

# Measuring Building Occupancy Using Existing Network Infrastructure

Ryan Melfi<sup>1</sup>, Ben Rosenblum<sup>2</sup>, Bruce Nordman<sup>2</sup>, and Ken Christensen<sup>1</sup>

<sup>1</sup>University of South Florida, Tampa, Florida, USA, rmelfi@mail.usf.edu, christen@csee.usf.edu,

<sup>2</sup>Lawrence Berkeley National Laboratory, Berkeley, California, USA, brosenblum@lbl.gov, bnordman@lbl.gov

**Abstract**—The primary focus of Green IT has been on reducing energy use of the IT infrastructure itself. Additional significant energy savings can be achieved by using the IT infrastructure to enable energy savings in both the IT and non-IT infrastructure. Our premise is that energy can be saved by driving building operation on information gleaned from existing IT infrastructure already installed for non-energy purposes. We call our idea *implicit occupancy sensing* where existing IT infrastructure can be used to replace and/or supplement traditional dedicated sensors to determine building occupancy. Our implicit sensing methods are largely based on monitoring MAC and IP addresses in routers and wireless access points, and then correlating these addresses to the occupancy of a building, zone, and/or room. Occupancy data can be used to control lighting, HVAC, and other building functions to improve building functionality and reduce energy use. We experimentally evaluate the feasibility of this dual-use of IT infrastructure and assess the accuracy of implicit sensing. Our findings, based on data collected from two facilities, show that there is significant promise in implicit sensing using the existing IT infrastructure present in most modern non-residential buildings.

**Keywords**—green IT; green buildings; occupancy detection; network infrastructure; implicit sensing

## I. INTRODUCTION

Occupancy detection is used in many buildings to reduce energy consumption [15]. Occupancy information can be used directly by building control systems to reduce the energy consumption of air conditioning, lighting, IT infrastructure, and other building systems [1, 7]. Occupancy detection can provide information to these building systems to allow them to operate proportionally to the number of occupants in the building. In [8] it is shown that for lighting there are potential energy savings of 50% with good occupancy detection. For HVAC, [6] shows that occupancy driven control could enable energy savings of 20%. Significant energy savings are possible by controlling desktop PC power on and sleep states as a function of usage. This is evidenced by a growing number of commercial products (for example, [18]) and research efforts in this area (for example, [17]). At the heart of these methods is occupancy detection of PC users.

Occupancy detection in current buildings is typically accomplished using passive infrared (PIR) and ultrasonic motion detectors. These dedicated occupancy sensors are installed expressly to determine occupancy and control building systems to reduce energy. Drawbacks to such explicit occupancy sensing include high cost of installation and use, limited accuracy, incorrect installation, failure after

installation (often unrecognized), and lack of networking capabilities for data collection. In [19] it has been shown that energy savings from using PIR sensors to control lights can vary from 10% to 45% depending on the type of room and the detector settings. PIR sensors are particularly effective for controlling lighting in infrequently occupied closed spaces such as storage rooms, but are ineffective for more open layouts such as offices [4, 19]. The lack of communication among these sensors precludes the sharing of occupancy information with other systems and fusion of information from multiple sensors. Additionally, some of these sensors provide limited occupancy information. For example, existing sensors cannot provide detailed location, movement, or identity information that hampers their ability to improve building functionality and reduce energy use [4, 19]. This is one of the reasons why occupancy sensors usually control only lighting and not other building systems.

This paper explores an approach to occupancy detection using existing IT infrastructure common to many (if not most) building types. This approach, first proposed in 2008 by Bruce Nordman and Alan Meier from Lawrence Berkeley National Laboratory [14] (unpublished), is called *implicit occupancy sensing*. Instead of directly measuring physical occupancy, the occupancy is derived from sources not originally intended for occupancy sensing. By using existing infrastructure, we believe occupancy information can be gathered inexpensively and with reduced energy overhead. We also believe implicit sensing can provide additional information about the occupants such as activity and identity. The contributions of this paper are:

- A formalization of occupancy measurement
- The proposal and evaluation of several implicit sensing methods

In the remainder of this paper, we present a model for understanding what information correlates to occupancy and how occupancy sensing can be evaluated. We then describe several new methods that have potential for providing occupancy information. Next, we evaluate these methods through experiments performed in two different buildings. Lastly, we describe related work and conclude.

## II. DEFINING OCCUPANCY

We define occupancy along the dimensions of *resolution* and *accuracy*.

