

# Enabling Data-Centric Energy-Fidelity Scalability in Wireless Body Area Sensor Networks

Mark A. Hanson, Harry C. Powell Jr., Adam T. Barth, John Lach

Charles L. Brown Department of Electrical and Computer Engineering  
University of Virginia, Charlottesville, VA 22904 USA

{mah6s, hcp7ad, atb4c, jlach}@virginia.edu

## ABSTRACT

Energy-fidelity tradeoffs are central to many battery-constrained systems, but they are essential in body area sensor networks (BASNs) due to energy and resource constraints, and the critical nature of many healthcare applications. On-node signal processing and compression techniques can save energy by greatly reducing the amount of data transmitted over the wireless channel, but lossy techniques can incur a reduction in application fidelity. In order to maximize system performance, these tradeoffs must be considered at run-time due to the variable nature of BASN application, including sensed data, operating environments, user actuation, etc. BASNs therefore require energy-fidelity scalability, so automated and user-initiated tradeoff decisions can be made dynamically.

This paper explores the utility of energy-fidelity scalability in BASNs from a data-centric perspective. Compression algorithms are identified that can be implemented on resource constrained BASN nodes and that have “knobs” capable of trading off compression ratios (and resulting transmission energy) with fidelity. To demonstrate the potential of energy-fidelity scalability on a real BASN and for a real application, the tradeoff space is established by adjusting these algorithms for different movement disorder data sets collected by a custom accelerometer-based BASN. Finally, mechanisms for energy-fidelity dynamic control are explored.

## Keywords

Body area sensor networks, energy-fidelity scalability, Haar wavelet compression, sample rate-resolution adaption, quality of service

## 1. INTRODUCTION

A confluence of advancements in diverse areas of research, including device integration, energy storage, sensor technology, and wireless communications, have facilitated the creation of body area sensor networks (BASNs) [1]. BASNs have made an early and significant impact in healthcare applications because of intrinsic and unique capabilities of the technology – networked body area sensor nodes continuously capture objective measures

of human physiology and performance both inside and outside of traditional healthcare settings [2]. For example, research efforts at the University of Virginia have resulted in the development of a BASN platform, TEMPO (Technology Enabled Medical Precision Observation), which is an integral component of three clinical studies related to movement disorder assessment [3-4]. BASNs like TEMPO are helping clinicians improve healthcare assessment accuracy and precision for better diagnosis, treatment, and assistance.

Continued evolution of this nascent field requires compelling new applications of the technology, as well as research to address formidable practical barriers to adoption, such as form factor, battery life, and quality of service (QoS). Additionally, BASNs must be minimally-invasive to users. Unlike typical wireless sensor network (WSN) implementations, BASNs have fewer, smaller nodes. The lack of sensor redundancy increases the importance of preserving the fidelity of sensed data, as do safety-critical considerations of some healthcare applications. Moreover, diminutive packaging bounds battery size and therefore, capacity. Battery lifetime and application fidelity are competing metrics, as the majority of the energy consumption in BASN applications is due to the wireless transmission of sensed data [5]. Therefore, BASN nodes either suffer from poor battery life or operate at reduced data rates to decrease time-on-radio and lower system energy needs. If the reduction of data rate is governed by lossy processes, fidelity likely suffers.

Thus, there exists a need for design space expansion in the BASN domain that creates efficient tradeoffs between energy and fidelity. For example, some sensing tasks require longer battery life at lower data rates and lower fidelity (e.g. activity monitoring). Other sensing tasks require much higher data rates and fidelity at the expense of battery life (e.g. gait analysis). However, many applications have dynamically changing requirements [6], which incur momentary increases in data rate over long periods of low data rate operation (e.g. fall detection). Furthermore, run-time factors, such as the variability of the body area wireless channel [7], the availability of harvested energy, and the desire for graceful degradation of function as the battery drains, necessitates dynamic control over the energy-fidelity tradeoff.

This paper explores the utility of energy-fidelity scalability in BASNs from a data-centric perspective – those that involve the processes of signal acquisition and digital signal processing – through the following contributions:

- **Identification of energy-fidelity scalable algorithms for BASNs:** The paper identifies two algorithmic data rate reduction techniques: Haar wavelet compression and sample rate-resolution adaption. Both algorithms

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*BodyNets'09*, April 1–3, 2009, Los Angeles, CA, USA.  
Copyright 2008 ICST 978-963-9799-41-7.

are capable of operating in severely resource constrained environments such as those found in BASNs. Within each approach, tunable “knobs” that enable the tradeoff of energy and fidelity are explored.

- **Demonstration of energy-fidelity tradeoff space:** Selected data rate reduction algorithms are implemented on a mixed signal processor found on the TEMPO BASN node. Demonstration of the tradeoff space is illustrated using real BASN data collected by TEMPO in a clinical study of Essential tremor.
- **Proposal of energy-fidelity dynamic control:** Two on-node and two off-node techniques are presented, which create opportunities for dynamic control over “knobs” associated with both schemes. The on-node approaches introduce fidelity-estimation and control-theoretic mechanisms. The off-node techniques extend the theme of feedback to the aggregator and system levels.

