

Real-World Implementation of Reinforcement Learning Based Energy Coordination for a Cluster of Households

Gargya Gokhale

IDLab, Ghent University – imec
Ghent, Belgium

Niels Tiben

IDLab, Ghent University – imec
Ghent, Belgium

Marie-Sophie Verwee

IDLab, Ghent University – imec
Ghent, Belgium

Manu Lahariya

IDLab, Ghent University – imec
Ghent, Belgium

Bert Claessens

IDLab, Ghent University – imec
Beebop.ai
Belgium

Chris Develder

IDLab, Ghent University – imec
Ghent, Belgium

ABSTRACT

Given its substantial contribution of 40% to global power consumption, the built environment has received increasing attention to serve as a source of flexibility to assist the modern power grid. In that respect, previous research mainly focused on energy management of individual buildings. In contrast, in this paper, we focus on aggregated control of a set of residential buildings, to provide grid supporting services, that eventually should include ancillary services. In particular, we present a real-life pilot study that studies the effectiveness of reinforcement-learning (RL) in coordinating the power consumption of 8 residential buildings to jointly track a target power signal. Our RL approach relies solely on observed data from individual households and does not require any explicit building models or simulators, making it practical to implement and easy to scale. We show the feasibility of our proposed RL-based coordination strategy in a real-world setting. In a 4-week case study, we demonstrate a hierarchical control system, relying on an RL-based ranking system to select which households to activate flex assets from, and a real-time PI control-based power dispatch mechanism to control the selected assets. Our results demonstrate satisfactory power tracking, and the effectiveness of the RL-based ranks which are learnt in a purely data-driven manner.

CCS CONCEPTS

• **Hardware** → **Smart grid**; • **Computing methodologies** → **Cooperation and coordination**; **Reinforcement learning**.

KEYWORDS

Demand Response, Reinforcement Learning, Building Cluster, Coordination, Advantage function

ACM Reference Format:

Gargya Gokhale, Niels Tiben, Marie-Sophie Verwee, Manu Lahariya, Bert Claessens, and Chris Develder. 2023. Real-World Implementation of Reinforcement Learning Based Energy Coordination for a Cluster of Households.

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. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

BuildSys '23, November 15–16, 2023, Istanbul, Turkey

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0230-3/23/11...\$15.00

<https://doi.org/10.1145/3600100.3625681>

In *The 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '23)*, November 15–16, 2023, Istanbul, Turkey. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3600100.3625681>

1 INTRODUCTION

The past decades have witnessed a steady increase in the integration of renewable energy sources (RES) into the power grid, a trend that is expected to continue in the coming years. This growth in RES deployment implies a rising need for energy flexibility to ensure reliable and secure grid operation. Conventionally, energy flexibility in power grids was obtained using fast ramping energy sources such as gas-powered generators. However, with growing emphasis on clean, CO₂-free energy, demand-side flexibility has been gaining significant attention, e.g., grid level batteries, industrial processes. In this respect, the built environment holds a significant potential as an as of yet (largely) untapped source of flexibility. Indeed, buildings account for over 40% of global energy consumption [7], and offer ample opportunities to manage their consumption flexibly. The latter typically entails using their intrinsic thermal mass as a passive storage element and shifting their consumption patterns whilst satisfying necessary comfort levels [25].

While buildings constitute a huge potential for contributing towards demand side flexibility, realizing a *control framework* capable of exploiting this flexibility is a highly challenging problem [13]. This is particularly due to the non-linear dynamics of the building and the stochastic operating conditions related to their dependence on outside temperatures, human behavior, solar irradiation, etc. Clearly, the challenge is exacerbated when extended to coordination of a group of buildings or an entire neighborhood. Consequently, most works in the well-researched area of thermal and HVAC control for buildings (see [10, 24] for recent overviews) focus on cost or power consumption minimization for a *single building*.

Those works adopt advanced control techniques to develop building-level controllers, with *Model Predictive Control (MPC)* [8] being a predominantly used strategy. Works such as [6, 12, 21] have demonstrated the deployment of MPC-based control in large commercial buildings and have shown significant improvement in energy flexibility of these buildings as compared to business-as-usual controllers. Such MPC solutions require a good quality model of the building(s) at hand, which is challenging in practice. This is one of the reasons for limiting the control to a single, typically

commercial¹ building rather than, e.g., a cluster. Given that such lack of scalability in MPC-based control forms a major bottleneck, recent research is moving towards data-driven control paradigms such as Reinforcement Learning (RL).

