



CHALMERS
UNIVERSITY OF TECHNOLOGY

Effects on district heating networks by introducing demand side economic model predictive control

Downloaded from: <https://research.chalmers.se>, 2026-04-05 16:25 UTC

Citation for the original published paper (version of record):

Håkansson, H., Önnheim, M., Gustavsson, E. et al (2024). Effects on district heating networks by introducing demand side economic model predictive control. *Energy and Buildings*, 309. <http://dx.doi.org/10.1016/j.enbuild.2024.114051>

N.B. When citing this work, cite the original published paper.



Effects on district heating networks by introducing demand side economic model predictive control

Henrik Håkansson^{a,b,*}, Magnus Önnheim^a, Emil Gustavsson^a, Mats Jirstrand^{a,b}

^a Fraunhofer-Chalmers Centre, Gothenburg, SE-412 88, Sweden

^b Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, SE-412 96, Sweden

ARTICLE INFO

Keywords:

District heating
Economic model predictive control
Load shifting
Multi-objective optimization

ABSTRACT

Using economic model predictive control in space heating systems, the heat demand can flexibly be adapted to varying scenarios of conflicting objectives such as thermal comfort, energy consumption, and peak demand. With a controller that minimizes the heating costs subject to thermal comfort constraints within occupancy hours, the resulting heat demand will depend on the cost mechanism. In the context of Swedish district heating networks, optimal demand side control is a multi-objective problem due to variable costs based on both energy consumption and peak demand. By simulating heat demand control using a gray-box model estimated from a Swedish space heating system, we investigate how the established price structures influence the heat load of a population of buildings with economic model predictive control. Our results suggest that by adjusting the incentives from the type of price structures commonly used today, the peak demand can often be reduced by 10-20% with a minor increase in consumption of 1-2%. We also show that by charging the peak demand for multiple buildings collectively, it is financially beneficial to cooperatively control buildings which can reduce the combined consumption and peak demand even further.

1. Introduction

District heating provides efficient heat distribution in urban dense areas and is believed to play an important role in the transition to future energy systems [1,2]. Yet, the International Energy Agency has pointed out that innovations must be exploited cleverly to facilitate the necessary advancements toward fully sustainable solutions [3]. In a present-day district heating plant, boilers can typically be categorized as invariable base capacity or more variable peak capacity, where employment of the peak capacity is more expensive and often less environmentally friendly compared to the base capacity. To minimize the need for peak capacity boilers, or to facilitate other heat sources such as industrial waste heat and intermittent energy sources, different practices of flexibility solutions have been proposed and implemented for optimizing heat production [4]. The different practices of flexibility can be associated with one of the different levels of control within a district heating network, spanning from plant operation and flow control as performed by a supplier, down to the heat demand control as performed by a consumer. Currently, the employed flexibility solutions are mostly based on supplier-controlled thermal energy storage, such as

hot water tanks [5] and the heat storage available in the network pipes [6].

Another area of conceivably cheap, but less exploited, flexibility techniques involves demand side management, or demand response, by utilization of the pre-existing but customer-controlled building thermal inertia. Load shifting, which can be described as sequentially alternating between overheating and underheating compared to the normal levels [7], is one extensively studied approach in this area. In parallel to the research on demand side management, there has been a plethora of works on applying model predictive control, and in particular economic model predictive control (EMPC), for various types of space heating (SH) systems. Conceptually, a heat demand EMPC can be formulated as providing indoor comfort at a minimal heating cost. In many different case studies, substantial savings in energy consumption have been achieved by replacing various types of rule-based controllers such as on-off, weather compensation, or PID [8]. This recent development can largely be attributed to the advancements of the Internet of Things (IoT), e.g., enabling reference signals from indoor temperature sensors and control via building energy management systems. Although the benefits of switching the heat demand control method to EMPC

* Corresponding author at: Fraunhofer-Chalmers Centre, Gothenburg, SE-412 88, Sweden.
E-mail address: henhak@chalmers.se (H. Håkansson).

depend on the efficiency of the preceding method, EMPC can be distinguished from the traditional rule-based methods due to its ability to adapt the heat demand to what is financially preferable.

In a survey performed in 2017, the price structures used among Swedish district heating suppliers were found to always include a tariff for the energy consumption (kWh) that was often accompanied by a tariff for the peak demand (kW) [9]. The relative unit cost of each of the two components varied however largely between different suppliers. Up to this point, the rationale for a supplier to include a peak demand component has been to cover the peak capacity costs. However, a peak demand component may incentivize load shifting to some extent, which may automatically be applied when consumers employ demand-side EMPC. Due to an ongoing trend of more demand-side EMPC employments in Swedish SH systems [10,11], the financial incentives from such a price structure may have a significantly larger impact on the resulting heat load in the network with this shift in heat demand control.

With the established Swedish price structures, the resulting heat demand as given by an EMPC is retrieved via a multi-objective optimization of the energy consumption and the peak demand, meaning that the relative unit cost of the price components will influence the obtained heat demand. The purpose of this paper is to investigate and enhance the understanding of how the relative unit cost of the Swedish district heating price structures affects the heat demand from EMPC-controlled buildings. We perform simulations of one heating season using a gray-box model with parameters estimated from an SH system in Örebro, Sweden, where EMPC is currently employed.

1.1. Outline

In Section 2 we present the background of SH systems and EMPC, and how these have been utilized for demand side management. In Section 3 we give a few numerical examples of how the price mechanism may affect the heat demand that motivates the simulation study. In Section 4 we present the methods for setting up the simulation with the gray-box modeling and the EMPC setup. The results from the simulation are presented in Section 5. In Section 6 we discuss how the results can be interpreted. Finally, in Section 7 we present our conclusions.

2. Background

This work has been preceded by advancements in new equipment and control-oriented studies of SH systems, which enable large-scale demand-side management. In Section 2.1 we present the basic components of typical Swedish SH systems, and how these nowadays are upgraded with computing and online metering tools. In Section 2.2 we briefly review the literature regarding demand-side management utilizing building thermal inertia for district heating. In Section 2.3, we present the current state of the district heating pricing mechanisms. Lastly, in Section 2.4, we connect these aspects to prior works in the literature on heat demand EMPC.

2.1. System and control components

As of 2021, district heating accounted for about 60% of all Swedish heat deliveries used for SH and domestic hot water (DHW), being the dominant heat source for multi-family dwellings and non-residential buildings [12]. The district heating network, i.e., the primary side, transports hot water and is hydraulically separated from the SH systems and DHW systems, i.e., the secondary side [6]. The heat transfer from the primary side to the secondary side is performed via a heat exchanger in the substation where the primary side flow is varied to control the supply temperature on the secondary side. Traditionally, the SH supply temperature has often been set by a pre-defined piece-wise linear function from the outdoor temperature, referred to as weather compensation. The final heat transfer to the indoor air is typically locally regulated in each heating zone by a thermostatic radiator valve

(TRV), which actuates the flow in a radiator to maintain a desired indoor temperature.

While these components and methods have conveniently been used for a long time, there are ongoing shifts in updating and integrating these systems with modern tools. One example is online connected TRVs, sometimes referred to as smart TRVs, which may allow for scheduling setpoints and remote control [13]. Another trend is the replacement of the weather compensation of the supply temperature with more sophisticated control based on feedback from indoor temperature sensors [10]. The main problem that is addressed with this development is that the weather compensation curve can only be shifted manually, which is normally only done after complaints from tenants experiencing a cold climate. Thus, TRVs are given a large control authority, but they do not always operate as desired due to malfunctioning valves and poorly balanced SH systems [14,15]. The consequence of poor TRV control is a fluctuating indoor temperature, and to avoid regularly overheated zones the weather compensation curve is often shifted upwards to give a disproportionately high supply temperature [16]. The average indoor temperature is therefore well above the lower comfort limit and the energy consumption is unnecessarily large. With the feedback signals from the indoor temperature sensors, a better adaptation of the secondary side supply temperature can assist the TRV control. In some cases, TRVs have even been abandoned and feedback control of the supply temperature control has been established as the single heat demand control mechanism [10].

2.2. Load shifting with building thermal inertia

In the context of district heating, most of the studies about load shifting using building thermal inertia have focused on reducing the peak load using various techniques [4], whose applicability has been demonstrated in various settings. In one field study by Kensby et al. [17] the flexibility potential of a typical building in a Swedish district heating network was demonstrated by periodic offsets from the normal level of the secondary side supply temperature. Although this scheme effectively resulted in alternating overheating and underheating of the building, the indoor climate was kept within an acceptable range. Other field studies have focused on constraining the total heat demand by simultaneous control of the secondary side supply temperature of 14 buildings [18] and prioritization of DHW over SH during peak demand hours [19]. Simultaneous control has also successfully been used for reducing morning peak heat loads, by distributing slots for the start-up of different SH systems after being turned off during the night [20]. The value of these activities from the perspective of a district heating network has also been suggested by numerous simulations, see for example Romanchenko et al. [21] or Dominkovic et al. [22].

