

Learnings from experiments with MPC for heating of older school building Building Simulation 2022 Conference, Copenhagen

Christian Ankerstjerne Thilker¹; Peder Bacher¹; Henrik Madsen¹

¹Technical University of Denmark, Department of Applied Mathematics and Computer Science,
DK-2800 Kgs. Lyngby, Denmark

Abstract

The paper presents the learnings from designing and running a model predictive control (MPC) of the heating system in a school building. Several real-life applications of MPC controlled heating have been presented in the literature. Most of them work by controlling the room temperature using a heating system and thus need a reference measured temperature in the building. Some have a single-zone temperature as the reference, while others use some kind of mean temperature of multiple rooms. In the present experiment, the MPC used the mean temperature of all rooms as the reference and was able to keep it within a lower and upper comfort bound, while minimizing the heat costs by responding to a heat price signal. However, the analyses of the temperature in each room revealed that the temperature bounds were heavily violated: some rooms were too cold and some too warm, while the mean was within the bounds. The main conclusion from the study is that, at least for buildings with different sized rooms and room radiator capacities, it's not reliable to use a mean room temperature – rather, the control must consider individual rooms in order to guarantee comfort.

Introduction

Control of HVAC systems in buildings is important for the green energy transition, both in order to decrease the heat demand and to increase flexibility for integration of variable renewables, like wind and solar. The level of activity in the field has been increasing over the last decades and focus has centered around the use of Model Predictive Control (MPC). The present paper describes an actual real-life experiment where MPC was applied to control the heating system in an older school building in a cold climate. The main focus of the paper is on the learnings achieved from the experiment, especially the way the room temperatures were taken into account in the MPC and the underlying model, and how they actually realized. In the MPC, a simple mean of all room temperatures was used, hence the control was carried out at a building level and this temperature was kept well within comfort temperature bounds. However at individual room level the temperatures violated the comfort bounds, some rooms got too cold and others too hot.

Recent overview papers provides a lot of insights into the techniques and challenges of MPC for HVAC in buildings. Killian and Kozek (2016) pose ten questions that

should be considered for MPC in buildings, however of the questions none related to building vs. room level temperature. They emphasize that a main problem "is the high modeling efforts, where currently no commercial tools exist to derive easily a suitable model for MPC design". Drgoňa et al. (2020) provides a very comprehensive review and overview of MPC of energy systems in building. Not much on multi-zone temperature control is included, though Table 5 lists several studies of modelling for control with multi-zone models, however none of these are implemented in real-life experiments.

The the volume of literature on real-life experiments with MPC in buildings is increasing. Liao and Dexter (2004) identifies a single-zone model for a three storage building where they model the mean room temperature of the entire building. Each floor was similar in terms of the heating equipment, which made it easy to model. The results were good in terms of controlling the mean temperature to different levels, but no analysis on room-level is presented. Široký et al. (2011) present an experiment with MPC of heating in a five-floor building block on a university campus. The room temperature was measured in only two reference rooms and parameters in a linear RC-model was estimated using data from those rooms. An MPC ran in closed loop for two weeks and energy savings were achieved over a rule based control. However, the room temperatures are not evaluated in detail in the paper and there is no information about temperatures in other rooms than the reference rooms. West et al. (2014) present an MPC implemented in a large office building. They consider varying costs in the objective function, however it was implemented as constant in the demonstration period. Indoor comfort were evaluated via feedback from occupants, however no detailed evaluation of the measured temperatures is presented. Huang et al. (2015) developed an MPC for control of the indoor air temperature of an airport terminal in Australia. They carry out a simulation and an experiment to test the model and MPC both showing increased comfort and cost reductions. De Coninck and Helsen (2016) presents the results of implementing an MPC in a two-storey office building. They do room temperature averaging: "To obtain a single-zone model, all room temperature measurements are arithmetically averaged into TZon.". The objective was to minimize heat demand, not with a varying price. A comparison to rule based control is included, which showed 20% to 30% per-

cent reduction in heat demand. No evaluation of temperatures on room level is included. Finck et al. (2019) present and implement EMPC for a Dutch building. The models for the building heat dynamics and the weather forecasts are based on artificial neural networks. They tested the controller for flexibility optimisation and to regulate on-site power generation and grid-consumption and feed-in. The results showed that the EMPCs increased the flexibility of the heat demand while maintaining the same heat costs.

