Applied Energy 106 (2013) 209–221

Contents lists available at SciVerse ScienceDirect

Applied Energy

journal homepage: www.elsevier.com/locate/apenergy

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

HIGHLIGHTS

- ▶ Investigates performance of energy-efficient building climate-control algorithms as a function of their complexity.
- ▶ Simulations are performed with multiple types of zone, occupancy, outside weather, and geographic location.
- ▶ Shows that significant amount of energy savings without sacrificing comfort are possible with only occupancy measurements.
- ▶ Feedback controller performs as well as optimal controller if only occupancy measurements are used.
- ▶ Availability of occupancy prediction leads to small energy savings.

ARTICLE INFO

Article history:
Received 5 August 2012
Received in revised form 6 January 2013
Accepted 14 January 2013
Available online 27 February 2013

Keywords:
Occupancy-based zone-climate control
Model predictive control
Energy-efficient buildings
Building thermal dynamics

ABSTRACT

We propose several control algorithms and 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 compute 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 occupancy predictions. 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 complexity of the MPC algorithms compared to the feedback control algorithm, we conclude that feedback control is the more suitable one for occupancy based zone-climate control. A related conclusion is that the difficulty of obtaining occupancy predictions does not commensurate with the resulting benefits; though these benefits are a strong function of ventilation standards.
© 2013 Elsevier Ltd. All rights reserved.

1. Introduction

Buildings are one of the main consumers of energy 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 replacing HVAC equipment with more efficient ones, it requires substantial investment to retrofit an existing building with improved HVAC equipment [4]. In contrast, improving control algorithms (that operate the HVAC system) to achieve higher efficiency is far more cost effective. Indeed, a number of recent papers have focused on improving energy efficiency in buildings through advanced control algorithms [5–11]. In this paper, we examine control algorithms that use occupancy information to control the climate of individual zones with reduced energy use compared to conventional control algorithms that do not use such information.

We limit ourselves to commercial buildings with variable-air-volume (VAV) systems. More than 30% of the commercial building floor space in the United States is served by VAV systems [12]. In a VAV system, a building is divided into a number of "zones", where a zone can be a single room or a collection of rooms. The flow rate

☆ This work has been supported by the National Science Foundation by Grants CNS-0931885 and ECCS-0955023.
* Corresponding author.
E-mail addresses: siddgoya@ufl.edu (S. Goyal), ingley@ufl.edu (H.A. Ingley), pbarooah@ufl.edu (P. Barooah).

0306-2619/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.apenergy.2013.01.039

Nomenclature

CLG	cooling set-point
D_H	humidity violation
$\overline{D}_H$	average humidity violation
D_T	temperature violation
$\overline{D}_T$	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
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 flow rate)
T	temperature
T^{set}	desire 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
m	mass flow rate
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

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)

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 determined are the mass flow rate and temperature of the supply air.

The conventional control strategies used at the VAV box use real-time temperature measurements but do not use real-time occupancy measurements. In this paper, "occupancy" is used to denote the number of people in a space. 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), a minimum airflow rate is maintained, which 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 [13]. This widely used control logic is inefficient in terms of energy use since even during unoccupied times it maintains the indoor environment as if it was occupied. Energy efficiency of HVAC systems can be improved by changing the indoor climate in response to occupancy change. Demand control ventilation (DCV) seeks to do so by changing the supply air flow rate based on measured or estimated occupancy [14].

However, with real-time occupancy measurements, it should be possible to do more to reduce energy use than merely controlling ventilation. For example, energy use can be further reduced by letting the temperature vary during unoccupied times in a wider range than that during occupied times. 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 temperature to come back to a comfortable range when the zone becomes occupied again. The same goes for humidity and IAQ. On the other hand, if future occupancy is known then one might be able to avoid such a scenario by bringing the temperature back up in time, which requires predictions from a dynamic model of the zone temperature and humidity. Thus, the dynamics of temperature, humidity, and IAQ have to be taken into account in designing such control algorithms. Moreover, the controller should also have some robustness to error in occupancy measurements/predictions.

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. With more information (prediction vs. measurement), 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 question of sensors to measure occupancy or models to predict occupancy is relevant to control systems seeking to use such measurements and/or predictions. Real-time occupancy measurements can be obtained from presence sensors—such as PIR and ultrasound sensors—that are inexpensive and work well in small office spaces where the nominal occupancy value is one [15]. For spaces occupied by more than one person, measuring occupancy is not trivial. Efforts in developing novel occupancy measurement technology are carried out by several researchers; see [16,17,15] and references therein. There is also considerable recent effort in developing models and algorithms for predicting future occupancy in real-time [18–21,6]. These predictions are distinct from—and expected to be more accurate than—predetermined occupancy schedules. In this paper, we assume that sensors and/or algorithms to obtain real-time occupancy measurement and/or prediction are available.

The rest of paper is organized as follows. In the remainder of the section, we discuss the related literature and the contribution of the paper. The models of a zone's hygro-thermal dynamics (i.e., the dynamics of temperature and humidity) and power consumption are described in Section 2. The proposed control algorithms and a baseline controller (a conventional controller used commonly in commercial buildings) are described in Section 3. Section 4 describes performance metrics related to thermal comfort and energy savings. Section 5 provides a description of the parameters chosen for the simulation study. Simulation results and their implications are discussed in Section 6. Section 7 concludes the paper with a discussion of the results and ways to extend this work.

1.1. Relation to prior work

A number of papers have investigated control algorithms that use occupancy information (either measurements or predictions), and compared their energy consumption to that of conventional controllers. Some of these are simple rule-based controllers while others are quite complex. By "conventional controller", we mean a controller that is commonly used in existing commercial buildings. Such a controller does not use occupancy measurements; though it may use predefined occupancy schedules. These control algorithms consist of "if-else" logics for higher level decision making and PID loops for lower level control, such as set point maintenance. Some of the controllers examined in the literature—as well as in practice—are rule-based controllers, while other rely on solving an optimization problem in real-time to make decisions. A "rule-based controller", like a conventional controller, uses "if-else" logics for decision making. They may use other forms of information—such as occupancy measurements—that the conventional controllers do not use. The optimization based controllers typically use a receding horizon optimization approach, which is also known as model predictive control (MPC).

