From Buildings to Smart Buildings – Sensing and Actuation to Improve Energy Efficiency

Thomas Weng, *Student Member, IEEE* Yuvraj Agarwal, *Member, IEEE*

***Abstract*—**
Commercial buildings consume a significant amount of energy today and are slated to consume even more in the future. This consumption necessitates the use of carbon-producing fuels such as coal and natural gas, both of which have significant environmental impacts. While renewable energy sources remain promising, it is expected that most of the electricity generation will still use fossil fuels in the near term. Therefore, improving the energy efficiency in commercial buildings is critical, and one of the central visions of "smart buildings" is to reduce their energy use while maintaining the same level of service. In this paper we examine how such a building can be designed, focusing on the central role of actuation within buildings. We describe what is required for effective control and actuation (such as sensing), how it can be accomplished, and show the results and findings from our own deployments in a commercial mixed-use building.

***Index Terms*—Energy-aware Buildings, Smart-Buidlings**

I. Introduction

BUILDINGS consume almost 70% of the total electricity generated in the US alone. Commercial buildings account for over half of this electricity usage, and their share of energy consumption is projected to increase even further as compared to residential buildings, industry, and transportation. Furthermore, commercial buildings are increasingly becoming mixed-use, that is, they now house both human occupants and significant energy-consuming IT equipment such as desktop computers, monitors, printers, and servers. A modern mixed-used building will therefore typically have four major energy-consuming subsystems: HVAC, lighting, IT equipment, and miscellaneous plug-load devices[2], [4], [16].

Figure 1 shows the breakdown of a typical mixed-use building (the CSE building) on our campus at UCSD. Energy consumption of lighting, often thought of as quite significant, is actually low, while HVAC and plug-load devices (which include IT devices) are the most dominant consumers. This particular building is only seven years old and has a modern centrally managed HVAC system[4]. A central air handler produces chilled air, which gets circulated through the ducts that span the entire building; VAV boxes reheat and send the air to the actual zones (which are the offices and labs). Static schedules determine the HVAC operation, with 6AM to 6PM being the operative hours. This simplistic policy is actually the standard way of controlling HVAC in commercial buildings, and causes significant energy waste. For example, most occupants might not get in until 8AM or 9AM, and even during the day there will be periods of vacancy (such as for lunch or meetings). The same is true for plug-load devices. Many devices are left powered on regardless of actual occupancy or needs. Computers and monitors might be on 24/7 even when they are not being actively used; this is also the case for many other plug-load devices, such as desk lamps, space heaters, and fans.

Authors are with the Department of Computer Science and Engineering, UC San Diego, La Jolla, CA 92093-0404. Email: {tweng, yuvraj}@cs.ucsd.edu

From Figure 1, it is apparent that both HVAC and plug-load subsystems peak at around mid-day, which is precisely when aggregate demand is at its highest throughout an entire region. This peak demand causes significant issues for local utilities and necessitates the use of very costly carbon-producing peaker plants particularly during warmer (or colder) than normal days. Currently the utilities absorb most of the cost for this peak electricity (which can be up to $1 per kw-h, as compared to the usual 5-7 cents per kw-h), but in the future, a major fraction of this cost will likely be passed on to building owners themselves. The high energy consumption of HVAC and plug load devices has consequences for both energy costs as well as the environment, and going forward this problem will only get worse.

Therefore, the future "smart building" must be able to reduce its energy usage, and fine-grained control over the HVAC system and plug-load devices is the critical component needed to meet this challenge. We define building actuation as the ability to directly control the operation of various building systems (e.g. turning off HVAC or turning off a plug-load device), with one of the key goals being the reduction of energy consumption by shutting off unused equipment and eliminating energy waste. This is important for two reasons. First, for energy efficiency reasons, simply powering off unnecessary loads will reduce overall energy usage while minimally impacting building occupants. Second, reducing energy loads during periods of peak-demand to the grid can lead to substantial monetary savings. This ability of reducing energy usage, called demand response, will become a priority for utilities going forward. Being able to actuate different building subsystems, both for energy efficiency and demand response reasons, is therefore of critical importance in the design of a smart building.

Actuation comprises two components - the *mechanisms* to control various building subsystems and the *policies* that determine when the control should be exercised. In order to actuate effectively, it is critical to know the operational status of the building at fine temporal and spatial granularities. This task essentially boils down to sensing various physical attributes – occupancy, internal and external environmental conditions, and energy usage – using either existing sensors or by augmenting the building with additional sensors.