The aforementioned field studies are examples of direct demand-side management, i.e., the heat demand is directly modified whenever it is useful for production. On the contrary, indirect demand side management refers to adjusting the heat demand by giving the consumers financial incentives, which has been extensively studied for electricity-powered heat sources [23]. Similar procedures have more recently also been studied for demand-side applications in the context of district heating, where the cost is minimized for a single building based on a time-varying heat price, as in Thilker et al. [13]. The large-scale implications of such demand-side optimal control with time-varying costs, using historic real-time marginal costs, have shown great potential in reducing the operational cost [24]. While time-varying prices can be an attractive approach to motivate consumers to install flexible heat demand control and enable demand-side management in the first place, there may be side effects. With the primary intention of peak load reduction, a dynamic price signal may be challenging to implement when many consumers can have a flexible heat demand. A relevant example is Hedegaard et al. [25], where optimal demand-side control of an entire urban area was simulated under a time-varying price signal. The price signal, designed by the authors, introduced occurrences of new

peaks during the hours of lower prices, which could only be mitigated by constraining the degree of load shifting.

2.3. Price mechanisms

Swedish district heating is a deregulated market, but to enhance transparency the heat suppliers must report the costs and profits to the authorities [26]. With the price mechanisms that are commonly used among the Swedish district heating companies, the cost is dominantly variable of the heat demand but sometimes accompanied by a small fixed cost [9]. The structure of the variable components varies between different suppliers, but there are conventionally three different types of components: energy consumption, load demand, and efficiency.

The energy consumption component, which is virtually used by all suppliers, corresponds to the energy consumed in kWh and may be related to the average operating costs. The load demand component, which is typically the other dominating price component besides energy consumption, is often included to cover the costs related to the necessary peak capacity. It has been implemented in many ways by different suppliers, and while some of them are based on averaged consumption, some suppliers use the largest peak demand in kW over a period [27]. The efficiency component, implemented by either flow or return temperature, is included by many suppliers to motivate consumers to have efficient equipment.

2.4. Heat demand EMPC

Model predictive control [28] is a generic framework where a process is controlled by repeated optimization over a receding horizon, where future values are predicted using a dynamic model of the process. When the objective function is based on a financial metric rather than the process itself, one uses to talk about economic model predictive control, EMPC. While its history is more profound in other industries, EMPC for SH demand has been extensively studied mainly from the 2010s and forward. Numerous experimental studies have shown remarkable energy-saving potential, even up to 40%, by employing EMPC [8,29]. Examples of building types where successful real-world implementations have been described in the literature include residential buildings, schools and universities, offices, and commercial buildings. The basic prerequisites for setting up a heat demand EMPC are computational resources for solving the optimization problem, a digital feedback signal from indoor temperature, the ability to control the heat demand, and a dynamic model.

In buildings supplied by district heating, the EMPC implementations have been enabled by the newly installed components as discussed in Section 2.1, including indoor temperature sensors. As for the varying system setups, the heat demand is actuated differently, which gives slightly different EMPC formulations. In the cases with controllable TRVs, the control signal is an indoor temperature setpoint [13,30,31], but with supply temperature control it is the water temperature [32–34]. In neither case, it is the actual heat demand that is controlled, although it can be estimated from the control signals and process values.

The dynamic model has frequently been pointed out as an obstacle in the employment of EMPC. There have been different approaches proposed, in a range spanning from detailed white-box models to neural networks-based black-box models [35]. Gray-box models have been an increasingly popular alternative since they often require less effort than white-box modeling, but still allow for a physical interpretation of parameters [36,13]. With a given dynamic model and a sequence of control inputs, a control plan, the future indoor temperature can be predicted. To integrate the dynamics into the optimization, predicted thermal comfort can be included as a constraint [31,37], as a penalty term in the objective function [38], or by combining both strategies [39].

The price mechanism will have a direct control impact since it is used in the EMPC objective function. In the aforementioned studies of indirect demand-side management in Section 2.2, time-varying prices have been considered. However, as we will see in the next section, even the current district heating price structures may incentivize load shifting via the load demand component.

3. Motivating examples

In this section, we discuss the heat demand EMPC under the price structures as presented in Section 2.3 with respect to load shifting. First, in Section 3.1, we establish the tradeoff between energy consumption and peak demand by an example. This perspective is extended in Section 3.2, where we introduce the occurrences of nighttime setbacks and how this may affect the heat load.

3.1. The tradeoff between consumption and peak demand

With the established district heating price components, consumption and load demand, the load demand component may incentivize load shifting when it is based on the peak demand. To exemplify this, we consider a single building, where the indoor temperature change is mainly driven by the supplied heat and the thermal loss by the heat transfer from the interior to the exterior. To provide the required indoor climate, the consumer sets a control plan, i.e., a sequence of heat demand over time, that compensates for the time-varying thermal loss. Due to the thermal inertia of a building, there is not a unique control plan solving this problem, but typically a large set of plans with different consumption and peak demand.

Narrowing the scope, we cannot typically obtain the least consumption and the least peak demand from the same control plan. Under the assumption of perfect knowledge of the exact thermal loss at any time generated by any control plan, we can retrieve the least consumption control plan by feeding an amount of heating power equaling the thermal loss. By feeding more heating power than the corresponding thermal loss, the indoor temperature will increase. To reduce the peak demand, a load shift procedure can be applied, meaning that the least consumption plan is modified by employing such an overheating procedure before a thermal loss peak. The heat demand at the time of the peak can then be reduced by utilizing the heat buffer that has been accumulated by the overheating. Since the thermal loss is driven by the difference in indoor and outdoor temperature, buffering heat will also increase the thermal loss per consumed heating power. Hence, there is a tradeoff where reduced peak demand also increases consumption, which is illustrated in Fig. 1.

To determine what is an optimal plan, the importance of reducing peak demand and consumption, respectively, must be defined. Reconnecting to Section 2.3, the relative unit costs between the energy consumption component and the load demand component, when implemented as the peak demand, will therefore dictate the optimal level of load shifting for buildings in the network. Thus, if a substantial number of customers employ optimal control, the price structure will have a significant impact on the total load.

3.2. Multi-building control with various scheduling

In the previous example, we discussed the control problem for a single building with time-constant requirements. Such requirements may be suitable in a dwelling, where the indoor climate should be comfortable at any time. For buildings such as workplaces or schools, when occupancy is limited to work hours, time-dependent requirements are conceivable. To exemplify such a case, we consider a workplace building where the indoor climate must be comfortable only on workdays from 8 a.m. to 5 p.m., i.e., we can have a nighttime setback to reduce energy consumption. When comparing optimal plans for this case in




Fig. 1. Illustration of the tradeoff between consumption and peak demand with three different control plans for a dwelling. From the top, the first panel shows the heating power, the second panel shows the simulated indoor temperatures given the control plans, and the third panel shows the outdoor temperature. The rightmost panel shows the peak demand, the maximum heating power, against the consumption averaged over the period. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

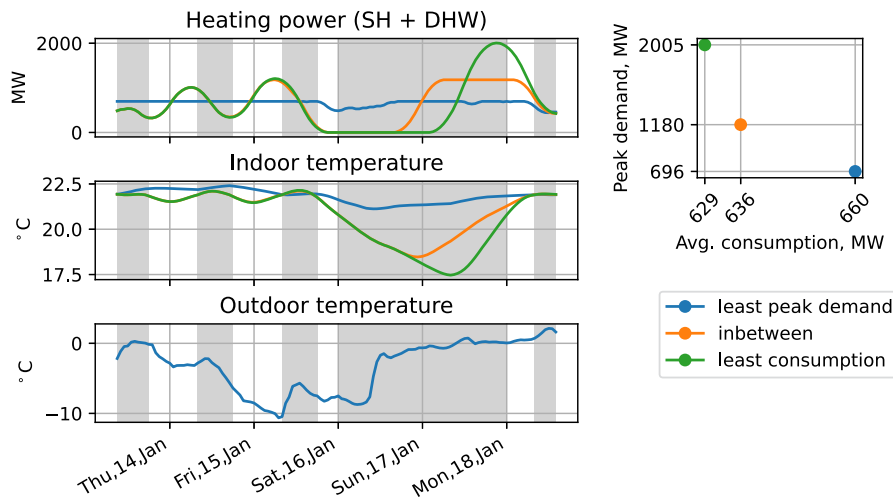


Fig. 2. Comparison of three simulated control plans for a workplace with varying consumption and peak demand. Time intervals marked with gray denote the inactive hours when the comfort requirement is not applicable. For more details, see the caption of Fig. 1.

