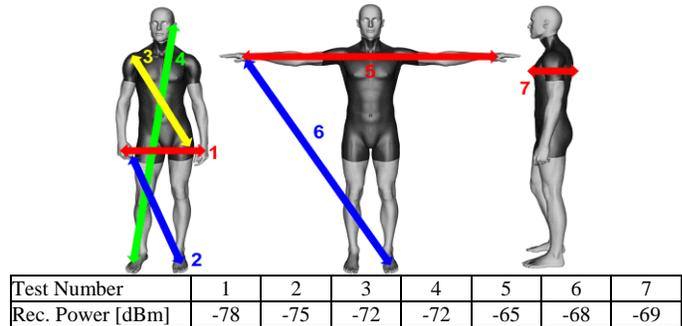


Figure 10. Receive power between nodes for varying positions using 2.4 GHz radio (-12 dBm TX).

fitting while second order statistics such as amplitude distribution, level crossing rate, and fade duration are measured as well [41]. Distributions such as lognormal, Rayleigh, Weibull, and Gamma, are all found to be a best fit in certain situations, but the most common fit is lognormal.

Due to the rhythmic nature of a body's motion, channel periodicity is observed when test subjects are walking, running, or even trying to stand still [42][43]. Channel periodicity is seen in the collected data and its impact is reflected in the distribution models. However, these models do not provide insight into the impact of normal, uncontrolled movements on the shape and frequency of the channel. Figure 11 reports LQI data recorded using Motorola wireless sensor nodes in the 2.4 GHz band. The dark blue line shows control data taken using a mechanical slider on a track, simulating someone's arm swinging at a constant rate while walking. The time domain data shows a relatively consistent periodicity which results in the expected lognormal probability distribution function to its right. The red line in Figure 11 is 3 minutes of LQI data taken while a subject was working at an office desk with a transmitter on the right wrist and a receiver at the left hip. Unlike most channel modeling experiments, the person in this test was *not* instructed to sit still. Due to the typical, random movement of the person during the 3 minute test, the channel model shows variation, but not the repetitive and predictable motion of the control setup. This is again reflected in the distribution to its right, but the binomial nature of the data makes a lognormal fit inaccurate.

Real world scenarios like the red line in Figure 11 motivate the need for an advanced channel modeling method beyond distribution fitting. This measurement illustrates that there may be long periods of time when the channel is exceedingly

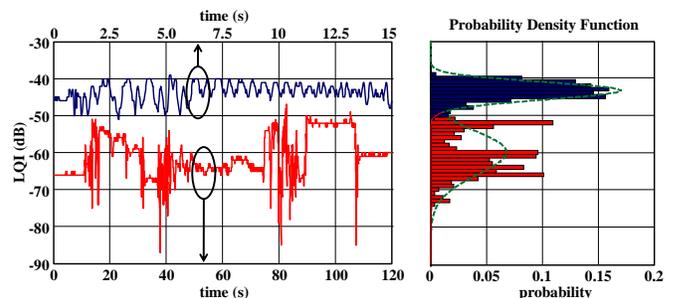


Figure 11. Measured LQI on a controlled track (dark blue), and when sitting at a desk (red), with corresponding histograms.

good, and a low-power radio with poor sensitivity is sufficient for communication. There are also long periods of time when the channel is exceedingly bad, and wireless transmission should not even be attempted. The periodic channels from bottom-up channel measurements depend on the top-down context, including movement, activities, and locations. Therefore, predictive channel models can be generated within a BSN by combining bottom-up channel measurements (e.g. a history of RSSI data) with top-down context information (e.g. a person is sitting vs. walking).

Once the predictive channel models are created, techniques to collect both top-down and bottom-up information for channel prediction can be developed. The BSN sensors can be used to detect body context, motion, and activities. The use of this information in predicting channel quality can be considered when deciding whether sensor data should be stored, compressed, or reported to the BSN hub node. For bottom-up channel sensing, the hub will sense the quality of each node-to-hub channel. For example a known preamble could be transmitted prior to each data packet that allows the receiver to measure RSSI, which provides the information to assess the quality of the link.

Once sufficient history of the channel quality is known, and a short-term future channel quality predicted by the BSN hub, the predictions must be sent to the relevant nodes via the asymmetric downlink. The rate at which nodes are updated can be determined by the BSN hub, and the estimate for next channel prediction can be included with the prediction itself, when possible. The mechanisms for making channel predictions should use both top-down context and bottom-up channel sensing. For example, history could be collected on channel data and used to predict future link quality. If the channel quality should get better or worse before the next update, the coding rate may become too high or low, bringing an unnecessary cost of lower data fidelity or lower packet success rates, respectively. Conversely, a highly dynamic channel is likely to remain so, such as a body area channel when the wearer is walking or running. In such a case, a node can either receive more frequent updates, leverage periodicity in the channel quality, or use conservative coding techniques as a function of the worst-case channel conditions. To maximize performance, a decision theoretic framework can be used that weighs the balance of conservative channel estimates versus the cost of increased update rates. These decisions will be based not only on performance characteristics, but also on current and future QoS requirements at the application level.

