

ViridiScope: Design and Implementation of a Fine Grained Power Monitoring System for Homes

Younghun Kim, Thomas Schmid, Zainul M. Charbiwala, and Mani B. Srivastava

Networked and Embedded Systems Lab.,
Electrical Engineering Department,
University of California, Los Angeles

{kimyh,thomas.schmid,zainul,mbs}@ucla.edu

ABSTRACT

A key prerequisite for residential energy conservation is knowing when and where energy is being spent. Unfortunately, the current generation of energy reporting devices only provide partial and coarse grained information or require expensive professional installation. This limitation stems from the presumption that calculating per-appliance consumption requires per-appliance current measurements. However, since appliances typically emit measurable signals when they are consuming energy, we can estimate their consumption using indirect sensing. This paper presents ViridiScope, a fine-grained power monitoring system that furnishes users with an economical, self-calibrating tool that provides power consumption of virtually every appliance in the home. ViridiScope uses ambient signals from inexpensive sensors placed near appliances to estimate power consumption, thus no inline sensor is necessary. We use a model-based machine learning algorithm that automates the sensor calibration process. Through experiments in a real house, we show that ViridiScope can estimate the end-point power consumption within 10% error.

ACM Classification Keywords

H.4 Information Systems Applications: Miscellaneous

General Terms

Algorithms, Design, Experimentation, Human Factors, Measurement

Author Keywords

Adaptive Sensor Calibration, Machine Learning, Nonintrusive and Spatially Distributed Sensing

INTRODUCTION

Natural resource preservation has recently become a significant concern, and has motivated research and development efforts to assist in both conservation and management. As

Chetty et. al. [8] observed last year, researchers in ubiquitous computing have a major role to play by developing technologies that encourage people to efficiently manage energy consumption at home. A fundamental component of such technologies is fine grained monitoring to measure real-time consumption of each domestic appliance [14, 18]. Many researchers and companies have begun addressing this problem. For example, Jiang et. al. [17] developed a wireless networked sensor that measures the power consumption at a power outlet; while other devices that sense the real-time electricity consumption for a whole house (e.g. Cent-a-Meter), per circuit (e.g. EM-2500 [1], TED [2]) or per outlet (e.g. Kill-A-Watt and Watts Up [11]) are now commercially available.

Although these devices are promising, each has some drawbacks. For example, Cent-a-Meter, EM-2500 and TED devices monitor the household power consumption but do not provide per-appliance level measurements. Kill-A-Watt and Watts Up devices provide finer granularity but require inline installation between a standard AC plug and the outlet. While it is possible to instrument many appliances in this way, some of the major energy consumers cannot be easily instrumented. For example, most heating and ventilation systems (HVAC) and electric boilers do not have standard AC plugs, or are hard-wired to the main power lines. Ceiling lights are another example. We conclude that the current generation of energy reporting devices can simultaneously provide only two of the following three features essential for true ubiquity:

- *Comprehensive Coverage*: The system should monitor all the appliances in the home, making information complete for users.
- *Fine-grained Reporting*: The system should be able to report individual consumption profiles for each appliance as this will enable targeted conservation.
- *Easy and Seamless Installation*: The system should be installable by a non-professional, e. g. no modifications to power lines or power cords should be necessary. Further, user intervention for configuration, calibration and maintenance should be minimal.

Simultaneously satisfying the above criteria precludes the sole use of traditional direct monitoring techniques, and a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UbiComp 2009, Sep 30 – Oct 3, 2009, Orlando, Florida, USA.

Copyright 2009 ACM 978-1-60558-431-7/09/09...\$10.00

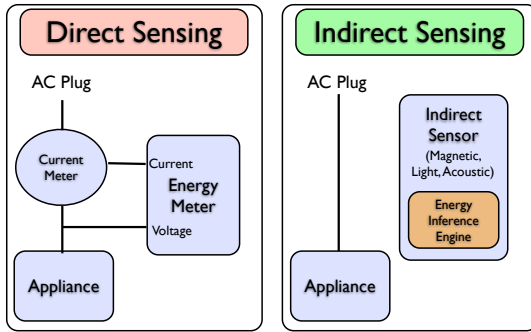


Figure 1: Direct power monitoring devices need to be installed in series with the appliance, whereas our indirect power monitoring concept senses signals emitted from appliances making it less invasive (Left adapted from [17]).

new monitoring dimension that does not require an in-line sensor is necessary. To this end we present ViridiScope, a spatially fine grained power monitoring system for residential spaces. The ViridiScope system provides real-time appliance-level power estimation by extensively leveraging ubiquitous sensing and computing devices, including magnetic, acoustic and light sensors. The *modus operandi* is a network of radio-enabled distributed sensors monitoring signals that appliances emit and forwarding them to a personal computer acting as a back-end fusion center. The fusion center collects data from the heterogeneous sensors and measurements from the main power meter, and runs a model-based machine learning algorithm that automatically learns and estimates power consumption of every appliance on-the-fly. The contributions of this paper are three fold:

Introducing Indirect Power Monitoring Concept:

Our approach is based on the fact that an *appliance emits measurable signals when it consumes energy*. By sensing these signals we can estimate power consumption (Figure 1). The core advantage is that one does not need to install a monitoring device in-line with the electric wire. Instead an indirect sensor near a power line or the appliance is enough for accurate power estimation. For example, we show in the Problem Description Section that simply placing a magnetic sensor near a power cord or a light intensity sensor near a ceiling light is sufficient to estimate their consumption profile.

