Energy 197 (2020) 117177

Contents lists available at ScienceDirect

Energy

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

Testing and performance evaluation of the STORM controller in two demonstration sites

Tijs Van Oevelen [a, b, *], Dirk Vanhoudt [a, b], Christian Johansson [c], Ed Smulders [d]

[a] VITO NV, Boeretang 200, 2400 Mol, Belgium
[b] EnergyVille, Thor Park 8310, 3600 Genk, Belgium
[c] NODA Intelligent Systems AB, Netport Science Park, Biblioteksgatan 4, 374 35, Karlshamn, Sweden
[d] Mijnwater BV, Valkenburgerweg 177, 6419 AT Heerlen, the Netherlands

ARTICLE INFO

Article history:
Received 28 October 2019
Received in revised form
29 January 2020
Accepted 15 February 2020
Available online 20 February 2020

Keywords:
District heating and cooling (DHC)
Demand-side management
STORM controller
Testing
Performance evaluation

ABSTRACT

District heating and cooling (DHC) network control systems have a big part to play in enhancing energy system integration, which is vital to allow the transition towards a more sustainable and affordable energy system. The STORM controller is an example of an advanced DHC network controller that allows activating the flexibility from building thermal capacity. The STORM controller is capable of shifting and managing thermal demands in time, to improve the operational performance of an entire DHC system. This article presents the basic properties of the STORM controller, and the results of a field test campaign in two operational demonstration networks: in Heerlen (The Netherlands) and Rottne (Sweden). The performance of three control strategies is evaluated and discussed. The peak shaving strategy led to 3.1% reduction of peak heat production in Rottne. The market interaction tests demonstrated the possibility to temporarily increase the heat load by up to 96%, by charging the buildings, while limiting the overall heat consumption increase to 5.8%. The cell balancing tests in Heerlen achieved a 37–49% increase of the system capacity and peak reduction in the range 7.5–34%. DHC system operators can benefit from the STORM controller in the form of savings on operational costs and CO_2 emissions, and increased system capacity.

1. Introduction

District heating and cooling (DHC) networks are recognized as a prominent technology for decarbonizing the energy demand of the built environment. They enable integration of local renewable or residual sources of thermal energy in the energy system, by providing the connection between heating and cooling demands with these local sustainable sources. This integration needs to be further intensified in order to make the upcoming energy transition possible. Therefore, energy system integration will need to extend beyond the spatial dimension. The temporal (e.g. shifting production or consumption in time) and sectorial (e.g. integrating heating, cooling and electricity) dimensions need to be considered as well [1–3].

The way forward in this transition to further integration is digitalisation. Digitalisation is a broad concept referring to the transition of daily operation processes to a digital environment. The purpose of digitalisation is to make these processes perform better and work faster, and even enable processes that would otherwise be impossible. A central role in digitalisation is played by data: ranging from data gathering (e.g. sensoring, collection) over data management (e.g. communication, storage) to data processing (e.g. analysis, control) and post-processing (e.g. reporting, visualization). The more data that can be used and the easier it can be processed, the more value can be created from digitalisation. Therefore, digitalisation is gaining more and more momentum as a result of ongoing evolutions. On the one hand, digital hardware is becoming cheaper: e.g. sensors, storage, communication. On its turn, this has boosted progress in data processing, such as analysis tools and artificial intelligence. As a result, on the other hand, sophisticated algorithms are now available to handle the growing amounts of data in an automatic – and thus efficient – way.

In the field of DHC technology, digitalisation has many roles to play. One of them is supporting the integration of renewable and residual sources of thermal energy into the energy system. In order to better balance the supply and demand of energy in space, time

* Corresponding author. VITO NV, Boeretang 200, 2400 Mol, Belgium.
E-mail address: tijs.vanoevelen@vito.be (T. Van Oevelen).

https://doi.org/10.1016/j.energy.2020.117177
0360-5442/

Nomenclature

AWS	Amazon Web Services
B2B	business-to-business
BMS	building management system
CHP	combined heat and power
DH(C)	district heating (and cooling)
HP	heat pump
MAE	mean absolute error
QoS	quality of service
RME	rapeseed methyl ester
STORM	Self-organizing Thermal Operational Resource Management
TCL	thermostatically controlled load
TEMPO	TEMPerature Optimization for low temperature district heating across Europe
vDER	virtual Distributed Energy Resource

and across sectors, more intelligent DHC control systems are required. Such an intelligent controller for DHC systems has been developed in the Horizon 2020 STORM (Self-organizing Thermal Operational Resource Management) project [4–6]. The objective of the STORM controller is to move from separate control systems active in DHC systems to one integrated control system that manages the entire DHC system in an optimal way. At present, the STORM controller is capable of steering distributed flexible thermal loads, such as thermally massive buildings. This allows controlling the thermal energy production indirectly – in an otherwise pure load-following system – towards supply-side objectives. During the course of the STORM project, this approach has been tested in two demonstration sites: in Heerlen (The Netherlands) and Rottne (Sweden).

The scope of this article is to evaluate the performance of the integrated STORM controller technology during the tests in the two demonstration sites. The purpose is to quantify the impact of the STORM controller in these two DHC systems, based on newly gathered operational data and novel methodologies. The next section describes the STORM controller in more detail. The two demonstration sites and the testing approach are presented in the third section. The fourth section focuses on the main test results and the evaluation of the STORM controller performance. These results are discussed in the subsequent section, focusing on lessons learned for further testing and evaluation. The major results and findings are summarized in the last section.

2. The STORM controller

2.1. The STORM project

During the Horizon 2020 STORM project an advanced self-learning controller for district heating and cooling (DHC) networks has been developed and demonstrated. STORM controls the DHC network by means of demand-side management, in which the thermal capacity of consumer buildings is activated in an intelligent way to satisfy system objectives. The general applicability is guaranteed by the following features:

- Three control strategies are included in the controller. Dependent of the network, one or more of these strategies can be activated.
- The controller is an add-on to many existing DHC network controllers and SCADA systems.
- Self-learning control techniques allow a wide-range implementation by minimizing the required expert knowledge.

The STORM project has been coordinated by VITO, the Flemish Institute for Technological Research (Belgium). The project consortium further consists of partners from various European countries: NODA Intelligent Systems (Sweden), Mijnwater (The Netherlands), Växjö Energi - VEAB (Sweden), Stichting Zuyd Hogeschool (The Netherlands), and Euroheat & Power (Belgium).

