UCLA

UCLA Previously Published Works

Title

Real-time, appliance-level electricity use feedback system: How to engage users?

Permalink

<https://escholarship.org/uc/item/58f85945>

Authors

Chen, VL
Delmas, MA
Kaiser, WJ

Publication Date

2014-02-01

DOI

10.1016/j.enbuild.2013.11.069

Peer reviewed

Real-time, appliance-level electricity use feedback system: How to engage users?

Victor L. Chen^{a,*}, Magali A. Delmas^b, William J. Kaiser^a

^a Electrical Engineering Department, University of California, Los Angeles, CA, United States

^b Institute of the Environment and Sustainability and Anderson School of Management, University of California, Los Angeles, CA, United States

ARTICLE INFO

Article history:

Received 31 May 2013

Received in revised form

15 November 2013

Accepted 25 November 2013

Keywords:

Energy use behavior
Information feedback
End-to-end system
Retrofit technology
Smart meter
Energy monitoring
Energy dashboard
Web analytics
Social impact
User engagement

ABSTRACT

Engage is a rapidly deployable, retrofit energy monitoring system developed for direct support of a novel energy use behavior investigation and a large-scale deployment in campus apartments. We describe the end-to-end system and report results related to web dashboard engagement during a year-long experiment. The objective was to determine user engagement with real-time and easily accessible information about personal energy consumption. Leveraging low-cost components, this system was designed to measure separately appliance plug load, heating and cooling, and lighting electrical load in dense-occupancy building environments. We developed and used an open source technology for measurement of plug load and developed signal processing algorithms to significantly improve measurement accuracy. We also developed proxy sensors to measure heating and cooling and lighting. Our results indicate that 90% of the dashboard activity was undertaken by 50% of the participants and that website engagement was more likely in mid-day and more effective in combination with email reminders. Energy conservation was achieved when combining the dashboard with public information about energy consumption.

Published by Elsevier B.V.

1. Introduction

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 [1,2]. This is not surprising considering that residents of the United States spend more than 90% of their lives in buildings [3]. Recent studies estimate that behavioral changes can reduce residential energy consumption significantly [4,5]. Delmas et al. reviewed 156 studies and found an average 7.4% reduction in energy consumption, with the largest reductions resulting from individualized feedback [6]. However, many of the studies in the literature suffer from methodological limitations including small sample sizes [7,8], short time periods [9], or low-granularity feedback (i.e. providing only total usage for the building versus per-appliance usage) [10–12].

While scholars argue that high-granularity information can facilitate energy reduction [13], 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. Plug-level devices, such as the Kill A Watt [14] and ACme [15], 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 [16,17]. 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 [18].

In this project, to achieve high-granularity, appliance-level feedback at the room level, we developed an end-to-end system architecture that included low-cost, wired appliance-level sensors and wireless plug-level meters, a remote gateway for

* Corresponding author. Tel.: +1 714 414 9355.

E-mail addresses: victor.l.chen@gmail.com (V.L. Chen), delmas@ioes.ucla.edu (M.A. Delmas), kaiser@ee.ucla.edu (W.J. Kaiser).

Fig. 1. Engage end-to-end system architecture.

local processing and data upload, and a backend for data storage, data processing, and web services. We monitored residents' energy usage over a 7-month period, providing high-resolution and high-granularity individualized information. We also observed dashboard engagement using website analytics. While several studies have assessed the effectiveness of various design components of feedback information [19–21], but to our knowledge none have used website analytics data to identify dashboard access patterns. Our study is therefore the first to test the effectiveness of an end-to-end feedback system on consumer engagement and energy consumption.

We deployed our end-to-end system 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 variations in design. This allowed us 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 allowed us to test the effectiveness of information in the absence of an inherent financial incentive to conserve electricity.

The remainder of the paper is organized as follows. We describe the system design and each of the components. We also describe in detail the signal processing for our energy meters as it significantly improves the capabilities of the reference design on which it is based. Finally, we describe the user web dashboard and provide results from the analytics data and energy consumption.

