Energy-efficient control of under-actuated HVAC zones in commercial buildings

Jonathan Brooks[a,1,*], Saket Kumar[a], Siddharth Goyal[b], Rahul Subramany[c], Prabir Barooah[a]

[a] *University of Florida, MAE-B 327, 633 Gale Lemerand Drive, Gainesville, Florida 32611, United States*
[b] *Pacific Northwest National Laboratory, 902 Battelle Blvd, Richland, Washington 99354, United States*
[c] *Lutron Electronics, 101 N.W. 100th Avenue Plantation, Florida 33324, United States*

Abstract

An occupancy-based feedback control algorithm is proposed for variable air volume HVAC systems that is applicable to the "under-actuated" case in which multiple rooms share the same HVAC equipment. The proposed algorithm is scalable to buildings of arbitrary size without increase in complexity. Experimental results in five rooms show 29-80% energy savings potential. Despite the inability to condition rooms independently due to the shared HVAC equipment, comfort was found to be well maintained—even when one room was warmer and another was cooler.

Keywords: energy efficiency, thermal comfort, HVAC, variable air volume, occupancy-based control

1. Introduction

Buildings consume a significant portion of total energy use, and approximately half of this consumption comes from heating, ventilation, and air-conditioning (HVAC) systems [1]. A great deal of work has been done on HVAC

*Corresponding author
Email addresses: JonathanBrooksUF@gmail.com (Jonathan Brooks), saketkumar@ufl.edu (Saket Kumar), siddelec@gmail.com (Siddharth Goyal), rahulsubu@gmail.com (Rahul Subramany), pbarooah@ufl.edu (Prabir Barooah)
[1] Phone: 407-617-5105

control algorithms to make buildings more energy-efficient [2, 3, 4, 5]. Most conventional HVAC control systems do not use real-time measurements of occupancy. Rather, they maintain indoor climate regardless of whether occupants are present or not. Some systems use a "nighttime setback" logic in which set points are relaxed at night when the building is presumed to be unoccupied, but such time-based strategies can cause discomfort to occupants should they be present during the "nighttime" period. In addition, nighttime setback is unable to take advantage of absence of occupants during daytime. Significant energy savings—as well as increased comfort and indoor air quality (IAQ)—can be achieved by incorporating real-time occupancy information into the control loop [6, 7, 8].

Most of the works that have used real-time occupancy measurements for HVAC control, however, are restricted to scenarios where the climate of each occupied space can be controlled independently, which is often not the case. For indoor climate control, a building is typically divided into zones. Some of these zones consist of only a single room—allowing for independent control of room climate. We call such zones "fully actuated". In many systems, however, a zone is comprised of multiple rooms. We call such zones "under-actuated". In an under-actuated zone, each room's climate cannot be controlled independently because multiple rooms share one set of HVAC equipment. For instance, if a variable air volume (VAV) terminal box is used to supply air to two rooms, cooling one room will necessarily cool the other. This poses challenges when climate conditions are different in different rooms within the same zone (e.g., one room is hot and one room is cold).

Conventional HVAC systems with under-actuated zones compute control actions for a zone based on either average values of measurements from the rooms or measurements from a single room in the zone. Both of these can result in poor climate control.

Despite the fact that many HVAC systems contain under-actuated zones, work on energy-efficient control of under-actuated zones is noticeably lacking in the literature. While some works have experimentally tested occupancy-

based control schemes in under-actuated HVAC zones, such as [9, 10], it is not clear whether the algorithms provide any guarantee for under-actuated zones or whether the results indicate that the algorithms are effective for the under-actuated case. For instance, in [9], it is unclear how the proposed control algorithm deals with under-actuated zones as it is described on the basis of rooms rather than zones. Similarly, the controller proposed in [10] "...will turn-off (or put into stand-by) zones that are currently unoccupied", but it is not clear how difference in occupancy between the rooms of an under-actuated zone is handled. Finally, the existing work does not provide analysis of the experimental results to show how the controllers performed in the under-actuated zones.

In this work, we propose a feedback control algorithm, MOBS^{ua} (Measured Occupancy-Based Setback for under-actuated zones), for energy-efficient control of both fully actuated and under-actuated zones. The MOBS^{ua} algorithm is an extension of the MOBS algorithm described in [11]. The MOBS algorithm utilizes real-time occupancy measurements to increase the energy efficiency of a VAV HVAC system. Significant energy savings were found through simulation, and these results were verified through experiments in [12]. The MOBS algorithm was designed, however, for fully actuated zones and cannot be implemented "as-is" in under-actuated zones. The MOBS^{ua} algorithm is designed to improve both energy efficiency and thermal comfort of occupants in under-actuated zones over conventional control logics that do not use occupancy measurements; it also ensures that ventilation constraints mandated by the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) are met in all rooms of a zone even if it is under-actuated. It is important to note as well that the MOBS^{ua} algorithm is scalable to buildings of arbitrary size because control decisions are made independently for each zone.

The algorithm was implemented in three zones in a real commercial building on the University of Florida campus (Pugh Hall) and tested continuously for a week-long period. A wireless sensor network (WSN) was deployed in the building to provide real-time measurements of occupant presence, which were used for the control, and of CO_2, humidity, and temperature. Results of the

experiment indicate that thermal comfort was maintained well—even in the worst case scenario when one room is hot and the other is cold. In addition to improving comfort, the MOBS^{ua} algorithm led to significant energy savings over the baseline controller. The control system achieved the improvement in energy use and thermal comfort at low cost since the required measurements for the control computations were obtained with minimal additions to the building's existing infrastructure.