2.2. State of the art before the project

The STORM controller has been built partially on established research performed in the field of cluster control of thermostatically controlled loads (TCLs) in electrical grids [7,8]. Examples of TCLs are water heaters, heat pumps, refrigerators, freezers, air conditioners or cooling machines. From a control perspective, these approaches are computationally challenging for large clusters due to the large number of control and state variables. This can be mitigated with distributed control methods. One example is distributed optimization where the control problem is decomposed and subsequently distributed over a cluster of local agents, jointly finding the optimal solution. This is however challenging practically due to the high requirements on communication and local computations. Therefore, an alternative approach is 'aggregate and dispatch' used in multi-agent systems [7,9]. An agent is capable of autonomous and independent actions, possibly based on interactions with other agents.

In DHC networks, model predictive control solutions [10] have been presented with excellent performance. However, these require expert-level customization of models, preventing scalable roll-out of intelligent control algorithms. Instead, automatic approaches relying on machine learning such as (batch) reinforcement learning [11,12] are considered more promising because of their model-free nature.

2.3. STORM controller functionality

The STORM controller was developed using innovative control algorithms and implemented in the existing control framework 'Smart Heat Grid' of NODA. The hierarchy of the STORM controller on NODA's platform is displayed in Fig. 1.

The STORM controller consists of four main components [6]. The *Forecaster* predicts the future heat demand profile and estimates the thermal flexibility available in the buildings. This information is used by the *Planner* to create an optimized heat load control plan, taking into account system constraints and the choice of control strategy. The following three control strategies for optimizing the heat demand profiles have been developed in the STORM project:

* **Peak shaving:**Reduction of heat load peaks above a prescribed peak heat production power threshold.
* **Market interaction:** Shifting heat loads according to electricity price signals in DHC systems with connections to the electricity grid (e.g. heat production with CHPs (combined heat and power) or heat pumps).
* **Cell balancing:** Increase of thermal energy exchange in grids providing heat and cold simultaneously, i.e. matching heat and cold demands.

Once a heat load control plan has been established, it is handed over to the *Tracker*. This will dispatch control signals towards individual building agents (*vDER*s) in order to try to follow the heat load control plan. Each vDER (virtual Distributed Energy Resource) interacts with the Tracker to negotiate how much it can contribute

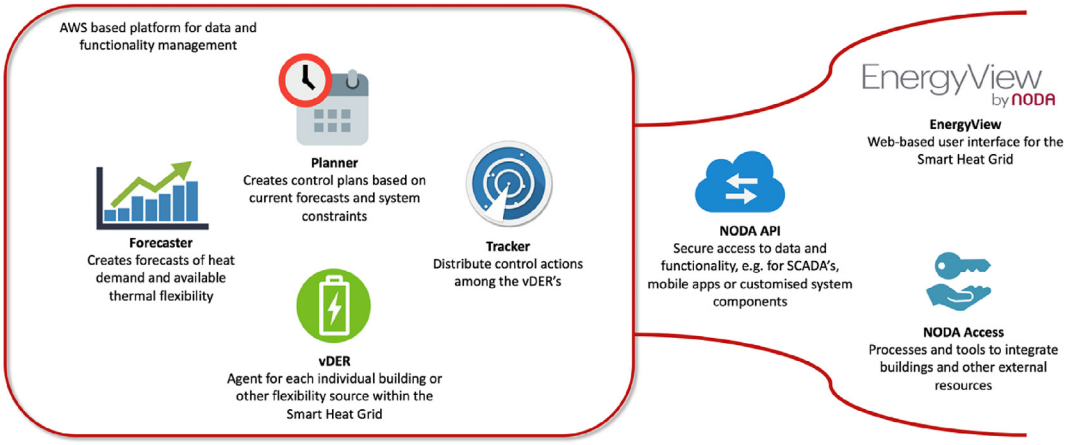

Fig. 1. Overview of the STORM ICT platform on NODA Smart Heat Grid.

to the heat load control plan, taking into account local constraints.

The ICT platform itself is based on the NODA Smart Heat Grid system. In the STORM project certain parts of this system have been further developed. The operational modules in the Forecaster, Planner and Tracker have been further developed using algorithms from STORM, while the other modules create the framework within which these algorithms operate (see Fig. 1). The platform also includes the capacity to interface with existing control and supervision systems of different kinds. The Access layer is a collection of processes and tools to integrate a large range of systems on the market. The STORM controller algorithms are deployed on the Amazon Web Services (AWS) cloud platform. After the two systems have been linked, the controller performs one or more control strategies based on predicted weather data making use of self-learning techniques to incorporate historical knowledge on the customer's operating system.

The performance of the STORM controller forecasting algorithms has been investigated separately in previous studies [13–15]. The novel contributions in the present manuscript are:

- a thorough description of the field tests in two demonstration sites: a low-temperature 4th generation network and a high-temperature 3rd generation network,
- a presentation of results obtained using three different control strategies during long-term demonstration in an operational environment, including more data and new analyses,
- a discussion of the results, elaborating on important lessons learned for the future.

3. Testing in demonstration sites

The STORM controller has been demonstrated in two existing networks: an innovative combined district heating and cooling network in Heerlen (The Netherlands), and a typical 3rd generation district heating network in Rottne (Sweden). The objective of implementing the STORM controller in these thermal networks has been to improve their efficiency by applying one or a combination of the aforementioned control strategies.

To connect a DHC system, the STORM controller is digitally linked to the control systems of the participating buildings as well as the production and distribution system. This link can be established using a) sensor override technology, b) gateway solutions, including building automation systems or c) full integration with an existing data management system. Furthermore, additional measurement devices, such as indoor sensors, can be added if required.

Both demonstration sites have been operational before the start of the STORM project and, apart from weather conditions, subject to constant changes during the course of the project, such as extra connections, changing thermal energy demand by customers, etc. It was therefore possible to obtain reference data for the evaluation of the STORM controller performance.

After a short description of the two demonstration sites, this section discusses how the STORM controller has been tested.

