Energy 301 (2024) 131690

Contents lists available at ScienceDirect

Energy

journal homepage: www.elsevier.com/locate/energy

Peak shaving at system level with a large district heating substation using deep learning forecasting models

Ulrich Trabert *, Felix Pag, Janybek Orozaliev, Ulrike Jordan, Klaus Vajen

Institute of Thermal Engineering, University of Kassel, Kurt-Wolters-Str. 3, Kassel, 34125, Hesse, Germany

ARTICLE INFO

Keywords:
Forecasting
Machine learning
Incremental learning
Large district heating substation
Thermal storage
Peak shaving

ABSTRACT

The decarbonisation of urban district heating (DH) systems requires increased heating grid flexibility. Therefore, this article examines the optimised operation of a tank thermal energy storage (TTES) on the secondary side of a new DH substation for an industrial site in a German city, in order to shave the peaks of the whole DH system and thus reduce the need for heat-only boilers (HOB). The accuracy of heat load and return temperature forecasts for both the industrial consumer and the DH grid is critical to the performance of the optimisation-based operating strategy of the TTES. Therefore, long short-term memory neural networks are used in combination with continuous model updates through incremental learning to create two forecasting scenarios, one using only preceding data for the forecasts and the other including future weather data. The results show that high forecasting accuracy is most relevant for reducing the annual maximum peak, with a reduction of 2.8% in the preceding data scenario, 4% with future weather data and 7% in a benchmark with perfect forecasts. The economic viability of the storage through HOB heat savings is primarily affected by lower forecasting accuracy when the additional cost of HOB heat is less than 60 €/MWh.

1. Introduction

The ongoing transformation of the energy system towards a renewable energy supply is accompanied by an increasing integration of the heat and electricity sectors. The efficient use of limited and intermittent renewable resources requires greater flexibility on both the production and consumption sides of the energy supply. In this context, the district heating (DH) sector offers the opportunity for relatively cheap storage options to support the concept of smart energy systems [1,2].

Tank thermal energy storage (TTES) is a well-known technology in DH systems today, but its potential for flexible use to support sector-coupled energy systems has not yet been fully exploited [3]. However, TTES play an important role in optimising the operation of DH systems by providing flexibility to balance supply and demand, operate in line with energy prices, or respond to weather conditions [4].

While DH systems are still predominantly centralised, the shift towards renewable heating technologies such as heat pumps, solar energy and industrial waste heat also brings a more decentralised approach with a need for distributed storage to exploit all the potentials of renewable heat [5].

This study investigates the impact of a novel operating strategy for a decentralised TTES to shave the peak loads of an urban DH system.

1.1. Forecasting in DH systems with machine learning

An important aspect of integrating the DH sector into a flexible, smart energy system is the forecasting of heat demand in order to optimise the control of the components of the DH system. In recent years, research has shown that the use of machine learning algorithms is very well suited to perform this task. The following literature review provides an overview of the various options available in this area. It is common for studies to test several different machine learning algorithms, as it is difficult to anticipate which method will deliver the best results for a particular application or data set.

Protić et al. [6] emphasise the importance of heat load forecasting for the implementation of model predictive control strategies for DH consumers. For this purpose, they developed a novel method based on the support vector machine algorithm that performed better in their experiments to forecast heat loads compared to genetic programming and artificial neural network models. The support vector machine algorithm also performs best in a comparison by Idowu et al. [7], but they find that feed forward neural networks and multiple linear regression were also suitable options to predict the thermal load in buildings.

Suryanarayana et al. [8] present methods based on linear, ridge, and lasso regression that require low computational effort, but they find

* Corresponding author.
E-mail address: solar@uni-kassel.de (U. Trabert).

https://doi.org/10.1016/j.energy.2024.131690
Received 20 December 2023; Received in revised form 2 April 2024; Accepted 16 May 2024
Available online 20 May 2024

Nomenclature

Latin symbols

$\dot{Q}$	Heat load (kW)
A	Area (m^2)
C	Cost (€)
c	Specific cost ($\frac{€}{MWh}$)
c_p	Specific heat capacity of water ($\frac{Wh}{kg\ K}$)
f	Factor (–)
Q	Heat (kWh)
T	Time period (h)
U	Heat transfer coefficient ($\frac{W}{m^2\ K}$)
V	Volume (m^3)

Greek symbols

ρ	Density of water ($\frac{kg}{m^3}$)
ϑ	Temperature (°C)

Abbreviations and subscripts

a	Annual
amb	Ambient
CAPEX	Capital expenditures
CHP	Combined heat and power
DH	District heating
fc	Forecast
HOB	Heat-only boiler
ind	Industry
LSTM	Long short-term memory
max	Maximum
min	Minimum
nom	Nominal
NRMSE	Normalised root mean square error
OPEX	Operational expenditures
opt	Optimised
org	Original
RL	Return line
sim	Simulated
SoC	State of charge
sto	Storage
t	Timestep
th	Thermal
TTES	Tank thermal energy storage

that the use of a deep neural net is more accurate in day-ahead forecasting of the load in DH systems. Saloux and Candanedo [9] confirm that machine learning algorithms such as decision trees, support vector machines, or artificial neural networks improve heat load forecasts compared to linear regression and see the potential for their inclusion in control applications. This aspect also plays a role for Potočnik et al. [10], who select the Gaussian process regression algorithm for an online forecasting solution for a DH utility in Ljubljana, mainly because the algorithm provides confidence intervals for the forecasts.

Huang et al. [11] and Wang et al. [12] are using graph neural networks to address spatiotemporal dependencies of data from heat meters in DH grids.

Another algorithm that has received a lot of attention recently due to its good performance in various studies is the Long Short-Term Memory (LSTM) algorithm. Leiprecht et al. [13] identify it as the best option when developing a web interface for heat load forecasting for a DH utility. Furthermore, Jesper et al. [14] show that LSTM models are also applicable to heat load profiles from the industrial and tertiary sectors when electricity consumption is included as an input parameter. Their intended application is the detection of anomalies in energy monitoring systems. In a new method developed by Chung et al. [15], the LSTM algorithm is combined with a convolutional neural network to take into account the importance of spatiotemporal characteristics of the input data.

In a comprehensive study of various artificial intelligence models, Runge and Saloux [16] find that the LSTM algorithm again provides the most accurate forecasts, although it requires the longest training time. In this study, they suggest the XGBoost algorithm as a less computationally intensive alternative. This result is confirmed by Xue et al. [17], where an XGBoost-based forecasting model performed best in a case study of a DH system in China.

The most frequently reported application for head load forecasts in the literature is to optimise the operation of heat production units. In this case, it is common that new measurement data is constantly available. In a concept called incremental learning, constant data streams are used to continuously update the machine learning models. Aragon et al. [18] apply this concept in combination with the LSTM algorithm to predict the electrical loads of a household for energy management systems.

From the literature review it is clear that machine learning algorithms and especially the LSTM algorithm appear to be an excellent fit for performing forecasting tasks in DH systems. However, the architecture of the models in most cases differs slightly and depends on the conditions and requirements of a particular case.

The challenges that this study addresses are how to improve results for small training data sets, but also how to account for changing heat load profile characteristics, which is common for developing industrial sites or DH consumers implementing energy efficiency measures.

1.2. Peak shaving in DH systems

There are various strategies to reduce the maximum peak load in DH systems described in the literature. Verda and Colella [19] show that a centralised TTES at the heat production site can be used to shave the peak loads in DH grids, reduce the use of heat only boilers (HOBs), and consequently reduce primary energy consumption and costs. Similar results are presented by Kauko et al. [20] and Knudsen et al. [21], who developed a model predictive control for the operation of a TTES at a heating plant that reduces the peak heating by 12%, thus minimising the use of HOB and maximising the use of industrial waste heat. Another common approach is to use the distribution grid as a heat storage by raising the supply temperature before the peak load. The underlying control strategies are studied by Laakkonen et al. [22] and Svensen [23]. A similar concept is implemented by van Oevelen et al. [24] to test a smart controller at a demo site in a subnetwork of the DH system of Brescia (Italy), leading to a reduction of the peak load energy supply. Despite the measures on the supply side of DH systems, several studies have also focused on the possibility of managing the demand side. As long as the thermal comfort is maintained, it is possible to use the thermal inertia of buildings to optimise the load in DH systems [25]. Guelpa and Marincioni [26] propose a control strategy for DH substations based on limiting the primary return temperature to reduce peak loads, while using a building model to keep indoor temperatures at the desired levels. A combination of reinforcement learning and agent-based modelling is used by Solinas et al. [27] to consider end-user acceptance of indoor temperature variations due to load shifting when optimising peak demand of the DH system. Combining demand side management by shifting the load of consumers with production optimisation through the use of thermal storage at the heat production site, Capone et al. [28] find that annual costs can be reduced by 3.8%. In a comprehensive literature review Guelpa and Verda [29] conclude that thermal peak loads can be reduced by up to

30% through demand response strategies in DH systems. In contrast to peak shaving at the system level, Li et al. [30] show how DH prosumers can optimise their own heat production from waste heat or renewable energies by using a thermal storage to reduce their individual peak load and maximise waste heat self-utilisation, achieving annual heating cost savings of up to 9%.

The strategies described in the literature for peak shaving in DH systems are mostly focused on either the system level or individual consumers. However, large DH consumers such as industrial sites bear the potential to have a major impact on the whole system, so this article explores an operational strategy for the individual consumer that supports the whole system, with the underlying case study described in Section 2. Furthermore, the operational optimisations in the mentioned studies have mostly been based on measured heat load profiles of the consumers. Whilst this provides a benchmark for the potential benefits of the strategies described, implementation in real systems relies on forecasts for heat loads and system temperatures and the accuracy of these can affect the performance of the proposed optimisations.