In addition to a novel algorithm and its experimental evaluation, a third contribution of this paper is a method to identify a "baseline" for comparing the performances of distinct controllers. While others have grappled with this issue and have developed ways to address it, the methods are not clearly described to be reproducible.

A preliminary version of this work was presented in [13] in which we compared predictive and feedback control algorithms to a common baseline algorithm via simulation of an under-actuated zone. Because the HVAC actuation considered in [13], specifically room heating, is different than the one in this work, the MOBS^{ua} algorithm described here is somewhat different from the one in [13]. The simulations reported in [13] indicated that additional benefits of the predictive control algorithm may not be enough to justify the increased cost of its deployment. This increase in deployment cost is due to the increase in computational complexity and the difficulty in obtaining occupancy predictions. Therefore we limit ourselves to feedback control in this paper both in algorithm development and experimental evaluation.

The rest of the paper is organized as follows. We describe both the baseline and MOBS^{ua} control algorithms in Section 2. In Section 3, we describe the evaluation metrics that we use to assess relative performance between the two control algorithms. The experimental setup is described in Section 4. We present the results of our experiments in Section 5. Finally, we conclude this work in Section 6.

2. Control Algorithms

The control algorithms described below are for a VAV HVAC system with terminal reheat. In VAV systems, a central air handling unit (AHU) conditions mixed air—consisting of a mixture of outside air and recirculated air from the zones—and then distributes the conditioned air to the zones that it serves; see Fig. 1. Each VAV box has an air damper to modulate airflow to the zone and a heating coil to reheat the air conditioned by the AHU. The actuation commands are the air flow rate through the VAV box and the reheat valve position; the latter decides the amount of reheating provided to the air passing through the box. If the zone is under-actuated, neither the airflow to the rooms nor the heating to them can be independently controlled; instead, they are determined by the relative size of the supply ducts from the VAV box to the rooms.

2.1. Baseline controller

The dual-maximum controller described in [14, chapter 47] is taken as a baseline controller. Even though single-maximum control [14] is more common in existing commercial buildings, dual-maximum is the more efficient of the two. In a dual-maximum scheme, the controller uses different operation modes when the room temperature is in different temperature bands; see Fig. 2. When the temperature is in the cooling band, no reheating is performed, and the flow rate is varied to return the temperature to the "dead band". In the dead band, no reheating is performed, and flow rate is set to its minimum value. When the temperature is in the heating band, flow rate is set to its minimum while heating is varied to return the temperature to the dead band. Finally, in the reheating band, heating is set to its maximum, and flow rate is varied. The baseline controller utilizes a nighttime setback in which each set point of the four temperature bands is relaxed (increasing the size of the dead band). Additionally, under-actuated zones are controlled using temperature measurements from only one room within the zone.

2.2. MOBS$^{\text{ua}}$

Fig. 3 shows a schematic of the MOBS$^{\text{ua}}$ control algorithm. The "control temperature", T_{control}, is the room temperature of the "active room". The active room is determined by Algorithm 1 from real-time occupancy and temperature measurements; see Algorithm 1 for the details of the decision logic. Roughly speaking, the priority in determining the active room is (Occupied,Uncomfortable) > (Unoccupied,Uncomfortable) > (Occupied,Comfortable) > (Unoccupied,Comfortable). The control temperature is then used as the zone temperature in the dual-maximum controller. If the zone happens to consist of only one room (i.e., it is fully actuated), that room is always the active room.

Occupancy measurements are also used to update the parameters T_{reheat}, T_{heat}, and T_{cool} in real time. Simultaneously, occupancy measurements are used to calculate the minimum supply airflow for each room according to ASHRAE Standard 62.1-2010 [15]. The minimum supply airflow for the entire zone is then computed so that the ventilation constraint for each room is met. To determine these minimum airflows, various parameters such as ratios of cross-sectional areas of ducts (duct distribution) to the rooms, outside-airflow ratio, room size, and room type are required. These are provided to the controller *a priori*.

Only presence/absence information about occupants is needed to execute Algorithm 1 and update the values of the parameters, but the entire MOBS$^{\text{ua}}$ algorithm needs measurements of occupancy count (number of people) of each room to compute minimum air flow rates. In our implementation, the following constant parameters for each room are required: (i) design occupancy, (ii) room size (floorspace), (iii) room type (e.g., office), and (iv) duct distribution (fraction of total zone air that is supplied to the room). We require the following real-time measurements for each room in our implementation: (i) occupied/unoccupied measurement, (ii) temperature, and (iii) outside-airflow ratio.

Algorithm 1: Under-Actuation Algorithm. MOBS$^{\text{ua}}$ logic for determining active room of a single zone consisting of n rooms.

$\mathrm{t} \leftarrow$ Current time

$\mathrm{T_i(t)} \leftarrow$ Room temperature of room i at time t, i=1,2,...,n

$\mathrm{Occ_i(t)} \in \{0,1\} \leftarrow$ Occupancy of room i at time t, i=1,2,...,n

$[\mathrm{T_{heat}(t)}, \mathrm{T_{cool}(t)}] \leftarrow$ Dead-band temperature interval at time t

$\mathrm{d(T_i(t), [T_{heat}(t), T_{cool}(t)])} \leftarrow$ Distance of temperature of room i at time t from the dead-band temperature interval at time t

