

Sub-threshold Operation and Cross-Hierarchy Design for Ultra Low Power Wearable Sensors

Benton H. Calhoun¹, Jonathan Bolus¹, Sudhanshu Khanna¹, Andrew D. Jurik², Alfred C. Weaver², Travis N. Blalock¹

¹Electrical and Computer Engineering, ²Computer Science, University of Virginia, Charlottesville, VA, USA
<bcalhoun, jbolus, sk4fs, jurik, acw, tbladlock>@virginia.edu

Abstract— This paper examines the requirements of wearable sensing applications and their implications for designing the next generation of body area sensors. We define key metrics for wearable sensors and discuss how body area sensors differ from generic wireless sensors. To explore the system level issues with a wearable node, we show measurements from a wearable electrocardiogram (ECG) sensor prototype. Using heart rate monitoring as an example, we show how ultra low power (ULP) circuit design must be applied to support the stringent energy and/or power demands of long-life wearable sensors. Specifically, sub-threshold operation of digital circuits creates opportunities for re-thinking the entire system. We conclude that we can only reach the lower limits of power consumption through cross-hierarchy design of the entire sensor node that leverages ULP digital circuits.

I. INTRODUCTION

Healthcare applications provide enormous opportunity for technology to create positive impact. As healthcare costs and life expectancy continue to rise [1], the incentive to find novel ways for people to age comfortably and safely within their own homes increases as well. Although the technology is still nascent, smart wearable sensors will improve diagnosis and treatment of an array of conditions, potentially warding off life-threatening illness. Wearable body area sensor networks (BASNs) create new data streams and enable novel applications by performing real-time processing and communication on, in, or near the body [2]. Barriers of cost, limited device lifetime, and uncomfortable form factors have prevented deployment of BASNs. Extreme low power ICs would enable BASNs to minimize power, thus enabling long-life operation and smaller batteries; most current body area sensors use commercial off-the-shelf (COTS) components and have lifetimes of 12 hours or less.

Advances in ultra low power (ULP) circuit design have recently demonstrated capabilities compatible with BASN needs. To maximize the benefits of extreme ULP techniques, we argue that sub-threshold (sub- V_T) logic is essential for digital circuit elements, and cross-layer optimization is necessary to maximize energy efficiency across the entire system. To further reduce power consumption and complexity, we specifically propose that highly integrated system-on-chip (SoC) solutions must fully utilize the resources provided by ULP digital logic to augment the operation of analog components. This will require appropriate interfaces to higher level software to prevent constrained operation that hampers power

management. For example, some BASN applications must intentionally sacrifice power for higher performance or fidelity when they detect an emergency situation. Non-emergency cases might reduce performance or fidelity to extend system life.

Using heart rate monitoring as an example, this paper outlines system requirements for wearable sensors. Next, we review ULP sub- V_T operation. Finally, we question whether circuit techniques alone can achieve the lower limits of system level low power design. Using our discrete prototype of a wireless ECG sensor, we show how integrating our design effort vertically across traditional hierarchical partitions can improve power management and power consumption of the wearable sensor.

II. WEARABLE SENSORS

Table 1 lists several key differences between BASNs and generic wireless sensors. The scale of generic wireless sensor networks is larger than BASNs because the human body limits the surface area available for sensors. The networking of body sensors differs from generic sensors because BASNs use fewer nodes that tend to be more sophisticated (i.e., the “cheap and unreliable” sensor paradigm is usually inappropriate). Beyond typical requirements like reliability and low power, BASNs also must satisfy other conditions that are harder to quantify such as wearability, privacy, and security.

For body sensors, wearability and power are interconnected because battery size directly relates to the degree of wearability and lifetime of the sensor node. Since wearable sensors are intended to be worn on the body, miniaturization and minimal weight are important. Energy scavenging has received much attention as a way to sustain sensor nodes, and body area sensors might take advantage of this emerging technology (e.g., through heel strikes, temperature gradients, vibration, etc.). These technologies impose more stringent limits on the maximum average power the sensor can draw, indicating the need for extreme ULP SoC designs. While a battery allows for very large power delivery at the expense of reduced lifetime, the long-term average power produced by energy scavenging circuits represents a maximum bound on the long-term average power consumption of the sensor.

