Sentinel: Occupancy Based HVAC Actuation using Existing WiFi Infrastructure within Commercial Buildings

Bharathan Balaji†, Jian Xu†, Anthony Nwokafor†, Rajesh Gupta†, Yuvraj Agarwal†‡

†University of California, San Diego ‡Carnegie Mellon University

†{bbalaji, jix024, aanwokafor, gupta}@cs.ucsd.edu, ‡yuvraj.agarwal@cs.cmu.edu

ABSTRACT

Commercial buildings contribute to 19% of the primary energy consumption in the US, with HVAC systems accounting for 39.6% of this usage. To reduce HVAC energy use, prior studies have proposed using wireless occupancy sensors or even cameras for occupancy based actuation showing energy savings of up to 42%. However, most of these solutions require these sensors and the associated network to be designed, deployed, tested and maintained within existing buildings which is significantly costly.

We present Sentinel, a system that leverages existing WiFi infrastructure in commercial buildings along with smartphones with WiFi connectivity carried by building occupants to provide fine-grained occupancy based HVAC actuation. We have implemented Sentinel on top of RESTful web services, and demonstrate that it is scalable and compatible with legacy building management. We show that Sentinel accurately determines the occupancy in office spaces 86% of the time, with 6.2% false negative errors. We highlight the reasons for the inaccuracies, mostly attributed to aggressive power management by smartphones. Finally, we actuate 23% of the HVAC zones within a commercial building using Sentinel for one day and measure HVAC electrical energy savings of 17.8%.

Categories and Subject Descriptors

[**Computer systems organization**]: Special purpose systems, sensors and actuators, real-time system architecture

General Terms

Sensing, Measurement, Control

Keywords

Sentinel, HVAC, occupancy, buildings, energy efficiency

1. INTRODUCTION

Commercial buildings contribute to 19% of the primary energy consumption within US, with Heating, Ventilation and Air Conditioning systems (HVAC) accounting for 39.6% of this usage [48]. As a result, improving the energy-efficiency of building HVAC systems is key from both a cost saving and sustainability standpoint. Prior research has shown that most modern buildings use static schedules to run HVAC systems, thereby wasting considerable energy in conditioning unoccupied spaces [8, 25, 28, 29]. Furthermore, while modern building HVAC systems use Variable Air Volume(VAV) control, which allows independent control of thermal zones [34], it is not leveraged effectively by facility managers in practice due to the absence of accurate occupancy information within physical spaces.

Using occupancy information for HVAC control has in fact been studied extensively [8, 27, 28, 29, 30, 31]. While CO_2 sensors are used to detect occupant density in large spaces [2, 5], the detection times for changes in concentration of CO_2 with occupancy were found to be too slow for use within commercial buildings [31]. Motion sensors used for lighting control in modern buildings are inadequate for HVAC control as they fail to detect relatively stationary occupants [9]. Recent works from Erickson et al. [28, 29] and Agarwal et al. [8] have therefore focused on deploying more accurate occupancy sensors within commercial environments, as well as actuating the HVAC system based on the near real-time occupancy information collected. They estimate that the energy use of HVAC systems can be reduced by 30% to 42% effectively in enterprise-scale buildings.

While these occupancy based HVAC actuation systems are indeed effective in terms of reducing HVAC energy usage, they require deployment of additional occupancy sensors and the design, setup and maintenance of the associated data collection network. To examine the upfront installation cost, Erickson et al. [28] report an expense of $147k for just the hardware for a three floor building, and even simple wireless motion sensors would cost over $120k for our five floor building testbed. Most importantly, the deployment and maintenance hurdles are particularly daunting in case of existing buildings with occupants already inhabiting them. Although wireless sensors help reduce the deployment costs to some extent, recent research has shown that it can be very difficult to deploy and maintain a large-scale wireless sensor network in reality [26, 1].

There is a tradeoff between accuracy of detection, cost of deployment and energy savings. This paper presents one such design point whose effectiveness we have quantified. Specifically, we present the design and implementation of *Sentinel*, a system that utilizes a building's existing WiFi network along with WiFi enabled smartphones carried by occupants of that building to infer occupancy and use that information to actuate the HVAC system. We show that even coarse grained information readily available from enterprise WiFi systems such as the Authentication, Authorization and Accounting (AAA) logs of WiFi clients is sufficient in most cases to determine occupancy of office spaces. In contrast to recent in-

SenSys '13 Rome, Italy

frastructure based occupancy solutions [32, 39], Sentinel augments the information collected from the AAA WiFi logs with metadata information such as occupant identity, WiFi MAC address and AP location within the building to improve the accuracy of occupancy detection further.

We have implemented Sentinel on top of BuildingDepot [11], a RESTful webservice that interfaces with legacy building management systems, and show that it is scalable and can actuate the HVAC system in our building effectively. We have deployed Sentinel in the Computer Science and Engineering(CSE) building, a 145,000 sqft enterprise-scale building at UC San Diego(UCSD). We show that Sentinel can effectively determine occupancy in office spaces, covering ∼40% of floor space in the CSE building. We demonstrate the feasibility of using WiFi as a sensing solution by observing the usage pattern of smartphones in CSE and studying the WiFi implementation in modern smartphone operating systems. We find that the requirement for continuous WiFi connectivity contradicts the aggressive WiFi sleep algorithms implemented in smartphones, and provide provisional solutions to maintain WiFi connectivity without significant affect on battery life. Based on ground truth occupancy collected for over 10 days we show that Sentinel accurately infers occupancy 86% of the time, with only 6.2% false negative occupancy detections in personal spaces (Actual=Occupied, Inferred=Unoccupied). We highlight the reasons for the inaccuracy, mostly attributed to aggressive power management by smartphones. Finally, we control 23% of the HVAC zones of our test building using Sentinel in a single day experiment, and measure savings of 17.8% in the HVAC electrical energy consumption.

2. BACKGROUND

Sentinel utilizes several key infrastructures prevalent in modern buildings for occupancy based HVAC control. Our building testbed, the CSE building at UCSD, was built in 2004 and consists of 466 rooms with 145,000 sqft of floor space. The HVAC system in CSE uses a Variable Air Volume(VAV) system to partition our building into independently controllable thermal zones, and a Building Management System(BMS) provides centralized control over the HVAC system. UCSD also employs a modern enterprise-class WiFi system to support the 48,000 strong community. We describe the relevant details of each of these subsystems to explain the framework leveraged by Sentinel. It is important to note that while we do describe these subsystems within the context of the CSE building, they are nevertheless common across modern buildings. For example, VAV based HVAC systems are commonplace as are BMS's, although there can be vendor and instance specific differences. Similarly, managed WiFi infrastructures are common within buildings as is some form of AAA system, although particular implementations may differ.

2.1 HVAC Systems

The HVAC system within our building uses a combination of hot and cold water pipes in conjunction with air-handler units(AHU) to maintain the appropriate thermal environment within the building. Given the size of our university, we employ a central utility plant for producing the hot ($\sim 325\,^{\circ}$F) and cold ($\sim 45\,^{\circ}$F) water distributed campus wide using separate loops as shown in Figure 1. The AHU in our building consists of variable speed drives which supply cold air (converted from the supplied cold water) using ducts to VAV boxes distributed throughout the building. The hot water loop is also connected to these VAV boxes using separate pipes. Each VAV box controls the amount of cold air to be let into an HVAC zone using dampers. A reheat coil, which uses supplied hot water, is

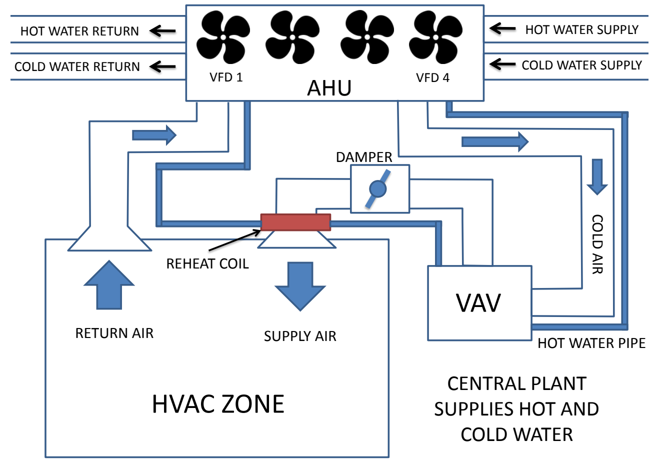

Figure 1: Overview of the HVAC System in commercial building on our campus. Hot water and cold air is pumped to different VAV boxes by AHU. VAV boxes provide independent control in each HVAC zone.

used to heat the cooled air to meet the appropriate HVAC settings for each zone.

Our building is divided into 237 thermal zones, each with its own VAV box to supply conditioned air to that zone. Each zone comprises of two or three small offices and in some cases multiple zones cover bigger rooms such as lecture halls and student labs. The supply air flow and temperature setpoint for each zone is predetermined by the building manager according to the size and maximum occupancy of the zone. Each zone has a thermostat, which allows limited control by allowing the occupant to change the zone temperature setpoint by up to $\pm 1\,^{\circ}$F.

2.2 Building Management System

A central BMS, managed by Johnson Controls, has supervisory control over the HVAC system and the various HVAC components are connected to the BMS via BACnet - a standard protocol for Building Automation and Control networks [19]. Each VAV box has sensors for measurement (zone temperature, air flow, damper position), virtual sensors for monitoring (occupancy status, heating and cooling temperature set points) and control (change set point, change minimum air flow, change occupancy status). Figure 2 gives an overview of BACnet connecting different HVAC subsystems. Sentinel gets access to the BMS by connecting to the network as a BACnet Foreign Device interface. Our facility manager has provided read access to all the sensors in CSE building via this interface.

The primary mode of control of each zone is done using the "Occupancy Command" BACnet point which supports three modes - *Unoccupied, Standby* and *Occupied*. In the *Occupied* mode, the temperature of the zone is maintained within a $4\,^{\circ}$F range with adequate airflow, in the *Standby* mode, the temperature range is increased to $8\,^{\circ}$F with minimal airflow, and in the *Unoccupied* mode, the range is further increased to $12\,^{\circ}$F. During weekdays, all the zones are set as follows: *Occupied* from 6am - 6.30pm, *Standby* from 6.30pm - 10pm, and *Unoccupied* for the rest of the night. On weekends, zones are *Unoccupied* by default, with the occupant required to press a button on their thermostat to mark the zone as *Occupied* for the next two hours.