1.3. Research objective and scope

The main purpose of this article is to develop an optimisation-based operating strategy for a TTES at large DH consumers and to assess the effect of varying TTES dimensions using the case study described below as an example. The objective of the operating strategy is to decrease the peak load in the DH system by shifting the heat demand of the industrial DH consumer on the primary side through the use of a TTES on the secondary side, thus decreasing the need for HOBs based on fossil fuels.

To optimise the operation of the TTES, it is essential to forecast the heat demand profiles for both the DH system and the industrial consumer. In addition, forecasts of the return temperature of the secondary grid of the industrial consumer are necessary to calculate the TTES capacity, which depends on its minimum and maximum temperatures. Therefore, the upstream aim of this study is to use deep learning models to create accurate forecasts and to evaluate the impact of forecasting accuracy on the operating strategy.

In a real system, the proposed operating strategy could be implemented with a model predictive control. However, implementing a control that benefits the system, but acts on the consumers secondary equipment presents several challenges, such as data and control signal exchange, ownership, and tariff structures. To overcome these challenges, and as a first step towards automated control, the optimised operating strategy could be implemented as a decision support system, with the actuators (e.g. valves and pumps) still being operated under human supervision. However, human interaction should be kept to a minimum, so that the third aim of this article is to study the effect that the length of the optimisation interval may have on TTES operation.

This article is structured into the following sections:

- Section 2: Description of the underlying case study.
- Section 3: Methodology for the forecasting model, the optimisation framework for TTES operation, the TTES simulation model, and the evaluation metrics for forecasting accuracy and the simulation.
- Section 4: Results showing the forecasting accuracy, the simulated operation of the TTES, the simulation evaluation, and a discussion of the results.

2. Case study

The methods developed in this work are applied to a case study of a large urban DH system in Germany. Fig. 1 shows a simplified diagram of the components, load and temperature profiles from the case study that are relevant to this article. The primary heat production unit of the DH system is a central combined heat and power (CHP) plant that is operated with the energy carriers natural gas, waste, and biomass. Additionally, heat-only boilers (HOBs) based on natural gas and oil are used at another location near the city outskirts to cover peak loads in the system. At the end of 2022, an industrial manufacturing facility located near the CHP plant has been connected to the DH system with a nominal maximum heat load of 15 MW_{th}, which is around 7% of the total heat load in the DH system. The load profile of the new industrial consumer is dominated by space heating with only a low share of process heat. So far, the substation is operating according to the demand in the secondary grid of the consumer. However, the infrastructure of the industrial site offers the opportunity to install a tank thermal energy storage (TTES) at the secondary side of the substation that could be used to implement demand response strategies to benefit the DH system.

The presented research is based on operational data in hourly resolution from 2019 to 2021 of the industrial consumer (supply, return, and ambient temperature, load profile) and from 2021 to 2022 of the DH system (ambient temperature, load profiles, and heat production profiles of the CHP and the HOBs).

3. Methods

The following sections provide a brief overview of the forecasting techniques, followed by an elaboration on the optimisation framework and the model for simulating the TTES operation.

3.1. Forecasting

The long-short-term memory (LSTM) algorithm, which was developed to improve the prediction of time series with recurrent neural networks, is selected to perform the forecasting tasks for this work. Several studies demonstrated that the algorithm performs well in forecasting sequential data such as heat demand or return temperature profiles, which is an aim of this study [14,16,31,32].

The framework for creating deep learning models for multistep forecasts is set up in Python [33]. The forecast horizon is chosen to be 24 h to meet the demands of the optimisation problem that deals with the short-term operation of a TTES.

In the first step, the data sets from the case study are preprocessed using the Python library Scikit-Learn [34]. The input parameters to predict the load and return temperature profiles include ambient temperature, the preceding values of the forecast variable, and the timestamp. All data are available with an hourly resolution. The ambient temperature is first converted into a standardised timeseries, so that the mean of the values is 0 with a standard deviation of 1. It is then transformed into a lagged series, which means that the values of the previous time period (lag interval of 12 h) are input parameters. In this context, it should be mentioned that weather forecasts for the forecast horizon may be considered when practically applied. As historic ambient temperature forecasts are not available for the study, the measured ambient temperature for the forecast horizon is included in one scenario to predict future time steps (referred to as 'future weather data' in Section 4.1) to estimate the potential benefit from including weather forecasts. The preceding values of the forecast variable (load or return temperature) are normalised, i.e. rescaled to the interval [0,1] and also transformed into a lagged series with a lag interval of 24 h. Finally, the timestamp is encoded into sequences of 1s and 0s, with each set of digits representing an hour of the day or a day of the week.

Only data from 2021 were available for both the industrial consumer and the DH system. Consequently, it is chosen as the test year. The load and return temperature models for the industrial consumer are trained with data from 2019 and 2020, while the load model of the DH system is trained with data from 2022. The models are created using the Keras deep learning library [35]. A wide set of parameters can be tuned to improve model performance, which is done manually here. The parameters used are given in Table 1 and will not be discussed

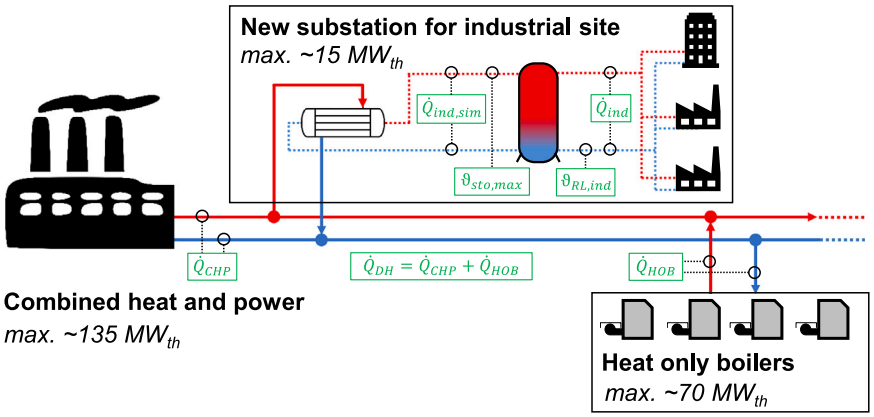

Fig. 1. Schematic representation of the case study.

Table 1
Parameters used for model training.

Profile	LSTM layers	No. of neurons	Batch size	Training epochs	Update epochs	Update interval
Industry	2	10	12	100	2	12 h
DH	2	10	24	100	2	24 h

in detail, since fine-tuning of the forecasting models is not the main topic of this article, but rather the effect of forecasting accuracy on the optimised operation of a TTES.

Another concept used in this work to improve the performance of the LSTM models is incremental learning. It is especially relevant for applications that involve constant data streams, such as the short-term optimisation of energy systems with data that is continuously recovered from the sensors of a physical system [18]. In this work the LSTM models are updated in a defined interval (12 or 24 h) with a data set that includes the most recent measurement data. A moving window is used to keep the length of the update training data set constant to account for the seasonality in the load and return temperature profiles.

The performance indicators that are used to evaluate forecasting accuracy are the coefficient of determination (R^2) and the normalised root mean square error (NRMSE) according to Eqs. (1) and (2):

$$R^2(y,\hat{y}) = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \overline{y})^2} \tag{1}$$

$$\mathrm{NRMSE}(y,\hat{y}) = \frac{\sqrt{\frac{1}{n}\cdot\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}}{\overline{y}} \tag{2}$$

where y is the measured value, $\hat{y}$ is the predicted value, $\overline{y}$ is the mean of the measured values in the test data set and n is the number of samples.

3.2. Storage operation at the industrial DH substation

The method for investigating the impact of a TTES at the industrial DH substation is partially similar to the concept of a model predictive control. The idea is to take an optimised operating strategy already into account when designing the system.

The basic procedure is depicted in Fig. 2. The starting point is the initial training of the LSTM models described in the previous section. The 24-hour forecasts of the three profiles industrial load $\dot{Q}_{ind,fc}$, industrial return temperature $\vartheta_{RL,ind,fc}$ and DH load $\dot{Q}_{DH,fc}$ are used in the subsequent optimisation (see Section 3.2.1) to create an operating schedule for the TTES. Afterwards, the actual storage operation is simulated using the measured load profile of the consumer $\dot{Q}_{ind}$ and a capacity model that was developed in Python (see Section 3.2.2). Finally, the LSTM models are fed with the preceding measured values for a new forecast. Optionally, the LSTM models are updated in the defined update interval. The new forecasts and the last storage state of charge (SoC) from the simulation are then used in the next optimisation. The chosen optimisation interval determines how often this process is repeated. The interval is varied between 1 and 24 h. The simulation was conducted for the year 2021, during which the data sets from the industrial consumer and the DH system overlapped.

3.2.1. Optimisation framework

The optimisation problem is formulated using the open-source Python library cvxpy [36,37] and solved with the embedded ECOS solver [38]. At the start of an optimisation cycle the storage capacity is estimated based on the temperature limits and the volume of the tank. The minimum storage temperature $\vartheta_{sto,min,fc}$ is estimated as an average of the return temperature of the industrial consumer $\vartheta_{RL,ind}$ in the previous 24 h and its respective forecast $\vartheta_{RL,ind,fc}$ for the next 24 h according to Eq. (3). This is done to account for the mixing of water inside the storage and water flowing into the storage from the return line of the secondary grid at the industrial site.

$$\vartheta_{sto,min,fc}(t) = \frac{\frac{\sum_{t=-23}^{t=0}\vartheta_{RL,ind}(t)}{24} + \frac{\sum_{t=1}^{t=24}\vartheta_{RL,ind,fc}(t)}{24}}{2} \tag{3}$$

The maximum storage temperature $\vartheta_{sto,max}$ is given by the storage charging temperature, which is assumed with a fixed value of 90 °C. This ensures an upper terminal temperature difference at the heat exchanger of the substation of at least 5 K, as the minimum supply temperature of the DH grid in summer is 95 °C. The storage charging temperature could also be assumed to be a function of ambient temperature or the supply temperature in the DH grid in order to increase storage capacity in winter. However, this would result in higher return temperatures on the primary side of the heat exchanger, which is not desired by the DH utility.

