Autonomous Pre-Conditioning and Improved Personalization in Shared Workspaces Through Data-Driven Predictive Control

Syed Ahsan Raza Naqvi, Koushik Kar and Sandipan Mishra
Rensselaer Polytechnic Institute, USA
Email: {naqvis2,kark,mishrs2}@rpi.edu

Abstract

This paper studies the problem of indoor zone temperature control in shared workspaces equipped with heterogeneous heating and cooling sources with the goal of increased energy savings and environment personalization. Shared workspaces typically witness distinct, pre-scheduled intervals when they are occupied or are unoccupied. In this work, we develop indoor climate control strategies for each of these intervals. For the interval when the workspace is unoccupied, we propose multiple time-bound control strategies for pre-conditioning the workspace in preparation for a scheduled activity (Phase I). For the interval when the workspace is occupied, we propose a separate control strategy which enhances the thermal comfort of the occupants by harnessing the spatial differentiation of the thermal environment to satisfy the different temperature preferences of the individuals (Phase II). Utilizing a physical test-bed and data-driven model learning, we show that our proposed pre-conditioning strategies in Phase I are less computationally expensive than conventional model predictive control (MPC). For Phase II, we use a low complexity quadratic program to minimize the thermal discomfort experienced by individuals based on their temperature preferences. The experimental results show that for Phase I, the proposed control policies can save a significant amount of energy and achieve the desired mean temperature in the space fairly accurately. We further note that for Phase II, the control scheme can achieve a significant spatial differentiation in temperature towards satisfying the occupants' thermal preferences.

NOMENCLATURE

List of Abbreviations

AHU Air handling unit

ASC Adaptive switching control

c-MPC Vanilla model predictive control with constant weights

FOC Fixed end point optimal control

HVAC Heating, ventilation and air conditioning

NN Neural network

PCS Personal comfort system

RBB Rule-based baseline

RLSE Recursive least squares estimation

TEPP Two phase pre-conditioning and environment personalization

v-MPC Vanilla model predictive control with variable weights

List of Symbols

α Weighting parameter for v-MPC.

α_0 Starting value of α for v-MPC.

Δ_i Desired temperature at sensor i.

$\hat{\mathbf{W}}$ Column vector, with element $\hat{W}_i$, which captures the effect of the ambient temperature and unmodeled dynamics.

μ Duration of each time instance k in minutes.

Φ_1 Metric for capturing the energy consumption for Phase I.

Φ_t Metric for capturing the total energy consumption for Phases I and II.

$\mathbf{A}$ Column vector, with element A_i, which maps the effect of the temperature readings in the preceding time instance on to those in the next time instance.

$\mathbf{B}$ Matrix, with element $B_{i,j}$, which maps the effect of the HVAC input on to the changes in temperature readings at each sensor.

$\mathbf{d}$ Vector of scalar multipliers, with element d_j, which captures the efficacy of HVAC input j for the same % valve opening.

$\underline{T}$ Nominal temperature for pre-conditioning.

I Total number of temperature sensors considered in this work.

i Index for an individual temperature sensor.

J Total number of heating/cooling inputs considered in this work.

j Index for an individual heating/cooling input.

K Total number of time instances for Phase I.

k Index for a single, discrete time instance.

r Rate of updating weights for v-MPC.

$\mathbf{U}[k]$ Column vector, with element $U_j[k]$, which denotes the input signal to the HVAC input j at time k.

$\mathbf{T}[k]$ Column vector, with element $Y_i[k]$, which denotes the temperature readings at sensor i at time k.

I. Introduction

It has been estimated that nearly 75% of all US electricity is consumed within buildings [1], with heating, ventilation and air conditioning (HVAC) systems accounting for approximately 30% of the total energy consumed by the commercial building sector [2]. Therefore, optimizing indoor thermal management operations in buildings could significantly reduce power demand, thereby curtailing the operation of power stations running on fossil fuels.

Shared workspaces, such as conference rooms, are generally occupied for only a part of the work day, according to a pre-determined schedule. The sporadic nature of occupancy means that implementing a business-as-usual approach for heating or cooling such workspaces is inefficient. Instead, HVAC operations can be scaled down when these spaces are unoccupied, only to be ramped up to achieve a nominal temperature when a scheduled activity is imminent.

The increasing desire for personal comfort and wellness among building occupants means that HVAC systems must strive to attain a greater degree of environment personalization in indoor spaces. Recent studies [3] posit that obtaining the occupants' thermal preferences is a suitable means for achieving personalized comfort in the workplace. In the presence of the occupants' temperature preferences, as stated in [4], mapping these demands to the operation of the HVAC system is challenging because the different zones in the shared space may be thermally correlated, and it may not be possible to satisfy the preferences of all occupants simultaneously.

Indoor climate control strategies in buildings can broadly be categorized into rule-based and optimal control-based techniques. Rule-based strategies employ simple algorithms to achieve temperature control and buildings making them popular among building operators, in spite of them offering limited energy savings [5]. In contrast, optimal control-based strategies, such as model predictive control (MPC), determine optimal control inputs that minimize a given objective objective over a time horizon [6]. Given an adequate thermal dynamics model of the building and a reasonably accurate forecast for the ambient conditions, MPC can cause a building's HVAC to act in an anticipatory fashion to help minimize the system cost over the optimization horizon. However, the performance of MPC in temperature management in a building is contingent upon

the fidelity of the thermal dynamics model. Moreover, for MPC to be adopted in buildings, practitioners may have to acquire additional expertise and on-site computational resources for this purpose. In this work, we address some of these hurdles to the modernization of building energy management systems by developing alternate control strategies that have a significantly lower computational complexity than conventional MPC.

The development of high fidelity thermal dynamic models for buildings is hindered primarily by three factors [7]: the gap between the behavior of the model and the building being considered, the seasonal dependence of the model, and the fact that such models are often building-specific. In this work, we hope to demonstrate to the reader that simple, linear models for characterizing thermal dynamics of indoor spaces can be sufficient to achieve the desired indoor climate control objectives. Furthermore, we present an interplay of simple online estimation and predictive control techniques to ensure that control signals to the HVAC system remain robust to seasonal changes in ambient conditions.

Although, a significant body of past research work has studied the use of MPC in building energy management systems, there is a need to design control frameworks to achieve efficient, time-bound, pre-conditioning of a workspace, in addition to satisfying individual thermal preferences. Therefore, this paper develops control strategies that combine data-driven learning with predictive control to improve efficiency as well as to achieve personalization in indoor spaces.

In this paper, we consider an instrumented test-bed which can be taken to represent a typical shared workspace. The workspace is equipped with *heterogeneous* heating/cooling inputs each with distinct thermal output. On a typical day, shared workspaces are sporadically occupied according to pre-determined schedules. A naïve approach to indoor climate control in this situation would involve the use of a single control approach which is agnostic to the occupancy patterns in the space. In contrast, this work develops separate control strategies, which use a data-driven, thermal dynamics model, for space heating and cooling operations to achieve multiple indoor conditioning objectives applicable to varying occupancy patterns in a shared workspace. Specifically, we consider two separate intervals where a shared workspace remains unoccupied prior to hosting a scheduled event, such as a work meeting. *Phase I* (pre-conditioning phase) considers the first interval when the workspace is unoccupied. Here, we propose efficient, time-bound pre-conditioning control strategies for a shared workspace. We further show that our proposed strategies are less computationally taxing than conventional MPC. *Phase II* (environment personalization phase) of our approach considers the subsequent interval when the scheduled

event is underway and the workspace is occupied. The attendees are assumed to have their own temperature preferences. Here, we propose a separate control strategy which enhances the thermal comfort of the occupants by harnessing the spatial differentiation of the thermal environment to satisfy the different individual temperature preferences (within known bounds). In the context of indoor climate control, our contributions are:

- We identify the need for developing separate, occupancy-based control strategies for HVAC operations in shared workspaces. To this end, we develop an autonomous, robust and low-complexity indoor pre-conditioning mechanism for indoor spaces, called *adaptive switching control (ASC)*, which only ramps up HVAC operations when a scheduled activity is imminent to achieve an adequate, nominal indoor temperature. Once the workspace is occupied, we employ a low complexity, quadratic programming-based approach for satisfying the temperature preferences of the occupants.
- The ASC strategy does not rely on optimization solvers, making it computationally less expensive and easier to scale as compared to MPC-based approaches. Unlike [8], we have also shown that, in addition to using ASC, the time-bound pre-conditioning operation can be solved by casting it as a fixed end point optimal control (FOC) problem.
- We implement the ASC in both simulation and on a physical test-bed and compare its performance to that of FOC and variants of the conventional MPC approach. We have shown that in addition to it being quicker to solve in real-time, ASC and FOC consume less energy as compared to the MPC-based approaches studied herein.
- We show how, for an indoor space without plug load measurement and with heterogeneous heating/cooling inputs, the linear, grey-box model characterizing the thermal dynamics can be manipulated to derive a metric for capturing the power consumption of each of the inputs, which can also be used to formulate the objective functions for the MPC-based strategies used for pre-conditioning the workspace.
- To the best of our knowledge, this work is the first interplay of recursive least squares estimation (RLSE) and predictive control techniques for indoor spaces with heterogeneous HVAC inputs. RLSE requires significantly less computational resources for building model identification as compared to reinforcement learning.
- We show how a low complexity optimal control framework, which can be solved in closed form, can be used to exploit spatial differences in the impact of the heterogeneous

heating/cooling inputs in an indoor space to satisfy disparate thermal preferences (within known bounds) when the workspace is occupied.