$\mathrm{T_{control}(t)} \leftarrow$ Temperature used for control at time t

$\mathrm{T_{set}} \leftarrow$ Room temperature set point

foreach k minutes **do**

  if $\{\mathrm{i} \mid \mathrm{Occ_i(t)} = 1,\ \mathrm{T_i} \notin [\mathrm{T_{heat}(t)}, \mathrm{T_{cool}(t)}]\} \neq \phi$ **then**

    $$\mathrm{T_{control}(t)} \triangleq \arg \max_{\{\mathrm{i} \mid \mathrm{Occ_i(t)}=1, \mathrm{T_i(t)} \notin [\mathrm{T_{heat}(t)}, \mathrm{T_{cool}(t)}]\}} \mathrm{d(T_i(t), [T_{heat}(t), T_{cool}(t)])}$$

  else if $\{\mathrm{i} \mid \mathrm{T_i(t)} \notin [\mathrm{T_{heat}(t)}, \mathrm{T_{cool}(t)}]\} \neq \phi$ **then**

    $$\mathrm{T_{control}(t)} \triangleq \arg \max_{\{\mathrm{T_i(t)} \mid \mathrm{T_i(t)} \notin [\mathrm{T_{heat}(t)}, \mathrm{T_{cool}(t)}]\}} \mathrm{d(T_i(t), [T_{heat}(t), T_{cool}(t)])}$$

  else if $\{i \mid \mathrm{Occ_i(t)} = 1\} \neq \phi$ **then**

    $$\mathrm{T_{control}(t)} \triangleq \arg \max_{\{\mathrm{T_i(t)} \mid \mathrm{Occ_i(t)}=1\}} |\mathrm{T_i(t)} - \mathrm{T_{set}}|$$

  else

    $$\mathrm{T_{control}(t)} \triangleq \arg \max_{\mathrm{T_i(t)}} |\mathrm{T_i(t)} - \mathrm{T_{set}}|$$

end

3. Evaluation Metrics

3.1. Thermal comfort and IAQ

To examine how well a controller maintains thermal comfort, we define a daily temperature deviation metric, ΔT, with units of °C·minutes for each room shown by (1).

$$\Delta \mathrm{T} = 24 \times 60 \times \int_0^1 \mathrm{d}(\mathrm{T(t)},\ [\mathrm{T_{heat}(t)}, \mathrm{T_{cool}(t)}])\mathrm{dt} \tag{1}$$

In (1), d(x,I) denotes the distance between the number x and the interval I, T(t) is the room temperature at time t where t is in the unit of days, and $[\mathrm{T_{heat}(t)}, \mathrm{T_{cool}(t)}]$ is the dead-band temperature interval at time t. Note that this interval is a function of time because the size of the interval is increased during unoccupied times. The multiplication preceding the integral converts units from °C·days to °C·minutes. A similar metric is defined for absolute humidity:

$$\Delta \mathrm{W} = 24 \times 60 \times \int_0^1 \max\{\mathrm{W(t)} - 0.012,\ 0\}\mathrm{dt},$$

which is chosen because [16] recommends that absolute humidity should not exceed 0.012.

Quantifying IAQ is much more difficult because there is no universally accepted metric [17]. In this work, we consider CO_2 levels under 1000 ppm as an indicator of good outside-air ventilation [16] and thus good IAQ of a zone.

3.2. Power consumption estimation

To estimate cooling power consumption for each zone, we perform an air-side enthalpy balance across the AHU's cooling coils using the zone-level flow rates:

$$\mathrm{P_i^{cool}(t)} = \dot{\mathrm{m}}_\mathrm{i}\mathrm{(t)}(\mathrm{h^{MA}(t)} - \mathrm{h^{ADA}(t)}), \tag{2}$$

where $\dot{\mathrm{m}}_\mathrm{i}$(t) is the mass air flow rate being supplied to zone i at time t, $\mathrm{h^{MA}(t)}$ is the specific enthalpy of the unconditioned mixed air before the AHU cooling coils at time t, and $\mathrm{h^{ADA}(t)}$ is the specific enthalpy of the AHU discharge air after the cooling coils at time t; see Fig. 1 for superscript definitions.

Heating is estimated by a simple air-side enthalpy balance across each VAV box's reheat coil:

$$\mathrm{P}_{\mathrm{i}}^{\mathrm{heat}}(\mathrm{t}) = \dot{\mathrm{m}}_{\mathrm{i}}(\mathrm{t})(\mathrm{h}^{\mathrm{ADA}}(\mathrm{t}) - \mathrm{h}_{\mathrm{i}}^{\mathrm{ZDA}}(\mathrm{t})), \tag{3}$$

where $\mathrm{h}_{\mathrm{i}}^{\mathrm{ZDA}}(\mathrm{t})$ is the specific enthalpy of the zone discharge air for zone i (after the heating coil in the zone's VAV box) at time t.

For fan power, we simply use the fractional fan power based on flow rates as shown by (4) where $\dot{\mathrm{m}}_{\mathrm{AHU}}(\mathrm{t})$ is the total mass flow rate through the AHU at time t and $\mathrm{P}_{\mathrm{AHU}}^{\mathrm{fan}}(\mathrm{t})$ is the power consumption of the AHU's supply fan at time t:

$$\mathrm{P}_{\mathrm{i}}^{\mathrm{fan}}(\mathrm{t}) = \frac{\dot{\mathrm{m}}_{\mathrm{i}}(\mathrm{t})}{\dot{\mathrm{m}}_{\mathrm{AHU}}(\mathrm{t})}\mathrm{P}_{\mathrm{AHU}}^{\mathrm{fan}}(\mathrm{t}) \tag{4}$$

For each of the specific enthalpies in (2), (3), and (4), we use the standard formula, h = 1.006T + W(2501 + 1.86T), provided in [16] where T is the air temperature in degrees Celsius and W is the humidity ratio of the air.

