An experimental investigation of occupancy-based control of commercial building climate

Jonathan Brooks, Siddharth Goyal, Rahul Subramany, Yashen Lin, Timothy Middelkoop, Laura Arpan, Luca Carloni, Prabir Barooah

***Abstract*— We present results from a week-long experimental evaluation of a scalable control algorithm for a commercial building heating, ventilation, and air-conditioning (HVAC) system that was proposed in our earlier work. The experiments showed that the proposed controller resulted in 37% energy savings over baseline without sacrificing indoor climate. In contrast to prior work that reports energy savings without a careful measure of the effect on indoor climate, we verify that the controller indeed maintains indoor climate as well as the building's baseline controller from measurements of a host of environmental variables and analysis of before-after occupant survey results. A complete system required to retrofit existing buildings with the controller is presented, which includes a wireless sensor network and a software execution platform in addition to control computation. Results show that there is a large variation in energy savings from zone to zone, which indicates that estimating energy savings potential of novel HVAC control systems is not trivial even from experiments—something that prior work with a uniformly positive message has failed to point out.**

I. Introduction

Buildings consume nearly 40% of the energy in the United States, and a significant fraction of this energy is due to heating, ventilation, and air-conditioning (HVAC) systems [1]. Due to the expense in refitting buildings with more efficient HVAC systems [2], a great deal of work in the literature has been dedicated to developing more efficient control algorithms for HVAC systems that can be used in existing buildings.

A number of recent papers have propounded the use of Model Predictive Control (MPC) for building applications. Experimental evaluations of MPC on various types of buildings have been reported in [3]–[6] with uniformly positive results: that the controller has been able to reduce energy use while maintaining constraints on space temperatures. One of the bottlenecks in applying MPC to building control is the variation from building to building in terms of equipment and dynamics, which requires building-specific tuning of models used in MPC computations. Consequently, such controllers may not be suitable for scaling up to a large number of buildings.

It has been shown in our prior work that it is possible to significantly reduce the energy used in maintaining indoor climate through the use of rule-based control (RBC) algorithms that use real-time measurements of occupancy, which are simpler to apply than MPC [7]. It was concluded through an extensive simulation study in [7] that the rule-based Measured Occupancy-Based Setback (MOBS) controller performed almost the same as a much more complex MPC-based controller in terms of both energy consumption and indoor climate conditioning. An independent study obtained similar conclusions [8]. Experiments in a single zone of a building were found to be consistent with the simulation results [9].

In contrast to prior work on MPC-based HVAC control, the MOBS controller does not require any building-specific or zone-specific tuning, which makes it highly scalable for deployment. The principle bottleneck in deploying the controller is obtaining real-time occupancy measurements. Integrating the control computations with the equipment controllers of the building is another challenge as building equipment is typically controlled through proprietary software and building automation systems.

This paper reports results of scaling up the implementation of the MOBS controller in [9] to several zones of a commercial building. Occupancy measurements were obtained with the help of low-cost, wireless sensor nodes. A custom software was developed for control execution that works with any BACnet-compatible equipment.

The results of the experiments confirm that the predicted performance of the MOBS controller from simulations and experiments on one zone do indeed persist when the control system is scaled up to multiple zones. Average energy savings of 37% were achieved without any adverse effects on the indoor climate as measured by temperature, humidity, and CO_2 levels as well as occupant feedback obtained from surveys before and during the tests.

Among the prior work, [10]–[14] have also proposed and experimentally evaluated energy-efficient RBC algorithms that are scalable to a large number of buildings. However, these papers only report measurements of temperature to claim that indoor climate is unaffected, while in practice both thermal comfort and indoor air quality (IAQ) must be maintained by the controller. Even thermal comfort is not completely determined by temperature alone; humidity is another important factor. In addition, the MOBS controller ensures that the ASHRAE ventilation standards 62.1-2010 [15] are met, while that is not the case for the other RBCs reported in the literature with the exception of [13], [14].

Our experiments consisted of twelve zones, and while

JB, YL, and PB are with the MAE Dept. at the U. of Florida, S.G. is with PNNL, RS is with Lutron Elec., Inc, TM is with the ISSE Dept. at the U. of Missouri, LC is with the CS dept. at Columbia U. and LA is with the College of Comm. and Info. at the Florida State U. Corresponding author email: brooks666@ufl.edu. This work was supported by the National Science Foundation through Grants CPS-0931885 and ECCS-0925534.

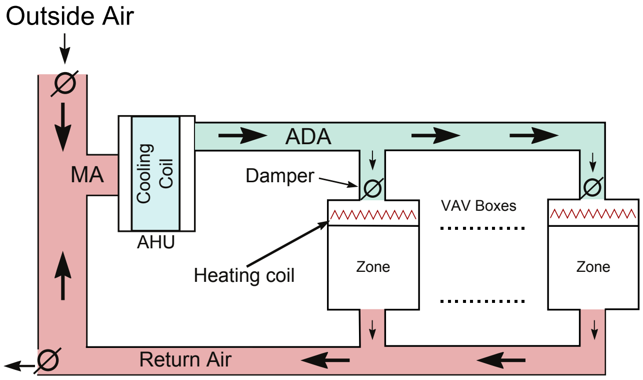

Fig. 1. A typical single-duct VAV HVAC system.

results were generally positive, both thermal comfort metrics and energy savings varied considerably among zones. Prior experimental work has been limited to a small number of zones so that variation among zones has not been noticed.

