

Detailed White-Box Non-Linear Model Predictive Control for Scalable Building HVAC Control

Filip Jorissen¹ Damien Picard¹ Kristoff Six¹ Lieve Helsen^{1,2}

¹Mechanical engineering, KU Leuven, Belgium filip.jorissen@kuleuven.be

²EnergyVille, Genk, Belgium

Abstract

Grey-box and black-box MPC approaches for building HVAC applications often use lumped, low-order models with a low level of detail. While such models require smaller computation times, their accuracy is limited and there are practical constraints related to data collection, how to deal with multi-zone buildings and they often do not explicitly model the building HVAC equipment. In this paper we present an alternative approach based on detailed white-box models. TACO, a custom toolchain that builds upon physics-based Modelica models and JModelica, is used to efficiently solve the resulting optimisation problems. This paper presents a realistic case study model of 79 zones and OCP results for this case study are discussed, demonstrating the feasibility of the approach and the underestimated potential of detailed white-box MPC.

Keywords: Optimal control of hybrid systems, HVAC, white-box modelling, building automation, TACO, JModelica, MPC

1 Introduction

Building Heating, Ventilation and Air Conditioning (HVAC) accounts for 15 % of the world final energy use (International Energy Agency 2019). While building design standards become stricter, the building energy use is to a large extent determined during operation, when control and the available building flexibility play an important role. Model Predictive Control (MPC) is a methodology for controlling the building HVAC equipment during this operational phase. The goal of MPC is typically to find the HVAC control set points (or control actions) that lead to the lowest (operational, environmental or other) cost while ensuring that comfort and other constraints are respected.

Most MPC research and companies use data-driven approaches such as grey-box and black-box modelling. These approaches fit model parameters using measurements from a real building and thus rely on the availability of qualitative data, which may not be easy to obtain in practice. Moreover, the main challenge of grey-box modelling still is the need for a robust parameter estimation method (Drgoňa et al. 2020). For data-driven approaches, building zones are therefore often lumped to limit the complexity during the training phase. This lumping phase inevitably leads to a loss of detail, which could

be problematic when different parts of the lumped zones have a different behaviour. E.g. the set points for a lumped zone with average heat load may not be sufficient to heat the two consisting zones that have low and high heat load respectively. Perhaps this infrequently poses problems in practice, but managing these kinds of problems takes time (which is expensive and a liability in a commercial context), and model simplifications are likely to cause reduced energy efficiency and thermal comfort throughout the building lifetime. Furthermore, simplified models typically require the development of a subcontroller that maps MPC set points to device set points, which can be cumbersome (Drgoňa et al. 2020).

Using models with a higher level of detail leads to additional advantages. For instance, the detailed physics-based model can be used to benchmark the actual system, it can be used for fault detection (Frank et al. 2016), model and results are easier to interpret ('Explainable AI'), the model could be easier to adjust and update since it is physically interpretable, it could be reused for retrofit analyses, or even to visualize the building using augmented reality. Sometimes white-box models are used to train a grey or black-box model. Why would you train a simplified model if you already have a detailed model? Often the answer to that question relates to computation time, or to the inability of the white-box tool to perform optimizations altogether. Our goal is to improve the solver, rather than to simplify the model. We start from physics-based Modelica models, and JModelica (Åkesson et al. 2010) to keep the original model accuracy, at acceptable computational cost. For a recent overview of MPC for building applications we refer to (Drgoňa et al. 2020).

Considering the above and also the economic context we present an MPC development workflow that is designed with scalability in mind. More specifically, it is designed to be robust against modelling errors (user error), fast to use, easy to maintain and to extend, generic for many types of HVAC devices and HVAC schematics, low-demanding with respect to expertise to implement and operate, and systematically applicable to (a class of) buildings. Our approach uses detailed white-box Modelica models as a starting point, which can be refined using measurement data during operation. The use of Modelica

unlocks efficient numerical algorithms that scale well with the problem size and it facilitates collaborations on model development (e.g. within IBPSA project 1). An often quoted disadvantage of white-box modelling is the effort required to describe building properties (Drgoňa et al. 2020). By using automated modelling workflows, this implementation effort is strongly reduced.

Few researchers have demonstrated white-box optimal control methodologies that scale to the size and complexity that is required for large buildings, assuming that the goal is a qualitative implementation that sets individual set points for individual actuators. (Sturzenegger et al. 2016) have presented a respectable physics-based modelling approach, BRCM, that uses bi-linear models. A 20-zone model with 300 states has been demonstrated and model reduction has been used to reduce the number of states to 55. While the total conditioned floor area was 6000 m², only one of the six floors (1000 m²) was actually modelled. In this work we present a white-box modelling approach that has been applied to case-study model that has 3421 state variables, 217 control inputs and 79 zones that span 10 000 m². Furthermore, multiple *non-linear* AHUs and other HVAC equipment are included instead of only using bi-linear HVAC. The resulting proof-of-concept Optimal Control Problem (OCP) demonstrates the feasibility of this approach even for complex multi-zone buildings. Furthermore, the presented optimisation results illustrate the potential of our approach. We present OCP results in this work instead of MPC results since results of a single optimisation with a long horizon are easier to interpret than a concatenation of the first intervals of a sequence of optimisations. We do not consider uncertainties on disturbance forecasts or modelling errors in this work. Clearly, our approach would have to deal with these uncertainties, as would any other approach. Note that the model can indeed be used for MPC too, since MPC is controlling the modelled building (located in Luxemburg) since a few months. A preliminary interpretation of the results suggests that forecasting errors of weather and occupancy dominate modelling errors.

