2020 IEEE/ACM Fifth International Conference on Internet-of-Things Design and Implementation (IoTDI)

Leveraging Fine-Grained Occupancy Estimation Patterns for Effective HVAC Control

Yukun Yuan*, Kin Sum Liu*, Sirajum Munir†, Jonathan Francis†, Charles Shelton† and Shan Lin*
*Stony Brook University, †Bosch Research and Technology Center
Email: {yukun.yuan, shan.x.lin}@stonybrook.edu, kiliu@cs.stonybrook.edu,
{sirajum.munir, jon.francis, charles.shelton}@us.bosch.com

***Abstract*—As occupancy sensing technologies become mature, various occupancy sensors are increasingly deployed in commercial buildings for pervasive occupancy monitoring. These sensors provide occupant-count data, which contains rich spatiotemporal information about occupancy patterns. With long-term occupant-count data collected from a commercial building, we design three different predictive models that capture the occupancy dynamics and examine how a model predictive control of the HVAC system benefits from actual occupancy count prediction. Our analysis reveals that mispredictions of occupancy states, especially false positives and false negatives, may introduce inefficient control that leads to energy waste or user discomfort. To address this issue, we take a step further to design an adaptive model predictive controller that minimizes inefficient control actions according to misprediction types and distributions. A comprehensive evaluation is performed in OpenBuild and EnergyPlus simulators to study the effectiveness of the proposed end-to-end control strategy. The evaluation shows that the proposed solution reduces energy consumption by 29.5% while improving the average weighted occupants comfort by 86.7% in Predicted Mean Vote (PMV) over the fixed schedule strategy.**

I. Introduction

Heating, ventilation, and air conditioning (HVAC) is a major source of energy consumption in the U.S. In 2017, approximately 30% of energy consumption for commercial buildings in the U.S. was used for HVAC [1]. Usually, building operators use a static schedule for controlling HVAC systems without having a deeper understanding of how many people use the building at different times of the day. In addition, many HVAC systems operate by assuming maximum occupancy in each room, which leads to significant energy waste, e.g., an HVAC system providing ventilation for 30 people when there are only ten people in a room [2]. Such widely-used HVAC control designs miss opportunities to perform more accurate and efficient control.

As occupancy sensing technologies become mature, occupancy sensors are increasingly deployed in modern commercial buildings. These sensors provide fine-grained occupancy count in real-time. Such information can benefit the HVAC control to significantly improve building energy efficiency, provide occupant thermal comfort, and enhance building management.

There have been a number of research works on occupancy-based sensing and control [3], [4], [5], [6]. However, most of these efforts focus on coarse-grained occupancy estimation, e.g., binary occupancy patterns (occupied or not). These works include recent studies employing PIR motion detectors [3], [4] and energy consumption data analysis [5], [7]. Compared to these approaches, this work leverages fine-grained occupancy estimation information (people count) from a large commercial office for the long-term (six months). With this dataset, we design predictive models that capture the spatiotemporal nature of occupancy dynamics, and with such predictive models, we design a model predictive control algorithm for HVAC control that achieves more significant energy savings.

There are a few works considering occupancy count for HVAC control [8], [9], [10], [11]. In [12], [13], heuristic-based HVAC control algorithms are introduced to control room temperature. In [14], [15], [16], stochastic model predictive control algorithms have been developed to infer the statistics of the disturbances and optimize control actions. Different from these works, our work takes a step further to systematically analyze the prediction error and the introduced extra control cost, addressing a common problem for existing predictive control designs. Our analysis shows that even with sophisticated predictive models, e.g., inhomogeneous Markov chain and sequential & contextual neural network, mispredictions still occur often due to the stochastic nature of the occupancy behaviors. These mispredictions usually introduce inefficient and ineffective HVAC control, especially in the morning and evening hours: the false positives may increase HVAC working time significantly when no occupants actually show up, whereas the false negatives decrease occupant comfort when they are still present.

To address these issues, it is essential to find a good trade-off between prediction accuracy and control cost. We design an adaptive model predictive controller to minimize inefficient control actions according to different types of misprediction and the occupancy states. Our control design provides probabilistic guarantees on the control cost under occupancy prediction uncertainty.

We implement the adaptive MPC control using OpenBuild and EnergyPlus and compare our approach against existing MPC control algorithms with different prediction models. Extensive real-world trace-driven simulations show that our approach outperforms the existing fixed schedule solution by 29.5% in energy-saving and 86.7% in thermal comfort.

The major contributions of this work are listed as follows:

- We perform a comprehensive study to understand occupancy patterns in a commercial space with long-term

978-1-7281-6602-5/20/$31.00 ©2020 IEEE
DOI 10.1109/IoTDI49375.2020.00016

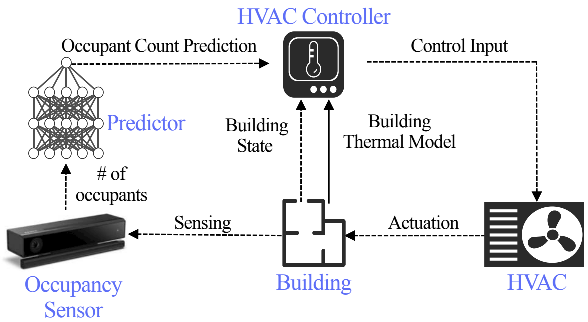

Fig. 1: Occupancy-aware HVAC Control System Overview

(nine months) fine-grained occupancy estimation data. We develop multiple predictive models to understand the effectiveness of future occupancy prediction for such a setting.

- We develop an adaptive model predictive HVAC controller that uses future occupancy prediction to control HVAC to reduce energy consumption and improve the thermal comfort of the occupants. Different from existing MPC control designs that rely on accurate prediction to perform precise control, our adaptive MPC dynamically adjusts the policy to minimize inefficient HVAC control actions due to misprediction under high occupancy uncertainty and provides probabilistic guarantees on the control cost.
- We perform a thorough evaluation using OpenBuild and EnergyPlus to evaluate the effectiveness of the proposed end-to-end control strategy in terms of energy-saving and thermal comfort. Our evaluation based on real-world occupancy data shows that the proposed solution reduces energy consumption by 29.5% while improving thermal comfort by 86.7% in terms of PMV over the existing fixed schedule control strategy.

II. Occupancy-aware MPC for HVAC System

A. Design Overview

Figure 1 shows the system overview of the occupancy-aware HVAC control system for building HVAC. Firstly, the HVAC controller imports a building thermal model according to the building data, such as floor plan, building material, and HVAC system deployment, and so on. Then the occupancy-aware HVAC control system adaptively adjusts the HVAC power to maintain desired occupants' comfort and save energy with two feedback inputs: predicted occupant-count and real-time building thermal states. The first one is the predicted number of occupants in each zone over a time slot based on the historical and sensed real-time occupants' information, requiring an occupancy count predictor. The building thermal state information can be detected via building monitoring systems [17].

The knowledge of real-time and future occupant-count in each zone is important to make efficient HVAC control in two aspects. Firstly, accurate binary occupancy status, e.g., occupied or not, helps HVAC controllers to determine when to turn on or off the HVAC system. Secondly, due to the occupants' heat emission, the controller with accurate occupant count can save energy and perform precise heating and cooling accordingly. In this framework, we deploy occupancy sensors to collect the real-time occupant counts of different zones, and then use such a dataset to design occupancy predictors and model predictive HVAC controllers.

For the model predictive control (MPC) based HVAC controller [15], [18], one MPC-based optimization problem is solved to determine the HVAC power for the future time slots. The MPC controller discretizes the timeline into multiple time slots and considers the power control for future N time slots at the beginning of time slot t. One time slot is index by k, $(k = t, ..., t + N - 1)$. Suppose there are n zones in one building and let $u(k)$ be a length n column vector to describe control input, i.e., heating/cooling power during time slot k for n zones.

B. Real-time Building Occupancy Detection

Estimating real-time building occupancy information is an active area of research. There are several techniques for occupancy estimation using different types of sensors, e.g., ultrasonic sensors [19], PIR sensors [20] and RGB cameras [21]. In this work, we implement a solution called FORK [22] (Fine grained Occupancy estimatoR using Kinect), which uses a depth sensor (Microsoft Kinect for XBOX One) mounted at the ceiling near a doorway to estimate occupancy. We deploy five depth sensors to cover all the doorways of a Bosch office to collect real-time depth frames that are fed into the FORK algorithm. FORK estimates and updates occupancy count when someone enters or exits through each door. FORK can accurately estimate occupancy even when multiple people move simultaneously as it achieves over 99% accuracy for occupancy estimation under realistic scenarios [22]. FORK estimated occupancy count is used as the actual occupancy count in this work. The floorplan and dataset can be found in [23]. One empirical study of occupancy estimation can be found in [24].

