White-box Model Predictive Control: Optimal Control and System Integration of Heat Pumps

Filip Jorissen, Damien Picard, Wim Boydens and Lieve Helsen, Belgium

Model Predictive Control (MPC) has a large sustainability potential for the optimal control of Heating, Ventilation and Air Condition in buildings. This article summarizes some of the main features of MPC, and the advantages and results (including real-life demonstrations) of our particular implementation, which uses detailed physics-based simulations and optimizations of both the building envelope and its HVAC. This research track has been developed at the Thermal Systems Simulation (The SySi) research group of KU Leuven over the past 12 years.

Building Heating, Ventilation, and Air Conditioning (HVAC) systems are becoming more complex due to the integration of renewable energy sources and heat pump-based technologies. We are evolving from a society where the building demand determines heating and cooling loads to a society where the availability of heat and cold, through price signals, determines when to heat or cool a building. Furthermore, renewable energy sources tend to use lower temperatures for heating and higher temperatures for cooling. E.g. heat pumps operate more efficiently at smaller temperature differences, and direct geothermal cooling is simply not available at temperatures lower than the soil temperature. Smaller temperature differences result in smaller thermal powers, meaning that sudden power peaks have to be spread over longer periods. In order to reach the building comfort set points in time, this typically means that heating and cooling has to start sooner, depending on the emission system inertia. Thermally massive systems such as Concrete Core Activation (CCA) can benefit from heating/cooling loads that are shifted to multiple hours before the heating/cooling is required.

The delayed temperature response is illustrated in Figure 1, as well as a significant lasting temperature influence during the days following the cooling peak. The model thus, in fact, predicts that the cooling peak introduces a heating demand for satisfying the lower temperature limit starting at hour 100, which illustrates the importance of dynamic simulations for the system design.

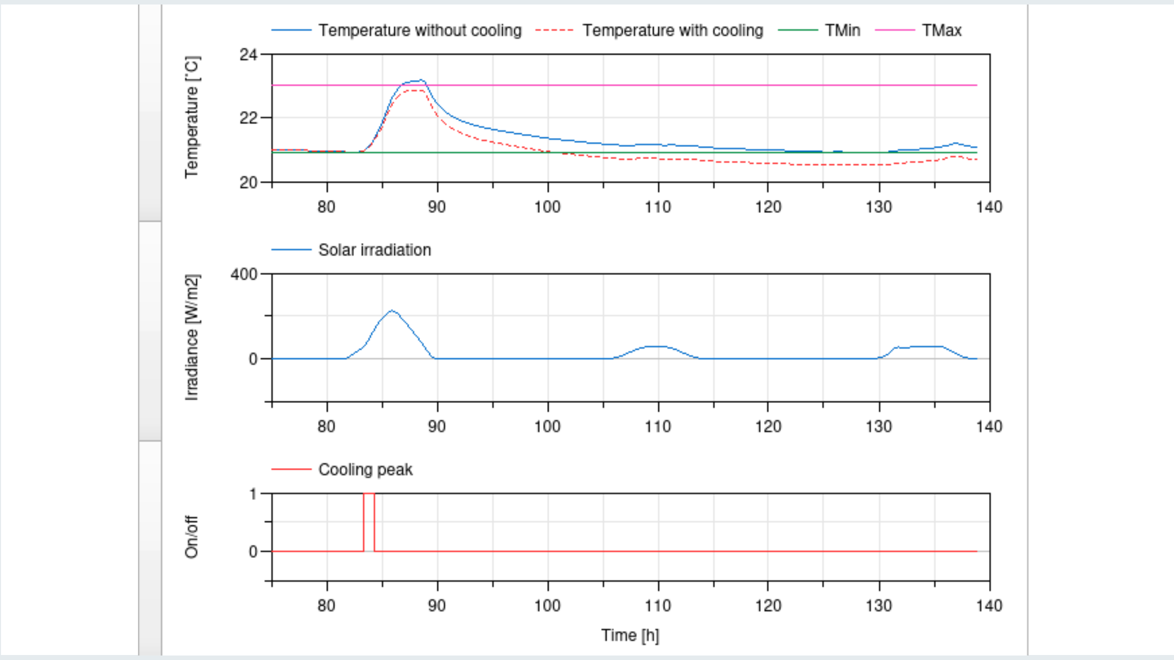

Figure 1. Delayed zone temperature response after a short cooling peak injection in a floor cooling slab.