The results are presented within the context of a derived processor energy model used for approximation of energy consumption, and for two measures of fidelity: mean squared error (MSE) and percent RMS distortion (PRD). As discussed in Section 2, these information-theoretic measures are not necessarily representative of application fidelity, but they offer examples for how scalability can be analyzed in a given BASN setting. The analysis uses data collected by the TEMPO system in clinical studies of movement disorder; however, these techniques are applicable to the larger domain of applications and BASN platforms.

The remainder of this paper is organized as follows. Section 2 discusses issues related to the identification of fidelity metrics. Section 3 presents the tunable compression techniques that enable the dynamic tradeoff of energy and fidelity. A generalized processor energy model is described in Section 4. Section 5 demonstrates the energy-fidelity tradeoff space using the identified algorithms and the collected TEMPO data, and Section 6 presents the mechanisms for dynamic control. Section 7 provides conclusions and opportunities for future work.

## 2. FIDELITY

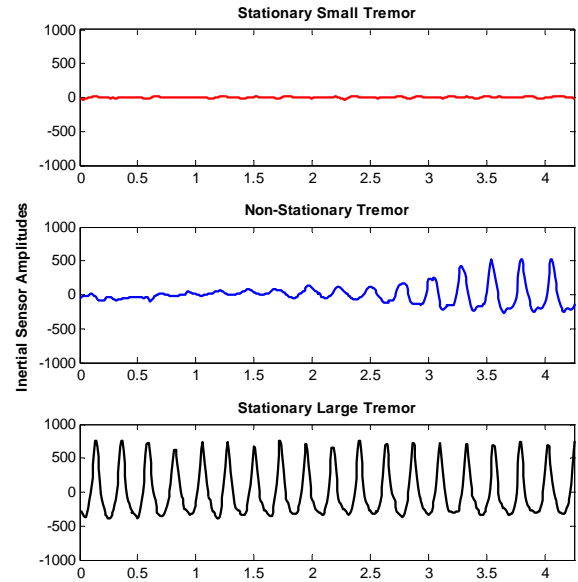
Data reduction is an important goal of BASN systems, as energy is saved when the radio does not have to transmit as frequently. Arbitrary reduction of data rate, however, will increase distortion and decrease fidelity. Although accepted measures of distortion and fidelity exist in information theory, the fidelity measures of concern are ultimately related to application fidelity.

For example, BASN applications in healthcare come with strict application-level requirements, including precision, accuracy, reliability, etc. But ultimately, application fidelity measures, such as misclassification rates for feature extraction, false positive and false negative rates for event detection, and a number of other quantifiable measures, are what define a system’s success. In other words, blindly using MSE or PRD as node-level fidelity metrics without determining their relationships to application fidelity could lead to poor system performance.

However, it is often difficult to evaluate – especially at the sensor node level – the impact that changes in signal processing and data reduction will have on such high-level fidelity metrics. Application-specific analyses are therefore required to explore how turning certain fidelity-compromising “knobs” affects

application fidelity metrics, thus enabling the creation of models between data reduction techniques and fidelity metrics that can be calculated at different levels in the BASN, such as MSE and PRD.

Given that this paper is providing an application-agnostic analysis of data-centric energy-fidelity scalability in BASNs, we must resort to information theoretic fidelity metrics. However, even the use of these metrics reveals complexities. For example, MSE and PRD often provide very different assessments of fidelity, which further emphasizes the need to identify the most appropriate fidelity metrics for each BASN application.



**Figure 1. Tremor Data Collected by TEMPO**

Tremor analysis provides an example of the impact that fidelity metric selection has on energy-fidelity scalability. The single-axis acceleration data in Figure 1 was collected by the TEMPO system (120 Hz sampling at 12-bit quantization) on patients diagnosed with Essential tremor. Each plot spans ~4.3 seconds of data (6144 bits) extracted from an hour long data set. The top plot represents a small stationary tremor, the bottom plot represents a large stationary tremor, and the middle plot illustrates the same patient transitioning from a small to a large tremor.

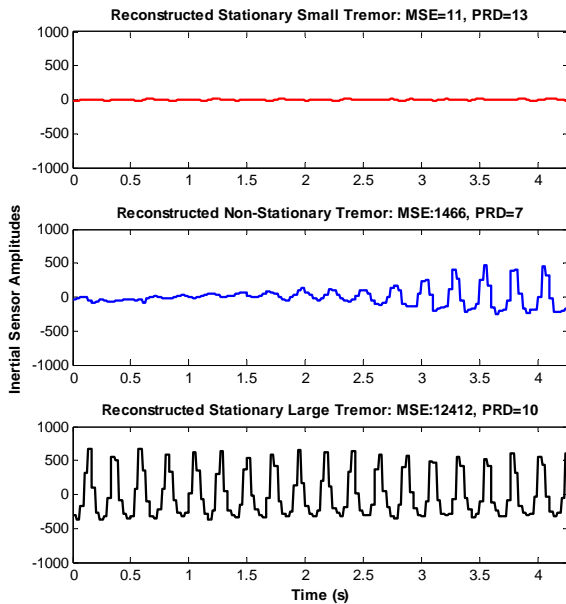
The signals in Figure 1 were compressed using the Haar wavelet compression algorithm described later in this paper. The algorithm yielded a nearly 4:1 compression ratio for all three signals. The resulting post-compression plots are shown in Figure 2 and were assessed with both MSE and PRD as follows:

$$MSE = \frac{1}{N} \sum_{n=1}^N [x_{uncompressed}(n) - x_{compressed}(n)]^2 \quad (1)$$

$$PRD = \frac{100 * MSE}{\frac{1}{N} \sum_{n=1}^N x_{uncompressed}^2(n)} \quad (2)$$

While both metrics are generally accepted measures of fidelity from an information theory perspective [9], their assessments of the waveforms in Figure 2 were completely different. Using MSE as the metric, the Haar compression hardly affected the small

stationary tremor and horribly distorted the large stationary tremor. Conversely, PRD assessed the small stationary tremor as being the most distorted by the compression.



