Energy-Efficient Control of Under-Actuated HVAC Zones in Buildings

Jonathan Brooks and Prabir Barooah

***Abstract*— We study control-oriented methods of improving energy efficiency of "under-actuated" heating, ventilation, and air-conditioning systems in commercial buildings where multiple rooms are served by a single variable air volume box. Two novel control algorithms are studied: one is a rule-based, feedback controller that uses real-time occupancy measurements, and the other is based on model predictive control (MPC) that uses predictions of occupancy and other exogenous inputs. In an under-actuated zone, flow rate and temperature of ventilation air to individual rooms in a zone cannot be independently controlled. These have not been studied in the literature but are quite common in commercial buildings. Our study finds that there is potential for significant energy savings relative to conventional controllers for under-actuated zones—even when occupancy in the rooms is vastly different—when independent heating is available for each room. Energy efficiency of simple, rule-based controllers that use occupancy measurements was found to be highly dependent on choice of temperature constraints during unoccupied times. Depending on this choice, MPC may or may not perform better than rule-based control.**

I. Introduction

The impact of buildings' heating, ventilation, and air-conditioning (HVAC) systems on the total energy consumption in the United States is well documented [1]. Reducing the energy consumption of building HVAC systems can have significant economic benefits. Upgrading older and less efficient HVAC systems with more efficient ones, however, can be prohibitively costly [2]. Another alternative to improve energy efficiency is to use smarter, more efficient control algorithms for indoor climate control. This has been the subject of many papers in the literature and is also the topic of this paper.

Climate control algorithms must maintain indoor air quality (IAQ) and thermal comfort. The American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) recommends a minimum outside air flow rate [3, ASHRAE Standard 62.1] and "comfortable" ranges for temperature and humidity [4, ASHRAE Standard 55]. When a room is unoccupied, however, some of these constraints may be relaxed. Most conventional HVAC control strategies maintain climate according to predefined occupancy schedules instead of actual occupancy—resulting in inefficient energy use. Energy efficiency of HVAC systems can be improved by basing the control actions on real-time measurements or predictions from models [5, 6].

This paper deals with developing such control algorithms for HVAC systems with variable air volume (VAV) boxes. VAV systems serve approximately one third of the commercial building floorspace in the U.S. [7]. In a VAV system, cooling is provided centrally by an air handling unit (AHU). Each VAV box then modulates the airflow into its zone to maintain the zone's temperature and IAQ.

JB and PB are with the MAE dept., University of Florida, Gainesville, FL 32611, USA.

This research is supported by NSF awards CPS-0931885 and ECCS-0925534.

Recent papers have been limited to "fully actuated" zones in which a single room is controlled by a single AHU or VAV box. In this paper, we focus on the "under-actuated" case in which a single VAV box serves multiple rooms—meaning the supply airflow to each room cannot be controlled independently. Under-actuated rooms are as common as fully actuated rooms in U.S. commercial buildings—if not more—and are an overwhelming majority in residential buildings. In the scenario we consider, heating is supplied by heaters in each room; the distribution of air into the individual rooms is determined by the duct sizes. This HVAC process is shown in Figure 1.

Fig. 1. Diagram of a standard AHU-VAV HVAC system. Some air from each zone is mixed with outside air. The mixed air is then conditioned and supplied to the zones. M represents the distribution of air from the VAV box to each room.

Many papers have examined algorithms to reduce energy consumption compared to conventional controllers that do not use occupancy measurements. Here, "conventional controllers" use simple "if-else" logic for higher-level decision-making. PID loops are then used for lower-level control actions such as set point maintenance. Rule-based controllers use additional information—such as occupancy measurements—that conventional controllers do not. The literature shows rule-based control (RBC) using occupancy measurements can lead to significant energy savings compared to conventional control [8–13]. Recently, a great deal of focus has been placed on model predictive control (MPC) for building HVAC control. MPC may result in more energy savings due to the optimization of control decisions, and its ability to maintain constraints is attractive, but it requires a model of the building's hygro-thermal dynamics and predictions of exogenous inputs, which are nontrivial to obtain [5, 14, 15]. It can also require significant computational power [16]. Many papers have compared the performances of MPC-based HVAC controllers and conventional controllers and reported potential for significant energy savings with MPC [13, 16–20].