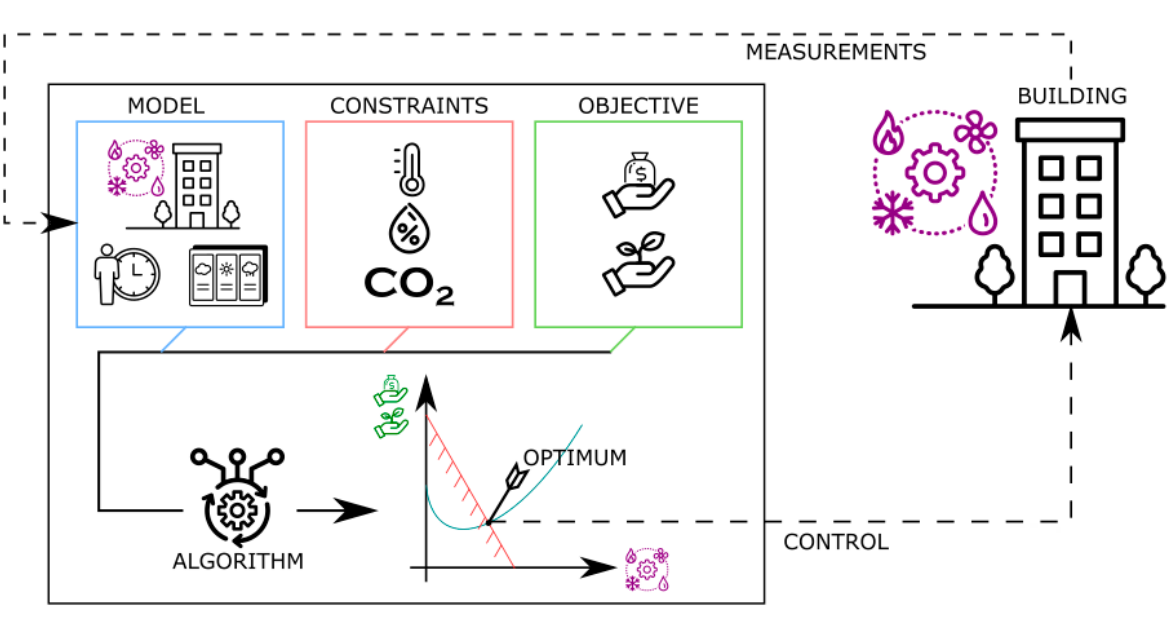

Figure 2. Illustration of the main MPC components and control loop.

Next to CCA, additional storage options such as thermal storage tanks could be considered, adding degrees of freedom where similar questions arise. How large should the tank be, and to what temperature levels should the tank be charged, and when?

This discussion firstly illustrates the need for adequate HVAC system selection and component sizing during the design stage. There no longer exists one correct design, rather a multitude of feasible designs that consider the local context and potential of the building and the sustainability ambitions of the stakeholders in the construction process.

Secondly, the flexible, timely and sustainable delivery of heating and cooling loads will require a predictive control solution. Heat supply to a building is an interplay of compressor speeds, pump speeds and valve openings, which will have to be coordinated with quarter-hourly changing electricity prices and the availability of renewable energy sources. The building use, and even climate changes throughout the building's lifetime calling for a tailored yet easily adjustable and reconfigurable control solution.

To tackle the aforementioned control challenges in commercial buildings, we propose white-box model predictive control (MPC). MPC is a predictive control methodology that relies on a mathematical model of a system to control that system, in this case, a building and its HVAC. The model considers weather forecasts, occupancy, the building envelope, and the HVAC devices that are connected to the building envelope. The model predicts the impact of the current control actions on the building energy use, emissions and comfort (KPIs) and on operational constraints during the coming days. It determines what control actions result in the best trade-off between these KPIs. Control signals are sent to the existing building management system, and 15 minutes later, the optimization is repeated. This is illustrated in Figure 2.

While this summarizes the main functionality of an MPC, various approaches exist for implementing the model, which has to be custom developed for a building. Differences are related mostly to 1) the number of components and 2) the level of detail of those components. Both aspects can strongly affect the computational effort for solving the model equations, causing problems to become practically insolvable by general-purpose solvers. Therefore, models are often either simplified by reducing the number of components, e.g. by representing a whole building using one or a few rooms only. Such implementations are consequently unable to provide individual set points for different rooms, which may have very different heating or cooling demands. Other model simplifications mostly focus on reducing the level of detail of the models. Such (linear) models generally have a strong mismatch between the model and the physical system, making it impractical to couple the model to the physical building. This is problematic from a business deployment point of view since every building requires an ad-hoc solution and consequently requires a high level of expertise to install the controller.

Our goal is to capture the full complexity of each building and exploit the full potential of its HVAC system. Therefore, we use models that are both detailed in the represented physics and in the spatial representation of the rooms of the building. Thanks to this, the model is aware of the full system complexity. While this approach is accurate and generic (applicable to all buildings), it leads to large models, which may be difficult to optimize and time-consuming to set up. Solving those issues has been the core of our research and development over the last eight years, and we have shown that our custom solver is able to optimize our complex models sufficiently quickly [1]. Our goal is to have no upper limit on the model complexity and size that our solver can optimize, starting from the practical and scalable model implementation methodology described below.