As a validation of the objective assessment of occupant comfort, we also include the results of surveys of occupants' comfort and perceptions of IAQ before and during the experiments for *subjective* assessment. While many works in the literature have implemented new control algorithms and analyzed their performances via varying degrees of energy and comfort analysis, to our knowledge none have methodically examined if the tested controller has led to any change (positive or negative) in the occupants' perceptions of thermal comfort and IAQ. The ultimate metric for comfort is occupants' opinions, and we provide these in our analysis.

This paper is organized as follows. Section II describes the building, HVAC system, control algorithms, and wireless sensor network (WSN) used for the experiments. We define our evaluation criteria in Section III. We analyze and discuss the results of the experiments in Section IV. Finally, Section V provides a summary of our results and avenues for future research.

II. System Architecture

Tests were carried out in Pugh Hall on the University of Florida campus, which is a LEED Silver certified building with a floor space of 40,000 sq. ft. and a variable air volume (VAV) HVAC system that has 3 air handling units (AHUs) and 65 VAV boxes. Each AHU conditions a mixture of outside and return air and then distributes the conditioned air to a number of VAV boxes through a single supply duct; see Figure 1. Each VAV box has an airflow damper and a reheat coil; it can modulate the flow rate of conditioned air delivered to its zone and re-heat the air. In the tests, twelve zones—each consisting of a single room—were controlled using the MOBS control algorithm described in Section II-B. The control commands at each VAV box were airflow damper position and re-heat valve position. Set points at the AHU level were *not* manipulated by the MOBS controller.

The control system architecture consists of the following components: (i) a wireless sensor network (WSN), (ii) a control algorithm for computing commands for the HVAC equipment, and (iii) a software infrastructure for data management and control execution. Figure 2 shows how the components of the system interact with one another. In the sequel, we describe each component starting with the third component.

A. Execution platform

The software system samples all of the "points" (sensor/actuator/set point values) from the building automation system (BAS) as well as from the wireless sensors and makes them available for real-time control and off-line analysis through a central database (RD_1 in Figure 2). Control commands from a control algorithm are communicated to the software by appending a row to a table in a relational database (RD_2). A scheduler, S in Figure 2, which checks for such updates every second, communicates these new commands to the equipment controllers through BACnet [16]. During the tests, commands that are otherwise computed by the BAS's control logics are overwritten by the software by using a higher priority in BACnet. More details of the software are described in [17], and the data logging component of the system is available publicly in `https://github.com/mtim/BacLog`.

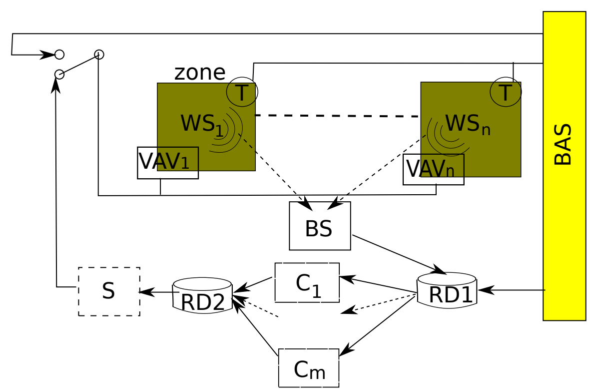

Fig. 2. Control system architecture.

This design, which mediates control computation and execution through databases but otherwise keeps the two separate, enables the use of arbitrarily complex control algorithms without bogging down the execution platform. In fact, receding horizon optimal control of a zone was performed by using the same software system, and the results are reported in [9].

Since the MOBS control algorithm is fully decentralized, information only from zone i is required to compute control actions for that zone. Control computation was performed in an off-site computer using MATLAB©. Commands to the controlled VAV boxes were updated every five minutes.

B. Control algorithm

The controller currently used in Pugh Hall (the baseline controller) is very close to a dual-maximum control scheme [18], though the exact nature of the controller is unknown due to its proprietary nature. Detailed descriptions of the dual-maximum and the MOBS control logic are available in [7]; here we provide a brief sketch for the sake of completeness.

The dual-maximum control logic has four modes of operation of a zone based on the measured zone temperature: (i) re-heating, (ii) heating, (iii) dead-band, and (iv) cooling. Each mode is activated if the room temperature remains within a certain temperature band for more than ten minutes: (i) re-heating: below the re-heating set point (T_{RTG}); (ii) heating: between the re-heating and the heating set point (T_{HTG}); (iii) dead-band: between the heating set point and the cooling set point (T_{CLG}); (iv) cooling: above the cooling set point. In re-heating mode, the temperature of the air supplied to the zone is set to its maximum value (T_{high}), and supply air flow rate is varied by a PID controller. In heating mode, supply air flow rate is set to its minimum value (V_{min}) while supply air temperature is varied via a PID controller. In dead-band mode, supply air flow rate is set to its minimum value, and no re-heating is performed (i.e., the re-heat valve remains closed). In cooling mode, supply air flow rate is varied using a PID controller, and no re-heating is performed.

The baseline controller in Pugh Hall also employs a nighttime setback in which temperature set points are relaxed during the night when the building is presumed to be unoccupied (10:30 PM-6:30 AM).

