UCLA
UCLA Electronic Theses and Dissertations

Title
Integrated, Interactive Sensing for Scalable Behavior Guidance: Health and Energy

Permalink
https://escholarship.org/uc/item/6jr1331d

Author
Chen, Victor

Publication Date
2015

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

University of California
Los Angeles

Integrated, Interactive Sensing for Scalable Behavior Guidance: Health and Energy

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Electrical Engineering

by

Victor Liu Chen

2015

© Copyright by

Victor Liu Chen

2015

Abstract of the Dissertation

Integrated, Interactive Sensing for Scalable Behavior Guidance: Health and Energy

by

Victor Liu Chen

Doctor of Philosophy in Electrical Engineering

University of California, Los Angeles, 2015

Professor William J. Kaiser, Chair

Deployments of sensors for real-world applications face critical challenges that can make large-scale adoption difficult or impossible to achieve due to prohibitive cost. Much of this cost can be reduced by enabling the adoption of more cost-effective behaviors by users and by guiding users on proper usage of technology solutions. Proper guidance for both subject behaviors and sensor usage training can reduce barriers to adoption by improving compliance and providing guarantees.

The aim of this research is to design and develop tractable and scalable monitoring system architectures for large-scale, real world usage that can guide users to adopt more beneficial behaviors. To achieve this aim, system designs were developed in direct support of monitoring applications to address problems in healthcare and energy conservation.

For energy conservation, the research used low-cost energy monitoring as a building block to test the effectiveness of various types of information feedback on real consumers in controlled experiments. In healthcare, we addressed critical unmet needs in post-operative patient care using interactive mobile technology solutions, integrated with

physiological sensing to implement sensor usage guidance. Scalability was enabled by the development of hardware and software systems that allow subjects to easily perform data collection with guarantees on reliability and validity. To provide the complete system, end-to-end software and hardware architectures were developed to combine the core research with the other necessary components and field validations have been conducted.

The dissertation of Victor Liu Chen is approved.

Magali A. Delmas

Chi On Chui

Gregory J. Pottie

William J. Kaiser, Committee Chair

University of California, Los Angeles

2015

This work is dedicated to my family and friends who supported my endeavors and helped me persevere.

TABLE OF CONTENTS

List of Figures

List of Tables

ACKNOWLEDGMENTS

First and foremost, I want to express my tremendous gratitude to my advisor, Dr. William J. Kaiser. I am fortunate to have worked with him for a great many years now and I cannot overstate how much of an impact he has had on my life as a teacher and mentor. He has been a positive role model for me not only professionally as a researcher and engineer but also as a kind and thoughtful person.

I would like to thank the members of my dissertation committee, Dr. Chi On Chui, Dr. Magali A. Delmas, and Dr. Gregory J. Pottie for their insight and advice. I thank Dr. Delmas for her expert guidance and support on the Engage studies in developing such unique and fascinating research. I would also like to thank Dr. Anne Lin for her guidance in developing PRIME.

I feel very fortunate to have worked with so many talented researchers from the Wireless Health Institute, the Institute of the Environment and Susintability, and elsewhere including Mahdi Ashktorab, Yeung Lam, Neil Lessem, Stephen Locke, and Eric Yuen. I have benefited greatly from their advice, encouragment, and more importantly, their friendship.

Finally, I could not have achieved anything with the love and support of my family and friends and I thank them for everything they have given me.

VITA

2005	B.S., Electrical Engineering, University of Californa, Los Angeles
2008	M.S., Electrical Engineering, University of Californa, Los Angeles

PUBLICATIONS

Aaron Dawes, Sarah Reardon, **Victor Chen**, William Kaiser, Marcia Russell, Clifford Ko, Anne Lin. Wireless Technology to Track Surgical Patients After Discharge: A Pilot Study. *American Surgeon*, 2015

Victor Chen, Magali Delmas, William Kaiser, Stephen Locke. What Can We Learn From High-Frequency Appliance-Level Energy Metering? Results From a Field Experiment. *Energy Policy*, 2015

Victor Chen, Magali Delmas, William Kaiser. Real-Time, Appliance-Level Electricity Use Feedback System: How to Engage Users? *Energy and Buildings*, 2014

CHAPTER 1

Introduction

1.1 Behavior-Technology Gap

Within the last decade, society has seen tremendous advances in a number of critical technology areas. These areas include networking and information technology, mobile computing platforms, sensing, backend computing and storage, open source software and software services, and algorithms and analytics. Many of these tools are available as commodoties; cloud computing and cloud storage have become reliable and cost-effective; mobile devices and broadband internet are affordable and accessible; data processing and system management methods and tools are widely available and often free. The following statistics further highlight the current technology state:

- Networking is now more pervasive than ever and a reported 70% of US homes now have access to high-speed internet, an increase by 40% from 10 years ago [ZS13].
- As of October 2014, mobile devices account for 55% of internet usage in the US, of which 47% came from mobile app usage marking the first time app usage has exceed PC usage [Smi15]. Remarkably, modern smart phones also possess 1000s of times more computer power than the computers used in the Apollo space missions.
- Storage now costs pennies per gigabyte and the primary costs for backend comput-

ing systems have shifted toward improving performance.

- There is now a remarkable diversity of high-quality open-source software and free software services. Its reported that 98% of businesses use open-source software and this amounts to savings of $60 billion in annual cost [The08].

These advancements have significantly reduced the barriers for application development. Researchers and developers can work more quickly and achieve more than was previously possible even with limited resources. Taken together, all of these technologies provide a foundation for new applications focused on human behavior guidance.

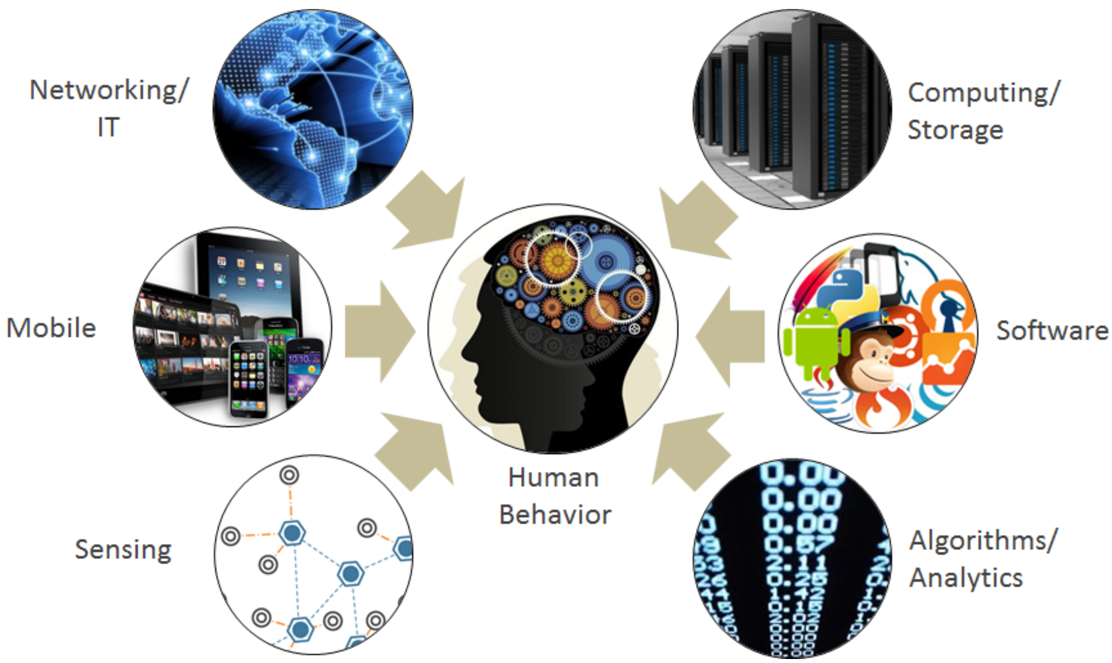

Human behavior research is not a new study and has traditionally been the domain of psychology and economics. However, these disciplines have focused on modeling human behavior without fully leveraging the capabilities of more recent advances in technology. At the other end of the spectrum, engineering has been focused on developing technology solutions without thoroughly considering the human response to technologies and the large-scale impacts across diverse populations and environments. Thus, there exists a gap between technology development and behavior science. This gap provides an opportunity for developing more effective solutions for a broad range of sociotechnical challenges where large-scale impact is a critical measure for success.

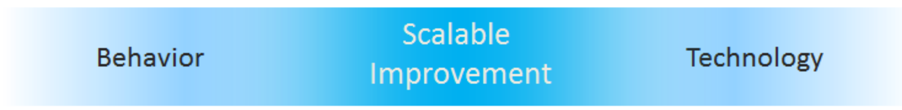

The aim of this research is to address critical unmet needs in deployment of monitoring technology in real applications in energy and healthcare. We focus on the intersection of technology and behavior and provide solutions based on understandings of system limitations and technology adoption challenges where system deployment can lead to higher cost relative to deployment scale. In these cases, data loss, corruption, or non-compliance due to improper use can result in failure of the mission and this is prohibitive to achieving large scale. The approaches undertaken entails development of novel end-to-end systems, validated with deployments in the field. In the process, new scientific understanding about user behavior was discovered with important policy implications.

1.2 Energy User Behavior Monitoring and Guidance

Global warming is considered to be one of the most significant problems our society is now facing and is estimated to carry a cost of \$1.9 trillion annually by 2100. Increased greenhouse gas emissions, particularly anthropogenic CO2 emissions, largely determine the magnitude of the climate change effects. Electricity generation accounts for over 40% of the carbon dioxide emitted by the United States with residential and commercial buildings collectively accounting for over two-thirds of total U.S. energy consumption. This is not surprising considering that residents of the United States spend more than 90 percent of their lives in buildings. Recent studies estimate that behavioral changes can reduce residential energy consumption significantly. Delmas et al. reviewed 156 studies and found an average 7.4 percent reduction in energy consumption, with the largest reductions resulting from individualized feedback. However, many of the studies in the literature suffer from methodological limitations including small sample sizes, short time periods, or low-granularity feedback (i.e. providing only total usage for the building versus per-appliance usage).

These challenges are addressed with both large-scale deployments of retrofit energy monitoring technology. This deployments are conducted in support of behavioral science objectives to understand the effectiveness of information feedback mechanisms for promoting conservation behavior. Load disaggergation is also presented, using computational methods for high-granularity inference by leveraging existing deployments of smart meter technology.

1.3 Wireless Health Monitoring and Guidance

Post-operative care is a critical problem in healthcare. Studies have shown that 30% of the cost of surgical care is incurred during the post-operative care phase. Complications arise during the post-discharge phase after patients have left the hospital requiring patient readmission. Overall, 12% of Medicare patients require readmission within the 30-day time frame following a procedure, costing hospitals $280 million in readmission penalties by the Centers for Medicare and Medicaid Services. There is a challenge in balancing the cost of hospital stay and the cost of readmission. Studies have also shown that the majority of readmissions occur within the first two weeks after discharge from the hospital marking this as a critical recovery period.

Currently, mechanisms and protocols to prevent readmissions involve traditional modes of self-care by providing patients and families with care instructions. Remote care is provided by means of periodic live assessment at home from visiting nurses. However, this type of episodic observation may not adequately detect complications based on temporal trends. Further, visiting nurses Recently, a limited number of mobile applications have been developed in support of specific health conditions such as inflammatory bowel disease but these are not comprehensive in scope and do not address the end-to-end requirements for large-scale adoption.

Despite the lack of current solutions, many of these readmission cases result from preventable complications having well-defined issues and symptoms that support earlier detection. These issues and symptoms may be assessed from dimensions of patient health status including wound care, pain management, recognition of complication symptoms, activity guidelines, personal care, fatigue, sleep quality, and cognition status. While self-assessment does guarantee reliable assessment versus the visiting nurse protocol, it can both complement the available data for specialists as well as augment the capabilities of the visiting nurses.

1.4 Summary of Contributions

The novel contributions of this dissertation are as follows:

- Rapid-retrofit, appliance-level energy monitoring system for shared-space building environments where electrical infrastructure precludes measurement for targeted reporting of energy consumption.

- Holistic system design methodology for large-scale sensor deployment to support assessment of behavioral effects resulting from information feedback strategies.
- Appliance load-disaggregation algorithms leveraging measurement capabilities of large-scale, currently-deployed smart meter technology.
- End-to-end remote patient monitoring and feedback system adaptable to diverse healthcare requirements.
- Sensor usage guidance framework to support proper use of wearable devices.
- Extensive validation of all systems and algorithms through simulations and large-scale deployments.

1.5 Organization

The rest of the chapters are organized as follows: Chapter 2, introduces the Engage energy monitoring and feedback system and describes a pilot study conducted to test information feedback strategies on energy usage behavior. Chapter 3, extends the scope of the Engage system to study a more typical user population living in rental apartments and expands on the study of information feedback. Chapters 2 and 3 provide the foundation for energy load disaggregation in Chapter 4, which addresses the need to achieve return on investment of large-scale remote monitoring. Chapter 5 demonstrates broader application of the end-to-end monitoring and guidance system design to a healthcare application for remote patient monitoring.

CHAPTER 2

Rapid-Retrofit, High-Granularity Energy Monitoring and Information Feedback

This chapter presents a pilot study in residential energy use behavior. The study was conducted over the course of one academic year in several undergraduate residence halls at UCLA. A custom energy monitoring system was designed for instrumenting rooms to measure disaggregated energy consumption. Information feedback was provided in the form of a personalized web dashboard for each room as well as public status. This chapter describes the experimental site, the system design, and results from the experiment.

2.1 Introduction

Electricity generation accounts for over 40 percent of the carbon dioxide emitted by the United States with residential and commercial buildings collectively accounting for over two-thirds of total U.S. energy consumption [U. a, U. b]. This is not surprising considering that residents of the United States spend more than 90 percent of their lives in buildings [EM98]. Recent studies estimate that behavioral changes can reduce residential energy consumption significantly [LEM09, GS08]. Delmas et al. reviewed 156 studies and found an average 7.4 percent reduction in energy consumption, with the largest reductions resulting from individualized feedback [DFA13]. However, many of the studies in the literature suffer from methodological limitations including small sample sizes [GT11, USST06], short time periods [PSJ+07], or low-granularity feedback (i.e. providing only total usage for the building versus per-appliance usage) [All11, BCO+10, Wol11].

While scholars argue that high-granularity information can facilitate energy reduction [EMDL+10, Fis08, NRB09], technological challenges make it unlikely to achieve large-scale deployment in the near future. The main constraint to obtain higher-granularity energy usage information in existing buildings is the infrastructure, which does not allow for direct measurement of distinct loads such as lighting or heating, ventilation and air conditioning (HVAC) for separate rooms. Current building-level meters do not provide high-resolution, high-granularity, real-time information at the room level, even with newer smart meters. This is because circuits in electrical panels in large buildings are not likely to be dedicated to room-specific, individual appliances. While plug-level and appliance-level approaches can provide energy use for specific appliances, they do not provide comprehensive monitoring. For example, plug-level devices, such as the Kill A Watt [Kil] and ACme [JDHDC09], cannot measure energy consumption from built-in appliances like recessed lighting. Appliance-level sensors measure indirect energy emissions (including light, sound, vibration, or electromagnetic fields) to determine appliance state but face scalability challenges [KSCS09, PRK+07]. Appliance load disaggregation methods which aim to provide appliance-level information from aggregate energy measurements are promising but their overall effectiveness is not proven, especially in large-building infrastructures [ZR11].

To achieve high-granularity, appliance-level feedback at the room level, an end-to-end system architecture was developed that included low-cost, wired appliance-level sensors and wireless plug-level meters, a remote gateway for local processing and data upload, and a backend for data storage, data processing, and web services. Residents' energy usage was monitored over a 7-month period, providing high-resolution and high-granularity individualized information. Dashboard engagement was also observed using website analytics. While several studies have assessed the effectiveness of various design components of feedback information [JTP12, Kar11, PTS10] but none have used website analytics data to identify dashboard access patterns. This study was therefore the first to test the effectiveness of and end-to-end feedback system on consumer engagement and energy consumption.

The energy monitoring systems were installed in 66 rooms in three high-rise residence halls on the UCLA campus over the course of one academic year. The population consisted of 102 undergraduate students living in single-, double-, and triple-occupancy rooms. These buildings were constructed in 2005 and 2006 as part of a single construction phase with only minor variation in design. It was possible then to isolate differences in energy consumption due to information feedback rather than infrastructure. Furthermore, because students do not pay for electricity, they are an ideal population to study behavior responses to various forms of information feedback. This setting provided an opportunity to test the effectiveness of information in the absence of an inherent financial incentive to conserve electricity.

The remainder of the chapter is organized as follows. First, the system design and each of the components is described. The signal processing for the energy meters is also described in detail since it significantly improves the capabilities of the reference design on which it is based. Finally, results for user web dashboard as well as both the analytics data and the behavior experiment are discussed.

2.2 Engage Residence Halls System Architecture

The objective of the study was to test the effect of access to personalized, appliance level, real-time and historical energy use information on energy consumption. Achieving high granularity measurement required an end-to-end system architecture, shown in Figure 2.1, that included low-cost sensors, a remote gateway for local processing and data upload, and a back-end for data storage, data post-processing, and web hosting. Three load categories were identified that were controllable by residents and also practical to measure: (1) the appliance plug load fromthe electrical outlets, (2) the overhead lighting, and (3) the heating, ventilation, and air conditioning (HVAC) system. This allowed participants to learn about the contribution of different types of appliances to overall energy consumption and adjust accordingly. For example, survey results revealed that residents consistently overestimated the share of energy from lighting use as being nearly equal to heating and cooling use whereas results from the Engage system indicated that heating and cooling comprised 72% of energy versus 5% for lighting [ADDdB10].

It was important to develop a solution that balanced many factors including cost, development time, reliability, accuracy, and deployment setting. The system also needed to be rapidly deploy-able, given the short time span allotted for installation by the administration between the end of summer occupancy and start offall occupancy. Further, budget constraints required a low-cost solution in order to reach a deployment scale and popula-tion sample size that would yield statistically significant behavioralanalysis. The total hardware cost per installation was under $200.

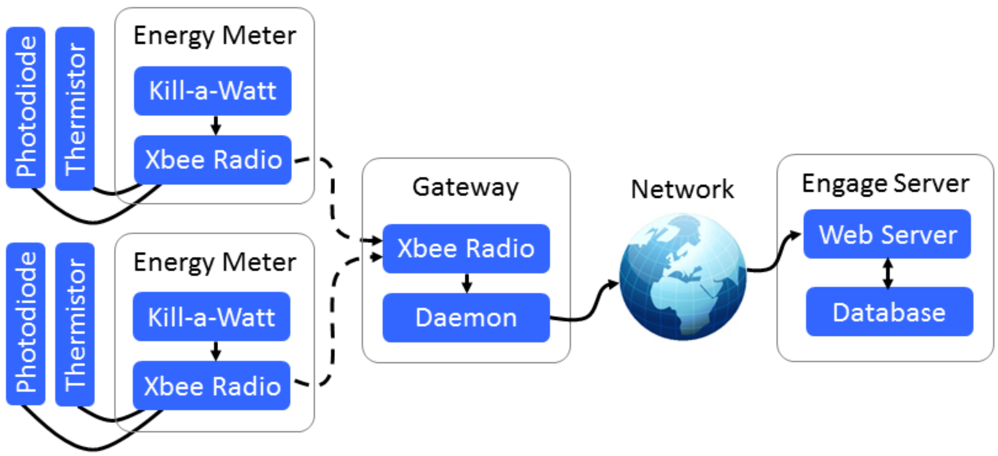

Figure 2.1: End-to-end data flow architecture. Power consumption or state signal is acquired by energy meters and sensors, respectively, and aggregated at the wireless gateway, then uploaded through the network to the server for storage and processing.

2.2.1 Hardware Installation Kit

The deployment hardware shown in Figure2.2 consisted of four components designed to collect information about electricity consumption and appliance state and transmit it to the database via the building's wired network: the energy meter, light and temperature proxy sensors, and gateway. Following is a description of each component.

2.2.1.1 Energy Meter

The energy meter is a modified Kill A Watt which allows measuremeant of the electrical plug load, interface with proxy sensors, and transmit measurements to the gateway. The augmentation is inspired by a popular open source modification called Tweet A Watt [22] which integrates a wireless XBee radio [23] into the Kill A Watt to enable the device to "tweet" energy usage data to the online microblogging service Twitter, via an internet-

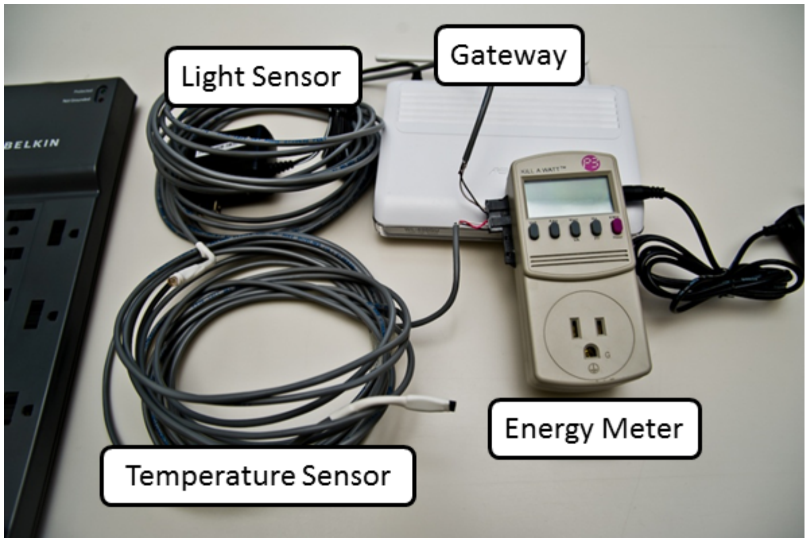

Figure 2.2: Engage energy monitoring system components including energy meter, proxy sensors, and gateway. A standard kit consists of (2) meters, (2) light sensors, (2) temperature sensors, and (1) wireless gateway, as well as requisite power and network cables and power strips.

connected computer acting as a gateway. The Tweet A Watt design leverages two of the six analog input channels on the XBee radio module to measure the current and voltage signals on the Kill A Watt. The Engage design makes use of the additional analog input channels on the XBee radio to interface with the proxy sensors.

Based on the location of electrical outlets, arrangement of appliances and convenience to the residents, two energy meters were installed in each room, typically on opposite walls. All other electrical outlets which were not connected to the meters would be covered over with security tape to discourage use. All electrical devices were plugged into the energy meters via the powerstrips. This way, the energy meters measured the total energy consumed by all the electrical devices (i.e. computer, microfridge, phone charger, TV, game consoles, etc.). To calculate total plug load, energy measurements from both meters were simply added.

2.2.1.2 Proxy Sensors

To allay cost and work with infrastructure constraints, proxy sensors were developed using photodiodes and linear active thermistors. The photodiodes were used to infer light state, which was then used to estimate to light energy consumption. Similarly, the thermistors were used to determine HVAC state and energy consumption. The component materials cost on the order of pennies to a dollar per sensor, compared to direct measurement sensors using current transducers which would have additionally required infrastructure modification. Cost was further reduced by designing the sensors as cable strands rather than wireless sensor platforms. While wireless sensors could have expedited the installation process, the hardware cost of a wireless platform with a limited budget would have been prohibitive to the deployment scale required for demonstrating behavioral significance.

An example floor plan and sensor installation is shown in Figure2.3. To the extent possible, the cables were installed along the room edges in order to reduce the potential for accidental tampering and also to be minimally conspicuous.

Depending on whether the room was designed for single- or double- and triple-occupancy, each room contained either one or two overhead lights, respectively, to which light sensors were affixed. The amount of ambient light on the sensor was minimized by either inserting the sensor into the light fixture or placing it against the fixture and covering it over with tape.