The existing literature presenting multi-zone control systems don't deal with flexible demand – only control of air temperatures. Scattolini (2009) explains and review hierarchical and distributed control. Moroşan et al. (2010) presents a simulation study demonstrating how different MPCs for multi-zone temperature control in a building perform. The focus is on the interaction between the rooms in form of heat exchange due to temperature differences. The results indicate, that either a centralized control, which has a full multi-zone model, or a distributed control, where the room models exchange information, is preferable over a fully decentralized control, which does not take the interactions into account. Elliott and Rasmussen (2013) present temperature control of multiple zones with a multi-evaporator vapor compression system. An architecture that is decentralized and modular, avoiding competing controllers and the practical difficulty of implementing a centralized controller, is presented. Eini and Abdelwahed (2019) presents a distributed control, which in a simulation study has a better performance over a centralized control. The model is a detailed multi-zone model, where the parameters are known in advance. Zong et al. (2019) present a case study of MPC-based BEMS for a multi-family residential building where a hierarchy controller design is applied. The performance of the decentralized controller strongly depends on the level of interactions between the subsystems: The distributed, as each controller knows about control actions of its neighbors, keeps the same performance as the centralized. Results are not presented in detail, only a plot for a single zone is presented.

In the modelling carried out for the present paper, a single-zone grey-box model was identified using the approach presented in Bacher and Madsen (2011) describes. The particular model used is the model presented by Thilker et al. (2021). Some development in terms of automatic model selection has been made, as presented by Andriamamonjy et al. (2019). Interesting multi-zone model identification studies have been carried out by Joe et al. (2020), who present a grey-box model of room temperature fitted for each room individually and compared to a model fitted to all at the same time. The total RMSE is smaller for the decentralized approach. Arroyo et al. (2020) presents a divide-and-conquer approach to grey-box multi-zone parameter estimation, where first the parameters are estimated on single-zone level and then used an initial guess in the multi-zone model parameter estimation.



Figure 1: Photo of the building in question.

From the literature, it's apparent that focus on MPC and the underlying data-driven modelling is increasing, however experimenting with MPC in real-life is in an early stage – especially the implementation and application of a price responsive control in real-life experiments is novel.

Main contributions of the paper

The main contributions of the present paper are, first, the presentation of a successful real-life experiment with a price responsive MPC, and second, to highlight some of the challenges encountered using a single-zone model. In particular, it's emphasized that by using the mean room temperature of all rooms as the reference, which must be kept inside a comfort bound, worked well on the single-zone (or building) level, but caused violations in the individual rooms: some rooms got too hot while others too cold, while the mean was still within the comfort bound.

First the building and data setup is described, and thereafter the two experiments carried out: the system identification and the MPC experiment. Second, the results are presented and discussed, and finally the conclusions are drawn.

Building, systems and data

The building

The building is located in Høje Taastrup, Denmark, and is a school with three floors. The uppermost floor is a part-refurbished roof attic.

Build in 1929, the building is not insulated up to modern standards. Figure 1 shows a photo of the building. It includes 10 classrooms that are ventilated by mechanical ventilation using an air handling unit for air circulation. The ventilation was not active during any of the experiments (due to absence of occupants). The facade and internal walls consist of solid bricks (300 mm and 180 mm thickness, respectively). The windows have wooden frames and double-paned low-E glazings. Floors are made from wood joists and the roof is a partly uninsulated and partly insulated slate roof. The building is connected to the district heating grid. The heating system is used for domestic hot water, air handling unit, and space heating. The space heating is a separate water-based circuit with dedicated pumps. Radiators of different types (cast-iron and plane conductors) with individual smart thermostats distribute the heat in the rooms. The thermostats work as

PI-controllers, which regulate the water flow into the radiator units to maintain a pre-defined *set point* by the user. See Bruun (2019) for further technical details about the building.

