

Minimizing Intrusiveness in Home Energy Measurement

David Lachut
Department of Computer Science
and Electrical Engineering
University of Maryland, Baltimore
County

dlachut@csee.umbc.edu

Lazeeb Choudhury, Kevin
Moran, Simon Piel, Yucheng
Xiong
Department of Computer Science
University of San Francisco

{lfchoudhury, ksmoran, spiel,
yxiong3}@usfca.edu

Nilanjan Banerjee
Department of Computer Science
and Electrical Engineering
University of Maryland, Baltimore
County

nilanb@csee.umbc.edu

Sami Rollins
Department of Computer Science
University of San Francisco

srollins@cs.usfca.edu

Abstract

The expanding deployment of renewable energy sources as well as the widespread deployment of smart meters enables and encourages demand management in homes. Like smart meters, most solar or other renewable deployments allow homeowners to carefully monitor energy supply and past energy consumption, however, using this information to drive demand management is still a manual process. The overarching goal of our work is to automate the process of adapting energy demand to meet supply, which requires a comprehensive understanding of home energy use. Though home energy measurement systems exist, they are often intrusive—requiring several physical components and using often limited resources including energy and bandwidth. In this work, we present the design of a system for comprehensive home energy measurement and analyze the resource requirements of the basic system. Using data collected from six deployments, including one in an off-grid home, we then present two techniques for reducing the resource requirements of the system. Our techniques reduce the energy footprint of the system as well as the amount of physical infrastructure required, making adoption of the system more attractive, particularly to those who live in homes powered by renewable energy sources.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

General Terms

Human Factors, Measurement

Keywords

energy, sustainability

1 Introduction

The expanding deployment of renewable energy sources as well as the widespread deployment of smart meters enables and encourages demand management in homes. Like smart meters, most solar or other renewable deployments allow homeowners to carefully monitor energy supply and past energy consumption. Using this information to drive demand management, however, is still a manual process; it is up to the consumer to determine how to reduce demand to meet the restrictions of the supply. Particularly in the case of renewables, traditional demand management techniques, such as reducing usage of high-power appliances between the hours of 7PM and 7AM, is insufficient [1]. One of our research subjects lives in an off-grid, solar-powered home and reports that it is an exceptionally cumbersome process to manually monitor the state of his home's energy generation and consumption and aggressively reduce energy usage when supply is low. We propose that automated and adaptive demand management systems will be the key to broader implementation of renewables.

The overarching goal of our work is to automate the process of adapting energy demand to meet supply. Though there are many products on the market that enable measurement of energy usage, either at the scale of the entire home or individual appliances, there is no comprehensive solution that provides a view of both how much energy a home consumes and how the energy consumption is broken down by appliance. The challenges of designing such a system include ensuring that the energy requirements of the system itself are minimal and that the physical components of the system are not intrusive. In a home powered by renewable sources, where residents eschew using toasters because of their energy consumption, a measurement system with a large energy footprint is unlikely to be adopted. Moreover, in any home a system comprised of several clunky devices attached to appliances all over the home is intrusive.

In this work, we explore the design, implementation, and

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.

Buildsys'12, November 6, 2012, Toronto, ON, Canada.
Copyright © 2012 ACM 978-1-4503-1170-0 ...\$10.00

deployment of a system that balances thorough data collection with minimally-intrusive system resource requirements and components. We have built our system using off-the-shelf components that measure both whole-home energy consumption and energy consumption of several individual devices in the home. The system has been deployed in six houses to date, including one off-grid home and one grid-tied home with renewables. The main contributions of this work are as follows:

- **System Design and Resource Analysis** - We overview the design of our system and present a thorough analysis of its resource usage, including its energy and bandwidth requirements.
- **Minimization of Energy Consumption** - Using data we have collected from our deployments, we present a technique to reduce the energy consumption of the system by minimizing the data collection rate, hence enabling component duty cycling.
- **Minimal Appliance Set Determination** - Using an Additive Factorial Hidden Markov Model [2], we show that the minimal set of appliances that must be monitored should include appliances that are sporadically used, however, appliances with regular usage patterns can be identified from aggregate data.

2 Motivation

The design of our measurement infrastructure is guided by the results of a survey that seeks to understand energy monitoring and usage in *green homes*—homes fully or partially powered by renewable energy sources. Though our system is designed for any home, green homes represent the extreme end of the spectrum. Energy supplied is limited and demand must be carefully managed. A system designed for this environment will also suit the needs of a grid-only user. We published the survey online and, through a press release, invited users of green homes to participate. We have, so far, received eight responses from California, Arkansas, Nebraska, Hawaii, Ontario (Canada), and Massachusetts, including four off-grid, and four grid-tied users. Two participants also donated data they have collected manually in their homes. There are two salient conclusions that can be drawn from the user survey.