Fig. 2, concerning the same tradeoff as in Fig. 1, we see that the workplace’s plans are considerably different from the dwelling’s.

One general remark from the two examples visualized in Figs. 1 and 2 is that the varying requirements cause the peak demand events of each building separately to not occur at the same time. Hence, the overall peak demand of both buildings does not equal the sum of the buildings’ peak demand separately. With this in mind, we consider a slightly different example, visualized in Fig. 3, where a dwelling and a workplace are connected to the same district heating network. We consider two different control plans where the consumption summed over the two buildings is the same for both control plans. The peak demand for control plan 2 is higher than for control plan 1 for each building separately, but considering the total load of both buildings, control plan 2 is lower than control plan 1.

As illustrated in this example, the control plan with the least peak demand of each building separately does not necessarily correspond to the one with the least total peak demand. Hence, control of all buildings must be included in one common computation to optimize the total

peak load for the entire system. With the present-day price models, a housing company owning both the dwelling and the workplace would still benefit from selecting control plan 1 over 2 in Fig. 3 where peak demand is charged by each district heating connection individually. In other words, optimal scheduling will not be applied even when it is possible, because there is no financial upside for the consumer. By instead charging all connections that the consumer controls together, we see from the example that optimal scheduling may be financially advantageous, which could contribute to reduced peak loads.

4. Methods

To investigate the influence of different price structures concerning the total heat load, we simulate EMPC heat demand control by using gray-box modeling of a template SH system located in Örebro, Sweden. The template SH system consists of 20 multi-family dwellings, with a total of almost 300 apartments, in Örebro, Sweden, which are all connected to a substation where an EMPC algorithm controls the SH heat




Fig. 3. Example of two sets of control plans for a dwelling and a workplace with equal consumption but different peak demand. The two upper plots show the heating power for each building of both control plans. The third plot shows the total load from both buildings, which is equal to the sum of the two upper plots at each time step. Dashed horizontal lines correspond to the maximum power, the peak demand, in each plot.

demand. More details about the template SH system are given in Appendix A.

The simulations are based on a population of two buildings, a dwelling and a workplace, that have identical dynamics but different occupancy scheduling. The dwelling must maintain a comfortable indoor climate all of the time, while the workplace indoor climate needs only to be comfortable during work hours. With variations of the price compositions and the corresponding sizes of the two types of buildings, which give different financial incentives, the heat demand control is simulated using the outdoor temperature for 14 weeks during winter 2020/2021. We do this in two different settings: *separable EMPC*, in which the cost corresponds to the existing price mechanisms where the peak demand is calculated building-wise, and *cooperative EMPC* in which the peak demand is instead calculated group-wise.

In Section 4.1, we present our gray-box approach for modeling the heat demand of the template system. In Section 4.2, we describe the implementation of the heat demand EMPC. In Section 4.3, we give more details on how the simulation was implemented.

4.1. Simulated indoor temperature dynamics

For simulating the heat dynamics of buildings, we use a linear gray-box model that outputs the average indoor temperature given inputs of outdoor temperature and heating power. The model can be described as a linear autoregressive model with exogenous inputs (ARX), whose parameters are recursively estimated. This model structure and identification approach corresponds to the predictive model in the EMPC that is employed in the template SH system. The same estimated dynamics are used for the simulation of both the dwelling and the workplace, although some modifications, which are presented in Appendix B, are implemented in the latter.

The basics of the gray-box model are given by a discrete-time linear dynamical system, which can be interpreted as the discrete-time counterpart of an energy balance equation with linear loss and production terms (i.e., the rate of change in temperature is proportional to the net heat flux into the system). Hence, in discrete time the temperature difference between two consecutive points in time can be modeled as a linear combination of the difference between indoor and outdoor temperature and the heating power. To allow for a more flexible model, we also let the linear combination include time-lagged heating power and a discrete-time signal with a period of 24 hours. This linear combina-

tion of terms is denoted f below and constitutes the right-hand side of the following temperature difference equation

$$\frac{x_t - x_{t-1}}{\Delta t} = f(y_{t-1}, x_{t-1}, u_{t-1:t-1-H}; \theta, c_{0:23}, v_{0:H}) \quad (1a)$$

$$= \left(\theta (y_{t-1} - x_{t-1}) + c_{\text{mod}(t,24)} + \sum_{n=1}^H v_n u_{t-n} \right), \quad (1b)$$

where θ is a coefficient for the heat transfer between the interior and the exterior of the building per hour, $y_t \in \mathbf{R}$ is the outdoor temperature in $^{\circ}\text{C}$, $x_t \in \mathbf{R}$ is the indoor temperature in $^{\circ}\text{C}$, $c_{1:24} \in \mathbf{R}^{24}$ is a vector of hour-wise constants in $^{\circ}\text{C}/\text{h}$ and mod denotes the modulo operator (i.e., each hour of the day has a different constant), $v_{1:H} \in \mathbf{R}^H$ are coefficients describing the heat transfer from the heating system into the building including inertia of the past H hours in $^{\circ}\text{C}/\text{Wh}$, and $u \in \mathbf{R}^H$ is the heat demand of the previous H hours in W . In the simulations, we used $H = 6$.

Although linear models with a similar structure as Eq. (1b) have been used in other works, see for example [40], the real dynamics may be non-linear [33,41]. Nevertheless, the temperature dynamics may be accurately described by a linear model in the short term for the template SH system. As an identification procedure we use a Kalman filter approach to estimate model parameters recursively [42]. All the time steps from the first hour of the first simulation week up to the last hour of the simulation week are used when computing the estimates for that particular week. More details about the identification procedure and the obtained coefficient estimates are given in Appendix B. The error of an open-loop simulation of the estimated dynamics for a week-long horizon, which is visualized in Fig. 4, is typically limited to $\pm 0.4^{\circ}\text{C}$. In Fig. 4, we also see that the median simulation error over different weeks stays close to zero over the full week-long horizon.

4.2. Heat demand EMPC

From the numerical examples in Section 3, we have seen that control plans can be characterized by a tradeoff between minimizing the energy consumption and minimizing the peak demand and that the peak load can be lowered by aggregating multiple buildings into one single optimization problem. In this section, we present two different EMPC formulations, called *separable* and *cooperative*, intended for modeling heat demand control of building populations. These EMPC formulations can be applied to different populations and different price tariffs to form




Fig. 4. Distribution of error from multiple week-long open-loop simulations, using the estimated dynamic models. The error distribution, over the simulation weeks, at a given time since the start of each week is visualized by the range between minimum and maximum, the range between the first quartile (Q_1) and the third quartile (Q_3), and the median.

the basis of the simulated heat demand control presented in the next section.

Both the separable EMPC and the cooperative EMPC are based on multi-objective optimization of energy consumption and peak demand, but they differ in the way the peak demand is calculated. In the separable EMPC, the peak demand is calculated as the maximum heat demand of the optimization horizon for a single building, while in the cooperative EMPC, we assume multiple buildings are controlled together and the peak demand is calculated as the maximum heat demand summed over all buildings. The separable EMPC corresponds to the setup of the EMPC in the template heating system, where the cost is optimized separately per district heating connection, i.e., per substation.

In a population \mathcal{P} of buildings, each building $b \in \mathcal{P}$ may have different requirements and dynamics. The heat demand is set in W as the control signal, denoted as u_t^b for time step t and building b . This control signal is updated every Δt hour. In the separable setting, a sequence of control signals (a control plan) $u^b \in \mathbf{R}^T$, for time horizon T , and building $b \in \mathcal{P}$ is retrieved by solving

$$z^b = \underset{u^b}{\text{minimize}} \lambda C^b + (1 - \lambda) D^b \quad (2a)$$

subject to

$$u_t^b \geq 0 \quad t = 1, \dots, T \quad (2b)$$

$$|u_{t-2}^b - 2u_{t-1}^b + u_t^b| \leq K \Delta t^2 \quad t = 3, \dots, T \quad (2c)$$

$$L_t^b \leq x_t^b \leq U_t^b \quad t = 1, \dots, T \quad (2d)$$

$$x_t^b - x_{t-1}^b = \Delta t f^b(y_{t-1}, x_{t-1}^b, u_{t-1:t-1-H}^b) \quad t = 2, \dots, T \quad (2e)$$

$$\frac{1}{T-H} \sum_{t=1}^{T-H} (u_t^b + w_t^b) \Delta t = C^b \quad (2f)$$

$$\max_{t=1, \dots, T} u_t^b + w_t^b = D^b \quad (2g)$$

where C^b is the averaged consumption in W , D^b is the peak demand in W , $0 \leq \lambda \leq 1$ is a weighting of the cost per W between consumption and peak demand, K is a smoothness constant in W/h^2 , L_t^b is the lower indoor temperature limit in $^\circ\text{C}$, x_t^b is the indoor temperature in $^\circ\text{C}$, U_t^b is the upper indoor temperature limit in $^\circ\text{C}$, f^b is the dynamic model of the building for simulating the indoor temperature change rate in $^\circ\text{C}/h$, H is the number of time lags of the heating power used in Eq. (1b), and w_t^b is the DHW usage in W . The cost function in Eq. (2a) involves a weighted summation of the averaged consumption and peak demand, as defined in Eqs. (2f) and (2g), respectively. Note that the objective value z^b needs to be scaled to obtain a cost corresponding to the prices

discussed in Section 2.3. Eqs. (2b) and (2c) ensure that the heating power is non-negative and smooth over time. The comfort constraint Eq. (2d) bounds the indoor temperature x_t^b within the lower and upper limits, as specified for the particular building at each time step. Eq. (2e) incorporates the simulation of the indoor temperature from the control plan u^b via f^b .