In this paper we describe several energy-saving architectures[2], [3], [16] that we have developed to actuate the HVAC system, IT equipment such as PCs, and other plug-loads within an actual building and highlight research efforts in this space from other groups. We also describe our deployment experiences and show preliminary results that demonstrate significant energy saving potential. Finally we discuss the challenges that future efforts will need to address to truly make smart buildings a reality.

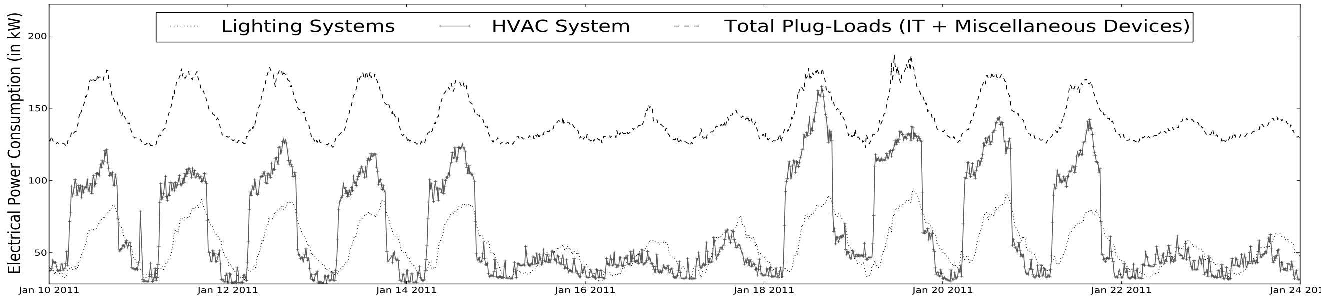

Fig. 1. Electrical power usage breakdown (in kW) for a typical mixed-use building over two weeks in 2011. HVAC and plug-loads dominate power usage as compared to lighting. Plug-loads include both IT related loads and miscellaneous devices, which we cannot differentiate between at the level of the entire building.

II. SENSING

Sensing is a key function of smart buildings, and therefore a significant amount of research has examined various sensing modalities and techniques. Of critical importance is the actual sensing infrastructure, and wireless sensor networks have emerged as enablers for delivering sensor data. Two current standards exist that are popular for wireless sensor networks in buildings - ZigBee and 6LowPan; both utilize 802.15.4 radios and are geared towards low-power wireless networks.

For our own building deployment, we built our network on top of ZigBee protocols in a three-tier topology, with a centralized smart building server collecting data from deployed basestations, which in turn receive data from nearby wireless sensor nodes. Our custom wireless boards are designed to be low-cost ($10 in large quantities) and allow sensors to be attached via I/O pins. Each board contains a TI CC2530 as the main RF chip, which runs a custom ZigBee stack and can communicate with the base stations. Our basestations consist of a Linux-based plug computer (SheevaPlug) connected to one of our wireless boards via USB. With the sensing fabric in place, the pertinent question becomes what building operations should be monitored. Occupancy and user context, environmental conditions, and energy usage are three important ones that should be sensed in order to know the building status in real-time.

A. Occupancy and User Context

In order to design effective actuation policies, it is necessary to detect occupancy at a fine granularity within a building. PIR sensors are the most ubiquitous form of sensing in modern buildings, but suffer from accuracy problems due to the fact that they are in reality motion sensors (and not presence sensors). Thus, false positives (when the sensor detects a person, but no one is actually there) and false negatives (when a sensor fails to detect a person in the room) are prevalent.

We have designed an occupancy sensor that improves upon PIR in order to eliminate a significant amount of this error[1]. Our occupancy sensor combines a PIR with a magnetic reed switch that can determine when an office door is closed or open. When the office door is open, the sensor assumes that the room is occupied (for most occupants in our building, leaving the door open means they are either inside or are away temporarily). When the office door is closed, the PIR sensor is used to determine if the room is occupied or unoccupied. If the PIR detects movement, then a person is still in the room, and if it does not, then it likely means that the occupant left the office. Figure 3 shows a picture of our occupancy nodes installed on a typical office wall. Based on a test deployment of our occupancy sensors on one floor of our building and comparing their output with ground truth data over a day we found that our sensors were 96% accurate, with the majority of the errors being caused by improper placement and installation[1].