There is a need for automated or semi-automated energy monitoring, visualization, and control systems. Survey participants emphasized both the importance and difficulty of tracking home energy usage, and all expressed interest in a real time monitoring, visualization, and control system on a smartphone device. One of the survey participants stated *“I’ve kept a home energy spreadsheet for many years. The spreadsheet allows us to tell where the energy is going and how much we are using. This simple spreadsheet has been the most useful energy saving device.”*, while another participant commented: *“Tracking all sources and uses of energy is a challenge.”* Though there are smart home automation systems on the market, anecdotally we have identified cost as one factor limiting their adoption. Additionally, we infer that energy consumption of the system is another limiting factor. One participant noted: *“I own a color laser printer*

from before I went off grid. This also a massive energy hog, so I hardly use it.” Limiting system energy consumption is a key design goal of our system.

There is a need for automated or semi-automated demand management systems. The households in our survey often resorted to manual demand management such as using vacuum cleaners during sunny days, or getting rid of appliances like toasters or electric heating. One participant claimed *“No longer operate a separate freezer as system will not support it. Have a Steca 12v fridge/freezer. Either but not both ? Would like a freezer ?”*. Moreover, another participant who uses a diesel generator as an auxiliary power source claimed that he has to depend on the generator 20% of the time. We conjecture that better demand management in terms of prediction and optimization of energy generation and consumption can lead to more comfortable living.

3 System Design and Evaluation

Based on analysis of our survey results, we have designed a system that provides comprehensive measurement, visualization, and real time appliance control. In this section we overview the basic system and performance, and in Sections 4 and 5 we present optimizations that reduce the intrusiveness of the system by minimizing system-wide energy consumption and reducing physical infrastructure required.

3.1 System Architecture

The basic measurement system consists of three key components as illustrated in Figure 1.

Home Components: Our client-side measurement system has several components. We use off-the-shelf Z-Wave [3] energy meters to track the energy consumption of several devices in each home. Our main client-side software is written using HomeOS [4] and runs on a low-power, Atom-based FitPC platform [5]. It uses WiFi to poll a dual-radio gateway (the Vera2 [6]) that, in turn, polls the individual meters to collect power readings using the Z-Wave protocol. Additionally, two of our deployments are in homes with solar panels and we use specialized components to collect data on energy generation by the panels and consumption by the home. We also use Z-Wave clamp meters to collect whole-home energy consumption data in two other deployments. We have chosen to use plug-n-play meters that are easy to install. The meters also act as switches that can be used to turn on and off appliances. Finally, our client-side software supports lamp dimmers and thermostats, and we plan to deploy them in the future. While custom hardware monitors can be used, we have opted for the Z-wave based system because of its popularity and wide adoption in the home automation domain.

Backend Server: We have implemented a standard servlet-based web service backed by a MySQL database. The server provides an extensive RESTful API that allows access to power draw and energy consumption data for each home. The API provides secure access to the data collected. We are currently extending the server to support additional analytics as well as to implement machine learning techniques for deriving suggested demand management actions based on energy generation and consumption patterns.

Smartphone Application and Web Interface: We have

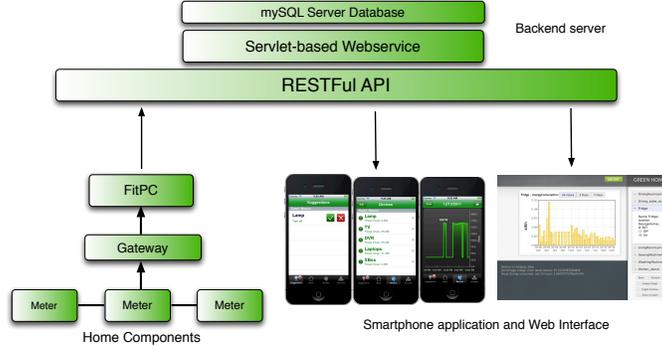


Figure 1. Overall architecture of our monitoring and visualization system.

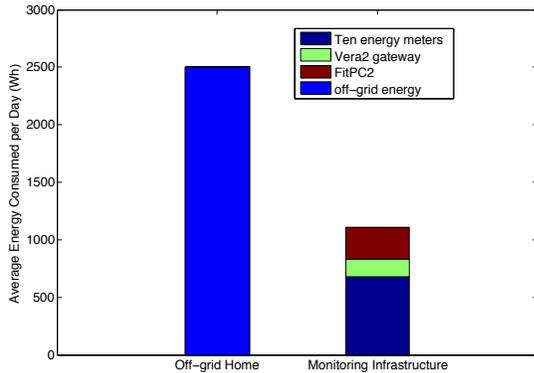


Figure 2. The figure illustrates the energy overhead of our monitoring system compared to the per day energy consumption of an off-grid home.

used the web API to develop a robust smart phone application. Users may access their energy consumption data through this application. The application provides a detailed view of power draw and energy consumption of individual appliances and, where available, the entire home energy consumption and battery levels. In addition, the application provides a control mechanism that allows users to switch devices on and off through the energy meters. We are extending the application with thermostat and dimmer controls, as well as the ability to collect contextual information on user activities, location, and schedules. This additional information will be used to develop a context-driven demand management system. We have developed a web interface that acts as a visualization and control portal for users for energy auditing, tracking energy consumption, and understanding energy bottlenecks and hotspots in homes.