This paper is organized as follows: Section II provides a review of relevant literature and an account of our contributions to the state-of-the-art. Section III provides an overview of the data-driven thermal model of the workspace, and formulates the proposed control policies as optimization problems. Section IV presents experimental and simulation results and their analysis. Finally, Section V summarizes the findings of this paper. Further details on the methodology and results presented herein may be found in a detailed technical report available online [9].

II. Literature Review

The building system models considered in the literature can broadly be categorized into three classes: white-box [10], black-box [11] and grey-box. The grey-box model is a hybrid of the white- and black-box model in that it is more physically intuitive than the black-box model and yet is much simpler than the white-box model. Past work, such as [12], has used grey box models to characterize building thermal dynamics. Our work uses a linear, grey-box model for this purpose as well. However, unlike [12], we not only develop a data-driven model to capture indoor thermal dynamics but also deploy predictive control strategies on a physical test-bed.

A significant body of past research work has studied the use of predictive control in tandem with data-driven learning in building energy management systems. The authors in [13] used a physics-based thermal model of an indoor space. Using a standard MPC formulation, the authors then proceeded to minimize the total energy and peak power consumption while keeping the space temperature within prescribed bounds. Subsequently, a stochastic MPC approach was presented to minimize the expected energy costs for temperature regulation within certain bounds. The work in [14] determined the optimal time to turn on the boiler in a building to achieve the desired temperature by a certain given time. The paper used the historical data representing two heating scenarios to develop a neural network (NN) for predictive control to achieve the desired indoor temperatures. The paper relied on offline learning which used the time required to achieve the temperature set-point as the output and the average indoor temperature, the ambient temperature and the water heating system temperature as inputs. Owing to the static nature of the training, the model was not robust to variations in outdoor winter temperatures.

In [15], the authors first developed a mathematical model for the thermal dynamics of a multi-zone building. As this model was unable to capture the effect of the variable air volume

(VAV) box air damper on the zonal temperature, a radial basis function (RBF) NN was used to characterize this relationship. The authors implemented control in a VAV HVAC system by first determining predictions for zonal temperatures at each time-step using RBF-NN. Using the deviation of the predicted temperatures from the set-point, the authors determined the air damper opening for minimizing this deviation. The paper implemented the proposed control approach in simulations. Unlike our work, [15] did not implement its predictive control strategy on a physical test-bed, nor did it consider heterogeneous heating/cooling inputs in individual zones. In [16], the authors used an autoregressive with exogenous inputs (ARX) model for predicting indoor temperature models. The model was trained offline on synthesized data-sets. The authors used an MPC framework in conjunction with the ARX prediction model to implement temperature control in simulations. In [17], the authors considered an indoor space equipped with an under-floor air distribution (UFAD) system for temperature control. The paper used data collected from a real test-bed to train a grey-box model for predicting indoor temperature. The authors compared the performance of an MPC framework with a feedback control approach for indoor temperature management. The authors in [18], as well as those in the references cited in [19], employed deep reinforcement learning to implement HVAC control in buildings. It is noteworthy that papers [13]–[18] primarily relied on simulation studies to evaluate the performance of their proposed data-driven learning control strategies. In contrast, our work performs both simulation and experimental studies. Furthermore, our work is aimed to serve buildings with limited computational power, which would be inadequate in implementing the deep learning strategies that have hitherto been discussed.

Some work in the domain of building climate control has involved the implementation of data-driven learning control approaches in physical spaces. In this context, the authors in [8] used data-driven learning and predictive control to efficiently pre-cool/pre-heat an indoor space based on an occupancy schedule. The authors in [20] used semi-parametric regression to map temperature changes in discrete time to HVAC inputs and employed a learning-based MPC to estimate the occupancy heating load and to adjust the control action accordingly. It is noteworthy that [20] considered only a single type of HVAC input, the air handling unit (AHU), while our work simultaneously utilizes the operation of *multiple heterogeneous* HVAC elements, i.e., AHUs and radiators to achieve the desired objectives. It must be noted that the authors in [8], [20] implemented predictive control using static, time-independent models for characterizing building thermal dynamics, which required retraining to account for seasonal changes in the ambient

conditions. Furthermore, unlike [8], we achieve pre-conditioning of the test-bed using a simpler control approach which has been shown to have significantly lower time complexity than MPC.

In order for indoor thermal management control policies to be robust to changes in ambient conditions, the thermal model for the physical space should be estimated in both offline and online settings. For instance, the authors in [21], [22] performed offline estimation of the parameters of the thermal model of an indoor space using Kalman filtering. Several papers in the domain of online estimation of thermal model parameters have also used reinforcement learning for this purpose [23], [24]. In [23], for instance, the authors used the behavior of the on-site building automation system to be an expert demonstration for reinforcement learning in the context of indoor thermal management. Another technique, the RLSE has also been used in literature for estimating the thermal model due to its low sensitivity to outliers [25]. In this work, we use RLSE for online estimation of the thermal model. Using an online estimator alongside an optimal control framework may require significant computational resources which are generally unavailable in most buildings. Therefore, there is a need to not only perform time-complexity studies of online predictive control strategies but also to reduce the computational complexity of such approaches.

Researchers in the domain of indoor climate control in buildings have previously explored the prospects of achieving temperature personalization in shared spaces by introducing personal comfort systems (PCSs) to satisfy the occupants' temperature preferences, under varying levels of automation. In [26], the authors used leg warmers as PCSs and studied their impact on the occupants' vital signs. It was shown experimentally that these PCSs could help maintain the occupants' key physiological indicators within healthy limits in winters. Similarly, in [27] occupants were provided with wearable face and neck cooling fans as PCSs to achieve their temperature preferences during summers. The work done in both [26] and [27] ascertained the impact on the occupants under fixed, pre-set operational settings of the PCSs. As such, these papers did not automate the process of temperature personalization for the occupants. However, in [28], the authors used machine learning techniques, alongside an experimental setup to determine the occupants' temperature preferences and implemented temperature personalization under varying degrees of automation.

It is noteworthy that the work done in [26]–[28] only focused on operating the PCSs and stopped short of developing strategies for coordinating central HVAC and PCS operations. In contrast, the authors in [29] used physics-based thermal models and the occupancy patterns of

indoor spaces to develop an optimization framework for a PCS-aware central HVAC system. The performance of the proposed strategy was evaluated through simulations only. In [30], the authors showed experimentally how an adjustable thermostat and a space heater may be used to meet the temperature requirements of two occupants in shared workspace. Finally, the authors in [31] used deep reinforcement learning to develop a model-free approach for optimizing energy consumption and achieving temperature personalization in a simulated indoor space which consisted of a central HVAC system along with PCSs at the occupants' work desks.

In contrast to past work in the domain of temperature personalization, our work develops a low-complexity approach for satisfying occupant temperature preferences without incurring the overhead of installing PCSs at each of the occupants' locations. Instead, we have evaluated our personalization approach using the existing heating/cooling inputs available in the test-bed. These inputs are located at fixed positions in the test-bed's walls and ceiling. Therefore, unlike [26]–[31], we have achieved temperature personalization without employing the highly localized PCSs considered in related literature. Moreover, the work in [29]–[31] achieved temperature personalization for workspaces where the occupants were effectively siloed from the impact of the neighboring PCSs. In contrast, our test-bed is an 'open' indoor space where the temperatures experienced by individual occupants are correlated with each other and are dependent on the operation of *all* the heating/cooling inputs in the space. Furthermore, unlike [31] which uses deep reinforcement learning, we use a linear grey-box model for characterizing temperature dynamics to develop a low-complexity personalization strategy which is suitable for buildings with limited computational resources at their disposal.

III. Problem Formulation

We consider a typical shared workspace which is equipped with heterogeneous heating and cooling sources. This indoor space is instrumented with temperature sensors at various locations. The daily activities in such shared spaces typically follow a fixed schedule, such as the timings for work meetings. Therefore, these workspaces may be occupied or unoccupied depending on the time of the day. A simple control strategy for managing indoor temperature in such spaces would involve the use of a single control approach which is agnostic to the occupancy patterns in a shared workspace. Instead, in this section, we propose a *two phase pre-conditioning and environment personalization (TEPP)* approach which develops separate control strategies for space heating and cooling operations to achieve multiple indoor conditioning objectives

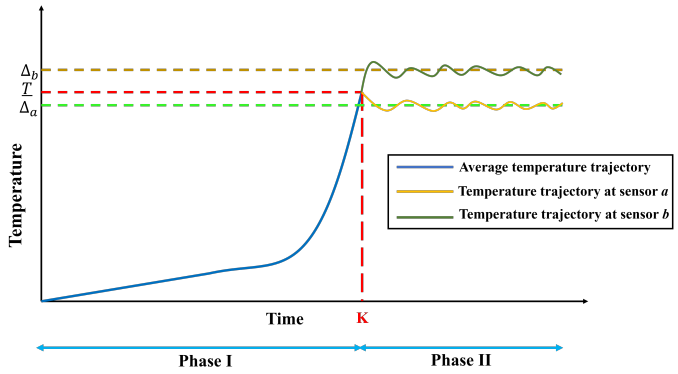

Fig. 1: HVAC operations in Phases I and II.

depending on the status of the occupancy in a shared workspace. Specifically, we consider two separate intervals where a shared workspace remains unoccupied prior to hosting a scheduled event, such as a work meeting. During the first interval, the workspace is unoccupied and our control strategies, as part of Phase I of our indoor climate control approach, aim to efficiently achieve a nominal temperature value by the end of a given deadline, K (see Fig. 1). Phase II of our approach considers the subsequent interval when the scheduled event is underway and the workspace is occupied. The attendees are assumed to have their own, disparate preferences for the temperature in the workspace. We propose a control strategy which enhances the thermal comfort of the occupants by exploiting the spatial differentiation of the thermal environment to satisfy the temperature preferences of multiple attendees (within known bounds). The objective of Phase II is also pictorially represented in Fig. 1.

