Occupancy-Based Zone-Climate Control for Energy-Efficient Buildings: Complexity vs. Performance ☆

Siddharth Goyal, Herbert A. Ingley, Prabir Barooah

Department of Mechanical and Aerospace Engineering, University of Florida, Gainesville, FL 32611, USA

Abstract

We propose several control algorithms to compare their performance and complexity through simulations; the control algorithms regulate the indoor climate of commercial buildings. The goal of these control algorithms is to use occupancy information to reduce energy use—over conventional control algorithms—while maintaining thermal comfort and indoor air quality. Three novel control algorithms are proposed, one that uses feedback from occupancy and temperature sensors, while the other two computes optimal control actions based on predictions of a dynamic model to reduce energy use. Both the optimal-control based schemes use a model predictive control (MPC) methodology; the difference between the two is that one is allowed occupancy measurements while the other is allowed long term occupancy prediction. Simulation results show that each of the proposed controllers lead to significant amount of energy savings over a baseline conventional controller without sacrificing occupant health and comfort. Another key finding is that the feedback controller performs almost as well as the more complex MPC-based controllers. In light of the informational/computational complexity of the MPC algorithms compared to the feedback control algorithm, we conclude that feedback control is more suitable for energy-efficient zone-climate control than MPC-based control, and that the difficulty of obtaining occupancy predictions is not commensurate with the resulting benefits.

☆This work has been supported by the National Science Foundation by Grants CNS-0931885 and ECCS-0955023. Author email addresses: {siddgoya,ingley,pbarooah}@ufl.edu.

Keywords: Occupancy based building climate control, model predictive control, energy-efficient buildings, building thermal dynamics.

Nomenclature

CLG	Cooling set-point
D_H	Humidity violation
$D_H^\star$	Average humidity violation
D_T	Temperature violation
$D_T^\star$	Average temperature violation
E_C	Energy consumed by controller C
E_{BC}	Energy consumed by the baseline controller
H	Relative humidity
HTG	Heating set-point
K	Number of steps chosen for prediction horizon during the optimization
P	Total power
P_F	Fan power
P_R	Re-heating power, i.e., power consumed in reheating at the variable-air-volume (VAV) box
P_U	Conditioning power, i.e., power consumed by chiller
Q^s	Rate of heat gain due to solar radiation
RTG	Re-heating set-point
R^{RA}	Return air ratio (ratio of return air to mixed air flowrate)
T	Temperature
T^{set}	Desired set-point

T_{RTG} Re-heating set-point

T_{high} Maximum temperature allowed in the zone

T_{low} Minimum temperature allowed in the zone

W Humidity ratio

W_{high} Maximum humidity ratio allowed in the zone

W_{low} Minimum humidity ratio allowed in the zone

Δt Discretization time

α IAQ factor of safety

h Enthalpy of air

m_z^A Amount of fresh outside air required per unit area

m_p^{OA} Amount of fresh outside air required per person

m_p^{SA} Amount of supply air required per person

m_{high}^{SA} Maximum amount of supply air during occupied or unoccupied time

m_{low}^{SA} Minimum amount of supply air during unoccupied time

n^p Number of people

u Controllable input vector

v Exogenous input vector

A Floor area

β Fan power constant

subscripts

d Designed

superscripts

z Zone

OA Outside air

occ During occupied time

$unocc$ During unoccupied time

CA Conditioned air: air being supplied by air handling unit (AHU)

SA Supply air (air leaving the VAV box)

1. Introduction

Buildings are one of the primary energy consumers worldwide. In the United States, they account for about 40% of the total energy consumption [1]. Heating ventilation and air-conditioning (HVAC) contributes to more than 50% of the energy consumed in buildings [1]. Poor design and inefficient operation of HVAC system cause a large fraction of energy used to be wasted [2, 3]. Though it is possible to improve energy efficiency through better HVAC system design, it requires substantial investment to retrofit an existing building with improved HVAC equipment. In contrast, improving control algorithms (that operate the HVAC system) to achieve higher efficiency is far more cost effective, as long as such a solution does not require expensive additional sensors. Indeed, a number of recent papers have focused on improving energy efficiency in buildings through advanced control algorithms that use occupancy information [4, 5, 6, 7, 8, 9, 10]. This is the subject of our paper as well; we consider control algorithms that use occupancy information to maintain the climate of individual zones at appropriate conditions with reduced energy use compared to conventional control algorithms. An important constraint is cost; one should be able to apply these control algorithms with minimal investment.

We limit ourselves to commercial buildings with variable-air-volume (VAV) system. More than 30% of the commercial building floor space in the United States is served by VAV systems [11]. In a VAV system, a building is divided into a number of "zones", where a zone can be single room or a collection of rooms. The flow rate of supply air, i.e., air supplied to a zone, is controlled through dampers in the VAV box of the respective zone. The conditioned air, which is the air supplied by an AHU, may be reheated at the VAV box before being supplied to the zone. We focus on control strategies that can be applied at each VAV box, where the control inputs that need to be decided are the mass flow rate and temperature of the supply air.

Typically, a simple rule-based feedback control strategy is used at the VAV box that does not use real-time occupancy measurements[1]. The controller determines the flow rate of air supplied to the zone, as well as any reheat to be applied, to maintain the temperature of the zone at specific ranges that are based on predetermined occupancy schedules. To maintain indoor air quality (IAQ), the minimum airflow rate is determined based on the occupancy schedules and building standards, such as ASHRAE (American Society of Heating, Refrigerating and Air-Conditioning Engineers) ventilation standard 62.1-2010 [12]. This minimum flow rate is usually 30–40% of the designed maximum. Hence, zones are typically over-ventilated, especially when the zone is not occupied but it is expected to be, e.g., in "daytime" mode. This causes wastage of energy.

Over-ventilation can be prevented by applying demand control ventilation (DCV), i.e., by changing the supply air flow rate based on real-time occupancy measurements or CO_2 measurements instead of a pre-defined schedule. Real-time occupancy measurements can be obtained from motion detectors such as PIR and ultrasound sensors, which are inexpensive and work well in small office spaces where the nominal occupancy value is one [5, 13]. In very large spaces, CO_2 measurements can be effectively used for DCV in lieu of an occupancy sensor. DCV is typically used in large spaces with the help of CO_2 sensors; its use in small zones (such as office rooms) is less common. In medium-sized spaces where the nominal occupancy is more than one but not very large, measuring occupancy is non-trivial. Efforts in developing occupancy measurement technology are carried out by several researchers; see [14] and references therein.

As sensors and/or algorithms for inexpensive yet reliable real-time occupancy measurement/estimation become available, it should be possible to do more to reduce energy use apart from controlling ventilation. For example, we can save energy by reducing the airflow rate as well as letting the temperature float during unoccupied times in a wider range than when it is occupied. Caution is required while developing a control algorithm to achieve that objective. For instance, if we let the temperature during unoccupied times deviate far away from what is considered comfortable, it might take a while for the zone temperature to come back to a comfortable condition when the zone becomes occupied again. Same goes for humidity and IAQ. Thus, the dynamics of temperature, humidity, and IAQ

[1] In this paper, "occupancy" is used to denote the number of people in a space.

have to be taken into account in designing such control algorithms. Moreover, the controller should also have some robustness to error in occupancy measurements.

In this paper, we examine how much energy can be saved by control algorithms that use information of occupancy and system dynamics, and how the savings depend on the fidelity of the information. As more fine-grained information is available, we may be able to save more, but the control algorithm may become more complex. Our focus is on control algorithms that can be used in VAV boxes of individual zones in existing (and new) commercial buildings: the controller has to decide the flow rate and temperature of the air supplied to the zone. It can vary the airflow rate between 0 and some upper bound, while the temperature can be only increased beyond the temperature of the conditioned air (air leaving the AHU) by using the reheat coil, but not decreased. Though it is possible to add additional actuation such as controllable window blinds, they require significant hardware upgrade, and therefore are not considered here.

The first and the simplest controller we propose is a rule-based feedback control law that decides the control inputs based on instantaneous measured occupancy. This control strategy is called *MOBS* (Measured Occupancy Based Setback), since it typically “sets back” the zone temperature set-points and flow rate to smaller values during unoccupied times. During occupied times, the zone temperature is maintained in the same range as that a conventional controller would do, and the flow rate is determined based on the measured occupancy.

