J. Low Power Electron. Appl. **2011**, *1*, 109-130; doi:10.3390/jlpea1010109

OPEN ACCESS

Journal of Low Power Electronics and Applications

ISSN 2079-9268

www.mdpi.com/journal/jlpea/

Review

Energy Efficient Design for Body Sensor Nodes

Yanqing Zhang [1,*], Yousef Shakhsheer [1], Adam T. Barth [1], Harry C. Powell Jr. [1], Samuel A. Ridenour [1], Mark A. Hanson [2], John Lach [1] and Benton H. Calhoun [1]

[1] The Charles L. Brown Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA 22908, USA; E-Mails: yas5b@virginia.edu (Y.S.); atb4c@virginia.edu (A.T.B.); hcp7ad@virginia.edu (H.C.P.); sar7f@virginia.edu (S.A.R.); jlach@virginia.edu (J.L.); bcalhoun@virginia.edu (B.H.C.)

[2] Wireless Health Interactive, LLC, Vienna, VA 22182, USA; E-Mail: mah6s@virginia.edu

* Author to whom correspondence should be addressed; E-Mail: yanqing@virginia.edu; Tel.: +1-434-466-3957; Fax: +1-434-924-8818.

Received: 24 November 2010; in revised form: 21 March 2011 / Accepted: 7 April 2011 / Published: 11 April 2011

Abstract: This paper describes the hardware requirements and design constraints that derive from unique features of body sensor networks (BSNs). Based on the BSN requirements, we examine the tradeoff between custom hardware and commercial off the shelf (COTS) designs for BSNs. The broad range of BSN applications includes situations where either custom chips or COTS design is optimal. For both types of nodes, we survey key techniques to improve energy efficiency in BSNs and identify general approaches to energy efficiency in this space.

Keywords: body sensor networks; sub-threshold circuits; wearable computing; energy efficient design

---

1. Overview of BSNs

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 sensor networks (BSNs). BSNs—networked body area sensor nodes that continuously capture objective

measures of human physiology and performance both inside and outside of traditional healthcare settings [1]—are making an impact in numerous applications and enhancing the possibilities of wireless health systems. For example, the TEMPO (Technology Enabled Medical Precision Observation) BSN platform (Figure 1) at the University of Virginia is an integral component of three clinical studies related to movement disorder assessment, including tremor in Parkinson's Disease patients, physical agitation in dementia patients, and fall risk in the elderly [2]. The TEMPO platform sensor node integrates MEMS inertial sensors (linear acceleration and rotational rate in six-degrees-of-freedom), a mixed-signal processor, Bluetooth wireless transmission, and a lithium battery in a wearable enclosure. BSNs like TEMPO are helping clinicians improve healthcare assessment accuracy and precision for better diagnosis, treatment, and assistance of movement disorders.

Figure 1. TEMPO packaged (cover removed) and wrist mounted.

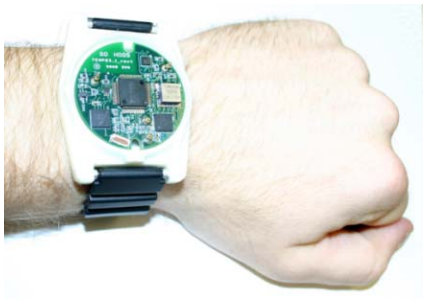

BSNs have made a significant impact in healthcare applications because of intrinsic and unique capabilities of the technology—networked body area sensor nodes capture and process precise and accurate data continuously and non-invasively. With manifold sensing opportunities paired with healthcare information technology infrastructure, body sensor networks also enable longitudinal and naturalistic monitoring, enhancing the assessment of human physiology and performance both inside and outside of traditional healthcare settings. Thus, BSNs address the weaknesses of traditional patient data collection, such as imprecision (qualitative human observation) and under-sampling (infrequent assessment). Beyond assessment, BSNs are finding use in medical actuation. For example, an artificial pancreas can regulate blood glucose via implanted biochemical sensors located in interstitial fluid [3]. Moreover, future applications of deep brain stimulation, pacemaker regulation, drug delivery, and prosthetic control will be made possible by body area sensor networks that can provide real-time assistance based on real-time assessments. Finally, BSNs are creating opportunities in the telehealth domain [4]. As medicine moves beyond the confines of hospitals and clinics, technological infrastructure will be necessary in the home. BSNs are poised to be primary enablers to this movement. While BSNs open new horizons to healthcare delivery, such technology is also helping to define new entertainment and wellness applications in the commercial space.

BSN Requirements

BSNs have tremendous potential to transform how people interact with and benefit from information technology, but their practical adoption must overcome formidable technical and social challenges (e.g., form factor, battery life, reliability, safety, security, privacy, interoperability, ease of use, *etc.*). These challenges have far-reaching implications but offer many immediate opportunities for

system design and implementation [5]. Although BSNs share many of these challenges and opportunities with general wireless sensor networks (WSNs)—and can therefore build off the body of knowledge associated with them—many BSN-specific research and design questions have emerged that require new lines of inquiry. Unlike generic WSNs that have many nodes doing the same thing, BSNs are likely to have a small number (<10) of nodes with each node dedicated to a specific task. For example, a sensor node monitoring acceleration at the ankle for gait analysis clearly cannot also measure brainwaves using an EEG since both the location and sensing hardware are so different. To achieve widespread adoption, BSN nodes must be extremely noninvasive, which means that the nodes must have a small form factor that is not overly inconvenient to use. Smaller nodes imply smaller batteries, creating strict tradeoffs between BSN node energy consumption and the fidelity, throughput, and latency requirements of BSN applications.

Therefore, while the diverse BSN application space results in wide ranging system requirements, all BSN applications—whether real-time or delay insensitive, continuous high data rate streaming or infrequent small packet bursts, *etc.*—demand energy efficiency while meeting data fidelity requirements. The battery size *versus* battery life tradeoff plays a major role in defining any BSN system, and applying design techniques to reduce energy consumption can improve both size and lifetime. If energy consumption can be reduced far enough, perpetual operation on harvested energy becomes a possibility. Thus, BSN node sensing, processing, storage, and wireless transmission must all be done in a way that reliably delivers the important data but with the lowest possible energy consumption, thus minimizing battery size (which dominates BSN node form factor) and maximizing time between battery recharge (which is a key factor in wearability), both of which can impact the performance and practicality of possible applications.

The best approach for optimizing the tradeoff between energy consumption and other requirements varies depending on the specific BSN application. To illustrate this, we can consider the tradeoff between battery lifetime (e.g., the application requirement for how long the system must work between re-charges) and effective wireless communication data rate (e.g., the average rate at which data from the node must reach the base station) across different applications. Figure 2 estimates how different applications map to this tradeoff space. Some applications like pulse oximetry (measures saturation of peripheral oxygen, SPO_2), ambulatory blood pressure (ABP), or electromyography (EMG) for muscle activity require monitoring lifetimes between an hour (e.g., in the clinic) and a day or two (assuming sensors can be recharged at night), but the quantity of data that must be transferred varies dramatically. Continuous glucose monitoring (CGM) sensors may need to have lifetimes approaching a month, but they do not need to send much data on average. Some RFID-like sensors may only need to work for a second after being queried by a base station acting as a reader, but some long term sensors implanted in the body or incorporated into clothing may need to last for years. This great variety in requirements defies a single solution to solve the energy constraint problem.

For life-critical applications that require continuous high fidelity sensed data for real-time assessment and intervention (e.g., fall detection, heart arrhythmia detection, *etc.*), which would be very costly to transmit wirelessly, reduction or elimination of wireless transmission may be necessary to meet longer battery life and wearability requirements. Such applications may need to make intervention/actuation decisions on-node and only employ wireless transmission when events of interest are detected. This system level design decision will help to reduce node power consumption

sufficiently to satisfy the other system requirements. BSNs for delay insensitive applications, such as those employed by clinicians to gather information in large volume, may alternatively leverage lower power on-node storage, rather than wireless transmission, to increase battery life. Such store-and-forward use cases, including Holter monitoring and activity logging, are capable of acquiring high fidelity data for later assessment off-node. In such cases, on-node processing is limited, as more resource rich or expert assessments are made off-node. Finally, real-time applications involving wireless transmission and high fidelity data (e.g., gait analysis, activity monitoring, gaming, *etc.*), combine on-node signal processing with radio management to meet battery life demands of hours to days.

Figure 2. Broad design space for BSN, but size limits energy for all applications.

