ACHIEVING IMPROVED PERSONALIZATION AND ENERGY EFFICIENCY IN COHABITED WORK-SPACES THROUGH DATA-DRIVEN PREDICTIVE CONTROL

Syed Ahsan Raza Naqvi, Zachary Nawrocki, Zaid Bin Tariq, Koushik Kar and Sandipan Mishra
Rensselaer Polytechnic Institute, USA
Email: {naqvis2,nawroz,tariqz,kark,mishrs2}@rpi.edu

ABSTRACT

This paper studies the problem of indoor zone temperature control in shared work-spaces equipped with heterogeneous heating and cooling sources with the goal of increased energy savings and environment personalization. We consider two scenarios to assess the performance of our control strategies. The first scenario requires time-bound pre-cooling/pre-heating of a shared space in preparation for a scheduled activity (Scenario A). The second scenario considers a cohabited work-space where occupants have different temperature preferences (Scenario B). Utilizing an on-campus smart conference room (SCR) as a test-bed, we use data-driven model learning to establish a relationship between the room's heating, ventilation and cooling (HVAC) operations and the zone temperatures. Next, we use a model predictive control (MPC)-based approach to achieve a desired average temperature while minimizing power consumption (for Scenario A) and minimize the thermal discomfort experienced by individuals based on their temperature preferences (for Scenario B). The experimental results show that for Scenario A, the proposed control policy can save a significant amount of energy and achieve the desired mean temperature in the space fairly accurately. We further note that for Scenario B, the control scheme can achieve a significant spatial differentiation in temperature towards satisfying the occupants' thermal preferences.

1 Introduction

Since HVAC systems account for approximately one-third of the total energy consumption in buildings in the United States [1], there is a need for making such systems more energy efficient. At the same time, the increasing desire for personal comfort and wellness among building occupants means that HVAC systems must strive to attain disparate temperatures within the same indoor environment. In the presence of the occupants' temperature preferences, as posited in [2], [3], mapping these demands to the operation of the HVAC system in the building involves several challenges. Firstly, the different zones in the shared space where the occupants are located may be thermally correlated, and it may not be possible to satisfy the preferences of all occupants simultaneously. Furthermore, the building may be equipped with multiple sources of heating and cooling and modeling their combined effect on the temperature evolution at various points within the indoor environment is challenging. In addition, identifying dynamically changing thermal zones within a large, open space is non-trivial.

A significant body of past research work has studied the use of MPC in building energy management systems. For instance, [4] used semi-parametric regression to map temperature changes in discrete time to HVAC inputs by estimating the heating load due to occupancy and solar heat in an indoor space. The authors employed a learning-based MPC to estimate the occupancy heating load and to adjust the control action accordingly. The authors in [5] 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. In [6], a probabilistic framework derived from historical data was used to model the uncertain load forecast in the thermal zones of an on-campus building. Subsequently, a stochastic MPC approach was presented to minimize the expected energy costs for temperature regulation within certain bounds. In [7], the authors used data-driven learning and predictive control to efficiently achieve desired temperatures in an occupied indoor space. Although related literature has extensively studied indoor thermal management, there is a need to design control frameworks to achieve efficient time-bound pre-cooling/pre-heating of a work-space, 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.

We propose a learning-based approach to estimate the model parameters by determining the dependence between the heterogeneous HVAC controls (the inputs) and the temperature at specific sensor locations (the outputs) within a shared space using a simple dynamical model. Additionally, we also show how the model can be further simplified by ignoring some dependencies as well as by temporal averaging. We utilize an on-campus SCR as a test-bed to train the model. Next, we use our dynamical model in conjunction with a MPC-based approach to achieve desired temperatures in the test-bed. We employ our control policy to pre-cool/pre-heat the space to a specified temperature prior to the work-space being occupied, while minimizing the total power consumed in the process. Furthermore, by exploiting spatial differences in the impact of the different thermal inputs, our control algorithm is also shown to satisfy disparate thermal preferences (within known bounds) during periods of occupancy.