Autonomous Sensor Calibration Framework:

The application of indirect sensing introduces a crucial technical challenge because the indirect sensors are inherently more prone to variations in ambient conditions, e.g. the actual sensor placement, surrounding materials [4]. To estimate the power consumption of an appliance based on signals that an indirect sensor is reporting, we need to find a unique mapping from the signals to power consumption. Because of uncertainties in installation, changes in ambient conditions, and sensor variability, it is virtually impossible to calibrate indirect sensors prior to installation. Therefore, our approach uses *in-situ* sensor calibration. Unfortunately, *in-situ* calibration often involves extensive manual interaction. To avoid this we develop an




Figure 2: The magnetic field change near a PC looks like noise, although we can infer whether the PC is active or not. The noisy region indicates that the PC is active.

autonomous sensor calibration framework. We formulate a model based regression problem that combines the aggregate power consumption and signals from indirect sensors. Our framework is able to incorporate with uninstrumented appliances as well as direct monitoring devices.

Prototype Deployment in a Realistic Setting:

We conduct a small scale deployment in a 2-person apartment unit to show the feasibility of this approach. We find that the indirect monitoring is feasible and results in less than 10% of estimation error.

MOTIVATION

Various studies have shown the necessity of fine grained energy monitoring to encourage conservation. For example, McMakin et. al. [21] find that to sustain changes in human behavior for energy conservation, continued awareness is as important as incentives and disincentives. Stern [24] suggests that information awareness has a synergistic effect with incentives in energy conservation. He states that awareness not only encourages people to actively participate and do the right thing, but also provides indirect monetary benefits because of reduced cost of resources. Darby [10] also concluded that feedback about consumption is essential for energy savings, and that immediate direct feedback can be extremely valuable in influencing behavior with savings in the range of 5-15%. The World Business Council for Sustainable Development in [27] noted, “Lack of awareness and information on energy consumption and cost - people are often not aware that they are wasting energy - prevents them from behaving efficiently.” They also note, “Technical devices to measure energy consumption and provide immediate feedback could help households cut energy consumption by as much as 20%. Direct and immediate feedback reveals the link between actions and their impacts. Well informed consumers choose actions to save energy with minimal impact on their comfort.” We believe that helping individuals know the *what*, *when*, and *where* of their energy usage, and identifying wasteful patterns of usage, encourages modifying wasteful behaviors.

PROBLEM DESCRIPTION AND SYSTEM DESIGN

To show the feasibility of indirect power monitoring, we present some experimental data that illustrates that the measured signals from magnetic, light and acoustic sensors are




Figure 3: Correlation between power consumption and the magnetic field variation near AC power wires of a PC set, a hotairgun and a message chair. It is clear that they are very closely correlated, and we exploit this in our algorithm.

good proxies for the appliance power consumption. We then propose a system architecture that consists of a two tiered information and sensing hierarchy. This architecture details the sensor configuration and information flow among ubiquitous sensing and computing devices. To make the monitoring system *self-configurable*, we formulate an autonomous calibration mechanism that *automatically* trains the indirect sensors so that it can compute the appliance-level power consumption in real-time.

Theory of Operation: Indirect Sensors and Power Estimation

The prototypical ViridiScope system infers per-appliance level power consumption using three types of indirect sensors: magnetic, acoustic, and light. A magnetic sensor near the power cord of an appliance is a good monitor of the current draw because the magnetic field variations near the power wires are strongly correlated with the current flow inside the wires [7]. In fact, magnetic field transducers are regularly used for current sensing [16, 17].

ViridiScope exploits the noisy magnetic field changes near an appliance (Figure 2). We use a statistical signature, *the standard deviation of the magnetic field* near the power wire. This signature has a strong correlation with the real power consumption (W). This is intuitive because when an appliance is active, electrons are flowing and generate magnetic field in the air [7]. Figures 3 (a), (b), and (c) demonstrate that the standard deviation is a good proxy for the power consumption of a device. This phenomena is universal to all AC powered appliances. Therefore, a sensor node equipped with a one dimensional magnetometer is adequate to capture the variable power consumption of an appliance.

Different from a magnetic sensor, a light intensity or an acoustic sensor can detect the internal power state of an appliance. Fortunately, most appliances have a very limited number of power states. Knowing these states and their average power consumption is sufficient to describe the overall power consumption of an appliance [15, 19, 20, 23]. For example, a simple light can only be *On* or *Off*. Detecting these on and off states is simple, and is enough information to in-




Figure 4: A refrigerator is a good example of an appliance having several power states. It has two main components: a door light and a compressor making four effective power states (Both off, Light on, Compressor on, and Both on).

fer the light’s power. A refrigerator, on the other hand, has four power states, *Compressor On*, *Inside Light On*, *Both On*, *Both Off*, since it has two active components: a compressor and a door light (see Figure 4). A combination of a light and acoustic sensor can effectively detect all four states. For example, a light intensity sensor inside the refrigerator can easily detect the on/off status of the door light (Figure 6). In addition, a very primitive microphone sampling at 4Hz is adequate for detecting the power state of a compressor in a typical refrigerator (Figure 5). All these examples support the concept of indirect power monitoring.