It is important to consider the variability in maximum ambient light levels across different housing units that result from a room's floor level and orientation relative to the path of the sun as well as external occlusions. However, testing across rooms revealed that photodiode output from the overhead lights was significantly higher than from ambient light. Thus light state was able to reliably determined using a common threshold. Figure2.4

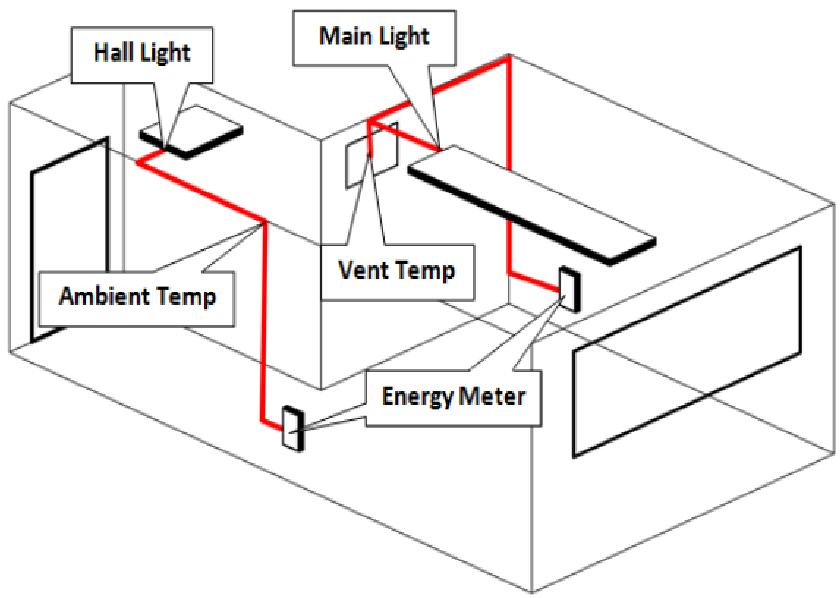

Figure 2.3: Representative physical installation of Engage energy monitoring system. The wireless gateway is not shown but would typically be colocated with one of the energy meters since the room's ethernet port for accessing the building network would always be found next to a wall outlet.

shows an example of the measured light intensity (thin black line) and the threshold (thick grey line) over the course of a day along with the light state estimation. Although ambient light from the sun can be detected by the sensor, light intensity from the target source is sufficiently greater as to be easily discernible. Lighting power is given by

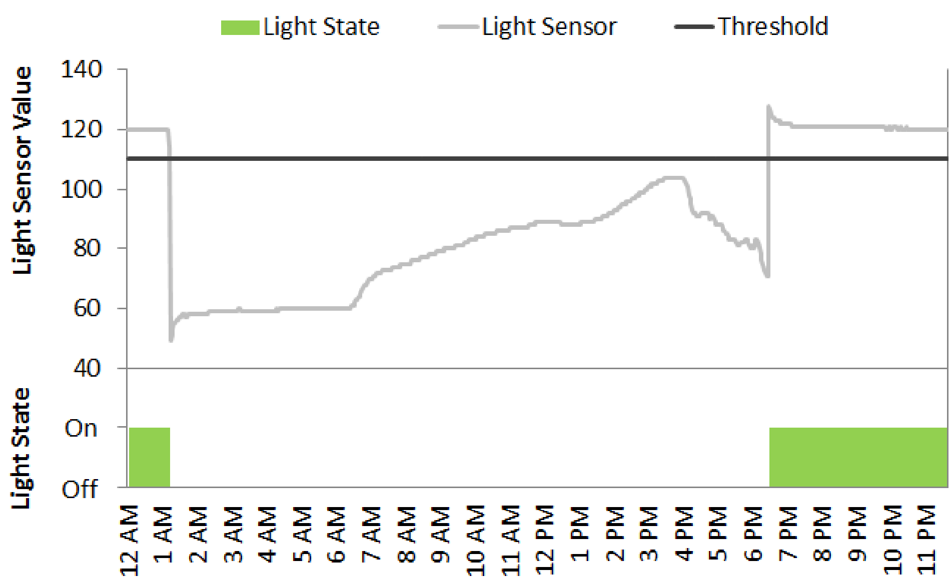

Figure 2.4: Light sensor signal and state classification over a day.

$$P_{light}(t) = \begin{cases} \rho_{light} & \text{if } L(t) \geq \lambda \\ 0 & \text{if } L(t) < \lambda \end{cases} \tag{2.1}$$

where ρ_{light} denotes the rated power consumption of the overhead lights (64W for the main light and 26W for the hall light), $L(t)$ denotes the light intensity signal, and λ denotes the light intensity threshold.

For temperature control, two adjacent rooms share an HVAC unit with dampers prioritycontrolled by individual thermostats based on their set points. To classify HVAC state, one temperature sensor was placed at the opening of the vent to measure exhaust air temperature and another at a distance to measure ambient air temperature. The differential temperature of the sensors was used to infer the HVAC state. It's important to note that floor level and room orientation can significantly affect temperature. For example, normal ambient temperatures would be higher for higher floors and for rooms facing the sun. To account for this variation, a dynamic temperature threshold used, defined to be the average differential over the last 100 minutes. Heating or cooling was classified as "On" when the absolute value of the current temperature differential exceeded the threshold.

With guidance from the housing administration and facilities, it was learned that a conservative figure for per room HVAC energy usage would be 2,000 watts. HVAC power consumption was calculated as follows:

$$P_{hvac}(t) = \begin{cases} \rho_{hvac} & \text{if } T_{diff}(t) \geq \kappa_\tau(t) \\ 0 & \text{if } T_{diff}(t) < \kappa_\tau(t) \end{cases} \tag{2.2}$$

$$T_{diff}(t) = \left| T_{vent}(t) - T_{amb}(t) \right| \tag{2.3}$$

$$\kappa_\tau(t) = \frac{1}{1-\tau} \int_\tau^t T_{diff}(t) dt \tag{2.4}$$

where ρ_{hvac} is the rated HVAC power consumption (estimated to be 2000 W), $T_{diff}(t)$ is the temperature differential between the ambient and vent sensors, and $\kappa_\tau(t)$ is the HVAC threshold constant calculated as the average differential over the last τ units of time (chosen to be 100 minutes).

Figure2.5 shows the absolute temperature differential and dynamic threshold temperature signals for one room over a day along with the HVAC state estimation. The ”On” state indicated by the algorithm corresponded well to the spikes in differential temperature resulting from heating cycles.

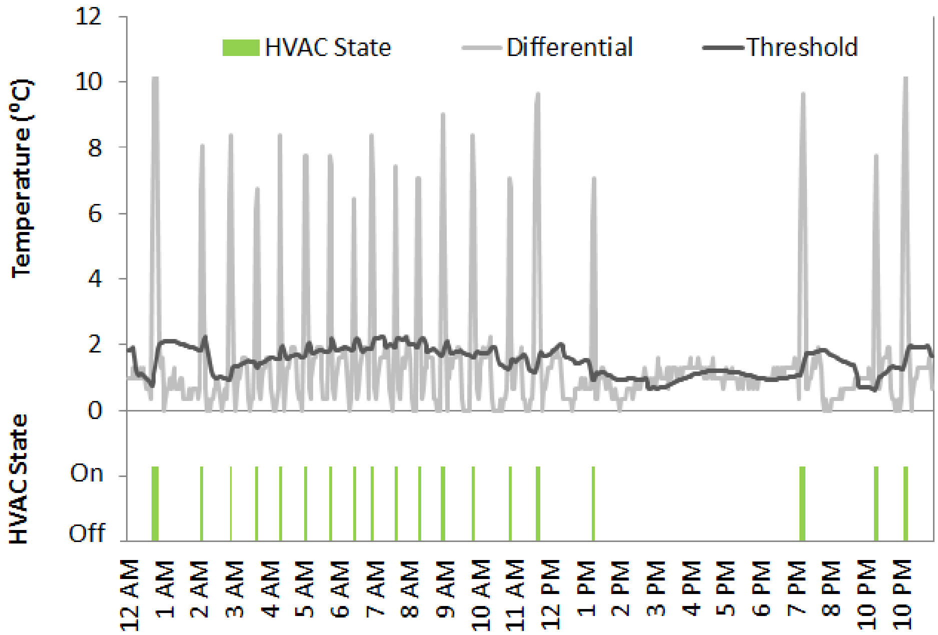

Figure 2.5: Temperature differential T_{diff}, dynamic threshold κ_{τ}, and HVAC state classification over a day.

2.2.1.3 Wireless Gateway

To reduce gateway development time and cost a commercial wireless router (Asus WL-520gU) was used. To make give it the capabilities required for the experiment, the router was flashed it with OpenWRT, an open source operating system optimized for wireless router devices. This enabled other critical software modules to be loaded onto the router.

The selection of this router was important for the following reasons: 1) the router supports firmware restoration, 2) OpenWRT supports the Broadcom BCM5354 processor used on the Asus router, 3) the router includes a USB port for adding USB flash drive storage, and 4) an RS-232 serial port.

Once OpenWRT is installed, the router can be loaded with additional open-source software packages including the Python programming language, OpenVPN virtual private networking to enable remote access, network time protocol (NTP) daemon, and USB flash filesystem support. The Asus router includes only 4MB of onboard flash memory so the USB port is necessary to mount a flash drive to expand the storage capacity. The serial port is used to interface with the XBee radio.

2.2.2 Signal Processing

This section describes the methods used for sampling the current and voltage and for processing the signals to obtain power factor estimates. This processing enables estimation of real power consumption, as is done for electric utility billing. The terms used are summarized in 2.1.

2.2.2.1 Voltage and Current Sampling

The Tweet A Watt design uses a 2-second cyclic sleep mode for the XBee radio's sampling and transmit functionality. In this mode, the XBee wakes from sleep every two seconds to sample the ADC and transmit before returning to sleep. This 2-second delay was chosen to conserve energy since the XBee subsystem was powered by a large capacitor charged by the Kill A Watt's internal power supply; the Kill A Watts internal power supply is current limited and cannot provide the 50mA burst needed for wireless transmission. Forcing a 2-

Table 2.1: Experiment design for Engage Pilot study.

V	Measured voltage (digital value)
$\overline{V}$	Average measured voltage (digital value)
$\widehat{V}$	Estimated actual voltage
I	Measured current (digital value)
$\overline{I}$	Average measured current (digital value)
$\widehat{I}$	Estimated actual current
α_V	Voltage scaling coefficient
α_I	Current scaling coefficient
P	Real power
S	Complex power
$\|S\|$	Apparent power
pf	Power factor
μ	Mean difference (of voltage and current)

second delay between wireless transmissions allows the capacitor sufficient time to charge for transmissions. In this design, the additional power required by the active thermistors required a larger current source so the Kill A Watt was modified to supply the XBee subsystem with an external 5V AC/DC adapter. Despite the removal of the power constraints, the 2-second cyclic sleep was retained instead of continuously streaming data.

The XBee radio can send at most 7 samples per channel per packet. While the actual signal is continuous and of the form shown in Figure 5 (left), discontinuities arise in the sampled waveform as shown in Figure 5 (right) due to the sleep delay between packets. Average power over a period is calculated as the product of the root-mean-square (RMS) voltage and current. Given the discontinuities in sampling, it is not possible to achieve accurate power calculation over very short time periods. However, it is acceptable to assume that power consumption within a short time window (on the order of tens of seconds) has reached steady state. The measured samples could be treated as a signal that has the same statistical distribution as a continuous signal of the same length. This is particularly important for power factor estimation described in the next section.

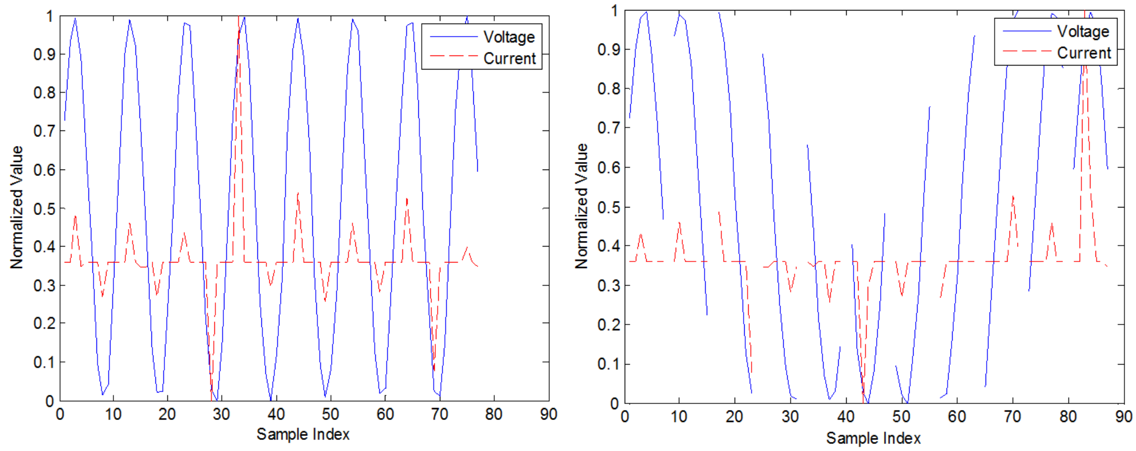

Figure 2.6: Example voltage and current waveforms produced by the energy meter (right) and as it would appear without a cyclic sleep mode (left).

This sampling approach requires that many packets of data need to be collected to accurately approximate the RMS voltage and current. The XBee radio uses a 10-bit ADC so voltage and current measurements vary between 0 and 1023. These digital values are then zero-adjusted and scaled to arrive at the analog voltage and current value.

$$\widehat{V} = \alpha_V(V - \overline{V}) \tag{2.5}$$

The RMS voltage is then calculated as

$$\widehat{V}_{RMS} = \sqrt{\frac{1}{n}\sum_i \widehat{V}_i^2} \tag{2.6}$$

To improve processing performance, the calculation can be distributed over time by preprocessing raw samples after each packet arrives. This can be done by combining Equations 2.5 and 2.6 to arrive at the form given in 2.7. RMS current is calculated similarly. Preprocessing becomes a simple task of keeping a running sum of the samples and running sum of squared samples. The scaling coefficients are determined empirically using known values of voltage and current. The full derivation is given in Appendix A.

$$\widehat{V}_{RMS} = \alpha_V \sqrt{\frac{1}{n}\sum +V_i^2 - \big(\frac{1}{n}\sum_i V_i\big)^2} \tag{2.7}$$

2.2.2.2 Power Factor Estimation

The power factor of an AC electric power system is the ratio of the real power P flowing to the load, also known as true power, over the apparent power $|S|$, which is the magnitude of the complex power S. The complex power is simply the vector sum of the real power and the reactive power Q. Power factor is critical in power measurement since different loads will exhibit different power factors. Only the real power component is responsible for electrical work while the reactive component does no work at the load and heats the electrical wires, wasting energy. For this reason, many utility providers generally bill customers only for the real power consumed, regardless of the power factor. Thus, real power was also provided in this study to reflect convention.

There are a number of methods for determining power factor in a load. The method applied here was based on the statistical dispersion between the voltage and current signals that characterizes a load with non-unity power factor and developed an empirical model to estimate power factor. This represents a significant improvement over the Tweet A Watt approach which measures only apparent power. now describe the approach for estimating the real power. Real power can be computed from apparent power using power factor as a scaling coefficient.

$$|P| \stackrel{\text{def}}{=} |S| \times pf \tag{2.8}$$

Apparent power is calculated as the product of the RMS voltage and RMS current given by

$$|S| \stackrel{\text{def}}{=} V_{\text{RMS}} \times I_{\text{RMS}} \tag{2.9}$$

where the RMS voltage and current can be computed directly from the sample measurements as described in the previous section. To estimate the power factor, voltage and current signals are first normalized to remove variation in voltage and current levels across locations and for different loads. For each voltage and current sample vector, V and I respectively, the normalized vectors are

$$V_{\text{norm}} = \frac{V - V_{\text{min}}}{V_{\text{max}} - V_{\text{min}}} \tag{2.10}$$

Then difference of the means of the normalized voltage and current is computed. A purely resistive load has unity power factor meaning the voltage will neither lead nor lag the current. For such a purely resistive load, the normalized voltage and current waveforms will be identical and thus the mean difference between the waveforms will be zero. Any deviation from unity in the power factor will correspond to a nonzero mean difference. This deviation can be modeled to derive the relationship to power factor. The mean difference is given by

$$\mu = \overline{V_{\text{norm}} - I_{\text{norm}}} \tag{2.11}$$

Comparing the power factor and observed mean difference values reported by the Kill A Watt meter (for a variety of appliances with varying power factors) shows a roughly quadratic model shown in Figure 2.7. The relationship was calculated to be

$$pf = -5\mu^2 + 0.83\mu + 0.93 \tag{2.12}$$

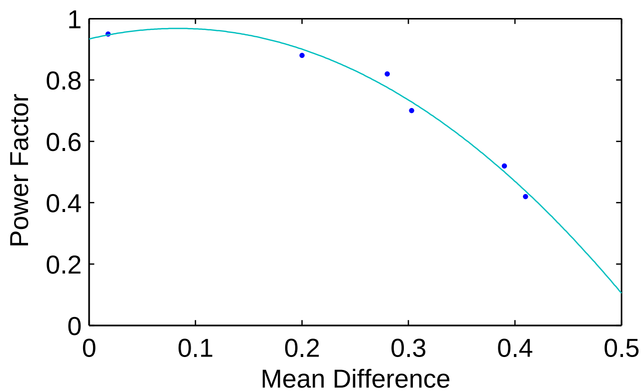

Figure 2.7: Quadratic regression model to estimate power factor.

This power factor model was then incorporated into the signal processing algorithms of the gateway daemon to compute real power. A variety of loads consisting of typical home appliances with various power factors was tested. The error in real power estimate compared to Kill A Watt's real power measurement was found to be less than 10%. These measurements are summarized in Table 2.2. The table highlights the fundamental limitation of the Tweet A Watt to produce an accurate real power estimate for general loads.

Table 2.2: Comparison of Kill A Watt (KAW) real power (RP), Tweet A Watt (TAW) apparent power (AP), and Engage real power measurements.

Appliance	KAW RP (W)	TAW AP (VA)	TAW Error (%)	Engage RP (W)	Engage Error (%)
LED Bulb	6.3	14.7	133.3	6.3	0
Incadescent Bulb	19.3	27.8	44	20.1	4.1
Wireless Router	25.4	66	159.8	23.6	-7.1
Laptop Computer	25.5	48.9	91.8	24.1	-5.5
Mini Fridge	51.2	62.5	22.1	48.2	-5.9
Blow Dryer	86.3	98.1	13.7	87.9	1.9
Flood Light	94.9	105.2	10.9	99.2	4.5
Air Conditioner	1254	1242.3	-0.9	1170.3	-6.7

2.2.2.3 End-Use Energy Calculation

As described above, the Engage system provides accurate estimation of the power consumption over time for plug load, heating and cooling use, and lighting use. Using these estimates, end-use energy statistics can be calculated for presentation on the dashboard. For a given room and given time window defined by start time t_1 and end time t_2, the total energy consumed for a load $\mathcal{L}$ in the set $\{\text{plug, hvac, light}\}$ is calculated as

$$E_{\mathcal{L}}(t_1, t_2) = \int_{t_1}^{t_2} P_{\mathcal{L}}(t)dt \tag{2.13}$$

The total load is the sum of the appliance loads given by

$$E_{\text{total}}(t_1, t_2) = E_{\text{plug}}(t_1, t_2) + E_{\text{hvac}}(t_1, t_2) + E_{\text{light}}(t_1, t_2) \tag{2.14}$$

This can further be used to calculate the total energy of this window for each room and group statistics such as average and quantile usage.

2.2.3 Software Systems

The software for the end-to-end system consisted of three components: 1) the gateway daemon, 2) software running on the server for the backend including data management and data processing scripts as well as administrative tools, and 3) software for the dashboard.

2.2.3.1 Gateway Daemon

The gateway daemon is designed to facilitate the transport of data from the energy meters to the database and then to the user. Each energy meter is configured with a unique identification number and automatically transmits analog data to its gateway. The gateway receives these packets and processes them to extract current, voltage, and power and uploads these measurements along with light and temperature sensor measurements to the server.

2.2.3.2 Engage Server System

The Engage Server system consists of a blade server hardware platform with a LAMP software stack (Linux operating system, Apache web server, MySQL database, PHP/Python/Perl programming languages). Configuration information for each energy meter was collected during installation and stored in the database. This included identification numbers of the meters installed in each room and which sensors were connected to each input channel on the meter. Software scripts also periodically processed sensor measurements into estimates of energy consumption.

2.2.3.3 Engage User Dashboard

The Engage User Dashboard, shown in Figure2.8, displays three key pieces of information: 1) usage summary information with current power consumption, daily energy usage projection, and average historical daily energy usage, 2) a bar chart showing energy usage from the past week compared to the average of other rooms as well as historical usage from the previous period, and 3) a pie chart showing the breakdown of the usage among the three load categories. The dashboard provided three levels of data resolution: Real-Time with 5-minute resolution for the past 3 hours, Hourly (as the default view) with 1-hour resolution for the past day, and Daily with 1-day resolution for the past week.

The residents of each room received a unique alphanumeric code which was used as an identifier in the URL to access their room energy dashboard. This avoided the need to develop a user authentication system since the dashboards would be openly accessible as long as the code was known. Although the code could be leaked, the dashboard contained no identifying information. This approach also allowed the dashboard to be bookmarked for easy access. Access codes were distributed via email to the residents at the beginning of the experimental treatment phase. A direct link to the dashboard was also included in individualized weekly email reports.

2.3 Experimental Design

A pilot study was conducted to investigate behavioral mechanisms leading to energy conservation in the absence of inherent pecuniary incentives. The residence hall population consisted of 102 undergraduate students living in single-, double-, and triple-occupancy rooms. These buildings were constructed as part of a single construction phase with only

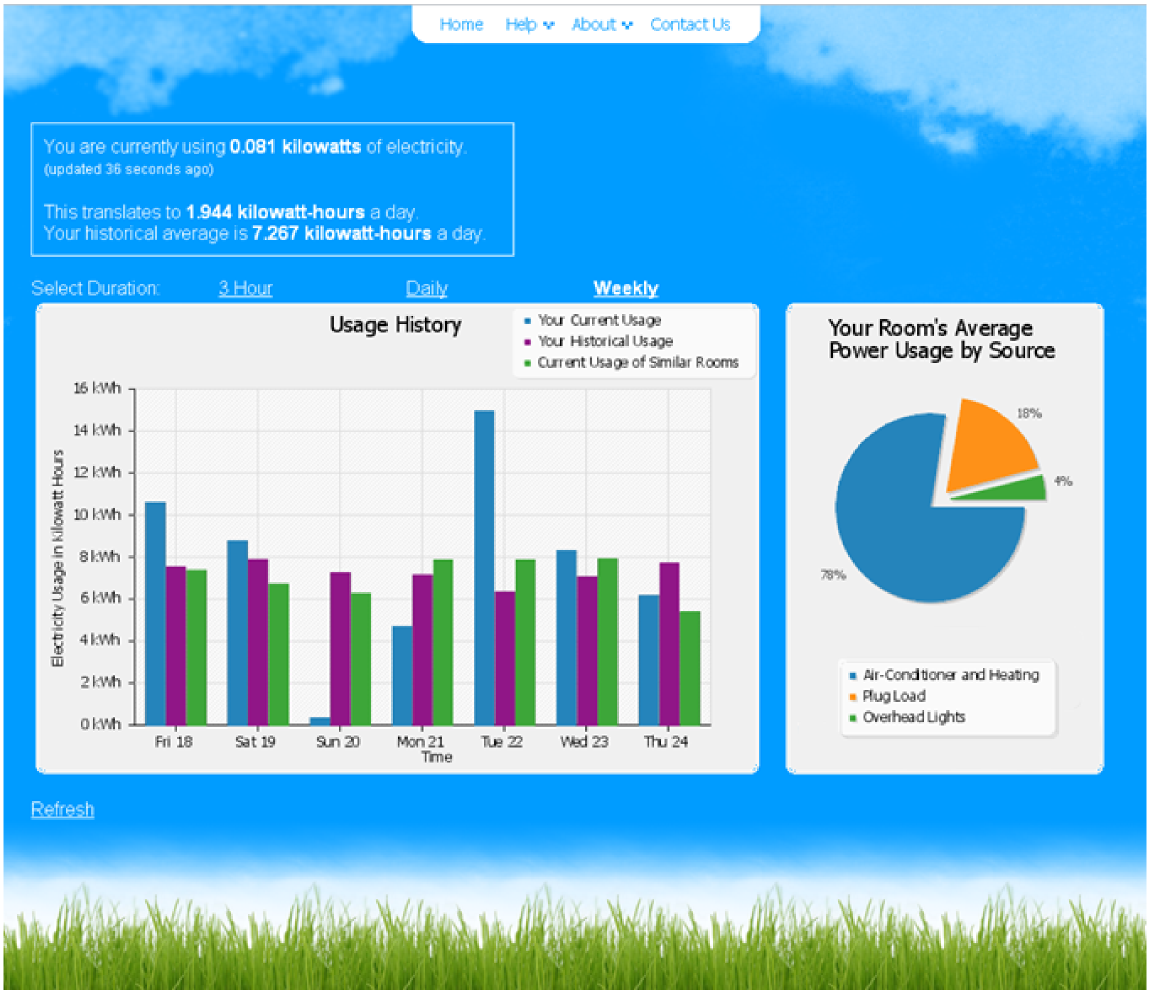

Figure 2.8: Personalized Engage energy dashboard

minor variation in design. This standardization allowed for evaluation of differences in energy consumption based on information feedback rather without contamination by differences in infrastructure. Furthermore, because students do not pay for electricity, they are an ideal population to study behavior responses to various forms of information feedback. The effectiveness of information could be studied in the absence of an inherent financial incentive to conserve electricity.

