Enabling smart control by optimally managing the State of Charge of district heating networks

Enabling smart control by optimally managing the State of Charge of district heating networks

Costanza Saletti[a*], Nathan Zimmerman[b], Mirko Morini[a,c], Konstantinos Kyprianidis[b], Agostino Gambarotta[a,c]

[a] *Department of Engineering and Architecture, University of Parma, Parco Area delle Scienze 181/A, 43124 Parma, Italy*
[b] *Department of Automation in Energy and Environment, School of Business, Society and Engineering, Mälardalen University, Box 883, Västerås, 72123, Sweden*
[c] *Center for Energy and Environment (CIDEA), University of Parma, Parco Area delle Scienze 42, 43124 Parma, Italy*

* Corresponding author: costanza.saletti@unipr.it

Abstract

Digitalization and smart control of district heating networks are emerging as key features to make these systems flexible and optimal. However, since effective and scalable methods for large-scale systems are currently unavailable, the implementation of smart controllers can be challenging and time-consuming. This is addressed herein by proposing a novel approach to include the thermal capacity of the connected buildings in the optimal control of large-scale heating networks. A reduced-order model of the aggregated regions supplied by the Västerås network is used to define their State of Charge, which is exploited to store or retrieve energy when convenient, while maintaining indoor comfort. This concept is included in a Model Predictive Controller that optimizes power plant management and heat distribution. The results show that the controller successfully shaves heat supply peaks to different regions up to 16% and reduces the difference between distribution and soil temperature up to 20%. At the same time, the return temperature is kept close to the set-point of 35 °C, which is lower than the historical operation and further reduces distribution heat losses. The procedure can be easily replicated to optimize systems of different sizes and to support their transition to efficient, smart district heating networks.

Keywords: District Heating Network; State of Charge; building heat capacity; Model Predictive Control; optimal management; scalability

1. Introduction

District Heating Networks (DHNs) are one of the most promising means to supply heat efficiently [1] and improve air quality in urban areas [2]. Indeed, by supplying hot water to consumers by means of a network of pipelines, DHNs offer several advantages [3]: (i) the economy-of-size, which determines a higher efficiency and lower cost of central heat supply devices compared to low-size boilers, (ii) the possibility to exploit local strategic resources, such as renewables and excess heat from high-temperature processes, (iii) flexibility, and (iv) a lower local environmental impact in terms of emissions of particulate matter and nitrogen oxides. The share of heat supply by DHNs in the heating sector is significant in many European countries (e.g. both Sweden and Finland have a 50% market share) but should be further improved in other areas (e.g. Italy has a market share lower than 5%). According to a report from the Heat Roadmap Europe project [4], district heating has the chance to provide 50% of heating demand in 14 countries and to allow 30% energy saving, showing its importance in the energy sector. However, DHNs have to be integrated in a global energy scenario that is rapidly changing [5], due to the growing penetration of discontinuous sources and distributed generation devices (e.g. heat pumps [6] and solar panels). This might lead to systems in which production and demand are not synchronous. Hence, it is necessary to exploit smart approaches and innovative control strategies to remove the barriers of location and time in thermal energy distribution networks [7] and to achieve system *flexibility*. According to Vandermeulen *et al.* [8], as far as energy systems are concerned, flexibility is "the ability to speed up or delay the injection or extraction of energy into or from a system". Hence, the flexibility requires the system to have a thermal capacity which acts as a buffer between energy production and actual delivery.

This concept is recalled by Hennessy *et al.* [9], who propose a review of research works that exploit short-term storage to improve the flexibility of thermal energy grids. The types of thermal energy storage considered in this study are centralized storage tanks and storage in the network pipes. The latter consists of increasing the temperature of the water in the distribution network to preload it and

optimize heat load generation, but it is based on an approximated calculation of the storage capacity. Moreover, this method goes against the actual trend of lowering the network temperature in both design [10] and optimization studies [11] in order to reduce heat losses – with consequent increased efficiency and more affordable costs – and achieve 4th generation DHNs [12]. Therefore, it is worth carrying out more research on these topics [13] and to investigate more preferable types of storage in energy networks.

Another review [14] classifies the storage solutions for DHNs according to the occurring physical phenomenon (sensible, latent or chemical storage), the duration (short-term or long-term) and the layout (distributed or centralized storage). The review mainly focuses on devices specifically designed to be included in the system (e.g. storage tanks), and introduces only marginally alternative methods such as exploiting the pipeline capacity or the building inertia, or modifying the building request (also known as demand side management [15]). The inclusion of these aspects in the optimal management of the building or district might lead to significant improvements in the energy efficiency of these systems.

While some works investigate the potential role of the water in the networks [16], a more promising direction is the analysis of the thermal capacity of buildings [17], exploited as individual devices that are able to modulate heating power [18] and provide more flexibility to the electrical grid [19], or as distributed storage means within DHNs. They can be exploited to achieve energy efficiency by peak shaving and valley filling, which consist of shaping production and thermal load in such a way that they are kept as constant as possible [20]. These are objectives of studies that develop a model-free technique to control a cluster of thermostatically controlled loads connected to a DHN [21], and that take into account the dynamics of a building that can be charged or discharged to a certain degree depending on the indoor comfort [22]. However, these methods are nor real-time-control-oriented neither applicable on a large-scale, which require novel modeling approaches, such as dealing with simplified energy balances with regional aggregated parameters for heat capacity and insulation [23].

1.1 Smart control

In large-scale DHNs, heat propagation is subjected to considerable delays that can reach several hours, depending on flows and system sizes. For this reason, the temperature control of these systems is typically performed manually according to operator experience, which is non-optimal [24]. Thus, smart control approaches [6] could foster the automation of DHNs and lead to optimal management, lower temperatures and reduced losses. For instance, predictive control showed great potential in multiple simulation studies for achieving the stable operation of the power plant supplying large-scale networks [25], and for scheduling the cold distribution in an integrated energy system comprising source, network and load [26]. An increase in the system scale challenges the computational efficiency of the approach in the latter work, in which the necessity to develop a more efficient solution for large communities is highlighted. Other recent articles propose the smart control of large-scale DHNs by exploiting different automation approaches, such as feed-forward predictive control to reduce the network return temperature [27], and artificial intelligence to operate peak load management and rationalize the use of the production units [28]. However, the proposed methods are frequently case-dependent and their extension to real scenarios is not always straightforward.

A predictive controller based on Model Predictive Control (MPC) with a Dynamic Programming optimization algorithm [29] has been developed [30], tested in simulation cases [31] and prototyped [32] for small-scale real applications such as school complexes and university campuses. In such systems, the reduced dimension of the problem allows each consumer to be individually modeled and identified without compromising the computational feasibility of the approach.

The extension to large-scale DHNs is not straightforward, as the number of buildings increases significantly, and therefore they cannot be considered as separate individual elements. Suitable reduced-order models [33] and optimization algorithms [34] are necessary to promote the aforementioned technique at city-scale. More studies that aim to address these challenges might be essential in the further improvement of DHNs.

1.2 Scope of the work

In the light of the research gaps outlined above, this paper proposes a method to enable the smart control and optimal management of large-scale DHNs by including the thermal capacity of the connected buildings. It is especially applicable to cases where not enough information has been provided about the actual network characteristics, or where the end-users are not monitored at an individual-building scale. Aggregated areas of large-scale DHNs are exploited as thermal storage by assigning them a State of Charge (SoC). This concept is included in an optimization problem which aims to achieve thermal peak shaving and supply temperature reduction. This procedure is finally embedded in a predictive controller that is tested on a detailed dynamic model of the network developed in [27]. The case study is the district heating network of the city of Västerås, located in central Sweden. Nonetheless, the method is suitable for any region of a DHN for which coarse data at the main substation heat exchangers are available.

1.3 Main contribution

The novelty of the work with respect to the current state-of-the-art research is determined by the following main contributions:

- *The State of Charge of district heating network*. According to this concept, the buildings connected to the network are equated to a battery which can be charged/uncharged when it is convenient for the system. The classical thermal-electrical analogy, typically exploited in heat transfer problems [35], is recalled here to explore the potential within DHNs.
- *A two-stage optimization procedure* that is able to manage (i) the thermal energy delivery to different DHN regions according to different cost functions, and (ii) the thermal energy production unit. Plus, it is fast, scalable and suitable for real-time control applications.
- *A smart controller based on Model Predictive Control* that uses the developed model to perform thermal peak shaving and network temperature reduction to increase network

efficiency. Since real-time control is typically applied to individual buildings or to production sites, a strategy that integrates production and distribution can contribute toward the advancement of the technology.

To the best of the authors' knowledge, these aspects have not been treated extensively or within an integrated procedure in the existing literature on DHNs. Such method fosters the exploitation of smart and digital technologies in order to bring low temperature, optimized DHNs to maturity.

2. Methods

This section summarizes the research methods used in this work to develop a smart controller for large-scale DHNs with the ability to exploit consumer thermal capacity to achieve peak shaving and reduce network temperature. This procedure is illustrated in the block diagram in Figure 1.

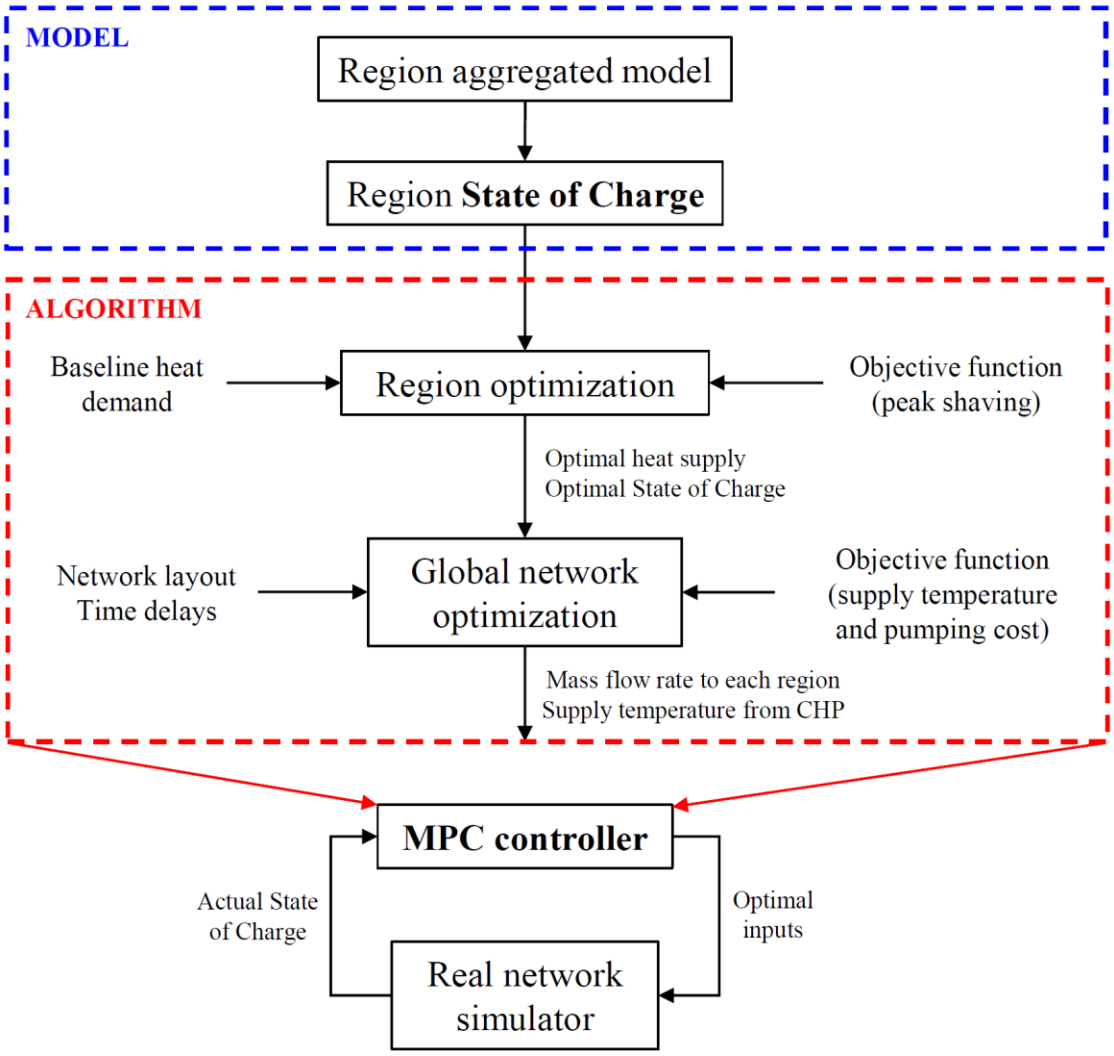

Figure 1. Block diagram of the method for the development of the smart controller for a large-scale DHN. The procedure includes the development of the region model and optimization algorithm, which are embedded in the MPC controller.

