

Multi-Objective Model Predictive Control Framework for Buildings

Krzysztof Arendt¹, Anders Clausen¹, Claudio G. Mattera¹, Muhyiddine Jradi¹,
Aslak Johansen¹, Christian T. Veje¹, Mikkel B. Kjægaard¹, Bo N. Jørgensen¹

¹Center for Energy Informatics, University of Southern Denmark, Odense, Denmark

Abstract

The aim of this paper is to present the implementation and performance of an MPC framework based on a multi-objective genetic algorithm. The framework optimizes building control by firstly identifying the Pareto frontier with respect to multiple objectives considered, and then selecting the final strategy based on the user-defined priorities for the respective objectives. Although the approach requires more computing resources than the more traditional constrained convex optimization, it is more flexible in terms of the optimization problem formulation. New objectives can be easily added, and the objective priorities altered during the operation of the system. This flexibility makes the framework attractive for global optimization of multiple systems, including systems based on on/off control. The framework is compatible with the Functional Mock-Up Interface and uses models exported to Functional Mock-Up Units. The framework performance is tested in a virtual experimental testbed using a building modeled in EnergyPlus.

Introduction

The building sector is responsible for approximately 40% of the energy consumption in U.S. and European Union (Cao et al. (2016)). In the residential and commercial buildings, approximately half of the energy is consumed by the heating and cooling systems (D&R International (2012)). These facts lead to ongoing efforts across the world to manage resources in buildings in a more efficient manner. One promising approach identified by researchers, is Model Predictive Control (MPC). MPC refers to a family of control strategies in which there is a direct use of physical modeling (García et al. (1989)). MPC uses model-based simulation to predict the effects of control strategies. This feature is especially attractive for systems with high inertia, like buildings with high thermal mass (Li and Malkawi (2016); Chen (2001)). MPC can also enable demand side flexibility to buildings.

Contrary to MPC, the traditional control systems (e.g. based on PID) are reactive in nature and their

function is limited to maintaining some set of desired indoor environment parameters, expressed through setpoints. In this approach, the building's thermal inertia is considered as a disturbing factor, having a negative effect on the control performance. Hence, traditional approaches are unable to take the advantage of building dynamics and optimize for passive heating and cooling that relies on building thermal mass. However, this can be achieved with MPC.

MPC can be implemented in buildings on two different levels: (1) low-level control of actuators, and (2) high-level control of setpoints. The low-level MPC, as in Karlsson and Hagetoft (2011), can potentially replace PID in controlling specific system components. The high-level MPC, as in Huang et al. (2015b) or Li and Malkawi (2016), is sometimes referred to as "supervisory", and can supplement Building Automation Systems (BAS).

The models used in MPC fall into three broad categories: (1) white-box models, (2) grey-box models, and (3) black-box models. In the white-box models the physical phenomena are captured explicitly in model equations. An exemplary application of white-box models (implemented using EnergyPlus) in MPC was presented in Ascione et al. (2016), Coffey et al. (2010), Kwak et al. (2015), and Kwak and Huh (2016). The white-box approach is the most physically sound category, but also the one with the highest requirements to data about the building. The amount of detailed data required by white-box models is one of the reasons why they are used mostly in virtual case studies. The grey-box models capture some basic physics, but some or all of the model parameters have to be estimated based on the measured data. The most frequently used grey-box modeling approach in MPC is the resistor-capacitor (RC) thermal network, as in De Coninck and Helsen (2016), Hazyuk et al. (2012), Hazyuk et al. (2014), and Maaoumy et al. (2014). Finally, the black-box models are entirely data-driven (e.g. statistical and machine learning models) as in Garnier et al. (2014, 2015) and Huang et al. (2015a). Sometimes the grey- and black-box models are combined, as in Huang et al. (2015b)

(RC and neural networks combined). In actual buildings, especially the grey- and black-box approaches are promising, since they can be generic and portable between buildings.

