Duty-Cycling Buildings Aggressively: The Next Frontier in HVAC Control

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

ABSTRACT

Buildings are known to be the largest consumers of electricity in the United States, and often times the dominant energy consumer is the HVAC system. Despite this fact, in most buildings the HVAC system is run using primitive static control algorithms based on fixed work schedules causing wasted energy during periods of low occupancy. In this paper we present a novel control architecture that uses occupancy sensing to guide the operation of a building HVAC system. We show how we can enable aggressive duty-cycling of building HVAC systems – that is, turn them ON or OFF – to save energy while meeting building performance requirements using inexpensive sensing and control methods.

We have deployed our occupancy sensor network across an entire floor of a university building and our data shows several periods of low occupancy with significant opportunities to save energy over normal HVAC schedules. Furthermore, by interfacing with the building Energy Management System (EMS) directly and using real-time occupancy data collected by our occupancy nodes, we measure electrical energy savings of 9.54% to 15.73% and thermal energy savings of 7.59% to 12.85% for the HVAC system by controlling just one floor of our four floor building.

Categories and Subject Descriptors

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

General Terms

Design, Management, Human Factors

Keywords

Occupancy Detection System, Wireless Sensor Network, HVAC Control System

IPSN'11, April 12–14, 2011, Chicago, Illinois.

1. INTRODUCTION

According to the US Department of Energy, buildings accounted for about 38.9% of US primary energy consumption in 2006, 74% of which is electrical energy[7]. This electrical usage is roughly divided equally between residential and commercial buildings. Consequently, several efforts by the DOE and the research community[2] have begun to analyze energy use within buildings to identify the dominant energy loads. Recent research shows that depending on the specific use modality of the building, the dominant electricity consumers can be lighting, computing infrastructure, or what is most often the case, heating ventilation and air-conditioning systems collectively referred to as HVAC[1, 2].

Depending on external environmental conditions, HVAC energy usage can easily dominate within buildings. As an example, the measured HVAC load for the CSE building on the UC San Diego campus is between 25% to 40% of its total electricity load over the entire year. This building is only six years old and has an efficient and optimized HVAC system that is centrally managed. In fact, all the buildings around the campus are managed by a central Energy Management System (EMS) and are operated on a static occupancy schedule, which is set to 5:15AM to 10:00PM during weekdays. This simplistic policy of HVAC management is actually predominant and widely used in commercial buildings in accordance with standard working hours.

While fixed schedule-based HVAC energy management is often used because of its simplicity, it causes significant wasted energy due to unnecessary cooling regardless of occupancy levels. To manage building operations, building operators will identify and divide building operations into thermal zones. A thermal zone represents a single domain of sensing and control. HVAC systems can use zonal control to, in principal, direct energy use based on occupancy in zones, activating and deactivating cooling depending upon the occupancy within the zones. This ideal HVAC system would then have a baseline energy usage and any additional energy use would depend on the number of zones occupied, leading to its energy consumption effectively being proportional to the number of people within the building. Obtaining real-time occupancy information on a per zone basis is therefore critical for fine-grained HVAC control strategies.

There are two key components in realizing this vision of an aggressively duty-cycled HVAC system. The first component relates to sensing – determining fine-grained and accurate occupancy information. While movement sensors such as passive infrared (PIR) devices can be used, we show that for accurate occupancy detection multiple sensors are re-

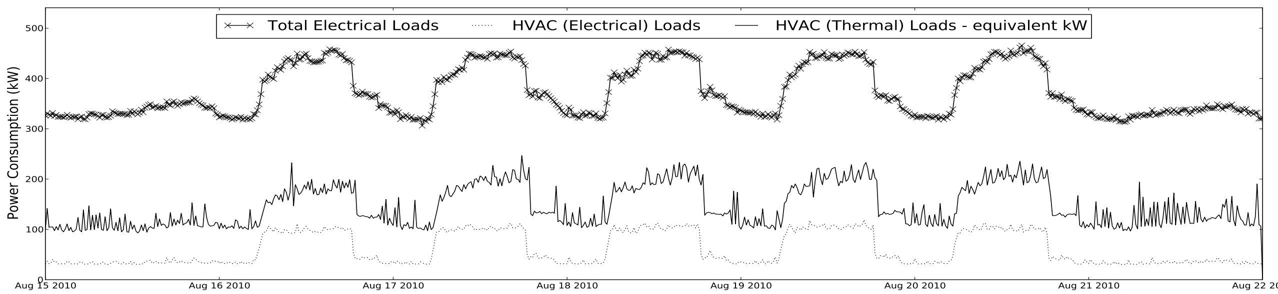

Figure 1: Power consumption breakdown for our building. The HVAC electrical load is between 25% to 33% of the total electrical load. The HVAC thermal load, as expressed in kW equivalent of cooling, is also significant.

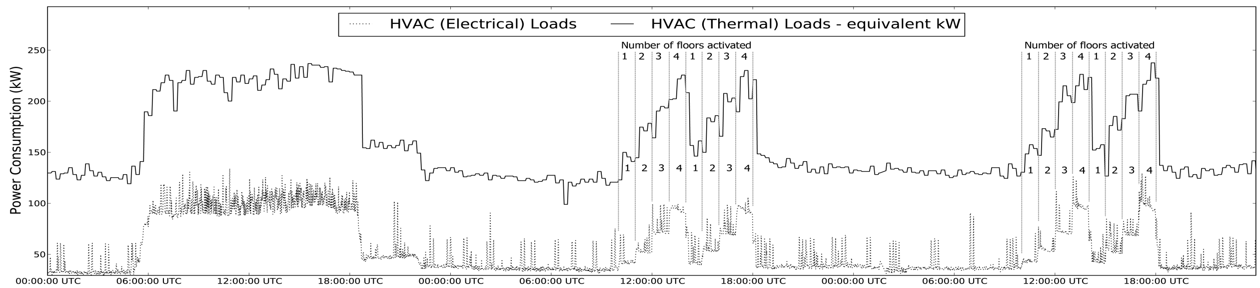

Figure 2: HVAC power consumption - Friday Oct 22nd to Sunday Oct 24th. The first day is a regular day (Friday) that shows the HVAC as is without any duty-cycling and all floors are ON as normal. The next day is Saturday when we turn on one floor at a time – 1 floor HVAC on (10AM – 11AM), then 2 floors (11AM – Noon), 3 floors (12PM – 1PM), 4 floors (1PM – 2PM), then turn all floors off and repeat from 2PM to 6PM. We repeat the same experiment on Sunday.

quired. Addressing deployment challenges of these sensors within existing buildings is also important. In particular, these occupancy sensors must be wireless and low power so that they are cost effective to deploy and can run on batteries for several years. Finally, the network architecture to collect the real-time occupancy information must be designed carefully to be robust and reliable.

The second component relates to control – determining mechanisms to interface with, and actuate, the building HVAC and EMS systems. This is particularly challenging since most building EMS systems use several proprietary protocols and access to these systems is difficult to obtain due to policy issues.

Several research efforts[6, 16] have investigated using occupancy sensors to control lighting and HVAC systems. Most of these efforts are either geared towards a home scenario with a single air-conditioning system[16], or consider much smaller settings with a single sensor and a single HVAC unit per room[10]. Our own recent work in this space has addressed obtaining accurate occupancy information[1] but does not address controlling an actual HVAC system or deploying sensors at a large scale.

This paper primarily focuses on the experimental design and deployment experiences of our system and makes the following contributions: the design and implementation of a low-cost and high-accuracy wireless occupancy sensor node; a control architecture to actuate individual HVAC zones based on occupancy information; system evaluation in terms of energy savings enabled by our system; and our experiences in deploying the occupancy nodes and controlling the building HVAC system.