It is noteworthy that [4] – [6] considered only a single type of input, the air handling unit (AHU), for heating and cooling, while formulating the indoor thermal model. In contrast, our control strategy simultaneously utilizes the operation of *multiple heterogeneous* HVAC elements (i.e. AHUs and radiators) to achieve the desired objectives. Additionally, in view of the time-bound nature of the efficient pre-cooling/pre-heating problem, we use a shrinking horizon MPC approach [8], unlike [4] – [6]. In contrast to the work in [5], [6], we evaluate the performance of our control policies *experimentally*. Moreover, unlike [7], this paper poses a fixed end-point problem to achieve efficient pre-cooling/pre-heating of a shared space. Also in contrast to [4] – [7], our work aims to satisfy the thermal preferences of individual occupants.

2 Problem Formulation

We consider a typical work-space occupied by multiple occupants and equipped with heterogeneous heating and cooling sources. This indoor space is instrumented with temperature sensors at various locations, which may be used to estimate the temperatures at the occupants' locations. This work uses a data-driven learning model to determine the effect of HVAC operations on the temperature at these locations.

2.1 Data-driven learning model

We use a black box model to establish a relationship between the temperature measurements at multiple sensors and the heterogeneous heating and cooling sources in the indoor space. We discretize our time horizon into K time instances. Each time instance k has duration μ minutes. The model used here aims to predict the temperature measured by each sensor at the beginning of each time instance, using the temperature readings and heating or cooling input from past instances. Furthermore, the formulation assumes that the heating or cooling input remains constant for μ minutes. Taking the R-C model in [9] as a motivation, we can use the following discrete, linear model to estimate the temperature evolution in the test-bed:

$$\mathbf{Y}[k+1] = \mathbf{A}\mathbf{Y}[k] + \mathbf{B}\mathbf{U}[k] + \mathbf{D}y_\infty[k] + \mathbf{W}[k], \tag{1}$$

where $\mathbf{Y}[k] \in \mathbb{R}^{I\times 1}$ is the vector of temperature measurements at I sensors at time k, $\mathbf{U}[k] \in \mathbb{R}^{J\times 1}$ is a vector of input signals to the J heating and cooling elements of the HVAC system and function $f(\mathbf{U}[k])$ is the instantaneous energy consumption of the system. Finally, $y_\infty[k]$ is the ambient temperature at time instance k. $\mathbf{A}, \mathbf{B}, \mathbf{D}$ and $\mathbf{W}$ are matrices that are to be estimated through learning. Here, $\mathbf{W}$ captures unmodeled dynamics such as heat from human bodies, the server and workstations, solar gains etc.

It is worth noting, however, that in (1), y_∞ only changes very slowly. Therefore, the training data collected for this quantity lacks the required richness to adequately estimate $\mathbf{D}$. Since we aim to achieve specific temperatures over a much smaller time scale (tens of minutes), we drop the term for the ambient temperature in our final model to obtain,

$$\mathbf{Y}[k+1] = \mathbf{A}\mathbf{Y}[k] + \mathbf{B}\mathbf{U}[k] + \hat{\mathbf{W}}[k], \tag{2}$$

where the bias $\hat{\mathbf{W}}$ captures the effect of the noise term $\mathbf{W}$, as well as that of the ambient temperature on $\mathbf{Y}[k+1]$. Details of the derivation of this data-driven model may be found in [7].

Given this thermal model of the shared space, we now proceed to formulating the control strategy to determine the series of HVAC operations required to achieve a desired temperature within a fixed time window in an efficient manner.

2.2 Time-bound pre-cooling/pre-heating of the test-bed

A previously unoccupied indoor space might need to be pre-cooled or pre-heated to a nominal temperature, $\tilde{Y}$, in preparation for some scheduled activity e.g. a work-related meeting. In this case, prematurely heating or cooling the space to $\tilde{Y}$ would be energy inefficient. Therefore, in order to save energy, it is vital that HVAC operations are controlled such that the average temperature in the space reaches approximately $\tilde{Y}$ immediately before the scheduled activity. This condition may be expressed as,