Subsequently, the minimum and maximum storage capacity Q_{sto} are determined depending on the TTES volume V_{sto} according to Eqs. (4) and (5).

$$Q_{sto,min,fc}(t) = f_{backup}\cdot V_{sto}\cdot c_p\cdot\rho\cdot(\vartheta_{sto,max} - \vartheta_{sto,min,fc}(t)) \tag{4}$$

$$Q_{sto,max,fc}(t) = f_{net}\cdot V_{sto}\cdot c_p\cdot\rho\cdot(\vartheta_{sto,max} - \vartheta_{sto,min,fc}(t)) \tag{5}$$

where c_p is the specific heat capacity and ρ is the density of water. f_{backup} (= 0.1) is a factor to maintain a backup capacity that is not used in the optimisation, but is available to the simulation to buffer differences between forecasted and actual load. f_{net} (= 0.9) is a factor to reduce the storage volume to account for dead volume in the tank.

The optimisation problem is formulated with the objective function in Eq. (6a) and the respective constraints in Eq. (6b) to (6i):

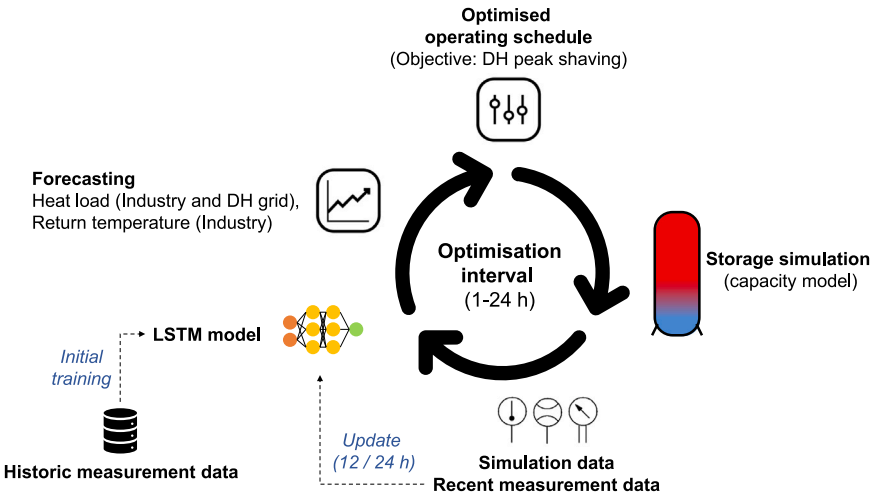

Fig. 2. Procedure diagram for the continuous optimisation.

$$\underset{\delta}{\text{minimize}} \quad \max(2 \cdot \Delta\dot{Q}_{DH} + \Delta\dot{Q}_{ind} + \max_{t=1\ldots24} \nabla\dot{Q}_{ind,opt}(t)) \tag{6a}$$

$$\text{subject to} \quad \Delta\dot{Q}_{DH} = \max_{t=1\ldots24} \dot{Q}_{DH,opt,fc}(t) - \min_{t=1\ldots24} \dot{Q}_{DH,opt,fc}(t), \tag{6b}$$

$$\Delta\dot{Q}_{ind} = \max_{t=1\ldots24} \dot{Q}_{ind,opt}(t) - \min_{t=1\ldots24} \dot{Q}_{ind,opt}(t), \tag{6c}$$

$$\dot{Q}_{DH,opt,fc}(t) = \dot{Q}_{ind,opt}(t) + \dot{Q}_{DH,2021,fc}(t), \tag{6d}$$

$$\dot{Q}_{ind,opt}(t) = \delta(t) \cdot \dot{Q}_{ind,2021,fc}(t), \tag{6e}$$

$$\dot{Q}_{ind,opt}(t) \geq 0, \tag{6f}$$

$$\dot{Q}_{ind,opt}(t) \leq \dot{Q}_{ind,nominal}, \tag{6g}$$

$$Q_{sto,t-1} + \sum_{t=1}^{24} (\dot{Q}_{ind,opt}(t) - \dot{Q}_{ind,fc}(t)) \geq Q_{sto,min,fc}(t), \tag{6h}$$

$$Q_{sto,t-1} + \sum_{t=1}^{24} (\dot{Q}_{ind,opt}(t) - \dot{Q}_{ind,fc}(t)) \leq Q_{sto,max,fc}(t) \tag{6i}$$

where $\delta(t)$ is the vector to modify the original load profile and $\dot{Q}$ represents the heat loads of the industry (ind) and the DH system (DH).

The objective function aims to minimise the maximum of a weighted sum. This sum is made up of the difference between the maximum and minimum value of the DH load $\Delta\dot{Q}_{DH,opt}$, the difference between the maximum and minimum value of the industry load $\Delta\dot{Q}_{ind}$, and its maximum gradient $\nabla\dot{Q}_{ind,opt}$ during the forecast horizon. The weights of the objective function (2, 1, 1) depend on the relative share of the industrial load in the total DH load and are selected by comparing operational results from simulations with different weighting factors. The criteria used in this process are, on the one hand, the (prioritised) reduction of the peaks and filling of the valleys of the DH load profile and, on the other hand, the avoidance of sudden load changes at the industrial consumer to keep the additional stress on the components of the substation low. For example, a stronger weight on the industrial load profile reduces the peak reduction in the DH load profile.

The constraints in Eqs. (6f) and (6g) keep the optimised industry load profile within the limits of the heat exchanger at the substation. The boundaries of the storage capacity are described through Eqs. (6h) and (6i).

3.2.2. Storage simulation

The detailed modelling of a thermal storage requires rather complex models that include several heat and mass flows within the storage and with its surroundings. This results in a relatively high computational effort. Therefore, a simplified simulation model for the TTES is developed in Python for this work, which is considered to be sufficiently accurate during the design phase of a system. The parameters used for the storage model are documented in Appendix A.2.

To calculate the minimum storage temperature and the maximum storage capacity, the same equations as described in Section 3.2.1 for the optimisation problem are used. However, the forecasted return temperature is replaced with the actual return temperature and the backup capacity, which was excluded in the optimisation problem, may be used in the simulation.

The storage is simulated for each timestep t depending on its state of charge with the following algorithms:

1. TTES is neither full nor empty (Algorithm 1),
2. TTES is almost full (Algorithm 2), and
3. TTES is almost empty (Algorithm 3).

The algorithms are described in detail through pseudocode in Appendix A.1. The deviations between forecasted and actual profiles require a control for the storage, that adjusts the charging or discharging load when the storage capacity is close to its upper or lower limit. This feature is implemented in Algorithm 2 and 3 as a damping factor and is inspired by temperature based controls in more complex models or real systems.

A large number of storage cycles and therefore a relatively short storage time are expected for a TTES used for daily peak shaving. This also means that the relative storage losses only play a minor role for the operating strategy, so they are not included in the optimisation problem and are estimated after the simulation according to Eqs. (7) and (8):

$$\vartheta_{sto,mean}(t) = \vartheta_{sto,min}(t) + (\vartheta_{sto,max} - \vartheta_{sto,min}(t)) \cdot \frac{Q_{sto}(t)}{Q_{sto,max}(t)} \tag{7}$$

$$\dot{Q}_{sto,loss}(t) = U_{sto} \cdot A_{sto} \cdot (\vartheta_{sto,mean}(t) - \vartheta_{amb,sto}) \tag{8}$$

where $\vartheta_{sto,mean}$ is the mean storage temperature of a fully mixed storage, $\dot{Q}_{sto,loss}$ is the heat loss, U_{sto} is the heat transfer coefficient of the storage envelope and $\vartheta_{amb,sto}$ is the ambient temperature at the TTES location.

The outcome of the simulation algorithms is a timeseries for the storage capacity Q_{sto}. Using Eq. (9) the resulting storage charging or discharging load $\dot{Q}_{sto}(t)$ can be calculated:

$$\dot{Q}_{sto}(t) = (Q_{sto}(t) - Q_{sto}(t-1)) \cdot \frac{1}{\Delta t} + \dot{Q}_{sto,loss}(t) \tag{9}$$

The simulated load $\dot{Q}_{ind,sim}$ is determined with the original load profile $\dot{Q}_{ind}$ and the storage load $\dot{Q}_{sto}$ according to Eq. (10):

$$\dot{Q}_{ind,sim}(t) = \dot{Q}_{ind}(t) + \dot{Q}_{sto}(t) \tag{10}$$

Finally, the original and the optimised DH load profiles $\dot{Q}_{DH,org}$ and $\dot{Q}_{DH,opt}$, which are compared to evaluate the operating strategy, are

given through Eqs. (11) and (12). It should be noted that the new substation had not yet been connected to the DH system in the simulation year 2021.

$$\dot{Q}_{DH,org}(t) = \dot{Q}_{DH,2021}(t) + \dot{Q}_{ind,2021}(t) \tag{11}$$

$$\dot{Q}_{DH,opt}(t) = \dot{Q}_{DH,2021}(t) + \dot{Q}_{ind,sim}(t) \tag{12}$$

3.2.3. Simulation evaluation

This study analyses the effects of varying forecasting accuracy, storage capacity (V_{sto} = [100, 200, …, 2500] m^3), and optimisation interval (T_{opt} = [1, 2, 3, 4, 6, 8, 12, 24] h) on TTES operation.