To obtain a variance for different outcomes of control plans we solve Eq. (2) for multiple consecutive week-long plans, rather than bunching the whole heating season into one single control plan. However, the terminal indoor temperature and the last H input signals may be unfavorable for the control plan for the consecutive week. To mitigate this side effect, we exclude the last H time steps in the consumption calculation in Eq. (2f). Thus, given a preceding control plan ending at $t = T - H - 1$, the first H hours of the consecutive control plan starting at $t = T - H$ can provide a comfortable climate as in Eq. (2d) without violating the smoothness assumption in Eq. (2c) and to a peak demand D^b that is at most the peak demand of the preceding control plan.

In the separable setting, we use the same kind of program Eq. (2) for each building, but L_t^b , U_t^b , and f^b vary for different buildings. The control of all buildings in the population is retrieved by

$$\begin{aligned} z^{\text{separable}} &= \underset{u^b, \forall b \in \mathcal{P}}{\text{minimize}} \sum_{b \in \mathcal{P}} z^b \\ &= \sum_{b \in \mathcal{P}} \underset{u^b}{\text{minimize}} z^b. \end{aligned} \quad (3)$$

In the cooperative setting, where we assume that the price is calculated for all buildings of the population jointly, the population peak demand D is used instead of the building-wise computed peak demand D^b . The cooperative EMPC problem is formulated to

$$z^{\text{cooperative}} = \underset{u^b, \forall b \in \mathcal{P}}{\text{minimize}} \lambda \sum_{b \in \mathcal{P}} C^b + (1 - \lambda) D \quad (4a)$$

subject to

$$\text{Eqs. (2b) to (2f)} \quad \forall b \in \mathcal{P} \quad (4b)$$

$$\max_{t=1, \dots, T} \left(\sum_{b \in \mathcal{P}} u_t^b + w_t^b \right) = D. \quad (4c)$$

4.3. Heat demand simulation

To simulate heat demand control, we use the gray-box model presented in Section 4.1, and the two EMPC formulations introduced in Section 4.2 with $\Delta t = 1$ hour. In the simulations, there are two types of buildings: a dwelling and a workplace. Both share most of the parameters of their corresponding gray-box models, but they differ mainly concerning occupancy scheduling. In the dwelling, the comfort requirements in Eq. (2d) are set to give a comfortable climate all the time, i.e., $L_t^b = 21.9^\circ\text{C}$ and $U_t^b = 23^\circ\text{C}$ for all $t = 1, \dots, T$. On the contrary, for the workplace, these limits are only applied during normal work hours Monday to Friday from 8 a.m. to 5 p.m., and outside these hours the indoor climate is unconstrained.

We simulate heat demand control for different populations of two buildings, one of each type, with varying relative sizes by introducing a population parameter $\gamma \in [0, 1]$. This parameter is used in Eqs. (2f), (2g) and (4c) for scaling the heating power at each time step according to the building size when calculating the consumption and peak demand. For example, at time t the heating power of the workplace is calculated by $\gamma (u_t^{\text{workplace}} + w_t^{\text{workplace}})$ and the heating power of the dwelling is calculated by $(1 - \gamma) (u_t^{\text{dwelling}} + w_t^{\text{dwelling}})$. We simulate 20 different populations, represented by 20 evenly spaced values, where the workplace heating power is multiplied by $\gamma = 0, 0.05, \dots, 0.95, 1$, and the dwelling heating power is multiplied with $1 - \gamma = 1, 0.95, \dots, 0.05, 0$.

For modeling the DHW usage of the dwelling, we estimate a daily usage profile based on data from the template SH system. This estima-




Fig. 5. Consumption and peak demand over varying populations of separable EMPC with either least peak demand ($\lambda \approx 0$ in Eq. (3)) in the left column plots, or the least consumption ($\lambda = 1$ in Eq. (3)) in right column plots. Each plot shows, for a population with a workplace size given by the horizontal axis, how the separable EMPC's consumption (upper row) or peak demand (lower row), relative to the baseline performance for the corresponding simulation week, is distributed over the 14 different simulation weeks. This distribution is visualized by the range between minimum and maximum, the range between the first quartile (Q_1) and the third quartile (Q_3), and the median.

tion procedure is further explained in Appendix A, and the DHW usage estimates correspond to the blue line in Fig. A.12. Thus, the simulated dwelling DHW varies throughout the day but is the same for a given time each day. For the workplace, we assume very low usage all the time, but there may still be excess heat due to DHW circulation. We therefore use the minimum average DHW usage of the dwelling, occurring at 2 a.m. of about 80 kW, for the workplace all of the time.

The peak demand component in the price model is based on a week-long period. To simulate the effect of varying price models, control plans for 100 values of $\lambda \in [0.001, 1]$ are computed. With $\lambda = 0$, meaning only peak demand is of interest in the optimization, the heat demand can be set as large as the peak demand all the time, without affecting the objective. However, by setting λ to a value only slightly larger than 0, the peak demand hardly changes but the consumption is lowered substantially. Hence, $\lambda = 0$ should not be of any practical interest, and therefore it is not included in the simulations.

To simulate control plans with EMPC we use Eqs. (3) and (4), with the heating power scaled according to the sizes as given by γ . Both formulations are simulated for 14 different weeks, spanning from December 2020 to March 2021. In both settings, the optimization problems are linear, meaning that a linear programming solver can be used. Each simulated control plan spans a single week-long simulation period, with the different gray-box parameter estimates, as explained in Section 4.1, for that particular week. In contrast to the EMPC running in the template SH system, where the employed control plan is retrieved by an hourly re-computation, we only use the control plan computed at the start of the simulation period. The smoothness constant in Eq. (2c) is set to $K = 0.4$. For the dwelling, the indoor temperature limits in Eq. (2d) are set to $L = 21.9^\circ\text{C}$ and $U = 23^\circ\text{C}$ all the time. For the workplace, the limits on 8 a.m. to 5 p.m. Monday to Friday are $L = 21.9^\circ\text{C}$ and $U = 23^\circ\text{C}$, and outside those hours they are set to $L = -\infty$ and $U = \infty$.

To summarize, we have 2 types of EMPCs, 20 populations, 14 simulation weeks, and 100 different price models. Therefore, we have in total $2 \cdot 20 \cdot 14 \cdot 100 = 56000$ linear programs, for which we use GLPK to solve for obtaining all control plans.

5. Results

In this section, we present the simulation results, focusing on the total heat demand of the simulated population. As a baseline, we consider a traditional control with weather compensation and TRVs without any nighttime setbacks in the workplace. Assuming all components are functioning properly, the baseline should still be energy-efficient and maintain a comfortable climate all of the time. In terms of our simulation, this translates to the population consisting of only a dwelling ($\gamma = 1$) and a price composition only charging consumption ($\lambda = 1$). However, since the relative sizes affect the accumulated DHW usage profile, and thereby the total heat demand, the baseline heat demand is different for the different populations.

In Fig. 5, we see that a separable EMPC can achieve a substantial reduction in peak demand up to 10% - 20% under the price composition with $\lambda \approx 0$ or up to a 4% reduction in consumption under $\lambda = 1$. There may however be downsides with both of these price compositions, as most visible with $\lambda = 1$ where the peak demand is about 100% of the baseline with the dwelling-only population and 230% of the baseline with the workplace-only population. With $\lambda \approx 0$, the consumption is comparable and often larger than the baseline, despite having a large workplace that allows for night-time setbacks and consumption savings. This can also be viewed the other way around: with approximately the same consumption as the baseline, a separable EMPC can reduce the peak demand substantially under $\lambda \approx 0$ for all populations.