Being data-driven, *RL-based control* relies on interacting with the building and iteratively improving the control policy. Spurred by advances in deep learning and compute capabilities, prior works have implemented different RL-based control strategies for optimization in individual buildings [1, 26], a group of buildings [4, 15] and for coordination of a cluster buildings [16, 19, 23]. While these works show the performance benefits of using RL-based controllers for different buildings and optimization objectives, most works are limited to simulation environments such as BOPTTEST, CityLearn [2, 20] or utilize a high-fidelity building model based on EnergyPlus or Modelica [5, 22]. Additionally, prior works on grid interactive buildings such as [16, 19, 23], primarily consider explicit storage assets such as electric vehicles (EVs), batteries or hot water storage tanks for flexibility. While effective, such storage assets need external installation and carry an investment cost. However, utilizing passive thermal storage in buildings (building thermal mass) [17] presents an (almost) cost-free, thermal storage alternative, a flexibility element that has often been overlooked in past works related to building coordination.

Our pilot study aims to contribute to this lesser-researched area of energy coordination in residential buildings, utilizing only their intrinsic thermal mass (i.e., without any explicit storage element). To the best of our knowledge, this work is among the first works to present a real-world RL-based deployment of an energy coordination framework for a group of households without using any prior information or simulator models for these buildings. The main contributions of our case study on controlling a cluster of residential housing units (§2) are that we:

- Implement a real-world proof-of-concept demonstration of a scalable, data-driven RL-based approach (§3) for coordinating a group of residential buildings;
- Demonstrate the power tracking capability of the group of residential buildings for providing short-term DR services (§4);
- Discuss the challenges encountered during this pilot study and present future research directions for unlocking DR-potential from residential buildings (§5).

2 PROBLEM STATEMENT

2.1 Case Study

The pilot study undertaken in this work relates to a group of residential apartments whose energy consumption is managed by a demand aggregator. Following works such as [12], we assume that the aggregator offers a power tracking service (similar to secondary frequency control) over a short time interval (sub-hour), when required. During normal operation (when no DR service is needed), the aggregator is responsible for maintaining the indoor room temperature in each household within the user defined comfort bounds. These comfort bounds refer to a narrow band of $\pm 1^\circ\text{C}$ around a target room temperature that is set by each user. This

normal operation represents the business-as-usual (BAU) scenario and the aggregator aims to maximize the comfort of the users by maintaining the room temperature as close to the target setpoint as possible.

However, in case of a DR event, the aggregator must switch from this BAU scenario and follow the required power signal requested by the system operator. To do so, the aggregator must adjust the power consumption of different households, to ensure that the aggregate power consumption of all households together follows the requested tracking signal. This entails that some of the houses need to deviate from their BAU scenario — either being heated slightly above the setpoint (power upregulation) or being allowed to cool down slightly below the requested setpoint (power downregulation) — whilst ensuring user comfort ($\pm 1^\circ\text{C}$ around the target temperature setpoint). Thus, during a DR event, the aggregator needs to identify suitable (flexible) households that must deviate from the BAU scenario and by how much.

2.2 Problem Formulation

This section formulates the energy coordination problem from the perspective of the demand aggregator. First, we model each household as a partially observable Markov Decision Process (MDP). MDP is a commonly used framework for modelling sequential decision-making problems and is characterized by a state space \mathbf{X} , an action space \mathbf{U} , a state transition function f and a cost function ρ [18]. In our case, the *state* for house h is defined as $\mathbf{x}_t^h = \{t, T_{o,t}^h, T_{r,t-k}^h, \dots, T_{r,t}^h\}$, i.e., the time-of-day t , the past k room temperature values, the current room temperature $T_{r,t}^h$ and outside air temperature $T_{o,t}^h$. An *action* u_t^h is the heating power (ON or OFF), and the *cost* function is the amount of energy consumed ($G_t^h(\mathbf{x}_t^h, u_t^h)$). The transition function ($f(\mathbf{x}_t^h, u_t^h)$) relates to the natural evolution of room temperature over time (duration of each time step is 15 minutes). Note that the actual system cannot be observed fully, with essential quantities such as building thermal mass being hidden, making it a partially observable MDP. To mitigate this lack of complete information, we use a (short) history of past room temperature values over the last k time steps [18]. For the BAU scenario, the control policy (π_b^h) used for maintaining the indoor temperature within user comfort levels is given by:

$$\pi_b^h(\mathbf{x}_t^h) = \begin{cases} 0 & : T_{r,t}^h > T_{r,t}^{h,\text{set}} \\ 1 & : T_{r,t}^h \leq T_{r,t}^{h,\text{set}} \end{cases} \quad (1)$$

where $T_{r,t}^{h,\text{set}}$ denotes the user defined target temperature for time step t . We define the aggregate energy consumption (G) of the entire portfolio of households during time step t as:

$$G(\mathbf{x}_t, \mathbf{u}_t) = \sum_{h=1}^N G_t^h(\mathbf{x}_t^h, u_t^h) \quad (2)$$

where, $\mathbf{x}_t, \mathbf{u}_t$ represent the collection of current states and actions of all the houses. During BAU scenario, individual actions u_t^h are obtained from the respective BAU policies as defined in Eq. (1).

However, during a DR event, we assume that the system operator conveys a tracking signal (κ_t) at the beginning of the DR event (t_d), for a fixed duration τ and the aggregator seeks to match

¹Since for that scale, constructing/obtaining a sufficiently accurate model is affordable.

it by deviating from BAU scenario. This is formulated as a bi-level optimization problem given by:

$$\begin{aligned} & \min_{\mathbf{u}_{t_d}, \dots, \mathbf{u}_{t_d+\tau}} \sum_{t=t_d}^{t_d+\tau} |\kappa_t - G(\mathbf{x}_t, \mathbf{u}_t)| \\ \text{s.t. } & \mathbf{u}^h \in \arg \min_{\mathbf{u}_t^h, \dots, \mathbf{u}_T^h} \left[\sum_{t=1}^T G_t^h(\mathbf{x}_t^h, \mathbf{u}_t^h) - G_t^h(\mathbf{x}_t^h, \pi_b^h(\mathbf{x}_t^h)) \right] \quad \forall h \quad (3) \\ & T_{r,t}^{h,\text{set}} - 1 \leq T_{r,t}^h \leq T_{r,t}^{h,\text{set}} + 1 \quad \forall h, t \end{aligned}$$

3 METHODOLOGY

This section describes the proposed RL-based coordination strategy to solve the problem described in §2. The demand aggregator follows a business-as-usual strategy (described by Eq. (1)) during normal operation. In case of a DR event, as discussed in §2, the aggregator needs to identify suitable households that can deviate from their BAU policy and dispatch them appropriately so that the aggregate power consumption follows the requested tracking signal. Closely following the work presented in [14], we model the aggregator's control strategy using a hierarchical control architecture comprising an RL-based ranker for identifying suitable households and a PI control based real-time dispatcher for dispatching houses based on the ranker's output. The real-time dispatcher sends out actions every minute, contrary to the 15 minute interval used during normal operation.

3.1 RL-based Ranking

As discussed in §2, the aggregator must choose households such that the energy deviation of each household due to the DR event is minimized. For this, the RL-based ranker uses the energy consumption of each household to rank different households and identify suitable houses during a DR-event. With each house modelled as an (episodic) MDP, the cumulative energy consumption over a day for a household following control policy π_b^h can be defined using the Q -function (or state-action value function) as:

$$\begin{aligned} Q^h(\mathbf{x}_t^h, u_t^h) &= G_t^h(\mathbf{x}_t^h, u_t^h) + \sum_{k=1}^{k=T} G_t^h(\mathbf{x}_{t+k}^h, \pi_b^h(\mathbf{x}_{t+k}^h)) \quad (4) \\ Q^h(\mathbf{x}_t^h, u_t^h) &= G_t^h(\mathbf{x}_t^h, u_t^h) + Q^h(\mathbf{x}_{t+1}^h, \pi_b^h(\mathbf{x}_{t+1}^h)) \end{aligned}$$