Three forecasting scenarios are considered:

1. Preceding data,
2. Future weather data, and
3. Perfect forecast.

The 'preceding data' scenario only considers recent measured data to forecast load and return temperature profiles. In the 'future weather data' scenario, the measured ambient temperature of the forecast horizon is used as an additional input parameter due to the lack of historical weather forecasts for the year 2021. In the 'perfect forecast' scenario the actual load and return temperature profiles are used in the optimisation to provide a benchmark for the best possible solution.

The nominal storage capacity $Q_{sto,nom}$ is calculated as shown in Eq. (13):

$$Q_{sto,nom} = V_{sto} \cdot c_p \cdot \rho \cdot (\vartheta_{sto,max} - \min_{t=1\ldots8760} \vartheta_{sto,min}(t)) \tag{13}$$

This is analogue to the calculation of the maximum storage capacity in Section 3.2.1, except that for the nominal capacity the full storage volume V_{sto} and the lowest minimum storage temperature $\vartheta_{sto,min}$ throughout the whole year is considered.

To draw more generalised conclusions from this work, the nominal storage capacity $Q_{sto,nom}$ is then converted from an energy amount to a time period $T_{sto,nom}$ according to Eq. (14):

$$T_{sto,nom} = \frac{Q_{sto,nom}}{\max_{t=1\ldots8760} \dot{Q}_{ind}(t)} \tag{14}$$

where $T_{sto,nom}$ is the storage capacity in hours and represents the time that a full storage could provide the maximum load of the industrial consumer $\dot{Q}_{ind,max}$.

The simulated peak shaving is evaluated with the relative reduction of the annual maximum peak of the DH load profile $d_{DH,max,peak}$ and the mean relative reduction of the daily maximum peaks of the DH load profile $d_{DH,daily,peak}$. Their calculation is described in Eqs. (15) and (16).

$$d_{DH,max,peak} = 1 - \frac{\max_{t=1\ldots8760} \dot{Q}_{DH,opt}(t)}{\max_{t=1\ldots8760} \dot{Q}_{DH,org}(t)} \tag{15}$$

$$d_{DH,daily,peak} = 1 - \frac{\sum_{i=0}^{364} \frac{\max_{t=(i\cdot24+1)\ldots(i\cdot24+24)} \dot{Q}_{DH,opt}(t)}{\max_{t=(i\cdot24+1)\ldots(i\cdot24+24)} \dot{Q}_{DH,org}(t)}}{365} \tag{16}$$

Additionally, it is analysed how the operating strategy can shift heat production from HOB to the CHP plant. This is done by determining the HOB heat $Q_{HOB,a}$, that is additionally required to supply the industrial consumer, using Eq. (17):

$$Q_{HOB,a} = \sum_{t=1}^{8760} \dot{Q}_{ind}(t) \cdot f_{HOB}(t) \tag{17}$$

where f_{HOB} represents the share of heat supplied through the HOBs to the industrial consumer. It is determined through Algorithm 4, which is shown in detail in Appendix A.3. It is based on the production schedule of the DH system in the reference year 2021 and describes the following rules:

- f_{HOB} is 1, if the HOBs are already in operation,
- f_{HOB} is 0, if the HOBs are not in operation and the new DH load is below the threshold for CHP supply ($\dot{Q}_{CHP,max}$ = 135 MW_{th}), and
- f_{HOB} is between 0 and 1, if the new DH load is just above $\dot{Q}_{CHP,max}$, but the HOBs are not in operation during the reference year.

Subsequently, the heat supplied through the HOBs with a TTES is compared to an unaltered load profile through the difference ΔQ_{HOB} following Eq. (18):

$$\Delta Q_{HOB,a} = Q_{HOB,a,org} - Q_{HOB,a,opt} \tag{18}$$

The technical evaluation is followed by a basic economic evaluation to determine an optimal storage capacity. The capital and operational expenditures (C_{CAPEX} and C_{OPEX}) for the TTES are estimated depending on the storage capacity with Eqs. (19) and (20) using the cost functions given in Große et al. [39]. The capital expenditures are adapted to the year 2023 with $f_{index,invest}$ (= 1.229) using an index for investment goods [40].

$$C_{CAPEX}(Q_{sto,nom}) = (-8\cdot(Q_{sto,nom}\cdot10^{-3})+23,000)\cdot Q_{sto,nom}\cdot10^{-3}\cdot f_{index,invest} \tag{19}$$

$$C_{OPEX}(Q_{sto,nom}) = 0.009 \cdot Q_{sto,nom} \tag{20}$$

The additional costs C_{HOB} that arise from using the HOBs instead of the CHP are calculated using Eq. (21):

$$C_{HOB} = \Delta Q_{HOB} \cdot 10^{-3} \cdot \Delta c_{HOB} \tag{21}$$

Since a distinct value for the additional specific cost Δc_{HOB} was not available for the study, Δc_{HOB} was varied between 20 €/MWh and 120 €/MWh.

Finally, the expenditures are compared to the costs by determining a static payback time $T_{payback}$ with Eq. (22):

$$T_{payback} = \frac{C_{CAPEX}}{C_{HOB} - C_{OPEX}} \tag{22}$$

4. Results and discussion

This section compares the results of the 'preceding data' and 'future weather data' forecasting scenarios and the effect of using model updates to implement incremental learning. The results of the storage simulation with the optimised operating strategy are then shown for the two forecasting scenarios and additionally with a 'perfect forecast' in terms of annual maximum peak reduction, mean daily peak reduction and HOB heat reduction.

4.1. Forecasting accuracy

In February 2021, there was a very cold winter period that caused high loads in the DH system and for the industrial consumer. These high loads are unique to the 2021 data set and are not included in the training data for the DH system model (2022) or the training data for the industrial consumer (2019 and 2020). Therefore, the load profile of the DH system (black line) from this period is shown in Fig. 3, as this period is well suited to show the differences between the forecasting scenarios. The load profile is normalised using its annual maximum value and plotted together with the forecasts for the next 24 h, which are renewed every eight hours (green lines with red markers at the start).

In the 'preceding data' scenario with no updates, the forecasts are approximately at the same level as the actual load (black line), but the shape of the daily profiles cannot be predicted (see subplot (a) in Fig. 3). This improves when the model is updated, but the forecasts still underestimate the actual load because the peak data was not included in the original training (see subplot (b) in Fig. 3). In the 'future weather

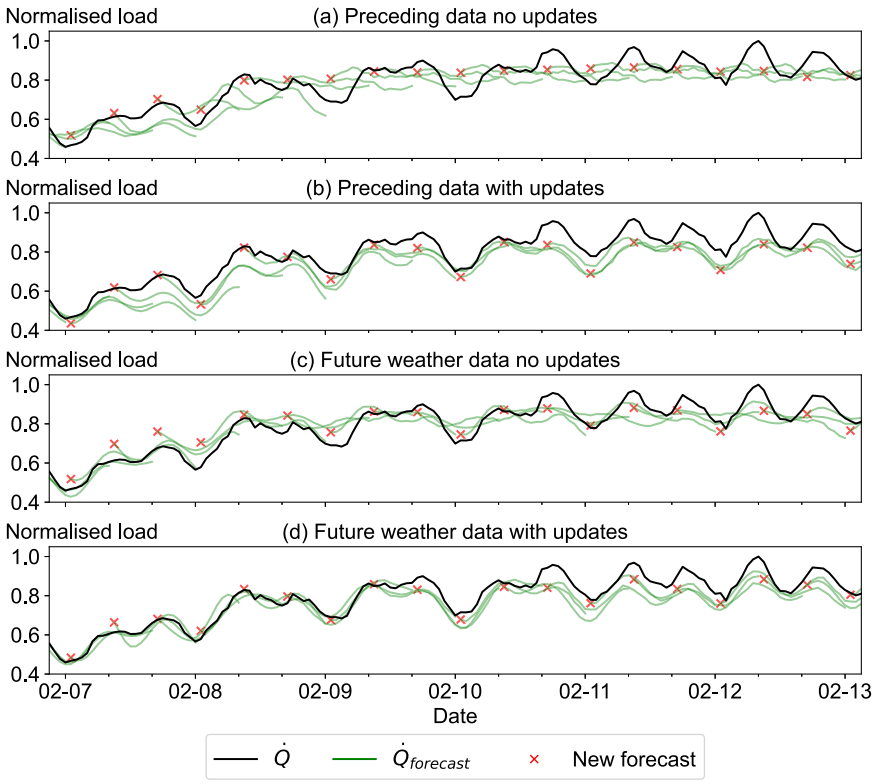

Fig. 3. Comparison of the normalised load of the DH system in the forecasting scenarios (a) Preceding data no updates, (b) Preceding data with updates, (c) Future weather data no updates, and (d) Future weather with updates from 7th to 13th of February 2021 (24-hour forecasts shown with new forecasts (x) every 8 h).

Table 2
Mean values of R^2 and NRMSE in the forecast horizon.

Forecasting scenario	Industrial load		DH load	
	R^2	NRMSE	R^2	NRMSE
Preceding data no updates	0.841	0.210	0.860	0.224
Preceding data with updates	0.924	0.144	0.945	0.139
Future weather data no updates	0.912	0.155	0.977	0.090
Future weather data with updates	0.967	0.096	0.986	0.072

data' scenario the same effect can be seen, but the forecasting accuracy is generally higher than in the 'preceding data' scenario (see subplots (c) and (d) in Fig. 3). The height of the peaks may still not be predicted exactly, but the time of the peak is very well matched.