3.2 Deployment and Benchmarks

Our system has been deployed in six homes in California and Arkansas, including one off-grid solar home and one grid-tied solar home. We measure between 5 and 10 devices in each home, including refrigerators, televisions, lighting, clothes washers, and computers and other electronics, and data is reported every 30 seconds. In the solar homes, we collect data on the energy generated by the solar infrastructure, the energy consumed by the entire home, and residual battery capacity. For two of the other houses in the study, we

Home	Type	Days of Data	Percentage Downtime
1	off-grid	92	9.4%
2	grid-tied	37	37.1%
3	grid-only	287	3.8%
4	grid-only	267	19.1%
5	grid-only	256	19.4%
6	grid-only	77	12.8%

Table 1. The table shows the number of days of data collected and system downtime. In addition, two survey participants have donated 686 and 382 days of energy generation and energy consumption data.

collect whole-home energy consumption using a clamp meter attached at the fuse box for the home, and for one home we have access to whole-home energy consumption data collected by a smart meter. Table 1 shows the duration of data collected from the six deployments and the percentage of days that the system did not report data. Our deployments are ongoing and we will continue to collect data.

System Robustness: To evaluate the robustness of the system, we examine the percentage of time for which the FitPC at a given home fails to report data for *all* meters in the home. We note that we have experienced server failures, but because the client-side software caches data during server outages we do not have gaps in the data. We also note that it is rare for the FitPC to report data for some meters and not others because the software that runs on the Vera gateway returns the latest power reading for all meters, even if some have not been recently polled. We have deployed a workaround for this problem, but do not yet have sufficient data to report results on individual meter failures. We have observed, however, that there are a few cases, particularly in the off-grid home, where meters are in series with a wall switch. The user uses the wall switch to control the measured device, for example a lamp, and so the meter does not report data when the device is off.

Table 1 shows that the total data collection downtime varies from 9.4%–37%. In Homes #2 and #6, a large fraction of the failures occurred when the FitPC was unplugged accidentally. Moreover, a fraction of the 19.4% failures in Home #5 occurred due to an automatic system reboot (Windows update installation), when the homeowner was traveling. Though the software components of the system have been robust, these results demonstrate that end-user hardware main-

tenance accounts for a large portion of failures. Addressing hardware failures, unfortunately, requires involvement of the end users, however, we believe that as we deploy additional system functionality users will monitor the system more frequently, for example to view suggested demand management actions, and will have more incentive to make sure it is functional.

System Overhead: To understand system resource usage, we benchmark two metrics: energy consumption and bandwidth usage. Because our goal is to intelligently manage home energy, understanding the energy requirements of the system itself is paramount. Figure 2 shows the energy consumption of the Z-Wave energy meters, the Vera2 gateway, and the FitPC when sampling appliances at a uniform rate of once every second. We use a 70 MHz Tektronix Oscilloscope to perform the measurements. The absolute power consumption of the entire system is close to 46 W for a deployment of 10 meters. This amounts to more than 1,140 Wh of energy consumption per day. This is small overhead in a grid-only home, however, in the off-grid home in our deployment, where the per-day power consumption is less than 2.5 KWh, this is a more than 44% overhead. We also measured the bandwidth of the system by sniffing packets on a home router. We found that for a single meter, uploading data to the server every second results in an overhead of 44 Kbits/second, which increases tenfold for our deployments with 10 meters. This is a substantial overhead for low-speed DSL or 3G connections characteristic in off-grid deployments.

Discussion: The system described above is designed to collect as much information as possible for each home. Our initial research subjects have been very generous in allowing us to deploy as many components as necessary, and have been tolerant of the energy and bandwidth usage of the system. We believe, however, that to make the system practical for broad deployment, especially in off-grid homes, it is necessary to reduce the intrusiveness of the system in terms of both physical infrastructure and resource usage. In Section 4 we examine the feasibility of reducing the amount of data collected from the system such that the home system components may be duty cycled to reduce energy usage. In Section 5 we explore techniques for determining the minimal set of appliances that must be measured to derive the necessary information from the system.

