

Stochastic Optimal Control of HVAC system for Energy-efficient Buildings

Yu Yang, *Student Member, IEEE*, Guoqiang Hu, *Senior Member, IEEE*, and Costas J. Spanos, *Fellow, IEEE*

Abstract—This paper studies the optimal *stochastic* policies for *energy-efficient* operation of HVAC systems: *maintaining better human thermal comfort with less energy use*. Our main contributions are outlined. *First*, we formulate the problem to optimize HVAC's energy cost while maintaining human thermal comfort as a Markov decision process (MDP). In particular, we involve the elaborate thermal comfort model, predicted mean vote (PMV) that directly reflects human thermal sensation in the problem; *Second*, we propose a gradient-based policy iteration (GBPI) method to learn the optimal *stochastic* policies for HVAC control subject to uncertain thermal demand. *Thrid*, we prove the method can converge to the optimal policies theoretically. The main characteristics of the proposed method are that: *i*) it can be implemented *off-line* by learning control policies from historical data (e.g., whether, occupancy, etc); that can reduce the high *on-line* computation cost, and *ii*) it can handle the nonlinear system dynamics and the non-analytical PMV model. The performance of the *stochastic* policies yield by the GBPI method is demonstrated through comparisons with the optimal policies obtained with assumed accurate predictions.

Index Terms—HVAC control, energy-efficient, Markov decision process (MDP), stochastic policy, predicted mean vote (PMV).

NOMENCLATURE

Notations:

α_w	The absorption efficient of the wall;
A_{gs}	The area of glass window [m^2];
A_{wl}/A_{wr}	The area of left/right wall [m^2];
C_p	Air specific heat [$\text{J}/(\text{kg} \cdot \text{K})$];
C_w	The wall capacity [$\text{J}/(\text{kg} \cdot \text{K})$];
t	Time index;
c_t	The electricity price at time t [$\text{s}\$/\text{kW}$];
η	The reciprocal of coefficient of performance of chiller;
$G_t^{\text{FAU}}/G_t^{\text{FCU}}$	The supply air flow rate of FAU/FCU at time t [kg s^{-1}];
$G^{\text{FCU, Rated}}$	The nominal air flow rate of FCU [kg s^{-1}];
$G^{\text{FAU, Rated}}$	The nominal air flow rate of FAU [kg s^{-1}];

$\underline{G}^{\text{FAU}}/\underline{G}^{\text{FCU}}$	The lower bound of the supply air flow rate by FAU/FCU [kg s^{-1}];
$\bar{G}^{\text{FAU}}/\bar{G}^{\text{FCU}}$	The upper bound of the supply air flow rate by FAU/FCU [kg s^{-1}];
H_t^o	Outdoor relative humidity at time t [%];
H_t^a	Indoor relative humidity at time t [%];
H_g	The average humidity generation rate per occupant [kg s^{-1}];
H_t^{FAU}	The relative humidity of the supply air by the FAU [%];
h_{gs}	The heat transfer coefficient of glass window [$\text{J}/(\text{m}^2 \cdot ^\circ\text{C})$];
h_w	The heat transfer coefficient of walls [$\text{J}/(\text{m}^2 \cdot ^\circ\text{C})$];
m^a	The mass of indoor air [kg];
m^{wl}/m^{wr}	The mass of the left/right wall [kg];
$P^{\text{FCU, fan, Rated}}$	The nominal fan power of FCU [kW];
$P^{\text{FAU, fan, Rated}}$	The nominal fan power of FAU [kW];
Q_o	The average internal heat generation rate per occupant [J s^{-1}];
Q_t^{dev}	The average heat generation rate of devices caused by per occupant at time t [J s^{-1}];
Q_t^w	The solar radiation density on the wall at time t [$\text{J}/\text{m}^2 \cdot \text{s}$];
T_t^o	Outdoor temperature at time t [$^\circ\text{C}$];
T_t^a	Indoor temperature at time t [$^\circ\text{C}$];
T_t^{wl}/T_t^{wr}	The temperature of the interior left (right) wall [$^\circ\text{C}$];
$T_t^{\text{FAU}}/T_t^{\text{FCU}}$	The set-point temperature of FAU/FCU at time t [$^\circ\text{C}$];
$\underline{T}^{\text{FAU}}/\underline{T}^{\text{FCU}}$	The lower bound of the set-point temperature of FAU/FCU [$^\circ\text{C}$];
$\bar{T}^{\text{FAU}}/\bar{T}^{\text{FCU}}$	The upper bound of the set-point temperature of FAU/FCU [$^\circ\text{C}$].

Acronyms:

<i>HVAC</i>	Heating, ventilation and air-conditioning;
<i>MDP</i>	Markov decision process;
<i>MPC</i>	Model predictive control;
<i>GBPI</i>	Gradient-based policy iteration.

This work is supported by the Republic of Singapore's National Research Foundation through a grant to the Berkeley Education Alliance for Research in Singapore (BEARS) for the Singapore-Berkeley Building Efficiency and Sustainability in the Tropics (SinBerBEST) Program. BEARS has been established by the University of California, Berkeley as a center for intellectual excellence in research and education in Singapore.

Yu Yang is with SinBerBEST, Berkeley Education Alliance for Research in Singapore, Singapore 138602 e-mail: (yu.yang@bears-berkeley.sg).

Guoqiang Hu is with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore, 639798 e-mail: (gqhu@ntu.edu.sg).

Costas J. Spanos is with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, 94720 USA email: (spanos@berkeley.edu).

I. INTRODUCTION

HEATING, ventilation and air-conditioning (HVAC) systems account for a large proportion of the energy use in buildings [1, 2]. That can be dramatically attributed to the *insensible* operation patterns of HVAC systems, e.g., timed on-off switch, fixed temperature settings, etc. However, regardless

of the high energy consumption cost paid by the buildings, a large majority of occupants are still not satisfied with their thermal environment [3]. Therefore, it's imperative to pursue *energy-efficient HVAC operation: maintaining better human thermal comfort with less energy use*; it has become a critical issue for building's energy system management.

A. Literature

There exist various optimization-based control methods for HVAC systems to save energy cost [4]. We report the results from two aspects: *complex models* and *uncertainties* that have been addressed.

In the literature, some typical methods for HVAC control include sequential quadratic programming (SQP) [5], mixed-integer linear programming (MILP) [6, 7], genetic algorithms [8], and rule-based strategies [9, 10]. These works have been mainly dedicated to addressing the computational challenges imposed by the *complex models*, especially the nonlinear system dynamics. In particular, some simplification or approximation techniques have been discussed for computation (see [5, 9, 10] for examples).

Another challenge related to HVAC control is the multiple *uncertainties*, such as the outdoor weather and the indoor occupancy that determine the dynamic thermal demand. To address the *uncertainties*, MPC formulations have been widely adopted in the literature [4]. The general idea is to compute the control inputs at each executive epoch based on the updated measurements and the short-term future predictions. These methods are generally computationally efficient, however the applications may suffer limitations in two aspects. *First*, the simplified or approximated models used fail to capture the real nonlinear system dynamics [6]; *Second*, the required accurate predictions are not available [11]. To compensate for the deficiencies, especially the dependence on prediction accuracy, stochastic MPC [12, 13] and explicit MPC [14, 15] have been explored to explicitly incorporate the uncertainties in HVAC control. They minimize the average energy cost under uncertainties while guaranteeing a confidential level of thermal comfort using chance constraints. These MPC variations can alleviate the conservatism of standard MPCs but at an expense of increased computational challenges, which is actually still a cumbersome issue. For stochastic MPCs, the general solution is to transform the stochastic optimization to a deterministic one with a selected group of scenarios representing the uncertainties [12]. In contrast to stochastic MPCs with high *on-line* computation cost, explicit MPCs seek to reduce *on-line* computation cost for HVAC control by exploring *off-line* structure construction of the optimal control [14]. MPC is a general formulation for the control of dynamic systems under uncertainties. However, their applications to HVAC systems are generally hampered from the complex system dynamics. Moreover, they are oriented towards finding deterministic control policies, whereas the optimal policy is often stochastic under uncertainties [16].

Apart from MPCs, Markov decision process (MDP) is another general framework for sequential decision-making under uncertainties [17, 18]; and it has been explored in HVAC

control (see [2, 19, 20] for examples). Compared with MPCs, one advantage of MDP is that it can accommodate the complex system dynamics for the optimal control policies are generally computed by performance evaluation not function optimization. However, their applications to HVAC control still face challenges as the traditional solution methods like dynamic programming (DP) [18, 18] is computationally intractable due to the large state and action spaces.

While the energy saving potential of HVAC systems has been extensively demonstrated, some underlying issues still need further investigations. One one hand, a computational framework that can handle the *nonlinear system dynamics* and *uncertainties* efficiently is required so as to achieve the energy saving target; on the other hand, the current practices mostly use static temperature ranges to indicate human thermal comfort (see [12, 13] and the references therein). Such settings can reduce complexity but tend to result in unsatisfied thermal environment [3]; for the occupants' thermal sensation is determined by multiple parameters: indoor air temperature, mean radiant temperature, humidity, air velocity, metabolic activity and clothing insulation etc. [21]. Therefore, it's imperative to integrate an elaborate thermal comfort model, such as the well-known predicted mean vote (PMV) model [21] to directly reflect human thermal sensation while studying energy cost saving. Moreover, it has been ascertained that the involvement of PMV model not only can enhance human thermal comfort but also can reduce energy cost by avoiding overcooling compared with conventional temperature oriented settings [1, 6, 15]. However, [1] focused on maintaining thermal comfort indicated by PMV metric not energy cost saving, which is different from our work to be addressed. [15] and [6] are some exceptions that integrated the PMV in HVAC control to reduce energy cost. However, the solution methods are complicated as various approximation techniques are required both to tackle the nonlinear system dynamics and the non-analytical PMV model.