Security in both wireless sensor networks and body sensor networks is necessarily limited due to resource constraints. Since in wireless sensor networks the nodes tend to be physically accessible to a potential adversary,

Table 1 – Wearable body-area sensors have unique requirements when compared to their generic wireless sensor counterparts.

	Multipurpose Wireless Sensors	Wearable Body Sensors
Scale of Network	10s to 1000s of nodes over a large area using sophisticated multi-hop routing protocols	Less than 10 nodes around the body – generally one hop to a base station.
Lifetime	Rely on many nodes to bypass dead nodes	May be more conducive to periodic recharging
Form factor	Not a significant constraint	Physically small, light, and wearable
Security	Physical access to nodes defeats many security protocols	Since nodes carry health information, secure transmission is critical

merely securing transmissions may not be effective. On the other hand, in BASNs the sensors are in close proximity to the wearer, making physical access by an adversary difficult. Since body sensors can collect and carry personal health information, the security of transmissions deserves more consideration than in an environmental monitoring application for wireless sensor networks, for example. Yet in a telemedicine system, it must be acknowledged that in an emergency situation, the lifetime of a node as well as security and privacy may be intentionally sacrificed in favor of data volume and rapid access to life-saving information.

While generic wireless sensor nodes and body area sensor nodes share many common attributes, the unique BASN applications require varying constraints and different solutions. We now examine how circuit techniques for ULP can enable breakthrough technologies for sensor nodes, and body area sensor nodes in particular.

III. SUB- V_T CIRCUITS FOR ULP DIGITAL LOGIC

Sub-threshold circuit design uses a power supply voltage (V_{DD}) less than the transistor threshold voltage (V_T) to lower energy. Since sub- V_T operation increases delay, it is most applicable when performance is not the most critical metric, as in BASNs. Wireless sensor nodes and implantable medical devices have used this circuit design technique [3][4]. Lowering V_{DD} provides quadratic savings in dynamic power (and energy) and linear savings in leakage power. The V_{DD} to minimize energy per operation (V_{opt}) occurs in the sub- V_T region [5], which explains why sub- V_T circuits may enable long-life wearable sensors.

From the analysis of our wireless ECG sensor and from previous sensor application scenarios, we find that most sensor applications have bursts of activity interspersed among varying periods of inactivity (sleep time). The sleep mode may be interrupted by periodic processing, scheduled data collection, detected emergencies, or user-driven interrupts. Sleep mode energy consumption is an important factor in energy minimization of wearable sensors.

Knowing that wearable sensors exhibit varying amounts of sleep time, we first investigate the impact of sleep mode energy on total energy at a given technology node. We then

examine how the technology choice affects energy for a particular amount of anticipated sleep time.

There are two basic strategies for dealing with sleep time. The first is to set V_{DD} for sleeping circuits independent of active mode considerations. In that case, active mode circuits will run at V_{opt} to minimize energy/op, and sleep V_{DD} will be set as low as possible to lower leakage power while avoiding data loss [5]. Even during sleep mode, some circuits (e.g., timers) usually will be active to determine when to wakeup. Separate V_{DD} optimization for sleeping circuits requires a dedicated DC-DC converter. The overhead cost of the extra converter (which often requires large, costly off-chip passives) may prevent its adoption for a body area sensor. Also, the V_{DD} rail capacitance for a large chip may be large, creating a long (e.g., 100s of usecs) settling time to the sleep voltage. This timing overhead may be unacceptable in BASN applications with frequent shorter sleep periods.

Power gating (cutting off current from V_{DD} with an on-chip transistor in series with the supply) is an alternative sleep mode power reduction technique that has been widely used in industry, and that appeals to BASNs due to low overhead. In the following analysis, we assume that power gating can lower the standby leakage current by 10x, a number that is readily achievable by appropriate sizing of power gating transistors.