**Figure 2. Reconstructed Tremor Data**

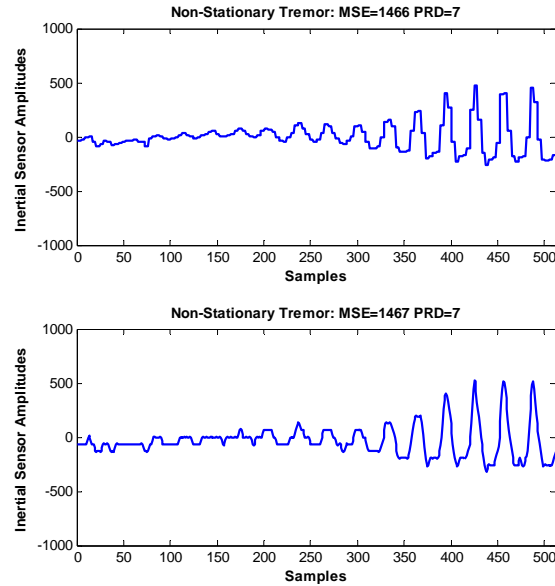
In addition, the use of different compression techniques can impact how metrics such as MSE and PRD are evaluated. Figure 3 shows the small to a large tremor transition after being compressed using the Haar wavelet algorithm (top) and a rate-resolution algorithm (bottom). Both signals have nearly identical MSE and PRD; however, the top signal appears much more quantized than the bottom.

As discussed above, application-specific analyses must be performed to determine if these or other information theory fidelity metrics correlate well to the application fidelity metrics that ultimately matter for users. Such analyses are central to ensuring that energy-fidelity scalability techniques achieve the desired results.

### 3. ALGORITHM IDENTIFICATION

Data rate reduction is a pervasive goal of networked embedded systems, but identifying algorithms that are suitable for highly resource constrained environments remains challenging. Many mixed signal microcontrollers lack the rich resources found even in portable consumer electronic devices. BASN node resource constraints include: limited memory, clock frequency, storage, and even lack of co-processors such as floating point units. Therefore, computationally complex, yet highly effective, transform based coding is not a viable option. Data rate reduction on BASN platforms must meet three criteria:

- capable of being implemented on resource-constrained BASN embedded processors,
- capable of executing in low-latency and soft real-time applications, and
- adjustable by key “knobs” to alter expected data reduction rates.



**Figure 3. Comparison of MSE to PRD for Haar Wavelet Compression (Top) and a Rate-Resolution Algorithm (Bottom)**

Data rate reduction is quite simple, but controlling how much data is reduced (and the resulting impact on the defined fidelity metric) is more challenging. For example, simple run-length encoding, a lossless form of compression that exploits repeated values, achieves different compression ratios for different input data streams. Other efforts [8] have evaluated lossless data compression (e.g. bzip2, compress, LZO, PPMd, etc.) and have identified memory footprints ranging from 3 KB to 6 KB to achieve compression ratios of 1.4 to 5 in times ranging from ~0.5 s to 12 s on the StrongArm SA-110 processor. These levels of resources and performance are not well suited to far less capable BASN processors operating in soft real-time applications. Identifying algorithms suitable for BASN implementation and capable of producing deterministic compression ratios based on tuning “knobs” is therefore a focus of this effort.

Perhaps the simplest scheme to reduce data rate involves decimation and quantization of data. Decimating will reduce the data rate by the decimation factor. Figure 4 demonstrates different data compression options pertaining to decimation and quantization. For instance, successive decimation at a rate of 2, results in equivalent 2X reduction of data if no quantization is present.

In a typical mixed signal processor execution environment, such as the Texas Instruments MSP430, integer multiple decimation factors are simple to implement and place minimal overhead on execution times. Before decimation, incoming signals are filtered to a rate of:

$$F_{in\_bw} < \frac{F_{sampling}}{2D} \quad (3)$$

where the maximum frequency present in the original sampled data ( $F_{in\_bw}$ ) is less than the sampling frequency ( $F_{sampling}$ ) divided by twice the decimation factor ( $D$ ). By taking advantage of the multiply accumulate (MAC) peripheral, a feature found in many

mixed signal microcontrollers, a high-order finite impulse response (FIR) low-pass filter is employed prior to decimation in order to reduce aliasing artifacts. The FIR filter is sub-sampled at the decimated rate, reducing the processor overhead by a proportional factor.