B. Our Contributions

This paper studies the optimal *stochastic* control policies to optimize HVAC's energy cost via MDP. In particular, the PMV model is involved to enhance indoor thermal condition. Our main contribution are outlined. *First*, we formulate the problem as a MDP. *Second*, we propose a gradient-based policy iteration (GBPI) method to learn the *stochastic* control policies from historical data. *Third*, the convergence of the proposed method is proved. The main characteristics of the proposed method are that: *i)* it can be implemented *off-line* thus reducing the high *on-line* computation cost; and *ii)* it can handle the complex system dynamics and the non-analytical PMV model. We demonstrate the performance of the method through comparison with the *optimal* solutions attained with assumed accurate predictions.

The remainder is outlined. In Section II, we give the MDP formulation. In Section III, we introduce the GBPI method and study the convergence. In Section IV, we evaluate the method through applications. In Section V, we conclude this paper.

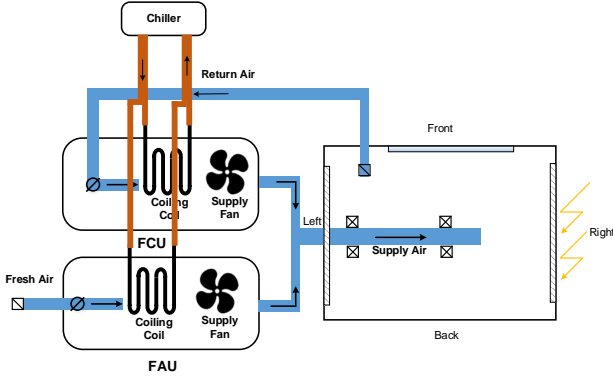


Fig. 1. The schematic of a typical HVAC system for a general room

II. THE PROBLEM

A. HVAC System

This paper focuses on one typical HVAC system mainly composed of air handling units: fresh air unit (FAU) and fan coil unit (FCU) and the chiller as depicted in Fig. 1 [2]. The air handling units first cools and dehumidifies the air to the set-points (temperature and humidity) and then force the supply air to the duct network by the fans (this paper investigates the cooling mode). The chiller provides chilled water to the cooling coils within the air handling units. In particular, this HVAC system uses FCU and FAU to handle the recirculated air and the fresh air separately, whereas some HVAC systems integrate them into one single air handling units [22]. This paper uses such HVAC system to study a general framework for HVAC control; however the formulation and the method can be extended to other types of HVAC systems like [22].

This paper investigates the HVAC control for an office room with the thermal demand mainly caused by the heat gain from outside and the indoor occupancy. We study the multivariable HVAC control including the supply air flow rates and the set-point temperature of the handling units systems with the objective to minimize daily energy cost while maintaining thermal comfort indicated by the PMV metric. Our problem is discussed in a discrete setting with a decision interval $\Delta t = 30$ mins (one day, $T = 48$ stages).

B. MDP

1) *System State*: The cooling demand of the room depends on the current indoor thermal condition, the outdoor weather and the thermal loads, therefore we define the system state as

$$S_t = [T_t^o, H_t^o, T_t^a, H_t^a, N_t^a]^\top$$

2) *Decision Variables*: We investigate the control of both the supply air flow rates and the set-point temperature for the FAU and FCU, thus we have the decision variables:

$$A_t = [G_t^{\text{FAU}}, T_t^{\text{FAU}}, G_t^{\text{FCU}}, T_t^{\text{FCU}}]^\top$$

3) *System Dynamics*: We adopt the widely-used gray-box models based on energy and mass conservation equations [2, 23] to capture the indoor thermal and humidity dynamics.

Besides, we establish some Markov chains to capture the outdoor weather and indoor occupancy uncertainties.

Indoor temperature: the indoor temperature is determined by the interplay of the HVAC operation and the thermal loads (the heat generations caused by the occupants and the heat gains from outside through the window and the walls), which can be described by the following equations [2, 23]:

$$\begin{aligned} C_p m^a (T_{t+1}^a - T_t^a) = & N_t^i Q_o \Delta t + P_t^{\text{dev}} \Delta t \\ & + h_{gs} A_{gs,i} (T_t^o - T_t^a) \Delta t \\ & + h_w A_{wl} (T_t^{wl} - T_t^a) \Delta t + h_w A_{wr} (T_t^{wr} - T_t^a) \Delta t \\ & + G_t^{\text{FAU}} (T_t^{\text{FAU}} - T_t^a) \Delta t + G_t^{\text{FCU}} (T_t^{\text{FCU}} - T_t^a) \Delta t. \end{aligned} \quad (1)$$

In (1), the first two terms capture the heat generations of occupants and electrical devices (e.g., laptops, monitors and desktops etc.). In particular, the electrical devices are manipulated by the occupants, therefore we use $P_t^{\text{dev}} = N_t^i Q_t^{\text{dev}}$ to estimate their heat generations. The subsequent three terms calculate the heat gains from outside through the glass window and the walls, and the last two terms indicate the cooling load supplied by the FAU and FCU of the HVAC system.

The thermal dynamics of the walls can be captured by

$$\begin{aligned} C_w m^{wl} (T_{t+1}^{wl} - T_t^{wl}) = & h_w A_{wl} (T_t^a - T_t^{wl}) \Delta t \\ C_w m^{wr} (T_{t+1}^{wr} - T_t^{wr}) = & h_w A_{wr} (T_t^a - T_t^{wr}) \Delta t \\ & + \alpha_w A_{wr} Q_t^w \Delta t \end{aligned} \quad (2)$$

Indoor humidity: Similarly, the indoor (relative) humidity dynamics can be described as [2]

$$\begin{aligned} m^a (H_{t+1}^a - H_t^a) = & N_t^a H_g \Delta t + G_t^{\text{FAU}} (H_t^{\text{FAU}} - H_t^a) \Delta t \\ & + G_t^{\text{FCU}} (H_t^{\text{FCU}} - H_t^a) \Delta t \end{aligned} \quad (3)$$

where the first term indicates the humidity generations of the occupants and the other two terms reflect the regulation of HVAC control. The (relative) humidity H_t^{FAU} of supply air by the FAU can be estimated by $H_t^{\text{FAU}} = \min(H_t^o, H_t^{\text{FAU,Sat}})$ with $H_t^{\text{FAU,Sat}}$ denoting the saturation point.

Indoor occupancy: Based on [24, 25], the following Markov chain is used to capture the indoor occupancy patterns:

$$\Pr(N_{t+1}^a = j | N_t^a = i) = p_t^{N,i,j}, \quad \forall i, j, \in \{1, \dots, L_A\}. \quad (4)$$

where $p_t^{N,i,j}$ is the transition probability from occupancy level i to level j at time t . L_A is the number of occupancy levels.

Outdoor weather: Though uncertain, the outdoor weather generally shows some dynamic patterns. This paper uses two Markov chains to capture the uncertain outdoor temperature and humidity dynamics:

$$\begin{aligned} \Pr(T_{t+1}^o = j | T_t^o = i) = & p_t^{T,i,j}, \quad \forall i, j, \in \{1, \dots, L_T\}. \\ \Pr(H_{t+1}^o = j | H_t^o = i) = & p_t^{H,i,j}, \quad \forall i, j, \in \{1, \dots, L_H\}. \end{aligned} \quad (5)$$

where $p_t^{T,i,j}$ and $p_t^{H,i,j}$ denote the transition probability of outdoor temperature and (relative) humidity from level i to level j . In particular, we equally discretize the ranges of outdoor temperature and (relative) humidity into L_T and L_H segments. The transition probabilities $p_t^{T,i,j}$ and $p_t^{H,i,j}$ are estimated based on the historical weather data in Singapore, which will be illustrated in Section IV of this paper.

4) *Objective Function*: To improve energy efficiency, the expected total HVAC cost is selected as the objective, i.e.,

$$J = \mathbb{E} \left\{ \sum_{t=0}^{T-1} c_t (\eta(C_t^{\text{FCU}} + C_t^{\text{FAU}}) + P_t^{\text{FCU, fan}} + P_t^{\text{FAU, fan}}) \Delta t \right\} \quad (6)$$

As indicated in (6), the power consumption of the HVAC system is mainly composed of *i*) the cooling power C_t^{FCU} , C_t^{FAU} , and *ii*) the fan power $P_t^{\text{FCU, fan}}$, $P_t^{\text{FAU, fan}}$. Wherein each part can be estimated by [2]

$$C_t^{\text{FCU}} = C_p G_t^{\text{FCU}} (T_t^a - T_t^{\text{FCU}}) + C_p G_t^{\text{FCU}} (H_t^a (2500 + 1.84 T_t^a) - H_t^{\text{FCU}} (2500 + 1.84 T_t^{\text{FCU}})) \quad (7a)$$

$$C_t^{\text{FAU}} = C_p G_t^{\text{FAU}} (T_t^o - T_t^{\text{FAU}}) + C_p G_t^{\text{FAU}} (H_t^o (2500 + 1.84 T_t^o) - (2500 + 1.84 T_t^{\text{FAU}})) \quad (7b)$$

$$P_t^{\text{FCU, fan}} = P^{\text{FCU, fan, Rated}} \cdot \left(\frac{G_t^{\text{FCU}}}{G^{\text{FCU, Rated}}} \right)^3 \quad (7c)$$

$$P_t^{\text{FAU, fan}} = P^{\text{FAU, fan, Rated}} \cdot \left(\frac{G_t^{\text{FAU}}}{G^{\text{FAU, Rated}}} \right)^3 \quad (7d)$$

5) *PMV Metric*: The fundamental functionality of buildings' HVAC system is to maintain human thermal comfort. To achieve such objective, this paper uses the well-known PMV metric [21] to indicate the occupants' thermal sensation. The PMV model is non-analytical and characterized by multivariables, and its calculation relies on numerical iterations. We use $PMV(\cdot)$ to implicitly describe such model which is constituted by a group of equations [6]:

$$\text{pmv}_t = PMV(M, W, T_t^a, H_t^a, t_t^r, v_t^a, I_{cl}) \quad (8)$$

where the input parameters of the PMV model include: metabolic rate M (W/m²), mechanic work intensity W (W/m²), indoor air temperature T_t^a (°C), relative humidity H_t^a (%), mean radiation temperature t_t^r (°C), indoor air velocity v_t^a (m/s), and clothing insulation I_{cl} (m²K/W). The PMV model reflects thermal comfort by establishing the mapping of indoor environment to the occupants' average satisfaction within the range $[-3, 3]$, with -3 , 0 , 3 indicating too cold, ideal, and too hot, respectively.

