

Design of a Real-Time Adaptive Power Optimal Sensor System

Ari Y. Benbasat and Joseph A. Paradiso

Responsive Environments Group, MIT Media Lab

Cambridge, MA 02139, USA

Email: {ayb,joep}@media.mit.edu

Abstract

Wireless sensors systems are currently being deployed in a wide variety of lightweight mobile applications such as distributed object tracking and wearable medical data collection. For such applications to enter the consumer mainstream, it is necessary for them to operate far more power efficiently than they do currently. Using a simple illustrative example, this paper explores some general design rules for techniques which can reduce power consumption by one to two orders of magnitude through a combination of real-time sensor selection based on system state and in-situ feature extraction before storage and/or transmission (if necessary). In the presented example, it is demonstrated that the use of both low and high accuracy sensors for a single parameter (acceleration) can allow system state identification to be far more power-efficient (in this case, by 94%) than with a single sensor alone. Use of sensors with short wakeup times will further improve this result.

Keywords

Adaptive sensing, power-optimal sensing, modular sensors, wireless sensors, wearable sensing, gait analysis.

INTRODUCTION

Wireless sensor nodes and networks are currently being used in a wide array of applications. These include, but are certainly not limited to, distributed sensing and processing arrays[11], mobile data collection nodes[6] and wearable medical analysis systems[1]. In most cases, such systems leverage the coincident decrease in both price and size of the necessary components, particularly MEMS based sensors.

However, such systems are often limited to prototype and experimental usage due to limited battery (and therefore system) life. The most common solutions to this problem are to tether the system to an unlimited power source[5] or limit the system to sampling at such a low rate that the lifespan is satisfactory[9]. Obviously, the full potential of wireless and wearable sensors is not being achieved through such systems, as their limited lifespans, sensing capabilities or update rates greatly reduce the utility to the end user, and therefore their marketability. Long-term medical monitoring, for example, is currently hindered by stationary sensors, which cannot provide a full picture of an active patient's movements, or body-worn sensors, which require large battery packs and/or frequent replacement.

While the progress of semiconductor technology will eventually bring lower power systems, battery life problems can be attacked today via a more careful approach to sensor system design. The outlines of such an approach are detailed in this paper.

GOALS AND MOTIVATION

The main goal of this work is to delineate a collection of design rules and techniques for the creation of power-optimal sensor systems. Such systems can be defined as those which extract the necessary or desired information for the smallest outlay of power. While the processor and RF transceiver are generally considered the largest power drains in wearable systems, sensor power has become equivalent if not greater in many recent lightweight systems[12].

We wish to create general approaches for two different methods of power reduction. The first is the use of multiple sensors to measure a single parameter of interest. The vast majority of sensor systems limit themselves (usually in the interests of simplicity or compactness) to a single sensor for each modality of interest. No matter how efficient such an implementation is for its intended goal, it is guaranteed to be power inefficient in states were less (or more) data is necessary to collect all the available information. A system which can tailor its sensing in real-time to the current state of the device can draw far less power on average.

The second approach is to use a number of different techniques for *in-situ* compression of data. Most sensor systems either transmit or cache all data without regard to expository value. Simple data processing and feature extraction can reduce this volume considerably for a small power investment. A further benefit is that the results of the compression operations could lead to a decision not to store and/or transmit the data at all.

The motivation for this work is two-fold. In the short term, a decrease in power consumption increases the system lifespan, which directly correlates to utility and ease of adoption. Systems which are currently confined to specialized laboratories could be used in unconstrained environments, opening a large number of potential applications. In an example pursued by our group, wearable gait analysis systems could provide real-time feedback to the user, allowing for corrections which could help avoid a number of different injuries[10]. In the long-term, these techniques will continue to be of benefit in battery-powered systems. While key thresholds in user utility (*e.g.* a full

day's uninterrupted use for a medical wearable) may have been exceeded, these techniques will still provide the same percentage increase in system life and will therefore continue to allow for a greater application range than non-adaptive systems.

RELATED WORKS

Shrivastava has written a number of papers (e.g.[13]) examining the power usage of mobile nodes under a variety of conditions (data collection, wireless transmission, sleep, etc) and comments on how to select modes for processors and RF transceivers which are most efficient for their particular task. However, his work expressly ignores the issue of sensor power.