This work uses a data-driven learning model to characterize the thermal dynamics in the test-bed. We will now present a description of this model and show how it can be made robust to seasonal changes in outdoor air temperatures.

A. Data-Driven Learning Model

Taking the R-C model in [32] as a motivation, we use the grey-box model employed in [8] which uses past temperature readings and input signals to the HVAC system to predict future temperature readings at individual sensors in the test-bed. We discretize the time horizon into K time instances. Each time instance k has duration μ minutes. Furthermore, the formulation assumes that the heating or cooling input remains constant for μ minutes. In this work, we

consider the following model to characterize the temperature dynamics in the shared workspace:

$$\mathbf{T}[k+1] = \mathbf{A}\mathbf{T}[k] + \mathbf{B}\mathbf{U}[k] + \hat{\mathbf{W}}, \tag{1}$$

where $\mathbf{T}[k] \in \mathbb{R}^{I\times 1}$ is the vector of temperature measurements at I sensors at time k, $\mathbf{U}[k] \in \mathbb{R}^{J\times 1}$, with element $U_j[k] \in [0, \bar{U}]$, is a vector of input signals (representing % valve opening in our test-bed) to the J heating and cooling elements of the HVAC system, and column vectors $\mathbf{A}$ and $\hat{\mathbf{W}}$, and matrix $\mathbf{B}$, need to be estimated through learning. The bias, $\hat{\mathbf{W}}$, captures the effect of the ambient temperature and unmodeled dynamics such as heat from human bodies, the server and workstations, solar gains etc. Details of the derivation of this data-driven model may be found in [8].

It is noteworthy that matrix $\mathbf{A}$ and column vector $\hat{\mathbf{W}}$ in (1) are greatly dependent on ambient conditions and have to be re-estimated periodically. Needless to say, repeating the 24-hour training periodically to re-estimate these parameters is both computationally burdensome and operationally infeasible. Therefore, for building HVAC control strategies to be truly robust to changes in the ambient conditions, these parameters need to be re-estimated on-the-fly. In this work, we use RLSE [33], [34] that uses successive temperature readings obtained μ minutes apart, for the online estimation of $\mathbf{A}$ and $\hat{\mathbf{W}}$ for all experimental studies performed in this paper. Matrix $\mathbf{B}$ maps the effect of individual HVAC elements on to the changes in temperature readings. Therefore, it can be taken to be independent of the ambient conditions. Hence, in all the control approaches described subsequently in this paper, we use the estimate of $\mathbf{B}$ as obtained from a single, 24-hour training run.

We will now provide a brief overview of the RLSE approach used in this work for estimating the thermal parameters of an indoor space.

B. RLSE

RLSE is a variant of linear regression which estimates system model parameters by using a forgetting factor which allocate greater weight to the more recently collected data values. This approach aims to minimize the following cost function [35]:

$$V(\hat{\theta}, m') = \frac{1}{2}\sum_{m=1}^{m'} \lambda^{m'-m}(y(m) - \phi^T(m)\hat{\theta}(m'))^2, \tag{2}$$

where m' is the total number of collected samples, λ is the forgetting factor, $\hat{\theta}(m')$ is the estimate of the ground truth at m', $\phi(\cdot)$ is the matrix of coefficients for the linear model, and $y(m)$ is

the ground truth. In our work, $\phi(\cdot)$ is taken to be a concatenation of vectors $\mathbf{A}$ and $\hat{\mathbf{W}}$, and $y(\cdot)$ represents the instantaneous temperature readings collected from the sensors in the test-bed. In the context of online estimation of thermal parameters of an indoor space, RLSE can be helpful in estimating, with low time complexity, the relationship between successive temperature readings as well as the dependency of instantaneous temperature on unmodeled heat gains for applications such as indoor space pre-conditioning, where there is limited HVAC operation expected in the near-term. As will be seen subsequently in this paper, the experiments conducted in this work showed that estimates obtained from RLSE using only a limited number of samples resulted in the objectives of the pre-conditioning control strategies being met. This low complexity re-estimation approach has helped make our proposed approach impervious to changes in ambient conditions on a diurnal and seasonal basis. However, the estimation of matrix $\mathbf{B}$ (which by its nature is minimally dependent on ambient conditions) requires the perturbation of multiple HVAC inputs. Such perturbation is energy inefficient during routine building operations, especially for space pre-conditioning applications. Therefore, the RLSE may not be the most adequate approach for estimating $\mathbf{B}$. Instead, this work obtained a sufficiently rich data-set after perturbing the test-bed with different combinations of HVAC inputs over the course of 24 hours to estimated matrix $\mathbf{B}$ using standard linear regression techniques.

Given the thermal model of the test-bed in (1), we now proceed to formulating the HVAC control strategies for the TEPP approach.

C. TEPP

This subsection presents the control strategies for each of the two phases of the proposed TEPP approach. We begin by laying out the control strategies for Phase I.

1) Phase I – Time-Bound Pre-Conditioning of the Test-Bed: A previously unoccupied indoor space might need to be pre-cooled or pre-heated to a nominal temperature, $\underline{T}$, in preparation for some scheduled activity, e.g. a work-related meeting. In this case, prematurely heating or cooling the space to $\underline{T}$ would be inefficient. On the other hand, waiting to ramp up the heating or cooling operations until the start of the activity would cause significant thermal discomfort to the occupants. Therefore, in order to save energy, it is vital that HVAC operations are controlled such that the average temperature in the space reaches close to $\underline{T}$ immediately before the scheduled

activity. This condition may be expressed as,

$$\frac{\sum_{i=1}^{I} T_i[K]}{I} = \underline{T}, \tag{3}$$

where K represents the time at which the scheduled activity begins and $T_i[K]$ is the temperature measurement at sensor i at that time.

The purpose of the time-bound pre-conditioning of the test-bed is to achieve a nominal temperature in the test-bed by the time a scheduled activity starts with minimum power consumption. Developing control strategies for such operations requires knowledge of the instantaneous power consumption of the HVAC inputs, which we denote by $f(\mathbf{U}[k])$. Therefore, in the absence of plug-in meters measuring the exact power consumption of each HVAC input in the test-bed, $f(\mathbf{U}[k])$ would need to be estimated as described next.

From (1), we notice that the knowledge of the power consumption of the HVAC elements relative to each other is embedded as scalar multipliers in $\mathbf{B}$ which we hope to extract and use to determine $\hat{U}_j$, the maximum allowable valve opening for HVAC element j, normalized to lie between 0% and 100%. The greater the value for $\hat{U}_j$, the greater the power rating of element j and, hence, the greater the change in zone temperature per unit change in valve positions.

In order to capture the relative power consumption of each HVAC element, we consider the product,

$$\mathbf{BU}[k] = \sum_{i=1}^{I} \sum_{j=1}^{J} B_{i,j} \cdot U_j[k], \tag{4}$$

where $\mathbf{B} = \{B_{i,j}\}, \forall i \in \mathbb{Z} \cap \{1, \cdots, I\}, j \in \mathbb{Z} \cap \{1, \cdots, J\}, \mathbf{B} \in \mathbb{R}_{I \times J}$. We can further express this product as,

$$\mathbf{BU}[k] = \tilde{\mathbf{B}}\tilde{\mathbf{U}}[k], \tilde{\mathbf{B}} = \left\{ \frac{B_{i,j}}{d_j} \right\} = \{\tilde{B}_{i,j}\}, \tilde{\mathbf{U}}[k] = \{U_j[k] \cdot d_j\} = \{\tilde{U}_j\}, \tag{5}$$

where d_j is a constant. Note that element $B_{i,j}$ denotes the change in temperature observed at sensor i brought about by a unit change of the valve opening of HVAC input j. We take $\mathbf{d}$ to be a vector such that $\mathbf{d} \in \mathbb{R}^{J \times 1}$, with element d_j such that $\frac{B_{i,j}}{d_j} \in [0, 1]$. Here, $d_{j_1} > d_{j_2}$ implies that for the same % valve opening, j_1 can effect a greater change in temperature in the workspace than j_2. Therefore, in the absence of plug-in meters measuring the exact power consumption of each HVAC input, $U_j[k] \cdot d_j = \tilde{U}_j[k]$ can provide a representative value for power consumption.