A natural choice for a control algorithm is one that minimizes energy consumption while satisfying constraints on the thermal comfort and IAQ. We next propose a controller that does so by solving an optimal control problem in a receding horizon fashion. That is, given a dynamic model relating the control inputs to the relevant outputs (temperature, humidity, etc.,), the controller computes the control inputs that will minimize the total energy consumption over some finite time interval, say ΔT, as well as maintain pre-specified constraints on zone temperature, humidity etc. It applies the resulting inputs for a time interval that is shorter than ΔT, and then re-computes the control inputs for the next interval of length ΔT by utilizing up-to-date measurements. The process is repeated. This method of computing the control inputs is called receding horizon control (RHC) or MPC, and is widely used in practice [15] due to its ability to efficiently compute optimal control solutions in problems that involve constraints. Several recent papers have proposed MPC-based controllers for efficient building control, which we discuss in Section 1.1. Application of MPC to energy minimization control requires a model of the dynamics of temperature, humidity, and IAQ (i.e., contaminants), as well as predictions of the exogenous inputs such as solar radiation, temperatures of surrounding spaces, and most importantly, occupancy. Occupancy affects not only the heat gains and humidity but also the constraints. For instance, the temperature of the zone can be allowed to vary in a wider range during unoccupied times compared to the range during occupied times. A model of the hygro-thermal dynamics, i.e., of temperature and humidity, are obtained using lumped parameter models [16]. However, IAQ dynamics are not modeled since there is no well accepted numerical metric for measuring IAQ. Instead constraints are posed on the air flow rate according to ASHRAE ventilation standard 62.1-2010 [12] so that IAQ is assured. The proposed MPC-based controller is called *POBO* (Predicted Occupancy Based Optimal), because it uses occupancy predictions, among other things, to calculate the control inputs.

Obtaining occupancy prediction is quite challenging. Therefore, we next consider the case when an MPC algorithm is sought that uses only occupancy measurements, not predictions. In this case, to use an MPC formulation, future occupancy is assumed to stay at the currently measured occupancy. We call the resulting control scheme the *MOBO* (Measured Occupancy Based Optimal) controller because it uses occupancy measurements—as opposed to prediction—to calculate the optimal control inputs. This controller is expected to retain some of the benefits of MPC, such as directly minimizing energy and maintaining constraints, while making the information requirements of the controller more realistic.

Performance of the proposed algorithms are compared through simulations to that of a conventional controller used in existing commercial buildings; the so-called "dual maximum" control [17, Chapter 47]. We henceforth refer to dual maximum as the *BL* (*baseline*) controller. The *BL* and *MOBS* control algorithms are pure feedback strategies. The *BL* controller uses only zone temperature measurements but not occupancy measurements. The *MOBS* controller uses measurements of both zone temperature and occupancy. Both the *MOBO* controller and the *POBO* controller need a model of the hygro-thermal dynamics of the zone to solve the underlying optimal control problems. While *MOBO* requires occupancy measurements, the *POBO* algorithm requires occupancy prediction. Thus, the complexity of the control algorithms increases in the order *BL*, *MOBS*, *MOBO*, *POBO*.

Simulations are performed for three different types of zones exposed to several types of outside weather and climates. The zone consists of single room in all simulations, with design occupancy varying between one and three. The main conclusions from the simulations are the following:

1. It is possible to obtain significant energy savings by using occupancy mea-

surement or prediction, with negligible effect on the zone IAQ and occupants thermal comfort. The proposed controllers that use occupancy measurement or prediction provide savings of $40 - 60\%$ over the baseline controller that does not use real-time occupancy measurement but only nighttime setback.

2. A simple rule-based feedback control algorithm can perform just as well as an MPC-based control algorithm, if only occupancy measurements (not predictions) are available.
3. If occupancy predictions are available, an MPC controller that uses these predictions can yield higher energy savings over controllers that use only occupancy measurements. However, the additional savings are small; about 1–13%. This is due to the minimum ventilation required by the current building standards, which prevents the controller from drastically reducing the airflow even when it is known that the zone will remain unoccupied.

These results show that by installing sensors capable of providing occupancy measurements and augmenting the control logic at the VAV boxes to use these measurements, substantial energy savings can be achieved. This study shows that the benefit of using MPC for energy-efficient zone-climate control is questionable; both the feedback controller and the MPC controller provide similar energy savings while the MPC controller is much more complex. The results also indicate that the effort required in obtaining occupancy prediction may not be commensurate with the benefit obtained. It should be noted that occupancy prediction is quite challenging; there are only a few papers on dynamic occupancy models [13, 5]. It is also not clear how easily such models can be calibrated to individual buildings and zones, and how accurate their predictions will be in general.

The rest of paper is organized as follows. In the remainder of the section, we discuss the related literature. The baseline control and the proposed control, along with the model of hygro-thermal dynamics and power consumption, are described in Section 2. Section 3 describes performance metrics related to thermal comfort and energy savings. Simulation results with the controllers are shown in Section 4. Section 5 concludes the paper with a discussion of the results and ways to extend this work.

1.1. Related work

A number of papers have investigated the use of occupancy information (either measurements or predictions) to reduce energy consumption in buildings. The papers [6, 7, 8, 9] compare MPC-based controllers that use occupancy prediction with conventional controllers that do not use such prediction, apart from day/night schedules. They report substantial energy savings with MPC compared to conventional controllers. However, these papers do not investigate how much energy savings are possible with any controller that is less complex than MPC and that uses occupancy measurements, which are easier to obtain than occupancy prediction.

A number of papers have proposed simple rule-based controllers that use occupancy measurements, and conclude that significant energy savings are possible with the rule-based controllers compared to the conventional controllers that do not use occupancy measurements [18, 19, 20, 21]. The controller in [18] uses occupancy measurements to turn off the HVAC system, while the controllers in papers [19, 20, 21] modulate only the ventilation rate based on measured occupancy. However, these papers do not compare rule-based control with complex control schemes such as MPC. While MPC may require more information (i.e., dynamic model and occupancy prediction) compared to rule-based control it may also lead to more energy savings. The paper [5] compares several rule-based controllers that use various types of occupancy information: two use occupancy prediction while one uses binary occupancy measurements (presence/absence). It is concluded that significant energy savings are possible with the rule-based feedback control that uses binary occupancy measurements compared to the baseline controller that does not. It also concludes that a small amount of additional energy savings are possible if the predictive rule-based controller is used instead of the feedback controller. However, it does not compare the predictive rule-based controller with complex predictive control algorithms such as MPC, which may result in more savings than the rule-based control.

While some of the previous work has compared either MPC or rule-based controllers with conventional controllers, they did not compare all three. The conventional controllers used for comparison were distinct, making such comparison harder. It is useful to know how performance (as measured by energy savings and/or comfort) varies with the complexity of the control algorithm. In particular, the value of occupancy measurement vs. occupancy prediction is not clear from the prior work. Since prediction is much more difficult to obtain than measurement, it is particularly useful to know their relative value. Though [5] compares performance of the rule-based controllers that use occupancy prediction with that of the feedback controller, the feedback controller uses only presence/absence measurement but not occupancy measurement.

In this paper we examine the performance of a baseline conventional controller (without occupancy measurement or prediction), a feedback controller (with occupancy measurement), and two MPC-based control schemes. One of the MPC-based control schemes uses occupancy prediction while the other uses occupancy measurement in lieu of prediction. This way, we are able to compare the performance of feedback control to that of MPC-based control when both are allowed only occupancy measurement. In short, we examine trade-off between energy savings achieved and the information requirements/complexity of the control algorithm in a unified manner. This is the key difference between our work and much of prior work. In addition, the papers mentioned above that propose MPC-based controllers do not take humidity into the control computation, while humidity is taken into account as part of thermal comfort constraints in the MPC schemes proposed here.

A preliminary version of this work appeared in [4]; which compares feedback, MPC, and baseline controllers. There are several significant differences between [4] and this paper. The baseline conventional controller used here is more energy-efficient than the one used in [4]. The feedback controller used in [4] modulates only the ventilation rate based on measured occupancy. However, the feedback strategy in this paper controls not only the ventilation rate but also the zone temperature, which results in high energy savings. The design parameters have been fine-tuned in this paper to get better performance from all the controllers. One of the controllers proposed in [4] allowed 0 flow rate when the zone was known to be unoccupied. In this paper, all controllers are designed to supply a minimum airflow rate in accordance with the latest ASHRAE ventilation standard 62.1-2010 [12]. This significantly changes some of the conclusions, especially one about the value of occupancy predictions. Moreover, this paper provides a more comprehensive simulation study of the performance of the controllers compared to [4]. While [4] considers one type of zone with three occupants exposed to only one type of outside weather, here we examine several types of zones with varying levels of occupancy that is exposed to multiple outside weather and climate conditions.