2. BACKGROUND AND MOTIVATION

In our previous work[2] we presented the detailed breakdown of energy use in a typical mixed-use building (the Computer Science and Engineering building at UCSD). The CSE building contains four floors and a basement level which houses several undergraduate labs. Through our extensive submetering, we have been able to break down the electrical and thermal energy usage of our building's HVAC system. The HVAC related energy use consists of components such as air handlers, pumps, fans and actuators. Furthermore, our building HVAC system is connected to a central campus-wide chilled water loop that acts as a heat exchanger to cool the air passing through the air handlers. We have thermal meters installed on the inlets and outlets of this chilled water loop to measure the 'thermal energy' in MMBTUs that can be converted to equivalent kW based on the aggregate energy sourcing architecture as supplied by our campus energy managers. Note that since the chilled water is produced partly as a by product of our campus natural gas fired power plant, the conversion factor into kW is an approximation and should not be taken literally as the energy it would take to chill water using electric power.

Figure 1 illustrates the power consumption breakdown for our building for a typical week in the summer (August 2010). As can be seen from the graph, HVAC Electrical loads (air handler units, pumps, fans) rise in the morning during the weekdays, stay high during the day and reduce only in the evening. This is due to the fact that our building, like most buildings across our campus is run on a fixed 5:15AM to 10:00PM schedule.

The HVAC Electrical loads are between 25% and 33% of the total building electrical load during the entire week.

The HVAC Thermal equivalent load (expressed in kW) is also significant and follows a similar pattern as the HVAC electrical load – it is higher during work hours on weekdays and is low during the weekends. The combined HVAC load (electrical and thermal) during nights and weekends is due to the fact that a minimum air flow must be maintained for the basement laboratories at all times. Furthermore, our building has a small server room that requires constant cooling. We point out that there are power spikes when the fans turn on or off to increase the air flow, but that these are brief compared to the base power draw.

While Figure 1 provides insight into the energy consumed by the HVAC system in our building, we wanted to know what *potential* impact variations in occupancy have on the energy consumption of the HVAC system. This would provide a bound on the potential energy savings a detailed occupancy driven HVAC system may have. To do this, we designed an experiment where we would initially turn off the HVAC in all four floors, and then incrementally turn on the HVAC system on each floor to see the change in energy consumption. Over a particular weekend (Oct. 22-24, 2010), we implemented this experiment on both Saturday and Sunday, turning on the HVAC on the first floor from 10AM to 11AM, the first two floors from 11AM to 12PM, the first three floors from 12PM to 1PM, and all four floors from 1PM to 2PM. We repeated this experiment from 2PM to 6PM. The weather for the two days were fairly mild, with a high of 77°F - this is typical of San Diego.

Figure 2 illustrates the results of our experiment. The figure also shows the HVAC electrical and HVAC thermal loads for the previous Friday as a baseline of how much energy is consumed during a normal weekday under the normal HVAC schedule. From the figure, we can clearly see that the energy consumed by the HVAC increases as additional floors are activated. Interestingly, the marginal energy consumed by the fourth floor is actually larger than the marginal energy consumed by the third floor. This means that the closer the HVAC system is to maximum cooling, the more costly cooling each additional zone becomes. This is partly due to the fact that power consumption of the air handlers is cubic with fan speeds. What this implies is that reducing HVAC loads during the workdays for even a few zones can potentially have significant savings in energy. In other words, the sensitivity of energy consumption to occupancy increases at higher occupancy levels, making a very strong case for dynamic HVAC control.

3. SYSTEM ARCHITECTURE

In this section, we describe several important design considerations for an occupancy-based HVAC control system. The main architectural components of our system are: (a) occupancy sensor nodes that are able to accurately detect presence of occupants in individual zones; (b) a wireless network infrastructure that is able to reliably collect data from occupancy nodes and relay it to our back-end server infrastructure for data collection and analysis; and (c) a control architecture that interfaces with the existing building EMS systems to actuate the HVAC system selectively based on occupancy. Figure 3 illustrates the overall architecture at an abstract level.

3.1 Occupancy Nodes

The design of our occupancy nodes is based on several

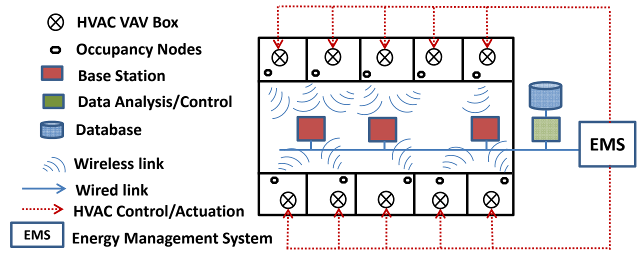

Figure 3: System Architecture - On an abstract level, these are the components of our occupancy sensing / HVAC control system.

requirements. First and foremost is ease of deployment, especially within existing buildings. Low power is another important consideration. The choice of the wireless radio also is crucial since it affects the range, cost, power consumption and data rates of the occupancy sensors.

The second requirement is detecting occupancy accurately. Passive infrared (PIR) sensors are commonplace, but these only detect movement and not necessarily occupancy. In many buildings they are used for localized lighting control, but PIR sensors by themselves do not work well for HVAC. By the time a person realizes their PIR sensor mistakenly detected them as absent, their HVAC has long since turned off and their thermal comfort might have deteriorated to uncomfortable levels. This necessitates the use of multiple sensors to eliminate false or inaccurate detection and once again, the choice of sensors needs careful consideration in light of the low power and cost requirements. An important factor in determining occupancy is the detection algorithms that can combine the output of these sensors and provide occupancy accurately. The specific details of our sensors and the occupancy algorithms are provided in Section 4. We note that while our sensing is currently limited to occupancy, it can easily be expanded to include variables such as temperature, humidity, light levels, etc.

3.2 Data Collection Network

The data sent by the occupancy nodes then needs to be relayed to the data analysis and control component. The wireless network infrastructure to relay this data back from the sensors requires careful design. While multi-hop networks are often proposed in sensor network applications, we envision deploying dedicated base stations in a star topology that can relay data from the sensor radios to the rest of the infrastructure network. This design choice is based on the assumption that most buildings already have some form of wired (Ethernet) or wireless (WiFi) IP-based connectivity that can be leveraged to form the back haul for these base stations. Using a carefully placed set of base stations also allows us to ensure good radio coverage to all occupancy nodes. Furthermore this decision makes the network more manageable and secure since the occupancy sensors can only communicate with the base stations, which are themselves IP addressable, but cannot be reached directly.

3.3 Data Analysis and Control

Once the occupancy data has been received by the base stations from the individual nodes, it needs to be analyzed. This data analysis can be performed on a separate PC or a server and can be as simple as actuating a particular HVAC zone purely based on occupancy in that zone. On the other hand the data analysis can also be more complex and intelligent requiring long term trending and sensor data storage. If the occupancy information for users or even individual zones can be stored and analyzed over longer periods, it is possible to recognize patterns that allow for a more predictive and pro-active HVAC management, such as pre-cooling a zone based on observed trends[16].

Based on the data analysis the final step is to control the actual HVAC system. Modern buildings typically use central HVAC systems which contain one or more large air-handlers to push chilled air through the air-conditioning ductwork[7]. Each zone in a building is directly controlled by a variable air volume (VAV) box, which controls the flow of the chilled air being provided by the air-handlers. This entire system is controlled and monitored by an Energy Management System (EMS) that can set the command for each zone (on, off, stand-by). To exercise control on individual thermal zones, our control engine needs to communicate with the EMS and specify the commands for each zone. Normally EMS run on fixed occupancy schedules. Our control architecture needs to override these fixed occupancy schedules and instruct the EMS what each zone's command should be.

4. IMPLEMENTATION

The actual implementation of our HVAC control system follows the overall architecture outlined in the previous section. We first describe the building specifics and its HVAC infrastructure. Next we describe the specific aspects of our deployment including our occupancy sensors, the wireless network for data collection, and the back-end infrastructure. Finally we discuss our deployment experiences.