6) *Constraints*: The operation of the HVAC system should comply with the operation limits: *(i)* the lower and upper bound of supply air flow rates imposed by the dampers within FAU and FCU (9a); *(ii)* the set-point temperature range of FAU and FCU determined by chiller capacity (9b).

$$\underline{G}^{\text{FAU}} \leq G_t^{\text{FAU}} \leq \overline{G}^{\text{FAU}}, \quad \underline{G}^{\text{FCU}} \leq G_t^{\text{FCU}} \leq \overline{G}^{\text{FCU}}. \quad (9a)$$

$$\underline{T}^{\text{FAU}} \leq T_t^{\text{FAU}} \leq \overline{T}^{\text{FAU}}, \quad \underline{T}^{\text{FCU}} \leq T_t^{\text{FCU}} \leq \overline{T}^{\text{FCU}}. \quad (9b)$$

We use $\underline{\text{pmv}}$ and $\overline{\text{pmv}}$ to capture the desirable environment indicated by the PMV metric, thus we have,

$$\underline{\text{pmv}} \leq \text{pmv}_t \leq \overline{\text{pmv}} \quad (10)$$

7) *Optimization Problem*: Overall, the optimal operation of the HVAC system subject to uncertain thermal demand can be depicted as a MDP:

$$\min_{\pi} J(S_0, \pi) = \mathbb{E}^{\pi} \left\{ \sum_{t=0}^{T-1} c_t (\eta(C_t^{\text{FCU}} + C_t^{\text{FAU}}) + P_t^{\text{FCU, fan}} + P_t^{\text{FAU, fan}}) \Delta t \right\} \quad (11)$$

subject to System dynamics: (1) – (5), Operation limits: (9), Thermal comfort: (10), $\forall t \in \{0, 1, \dots, T-1\}$.

where $\pi = (\pi_0, \pi_1, \dots, \pi_{T-1})$ denotes the *stochastic* control policy for HVAC system over the optimization horizon. S_0 is the initial system state. At each time t , the control rule $\pi_t: S_t \rightarrow \mathcal{A}_t$ establishes a mapping from the state space S_t to the action space \mathcal{A}_t . \mathbb{E}^{π} denotes the expectation under policy π .

It's nontrivial to search for an optimal policy for problem (11) concerning that:

- (i) the multiple uncertainties make it intractable to analytically evaluate the policies under expectation;
- (ii) the large state and the action space of the multi-stage decision problem result in intensive computation burden.

The standard solution methods for finite-stage MDPs like dynamic programming (DP) [18] are computationally intractable as they require to traverse the Q-factors for the large state and action spaces.

III. GRADIENT-BASED POLICY ITERATION

To tackle problem (11), this section proposes a gradient-based policy iteration (GBPI) method based on performance gradients. Generally, the main idea is to iteratively update a *stochastic* policy based on the performance gradients estimated using Monte Carlo (MC) simulation [26]. In particular, two strategies are established to reduce computation burden based on data analysis. The remainder of this section introduces the main notations and the establishment of the method.

A. Notations

We use the lower cases s_t and a_t (b_t , c_t) to represent a state and action instances at time t . We use the integer sets $S_t \triangleq \{1, 2, \dots, |S_t|\}$ and $\mathcal{A}_t \triangleq \{1, 2, \dots, |\mathcal{A}_t|\}$ to represent the *feasible* state and action spaces at time t . We use $|\cdot|$ to indicate the cardinality of a set. $\theta = (\theta_0, \theta_1, \dots, \theta_{T-1})^T$ denotes a stochastic policy; $\theta_t \in \mathbb{R}^{|S_t| \times |\mathcal{A}_t|}$ establishes a mapping from the state space S_t to the action space \mathcal{A}_t with $\theta_t(s_t, a_t) \in [0, 1]$ denoting the probability to take action a_t at state s_t under policy θ . The lower cases $p_t(s_{t+1}|s_t, a_t)$ and $p_t^{\theta}(s_{t+1}|s_t)$ indicate the transition probability from state s_t to state s_{t+1} under action a_t or policy θ , and $P_t^{\theta} = [p_t^{\theta}(s_{t+1}|s_t)]_{|S_t| \times |S_{t+1}|}$ indicate the transition probability matrix. The superscript k denotes the iteration.

B. GBPI

Since there exist various constraints in problem (11), we first eliminate the constraints by penalizing the violations of constraints in the objective function via some indicator functions $\mathbb{I}(\cdot)$. Therefore, we have the augmented stage-cost:

$$r_t(s_t, a_t) = c_t (\eta(C_t^{\text{FCU}} + C_t^{\text{FAU}}) + P_t^{\text{FCU, fan}} + P_t^{\text{FAU, fan}}) \Delta t + \mathbb{I}_t(\mathcal{X}_t) \quad (12)$$

where we use \mathcal{X}_t to indicate the constraints at each t constructed by (1)-(5), (9) and (10) as clarified in problem (11). For the indicator function, we $\mathbb{I}(A) = 1$ with the condition A true, otherwise $\mathbb{I}(A) = 0$.

Remark 1. We want to articulate that this penalization of constraint violations is for the subsequent analysis, whereas in practical implementation, the feasibility issue is handled

by alternatively evaluating and updating the policies while performing the algorithm, which is to be addressed later.

Our method is stimulated by the performance difference equation for any two stochastic policies σ and μ [27]:

$$J(\mu; S_0) - J(\sigma; S_0) = \sum_{t=0}^{T-1} \pi_t^\mu \left[(r_t^\mu - r_t^\sigma) + (P_t^\mu - P_t^\sigma) V_{t+1}^\sigma \right] \quad (13)$$

where $\pi_t^\mu = (\pi_t^\mu(1), \pi_t^\mu(2), \dots, \pi_t^\mu(|S_t|))^T$ denotes the state distribution at time t under policy μ . $r_t^\theta = (r_t^\theta(1), r_t^\theta(2), \dots, r_t^\theta(|S_t|))^T$, $V_{t+1}^\theta = (V_{t+1}^\theta(1), \dots, V_{t+1}^\theta(|S_{t+1}|))^T$ and P_t^θ denote the one-step cost, the performance potentials and the transition probability matrices under a specific policy θ ($\theta \in \{\sigma, \mu\}$). In particular, the performance potential is captured by

$$V_{t+1}^\theta(s_{t+1}) = \mathbb{E}^\theta \left\{ \sum_{\tau=t+1}^{T-1} r_\tau(s_\tau, a_\tau) \right\}, \text{ with } a_\tau = \theta(s_\tau). \quad (14)$$

We note that (13) can quantify the performance gap for any two policies μ and σ . In terms of perturbation analysis (PA), we can view σ and μ as the base and perturbed policy. We find that the performance of the perturbed policy μ can be calculated only if the performance potentials (V_{t+1}^σ) under the base policy σ and the state distribution (π_t^μ) under the perturbed policy μ can be identified. However, as the latter is generally unknown *a priori*, it's impractical to update the policy based on (13).

To further explore the information in (13) for performance improvement, we define a structured random policy δ that adopts policy σ with probability δ and policy μ with probability $1-\delta$. For policy δ , we can figure out the state transition probability as $P_t^\delta = P_t^\sigma + \delta \Delta P_t$ with $\Delta P_t = P_t^\mu - P_t^\sigma$, and the stage-cost vectors $r_t^\delta = r_t^\sigma + \delta \Delta r_t$ with $\Delta r_t = r_t^\mu - r_t^\sigma$. Thus, the performance difference equation (13) amounts to

$$J(\delta; S_0) - J(\sigma; S_0) = \sum_{t=0}^{T-1} \pi_t^\delta \left[\Delta r_t + \Delta P_t V_{t+1}^\sigma \right] \quad (15)$$

We thus have the performance differential equation with $\delta \rightarrow 0$:

$$\frac{dJ(\delta; S_0)}{d\delta} = \lim_{\delta \rightarrow 0} \sum_{t=0}^{T-1} \pi_t^\delta \left[\Delta r_t + \Delta P_t V_{t+1}^\sigma \right] \quad (16)$$

Equivalently, we have

$$\begin{aligned} \frac{\partial J(\sigma; S_0)}{\partial \sigma_t(s_t, a_t)} &= \pi_t^\sigma(s_t) \left[\frac{\partial r_t^\sigma(s_t)}{\partial \sigma_t(s_t, a_t)} \right. \\ &\quad \left. + \frac{\partial p_t^\sigma(s_{t+1}|s_t)}{\partial \sigma_t(s_t, a_t)} V_{t+1}^\sigma(s_{t+1}) \right] \end{aligned} \quad (17)$$