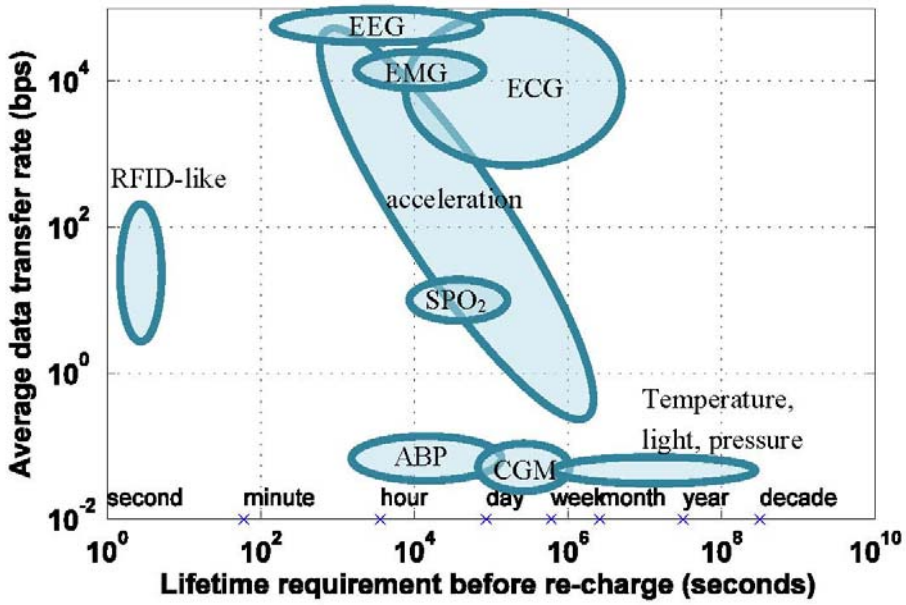

Value to the user will ultimately determine each technology's success. BSNs must effectively transmit and transform sensed phenomena into valuable information and do so while meeting other system requirements, such as energy efficiency. The value of a BSN therefore rests in large part on its ability to selectively process and deliver information at fidelity levels and rates appropriate to the data's destination, whether that is to a runner curious about her heart rate or a physician needing a patient's electrocardiogram. These disparate application requirements require the ability to aggregate hierarchical information and integrate BSN systems into the existing information technology infrastructure. Increased value of the BSN to the user will also increase user tolerance of non ideal wearability or other technological difficulties.

In this paper, we describe methods for developing efficient hardware within the unique set of requirements of BSNs for different parts of the BSN application design space. Due to the ubiquitous and strict energy constraint on all BSNs, we focus on energy efficiency. In addition, the approaches for achieving energy efficiency in BSNs designed with COTS components sometimes differ from those designed with custom hardware, and this paper explores both paradigms. Finally, this analysis is done within the context of current and projected BSN applications and use cases.

2. Hardware Selection—Commercial off the Shelf or Custom?

As we will describe in more detail in Section 3, flexible hardware is inherently less efficient than hardware targeted for a specific application. Considering the broad range of requirements in Figure 2 combined with a strict energy constraint, there is little hope for a single system platform to support the full range of applications. General purpose sensing platforms like the motes of the WSN community are typically too large in size or consume too much power to meet application requirements in the BSN domain.

The first implementation choice for a BSN application is whether to use commercial off the shelf (COTS) components or whether to design custom integrated circuits (ICs). We observe that both COTS and custom silicon solutions are optimal in certain parts of the BSN design space. Custom ICs can invariably provide a higher energy efficiency than COTS when they are tailored for a specific application. For this reason, BSN applications with lifetimes exceeding roughly one to two weeks (on the right in Figure 2) will most likely require custom hardware. Size also plays a role. For example, COTS systems can provide longer lifetimes if a larger solution (providing larger batteries) is acceptable for a given application. We illustrate this relationship in Figure 3. Users are likely to remove certain types of BSN sensors on a nightly basis, providing an excellent opportunity to recharge the battery and limiting the required lifetime to less than a day. For systems with short lifetime requirements between re-charges (e.g., a few days or less), COTS systems can provide reasonable size, performance, and lifetime for many BSN applications. These generalizations do not account for the fact that each application also has requirements based upon sensed data rates, communication and real-time requirements, sensor properties, and other factors that influence the boundary between using COTS *versus* custom in Figure 2 and Figure 3.

Figure 3. Hardware design space for COTS *versus* custom circuits.

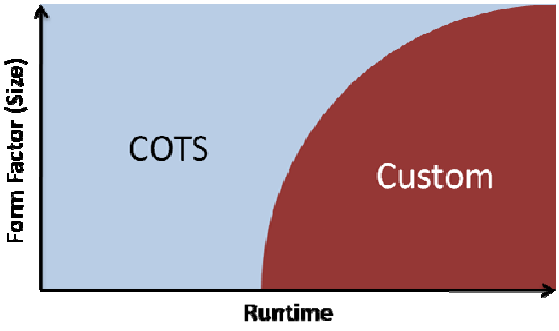

There are many applications for which there may be no clear line of demarcation in the decision process for selecting a COTS or custom ASIC (application specific integrated circuit) design. For example, ECG measurements require no transducer element other than a contact with the skin and may be implemented as a single purpose device with a clear and well defined scope of operation—all characteristics amenable to ASIC solutions. For long-term continuous ECG monitoring, for example to screen for cardiac arrhythmia, very low energy operation is essential, demanding an ASIC solution. However, Holter monitors are an example of a very successful and widely used 24-hour ECG recording system that can easily use COTS and have no need for an ASIC solution. Inertial measurements require the use of a MEMS device as the sensor and are not easily integrated into a

single chip ASIC design due to the fundamentally different manufacturing processes required for MEMS *versus* electronics. BSN sensors that measure acceleration may need to incorporate multiple ASICs or a mixture of ASICs and COTS.

The final decision between a custom design *versus* a COTS design must account for the previous points combined with the economics of the intended application. Designing a COTS system is orders of magnitude faster and cheaper than building a custom IC based node, and COTS nodes provide excellent solutions in many lower lifetime BSN scenarios. For example, low volume research platforms or nodes intended for short term clinical monitoring applications may be more economically produced via a COTS design. In such applications, the final device operational characteristics are much less well defined, and engineering costs are ongoing. In this case the economies of reducing such costs via the employment of a flexible platform outweigh the benefits of extra efficiency that an ASIC solution would offer. For example the TEMPO3 system mentioned in Section 1 may be reprogrammed to operate in a clinical environment in which continuous data streaming is a requirement, or in a more longitudinal study in which data may be stored on node and offloaded after an extended measurement session. Additionally, COTS devices have steadily been improving in computing performance. When TEMPO1 was introduced in 2006, the processor employed had 48 kB of flash memory, 2 kB of RAM, and operated at a maximum clock frequency of 8 MHz. There are now available pin compatible drop-in devices from the same family that have over 100 kB of flash, 8 kB of RAM, and are capable of operating at 20 MHz within similar power budgets. This clearly leads to an expanded application space for a given sensing technology, with little or no non-recurring engineering (NRE) costs for hardware design. High volume, single purpose, mass market devices favor ASIC approaches in which the NRE costs are amortized over many units. Even if COTS systems provide a weaker solution (e.g., by limiting lifetime) than ASICs, simple economics will make COTS the better choice for many BSN applications that cannot provide the volume required to justify an ASIC solution.

3. General Strategies for Energy Efficient BSN Hardware

We have emphasized that many design decisions depend on the specific BSN application in question, but we can also identify general strategies that should influence any BSN design. In this section, we examine several key tradeoffs that affect BSN design and that provide important opportunities for saving energy regardless of the specific BSN application. Specifically, we examine balances between on-node computation and communication, flexibility and efficiency, and data fidelity and energy consumption. Before describing these tradeoffs, we introduce supply voltage management as a means of energy minimization in circuits, which provides an important foundation for custom energy efficient circuit design of energy constrained systems like BSNs.

3.1. Supply Voltage Management

Lowering the supply voltage to a circuit is a well known approach for reducing energy. In this section, we first discuss the limit of lowering voltage to reduce energy consumption, and then describe how dynamic voltage scaling can allow us to tradeoff energy and performance.

For digital circuits, energy of computation varies as the square of the supply voltage (V_{DD}), which makes it desirable to operate at the lowest possible voltage while preserving functionality and meeting

timing constraints. Taking this principle to the extreme, we observe that sub-threshold (sub-V_T) operation of digital integrated circuits provides one important option for energy efficient processing. Sub-V_T circuits use a V_{DD} that is below the threshold voltage, V_T, of the transistors. This makes the transistors "off" by conventional definitions, but the change in transistor gate-to-source voltage (V_{GS}) produces a difference in sub-V_T conduction current that allows static digital circuits to operate robustly, although slower than they would be at higher voltage. The lower speeds are still more than sufficient for many BSN operations (up to 10's of MHz). Both the off-current and the on-current of the transistors vary exponentially with V_{DD} in the sub-V_T region ($V_{GS} < V_T$). Nevertheless, the on-current in sub-V_T remains larger than the off-current by enough (1000× or so) to enable proper functionality of the digital gates. Due to the quadratic relationship between energy and V_{DD}, the main advantage of sub-V_T operation is a reduction in energy consumption of over 10× compared to traditional circuit implementations. In fact, sub-V_T operation has been shown to minimize energy per operation in conventional CMOS circuits [6]. For this reason, sub-threshold operation will play an important role in custom hardware for BSNs.