We installed electricity metering systems in 66 rooms in 3 separate undergraduate residence halls for one academic year (September 2010 to May 2011). The 3 residence halls were selected based on their similiarity as they were constructed around the same time as part of a sinigle development phase. Rooms are standardized across the buildings and use similar infrastructure designs and floorplans. Each room is equipped with a programmable thermostat, operable window, and one or two overhead florescent lights.

Table 2.3: Experiment design for Engage Pilot study.

	Baseline (6 weeks)	Phase 1: Treatment (5 weeks)	Phase 2: Treatment (7 weeks)
Control (23 rooms)	None	None	None
Treatment 1 (22 rooms)	None	Private Information	Private Information
Treatment 2 (21 rooms)	None	Private Information	Private Information Public Information

Participants in Group 2 were exposed to a public information treatment where they were rated as using above or below the average energy usage. This relative rating system was used to protect the privacy of participants and prevent contamination of the control group with electricty usage information. Ratings were made on a weekly basis where the weeks corresponded to the calendar weeks of an academic quarter. For a given week, if a room used less electricity than similar rooms based on occupancy, it was given a green sticker to indicate that it was an äbove average energy energy conserver. If a room use more electricity then it was given a red stick to indicate that it was a b̈elow average energy conserver. The language of the poster was chosen to enourage positive behavior by calling particpants energy conservers as opposed to energy users.

Ratings were made public using large posters promimently displayed in front of the elevators on each floor where a Group 2 participant resided. An example of the public information poster is shown in Figure 2.9. Posting in front of the elevators ensured that all students would pass by the posters several times per day.

Since the poster group comprised one third of the entire group, it was possible for all rooms in the public information treatment to conserve and be awarded green stickers. This was made clear to the students through several emails as well as a note on the poster. To increase exposure, a copy of the poster was emailed to each participant in the experiment whether they were in the treatment or control groups. Public information transformed the previously invisible prosocial action of energy conservation to be publicly visible, which created additional motivation for participants to conserve by creating or maintaining a green reputation.

2.4 Experiment Results

This section summarizes the results of the energy use and dashboard access data. The full description of the experimental results of the behavioral study can be found in the paper by Delmas and Lessem [DL14].

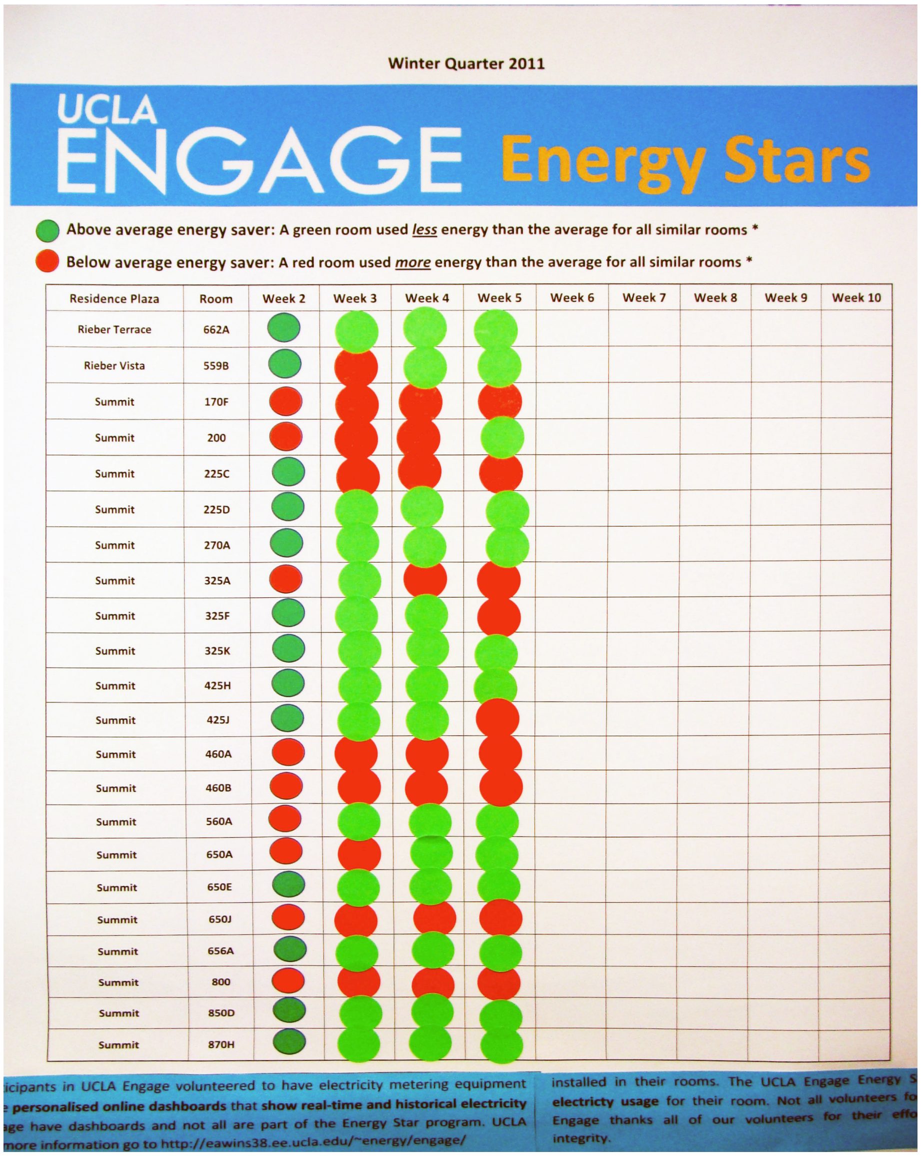

Figure 2.9: Engage public information poster. Group 2 participants' weekly energy consumption status was displayed publicly in front of the elevators on each floor where a Group 2 participant resided.

2.4.1 Energy Use

Figure 2.10 shows a stacked area plot of the daily total energy use averaged across all rooms and divided by load category (lighting, plug load, and HVAC). While lighting and plug load are relatively constant throughout the experimental period, HVAC use is more volatile. Gaps in the data are the result of network or server failure events whereas dips in electricity consumption result from academic breaks.

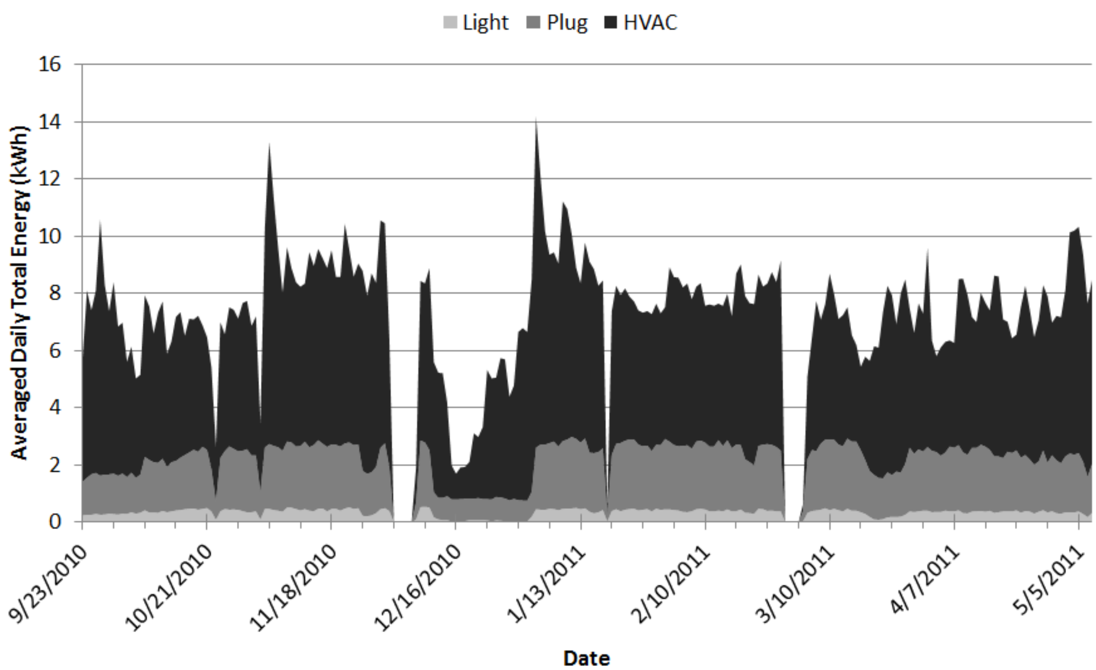

Figure 2.10: Total daily energy use averaged across all rooms, shown as a stacked area plot.

Despite this outcome, dashboard access is still higher when residents received weekly email reports.

2.4.2 User Dashboard Analytics

Google Analytics was used to track website access based on room, time of access, and display type where display type refers to the weekly view, daily view, or 3-hour view. These analytics provide a wealth of information about users' response to the information provided.

Figure2.11 shows a histogram of user access to the dashboard over the course of the experiment. While every dashboard was accessed at least once, there are only a few residents who visited the dashboard more than 80 times and many residents that visited the dashboard less than 10 times. These results follow a power law distribution and where 80% of the dashboard activity was generated by 25% of the residents, approximating the Pareto principle which states that roughly 80% of the effects of from 20% of the causes for many natural phenomena.

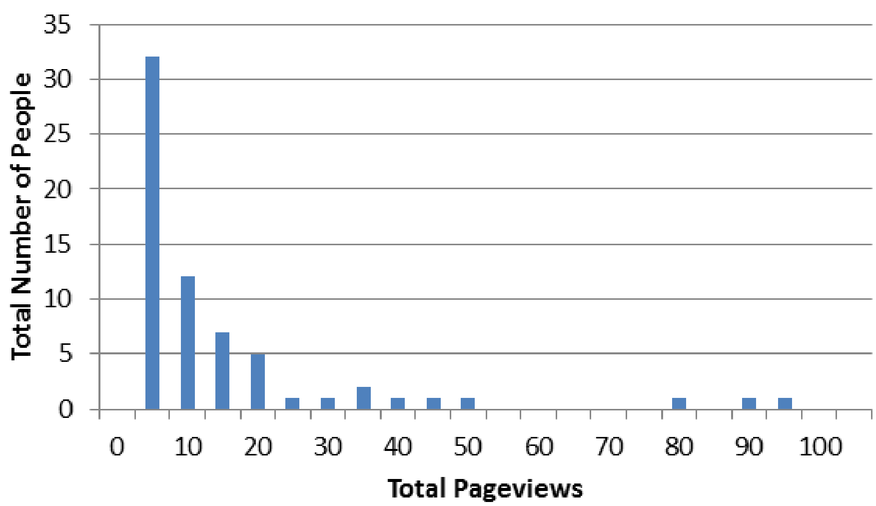

Figure 2.11: Distribution of users over total pageview count.

Another effect is the rapid drop-off in dashboard access shown in Figure2.12. Access in the first week (following the launch) starts off very high as everyone explores this new tool. However, especially after the first day (when the launch announcement was emailed) there is a large drop in activity as the novelty wears off.

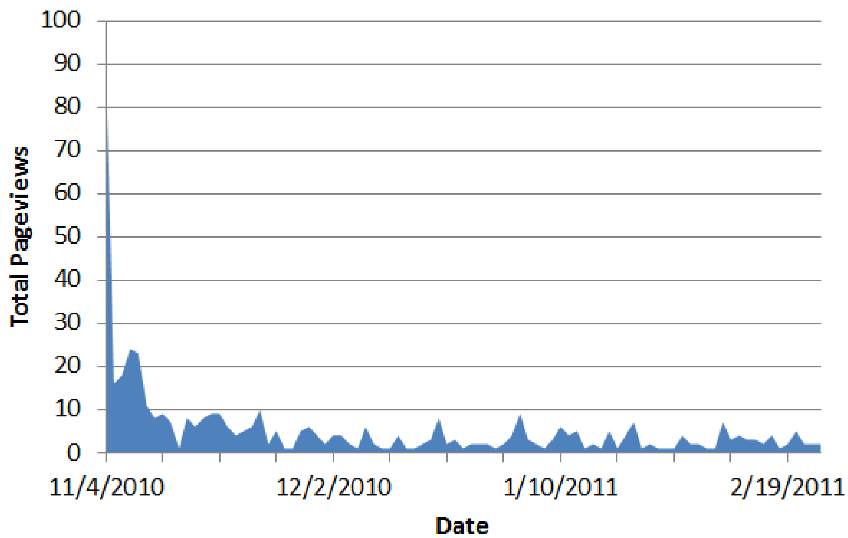

Figure 2.12: Total unique pageviews per day.

Figure 2.13 shows a histogram of total pageview count for each day of the week. Weekly reports containing summary information and a link to the dashboard were emailed on Mondays and, correspondingly, pageview count is highest for Mondays. This demonstrates that the regular, periodic emails play an important role in reminding users about the study and drawing them to the website to view their dashboards.

Figure 2.14 shows total pageviews as a function of the time of day the dashboard was accessed. Dashboard access follows a daily cycle that peaks around midday with a second peak period late at night. Presumably these are periods when people find free time from work or school to view their dashboards. Dashboard viewership could potentially be increased by sending weekly emails at daily peak access times.

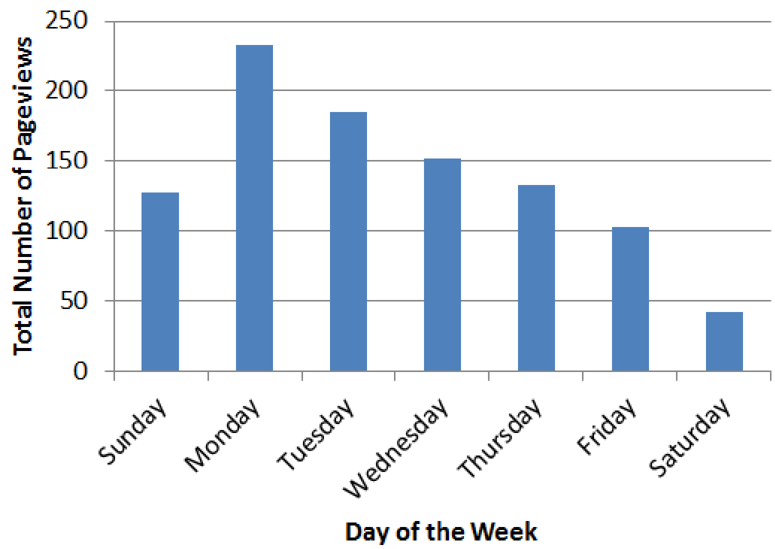

Figure 2.13: Total number of pageviews per day of the week, averaged over the full study period.

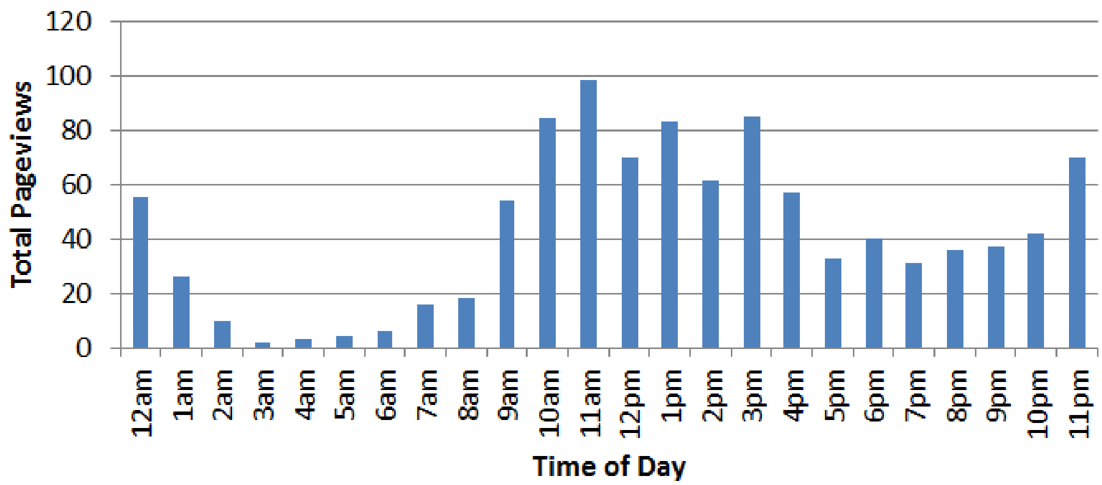

Figure 2.14: Total number of pageviews for each hour of the day, averaged over the full study period.

In order to learn what aspects of the dashboard were more useful, dashboard access was broken down by the view type (real-time, hourly, or daily), as shown in Figure2.15. It was found that there were more hourly page views (41%) as compared to daily (29%) and real-time views (30%). This might be due to the design of the dashboard since the hourly view is the default page and visitors may not be exploring beyond the hourly page. However, the view count is not heavily skewed toward any one view type. In conjunction with the power law distribution of page viewership, this suggests that people who are highly engaged with the dashboard are likely to explore the different views.

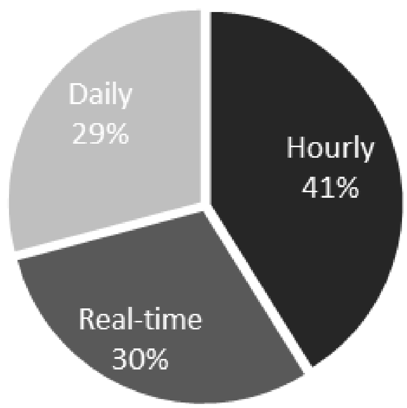

Figure 2.15: Percent of the total pageview count for each view type.

Users also spend the majority of the total access time on the hourly page rather than the real-time or daily views. Whereas this could be realistic, it is also possible that this is a limitation of how Google Analytics calculates pageview duration. Google Analytics calculates viewing duration as the difference between timestamps of subsequent pageviews. Thus, for the last or only page viewed in a session, there will not be a duration value calculated for that pageview. Therefore, even though users may be subsequently viewing the real-time and daily pages, data could skewed toward recording viewing time for the hourly page since it is the default view. This is reinforced by the pageview count breakdown in

Figure2.15, which shows similar levels of access across view types. This limitation could be overcome using more sophisticated methods for monitoring access such as a software timer to record viewing duration.

2.4.3 Energy Behavior Impact

This section summarizes several findings on the behavioral effects of dashboard information on energy consumption. Delmas and Lessem [DL14] present a thorough analysis and explanation of the results.

First, exit survey was conducted to assess the efficacy of the dashboard and what if any conservation methods were undertaken. 58 of the 102 experiment participants completed the survey, accounting for 52 of the 66 experiment rooms. The main findings of this survey indicate that 78% of participants felt encouraged to conserve energy from viewing the dashboard. The exit survey was combined with a series of focus groups. Testimonies indicated that students learned from the dashboard and became more aware of their consumption. For example, overhead lighting use was reduced by 78 watt-hours per day, or 80 minutes per day, representing a 20% reduction compared to the control group. The reported energy consumption on the dashboard was combined with weekly emails and public information about above- and below-average consumption. In addition to having the dashboard, participants whose consumption was made public were more likely to reduce their overall consumption by 20%.

2.5 Conclusion

This chapter presented Engage, a rapidly-deployable, retrofit, end-to-end system for monitoring per room electricity consumption. The technology was designed to address the infrastructure constraints encountered in a large-building environment while providing individualized, appliance level, real-time and historical feedback. The system could easily be applied to any large building where per room energy monitoring is desired but the electrical infrastructure does not support centralized approaches to monitoring. The experimental objectives and the deployment challenges informed the design of the system, along with cost and development constraints.

The design of energy meter was discussed in detail as being based on an open-source modification of commercial hardware and improved signal processing enabled estimation of real power consumption. The design of the proxy sensors, used for measurement of light and HVAC usage, was also described.

One limitation of this system design is the difficulty of performing system maintenance compared to systems that operate outside of the residence and do not have to living space considerations. However, software improvements can improve overall system reliability and minimize the need for on-site maintenance. Startup diagnostics procedures on the gateway, such as filesystem checks coupled with a self-reboot policy could help avoid flash drive corruption that leads to system failure as well as enable system recovery. Automatic reporting of system analytics could reduce system downtime and data loss.

From the dashboard analytics, the weekly summary report emails acted as a trigger to view the dashboard. A power law distribution was observed in the dashboard access where most users were minimally engaged but there was a small number of highly engaged users. Dashboard information resulted in energy conservation of 20% when combined with public information about energy consumption.

In the next chapter, we detail the deployment at the graduate student family housing complexes. There, residents are married couples or have children. The residents also pay for their electricity consumption, and are thus a better representation of the general population. The Engage systems have been adapted to include circuit level sensing to leverage the electrical infrastructure of these apartments, each of which uses a distribution panel. Software features have also been implemented to improve gateway reliability and system diagnostics. Based on the results of the pilot study, the dashboard designs have also been modified to improve engagement and facilitate learning.

CHAPTER 3

Residential Energy Monitoring and Behavior Guidance

The previous chapter described the Engage methodology and the energy monitoring pilot study conducted in the UCLA undergraduate residence halls. The pilot study revealed interesting behavioral responses to information feedback. For example, public status information proved more effective than private information at motivating energy conservation. However, to understand these conservation effects, it's important to consider the circumstances of the study. The population consisted entirely of undergraduate students living in a dense-occupancy housing environment, where the social network should be more highly connected than typical residental areas. On top of that, utility costs were standard and paid upfront removing personal financial incentives for conservation. So while the study was successful and revealed significant results, the population was arguably unique.

A broader study was design which would be more representative of the average resident in Los Angeles. The field experiment site, University Village, is an apartment complex for graduate student families. Eligibility for housing at University Village requires that at least one occupant is a currently-enrolled graduate student at UCLA and is either married or has children or both. Thus, single-parent families may also reside there. Residents also pay their own utility bills, including electricity, gas, cable and internet. This is an important difference from the undergraduate residence halls population which paid a flat housing fee that abstracted individual consumption. The characteristics make the graduate student family population much more representative of the average California resident so that experimental results are more generalizable.

The site comprises two sites with 1,102 one-, two-, and three-bedroom rental apartment units. Of these, 124 apartments were occupied by residents who agreed to participate in the experiment, also known as the ENGAGE project, and were equipped with an electricity metering system that allows electricity usage to be recorded in real time. During the study period, some participants moved out of the apartment complex and the new occupants of their apartments agreed to participate in the ENGAGE study. This led to a sample of 137 unique households. Each apartment is equipped with heating and cooling systems and a full kitchen including a refrigerator, microwave, stove, dishwasher, and garbage disposal. Except for variations in size and floor plan, apartments are standardized with the same major appliances and amenities. This consistency ensures that variation in electricity usage results from household behaviors and lifestyles, not differences in apartment or appliance features. Furthermore, circuits in University Village are fairly standardized with only minor variations which allowed for a hardware installation kit that would accommodate all of the circuit breaker panels without any hardware reconfiguration.

3.1 Introduction

Providing more detailed feedback to consumers about their energy usage at the appliance level can potentially encourage such behavioral changes [EMDL+10, Fis08, NRB09]. However, currently, the majority of residents in the United States and around the world do not receive such feedback. Consumers electricity bills report total consumption rather than consumption by each appliance, and do not provide information about which appliances offer the consumer the highest potential for energy savings. Kempton and Layne analogize a households electricity bill to getting a grocery-shopping receipt each month without knowing how much each good contributed to the dollar amount they must pay [KL94]. The planned deployment of more than 65 million digital electricity meters by 2015 [Edi] will allow utilities to provide a wealth of new information to more than half of the nations electricity accounts, unlocking new conservation potential [AGSA13]. While this new information could help consumers make better decisions about their appliance use, little is known about energy consumption patterns by appliance and the behavioral component of appliance energy use.

In this chapter, high-frequency appliance-level electricity consumption data was collected from 124 apartments over 24 months to answer the following questions: Are consumers cognizant of their electricity usage across different appliances? Are there important differences in the use of the same appliances across households and what is the behavioral component of appliance energy use? Which individual appliances are contributing to peak demand usage? How do the savings from installing new appliances compare to savings from behavioral changes? The answers to these questions have important implications for the design of more effective policies to encourage energy conservation behavior.