$$\frac{\sum_{i=1}^{I} Y_i[K]}{I} = \tilde{Y}, \tag{3}$$

where K represents the time at which the scheduled activity begins and $Y_i[K]$ is the temperature measurement at sensor i at that time. In order to achieve time-bound pre-cooling/pre-heating, we solve the following fixed end-point problem:

$$\min_{\mathbf{U}}. \sum_{k=1}^{K}\sum_{j=1}^{J} U_j[k], \text{ s.t. } (2), (3), \tag{\textbf{P1}}$$

where $U_j[k]$ represents the energy used by HVAC input j. The decision variables for the optimization problem are the elements in matrix $\mathbf{U}$, where $\mathbf{U} \in \mathbb{R}^{J\times K}$. The constraint (3) ensures that the mean sensor readings at the end of time K must ideally equal $\tilde{Y}$.

Given a deadline K, a shrinking horizon MPC approach [8] is used to determine the desired trajectory of the heating and cooling inputs. After every μ minutes, our control strategy uses the temperature readings at all sensors as well as the heating and cooling inputs at that time to solve **P1**. However, at each instance k, the problem is solved over $(K-k)$ time instances. Solving the pre-heating/pre-cooling problem using MPC for a progressively diminishing optimization window was seen to produce the desired results, as will be seen in Section 3.

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.

2.3 Satisfying individual temperature requirements

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

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

where K' represents the size of the optimization window. Δ_i is the desired temperature at sensor i's location, which is computed based on the occupants' preferences and their locations relative to the sensor. Unlike the policy for pre-cooling/pre-heating, 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 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 $\bar{\mu} < \mu$.

3 Experimental Evaluation

3.1 Test-bed layout

Fig. 1 shows the layout of the SCR 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 SCR is instrumented with five wall-mounted temperature sensors located at various positions, as shown in Fig. 1a. Temperature readings, taken once every minute, are wirelessly transmitted to a server in the room.

3.2 Results

The experimental evaluation of the dynamical model and the control strategy considers two scenarios. The first (Scenario A) requires time-bound pre-cooling or pre-heating of the shared space in anticipation of a scheduled work activity. The second scenario (Scenario B) considers a cohabited work-space where occupants at two locations in the shared space have two different temperature preferences. In view of the facts that the control policies for the aforementioned scenarios operate at a relatively shorter time scale and that the space has significant thermal inertia, μ was chosen to be five minutes for training the model in (2). The time till the scheduled activity in Scenario A was taken to be sixty minutes. The experiments were conducted in Feb. 2020.

3.2.1 Scenario A For Scenario A, we use an MPC-based approach to achieve a desired average temperature while minimizing power consumption. Figs. 2a–2c show the evolution of the mean temperature (averaged across all five sensors in the test-bed) when the space needs to be pre-heated or pre-cooled over a period of an hour. All three tests were conducted at night. It can be observed that the desired temperatures can be achieved to within approximately 0.25°C by the end of the hour, which represents the beginning of the scheduled activity. These results also demonstrate that the discrete, linear model in (2) is able to satisfactorily capture the relationship between the HVAC operations and the change in temperature. Figs. 2d–2f show how the optimization in **P1** ensures that the heating and cooling sources are only utilized a short period before the start of the planned work activity, thereby reducing the energy costs incurred by the building operator. It may also be seen that the data-driven approach adequately models the test-bed's rate of heat loss to the external environment. Therefore, as can be observed in Fig. 2d, the time-bound pre-cooling strategy exploits not only the test-bed's HVAC operations but also the colder ambient conditions. Once the heat loss to the ambient alone results in a fall of 0.32°C by the forty minutes mark, the AHU cooling input 'turns on' for ten minutes to accelerate the cooling operation. Furthermore, the HVAC system relies on the AHU to achieve pre-heating in Figs. 2e, 2f, as it can heat up the space quicker than the radiators. Also, turning on the radiators would have had a greater impact on the readings at sensors 1 and 5 (see Fig. 1a), whereas our aim is to drive the *average* temperature readings of all five sensors to $\tilde{Y}$.

