31st Annual International Conference of the IEEE EMBS
Minneapolis, Minnesota, USA, September 2-6, 2009

Episodic Sampling: Towards Energy-efficient Patient Monitoring with Wearable Sensors

Lawrence K. Au, Maxim A. Batalin, Thanos Stathopoulos, Alex A.T. Bui, and William J. Kaiser

***Abstract*— Energy efficiency presents a critical design challenge in wireless, wearable sensor technology, mainly because of the associated diagnostic objectives required in each monitoring application. In order to maximize the operating lifetime during real-life monitoring and maintain sufficient classification accuracy, the wearable sensors require hardware support that allows dynamic power control on the sensors and wireless interfaces as well as monitoring algorithms to control these components intelligently. This paper introduces a context-aware sensing technique known as episodic sampling – a method of performing context classification only at specific time instances. Based on Additive-Increase/Multiplicative-Decrease (AIMD), episodic sampling demonstrates an energy reduction of 85 percent with a loss of only 5 percent in classification accuracy in our experiment.**

I. Introduction

Wireless, wearable sensor technology has recently received significant attention in healthcare because it offers continuous monitoring capability that traditional methodology lacks [1]. Coupled with the prevalence of mobile devices (e.g., cell phones), wearable sensors have demonstrated their utility in healthcare applications such as physical rehabilitation ([2]), chronic disease management ([3], [4]), and preventive care ([5], [6]). Many practical issues surround the use of wearable sensors, and they should be examined and resolved for any real-world deployments.

From a clinical perspective, the wearable sensors must be able to recognize a patient's context with sufficient classification rate. Consider a patient with chronic obstructive pulmonary disease (COPD). The patient is equipped with the wearable sensors as shown in Figure 1. Diagnostic information that may be significant or valuable to physicians include:

1) Event Triggers: Alerts relevant individuals whenever the breathing rate of the patient exceeds a threshold.

2) Context Snapshots: Selectively collects physiological signals of the patient during specific activity context [7].

3) Evaluation of Quality of Life: Recognizing various activities of daily living (ADL) helps evaluate a COPD patient's well-being [8].

These operations require algorithms that accurately classify contexts as specified by medical experts. Wearable sensors must be able to perform these operations reliably in order to be clinically useful over an extended period of time.

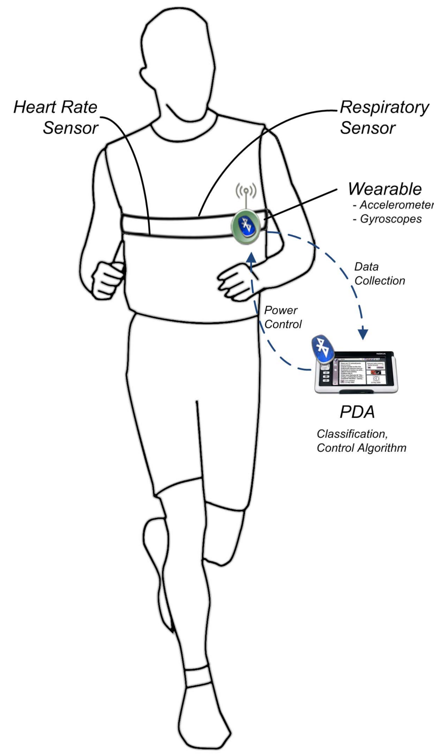

Fig. 1: Wearable Device and Sensors

From a technical perspective, the combination of sensor data acquisition and signal processing algorithms must provide adequate diagnostic and classification information to the patients or medical experts when needed. The wearable devices must also be lightweight enough to be worn without much inconvenience [9].

Nevertheless, the most stringent design parameter in wearable sensors is the operating lifetime of these wearables [10], [11]. Battery power is often limited in wearable sensors, yet constant context recognition and wireless communications demand significant sensing and processing power on both wearable sensors and mobile devices. For example, in patient monitoring, *raw* sensor data are often transmitted continuously to a central location for real-time processing and data fusion; such practice is not necessarily the most *energy-efficient* [12]. Many existing wearable sensors applications in healthcare do not address this problem directly because the

This work is supported by National Institutes of Health/National Library of Medicine under grant NLM T15 LM07356.