2 Methodology

Our methodology uses detailed, physics-based (white-box) models. We start by mapping each physical component to a respective model. For the building envelope, zones are grouped such that rooms that can be controlled individually, are also modelled individually. For the building HVAC, components can be mapped to models one-to-one as illustrated in Figure 1, or each group of components that serves the same function is modelled using one component model. E.g. each pump, valve, heat exchanger, etc. is modelled individually, but a set of two redundant pumps, or a set of 20 solar collectors has 1 model per set. The component models are implemented using the modelling language Modelica. For a description

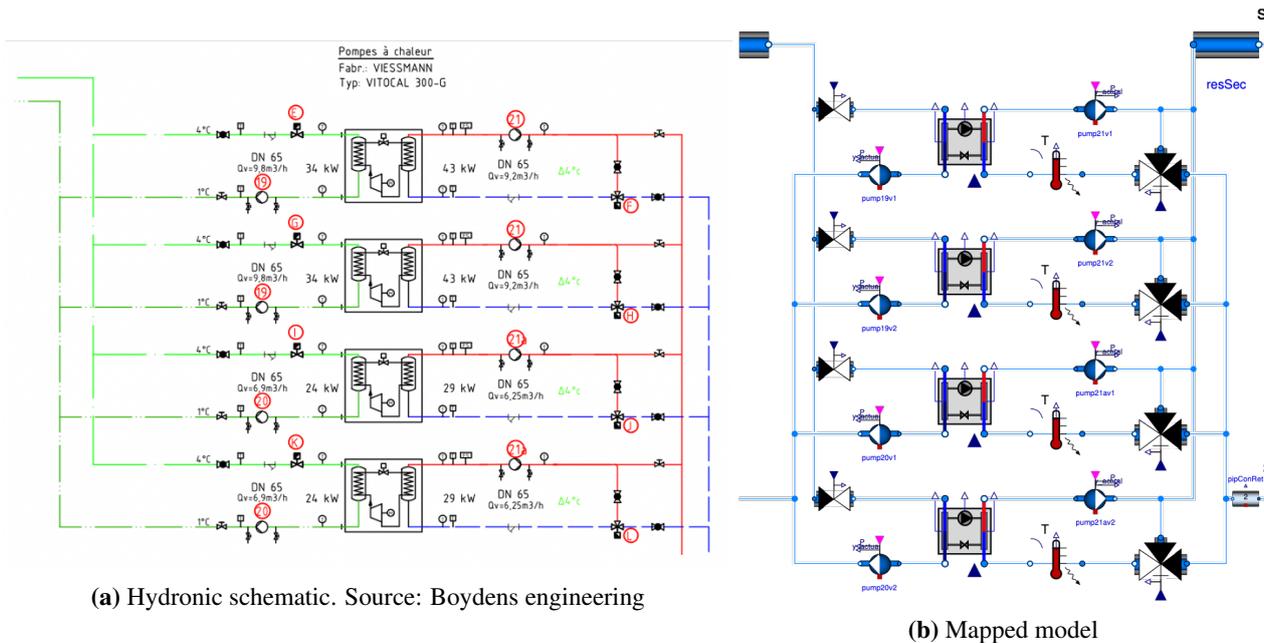
of the building envelope models and equations, we refer to the Modelica IDEAS library (Filip Jorissen, Reynders, et al. 2018; KU Leuven and 3E 2012) and the Modelica Buildings library (Michael Wetter et al. 2014; Wetter et al. 2019). The models are parameterised using parameters that are commonly available in the technical specifications of the building HVAC equipment. Instantiating the Modelica component models and making the required connections between them is a tedious and error-prone process. Therefore a browser-based tool has been developed that automatically generates the required Modelica/IDEAS code. This graphical user interface (GUI) only allows valid building geometries to be specified, thereby avoiding user errors. The GUI export of this geometrical information is automatically mapped into the Modelica model as illustrated in Figure 2. Building geometry information, such as orientations and surface areas, is automatically deduced from the export, while the user specifies additional information such as material layers through the GUI options.

Code generation of connections is automated, which rules out errors such as unconnected ports or ports with too many connections. Furthermore, sanity checks are performed such as identification of unconnected devices, (un)realistic thermal insulation values in outer walls, missing parameters, etc.