In addition to sensing occupancy, knowing user context will be very useful in optimizing building operations. For example, knowing what users are doing (such as if they are actively using their IT equipment and other devices) can lead to better policies for plug load device and IT management. Software that runs on users' computers can monitor such information. This can be combined with binary occupancy to further improve accuracy (if a room has been marked as unoccupied, but the computer is being physically used, then obviously the unoccupied reading was incorrect) or reveal additional useful information for more effective actuation. We briefly note that there are privacy implications for having this much data on user occupancy and activity. For our own system, we do not allow anyone to have access to this data - however, every organization must determine their own policies regarding privacy and access to the occupancy statistics.

Other efforts have looked at other ways to measure occupancy. Camera sensing for example has been used by researchers from UC Merced[5]. In this case, cameras mounted in hallways of a building record when people transition from one area to another. Using statistical models to augment the data collected from the camera sensors, their system attempts to model the occupancy of the building. While PIR sensors are the most popular deployed sensors in buildings, ultrasonic sensors can also be used - these utilize sound waves and the Doppler principle to detect movement in an area.

B. Environmental Sensing

Another important area is environmental sensing, namely the temperatures and other variables for the zones in a building. Most modern buildings come installed with building management systems (BMS) that monitor (and control) these conditions. We have come up with our smart building server that can interface with existing BMSes and retrieve these values through standardized protocols such as BACNet. Through interfacing with the existing BMS, we can obtain on a per-zone basis the temperature, damper position, fan speed, cooling setpoints, and user override settings. This information is important for optimized control over HVAC settings. The environmental

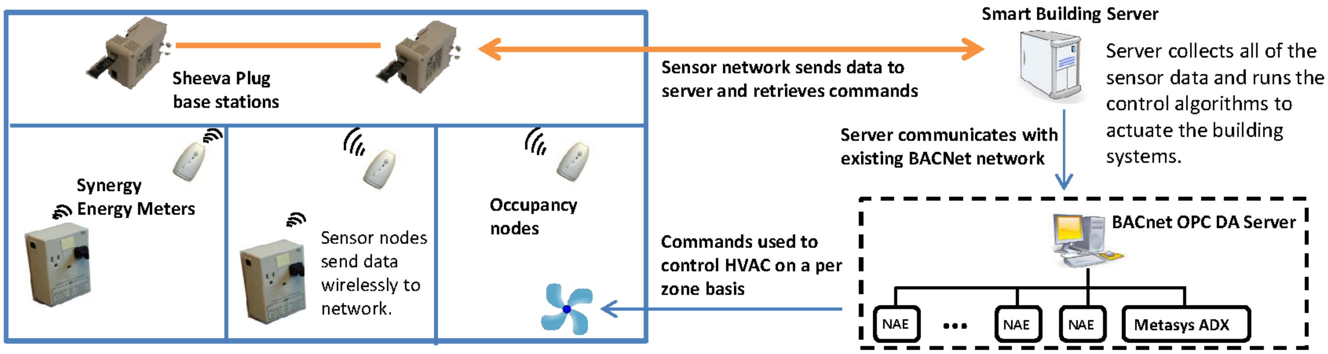

Fig. 2. Overall architecture of the building actuation system, including the smart building server, wireless sensor nodes, basestation, and connection to the facilities' BACNet network.

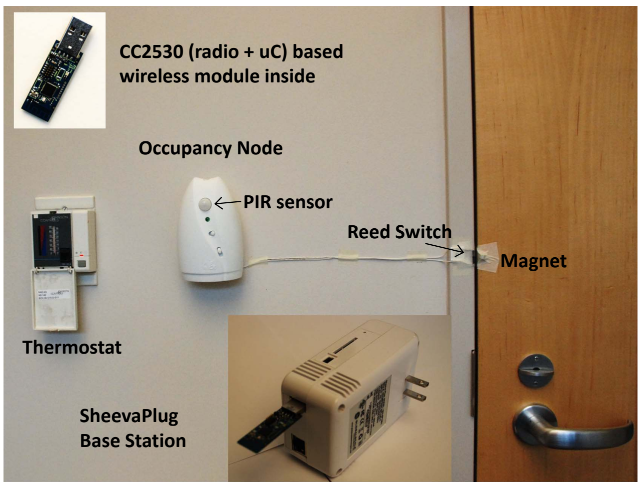

Fig. 3. Occupancy node deployed on the wall of an office, with its reed switch, PIR sensor and our CC2530 based radio module also shown.

sensing component consists of an OPC Tunneler (managed by us) communicating using OPC protocols to the BACNet OPC Data Acquisition Server (hosted by the campus facilities group). The OPC Tunneler retrieves values from the facilities' OPC DA Server, which reads the actual environmental data from the Network Automation Engines (NAE) installed throughout the building.

