A Model-Predictive Controller for Joint Smart Building Systems



Prashanth Ganesh

Electrical Engineering and Computer Sciences University of California at Berkeley

Technical Report No. UCB/EECS-2019-51 http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-51.html

May 16, 2019

Copyright © 2019, by the author(s). All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

I would like to thank Professor Costas Spanos for his mentorship and advice over the past 2 years. Additionally, I would like to thank Ruoxi Jia, who mentored me while I was an undergraduate student and brought me into CREST. Finally, I would like to thank Professor Kameshwar Poolla for his advice in the completion of this report.

A Model-Predictive Controller for Joint Smart Building Systems

Prashanth C. Ganesh

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science**, **Plan II**.

Approval for the Report and Comprehensive Examination:

Committee

C.Z.STANOT

Professor Costas J. Spanos Research Advisor

05/17/2019

(Date)

* * * * * * *

Professor Kameshwar Poolla Second Reader

05/17/2019

(Date)

Abstract

In recent years, buildings have become a major source of global energy consumption. A significant portion of that can be attributed to the heating, ventilation, and air conditioning (HVAC) and electric lighting systems. With the advent of new smart building technologies that can, among other things, monitor environmental data, occupancy, and remotely control various building systems, it becomes possible to implement intelligent systems that can alleviate these issues. In this work, we propose a lightweight, decentralized controller that minimizes building energy consumption while maintaining occupant comfort by integrating the aforementioned smart building technologies with EnergyPlus simulation and programmatic building control in a model-predictive control loop. We demonstrate its efficacy on two real-world implementations in the CREST lab in Cory Hall and illustrate that with a short data collection period, it is rapidly deployable and can offer significant energy savings and presents extensions to differential privacy while achieving the desired comfort objectives.

Contents

List of Symbols					
1	1 Introduction				
2 Related Work				6	
3	Methodology				
	3.1	Model	-Predictive Controller	8	
	3.2	Tempe	erature and Brightness Models	9	
4	Experimental Implementation				
	4.1	CREST	Conference Room (HVAC Only with Privacy-Preservation)	10	
		4.1.1	Temperature Model	10	
		4.1.2	Occupancy Detection and Privacy Preservation	12	
		4.1.3	Results	12	
	4.2	CREST	Iab implementation	14	
		4.2.1	MPC Design	14	
		4.2.2	EnergyPlus Model	15	
		4.2.3	Temperature and Brightness Models	16	
5	Results				
	5.1	Energy	y Expenditure	19	
	5.2	Lumin	escence	19	
	5.3	Tempe	erature	19	
6	Discussion and Future Work				
Bibliography					

List of Symbols

U_k^n	Control inputs for K timesteps across N zones
J	Cost function
$P_{h,t}^z, P_{d,t}^z$	HVAC and electric lighting power consumption (W)
T_k^z	Temperature (°C) in zone z at timestep k
B_k^z	Brightness (lx) in zone z at timestep k
S_k^z	Supply air temperature (°C) in zone z at timestep k
L_k^z	Dimmer level $[0, 1]$ in zone z at timestep k
θ_k^z	Blind angle [0, 100] in zone z at timestep k
O_k^z	Occupancy (\mathbb{N}) in zone z at timestep k
B_k	Solar radiation (lx) at timestep k

1 Introduction

In 2013, the US Energy Information Administration projected that world energy consumption will increase 56% by 2040 [1]. Buildings account for nearly 40% of global energy usage, of which 40% and 15% are consumed by HVAC and lighting, respectively. Thus, it is of paramount importance to optimize the energy consumption of buildings in pursuit of a more sustainable energy future. More recently, the rapid growth of smart building technologies present valuable opportunities to closely couple information processing and optimal control with the operation of building energy systems [2]. These include technologies such as Building Operating System Services (BOSS), a distributed operating system for buildings, and Building-in-Briefcase (BiB), a rapidly deployable portable sensor platform to monitor a variety of environmental data [3, 4]. Such developments offer platforms on which intelligent control schemes can be deployed to optimize building energy usage.

