Managing Plug-Loads for Demand Response within Buildings

Thomas Weng, Bharathan Balaji, Seemanta Dutta, Rajesh Gupta, Yuvraj Agarwal
Department of Computer Science and Engineering, UC San Diego
{tweng, bbalaji, sedutta, gupta, yuvraj}@cs.ucsd.edu

Abstract

Detailed and accurate energy accounting is an important first step in improving energy efficiency within buildings. Based on this information, building managers can perform active energy management, especially during demand response situations that require load shedding over short time scales. While individual plug-loads are an important target for demand response, they pose significant challenges due to their distributed nature and the significant diversity of devices that are deployed.

This paper presents the design and implementation of our energy accounting and management system which is specifically geared towards managing plug-loads within enterprise buildings. Our system provides fine-grained visibility and control of plug-loads to building managers, allowing them to deal with demand response situations through user-specified actuation policies. At its core, our system consists of our wireless smart energy meter with actuation capabilities, ZigBee-based wireless network infrastructure, and our Demand Response Server, an analysis engine that provides interfaces for initiating load-shedding policies. Our micro-benchmarks show the different methods that building managers can utilize to efficiently manage devices during demand response events.

Categories and Subject Descriptors

C.3 [**Special-Purpose and Application-Based Systems**]: Real-time and Embeded Systems; J.7 [**Computers in Other Systems**]: [Industrial control]

General Terms

Design, Management, Human Factors

Keywords

Energy Metering, Wireless Sensor Network, Energy Management, Plug-Loads Management

BuildSys '11, November 1–4, 2011, Seattle, WA, USA.

1 Introduction

Managing the electrical energy consumption within buildings has become increasingly important in recent years given that they are the dominant consumer of electricity. Within the United States, buildings constitute more than 70% of the total electricity consumption[5]. Importantly, buildings are the prime contributers to peak load demand, and this peak load is incredibly costly from a grid perspective. Consequently, with time-of-day pricing, peak electricity is very expensive to consumers as well. In fact, controlling demand is vital when the grid is near capacity - too much additional demand can cause major disruption and potential blackouts. Buildings, as the largest consumers of grid electricity, have a major role to play in reducing peak demand, and one of the main mechanisms is through demand response. Demand response (DR) is the ability of a system to respond to requests to reduce energy consumption. While lighting, HVAC, and large appliances have been examined for DR, plug-loads have been largely neglected[12]. Previous efforts have demonstrated that plug-load devices consume a significant amount of electricity. For example, we previously measured that between plug-loads, HVAC, and lighting, plug-loads accounted for more than one third of the total power consumption in an enterprise building[3].

Thus given their potentially significant energy consumption, managing plug-loads can be a very effective tool for building managers (BM) to reduce energy usage during DR events. However, in order to consider plug-loads for DR two key challenges must be overcome – knowing how much energy each device is consuming (energy accounting) and having the ability to turn off devices when necessary (energy management). Both are challenging due to the fact that these loads are widely deployed within every physical space in a building, and it is not immediately known what these loads are. Unlike HVAC, which can be managed centrally by a building management system, plug-load devices must be controlled on a per-device basis. For many devices, such as desktop PCs, actuation at the outlet-level might not be an ideal or even feasible solution.

However, for many other types of devices, having actuation ability can make sense, giving BMs an additional tool in handling energy emergencies and DR events. An examination of our own building revealed a wide variety of devices used in personal offices, including phone chargers, laptops, desk lamps, space heaters, fans, and microwaves, as well as shared equipment, including copiers, vending machines, refrigerators, water coolers, coffeemakers, and televisions. Such devices can be useful to turn off for load shedding purposes during DR events. In the case of shared devices, temporarily shutting off things like coffeemakers and vending machines may not even inconvenience users significantly.

To realize this vision of better control of plug-loads, we have designed an energy metering and management system specifically targeted for DR. Our energy meter is incrementally deployable (the part cost is under $17) and has an extended set of actuation abilities that allow BMs to quickly shed loads when needed. Controlling the meters is our Demand Response Server (DRS), which contains an administrative interface that gives BMs the ability to set high-level DR policies. The DRS also gives BMs energy awareness and visibility of the plug-loads in the building. This paper focuses on the design and implementation of our DR system. Furthermore, using several micro-benchmarks we evaluate its use as a tool in managing multiple DR scenarios.

