TinyEARS: Spying on House Appliances with Audio Sensor Nodes

Z. Cihan Taysi, M. Amac Guvensan*
Department of Computer Engineering
Yildiz Technical University
Istanbul, TURKEY

{cihan, amac}@ce.yildiz.edu.tr

Tommaso Melodia
Department of Electrical Engineering
State University of New York at Buffalo
Buffalo, NY 14260 USA

tmelodia@buffalo.edu

Abstract

Fine-grained awareness on how and where energy is spent is being increasingly recognized as the key to conserve energy. While several solutions to monitor the energy consumption patterns for commercial and industrial users exist, energy reporting systems currently available to *residential* users require time-consuming and intrusive installation procedures, or are otherwise unable to provide device-level reports on energy consumption. To fill this gap, this paper discusses the design and performance evaluation of the Tiny Energy Accounting and Reporting System (TinyEARS), a fine-grained energy monitoring system that generates device-level power consumption reports primarily based on the acoustic signatures of household appliances. Experiments demonstrate that TinyEARS is able to report the power consumption of individual household appliances within a 10 % error margin.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Measurement, Design, Experimentation

Keywords

Energy Monitoring, House Appliances, Audio Data Classification, Wireless Audio Sensor Networks

1 Introduction

Residential spaces account for approximately 21% of the total energy consumption in the United States [1], with raising figures worldwide. This motivates the growing interest in smart building technology to automate and otherwise

---

* The research described in this paper was carried out while Z. C. Taysi and A. Guvensan were visiting Ph.D. students at the State University of New York at Buffalo.

promote energy conservation in residential and commercial spaces. However, while businesses rely increasingly on procedures and best practices to save energy and reduce costs, most home users do not have any means to control their energy usage patterns. The main resources in a typical household are electricity, water, natural gas, and heating oil. Saving a small portion of each in each residential space could have a significant impact on reducing costs, energy consumption, and impact on the environment. Several studies [2][3] have shown the necessity of fine-grained energy monitoring to encourage conservation.

Motivated by this, in this paper we discuss the design and implementation of a wireless monitoring system for residential spaces based on a multi-layer decision architecture, called Tiny Energy Accounting and Reporting System (TinyEARS). The objective of TinyEARS is to detect and classify "on" devices based on their acoustic signatures and report device-level energy consumption by correlating node decisions with time and power information obtained from a real-time power meter.

The paper discusses the architecture, signal acquisition and processing, and algorithmic details of TinyEARS. The key contributions of our work are as follows:

- The multi-layer architecture of TinyEARS enables fine-grained energy monitoring at the device level by leveraging the acoustic signatures of house appliances. This information is correlated with the overall power usage information of the house obtained from a real-time power meter;

- While most of the existing solutions for device-level monitoring require one sensor node per appliance, in TinyEARS a single sensor node monitors house appliances in a room based on their acoustic signatures. Deploying one sensor node per room reduces both the overall cost and communication burden of the wireless sensor network;

- TinyEARS, being based on a limited number of sensor nodes, is an easily deployable and maintainable system;

- Our study shows that house appliances can be recognized with an overall success rate of 94% by their acoustic signatures with relatively simple processing algorithms implemented on the motes;

- Finally, we discuss the main system design challenges

and their solutions in implementing an audio classification process on an Imote2 sensor node.

The rest of this paper is organized as follows. In Section 2, we review existing solutions for energy monitoring in residential spaces. In Section 3, we introduce the system architecture of TinyEARS. Section 4 describes the details of the house appliance sound recognition system (HASRS). In Section 5, we describe the data acquisition and correlation process of TinyEARS. In the same section, we also discuss the issues and challenges of implementing the HASRS on Imote2. Section 6 describes our experimental environment and presents the test results obtained using Imote2 and the TED5000 power meter.

2 Related Work

Previous work on monitoring energy consumption in residential spaces can be broadly classified into two categories. The first group is concerned with systems designed to measure the overall energy consumption with a single sensor, usually located in a power box. The second approach is to monitor each household appliance individually, with fine-grained consumption feedback.