Technical Challenges

The biggest technical challenge with indirect sensing is calibration. Even if the form of the *calibration function* is known, the calibration parameters for each sensor must be identified after the sensor is installed. Thus, our approach requires *in-situ* sensor calibration. *In-situ* calibration traditionally requires a well trained technician [3, 26].

To automate sensor calibration, ViridiScope exploits an additional piece of information from the existing infrastructure: the main power meter that provides real-time aggregated power consumption for the whole household. Since most homes have one main power-line, this information is easily available. Using this information, we develop an *au-*




Figure 5: An acoustic sensor at 4Hz sampling rate can detect the on/off status of the compressor.




Figure 6: Light intensity in a refrigerator goes up when the door is open and the inside light is on.

onomous calibration algorithm that calibrates the indirect sensors with minimal human intervention.

Another problem is that indirect sensors are generally more prone to external noise, which makes reliable detection difficult. This is resolved by adding adaptive filtering mechanisms. We discuss and analyze these in the Practical Consideration Section.

System Architecture

Information Hierarchy

ViridiScope consists of a two tiered information hierarchy that contains one utility installed meter (or a main meter monitor), and multiple uncalibrated indirect sensors for energy consuming end-points (Figure 7). The first tier provides the aggregated power consumption, $y_0(t)$. The second tier sensors monitor energy consuming related signals, $s_i(t)$.

The characteristics of this hierarchy are as follows. (1) The first tier is reliable and the utility provider maintains it. (2) The calibration of the non-intrusive indirect sensors can be automated using correlations between the first and second tier sensors. This makes the monitoring system readily deployable and can be installed by a nonprofessional without any technical background. (3) The first tier is used as “ground truth” that is always available. This allows the system to seamlessly track its performance.

Monitoring the Main Power Meter

The first tier information is the aggregated power consumption of a household. This is obtained by monitoring the main power meter in real-time. Although the traditional




Figure 7: A prototype implementation of ViridiScope consists of several non-intrusive sensors: a main meter monitor, light intensity sensors, magnetometers, and microphones. This implementation uses CrossBow MicaZ motes, HMC1002 magnetic sensors, and MTS310 sensor boards. Each battery-powered sensor node monitors energy consumption related signals, and sends them back to the fusion center. The fusion center is a PC that solves a numerical optimization problem using CVX tools, combines data from the distributed sensors, and profiles appliance-level power consumption according to unique node IDs. It’s important to note that the system doesn’t require in-line sensors.

main power meter may not have an interface to get real-time power consumption, many companies sell easy-to-use devices that can monitor the main power meter in real-time. For example, the Energy Detective (TED) and The Power Cost monitor are readily available in the U.S. [2, 25] (See Figure 8). Furthermore, a better and systematic solution is expected in the very near future. The Automatic Meter Reading Association is developing a system that relays real-time meter information wirelessly for billing purposes. Following the trend of utility companies simplifying their billing systems through wireless technologies, we believe that extracting data from the main power meter of a household in real-time will soon be commonplace as the Smart Grid technology is deployed at a large scale in the US [6] and elsewhere.

Power Monitoring Using Indirect Sensing Subsystem

The second tier sensors monitor a signal related to energy consumption, which is used to infer the power consumption of an appliance. Our prototypical design uses HMC1002 magnetic sensors, CdSe photocell light intensity sensors and simple piezoelectric microphones. Magnetic sensors monitor raw magnetic field changes near power wires and the standard deviation of the magnetic field change is used to estimate the power consumption. Internal power states of appliances is inferred from acoustic and light intensity sensors that monitor noise patterns and light intensity changes near appliances. Our implementation uses Crossbow MicaZ wireless sensor nodes that run TinyOS, HMC1002 magnetic sensors, and MTS310 sensor boards (Figure 7). Note that an instantiation of ViridiScope is not limited to this particular choice.




Figure 8: Real-time power trace of a 2-person apartment: A TED device, installed on the main power line, monitors the real-time power consumption. Although this information gives a good sense of the energy usage at home, it is only giving a partial view of where the energy gets used. For example, in the evening, many appliances are simultaneously on, and it is hard to disambiguate where the energy is actually used.

Information Fusion via Numerical Optimization Toolbox and Data Aggregation

The data collected by the two tiers is sent to a personal computer (fusion center) where the indirect signals and the aggregated power consumption from the main meter are “fused”. The PC fuses the data by solving a numerical optimization problem introduced in the following section. The solution calculates the calibration parameters for the second tier sensors, which subsequently are used to compute the real-time power consumption of each appliance. Our implementation uses the CVX toolbox [9], an open source convex optimization tool, to solve the autonomous calibration problem (Equation 5). Alternative tools, such as MOSEK [22], could be used as well.