There are some challenges to making sub-V_T digital circuits work. Most notably, the reduced I_{on}/I_{off} ratio combines with process variations in the threshold voltage to increase the potential for circuit failure. Sub-V_T circuits also must be level converted to interface with super-V_T design, such as radios or sensors. Additionally, design of sub-V_T circuits is not yet commonplace. Standard cells used in designs are rarely designed for this voltage of operation, in which transistor strengths change. Nevertheless, sub-V_T operation is an emerging approach that is very useful for BSN nodes [7].

Figure 4. Energy-workload curve of normal operation and dynamic voltage scaling (DVS).

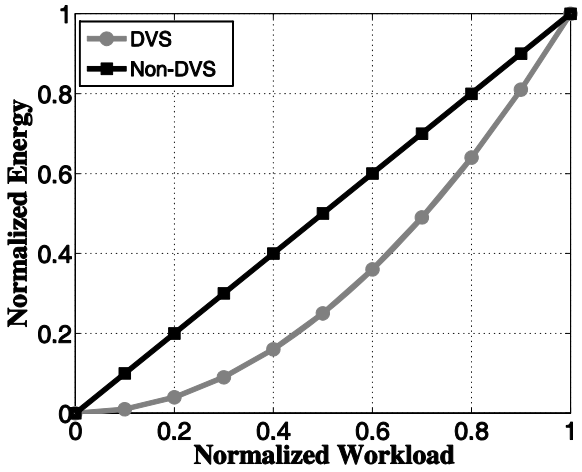

Operating at a low voltage all of the time may not be a viable option for all BSNs, because lower voltages slow down circuit speed. Given that a BSN's processing latency and throughput requirements may change during execution in response to real-time data and mode changes, dynamic voltage scaling (DVS) can be employed to minimize V_{DD} given those requirements. When high performance is necessary to meet system level requirements, the circuits can operate at the energy-costly higher voltage level. By reducing the circuit's V_{DD}, quadratic energy savings can be achieved instead of just the linear savings obtained through power gating (Figure 4). Different DVS schemes propose different approaches to scaling in terms of the circuit topology and interval at which the voltage is changed, and

the overhead of most schemes are minimal compared to the energy savings accomplished, especially when that scaling includes dropping to sub-V_T levels when permissible [8].

Some COTS chips provide built-in DVS capabilities or allow for development of DVS schemes. For instance, the TI MSP430 and other similar microcontrollers (MCUs), have on-board clock generation hardware that allows the MCU to programmatically change the operating clock frequency. This is accomplished in the MSP430 through the use of a Digitally Controlled Oscillator (DCO) that may be calibrated using a low frequency (32 kHz) watch crystal as a reference. Frequency agility is accomplished by switching different programmable constants into two clock control registers. The actual change in clock frequency occurs within approximately 10 µs. Furthermore, this microcontroller operates over a wide range of voltages. The clock oscillator may be varied over a 16 to 1 range and the supply voltage over a 2 to 1 range. Within this envelope, a combination of processing rate and power requirements exist, making COTS embedded processors of this type ideal candidates for inclusion in a DVS scheme for BSNs. Figure 5 shows potential DVS operating points measured for the MSP430F2131 processor as explored in [9].

Figure 5. Operating points for a COTS MCU [9].

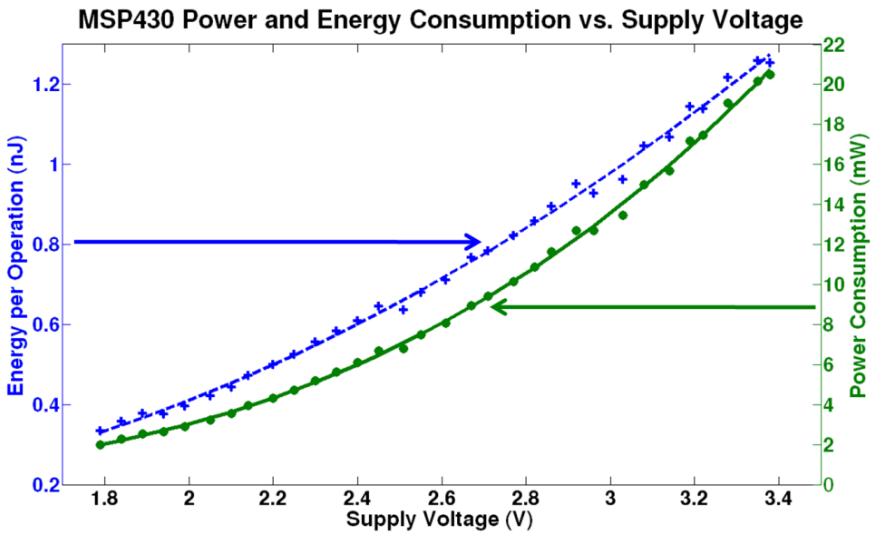

For longer lifetime BSN applications where the savings from Figure 5 are still inadequate, a similar DVS scheme can apply to a custom chip. Figure 6 compares a custom MCU design [7] to the MSP430. The custom design offers a 100× improvement in energy per instruction. However, this does not come free of tradeoffs. In this case, the custom built MCU does not have its own clock generation hardware, and frequency agility is not as straight forward since there is only one, single frequency main clock. Though this custom designed MCU also operates over a wide range of voltages and is capable of supporting DVS, additional design effort is required to build in these operating modes. What's more, custom designed hardware does not enjoy the complete suite of mature and compatible peripherals as COTS components, which degrades custom hardware's flexibility.

Figure 6. Energy-delay curves for DVS in a COTS microcontroller and a custom design.

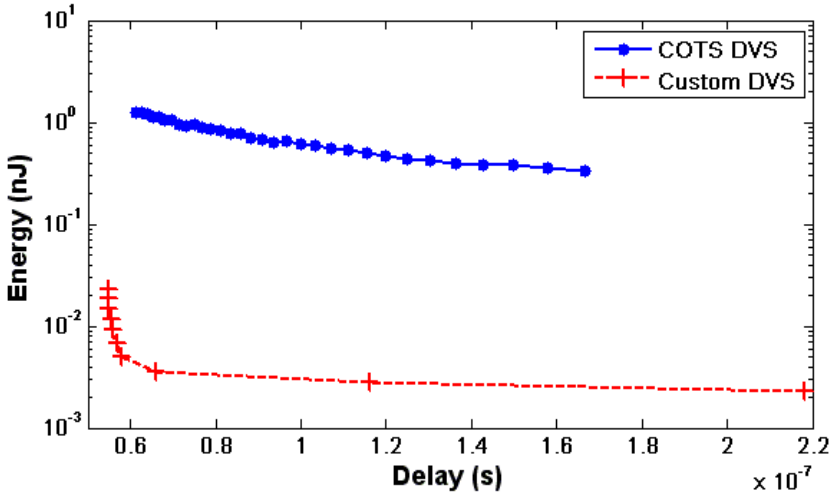

3.2. Communication versus *Computation*

As is the case in most WSNs, wireless transmission of sensed data is the largest power consumer in most current BSNs [9]. This problem is particularly acute in medical BSN applications, in which sensor data rates may be high relative to many WSN applications. Figure 7 illustrates this relationship with the COTS TEMPO platform as an example, where the high power consumption of the Bluetooth transceiver swamps the low power consumption of the TI MSP430 microcontroller during raw data transmission. We could improve this situation by using a lower power radio (e.g., COTS implementing a different protocol, or a custom design), by duty cycling and sending data in bursts, or by other strategies. In this section, however, we focus on the strategy of using computation on the node to reduce the cost of communication, which can influence all types of BSN design regardless of hardware choice.

Figure 7. TEMPO 3.1 power consumption breakdown (with gyroscopes off).