3.3. Supply air humidity estimation

The sensor measuring the humidity of the AHU discharge air is not accurate at high relative humidities. It was found that the humidity sensor was often reporting higher AHU discharge-air humidity ratio than mixed-air humidity ratio, which is not possible. This error can cause significant error in air enthalpy calculation and thus in power consumption calculation. An affine humidity model was therefore identified for the AHU discharge air:

$$\mathrm{W}^{\mathrm{ADA}}(\mathrm{t}) = \mathrm{W}^{\mathrm{MA}}(\mathrm{t}) - \alpha \cdot \max\{\mathrm{T}^{\mathrm{MA}}(\mathrm{t}) - \mathrm{T}^{\mathrm{ADA}}(\mathrm{t}),\ 0\}$$

where the value of the coefficient α is $1.83{\times}10^{-4}$, which was obtained performing a least-squares fit between the predicted power consumption and the actual power consumption of an entire AHU using a year of historical data. The model ensures that the AHU discharge-air humidity ratio is equal to the mixed-air humidity ratio when the mixed air is colder—i.e., there is no condensation across the cooling coils when no cooling is delivered.

3.4. Comparison between controllers tested on distinct days

Comparison between HVAC controllers tested on different days is difficult because external conditions (e.g., occupancy, weather, etc.) can never be exactly the same. We define a distance function, $d_o(i,j)$ between two days i and j that measures how similar the external conditions were between them:

$$d_o(i,j) \triangleq \int_0^1 \Big[|T_i^{OA}(t) - T_j^{OA}(t)| + |h_i^{OA}(t) - h_j^{OA}(t)| \Big] dt; \tag{5}$$

where $T_i^{OA}(t)$ is the outside-air temperature at time t on day i, and $T_j^{OA}(t)$ is the outside-air temperature at time t on day j. Similarly, $h_i^{OA}(t)$ and $h_j^{OA}(t)$ are the (volume-specific) outside-air enthalpies on days i and j, respectively, at time t. In (5), $d_o(i,j)$ defines a metric on the search space.

To ensure occupancy schedules are similar, we only compare the same days of the week (e.g., only baseline Mondays are compared to the Monday of the test). Additionally, holidays are not considered. We then perform a search through each week of the 2013 and 2014 years.

Given a day i during which the MOBS^{ua} controller was active, the day used for the baseline comparison, $j^*(i)$, is then defined by (6).

$$j^*(i) := \arg\min_j d_o(i,j) \tag{6}$$

4. Experimental Setup

Experiments were performed in the University of Florida's Pugh Hall, which is a LEED Silver-certified building and has a floor space of 40,000 square feet. The rooms/days in/during which the test was conducted are called the test rooms/days. Pugh Hall is equipped with a VAV HVAC system with 3 AHUs and 65 VAV boxes, most of which have terminal reheat. Some of Pugh Hall's zones are fully actuated, and some are under-actuated. In the experiments, 3 zones consisting of 5 rooms were controlled—1 fully actuated zone and 2 under-actuated zones. Each under-actuated zone consists of 2 rooms. Three of these rooms are located on the second floor, and two of these rooms are on the third floor in the same wing of the building; see Fig. 4.

Pugh Hall's BAS is equipped with a temperature sensor in each zone but not in each room. Furthermore, humidity and CO_2 sensors were not present in the building's existing sensor suite. Therefore a WSN was designed and deployed to measure these quantities. Each sensor node has a temperature sensor, humidity sensor, CO_2 sensor, and a PIR sensor to determine whether a room is occupied or not. The PIR sensor acts as a motion detector. If any motion has been detected in the past five minutes, the room is considered to be occupied for the next five minutes. A radio transceiver on each sensor node transmits data to a base station, which transmits the data to an off-site database via the Internet. Fig. 5 shows the wireless nodes deployed in Pugh Hall, and Fig. 6 shows an example of the network structure. For more information regarding the wireless sensor network, the interested reader is referred to [18, 19].

A PIR sensor can only provide information on presence or absence of occupants, but the MOBS^{ua} needs measurements of occupancy count for a part of the control computations. Occupancy count of a room was estimated by assuming that if a room is occupied then its occupancy count is equal to its design occupancy. The design occupancy of a room is the number of people expected to be present during normal working hours.

The baseline controller that Pugh Hall's existing HVAC system uses by default is similar to a dual-maximum controller, but the specific logic is unknown due to its proprietary nature. Table 1 shows the set points used in the experiments. For ease of comparison, the set points for the MOBS^{ua} controller are chosen to be the same as those used by the baseline controller.

5. Results and Discussion

The MOBS^{ua} controller was implemented for a full week from 00:00 Tuesday, February 4, 2014 through 24:00 Monday, February 10, 2014. We refer to these days as *test days*, and we refer to the corresponding days found by the method described in Section 3.4 as *baseline days*. The evaluation metrics of Sections 3.1 and 3.2 are used to compare the control performance of each day.

5.1. Effect on thermal comfort and IAQ

The MOBS^{ua} controller regulated thermal comfort and IAQ well even in "worst" cases. CO_2 levels never exceeded 1000 ppm in any room—indicating adequate IAQ. Furthermore, relative humidity rarely exceeded 50% and never exceeded 60% for any room, and absolute humidity never exceeded 0.012 (i.e., $\Delta W \equiv 0$).

The only thermal comfort constraint violated was temperature, but these violations were minimal. Fig. 7 shows box plots of the daily ΔT for each room during the baseline days and during the test days, respectively. Because Pugh Hall did not have temperature sensors in every room before the installation of the WSN, ΔT data are not available for two of the rooms during the baseline days. The increased performance by the MOBS^{ua} controller is obvious. The larger ΔT of room 3 in Fig. 7(b) was due to actuator saturation. This room was often too warm, and despite the MOBS^{ua} controller's commands for maximum flow rate, the VAV box was not able to effectively cool the room. The large ΔT values in Fig. 7(a), however, were not due to actuator saturation. For each under-actuated zone under MOBS^{ua} control, temperatures in or close to the dead band were maintained for both rooms—even when one room warmer and the other was cooler; see Fig. 8, which shows the room temperatures in zone 2 during the test.