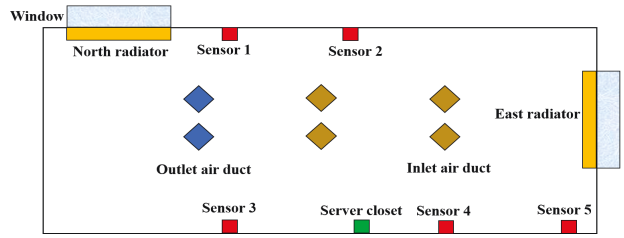

(a)

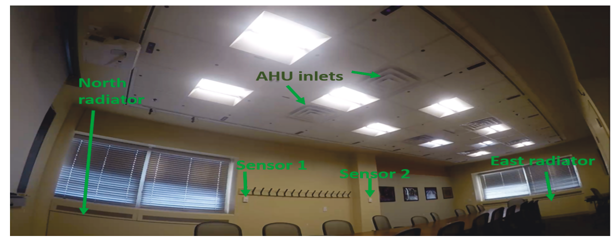

(b)

FIGURE 1: (1a) Test-bed layout; (1b) Photograph of the SCR.

Table 1 records the performance of control strategy **P1** for various values of $\tilde{Y}$ in terms of the mean final temperature deviation and the duration for which the test-bed's HVAC elements are on (non-zero % valve opening). It may be seen that all experiments resulted in mean final temperature deviations of less