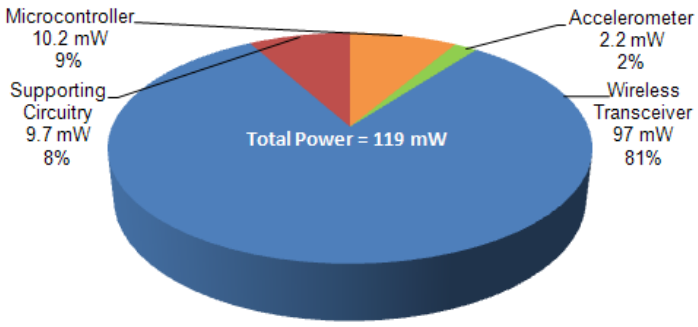

Significant power reduction can be achieved through the development of on-node signal processing and data management which can dramatically reduce the number of bits to be transmitted. By reducing the number of bits to transmit, we effectively allow more substantial duty cycling of the radio (e.g., leaving it off for a larger fraction of the time). Methods to reduce communication data include traditional compression along with advanced signal processing techniques such as pattern classification and feature detection algorithms. Low power signal processing therefore becomes increasingly

important to BSN power efficiency. We can quantify the impact of this tradeoff on the overall node energy using a simple energy model. Assume that E_r is a ratio of the average energy to transmit one bit (E_{tx}) to the average energy to process one bit (E_{proc}). This ratio is typically large (*i.e.*, $E_r >> 1$) and is determined by a number of factors, including processor energy per operation, the signal processing algorithm and implementation, the packet organization and coding, the networking protocol, transmit power, *etc.* Also assume that the compression ratio (CR) achieved by on-node signal processing is the ratio of the number of raw bits to the number of transmitted bits. The ratio of average processing energy (E_{proc}) to average total energy (E_{total}) is therefore:

$$\frac{E_{proc}}{E_{total}} = \frac{E_{proc}}{\frac{E_{tx}}{CR} + E_{proc}} = \frac{1}{\frac{E_r}{CR} + 1} \quad (1)$$

Figure 8 plots this ratio as a function of CR for different values of E_r. It is clear that the importance of low power signal processing increases with more effective pre-transmission compression techniques, even at high E_r ratios. As a point of reference, 25 nJ/bit is typical for state of the art custom Bluetooth radios targeting 1 Mb/s [10,11], and an MSP 430 consumes roughly 1 nJ/bit. If the system had no other processing costs (e.g., ignoring memory, *etc.*), it would have an E_r of only 25, indicating that processing energy becomes quite important if the CR is even 10. Applying simple generic compression schemes could compress raw data streams by this amount. Even more substantial compression is possible by extracting important features on chip and only transmitting those instead of the raw data. This motivates the need for reducing the hardware energy costs of on-node computation, especially with custom radio solutions. As we described above, sub-threshold operation is one excellent method for decreasing E_{proc} by over 10× compared to operation at the nominal V_{DD}.

Figure 8. Percentage of total energy contributed by on-node signal processing for different $E_r = E_{tx}/E_{proc}$ ratios as a function of the bit compression ratio (CR) [9].

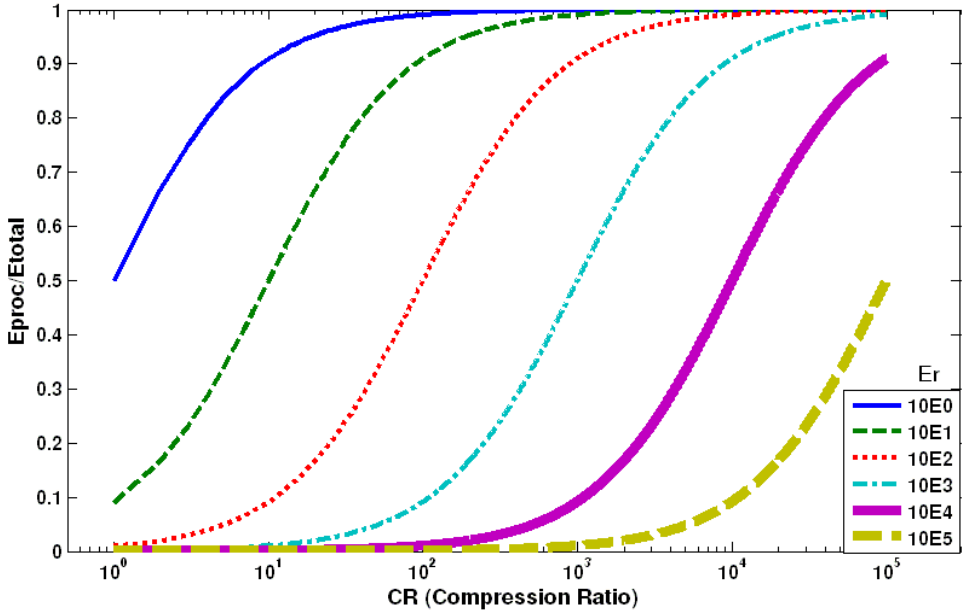

The power required to transmit data wirelessly can be high even with reduced duty cycle, so the radio remains a critical component even with on-die data compression. Therefore, it is worthwhile to understand the requirements for BSN radios. BSNs have slightly different radio requirements compared to radios in typical WSNs with low data rate communication for monitoring environmental conditions. For BSN applications, the network is probably arranged as a star-hub topology with the hub acting as a base station [1]. All communication from the nodes is to the hub, which can be assumed to have substantially more resources than the nodes (e.g., the hub may be a smart phone). Therefore, the nodes need a small, low power, and short range (1–2 m) radio. This means that a BSN radio could be optimized to operate at a much lower transmission power compared with radios designed for WSN applications. Judging from Figure 2, the ability to accommodate variable communication data rates will be important. Also, since we have observed that processing to reduce the communication data rate allows us to save power by using the radio less often, the radio should save energy when operating at lower data rates or should allow energy efficient transitions to and from active mode. We note that turning off the radio for longer times creates the need to re-synchronize the radio when it turns back on. Since the hub has more resources than the nodes, it can remain active permanently to listen for communication from the nodes. Most radio traffic in a BSN system is from the nodes to the hub, alleviating the need for the nodes to run a receiver continuously looking for messages from the hub. However, the cost of synchronization may still be significant depending on the specific BSN application and communication protocol.

BSN applications can span a large range of data rates from a few bits per minute to almost 1 Mbps depending on the application [5]. Currently, there are a few low power radio and radio protocols such as ANT and Zigbee that are commonly used in wireless sensor networks. However, these radios and protocols can only operate with data rates of 10 kbps and 150 kbps, respectively. This severely limits their usefulness for the upper range of BSN applications such as motion assessment, ECG (electrocardiogram), EMG (electromyogram), and EEG (electroencephalogram). Conversely, the high data rate COTS radios and protocols, such as Wi-Fi have data rates which easily cover the entire span of BSN applications. However, these radios and protocols consume so much energy that they are impractical for use on BSNs with longer battery life requirements. Bluetooth is a radio and protocol that sits somewhere in between high data rate and low data rate radios. The protocol uses a large amount of energy due to the fact that it was designed as a very general purpose radio for applications spanning outside of the area of BSN. Bluetooth is convenient for BSN development platform purposes because of its widespread adoption, but its relatively poor energy efficiency leaves room for optimization with custom radios and protocols. Nevertheless, Bluetooth is a convenient and viable option for lower lifetime BSN applications. Due to inefficiencies in existing radios and protocols, there are other protocols that are being developed to accommodate the area of BSNs such as 802.15.6. This protocol specifically targets body sensor devices and the medical applications that can span a wide range of data rates [12].This new protocol supports data rates greater than 850 kbps and allows flexibility for the PHY layers supporting ultra wide band, narrow band, medical implant communication bands, and human body communication PHY. While this new protocol is not a standard yet and is still in working group, it is being designed to be more efficient for BSNs compared to existing options, and their main challenges will be providing efficient access for the broad range of

BSN applications (see Figure 2) and achieving the pervasiveness that Bluetooth and Wi-Fi have achieved in smart phones and other personal computing devices.

While there are many COTS radios and protocols available today that are serviceable for BSN applications, there is still a large opportunity for custom radios that provide better energy efficiency for the applications of the BSN community. These radios can take advantage of the small transmission distance and asymmetry of the channel and provide a substantial benefit to the power consumption of the devices. For example, [13] presents a 830 pJ/bit 2.4 GHz radio that can transmit with a data rate of 500 kbps. [14] presents a 2 Mbps low power receiver that consumes 0.18 nJ/b and [15] shows a 0.65 nJ/b 100 kbps receiver at 1.9 GHz. For sub 1 GHz transmission, [16] shows a 1 Mbps OOK transceiver that operates at 10 nJ/bit with very fast startup time of 2.5 μs to allow for efficient duty cycling. All the previously mentioned low power transmitters and receivers take advantage of short range requirements of BSNs and consume much less energy compared to common COTS radios such as Bluetooth and Zigbee. With these improvements in energy consumption, BSNs can run much longer on a single battery charge or the device can be made smaller by allowing the same runtime but with a smaller battery. It is worth noting that since standards for BSNs are still under development, a concrete guideline for low power radios is not readily available.