The MOBS controller, which was proposed in [7], is similar to that of the dual-maximum controller. Instead of pre-specified, building-wide times for unoccupied mode, real-time occupancy measurements are used to determine unoccupied times for each room individually. When a zone has been unoccupied for more than five minutes, it switches to the unoccupied mode; otherwise, it remains in the occupied mode. The heating/cooling/re-heating set points used in the MOBS controller were chosen to be the same as those used by the baseline controller (Table I).

Apart from the zone temperature bounds, the MOBS controller determines the minimum air flow rate and minimum outside air flow rate based on whether the zone is in occupied or unoccupied mode in accordance with ASHRAE ventilation standards [15]. The air flow rates are computed based on the measured occupancy count—i.e., the number of occupants.

TABLE I
CONTROLLER SET POINTS (° F)

T_{set}	T^{unocc}_{RTG}	T^{occ}_{RTG}	T^{unocc}_{HTG}	T^{occ}_{HTG}	T^{unocc}_{CLG}	T^{occ}_{CLG}
72.0	69.7	71.2	70.0	71.5	74.5	76.0

C. Wireless sensor network

The MOBS controller requires real-time measurements of occupancy count in each zone. Even though not used in control computation, measurements of humidity and CO_2 concentration are desired for post facto analysis to determine a controller's effect on thermal comfort and IAQ. These sensors were not available in the building.

A wireless sensor network was therefore deployed to obtain these measurements. Each wireless sensor node (shown in Figure 3) measures occupant presence (through a PIR sensor), CO_2 concentration, temperature, and relative humidity. Each node is equipped with a TI MSP430 microprocessor and a CC2500 radio operating in the 2.4 GHz band. Nodes use their radio transceivers to transmit their measurements to base stations throughout the building using TI's SimpliciTI communication protocol. The base stations were Dreamplug computers running Linux with their own radio transceivers. The base stations write the sensor data to a database through the Internet. For more information about the WSN and its design, the reader is referred to [19].

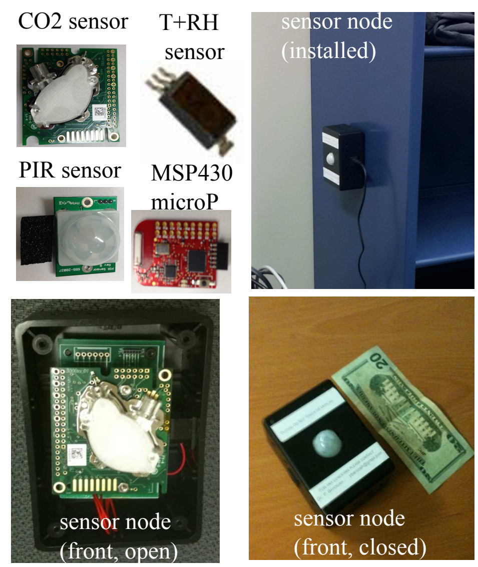

Fig. 3. Wireless sensor node deployed in Pugh Hall [19].

A PIR sensor only provides presence/absence measurements, but the MOBS controller requires occupancy count measurements. Occupancy count was estimated by assuming that whenever a room is occupied, it is occupied by its design occupancy. All of the rooms except the conference rooms had a design occupancy in the range of 2 to 4.

III. EVALUATION CRITERIA

Performance of a controller is measured in terms of how well it maintained indoor climate for occupants and how much energy it consumed per day.

A. Indoor climate

We evaluate indoor climate through thermal comfort and IAQ. Thermal comfort is determined by many factors; here we consider temperature and humidity. ASHRAE standard 55.1 specifies a range of temperature-humidity values in which a majority of building occupants are believed to be comfortable, which is shown in Figure 5. To perform an