### A. Resolution

A traditional occupancy sensor such as a PIR sensor outputs a binary value indicating that the area it is monitoring has people in it or not. This simple notion of occupancy is not adequate to compare more sophisticated

sensors which provide additional occupancy related information such as the number of occupants. Such sensors may facilitate control strategies not possible with PIR sensors. For example, consider an area with multiple PCs used by multiple users. A PIR sensor will not differentiate among occupants and so could not control individual workstations. Occupancy sensing that identifies occupants could power up those occupants' specific workstations. The notion of occupancy measurement should include information about the 1) space, 2) occupants, and 3) time span.

The quality of information provided by different types of sensors varies widely and can be thought of as the resolution of the sensor (see Fig. 1). Generally, as measured resolution increases the space becomes smaller, the occupant becomes more defined, and the information is available more quickly. For instance, a low resolution sensor might indicate that a building was occupied by one or more unidentified people in the last hour. A high resolution sensor might indicate that a specific room was occupied by three identified people in the last minute.

The spatial resolution of occupancy is easily defined in terms of building structures (that is, floors and rooms) which are referred to as zones in this paper. On the other hand, occupant resolution is more ambiguous. We define four levels of occupant resolution:

- *Occupancy* – a zone has at least one person in it
- *Count* – how many people are in a zone
- *Identity* – who they are
- *Activity* – what they are doing

In addition to spatial and occupant resolution, there is temporal resolution. Temporal resolution refers to the smallest time span in which changes in spatial and occupant resolution can be reported by a given sensor.

### B. Accuracy

Occupancy accuracy or error indicates how far from the ground truth any measurement is. Given a number of readings from a sensor and the ground truth occupancy, accuracy is the relative error of the measured value from the ground truth.

High levels of accuracy can be important so that control strategies can make correct decisions. The inaccuracies in PIR sensors are apparent when lights are turned off when an occupant is present or turned on when there are no occupants. If the consequences of an incorrect decision are acceptable (as is the case with PIR controlled lighting) then low accuracy is also acceptable. Sensor accuracy affects what types of control strategies are implemented using the sensor. The inaccuracies of a PIR sensor therefore make it acceptable for use in lighting, but not for controlling a PC.

While some applications require high accuracy others, such as HVAC, often do not. For example, ventilation rates can be keyed to whether no one is present, or to a rough estimate of the number of people around. Commonly no dynamic information is available – only a fixed schedule is used (typically with only maximum occupancy assumed). For lighting or HVAC, the start of the day mode can be keyed to when enough people actually show up on a




Figure 1. Occupancy resolution

particular day (and the same for initiating night mode) rather than being keyed to a particular time. For HVAC, temperature set points could be keyed to the number of people present, so that when fewer people are present larger deviations from the “ideal” set points are allowed. One feature of implicit occupancy sensing is that holidays and other low-occupancy days are automatically detected; there is no need to rely on a person to input such information on an ongoing basis. In many of these cases, the relevant comparison is not widely distributed traditional sensors – the practical alternative to implicit occupancy sensing is no occupancy sensing at all.

### III. IMPLICIT OCCUPANCY SENSORS

Implicit occupancy sensing is the use of existing building infrastructure (that is not originally intended for occupancy detection) to measure occupancy. Implicit sensing relies on the notion that the effects occupants have on building systems can be used to determine occupancy information. We refer to such effects as implicit measures of occupancy and the building systems as implicit occupancy sensors. Some building systems can be modified to add implicit occupancy sensors. We use the level of modification to classify the implicit occupancy sensors into three tiers:

- *Tier I* requires no modification to existing systems other than a collection and processing point.
- *Tier II* involves the addition of software to existing infrastructure to make existing occupancy related data available.
- *Tier III* involves the addition of software and hardware to introduce new sources of occupancy data to existing systems.

We have identified several implicit occupancy measures and observed how these measures relate to occupancy at two locations: the ENB building at the University of South Florida and Building 90 (B90) at Lawrence Berkeley National Laboratory (LBNL). ENB is the main engineering building at the USF Tampa campus and is a four-floor multi-purpose building that houses single occupant offices, conference rooms, classrooms, large communal areas, and research labs. B90 is a four-floor multi-purpose office




Figure 2. DHCP, ARP, and electricity use B90

building on the LBNL campus that includes private offices, workstations in communal areas, and conference rooms.

#### A. Tier I

The Address Resolution Protocol (ARP) is used by routers to determine the Ethernet or Wi-Fi Medium Access Control (MAC) address of all hosts on a subnet and associate it with the host's IP address. Routers use the ARP protocol to maintain a table, or ARP cache, of associated IP and MAC addresses to enable forwarding of packets from a router to a host. The ARP table is a measure of the number of active hosts in the local network. ARP data collected from routers and Wi-Fi Access Points (AP) within a building can serve as implicit occupancy measures.