C. Dynamic Building Thermal Model

According to [25], [26] and [15], we have the following dynamic building thermal model to describe how the building thermal states change with HVAC power and number of occupants:

$$x(k + 1) = Ax(k) + B_u u(k) + B_d d(k), \quad y(k) = Cx(k)$$

where $x_i(k)$ is the state vector (containing the temperatures, surfaces, and internal nodes) by the end of time slot k of zone i. $d(k)$ is a column vector to represent the external (e.g., outside temperature and solar gains) and internal (e.g., occupants) gains disturbance vector during time slot k. $y(k)$ is a length n column vector denoting the indoor temperature by the end of time slot k for n regions. A, B_u, B_d and C are fixed parameter matrices which are specified by the building configurations. It is noted that these fixed parameters not only describe how the state of one zone i is affected by the state, action and disturbance of this zone, but also consider the heat exchange among zones, e.g., $x_i(k+1)$ is determined by a linear model of $x(k)$, $u(k)$ and $d(k)$, where the parameters are the i-th row of matrix A, B_u and B_d. In detail, some elements of $d(k)$ represent the heat load due to occupant heat emission and the other ones describe the heat exchange with surrounding environment. Suppose each occupant emits the same amount of heat, and the indoor human heat emission is formulated as $e_{heat} \times Occ_i(k)$, where e_{heat} is the heat emitted by an occupant during one time slot, $Occ_i(k)$ is the actual number of occupants in zone i for time slot k. For a zone that does not exchange heat directly with the ambient, the corresponding elements in B_d are zero.

D. occupant-count Predictor

In the Model Predictive Control (MPC) framework, occupant-count is a forcing-function for the system and does not depend on any state variables except time. The effect of occupants appears in the dynamics as an internal heat gain to the thermal state of the building. If we are able to predict the number of occupants accurately for the MPC's time horizon, the controller can make informed decisions to counteract this disturbance and stabilize the thermal condition of the building. In this section, we show how statistical models that are trained on real-world data can be leveraged to predict occupancy changes: at the end of the current time slot t and with previous occupancy counts, we want to predict the future occupancy counts in zone i for the next N time slots, denoted by $\widehat{Occ}_i(t+1), \widehat{Occ}_i(t+2), \ldots, \widehat{Occ}_i(t+N)$.

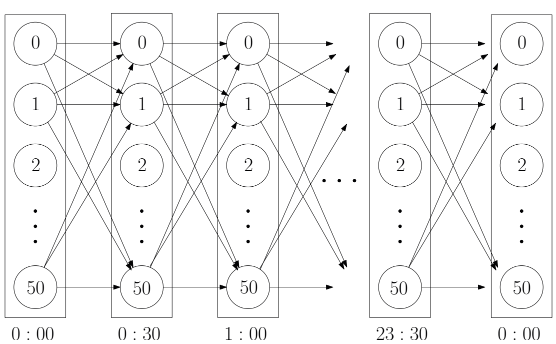

Fig. 2: Occupancy dynamics as Time Inhomogeneous Markov Chain unrolled for one day

1) Time-Inhomogeneous Markov Chain: For the baseline, we model the occupancy dynamics with a Time-Inhomogeneous Markov Chain, where the state represents the occupant count and the transition between states refers to occupancy change temporally. Given the nature of occupancy, the Markov Chain is naturally time inhomogeneous: the probability distribution of the next state depends on the current state and the current time. Therefore at time slot t, the probability of the future occupant count at the next time slot being m_2 given the current count m_1 is $P(\widehat{Occ}_i(t+1) = m_2 | Occ_i(t) = m_1)$.

If we fix time slot duration to be 30 minutes and occupant count to range from 0 to 50 (max count), the Time-Inhomogeneous Markov chain can be unrolled temporally as in Figure 2 in which the occupancy state of a whole day is represented as a chain of 48 states. Here, we assume the state dynamics do not differ from day to day. So the structure of the chain repeats itself after 48 states.

After we use historical occupancy data to estimate the transition probability, an N-step prediction is made by taking the expectation of future counts after N steps: $\widehat{Occ}_i(t+N) = \sum_{m_2} m_2 \times P(Occ_i(t+N) = m_2 | Occ_i(t+N-1)) \cdots P(Occ_i(t+1) | Occ_i(t) = m_1)$. It is possible that some states do not appear once in the training data but occur after training. In that case, we pick the closest state transition.

2) Linear Regression: The Markov Chain requires that the occupancy dynamics obey such assumptions as the Markov property, which may not be true in practice. Therefore, we also implement a linear regression model to make predictions: $\widehat{Occ}_i(t+1) = \sum_j^T \beta_j Occ_i(t-j) + \text{bias}$. The predicted occupancy at $t+1$ can be fed into the regression function again to predict the occupancy at $t+2$. After unrolling for N steps, N future occupant counts are predicted. In this model, we have T coefficients β_j and a bias term as the parameter to be trained on.

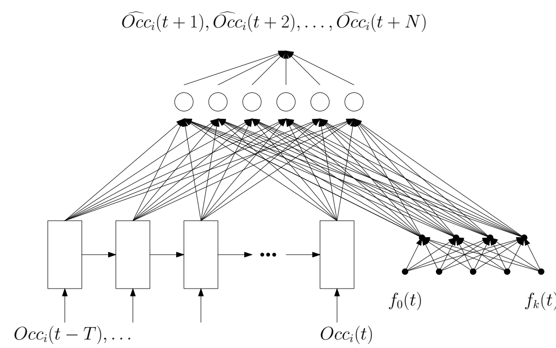

Fig. 3: Proposed network that combines temporal and contextual features to predict the next N occupanct count

3) Sequential & Contextual Neural Network: Based on the preliminary analysis in [27], the occupancy dynamics are inherently non-linear. With the prolonged period of data we collected, we are able to use statistical models that have higher representative power, meanwhile generalizing well without over-fitting.

With a neural network, we cast the prediction as a regression problem. The recurrent part of a neural network is a function that takes the occupant count in T previous time slots as input and connects to the final N neurons, which represented the value of the future occupancy counts of the next N time slots. This recurrent function is shown as the rectangular box in Figure 3, which consists of two hidden layers with 32 units (i.e., neurons) each and ReLU activation. Besides the sequential occupant count, we also leverage other contextual information such as time of the day and day of the week in the prediction. Therefore, we have a fully-connected layer to transform the categorical feature $f_0(t), \ldots, f_k(t)$ to a one-hot encoding (e.g., the time of the day feature is discretized into 24 columns as hour) and again connects to the final N units. So the pyramid-like architecture in Figure 3 illustrates how we can combine the temporal and contextual information to make the final occupancy prediction at the top level. This architecture is flexible to incorporate a more categorical feature that may be available in different applications and usage of the building space.

E. Occupancy Aware Model Predictive Controller Design

Our occupancy-aware model predictive controller has two objectives: minimizing HVAC energy consumption and optimizing occupant thermal comfort. At the beginning of time slot t, the model predictive controller considers the HVAC control for future N time slots and computes the HVAC power sequence in this horizon to optimize these two objectives.

Energy Consumption: Occupancy-aware MPC reduces energy waste based on occupancy prediction, especially when zero occupancy is correctly predicted during working hours, e.g., occupants arrive late in the morning or leave early in the evening. We constrain that the HVAC power for one zone i during time slot k as:

$$\underline{U_i} \leq u_i(k) \leq \overline{U_i}$$

where $\underline{U_i}$ and $\overline{U_i}$ represent the cooling and heating capacity of the building HVAC system, respectively. If $u_i(k)$ is positive, it means the HVAC system is heating the zone; otherwise, the HVAC system is cooling the zone. The $u_i(k)$ refers to the amount of heat flux from the HVAC system that is acting on the zone. Because larger values of $u_i(k)$ would imply more HVAC energy consumption, we use the L1-norm of $u_i(k)$ as a proxy of energy consumption, which is also used in [28]. The total power consumption is represented as $\sum_{k=t}^{t+N-1} \sum_{i=1}^{n} |u_i(k)|$.

Occupant Thermal Comfort: Predicted Mean Vote (PMV) is a common comfort measurement, which is standardized in ISO 7730 [29]. The PMV model estimates the average occupants' comfort level using a function $PMV(\cdot)$ [30]: $PMV(M, T_a, T_r, v, P_a, I_{cl})$, where M is the metabolic rate of the occupant; T_a is the air temperature; T_r is the mean radiant temperature (set equal to T_a); v is the relative air velocity; P_a is the relative humidity; and I_{cl} is the clothing insulation factor of the occupant. The range of PMV is between -3 (cold) and 3 (hot), where 0 is neutral. To simplify the description, let $PMV_i(k)$ to describe the individual occupant comfort for zone i in time slot k.