3.1. Heerlen, the Netherlands

The Mijnwater DHC system in Heerlen (see Fig. 2) is an ultra-low temperature 4th generation DHC grid connected to geothermal storage in the mine water reservoirs from abandoned coal mines. At the moment the total connected surface amounts to 177,291 m^2 with a wide variety of owners and buildings.

Originally, the Mijnwater system just aimed to gain geothermal energy from the mine reservoirs. Thereafter the system has evolved into a smart DHC grid, organized into clusters of buildings. This upgraded Mijnwater grid exchanges thermal energy, provides additional buffers and is fed by multiple renewable sources. Each building is connected by a decentralized thermal energy station, where the temperature demand is ensured by electric heat pumps. The heating of buildings generates cooling for data centers, stores and offices and vice versa. If an imbalance between heat and cold demands occurs in the cluster grid, the Mijnwater backbone transfers this residual thermal energy to other clusters or -in the end-to the mine water reservoirs. The thermal energy flows are optimized and controlled by a fully-automatic demand-driven process-control system.

Currently the Mijnwater-system consists of two hot and two cold bidirectional wells. These wells are connected by means of a pipeline system (the so-called backbone) through which, depending on the demand, mine water flows. The thermal energy exchange with the clusters takes place through heat exchangers in the cluster installations. At the moment there are three operational cluster grids (A, B and C) while the development of cluster D is in progress. The cluster grids are conventional two-pipe systems to which the buildings are connected.

Each building has its own thermal energy station with heat pumps. As a result, there is no central thermal energy plant in the grid. In fact the Mijnwater grid is connecting a cloud of heat pumps in buildings, which provide the right building conditions while also maintaining the temperature conditions in the grid (utilizing the backside of the heat pumps). The buildings connected to the

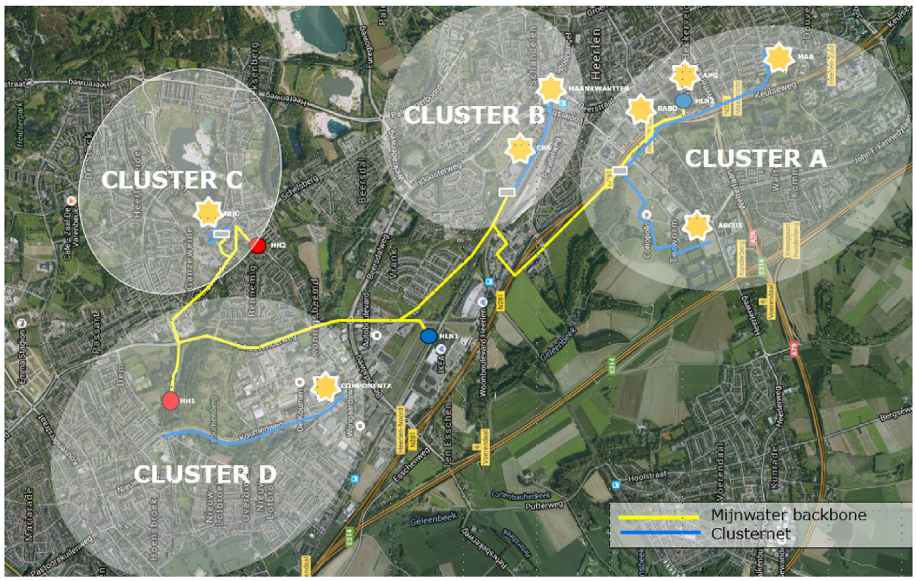

Fig. 2. The Mijnwater backbone with its connected clusters in Heerlen (The Netherlands).

Mijnwater system have appropriate insulation quality and/or heat emission systems such as floor heating, such that heating with heat pumps is feasible.

The Mijnwater system has no thermal production units as such. Basically the temperature difference of 12–15 °C between the warm and cold wells is used to provide the customers with the desired heat and cold. In fact, the heat pumps in the customer power stations provide clients with the desired temperatures and required return temperatures so that thermal energy flows can be exchanged as much as possible in the system.

The Mijnwater system is entirely powered by electrical energy (for heat pumps, wells, transportation pumps, etc.). Currently, there are three sites with photovoltaic installations (total capacity 208.7 kWp) which provide sufficient renewable electricity for all the thermal energy stations of the connected customers. The electricity generated by the photovoltaic panels goes directly to either the heat pumps or the customer. The remainder is supplied to the public electricity grid.

The ratio between heat and cold delivery in the Mijnwater system is almost equal over the years 2014–2018. By linking as many cold receivers as possible (such as data centers and supermarkets) to heat users, Mijnwater tries to optimize the balance between cold and heat supply as much as possible in order to reach a high level of thermal energy exchange between customers. While the highest degree of exchange between customers is reached in cluster A (64%), at the moment the year-round average exchange level is 44% for the entire system.

There are clear demand peaks during the winter months for heating, as well as for cooling during the summer months. The intermediate periods (such as spring and autumn) will be the best for reaching the highest exchange rates. The objective of the STORM controller in the Mijnwater system is to improve the exchange rate, i.e. using the cell balancing control strategy to balance heat and cold demands, in order to reduce the net heat/cold demand from the higher-up hierarchical level: the backbone from a cluster perspective and the mine water wells from the backbone perspective. The aim is to reduce electricity costs, increase system capacity and thus make the operation more profitable.

All STORM controller hardware had to be installed in Heerlen during the project and above that all connections had to be programmed between the PRIVA operating system and the NODA system to allow the STORM controller to be active. During the course of the project there has been testing performed in the Heerlen grid to evaluate different ways to influence the control points. This testing has been done both during heating periods (i.e. winter) and cooling periods (i.e. summer). Unlike the Rottne demonstration site in Sweden, the Heerlen grid has control points both in buildings and at so called cluster stations. These cluster stations are located between each cluster and the backbone, and basically act as a substation for that specific cluster. Based on these test periods it has been concluded that it is most beneficial for the Heerlen grid to focus on the cluster stations rather than the building as active control points. The reason for this is the special set-up in the buildings using heat pumps and other components controlled by the BMS (Building Management System) within the building heating system that are not possible to control externally (e.g. the STORM controller). Furthermore, the intermittent behaviour of the building interactions with the cluster makes it hard to generate robust models of their operational pattern. On the other hand, the cluster stations display more predictable patterns related to time of day and weather. This makes them easier to handle within the STORM framework of algorithms.