4 Minimizing Energy Usage for Measuring Devices

The aggregated energy consumption of the always-on system we describe in Section 3.1 becomes intrusive in an energy-limited home. Because each individual component of the system is, in fact, quite power efficient, we propose that system-wide duty cycling is necessary to reduce its overall energy footprint. Ideally, the meters, dual-radio gateway, and FitPC running the client-side software will remain in a low-power state when the power draw of the measured devices remains constant. When a device changes power state, for example a lamp is switched on, the system will wake to report the state change to the central server. Implementing system-wide duty cycling requires knowing when to wake

the system to sample the power draw of the measured devices. In this section, we explore the feasibility of using a Fast Fourier Transform-based approach to deriving sampling intervals for the devices measured in our deployments.

Algorithm and Sampling Intervals: To derive the sampling interval we compute the Fast Fourier Transform (FFT) to identify the maximum frequency of the signal given by the raw power consumption readings. For each device, we first construct a FFT and determine the *band-gap*. The band-gap is the difference between 0 and the maximum frequency in the frequency spectrum of the device. The sampling frequency is taken as the band-gap*2, per Nyquist’s theorem. Nyquist’s theorem, however, applies to noiseless signals. Appliance power data has several elements of noise due to the measurement circuit, EMI from other appliances, and other factors, hence, the frequency spectrum contains several frequency elements with low but non-zero amplitude. To filter this noise, we only consider frequencies with amplitude between the maximum amplitude and two orders of magnitude below it. For instance, if the maximum power observed is 1000, we consider all frequencies with power values between 10 and 1000. Any frequency with power value below 10 is excluded. The derived sampling frequency determines how often the meters must be sampled, suggesting when the gateway and data collection components must be duty cycled.

Device	Sampling Interval (Minutes)
Refrigerator	0, 3, 5, 7
Lamp	15, 17, 25, 59, 60, 140, 432, 605, 1080
TV/Media Center	2, 14, 30, 45, 127, 144, 180
Laptop	20, 85
Kitchen Island	945
Kitchen Sink	1080
Computer	47, 630
Electronics	432
DVR	20
Monitor	1
Oven	1
Sewing Machine	1
Washing Machine	1, 1, 3
Microwave	1, 1, 2
XBox	0

Table 2. Sampling intervals derived using basic algorithm.

To evaluate the FFT algorithm, we use four weeks of data collected from 5 homes between June 7, 2012 and July 5, 2012. We use the first two weeks as our training set and the second two weeks as our test set. We construct the FFT over the training set and derive the sampling intervals shown in Table 2. The left column describes the device and the right column shows the sampling interval. Some devices are measured in several of the homes in our study (e.g., refrigerators), and some homes measure several of the same device (e.g., lamps). Most of the device names accurately describe the connected device, though there are a couple worthy of further explanation. The Media Center is one meter measuring the TV, DVD, Roku media streamer, and PlayStation 2 console. One of the Laptop devices is actually a meter connected to a power strip that powers the chargers for two laptops and a printer. Similarly, the Electronics device is a

power strip that powers a modem, access point, low-power PC, and a printer. The Kitchen Sink and Kitchen Island correspond to lamps in an off-grid home, and the Sewing Machine measures a computer-aided sewing machine.

We draw several conclusions from the raw sampling intervals derived:

- The profile of a given device class (e.g., lamp) varies across homes as well as across different instances of the same type of device within a home. This is not surprising; some users may watch more TV than others, for example. This confirms that the approach of calculating a sampling interval per appliance rather than using a common interval for the home is warranted.
- For some devices, the period used for training was insufficient. The sampling interval of 2 for a TV caused us to delve further into the data to discover that the device was simply not used enough during the training period for our algorithm to produce meaningful results. This is a concern for devices that do not have a regular use pattern, and could even be a result of a user going on vacation during the period we decide to use to train the system. This points to the need to modify our algorithm such that it will retrain itself if a device’s usage changes significantly.
- This basic approach is inappropriate for some devices including the monitor, oven, sewing machine, washing machine, microwave, and XBox. The sampling intervals for these devices are very small, yet most of the devices are used infrequently. The washing machines, for example, were used on no more than 3 of 14 days analyzed. The sampling interval correctly tells us that when the device is in use the power state is likely to change very frequently, however sampling at a rapid rate of every 2 minutes would certainly be wasteful if the device is only used once per week.

Device	Hours Per Day Sampled
Oven	6
Washer	12, 12, 24
Microwave	6, 18, 24
Sewing Machine	18
Monitor	24

Table 3. Hours per day a device must be sampled according to the enhanced algorithm.

To accommodate the devices that have infrequent use patterns, we modify the basic algorithm to capture time of day characteristics. Our goal is to derive 4 sampling intervals for each device—one for the period from midnight to 6AM, one for the period from 6AM to noon, and so on. If a device is not used during a particular time period, it need not be sampled during that period. We separate each of the 14 days of data in our training set into 4 blocks of 6 hours. We then create a signal that is the aggregate of the blocks for the same time frame. In other words, we create one signal comprised of midnight to 6AM of day 1 followed by midnight to 6AM of day 2 and so on. Next we calculate the FFT over each of the 4 individual signals and derive the sampling interval

as described above. The algorithm produces four sampling intervals for each device. The oven, for example, has a sampling interval of once per minute between 6PM and midnight and is not sampled at any other time during the day.