To make HVAC work efficiently, we consider the weighted occupants' comfort, meaning: (i) we ensure occupants' thermal comfort if the zone is occupied, and (ii) the more more occupants there are in the zone, the more comfortable the indoor environment becomes. Our objective for providing group occupants' comfort is formulated as: $\sum_{k=t}^{t+N-1} \sum_{i=1}^{n} Occ_i(k)|PMV_i(k)|$.

There exists a trade-off between the two objectives. For example, if the outside weather is cold, e.g., below 0 °C, to make occupants comfortable, HVAC needs to heat the zones. However, this would induce high energy cost. Therefore, we use one weight β, to sum up and balance the two objectives, and the problem for our model predictive HVAC controller is formulated as:

$$\begin{aligned} \min_{u(k),x(k)} \; & \sum_{k=t}^{t+N-1} \sum_{i=1}^{n} \left(|u_i(k)| + \beta \; Occ_i(k) \; |PMV_i(k)|\right) \\ \textbf{s.t. } & x(k+1) = Ax(k) + B_u u(k) + B_d d(k) \\ & y(k) = Cx(k) \\ & \underline{U_i} \leq u_i(k) \leq \overline{U_i} \end{aligned} \tag{1}$$

where $\underline{U_i} \leq u_i(k) \leq \overline{U_i}$ constrains the HVAC power for each zone i. Let $\mathbf{P}(Occ(t))$ denote the above Problem (1) with actual occupant-count $Occ(t) = \{Occ_i(k)|1 \leq i \leq n, t \leq k \leq t+N-1\}$. However, at the beginning of time slot t, the actual occupant-count $Occ(t)$ is unknown for our controller. Therefore, predicted occupant-count is used as the input parameters of the MPC to determine the HVAC power and the model predictive HVAC control problem is denoted as $\mathbf{P}(\widehat{Occ}(t))$, where $\widehat{Occ}(t) = \{\widehat{Occ}_i(k)|1 \leq i \leq n, t \leq k \leq t+N-1\}$. In this work, we use the previous three predictors: time-inhomogeneous Markov chain, linear regression, and sequential & contextual neural network to predict $\widehat{Occ}(t)$ based on the historical collected data and real-time occupancy information at the beginning of time slot t.

Although there are some other papers [31], [32], [33] proposing model predictive HVAC controllers using occupancy prediction, their objectives are only reducing the HVAC energy consumption while bounding the indoor temperature within one range, which usually sets the temperature as the feasible extreme values and does not provide satisfied human comfort.

Since our objective and constraints are convex functions, it can be solved by existing solvers, such as Gurobi [34], which is used in OpenBuild [25].

III. Adaptive Model Predictive HVAC Controller Design

A. Empirical MPC Control Performance Analysis

Based on the previous description, the occupancy-aware model predictive HVAC controller determines the power for each zone based on the predicted occupancy count for future

		Prediction of occupancy	
		unoccupied	occupied
Truth of occupancy	unoccupied	True negative	False positive
	occupied	False negative	True positive

TABLE I: Misprediction classification

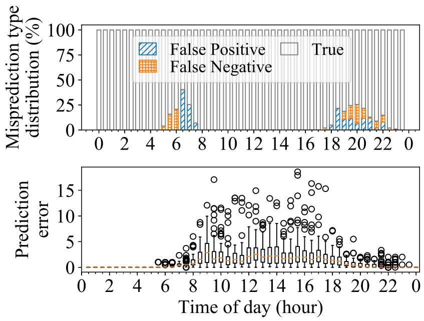

Fig. 4: Performance of Time-Inhomogeneous Markov Chain Predictor

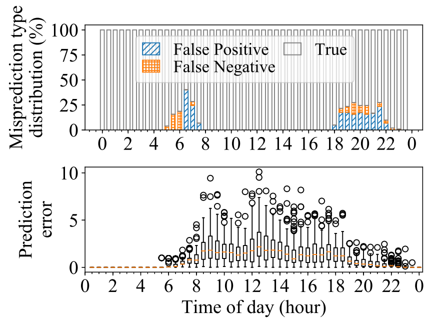

Fig. 5: Performance of Linear Regression Predictor

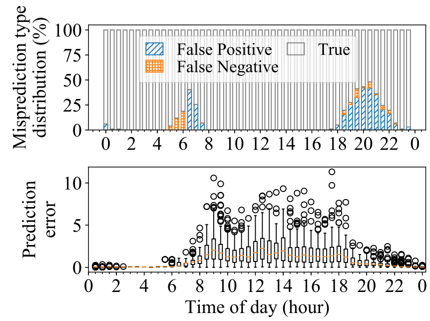

Fig. 6: Performance of Sequential & Contextual Neural Network Predictor

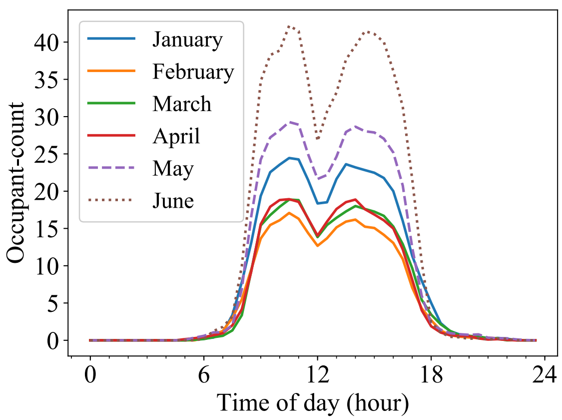

Fig. 7: Average occupant-count over one day for six months

time slots. However, its performance is significantly affected by prediction accuracy due to occupancy pattern uncertainty. To understand its impact, we conduct an empirical analysis using the real-world data collected.

1) Prediction accuracy: We first define misprediction and misprediction on occupant-count. Misprediction means the predictor mispredicts the occupied or unoccupied status for the future time slots, and the misprediction on occupant-count is where the predictor estimates the number of occupants in one zone for the future time slots incorrectly. Then we have two measurement metrics: misprediction type distribution and prediction error to measure the performance of one predictor. We classify the misprediction into four categories: false positive, false negative, true positive and true negative, which is shown in Table I. If both truth and prediction of occupancy are occupied, the misprediction type is marked as true positive. If both truth and prediction of occupancy are unoccupied, the misprediction type is marked as true negative. If the predictor mispredicts occupied as unoccupied or unoccupied as occupied, we note the misprediction as false negative or false positive, respectively. The prediction error is used to measure the misprediction on occupant-count and it is equal to the absolute value of the difference between the predicted number of occupants and the ground truth.

Figures 4, 5, and 6 show the performance of three predictors, respectively. Since predicting the occupancy status correctly does not influence the control performance, we combine true positives and true negatives as true. It can be observed that each predictor incorrectly predicts the occupancy index at the beginning and end of the day with a probability of up to 40.2%. Meanwhile, during the working hours of a day, the predicted occupancy index is always true, whereas, the predictor error is small at the beginning and end of the day and is large during the working hours. For example, the actual occupant-count over one day is shown in Figure 7 and we see that there are few occupants in the zone during the large misprediction probability period, which is why the predictor mispredicts the occupancy status with a higher probability compared with the probability during the high occupant density period. The occupancy pattern variation over months in Figure 7 can be explained as follows. Interns joining in the Summer causes an increase of occupancy count from April to May to June. The reason for smaller occupancy count in February, March, and April compared to January is because employees return to work in January after holidays. However, in February, March, and April, some of them need to do traveling to different offices and conferences. Sometimes they need to work from home due to winter conditions. The occupancy count is also affected due to incoming visitors and events hosted in the office.

2) HVAC control performance assessment : After analyzing the prediction accuracy of the three predictors, we study the performances of an occupancy-aware MPC controller with any of the three predictors.