Considering the consumption and the peak demand for the separable EMPC, all price compositions for a single simulation week, i.e., the Pareto front as visualized in Fig. 6, there is a larger sensitivity with respect to consumption and peak demand as the workplace size increases. Comparing the cooperative EMPC to the separable EMPC, we see in Fig. 6 that the cooperative EMPC can produce lower peak demand given a consumption and vice versa. This however only applies to the populations not fully dominated by one of the building types, i.e., $\gamma \neq 0$ and $\gamma \neq 1$. For the cases with $\gamma = 0$ and $\gamma = 1$, the cooperative EMPC and separable EMPC boil down to the same problem since the peak demand component is calculated for the only building of the population.

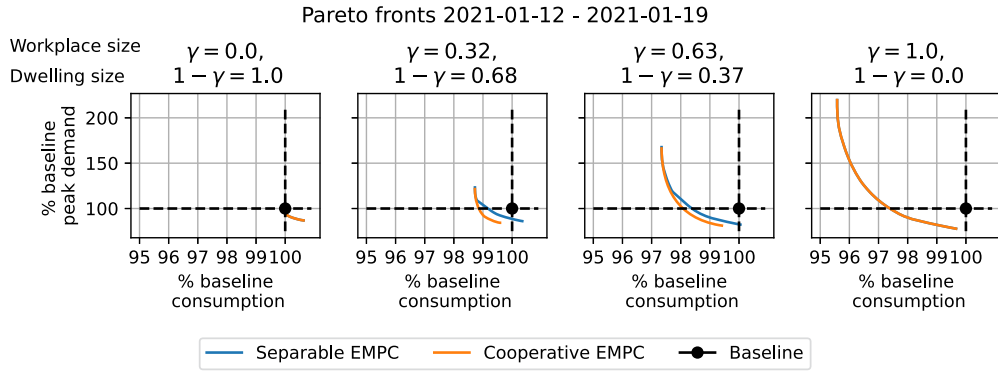


Fig. 6. Consumption and peak demand for all price compositions of four different populations with separable EMPC and cooperative EMPC for one week. The scales of the axes are relative to the baseline consumption on the horizontal and the peak demand on the vertical axis, respectively. The baseline solutions, i.e., 100%, are marked by the dot and dashed lines.

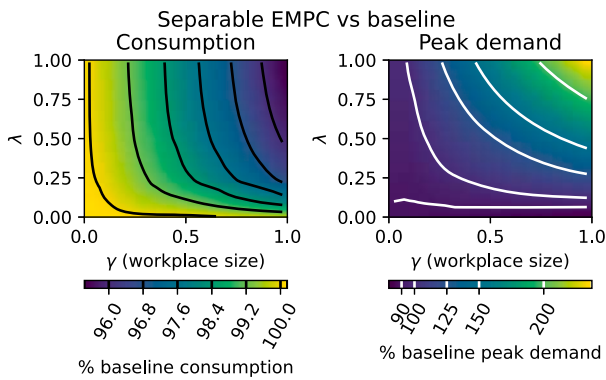


Fig. 7. Consumption and peak demand for separable EMPC compared with baseline for all populations and price compositions. In both plots, the population is given as the workplace size γ on the horizontal axis, and the price composition of the separable EMPC, indicated by λ in Eq. (3), is given by the vertical axis. The color is given by the median consumption or the peak demand, relative to the baseline, over the 14 simulation weeks. The black and white lines show the contour levels of consumption and peak demand, respectively, as indicated by the color bars.




Fig. 8. Reduction in consumption and peak demand with cooperative EMPC compared to separable EMPC for the different populations and price compositions. The black and white lines correspond to the contour levels of Fig. 7.

The overall trend of the median simulation week for all populations and price composition, which is depicted in Fig. 7, is consistently an interpolation between the least consumption EMPC and the least peak demand EMPC, as visualized in Fig. 5. Similarly to Fig. 6, the reduction with cooperative EMPC visualized in Fig. 8 is most apparent among the inhomogeneous populations where the workplace size is 0.2-0.5. The reduction in peak demand is larger than for consumption, up to 7% of the baseline peak demand and up to 0.5% of the baseline consumption. This gain does however also depend on the price composition of the separable EMPC. With a consumption component that is large but not exclusively dominating the price composition, e.g., $\lambda \approx 0.8$, the peak demand is reduced more compared to price compositions with both a smaller and a larger λ . For the consumption reduction, we have the inverse trend that reductions are larger when the peak demand component is larger.

To illustrate the difference between a separable and cooperative EMPC, we consider the control plans produced in one week in January 2021 and a population with a workplace size of 0.32, as shown in Fig. 9. In both of these plans, the peak demand is minimized while constraining the consumption to 99.5% of the corresponding baseline consumption of this week, meaning they correspond to solutions intersecting a vertical line at 99.5% of the upper right plot in Fig. 6. Overall, the complexity of optimal planning makes it hard to explain the rea-

son behind every single action of the controllers by visual inspection. However, one noticeable difference is that the cooperative EMPC utilizes more overheating in the dwelling and utilizes a faster re-heating of the workplace after the weekend. On the other hand, the peak demand hours for the separable EMPC that are avoided with the cooperative EMPC are more likely to be connected to the cold outdoor temperature during the night between Thursday and Friday.

6. Discussion

With these simulations, we have studied the impact of the financial incentives given by the price composition on EMPC-controlled building populations. While the studied populations only consist of two buildings, the results can also be interpreted through populations with multiple dwellings and workplaces where γ denotes the relative size of all workplaces with respect to the total size of the building population. Due to the linearity of the problem, and by assuming the same dynamic model for all buildings of the same type, Eqs. (3) and (4) still gives the optimal heat demand control for each type of building. To account for the different sizes of buildings of the same type, the control plan may then be scaled by the corresponding size.

6.1. Modeling limitations

In many aspects, the modeling choices of this work are characterized by simplifications which may reduce the accuracy, but also enable more in-depth analysis and eliminate potential side-effects. The most obvious




Fig. 9. Comparison of plans from separable EMPC and cooperative EMPC for a week in January 2021 for a population with $\gamma = 0.32$. The upper plot shows the outdoor temperature, the middle plot shows the simulated indoor temperature and the lower plot shows the heating power fed to each group of buildings at a given time with the different controllers. Gray-colored time slots mark the inactive hours, i.e., when the comfort requirement is not applied.

example of such a simplification is to not use a sophisticated model of the template SH system, but only a gray-box model. Realistically, forecasting errors may affect the control and re-plannings of the EMPC, but these are hence omitted in our approach. Still, the gray-box model predicts the dynamics of the average indoor temperature accurately most of the time, with an error of ± 0.2 °C in 50% of the simulation weeks as seen in Fig. 4. It should be reemphasized that our analysis focuses on thermal properties on an aggregated level of multiple buildings, and more comprehensive characteristics such as room-level climate are left out in the modeling. While these modeling simplifications facilitate our exhaustive investigation of different populations and price compositions, a natural extension from this work would be to consider more fine-grained models. That could be detailed and accurate white-box models, but also more complex, potentially non-linear, gray-box models that address the varying behavior of different thermal zones.

Another simplification is the DHW usage profiles, which are assumed to be the same every week, although there is a certain variance. Also, the DHW usage of the workplace is based on assumptions, rather than modeling of a real counterpart. Still, the relative power between SH and DHW should be representative of a real scenario, and despite the presence of the non-avoidable DHW heat demand, the SH load shifting is applicable to affect the total heat demand.

In the EMPC formulations in Eq. (2), there is no upper limit of heating power, only a lower bound. In practice, such an upper bound may be given by the installed capacity and the primary supply temperature. Still, many SH systems are oversized since they typically operate well below their maximum installed capacity. This affects the results through control plans with an enormously large peak demand, for example in Fig. 5 where the peak demand is 230% of the baseline for the workplace-only population. Considering an SH system in a workplace where the baseline control is replaced by EMPC, the installed capacity will in most cases constrain the realizable peak demand. Nevertheless, our simulation still shows what is optimal for the consumer given these incentives and an infinite capacity, and the high peak demand can be avoided by another price composition.

6.2. Implications for district heating networks

To this date, load shifting has rarely been employed in the template SH system since the financial incentives given by the heat supplier are not sufficient. Our simulations show that substantial load shifting

may be applicable under a sufficiently large peak demand component, even with the dwelling-only population. While this work has focused on such an incentive, how the heat supplier may design the actual price mechanism is of course a broader question. In the first place, it may be problematic for the supplier to use a peak demand component, that is implicitly temperature-dependent, for covering the peak capacity costs since much of those costs are present every year, but the income will be less in a warm winter than in a cold winter. Also, during a mild week, a load-shifting strategy may just give unnecessary consumption if the peak capacity is not needed even without load shifting. One solution could be a slightly different price mechanism used by some Swedish district heating companies called “power subscription”, where the consumer selects a maximum capacity in W that the heating power cannot exceed at all during the heating season.