Few papers, however, have compared MPC with occupancy-based RBC for HVAC control. The paper [13] has found that, indeed, RBC can perform almost as well as

MPC while being much simpler. In addition, the sensitivity of MPC and RBC to various design parameters has not been studied extensively. The effect of certain parameters in MPC robustness of MPC to exogenous input information have been studied in [20] and [21], respectively, but these references did not examine the effects of several key parameters (such as constraints on temperatures) on the relative performances of these control schemes.

In this paper, we propose two types of room-level, climate control algorithms for the under-actuated case. One is a rule-based controller that uses occupancy measurements (similar to the Measured Occupancy-Based Setback (MOBS) controller in [13]). The other is an MPC algorithm that uses occupancy and weather predictions. We then conduct a thorough comparison of the performance of the proposed controllers with a commonly used baseline controller that uses predefined occupancy schedules.

A motivation of this study is the high cost of implementing MPC, which has challenging requirements such as occupancy and model predictions. This necessitates determination of whether the additional benefit is worth the cost. Earlier, limited work in this direction seemed to suggest that the benefit of the additional cost is small and may even be negligible [13]. This study shows that, for some scenarios at least, the cost of MPC may be justified by the higher savings.

II. Dynamic Model

The zone we analyze consists of two adjacent rooms controlled by the same VAV box. Each room has one exterior wall. A standard resistor-capacitor (RC) model was used for the zone thermal dynamics with nonlinear terms for enthalpy exchange (see [22]). The resistance and capacitance values were chosen to be representative of commercial buildings in the eastern United States that experience both hot and cold weather. The heaters are able to supply any heating up to their heating capacity.

The states of the coupled ODE model consist of the temperature and humidity of each room and the temperature of each internal wall (including floor and ceiling). We have three control inputs: supply air flow rate from the VAV box and each room's heater command. Exogenous inputs include the condition of air from the AHU, internal heating (from occupants and electrical equipment), solar heat gain, ambient weather conditions, and the surrounding room temperatures.

The instantaneous power consumption is the sum of the power consumptions of the supply fan, the AHU cooling coils, and the heater in each room. The power consumptions of the fan and AHU are defined in (1) and (2), respectively:

$$P_{fan}(i) = \alpha\big(\dot{m}_{SA}(i)\big)^3 \tag{1}$$

$$P_{AHU}(i) = \dot{m}_{SA}(i)|(h_{MA}(i) - h_{SA}(i)| \tag{2}$$

where α is a coefficient of proportionality; $h_{MA}(i)$ and $h_{SA}(i)$ are the specific enthalpies of the mixed and supply air, respectively, at time i; and $\dot{m}_{SA}(i)$ is the mass flow rate of the supply air. The heating power P_h is simply the heating supplied to the rooms. The coefficient α used in the numerical studies is estimated from data collected from Pugh Hall at the University of Florida [?].

III. Control Schemes

A. Baseline Controller

We compare our proposed algorithms against a baseline (BL) algorithm commonly used in conventional HVAC systems. This controller assumes the rooms are occupied during the day and unoccupied during the night. During daytime, the controller maintains each room's air temperature between a lower and upper bound. During nighttime, the bounds are loosened.

This controller has two main modes of operation: cooling mode and non-cooling mode. If either room's temperature has been above the upper temperature bound for a certain amount of time, cooling mode is activated. If both rooms' temperature have been beneath the upper temperature bound for a certain amount of time, non-cooling mode is activated. In cooling mode, a PI controller determines the supply air flow rate based on the temperature deviation. If one room is too warm but the other is not, the warm room is used to determine the supply flow rate. In non-cooling mode, the minimum air flow rate allowed by ASHRAE Standard 62.1 is supplied. Heating commands are determined for each room by individual PI controllers.

B. Rule-Based, Feedback Controller with Occupancy Measurements

For rule-based control, we use a modified form of the MOBS controller proposed in [13]. The original MOBS controller was for fully actuated rooms. The proposed controller is called MOBS$^{\mathrm{ua}}$. It is similar to the BL controller, but occupancy measurements are used instead of a day-night schedule. If one room is occupied and the other is not, the occupied room is used to determine the minimum flow rate. During unoccupied times, relaxed temperature constraints are used. The temperature constraints during occupied times are the same as those used by the BL controller during daytime.