We investigated energy numbers for a 32-bit Kogge-Stone adder implemented in a set of predictive technology models (PTMs) that we optimized for low power applications (LP-PTMs), since published PTMs are for high performance devices. At the 45nm node, we first examine V_{opt} assuming sleep V_{DD} is optimized separately. The zero sleep curve in Figure 1 shows how total energy varies with V_{DD} at 45nm, and V_{opt} is roughly 250mV.

Now we consider the impact of energy during sleep. To do so, we assume that the total T_{on} plus T_{sleep} is fixed (i.e., dictated by a periodic measurement). Figure 1 shows how the sleep energy affects the total energy for T_{on} plus T_{sleep} values of 0.1ms and 1ms. At lower V_{DD} s, the total sleep time decreases for a fixed measurement cycle time since

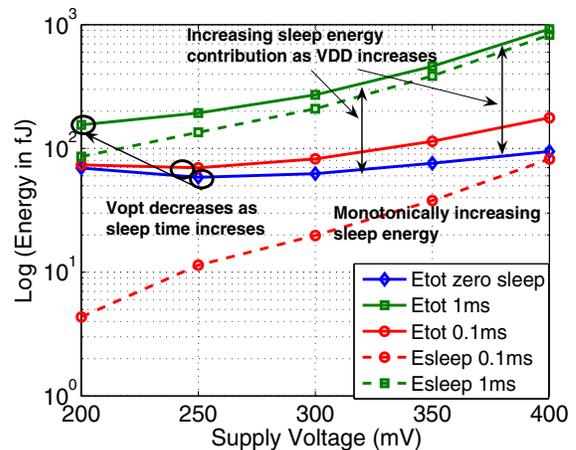


Figure 1 – Total energy (fJ) for a 32b adder at 45nm vs. V_{DD} for measurement cycles: $T_{sleep}=0$, $T_{on}+T_{sleep}=0.1ms$, and $T_{on}+T_{sleep}=1ms$.

Table 2 – Minimum energy (in fJ) points showing the trends with technology scaling for three different application scenarios. $V_{DD} = \max(V_{opt}, V_{min})$ for each cell in the table.

PTM (nm)	$T_{sleep}=0$	$T_{on} + T_{sleep} = 0.1ms$	$T_{on} + T_{sleep} = 1ms$
90	107	107	129
65	77.7	85.5	147
45	58.4	69.8	193
32	47.2	84.0	428
22	41.2	222	1860

T_{on} increases, and the leakage power decreases, so sleep mode energy decreases monotonically with V_{DD} . The plot shows that longer sleep times (e.g., 1ms measurement period) increase the minimum energy while decreasing V_{opt} , which now depends on the sleep energy as well. This data assumes 10x leakage reduction from power gating, but different power gating efficiencies will simply offset the sleep energy curves vertically on the plot.

When sleep time greatly exceeds on-time, sleep energy will dominate total energy so that V_{opt} becomes limited by functionality issues at low voltage (e.g., noise margin). For example, in the case of 1ms measurement cycles at 22nm, V_{opt} is 200mV but the minimum allowable V_{DD} (V_{min}) to maintain sufficient 3σ noise margin is 350mV. Due to increased variation, V_{min} increases at smaller process nodes.

We now consider the impact of sleep time on total energy consumption across technology nodes. Table 2 shows the minimum total energy for an adder (i.e., energy at V_{DD} equal to $\max(V_{opt}, V_{min})$) for three measurement cycle times.

There are two important observations from Table 2. First, working at the same node, higher sleep time increases total energy due to increasing leakage current. BASN applications with long sleep cycles would thus benefit from standby leakage reduction circuits. At 90nm, there is no energy increment going to 0.1ms cycle time because the 90nm active delay essentially fills the entire measurement cycle time. Because of decreased active delay, an application with a fixed measurement cycle time will see more sleep time at lower technology nodes.

The second critical observation is that increased sleep time causes the technology that minimizes total energy to increase [6] (highlighted cells in Table 2). Leakage energy dominates the sleep energy, and older technologies are much less leaky. For the 1ms measurement cycle time, 90nm saves over 10x energy relative to 22nm. By knowing application-driven on-time to sleep-time ratios, one can select an appropriate technology to lower total energy dramatically. The topology of the circuit will also affect this decision; the adder in our example is a fairly active circuit. Bigger circuits with a larger amount of relative leakage would benefit even more by using older processes to reduce sleep power consumption.