4.1 Building Testbed

Our system is deployed in the four-story CSE building on the UCSD campus. It is a six-year old building, and is among the more energy efficient buildings. Our HVAC system is connected to the campus chilled and hot water loops which are used by four large air handlers in the basement to provide cool air for the building. The air is circulated through a network of air ducts at 55°F. There are more than 300 thermal zones (each of which includes one to three rooms) that can be independently controlled. Each thermal zone is served by a VAV unit which adjust dampers to control the amount of air flows in its zone and reheats the air as needed. There are three basic command modes - `on`, where chilled air is released to satisfy cooling requirements; `stand-by`, where a minimum amount of air flow is maintained, and `off`, where the dampers are completely closed. The required air flow that the air-handlers need to provide is controlled by a PID controller using a pressure sensor in the main trunk of the air duct as its feedback element. As additional VAV units release more cold air into the building, the pressure drops in the air ducts forcing the air handlers to throttle up. Conversely, as VAV dampers close, the air pressure builds up, signaling the air handlers to slow down.

The building EMS is handled by the campus facilities group. Each building on campus is controlled through a BACNet control network which provides access to each thermal zone. This control network is centrally managed by a Metasys ADX server which provides access to all of the elements in the network. Each zone provides current temperature through a connected thermostat and can be set with various control parameters, such as cooling setpoint temperature and the operation command. The operation command, as mentioned earlier, is on, off, or stand-by.

Normally, the building HVAC is scheduled statically with occupied mode set starting at 5:15AM for the fourth floor, 5:30AM for the third floor, 5:45AM for the second, and 6:00AM for the first floor. The start times for the floors are staggered to avoid causing too much stress for the air handler units. Occupied mode lasts until 6:30PM, with stand-by mode set from 6:30PM to 10:00PM, and unoccupied mode set from 10:00PM to 5:15AM. When the building is in unoccupied mode, occupants can still turn on the HVAC by pressing a button on the thermostat; however, if an occupant does not have access to the room that has the thermostat for that zone, they will not be able to turn it on (since a zone can consist of more than one room).

The HVAC system is set to unoccupied mode during weekends. Discussions with the facilities group have revealed several interesting facts that are important in understanding how building operations can be improved to use less energy. The static schedule was set to best accommodate the diverse work times of university employees, some of who come in early, since real-time occupancy information is not available. The facilities are often overcooled since that generates, in general, less complaints than a hot building in our temperate climate zone. This type of schedule is common among modern buildings. The system is optimized for cooling energy use since the weather in our location is mild and warm year-round, and the heating component of the HVAC network is only utilized during the colder days. Additionally, the air that flows into each room is closer to the 55°F chilled air, as the VAV will minimally reheat the air to save energy. This means that many rooms are actually cooler than their cooling setpoint.

For our deployment, we were allowed to control the 2nd floor of the CSE building. This floor contains a total of 81 rooms and has a mixture of faculty, graduate students, and staff workers. There are 11 graduate student labs housing over 100 graduate students, which become occupied when the first graduate student comes in (usually between 8:30AM to 10:30AM). There are 23 faculty offices, of which four are currently unoccupied and eight are occupied only occasionally. The occupancy for the remaining offices varies greatly depending on the schedule of the professor or researcher. There are 11 affiliate offices that are occasionally occupied when affiliates come in or when there is an event in the building and 17 staff offices of which 5 are currently empty. Staff workers start entering between 8:30AM to 9:30AM and leave between 4:30PM to 6PM. The rest of the rooms include 1 classroom, 5 hardware labs that are occupied as needed, 2 computer labs, 4 conference rooms, 3 kitchenettes, and several equipment rooms and storage that do not contain occupants. The cooling setpoint for the rooms are statically set and vary from 71°F to 78°F (for storage rooms). Most rooms are set to 72°F. The heating setpoint varies from 66°F to 68°F, depending on the room. Each of these 81 rooms falls under one of 55 thermal zones on this floor. The hallways are divided up into two zones, each of the labs and conference rooms is its own zone, several of the office rooms are also their own zone, and the remaining zones represent two to three office rooms.

4.2 Wireless Occupancy-Node Network

We have implemented our occupancy-driven HVAC control architecture around three components: our own custom designed wireless occupancy nodes; several wireless base stations deployed to collect the data sent by the occupancy nodes; and a back-end server to process and analyze the occupancy data sent by the base stations and control the HVAC system.

Both the occupancy nodes and the base stations are based around a custom wireless sensor module that we have implemented. Our custom RF module is based on the TI CC2530/CC2531 reference design modified to incorporate our sensors. The CC2530/CC2531 is an SoC from TI that incorporates a 802.15.4 RF module and a 8051 micro controller core in a compact package. The CC2530 solution also includes a complete Zigbee compliant stack. For cost reasons, and to keep power consumption low, we based our design on a printed PCB antenna rather than as a separate external antenna. While this choice limits the radio range of our nodes, our experience with deploying them showed that the range was sufficiently large that we only needed a handful of base stations to cover an entire floor. A smaller range also has the benefit of spatial frequency reuse across the various base stations to reduce interference. The total cost of our custom wireless board was less than USD $10 in quantities of a thousand including parts and assembly.

Based on the range tests in our initial position paper[1] (which contained a single base station and a limited deployment of eight rooms) we first deployed four base stations in the hallways based on the radio ranges. However, we discovered dead spots and that some nodes were just outside the range of these base stations. In the end we ended up deploying a total of nine base stations for complete coverage. We implemented the base stations using inexpensive $100 Linux based plug computers called the GuruPlug and SheevaPlug, both of which have a 1.2Ghz ARM class processor, 512MB of memory, flash storage and Ethernet and USB ports. We chose the plug computers for their low cost, small form factor, low power (typically less than 5W) and the availability of several expansion ports. As mentioned earlier each base station is connected to our custom Zigbee based wireless module over USB, which the end occupancy nodes can then associate to (Figure 4).

We chose this single-hop star architecture rather than a multi-hop mesh network with relay/router nodes for several reasons. First, we plan to use the same Zigbee network for multiple smart building projects in which data rate is important. Multi-hop networks require more traffic for each data packet. Furthermore, for ease of management, dealing with a few dedicated base stations is easier than keeping track of many relay nodes in a mesh-network. Finally, since we already have building-wide Ethernet and WiFi, connecting these base stations to the network infrastructure was relatively easy. In buildings without such existing infrastructure, a multi-hop mesh network may be the only choice.

To further ease deployment, we configured our nodes to automatically connect to the nearest base station with the correct extended PAN ID and to start sending data to its parent. Using the Zigbee stack allows us to leverage many of the features of that stack, such as authentication and AES encryption for security. Once the wireless nodes are connected, the nodes will send event messages whenever an occupancy event happens. In addition, the nodes transmit a heartbeat status message every 15 minutes so that the base station and central server can determine if a node has fallen off the network. This heartbeat period can be made even longer to conserve energy and further increase node lifetime.

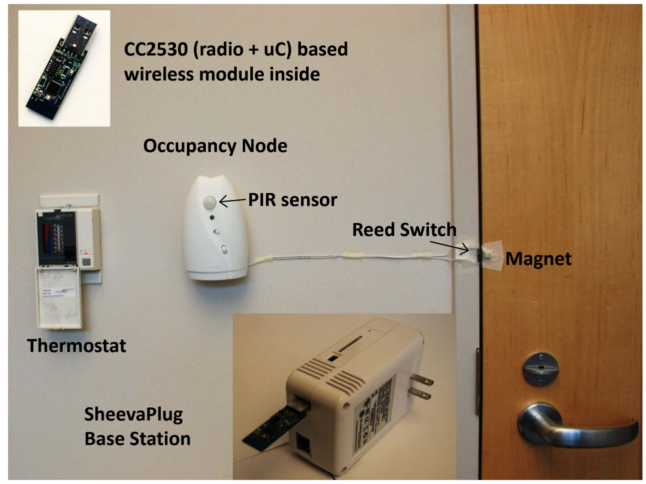

Figure 4: Occupancy node deployed on the wall of an office. The reed switch, PIR sensor and our CC2530 based radio module inside the occupancy node are also shown.