At first, we define one measurement metric: the energy efficiency times effective PMV improvement to measure the combination of energy efficiency and how much occupant comfort is offered. We conduct several steps to calculate energy efficiency and effective PMV improvement for time slot k: (i) based on the actual detection of occupant-count in time slot k, $Occ_i(k)$ and the initial building state $x(k-1)$, we solve the one time slot version of Problem (1) to decide the optimal HVAC power $u_i(k)$ and comfort value $PMV_i(k)$; (ii) we calculate the energy efficiency and effective PMV

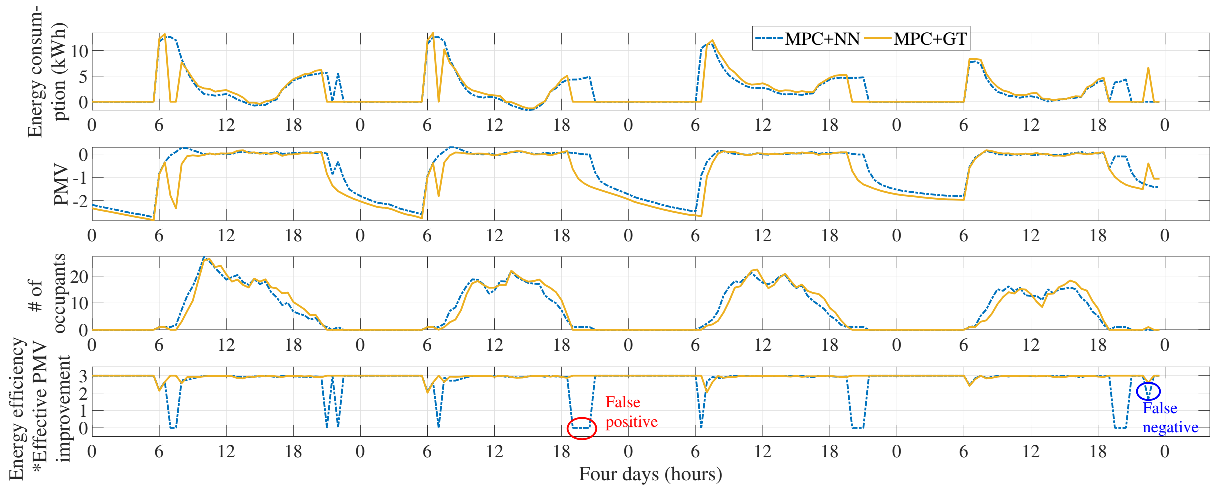

Fig. 8: Performance of MPHC using predicted occupant-count (neural network) and ground truth over 4 days

improvement as follows:

$$\text{Energy efficiency} = \begin{cases} \frac{\min\{u_i(k), \widehat{u}_i(k)\}}{\widehat{u}_i(k)} & \text{if } \widehat{u}_i(k) \neq 0 \\ 1 & \text{if } \widehat{u}_i(k) = 0 \end{cases}$$

$$\begin{array}{c}\text{Effective PMV}\\ \text{improvement}\end{array} = \begin{cases} |\widehat{PMV}_i(k) - F| & \text{if } Occ_i(k) \neq 0 \\ 3 & \text{if } Occ_i(k) = 0 \end{cases}$$

where F is the extreme value of PMV (-3 for cold and 3 for hot). The idea behind this measurement metric is that if HVAC consumes the energy while the zone is unoccupied, its efficient energy consumption ratio is 0 because the optimal energy consumption is 0, and if there are occupants, the control performance is determined by the PMV improvement and energy efficiency regarding the model predictive HVAC control using true occupant-count.

3) Performance of MPHC using three predictors: We show the control traces of the MPC-based HVAC Controller (MPHC) using predicted occupant-count and ground truth over four days in Figure 8.

Due to the space limit, we only show the performance of using the Sequential & Contextual Neural Network (NN)

	Large misprediction probability period		Small misprediction probability period	
	Energy (kWh)	PMV	Energy (kWh)	PMV
GT	54.26	-0.159	36.08	-0.052
MC	71.08	-0.130	28.66	-0.050
LR	72.54	-0.099	27.50	-0.043
NN	79.03	-0.094	25.50	-0.041

TABLE II: HVAC controller performance comparison using different predictors for one day (GT: ground truth; NN: sequential & contextual neural network; LR: linear regression; MC: Time-Inhomogeneous Markov Chain)

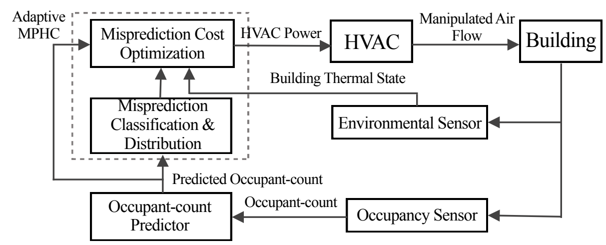

Fig. 9: Adaptive Model Predictive HVAC Control Architecture

predictor, and the details of the experiment setting are described in Section IV. It can be observed that due to the false positive prediction, the controller heats the zone when it is unoccupied, wasting energy, as highlighted in red in the bottom sub-figure. Meanwhile, the prediction-based controller also sacrifices occupants' comfort by closing the HVAC when it makes the false negative prediction, as highlighted in blue in the bottom sub-figure. We partition one day into two periods: large misprediction probability period and small misprediction probability period according to the misprediction probability distribution. Table II shows the performance of the HVAC controller using ground truth or different predictors in terms of energy consumption and average weighted PMV during a different period of one day. During the period with large misprediction probability, using ground truth saves more energy and gives slightly worse occupants comfort compared with using prediction. Based on Figure 8 and Table II, it is concluded that the prediction-based HVAC controller wastes the energy and provides a little better occupants comfort due to its large false positive or false negative misprediction probability in some time periods.

B. Adaptive Controller Design

As described in Section II-E, there are two objectives of HVAC control, minimizing energy consumption and maximizing group occupants comfort, and there exists a trade-off between these two objectives. Ideally, an optimal controller should aim at minimizing energy consumption if there is no occupant in one time slot and consider minimizing one weighted sum of energy consumption and group occupant comfort for a time slot with occupants. However, due to the misprediction, the existing controller using predicted occupant-count optimizes HVAC control with incorrect objectives. Therefore, in order to handle the misprediction, especially the random false positive/negative case, the controller should determine the HVAC power, which is robust to the misprediction and introduces the minimum misprediction cost.

To make our control decisions robust to the random misprediction, we design an adaptive model predictive HVAC controller (MPHC), adapting to the misprediction type distribution and corresponding misprediction cost in the different time slot. The control diagram is shown in Figure 9, which has a closed control loop. The adaptive MPHC determines the HVAC power, which further changes the temperature of manipulated airflow and building thermal states. The occupancy sensors sense the occupants' events and calculate the actual occupant-count, which is used to predict the future number of occupants by the predictor. The environmental sensors detect and output the real-time building thermal state. This information from occupant-count predictor and environmental sensors is forwarded to our adaptive MPHC as feedback to improve control efficiency. The main idea behind our adaptive MPHC is for one upcoming time slot we first generate the misprediction type distribution by sampling the historical prediction and true occupant-count information. Then based on the misprediction type distribution and predicted occupant-count information for a given time slot, we determine the HVAC power by minimizing the misprediction cost expectation.

Misprediction classification and distribution: This part mainly updates misprediction type distribution in the different time slots of the day based on the real-time collected prediction and actual occupant-count data. For the upcoming N time slots, we first sample the historical prediction and true occupancy status data at the same time of day and then count the frequency of different types of misprediction. Let $p_i^{tn}(k), p_i^{fp}(k), p_i^{fn}(k)$ and $p_i^{tp}(k)$ be the probability that the misprediction types true negative, false positive, false negative and true positive exist in zone i for time slot k, respectively. The distribution of these four types of misprediction is forwarded to the misprediction cost optimization part.

Misprediction cost optimization: As shown in Section III-A, the MPHC solving problem $\mathbf{P}(\widehat{Occ}(t))$ is sensitive to the misprediction when there is a large probability of false positive/negative misprediction due to the incorrect weight between two objectives. Hence, our misprediction cost optimization part aims at determining the HVAC power to minimize the misprediction cost expectation based on the probability of predicting the occupancy status correctly or incorrectly for a given predicted status.

Let us recall our model predictive HVAC control objective with true occupant-count $Occ_i(k)$ for zone i in time slot k, defined as

$$J_i(k) = |u_i(k)| + \beta Occ_i(k)|PMV_i(k)|$$

However, at the beginning of time slot k, we only have the predicted occupant-count for future N time slots and n zones.

Let $\widehat{u}_i(k)$ and $\widehat{PMV}_i(k)$ be the energy consumption and PMV when using $\widehat{Occ}_i(k)$ to solve problem (1) for zone i and time slot k. Here, the control objective value with $\widehat{u}_i(k)$ and $\widehat{Occ}_i(k)$ is defined as

$$\widehat{J}_i(k) = |\widehat{u}_i(k)| + \beta Occ_i(k)|\widehat{PMV}_i(k)|$$

Then the misprediction cost is defined as $\Delta J_i(k) = \widehat{J}_i(k) - J_i(k)$.