As BASN design matures, we anticipate increased levels of integration with more on-node signal processing. The next section explores how SoC designs can further reduce the total energy by leveraging ULP sub- V_T digital circuits like those we have described.

IV. A WEARABLE ELECTROCARDIOGRAM SENSOR

A. Wireless ECG sensor

To better understand BASN design trade-offs, we constructed a wireless ECG sensor prototype using COTS components. This prototype allows us to explore the wireless sensor design space and to observe limitations in commercially available components. Figure 2 shows a block diagram of the prototype. The patient attaches the sensor prototype to the chest via three disposable electrodes (to obtain a differential measurement of the cardiac signal and a reference). The differential measurement leads to high common-mode rejection, which reduces many forms of interference. We chose a two-stage amplifier topology with the first stage consisting of an AD623 instrumentation amplifier with adjustable DC offset, and the second stage consisting of a single-ended amplifier with adjustable gain.

A Texas Instruments MSP430 microcontroller with an integrated 12-bit analog-to-digital converter (ADC) digitizes the ECG signal and adjusts the DC offset and gain of the amp so the cardiac signal occupies the full dynamic range of the ADC. The digitized signal is then transmitted via Bluetooth to a PDA or other Bluetooth-enabled device. A 430-mAh polymer lithium-ion battery powers the prototype and radio.

While transmitting the complete ECG waveform, the prototype consumes 94 mW, 87 mW of which goes to the Bluetooth radio. While custom radios can reduce power relative to COTS parts, the radio will likely remain the highest power component of any wearable sensor. To keep this high power from dominating the energy budget, we must design the system to minimize the amount of data communicated wirelessly [7]. Since wearable sensors in healthcare applications may operate in life-threatening situations, the system must also allow dynamic adjustments to the tradeoff between power and signal fidelity to accommodate changing conditions such as emergencies. By

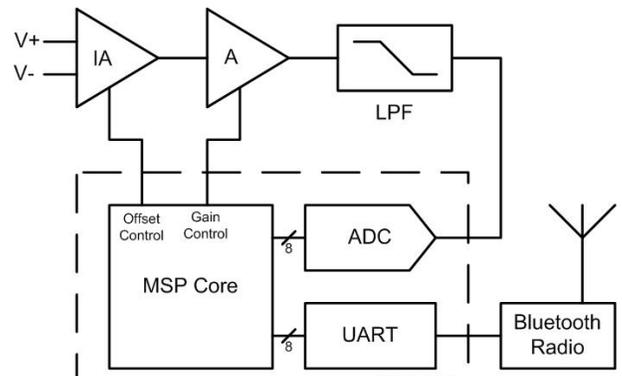


Figure 2 – Block diagram of our discrete prototype of a wearable ECG sensor.

implementing multiple operating modes, the sensor can spend most of its time in the low data rate modes, and switch to the higher data rate modes only when required.

Figure 3 shows our working prototype, which supports two operating modes. In full ECG mode, the entire digitized ECG waveform is transmitted at a data rate equal to the product of the sampling rate and the bit depth of the ADC. In our application, the sampling rate is 1 kHz and the ADC has 8-bit resolution, resulting in a data rate of 8 kbps. In the second operating mode, only information related to the patient's heart rate is transmitted. The computation of the inter-beat intervals reduces the required data rate to the product of the heart rate and the bit-depth of the interval measurement. For a person with an average heart rate of x beats per minute (bpm) and with 16 bits to represent the duration of each inter-beat interval, the bandwidth requirement reduces to x beats/minute \times 1 minute/60 seconds \times 16 bits/beat = $4x/15$ bits per second. For a heart rate of 70 bpm, sending inter-beat intervals achieves a 429:1 reduction in the wireless data rate. This corresponds to a 55% reduction in the wireless channel power of our prototype, from 87 mW to 48 mW, which is not as large as the reduction in the data rate due to overhead in the Bluetooth protocol. Radios for BASNs would ideally use low overhead protocols that allow transmit power to efficiently scale with data rate.

