Stripping off the implementation complexity of physics-based model predictive control for buildings via deep learning

Ján Drgoňa[1,2], **Lieve Helsen**[2,3], and **Draguna L. Vrabie**[1]

[1]Pacific Northwest National Laboratory, Richland, WA, USA
{jan.drgona, draguna.vrabie}@pnnl.gov
[2]Department of Mechanical Engineering, KU Leuven, Belgium
{jan.drgona, lieve.helsen}@kuleuven.be
[3] EnergyVille, Thor Park, Waterschei, Belgium

Abstract

Over the past decade, model predictive control (MPC) has been considered as the most promising solution for intelligent building operation. Despite extensive effort, transfer of this technology into practice is hampered by the need to obtain an accurate controller model with minimum effort, the need of expert knowledge to set it up, and the need of increased computational power and dedicated software to run it. A promising direction that tackles the last two problems was proposed by approximate explicit MPC where the optimal control policies are learned from MPC data via a suitable function approximator, e.g., a deep learning (DL) model. The main advantage of the proposed approach stems from simple evaluation at execution time leading to low computational footprints and easy deployment on embedded HW platforms. We present the energy savings potential of physics-based (also called 'white-box') MPC applied to an office building in Belgium. Moreover, we demonstrate how deep learning approximators can be used to cut the implementation and maintenance costs of MPC deployment without compromising performance. We also critically assess the presented approach by pointing out the major challenges and remaining open-research questions.

1 Introduction

Nowadays buildings use roughly $40\,\%$ of the global energy (approx. 64 PWh), a large portion of which is being used for heating, cooling, ventilation, and air-conditioning (HVAC) [1]. The energy efficiency of buildings is thus one of the priorities to sustainably address the increased energy demands and reduction of CO_2 emissions in the long term [2].

It has been shown that smart control strategies like model predictive control (MPC) can maximize system-level efficiency for existing built environments, thus reducing the emissions of greenhouse gases, and can improve the thermal comfort of the occupants, with reported energy use reductions of $15\,\%$ up to $50\,\%$ [3, 4, 5].

Despite this, the practical implementations of MPC are hampered by the challenge of obtaining an accurate controller model with minimum effort, the need of expert knowledge to set it up, and the need of increased computational power and dedicated software to run it [6]. Every building represent a unique system which requires tailored modeling and control design.

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

The MPC in this work is based on detailed physical modeling of a real-life office building which provides an accurate prediction of the building's thermal behavior and high control performance. On the other hand, the disadvantage of such high-fidelity MPC approach lies in its computational demands and software dependencies. Here we are exploring the use of DL to learn the optimal control policies from MPC data. The main advantage of the proposed method stems from its low computational footprint, minimal software dependencies and easy deployment even on low-level hardware without compromising control performance. The advantage compared to reinforcement learning is its sample efficiency, because policies are learned with supervision from pre-computed optimal control trajectories in realistic operational scenarios.

2 Optimization-based model predictive control using physical models

The office building considered in the experimental and simulation study called *Hollandsch Huys* is located in Hasselt, Belgium. Hollandsch Huys represents a so-called *GEOTABS* building with slow dynamics and complex heating ventilation and air conditioning (HVAC) system [7]. The building's layout consists of five floors divided into 12 thermal zones. For detailed description of the building and physics-based modeling in Modelica language we refer to [8]. The main advantage of such a high fidelity physics-based "digital twin" model stems from its potentially high prediction accuracy, interpretability, and reliability. Based on the approach described in [9] the physics-based model can be transformed to state-space representation with 700 states x, 301 disturbance signals d, 12 thermal zones y and 20 control inputs u.

Fig. 1 shows the corresponding control configuration. The optimization-based MPC (OB-MPC) computes the optimal control actions u, based on estimated states x via Kalman Filter (KF), for details see [10, 11]. The MPC problem is solved using a state-of the art optimization solver Gurobi [12] running in the MATLAB environment. The non-linear weather forecaster model is running in the Dymola environment and computes the forecasts of disturbances d (weather, occupancy), and reference r trajectories based on actual weather data w obtained from the Dark Sky API [13]. Optimal control actions at the current time step u_0 represent the heat flows to be delivered to the building and are re-computed once per sampling time in so-called receding horizon control (RHC) fashion.