Data

All main rooms are each equipped with a temperature sensor (uncertainty is $\pm 0.1^\circ\text{C}$) to measure the indoor air temperature. All radiators in rooms with a temperature sensor are equipped with smart thermostats where a temperature set point can be set remotely. Each thermostat have their own temperature sensor, hence they are not controlled using the measured temperature included in the analysis. The supply and return temperature of the water of the building's heating system is measured together with the actual heat usage of the building.

Weather forecasts for the location are available with 6 hours delay and 48 hours ahead. The sampling time of the forecasts is 1 hour. The forecasts were available in real-time and were used in the filter update in the MPC. To see how the weather was during the period, Figure 2 shows the observed ambient temperature and global solar radiation. It can be seen that the weather conditions include both cold and mild days, as well as both sunny and cloudy days.

System identification

The model used for the heat dynamics of the building is based on stochastic differential equations. The identification method is based on a maximum likelihood method using a variant of the Kalman filter to compute the transition densities. The modelling procedure is thoroughly described in Thilker et al. (2021). The model includes the following states

$$\mathbf{x} = [T_i, T_w, \Phi, T_h, T_{ret}]^T,$$

where T_i is the indoor air temperature, T_w is the temperature of the building's wall, Φ is the water flow of the heat system, T_h is the temperature of the average surface temperature of the radiators in the building, and T_{ret} is the return water temperature.

The model is a significant simplification of the heat dynamics of the building. We use the arithmetic mean of the air temperature of the rooms as a measure of the indoor air temperature in the building as a whole. Since we didn't have knowledge about the heat released in the individual rooms, we cannot easily employ a multi-zone RC-based model. Also, the dimensionality of the model increases drastically if multiple rooms were modelled, which complicates real-time MPC due to a bad scaling in computational requirements to solve the optimal control problem. For these reason, we used a single-zone model, well knowing that it might cause issues in the individual rooms.

System identification experiment

The system identification experiment carried out was planned in advance and designed to generate data suitable for system identification. The main focus was to change the control input, the thermostat set point, such that infor-

mation about the essential dynamics of the system can be estimated. A sequence of the set point was designed with four different parts. First part, contained a few long steps to get information about the slow dynamics governing the system. Second part, was a multilevel signal, where extremes were kept for longer periods than values closer to 20°C . Third part, contained short periods with drops to a minimum from the base temperature. Finally, a step sequence where the set point was stepped down and back up again. The forward temperature of the space heating water was set constant to 55°C at all times. The entire sequence was slightly shorter than 7 days and was executed during the Christmas vacation, where the building was unoccupied. The experiment is described in more details by Thilker et al. (2021).

MPC experiment

The MPC experiment was carried out during January and February 2021. The building was no used in the period, due to the pandemic lockdown. This section introduces a direct multiple-shooting method for solving the particular non-linear MPC problem. It also discusses a method to discretise the optimisation problem to make it numerically tractable. The optimisation problem results in the set points applied for the radiator thermostats. However, solving the optimisation problem requires us to know the entire state of the system. For reconstructing the system states the continuous-discrete extended Kalman filter is used Kristensen et al. (2004).