We use a library of physics-based (white-box) component models where each component model (e.g. a heat pump) has a set of characteristic parameters (e.g. the nominal heat flow rate and a COP curve). Physics do not change, and consequently, the main model equations have to be developed only once. These models are implemented using the equation-based modeling language Modelica and are available in the open-source Modelica library IDEAS (https://github.com/open-ideas/IDEAS) [2]. Model parameters are configured using available manufacturer specifications. Each component has connectors that correspond to relevant physical quantities (e.g. four fluid ports and an electrical power inlet for a heat pump model). Building hydronic and aerolic schematics are used to specify how all components are interconnected. Building plans are used to define the building geometry. The resulting modeling process is a simple mapping of physical components into component models, see Figures 3 and 4 and is easy to understand and execute. Ongoing developments related to Building Information Modeling (BIM) could even lead to further automation of this work. Differentiation between heat pumps could be facilitated by integrating manufacturer-specific performance curves. For existing buildings, the required as-built information may not be available. Extensions of the approach are being developed for these buildings.

Within the EU H2020 hybrid GEOTABS project and in collaboration with stakeholders such as Boydens engineering, part of Sweco, UGent, DTU, EnergoKlastr, Viessmann, Uponor, KU Leuven has demonstrated this technology within three large buildings (e.g. Figure 5, Ter Potterie), thanks to the building owners Fluvius, Mintus, Progroup and Schuler. These buildings had been thoroughly commissioned within the construction phase, but MPC was nevertheless able to achieve similar or improved thermal comfort, and reduced the energy use and associated CO_2 emissions. In some cases, the MPC relied on the existing sub-controllers of the building management system for conveying its set point to the system. A mismatch between the expected and observed system behaviour even pointed out some errors in the existing system and enables a higher degree of commissioning. Such errors could be identified automatically in the future, or could even be avoided entirely by no longer relying on these sub-controllers, thereby also removing the cost for implementing them.

One of the demonstration buildings is the four story, 3000 m² office building of Fluvius and Boydens engineering, part of Sweco in Brussels. In this building, energy savings of more than 40% were observed during long periods of the year. The building uses a geothermal heat pump, CCA, an Air Handling Unit (AHU) with a heating coil (HC1), heat recovery, and multiple VAVs with individual heating coils (HC2). The substantial savings could be explained as follows. For the relatively colder bottom floor, MPC managed to shift a larger fraction of the heating load to the CCA, which allowed the AHU flow rates and the supply air temperature for the whole building to decrease. Instead of supplying 100% air flow to the ground floor, about 50% sufficed, which could be heated up to the desired supply air tempera-

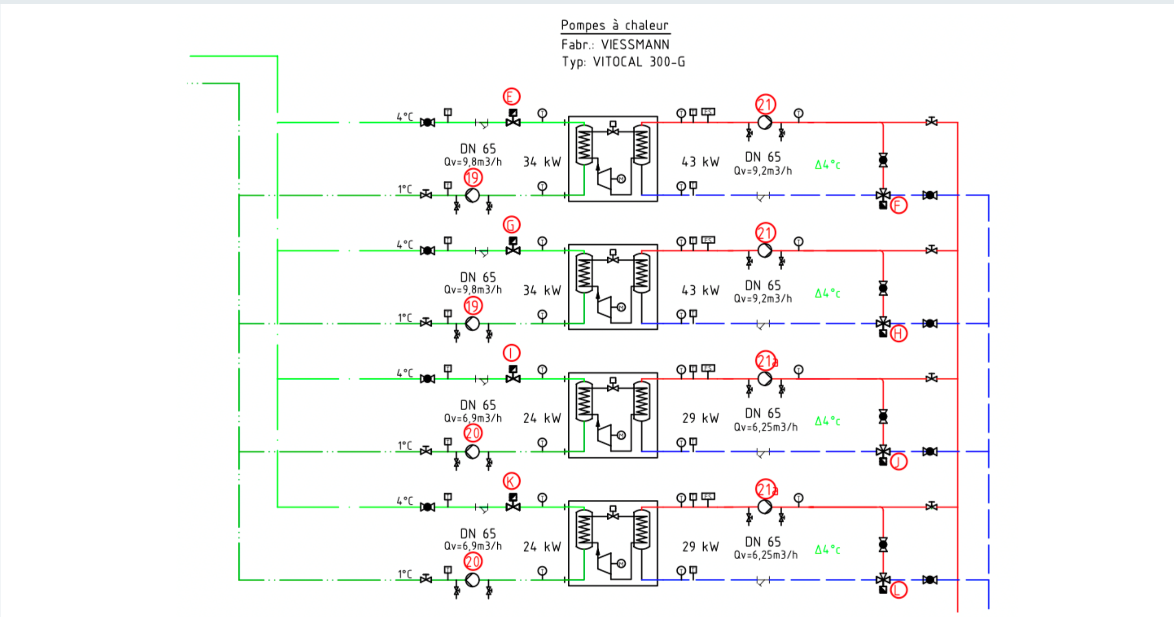

Figure 3. Schematic of a heat pump system. Source: Boydens Engineering, part of Sweco.