Furthermore, to ensure that $\tilde{U}_j[k]$ lies within the prescribed operational range, we will consider,

$$\hat{\mathbf{U}}[k] = \left\{ \frac{\tilde{U}_j[k]}{\max(\mathbf{d})} \right\}, \hat{U}_j[k] \in \left[0, \bar{U} \cdot \frac{d_j}{\max(\mathbf{d})} \right], \forall j, \tag{6}$$

where $\hat{\mathbf{U}}[k]$ is taken to be the vector of controllable inputs for our proposed control strategies which will be presented later in this section. Therefore, the temperature dynamics in the test-bed can now be given by,

$$\mathbf{T}[k+1] = \mathbf{A}\mathbf{T}[k] + \tilde{\mathbf{B}}\hat{\mathbf{U}}[k] + \hat{\mathbf{W}}. \tag{7}$$

Finally, the efficient, time-bound, pre-conditioning operation in Phase I can be expressed in the form of the following optimization problem:

$$\min_{\hat{\mathbf{U}}[.]} . \sum_{k=1}^{K-1} \sum_{j=1}^{J} \hat{U}_j[k], \text{ s.t. } (3), (7). \quad \textbf{(P1)}$$

In this work, we propose two approaches for achieving energy efficient pre-conditioning of a shared workspace: the ASC and FOC. We will now provide an exposition of these two pre-conditioning control strategies considered in this work.

a) TEPP-ASC: We aim to develop an approach for time-bound pre-conditioning of the test-bed which is not only adaptive but also computationally less burdensome. We refer to our first approach as ASC. It utilizes the underlying properties of the linear framework in (**P1**) for this pre-conditioning operation to determine the time instance at which the cooling/heating inputs need to be engaged to achieve a nominal temperature by the end of a given deadline.

The framework (**P1**) is a linear programming problem. Hence, its solution, $\hat{\mathbf{U}}^*[k] \ \forall k$, with elements, $\hat{U}_j^*[k] \ \forall j, k$, lies at the vertices of the feasible region [36]. This implies that in the optimal solution, the valves for each HVAC element can either be completely closed or fully open (or to the extent permissible). For a sufficiently small value of μ, $\hat{U}_j^*[k] \in \{0, \bar{U} \cdot \frac{d_j}{\max(\mathbf{d})}\}$. This observation reduces the pre-conditioning problem to one that determines

- whether the test-bed needs to be cooled or heated by the end of time period K, and
- the instance when the valve openings for the HVAC element are set to the maximum value.

Estimates for $\mathbf{A}$ and $\hat{\mathbf{W}}$ along with (7) may be used to determine the trajectory of the temperature readings up till time K when none of the HVAC inputs are switched on. Depending on whether the final value in this trajectory is greater/less than $\underline{T}$, the ASC determines if cooling/heating operations will be required to pre-condition the test-bed. The trajectory of the

mean temperature in the test-bed when the heating/cooling inputs are at their maximum values for all k may be obtained by back-solving (7) for a desired $\underline{T}$. The point of intersection of the two trajectories, $\underline{k}$, is an estimate of the time instance when the appropriate heating/cooling inputs need to be engaged to approximately achieve (3). These trajectories are re-calculated when $\mathbf{A}$ and $\hat{\mathbf{W}}$ are re-estimated every μ minutes. Once the instance in time represented by the intersection of the two trajectories is reached, the relevant HVAC elements are switched on. If the mean temperature over-shoots (in the case of pre-heating) or under-shoots (in the case of pre-cooling) from $\underline{T}$, the trajectories of the mean temperature with and without the intervention of an HVAC element are obtained again using the most recent estimates of $\mathbf{A}$ and $\hat{\mathbf{W}}$ to allow the system to correct course and bring $\frac{\sum_{i=1}^{I} T_i[K]}{I}$ closer to $\underline{T}$. This approach is summarized in Algorithm 1.

Algorithm 1 TEPP-ASC

Data: Temperature readings from I temperature sensors and estimate for matrix $\mathbf{B}$

Result: Implementing TEPP-ASC for time-bound pre-conditioning

while $k < K$ **do**

  Determine mean instantaneous temperature, $\bar{T}[k]$

  if $k < \underline{k}$ ***and*** $\bar{T}[k] \neq \underline{T}$ **then**

    Estimate $\mathbf{A}$ and $\hat{\mathbf{W}}$ using past temperature readings through RLSE

    Taking $\hat{\mathbf{U}} = \mathbf{0}$, use (7) to predict mean temperature trajectory under no HVAC input

    Set $\tilde{T}$ to be the no-input terminal temperature

    if $\tilde{T} < \underline{T}$ **then**

      Set heating inputs to $\bar{U} \cdot \frac{d_j}{\max(\mathbf{d})}$

    else

      Set cooling input to $\bar{U} \cdot \frac{d_j}{\max(\mathbf{d})}$

    end

    Back-solve (7) to predict mean temperature trajectory when HVAC is engaged

    Determine the intersection between the temperature trajectories to update $\underline{k}$

  else

    Engage HVAC inputs in the test-bed to begin pre-conditioning operation

  end

end

b) TEPP-FOC: This approach achieves time-bound pre-conditioning by solving the linear, fixed end-point problem (**P1**) using a mathematical solver. Given a deadline K, an FOC-based approach [37] is used to determine the signals to the HVAC inputs. After every μ minutes, this control strategy uses the sensor readings and the instantaneous HVAC inputs to solve (**P1**). However, at each instance k, the problem is solved over a shrinking horizon of $(K-k)$. This approach is summarized in Algorithm 2.

Algorithm 2 TEPP-FOC

Data: Temperature readings from I temperature sensors, estimates for **A**, **B** and $\hat{\mathbf{W}}$

Result: Implementing TEPP-FOC for time-bound pre-conditioning

while $k < K$ **do**

Solve (**P1**) over $K-k$ time instances

Using the solution to (**P1**) and (7), obtain predicted temperature readings at each sensor

Update control horizon

end

2) Phase II – Satisfying Individual Temperature Requirements: Next, we develop a protocol for the indoor environment's HVAC operations for the case when it is already occupied and the workers are seated at known locations. Individuals at each of these positions may have different temperature preferences. We present a control mechanism that aims to enhance wellness for all occupants by minimizing their thermal discomfort.

We can minimize the individuals' thermal discomfort by optimizing the following objective:

$$\min_{\mathbf{U}[.]} \sum_{k=1}^{K'} \sum_{i=1}^{I} (T_i[k] - \Delta_i)^2, \ \forall i \in \{1 \dots I\}, \ \text{ s.t. } (1), \ \textbf{(P2)}$$

where K' represents the size of the optimization window. Δ_i is the desired temperature at sensor i's location, which is determined based on the occupants' preferences and their locations relative to the sensor. Observing (7), it may be noted that (**P2**) is a low complexity quadratic program which can be solved explicitly in closed form. Owing to the limited number of sensors in the test-bed, (**P2**) assumes that the sensor which is located closest to an occupant accurately gives the temperature experienced by that occupant. Unlike the policy for pre-conditioning, we aim to minimize the occupants' thermal discomfort as quickly as possible. Ideally, our control policy for meeting individual thermal requirements should also prevent significant overshoots or

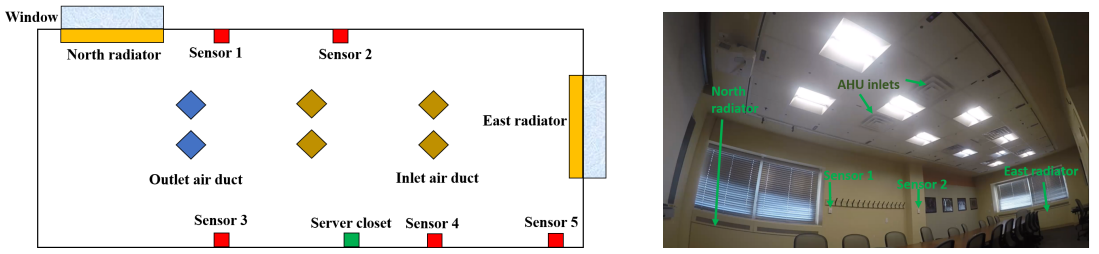

(a) Test-bed layout.

(b) Photograph of the on-campus test-bed.

Fig. 2: The physical test-bed used for experimental evaluation.

undershoots about Δ_i. Therefore, at each time instance k, (**P2**) is solved by updating the starting temperature and heating/cooling input values every $\bar{\mu}$ minutes where $\bar{\mu} < \mu$.

IV. Performance Evaluation

A. Test-Bed Layout

Fig. 2a shows the layout of the test-bed used for our experiments. The room is equipped with three controllable heating sources and one cooling source. The heating sources include two radiators, that are attached to the test-bed's walls, and an AHU in the ceiling. The cooling operation takes place through a separate AHU in the ceiling. Furthermore, the test-bed is instrumented with five wall-mounted temperature sensors located at various locations. The test-bed was equipped with Raspberry Pi-based panels where occupants could declare their temperature preferences, with the test-bed tracking the location of the occupants through the use of time-of-flight sensors.