2 Background and Related Work

Energy management for buildings has been an important research topic over the past few years. Multiple projects have targeted the HVAC and lighting subsystems of buildings for improvements in energy efficiency[1, 4, 6, 10]. Several efforts have also identified the role of IT equipment as dominant energy consumers[3]. Recent work showed that PCs can be put to sleep while maintaing network connectivity[2].

Given that plug-loads are also pervasive in buildings, multiple efforts have looked at metering them. These efforts can generally be classified into two broad themes: *direct sensing*, where energy is measured directly using individual energy meters[7], and *indirect sensing* (often referred to as Non-Intrusive Load Monitoring), which uses either load disambiguation techniques[11] or indirect inference using additional sensors[8, 9, 14]. Direct sensing has been well-examined in both academia (ACme[7]) and industry[16]. However, incorporating actuation abilities to handle demand response scenarios has not been explored in depth.

Demand response (DR) is significant to both grid operators and building managers (BM). Peak demand is extremely costly - a study found that a 1% decline in peak demand would lead to 3.9% monetary savings[15]. Even more importantly, reducing peak consumption can help stablize the grid when generation is close to maximum and additional demand might cause a grid failure. Currently, DR is typically implemented by time-of-day pricing signals with the hope that higher prices will provide incentives to users to shed loads. Recent work has looked at designing the IT infrastructure needed to send such signals to buildings[13]. Lighting solutions for DR exist that utilize ballasts to shut off when signaled. Many buildings have centralized HVAC controls, and thus can change cooling set points or shut down the entire HVAC system if needed. Recently, proposals for the smart grid have called for certain appliances such as washers and dryers to have load control capabilities that can be used to schedule them. However, such load control switches are designed for heavy appliances and do not offer BMs an efficient way of controlling many smaller plug-load devices.

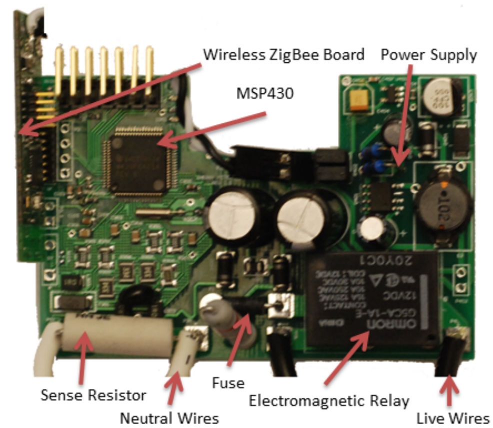

Figure 1. Picture of our energy meter with various components marked.

3 Architecture

Our system consists of our Smart Energy Meter, our wireless network, and our backend DR Server (DRS) that collects the data from the meters and allows for DR control.

3.1 Smart Energy Meter

At the center of our DR load shedding system is our Smart Energy Meter (SEM). While there are already several commercial[16] and research[7] plug-load meter designs, we developed our own meter to facilitate our DR vision.

3.1.1 Hardware Design

Our SEM node comprises five components: voltage and current sensing circuitry, energy measurement unit, power supply, a wireless radio, and a relay for switching loads on or off. The overall logical layout is shown in Figure 1. The voltage and current sense circuitry is responsible for converting the line voltages and current measurements to appropriate voltage levels for sampling by the energy metering IC (an MSP430 with ESP subprocessor). The energy metering IC then calculates various parameters, such as power and power factor and also maintains averages over time. These average values are then periodically transmitted to a base station over a ZigBee wireless radio. A mechanical relay is connected to the energy metering IC and can actuate the electrical load plugged into the SEM.