4.3 Occupancy Sensor Modules

Our wireless occupancy node builds upon our previous occupancy node design with increased accuracy, improved robustness and lower power consumption[1]. As mentioned earlier, the core of our occupancy sensor is our CC2530-based custom wireless module. This module is connected to an interposer sensor board that includes a PIR sensor to detect movement and a magnetic reed switch to detect door open/close events. Both these sensors are connected to an interrupt enabled GPIO pin of our wireless module. To deploy widely within all the offices on our floor, we needed to package the occupancy sensor in a clean enclosure. We identified a commercial product, the AirWick, that has a PIR sensor which dispenses room freshener on movement and costs less than $4. We dismantled the AirWick case to remove all components and salvaged its PIR sensor for our own board since we found it to be much more reliable in detecting movement than several more expensive PIR sensors we tested. Our installation as shown in Figure 4 consists of a magnetic reed switch along the door frame. A small magnet is fitted on the actual door so that when the door is closed, the magnet and reed switch become adjacent.

Based on the typical occupancy modalities for our building and other buildings around our campus, we observed that almost everyone closes their office door when they are either leaving for the day or when they are going to be out for more than a few minutes. Our occupancy detection therefore works as follows. The magnetic reed switch detects when the door to a particular office is open or closed. When a door-open event occurs, we mark the particular room as occupied. This choice may lead to some lost opportunities in terms of energy savings if a person leaves their door open even when they are not there. A door-close event, however, may imply either the occupant is leaving (room unoccupied) or is merely closing the door to remain in the room (room still occupied). In order to differentiate between these two possibilities, we leverage the PIR sensor to detect whether or not there is movement immediately following the door-close event denoting that the person is still in the room. Another situation arises when someone leaves a room and closes the door behind them, leaving another occupant still in the room. In such a scenario, our occupancy node may register the room as unoccupied if the person does not move, even though it is still occupied. To account for this, we enable the PIR interrupt so that any future movement will trigger a brief polling mode to determine if an occupant is still in the room. Thus, there are three types of event messages that the node can send (in addition to the status messages): `open-door occupied`, `closed-door occupied`, and `closed-door unoccupied`.

An important question is whether a PIR only solution would suffice for occupancy detection. Our previous work[1] compared occupancy sensors and showed that the accuracy of the PIR-only sensor was worse than the Reed-switch + PIR solution. A PIR only solution had both false positives (leads to wasted energy) and false negatives (leads to occupant discomfort). As we report in our evaluations (Section 5), our occupancy sensor is accurate 96% of the time and more importantly leads to very few false negatives. Of the 81 rooms, only 64 are ever occupied, and we were able to achieve a deployment of 47 rooms with our occupancy nodes.

An important design goal for our wireless occupancy node was to make it battery powered, necessitating aggressive energy management. Our choice of using the CC2530 was in part because of its low power consumption. However, since our occupancy node combines several sensors with the CC2530 we wanted to measure accurately its power draw in different modes and estimate total battery lifetime. We use a high sample rate Data Acquisition card from National Instruments USB-6210 to measure the total current draw across a sense resistor, and consequently calculate the power draw. The maximum current is during data transmission and is 30mA. The CC2530 supports multiple sleep states, the lowest of which consumes less than 0.045mA. All current measurements are at 3.6V. The average daily energy consumption depends on various factors that determine which power state the occupancy node is in. Assuming a periodic heartbeat of once every 15 minutes and over 100 occupancy events transmitted per day, we calculate the total energy drain to be 3.37mAh. Assuming standard alkaline batteries with 2850mAh, this current draw translates to a lifetime of over 2 years. The low power draw of our occupancy sensor makes it perhaps possible to employ energy harvesting using indoor solar cells, and use super-capacitor based designs to makes these nodes almost perpetually powered[4, 17].

4.4 Back-end Infrastructure

The base stations send the wireless sensor data along with the status messages to the central server, which we call the Occupancy Data Analysis server (ODAS). The core component of the server software is the database that stores the information and a collection of Python programs that read from the database and perform actuations based on the data.

The ODAS is connected to a Windows server machine that runs an OPC tunneler. OPC is a common standard that allows for process control and communication between industrial devices. This machine is connected directly to two OPC Data Access servers managed by the facilities group on campus, with one providing real time energy usage data in our building, and the other providing access to the BACNet network. We developed OPC client applications to interface with the OPC servers and obtain the data points that we are interested in, which include temperature for each thermal zone and energy usage of the HVAC system. In addition, we were given write access privileges for setting the occupancy command for each zone in the building.

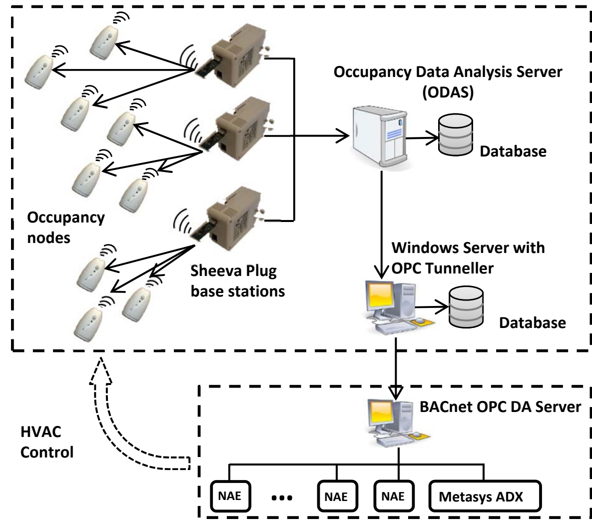

Figure 5: Implementation Diagram.

Our ODAS runs a process that retrieves this data (zone temperatures) from the Windows server and stores it in the ODAS database. It is important to note that the temperature readings can be delayed up to 10 minutes. To control the building HVAC, the ODAS sends a zone HVAC command (e.g. turn zone 2121 to unoccupied) back to the Windows server. The *actuator* OPC client application will scan incoming commands at a rate of once a second to see if a new one has arrived. If so, the actuator client will write the appropriate value for the OPC item to the BACNet OPC server. The BACNet OPC server has a higher priority than the static schedules, which allows it to override any previous command. (Higher priorities do exist, such as for emergencies). This set up provides us with an an ideal test bed to experiment with different HVAC control algorithms revolving around using occupancy information along with other information sources. Figure 5 shows our overall system design with the individual components marked.

4.5 HVAC Control

The ODAS database comprises of several tables for our HVAC control including ones that contain all of the rooms on each floor, the thermal zones, mappings between rooms and thermal zones, and temperature of each zone. Using this information, our HVAC control algorithms improve upon the static schedules set by the facilities management. Due to our location in a mild climate zone, typical days generally do not exceed a high of 76°F (except in a few weeks during the summer). Temperatures do not get very low either, with lows of the mid 50s. When we looked at the indoor temperatures for each zone during a normal warm weekend, we noticed that even with HVAC off, most of the rooms would not go above 75°F. The exception was one set of sparsely occupied offices

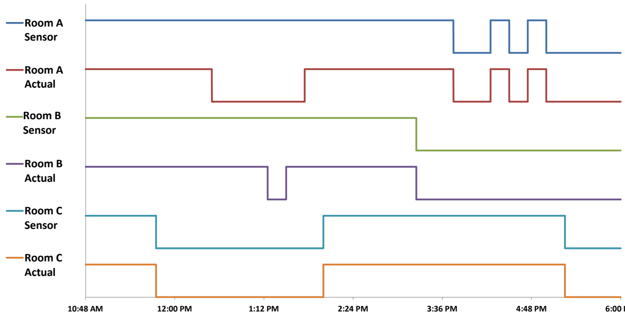

Figure 6: Accuracy test of three representative rooms over seven hours.