The method involves neither extensive geographical details nor country-specific regulations, thus it can be easily applied to systems in different countries and areas. In this work, its demonstration is carried out on the DHN of the city of Västerås, in central Sweden. The Västerås DHN is supplied by a centralized production site – comprising a waste-to-energy Combined Heat and Power plant (CHP),

back-up boilers and thermal energy storage tanks – and distributes hot water to the city center and to six main peripheral areas within the county, namely Surahammar, Skultuna, Rönnby, Tillberga, Barkarö and Hallstahammar. A schematic representation of the system is given in Figure 2. It is possible to notice that the distribution network can be split into nine main pipelines. The total pipeline lengths to reach the external regions with reference to the central power plant are reported in Table 1.

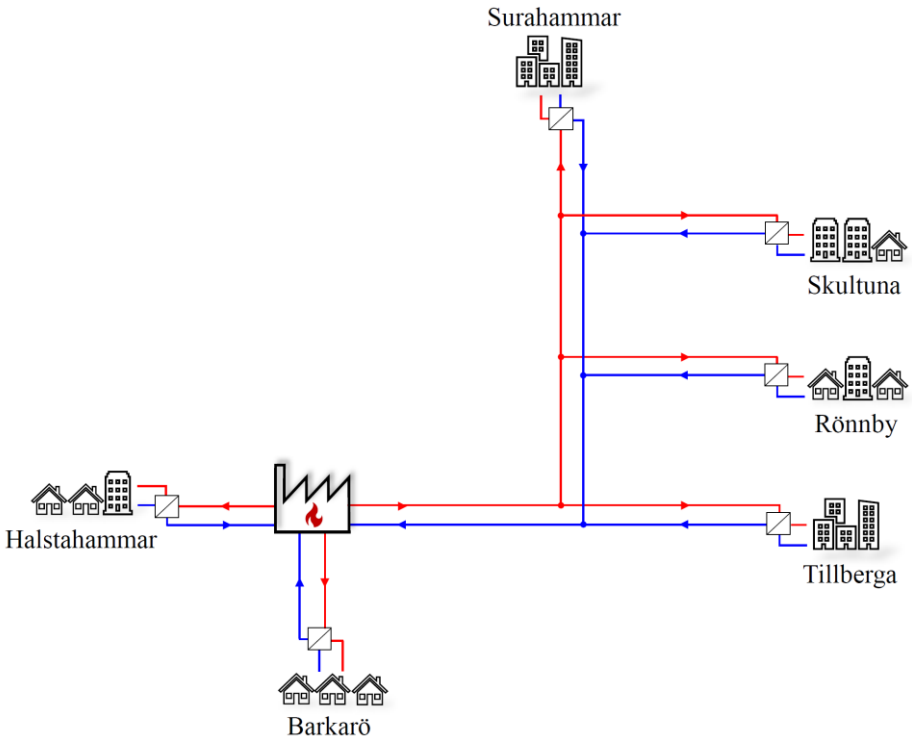

Figure 2. Schematic representation of the Västerås district heating network.

Table 1. External regions of the Västerås district heating network and related pipeline length with reference to the centralized production site.

Region	Pipeline length [km]
Surahammar	31
Skultuna	18.5
Rönnby	11.5
Tillberga	14.5
Barkarö	5.5
Hallstahammar	9

According to the scale-free, simplified model developed in [23], each peripheral region of the network is considered as a single aggregated user with a given mass at a uniform temperature *T*, which is an equivalent representation of its energy content. The model is represented by the thermal power balance in Eq. (1):

$$\frac{dT}{dt} = -\frac{U}{C}(T - T_{\text{ext}}) + \frac{1}{C}\dot{Q} \tag{1}$$

where T_{ext} is the outdoor temperature and $\dot{Q}$ is the thermal power actually transferred to each region through the substation heat exchanger. The aggregated heat transfer coefficient *U* and the aggregated heat capacity *C* of the region are identified with historical data available on an hourly basis at the main substations as reported in [23]. Since an optimization algorithm has to compute the system model multiple times to find the optimal solution, such a simplified model is required to implement an MPC controller in real cases due to its low computational effort. The model is used to calculate the heat storage potential of the region, by considering the amount of heat that can be added or retrieved from the baseline load in order to have an increase or decrease in temperature that is acceptable to the users, and by translating it into the SoC of the aggregated region. This constitutes the "model block" (Figure 1).

The historical heat consumption of the city center is also available. However, since the center is relatively close to the power plant, it is easier to control its heat supply and comfort fulfillment, compared to the farthermost areas. Therefore, it is considered to be supplied according to the historical demand and is not included in the optimization.

The "algorithm block" (described in Section 3) regards the formulation and analysis of an optimization algorithm that comprises two steps:

1. Optimization of the thermal power to be delivered to each separate region with the goal of shaving the supply peaks over a future prediction horizon (e.g. three days). This is achieved

by exploiting each region as a battery and, therefore, by varying its SoC. The optimal heat that has to be supplied to each region is passed to step 2.

2. Optimization of (i) the water mass flow rates in the various branches of the network and (ii) the supply temperature of the CHP plant, with the objective of minimizing the network temperature while keeping pumping costs down. In this case, the optimal heat calculated at the previous step has to be supplied to each region as a constraint of the optimization problem over the entire prediction horizon.

The model block and algorithm block are incorporated in an economic MPC strategy [36] for the actual real-time control of the network. This consists of predicting the behavior of the system in a future time prediction horizon by means of the simplified model of the system itself and optimizing system management over the given horizon by solving a dynamic optimization problem [37]. The system is actually controlled with the first element of the optimal control sequence. At every given time-step, the system variables and the prediction of the disturbances are updated and the problem is solved again [30]. This continuous updating makes it possible to cope with the model approximations and with the influence of the disturbances.

Similarly to the small-scale case study presented in [31], the feasibility of the developed method is finally tested in a Model-in-the-Loop (MiL) application in MATLAB®/Simulink®. The MPC block represents the actual controller and is used to control a detailed dynamic model of the real distribution network, which is developed and validated in [27]. This simulates the heat propagation of the real network and, at every calculation step, returns the estimation of the actual SoC for each region, which is then used as a new initial condition for solving the next optimization problem at the following step.

3. Double-stage optimization problem

This section outlines the optimization procedure of the large-scale DHN. Firstly, the new concept of SoC of a region is explained and used to represent the potential of storage in the consumer heat capacity. Secondly, the heat supply to each region is optimized by means of a linear programming algorithm with the aim of shaving the heat distribution peaks. Lastly, the global optimization problem that calculates the optimal supply temperature and mass flow rates in the whole network is described.

3.1 State of Charge concept

A straightforward way to verify the indoor thermal comfort of buildings is to implement internal temperature monitoring which, in large networks, can be invasive and expensive. For this reason, other approaches are required to tackle and optimize systems such as the Västerås DHN.

Numerous methods to accurately predict the thermal load of the consumers in a DHN are available in the literature: from a simple white-box model that describes heat demand as a function of outdoor temperature and social component [38] to modern black-box models with machine learning techniques (e.g. neural networks [39] and deep learning [40]) based on historical data. The advantages and disadvantages of these main model approaches [41] are briefly introduced in Table 2.

It is possible to consider the historical or predicted load of each region as the baseline thermal power $\dot{Q}_{\text{base}}$, which guarantees the thermal comfort. Therefore, the following fundamental assumption has to be made in this work: the thermal load – represented in this case through hourly historical data of the main substations of the network – has the effect of maintaining precisely the required indoor comfort temperature in the aggregated region. The possibility to store/retrieve thermal energy in/from the consumer mass is based on the deviation from this comfort temperature. This approach is similar to that proposed in [42], in which the DHN heat load is an exogenous input and is required to maintain the indoor temperature at a set-point. Only the deviations from this heat load are considered to have an effect on the temperature variations.

Table 2. Main model approaches for the prediction of thermal load in district heating networks, with advantages and disadvantages.

Approach	Description	Advantages	Disadvantages
White-box	Thermal load forecast based on physical principles of buildings (conservation equations and detailed characteristic)	Detailed dynamic simulation in each possible condition	Costly, time consuming, requires high computational time and detailed information on all buildings
Black-box	Thermal load forecast based on available data of building historical behavior with statistical and machine learning methods	Low computational time, good accuracy	Requires large dataset, not suitable for unpredictable events or for conditions other than training set
Gray-box	Thermal load forecast based on parametrized models identified with data, combining physical and data-driven knowledge	Good accuracy, feasible computational time, model parameters have physical meaning	Requires data and expert knowledge on system

Firstly, the model equation Eq. (1) is rewritten with the incremental ratio to give Eq. (2):

$$\dot{Q} = \mathrm{C}\frac{\Delta T}{\Delta t} + \mathrm{U}(T - T_{\mathrm{ext}}) \tag{2}$$

Secondly, the supplied thermal power is split into two contributions: the baseline load $\dot{Q}_{\mathrm{base}}$ which leads to a rise in equivalent temperature equal to ΔT_{base}, and the stored thermal power $\dot{Q}_{\mathrm{stored}}$ which leads to an additional rise – compared to baseline – equal to $\Delta T_{\mathrm{stored}}$. Thus, Eq. (3) is obtained:

$$\dot{Q}_{\mathrm{base}}\Delta t + \dot{Q}_{\mathrm{stored}}\Delta t = C(\Delta T_{\mathrm{base}} + \Delta T_{\mathrm{stored}}) + U\Delta t(T - T_{\mathrm{ext}}) \tag{3}$$

The heat that is actually stored in the region (i.e. Q_{stored}), compared to the baseline heat (i.e. Q_{base}), is calculated in Eq. (4):

$$Q_{\mathrm{stored}} = \dot{Q}_{\mathrm{stored}}\Delta t = \left(\dot{Q} - \dot{Q}_{\mathrm{base}}\right)\Delta t = Q - Q_{\mathrm{base}} \tag{4}$$

where Q is the actual heat supplied by the DHN through the substation heat exchanger.

The maximum heat that can be stored/taken in/from each region is defined by the maximum deviation from the temperature baseline that is acceptable to the users (i.e. that does not compromise the indoor comfort), namely $\Delta T_{\text{stored,max}}$, as expressed in Eq. (5):

$$Q_{\text{stored,max}} = C\Delta T_{\text{stored,max}} \tag{5}$$

In the work by Romanchenko *et al.* [42], the deviation above the set-point to estimate the heat storage capacity of the connected buildings is limited to 1 °C. In the present work, instead, a more conservative assumption is adopted: a maximum temperature deviation $\Delta T_{\text{stored,max}}$ of 0.5 °C is assumed for the calculation of the heat storage potential of the aggregated consumer, so that thermal comfort is not jeopardized. Nevertheless, it is possible to repeat the analysis for different temperature deviations depending on the indoor comfort requirements and legislations related to each specific case.

The SoC of each region is therefore defined by Eq. (6):

$$\text{SoC} = \frac{Q_{\text{stored}}}{Q_{\text{stored,max}} - (-Q_{\text{stored,max}})} + 0.5 = \frac{Q_{\text{stored,max}} + Q_{\text{stored}}}{2Q_{\text{stored,max}}} \tag{6}$$

In this way, the SoC is equal to:

- zero, when the maximum heat in absolute value has been retrieved and a greater subtraction of heat would lead to the violation of thermal comfort constraints;
- 0.5, when no additional heat (compared to baseline) has been stored;
- 1, when the maximum heat has been stored and a further indoor temperature increase cannot be accepted.