Fig. 4 shows the same trend as just described for the peak load in February when comparing the scores (R^2 and NRMSE) of the forecasting scenarios for the whole year for each hour of the forecast horizon. While the model updates for the industrial profile result in a vertical shift of the score curves in the 'preceding data' scenario, they mainly lead to an increase in accuracy in the later hours of the forecast horizon in the 'future weather data' scenario. Another effect of model updates can be seen when looking at the scores for the DH load profile. The score curves do not fall monotonically with an increasing hour of the forecast, but this is achieved by including the model updates. This is an indication that the training data set of only one year for the DH load impairs the robustness of the model. However, this effect is relatively small in the 'future weather data' scenario, as the scores are already very good due to the high dependence of the DH load profile on ambient temperature.

The mean values for R^2 and NRMSE within the forecast horizon for the different variants are given in Table 2. Since the scores when using model updates are always higher, the scenarios without updates are disregarded for the evaluation of storage operation in the following Section 4.2.

4.2. Evaluation of storage operation

This section analyses the impact of the varied parameters (forecast scenario, storage capacity, and optimisation interval) on the simulation results. The storage capacity in hours is calculated according to Eq. (14) in Section 3.2.3. An example of the fundamental operating principle behind the optimisation is given in Fig. 5. It depicts the normalised DH load profile (a), the normalised industry load profile (b) and the corresponding storage SoC (c) for the first week of March. The normalisation for both profiles is done using the annual maximum value of the DH load as reference. The original load profiles are shown in black (DH) and blue (industry), while the profiles modified by the optimisation are shown in red (DH) and magenta (industry). When the modified profile of the industrial consumer deviates from the original profile, the storage is charged (green areas) or discharged (red areas). Charging is done primarily in the valleys of the DH load profile, while discharging reduces the peaks. If the load forecasts (grey and light blue lines) deviate strongly from the actual profile, the optimisation does not lead to the desired operation and the peak is not reduced (i.e. end of day on 2nd of March).

The storage is used extensively during this period with several storage cycles per day. In this specific example, 290 storage cycles are reached within the whole year, while the relative storage losses are only 0.7%. When the storage limits are reached, the simulated industry load (magenta line) deviates from the optimised operating schedule (dashed magenta line). A storage SoC of 100% corresponds to the maximum temperature difference at the top and the bottom of the storage throughout the year, which is reached when the return temperature of the industrial consumer is at its minimum during winter. In contrast, low loads tend to be accompanied by high return temperatures, so that the usable storage capacity is reduced in summer.

The two main reasons for seeking lower annual maximum peaks in a DH grid are to minimise the back-up capacities that a utility needs to provide for heat production, and to increase the amount of heat

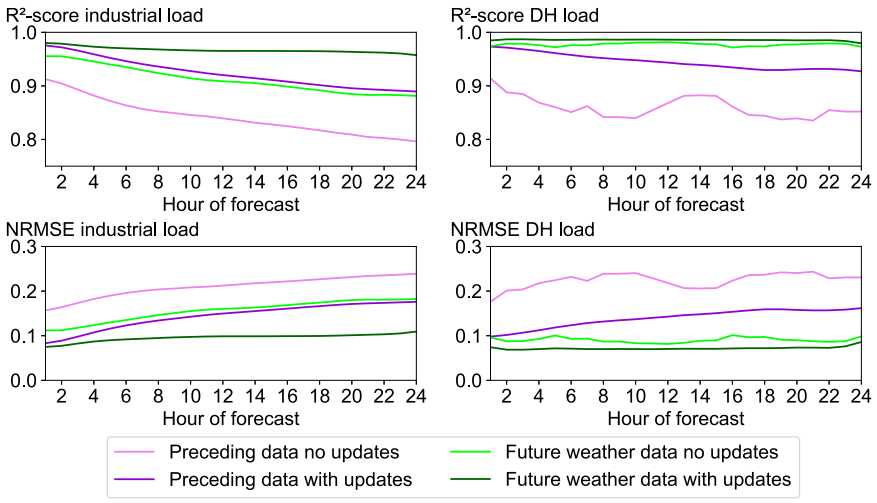

Fig. 4. Comparison of forecasting accuracy for one year: Coefficient of determination (R^2) and normalised root mean square error (NRMSE) for each hour of the forecast horizon.

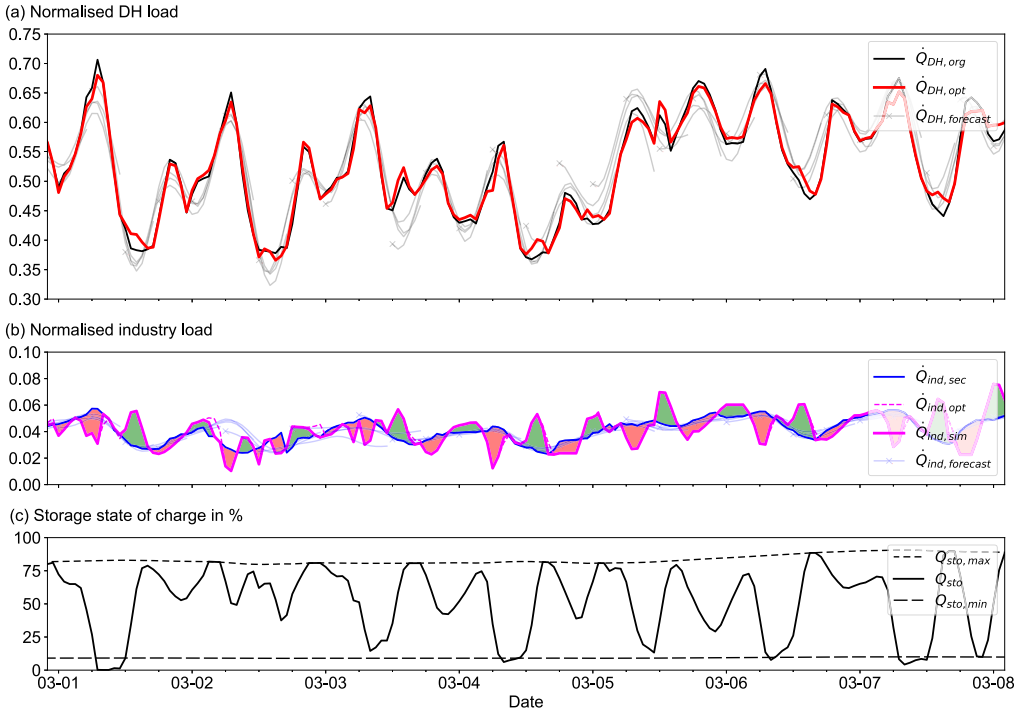

Fig. 5. Simulation results for the optimised storage operation from 1st to 7th of March 2021 with a storage capacity of 1.5 h ($V_{sto} = 600$ m^3) and an optimisation interval of six hours.

transported through the grid by keeping seasonal variation low. As these peaks are rare and were not included in the training data for the forecasting models, heat load forecasts during these periods are less accurate in this study. Fig. 6 shows the relative annual maximum peak reduction $d_{DH,max,peak}$ in the three forecasting scenarios. Theoretically, with perfect forecasts (a), the annual peak of the DH load profile could be reduced by around 7% with a storage capacity of four hours. This value is at around 4% in the future weather data scenario (b) and around 2.8% in the preceding data scenario (c) at the same storage capacity. The reduction depends on the load profile of one specific day (12.02.2021) within the whole year. The optimisation interval is therefore a sensitive parameter, as it determines at which hours of this day the optimisation is performed in respect to the peak hour. Especially in the preceding data scenario (c), it is observed that the simulated peak reduction is lower, when an optimisation is performed during the peak hour (here: 8 a.m.), which is true for T_{opt} = [1, 2, 4, 8] h. In the future weather data scenario (b) this is only visible at T_{opt} = 8 h, so that improved forecasting accuracy and consequently less underestimation of the load in the hours before the peak seem to cancel out this effect (see also plot (b) and (d) in Fig. 3 to compare the forecasts during the peak day 12.02.2021). In general, a low forecasting accuracy during these peak hours affects the ability to operate the TTES optimally. Nevertheless, the trend of the results shows that a short optimisation interval in combination with the use of weather forecasts are recommended to reduce annual maximum peaks.

The mean daily peak reduction $d_{DH,daily,peak}$, which is shown in Fig. 7, is much less sensitive to the optimisation interval and depends mainly on the storage capacity. However, the 3D plots in Fig. 7 show a plateau, which means that at some point increasing the storage size no longer leads to a significant improvement of $d_{DH,daily,peak}$. The plateau starts with a storage capacity of three hours and is around 6% to 7% with perfect forecasts (a), while it starts with a storage capacity of 4 h

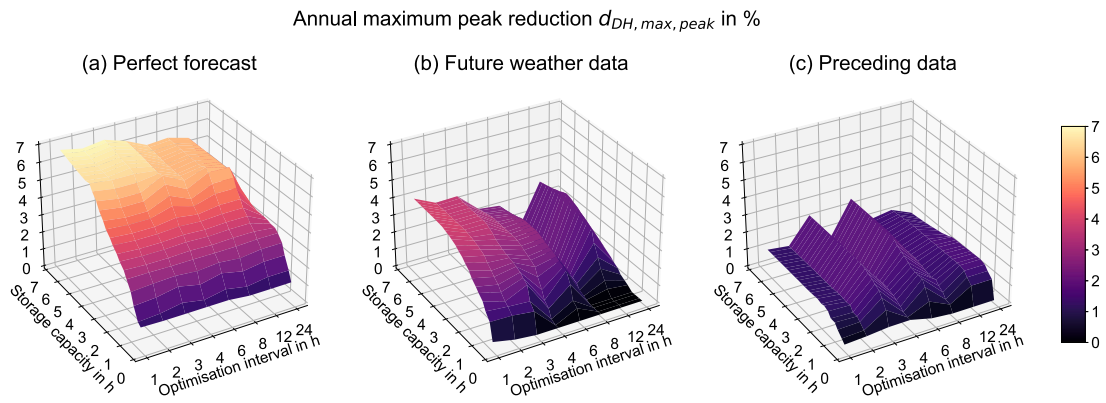

Fig. 6. Relative annual maximum peak reduction $d_{DH,max,peak}$ in the three forecasting scenarios (a) to (c) for varying storage capacities and optimisation intervals.