L.K. Au, M.A. Batalin, T. Stathopoulos and W.J. Kaiser are with the Department of Electrical Engineering, University of California, Los Angeles, CA 90095 USA (email: au@ee.ucla.edu; maxim.batalin@ucla.edu; thanos@cs.ucla.edu; kaiser@ee.ucla.edu)

A.A.T. Bui is with the Department of Radiological Sciences, UCLA, Los Angeles, CA (email:buia@mii.ucla.edu)

978-1-4244-3296-7/09/$25.00 ©2009 IEEE

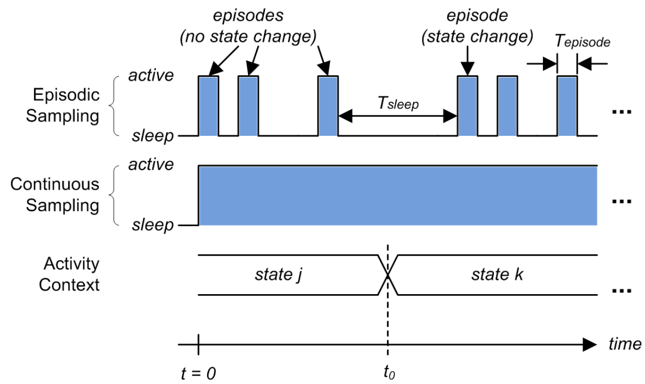

Fig. 2: Episodic Sampling vs. Continuous Sampling

length of their experiments and pilot studies are relatively short and the depicted scenarios are relatively controlled (e.g., a laboratory setting without significant wireless interruptions). Such scenarios may not be applicable to real-world deployments: Battery power can run out within hours in wearable sensors, and recharging is often unavailable while the devices are in use.

As a result, wearable sensors needs to become more energy-efficient. Specifically, they must have the ability to power on or shut off various sensors and wireless interface *on-demand*, as dictated by the context and diagnostic objectives. Furthermore, monitoring algorithms should possess knowledge about the application in question and control these wearable sensors intelligently [11]. Such hardware-software system approach has been shown to be effective in other applications [13]. In summary, wearable sensors must perform operations as needed by the monitoring objectives through careful energy management.

This paper proposes a control algorithm that reduces energy consumption (and thus improves overall energy efficiency) by performing context classification only at specific time instances, or *episodes*. At all other times, the wearable sensors are placed in a low-power or sleep mode. To monitor the activity pattern, the second generation of MicroLEAP (ULEAP2) wearable sensor system is used in the subsequent experiment [14], [15]. It is built with an energy management unit that provides an interface to monitor real-time energy consumption of the wearable and to power-cycle sensors and wireless interface. This unique capability allows monitoring algorithms to dynamically optimize their behavior based on the context information. The proposed control algorithm, which is based on Additive-Increase/Multiplicative-Decrease (AIMD), determines the time instances (episodes) for context classification.

The subsequent sections will elaborate on the idea of episodic sampling, the rationale behind it, and the required system support (Section II), provide experimental results based on the proposed algorithm (Section III), and discuss some of the implications of the results (Section V).

II. Methodology

Episodic sampling is a context-aware scheduling technique where context classification occurs only episodically, and the time for the next episode depends on prior context information. This takes advantage of the observation that human activities typically vary slowly and constant state classification is not necessary in most situations [16]. Figure 2 compares episodic sampling with the traditional continuous sampling. In continuous sampling, the system performs constant classifications in the active state. In episodic sampling, the system only executes classification at specific episodes and remains in a sleep state between episodes. $T_{episode}$ represents the duration of an episode, and it mainly depends on the amount of sensor data required to produce classifier outputs. In each episode, a feature vector, $\vec{F}$, is extracted from the sensor data. The classifier, $C(\vec{F})$, infers the present context. T_{sleep} represents the sleep time between consecutive episodes; this is the time interval when the wearable sensor spends in low-power mode.