The voltage drop across a sense resistor (4mohm/4W) on the neutral line is used to measure current, while the supply voltage is measured across the live and neutral lines brought down using a voltage divider network. The Analog Front End in the MSP430 samples the voltage and current signals using a 16-bit ADC at 4kHz samples per second. This IC internally processes these samples to calculate active power, reactive power, power factor, and RMS values into designated registers periodically. The IC also contains an MSP430 core, which is a 16 bit RISC processor with up to 32KB of flash and 1KB RAM. The USART of this IC is connected to a CC2530 system-on-chip, 802.15.4 compatible radio from TI. One of the GPIOs of the MSP430 is used for switching the load using a mechanical relay (Omron G5CA). Another GPIO is connected to a switch for manual overrides. The power supply unit of the SEM is based on direct-rectification


(a) (b) (c) (d)

Figure 2. Picture of our energy meter (a, b), our SheevaPlug base station (c) that is deployed in the hallways, and the CC2530-based wireless module (d) used in our base station and energy meters.

with a buck boost converter IC (LNK304), providing up to 120mA at +12VDC. The 12V rail is used to power the mechanical relay. We use another LDO regulator to drop the +12V down to +3.3V to power the rest of the circuit.

We calibrated the SEM using a WattsUp Pro[16] and tested the power meter values to be 99% accurate for loads (both resistive and inductive) up to 1kW. The bill of material for our complete SEM is less than \$17 (in quantities of 1000). This cost does not include PCB manufacturing and component stuffing, estimated to be less than \$5.

3.1.2 Software Design and API

The software on the MSP430 controls the core functionality of the SEM. Our software comprises four basic tasks - the energy metering task, the command task, the serial task, and the safety task. The energy metering task handles retrieving values from the ESP subprocessor and calculating the final outputs that will be sent to the DRS. The serial task handles serial communications with the CC2350 radio. The command task controls the operation of the meter by recording what options and parameters the meter is configured with. The SEM supports different send modes, which allow the building manager (BM) to control the rate at which data messages are sent, including using threshold compression and averaging. These modes are set by configuration messages that come from the DRS. The safety task monitors the temperature and will shut down the connected load if the temperature goes over predefined safe values.

Our DR functionality is supported using several parameters that control how our SEM operates. First, the SEM can be set with the device type that it is connected to. Currently, this device type needs to be set by the user. The SEM can also be set with two priority levels (one for day time and one for night time) that dictate when the connected device should be turned off. These priority levels can be set by the BM, the user, or automatically (based off the device type), and essentially determine how important it is that this device stay on. These parameters will be set by the end-users or the BM via the DRS. After registering the device, the DRS will send a configuration message to the SEM specifying what type of device it is connected to, and what the priority levels are for that device. The manual override switch in each SEM can be used to turn on devices that have been remotely turned off if the BM enables that option as well.

3.2 Wireless Data Collection Network

The next component is our wireless data collection network. The network architecture is tiered with the SEMs sending data to base stations, which then relay the data to the DRS. Our base stations are equipped with both ZigBee radios and WiFi/Ethernet interfaces for connectivity to the building LAN. Our ZigBee network allows for bi-directional communication between SEMs and the base station. Based on our preliminary tests each base station ZigBee node can sustain more than 20 SEM devices transmitting data to it at once per second without any noticeable packet loss. The base stations are plug PCs (\$100) and are attached to two ZigBee nodes (\$10 each). Thus each base station can handle more than 40 nodes, so a complete set for 40 SEMs costs \$1000 (average of \$25 each).

Security is extremely important due to the fact that the SEMs can control their connected devices. Accepting forged messages would allow hackers to remotely shut down devices. ZigBee has a defined set of security features including AES encryption, trust center authorization, and key establishment and transport. Using these features it is possible to ensure that messages and nodes are authenticated and limit the possibility of fraudulent commands.

3.3 Demand Response Server

The DRS collects the energy data from the base stations and stores it in a database. The database records energy measurements as well as metadata about all SEMs. A collection of Python applications reside on the DRS that enable the actuation controls and demand response functionality. The DRS has a web-based interface for both end-users (called `MyDashboard`) and BMs. MyDashboard allows users to add their energy meters to the site and view them. When the user registers their meter with the DRS, they select the device type, such as “lamp” or “laptop” from a set of categories. The user may also change the priority level of the device if the BM allows it. Using the site, users are able to track their devices and see how much energy they are using.