3.2 Background

There is a growing interest in reducing energy consumption and the associated greenhouse gas emissions in every sector of the economy. According to the International Energy Agency, the continuing demand for newer appliances with improved functionality and more power is leading to an increase in electricity consumption even though appliances are becoming more energy efficient. This increasing residential energy consumption level warrants a detailed understanding of the residential sectors consumption characteristics to prepare for and help guide the sectors energy consumption.

Studies of the effect of different types of energy feedback on energy savings indicate that information on real-time appliance-level energy consumption data has the potential to empower consumers to effectively manage their household energy consumption and encourage conservation [DL14, EMDL$^+$10, NRB09].

The current study goes beyond previous work analyzing appliance-level consumption in five ways. First, scholars have argued that households are unaware of how much electricity is used by specific appliances and the potential for energy savings from each appliance [ADDdB10]. However, so far, the evidence presented was mostly based on surveys and expert recommendations - not on observed household electricity usage. This study compares each households actual electricity usage with the estimated usage they stated at the beginning of the study. This allows precise evaluation of households knowledge of energy use for each appliance.

Second, studies that have shown variation in usage across households focus on total usage or on usage in a particular subset of appliances. For example, [Lut93] notes that even in energy consumption studies that use nearly identical units, electricity usage can vary as much as 200-300% but did not differentiate among appliances. Other appliance studies include the research by Wood and Newborough [WN03], who focused mainly on cooking appliances, Coleman et al. (2012), Rosen and Meier (1999, 2000), and Rosen et al. (2001), who focused on entertainment appliances, and Isaacs et al. (2010), who studied space heating. In contrast, this study includes a broad set of appliances and end uses that are found in most homes. The study site also consists of apartments with little variation in design and identical major appliances, something that most previous studies were unable to provide [Par03, PCCR93]. Additionally, since all the apartments are in the same complex, the results are not affected by variations in weather [HdD04].

Third, recent studies highlight the growing share of non-HVAC sources in electricity consumption [DAFSF11, SBM12]. However, neither study was conducted in the United States and both had other shortcomings such as small sample size and using a non-standard set of appliances across countries. This Engage study compares electricity usage for appliances throughout the day and assesses lighting and plug load usage during peak demand hours using a large sample size with a common set of appliances across households.

Fourth, previous studies have attempted to estimate the variation in appliance usage across household types [BK08, PCCR93]; however, they were unable to do this for a standardized set of major appliances. The study site used here allows an identical set of major appliances to be compared across households in apartments with little variation in design. Finally, this is the first study that uses real-world observations to estimate energy savings from the installation of a new appliance; previous research relied on simulation techniques to estimate energy savings [DAFSF11].

3.3 Engage University Village System Architecture

Since the original metering system at University Village did not provide detailed information about electricity usage, an end-to-end system architecture was designed to measure real-time, appliance-level data and provide feedback to households. The elements of the Engage energy metering system are displayed in Figure 3.1. The electrical panel distributes electricity from the main power lines through different circuits to various outlets and appliances throughout the home. Commercial energy meters are installed inside the panel to measure the energy consumption on each circuit using current and voltage sensors. The energy measurements are transmitted wirelessly to a gateway connected to the home network. The gateway transmits the data to be processed on the backend system which then also provides feedback to residents via a web dashboard and email reports. Each component is described in detail in the following subsections.

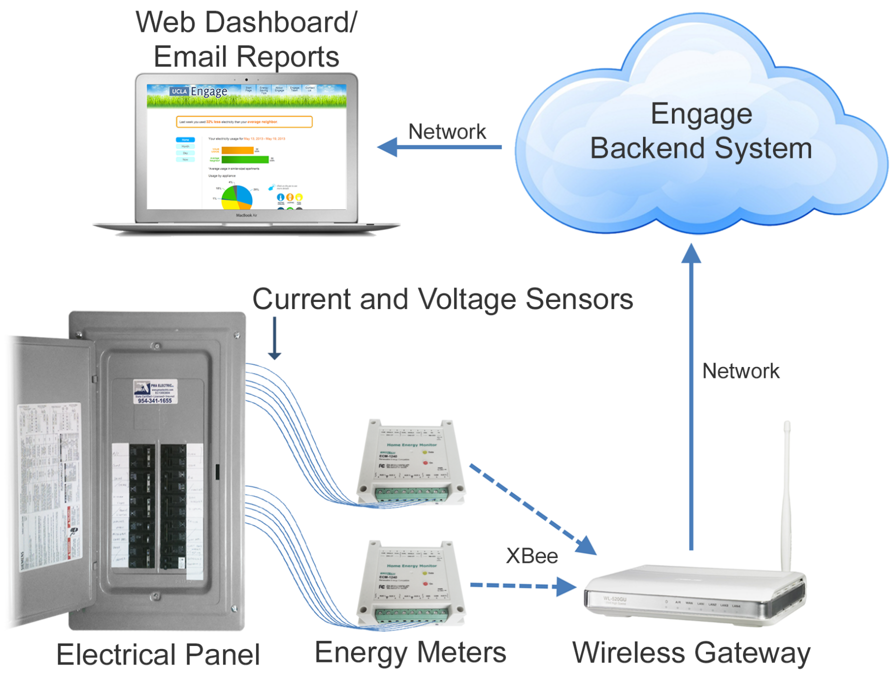

Figure 3.1: Engage apartment energy monitoring system components and data flow.

3.3.1 Electrical Infrastructure and Circuit Configuration

University Village uses a 208/120 volt 3-phase electrical service. Each apartment receives two legs of this system, which will be referred to as phase A and phase B. Each apartment at University Village uses its own electrical panel through which power is routed from mains electricity to the various loads in the apartment. Many of these loads are powered by dedicated circuits. For example, recessed lighting in the kitchen and bathrooms are powered by dedicated lighting circuits. Similarly, the refrigerators are powered by a dedicated refrigerator circuit. Thus, it is convenient to acquire high-granularity measurement of individual appliance power consumption from the electrical distribution panel. Although this provides a significant advantage compared to large building infrastructures, as in the undergraduate residence halls from Chapter 1, all wall outlets are fed by one or two circuits. Thus, it is not possible to isolate power consumption for plug-in devices.

Another advantage for both system design is configuration standardization. Major appliances in University Village are highly uniform with only some minor variations. For example, the heating and cooling system in each apartment is typically powered by four circuits and the refrigerator and microwave are always on their own dedicated circuits, etc. This allowed for the design of a hardware installation kit that would accommodate all of the circuit breaker panels without any hardware reconfiguration. Circuit configurations were also largely consistent. Most electrical panels had the same circuit layout where, for example, the refrigerator would be on the first circuit on phase A. However, another apartment might have the refrigerator on the fifth circuit on phase B. Most apartments fell into one or two templates, while there several other rare templates. The variation is a result of different electricians performing the initial wiring installation during building construction. This created some challenges and made it necessary to record each circuit's configuration to be stored in the database for use in appliance load calculations. Even though labels inside the panel provided a key to determine the load for each circuit, many of these labels would be incorrect and would not be detected until preliminary data collection and load signal analysis.

3.3.2 Hardware Installation Kit

The panel instrumentation and end-to-end system design follows an architecture similar to what was developed for the undergraduate residence halls. The hardware consists of wireless energy meters for measuring energy consumption in the panels. The meters transmit data to a wireless gateway for local processing and transfer to the backend system. A commercial energy metering device was used which provides current measurement channels for individually measuring up to seven circuits. The meters are designed for only single-phase measurement and since the buildings use a two-phase electrical service, two meters were used to fully instrument the electrical panel. The energy meters use personal area network (PAN) radios (called XBee radios) to enable wireless communication between the energy meters and the gateway. The gateway is built from an Asus 520-gU wireless router which was modified to interface with an XBee radio on its serial port for communicating with the energy meters.

3.3.3 Gateway and Server Software Systems

On the gateway, a software program which continuously runs in the background (known as a daemon) was developed to manage reading and local processing of the meter data as well as uploading data to the server. The program is executed via a boot script to enable automatic recovery in the event of a hard reset. A status program is also executed periodically using a job scheduling utility to ensure that the daemon process is always running. If a fault is detected, the status check program issues a kill signal or the process may be killed forcefully, after which the daemon is restarted. The daemon is also responsible for ensuring reliable upload in the event of a network failure, server processing delay, or server crash. Other utilities such as network time protocol (NTP) were added to manage time synchronization of the on-board clocks.

Due to the distributed and remote nature of energy monitoring, software for the gateways and server was designed to facilitate system management. Data reliability is critical given the time and cost of deployment. Consequently, it is necessary to detect failures and troubleshoot as quickly as possible. The gateways are installed as clients on the residents' own routers. This puts them behind a network address translation (NAT) wall and they are not directly accessible outside of the local network. To circumvent this limitation, a virtual private network (VPN) using open-source VPN software enabled remote access to the gateways from the ENGAGE server. Similar to the processing daemon, a VPN client on the gateway is executed using a boot script and a status check script ensures that the VPN connection is maintained 3.2. From the server, software was developed to perform remote status checks on each gateway. As with any software development, bugs can be discovered after release or improved features may be developed. A script was developed to use the VPN IP addresses of the gateways to remotely check the version of the processing daemon and update to the latest version if a MD5 checksum mismatch was detected.

Remote system status was also monitored by analyzing the data upload history and an administrative dashboard was developed to provide these analytics. These analytics included a summary of hourly sample counts for each apartment, last upload timestamps, as well as VPN IP addresses and all hardware ID numbers for meteres and gateways. These administrative tools facilitated debugging and repair and and enabled effective system management. Given the size of the deployment, as wellas limited time and human resources, it was absolutely critical for the success of the experiment.

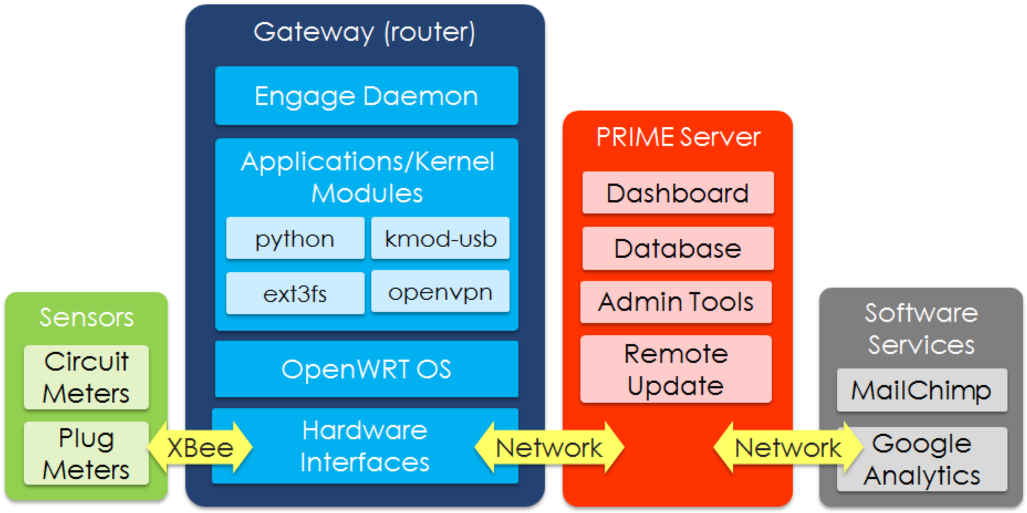

Figure 3.2: End-to-end architecture and software systems.

3.3.4 Device Assignment and Data Processing

Data from the meters consist of energy measurements in units of watt-seconds for each of seven channels, identified by meter ID's. The meters record energy using counters, analogous to the dials on analog utility meters. These counters increase monotonically until a maximum byte value is reached at which point the counters start over at 0. Data packets are transmitted wirelessly to the gateway at 1 Hz. A custom Python daemon running on the gateway receives and parses the packets, preprocesses the data, and then uploads measurements to the gateway along with identifiers.

With 4 bytes per channel, 7 channels per meter along with timestamp and identifiers, 2 meters per apartment, and 124 apartments recruited in the study, the system would generate almost 300 GB per year. To reduce the amount of data stored, the data is down sampled by the gateway at 1/30 Hz (1 sample every 30 seconds). The incoming data packets are monitored constantly and at the end of the 30 second window the total energy per channel is computed as the difference of the energy from the first and last packets received. This energy is divided by time (nominally 30 seconds) to produce power measurements.

$$P_{m_i,c_j}(t) = \frac{E_{m_i,c_j}(t) - E_{m_i,c_j}(t-T)}{T} \tag{3.1}$$

The gateway then uploads the power measurements to a database on the Engage server via HTTP POST. Once every hour, a software script processes the new power data into hourly energy measurements as follows. Let $A = \{a_1, a_2, ..., a_M\}$ be the set of apartments. For each apartment $a \in A$, there are a set of meters $M_a = \{m_1, m_2, ..., m_N\}$. For each meter m, there are a set of measurement channels $C_m = \{c_1, c_2, ..., c_Q\}$. Each meter provided 7 circuit measurement channels. Two meters were necessary to fully instrument each panel since the electrical system used two legs of the three phase electrical service. Two meters were also sufficient to instrument each circuit in the electrical panel since each apartment typically had 14 circuits including the two mains. In apartments with more than 14 circuits, circuit measurements were simply combined into single channels, e.g. multiple HVAC circuits could be measured together.

For each apartment $a \in A$, the power for channel i on meter j is given by $P^a_{m_j,c_i}$. Next, a set of appliance loads $L = \{l_1, l_2, ..., l_R\}$ were defined corresponding to the following categories: “plug load”, “lighting”, “HVAC”, “refrigerator”, “dishwasher”, and “other kitchen”. The power for each load $l \in L$ in apartment a is given by

$$P^a_l = f^a_l(P^a_{m,c}) \tag{3.2}$$

for each $m \in M_a$ and each $c \in C_m$ where f^a_1 is some known function based on the circuit configuration of a particular apartment. For example, in apartment 305, the HVAC circuits may comprise channels 3 and 4 on meter 1 and channel 5 and 6 on meter 2 so the HVAC power function would be $f^{305}_{\text{HVAC}}(P^{305}_{m,c}) = P^{305}_{1,3} + P^{305}_{1,4} + P^{305}_{2,5} + P^{305}_{2,6}$. Each load power function is defined and stored in the database. Conveniently, circuit configurations were not completely unique for each apartment and most used a few templates. The templates functioned as a higher level abstraction of a set of common load power functions and was also stored in the database. Despite some level of standardization, a great deal of effort was needed to collect the configuration information and still resulted in some errors which required rechecking.

3.3.5 Web Dashboard and Information Feedback

Information feedback was implemented using individualized web dashboards and weekly email reports. An example dashboard is shown in Figure 3.3. The dashboard was designed to inform the residents about various aspects of their energy consumption using different graphical elements and provided them tips on how to reduce their energy consumption. Weekly summaries of total energy usage and appliance-level energy breakdowns with comparison to a reference group were provided. Several temporal breakdowns of energy consumption were provided: (1) daily total energy for the past four weeks, (2) hourly total energy for the past day shown in Figure 3.4, and (3) real-time total power consumption shown in Figure 3.5.

Each chart also provided interactivity so that users could better engage with the dashboard to make behavioral inferences. Hovering over a chart element would generate a tooltip with numerical values and date and time information. Hovering over each slice of the pie chart would cause an energy savings tip[1] to appear. These tips were drawn from a collection of tips provided by the U.S Department of Energy [U.S]. Clicking on a each

[1] An example for reducing heatign and cooling energy consumption is "Close your doors and windows. The air conditioner wont work so hard to heat and cool your apartment."

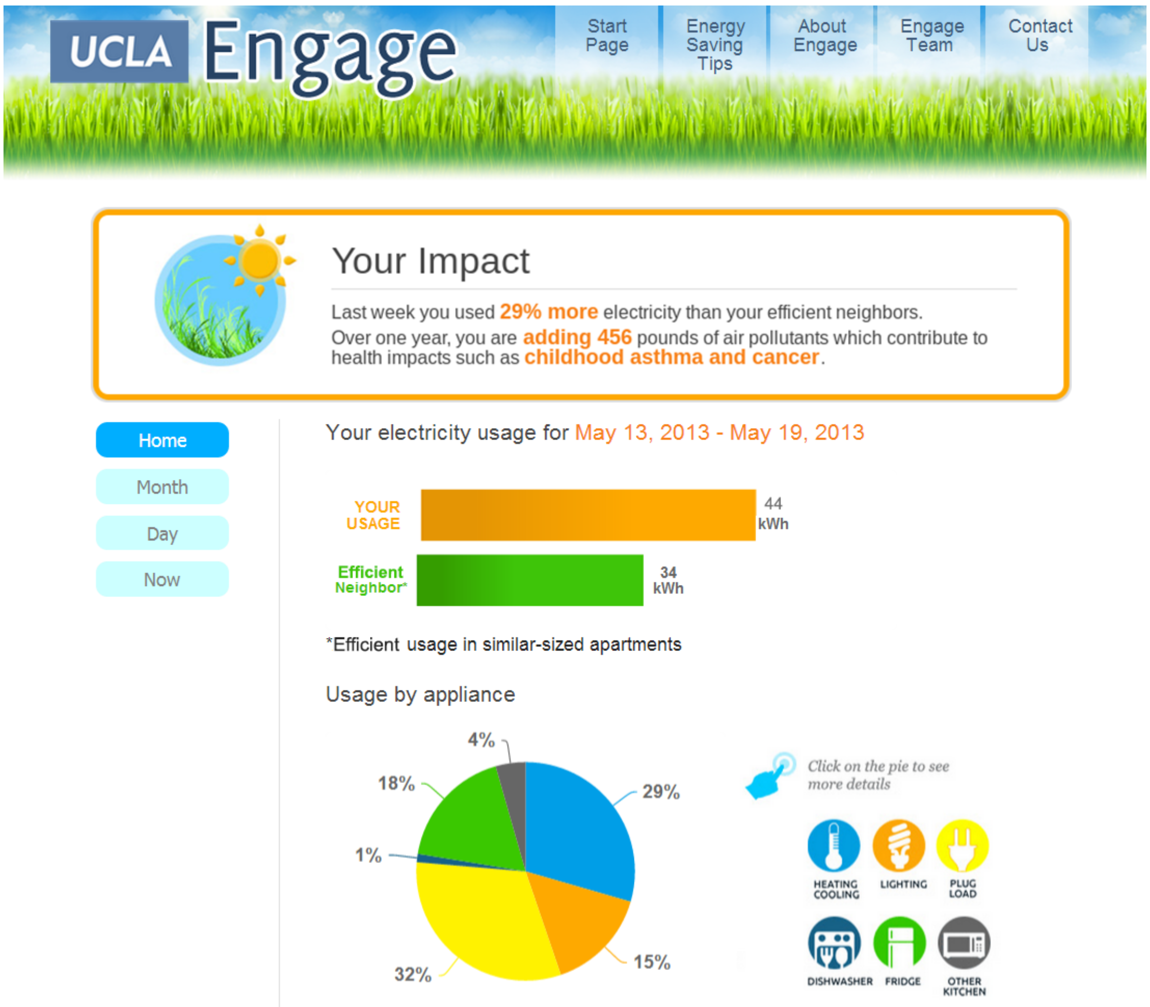

Figure 3.3: Main Engage dashboard view with weekly summary information including summary message, total energy use and reference comparison, and appliance level pie chart.

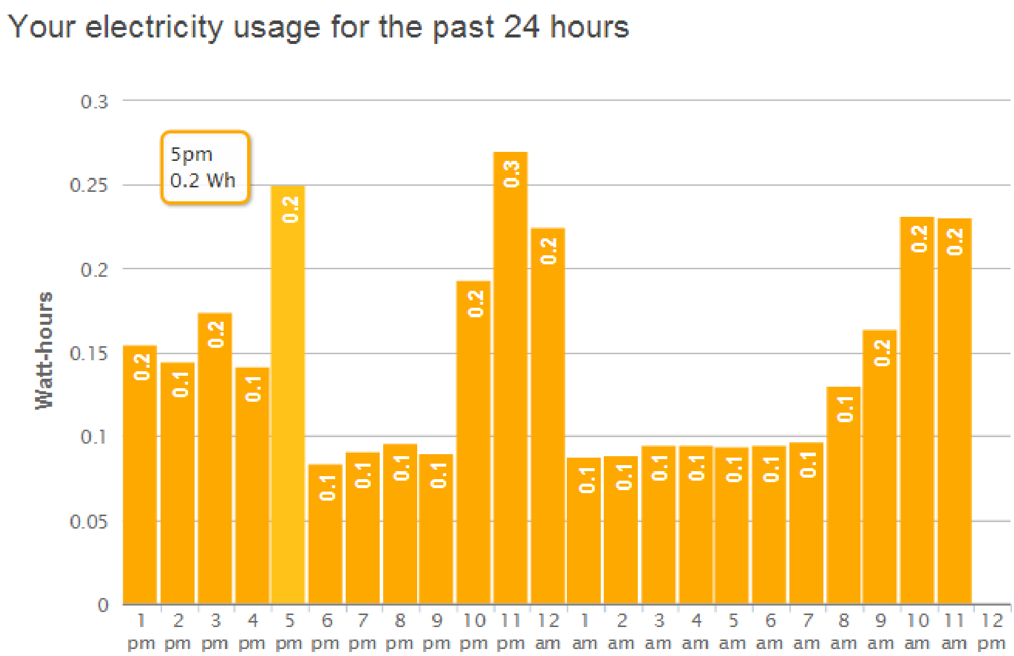

Figure 3.4: Hourly total energy consumption for the past 24 hours.

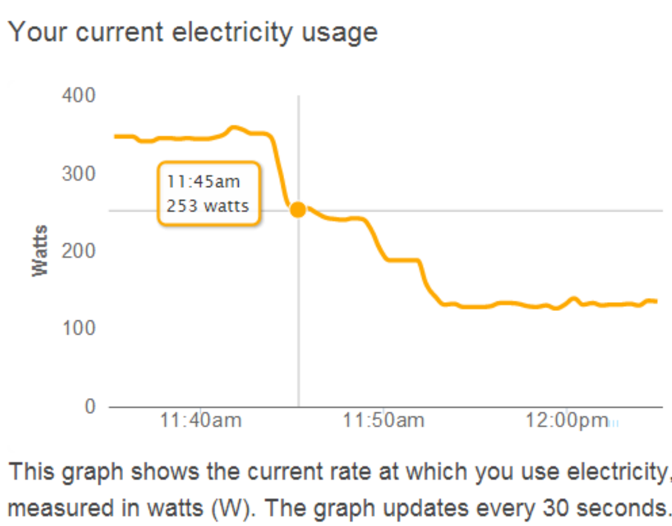

Figure 3.5: Real-time display of apartment power consumption. The view automatically refreshes every 30 seconds.

slice of the pie chart would generate a popup showing that appliance category's recent power as shown in Figure 3.6.

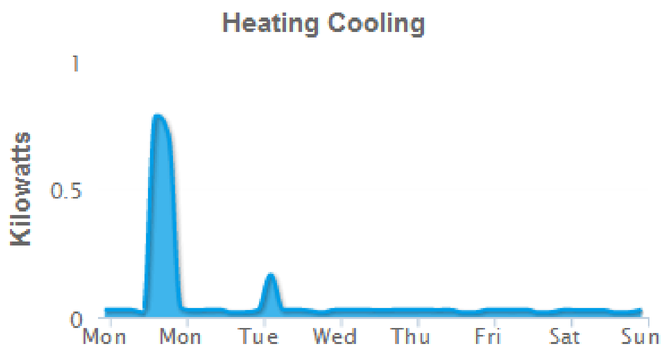

Figure 3.6: Example of appliance-specific daily power consumption history for heating and cooling usage.

The dashboard structure, or sitemap, is shown in Figure 3.7. The information feedback elements were structured hierarchically with the weekly summary as the top-level information. The other displays were easily accessible using the side menu for users who wish to drill down and learn more about their energy consumption. The dashboard was structured this way to avoid overwhelming the user with information and placing summary information first. The other information was easily accessible if a user was so inclined.