The use of control schemes ranging in complexity from rule-based control actions to model predictive control for HVAC and lighting have been shown to be effective in reducing the energy cost of buildings. However, the controls over different building systems have remained largely separated. There have been a few studies that explore the synergistic benefits of joint control of multiple building systems, while the benefits have been mostly validated in simulation environments. In this work we develop a joint optimization-based control scheme for the HVAC, electric lighting, and blind systems that is designed to be fairly lightweight and decentralized. The objective is to develop a software platform for robust and rapid deployment of integrated controls over different building systems and demonstrate possible energy and non-energy benefits of integrated control on real-world sites. We pair this system with the aforementioned BiB sensor network to validate its effectiveness in optimizing energy use while maintaining occupant comfort through a realization in the CREST lab in Cory Hall at UC Berkeley.

2 Related Work

Extensive recent research has been conducted in intelligent control schemes for HVAC and lighting systems. These range from simple rule-based methods to model-predictive control techniques (such as that utilized in this study) and advanced methods that utilize deep reinforcement learning.

Early work includes simple heuristic models such as a rule-based shading design to select glazing areas and shading properties for minimizing energy demand from HVAC and lighting systems developed by [5]. Similarly, [6] developed rule-based free day cooling based on temperature as well as shading control based on temperature and irradiation. While effective in practice, neither of these studies utilized optimal control techniques. The authors of [7] present a stochastic receding horizon control and energy forecasting method to minimize the expected value and variance of primary energy consumption over a series of time windows in the control horizon. They forecasted electricity and hot-water demand over numerous scenarios, but did not account for occupant thermal loads. [8] introduces a stochastic model-predictive control scheme for building HVAC systems using an analytical energy expenditure model and simple occupancy dynamics to maintain comfort and operational constraints, and was experimentally successful in minimizing energy cost and maintaining thermal comfort. Similar optimal control strategies for HVAC systems are investigated in [9], which utilizes swarm intelligence to determine energy expenditure in the various HVAC systems of interest, and [10], which investigates the use PID and LOR controllers for this purpose. A token-based method is proposed in [11] which specifies zone cooling needs in terms of token requests, which are gathered by a central scheduler that allocates tokens in a model-predictive control framework in such a way as to minimize overall energy consumption. This method is scalable to large commercial buildings and has been shown to be effective when compared to centralized strategies.

In addition, numerous studies have used more recent deep reinforcement-learning techniques to optimally control HVAC systems, including [12] which formulated control actions as a Markov-decision process and used an EnergyPlus building model for offline training and validation. It proved to be effective in maintaining thermal comfort while achieving energy cost savings. Similarly, [13, 14] also used deep reinforcement learning to achieve significant energy savings for HVAC and hot water heaters, respectively. Furthermore, as in some of the aforementioned studies, [15, 16]

use the EnergyPlus simulation tool to develop a model-predictive control system for HVAC systems and demonstrably reduce energy expenditure compared to baseline approaches. These studies thus validate the effectiveness of using EnergyPlus in a model-predictive control scheme but do not integrate lighting and shading systems. Moreover, some do not experimentally validate the methods in real-world spaces. A more comprehensive study of integrated control of all of the aforementioned systems (HVAC, electric lighting, and shading) was conducted by [17], who developed a rigorous model-predictive optimization. They modeled interactions of these building functions to minimize the total daily energy cost. They also recognized the computational intractability of such an optimization and further proposed a system based on stochastic dynamic programming and the rollout technique to simplify the optimization. Their system saved computation time and offered energy cost advantages over baseline methods. However, it was done entirely using simulations with the DeST building software.

As we can see, there has been extensive research in the individual optimal control of building HVAC, electric lighting, and blind systems, but limited study as to the feasibility and effectiveness of a joint controller over all of these systems. In this work we present such a system that utilizes joint control while also incorporating the aforementioned techniques such as EnergyPlus simulation in a model-predictive control loop.

3 Methodology

There are several key components of the described system, including a collection of brightness and temperature dynamics models for the sub-areas of interest, an occupancy detection mechanism, an EnergyPlus model of the space that accurately captures the energy consumed by the aforementioned systems, and a model-predictive controller to solve an optimization problem to obtain the HVAC and lighting control actions required.

