Finding the Best Models for Model Predictive Control in Buildings

B. Schubnel, R. E. Carrillo, P.-J. Alet, A. Hutter, T. Gorecki

As a complement to thermal renovation, integration of intelligent controllers in buildings can significantly reduce the CO_2 footprint of heating and cooling systems by increasing their usage efficiency. Within the European project SABINA, CSEM investigated what could be the best machine learning architecture to represent and control complex buildings. Representation power does not go hand in hand with controllability, and our results suggest that linear state space models with simple non-linear regressors give better control results than RNN architectures for model predictive control. The best architecture found in our study has been deployed on the CSEM site in Neuchâtel.

Deep learning architectures have been increasingly popular to model non-linear physical systems. They are known to provide a net gain in prediction accuracy as compared to state of the art (linear state space models). However, if deep models are used within an optimization problem, their non-convexity creates issues, and their higher accuracy does not necessarily translate into a net gain for optimal controllers built on top of them. Energy management in buildings is a representative example of systems where such models can be used, as recent technical equipment (heat pump, HVAC) or low-level control loops are a source of non-linearities.

CSEM recently investigated deep architectures called recurrent encoder-decoder networks (or "seq2seq" models), that show very good representation power on a large set of physical systems. We have shown [1] that they could effectively be used as models to control complex physical systems, in particular large building facilities, using as controllers neural networks trained with state-of-the-art reinforcement learning (RL) algorithms. Using model-based RL algorithms for optimal control is a promising path for such non-linear architectures but is still in its infancy for industrial applications due, notably, to sample inefficiency and difficulties in reward shaping.

In our recent work [2], we investigated how encoder-decoder models were performing with respect to linear state space models coupled with non-linear regression (LSS-NL), in the case were traditional optimization techniques (here sequential quadratic programming, SQP) rather than RL are used to solve the optimal control problem. Our study was made on a simulated building with four apartments, eight thermal zones and solar panels on its roof (Figure 1).

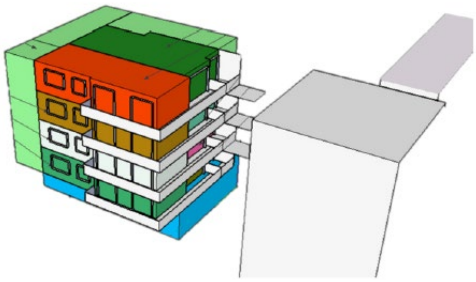

Figure 1: Overview of the test building with its four floors.

The simulated building incorporated two main sources of non-linearities. Its first source is a geothermal heat pump that provides heat to the living spaces and the hot water tanks; see Figure 2. Its second source is a low-level control loop (that cannot be changed, as often in real buildings) that triggers the opening of the thermostatic valves and the on-off cycles of the pump, based on the room temperature setpoints.

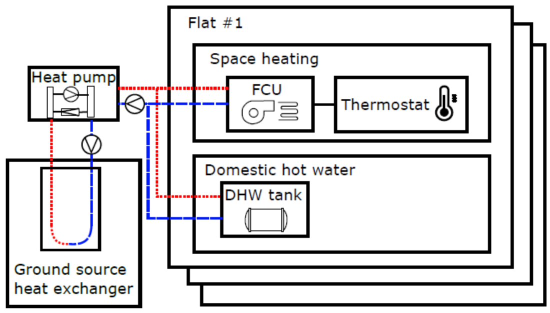

Figure 2: Schematic view of the heating and distribution system.

The purpose of the optimal controller was to respect the temperature constraints in the building (the zone temperatures had to stay between 19°C and 23°C) while minimizing the energy exchanged with the grid (i.e., consuming as much as possible its own PV production). Both architectures were used to solve the same optimization problem under the same conditions with SQP.

As expected, encoder-decoder architectures need more samples to carry out proper system identification but are more accurate once enough data are used. More surprisingly, the objective minimization with SQP was best performed with the simpler architecture, even if it slightly underperformed on the comfort constraints. Even if we tried to optimize the computation time with the encoder-decoder, and performed gradient computation with the GPU, we could not achieve the same computation time as with simpler models.

Based on these results, CSEM has developed a library for model predictive control with the LSS-NL architecture and deployed it successfully in the CSEM building ("Ecrivain" and "Dessinateur" meeting rooms).

This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement no 731211, project SABINA.

[1] B. Schubnel, R. E. Carrillo, P.-J. Alet, A. Hutter, "A Hybrid Learning Method for System Identification and Optimal Control," IEEE Transactions on Neural Networks and Learning Systems (2020) [doi:10.1109/TNNLS.2020.3016906].

[2] B. Schubnel, R. E. Carrillo, P. Taddeo, L. C. Casals, J. Salom, Y. Stauffer, P.-J. Alet, "State-space models for building control: how deep should you go?," Journal of Building Performance Simulation 13(6), 707–719, Taylor & Francis (2020) [doi:10.1080/19401493.2020.1817149].