The resulting workflow leads to a set of interconnected Modelica component models that describe the relevant physics:

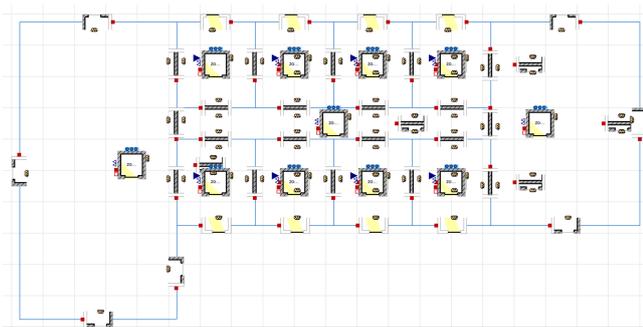
- Thermal conduction, convection and radiation within and outside of the building envelope,
- Thermal inertia of the building,
- Solar heat gains (considering glazing type and shading) and internal heat gains from occupants,
- Pressure-driven flow rates for aerolic, hydronic system, including non-linear valve and damper models,
- Pumps and fan powers,
- Temperature dependent and mass flow rate dependent heat flow rates in emission systems,
- Temperature dependent and flow rate dependent efficiencies in production and distribution systems.

The resulting model is non-linear. The most important non-linearities are the relation between flow rate and pressure in fluid flow networks, and the strongly non-linear relation between fan/pump power and flow rate. For more details about these models see (Filip Jorissen, Reynders, et al. 2018; F. Jorissen, Boydens, and L. Helsen 2019; Filip Jorissen, Michael Wetter, and Lieve Helsen 2018; F. Jorissen, Boydens, and L. Helsen 2017; Filip Jorissen 2018).



(a) Hydraulic schematic. Source: Boydens engineering

(b) Mapped model

Figure 1. Illustration of one-to-one component to model mapping**Figure 2.** Automatically generated Modelica model. The drawn icons depict zones, windows, exterior walls and interior walls. Note that window icons are drawn on top of wall icons.

That Modelica model is translated into an optimisation code using our Toolchain for Automated Control and Optimisation (TACO) (F. Jorissen, Boydens, and L. Helsen 2018), which extends the JModelica framework (Åkesson et al. 2010). TACO identifies what equations must be solved for what variables, splits variables that depend on optimisation variables from those that do not, and performs preprocessing on linear state¹ interdependencies to speed up code evaluation. The continuous time problem is discretised at a user-defined non-equidistant set of points in time. CasADi (Andersson et al. 2019) computes the equation derivatives and generates C-code for evaluating the objective, constraints, derivatives and other outputs. The compiled code is coupled to a gradient-based NLP solver and is portable to other (linux) machines. For more

¹We distinguish state variables that are computed by a differential equation from algebraic variables that are computed from an algebraic equation.

details about the translation process that TACO performs we refer to (F. Jorissen, Boydens, and L. Helsen 2018).

3 Case study

In this paper we present a case study building, Solarwind, on which the presented methodology is applied. The office building has a conditioned floor surface area of 10 000 m² and was designed to be an exemplary showcase towards Luxembourg and the European design and construction industry of a holistic and integrated operation and design sustainability approach for future oriented office buildings. It uses geothermal heat pumps, concrete core activation (CCA), solar collectors, solar PV, passive cooling, indirect evaporative cooling, a pellet boiler, a large storage tank, etc. This building has been described in detail in Chapter 2 of (Filip Jorissen 2018). Chapter 10 describes an MPC for (a part of) the same building. That model is substantially smaller and less detailed and complex than the OCP that we present here. It also required manual building-specific simplifications in the HVAC models that conflict with the 1-to-1 mapping philosophy and the expertise requirements that were outlined above. We now explain the building and its model together with the main differences from the earlier MPC implementation.

3.1 Building envelope

The new model consists of six floors (instead of four) that are modelled using 79 zones (instead of 32). We assume that 6 people are present in each zone, between 7:00 and 19:00 on week days. Three zones have 9 occupants instead of 6, to show the influence on the CO₂ concentration in the result section. Solar shading is not modelled to artificially increase the heat load and to make the optimisation problem more challenging to solve. A weather data

file for Uccle, Belgium is used. The building uses triple glazing and is strongly insulated with a U-value of about $0.1 \text{ W/m}^2\text{K}$.

3.2 Emission system

Ventilation is provided using six air handling units (instead of 2), for which the supply and extraction fan pressure, the heating coil valve, humidifier, heat recovery bypass (2 dampers), indirect evaporative heat exchanger control signal and active chiller control signal are optimised. Most zones have a supply Variable Air Volume (VAV) and an extraction VAV for which one control signal is optimised. A VAV control signal of 0 % corresponds to a *set point* of 50 % of the nominal flow rate². Furthermore, each VAV has a heating coil for which the heating fluid flow rate is optimised using a two-way valve. The supply water temperature of the VAV coil collector is controlled using a three-way valve.

The concrete ceilings are heated or cooled using CCA, except for the top floor, which uses chilled ceilings (CC) instead. Each floor is subdivided in about six CCA or CC sections. Each section spans one or more zones, is connected to one of three collector connections (south, north, or top floor), and its total flow rate is controlled using one two-way valve. The supply temperature of each connection (i.e. group of sections) is controlled at the collector using a three-way valve.

3.3 Production system

The main collector draws water from either a geothermal heat pump, or the geothermal borefield using a series of pumps and heat exchangers (see Figure 2.3 in (Filip Jorissen 2018) for more details).

A pellet boiler and solar collector feed hot water in a 20 m^3 storage tank, which is used by the AHU heating coils. The corresponding fluid loop consists of 2 heat exchangers, 6 pumps, 3 three-way valves and a pressure-independent valve. Three domestic hot water tanks that are present in the building are not modelled for this study since realistic load profiles are not available.