Installing additional environmental sensors can also be beneficial. Researchers from UVa have examined using light sensors mounted on the outside of a building to provide accurate information on current external light conditions, which they used to adjust variable lightning to save energy[12], while researchers from UCLA have developed a water monitoring system to measure fine-grained water consumption in a building through the use of pipe-attached vibration sensors[9].

C. Energy Sensing

The last important area is detailed energy monitoring, that is, identifying the energy consumption of specific building subsystems. The most accurate way of determining macro breakdowns is through the use of industrial mains meters. We have installed Schneider Electrical power meters throughout the CSE building to measure the energy consumption of HVAC, plug-loads, and lighting in our building. Doing so gives us a clear picture on the relative energy consumption by each subsystem. Figure 1 from the first section shows a breakdown of our particular building.

The breakdown reveals an interesting fact - plug-load devices are a significant consumer of energy; therefore knowing where the energy is specifically going is important. Two broad classes of plug-load energy detection exist - direct sensing, which connects a meter directly inline with the device, and indirect methods, which attempt to measure energy usage without having to deploy a meter for every device. Indirect sensing in particular has seen a significant amount of research, and we highlight some of the approaches that other researchers have developed. Non-intrusive Load Monitoring techniques are varied - one of the original attempts sought to disambiguate energy loads from a central energy meter[7]. Since then, more sophisticated attempts have been tried that use learning algorithms[14], [6]. Deployment of indirect sensors has also been popular. Such techniques utilize sensors to detect the magnetic field variations[10] and magnetic state transitions[15] that occur near plug-load devices.

While indirect methods have deployment advantages, direct methods are more accurate, and several plug-load meters have been developed that are able to measure energy usage from a connected device[11]. Researchers at UC Berkeley have developed a wireless meter called ACme, which uses 6LowPan to communicate[8], while several commercial meters exist as well. One drawback however is that these meters tend to be expensive. We have developed a wireless plug-load meter that is low cost and allows the measuring of a multitude of variables[16]. The core of our Synergy Energy Meter (SEM) is a TI MSP430 chip with built-in energy meter front end. Our wireless sensor node is connected to the main energy meter board via a serial port. The SEM measures the voltage, current, watts, and power factor of the connected plug-load once every second and will transmit it through the wireless node to the parent basestation. Figure 4 shows a picture of our energy metering system, including the basestation and wireless node. We tested our SEM and found it to be up to 99% accurate for loads (resistive and inductive) up to 1kW.

III. Actuation

With the sensing aspect addressed, we now focus on actuation. The ability to control the building operations is essential in improving energy efficiency, and the key challenges that must be addressed relate to exercising control over building subsystems (mechanisms) and determining when to control them

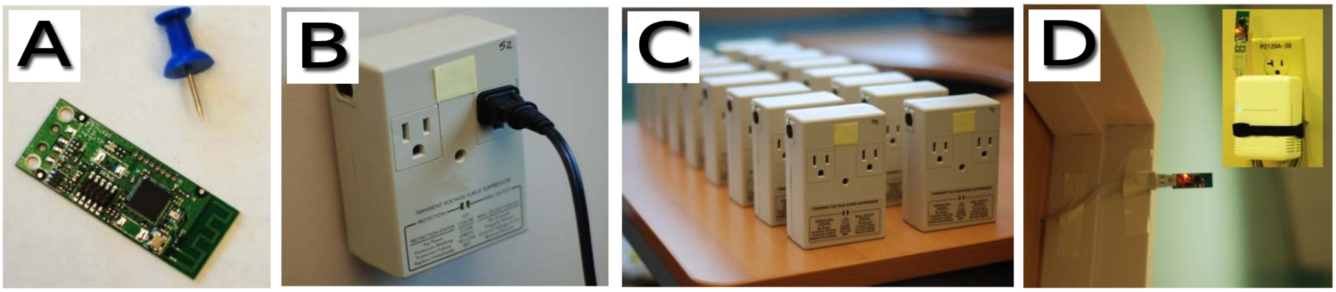

Fig. 4. Picture of the CC2530-based wireless module (a) used in our base station and energy meters (b,c). The data from the sensors is collected by various SheevaPlug base stations (d) deployed in the hallways.

