

POSTER ABSTRACT P3

A Concept for Standardized Benchmarks for the Evaluation of Control Strategies for Building Energy Management

David Wölfle^{1*}, Kevin Förderer¹ and Hartmut Schmeck^{1,2}

*Correspondence: woelfle@fzi.de
¹FZI Research Center for Information Technology, Haid-und-Neu-Str. 10-14, Karlsruhe, 76131, Germany
Full list of author information is available at the end of the article

Energy Informatics 2019, 2(Suppl 2):P3

Abstract

Given the expected high penetration of renewable energy production in future electricity systems, it is common to consider buildings as a valuable source for the provisioning of flexibility to support the power grids. Motivated by this concept, a wide variety of control strategies for building energy management has been proposed throughout the last decades. However, these algorithms are usually implemented and evaluated for very specific settings and considerations. Thus, a neutral comparison, especially of performance measures, is nearly impossible. Inspired by recent developments in reinforcement learning research, we suggest the use of common environments (i.e. benchmarks) for filling this gap and finally propose a general concept for standardized benchmarks for the evaluation of control strategies for building energy management.

Keywords: building energy management; building control; environment; benchmark; evaluation; reinforcement learning

Introduction

In order to limit the unpredictable and potentially devastating effects of global warming, the governments of most countries have committed themselves to drastically cut CO_2 emissions by 2050 [1]. Consequently, the federal government of Germany has enacted the "Klimaschutzplan" [2] (literally "climate protection plan") defining concrete measures to achieve this ambitious target, including a complete decarbonisation of the German electricity production and an extensive electrification of the heat and mobility sector ("Sektorenkopplung"). While it seems to be consensus that high shares of renewable energy production will lead to an increased demand for energy flexibility [3, 4, 5], it also appears that buildings, in particular the heating, ventilation and air conditioning (HVAC) components of these, are especially suited to provide such energy flexibility [2, 6].

Thus, it is not very surprising that the optimization of energy consumption patterns of buildings is a common research topic. A review on control systems for building energy management [7] lists 121 publications, of which the vast majority optimize HVAC components. Albeit these control strategies appear unrelated to the provisioning of energy flexibility at first glance, both concepts can be connected by dynamic pricing strategies, as suggested, for example, in [8, 9].

Given the need for energy flexibility and the vast and diverse set of potential control strategies to provide such with buildings, the question arises how to evaluate which approach is best suited for a particular task. Following up we will thus

investigate the current state of the art regarding the comparison and evaluation of control approaches for building energy management.

State of the Art

While reviewing the existing literature, it was possible to identify qualitative and quantitative approaches for the comparison of control strategies for building energy management. General reviews like [7, 10, 11] use rather qualitative measures and thus focus on a meta level. The contrasting of control strategies is carried out using general topics, such as the implemented algorithms, the applied control schemas, the utilized simulation tools, and/or the considered devices. Albeit these publications are very useful in general, they provide no objective comparison of the achieved performance of the reviewed control strategies.

On the other hand, in publications following the second approach, it is common that the authors evaluate one or more control strategies in a very specific setting using rather quantitative metrics. For example, Salpakari & Lund [12] developed a complex model consisting of a heat pump with storage, photovoltaic panels, a battery and smart appliances inspired by an existing low energy house. Using that model, the authors compared a rule based control (RBC) strategy to cost optimal behavior with respect to energy costs and self-consumption rates. Both Lösch et al. [13] and Faßnacht et al. [14] developed an algorithm for the scheduling of heat pumps based on a dynamic electricity tariff. The algorithms have been compared against the default hysteresis control strategy in a simulation with energy cost as performance measure. Oldewurtel et al. [15] compared RBC with two MPC approaches utilizing a custom developed simulation model on 1280 test cases with respect to energy usage, thermal comfort and temperature dynamics.

We were able to identify numerous publications following a similar schema to the ones introduced above. However we refrain from giving a more extensive review at this point, as it is out of this paper's scope. Concerning the main topic of this work, the evaluation of control strategies for building energy management, we identify two major issues. First and although many publications use similar metrics like energy usage or energy costs, it is not feasible to compare different approaches in a quantitative manner. This is due to the usual procedure of not reusing the simulation models of others, which may also be caused by the uncommonness of open source publications of these. The second issue arises if one considers the provisioning of electric flexibility through buildings on larger scales, which obviously leads to the utilization of numerous, potentially very diverse, buildings. However, the current procedure of evaluating control approaches against one simulation model alone leaves no indication how well the evaluated strategies generalize to other buildings.