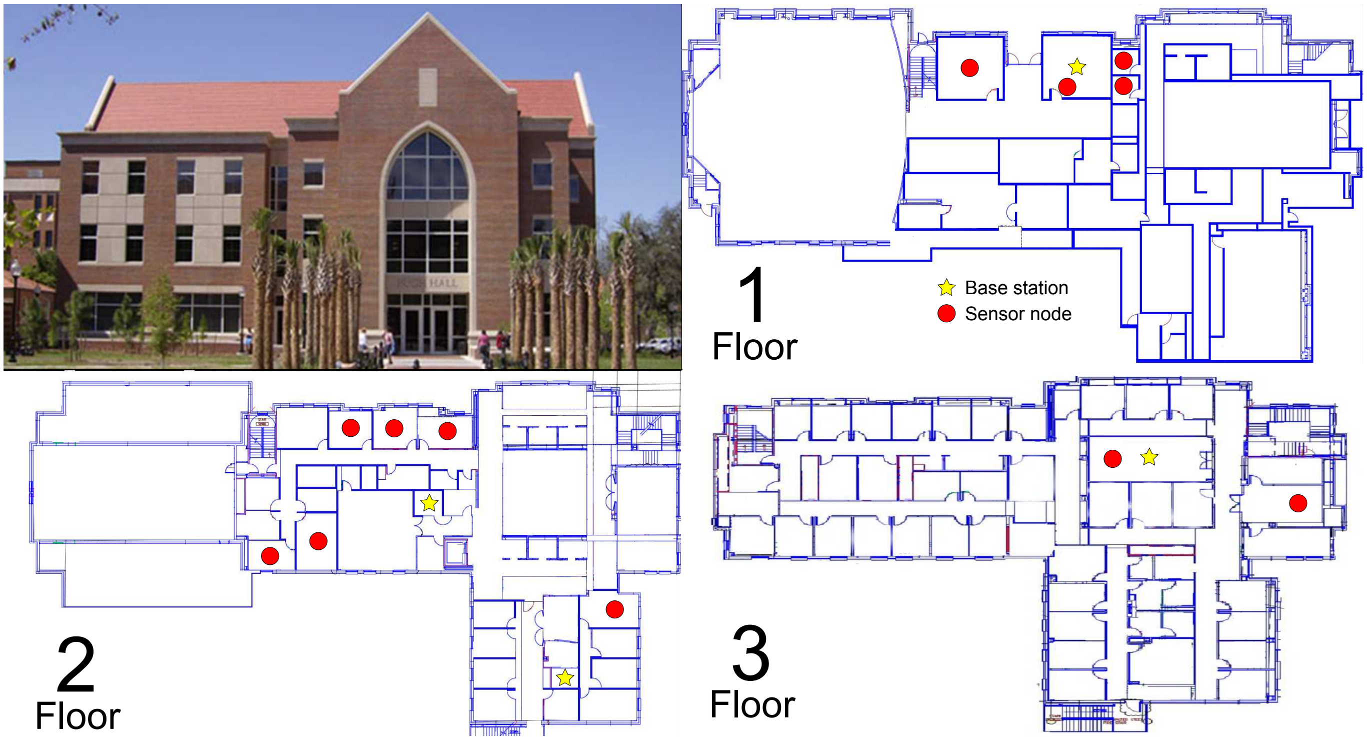

Fig. 4. Pugh Hall and the network structure on each floor. Base stations are represented by stars. Sensor nodes are represented by circles.

objective comparison of temperature constraint maintenance, we define the following measure, called *daily temperature deviation*, with a unit of °F·minutes:

$$\Delta T = \int_0^{24\times 60} d(T(t), [T_{\text{low}}(t),\ \ T_{\text{high}}(t)]dt \qquad (1)$$

where T_{low} and T_{high} are the lower and upper bounds on the space temperature, respectively, and $d(x, I)$ indicates distance between x and the interval I. A similar metric for humidity can also be defined, but we do not do so because in the tests humidity never exceeded the bounds, so the metric would turn out to be uniformly zero.

IAQ is trickier to evaluate as there is no universally accepted measure of IAQ. ASHRAE's ventilation standards are designed to ensure IAQ on average. We take CO_2 concentration as a rough measure. In particular, a CO_2 concentration below 1000 ppm is considered to indicate good outside air ventilation [18] and therefore satisfactory IAQ.

Physical variables do not capture all factors that affect thermal comfort. The ultimate measure of comfort is occupants' perceptions of comfort. To determine the controllers' effects—if any—on occupants' perceptions of comfort, web-based surveys were conducted during both baseline and test weeks that asked occupants to rate if their workspaces felt un/comfortable overall and stuffy/fresh.

B. Energy consumption estimation and comparison

There are three sources of power consumption: cooling power, heating power, and mechanical (fan) power. We estimate the cooling power consumed by each zone through an air-side specific enthalpy balance across the cooling coil at the AHU and the mass flow rate into the zone. The total cooling power consumed by the zones in which the MOBS controller was tested was then estimated by summing the power of each zone. A similar method was used to estimate the heating and the fan power consumptions.

Ideally, comparisons between two distinct controllers should be made by testing them under identical conditions, which is impossible to perform in practice since weather and occupant-related loads are never exactly the same for any two days. As an alternative, we compare each day of the week of experiments with the same day of another, baseline-controlled week with similar weather conditions.

For each day of the test, we must identify a day during which the baseline controller was operational and exogenous inputs were similar. To measure the degree of similarity, we define a metric, R, to compare the temperature and (volume-specific) enthalpy of the outside air of each day:

$$R = \int_0^{24\times 60} \Big[d\big(T^c(t), T^b(t)\big) + d\big(h^c(t), h^b(t)\big)\Big] dt \qquad (2)$$

where $T^b(t)$ and $h^b(t)$ are the temperature and enthalpy, respectively, at time t on the baseline day; similarly, $T^c(t)$ and $h^c(t)$ are the temperature and enthalpy at the same time t on the test day, and $d(x, y)$ denotes the absolute difference between x and y.

We perform a search over all days of 2013 and 2014, and the day with the lowest rating is used as the baseline day. We consider both enthalpy and temperature as evaluation criteria because the power consumption of the AHU is a function of the outside air enthalpy but the decisions of the baseline controller are based on the temperature. In the search, only the same days of the week are matched (e.g., only Mondays are compared to the Monday of the test week). This decreases the chance of vastly different occupancy schedules between the test and baseline days. Additionally, holidays are not considered.

IV. Results and Discussion

The MOBS controller was implemented for six days from 00:00 hours April 21, 2013 through 24:00 hours, April 26, 2013 in twelve zones of the building. The rooms/days in/during which the test was conducted are called the test rooms/days. The locations of the test rooms are shown in Figure 4.