2. System design

Our objective was to test the effect of access to appliance level, real-time and historical energy use information on energy consumption. To achieve high granularity, we developed an end-to-end system architecture shown in Fig. 1 that included low-cost sensors, a remote gateway for local processing and data upload, and a backend for data storage, data post-processing, and web hosting. We identified three load categories that were controllable by residents, and were practical to measure: (1) the appliance plug load from the 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, we found from survey results that residents consistently overestimated the share of energy from lighting use as being nearly equal to heating and cooling use whereas results from our system indicated that heating and cooling comprised 72% of energy versus 5% for lighting.

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 deployable, given the short time span allotted for installation by the

administration between the end of summer occupancy and start of fall occupancy. Further, budget constraints required us to strive for a low cost solution in order to reach a deployment scale and population sample size that would yield statistically significant behavioral analysis. The total hardware cost per installation was under \$200.

2.1. Hardware

The deployment hardware consisted of four components designed to collect information about electricity consumption and appliance state and transmit it to our database via the building's wired network: the energy meter, light and temperature proxy sensors, and gateway.

2.1.1. Energy meter

The energy meter was a modified Kill A Watt which allowed us to measure the electrical plug load, interface with proxy sensors, and transmit measurements to the gateway. The augmentation was 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-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. Our design made use of the additional analog input channels on the XBee radio to interface with the proxy sensors.

For each room, we installed two energy meters to capture the energy usage. This number was based on the room design such as the placement of electrical outlets, arrangement of appliances and convenience to the residents. We installed the meters on two selected electrical outlets and the power strips were plugged into the meters with other outlets covered over with 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.). Energy measurements from the two meters were then added together and we refer to this as the plug load.

2.1.2. Proxy sensors

To allay cost and work with infrastructure constraints, we developed proxy sensors using photodiodes and linear active thermistors. The photodiodes were used to determine light state, which we then converted to light energy consumption. Similarly, the thermistors were used to determine HVAC state and energy consumption. The component materials cost only a few dollars and are less expensive than direct energy measurement using current transducers which would have also required infrastructure modification. We further reduced cost by designing the sensors as cable strands rather than as wireless sensor platforms. While wireless sensors would have helped expedite the installation, the hardware cost of a wireless platform using our limited budget would have been prohibitive to the deployment scale we required.

An example floor plan and sensor installation is shown in Fig. 2. To the extent possible, the cables were installed along the corners of the participants' rooms to be minimally conspicuous and reduce the potential for accidental tampering.

Each room contained either one or two overhead lights. We affixed light sensors to the light source so as to minimize the amount of ambient light. It was 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. We were

Fig. 2. Typical double- and triple-occupancy room floor plan and sensor installation.

thus able to reliably determine light state from the light sensor using a common threshold. Fig. 3 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 the sensor could detect ambient light from the sun, light intensity from the target source was sufficiently greater as to be easily discernible. Lighting power is given by

$$P_{\text{light}}(t) = \begin{cases} \rho_{\text{light}}, & L(t) \geq \lambda \\ 0, & L(t) < \lambda \end{cases} \quad (1)$$

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

For temperature control, two adjacent rooms shared an HVAC unit with dampers which were priority-controlled 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 sensor was placed some distance from the vent to measure ambient temperature. We used the differential temperature of the sensors to infer the HVAC state. However, 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, we computed the average differential over the past 100 min as a dynamic temperature threshold. Heating or cooling was classified as “On” when the current absolute temperature differential exceeded this threshold. With guidance from the housing administration and facilities, we learned that a conservative figure

Fig. 3. Light sensor signal for one room over the course of a day.

Fig. 4. Temperature differential T_{diff} , dynamic threshold κ_{τ} , and HVAC state estimates for one room over one day.

for per room HVAC energy usage was 2000 W. We calculated HVAC power consumption as follows:

$$P_{\text{HVAC}}(t) = \begin{cases} \rho_{\text{HVAC}}, & T_{\text{diff}}(t) \geq \kappa_{\tau}(t) \\ 0, & T_{\text{diff}}(t) < \kappa_{\tau}(t) \end{cases} \quad (2)$$

$$T_{\text{diff}}(t) = |T_{\text{vent}}(t) - T_{\text{amb}}(t)| \quad (3)$$

$$\kappa_{\tau}(t) = \frac{1}{t - \tau} \int_{\tau}^t T_{\text{diff}}(t) dt \quad (4)$$

where ρ_{HVAC} is the rated HVAC power consumption (estimated to be 2000 W), T_{diff} 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 (100 min in our case).