Fig. 8 shows indoor climate measurements of rooms 2 and 3 (composing a single under-actuated zone) during the test week. As an illustrative example, we focus on Tuesday (first day of the test). The temperature of room 3 increased beyond the dead band. As a result, the MOBS^{ua} controller considered room 3 to be the active room and commanded larger airflow to the zone in order to cool room 3; see Fig. 9. The larger amount of cooling being supplied to the zone by the increased airflow also drove the temperature of room 2 down. Eventually, room 2 left its own dead band. At that time, room 3 became unoccupied, and room 2 was then used as the active room briefly—resulting in the reheating valve being opened to return the room to its dead band.

5.2. Effect on energy consumption

The average daily energy savings for the under-actuated zones were 50±33%. For the under actuated zones, the lowest energy savings by percent were on Tuesday (29%); the highest energy savings by percent were on Saturday (80%). Although Saturday had the highest energy savings by percent, more energy savings (in kJ) were achieved each on Wednesday, Thursday, and Friday. Overall, the three test zones had a reduction in energy consumption of 62% over baseline; see Fig. 10, which shows the aggregate and component-wise energy consumptions for the test zones.

While generally consumption from cooling was reduced, rooms 2 and 3 consumed more cooling energy during two test days (Tuesday and Monday) under MOBS^{ua} control compared to baseline. The AHU discharge-air temperature was much lower on the Tuesday and Monday of the test week than on the baseline Tuesday and Monday. This resulted in more cooling energy consumption to reduce the mixed air's temperature to a lower set point. This is why the cooling energy percent savings were smaller on Tuesday and Monday in Fig. 10.

Heating energy accounted for 40% of the total energy consumed by the *baseline* controller, but 50% of the *total* energy savings came from heating. The baseline controller supplies large amounts of cooling throughout the day. As a result, re-heating is required when rooms become too cold. Since less cooling was supplied to the rooms during MOBS^{ua} control, the room temperatures did not decrease as much compared to baseline. This in turn required less heating to maintain comfortable room temperatures.

The MOBS^{ua} controller also significantly reduced fan power due to its large flow rate reduction; the fan consumed 75% less energy under MOBS^{ua} control. However, total energy savings from the supply fan were 12% because the fan accounts for less than 15% of the total energy consumption.

In [12], it was shown that there is a large variation in energy consumption (and savings) among zones. The AHU that serves the three zones controlled by the MOBS^{ua} controller also serves 22 other zones. Fig. 11 shows the percent energy reduction (between the baseline days and the test week) of the 3

MOBS^{ua}-controlled and the 22 baseline-controlled zones. As is evident, variations in load between the baseline and test days resulted in variations in total energy consumption for the zones controlled by the baseline controller during both baseline and test days. These variations, however, are small. Conversely, the zones which were controlled by the MOBS^{ua} controller clearly show significant energy savings.

Fig. 12 is a box plot of the reduction in energy consumption from each zone served by the AHU. Zones 4-25 are the other zones served by the same AHU and are controlled by the baseline controller on all days. As is clear from the figure, the zones under MOBS^{ua} control consumed significantly less energy than did those under baseline control. The reduction in zones 4-25 can be interpreted as the change in consumption due to differing conditions between the baseline and test days (such as weather). The additional savings in zones 1-3 are the changes in consumption due to the MOBS^{ua} control logic. Note that zone 2 consumed more energy on Tuesday and Monday as discussed above—resulting in a taller box in the plot.

The MOBS^{ua} controller also reduced peak power consumption of two of the three test zones compared to baseline. This is critical for energy efficiency because many commercial buildings are charged for their peak power consumptions in addition to overall energy consumptions. It should be noted, however, that the peak power consumption in zone 2 increased on Tuesday, Wednesday, and Monday; see Fig. 13. This is a result of the increased flow rate of the MOBS^{ua} controller; compare with Fig. 9.

6. Conclusions

The MOBS^{ua} controller was shown to offer significant energy savings (29-80%) for under-actuated zones. A method for determining a baseline for comparing the performances of distinct controllers was proposed as well. Despite the test week having extraordinarily cold weather, heating energy was reduced significantly, and 50% of the total energy savings came from heating.

Cooling energy consumption was found to be dependent on AHU discharge-air temperature. This suggests that control of AHU-level variables may be required for energy savings guarantees because the AHU controller may "fight" the MOBS^{ua} controller by changing the AHU discharge-air temperature set point. The MOBS^{ua} controller also resulted in significantly fewer temperature constraint violations compared to the baseline controller as well—even when one room of an under-actuated zone was hot and the other was cold.

These positive results were achieved by the MOBS^{ua} algorithm without using any model of the building or its zones. This is an important point as modeling a building's hygro-thermal dynamics is non-trivial and requires extensive time and labor. Only building parameters such as duct ratios, outside-airflow ratio, room sizes, and room types are needed *a priori*. Furthermore, because control actions are decided independently for each zone, the algorithm is fully scalable to a large number of zones without any increase in complexity. Additionally, no expensive occupancy estimation algorithm or detection equipment was used—a simple estimate of occupancy count using the design occupancy was satisfactory. On the other hand, measurements of the exact number of occupants in a room can increase savings even further due to ASHRAE Standard 62.1.

7. Acknowledgments

The authors thank Dr. Timothy Middelkoop and UF's Physical Plant Division for their help. This research is supported in part by the National Science Foundation awards CPS-0931885 and ECCS-0925534.