B. Test Scenarios

The experimental evaluation of the dynamical model in (7) and our proposed HVAC control strategies considers two contiguous time intervals. The first interval, Phase I (pre-conditioning phase), takes the shared workspace to be unoccupied and requires time-bound pre-conditioning of the shared space in anticipation of a scheduled work activity. The second interval, Phase II (environment personalization phase), considers the situation when the workspace is occupied and where occupants at two locations in the shared space have two different temperature preferences. In view of the test-bed having a significant thermal inertia, μ was chosen to be five minutes for training the model in (7). The duration of Phase I was taken to be 60 minutes, i.e., $K = 12$.

C. Performance Benchmarks

This work considers three baseline approaches. The first two use a standard MPC framework with some slight differences. We term the first of these baselines as vanilla MPC with constant weights (c-MPC), and the second as vanilla MPC with variable weights (v-MPC). Unlike the proposed TEPP approach, these vanilla MPC strategies can be used for both pre-conditioning as well as for HVAC operations during a scheduled activity with some minor modifications. The third baseline considered herein is the rule-based baseline (RBB). We will now provide brief descriptions of these baseline approaches.

1) c-MPC: This approach minimizes the weighted sum of the total energy consumed for pre-conditioning the test-bed and the deviation of the temperature from the desired value at individual sensors, expressed as,

$$\min_{\hat{\mathbf{U}}[.]} \sum_{k=1}^{2K-1} \Big(\alpha \sum_{i=1}^{I} (T_i[k] - \underline{T})^2 + (1-\alpha) \sum_{i=1}^{I} \sum_{j=1}^{J} \hat{U}_j[k] \Big) \quad \text{s.t. } (7), \quad (\mathbf{P3})$$

where $0 \leq \alpha \leq 1$. Unlike FOC, the vanilla MPC approach incorporates the thermal discomfort term in the objective, causing the α parameter to be re-tuned for different values of $\underline{T}$. Furthermore, vanilla MPC has a prediction window of a *fixed* length, K, unlike FOC.

2) v-MPC: This approach uses the same objective as in (**P3**). However, it considers α to be a function of time for pre-conditioning operations. Specifically, we take $\alpha[k+1] = \min\Big(\alpha[k] + r\Big(\frac{1-\alpha_0}{K}\Big), 1\Big)$, where α_0 is the value of α at the beginning of the pre-conditioning operation and the constant r determines how quickly α approaches 1. Taking α to be linearly varying (unlike the case for c-MPC) can help ensure that the heating/cooling inputs in the test-bed are engaged closer to the pre-conditioning deadline, thereby saving energy.

3) RBB: This approach presents a naïve pre-conditioning strategy where the valve openings of the required heating/cooling inputs are set to the maximum possible value x minutes before the scheduled activity commences, i.e., at $(K\mu - x)$ minutes. Prior to this instance, these valves are completely shut. Upon achieving $\underline{T}$, this baseline approach maintains the mean temperature close to this level using hysteretic control. The RBB approach can result in non-trivial deviations from $\underline{T}$ and is highly dependent on the starting temperatures in the indoor space.

D. Evaluation Results - Phase I

This subsection presents the results for pre-conditioning the test-bed using the ASC, FOC and the vanilla MPC approaches. Each approach uses the same matrix **B** that was estimated offline

by performing linear regression on a data-set obtained by randomly perturbing each of the four HVAC inputs over the course of a 24-hour period. The matrix $\mathbf{A}$ and the column vector $\hat{\mathbf{W}}$ were continually re-estimated by RLSE using the temperature readings up till that particular instance. The re-estimations were aimed at determining the thermal characteristics of the test-bed on a particular day without any influence from the heating/cooling elements in the room. As given by the expression for $\hat{U}_j[k]$ in Section III, the maximum values of the control signals to the AHU heat, AHU cool, North radiator and East radiator are given by the vector, $[100, 73.5, 17.9, 40.8]$.

We will now present the dynamics of the mean temperature readings in the test-bed for different values of $\underline{T}$ for our proposed control strategies implemented over a 60-minute horizon for Phase I. For each value of $\underline{T}$, we include plots for (i) the simulated trajectory of the mean temperature in the test-bed, and (ii) the change in the control signals to the heating/cooling elements over the 60-minute period. Additionally, we also include the experimental results for the trajectory of the temperature and the change in the control signals for the ASC approach. The simulations results presented herein have been obtained from MATLAB routines that implement ASC, FOC, c-MPC and v-MPC using the linear state space model in (7). The matrix $\mathbf{B}$ used in the simulations is the same as the one used for obtaining the experimental results. We also include tabulated comparisons of the energy consumption for the pre-conditioning strategies studied here as well as the time taken to execute each of them for a single time step.

1) Using TEPP-ASC: Fig. 3a shows the evolution of the mean temperature in the test-bed obtained from simulation and experimental results for ASC when $\underline{T} = 23°$C. Fig. 3b shows the control signals to the HVAC elements to achieve this $\underline{T}$ during experimental runs on the test-bed. Finally, Fig. 3c shows the status of the HVAC elements as recorded in simulations.

It may be seen that for the value of $\underline{T}$ considered here under experimental settings, ASC successfully achieved the desired mean temperature in the test-bed by the end of the 60-minute deadline to within 0.2°C. Furthermore, Fig. 3 show that the proposed ASC strategy achieved the desired temperature for $\underline{T} = 23°$C by the deadline without any course correction, i.e., the relevant HVAC inputs, once engaged, did not need to be switched off before the completion of the hour. This result demonstrates that the discrete, linear model in (7) is able to satisfactorily capture the relationship between the HVAC operations and the change in temperature.

2) Using TEPP-FOC: Here we use the FOC approach to achieve a desired average temperature while minimizing power consumption. Fig. 4a shows the simulated results for pre-conditioning the test-bed for all values of $\underline{T}$. These results have been obtained using the matrix

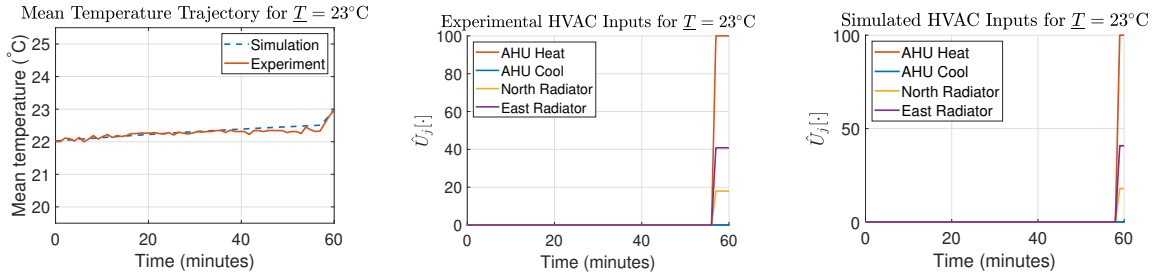

(a) Simulation and experimental results for trajectories of mean temperature in the test-bed for $\underline{T} = 23°C$ using ASC.

(b) Experimental results for heating/cooling inputs for $\underline{T} = 23°C$ using ASC.

(c) Simulation results for heating/cooling inputs for $\underline{T} = 23°C$ using ASC.

Fig. 3: Trajectory of mean temperature and heating/cooling input dynamics for ASC when $\underline{T} = 23°C$.

A and column vector $\hat{\mathbf{W}}$ estimated during the experimental runs of the ASC strategy for the corresponding values of $\underline{T}$. We will now compare the simulated HVAC operation for ASC (Fig. 3c) with that of FOC. As evident from the trajectories of the mean temperatures in Fig. 4a (e.g. beyond the 55^{th} minute for $\underline{T} = 23°C$, see Fig. 4d), it may be seen that the heating/cooling elements are engaged at approximately similar instances in time for both ASC and FOC. As the FOC solves a convex optimization problem, therefore, the ASC strategy can be conservatively stated to be a near-optimal power minimization strategy for Phase I.

3) Using c-MPC: This approach solves the optimization framework for vanilla MPC as presented in Section III. Specifically, it uses a fixed value of $\alpha = 0.99999$. Consequently, as seen in Fig. 4b, for each value of $\underline{T}$, this approach achieves the desired temperature much before the 60-minute deadline and subsequently maintains the average temperature close to this value. Fig. 4e shows the heating/cooling inputs for this strategy when $\underline{T} = 23°C$. Unlike the plots in Figs. 3c and 4d, it may be seen here that c-MPC engages multiple HVAC inputs much earlier than is necessary to achieve $\underline{T}$. Consequently, c-MPC can be significantly inefficient as compared to ASC and FOC unless the weights in the optimization framework are suitably tuned for each value of $\underline{T}$. The ASC and FOC approaches do not require such tuning for their use in Phase I.

4) Using v-MPC: Here, we take $\alpha_0 = 0.99, r = 2$. The simulated trajectories of mean temperatures in the test-bed for v-MPC are presented in Fig. 4c. The heating/cooling inputs employed by this strategy when $\underline{T} = 23°C$ are given in Fig. 4f. It may be seen from Figs. 4c and 4f that this approach restricts the use of heating/cooling inputs till the instance when

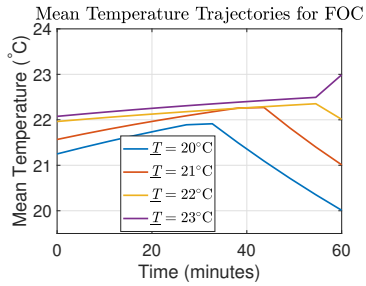

(a) Simulation results for trajectories of mean temperature in the test-bed for $\underline{T} \in \{20, 21, 22, 23\}^{\circ}$C using FOC.