Once a node has been notified of its predicted channel quality, it must adapt the communication parameters. These adaptation schemes should involve high-level network protocols and distributed data fusion algorithms, as discussed in Section III.B. In situations where a constant packet transmission rate must be maintained, adjusting coding rates based on channel quality would affect data processing output rates and, as discussed in Section III.D., the resulting data fidelity. This competing relationship is central to the cognitive adaptation of a node to its channel condition.

Regardless of the application requirements and operating environment for a given system, it is possible to use a buffer (e.g. FIFO) between the data processor and the code generator to manage these relationships. The processor pushes data into the buffer at its data processing output rate, and the code generator pulls it out based on its coding rate. It is desirable to maintain a specific buffer occupancy set-point throughout operation. Therefore, if the channel quality update results in a coding rate increase to maintain a high packet success rate across a bad channel, the rate at which data is pulled from the buffer will be reduced, forcing the processor to reduce its data processing output rate. Conversely, if the data processing output rate should increase based on the dynamic input data, the increased buffer occupancy would need to be placated by a reduced coding rate, potentially reducing the packet success rate. An overarching node controller would need to manage such a competition. Such policies could be incorporated into node-level operating systems, such as the PixieOS for wireless sensor networks [44].

D. Adaptable Sensing, Coding, and Signal Processing

The ultimate efficacy of a BSN application is dependent on the fidelity of the collected data, but the collection and wireless transmission of this data pose the greatest challenges for achieving the battery lifetimes required by the same applications. It, therefore, becomes necessary to model the relationship between energy consumption and fidelity in a way that enables system designers to select sensing, coding, and processing modes that best achieve the desired tradeoff between these metrics. In addition, given the system dynamics of variable sensor data, channel characteristics, and application context, the energy-fidelity relationship changes at runtime, require dynamic mode adaptations. In certain contexts, a particular sensor may provide essential data that must be timely and of the highest quality, while other contexts may enable that sensor to be turned off to save energy. Such top-down contextual information can be combined with bottom-up analysis in which a BSN node is empowered to determine how important real-time data is to an application. Overall, the system must continuously balance fidelity requirements with resource efficiency.

To illustrate, consider on-node lossy compression of inertial data collected on a Parkinson’s Disease patients with the TEMPO BSN in a human subjects study on the efficacy of deep brain stimulation for tremor control [45][46]. As illustrated in Figure 12, lower data rates result in higher distortions, which correspond to lower energy consumption (fewer bits to transmit) but lower data fidelity. The figure also shows that the energy-fidelity relationship is a function of the data. For example, scaling the “Stationary Small Tremor” data rate to 1-bit (from 12-bits, normalized per sample), equates to a 12X reduction in data rate – and a substantial savings in BSN node energy – with minimal resultant mean squared error (MSE) distortion, but scaling the “Stationary Large Tremor” to 1-bit, would result in a much larger distortion and perhaps an unacceptable loss of application fidelity.

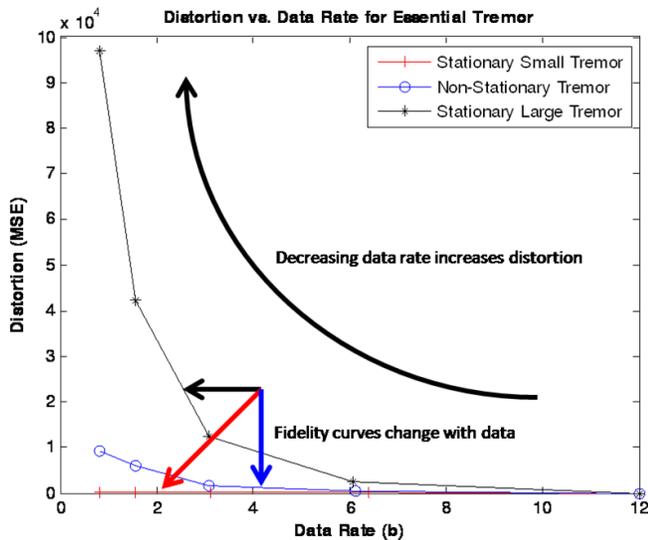


Figure 12. Example rate-distortion for compressed tremor data [50].

This observation points to the need for application- and individual-specific profiling to effectively manage the energy-fidelity tradeoff. Moreover, it is interesting to note that Figure 12 presents data collected from a single patient over the course of a single clinical visit, which elucidates the need for system adaptation based on dynamic data. To illustrate further, Figure 13 depicts a time domain distortion plot for fixed data compression, yielding a compression ratio (CR) of approximately 18, for a 40 minute tremor dataset. If an $MSE \leq 100$, for example, were required for application fidelity to remain acceptable, then any distortion below this level would be considered energy inefficient (marked as the lower region in Figure 13) because data rate could be further reduced to meet the application requirement; and data above this level would not have high enough fidelity to meet the requirement (upper region). Only by dynamically adjusting a data rate knob would the node operate in an application-specific energy-fidelity optimized range (middle region). Similarly, variable wireless channel quality can affect the energy cost of transmission (due to retransmissions, stronger ECC, higher transmission strengths, etc.) and contextual information can affect the measure of fidelity (different information may be of value for different locations, activities, incidents, etc.), pointing to another need for dynamic adaptation.