A. Thermal comfort and IAQ

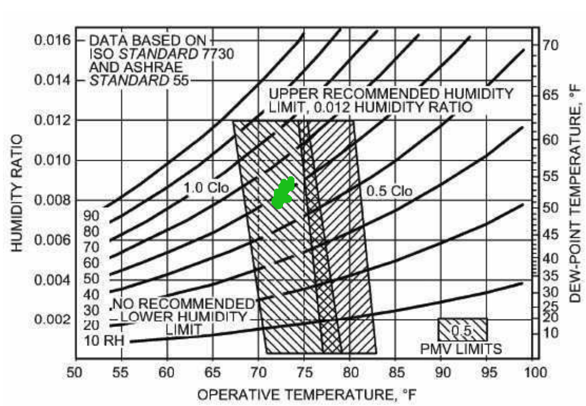

Fig. 5. Green dots: measurements of temperature and humidity in one of the test rooms during April 24, 2013. The dashed region is where most occupants are assumed to be comfortable [18].

All rooms' absolute humidities remained beneath 0.012 at all times, and the MOBS controller generally maintained temperature well. Figure 5 shows measurements of temperature and humidity in one of the test rooms for a test day. Note that these pairs remained within the comfortable region at all times during that day. For this room and day, the daily temperature deviation ΔT defined in (1) was 0. This was not the case for all rooms and days, see Figure 6 that shows a box plot of ΔT for each of the test rooms during the test period as well as the corresponding baseline days. As with the baseline controller, ΔT was small for most of the zones on the test days, but a few zones had large occasional violations.

The large daily temperature violations that occurred in rooms 3 and 4 did so due to actuator saturation: the maximum air flow rates allowed in those rooms were not enough to service the large thermal loads experienced by them. Rooms 3 and 4 were small study rooms and were subjected to large thermal loads due to students and their electronics during the test week as it was close to the final exams of the semester.

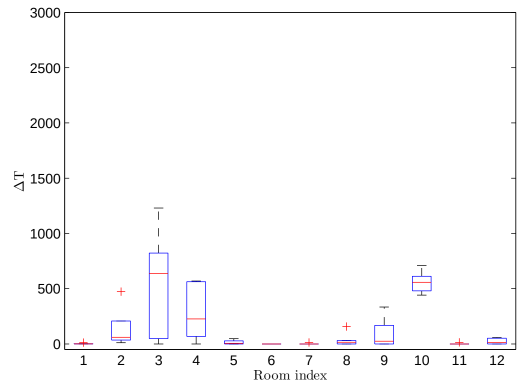

(a) Baseline

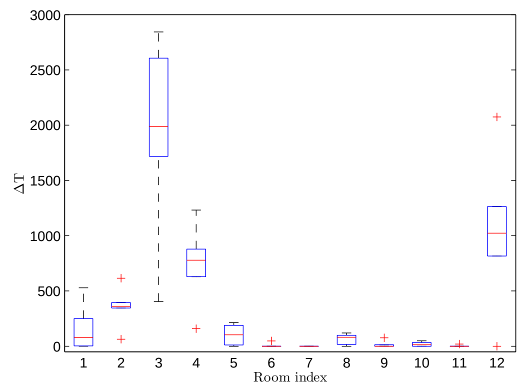

(b) MOBS

Fig. 6. Box plot of daily temperature violation of each room during the test (six samples per column).

Figure 7 shows the temperature, humidity, and CO_2 concentration measurements in room 11 during a test day (Wednesday, April 24, 2013) and the corresponding baseline day (Wednesday, May 15, 2013). Included in the figure are the zone temperature bounds for each controller. For the MOBS controller, these bounds are computed in real-time based on occupancy measurements provided by the wireless sensors. Notice that the bounds are stricter during the occupied periods than during the unoccupied ones. The baseline controller changes its bounds only twice during the day due to the nighttime setback. Note also the spike in CO_2 around 7 PM. Room 11 is a heavily used conference room, and the spike may have been caused by the presence of an extraordinarily large group (much larger than the design occupancy) just before. This was the only instance for any room in which the CO_2 exceeded 1000 ppm, so we consider the MOBS controller as having maintained IAQ well.

B. Energy savings

Figure 8 shows the total daily energy consumption of both the baseline and MOBS controllers for the test days. Figure 9

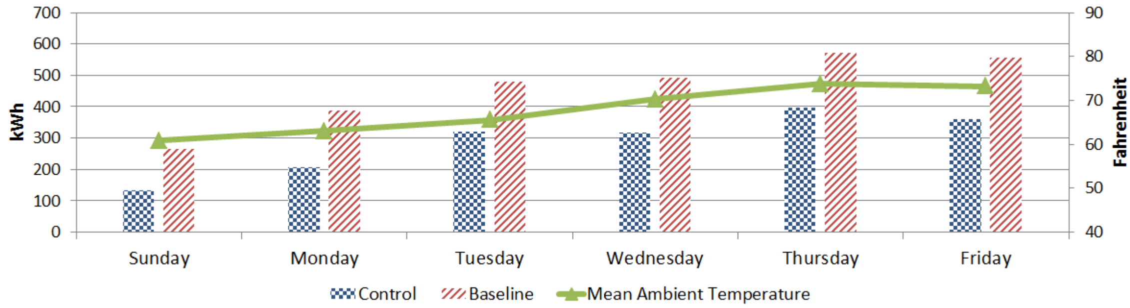

Fig. 8. Daily energy consumption of MOBS and baseline controllers and daily mean ambient temperature.