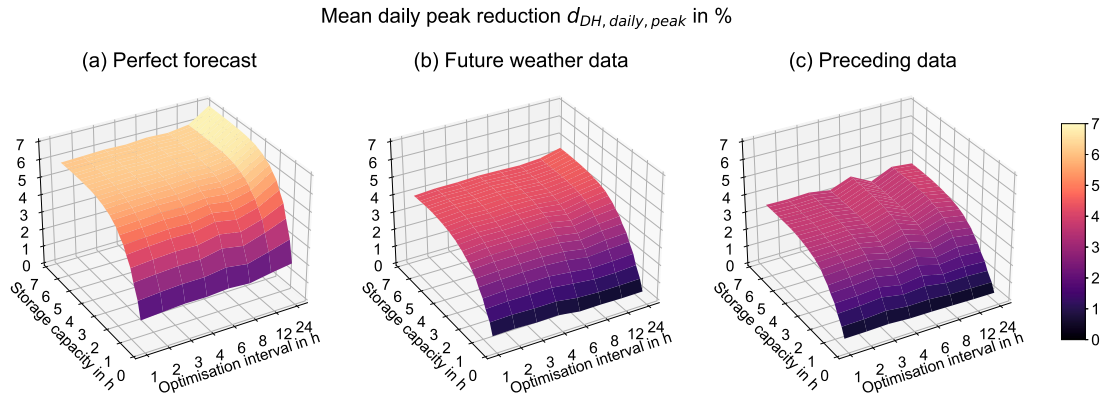

Fig. 7. Relative mean daily peak reduction $d_{DH,daily,peak}$ in the three forecasting scenarios for varying storage capacities and optimisation intervals.

and is around 4% to 5% in the future weather data scenario (b) and around 3% to 4% in the preceding data scenario (c). It is also noticeable that an optimisation interval of 24 h on average seems to work best to reduce the daily peaks. This is because the optimisation interval in this case is the same as the forecast horizon and the optimisation is always performed for a period of a single day. If the optimisation interval does not match the forecast horizon, it is possible that a new daily peak is created to charge the storage at the end of a day to reduce the peak of the next day.

The evaluation of storage operation is concluded with the analysis of how much heat from the DH system for the industrial consumer can be shifted from HOB to CHP supply. The 3D plots in Fig. 8 show that with decreasing forecasting accuracy ΔQ_{HOB} also decreases from up to 3 GWh/a with perfect forecasts (a) down to 2.4 GWh/a in preceding data (c). Once again, the benefit of a larger storage diminishes at a storage capacity of around three to four hours. The optimisation interval only has a minor influence on ΔQ_{HOB}. However, for future weather data (b) and preceding data (c), 8 and 24 h optimisation intervals are slightly out of trend, which can be caused by load underestimations in the forecasts before the daily peak hour (on average 9 am). In general, the results indicate that longer optimisation intervals do not affect the performance of the operating strategy.

Finally, ΔQ_{HOB} is used to calculate a static payback period $T_{payback}$ for the installation of a TTES for different specific values for the additional cost Δc_{HOB} caused by HOB heat compared to CHP heat. Fig. 9 indicates that the economic optimum is reached with storage capacities between one and two hours. The influence of forecasting accuracy on $T_{payback}$ increases with larger storage capacities and lower additional costs Δc_{HOB}. Assuming a technical lifetime of 25 years, the TTES can be an economically viable solution for the DH utility starting from specific additional costs of 40 €/MWh for HOB heat even in the scenario with the lowest forecasting accuracy.

4.3. Discussion

The LSTM models developed in this work for forecasting achieve good scores in terms of the coefficient of determination R^2 and the normalised root mean square error NRMSE. They are on the same level or better compared to similar studies on the use of machine learning models to predict head loads [6,7,14–17]. However, the models are tuned to work well on the specific data set that was available from the case study, so that an application with new data may require re-tuning of the model parameters and a new selection of input parameters. In general, the results of the 'future weather data' scenario are overestimated because in a real application the measured ambient temperature would have to be replaced by weather forecasts. It is expected that this will result in a forecasting accuracy between the 'preceding data' and the 'future weather data' scenario. In contrast, the forecasts in this study are only based on one year of data for the DH load and two years of data for the industry load. The forecasts and subsequent benefits are expected to increase with a larger database.

In principle, the results of this study can be transferred to other cases. The benefits of the proposed TTES in combination with a large consumer in a DH grid depend on the ratio of the peak demands of the DH system and the industrial (or other large) consumer in combination with the ratio of the installed capacities of CHP and HOBs.

The formulation of the objective function is a critical step in the design of multi-objective optimisation problems, and in the case of TTES in DH systems it is often determined by the individual framework conditions given by a case study. While Capone et al. [28] and Nakama et al. [41] include the operational costs for peak heating in their objective function, this study only considers the load profiles due to uncertain economic framework conditions. Knudsen et al. [21] use a similar objective function in a model predictive controller that

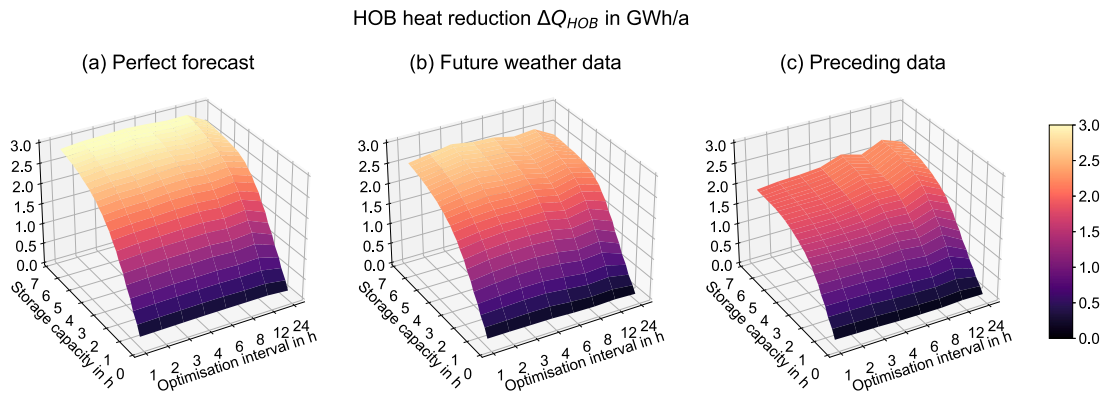

Fig. 8. HOB heat reduction ΔQ_{HOB} in the three forecasting scenarios (a) to (c) for varying storage capacities and optimisation intervals.

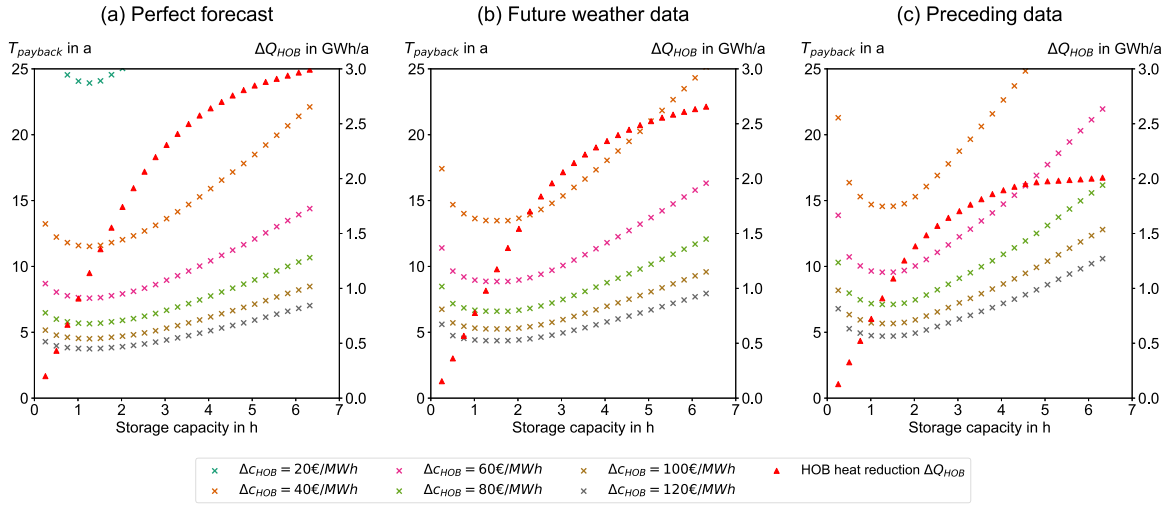

Fig. 9. Payback period $T_{payback}$ and HOB heat reduction ΔQ_{HOB} in the three forecasting scenarios (a) to (c) for different specific additional costs Δc_{HOB} and storage capacities with $T_{opt} = 1$.

minimises peak heating. However, the results are difficult to compare in terms of peak heating reduction due to different heat generation technologies and the location of the TTES at the production plant as opposed to the decentralised TTES considered in this study.

The simulation model developed for the TTES is a simple approach that considers the storage as a capacity. Details such as the mass flows at inlet and outlets and within the storage, as well as temperature stratification, are neglected in order to keep computational effort low and to perform a wide parameter variation. These details should be addressed in the further planning process for the implementation of the proposed operating strategy and in the translation of the methods into control signals for the actual system. Preliminary results of detailed thermal simulations of the substation heat exchanger and the TTES, which are part of future work, indicate that the absolute values of the considered KPIs will differ, but the determined relatively broad economic optimum of a storage capacity between 1 and 2 h should be in the same range.