C. Model Predictive Control

MPC minimizes total energy consumption over its time horizon. The dynamic model's state equations and the thermal comfort region are used as constraints. In this study, we choose:

$$[T^{\mathrm{MOBS^{ua}}}_{\mathrm{low,unocc}}, T^{\mathrm{MOBS^{ua}}}_{\mathrm{high,unocc}}] \subseteq [T^{\mathrm{MPC}}_{\mathrm{low,unocc}}, T^{\mathrm{MPC}}_{\mathrm{high,unocc}}] \tag{3}$$

We make this choice because MPC can predict future temperature constraints and take control actions to reach the comfortable range before occupants arrive but MOBS$^{\mathrm{ua}}$ can only *react* to stricter temperature constraints *after* an occupant has entered the room.

The problem formulation for the MPC optimization, for a prediction horizon of K time steps, is:

$$U^* := \arg\min_{U} J(U) \tag{4}$$

where $U = [u^T(k), \ldots, u^T(k+K)]^T$, $J(U) = \sum_{i=k}^{k+K} E(i)$, and $E(i)$ is the energy consumption between the i^{th} and $(i+1)^{th}$ time steps and is estimated as:

$$E(i) \approx \Delta t \Big[P_{fan}(i) + P_{AHU}(i) + P_h(i) \Big] \tag{5}$$

The optimization problem (4) is subject to the room temperature and humidity constraints $T(i) \in [T_{\text{low}}^{\text{MPC}},\ T_{\text{high}}^{\text{MPC}}]$ and $W(i) \in [W_{\text{low}},\ W_{\text{high}}]$, respectively. The flow rate and temperature constraints during occupied times are the same as those for the other controllers. The humidities for the comfortable range were chosen using [23].

In theory, MPC should predict and compensate for constraint violations, but sometimes these are unavoidable due to physical limitations of the actuators or of the algorithm (e.g., prediction horizon). When a violation occurs, control commands are chosen to return the system to the feasible range as quickly as possible.

IV. Simulation Setup

All simulations were performed using MATLAB. The optimization for MPC was performed using the BFGS optimization method in IPOPT [24]. One day was simulated for winter, spring, and summer ambient weather conditions in Gainesville, FL using historical data. Two different occupancy profiles were used. In the first profile, each room is occupied by a single occupant from 8AM–5PM. Each occupant takes a one-hour lunch break at noon. In the second profile, one room is unoccupied at all times, and the other room's occupancy is identical to profile 1.

We assume symmetric ducts that distribute airflow evenly to each room. Because of this, occupancy profile 1 mimics the fully actuated case (i.e., room dynamics will be identical due to boundary conditions and internal gains). Occupancy profile 2 then represents the under-actuated case in which the flow to the unoccupied room is bound to the occupancy and condition of the occupied room.

Table I lists the building and simulation parameters of interest. All states and exogenous inputs are known deterministically by each controller. The supply air has constant temperature (55 °F) and relative humidity (90%).

TABLE I
Relevant Building and Simulation Parameters

Order of R values	10^{-1} K/W
Order of C values	10^5 J/K
Fan power coefficient, α	36.8 W/(kg/s)^3
Area of each room	110 ft^2
Room heater capacity	1500 W
Internal heating when occupied	450 W
Internal heating when unoccupied	50 W
Ratio of return air to outside air	4:1
Surrounding room temperatures	72 °F

During daytime for the BL controller and during occupied times for MOBS$^{\text{ua}}$ and MPC, the room temperatures are subject to the constraint $T(i) \in [71.5\ °\text{F},\ 74.5\ °\text{F}]$. During nighttime, room temperatures are constrained to $T^{\text{BL}}(i) \in [64\ °\text{F},\ 80\ °\text{F}]$ for the BL controller. When a room is unoccupied, the MOBS$^{\text{ua}}$ and MPC controllers use the constraints $T^{\text{MOBS}^{\text{ua}}}(i) \in [70\ °\text{F},\ 76\ °\text{F}]$ and $T^{\text{MPC}}(i) \in [60\ °\text{F},\ 90\ °\text{F}]$, respectively. MPC also places the constraint $W(i) \in [0.0074,\ 0.01\ °\text{F}]$ on the absolute humidity in each room.