Fig. 4 shows the absolute temperature differential and dynamic threshold temperature signals for one room over a day along with the HVAC state estimation. The “On” states indicated by the algorithm corresponded well to the spikes in differential temperature resulting from heating cycles.

2.1.3. Gateway

To reduce gateway development time and cost we leveraged an Asus WL-520gU wireless router and flashed it with OpenWRT, an open source operating system optimized for wireless router devices. The selection of this router was important for a number of reasons: (1) the router supports operating system upgrade, (2) OpenWRT supports the Broadcom BCM5354 processor used on the Asus router, (3) the router includes a USB port, and (4) an RS-232 serial port.

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

2.2. Signal processing

We now describe the methods used for sampling and processing the current and voltage signals to obtain power factor estimates. This processing enabled estimation of real power consumption, as is done for electric utility billing.

Table 1 summarizes the terms used.

Table 1
Summary of terms.

V	Measured voltage (digital value)
$\bar{V}$	Average measured voltage (digital value)
$\hat{V}$	Estimated actual voltage
I	Measured current (digital value)
$\bar{I}$	Average measured current (digital value)
$\hat{I}$	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)

2.2.1. Voltage and current sampling

The Tweet A Watt design uses a 2-s 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-s 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 Watt's internal power supply is current limited and cannot provide the 50 mA burst needed for wireless transmission. Forcing a 2-s delay between wireless transmissions allows the capacitor sufficient time to charge for transmissions. In our design, the additional power required by the active thermistors required a larger current source so we modified the Kill A Watt so the XBee subsystem was powered with a few 5 V AC/DC adapter. Despite the removal of the power constraints, we retained the 2-s cyclic sleep 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 Fig. 5 (left), discontinuities arise in the sampled waveform as shown in Fig. 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, we could not achieve accurate power calculation over very short time periods on the order of seconds. However, for our purposes we assumed that power consumption within a window of tens of seconds would have reached steady state. We then treated the measured samples as a signal that had the same statistical distribution as a continuous signal of the same length. This was particularly important for power factor estimation described in the next section.

This sampling approach required that many packets of data to be collected to accurately approximate the RMS voltage and

current. To improve computation, the processing was distributed in time by preprocessing raw samples after each packet arrived. Raw digital values were first converted to volts and amps using a scaling coefficient. Then the actual value corresponding to each sample was calculated as the scaled, zero-mean adjusted sample value

$$\hat{V} = \alpha_V (V - \bar{V}) \quad (5)$$

The RMS voltage was then calculated as

$$\begin{aligned} V_{\text{RMS}} &= \sqrt{\frac{1}{n} \sum_i \hat{V}_i^2} = \alpha_V \sqrt{\frac{1}{n} \sum_i (V_i - \bar{V})^2} \\ &= \alpha_V \sqrt{\frac{1}{n} \sum_i V_i^2 - \left(\frac{1}{n} \sum_i V_i \right)^2} \end{aligned} \quad (6)$$

and RMS current was calculated similarly. Using the last form in Eq. (6), preprocessing became a simple task of keeping a running sum of the samples and running sum of squared samples. The scaling coefficients were determined empirically using known values of voltage and current.

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, utility providers generally bill customers only for the real power consumed, regardless of the power factor. We wanted to also provide real power to reflect convention.

There are a number of methods for determining power factor in a load. We based our method 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. We now describe our approach for estimating the real power.

Real power needed can be computed from apparent power using power factor as a scaling coefficient.

$$P_{\text{def}} |S| \times pf. \quad (7)$$

Fig. 5. Real (left) and sampled (right) voltage and current waveforms.

Fig. 6. Quadratic regression for power factor estimation.

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

$$|S|_{\text{def}} V_{\text{RMS}} \times I_{\text{RMS}} \quad (8)$$

where the RMS voltage and current was computed directly from the sample measurements as described in the previous section.