Up to recently, the computational demand of models and the technical complexity of the resulting cyber-physical system were considered prohibitive for real implementation of MPC. MPC adoption suffers also from the difficulty to automate building model derivation and the lack of adequate data for real buildings. The complexity of MPC is especially significant if all building systems are to be controlled in a coordinated manner. The difficulty of this overall approach lies not only in the need for multi-domain modeling, but also in the integration of all communication protocols, many of which are proprietary (Ogunsola et al. (2014)). Therefore, in the majority of studies the MPC performance was tested on virtual buildings (Ascione et al. (2016); Bianchini et al. (2016); Coffey et al. (2010); Garnier et al. (2014, 2015); Li and Malkawi (2016); Morosan et al. (2010); Chen (2001); Hazyuk et al. (2012, 2014); Killian et al. (2016); Kwak et al. (2015); Kwak and Huh (2016)). MPC has been implemented only in a handful of real buildings so far, e.g. De Coninck and Helsen (2016), Huang et al. (2015b), Dong and Lam (2014).

In most of the MPC studies the optimization problem is described in terms of the minimization of a cost function, which is defined as the weighted sum of different objectives (Bianchini et al. (2016); Coffey et al. (2010); De Coninck and Helsen (2016)). The optimization problem is often solved using linear programming (e.g. Schirrer et al. (2016)), nonlinear programming (e.g. De Coninck and Helsen (2016)), or meta-heuristic evolutionary algorithms, e.g. genetic algorithms (e.g. Coffey et al. (2010)) or particle swarm optimization (e.g. Li and Malkawi (2016)).

The actual energy saving potential depends on the quality of the reference control strategy. In example, the MPC system tested on a virtual building by Garnier et al. (2015) was compared with five reference non-predictive strategies and it was found that two of them provided up to 20% lower energy consumption than MPC, while three of them provided higher energy consumption (up to the factor of three). MPC outperformed, however, the traditional strategies in terms of thermal comfort. For reasonable reference strategies, the MPC potential for saving energy has frequently been reported to be below 30% (Huang et al. (2015a); Li and Malkawi (2016); Morosan et al. (2010); Chen (2001); Hazyuk et al. (2012); Zhao et al. (2015)), based on virtual test cases. As for real buildings, the cost savings due to implementation of MPC were reported to be around 34-40% (De Coninck and Helsen (2016)), 18-30% (Dong and Lam (2014)), and 13% (Huang et al. (2015b)). Huang et al. (2015b), however, stated that higher cost savings could be pos-




Figure 1: Virtual test framework.

sible if a detailed occupancy profile prediction was available.

Summarily, the conclusion from the presented literature review is that MPC has the potential to decrease building energy consumption, but it is technically difficult to implement due to the highly multi-domain, multi-disciplinary, comprehensive approach needed. In addition, typically the MPC systems deployed in particular buildings are not easily portable to other ones.

In this paper we present a new multi-objective MPC framework which addresses some of the issues reported in the above-mentioned studies. First, thanks to the multi-objective optimization, there is no need to normalize and lump many objectives into a single cost function. Secondly, the framework uses generic zone models based on the gray-box approach and automatic parameter estimation, making it easier to deploy in a new building. Finally, due to the adoption of the Functional Mock-up Interface (FMI) (Blochwitz et al. (2011)), the framework is not tied to any particular modeling approach or software environment. The performance of the framework is planned to be conducted in both virtual and physical setups. This paper presents the initial results of the virtual test case.

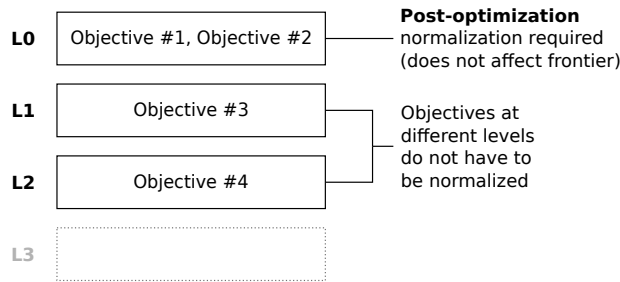
Optimization framework

The optimization framework consists of the following parts (Figure 1):

- Controlem – optimization toolbox (Sørensen and Jørgensen (2017)),
- ModestPy – Functional Mock-up Unit parameter estimation toolbox (Arendt et al. (2018)),
- simulation interface compliant with the Functional Mock-up Interface,
- archiver with the sMAP interface (Dawson-Haggerty et al. (2010)).