The misprediction cost expectation for zone i and region k is defined as:

$$\begin{aligned} \mathbf{E}(\Delta J_i(k)) &= \frac{p_i^{fp}(k)\Delta J_i^{fp}(k)}{p_i^{fp}(k)+p_i^{tp}(k)} + \frac{p_i^{tp}(k)\Delta J_i^{tp}(k)}{p_i^{fp}(k)+p_i^{tp}(k)} \\ \textbf{or} \quad &= \frac{p_i^{tn}(k)\Delta J_i^{tn}(k)}{p_i^{tn}(k)+p_i^{fn}(k)} + \frac{p_i^{fn}(k)\Delta J_i^{fn}(k)}{p_i^{tn}(k)+p_i^{fn}(k)} \end{aligned} \tag{2}$$

According to the predicted occupancy status information, we have the different mathematical equations of misprediction cost expectation. $\Delta J_i^{fp}(k), \Delta J_i^{fn}(k), \Delta J_i^{tp}(k)$ and $\Delta J_i^{tn}(k)$ are the misprediction cost of different type of misprediction, and we will discuss how to calculate them one by one.

As shown in Table I, for both true negative and false positive misprediction, there is no occupant in the zone, so the misprediction cost should be the wasted energy, i.e.,

$$\Delta J_i^{tn}(k) = |\widehat{u}_i(k)|, \quad \Delta J_i^{fp}(k) = |\widehat{u}_i(k)| \tag{3}$$

For true positive prediction, the objective value determined by our control decision $\widehat{u}_i(k)$ is $\widehat{J}_i(k) = |\widehat{u}_i(k)| + \beta\widehat{Occ}_i(k)|\widehat{PMV}_i(k)|$, where we use the predicted value $\widehat{Occ}_i(k)$ as the actual value since the predictor makes true positive prediction. If the occupancy status is correctly predicted as occupied, the optimal HVAC power in zone i during time slot k is also influenced by the other $nN-1$ decision variables in MPC. Therefore, we define the optimal objective value of zone i and time slot k with true positive prediction as the expected optimal objective value with 2^{nN-1} possible cases. For case $1 \leq j \leq 2^{nN-1}$, let P_j be the probability that this case exists. Given the case with certain correct or incorrect prediction information, we can determine the MPHC problem formulation and calculate the optimal objective value, denoted by $J_{i,j}(k)$. Then we have $J_i^{tp}(k) = \sum_{j=1}^{2^{nN-1}} P_j J_{i,j}(k|I_i(k) =$ true positive), where $I_i(k)$ is one indicator function to show the misprediction type. The misprediction cost is defined as:

$$\begin{aligned} \Delta J_i^{tp}(k) = {} & |\widehat{u}_i(k)| + \beta\widehat{Occ}_i(k)|\widehat{PMV}_i(k)| \\ & - \sum_{j=1}^{2^{nN-1}} P_j J_{i,j}(k|I_i(k) = \text{true positive}) \end{aligned} \tag{4}$$

In the above function, $J_i^{tp}(k)$ is one pre-computed constant value for given prediction and misprediction distribution of

Algorithm 1: ***Adaptive model predictive HVAC controller real-time HVAC power control***

Input: Time horizon N time slots; number of zones n; weight decided by building managers to balance two objectives β.
Output: Control decision: $u_i(k)$, $1 \le i \le n$, $t \le k \le t+N-1$

1: **while** At the beginning of every time slot t, denoted as t-th time slot **do**
2: Update the initial building thermal state $x(t-1)$;
3: Update the predicted occupant-counts in the upcoming N time slots of n zones, denoted as $\widehat{Occ}(t)$;
4: Update the misprediction type distribution $p_i^{tn}(k), p_i^{fp}(k), p_i^{fn}(k)$ and $p_i^{tp}(k)$ for the future N time slots and n zones.
5: According to predicted occupant-count $\widehat{Occ}_i(k)$ and misprediction type distribution, update the total misprediction cost expectation $\sum_{i,k} \mathbf{E}(\Delta J_i(k))$ by Equation (2).
6: Solve problem (6) to determine the HVAC power $\widehat{u}_i(k)$ for $1 \le i \le n$ and $t \le k \le t+N-1$.
7: **end while**
8: **return** HVAC power decisions

n zones and N time slots and $\widehat{PMV}_i(k)$ is also one linear function related to $\widehat{u}_i(k)$. Therefore, $\Delta J_i^{tp}(k)$ is one linear function of $\widehat{u}_i(k)$.

For false negative prediction, $\widehat{J}_i(k) = |\widehat{u}_i(k)| + \beta|\widehat{PMV_i(k)}|$, where we use $Occ_i(k) = 1$ since the predictor infers the status as unoccupied. The optimal objective value of time slot k and zone i, denoted as $J_i^{fn}(k)$, is also one expected optimal objective value. Its equation is:

$$\Delta J_i^{fn}(k) = |\widehat{u}_i(k)| + \beta|\widehat{PMV}_i(k)| - \sum_{j=1}^{2^{nN-1}} P_j J_{i,j}(k|I_i(k) = \text{false negative}) \quad (5)$$

This equation is also one linear equation related to HVAC power $\widehat{u}_i(k), 1 \le i \le n, 1 \le k \le N$.

Our objective is minimizing the total misprediction cost expectation over N time slots and n zones, defined as $\sum_{i=1}^{n}\sum_{k=t}^{t+N-1}\mathbf{E}(\Delta J_i(k))$. According to Equation (2), $\mathbf{E}(\Delta J_i(k))$ is linear to $\Delta J_i^{fp}(k)$ and $\Delta J_i^{tp}(k)$ or $\Delta J_i^{tn}(k)$ and $\Delta J_i^{fn}(k)$. Based on the previous definition of these four variables, they are linear to the control decision variables $\widehat{u}_i(k)$. Therefore, our objective function is linear to HVAC power $\widehat{u}_i(k)$. The problem of determining the HVAC power to minimize the total misprediction cost is formulated as follows:

$$\begin{aligned} \min_{\widehat{u}(k),\widehat{x}(k)} \quad & \sum_{k=t}^{t+N-1}\sum_{i=1}^{n} \mathbf{E}(\Delta J_i(k)) \\ \textbf{s.t.} \quad & \widehat{x}(k+1) = A\widehat{x}(k) + B_u\widehat{u}(k) + B_d\widehat{d}(k) \\ & \widehat{y}(k) = C\widehat{x}(k), \quad \underline{U_i} \le \widehat{u}_i(k) \le \overline{U_i}, \quad (2) \sim (5) \end{aligned} \quad (6)$$

Since the objective function of the total misprediction cost minimization problem is linear to our decision variables and all the constraints are also linear, problem (6) is convex and can be solved using the convex optimizer.

The pseudo-code of the adaptive model predictive HVAC controller algorithm is shown in Algorithm 1. At the beginning of each time slot t, we first update the building thermal state and predict the occupant-count for the future N time slots in n zones. Then the misprediction classification and distribution part updates the probability of the different types of misprediction for N time slots and n zones. Thirdly, according to the predicted occupant-count and misprediction type distribution, we update our objective of minimizing the total misprediction cost expectation. Finally, we solve the problem (6) to determine the optimal HVAC power, which is robust to the random misprediction.

Based on the empirical performance analysis of predictors and MPC control, during the time slots with small misprediction probability, the occupancy-aware MPC control shows very close performance compared with MPC using true occupant-count. However, for the time slots with large misprediction probability, due to the small true number of occupants, it is hard to estimate the occupancy status with high accuracy. Our probability-based adaptive MPHC essentially offers a probabilistic guarantee that if ϵ percentage of all misprediction type data follows the sampled misprediction type distribution, our solution can minimize the total misprediction cost of ϵ percentages of future time slots. According to the law of large numbers, with long-term occupancy data, the ϵ percentage of time slots' mispredictions gets close to the true distribution. Therefore, our solution can minimize the misprediction cost expectation over time.

IV. Evaluation

A. Evaluation of Prediction Accuracy

We deployed occupancy sensors in a commercial building to collect data for learning and prediction of occupancy. The description of the dataset is shown in Table III. After collecting over nine months of data, there were around 90,000 events in total. Each event is represented by a tuple (date, time, occupancy count) in the dataset.

To investigate the effectiveness of the various occupancy prediction models, we first perform an offline training that uses one month of occupancy data as training data. Then we evaluate the trained model with six months of testing data. The mean average error (MAE) and root mean square error (RMSE) is reported in Table IV. The model of neural network is: input→dense (32 units)→ReLu→dense (32 units)→ReLu→output. The number of epochs is 30, the loss function is mean square error and the optimizer is Adam. Our proposed neural network is able to perform very well in both

Collection Period	2015/08/26∼2016/06/07
Events	>90,000
Format	[date,time,real-time occupant-count]

TABLE III: Dataset of occupants

Granularity= 30 minutes, Horizon= 4 slots (2 hours)			
Models		Main office	Warhol
MC	MAE	1.64	**0.41**
	RMSE	3.09	0.98
LR	MAE	1.46	0.44
	RMSE	2.69	0.92
NN	MAE	**1.43**	0.42
	RMSE	**2.60**	**0.9**

TABLE IV: Prediction accuracy of Time Inhomogeneous Markov Chain (MC), Linear Regression (LR) and Sequential Contextual Neural Network (NN)

metrics. Note that the main office in our dataset is a medium-sized commercial space that can possibly host 50 people. This accurate prediction made by our models, will be beneficial for the model predictive controller, as shown in the later evaluation section.