The way that the STORM controller interacts with the control points has also been evaluated during the course of the project. The original plan was to interact with the control points in Heerlen in a way that is similar to the generic sensor override concept utilised in the Swedish demonstration site in Rottne which uses manipulation of the outdoor temperature signal for the building controller. However, since the buildings in Heerlen display a behavioural pattern very loosely correlated to changes in outdoor temperature, this makes such a control signal highly inefficient. During the testing it was also concluded that, even in those cases where such a correlation was detected, it was still inefficient due to delays in the internal building heating and cooling system.

In order to address this issue, an alternative way to interact with control points was devised and consequently integrated into the STORM controller by Mijnwater and NODA. This control scheme is based on the control signal setting restrictions on maximum flow levels, which then gives the STORM controller a way to impose control behaviour at the control point. The existing system is still

free to operate below those levels of restrictions, which means that the basic control behaviour of the system was not affected. This means that the system was more straightforward to implement and consequently more robust in its operational behaviour. The STORM controller was then able to influence the behaviour in real-time using a control signal that in turn dynamically sets this maximum flow level. This, in turn, leads to a larger temperature difference between the supply and return temperature on the backbone side of the cluster station and in the end leads to a lower cluster supply temperature in heating mode and a higher cluster supply temperature in cooling mode because the heat pumps in the buildings are forced to operate on more optimal levels.

The intermittent behaviour of the control points made the construction of the models more difficult, and not achievable using the normal STORM set of algorithms. In particular, changes in system behaviour has fragmented the data into shorter time periods ill-suited for time series forecasting, so the current models focus on recurring behaviours rather than on system dynamics.

The focus of the testing has been on cluster A and cluster B in the Heerlen grid. For this purpose, three Trackers are deployed for the clusters themselves and then also the set of controllable buildings within the two clusters. However, as per the conclusions in the previous section, the primary focus of the tests have been on the cluster A and cluster B control points.

3.2. Rottne, Sweden

Rottne is a small city located about 19 km north of Växjö city and has 2427 inhabitants. The total land surface area in Rottne is 1,930,000 m^2 , of which 201,137 m^2 is built area, visualized in Fig. 3. Rottne consists mostly of built-up area, but also some green areas. The district heating system in Rottne, operated by VEAB, is a traditional 3rd generation system with a high temperature distribution network. The production plant in Rottne was put into commission in September 1998 and extended in 2004. In 2012, the production plant in Rottne became completely fossil free with the introduction of biofuels.

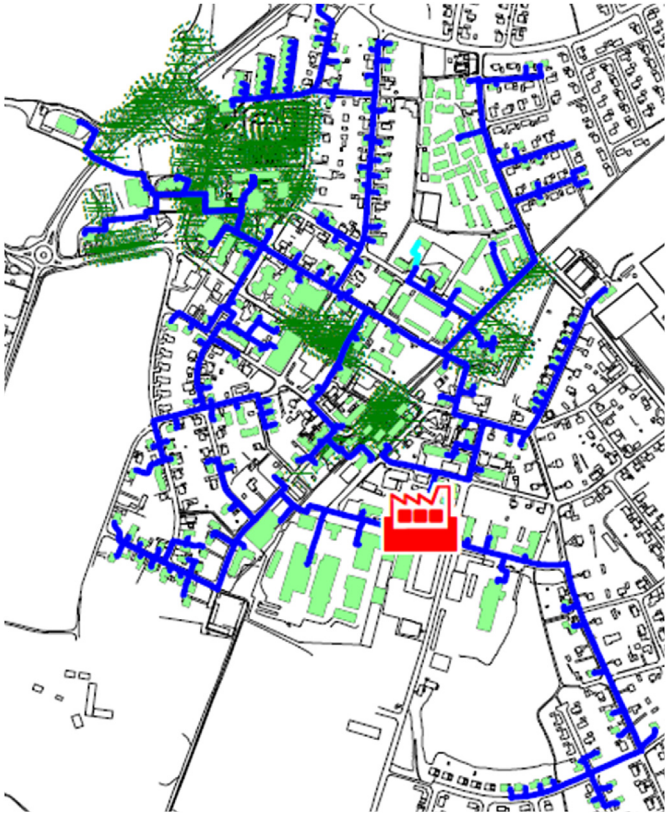

Fig. 3. Overview of the DH distribution system in Rottne (Sweden).

The base load heat production in Rottne is based on two biomass boilers of 1.5 MW and 1 MW maximal capacity. The biomass originates from the surrounding forests and wood industry, consisting of wood chips, branches and peaks. An additional boiler with a capacity of 3 MW, running on rapeseed methyl ester (RME), is used for peak loads and backup. The production plant in Rottne generates approximately 12.8 GWh a year.

The RME peak load boiler is activated when the heat demand surpasses the capacity of the two biomass boilers. This happens approximately when the outdoor temperatures is about −1 °C. During mid-winter it is common that the RME boiler kicks in due to switching outdoor temperatures. In order to reduce the amount of heat produced with the costly RME fuel, the heat demand needs to be controlled to avoid demand peaks.

There are 180 connections to the grid and the same number of substations, owned by the customers. 70% of all connections are for villas, small-houses. However, 53 connections are designed for other buildings such as multi-family houses, industries, public buildings and offices. Here the same connection can be used for several buildings, since building owners may have their own network to their buildings but obtain the heat from one substation. This gives VEAB 71 more buildings using district heat in Rottne. The total consumed heat per year is about 10.4 GWh.

The district heating network in Rottne implemented the STORM controller in 2014. The targeted customer segment for the STORM controller consists of large building owners (B2B), such as non-residential building owners and housing cooperatives. The STORM system is installed in nine of the largest customer substations in the network of Rottne, representing 34% of the total heat consumption in Rottne and a heat load of 1.7 MW during an outdoor temperature of −14 °C, when the total heat load of the entire grid would be 4.4 MW. The control hardware and a basic demand-side management system were already present before the start of the project. In a first step, the STORM Forecaster was implemented. Manual control actions were based on these forecasts. The second step was to implement the STORM Planner, using the peak shaving control strategy. This strategy has been active since 2018-02-14 in the latest version. Thereafter it was active for testing during the later stages of the heating season. Some settings were updated for the 2018/19 heating season and the Planner was fully active throughout the full heating season of 2018/19.