Dynamic Host Control Protocol (DHCP) can be used as a similar measure as ARP of the number of hosts on a network. DHCP is used by non-permanent hosts to lease an IP address for a temporary period of time. A DHCP server has a pool of IP addresses and maintains a table of which addresses are currently leased-out. Thus, the DHCP lease data is a measure of the number of active hosts on the network not including the hosts with fixed IP addresses, typically servers). In wireless networks the majority of (if not all) hosts use DHCP to obtain an IP address. Additional measures related to DHCP include the number of new hosts, the number of hosts which hop between wireless access points (APs), and the number of hosts whose leases have expired. These DHCP measures may relate to occupancy measures of people entering the building, moving within the building, and leaving the building within the last DHCP timeout period.

For both the ARP table and DHCP lease table, hosts that have powered-down or otherwise left the network are detected and deleted. For DHCP the mechanism is a timeout set by the network administrator. If a host does not renew its lease within the timeout period that host is assumed to be gone and its IP address is returned to the IP pool. ARP cache entries are removed after there have been no packets sent to an address for an administrator defined ARP timeout period.

The number of hosts in a building as indicated by ARP and DHCP tend to follow trends similar to the number of building occupants (see Fig. 2 and Fig. 3). Fig. 2 shows variation in the total number of hosts as indicated by ARP and DHCP records from B90 routers and wireless APs and building electricity use as a function of time of day. The X-axis of the graph shows time beginning on Monday and the primary Y-axis shows the number of unique MAC addresses seen within 2 hour intervals for ARP and 10 minutes for DHCP. The secondary Y axis shows the electricity use of B90. The changes in ARP, DHCP, and energy are similar to those one would expect from occupancy in a laboratory or office; increase in the morning, stabilization around lunch, a decrease towards evening, and lower occupancy on the weekend than during week days (2/6/11 and 2/7/11).

Fig. 3 shows the number of measured active Wi-Fi hosts and the building energy use for ENB as measured by the electricity meter for ENB from December 17, 2010 to January 21, 2011. The figure shows the number of unique Wi-Fi hosts as a function of day on the left Y-axis graph and the actual measured electricity use of ENB on the right Y-axis. The correlation between these values can be seen in the scatter graph of Fig. 4 and is quantified by the Pearson coefficient, which measures linear correlation between two data sets with  $\pm 1$  being a perfect linear correlation. For the data in Fig. 3 the Pearson coefficient is 0.59 indicating a partial positive linear correlation between the energy and Wi-Fi hosts (that is, as the number of Wi-Fi hosts increases so does energy use). A similar data set was collected for the week of September 7, 2010 to September 13, 2010 (a non-holiday period) where the Pearson coefficient was 0.89.

The energy saving control strategies for ENB are currently limited. The HVAC is running 24/7 as is the lighting in communal areas, including hallways. This suggests that the increase in energy consumption during the day must be dependent on environmental changes and




Figure 3. DHCP and electricity use ENB

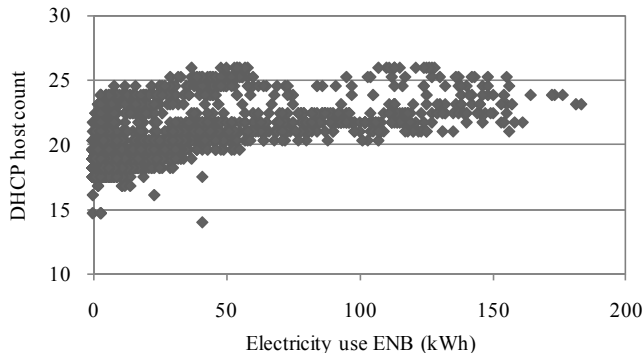


Figure 4. DHCP vs electricity ENB

occupancy, and not regulation of HVAC or lighting. Since the changes in energy are at least partially a function of occupancy, the correlation between Wi-Fi data and energy indicates that Wi-Fi host count is an indicator of building occupancy.

Other Tier I measures are telephone calls, security door sensors being triggered, and badge in/out systems. Fig. 5 shows the number of outgoing telephone calls from B90 over a period from mid December to late January. Each spike in the number of calls corresponds to a day. It can be seen that the trends in the number of telephone calls include increases during the weekdays and small increases during the weekends.

These measures have the ability to reveal detailed information about the occupancy level of a building. This is opposed to simpler occupancy schedules used in building control which only define periods when the building is expected to be occupied (for example, 8am to 6pm) and when it is unoccupied (for example, 6pm to 8am). It can be seen from the measures that the occupancy patterns in ENB have more complexity than a simple occupied/unoccupied model. In addition to normal occupancy patterns these measures can reveal occupancy patterns that are unusual due to special events or holidays. For instance, Martin Luther King Day is apparent in the DHCP and phone call data on January 17th. Also apparent is the end-of-year Winter holiday for both ENB and B90.