A. System Support

The wearable used in this paper is MicroLEAP2 (ULEAP2), the second generation of the energy-aware wearable sensor system [14]. It is equipped with a Texas Instruments's MSP430 embedded processor, a Bluetooth interface, and an external ADC that accommodates up to eight analog sensor inputs. ULEAP2 provides a mini-USB port for programming and battery recharging.

The energy management unit in ULEAP2 provides a software-based interface for real-time energy monitoring and power control on the wireless interface and sensors with the use of current-sensing circuits and power switches. This unique feature in ULEAP2 allows local energy optimization based on both context and energy consumption.

Additionally, ULEAP2 incorporates various sensors for physiological monitoring:

1) Motion Sensors: The wearable is integrated with motion sensors that can capture tri-axial accelerations (accelerometer: ADXL330 [17]) and angular rotations (gyroscope: IDG-300 [18]). These sensors provide measurements of motions, capable of capturing various body movements related to activities of daily living.

2) Respiratory Sensor: The respiratory sensor is ADInstrument's MLT1132 [19], a piezoelectric belt that generates a small voltage in response to mechanical stretching. The sensor is completely passive; thus it is safe to be worn non-invasively. ULEAP2 amplifies and conditions the voltage generated by the respiratory sensor. Figure 3 shows the sample waveforms collected after filtering.

3) Heart Rate Sensor: For heart rate monitoring, ULEAP2 includes the Polar chest strap that is commercially available [20]. ULEAP2 collects the heart beats from the chest strap through a separate receiver and locally computes the corresponding heart rate.

ULEAP2 runs on MicroC/OS-II, a preemptive real-time operating system [21]. It provides a simple mechanism to multi-task sampling, processing, and data transmission.

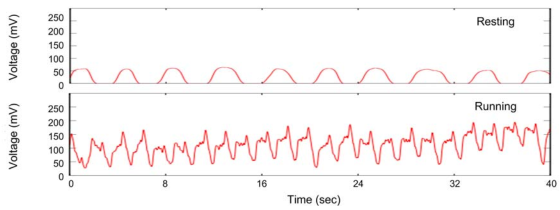

Fig. 3: Raw Waveforms from Respiratory Sensor

Sensor data are collected by the mobile device (a Nokia N770 tablet PDA), which performs feature extraction and classification in a Linux-based, multi-threaded environment.

B. Feature Extraction and Classification

Selecting good feature extraction algorithms is critical to accurate context classification. For the activities in the experiment, the N-point Fast Fourier Transform (FFT) is an effective algorithm because all the activities of interest exhibit certain periodicity. Specifically, the peak magnitude, F_{peak}, and the energy, F_{energy} of the N-point FFT, $X(k)$, from the accelerometer signals form the feature vector, $\vec{F}$, as well as the input of the classifier, $C(\vec{F})$:

$$F_{peak} = max_k \|X(k)\| \tag{1}$$

$$F_{energy} = \sum_{k=1}^{N/2} X(k) * X(k) \tag{2}$$

Both of these features are highly correlated to body movements [22], and a sampling rate of 128 Hz captures major human motions with sufficient accuracy. For each N-point FFT operation, a sliding window of size 256 data points (N = 256) is used to ensure sufficient frequency resolution. A new $\vec{F}$ can be derived every second; each operation takes the newest 128 samples and discards the oldest 128 samples.

C. Control Algorithm

Intuitively, classification should be performed when a significant variation in activity pattern appears. On the other hand, if the activity pattern does not vary frequently, classification can be minimized. By adaptively adjusting the frequency of context classification, one can reduce the amount of time wearable sensors spend in active power state performing constant context classification. The apparent tradeoff is the classification accuracy; as a result, it can be viewed as a control problem that adapts to the dynamics of the most recent activity pattern.