3.3. Testing approach

The three control strategies developed during the STORM project have been tested according to the test scheme in Table 1. Due to the differences between the two demonstration sites, not all strategies have been tested on each location.

Testing the STORM controller technology in operational district

Table 1
Test scheme for the three control strategies in the demonstration sites.

Control strategy	Peak shaving	Market interaction	Cell balancing
Demo-site Heerlen	△	–	□
Demo-site Rottne	■	□	–

■ Fully automated testing.
□ Partially manual testing.
△ Performance evaluated as side effect.
– Strategy not tested.

heating and cooling systems means that quality of service (QoS) needs to be ensured. Therefore, both in Rottne and in Heerlen, only a limited control over the heat/cold demand could be achieved. For example, very conservative comfort constraint margins have been used during testing in order to avoid complaints from building occupants. This means that the availability flexibility in the buildings may not have been fully exploited. Furthermore, in the Rottne demo site, only a small part of the buildings within the Rottne DH system is connected to the STORM controller, allowing only about 34% of the heat load to be controlled. In the Heerlen demo site, although all buildings as well as all system components were linked via the building management system (BMS) to the STORM controller, it appeared to be impossible to fundamentally intervene in the buildings' climate control. This was due to the lack of permissions by the building owners.

The data of the STORM controller tests in both demonstration sites was accessed through the NODA EnergyView dashboard. In the Mijnwater system, also data from the PRIVA system was available.

4. Test results and performance evaluation

The test results of the STORM controller in the demonstration sites are presented and discussed in this section. Each control strategy has been tested separately in the applicable demonstration site(s), see Table 1. Consequently, the performance of each control strategy is evaluated individually.

4.1. Peak shaving

The peak shaving control strategy aims to reduce the amount of heat produced by peak units with higher fuel costs. It assumes that base load units are activated first until their full capacity, and that peak units deliver the heat above this power level. In other words, the peak shaving control strategy tries to shift heat loads above the base load capacity threshold towards times with lower heat load. The peak shaving control feature has been tested in the Rottne demonstration site, but the cell balancing tests in Heerlen effectively also brought about peak shaving behaviour (see Section 4.3).

4.1.1. Methodology

The methodology for evaluating the peak shaving control strategy is based on the heat load-duration curve of the test period, in comparison with the behaviour during a historical reference period without active STORM controller. For this, hourly data is used for the following variables:

- total heat load of the network, measured at the heat production plant (Ptot)
- combined heat load of the controllable buildings, aggregated from heat meters of the controllable buildings. Due to data logging and communication problems over the course of the project, only a subset of four controllable buildings is used in the analysis (Pset).
- outdoor temperature, calculated as the mean of the outdoor temperatures measured at each of the controllable buildings (Tout)
- controller activity, aggregated from control signals towards all controllable buildings

The data used in this evaluation has been recorded in the Rottne demonstration site from July 1st, 2015 to January 31st, 2019. Times where any of the variables is missing are excluded from the analysis, as well as the data from January 23rd, 2018 because of a mechanical incident with one of the wood chip boilers that day. The resulting pre-processed data set is visualized in Fig. 4.

Because heat loads are affected by other variables than only the presence or absence of STORM controller activity, a direct comparison between test period and reference period is not possible. Instead, the test results are evaluated indirectly in the present study, by comparing with a reference model derived from historical data. The reference heat load model is built as follows:

1. From the historical data set, all hours affected by control actions are removed. It is assumed that all hours in the range between 1 h before and 23 h after an hour with control activity are potentially affected. Note that the STORM controller was developed on top of an existing operational ICT platform, so that no recent period without control activity was available.
2. The emerging reference data set is analyzed statistically as a function of two prediction variables: outdoor temperature and

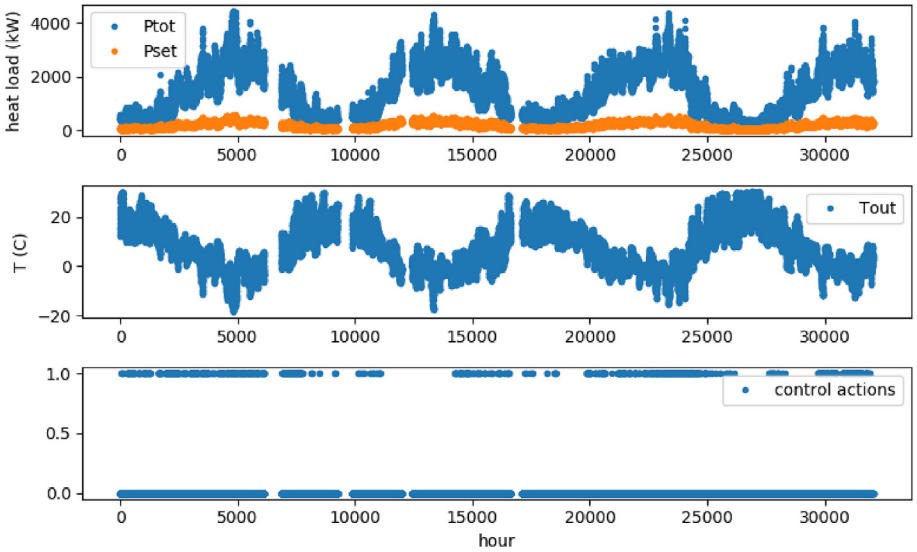

Fig. 4. Pre-processed data from Rottne demonstration site (2015-07-01 to 2019-01-31) used for evaluation of STORM controller peak shaving performance. Ptot: total heat load; Pset: summed heat load of 4 controllable buildings; Tout: outdoor temperature.

hour of the day. This is done in the form of a two-dimensional look-up table with resolution 1 °C × 1 h (see Fig. 5). It is assumed that these two variables are sufficient to model the variance in the hourly heat load. This choice involves a trade-off between accuracy, complexity, data richness and representativeness. More advanced heat load modeling is possible, but would lack transparency.

3. The reference heat load is modeled by the mean value of the reference data corresponding to the outdoor temperature and hour of the day.