(policy). The significant energy consumers in a building are HVAC equipment, IT equipment, lighting, and miscellaneous plug-loads. Each must be approached differently in terms of mechanisms for actuation, and we will discuss our research and deployments for actuating HVAC and miscellaneous plug-load devices. Figure 2 shows the overall architecture of the system and its components.

A. HVAC System

Modern buildings with BMS systems not only sense the environmental conditions, but also control the per-zone environmental settings. These controls include setting the thermal setpoints for each zone (the temperature that the HVAC system should meet) and the command state for each zone (whether the HVAC system should be on or off). A smart building must be able to access these controls in order to actuate the HVAC system on a fine-grained basis.

Our smart building server achieves actuation ability with the BMS through the same protocols that it uses to receive environmental information. The server will write values back to the OPC server, which will send the commands to the NAE boxes. The two actuation parameters that can be set using this method are `setpoint` and `command state`. The latter essentially puts a zone into one of two possible HVAC states - on, in which case the HVAC system will open the dampers and send conditioned air to meet the setpoint, and off, in which case the damper is minimized and only a minimal amount of airflow is released. In both cases, our building server does not affect the specific operation of the HVAC system - only final zone commands are set, which means that the BMS is still in charge of determining the amount of cool air to produce and what adjustments it needs to make to the dampers.

With the mechanism in place, we need to determine what zones to actuate and in what situations to do so. Having occupancy as an input allows us to dynamically control the HVAC settings per zone based on if the room is occupied or not. We have deployed an occupancy-based HVAC control system in our building, and we opted to use a simple reactive system that sets the command state of a zone to "on" when the zone becomes occupied, and to "off" when the zone becomes unoccupied. More complex policies can be designed, such as using predictive models along with real-time occupancy status[13]. In all such cases, occupancy is still the central driver for control of the HVAC system. The occupancy information is collected in the smart building server and stored in a database. When a change in occupancy is detected, the smart building server will send a corresponding command to turn on or off that room's HVAC.

Figure 5 and Figure 6 highlight the potential savings that our system can achieve during an experiment we ran in our building over two warm days[2]. The first test day was a control baseline day where the HVAC system ran as normal (using the normal static schedule where the system fully ramps up by 6AM and goes into "off" mode at 6:30PM). During the experimental day, we deployed our occupancy sensors across one floor of our four-floor building and controlled the per-zone HVAC using real-time occupancy. Most building occupants arrive anywhere from 8AM to 10AM, thus the HVAC systems started much later during the test day. Additionally, many people left their offices during the day for meetings and lunch breaks, and thus the HVAC system in those zones was turned off for those periods even during working hours. Finally, many of the office staff workers left at 4PM (rather than 6:30PM) enabling the HVAC system in certain zones to be turned off earlier. When combined, these periods of absence led to significantly reduced HVAC usage. The total HVAC electrical load for test day 1 was 1556.11kW-H while the total HVAC electrical load for the baseline day was 1759.9 kW-H. Therefore, in terms of electricity, our HVAC control scheme saved a significant 11.59%, despite only controlling one floor in a four story building. We also note that the thermal load consumption was less than the base-line as well, saving 12.41% in thermal cooling loads and 9.59% in thermal heating loads. We estimate savings in excess of 30% if our system is deployed across the entire building.

This reactive policy can save a significant amount of energy, but more sophisticated policies can be used. Other researchers have examined using predictive approaches in order to pre-cool and pre-heat rooms before occupants return. These approaches have been shown to save energy in building simulators[5], [13]. Going forward, more advanced control policies should be developed on top of the HVAC control mechanisms in order to minimize energy consumption while maintaining occupant comfort.

B. Miscellaneous Plug-Load Devices

Actuating plug-load devices is also important towards improving building energy efficiency. Many devices are left on, even when the user is not present, and thus significant aggregate energy is wasted when all of these devices are accounted for. A natural way to remotely actuate devices is at the level of each