Our simulation results may suggest that both separable and cooperative EMPC are useful for reducing the peak demand by utilizing load shifting. We should however recall that this is compared to a baseline where the consumption is already minimal for the case of a dwelling, but EMPC has been employed by many housing companies primarily as an energy-saving measure. For the case of the template SH system, the energy savings by switching from weather compensation to EMPC were estimated to be 15%.

Encouraging load shifting using the building inertia cannot be achieved by having a price model with only a consumption component. But when combined with another load-shifting incentive mechanism, e.g., a power subscription or a peak demand component, the multi-objective dilemma between consumption and peak demand arises. Even with our simulations, where a plethora of crude simplifications have been made, the results are very sensitive to the price composition. This observation is particularly relevant for populations where nighttime setbacks are common, corresponding to a large γ in our simulations. The increase in consumption that a large peak demand component brings compared to the baseline is however much smaller than many of the reported savings by employing EMPC in the first place. The worst case on the other end, a very small peak demand component, appears more problematic with a substantially increased peak demand. Thus, from this perspective of the financial incentive, a low peak demand tariff is more questionable than a high.

Another new aspect for the heat supplier that comes with penalizing high peak demand is that prediction of the heat load may be more challenging. With the traditional control methods, the heat load may

accurately be described by the current outdoor temperature and a daily pattern of DHW use. By introducing demand-side load shifting and occupancy schedules, forecasting the heat load by using the same variables may no longer be sufficient for production planning. How this should be handled is however a broad question that is left for future work.

6.3. Benefits of cooperative control

Based on our simulations, the cooperative EMPC is advantageous in populations where a fraction of the buildings have varying occupancy. When comparing the results in Fig. 8 with the statistics of the Swedish district heating systems presented in Section 1, where buildings with a varying occupancy may account for 40% of the buildings, there may be potential for applying cooperative EMPC control. However, with the current price mechanisms, there is no reason for consumers to employ such a strategy. By charging all substations owned by the same housing company as a group, instead of each district heating connection separately, the housing company could potentially benefit from cooperative control of their buildings. Still, there are typically many different housing companies in a single district heating system, and it is not obvious that they have access to a sufficiently heterogeneous set of buildings to make any significant gains from cooperative control. Therefore, this change in billing strategy may be weakly reflected by our simulations where the cooperative EMPC always has access to all buildings connected to the system. In summary, cooperative control may be beneficial, but how it should be implemented to encourage consumers to employ optimal schemes with respect to the full system is left for future work.

7. Conclusions

We have investigated the financial incentives that affect the heat load when EMPC is used for SH systems connected to a district heating network. Our simulations indicate that tweaks to the established price mechanisms used by companies today can be used as a tool for motivating load shifting using the pre-existing building thermal inertia. The crucial element for achieving this is to have a sufficiently large tariff for penalizing a high peak demand, which will increase consumption slightly as a side effect. By charging the peak demand for multiple SH systems together, it is possible and financially favorable to reduce the overall peak demand further via cooperative control of the multiple SH systems.

We obtain these results by simulating the heat demand of different populations of two buildings, one dwelling and one workplace, where a demand-side EMPC is employed in each of them. The dwelling must maintain a comfortable climate all of the time, while the workplace must only do so 8 a.m. to 5 p.m. during the workdays. By varying the relative size of the dwelling and the workplace, we obtain different populations. We also consider different price compositions consisting of a consumption component and a peak demand component, where the relative tariffs between the two are varied. We also implement two different calculation methods for the peak demand component: one as the weekly maximum per building, where the separable EMPC formulation is used, or the weekly accumulated maximum, where the cooperative formulation is used.

A peak demand component that is large relative to the consumption component incentivizes indirect demand side management via load shifting using the building inertia. On the other way around, a large consumption component motivates full utilization of the allowed nighttime setback in the workplace, which also comes with a high peak demand. In our simulations, a price composition that primarily penalizes high peak demand gives a 10%-20% reduction in peak demand relative to a baseline, to only a slightly larger consumption than our baseline. With a price composition that only penalizes high consumption, corresponding to what is implemented among many Swedish district heating compa-

nies, the consumption is reduced at the cost of large peak demand, especially when the workplace is of substantial size.

CRedit authorship contribution statement

Henrik Håkansson: Software, Visualization, Writing – original draft, Formal analysis, Investigation. **Magnus Önnheim:** Conceptualization, Supervision, Writing – review & editing, Investigation, Methodology, Software. **Emil Gustavsson:** Funding acquisition, Supervision, Investigation, Writing – review & editing. **Mats Jirstrand:** Funding acquisition, Supervision, Writing – review & editing, Investigation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors do not have permission to share data.

Acknowledgements

This work was financially supported by the Swedish Energy Agency (contract 50241-1: *Optimization of building heating systems using artificial intelligence and model predictive control*). We also acknowledge Jonas Tannerstad and Örebrobostäder AB for providing access to building population data and for fruitful discussions.

Appendix A. Template SH system

The template SH system used in this study is located in Örebro, Sweden, and was constructed in the early 1970s. It consists of multiple buildings connected to the same district heating substation, which is schematically visualized in Fig. A.10. The substation consists of multiple heat exchangers so that heating power is supplied for both SH and circulation of DHW. In Fig. A.10 we also see the measurement points that allow sampling the total heating power, of both the SH and DHW demand, via the heat meter on the primary side, together with the SH supply temperature and SH return temperature.

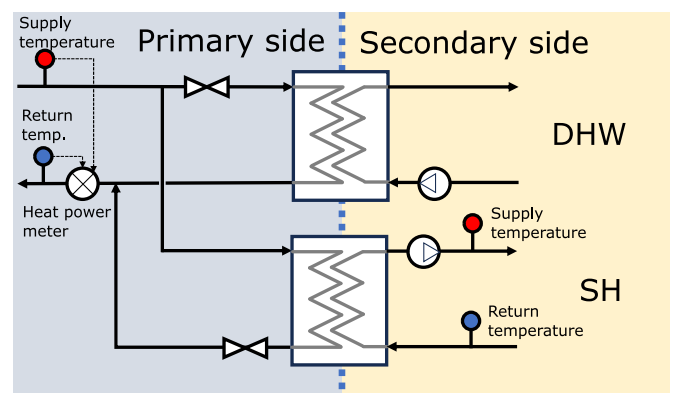


Fig. A.10. Schematic of the template building heating system and the substation. The district heating supply is used for both DHW and SH on the secondary side, with two separate heat exchangers and valves controlling the primary flow. The supply temperature and the return temperature are measured on both the primary side and the SH on the secondary side. The total heating power, for both DHW and SH, is measured by a heat meter on the primary side, multiplying the flow rate with the difference of primary supply and return temperature.




Fig. A.11. Left plot: Scatter plot of sampled supply temperature on the x-axis and return temperature on the y-axis at the same time steps. Right plot: scatter plot (blue dots) of the difference between the supply temperature and the return temperature of the SH system on the x-axis and the heating power, measured on the primary side, on the y-axis. The red line is a linear regression line fitted to the data with least squares.




Fig. A.12. Upper plot: Scatter of samples during the heating season of estimated DHW usage calculated from the residual of the linear regression in the right-hand side panel of Fig. A.11, with the corresponding hour of the day on the x-axis. Lower plot: Scatter of samples of the measured heat demand during summer, when SH is turned off, with the corresponding hour of the day on the x-axis. In both plots, the blue line shows the hour-wise mean of the estimated DHW usage of the upper plot, and the red line is the hour-wise mean of the summer heat demand in the lower plot.

In the SH system, there are no TRVs and the flow is close to constant. Thus, the variation of the SH supply temperature is the only active heat demand control mechanism. There is a significant spread in the indoor temperature among different apartments, typically with a standard deviation of about 0.4 °C, but the supply temperature can still be used to control the total heat demand, which can be seen in Fig. A.11. From the left-hand side plot in Fig. A.11, we see that the SH return temperature can accurately be modeled as a linear function from the SH supply temperature. In the right-hand side plot of the same figure, we see similarly that the total heating power is well correlated with the difference between the SH supply temperature and the SH return temperature. However, there is a remarkable variance in the second regression, which can mainly be explained by the DHW usage. This explanation can be justified by inspecting Fig. A.12, where the daily patterns of the DHW usage as estimated by the residuals of the regression are compared to the total heat demand when the SH is turned off during the summer. In both of these cases, similar daily patterns arise, although with some variance from day to day.