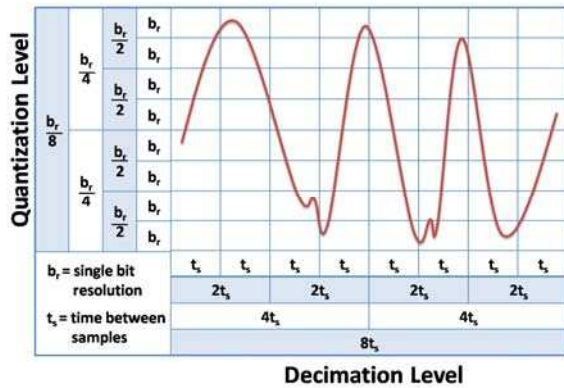


Figure 4. Sample Rate-Resolution Adaption

Beyond sample rate scaling, data rate may be further reduced by changing the sample resolution (quantization), as shown in Figure 4. Computationally efficient on a mixed signal microprocessor, this form of scaling removes successive least significant bits from each sample (e.g., retain the 8 most significant bits and discard the 4 least significant bits in a 12 bit sample). It is assumed for both evaluation purposes and practical implementation that a more resource-rich aggregator or other post-processing apparatus would perform an interpolation to restore the signal to its original rate. Calculation of fidelity in Section 5 is performed using a cubic spline interpolation on the compressed data.

A second form of compression utilizes the Haar discrete wavelet transform (DWT) discussed in Section 2, which performs successive sub-band decompositions using quadrature mirror filters [10]. The Haar DWT is attractive for data compression in a resource-constrained environment because coefficients may be readily computed using only basic arithmetic operations such as addition, subtraction, and logical shifts. The algorithm completes in  $O(N)$  operations. Furthermore, the Haar DWT creates opportunities for multi-resolution access and spatio-temporal data mining on-node [11], which are desirable for context awareness and highly selective data reduction in-network.




Figure 5. Haar DWT Compression

Decompositions are performed iteratively to increase frequency resolution. To limit the size of output coefficients, successive transforms are only performed on the approximation coefficients representing lower frequency information.

Cascaded as described, the Haar DWT produces a joint-time frequency representation of the data with an output equal in size

to the input (Figure 5). To achieve deterministic data rate reduction, low-order detail coefficients (starting at  $d_1$ ) are zeroed. Following the Haar DWT with a run-length encoding (RLE) scheme removes the repeated zero coefficients. Thus, by increasing the level of decomposition, more compression is possible, but at the expense of fidelity.

To summarize, data rate-resolution scaling is a simplistic and computationally efficient form of data rate reduction that features two knobs: decimation factor and quantization bits removed. Haar wavelet compression also reduces data rate by adjusting the level of decomposition required along with the number of detail coefficients removed. All four knobs produce deterministic compression ratios and are therefore more easily modeled within an energy-fidelity scalability framework.

#### 4. ENERGY MODEL

With fidelity metrics and compression algorithms identified, the energy model is now formalized to complete the energy-fidelity scalability analysis. The model is necessary to evaluate the performance of our energy-fidelity scheme with previously collected medical data from our TEMPO platform, and to isolate individual components (e.g. radio, sensors, mixed signal processor) that contribute overheads to the total energy function, yet cannot be measured in isolation on our integrated platform. Additionally, this energy model is generalizable to other BASN platforms.

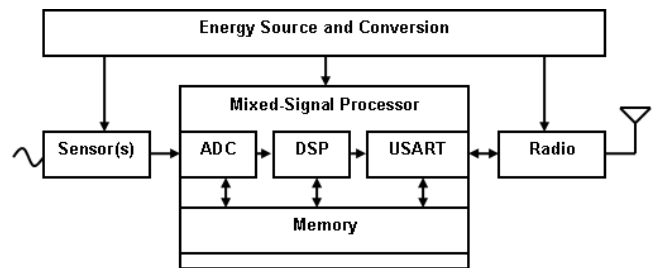


Figure 6. BASN Node Architecture

Similar to WSNs, typical BASN node implementations are designed to encompass sensing, signal processing, and communication functionality (Figure 6). The central component of the BASN node architecture is the mixed-signal processor. This resource-constrained microprocessor may integrate a multi-channel analog to digital converter (ADC) for sensor output conversion, digital signal processing (DSP) capability, and a Universal Synchronous-Asynchronous Receiver/Transmitter (USART) peripheral for radio interfacing. Control logic is typically coordinated by a lightweight operating system.

Unlike WSNs, BASN devices are physically constrained by packaging capable of being worn, implanted, or ingested. Therefore, battery capacity is limited by the small form factor of the enclosure. As an example, a mote-like WSN node powered by two AA batteries will have a capacity of  $\sim 20$  kJ; however, a BASN node inside of a wrist watch enclosure could be limited to as few as  $\sim 2$  kJ using a common lithium coin cell battery. This places additional demands on BASN designs to tame conspicuous power consumption.

In prevailing wireless sensor implementations, power-consumption is dominated by the radio [5]. For a BASN with a star topology network and nodes communicating directly to a base

station, the nature of communication is usually asymmetrical – i.e. the nodes generate and transmit the majority of network traffic. Therefore, modeling radio energy consumption for most BASN applications can be simplified by neglecting receiver power and assuming fixed-interval, round-robin polling. The energy consumed by the radio (Figure 7) can be modeled as a follows [12]:

$$E_{radio} = P_{tx} \left( \frac{L_{tx}}{R_{tx}} + T_{tx\_st} \right) + P_{out} \left( \frac{L_{tx}}{R_{tx}} \right) \quad (4)$$

where  $P_{tx}$  is the power consumed by the transmitter circuitry,  $P_{out}$  is the power consumed by the power amplifier,  $L_{tx}$  is the bit length of the packet to transmit,  $R_{tx}$  is the bit rate, and  $T_{tx\_st}$  is the start-up time for the modulation circuitry to settle.