By substituting $r_t^\sigma(s_t) = \sum_{a_t=1}^{|A_t|} p_t^\sigma(a_t|s_t) r_t(s_t, a_t)$, $p_t^\sigma(s_{t+1}|s_t) = \sum_{a_t=1}^{|A_t|} p_t^\sigma(a_t|s_t) p(s_{t+1}|s_t, a_t)$ into (17), we have

$$\begin{aligned} \frac{\partial J(\sigma; S_0)}{\partial \sigma_t(s_t, a_t)} &= \pi_t^\sigma(s_t) \left[\sum_{b_t=1}^{|A_t|} \frac{\partial p_t^\sigma(b_t|s_t)}{\partial \sigma_t(s_t, a_t)} r_t(s_t, b_t) \right. \\ &\quad \left. + \sum_{s_{t+1} \in S_{t+1}} \sum_{b_t=1}^{|A_t|} \frac{\partial p_t^\sigma(b_t|s_t)}{\partial \sigma_t(s_t, a_t)} p_t(s_{t+1}|s_t, b_t) V_{t+1}^\sigma(s_{t+1}) \right] \\ &= \pi_t^\sigma(s_t) \left[\sum_{b_t=1}^{|A_t|} \frac{\partial p_t^\sigma(b_t|s_t)}{\partial \sigma_t(s_t, a_t)} (r_t(s_t, b_t) + V_{t+1}^\sigma(s_t, b_t)) \right] \end{aligned} \quad (18)$$

where $V_{t+1}^\sigma(s_t, b_t) = \sum_{s_{t+1} \in S_{t+1}} p_t(s_{t+1}|s_t, b_t) V_{t+1}^\sigma(s_{t+1})$.

As $p_t^\sigma(b_t|s_t) = \frac{\sigma_t(s_t, b_t)}{\sum_{c_t=1}^{|A_t|} \sigma_t(s_t, c_t)}$, we have

$$\frac{\partial p_t^\sigma(b_t|s_t)}{\partial \sigma_t(s_t, a_t)} = \begin{cases} \frac{\sum_{c_t=1}^{|A_t|} \sigma_t(s_t, c_t) - \sigma_t(s_t, a_t)}{[\sum_{c_t=1}^{|A_t|} \sigma_t(s_t, c_t)]^2}, & b_t = a_t \\ \frac{-\sigma_t(s_t, b_t)}{[\sum_{c_t=1}^{|A_t|} \sigma_t(s_t, c_t)]^2}, & b_t \neq a_t \end{cases} \quad (19)$$

The equation (18) can be interpreted as the performance gradients of policy σ and we thus establish iterative procedure to update policies as

$$\sigma^{k+1} = \sigma^k - \gamma^k \cdot \nabla_{\sigma^k} J(\sigma^k; S_0) \quad (20)$$

where we have $\nabla_{\sigma^k} J(\sigma^k; S_0) = \left[\frac{\partial J(\sigma^k; S_0)}{\partial \sigma_t^k(s_t, a_t)} \right]$. $\gamma^k = [\gamma_t^k(s_t, a_t)]$ denotes the step-size at iteration k .

Observe (20), we note two problems to be addressed. (i) computing the performance gradients $\nabla_{\sigma^k} J(\sigma^k; S_0)$; (ii) determining the step-size γ^k . As there exists randomness, it's difficult to analytically estimate the performance gradients under expectation. To overcome this difficulty, we use the MC simulation technique [26] to estimate the performance gradients $\nabla_{\sigma^k} J(\sigma^k; S_0)$ for any given policy σ . The main procedures are structured in **Algorithm 1**. We use $\Omega_t(\cdot)$ to denote the sets of states or actions and $\mathcal{I}_t(\cdot)$ indicates the set of sample path indexes. The algorithm consists of two main steps: i) generate a number of sample paths by executing the policy σ via MC simulation; ii) estimate the buildings blocks $\pi_t^\sigma(s_t)$, $r_t(s_t, b_t)$ and $V_{t+1}^\sigma(s_t, b_t)$ of the performance gradients based on the sample paths. For the step-size γ^k , it is closely related to the convergence and we have the main results in **Theorem 1**.

Theorem 1. For any given initial policy σ^0 , the GBPI method can converge to an optimal policy of problem (11) with the selected stepsize $\gamma^k(s_t, a_t) = \frac{\sigma_t^k(s_t, a_t)}{\sum_{c_t=1}^{|A_t|} \sigma_t^k(s_t, c_t)}$ ($\forall s_t \in S_t, a_t \in A_t$) and the performance gradients $\nabla_{\sigma^k} J(\sigma^k; S_0)$ estimated accurately enough.

The proof refers to **Appendix A**.

Remark 2. As illustrated in **Theorem 1**, the convergence requires some estimation accuracy of the performance gradients. Generally, any sufficiently accurate estimation can be approached by increasing the number of sample paths. However, the practical implementation only requires an estimations that preserve the (performance) order of the (action) candidates.

An example with two action candidates a and b , we only require an estimation $\tilde{V}(a) \leq \tilde{V}(b)$ if $V(a) \leq V(b)$ where $V(\cdot)$ and $\tilde{V}(\cdot)$ denote the real and predicted value.

Remark 3. The convergence of the method does not depend on the initial stochastic policy σ^0 . This can be identified from the proof. However, a better initial policy can speed up the convergence and reduce computation cost in implementation.

The overall procedures to perform GBPI to learn the optimal stochastic policy for problem (11) are organized in **Algorithm 2**. We select the stopping criterion as $\|\nabla_{\sigma^k} J(\sigma^k; S_0)\|_2 \leq \epsilon$ (ϵ is a positive threshold). Particularly, for theoretical analysis, we penalize constraint violations in the objective via indicator functions; whereas for practical implementation (Algorithm 2),

the feasibility issue is pursued by alternatively checking the feasibility (only the PMV metric) and updating the policy σ : set $\sigma_t(s_t, \mathbf{a}_t) = 0$ for any visited state-action pairs (s_t, \mathbf{a}_t) that violate the PMV comfortable range while performing the MPC simulation (Step 3). The underlying idea is to assign low or zero probability for the actions that result in discomfort.

One may note one advantage of the proposed method is that it can be implemented *off-line* based on the historical data, thus greatly reducing the high *on-line* computations in implementation. More specifically, as the control policies have been learned *off-line*, the main computation for the *on-line* implementation is to look up the policy table and figure out the right action based on the measured system state (i.e., indoor/outdoor temperature, humidity and occupancy).

From **Algorithm 2**, we note that one advantage of the proposed method is that it can be implemented *off-line* thus reducing the high *on-line* computation cost. More specifically, if the control policy is learned *off-line*, the *on-line* implementation only requires identifying the right action based on the system state.

Algorithm 1 Estimate performance gradients by MC method

- 1: **Input:** a given policy σ .
- 2: Generate $|\mathcal{W}|$ *feasible*¹ sample paths by executing policy σ and index each sample path as

$$\{s_0^\omega, \mathbf{a}_0^\omega, r_0^\omega, s_1^\omega, \mathbf{a}_1^\omega, r_1^\omega \dots, s_{T-1}^\omega, \mathbf{a}_{T-1}^\omega, r_{T-1}^\omega\}, \forall \omega \in \mathcal{W}.$$

- 3: **For** $t \in \{0, 1, \dots, T-1\}$ **do**
- 4: Count the occurrence of states s_t^ω and actions \mathbf{a}_t in the sample paths by

$$\begin{aligned} \Omega_t(s_t) &= \{s_t^\omega \mid \omega \in \mathcal{W}\}. \\ \Omega_t(\mathbf{a}_t | s_t) &= \{\mathbf{a}_t^\omega \mid s_t^\omega = s_t, \omega \in \mathcal{W}\}, \forall s_t \in \Omega_t(s_t). \\ \mathcal{I}_t(s_t) &= \{\omega \mid \omega \in \mathcal{W}, s_t^\omega = s_t\}, \forall s_t \in \Omega_t(s_t). \\ \mathcal{I}_t(s_t, \mathbf{a}_t) &= \{\omega \mid \omega \in \mathcal{W}, s_t^\omega = s_t, \mathbf{a}_t^\omega = \mathbf{a}_t\}, \\ &\quad \forall s_t \in \Omega_t(s_t), \mathbf{a}_t \in \Omega_t(\mathbf{a}_t | s_t). \end{aligned}$$

- 5: Estimate the building blocks $\pi_t^\sigma(s_t)$, $r_t(s_t, \mathbf{a}_t)$ and $V_t^\sigma(s_t, \mathbf{a}_t)$ according to

$$\begin{aligned} \pi_t^\sigma(s_t) &\approx |\mathcal{I}(s_t)|/|\mathcal{W}|, \forall s_t \in \Omega_t(s_t). \\ r_t(s_t, \mathbf{a}_t) &\approx \frac{1}{|\mathcal{I}(s_t, \mathbf{a}_t)|} \sum_{\omega \in \mathcal{I}(s_t, \mathbf{a}_t)} r_t^\omega, \\ &\quad \forall s_t \in \Omega_t(s_t), \mathbf{a}_t \in \Omega_t(\mathbf{a}_t | s_t). \\ V_t^\sigma(s_t, \mathbf{a}_t) &\approx \frac{1}{|\mathcal{I}(s_t, \mathbf{a}_t)|} \sum_{\omega \in \mathcal{I}(s_t, \mathbf{a}_t)} \sum_{\tau=t+1}^{T-1} r_\tau^\omega, \\ &\quad \forall s_t \in \Omega_t(s_t), \mathbf{a}_t \in \Omega_t(\mathbf{a}_t | s_t). \end{aligned}$$

- 6: Compute $\nabla_\sigma J(\sigma; S_0)$ using (18) and (19).
 - 7: **EndFor**
-

¹While generating the sample paths, we repeatedly check the thermal comfort constraints (10). If the thermal comfort is not satisfied at state s_t while taking action \mathbf{a}_t , we set $\sigma_t(s_t, \mathbf{a}_t) = 0$ and regenerate a new action based on the updated policy until the thermal comfort is satisfied.