Estrin's[3] survey paper on pervasive networks mentions many of the techniques, both hardware and software, explored in our work, though only at the theoretical level. Specifically of note are comments on frequency of appearance of interesting data and the benefits of multiple modes of operation.

Also, many sensor networks[8] use rotating cluster heads – an example of dynamical power optimality. However, such systems merely attempt to maximize network life by distributing power draw among the nodes rather than through increasing the life of each node.

Similarly, systems employing passive wakeup[7] partially achieve the goal of state-based sensing, though they take a very binary approach to context. Our system is able to consider a much finer concept of state, and therefore can both respond more quickly and more appropriately.

WORK TO DATE

System Design

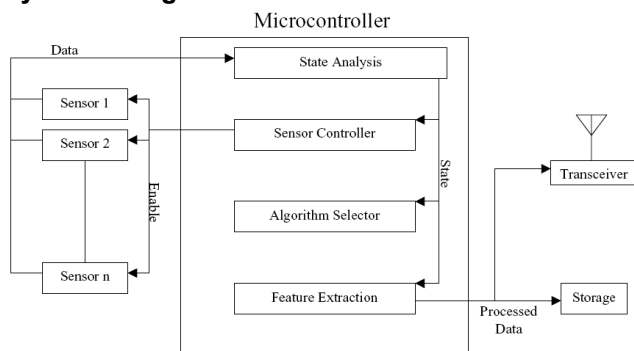


Figure 1: Flowchart of Data Flow in Real-Time Power Optimal System

The general design for our system is shown in Figure 1. This system is centered around the concept of what we term “groggy”, or semi-active, wakeup. Rather than have a sensor systems which is either fully active, collecting all possible data, or fully asleep, we envision a system with a number of different levels of activity. Each of these is designated by a state, which comprises the currently active sensors, their sampling rate and accuracy, and algorithms to describe both the state transitions and any processing of the

data (if desired). This data is then either stored locally or wirelessly transmitted to a basestation.

From a starting state, the system will analyze data and switch states accordingly. At any given time, it will attempt to collect as little data as possible to extract the available information (as defined by the application) from the environment. Further, the processor and sensors are put into a sleep mode between samples if possible. By definition, these two characteristics should provide power optimal sensing¹. To allow for further power savings, compression algorithms can be applied to the data prior to storage or transmission. These algorithms will exploit the structure in the data with respect to the application to condense it to relevant features. Beyond this point, further data analysis is left to the application designer or an expert in the field.

A concern with any wakeup based system is the potential to miss fast transient events - either entirely or during the wakeup/state analysis procedure. While this is a problem in general, it shouldn't be in this case as we will concentrate on wearable sensor systems. Since most human activities take place on the order of seconds and state analysis requires on the order of tens of milliseconds, the chance of missing an event is slim. Also, a concern with any state based system is that important information will be not be gathered if it doesn't fall within the data collection state(s). While can be avoided by making the states as broad as possible (without being inefficient), a better approach is to consider unsupervised state creation algorithms[3] in the design process. While more complicated, it should guarantee that the data collection states include all information of interest.

Hardware




Figure 2: Assembled modular sensor stack

We have developed a compact wireless modular sensor architecture (Fig. 2). The system itself is comprised of boards (panes) 1.4 inches square and 0.4 inches high which are electrically and mechanically connected at the corners. Each pane instantiates a major function, with most

¹ Predicated on the assumption that the data is being collected in the most efficient way possible. We have separately tabulated[unpublished technical report] the best sensor choices given the parameter to measure and the number of bits desired.

Table 1: Summary of states and power usage (μC : MSP430F147, Gyro: ADXRS300, Accel: ADXL202)

State	% of Time	Sensors	Rate	Sensor Power	Processor Power	Total Power	Weighted Contribution
Still	90	Tilt	15 Hz	Passive	1.1 μW	1.1 μW	0.11 μW
Non-Ambulatory	8	Tilt, Gyro	15 Hz	15mW	2.3 μW	15mW	1.2mW
Walking	2	Tilt, Gyro, Accels (x2)	100 Hz	32mW	2.3 μW	32mW	0.64mW

Total Average Power 1.84mW

parameters being measured in multiple fashions. Key panes to this discussion are:

- Inertial measurement unit board with three accelerometers (2 x Analog Devices ADXL202), three gyroscopes (1 x Analog Devices ADXRS300, 2 x Murata ENC-03J) and a four-direction passive tilt switch (ALPS SPSF100100).
- Processor board with a Silicon Labs C8051F206 22 MIPS processor and a RFM TR1000 115.2 kbps transceiver. This processor will be replaced by a TI MSP430F147, which has far lower power sleep modes and upon which the calculations below are based.