3.3. Flexibility versus *Efficiency*

The tradeoff between flexibility and efficiency in hardware is well known and very prominent in a comparison of conventional hardware paradigms [17,18]. The most flexible category of hardware is general purpose processors (GPPs). GPPs exhibit poor energy efficiency due to the overhead of fetching and decoding the instructions that are required to perform a given operation in the datapath. For low power embedded applications like BSNs, general purpose computation is generally performed in fairly simple microcontrollers [7,19–21]. Sophisticated operations like a fast Fourier transform (FFT) or data processing algorithm will thus require numerous instructions in the simple core. For example, several sub-threshold processors provide energy per instruction nearing 1 pJ per operation, but they also tend to use small instruction sets and thus result in more instructions to run an operation [7,19–21].

The most efficient hardware is hardwired to do its specific task or tasks (e.g., ASIC). ASICs achieve very efficient operation, but they can only perform the function for which they were originally defined. Examples of hardwired implementations in sub-threshold circuits include [22–25]. Different types of hardware in sub-threshold systems reveal a similar trend as their above-threshold counterparts. Microcontrollers like the one in [19] consume as low as 2.6 pJ/instruction and provide excellent flexibility since they can be reprogrammed for arbitrary tasks. The ASIC implementation of a JPEG co-processor in [24] consumes 1.3 pJ/frame for VGA JPEG encoding. The numbers for energy/operation are similar, but the individual operations on the microcontroller (e.g., instructions) are simple integer computations like addition. Executing a complete JPEG encoding would take many (100s or 1000s) instructions on such a light weight processor, making the total energy per frame much higher than on the ASIC. Of course, the GPP can perform a much broader range of tasks than the JPEG encoder, so this comparison exemplifies the tradeoff between energy efficiency and flexibility.

Some BSN nodes may be implemented as complete ASICs like the JPEG processor, but more commonly, ASICs may appear in BSNs as auxiliary hardware accelerator modules, performing commonly occurring functions in the context of a larger system on chip (SoC). Good examples of hardware acceleration are multipliers, floating point units, or FIR filters. These operations can take several instructions over many clock cycles to complete using a GPP, consuming a large amount of energy and time. A hardware accelerator can process data quickly and efficiently. Here, these commonly used components take advantage of the energy and computational efficiencies of the accelerators, whilst their designs need not change. Hardware accelerators provide an opportunity to process data in very specific ways more efficiently than on accompanying programmable hardware.

Microprocessor operations are largely inefficient, as we described above. Field Programmable Gate Arrays (FPGAs) are reprogrammable hardware that provide an intermediate choice between ASICs and processors in terms of flexibility and efficiency. An FPGA is configured to act like specific hardware, similar to an ASIC, but the configuration can be changed an arbitrary number of times. The cost of this flexibility is that FPGAs consume 10~100 times more energy than an ASIC due to energy overhead from interconnects, which may account for 85% of the total energy consumption. Most commercial FPGAs target high performance applications to compete with processors, but a sub-threshold FPGA [26] demonstrates that custom FPGA implementations can offer a good tradeoff for flexibility and energy efficiency for energy constrained applications like BSNs.

To demonstrate the performance of different hardware platforms in the context of BSN applications, we simulated a typical heart rate (R-R) extraction algorithm that calculates the heart rate of a user based on the raw data of an ECG, which was run on the three different platforms designed in the same technology operating at the same operating voltage (0.4 V) while targeting the same data rate. The results are shown in Table 1.

We make two observations from Table 1. First, not only is ASIC > FPGA > GPP with respect to energy efficiency, but ASIC > FPGA > GPP in terms of potentially speed and performance capacity. The second observation is that there is a drastic improvement in efficiency (>100×) between GPPs and FPGA/ASICs. Therefore, it makes sense to assign on-node processing to FPGA and ASIC platforms, while using GPPs strictly for control or rarely occurring operations.

Table 1. Comparison of different hardware platforms.

	Energy per Instruction	**Energy per Processed Sample**	**Delay per Sample**	**Estimate Max Achievable Data Rate**	**GOPS/W**
GPP (from [7])	2.62 pJ	210 pJ	8 μs (80 clock cycles)	125 kHz	4.76
FPGA (from [26])	N/A	2.22 pJ	94.5 ns (1 clock cycle)	10 MHz	450
ASIC	N/A	0.23pJ	6.18 ns (1 clock cycle)	150 MHz	4348

Given the large space of BSN nodes and their applications, there is no obvious optimal platform for all nodes. Though ASICs are extremely efficient in terms of energy minimization and computational capability, they are highly inflexible as their functionality is set. Thus, they must be revisited and

redesigned whenever the functionality changes. This is a major drawback, as it leads to increased design time and design cost. Furthermore, ASICs are limited to a certain application space. Therefore, flexibility is another requirement for BSNs that must be examined during the design of a node for a specific application or set of applications.

On the other end of the spectrum, GPPs offer a highly flexible option for on-node processing. Along with popular peripherals, such as the aforementioned floating point unit or multiplier, GPPs are able to perform almost any job and run any processing algorithm for the BSN node. Thus, they are useful in building most nodes, serving as a central controller for the node. The flexibility advantage is most noticeable in generic nodes, where the specific algorithm or signal processing requirements are not pre-determined, but coded into instruction memory. However, this advantage comes at the cost of energy efficiency. GPPs are highly inefficient because of unused logic components and resources within the GPP for each instruction executed. Also, given the instruction per cycle limitations of GPPs, programs cannot fully take advantage of instruction parallelism, resulting in greater latency and energy consumption per package of data processed. State-of-the-art low power COTS GPPs can meet energy and speed requirements for many BSN applications. For example, TI's MSP430 supports a wide range of applications, running on clock frequencies up to 25 MHz while consuming 165 μA/MHz [27]. Custom GPPs will be even more efficient but will incur the development costs of an ASIC.

In summary, increasing the flexibility of processing to cover more scenarios will sacrifice energy efficiency. This means that platforms encircling larger regions of Figure 2 will necessarily be less efficient than more targeted solutions, resulting in shorter lifetimes and/or larger form factors.

3.4. Data Fidelity versus *Energy*

The last key tradeoff we will explore involves looking into how much processing and communication is necessarily needed and relevant in an application. Previous work has shown the existence of an energy-fidelity tradeoff in BSNs with digital signal processing employed to examine tremor in a Parkinson's patient [28]. This research used Haar wavelet compression and rate-resolution scaling as an example lossy data reduction scheme for use in exploring the tradeoff space since it met the following three criteria:

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

Mean Squared Error (MSE) was used to assess fidelity as is commonly done in the signal processing community. The results indicated there is a large energy-fidelity exploration space possible in BSNs. Figure 9 shows a small portion of this space using the Haar wavelet transform and run length encoding for data compression and highlights another interesting fact: the input signal characteristics change the possible energy-fidelity operating points.

Moreover, it is interesting to note that the data shown in Figure 9 is from a single patient over the course of a single clinical visit. The amount of "information" present in the sensor signals changes over time along with the rate-distortion curve pointing to the need for dynamic management of energy-fidelity tradeoffs in these embedded environments. To illustrate further, Figure 10 depicts a

time domain distortion plot for fixed data compression, yielding a compression ratio (CR) of approximately 18, for a 40 minute tremor dataset. Thus, merely choosing a static operating point on a curve of Figure 9 is not sufficient for application fidelity regulation or energy efficiency. Instead, runtime adjustment of processing methods should be performed for more optimal, data-centric operation.

Figure 9. Preliminary energy-fidelity tradeoff results [29].