### B. Tier II

Wi-Fi can be used to locate occupants in a building, this ability has been explored in [12] where software on user hosts report all the available APs and their respective signal strengths. This information can be used to better triangulate occupants since all the nodes available to that occupant are known as opposed to only knowing which one was connected to through DHCP.

Another indication of occupancy is keyboard and mouse use. Operating systems provide functions to detect input activity but additional software is needed to collect this activity and make it available remotely. We have implemented a software method to gather this information called the PC Activity Monitor. The PC Activity Monitor calls the `GetLastInputInfo()` function from the Windows API which returns the time of the last keyboard or mouse activity. If there was activity within the last minute then PC

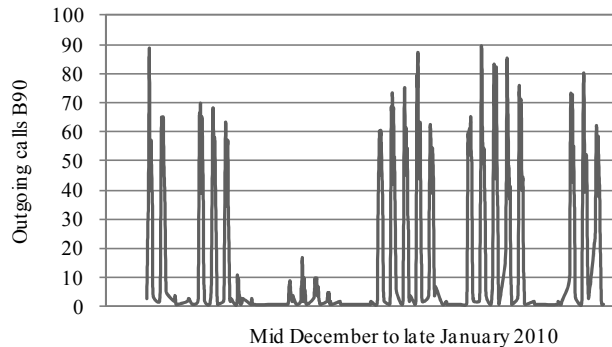


Figure 5. Outgoing phone calls B90

Activity Monitor logs 1 otherwise it logs 0. This information is made available remotely using the Simple Network Management Protocol (SNMP). SNMP provides a standard way to centrally gather the information recorded by the PC Activity Monitor on various computers. In [8] occupancy patterns were observed in several offices, and it was found that employees spend 76% of their time in their private office where their PCs are located. In this case, monitoring computer activity would be a good predictor of total building occupancy.

Other Tier II methods include the use of existing microphones and webcams to detect occupancy. Microphones on PC's, notebooks, and phones could provide occupant counts and identification based on sound levels or conversations. In [5] work has been done in relating acoustics to occupancy using dedicated sensors. The use of a webcam to detect occupants in front of a computer has already been demonstrated in [3].

### C. Tier III

Tier III methods require the addition of hardware and software to existing systems. The difference between this and dedicated sensors is that it makes use of the existing IT infrastructure to enable communication and processing. Furthermore, much of the additional hardware may in the future be standard for other purposes. Tier III methods could be used, for example, for measuring environmental conditions (temperature, etc.) as a measure of occupancy.

## IV. EXPERIMENTAL EVALUATION OF IMPLICIT SENSORS

We conducted an experimental evaluation of several implicit occupancy information sources in both ENB and B90. The experiments were intended to answer the following four questions:

- 1) What is the accuracy of implicit sensing as a measure of ground truth occupancy levels?
- 2) What is the accuracy of implicit sensing as a measure of individual occupant location?
- 3) Can implicit sensing measures reveal information about occupant activity?
- 4) How does the accuracy of implicit occupancy measures compare to that of explicit sensors?

The four experiments evaluated DHCP and occupancy counts, DHCP and occupant localization, DHCP and

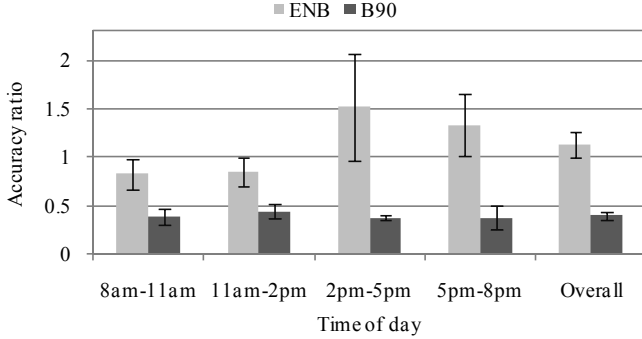


Figure 6. Accuracy of DHCP count

TABLE I. NUMBER OF SAMPLES FOR COUNT EXPERIMENT

Time of day	No. of samples ENB	No. of samples B90
8am to 11am	16	10
11am to 2pm	14	9
2pm to 5pm	5	11
5pm to 8pm	7	4
Overall	42	34

occupant activity, and comparison of explicit and implicit sensing methods.