Other panes include a tactile board with inputs for bend and pressure sensors; an ambient board with a cell phone camera, PIR motion detector and IR phototransistors; and a sonar board for two dimensional acoustic ranging. The system has been used to implement a number of applications, including a compact on-shoe gait analysis system[10].

This architecture is easily extensible with new panes, allowing for rapid prototyping of systems with numerous sensors for each modality, which will be used for testing of our designs and assumptions. As a preliminary test, we have repeated a small amount of data collection from the gait analysis system using the two boards detailed above. This data will be analyzed to determine the benefits of our proposed design.

RESULTS AND DISCUSSIONS

For the case of a shoe-mounted inertial measurement system, we have examined the benefits of dividing the operation of the device into three states: still, non-ambulatory motion (twitching, swinging foot, jumping) and walking. In case each, we considered the data analysis necessary to determine transitions to the next state in the hierarchy. The still state is trivial, relying solely on whether the passive tilt sensors show any deviation from their previous state. This data is currently sampled at 15Hz, though transitions could also be used to trigger interrupts to wake the processor from a sleep state. Once motion is seen, it needs to be determined whether there is structure therein. This is achieved via analysis of the gyroscope data, again at 15Hz. The data is analyzed first for motion beyond a fixed threshold (based on its variance). If this threshold is exceeded, the data is then examined for the regular three-peaked structure created by the foot lifting off the ground, swinging through its arc and returning to the ground. This

data is separated both through peak counting and by confirming that the area under that segment of the graph is close to zero (since the foot starts and stop in the same orientation). More details on this technique can be found in [2]. If walking motion is found, the system is switched into a third state, where is collected data from each of the sensors at 100Hz. Note that the state determination analysis is still the same as before and therefore requires the same amount of power. These same tests are used in the reverse order to transition back down the states.

Figure 3 shows a parsed data stream. The stream includes a number of different common motions such as walking and swinging the foot. We note that the system accurately finds the walking states with no false identifications². Despite the fact that these preliminary results are based on a short scripted interaction, we are confident that the analysis techniques should generalize to longer, freeform streams based on positive experiences with these algorithms in previous work[2].

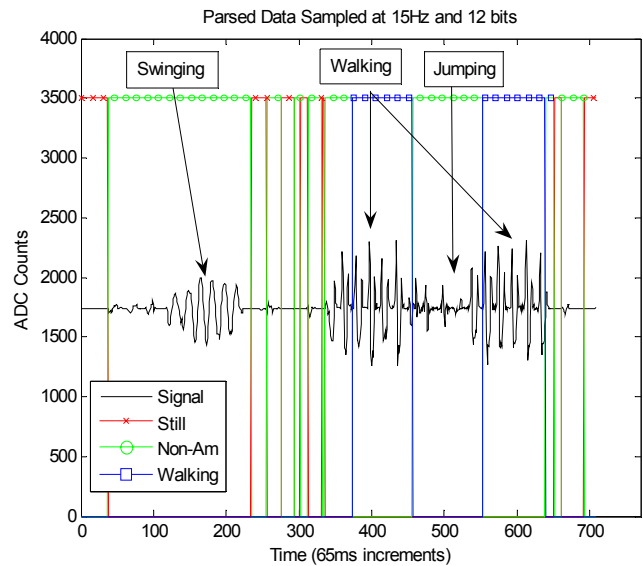


Figure 3: Parsed Data Stream

Table 1 summarizes the system states. The percentage of time that the user is assumed to be in each of the states is an estimate based on personal experience and assumes the system is never turned off (even when not worn). We note a few points of interest. First, the system achieves a savings

² The recognition appears to miss the first step of each walking state because one cycle is necessary to confirm the transition.