We apply the enhanced algorithm to the devices with original sampling intervals below 3 minutes, excluding the anomalous refrigerator and TV devices. Table 3 illustrates the number of hours per day each device is sampled. The enhanced algorithm reduces the number of samples per day over the basic algorithm, in some cases as significantly as 1,080. Though one microwave and one washer are sampled 24 hours per day, the microwave is only sampled once per hour from midnight to noon and the washer is only sampled once per hour from 6AM to noon. This demonstrates that, in many cases, time of day can be effectively used to determine when a device needs to be measured.

Metrics: To evaluate the FFT-based algorithm we want to understand whether our sampling interval is too frequent, or not frequent enough. If it is too frequent, then we will sample to find that there has been no power state change. If it is not frequent enough then we will miss power state changes that happen between samples. In other words, if the interval is too small then we waste energy by sampling when a device remains on or off, while if the interval is too large then we may miss a device turning on and then off again. To evaluate our algorithms, we adapt the metric of *recall*. Recall tells us how many of the power state changes that appear in the raw data are captured in our sampled trace. If the difference in power consumption of a device between sample s_1 and sample s_2 differs by more than 5 W—a configurable parameter—we record a power state change. We calculate the number of power state changes in the raw data set (with samples every 30 seconds) and in the sampled data set. Recall is the number of state changes in the sampled data set divided by the number of state changes in the raw data set.

Results: In Figures 3 and 4 we illustrate recall using the basic sampling intervals and the enhanced sampling intervals, respectively. We note that we exclude 6 devices from one home from our results because the home was unoccupied during the test period. We also exclude several devices that were not used or were used rarely during the test period. The bars show the mean recall for all devices in a particular device class (e.g., lamps) and the confidence intervals show the minimum and maximum values.

We make several observations from these results. Recall is high for many of the devices in our data set. Most of the lamps perform similarly as do most of the fridges. This is not surprising as one would expect the usage of these devices to be similar across homes and users. For the fridges, high recall is likely the result of a very short sampling interval—frequent sampling ensures that most state changes are captured. For the lamps, this is likely the result of limited and regular usage each day, which is accurately captured by the FFT algorithm. Some of the devices exhibiting low recall have few state changes in the data set. Other devices exhibiting low recall suffer from the opposite problem. The computers are always on and, in one case, power draw fluctuates between 38–48 W every few minutes. The TV exhibiting the

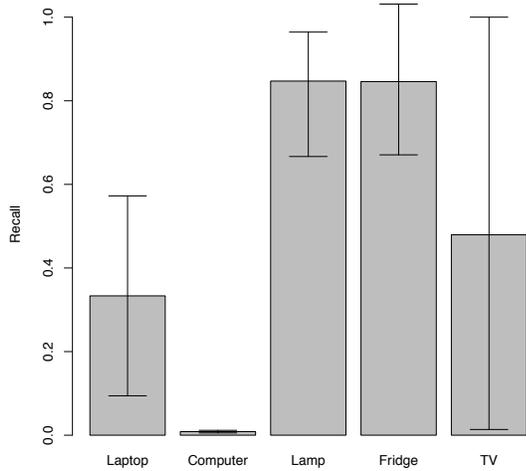


Figure 3. The figure evaluates the basic sampling algorithm for various appliances.

poorest performance shows similar behavior. Recall that we use 5 W as the cutoff to determine that a device has experienced a power state change. We are currently investigating whether using the power states observed from a particular device to determine this cutoff is a better choice, particularly for devices like computers.

The performance of the enhanced algorithm is similar for all devices, suggesting that they all exhibit time of day characteristics. As noted above, the key benefit of the enhanced algorithm is that it reduces the number of samples per day over the basic algorithm in some cases up to 1,080. Further investigation of the data, however, reveals that in some cases the enhanced algorithm misses power state changes because there is a mismatch between the hours the devices are used during the training period and the hours the devices are used during the testing period. This points to the need to improve our training approach.