For occupancy based control, we set the zone to *Occupied* mode

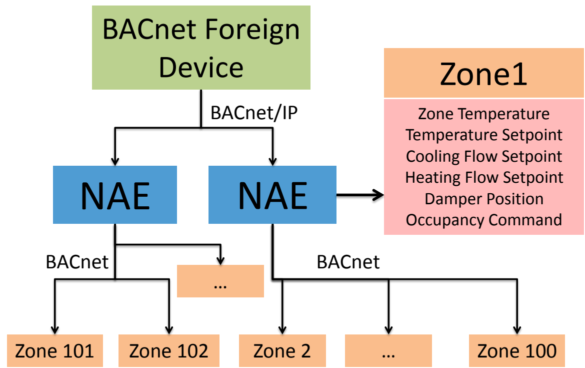

Figure 2: Overview of control system of HVAC using BACnet. Sentinel connects to the BACnet as a Foreign Device and sends commands using BACnet Read Property and Write Property for control of HVAC system.

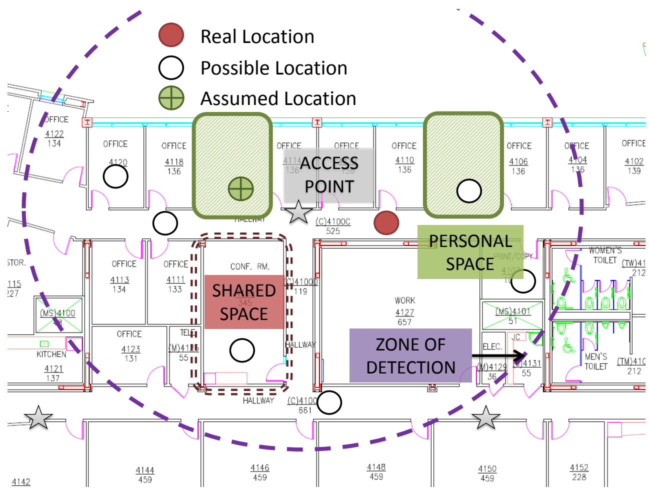

Figure 3: Example of occupancy inference using WiFi connectivity. The occupant is assumed to be in her personal space whenever she is within the associated AP's zone of detection, as denoted by "Assumed Location".

when we detect a zone to be occupied, and set it to *Standby* mode otherwise. We chose a shallow setback temperature for our control to reduce any discomfort to the occupants due to misdetection by Sentinel. Prior research has shown that increased energy savings can be achieved by deeper setback temperature and modulation of ventilation rate based on the number of people in a zone [29, 33, 45]. Thus, the energy savings we demonstrate is a conservative estimate of the savings that could be obtained using advanced control methods. In the rest of the paper, we refer to an HVAC zone being turned *On* and *Off*, which is equivalent to the HVAC zone being set to *Occupied* and *Standby* modes respectively.

We have instrumented our building to measure both the electrical and thermal energy consumption of the HVAC system [12], and and our data shows that the HVAC system in our building consumes 25% - 40% of the total building electricity consumption. We utilize this underlying instrumentation to quantify the effect of our occupancy based actuation.

2.3 WiFi Infrastructure

The enterprise WiFi network in UCSD consists of three SSIDs, one open network - `UCSD-GUEST`, and two secured networks - `eduroam` and `UCSD-PROTECTED`. The two secured networks are mostly identical, and henceforth, we refer to them as the *protected* network. The protected network employs WPA2-E/802.1x for encryption, and authorized users login using their Active Directory username and password. It is common in our building, as we will show in Section 5.1, for occupants to connect to the *protected* network for regular usage. `UCSD-GUEST`, on the other hand, is generally used by visitors of the campus and is insecure with limited access. We describe the specific details of the WiFi logs collected and used by Sentinel in Section 3.2.

3. SENTINEL: SYSTEM DESIGN

Our initial goal was to determine the occupancy of each zone in our building using existing infrastructure without requiring additional sensors or installing any software on our occupants phones. Although we do not achieve this goal completely, we show that it is indeed possible to infer occupancy information for approximately half the zones in our building using WiFi network logs with minimal functionality on client devices.

3.1 Occupancy Inference Algorithm

The idea of localization using wireless radios is well known [16]. Turner et al. [47] studied the performance of established self calibrating WiFi localization algorithms within the CSE building and found that the median and the 95th-percentile error distance of the algorithms to be worse than simple nearest access point location algorithm. The errors were attributed to signal reflection and RSSI variations with time. The accuracy of indoor localization could be improved with fingerprinting algorithms at the cost of significant manual effort [50] or with use of compute intensive algorithms [21]. For our application, we need to localize up to a thousand people in our building for real-time actuation of HVAC zones. Furthermore, we want to develop an occupancy detection solution that relies on minimal information from the network infrastructure. Therefore, for simplicity and scalability, we concentrate on easily obtainable coarse-grained location of client devices, without employing complex localization techniques that may be more accurate. Thus, when a client sends a packet to an access point(AP), we assume that the client is located in a zone within the range of the AP. We show, with the occupancy model described below, that it is possible to make inferences about the occupancy of users in the building even with such coarse-grained information.

3.1.1 Personal and Shared Spaces

We classify physical spaces into two categories: *personal spaces* and *shared spaces*. We define a *personal space* as an area with a designated owner such as individual offices assigned to faculty, or desks assigned to students in a lab. There is no restriction on the size or type of a personal space, so it includes single person offices, cubicle spaces and rooms shared by multiple people. *Shared spaces* on the other hand includes the rest of the building, which essentially have no designated occupant or owner such as restrooms, conference rooms, cafeteria, etc.

Consider an occupant with a WiFi enabled device located within the building as depicted in Figure 3. As the device is associated with one of the access points(APs) in the building, it can be located anywhere in the range of the AP. The occupant could be in her office, or visiting a colleague's office, or in a shared space. We assume that the occupant does not visit a colleague's office unless

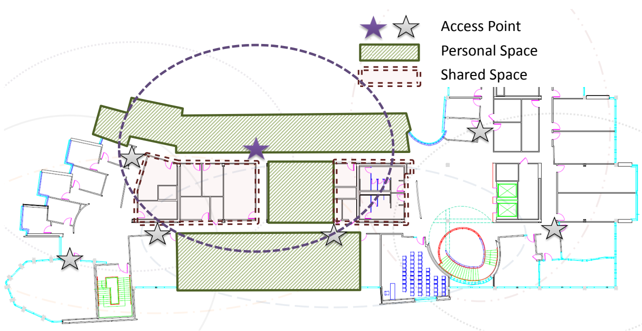

Figure 4: Example of an AP with its associated personal spaces. The network coverage of APs are marked conservatively to reduce false negative errors.

the colleague herself is present in the office. Thus, a personal space cannot be occupied unless the owner is present in the space. If we can detect the presence of owners in their respective personal spaces, then we can effectively monitor the occupancy of all personal spaces in the building.

Shared spaces, on the other hand, can be visited by anyone in the building without restriction. Thus, for inferring occupancy of shared spaces correctly, we would need to detect the entry and exit of **each** person in the shared space accurately. Since we do not employ fine-grained localization information, we **do not** aim to detect occupancy in shared spaces and assume that they are always occupied.

Note that one person can be allocated to more than one personal space, and any number of personal spaces can exist within an HVAC zone. Thus, the personal and shared space division can be applied to a wide variety of buildings and occupancy patterns.

3.1.2 Zone of Detection

We refer to the physical area covered by a WiFi Access Point (AP) as its *zone of detection*. An AP is *affiliated* with a personal space, if the personal space falls within an AP's zone of detection. There can be multiple APs affiliated to a personal space. If the owner of the personal space is connected to an affiliated AP, then she is considered to be present in the personal space.

Smaller zones of detection will naturally lead to more precise occupancy inferences, while larger zones of detection causes loss in accuracy. In our building, we found that the zone of detection of an AP typically covers up to 10 HVAC zones. This lack of precision means that we sacrifice potential energy savings when an occupant is just outside their personal space but inside the zone of detection. Note that even WiFi localization methods will not help as the 95th-percentile error distance from AP was found to be **worse** than the nearest AP algorithm [47]. Figure 3 gives an example of occupancy inference of an occupant who is within the zone of detection of an AP. In this case, the occupant is assumed to be in her office irrespective of their actual location within that zone. This assumption resolves the discrepancy between the area covered by zone of detection of APs and HVAC zones.

We conservatively mark the boundaries of zone of detection of each AP as well beyond the points at which a typical client handoff takes place. We also assume there is no cross floor interference between the AP coverage as it was never observed in practice. For our building, each personal space was associated with at the most four APs. Figure 4 shows an example of the personal spaces associated with one of the APs in the building.

3.1.3 Identity

When the WiFi logs indicate that a client device is connected to a particular AP, we infer that the client is within the AP's zone of detection. In order to make a relation with the personal spaces within the zone of detection of the AP, the client needs to be mapped to the owner of her personal space. Therefore, an accurate mapping of owners to personal space, i.e. occupant to office number, has to be maintained by our system. Further, information of all wireless capable devices used by a building occupant also has to be maintained. As we are using the AAA logs from the WiFi network for inferring occupancy, the wireless device to actual building occupant mapping is available to us.

3.2 Capturing WiFi Data

We use AAA logs from the WiFi network to collect relevant information from the occupant devices. AAA logs only collect the connection, disconnection and periodic live packets from the client devices, which provides us with enough information for occupancy inference. An alternative is to collect data at the AP level and process each packet sent by the device. However, the additional information does not help to improve the accuracy of detection as we show in Section 3.4, but increases the burden of data processing by several orders of magnitude and also intrudes on the privacy of the occupants.

We use the requests received by the RADIUS server as part of the WPA2/802.1x protocol for acquiring information on the WiFi devices in CSE. A WiFi device sends an authentication request to the AP when it first tries to make a connection. The AP forwards the request to the RADIUS server, which has information on the client MAC address, the AP MAC address, the SSID to which connection was requested for, as well as the client username and password. After successful authentication, the AP sends an accounting packet indicating the "Start" of the connection to the server.