3.1 Model-Predictive Controller

Given a cost function *J* and a set of control inputs $U_{1:K}^{1:N} = \{U_k^n | k = 1, ..., K, n = 1, ..., N\}$ for *K* timesteps across *N* zones as described in [18], the general formulation of the optimization problem for a model-predictive controller is given by

$$\min_{U_{1:K}^{1:N}}\sum_{n=1}^{N}J^{n}$$

subject to a set of inequality and/or equality constraints on U.

Specifically, in order to design an MPC to control the HVAC, electric lighting, and blind systems in the lab, the area must be divided into zones, each with their respective temperature and brightness models. Additionally, an energy cost function must be developed for the optimization problem at hand. The total energy cost from time t = 1, ..., T from zones z = 1, ..., Z can be given as:

$$J_T = \sum_{z=1}^{Z} \sum_{t=1}^{T} P_{h,t}^z + P_{d,t}^z \Delta t$$

Where $P_{h,t}^z$ corresponds to the HVAC power and $P_{d,t}^z$ corresponds to the power consumed by the dimmer-controlled electric lighting system in that zone. We assume the energy consumed by the blinds controller to be negligible. We minimize the above cost with respect to the supply air temperature and mass flow rate (using the zone setpoint as a proxy), electric lighting dimming level, and blind levels for each zone at each interval. In this implementation, the optimization

problem is solved for 15 minute timesteps over 4 hours (to obtain 16 control actions), and the first one is applied.

We obtain the optimal control actions for the above systems subject to the following constraints on thermal comfort, visual comfort, and physical limitations of the systems of interest:

 $T_{min,occ} \leq T_k^z \leq T_{max,occ} \qquad \text{Occupied} \\ T_{min,empty} \leq T_k^z \leq T_{max,empty} \qquad \text{Unoccupied} \\ B_{min,occ} \leq B_k^z \qquad \text{Occupied} \end{cases}$

3.2 Temperature and Brightness Models

We must also develop dynamics models for the temperature and brightness in the zones, with dimmer level *L*, blind angle θ , occupancy *O*, supply air temperature *S* and external solar radiation *W*:

$$T_{k+1}^{z} = f(T_{k}^{1}, \dots, T_{k}^{Z}, S_{k}^{z}, L_{k}^{z}, \theta_{k}^{z}, O_{k}^{z})$$

and

$$B_{k+1}^{z} = g(W_{k}^{z}, L_{k}^{1:z}, \theta_{k}^{1:z})$$

In the experimental implementations below, the models were constructed by fitting linear and multi-layer perceptron models to environmental data that was collected over the course of 2-3 days under varying temperature, brightness, and weather conditions.

4 Experimental Implementation

4.1 CREST Conference Room (HVAC Only with Privacy-Preservation)

Located in the center of the CREST space is a conference room that lends itself to the implementation of a simplified version of the MPC described above. The objectives behind this particular implementation were to improve the energy efficiency of the CREST conference room, maintain thermal comfort, and validate the privacy mechanism described in [18] using a real-world testbed. This privacy preservation architecture distorts the occupancy such that individual occupant information is hidden while maintaining a control performance guarantee regarding HVAC performance and comfort.

The model-predictive controller used in this implementation focused solely on optimizing the room temperature for energy cost while maintaining comfort. Thus, a simplified energy consumption proxy was used based on the idea that the energy used to heat or cool the room was proportional to the difference between the desired supply air temperature and the outdoor temperature (which is heated or cooled to produce the supply air temperature:

$$J_T = \sum_{t=1}^{T} |T_{supply,t} - T_{outside,t}|$$

4.1.1 Temperature Model

In order to implement this system, a temperature model was developed by collecting temperature dynamics in the conference room over the course of several hours with varying levels of occupancy (thus capturing the impact of occupancy on the temperature in the room). The temperature at the next time-step was modeled as a linear function of the current temperature, the temperature of the area of the CREST lab outside of the conference room, the supply air temperature, and the occupancy of the room.

$$T_{k+1} = f(T_k, T_{k, external}, S_k, O_k^z)$$



Figure 4.1: Conference Room Temperature Model

4.1.2 Occupancy Detection and Privacy Preservation

In order to accurately collect the occupancy of the conference room in real time, an occupancy counting module was built using 2 laser diode/light dependent resistor pairs. If an occupant breaks one beam before another, it indicates that they have travelled in one direction in relation to the room which can be interpreted as entering or exiting the room.