Four scenarios for MPC were examined. MPC algorithms using a 30-minute prediction horizon are compared to MPC algorithms using a 120-minute prediction horizon, and, for each prediction horizon, MPC *with* weather predictions is compared to MPC *without* weather predictions. Weather conditions used for control are ambient temperature and humidity. Without weather predictions, current weather is assumed to remain constant.

V. Simulation Results

We use two metrics for each controller's performance: total energy consumption and average temperature constraint violation (ATCV) during occupied times. The ATCV during occupied times is the mean magnitude of the room temperatures' deviations from the comfortable range during occupied times.

Table II shows the total energy consumptions for the baseline controller and the percent savings of each proposed control method compared to the baseline. The MOBS$^{\text{ua}}$ and MPC algorithms all resulted in significant savings. Table III shows the ATCVs when the rooms were occupied. A "—" represents a value less than 0.005. None of the controllers yielded any humidity constraint violations.

TABLE II
Daily Energy Consumption (kWh) of Baseline Controller and % Improvement over baseline of Proposed Controllers, for Occupancy Profile 1/2

	Winter	Spring	Summer
Baseline (kWh)	38.1/39.2	49.9/50.8	67.9/69.0
MOBS$^{\text{ua}}$ (%)	10.2/12.2	11.0/14.0	11.8/15.2
MPC30 (weather) (%)	35.7/48.0	22.6/35.2	16.8/23.5
MPC30 (no weather) (%)	37.0/48.2	26.4/31.9	19.1/25.6
MPC120 (weather) (%)	28.3/43.4	22.8/38.0	15.8/25.6
MPC120 (no weather) (%)	28.3/43.6	23.0/28.3	13.8/26.1

A. *Occupancy Profile 1 vs. Occupancy Profile 2*

The power consumptions for the second occupancy profile were generally lower than those for the first. ATCVs were nearly identical. The BL controller used slightly more energy for profile 2. This is because, when the room is occupied, the occupant provides a small amount of heating, but the BL

TABLE III

AVERAGE OCCUPIED TEMPERATURE VIOLATION (°F) FOR OCCUPANCY PROFILE 1/2

	Winter	Spring	Summer
Baseline	0.05/0.05	0.06/0.07	0.05/0.05
MOBS$^{\mathrm{ua}}$	0.08/0.08	0.09/0.07	0.12/0.11
MPC30 (weather)	0.27/0.27	0.22/0.22	0.13/0.13
MPC30 (no weather)	0.27/0.27	0.22/0.22	0.13/0.13
MPC120 (weather)	—/0.01	—/—	—/0.01
MPC120 (no weather)	—/0.01	—/—	0.01/0.01

controller must supply this additional heating to maintain the temperature constraint when the room is unoccupied.

MOBS$^{\mathrm{ua}}$ for occupancy profile 2 resulted in more energy savings as a percentage of BL than for profile 1. For profile 2, there was a small increase in nominal consumption during winter compared to profile 1. This is because the colder air in the empty room (due to relaxed constraints) reduces the mixed air temperature at the AHU—requiring less cooling. Just like BL, however, the controller must provide the heating normally supplied by the occupant. Cold weather amplifies this effect (resulting in more consumption).

All MPC scenarios resulted in more energy savings (both nominal and percentage of BL) for the second occupancy profile than for the first. This is again due to the temperature of the mixed air and reduced heating. The MPC algorithms allow the temperature of the unoccupied room to float in a much wider range–causing the temperature to drop significantly.

B. MOBS$^{\mathrm{ua}}$ vs. Baseline

The MOBS$^{\mathrm{ua}}$ controller resulted in consistent energy savings (10–15%) compared to the BL controller for both occupancy profiles. The MOBS$^{\mathrm{ua}}$ controller did, however, result in consistently more ATCVs during periods of occupancy. Each of these results is caused by two separate differences between the controllers.

First, the BL controller uses an occupancy schedule rather than actual measurements to determine constraints. Therefore it "over-conditions" the rooms for long periods (e.g., 5PM–10:30PM). The MOBS$^{\mathrm{ua}}$ controller, however, only uses the stricter constraints when the room is actually occupied—resulting in less energy consumption.