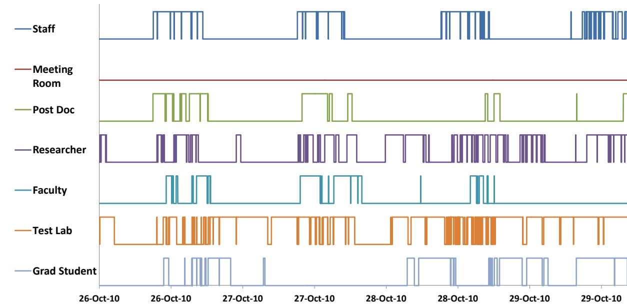

Figure 7: Occupancy for a representative set of seven occupants across four days. The data shows significant diversity in their occupancy patterns.

on the wing facing the sun containing some IT equipment. These offices would climb up to 77°F, even on mild days, due to the solar effects and the heat-generating computers.

Our HVAC control system implementation comprises of several programs. The first program checks the occupancy of every zone and will turn-off (or put into stand-by) zones that are currently unoccupied. We rate-limit this to once a minute in order to prevent thrashing for the dampers. Damper power consumption is, however, quite low measuring around 100W over 20 seconds. In addition, we have a program running that will check the temperature for every occupied zone and will turn on HVAC if the temperature goes over 76°F or under 66°F.

4.6 Deployment Experiences

As reported by other efforts, difficulties and challenges arise in any large wireless sensor network deployment[11]. Socially, invasion of privacy is a big concern, though it is beyond the scope of this paper. On issues related to building performance, in our 20 informal interviews with the occupants, the most frequent complaint with existing environmental conditions was excessive cooling, which is a side effect of a system designed to provide frequent cooling rather than heating actuations.

As mentioned earlier, we were able to cover most of the entire second floor of our building with our occupancy sensors, except for a few rooms whose occupants were uncomfortable with a sensor deployed in their office. We also did not cover several storage rooms that do not house any people. In addition, several rooms happened to be unoccupied during the time period in which we were able to run our system.

5. EVALUATION

We now evaluate our test deployment for accuracy of occupancy detection, show how occupancy patterns vary across people, and demonstrate the potential energy savings for running our dynamic HVAC control scheme.

5.1 Occupancy - Accuracy and Patterns

In our previous work[1], we compared the accuracy of a PIR-only node versus our initial occupancy node, and the improvements we have made on our node since then have led to an even higher accuracy. Over the course of a weekday we compared the ground truth (obtained by going around the floor every 15 minutes and verifying actual occupancy) vs. what our occupancy nodes were registering. Figure 6 displays the comparison between ground truth and sensor results over a day for three representative nodes. These results show that our occupancy nodes are indeed very accurate. Of the 33 nodes that we tested, 29 were 96% accurate over the day, with four having lower accuracy. A closer analysis of the erroneous sensors revealed that detection inaccuracies were due to false detection since we placed the sensors in a poor position (where gusts of air would activate them). When we fixed that by increasing the distance between door and sensor, this error reduced considerably. Note that the rooms were never marked unoccupied when a person was present in the room (false negatives). Therefore the subject never experiences any discomfort because of our occupancy detection system. Recall, the only downside to false positives is extra energy consumed by the HVAC system.

Our occupancy data reveals several interesting trends. It is important to note that as this is a building on a university campus, occupancy patterns are extremely dynamic when compared to a typical 9-5 office building. Figure 7 shows occupancy patterns of seven representative rooms over four days. We can observe that the staff worker has a fixed schedule from 8:30AM to 4:30PM everyday. The ad hoc meeting room (that is often times not used) was empty during these 4 days. The faculty and post doc have sporadic occupancy patterns, mainly because of different commitments outside of their offices. The gaps in the occupancy of the graduate student perhaps indicate that he/she was attending classes.

5.2 Energy Savings

We ran our experiment in two sets, once in late October 2010, and the other in late February 2011, giving us data over two seasons (fall and spring). Our first set occurred over Oct. 27 and Oct. 28 (Wednesday and Thursday) 2010. For this set, we only had a partial deployment with our our occupancy sensors in about half of the rooms on one floor. For the zones that did not have sensors, we scheduled the HVAC such that they would activate at a time closer to when the occupants arrived (instead of 5:45AM). This was usually between 8:45AM to 10:00AM depending on the room. We were very aggressive in this set with reducing energy consumption. For the second set of tests, which we ran on Feb. 23 and Feb. 24 (Wednesday and Thursday), we were able to achieve 75% deployment, with the only exceptions being people who were uncomfortable with the sensors and supply rooms that were never extensively occupied. In both cases, we set all zones to stand-by at 6:30PM and turned off all

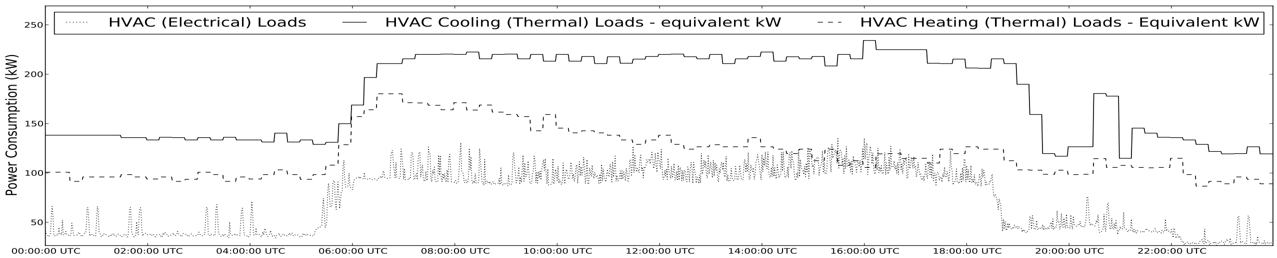

Figure 8: The energy consumption of HVAC during our baseline day. We show HVAC electrical loads as well as the HVAC thermal loads for both cooling and heating (as equivalent kW).

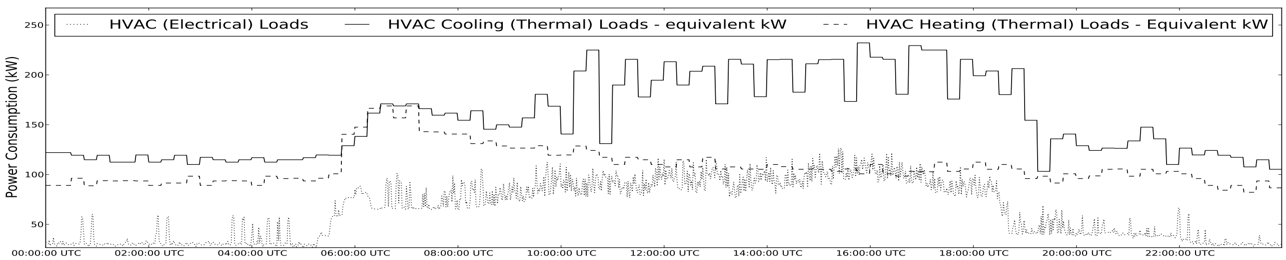

Figure 9: The energy consumption of HVAC during our first test day. We graph both electrical loads and thermal loads (as equivalent kW). The HVAC-electrical savings compared to baseline shown in Figure 8 are 11.59% while the HVAC-thermal savings are 12.41% and 9.59% for cooling and heating loads respectively.