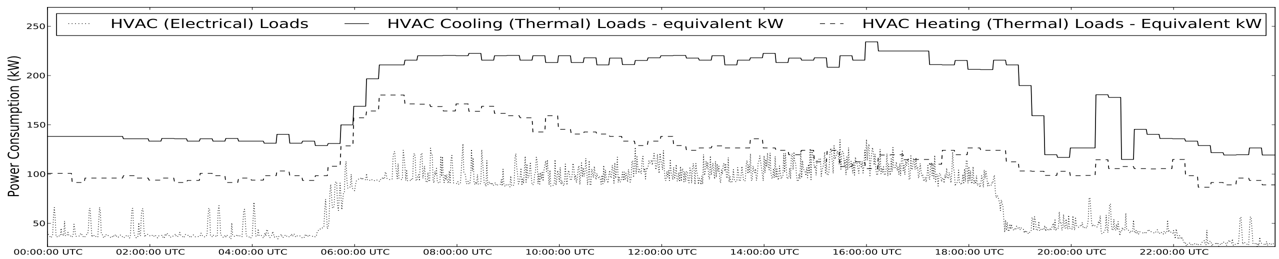

Fig. 5. The energy consumption of HVAC during our baseline day. We show HVAC electrical loads as well as the HVAC thermal loads for both cooling and heating (as equivalent kW).

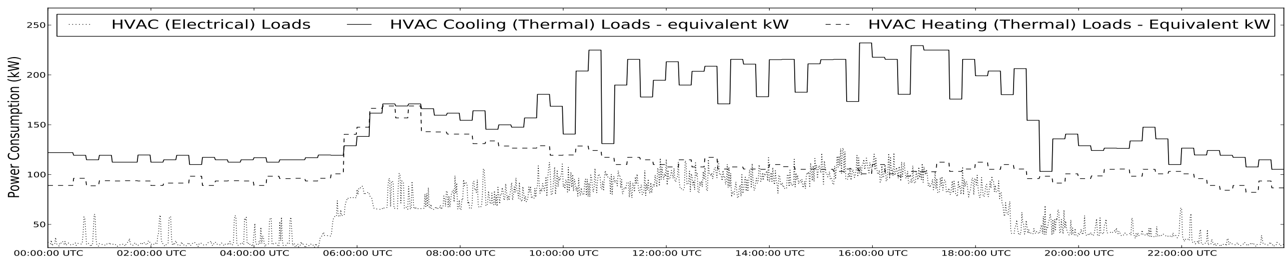

Fig. 6. The energy consumption of HVAC during our test day. We graph both electrical loads and thermal loads (as equivalent kW). The HVAC-electrical savings compared to baseline shown in Figure 5 are 11.59% while the HVAC-thermal savings are 12.41% and 9.59% for cooling and heating loads respectively.

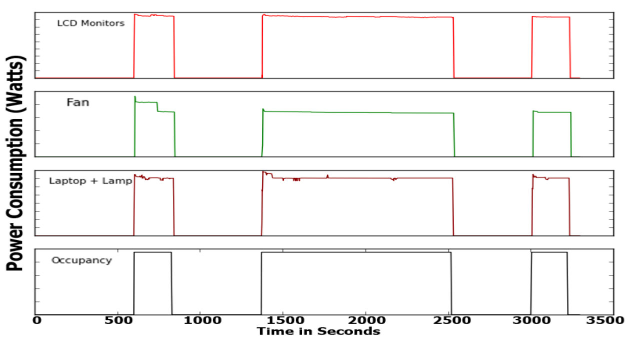

Fig. 7. Results of using occupancy information to control devices. Notice how the devices turn off and on after an occupancy event.

outlet; a connected plug-load meter such as our Synergy Energy Meter (SEM) can disconnect the electricity and shut off the device. A mechanical relay is connected to the MSP430 in our SEM that allows it to turn on or off the electrical load plugged into the meter, while our software allows remote management over the wireless network.

The smart building server can utilize occupancy to determine when to control plug-load devices. For certain types of devices, this can make sense - lamps and monitors for example can be shut off when their locations become unoccupied without any significant impact. Figure 7 shows the results of using occupancy to remotely actuate devices in an office room for an hour. As the occupancy changes are detected, the smart building server will send commands to turn off or on the connected device to the energy meter. Various policies, other than occupancy-based actuation, can also be used to control plug-loads[16].

C. IT Equipment and Lighting

We briefly discuss actuation for IT equipment and lighting. Actuation for IT equipment such as computers might not be best done through direct plug-load control. Instead software can be installed on individual computers that allows access to the power management features of the system. Specifically, the software will be able to put the computer to power saving sleep modes.