There exists several commercially available products to measure the overall energy consumption of a household, including Power Cost Monitor [4], Wattson [5], and TED-xxxx [6]. Although it is moderately difficult to install these devices, they do not require any maintenance once deployed. However, while these products are able to present the overall energy consumption and detect anomalies in energy usage, they cannot provide per-appliance energy measurement. A different product, the Google Power Meter [6], is an opt-in software tool that allows users to visualize detailed home energy information.

Device-level monitoring has recently received attention in the literature because it is able to provide fine-grained feedback on energy consumption. To monitor each device individually, two approaches have been considered in the literature, i.e., (i) installing electrical current sensors inline with each appliance, and (ii) deploying multiple sensors throughout the household. Some commercial products, based on electrical current sensors, are already available, e.g., Plogg [7], Kill-a-Watt [6], and Watts Up [6]. Although these products provide fine-grained energy monitoring, they require inline installation between a standard AC plug and the outlet. Therefore, some appliances such as heating and ventilation systems (HVAC), and ceiling lights can not be easily instrumented, since they lack AC plugs and have wired connections.

Researchers have recently proposed deploying sensor networks in residential spaces to achieve device-level monitoring of energy consumption. In [6], the authors developed Viridiscope, a system designed to provide device-level power consumption feedback by using magnetic, acoustic and light sensors. Ambient signals from sensors placed near appliances are used to estimate power consumption.

Schoofs et al. [8] propose ANNOT (Automated Electricity Data Annotation), a system to automate electricity data annotation leveraging cheap wireless sensor nodes such as temperature, sound, light and accelerometer. ANNOT automatically acquires appliance signatures, collects training
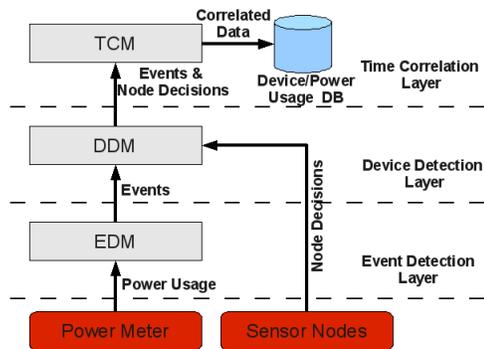
Figure 1. Multi-layer decision architecture of TinyEARS.

data, validates the monitoring output without human supervision, and is integrated within the RECAP (Recognition and Profiling of Appliances) appliance load monitoring system.

Marchiori and Han [9] explore an alternative approach to monitor household energy usage, including small devices. They propose using circuit-level energy measurements as a compromise between the two aforementioned approaches. The drawback of this approach is that the accuracy of the system decreases dramatically with increasing number of house appliances on each circuit.

Finally, several studies [10][11][12] use wireless sensor networks for energy efficiency without directly addressing energy monitoring in residential spaces. For example, in [10] the authors design a system where system parameters are automatically set according to user profiles to minimize the energy consumption while guaranteeing a desired comfort level. Erickson et al. [12] achieve a 14% reduction in HVAC energy usage by having an optimal control strategy based on occupancy estimates and usage patterns.

TinyEARS differs from existing solutions by combining information from acoustic signatures of house appliances and power meter readings through a multi-layer architecture. The system allows deploying one single sensor node per room. TinyEARS can potentially combine the individual features of different types of sensors such as light, temperature and magnetic to improve the accuracy of energy consumption in a residential space. However, this paper is focused on the detection of household appliances emitting acoustic signals.

3 System Design

TinyEARS is composed of a multi-layer decision architecture that includes (i) an event detection layer, (ii) a device detection layer and (iii) a time correlation layer, as shown in Figure 1. Operations within each layer are handled by individual software modules running on the Data Fusion Center (DFC). The Event Detection Module (EDM) communicates with the real-time power meter and uses a basic filtering mechanism to detect changes in the use of house appliances. When an event is detected, the EDM alerts the Device Detection Module (DDM). The DDM controls communication with the sensor nodes through control packets that trigger sensor nodes to start collecting audio samples. Based on the collected samples, each individual sensor decides which (if any) house appliances are active. Each sensor node reports

Figure 2. MTU Connection.