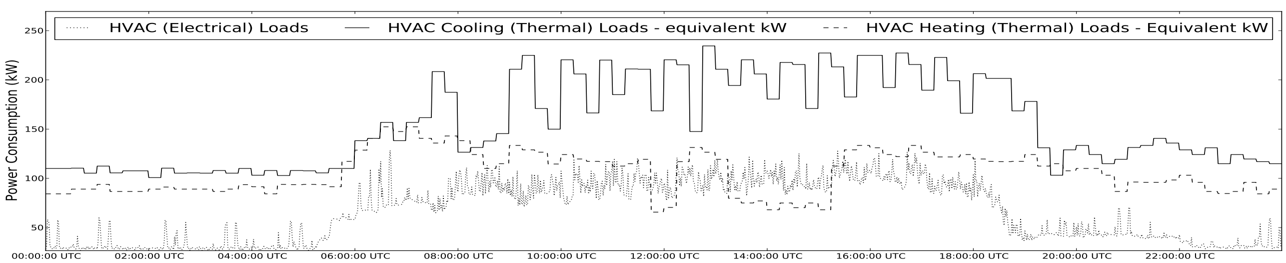

Figure 10: The energy consumption of HVAC during our second test day. We graph both electrical loads and thermal loads (as equivalent kW). The HVAC-electrical savings compared to baseline shown in Figure 8 are 9.54% while the HVAC-thermal savings are 12.85% and 11.51% for cooling and heating loads respectively.

HVAC at 10PM, identical to the normal HVAC schedule. Measuring energy savings is difficult because every day has its own environmental parameters that affect how much energy must be spent to cool or warm a building, so we found baseline days that were very close in time and in terms of weather to our test dates to compare against.

Running the experiments over these two seasons also offered us the opportunity to determine how our system is effected by different outside temperatures. The fall season in San Diego is fairly warm, while the winter/spring season is fairly cool. In the former, the main mode of operation is in cooling the building, while in the latter it is in both warming and cooling the building. In both cases however, the mild temperature and modern building enclosure meant that even without HVAC, temperatures never got too extreme.

5.2.1 *Fall 2010 Experiments*

Because the offices often had moderate indoor temperatures, we were extremely aggressive in reducing HVAC loads for our first test day. Of the rooms that we were unable to deploy our nodes, we either set the HVAC to on or off the entire day, depending on whether the rooms were currently vacant (or vacant for the day). For rooms that were occupied, we set their occupied commands at a time closer to when the occupants arrived (instead of 5:45AM), usually between 8:45AM to 10:00AM depending on the room.

For the rooms that had our occupancy sensors installed and were their own zone, we simply cycled the occupancy command based on actual room occupancy. However, larger zones that contained multiple rooms were more challenging to control, and here we opted to save energy. If the other occupants of the zone stated that they felt the cooling was too high, we simply put the HVAC to stand-by when one occupant was gone. Combined with the fact that cooling from other zones would seep in anyway, and the fact that the days rarely got hot, we believed that it would be enough to maintain comfort. For unoccupied rooms, we simply turned off the HVAC. After we ran our experiment on our first day, we changed our control procedures to opt for a more conservative approach on the second day.

The first test day was typical for the location, with mild temperatures and a high of 75°F. The second day was warmer, hitting a high of 82°F. Comparing energy consumption for HVAC across multiple days is difficult, as the exact weather patterns are difficult to reproduce, and HVAC loads are directly impacted by temperature and solar radiance. However, aside from this difference, a reasonable comparison across multiple days can still be made. Our test days were on October 27th (Wednesday) and October 28th (Thursday) of a typical work week, and we set as our baseline October 25th which falls on the Monday of that same week. This day has a fairly representative HVAC energy pattern for a mild day in our location, with a high of 73°F. We note that this baseline day was much cooler than our two test days, therefore our energy savings are somewhat conservative and would have been even higher had we compared with a similar warmer day. Figure 8 shows the HVAC energy trace of our baseline day, including both electrical consumption of the air handlers as well as the thermal loads for the building (given as equivalent kW). As mentioned earlier, equivalent kW is merely an approximation, and should not be taken as an exact conversion.

5.2.2 Fall 2010 – Day 1 Results

Figure 9 shows the HVAC power consumption traces for test day 1 (Oct 27). The graphs show the energy consumed by the entire building's HVAC system, not just the second floor (which we controlled). A normal static schedule day will start up the four floors in sequence starting at 5:15AM (for the fourth floor) until 6:00AM (for the first floor). A close look at the comparison day (Figure 8) from 5:15AM - 6:00AM shows the spikes that each floor causes. The effect on energy consumption is apparent, as the average power consumption hits past 100 kW. The rapid succession of open dampers causes the air handler to have to ramp up its fan speed, and this causes an even greater energy load.

For our test day 1 control scheme, we actually started our energy control scheme for floor 2 at 6:05AM since we wanted the normal initialization procedures to start up first. At 6:05AM our control commenced. Because it was before 6:30AM (the earliest time that we set for HVAC initialization), the system immediately set all the second floor zones to unoccupied. The energy savings this had were surprising. Rather that seeing energy consumed go up to an average of over 100kW, the energy consumed only went up to 80 kW, and settled down at an average 68kW for the morning. At 6:30AM our system went to "on" mode, which meant that it will only turn on the HVAC when the occupant arrived (for rooms that have the occupancy sensor), or at a later time that we set statically (typically 8:30AM to 9:30AM, based on when we observed the room usually becoming occupied). Starting at 8:00AM more rooms became occupied, and at 8:30AM the rest of the rooms (ones that we did not sense) were turned on.

The energy consumption rose as the day went on, but was still 5kW to 10kW less than the comparison day until the mid-afternoon (around 3PM). The effects from duty cycling (setting rooms to unoccupied mode when they were absent) had some effect, but the occupancy patterns for many rooms were quite static (long periods of occupancy). Even when multiple rooms were part of the same zone, our control scheme would opt to set the entire zone to unoccupied until it hit 75°F when we would turn the HVAC on. Many of these staff rooms however were not facing the sun in the afternoon, and therefore actually never floated that high.

Day	Electricity	Cooling	Heating
Baseline	1760 kW-H	4302 kW-H	2877 kW-H
Day 1	1556 kW-H	3768 kW-H	2601 kW-H
Day 1 Savings	11.59%	12.41%	9.59%
Day 2	1592 kW-H	3749 kW-H	2546 kW-H
Day 2 Savings	9.54%	12.85%	11.51%

Table 1: Fall Tests - Energy consumption for electricity and thermal cooling and heating (as equivalent kW).

The late afternoon also showed a period of significant savings. Staff workers tended to leave from 4:30PM to 6PM. The normal schedule system puts all the floors into standby mode at 6:30PM, whereas our control scheme started labeling rooms as unoccupied when they left. The effects are significant, as due to our control the HVAC loads started dropping towards 80 kW at 5PM. In comparison, the static schedule averaged over 90 kW until 6:30PM.

The total HVAC electrical load for test day 1 was 1556 kW-H. The total HVAC electrical load for the baseline day was 1760 kW-H. Therefore, in terms of electricity, our HVAC control scheme saved a significant 11.59%, despite only controlling one floor in a four floor building. We also note that the thermal load consumption was less than the baseline as well, saving 12.41% in thermal cooling loads and 9.59% in thermal heating loads (results summarized in Table 1).

5.2.3 Fall 2010 – Day 2 Results

The first day we were very aggressive in cutting off HVAC cooling to as many rooms as possible. However, given that we did not actually detect occupancy in half of the rooms, setting them as completely unoccupied could have potentially unintended consequences in terms of higher temperatures if an occupant did happen to come into that room. Therefore, for day 2, we ran a less aggressive cooling control pattern where we would actively monitor the temperature and turn on cooling whenever the temperature would rise past 75°F, regardless of whether or not we had a sensor node in the room or not.

Another change was to start the day off with full control, as opposed to letting the normal static procedure initiate. This meant that the entire 2nd floor would not be set to occupied at 5:45AM, and instead be turned on at 6:30AM. The effects of this were immediate, as the average power consumption hovered near 80 kW for most of the morning. As the second floor started to become occupied, the power rose to an average of 100 kW, not entirely dissimilar to our baseline day. However, looking closely, we observe that the average power consumption of test day 2 was still slightly lower than the baseline day. Similar to test day 1, energy consumption started falling rapidly as the work day ended.

The total energy consumed for test day 2 was 1591.68 kW-H. Test day 2 was also much warmer, resulting in higher than normal HVAC loads, but our conservative approach likely added additional energy consumption over test day 1. Compared to our baseline day, test day 2 saved 10.5% in electricity. We observed 12.85% savings in thermal cooling loads, with the savings mostly concentrated in the morning, and 11.59% in heating loads. Table 1 shows the results for test day 1 and test day 2 compared to our baseline day.