Using the same 40 minute tremor dataset in Figure 13 and the BSN node energy model and distortion profiling techniques presented in [47], the average MSE and energy per bit

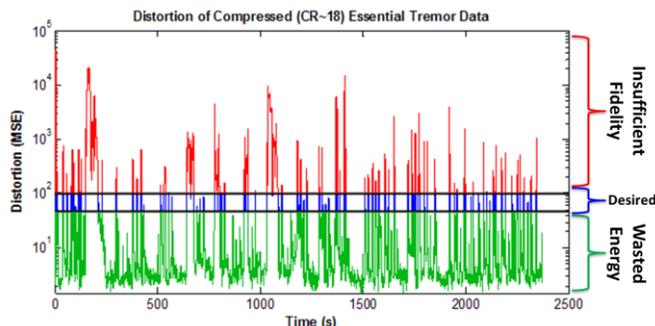


Figure 13. Dynamic distortion of tremor data for a fixed CR [50].

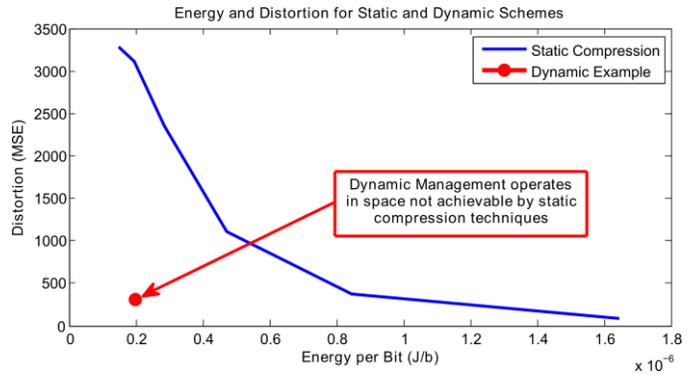


Figure 14. Static vs. dynamic Haar compression [50].

were calculated for scenarios where the compression ratio (Haar level for this example) is statically fixed and where it is varied dynamically based on real-time measures of variance, which was shown to be a good proxy to determine the instantaneous energy-fidelity relationship. Processing overheads for performing Haar compression and variance calculation using 128-sample windows are included in the energy model, along with power management schemes that put the processor to sleep when inactive. The static setting is swept across all possible Haar levels, while dynamic management results in a single outcome for the given data set and optimization objective. As shown in Figure 14, the dynamic compression scheme operates in a region that is well under the Pareto-optimal curve of the static case.

New work can explore data-specific energy-fidelity relationships provided by not only on-node processing (which also includes classification, event detection, feature extraction, etc.), but also for sensing (i.e. what sensors to have on when and at what sampling rate and quantization bit depth) and coding (i.e. how to best code for wireless transmission). Based on these relationships, control schemes can be devised that dynamically manage energy and fidelity based on real-time measures of sensed data, channel quality, and context.

IV. OPEN RESEARCH CHALLENGES

The vision for BSNs presented above addresses many of the CPS challenges relating to application demands, hardware platforms, and wireless communications. We now briefly present several overarching and open challenges for BSNs.

A. QoS and Fidelity

A formal methodology for the design and runtime management of QoS requires identifying quantitative contributing metrics and techniques for combining them to form a single QoS assessment. Ultimately, the selection of metrics and the weight they are given is application dependent, but nearly all candidate BSN applications share requirements for high application fidelity and long battery life, and many call for high throughput and low latency.

In essence, a high QoS must be achieved under severe resource constraints, complex environments, and varying application semantics and context. While a layered approach to QoS management is appropriate for more resource-heavy systems, BSNs require an integrated, multi-scale approach

capable of finding operating points across system levels to maximize QoS metrics based on application-specific requirements and dynamic operating conditions. This is especially important considering that each system level in a BSN affects all of the other levels, making the level-independent selection of QoS operating points suboptimal. QoS management from a multi-scale perspective must also be explored by creating BSN system models that not only leverage the optimizations that can be performed at individual system levels (node-level and intra- and extra-BSN networking), but also consider how such optimizations affect the levels above and below it. It is also necessary to validate and improve QoS optimization strategies by experimenting on physical BSN systems in real application settings to ensure realistic BSN system models.

High reliability and application fidelity are essential for many BSN applications – especially where “medical-grade” reliability is required – but can be quite difficult to achieve. One key challenge arises from the communication difficulties due to body motions, user mobility, co-existence of many BSNs in the overall system, interference from other wireless devices found in the environment, and other dynamics caused by realistic environments with obstacles and movement. It can also be expected that individual BSNs may become “detached” from the infrastructure system at times, so delay tolerant networking solutions may be required.