#### A. Measurement Testbeds

Our experiments used the wireless network infrastructure in ENB and B90. ENB’s wireless network covers all four floors and consists of 44 APs. The AP names correspond to the floor and room in which the AP is found. B90 has three APs per floor except for the fourth floor which has one AP for a total of 10 APs.

A location in ENB called the Fishbowl and the 4th floor of B90 have been isolated for purposes of counting ground truth occupancy. The Fishbowl is furnished with a number of circular tables which are often used by groups of students to study and socialize. The area generally has 5 to 50 occupants. Wireless coverage for this area is redundant with around seven accessible APs. The fourth floor of B90 has 80 offices, a number of them private which surround an open office area in the middle. The floor generally has 20 to 60 occupants. There is only one AP on the fourth floor, but it is possible to connect to the third floor as well. For experiments in ENB the Wi-Fi device used was a BlackBerry mobile phone and in B90 the device used was an iPhone 4.

#### B. Count Experiment

For this experiment we measured the number of hosts on the seven APs which serve the Fishbowl area of ENB as indicated by DHCP leases. We also visually counted the number of occupants to establish the ground truth occupancy. We repeated this measurement on the fourth floor of B90 for the single AP. The procedure for this experiment was as follows:

- 1) At a given time visually record the number of occupants and count the number of hosts as indicated by DHCP leases.
- 2) Repeat step (1) approximately 10 times for time intervals 8am to 11am, 11am to 2pm, 2pm to 5pm, and 5pm to 8pm for four weeks.

#### C. Localization Experiment

For this experiment, a walking path through a building was selected where it is believed that Wi-Fi network data would show the occupant in all areas of the path. The walking path was about 15 minutes and included stops (of duration several minutes) in multiple rooms. The experiment procedure was as follows:

- 1) Walk the predetermined path with a Wi-Fi device on person. Log the actual locations and the AP locations as indicated by the devices DHCP leases by time.
- 2) Repeat step (1) at least three times for each path.

#### D. Activity Experiment

For this experiment we selected an area within a building and monitored the people who passed through it in order to measure the number of moving occupants. For ENB, four exits at the rear of ENB were selected. The procedure for the experiment was as follows:

- 1) Count and log the number of people that pass through the area for two minutes. If the area has defined directions of motions (that is, in and out) record these separately.
- 2) For the same two minutes count the number of new DHCP leases, the number of expired leases, and the number of AP transitions for the entire building.
- 3) Reset all counts.
- 4) Repeat steps (1) through (3) for one hour.

#### E. Comparison of Implicit and Explicit Sensors

To understand how implicit and explicit occupancy measures relate, a work area was selected which had Wi-Fi coverage, a desktop PC was instrumented with the PC Activity Monitor software (see Section III.B) installed and a PIR sensor was connected to the monitor of the PC facing the user’s seat. The sensor used was a USB PIR motion sensor, (part number CGUSBPIR) from CircuitGizmos.com. The sensor is a PIR module combined with a U421 USB interface which allows recording via a PC through the USB interface. The PIR sensor has a range of about 20 feet. The cost of the device was around \$50. The measurements from the PC Activity Monitor software, the PIR motion sensor, and the DHCP leases were recorded while the following procedure was executed four times for  $t = 5$  minutes, two times for  $t = 10$ , and one time for  $t = 15$ :

- 1) An occupant carrying a Wi-Fi enabled device approaches the instrumented desktop PC and logs the time he or she sat down. The Wi-Fi device should remain active and auto-connect to APs.
- 2) The occupant remains seated for  $t$  minutes while interacting with the PC.
- 3) The occupant logs the time then leaves the area for two minutes.

## V. RESULTS

The results for the count experiments at ENB are in Fig. 6, which shows the average error for DHCP as a measure of occupancy for several time periods during the day. It also shows a 95% confidence interval for each error as an indication of how much the error fluctuated within each time

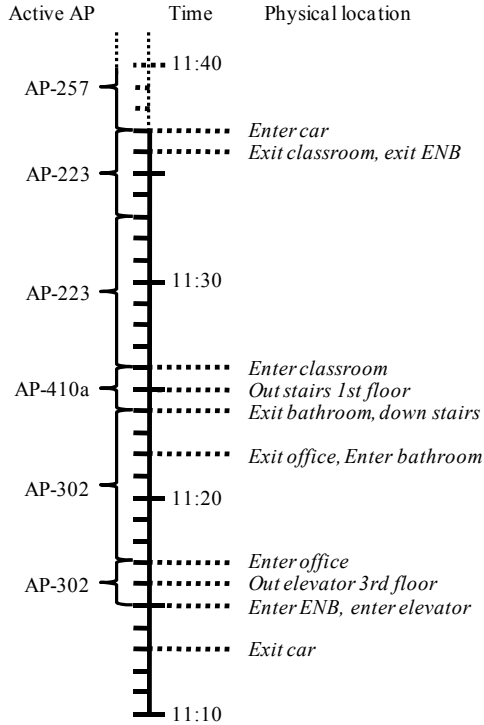


Figure 7. Walking path in USF ENB

TABLE II. LOCALIZATION ACCURACY OF DHCP SENSOR

Trial	ENB		B90	
	Building	Floor	Building	Floor
1	71%	31%	75%	41%
2	71%	42%	78%	50%
3	94%	84%	78%	44%