Standardized Benchmarks in Reinforcement Learning

In reinforcement learning, a field closely related to control theory [16], several benchmarks have been published for the development, evaluation and comparison of algorithms. These benchmarks, commonly referred to as environments, have been widely adopted, which allows the direct performance comparison of newly developed algorithms with existing approaches, e.g. in [17, 18, 19]. Popular examples are often based on computer games, like the Arcade Learning Environment [20]

while other environments focus on the control of robots [21] or implement "classic control" problems like balancing a pole on a cart [22]. It appears especially noteworthy that the Arcade Learning Environment in fact consists of 57 games and that reinforcement learning algorithms, e.g. [17, 18, 19], are often routinely evaluated against all of those. The distribution of scores is thereby considered a measure for the generalization abilities of the proposed algorithms, that is, an indication of the expected performance on related tasks.

The interaction with the environments follows the schema of observations, actions and rewards, which is generally well suited for control problems [16] and has been applied successfully to challenging domains, like e.g. aerobatic helicopter flight [23]. Consider the Atari game of space invaders as an example, in which the player can control a space ship at the lower end of the screen and receives points for shooting alien space ships approaching from the top, which is actually part of the Arcade Learning Environment. As common in reinforcement learning, the environment is processed in discrete time steps. In every time step, the environment emits a frame of the gameplay, called observation, as well as the current score referred to as reward. Observation and reward are then passed to the evaluated algorithm, which is usually named agent in reinforcement learning. The agent is queried for the next action, i.e. the control input to the environment, like e.g. move left, move right, shoot or do nothing. Once the action is passed to the environment, it will advance one time step, emitting a new observation and reward. These steps are usually run as a loop until a terminal state is reached, i.e. game over, while the accumulated reward is commonly used as metric to evaluate the performance of the algorithm.

Proposal

In order to overcome the issues identified in the current [State of the Art](#), we propose the establishment of standardized and shared benchmarks for the evaluation of control strategies for building energy management following the example of environments used in reinforcement learning. We propose the following:

- To allow the development of building energy management approaches that can be applied to large number of diverse buildings, the benchmarks should be a collection of distinct building simulation models. These should focus on different optimization targets, e.g. heating, cooling or appliances, as well as varying objectives like e.g. own consumption or energy costs.
- All benchmarks should be published as open source projects to allow widespread usage and verification.
- The communication between algorithm and benchmarks should be standardized in order to limit the effort to execute a benchmark to the necessary minimum. The interface should follow the schema of observations, rewards and actions, as it is well established and allows the usage of available reinforcement learning algorithms for comparison.

Acknowledgements

The authors would like to thank the anonymous reviewers for their valuable feedback.

Funding

This research has been partially founded by the Helmholtz Association's Supercomputing and Big Data program.

Availability of data and materials

Not applicable.

Author's contributions

DW developed the initial concept and wrote this article. KF and HS contributed discussion, feedback and revision.

Competing interests

The authors declare that they have no competing interests.

Author details

¹FZI Research Center for Information Technology, Haid-und-Neu-Str. 10-14, Karlsruhe, 76131, Germany.

²Karlsruhe Institute of Technology, Institute AIFB, Kaiserstr. 12, Karlsruhe, 76131, Germany.