Another QoS challenge relates to the quantitative evaluation of reliability and application fidelity. Current level-specific mechanisms for these metrics include signal-to-noise ratio (SNR), inverse bit error rate (1/BER), and packet reception ratio (PRR). While these are useful measures for some aspects of QoS assessment, they do not necessarily directly correlate with application fidelity because they are inherently application-independent. For example, an application focused on tremor assessment in Parkinson’s Disease patients is ultimately most interested in the resulting rates of correct diagnoses, courses of treatment, etc. Those are ultimately the kinds of quantitative reliability and fidelity metrics that must be incorporated into QoS design and runtime management.

B. Robustness

Traditionally, the majority of sensor based systems have been closed systems. Creating robust systems in these settings, while not easy, has been the subject of research for many years. However, BSNs are not closed systems. They move around in arbitrary environments and upload various physiological data to doctors in real-time. Current software/hardware composition techniques, associated analysis techniques, and tools need to be re-thought and developed to account for BSNs operating in open and heterogeneous environments. New unified communications interfaces will be required to enable efficient information exchange across diverse systems and nodes.

After initialization, BSNs have certain properties. For example, the nodes of the BSN know their locations, have synchronized clocks, know their neighbors or aggregator, and have a coherent set of parameter settings such as consistent

sleep/wake-up schedules, appropriate power levels for communication, and pair-wise security keys. However, over time these coherent states can deteriorate. The most common example of this deterioration problem is clock drift, which causes nodes to have different enough times to result in application failures. For robustness with respect to time, clock synchronization must re-occur. Re-running protocols to re-establish BSN properties is an important robustness issue. Robustness must also make use of formal methods to develop reliable code, use in-situ debugging techniques, and provide on-line fault tolerance, in-field-maintenance, and general health monitoring services. Problems are exacerbated due to the unattended operation of the system, the need for a long lifetime, the openness of the systems, and the realities of the physical world.

C. Safety

The safety of sophisticated new health care systems that include BSNs has not received the attention that developing novel functionality has. What is required is a systematic, multi-disciplinary research program to tackle the safety challenges raised by the increasing synergy of health care and engineering. Of particular concern are the many medical technologies that are built as stove piped systems. Each addresses aspects of safety to some limited extent, but there is no coordinating blueprint for ensuring interoperable and safe operation of the resulting system of systems. Further, a looming danger is the market pressure that favors functionality over demonstrated quality as exemplified by major commercial software applications that ship with known defects.

In order to create and maintain safety in next-generation medical environments we must have radically new solutions that are aware of the complexities, intricacies and dynamics of the physical world, as well as being aware of human behaviors, disease information, and external knowledge such as product safety studies. Explicit demonstration of safety on an ongoing basis is also essential because of the dynamic nature of engineered healthcare systems.

D. Security and Privacy

A fundamental problem that must be solved in BSNs is dealing with security attacks [48]. Security attacks are problematic for BSNs because of the minimal capacity devices being used and the openness of the systems, including the fact that most devices will communicate wirelessly. The security problem is further exacerbated because transient and permanent random failures are commonplace and failures are vulnerabilities that can be exploited by attackers. To meet realistic system requirements that derive from long lived and unattended operation, BSNs must be able to continue to operate satisfactorily in the presence of, and to recover effectively from security attacks. Cryptographic keys must be distributed securely. The system must also be able to adapt to new attacks unanticipated when the system was first deployed. It is an open question as to the type of hardware support required and how the layers of software can deal with these difficult problems.

BSNs provide many useful services for individuals, but also create many opportunities to violate privacy. To solve the privacy problem created by wireless BSNs, the privacy policies for a BSN must be specified. Once specified the BSN system must enforce privacy. Consequently, the system must be able to express users' requests for data access and the system's policies such that the requests can be evaluated against the policies in order to decide if they should be granted or denied. One of the more difficult privacy problems is that systems may interact with other systems, each having their own privacy policies. Consequently, inconsistencies may arise across systems. On-line consistency checking and notification and resolution schemes are required. Since BSNs communicate wirelessly, privacy can be violated by eavesdropping. It has been shown [49] that encryption alone is not sufficient to ensure privacy due to eavesdropping. New cost effective techniques are needed.

V. SUMMARY

BSNs are important cyber-physical systems that promise to improve quality of life through improved health, augmented sensing and actuation for the disabled, independent living for the elderly, and reduced healthcare costs. We described the state of art that has primarily investigated several key applications and off-the-shelf hardware and communication quality. We also described a new BSN hardware platform that integrates novel circuit designs and cutting-edge technologies to reduce the cost of communication and computation by several orders of magnitude. New cross-layer solutions that incorporate the unique characteristics of this platform at all layers of the BSN application stack were presented. Such a design produces techniques to adaptively balance changing top-down QoS demands and bottom-up resource and energy availability. We also described techniques to create bi-directional information flow across multiple layers of the system from the HW to the application, allowing all layers to exploit both high-level information about physical context (such as location or physical activity) and low-level system information (such as direct sensing of the wireless channel or physical environment). The value of adaptive sensing, communication, and data fusion were also discussed. Finally, a summary of overarching challenges was presented.