Similar authentication and accounting packets are sent to the RADIUS server when a client migrates from one AP to another in the same network, and when the client disconnects from the network. In addition, the AP sends "Alive" accounting packets to indicate the client is still connected to the network. If the AP does not hear from the client for a fixed period of time (1000 seconds in our network), it terminates the connection with the client and sends a "Stop" accounting packet to the RADIUS server.

When the RADIUS packets indicate that the client has connected to one of the APs near the personal space of the occupant, then Sentinel marks that personal space as occupied. When the client migrates to APs in other areas of the building, or gets disconnected from the network, Sentinel marks that personal space as unoccupied.

3.3 Phone Detection Algorithm

There are many WiFi enabled devices popular today - laptops, smartphones, tablets, and it is possible that a building occupant owns more than one WiFi device. When the occupant is moving in and out of her personal space, she may not carry all her WiFi devices. For accurate inference of occupancy, it is important that the system knows the MAC address of the device which is representative of the current location of the occupant. For most occupants in our building, this WiFi device was their smartphone, and henceforth, we refer to the *phone* as the location representative device.

The RADIUS server gets a packet when a client migrates from one AP to another. When an occupant is moving inside the building, the phone gets handed-off between many APs. Over a period of time, the phone would send more number of requests of authentication to the RADIUS server than other devices. Thus, we mark the device with the highest number of requests to be the occupant's phone.

The algorithm fails when an occupant buys a new phone. As the new phone starts off with zero requests, it would be ignored even if it best represents the location of the user. Such an event cannot be ignored at the scale of a thousand occupants, as there could always be a few occupants who have a new device.If we do not see any access request from the device with highest number of requests for 48 hours, we reset the number of requests of all device owned by the occupant. The 48 hour resets also increased the robustness of the system to the changing usage patterns of the occupant.

We verified the accuracy of the algorithm by identifying the MAC addresses used by Sentinel for changing the occupancy status of a personal space. 44 occupants were chosen at random for manual verification, and for 40 of them, the phones were identified correctly. The algorithm worked well for all types of devices despite the aggressive WiFi sleep policies employed (Section 3.4).

We found that Mac OS X devices connected and disconnected from the WiFi network despite being put to sleep mode. Thus, when a Mac OS X computer is left in sleep mode over a weekend, the number of access requests of the computer exceeds those of the occupant's phone, and our system detects the room as occupied. We observed this on four occasions during our experiments, and it can be avoided by incorporating the unique number of access points connected to by a device into the algorithm.

3.4 Perpetual WiFi Connectivity

Sentinel assumes that the phone is continuously connected to the protected network when the occupant is in the building. However, this may not happen in practice because of various reasons - the occupant may not own a smartphone, the occupant may have forgotten her phone at home, the phone may run out of battery, WiFi network coverage may be poor within the office, or there may be a network outage. These problems are associated with any system which seeks to use WiFi clients as a sensor, and we do not handle them as part of this work. If the entire building is affected, Sentinel falls back to the default schedule. If an individual occupant is affected, alternate means of informing the occupancy of an HVAC zone can be provided. In Sentinel, the occupants indicate their presence by pressing a button on the thermostat. We also provide a web interface for indicating user occupancy and preference, similar to the personalized building control system developed by Krioukov et al. [36].

With smart devices permeating every part of our lives, we hope that WiFi connectivity will become part of the essential infrastructure provided in commercial buildings, and the connectivity issues would become a rare event in a few years. Further, as offices typically have abundant power supply, we assume that the occupant would connect the phone to a charger once it indicates low battery. However, battery powered smartphones employ a number of power saving strategies, and the specifics of WiFi sleep algorithm depend on the type of operating system and the model of the device.

We consider three popular variants of smartphones - Android, iOS and Windows Phone. Both Android and Windows Phone provide options for WiFi power management when the device is in sleep mode, and the user can opt to keep the WiFi radio awake even when the device is not in active use. iOS, on the other hand, employs aggressive sleep algorithm as soon as the screen is locked. On studying the network traces of a WiFi-only iPad2 using iOS v6.1.2, we observed that when the device screen is locked, it only keeps the TCP port to Apple Push Notification Service open, and does not respond to other network packets. When the device does not get a push notification for a period of time, the WiFi radio is turned off and woken up at 30 minute intervals. In order to avoid errors in occupancy detection, we request the occupants of the building to change their settings to fetch mail every 15 minutes, thereby ensuring that we get some information coming from them over WiFi.

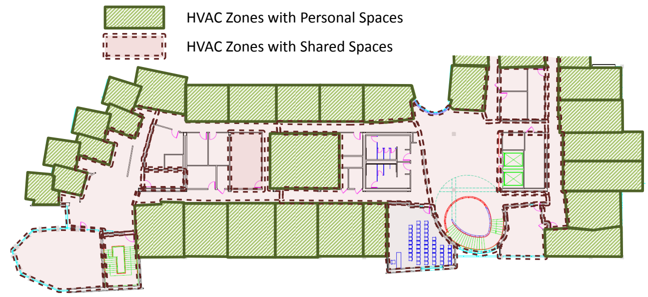

Figure 5: Partitioning of one wing of a floor based on shared and personal spaces. Personal spaces which have a common zone with shared spaces are marked as shared.

3.5 Partitioning the Building

As we explained in Section 3.1.1, we need to divide the building into personal and shared spaces. As Sentinel can only infer occupancy of personal spaces, the energy savings obtained are lower than when actuating entire building HVAC based on occupancy.

In our building, personal space consists of single room offices and multi-person shared offices. The shared space consists of computer labs, cafeteria, conference rooms, etc. In addition, there are storage rooms that are rarely visited, and we mark them as unoccupied for actuation. The HVAC zones in the building, however, do not follow the personal and shared space partitioning. For example, there are several zones which condition a personal space as well as the hallway connected to it. As Sentinel needs to run shared spaces in static schedule, the personal spaces which share its HVAC zone with a hallway or lobby are marked as shared spaces as well. Figure 5 shows an example of shared and personal zone mapping for a section of our building.

Table 1 shows the area covered by each kind of space in our building. Some of the shared spaces like staircases and small hallways are not covered by HVAC zones. Hence, the HVAC power consumption of personal and shared spaces is not proportional to the area covered. To measure the contribution of each type of space to the total HVAC power consumption, we operated the HVAC system with all the zones turned on for one hour, then turned off all the personal spaces for two hours, then switched the personal spaces back on, and finally, turned off all the shared spaces for two hours. We conducted this experiment overnight, as the outdoor temperature is stable at San Diego. On the night of the experiment - March 20, 2013, the outdoor temperature was at $61 \pm 1.7^{\circ}$F.

Table 1 shows the electrical and thermal energy savings obtained turning off shared and personal spaces. The personal and shared spaces contributed 63.9% and 66.9% to the electrical power consumption respectively. Thus, the personal spaces contribute to roughly half of the total HVAC electricity consumption. As the shared spaces remain conditioned in our system, the electrical power savings we can obtain by occupancy based conditioning of personal spaces is $\sim$33% for our building. The heating and cooling thermal power consumption do not follow similar trends, and we examine them in detail in Section 5.6.

	Area	Electrical	Cooling	Heating
Personal	37.5%	63.9%	96.0%	108.0%
Shared	58.3%	66.9%	96.4%	90.0%
Storage	4.2%	-	-	-

Table 1: Contribution of personal and shared building spaces by area and by HVAC power consumption. Actuating only personal spaces can lead to at most ∼33% electricity savings.

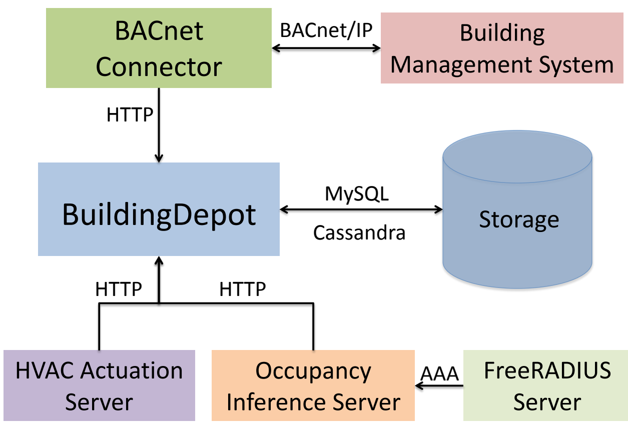

Figure 6: System Architecture of Sentinel

4. IMPLEMENTATION

Sentinel's system architecture follows the principles proposed for management of sensors in commercial buildings in recent literature [13, 24, 25, 44]. Figure 6 provides an overview of Sentinel. BuildingDepot(BD) [11] acts as a central authority for collection of sensory data of the building and provides access control to users and applications for analyzing sensor data and controlling the building actuators. The BACnet Connector acts as a gateway between the sensors which use the BACnet protocol and BD. The Occupancy Inference Server receives a copy of the packets received at the RADIUS server, and processes the packets to infer occupancy for the various HVAC zones in the building. The HVAC Actuation Server processes this occupancy information, and actuates the HVAC system.

4.1 BuildingDepot

BuildingDepot(BD) is the central repository which collects and stores data from different sensors across the building [11]. A connector for each type of sensor protocol translates vendor specific information to a uniform format, and a RESTful API provides access to sensor data and metadata. The sensor data is stored in a Cassandra timeseries database, and sensor metadata, access control lists and user/application specific information is stored in a MySQL database. Each sensor can be queried using contextual information such as location, type and sensor ID.

Each application has to register with BD to gain access to the system. Depending on the permissions provided by the administrator, an application can create/delete sensors, read/write to specific sensors or sensor groups and subscribe to sensor changes. For Sentinel there are three different applications - the BACnet Connector, the Occupancy Inference Server and the HVAC Actuation Server. BD has been designed for enterprise level management of buildings, and can be implemented in a distributed manner. For the CSE building, we have implemented BD in a virtual machine running on top of the Xen VMM. The HTTP server is implemented using Nginx as the web server, and uWSGI is used as the interface between the web server and python application, which is implemented with the Flask framework.