Table 1. Comparison of different types and number of features.

	Physical Features	MFCC (5)	MFCC (9)	MFCC (13)
Recognition Ratio(%)	76.78	89.65	93.12	94.35
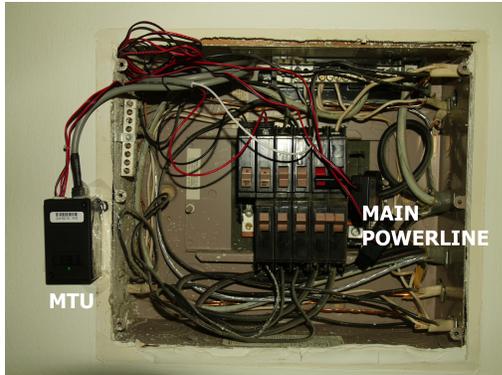
its decision back to the DDM along with a confidence value. The DDM sends event information and node decisions to the Time Correlation Module (TCM). The TCM may decide to rely on the node decision or override it with a better match. The TCM creates alternative decisions by correlating time, power usage and node decisions. A set of rules have been defined for the correlation operation, whose details will be discussed in Section 5.5.

TinyEARS consists of three main components: a real-time power meter, a DFC, and sensor nodes. The real-time power meter is used to monitor changes in power usage in the household. There are several off-the-shelf power meters with different capabilities and prices. We have used the TED 5000 produced by Energy Inc., because of its simple installation procedure, very accurate readings, and open development API. The TED 5000 system consists of three units: a measuring transmitting unit (MTU), a gateway, and a user display unit. The MTU is mounted on the main electric panel of the house as shown in Figure 2. It measures and transmits energy, power, and voltage information to the gateway, which can be plugged to any outlet in the house. Real-time and historical data can be accessed by utilizing the gateway's Ethernet interface. The DFC is a typical PC that runs EDM, DDM, TCM modules, and a configuration utility to manage initial training of TinyEARS.

In TinyEARS, sensor nodes are required to perform signal processing operations including Fast Fourier Transform (FFT), feature extraction and classification. Therefore, we have used the Imote2 [13] sensor node platform that can satisfy the processing and memory requirements of these algorithms [14].

4 On-Board Real-Time Device Classification

Discriminating devices based only power meter readings or on electrical noise on power lines is a complex task. Therefore, we propose an advanced architecture, TinyEARS, where the key idea is to infer the energy consumption of house appliances by primarily relying on their acoustic signature. The fact that most house appliances (e.g, refrigerator, dishwasher, exhauster, blender, vacuum cleaner, washer, and hair dryer) contain a motor suggests that it is possible to recognize them and infer their state (active or inactive) based on their acoustic signature [15]. To the best of our knowledge,

however, the sound characteristics of house appliances have not been explored yet. We have therefore concentrated on feature extraction and classification algorithms that had been successfully applied before to mechanical sound classification problems such as vehicle classification.

4.1 Feature Extraction

In this study, seven different audio features are primarily considered, including zero-crossing rate (ZCR), short-time energy (STE), band-level energy (BLE), spectral-centroid (SC), spectral roll-off (SRO), spectral flux (SF), and mel-frequency cepstral coefficients (MFCC) [15] . To obtain a high classification accuracy, we have selected effective and robust feature combinations for discriminating the acoustic signatures of household appliances. The preliminary analysis of power spectral densities of house appliances obtained via an audio signal processing tool AudaCity shows that the discriminative frequency band is between 0 and 1 kHz. Before implementing our algorithms on the sensor nodes, we ran extensive MATLAB tests whose results are summarized in Table 1 to validate our feature set.