The administrative interface for the BM allows them to view all the connected meters and gives them control over the entire network. Real-time energy usage information can be useful to BMs, and the system provides this information along with aggregate statistics such as total energy used across all the devices. Importantly, the BM can set a default priority level for the various device types in a building for both day and night times. Some devices have a natural higher priority during the working day, such as printers and kitchen appliances, while other devices might have higher priorities at night, such as lamps. These defaults allow automated priority level setting for most devices that are connected to the system, and will be set when the user indicates what type of device is connected to the SEM. BMs can exercise finer-grained control by setting priority levels for individual devices when needed. For example, some shared devices might have higher priority than other devices of the same type, such as a shared printer versus a private one.

This interface gives the BM the ability to set a demand response policy by setting a DR command with the desired parameters. The BM sets the parameters of the devices that they want disabled for DR purposes, e.g. all devices of a certain device type, priority level, or combination thereof. The DRS will then send out commands to the meters to shut off. Actuation by device type and priority level is extremely efficient due to the fact that the system sends out a broadcast message that labels which parameters need to shut down. Each SEM checks its own parameters (device type and priority level) against the broadcast message to determine if it needs to disable or enable the connected device. We describe the usage modalities for setting demand response policies in the following section.

4 Usage Model

The utility of our Demand Response (DR) system is in the interface that building managers (BM) can use to quickly set DR commands. There are several methods that BMs can use to enact DR events. The different options reflect the fact that a DR event might be economic, in which case reducing low priority loads is sufficient, or an emergency, in which case shedding significant loads is required.

4.1 Direct Actuation

The device parameters allow BMs to send a single command to all the SEMs for actuation. The simplest method is to turn off all devices of a certain device type, e.g. laptops. A broadcast message specifying an actuation command for the requested device type will be sent to all SEMs. If a SEM observes that its device type matches, it will shut off the device connected to it. When the DR event passes, the BM can send a message instructing all SEMs to re-enable the devices of the specified type.

Another DR mechanism is to actuate based on priority levels. Our system can send a broadcast message to all SEMs stating which priority levels to turn off. Each SEM will compare its own priority level with the one in the message, and if its priority level is lower the SEM will turn its connected load off. Turning off a device inconveniences users who might be present in the DR event - it is up to the BM to determine the best way to notify users. It is important to note that remote actuation can be intrusive to users, so the BM should make it clear that any actuation will be for DR events only.

The priority levels allow the BM to conserve energy with minimal impact by selecting lower priority devices to shut off. For example, the BM can decide to turn off devices with a priority level of 3 or less, which might include devices such as vending machines, coffeemakers, space heaters, laptops, and phone chargers. In this case, every SEM will check its priority level and compare it to the message - if it is 3, 2, or 1 it will shut off the connected device. A red LED on the meter will blink to notify the users that this meter is in a DR shut-off state. When the DR event passes, the BM can set a restore command, which sends a broadcast message asking all devices of that priority level and higher to return to normal operation. For example, if a device with priority level 4 has been turned off, and receives a restore message with priority level 5, it will not re-enable, but if it receives a restore message with priority level 3, it will. This allows the BM to stagger re-enabling devices by priority level. The BM can even combine priority level and device types, such as if the BM wants only laptops in priority level 3 shut down.

Turning on all the devices of a priority level might cause an unwanted spike, especially if there are many devices. Thus, the BM has the option of specifying a random time value (in seconds) along with the actuation command. Actuated devices will wait for a random amount of time up to the time value until they actually turn on, ensuring the staggering of devices within a priority level.

Additionally, it is possible to actuate specific devices by turning them ON/OFF, although this requires sending a message to each device individually. There are actually two different types of individual actuation messages that can be sent to the meters. The first is a direct actuation command that will shut down (or turn on) the connected device. The SEM will record this as its current mode of operation. If later the SEM receives a priority level restore command, it will not re-enable the device since its normal operating mode was already set to ”off.” The other is a demand response individual actuation command. If the SEM later receives the priority level restore command, it will turn on the device if the specified priority level matches. This enables BMs to quickly re-enable all devices that were shut off due to demand response reasons, while keeping other devices that the BM wanted off in a shut down state.