ACKNOWLEDGEMENT

This material is based upon work supported by the National Science Foundation under Grant Nos. 0901686, 1035303, and 1035771.

REFERENCES

- [1] A.T. Barth, M.A. Hanson, H.C. Powell Jr., J. Lach, "TEMPO 3.1: A Body Area Sensor Network Platform for Continuous Movement Assessment," International Workshop on Body Sensor Networks, pp. 71-76, 2009.
- [2] M.A. Hanson, H.C. Powell, A.T. Barth, K. Ringgenberg, B.H. Calhoun, J.H. Aylor, J. Lach, "Body Area Sensor Networks: Challenges and Opportunities," IEEE Computer, vol. 42, no. 1, pp. 58-65, January 2009.
- [3] A. D'iaz, M. Prado, L. M. Roa, J. Reina-Tosina, and G. Sanchez. Preliminary evaluation of a full-time falling monitor for the elderly. In Proceedings of the 26th Annual International Conference of the IEEE EMBS, pages 2180–2183, San Francisco, USA, Sept 2004.
- [4] M. Kangas, A. Konttila, I. Winblad, and T. J'ams'a. Determination of simple thresholds for accelerometry-based parameters for fall detection. In Proceedings of the 29th Annual International Conference of the IEEE EMBS, pages 1367–1370, Lyon, France, Aug 2007.
- [5] A. K. Bourke, J. V. O'Brien, G. M. Lyons. Evaluation of a threshold-based tri-axial accelerometer fall detection algorithm. *Gait and Posture*, 26:194–199, 2007.
- [6] N. Noury, T. Herv'e, V. Rialle, G. Virone, E. Mercier, G. Morey, A. Moro, and T. Porcheron. Monitoring behavior in home using a smart fall sensor and position sensors. In Proceedings of the 1st International IEEE EMBS Special Topic Conference on Microtechnologies in Medicine and Biology, pages 607–610, Lyon, France, Oct 2000.
- [7] N. Noury, P. Barralon, G. Virone, P. Boissy, M. Hamel, and P. Rumeau. A smart sensor based on rules and its evaluation in daily routines. In Proceedings of the 25th Annual International Conference of the IEEE EMBS, pages 3286–3289, Cancun, Mexico, Sept 2003.
- [8] Q. Li, J.A. Stankovic, M.A. Hanson, A.T. Barth, J. Lach, "Accurate, Fast Fall Detection Using Gyroscopes and Accelerometer-Derived Posture Information," International Workshop on Body Sensor Networks, pp. 138-143, 2009.
- [9] M.A. Hanson, H.C. Powell Jr., A.T. Barth, J. Lach, M. Brandt-Pearce, "Neural Network Gait Classification for On-Body Inertial Sensors," International Workshop on Body Sensor Networks, pp. 181-186, 2009.
- [10] T.E. Lockhart, A.T. Barth, X. Zhang, R. Songra, E. Abdel-Rahman, J. Lach, "Portable, Non-Invasive Fall Risk Assessment in End Stage Renal Disease Patients on Hemodialysis," *Wireless Health*, 84-93, 2010
- [11] T. Schlebusch, L. Rothlingshofer, S. Kim, M. Kony, and S. Leonhardt, "On the Road to a Textile Integrated Bioimpedance Early Warning System for Lung Edema," 2010 International Conference on Body Sensor Networks, 2010, pp. 302-307.
- [12] K. S.-H. Ong, Y. Siew-Peng, and L. Keck-Voon, "Implementation of Fast Fourier Transform on Body Sensor Networks," 2010 International Conference on Body Sensor Networks, 2010, pp. 197-202.
- [13] E. Mitchell, S. Coyle, N. E. O'Connor, D. Diamond, and T. Ward, "Breathing Feedback System with Wearable Textile Sensors," 2010 International Conference on Body Sensor Networks, 2010, pp. 56-61.
- [14] E. Farella, A. Pieracci, L. Benini, and A. Acquaviva, "A Wireless Body Area Sensor Network for Posture Detection," 11th IEEE Symposium on Computers and Communications, Sept. 2006, pp. 454-459.
- [15] A. Chhikara, A. H. McGregor, L. Hadjilucas, F. Bello, and A. S. Rice, "Quantitative Assessment of the Motion of the Lumbar Spine and Pelvis with Wearable Inertial Sensors," 2010 International Conference on Body Sensor Networks, 2010, pp. 9-15.
- [16] K.-I. Wong, "Rapid Prototyping of a Low-Power, Wireless, Reflectance Photoplethysmography System," 2010 International Conference on Body Sensor Networks, 2010, pp. 47-51.
- [17] M. Kony, W. Walter, T. Schlebusch, and S. Leonhardt, "An RFID Communication System for Medical Applications," 2010 International Conference on Body Sensor Networks, 2010, pp. 71-75.
- [18] J. Penders, V. Pop, L. Caballero, J. van de Molengraft, R. van Schaijk, R. Vullers, and C. Van Hoof, "Power Optimization in Body Sensor Networks the Case of an Autonomous Wireless EMG Sensor Powered by PV-cells," International Conference of the IEEE Engineering in Medicine and Biology Society, pp. 2017-2020, 2010.
- [19] G. Chen, H. Ghaed, R. Haque, M. Wiecekowski, Y. Kim, G. Kim, D. Fick, D. Kim, M. Seok, K. Wise, D. Blaauw, and D. Sylvester, "A 1 cubic millimeter energy-autonomous wireless intraocular pressure monitor," International Solid-State Circuits Conference, 2011, pp. 310-311.
- [20] Y. Liao, H. Yao, B. Parviz, and B. Otis, "3μW wirelessly powered CMOS glucose sensor for an active contact lens," International Solid State Circuits Conference, 2011, pp. 38-39.
- [21] J. Yoo, L. Yan, S. Lee, Y. Kim, and H-J. Yoo, "A 5.2 mW Self-Configured Wearable Body Sensor Network Controller and a 12 uW Wirelessly Powered Sensor for a Continuous Health Monitoring System," IEEE Journal of Solid-State Circuits, vol. 45, no. 1, pages 178-188, January 2010.
- [22] L. Yan, J. Yoo, B. Kim, and H-J. Yoo, "A 0.5-uVrms 12-uW Wirelessly Powered Patch-Type Healthcare Sensor for Wearable Body Sensor Network," IEEE Journal of Solid-State Circuits, vol. 45, no. 11, pages 2356-2365, November 2010.