The SoC is the state of the system and is influenced by the incoming thermal power that is stored and by the heat losses to the outside. Its variation in time is expressed by the following differential equation:

$$\frac{d(\mathrm{SoC})}{dt} = \frac{\dot{Q}_{\mathrm{stored}}}{2Q_{\mathrm{stored,max}}} - \frac{\dot{Q}_{\mathrm{loss}}}{2Q_{\mathrm{stored,max}}} \tag{7}$$

The second term $\dot{Q}_{\mathrm{loss}}$ represents the additional heat losses compared to baseline – it is positive if the SoC is higher than the baseline, and negative if it is lower. This is due to the fact that the indoor temperature, together with the SoC, is increased or decreased by $\Delta T_{\mathrm{stored}}$. It is possible to formulate these additional heat losses by combining Eqs. (5) and (6) and knowing that the stored heat is equal to the heat capacity coefficient C multiplied by $\Delta T_{\mathrm{stored}}$. Eqs. (8) and (9) are obtained as follows:

$$\Delta T_{\mathrm{stored}} = \Delta T_{\mathrm{stored,max}}(2 \cdot \mathrm{SoC} - 1) \tag{8}$$

$$\dot{Q}_{\mathrm{loss}} = U\Delta T_{\mathrm{stored}} = U\Delta T_{\mathrm{stored,max}}(2 \cdot \mathrm{SoC} - 1) \tag{9}$$

The state function Eq. (7), together with Eqs. (4), (5) and (9), can be written in a discretized form at time-step k as in Eq. (10):

$$\mathrm{SoC}_{\mathrm{k+1}} = \mathrm{SoC}_{\mathrm{k}} + \frac{\left(\dot{Q}_{\mathrm{k}} - \dot{Q}_{\mathrm{base,k}}\right)\Delta t}{2C\Delta T_{\mathrm{stored,max}}} - \frac{U\Delta t}{C}\left(\mathrm{SoC}_{\mathrm{k}} - \frac{1}{2}\right) \tag{10}$$

To sum up, the assumptions that define the SoC concept are detailed as follows:

- each region is an aggregated set of end-users characterized by an equivalent temperature representing its energy content;
- the baseline thermal load maintains the set-point comfort temperature, which can vary over the day depending on customer requirements, corresponding to a region SoC equal to 0.5;
- maximum acceptable deviations from the comfort temperature are limited to ±0.5 °C, related to variations in the SoC in the range 0 to 1;
- these variations are obtained by regulating the heat supplied to each region and storing energy in or retrieving energy from its heat capacity, without altering thermal comfort.

In the light of this, the discretized state function Eq. (10) representing the evolution of the SoC of each region is the core of the optimization problem formulation outlined in Section 3.2.

3.2 Aggregated region optimization

3.2.1 Problem definition

The first stage of the optimization problem aims to optimize the thermal power supplied separately to each region of the DHN. The problem is written in a state-space form for a given prediction horizon. Since the thermal dynamics in large-scale DHNs such as the Västerås network is relatively long and can reach several hours, a time horizon of at least one day is required to capture all its effects. In order to illustrate the results of the optimization algorithm more extensively, a prediction horizon of three days is considered in this section. The discretized state function that represents the variation of the SoC for a time-step is given by Eq. (10). The initial condition of the problem is the known value of the SoC at the beginning of the prediction horizon SoC_0. The input (i.e. manipulated variable) of the system is the actual thermal power sent to the region $\dot{Q}_\mathrm{k}$ over the prediction horizon, while the baseline load (which depends on the external conditions) $\dot{Q}_{\mathrm{base,k}}$ is the disturbance.

Since the SoC, by definition, cannot be lower than 0 or higher than 1 and the supplied thermal power must be positive, it is necessary to include the following constraints:

$$0 \leq \mathrm{SoC} \leq 1 \tag{11a}$$

$$\dot{Q} \geq 0 \tag{11b}$$

By renaming SoC as x (system state), Q as u (system input) and Q_{base} as d (disturbance), and by using the following coefficients for the sake of readability:

$$\alpha = \frac{\Delta t}{2C\Delta T_{\mathrm{stored,max}}} \tag{12a}$$

$$\beta = \frac{U\Delta t}{C} \tag{12b}$$

it is possible to write the dynamic problem given by Eqs. (10–12) in the state-space form:

$$\begin{cases} x_{k+1} = (1-\beta)x_k + \alpha u_k - \left(\alpha d_k - \frac{\beta}{2}\right) \\ x_0 = \mathrm{SoC}_0 \\ 0 \le x_k \le 1 \\ u_k \ge 0 \end{cases} \qquad \forall\, k = 1 \dots N \tag{13}$$

where *N* is the number of time-steps of the prediction horizon.

This is a Dynamic Linear Programming problem that can be transformed [43] into the following Linear Programming (LP) problem with the vector of variables [u_0, x_1, u_1, x_2 … x_{N-1}, u_{N-1}, x_N]:

$$\begin{cases} \min\limits_{u,x} f(u,x) \\ -\alpha u_0 + x_1 = (1-\beta)x_0 - \left(\alpha d_0 - \frac{\beta}{2}\right) \\ (1-\beta)x_k - \alpha u_k + x_{k+1} = -\left(\alpha d_k - \frac{\beta}{2}\right) \\ -x_k \le 0 \\ x_k \le 1 \\ -u_k \le 0 \end{cases} \qquad \forall\, k = 1 \dots N-1 \tag{14}$$

The problem consists of calculating the values of the states and inputs for the time-steps of the prediction horizon that minimize the cost function *f(u,x)*. It can be readily solved with the standard algorithms of LP (e.g. simplex algorithm).

3.2.2 Analysis

The LP problem represented in Eq. (14) is analyzed by implementing different cost functions and by comparing the results. The problem is solved for three representative days (i.e. January 2017) with time-steps of one hour for the region of Surahammar, though similar considerations can be drawn for the other areas taken separately.

Firstly, a preliminary test, namely Case 0, is performed by minimizing the total thermal energy supplied to the region over the prediction horizon. Figure 3 illustrates (i) the optimal thermal power according to this objective function compared to the historical data and (ii) the SoC in both cases. It is worth remembering that the latter is constant and equal to 0.5 in the historical condition, as it

represents the baseline, in which the indoor temperature is maintained at exactly the set-point required for comfort. The obtained solution is trivial: the SoC of the region is brought to the lower boundary as soon as the simulation starts, in order to reduce the required heat as well as the thermal losses.

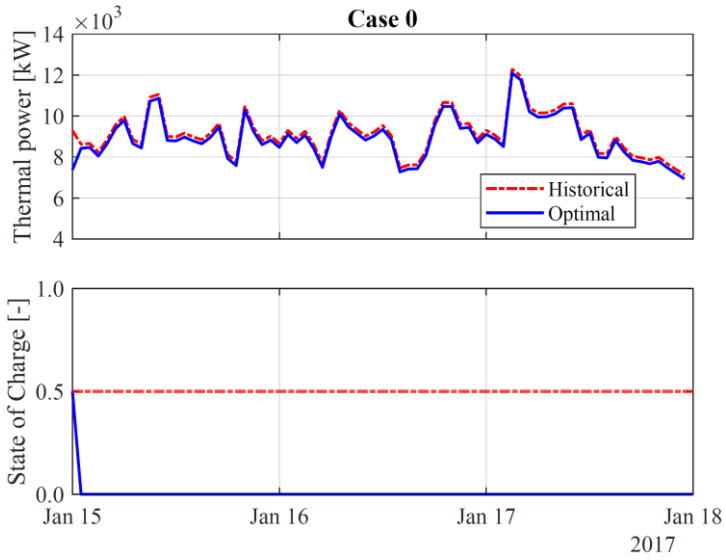

Figure 3. Historical and optimal thermal power and State of Charge as solutions of the Linear Programming problem. Case 0 considers the minimization of total heat.

This is equivalent to lowering the set-point temperature by 0.5 °C, and is impracticable and non-relevant.

In this work, the main interest is to achieve peak shaving but also a smooth evolution of the thermal power supplied, in order not to put the production and distribution systems under stress with rapid changes in the operating parameters. For this purpose, four different cost functions are implemented as follows:

1. minimization of the maximum heat over the prediction horizon, with no regard to the input sequence over time;
2. minimization of the maximum thermal power variation between consecutive time-steps (in absolute value), in order to reduce sudden changes;
3. minimization of the maximum range of input variation, which is defined as the difference between the maximum and minimum values of the thermal power over the prediction horizon;

4. minimization of the total variation of the inputs, obtained from the total of all the thermal power variations (in absolute value) between consecutive time-steps.

These cases and the related results are reported in Table 3 (which also includes Case 0) and illustrated in Figure 4. The percentage peak shaving and total energy reduction, calculated with reference to the historical baseline, are the indicators considered for comparing the outcomes of the analysis. It should be remembered that the objective is only related to the minimization of the thermal peak or of the thermal power variation, and not to energy reduction.

For all the cases, the sequences of optimal inputs (thermal power supplied) and optimal states (SoC of the region) are represented. It is possible to state that the optimization allows the thermal peaks to be reduced to different degrees by exploiting the heat capacity of the regions. As a matter of fact, the SoC fluctuates between the lower and upper boundaries when it is necessary to store or retrieve heat. Cases 1, 3 and 4 present a reduction in peak demand of more than 16%. Besides, it is interesting to note that the objective is achieved with no additional energy consumption (with the exception of Case 2 for which, nevertheless, the increase in consumption is negligible), demonstrating the feasibility of the approach.

Table 3. Results of the optimization of the region heat supply with different objective functions.

Case	Objective	Function	Peak shaving	Energy reduction
0	Minimize the total heat	$\min_{u,x} \sum_k u_k$	1.59%	2.26%
1	Minimize the maximum heat over the horizon	$\min_{u,x} \max_k u_k$	16.21%	0.13%
2	Minimize the maximum absolute value	$\min_{u,x} \max_k \lvert u_k - u_{k-1} \rvert$	13.17%	-0.37%
3	Minimize the maximum range of the input variation	$\min_{u,x} \left(\max_k u_k - \min_k u_k \right)$	16.21%	0.42%
4	Minimize the total variation of the input	$\min_{u,x} \sum_k \lvert u_k - u_{k-1} \rvert$	16.21%	0.25%