Because of this occupancy information, however, the MOBS$^{\mathrm{ua}}$ controller allows the temperature to begin floating during the lunch break. When the occupants return, the MOBS$^{\mathrm{ua}}$ controller must then bring the temperature back from some "uncomfortable" point while the BL controller has already been maintaining thermal comfort.

C. MPC vs. Baseline

All of the MPC schemes also resulted in significant energy savings (14–48%) compared to baseline. The availability of weather predictions had a small effect on average temperature violations. For each MPC algorithm without weather predictions, the addition of weather predictions resulted in less than a 0.02 °F change in the ATCV during occupied times for both profiles. For power consumption, weather predictions generally resulted in a negligible change in nominal energy consumption.

Fig. 2. Room temperature for MPC with different prediction horizons. The MPC with a 30-minute horizon is not able to supply enough heating to obey the temperature constraints in half an hour. The MPC with a 120-minute horizon has ample time to successfully raise the room's temperature to the comfortable range.

Prediction horizon also played a role in the performance of MPC in maintaining thermal comfort. For both profiles, using MPC with a 30-minute prediction horizon resulted in more ATCVs than using the BL controller; conversely, using MPC with a 120-minute horizon consistently resulted in fewer ATCVs. Due to its schedule, the BL controller begins conditioning the rooms using the stricter constraints 90 minutes before the room is actually occupied. In contrast, the short-horizon MPC only "knows" about the stricter constraints 30 minutes beforehand, but it takes more than an hour for the heaters to bring the temperatures to the comfortable range. Figure 2 shows these differences between the MPC prediction horizons. Still, the magnitudes of the ATCVs for the 30-minute MPC are small (less than 0.3 °F).

The total energy consumption was, however, not highly sensitive to the prediction horizon used. In general, a longer horizon typically resulted in slightly higher energy consumption than a shorter horizon. This is expected since the longer horizon MPC uses more actuation to ensure satisfaction of future constraints.

D. MOBS vs. MPC

The MPC controllers used less energy than the MOBS$^{\mathrm{ua}}$ controller in all simulations. However, MOBS$^{\mathrm{ua}}$ had fewer ATCVs than did the 30-minute MPC in all simulations. Both of these results are due to the choice of the temperature constraints during unoccupied times for the two controllers. Constantly maintaining the stricter constraints, however, requires considerably more control actions by MOBS$^{\mathrm{ua}}$ and therefore consumes more energy. MPC particularly outperforms MOBS$^{\mathrm{ua}}$ during the winter. This is again because the cold weather requires MOBS$^{\mathrm{ua}}$ to provide more heating to meet its stricter constraints.

Effect of temperature constraints: After relaxing the temperature constraints used by MOBS$^{\mathrm{ua}}$, significant energy savings were found. In fact, when the MOBS$^{\mathrm{ua}}$ unoccupied constraints were relaxed to those used for BL at night, MOBS$^{\mathrm{ua}}$ used less energy than any of the MPC controllers, but the ATCVs increased.

Effect of design occupancy: The result that MPC control is able to reduce energy use considerably more than MOBS$^{\mathrm{ua}}$ is apparently in conflict with those in [13] that report comparable energy savings between controllers (in the order of 50%). This discrepancy is caused by the number of occupants expected in each room. In [13], the authors consider rooms designed for 3 occupants. The two rooms in this paper are designed for single occupants. Because the BL controllers in both [13] and this paper assume maximum occupancy to calculate the minimum flow rate, the BL controller is even more inefficient for higher design occupancy. Both $\mathrm{MOBS}^{\mathrm{ua}}$ and MPC in that case improve significantly over baseline by reducing airflow when the room is unoccupied so that the relative difference in their performances becomes less prominent. In fact, upon repeating the simulation studies in this paper with a design occupancy of 3 for each room, it turned out that $\mathrm{MOBS}^{\mathrm{ua}}$ and MPC have approximately 50% and 50-70% energy savings compared to BL, respectively—which are consistent with previous findings. Thus, estimates of energy savings also depend on the number of occupants for which the rooms are designed.

VI. Discussion and Future Work

Both RBC with occupancy measurements ($\mathrm{MOBS}^{\mathrm{ua}}$) and MPC with perfect occupancy predictions were shown to offer significant potential energy savings (10-48%) compared to conventional control methods for the under-actuated zone studies. The scenario in which one room is occupied and the other is unoccupied resulted in more savings potential than when both rooms are occupied. This is largely due to the looser temperature constraints for the $\mathrm{MOBS}^{\mathrm{ua}}$ and MPC controllers (as well as the occupancy information available to each).