period. The number of samples for each time period and location can be seen in Table I. The overall accuracy and confidence intervals for all the samples at ENB and B90 were respectively  $1.13 \pm 0.13$  and  $0.40 \pm 0.036$  (this is roughly a  $\pm 10\%$  interval around the ground truth occupancy count).

The results from one trial of the localization experiment at ENB can be seen in Fig. 7. The leftmost column shows the time, location, and duration of DHCP leases and the rightmost shows the actual time and locations of the Wi-Fi device. The results from all trials and both locations can be seen in Table II. Here the accuracy is the percentage of time that the DHCP logs correctly localized the occupant in the building, the floor, and individual rooms. For example, in the first trial DHCP correctly detected the occupant in or not in the building 71% of the time.

The results from the activity experiment can be seen in Fig. 8. The top graph shows the number of new DHCP leases and the number of people entering the four doors for each time interval. The bottom graph shows the number of expired leases and the number of people exiting the building.

The sensor reading for one trial of this experiment at ENB can be seen in Fig 9. A reading of 1 for each sensor indicates occupancy and a 0 indicates a vacancy. The DHCP

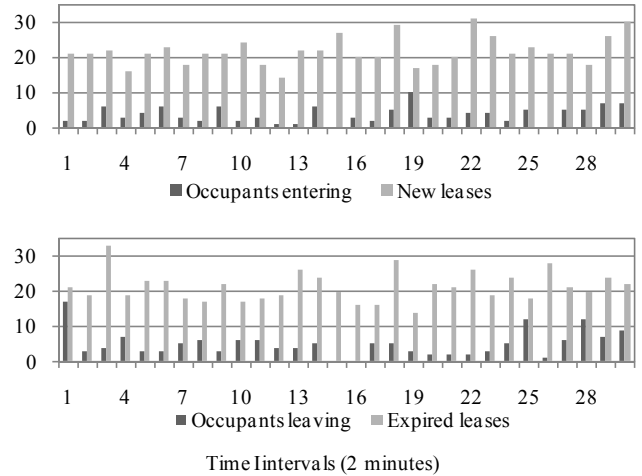


Figure 8. DHCP and occupants entering and leaving

TABLE III. ACCURACY OF IMPLICIT AND EXPLICIT SENSING

Device/ Accuracy	Vacancy	Occupancy	Overall
PIR	61%	100%	91%
PC Activity	100%	85%	89%
DHCP	61%	59%	59%

measure represents occupancy at a floor level resolution. The accuracy for detecting occupancy, vacancy, and total accuracy can be seen in Table III. Occupancy accuracy here is the percentage of time the seat was occupied and the sensors detected occupancy. Vacancy accuracy is the percentage of the time the seat was unoccupied and the sensors did not detect occupancy. Overall accuracy is the weighted average of vacancy and occupancy accuracy based on the time spent in each state.

#### A. Discussion

There are two main issues which affect the accuracy of DHCP as an implicit occupancy measure in the count and localization experiments: 1) overlap of AP coverage, and 2) inconsistent Wi-Fi connectivity of mobile phones. Issue (1) results in devices connecting to APs that are not in the same area as the occupant. For instance, a device may connect to an AP on the second floor when the occupant is on the third floor. This occurred frequently in experiments involving occupant localization. Although accuracy in measuring occupancy within the building was high, floor and room level accuracy was greatly reduced due to this issue.

Issue (1) also creates difficulty in mapping a physical area to a set of APs because AP coverage is not bounded by physical spaces such as rooms or floors. This causes inaccuracies in the measurement of occupant levels for specific floors and rooms. This was apparent in the count experiments performed at ENB where the APs available in the Fishbowl covered an area much greater than the Fishbowl itself. This could account for the change in accuracy for night hours, as the additional area could have had occupancy patterns different from the Fishbowl. Given the data, it is likely that an area with high nighttime occupancy was counted by the DHCP measure. For the




Figure 9. Accuracy of explicit and implicit sensing

occupancy count experiment at B90 a single AP was selected whose coverage closely matched the fourth floor while other APs (though available to the fourth floor) were ignored. This may account for the low variability in the accuracy for occupancy count measurements at B90.