To estimate the power factor, we first normalized the voltage and current signals 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 were calculated as

$$V_{\text{norm}} = \frac{V - V_{\min}}{V_{\max} - V_{\min}} \quad (9)$$

We then computed the difference of the means of the normalized voltage and current. 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. Thus, we could model the deviation to derive its relationship to power factor. We computed the mean difference as

$$\mu = (\overline{V_{\text{norm}}} - \overline{I_{\text{norm}}}) \quad (10)$$

We then developed an empirical power factor model to fit observed mean difference values to the power factor reported by

Table 2

Comparison of Kill A Watt real power (RP), Tweet A Watt apparent power (AP), and engage real power (RP) measurements.

KAW RP	TAW AP	TAW (% error)	Engage RP	Engage (% error)
6.3	14.7	133.3	6.3	0
19.3	27.8	44	20.1	4.1
25.4	66	159.8	23.6	−7.1
25.5	48.9	91.8	24.1	−5.5
51.2	62.5	22.1	48.2	−5.9
86.3	98.1	13.7	87.9	1.9
94.9	105.2	10.9	99.2	4.5
1254	1242.3	−0.9	1170.3	−6.7

the Kill A Watt meter for a variety of appliances with varying power factors and found a roughly quadratic as follows given by (Fig. 6)

$$pf = -5\mu^2 + 0.83\mu + 0.93 = \varepsilon \quad (11)$$

This power factor model was then incorporated into the signal processing algorithms of the gateway daemon to compute real power. We tested a variety of loads consisting of typical home appliances with various power factors and found the error in real power estimate compared to Kill A Watt's real power measurement to be less than 10 percent. These measurements are summarized in Table 2 to highlight the fundamental limitation of the Tweet A Watt to produce an accurate real power estimate for general loads.

2.2.3. End-use energy calculations

As described above, we have developed a means for accurately estimating the power consumption over time for plug load, heating and cooling use, and lighting use. Using these estimates, we can calculate end-use energy statistics for presentation on the dashboard. For any given room and any given window of time specified by start time t_1 and end time t_2 , the total energy consumed for a load ℓ in the set {plug, hvac, light} was calculated as

$$E_{\ell}(t_1, t_2) = \int_{t_1}^{t_2} P_{\ell}(t) dt \quad (12)$$

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

$$E_{\text{total}}(t_1, t_2) = E_{\text{plug}}(t_1, t_2) + E_{\text{HVAC}}(t_1, t_2) + E_{\text{light}}(t_1, t_2). \quad (13)$$

Fig. 7. Engage energy dashboard.

Fig. 8. Total daily energy use averaged across all rooms, shown as a stacked area plot.

We computed the total energy of this window for each room and group statistics such as average and quantile usage.

2.3. Backend system

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.

2.3.1. Gateway daemon

The gateway daemon was designed to facilitate the transport of data from the energy meters to the database and then to the user. Each energy meter was configured with a unique identification number and automatically transmitted analog data to its gateway. The gateway received these packets and processed them to extract current, voltage, and power and finally uploaded these measurements along with light and temperature sensor measurements to the server.

2.3.2. Engage server system

The Engage Server system consisted 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.

2.3.3. Engage user dashboard

The engage user dashboard, shown in Fig. 7, displayed 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. We provided three levels of data resolution: Real-Time with 5-minute resolution for the past 3 h, Hourly (as the default view) with 1-h 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 that 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. We distributed the codes via email to the residents at the beginning of the experimental treatment phase. We also included a direct link to the dashboard in individualized weekly email reports.

3. Results

We now describe results of the energy use and dashboard access data. Please refer to Delmas and Lessem [24] for a full description of the experimental results of the behavioral study.

Fig. 9. Number of people per total pageviews.

Fig. 10. Total unique pageviews per day.

3.1. Energy use

Fig. 8 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 were relatively constant throughout the experimental period, HVAC use was more volatile. Gaps in the data were the result of network or server failure events whereas dips in electricity consumption resulted from academic breaks.

3.2. Engage user dashboard analytics

We used Google Analytics 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 provided a wealth of information about users' response to the information we were providing.

Fig. 9 shows a histogram of user access to the dashboard over the course of the experiment. While every dashboard was accessed at least once, were a few residents visited the dashboard more than 80 times and many residents who visited the dashboard less than 10 times. These results follow a power law distribution with 80% of the dashboard activity generated by 25% of the residents.