3.4 Objective and constraints

We minimize the electrical power use and the pellet boiler thermal power, the latter being scaled by 1/3 to consider that pellet fuel has a different price than electricity. Various constraints are enforced, among them the minimum and maximum supply air temperature ($16 \text{ }^\circ\text{C}$ - $26 \text{ }^\circ\text{C}$), (building owner specified) zone temperature limits ($22 \text{ }^\circ\text{C}$ - $24 \text{ }^\circ\text{C}$) and a CO_2 concentration upper limit of 1000 ppm.

3.5 Horizon

For the purpose of this paper, we discretise the model in 720 intervals of 1 hour (i.e. 1 month). The resulting OCP is instantiated 12 times to optimise a full year. We use an OCP since it simplifies interpretation of the results compared to an MPC running in a receding horizon fashion.

²The set point may not be obtained if the fan pressure is too low.

Each OCP has the same initial state. For easier result interpretation we assume that all HVAC equipment is enabled 24/7.

4 Results

The resulting model has 217 control inputs, 3421 state variables and 41 577 algebraic variables as reported by the Modelica simulation software Dymola 2020. Since the OCP has 720 control intervals, 156 240 control inputs are optimised in each OCP. Computation time for one OCP (single-core, 2 GHz) is a few days, depending on the chosen convergence tolerance. Note that an MPC using this model is orders of magnitude faster since its horizon is shorter and warm starting can be used. At the time of writing, a revised version of this model is successfully controlling the modelled building, where the MPC is updated every 15 minutes. Computation speed is thus not a problem, even for white-box models of this size.

Figures 3 - 6 present OCP results for January, April and August. We use these figures to illustrate some of the strengths of our detailed white-box MPC approach, without discussing each sub-plot in detail. Despite that these are in fact OCP results instead of MPC results, we consider that the results are representative for MPC. The presented results are direct outputs of the OCP. This is a first strength: since the model is detailed, the outputs are detailed too, which allows a thorough analysis of the results without requiring additional simulations to see the impact on subsystems that are lumped in the optimisation.

4.1 System coordination and constraints

In the top sub-plot of each figure, relevant zone temperatures are indicated, while CO_2 constraints are also indicated in the top of Figure 3. The results show that the operational constraints of $22 \text{ }^\circ\text{C}$ - $24 \text{ }^\circ\text{C}$ are respected, as well as the upper limit of 1000 ppm CO_2 . At the same time, the third sub-plot in Figure 3 indicates that the AHU fans usually operate at low pressures of 30 Pa while the nominal system pressure is around 300 Pa. At the same time, most VAVs are at or around the minimum opening of 0%. However, the VAVs that correspond to zones with a larger occupancy are occasionally opened to avoid violating the CO_2 constraint. During the summer period, the fan pressures are occasionally increased to accommodate the peak cooling load due to solar heat gains. This illustrates the second strength of MPC: the coordination of multiple devices to ensure that constraints are achieved at minimum cost.

4.2 System dynamics

A third strength of MPC is illustrated in Figure 4: its ability to anticipate heating and cooling loads and to deal with them accordingly using fast (VAV) and slow (CCA) reacting systems. During the August period, large heat gains are present since we deliberately omitted solar shading in the model. This causes both the upper and lower temperature bounds to be reached within the same day (and in