3.4 Experimental Design

The experiment was conducted from October 2011 to June 2012. Baseline measurements were taken starting on October 5. The treatments started on March 20, 2012. The first treatment period lasted until May 8, 2012 (7 weeks), and the second treatment period lasted until June 19, 2012. In treatment period 1, there were two treatment groups and a control group. Participants were randomly assigned to these groups. The control group consisted

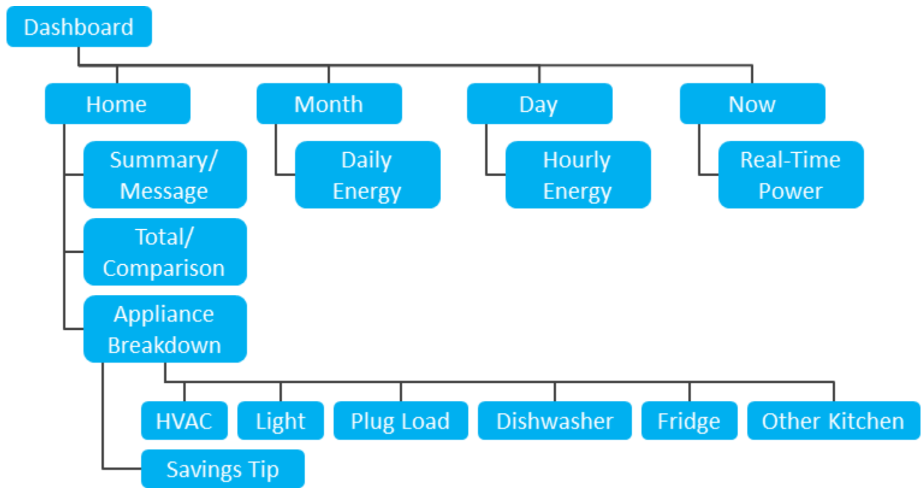

Figure 3.7: Sitemap for the Engage personalized energy use dashboard.

of 38 participants. These apartments had the energy metering system installed, but tenants did not receive any information about their energy usage. Both treatment groups received weekly e-mails with information about the past weeks energy usage as well as energy savings tips. All treated participants also had access to the web dashboard. Between the two treatment groups, the message describing their impact was varied. The 44 participants in the health treatment group received information about their emissions while the 43 participants in the financial treatment group received information about the money they saved (or spent). Figure 3.8 depicts the actual treatment messages.

After seven weeks, half of the participants started receiving an additional treatment. They were publicly rated as energy efficient (10% most efficient star), below average users (50% most efficient green dot), or above average users (50% least efficient red dot). The rating was added to participants personal websites and weekly reminder e-mails. In addition, posters were displayed at the entrance to each building. The posters listed scores for each participant in the public condition and were updated weekly (see Figure 3). To avoid polluting the experiment, the public condition was applied only on one side of the

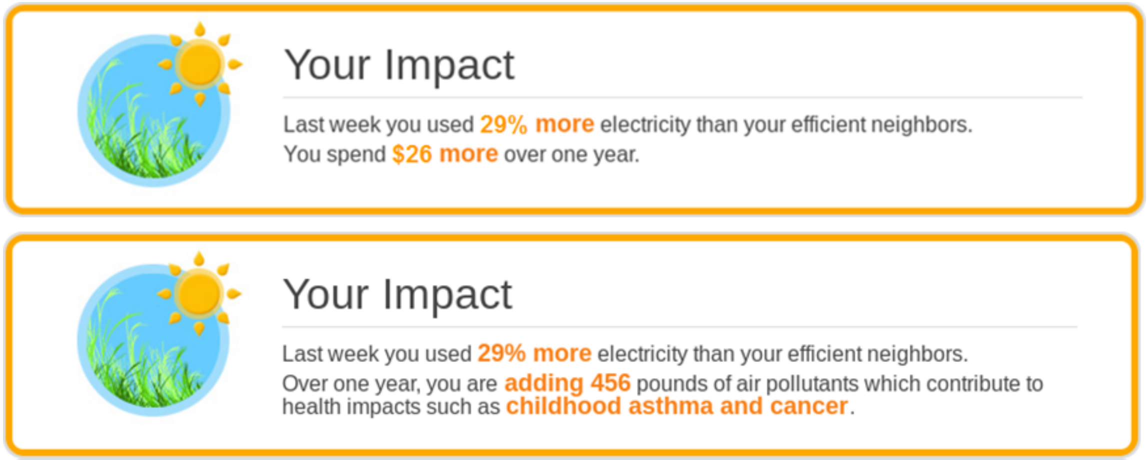

Figure 3.8: Messages provided to the financial (top) and health (bottom) treatment group.

complex. After the experiment ended, persistence measurements were taken throughout the fall semester for all participants who had not moved out.

Table 3.1: Experiment design for Engage study.

	Baseline (22 weeks)	Treatment Period 1 (7 weeks)	Treatment Period 2 (7 weeks)
Control (38 homes)	None	None	None
Treatment 1 (44 rooms)	None	Finiancial Information	Financial Information Public Information
Treatment 2 (43 rooms)	None	Health Information	Health Information Public Information

3.5 Results

Data was collected from January 1, 2012 through December 31, 2013 from 124 apartments that accounts for 137 unique households. This time period was chosen so that there would be a sufficient amount of time to determine each households electricity consumption behavior and also to allow for variation in temperature that will affect how a household uses particular appliances. While this paper does not discuss the results from the information strategies that were used in the ENGAGE project, it is important to mention that during this time period some participants were part of a treatment group that did lead to changes in their electricity usage. Results from the time period used in this study were compared to those that use data that was collected before any treatments were used and no significant differences were found. Because of turnover in theapartments, two full years of data could not collected for everyone who participated in the study.

3.5.1 Appliance-Level Trends

The data enables many types of visualization and analysis to be performed. For example, appliance-specific daily trends can be visualized. Figure 3.9 shows the energy consumption of the HVAC, lighting, plug load, and refrigerator averaged over all apartments over the full experimental period. The profiles confirm intuition about the usage behavior of these appliances. For plug load and lighting, the energy consumption peaks in the mid-morning and late evening when residents are most likely to be present in the apartment. The refrigerator shows a similar pattern resulting from both active use and from regulation with ambient temperature (as ambient temperature increases through the day, the refrigerator will need to use more energy to maintain its set point, and as the temperature decreases through the night, the refrigerator will use less). Since the majority of the experiment took place during the winter and spring period, the HVAC shows an opposite trend, where heating was used during the night to maintain a set temperature.

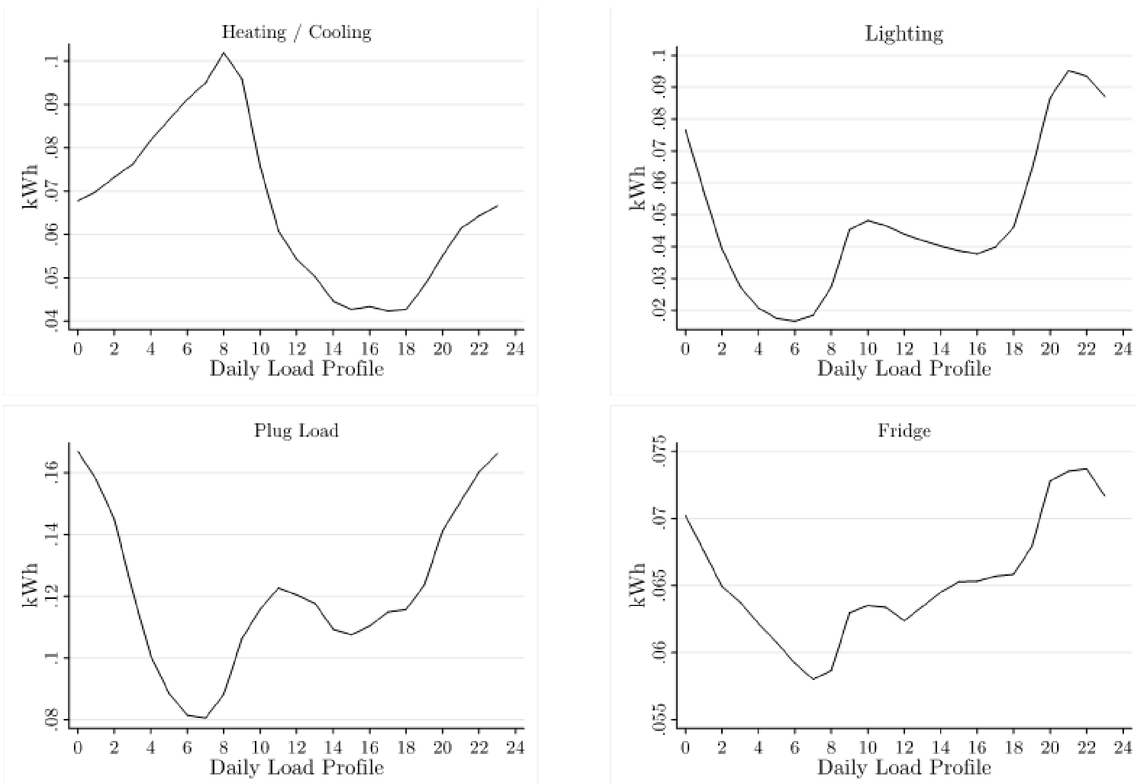

Figure 3.9: Average daily load profiles for various appliances.

3.5.2 Energy Behavior Impact

Energy data for each monitored end use was recorded along with basic demographic information so that differences in usage caused by variation in household characteristics could be accounted for. The average head of household is 31 years old, approximately 36% of households have children, and 9% of the households were members of an environmental organization. The average apartment is 863 square feet, 61% of the apartments have two bathrooms, 75% of the apartments are on the second or third floor, and about 57% face the south. The average household used 7.58 kWh of electricity per day with the majority of this coming from HVAC (1.77 kWh), plug load (2.42 kWh), and the refrigerator (1.43 kWh). The data shows that a user in the 75th percentile uses nearly twice as much electricity as a user in the 25th percentile. Variation in each end use is discussed in the next section.

Health and environment messages, which communicate the public health externalities of electricity production such as childhood asthma and cancer, outperform monetary savings information as a driver of behavioral change in the home. The results are shown in Figures 3.10 and 3.11. Figure 3.10 shows the mean kWh energy consumption for each experimental group over the duration of the experiment. Interestingly, there is initially high variance in the baseline period as the sample size is small because apartments are gradually being instrumented. There were also some bugs in the software that resulted in data corruption but as these problems were resolved, the means converged. Over the course of the experiment, the mean energy consumptions diverge for the treatment groups.

As Participants who received messages emphasizing air pollution and health impacts associated with energy use reduced their consumption by 8.2% over the 100-day experimental monitoring period versus the control group. Health and environment messaging

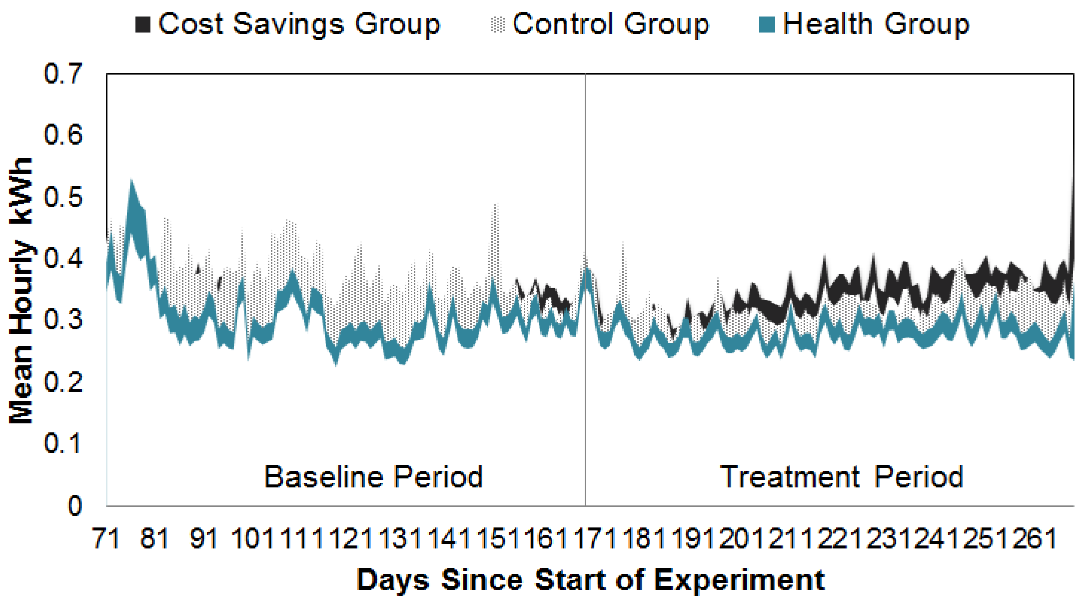

Figure 3.10: Messages provided to the financial (top) and health (bottom) treatment group.

was particularly effective on families with children, who collectively achieved up to 19% energy savings in the target population. This is shown in Figure 3.11. These net energy savings, which invoke considerations of health damages as a psychological mechanism, are at the high end of prior/nonprice strategies using social comparisons. Here, children had an amplifying effect on the base behavior response where households had their children's interest as a primary motivator. When the financial savings was small, residents were encourage to consume more electricity to provide a comfortable environment whereas when the health message encouraged conservation as a means of protecting their children's health.

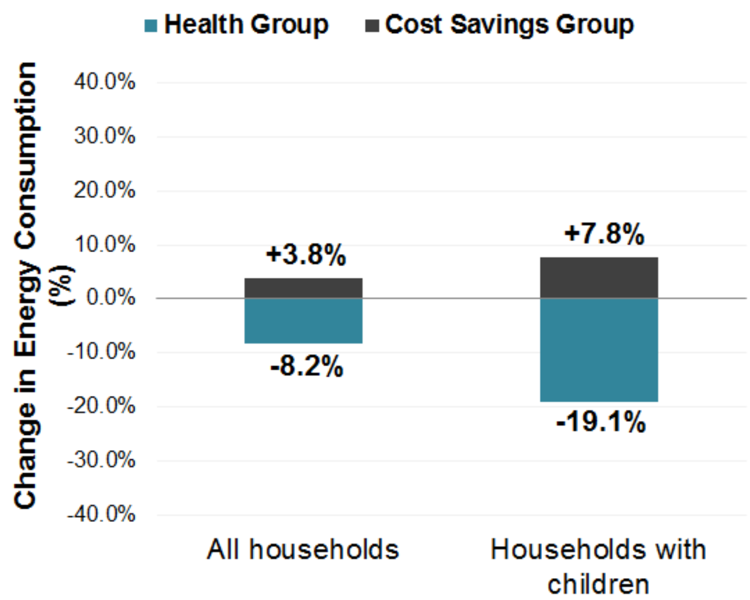

Figure 3.11: Energy behavior experiment treatment effects.

3.6 Conclusion

This chapter presented the implementation and deployment of a high-resolution, high-granularity, end-to-end energy monitoring system designed to enable information feedback for guidance of energy conservation behavior. The system was installed in residential apartments and a field study was conducted to test various modes of information feedback. Data from the experiment was used to answer important questions about information feedback for energy conservation that previous research could not answer due to technological and data limitations. The results highlight the potential for energy behavior research to inform policy in energy conservation and demand-response, for example.

CHAPTER 4

Time-Domain Energy Load Disaggregation

The previous chapters presented results from studies that demonstrated the effect of various information feedback mechanisms on energy conservation behavior. These studies emphasize the importance of large-scale deployment for understanding the effectiveness of different information feedback mechanisms for behavior guidance, particularly with regard to hetergeneous populations and environments. Despite the successful outcomes, return on investment remains a critical factor in deployment of metering equipment. The material and installation cost of high-granularity energy monitoring systems limits the ability to deploy at large-scale. Within the graduate student family population, the cost difference between the energy consumption of an average energy consumer and that of high or low consumers was less than $100 per year while the material cost of the current energy monitoring system for a single apartment unit was approximately $500 dollars. Return on investment would have already been difficult to achieve without factoring in labor costs, which includes equipment assembly, testing, installation, and maintenance. For these reasons, large scale deployment does not appear to be a very tractable proposition for consumers and businesses.

This chapter presents a time-domain approach for energy load disaggregation. The best alternative to direct current measurement for monitoring energy consumption with high-granularity is to perform load disaggregation using analytical tools. More complex signal processing methodologies are required for robust and scalable solutions. The complexity depends on the approach but the hardware systems can be made much simpler. With simpler hardware requirements, it becomes more compelling to produce a low-cost product and such a product may even be implemented as a low-cost modification to current smart meter systems.

4.1 Introduction

There are extensive load disaggregation studies in the literature [AGSA13, ZGIR12]. The goal of load disaggregation is to separate the individual appliance-level consumptions from the aggregate signal. The approaches generally use feature matching to do disaggregation. More specifically, a database of signatures is first constructed or trained using the features of different appliances, and then new signals of unidentified appliances can be compared in the database to find the best match. Various methods (e.g., clustering [Har92], pattern recognition [FZ99], genetic algorithms [BV04], neural networks [SNL06], hidden Markov model [KJ12], etc.) have been proposed using different features (e.g., real/reactive power, voltage, current, transients, harmonics, etc.). The existing studies are briefly summarized below.

Steady-state power is one of the most commonly-used features as it is simple to obtain without high sampling rates. Hart conducted one of the earliest studies of nonintrusive monitoring [Har92]. His system used customized hardware to collect 1Hz real power and reactive power. Steady-state power changes are used as the features to construct the database and classification is done through matching in the 2-D (real power-reactive power) space. 86% accuracy is reported for ON/OFF appliances (¿150W). Similar methods are proposed using different steady-state features [Sul91, BV04, MZ00, FZ99] or harmonics of steady-state power or voltage [SNL06, $LLC^{+}03$, CA00] to detect lower-power and more complex loads. Some others extend the linear appliance model to deal with continuously-varying appliances and multi-state appliances [DMXF12].

Another category of feature used in load disaggregation is transient characteristics, i.e. the fluctuations in power or current associated with the on/off events. Norford and Leeb conducted one of the first studies to use transients in load disaggregation [NL96]. It uses the shape of trainset events in the spectral envelopes as the appliance signatures. Other transient features such as energy during a transient event, power spikes or overshoots, and voltage noises are also used for load disaggregation [SLNC08, CYL08, CA98]. The main limitation of these approaches using transient features is that they generally require high sampling rate and the features may be appliance specific and sensitive to the wiring.

There are a few studies that aim to combine different features and algorithms together to achieve better overall performance. These load disaggregation platforms use committee decision mechanism (CDM) to obtain the disaggregation results from multiple prior algorithms and pick the best one by different mechanisms such as maximum likelihood, most occurrence, and minimum residue [LNKC10a, LNKC10b]. Unfortunately, most of the proposed load disaggregation methods in the literature are implemented in laboratory settings using customized hardware, which may not be suitable for real-world applications.

Therefore, this study focuses on developing a load disaggregation scheme leveraging the currently available smart meters for real households. Sampling rate is one of the main factors in load disaggregation that impact its applicability in the real world as it is limited by the on-site sensing hardware. Based on the frequency used in load disaggregation, the prior research can be divided into three categories: low frequency (periods between 1 hour and 5 minutes), high frequency (once per minute to 2 kHz), and ultra-high frequency (greater than 2 kHz). Different appliances are recognizable at different frequencies. Generally speaking, higher frequencies can provide more features (e.g., transients, harmonics) for the disaggregation algorithm and thus can recognize more appliances and improve accuracy but at a higher cost. Complex and expensive hardware is needed in order to obtain ultra-high frequency data. On the other hand, the widely deployed smart meters already provide certain capabilities for doing load disaggregation. It is cost effective and more feasible to leverage the existing smart meters instead of installing completely new and expensive sensors in order to achieve load disaggregation. Currently, the utilities are collecting low frequency smart meter data (1hr to 5min). Although it is possible to use the 1hr to 5min data directly for disaggregation [KBN10], it is desirable to use more fine-grained data within the capabilities of the smart meters for better accuracy. Fortunately, the smart meters deployed in the field can provide 1Hz real power data without hardware modifications [AGSA13].

Therefore, in this study, the focus is on using the 1 Hz real power data for load disaggregation. The study by Kolter et. al is the most relevant work to the study as it utilizes hourly metering data of 590 real houses for more than two years [KBN10]. It uses machine learning algorithms to predict the individual appliance consumptions of the houses that are not in the training set, which is different from predicting the consumptions of the same appliances in a single house after the training process. It is able to disaggregate into 10 broad general categories and achieve an overall accuracy of 55%.

4.2 Load Disaggergation Methodology

The objective in this work is to identify appliance events in terms of a positive change in power and subsequent negative change in power, which correspond to the appliance On state transition and Off state transition, respectively. These power changes will be referred to as as edges. These edges must be extracted from an aggregate energy signal which contains information from a multitude of appliance events. Extracting representative edges is challenging since appliances exhibit a great deal of variability in their power consumption patterns. The positive and negative edges must then be matched, potentially among a multitude of edges resulting from overlapping events.

One problem is noise which may manifest as gaussian noise or as burst noise of variable duration. Another problem is edge asymmetry which can result from input surge current, variable On state power consumption, and slow state transitions. Input surge current is a large instantaneous current drawn when an electrical device is first power on and is typical among electric motors, for example in refrigerators and HVAC systems. Additionally, many appliances employ feedback control which causes power consumption to change over time as the system reaches its set point. With slow state transitions, it may be unclear which sample measurement best represents the transition and this likely results in assymtery with the corresponding Off transition which consequently reduces matching accuracy.

Another problem is overlappinig events. The simplest case for event detection is a single event that occurs in isolation. In this case, it is trivial to identify the boundaries of the event and to compute the energy consumed. When events overlap and especially when different events have similar power consumption, it may not be clear which positive and negative edges should match. This may, for example, match edges too far apart in time resulting in overestimated energy conusmption calculation. Furthermore, mismatching may leave some edges unmatched which events unaccounted for.

The rest of this section addresses these challenges and details the algorithm design.

The approach is outlined as follows:

1. Preprocessing
   - Signal Imputation
   - Median Filtering
2. Edge Characterization
   - Edge Detection
   - Edge Concentration
   - Edge Combination
3. Edge Matching
   - Bipartite Formulation
   - Hungarian Matching
4. Event Classification
   - Feature Extraction
   - Load Classification
5. Signal Reconstruction

4.2.1 Signal Imputation

To measure current and voltage inside the apartment's electrical distribution panel, the system uses a wireless metering device to low-power radio. Apartments tend to be environments with many walls and obstructions resulting in multipath interference. In addition, there is interference from other wireless devices including the other energy monitoring systems installed in neighboring apartment units. The low-power radios used by the energy meters operate on the 2.4 GHz industrual, scientific, and medical (ISM) radio band, which is shared by wireless routers, cordless telephones, and even microwaves ovens. In fact, gaps in data acquisition were observed to correspond to times of microwave operation.

Missing data is a common problem in statistics and signal processing, particularly when working with sensor data. Algorithms may not be robust to the gaps in a signal. It is important to note that signal fidelity and problems of missing data are not universal. Rather, they are dependent on the data acquisition methods used and could thus be avoided. For example, smart meters installed at the utility meter may suffer no data loss due to minimal communication interference and reliable data transfer protocols. Such argues for the benefit of single-point energy measurement at the mains and also for the need for load disaggregation algorithms.

For the purposes of algorithm development, missing data is something that must be addressed to demonstrate the highest possible accuracy in load disaggregation. Imputation is the process of replacing missing data with substituted values. There are many ways to impute data, depending on assumption about the nature of the signal. One way to impute data is to perform simple linear interpolation between known data points. While linear interpolation is acceptable for steady-state power consumption, it does not satisfy assumptions about the transient properties of appliance events and thus a more sophisticated algorithm is needed.

The approach, outlined in Algorithm 4.1, can be summarized as that of determining whether the transition is postive or negative and whether it occurrs at the beginning or end of the gap. It begins by finding gaps in time greater than a threshold, chosen to be 10 seconds. The normalized change in power, s is then calculated for a small window around the missing data points. The total normalized change in power either before or after the gap (but not both, hence the use of xor) should be approximately 1. This indicates that there is a power difference before and after the gap but also accounts for gaps that occur in the middle of a transition. The change in power before and after the gap determines whether a concave or convex form is imputed at the edge, that is, what power the imputed point should take.

Algorithm 4.1: Impute Missing Data

Input: x, y

Output: x_{imp}, y_{imp}

for $i = 3 : m - 2$ **do**

  if $dx_i \geq 10$ **then**

    $s = \frac{dy_t}{dx_t}, t \in [i-2, i+2]$

    $s = \frac{s}{||s||}$

    if $(round(s_1 + s_2) == 1) \oplus (round(s_4 + s_5) == 1)$ **then**

      if $|y_i| < |y_i|$ **then**

        $j = \arg\min_t \{s_1, s_2\}$

        insert(y, y_{i+j-3})

      else

        $j = \arg\min_t \{s_4, s_5\}$

        insert(y, y_{i+j-1})

This is illustrated by the examples in Figure 4.1. For comparison, linear interpolation produces the profile with the dotted lines when in fact the microwave power consumption, like that of most appliances, exhibits a profile that approximates a step function shown by the thin line. The algorithm imputes gaps on both positive and negative state transitions equally well.