The total heat load of the network corresponds to the sum of the three heat production units in the Rottne district heating system: two wood chip boilers and one bio-oil boiler. Unfortunately, this value can't be broken down into individual heat production rates due to the absence of individual meters. In order to resolve this for the evaluation of peak shaving, it is assumed that the RME peak boiler is always activated last, and exactly delivers the heat load exceeding 2.5 MW. This threshold value is derived from the nominal capacities of the two wood chip base load boilers, and has been validated by manual observation. The same threshold value is used in the controller optimization objective function, so it is an appropriate choice. Furthermore, it could be argued that the actual heat load shares of the heat production units are irrelevant for the present evaluation of the STORM controller, because it focuses only on demand-side management and not on the supply side.

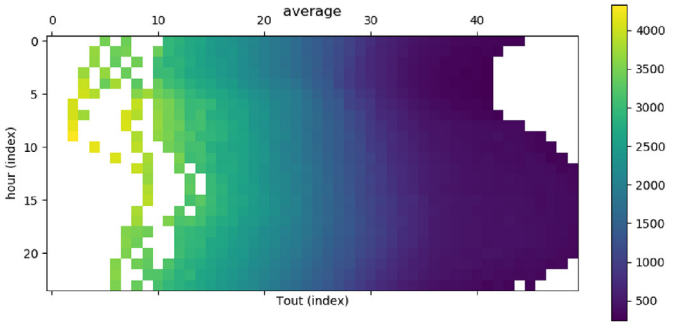

(a)

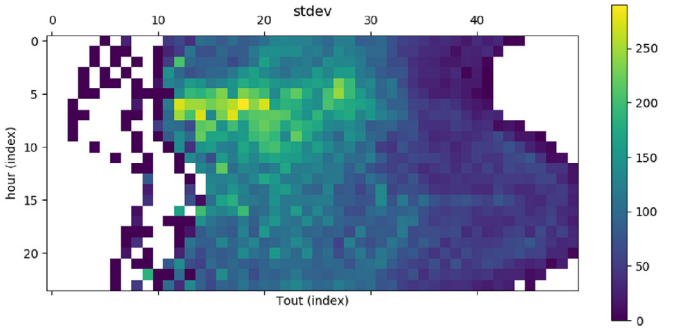

(b)

Fig. 5. Reference heat load statistics: (a) mean and (b) standard deviation, as a function of outdoor temperature and hour of the day. Data from Rottne demonstration site (2015-07-01 to 2019-01-31), excluding times affected by STORM controller activity. Hour index corresponds to hour of the day; Tout index corresponds to outdoor temperature with 20 °C offset. White/empty cells represent conditions that have not been encountered.

Heat load-duration curves are typically calculated from heat load profiles by sorting values in descending order. However, in this case the reference heat load profile results from a probabilistic model. The classical calculation method for the heat load-duration curve is therefore unsuitable, because it would be biased towards the mean. This is especially important for peak shaving evaluation because the classical method underestimates the heat load peaks. This bias is exemplified by the hypothetical case where the heat load is modeled by a simple normal distribution with constant mean and standard deviation, i.e. independent of outdoor temperature and hour of the day. The classical method would yield a constant load-duration curve at the mean value for all durations. But in reality, the S-like shape typical for load-duration curves would emerge, as could be shown by for example Monte Carlo simulations.

What is needed instead, is a curve depicting heat load with respect to *expected* duration, in a statistical sense. The expected duration that the load is higher than (or equal to) a specific heat load level l^* is calculated as follows (see Appendix A for derivation, meaning of symbols and explanation):

$$E[d_L](l^*) = \int_{t \in T} (1 - F_{L(t)}(l^*))\mathrm{d}t. \quad (1)$$

Furthermore, it is assumed in the calculations that the heat load at time t is independent of other times and is normally distributed. Then, the corresponding cumulative distribution function $F_{L(t)}$ of the random heat load profile $L(t)$ is available in standard statistical software packages.

This statistical approach for calculating the load-duration curve has been validated with the historical reference data and compared with the classical approach. It is concluded that the proposed approach predicts the actual load-duration curve with higher accuracy than the classical approach. The mean absolute error (MAE) of the duration of the total heat load is only 7.9h (relative to a dataset of about 18,000h) when calculated using the statistical approach, representing a very high accuracy (see Fig. 6).

4.1.2. Results

The evaluation period of the peak shaving control strategy in the Rottne demonstration site ranges from March 2018 to January 2019.

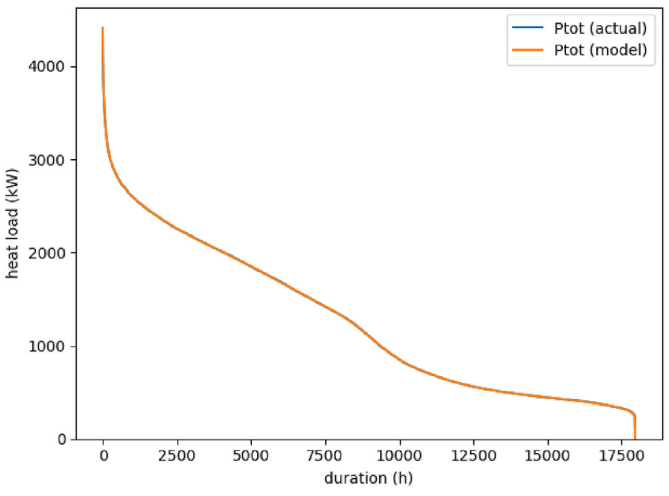

Fig. 6. Total heat load-duration curve of the historical reference data and validation of the reference model (MAE: 7.9h). Note: the actual and modeled load-duration curves are practically indistinguishable in this graph.

An individual evaluation of each month as well as an evaluation of the overall period are performed. The warmer months May to October 2018 are not considered, because the heat loads are too low to trigger the peak heat boiler.

First, the detailed results of a typical winter month, i.e. December 2018, are presented and discussed. Fig. 7 depicts the total heat load-duration curve for December 2018 in comparison with the reference model. The change in heat load duration from the reference (without STORM) to the evaluation period (with STORM) is separately shown in Fig. 8. It can be seen from both figures that the heat production is partly shifted from loads above 2.5 MW to loads below 2.5 MW. This means that the expensive peak heat production has decreased, and has been replaced by cheaper baseload heat production.