4.2 BACnet Connector

The BACnet Connector(BC) creates a virtual sensor in BD for each BACnet datapoint in the building. The metadata for the sensors are gathered from BACnet object properties, which include sensor type, location, and BACnet specific ID. The connector polls the sensors which are relevant for HVAC zone control, and posts the value to the BD.

For actuation, the BACnet protocol provides a priority array to resolve contention between applications which send actuation commands to BACnet objects. Our BC is assigned a higher priority over the default BMS schedule for actuation of HVAC zones, and any commands sent by the BC will override the default schedule being used by BMS. BACnet also provides a way to relinquish control, so the system switches back to the default schedule when BC does not control the HVAC system.

We have implemented our BC on a desktop machine, which is registered to the BACnet network as a Foreign Device. The BC server is added to the VLAN dedicated to BMS for controlled access to the BACnet/IP network. The connector has been implemented in C, on top of the open source BACnet Stack [3].

4.3 Occupancy Inference Server

The Occupancy Inference Server(OIS) receives a copy of each RADIUS packet sent by the APs in CSE building. OIS processes the incoming packets to infer personal space occupancy as described in Section 3.1.

For inferring occupancy, the OIS maintains several metadata information - a mapping between occupant to their phone MAC address, between the occupants and their office numbers, between offices and the APs in the building, and finally, a mapping between HVAC zones and offices. OIS creates a virtual sensor in BD for indicating occupancy of each HVAC zone in the building, and key information from each incoming packet is stored in a local MySQL database for debugging and future analysis. The usernames are anonymized in the database for preserving the privacy of the occupants.

Several levels of checks need to be made before deciding that an HVAC zone is occupied or not. The incoming packets are filtered for the registered occupants of the building, and then checked if the packets are coming from a "phone"(Section 3.3). If the phone is connected to an AP near the office of the owner, the corresponding personal space is marked as occupied, and otherwise, its marked as unoccupied. If all the other personal spaces in the same HVAC zone is unoccupied, the occupancy status of the zone is updated to occupied and the information is sent to BD.

We have implemented the OIS on top of an open source RADIUS client - `pyrad` [7].

4.4 HVAC Actuation Server

The HVAC Actuation Server(HAS) acts as a layer of abstraction between the occupancy information supplied by OIS and the HVAC control using BACnet. During normal operation, HAS converts the occupancy changes from the OIS to the appropriate commands for HVAC control. HAS was also used for experiments on HVAC control which we describe in Section 5.6.

Currently, we control the HVAC system in a reactive manner, i.e., we control the ventilation of a zone when its occupancy changes. Literature has shown that predictive control with deep setpoints can lead to higher energy savings in HVAC systems [14, 29, 33, 40].

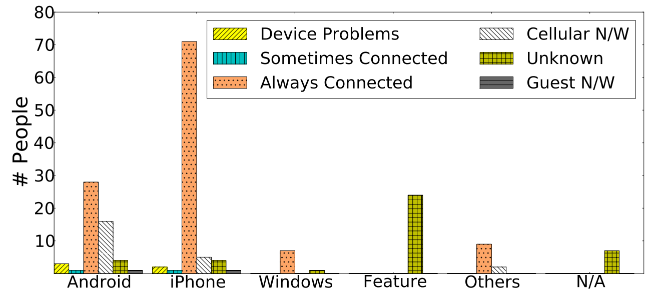

Figure 7: Distribution of smartphones and their WiFi usage patterns by the occupants in our building.

However, the setback temperature setpoints allowed in our building are conservative, and the temperature of unoccupied zones is kept within the range of 70°F to 78°F. Goyal et al. [33] find that the energy savings obtained by both predictive and reactive systems are similar when the setback temperature setpoints are set as per the ASHRAE standard. They also show that reactive systems have negligible effect on the comfort of the occupants as the setback temperature setpoints are conservative. Sentinel is not restricted to reactive control, and we will explore model predictive control as part of our future work.

5. EVALUATION

Sentinel has been operational for three weeks at the time of writing this paper, in the five floor, 145,000 sqft CSE building at UCSD. To show the feasibility of a building-wide deployment of Sentinel, we show the distribution of smartphone usage in the building. We evaluate the accuracy of occupancy detection using Sentinel over a period of 10 days. We then show the occupancy patterns of 38 smartphone users in our building across a week, and identify periods of inoccupancy which could save HVAC energy. We have run over 35 experiments on the HVAC system in our building testbed, and present the HVAC power consumption versus occupancy trends to demonstrate the potential energy savings using an occupancy based HVAC actuation system. Finally, we present the energy savings obtained by controlling 55 of the 237 HVAC zones in the building for one day.

5.1 User Study

We surveyed 187 of the 415 registered occupants in our building. The surveys were short, intended to garner interest in WiFi based control technology. We asked the occupants if they would be interested in using such a technology, the kind of smartphone they use, whether they connected their smartphone to the protected WiFi network in the building on a regular basis, and if they would participate in WiFi based actuation of HVAC system in their office space.

Majority of the occupants surveyed showed interest in controlling the HVAC system based on WiFi connectivity. Over 64% of the occupants owned a smartphone, and only 10% of the occupants did not connect to the internet using WiFi. Figure 7 shows the usage trend of the WiFi devices in the building. Despite the prevalence of WiFi devices and network coverage across the building, many people reported that they did not connect to WiFi due to various reasons - poor WiFi coverage in their offices, adequate data capacity available from cellular network, connectivity problems with the WPA2/802.1x protocol and battery problems.

It should be noted that there is little incentive for occupants of the building to stay connected to WiFi using smartphones in an IT building. Most of the occupants have a desktop computer with ethernet, and many occupants use their laptop for internet connectivity. Several occupants indicated that they would connect to WiFi using their phone if it provided automated control of HVAC system without significant effect on battery life. Problems with network coverage can be solved by careful placement of APs within the building, and device connectivity issues would get solved over time by software/hardware updates to the smartphones. There would always be a few occupants who do not, or cannot connect to the protected network for various reasons. In our experiments, occupants need to indicate their presence by manual press of a button on the thermostat as on weekends. We later added a web based control of access to HVAC system similar to that proposed by Krioukov et al. [36] as a failsafe option.

5.2 Occupancy Accuracy

Accuracy of detecting occupancy using WiFi connectivity has been shown to be noisy and inaccurate in prior work [32, 39, 46]. However, by restricting the occupancy detection of Sentinel to personal spaces, and by using additional metadata information like occupant identity and AP location, Sentinel improves the overall accuracy of occupancy detection significantly. We demonstrate the accuracy of Sentinel based on data collected for 116 of the 415 building occupants over a 10 day period.

57% of the smartphones used by the building occupants are iPhones, and as explained in Section 3.4, iOS devices turn off the WiFi radio when it is not in active use. To participate in WiFi based HVAC control experiments, we requested occupants to keep their iOS device connected to WiFi and to change device settings to fetch emails every 15 minutes. We requested the Android and Windows Phone users to enable WiFi and to change the settings to disable the WiFi aggressive sleep option. The change in device settings were enforced for two days, and the occupants were given the option to change back to their default settings if needed.

We define an *event* as a change in occupancy of a personal space, either as detected by Sentinel, or as seen in ground truth measurements. We use the number of events correctly identified by Sentinel as a measure of the occupancy accuracy. If Sentinel incorrectly marks a personal space as occupied, we classify the error as a *false positive*, and if the system incorrectly marks a personal space to be unoccupied, we classify it as a *false negative*. On a false positive error, we incur a penalty in the energy savings obtained as the HVAC system would ventilate the personal space unnecessarily. A false negative, on the other hand, would lead to discomfort to the occupants as the HVAC system would be put to "Standby" mode. For ground truth comparison, we note the occupancy in each office across the building, and compare Sentinel logs for occupancy status at the corresponding timestamp. We also inspect the latest logs from Sentinel, and examine the occupancy status of the respective zones. In case of discrepancy, we try our best to identify the underlying cause. We ignore the errors that occur when the occupant leaves her personal space for less than five minutes, and since the timeout period in RADIUS protocol is $\sim$ 17 minutes, we accept a delay of up to 20 minutes in detection when the occupant is leaving her personal space.

We measured 436 events during the 10 test days, of which 330 events were recorded in the first two days, and Sentinel accurately identified personal space occupancy 83% of the time. The false positives and the false negatives were 9.4% and 7.5% respectively. After the first two days, the ground truth was collected only for occupants known to be still using the modified phone settings. Figure 8 gives a breakdown of the causes of the errors in detection.

Majority of the false positive errors by Sentinel were caused due to an error in identifying the appropriate device by the phone detection algorithm. As many of the occupants were enabling their WiFi devices for our experiments, we reset the access count request of all the recorded occupant devices. As this was done early in the morning, all the WiFi enabled devices in the building were identified as phones by Sentinel, and the errors in detection increased. The phone detection algorithm corrected itself as occupants came in, and the incorrect device errors died down by midday.

System errors constitute the errors caused due to mistakes in metadata information stored in Sentinel. Some of the errors included incorrect mapping of the occupant to their personal space, incorrect authentication username, and incorrect mapping of APs to personal spaces. We corrected the errors after the first day of ground truth data collection. If we remove the temporary errors caused due to incorrect device detection and system configuration errors, the accuracy of Sentinel improves to 86%, and the false negative errors reduce to 6.2%.

The aggressive WiFi sleep mechanism used in iOS devices led to intermittent WiFi connectivity, and was the cause of majority of the false negative errors. Although, we used various mechanisms to keep the WiFi radio active, there were still circumstances in which the connectivity was not persistent. "iOS Start" errors indicate that the occupant has entered her personal space, but Sentinel could not detect the occupant as the iOS device did not switch from the cellular data network to the WiFi network. We noticed a maximum delay of 23 minutes in iOS Start errors. "iOS Stop" errors occur when the iOS device turned off the WiFi radio when the occupant was in her personal space. This behavior was observed among phones which were not in use for a long period of time, and as much as 3 hour periods of disconnection were observed. However, on most cases, the iOS devices woke up within 10 minutes of timeout. The device errors were mainly caused due to late detection of arrival of occupants in to their personal spaces. The late detection was observed among the Android devices, as it sometimes took longer than usual to detect WiFi networks in its vicinity. The inaccuracies due to device connectivity constitute 5% of the error and can be improved by the use of an app on the phones. We provide the details of an iOS app which addresses this issue in Section 6.