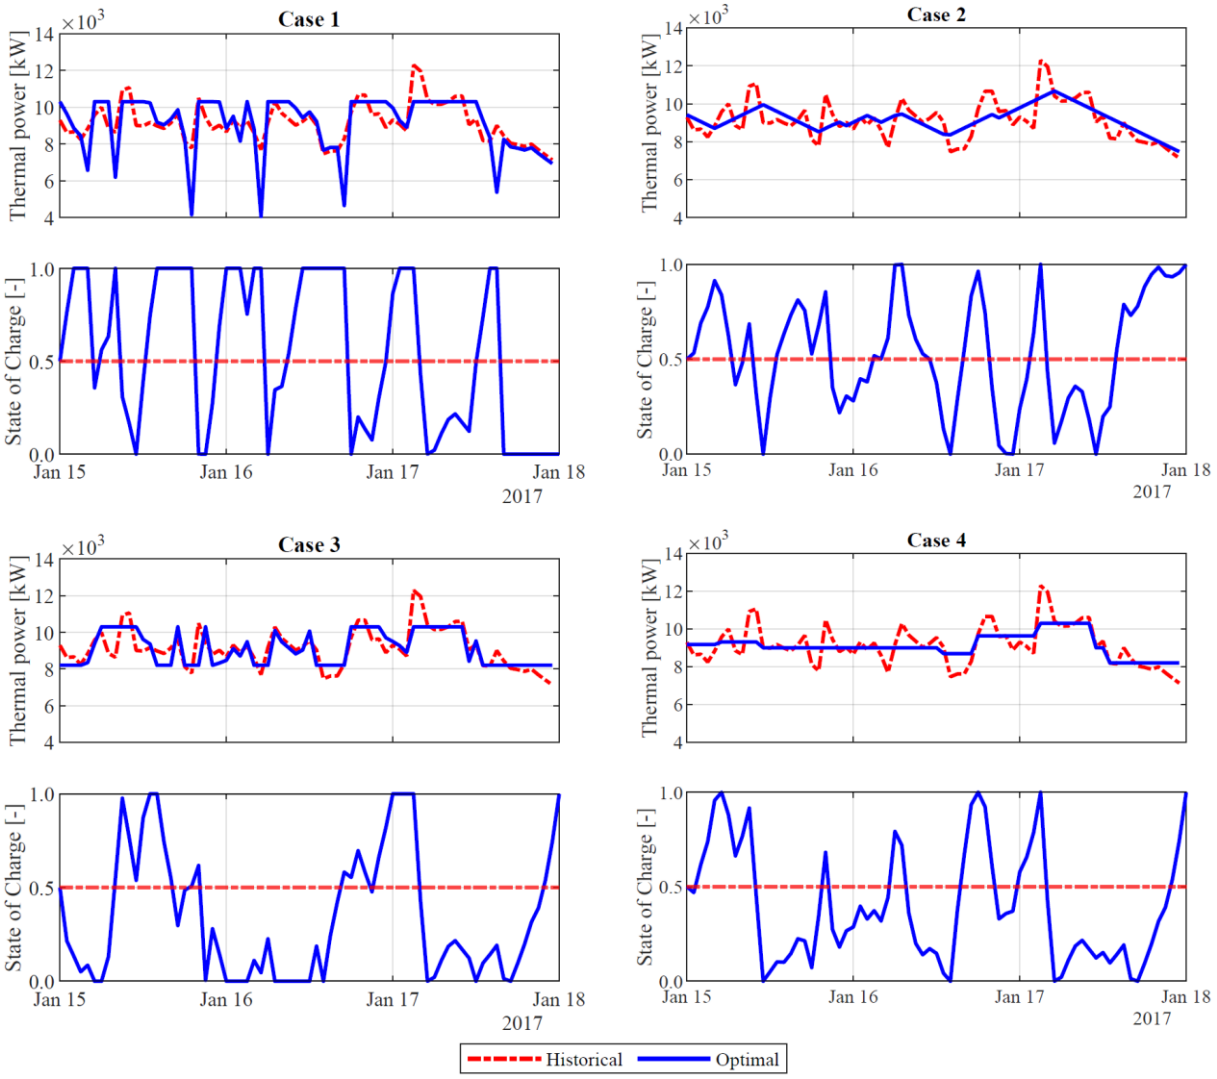

Figure 4. Historical and optimal thermal power and State of Charge as solutions of the Linear Programming problem. Cases 1 to 4 (the details are reported in Table 3).

In the former case, the input and, consequently, the system state are subjected to rapid fluctuations, such as the modulation of thermal power to 40% of the maximum value in a few hours. This is coherent with the aim of minimizing maximum input without including any information on its variation in the cost function. This behavior, however, is not preferable due to the potential difficulties in performing such regulation.

In Case 3, such fluctuations are reduced to 80% of the maximum input in one-hour steps, since the scope is to keep the range of variation of the thermal power as low as possible. In the latter case, instead, the evolution of the optimized thermal power is smoother than before, providing slower

changes in supply by regulating the SoC accordingly. Plus, it is possible to notice that Case 4 achieves both the same peak shaving and the same range of variation as Case 3, while showing more benefits. On the other hand, the sequence of inputs in Case 2 avoids fast changes of the parameters from one step to the next, but the peak shaving obtained is lower than the previous cases. Overall, the results shown in the graphs are coherent with the cost functions.

In the light of these reasons, the cost function of Case 4 is chosen in order to proceed with the implementation. Nonetheless, it is beneficial to show that the problem can be adapted with different objectives – even a combination of the previous goals – depending on the need.

The optimization, extended to the whole network, can prevent the occurrence of such demand peaks and, consequently, allow the system operator to avoid turning on back-up boilers at the production site, leading to a significant reduction in operating costs and energy consumption. Hence, the LP problem, which constitutes the first stage of the global optimization algorithm, is applied to all the six regions of the Västerås DHN. The sequence of optimal values of thermal power that should be supplied at the substation heat exchanger – hereinafter defined as $\dot{Q}_{LP}$ – and the optimal SoC of each region – defined as SoC_{LP} – are obtained and given as inputs to the second stage of the optimization algorithm (Section 3.3).

3.3 Global network optimization

The second stage of the optimization problem aims to optimize the control variables of the global DHN in order to reduce the supply temperature and move toward 4th generation district heating. This has to be achieved in compliance with the optimal heat supply to each region as calculated previously, i.e. $\dot{Q}_{LP}$ is a constraint of the new problem.

3.3.1 Problem definition

The states of the network, interpreted as the SoC of the regions, have been optimized in the first stage explained above (i.e. SoC_{LP}). The problem at this second stage is characterized as follows:

- The optimization variables are the mass flow rates sent to the six regions and the supply temperature at the power plant at each step of the prediction horizon, which comprises *N* time-steps as before. Therefore, the number of variables is (6 + 1)•N.
- As previously mentioned, the equality constraints are the actual supply of the optimal heat $\dot{Q}_{LP,i,k}$ for each time-step *k* and each region *i*. This is defined by Eq. (15):

$$\dot{m}_{i,k} c \left(T_{S,i,k} - T_{R,SP} \right) - \dot{Q}_{loss,i,k} = \dot{Q}_{LP,i,k} \tag{15}$$

where $\dot{m}_{i,k}$ is the mass flow rate to region *i* and time-step *k* and $T_{S,i,k}$ is the temperature that the water reaching region *i* at time-step *k* had when it left the power plant, some time before. The return temperature from the substation heat exchanger $T_{R,SP}$ is assumed equal to a set-point of 35 °C, as achieved in [27]. Similarly to [30], the heat losses from the network pipelines $\dot{Q}_{loss,i,k}$ are given by Eq. (16):

$$\dot{Q}_{loss,i,k} = \frac{\left(T_{S,i,k} - T_{soil} \right)}{R_{tot,i}} \tag{16}$$

where T_{soil} is the soil temperature, as the pipelines are typically underground, and $R_{tot,i}$ is the thermal resistance of the aggregated pipeline that conducts to the substation heat exchanger [27,30]. In this regard, the actual configuration and properties of the network (e.g. diameters, lengths and insulation) are collected in a database and considered in this calculation.

- The inequality constraints are that, for each pipe segment *j* of the network (Figure 2), the circulating mass flow rate $\dot{m}_{j,k}$ has to be positive but lower than the maximum mass flow rate allowed (Eq. (17a)). Moreover, the supply temperature has to be higher than the return temperature (Eq. (17b)).

$$0 \leq \dot{m}_{j,k} \leq \dot{m}_{j,max} \tag{17a}$$

$$T_{\mathrm{S,i,k}} > T_{\mathrm{R,SP}} \tag{17b}$$

- The cost function to be minimized is the maximum value of supply temperature over the prediction horizon. This leads to a reduction in heat losses to the soil while meeting the constraints. Under the same heat supply, as the temperature decreases, the mass flow rate increases and – as a consequence – so do the pumping costs, which depend on the third power of the mass flow rate itself. For this reason, a penalty factor proportional to the pumping power, calculated per region *i* at time-step *k* by Eq. (18) as in [30], is added to the objective function in order to discourage impracticable flow rates.

$$P_{\mathrm{pump,i,k}} = \sum\nolimits_{\mathrm{j}} \left[\frac{8 f_{\mathrm{j}} L_{\mathrm{j}}}{\pi^2 \rho^2 {D_{\mathrm{j}}}^5} \cdot \frac{1}{\eta_{\mathrm{pump}}} \cdot {\dot{m}_{\mathrm{j,k}}}^3 \right] \tag{18}$$

  where *j* is the index representing the pipeline segments that reach the *i*-th region, f_{j}, L_{j} and D_{j} are friction factor, length and diameter the *j*-th pipeline segment, respectively, ρ is the water density, and η_{pump} is the pump efficiency.

- The problem is nonlinear. Since it is a Non Linear Programming (NLP) case, it is solved with the Interior Point algorithm available in MATLAB®.

It is important to notice that the temperature that is actually supplied to each region at time *k* is not the supply temperature at the power plant at the same time *k*, since there are significant time delays (e.g. several hours) from production to actual delivery in large-scale networks. In the literature, a large number of papers considers this issue with different approaches. For instance, Laakkonen *et al.* [44] propose an optimization method for a DHN supply temperature based on brute force and calculates the heat load and return temperatures by means of a neural network. Additionally, the time delay from production to supply are based on the actual distance between the consumer and the plant. Otherwise, the temperature propagation can be treated by means of a phase and amplitude variation due to time delays and heat loss [45]. Similarly, in this work the simplified model of the network, which has to be incorporated into the global optimization problem embedded in an MPC controller,

includes a simplified representation of the phenomena to avoid an increase in computational complexity. While the authors of [27] adopt a fixed time delay selected as an approximation derived from historical data, the present paper implements a procedure to calculate the actual time delays to each region, which depend on the mass flow rates in the different positions of the network and at different time-steps. The basic assumption is that the temperature front moves at the same speed as the water inside the pipelines, calculated by Eq. (19):

$$w = \frac{4\dot{m}}{\rho\pi D^2} \tag{19}$$

At each calculation step, the position in the network of the supply temperature front sent at the previous time-steps is saved and continuously updated in a global matrix. Once this is known, it is possible to evaluate the actual temperature of the water that reaches the different regions and, therefore, to create the proper constraints for the optimization.

3.3.2 Analysis

The double-stage optimization algorithm that involves the entire network is analyzed by solving the problem for three representative days (i.e. January 2017). The LP problem and the NLP problem are solved in sequence as summarized in Section 2. The prediction horizon is three days and the time-step is one hour.

Firstly, it is of key importance to verify the applicability of the algorithm by checking that all the constraints are met. Figure 5 represents (i) the supplied thermal power according to the historical dataset and (ii) the thermal power that is actually delivered to the six regions with the mass flow rates and supply temperature calculated by the optimization algorithm. The maximum difference between the actual thermal power supplied with these parameters and that separately calculated by the LP and given as a constraint to the NLP is, in all cases, lower than 0.1%. Moreover, as expected, for all

network regions the peak of the actual power is significantly reduced compared to that before the optimization, if the given prediction horizon is considered.

The actual values of peak shaving for each region vary from 7.3% to 16.2% and are reported in Table 4. The table also reports the percentage reduction in the thermal power variation range (*RVR*), calculated by Eq. (20) as the difference between maximum and minimum values over the given horizon, compared to the historical case:

$$RVR = \left(1 - \frac{\dot{Q}_{\text{max,opt}} - \dot{Q}_{\text{min,opt}}}{\dot{Q}_{\text{max,hist}} - \dot{Q}_{\text{min,hist}}}\right) \cdot 100 \tag{20}$$

The optimization leads to a smoother and more regular operation (i.e. the range is reduced by 44% to 77% depending on the region), which is desirable for this kind of system. Hence, these outcomes show that the results of the optimization are coherent with the constraints and objectives.

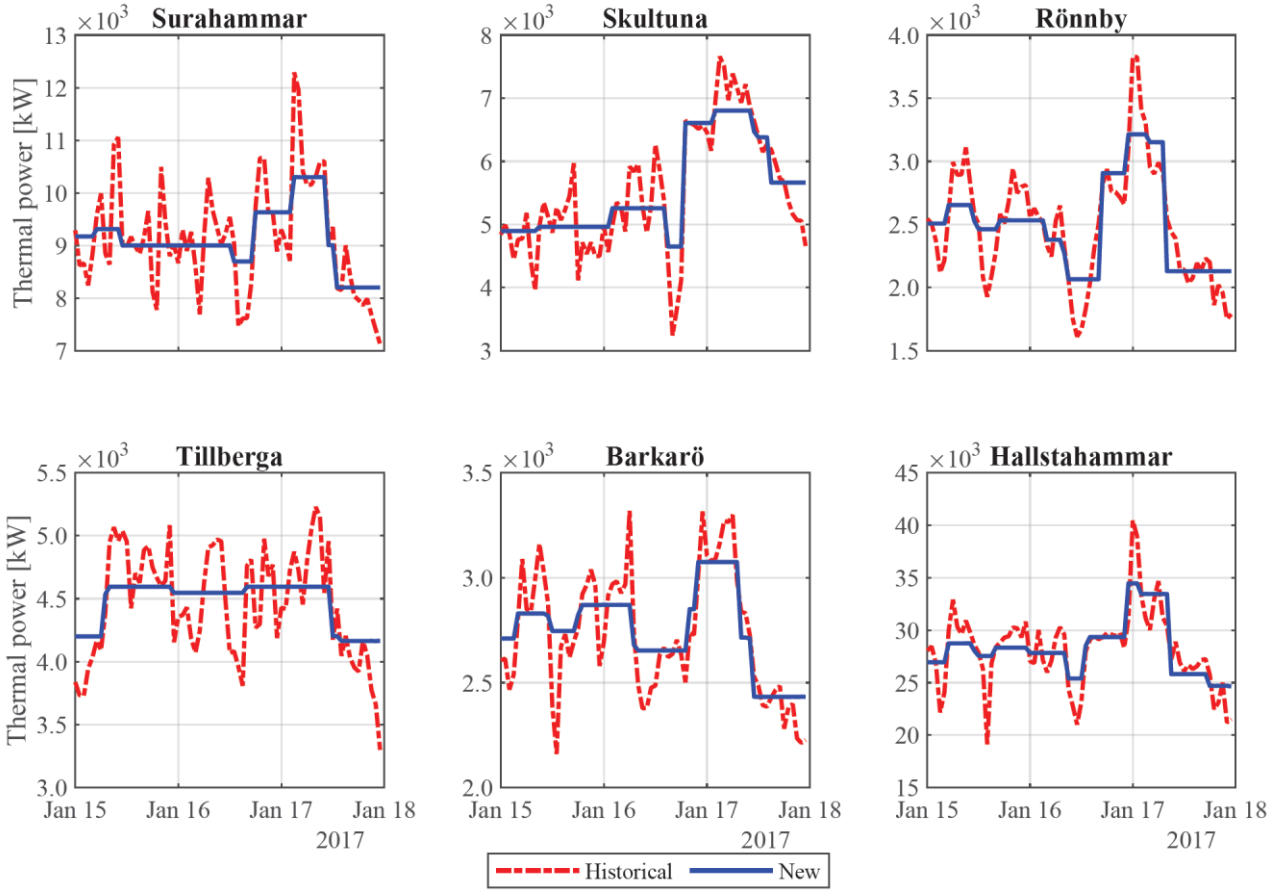

Figure 5. Thermal power supplied according to the historical dataset and to the optimization algorithm (new). Thermal peak shaving and reduction in the heat variation range are achieved in all regions.