8. References

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

[2] J. Široký, F. Oldewurtel, J. Cigler, S. Privara, Experimental analysis of model predictive control for an energy efficient building heating system, Applied Energy 88 (9) (2011) 3079–3087.

[3] D. Kolokotsa, A. Pouliezos, G. Stavrakakis, C. Lazos, Predictive control techniques for energy and indoor environmental quality management in buildings, Building and Environment 44 (9) (2009) 1850–1863.

[4] S. Bengea, A. Kelman, F. Borrelli, R. Taylor, S. Narayanan, Model predictive control for mid-size commerical building HVAC: Implementation, results and energy savings, in: 2nd International Conference on Building Energy and Environment, 2012.

[5] S. Goyal, P. Barooah, Energy-efficient control of an air handling unit for a single-zone VAV system, in: 52nd IEEE Conference on Decision and Control, 2013, pp. 4796 – 4801. doi:10.1109/CDC.2013.6760641.

[6] S. Goyal, P. Barooah, T. Middelkoop, Experimental study of occupancy-based control of HVAC zones, Unpublished results.
URL http://humdoi.mae.ufl.edu/~pbarooah/PBpublication.html

[7] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, 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.

[8] P. X. Gao, S. Keshav, SPOT: a smart personalized office thermal control system, in: Proceedings of the 4th International Conference on Future Energy Systems, ACM, 2013, pp. 237–246.

[9] V. L. Erickson, S. Achleitner, 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.

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

[11] S. Goyal, H. Ingley, P. Barooah, Occupancy-based zone climate control for energy-efficient buildings: Complexity vs. performance, Applied Energy 106 (2013) 209–221. doi:10.1016/j.apenergy.2013.01.039.

[12] J. Brooks, S. Goyal, R. Subramany, Y. Lin, T. Middelkoop, L. Arpan, L. Carloni, P. Barooah, An experimental investigation of occupancy-based energy-efficient control of commercial building indoor climate, in: 53rd IEEE Conference on Decision and Control, 2014, in press.
URL http://humdoi.mae.ufl.edu/~pbarooah/PBpublication.html

[13] J. Brooks, P. Barooah, Energy-efficient control of under-actuated HVAC zones in buildings, in: Proceedings of the 2014 American Control Conference, 2014, pp. 424–429.

[14] American Society of Heating, Refrigerating and Air Conditioning Engineers, The ASHRAE Handbook : Applications (SI Edition) (2011).

[15] American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., ANSI/ASHRAE Standard 62.1-2010, ventilation for acceptable air quality (2010).
URL www.ashrae.org

[16] ASHRAE, The ASHRAE Handbook Fundamentals (SI Edition) (2009).

[17] H. R. Bohanon, Good IAQ practices, ASHRAE Journal 54 (2012) 106–107.

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

[19] S. Kumar, Improved wireless sensor network for HVAC control, Master's thesis, University of Florida (April 2014).

List of Figures

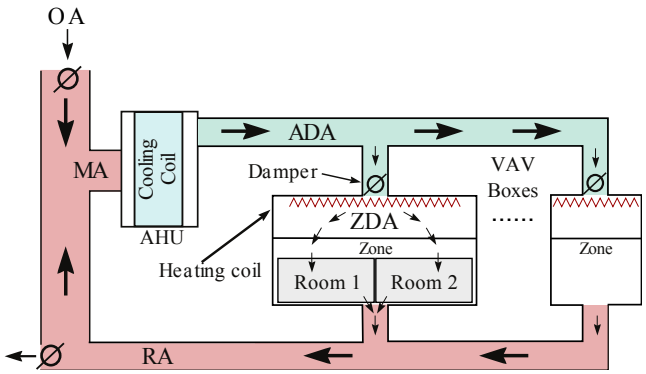

Figure 1: A typical single-duct VAV HVAC system with an under-actuated zone shown. ADA represents the AHU discharge air; ZDA represents the zone discharge air after the heating coil; RA represents the recirculated air from the zones; OA represents the outside air; MA represents the mixed air.

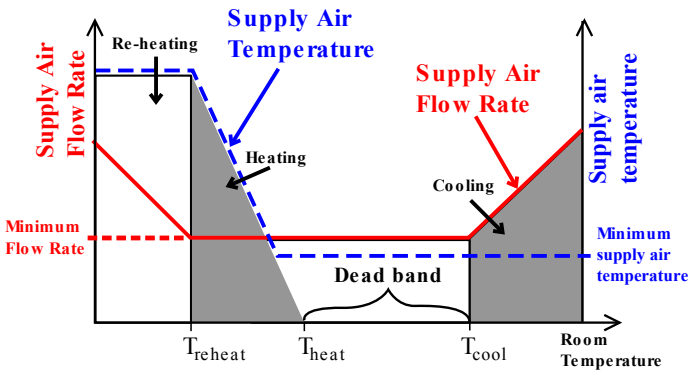

Figure 2: Diagram of dual-maximum control logic and operation modes.

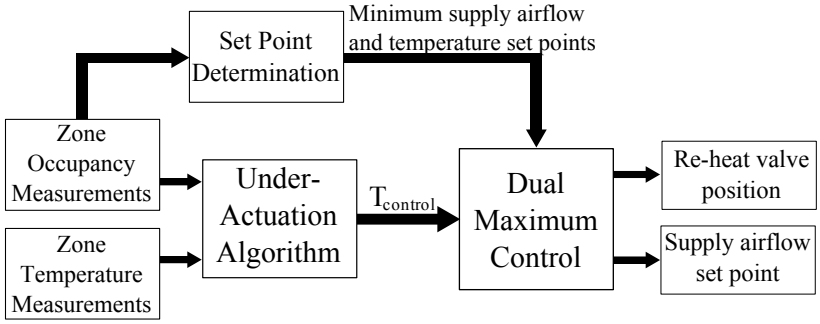

Figure 3: MOBS^{ua} control logic.