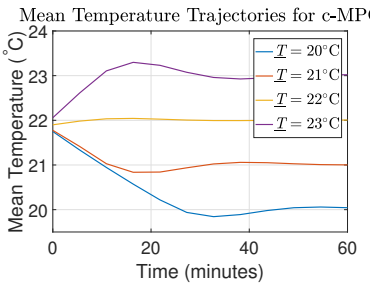

(b) Simulation results for trajectories of mean temperature in the test-bed for $\underline{T} \in \{20, 21, 22, 23\}^{\circ}$C using c-MPC.

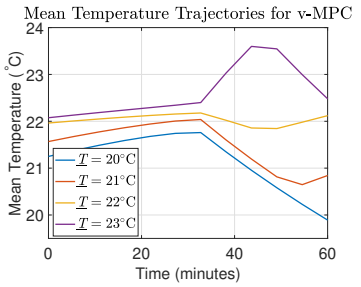

(c) Simulation results for trajectories of mean temperature in the test-bed for $\underline{T} \in \{20, 21, 22, 23\}^{\circ}$C using v-MPC.

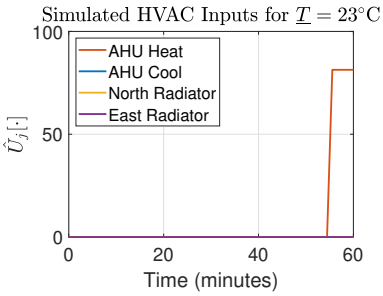

(d) Simulation results for heating/cooling inputs for $\underline{T} = 23^{\circ}$C using FOC.

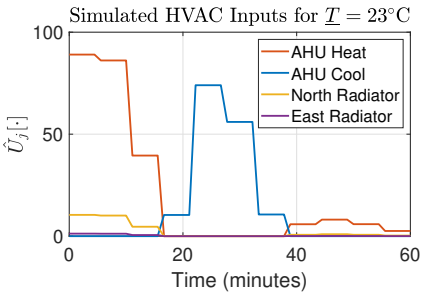

(e) Simulation results for heating/cooling inputs for $\underline{T} = 23^{\circ}$C using c-MPC.

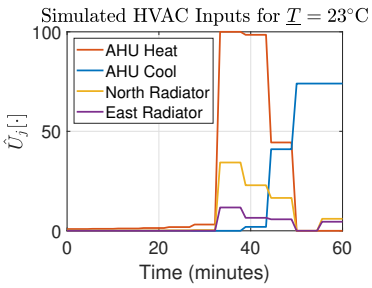

(f) Simulation results for heating/cooling inputs for $\underline{T} = 23^{\circ}$C using v-MPC.

Fig. 4: Simulation results for FOC, c-MPC and v-MPC.

$\alpha \approx 1$, which, in this series of simulations, occurs just after the 30-minute mark. Beyond this point in time, $\hat{U}_j[\cdot]$ takes greater values than before to reduce the squared difference between the instantaneous average temperature in the test-bed and $\underline{T}$.

Table I records the values of $\Phi_1 = \sum_{j=1}^{J} \sum_{k=1}^{K} \hat{U}_j[k]$ for all control strategies studied in this paper. In the absence of explicit power measurements for each HVAC input in the test-bed, we take Φ_1 to be an alternate metric for capturing the energy consumption of each of the control strategies for Phase I. All values in Table I have been obtained from simulations using identical $\mathbf{A}$, $\mathbf{B}$ and $\hat{\mathbf{W}}$ for each control strategy to ensure a fair comparison. The table shows that c-MPC and v-MPC consumed significantly greater energy than ASC and FOC did. This observation can be attributed to the framework used for vanilla MPC. Here, α needs to be re-tuned for each $\underline{T}$, to achieve a mean temperature close to $\underline{T}$ only at the end of the 60-minute deadline. No such

TABLE I: Values of Φ_1 for Phase I.

Control Strategy	**20**°C	**21**°C	**22**°C	**23**°C
ASC	1323	367.4	158.8	635.0
FOC	2078	736.8	287.6	976.1
c-MPC	5849	5018	3262	3173
v-MPC	3155	3092	3117	3563

TABLE II: Execution time of a single iteration.

Control Strategy	**Execution Time (s)**
ASC	0.023
FOC	1.67
c-MPC	8.59
v-MPC	8.61
RBB	0.030

TABLE III: Values of Φ_t with $\underline{T} = 23$°C.

Control Strategy for Phase I	Φ_t
ASC	2040
FOC	2598
c-MPC	4467
v-MPC	4882

re-tuning is required for ASC and FOC, thereby simplifying the pre-conditioning operations for the building operator.

It may be observed from Table I that although ASC and FOC solve the same optimization problem, the Φ_1 values for ASC are consistently lower than those for FOC. This is because the solution to the optimization problem for FOC is determined at intervals of length μ minutes. This means that the control signals to each heating/cooling input remain constant for μ minutes. In contrast, ASC determines the switching instance (which may not be a multiple of μ) for the HVAC inputs by determining the point of intersection of the trajectories of the mean temperature in the test-bed when all HVAC inputs are switched off and when all relevant HVAC inputs are engaged, to solve the pre-conditioning problem in Phase I.

Table II shows the time taken to run the four control strategies in a single iteration. It may be seen that ASC offers superior run-time performance as compared to FOC and the vanilla MPC approaches. This is because ASC uses an analytical approach for solving the pre-conditioning problem for Phase I, unlike the other control strategies presented previously, which use mathematical solvers (which are more computationally intensive) to obtain their respective solutions. The execution times for the vanilla MPC strategies are significantly higher than those for ASC and FOC. This is because the prediction horizon for the vanilla MPC strategies is K instances long, whereas this duration is $(K - k)$ instances long for FOC.

Table IV records the values of Φ_1 and the deviation of the mean temperature in the indoor space from $\underline{T}$ at the end of Phase I for the RBB approach when $x = 5$. Table V records these values for the case when $x = 10$. It may be seen that the energy consumed by ASC and FOC (see Table I) is significantly lower for most values of $\underline{T}$ as compared to those for RBB. Furthermore,

TABLE IV: Values of Φ_1 and deviation from $\underline{T}$ of the mean temperature at the end of the 60 minutes for the RBB when $x = 5$.

	20°C	**21°C**	**22°C**	**23°C**
Φ_1	735	735	1587	1587
Deviation from $\underline{T}$	1.3549	0.2469	0.5139	-0.1171

while ASC and FOC were seen to achieve $\underline{T}$ at the end of Phase I, RBB exhibited significant deviations from this nominal temperature as it necessarily engaged the heating/cooling inputs x minutes before the pre-conditioning deadline, following which it implemented hysteretic control to maintain the mean temperature in the indoor space within a specified range.

It may further be seen from Table IV, that the deviation of the mean temperature from $\underline{T}$ for $x = 5$ under cooling operation (i.e., when $\underline{T} = 20°$C and $\underline{T} = 21°$C) was greater than the corresponding values for RBB with $x = 10$. This is because the indoor space being studied was equipped with a single cooling input. Therefore, 5 minutes were insufficient to pre-cool the space to the desired temperatures. In contrast, RBB with $x = 5$ performed better than RBB with $x = 10$, both in terms of energy consumption and temperature deviation, when the indoor space needed to be heated up to $\underline{T} = 22°$C and $\underline{T} = 23°$C. This is due to the fact that RBB engages all three heating inputs in the indoor space. Operating these inputs 10 minutes prior to the scheduled activity is not only energy inefficient, but also a cause of significant overshoots from $\underline{T}$ (see Table V). On the other hand, using RBB with $x = 5$ for heating operations resulted in significantly lower deviations from $\underline{T}$. Therefore, it is noteworthy that the parameter x in the RBB approach needs to be continually tuned depending on the pre-conditioning requirements of the indoor space. It follows that the RBB's performance is greatly dependent on the temperatures in the indoor space at time $k = 1$ as well as on the building thermal dynamics.

It may be observed that the achievable comfort metrics under heating/cooling operations using the RBB approach, as seen in Tables IV and V exhibit a great deal of variability. This indicates that strategies using data-driven learning and optimal control, such as those proposed in this study, are better equipped than rule-based approaches to achieve robust, time-bound pre-conditioning in indoor spaces. Therefore, it makes sense to incorporate data-driven learning to implement pre-conditioning in shared workspaces.

TABLE V: Values of Φ_1 and deviation from $\underline{T}$ of the mean temperature at the end of the 60 minutes for the RBB when $x = 10$.