References

1. United Nations Framework Convention on Climate Change: Report of the Conference of the Parties on its twenty-first session, held in Paris from 30 November to 13 December 2015 Addendum Part two: Action taken by the Conference of the Parties at its twenty-first session. Technical report, United Nations Framework Convention on Climate Change (2016). <http://unfccc.int/resource/docs/2015/cop21/eng/10a01.pdf>
2. Bundesministerium für Umwelt, Naturschutz, B.u.R.B.: Klimaschutzplan 2050 – Klimaschutzpolitische Grundsätze und Ziele der Bundesregierung. Technical report (2017). https://www.bmu.de/fileadmin/Daten_BMU/Download_PDF/Klimaschutz/klimaschutzplan_2050_bf.pdf
3. Kondziella, H., Bruckner, T.: Flexibility requirements of renewable energy based electricity systems – a review of research results and methodologies. *Renewable and Sustainable Energy Reviews* **53**, 10–22 (2016)
4. Papaefthymiou, G., Dragoon, K.: Towards 100% renewable energy systems: Uncapping power system flexibility. *Energy Policy* **92**, 69–82 (2016)
5. Alizadeh, M.I., Parsa Moghaddam, M., Amjady, N., Siano, P., Sheikh-El-Eslami, M.K.: Flexibility in future power systems with high renewable penetration: A review. *Renewable and Sustainable Energy Reviews* **57**, 1186–1193 (2016)
6. Jensen, S.Ø., Marszal-Pomianowska, A., Lollini, R., Pasut, W., Knotzer, A., Engelmann, P., Stafford, A., Reynders, G.: IEA EBC Annex 67 Energy Flexible Buildings. *Energy and Buildings* **155**, 25–34 (2017)
7. Shaikh, P.H., Nor, N.B.M., Nallagownden, P., Elamvazuthi, I., Ibrahim, T.: A review on optimized control systems for building energy and comfort management of smart sustainable buildings. *Renewable and Sustainable Energy Reviews* **34**, 409–429 (2014)
8. Samadi, P., Mohsenian-Rad, H., Wong, V.W.S., Schober, R.: Real-Time Pricing for Demand Response Based on Stochastic Approximation. *IEEE Transactions on Smart Grid* **5**(2), 789–798 (2014)
9. Lübker, T., Venzke, M., Turau, V.: Calculating retail prices from demand response target schedules to operate domestic electric water heaters. *Energy Informatics* **1**(S1), 31 (2018)
10. Wang, S., Ma, Z.: Supervisory and Optimal Control of Building HVAC Systems: A Review. *HVAC&R Research* **14**(1), 3–32 (2008)
11. Dounis, A.I., Caraiscos, C.: Advanced control systems engineering for energy and comfort management in a building environment—A review. *Renewable and Sustainable Energy Reviews* **13**(6-7), 1246–1261 (2009)
12. Salpakari, J., Lund, P.: Optimal and rule-based control strategies for energy flexibility in buildings with PV. *Applied Energy* **161** (2016)
13. Loesch, M., Hufnagel, D., Steuer, S., Fabnacht, T., Schmeck, H.: Demand side management in smart buildings by intelligent scheduling of heat pumps. In: 2014 IEEE International Conference on Intelligent Energy and Power Systems (IEPS), pp. 1–6 (2014). <http://ieeexplore.ieee.org/document/6874181/>
14. Faßnacht, T., Loesch, M., Wagner, A.: Simulation Study of a Heuristic Predictive Optimization Scheme for Grid-Reactive Heat Pump Operation. In: Proceedings of REHVA Annual Conference "Advanced HVAC and Natural Gas Technologies", vol. 0, p. 132. Riga Technical University, Riga (2015). <https://journals.rtu.lv/index.php/AHNGT/article/view/rehvaconf.2015.019>
15. Oldewurtel, F., Parisio, A., Jones, C.N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., Morari, M.: Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings* **45**, 15–27 (2012)
16. Sutton, R.S., Barto, A.: Reinforcement Learning - An Introduction, 2nd edn., p. 526 (2018). <http://incompleteideas.net/book/the-book.html>
17. Gruslys, A., Dabney, W., Azar, M.G., Piot, B., Bellemare, M., Munos, R.: The Reactor: A fast and sample-efficient Actor-Critic agent for Reinforcement Learning (2017). [1704.04651](https://arxiv.org/abs/1704.04651)
18. Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T.P., Harley, T., Silver, D., Kavukcuoglu, K.: Asynchronous Methods for Deep Reinforcement Learning (2016). [1602.01783](https://arxiv.org/abs/1602.01783)
19. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D.: Human-level control through deep reinforcement learning. *Nature* **518**(7540), 529–533 (2015)
20. G. Bellemare, M., Naddaf, Y., Veness, J., Bowling, M.: The Arcade Learning Environment: An Evaluation Platform for General Agents. *Journal of Artificial Intelligence Research* **47** (2012)
21. Plappert, M., Andrychowicz, M., Ray, A., McGrew, B., Baker, B., Powell, G., Schneider, J., Tobin, J., Chociej, M., Welinder, P., Kumar, V., Zaremba, W.: Multi-Goal Reinforcement Learning: Challenging Robotics Environments and Request for Research (2018). [1802.09464](https://arxiv.org/abs/1802.09464)
22. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: OpenAI Gym (2016)
23. Abbeel, P., Coates, A., Quigley, M., Ng, A.Y.: An Application of Reinforcement Learning to Aerobic Helicopter Flight. *Advances in Neural Information Processing Systems* **19** (NIPS 2006) (2006)