We also evaluate how the freshness of the data will affect the prediction result in an online setting. To perform online training, we re-train our model every day using all the data we have by this day. The prediction accuracy over months of the main office is shown in Table V. It is observed that the neural network-based model improves the RMSE from 2.23 to 2.03. However, it performs worse in May and June, which is because the occupancy pattern changes a lot from May, meaning that the previous knowledge does not hold accurately for these two months. This observation aligns with the empirical analysis that the distribution of occupancy in the office space may shift in the year and we need to update our belief regularly.

B. Evaluation of Adaptive Model Predictive HVAC Controller

1) Evaluation Setup: We perform a data-driven simulation to evaluate and compare the quantitative performance of the following controllers: (i) OBSERVE [12]: at first, it estimates the future occupancy information using Markov Chain model, and then sets up the temperature set point based on the different predicted occupancy status. It works from 5 am to the end of the day. It aims to make PMV be close to 0 if the zone is occupied; otherwise, the temperature is set to make PMV be close to -0.5. (ii) Fixed Schedule: At the beginning of each time slot, it detects the current indoor temperature and decides the HVAC power during this time slot to make the future PMV value close to 0. The controller works between 6 am and 11 pm, and it does not predict any future binary occupancy

	NN		MCT		LRT	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
Jan.	1.28	2.23	1.32	2.44	1.23	2.19
Feb.	1.26	2.21	1.47	2.69	1.24	2.27
Mar.	1.25	2.11	1.40	2.53	1.18	2.04
Apr.	1.21	2.03	1.33	2.32	1.18	2.01
May	1.83	3.37	1.84	3.44	1.57	2.88
Jun.	3.11	5.47	3.83	6.64	2.46	4.37

TABLE V: Prediction accuracy over months of three predictors

		Energy (kWh)	PMV
MPHC +OC	NN	104.53	-0.046
	MC	99.74	-0.058
	LR	100.04	-0.048
Adaptive MPHC	NN	89.35 (15.3%)	-0.052 ($-0.026\,°$C)
	MC	88.01 (11.8%)	-0.068 ($-0.037\,°$C)
	LR	87.59 (12.4%)	-0.056 ($-0.029\,°$C)

TABLE VI: Average one-day energy consumption and PMV over six months with different predictors

information or occupancy count. (iii) MPHC+Binary: At the beginning of each time slot, it first predicts whether the zone is occupied or unoccupied in the following time horizon. Then it determines the HVAC power by MPHC assuming the number of occupants is 1 if the zone is predicted as occupied in one upcoming time slot. (iv) MPHC+OC: it determines the HVAC power by solving MPHC with predicted number of occupants information. (v) Adaptive MPHC: our design that determines the HVAC power adaptively to the misprediction type distribution.

To measure the occupants comfort of zone i over one day, we use the weighted average PMV value defined as $\sum_{k=1}^{48} Occ_i(k)PMV_i(k)/\sum_{k=1}^{48} Occ_i(k)$. We also evaluate the total energy consumption of one day.

The building in the simulation environment EnergyPlus [35] has three thermal zones with sizes $250m^3$, $100m^3$, and $250m^3$, respectively. We use the default weather data in EnergyPlus as the environmental conditions. We set the power range for zone 2 to be $-5 \leq u(k) \leq 5$ and that for zone 1 and 3 is $-15 \leq u(k) \leq 15$. The time slot is 30 minutes and prediction horizon N is 4. To simulate PMV model, we use the thermal comfort tool [36] to estimate the PMV value for given indoor temperature. Several parameters are set as follows: $M = 1$ met, $v = 0.2m/s$, $P_a = 50\%$ and $I_{cl} = 1$ col, where we set metabolic rate and clothing level by assuming an occupant wears typical winter indoor clothing and takes a seat according to the instructions in [36]. The default β value is 5.

2) Experimental Results: We summarize the main evaluation results as follows:

- Occupancy patterns obtained from the historical data are useful to improve the energy efficiency and occupants comfort, e.g., MPHC+binary reduces the energy consumption by 9.9% and increases occupant comfort by 4.5% over the fixed schedule control strategy.
- Fine-grained occupant-count allows more efficient and precise heating/cooling control than using binary occupancy status information. For instance, MPHC+OC consumes 10.7% less energy and provides 88.1% better thermal comfort than MPHC+binary does.

	Energy (kWh)	PMV
OBSERVE	132.17	0.350
Fixed Schedule	124.30	0.422
MPHC+Binary	111.97	0.403
MPHC+OC	100.04	-0.048
Adaptive MPHC	87.59	-0.056

TABLE VII: Average one-day energy consumption and PMV over six months by different solutions

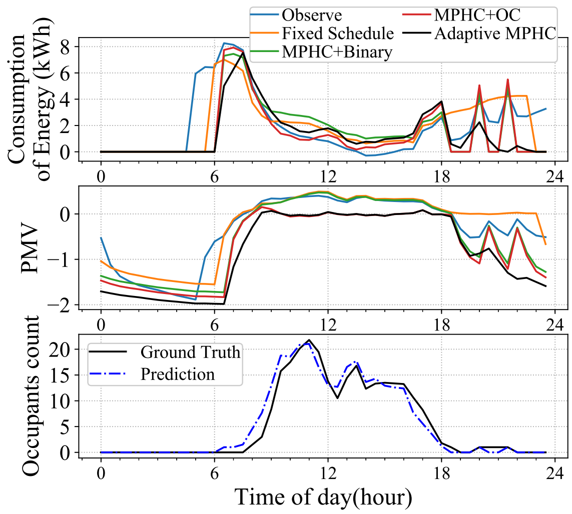

Fig. 10: Control Traces of different solutions over one day

- Due to the false positive and false negative misprediction at the beginning and end of the day, only occupancy-aware MPHC experiences energy waste and occupant discomfort. Our adaptive MPHC based on historical misprediction can minimize the total energy consumption of one day while only introducing a slight decrease of occupant comfort.

First, we compare the performance of MPHC+OC and adaptive MPHC using one of three predictors that we proposed in Section II-D and show the average one-day energy consumption and PMV over six months in Table VI. The key observation is that compared with MPHC+OC, our adaptive MPHC reduces the average one-day energy consumption significantly, and meanwhile, a slight decrease of average weighted PMV is introduced for any predictor. For instance, when using a neural network-based predictor, our solution consumes 15.3% less energy on average for one day and only introduces a temperature decrease of 0.026 °C. Another observation is that no matter which predictor our adaptive MPHC uses, it always outperforms MPHC+OC. In the rest of the evaluation, we compare our adaptive MPHC using LR against the other three solutions since it consumes the least energy compared with using the other two predictors.

Secondly, in Table VII, we show the performance of five solutions in terms of energy consumption and weighted average PMV for one-day when conducting a six-month trace-driven evaluation. Our adaptive MPHC solution reduces the average one-day energy consumption by 33.7%, 29.5%, 21.8%, and 12.4%, respectively, compared with OBSERVE, fixed schedule, MPHC+binary, and MPHC+OC. Meanwhile, it decreases the indoor temperature by 0.044 °C against the fixed schedule, which provides the most comfortable indoor environment. One more observation is that OBSERVE, fixed schedule and MPHC+binary solutions overheat the zones during the daytime, e.g., fixed schedule makes the average weighted indoor temperature be 0.044 °C above the most comfortable value. The reason is that they do not consider the human heat emission, resulting in more energy consumption to heat the zone.

	Large misprediction probability period		Small misprediction probability period	
	Energy	PMV	Energy	PMV
OBSERVE	94.92	0.139	37.25	0.373
Fixed Schedule	92.33	0.186	31.97	0.447
MPHC +binary	73.44	0.157	38.53	0.428
MPHC +OC	72.54	-0.100	27.50	-0.043
Adaptive MPHC	52.20	-0.180	35.39	-0.043

TABLE VIII: Average energy consumption and PMV during different periods of one day over six months

To better understand the control strategy of the different solutions, we plot their control traces over one day in Figure 10, and we have several observations. (i) The first one is both MPHC+Binary and MPHC+OC turn on the HVAC to heat the zones when the zone is predicted as occupied for the future time slots. (ii) The second observation is OBSERVE, fixed schedule and MPHC+binary solutions consume more energy and overheat the zones due to no consideration of occupants' heat emission. (iii) The last one is our adaptive MPHC consumes the power conservatively when one time slot experiences high probability that the ground truth is unoccupied for a given prediction as occupied, e.g., 6:30 and 20:00, and heats the zone aggressively when there is a probability that the ground truth is occupied for a given prediction as unoccupied, e.g., 18:30, 19:00 and 20:30.