5.2.4 *Additional Observations for Fall 2010 Tests*

Looking at the fall results, it is clear that a significant source of the energy savings comes from starting the HVAC when the users arrive rather than as early as 5:45AM. A side benefit of this is that the load on the air handlers is more staggered. Since our building is located in a mild climate zone, it is not necessary to aggressively pre-cool, however in other climates, we would opt for a learning algorithm to predict when users arrive and initiate cooling accordingly. Since occupancy patterns for a given individual tend to be similar, this is likely to be an effective strategy.

We also note that occupancy patterns tend to disallow a great deal of online duty cycling, as people tend to be in their offices for long periods of times. This is especially true for the graduate laboratories, which always had some occupants throughout the day. The effects on air flow when one zone is being cooled and another is not is also significant and these thermal effects may have a non-trivial impact on overall energy consumption.

It is important to factor in the effect on temperatures through the zones. One interesting observation was that the location of the room significantly affects how warm it will be. As mentioned previously, one side of the wing of our building faces the sun during the afternoon, and thus gets much warmer than the other rooms. We noticed that the majority of the other rooms would stay constantly under 75°F even with the absence of cooling, but these rooms would rise to 77°F during the afternoons. We ran an experiment to control the HVAC in one of these unoccupied room with several computers. Over the course of a warm day, we enabled and disabled HVAC to see how temperature rises and falls, and importantly, how quickly the temperature adjusts. The temperature readings are read from the BACnet OPC DA server and thus exhibit some amount of delay and discrete jumps.

Figure 11 shows how this particular room reacts to turning off and on HVAC over the day. The HVAC system was able to send a significant amount of cold air into the room to rapidly cool it to acceptable temperatures. We note that it takes less than ten minutes for the HVAC system to cool the room below 75°F. This suggests that allowing temperatures to float will have minimal impact on comfort levels; the effect of this scheme on energy consumption however still needs to be studied further. We do note that the cold air in the ducts is at 55°F normally, and is reheated to maintain temperatures in each zone. When a zone is warm, the reheat required is reduced. This suggests that cooling a warmer room might not significantly impact HVAC energy consumption more than maintaining a room at its setpoint.

5.2.5 *Spring 2011 Experiments*

We ran our HVAC control experiment again in spring 2011, and were able to achieve a much larger deployment of occupancy nodes. We also adjusted our test parameters based on our discoveries from the first set. Since we had a much better deployment, we were able to monitor most of the rooms, including some unoccupied rooms, and actuate based on that. We adjusted the timing and opted for a later start time of 8:45AM to turn on the common areas and unmonitored offices, as this was when most people came in. Rather than turn off HVAC as in our first set of tests, we opted to put unoccupied zones to stand-by instead to maintain some airflow. We also were more conservative

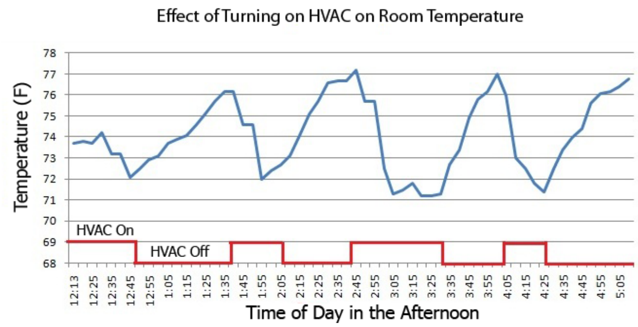

Figure 11: Effect of actuating HVAC for an IT-heavy room facing towards the sun during a warm day.

with placing zones to unoccupied for this set of tests as we had more monitors - if a single room in a zone was occupied, we turned on the HVAC.

For this test, our baseline day was February 17th (Thursday) of a typical work week. It was mildly cloudy with temperatures ranging from 53°F to 62°F. Our test days were Feb. 23 and 24, days with similar weather to the baseline day. The temperature on Feb. 23 varied between 54°F to 63°F, with the temperatures not varying by more than 3°F between the two days. The weather on Feb. 24 was generally cloudy, with temperature varying from 51 deg F to 60 deg F. The weather was much colder than the fall tests, meaning we would be able to test how our system handles conditions where it must sometimes warm rather than cool. From observing the data over the previous weekends, we noticed that the temperatures in the building would typically range from 65°F to 73°F when the HVAC system was completely off.

5.2.6 *Spring 2011 Results*

Due to lack of space, we were unable to include the spring 2011 graphs, but the summary of our spring results is listed in Table 2. Comparing the HVAC electrical load on the test day with that on our baseline day, we see that the load increases more gradually during our experiment days. The load at 6AM on the test day is 70kW compared to 105kW. The load increases slowly to 105kW at about 8:30AM, while it remained constant during this period on the base day. The peak load on the test day was 165kW compared to 180kW on the baseline. The power consumption starts decreasing at 3:30PM compared to 4:30PM on the baseline. It remains at 120kW until 6PM compared to 140kW, mapping closer to the occupancy levels in the second floor. The total energy consumed by the HVAC electrical on the test day was 1977 kWh compared to 2187kWh on the base day. The energy saved during the period was 9.60%, with similar savings in thermal-cooling. Thermal heating energy consumption was almost the same however.

The second day had similar temperatures but was more overcast. The HVAC electrical consumption stabilized around 75W at 6AM, and stayed there until 8:15AM. The power consumption then started to increase gradually until it reached a peak of 145kW at 3:30PM. The total HVAC electrical power consumption for the day was 1843 kWh. We saved 15.73% compared to the baseline day. This day was cooler however, which meant that reducing the HVAC would save more in energy because of less heating.

Day	Electricity	Cooling	Heating
Baseline	2187 kW-H	3137 kW-H	2124 kW-H
Day 1	1977 kW-H	2885 kW-H	2128 kW-H
Day 1 Savings	9.60%	8.03%	-0.18%
Day 2	1843 kW-H	2899 kW-H	2021 kW-H
Day 2 Savings	15.73%	7.59%	4.85%

Table 2: Spring 2011 Tests - Energy consumption for electricity and thermal cooling and heating (as equivalent kW).

As our original system was designed for warm days, running it on a colder day proved enlightening. Temperatures in the 60s is about as cold as it gets in San Diego, and thus our building was forced to warm the rooms. We did notice though that rooms got as low as 66°F in the morning. Rooms where HVAC was turned on were warmer at about 70°F-72°F, while rooms that had HVAC turned off were cooler at about 67°F-69°F in the late mornings. By the afternoons, most of the rooms were above the heating setpoint and some were even being cooled. In retrospect, we perhaps needed to better optimize heating strategies, as pre-heating is perhaps more important than pre-cooling for these situations. It would also have been instructive to take into consideration warm air spreading from warmer zones to non-warm zones. We did not however receive any complaints about the temperatures; the people we asked did not even notice any change.

6. RELATED WORK

Prior work has looked at ways to access sensors and actuators within buildings in an easy and efficient way for various reasons including improving energy efficiency[5, 18, 20]. Researchers have even proposed pushing TCP/IP down to individual sensors to make them IP addressable and directly accessible to existing IT infrastructure and using web services[20] to access these sensors.

Using occupancy information as an input to drive energy savings in buildings has been a well studied topic both in the research community as well as in commercial products. The most common is to use motion sensors, usually based on PIR technology, and turn on and off lights depending on occupancy[6]. These PIR sensors are almost always hard-wired during the commissioning of new buildings and are connected directly to the appropriate lights but not to the HVAC systems. Recent commercial products have started to use wireless radios from the PIR sensors installed in wall switches to control ceiling lighting, but these radios are short range and more importantly do not control HVAC systems. A recent product released by Honeywell[10] uses a simple motion sensor to communicate wirelessly with a single HVAC unit (for a single room) to control it, but is not meant for large central HVAC systems.