The implemented optimal control problem has the following (Lagrange) form

$$\min_{\mathbf{x}, u, s} \varphi_k = \int_{t_k}^{t_k+T} \ell(\mathbf{x}(t), u(t), \mathbf{d}(t), s(t)) dt, \quad (1a)$$

$$s.t. \quad \mathbf{x}(t_k) = \hat{\mathbf{x}}_{k|k}, \quad (1b)$$

$$d\mathbf{x}(t) = f(\mathbf{x}(t), u(t), \mathbf{d}(t)) dt, \quad t \in \mathcal{T}_k, \quad (1c)$$

$$u_{\min}(t) \leq u(t) \leq u_{\max}(t), \quad t \in \mathcal{T}_k, \quad (1d)$$

$$\Delta u_{\min}(t) \leq \Delta u(t) \leq \Delta u_{\max}(t), \quad t \in \mathcal{T}_k, \quad (1e)$$

$$T_{\min}(t) \leq T_i(t) + s(t) \leq T_{\max}(t), \quad t \in \mathcal{T}_k, \quad (1f)$$

where $\hat{\mathbf{x}}_{k|k}$ is the reconstructed system state, T is the prediction horizon, ℓ is the cost function, L is the terminal cost, and f is the model dynamics. The time set is $\mathcal{T}_k = [t_k, t_k + T[$. $s(\cdot)$ is a slack variable that softens the temperature constraints. In the next section, we describe how we penalise a non-zero slack.

Optimal control problem

To make the optimal control problem in (1) numerically tractable, we use a multiple shooting method to discretise the problem. The problem is discretised in the sense that the system considers \mathbf{x} at discrete time points $t_k < t_{k+1} < \dots < t_{k+N} = t_k + T$. Now, define a function $\phi(\mathbf{x}, u, \mathbf{d})$ that computes the solution at time t_{k+1} to the following initial value problem

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), u(t), \mathbf{d}(t)), \quad (2a)$$

$$\mathbf{x}(t_k) = \mathbf{x}_k. \quad (2b)$$




Figure 2: Observed weather during the period.

Thus, $\phi(\mathbf{x}, u, \mathbf{d}) = \mathbf{x}(t_{k+1})$ integrates the system forward to the next time instance. Furthermore, we assume the disturbances and input be constant between control points

$$u(t) = u_k, t \in [t_k, t_{k+1}[, \quad (3a)$$

$$\mathbf{d}(t) = \mathbf{d}_k, t \in [t_k, t_{k+1}[. \quad (3b)$$

The discretised optimal control problem can thus be written as

$$\min_{\mathbf{x}_i, u_i, \mathbf{s}_i} \varphi_k = \sum_{i=k}^{k+N-1} L_i(\mathbf{x}_i, u_i, \mathbf{d}_i, \mathbf{s}_i), \quad (4a)$$

$$s.t. \quad \mathbf{x}_k = \hat{\mathbf{x}}_{k|k}, \quad (4b)$$

$$\mathbf{x}_{i+1} = \phi(\mathbf{x}_i, u_i, \mathbf{d}_i), \quad i \in \mathcal{N}, \quad (4c)$$

$$u_{\min,i} \leq u_i \leq u_{\max,i}, \quad i \in \mathcal{N}, \quad (4d)$$

$$\Delta u_{\min,i} \leq \Delta u_i \leq \Delta u_{\max,i}, \quad i \in \mathcal{N}, \quad (4e)$$

$$T_{\min,i} \leq T_{i,i} + s_i \leq T_{\max,i}, \quad i \in \mathcal{N} \quad (4f)$$

where $L_i(\cdot) = \int_{t_{k+i}}^{t_{k+i+1}} \ell(\cdot) dt$ and the index set is $\mathcal{N} = \{0, 1, \dots, N-1\}$. To approximate L_i , we use a fourth-order Runge-Kutta method with fixed step size of 3 minutes. The sampling time between control points is fixed and is $T_s = t_{k+1} - t_k = 0.5$ h.

We implement the optimal control problem using CasADi, which offers easy numerical implementation and automatic differentiation for optimal control problems as the above.