Figure 7. Radio Power Model

Duty cycling of the radio is essential for decreasing power consumption in BASN nodes. A power-gated radio can be duty-cycled between full-on and full-off states. Therefore, energy is saved when the radio is not in use, and as a result, not consuming power. Typical values for  $P_{tx}$ ,  $P_{out}$ , and  $T_{tx\_st}$  derived from a commercial 2.4 GHz RF transceiver [13] and normalized for packet length, illustrate that radio energy is minimized by transmitting larger packets at higher bit rates; thus, reducing radio start-up overhead (Figure 8).

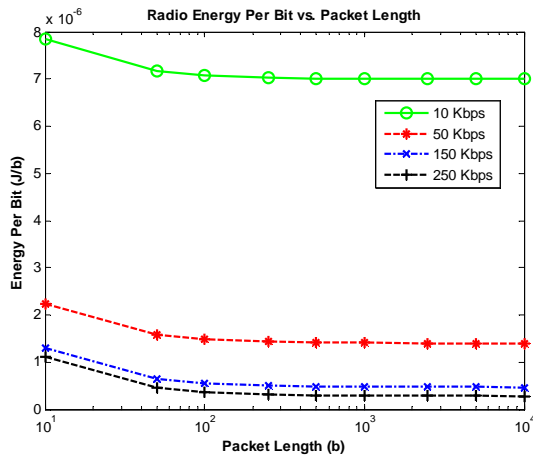


Figure 8. Energy Consumed by a 2.4 GHz IEEE 802.15.4 Radio Transceiver

BASNs in the healthcare domain are often designed to capture and forward data to a base station for post-processing into clinically-significant, non-obvious information. However, in other application domains, data reduction may be equally important. The volume of information extracted may only represent a fraction of the volume of data collected in-network. Therefore, reducing the quantity of transmitted data has the potential to increase node battery life without sacrificing relevant information.

Such a reduction can be lossless or lossy. Although lossy source coding increases distortion of the input signal post reconstruction, it also yields higher compression ratios, on average, than lossless alternatives (e.g. run-length encoding, Huffman encoding). Lossy source encoding must be carefully implemented to reduce data rates for wireless transmission without significantly sacrificing application fidelity due to distortion. The implementation must analyze the rate-distortion performance over the application data to make appropriate decisions about the design of a source coder.

The energy needed to compress data using a mixed-signal processor ( $E_{dsp}$ ) adds to the total node energy dissipation:

$$E_{total} = E_{radio} + E_{dsp} \quad (5)$$

As energy is dissipated in processing, data is compressed, and like its radio counterpart, the processor can sleep when idle. Dynamic power management (DPM) is also used by the processor to move between ultra-low power sleep modes and active mode power.

Data reduction is measured by a compression ratio (CR), defined as the size of the uncompressed message ( $L_{uncompressed}$ ) divided by the compressed message ( $L_{compressed}$ ):

$$CR = \frac{L_{uncompressed}}{L_{compressed}} \quad (6)$$

The time needed to compress an arbitrary-size message, or collection of samples on the microprocessor, is related to the CR for specific compression algorithms. This time ( $T_{\mu\_compress}$ ) is an aggregation of the time to capture samples from the ADC ( $T_{capture}$ ), the time to process the samples ( $T_{process}$ ), the time spent waking the microprocessor from its sleep state ( $T_{st\_mic}$ ), and the time to transfer the message to the radio FIFO buffer over USART ( $T_{tx\_fifo}$ ).

$$T_{\mu\_compress} = T_{capture} + T_{process} + T_{st\_mic} + T_{tx\_fifo} \quad (7)$$

The USART time is non-trivial, and for some modern low-power mixed signal microprocessors is an order of magnitude slower than that of the supported radio data rate. This overhead must be considered when designing low-power systems.

The processing energy ( $E_{dsp}$ ) is therefore shown in Equation 8:

$$E_{dsp} = P_{\mu\_active} (T_{\mu\_compress}) + \dots \dots + P_{\mu\_sleep} \left( \frac{N_{samples}}{R_{sampling}} - T_{\mu\_compress} \right) \quad (8)$$

where ( $P_{\mu\_sleep}$ ) is the sleep mode power consumption of the microprocessor and ( $N_{samples}$ ) is the number of samples taken over a given sampling rate ( $R_{samples}$ ).

Finally, the following relation between the length of the transmitted packet and the compression ratio is evident, which includes the ADC sample resolution ( $N_{bits/sample}$ ):

$$L_{tx} = \frac{N_{samples} * N_{bits/sample}}{CR} \quad (9)$$

In summary, a higher CR, larger packet size, and higher active mode radio transmission bit rates will yield the largest energy

savings. Despite the energy advantages of data compression, high fidelity reconstruction and high CR are competing objectives for lossy source coding.