For MPC, decreasing the prediction horizon increases temperature violations. This is due to limitations of the control actuation. However, for the scenarios considered, average violations were so small for the short-horizon MPC that it is difficult to justify using much longer horizons to reduce them further because a longer horizon significantly increases computational complexity. Furthermore, MPC with a longer prediction horizon does not appear to yield significantly more energy savings for the scenarios considered. In fact, a longer horizon may increase energy consumption to pre-condition the zone to obey future constraints.

In this study weather predictions were found to have little effect on the performance of MPC. The authors of [20], however, report significant improvement with the use of weather predictions. In [20] the HVAC system has a great deal of actuation (including the use of a cooling tower for free cooling, which is highly dependent on weather conditions). This additional actuation may increase the value of weather predictions. Examining how much weather predictions matter for buildings that do not have access to such control actuation requires a more extensive study, which is a topic of future research.

Most of the savings by both $\mathrm{MOBS}^{\mathrm{ua}}$ and MPC come from reduced flow rate. Because these controllers have occupancy information, they reduce the supply airflow when the rooms are unoccupied. This lower airflow requires less air to be conditioned by the AHU—which is the largest consumer of power in the HVAC system examined. However, MPC also resulted in significantly more savings than did $\mathrm{MOBS}^{\mathrm{ua}}$. This is caused in part by the number of occupants for which the rooms are designed. In this study, personal offices designed for single occupants are considered. When the design occupancy is increased, the baseline controller consumes much more energy due to ASHRAE Standard 62.1, and $\mathrm{MOBS}^{\mathrm{ua}}$ and MPC perform more comparably.

In addition, a key observation of this study is that when MPC outperformed the rule-based controller, it did so mostly due to the looser temperature constraints used in MPC during unoccupied times, which required less actuation. The looser constraints also allowed the room temperature of unoccupied rooms to decrease. Lower room temperatures reduce the difference between the mixed air and supply air at the AHU level—which greatly decreases energy consumption. When temperature constraints are loosened for $\mathrm{MOBS}^{\mathrm{ua}}$ as well, the energy savings performance of both controllers become similar. Widening the temperature constraints in $\mathrm{MOBS}^{\mathrm{ua}}$, however, led to larger temperature deviations from the comfortable range during occupied times. If more powerful control actuation is available (e.g., heaters with larger capacities), these temperature violations can be reduced. This observation is significant for the practitioner because rule-based control is much easier to implement than MPC, and obtaining occupancy predictions is much more challenging than obtaining occupancy measurements.

Another key result of this paper is that under-actuated zones can still offer significant energy savings over baseline. In fact, under-actuated zones with both occupied and unoccupied rooms can even offer more energy savings than those with only occupied rooms. These savings are largely due to relaxed temperature constraints during unoccupied periods as well as the ability to provide heating to each room independently. If heating were supplied by the VAV only, then having some occupied and some unoccupied rooms may result in similar savings to having only occupied rooms. This scenario requires further research. Additionally, a more direct comparison between fully actuated and under-actuated zones is of interest.

In this paper, all AHU variables are assumed constant, but the ratio of return air to outside air fed into the AHU and the temperature to which the air is conditioned by the AHU significantly affect the cooling power required. More research is required to develop energy-efficient control algorithms when AHU variables are also available for manipulation [25]. Experimental validation of the results presented in this paper is currently ongoing.