Objective functions

The objective function in an optimal control problem has the purpose of making solutions comparable. The objective function should thus reflect all considerations towards the desired behaviour of the controller. We use the following objective function

$$\ell = c \cdot P_h + \rho \cdot s^2, \quad (5)$$

where c is the heat price, P_h is the heat load, ρ is slack penalty and s is the slack. The objective function is non-linear due to the heat term, $P_h = \Phi c_w (T_{\text{for}} - T_{\text{ret}})$, which

depends on the product of two states, Φ and T_{for} , and the quadratic slack.

In the experiment carried out, we wanted to shift the heat load away from peak hours (also known as peak shaving). Therefore, we constructed a price signal accordingly:

$$c(t) = \begin{cases} 1000 & \text{if } t \in \text{PEAKHOURS} \\ 10 & \text{otherwise} \end{cases} \quad (6)$$

where we define the peak hours to be

$$\text{PEAKHOURS} = [06\text{AM}, 10\text{AM}[\cup [5\text{PM}, 9\text{PM}[\quad . \quad (7)$$

The price for heating is thus expensive during the morning and evening hours where the district heating peak hours usually are.

MPC tuning

During the experiment we tuned several parameters of the objective function in an iterative process as we learned how the MPC behaved. The objective function consists of two terms that need to be weighted such that the controller prioritises in an appropriate manner. We found $\rho = 10$ to be suitable together with the price in Equation (6). The reason for the quadratic slack penalty (instead of linear) is that it ensures a smooth objective function. We found that a linear slack penalty caused a more sensitive solution when it is close to the temperature bounds.

To avoid too large variations in the input signal, dampen oscillations and make the solution more robust, we choose to restrict the allowed variation between control points, furthermore, too large increases in the set point will cause large increases in the return temperature, which in general is not desired. We choose the maximum allowed absolute change to be 3 °C per hour. With $T_s = 0.5$, Equation (4e) becomes $-1.5 \leq \Delta u_i \leq 1.5$.

We chose the maximum and minimum bounds of the input to be the maximum and minimum temperature bounds, $u_{\min,k} = T_{\min,k}$ and $u_{\max,k} = T_{\max,k}$. The rationale behind this, is to avoid overheating in some of the fast reacting




Figure 3: Results from the period running with the tuned MPC. Upper plot is of the price and heat demand. Mid plot is of the temperatures considered by the MPC. Lower plot is of the individual room temperatures and the mean temperature. The latter plots also contain the temperature constraints.

rooms to limit the chance that they violate the temperature bounds.

Results

The MPC was tuned for a period and thereafter it was run for a period of nine days. The results from the nine days period is presented and analysed in the following. Selected variables recorded during the period are plotted in Figure 3. The upper-plot shows the price signal and the realized heat demand. It clearly shows that the MPC was able to lower the heat demand in the high price periods, although it was not decreased all the way to zero. The mid-plot shows the lower and upper temperature bounds, together with the set point and mean room temperature. It's clearly seen that the MPC managed to keep mean room temperature within the bounds, except during the first four days in the morning hours where the lower temperature bound is stepped up. The lower-plot shows the individual room temperatures. It's easy to see that there was a huge spread in the temperatures among the rooms. Some responded very fast and became very warm when the set point was increased, others responded slow and the coldest rooms didn't even to reach the set point – these issues will be discussed in detail later. Finally, it's noticed that there was an increasing trend in the mean tempera-

ture over the period, which was caused by the increase in outdoor air temperature over the period, as pointed out in previously. The pattern in the heat demand and temperature response to the price signal is very regular. In order to get more insight into the details, a zoom on the two first days is shown in Figure 4.

From the two presented plots of the results, it becomes clear that the MPC was capable of controlling the heat demand in response to the price signal. However, there is a potential for improvements. Firstly, the heat demand was not able to decrease fully to zero in the periods of high price, especially in the morning price peak. We identify the two main reasons for that as:

- Some hours before the price peak in the morning the set point was stepped down, which can seem to be too early to be optimal, however it's a compromise between decreasing the demand during the peak and the temperature lower bound violation. This could most likely be tuned to achieve a lower demand during the morning peak.
- Due to technical issues, not all the radiators were controlled, so there was a lower limit to the heat demand – probably around the heat level in the afternoon price peak.