## 5. DEMONSTRATION OF SCALABILITY

To demonstrate the feasibility of our data-centric, energy-fidelity scalable algorithms, we employed the TEMPO BASN platform at the University of Virginia. TEMPO 2.0 (Figure 9) is custom BASN node, based on the previously described reference architecture (Figure 6). The device consists of a TI MSP430F1611 mixed-signal processor, a Mitsumi 2.4 GHz radio transceiver, a Kionix KMX52-1050 monolithic tri-axial MEMS inertial sensor, and front-end signal conditioning circuitry. The filtered sensor outputs are captured by the on-board 12-bit ADC. Processed data is forwarded to the FIFO buffer of the USART-interfaced radio transceiver for base-station offload.




Figure 9. TEMPO 2 BASN Node

The two data reduction schemes described above were tested using the three data sets collected by the TEMPO system for a study of deep-brain stimulation efficacy in tremor control (Figure 1). The 512 sample data sets, captured at 12 bits of resolution, represent an upper-bound for message length capable of being held in memory by a typical BASN microprocessor such as the MSP430.

Data rate reduction was measured by a compression ratio and included in measures of average energy-per-bit. Distortion of the data was measured by mean squared error (MSE). To compare two compression schemes, we adjusted the following “knobs”: Haar wavelet decomposition level and number of zeroed coefficients, sampling decimation rate, and sampling resolution. Finally, both data reduction methods were implemented on the TEMPO platform’s TI MSP430F1611 microprocessor. Processing energy was calculated from the previous energy model (Equation 8). Delays associated with processing (Equation 7) were measured by a Tektronix TPS2024 triggering on toggled port pins between sections of the data reduction firmware. Fixed delays are:  $P_{\mu c\_active}=8.1mW$ ,  $P_{\mu c\_sleep}=24.6\mu W$ ,  $T_{capture}=13.2\mu s$ , and  $T_{st\_\mu c}=6\mu s$ . The processor was operated at a 4 MHz clock rate, while the supply voltage was set to 3.3 V. The wireless transmission energy-per-bit of  $2.82 \mu J$  was estimated for a 2.4 GHz transmitter using the 250 Kbps plot and associated packet lengths above (Figure 8). This value assumes that the radio buffers the entirety of a message before sending it across the body area channel.

Figure 10 illustrates two plots: the fidelity vs. total energy per incoming bit for both Haar wavelet compression and sample rate-resolution adaption. All three tremor signals are plotted for both schemes. The top plot shows the highest energy per incoming bit

existing at approximately  $2 \mu J$  for all three signals. These data points correspond to a zeroing of only the  $d1$  coefficients in a level-8 Haar decomposition. As successively more  $d$  coefficients are removed, the energy-per-bit is reduced, however, so is the fidelity, as measured by the inverse MSE. There are diminishing returns in reducing energy-per-bit as the higher level  $d$  coefficients have fewer bits than the lower coefficients (Figure 5). By adjusting the number of zeroed bits, a knob in this scheme, a BASN designer would be able to move along the energy fidelity curve for the Haar scheme. By zeroing the first 4  $d$  coefficients, a designer would like save half the total energy-per-bit of a system, resulting in an increase in battery life. These results are comparable to results of other energy/resource constrained compression schemes (i.e. S-LZW) [14], yet the Haar DWT also enables multi-resolution processing on-node, completes in less time ( $\sim 12$  ms), and has a smaller computational and memory footprint.

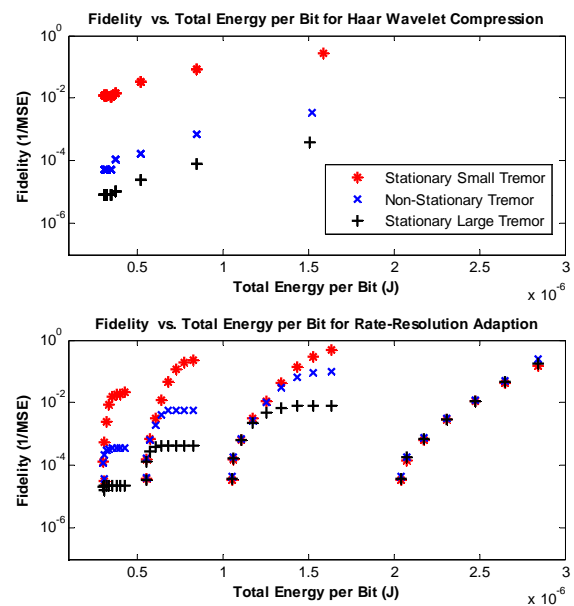


Figure 10. Demonstration of Energy-Fidelity Scaling

The bottom plot for of Figure 10 illustrates four distinct energy-fidelity curves that correspond to each level of decimation. Curves start at the right hand side of the plot and are decimated by a factor of 2, 4, and 8 respectively. At the top of each decimation curve, are samples quantized to full resolution (i.e. 12 bits). Decreasing the resolution of the samples moves the energy-fidelity point down within a decimation level curve. This result suggests that BASN designers should decimate as much as possible while remaining cognizant of Nyquist bounds for their application. Once at the appropriate level of decimation, the sampling resolution can be tuned.

Both plots reveal two important observations about enabling data-centric energy-fidelity scalability in BASNs:

- **A large energy-fidelity design space for BASN nodes exists and is ripe for exploration.** Placed in context, the results clearly demonstrate that energy-fidelity scalability is a feasible option for resource-constrained BASN nodes and real data. Knobs for Haar wavelet compression and for sample-rate resolution adaption

move BASN operation points to new levels of fidelity and energy-per-bit. More work is needed to characterize the sensitivity of the knobs and to create control schemes that are appropriate for measures of application fidelity and not necessarily MSE.