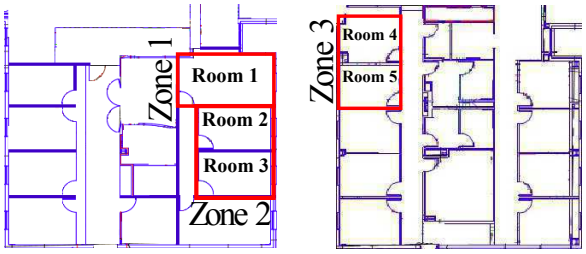

(a) Second floor (b) Third floor

Figure 4: Layout of the five test rooms.

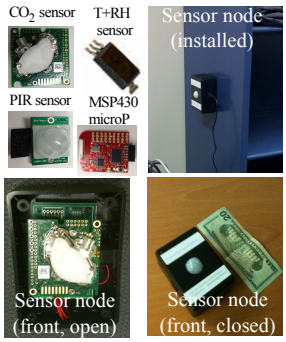

Figure 5: Wireless sensor node.

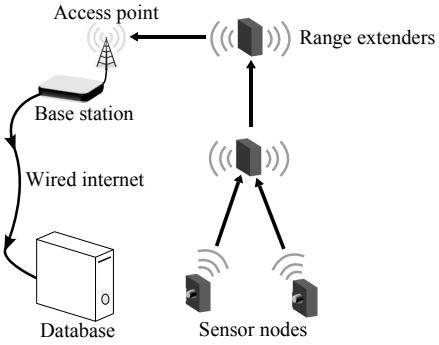

Figure 6: Example of network structure.

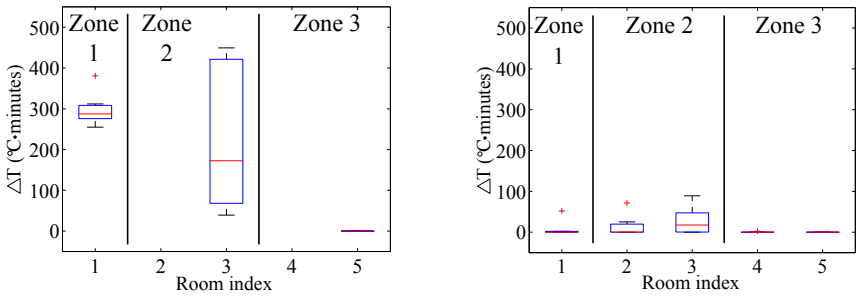

(a) Baseline controller

(b) MOBS$^{\text{ua}}$ controller

Figure 7: Box plot of daily ΔT (seven data points each.)

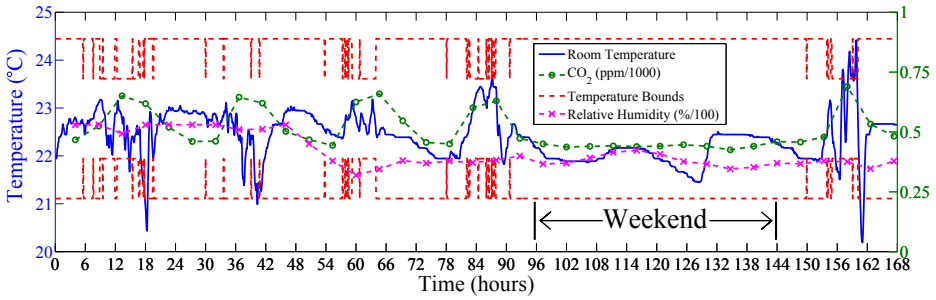

(a) Room 2

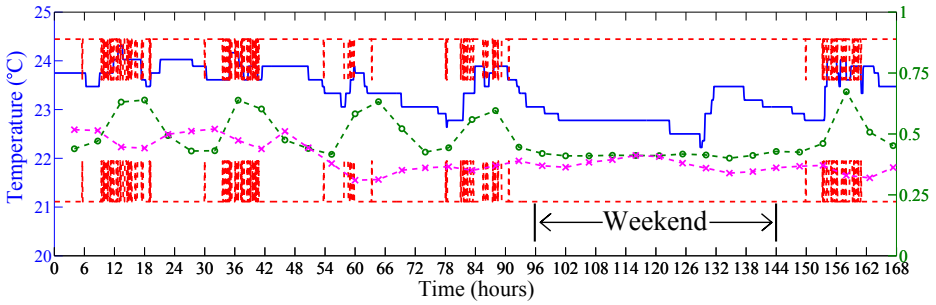

(b) Room 3

Figure 8: Temperature (left axis) and CO_2 and relative humidity (right axis) of two rooms served by the same VAV box. Note that the temperature constraints are stricter when the room is occupied. Despite the rooms having very different temperatures, comfort constraints were generally still maintained.

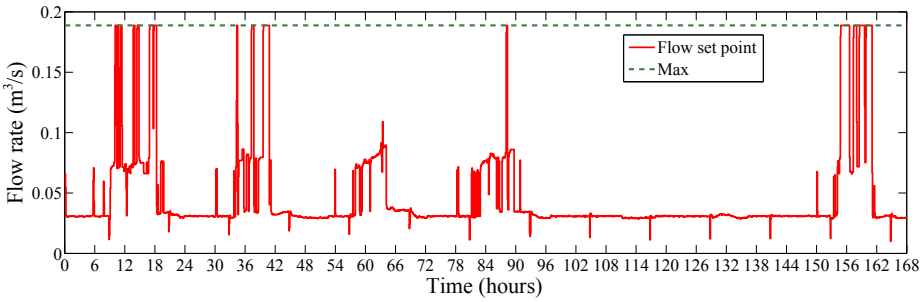

Figure 9: Flow rate for zone 2 (rooms 2 and 3).