- [23] M. Seok, S. Hanson, Y.-S. Lin, Z. Foo, D. Kim, Y. Lee, N. Liu, D. Sylvester, and D. Blaauw, "The Phoenix Processor: A 30pW platform for sensor applications," Symposium on VLSI Circuits, 2008, pp. 188-189.
- [24] X. Zou, X. Xu, L. Yao, and Y. Lian, "A 1-V 450-nW fully integrated programmable biomedical sensor interface chip," IEEE Journal of Solid-State Circuits, vol.44, no.4, Apr. 2009, pp.1067-1077.
- [25] S. Jocke, J. Bolus, S.N. Wooters, A.D. Jurik, A.C. Weaver, T.N. Blalock, and B. H. Calhoun, "A 2.6- μ W sub-threshold mixed-signal ECG SoC," Symposium on VLSI Circuits, 2009, pp. 60-61.
- [26] N. Verma, A. Shoeb, J. Bohorquez, J. Dawson, J. Guttag, and A. Chandrakasan, "A Micro-Power EEG Acquisition SoC with Integrated Feature Extraction Processor for a Chronic Seizure Detection System," IEEE Journal of Solid-State Circuits, vol. 45, no. 4, pages 804-816, April 2010.
- [27] N.M. Pletcher, S. Gambini, J. Rabaey, "A 52 μ W Wake-Up Receiver With -72 dBm Sensitivity Using an Uncertain-IF Architecture," IEEE Journal of Solid-State Circuits, vol. 44, no. 1, pp. 269-280, January 2009.
- [28] W.W. Si, et al., "A Single-Chip CMOS Bluetooth v2.1 Radio SoC," IEEE Journal of Solid-State Circuits, vol. 43, no. 12, pp. 2896-2904, December 2008.
- [29] J. R. Fernandes, D. D. Wentzloff, "Recent Advances in IR-UWB Transceivers: An Overview," IEEE International Symposium on Circuits and Systems, June 2010, pp. 3284-3287.
- [30] H. Zhang, V. Prabhu, V. George, M. Wan, M. Benes, A. Abnous, J.M. Rabaey, "A 1-V Heterogeneous Reconfigurable DSP IC for Wireless Baseband Digital Signal Processing," IEEE Journal of Solid-State Circuits, vol. 35, no. 11, pp. 1697-1704, November 2000.
- [31] A. Wang, B.H. Calhoun, A. Chandrakasan, Sub-threshold Design for Ultra-Low-Power Systems, Springer, 2006.
- [32] B.H. Calhoun, A. Chandrakasan, "A 256kb Sub-threshold SRAM in 65nm CMOS," IEEE International Solid-State Circuits Conference, pp. 628-629, 2006.
- [33] B.H. Calhoun, A. Chandrakasan, "Ultra-Dynamic Voltage Scaling (UDVS) Using Sub-Threshold Operation and Local Voltage Dithering," IEEE Journal of Solid-State Circuits, vol. 41, no. 1, pp. 238-245, January 2006.
- [34] A.P. Chandrakasan, F.S. Lee, D.D. Wentzloff, V. Sze, B.P. Ginsburg, P.P. Mercier, D.C. Daly, R. Blazquez, "Low-Power Impulse UWB Architectures and Circuits," Proceedings of the IEEE, vol. 97, no. 2, pp. 332-352, Feb. 2009.
- [35] J. K. Brown, D. D. Wentzloff, "A 1900MHz-Band GSM-Based Clock-Harvesting Receiver with -87dBm Sensitivity," IEEE Radio Frequency Integrated Circuits Symposium, June 2011.
- [36] K. Whitehouse, C. Sharp, E. Brewer, D. Culler, "Hood: a Neighborhood Abstraction for Sensor Networks," International Conference on Mobile Systems, Applications, and Services, 2004.
- [37] K. Whitehouse, G. Tolle, J. Taneja, C. Sharp, S. Kim, J. Jeong, J. Hui, P. Dutta, D. Culler, "Marionette: Using RPC for Interactive Development and Debugging of Wireless Embedded Networks," International Conference on Information Processing in Sensor Networks, 2006.