- **Energy-fidelity curves will change over time as captured signals fluctuate.** The three tremor signals each have different energy-fidelity curves for both data reduction schemes. These signals, however, are in actuality a single axis of inertial data from one BASN node. Therefore, there exists a need to dynamically tune knobs to adjust to situational observations. Moreover, fidelity estimation at the node may create opportunities to adapt to changes in signal characteristics.

The total BASN node energy consumption is often assumed to be dominated by the radio. Figure 11 plots the total power consumed by the processor to achieve various compression ratios for both data reduction schemes. The figure paints a new picture of BASNs energy use for data-centric, scalable operation. At compression ratios between 2 and 20, the processor begins to consume more than 50% of the node's total energy. This finding implies that there exists a need to explore both energy efficient processing as well as energy efficient transmission in current and future BASN designs.

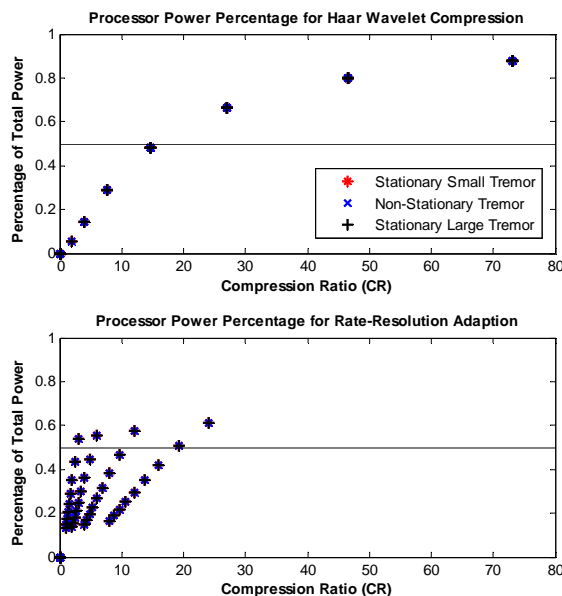


Figure 11. Processor Power Percentage

## 6. DYNAMIC CONTROL

A key aspect of achieving optimal energy and fidelity tradeoffs in BASN nodes is the inclusion of a dynamic control scheme which adapts both to the incoming data from the sensors, as well as the highly variable radio frequency medium employed for data transmission. The locus of this control mechanism may reside both on-node and off-node. For example, in a typical BASN environment, several sensor nodes may feed data to an aggregator at a central location on the body (i.e. the belt or pocket). The aggregator node usually has a larger form factor, more energy

availability, and a more capable processor than that found on a sensor node. This resource asymmetry suggests employing a hierarchical scheme in which the aggregator may provide feedback to the sensor nodes in an effort to promote an optimal energy versus fidelity strategy.

First, consider data-centric control strategies that originate on-node. With proper input scaling, one node could simply make a decision based on the amplitude or frequency of a sensed variable and periodically broadcast a short status message indicating that no significant signal is present. However, when significant amplitude is reached then control strategies must exhibit a higher degree of intelligence.

Another simplistic strategy may consider the relative stationarity of the observed data. Haar wavelet coefficients may be compared in successive time periods and at various decomposition levels. With little to no significant change in coefficients from one block of data to the next, it may be considered stationary in a wide-sense. The coefficients may be sent once, and then follow up messages again indicate that there has been no significant change since the previous transmission. This method has the advantage of only transmitting a representation of the signal infrequently, with a high degree of fidelity, and if the signal remains relatively stationary, then little extra energy is required.

As the input signal transitions from stationary to non-stationary, the control strategy must again be changed. A possible solution may be to employ wavelet/run length compression or bit depth scaling. Such a decision must be driven, in part by the availability of bandwidth for transmission, and any pre-established priorities on reconstructed signal fidelity. This may call for making on-node control decisions according to the availability of bandwidth in the radio transmission system. One possible software approach for implementing the interface between data acquisition and radio transmission is the use of a first-in-first-out (FIFO) queue. Conveniently, the availability of the transmission medium may be determined by measuring the level of queue occupancy.

The availability of this variable facilitates a decision strategy based on control-theoretic approaches. One may establish a preferred level of queue occupancy, and use the current level of occupancy to develop an error indication that in turn drives a conventional control loop. The rate of change of this queue occupancy error may also be easily determined, which facilitates a control strategy based on the common proportional-integral-derivative (PID) method, which is both very effective as well as implementable in a typical microcontroller execution environment. The magnitude of the feedback signal derived from these three terms may be used to drive further data compression, or increase radio power, providing that the limit in power has not been reached.

The aggregator may be considered a second locus of control for driving decisions based on energy and fidelity requirements. This is especially attractive as it alleviates algorithmic complexity on-node and may serve as a further check and feedback mechanism for testing the efficacy of on-node decisions as well as adding sensor fusion information to and control decisions.