The decision whether to install a TTES is highly dependent on the economic framework conditions of a DH utility and its relationship with its customer. The proposed application requires good cooperation between the parties involved in order to achieve data transfers, exchange of control signals, and a mutually beneficial tariff structure. If this is the case, the economic evaluation should take into account more details such as the interest rate, an investment estimate based on manufacturers quotes, and the actual additional cost of HOB heat. The results of this study suggest that the demand response strategy outlined could be a viable business model in the DH sector.

5. Conclusions and future work

In this work, a novel operating strategy was developed for the integration of a TTES at a large industrial DH consumer with the aim of reducing peaks in an urban DH system. An important aspect of the analysis was the development of accurate forecasting models, which are necessary for the operating strategy to reach its full potential. The strategy is based on an optimisation problem that aims to modify the load profile of the industrial DH consumer by charging the TTES during the valleys and discharging it during the peaks of the DH system's load profile. This helps to reduce the use of carbon-intensive HOBs and increases the use of CHP heat.

The data used in this study is historical measurement data from a real system in Germany with the described design. It includes load, supply, return, and ambient temperature profiles as well as heat production data. The data is used to train LSTM machine learning models to produce 24-hour forecasts of the heat load and return temperature of the industrial consumer and the heat load of the DH system. Forecasting accuracy was assessed using the coefficient of determination R^2 and the normalised root mean square error as metrics. Applying continuous model updates through incremental learning significantly improves the forecasts only based on preceding data, reaching an R^2 of 0.924 for the industrial profile and 0.945 for the DH load profile, and a NRMSE of 0.144 and 0.139 respectively. In a scenario that included 'future weather data', further improvement of forecasting accuracy through the inclusion of weather forecasts as an input parameter was outlined.

The operating strategy was tested with a simulation of the TTES operation for the year 2021. A simplified capacity-based storage model was developed to allow for a wide parameter variation, including forecasting scenarios, storage capacity and the optimisation interval. The results showed that accurate forecasting and a short optimisation interval are particularly important for reducing the annual maximum peak. With a storage capacity of four hours the peak in the DH load profile can be reduced by about 4% in the 'future weather data' scenario just by modifying the industrial load profile. It was also found that a storage capacity above four hours did not lead to a further reduction in daily peaks. Longer optimisation intervals of up to 24 h have no significant impact on the performance of the operating strategy in terms of annual mean daily peak reduction, nor on the shift from costly HOB to CHP heat in the DH system. A storage capacity between one and two hours (400 to 700 m^3) was determined as the economic optimum. The installation of a TTES in combination with the implementation of the proposed operating strategy can be an economically viable solution starting from an additional cost of HOB heat compared to CHP heat of 40 €/MWh. The forecasting accuracy is important for economic viability, if the additional costs for HOB heat are less than 60 €/MWh.

As a next step, the results of the study can serve as a starting point for the development of a model predictive control or an on-line decision support system for the potential operators of a TTES at the industrial site investigated in the case study. This requires a more detailed and complex simulation model that includes, in addition to the TTES, other components at the DH substation such as the heat exchanger, pumps, and control valves. In addition, it should be investigated whether the integration of an electricity price-driven heat pump to utilise industrial waste heat can transform the substation into a prosumer, supporting the transformation of the urban DH system towards renewable energy supply.

CRediT authorship contribution statement

Ulrich Trabert: Writing – original draft, Visualization, Software, Methodology, Formal analysis, Conceptualization. **Felix Pag:** Writing – review & editing, Project administration, Funding acquisition, Conceptualization. **Janybek Orozaliev:** Writing – review & editing, Funding acquisition, Conceptualization. **Ulrike Jordan:** Supervision, Funding acquisition. **Klaus Vajen:** Supervision, Funding acquisition.

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 data that has been used is confidential.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used 'DeepL - Write' in order to check grammar and improve readability of the text. After using this service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Acknowledgements

This work was supported by the German Federal Ministry for Economic Affairs and Climate Action [Funding Code: 03EN3023A].

Appendix. Algorithms and parameters for modelling

A.1. Simulation algorithms

Algorithm 1 Storage neither full nor empty

$\dot{Q}_{\mathrm{sto,theo}}(t) = \dot{Q}_{ind,opt}(t) - \dot{Q}_{ind}(t)$
for $t = 1$ to T_{opt} **do**
  if $Q_{\mathrm{sto}}(t-1) \le f_{full} \cdot Q_{\mathrm{sto,max}}(t)$
    and $Q_{\mathrm{sto}}(t-1) \ge f_{empty} \cdot Q_{\mathrm{sto,max}}(t)$
    and $Q_{\mathrm{sto}}(t-1) + \dot{Q}_{\mathrm{sto,theo}}(t) \le Q_{\mathrm{sto,max}}(t)$
    and $Q_{\mathrm{sto}}(t-1) + \dot{Q}_{\mathrm{sto,theo}}(t) \ge 0$ **then**
    $Q_{\mathrm{sto}}(t) \leftarrow Q_{\mathrm{sto}}(t-1) + \dot{Q}_{\mathrm{sto,theo}}(t) \cdot \Delta t$
  end if ▷ to be continued

Algorithm 2 Storage almost full (continued)

  if $Q_{\mathrm{sto}}(t-1) > f_{full} \cdot Q_{\mathrm{sto,max}}(t)$
    or $Q_{\mathrm{sto}}(t-1) + \dot{Q}_{\mathrm{sto,theo}}(t) > Q_{\mathrm{sto,max}}(t)$ **then**
    if $\dot{Q}_{\mathrm{sto,theo}}(t) > 0$ **then**
      $f_{damping} \leftarrow \frac{Q_{\mathrm{sto,max}}(t-1) - Q_{\mathrm{sto}}(t-1)}{(1-f_{full}) \cdot Q_{\mathrm{sto,max}}(t-1)}$
      if $Q_{\mathrm{sto}}(t-1) + f_{damping} \cdot \dot{Q}_{\mathrm{sto,theo}}(t) \cdot \Delta t \le Q_{\mathrm{sto,max}}(t)$ **then**
        $Q_{\mathrm{sto}}(t) \leftarrow Q_{\mathrm{sto}}(t-1) + f_{damping} \cdot \dot{Q}_{\mathrm{sto,theo}}(t) \cdot \Delta t$
      else
        $Q_{\mathrm{sto}}(t) \leftarrow Q_{\mathrm{sto,max}}(t)$
      end if
    else
      $Q_{\mathrm{sto}}(t) \leftarrow Q_{\mathrm{sto}}(t-1) + \dot{Q}_{\mathrm{sto,theo}}(t) \cdot \Delta t$
    end if
  end if ▷ to be continued

Algorithm 3 Storage almost empty (continued)

  if $Q_{\mathrm{sto}}(t-1) < f_{empty} \cdot Q_{\mathrm{sto,max}}(t)$
    or $Q_{\mathrm{sto}}(t-1) + \dot{Q}_{\mathrm{sto,theo}}(t) < 0$ **then**
    if $\dot{Q}_{\mathrm{sto,theo}}(t) < 0$ **then**
      $f_{damping} \leftarrow \frac{Q_{\mathrm{sto}}(t-1)}{f_{empty} \cdot Q_{\mathrm{sto,max}}(t)}$
      if $Q_{\mathrm{sto}}(t-1) + f_{damping} \cdot \dot{Q}_{\mathrm{sto,theo}}(t) \ge 0$ **then**
        $Q_{\mathrm{sto}}(t) \leftarrow Q_{\mathrm{sto}}(t-1) + f_{damping} \cdot \dot{Q}_{\mathrm{sto,theo}}(t)$
      else
        $Q_{\mathrm{sto}}(t) \leftarrow 0$
      end if
    else
      $Q_{\mathrm{sto}}(t) \leftarrow Q_{\mathrm{sto}}(t-1) + \dot{Q}_{\mathrm{sto,theo}}(t) \cdot \Delta t$
    end if
  end if
end for

$\dot{Q}_{\mathrm{sto,theo}}(t)$: Theoretical storage charging/discharging load; $\dot{Q}_{ind,opt}(t)$: Optimised industry load; $\dot{Q}_{ind}(t)$: Original industry load; t: Timestep; T_{opt}: Optimisation interval; $Q_{sto}(t)$: Storage capacity; $Q_{sto,max}(t)$: Maximum storage capacity; f_{full}: Percentage of the state of charge from which on the storage charging is damped; f_{empty}: Percentage of the state of charge from which on the storage discharging is damped; $f_{damping}$: Damping factor applied to storage charging or discharging load

A.2. Simulation parameters

See Table A.3.

Table A.3
Parameters used for the storage model.