2. Control Algorithms

A schematic of a typical multi-zone commercial building with a VAV-based HVAC system and a conceptual representation of a control algorithm that can be implemented in a zone is shown in Figure 1. Part of the air removed from the zones, which is called return air, is mixed with the outside air before being conditioned at the AHU to temperature T^{CA} and humidity ratio W^{CA}. The conditioned air, which is usually cold and dry, is distributed to the VAV boxes at the zones through the ductwork. The air supplied to a zone by its VAV box can be heated using the reheat coils at the box. The amount of return air and outside air that needs to be mixed is decided by the return air ratio R^{RA}. The humidity ratio of the supply air (W^{SA}) is same as the humidity ratio of air being supplied by the AHU, i.e., ($W^{SA} = W^{CA}$), since reheating does not change the humidity ratio. The parameters T^{CA}, W^{CA} and R^{RA} are assumed constant in this paper.

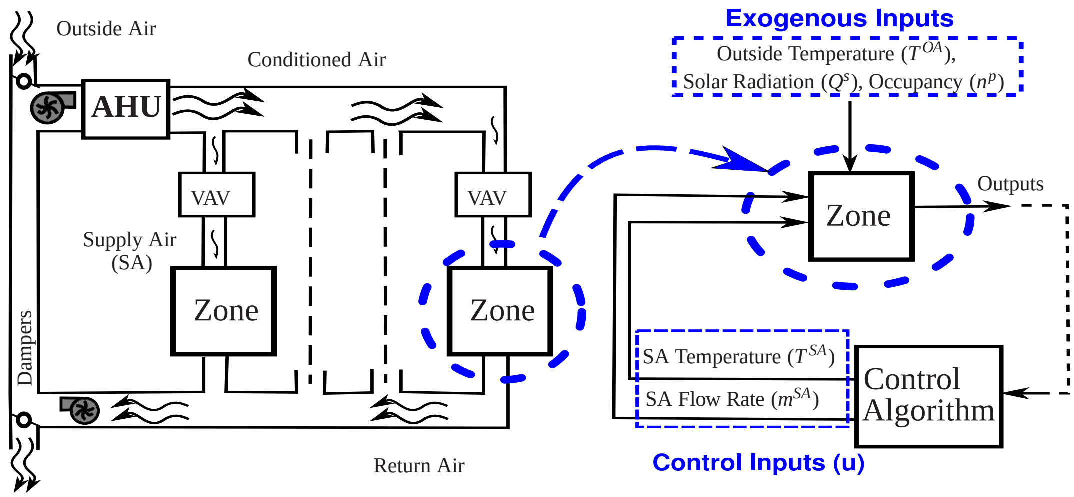

Figure 1: Generic scheme for the implementation of a zone-level control algorithm.

The task of a zone-climate control algorithm is to decide the control inputs in such a way that thermal comfort and IAQ are maintained in that zone. The control inputs are temperature (T^{SA}) and flow rate (m^{SA}) of the air supplied to that zone by its VAV box. The control algorithm may require certain measurements and/or predictions to compute the control inputs, which vary depending on the control algorithm. For instance, the commonly used single maximum and dual maximum control logics [17, Chapter 47] require only zone temperature measurements. However, an MPC-based controller, such as that proposed in [4] or in this paper, requires measurement and prediction of outside temperature, solar radiation, occupancy, zone temperature and humidity.

We now describe the *BL* (baseline) controller, and three proposed control algorithms, *MOBS* (Measured Occupancy Based Setback), *MOBO* (Measured Occupancy Based Optimal) and *POBO* (Predicted Occupancy Based Optimal).

*2.1. Baseline (*BL*) Controller*

Among the common control logics used at the VAV boxes to maintain IAQ and temperature in a zone, we choose the dual maximum [17, Chapter 47] as the *baseline controller*. Even though the single maximum control [17, Chapter 47] is more common in existing commercial building, dual maximum is the more efficient of the two. In this scheme, the control logic is divided into four modes based on the zone temperature: (i) Re-heating (ii) Heating (iii) Dead-Band and (iv) Cooling, which are shown schematically in Figure 2. If the zone temperature stays below the "Re-heating Set-Point (RTG)" for more than 10 minutes, the re-heating mode is turned on. Similarly, if the zone temperature remains above the "Cooling Set-Point (CLG)" for more than 10 minutes, the cooling mode is turned on. If the zone temperature stays between RTG and "Heating Set-Point (HTG)" for more than 10 minutes, the heating mode is turned on. If the zone temperature stays between HTG and CLG for more than 10 minutes, the dead-band mode is turned on. In the re-heating mode, the supply air temperature is set to maximum possible value (T^{SA}_{high}), and the supply air flow rate is varied using a PID controller to maintain the zone temperature to a desired set-point T^{set}. In the heating mode, the supply air flow rate is set to the minimum allowed value, and the supply air temperature is controlled by a PID controller so that the zone temperature is maintained close to the set-point (T^{set}). The minimum allowed value for the flow rate is determined as follows

$$\begin{aligned} &\text{Minimum Allowed Flow Rate} = m_p^{SA} n_d^p + \alpha m_{low}^{SA}, \\ &\text{where } m_p^{SA} = m_p^{OA}/(1-R^{RA}),\ m_{low}^{SA} = m_z^A A_z/(1-R^{RA}). \end{aligned} \tag{1}$$

When $\alpha = 1$, these calculations yield the minimum airflow requirements specified by ASHRAE ventilation standard 62.1-2010 [12]. Since the baseline controller does not use occupancy measurement, the minimum allowed flow rate is calculated using the designed occupancy n_d^p, which is assumed constant. We usually choose $\alpha > 1$ to make IAQ robust to mismatches between actual and designed occupancy. In the dead-band mode, no re-heating is performed, i.e., $T^{SA} = T^{CA}$, and supply air flow rate is set to the minimum allowed value (1). In the cooling mode, no heating or re-heating is performed, i.e., $T^{SA} = T^{CA}$, but the supply flow rate is varied to maintain the desired set-point T^{set} in the zone.

The desired set-point T^{set} used by the PID controllers in the re-heating, heating and cooling modes is usually the temperature preferred by the occupants. If the temperature preferred by the occupants is not known, then there are several other ways to decide the value of T^{set}. One way is to choose T^{set} as RTG, HTG and CLG during the re-heating, heating, and cooling modes, respectively. Another way is to chose T^{set} as an average of HTG and CLG during all the modes. We choose T^{set} as the average of HTG and CLG in this paper, i.e., $T^{set} = \frac{HTG+CLG}{2}$.

Note that the baseline controller uses nighttime setback: the set-points RTG and HTG are decreased while the set-point CLG is increased during a pre-specified period deemed “nighttime”. The set-points are changed based on the assumption that the zone is not occupied during the night, which results in reduced energy usage.

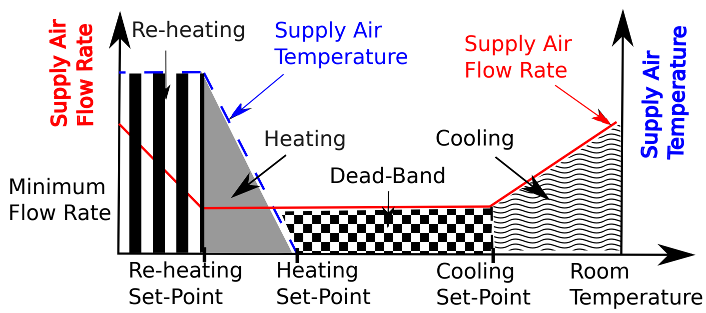

Figure 2: Schematic representation of the baseline control strategy (“dual maximum”) used at the VAV terminal boxes of commercial buildings.