Table 4. Results of peak shaving and range of variation reduction of the six external regions of the Västerås district heating network as obtained with the operating parameters calculated by the optimization algorithm.

Region	Peak shaving	Variation range reduction (RVR)
Surahammar	16.2%	59.3%
Skultuna	11.2%	51.3%
Rönnby	16.1%	48.3%
Tillberga	12.1%	77.6%
Barkarö	7.3%	44.5%
Hallstahammar	14.8%	53.9%

The range of variation of the mass flow rates in the nine main distribution pipelines of the network, as well as their boundaries given by historical data, are represented in Figure 6. It is confirmed that the operating constraints are respected, and the flows are not subjected to impracticable variations.

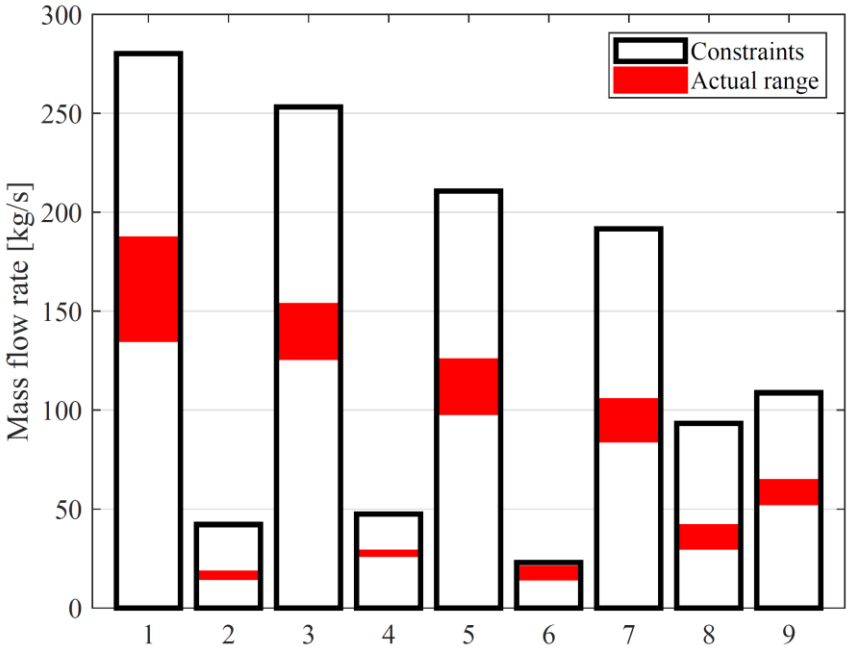

Figure 6. Actual range of variation of the mass flow rate compared to the constraints for nine pipeline segments of the district heating network. All values of mass flow rates respect the operating constraints.

The entire two-stage optimization algorithm with prediction horizons of one, two and three days has a computational time of 10 s, 50 s and 120 s, respectively. Hence, its implementation as part of an MPC controller that updates its variables and performs the calculation every hour is feasible. The application and testing of this integrated method in an MiL configuration is discussed in Section 4.

4. Global network predictive control

The practicability of the developed method is tested in an MiL application in MATLAB®/Simulink®. As represented in Figure 1, the algorithm is included in an MPC block which controls a detailed dynamic model of the Västerås DHN. This serves as a test bench for analyzing the results and, eventually, suggesting smarter strategies for control, operation and design. Firstly, the detailed model and the MPC control implementation are described. Secondly, the discussion of their results is reported.

4.1 Global network detailed model

The detailed dynamic model of the global network [27] was originally developed in Dymola [46], which exploits the programming language Modelica and is widely used in simulating dynamic systems [47]. The individual components of a DHN, such as pipes, valves and pumps, are modeled by considering the governing physical phenomena. For instance, the pipe model is represented by the energy balance dynamic equation applied to the pipe wall and to the water contained in each control volume in which the pipe is divided. The heat transferred to the consumer and the heat lost to the environment are included in the equation. These components are connected according to the network configuration to form the global model. Its validation with historical network data (i.e. temperature and mass flow rate supplied from the CHP, soil temperature and thermal load), provided by the system operator, is reported in detail by Zimmerman *et al.* [27]. This validation procedure shows that the model is reliable and effectively represents network temperature dynamics, heat propagation and losses over pipeline lengths of several kilometers, as is the case for the present network (Table 1). Indeed, the simulated supply and return temperatures are in good agreement with the actual trends.

4.2 Problem definition

The main peculiarity of MPC is the receding time horizon strategy, according to which, at each time-step, the actual measurement of the system state and the new prediction of the disturbances are considered for updating and solving a new optimization problem. In this way, the influence of the prediction uncertainty and model inaccuracies on the actual control can be reduced step-by-step.

In this work, the model outlined in Section 4.1 is controlled with the optimal inputs calculated by the optimization module (i.e. mass flow rates to the regions and supply temperature from the power plant). The model simulates the heat propagation of the real network and, at every calculation step, has to return the estimation of the actual SoC for each region. These are used as new initial conditions for the next optimization problem. However, the current system configuration does not allow the end-user indoor temperatures to be monitored. Hence, another fundamental assumption, explained below, is necessary to update the SoCs of the regions and correctly perform real-time control. The thermal power contributions also reported below are schematized in a qualitative way in Figure 7.

The optimization algorithm of the MPC receives the historical or predicted heat load as a disturbance – $\dot{Q}_{\mathrm{base,i}}$ for each region i – and calculates the network operating parameters that deliver the optimal thermal power – $\dot{Q}_{\mathrm{LP,i,k}}$ for region i and time-step k – which in turn includes the optimal storage by considering a set-point return temperature of 35 °C [27]. Nonetheless, there is typically a difference between the predicted load and that actually consumed, due to various factors:

- errors in the prediction of the external temperature and, consequently, the thermal demand;
- different behavior of part of the consumers (e.g. a building is not occupied and does not require heating).

Additionally, part of the consumers may decide not to actuate the storage, for instance by exploiting the space heaters differently or slightly varying the indoor comfort.

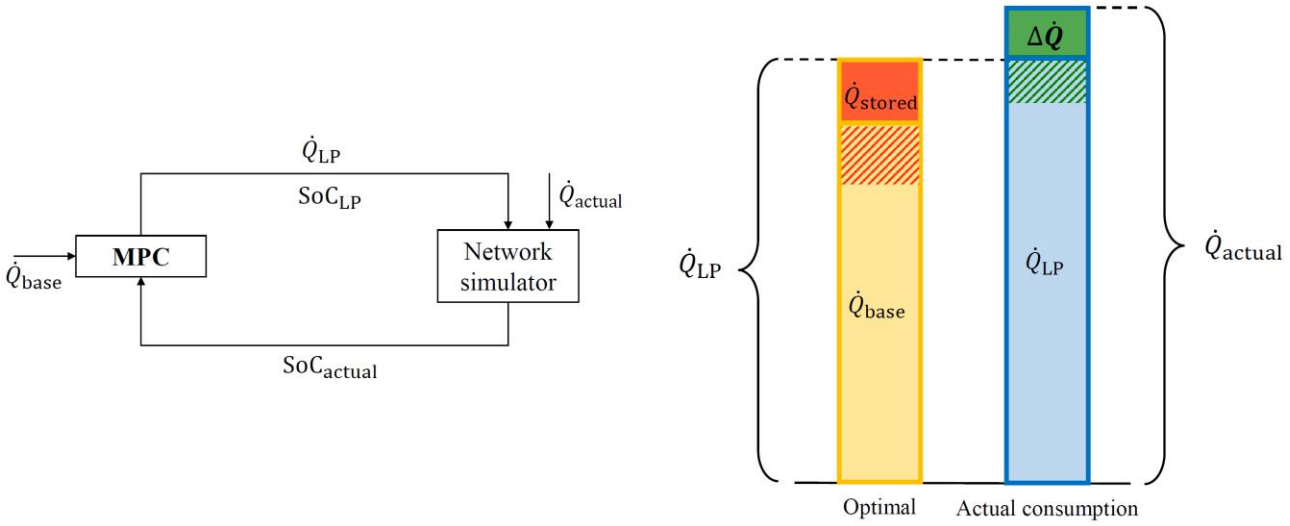

Figure 7. Qualitative representation of the contributions to the optimal and actually consumed thermal power. The quantities $\dot{Q}_{stored}$ and $\Delta\dot{Q}$ can be positive (full area) or negative (dashed area). The total heat stored in the SoC is represented by the sum of $\dot{Q}_{stored}$ and $\Delta\dot{Q}$ in the current time-step.

These random contributions are included in the current analysis by estimating the difference $\Delta\dot{Q}_{i,k}$ between the predicted thermal power and the thermal power actually consumed as a fraction of $\dot{Q}_{base}$ according to a uniform distribution in the range ± 5%. This implies a return temperature – $T_{R,i,k}$ from region i at time-step k – that is different from the abovementioned set-point and is returned by the detailed network model once the actual thermal power $\dot{Q}_{actual,i,k}$ has been consumed. Eq. (21) expresses the relation between the thermal power actually consumed and this new return temperature. It is worth noting that this balance equation is analogous to Eq. (15).

$$\dot{m}_{i,k}c\left(T_{S,i,k} - T_{R,i,k}\right) - \dot{Q}_{loss,i,k} = \dot{Q}_{actual,i,k} \tag{21}$$

The subtraction of Eq. (15) from Eq. (21) leads to an expression of $\Delta\dot{Q}_{i,k}$ in terms of the difference between the actual return temperature and the set-point return temperature as follows:

$$\Delta\dot{Q}_{i,k} = \dot{Q}_{actual,i,k} - \dot{Q}_{LP,i,k} = \dot{m}_{i,k}c\left(T_{R,SP} - T_{R,i,k}\right) \tag{22}$$

In the current real system configuration, it is not possible to quantify and separate the random contributions listed above. At the same time, it is necessary to update the SoC with a new estimation.

Therefore, it is assumed that the sum (with its corresponding sign) of the random contribution $\Delta\dot{Q}_{i,k}$ is entirely stored in the heat capacity of the aggregated consumer and modifies the SoC foreseen by the optimization module. According to Eq. (22), the difference in the return temperature gives a measurement of this additional heat which, in turn, effectively modifies the predicted SoCs of the network – $SoC_{LP,i,k}$ for region *i* and time-step *k* – as in Eq. (23):

$$SoC_{actual,i,k} = SoC_{LP,i,k} + \frac{\dot{m}_{i,k} c \left(T_{R,SP} - T_{R,i,k}\right) \Delta t}{2 C \Delta T_{stored,max}} \quad (23)$$

where $SoC_{actual,i,k}$ represents the new measurement of the state of region *i*. It is also the new initial condition for the optimization problem at time-step *k* + 1, according to the receding time horizon strategy. This is applied to all the regions and the simulation proceeds for the desired period.