The proposed control algorithm is based on Additive-Increase/Multiplicative-Decrease (AIMD). AIMD has been used most notably in TCP congestion avoidance [23]. In TCP congestion avoidance, AIMD adapts its window size to accommodate the available bandwidth. When the window size is less than the available bandwidth, the algorithm gradually increases its window size. When it detects packet drops or receives no acknowledgement, the window size is halved.

Algorithm 1 Episodic Sampling

$T_{sleep} = 0$
while in each episode **do**
  Collect Sensor Data
  Extract Features, $\vec{F} = (F_1, F_2, ...)$
  $state \Leftarrow C(\vec{F})$
  if $state$ has changed since last classification **then**
    $T_{sleep} \Leftarrow a \cdot T_{sleep}$, where $0 \leq a \leq 1$
  else
    $T_{sleep} \Leftarrow T_{sleep} + T_{incr}$
    if $T_{sleep} > T_{sleep,max}$ **then**
      $T_{sleep} \Leftarrow T_{sleep,max}$
    end if
  end if
end while

This paper adopts the concept and applies it to episodic sampling; similar control strategy has been used before [16]. Consider Figure 2: when the state (context) remains the same, the algorithm gradually increases T_{sleep}. When a state change is detected, the algorithm aggressively decreases T_{sleep} to avoid future misses. Algorithm 1 illustrates the pseudocode. During each episode, sensor data are collected for classification. If the pattern exhibits constant state changes (e.g. during interval training), the algorithm decreases T_{sleep} by a constant factor, a. In the limiting case (i.e., when T_{sleep} equals zero), the algorithm reduces to continuous sampling. When the state remains identical for an extended period, the algorithm increases T_{sleep} by an amount T_{incr}. To ensure T_{sleep} does not become unbounded, the algorithm caps T_{sleep} at some maximum value, $T_{sleep,max}$. Such a heuristic algorithm has a constant run-time (i.e., $O(1)$). The effectiveness of the algorithm mainly depends on the accuracy of the classifier and the parameters described.

1) Selection of T_{incr}: Depending on the prior knowledge of the activity pattern, T_{incr} can be a constant or a random variable. In this experiment, T_{incr} is a uniformly distributed random variable between 0 and 30 seconds. Such an approach has been used in similar experiments [16].

2) Selection of a: The multiplicative factor, a, is a constant that describes how 'sensitive' the algorithm should respond to a change in context. When a is small (large), the algorithm computes the classifier more (less) frequently and the energy consumption increases (decreases). In this experiment, a equals 0.5.

In this paper, raw sensor data are transmitted to the mobile device in real time. This allows for the comparison of context classification between continuous and episodic sampling. In practice, the episodic sampling algorithm can be implemented on wearable sensors, and thus able to carry out local classification and signal the mobile device only when a change of context is detected, further improving the energy efficiency in computing the overall activity pattern.