2.2. *Measured Occupancy Based Setback (*MOBS*) Controller*

The proposed *MOBS* control strategy requires occupancy measurements in addition to the zone temperature measurements. It is quite similar to the *BL* controller described in Section 2.1, except for two key differences. First, the minimum allowed flow mentioned in (1) is calculated based on the measured occupancy instead of the design occupancy, which is expressed as

$$\text{Minimum Allowed Flow Rate at time } t = m_p^{SA} n^p(t) + \alpha m_{low}^{SA}, \tag{2}$$

where $n^p(t)$ is the occupancy measured at time t, and m_p^{SA}, m_{low}^{SA} are computed using (1). Second, the temperature set-points are determined based on whether the zone is occupied or not:

$$\left.\begin{aligned} RTG(t) &= T_{RTG}^{unocc} \\ HTG(t) &= T_{low}^{unocc} \\ CTG(t) &= T_{high}^{unocc} \end{aligned}\right\} \text{if } n^p(t) = 0, \quad \left.\begin{aligned} RTG(t) &= T_{RTG}^{occ} \\ HTG(t) &= T_{low}^{occ} \\ CTG(t) &= T_{high}^{occ} \end{aligned}\right\} \text{if } n^p(t) \neq 0. \tag{3}$$

The choice of design variables $T^{unocc}_{RTG}, T^{occ}_{RTG}, T^{unocc}_{low}, T^{occ}_{low}, T^{unocc}_{high}, T^{occ}_{high}$ involves a trade-off between energy savings and thermal comfort. Clearly, the range $[T^{occ}_{low}, T^{occ}_{high}]$ should be chosen to ensure that occupants are comfortable if the zone temperature is within this range. A wider range will in general reduce energy consumption, since the controller may be able to reduce reheating during low thermal load conditions and reduce the airflow during high thermal load conditions. Too wide a range will, however, lead to discomfort on the occupants part. As a general rule, the parameters for the unoccupied periods should be chosen so that

$$[T^{occ}_{low}, T^{occ}_{high}] \subseteq [T^{unocc}_{low}, T^{unocc}_{high}], \tag{4}$$

i.e., the temperature is allowed to vary within a wider range of values during unoccupied periods than in occupied ones. This is expected to lead to energy savings as well. However, even in unoccupied times it is not advisable to let the temperature deviate too far from what is allowed during occupied times. Otherwise, when the zone becomes occupied again, it will take a long time to bring the temperature back to the range allowed during the occupied time, which will cause discomfort to the occupants. In addition, letting the temperature become too low may cause condensation on surfaces leading to mold growth. Similarly, choosing the reheating set-points ($T^{unocc}_{RTG}, T^{occ}_{RTG}$) far from the heating set-points (T^{unocc}_{low}, T^{occ}_{low}) is likely to lead to not only more the energy savings but also more discomfort.

The algorithm described above is termed *MOBS* (Measured Occupancy Based Setback) control because, in general, it sets back the temperature set-points (RTG, HTG, and CLG) and the airflow rate when the zone is not occupied.

2.3. MPC-based Controllers

In this section, we propose two MPC-based control algorithms: *MOBO* and *POBO*. The block diagram of the implementation of the *MOBO* and *POBO* controllers is shown in Figure 3. Time is measured with a discrete index $k = 0, 1, \dots,$ where the time period between k and $k+1$ is denoted by Δt. Both the controllers compute the control inputs $(T^{SA}(k), m^{SA}(k))$ over K time indices by solving an optimization problem which minimizes total energy consumption over that period while maintaining thermal comfort and IAQ. The control inputs are applied at the current time index k. The optimization problem is solved again at time index $k+1$ to compute the control inputs for the next K time instants. The whole process is repeated ad infinitum.

To solve the underlying optimization problem, the controllers need (i) predictions of the exogenous inputs such as T^{OA}, W^{OA}, Q^s and n^p, over the time horizon

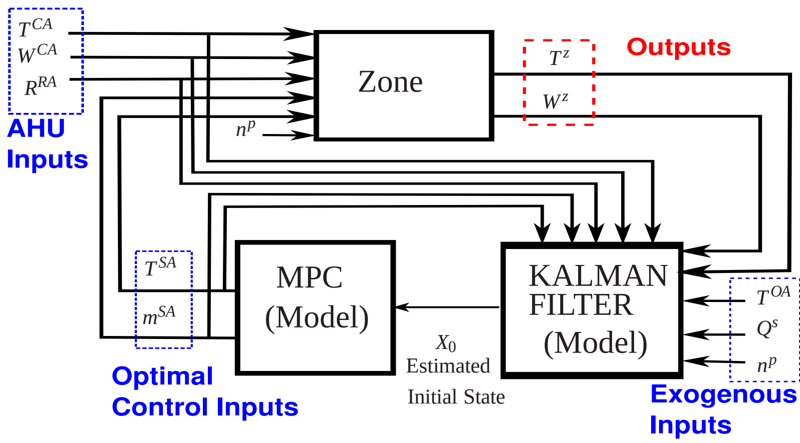

Figure 3: Schematic representation of MPC-based controllers (*MOBO* and *POBO*) implementation for a zone-level control.

of optimization, and (ii) a model of the zone hygro-thermal dynamics and the initial state of the hygro-thermal dynamics model. Prediction of T^{OA}, W^{OA} and Q^s is assumed available from weather forecasts. Obtaining occupancy prediction (n^p) is explained later when both the controllers are explained in detail. The model of building hygro-thermal dynamics and power used by the controllers is explained next. An EKF (Extended Kalman Filter)-based state observer is employed to estimate the initial state of the model at the start of the optimization.

The model of the zone thermal dynamics is constructed by combining elemental models of conductive interaction (RC networks) between two spaces separated by a solid surface such as a wall, as well as heat exchange due to the supply and return air. Humidity dynamics are derived from mass balance. The resulting model of the hygro-thermal dynamics of the zone is a set of coupled ODEs. We refer the reader to [16] for the details of the model. The continuous-time coupled ODE model is discretized using Euler's forward method to obtain a discrete-time model, which can be expressed as

$$\begin{aligned}\mathbf{T}(k+1) &= A\mathbf{T}(k) + Bv(k) + f(T^z(k), W^z(k), u(k), v(k)),\\ W^z(k+1) &= g(T^z(k), W^z(k), u(k), v(k)),\end{aligned} \tag{5}$$

where the vector $\mathbf{T}(k) \in \mathbb{R}^n$ consists of the zone temperature $T^z(k)$ and the temperatures of the nodes interior to the walls. The vector $v(k)$ consists of exogenous inputs, while the vector $u(k)$ consists of the control inputs ($m^{SA}(k), T^{SA}(k)$), i.e., $u(k) = [m^{SA}(k), T^{SA}(k)]^T$. The interior nodes come from the resistance-capacitance network used to model conduction. The parameters in such a model, in particular, the resistances and the capacitances of the walls and windows depend on their construction, and can be determined from properties listed in [22] and methods described in [23].

The total power consumption $P(k)$ at the time index k, which consists of fan power $P_F(k)$, reheating power $P_R(k)$, and conditioning power $P_U(k)$, is given by

$$P(k) \triangleq P_F(k) + P_U(k) + P_R(k). \tag{6}$$

We write the total power consumption as $P(u(k)$ when we want to emphasize its dependency on control inputs. Since the dynamics of the AHU are much faster than the thermal dynamics of a zone, we ignore the AHU dynamics. As a result, the power consumed in conditioning the air is a function of the instantaneous temperature and humidity. The fan power, the reheating power, and the conditioning power are given by

$$P_U = m^{SA}(h^{OA} - h^{CA}),\ P_F = \beta m^{SA}, P_R = m^{SA}(h^{SA} - h^{CA}), \tag{7}$$

where β is a system dependent constant. We refer the interested reader to [16] for details about the enthalpy terms h^{CA}, h^{OA}, and h^{SA}. The energy $E(k)$ consumed during the time $[(k-1)\Delta t,\ k\Delta t]$ is estimated as:

$$E(k) = \Delta t P(u(k)). \tag{8}$$

2.3.1. POBO *(Predicted Occupancy Based Optimal) Controller*

In this control algorithm, we assume that prediction of occupancy is available from the time index k to $k+K$, and the optimal control inputs for the next K time indices are obtained by solving the following optimization problem:

$$U^\star := arg \min_{U \in \mathbb{R}^{2K}} G(U), \tag{9}$$

where $U = [u(k)^T, \cdots, u^T(k+K)]^T$ and $G(U) = \sum_{i=k}^{k+K} \Delta t P(u(i))$, subject to the following constraints:

$$\left.\begin{array}{c} T^{occ}_{low} \leq T^z(i) \leq T^{occ}_{high},\ \text{if } n^p(i) \neq 0 \\ W^{occ}_{low} \leq W^z(i) \leq W^{occ}_{high},\ \text{if } n^p(i) \neq 0 \\ T^{CA} \leq T^{SA}(i) \leq T^{SA}_{high} \\ m^{SA}_p n^p(i) + \alpha m^{SA}_{low} \leq m^{SA}(i) \leq m^{SA}_{high} \end{array}\right\} \forall i = k, \ldots, k+K. \tag{10}$$

The first two constraints mean that the zone temperature and humidity ratio are allowed to vary in the range of $[T^{occ}_{low}, T^{occ}_{high}]$ and $[W^{occ}_{low}, W^{occ}_{high}]$, respectively, during the occupied time, while there are no constraints on the zone temperature and humidity ratio when the zone is not occupied. The third constraint is simply to take into account actuator capabilities, since the VAV box can only increase the temperature of the supply air above the conditioned air temperature. In addition, there is an upper bound on the amount by which the reheat coil can increase the temperature of the supply air. The fourth constraint means that there is a lower and upper bound on the flow rate entering the zone (m^{SA}). The lower bound on the flow rate is same as (2), while the upper bound m^{SA}_{high} reflects the maximum flow rate possible when the dampers in the VAV box are completely open.