The load-duration curve of the subset of controllable heat load in December 2018 is shown in Fig. 9. Again, the shift of heat load from high to low loads can be clearly seen. There are two remarkable qualitative differences with the results on the total network level. First, the region with reduced 'peak' heat load duration is more spread out. The reason for this is that there is no specific peak

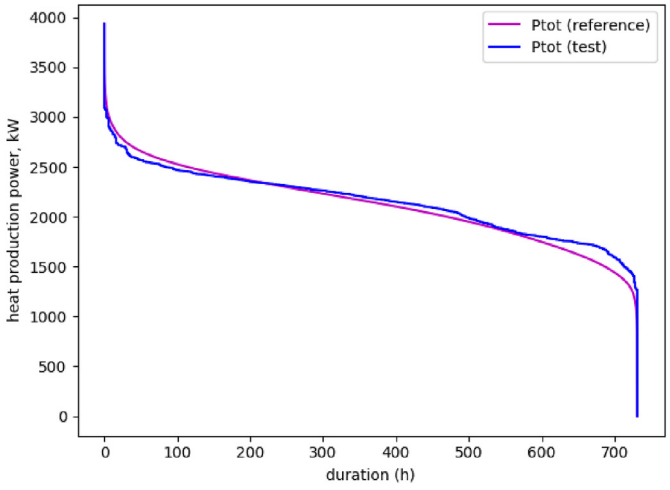

Fig. 7. Load-duration curve of total network heat load (Ptot) in Rottne demonstration site during December 2018 (blue), compared to reference model (purple). (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

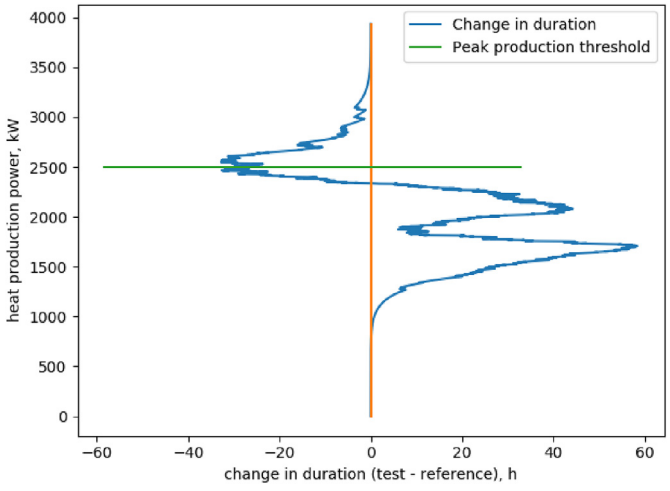

Fig. 8. Change in the load-duration curve of total network heat load in Rottne demonstration site from the reference to December 2018.

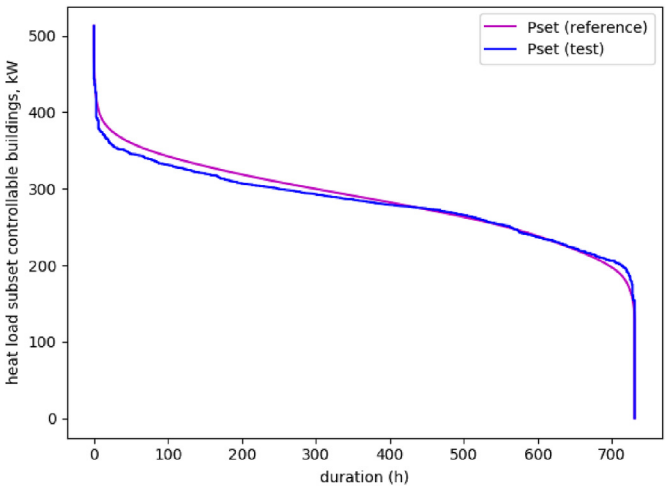

Fig. 9. Load-duration curve of a subset of the controllable heat load (Pset) in Rottne demonstration site during December 2018 (blue), compared to reference model (purple). (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

threshold value related to the controllable heat load. The peak shaving objective is defined on the total network level, and not on the level of the controllable buildings. Second, the STORM controller has reduced the average heat load of the controllable buildings. Despite this, the average total network heat load has increased. This is attributed to the unobservable behaviour of the remaining buildings in the district heating grid.

The quantitative results for all evaluated months are summarized in Fig. 10 and Table 2. The following observations are made:

- The controllable heat load was lower than the reference in all months except April. Overall, the controllable heat load decreased by 12.7 MWh.
- The total heat load was higher than the reference in all months except November. Overall, the total heat load increased by 69.1 MWh, as result of an increase of the uncontrollable heat demand of 81.8 MWh. Unfortunately, this interferes with the peak shaving testing and disturbs the evaluation.
- Despite the overall heat load increase, the overall peak heat production was reduced by 7.4 MWh (−3.1%) compared to the reference period without STORM. The peak heat production was lower in all months except January, when the peak heat production increased inexplicably by 12.4 MWh (together with an overall heat demand increase of 48.4 MWh). For the other test months, absolute peak load reductions up to 7.9 MWh have been obtained on a monthly basis. Together, the reduction in peak heat production is 19.8 MWh (12.7%) during these months, which paints a brighter picture than the currently reported result.

To summarize, it appears that the STORM controller peak shaving control strategy performs in the way it should in Rottne. The peak heat production is reduced in general, which is the desired result. The overall heat production seems to have increased, despite the reduction in the heat load of the controllable buildings.

4.2. Market interaction

The market interaction control strategy aims to maximize income from electricity production, e.g. when heat is produced by combined heat and power (CHP), or minimize expenses for

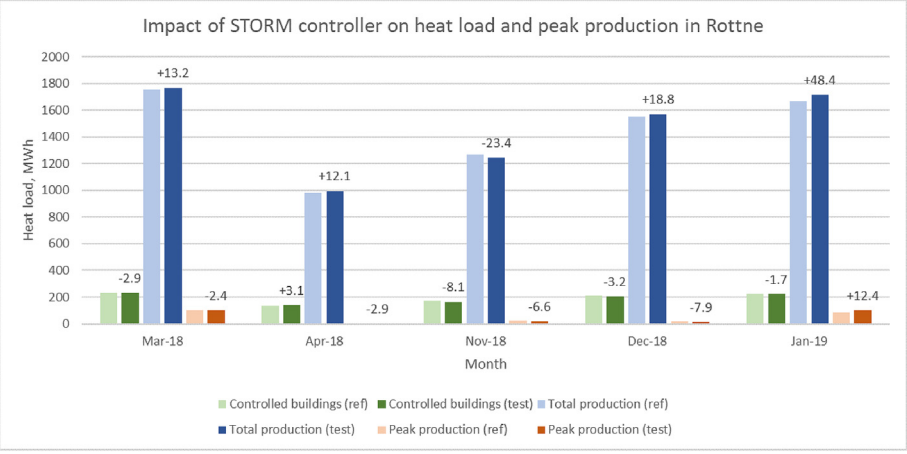

Fig. 10. Quantitative results of the peak shaving control strategy in Rottne on monthly basis.