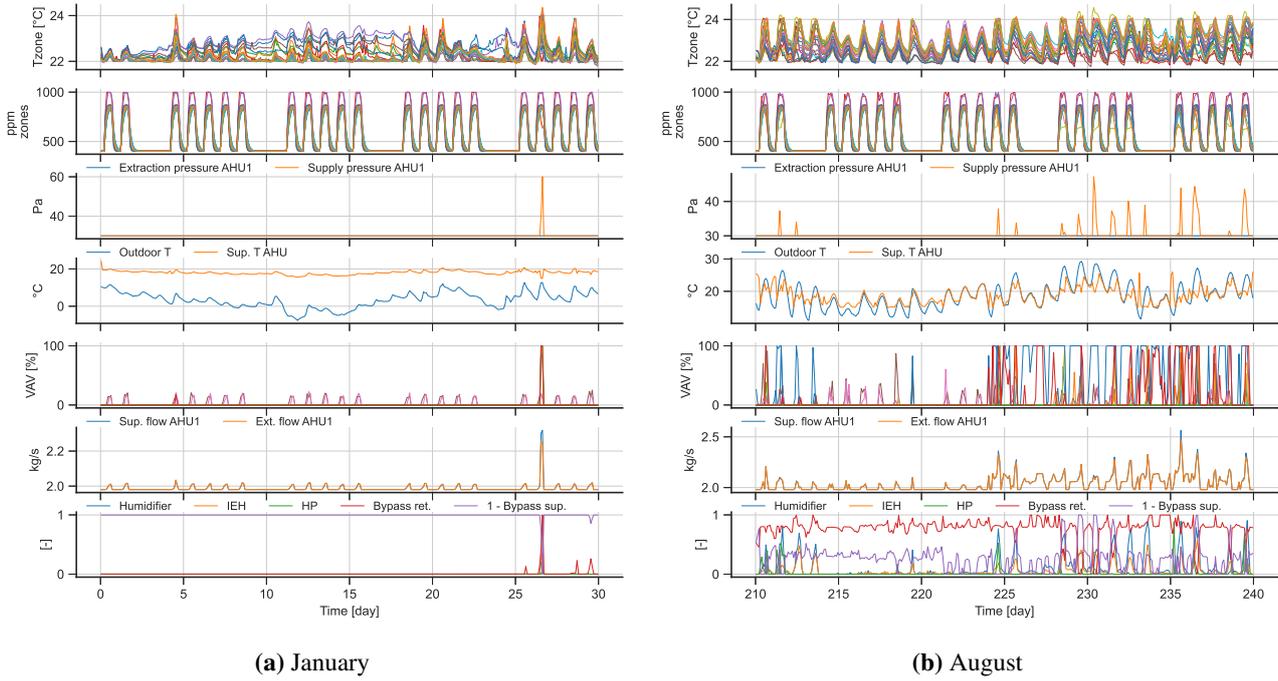


Figure 3. Results that are relevant to air handling unit 1

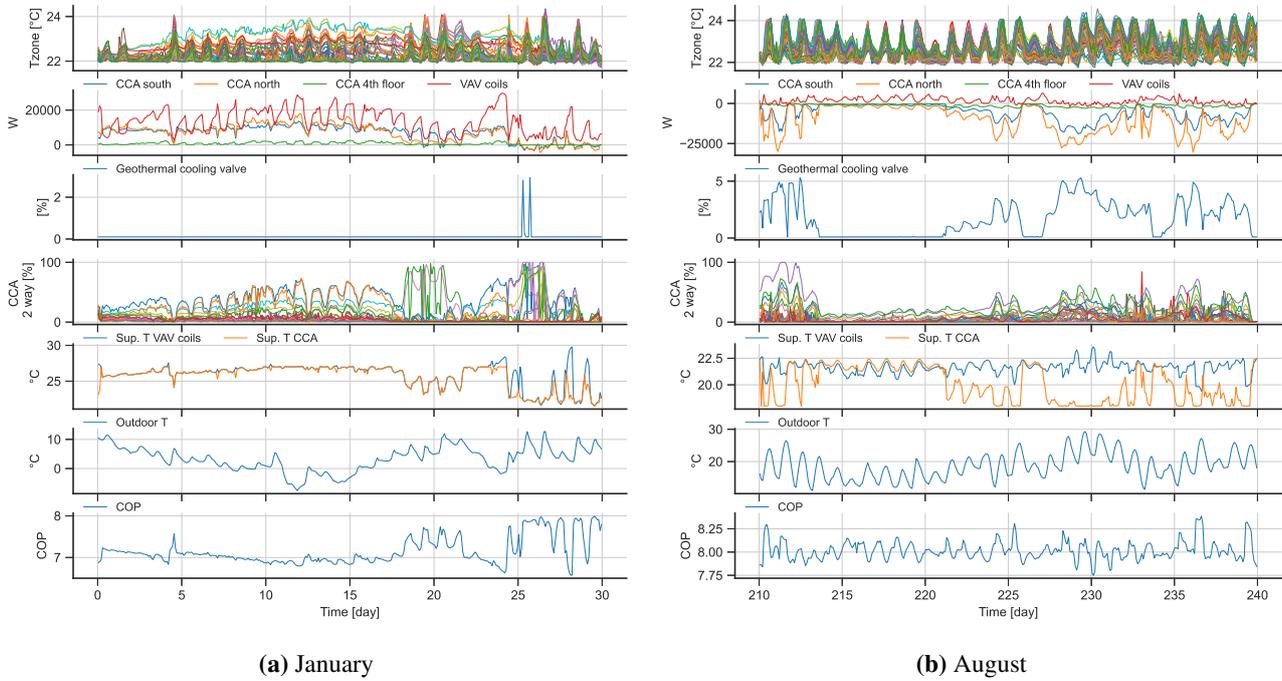


Figure 4. Heat pump, geothermal cooling, CCA and VAV operation.

different zones). The cooling is spread over the entire day and the cooling peak is out of phase with the temperature peaks, see sub-plots 1 and 2. The building is thus pre-cooled, thereby considering the emission system delay.

This is illustrated once more during the winter period around day 25, when there is a period of large solar irradiation. During the three days preceding the solar peak, the CCA thermal power is relatively low and shifted towards the VAV coils to avoid excess heat being stored in the building thermal mass, which would overheat the building during the sunny period.

4.3 Hybrid systems

A fourth strength of MPC is its ability to coordinate between multiple heating and cooling sources and to use the one with the lowest cost. This is illustrated in Figure 3. During winter, the AHU bypasses (see last sub-plot) are both closed since heat is valuable during winter. During summer, 4 AHU cooling options exist (free-cooling bypasses, humidifier³, indirect evaporative heat exchanger (IEH), chiller (HP)), which are all used. Even when no immediate cooling is required and the outdoor temperature is low, the AHU bypasses are opened. This cooling option *reduces* the fan power since the AHU internal pressure drop decreases. Note that the lower supply air temperature limit of 16 °C is respected. When the outdoor temperature is higher than the indoor temperature, *cold* is recovered by *closing* the bypasses. When this does not suffice, the humidifier and IEH are used. They have the lowest cost since they only require a pump to be operated. In some cases, the chiller is used, which is the most expensive cooling option. Thanks to the good coordination of all other devices, this rarely happens.