Issue (2) results in devices losing network presence, which causes false negatives. It also means that fine grain movement (room to room) cannot be tracked with DHCP due to delays in reconnecting to the wireless network after an occupant has moved to another area.

The DHCP and occupant activity experiment shows changes in the amount of occupants entering and leaving the building and the number of new and expired DHCP leases however are not well correlated. Some peaks of movement and leases seem to align while others do not. This may be because the actual occupant movement was measured for only a portion of the building while the leases were taken over the entire building. Also, a delay is likely in the DHCP measurement given that entering occupants do not immediately connect to the network and exiting occupants' leases may not expire until some period after the occupant has left. Additional experiments are needed to conclude on the ability of DHCP to measure occupant movement.

Results from the sensor comparison experiment indicate that the PIR sensor was most accurate in measuring overall occupancy. It was however prone to false positives which were likely due to other occupants entering the space. On the other hand the PC Activity Monitor software was prone to false negatives indicating that the space was in use while the PC mouse and keyboard were not. It is important to note that the implicit measures provided occupancy information including occupant activity and identity not generated by the PIR sensor. It also indicates that sensor redundancy may be necessary for a complete measurement of occupancy and more complex energy saving control strategies.

## VI. RELATED WORK

There have been several efforts to improve the accuracy of dedicated occupancy sensors including of PIR and IR people counters by allowing the sensors to communicate [4, 10]. Other researchers have tried to create a general occupancy model which is capable of incorporating occupancy information from multiple sources [13, 16]. Bayesian Modeling is used in [4] to improve sensor accuracy through sensor communication. A convex optimization algorithm named SUN is developed in [16] that incorporates prior information, real time sensor readings, and the network configuration to estimate occupancy. Markov models are used in [13] to improve the accuracy of a network of traditional sensors. Each of these papers tries to improve dedicated sensors through modeling and communication. Finally, in [11] the relationship of measured zone energy use as a function of IP traffic is explored. This work in [11] is the closest to our work of any related work we have found.

Other implicit sensors have been explored in [2, 3, 9, 12]. In [3] a program called Face Off which uses a PC's webcam in order to detect occupancy is developed. In [9] a Bluetooth device is used as a proximity sensor to turn on and off a user's computer. In [2, 12] a wireless card and Wi-Fi routers are used to calculate the position of the device. This localization method is more accurate than DHCP leases since it is active on the part of the wireless hosts and may resolve the two issues found in our experimentation (see Section V.A). These methods are all Tier II in that they require additional software to gather information from the devices.

There has also been some work in the area of Tier III sensing, which incorporates environmental sensors (additional hardware) to detect occupancy. In [5] a wireless sensor network (WSN) test-bed of air quality and CO<sub>2</sub> sensors in an office space was implemented. Using Hidden Markov models to process the information from the sensors they could count the number of occupants (ranging from 1 to 5) with an average accuracy of 75%. It was concluded that the factors most related to occupancy in this situation were acoustics and CO<sub>2</sub> levels. Data mining analysis on sensor information for a WSN containing temperature, light, and humidity sensors was performed in [20]. It was concluded that measurements from these sensors contain noise due to inaccuracies and imprecision in the sensors. It was also found that many of the sensors exhibited spatial and temporal correlation, which has implications for applying predictive models such as Bayesian Networks to these sensor networks. Since these methods require costly WSNs it is important to address whether these Tier III sensors can provide additional accuracy over the less expensive Tier I and II counterparts.

## VII. CONCLUSIONS

In this paper we have shown that low cost implicit occupancy sensors are already available within existing building infrastructure. We have evaluated the accuracy of several implicit sensors in terms of occupancy counts, occupant location, and occupant activity. We have also demonstrated the feasibility of implicit occupancy sensing.

Implicit sensing was shown to have several advantages over sensing with dedicated sensors. These advantages include:

- No additional cost in terms of hardware, installation, operation, or maintenance of sensors.
- Availability of sensor readings over existing IT networks.
- Providing better information (such as occupancy resolution) not available from dedicated sensors.

Drawbacks to implicit sensing include:

- The requirement of a conversion between implicit measures and occupancy, which is not yet well understood.
- The accuracy of implicit sensing is dependent upon this conversion factor.
- Effective control strategies using implicit methods have not yet been demonstrated.

We believe that we have demonstrated the promise of the concept of implicit occupancy sensing. While some applications require high accuracy occupancy information to be useful, others, particularly HVAC, can utilize fairly inaccurate data and still save energy because the alternative is non dynamic fixed building control schedules.