As mentioned earlier, the occupancy of the room was masked to preserve occupant privacy while also maintaining HVAC performance. In order to do so, the optimization problem was solved several times using several occupancies in the neighborhood of the true occupancy. A set was constructed of those false occupancies that resulted in a control action within a 1 degree tolerance of that of the true occupancy. One occupancy was selected from this "equivalence" set and reported to the controller to use to determine the optimal control action. This method used the key concepts described in [18] but was also computationally tractable to use in this practical physical implementation. It is important to note that in true privacy preservation, we want to stem the use of the true occupancy as early as possible (ideally, at the sensor itself). This implementation assumes that there is a safe place to perform the computations described earlier, which is a decision to be made as part of using a privacy-masking scheme.

4.1.3 Results

Figure 4.2 summarizes the results of the implementation described above in the CREST conference room. These experiments were conducted over the course of two days- the first day reported the true occupancy to the controller, and the second day reported the masked occupancy. The second row of graphs illustrates the true occupancy, masked occupancy, and the size of the equivalence set (note that this set is not in fact an occupancy but rather the number of possible values to report to the controller; since the true occupancy of the room is limited and on the same scale as size of the set, we include them on the same graph for visualization). The last row of graphs demonstrates the resulting thermal comfort levels and real room temperature over the course of the experimentation periods. As we can see, using the true occupancy resulted in the controller successfully maintaining thermal comfort (note that the thermal constraints are illustrated by the dashed lines). Significantly, the masked occupancy experiment was also successful in this objective.



Figure 4.2: Conference Room Controller Results



Figure 4.3: CREST Floor Plan

4.2 CREST lab implementation

Figure 4.3 is a diagram of the Center for Research in Energy Systems Transformation (CREST) floor-plan. The space is equipped with two HVAC units, eight electric lighting units, and three large windows with controllable blinds. As we can see, there are 6 desk areas in an L-shape, each capable of seating four researchers. To cover this area there is a Sensortag network described in [4] collecting a wide array of environmental data including temperature and brightness.

4.2.1 MPC Design

The design of the model-predictive controller in this implementation is similar to that of the one used in the above conference room experiment, with some key differences. The control action U^t is composed of the supply air temperatures for the north and west zones $S_Z^t \in [0, 100^o F]$, the electric lighting levels for each of the 8 lighting units $L_u^t \in [0, 1]$, and shade levels for the 3 controllable shades $\theta_s^t \in [0, 100]$. Every 15 minutes, the MPC attempts to solve for 16 such control actions to be applied every 15 minutes over a horizon of 4 hours. Over this entire 4 hour period, the constraints on the temperatures of each zone are given by the following to ensure thermal comfort, and the constraints on the brightness in each zone are also enforced to ensure visual comfort. The MPC uses the temperature models of each zone (mentioned above and described below) to forecast the result of applying these control actions.

$24^{o}C \le T_k^z \le 26^{o}C$	Occupied
$23^{o}C \le T_k^z \le 27^{o}C$	Unoccupied
$150lx \le B_k^z$	Occupied

In addition, there are bounds on the maximum and minimum allowable supply air temperatures $(30^{\circ}C \text{ and } 10^{\circ}C \text{ respectively}).$

4.2.2 EnergyPlus Model

As mentioned earlier, an integral part of the model-predictive control loop is the cost function, which captures the energy consumed by the HVAC system and electric lighting. Two options emerge when considering how to measure the energy consumption of these systems- a physical metering system, or a model that accurately captures the interactions of the various systems and loads in a space. It would be relatively straightforward to install a meter to measure electric lighting consumption. However, since the CREST lab is part of a larger building, it is difficult and expensive to directly measure HVAC energy consumption as this would involve installing sensors within the building's system to determine how much of the heating and cooling costs are due to the space at hand. Therefore, we chose to utilize the Department of Energy's EnergyPlus program- a building energy simulation program that allows for the construction and simulation of custom buildings and spaces. As mentioned earlier, the use of EnergyPlus is well validated in related works on optimal control of building HVAC systems.