As in the Measured Occupancy Based Setback controller, the choice of the design variables T^{occ}_{low}, T^{occ}_{high}, W^{occ}_{low}, W^{occ}_{high} involve a trade-off between energy savings and potential occupant discomfort. The greater the range that the temperature and humidity are allowed to vary in, both the potential energy savings and occupant discomfort are larger.

After solving the optimization problem (9)–(10) at time k, only the part of U^* corresponding to the current time index k is implemented.

2.3.2. MOBO *(Measured Occupancy Based Optimal) Controller*

The proposed *MOBO* controller is also an MPC-based control strategy similar to the POBO controller, but with an important difference. The MOBO controller only has access to instantaneous occupancy measurements, not predictions. Since MPC requires predictions of all exogenous inputs to perform the optimization involved in computing the control inputs, some form of occupancy prediction must be provided to the controller. Moreover, occupancy prediction decides the range in which the zone temperature is allowed to stay based on whether the zone is occupied or not. Since only occupancy measurements are available, the predicted occupancy for the next K time indices is assumed to be the same as the measured occupancy at the k-th time period: $n^p(i) = n^p(k), i \geq k$.

The control logic is divided into two modes: (i) Occupied, and (ii) Unoccupied, which are explained below in detail.

Occupied Mode: The controller operates in the occupied mode if the measured occupancy at the k-th time index, i.e., at the beginning of the time interval $[k\Delta t,\ (k+1)\Delta t]$, is at least 1. The optimal control inputs for the next K time indices are obtained by solving the optimization problem (9)–(10).

Unoccupied Mode: If the measured occupancy at the time index k, i.e., at the beginning of the k-th time period, is observed to be 0, then the controller operates in the unoccupied mode. At time k, the optimal control inputs for the next K time indices are obtained by solving the following optimization problem:

$$U^{\star} := arg \min_{U \in \mathbb{R}^{2K}} G(U), \tag{11}$$

subject to the following constraints:

$$\left.\begin{array}{c} T_{low}^{unocc} \leq T^{z}(i) \leq T_{high}^{unocc} \\ W_{low}^{unocc} \leq W^{z}(i) \leq W_{high}^{unocc} \\ \alpha m_{low}^{SA} \leq m^{SA}(i) \leq m_{high}^{SA} \\ T^{CA} \leq T^{SA}(i) \leq T_{high}^{SA} \end{array}\right\} \forall i = k, \ldots, k+K. \tag{12}$$

The reason for these constraints is the same as that explained previously. The constraints on the zone temperature and humidity ratio in the unoccupied mode, however, are chosen to be such that $[T_{low}^{unocc}, T_{high}^{unocc}] \supseteq [T_{low}^{occ}, T_{high}^{occ}]$, and $[W_{low}^{unocc}, W_{high}^{unocc}] \supseteq [W_{low}^{occ}, W_{high}^{occ}]$. This allows the controller greater flexibility in reducing energy consumption by letting the temperature and humidity ratio to vary in a wide range when the zone is unoccupied. The choice of the parameters for the unoccupied times also involves a trade-off. The farther they are from their counterparts for the occupied mode, greater is the energy savings potential, but also greater is the risk of occupant discomfort when occupancy changes.

Remark 1. *By choosing $\alpha > 1$, we ensure that for all the controllers the minimum flow rate during unoccupied times is greater than that prescribed by ASHRAE ventilation standard 62.1-2010 [12]. One reason for doing so is to make the resulting IAQ robust to the errors in occupancy measurements or predictions. It also makes the IAQ robust to the uncertainty in the measured flow rate and damper position. By ensuring good IAQ even during times when the zone is predicted to be unoccupied (whether correctly or not), we eliminate the problem of predicting the effect of control inputs on IAQ for the proposed controllers.*

3. Performance Metrics

The energy consumed by a controller C over a period ΔT is $E_C = \sum_{i=1}^{i=\frac{\Delta T}{\Delta t}} E_c(i)$, where $E_C(i)$ is the energy consumed by the controller C during the time $[(i-1)\Delta t,\ i\Delta t]$, calculated using (8). An energy related performance metric is the % *savings over the baseline controller*, which is defined as

$$\%\ \text{Savings} = \frac{E_{BC} - E_C}{E_{BC}}, \tag{13}$$

where E_C and E_{BC} are the energy consumed by the controller C and the baseline controller, respectively, over the same time period. The parameter ΔT is chosen as 24 *hrs* in this paper.

Two metrics are chosen for analyzing the thermal comfort related performance of the controllers: (i) Temperature Violation D_T, and (ii) Humidity Violation D_H, which are defined as

$$D_T = \left\{ \begin{array}{c} -T^z(t) + T^{occ}_{low}, \text{ if } T^z(t) < T^{occ}_{low} \text{ and } n^p(t) \neq 0 \\ T^z(t) - T^{occ}_{high}, \text{ if } T^z(t) > T^{occ}_{high} \text{ and } n^p(t) \neq 0 \\ 0, \text{ otherwise} \end{array} \right\},$$

$$D_H = \left\{ \begin{array}{c} -W^z(t) + W^{occ}_{low}, \text{ if } W^z(t) < W^{occ}_{low} \text{ and } n^p(t) \neq 0 \\ W^z(t) - W^{occ}_{high}, \text{ if } W^z(t) > W^{occ}_{high} \text{ and } n^p(t) \neq 0 \\ 0, \text{ otherwise} \end{array} \right\}.$$

These metrics measure the deviation of the zone temperature/humidity from the allowed range during occupied times. During the unoccupied times, both the temperature and humidity violations are considered 0 since there is no one in the zone. The *average temperature violation* ($D_T^\star$) and the *average humidity violation* ($D_H^\star$) during time period ΔT are defined as

$$D_T^\star = \frac{1}{\Delta T}\int_0^{\Delta T} D_T(t)dt \approx \frac{1}{L}\sum_{k=1}^{L} D_T(k), \quad D_H^\star = \frac{1}{\Delta T}\int_0^{\Delta T} D_H(t)dt \approx \frac{1}{L}\sum_{k=1}^{L} D_H(k). \tag{14}$$

where $L = \Delta T/\Delta t$. According to ASHRAE [22, Chapter 8], as long as people are wearing clothing of thermal resistance between $0.0775\, m^2K/W$ and $0.155\, m^2K/W$, doing primarily sedentary activity, and the air speed in the zone is less than $0.2\, m/s$, then ensuring that the temperature and humidity of the zone stays within certain range ensures thermal comfort of occupants (see Figure 6 in Section 4.2). Therefore, with appropriate choice of the parameters $T^{occ}_{(\cdot)} and W^{occ}_{(\cdot)}$, the temperature violation and the humidity violations defined above can be used as metrics for thermal comfort. Though Predicted Mean Vote (PMV) [22, Chapter 8] is a widely used metric to evaluate thermal comfort, it is a function of complex factors such as metabolism rate, clothes worn by the occupant, etc., which is quite difficult to compute in real-time. Therefore, we use temperature violation and the humidity violation to evaluate the thermal comfort, which are simpler to compute as well as more robust to assumptions made about the occupants.

Though IAQ is as important a concern as thermal comfort, if not more, we do not define a metric to measure "IAQ performance" of the controllers. Though CO_2 and volatile organic compounds contribute to poor IAQ, there is no well defined numerical measure to calculate IAQ [24]. Instead, we impose constraints on the minimum flow rate such that IAQ is maintained by all the controllers, even during unoccupied times (see also Remark 1).