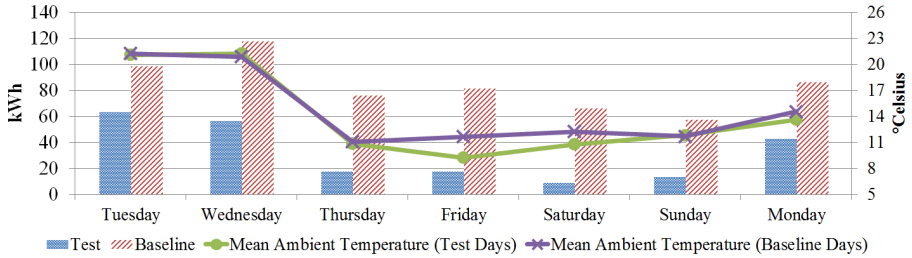

(a) Aggregate energy consumption

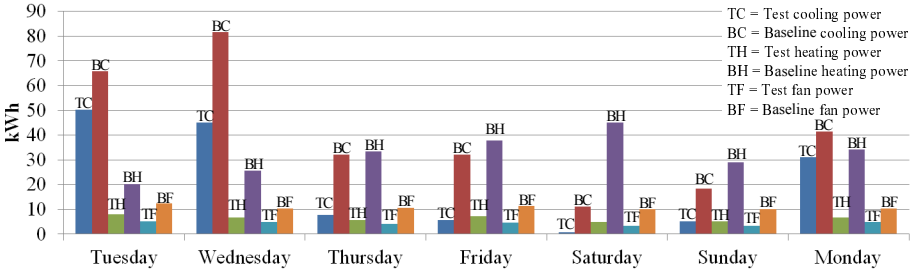

(b) Energy consumption by component

Figure 10: Aggregate and component-wise energy consumptions during test week and baseline days.

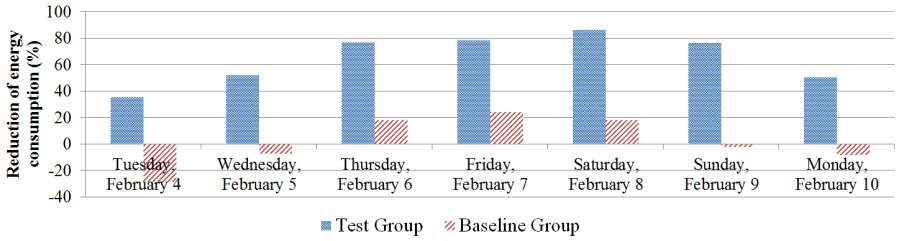

Figure 11: Percent energy reduction for test group and baseline group on the test days over the baseline days. Test group (left) was under MOBS$^{\text{ua}}$ control during the test week; baseline group (right) was under baseline control during the test week.

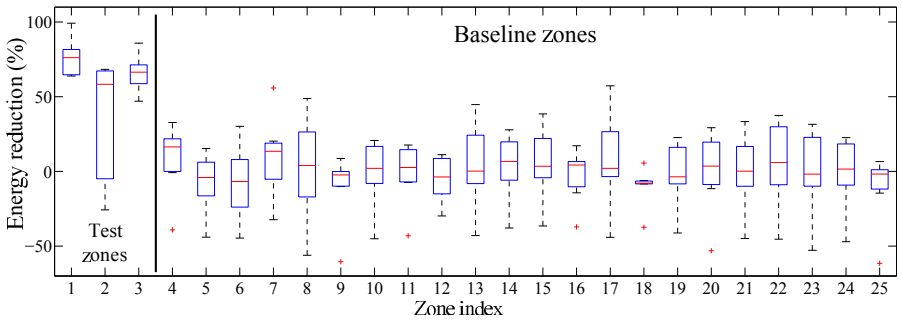

Figure 12: Daily energy reduction in each zone between the test days and the baseline days. Zones 1-3 were under MOBS$^{\text{ua}}$ control during the test week and under baseline control during the baseline days. Zones 4-25 were under baseline control for both the test week and the baseline days.

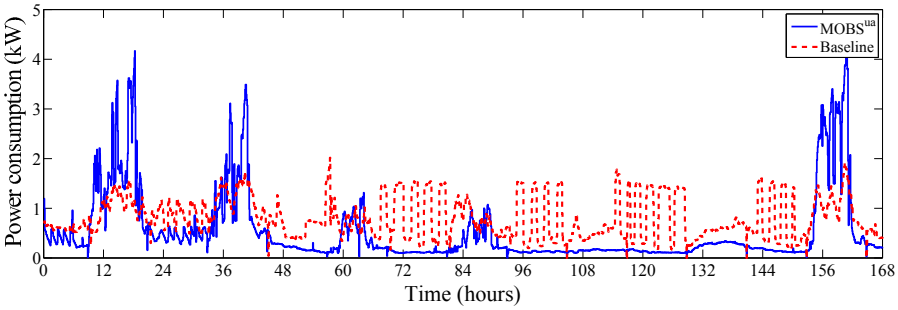

Figure 13: Power consumption of zone 2 for both MOBS$^{\text{ua}}$ and baseline control.

List of Tables

Table 1: Temperature set points used by MOBS^{ua} and the baseline controller in the experiments.

	T_{cool}	T_{heat}	T_{reheat}
Occupied	23.6 °C	21.9 °C	21.8 °C
Unoccupied	24.4 °C	21.1 °C	20.9 °C