It should be noted that the algorithm is limited by the input signal fidelity and may still generate errors in reconstruction. If there is a transition that is completely missed such that the power is observed to change from one steady state power level to another over some gap of time, it is not possible to know with certainty whether the transition occurred at the

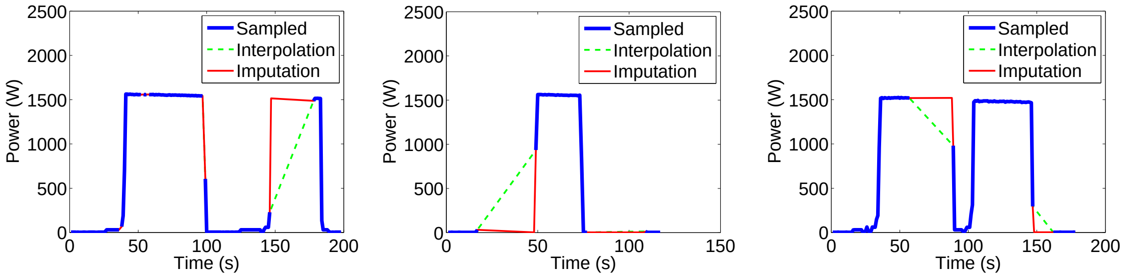

(a) Positive edge, convex (b) Positive edge, concave (c) Negative edge, convex and concave

Figure 4.1: Example of missing data from several microwave events and imputation results.

start of the gap, the end, or somewhere in the middle. However, this is acceptable since the objective at this step is merely to produce the most suitable representation of the signal for subsequent load disaggregation steps while accounting for as much of the observed information as possible.

4.2.2 Median Filtering

The median filter is a common signal processing technique used to reduce noise while preserving edges. For each sample, the filter computes the median value for a ”window” of samples around the target sample. The window must be selected carefully as there will be a tradeoff between smoothing and complete removal of events that are brief and thus appear to be noise. For example, if the window is 10 seconds, appliances like the garbage disposal or the light from opening the refrigerator door, may be filtered since these events are typically short in duration. While this is undesirable, loss of such information may be acceptable to an end user since idenfitication of such events is not likely to be as meaningful for behavior as knowing how much energy the heating and cooling system consumes.

4.2.3 Burst Noise Filtering

Most modern electronic devices use switched-mode power supplies (SMPS). These devices often produce negative burst noise that appears as step changes of variable duration. Applying an edge detection algorithm on these burst noise signals can produce a large number of edges which do not correspond to appliance state transition and can lead to errors in matching. To filter the noise, the first step is to compute the change in power over time $P'(t) = \frac{dP(t)}{dt}$. From this, the start of a noise episode can be detected if $P'(t)$ is less than some threshold. A window, Δ_{burst}, can be expanded to search for the boundary using several conditions. If the total change in power over the window, $\int_{\Delta_{\text{burst}}} P'(t)$, is less than some threshold (i.e. if the power decreased and then returned to the initial level), then if this window contains burst noise, $P_{\Delta_{\text{burst}}}(t)$ is replaced with the linear interpolation between the window boundaries, $interp(P(\Delta^-_{\text{burst}}), P(\Delta^+_{\text{burst}}))$. Conversely, if the previous condition has not been met and the difference between the power at the boundaries $P(\Delta^+_{\text{burst}}) - P(\Delta^-_{\text{burst}})$ is greater than some threshold, or if the window exceeds some size, then the search is exited.

Figure 4.2 demonstrates the filter applied to a section of a plug load signal. This approach performed well for the burst noise observed in the data. However, fast switching events (such as a light switched being flipped On and Off repeatedly) would generate many valid edges that would be removed by this filter. Fortunately, such behavior is not typical and the loss of this information should be accepable for most user applications.

4.2.4 Edge Detection

A simple first-order difference function is used to detect edges. Differencing calculates the change in power $P'(t)$ which provides an indication that an edge may be present.

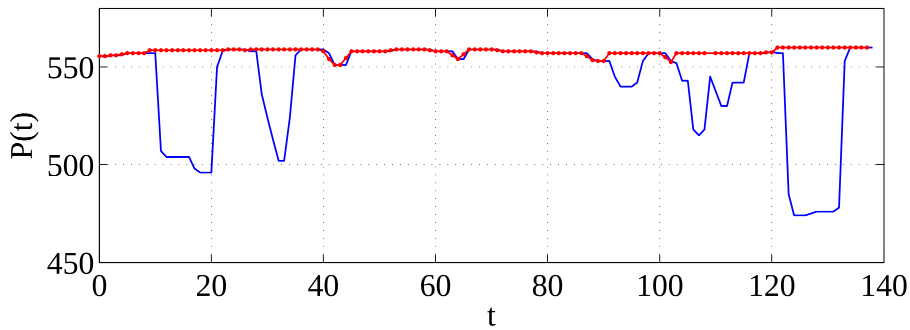

Figure 4.2: Example of burst noise filter smoothing.

However, this is not sufficient to systematically characterize state transitions in a way that yields accurately matched edges, as will become evident. A series of refinements must then be performed which will be described in the next steps.

4.2.5 Edge Concentration

For slow state transitions, the total change in power is divided over the samples comprising the transition period. In fact, the speed of the transition is relative to the sampling frequency. For example, a 2-second transition could be captured by a single sample at 0.5 Hz while a 0.5-second transition would be captured over 5 samples at 10 Hz. In any case, when a state transition occurs over multiple samples, no single sample represents the full power change. asymmetry between the corresponding positive and negative edges which will lead to inaccurate matching.

This problem is addressed by concentrating (i.e. integrating) the power change within the state transition window Δ to form an impulse representing the total change in power $\int_{\Delta} P'(t)$. First, the boundaries of Δ are determined by detecting if a state transition has occurred if the change in power $P'(t)$ exceeds a threshold. Then, using a sliding window w, if $\int_{w} P'(t)$ is less than some threshold or if a zero crossing is detected then the boundary has been found.

Figure 4.3 illustrates the edge concentration results. Figure 4.3a shows a single edge concentration in detail. Figure 4.3b shows the change in power, $P'(t)$, for two consecutive microwave events with the solid points indicating measurements detected within the transition window, Δ. The circles represent the total change in power, $\int_\Delta P'(t)$. Concentration, in Δ_{On} and the minimum change in power from Δ_{Off}. After concentration, $\int_{\Delta_{\text{On}}} P'(t) \approx \left|\int_{\Delta_{\text{Off}}} P'(t)\right|$.

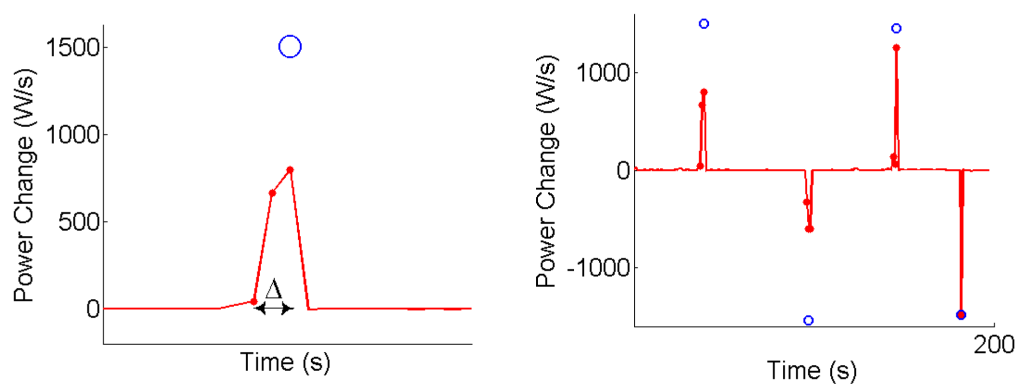

(a) (b)

Figure 4.3: Example of edge concentration on consecutive microwave events.

4.2.6 Edge Combination

Edge combination addresses the problem created by superfluous transient behavior that is not removed by the median filter. Often with appliances like refrigerators and HVAC systems, there is a large surge current during the motor start-up. Edge detection on these transient phases results in a large positive edge followed by a negative edge representing the end of the surge. Taken directly, the positive edge does not a closely match with the negative edge and may not generate a matched event in later stages. When a positive and negative edge are detected in close proximity, there is a high likelihood that they represent power changes due to surge current. The combination of the positive and negative edges produces a third edge whose power is more closely aligned with a subsequent negative edge. Figure 4.4 shows an example of an HVAC cycle with a large input surge current. The positive edge at 1868 W has been added as a combination of the peak power of 2905 W due to surge current and the subsequent drop in surge current of -1037 W. This 1868 W edge is a closer match with the later negative edge of -1918 W.

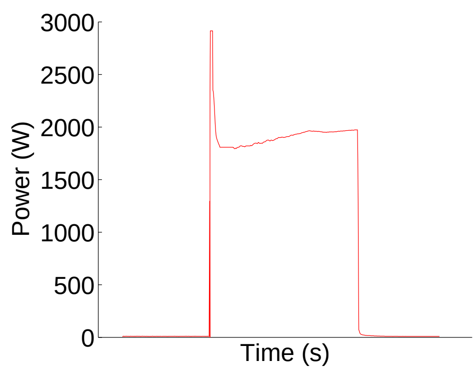

(a)

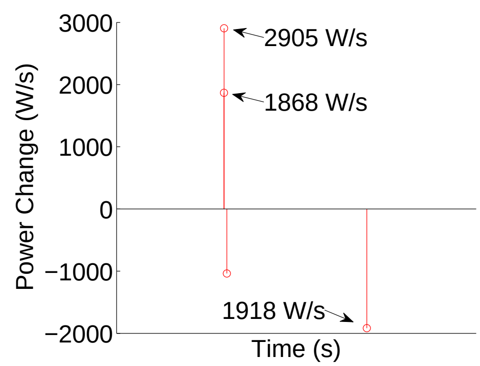

(b)

Figure 4.4: Example edge combination on HVAC event. Figure 4.4a shows the power signal and Figure 4.4b shows the edges.

4.2.7 Bipartite Formulation

After edges have been extracted from the signal, the positive and negative edges need to be matched to so appliance events can be bounded and characterized. The matching problem is formulated as the minimum-cost matching of a weighted bipartite graph. A bipartite graph is a graph whose vertices can be divided into two disjoint sets.

In this case, there is a set of positive state transition edges U and a set of negative state

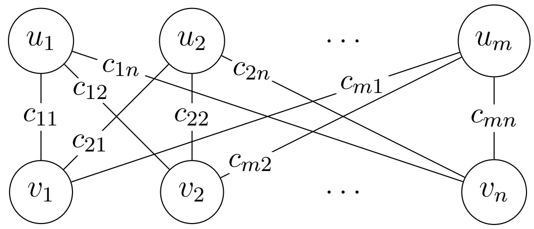

Figure 4.5: Edge matching problem represented as a bipartite graph. Positive and negative transition edges are represented as graph nodes u and v, respectively, and costs are represented as graph edges c.

transition edges V. As shown in Figure 4.5, each transition edge $u \in U$ and $v \in V$ is represented as a vertex in the bipartite graph. A weight is assigned to every graph edge that represents the cost of matching the two transition edges. The total cost c is based on a time cost τ, a power cost δ, and an energy cost ϵ.

The time cost is the duration of the event given by

$$\tau_{uv} = \begin{cases} log(t_v - t_u) & 0 < (t_v - t_u) < T_{max} \\ \infty & \text{otherwise} \end{cases} \tag{4.1}$$

where the duration should be less than T_{max} and Off-On events are disregarded by assigning infinite cost.

The power cost is the difference between the magnitudes of the positive edge power and negative edge power and is given by

$$\delta_{uv} = \begin{cases} \left(\frac{P_u + P_v}{P_u}\right)^2 & \rho_{\text{min}} < \frac{P_u}{|P_v|} < \rho_{\text{max}} \\ \infty & \text{otherwise} \end{cases} \tag{4.2}$$

where events whose power difference exceeds some bounds indicated by ρ_{min} and ρ_{max}, are disgregarded.

To compute the energy cost, the approximate energy of the feasible match $\widehat{E}_{uv}$ is calculated as the product of the estimated average power level $\frac{1}{2}(P_u + |P_v|)$ and the event duration $(t_v - t_u)$. the measured energy E_{uv} is calculated by integrating the measured power over the event window.

$$\widehat{E}_{uv} = \tfrac{1}{2}(P_u + P_v)(t_v - t_u) \tag{4.3}$$

$$E_{uv} = \int_{t_u}^{t_v} P(t) \tag{4.4}$$

To calculate the energy cost ϵ_{uv}, a constraint is imposed where the actual power level cannot be less than 50% of the estimated average power for more than 10% of the event time. This is denoted using Iverson bracket notation which converts the boolean value from the condition inside the brackets into an integer value and allows counting to be represented as a summation. The 50% and 10% thresholds were selected to allow for some signal noise, measurement error, and algorithm performance limitations. In this case, the match is infeasible. Otherwise, the energy cost is the absolute difference between the estimated and measured energy.

$$\epsilon_{uv} = \begin{cases} \infty & \big(\sum_{t=t_u}^{t_v} [P(t) < \tfrac{1}{2}(P_u + P_v)]\big) > \tfrac{1}{10}(t_v - t_u) \\ |\widehat{E}_{uv} - E_{uv}| & \text{otherwise} \end{cases} \tag{4.5}$$

The total cost of a feasible match as

$$c_{uv} = \tau_{uv} + \delta_{uv} + \epsilon_{uv} \tag{4.6}$$

4.2.8 Hungarian Matching

The costs are represented as a matrix C of edge weights between the positive edges U and the negative edges V.

$$\begin{array}{cc} C & \begin{array}{cccc} v_1 & v_2 & \dots & v_n \end{array} \\ \begin{array}{c} u_1 \\ u_2 \\ \vdots \\ u_m \end{array} & \begin{pmatrix} c_{11} & c_{12} & \dots & c_{1n} \\ c_{21} & c_{22} & \dots & c_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ c_{m1} & c_{m2} & \dots & c_{mn} \end{pmatrix} \end{array}$$

The Hungarian algorithm[1] is used to obtain the minimum-cost matching of On and Off edges which represent the most likely bounds of each appliance event.

4.2.9 Feature Extraction

The classification stage of the load disaggregation algorithm uses the following features:

- P_+: On power level
- P_-: Off power level
- T: Event duration
- P_S: Input surge current level

[1] The Hungarian algorithm is combinatorial optimization algorithm which solves the assignment problem. It was originally developed by Harold Kuhn [Kuh55] and James Munkres [Mun57] and had a time complexity of $O(n^4)$. It has since been modified to achieve $O(n^3)$ running time [EK72, Tom71].

These features are easily extracted from the original power signal based on the matched event data as shown in Figure 4.6. While these features are convenient, they are limited in their ability to classify as shown in Figure 4.7. Although some appliance, like the HVAC and refrigerator are well separated in feature space, other events like lights are mixed with many of the plug-in devices. It is not possible to determine from this dataset what appliances the plug load is composed of since that ground truth does not exist. Based on earlier studies, it is suggested that classifying these appliances can benefit from additional metrics such as reactive power or waveform measurement.

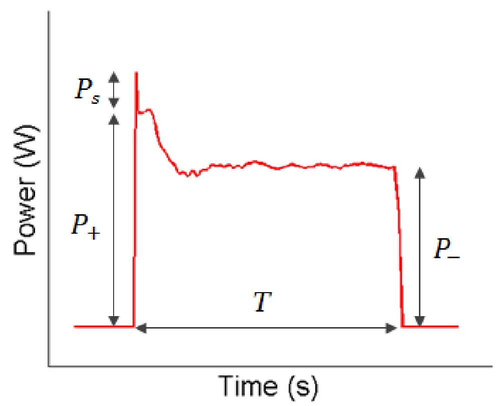

Figure 4.6: Extracted features for HVAC event.

4.2.10 Classification

For classification of the appliance loads, a training set is first constructed using sample datasets from a percentage of randomly selected apartments. The data provides ground truth energy consumption for several individual appliances. The load disaggregation algorith can be applied to characterize appliance events from the disaggregated appliance signals. Classification on test data is performed using a k-nearest neighbors search with

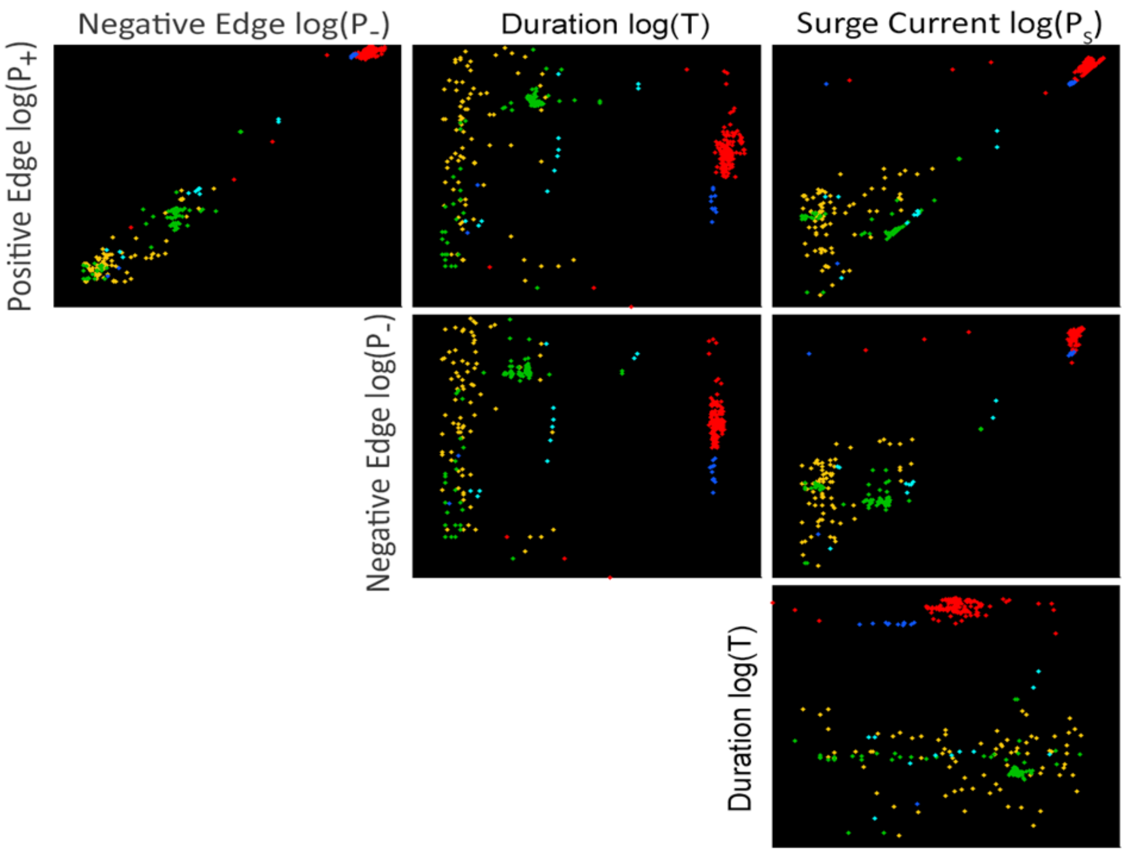

Figure 4.7: Feature space correlations (axes are in log scale).

the training set.

4.2.11 Reconstruction

Reconstruction is a multi-stage process that begins with filtering. The event matching and classification algorithms that precede reconstruction may identify false positive events. Reconstruction tries to eliminate these by evaluating the distance metrics from classification for each appliance. The idea is that a single appliance may not exhibit events that overlap in time. For example, the refrigerator may not have two commpressor cycles overlapping. This is a reasonable assumption, as long as ground truth exists for individual appliance event types for classification. Taking the refrigerator as a further example, an interior light event may overlap with the compressor cycle but the refrigerator only has one compressor so there should be no overlap. If there are overlapping feasible matches for a single appliance, the algorithm removes matches which have a larger distance metric.

To reconstruct the power consumption signal, the algorithm first evaluates the ratio of the estimated energy consumption defined by the state transitions with the raw energy of the event window. If the estimated energy is approximately equal to the original energy (total energy signal), then the original signal is used for reconstruction. Otherwise the estimated signal is used, calculated as the linear interpolation of the power levels for the edges (although other interpolation models may be used depending on known appliance dynamics). Figure 4.8 illustrates the reconstruction of the refrigerator power consumption where the event in the middle overlapped with other appliance events such as the HVAC or microwave. For that event, taking the raw power consumption signal would have included power consumption for these other appliances and may not have been an accurate representation depending on the magnitude of energy consumption from the overlapping appliance events. The events on the side were isolated without overlapping events from other appliances and were thus reconstructed from the original signal.

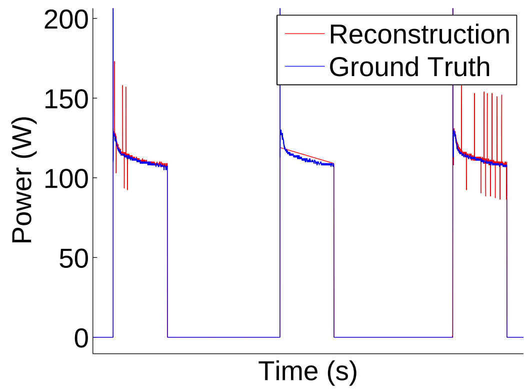

Figure 4.8: Example of refrigerator power consumption signal reconstruction.

4.2.12 Process Segementation

The code is structured to process chunks of data delimited by the periods that are likely to have only baseline power consumption. The baseline power P_{base} can be learned by observing the total power consumption over hours or days and calculating the mode of the signal. This is based on the assumption that appliances are turned off more than they are turned on and there should be ample periods where no appliances are turned on or are always on. While simplistic, this approach was validated experimentally. Segementing the processing in this way enables the algorithm to generate results in in pseudo-real time since the data set will contain only the most recent data since the last time of baseline power consumption. It also results in improvements in performance, since the data set is kept relatively small rather than days or weeks, as well as accuracy, since collection of partial events will be avoided.

4.3 Results

The data set consists of over 1 Hz power signals for 20 apartments and over 2 apartment-years of data. Figure 4.9 shows the energy estimation for HVAC, refrigerator, and microwave appliances averaged across 15 apartments using a training data set of 1, 3 and 5 apartments for classification. The refrigerator reconstruction is very accurate and converges as the training set is increased. This is due to the uniformity of the refrigerators across apartments. The HVAC is accurate but diverges as training is increased. Although the exact reason is unclear, one hypothes is the HVAC power consumption is based more on user behavior than other appliances and there is more variation across apartments, so a larger training set is required before the variance converges. The microwave accuracy should be high since the event profile for microwaves is unique and earlier tests had indicated high accuracy in microwave classification and reconstruction. However there is likely some parameter in the algorithm that is corrupting the process for microwave events and will require tuning.

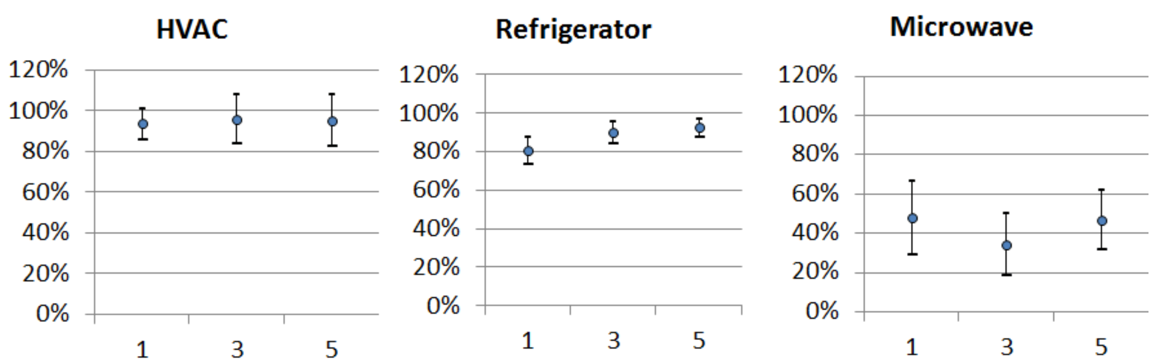

Figure 4.9: Energy esimations and confidence intervals from reconstructions.

4.4 Discussion

The time-domain load disaggregation method outlined in this chapter can result in errors for complex sequences of highly heterogenous events. A better approach might use a dynamic programming adaption of the Hungarian algorithm to recursively approximate the error and determine an optimal matching over an expanding subset of edges. Further, the algorithm currently leverages only real power and temporal dynamics whereas reactive power provides an orthogonal dimension which can improve detection accuracy. More to the point, the data contains disaggregated information for major home appliances including the refrigerator, microwave, recessed lighting, and HVAC but does not disaggregate plug-in devices like laptops, computers, temlevisions and other home electronics. For such appliances, it may be necessary to include additional features, such as higher order harmonics or complex power, in order to correctly classify them.