Figure 4: Two days plot of the MPC results. Upper plot is of the price and heat demand. Mid plot is of the temperatures considered by the MPC. Lower plot is of the individual room temperatures.

However, the biggest issue encountered with the implemented MPC is the resulting huge spread in room temperatures. As seen in the lower plots of Figure 3 and 4, the spread of the realized room temperatures was huge. This pose a real problem, since the comfort of occupants would have been compromised – essentially the temperature bounds in the individual rooms cannot be guaranteed when a mean temperature over multiple rooms is used as reference. It is noted here, that the thermostats were not controlling using the measure temperature, they had each a sensor. In order to see more details of the rooms temperature response a two days plot of the temperatures is presented in Figure 5. The rooms are divided between the types of rooms in order to see if there are any similarities because of the type of room.

The main findings from this plot are:

- Clearly, the temperature bounds are violated in nearly all rooms – some gets too hot and some gets too cold.
- Similarities due to the type of room is mainly seen for the hallways.
- The response to the temperature set point is very different among the rooms. Some rooms respond very fast, indicating that the radiators heating power is high relative to the room size and heat losses. Some

rooms respond very slow, indicating that the radiator’s power is not sufficient to heat the rooms under the conditions during the period.

Discussion

From the presented results, it is clear that there are various trade-offs between controllability, stability, and performance. The MPC input was quite regularized, which resulted in higher heat usage during the morning peak hours (since the set point could not be lowered further). It overall affected the ability of the MPC to adjust demand when the price changes – and therefore affected the controllability. However, the regularisation was necessary for the MPC to increase stability of the solution. The slower-changing set point resulted in a lower return temperature, since the temperature in the rooms has more time to adapt to the increased set point.

The performance is of course also affected by the regularisation and the model’s ability to predict the air temperature, which could be improved. This could be done using an adaptive parameter estimate update as observations become available.

From the presented room temperature plots, it is apparent that their individual properties, such as radiator capacities, dimensions, etc., have a huge impact in the rooms’