This functionality can be combined with occupancy sensors to reactively put the computer to sleep when the user is no longer occupying the room. One issue that arises when putting computers to sleep is that they lose all network connectivity when in sleep mode. This problem discourages people from putting their computers into sleep. To solve this, we have developed the SleepServer[3] architecture that allows computers to go to sleep while maintaining network connectivity.

Techniques to manage lighting loads have already been extensively examined in both industry and in research. Electronic ballasts have been used to facilitate demand response management for light fixtures while centralized lighting controls can be used to remotely actuate installed lights throughout the building. Coupled with occupancy and user context, the smart building server can also turn off or reduce lighting when a room becomes unoccupied.

IV. Discussion and Future Work

In this paper we have discussed the importance of actuation in making buildings more energy efficient. Our research efforts have led to the development of several mechanisms for actuating the HVAC system and controlling miscellaneous plug-loads, as well as an occupancy-based policy to turn off these systems when the rooms are no longer being occupied. Many other areas of smart building research, such as modeling and prediction of building operations, can be used to augment and improve the control over a building. For example, being able to predict when occupants plan on going into their offices means that the system can pre-cool the area prior to the occupant's arrival. Improved analysis of the sensed information can also lead to better understanding of the building processes and their inter-dependencies. Actuating one building process might have unintended effects on another, and only through a combined actuation-sensing approach can we capture this relationship.

One important consideration is the economics of deploying smart building technologies. Smart building systems (such as the sensors and actuators) cost money to deploy and maintain; thus the return on the investment must be adaquate for any organization to pay for the sensor systems. An estimation for our own building reveals that the total cost for installation is approximately the same as the yearly energy savings at $.13/kW-H. This means a building can recover the installation costs in one year through the reduction in energy usage. While this will be different for every building, we believe that smart building systems are extremely economical and provide monetary value for buildings that choose to invest in them.

Going forward, we envision a smart building system capable of holistically controlling all of the building processes. Such building processes would include not only HVAC and plug load devices, but also IT and lighting. User context can also be used to control various building systems. The thermal conditions of each room would be set automatically according to each individual's preferences. Developing control algorithms that can optimally control all of the building processes is an ongoing technical challenge. Ultimately, the result of these technologies will be a building that not only significantly reduces energy consumption, but also improves the quality of service for every occupant.

Many non-technical challenges remain however. These technologies must eventually be implemented by building designers, and thus must mature beyond research prototypes. For example, software systems must present intuitive user interfaces for building administrators and technicians who may not be fluent with information technology. Aggressive actuation to reduce energy can potentially affect occupants adversely if the algorithms are not carefully designed. Policy challenges must also be addressed - building administrators must determine how best to control the building in order to meet the multiple demands of energy conservation and occupancy comfort. While the technical challenges in developing the energy-efficient smart building are important for researchers to solve, it is vital that these other issues are considered as well.

References

[1] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng. Occupancy-Driven Energy Management for Smart Building Automation. In *ACM Workshop on Embedded Sensing Systems For Energy-Efficiency In Buildings*, 2010.

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

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

[4] 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.

[5] V. Erickson, M. Carreira-Perpinan, and A. Cerpa. OBSERVE: Occupancy-based system for efficient reduction of HVAC energy. In *IPSN*, 2011.

[6] S. Gupta, M. S. Reynolds, and S. N. Patel. ElectriSense: Single-point Sensing using EMI for Electrical Event Detection and Classification in the Home. In *Ubicomp*, pages 139–148. ACM, 2010.

[7] G. Hart. Nonintrusive Appliance Load Monitoring. *Proceedings of the IEEE*, 1992.

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

[9] Y. Kim, T. Schmid, Z. Charbiwala, J. Friedman, and M. B. Srivastava. NAWMS: Nonintrusive Autonomous Water Monitoring System. In *SenSys*, 2008.

[10] 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 *Ubicomp*, 2009.

[11] J. Lifton, M. Feldmeier, Y. Ono, C. Lewis, and J. A. Paradiso. A Platform for Ubiquitous Sensor Deployment in Occupational and Domestic Environments. In *Proceedings of IPSN*, 2007.

[12] J. Lu, D. Birru, and K. Whitehouse. Using simple light sensors to achieve smart daylight harvesting. In *BuildSys*, 2010.

[13] 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.

[14] 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.

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

[16] T. Weng, B. Balaji, S. Dutta, R. Gupta, and Y. Agarwal. Managing Plug-Loads for Demand Response within Buildings. In *ACM Workshop on Embedded Sensing Systems For Energy-Efficiency In Buildings*, 2011.