Problem Formulation

The ViridiScope hardware provides the total power consumption, $y_0(t)$, and the signals, $s_i(t)$, that are correlated with the energy consumption of the individual appliances. We can now formulate a problem statement that allows the system to automatically learn the appliance-level power consumption. We first model the explicit calibration functions that are mappings from the measured signals to the power consumption ($f_i : s_i(t) \rightarrow p_i(t)$). Then we set up a numerical optimization problem that solves for f_i by exploiting the fact that the total power consumption of a house is the sum of the power consumption of all the appliances.

Calibration Function Modeling

Magnetic Sensors

It is very well known that magnetic fields are coupled to AC currents (Maxwell’s Equation [7]). Many types of magnetic sensors are used for current measurement methods [7, 16]. Power consumption can be calculated through current measurement if the consumption is purely resistive. But the power consumption is usually both resistive and inductive, and so both voltage and current is needed to calculate power consumption. Also, clamp type current sensors are not adequate for our needs because they need to be clamped on




Figure 9: A simple light intensity sensor can detect the on/off status of a light near a window even during sunset, this plot shows a sufficient noise to signal ratio to infer on/off status very reliably.

one of the two AC wires, thus requiring non-trivial power wire modification. Clamping both wires results in the useless measurement of the net current of zero.

As shown in Figure 3, the standard deviation of the magnetic field change has a strong correlation with the power consumption. We can setup a simple calibration curve that is an affine mapping from the magnetic field variation to the power consumption:

$$p_i(t) \stackrel{def}{=} \alpha_i s_i(t) + \beta_i, \quad (1)$$

where $p_i(t)$ is the power consumption, α_i, β_i are calibration parameters, and $s_i(t)$ is the standard deviation of the magnetic field change. In our implementation, a node samples the magnetic field change at 100Hz and calculates the sample standard deviation, $s_i(t)$, over a 1 second sliding window.

Acoustic and Light Sensors

Because many appliances have a limited number of power states, estimating the average power consumption of the internal states is enough. Since internal power states often have an almost constant power consumption (Figure 4), we can infer the power consumption if we know the states and their respective power consumption.

Although this approach compromises time-resolution and does not capture transient power consumption, its resolution is generally good enough for this non-critical domestic application. The main benefit comes in total system cost. For example, a simple light sensor can detect the on/off status of light (Figure 6) and costs much less than a magnetic sensor. For example, a HMC1002 magnetometer is \$20 whereas a simple light sensor is less than \$1. Similarly, a simple acoustic sensor is much cheaper than the magnetic sensor, and can easily identify the internal power states of an appliances (Figure 5).

We now define a simple calibration model. Let $s_{i,j}(t)$ be an indicator function indicating the j -th internal state of the i -th appliance. This value is the output from the indirect sensors detecting the internal power state of appliances. Then the power consumption of an appliance can be described by the following equation

$$p_i(t) \stackrel{def}{=} \sum_{j=1}^{K_i} P_{i,j} s_{i,j}(t), \quad (2)$$

where K_i is the number of internal power states of the i -th appliance, $P_{i,j}$ is the average power consumption of the j -th power state of the i -th appliance, and $s_{i,j}(t)$ is a boolean indicator function that indicates the on/off status of the j -th power state of the i -th appliance.

Uninstrumented Appliances

ViridiScope needs to handle the case where several appliances are not instrumented with any sensors. This is likely in a real house, since a user might have forgotten one or more devices, or may not want to install sensors everywhere [4]. We define an artificial appliance that is always on to account for the uninstrumented appliances, and thus takes over for all the uninstrumented appliances (also called ghost power consumer). We denote the average power consumption of this artificial appliance as P_i . This gives:

$$p_i(t) \stackrel{\text{def}}{=} P_i s_i(t), \quad (3)$$

where P_i is the power consumption for all uninstrumented appliances and s_i is the indicator function that indicates if the system has uninstrumented appliances in the configuration or not.

Direct Meters

Direct appliance power meters are commercially available [11, 17] and it is reasonable to include these meters in the ViridiScope system configuration. Because these meters directly monitor $p_i(t)$, we can directly use that measurement in our formulation. For notational consistency we define:

$$p_i(t) \stackrel{\text{def}}{=} \tilde{p}_i(t). \quad (4)$$

where $\tilde{p}_i(t)$ is the data from a direct meter.

Autonomous Calibration Formulation

Equations 1, 2, and 3 state that if we know all the calibration parameters, we can calculate the individual power consumption of every appliance. Unfortunately, the calibration parameters are not known and must be learned after the sensors are deployed. This could be done by manual intervention, measuring and calibrating each sensor. However, this introduces a significant burden on the person deploying the system, and thus we want to automate this step.

The total power consumption of a household is the sum of the power consumption of all the appliances. Each of Equations 1, 2, 3, and 4 represents the power consumption of an individual appliance. We can formulate an explicit numerical optimization problem because $y_0(t)$ has to be equal to $\sum_{i=1}^N p_i(t)$, where N is the total number of appliances. And thus,

$$\begin{aligned} & \min \|y_0(t) - \sum_{i=1}^N p_i(t)\| \\ \text{where} & \\ p_i(t) = & \begin{cases} \alpha_i s_i(t) + \beta_i : \text{magnetometers} \\ \sum_{j=1}^{K_i} P_{i,j} s_{i,j}(t) : \text{light/acoustic sensors} \\ P_i s_i(t) : \text{uninstrumented} \\ \tilde{p}_i(t) : \text{direct meter input.} \end{cases} \end{aligned} \quad (5)$$




Figure 10: A single acoustic sensor is capturing ambient noise in addition to the sound from the refrigerator compressor. This makes detection of the compressor state difficult.