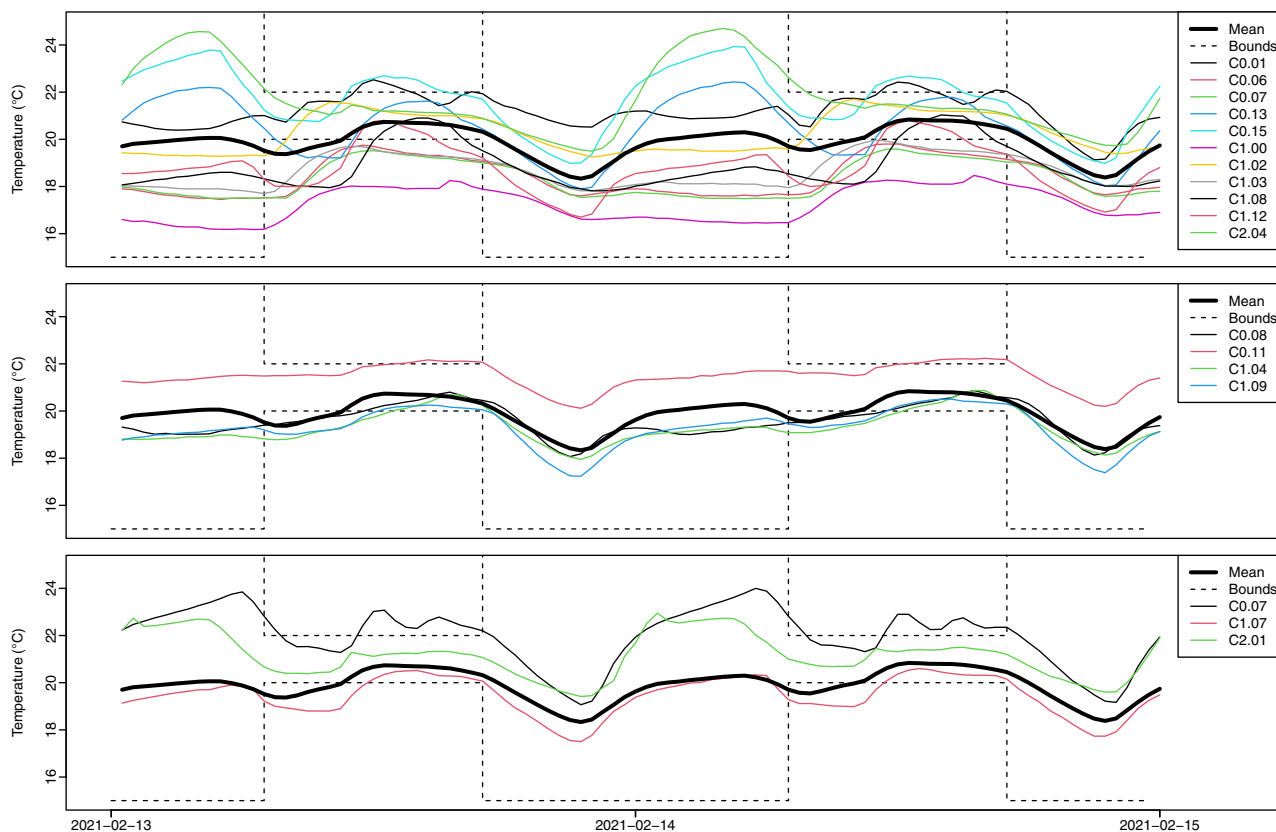


Figure 5: Plot of the temperature set point, the mean room temperature and the individual room temperatures. The upper-plot contains the classrooms, the mid-plot the hallways and the lower-plot the smaller rooms.

individual response to set point changes. Potentially, another interaction may be due to the thermostats opening and closing at about the same time in all rooms, which can create pressure losses, such that the radiators at the far end of the heating circuit cannot supply enough heat. This is a balancing issue, which may be fixed by tuning the max valve openings of each radiator and installing additional radiator capacity.

To which extend these results generalise to other buildings cannot be concluded from the present results. Newer and more well-insulated buildings are probably less likely to suffer from dimensioning issues since less heat is needed Knudsen et al. (2021).

Further work

Regarding the significant room temperature differences, it's a real fundamental issue, which must be solved for MPC in buildings to be usable in practice. It's possible to make a few simple changes to the constraints, for example keeping the upper limit of the set point to e.g. 23 °C could avoid overheating rooms, however it also limits the control capabilities.

Another approach could be to introduce a hierarchical and distributed MPC (Scattolini, 2009; Moroşan et al., 2010):

- Constant room temperature model: Simply a room temperature set point offset can be learned for each

room. This does not take the individual room dynamics into account.

- Dynamic room temperature model: The dynamics of the rooms taken into account by individual room temperature models. One idea is to use an ARX model (which are more black-box models compared to the resistor-capacitor model in this paper) for each room and have individual MPCs run each room.

Conclusion

An MPC for price flexible heat demand was demonstrated in an experiment. The results illustrate the ability of the MPC to respond to a varying price and lower the heat demand of the building in the high price periods, however, there was potential for improvements.

Using the arithmetic mean temperature as a representative for all rooms led to a high spread in room temperatures between rooms, thus violating the temperature comfort bounds. The rooms and radiator power in the building were not uniformly sized, hence this behaviour is not surprising, but pose a real problem for MPCs, which does not take individual room's dynamics into account. To which degree this phenomena can be generalized to other buildings can of course not be concluded with the present study. However, it is clear that when using temperature constraints on a mean temperature over multiple rooms, the

constraints cannot be guaranteed for the individual rooms.