This Q -function is then used to obtain an advantage function (A^h) defined as:

$$A^h(\mathbf{x}_t^h, u_t^h) = Q^h(\mathbf{x}_t^h, u_t^h) - V^h(\mathbf{x}_t^h) \quad (5)$$

where V^h denotes the value function and quantifies the value of being in state \mathbf{x}_t^h . As discussed in [14], this advantage function quantifies the effect of deviating from the BAU policy for the given state \mathbf{x}_t^h . Combining this with Eq. (4), we observe that, for our case, the advantage function approximates the deviation in cumulative energy consumption if an action u_t^h is chosen instead of the business-as-usual action ($\pi_b^h(\mathbf{x}_t^h)$). Thus, for each household, the advantage value ($A^h(\mathbf{x}_t^h, u_t^h)$) provides an estimate of the "cost" of deviating from the BAU policy and can be used as a method for ranking different households.

To obtain these advantage functions and hence the ranks, we first need to approximate the Q -function for each household. This represents a standard RL problem of policy evaluation [18] and is resolved by training Q -networks for each household following the fitted Q -iteration algorithm [9]. The Q -networks for each household are trained on a batch of previous 30 days of data (~2.8k samples) and are retrained at the beginning of each day to include the latest data.

3.2 Real-time Dispatcher

During a DR event, the real-time dispatcher is responsible for following the system operator tracking signal as closely as possible. For this, the dispatcher carries out three basic steps: (1) Collect advantage functions for the current states of each household; (2) Create a ranking hash table using these advantage values; and (3) Dispatch control actions for each household using this ranking table and a PI controller. To ensure good signal tracking, the dispatcher sends out actions every minute, dispatching houses with the lowest rank first and gradually increasing the number of dispatched houses based on the PI controller. Additionally, to ensure that thermal comfort of users is always maintained, the dispatcher actively filters control actions, overriding actions that can lead to comfort violations (i.e., room temperature exceeding the $\pm 1^\circ\text{C}$ comfort band) for individual households.

4 RESULTS AND DISCUSSION

In this section, we describe the pilot study implemented and present the results obtained. The pilot study was conducted in collaboration with anonymous company. This proof-of-concept study comprised 8 residential apartments that were controlled over a period of 4 weeks (20 weekdays) during February this year (2023). The software for the control framework described in §3 was implemented in Python and hosted on an in-premise Kubernetes cluster, interacting with the company API for sending control actions and receiving households' states.

The main objective of this pilot study was to demonstrate a scalable, data-driven RL-based coordination mechanism for residential households without using external simulators or building models. This includes answering two main questions: (Q1) How closely can the aggregate energy consumption of the cluster follow the tracking signal? (Q2) Does the RL-based ranking system learn any useful patterns to select households as compared to random selection? The following results aim to address these two questions.

4.1 Overall Response

To address Q1 regarding the signal tracking capacity of our control framework, we set a square wave DR signal of 40 minutes duration (twice per day) and observed the behavior of the system over 10 days.

Figure 1 presents consolidated results over the 10 days. It can be observed that the median value of aggregate power of the entire cluster is able to follow the trend of the square wave tracking signal. However, it falls short of accurately following the signal, especially during the second peak. This is due to the lack of flexibility available during the DR event and can be largely contributed to the small number of households (8) available in this case study. However,




Figure 1: Consolidated performance of DR events. The solid lines represent the median values for observed aggregated power and the shaded area represents the standard deviation across all DR events performed. The target values represent the DR signals provided to the aggregator for tracking.

while the small number of households limit the accuracy, the overall trend demonstrates that our control framework can steer the aggregate energy consumption of the cluster to satisfactorily follow an external tracking signal.