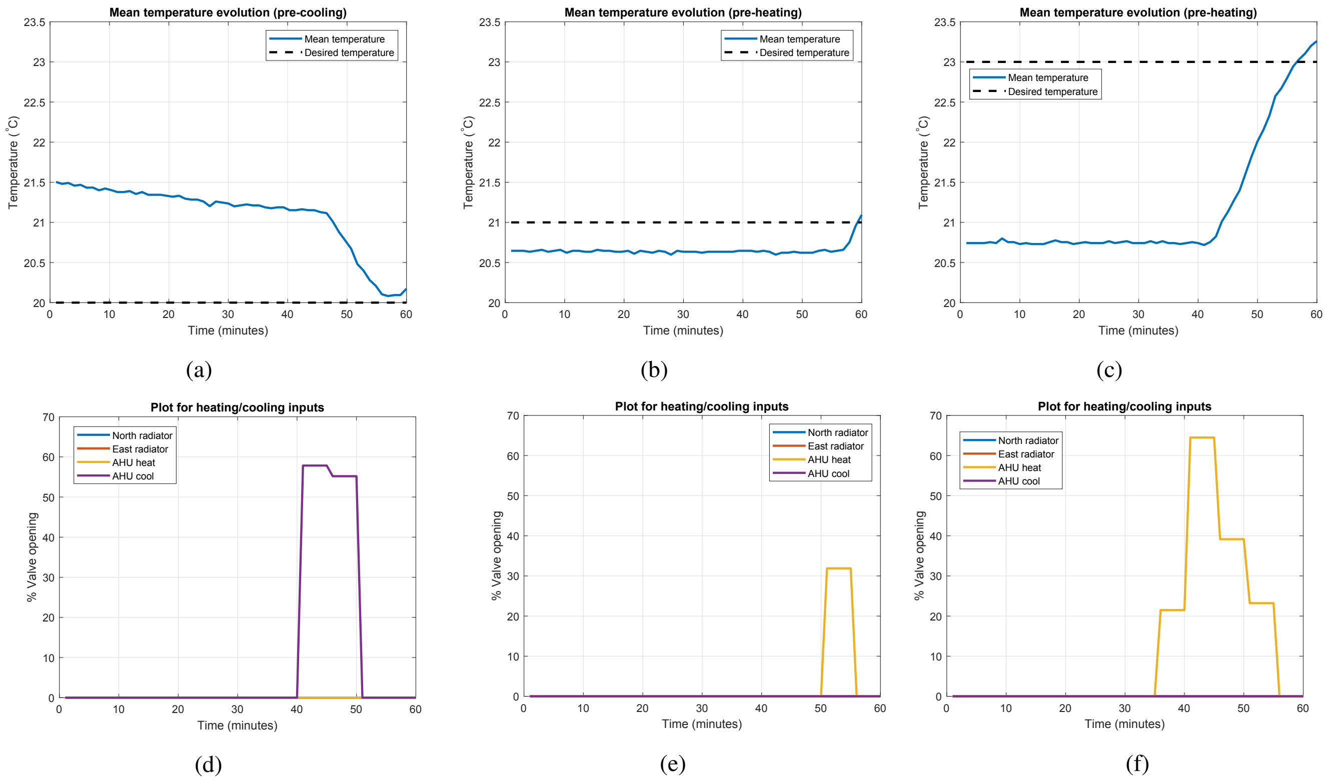

FIGURE 2: (2a): Pre-cooling the test-bed to $\tilde{Y} = 20°C$; (2b): Pre-heating the test-bed to $\tilde{Y} = 21°C$; (2c): Pre-heating the test-bed to 23°C; (2d): The HVAC operation required to pre-cool the test-bed to $\tilde{Y} = 20°C$; (2e): The HVAC operation required to pre-heat the test-bed to $\tilde{Y} = 21°C$; (2f): The HVAC operation required to pre-heat the test-bed to $\tilde{Y} = 23°C$.

Performance \ $\tilde{Y}$	**19°C**	**21°C (II)**	**26°C**
Starting temperature (°C)	20.75	22.51	21.86
Temperature deviation (°C)	0.22	0.09	0.38
Operation	Cooling	N/A	Heating
$\frac{\text{Duration HVAC elements are on}}{\text{60 min.}}$	0.42	0	0.67

TABLE 1: Performance of **P1** for various values of $\tilde{Y}$.

HVAC element \ Strategy	**P1**	**B1**	**B2**	**B3**
North Radiator (%)	0	8.25	0.61	21.82
East Radiator (%)	0	6.20	1.95	21.78
AHU heating (%)	2.66	8.04	1.09	21.76
AHU cooling (%)	0	3.08	4.55	1.13

TABLE 2: Comparison of the average valve openings for the proposed and the benchmark approaches for pre-heating to 21°C.

than 0.4°C, despite the HVAC elements not being on for the entire hour. The test labeled '**21°C(II)**' was conducted during the day to contrast with the experiment in Figs. 2b, 2e, which was run at night. The results show that our thermal model correctly estimated the rate of heat loss to the surroundings and thus did not activate any of the HVAC elements. The heat loss to the ambient alone resulted in the mean final temperature deviation to be less than 0.1°C.

We now compare the performance of our proposed pre-cooling/pre-heating control policy with three benchmark strategies, **B1**, **B2** and **B3**. **B1** uses an optimization framework similar to that in [7] to minimize the deviation of the space's temperature from $\tilde{Y}$ for the entire hour prior to the scheduled activity. **B2** solves the same objective as **B1**, however, in view of the results in Fig. 2c, it comes into effect at the 50 minutes mark for a fairer comparison with **P1**. Similarly, **B3** represents the ten-minute operation of the set-point-based building management system (BMS), manufactured by Johnson Controls, which is currently installed in the SCR. Table 2 records the valve positions for each HVAC element (averaged over sixty minutes) for pre-heating the test-bed to 21°C. Since, the average % valve openings for each HVAC element (other than AHU heating for **B2**) is higher for the benchmark approaches than **P1**, it can be concluded that our proposed time-bound pre-cooling/pre-heating policy is more ef-

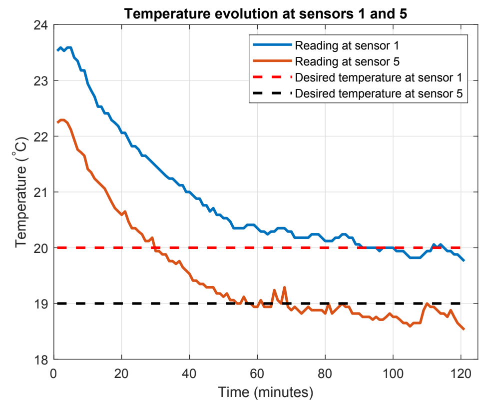

FIGURE 3: Temperature differentiation in the shared space when the desired temperatures are 20°C and 19°C.