References

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

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

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

[4] ——, “ANSI/ASHRAE standard 55-2004, thermal environmental conditions for human occupancy,” 2004.
[5] C. Liao, Y. Lin, and P. Barooah, “Agent-based and graphical modeling of building occupancy,” *Journal of Building Peformance Simulation, Special Issue on Modelling Occupants' Presence and Behaviour Part II*, vol. 5, pp. 5–25, 2012.
[6] S. Meyn, A. Surana, Y. Lin, S. M. Oggianu, S. Narayanan, and T. A. Frewen, “A sensor-utility-network method for estimation of occupancy distribution in buildings,” in *48th IEEE Conference on Decision and Control*, December 16-18 2009.
[7] “Commercial buildings energy consumption survey (CBECS): Overview of commercial buildings, 2003,” Energy information administration, Department of Energy, U.S. Govt., Tech. Rep., December 2008. [Online]. Available: http://www.eia.doe.gov/emeu/cbecs/cbecs2003/overview1.html
[8] Y. Agarwal, B. Balaji, S. Dutta, R. K. Gupta, and T. Weng, “Duty-cycling buildings aggressively: The next frontier in HVAC control,” in *Information processing in sensor networks (IPSN)*, 2011, pp. 246–257.
[9] A. Persily, A. Musser, S. Emmerich, and M. Taylor, “Simulations of indoor air quality and ventilation impacts of demand controlled ventilation in commercial and institutional buildings,” August 2003.
[10] Y. Tachwali, H. Refai, and J. E. Fagan, “Minimizing hvac energy consumption using a wireless sensor network,” in *33rd Annual Conference of the IEEE Industrial Electronics Society*, November 2007.
[11] V. Dhummi, D. Demetriou, H. J. Palanthandalam-Madapusi, H. E. Khalifa, and C. Isik, “Robust occupancy-based distributed demand control ventilation,” *International journal of ventilation*, vol. 9, no. 4, pp. 359–369, March 2011.
[12] V. Erickson, M. Carreira-Perpiñán, and A. Cerpa, “OBSERVE: Occupancy-based system for efficient reduction of HVAC energy,” in *10th International Conference on Information Processing in Sensor Networks (IPSN)*, April 2011.
[13] S. Goyal, H. Ingley, and P. Barooah, “Occupancy-based zone climate control for energy efficient buildings: Complexity vs. performance,” *Applied Energy*, vol. 106, pp. 209–221, June 2013.
[14] Y. Lin, T. Middelkoop, and P. Barooah, “Issues in identification of control-oriented thermal models of zones in multizone buildings,” in *IEEE Conference on Decision and Control*, December 2012, pp. 6932 – 6937.
[15] C. Liao and P. Barooah, “An integrated approach to occupancy modeling and estimation in commercial buildings,” in *2010 American Control Conference*, June 30-July 2 2010.
[16] P. Morosan, R. Bourdais, D. Dumur, and J. Buisson, “Building temperature regulation using a distributed model predictive control,” *Energy and Buildings*, vol. 42, no. 9, pp. 1445–1452, September 2010.
[17] S. Prívara, J. Široký, L. Ferkl, and J. Cigler, “Model predictive control of a building heating system: The first experience,” *Energy and Buildings*, vol. 43, no. 2-3, pp. 564–572, February-March 2011.
[18] P. Ferreira, A. Ruano, S. Silva, and E. Conceição, “Neural networks based predictive control for thermal comfort and energy savings in public buildings,” *Energy and Buildings*, vol. 55, no. 12, pp. 238–251, December 2012.
[19] S. Bengea, A. Kelman, F. Borrelli, R. Taylor, and S. Narayanan, “Model predictive control for mid-size commerical building HVAC: Implementation, results and energy savings,” in *Second International Conference on Building Energy and Environment*, August 2012.
[20] F. Oldewurtel, A. Parisio, C. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari, “Use of model predictive control and weather forecasts for energy efficient building climate control,” *Energy and Buildings*, vol. 45, no. 2, pp. 15–27, February 2012.
[21] S. Goyal, H. Ingley, and P. Barooah, “Effect of various uncertainties on the performance of occupancy-based optimal control of HVAC zones,” in *IEEE Conference on Decision and Control*, December 2012, pp. 7565 – 7570.
[22] S. Goyal and P. Barooah, “A method for model-reduction of nonlinear building thermal dynamics,” *Energy and Buildings*, vol. 47, pp. 332–340, April 2012. [Online]. Available: http://dx.doi.org/10.1016/j.enbuild.2011.12.005
[23] American Society of Heating, Refrigerating and Air Conditioning Engineers, “The ASHRAE handbook fundamentals (SI Edition),” 2005.
[24] A. Wächter and L. B, “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming,” *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, March 2006.
[25] S. Goyal and P. Barooah, “Energy-efficient control of an air handling unit for a single-zone VAV system,” in *IEEE Conference on Decision and Control*, 2013, pp. 4796 – 4801.