A number of papers have proposed 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 [22–25]. The controller in [22] uses occupancy measurements to turn off the HVAC system, while the controllers in papers [23–25] 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 predictions) compared to rule-based control it may also lead to more energy savings. The paper [6] compares several rule-based controllers that use various types of occupancy information: two use occupancy predictions while one uses measurements of 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.

The papers [7–10] compare MPC-based controllers with conventional controllers. The paper [26] examines use of MPC in the design stage to decide on HVAC control strategies. The MPC controllers mentioned in [7–10] use occupancy predictions while the conventional controllers use only 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 a controller that is less complex than MPC by using occupancy measurements, which are easier to obtain than occupancy predictions.

The paper by Oldewurtel et al. [27] is closest in spirit to our work; it also examines the effect of occupancy information fidelity on controller performance. However, it does not compare performance of the MPC strategies with simpler occupancy measurement based feedback controllers. The MPC strategies in [27] have a long prediction horizon (in the range of days). However, in this paper, one of the MPC controller has a short prediction horizon (30 min) while the other MPC controller has a long prediction horizon (24 h). The HVAC systems considered in [27] are not VAV systems, and the actuation available to the controllers include blind positions. In contrast, we limit ourselves to the HVAC systems with VAV boxes, where the only actuation available to the controller are the supply air mass flow rate and supply air temperature. These systems are far more prevalent in the US than systems where blind positions can also be commanded. Another difference is that [27] does not consider humidity in the problem formulation while in this paper humidity constraints are incorporated into control design.

In summary, 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 measurements vs. occupancy predictions is not clear from the prior work. Since predictions are much more difficult to obtain than measurements, it is particularly useful to know their relative value. Though [6] compares performance of the rule-based controllers that use occupancy predictions with that of the feedback controller, the feedback controller uses only presence/absence measurements but not occupancy measurements.

1.2. Contributions

In this paper we examine the performance of a conventional controller that does not have occupancy measurements or predictions to that of three proposed controllers that use varying degree of occupancy information and are of varying degree of computational complexity. The purpose is twofold. The first is to determine how much energy savings can be achieved by using occupancy information over conventional control schemes, and to examine how the savings vary with the fidelity of occupancy information: real-time measurements vs. perfect predictions of future occupancy. The reason for considering perfect occupancy predictions—though infeasible in practice—is to establish the limits of achievable performance. The second aim is to determine how the savings achieved by these controllers vary with the complexity of the controllers, from a simple feedback scheme (aka, "rule-based") to complex optimization based schemes that uses a model of the zone's temperature and humidity dynamics. In short, we examine trade-off between energy savings achieved and the information requirements/computational complexity of the control algorithm in a unified manner. This is the key difference between our work and much of prior work. Another difference is that the papers mentioned in Section 1.1 that propose MPC-based controllers do not take humidity into account, while humidity is taken into account as a part of thermal comfort constraints in the MPC schemes proposed here.

The results of the study provide guidelines to control system engineers in choosing an appropriate control scheme that strikes a balance between cost (complexity) and benefit (energy efficiency). Our study shows that when real time occupancy measurements are available, substantial energy savings can be obtained by a simple feedback control scheme, and that the additional savings from using complex MPC-based controllers are small. Thus, investing in implementation of complex MPC-based schemes may not be economically justifiable. Another useful result of the study is that the additional savings that can be obtained by using occupancy predictions—even when perfect predictions are available—are small. This shows that the substantial effort required in obtaining occupancy predictions is of questionable benefit. Interestingly, the reason turns out to be the recent ASHRAE ventilation standards. If lower ventilation is allowed, as earlier standards did, then the value of occupancy predictions increases substantially.

A preliminary version of this work appeared in [5]; which compares feedback, MPC, and baseline controllers. There are several significant differences between [5] and this paper. The baseline conventional controller used here is more energy-efficient than the one used in [5]. The feedback controller used in [5] 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 [5] 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 [13]. 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 [5]. While Ref. [5] 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. System description and models

A schematic of a typical multi-zone commercial building with a VAV-based HVAC system is shown in Fig. 1. Part of the air removed from the zones (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. The task of a zone-climate control algorithm is to determine 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. A conceptual representation of a control algorithm that operates the VAV box is also shown in Fig. 1.

For simulation studies of a control scheme, we need a model of the hygro-thermal (humidity and temperature) dynamics of the zone as well as a model of power consumption as a function of control signals and exogenous inputs. The model of the thermal dynamics is constructed by combining RC (resistance–capacitance) networks models of heat transfer between two spaces separated by a solid surface (such as a wall) with that of heat exchange due to supply and return air. Humidity dynamics are derived from mass balance. We refer the reader to [28] for the details of the model, which presents a model for the general case of multiple zones. The same model, but for the special case of a single zone, is used in this paper. Here we only mention some of the salient features of the model that are relevant to the subsequent discussion.

The resulting model of the hygro-thermal dynamics of the zone is a set of coupled ODEs, and thus can be expressed as

$$\begin{aligned} \dot{\mathbf{T}}(t) &= f_c(\mathbf{T}(t), u(t), v(t)), \\ \dot{W}(t) &= g_c(\mathbf{T}(t), W(t), u(t), v(t)), \end{aligned} \tag{1}$$

where $u(t) = [m^{SA}(t), T^{SA}(t)]^T \in \mathbb{R}^2$ is the control input (command), while the exogenous inputs vector $v(t)$ consists of outside temperature, outside humidity ratio, solar heat gain, and occupancy, i.e., $v(t) = [T^{OA}(t), W^{OA}(t), Q^s(t), n^p(t)]^T \in \mathbb{R}^4$. We use a 3R2C model for each wall of the zone and a 1R model for a window. For a zone consisting of four walls and one window, the resulting hygro-thermal model of the zone has ten states: nine for temperature ($\mathbf{T} \in \mathbb{R}^9$) and one for humidity ($W \in \mathbb{R}$). One of the temperature states is the temperature of the zone. The additional temperature states can be thought of as temperatures inside the walls, they come from the internal states of the RC network models of walls. The simula-