Table 2: Comparison of Power Cost of Common Data Operations (64 byte block, μ C: MSP430F147)

Operation	Example Use	Output Size	Energy Consumption
Wireless transmission	Raw data transfer	N/A	154 μ J
FFT	Find spectrum for state calculation	64 bytes	1.15 μ J
Windowed Variance	Feature extraction	1 byte	291 nJ
N-tap Feedback Network	Generalized filtering	64 bytes	190 nJ per tap

of 94% relative to operating in the walking state at all times. The system also realizes a 43% savings compared to a binary state system. Both of these savings are constrained by the fairly high power usage of the non-ambulatory state, which is due to the slow wake-up time of the ADXRS300 gyro. Other parts, such as the ENC-03J, may provide better results, but their documentation was not complete enough to allow for calculations. Also, we note that the processor power used is three orders of magnitude less than that of the sensors. Processing equivalent to approximately 10,000 cycles per data point would be necessary to place the two power draws at the same order of magnitude.

Generally, there are two main lessons here. The first is the key importance of wake-up time in any multi-state system. The slower the wake-up time, the smaller the benefit of power-cycling the device in the case of low update rates. The second is that common MEMS sensors draw far more power than the processing necessary to determine if their data is of value, which bodes well for the future success of this work.

FUTURE WORK

We consider two directions for future work. The first is to extend the analysis of the previous section. While strong results were achieved, the state assignment and data analysis were very *ad-hoc*. Using those results as a starting point, we wish to establish a set of general design rules and techniques, which can be used to create similar results for other types of data. Most straight-forward is generalizing the data stream analysis techniques such that they can be used with non-inertial data. More useful is to use the techniques themselves as a guide to create an unsupervised learning algorithm, which should be able to find similar states to those chosen by the authors. At very least, it should be able to find a number of interesting states which an application designer can then choose among.

The second direction for future work is to consider algorithms for *in-situ* compression of data. This area again has the potential for large power savings, as the results should reduce transmission/storage or lead to a decision to forego it entirely. However, the compression algorithms first need to be defined and then tested for utility (bytes saved versus data lost). Table 2 gives examples of the power cost of a number of different operations on a 64 byte window, compared to the cost of transmitting that data outright.

CONCLUSIONS

We have created a general design for adaptive power-optimal sensor systems. It is based upon a system with a

number of different states, each using different sensors and analysis algorithms, and the use of multiple sensors for each parameter of interest. A preliminary test of a wearable gait analysis system, using our modular sensor platform and an *ad-hoc* assignment of states and algorithms, achieved very strong results and provided important directions for future improvement and generalization of our system design.

ACKNOWLEDGMENTS

The authors would like to acknowledge the support of the Things that Think Consortium, as well as all the sponsors of the MIT Media Lab. Mr. Benbasat acknowledges the support of the MIT Center for Bits and Atoms and Media Lab Europe.

REFERENCES

- [1] A. Benbasat, S. Morris and J. Paradiso. "A Wireless Modular Sensor Architecture and its Application in On-Shoe Gait Analysis." IEEE Sensors 2003. pp. 1086-1091.
- [2] A. Benbasat. An Inertial Measurement Unit for User Interfaces. SM Thesis, MIT. September 2000.
- [3] R. Duda, P. Hart and D. Stork. Pattern Recognition. 2nd Ed. Wiley: New York. 2001.
- [4] D. Estrin et al. "Connecting the Physical World with Pervasive Networks." IEEE Pervasive Computing. Jan 2002. pp. 59-69.
- [5] S.S. Intille. "Designing a Home of the Future." IEEE Pervasive Computing. April 2002. pp. 80-86.
- [6] M. Laibowitz and J. Paradiso. "Parasitic Mobility in Dynamically Distributed Sensor Networks." MobiSys 2004.
- [7] H. Ma and J. Paradiso. "The FindIT Flashlight: Responsive Tagging Based on Optically Triggered Microprocessor Wakeup" UbiComp 2002. pp. 160-167.
- [8] A. Mainwaring et al. "Wireless Sensor Networks for Habitat Monitoring." WSNA 2002. pp. 88-97.
- [9] Millennial Net. <http://www.millennial.net/>
- [10] S. J. Morris. A Wireless Wearable System for Gait Analysis. ScD thesis, Program in Health, Science and Technology, MIT. June 2004.
- [11] J. Paradiso, J. Lifton and M. Broxton. "Sensate Media -- multimodal electronic skins as dense sensor networks." BT Technology Journal. October 2004.
- [12] M. Srivastava. Sensor Node Platforms & Energy Issues. MobiComp 2002 Tutorial Notes.
- [13] V. Raghunathan et al. "Energy-aware Wireless Microsensor Networks." IEEE Signal Processing Magazine. March 2002. pp.40-50.