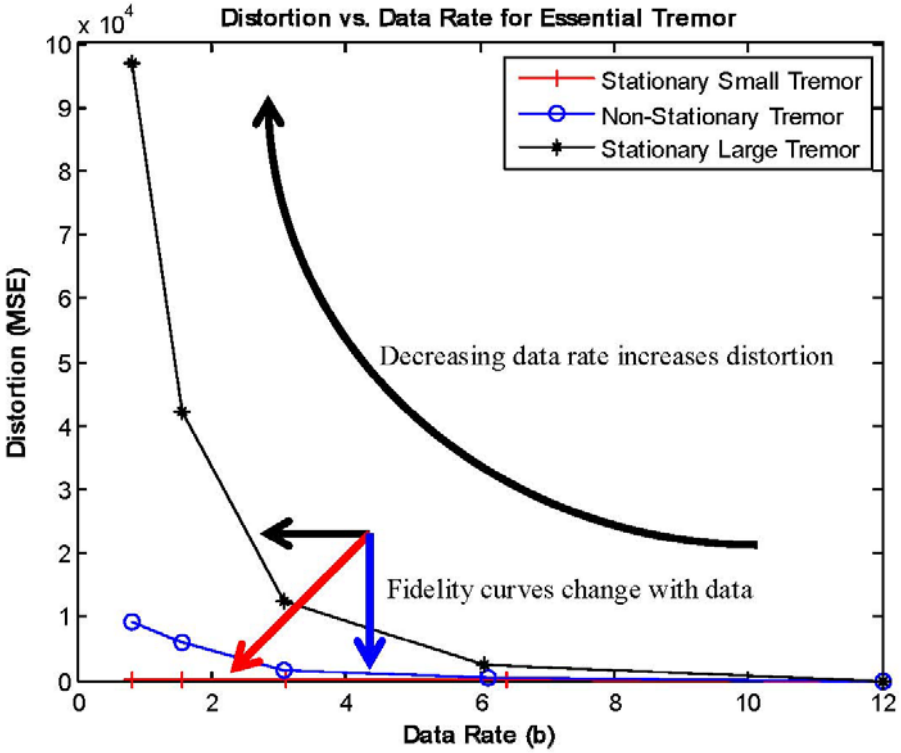

BSN devices must therefore possess energy awareness (knowledge of how much energy has been consumed), data awareness (knowledge of how compression affects current data), and computational resource awareness (knowledge of how algorithm execution affects processing and memory resources) to effectively tradeoff runtime and output fidelity in a way that is executable on resource constrained platforms and that meets real-time requirements. These tradeoff decisions can be made based on efficiently meeting requirements (e.g., maximum lifetime for a given minimum fidelity, maximum fidelity for a given minimum lifetime, *etc.*) or minimizing bounded cost functions (e.g., minimizing $\text{lifetime}^{-\alpha} \cdot \text{fidelity}^{-\beta}$ given minimum lifetime and fidelity requirements, where α and β are determined based on metric priorities).

BSN devices also need adaptable and efficient data rate scaling mechanisms to fully exploit energy-fidelity tradeoffs at the node-level in real-time. For instance, if a MSE ≤ 100 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 10) 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 (marked as the upper region in Figure 10). Only by adjusting a data rate knob at runtime would the node operate in an application-specific energy-fidelity optimized range (marked as the middle region in the shaded box of Figure 10).

Many BSNs also rely on event-detection where data is relatively unimportant until a particular "event" occurs. Once that event occurs however, it becomes critical to forward the medical information to an aggregator, and possibly off-body. In this situation, the ability to quickly detect the "event" and immediately change operating modes to one of high-fidelity throughput is essential. Runtime control of the various compression knobs available should be pushed to the periphery of the network (the BSN sensing nodes themselves) to facilitate the highest efficiency and quickest reaction time possible. Therefore, it is desirable to use embedded techniques capable of being executed on resource constrained microcontrollers typically used in BSN applications.

While it is convenient and important to use standard measures for signal fidelity (MSE, PRD, *etc.*) when investigating issues related to compression-fidelity tradeoffs, it is important to note that BSN pervasiveness relies instead on application-specific fidelity measures. Since the "information" present in BSN signals changes over the course of a single data-taking session, data reduction and compression techniques should be adjusted and managed at runtime to promote energy efficiency and enable new applications.

Figure 10. Dynamic distortion of movement data for a fixed compression ratio [29].

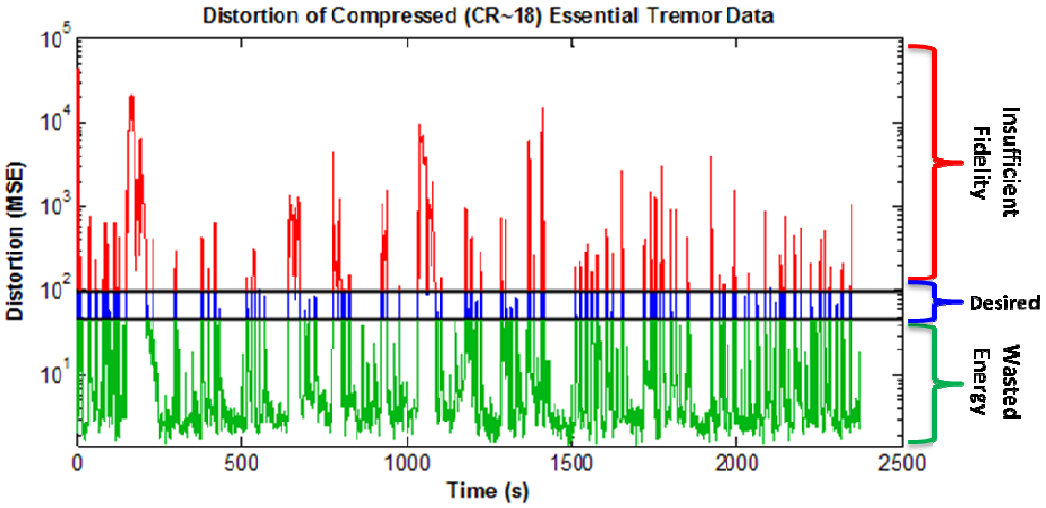

4. Case Studies

We will now examine the development of a COTS based BSN and a custom ASIC for BSNs to show how the important tradeoffs and energy saving strategies that we mentioned affect real BSN systems. First, we summarize the hardware selection approach as follows. Identify the region(s) of the application design space you want to cover, then map the system requirements into hardware constraints. Important factors that influence the decision include lifetime, form factor, cost, data rate (computation intense or communication intense?), and node purpose (research prototype or potential product?). The flexibility *versus* efficiency tradeoff will influence the size of the design space that one hardware platform can cover. Based on the communication intensity, lifetime requirement, and effective data rate of the node determine whether a COTS or custom radio is needed. This decision incorporates both the computation *versus* communication tradeoff and the fidelity *versus* energy tradeoff. Determine the appropriate hardware platforms for the on-node processing and on-node

control portions of the BSN node, based on the flexibility *versus* efficiency discussion. Decide whether a COTS component or custom built hardware is more suitable based on application, flexibility, cost, and purpose. To provide examples for this process, we present two designs for BSN nodes.

4.1. Case Study of COTS System: TEMPO

TEMPO 3.2 is the latest version of the TEMPO platform which has been designed for a range of BSN applications and illustrates the aforementioned tradeoffs between flexibility and efficiency in a COTS based system platform. Specifically, TEMPO was designed to meet requirements for human motion analysis: a broad category that can contain many specific applications. With this application area in mind, a low power non-invasive device is needed that is still flexible enough to address applications that vary from gait analysis to tremor assessment and activity detection.

TEMPO 3.2 uses MEMS accelerometers and gyroscopes to perform inertial sensing to measure and study human motion and wirelessly transmit this data to central aggregator such as a smart phone or PDA. Accelerometers and gyroscopes were chosen for inertial measurement because they are small in size, self contained, and inexpensive when compared with other technologies like optical motion capture or magnetic localization. In order to enable communication to PDAs and smart phones while still keeping power consumption as low as possible, Bluetooth was selected as the communication protocol. While Bluetooth and other standard protocols allow for interoperability, they are not the optimal choice for BSNs, and a networking protocol tailored to these types of systems may be necessary if systems are to be effective and energy efficient. However, standard protocols such as Bluetooth enable quick prototyping for initial data collection which can be beneficial for showing the value of new and emerging BSN technologies.

Also, since the devices need to be a small form factor and wearable, TEMPO 3.2 was created in the form of a large wristwatch. This enables it to be wearable and flexible from design perspective, but puts other significant limitations on the system. The size of the device not only puts restrictions on the size of the electronics, but also has a significant impact on the size and capacity of the battery that powers the device. Therefore, efficiency becomes a large concern for TEMPO 3.2 as it is required to have a runtime of several hours up to several days.

In order to meet the runtime constraints mentioned above, TEMPO 3.2 uses the ultra-low power TI MSP430 microcontroller that still provides the ability to program and load a wide range of functions and digital signal processing techniques. The microprocessor gives the ability to optimize power consumption of the system by compressing data or performing the digital signal processing techniques that let us transmit less data over the radio (the main consumer of power).

However, as is common with many systems, there is a desire to be able to adapt as new technology emerges and as application requirements change. TEMPO 3.2 remains flexible by including a daughter board connector that allows the addition of another sensor for addressing a wider range of applications. Likewise, radio technology is a constantly changing field as newer and lower power technologies are being developed. So TEMPO 3.2 includes the option of taking out the Bluetooth radio and replacing it with a different radio that communicates over the UART or SPI protocols. This leaves TEMPO 3.2 as a general platform that can address a wide range of applications, but does not address a specific application as power and size efficiently as custom hardware could.