	20°C	**21°C**	**22°C**	**23°C**
Φ_1	1102.5	1102.5	1954.5	2380.5
Deviation from $\underline{T}$	1.0008	-0.0945	1.3032	0.6537

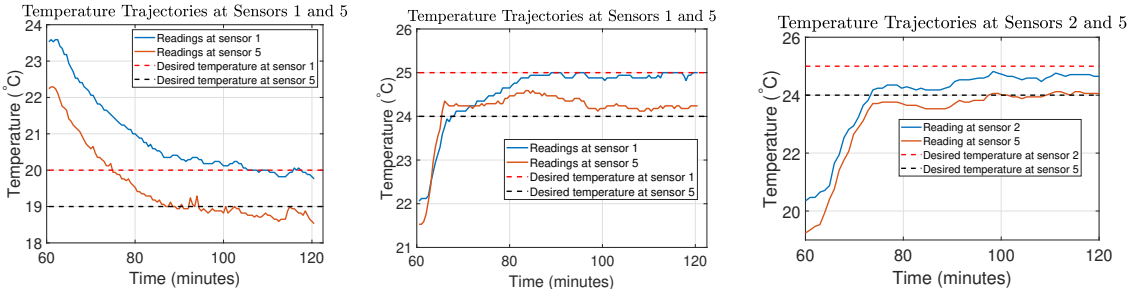

(a) Temperature differentiation in the shared space when $\Delta_1 = 20°$C and $\Delta_5 = 19°$C.

(b) Temperature differentiation in the shared space when $\Delta_1 = 25°$C and $\Delta_5 = 24°$C.

(c) Temperature differentiation in the shared space when $\Delta_2 = 25°$C and $\Delta_5 = 24°$C.

Fig. 5: Experimental results for Phase II.

E. Evaluation Results - Phase II

Next, we consider Phase II or the environment personalization phase. This phase considers the situation where the test-bed's HVAC has achieved the desired $\underline{T}$ by the beginning of this meeting. Here, we consider that the individuals requesting different temperature set-points are seated close to sensor 5 and either sensor 1 or sensor 2 in the test-bed (see Fig. 2a). We assume that the temperatures recorded by each of these sensors are the temperatures experienced by the nearest occupants. Figs. 5a, 5b and 5c depict how the HVAC operations determined using (**P2**) can successfully achieve a spatial temperature differentiation of up to 1°C for heating and cooling operations at different locations within the test-bed, independent of the starting temperatures. It may be observed in Fig. 5a that it takes longer to cool the test-bed to the desired temperatures, as the test-bed is equipped with three heating elements and only one cooling element. Our results indicate that our control policy can achieve, within reasonable bounds, disparate temperatures in the same indoor space, thereby 'personalizing' the workspace.

In the absence of clear demarcations of workstations for individual occupants, such as cubicles,

TABLE VI: Mean thermal discomfort experienced when nominal temperatures $\underline{T} \in \{20, 20.5, 21, 21.5, 22\}^{\circ}$C are maintained during the scheduled activity compared to the situation where temperature personalization is implemented.

	20°C	**20.5°C**	**21°C**	**21.5°C**	**22°C**
%age difference from the thermal discomfort for Phase II	38.1	13.1	13.5	52.2	105.6

in the indoor space considered in this work, the temperatures recorded at different locations are highly correlated with each other. Moreover, they are impacted, to varying degrees, by each of the heating/cooling inputs in the indoor space. Therefore, our proposed personalization approach was seen to perform at its best when two different temperature preferences were provided to the building energy management system. However, for the sake of completion, we now study the case where occupants request five different temperature preferences.

We define $\mathbf{\Delta}$ to be the vector of temperature preferences, with element Δ_i, such that $\mathbf{\Delta} = [20, 20.5, 21, 21.5, 22]$. We compare the performance of the temperature personalization control framework presented in this work with a baseline approach which maintains the mean temperature in the indoor space close to nominal temperature, $\underline{T}$, achieved at the end of Phase I while a scheduled activity is underway in the shared space. Table VI records the %age difference of the thermal discomfort when the aforementioned baseline approach is used, with respect to the thermal discomfort when our proposed personalization approach is implemented to satisfy the temperature preferences in $\mathbf{\Delta}$. The table records results for $\underline{T} \in \{20, 20.5, 21, 21.5, 22\}^{\circ}$C. It may be observed that our personalization approach resulted in a significantly lower thermal discomfort than was the case with the baseline approach. Specifically, our personalization approach resulted in an average instantaneous thermal discomfort of 0.43°C. In contrast, this value for the baseline approach when $\underline{T} = 22^{\circ}$C was 0.88°C. Therefore, despite the limitations of our test-bed, which resulted in the temperature readings at different locations in the space to be highly correlated, using our temperature personalization framework in (**P2**) could still be beneficial to the occupants' well-being, even when five different temperature preferences are provided to the building energy management system.

We now consolidate our control strategies for Phases I and II and present simulation results for a real-life application. Here, we consider two hour-long intervals, each representing Phases

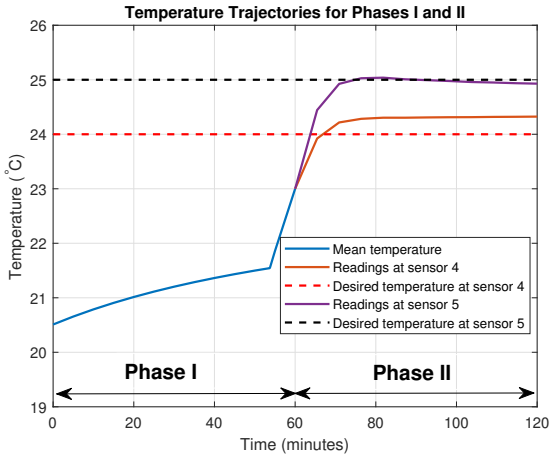

Fig. 6: Simulation results for the trajectory of the mean temperature and the temperature readings recorded at sensors 4 and 5 for Phases I and II.

I and II. For Phase I, we take $\underline{T} = 23°\text{C}$. In Phase II, the participants are assumed to be seated close to sensors 4 and 5 having temperature preferences of $24°\text{C}$ and $25°\text{C}$, respectively. Table III presents the values of $\Phi_t = \sum_{j=1}^{J} \sum_{k=1}^{2K} \hat{U}_j[k]$, which represent the total energy consumed for both Phases I and II. It may be observed that the values generally follow a trend similar to the one seen in Table I for $\underline{T} = 23°\text{C}$. Since the control strategy for Phase II is identical in each case, the value of Φ_t for each of the pre-conditioning strategies increases by roughly the same amount from the corresponding entries in Table I. Slight differences in this increase can be attributed to the disparate valve positions as well as the temperatures at each sensor at the end of Phase I for each the four pre-conditioning strategies.

Fig. 6 presents the trajectories of the temperature readings for Phases I and II when ASC is used for pre-conditioning the space to $\underline{T} = 23°\text{C}$ under simulated conditions. The occupants are assumed to be seated close to sensors 4 and 5 having temperature preferences of $24°\text{C}$ and $25°\text{C}$. It may be seen that our proposed quadratic programming-based approach for Phase II successfully achieves temperature differentiation in the indoor space. Furthermore, the temperatures at each of the locations in the indoor space reached close to the respective values of Δ_4 and Δ_5 early on in Phase II, thereby reducing the occupants' thermal discomfort.

V. Conclusion

In this paper, we used a linear, discrete formulation to model the temperature evolution in an indoor space as a function of the heating and cooling inputs, exogenous heat gains and past zone temperatures. Aided by this *data-driven* learning model, we developed control policies for *heterogeneous* HVAC elements that operate in tandem to: (i) achieve time-bound pre-conditioning of workspaces and (ii) create spatial differentiation in the thermal environment based on the occupants' individual preferences. The proposed energy efficient pre-conditioning techniques may be applied to classrooms, meeting spaces or any indoor space whose occupancy schedule can be well-estimated in advance. Specifically, we developed an alternate pre-conditioning approach (ASC) which was tailored for indoor spaces with limited computational resources. Our results showed that data-driven modeling, coupled with our proposed predictive control formulations, can not only make building operations more efficient, but also result in increased personalization in situations where an indoor space is shared by occupants with temperature preferences. The personalization approach proposed in this work assumed that the occupants were located close to the temperature sensors in the test-bed. An alternative to this assumption could involve taking the temperature experienced by an occupant to be a function of their distance from each sensor and estimating it to be the weighted average of the five sensor readings. However, following preliminary tests, we noted minimal benefits of adopting such an approach. Therefore, the temperature personalization framework considered in this work approximated an occupants' position to that of its closest sensor. In order to further enhance occupant comfort, the temperature personalization approach, as represented by (**P2**), could be extended by incorporating occupant feedback to update the set-point requirements at different locations in the test-bed. This could involve gathering qualitative feedback from occupants requesting more heating/cooling, such as that proposed in [4], thereby providing finer adjustments to the solutions obtained from (**P2**). This extension has been left as future work. The achievable personalization observed through the experiments performed as part of this work was limited to about 1°C owing to the test-bed being an *open* space, which caused the temperature at different locations in the indoor space to be highly correlated. The introduction of cubicles to the test-bed would not only better replicate the setup of a shared workspace but also allow increase the achievable temperature differentiation in the space.

VI. Acknowledgement

This work was supported by the National Science Foundation through Award No. 1827546.