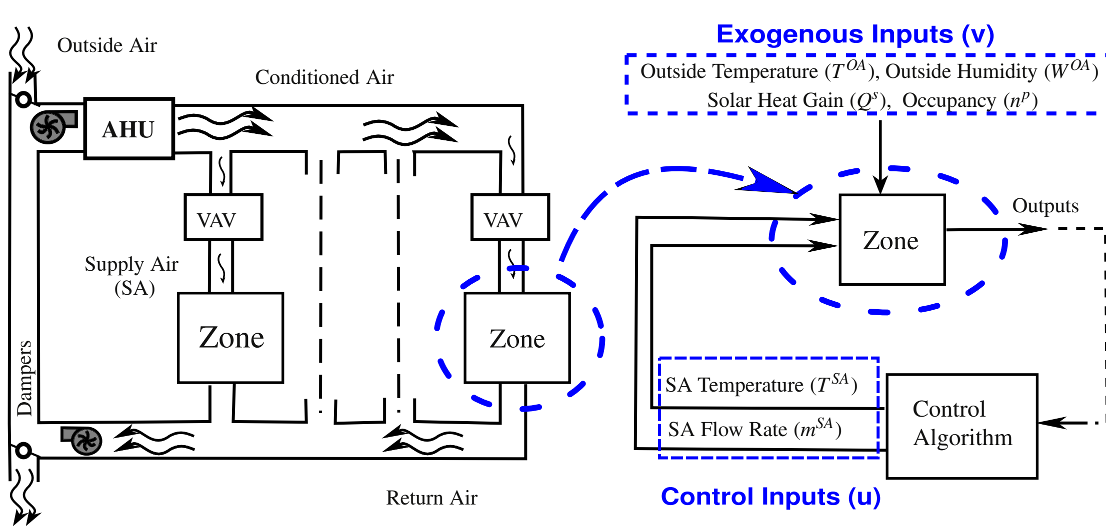

Fig. 1. Generic scheme for the implementation of a zone-level control algorithm.

tion studies reported in Section 6 are based on such a model. The parameters of the model, in particular, the resistances and the capacitances of the walls and windows depend on their geometry and the material they are constructed from. The continuous-time coupled ODE model (1) is discretized using Euler's forward method to obtain a discrete-time model, which can be expressed as

$$\begin{aligned}\mathbf{T}(k+1) &= f_d(\mathbf{T}(k), W(k), u(k), v(k)),\\ W(k+1) &= g_d(\mathbf{T}(k), W(k), u(k), v(k)),\end{aligned} \tag{2}$$

where $x(k) = x(t)|_{t=k\Delta t}$ for any signal $x(t)$, with Δt being the sampling interval and k being the discrete time index.

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

$$P(k) \triangleq P_F(k) + P_U(k) + P_R(k). \tag{3}$$

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

$$P_U = m^{SA}(h^{OA} - h^{CA}), \quad P_F = \beta m^{SA}, \quad P_R = m^{SA}(h^{SA} - h^{CA}), \tag{4}$$

where β is a system dependent constant. We refer the interested reader to [28] 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:

$$E(k) = \Delta t P(u(k)). \tag{5}$$

3. Control algorithms

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), which vary in complexity and the required information as shown in Table 1. Recall that the actuation signals that the controllers have to decide are the flow rate and temperature of the supply air.

3.1. BL (Baseline) controller

Among the conventional control logics used at the VAV boxes to maintain IAQ and temperature in a zone, we choose the dual maximum [29, chapter 47] as the *baseline controller*. Even though the single maximum control [29, 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 Fig. 2. If the zone temperature stays below the "Re-heating Set-Point (RTG)" for more than 10 min, the re-heating mode is turned on. Similarly, if the zone temperature remains above the "Cooling Set-Point (CLG)" for more than 10 min, the cooling mode is turned on. If the zone temperature stays between RTG and "Heating Set-Point (HTG)" for more than 10 min, the heating mode is turned on. If the zone temperature stays between HTG and CLG for more than 10 min, the dead-band mode is turned on. In the re-heating mode, the supply air temperature is set to maximum possible value $\left(T^{SA}_{high}\right)$, 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

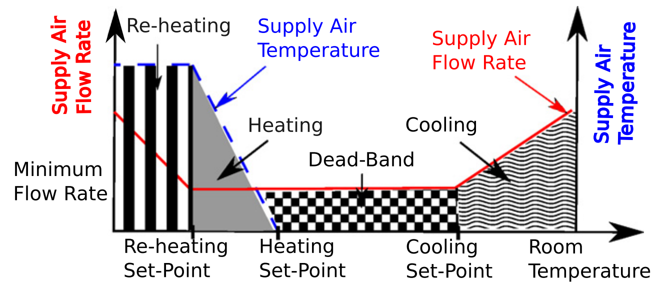

Fig. 2. Schematic representation of the baseline control strategy ("dual maximum") used at the VAV terminal boxes of commercial buildings.

$$\begin{aligned}&\text{Minimum allowed flow rate} = m^{SA}_p n^p_d + \alpha m^{SA}_{low},\\ &\text{where } m^{SA}_p = m^{OA}_p/(1-R^{RA}),\ m^{SA}_{low} = m^A_z A_z/(1-R^{RA}).\end{aligned} \tag{6}$$

When $\alpha = 1$, these calculations yield the minimum airflow requirements specified by ASHRAE ventilation standard 62.1-2010 [13]. Since the baseline controller does not use occupancy measurements, the minimum allowed flow rate is calculated using the designed occupancy n^p_d, which is assumed constant. 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 (6). 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.

Table 1
Overview of the control algorithms in terms of the amount of information required and complexity. The overall complexity of the control algorithms increases in the order. (1) *BL*, (2) *MOBS*, (3) *MOBO*, and (4) *POBO*.