In summary, this platform utilizes the advantages of COTS based systems such as hardware expandability and interoperability with other commercial devices. Hardware capabilities may be swapped out with pin-compatible ICs and daughter boards to facilitate application reuse along with changing technology standards and software functionality is easily modified and tested which can be beneficial for applications in which the processing requirements are still unknown. These advantages are typical of COTS platforms. However, careful consideration must be given to ensure form factor is kept relatively small even when modifications are done to fully make an attractive BSN node, which is exemplified in [30–32]. The Mica mote platform presented in [30] sits atop two AA batteries side by side resulting in a form-factor difficult to place on the human body. It contains an 802.15.4 radio with a maximum data rate of 250 kbps and a small 8-bit Atmel processor. No sensors exist on the main circuit board, but a 51-pin expansion connector allows for easy expandability at the expense of wearability. The Telos mote platform presented in [31] provides similar functionality to the Mica mote, comes with a commercial Texas Instrument microcontroller, 16-pin expansion, and optional light, temperature, and humidity sensors on the main board. However, the Telos platform still sits atop two AA batteries with a similar form factor which makes it undesirable for many BSN applications. The BSN node presented in [32] contains the same processor and radio as Telos, but focuses on form factor more extensively. Sensors must be added via a daughter board and the expansion connector with 6 analog channels and two serial ports which adds size. The main board however, is only 26 mm × 26 mm which promotes better wearability.

4.2. Case Study of Custom IC System

One example of a custom built BSN is the ECG chip presented in [7]. It is a 0.13-μm bulk CMOS sub-threshold (sub-V_T) mixed-signal system-on-chip (SoC) that acquires and processes an ECG signal for wireless ECG monitoring. The die photo is shown in Figure 11. The system consists of an adjustable gain instrumentation amplifier (IA), an 8-bit analog to digital converter (A/D), a microprocessor that operates in the sub-threshold region (sub-V_T), and a universal synchronous receiver/transmitter (UART) to communicate with an external radio. The SoC uses a sub-threshold digital microcontroller (μC) for adaptive control of the sub-V_T biased analog components and for processing the ECG data. The microcontroller core is a customized variant of the Microchip PIC 16C5X [33]. This base unit has 33 instructions and memory sizes of 24 to 73 bytes of RAM. A simple differential IA topology was chosen for the ECG amplifier. Because the amplitude of an ECG signal varies depending on the placement of the recording electrodes and physiological variations between individuals, the IA has a digitally adjustable gain. The 8-bit A/D digitizes the amplified ECG signal at a 1 kHz sampling rate. The A/D uses a dual-slope, integrating architecture. This architecture was chosen for its simplicity, low power consumption, and its insensitivity to device variation. By adjusting the A/D supply voltage, we can trade off power consumption with resolution. This allows for lower system power, since the A/D power can be reduced when the system does not require the full fidelity capabilities of the A/D.

Figure 11. Die photograph of the ECG SoC. The analog front end (instrumentation amp (IA) and A/D) and microcontroller (μC) comprise only 0.0633 mm^2 of active area [7].

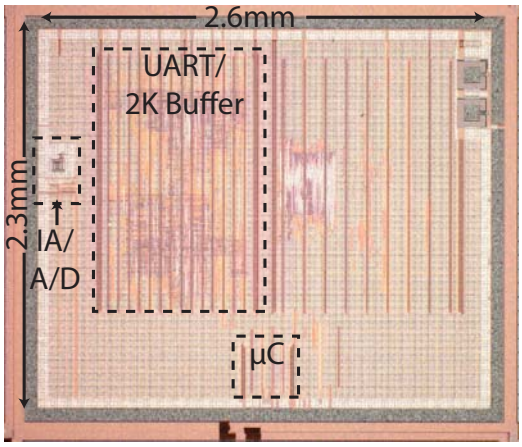

Level converters were used to transition data signals from low voltage domains to higher voltage domains. This is not a simple problem due to the large difference in voltage between a sub-threshold region supply and the nominal V_{DD} of 1.2 V. Most previous implementations of level converters supporting sub-threshold inputs use intermediate voltages to perform the up-conversion over multiple stages. A custom level converter for our SoC can convert from an input voltage of 160 mV up to 1.2 V.

The microcontroller, based on a PIC architecture, operates from 0.24 V to 1.2 V and consumes as little as 1.51 pJ per instruction at its minimum energy voltage of 0.28 V. The entire SoC (analog front end, ADC, and digital processor) consumes only 2.6 μW while providing raw ECG data or processed heart rate data. This level of energy efficiency far exceeds the abilities of COTS implementations and makes the idea of an energy harvesting ECG sensor feasible.

When only heart rate information is required, the onboard computation of heart rate reduces the wireless channel data rate by a ratio of 500:1, which allows complete beat-by-beat heart rate information to be communicated with much less energy expended in the radio. This chip exemplifies how low energy processing can be used to increase the effective CR in a BSN by extracting the important information from raw data prior to communication.

This custom SoC platform takes advantage of the fact that the application is well defined, and therefore platform flexibility can be traded off for optimization in energy efficiency and form factor. This platform utilizes the general strategies of communication *versus* computation and fidelity *versus* energy to obtain better energy efficiency. Other examples of custom platforms include [34–37]. [34] is a 0.5 × 1.5 × 2 mm^3 size, 5.3 nW intraocular pressure sensor with microprocessor and transmitter that is used to detect glaucoma. It achieves low power by duty cycling, the use of low power clocks, and on board processing. [35] integrates a glucose sensor with a wireless transmitter in a contact lens for diabetes monitoring with a power of 3 μW that is wirelessly transmitted. It utilizes a sub-μW low-power regulator and bandgap reference to achieve its low power profile. [36] is a fully integrated platform that processes heart rate detection and ECG for 445 nW and 895 nW, respectively. It lowers its power profile through the ability to utilize a low power, less precise clock or a higher power, more accurate clock as well as power-efficient biasing analog components. [37] is designed to be used in fabric to monitor vital signs, utilizing only 12 μW. [37] selects their topologies of their low-drop out

regulator, analog front end, and A/D carefully to remain under their power budget. As can be seen, design efforts for energy efficiency and very small form factor are two features common for custom platforms.

5. Conclusions

In this paper, we have explored strategies and methodologies for energy efficient design of BSN nodes. Starting from the characteristics of BSNs that arise from their application space and make them unique (including significant differences from traditional WSNs), we have identified the tradeoff metrics available for design optimization. We then elaborate on general strategies for designing energy efficient hardware, focusing on the tradeoffs of computation *versus* communication, flexibility *versus* efficiency, and data fidelity *versus* energy. We examine key tradeoffs in the BSN space that ultimately may lead to the decision between a COTS based platform or a custom IC design. Finally, we present two cases of previous work to show examples of a COTS based node and a custom designed hardware node. As the field of BSNs continues to grow, we anticipate that a rich selection of design techniques will lead to creative solutions leveraging both types of hardware design and resulting in numerous successful BSN deployments.

References

1. Yang, G.Z., Ed. *Body Sensor Networks*; Springer-Verlag: London, UK, 2006.
2. Barth, A.T.; Hanson, M.A.; Powell, H.C., Jr.; Lach, J. TEMPO 3.1: A Body Area Sensor Network Platform for Continuous Movement Assessment. In *Proceedings of the 2009 Sixth International Workshop on Wearable and Implantable Body Sensor Networks*, Berkeley, CA, USA, 3–5 June 2009; pp. 71–76.
3. Bequette, B.W. A critical assessment of algorithms and challenges in the development of a closed-loop artificial pancreas. *Diabetes Technol. Ther*. **2005**, *7*, 28–47.
4. Jurik, A.; Weaver, A. Remote medical monitoring. *IEEE Comput.***2008**, *41*, 96–99.
5. Hanson, M.A.; Powell, H.C., Jr.; Barth, A.T.; Ringgenberg, K.; Calhoun, B.H.; Aylor, J.H.; Lach, J. Body area sensor networks: Challenges and opportunities. *IEEE Comput.* **2009**, *42*, 58–65.
6. Wang, A.; Chandrakasan, A.; Kosonocky, S. Optimal Supply and Threshold Scaling for Sub-threshold CMOS Circuits. In *Proceedings of the IEEE Computer Society Annual Symposium on VLSI 2002*, Pittsburgh, PA, USA, 25–26 April 2002; pp. 7–11.
7. Jocke, S.; Bolus, J.; Wooters, S.N.; Jurik, A.D.; Weaver, A.C.; Blalock, T.N.; Calhoun, B.H. A 2.6-μW Sub-threshold Mixed-signal ECG SoC. In *Proceedings of the Symposium on VLSI Circuits 2009*, Kyoto, Japan, 16–18 June 2009; pp. 60–61.
8. Calhoun, B.H.; Chandrakasan, A.P. Ultra-dynamic voltage scaling (UDVS) using sub-threshold operation and local voltage dithering. *IEEE J. Solid-State Circuits* **2006**, *41*, 238–245.
9. Powell, H.C., Jr.; Barth, A.T.; Lach, J. Dynamic Voltage-frequency Scaling in Body Area Sensor Networks Using COTS Components. In *Proceedings of the 4th International Conference on Body Area Networks*, Los Angeles, CA, USA, 1–3 April 2009.