4.2 Using Occupancy Information

We have also experimented with incorporating information from occupancy sensors, based on our sensors for HVAC control[1], to augment our demand response load shedding system. These sensors utilized PIR with a reed switch to detect occupancy. Unoccupied rooms might be better candidates for load shedding and our system can take that information into consideration. For example, rather than having all priority level 2 (and lower) devices turn off, the BM can set a load shedding command for all priority level 2 devices that are in unoccupied rooms. As the DR Server will receive real-time occupancy information, the system will re-enable the devices when occupants return to their room, which effectively leads to occupancy-based energy actuation. This allows rooms that are unoccupied to continue having their devices shut down, while rooms that become occupied will have their devices restored.

4.3 Targeted Load Demand Response

Finally, a more sophisticated method for DR is to set a target kW that needs to be shed, and the system will deter-

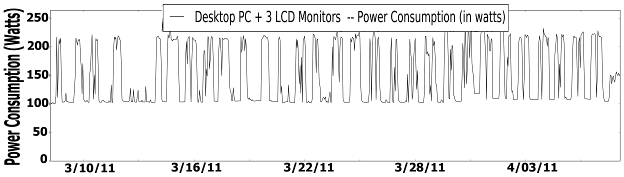

Figure 3. Power consumption of a desktop PC + 3 LCD monitors for over a week.

mine which devices to shut down to hit that target and minimize inconvenience among all of the users. The BM sets the kW target and the maximum priority level that they will allow to shut down. The system will first examine all the current loads of the lowest priority level. It looks at the past five seconds for a quick average on how much each device is currently using. It adds up the current energy consumed by all such devices for that priority level, and compares it to the target. If the target is higher, then the system knows all devices of that priority level need to be shut down. It then repeats with the next priority level, adding the energy consumed of that priority level with the energy consumed of the lower priority levels and comparing this sum against the target kW. When it gets to a priority level where it does not need to shut down all of the devices (e.g. level 5), it will send out the broadcast message to shut down all the devices of the previous priority level (e.g. level 4) and then selectively shut down individual devices of priority level 5, starting from the highest consuming device, until it hits the requested target. The system will only do this until the specified max priority level - if the target is higher than what can be achieved, the system will just shut down all lower priority levels.

5 Evaluation

We evaluate our system using micro-benchmarks for basic functionality and efficient demand response load shedding. We have deployed more than 20 of our SEMs on one floor of our building across multiple individual offices. We have collected energy usage data from different types of devices, with a majority being IT related loads including LCD screens, laptops, desktop PCs and printers.

5.1 Data Collection Results

We show a sample trace from a SEM in our deployment in Figure 3, which graphs the energy consumption over a month for a computer and three monitors, combined as a single load. The reduction in energy when the monitors go to standby mode can be seen, while the computer remains on the entire time. Several other research efforts have presented energy traces to demonstrate the diversity in energy consumption loads across different device types[7, 14].

5.2 Priority Levels Demonstration

We demonstrate the results of turning off devices based on priority levels. We test our priority level actuation in a single-person office. We have seven loads, each with a different priority level (in parentheses) - a fan (1), phone charger (1), laptop (4), lamp (5), two monitors (6), and a desktop PC (10). The priority levels were set according to what a building manager might set - chargers and fans have low priority as they can be shut off without too much inconvenience, laptops can be shut off too because they typically have batteries, and desktop computers have the highest priorities because shutting them off can have a huge adverse effect on users. We stagger turning off each priority level a few minutes apart, starting from the lowest priority level and moving up to the highest priority level.

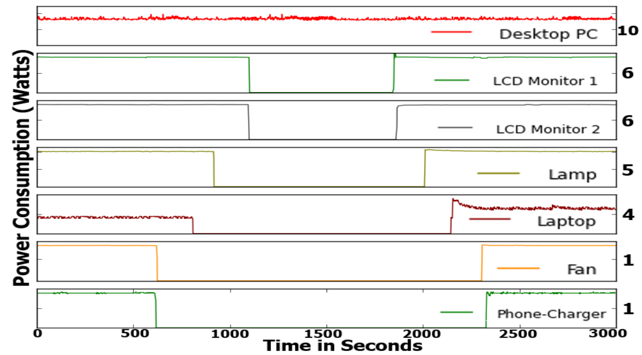

Figure 4. Priority level actuation: Notice how devices of the same priority level turn off and on at the same time (priority level listed on the right side).