Another effect was the rapid drop-off we observe in dashboard access shown in Fig. 10. Access in the first week (following the launch) started off very high as participants explored their dashboards. However, especially after the first day (when the launch announcement was emailed) there was a large drop in activity as the novelty wore off.

Despite this outcome, dashboard access was still higher when residents received weekly email reports. Fig. 11 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 played an important role in reminding users about the study and drawing them to the website to view their dashboards.

Fig. 12 shows total pageviews as a function of the time of day the dashboard was accessed. Dashboard access followed a daily cycle that peaked around midday with a second peak period late at night.

Fig. 11. Total number of pageviews per day of the week.

Fig. 12. Total number of pageviews for each hour of the day, over entire study.

Presumably these were periods when people found 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.

In order to learn what aspects of the dashboard were more useful, we broke down dashboard access by the view type (real-time, hourly, or daily), as shown in Fig. 13. We observed 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 was the default page and visitors may not have been exploring beyond the hourly page. However, the view count was not heavily skewed toward any one-view type. In conjunction with the power law distribution of dashboard viewership, this suggests that people who were highly engaged with the dashboard were likely to explore the different views more thoroughly.

We also observed that users spent 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 was a limitation of how Google Analytics calculated pageview duration. Google Analytics calculated viewing duration as the difference between timestamps of subsequent pageviews. Thus, for the last or only page viewed in a session, there would not be a duration value calculated for that pageview. Therefore, even though users may have been subsequently viewing the real-time and daily pages, data could have been skewed toward recording viewing time for the hourly page since it was the default view. This is reinforced by the pageview count breakdown in Fig. 13, 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.

3.3. Energy behavior impact

We now summarize several findings on the behavioral effects of dashboard information on energy consumption. Delmas and Lessem [24] present a thorough analysis and explanation of the results. First, we conducted an exit survey to ask about the efficacy of the dashboard and whether any energy conservation actions 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

Fig. 13. Percent of the total pageview count for each view type.

encouraged to conserve energy after 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 W-h/day, or 80 min/day, representing a 20% reduction compared to the control group. We also reported energy consumption on the dashboard in combination with weekly emailing and making public the 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%.

4. Conclusion

We presented Engage, a rapidly-deployable, retrofit, end-to-end system for monitoring per room electricity consumption. The technology we developed was designed to address the infrastructure constraints we encountered in our setting while providing individualized, appliance level, real-time and historical feedback. Our system could easily be applied to any large building where per room energy monitoring is desired but where the electrical infrastructure does not support centralized approaches to monitoring. We described our experimental objectives and the deployment challenges. These informed the design of our system, along with our cost and development constraints. We discussed the energy meter in detail as a device based on a popular open-source modification of commercial hardware and demonstrated how improved signal processing enabled estimation of real power consumption. We also described the proxy sensors used for measurement of light and HVAC usage.

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 status could reduce system downtime and data loss.

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

Our current work has focused on deployment at graduate student family housing complexes. Here, residents are married couples or have children and also pay for their electricity consumption, providing a better representation of the general population. The Engage systems have been adapted to include circuit level sensing since each apartment has an electrical distribution panel. Software features have also been implemented to improve gateway reliability and system diagnostics. Dashboard designs have been modified and tested to determine how to optimize dashboard information as a learning mechanism.

Acknowledgments

This work was made possible by a grant from UCLA's The Green Initiative Fund (TGIF) and the University of California Center for

Energy and Environmental Economics (UCE³). We would like to express our gratitude to Dr. Robert Gilbert for his guidance and for enabling and this work and to Dr. Neil Lessem for guiding and managing the project. We are also grateful to UCLA housing services for allowing us to conduct this project.