4.3 Results

The MiL application comprising the detailed network model and the MPC controller is illustrated in Figure 8. The results of the real-time control with the developed MPC strategy are shown for the first week of February 2017. The outdoor temperature, which greatly affects the baseline heat demand, is depicted in Figure 9.

A prediction horizon of one day is adopted. The time delays to reach the different regions of the network are in the range of 2 h to 10 h, with the considered operating parameters. Hence, such a prediction horizon includes the propagation of the thermal power produced in the power plant and captures the main system dynamics, while reducing the simulation time. Furthermore, since the model incorporates the distribution line to the furthest reaching areas in the supply network, the dynamics of the central area of Västerås remains within acceptable limits, and therefore indoor comfort is achievable as well.

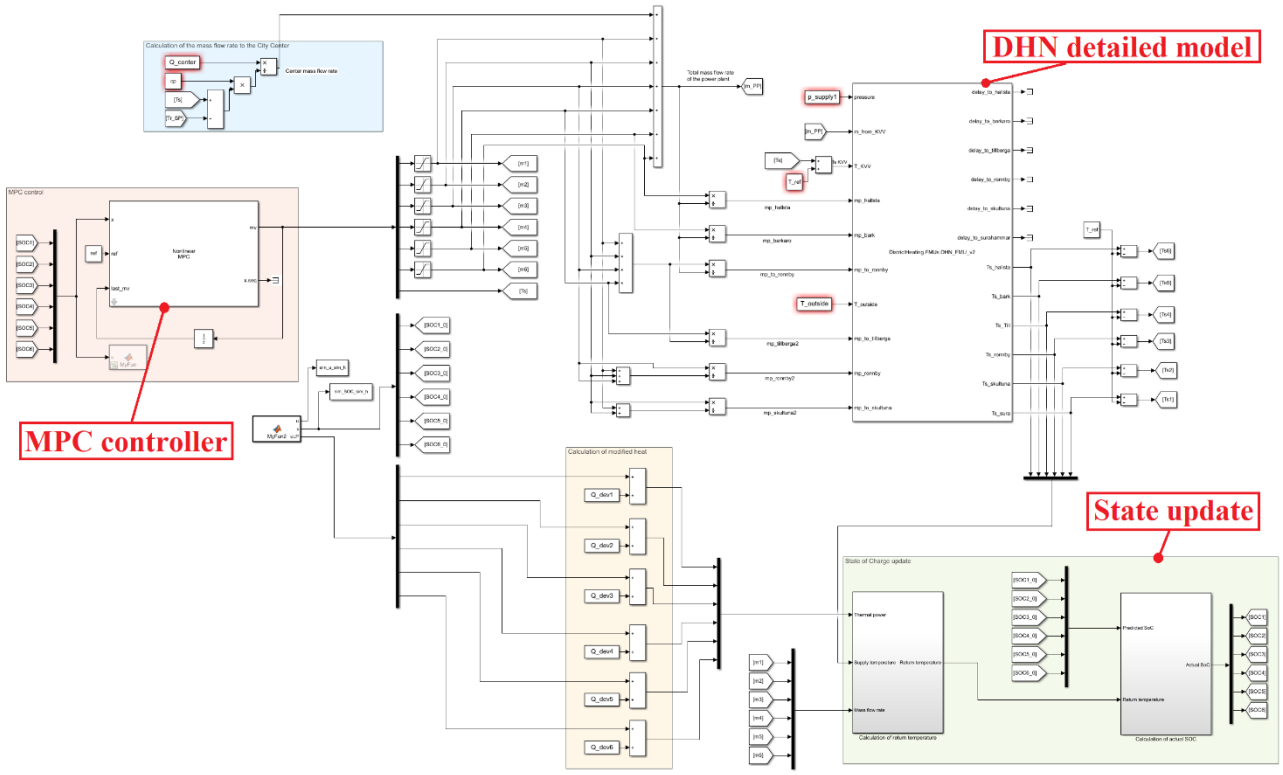

Figure 8. Model-in-the-Loop application of the Västerås district heating network controlled with the proposed control strategy in the MATLAB®/Simulink® environment. The MPC controller, detailed network model and state update sections are highlighted.

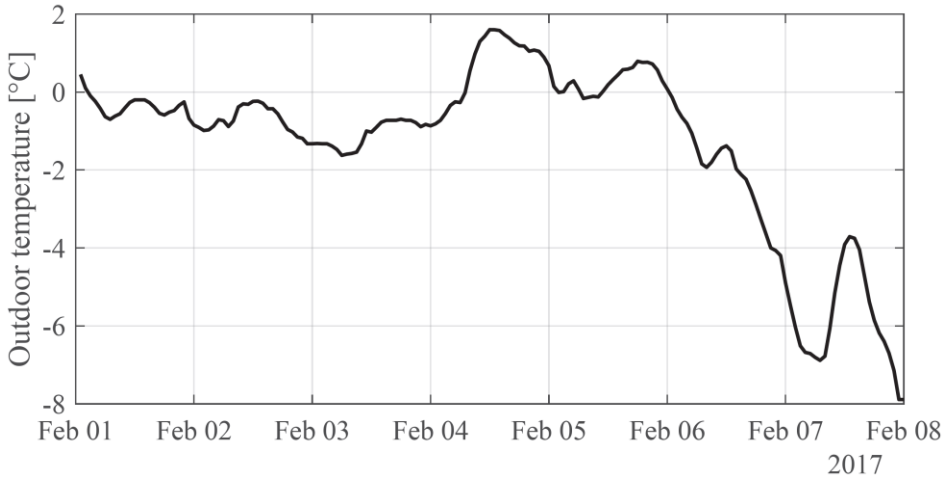

Figure 9. Historical outdoor temperature over the simulation period.

Figure 10 and Figure 11 depict the operating parameters of the region of Skultuna, taken as a representative example, in terms of mass flow rates and temperatures supplied directly to the main substation heat exchanger. The comparison between the historical and new values shows that the MPC controlled variables decrease: this leads to lower consumption for the pump and lower heat

losses from the pipelines. For instance, the mass flow rate to Skultuna is decreased by 23% but similar results are obtained in the other cases. Moreover, the optimal mass flow rate is subjected to smaller variations over the simulation period, unlike the previous management conditions in which the operation of the pump changes and fluctuates more significantly. The supply temperature is also more constant. The same considerations can be drawn for the other regions of the DHN.

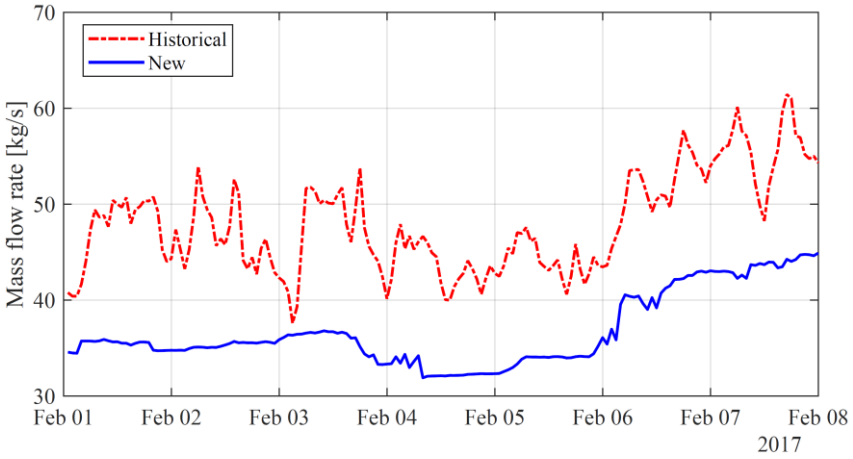

Figure 10. Historical and new water mass flow rates to Skultuna. The MPC reduces the mass flow rate.

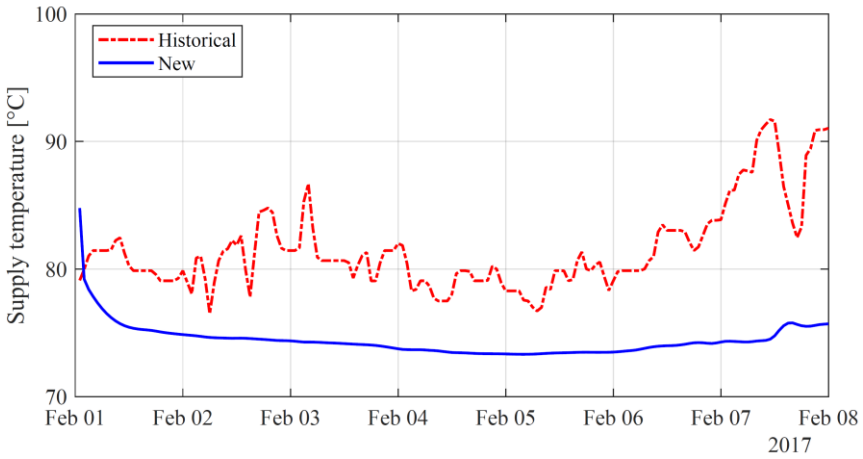

Figure 11. Historical and new supply temperatures to Skultuna. The MPC reduces the distribution temperature and, thus, the heat losses.

The first part of the simulation shows a higher temperature that derives from the fact that the optimal supply temperature reaches the substation heat exchanger some hours after it outflows from the power plant, due to the network time delays. This diminishing evolution in the first simulation time-steps is different for the various regions, depending on the distance and water velocity. Indeed, it is underlined

that, even though they are supplied by the same power plant, the actual temperatures that reach the substation heat exchanger are different due to the different rates of heat transfer from the pipelines to the soil, which depend on the network topology, length and demand.

In regard to the effects that the outdoor temperature imposes upon the controller, Figure 12 illustrates the power plant supply temperature, which is valid for all regions, as a function of external environment temperature. The trend obtained with the MPC, highlighted by the regression fit lines added to the graph, is significantly decreased compared to baseline. Furthermore, the new supply temperature is maintained more uniform, especially when the outdoor temperature is relatively low. Indeed, the absolute value of the slope decreases from 1.42 to 0.12. As far as this simulation period is concerned, the water is supplied at an optimal temperature lower than 80 °C, while in the original management strategies it varies between 80 °C and 95 °C. The relative reduction in the difference between supply temperature and soil temperature (which is assumed as 10 °C in the simulations), which gives the perception of a reduction in heat losses from the pipelines, ranges from a minimum of 3.1% for higher outdoor temperatures to a maximum of 20% for lower outdoor temperatures, with an average reduction of 8.8% over the considered time period.

Other simulation periods can lead to further improvements, since the historical supply temperature – according to the network data – can reach values up to 110 °C. Hence, the secondary goal of the optimization, which consists of reducing the heat losses, is successfully implemented. This makes the roadmap toward low-temperature district heating clearer and more straightforward.

Figure 13 represents the real return temperature from the substation heat exchanger of the region of Skultuna, compared to the baseline case and to the set-point considered for the optimization. The return temperature fluctuates around the set-point, proving the acceptability of the assumption made. In addition, it is lower than the historical data, giving its contribution to the general trend of decreasing the pipeline temperature. It is also possible to notice that the return temperature is always maintained

above the threshold of 30 °C, which is preferable for the global system operation. Similarly, these results are valid for the other regions of the DHN.