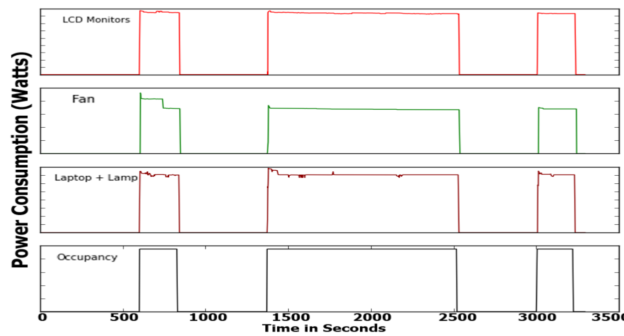

Figure 5. Results of using occupancy information along with priority levels. Notice how the devices turn off and on after an occupancy event.

As can be seen from Figure 4, the fan and phone charger (with the lowest priority of 1) both turn off simultaneously first, and the laptop and lamp follow. We restore all the devices afterwards starting from the highest priority. The higher priority monitors turn on at the same time, followed by the laptop, lamp, and finally the fan and phone charger.

5.3 Using Occupancy Information

A BM can set a load shed action for unoccupied rooms of a certain priority level. In this experiment, we deployed an office room with four devices (a lamp, LCD, fan, and laptop) in addition to an occupancy sensor. Our occupancy system is based on the design by Agarwal et al.[1] since it claimed an accuracy of 96%. The DR priority level is set so all the devices in the room will be actuated, but with the occupancy option enabled. This effectively will mean that the devices in the room will be actuated according to the occupancy status in that room. Figure 5 shows the results for an hour.

As the occupant leaves, the Demand Response Server (DRS) is notified by the occupancy sensor and will send a command to shut down all five devices, and as the occupant returns, the DRS will send a command to turn on the devices. Because of how our occupancy sensor works, there is a short

Device	Watts	PL	Device	Watts	PL
Desktop	75.6	10	Phone	3.1	9
Monitor 1	17.5	8	Monitor 2	25.7	8
Lamp 1	92.1	4	Lamp 2	33.9	4
Table Fan	29.7	3	Laptop	46.8	2
Cell Phone	2.4	1	Speaker	3.8	1

Table 1. Devices, their average power, and priority level (PL) for targeted DR test. Ten devices were tested.

delay of 15 seconds between a person leaving a room and our system registering the event. Therefore, it takes at least 15 seconds to actually turn off a device. Turning on a device however is immediate. The drawback of this approach is that messages must be sent for every device individually. We note that not all devices (e.g. laptops) are good candidates to be actuated through occupancy; it is up to the BM to determine the proper policies.

5.4 Target-based Demand Response Scenario

We demonstrate our system using target-based load shedding and allowing the DRS to automatically decide the devices to turn off. We have 10 devices over multiple rooms, listed in Table 1 with their average power and priority levels.

Figure 6 shows the combined power trace of all 10 devices for our experiment. Initially all devices are active and consuming about 350 watts. At time A, we set a target for 175 watts with max priority level 9. The DRS turned off priority level 3 devices (total consumption of about 108 watts) and lamp 1, which was the highest consuming device with priority level 4 at 92.1 watts. As can be seen, the energy consumption went from 350 watts to 150 watts, a reduction of 200 watts. We then re-enabled all of the devices by sending a priority level 1 restore command. At time B, we set a target of 300 watts with max priority level 9. The system was only able to reduce down to 75 watts however because the desktop computer had a priority of 10. With a max priority of 9 for this targeted load shedding command, the system sent out a priority level 9 shut off message, turning off all devices except the desktop.

6 Conclusions and Future Work

We have presented our energy accounting and management system for demand response control of plug-load devices. We have designed a smart energy meter (SEM) that is able to actuate its connected device, and demonstrated how our SEM can be used for handling demand response events. Tied to the energy meter is our Demand Response Server that has a web-based user interface to allow building managers the ability to visualize and control the meters. We outline different methods for enabling building managers to quickly deal with demand response events for plug load devices.

Going forward, we seek to extend the concept of placing intelligence in the SEMs and devise new ways to facilitate energy management. We are experimenting with signature detection algorithms to automatically classify device types, which will help simplify the deployment process. We hope to better minimize inconvenience to users by optimizing our control schemes. Finally, we are working with consultants to make our SEM UL Certified.