The objective function is given by an EnergyPlus model rather than a temperature difference proxy as described above. With two HVAC zones and several electric lighting regions it is an inherently more complex area and thus EnergyPlus offers a flexible and representative (to be discussed) model of the energy consumption of the space. The EnergyPlus model for the CREST lab involves constructing a model that mimics the physical space illustrated in Figure 4.3 below. There are eight independently controlled overhead electric lighting systems above the desk areas. Additionally, In order to model the HVAC system, we use the Ideal Loads Air System in EnergyPlus since it is not connected to a central air system and instead supplies heating or cooling to a zone to meet the load. These two controllable HVAC zones are modeled as ideal VAV terminal units with variable supply temperature and humidity on the North and West areas of the lab. As seen in Figure 4.3, desk areas 1 and 2 belong to the North zone, areas 4, 5, and 6 to the West zone, and area 3 is on the boundary between the two. The zones freely mix in the open desk area in the lab. The model of the lab also includes the 3 large windows controllable with blind systems and is updated with the most recent occupancy measurements during simulation, as these add loads to the HVAC system that must be accounted for when computing the energy costs.

The input to the EnergyPlus model is then given as $(S_Z^{t:t+16}, L_u^{t:t+16}, \theta_s^{t:t+16})$, and the values for the setpoints, lighting levels, and shade levels are filled into the compact schedules for each HVAC and lighting zone.

A key advantage of using EnergyPlus to model the space is the ability to avoid installing expensive metering equipment for the HVAC system; however, we must ensure that the energy expenditure values given by the model are a reasonable proxy for the real costs. Prior to closer verification, the EnergyPlus predictions for the temperature dynamics of the space were wildly inaccurate. However, upon testing several parameters of the HVAC model used, it was found that setting the maximum heating supply air temperature ($30^{\circ}C$), the minimum cooling supply air temperature ($10^{\circ}C$), and the maximum heating/cooling airflow rate (to $0.377m^3/s$) were the most effective in bringing the predicted EnergyPlus temperature dynamics closer to those of the actual space (as opposed to using the default values for the supply air temperatures and setting no limit upon the airflow rate). This can be seen in Figure 4.4. Even with the adjustments to the EnergyPlus model to more closely mimic the HVAC system parameters, the transitions in temperature are quite sharp in response to control inputs. It may be possible to smoothen this behaviour, but could come at the expense of setting HVAC parameters that are too different from the real settings such that the energy calculations are no longer representative.

4.2.3 Temperature and Brightness Models

In order to model the temperature dynamics of the CREST lab, similar data collection methods as described for the conference room were utilized. A range of setpoints were tested for both the north zone and the west zone, as well as a range of lighting levels for all eight of the overhead electric lighting units. Thus, the temperature dynamics of each zone are given by $T_{k+1}^z = f(T_k^1, \ldots, T_k^6, S_k^z, L_k^z, \theta_k^z, O_k^z)$ where f is a linear function.

In addition, the lighting data collection was undertaken at different times of the day with varying levels of window shading in order to construct a more complete model of the illumination of the space. The brightness of each zone is given by $B_{k+1}^z = g(W_k, L_k^{1:8}, \theta_k^{1:3})$ where g is a multi-layer perceptron; Figure 4.5 are predictions versus real measurements of the brightness in two zones over the course of several hours.



Figure 4.4: EnergyPlus Temperature Simulation



Figure 4.5: Brightness Model Examples

5 Results

Experiments to verify the efficacy of the model-predictive controller were conducted using two strategies. The CREST lab itself implements a simple baseline strategy that controls the HVAC, lights, and blinds. It is a rule-based approach that sets the setpoint to the desired value, raises or lowers the shades depending on the time of day, and always fully utilizes the electric lighting. The second set of experiments uses the MPC setup described earlier. Both experiments were conducted over periods of about 4 hours at similar times of the day during weekdays. As given in the legend, the orange dashed line represents the upper comfort range of the brightness/temperature as given above, and the green line is the corresponding lower comfort range. As the desk area becomes occupied or unoccupied, the constraints on the temperature and brightness for the MPC change accordingly.