4.2 Quality of RL-based Ranking

To analyze what patterns the RL-based ranking system has learnt (i.e., Q2), we conducted a smaller experiment using only 2 households. The idea was to observe the relevance of ranks obtained in this reduced size test and provide an intuition about the RL-based based ranks. The underlying advantage functions for both households are visualized as a heatmap in Fig. 2. The heatmap presents the difference between the advantage values for switching ON ($u_t = 1$) and switching OFF ($u_t = 0$) for a given state (\mathbf{x}_t). The time and room temperature for state (\mathbf{x}_t) are depicted by the grey box, which is the exact time when a DR event was activated. Through this visualization, we observe that the heatmap for house 2 shows a clear pattern that is strongly correlated to the user-defined room temperature setpoint. Specifically, the Q-network estimates that ‘switching ON’ at any time before 15h is not advantageous and is depicted by darker shades of blue. Similarly, around 17h, the Q-network for house 2 expects a change in user-defined temperature setpoint and hence values switching ON almost equal to or sometimes lower than switching OFF, indicated by the warmer shades in the heatmap. This trend corresponds to the regular pattern of room temperature setpoints observed for this household. Contrary to this, for house 42, we observe a significantly different trend, where the Q-network learns an almost flat trend of always estimating a higher cost for switching ON (compared to switching OFF) unless the room temperature falls below a certain value (20°C) corresponding to the mean value of past room temperature setpoint. This trend also correlates to a relatively flat room temperature setpoints set by users of this house. Note that the room temperature setpoint is a user-defined quantity which is not used as input for training the Q-network. The trends in Fig. 2 indicate that the Q-networks are




Figure 2: Heatmaps depicting the ‘dis’-advantage of switching ON over switching OFF action (i.e., $A^h(\mathbf{x}_t^h, 1) - A^h(\mathbf{x}_t^h, 0)$). Blue regions indicate that the Q-network believes it is better to remain switched off or switching ON is expected to be more costly (in terms of energy consumption). Conversely, red regions imply a big advantage to switching ON. The dashed line shows the user-defined temperature setpoint for the given day and the grey box denotes the state when the DR event corresponding to Table 1 occurred.

Table 1: Ranks assigned during a two-household upwards DR event. The action values represent non-BAU control actions.

	Action	Advantage	Rank
House 2	ON	0.0963	1
House 42	ON	0.3460	2

able to learn these user-specific patterns and assign ranks based on such patterns.

Further, Table 1 shows the ranks for both these houses for an upward DR event, requiring an increase in power consumption, occurring at around 15h. Using Fig. 2, these ranks can be justified since house 2 is expected to switch ON heating in any case and hence it is ranked lower than the household that is expected to keep their heating OFF for the foreseeable future. Based on these ranks, the real-time dispatcher will first switch ON heating for House 2, and if insufficient, it will move forward to house 42. These results indicate that the RL-based ranking system learns useful patterns related to the user-behavior and assigns ranks that are intuitive and better than a random ranking strategy.

5 CONCLUSION

This work presented a real-world, proof-of-concept demonstration of a scalable, data-driven RL-based control strategy for coordinating the power consumption of a group of residential buildings to provide demand response services. As presented in §4, our proposed coordination strategy shows good power tracking capacity with the RL-based ranker producing intuitive and reasonable ranks by solely relying on observed data from different households. The aim of this work was to deliver a proof-of-concept feasibility study and hence we focused on a relatively small set of households. While

this small size had practical benefits including ease of deployment, it has led to a few limitations that we aim to address in future work.

5.1 Limitations and Future Work

As discussed earlier, our current work presented a pilot study of a relatively small size (in terms of both the number of households as well as the deployment period). While this had practical benefits, future work would involve expanding the size of this study by adding more households and testing the control strategy over an entire winter season. This would allow testing the proposed coordination strategy for tracking arbitrary signals and analyzing the commercial potential of this technique. Besides the size of the study, we have identified two key areas for improvement. The first relates to the need for a validation and backtesting framework for the trained Q -functions and advantage functions. Future work will research policy evaluation methods [3] and data-driven, physics informed simulators [11] for data-driven backtesting strategies. The second improvement area involves using transfer learning-based methods to leverage data from different households to pre-train base models that can be finetuned in a data-efficient way to obtain the necessary Q -functions. This will ensure robust training performance and allow us to use advanced deep learning architectures such as transformers that can improve the overall accuracy and scalability of the learnt Q -functions.

ACKNOWLEDGMENTS

This pilot study was performed in collaboration with domX — a Greek IoT-based energy management company — and received funding through the European Union’s Horizon 2020 Projects BiGG (grant no. 957047) and BRIGHT (grant no. 957816). We thank Polychronis Symeonidis, Stratos Kerandis and the entire domX team for their support throughout this study. We also thank Matthias Strobbe for his support during this study.