Acknowledgments

The work was carried with funding from the following projects; *Sustainable plus energy neighbourhoods (syn.ikia)* (H2020 No. 869918), FME-ZEN (Research Council of Norway - 257660), Top-up (Innovation Fund Denmark 9045-00017B), SCA+ (Interreg Öresund-Kattegat-Skagerrak) and *Flexible Energy Denmark (FED)* (IFD 8090-00069B).

References

- Andriamamonjy, A., R. Klein, and D. Saelens (2019). Automated grey box model implementation using bim and modelica. *Energy and Buildings* 188, 209–225.
- Arroyo, J., F. Spiessens, and L. Helsen (2020). Identification of multi-zone grey-box building models for use in model predictive control. *Journal of Building Performance Simulation* 13(4), 472–486.
- Bacher, P. and H. Madsen (2011). Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings* 43(7), 1511–1522.
- Bruun, C. G. (2019). Optimization of Building Operation Using High-resolution Sensor Data.
- 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.
- Drgoňa, J., J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, and L. Helsen (2020). All you need to know about model predictive control for buildings. *Annual Reviews in Control* 50, 190–232.
- Eini, R. and S. Abdelwahed (2019). Distributed model predictive control based on goal coordination for multi-zone building temperature control. In *2019 IEEE Green Technologies Conference (GreenTech)*, pp. 1–6. IEEE.
- Elliott, M. S. and B. P. Rasmussen (2013). Decentralized model predictive control of a multi-evaporator air conditioning system. *Control Engineering Practice* 21(12), 1665–1677.
- Finck, C., R. Li, and W. Zeiler (2019). Economic model predictive control for demand flexibility of a residential building. *Energy* 176, 365 – 379.
- Huang, H., L. Chen, and E. Hu (2015). 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.
- Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States) (2020). *DECENTRALIZED APPROACH TO MULTI-ZONE GREY-BOX MODELING FOR MODEL-BASED PREDICTIVE CONTROL*.
- Killian, M. and M. Kozek (2016). Ten questions concerning model predictive control for energy efficient buildings. *Building and Environment* 105, 403–412.
- Knudsen, M. D., L. Georges, K. S. Skeie, and S. Petersen (2021). Experimental test of a black-box economic model predictive control for residential space heating. *Applied Energy* 298, 117227.
- Kristensen, N. R., H. Madsen, and S. B. Jørgensen (2004). Parameter estimation in stochastic grey-box models. *Automatica* 40(2), 225–237.
- Liao, Z. and A. Dexter (2004). A simplified physical model for estimating the average air temperature in multi-zone heating systems. *Building and Environment* 39(9), 1013–1022.
- Moroşan, P.-D., R. Bourdais, D. Dumur, and J. Buisson (2010). Building temperature regulation using a distributed model predictive control. *Energy and Buildings* 42(9), 1445–1452.
- Scattolini, R. (2009). Architectures for distributed and hierarchical model predictive control—a review. *Journal of process control* 19(5), 723–731.
- Thilker, C. A., P. Bacher, H. G. Bergsteinsson, R. G. Junker, D. Cali, and H. Madsen (2021). Non-linear grey-box modelling for heat dynamics of buildings. *Energy and Buildings* 252, 111457.
- West, S. R., J. K. Ward, and J. Wall (2014). Trial results from a model predictive control and optimisation system for commercial building hvac. *Energy and Buildings* 72, 271–279.
- Zong, Y., W. Su, J. Wang, J. K. Rodek, C. Jiang, M. H. Christensen, S. You, Y. Zhou, and S. Mu (2019). Model predictive control for smart buildings to provide the demand side flexibility in the multi-carrier energy context: Current status, pros and cons, feasibility and barriers. *Energy Procedia* 158, 3026–3031.
- Široký, J., F. Oldewurtel, J. Cigler, and S. Prívvara (2011). Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy* 88(9), 3079–3087.