In this problem, the control variables are the calibration parameters, α_i , β_i , P_i and $P_{i,j}$. Note that $\|\cdot\|$ is an arbitrary norm, making the formulation a convex program. For example, Equation 5 can be easily cast as a linear programming problem given that the norm is l_1 , which can be solved in known polynomial time using very efficient numerical solvers [5]. Or if the norm is l_2 , we can solve it as a least square problem. Equation 5 shows that a completely non-calibrated function known to have a monotonic relation with the phenomena of interest can be calibrated with the help of an additional sensor that provides the aggregated information, y_0 . Thus, once we deploy indirect sensors near appliances of interest the system can automatically calculate calibration parameters used to compute real-time power consumption of the individual appliances. Although not explicit, by letting $s_i(t)$ be an on/off indicator function for an individual appliance the formulation estimates the average power consumption of an appliance over time.

We use the l_1 norm in the ViridiScope system, because it is more robust to outliers than the l_2 norm. It is well known that the l_1 regression, or Least Absolute Value Regression, solves for the median value. The l_2 regression, or Least Mean Square Regression, solves for the average value [5]. The median operator is more robust to outliers than the mean operator.

In addition, a real-time measure of the system performance makes the system more autonomous and adaptive. Equation 5 implies that a performance measure can be the normalized estimation error, $\text{Performance} \stackrel{\text{def}}{=} \left| \frac{\sum_{i=1}^N \hat{p}_i(t) - y_0(t)}{y_0(t)} \right|$, where $\hat{p}_i(t)$ is the estimated power consumption of the i -th appliance. This measure is always available because $y_0(t)$ is always available. By monitoring this value, the system can adapt to performance degradation by re-solving Equation 5.

PRACTICAL CONSIDERATION

Because indirect sensors measure ambient conditions to detect power states, they are prone to external noise sources that also influences the ambient condition. Although l_1 regression is intrinsically robust to the outlier problem, it performs better when the system has a better signal-to-noise ratio. Thus, in a real system, we want to suppress such noise.




Figure 11: A simple thresholding of a statistical signature with an adaptive filter helps to detect states more reliably.

Rejecting Undesired Noise via Adaptive Filtering: Acoustic Case

Figure 6 shows that a light sensor in a dark room deems to have a very good signal to noise ratio. Even if the light sensor is exposed to ambient light, by placing the sensor close to the lamp we can still achieve an excellent signal to noise ratio (Figure 9). Therefore, no sophisticated filters are necessary in many cases.

A simple microphone can detect the compressor’s power state (see Figure 5). However, it also captures other ambient noise sources adjacent to the refrigerator, including cooking noise, human voices, and TV sound (Figure 10). Thus, the raw signal is not very reliable to detect the compressor’s on/off status.

The story changes quickly when using one additional microphone. Noting a high spatial correlation among the sampled signals, a simple adaptive filter can reject the excess noise. For example, if one of the microphones is placed close to the compressor and the other one on the counter top then a simple scale estimator and subtraction of the two sampled signals is enough to clean up the signal (Figure 11).

One possible form of an adaptive filtering for acoustic based detection is,

$$\begin{aligned} s_1 &= s + n && \text{: on the refrigerator} \\ s_2 &= \kappa n && \text{: on the counter top,} \end{aligned} \quad (6)$$

where s is the sound signature from the compressor, n the ambient noise, and κ is a scale factor. When the compressor is off, s is 0 and simple training for κ is possible by solving

$$\begin{aligned} \sigma^2(s_2) &= \kappa^2 \sigma^2(n) = \kappa^2(s_1) \\ \kappa^2 &= \frac{\sigma^2(s_2)}{\sigma^2(s_1)}, \end{aligned} \quad (7)$$

where $\sigma^2(\cdot)$ is the variance operation. Given κ , a simple adaptive filter effectively rejects the ambient noise and thus allows a reliable detection of the compressor state (Figure 11).

EVALUATION

Evaluation Setting

To test and validate the ViridiScope design concept, we conducted several experiments in a 2-person apartment. We chose three cases with increasing complexity (Table 1). In **Case I**, we estimate power consumption of a desktop computer, a table lamp, and a refrigerator using two different

Cases	Participating Appliances
I	PC, refrigerator, table lamp
II	PC, refrigerator, table lamp, alarm clock, router, settop box
III	2 laptops, massage chair, TV, water heater, 2 lights, 2 table lamps

Table 1: Evaluation Scenarios

sensor configurations: (1) A magnetometer on the power cord to the computer, a light sensor near the table lamp, and a light and acoustic sensor to monitor the refrigerator. (2) A magnetometer on the power cord of each of the three appliances. For **Case II**, we add an alarm clock, a wireless router, and a cable TV settop box to the Case I configuration. The goal is to test the system with unmonitored appliances, and thus the power states for these three additional devices was not monitored. In **Case III**, we let the ViridiScope system estimate the average power consumption of 9 different appliances by monitoring only on/off status of the appliances.