4. Simulation Results

4.1. Model Calibration and Validation

Data from room 247 in Pugh Hall at the University of Florida, Gainesville, FL, USA is used to calibrate the model (5). The thermal capacitance per unit area and thermal resistance per unit area of external walls for all the zones obtained from [22, Chapter 39] are 369 $kJ/(m^2K)$ and 2.69 (m^2K/W), respectively. Measurements of the zone temperatures, supply air temperatures and flow rates are obtained from the Building Automation System at 10-minute intervals. The model is calibrated by tuning the total thermal resistance per unit area of the *internal walls* to minimize the error between the measured temperature and the predicted temperature of the zone. Data for a 48 hour long period (Jan 29-Jan 30, 2011) is used to calibrate the model. Since this time corresponds to a weekend, it is assumed that there are no occupants during this time. The comparison between the measured and predicted temperatures with the calibrated model are shown in Figures 4(a)–4(b). The validation data set (midnight Feb 5th through midnight of Feb 6th, 2011) also is from a weekend. It is clear from the figure that the temperature predictions by the model are close to the measured values.

4.2. Choice of parameters

Simulations are carried out for a model of three types of zones. All the zones have one external wall, one window and three internal walls. The internal walls are of the same type. It is assumed that the floor and the celling are perfectly insulated, and the window has negligible thermal capacitance. Each zone has the same window and same external wall construction, but the internal walls vary from zone to zone. A type-1 zone has internal walls of high thermal resistance and low thermal capacitance. The internal walls of a type-2 zone have low thermal resistance and high thermal capacitance. The internal walls of a type-3 zone have low thermal resistance and low thermal capacitance. We do not consider a zone with internal walls of high thermal capacitance and high thermal resistance, since this is unusual. The calibration and validation for the dynamic model of room 247, which

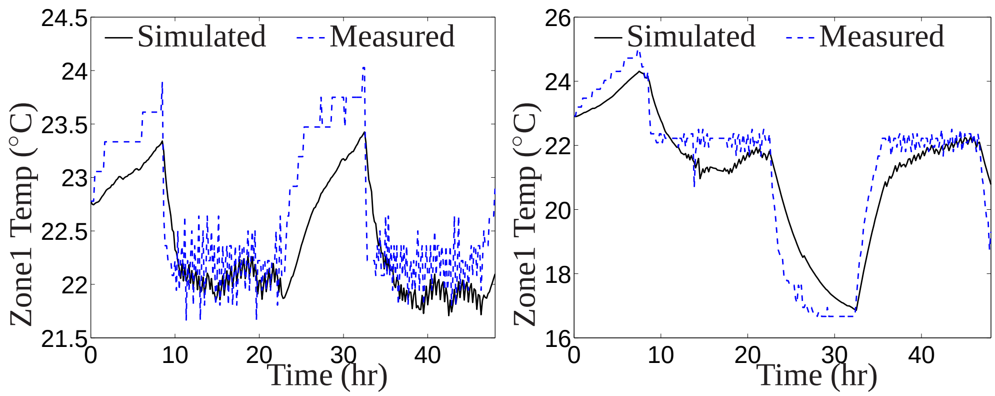

(a) Calibration

(b) Validation

Figure 4: Comparison of predicted and measured temperature in the room 247 in the Pugh Hall at the University of Florida, Gainesville, FL, USA.

is of zone type-3, is shown earlier in Section 4.1. The total thermal resistance and capacitance of internal walls of this zone are increased to construct models of type-1 and type-2 zones. The resulting resistance and capacitance values are shown in Table 1.

Table 1: Total thermal resistance and capacitance of the window and the walls (internal and external) of three types of zones.

Zone Type	Internal Wall		External Wall		Window
	Total Thermal Resistance ($\frac{m^2K}{W}$)	Total Thermal Capacitance ($\frac{kJ}{m^2K}$)	Total Thermal Resistance ($\frac{m^2K}{W}$)	Total Thermal Capacitance ($\frac{kJ}{m^2K}$)	Total Thermal Resistance ($\frac{m^2K}{W}$)
1	2.7	31	2.7	368	0.5
2	0.5	368			
3	0.5	31			

The boundaries of each zone that are separated from the zone by the internal walls are assumed to have a constant temperature of 22.2°C. The external wall separates a zone from outside weather, and three types of outside weather conditions are considered: cold, hot and pleasant. Figure 5 shows the temperature and humidity data for the cold (Jan 14, 2011), hot (Jul 31, 2011), and pleasant (Mar 16, 2011) days in Gainesville, FL, USA. "Pleasant weather" is non-standard terminology; we use it to denote weather that is neither neither hot nor cold.

The maximum flow rate for all the controllers is chosen as 0.125 kg/s. From

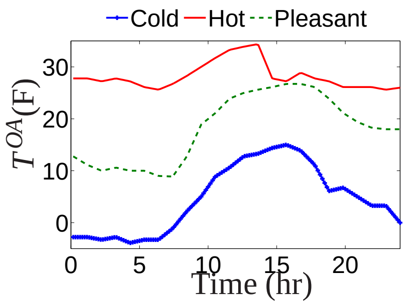

(a) Outside Temperature

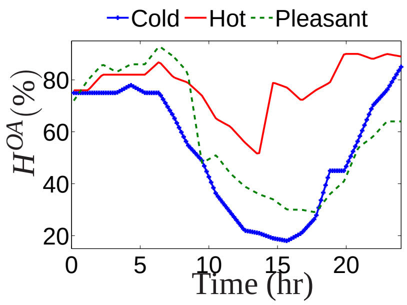

(b) Relative Humidity

Figure 5: Outside temperature (T^{OA}) and relative humidity (H^{OA}) for the cold (Jan 14, 2011), hot (Jul 31, 2011), and pleasant (Mar 16, 2011) day in Gainesville, FL, USA.

ASHRAE ventilation standard 62.1-2010 [12] requirements and return air ratio shown in Table 2, it turns out that $m_p^{SA} = 0.005\ kg/s$ and $m_{low}^{SA} = 0.015\ Kg/s$ and. These values are computed using (1), with $A_z = 25\ m^2$. For the *BL* controller, the Minimum Allowed Flow Rate is chosen as $0.05\ kg/s$, which corresponds to a designed occupancy of approximately 5 persons for the given zone. This is also the minimum flow rate that is currently being used by the existing control logic in room 247 of Pugh Hall. The IAQ factor of safety is chosen as $\alpha = 1.7$, so that the minimum flow rate for the *MOBS*, *MOBO*, and *POBO* controllers during the unoccupied mode turns out to be $\alpha m_{low}^{SA} = 0.0255\ Kg/s$. For the *BL* controller, the temperatures: RTG, HTG, and CLG are set to 21.8°C, 21.9°C, and 23.6°C, respectively, from $6:30$ a.m. to $10:30$ p.m. During the time $10:30$ p.m.–$6:30$ a.m., the temperatures: RTG, HTG, and CLG for the *BL* controller are chosen as 20.9°C, 21.1°C, and 24.4°C, respectively. This nighttime setback is currently used in the Pugh Hall.

Other design parameters are shown in Table 2. It is shown in table 2 that the set-points (RTG, HTG, and CTG) are changed symmetrically around the set-point T^{set} based on whether the zone is occupied or not. Since $T^{set} = \frac{RTG+CLG}{2}$ as mentioned in the Section 2.1, the desired set-point T^{set} stays constant.

The comfort envelope (which is defined by the constraints on the zone temperature and humidity ratio) used in this paper during the occupied and unoccupied times are shown in Figure 6. As long as certain assumptions on occupants clothing etc., are satisfied (see Section 3), thermal comfort is ensured if temperature and humidity ratio are maintained within the shaded regions shown in the figure. The constraints on the zone temperature and humidity ratio are chosen so that when they are met, the zone-climate meets the ASHRAE mandated conditions [22].

Table 2: The design parameters used in the various controllers.

Design Parameters									
Temperature Parameters									
T^{set} ($^\circ C$)	T^{SA}_{low} ($^\circ C$)	T^{SA}_{high} ($^\circ C$)	T^{unocc}_{RTG} ($^\circ C$)	T^{occ}_{RTG} ($^\circ C$)	T^{occ}_{low} ($^\circ C$)	T^{occ}_{high} ($^\circ C$)	T^{unocc}_{low} ($^\circ C$)	T^{unocc}_{high} ($^\circ C$)	T^{CA} ($^\circ C$)
22.8	12.8	30.0	20.9	21.8	21.9	23.6	21.1	24.4	12.8
Humidity and Other Parameters									
W^{unocc}_{low} ($\frac{g}{kg}$)	W^{occ}_{low} ($\frac{g}{kg}$)	W^{unocc}_{high} ($\frac{g}{kg}$)	W^{occ}_{high} ($\frac{g}{kg}$)	W^{CA} ($\frac{g}{kg}$)	K	Δt (min)	ΔT (hr)	R^{RA} (%)	n^p_d
7.4	7.4	10	10	7.4	3	10	24	40	5

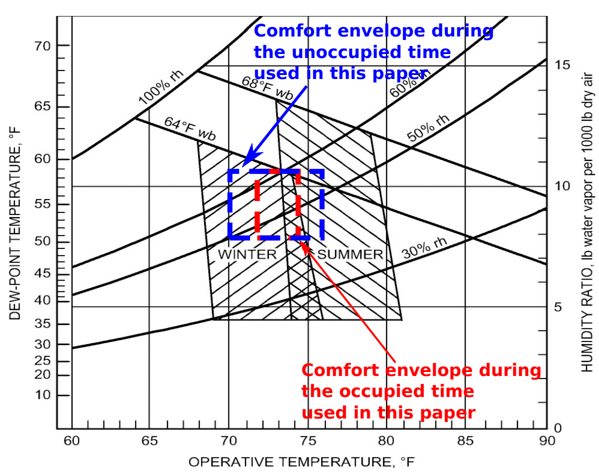

Figure 6: Comfort envelope specified in [22, Chapter 8], shown in the striped black area, and the envelope chosen in this paper during the occupied and unoccupied time, shown in dashed red and blue boxes, respectively.