Before the winter of 2019/2020, the supply temperature was controlled using weather compensation. By utilizing temperature sensors




Fig. A.13. Average indoor temperature of the template SH system before and after switching from weather compensation to EMPC, which occurred on 12/24/2019.

inside the buildings, one in each apartment, and remote control of the supply temperature an EMPC algorithm was employed in December 2019. After the employment of the EMPC, which is visualized in Fig. A.13, the average indoor temperature could be reduced and kept on a steady level compared to the situation with weather compensation.

Appendix B. Simulation model

The recursive identification of parameters in Eq. (1b) uses a Kalman filter approach to obtain the parameter estimate $\hat{\Theta}_t = [\theta, c_{0:24}^T, v_{0:H}^T]^T$ at time t according to

$$\hat{\Theta}_t = \hat{\Theta}_{t-1} + K_t \varepsilon_t,$$

where K_t is the Kalman gain matrix and

$$\varepsilon_t = x_t - x_{t-1} - f(y_{t-1}, x_{t-1}, u_{t-1:t-1-H}; \theta, c_{0:24}, v_{0:H}) \Delta t.$$

The Kalman gain is given by

$$K_t = \frac{P_{t-1} \hat{\Theta}_t}{R + \hat{\Theta}_t^T P_{t-1} \hat{\Theta}_t}, \quad P_t = P_{t-1} - \frac{P_{t-1} \hat{\Theta}_t \hat{\Theta}_t^T P_{t-1}}{R + \hat{\Theta}_t^T P_{t-1} \hat{\Theta}_t} + Q,$$

where $Q = 10^{-13} \cdot \mathbb{1}$, and $R = 10^{-3}$. Θ_0 is obtained by ordinary least squares (OLS) estimate using data from one month before the first simulation week and P_0 is the covariance matrix of that OLS estimate. This strategy is applied recursively for each time step, meaning that new parameter estimates are computed every time step. However, we use only 14 different sets of parameter estimates: one corresponding to the last date or each simulation week.

In Fig. B.14, we see that the estimated heat transfer coefficients \hat{v} using this estimation method become negative for some of the lags and therefore fail to preserve a fully realistic physical interpretation per lag. However, the sum of all lag coefficients, corresponding to an integrated impulse response of the heating, is still positive. For the outdoor temperature coefficient, the estimate remains positive for all weeks.

Comparing a dwelling and a workplace, we expect a difference in the daily seasonality, i.e., the hour-wise constants. Therefore, the workplace model constants are modified to have one value for every hour during the work hours and another value outside these hours. These values are selected so that the mean and standard deviation of the constants over the day are the same as for the dwelling constants. The other model parameters have the same value for both the workplace and the dwelling. The estimated hour-wise constants of the dwelling model and the modified constants of the workplace are shown in Fig. B.15.




Fig. B.14. Heat transfer coefficients of the simulation model in Eq. (1b) per week for the 6 different lags. The integrated impulse response is given by summing the coefficients of all lags of the same week.

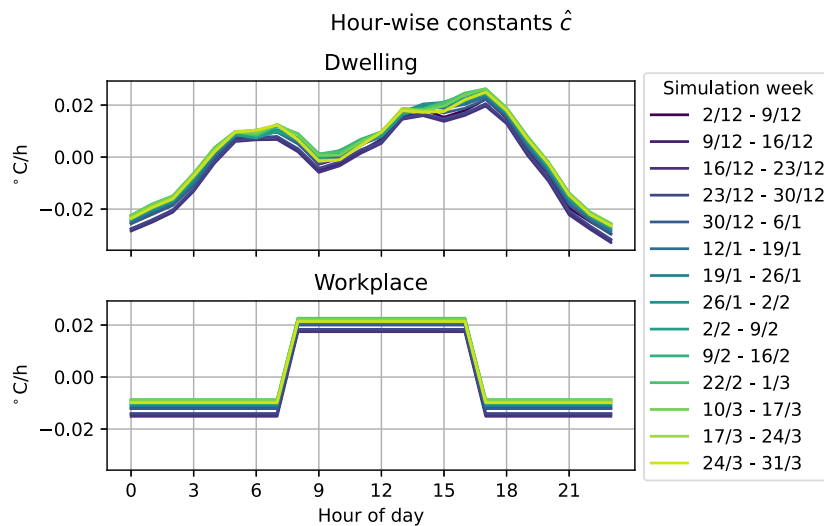


Fig. B.15. Hour-wise constants \hat{c} of the simulation model Eq. (1b) for the dwelling and the workplace for all simulation weeks.

References

- [1] S. Werner, District heating and cooling, in: C.J. Cleveland (Ed.), Encyclopedia of Energy, Elsevier, New York, 2004, pp. 841–848, <https://www.sciencedirect.com/science/article/pii/B012176480X00214X>.
- [2] H. Lund, S. Werner, R. Wiltshire, S. Svendsen, J.E. Thorsen, F. Hvelplund, B.V. Mathiesen, 4th Generation District Heating (4GDH): integrating smart thermal grids into future sustainable energy systems, *Energy* 68 (2014) 1–11, <https://doi.org/10.1016/j.energy.2014.02.089>, <https://www.sciencedirect.com/science/article/pii/S0360544214002369>.
- [3] C. Delmastro, F. Brience, F. Voswinkel, Y. Monschauer, R. Martinez-Gordon, J. Lynch, District heating - energy system, <https://www.iea.org/energy-system/buildings/district-heating>, 2023.
- [4] A. Vandermeulen, B. van der Heijde, L. Helsen, Controlling district heating and cooling networks to unlock flexibility: a review, *Energy* 151 (2018) 103–115, <https://doi.org/10.1016/j.energy.2018.03.034>, <https://www.sciencedirect.com/science/article/pii/S0360544218304328>.
- [5] H. Gadd, S. Werner, 21 - thermal energy storage systems for district heating and cooling, in: L.F. Cabeza (Ed.), Advances in Thermal Energy Storage Systems, second edition, in: Woodhead Publishing Series in Energy, Woodhead Publishing, 2021, pp. 625–638, <https://www.sciencedirect.com/science/article/pii/B9780128198858000218>.
- [6] S. Frederiksen, S. Werner, District Heating and Cooling, Studentlitteratur AB, Lund, 2013.
- [7] E. Guelpa, V. Verda, Demand response and other demand side management techniques for district heating: a review, *Energy* 219 (2021) 119440, <https://doi.org/10.1016/j.energy.2020.119440>, <https://www.sciencedirect.com/science/article/pii/S0360544220325470>.
- [8] G. Serale, M. Fiorentini, A. Capozzoli, D. Bernardini, A. Bemporad, Model Predictive Control (MPC) for Enhancing Building and HVAC System Energy Efficiency: Problem Formulation, Applications and Opportunities, *Energies*, vol. 11, 2018, p. 631, <https://www.mdpi.com/1996-1073/11/3/631>.
- [9] J. Song, F. Wallin, H. Li, District heating cost fluctuation caused by price model shift, *Appl. Energy* 194 (2017) 715–724, <https://doi.org/10.1016/j.apenergy.2016.09.073>, <https://www.sciencedirect.com/science/article/pii/S0360261916313757>.
- [10] D. Olsson, Behovsanpassad värmereglering, prestudy 2022:052, bebostad, <https://www.bebostad.se/projekt/avslutade-projekt/2022/2022-behovsanpassad-varmereglering-forstudie>, 2022.
- [11] D. Olsson, Modellbaserad styrning av värmesystem baserat på prognostiserat väder - En jämförelse med andra reglerstrategier, Ph.D. thesis, Chalmers University of Technology, 2014, <https://research.chalmers.se/en/publication/209907>.
- [12] Statistiska Centralbyrån (SCB), Årlig energistatistik (el, gas och fjärrvärme) 2020, <https://www.scb.se/en0105>, 2022.
- [13] C.A. Thilker, P. Bacher, H. Madsen, Learnings from experiments with MPC for heating of older school building, *E3S Web Conf.* 362 (2022) 12004, <https://doi.org/10.1051/e3sconf/202236212004>, https://www.e3s-conferences.org/articles/e3sconf/abs/2022/29/e3sconf_bsn2022_12004/e3sconf_bsn2022_12004.html.
- [14] D.S. Østergaard, S. Svendsen, Experience from a practical test of low-temperature district heating for space heating in five Danish single-family houses from the 1930s, *Energy* 159 (2018) 569–578, <https://doi.org/10.1016/j.energy.2018.06.142>, <https://www.sciencedirect.com/science/article/pii/S0360544218312076>.
- [15] S. Månsson, P.-O. Johansson Kallioniemi, M. Thern, T. Van Oevelen, K. Sernhed, Faults in district heating customer installations and ways to approach them: experiences from Swedish utilities, *Energy* 180 (2019) 163–174, <https://doi.org/10.1016/j.energy.2019.04.220>, <https://www.sciencedirect.com/science/article/pii/S0360544219308606>.
- [16] T. Benakopoulos, R. Salenbien, D. Vanhoudt, S. Svendsen, Improved Control of Radiator Heating Systems with Thermostatic Radiator Valves without Pre-Setting Function, *Energies*, vol. 12, 2019, p. 3215, <https://www.mdpi.com/1996-1073/12/17/3215>.
- [17] J. Kensby, A. Trüschel, J.-O. Dalenbäck, Potential of residential buildings as thermal energy storage in district heating systems - results from a pilot test, *Appl. Energy* 137 (2015) 773–781, <https://doi.org/10.1016/j.apenergy.2014.07.026>, <https://www.sciencedirect.com/science/article/pii/S0360261914007077>.
- [18] F. Wernstedt, P. Davidsson, C. Johansson, Demand side management in district heating systems, in: Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS '07, Association for Computing