REFERENCES

- [1] Donald Azuatalam, Wee-Lih Lee, Frits de Nijs, and Ariel Liebman. 2020. Reinforcement learning for whole-building HVAC control and demand response. *Energy and AI* 2 (Nov. 2020), 100020. <https://doi.org/10.1016/j.egyai.2020.100020>
- [2] David Blum, Javier Arroyo, Sen Huang, Ján Drgoňa, Filip Jorissen, Harald Taxt Walnum, Yan Chen, Kyle Benne, Draguna Vrabie, Michael Wetter, et al. 2021. Building optimization testing framework (BOPTTEST) for simulation-based benchmarking of control strategies in buildings. *Journal of Building Performance Simulation* 14, 5 (2021), 586–610. <https://doi.org/10.1080/19401493.2021.1986574>
- [3] Bingqing Chen, Ming Jin, Zhe Wang, Tianzhen Hong, and Mario Bergés. 2020. Towards off-policy evaluation as a prerequisite for real-world reinforcement learning in building control. In *Proceedings of the 1st International Workshop on Reinforcement Learning for Energy Management in Buildings & Cities*. 52–56.
- [4] Davide Coraci, Silvio Brandi, Tianzhen Hong, and Alfonso Capozzoli. 2023. Online transfer learning strategy for enhancing the scalability and deployment of deep reinforcement learning control in smart buildings. *Applied Energy* 333 (March 2023), 120598. <https://doi.org/10.1016/j.apenergy.2022.120598>
- [5] Drury B Crawley, Linda K Lawrie, Frederick C Winkelmann, Walter F Buhl, Y Joe Huang, Curtis O Pedersen, Richard K Strand, Richard J Liesen, Daniel E Fisher, Michael J Witte, et al. 2001. EnergyPlus: creating a new-generation building energy simulation program. *Energy and buildings* 33, 4 (2001), 319–331. [https://doi.org/10.1016/S0378-7788\(00\)00114-6](https://doi.org/10.1016/S0378-7788(00)00114-6)
- [6] Roel De Coninck and Lieve Helsen. 2016. Practical implementation and evaluation of model predictive control for an office building in Brussels. *Energy and Buildings* 111 (Jan. 2016), 290–298. <https://doi.org/10.1016/j.enbuild.2015.11.014>
- [7] U Doe. 2015. Chapter 5: Increasing efficiency of building systems and technologies. *Quadrennial Technology Review: An Assessment of Energy Technologies and Research Opportunities* (2015), 143–181.
- [8] Ján Drgoňa, Javier Arroyo, Iago Cupeiro Figueroa, David Blum, Krzysztof Arendt, Donghun Kim, Enric Perarnau Ollé, Juraj Oravec, Michael Wetter, Draguna L. Vrabie, and Lieve Helsen. 2020. All you need to know about model predictive control for buildings. *Annual Reviews in Control* 50 (2020), 190–232. <https://doi.org/10.1016/j.arcontrol.2020.09.001>
- [9] Damien Ernst, Pierre Geurts, and Louis Wehenkel. 2005. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research* 6 (2005).
- [10] Qiming Fu, Zhicong Han, Jianping Chen, You Lu, Hongjie Wu, and Yunzhe Wang. 2022. Applications of reinforcement learning for building energy efficiency control: A review. *Journal of Building Engineering* 50 (June 2022), 104165. <https://doi.org/10.1016/j.jobe.2022.104165>
- [11] Gargya Gokhale, Bert Claessens, and Chris Develder. 2022. Physics informed neural networks for control oriented thermal modeling of buildings. *Applied Energy* 314 (2022), 118852. <https://doi.org/10.1016/j.apenergy.2022.118852>
- [12] Tomasz T. Gorecki, Luca Fabietti, Faran A. Qureshi, and Colin N. Jones. 2017. Experimental demonstration of buildings providing frequency regulation services in the Swiss market. *Energy and Buildings* 144 (June 2017), 229–240. <https://doi.org/10.1016/j.enbuild.2017.02.050>
- [13] Rongling Li, Andrew J. Satchwell, Donal Finn, Toke Haunstrup Christensen, Michaël Kummert, Jérôme Le Dréau, Rui Amaral Lopes, Henrik Madsen, Jaume Salom, Gregor Henze, and Kim Wittchen. 2022. Ten questions concerning energy flexibility in buildings. *Building and Environment* 223 (Sept. 2022), 109461. <https://doi.org/10.1016/j.buildenv.2022.109461>