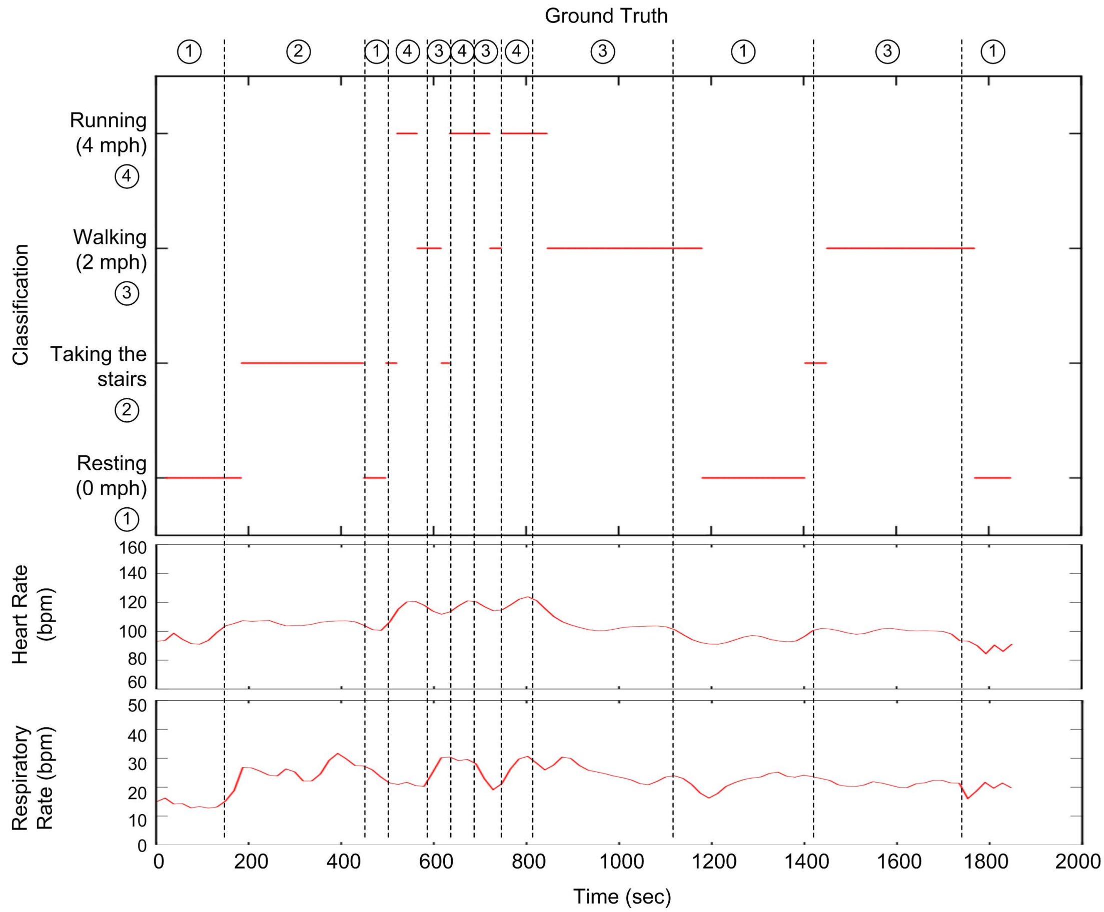

Fig. 4: Classification with Episodic Sampling

III. Experimental Setup

An experiment has been performed to identify an activity pattern with episodic sampling. The activity pattern involves exercise warmup, interval training, and gradual cool-down (see "Ground Truth" in Figure 4). The activities include resting, taking the stairs, walking (at 2 MPH), and running (at 4 MPH). Such an activity pattern resembles a subset of those typically performed by patients with COPD. The entire experiment lasts over 30 minutes. This activity pattern involves some behaviors with low dynamics (e.g. 5 minutes of rest) and some with high dynamics (e.g. interval training with 1-minute duty cycle).

The heart rate chest strap, the respiratory sensor, and the wearable device are worn on the chest of the body (see Figure 1). A single wearable device collects all the sensor data and transfers to the mobile device for post-processing and classification. For all eight channels of sensor data (three accelerometer outputs, three gyroscope outputs, the respiratory waveform, and the heart rate), each sample is 16-bit and acquired at 128 Hz.

Sensor data are first collected in real time continuously. The episodic sampling algorithm is then performed on the classification results to calculate the corresponding classification rate. T_{incr} is randomly chosen between 0 and 30 seconds in each iteration. The algorithm is performed on the same data set one hundred times, and Figure 4 corresponds to the classification result with the average T_{incr}.

IV. Results and Discussion

Figure 4 shows the classification result of the experiment with episodic sampling. One may observe that the classification result generated by episodic sampling resembles a shifted version of the ground truth. Furthermore, the algorithm is weak in tracking patterns with relatively high dynamics, which is primarily due to the choice of T_{incr}. Since T_{incr} reflects the dynamics of activities, one can determine T_{incr} more analytically (as opposed to a uniform distribution of values) through the use of statistical models such as Hidden Markov Models (HMMs). Nevertheless, the episodic sampling algorithm performs reasonably well with little implementation overhead.