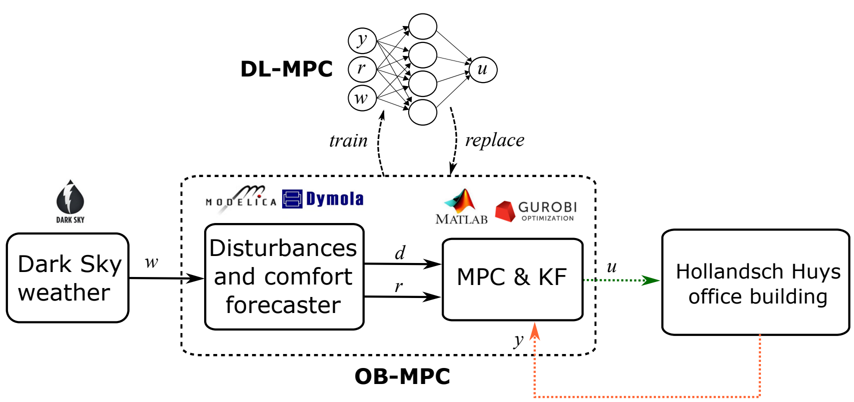

Figure 1: Optimization-based MPC methodology with deep learning-based policy approximator.

3 Deep learning-based approximation of MPC policies

The central idea here is based on learning the optimal control policies from optimal trajectories generated by OB-MPC via deep learning model in an imitation learning fashion, as shown in Fig. 1. A detailed description of the applied methodology can be found in [10]. After training, the DL-MPC policy replaces the computationally heavy and costly OB-MPC implementation. We use MATLAB's neural network toolbox for the design and training of the three-layer time-delayed neural network on 330 days of simulated operation of the original OB-MPC.

4 Experimental and simulation results

The real operational performance of the physics-based OB-MPC is compared to the conventional rule-based controller (RBC) on a dataset of 72 days (31 for MPC, 41 for RBC) during the transient season (intermediate between spring and summer). The mean ambient temperature for the MPC dataset is 17.3 °C, and for RBC it is 18.8 °C. The corresponding HP energy savings of OB-MPC are equal to 50.4 %, with a thermal comfort improvement of 50.5 %. However, it is essential to mention that that these are preliminary results for the transient season, that can not be generalized over all seasons. Nevertheless, these results are encouraging and provide a glimpse of the energy-saving potential of the proposed physics-based predictive control strategy in a real setting.

Subsequently, we evaluate the control performance on a simulated 30-days test set, together with the deployment cost reduction of the proposed DL-MPC with respect to OB-MPC. The simulation setup is idealized, as no uncertainty in the feature space of both OB-MPC and DL-MPC is considered. As a result, DL-MPC kept very high comfort satisfaction close to 100 %, but it slightly increased the energy use roughly by 3 % w.r.t. high-fidelity OB-MPC. Yet, DL-MPC kept high energy saving potential compared to the classical RBC. However, in contrast to the runtime and deployment cost of OB-MPC, the presented neural policies require only a fraction of computational and memory resources without the need for expensive software dependencies. In this case, we observed that DL-MPC is roughly 50 000-times faster and consumes 638-times less memory. The overall control performance, average CPU evaluation time per sample [1], memory footprint [2], together with the cost associated with commercial software licenses [3] are summarized in Tab. 1.

Table 1: Comparison of OB-MPC and DL-MPC. Performance indicators: simulation performance on 30-days test set, computational and memory footprint, and software deployment cost.

Method	**Discomfort** [K h]	**Energy use** [kW h]	**CPU time** [1×10^{-3} s]	**Memory** [MB]	**SW Deployment Cost** [$]
OB-MPC	0.0	801.2	26 843	415	18, 050
DL-MPC	0.15	824.5	0.528	0.65	0

5 Conclusions, challenges and future work

In this work, we demonstrated the preliminary results of the energy-saving potential of the optimization-based model predictive control (OB-MPC) based on a physical model in the operation of the real office building in Belgium. Additionally, we showed on simulation results, how deep learning technology could be used to reduce the deployment cost of such advanced control strategies, maintaining high control performance, while using only a fraction of computational resources.