4.3. Performance Comparison

In this section, we compare the performance of *BL*, *MOBS*, *MOBO*, and *POBO* control algorithms that are described in Section 2. Simulations are performed using MATLAB; while IPOPT [25] is used to solve the optimization problems for the *MOBO* and *POBO* control algorithms.

Each zone is occupied by a person from 8 : 00 a.m. to 12 : 00 p.m., and 1 : 00 p.m. to 5 : 00 p.m., everyday. The total daily energy consumption, average temperature violation, average humidity violation, and % savings over the baseline controller are shown in Table 3. We see from the table that depending on the zone type and outside weather, the *MOBS* and *MOBO* controllers result in 42–59% and 45–59% energy savings, respectively, over the baseline controller. Recall that both the *MOBS* and *MOBO* controllers use occupancy measurements; not predictions. The table also shows that the *POBO* controller—which requires occupancy predictions—can result in additional energy savings over the *MOBS* and *MOBO* controllers by an amount varying from 1% to 13%, again depending on zone type and weather. All the controllers have very small average temperature violation, and uniformly zero average humidity discomfort, irrespective of the type of zone or weather. Recall that IAQ is maintained at all times by the constraint on the minimum airflow rate. The results thus indicate that the energy savings from the proposed controllers are achieved with minimal impact on either thermal comfort or IAQ.

Table 3: Energy consumption, average temperature violation, average humidity violation, and % savings over a 24-hour period for single zone with various controllers. The three weather conditions are chosen for Gainesville, Fl, USA.

Zone Type	Control Scheme	Cold				Hot				Pleasant			
		E MJ	Savings %	$D_T^\star$ $^\circ C$	$D_H^\star$ $\frac{g}{kg}$	E MJ	Savings %	$D_T^\star$ $^\circ C$	$D_H^\star$ $\frac{g}{kg}$	E MJ	Savings %	$D_T^\star$ $^\circ C$	$D_H^\star$ $\frac{g}{kg}$
1	*BL*	93.4	-	0.007	0	179.4	-	0.003	0	78.3	-	0.004	0
	MOBS	53.5	42.7	0.026	0	97.5	45.6	0.014	0	41.5	47.0	0.018	0
	MOBO	50.6	45.8	0.006	0	93.7	47.7	0.004	0	39.0	50.1	0.006	0
	POBO	41.5	55.6	0	0	83.9	53.2	0	0	33.6	57.1	0	0
2	*BL*	86.8	-	0.005	0	173.7	-	0.001	0	72.2	-	0.003	0
	MOBS	42.1	51.4	0.016	0	79.6	54.2	0.001	0	29.9	58.6	0.008	0
	MOBO	40.2	53.7	0.004	0	80.0	54.0	0	0	30.2	58.2	0.001	0
	POBO	35.9	58.7	0	0	78.9	54.6	0	0	28.4	60.7	0	0
3	*BL*	91.9	-	0.007	0	178.4	-	0.002	0	76.8	-	0.004	0
	MOBS	49.7	45.9	0.023	0	92.2	48.3	0.013	0	38.4	49.9	0.021	0
	MOBO	47.3	48.5	0.006	0	90.0	49.5	0.002	0	36.2	52.8	0.005	0
	POBO	40.5	56.0	0	0	83.3	53.3	0	0	32.9	57.2	0	0

The energy savings come from the reduction of supply air flow rate and the increase in the allowable temperature range when the zone is not occupied. Reduction in the flow rate decreases fan-, conditioning-, and reheating-energy consumption. Increasing the allowable temperature range results in less reheating energy consumption at the VAV box, because the zone temperature is allowed to be lower during unoccupied times than what the baseline controller allows. For every zone, the total energy consumption is maximum during hot weather because more energy is consumed by the AHU to condition the hot and humid outside air than to condition the cold dry air. Among the three weathers, pleasant weather leads to the minimal energy consumption because apart from small conditioning energy requirements in such a weather, only a small amount of reheating energy is required. For a fixed zone, the fan energy is approximately same during all the weather conditions.

Given a controller and outside weather, we observe that $E_{zone\ type-2} < E_{zone\ type-3} < E_{zone\ type-1}$. Among the three types of zones, the type-2 zone consumes the least amount of energy. This is because the zone type-2 walls have low thermal resistance and high thermal capacitance, and the surrounding spaces of the zone that are separated by the internal walls are maintained at 22.2°C. The low thermal resistance helps maintain the zone temperature close to 22.2°C by fast transfer of energy through the internal walls from the surroundings, without the controller having to expend much energy. In addition, the high thermal capacitance causes the internal walls to store energy, which helps in maintaining the zone temperature. Type-1 zone consumes the maximum amount of energy because of the high thermal resistance and low thermal capacitance of the internal walls. The high thermal resistance does not allow easy transfer of energy from the surroundings through the internal walls, which, since they are maintained at 22.2°C, could have helped the control maintain the zone temperature around 22.2°C with less effort. In addition, the low thermal capacitance does not help in storing energy as in the case of type-2 and type-3 zone.

The average temperature violation $D_T^\star$ with either the *BL* controller or the *MOBS* controller is more than the average temperature violation with the *MOBO* controller for a fixed zone. It occurs because the *BL* and *MOBS* controllers wait for 10 minutes to turn on the heating/cooling mode. Among all the controllers, the average temperature violation is maximum for the *MOBS* controller. Since the *MOBS* controller increases the temperature range during the daytime if unoccupied, it takes some time for the zone temperature to come back to the allowable range when the zone becomes occupied again. However, the *BL* controller does not increase the allowable temperature range during the daytime even if it is not occupied. Therefore, the average temperature violation with the *MOBS* controller is more than that with the *BL* controller.

The simulation results shown above are for the case when occupancy varies between 0 and 1, and for the Gainesville, FL, USA location. We have also conducted simulations for three more cases: 1) occupancy varies between 0 and 3; location: Gainesville, FL, USA, ii) occupancy varies between 0 and 1, location: Phoenix, AZ, USA, and ii) occupancy varies between 0 and 3, location: Phoenix, AZ, USA. The weather days for Phoenix are chosen to be the same as those for Gainesville; see Section 4.2. Very similar % savings over the baseline controller, and average temperature/humidity violations, are obtained for all the cases. The results are not shown due to space limits.

MPC vs. feedback, with occupancy measurements: While the *MOBS* controller uses simple rule-based feedback control based on temperature and occupancy measurement, the *MOBO* controller is a much more complex MPC-based control scheme that requires prediction of relevant state variables and exogenous signals. Yet, the results above show that the performance of the *MOBS* and *MOBO* controllers are quite similar, both in terms of energy savings and thermal comfort. This is due to the fact that without occupancy prediction, the MPC-based controller cannot really take advantage of its powerful optimization algorithm. If predictions are available, the optimization routine may be able to reduce the airflow and let the temperature "float", thus saving energy, and then bring it back up right before the zone is about to be occupied. In the absence of such prediction, the MPC-controller can only do what a *well-designed* feedback controller will also do, that is, set back the zone temperature when the zone is unoccupied, but not too much so that it can be changed quickly when occupancy changes, and maintain some minimum airflow to ensure good IAQ.

One concern during the initial stages of the research was that the slow thermal dynamics of a typical zone, along with the limitations of the actuators, will make the response of the closed-loop control system too slow to ensure occupant comfort during the transition period when occupancy changes. However, the results reported here show that this concern can be mitigated by appropriate choice of the temperature and humidity bands.