Another example is shown in Figure 5, which shows the operation of the solar collector, pellet boiler and storage tank. During winter, the sun is at a low altitude due to which the collector heat losses are often larger than the solar heat gains. The solar thermal collector valve (see sub-plot 5) is therefore only opened when the sun intensity is sufficiently strong to reach a positive thermal power (see sub-plot 2). The remaining high-temperature heat load is provided by the pellet boiler. During the summer period, there is abundant (free) heat available from the solar collector due to which the pellet boiler is never activated.

4.4 Operational limits

A fifth strength of MPC is its ability to operate the system at its limits. For instance, the geothermal borefield can be used to passively cool the building using the CCA. A borefield temperature of 15 °C is assumed. Directly using this low temperature would be the least costly since less mass flow rate (and thus pump power) is required to achieve the same heat flow rate. However, condensation of moist air can occur on CCA when its temperature is low. Therefore, a minimum supply water temperature of 18 °C

³The OCP does not contain a humidity constraint due to which the humidifier can be used to cool.

was set. Figure 4 clearly shows that this minimum supply water temperature is used during periods of large cooling load (see sub-plot 5).

4.5 Efficient operation

A sixth strength of MPC is its ability to operate the available equipment at an efficient operating point. Figure 4 illustrates this for the heat pump operation in the three last sub-plots. For lower outdoor temperatures, larger heating powers are required. This increases the required heat pump supply water temperature, which reduces the heat pump COP (excluding pump power). Note day 4, where a day of large solar intensity (see sub-plot 2 of Figure 5) ‘charges’ the building, immediately reducing the heating requirements, increasing the COP and also increasing the COP during the days after the heating event, despite the decreasing outdoor temperatures.

4.6 Exploiting flexibility

The seventh MPC strength is to shift heating loads using the available system flexibility. We already discussed pre-cooling using the CCA. Note that this effect could have been more pronounced if the building solar heat gains were smaller or if the comfort band were larger than 2 K, allowing more drift of the indoor temperature. Hitting both the upper and the lower comfort bound within the same day limits the available system flexibility.

Additionally, the hot water storage tank flexibility is used, which is allowed to fluctuate between 50 °C and 90 °C. E.g. during the winter period in Figure 5 solar heat is accumulated for three weeks until the last days of the month when the heat is most useful. During summer, solar heat is abundantly available and is even dissipated by not closing the solar collector valve at night. Otherwise the storage tank upper temperature limit of 90 °C would be violated at day 233. Figure 6b shows the solar collector operation in April, where solar heat is stored between sunny (see sub-plot 2) days 95 and 107 and used during overcast days 108, 109, 117 and 118.

4.7 Reliable performance

Finally, we discuss the system operation during the intermediate season, when well insulated buildings are often hard to control. For instance, on cold but sunny days the building can be overheated by conventional heating curve-based controllers that increase heating set points despite substantial solar heat gains. It is hard to determine whether the building should be in heating mode, cooling mode, or perhaps even in both within the same day. It is interesting to see how MPC copes with such scenarios, which is illustrated for the month of April in Figure 6. Our MPC controller does not formally define a ‘mode’ but based on the second sub-plot left we conclude that MPC is in heating mode during the first two days and in cooling mode for three days. During the remainder of the month, the system is in a neutral mode despite the outdoor temperature ranges from 0 °C to 20°C. In this neutral mode,