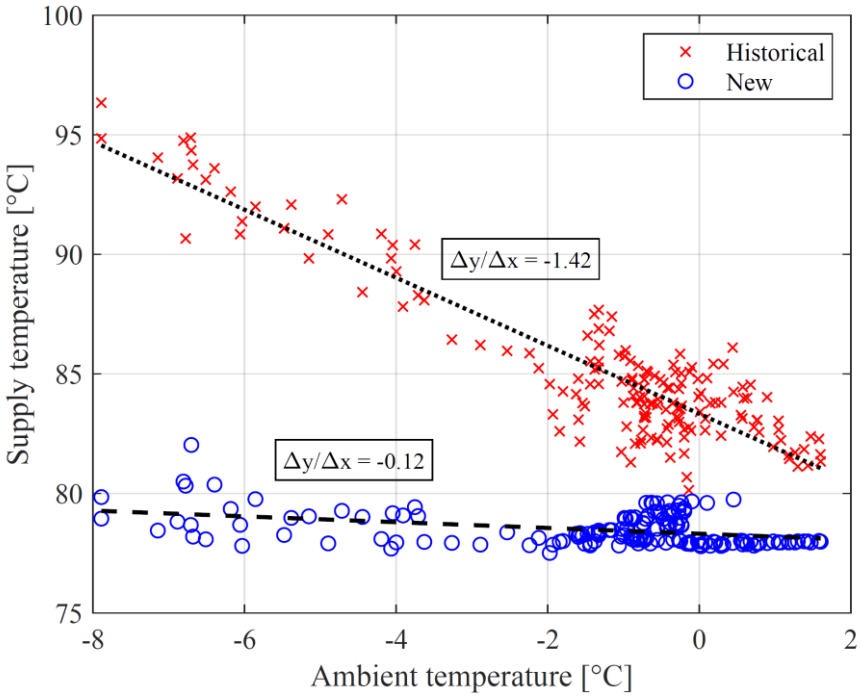

Figure 12. Historical and new data of the supply temperature from the power plant with the related regression fit lines. The minimum, average and maximum reductions in the difference between supply and soil temperature are 3.1%, 8.8% and 20%, respectively.

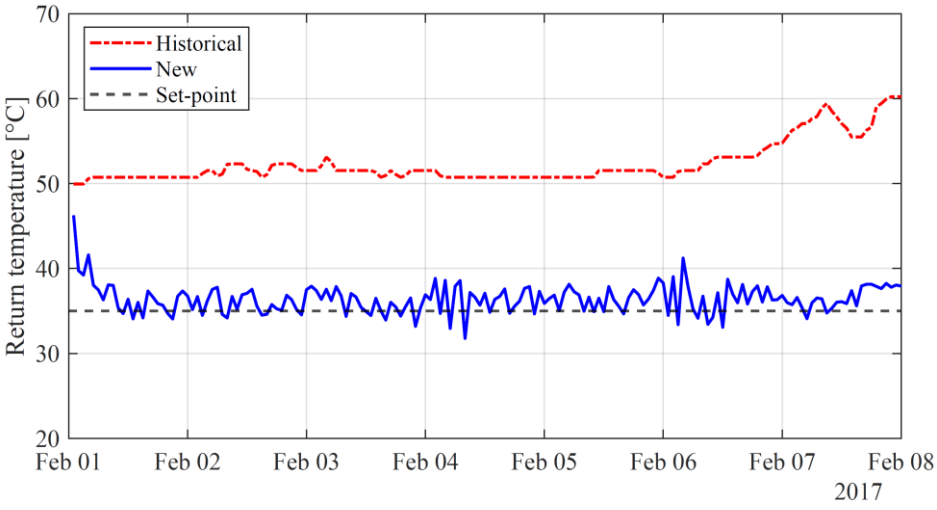

Figure 13. Historical and new return temperatures from Skultuna, compared to the set-point considered for the optimization.

Overall, the optimal control parameters determined by the MPC are able to achieve the primary objective of the analysis, which is to shave the peaks of the thermal power behavior by exploiting the building heat capacity and without endangering the indoor comfort. This is demonstrated by Figure 14, which illustrates the historical and new thermal power actually delivered. The heat supplied is subjected to greater fluctuations compared to the analysis reported in Section 3.3.2 because the real-time MPC control considers the deviation of the heat consumed from that predicted, which adds uncertainty to the overall performance of the proposed approach. This happens also due to the receding time horizon strategy, according to which the optimization is repeated at each time-step with a continuous update of the variables and disturbance prediction.

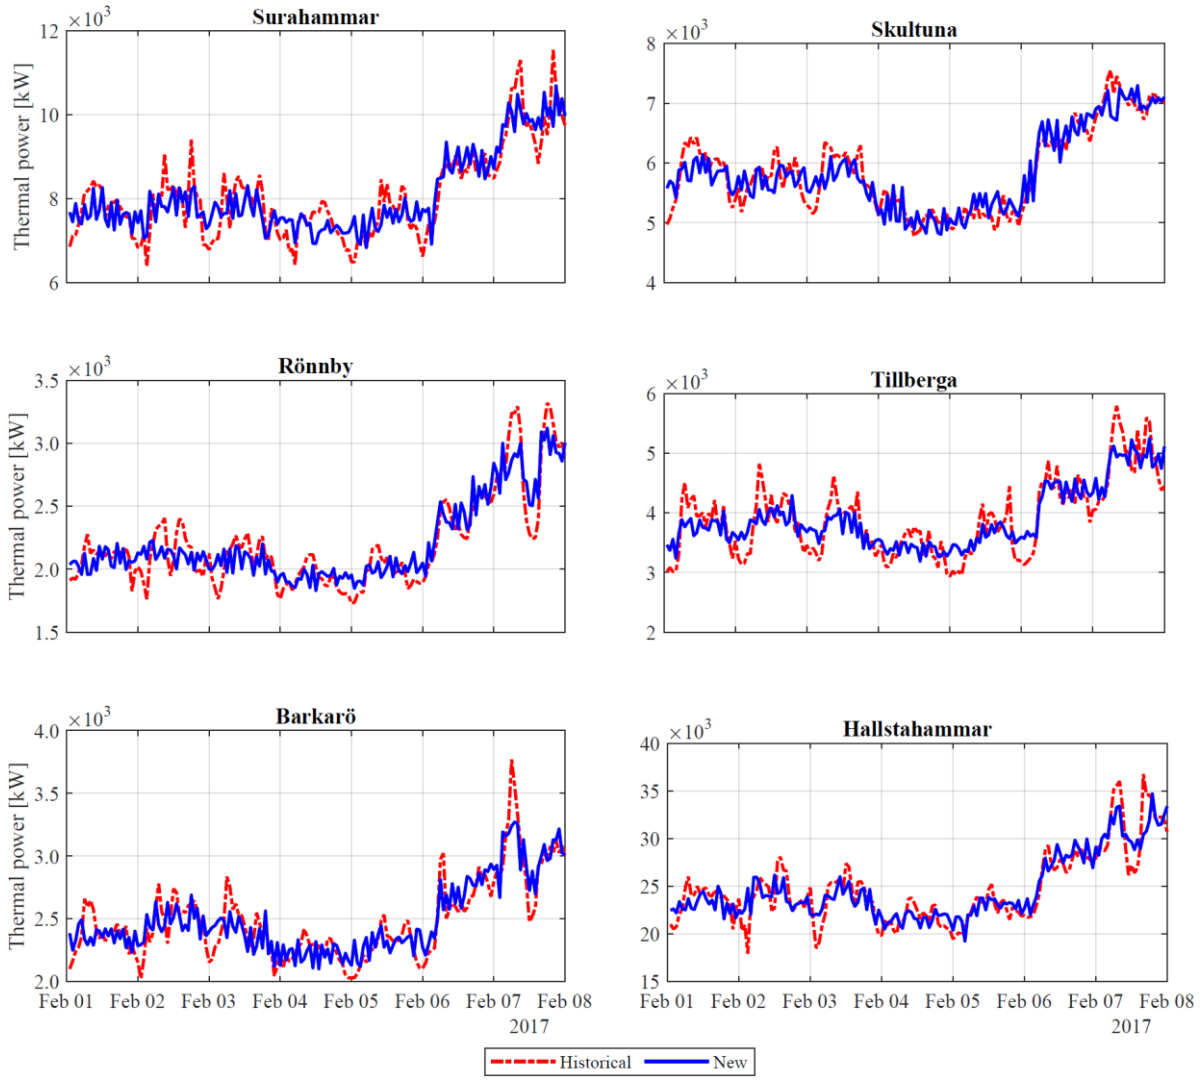

Figure 14. Historical and new thermal power to the regions. The MPC achieves peak shaving in all regions.

Nevertheless, these results are achieved by adopting a conservative estimation of the heat capacity of the regions [23] and a safe maximum deviation of their equivalent indoor temperature, i.e. 0.5 °C. A more significant peak shaving could be obtained by increasing this parameter or by testing more intense demand side management measures. Furthermore, this integrated approach and the presented set of tools (aggregated model of the region, optimization algorithm and network simulator) can lead to promising results also in terms of the exploitation of distributed thermal storage devices. For instance, the region of Skultuna achieves the lowest level of peak shaving and might benefit from a bigger storage potential. The placement of a storage tank in this area, as well as in other key portions of the system, might positively affect the overall network management.

In summary, these considerations can be useful to DHN companies and system operators regarding efficiency improvement and energy saving. Indeed, the presented results can be easily exported to real applications by implementing the proposed MPC as network supervisory controller, without the need to modify the network configuration or carry out extensive investments.

Nevertheless, this work has some limitations that can be addressed in future research:

- The analysis is performed at main distribution network level, which means that only the main branches of the network supplying the peripheral regions are considered. The allocation of this thermal power to each individual building comprised in the region has to be regulated by another level of control.
- The SoC of each region is updated by assuming that the uncertainty about the thermal demand affects the SoC itself and not the indoor comfort. It is possible to install measurement devices in a small portion of the network and conduct experimental tests in order to better validate these results, which can be then extended to the entire system.

5. Conclusions

The heating and cooling sector has the chance to become more efficient and to reduce carbon emissions in populated areas if innovative management and control approaches are adopted. A promising strategy for district heating networks consists of exploiting the thermal capacity of the buildings to optimize energy supply and reduce peak demand.

In this work, a Model Predictive Controller exploiting the receding time horizon approach for large-scale district heating systems was developed and tested on a digital twin of the network of Västerås, Sweden. This controller is fast and scalable, and it is able to include any section of a network for which basic consumption data is available, since it does not require a long model tuning procedure. Moreover, it updates the global network optimization at every given time-step (i.e. one hour) in order to reject prediction errors and improve control efficiency. It performed (i) the minimization of the peaks of the heat supplied to each region by considering it as a thermal battery with a State of Charge, and (ii) the minimization of the network temperature. The analysis of the approach and the results showed that the method is applicable, and the goals have been achieved. The peak shaving is up to 16% and the mass flow rates are reduced by up to 23% with a consequent significant reduction in pumping costs. Furthermore, the supply temperature is decreased to values lower than 80 °C with the difference between distribution and soil temperature being reduced by up to 20%, in turn leading to lower losses from the pipelines.

The outcomes of this work can have multiple benefits for district heating operators and stakeholders in different countries:

- to assess the storage potential in systems that lack extensive data and information on the building characteristics;
- to propose new planning and management strategies of the thermal power distribution in order to reduce the thermal peaks in the plant operation;

- to suggest the installation of new storage devices (e.g. water storage tanks), if considered beneficial in increasing the heat capacity and, consequently, the storage potential of a defined area and in further improving the operation and management;
- to understand the state of the network and to control the heat supply to the furthest neighborhoods.

The implementation of these aspects will be investigated in future studies. It will also be possible to exploit the potentialities of the approach to a greater extent by means of an improved monitoring of the network State of Charge. Furthermore, the proposed procedure will be scaled and implemented fractally, in order to demonstrate a single methodology potentially applicable to any kind of heating network, from the micro-scale (i.e. building level) to the macro-scale (i.e. city level).

This integrated methodology can encourage (i) the transition of existing district heating toward low temperature, optimally managed and automatized networks, and (ii) the spread of smart (e.g. predictive) control architectures, which enable the digitalization of the heating sector. This could lead to a globally connected system in which the energy is produced, stored and distributed in an optimal way.

Acknowledgements

This work was supported by the "DISTRHEAT – Digital Intelligent and Scalable conTrol for Renewables in HEAting neTworks" project, which received funding in the framework of the joint programming initiative ERA-Net Smart Energy Systems' focus initiative Integrated, Regional Energy Systems, with support from the European Union's Horizon 2020 research and innovation programme under grant agreement No 775970. It was also supported by the "Bando Leonardo da Vinci 2019 – Azione 2" funded by the Italian Ministry of Education, University and Research (MIUR).

The authors would also like to thank their industrial partner Mälarenergi AB for providing the data of the district heating network of Västerås.

References

[1] Connolly D, Lund H, Mathiesen BV, Werner S, Möller B, Persson U et al. Heat Roadmap Europe: Combining district heating with heat savings to decarbonise the EU energy system. Energy Policy 2014;65:475–489. https://doi.org/10.1016/j.enpol.2013.10.035

[2] Hast A, Syri S, Lekavičius V, Galinis A. District heating in cities as a part of low-carbon energy system. Energy 2018;152:627–639. https://doi.org/10.1016/j.energy.2018.03.156

[3] Frederiksen S, Werner S. District Heating and Cooling, Studentlitteratur AB, Lund, 2013.