For all the experiments, the main power line was monitored with a commercial power meter that measures the real power consumption. We used a Crossbow MicaZ mote connected to this power meter to send the measurements $y_0(t)$ to the fusion center. Ground truth data was collected with a WattsUp Pro [11] connected to every device. These measurements were used to evaluate the accuracy of our estimation.

Case Study I: Three Appliances

Figure 12 illustrates the information processing of the ViridiScope system for the case of one magnetic, one light, and one sound/light combination sensor. As mentioned earlier, there is a small error introduced in the power-state estimation, because it assumes that the power consumption of a device in one power state is constant. Figure 13 depicts this effect. While the error for the magnetic sensor power estimation quickly varies around 0, the error for the refrigerator and the lamp slowly varies because their power consumption patterns are monotonic. We note that the average power consumption is accurate, and the accumulated errors are small (see Figure 14).

When we use only magnetometers, ViridiScope performs very well with estimation accuracy almost the same as in Figure 12. In summary for only the magnetometers, the mean errors(the standard deviations) for the PC, the lamp, and the refrigerator are 1.29%(2.74%), 0.04%(0.05%), and -1.05%(4.24%), respectively. The estimation accuracy for the lamp is particularly good because it consumes almost constant power.

Case Study II: With Uninstrumented Appliances

In the case where uninstrumented appliances are present, we have to introduce an artificial appliance that is always on. The power consumption of this artificial appliance will contain the accumulated power consumption of all the uninstrumented devices. In our test scenario, we added an alarm clock, a WiFi access point, and a cable TV set-top box. The power consumption of the three appliances is about 2.7W,

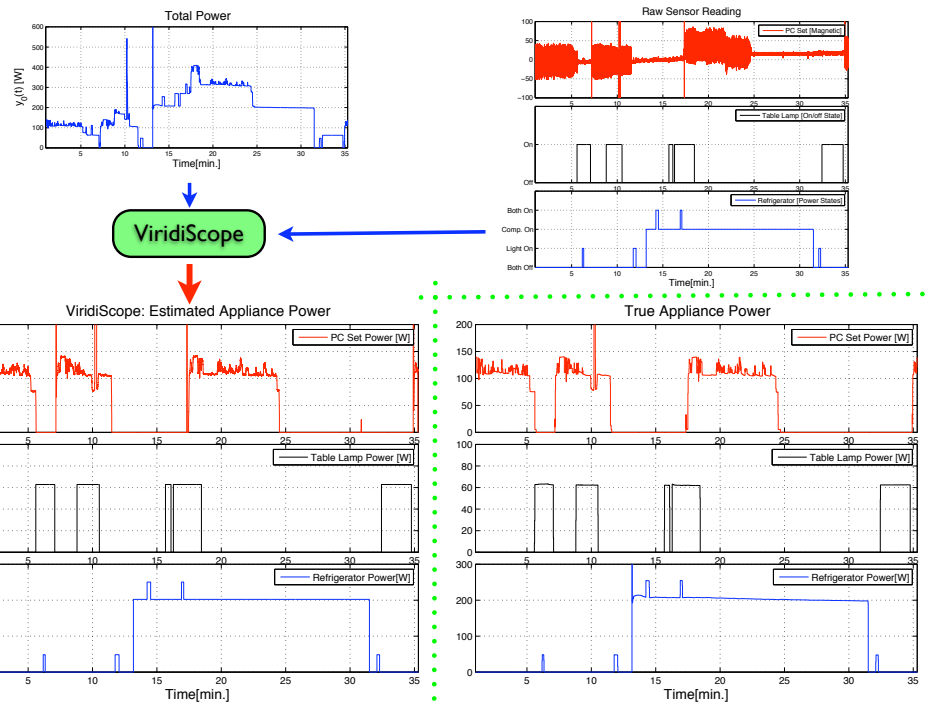


Figure 12: The ViridiScope system takes the total power consumption, magnetic field, and internal power states information from heterogeneous sensors including magnetic, light and acoustic sensors. It then solves the l_1 norm minimization problem (Equation 5) to compute calibration parameters. The calculated calibration parameters are used to estimate appliance-level power consumption. The estimated power consumption(left bottom) tracks the true power consumption(right bottom) very well.




Figure 13: Case I: The estimation error of Figure 12. The estimated system power consumption of each individual appliance has very good accuracy, even when they are simultaneously on.




Figure 14: Case I: Comparison among true accumulated energy consumption of the appliances and their estimation.

6.3W, and 22W respectively. This is a total of 32W. Figure 15 illustrates the estimated, and true power consumption for this scenario. We can see that, except in some very small transition cases, the estimated power consumption is close to the real power consumption. As a matter of fact, ViridiScope estimates the ghost power consumption at 32.1W.