5.1 Energy Expenditure

For the control actions (supply air temperature, electric lighting levels, blind levels) given by both of the strategies above, Figure 5.1 summarizes the total energy expenditure (kJ) over each experimentation period.

5.2 Luminescence

Figure 5.2 presents the results of the brightness levels in the space using the optimal controller as compared to the baseline controller.

5.3 Temperature

Figure 5.3 presents the results of the temperature levels in the space using the optimal controller as compared to the baseline controller, and also gives the model predictions for the temperature.



(a) HVAC Energy Expenditure



Total Lighting Energy Expenditure (kWh)

(b) Lighting Energy Expenditure

Figure 5.1: Experimental Energy Expenditures over 4 Hour Experimentation Period



Figure 5.2: Brightness Controller versus Baseline Strategy



Figure 5.2: Brightness Controller versus Baseline Strategy (cont.)



Figure 5.3: Temperature Controller versus Baseline Strategy



Figure 5.3: Temperature Controller versus Baseline Strategy (cont.)

6 Discussion and Future Work

Upon inspection of the experimental results above, we can see that the model-predictive controller is quite effective in maintaining thermal and visual comfort while minimizing overall energy expenditure.

The baseline strategy was effective in maintaining the luminescence within the comfort rangeits primary strategy was a rule-based control of the shades with little control of the electric lighting levels. It was able to maintain the comfort levels of brightness 89.8% of the time during the experimentation period. The model-predictive controller was also effective but did not offer much of an advantage over the baseline from the standpoint of enforcing brightness constraints. It maintained the comfort levels of brightness 81.2% of the time during the experimentation period (it did offer significant energy consumption savings to be discussed further). However, the baseline strategy for the HVAC system was quite ineffective in maintaining the temperature within the comfort ranges. It essentially fixed a setpoint without regard for occupancy, solar radiation, and other significant factors, resulting in an environment that did not align with the varying temperature constraints. This resulted in the baseline strategy maintaining thermal comfort levels for just 9.9% of the time during the experimentation period. The model-predictive controller was much more effective in this regard. Thermal comfort was maintained for 82.4% of the experimentation period and the temperature was always within the comfort zone for nearly all of the desk areas except for desk area 1 (likely due to the mixing of the two different HVAC zones at this area which resulted in one system overpowering the other).

A very interesting result is that the model-predictive controller was quite effective from an energy conservation standpoint. This is primarily due to the savings offered by the optimal lighting strategy. While the baseline strategy maintained visual comfort, it did not take full advantage of the free day-lighting offered by the 3 large windows and instead controlled the blinds on a set schedule. The model predictive controller was able to turn off several of the eight available overhead lights during much of the experimentation period, instead opting to raise the blinds to achieve the same lighting levels. Paired with the fact that the temperature models for each zone accounted for solar radiation from the windows and potentially offered some mild advantages in HVAC savings, the energy consumption reduction from using significantly fewer of the overhead

lights in favor of day-lighting resulted in much lower overall expenditure over the experimentation period.

There are several interesting extensions to this work with regards to predicting occupancy patterns, utilizing occupant preferences, and privacy. A reinforcement learning approach to learn the preferred thermal and brightness ranges for the occupants of the space could be helpful in personalizing the environment and improving comfort. Additionally, there has been extensive work in predicting occupancy trends and incorporating this into the model-predictive control loop that could improve its efficacy. Moreover, while we used the privacy-preserving occupancy scheme in the conference room, we did not do so for the larger lab implementation. Since the latter uses significantly more occupancy information it could benefit from the scheme described and implemented earlier.

In this work we present a lightweight and flexible model-predictive control scheme to jointly control the HVAC, electric lighting, and blinds system in a space to make it more energy efficient. By combining these different systems into one optimization problem we were able to achieve improved thermal and visual comfort while decreasing energy consumption in a real-world implementation. This work can serve as a platform for future research in enabling robust and effective retrofits towards the development of smart and energy-efficient buildings.