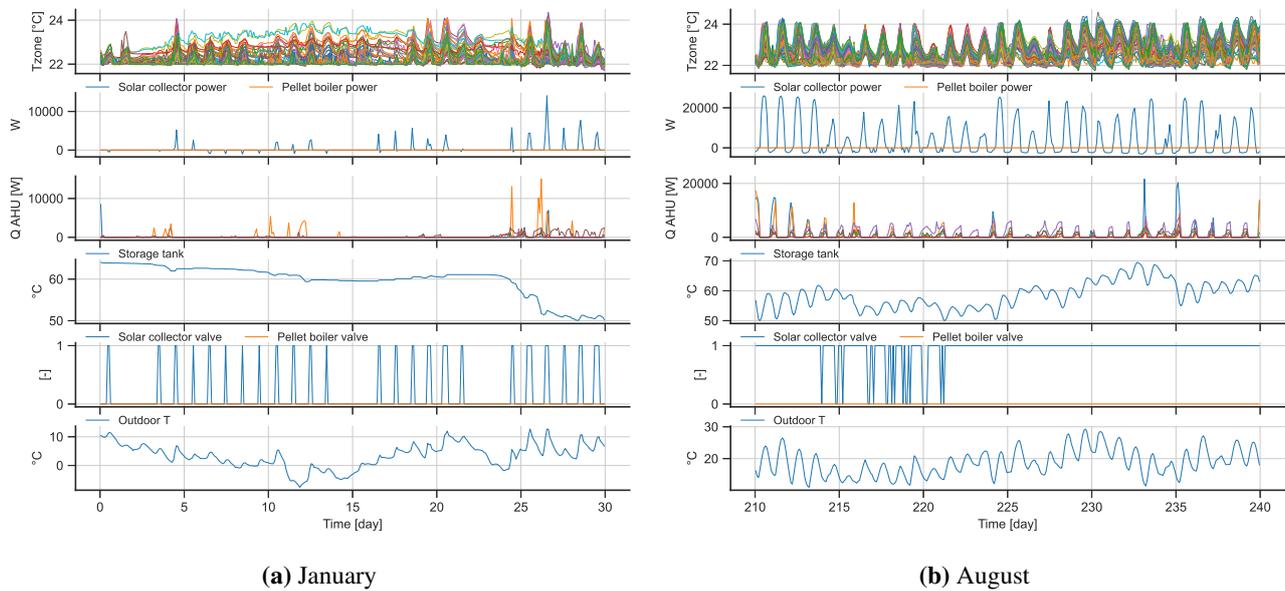


Figure 5. Solar collector and pellet boiler operation.

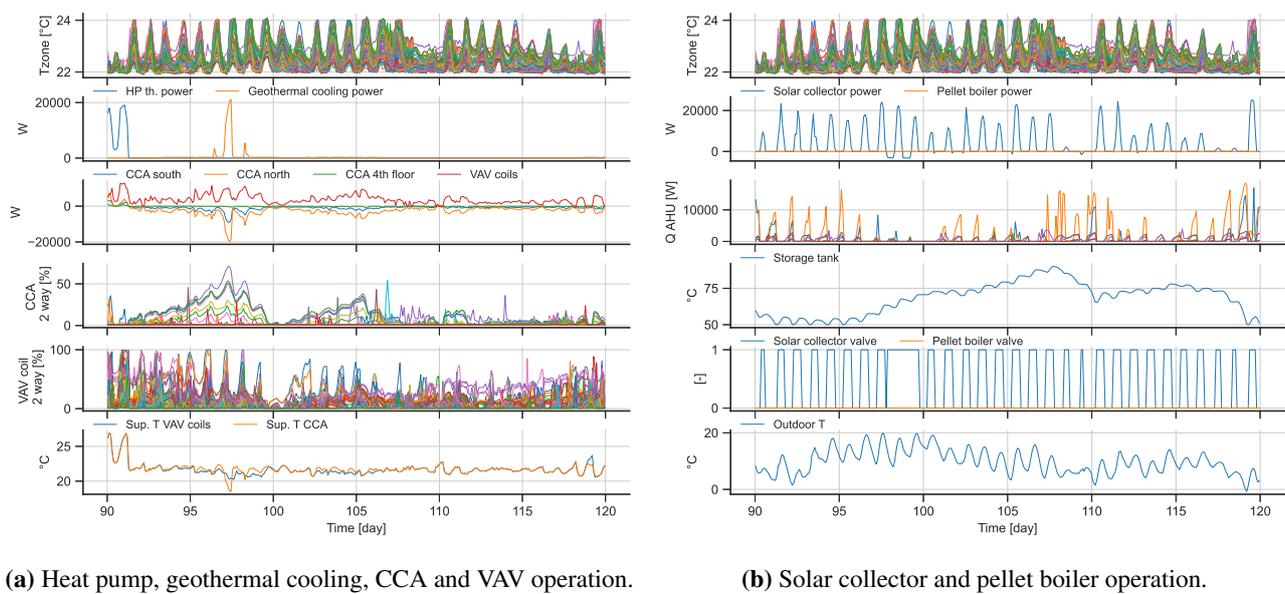


Figure 6. System operation in April.

the system is *simultaneously heating and cooling by exchanging heat within the system* (see sub-plot 3 in Figure 6a). More specifically, the relatively warm CCA return water is recirculated through the VAV heating coils, heating the supply air. The fifth sub-plot shows that the CCA heat extraction is focussed on just a few CCA circuits and the sixth sub-plot shows that heat dissipation does not use all VAV circuits. This operating mode is particularly efficient since fewer pumps and even the heat pump need not be enabled. Perhaps more importantly, the low temperature differences allow simultaneous heating and cooling in different zones and the comfort constraints are respected. Furthermore, heat is also exchanged between warm and cold CCA circuits, but this is not indicated in the graph. Therefore, the total heat exchange is even larger than indicated. Interestingly, this operating mode is automatically discovered thanks to the level of detail of our white-box MPC.

5 Conclusion

This paper presents a detailed white-box MPC approach for buildings that is designed for commercial MPC applications, including large and complex buildings. The approach maps physical objects and devices into their respective Modelica models using a custom browser-based graphical user interface. The resulting Modelica model is translated into an efficient MPC code using TACO. The approach is applied to a case study office building of 10 000 m², resulting in a model of 79 zones. A discussion of OCP results for three individually optimised months shows the strengths of the toolchain. All presented strengths are automatically achieved by our white-box OCP, without substantial tuning or training. Each optimisation consists of 720 intervals of one hour, resulting in 156 240 control inputs and 29 935 440 algebraic variables, which demonstrates the feasibility of large-scale white-box optimisation. The presented case study is currently being controlled by MPC, hence also demonstrating the computational feasibility of this approach for direction optimisation. Furthermore, the presented OCP approach can be compared to other control methodologies, e.g. within the frame of IBPSA Project 1, WP 1.2 BOPTTEST (Blum et al. 2019). Future work will present real-life, operational results of our MPC approach.