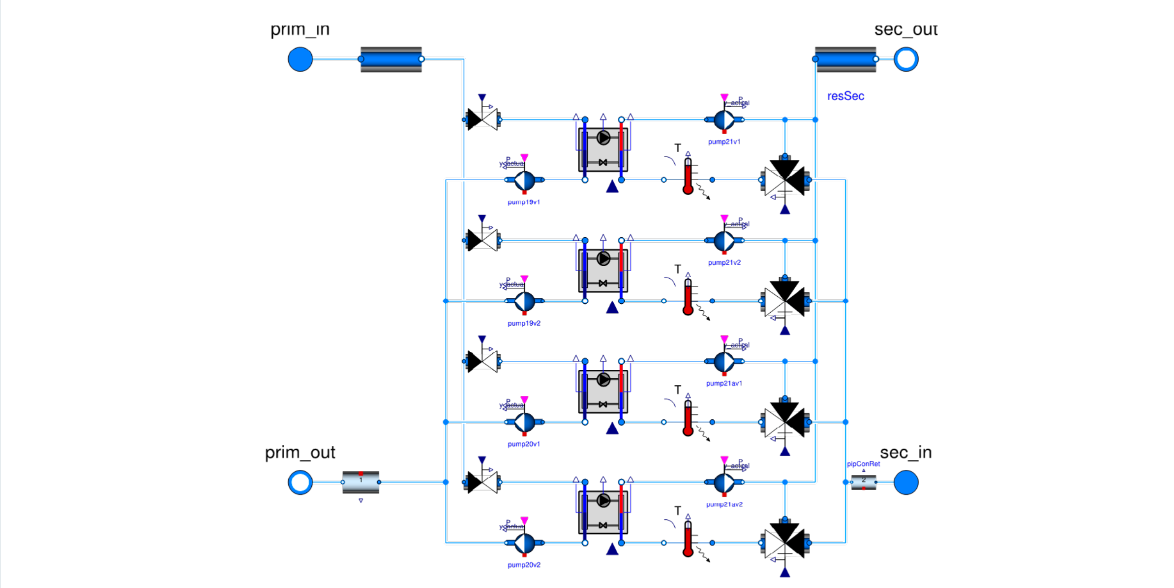

Figure 4. Mapped model of the heat pump system. Illustrated icons are components models.

ture by using only HC2, no longer requiring HC1, which would have increased temperatures supplied to the whole building. This removed a substantial heat load from HC1.

Furthermore, the smaller air flow rates allowed the supply and return air pressures to be lowered, which resulted in significantly reduced electrical power use for the AHU. Finally, the heat pump COP could be increased by lowering the condenser supply water temperature, albeit only slightly, since HC2 did, in fact, require relatively high supply water temperatures. Note that in this particular case, MPC decided not to significantly reduce the condenser water temperature, since that would have required the AHU flow rates and HC1 to pick up additional heat load again, which has a worse effect than the COP decrease. The strength of our detailed MPC implementation is that it was able to deal with these complex interactions of multiple zones and components, while at first sight, we, in fact, believed that the elevated condenser temperatures were a bug. Furthermore, these results were obtained without requiring substantial tuning of the controller.

Conclusions

Traditional rule-based controllers implement static control rules, which rigidly focus on tracking a particular set point. When tuned correctly, this can result in the desired comfort. However, many set points can lead to the same comfort level at different costs and CO_2 emissions. Practice has shown that in modern buildings, even tuning these parameters can be a challenge, let alone tuning them under varying occupancy levels, climates and time-dependent prices. Our results illustrate the true potential of MPC. It optimizes the operation on system level. It is bound by the real system constraints only, no extra rules, and is, therefore, able to identify control solutions that a human might not discover, for working conditions we did not expect, and a future climate we cannot yet predict.

Yet, the future climate is approaching, and we have to invest in our planet today to avoid crossing environmental limits and associated costs in the future, in fact, quite similar to Figure 1. We and the KU Leuven spin-off Builtwins believe that the outlined technology will play an important role as part of our mission towards climate-neutral building operation. We thank the Flemish government, KU Leuven and the European Commission for funding our research.

Figure 5. Photograph of elderly care home Ter Potterie. Source: Boydens Engineering, part of Sweco.

FILIP JORISSEN
KU LEUVEN
Belgium
filip.jorissen@gmail.com
https://doi.org/10.23697/4934-p272