Case Study III: 9 Appliances with On/off detection

If we are only interested in the average power consumption of each appliance, it is enough to know only the on/off status of an appliance. Therefore, we use light or acoustic sensors

RELATED WORK

Resource Monitoring in Homes

Directly metering the consumption of every electrically operated light, device, or appliance requires current and voltage sensing at every end-point or circuit leading to that end-point. Such sensors are costly and often require installation by an electrician. An exception to this are plug-in power meters for appliances that are plugged into wall sockets such as Kill-A-Watt and Watts UP that can be installed in series. Moreover, such technology is limited in its ability to aggregate data from multiple sensors in real-time for fusion, analysis, and visualization. Products based on such technology

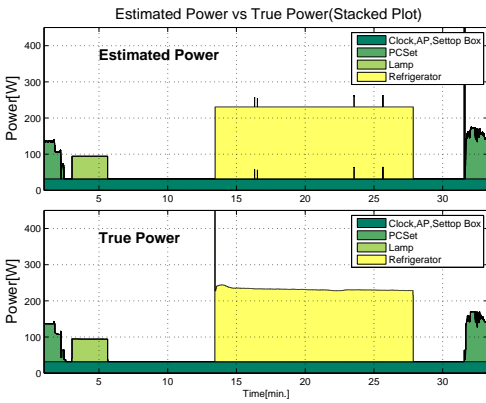


Figure 15: Case II: In the presence of uninstrumented appliances, ViridiScope still succeeds in determining the individual appliance power consumption. On average, the uninstrumented devices (alarm clock, WiFi access point, and set-top box) consume 32W. ViridiScope estimates their power consumption to be 32.1W.

are often designed as data loggers and limited to displaying and storing data on the unit itself, or sending one or two channels of sensor data to remote display units or to a PC via serial, USB, or Ethernet cables [1, 2]. Recently developed wireless sensors that can measure the power consumption of a power outlet have been presented in research setups [17]. Even though they are a promising step toward in ease of use and management, they have the same cost and labor-intensive installation problem as already existing solutions. Another problem is that such in-line monitoring momentarily interrupts services. Even worse, appliances that do not have standard AC connectors or that are hard-wired into the power distribution network (e.g. ceiling lights, water heater, etc.) can not be monitored with such direct sensors [15].

Similar to our approach is the work on non-intrusive appliance load monitoring (NALM) [12, 15]. Instead of a per-outlet sensor, a single central sensor is envisioned to monitor a circuit with multiple electrical loads that operate independently. In the extreme case, this single sensor is situated at the electrical meter for the entire building. Sophisticated statistical signature detection algorithms are used to analyze the current and voltage waveforms to separate out the individual loads, their state transitions, and their energy consumptions. Models are used to describe individual appliances or groups of appliances to assist in the disaggregation of the total load into its constituents. Issues of cost, size due to circuit complexity, and the accuracy of appliance detection and monitoring were identified as problems [12]. In particular, NALM has problems in coping with multi-state appliances that change their power profile over time, as well as certain two-state appliances. The training phase for NALM is quite intrusive and labor intensive in the sense that someone must feed signature, patterns, or similar information to the system so that the system can infer what’s going on based on the signature. A distinctive aspect of ViridiScope is that it makes power monitoring almost autonomous, is able to monitoring variable power consumption, and can monitor power consumption of multiple appliances that are simultaneously on.

Appliance	True Power[W]	Estimate[W](Error)
Light 1	59.8	58.4(-2.32%)
Light 2	31.9	34.2(7.02%)
Massage Chair	46.2	45.8(-1%)
Table Lamp 1	68.1	69.3(1.73%)
Table Lamp 2	14.7	15(2.54%)
Water Heater	1623.6	1620.6(-0.18%)
TV	100.1	101(0.94%)
Laptop 1	38.9	38.3(-1.45%)
Laptop 2	31.4	29.2(-7%)

Table 2: Case III evaluation

Home Infrastructure Monitoring

Several recent studies have shown monitoring infrastructure leads to interesting conclusions of what is going on in a household and provides resource consumption relevant signals. Often times this information can be extracted through simple interfaces and means. For example, Patel et. al. [23] monitor the electrical noise within the power-lines of a house. They exploit the fact that each appliance introduces a unique noise signature. They can infer if an appliance is on or off, by detecting and identifying this signature, but not its actual power consumption. Our approach is complementary to their result, and our algorithm is able to incorporate results from their system to infer the power consumption of an appliance. In a different context, Fogarty et. al. [13] investigated monitoring the plumbing system by using microphones on pipes to infer water activity in a household. Both systems [13, 23] are easy to install but employ complex calibration mechanisms in order to learn the detection patterns. Though these systems capture user behavior, we consider them complementary for our approach because our ViridiScope formulations can incorporate with them to estimate actual consumption numbers with a slight modification.

LIMITATIONS AND FUTURE WORK

In the current implementation, we assume that uninstrumented devices consume constant power. A better model assumes that the ghost power varies. One approach is to opportunistically calibrate the sensors so that they can track the power consumption of the instrumented appliances. Once they are calibrated, it is possible to infer a change in the ghost power consumption by subtracting the sum of the power estimation from the main power consumption. The challenge is to figure out when to learn and when to estimate.