Controlem is a multi-objective genetic algorithm (MOGA) that constructs a Pareto frontier with respect to the considered objectives (Figure 2). Each

Priority levels:



Recursive policy selection:

1. Select L0 frontier (subset of the full frontier)
2. Select L0+L1 frontier (subset of L0)
3. Select L0+L1+L2 frontier (subset of L0+L1)
4. ...
5. One or more policies left (equally optimal solutions)

Pareto frontier:

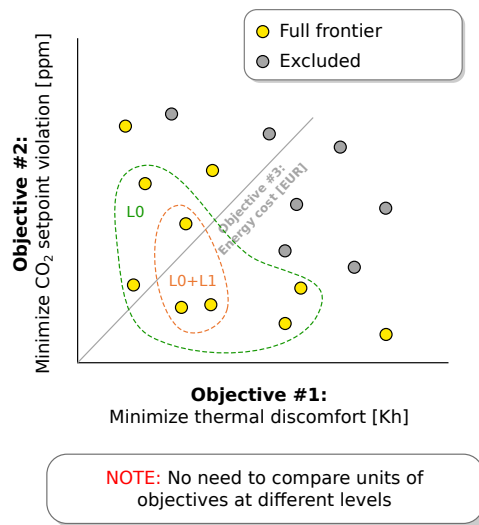


Figure 2: Pareto frontier in the multi-objective genetic algorithm optimization.

individual in the population represents a specific control policy, e.g. specific heating and ventilation profiles, that is tested in simulation. The Pareto frontier construction is iterative and based on the genetic algorithm operations: selection, crossover, and mutation. Whenever a new individual appears in the evolution, with a better fitness with respect to at least one of the objectives, it joins the Pareto frontier. When one of the stopping criteria is met, the evolution stops and the algorithm proceeds to the second phase in which the final policy is selected. There are two stopping criteria in use: (1) no improvement in the Pareto frontier for a defined number of generations, (2) maximum computational time reached. According to the authors' knowledge this work is the first example of applying Controlem to MPC in buildings. It is also one of very few examples of applying MOGA to MPC in buildings in general.

The policy selection in the second phase is conducted recursively. In each step i the subset of individuals optimal with respect to the priority level L_i (Figure 2) is selected. Each priority level can contain either one or more objectives. In the case there are two or more objectives at the same level, the objectives are normalized in order to identify the optimal population subset. However, no normalization is required for priority levels with a single objective. Finally, after traversing through all the levels, one or more equally optimal policies are left, out of which one is selected randomly.

The developed optimizer (MOGA) communicates with the archiver through the sMAP interface (Dawson-Haggerty et al. (2010)). sMAP allows for a two-way communication with the building, i.e. for reading measurements and setting new setpoints. Therefore, sMAP provides a higher level of abstrac-




Figure 3: (a) Simplified test model limited to 7 zones, (b) actual OU44 building.

tion and from the optimizer point of view, there is no difference between communicating with an EnergyPlus model (as in this work) or with a real building. Due to the FMI-compatibility and the use of MOGA the framework is essentially model independent. The models can be implemented in any FMI-compliant tool, and they can be non-linear, non-differentiable or even non-continuous. In addition, since the optimization is not based on a cost function, adding new objectives is straightforward. The objectives do not have to be normalized with respect to one another. Therefore, Controlem is potentially more flexible than MPC systems based on collocation (Magnusson and Åkesson (2015)) or shooting methods (Lazutkin et al. (2014)). Such features might be especially relevant for building systems, which are often non-linear (e.g. HVAC) and non-continuous (e.g. on/off controllers). On the other hand, MOGA is expected to be more computationally demanding than the gradient-based methods.

Experimental setup

The framework was tested on a virtual case study building modeled in EnergyPlus (similarly to e.g. Ascione et al. (2016) or Bianchini et al. (2016)). The

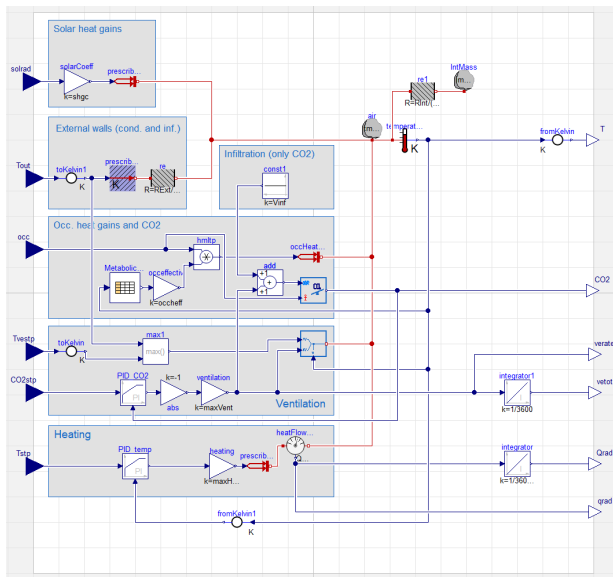


Figure 4: Gray-box model of a zone based on the R2C2 thermal network (Modelica).