Another limitation is the dataset which does not include ground truth information plug-in devices including consumer electronics and kitchen appliances. This is a result of the infrastructure of the typical apartment, where the electrical infrastructure tends to source wall outlets from just a few circuits. Without appliance-level instrumentation, it is difficult to acquire ground truth on appliance-level power consumption.

A third limitation is identification of clustered appliances which are fed by a power strip. In this case, the appliances could all be turned on at the same time but turned off individually. Similarly, devices with multiple, variable operating states such as computers, fans, and hair driers may escape event detection scheme. In some of these cases, expanding the combinatorial possiblities for matching (by splitting rather than combining edges) may improve robustness. For other appliances like dishwashers and clothes washing machines with multiple but consistent operating states, the event detection approach would still be able to isolate the discrete cycles and then sequence matching could be used for classification.

4.5 Conclusion

This chapter presented a load disaggregation algorithm using time-domain analysis of energy consumption data sampled at low frequency (1 Hz). Although futher testing is required to refine the algorithm paramters, it has the advantage of potentially being able to achieve immediate and large-scale application for existing smart meter deployments, while other techniques require additional hardware installation.

CHAPTER 5

Post-Operative Remote Patient Monitoring and Guidance

The previous chapters demonstrated behavior guidance applications in residential energy consumption. Healthcare faces similar challenges with inefficient operations leading to avoidable cost. Failure to detect changes in patients' post-operative health status increases the risk of adverse outcomes, including readmission. While inpatient monitoring after surgery is routine, there is no current method for detecting complications arising between discharge and first clinic visit.

This chapter presents the design and implementation of a real-time monitoring system for post-operative colorectal surgery patients utilizing wireless health technology. Participants were assigned a pre-programmed tablet computer after surgery between post-operative days #1-3on post-operative day #1, and asked to complete a daily survey regarding their post-operative health until their first clinic visit. Completed surveys were then transmitted wirelessly to a secure database for review by the clinical team. As a pilot study, the clinical team monitored the first 33 consecutively-enrolled patients for a total of 390 patient-days: 134 out of 181 (74%) in the inpatient setting and 120 out of 209 (57%) after discharge. Overall patient compliance was 65% (data available for 254 of the 390 days), but varied by patient from 32-100% of monitored patient-days. The system was able to reliably collect basic data on post-operative health status as well as patient-reported outcomes not previously captured by standard assessment techniques. Qualitative data from multiple respondents suggest that the experience strengthened their relationship with their surgeon and aided in their post-operative recovery. The team found that remote monitoring of post-operative patients using wireless technology is feasible, and provides more detailed and complete information to the clinical team. Wireless health technology represents an opportunity to close the information gap between discharge and first clinic visit, and, eventually, to improve patient-provider communication, increase patient satisfaction, and prevent unnecessary readmissions.

5.1 Introduction

Patients undergoing elective colectomy are at increased risk for post-operative complications and readmission [GMA$^+$07, WSH$^+$11, KOP$^+$12, LHL$^+$13, LHL$^+$14]. With financial incentives and fast track clinical pathways leading to shorter lengths of stay, many complications that would previously have occurred in the hospital are developing, potentially unnoticed, at home [DFS$^+$01]. As many surgical readmissions are related to modifiable risk factors, including dehydration and surgical site infection, opportunities for intervention may exist if complications can be identified and addressed in a timely manner [DSR$^+$14, SDR$^+$14]. The vast majority of hospitals, however, have no current method for identifying complications arising between discharge and first clinic visit.

Wireless health technology has emerged as a way of capturing data not previously available to the clinical team. Although initially more common among medical patients, particularly in the treatment of diabetes [RMJS$^+$09] and heart failure [LWJ$^+$01] wireless technology is expanding into peri-operative care, particularly among cardiothoracic and transplant patients [YGHF08, PHL$^+$08, BZST09, DHP$^+$12]. Modern devices may be equipped to transmit complex information regarding patients' mobility [CTP$^+$13, AWL$^+$14], wound healing [VTG+08], or physiologic function, such as electrocardiographic wave forms and spirometric flow rates [DHP+12, MKPE02]. Beyond simply capturing patient-reported health outcomes, many have suggested that wireless technology could be used to influence behavior and help guide their path toward recovery from surgery.

PRIME (Patient Remote Interactive Monitoring Enterprise) is a system for developed for post-operative patient condition monitoring. The primary mechanism in PRIME is patient-reported responses to standardized survey instruments covering many dimensions of health. PRIME surveys are based on PROMIS v1.0 instruments which have been validated in multiple NIH-sponsored studies for depression, back and leg pain, chronic obstructive pulmonary disease (COPD), congestive heart failure (CHF), and rheumatoid arthritis (RA). The PROMIS®instruments divide health assessment into a comprehensive array of domains such as alcohol use and pain interference. They additionally provide targeted survey instruments for pediatric assessments.

Different surgical service lines have different risk factors for readmission. For colorectal surgery patients, dehydration is a leading complication and thus there are measures focused on capturing fluid intake and urine output. For bariatric surgery patients, monitoring blood glucose levels is critical. PRIME allows survey instruments to be adapted to address the specific needs of a service to customize the language or the parameters. Instruments can be added or removed as needed and new instruments can be created.

Functions can also be customized to address usage. The alarm function generates an audible ring to alert the patient to complete a task. The colorectal service line has determined 10AM as the best time to alert their patients.

5.2 System Architecture

A mobile application development framework was used to build the PRIME mobile application. This framework enables cross-platform compatibility by supporting application development using markup languages (HTML and CSS for interface structure and style) and programming languages (Javascript for logic) rather than native code such as Java for Android and Objective-C for iOS.

Figure 5.1 illustrates the PRIME application architecture. The Application Framework renders the PRIME application to the user via a WebView. This Webview communicates to the Mobile Operating System and access device features such the camera, GPS, network, and filesystem through a plugin library and platform specific API's. The system uses custom plugins to extend functionality including a task scheduler and usage log. The task scheduler notifies the patient to perform a requested action such as to take the survey, take a photo of the wound, or take medication. The usage log records when and how the patient has used the system and would be useful not only for verification of compliance but for analysis of weaknesses in user experience. For example, many users spend a long time with a given survey question indicating that the question may not be clear or that the input mechanism is not responsive.

All data is uploaded via the network interface to the PRIME server for storage, analysis, and reporting. The features include the survey designer, task scheduler, usage analyzer, and dashboard for patient status reports. The PRIME server also interfaces with external medical data enterprises such as CareConnect which houses electronic medical record (EMR) data and RedCap which is used for clinical trials management.

A dashboard interface was developed for both care providers and patients to access data and interface with the application. An example dashboard report is shown in 5.2.

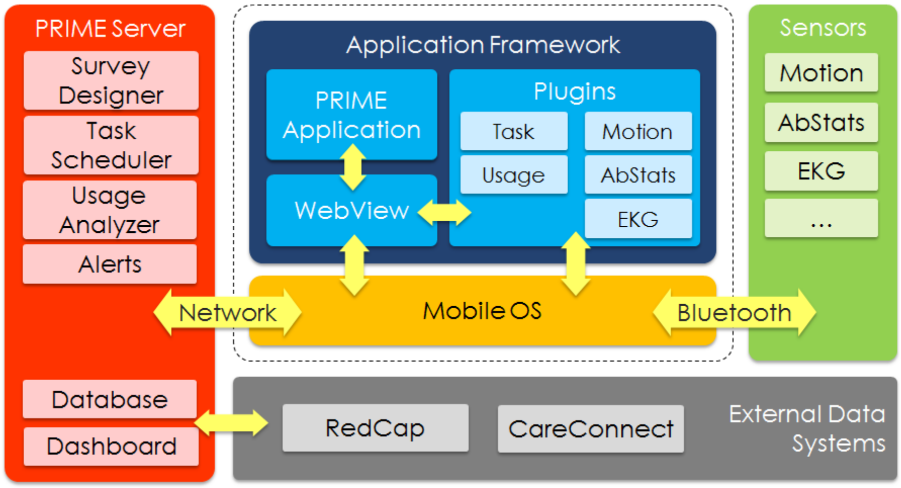

Figure 5.1: PRIME software architecture and system interaction between wearable sensors, mobile app, backend, and external party data system.

Patterns can be visualized to show trends, such a patient having high pain levels on post-operative day 1, shortly after surgery. It was known from the medical record that the patient took pain medication which caused the pain levels to decrease until after discharge (indicated by the pink-shaded day) when the pain increased again. The pain levels did not warrant medical attention but such information can be used in an alert system to notify care providers that an intervention may be required.

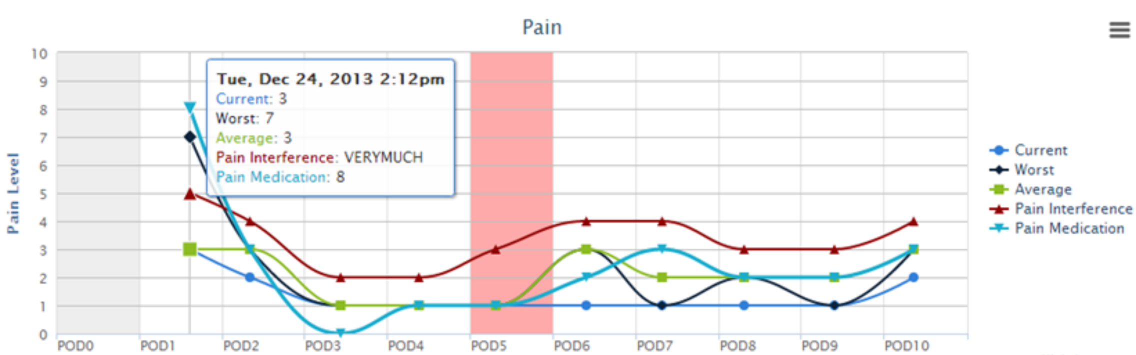

Figure 5.2: Patient-reported pain measures over post-operative period.

The dashboard also presents full summaries of the patient self-report responses on each day. An example is shown in Figure 5.3. These summaries show the trends for every self-report measure along with highlighting for responses that satisfy physician-defined alert conditions. This allows for rapid asssessment of patient condition and recovery trend.

QuestionID	Question	2014-10-09	2014-10-10	2014-10-11
pain-1	What is your level of pain right now?	3	2	5
pain-2	How intense was your pain at its worst during the past 24 hours?	7	7	5
pain-3	How intense was your average pain during the past 24 hours?	5	4	3
pain-4	How much did pain interfere with your activities during the past 24 hours?	QUITEABIT	SOMEWHAT	ALITTLEBIT
fatigue-1	How fatigued were you on average during the past 24 hours?	SOMEWHAT	SOMEWHAT	ALITTLEBIT
fatigue-2	To what degree did your fatigue interfere with your physical functioning during the past 24 hours?	SOMEWHAT	SOMEWHAT	ALITTLEBIT
fatigue-3	How was your sleep quality in the past 24 hours?	FAIR	FAIR	FAIR
fatigue-4	Were you satisfied with your sleep in the past 24 hours?	ALITTLEBIT	SOMEWHAT	ALITTLEBIT

Figure 5.3: Example PRIME patient self-report summary.

The PRIME application can interface with a variety of Bluetooth sensors to monitor human motion, abdominal sounds, and heart rate, for example. Communication to and control of these devices are implemented via custom application plugins.

5.3 Results

Table 5.1 summarizes the patient characteristics from the first 33 patients. The majority of the sample was male (52%) with a median age of 56 years. The most common operation was low anterior resection (40%) followed by right/subtotal colectomy (34%); one-quarter of cases were performed laparoscopically. Nearly all patients underwent an operation for cancer. The average length of stay was 7 days and the median time from discharge to first clinic visit was 14.5 days.

Table 5.1: Patient Characteristics (N=33)

	N (%)
Age-Mean (Std. dev.)	56 (10.9)
Male Gender-N (%)	17 (52%)
Length of stay-Mean (Std. dev.)	7 (4.3)
Type of surgery-N (%)	
Low anterior resection	14 (40%)
Right or subtotal colectomy	11 (34%)
Abdominoperitoneal resection	5 (17%)
Colostomy revision or takedown	3 (9%)
Laparoscopic	17 (52%)
Indication for surgery-N (%)	
Cancer	28 (82%)
Polyp	2 (12%)
Ulcerative colitis	1 (4%)
Rectal prolapse	1 (4%)
Benign disease requiring ostomy	1 (4%)

5.3.1 Patient Compliance

There is currently no widely accepted method for measuring patient compliance with electronic submission of patient-report outcome information. The question of what is ultimately effective in improving outcomes for mobile application engagement may be different for other measures such as pharmaceutical compliance which has fewer situational dependencies and drug regimens are

Compliance days were measured based on whether the patient completed any self-report tasks for that day, i.e. if they completed a daily survey or took a wound photo. Patients would typically start using the tablet between post-op day 1 and 3 compliance measurement started with the first instance of either a completing a survey or taking a photo. but in some cases there were delays in the provisioning the tablet or in the patient feeling well enough to use it. Since there is no standard for evaluating patients' ability to use a mobile device, it did not make sense to base

Over the pilot period, the 33 participants submitted data during 254 of the 390 monitored patient-days (65% compliance). Individual device use varied between 6 and 24 days, based on the time bestween operation and first clinic visit. Individual compliance rates for survey completion ranged from 32-100%. Compliance rates were higher during the inpatient time period (134 entries in 181 monitored days, 74%) than the outpatient period (120 in 209, 57%).

In addition, significant improvements were made during the pilot study to both the user experience design and workflow integration of PRIME for both patients and physicians. This led to improvements in patient compliance summarized in Tables 5.2 and 5.3. These changes included improvements to the user interface with better responsiveness and navigation. More importantly, patient alert summaries were generated on a daily basis and e-mailed to providers to provide a snapshot of each patient's responses and potential complicating factors. This targeted and actionable feedback enabled the providers to follow up with patients more effectively.

Table 5.2: Compliance Results for Patients 1-20

	Surveys	Days	Compliance (%)
Hospital	98	135	73%
Home	68	128	53%
Total	166	263	63%

Table 5.3: Compliance Results for Patients 21-33

	Surveys	Days	Compliance (%)
Hospital	36	46	78%
Home	52	81	64%
Total	88	127	69%

5.3.2 Examples of Survey Responses

Trends were examined in survey responses both for individual patients and for the sample as a whole, standardized by post-operative day. As an example, pain data is traditionally captured in the inpatient setting while patient satisfaction with their ability to perform light household duties is not typically addressed by standard care processes.

Figure 5.2, shown earlier in the system architecture, shows an example of the web-based interface used by the surgical team to track real-time patient responses. Pain was measured in five ways, as displayed in the corresponding legend: current pain level (1-10), worst level for the day (1-10), average level for the day (1-10), level of interference with daily activities (Not at all, little bit, Somewhat, Quite a bit, Very Much), and the use of pain medications (number of doses in past 24 hours). The shaded column represents the day of discharge. A similar display was available for each patient in each of the eight measured domains.

The information collected by PRIME can be used to generate new analytics. For example, it is possible to track average levels of satisfaction with physical functioning for the entire sample over the course of their post-operative care, shown in Figure 5.4. The percentage of patients reporting that they were not at all satisfied decreased steadily over the first week while the percentage who were somewhat or quite a bit satisfied increased during the same period. This is perhaps not surprising since patients tend to experience significant discomfort immediately following surgery but it suggests that perhaps greater measures can be taken to improve patients' experiences during this period. More broadly, other insights can be obtained with similar analyses of other recovery measures.

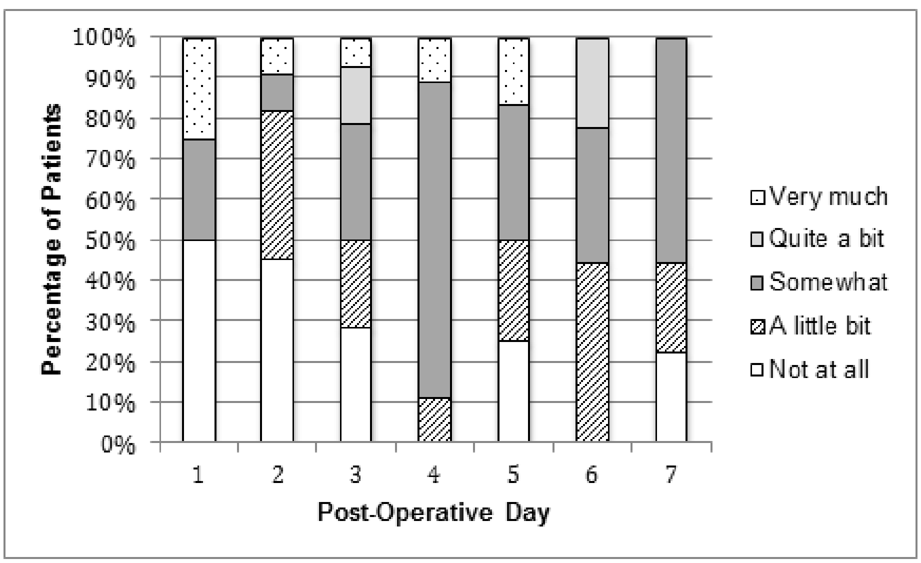

Figure 5.4: Average satisfaction with physical functioning by post-operative day.

5.3.3 Patient Experience

All five patients who completed the questionnaire reported that the device was helpful and made them more involved with their post-operative recovery. A summary of qualitative themes and example quotations is presented in Table 5.4. Multiple patients gave examples of post-operative instructions (e.g. maintaining adequate water intake) that were reinforced by using the device to record their performance. Multiple respondents suggested that the experience strengthened their relationship with their surgeon and aided in their post-operative recovery.

5.4 Discussion

Complications after abdominal colorectal operations may go unnoticed in patients who are discharged home without sufficient follow-up. It was found that remote home monitoring using wireless health technology is feasible, can include patient-reported outcomes not typically captured in the post-operative setting, and may even improve patient experience.

Although there is no widely-accepted cutoff for survey response rates, 60% or higher is generally cited as a rule-of-thumb for producing generalizable results [LW12]. Unfortunately, the overall compliance rate of 63% fell below this mark, with only slightly higher rates during inpatient monitoring. Much of this may be due to early participants receiving less attention and coaching while the device and network were actively being developed and improved. More than simply collecting reportable data, improving compliance represents an opportunity to increase patient engagement in the recovery process. Other studies have found that automated reminders, contracts with family members, and financial incentives can nearly double compliance with remote home monitoring program [SSR+14].

Table 5.4: Themes and examples of behavior change and patient experience using PRIME.

Theme	Example Quotations
Awareness of recovery process	• The ability to look at my wound daily made me more cognizant of my wound and to ensure it was clean. • The real benefit was being connected to my body. • I felt more connected to myself and aware of my healing process as a result of using the device.
Activation and participation in post-operative recovery	• The tablet was my daily guide. It reinforced the discharge orders I received from the hospital . . . when I left the hospital was when the tablet became most important. Just knowing you had to report all that you were doing, made you aware of what you were supposed to be doing. • I did a much better job paying attention to the output from my ostomy than I would have otherwise. The survey questions also reminded me daily that I need to be drinking more water. • One day I didn't feel like walking at all. But knowing I had to report it on the daily survey made me actually get up and walk.
Shared goals and understanding of recovery period	• When I called about skin irritation around the stoma they could see exactly what I was talking about. They noticed from the photos that my wound seemed to be taking a long time to heal and contacted me about how to pack it differently.
Sense of connection to the surgical team	• I felt when using the tablet as if I was in [the clinic] . . . it was IMPERATIVE to my health and recovery. • I was more aware of being monitored on a daily basis which in turn made me feel better connected to my health care providers.
Improved patient experience	• Knowing your doctor has access to you on a daily basis was extremely helpful. Invaluable, actually. • The doctors were really interested in learning what their patients are going through, which doesn't always get reported in office visits. • This was a very positive experience for me.

Automated reminders have since been activated for the patient to participate and anticipate improved responses. Future efforts should utilize these and other incites from behavioral economics to encourage patient participation in post-operative care.

Differences were found in both compliance rates and self-reported outcomes between the inpatient and outpatient setting. Compliance was much higher in the inpatient setting, which may be due to either reinforcement from the surgical team or fewer distractions in the hospital than patients typically face after returning home. While some information (e.g. pain level) is already captured during nursing or physician assessments, there is additional value for wireless health technology in the inpatient setting. First, learning to use the device prior to discharge may improve outpatient compliance by priming patients to the experience and incorporating data entry into their daily routine. Second, patient-reported outcomes may be more accurate or valid than those collected by a provider whose presence may intentionally or unintentionally affect the response. Finally, less traditional components of the patient experience (e.g. cognitive function) may not be adequately captured by current assessment techniques. Creating a simple, reliable conduit for patient-reported outcomes may improve the attentiveness of the surgical team to important but overlooked patient issues while increasing patient voice.

Finally, patients overwhelmingly reported a positive experience using the device, and described both improved communication with the surgical team and increased participation in their own post-operative recovery. Anecdotally, one patient was called into the clinic and eventually readmitted to the hospital after the surgical team recognized high and increasing ileostomy output from his daily survey responses; other patients reported that their survey responses facilitated a discussion with the surgical team about specific concerns that they may have otherwise neglected to mention. While the ability to provide real-time monitoring and feedback will undoubtedly improve the management of post-operative complications, the larger impact of wireless health technology on recovery may be its ability to influence patient behavior outside of traditional healthcare settings. Asch and colleagues describe the process as automated hovering: the use of technology to provide guidance and influence decision-making between doctor's office visits [AMV12]. Even in the small sample, patients reported being more vigilant about fluid intake, wound care, and ambulation because they knew they would have to report on these items in their daily survey. Therefore, incorporating wireless health technology into the standard discharge plan may allow the surgical team to guide the recovery process in ways beyond simply monitoring patients for early signs of complications.

The study had several limitations. First, the sample size in this pilot was small and drawn from a single academic institution. Second, the participants were pulled from a convenience sample of patients undergoing abdominal colorectal surgery. Patients were selected based on their self-reported familiarity with tablet devices, it is possible that a larger sample may not have used the device as readily or reported as positive a post-operative experience. However, the patient and operation characteristics were similar to the majority of patients undergoing colorectal surgery, and the first 33 consecutive participants were examined in order to increase generalizability. Third, both compliance and patient-reported outcomes varied across the sample, suggesting that the device may have a different impact in different subsets of the population. However, given the small sample size and the importance of unmeasured covariates (e.g. use of a tablet device at home), it was unable to be determined which groups may benefit the most from using the device. Finally, all data in the pilot phase were entirely patient-reported. Prior to larger-scale data collection efforts, there are plans to validate inpatient responses against other sources of data to ensure the work generates meaningful results.

5.5 Conclusion

These limitations notwithstanding, this pilot study demonstrates the potential of wireless health technology to impact recovery after elective colorectal surgery. Once impenetrable to the clinical team, the period between discharge and first clinic visit may soon seem like an extension of inpatient monitoring without the need for resource-intensive hospital days. By incorporating wireless health technology into standard post-operative recovery, it becomes much more feasible to improve patient-provider communication, increase patient satisfaction and engagement, and, eventually, decrease both adverse clinical events and unnecessary readmissions.

5.6 Mental Health Extension

Although PRIME has been developed for post-surgical patient monitoring and hospital readmissions is a critical problem, depressive disorders may represent the greatest challenge to global health in the 21st century:

- Depression is the single most common cause of disability worldwide, affecting more than 350 million people. By 2030, it will be the single largest contributor to the global burden of disease.
- Depression is the strongest risk factor for suicide. In 2010, worldwide deaths from suicide outnumbered deaths from war, natural disasters, and murder.
- Depression affects individuals of all ages and backgrounds. The symptoms make ordinary daily activities seem impossible.

- Depression worsens the outcome of other common diseases such as heart disease, cancer, stroke, Alzheimers disease, and Parkinsons disease.

- Depression has a devastating worldwide economic impact, including $116 billion in medical and long-term care costs in the U.S. alone in 2010. This sum does not take into account the lost productivity of affected individuals or the impact on their families and communities.

- The majority of depressed individuals do not seek or receive any treatment, in large part due to self-blame and social stigma.

- Existing treatments are inadequate; 50% of depressed individuals do not meaningfully respond.