4.1.1 Mel Frequency Cepstral Coefficients

The MFCC, introduced in [15] as a candidate feature extraction method for stationary audio signals, perform better than other feature extraction methods such as Fourier transform, human factor cepstral coefficients, fast wavelet transform and short-time Fourier transform. Since the sounds of house appliances can be considered to be stationary, we applied the MFCCs for our audio classification problem. There are five main processing steps for obtaining MFCC features of an audio signal. First, the audio signal is fragmented into frames of predefined length. Then, each frame is multiplied with a Hamming window to maintain the continuity between the first and the last points in the frame. Then, the FFT is applied to the signal and smoothened by a series of triangular filters, including 13 linear filters below 1 kHz and 27 logarithmic filters within $1 - 6.4$ kHz. The MFCCs are finally calculated. The complete coefficient extraction procedure is described in [16].

Since discriminative frequencies of each house appliance are concentrated below 1 kHz, we omitted logarithmic filters that are applied only to the frequencies above 1 kHz and evaluated our audio classification process for different number of cepstral coefficients to compare the recognition success rate.

4.2 Classification

In the classification step, the objective is to recognize "on" devices based on the MFCC features. Several classification algorithms, including support vector machines (SVM) and k-nearest neighbor (k-NN), have been previously proposed for environmental sound classification [15].

During our preliminary design phase, we compared the household appliance recognition success rates of MDC, k-

Table 2. Comparison of k-NN, SVM and MDC with 9 MFCC features of 7 house appliances.

	MDC	k-NN	SVM
Success Rate (%)	93.69	97.14	97.74

NN and SVM through Matlab simulations. Our findings, reported in Table 2, show that the success rates do not vary significantly for the considered application. The performance of k-NN and SVM is only slightly better than MDC. Therefore, even though Imote2 motes have the capability to run complex classification algorithms, we decided to use a simple classifier as MDC, characterized by a relatively low computational cost, to speed up the decision process on each sensor node and reduce the energy consumption.

4.2.1 Minimum Distance Classifier

The minimum distance classifier algorithm aims at finding the closest class centroid to the given test sample.

Let $\mathbf{x}$ be an unknown sample to be classified, and $\mathbf{y}_i$, $i = 1, ..., n$ be a prototype for class (c_i). Both $\mathbf{x}$ and $\mathbf{y}_i$ are m-dimensional vectors in the feature space, n is the number of classes, and m is the number of dimensions of the feature space. The minimum distance classifier is defined as

$$\mathbf{x} \in (c_i) \equiv d(\mathbf{x}, \mathbf{y}_i) = \min_{j \in 1, ..n} d(\mathbf{x}, \mathbf{y}_j),$$

where $d(a, b)$ represents the Euclidean distance function.

4.2.1.1 Training Process

The system needs to be trained using the sound of each house appliance to generate their acoustic signatures. The training process is done only once after deploying the sensor nodes. This process takes place on the DFC. The acoustic signatures of each house appliance are generated at the end of this process. After training the system on the DFC, the mean values and scale coefficients of each class, which are used to normalize the feature set, are delivered to the corresponding sensor node in the house.

4.2.1.2 Test Process

In the test phase, after sampling audio data, the sensor node extracts the MFCC features of the signal. Then, it calculates the Euclidean distance to each class. The class with minimum distance and the distance are sent to the DFC.

5 Data Acquisition and Processing

TinyEARS requires real-time power usage information and node decisions to estimate the power consumption of individual house appliances.

5.1 Interfacing With the Power Meter

Energy Inc. provides an XML-based development API to retrieve real-time and historical data from the gateway. Interfacing with the real-time power meter involves two operations. First, an HTTP request is sent to the TED5000 gateway. The gateway replies back with a XML file that contains information including a time stamp, voltage, and power from several MTUs. In the second step, this XML file is parsed to extract relevant information, i.e., time stamps and power usage of a particular MTU.

Table 3. Sample Data Fusion Operation.