	Control algorithms	Type of occupancy information required	Model required	Computation required	Overall complexity
1	*BL*	None	No	Low	Low
2	*MOBS*	Measurements	No	Low	Medium
3	*MOBO*	Measurements	Yes	High	High
4	*POBO*	Predictions	Yes	High	Very High

3.2. MOBS (Measured Occupancy Based Setback) 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 3.1, except for two key differences. First, the minimum allowed flow mentioned in (6) is calculated based on the measured occupancy instead of the design occupancy, which is expressed as

$$\text{Minimum allowed flow rate at time } t = m_p^{SA} n^p(t) + \alpha m_{low}^{SA}, \tag{7}$$

where $n^p(t)$ is the occupancy measured at time t, and m_p^{SA}, m_{low}^{SA} are computed using (6). Second, the temperature set-points are determined based on whether the zone is occupied or not:

$$\left.\begin{array}{l} RTG(t) = T_{RTG}^{unocc} \\ HTG(t) = T_{low}^{unocc} \\ CTG(t) = T_{high}^{unocc} \end{array}\right\} \text{ if } n^p(t) = 0, \quad \left.\begin{array}{l} RTG(t) = T_{RTG}^{occ} \\ HTG(t) = T_{low}^{occ} \\ CTG(t) = T_{high}^{occ} \end{array}\right\} \text{if } n^p(t) \neq 0. \tag{8}$$

The choice of design variables T_{RTG}^{unocc}, T_{RTG}^{occ}, T_{low}^{unocc}, T_{low}^{occ}, T_{high}^{unocc}, T_{high}^{occ} involves a trade-off between energy savings and thermal comfort. Clearly, the range $\left[T_{low}^{occ}, T_{high}^{occ}\right]$ 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

$$\left[T_{low}^{occ}, T_{high}^{occ}\right] \subseteq \left[T_{low}^{unocc}, T_{high}^{unocc}\right], \tag{9}$$

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_{RTG}^{unocc}, T_{RTG}^{occ}) far from the heating set-points (T_{low}^{unocc}, T_{low}^{occ}) 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.

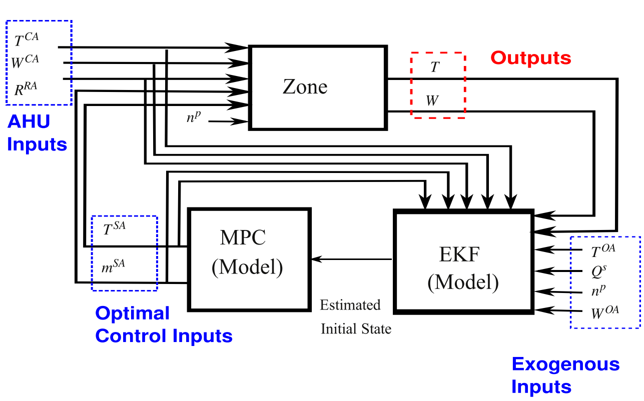

Fig. 3. Schematic representation of MPC-based controllers (*MOBO* and *POBO*) implementation for a zone-level control.

3.3. 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 Fig. 3. Both the controllers compute the control input $u(k)$ 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 input $\upsilon(k)$ over the time horizon of optimization and (ii) a model of the zone hygro-thermal dynamics as well as its initial state. Predictions of T^{OA}, W^{OA}, and Q^s (part of $\upsilon(k)$) are assumed available from weather forecasts. Obtaining occupancy predictions is explained later when both the controllers are explained in detail. The models for power consumption and hygro-thermal dynamics used by the controller are the ones presented in Section 2. An EKF (Extended Kalman Filter)-based state observer is employed to estimate the initial state of the model during optimization.

3.3.1. MOBO (Measured Occupancy Based Optimal) controller

The proposed *MOBO* controller is an MPC-based control strategy. In this control algorithm, we assume that the instantaneous occupancy measurements are available at the time index k. Since MPC requires predictions of all exogenous inputs to perform the optimization involved in computing the control inputs, some form of occupancy predictions over the prediction horizon K must be provided to the controller. Moreover, occupancy predictions decide 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 kth time period: $n^p(i) = n^p(k)$, $i \geqslant k$.

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 kth 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 following optimization problem:

$$U^{\star} := \underset{U}{\operatorname{argmin}}\, G(U), \tag{10}$$

where $U = [u^T(k), \ldots, u^T(k+K)]^T \in \mathbb{R}^{2(K+1)}$ and $G(U) = \sum_{i=k}^{k+K} E(i)$, subject to the following constraints:

$$\left.\begin{array}{l} T_{low}^{occ} \leqslant T(i) \leqslant T_{high}^{occ}, \\ W_{low}^{occ} \leqslant W(i) \leqslant W_{high}^{occ}, \\ T^{CA} \leqslant T^{SA}(i) \leqslant T_{high}^{SA} \\ m_p^{SA} n^p(i) + \alpha m_{low}^{SA} \leqslant m^{SA}(i) \leqslant m_{high}^{SA} \end{array}\right\} \quad \forall i = k, \ldots, k+K.$$

The first two constraints mean that the zone temperature and humidity ratio are allowed to vary in the range of $\left[T_{low}^{occ}, T_{high}^{occ}\right]$ and $\left[W_{low}^{occ}, W_{high}^{occ}\right]$, respectively. 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 (7), 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 (10) at time k, only the part of U^* corresponding to the current time index k is implemented.

Unoccupied mode: If the measured occupancy at the time index k, i.e., at the beginning of the kth 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:

$$U^{\star} := \arg\min_{U} G(U), \tag{11}$$

subject to the following constraints:

$$\left.\begin{array}{l} T^{unocc}_{low} \leqslant T(i) \leqslant T^{unocc}_{high} \\ W^{unocc}_{low} \leqslant W(i) \leqslant W^{unocc}_{high} \\ \alpha m^{SA}_{low} \leqslant m^{SA}(i) \leqslant m^{SA}_{high} \\ T^{CA} \leqslant T^{SA}(i) \leqslant T^{SA}_{high} \end{array}\right\} \quad \forall i = k, \ldots, k+K.$$

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 $\left[T^{unocc}_{low}, T^{unocc}_{high}\right] \supseteq \left[T^{occ}_{low}, T^{occ}_{high}\right]$, and $[W^{unocc}_{low}, W^{unocc}_{high}] \supseteq \left[W^{occ}_{low}, W^{occ}_{high}\right]$. 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 during the occupied mode, greater is the energy savings potential, but also greater is the risk of occupant discomfort when occupancy changes.