Algorithm 2 Gradient-based Policy Iteration (GBPI)

- 1: **Initialization:** $k \rightarrow 0, \sigma^0$.
- 2: **Iteration:**
- 3: Estimate the gradients $\nabla_{\sigma^k} J(\sigma^k; S_0)$ using **Algorithm 1**.
- 4: Policy Update:

$$\sigma^{k+1} = \sigma^k - \gamma^k \cdot \nabla_{\sigma^k} J(\sigma^k; S_0) \quad (21)$$

- 5: Stop if the **stopping criterion** (??) is reached, otherwise $k = k + 1$ go to **Step 3**.
-

C. Computation Reduction

We discuss two strategies to reduce computation of the GBPI for HVAC control.

Strategy I: Though the temperature and humidity spread wide ranges (i.e., temperature: $[22, 34]^\circ\text{C}$ and relative humidity $[40, 100]\%$) within a whole day, they are generally concentrated within relative narrow ranges at each time period (e.g., at 6:00 am, the temperature and humidity are almost within $[26, 28]^\circ\text{C}$ and $[75, 95]\%$). These can be observed from the data analysis in Section IV-A. This implies that while estimating the performance gradients for state-action pairs to update policy at each stage, we can concentrate our computation only on the states within such narrow ranges. Since the states out of such ranges are extreme cases, this can greatly reduce computation cost while not compromising the performance of the policies.

Strategy II: the main computation to learn control policy using the GBPI method is to estimate the performance gradients for the state-action pairs. The operation of the HVAC system should maintain thermal comfort (avoid the occurrence of uncomfortable states), therefore we can initialize all the entries of σ^0 to *zero* corresponding to state-action pairs that result in human discomfort (i.e., indoor temperature out of $[23, 28]^\circ\text{C}$ or relatively humidity out of $[40, 70]\%$). The underlying idea is avoid the waste of computation and speed up the convergence by starting with a better initial policy.

IV. APPLICATIONS

The section illustrates the performance of the GBPI method applied to HVAC control. We organize this section into two parts. *First*, we study the characteristics of weather in Singapore based on historical data. In particular, we establish two Markov chains to capture the uncertainties of outdoor temperature and humidity, which are to be employed in the GBPI method to learn the *stochastic* control policy. *Second*, we compare the results of the *stochastic* policies yield by the GBPI with the optimal solution obtained with assumed accurate predictions in a number of case studies.

A. Data Analysis

This section studies the characteristic of weather based on the historical meteorological data in Singapore (from 2019/09/01 to 2019/10/13, 43 days). This data set records the temperature and humidity per minute. As an example, we plot the temperature and humidity curve on a typical

day (2019/09/24) in Fig. 2. The figure reflects some general characteristics of weather in tropical countries. Specifically, the temperature and (relative) humidity are within the ranges of $[22, 34]^\circ\text{C}$ and $[40, 100]\%$. The outdoor temperature is usually low in the early morning, but will gradually rise to the peak level in the midday hours. It then gradually decreases and return to its lowest point at late night. By contrast, the (relative) humidity generally presents an almost opposite trend.

As discussed in Section II, we establish two Markov chains to capture the patterns of outdoor temperature and (relative) humidity. It mainly includes two steps. We first equally discretize the temperature and humidity in the data set with a resolution of 1°C and 5% , respectively. After that we estimate the transition probabilities by

$$\begin{aligned} \text{Temperature: } p_t^{T,ij} &\approx \frac{\sum_{\omega=1}^{Day} \mathbb{I}(T_t^{o,\omega} = i, T_{t+1}^{o,\omega} = j)}{\sum_{i=1}^{Day} \mathbb{I}(T_t^{o,\omega} = i)}, \\ \forall i, j \in \{1, 2, \dots, L_T\}, t \in \{1, \dots, T-1\}. \\ \text{Relative humidity: } p_t^{H,ij} &\approx \frac{\sum_{\omega=1}^{Day} \mathbb{I}(H_t^{o,\omega} = i, H_{t+1}^{o,\omega} = j)}{\sum_{i=1}^{Day} \mathbb{I}(H_t^{o,\omega} = i)}, \\ \forall i, j \in \{1, 2, \dots, L_H\}, t \in \{1, \dots, T-1\}. \end{aligned}$$

where $\mathbb{I}(\cdot)$ is the indicator function. ω denotes the index of the day, and we have $Day = 43$.

Besides, we perform some statistical analysis on the distribution of outdoor temperature and (relative) humidity. As shown in Fig. 3, we observe some special characteristics: though the temperature and humidity spread wide ranges (i.e., temperature: $[22, 34]^\circ\text{C}$ and relative humidity $[40, 100]\%$) within a whole day, they are generally concentrated within relative narrow ranges at each time period (e.g., at 6:00 am, the temperature and humidity are almost within $[26, 28]^\circ\text{C}$ and $[75, 95]\%$). This implies that while performing the GBPI method to update policy at each stage, we can concentrate our computation only on the (weather) states within such narrow ranges. This motivates our *Strategy II* to reduce computation burden of the GBPI method applied to HVAC systems.




Fig. 2. (a) Outdoor temperature for a typical day. (b) Outdoor relative humidity for a typical day.

B. Case Studies

Parameter settings: We consider an office room of size $6\text{m} \times 5\text{m} \times 4\text{m}$. The maximum number of occupants is 5 and discretized into $L_A = 5$ levels. The comfortable PMV range ($[-0.5, 0.5]$) and the static inputs of the PMV model are typical and refer to the ANSI/ASHRAE Standard [28] (as shown in TABLE II). In particular, the mean radiant temperature

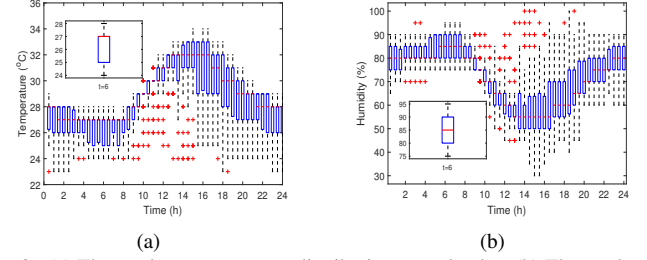


Fig. 3. (a) The outdoor temperature distribution over the day. (b) The outdoor relative humidity distribution over the day.

t_t^T is estimated by 2°C higher than the instantaneous indoor air temperature according to the standard [28]. The structural parameters of the room and thermal demand of the occupants refer to [2] and are gathered in TABLE I. The nominal parameters of the HVAC system are covered in TABLE II. We use the time-of-use (TOU) price in Singapore [6].

TABLE I
ROOM & OCCUPANT SETTINGS

Param.	Value & Units	Param.	Value & Units
C_p	$1012\text{J}/(\text{kg} \cdot \text{K})$	m^a	144.6kg
h_{gs}	$2.5\text{W}/\text{m}^2$	m^{wl}	$7.2 \times 10^3\text{kg}$
A_{gs}	10m^2	m^{wr}	$8.64 \times 10^3\text{kg}$
h_w	$0.8\text{W}/\text{m}^2$	C_w	$1.05 \times 10^3\text{J}/(\text{kg} \cdot \text{K})$
a_w	0.4	Q_o	40 J s^{-1}
A_{wl}	20m^2	H_g	0.03 g s^{-1}
A_{wr}	24m^2		

TABLE II
HVAC & PMV PARAMETERS

HVAC		PMV	
Param.	Value & Units	Param.	Value & Units
$P_{\text{FAU, fan, Rated}}$	0.1 K W	v^a	0.2m/s
$G_{\text{FAU, Rated}}$	0.01 kg s^{-1}	M	1.0 met
$P_{\text{FCU, fan, Rated}}$	0.1 K W	W	0
$G_{\text{FCU, Rated}}$	0.05 kg s^{-1}	I_{cl}	0.155 clo
COP	2.7	P_a	$1.01 \times 10^5\text{ Pa}$

In this part, we compare our results with the optimal solution in an MPC formulation with assumed accurate information over each planning horizon H (not available in practice); Such solution can be obtained by sequentially solving the deterministic problem (22):

$$\begin{aligned} \min_{G_t^{\text{FAU}}, T_t^{\text{FAU}}, G_t^{\text{FCU}}, T_t^{\text{FCU}}} J(t_k) &= \sum_{t=t_k+1}^{t_k+H-1} \{c_t(\eta(C_t^{\text{FCU}} + C_t^{\text{FAU}})) \\ &\quad + P_t^{\text{FCU, fan}} + P_t^{\text{FAU, fan}}\} \Delta t \\ \text{subject to} \quad &\text{System dynamics: (1) - (3),} \\ &\text{Operation limits: (9),} \\ &\text{Thermal comfort: (10),} \\ &\forall t \in \{t_k, t_k + 1, \dots, t_k + H - 1\}. \end{aligned} \quad (22)$$

It's nontrivial to solve problem (22) due to the nonlinear system dynamics and the constraints on the PMV index. We use the approximated solution method proposed in [6, 7] as it is one of the scarce methods available and capable to our best knowledge. The main idea is to reformulate problem (22) as a MILP by using various linearization and approximation techniques. We refer readers to [6, 7] for details.

For the two methods, we compare the results in three cases with different settings. In particular, the discretization of state

variables are only employed in the GBPI method not in the MPC approach.

Case I: The discretization intervals for temperature and relative humidity are 2°C and 10% , respectively (GBPI). The set-point temperature range of FAU and FCU are $\{12, 14, 16\}^{\circ}\text{C}$, and their supply air flow rates are equally divided into 3 levels.