Time	10:10:54	10:14:41	10:21:12
Power Usage	348.9	410.4	335.5
Distance Value	0.105	0.165	0.129
Node Decision	Ref.	Exh.(H)	Ref.
Correlated Data	Ref.	Ref & Exh.(H)	Ref.
Real Data	Ref.	Ref & Exh.(H)	Ref.

5.2 Data/Event Filtering

The DFC monitors power changes in the household and decides whether a device has been turned on or off. In our experiments, we have noticed two types of anomalies related to the power consumption of house appliances. First, the power consumption of some devices in the household varies in time. Second, some devices cause a peak in power consumption when they are turned on and then descend to a lower power consumption level. Thus, basic approaches like using the difference between the last two readings fail to correctly detect changes in device states. In our system, we employ two basic filtering techniques to overcome this problem. First, when a change occurs between the last two readings, the difference in power usage is tested against a pre-defined threshold. The purpose of this threshold is to eliminate the effect of devices with varying power consumption. If the change is greater than the threshold, a second filter is applied. The DFC keeps monitoring the power usage for five additional seconds and calculates the average of these values. This low-pass filter is used to remove peaks and estimate the actual steady-state change in current absorption caused by the device. This value is also used by the TCM to validate and (possibly) override node decisions.

5.3 Audio Sampling at the Sensor Node

The ITS400 [13] sensor board is designed to be interfaced with Imote2. It contains a three-axis accelerometer, an advanced temperature/humidity sensor, a light sensor and a 4-channel Analog/Digital Converter (ADC). The ADC on this board can be used to interface with different analog sensors. In our application, one of the four available ADC channels is used to interface an electret microphone. The ADC on the ITS400 sensor board can provide 10- and 12-bit samples. To preserve the quality of the captured sound, we use 12-bit samples. Figure 3 shows the microphone connected to the Imote2 sensor node.

5.4 Implementing Device Classification on Imote2

In our study, we have implemented the house appliance sound recognition system (HASRS) component of TinyEARS on imote2 sensors. Each mote is responsible for its individual room/space. Motes wait in the idle state until they receive a command from the DFC. In our tests, the mote is always in the "on" state and monitors the channel for command messages. After receiving the command, the mote performs audio classification including sampling audio data, extracting features (MFCC), processing minimum distance
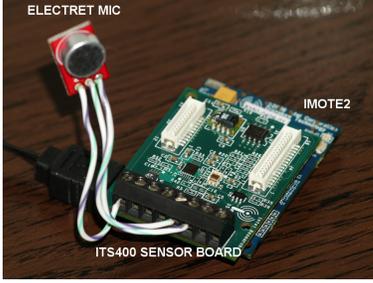
Figure 3. Imote2, ITS400 and electret microphone.
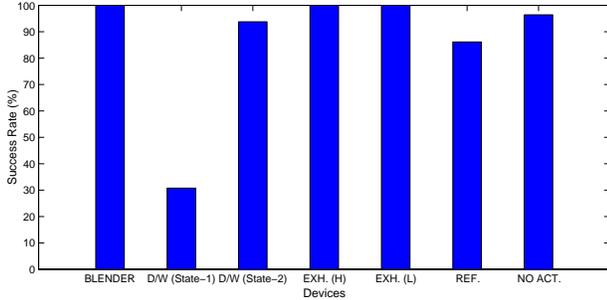
Figure 4. Success Rate of Audio Classification.

classifier. After finalizing a decision, each mote transmits the decision class label and the distance value to the DFC.

We faced a number of challenges in implementing a fully functional HASRS component. The .Net Micro Framework has some restrictions that need to be worked around to implement complex applications like MFCC. This includes the lack of mathematical methods such as $log()$, $sin()$, $cos()$, which led us to override the existing math library with a new library from ALGLIB [17].

5.5 Data Fusion Rules

When a sensor node performs classification of the sampled audio, it calculates the distance between the sample and all possible classes. Then, it classifies the sample as belonging to the class with minimum distance. In the multilayer decision architecture of TinyEARS, the distance value is used by the TCM module. In the case where only a single device is working, the HASRS identifies the sample as belonging to one of the classes in the system with very low distance. The distance is higher when the sampled audio contains superimposed sounds of multiple devices. The distance value is then used by the TCM to discriminate and make decisions in situations when multiple devices are active at the same time.