Despite the enormous costs associated with the condition, research on depression is grossly underfunded world-wide. Of these costs, $26.1 billion dollars were direct medical costs, $5.4 billion were suicide-related mortality costs, and $51.5 billion were workplace costs. These figures do not include costs associated with co-existing psychiatric and medical conditions [GKB+03]. Even within the colorectal patient pilot study, strong correlations between mental health and post-discharge recovery have been observed. Anecdotal evidence suggests that patients who are non-compliant or who exhibit low compliance may be suffering from mental illness such as depression. This is not surprising, when considering these patients had been diagnosed with life-threatening disease. For many patients, the disease and the surgery are life-changing and they may be ill-equipped to manage the transition and adjust to changes in lifestyle. Nurses using PRIME to manage patient care have noted that in follow-up phone conversations with non-compliant patients, the patients have revealed difficulty with emotional problems. Other patients who are more compliant have responded affirmatively to self-report measures about emotional difficulty, suggesting that the challenges may be even more widespread. Addressing mental health challenges is of critical importance not purely for mental health professionals as a complementary component to traditional hospital services that will facilitate a more holistic approach to healthcare. PRIME enables novel ways to use new or existing sensors and mobile devices to remotely assess patients and adjust treatments. Taken together, these mobile health tools will make it possible to reduce the frequency and level of patient visits for assessment and intervention, thus reducing the costs for treatment of depression.

CHAPTER 6

Conclusion

This dissertation addressed the challenges of developing scalable solutions for behavior guidance. This was demonstrated for critical applications in healthcare and energy. Chapter 2 presented design and development of a methodology for energy use behavior guidance, implemented in the Engage sytem. Undergraduate residence halls were instrumented with retrofit energy monitoring systems and residents were provided feedback about their energy consumption. Results showed that residents whose energy consumption was made public were motivated to conserve electricty by 20% whereas private information had little conservation effect.

Chapter 3 presented an update to the Engage system design and application in a broader experimental setting. University apartments were instrumented and participants consisting of graduate student families were provided feedback about their energy consumption with different message framing. The results showed that health messaging had a greater conservation effect and financial messaging actually increased energy consumption. This contrasted with residents' own initial beliefs that financial information would be more effective. When demographic information was incorporated, it was found that households with children saw an amplification of these message treatment effect. On the other hand, public information, as was implemented in the residence hall study, had no effect and the reason is speculated to be due to differences in social connectedness.

In both the residence hall study and the university apartment study, end-to-end remote monitoring, gateway, backend management, and frontend system had to be developed to enable the experimental design objectives. These systems required a full stack understanding of the application from the building electrical infrastructure to user experience design. Tools had to be developed to manage the deployment for troubleshooting in order to ensure the experimental objectives could be realized.

Chapter 4 addressed the challenges in scalability of energy behavior guidance applications. Return on investment of energy behavior guidance given the cost of equipment, installation, and management is prohibitive to large-scale deployment. The solution is develop computational methods to infer high-granularity information from aggregate consumption signals. The presented approach used a time-domain methodology to leverage existing smart meter deployments along with their limitations in data acquisition. The preliminary results show potential although further testing and development with expanded ground truth is needed.

Chapter 5 transitioned from residential energy use behavior and conservation to the problem of hospital readmissions and patient guidance. There are many similarities between these applications in both technology design and information feedback theory. The chapter presented PRIME, an end-to-end system for remote patient monitoring during post-operative recovery. The system uses patient self-report of comprehensive outcome measures to alert physicians to potential complications. A study was conducted which shows strong indications of this paradigm to guide both providers and patients towards improved patient care and readmissions reduction.

To summarize, the contributions of this dissertation are the following:

- Multiple large-scale deployments and field studies of energy monitoring behavior guidance systems challenging conventional beliefs of price/nonprice incentives, generating worldwide media coverage

- Novel time-domain, combinatorial optimization approach for energy load disaggregation leveraging existing smart meter capabilities for improved scalability
- Low-cost, integrated, end-to-end remote patient monitoring system emphasizing provider workflow integration and demonstrating readmission reductions potential, with rapidly expanding service scope, adoption into UCLA Healthcare system, and multi-phase integration into EPIC EHR

As technology further evolves and our society becomes more connected, the barriers to entry for developing behavior guidance will fall. There is tremendous opportunity for technology to play a critical role as a social tool for learning, to optimize how we interact with our environment and each other, and shape ourselves towards our better nature.

APPENDIX A

Estimated RMS Voltage Derivation

Here we derive the formula for the estimated RMS voltage for the Engage plug level energy meter. To review, the sampling approach requires that many packets of data need to be collected to accurately approximate the RMS voltage and current. The following equations are used to compute the voltage estimate from the digital measurements and then to compute the RMS voltage from the voltage estimates.

$$\widehat{V} = \alpha_V(V - \overline{V}) \tag{A.1}$$

$$\widehat{V}_{RMS} = \sqrt{\frac{1}{n}\sum_i \widehat{V}_i^2} \tag{A.2}$$

To improve processing performance, we can distribute the calculation over time by preprocessing raw samples after each packet arrives.

$$
\begin{aligned}
\widehat{V}_{RMS} &= \sqrt{\frac{1}{n}\sum_i \widehat{V}_i^2} \\
&= \alpha_V \sqrt{\frac{1}{n}\sum_i (V_i - \overline{V})^2} \\
&= \alpha_V \sqrt{\frac{1}{n}\sum_i V_i^2 - \frac{2}{n}\sum_i \left(V_i\overline{V}\right) + \left(\frac{1}{n}\sum_i \overline{V}^2\right)} \\
&= \alpha_V \sqrt{\frac{1}{n}\sum_i V_i^2 - 2\overline{V}\left(\frac{1}{n}\sum_i V_i\right) + \overline{V}^2\left(\frac{1}{n}\sum_i 1\right)} \\
&= \alpha_V \sqrt{\frac{1}{n}\sum_i V_i^2 - 2\overline{V}^2 + \overline{V}^2} \\
&= \alpha_V \sqrt{\frac{1}{n}\sum_i V_i^2 - \overline{V}^2} \\
&= \alpha_V \sqrt{\frac{1}{n}\sum +V_i^2 - \left(\frac{1}{n}\sum_i V_i\right)^2}
\end{aligned} \tag{A.3}
$$

Preprocessing becomes a simple task of keeping a running sum of the samples and running sum of squared samples.

References

[ADDdB10] Shahzeen Z Attari, Michael L DeKay, Cliff I Davidson, and Wändi Bruine de Bruin. Public perceptions of energy consumption and savings. *Proceedings of the National Academy of Sciences*, 107(37):16054–16059, 2010.

[AGSA13] K Carrie Armel, Abhay Gupta, Gireesh Shrimali, and Adrian Albert. Is disaggregation the holy grail of energy efficiency? the case of electricity. *Energy Policy*, 52:213–234, 2013.

[All11] Hunt Allcott. Social norms and energy conservation. *Journal of Public Economics*, 95(9):1082–1095, 2011.

[AMV12] David A Asch, Ralph W Muller, and Kevin G Volpp. Automated hovering in health carewatching over the 5000 hours. *New England Journal of Medicine*, 367(1):1–3, 2012.

[AWL+14] L Atallah, A Wiik, B Lo, JP Cobb, AA Amis, and GZ Yang. Gait asymmetry detection in older adults using a light ear-worn sensor. *Physiological measurement*, 35(5):N29, 2014.

[BCO+10] Marthinus J Bekker, Tania D Cumming, Nikola KP Osborne, Angela M Bruining, Julia I McClean, and Louis S Leland. Encouraging electricity savings in a university residential hall through a combination of feedback, visual prompts, and incentives. *Journal of applied behavior analysis*, 43(2):327–331, 2010.

[BK08] Mats Bladh and Helena Krantz. Towards a bright future? household use of electric light: A microlevel study. *Energy Policy*, 36(9):3521–3530, 2008.

[BV04] Michael Baranski and Jürgen Voss. Genetic algorithm for pattern detection in nialm systems. In *Systems, Man and Cybernetics, 2004 IEEE International Conference on*, volume 4, pages 3462–3468. IEEE, 2004.

[BZST09] Susan Barnason, Lani Zimmerman, Paula Schulz, and Chunhao Tu. Influence of an early recovery telehealth intervention on physical activity and functioning after coronary artery bypass surgery among older adults with high disease burden. *Heart & Lung: The Journal of Acute and Critical Care*, 38(6):459–468, 2009.

[CA98] Agnim I Cole and Alexander Albicki. Data extraction for effective non-intrusive identification of residential power loads. In *Instrumentation and Measurement Technology Conference, 1998. IMTC/98. Conference Proceedings. IEEE*, volume 2, pages 812–815. IEEE, 1998.

[CA00] Agnim Cole and Alexander Albicki. Nonintrusive identification of electrical loads in a three-phase environment based on harmonic content. In *Instrumentation and Measurement Technology Conference, 2000. IMTC 2000. Proceedings of the 17th IEEE*, volume 1, pages 24–29. IEEE, 2000.

[CTP+13] David J Cook, Jeffrey E Thompson, Sharon K Prinsen, Joseph A Dearani, and Claude Deschamps. Functional recovery in the elderly after major surgery: assessment of mobility recovery using wireless technology. *The Annals of thoracic surgery*, 96(3):1057–1061, 2013.

[CYL08] Hsueh-Hsien Chang, Hong-Tzer Yang, and Ching-Lung Lin. Load identification in neural networks for a non-intrusive monitoring of industrial electrical loads. In *Computer Supported Cooperative Work in Design IV*, pages 664–674. Springer, 2008.

[DAFSF11] Anibal De Almeida, Paula Fonseca, Barbara Schlomann, and Nicolai Feilberg. Characterization of the household electricity consumption in the eu, potential energy savings and specific policy recommendations. *Energy and Buildings*, 43(8):1884–1894, 2011.

[DFA13] Magali A Delmas, Miriam Fischlein, and Omar I Asensio. Information strategies and energy conservation behavior: A meta-analysis of experimental studies from 1975 to 2012. *Energy Policy*, 61:729–739, 2013.

[DFS+01] CP Delaney, VW Fazio, AJ Senagore, B Robinson, AL Halverson, and FH Remzi. fast trackpostoperative management protocol for patients with high co-morbidity undergoing complex abdominal and pelvic colorectal surgery. *British Journal of Surgery*, 88(11):1533–1538, 2001.

[DHP+12] Lynn V Doering, Kathleen Hickey, David Pickham, Belinda Chen, and Barbara J Drew. Remote noninvasive allograft rejection monitoring for heart transplant recipients: study protocol for the novel evaluation with home electrocardiogram and remote transmission (new heart) study. *BMC cardiovascular disorders*, 12(1):14, 2012.

[DL14] Magali A Delmas and Neil Lessem. Saving power to conserve your reputation? the effectiveness of private versus public information. *Journal of Environmental Economics and Management*, 67(3):353–370, 2014.

[DMXF12] Ming Dong, Paulo CM Meira, Wilsun Xu, and Walmir Freitas. An event window based load monitoring technique for smart meters. *Smart Grid, IEEE Transactions on*, 3(2):787–796, 2012.

[DSR+14] Aaron J Dawes, Greg D Sacks, Marcia M Russell, Anne Y Lin, Melinda Maggard-Gibbons, Deborah Winograd, Hallie R Chung, Areti Tillou, Jonathan R Hiatt, and Clifford Ko. Preventable readmissions to surgical services: Lessons learned and targets for improvement. *Journal of the American College of Surgeons*, 219(3):382–389, 2014.

[Edi] Edison Foundation. Utility-scale smart meter deployments, plans & proposals. `http://www.edisonfoundation.net/iei/Documents/IEE_SmartMeterRollouts_0512.pdf`.

[EK72] Jack Edmonds and Richard M Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM (JACM)*, 19(2):248–264, 1972.

[EM98] Gary W Evans and Janetta Mitchell McCoy. When buildings dont work: the role of architecture in human health. *Journal of Environmental psychology*, 18(1):85–94, 1998.

[EMDL+10] Karen Ehrhardt-Martinez, Kat A Donnelly, Skip Laitner, et al. Advanced metering initiatives and residential feedback programs: a meta-review for household electricity-saving opportunities. *American Council for an Energy-Efficient Economy*, 2010.

[Fis08] Corinna Fischer. Feedback on household electricity consumption: a tool for saving energy? *Energy efficiency*, 1(1):79–104, 2008.

[FZ99] Linda Farinaccio and Radu Zmeureanu. Using a pattern recognition approach to disaggregate the total electricity consumption in a house into the major end-uses. *Energy and Buildings*, 30(3):245–259, 1999.

[GKB+03] Paul E Greenberg, Ronald C Kessler, Howard G Birnbaum, Stephanie A Leong, Sarah W Lowe, Patricia A Berglund, and Patricia K Corey-Lisle. The economic burden of depression in the united states: how did it change between 1990 and 2000? *Journal of clinical psychiatry*, 64(12):1465–1475, 2003.

[GMA+07] D Guinier, GA Mantion, A Alves, F Kwiatkowski, K Slim, Y Panis, et al. Risk factors of unplanned readmission after colorectal surgery: a prospective, multicenter study. *Diseases of the Colon & Rectum*, 50(9):1316–1323, 2007.

[GS08] Gerald T Gardner and Paul C Stern. The short list: The most effective actions us households can take to curb climate change. *Environment: science and policy for sustainable development*, 50(5):12–25, 2008.

[GT11] Alice Grønhøj and John Thøgersen. Feedback on household electricity consumption: learning and social influence processes. *International Journal of Consumer Studies*, 35(2):138–145, 2011.

[Har92] George William Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 1992.

[HdD04] Melissa Hart and Richard de Dear. Weather sensitivity in household appliance energy end-use. *Energy and Buildings*, 36(2):161–174, 2004.

[JDHDC09] Xiaofan Jiang, Stephen Dawson-Haggerty, Prabal Dutta, and David Culler. Design and implementation of a high-fidelity ac metering network. In *Information Processing in Sensor Networks, 2009. IPSN 2009. International Conference on*, pages 253–264. IEEE, 2009.

[JTP12] Rishee K Jain, John E Taylor, and Gabriel Peschiera. Assessing eco-feedback interface usage and design to drive energy efficiency in buildings. *Energy and buildings*, 48:8–17, 2012.

[Kar11] Sami Karjalainen. Consumer preferences for feedback on household electricity consumption. *Energy and Buildings*, 43(2):458–467, 2011.

[KBN10] J Zico Kolter, Siddharth Batra, and Andrew Y Ng. Energy disaggregation via discriminative sparse coding. In *Advances in Neural Information Processing Systems*, pages 1153–1161, 2010.

[Kil] Kill A Watt. `http://www.p3international.com/products/p4400.html`.

[KJ12] J Zico Kolter and Tommi Jaakkola. Approximate inference in additive factorial hmms with application to energy disaggregation. In *International conference on artificial intelligence and statistics*, pages 1472–1482, 2012.

[KL94] Willett Kempton and Linda L Layne. The consumer's energy analysis environment. *Energy Policy*, 22(10):857–866, 1994.

[KOP+12] Michael T Kassin, Rachel M Owen, Sebastian D Perez, Ira Leeds, James C Cox, Kurt Schnier, Vjollca Sadiraj, and John F Sweeney. Risk factors for 30-day hospital readmission among general surgery patients. *Journal of the American College of Surgeons*, 215(3):322–330, 2012.

[KSCS09] Younghun Kim, Thomas Schmid, Zainul M Charbiwala, and Mani B Srivastava. Viridiscope: design and implementation of a fine grained power monitoring system for homes. In *Proceedings of the 11th international conference on Ubiquitous computing*, pages 245–254. ACM, 2009.

[Kuh55] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.

[LEM09] John A Laitner and Karen Ehrhardt-Martinez. Examining the scale of the behaviour energy efficiency continuum. *European Council for an Energy Efficient Economy*, 2009.

[LHL+13] Elise H Lawson, Bruce Lee Hall, Rachel Louie, Susan L Ettner, David S Zingmond, Lein Han, Michael Rapp, and Clifford Y Ko. Association between occurrence of a postoperative complication and readmission: implications for quality improvement and cost savings. *Annals of surgery*, 258(1):10–18, 2013.

[LHL+14] Elise H Lawson, Bruce Lee Hall, Rachel Louie, David S Zingmond, and Clifford Y Ko. Identification of modifiable factors for reducing readmission after colectomy: A national analysis. *Surgery*, 155(5):754–766, 2014.

[LLC+03] Christopher Laughman, Kwangduk Lee, Robert Cox, Steven Shaw, Steven Leeb, Les Norford, and Peter Armstrong. Power signature analysis. *Power and Energy Magazine, IEEE*, 1(2):56–63, 2003.

[LNKC10a] Jian Liang, Simon KK Ng, Gail Kendall, and John WM Cheng. Load signature studypart i: Basic concept, structure, and methodology. *Power Delivery, IEEE Transactions on*, 25(2):551–560, 2010.

[LNKC10b] Jian Liang, Simon KK Ng, Gail Kendall, and John WM Cheng. Load signature studypart ii: Disaggregation framework, simulation, and applications. *Power Delivery, IEEE Transactions on*, 25(2):561–569, 2010.

[Lut93] Loren Lutzenhiser. Social and behavioral aspects of energy use. *Annual Review of Energy and the Environment*, 18(1):247–289, 1993.

[LW12] Edward H Livingston and Joseph S Wislar. Minimum response rates for survey research. *Archives of Surgery*, 147(2):110–110, 2012.

[LWJ+01] Simon Lusignan, Sally Wells, Paul Johnson, Karen Meredith, and Edward Leatham. Compliance and effectiveness of 1 year's home telemonitoring. the report of a pilot study of patients with chronic heart failure. *European Journal of Heart Failure*, 3(6):723–730, 2001.

[MKPE02] Birgit Morlion, Christiane Knoop, Manuel Paiva, and Marc Estenne. Internet-based home monitoring of pulmonary function after lung transplantation. *American journal of respiratory and critical care medicine*, 165(5):694–697, 2002.

[Mun57] James Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial & Applied Mathematics*, 5(1):32–38, 1957.

[MZ00] Medgar Louis Marceau and R Zmeureanu. Nonintrusive load disaggregation computer program to estimate the energy consumption of major end uses in residential buildings. *Energy Conversion and Management*, 41(13):1389–1403, 2000.

[NL96] Leslie K Norford and Steven B Leeb. Non-intrusive electrical load monitoring in commercial buildings based on steady-state and transient load-detection algorithms. *Energy and Buildings*, 24(1):51–64, 1996.

[NRB09] B Neenan, J Robinson, and RN Boisvert. Residential electricity use feedback: A research synthesis and economic framework. *Electric Power Research Institute*, 2009.

[Par03] Danny S Parker. Research highlights from a large scale residential monitoring study in a hot climate. *Energy and Buildings*, 35(9):863–876, 2003.

[PCCR93] RG Pratt, CC Conner, BA Cooke, and EE Richman. Metered end-use consumption and load shapes from the elcap residential sample of existing homes in the pacific northwest. *Energy and buildings*, 19(3):179–193, 1993.

[PHL+08] Sarah J Pangarakis, Kathleen Harrington, Ruth Lindquist, Cynthia Peden-McAlpine, and Stanley Finkelstein. Electronic feedback messages for home spirometry lung transplant recipients. *Heart & Lung: The Journal of Acute and Critical Care*, 37(4):299–307, 2008.

[PRK+07] Shwetak N Patel, Thomas Robertson, Julie A Kientz, Matthew S Reynolds, and Gregory D Abowd. *At the flick of a switch: Detecting and classifying unique electrical events on the residential power line (nominated for the best paper award)*. Springer, 2007.

[PSJ+07] John E Petersen, Vladislav Shunturov, Kathryn Janda, Gavin Platt, and Kate Weinberger. Dormitory residents reduce electricity consumption when exposed to real-time visual feedback and incentives. *International Journal of Sustainability in Higher Education*, 8(1):16–33, 2007.

[PTS10] Gabriel Peschiera, John E Taylor, and Jeffrey A Siegel. Response–relapse patterns of building occupant electricity consumption following exposure to personal, contextualized and occupant peer network utilization data. *Energy and Buildings*, 42(8):1329–1336, 2010.

[RMJS+09] Elizabeth Russell-Minda, Jeffrey Jutai, Mark Speechley, Kaitlin Bradley, Anna Chudyk, and Robert Petrella. Health technologies for monitoring and managing diabetes: a systematic review. *Journal of diabetes science and technology*, 3(6):1460–1471, 2009.

[SBM12] Neil Saldanha and Ian Beausoleil-Morrison. Measured end-use electric load profiles for 12 canadian houses at high temporal resolution. *Energy and Buildings*, 49:519–530, 2012.

[SDR+14] Greg D Sacks, Aaron J Dawes, Marcia M Russell, Anne Y Lin, Melinda Maggard-Gibbons, Deborah Winograd, Hallie R Chung, James Tomlinson, Areti Tillou, Stephen B Shew, et al. Evaluation of hospital readmissions in surgical patients: do administrative data tell the real story? *JAMA surgery*, 149(8):759–764, 2014.

[SLNC08] Steven R Shaw, Steven B Leeb, Leslie K Norford, and Robert W Cox. Nonintrusive load monitoring and diagnostics in power systems. *Instrumentation and Measurement, IEEE Transactions on*, 57(7):1445–1454, 2008.

[Smi15] Aaron Smith. U.s. smartphone use in 2015. Technical report, Pew Research Center, 2015.

[SNL06] D Srinivasan, WS Ng, and AC Liew. Neural-network-based signature recognition for harmonic source identification. *Power Delivery, IEEE Transactions on*, 21(1):398–405, 2006.

[SSR+14] Aditi P Sen, Taylor B Sewell, E Brooks Riley, Beth Stearman, Scarlett L Bellamy, Michelle F Hu, Yuanyuan Tao, Jingsan Zhu, James D Park, George Loewenstein, et al. Financial incentives for home-based health monitoring: A randomized controlled trial. *Journal of general internal medicine*, 29(5):770–777, 2014.

[Sul91] F Sultanem. Using appliance signatures for monitoring residential loads at meter panel level. *Power Delivery, IEEE Transactions on*, 6(4):1380–1385, 1991.

[The08] The Standish Group. Trends in open source. Technical report, The Standish Group, 2008.

[Tom71] N Tomizawa. On some techniques useful for solution of transportation network problems. *Networks*, 1(2):173–194, 1971.

[U. a] U. S. Energy Information Agency. Retail sales of electricity to ultimate customers: Total by end-use sector. epa. `http://www.eia.doe.gov/cneaf/electricity/epm/table5_1.html`.

[U. b] U. S. Environmental Protection Agency. Inventory of u.s. greenhouse gas emissions and sinks: 1990 - 2008. `http://www.epa.gov/climatechange/ghgemissions/usinventoryreport/archive.html`.

[U.S] U.S. Department of Energy. Energy saver guide: Tips on saving money and energy at home. `http://energy.gov/sites/prod/files/2014/09/f18/61628_BK_EERE-EnergySavers_w150.pdf`.

[USST06] Tsuyoshi Ueno, Fuminori Sano, Osamu Saeki, and Kiichiro Tsuji. Effectiveness of an energy-consumption information system on energy savings in residential houses based on monitored data. *Applied Energy*, 83(2):166–183, 2006.

[VTG+08] Prashanth Varkey, Ngian Chye Tan, Riccardo Girotto, Wen-Ruay Tang, Yi-Tein Liu, and Hung-Chi Chen. A picture speaks a thousand words: The use of digital photography and the internet as a cost-effective tool in monitoring free flaps. *Annals of plastic surgery*, 60(1):45–48, 2008.

[WN03] G Wood and M Newborough. Dynamic energy-consumption indicators for domestic appliances: environment, behaviour and design. *Energy and Buildings*, 35(8):821–841, 2003.

[Wol11] Frank A Wolak. Do residential customers respond to hourly prices? evidence from a dynamic pricing experiment. *The American Economic Review*, pages 83–87, 2011.

[WSH+11] Elizabeth C Wick, Andrew D Shore, Kenzo Hirose, Andrew M Ibrahim, Susan L Gearhart, Jonathan Efron, Jonathan P Weiner, and Martin A Makary. Readmission rates and cost following colorectal surgery. *Diseases of the Colon & Rectum*, 54(12):1475–1479, 2011.

[YGHF08] Hojung J Yoon, Hongfei Guo, Marshall Hertz, and Stanley Finkelstein. Adherence to home-monitoring and its impact on survival in post-lung transplantation patients. In *AMIA Annual Symposium Proceedings*, volume 2008, page 835. American Medical Informatics Association, 2008.

[ZGIR12] Ahmed Zoha, Alexander Gluhak, Muhammad Ali Imran, and Sutharshan Rajasegarar. Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. *Sensors*, 12(12):16838–16866, 2012.

[ZR11] Michael Zeifman and Kurt Roth. Nonintrusive appliance load monitoring: Review and outlook. *IEEE Transactions on Consumer Electronics*, pages 76–84, 2011.

[ZS13] Kathryn Zickuhr and Aaron Smith. Home broadband 2013. Technical report, Pew Research Center, 2013.