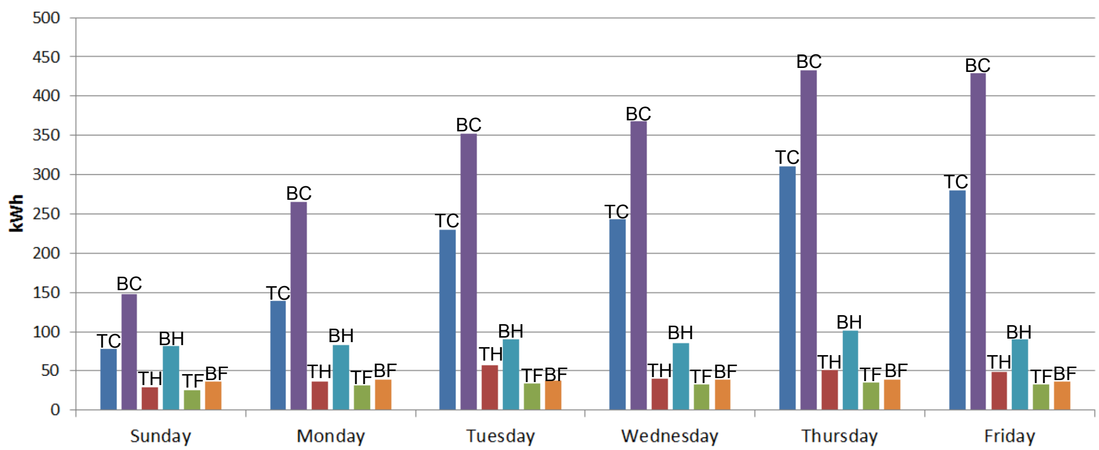

Fig. 9. Breakdown of daily energy consumption due to cooling (TC, BC), heating (TH, BH), and fan (TF, BF) for test (T) and baseline (B) days.

shows the breakdown of this energy consumption into each component—cooling, heating, and fan. The MOBS controller consistently reduced total daily energy consumption for each component compared to the baseline controller. For the entire week, the MOBS controller resulted in the consumption of 1.73 MWh, while the baseline controller consumed 2.75 MWh, indicating the MOBS controller reduced energy use by 37%.

The energy savings were comparable to the reduction in flow rate (32%). By using occupancy measurements rather than a fixed schedule, flow rate could be significantly reduced during the day. Figure 11 shows this reduction in the total flow rate for the test rooms.

While heating energy is only a fraction of the cooling energy, 30% of the energy savings came from heating. This is because the MOBS controller reduced overall heating energy consumption by more than 40%. A large amount of savings was possible because the baseline controller turns heating on and off rapidly throughout the day as a result of high flow rate. The reduced flow rate of the MOBS controller translates to less cooling supplied to the room. Because of this, the room temperature does not decrease as rapidly, and less heating is required then to maintain temperature. Most of the time during MOBS control, the heating valves remained closed—only opening infrequently for usually brief periods.

Just as with temperature regulation, there was a large variation in percent energy savings between rooms. Figure 10 is a box plot of each room's energy savings as percent of average baseline consumption for each room. As can be seen, the MOBS controller occasionally resulted in increased energy consumption over baseline, but positive percent savings were generally larger than negative percent savings. This is even more so when looking at nominal savings. The largest increase in nominal, daily energy consumption was less than 14 kWh, but the largest nominal decrease in daily energy consumption was greater than 45 kWh. The increase in energy consumption of the MOBS controller occurred in the rooms in which the baseline controller was supplying inadequate outside air, so that the MOBS controller had to increase airflow to meet ventilation standards.

C. *Occupants' perception of thermal comfort and IAQ*

Apart from variables such as temperature and humidity, the perceived thermal comfort of an occupant depends on a host of other variables such as air velocity, clothing, radiant heat, etc. [20]. Many—if not most—of these variables are nearly impossible to measure. Finally, there are subjective components to a person's perception of comfort as well.

To assess the impact of the MOBS control strategy on the occupants' perceptions of comfort and IAQ, web-based surveys were conducted. All occupants of the building were

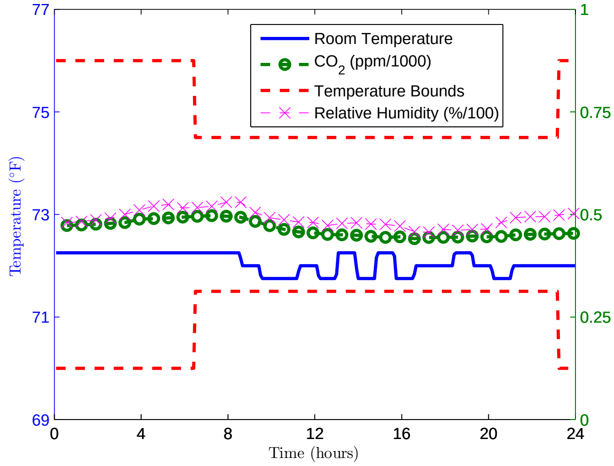

(a) Baseline, room 11

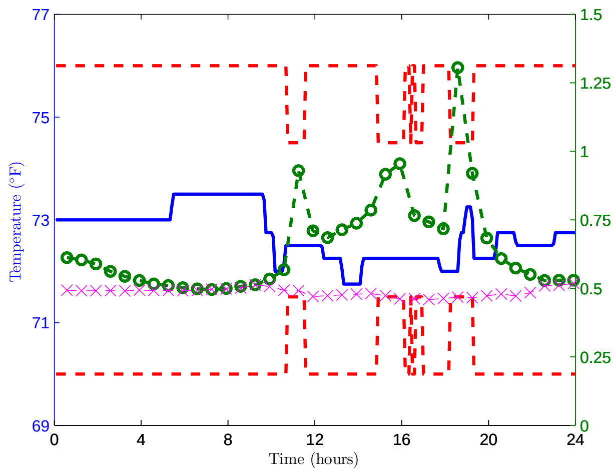

(b) MOBS, room 11

Fig. 7. Results of the MOBS controller on April 24, 2013, and of the baseline controller on May 15, 2013 (best matching baseline day for April 24), for a specific room. The dashed lines represent the upper and lower temperature bounds (left axis).