Case II: The discretization of temperature and relative humidity is the same as *Case I* (GBPI). However, the set-point temperature of FAU and FCU are fixed as 15°C , and their supply air flow rates are equally divided into 5 levels.

Case III: The discretization intervals for temperature and the relative humidity are 1°C and 5% , respectively (GBPI). The settings for FAU and FCU is the same as *Case II*.

In the GBPI method, we use 1000 (*Case I*), 2000 (*Case II*) and 5000 (*Case III*) sample paths for performance gradients estimation. As there exist uncertainties, we compare the distributions of HVAC cost yield by the two methods under 100 randomly generated scenarios as shown in Fig. 4 (*Case I*), Fig. 5 (*Case II*) and Fig. 6 (*Case III*). The results imply that the average performance gaps in energy cost are around 11.7% (*Case I*), 12.9% (*Case II*) and 6.5% (*Case III*) in the three cases. These performance gaps are mainly attributed to three aspects: *i*) the discretization of state variables in the GBPI method to reduce computation (not in the MPC method); *ii*) the estimation accuracy of performance gradients limited by the number of sample paths used; *iii*) the unequal information used in the two methods (accurate vs. stochastic information). Therefore, the performance gaps may be somehow lessened by decreasing discretization stepsize or increasing the number of sample paths. This has been illustrated by inspecting the results in *Case II* and *Case III*. Specially, while the settings are the same except the discretization interval, we observe an apparent reduction in the HVAC cost with finer discretization (*Case III*); However, that is achieved at an expense of higher computation cost due to the larger state space and the increased number of sample paths required to estimate performance gradients. Besides, we imply from the comparative case studies (*Case I, II*) that the proposed method can be applied to both constant and variable temperature setpoints for FAU and FCU.




Fig. 4. The histograms of HVAC cost for *Case I*.

We investigate the thermal condition of the two methods by inspecting the indoor thermal condition indicated by temperature and humidity ranges as well as the PMV metric in a randomly selected scenario in *Case II*. As shown in Fig. 7, the indoor temperature and relative humidity are both maintained in the typical comfortable ranges $[24, 27]^{\circ}\text{C}$ and $[40, 70]\%$.




Fig. 5. The histograms of HVAC cost for *Case II*.




Fig. 6. The histograms of HVAC cost for *Case III*.

Moreover, the indoor thermal comfort are also confirmed by the PMV metrics ($[-0.5, 0.5]$) as in Fig. 8.

The above results are for a specific scenario where the PMV index is totally maintained within the comfortable range $[-0.5, 0.5]$ by the proposed GBPI method. However, that is not always the case due the uncertainties. To illustrate the confidential level of thermal comfort for the GBPI, we investigate the results for 100 randomly generated scenarios and plot the distribution of PMV metrics in Fig. 9. We observe that the PMV metrics are located within $[-0.5, 0.6]$; that implies that the thermal comfort can almost be guaranteed. More specifically, we find the thermal comfort (PMV $[-0.5, 0.5]$) is guaranteed with a probability of 93% in this case study.

For this specific scenario, we also contrast the control inputs of FAUs and FCUs in Fig. 10 and Fig. 11. We observe quite




Fig. 7. The indoor temperature and relative humidity curves for a randomly selected scenario under the two methods.




Fig. 8. The PMV curves for a randomly selected scenario under the two methods.




Fig. 9. The distribution of PMV metric of the GBPI in Case II.




Fig. 10. The control of FAU for a randomly selected scenario under the two methods.




Fig. 11. The control of FAU for a randomly selected scenario under the two methods.




Fig. 12. The convergence rate of GBPI (Case II)

similar operation patterns of FAUs and FCUs in the two methods. That further demonstrates the desirable performance of the proposed method compared with the optimal solution. Besides, we can observe some particular phenomenon in Fig. 10. *First*, we find that the FAUs are mostly operated in quite lower level compared with the FCUs. This is rational due to the generally higher temperature of outdoor fresh air (FAUs) compared than the recirculated air (FCUs), therefore the HVAC system tends to regulate the FCUs to maintain thermal comfort to save energy. Besides, we observe that the operation patterns, especially the FCUs correspond well to the typical weather patterns (i.e., lower outdoor temperature in the early morning and late night, and higher temperature in the noon) and the occupancy patterns (i.e., high occupancy during the working hours and low occupancy during non-working hours).

We also study the convergence of the GBPI in Case II. Specifically, we inspect the convergence rate of the GBPI by applying the *stochastic* policies obtained at each iteration of **Algorithm 2** to 100 randomly generated scenario. Take the average HVAC cost as an indicator, the convergence rate of the GBPI in Case II is exhibited in Fig. 12. We observe that a suboptimal *stochastic* policy can be approached within 10 iterations.

As an example, we visualize the optimal *stochastic* policy (mapping from the state space to the action space) at 9:00 a.m. in Fig. 13 (only for the states visited in the sample paths). One may note that for some states, the probability (to take action) is generally distributed among a small groups of actions. However, for some states like $\{171, 296, 297, 321, 322, 696, 697, \dots\}$, the probability distributions scatter among a larger group of actions. This is attributed to their low occurrence frequency in the sample paths while performing the GBPI to learn policy, which is confirmed by investigating the occurrence of those states as shown in Fig. 14. However, this also implies that those states are rare case and deserve less computation.

V. CONCLUSION

This paper studied the optimal *stochastic* control of HVAC system to optimize energy cost while maintaining comfortable thermal condition via Markov decision process (MDP). In particular, the predicted mean vote (PMV) model is involved to directly reflect human thermal sensation in the optimization.

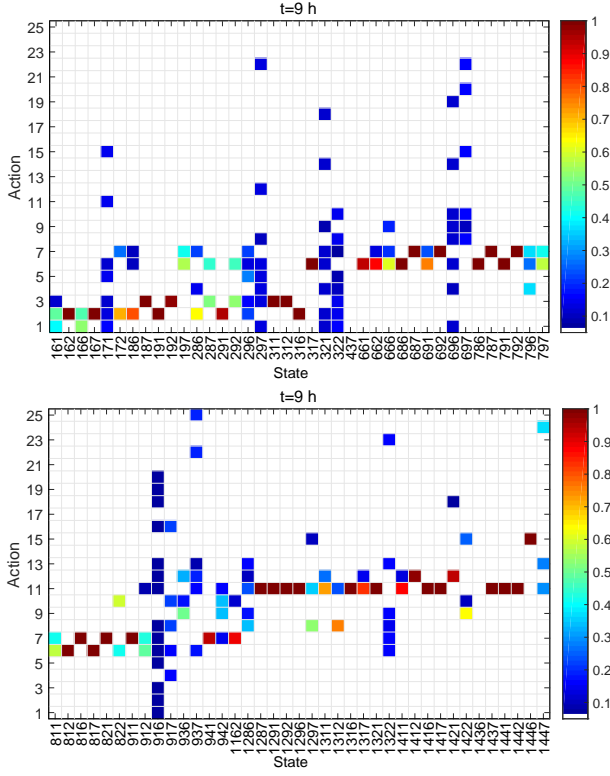


Fig. 13. The random policy at 9:00 a.m. for Case II).




Fig. 14. The state distribution at 9:00 a.m. for Case II.

We proposed a gradient-based policy iteration (GBPI) method that can be used to learn the optimal *stochastic* policies of HVAC control and proved the convergence theoretically. The main advantages of the proposed method are that *i)* it can be implemented *off-line* by learning the policies use historical data thus reducing the high *on-line* computation cost; *ii)* it can handle the nonlinear system dynamics and the non-analytical PMV model efficiently. The performance of the proposed method was demonstrated though comparison with the optimal solution obtained with assumed accurate predictions.

This paper mainly focused on single room case. One potential future direction is to extend the proposed method to multi-zone buildings and develop a distributed learning framework. However, the extension is nontrivial and not straightforward considering the zone interactions. One possible solution is to adopt the “*one-agent-at-a-time*” idea (i.e, sequentially learning the control policies for each individual zones) proposed in [29] to overcome the computational issue.

APPENDIX A PROOF OF THEOREM 2

Proof. According to (13), we have the following performance difference equation over two successive iterations:

$$\begin{aligned}
 & J(\sigma^{k+1}; S_0) - J(\sigma^k; S_0) \\
 &= \sum_{t=0}^{T-1} \pi_t^{\sigma^{k+1}} \left[(r_t^{\sigma^{k+1}} - r_t^{\sigma^k}) + (P_t^{\sigma^{k+1}} - P_t^{\sigma^k}) V_{t+1}^{\sigma^k} \right] \\
 &= \sum_{t=0}^{T-1} \sum_{s_t \in S_t} \pi_t^{\sigma^{k+1}}(s_t) \left[(r_t^{\sigma^{k+1}}(s_t) - r_t^{\sigma^k}(s_t)) \right. \\
 &\quad \left. + \sum_{s_{t+1} \in S_{t+1}} (P_t^{\sigma^{k+1}}(s_{t+1}|s_t) - P_t^{\sigma^k}(s_{t+1}|s_t)) V_{t+1}^{\sigma^k}(s_{t+1}) \right]
 \end{aligned} \tag{23}$$