In one scenario the aggregator may request the node to send a full uncompressed block of data as well as the current compressed version for comparison. The aggregator could then further classify these results as being satisfactory or not for the type of signal being received. For example, it may be critical not to lose

information pertinent to a potential fall and should the aggregator sense this condition, it may advise the node to employ less compression and more radio power until the situation is ameliorated. If the recovered movement data is clearly in the normal range, then the aggregator may indicate that the node should go back to higher compression levels. In both of these scenarios, advantage is made of the higher computational abilities of the aggregator to perform signal classification and use the results of this classification to feed information down the hierarchy, and augment on-node strategies. The aggregator may even send its classification results back to the node, enabling a rudimentary form of self adjustment.

We envision a system of interacting control strategies as being desirable. For example the aggregator may assist the node in making decisions to change the level of compression, while the node may inform the aggregator of the availability of on-node energy resources, i.e. the battery, in making decisions to increase radio power. Optimizing the decision space for these variables remains an area of active research.

## 7. CONCLUSION

BASNs represent a highly dynamic execution environment and as such, provide unique challenges to researchers. Demands for delivering ever increasing battery life while delivering high quality data represent opposing factors in their design; additional form factor requirements are also an imposing restriction, directly affecting battery size and therefore energy availability.

In order to meet these challenges, a dynamically adaptable execution model is necessary. The ability to alter compression ratios while maintaining adequate fidelity of the recovered signal requires careful attention to both distortion as a function of the algorithms employed as well as the capabilities and limitations of the microcontroller employed. Computationally efficient lossy compression techniques such as Haar wavelet transform compression and rate-resolution adaption are capable of acceptable distortion figures when appropriately applied on a dynamic basis to variable signals. Moreover they provide energy savings on-par with other resource-constrained compression techniques, yet they may also enable context awareness.

In order to maximize the effectiveness of a dynamically flexible signal processing and transmission chain, the control system must possess a great deal of intelligence. We envision hierarchical feedback control allowing nodes and aggregators to coordinate – with the aggregator having the ability to request operational level changes from individual nodes, and the nodes having the ability to respond to the aggregator with current functional abilities, such as remaining battery life.

Future research will encompass rigorous characterization and optimization of the control scenarios. In addition energy-fidelity efficient partitioning of the processing load between the node and aggregator remains an open research question.

## 8. ACKNOWLEDGEMENTS

The authors would like to thank the anonymous reviewers for their helpful suggestions. This work is supported in part by the National Science Foundation under grant Nos. CBET-0756645 and IIP-0646008, the ARCS Foundation, Philips Research North America, and the University of Virginia Institute on Aging.

## 9. REFERENCES

- [1] G. Yang, Ed., *Body Sensor Networks*, New York: Springer, 2006.
- [2] E. Jovanov et al., “A wireless body area network of intelligent motion sensors for computer assisted physical rehabilitation,” *Journal of NeuroEngineering and Rehabilitation*, vol. 2, 2005.
- [3] M.A. Hanson et al., “Teager energy assessment of tremor severity in clinical application of wearable inertial sensors,” *Life Science Systems and Applications Workshop (LISA)*. IEEE/NIH, 2007, pp. 136-139.
- [4] H.C. Powell, M.A. Hanson, and J. Lach, “A wearable inertial sensing technology for clinical assessment of tremor,” *Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 2007, pp. 9-12.
- [5] O. Landsiedel, K. Wehrle, and S. Gotz, “Accurate prediction of power consumption in sensor networks,” *The Second Workshop on Embedded Networked Sensors (EmNetS-II)*. IEEE, 2005, pp. 37-44.
- [6] V. Raghunathan, P. Spanos, and M. Srivastava, “Adaptive power-fidelity in energy-aware wireless embedded systems,” *22nd Real-Time Systems Symposium (RTSS 2001)*. IEEE, 2001, pp. 106-115.
- [7] D. Neiryneck, “Channel characterisation and physical layer analysis for body and personal area network development,” Ph.D. dissertation, University of Bristol, Bristol, England, 2006.
- [8] K.C. Barr and K. Asanović, “Energy-aware lossless data compression,” *ACM Transactions on Computer Systems*, vol. 24, 2006, pp. 250-291.
- [9] S.M. Ahmeda and M. Abo-Zahhad, “A new hybrid algorithm for ECG signal compression based on the wavelet transformation of the linearly predicted error,” *Medical Engineering & Physics*, vol. 23, Mar. 2001, pp. 117-126.
- [10] Y.T. Chan, “Wavelet Basics,” Kluwer Academic, Boston, 1995, pp. 25-50.
- [11] D. Ganesan, D. Estrin, and J. Heidemann, “Dimensions: why do we need a new data handling architecture for sensor networks?,” *ACM SIGCOMM Computer Communications Review*, vol. 33, 2003, pp. 143-148.
- [12] E. Shih, S. Cho, N. Ickes, R. Min, A. Sinha, A. Wang, and A. Chandrakasan, “Physical layer driven protocol and algorithm design for energy-efficient wireless sensor networks,” *7th Annual International Conference on Mobile Computing and Networking (MobiCom)*. ACM, 2001, pp. 272-287.
- [13] Texas Instruments, CC2500 Low-cost low-power 2.4 GHz RF transceiver, Datasheet, 2007. Available: <http://www.ti.com/lit/gpn/cc2500>.
- [14] C.M. Sadler and M. Martonosi, “Data compression algorithms for energy-constrained devices in delay tolerant networks,” *4th international conference on Embedded networked sensor systems (SenSys)*. ACM, 2006, pp. 265-278.