- [38] T. Hnat, T. Sookoor, P. Hooimeijer, W. Weimer, K. Whitehouse, "MacroLab: A Vector-based Macroprogramming Framework for Cyber-Physical Systems," Conference on Embedded Networked Sensor Systems, 2008.
- [39] K. Whitehouse, J. Liu, F. Zhao, "Semantic Streams: a Framework for Composable Inference over Sensor Data," European Workshop on Wireless Sensor Networks, 2006.
- [40] Khan, Mohammad Monirujjaman; Alomainy, Akram; Hao, Yang; , "Off-Body Radio Channel Characterisation Using Ultra Wideband Wireless Tags," Body Sensor Networks (BSN), 2010 International Conference on , vol., no., pp.80-83, 7-9 June 2010
- [41] Chaganti, V.G.; Smith, D.B.; Hanlen, L.W.; , "Second-Order Statistics for Many-Link Body Area Networks," Antennas and Wireless Propagation Letters, IEEE, , vol.9, no., pp.322-325, 2010
- [42] K.S. Prabh, J.-H. Hauer, "Opportunistic Packet Scheduling in Body Area Networks," European Conference on Wireless Sensor Networks, pp. 114-129, 2011.
- [43] Smith, D.; Hanlen, L.; Jian Zhang; Miniutti, D.; Rodda, D.; Gilbert, B.; , "Characterization of the Dynamic Narrowband On-Body to Off-Body Area Channel," Communications, 2009. ICC '09. IEEE International Conference on, vol., no., pp.1-6, 14-18 June 2009.
- [44] K. Lorincz, B. Chen, J. Waterman, G. Werner-Allen, M. Welsh, "Resource Aware Programming in the Pixie OS," Conference on Embedded Networked Sensor Systems, pp. 211-224, 2008.
- [45] M.A. Hanson, H.C. Powell Jr., R.C. Frysinger, D.S. Huss, W.J. Elias, J. Lach, "Teager Energy Assessment of Tremor Severity in Clinical Application of Wearable Inertial Sensors," IEEE-NIH Life Science Systems and Applications Workshop, pp. 191-194, 2007.
- [46] H.C. Powell Jr., M.A. Hanson, J. Lach, "On-Body Inertial Sensing and Signal Processing for Clinical Assessment of Tremor," IEEE Transactions on Biomedical Circuits and Systems, vol. 3, no. 2, pp. 108-116, April 2009.
- [47] M.A. Hanson, H.C. Powell Jr., A.T. Barth, J. Lach, "Enabling Data-Centric Energy-Fidelity Scalability in Wireless Body Area Sensor Networks," International Conference on Body Area Networks, 2009.
- [48] C.C.Y. Poon, Y. Zhang, and S. Bao, A Novel Biometrics Method to Secure Wireless Body Area Networks for Telemedicine and m-Health, IEEE Communications Magazine, 2006.
- [49] V. Srinivasan, J.A. Stankovic, K. Whitehouse, "Protecting Your Daily In-home Activity Information from a Wireless Snooping Attack," International Conference on Ubiquitous Computing, 2008.
- [50] A.T. Barth, M.A. Hanson, H.C. Powell Jr., J. Lach, "Online Data and Execution Profiling for Dynamic Energy-Fidelity Optimization in Body Sensor Networks," International Conference on Body Sensor Networks, 213-238, 2010.
- [51] L. Brooks, H.S. Lee, "A Zero-Crossing-Based 8b 200MS/s Pipelined ADC," IEEE International Solid-State Circuits Conference, Feb. 2007, pp. 460-461.
- [52] A. Wang and A. Chandrakasan, "A 180mV FFT Processor Using Subthreshold Circuit Techniques," IEEE International Solid-State Circuits Conference, pp. 292-293, February 2005.