For brevity, we define an operator as $\Sigma^k(s_t) = \sum_{a_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, a_t)$ and we have

$$\begin{aligned}
 & r_t^{\sigma^{k+1}}(s_t) - r_t^{\sigma^k}(s_t) \\
 &= \sum_{b_t=1}^{|\mathcal{A}_t|} [p_t^{\sigma^{k+1}}(b_t|s_t) - p_t^{\sigma^k}(b_t|s_t)] r_t(s_t, b_t) \\
 &= \sum_{b_t=1}^{|\mathcal{A}_t|} \left[\frac{\sigma_t^{k+1}(s_t, b_t)}{\Sigma^{k+1}(s_t)} - \frac{\sigma_t^k(s_t, b_t)}{\Sigma^k(s_t)} \right] r_t(s_t, b_t)
 \end{aligned} \tag{24}$$

$$\begin{aligned}
 & P_t^{\sigma^{k+1}}(s_{t+1}|s_t) - P_t^{\sigma^k}(s_{t+1}|s_t) \\
 &= \sum_{b_t=1}^{|\mathcal{A}_t|} p_t(s_{t+1}|s_t, b_t) (p_t^{\sigma^{k+1}}(b_t|s_t) - p_t^{\sigma^k}(b_t|s_t)) \\
 &= \sum_{b_t=1}^{|\mathcal{A}_t|} p_t(s_{t+1}|s_t, b_t) \left[\frac{\sigma_t^{k+1}(s_t, b_t)}{\Sigma^{k+1}(s_t)} - \frac{\sigma_t^k(s_t, b_t)}{\Sigma^k(s_t)} \right]
 \end{aligned} \tag{25}$$

By substituting (24) and (25) into (23), we have

$$\begin{aligned}
& J(\sigma^{k+1}; S_0) - J(\sigma^k; S_0) \\
&= \sum_{t=0}^{T-1} \sum_{s_t \in S_t} \left[\sum_{b_t=1}^{|\mathcal{A}_t|} \left[\frac{\sigma_t^{k+1}(s_t, b_t)}{\Sigma^{k+1}(s_t)} \right. \right. \\
&\quad \left. \left. - \frac{\sigma_t^k(s_t, b_t)}{\Sigma^k(s_t)} \right] (r_t(s_t, b_t) + V_t^{\sigma^k}(s_t, b_t)) \right] \\
&= \sum_{t=0}^{T-1} \sum_{s_t \in S_t} \sum_{b_t=1}^{|\mathcal{A}_t|} \left[\frac{\sigma_t^{k+1}(s_t, b_t)}{\Sigma^{k+1}(s_t)} - \frac{\sigma_t^k(s_t, b_t)}{\Sigma^k(s_t)} \right] A_t^{\sigma^k}(s_t, b_t)
\end{aligned} \tag{26}$$

where we have $A_t^{\sigma^k}(s_t, b_t) = r_t(s_t, b_t) + V_t^{\sigma^k}(s_t, b_t)$.

As indicated in (20), we have

$$\begin{aligned}
\sigma^{k+1}(s_t, a_t) &= \sigma_t^k(s_t, a_t) - \gamma^k(s_t, a_t) \Delta \sigma_t^k(s_t, a_t), \\
&\quad \forall a_t \in \mathcal{A}_t.
\end{aligned}$$

where we have

$$\begin{aligned}
\Delta \sigma_t^k(s_t, a_t) &= \frac{\partial J(\sigma^k; S_0)}{\partial \sigma_t^k(s_t, a_t)} \\
&= \pi_t^{\sigma^k}(s_t) \left[\sum_{b_t=1}^{|\mathcal{A}_t|} \frac{\partial p^{\sigma^k}(b_t|s_t)}{\partial \sigma_t^k(s_t, a_t)} (r_t(s_t, b_t) + V_t^{\sigma^k}(s_t, b_t)) \right] \\
&= \pi_t^{\sigma^k}(s_t) \left[\sum_{b_t=1, b_t \neq a_t}^{|\mathcal{A}_t|} \frac{-\sigma_t^k(s_t, b_t)}{[\Sigma^k(s_t)]^2} A_t^{\sigma^k}(s_t, b_t) \right. \\
&\quad \left. + \frac{\Sigma^k(s_t) - \sigma_t^k(s_t, a_t)}{[\Sigma^k(s_t)]^2} A_t^{\sigma^k}(s_t, a_t) \right] \\
&= \frac{\pi_t^{\sigma^k}(s_t)}{[\Sigma^k(s_t)]^2} \left[\Sigma^k(s_t) A_t^{\sigma^k}(s_t, a_t) - \sum_{b_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, b_t) A_t^{\sigma^k}(s_t, b_t) \right]
\end{aligned}$$

Besides, regarding (26), we have

$$\begin{aligned}
& \frac{\sigma_t^{k+1}(s_t, a_t)}{\Sigma^{k+1}(s_t)} - \frac{\sigma_t^k(s_t, a_t)}{\Sigma^k(s_t)} \\
&= \frac{\Delta \sigma_t^k(s_t, a_t) \cdot \Sigma^k(s_t) + \Delta \Sigma^k(s_t) \cdot \sigma_t^k(s_t, a_t)}{\Sigma^k(s_t) \cdot \Sigma^{k+1}(s_t)}
\end{aligned} \tag{27}$$

where $\Delta \Sigma^k(s_t) = \Sigma^{k+1}(s_t) - \Sigma^k(s_t)$.

From (27), we observe that if the stepsizes are selected as

$$\gamma_t^k(s_t, b_t) = \frac{\sigma_t^k(s_t, b_t)}{\Sigma^k(s_t)} \quad (\forall b_t \in \{1, 2, \dots, |\mathcal{A}_t|\}), \text{ we have}$$

$$\begin{aligned}
\Delta \Sigma^k(s_t) &= \sum_{b_t=1}^{|\mathcal{A}_t|} \gamma_t^k(s_t, b_t) \Delta \sigma_t^k(s_t, b_t) \\
&= \frac{\pi_t^{\sigma^k}(s_t)}{[\Sigma^k(s_t)]^2} \left[\sum_{b_t=1}^{|\mathcal{A}_t|} \gamma_t^k(s_t, b_t) \Sigma^k(s_t) A_t^{\sigma^k}(s_t, b_t) \right. \\
&\quad \left. - \sum_{b_t=1}^{|\mathcal{A}_t|} \gamma_t^k(s_t, b_t) \sum_{b_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, b_t) A_t^{\sigma^k}(s_t, b_t) \right] \\
&= 0
\end{aligned} \tag{28}$$

The equation (28) implies $\Sigma^k(s_t) = \Sigma^{k+1}(s_t)$ ($\forall s_t \in S_t$) with the selected stepsize $\gamma_t^k(s_t, b_t) = \frac{\sigma_t^k(s_t, b_t)}{\Sigma^k(s_t)}$ ($\forall s_t \in S_t, b_t \in \mathcal{A}_t$).

By substituting (28) into (27), we have

$$\begin{aligned}
& \frac{\sigma_t^{k+1}(s_t, a_t)}{\Sigma^{k+1}(s_t)} - \frac{\sigma_t^k(s_t, a_t)}{\Sigma^k(s_t)} = \frac{-\gamma_t^k(s_t, a_t) \Delta \sigma_t^k(s_t, a_t)}{\Sigma^k(s_t)} \\
&= \frac{-\pi_t^{\sigma^k}(s_t)}{[\Sigma^k(s_t)]^3} \left[\Sigma^k(s_t) \gamma_t^k(s_t, a_t) A_t^{\sigma^k}(s_t, a_t) \right. \\
&\quad \left. - \gamma_t^k(s_t, a_t) \sum_{b_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, b_t) A_t^{\sigma^k}(s_t, b_t) \right]
\end{aligned} \tag{29}$$

Then by substituting (29) into (26) we have

$$\begin{aligned}
& J(\sigma^{k+1}; S_0) - J(\sigma^k; S_0) \\
&= \sum_{t=0}^{T-1} \sum_{s_t \in S_t} \sum_{a_t=1}^{|\mathcal{A}_t|} \left\{ \frac{-\pi_t^{\sigma^k}(s_t)}{[\Sigma^k(s_t)]^3} \left[\Sigma^k(s_t) \gamma_t^k(s_t, a_t) A_t^{\sigma^k}(s_t, a_t) \right. \right. \\
&\quad \left. \left. - \gamma_t^k(s_t, a_t) \sum_{b_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, b_t) A_t^{\sigma^k}(s_t, b_t) \right] A_t^{\sigma^k}(s_t, a_t) \right\} \\
&= \sum_{t=0}^{T-1} \sum_{s_t \in S_t} \frac{-\pi_t^{\sigma^k}(s_t)}{[\Sigma^k(s_t)]^2} \left[\sum_{a_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, a_t) [A_t^{\sigma^k}(s_t, a_t)]^2 \right. \\
&\quad \left. - \sum_{a_t=1}^{|\mathcal{A}_t|} \frac{\sigma_t^k(s_t, a_t)}{\Sigma^k(s_t)} A_t^{\sigma^k}(s_t, a_t) \sum_{b_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, b_t) A_t^{\sigma^k}(s_t, b_t) \right]
\end{aligned} \tag{30}$$

We can set $\Sigma^0(s_t) = 1$ ($\forall s_t \in S_t$) for the initial σ^0 in the GBPI method, which is trivial. Therefore, we have

$$\begin{aligned}
& J(\sigma^{k+1}; S_0) - J(\sigma^k; S_0) \\
&= \sum_{t=0}^{T-1} \sum_{s_t \in S_t} \sum_{a_t=1}^{|\mathcal{A}_t|} \left\{ \frac{-\pi_t^{\sigma^k}(s_t)}{[\Sigma^k(s_t)]^3} \left[\Sigma^k(s_t) \gamma_t^k(s_t, a_t) A_t^{\sigma^k}(s_t, a_t) \right. \right. \\
&\quad \left. \left. - \gamma_t^k(s_t, a_t) \sum_{b_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, b_t) A_t^{\sigma^k}(s_t, b_t) \right] A_t^{\sigma^k}(s_t, a_t) \right\} \\
&= \sum_{t=0}^{T-1} \sum_{s_t \in S_t} \frac{-\pi_t^{\sigma^k}(s_t)}{[\Sigma^k(s_t)]^2} M(s_t)
\end{aligned} \tag{31}$$

where we have $M(s_t) = \sum_{a_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, a_t) [A_t^{\sigma^k}(s_t, a_t)]^2 - \sum_{a_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, a_t) A_t^{\sigma^k}(s_t, a_t) \sum_{b_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, b_t) A_t^{\sigma^k}(s_t, b_t)$.