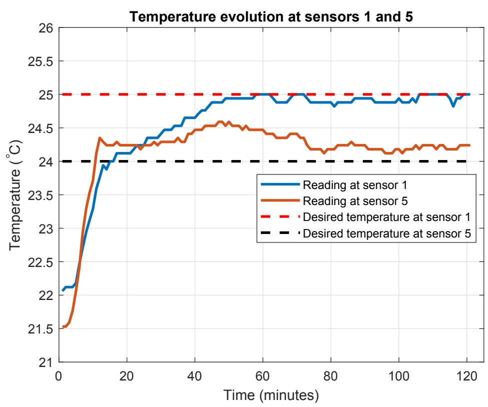

FIGURE 4: Temperature differentiation in the shared space when the desired temperatures are 25°C and 24°C.

ficient than all of the considered benchmarks. The results also show that thermal management can be improved by employing spatially distributed sensors, rather than basing policies on the readings of a single sensor as done in **B3**.

3.2.2 Scenario B Next, we consider Scenario B where we aim to minimize the thermal discomfort experienced by two individuals already present in the shared space, based on their temperature preferences. The occupants are located close to sensors 1 and 5, respectively. Thus, we assume that the temperatures recorded by each of these sensors are the temperatures experienced by the occupants. Figs. 3 and 4 depict how the HVAC operations determined using **P2** can successfully achieve a spatial temperature differentiation of upto 1°C for heating and cooling operations, independent of the starting temperatures. It may be observed in Fig. 3 that it takes longer to cool the test-bed to the desired temperatures. This may be attributed to the fact that the test-bed is equipped with three heating elements and only one cooling element. Our results indicate that through data-driven learning, our control policy can achieve, within reasonable bounds, disparate temperatures in the same indoor space, thereby 'personalizing' the work-space.

4 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 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-cooling/pre-heating of work-spaces and (ii) create spatial differentiation in the thermal environment based on the occupants' individual preferences. Our results show that data-driven modeling, coupled with MPC-based formulations, can not only make building operations more efficient, but also result in increased personalization of the work-space.

5 Acknowledgments

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

REFERENCES

[1] U.S. Department of Energy, "Energy savings potential and RD & D opportunities for commercial building HVAC systems," Dec. 2017, [Online] Available: https://www.energy.gov/sites/prod/files/2017/12/f46/bto-DOE-Comm-HVAC-Report-12-21-17.pdf

[2] V. Erickson and A. Cerpa, "Thermovote: Participatory sensing for efficient building HVAC Conditioning," *ACM BuildSys*, NY, USA, 2012, pp. 9-16.

[3] S. Gupta *et al.*, "BEES: Real-time occupant feedback and environmental learning framework for collaborative thermal management in multi-zone, multi-occupant buildings," *Energy and Buildings*, vol. 125, 2016, pp. 142-152.

[4] A. Aswani *et al.*, "Identifying models of HVAC systems using semiparametric regression," *American Control Conference (ACC)*, Montreal, QC, 2012, pp. 3675-3680.

[5] Y. Ma, A. Kelman, A. Daly and F. Borrelli, "Predictive control for energy efficient buildings with thermal storage: Modeling, simulation, and experiments," in *IEEE Control Systems Magazine*, vol. 32, no. 1, , Feb. 2012, pp. 44-64.

[6] Y. Ma, J. Matuško and F. Borrelli, "Stochastic model predictive control for building HVAC systems: Complexity and conservatism," in *IEEE Transactions on Control Systems Technology*, vol. 23, no. 1, Jan. 2015, pp. 101-116.

[7] Z. Tariq, M. Imam, K. Kar and S. Mishra, "Experimental evaluation of data-driven predictive indoor thermal management," *ACM e-Energy*, Phoenix, AZ, June 2019, pp. 531-535.

[8] J. Skaf, S. Boyd and A. Zeevi, "Shrinking-horizon dynamic programming," in *Int. J. Robust Nonlinear Control*, vol. 20, no. 17, 2010, pp. 1993-2002.

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