- [14] Mingxi Liu, Stef Peeters, Duncan S. Callaway, and Bert J. Claessens. 2019. Trajectory Tracking With an Aggregation of Domestic Hot Water Heaters: Combining Model-Based and Model-Free Control in a Commercial Deployment. *IEEE Transactions on Smart Grid* 10, 5 (Sept. 2019), 5686–5695. <https://doi.org/10.1109/TSG.2018.2890275>
- [15] Kingsley Nweye, Siva Sankaranarayanan, and Zoltan Nagy. 2023. MERLIN: Multi-agent offline and transfer learning for occupant-centric operation of grid-interactive communities. *Applied Energy* 346 (2023), 121323. <https://doi.org/10.1016/j.apenergy.2023.121323>
- [16] Giuseppe Pinto, Marco Savino Piscitelli, José Ramón Vázquez-Canteli, Zoltán Nagy, and Alfonso Capozzoli. 2021. Coordinated energy management for a cluster of buildings through deep reinforcement learning. *Energy* 229 (2021), 120725. <https://doi.org/10.1016/j.energy.2021.120725>
- [17] Glenn Reynders, Rui Amaral Lopes, Anna Marszal-Pomianowska, Daniel Aelenei, João Martins, and Dirk Saelens. 2018. Energy flexible buildings: An evaluation of definitions and quantification methodologies applied to thermal storage. *Energy and Buildings* 166 (2018), 372–390. <https://doi.org/10.1016/j.enbuild.2018.02.040>
- [18] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [19] Jose R Vazquez-Canteli, Gregor Henze, and Zoltan Nagy. 2020. MARLISA: Multi-agent reinforcement learning with iterative sequential action selection for load shaping of grid-interactive connected buildings. In *Proceedings of the 7th ACM international conference on systems for energy-efficient buildings, cities, and transportation*. 170–179. <https://doi.org/10.1145/3408308.3427604>
- [20] José R. Vázquez-Canteli, Jérôme Kämpf, Gregor Henze, and Zoltan Nagy. 2019. CityLearn v1.0: An OpenAI Gym Environment for Demand Response with Deep Reinforcement Learning. In *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation* (New York, NY, USA) (*BuildSys '19*). Association for Computing Machinery, New York, NY, USA, 356–357. <https://doi.org/10.1145/3360322.3360998>
- [21] Dan Wang, Yangzhe Chen, Wei Wang, Cheng Gao, and Zhe Wang. 2023. Field test of Model Predictive Control in residential buildings for utility cost savings. *Energy and Buildings* 288 (June 2023), 113026. <https://doi.org/10.1016/j.enbuild.2023.113026>
- [22] Michael Wetter, Wangda Zuo, Thierry S Nouidui, and Xiufeng Pang. 2014. Mod-lica buildings library. *Journal of Building Performance Simulation* 7, 4 (2014), 253–270. <https://doi.org/10.1080/19401493.2013.765506>
- [23] Jiahua Xie, Akshay Ajagekar, and Fengqi You. 2023. Multi-Agent attention-based deep reinforcement learning for demand response in grid-responsive buildings. *Applied Energy* 342 (July 2023), 121162. <https://doi.org/10.1016/j.apenergy.2023.121162>
- [24] Ye Yao and Divyanshu Kumar Shekhar. 2021. State of the art review on model predictive control (MPC) in Heating Ventilation and Air-conditioning (HVAC) field. *Building and Environment* 200 (Aug. 2021), 107952. <https://doi.org/10.1016/j.buildenv.2021.107952>
- [25] Huiliang Zhang, Sayani Seal, Di Wu, François Bouffard, and Benoit Boulet. 2022. Building Energy Management With Reinforcement Learning and Model Predictive Control: A Survey. *IEEE Access* 10 (2022), 27853–27862. <https://doi.org/10.1109/ACCESS.2022.3156581>
- [26] Zhiang Zhang, Adrian Chong, Yuqi Pan, Chenlu Zhang, and Khee Poh Lam. 2019. Whole building energy model for HVAC optimal control: A practical framework based on deep reinforcement learning. *Energy and Buildings* 199 (Sept. 2019), 472–490. <https://doi.org/10.1016/j.enbuild.2019.07.029>