Benton H. Calhoun received his B.S. degree in electrical engineering from the University of Virginia, Charlottesville, VA, in 2000. He received the M.S. degree and Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, in 2002 and 2006, respectively. In January 2006, he joined the faculty at the University of Virginia as an Assistant Professor in the Electrical and Computer Engineering Department. His research interests include low power digital circuit design, sub-threshold digital circuits, SRAM design for end-of-the-roadmap silicon, low voltage FPGAs, variation tolerant circuit design methodologies, BSNs, and low energy electronics for wireless health. Dr. Calhoun is a co-author of Sub-threshold Design for Ultra Low-Power Systems (Springer, 2006).

John Lach received the B.S. degree in Science, Technology, and Society from Stanford University in 1996 and the M.S. and Ph.D. degrees in Electrical Engineering from UCLA in 1998 and 2000, respectively. Since 2000, he has been a faculty member in the Electrical and Computer Engineering Department at the University of Virginia. He is a Senior Member of the IEEE and is a former Associate Editor for the IEEE Transactions on Computers and the IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems. He is a Co-Founder and Steering Committee member for the Wireless Health conference series and is a Co-Founder and Co-Director of the UVA Center for Wireless Health. His primary research interests include wireless health, BSNs, and digital system design methodologies.

John Stankovic is the BP America Professor in the Computer Science Department at the University of Virginia. He served as Chair of the department for 8 years. He is a Fellow of both the IEEE and the ACM. He won the IEEE Real-Time Systems Technical Committee's Award for Outstanding Technical Contributions and Leadership. He also won the IEEE Technical Committee on Distributed Processing's Distinguished Achievement Award (inaugural winner). He was the Editor-in-Chief for the IEEE Transactions on Distributed and Parallel Systems and was founder and co-editor-in-chief for the Real-Time Systems Journal. His research interests are in real-time systems, distributed computing, wireless sensor networks, and cyber physical systems. Prof. Stankovic received his PhD from Brown University.

David D. Wentzloff received the B.S.E. degree in Electrical Engineering from the University of Michigan, Ann Arbor, in 1999, and the S.M. and Ph.D. degrees from MIT in 2002 and 2007, respectively. Since August, 2007 he has been with the University of Michigan, Ann Arbor, where he is currently an Assistant Professor of Electrical Engineering and Computer Science. He has served on the technical program committee for ICUBW 2008-2010 and ISLPED 2011, and as a guest editor for the IEEE T-MTT, the IEEE Commu-

nications Magazine, and the Elsevier Journal of Signal Processing: Image Communication.

Kamin Whitehouse earned a BS in Electrical Engineering and a BA in Cognitive Science and Philosophy from Rutgers University in 1999. He received his M.S. and Ph.D. in Computer Science from UC Berkeley in 2003 and 2006, respectively, and has been on the faculty of the Computer Science Department at the University of Virginia since 2006. His research focuses on wireless embedded systems and has been recognized through the NSF CAREER award and NSF, NDSEG, GOF, Diversity, and Siebel fellowships.

Adam T. Barth received the B.E. degree in Electrical Engineering from Vanderbilt University in 2006, and the M.S. and Ph.D. degrees from the University of Virginia in 2009 and 2011, respectively. He is a student member of IEEE, and his current research interests include research and designed of embedded systems for use in the wireless health and BSN fields, and low-power electronic design for wearable, implantable, or unobtrusive assessment and treatment of medical disorders.

Jonathan K. Brown received the B.S.E. and M.S.E. degrees from the University of Michigan, Ann Arbor, in 2007 and 2008, respectively. He is currently pursuing his Ph.D. degree in electrical engineering at the University of Michigan, Ann Arbor. He is a recipient of the 2011 DAC/ISSCC Student Design Contest Award. His research interests include the design of low-power wireless integrated circuits and systems.

Qiang Li received his B.S. and M.S. degrees from Tsinghua University, China, where he worked on multimedia processing and embedded systems. He is currently pursuing his Ph.D. degree in Computer Science at the University of Virginia under the guidance of Prof. John Stankovic. His research is mainly on developing reliable and scalable solutions for BSNs.

Seunghyun Oh received the B.S. degree in electrical engineering from Seoul National University, Seoul, Korea, in 2007, and the M.S. degree in electrical engineering from the University of Michigan, Ann Arbor, in 2009. His research interests include RF integrated circuits and communication in body area networks. Mr. Oh is the recipient of the Korea Foundation for Advanced Studies Fellowship.

Nathan E. Roberts received the B.S degree in electrical engineering from the University of San Diego, in 2006, and the M.S. degree in electrical engineering from the University of Michigan, Ann Arbor, in 2011, where he is currently pursuing a Ph.D. degree. From 2006 to 2009, he held a position with the Product Development Division at Lattice Semiconductor in Hillsboro, OR. His research interests include the design of low power wireless integrated circuits.

Yanqing Zhang joined the RLP-VLSI group at the University of Virginia in 2009 after receiving his BS degree from Fudan University, Shanghai, China. He is in his third year working toward a PhD in Electrical Engineering. His research interests include low power SoCs, SoC architecture optimization, sub-threshold clock tree synthesis and timing closure, sub-threshold synthesis design flow optimization, and low power, robust block level design methods.