Future work should focus primarily on addressing these drawbacks. Better understanding the conversion between implicit occupancy measures and occupancy is crucial if accuracy better than approximately  $\pm 10\%$  is needed for building control (and it is not clear that better than  $\pm 10\%$  is needed). This conversion may exist in probabilistic models as in [4, 10, 13, 16]. Finally, energy saving control strategies making use of implicit occupancy sensing should be developed and evaluated.

#### ACKNOWLEDGMENT

This work was completed with the support of a Florida Energy Systems Consortium (FESC) grant titled, "Reducing the Energy Consumption of IT: A Focus on Data Centers and End User Devices." The authors thank the anonymous reviewers for their helpful comments that have improved the quality of this paper.

#### REFERENCES

- [1] California Energy Commission CEC, "Advanced lighting guidelines," CEC Publication 400-93-014, California Energy Commission, 1993.
- [2] K. Chintalapudi, A. Dhariwal, R. Govindan, and G. Sukhatme, "Ad-hoc localization using ranging and sectoring," In Proceedings of IEEE INFOCOM '04, Hong Kong, March 2004, pp. 2662-2672.
- [3] A. Dalton and C. Ellis, "Sensing user intention and context for energy management," Proceedings of HotOS IX: The 9th Workshop on Hot Topics in Operating Systems, May 2003, pp. 151-156.
- [4] R. Dodier, G. Henze, D. Tiller, and X. Guo, "Building occupancy detection through sensor belief networks," Energy and Buildings, vol. 38, issue 9, September 2006, pp. 1033-1043.
- [5] B. Dong, B. Andrews, K. Poh Lam, M. Hoynck, R. Zhang, Y. Chiou, and D. Benitez, "An information technology enabled sustainability test-bed (ITEST) for occupancy detection through an environmental sensing network," Energy and Buildings, vol. 42, issue 7, July 2010, pp. 1038-1046.
- [6] V. Erickson and A. Cerpa, "Occupancy based demand response HVAC control strategy," Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building (BuildSys 2010), pp 7-10, 2010.
- [7] V. Garg and N. Bansal, "Smart occupancy sensors to reduce energy consumption," Energy and Buildings, vol. 32, 2000, pp. 81-97.
- [8] R. Harle and A. Hopper, "The potential for location-aware power management," Proceedings of UbiComp, 2008.
- [9] F. C. Harris and V. Cahill, "Power management for stationary machines in a pervasive computing environment," Proceedings of the 38th Hawaii International Conference on System Sciences, 2005.
- [10] J. Hutchins, A. Ihler, and P. Smyth, "Modeling count data from multiple sensors: a building occupancy model," Computational Advances in Multi-Sensor Adaptive Processing, December 2007, pp. 241-244.
- [11] Y. Kim, R. Balani, H. Zhao, and M. Srinivastava, "Granger causality analysis on IP traffic and circuit-level energy monitoring," Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency In Buildings (BuildSys 2010), November 2010.
- [12] A. Kushki, K. Plataniotis, and A. Venetsanopoulos, "Kernel-based positioning in wireless local area networks," IEEE Transactions on Mobile Computing, vol. 6, no. 6, June 2007, pp. 689-705.
- [13] C. Liao and P. Barooah, "An integrated approach to occupancy modeling and estimation in commercial buildings," American Control Conference, July 2010, pp. 3130-3135.
- [14] B. Nordman and A. Meier, "Using existing networks For energy purposes," Internal Document, October 7, 2009.
- [15] "Measurement science roadmap for net-zero energy buildings workshop summary report", NIST Building and Fire Research Laboratory, March 2010.
- [16] S. Meyn, A. Surana, Y. Lin, S. Oggianu, S. Narayanan, and T. Frewen, "A sensor-utility-network method for estimation of occupancy distribution in buildings," 48th IEEE Conference on Decision and Control, December 2009.
- [17] J. Reich, M. Goraczko, A. Kansal, and J. Padhye, "Sleepless in Seattle No Longer," Proceedings of the USENIX Annual Technical Conference, June 2010.
- [18] "Verdiem::SURVEYOR," 2011. URL: <http://www.verdiem.com/surveyor.aspx>.
- [19] B. Von Neida, D. Maniccia, and A. Tweed, "An analysis of the energy and cost savings potential of occupancy sensors for commercial lighting systems", Journal of the Illuminating Engineering Society of North America, 2001, pp. 111-125.
- [20] S. Wu and D. Clements-Croome, "Understanding the indoor environment through mining sensory data – A case study," Energy and Buildings, vol. 39, Issue 11, November 2007, pp 1183-1191.