virtual building is a downscaled version of the OU44 teaching building at the SDU Campus Odense (Figure 3). It has the same HVAC system type and the building envelope construction as the actual building, but it is limited to just 7 thermal zones. The thermal zones have a similar geometry as the classrooms in OU44. The building is equipped with a hydronic heating loop and a mechanical ventilation system with pre-heating and a heat exchanger unit.

Controleum uses 7 gray-box zone models to simulate the effects of control policies on the thermal zones in the building. The zone models are based on the RC thermal network, and each zone model has the same structure (Figure 4), but different parameters. The models are implemented in Modelica (Elmqvist et al. (1998)). The zone model parameters were estimated by minimizing the error in indoor temperature and CO2 compared to the EnergyPlus outputs using the ModestPy toolbox (Arendt et al. (2018)).

The performance of the framework is compared with the rule-based control (RBC) in a one-month long simulation, based on the climate data for January from Typical Meteorological Year for Copenhagen. It is assumed that the framework has control over the room temperature setpoints (each room can have a different setpoint) and has access to room occupancy schedules. In the real applications the occupancy schedules would be replaced with occupancy predictions. The influence of the quality of the occupancy predictions on the MPC performance is not considered in this study. Two scenarios are considered:

1. CTRL-EE – optimization of temperature setpoints to minimize energy consumption and maintain indoor thermal comfort,
2. CTRL-DK1 – optimization of temperature set-

points to minimize energy cost and maintain indoor thermal comfort.

In both scenarios the maintenance of indoor thermal comfort has the highest priority. In the RBC strategy the temperature setpoints are scheduled to 20 °C during weekdays between 5:00-16:00, and 15 °C otherwise. The RBC strategy is implemented directly in EnergyPlus. The occupancy schedules for the seven zones were generated based on the reference schedule for office buildings available in OpenStudio, with additional time/value offsets, so that there are no two same schedules in the building.

Results and discussion

Compared to the implemented RBC, the implemented MPC used around 24-25% less energy for heating (Table 1). Similar savings were achieved in both scenarios, CTRL-DK1 and CTRL-EE.

Table 1: Total absolute and relative (to RBC) heating energy consumption per scenario.

Scenario	Q_{tot} [kWh]	Q_{rel} [% of RBC]
CTRL-EE	4050.43	75.4
CTRL-DK1	4091.90	76.2
RBC	5368.66	100.0

Based on the indoor temperature profiles (Figures 5-6) it can be concluded that most of the energy savings were due to the demand driven heating, as opposed to the fixed schedule-based behavior in the case of RBC. The indoor temperature profiles in RBC were repetitive and, in many periods, not reflecting the actual occupancy, e.g. see large deviations between RBC and CTRL-EE/CTRL-DK1 on January 4, zones 2-7 in Figure 5. In addition to the demand-driven behavior, in most cases MPC was able to preheat the zones before the actual occupancy occurred, with some exceptions when it did not start the preheating early enough, e.g. on January 4, zone 4 in Figure 5. Finally, the positive spike in indoor temperature in all zones around midnight January 4 (Figure 5) was due to the negative energy price (Figure 8).

The similar monthly profiles of indoor temperature in CTRL-EE and CTRL-DK1 (Figure 6) suggest that the highest priority objective, i.e. the thermal comfort maintenance, dominated the solution.

Controleum MPC reduced the discomfort (measured in Kh) by around 70% (Figure 7). The discomfort metric used in the study was calculated as the product of the temperature difference between the setpoint of 20 °C and the actual temperature and the time in which the difference was observed. Only occupancy periods and only the negative temperature differences were taken into account, i.e. when the indoor temperature was lower than 20 °C. E.g. 1 Kh means that the temperature was below the setpoint by 1 degree during 1 h of occupancy. The obtained discomfort metrics were 287.85, 42.18, 41.58