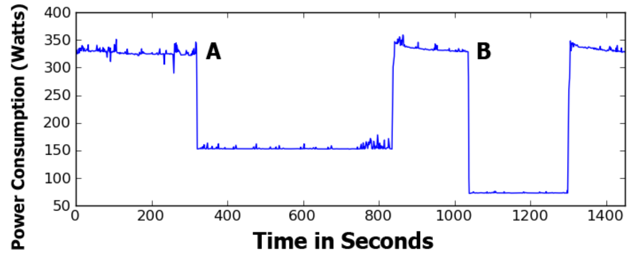

Figure 6. The targeted demand response mechanism. This is total power consumed by all of the devices.

7 Acknowledgments

We wish to thank the UCSD PPS personnel for their continued support to access the UCSD EMS. This work is supported in part by the Multiscale Systems Center (MuSyc) under the Focus Center Research Program (FCRP) supported by DARPA/MARCO, NSF grants SHF-1018632, CCF-1029783, CPS-0932360, and a Von Liebig Center CleanTech grant.

8 References

[1] Y. Agarwal, B. Balaji, S. Dutta, R. K. Gupta, and T. Weng. Duty-Cycling Buildings Aggressively: The Next Frontier in HVAC Control. In *IPSN*, 2011.

[2] Y. Agarwal, S. Savage, and R. Gupta. SleepServer: A Software-Only Approach for Reducing the Energy Consumption of PCs within Enterprise Environments. In *USENIX Annual Technical Conference*, 2010.

[3] Y. Agarwal, T. Weng, and R. Gupta. The Energy Dashboard: Improving the Visibility of Energy Consumption at a Campus-Wide Scale. In *ACM BuildSys*, 2009.

[4] D. T. Delaney, G. O'Hare, M. P. Gregory, and A. G. Ruzzelli. Evaluation of energy-efficiency in lighting systems using sensor networks. *ACM BuildSys*, 2009.

[5] Department of Energy (DOE). Buildings Energy Data Book, March 2009. `http://buildingsdatabook.eren.doe.gov/`.

[6] V. L. Erickson, Y. Lin, A. Kamthe, R. Brahme, A. Surana, A. E. Cerpa, M. D. Sohn, and S. Narayanan. Energy Efficient Building Environment Control Strategies using Real-Time Occupancy Measurements. *ACM BuildSys*, 2009.

[7] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and Implementation of a High-Fidelity AC Metering Network. In *ACM IPSN*, 2009.

[8] D. Jung and A. Savvides. Estimating Building Consumption Breakdowns using ON/OFF State Sensing and Incremental Sub-Meter Deployment. In *ACM SenSys*, 2010.

[9] Y. Kim, T. Schmid, Z. M. Charbiwala, and M. B. Srivastava. ViridiScope: Design and Implementation of a Fine Grained Power Monitoring System for Homes. In *ACM Ubicomp*, 2009.

[10] J. Lu, T. Sookoor, V. Srinivasan, G. Ge, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes. In *ACM SenSys*, 2010.

[11] A. Marchiori, D. Hakkarinen, Q. Han, and L. Earle. Circuit-Level Load Monitoring for Household Energy Management. *IEEE Pervasive Computing, Special Issue on Smart Energy Systems*, 2011.

[12] N. Motegi, M. A. Piette, D. S. Watson, S. Kiliccote, and P. Xu. Introduction to Commercial Building Control Strategies and Techniques for Demand Response. *Lawrence Berkeley National Laboratory, Berkeley, 59975*, 2007.

[13] M. Piette, S. Kiliccote, and G. Ghatikar. Design and implementation of an open, interoperable auto-dr infrastructure.

[14] A. Rowe, M. Berges, and R. Rajkumar. Contactless Sensing of Appliance State Transitions Through Variations in Electromagnetic Fields. In *ACM BuildSys*, 2010.

[15] K. Spees and L. Lave. Ceic-07-02: Impacts of responsive load in pjm: Load shifting and real time pricing.

[16] WattsUP Energy Meters. `http://wattsupmeters.com`.