Discussion: Our results suggest that using a FFT-based algorithm to derive sampling intervals that will enable system-wide duty cycling while accurately capturing device power state changes is feasible in most cases. We note several possible enhancements that will likely improve performance of many of the low-performing devices, including using the observed power states of the devices in the algorithm, and improving the training procedure. Time of day is a useful contextual cue that can be used to improve performance, and we are exploring other context-based enhancements including location of the user (tracked by our mobile phone app). We have also observed that a single sampling interval is inappropriate for some devices, for instance a lamp that is on for 2 hours and then off for 20 hours should not be sampled every 2 hours, nor every 20. We are working on a strategy for addressing this issue. Finally, our goal with this analysis was to perform an initial feasibility study to determine the accuracy of this approach. We are working to implement an

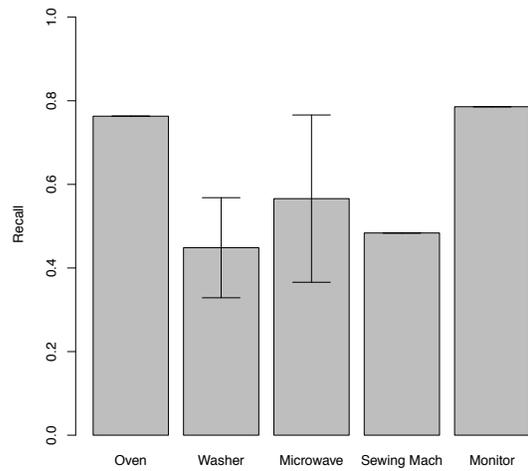


Figure 4. The figure evaluates the enhanced sampling algorithm for various appliances.

integrated algorithm for duty cycling the system based on the sampling intervals derived for all devices in a home.

5 Monitored Appliance Set Determination

Device	TPT, TNT
Refrigerator (off-grid)	79.7%, 80.1%
Sewing Machine (off-grid)	69.5%, 74.7%
Washer (off-grid)	70.8%, 78.7%
Microwave (grid-only)	2.1%, 1.8%
Refrigerator (grid-only)	59.7%, 60.7%
Washer (grid-only)	18.7%, 69.9%
Oven (grid-only)	19.2%, 70.2%

Table 4.

The table illustrates the TPT and TNT values calculated for appliances using AFHMM [2].

Accurate demand management requires usage profiles for individual appliances [7], however reducing system intrusiveness requires minimizing the number of components that comprise a measurement system. In this section, we explore the feasibility of using disaggregation techniques to reverse engineer device usage data, thus reducing the number of energy meters an end user must install. We argue that there are two necessary pieces of the system: (1) a meter to collect whole-home, aggregated energy usage, and (2) meters attached to any device a user wants to control remotely. We have found remote control of devices, such as lighting and televisions, via our mobile app is a key feature of our system, and to enable remote control we must deploy a switch that can also act as an energy meter. Our goal, however, is to minimize measurement of devices such as refrigerators, which are never directly controlled by the user, and clothes washers, microwaves, and ovens, which cannot be effectively controlled remotely.

Algorithm: For our analysis we use an unsupervised disaggregation technique called Additive Factorial Hidden

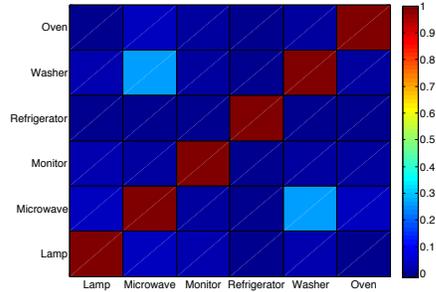


Figure 5. Pearson's coefficient for pairs of appliances.

Markov Model (AFHMM) [2]. The model has been shown to accurately reverse engineer appliance usage from aggregated power consumption data. Each appliance is modeled as a HMM (Hidden Markov Model) with an ON and OFF state. We determine the transition probabilities between the states using two weeks of data from individual appliances. The initial state occupancy distribution is assumed to be uniform over the two states. We perform a disaggregation analysis on appliances that are not directly or remotely controlled by the user as described above. Our aggregated data is the sum of the power consumption (collected every 30 seconds) from these individual appliances. We apply the AFHMM model with Exact MAP Inference [2] to predict individual appliance usage from the aggregated data.

Metrics: To evaluate the efficacy of the algorithm, we use two metrics (1) True Positive Transitions (TPT) = $TP/(TP+FP)$ where TP is the number of correct OFF to ON transitions and FP is the number of incorrect OFF to ON transitions inferred by the algorithm; and (2) True Negative Transitions (TNT) = $TN/(TN+FN)$ where TN is the number of correct ON to OFF transitions and FN is the number of incorrect ON to OFF transitions predicted by the algorithm. If the transitions in the predicted usage occurs within 2 minutes of the actual transition, we assume that it is a correct transition. We are interested in transitions since the usage of an appliance is determined by when the user turns the appliance on and off.

Results: Table 4 shows the TPT and TNT values for appliances in an off-grid and a grid-only home. The technique performs well for all three devices in the off-grid home, however in the grid-only home the refrigerator is the only device that is consistently identified. A closer look at the data reveals that the poor-performing devices in the grid-only home are used sporadically and their power signatures are confused with the power signatures of other appliances. In the off-grid home, however, use of devices is more periodic.