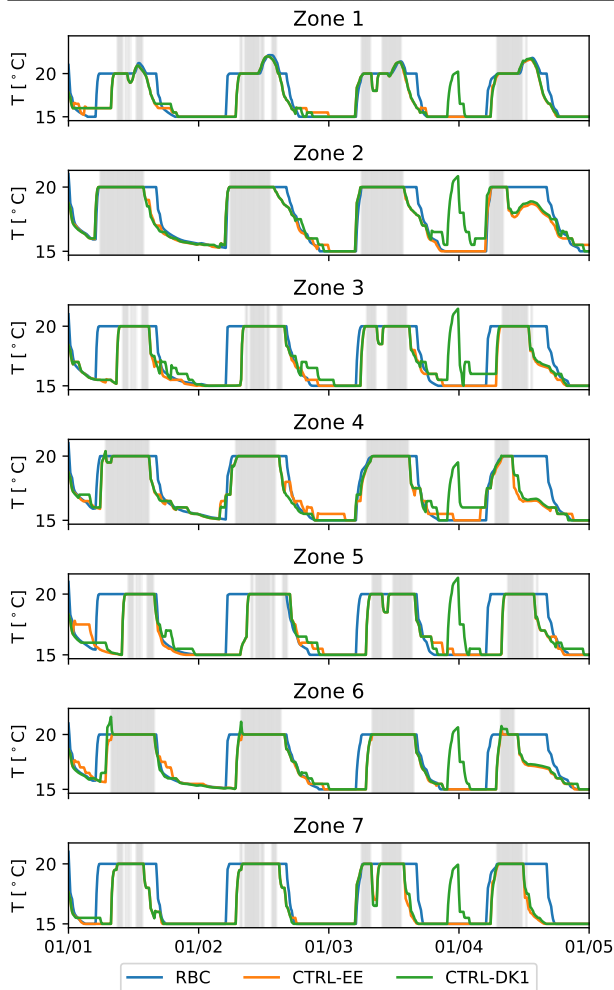


Figure 5: Indoor temperature profiles within the first five days of simulation. The areas shaded in gray mark occupancy periods.

for RBC, CTRL-EE, CTRL-DK1, respectively. The result depends on the chosen reference RBC schedules, however the discomfort in RBC could be only decreased at the cost of increased energy consumption. We believe that the slightly lower discomfort in CTRL-DK1 as compared to CTRL-EE is likely due to the stochastic nature of the optimization algorithm. Since Controlem had access to 100% accurate occupancy “predictions”, it could theoretically minimize the discomfort to 0 Kh. The fact that the discomfort metric was non-zero indicate that either the solution was suboptimal (e.g. due to the maximum CPU time reached) or the final solution is a result of the interplay between different objectives. The suboptimality of the solution is at least partially true as can be seen in the case of zone 5, CTRL-EE, January 3 in Figure 5, where the indoor temperature setpoint is slightly above actually needed. The influence of the objectives hierarchy and optimization settings should be investigated further in the future.

The indoor heating profiles and the energy price during a subperiod of the analyzed month can be com-




Figure 6: One-month indoor temperature profiles for the three considered scenarios (RBC, CTRL-EE, CTRL-DK1). Each column represents a single zone (scenario - zone number).

pared in Figure 8. As in the case of temperature, CTRL-EE and CTRL-DK1 followed a similar trend with one major exception on January 4 when the energy price was negative for a short period of time. The controller in CTRL-DK1 decided to consume as much energy as possible in that time, meaning that the second priority objective came into play. However, for most of the time the price signal had no influence on the solution. Possibly higher price variations or different objective hierarchy would be needed to effectively optimize for the total energy cost in a real application.

Contrary to the expectations the CTRL-DK1 yielded




Figure 7: Indoor comfort violations for the three considered scenarios.




Figure 8: (a) Total heating profiles for the three scenarios vs. (b) energy price.




Figure 9: Total energy cost per scenario.

slightly higher energy cost than CTRL-EE (Figure 9). However, as argued before, this is likely due to the stochastic nature of the optimization algorithm and the dominant role of the highest priority objective of the thermal comfort maintenance.