3.3.2. POBO (Predicted Occupancy Based Optimal) controller

The proposed *POBO* controller is also an MPC-based control strategy similar to the *MOBO* controller, but with two important differences. First, the *POBO* controller has a long prediction horizon as opposed to the *MOBO* controller. Second, the *POBO* controller has access to occupancy predictions from the time index k to $k + K$. The purpose of assuming availability of such predictions is to establish the limit of achievable performance. The optimal control inputs for the next K time indices are obtained by solving the following optimization problem:

$$U^{\star} := \arg\min_{U} G(U), \tag{12}$$

subject to the following constraints:

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

The first two constraints mean that the zone temperature and humidity ratio are allowed to vary in the range of $\left[T^{occ}_{low}, T^{occ}_{high}\right]$ and $\left[W^{occ}_{low}, W^{occ}_{high}\right]$, respectively, during the occupied time, while there are no constraints on the zone temperature and humidity ratio when the zone is not occupied. The last two constraints are same as the last two constraints of the optimization problem (10). Once the optimization problem (12) is solved at time k, only the part of U^* corresponding to the current time index k is implemented.

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 [13]. One reason for doing so is to make the resulting IAQ robust to the errors in occupancy measurements or predictions. It also makes 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.

Remark 2. The optimization problems (10) and (11), or (12) that the MPC controllers need to solve are infeasible when the constraints on the zone temperature/humidity are violated. This happens when the occupancy changes from 0 to 1 and the zone temperature/humidity during that time is either too high or too low. In such a scenario, it is not possible to bring the zone temperature back to comfortable range even with the maximum heating or cooling capabilities of the controller. However, the zone temperature/humidity discomfort can be minimized if the controller delivers maximum heating or maximum cooling based on the zone temperature. Therefore, we choose the control inputs at the time index k, in case of an infeasible solution, as:

$$\begin{aligned} & m^{SA}(k) = m^{SA}_{high}, \\ & T^{SA}(k) = T^{SA}_{high} \text{ if } T(k) < T^{occ}_{low}, \\ & T^{SA}(k) = T^{CA} \text{ if } T(k) > T^{occ}_{high}. \end{aligned} \tag{13}$$

Remark 3. If the optimization problems (10) and (11), or (12) are infeasible due to any reason other than the reasons mentioned in Remark 2, which can be due to numerical issues, etc., then the control inputs at the time index k are chosen as:

$$\begin{aligned} & m^{SA}(k) = m^{SA}_{high}, \\ & T^{SA}(k) = T^{set}. \end{aligned}$$

4. 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 (5). An energy related performance metric is the *% savings over the baseline controller*, which is defined as

$$\%\text{Savings} = \frac{E_{BC} - E_C}{E_{BC}}, \tag{14}$$

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 h 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

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

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

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* ($\overline{D}_T$) and the *average humidity violation* ($\overline{D}_H$) during time period ΔT are defined as

$$\overline{D}_T = \frac{1}{\Delta T}\int_0^{\Delta T} D_T(t)dt \approx \frac{1}{L}\sum_{k=1}^{L} D_T(k), \quad \overline{D}_H = \frac{1}{\Delta T}\int_0^{\Delta T} D_H(t)dt \approx \frac{1}{L}\sum_{k=1}^{L} D_H(k). \tag{15}$$

where $L = \Delta T/\Delta t$. According to ASHRAE [30, chapter 8], as long as people are wearing clothing of thermal resistance between 0.0775 m^2 K/W and 0.155 m^2 K/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 Fig. 5. 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) [30, chapter 8] is a widely used metric to evaluate thermal comfort, it is a function of complex factors such as metabolism rate and clothes worn by the occupant, 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 [31]. 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).

5. Choice of parameters

5.1. Parameters of the zone's dynamic model

Simulations are carried out for three different models of a zone, which are referred to as zone type 1–3. They correspond to a zone with a given geometry but different R,C parameter values. The dimensions of the zone is 5.2 m $\times$ 4.8 m $\times$ 2.7 m, with an window of area 2.8 m^2. The area of the external wall (the wall separates the zone to the outside) is 11.24 m^2. All the other walls of the zone are called internal walls. A hallway and two adjacent rooms are separated from the zone with internal walls. These dimensions are that of a typical zone in Pugh Hall, a modern office building in the University of Florida campus that has a VAV-based HVAC system. Each zone type has the same window and same external wall construction, but the internal walls' *R*, *C* values vary from one zone type to another. 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 resistance and capacitance values are shown in Table 2.

To decide the *R*, *C* parameters, we first examine construction information for the external walls in the Pugh Hall. It turns out that the external walls of all the rooms in the building are of same construction. The thermal capacitance per unit area and thermal resistance per unit area of external walls are obtained from ASHRAE handbook [30, chapter 39], which are shown in Table 2. The resistance values of internal walls obtained from handbook, however, are less reliable due to uncertainties such as air leakage through cracks. External walls are constructed to be highly insulating, so they do not suffer from this level of uncertainty. To obtain more reliable parameter values of internal walls, we perform model calibration. We pick a zone in the Pugh Hall (room 247) that is special in its layout; it is completely enclosed by other rooms and thus consists only of internal walls. This means that the room experiences no solar heat gain, which makes the model calibration easier since solar heat gain is one of the exogenous inputs and is not easy to measure accurately. Once the model of this room is calibrated and validated (to be described soon), the resulting *R*, *C* parameters of its walls are chosen as those for the internal walls of the nominal zone model. This fixes the model acceptable of the nominal zone, which we call zone type-3. Not surprisingly, the *R*, *C* parameters of the internal walls of this zone are observed to be on the low side when compared to those of the external wall. We then construct two other zone types by assigning to their internal walls various combinations of *R*, *C* values between these low values and the high values obtained previously for the external wall. The purpose for testing multiple zone types is to test the variation of controller performance with variation in zone's thermal interaction with the surroundings.

For model calibration of room 247 in the Pugh Hall, measurements of the zone temperature, supply air temperatures and flow rates, and temperatures of the surrounding spaces are obtained from the Building Automation System at 10-min intervals. The model is calibrated by tuning the total thermal resistance per unit area of the walls to minimize the error between the measured temperature and the predicted temperature of the zone. Data for a 48 h long period (January 29–30, 2011) is used to calibrate the model. Since this time corresponds to a weekend, we assume that there are no occupants—and therefore no occupant induced heat loads—during this time. The comparison between the measured and predicted temperatures with the calibrated model are shown in Fig. 4a–b. The validation data set (midnight February 5th through midnight of February 6th, 2011) also is from a weekend. The figure shows that the temperature predictions by the model are within 1° of the measured values. There are several degree Celsius variation of temperature in the room; the acceptable air temperature varies from 13° to 32° depending on the operation of the reheat coil in the VAV box. In addition, the temperature sensors have an uncertainty of around 0.5°. In view of these variations, a 1° variation between measurement and prediction is deemed acceptable for the purpose of calibration.