When the occupant has left her personal space, but is still within the zone of detection of the nearby AP, Sentinel incorrectly marks the space as occupied. We call such false positives as "zone of detection error". A similar false positive is incurred when occupant leaves her personal space but does not carry her phone with her. We classify such error under "people error". Occupants also sometimes forgot to enable WiFi on their phones, or connected it to the guest network, which leads to false negatives. We classify such errors as people error as well. Both zone of detection and people errors account for 6.9% error in occupancy detection, and are inherent to the occupancy inference algorithm used by Sentinel. People errors can only be reduced using wearable devices, and zone of detection errors can be reduced using accurate localization methods.

5.3 Occupancy Trends

We have collected the occupancy information inferred from the RADIUS logs for all the occupants for three weeks at the time of writing this paper. Figure 9 shows the occupancy of the 38 users who are always connected to the protected network, and have disabled WiFi sleep by default. The occupancy trend is shown for the week of March 18 to March 24, 2013 - one of the busier weeks in our building due to exams. Note that occupancy here refers to occupancy of personal spaces, rather than the whole building.

The most interesting part of Figure 9 is that the peak of the graph

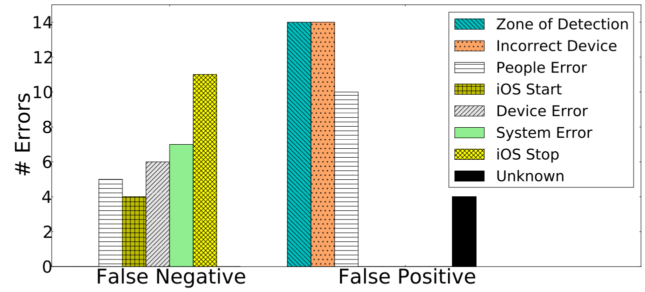

Figure 8: Distribution of occupancy detection errors as observed over 436 events and 10 days. The occupancy detection was accurate 83% of the time.

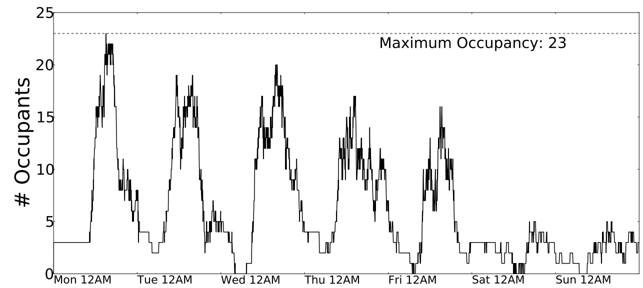

Figure 9: Occupancy trends of 38 occupants in our building who keep their smartphones always connected to WiFi as measured by Sentinel for the week of March 18-24, 2013

is at 23 people, only 57% of the maximum 38. Another point of interest is that the general occupancy decreases as the week progresses, indicating peak of productivity on Monday, and a maximum of just 15 people on Friday.

On most days, there is a fall in the occupancy during the middle of the day, indicating people leaving their offices for lunch, meetings and discussions. The graph clearly demonstrates the opportunity of energy savings that could be obtained by controlling the HVAC system based on occupancy.

On nights and weekends, the occupancy is understandably low, however it is not zero, as assumed by the static schedules used for HVAC control. The occupants are left to manually indicate their presence if they are in the building during off hours. WiFi based occupancy detection can easily detect the presence within an HVAC zone, and provide automated thermal comfort to the occupants.

5.4 Impact on Device Battery Life

Battery life of a device is dependent on the WiFi radio chip, the network coverage, the applications using WiFi, the usage pattern and potentially other factors. Prior works on WiFi and cellular radio power measurements [17, 20, 51] indicate that WiFi sleep power is about 2x the sleep power of cellular technologies such as 3G, the data transmission in WiFi is about an order of magnitude more efficient than cellular, the energy spent by WiFi radio to scan and associate to an AP is 5x the energy spent for 50KB data transfer, and the energy consumption of cellular radio varies significantly with signal strength. To reduce the impact of higher WiFi sleep power, Rahmati et al. [43] and Agarwal et al. [10] suggest waking up WiFi only when data transfer is required, and iOS follows a similar model based on our observations (Section 3.4). However, this strategy may not lead to power savings if the phone keeps switching between WiFi and cellular radios frequently or if

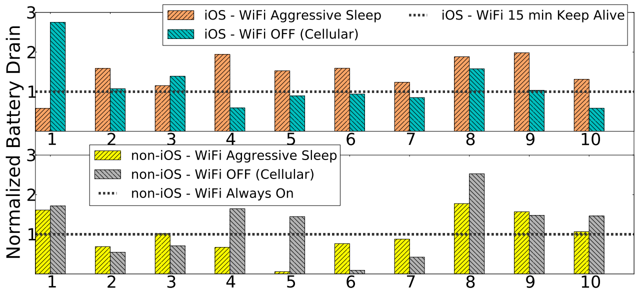

Figure 10: Distribution of smartphone battery consumption of 20 participants over 3 days with WiFi "always on", with WiFi aggressive sleep enabled and with WiFi off.

the apps installed on the phone require frequent data transfer. Thus, the impact on battery life would actually depend on the usage pattern of the phone.

Instead of measuring battery consumption in a controlled environment, we measure battery drain as seen by phone owners during their regular usage. We choose 20 participants, not necessarily building occupants, and measure their smartphone battery performance over three days. There were 10 iPhones, 9 Androids and 1 Windows Phone in the collection. On the first day, the smartphones were put to WiFi "always on" mode, by disabling the sleep mode in non-iOS phones, and fetching email every 15 minutes in iOS phones. On the second day, WiFi was enabled, with aggressive sleep mode enabled. On the third day, WiFi was switched off completely. The participants were requested to try and keep similar usage pattern across these three days and report any significant differences in usage. We normalize the battery drain during three days by the battery drain observed with WiFi "always on" option, and the combined result is shown in Figure 10.

As can be observed from Figure 10, there are no clear trends across the three WiFi modes for these devices. However, we do make several observations. First, in many cases (particularly for iOS devices) WiFi aggressive sleep leads to lower battery lifetime than keeping WiFi on probably due to the constant mode switches. Second, turning the WiFi off completely to use only 3G does not lead to significantly better battery life as compared to keeping WiFi on, or the 15-min Keep Alives mode for iOS. The Android device for which this is not the case (Device 6) were verified to be an anomaly since the user reported that they don't use the 3G data radio. Therefore, based on our current data, we have not seen conclusive evidence whether using the aggressive sleep modes for WiFi actually provides significant battery life improvements than the less aggressive WiFi on settings. However, given the variations we observed in battery consumption more extensive data collection would provide better insight into the effect on battery life due to continuous WiFi connectivity.

5.5 Actuation Latency

Unlike prior occupancy based control systems [8, 28], we have implemented Sentinel on top of RESTful web services as recommended in recent literature [25, 13] using our BuildingDepot(BD) system [11]. BD is designed to support different types of building applications, is compatible with existing building management solutions and scales well with number of users, applications and sensors. Similar RESTful frameworks are also being adopted by industry and academia for building automation applications such as plug level energy meter [4, 35] and wireless lighting system [6]. Sentinel is one of the first RESTful systems to be deployed at the scale of an enterprise-scale commercial building, and actuaion latencies for such systems have not been measured in the literature so far.

Operation	Latency (in ms)
OIS $\rightarrow$ BD	194.26 $\pm$ 50.6
BD $\rightarrow$ HAS	67.18 $\pm$ 13.6
HAS $\rightarrow$ BD	158.25 $\pm$ 61.3
BD $\leftrightarrow$ BC	185.35 $\pm$ 113.4
BD $\rightarrow$ HAS	126.35 $\pm$ 30.6
Total	**731.57 $\pm$ 125.4**

Table 2: Breakdown of latency of Sentinel from the time of reception of RADIUS packet from WiFi device to the time of sending actuation commands to HVAC.

Table 2 provides a detailed breakdown of latency to send an actuation command to the HVAC zones, from the time of detection of occupancy to the time to get the acknowledgment of the command completion. We have an actuation latency of ~750ms, which is fast enough for actuating HVAC systems. However, when the access controls extend to plug loads and lighting systems, the actuation latency would need to be reduced further so that the occupants do not notice the delay.

5.6 Potential Energy Savings

Prior work has focused on estimating the energy savings obtained by occupancy based actuation of HVAC system using simulations on calibrated EnergyPlus building models [23], and it has been shown that significant savings can be obtained across different seasons and geographical locations [28, 29, 33]. Goyal et al. [33] show that the amount of energy savings obtained remain almost the same for both reactive and predictive strategies for different outdoor conditions if the set back temperatures are conservative as per ASHRAE standards.

Instead of simulations, we measure the actual energy savings obtained at different levels of occupancy by conducting experiments directly on our building. We perform our experiments during night time, as there are only a few people present in the building, and the night temperature at San Diego was relatively stable at the time of our experiments. All the experiments were conducted during the month of March, 2013, when the night temperature was recorded between 55 °F and 60 °F. Note that compared to the day, the load on the HVAC system during night is lower due to reduced outdoor temperature and lesser number of people and machines in operation. The energy savings measured represent a constant load HVAC system, and is a conservative estimate of actual energy savings possible.

To determine the energy savings obtained with change in occupancy in the building, we randomly choose a fixed percentage of HVAC zones, and turn them off for a period of two hours. To allow for variations with respect to outdoor conditions, we choose the same set of zones, and repeat the experiment. Figure 11 shows the electricity consumption of the HVAC system when we actuated 25% of the zones in the building. The experiment was started at ~10pm, and all the zones in the building were gradually turned on with an interval of 10 seconds between each actuation command. The zones were allowed to stabilize for an hour, and then 25% of the zones were gradually turned off for a period of two hours. We turn back on the switched off zones after two hours, and repeat the process once more. We repeat the experiment for at least two nights for each level of occupancy.