Recognizing the limitations of PIR motion sensors, various researchers have proposed using other technologies like cameras and computer vision, CO2 sensors, sonar based methods or even radars [8, 9]. However, all these technologies have their own set of limitations including significant delays in detecting new people, higher costs, and higher power consumption. The use of networked cameras for occupancy detection at the coarse level of entire floors, while promising, run into privacy and cost issues.

Within the context of a home, recent work has used coarse grained cues providing occupancy information such as a person leaving his home or a person arriving back to drive a smart thermostat[16]. The authors however have focused more towards the algorithm aspects of the self programming thermostat instead of determining occupancy accurately. Similarly Barbato et al.[3] consider a home scenario for occupancy detection using PIR sensors while focusing on algorithms that can help extract user profiles. More recent work has also shown that within a limited context of a home, appropriately placed ultrasonic sensors can be used not only to detect occupancy but also to distinguish between individuals based on differences between their heights[19]. This capability can be useful to learn and set HVAC preferences for individuals in a home. In all of these cases, the techniques are geared more towards smaller home scenarios and not towards large buildings with many occupants.

Finally, there has been a significant amount of interest of late to identify and breakdown where the energy is consumed within large buildings[2, 15, 11, 12] and homes[14, 13]. Jiang et al.[11] have implemented and deployed individual plug level energy meters widely which can both measure and manage power use by switching off appliances. Jung et al. [12] look at combining the knowledge of the ON/OFF states of different appliances with a limited deployment of energy meters to estimate energy consumption but do not consider HVAC loads. Our own previous works[15, 2] have shown that IT loads and HVAC loads are both the most significant consumers of total energy, especially when considering both electrical and thermal loads of the HVAC system.

7. CONCLUSIONS AND FUTURE WORK

This paper describes our experiences with the design and deployment of a wireless sensor network for monitoring occupancy, and its use in the operation of HVAC resources in the CSE building at UCSD. By providing real-time occupancy information to the campus-level management network, we have implemented an HVAC control scheme that obtains 9.54% to 15.73% savings in HVAC electrical energy use and 7.59% to 12.85% savings in HVAC thermal energy use through controlling just a single floor of our four floor computer science building; controlling all four floors will lead to even more savings.

There are several applications of the sensing and data collection architecture described here and its implications on how commercial buildings can be operated. Accurate and real-time occupancy information is critical to emergency response especially with occupants with limited mobility or accessibility. Combined use of HVAC and IT resources, driven by occupancy information, presents another optimization opportunity that needs to be explored. Further, presence of in-building information storage and processing resources would allow for use of longitudinal weather data to devise strategies for learning to enable pre-cooling or other load-shifting strategies. We are actively working with the facilities management in order to obtain additional control parameters over the CSE building. We are also allowing users to adjust their preferences in real time so that they will be able to remotely adjust their HVAC preferences.

Because we are actually controlling HVAC in a building, we opted to use a rather naive and direct control algorithm to guard against any excessive changes to HVAC settings that could generate complaints from building occupants. Our future plans however are to use more sophisticated control methods. In fact, our results suggest that such algorithms, coupled with greater use of available data, can both drive very optimal HVAC schedules that conserve energy and increase comfort. We actively monitored all of the zones during the entire run of the experiment as well. It is also important to note that our experiments were on reducing energy consumption, and thus we have not fully examined the actual effects on comfort that pre-cooling or pre-heating provides. To the extent that this trade-off is beneficial requires more studies and more sophisticated techniques will be necessary going forward in optimizing HVAC control based on occupancy.

8. ACKNOWLEDGMENTS

We wish to thank Robert Austin, Bob Caldwell, John Dilliott and Walter Richardson at UCSD Physical Plant Services for their continued support to access the UCSD EMS. We also wish to acknowledge Sathya Kuppuswamy and Jawon Lee for their invaluable help in the sensor deployment. Finally we thank our anonymous reviewers and our shepherd, Dimitrios Lymberopoulos, for their comments and feedback to improve our paper. 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 CCF/SHF-1018632, CPS-0932360 and the Von Liebig Center CleanTech grant.

9. REFERENCES

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

[2] Y. Agarwal, T. Weng, and R. Gupta. The Energy Dashboard: Improving the Visibility of Energy Consumption at a Campus-Wide Scale. In *First ACM Workshop on Embedded Sensing Systems For Energy-Efficiency In Buildings*, 2009.

[3] A. Barbato, L. Borsani, A. Capone, and S. Melzi. Home energy saving through a user profiling system based on wireless sensors. *Conference On Embedded Networked Sensor Systems*, pages 49–54, 2009.

[4] C.-Y. Chen and P. Chou. DuraCap: A Supercapacitor-based, Power-bootstrapping, Maximum Power Point Tracking Energy-Harvesting System. In *Proceedings of the 16th ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED)*, pages 313–318. ACM, 2010.

[5] S. Dawson-Haggerty, J. Ortiz, X. Jiang, S. Shankar, and D. Culler. Enabling Green Building Applications. In *Sixth Workshop on Hot Topics in Embedded Networked Sensors (HotEmNets)*, 2010.

[6] D. T. Delaney, G. O'Hare, M. P. Gregory, and A. G. Ruzzelli. Evaluation of energy-efficiency in lighting systems using sensor networks. *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 61–66, 2009.

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

[8] P. Dutta, A. Arora, and S. Bibyk. Towards Radar-Enabled Sensor Networks. In *Proceedings of the 5th international conference on Information processing in sensor networks*, pages 467–474. ACM, 2006.

[9] 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. *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 19–24, 2009.

[10] Honeywell:. WSK-24 Wireless Occupancy Solution. `http://specifyhoneywell.com/customer/techlit/pdf/63-0000s/63-4519.pdf`.

[11] X. Jiang, M. V. Ly, J. Taneja, P. Dutta, and D. Culler. Experiences with a High-Fidelity Wireless Building Energy Auditing Network. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, pages 113–126. ACM, 2009.

[12] D. Jung and A. Savvides. Estimating Building Consumption Breakdowns using ON/OFF State Sensing and Incremental Sub-Meter Deployment. In *Proceedings of the 8th ACM conference on Embedded Network Sensor Systems (SenSys)*, 2010.

[13] 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 *Proceedings of the 11th International Conference on Ubiquitous Computing*, pages 245–254. ACM, 2009.

[14] Y. Kim, T. Schmid, Y. Wang, and M. B. Srivastava. Challenges in Resource Monitoring for Residential Spaces. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 1–6. ACM, 2009.

[15] J. Kleissl and Y. Agarwal. Cyber-physical Energy systems: Focus on Smart Buildings. In *Design Automation Conference (DAC), 2010 47th ACM/IEEE*, pages 749–754. IEEE, 2010.

[16] J. Lu, T. Sookoor, V. Srinivasan, G. Ge, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The Smart Thermostat: Using Wireless Sensors to Save Energy in Homes. In *Proceedings of the 8th ACM conference on Embedded Network Sensor Systems (SenSys)*, 2010.

[17] V. Raghunathan, A. Kansal, J. Hsu, J. Friedman, and M. Srivastava. Design Considerations for Solar Energy Harvesting Wireless Embedded Systems. In *Information Processing in Sensor Networks (IPSN '05)*, pages 457–462. IEEE, 2005.

[18] A. Rowe, M. Berges, G. Bhatia, E. Goldman, R. Rajkumar, L. Soibelman, J. Garrett, and J. M. F. Moura. Sensor Andrew: Large-Scale Campus-Wide Sensing and Actuation, CMU Tech Report, ECE-TR-08-11, 2008.

[19] V. Srinivasan, J. Stankovic, and K. Whitehouse. Using Height Sensors for Biometric Identification in Multi-resident Homes. *Pervasive Computing*, pages 337–354, 2010.

[20] D. Yazar and A. Dunkels. Efficient application integration in IP-based sensor networks. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 43–48. ACM, 2009.