Utility of occupancy prediction: One surprising observation is that the additional % savings of the *POBO* controller over the *MOBS* and *MOBO* controllers are small, 1–13%, even though it uses occupancy predictions while the other two only uses measurements. One could expect that since occupancy predictions are available, the controller can turn the airflow rate quite low, thereby resulting in large energy savings. The small additional savings are due to the ventilation requirements. ASHRAE ventilation standard 62.1-2010 [12] requires a certain amount of outside air that depends on the floor area even when the zone is unoccupied. For a medium sized office with a small design occupancy (1-5 people), the resulting minimum flow rate turns out to be a significant fraction of the nominal airflow rate during occupied periods. Savings would be higher if the ventilation rates during the unoccupied times were to be smaller than what are prescribed by current standards. For instance, the older ASHRAE ventilation standard 62.1-2001 [26] did not require outside air supply during unoccupied times. We performed simulations with a minimum airflow rate of 0 during unoccupied times. In that case the savings with the *POBO* controller increases up to about 80% over the baseline controller. That is, the additional savings possible with occupancy prediction—compared to occupancy measurement—is now about 40%.

5. Discussion and Future Work

We examine how a controller performance is affected by its complexity, where the goal of the controller is to minimize energy consumption while maintaining comfort level in a zone in a commercial building with a variable-air-volume HVAC system. For that purpose, we propose three control strategies of varying complexity and requiring varying fidelity of information: *MOBS*, *MOBO* and *POBO*. The performance of the proposed controllers are compared through simulations with that of a conventional baseline controller. The baseline controller uses temperature feedback but not real-time occupancy information. In contrast, the proposed *MOBS* and *MOBO* controllers require occupancy measurements, and the *POBO* controller requires occupancy predictions. While *MOBS* controller is a feedback control algorithm, the *MOBO* and *POBO* controllers are MPC-based algorithms. Simulation results show that all three controllers lead to substantial improvement in energy savings (about 50% on average depending on zone type, weather, climate, design occupancy, etc.,) with negligible impact on IAQ or thermal comfort.

The study shows that even a simple feedback-based algorithm can perform as well as an MPC-based algorithm, if only occupancy measurements are available. In the absence of occupancy prediction, MPC simply sets back the zone temperature to save energy; while the feedback controller is designed to mimic that behavior as well. Another conclusion of the study is that the additional savings with an MPC-based control that uses occupancy prediction - over one that only uses measurements - is small. The small additional savings are due to the restriction on the minimum airflow, which come from current ASHRAE ventilation standard 62.1-2010 [12]. If lower ventilation rates are allowed during unoccupied times, as earlier standards did, it is possible to save significantly more energy by using occupancy prediction; assuming of course that such predictions can be obtained. However, with the current standards, MPC-based control does not provide significant energy savings over much simpler feedback-based schemes, even when occupancy predictions are available. At the same time, considerable effort is required in developing/calibrating/validating dynamic models required by the controller, and the numerical optimization involved make the controller computationally complex. Thus, the use of MPC-based zone-climate control of existing VAV systems may not be economically justified. A feedback controller is the most appropriate control algorithm to be used at the zone level since it is simple, computationally fast, requires minimal investment in hardware and software, and delivers energy savings quite similar to that of much more complex control algorithms.

The study shows that occupancy measurement is a key component of energy-efficient zone-climate control. When the zone is designed for a single person, such as an office, a motion detector can be used to measure occupancy. However, if the zone is designed for multiple occupants, measuring occupancy is not trivial. Development of reliable yet inexpensive occupancy measurement technology will greatly facilitate the deployment of occupancy-based energy-efficient building control. The controllers proposed in this paper have some robustness to errors in occupancy measurements due to their higher-than-needed minimum airflow. A detailed study of their performance with varying levels of measurement error is planned as part of future work.

There are several additional avenues for further exploration. All the proposed control algorithms require choice of several parameters, which involve a trade-off between energy savings and potential discomfort. This trade-off needs to be more carefully examined to determine a set of guidelines on how to choose these parameters. Implementing the proposed controllers in a real building is required to verify the simulation results. Work on experimental verification is ongoing. In this paper, we have assumed that a zone consists of single room. The control algorithms can be extended in a straightforward manner to be applicable to a zone that consists of multiple rooms. Their performance in such a scenario, though, needs to be studied.

In this paper, the AHU control inputs (such as conditioned air temperature, flow rate and return air dampers position) are assumed constant and treated as exogenous inputs. It is possible that through a coordinated control among the AHU and multiple zones, more energy efficiency can be achieved than what can be achieved by keeping the AHU controller and zone-level controllers independent. This is another interesting direction to pursue.

References

[1] US EIA, Annual energy review (October 2011).

[2] W. Angel, HVAC Design Sourcebook, McGraw-Hill Engineering, 2012.

[3] M. Brambley, D. Hansen, P. Haves, D. Holmberg, S. McDonald, K. Roth, P. Torcellini, Advanced sensors and controls for building applications: Market assessment and potential *R&D* pathways, Tech. rep., Pacific Northwest National Laboratory (April 2005).

[4] S. Goyal, H. Ingley, P. Barooah, Zone-level control algorithms based on occupancy information for energy efficient buildings, in: American Control Conference (ACC), 2012, pp. 3063–3068.

[5] V. Erickson, M. Carreira-Perpinan, A. Cerpa, OBSERVE: Occupancy-based system for efficient reduction of HVAC energy, in: Information Processing in Sensor Networks (IPSN), 2011, pp. 258–269.

[6] Y. Ma, G. Anderson, F. Borrelli, A distributed predictive control approach to building temperature regulation, in: American Control Conference (ACC), 2011, pp. 2089–2094.

[7] A. Aswani, N. Master, J. Taneja, D. Culler, C. Tomlin, Reducing transient and steady state electricity consumption in HVAC using learning-based model-predictive control, Proceedings of the IEEE 100 (1) (2012) 240–253.

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

[9] F. Oldewurtel, A. Parisio, C. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, M. Morari, Use of Model Predictive Control and Weather Forecasts for Energy Efficient Building Climate Control, Energy and Buildings 45 (2012) 15–27.

[10] S. Afshari, S. Mishra, A. Julius, F. Lizarralde, J. Wen, Modeling and feedback control of color-tunable LED lighting systems, in: American Control Conference (ACC), 2012, pp. 3663–3668.

[11] US EIA-Department of Energy, CBECS detailed tables (2003).

[12] ASHRAE, ANSI/ASHRAE standard 62.1-2010: Ventilation for acceptable air quality (2010).

[13] C. Liao, P. Barooah, An integrated approach to occupancy modeling and estimation in commercial buildings, in: American Control Conference, 2010, pp. 3130–3135.

[14] T. Teixeira, G. Dublon, A. Savvides, A survey of human sensing: Methods for detecting presence, count, location, track and identity, www.eng.yale.edu/enalab/publications/human_sensing_enalabWIP.pdf, unpublished.

[15] S. Qin, T. Badgwell, A survey of industrial model predictive control technology, Control Engineering Practice 11 (2003) 733–764.

[16] S. Goyal, P. Barooah, A method for model-reduction of nonlinear building thermal dynamics of multi-zone buildings, Energy and Buildings 47 (2012) 332–340.

[17] ASHRAE, The ASHRAE handbook - HVAC applications (SI Edition) (2011).

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

[19] A. Persily, A. Musser, S. Emmerich, M. Taylor, Simulations of Indoor Air Quality and Ventilation Impacts of Demand Controlled Ventilation in Commercial and Institutional Buildings, U.S. Dept. of Commerce, Technology Administration, National Institute of Standards and Technology, 2003.

[20] Y. Tachwali, H. Refai, J. Fagan, Minimizing HVAC energy consumption using a wireless sensor network, in: Industrial Electronics Society, 2007. IECON 2007. 33rd Annual Conference of the IEEE, 2007, pp. 439–444.

[21] V. Dhummi, D. Demetriou, H. Palanthandalam-Madapusi, H. Khalifa, C. Isik, Robust occupancy-based distributed demand control ventilation, International Journal of Ventilation 9 (5) (2011) 359–369.

[22] ASHRAE, The ASHRAE handbook fundamentals (SI Edition) (2005).

[23] M. Gouda, S. Danaher, C. Underwood, Building thermal model reduction using nonlinear constrained optimization, Building and Environment 37 (2002) 1255–1265.

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

[25] A. W*ä*chter, L. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, Mathematical Programming 106 (2006) 25–57.

[26] ASHRAE, ANSI/ASHRAE standard 62.1-2001: Ventilation for acceptable air quality (2001).