Figure 12 shows the changes in electrical power consumption

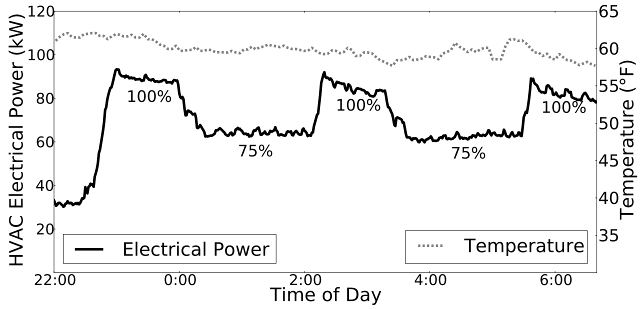

Figure 11: Measurement of HVAC electrical power consumption with 25% of the HVAC zones randomly chosen to be alternatively turned on and off on the night of March 16, 2013.

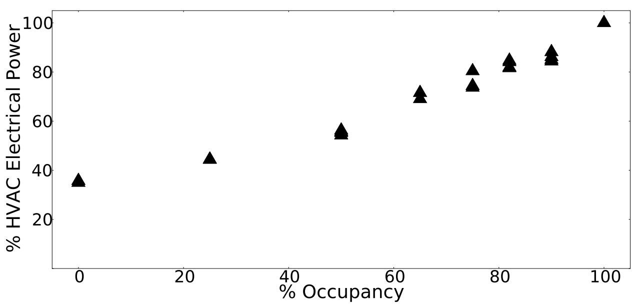

Figure 12: HVAC electrical power consumption with change in occupancy levels in the building.

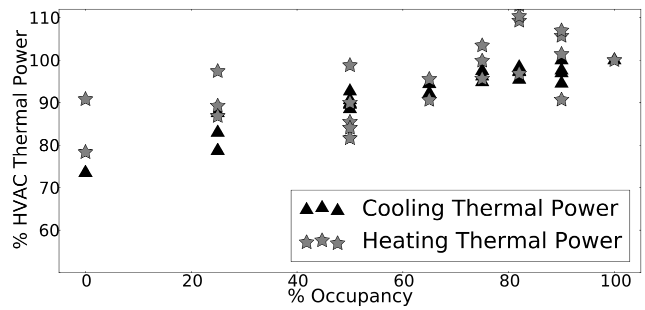

Figure 13: HVAC thermal power consumption with change in occupancy levels in the building

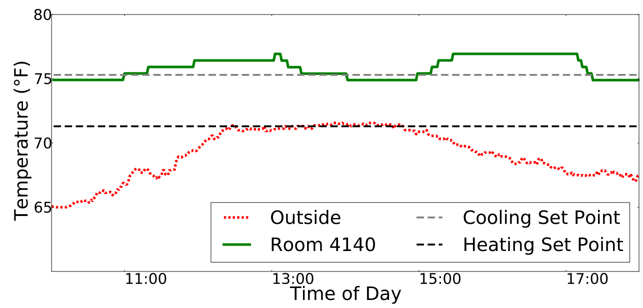

Figure 14: Temperature profile of an HVAC zone during daytime when it was turned on and off every two hours. Heating and cooling setpoints are $71\,^{\circ}\mathrm{F}$ and $75\,^{\circ}\mathrm{F}$ respectively.

of the HVAC system with increase in occupancy of the building. There is a clear increase in the electrical power consumption as the occupancy of the building increases. Although its not prominent in the figure shown, the drop in electrical energy is not directly proportional to the fall in occupancy within the building. The electrical power consumption is dominated by the fans in the Air Handler Unit(AHU) of the building, and power consumption of the fans are proportional to the cube of the fan rotation speed. Thus, as the occupancy of the building increases, the fan rotates at a higher speed, leading to disproportional increase in power. Thus, the energy savings are maximum when the occupancy of the building drops from 100%, and follows the pattern of diminishing returns as the occupancy further reduces.

Figure 13 shows the thermal power consumption of the HVAC system with increase in occupancy. Both cooling and heating thermal power decrease gradually with decrease in occupancy of the building. The trends in heating thermal power is not as clear as cooling thermal power or electrical power because the supplied hot water is not in continuous use by the HVAC system. The cold water is converted to cold air, and is used for ventilation by the VAV boxes. The amount of cold air is regulated by the VAV box using a damper, but a minimum amount of ventilation is maintained by the VAV even when the zone is unoccupied. Hot water, on the other hand, is used intermittently by the VAV box to reheat the cold air when needed. The intermittent usage of hot water translates to different heating thermal power consumption from day to day, and thus, we do not see any clear trends with change in occupancy.

Even when the building is completely unoccupied, electrical power consumption is $\sim$35% of the power consumption at full occupancy, and heating and thermal power is at $\sim$70%. As the building is put in to "Standby" mode when it is unoccupied, the HVAC system still tries to maintain minimum thermal comfort within the building. For our building, the temperature guardband is increased by $2\,^{\circ}\mathrm{F}$ on both cooling and heating setpoints with respect to the setpoints in "Occupied" mode.

The thermal power consumption is still high compared to electrical power when the building is fully unoccupied. This is because the cold water is used for cooling the server room in CSE, and the hot water is used for domestic water heating. Also, recall that our building receives its hot and cold water from a central utility plant(Section 2.1), and thus, the reduction in thermal energy observed is due to the decrease in the demand for hot and cold water. However, as the hot water and cold air still circulate through the building, there is still a drop in temperature in the returned hot and cold water. The thermal power consumption is measured as the energy spent due to the loss in the temperature difference between the supply and return water. Hence, even at zero percent occupancy, significant amount of energy is spent for thermal needs.

5.7 Thermal Comfort

Prior work suggests that reactive control of HVAC system does not lead to occupant discomfort when the setback temperature is conservative [8, 33]. To test this in our building, we performed a controlled experiment on a subset of HVAC zones. We chose 12 HVAC zones, each of them having different characteristics in terms of size, location, and number of rooms. Each of the zones were alternated between "Occupied" and "Standby" modes for two hour periods over a total period of 8 hours during the day on a weekend.

One of the HVAC zones had a faulty sensor, and we do not consider its temperature data. Figure 14 shows the variation in temperature of the HVAC zone which showed the *maximum* thermal discomfort among the remaining 11 zones in the experiment. The

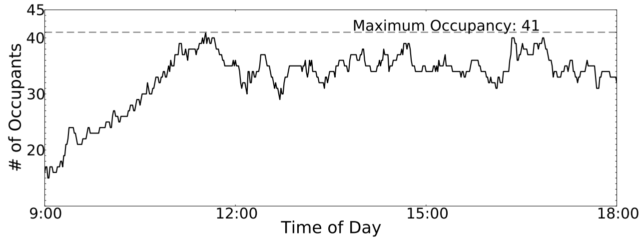

Figure 15: Occupancy trends of 116 volunteers on March 26, 2013, the Experiment Day.

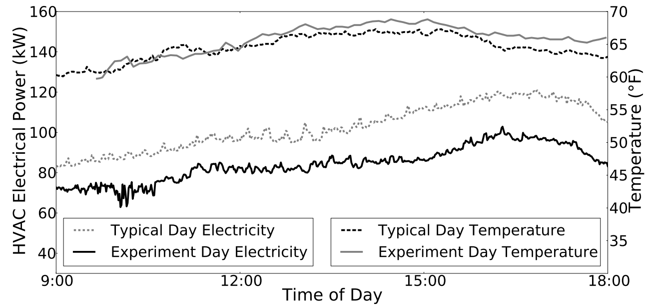

Figure 16: Comparison of outdoor temperature and HVAC electrical power consumption of the Typical Day and Experiment Day. Total savings of 17.8% in electrical energy was obtained for the duration of the experiment.

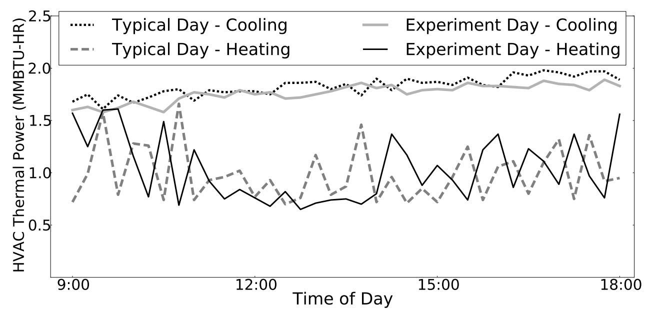

Figure 17: Comparison of HVAC heating and cooling thermal power comparison of the Typical Day and Experiment Day. No clear trends can be observed, and only 0.8% energy was saved for the duration of the experiment.

heating and cooling temperature set points of the "Occupied" mode for this zone was at 71 °F and 75 °F respectively, and the corresponding set points of the "Standby" mode was 69 °F and 77 °F respectively. Unfortunately, the outside temperature at the time of the year is temperate, and does not change the temperature of the zone significantly, even when it is in the Standby mode. It is clear from Figure 14 that the temperature of the HVAC zone never exceeds 77 °F, and quickly drops to 75 °F as soon as the zone is switched to "Occupied" mode. Thus, we confirm that the finding by Goyal et al. [33] by real temperature measurements that the thermal comfort is minimally effected when the setback temperature setpoints are conservative.

5.8 Energy Savings with Sentinel

We controlled the HVAC system of our building testbed using Sentinel for the 116 volunteers from 9am to 6pm on March 26, 2013. Of a total of 237 HVAC zones, we controlled 55 zones distributed across three of the five floors in the building.

As HVAC zones are often shared between rooms, the actuation policy of the occupants located within an HVAC zone needs to be the same. As a result of this sharing, some of the personal spaces needed to be converted to shared spaces, as explained in Section 3.5. Similarly, the occupants who could not participate in the experiment, share their HVAC control policy with our volunteer occupants. Therefore, a single non-eligible participant in an HVAC zone forces us to treat the entire zone as a shared space. Despite this limitation, we control 55 out of 237 HVAC zones in the building for our actuation experiment. As we are requesting the occupants of the building to shift from their regular usage patterns, we had to limit our control experiment to just one day. Of the 55 zones covered by the experiment, 12 zones were known to be unoccupied apriori on the day of the experiment, and we turned them off for the duration of the experiment.