Bibliography

- 1. A. Sieminski. "International energy outlook 2013". In: US Energy Information Administration (EIA) Report Number: DOE/EIA-0484, 2013.
- 2. J. Kleissl and Y. Agarwal. "Cyber-physical energy systems: Focus on smart buildings". In: *Design Automation Conference*. IEEE. 2010, pp. 749–754.
- S. Dawson-Haggerty, A. Krioukov, J. Taneja, S. Karandikar, G. Fierro, N. Kitaev, and D. Culler. "{BOSS}: Building operating system services". In: Presented as part of the 10th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 13). 2013, pp. 443–457.
- 4. K. Weekly, M. Jin, H. Zou, C. Hsu, A. Bayen, and C. Spanos. "Building-in-briefcase (bib)". In: *arXiv* preprint arXiv:1409.1660, 2014.
- 5. A. Tzempelikos and A. K. Athienitis. "The impact of shading design and control on building cooling and lighting demand". In: *Solar energy* 81:3, 2007, pp. 369–382.
- 6. G. Van Moeseke, I. Bruyère, and A. De Herde. "Impact of control rules on the efficiency of shading devices and free cooling for office buildings". In: *Building and environment* 42:2, 2007, pp. 784–793.
- A. Yoshida, J. Yoshikawa, Y. Fujimoto, Y. Amano, and Y. Hayashi. "Stochastic receding horizon control minimizing mean-variance with demand forecasting for home EMSs". In: *Energy and Buildings* 158, 2018, pp. 1632–1639.
- 8. 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* 23:1, 2015, pp. 101–116.
- R. Yang and L. Wang. "Optimal control strategy for HVAC system in building energy management". In: *PES T&D 2012*. IEEE. 2012, pp. 1–8.
- 10. M. M. Haghighi and A. L. Sangiovanni-Vincentelli. "Modeling and optimal control algorithm design for HVAC systems in energy efficient buildings". In: *Technical Report.* Citeseer, 2011.
- 11. N. Radhakrishnan, Y. Su, R. Su, and K. Poolla. "Token based scheduling for energy management in building HVAC systems". In: *Applied energy* 173, 2016, pp. 67–79.
- 12. T. Wei, Y. Wang, and Q. Zhu. "Deep reinforcement learning for building HVAC control". In: *Proceedings* of the 54th Annual Design Automation Conference 2017. ACM. 2017, p. 22.

- 13. Z. Zhang, A. Chong, Y. Pan, C. Zhang, S. Lu, and K. P. Lam. "A deep reinforcement learning approach to using whole building energy model for hvac optimal control". In: *2018 Building Performance Analysis Conference and SimBuild*. 2018.
- H. Kazmi, F. Mehmood, S. Lodeweyckx, and J. Driesen. "Gigawatt-hour scale savings on a budget of zero: Deep reinforcement learning based optimal control of hot water systems". In: *Energy* 144, 2018, pp. 159–168.
- 15. J. Zhao, K. P. Lam, and B. E. Ydstie. "EnergyPlus model-based predictive control (EPMPC) by using MATLAB/SIMULINK and MLE+". In: 2013.
- 16. J. Zhao, K. P. Lam, B. E. Ydstie, and O. T. Karaguzel. "EnergyPlus model-based predictive control within design–build–operate energy information modelling infrastructure". In: *Journal of Building Performance Simulation* 8:3, 2015, pp. 121–134.
- B. Sun, P. B. Luh, Q.-S. Jia, Z. Jiang, F. Wang, and C. Song. "Building energy management: Integrated control of active and passive heating, cooling, lighting, shading, and ventilation systems". In: *IEEE Transactions on automation science and engineering* 10:3, 2013, pp. 588–602.
- R. Jia, R. Dong, S. S. Sastry, and C. J. Sapnos. "Privacy-enhanced architecture for occupancy-based HVAC control". In: 2017 ACM/IEEE 8th international conference on cyber-physical systems (ICCPS). IEEE. 2017, pp. 177–186.