5.2. Controller parameters

The maximum flow rate for all the controllers is chosen as 0.125 kg/s. From ASHRAE ventilation standard 62.1-2010 [13] requirements and return air ratio shown in Table 3, it turns out that $m^{SA}_p = 0.005$ kg/s and $m^{SA}_{low} = 0.015$ Kg/s and. These values are computed using (6), 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 five persons for the given zone. This is also the minimum flow rate that is currently being used in a single occupancy room in the Pugh Hall. The IAQ factor of safety is chosen as α = 1.7, so that the minimum flow rate for the *MOBS*, *MOBO*, and *POBO* controllers during the unoccupied mode turns out to be $\alpha m^{SA}_{low} = 0.0255$ Kg/s. For the *BL* controller, the tempera-

Table 2
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 $\left(\frac{m^2\,K}{W}\right)$	Total thermal capacitance $\left(\frac{kJ}{m^2\,K}\right)$	Total thermal resistance $\left(\frac{m^2\,K}{W}\right)$	Total thermal capacitance $\left(\frac{kJ}{m^2\,K}\right)$	Total thermal resistance $\left(\frac{m^2\,K}{W}\right)$
1	2.7	31			
2	0.5	368	2.7	368	0.5
3	0.5	31			

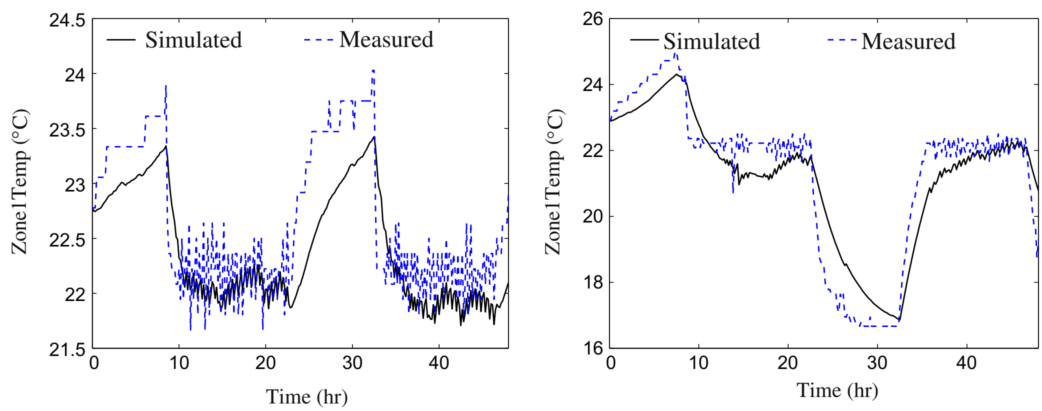

(a) Calibration (b) Validation

Fig. 4. Comparison of predicted and measured temperature in room 247, Pugh Hall, both for the data set used for calibrating the model and another data set that is not used for calibration.

Table 3
The design parameters used in the various controllers.

Design parameters									
T^{set} (°C)	T^{SA}_{low} (°C)	T^{SA}_{high} (°C)	T^{unocc}_{RTG} (°C)	T^{occ}_{RTG} (°C)	T^{occ}_{low} (°C)	T^{occ}_{high} (°C)	T^{unocc}_{low} (°C)	T^{unocc}_{high} (°C)	T^{CA} (°C)
Temperature parameters									
22.8	12.8	30.0	20.9	21.8	21.9	23.6	21.1	24.4	12.8
$W^{unocc}_{low}\left(\frac{g}{kg}\right)$	$W^{occ}_{low}\left(\frac{g}{kg}\right)$	$W^{unocc}_{high}\left(\frac{g}{kg}\right)$	$W^{occ}_{high}\left(\frac{g}{kg}\right)$	$W^{CA}\left(\frac{g}{kg}\right)$	*K* (*MOBO*, *POBO*)	Δt (min)	ΔT (h)	R^{RA} (%)	n^p_d
Humidity and other parameters									
7.4	7.4	10	10	7.4	30, 86400	1	24	40	5

tures: RTG, HTG, and CLG are set to 21.8 °C, 21.9 °C, and 23.6 °C, respectively, from 6:30 am to 10:30 pm During the time 10:30 pm–6:30 am, 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 3. It is shown in Table 3 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 3.1, the desired set-point T^{set} stays constant.

The prediction horizons used by the *MOBO* and *POBO* controllers to solve the optimization problem are 30 min and 24 h, respectively. To reduce computational complexity, the control commands are solved for and updated every 10 min. The time step Δt used for discretization of the model is 1 min.

The constraints on the zone temperature and humidity ratio used in this paper during the occupied and unoccupied times are shown in Fig. 5. As long as certain assumptions on occupants clothing, etc., are satisfied (see Section 4), 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 the constraints are met, the zone-climate meets the ASHRAE mandated conditions [30].

6. Simulation results and implications

We now compare the performance of *BL*, *MOBS*, *MOBO*, and *POBO* control algorithms through simulations. Simulations are performed using MATLAB; while IPOPT [32] is used to solve the optimization problems for the *MOBO* and *POBO* control algorithms.

The boundaries of each zone that are separated by the internal walls are assumed to have a constant temperature of 22.2°. Three types of outside weather conditions are considered: cold, hot and pleasant. Fig. 6 shows the temperature and humidity data for the cold (January 14, 2011), hot (July 31, 2011), and pleasant (March 16, 2011) days in Gainesville, FL, USA. “Pleasant weather” is non-standard terminology; we use it to denote weather that is neither hot nor cold. During simulations, the initial conditions for all the temperatures states (i.e., zone temperature and temperature corresponding to the interior of the walls) and zone humidity ratio are taken as 22.2 °C and 0.009, respectively.