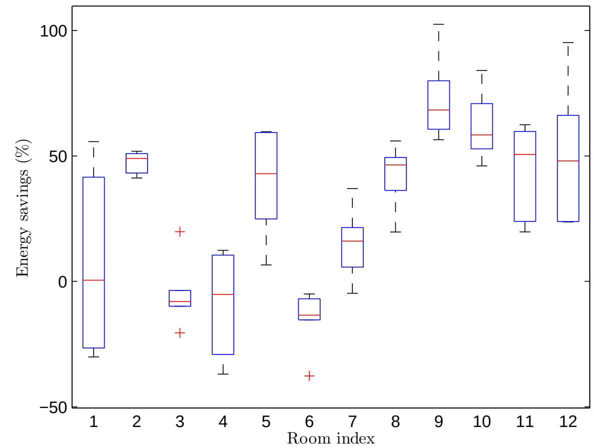

Fig. 10. Box plot of daily energy savings as percent of average baseline.

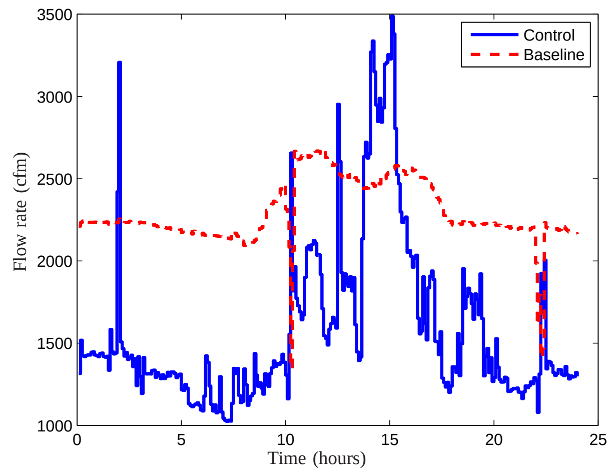

Fig. 11. Total flow rate of the 12 test zons during April 24, 2013 (test) and May 15, 2013 (baseline).

e-mailed a link to a web-based questionnaire that asked them to rate their overall comfort and air quality in the last five minutes, both on a numerical scale of 1 to 5, from very uncomfortable to very comfortable and very stuffy to very fresh.

Three waves of surveys were conducted: Wave 1 was during the installation of the wireless sensors in the building; Wave 2 occurred after the installation was complete but before the MOBS control tests were conducted; and Wave 3 occurred during the week when the control tests were in effect. Twenty-seven participants completed questionnaires from all three Waves. In general, the occupants responded positively to the MOBS controller, with more than 84% of participants rating the air as fresh (above the midpoint of the scale) and room conditions as comfortable (above the midpoint of the scale) in all three waves.

Paired samples *t*-tests indicated no significant differences in response to the item measuring overall air comfort (un/comfortable) in the workspace between the control test period, Wave 3 (M = 2.75, SD = 2.00), and the baseline period of Wave 1 (M = 2.71, SD = 1.92), t(23) = -.07, p = .94. Similar results were found for paired sample *t*-tests comparing perceived comfort between Wave 3 and Wave 1 (M = 2.95, SD = 1.47), with *t*(18) = .53, p = .60. In terms of the perceived air freshness measure (stuffy/fresh), it is interesting to note that participants perceived greater freshness during the control period (Wave 3), M = 2.82, SD = 1.59, than during the first baseline period (Wave 1), M = 1.64, SD = 1.65, *t*(21) = -3.78, $p < .010$. However, there were no significant differences in perceived air freshness during the Wave 2 baseline period, M = 2.70, SD = 1.59, and the time of the control test (Wave 3)M = 2.65, SD = 1.87, *t*(18) = -.53, p = .6. In short, the MOBS controller was not associated with decreases in reported occupant comfort and air freshness.

V. Conclusion and Future Work

The primary savings by the MOBS controller resulted from reduction of flow rate during unoccupied periods. The baseline controller uses a large flow rate throughout the day to ensure IAQ is maintained. Using occupancy measurements, the flow rate can be reduced during unoccupied times.

An important observation is that significant energy savings were achieved with PIR sensors that provide binary presence/absence measurements—not occupancy count. Accurately estimating the exact number of occupants in a room is still an open problem, but our results show that, for small office areas and even conference rooms, binary occupancy measurements can yield significant savings without the need for expensive hardware or complex estimation algorithms.

The overall cost of each sensor node was $215. The CO_2 sensor accounted for 30% of this cost. ASHRAE Standard 62.1 was designed to regulate CO_2 levels, so a CO_2 sensor is not explicitly necessary. Even with CO_2 sensors, the wireless network would cost approximately $14,000 to cover the entirety of Pugh Hall. The annual energy cost of Pugh Hall is approximately $100,000, so annual energy savings up to $40,000 may be expected. This means the WSN will have paid for itself in less than a year! It should be noted, however, that peak power consumption was increased on certain days. This increase may actually lead to an increase in the overall energy bill depending on the pricing scheme. Addressing and reducing this peak is the subject of future work.