Acknowledgements

We thank Boydens Engineering, Schuler SàRL, DRC Technology Sa, Menenga nv and Climalux Sa for enabling this research. We acknowledge the funding through the VLAIO spin-off fellowship HBC.2018.2092 and through the EU within the H2020-EE-2016-RIA-IA programme for the project ‘Model Predictive Control and Innovative System Integration of GEOTABS;-) in Hybrid Low Grade Thermal Energy Systems - Hybrid MPC GEOTABS’ (grant number 723649). We acknowledge the university of Leuven (KU Leuven) for funding the development

and valorization track of white-box MPC through the C3 project C32/18/063. This work emerged from IBPSA Project 1, an international project conducted under the umbrella of the International Building Performance Simulation Association. Project 1 develops and demonstrates a BIM/GIS and Modelica Framework for building and community energy system design and operation.

References

- Åkesson, J. et al. (2010). “Modeling and optimization with Optimica and JModelica.org – Languages and tools for solving large-scale dynamic optimization problems”. In: *Computers & Chemical Engineering* 34.11, pp. 1737–1749. DOI: 10.1016/j.compchemeng.2009.11.011.
- Andersson, Joel A. E. et al. (2019). “CasADi: a software framework for nonlinear optimization and optimal control”. In: *Mathematical Programming Computation* 11.1, pp. 1–36. DOI: 10.1007/s12532-018-0139-4.
- Blum, D. et al. (2019). “Prototyping the BOPTTEST framework for simulation-based testing of advanced control strategies in buildings.” In: *Proceedings of the 16th International Conference of IBPSA*. Rome, Italy.
- Drgoňa, Ján et al. (2020). “All you need to know about model predictive control for buildings”. In: *Annual Reviews in Control* 50, pp. 190–232. ISSN: 1367-5788. DOI: 10.1016/j.arcontrol.2020.09.001.
- Frank, Stephen et al. (2016). “Hybrid model-based and data-driven fault detection and diagnostics for commercial buildings”. In: *2016 ACEEE Summer Study on Energy Efficiency in Buildings*. Pacific Grove, CA.
- International Energy Agency (2019). *Global Status Report for Buildings and Construction: Towards a Zero Emissions, Efficient and Resilient Buildings and Construction Sector*. Tech. rep.
- Jorissen, F., W. Boydens, and L. Helsen (2017). “Validated air handling unit model using indirect evaporative cooling”. In: *Journal of Building Performance Simulation* 11.1, pp. 48–64. DOI: 10.1080/19401493.2016.1273391.
- Jorissen, F., W. Boydens, and L. Helsen (2018). “TACO, an Automated Toolchain for Model Predictive Control of Building Systems: Implementation and Verification”. In: *Journal of Building Performance Simulation* 12.2, pp. 180–192. DOI: 10.1080/19401493.2018.1498537.
- Jorissen, F., W. Boydens, and L. Helsen (2019). “Implementation and Verification of the Integrated Envelope, HVAC and Controller Model of the Solarwind Office Building in Modelica”. In: *Journal of Building Performance Simulation* 12.4, pp. 445–464. DOI: 10.1080/19401493.2018.1544277.
- Jorissen, Filip (2018-04). “Toolchain for Optimal Control and Design of Energy Systems in Buildings”. PhD thesis. Arenberg Doctoral School, KU Leuven.
- Jorissen, Filip, Glenn Reynders, et al. (2018). “Implementation and Verification of the IDEAS Building Energy Simulation Library”. In: *Journal of Building Performance Simulation* 11.6, pp. 669–688. DOI: 10.1080/19401493.2018.1428361.
- Jorissen, Filip, Michael Wetter, and Lieve Helsen (2018). “Simplifications for Hydronic System Models in Modelica”. In: *Journal of Building Performance Simulation* 11.6, pp. 639–654. DOI: 10.1080/19401493.2017.1421263.
- KU Leuven and 3E (2012). *IDEAS*. <https://github.com/openideas/IDEAS>.

- Sturzenegger, David et al. (2016). “Model Predictive Control of a Swiss Office Building: Implementation, Results, and Cost-Benefit Analysis”. In: *IEEE Transaction on Control Systems Technology* 24.1, pp. 1–12. DOI: 10.1109/TCST.2015.2415411.
- Wetter, Michael et al. (2014). “Modelica Buildings library”. In: *Journal of Building Performance Simulation* 7.4, pp. 253–270. DOI: 10.1080/19401493.2013.765506.
- Wetter, M et al. (2019-09). “IBPSA Project 1: BIM/GIS and Modelica framework for building and community energy system design and operation – ongoing developments, lessons learned and challenges”. In: *IOP Conference Series: Earth and Environmental Science* 323, p. 012114. DOI: 10.1088/1755-1315/323/1/012114.