However, several open-research problems remain unanswered. For example, what is the optimal topology and hyperparameter setup for efficient representation of such problems? How to guarantee satisfactory control performance far from the optimal trajectory? How sensitive is the policy to uncertainty in weather forecast? Does the policy stabilize the closed-loop system? How to explicitly include constraint handling properties of OB-MPC into DL-MPC policies? How can we use predictive models and state estimation algorithms to further improve policy performance based on feedback? How can we verify the policies using physics-based models? Can we parametrize the policies based on physical parameters of the buildings to be used in a transfer learning fashion? Can we create synthetic training datasets using generative models with the aid of physics-based modeling? Can we use generative models to synthesize the policies directly from the building parameters?

Future work of the authors, includes deployment of trained DL-MPC policies in a real office building. As the step towards computationally efficient and interpretable neural network policies for real-world systems, the authors are focusing on the development of novel deep neural topologies inspired by the sparse structure of the physics-based models and optimal control problems.

[1] In case of OB-MPC the average runtime is the sum of 24.534 s for the non-linear weather forecaster model running in Dymola and 2.309 s for the MPC solution via Gurobi.

[2] In case of OB-MPC, only the implementation code and actively used libraries are evaluated. We are omitting the memory requirements of the MATLAB and the Dymola environments themselves.

[3] Overall costs are computed as aggregate cost of MATLAB perpetual license (2, 150 $), Gurobi single user license (10, 000 $), and Dymola standard license (5, 900 $).

References

[1] IEA International Energy Agency and International Partnership for Energy Efficiency Cooperation. Building energy performance metrics - supporting energy efficiency progress in major economies. Technical report, IEA Publications, 2015.

[2] David Rolnick, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, Alexandra Luccioni, Tegan Maharaj, Evan D. Sherwin, S. Karthik Mukkavilli, Konrad P. Körding, Carla Gomes, Andrew Y. Ng, Demis Hassabis, John C. Platt, Felix Creutzig, Jennifer Chayes, and Yoshua Bengio. Tackling climate change with machine learning. *CoRR*, abs/1906.05433, 2019.

[3] Jan Širokỳ, Frauke Oldewurtel, Jiří Cigler, and Samuel Prívara. Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy*, 88(9):3079–3087, 2011.

[4] D. Sturzenegger, D. Gyalistras, M. Morari, and R. S. Smith. Model predictive climate control of a swiss office building: Implementation, results, and cost–benefit analysis. *IEEE Transactions on Control Systems Technology*, 24(1):1–12, Jan 2016.

[5] Y. Ma, F. Borrelli, B. Hencey, B. Coffey, S. Bengea, and P. Haves. Model predictive control for the operation of building cooling systems. *IEEE Transactions on Control Systems Technology*, 20(3):796–803, 2012.

[6] J. Cigler, D. Gyalistras, J. Široký, V. Tiet, and L. Ferkl. Beyond Theory: the Challenge of Implementing Model Predictive Control in Buildings. In *Proceedings of 11th Rehva World Congress, Clima*, 2013.

[7] Elisa Van Kenhove, Jelle Laverge, Wim Boydens, and Arnold Janssens. Design Optimization of a GEOTABS Office Building. *Energy Procedia*, 78:2989 – 2994, 2015. 6th International Building Physics Conference, IBPC 2015.

[8] D. Picard. Modeling, Optimal Control and HVAC Design of Large Buildings using Ground Source Heat Pump Systems, PhD Thesis, KU Leuven, Belgium. 2017.

[9] Damien Picard, Filip Jorissen, and Lieve Helsen. Methodology for obtaining linear state space building energy simulation models. In *Proceedings of the 11th International Modelica Conference*, pages 51–58, Paris, France, 2015.

[10] J. Drgoňa, D. Picard, M. Kvasnica, and L. Helsen. Approximate model predictive building control via machine learning. *Applied Energy*, 218:199 – 216, 2018.

[11] D. Picard, J. Drgoňa, M. Kvasnica, and L. Helsen. Impact of the controller model complexity on model predictive control performance for buildings. *Energy and Buildings*, 152:739 – 751, 2017.

[12] Inc. Gurobi Optimization. Gurobi optimizer reference manual, 2012.

[13] LLC. The Dark Sky Company. Dark Sky API.