- Machinery, New York, NY, USA, 2007, pp. 1–7, <https://dl.acm.org/doi/10.1145/1329125.1329454>.
- [19] P. Ala-Kotila, T. Vainio, J. Heinonen, Demand Response in District Heating Market—Results of the Field Tests in Student Apartment Buildings, Smart Cities 3, Multidisciplinary Digital Publishing Institute, 2020, pp. 157–171, <https://www.mdpi.com/2624-6511/3/2/9>.
- [20] E. Guelpa, L. Marincioni, S. Deputato, M. Capone, S. Amelio, E. Pochettino, V. Verda, Demand side management in district heating networks: a real application, Energy 182 (2019) 433–442, <https://doi.org/10.1016/j.energy.2019.05.131>, <https://www.sciencedirect.com/science/article/pii/S036054421931014X>.
- [21] D. Romanchenko, J. Kensby, M. Odenberger, F. Johnsson, Thermal energy storage in district heating: centralised storage vs. storage in thermal inertia of buildings, Energy Convers. Manag. 162 (2018) 26–38, <https://doi.org/10.1016/j.enconman.2018.01.068>, <https://www.sciencedirect.com/science/article/pii/S0196890418300803>.
- [22] D.F. Dominković, P. Gianniu, M. Münster, A. Heller, C. Rode, Utilizing thermal building mass for storage in district heating systems: combined building level simulations and system level optimization, Energy 153 (2018) 949–966, <https://doi.org/10.1016/j.energy.2018.04.093>, <https://www.sciencedirect.com/science/article/pii/S0360544218307060>.
- [23] A. Kathirgamanathan, M. De Rosa, E. Mangina, D.P. Finn, Data-driven predictive control in urban district heating networks, Appl. Energy 230 (2018) 506–518, <https://doi.org/10.1016/j.apenergy.2018.08.105>, <https://www.sciencedirect.com/science/article/pii/S0306261918312765>.
- [24] H. Cai, C. Ziras, S. You, R. Li, K. Honoré, H.W. Bindner, Demand side management in urban district heating networks, Appl. Energy 230 (2018) 506–518, <https://doi.org/10.1016/j.apenergy.2018.08.105>, <https://www.sciencedirect.com/science/article/pii/S0306261918312765>.
- [25] R.E. Hedegaard, M.H. Kristensen, T.H. Pedersen, A. Brun, S. Petersen, Bottom-up modelling methodology for urban-scale analysis of residential space heating demand response, Appl. Energy 242 (2019) 181–204, <https://doi.org/10.1016/j.apenergy.2019.03.063>, <https://www.sciencedirect.com/science/article/pii/S0306261919304726>.
- [26] S. Werner, District heating and cooling in Sweden, Energy 126 (2017) 419–429, <https://doi.org/10.1016/j.energy.2017.03.052>, <https://www.sciencedirect.com/science/article/pii/S0360544217304140>.
- [27] J. Song, F. Wallin, H. Li, B. Karlsson, Price models of district heating in Sweden, Energy Proc. 88 (2016) 100–105, <https://doi.org/10.1016/j.egypro.2016.06.031>, <https://www.sciencedirect.com/science/article/pii/S1876610216300959>.
- [28] E.F. Camacho, C. Bordons, Model Predictive Control, Advanced Textbooks in Control and Signal Processing, Springer London, London, 2007, <http://link.springer.com/10.1007/978-0-85729-398-5>.
- [29] J. Dragoña, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E.P. Ollé, J. Oravec, M. Wetter, D.L. Vrabie, L. Helsen, All you need to know about model predictive control for buildings, Annu. Rev. Control 50 (2020) 190–232, <https://doi.org/10.1016/j.arcontrol.2020.09.001>, <https://www.sciencedirect.com/science/article/pii/S1367578820300584>.
- [30] V. Amato, R.E. Hedegaard, M.D. Knudsen, S. Petersen, Room-level load shifting of space heating in a single-family house – a field experiment, Energy Build. 281 (2023) 112750, <https://doi.org/10.1016/j.enbuild.2022.112750>, <https://www.sciencedirect.com/science/article/pii/S0378778822009215>.
- [31] M.D. Knudsen, L. Georges, K.S. Skeie, S. Petersen, Experimental test of a black-box economic model predictive control for residential space heating, Appl. Energy 298 (2021) 117227, <https://doi.org/10.1016/j.apenergy.2021.117227>, <https://www.sciencedirect.com/science/article/pii/S0306261921006498>.
- [32] N. Aoun, R. Bavière, M. Vallée, A. Aurousseau, G. Sandou, Modelling and flexible predictive control of buildings space-heating demand in district heating systems, Energy 188 (2019) 116042, <https://doi.org/10.1016/j.energy.2019.116042>, <https://www.sciencedirect.com/science/article/pii/S0360544219317360>.
- [33] J. Hou, H. Li, N. Nord, Nonlinear model predictive control for the space heating system of a university building in Norway, Energy 253 (2022) 124157, <https://doi.org/10.1016/j.energy.2022.124157>, <https://www.sciencedirect.com/science/article/pii/S036054422201060X>.
- [34] H.T. Walnum, I. Sartori, M. Bagle, Model predictive control of district heating substations for flexible heating of buildings, in: International Conference Organised by IBPSA-Nordic, 13th–14th October 2020, OsloMet, BuildSIM-Nordic 2020, Selected Papers, SINTEF Academic Press, 2020.
- [35] Z. Afroz, G. Shafiqullah, T. Urme, G. Higgins, Modeling techniques used in building HVAC control systems: a review, Renew. Sustain. Energy Rev. 83 (2018) 64–84, <https://doi.org/10.1016/j.rser.2017.10.044>, <https://www.sciencedirect.com/science/article/pii/S1364032117314193>.
- [36] P. Bacher, H. Madsen, Identifying suitable models for the heat dynamics of buildings, Energy Build. 43 (2011) 1511–1522, <https://doi.org/10.1016/j.enbuild.2011.02.005>, <https://www.sciencedirect.com/science/article/pii/S0378778811000491>.
- [37] I. Hazyuk, C. Ghiaus, D. Penhouet, Optimal temperature control of intermittently heated buildings using model predictive control: Part I – building modeling, Build. Environ. 51 (2012) 379–387, <https://doi.org/10.1016/j.buildenv.2011.11.009>, <https://www.sciencedirect.com/science/article/pii/S0360132311003933>.
- [38] P.-D. Moroşan, R. Bourdais, D. Dumur, J. Buisson, Building temperature regulation using a distributed model predictive control, Energy Build. 42 (2010) 1445–1452, <https://doi.org/10.1016/j.enbuild.2010.03.014>, <https://www.sciencedirect.com/science/article/pii/S0378778810000915>.
- [39] S. Přívara, J. Šíroký, L. Ferkl, J. Cigler, Model predictive control of a building heating system: the first experience, Energy Build. 43 (2011) 564–572, <https://doi.org/10.1016/j.enbuild.2010.10.022>, <https://www.sciencedirect.com/science/article/pii/S0378778810003749>.
- [40] I. Hazyuk, C. Ghiaus, D. Penhouet, Optimal temperature control of intermittently heated buildings using model predictive control: Part II – control algorithm, Build. Environ. 51 (2012) 388–394, <https://doi.org/10.1016/j.buildenv.2011.11.008>, <https://www.sciencedirect.com/science/article/pii/S0360132311003921>.
- [41] C.A. Thilker, P. Bacher, H.G. Bergsteinsson, R.G. Junker, D. Cali, H. Madsen, Non-linear grey-box modelling for heat dynamics of buildings, Energy Build. 252 (2021) 111457, <https://doi.org/10.1016/j.enbuild.2021.111457>, <https://www.sciencedirect.com/science/article/pii/S0378778821007416>.
- [42] L. Ljung, Recursive identification algorithms, Circuits Syst. Signal Process. 21 (2002) 57–68, <https://doi.org/10.1007/BF01211651>.