10. Si, W.W.; Husted, P.; Weber, D.; Abdollahi-Alibeik, S.; Lee, M.; Chang, R.; Dogan. H.; Gan, H.; Rajavi, Y.; Luschas, S.; Ozgur, S.; Zargari, M. A single-chip CMOS Bluetooth v2.1 radio SoC. *IEEE J. Solid-State Circuits* **2008**, *43*, 2896–2904.
11. van Zeijl, P.; Eikenbroek, J.W.; Vervoort, P.P.; Setty, S.; Tangenberg, J.; Shipton, G.; Kooistra, E.; Keekstra, I.; Belot, D. A Bluetooth radio in 0.18-μm CMOS. *IEEE J. Solid-State Circuits* **2002**, *37*, 1679–1687.
12. Kohno, R.; Hamaguchi, K.; Li, H.; Takizawa, K. R&D and Standardization of Body Area Network (BAN) for Medical Healthcare. In *Proceedings of the International Conference on Ultra-Wideband*, Hannover, Germany, 10–12 September 2008; pp. 5–8.
13. Drago, S.; Leenaerts, D.; Sebastiano, F.; Breems, L.; Makinwa, K.A.A.; Nauta, B. A 2.4 GHz 830 pJ/bit Duty-cycled Wake-up Receiver with −82 dBm Sensitivity for Crystal-less Wireless Sensor Nodes. In *Proceedings of the International Solid-State Circuits Conference*, San Francisco, CA, USA, 7–11 February 2010; pp. 224–225.
14. Ayers, J.; Mayaram, K.; Fiez, T.S. An ultralow-power receiver for wireless sensor networks. *IEEE J. Solid-State Circuits* **2010**, *45*, 1759–1769.
15. Pletcher, N.; Gambini, S.; Rabaey, J. A 65 μW, 1.9 GHz RF to Digital Baseband Wakeup Receiver for Wireless Sensor Nodes. In *Proceedings of the IEEE Custom Integrated Circuits Conference*, San Jose, CA, USA, 16–19 September 2007; pp. 539–542.
16. Daly, D.C.; Chandrakasan, A.P. An energy-efficient OOK transceiver for wireless sensor networks. *IEEE J. Solid-State Circuits* **2007**, *42*, 1003–1011.
17. Rabaey, J.M.; Abnous, A.; Ichikawa, Y.; Seno, K.; Wan, M. Heterogeneous Reconfigurable Systems. In *Proceedings of the IEEE Workshop on Signal Processing Systems*, *SiPS 97 Design and Implementation formerly VLSI Signal Processing*, Leicester, UK, 3–5 November 1997; pp. 24–34.
18. Zhang, H.; Prabhu, V.; George, V.; Wan, M.; Benes, M.; Abnous, A.; Rabaey, J.M. A 1-V heterogeneous reconfigurable DSP IC for wireless baseband digital signal processing. *IEEE J. Solid-State Circuits* **2000**, *35*, 1697–1704.
19. Zhai, B.; Nazhandali, L.; Olson, J.; Reeves, A.; Minuth, M.; Helfand, R.; Pant, S.; Blaauw, D.; Austin, T. A 2.60 pJ/inst Subthreshold Sensor Processor for Optimal Energy Efficiency. In *Proceedings of the Symposium on VLSI Technology and Circuits*, Hawaii, HI, USA, 13–17 June 2006; pp. 154–155.
20. Seok, M.; Hanson, S.; Lin, Y.S.; Foo, Z.; Kim, D.; Lee, Y.; Liu, N.; Sylvester, D.; Blaauw, D. The Phoenix Processor: A 30pW Platform for Sensor Applications. In *Proceedings of the Symposium on VLSI Technology and Circuits*; Hawaii, HI, USA, 14–16 June 2008; pp. 188–189.
21. Kwong, J.; Ramadass, Y.; Verma, N.; Chandrakasan, A. A 65 nm sub-Vt microcontroller with integrated SRAM and switched capacitor DC-DC converter. *IEEE J. Solid-State Circuits* **2009**, *44*, 115–126.
22. Calhoun, B.H.; Wang, A.; Chandrakasan, A. Modeling and sizing for minimum energy operation in sub-threshold circuits. *IEEE J. Solid-State Circuits* **2005**, *40*, 1778–1786.
23. Wang, A. Chandrakasan, A. A 180 mV FFT Processor Using Subthreshold Circuit Techniques. In *Proceedings of the IEEE International Solid-State Circuits Conference*, San Francisco, CA, USA, 6–10 February 2005; pp. 292–293.

24. Pu, Y.; de Gyvez, J.P.; Corporaal, H.; Ha, Y. An Ultra-low-energy/frame Multi-standard JPEG co-processor in 65nm CMOS with Sub/near-threshold Power Supply. In *Proceedings of the IEEE International Solid-State Circuits Conference*, San Francisco, CA, USA, 8–12 February 2009; pp. 146–147.
25. Kim, C.; Soeleman, H.; Roy, K. Ultra-low-power DLMS adaptive filter for hearing aid applications. *IEEE Trans. Very Large Scale Integr. Syst.* **2003**, *11*, 1058–1067.
26. Ryan, J.F.; Calhoun, B.H. A Sub-threshold FPGA with Low-swing Dual-VDD Interconnect in 90nm CMOS. In *Proceedings of the Custom Integrated Circuits Conference*, San Jose, CA, USA, 18–21 September 2010.
27. Texas Instruments MSP430F5172 Datasheet. Available online: http://www.ti.com (accessed on 24 January 2011).
28. Hanson, M.A.; Powell, H.C., Jr.; Barth, A.T.; Lach, J. Enabling Data-centric Energy-fidelity Scalability in Wireless Body Area Sensor Networks. In *Proceedings of the 4th International Conference on Body Area Networks*, Los Angeles, CA, USA, 1–3 April 2009.
29. Barth, A.T.; Hanson, M.A.; Powell, H.C., Jr.; Lach, J. Online Data and Execution Profiling for Dynamic Energy-fidelity Optimization in Body Sensor Networks. In *Proceedings of the International Conference on Body Sensor Networks*, Biopolis, Singapore, 7–9 June 2010; pp. 213–218.
30. Hill, J.L.; Culler, D.E. Mica: A wireless platform for deeply embedded networks. *IEEE Micro* **2002**, *22*, 12–24.
31. Polastre, J.; Szewczyk, R.; Culler, D. Telos: Enabling Ultra-low Power Wireless Research. In *Proceedings of the 4th International Conference on Information Processing in Sensor Networks*, Los Angeles, CA, USA, 25–27 April 2005; pp. 364–369.
32. Lo, B.; Thiemjarus, S.; King, R.; Yang, G.Z. Body Sensor Network—a Wireless Sensor Platform for Pervasive Healthcare Monitoring. In *Proceedings of the 3rd International Conference, PERVASIVE 2005*, Munich, Germany, 8–13 May 2005.
33. Microchip PIC16C5X Datasheet. Available online: http://www.microchip.com (accessed on 24 January 2011).
34. Chen, G.; Ghaed, H.; Haque, R.; Wieckowski, M.; Kim, Y.; Kim, G.; Fick, D.; Kim, D.; Seok, M.; Wise, K.; Blaauw, D.; Sylvester, D. A 1 Cubic Millimeter Energy-autonomous Wireless Intraocular Pressure Monitor. In *Proceedings of the International Solid-State Circuits Conference*, San Francisco, CA, USA, 20–24 February 2011; pp. 310–311.
35. Liao, Y.; Yao, H.; Parviz, B.; Otis, B. 3 μW Wirelessly Powered CMOS Glucose Sensor for an Active Contact Lens. In *Proceedings of the International Solid-State Circuits Conference*, San Francisco, CA, USA, 20–24 February 2011; pp. 38–39.
36. Zou, X.; Xu, X.; Yao, L.; Lian, Y. A 1-V 450-nW fully integrated programmable biomedical sensor interface chip. *IEEE J. Solid-State Circuits* **2009**, *44*, 1067–1077.
37. Yan, L.; Yoo, J.; Kim, B.; Yoo, H. A 0.5-μV_{rms} 12-μW wirelessly powered patch-type healthcare sensor for wearable body sensor network. *IEEE J. Solid-State Circuits* **2010**, *45*, 2356–2365.