References

- [1] Retail Sales of Electricity to Ultimate Customers: Total by End-Use Sector, U. S. Energy Information Agency. Available: <http://www.eia.doe.gov/cneaf/electricity/epm/table5.1.html>
- [2] Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990 – 2008. U. S. Environmental Protection Agency. Available: <http://www.epa.gov/climatechange/ghgemissions/usinventoryreport/archive.html>
- [3] G.W. Evans, J.M. McCoy, When buildings don't work: the role of architecture human health, *Journal of Environmental Psychology* 18 (1998) 85–94.
- [4] J. Laitner, K. Ehrhardt-Martinez, V. McKinney, Examining the scale of the behavior energy efficiency continuum, in: *Proceedings of 2009 ECEEE Summer Study*, 2009, pp. 217–223.
- [5] G. Gardner, P. Stern, The short list: the most effective actions U.S. households can take to curb climate change, *Environment* 50 (2008) 12–24.
- [6] M.A. Delmas, M. Fischlein, O.I. Asensio, Information Strategies and Energy Conservation Behavior: A Meta-analysis of Experimental Studies from 1975 to 2012, *Energy Policy* 61 (2013) 729–739.
- [7] A. Gronhoj, J. Thøgersen, Feedback on household electricity consumption: learning and social influence processes, *International Journal of Consumer Studies* 35 (2011) 138–145.
- [8] T. Ueno, F. Sano, O. Saeki, K. Tsuji, Effectiveness of an energy-consumption information system on energy savings in residential houses based on monitored data, *Applied Energy* 83 (2) (2006) 166–183.
- [9] J.E. Petersen, V. Shunturov, K. Janda, G. Platt, K. Weinberger, Dormitory residents reduce electricity consumption when exposed to real-time visual feedback and incentives, *International Journal of Sustainability in Higher Education* 8 (2007) 16–33.
- [10] H. Alcott, Social norms and energy conservation, *Journal of Public Economics* 95 (9–10) (2011) 1082–1095.
- [11] M.J. Becker, T.D. Cumming, N.K.P. Osborne, A.M. Bruining, J.I. McClean, L.S. Leland Jr., Encouraging electricity savings in a university residential hall through a combination of feedback, visual prompts, and incentives, *Journal of Applied Behavior Analysis* 43 (2) (2010) 327–331.
- [12] F. Wolak, Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment, *American Economic Review* 101 (3) (2011) 83–87.
- [13] C. Fischer, Feedback on household electricity consumption: a tool for saving energy? *Energy Efficiency* 1 (2008) 79–104.
- [14] Kill A Watt. P3 – Kill A Watt. <http://www.p3international.com/products/P4400.htm>. (Accessed: Oct. 2, 2013).
- [15] X. Jiang, S. Dawson-Haggerty, J. Taneja, P. Dutta, D. Culler, Creating greener homes with IP-based wireless AC energy monitors, *SensSys* (2008) 355–356.
- [16] Y. Kim, T. Schmid, Z.M. Charbiwala, M.B. Srivastava, ViridiScope: design and implementation of a fine grained power monitoring system for homes, *UbiComp* (2009) 245–254.
- [17] S.N. Patel, T. Robertson, J.A. Kientz, M.S. Reynolds, G.D. Abowd, At the flick of a switch: detecting and classifying unique electrical events on the residential power line, *UbiComp* (2007) 29–32.
- [18] M. Zeifman, K. Roth, Nonintrusive appliance load monitoring: review and outlook, *IEEE Transactions on Consumer Electronics* 57 (1) (2011) 76–84.
- [19] R.K. Jain, J.E. Taylor, G. Peschiera, Assessing eco-feedback interface usage and design to drive energy efficiency in buildings, *Energy and Buildings* 48 (2011) 8–17.
- [20] S. Karjalainen, Consumer preferences for feedback on household electricity consumption, *Energy and Buildings* 43 (2) (2011) 458–467.
- [21] G. Peschiera, J.E. Taylor, J.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 (2010) 1329–1336.
- [22] Tweet-a-Watt! A Safe and Simple Wireless Power Monitor, Ladyadanet Blog. <http://www.ladyadanet.net/make/tweetawatt>. (Accessed: Oct. 2, 2013).
- [23] Need Wireless? Choose XBee. *Digi XBee® Wireless RF Modules*. <http://www.digi.com/xbee>. (Accessed: Oct. 2, 2013).
- [24] M. Delmas, N. Lessem, Saving power to conserve your reputation? in: *The Effectiveness of Private Versus Public Information*, 2012, (<http://www.environment.ucla.edu/ccep/research/article.asp?parentid=13799>) (accessed Apr 17, 2013).