Parameter	Value	Parameter	Value	Parameter	Value
f_{backup}	0.1	$\vartheta_{sto,max}$	90 °C	U_{sto}	0.18 $\frac{W}{m^2 \cdot K}$
f_{net}	0.9	$\vartheta_{amb,sto}$	15 °C	A_{sto}	0.75 $\frac{1}{m} \cdot V_{sto}$
f_{full}	0.9			ρ	975 $\frac{kg}{m^3}$
f_{empty}	0.25			c_p	1.16 $\frac{Wh}{kg \cdot K}$
$f_{index,invest}$	1.229				

A.3. Determining HOB share

Algorithm 4 Determining HOB share

for $t = 1$ to 8760 **do**
  if $\dot{Q}_{HOB,org} > 0$ **then**
    $f_{HOB} \leftarrow 1$
  end if
  if $\dot{Q}_{DH,org/opt} \geq \dot{Q}_{CHP,max}$ **then**
    if $\dot{Q}_{DH,org/opt} - \dot{Q}_{ind,org/opt} < \dot{Q}_{CHP,max}$ **then**
      $f_{HOB} \leftarrow \frac{\dot{Q}_{DH,org/opt} - \dot{Q}_{CHP,max}}{\dot{Q}_{ind,org/opt}}$
    else
      $f_{HOB} \leftarrow 1$
    end if
  else
    $f_{HOB} \leftarrow 0$
  end if
end for

$\dot{Q}_{HOB,org}$: Original heat generation of HOB; f_{HOB}: Share of industrial heat load provided through HOB; $\dot{Q}_{DH,org/opt}$: Original/optimised DH heat load; $\dot{Q}_{ind,org/opt}$: Original/optimised industry heat load; $\dot{Q}_{CHP,max}$: Maximum heat generation of CHP (= 135 MW_{th})

References

[1] Lund H, Østergaard PA, Connolly D, Ridjan I, Mathiesen BV, Hvelplund F, et al. Energy storage and smart energy systems. Int J Sustain Energy Plan Manag 2016;11:3–14. http://dx.doi.org/10.5278/IJSEPM.2016.11.2.

[2] Lund H, Østergaard PA, Connolly D, Mathiesen BV. Smart energy and smart energy systems. Energy 2017;137:556–65. http://dx.doi.org/10.1016/j.energy.2017.05.123.

[3] 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;158:2430–4. http://dx.doi.org/10.1016/j.egypro.2019.01.302.

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

[5] Roberto R, de Iulio R, Di Somma M, Graditi G, Guidi G, Noussan M. A multi-objective optimization analysis to assess the potential economic and environmental benefits of distributed storage in district heating networks: a case study. Int J Sustain Energy Plann Manag 2019;20:5–20. http://dx.doi.org/10.5278/IJSEPM.2019.20.2.

[6] Protić M, Shamshirband S, Petković D, Abbasi A, Mat Kiah ML, Unar JA, et al. Forecasting of consumers heat load in district heating systems using the support vector machine with a discrete wavelet transform algorithm. Energy 2015;87:343–51. http://dx.doi.org/10.1016/j.energy.2015.04.109.

[7] Idowu S, Saguna S, Åhlund C, Schelén O. Applied machine learning: Forecasting heat load in district heating system. Energy Build 2016;133:478–88. http://dx.doi.org/10.1016/j.enbuild.2016.09.068.

[8] 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;157:141–9. http://dx.doi.org/10.1016/j.energy.2018.05.111.

[9] Saloux E, Candanedo JA. Forecasting district heating demand using machine learning algorithms. Energy Procedia 2018;149:59–68. http://dx.doi.org/10.1016/j.egypro.2018.08.169.

[10] Potočnik P, Škerl P, Govekar E. Machine-learning-based multi-step heat demand forecasting in a district heating system. Energy Build 2021;233:110673. http://dx.doi.org/10.1016/j.enbuild.2020.110673.

[11] Huang Y, Zhao Y, Wang Z, Liu X, Liu H, Fu Y. Explainable district heat load forecasting with active deep learning. Appl Energy 2023;350:121753. http://dx.doi.org/10.1016/j.apenergy.2023.121753.

[12] Wang Z, Liu X, Huang Y, Zhang P, Fu Y. A multivariate time series graph neural network for district heat load forecasting. Energy 2023;278:127911. http://dx.doi.org/10.1016/j.energy.2023.127911.

[13] Leiprecht S, Behrens F, Faber T, Finkenrath M. A comprehensive thermal load forecasting analysis based on machine learning algorithms. Energy Rep 2021;7:319–26. http://dx.doi.org/10.1016/j.egyr.2021.08.140.

[14] Jesper M, Pag F, Vajen K, Jordan U. Heat load profiles in industry and the tertiary sector: Correlation with electricity consumption and ex post modeling. Sustainability 2022;14(7):4033. http://dx.doi.org/10.3390/su14074033.

[15] Chung WH, Gu YH, Yoo SJ. District heater load forecasting based on machine learning and parallel CNN-LSTM attention. Energy 2022;246:123350. http://dx.doi.org/10.1016/j.energy.2022.123350.

[16] Runge J, Saloux E. A comparison of prediction and forecasting artificial intelligence models to estimate the future energy demand in a district heating system. Energy 2023;269:126661. http://dx.doi.org/10.1016/j.energy.2023.126661.

[17] Xue P, Jiang Y, Zhou Z, Chen X, Fang X, Liu J. Multi-step ahead forecasting of heat load in district heating systems using machine learning algorithms. Energy 2019;188:116085. http://dx.doi.org/10.1016/j.energy.2019.116085.

[18] Aragon G, Puri H, Grass A, Chala S, Beecks C. Incremental deep-learning for continuous load prediction in energy management systems. In: 2019 IEEE milan powerTech. IEEE; 2019, p. 1–6. http://dx.doi.org/10.1109/PTC.2019.8810793.

[19] Verda V, Colella F. Primary energy savings through thermal storage in district heating networks. Energy 2011;36(7):4278–86. http://dx.doi.org/10.1016/j.energy.2011.04.015.

[20] Kauko H, Rohde D, Knudsen BR, Sund-Olsen T. Potential of thermal energy storage for a district heating system utilizing industrial waste heat. Energies 2020;13(15):3923. http://dx.doi.org/10.3390/en13153923.

[21] Knudsen BR, Rohde D, Kauko H. Thermal energy storage sizing for industrial waste-heat utilization in district heating: A model predictive control approach. Energy 2021;234:121200. http://dx.doi.org/10.1016/j.energy.2021.121200.

[22] 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. http://dx.doi.org/10.1016/j.egypro.2017.05.076.

[23] Svensen JL. Peak shaving in district heating utilizing adaptive predictive control. Energies 2022;15(22):8555. http://dx.doi.org/10.3390/en15228555.

[24] van Oevelen T, Neven T, Brès A, Schmidt R-R, Vanhoudt D. Testing and evaluation of a smart controller for reducing peak loads and return temperatures in district heating networks. Smart Energy 2023;10:100105. http://dx.doi.org/10.1016/j.segy.2023.100105.

[25] Guelpa E, Barbero G, Sciacovelli A, Verda V. Peak-shaving in district heating systems through optimal management of the thermal request of buildings. Energy 2017;137:706–14. http://dx.doi.org/10.1016/j.energy.2017.06.107.

[26] Guelpa E, Marincioni L. Demand side management in district heating systems by innovative control. Energy 2019;188:116037. http://dx.doi.org/10.1016/j.energy.2019.116037.

[27] Solinas FM, Bottaccioli L, Guelpa E, Verda V, Patti E. Peak shaving in district heating exploiting reinforcement learning and agent-based modelling. Eng Appl Artif Intell 2021;102:104235. http://dx.doi.org/10.1016/j.engappai.2021.104235.

[28] Capone M, Guelpa E, Mancò G, Verda V. Integration of storage and thermal demand response to unlock flexibility in district multi-energy systems. Energy 2021;237:121601. http://dx.doi.org/10.1016/j.energy.2021.121601.

[29] Guelpa E, Verda V. Demand response and other demand side management techniques for district heating: A review. Energy 2021;219:119440. http://dx.doi.org/10.1016/j.energy.2020.119440.

[30] Li H, Hou J, Tian Z, Hong T, Nord N, Rohde D. Optimize heat prosumers' economic performance under current heating price models by using water tank thermal energy storage. Energy 2022;239:122103. http://dx.doi.org/10.1016/j.energy.2021.122103.

[31] Hochreiter S, Schmidhuber J. Long short-term memory. Neural Comput 1997;9(8):1735–80. http://dx.doi.org/10.1162/neco.1997.9.8.1735.

[32] Gers FA, Schmidhuber J, Cummins F. Learning to forget: continual prediction with LSTM. Neural Computat 2000;12(10):2451–71. http://dx.doi.org/10.1162/089976600300015015.

[33] Oliphant TE. Python for scientific computing. Comput Sci Eng 2007;9(3):10–20. http://dx.doi.org/10.1109/MCSE.2007.58.

[34] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Müller A, Nothman J, Louppe G, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay É. Scikit-learn: Machine learning in Python. J Mach Learn Res 2012;12:2825–30. http://dx.doi.org/10.48550/arXiv.1201.0490.

[35] Chollet F. Keras. 2015, URL https://keras.io.

[36] Diamond S, Boyd S. CVXPY: A python-embedded modeling language for convex optimization. J Mach Learn Res 2016;17(83):1–5.

[37] Agrawal A, Verschueren R, Diamond S, Boyd S. A rewriting system for convex optimization problems. J Control Decis 2018;5(1):42–60. http://dx.doi.org/10.1080/23307706.2017.1397554.

[38] Domahidi A, Chu E, Boyd S. ECOS: An SOCP solver for embedded systems. In: European control conference. 2013, p. 3071–6.
[39] Große R, Binder C, Wöll S, Geyer R, Robbi S. Long term (2050) projections of techno-economic performance of large-scale heating and cooling in the EU. Luxembourg: European Comission, Joint Research Centre; 2017, http://dx.doi.org/10.2760/24422, URL https://data.europa.eu/doi/10.2760/24422,
[40] Erzeugerpreise gewerblicher Produkte - ausgewählte Indizies: Investitionsgüter. 2023, URL https://www.destatis.de/DE/Home/.
[41] Nakama CS, Knudsen BR, Tysland AC, Jäschke J. A simple dynamic optimization-based approach for sizing thermal energy storage using process data. Energy 2023;268:126671. http://dx.doi.org/10.1016/j.energy.2023.126671.