As Beckmann et. al. [4] pointed out, five design principles for end-user sensor deployment challenges need to be considered. While the ViridiScope system design addresses the sensor calibration issue, further researches need to be done in regard to (1) user conceptual models for familiar technologies, (2) balancing installation usability with domestic concerns, (3) detection mechanisms of incorrect installation of sensors, and (4) general end-user education. A possible first step is to exploit the performance metric. Since it is a measure of the normalized error, this information can be useful to users. For example, by observing this metric they can try different types of indirect sensors on appliances, adjust sensor placement, or try to reduce the number of unin-

strumented appliances to achieve better accuracy. Moreover, this information can be used to detect the incorrect installation of sensors and identify faults once a machine learning mechanism is employed.

CONCLUSION

We present ViridiScope, a fine grained power monitoring system for residential spaces using the combination of existing infrastructure and indirect sensors. By introducing indirect sensors, we extend the traditional power monitoring dimension. Indirect sensing introduces a new sensor calibration challenge. By exploiting existing infrastructure, we develop an autonomous sensor calibration scheme that automates the sensor calibration procedure. Experiments show that the system tracks the per-appliance power consumption to less than 10% error. In addition, ViridiScope can easily monitor the power consumption of appliances with multiple simultaneously active appliances as well as variable power consumption.

The challenges of sustainability is an impending issue to the global society. We believe that ViridiScope is a step towards energy efficient homes, because measurement is one of the most critical components. ViridiScope leverages ubiquitous technologies and their networking capabilities to extract valuable information from else meaningless data. It is just one example showing that a shrewd combination of various information sources is the key, and ubiquitous networked sensing and computation devices make this a feasible concept.

ACKNOWLEDGEMENT

The authors like to thank to our shepherd, James Scott, for his guidance in preparation the final version of the paper, Roy Shea for helping revise the paper, and the anonymous reviewers for their helpful comments. This material is based upon work supported by the NSF under award # CNS-0627-084 and CCF-0820061, and by the UCLA Center for Embedded Networked Sensing. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies.

REFERENCES

1. Energy monitoring technologies inc, <http://www.energymonitor.com/>.
2. TED: The Energy Detective, <http://www.theenergydetective.com>.
3. L. Balzano and R. Nowak. Blind calibration of sensor networks. In *ACM IPSN 2009*.
4. C. Beckmann, S. Consolve, and A. LaMarca. Some Assembly Required: Supporting End-User Sensor Installation in Domestic Ubiquitous Computing Environments, In *UbiComp 2004*.
5. S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
6. A. Burger. Meter, meter on the wall: Giving & taking from a smarter grid. *triplepundit*, 2007.
7. D. K. Cheng. *Field and Wave Electromagnetics*. Prentice Hall, 1989.
8. M. Chetty, D. Tran, and R. E. Grinter. Getting to green: understanding resource consumption in the home. In *UbiComp 2008*.
9. J. Dahl and L. Vandenberghe. CVXOPT. <http://abel.ee.ucla.edu/cvxopt>.
10. S. Darby. The effectiveness of feedback on energy consumption: A review for DEFRA of the literature on metering, billing and direct displays. Tech. report, Environmental Change Institute, University of Oxford, 2006.
11. Watts Up? Pro. <http://www.wattsupmeters.com/>.
12. EPRI. Nonintrusive appliance load monitoring system (NIALMS): Beta-test results. Tech. report, 1997.
13. J. Fogarty, C. Au, and S. E. Hudson. Sensing from the basement: A feasibility study of unobtrusive and low-cost home activity recognition. In *ACM UIST 2006*.
14. J. Froelich, K. Everitt, J. Fogarty, S. Patel, and J. Landay. Sensing opportunities for personalized feedback technology to reduce consumption. In *the CHI workshop on Defining the Role of HCI in the Challenge of Sustainability*, 2009.
15. G. Hart. Nonintrusive appliance load monitoring. In *Proc. of the IEEE*, 80(12):1870–1891, Dec 1992.
16. Honeywell. Hall effect sensing and application. Tech. report.
17. X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and implementation of a high-fidelity AC metering network. In *ACM IPSN/SPOTS 2009*.
18. Y. Kim, T. Schmid, Z. M. Charbiwala, J. Friedman, and M. B. Srivastava. NAWMS: Nonintrusive autonomous water monitoring system. In *ACM SenSys 2008*.
19. K. D. Lee. *Electric load information system based on non-intrusive power monitoring*. PhD thesis, MIT, 2003.
20. S. B. Leeb, S. R. Shaw, and J. R. Kirtley Jr. Transient event detection in spectral envelope estimates for nonintrusive load monitoring. *IEEE Trans. on Power Delivery*, 10(3):1200–1210, Jul 1995.
21. A. H. McMakin, E. L. Malone, and R. E. Lundgren. Motivating residents to conserve energy without financial incentives. *Environment and Behavior Journal*, 2002.
22. MOSEK Inc. <http://www.mosek.com>.
23. S. N. Patel, T. Robertson, J. A. Kientz, M. S. Reynolds, and G. D. Abowd. At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. In *UbiComp 2007*.
24. P. C. Stern. Information, incentives, and proenvironmental consumer behavior. *Journal of Consumer Policy*, 1999.
25. The Blue Line Company. The power cost monitor.
26. K. Whitehouse and D. Culler. Calibration as parameter estimation in sensor networks. In *WSNA 2002*.
27. World Business Council for Sustainable Development. Energy efficiency in buildings: Business realities and opportunities. Tech. report, 2008.