Figure 5 shows the energy and storage savings with episodic sampling - 85% in energy and over 97% in data size respectively - with only 5% decrease in classification accuracy (85% with continuous, 80.5% average with episodic). Furthermore, with episodic sampling, the battery life can be extended to almost four days (with an average power of 31.6 mW), a significant improvement compared to 13.9 hours for continuous sampling (with an average power of 213.5 mW).

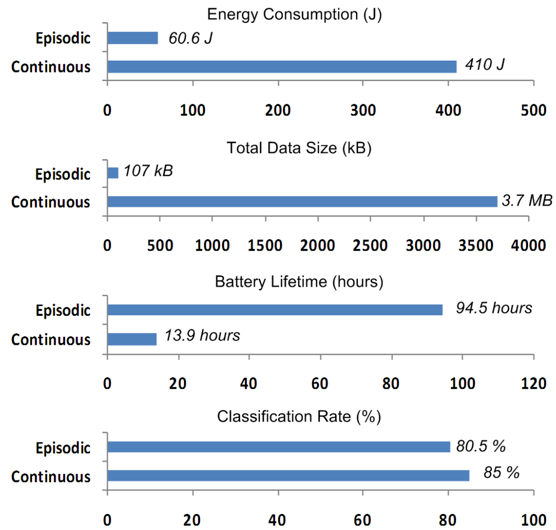

Fig. 5: Experimental Results - Continuous vs. Episodic Sampling

Results presented here are encouraging, and the authors plan to evaluate the algorithm more extensively in future studies with more representative patterns.

V. Conclusion

Episodic sampling offers a solution in enabling long-term activity monitoring related to wireless healthcare applications. With the support of both hardware system and software algorithms, the wearable sensors can operate efficiently by selectively choosing when to compute context information. The performance of the algorithm mainly depends on the context and the quality of the classifier. Nonetheless, initial results show a 85% saving in energy and over 97% saving in data size, with only 5% decrease in classification accuracy with a preliminary experiment. Future work includes the extension of the system to a model-based approach for describing the activity pattern of a subject and a corresponding model for describing the various operating states of the wearable sensors. For instance, the model may determine that capturing the physiological signals during interval training is important. To maximize classification accuracy, the wearable sensors may be left in continuous sampling during the period. This future approach will apply machine-learning techniques to select sampling episodes according to the model.

The granularity of the activities (context) also affects the performance of the algorithm because the variation in typical daily activities would be more dynamic. Future studies will include a much larger set of activities for longer periods of time. Lastly, the real-time energy information in ULEAP2 are not incorporated in the episodic sampling algorithm; future algorithms will include these energy information to achieve true online, context-aware, energy-efficient monitoring.

References

[1] P. Bonato, "Wearable Sensors/Systems and their Impact on Biomedical Engineering," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 22, no. 3, pp. 18–20, May-June 2003.

[2] P. Bonato, "Advances in Wearable Technology and Applications in Physical Medicine and Rehabilitation," *J Neuroengineering Rehabil*, vol. 2, no. 2, 2005.

[3] M. Moy, S. Mentzer, and J. Reilly, "Ambulatory Monitoring of Cumulative Free-Living Activity," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 22, no. 3, pp. 89–95, May-June 2003.

[4] W. Wu, M. Batalin, L. Au, A. Bui, and W. Kaiser, "Context-aware Sensing of Physiological Signals," in *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, Aug. 2007, pp. 5271–5275.

[5] T. Gao, D. Greenspan, M. Welsh, R. Juang, and A. Alm, "Vital Signs Monitoring and Patient Tracking Over a Wireless Network," *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, pp. 102–105, Jan. 2005.