[4] Paardekooper S, Lund RS, Mathiesen BV, Chang M, Petersen UR, Grundahl L et al. Heat Roadmap Europe 4: Quantifying the Impact of Low-Carbon Heating and Cooling Roadmaps. Aalborg Universitetsforlag 2018. Available at: https://vbn.aau.dk/en/publications/heat-roadmap-europe-4-quantifying-the-impact-of-low-carbon-heatin [accessed 20/08/2020]

[5] Chambers J, Narula K, Sulzer M, Patel MK. Mapping district heating potential under evolving thermal demand scenarios and technologies: A case study for Switzerland. Energy 2019;176:682–692. https://doi.org/10.1016/j.energy.2019.04.044

[6] Arabkoohsar A, Alsagri AS. A new generation of district heating system with neighborhood-scale heat pumps and advanced pipes, a solution for future renewable-based energy systems. Energy 2020;193:116781. https://doi.org/10.1016/j.energy.2019.116781

[7] Mazhar AR, Liu S, Shukla A. A state of the art review on the district heating systems. Renewable and Sustainable Energy Reviews 2018;96:420439. https://doi.org/10.1016/j.rser.2018.08.005

[8] Vandermeulen A, Van der Heijde B, Helsen L. Controlling district heating and cooling networks to unlock flexibility: A review. Energy 2018;151:103–115. https://doi.org/10.1016/j.energy.2018.03.034

[9] Hennessy J, Li H, Wallin F, Thorin E. Flexibility in thermal grids: a review of short-term storage in district heating distribution networks. Energy Procedia 2019;158:2430–2434. https://doi.org/10.1016/j.egypro.2019.01.302

[10] Harney P, Gartland D, Murphy F. Determining the optimum low-temperature district heating network design for a secondary network supplying a low-energy-use apartment block in Ireland. Energy 2020;192,116595. https://doi.org/10.1016/j.energy.2019.116595

[11] Chicherin S, Junussova L, Junussov T. Minimizing the supply temperature at the district heating plant – dynamic optimization. E3S Web of Conferences 2019;118:02004. https://doi.org/10.1051/e3sconf/201911802004

[12] Lund H, Werner S, Wiltshire R, Svendsen S, Thorsen JE, Hvelplund F et al. 4th Generation District Heating (4GDH) Integrating smart thermal grids into future sustainable energy systems. Energy 2014;68:1–11. https://doi.org/10.1016/j.energy.2014.02.089

[13] Lund H, Østergaard PA, Chang M, Werner S, Svendsen S, Sorknæs P et al. The status of 4th generation district heating: Research and results. Energy 2018;164:147–159. https://doi.org/10.1016/j.energy.2018.08.206

[14] Guelpa E, Verda V. Thermal energy storage in district heating and cooling systems: A review. Applied Energy 2019;252:113474. https://doi.org/10.1016/j.apenergy.2019.113474

[15] Guelpa E, Marincioni L, Deputato S, Capone M, Amelio S, Pochettino E et al. Demand side management in district heating networks: A real application. Energy 2019;182:433–442. https://doi.org/10.1016/j.energy.2019.05.131

[16] Kouhia M, Laukkanen T, Holmberg H, Ahtila P. District heat network as a short-term energy storage. Energy 2019;177:292–303. https://doi.org/10.1016/j.energy.2019.04.082

[17] Li X, Li W, Zhang R, Jiang T, Chen H, Li G. Collaborative scheduling and flexibility assessment of integrated electricity and district heating systems utilizing thermal inertia of district

heating network and aggregated buildings. Applied Energy 2020;258:114021. https://doi.org/10.1016/j.apenergy.2019.114021

[18] Leśko M, Bujalski W, Futyma. Operational optimization in district heating systems with the use of thermal energy storage. Energy 2018;165:902–915. https://doi.org/10.1016/j.energy.2018.09.141

[19] Le Dréau J, Heiselberg P. Energy flexibility of residential buildings using short term heat storage in the thermal mass. Energy 2016;111:991–1002. https://doi.org/10.1016/j.energy.2016.05.076

[20] Arnaudo M, Topel M, Puerto P, Widl E, Laumert B. Heat demand peak shaving in urban integrated energy systems by demand side management - A techno-economic and environmental approach. Energy 2019;186:115887. https://doi.org/10.1016/j.energy.2019.115887

[21] Claessens BJ, Vanhoudt D, Desmedt J, Ruelens F. Model-free control of thermostatically controlled loads connected to a district heating network. Energy and Buildings 2018;159:1–10. https://doi.org/10.1016/j.enbuild.2017.08.052

[22] Fazlollahi S, Schüler N, Maréchal F. A solid thermal storage model for optimization of buildings operation strategy. Energy 2015;88:209–222. https://doi.org/10.1016/j.energy.2015.04.085

[23] Saletti C, Zimmerman N, Morini M, Kyprianidis K, Gambarotta A. A Scale-Free Dynamic Model for District Heating Aggregated Regions. Preprints 2020;2020060320. https://doi.org/10.20944/preprints202006.0320.v1

[24] Korpela T, Kaivosoja J, Majanne Y, Laakkonen L, Murmoranta M, Vilkko M. Utilization of District Heating Networks to Provide Flexibility in CHP Production. Energy Procedia 2017;116:310–319. https://doi.org/10.1016/j.egypro.2017.05.077

[25] Birk W, Atta KT, Uden F. Improving district heating system operation through control configuration selection and adaptive control. 2019 18th European Control Conference, ECC 2019, article number 8795742, pages 2944–2949. https://doi.org/10.23919/ECC.2019.8795742

[26] Lv C, Yu H, Li P, Wang C, Xu X, Li S et al. Model predictive control based robust scheduling of community integrated energy system with operational flexibility. Applied Energy 2019;243:250–265. https://doi.org/10.1016/j.apenergy.2019.03.205

[27] Zimmerman N, Kyprianidis K, Lindberg CF. Achieving Lower District Heating Network Temperatures Using Feed-Forward MPC. Materials 2019;12(15), 2465. https://doi.org/10.3390/ma12152465

[28] Johansson C, Vanhoudt D, Brage J, Geysen D. Real-time grid optimisation through digitalisation – results of the STORM project. Energy Procedia 2018;149:246–255. https://doi.org/10.1016/j.egypro.2018.08.189

[29] Gambarotta A, Morini M, Saletti C. Development of a Model-based Predictive Controller for a heat distribution network. Energy Procedia 2019;158: 2896–2901. https://doi.org/10.1016/j.egypro.2019.01.944

[30] Saletti C, Gambarotta A, Morini M. Development, analysis and application of a predictive controller to a small-scale district heating system. Applied Thermal Engineering 2020;114558. https://doi.org/10.1016/j.applthermaleng.2019.114558

[31] Cadau N, De Lorenzi A, Gambarotta A, Morini M, Saletti C. A Model-in-the-Loop application of a Predictive Controller to a District Heating system. Energy Procedia 2018;148:352–359. https://doi.org/10.1016/j.egypro.2018.08.088

[32] De Lorenzi A, Gambarotta A, Morini M, Rossi M, Saletti C. Setup and testing of smart controllers for small-scale district heating networks: an integrated framework. Energy 2020;205:118054. https://doi.org/10.1016/j.energy.2020.118054

[33] Bavière R, Vallée M. Optimal Temperature Control of Large Scale District Heating Networks. Energy Procedia 2018;149:69–78. https://doi.org/10.1016/j.egypro.2018.08.170

[34] Olsthoorn D, Haghighat F, Mirzaei PA. Integration of storage and renewable energy into district heating systems: A review of modelling and optimization. Solar Energy 2016;136:49–64. https://doi.org/10.1016/j.solener.2016.06.054

[35] Çengel Y. Heat transfer: a practical approach. McGraw-Hill; 2003. ISBN 9780072458930.

[36] Ellis M, Durand H, Christofides PD. A tutorial review of economic model predictive control methods. Journal of Process Control 2014;24:1156–1178. https://doi.org/10.1016/j.jprocont.2014.03.010

[37] Mayne DQ. Model predictive control: Recent developments and future promise. Automatica 2014;50:2967–2986. https://doi.org/10.1016/j.automatica.2014.10.128

[38] Dotzauer E. Simple model for prediction of loads in district-heating systems. Applied Energy 2002;73:277–284. https://doi.org/10.1016/S0306-2619(02)00078-8

[39] Koschwitz D, Frisch J, Van Treeck C. Data-driven heating and cooling load predictions for non-residential buildings based on support vector machine regression and NARX Recurrent Neural Network: A comparative study on district scale. Energy 2018;165:134–142. https://doi.org/10.1016/j.energy.2018.09.068

[40] Suryanarayana G, Lago J, Geysen D, Aleksiejuk P, Johansson C. Thermal load forecasting in district heating networks using deep learning and advanced feature selection methods. Energy 2018;157:141–149. https://doi.org/10.1016/j.energy.2018.05.111

[41] Benzaama MH, Rajaoarisoa LH, Ajib B, Lecoeuche S. A data-driven methodology to predict thermal behavior of residential buildings using piecewise linear models. Journal of Building Engineering 2020;101523. https://doi.org/10.1016/j.jobe.2020.101523

[42] Romanchenko D, Kensby J, Odenberger M, Johnsson F. Thermal energy storage in district heating: Centralised storage vs. storage in thermal inertia of buildings. Energy Conversion and Management 2018;162:26–38. https://doi.org/10.1016/j.enconman.2018.01.068

[43] Propoi A, Willekens F. A Dynamic Linear-Programming Approach to the Planning of National Settlement Systems. Environment and Planning A: Economy and Space 1978;10:561–576. https://doi.org/10.1068/a100561

[44] Laakkonen L, Korpela T, Kaivosoja J, Vilkko M, Majanne Y, Nurmoranta M. Predictive Supply Temperature Optimization of District Heating Networks Using Delay Distributions. Energy Procedia 2017;116:297–309. https://doi.org/10.1016/j.egypro.2017.05.076

[45] Yuan X, Yali X, Qiongyao W. Dynamic temperature model of district heating system based operation data. Energy Procedia 2019;158:6570–6575. https://doi.org/10.1016/j.egypro.2019.01.073

[46] DYMOLA Systems Engineering website. https://www.3ds.com/products-services/catia/products/dymola/ [accessed 20/08/2020]

[47] Hermansson K, Kos C, Starfelt F, Kyprianidis K, Lindberg CF, Zimmerman N. An Automated Approach to Building and Simulating Dynamic District Heating Networks. IFAC-PapersOnLine 2018;51:855–860. https://doi.org/10.1016/j.ifacol.2018.04.021

Nomenclature

C	aggregated heat capacity coefficient [kJ °C^{-1}]
D	pipe diameter [m]
d	system disturbance [kW]
f	pipe friction factor [-]
L	pipe length [m]
$\dot{m}$	mass flow rate [kg s^{-1}]
N	number of time-steps [-]
P	power [kW]
$\dot{Q}$	thermal power [kW]
Q	heat [kJ]
R_{tot}	thermal resistance [°C kW^{-1}]
RVR	reduction in thermal power variation range [%]
SoC	State of Charge [-]
T	temperature [°C]
t	time [s]
U	aggregated heat transfer coefficient [kW °C^{-1}]
u	system input [kW]
w	water speed [m s^{-1}]
x	system state [-]
$\Delta\dot{Q}$	thermal power difference [kW]

ΔT	difference of temperature [°C]
α	first coefficient for the state-space dynamic problem [kW^{-1}]
β	second coefficient for the state-space dynamic problem [-]
ρ	water density [$kg\ m^{-3}$]
η	efficiency [-]

Subscripts

actual	actual
base	baseline (to keep the comfort conditions)
ext	outdoor
hist	historical
i	region i
j	pipeline segment j
k	time-step k
loss	heat loss
LP	referred to the solution of the Linear Programming problem
max	maximum
min	minimum
opt	optimal
pump	pump
R	return

S	supply
soil	soil
stored	stored in the thermal mass

Acronyms

CHP	Combined Heat and Power
DHN	District Heating Network
LP	Linear Programming
MiL	Model-in-the-Loop
MPC	Model Predictive Control
NLP	Non Linear Programming
SoC	State of Charge