Since we have $\sum_{a_t=1}^{|\mathcal{A}_t|} \sigma_t^k(s_t, a_t) = \Sigma^k(s_t) = 1$ ($\forall s_t \in S_t$) and the function $f(x) = x^2$ convex, we conclude $M(s_t) \geq 0$ ($\forall s_t \in S_t$) based on the properties of convex functions. Further, we have

$$J(\sigma^{k+1}; S_0) - J(\sigma^k; S_0) \leq 0 \tag{32}$$

The inequality (32) implies that the performance function $J(\sigma^k; S_0)$ is non-increasing w.r.t. iteration k .

As problem (11) is bounded (i.e., the action space is bounded), we imply that the GBPI method can converge to a local optima $\bar{\sigma}$ with $k \rightarrow \infty$. The remainder illustrates that the method can converge to a global optima.

We assume that at least one global optimal policy σ^* exists for problem (11). It is straightforward that

$$J(\sigma^*; S_0) \leq J(\bar{\sigma}; S_0).$$

We define a parameterized random policy as $\sigma^r = (1-r)\bar{\sigma} + r\sigma^*$ ($\delta \in [0, 1]$). Since $\bar{\sigma}$ is a local optima, for any r sufficiently small, we have

$$J(\sigma^r; S_0) - J(\bar{\sigma}; S_0) \geq 0 \tag{33}$$

According to (26) and $\Sigma(s_t) = 1$ ($\forall s_t \in \mathcal{S}_t$), we have

$$\begin{aligned} J(\sigma^r; S_0) - J(\bar{\sigma}; S_0) \\ &= \sum_{t=0}^{T-1} \sum_{s_t \in \mathcal{S}_t} \left[\sum_{a_t=1}^{|\mathcal{A}_t|} (\sigma_t^r(s_t, a_t) - \bar{\sigma}_t(s_t, a_t)) A_t^{\bar{\sigma}}(s_t, a_t) \right] \\ &= r \sum_{t=0}^{T-1} \sum_{s_t \in \mathcal{S}_t} \left[\sum_{a_t=1}^{|\mathcal{A}_t|} (\sigma_t^*(s_t, a_t) - \bar{\sigma}_t(s_t, a_t)) A_t^{\bar{\sigma}}(s_t, a_t) \right] \geq 0 \end{aligned} \quad (34)$$

On the other hand, since σ^* is global optima, we have

$$\begin{aligned} J(\sigma^*; S_0) - J(\bar{\sigma}; S_0) \\ &= \sum_{t=0}^{T-1} \sum_{s_t \in \mathcal{S}_t} \left[\sum_{a_t=1}^{|\mathcal{A}_t|} (\sigma_t^*(s_t, a_t) - \bar{\sigma}_t(s_t, a_t)) A_t^{\bar{\sigma}}(s_t, a_t) \right] \leq 0 \end{aligned} \quad (35)$$

By contrasting (34) and (35), we conclude $\bar{\sigma} = \sigma^*$, otherwise they contradict with each other. This implies the GBPI method can converge to the global optima of problem (11). \square

REFERENCES

- [1] K. Ku, J. Liaw, M. Tsai, T. Liu, *et al.*, “Automatic control system for thermal comfort based on predicted mean vote and energy saving,” *IEEE Trans. Automation Science and Engineering*, vol. 12, no. 1, pp. 378–383, 2015.
- [2] B. Sun, P. B. Luh, Q.-S. Jia, Z. Jiang, F. Wang, and C. Song, “Building energy management: Integrated control of active and passive heating, cooling, lighting, shading, and ventilation systems,” *IEEE Transactions on automation science and engineering*, vol. 10, no. 3, pp. 588–602, 2013.
- [3] M. Frontczak, S. Schiavon, J. Goins, E. Arens, H. Zhang, and P. Wargocki, “Quantitative relationships between occupant satisfaction and satisfaction aspects of indoor environmental quality and building design,” *Indoor air*, vol. 22, no. 2, pp. 119–131, 2012.
- [4] A. Afram and F. Janabi-Sharifi, “Theory and applications of HVAC control systems—a review of model predictive control (MPC),” *Building and Environment*, vol. 72, pp. 343–355, 2014.
- [5] A. Kelman and F. Borrelli, “Bilinear model predictive control of a HVAC system using sequential quadratic programming,” in *IFAC world congress*, vol. 18, pp. 9869–9874, 2011.
- [6] Z. Xu, G. Hu, C. J. Spanos, and S. Schiavon, “PMV-based event-triggered mechanism for building energy management under uncertainties,” *Energy and Buildings*, vol. 152, pp. 73–85, 2017.
- [7] Z. Xu, Q.-S. Jia, and X. Guan, “Supply demand coordination for building energy saving: Explore the soft comfort,” *IEEE Transactions on Automation Science and Engineering*, vol. 12, no. 2, pp. 656–665, 2015.
- [8] N. Nassif, S. Kajl, and R. Sabourin, “Optimization of HVAC control system strategy using two-objective genetic algorithm,” *HVAC&R Research*, vol. 11, no. 3, pp. 459–486, 2005.
- [9] J. Brooks, S. Goyal, R. Subramany, Y. Lin, T. Middekoop, L. Arpan, L. Carloni, and P. Barooah, “An experimental investigation of occupancy-based energy-efficient control of commercial building indoor climate,” in *53rd IEEE Conference on Decision and Control*, pp. 5680–5685, IEEE, 2014.
- [10] J. H. Yoon, R. Baldick, and A. Novoselac, “Dynamic demand response controller based on real-time retail price for residential buildings,” *IEEE Transactions on Smart Grid*, vol. 5, no. 1, pp. 121–129, 2014.
- [11] Y. Ma, A. Kelman, A. Daly, and F. Borrelli, “Predictive control for energy efficient buildings with thermal storage: Modeling, stimulation, and experiments,” *IEEE Control Systems*, vol. 32, no. 1, pp. 44–64, 2012.
- [12] B. Kouvaritakis and M. Cannon, “Stochastic model predictive control,” *Encyclopedia of Systems and Control*, pp. 1350–1357, 2015.
- [13] F. Oldewurtel, C. N. Jones, A. Parisio, and M. Morari, “Stochastic model predictive control for building climate control,” *IEEE Transactions on Control Systems Technology*, vol. 22, no. 3, pp. 1198–1205, 2013.
- [14] A. Parisio, L. Fabbietti, M. Molinari, D. Varagnolo, and K. H. Johansson, “Control of HVAC systems via scenario-based explicit MPC,” in *53rd IEEE conference on decision and control*, pp. 5201–5207, IEEE, 2014.
- [15] M. Klaučo and M. Kvasnica, “Explicit MPC approach to pmv-based thermal comfort control,” in *53rd IEEE conference on decision and control*, pp. 4856–4861, IEEE, 2014.
- [16] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour, “Policy gradient methods for reinforcement learning with function approximation,” in *Advances in neural information processing systems*, pp. 1057–1063, 2000.
- [17] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [18] M. L. Puterman, “Markov decision processes: Discrete stochastic dynamic programming,” 1994.
- [19] Z. Wu, Q.-S. Jia, and X. Guan, “Optimal control of multiroom HVAC system: An event-based approach,” *IEEE Transactions on Control Systems Technology*, vol. 24, no. 2, pp. 662–669, 2015.
- [20] Q.-S. Jia, J. Wu, Z. Wu, and X. Guan, “Event-based HVAC control—a complexity-based approach,” *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 4, pp. 1909–1919, 2018.
- [21] P. O. Fanger *et al.*, “Thermal comfort. analysis and applications in environmental engineering,” *Thermal comfort. Analysis and applications in environmental engineering.*, 1970.
- [22] N. Radhakrishnan, S. Srinivasan, R. Su, and K. Poolla, “Learning-based hierarchical distributed HVAC scheduling with operational constraints,” *IEEE Transactions on Control Systems Technology*, vol. 26, no. 5, pp. 1892–1900, 2017.
- [23] Z. Wu, Q.-S. Jia, and X. Guan, “Optimal control of multiroom HVAC system: An event-based approach,” *IEEE Transactions on Control Systems Technology*, vol. 24, no. 2, pp. 662–669, 2016.
- [24] R. Jia, R. Dong, S. S. Sastry, and C. J. Sap-

- nos, “Privacy-enhanced architecture for occupancy-based HVAC control,” in *Cyber-Physical Systems (ICCPS), 2017 ACM/IEEE 8th International Conference on*, pp. 177–186, IEEE, 2017.
- [25] W. Shen, G. Newsham, and B. Gunay, “Leveraging existing occupancy-related data for optimal control of commercial office buildings: A review,” *Advanced Engineering Informatics*, vol. 33, pp. 230–242, 2017.
- [26] Q.-S. Jia, J.-X. Shen, Z. Xu, and X. Guan, “Simulation-based policy improvement for energy management in commercial office buildings,” *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 2211–2223, 2012.
- [27] Y. Zhao, *Optimization theories and methods for Markov decision processes in resource scheduling of networked systems*. PhD thesis, Tsinghua University, 2010.
- [28] “Thermal Environmental Conditions for Human Occupancy,” standard, Standing Standard Project Committee (SSPC), Mar. 2017.
- [29] D. Bertsekas, “Multiagent rollout algorithms and reinforcement learning,” *arXiv preprint arXiv:1910.00120*, 2019.