If the distance value is lower than the predefined threshold, the DFC accepts the decision of the node as is. If the distance value is greater than the threshold, the DFC combines time and power consumption information and may override the node decision. This override operation is performed based on a set of predefined rules.

TCM identifies node decisions with low confidence and calculates the change in power usage Δpm. A positive change means that a device has been turned on, whereas a negative change indicates that a device has been turned off.

Table 4. Real Working Time vs. Estimated Working Time.

Device Name	Working Time		Success Rate(%)
	Real [min]	Estimated [min]	
Blender	3	3	100
D/W (State 1)	5.3	7.15	65
D/W (State 2)	33.5	32.05	95.67
Exh. (High)	14.46	14.46	100
Exh. (Low)	7	3.5	50
Refrigerator	154.36	158	97.64

When Δpm is positive, a list of possible devices is generated based on Δpm and known levels of power consumption of devices. Then, unfeasible options are eliminated based on previous node decisions. When Δpm is negative, the power consumption of each device in the previous time slot is compared with Δpm, and any matching device is marked as "turned off" for the current time slot. Table 3 shows an example of the time correlation operation. The TCM considers node decisions with a distance value below the threshold (set to 0.140 in the experiments reported).

6 System Evaluation
6.1 Deployment

To test and validate our design, we have deployed TinyEARS in a two bedroom apartment in Buffalo, NY. During deployment, a TED5000 power meter was connected to the electric panel, a laptop configured as DFC, and a sensor node deployed in the kitchen. In-situ training phase of TinyEARS enables a flexible deployment. Thus, sensor node can be placed anywhere in the room. The kitchen was selected since it is the most challenging room as multiple appliances are often used there concurrently.

During the tests, occupants of the house took accurate notes of "turn on" and "turn off" times of each device in the house. A ground truth for working devices and their power consumption was constructed by combining these notes and real-time power readings gathered from power meter on every second. This ground truth is only used to evaluate the system performance during the test phase. This experimental deployment was held for three days. Audio samples from the first day were used as a training set for the system. The performance of TinyEARS was evaluated in the next two days.

6.2 Test Results

To evaluate the performance of TinyEARS, we have observed four house appliances and their different run levels, i.e., *refrigerator, dishwasher (state 1 and state 2), exhauster (low and high) and blender*. Since TinyEARS is based on a multi-layer decision architecture, we assess the success rates of each layer individually.

The first layer of TinyEARS is the node level decision. The audio classification module was tested for 104 different activities. There might occur 7 different individual activities in the kitchen, i.e., *no activity, blender, dishwasher (state1, state 2), exhauster (low, high), and refrigerator*. The success rate of the recognition process for each appliance and its run level is shown in Figure 4.

Table 5. Success Rate of TinyEARS.

Device Name	Power Consumption		
	Real [W]	Estimated [W]	Success Rate(%)
Refrigerator	411.62	421.33	97.69
D/W (State 1)	36.66	47.67	76.90
D/W (State 2)	502.5	480.75	95.67
Exh. (High)	26.51	26.51	100
Exh. (Low)	8.16	4.08	50
Blender	15	15	100
Overall	1000.45	995.34	99.49

The overall system performance was evaluated by comparing power consumption information reported by TinyEARS to the ground truth data of two-day long activities. Table 5 shows the success rate of TinyEARS in estimating the power consumption of each house appliance.

At the second layer, the DFC correlates node decisions, time and power usage information to detect working house appliances. Table 4 shows the estimated working time of the monitored house appliances based on this correlation and their real working time.