The presented results were computed using code that was only partially parallelized, e.g. Pareto frontier detection was performed on multiple cores. However, the main bottleneck with respect to the CPU time was the simulation of the gray-box zone models, which was performed on a single core. Due to the nature of MOGA, the zone model simulations need to be repeated thousands of times. In this setup a maximum allowed optimization time per each 7h optimization horizon was five minutes. After each optimization, a one-minute time slot was used to synchronize the measurements and control strategy between Controlem and the virtual building. The optimization was repeated every 1h of the virtual building's time. In total, around 3 days of real time were needed to perform a one-month emulation of the MPC-controlled virtual building (for one scenario). Although the computational requirements of Controlem are considerable, they are feasible for implementation in real buildings. However, implementa-

tion in large buildings (with hundreds of zones) may require parallelization of the zone model simulations. In general, more investigations regarding the scalability of the framework are required.

Conclusions

In this paper the results of the MPC framework based on the multi-objective genetic algorithm are compared with the traditional rule-based control implemented in EnergyPlus. The initial results show energy consumption reduction of around 25%. The computational requirements of the framework are feasible for real implementation, at least in medium-sized buildings. However, there is a potential for further parallelization of the code, possibly extending its applicability to larger buildings.

The analysis conducted in this study was limited in a number of ways. Most importantly, only two scenarios with respect to the chosen objectives were considered, both with the thermal comfort maintenance as the highest priority, and with the minimization of the energy consumption and minimization of the energy cost as the second priority objective, respectively. The minimal difference in results between both scenarios suggest that the highest priority objective, i.e. thermal comfort, dominated the solution. More simulations are required in order to study the robustness of the framework under various objective hierarchies.

In addition, although the achieved savings are significant, it seems that they are mostly due to the demand-driven heating and not due to the utilization of the building dynamics, as commonly cited in MPC papers. This may be due to either a low thermal inertia of the studied building (it takes short time to heat up) or due to the deficiency of the dynamic optimization method. Again, more simulations are required in order to investigate this behavior, possibly with different levels of thermal mass in the building.

Acknowledgment

This work was supported by the Innovation Fund Denmark for the project COORDICY (4106-00003B).

References

- Arendt, K., M. Jradi, M. Wetter, and C. Veje (2018, 10). Modesty: An open-source python tool for parameter estimation in functional mock-up units. In M. Tiller, H. Tummescheit, and L. Vanfretili (Eds), *Proceedings of the 1st American Modelica Conference*, pp. 121–130. Modelica Association and Linköping University Electronic Press.
- Ascione, F., N. Bianco, C. D. Stasio, G. M. Mauro, and G. P. Vanoli (2016). Simulation-based model predictive control by the multi-objective optimization of building energy performance and thermal comfort. *Energy and Buildings* 111, 131 – 144.