Compared to prior work, we also showed that there is considerable variation from zone to zone—both in energy savings and indoor climate performance. This means predicting yearly energy savings from the controller for an entire building is not easy and requires further work.

It was shown in [21] that the selection of temperature set points has a large impact on energy consumption. Relaxing the temperature set points further during unoccupied times can lead to additional savings at the cost of initial comfort when the rooms become occupied again. An interesting avenue of future work in this direction is real-time control of the set points based on feedback from occupants perhaps provided through smart phones. Because preferences vary from person to person, one set point often does not fit all. Using occupant opinions as part of the feedback loop may reduce overall discomfort.

Acknowledgements

The authors would like to thank Dr. John Shea for help with the wireless sensor network, Dr. Herbert Ingley for his help and advice regarding HVAC systems, and UF administration & Physical Plants Department—in particular, Skip Rockwell, Peder Winkel, and John Lawson—for their help and support to the project.

References

[1] United States Energy Information Administration, "Annual energy review 2011," September 2012.

[2] M. Rahman, M. Rasul, and M. Khan, "Energy conservation measures in an institutional building in sub-tropical climate in australia," *Energy and Buildings*, vol. 87, no. 10, pp. 2994–3004, October 2010.

[3] S. Prívara, J. Široký, L. Ferkl, and J. Cigler, "Model predictive control of a building heating system: The first experience," *Energy and Buildings*, vol. 43, no. 2-3, pp. 564–572, February-March 2011.

[4] Y. Ma, F. Borrelli, B. Hencey, B. Coffey, S. Bengea, and P. Haves, "Model predictive control for the operation of building cooling systems," *IEEE Transactions on Control Systems Technology*, vol. 20, no. 3, pp. 796–803, May 2012.

[5] P. Ferreira, A. Ruano, S. Silva, and E. Conceição, "Neural networks based predictive control for thermal comfort and energy savings in public buildings," *Energy and Buildings*, vol. 55, no. 12, pp. 238–251, December 2012.

[6] S. Bengea, A. Kelman, F. Borrelli, R. Taylor, and S. Narayanan, "Model predictive control for mid-size commerical building HVAC: Implementation, results and energy savings," in *Second International Conference on Building Energy and Environment*, August 2012.

[7] S. Goyal, H. Ingley, and P. Barooah, "Occupancy-based zone climate control for energy efficient buildings: Complexity vs. performance," *Applied Energy*, vol. 106, pp. 209–221, June 2013.

[8] M. Gruber, A. Trüschel, and J.-O. Dalenbäck, "Model-based controllers for indoor climate control in office buildings–complexity and performance evaluation," *Energy and Buildings*, vol. 68, pp. 213–222, 2014.

[9] S. Goyal, P. Barooah, and T. Middelkoop, "Experimental study of occupancy-based control of HVAC zones," 2014, under review.

[10] Y. Agarwal, B. Balaji, S. Dutta, R. K. Gupta, and T. Weng, "Duty-cycling buildings aggressively: The next frontier in HVAC control," in *Information processing in sensor networks (IPSN)*, 2011, pp. 246–257.

[11] P. X. Gao and S. Keshav, "Spot: a smart personalized office thermal control system," in *Proceedings of the fourth international conference on Future energy systems*. ACM, 2013, pp. 237–246.

[12] K. Padmanabh, A. Malikarjuna V, S. Sen, S. P. Katru, A. Kumar, S. K. Vuppala, S. Paul, *et al.*, "isense: a wireless sensor network based conference room management system," in *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. ACM, 2009, pp. 37–42.

[13] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, and Y. Agarwal, "Sentinel: occupancy based hvac actuation using existing wifi infrastructure within commercial buildings," in *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2013, p. 17.

[14] V. L. Erickson, S. Achleitner, and A. E. Cerpa, "Poem: Power-efficient occupancy-based energy management system," in *Proceedings of the 12th international conference on Information processing in sensor networks*. ACM, 2013, pp. 203–216.

[15] American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., "ANSI/ASHRAE standard 62.1-2010, ventilation for acceptable air quality," 2010. [Online]. Available: www.ashrae.org

[16] C. Hubner, T. Hansemann, and H. Merz, *Building Automation*, ser. Signals and Communication Technology. Springer Berlin Heidelberg, 2009, pp. 185–273.

[17] T. Middelkoop, "High-resolution data collection for automated fault diagnostics," in *Automated Diagnostics*, B. L. Capehart and M. Brambley, Eds. Atlanta, GA: Fairmont Press, 2014, in press.

[18] ASHRAE, "The ASHRAE handbook fundamentals (SI Edition)," 2009.

[19] R. Subramany, "Wireless sensor network for HVAC control," Master's thesis, University of Florida, April 2013.

[20] P. O. Fanger *et al.*, "Thermal comfort. analysis and applications in environmental engineering." *Thermal comfort. Analysis and applications in environmental engineering.*, 1970.

[21] J. Brooks and P. Barooah, "Energy-efficient control of under-actuated HVAC zones in buildings," in *Proceedings of the American Control Conference*, 2014, accepted.