We compare the energy consumption on the day of our experiment (March 26, 2013) with the energy consumption on March 22, 2013, as the temperature profiles of the two days were similar. We refer to the day we controlled the HVAC system using Sentinel as "Experiment Day", and refer to the day of comparison as the "Typical Day". Other close days were cloudy, and we could not use them. Figure 16 show the electrical power consumption from 9am to 6pm on the Experiment and Typical days, Figure 17 shows the thermal power consumption of HVAC in the same time frame, and finally, Figure 15 shows the occupancy of the 116 volunteers on the Experiment Day as measured by Sentinel.

We saved 17.8% of electrical energy on the Experiment Day, as compared to the Typical Day. Occupancy trends from Figure 15 shows that the building occupancy gradually increases from 9am to 11am, and remains roughly constant till 6pm. However, the occupancy peaks at 40 people, indicating most of the volunteer occupants were not present in the building during the period of experimentation. The relative inoccupancy was expected, as the Experiment Day was the second day of the spring break at our university.

The occupancy trend is clearly reflected in the electrical power consumption of the HVAC system, as it initially starts off lower than the typical day at 9am due to the reduced number of occupants in the building. As the occupancy within the building increases, the power consumption also increases gradually until 11am. From 11am to 6pm, the electrical power consumption of both the days follow the same pattern, in accordance with the changing outdoor weather conditions. The energy savings from 11am to 6pm is mainly obtained because of the occupants who did not come in to their personal spaces on the Experiment Day. The 17.8% electrical energy savings obtained is in accordance with electrical power consumption trends shown in Figure 12, where the corresponding building occupancy is $\sim$90%.

As our Experiment Day falls on university spring break, but our Typical Day is during exam week, it is possible that part of energy savings occur due to reduced activity in the building. We compared the HVAC electrical power consumption on Experiment Day with two other spring break days (March 27 and 28, 2013) with cloudy weather conditions when the HVAC was under static schedule based control, and still measured electrical energy savings of 7.5% and 11.8% respectively.

The trends in thermal power consumption on the Experiment Day were not as clear. Cooling thermal energy consumption decreased by 2.2%, but the heating thermal energy actually increased by 1.5%. Figure 13 indicates that the thermal energy consumption is also consistent with our night time trending experiments, and the heating thermal power consumption sometimes increased despite a reduction in building occupancy.

The experiment provides an example of the energy savings that could be obtained across one particular day by controlling 23% of the HVAC zones in CSE. However, the long term energy savings will be different due to varying weather conditions or occupancy patterns. As long term occupancy patterns are not available, we do not attempt to project the energy savings obtained by simple extrapolation of trends we see for one day.

6. DISCUSSION

The occupancy inference algorithm proposed in this paper uses the metadata information available and typical occupancy patterns within offices to mitigate the inaccuracies associated with locating a WiFi enabled device with respect to its AP. The algorithm can be adapted to a wide range of office spaces independent of its building topology, or usage patterns. To infer occupancy using WiFi we use key metadata relating authorized occupants with their personal office space, their WiFi device MAC address, network logs to determine the current status of network connectivity with occupant devices and the location of APs within the building.

For adoption of our solution in a commercial building, a dependable and easily accessible fallback solution needs to be provided to the building occupants. Occupants should be able to inform the BMS of their presence easily in case they forget their phones at home, or need to lend their office to a visitor. The personalized building control proposed by Krioukov et al. [36] provides a good platform for user feedback, and we have implemented a similar web based interface for CSE. Automated tools for keeping track of occupants in personal spaces, mapping of APs to personal spaces and HVAC zones to office spaces would also help in quick deployment.

Reliable WiFi connectivity from the users phones is the only requirement from the occupants of the building for the proposed algorithm. However, as we saw in Sections 3.4 and 5.4, it is difficult to maintain perpetual connectivity in iOS devices, and there may be an effect on battery life of devices when they are always connected to WiFi. IEEE 802.11ah standard [15] is being designed specifically for low power, low data rate applications, and would enable applications like Sentinel without affecting battery life. In the meantime, we plan to develop mobile apps which would maintain WiFi connectivity and still have minimal effect on battery life. The apps would break the non-intrusive model of deployment, but can be integrated with the personalized building control system [36]. Alternatively, prediction mechanisms can be used to eclipse the intermittent connectivity of WiFi devices.

As most of the false negative occupancy detection errors in Sentinel is caused by iOS devices, we have already developed an iOS app. The app creates a geofence on the building, and wakes up the device when it enters the geofenced area. The app keeps the device awake until it connects to WiFi, and then allows the device to go to sleep. Periodic push notifications from the app wake up the WiFi radio, and the notifications are turned off when the device leaves the geofenced area. However, we have not yet evaluated the app extensively to present its performance results here.

Sentinel only targets personal spaces in office buildings. To improve HVAC energy efficiency further, shared spaces should also be regulated according to occupancy. One option is to install wireless sensor network solutions [8, 28] just for the shared spaces. Use of calendars has been proposed as a proxy for occupancy [25], however it is not applicable to several kinds of shared spaces like lobby, cafeteria, etc. Indoor localization has the potential to reduce the zone of detection enough for occupancy inference in shared spaces. We plan to explore infrastructure based localization techniques as part of future work.

The HVAC zones in modern buildings are not designed for occupancy based actuation. Although VAV systems have become commonplace since the late 1990s [34], the zones normally map several individual rooms. If only one of the rooms within a zone is occupied, the remaining rooms within the zone are unnecessarily ventilated. Further, sharing of HVAC zones between shared and personal spaces, requires conditioning of personal spaces whenever the shared space is occupied. Smaller and more insulated HVAC zones would lead to more savings based on occupancy control in lieu of higher installation cost. If the architects of the HVAC system incorporate occupancy based control into their design for next generation buildings, there could be a significant reduction in the running cost of the system.

7. RELATED WORK

Occupancy based HVAC control has been studied extensively for improving building energy efficiency [8, 18, 27, 30, 29, 28]. Extensive simulation studies and practical deployments in commercial buildings have shown that 15% - 42% energy savings can be obtained using occupancy based control, depending on weather conditions, building type and occupancy variation.

Several occupancy detection mechanisms have been developed over the years for HVAC control. CO_2 sensors are used for occupancy based control of high capacity spaces such as auditoriums and conference rooms [2, 5], but have been found to be too slow to respond to change in occupancy for smaller rooms found in commercial buildings [31]. Passive infrared(PIR) motion sensors have been used in modern buildings for actuation of lighting systems. PIR sensors often fail to detect occupants when they are relatively motionless, such as while reading or typing. Further, they are vulnerable to calibration errors, external triggers by sunlight or air draft and only provide binary occupancy information. These limitations make it challenging to use PIR sensors for HVAC control. Our own work improved upon these limitations with the addition of door sensors to obtain occupancy accuracy of 96% and demonstrated up to 15% savings in HVAC electrical energy for one floor deployment in CSE [8]. However, the occupancy detection mechanism is only accurate for single person offices, and depend on the occupants to close the door while exiting the office.

The POEM system [28] uses a combination of ceiling mounted camera and motion sensors to obtain 94% accuracy in occupancy detection. Erickson et al. use the near real-time occupancy information from the sensors for predictive control of 30% of the HVAC zones in an office building and demonstrate up to 26% energy savings. The cameras used in POEM exploit the hallway topology for occupancy detection. If the office spaces are located around a circular hallway, or use open cubicle spaces, the image processing algorithms would have to be modified and re-calibrated. Further, use of battery powered wireless sensor nodes in POEM involves changing batteries every 45 days.

In contrast, Sentinel provides a solution for actuation of HVAC system using near real-time occupancy derived from existing WiFi infrastructure. Sentinel neither makes any assumption regarding the topology of the building, nor requires careful calibration of sensors. Leveraging existing infrastructure allows Sentinel to be quickly deployed and easily maintained. The monetary and ease of use benefits of Sentinel comes at the cost of assumptions on usage patterns of WiFi devices by building occupants, and works only for personal spaces. Additional sensors or localization techniques still need to used for occupancy detection in shared spaces.

Numerous methods for occupancy detection have been developed that leverage existing infrastructure such as powerline [42], speakers [37], WiFi [16, 50, 21], geo-magnetism [22], HVAC ductwork [41], or a combination of these [49]. However, many of these solutions do not work well for HVAC control in commercial buildings due to issues of scale [42, 21], use of specialized sensors [41, 37], extensive war driving [50, 22] or complex functionality in client devices [49].

Existing WiFi infrastructure potentially provides the scalability needed for commercial buildings, does not rely on client device functionality and eases deployment and maintenance. Ghai et al. [32] use a combination of WiFi signals, calendar schedules, personal computer activity and instant-messaging client status to infer the occupancy within cubicles with an accuracy of up to 91%. The algorithms have been evaluated for just 5 volunteers, and do not evaluate scalability. In contrast, we only use WiFi information, and show the efficacy of our algorithm over 116 occupants of our building. Melfi et al. [39] use DHCP leases within a real building for occupancy inference and found the accuracy to be low - 31% to 84%. The inaccuracies of their system were attributed to unpredictable coverage provided by APs and intermittent connectivity of the WiFi devices. We overcome the limitations of WiFi sensing by using additional known information such as occupant identity, occupant office location and focus on personal spaces. Martani et al. [38] use WiFi logs to determine the live WiFi connections within a building, and provide a breakdown of the WiFi connections on a floor and room basis. They show that WiFi connections correspond well with the HVAC energy consumption of a building at MIT. However, they make no attempt to correlate WiFi connections with the ground truth occupancy.

8. ACKNOWLEDGEMENTS

We would like to thank the UCSD Facilities Management and Academic Computing and Telecommunications department, who have helped us get access to the buildings and to the wireless network traces. We would like to thank Thomas Weng for providing access to power measurements and the occupants of Computer Science and Engineering building for participating in the experiments. We would like to acknowledge Roshni Chandrashekhar, for her comments on the early drafts of this paper. We also wish to thank our shepherd Mahadev Satyanarayanan and our anonymous reviewers for their feedback to help improve the paper. This work was supported in part by NSF grant SHF-1018632 and a UC San Diego facilities grant for smart building deployment.