7 Conclusions and Future Work

In this paper, we presented TinyEARS, a fine-grained energy monitoring system for residential spaces using indirect audio sensors. TinyEARS is designed to eliminate the complexity of the installation and maintenance procedures of existing power metering solutions. TinyEARS employs a multi-layer decision architecture that exploits the true capabilities of sensor nodes by processing sensor data on the node itself, thus limiting transmissions over the wireless channel. This architecture combines various sources of information to reach an accurate estimation of the consumption patterns. Experiments show that the multi-layer architecture of TinyEARS enables monitoring device-level power consumption with less than 10 % error. In addition, TinyEARS can easily monitor the power consumption of multiple simultaneously active appliances as well as appliances with variable power consumption.

8 References

[1] Y. Kim, T. Schmid, M.B. Srivastava, and Y. Wang. Challenges in resource monitoring for residential spaces. In *ACM Workshop On Embedded Sensing Systems For Energy-Efficiency in Buildings (BuildSys 2009)*, San Francisco, USA, November 2009.

[2] S. Darby. The effectiveness of feedback on energy consumption: a review of defra of the literature on metering, billing and direct displays. http://www.eci.ox.ac.uk/research/energy/electric-metering.php, 2002.

[3] P. C. Stern. Information, incentives, and proenvironmental consumer behavior. *Journal of Consumer Policy*, 22(4):393–509, December 1999.

[4] Power Cost Monitor Specifications. http://www.bluelineinnovations.com.

[5] WATTSON Specifications. http://www.diykyoto.com.

[6] Y. Kim, T. Schmid, Z. Charbiwala, and M. Srivastava. Viridiscope: Design and implementation of a fine grained power monitoring system for homes. In *ACM Proc. of Intl. Conf. on Ubiquitous Computing (UbiComp 2009)*, pages 245–254, Orlando, USA, September 2009.

[7] Energy Optimizers Ltd. http://www.plogginternational.com.

[8] A. Schoofs, A. Guerrieri, D.T. Delaney, G. O'Hare, and A.G. Ruzzelli. Annot: Automated electricity data annotation using wireless sensor networks. pages 1–9, Boston, MA, USA, June 2010.

[9] A. Marchiori and Q. Han. Using circuit-level power measurements in household energy management systems. In *ACM Workshop On Embedded Sensing Systems For Energy-Efficiency in Buildings (BuildSys 2009)*, San Francisco, USA, November 2009.

[10] A. Barbato, L. Borsani, A. A. Capone, and S. Melzi. Home energy saving through a user profiling system based on wireless sensors. In *ACM Workshop On Embedded Sensing Systems For Energy-Efficiency in Buildings (BuildSys 2009)*, San Francisco, USA, November 2009.

[11] G. Gao and K. Whitehouse. The self-programming thermostat: Optimizing setback schedules based on home occupancy patterns. In *ACM Workshop On Embedded Sensing Systems For Energy-Efficiency in Buildings (BuildSys 2009)*, San Francisco, USA, November 2009.

[12] V. L. Erickson, Y. Lin, A. Kamthe, R. Brahme, A. Surana, A. E. Cerpa, M. D. Sohn, and S. Narayanan. Energy efficient building environment control strategies using real-time occupancy measurements. In *ACM Workshop On Embedded Sensing Systems For Energy-Efficiency in Buildings (BuildSys 2009)*, San Francisco, USA, November 2009.

[13] Crossbow Imote2,ITS400 Sensor Board Specifications. http://www.xbow.com.

[14] I. F. Akyildiz, T. Melodia, and K. R. Chowdhury. Wireless multimedia sensor networks: Applications and testbeds. *Proceedings of the IEEE*, 96(10):1588–1605, Oct. 2008.

[15] M. Cowling. *Non-Speech Environmental Sound Classification System for Autonomous Surveillance*. PhD Thesis, Faculty of Engineering and Information Technology, Griffith University, 2004.

[16] J.-C. Wang, Jhing-Fa, and Y.-S. Weng. Chip design of MFCC extraction for speech recognition. *Integr. VLSI J.*, 32(1-3):111–131, 2002.

[17] ALGLIB Specifications. http://www.alglib.net/.