[6] E. Jovanov, A. Milenkovic, C. Otto, P. D. Groen, B. Johnson, S. Warren, and G. Taibi, "A WBAN System for Ambulatory Monitoring of Physical Activity and Health Status: Applications and Challenges," in *Proc. International Conference on Engineering in Medicine and Biology Society (IEEE-EMBS)*, 2005, pp. 3810–3813.

[7] S. Park and S. Jayaraman, "Enhancing the Quality of Life Through Wearable Technology," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 22, no. 3, pp. 41–48, May-June 2003.

[8] S.-P. Tu, M. B. McDonell, J. A. Spertus, B. G. Steele, and S. D. Fihn, "A New Self-Administered Questionnaire to Monitor Health-Related Quality of Life in Patients With COPD," *Chest*, vol. 112, no. 3, pp. 614–622, 1997. [Online]. Available: http://www.chestjournal.org/content/112/3/614.abstract

[9] U. Anliker, J. Beutel, M. Dyer, R. Enzler, P. Lukowicz, L. Thiele, and G. Trster, "A Systematic Approach to the Design of Distributed Wearable Systems," *IEEE Trans. Comput.*, vol. 53, no. 8, pp. 1017–1033, Aug. 2004.

[10] T. Starner, "The Challenges of Wearable Computing: Part 1," *Micro, IEEE*, vol. 21, no. 4, pp. 44–52, Jul/Aug 2001.

[11] M. Satyanarayanan, "Pervasive Computing: Vision and Challenges," *Personal Communications, IEEE*, vol. 8, no. 4, pp. 10–17, Aug 2001.

[12] Pottie, Greg and Kaiser, William, *Principles of Embedded Networked Systems Design*. Cambridge University Press, 2005.

[13] D. McIntire, K. Ho, B. Yip, A. Singh, W. Wu, and W. J. Kaiser, "The Low Power Energy Aware Processing (LEAP) Embedded Networked Sensor System," in *IPSN '06: Proceedings of the fifth international conference on Information processing in sensor networks*. New York, NY, USA: ACM Press, 2006, pp. 449–457.

[14] L. Au, W. Wu, M. Batalin, D. Mclntire, and W. Kaiser, "MicroLEAP: Energy-aware Wireless Sensor Platform for Biomedical Sensing Applications," in *Biomedical Circuits and Systems Conference, 2007. BIOCAS 2007. IEEE*, Nov. 2007, pp. 158–162.

[15] L. K. Au. MicroLEAP2 (ULEAP2) Documentation. [Online]. Available: http://www.ascent.ucla.edu/wiki/index.php/Uleap/Hardware

[16] A. Krause, M. Ihmig, E. Rankin, D. Leong, S. Gupta, D. Siewiorek, A. Smailagic, M. Deisher, and U. Sengupta, "Trading off Prediction Accuracy and Power Consumption for Context-aware Wearable Computing," *Wearable Computers, 2005. Proceedings. Ninth IEEE International Symposium on*, pp. 20–26, Oct. 2005.

[17] ADXL330 Data Sheet. Analog Devices. [Online]. Available: http://www.analog.com

[18] IDG-300 Integrated Dual-Axis Gyroscope. InvenSense. [Online]. Available: http://www.invensense.com

[19] MLT1132 Piezo Respiratory Belt Transducer. ADInstruments. [Online]. Available: http://www.adinstruments.com

[20] Polar T31 Transmitter Belt. Polar Electro. [Online]. Available: http://www.polarusa.com

[21] J. J. Labrosse, *MicroC/OS II: The Real-Time Kernel*. Burlington, MA: CMP Books, 2002.

[22] L. Bao and S. Intille, "Activity Recognition from User-Annotated Acceleration Data," *Lecture Notes in Computer Science*, vol. 3001, pp. 1–17, March 2004.

[23] D. E. Comer, *Internetworking with TCP/IP (5th ed.)*. Upper Saddle River, NJ: Prentice Hall, 2005.