We next explore whether the algorithm may be enhanced by using additional contextual features. We consider correlation of usage across devices and time of day the device is used. Figure 5 plots the Pearson's coefficient for pairs of appliances in the grid-only home. Unlike previous work [8], we find that the correlation between the usage of appliances we seek to identify is low, thus we argue that it is not a reliable feature for this analysis. Figure 6, however, shows the fraction of time a microwave is ON at a given time interval

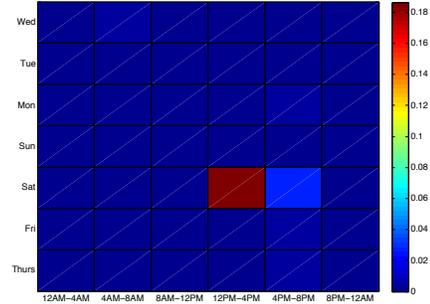


Figure 6. Times and days when the microwave was used.

during the week. The microwave does show higher usage at certain times and days, and as future work we are exploring ways to augment the AFHMM model with time of the day and week information.

Discussion: Our preliminary analysis shows that disaggregation techniques can be effectively used to identify energy consumption patterns of devices that are used regularly or periodically. In some homes, a large number of devices have a regular usage pattern, thus a measurement system need only have few physical components to provide a comprehensive view of home energy consumption. For homes that have devices with more inconsistent usage patterns, we are currently working to extend our technique to include contextual features such as time of day.

6 Related Work

Our research builds on previous work on home energy deployment studies, non intrusive load monitoring, and demand side energy management.

Home Energy Deployment Studies: There are several ongoing research efforts on deploying energy meters in a home or commercial building setting [9, 10, 11, 12]. Some deployments use a single sensor to collect energy consumption data on the total home energy consumption [13, 14, 15], while others collect data from individual appliances and circuits [16]. Several studies have tried to collect data on human occupancy and context such as ambient light and humidity [17]. While most deployments are performed over a short time period, there are some recent attempts to collect longitudinal data on energy consumption for disaggregation analysis and home side demand management systems [9]. In contrast to existing deployments, our deployment is unique as it involves three different types of homes from a renewable energy perspective — off-grid (completely driven by renewables), grid-tied (with renewables), and grid-only (completely dependent on the grid). Our goal is to understand the similarities and differences between these homes to develop better demand management systems that can work across the spectrum of homes.

Energy Disaggregation: Energy disaggregation from a non intrusive load monitoring perspective is a well studied area. Most of the focus in this area is on using a single energy meter (a clamp meter, for instance) to measure the energy consumption of an entire home, and then use machine learning and signal processing techniques to reverse engineer appliance usage profiles in the house [18]. There are

also attempts to use auxiliary techniques such as powerline noise [18] to detect which appliances are in use. Most studies rely on supervised machine learning, that uses labeled training data on appliance profiles to infer appliance usage from the aggregate energy consumption [19, 7]. Recent attempts also use generic profiles, and unsupervised learning techniques [8] that does not require any apriori training, to reverse engineer appliance usage. Additionally, using context from other utilities like water or gas consumption has been used to improve the accuracy of the learning algorithms [20]. Our goal is complimentary to the unsupervised learning techniques. We use a single sensor that collects aggregate energy consumption of the entire home, energy readings from appliances that users want to control, and unsupervised learning techniques to determine the *minimal* set of appliances that should be monitored, so that accurate usage profiles for all appliances can be inferred.

Demand Response Systems: The overarching goal of our project is to develop better demand management systems for a wide spectrum of homes (with or without renewables). We plan to build on previous work in the grid-only space [21, 22, 23], such as optimization frameworks that can flatten peak energy consumption [13] and minimize total energy consumption of a home. Our demand management system would be complementary to existing systems in the off-grid [24, 1, 25] and grid-tied space [26].

7 Future Work

This paper presents the design and deployment of a home energy measurement infrastructure and explores two techniques for minimizing the intrusiveness of the system. We have deployed the system in six homes, including one off-grid home and one grid-tied home. We are currently working to integrate and deploy our energy minimization and disaggregation techniques. We then plan to extend the system to provide automated and adaptive demand management suggestions to end users, enabling more efficient energy usage, particularly in green homes.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. 1115798, 1055061, 1018112. Any opinions, findings, and conclusions are those of the authors and do not necessarily reflect the views of the NSF.