9. CONCLUSION

We have presented the design and implementation of Sentinel- an occupancy based HVAC actuation system that leverages existing WiFi infrastructure and occupants with WiFi enabled smartphones within commercial buildings to reduce HVAC energy usage. In contrast to prior occupancy sensing solutions which required installation of additional sensors and associated wireless sensor networks, utilizing existing infrastructure for occupancy sensing reduces the costs and effort of deployment and maintenance significantly. We reduce the inaccuracies in occupancy sensing using noisy WiFi signals by using metadata information about the occupants, access points and the HVAC zones in the building. We have deployed Sentinel in a 145000 sqft commercial building, and show the accuracy of occupancy detection within office spaces to be 86%, with only 6.2% false negative errors. Furthermore, we provided a detailed analysis of the reasons for these inaccuracies, largely due to aggressive power management by smartphones. Based on our battery lifetime measurements across a number of devices we show that using less aggressive WiFi power modes, which improve accuracy of Sentinel, do not necessarily lead to significantly reduced battery life. We also discuss potential solutions, such as an App on users phones, that can increase the accuracy of WiFi based occupancy detection even further. Finally, we demonstrate occupancy based control of 23% of the HVAC zones of our building testbed using Sentinel and measure electrical energy savings of 17.8% in the HVAC system compared to the static scheduling based control used across the buildings on our campus.

10. REFERENCES

[1]

[2] Aereco - Demand Controlled Ventilation. `http://www.aereco.com/ventilation-systems/demand-controlled-ventilation`, Mar. 2013.

[3] BACnet Stack. `http://bacnet.sourceforge.net/`, Mar. 2013.

[4] Enmetric Systems. `http://www.enmetric.com/`, Mar. 2013.

[5] FPL - Demand Controlled Ventilation. `http://www.fpl.com/business/energy_saving/programs/interior/dcv.shtml`, Mar. 2013.

[6] Philips Hue. `https://www.meethue.com/`, Mar. 2013.

[7] pyrad 2.0. `https://pypi.python.org/pypi/pyrad`, Mar. 2013.

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

[9] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng. Occupancy-Driven Energy Management for Smart Building Automation. In *Proc. of ACM BuildSys*, 2010.

[10] Y. Agarwal, R. Chandra, A. Wolman, P. Bahl, K. Chin, and R. Gupta. Wireless Wakeups Revisited: Energy management for VoIP over WiFi smartphones. In *Proc. of ACM MobiSys*, 2007.

[11] Y. Agarwal, R. Gupta, D. Komaki, and T. Weng. BuildingDepot: An Extensible and Distributed Architecture for Building Data Storage, Access and Sharing. In *Proc. of ACM BuildSys*, 2012.

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

[13] P. Arjunan, N. Batra, H. Choi, A. Singh, P. Singh, and M. B. Srivastava. SensorAct: A Privacy and Security Aware Federated Middleware for Building Management. In *Proc. of ACM BuildSys*, 2012.

[14] A. Aswani, N. Master, J. Taneja, D. Culler, and C. Tomlin. Reducing Transient and Steady State Electricity Consumption in HVAC using Learning-Based Model-Predictive Control. *Proc. of IEEE*, 2012.

[15] S. Aust, R. V. Prasad, and I. G. Niemegeers. IEEE 802.11 ah: Advantages in standards and further challenges for sub 1 GHz Wi-Fi. In *Proc. of IEEE ICC*, 2012.

[16] P. Bahl and V. N. Padmanabhan. RADAR: An In-Building RF-Based User Location and Tracking System. In *Proc. of IEEE Infocom*, 2000.

[17] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani. Energy consumption in mobile phones: a measurement study and implications for network applications. In *Proc. of ACM SIGCOMM IMC*, 2009.

[18] M. J. Brandemuehl and J. E. Braun. The Impact of Demand-Controlled and Economizer Ventilation Strategies on Energy Use in Buildings. Technical report, Univ. of Colorado, Boulder, CO (US), 1999.

[19] S. T. Bushby. BACnet$^{\text{TM}}$: A standard communication infrastructure for intelligent buildings. *Automation in Construction*, 1997.

[20] A. Carroll and G. Heiser. An analysis of power consumption in a smartphone. In *Proc. of USENIX ATC*, 2010.

[21] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan. Indoor localization without the pain. In *Proc. of ACM MobiCom*, 2010.

[22] J. Chung, M. Donahoe, C. Schmandt, I.-J. Kim, P. Razavai, and M. Wiseman. Indoor location sensing using geo-magnetism. In *Proc. of ACM MobiSys*, 2011.

[23] D. B. Crawley, L. K. Lawrie, F. C. Winkelmann, W. F. Buhl, Y. J. Huang, C. O. Pedersen, R. K. Strand, R. J. Liesen, D. E. Fisher, M. J. Witte, et al. EnergyPlus: Creating a New-generation Building Energy Simulation Program. *Energy and Buildings*, 2001.

[24] S. Dawson-Haggerty, X. Jiang, G. Tolle, J. Ortiz, and D. Culler. sMAP: A Simple Measurement and Actuation Profile for Physical Information. In *Proc. of ACM SenSys*, 2010.

[25] S. Dawson-Haggerty, A. Krioukov, J. Taneja, S. Karandikar, G. Fierro, N. Kitaev, and D. Culler. BOSS: Building Operating System Services. In *Proc. of USENIX NSDI*, 2013.

[26] S. Dawson-Haggerty, S. Lanzisera, J. Taneja, R. Brown, and D. Culler. @ scale: Insights from a Large, Long-Lived Appliance Energy WSN. In *Proc. of IEEE IPSN*, 2012.

[27] S. Emmerich, J. Mitchell, and W. Beckman. Demand-Controlled Ventilation in a Multi-zone Office Building. *Indoor and Built Environment*, 1994.

[28] V. Erickson, S. Achleitner, and A. Cerpa. POEM: Power-Efficient Occupancy-Based Energy Management System. In *Proc. of IEEE IPSN*, 2013.

[29] V. Erickson, M. Carreira-Perpiñán, and A. Cerpa. OBSERVE: Occupancy-Based System for Efficient Reduction of HVAC Energy. In *Proc. of IEEE IPSN*, 2011.

[30] W. J. Fisk and A. T. De Almeida. Sensor-Based Demand-Controlled Ventilation: A Review. *Energy and buildings*, 1998.

[31] W. J. Fisk, D. Faulkner, and D. Sullivan. A Pilot Study of the Accuracy of CO_2 Sensors in Commercial Buildings. *Lawrence Berkeley National Laboratory Paper LBNL E*, 2008.

[32] S. K. Ghai, L. V. Thanayankizil, D. P. Seetharam, and D. Chakraborty. Occupancy Detection in Commercial Buildings using Opportunistic Context Sources. In *In IEEE Percom Workshops*, 2012.

[33] S. Goyal, H. A. Ingley, and P. Barooah. Occupancy-based zone-climate control for energy-efficient buildings: Complexity vs. performance. *Applied Energy*, 2013.

[34] M. Hydeman. *Advanced Variable Air Volume: System Design Guide: Design Guidelines*. California Energy Commission, 2003.

[35] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and implementation of a high-fidelity ac metering network. In *Proc. of IEEE IPSN*, 2009.

[36] A. Krioukov, S. Dawson-Haggerty, L. Lee, O. Rehmane, and D. Culler. A Living Laboratory Study in Personalized Automated Lighting Controls. In *Proc. of ACM BuildSys*, 2011.

[37] P. Lazik and A. Rowe. Indoor pseudo-ranging of mobile devices using ultrasonic chirps. In *Proc. of ACM SenSys*, 2012.

[38] C. Martani, D. Lee, P. Robinson, R. Britter, and C. Ratti. ENERNET: Studying the Dynamic Relationship between Building Occupancy and Energy Consumption. *Energy and Buildings*, 2011.

[39] R. Melfi, B. Rosenblum, B. Nordman, and K. Christensen. Measuring Building Occupancy using Existing Network Infrastructure. In *Proc. of IEEE IGCC*, 2011.

[40] F. Oldewurtel, A. Parisio, C. N. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and K. Wirth. Energy Efficient Building Climate Control using Stochastic Model Predictive Control and Weather Predictions. In *Proc. of IEEE ACC*, 2010.

[41] S. N. Patel, M. S. Reynolds, and G. D. Abowd. Detecting human movement by differential air pressure sensing in hvac system ductwork: An exploration in infrastructure mediated sensing. In *Pervasive Computing*. 2008.

[42] S. N. Patel, K. N. Truong, and G. D. Abowd. Powerline positioning: A practical sub-room-level indoor location system for domestic use. In *In Proc. of UbiComp*. 2006.

[43] A. Rahmati and L. Zhong. Context-for-wireless: context-sensitive energy-efficient wireless data transfer. In *Proc. of ACM MobiSys*, 2007.

[44] A. Rowe, M. Berges, G. Bhatia, E. Goldman, R. Rajkumar, J. Garrett, J. Moura, and L. Soibelman. Sensor Andrew: Large-Scale Campus-Wide Sensing and Actuation. *IBM Journal of Research and Development*, 2011.

[45] J. Taneja, A. Krioukov, S. Dawson-Haggerty, and D. E. Culler. Enabling Advanced Environmental Conditioning with a Building Application Stack. 2013.

[46] L. Thanayankizil, S. Ghai, D. Chakraborty, and D. Seetharam. Softgreen: Towards Energy Management of Green Office Buildings with Soft Sensors. In *In Proc. of IEEE COMSNETS*, 2012.

[47] D. Turner, S. Savage, and A. C. Snoeren. On the Empirical Performance of Self-Calibrating Wifi Location Systems. In *Prof. of IEEE LCN*, 2011.

[48] US Department of Energy. Buildings Energy Data Book. `http://buildingsdatabook.eren.doe.gov/`, Aug. 2012.

[49] H. Wang, S. Sen, A. Elgohary, M. Farid, M. Youssef, and R. R. Choudhury. No need to war-drive: Unsupervised indoor localization. In *Proc. of ACM MobiSys*, 2012.

[50] M. Youssef and A. Agrawala. The horus wlan location determination system. In *Proc. of ACM MobiSys*, 2005.

[51] L. Zhang, B. Tiwana, Z. Qian, Z. Wang, R. P. Dick, Z. M. Mao, and L. Yang. Accurate online power estimation and automatic battery behavior based power model generation for smartphones. In *Proc. of ACM CODES+ISSS*, 2010.