- Bianchini, G., M. Casini, A. Vicino, and D. Zarrilli (2016). Demand-response in building heating systems: A model predictive control approach. *Applied Energy* 168, 159 – 170.
- Blochitz, T., M. Otter, M. Arnold, C. Bausch, C. Clauß, H. Elmqvist, A. Junghanns, J. Mauss, M. Monteiro, T. Neidhold, D. Neumerkel, H. Olsson, J. v. Peetz, S. Wolf, A. S. GmbH, Q. Berlin, F. Scai, and S. Augustin (2011). The Functional Mockup Interface for Tool independent Exchange of Simulation Models. In *In Proceedings of the 8th International Modelica Conference*.
- Cao, X., X. Dai, and J. Liu (2016). Building energy-consumption status worldwide and the state-of-the-art technologies for zero-energy buildings during the past decade. *Energy and Buildings* 128, 198 – 213.
- Chen, T. (2001). Real-time predictive supervisory operation of building thermal systems with thermal mass. *Energy and Buildings* 33(2), 141 – 150.
- Coffey, B., F. Haghighat, E. Morofsky, and E. Kutrowski (2010). A software framework for model predictive control with genopt. *Energy and Buildings* 42(7), 1084 – 1092.
- Dawson-Haggerty, S., X. Jiang, G. Tolle, J. Ortiz, and D. Culler (2010). sMAP: A Simple Measurement and Actuation Profile for Physical Information. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, New York, NY, USA, pp. 197–210. ACM.
- De Coninck, R. and L. Helsen (2016). Practical implementation and evaluation of model predictive control for an office building in brussels. *Energy and Buildings* 111, 290 – 298.
- Dong, B. and K. P. Lam (2014, Feb). A real-time model predictive control for building heating and cooling systems based on the occupancy behavior pattern detection and local weather forecasting. *Building Simulation* 7(1), 89–106.
- D&R International (2012). *2011 Buildings Energy Data Book*. U.S. Department of Energy.
- Elmqvist, H., D. Ab, S. E. Mattson, and M. Otter (1998). Modelica - the new object-oriented modeling language. In *In Proceedings of The 12th European Simulation Multiconference*, pp. 127–131.
- García, C. E., D. M. Prett, and M. Morari (1989). Model predictive control: Theory and practice—a survey. *Automatica* 25(3), 335 – 348.
- Garnier, A., J. Eynard, M. Caussanel, and S. Grieu (2014). Low computational cost technique for predictive management of thermal comfort in non-residential buildings. *Journal of Process Control* 24(6), 750 – 762. Energy Efficient Buildings Special Issue.
- Garnier, A., J. Eynard, M. Caussanel, and S. Grieu (2015). Predictive control of multizone heating, ventilation and air-conditioning systems in non-residential buildings. *Applied Soft Computing* 37, 847 – 862.
- Hazyuk, I., C. Ghiaus, and D. Penhouet (2012). Optimal temperature control of intermittently heated buildings using model predictive control: Part i – building modeling. *Building and Environment* 51, 379 – 387.
- Hazyuk, I., C. Ghiaus, and D. Penhouet (2014). Model predictive control of thermal comfort as a benchmark for controller performance. *Automation in Construction* 43, 98 – 109.
- Huang, H., L. Chen, and E. Hu (2015a). A neural network-based multi-zone modelling approach for predictive control system design in commercial buildings. *Energy and Buildings* 97, 86 – 97.
- Huang, H., L. Chen, and E. Hu (2015b). A new model predictive control scheme for energy and cost savings in commercial buildings: An airport terminal building case study. *Building and Environment* 89, 203 – 216.
- Karlsson, H. and C.-E. Hagentoft (2011). Application of model based predictive control for water-based floor heating in low energy residential buildings. *Building and Environment* 46(3), 556 – 569.
- Killian, M., B. Mayer, and M. Kozek (2016). Cooperative fuzzy model predictive control for heating and cooling of buildings. *Energy and Buildings* 112, 130 – 140.
- Kwak, Y. and J.-H. Huh (2016). Development of a method of real-time building energy simulation for efficient predictive control. *Energy Conversion and Management* 113, 220 – 229.
- Kwak, Y., J.-H. Huh, and C. Jang (2015). Development of a model predictive control framework through real-time building energy management system data. *Applied Energy* 155, 1 – 13.
- Lazutkin, E., A. Geletu, S. Hopfgarten, and P. Li (2014). Modified Multiple Shooting Combined with Collocation Method in JModelica.org with Symbolic Calculations. In *Proceedings of the 10th International Modelica Conference*. Lund (Sweden), 10-12 March 2014.
- Li, X. and A. Malkawi (2016). Multi-objective optimization for thermal mass model predictive control

in small and medium size commercial buildings under summer weather conditions. *Energy* 112, 1194 – 1206.

Maasoumy, M., M. Razmara, M. Shahbakhti, and A. S. Vincentelli (2014). Handling model uncertainty in model predictive control for energy efficient buildings. *Energy and Buildings* 77, 377 – 392.

Magnusson, F. and J. Åkesson (2015). Dynamic Optimization in JModelica.org. *Processes* 3(2), 471–496.

Morosan, P.-D., R. Bourdais, D. Dumur, and J. Buisson (2010). A distributed mpc applied to multi-source temperature regulation in buildings. *IFAC Proceedings Volumes* 43(19), 91 – 96. 2nd IFAC Workshop on Distributed Estimation and Control in Networked Systems.

Ogunsola, O. T., L. Song, and G. Wang (2014). Development and validation of a time-series model for real-time thermal load estimation. *Energy and Buildings* 76, 440 – 449.

Schirrer, A., M. Brandstetter, I. Leobner, S. Hauer, and M. Kozek (2016). Nonlinear model predictive control for a heating and cooling system of a low-energy office building. *Energy and Buildings* 125, 86 – 98.

Sørensen, J. C. and B. N. Jørgensen (2017). An extensible component-based multi-objective evolutionary algorithm framework. In *Proceedings of the 6th International Conference on Software and Computer Applications*, ICSCA '17, New York, NY, USA, pp. 191–197. ACM.

Zhao, J., K. P. Lam, B. E. Ydstie, and O. T. Karaguzel (2015). Energyplus model-based predictive control within design–build–operate energy information modelling infrastructure. *Journal of Building Performance Simulation* 8(3), 121–134.