8 References

- [1] Nilanjan Banerjee, Sami Rollins, and Kevin Moran. Automating Energy Management in Green Homes. In *ACM Sigcomm Workshop on Home Networks*, August 2011.
- [2] J. Zico Kolter and Tommi Jaakkola. Approximate Inference in Additive Factorial HMMs with Application to Energy Disaggregation. In *Journal of Machine Learning Research*, 2012.
- [3] <http://www.aeon-labs.com/site/products/view/5/>. Aeon labs smart energy switch.
- [4] Colin Dixon, Ratul Mahajan, Sharad Agarwal, AJ Brush, Bongshin Lee, Stefan Saroiu, and Paramvir Bahl. An operating system for the home. In *NSDI 2012*.
- [5] <http://www.fit-pc.com/web/>. Fitpc 2.
- [6] <http://micasaverde.com/vera.php>. Micasaverde vera 2.
- [7] J. Froehlich, E. Larson, S. Gupta, G. Cohn, M. Reynolds, and S. Patel. Disaggregated end-use energy sensing for the smart grid. *Pervasive Computing, IEEE*, 10(1):28–39, jan.-march 2011.
- [8] Hyungsul Kim, Manish Marwah, Martin F. Arlitt, Geoff Lyon, and Jiawei Han. Unsupervised disaggregation of low frequency power measurements. In *SDM*, pages 747–758. SIAM / Omnipress, 2011.
- [9] J. Zico Kolter and Matthew J. Johnson. A public data set for energy disaggregation research. In *SustKDD workshop on Data Mining Applications in Sustainability*, August 2011.
- [10] Maria Kazandjieva, Brandon Heller, Omprakash Gnawali, Philip Levis, and Christos Kozyrakis. Green Enterprise Computing Data: Assumptions and Realities. In *Third International Green Computing Conference*, 2012.
- [11] Yuvraj Agarwal, Thomas Weng, and Rajesh Gupta. Understanding the role of buildings in a smart microgrid. In *In Proceedings of the ACM/EDAC Conference on Design Automation and Test in Europe*, March 2011.
- [12] Timothy W. Hnat, Vijay Srinivasan, Jiakang Lu, Tamim I. Sookoor, Raymond Dawson, John Stankovic, and Kamin Whitehouse. The hitchhiker’s guide to successful residential sensing deployments. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, SenSys ’11, pages 232–245, New York, NY, USA, 2011. ACM.
- [13] Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeanne Albrecht. SmartCap: Flattening Peak Electricity Demand in Smart Homes. In *IEEE International Conference on Pervasive Computing and Communications*, March, 2012.
- [14] Andres Molina-Markham, Prashant Shenoy, Kevin Fu, Emmanuel Cecchet, and David Irwin. Private Memoirs of a Smart Meter. In *2nd ACM Workshop On Embedded Sensing Systems For Energy-Efficiency In Buildings*, November, 2010.
- [15] Christopher Laughman, Douglas Lee, Robert Cox, Steven Shaw, Steven Leeb, Les Norford, and Peter Armstrong. Advanced non-intrusive monitoring of electric loads. *IEEE Power and Energy Magazine*, pages 56–63, March/April 2003.
- [16] Alan Marchiori and Qi Han. Using circuit-level power measurements in household energy management systems. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys ’09, pages 7–12, New York, NY, USA, 2009. ACM.
- [17] Thomas Schmid, David Culler, and Prabal Dutta. Meter any wire, anywhere by virtualizing the voltage channel. In *BuildSys*, 2010.
- [18] Shwetak N. Patel, Sidhant Gupta, and Matthew S. Reynolds. The design and evaluation of an end-user-deployable, whole house, contactless power consumption sensor. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI ’10, pages 2471–2480, New York, NY, USA, 2010. ACM.
- [19] Chao Chen and Diane Cook. Energy Outlier Detection in Smart Environments. In *AAAI*, 2011.
- [20] Vijay Srinivasan, John Stankovic, and Kamin Whitehouse. WaterSense: Water Flow Disaggregation using Motion Sensors. In *BuildSys*, 2011.
- [21] S Firth, K Lomas, A Wright, and R Wall. Identifying trends in the use of domestic appliances from household electricity consumption measurements. *Energy & Buildings*, Jan 2008.
- [22] Thomas Weng, Bharathan Balaji, Seemanta Dutta, Rajesh Gupta, and Yuvraj Agarwal. Managing plug-loads for demand response within buildings. In *In Proceedings of the ACM Workshop on Embedded Sensing Systems For Energy-Efficiency*, Seattle, October 2011.
- [23] Younghun Kim, Thomas Schmid, Mani B. Srivastava, and Yan Wang. Challenges in resource monitoring for residential spaces. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys ’09, pages 1–6, New York, NY, USA, 2009. ACM.
- [24] H. Hu. *Risk-Conscious Design of Off-Grid Solar Energy Houses*. PhD dissertation, The Georgia Institute of Technology, Department of Architecture, November 2009.
- [25] Tanuja Bapat, Neha Sengupta, Sunil K. Ghai, Vijay Arya, Yedendra B. Shrinivasan, and Deva Seetharam. User-sensitive scheduling of home appliances. In *GreenNets*, 2011.
- [26] Ting Zhu, Aditya Mishra, David Irwin, Navin Sharma, Prashant Shenoy, and Don Towsley. The case for efficient renewable energy management for smart homes. In *BuildSys*, 2011.