B. System level lessons from the ECG sensor

ULP digital logic allows us to implement increasingly complex forms of data analysis directly on the sensor for a negligible increase in sensor power. Possible applications include algorithms that detect cardiac arrhythmias, such as fibrillation events or bradycardia [8]. In this operating mode the sensor spends the majority of its time silent, only broadcasting an alert on the wireless channel when an abnormal condition is detected, and then transitioning to the full ECG operating mode. This technique can reduce the average wireless channel data rate by orders of magnitude, which dramatically lowers transmit power.

ULP digital logic can also be used to tune or augment the function of analog systems in the interest of energetic efficiency and reduced complexity. For example, the analog input's DC offset is eliminated by using programmable feedback from the MSP430 rather than having to resort to more complex forms of analog offset cancellation. Similarly, the amplifier gain is digitally controlled by the MSP430, eliminating the need for a complex analog gain control system.

The ideal sensor transmits only as much information as necessary at a particular moment in time, and performs the maximum amount of pre-processing and analysis prior to transmission. Counter to intuition, a "dumb" sensor actually increases system power. Making sensors more intelligent allows a significant reduction in overall system power with a simultaneous improvement in system lifetime. This requires taking a careful, holistic look at the



Figure 3 – ECG sensor sending waveform over a Bluetooth connection to a PDA base station.

system and application, and ensuring that interfaces are included to allow the upper levels of the system to adjust the lower-level operating modes of the system.

V. CONCLUSIONS

ULP techniques coupled with system-level optimizations and bursts of operation in between sleep periods work in concert to lower power consumption in body area sensors. We have shown how sub- V_T circuits lower energy by using V_{opt} for a given process. We analyzed the impact of sleep mode leakage from 90nm-22nm and concluded that choosing the best process node for a given application based on the amount of anticipated sleep time can lower total digital system energy by over 10x. Our ECG prototype shows how user-specific control of a BASN sensor can maintain low power operation in the common case and intelligently manage power in an alternate scenario, such as an emergency. Pre-processing techniques such as data buffering, burst transmissions, and compression further reduce total system power consumption by leveraging ULP digital circuits. Extreme ULP ICs for BASNs will use sub- V_T digital blocks and cross-hierarchy design for holistic power minimization.

REFERENCES

- [1] P. Ross, "Managing care through the air," *IEEE Spectrum*, vol. 41, No. 12, pp. 26-31, December 2004.
- [2] M. A. Hanson, et al., "Body Area Sensor Networks: Challenges and Opportunities," *IEEE Computer*, vol.42, no. 1, pp. 58-65, Jan. 2009.
- [3] A. Wang, A. P. Chandrakasan, and S. V. Kosonocky, "Optimal supply and threshold scaling for subthreshold CMOS circuits," in *Proc. IEEE Annu. Symp. VLSI*, 2002, pp. 5-9.
- [4] M. Seok, et al., "The Phoenix processor: a 30pW platform for sensor applications," *Symp. on VLSI Circuits*, pp. 188-189, 2008.
- [5] B.H. Calhoun, A. Wang, and A. Chandrakasan "Modeling and Sizing for Minimum Energy Operation in Sub-threshold Circuits," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 9, pp. 1778-1786, September 2005.
- [6] M. Hemstead, N. Tripathi, P. Mauro, G. Y. Wei, and D. Brooks, "An ultra-low power system architecture for sensor network applications," *ISCA*, pp. 208-219, 2005.
- [7] E. Shih, S. Cho, F.S. Lee, B.H. Calhoun, and A. Chandrakasan "Design Considerations for Energy-Efficient Radios in Wireless Microsensor Networks," *Journal of VLSI Signal Processing*, vol. 37, No. 1, pp. 77-94, November 2004.
- [8] G.D. Clifford, F. Azuaje, and P.E. McSharry, *Advanced Methods and Tools for ECG Data Analysis*. Boston: Artech House, 2006, pp. 55-99.