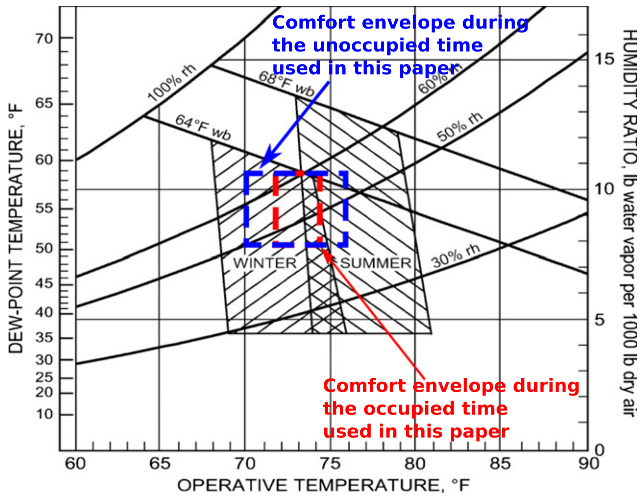

Fig. 5. Comfort envelope specified in [30, chapter 8], shown in the striped black area, and the envelope chosen in this paper during the occupied and unoccupied time, shown in dashed red and blue boxes, respectively.

The following occupancy profile is used during the initial set of simulations for all the controllers: the zone is occupied by a person from 8:00 am to 12:00 pm, and 1:00 pm to 5:00 pm. For a controller that uses real-time measurements of occupancy, instantaneous value of occupancy is provided to the controller, which simulates an occupancy sensor with no error. The entire occupancy profile for the 24 h period is provided to the *POBO* control algorithm ahead of time, which simulates a perfect occupancy prediction capability.

The total daily energy consumption, average temperature violation, average humidity violation, and % savings over the baseline controller are shown in Table 4. 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—that 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 violation, 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.

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 $E_{zone\ \text{type-2}} < E_{zone\ \text{type-3}} < E_{zone\ \text{type-1}}$. That is, the type-2 zone consumes the least amount of energy among the three types of zones. The reason for this is that 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°, which is close to the average of the allowable temperature range. The low thermal resistance helps maintain the zone temperature close to 22.2° 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°, could have helped the control maintain the zone temperature around 22.2° 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 $\overline{D}_T$ 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 min 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

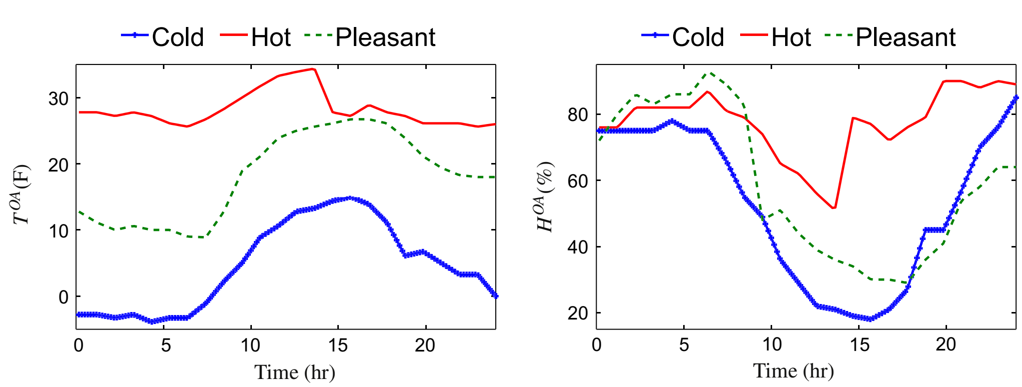

(a) Outside Temperature (b) Relative Humidity

Fig. 6. Outside temperature (T^{OA}) and relative humidity (H^{OA}) for the cold (January 14, 2011), hot (July 31, 2011), and pleasant (March 16, 2011) day in Gainesville, FL, USA.

Table 4
Energy consumption, average temperature violation, average humidity violation, and % savings over a 24-h 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 (%)	$\overline{D}_T$ (°C)	$\overline{D}_H$ $\left(\frac{g}{kg}\right)$	E (MJ)	Savings (%)	$\overline{D}_T$ (°C)	$\overline{D}_H$ $\left(\frac{g}{kg}\right)$	E (MJ)	Savings (%)	$\overline{D}_T$ (°C)	$\overline{D}_H$ $\left(\frac{g}{kg}\right)$
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

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 discussed 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: (i) occupancy varies between 0 and 3; location: Gainesville, FL, USA, (ii) occupancy varies between 0 and 1, location: Phoenix, AZ, USA, and (iii) occupancy varies between 0 and 3, location: Phoenix, AZ, USA. Weather data for Phoenix on January 14, July 31, and March 16 of 2011 are used in the simulations with Phoenix weather, as in the simulations with Gainesville weather. A 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 the space limit.

6.1. Implications

MPC vs. feedback, with occupancy measurements: While the *MOBS* controller uses simple rule-based feedback control that uses temperature and occupancy measurements, the *MOBO* controller is a much more complex MPC-based control scheme that requires predictions 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 predictions, 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 predictions, 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 predictions: 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 perfect occupancy predictions while the other two only use measurements. One could expect that since occupancy predictions are available, the controller can turn the airflow rate quite low during unoccupied times, resulting in large energy savings. The reason that this does not happen is due to the ventilation requirements. ASHRAE ventilation standard 62.1-2010 [13] 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 [33] 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 predictions—compared to occupancy measurements—is now about 40%. These results indicate that the effort required to obtain future occupancy predictions in real-time, say, by using occupancy models [6,20], are of questionable benefit with the current standards. If ventilation standards change again, it might be worthwhile to invest in developing techniques for occupancy prediction.