Finally, we analyze the average performance of five solutions during different time periods of one day and show the results in VIII. It is observed that our adaptive MPHC consumes the least energy during the large misprediction probability period than the other four solutions do by 45.0%, 43.5%, 28.9%, and 28.0%, respectively. The reason is that our controller determines the HVAC power adaptively to misprediction type distribution, which saves the energy. Meanwhile, it reduces the average weighted indoor temperature by 0.30 °C compared with MPHC+OC. Our controller consumes the more energy during the small misprediction probability period since it needs to compensate the saved energy during the previous large misprediction probability period in order to offer occupant comfort with more occupants. The OBSERVE solution also consumes more energy during the small misprediction probability period since it only aims at making PMV be close to 0 without considering the energy consumption. However, it still saves the energy for the overall day, as shown in Table VII.

V. Discussion

In this work, we deploy depth sensors to detect and estimate occupancy in a commercial space. Depth data from a similar deployment can be found in [37]. Compared with the traditional PIR based motion detection systems, our solution can accurately estimate the number of occupants entering or exiting in different rooms. Our solution achieves over 99% accuracy in estimating space occupancy [22] and it is not as privacy invasive as RGB cameras.

The collected long-term building occupancy count information is used to predict the future occupancy count and fed into the adaptive model predictive HVAC controller. In this work, we focus on the problem that the misprediction of future occupancy information can result in ineffective and inefficient HVAC control. We design an adaptive MPHC to address this challenge. Our proposed solution is agnostic of underlying occupancy estimation solution as it will work with other sensing technologies as long as they provide real-time occupancy estimation information. However, the more accurate the occupancy estimation system is, the better our solution performs. Our solution expects building managers to feed such occupancy data, building and HVAC system information to our solution in order to control the HVAC system.

VI. Related Work

Occupancy-aware HVAC control exploits real-time or predicted future occupancy states to determine the HVAC power, aiming to minimize the HVAC energy consumption. There are three types of occupancy-aware HVAC controllers: reactive controller, condition-based predictive controller, and model predictive controller. (i) The reactive controller uses real-time occupancy information to conduct HVAC power optimization without any prediction of future occupancy states [38]. (ii) The condition-based controllers [12] predict the schedule of occupancy and then adjust the HVAC temperature set-point according to some pre-defined logic to meet occupant comfort, rather than determining the HVAC power directly. (iii) The occupancy-aware model predictive HVAC controller employs a model describing the dynamic building thermal states with predicted knowledge of building occupancy and weather, and then solves an optimization to determine the best HVAC power decisions [8], [9], [10], [13], [15], [16], [14], [11], [39]. There are two differences between these papers and our work: (i) we investigate the sequence of false positive or negative misprediction and propose an adaptive controller to handle such cases, and (ii) our work intends to optimize both energy consumption and occupant comfort, whereas, the related works only constrain the indoor temperature within a bound.

Although some other papers fall within the scope of HVAC control, their focus is not determining the HVAC power directly. [31] concentrates on designing occupancy prediction algorithms for automatically setting up the target temperature that is used for MPC-based HVAC control. [40] considers combining a model predictive control HVAC system with free cooling, such as natural ventilation, minimizing energy consumption while maintaining occupant comfort. [6] studies how to use low-cost sensing technology to detect occupancy and sleep patterns. Moreover, based on such patterns, one strategy for automatically turning off the resident HVAC system is designed. [41] studies how to distort the occupancy data aiming at hiding individual occupant location information while bounding the HVAC system performance.

Some related works [42], [43], [44], [45], [46], [47] argues that traditional thermal comfort measurement methods introduces a high load of sensory input. Therefore, they propose several feedback mechanisms for receiving the occupants' feedback, e.g., their vote on thermal comfort within a space via some web applications.

There are several papers [48], [49] focusing on deploying some infrastructure, such as Wi-Fi, passive infrared sensors and motion (CO_2, sound, ambient light) detectors, etc. to collect real-time building states [50] and then inferring the real-time occupancy that is employed for predicting occupancy states and setting HVAC target temperature.

VII. Conclusion

With a long-term fine-grained real-world occupancy dataset, we conducted a comprehensive study on occupancy prediction and its benefits and constraints for HVAC control in commercial buildings. Our analysis confirmed the results from previous studies that occupancy patterns obtained from historical occupancy data can be used to improve energy efficiency and occupant comfort in buildings. We also reveal that using fine-grained occupancy count allows more efficient and precise control than binary occupancy count. However, mispredictions in the morning and evening often occur and introduce energy waste and user discomfort. Our solution to this problem is an adaptive control design that minimizes misprediction costs based on historical distributions of different types of mispredictions. A comprehensive evaluation is performed in OpenBuild and EnergyPlus, and the evaluation result shows that the proposed solution reduces energy consumption by 29.5% while improving the average weighted occupants comfort by 86.7% in terms of Predicted Mean Vote (PMV) over the fixed schedule control strategy.

Acknowledgment

We would like to thank the anonymous reviewers and shepherd for the insightful feedback. This work was supported, in part, by DOE grant DE-EE0007682 and NSF CNS-1553273. The opinions expressed here are those of the authors and do not necessarily reflect the views of the DOE.

References

[1] U.S. DOE, "Energy savings potential and rdd opportunities for commercial building hvac systems", 2017.

[2] Michael J Brandemuehl and James E Braun, "The impact of demand-controlled and economizer ventilation strategies on energy use in buildings", *ASHRAE Transactions*, vol. 105, pp. 39, 1999.

[3] William A Hoff and James W Howard, "Activity recognition in a dense sensor network.", in *SNA*, 2009, pp. 67–72.

[4] James Scott, AJ Bernheim Brush, John Krumm, Brian Meyers, Michael Hazas, Stephen Hodges, and Nicolas Villar, "Preheat: controlling home heating using occupancy prediction", in *Ubicomp '11*. ACM, 2011.

[5] Christian Beckel, Wilhelm Kleiminger, Romano Cicchetti, Thorsten Staake, and Silvia Santini, "The eco data set and the performance of non-intrusive load monitoring algorithms", in *BuildSys 14*. ACM, 2014.

[6] Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse, "The smart thermostat: Using occupancy sensors to save energy in homes", in *SenSys '10*. 2010, ACM.

[7] Corinna Fischer, "Feedback on household electricity consumption: a tool for saving energy?", *Energy efficiency*, vol. 1, no. 1, pp. 79–104, 2008.

[8] Anil Aswani, Neal Master, Jay Taneja, David Culler, and Claire Tomlin, "Reducing transient and steady state electricity consumption in hvac using learning-based model-predictive control", *Proceedings of the IEEE*, vol. 100, no. 1, pp. 240–253, 2011.

[9] Siddharth Goyal, Herbert A Ingley, and Prabir Barooah, "Occupancy-based zone-climate control for energy-efficient buildings: Complexity vs. performance", *Applied Energy*, vol. 106, pp. 209–221, 2013.

[10] Siddharth Goyal, Herbert A Ingley, and Prabir Barooah, "Zone-level control algorithms based on occupancy information for energy efficient buildings", in *ACC*, June 2012, pp. 3063–3068.

[11] Yudong Ma and Francesco Borrelli, "Fast stochastic predictive control for building temperature regulation", in *ACC*. IEEE, 2012.

[12] Varick L Erickson, Miguel Á Carreira-Perpiñán, and Alberto E Cerpa, "Observe: Occupancy-based system for efficient reduction of hvac energy", in *IPSN '11*. IEEE, 2011.

[13] Varick L. Erickson, Stefan Achleitner, and Alberto E. Cerpa, "Poem: Power-efficient occupancy-based energy management system", in *IPSN '13*. 2013, ACM.

[14] Alessandra Parisio, Marco Molinari, Damiano Varagnolo, and Karl Henrik Johansson, "A scenario-based predictive control approach to building hvac management systems", in *2013 CASE*. IEEE, 2013, pp. 428–435.

[15] Alessandra Parisio, Damiano Varagnolo, Marco Molinari, Giorgio Pattarello, Luca Fabietti, and Karl H Johansson, "Implementation of a scenario-based mpc for hvac systems: an experimental case study", *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 599–605, 2014.