References

[1] B. Goetzler, M. Guernsey, T. Kassuga, J. Young, T. Savidge, A. Bouza, M. Neukomm, and K. Sawyer, "Grid-Interactive Efficient Buildings Technical Report Series: Heating, Ventilation, and Air Conditioning (HVAC); Water Heating; Appliances; and Refrigeration," National Renewable Energy Laboratory (NREL), Golden, CO (United States), Tech. Rep., Dec. 2019. [Online]. Available: http://www.osti.gov/servlets/purl/1577967/

[2] W. Goetzler, R. Shandross, J. Young, O. Petritchenko, D. Ringo, and S. McClive, "Energy Savings Potential and RD&D Opportunities for Commercial Building HVAC Systems," EERE Publication and Product Library, Tech. Rep., Dec. 2017. [Online]. Available: http://www.osti.gov/servlets/purl/1419622/

[3] E. Azar and C. Menassa, "Agent-Based Modeling of Occupants and Their Impact on Energy Use in Commercial Buildings," *Journal of Computing in Civil Engineering*, vol. 26, no. 4, pp. 506–518, Jul. 2012.

[4] S. Gupta, S. Atkinson, I. O'Boyle, J. Drogo, K. Kar, S. Mishra, and J. Wen, "BEES: Real-time occupant feedback and environmental learning framework for collaborative thermal management in multi-zone, multi-occupant buildings," *Energy and Buildings*, vol. 125, pp. 142–152, Aug. 2016.

[5] C. Aghemo, J. Virgone, G. Fracastoro, A. Pellegrino, L. Blaso, J. Savoyat, and K. Johannes, "Management and monitoring of public buildings through ICT based systems: Control rules for energy saving with lighting and hvac services," *Frontiers of Architectural Research*, vol. 2, no. 2, pp. 147–161, 2013.

[6] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, and L. Helsen, "All you need to know about model predictive control for buildings," *Annual Reviews in Control*, vol. 50, pp. 190–232, 2020.

[7] F. Smarra, A. Jain, T. de Rubeis, D. Ambrosini, A. D'Innocenzo, and R. Mangharam, "Data-driven model predictive control using random forests for building energy optimization and climate control," *Applied Energy*, vol. 226, pp. 1252–1272, 2018.

[8] Z. Tariq, T. Imam, K. Kar, and S. Mishra, "Experimental Evaluation of Data-Driven Predictive Indoor Thermal Management," in *Proceedings of the Tenth ACM International Conference on Future Energy Systems*. New York, NY, USA: ACM, Jun. 2019, pp. 531–535.

[9] https://app.box.com/s/5s8yyj6e4y0461zvuxb8odtl1uak8838, [Online; accessed Dec. 22, 2022].

[10] Y. Li, Z. O'Neill, L. Zhang, J. Chen, P. Im, and J. DeGraw, "Grey-box modeling and application for building energy simulations - a critical review," *Renewable and Sustainable Energy Reviews*, vol. 146, p. 111174, 2021.

[11] P. Arjunan, K. Poolla, and C. Miller, "BEEM: Data-driven building energy benchmarking for Singapore," *Energy and Buildings*, vol. 260, p. 111869, 2022.

[12] P. Bacher and H. Madsen, "Identifying suitable models for the heat dynamics of buildings," *Energy and Buildings*, vol. 43, no. 7, pp. 1511–1522, 2011.

[13] 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, Feb. 2012.

[14] M. Macarulla, M. Casals, N. Forcada, and M. Gangolells, "Implementation of predictive control in a commercial building energy management system using neural networks," *Energy and Buildings*, vol. 151, pp. 511–519, 2017.

[15] L. Lei and W. Liu, "Predictive control of multi-zone variable air volume air-conditioning system based on radial basis function neural network," *Energy and Buildings*, vol. 261, p. 111944, 2022.

[16] H. Lee and Y. Heo, "Simplified data-driven models for model predictive control of residential buildings," *Energy and Buildings*, vol. 265, p. 112067, 2022.

[17] J. Joe, "Investigation on pre-cooling potential of UFAD via model-based predictive control," *Energy and Buildings*, vol. 259, p. 111898, 2022.

[18] L. Yu, Y. Sun, Z. Xu, C. Shen, D. Yue, T. Jiang, and X. Guan, "Multi-agent deep reinforcement learning for hvac control in commercial buildings," *IEEE Transactions on Smart Grid*, vol. 12, no. 1, pp. 407–419, 2021.

[19] L. Yu, S. Qin, M. Zhang, C. Shen, T. Jiang, and X. Guan, "A review of deep reinforcement learning for smart building energy management," *IEEE Internet of Things Journal*, vol. 8, no. 15, pp. 12 046–12 063, 2021.

[20] A. Aswani, N. Master, J. Taneja, V. Smith, A. Krioukov, D. Culler, and C. Tomlin, "Identifying models of HVAC systems using semiparametric regression," in *2012 American Control Conference (ACC)*. IEEE, Jun. 2012, pp. 3675–3680.

[21] S. Fux, A. Ashouri, M. Benz, and L. Guzzella, "EKF based self-adaptive thermal model for a passive house," *Energy and Buildings*, vol. 68, pp. 811–817, Jan. 2014.

[22] J. Joe and P. Karava, "A model predictive control strategy to optimize the performance of radiant floor heating and cooling systems in office buildings," *Applied Energy*, vol. 245, pp. 65–77, 2019.

[23] B. Chen, Z. Cai, and M. Bergés, "Gnu-RL: A precocial reinforcement learning solution for building HVAC control using a Differentiable MPC policy," in *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. New York, NY, USA: ACM, Nov. 2019, pp. 316–325.

[24] S. Liu and G. Henze, "Experimental analysis of simulated reinforcement learning control for active and passive building thermal storage inventory," *Energy and Buildings*, vol. 38, no. 2, pp. 148–161, Feb. 2006.

[25] D. Kim, K. Lee, S. Yoo, J. Kim, and M. Lim, "Thermal Model Parameter Estimation for HVAC Facility using Recursive Least Square Method," in *Proceedings of the 2016 International Conference on Mechanics, Materials and Structural Engineering*. Paris, France: Atlantis Press, 2016. [Online]. Available: http://www.atlantis-press.com/php/paper-details.php?id=25854571

[26] H. Wang, J. Wang, W. Li, and S. Liang, "Experimental study on a radiant leg warmer to improve thermal comfort of office workers in winter," *Building and Environment*, vol. 207, p. 108461, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S036013232100857X

[27] B. Yang, T.-H. Lei, P. Yang, K. Liu, and F. Wang, "On the use of wearable face and neck cooling fans to improve occupant thermal comfort in warm indoor environments," *Energies*, vol. 14, no. 23, 2021. [Online]. Available: https://www.mdpi.com/1996-1073/14/23/8077

[28] A. Aryal, B. Becerik-Gerber, G. M. Lucas, and S. C. Roll, "Intelligent agents to improve thermal satisfaction by controlling personal comfort systems under different levels of automation," *IEEE Internet of Things Journal*, vol. 8, no. 8, pp. 7089–7100, 2021.

[29] R. Kalaimani, M. Jain, S. Keshav, and C. Rosenberg, "On the interaction between personal comfort systems and centralized hvac systems in office buildings," *Advances in Building Energy Research*, vol. 14, no. 1, pp. 129–157, 2020.

[30] Y. He, N. Li, J. Lu, N. Li, Q. Deng, C. Tan, and J. Yan, "Meeting thermal needs of occupants in shared space with an adjustable thermostat and local heating in winter: An experimental study," *Energy and Buildings*, vol. 236, p. 110776, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378778821000608

[31] L. Yu, Z. Xu, T. Zhang, X. Guan, and D. Yue, "Energy-efficient personalized thermal comfort control in office buildings based on multi-agent deep reinforcement learning," *Building and Environment*, vol. 223, p. 109458, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360132322006898

[32] H. Boyer, J. Chabriat, B. Grondin-Perez, C. Tourrand, and J. Brau, "Thermal building simulation and computer generation of nodal models," *Building and Environment*, vol. 31, no. 3, pp. 207–214, May 1996.

[33] "Recursive Least Squares Parameter Estimation for Linear Steady State and Dynamic Models," https://cse.sc.edu/~gatzke/cache/edgar-recursive-estimation.pdf, [Online; accessed May 16, 2021].

[34] J. Jiang and Y. Zhang, "A revisit to block and recursive least squares for parameter estimation," *Computers Electrical Engineering*, vol. 30, no. 5, pp. 403–416, 2004.

[35] "Recursive Least Squares with Forgetting for Online Estimation of Vehicle Mass and Road Grade: Theory and Experiments," http://www-personal.umich.edu/~annastef/papers_Long_ctrl/JournalPaperMassGrade_Final.pdf, [Online; accessed Jan. 28, 2023].

[36] "Geometry and visualizations of linear programs," https://ocw.mit.edu/courses/sloan-school-of-management/15-053-optimization-methods-in-management-science-spring-2013/lecture-notes/MIT15_053S13_lec3.pdf, 2013, [Online; accessed May 16, 2021].

[37] J. Skaf, S. Boyd, and A. Zeevi, "Shrinking-horizon dynamic programming," *International Journal of Robust and Nonlinear Control*, vol. 20, no. 17, pp. 1993–2002, Nov. 2010. [Online]. Available: http://doi.wiley.com/10.1002/rnc.1566