7. 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 and IAQ 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 measurements. 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 to the controller. In the absence of occupancy predictions, 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 controller that uses perfect occupancy predictions—over one that only uses real-time occupancy measurements—are small. The small additional savings are due to the restriction on the minimum airflow during the unoccupied times, which come from current ASHRAE ventilation standard 62.1-2010 [13]. The minimum ventilation rate requirements during the unoccupied time are to take out the water vapor released by the equipments, furniture, etc., which contribute a significant amount to the minimum ventilation requirements during the occupied times for a medium sized office. If lower ventilation rates are allowed during unoccupied times, as earlier standards did, it is possible to save significantly more energy by using occupancy predictions; assuming of course that such predictions can be obtained. In that case, the significant energy savings will come from the reduction in the energy consumed to condition the air at AHU especially during the unoccupied time, since the minimum allowed flow rates are low during that time. However, with the current standards, MPC-based control does not provide significant energy savings over much simpler feedback-based schemes, even when perfect occupancy predictions ahead of time are available. One should note that considerable effort is required in developing/calibrating/validating dynamic models required by the MPC-based controllers, and the numerical optimization involved makes 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, approximately 50% energy savings over a conventional controller can be obtained by using occupancy measurements. 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, obtaining accurate instantaneous occupancy is not trivial. Work on developing sensors and algorithms for occupancy measurements is being carried out by several researchers and commercial entities [6,20,17], but more work is needed. Development of reliable yet inexpensive occupancy measurement technology will greatly facilitate the deployment of occupancy-based energy-efficient building control. All occupancy measurements are error free in the simulations reported here. Some robustness to occupancy measurement errors are present in the proposed controllers due to the higher-than-needed minimum airflow. Still, a detailed study of their performance with varying levels of measurements error is required, and 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. 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] Angel W. HVAC design sourcebook. McGraw-Hill Engineering; 2012.
[3] Brambley M, Hansen D, Haves P, Holmberg D, McDonald S, Roth K, et al. Advanced sensors and controls for building applications: market assessment and potential *R&D* pathways, Tech. rep., Pacific Northwest National Laboratory; April 2005.
[4] Rahman M, Rasul M, Khan M. Energy conservation measures in an institutional building in sub-tropical climate in australia. Appl Energy 2010;87:2994–3004.
[5] Goyal S, Ingley H, Barooah P. Zone-level control algorithms based on occupancy information for energy efficient buildings. In: American control conference (ACC); 2012. p. 3063–8.
[6] Erickson V, Carreira-Perpinan M, Cerpa A. OBSERVE: occupancy-based system for efficient reduction of HVAC energy. In: Information processing in sensor networks (IPSN); 2011. p. 258–69.
[7] Ma Y, Anderson G, Borrelli F. A distributed predictive control approach to building temperature regulation. In: American control conference (ACC); 2011. p. 2089–94.
[8] Aswani A, Master N, Taneja J, Culler D, Tomlin C. Reducing transient and steady state electricity consumption in HVAC using learning-based model-predictive control. Proc IEEE 2012;100(1):240–53.
[9] Siroky J, Oldewurtel F, Cigler J, Privara S. Experimental analysis of model predictive control for an energy efficient building heating system. Appl Energy 2011;88:3079–87.
[10] Oldewurtel F, Parisio A, Jones C, Gyalistras D, Gwerder M, Stauch V, et al. Use of model predictive control and weather forecasts for energy efficient building climate control. Energy Build 2012;45:15–27.
[11] Afshari S, Mishra S, Julius A, Lizarralde F, Wen J. Modeling and feedback control of color-tunable LED lighting systems. In: American control conference (ACC); 2012. p. 3663–8.
[12] US EIA-Department of Energy. CBECS detailed table; 2003.
[13] ASHRAE. ANSI/ASHRAE standard 62.1-2010: ventilation for acceptable air quality; 2010.
[14] Lawrence TM, Braun JE. Calibrated simulation for retrofit evaluation of demand-controlled ventilation in small commercial buildings. ASHRAE Trans 2007:227–40.
[15] Subramany R, Liao C, Barooah P. Performance comparison of sensing systems for building occupancy measurement. Energy Build Rev. <http://humdoi.mae.ufl.edu/prabirbarooah/PBpublication.html>.
[16] Kamthe A, Jiang L, Dudys M, Cerpa A. SCOPES: smart cameras object position estimation system. In: Proceedings of the 6th European conference on wireless sensor networks (EWSN). Cork (Ireland): Springer-Verlag; 2009. p. 279–95.
[17] Teixeira T, Dublon G, Savvides A. A survey of human sensing: methods for detecting presence, count, location, track and identity; in press. <http://www.eng.yale.edu/enalab/publications/human_sensing_enalabWIP.pdf>.
[18] Page J, Robinson D, Morel N, Scartezzini J-L. A generalised stochastic model for the simulation of occupant presence. Energy Build 2008;40:83–98.
[19] Meyn S, Surana A, Lin Y, Oggianu SM, Narayanan S, Frewen TA. A sensor-utility-network method for estimation of occupancy distribution in buildings. In: 48th IEEE conference on decision and control; 2009.
[20] Liao C, Barooah P. An integrated approach to occupancy modeling and estimation in commercial buildings. In: American control conference; 2010. p. 3130–5.
[21] Liao C, Lin Y, Barooah P. Agent-based and graphical modeling of building occupancy. J Build Perf Simulat, Spec Issue Model Occupants' Presence Behav Part II 5; 2012. http://dx.doi.org/10.1080/19401493.2010.531143.
[22] Agarwal Y, Balaji B, Dutta S, Gupta R, Weng T. Duty-cycling buildings aggressively: the next frontier in HVAC control. In: Information processing in sensor networks (IPSN); 2011. p. 246–57.
[23] Persily A, Musser A, Emmerich S, Taylor M. 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.
[24] Tachwali Y, Refai H, Fagan J. Minimizing HVAC energy consumption using a wireless sensor network. In: 33rd Annual conference of the IEEE industrial electronics society, 2007. IECON 2007; 2007. p. 439–44.
[25] Dhummi V, Demetriou D, Palanthandalam-Madapusi H, Khalifa H, Isik C. Robust occupancy-based distributed demand control ventilation. Int J Ventilat 2011;9(5):359–69.

[26] Petersen S, Svendsen S. Method for simulating predictive control of building systems operation in the earl stages of building design. Appl Energy 2011;88:4597–606.
[27] Oldewurtel F, Sturzenegger D, Morari M. Importance of occupancy information for building climate control. Appl Energy 2013. <http://www.sciencedirect.com/science/article/pii/S0306261912004564>.
[28] Goyal S, Barooah P. A method for model-reduction of nonlinear building thermal dynamics of multi-zone buildings. Energy Build 2012;47:332–40.
[29] ASHRAE. The ASHRAE handbook – HVAC applications (SI Edition); 2011.
[30] ASHRAE. The ASHRAE handbook fundamentals (SI Edition); 2005.
[31] Bohanon H. Good IAQ practices. ASHRAE J 2012;54:106–7.
[32] Wächter A, Biegler L. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math Program 2006;106:25–57.
[33] ASHRAE. ANSI/ASHRAE standard 62.1-2001: ventilation for acceptable air quality; 2001.