[16] Alessandra Parisio, Damiano Varagnolo, Daniel Risberg, Giorgio Pattarello, Marco Molinari, and Karl H. Johansson, "Randomized model predictive control for hvac systems", in *BuildSys'13*. 2013, ACM.

[17] Muhammad Waseem Ahmad, Monjur Mourshed, David Mundow, Mario Sisinni, and Yacine Rezgui, "Building energy metering and environmental monitoring–a state-of-the-art review and directions for future research", *Energy and Buildings*, vol. 120, pp. 85–102, 2016.

[18] Naren Srivaths Raman, Karthikeya Devaprasad, and Prabir Barooah, "Mpc-based building climate controller incorporating humidity", in *2019 American Control Conference (ACC)*, July 2019, pp. 253–260.

[19] Timothy W. Hnat, Erin Griffiths, Ray Dawson, and Kamin Whitehouse, "Doorjamb: Unobtrusive room-level tracking of people in homes using doorway sensors", in *SenSys '12*, 2012.

[20] Sujay Narayana, R. Venkatesha Prasad, Vijay S. Rao, T. V. Prabhakar, Sripad S. Kowshik, and Madhuri Sheethala Iyer, "Pir sensors: Characterization and novel localization technique", in *IPSN 15*, 2015.

[21] Danny B Yang, Leonidas J Guibas, et al., "Counting people in crowds with a real-time network of simple image sensors", in *ICCV*. IEEE, 2003.

[22] Sirajum Munir, Ripudaman Singh Arora, Craig Hesling, Juncheng Li, Jonathan Francis, Charles Shelton, Christopher Martin, Anthony Rowe, and Mario Berges, "Real-time fine grained occupancy estimation using depth sensors on arm embedded platforms", in *RTAS*. IEEE, 2017.

[23] Kin Sum Liu, Elvin Vindel Pinto, Sirajum Munir, Jonathan Francis, Charles Shelton, Mario Berges, and Shan Lin, "Cod: A dataset of commercial building occupancy traces", in *BuildSys '17*, 2017.

[24] Kin Sum Liu, Sirajum Munir, Jonathan Francis, Charles Shelton, and Shan Lin, "Long term occupancy estimation in a commercial space: an empirical study", in *IPSN*. IEEE, 2017, pp. 307–308.

[25] Tomasz T Gorecki, Faran A Qureshi, and Colin N Jones, "Openbuild: An integrated simulation environment for building control", in *2015 IEEE Conference on Control Applications (CCA)*. IEEE, 2015, pp. 1522–1527.

[26] Luca Fabietti, Tomasz T Gorecki, Faran A Qureshi, Altuğ Bitlislioğlu, Ioannis Lymperopoulos, and Colin N Jones, "Experimental implementation of frequency regulation services using commercial buildings", *IEEE Transactions on Smart Grid*, vol. 9, no. 3, pp. 1657–1666, 2016.

[27] Kin Sum Liu, Sirajum Munir, Shan Lin, and Charles Shelton, "Understanding occupancy patterns in a commercial space", *IEEE Technical Committee on Cyber-Physical Systems*, vol. 3, no. 1, 2018.

[28] Bingqing Chen, Zicheng Cai, and Mario Bergés, "Gnu-rl: A precocial reinforcement learning solution for building hvac control using a differentiable mpc policy", in *BuildSys 19*. 2019, ACM.

[29] Bjarne W Olesen and KC Parsons, "Introduction to thermal comfort standards and to the proposed new version of en iso 7730", *Energy and buildings*, vol. 34, no. 6, pp. 537–548, 2002.

[30] Khadija Hafeez, Yasra Chandio, Abu Bakar, Ayesha Ali, Affan A. Syed, Tariq M. Jadoon, and Muhammad Hamad Alizai, "Inverting hvac for energy efficient thermal comfort in populous emerging countries", in *BuildSys '17*. 2017, ACM.

[31] Jin Dong, Christopher Winstead, James Nutaro, and Teja Kuruganti, "Occupancy-based hvac control with short-term occupancy prediction algorithms for energy-efficient buildings", *Energies*, vol. 11, no. 9, pp. 2427, 2018.

[32] Alex Beltran and Alberto E. Cerpa, "Optimal hvac building control with occupancy prediction", in *BuildSys '14*. 2014, ACM.

[33] Jie Zhao, Ray Yun, Bertrand Lasternas, Haopeng Wang, Khee Poh Lam, Azizan Aziz, and Vivian Loftness, "Occupant behavior and schedule prediction based on office appliance energy consumption data mining", in *CISBAT 2013 Conference-Clean Technology for Smart Cities and Buildings*, 2013, pp. 549–554.

[34] Gurobi, 2018, http://www.gurobi.com/.

[35] Drury B Crawley, Linda K Lawrie, Frederick C Winkelmann, Walter F Buhl, Y Joe Huang, Curtis O Pedersen, Richard K Strand, Richard J Liesen, Daniel E Fisher, Michael J Witte, et al., "Energyplus: creating a new-generation building energy simulation program", *Energy and buildings*, vol. 33, no. 4, pp. 319–331, 2001.

[36] "CBE Thermal Comfort Tool. http://comfort.cbe.berkeley.edu/", 2019.

[37] Fabricio Flores, Sirajum Munir, Matias Quintana, Anand Krishnan Prakash, and Mario Bergés, "Dataset: Occupancy detection, tracking, and estimation using a vertically mounted depth sensor", in *Proceedings of the 2nd Workshop on Data Acquisition To Analysis*, New York, NY, USA, 2019, DATA19, ACM.

[38] KE Mary Reena, Abraham T Mathew, and Lillykutty Jacob, "Real-time occupancy based hvac control using interval type-2 fuzzy logic system in intelligent buildings", in *2017 12th IEEE Conference on Industrial Electronics and Applications (ICIEA)*, June 2017, pp. 455–460.

[39] Zhanhong Jiang, Jonathan Francis, Anit Kumar Sahu, Sirajum Munir, Charles Shelton, Anthony Rowe, and Mario Bergés, "Data-driven thermal model inference with armax, in smart environments, based on normalized mutual information", in *ACC*. IEEE, 2018.

[40] Jianjun Hu and Panagiota Karava, "Model predictive control strategies for buildings with mixed-mode cooling", *Building and Environment*, vol. 71, pp. 233–244, 2014.

[41] Ruoxi Jia, Roy Dong, S. Shankar Sastry, and Costas J. Spanos, "Privacy-enhanced architecture for occupancy-based hvac control", in *ICCPS'17*. 2017, ACM.

[42] Daniel A. Winkler, Alex Beltran, Niloufar P. Esfahani, Paul P. Maglio, and Alberto E. Cerpa, "Forces: Feedback and control for occupants to refine comfort and energy savings", in *UbiComp '16*. 2016, ACM.

[43] Varick L. Erickson and Alberto E. Cerpa, "Thermovote: Participatory sensing for efficient building hvac conditioning", in *BuildSys '12*. ACM.

[44] Bharathan Balaji, Hidetoshi Teraoka, Rajesh Gupta, and Yuvraj Agarwal, "Zonepac: Zonal power estimation and control via hvac metering and occupant feedback", in *BuildSys'13*. 2013, ACM.

[45] Farrokh Jazizadeh and Burcin Becerik-Gerber, "Toward adaptive comfort management in office buildings using participatory sensing for end user driven control", in *BuildSys '12*. 2012, ACM.

[46] Parisa Mansourifard, Farrokh Jazizadeh, Bhaskar Krishnamachari, and Burcin Becerik-Gerber, "Online learning for personalized room-level thermal control: A multi-armed bandit framework", in *BuildSys'13*. 2013, ACM.

[47] Jonathan Francis, Matias Quintana, Nadine von Frankenberg, Sirajum Munir, and Mario Bergés, "Occutherm: Occupant thermal comfort inference using body shape information", in *BuildSys 19*. 2019, ACM.

[48] Alex Beltran, Varick L. Erickson, and Alberto E. Cerpa, "Thermosense: Occupancy thermal based sensing for hvac control", in *BuildSys'13*. 2013, ACM.

[49] Bharathan Balaji, Jian Xu, Anthony Nwokafor, Rajesh Gupta, and Yuvraj Agarwal, "Sentinel: Occupancy based hvac actuation using existing wifi infrastructure within commercial buildings", in *SenSys'13*. 2013, ACM.

[50] Meiyi Ma, Sarah M. Preum, Mohsin Y. Ahmed, William Tärneberg, Abdeltawab Hendawi, and John A. Stankovic, "Data sets, modeling, and decision making in smart cities: A survey", *ACM Trans. Cyber-Phys. Syst.*, vol. 4, no. 2, Nov. 2019.