Table 2
Overall quantitative results of the peak shaving control strategy in Rottne (values in MWh). E_{set}: heat consumed by subset of controllable buildings; E_{tot}: total heat produced; E_{peak}: peak heat produced.

Period	E_{set}	E_{tot}	E_{peak}
Reference	974.8	7220.2	241.5
Test	962.1	7289.3	234.1
Change	−12.7	+69.1	−7.4 (−3.1%)

electricity consumption, e.g. when heat is produced by heat pumps (HP). The business case for load shifting with this control strategy is driven by time-dependent electricity prices. In that context, load shifting creates value by shifting the heat load towards times with favourable electricity prices (considered as input from an external source), i.e. high for production (CHP), low for consumption (HP). The testing during the STORM project is focused towards optimizing CHP operation. Analysis out of other projects, looking at the average volatility during 24-h periods on the spot-price market (both intraday and day-ahead), show that the intraday market provides most financial benefit, simply due to the higher volatility.

The primary difference between the peak load control strategy and the market interaction control strategy, is that the latter uses both charging and discharging. Using only discharging is not a problem in a grid, because the impact for the customer is to reduce heat consumption. However, if charging is also used, then the STORM controller has to be smart enough to not cause increased heat consumption. In certain circumstances this can be acceptable, e.g. in a situation where the heat source can be shifted from fossil to renewable and where the customer pays by a fixed sum. However, in most practical situations there needs to be a way to ensure that the customers balance the heat consumption over the period of time, even though the demand profile is changed. In the STORM controller this is handled by the zero-sum functionality. This is a setting that can be between 1 and 0, where 1 requires a full zero-sum compliance of demand while a number lower than 1 relaxes this requirement.

There are no CHPs in either of the demonstration sites, so there is no possibility to test the market interaction control strategy in practice. However, the heat production plants in Rottne have been emulated as being CHPs. The algorithms used for market interaction are the same as used for the other applications in the STORM Planner and Tracker, and these have been extensively tested throughout the project. Therefore, for the purpose of evaluating the feasibility of the market interaction control strategy, the potential for charging, in addition to discharging, has been specifically tested in the Rottne grid during the STORM project.

4.2.1. Methodology

The primary focus of the market interaction evaluation is the heat load shifting, while also considering the system temperatures. The variables used in the evaluation of market interaction in the Rottne demonstration site are:

- total instantaneous heat load of the network
- outdoor temperature
- control signals (offset outdoor temperature)
- supply temperature on the secondary side

The evaluation focuses on the impact of discharge and charge in relation to each other, and specifically on the charging since this hasn't been part of the more extensive testing of the peak shaving control strategy.

4.2.2. Results

An example of the possibility of charging a group of buildings in Rottne is shown in Fig. 11. It can be clearly seen that the cluster of buildings reacts to the control signals, and consequently increases the heat load. Table 3 shows the average results of a number of 2-h control actions for testing charging in the Rottne grid. All tests were done in the same range of outdoor temperatures. The results are in line with similar tests for discharging, and demonstrate the feasibility of the market interaction control strategy.

During the testing of charging, it was also tested to discharge in order to restore the zero-sum as described in the previous section. In that case the results show zero-sum achieved to a level of 5.8%. These results are in line with results shown in commercial applications of the STORM controller where the market interaction control strategy has been active for longer periods of time.

4.3. Cell balancing

The cell balancing control strategy aims to maximize the thermal energy exchange in a network with combined heat and cold supply, in order to reduce the need for external supply of thermal energy. This is done by shifting both heat and cold demands in time, such that they are can cover each other as much as possible. As a result, the network needs to rely less on external sources of heat and cold. The objective in Heerlen has been to balance heat and cold demands within each cell, as well as between cells, as much as

Fig. 11. Example of the reaction of the controllable heat load in kW (red, right y-axis) in the Rottne demonstration site to a coordinated charging control action, i.e. negative offset outdoor temperature in °C (green/pink, left y-axis). (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

Table 3
Average results of 2-h charging tests in Rottne. Heat load right before the beginning of the test and during the test are compared.

Offset (°C)	Reference (kW)	Test (kW)	Difference (kW)	Difference (%)
−5	805	1176	371	46
−6	700	1080	380	54
−8	582	1141	559	96

possible.

The STORM controller implements cell balancing by means of hydraulic peak shaving for different clusters. The Trackers strive to minimize the sum of squares of the flow, thus implementing hydraulic peak shaving on cell level, while the vDERs for each control point strive to uphold safety constraints while allocating flexibility in an optimal way with respect to their individual models of the system. The resulting cell balancing behaviour is effectively accompanied by peak shaving performance.

As discussed before in Section 2, the STORM controller implementation in Heerlen is a special case. On the one hand, the buildings connected to the Mijnwater grid were difficult to influence in a reliable and predictable way. The focus of the tests was therefore on the cluster stations, which connect the backbone network to the sub-networks of the different clusters. On the other hand, the outdoor temperature sensor override method for influencing the behaviour of the energy stations could not be applied. Instead, the energy stations were controlled by setting maximum flow levels.

4.3.1. Methodology

The operational behaviour of the fully deployed STORM controller is evaluated based on the available flow and temperature measurements in the cluster stations:

- supply temperature on the backbone side
- return temperature on the backbone side
- flow on the backbone side
- flow restriction level setpoint on the backbone side
- control signal intensity

These values are evaluated in relation to the control signals sent by the STORM controller in order to quantify the impact and benefit for the grid.

4.3.2. Results

The results obtained in cluster A will be focused on in this section, as well as some highlights from cluster B. Fig. 12 shows a period of time when the STORM controller was fully active in cluster A. The green line shows flow over the cluster station, while the red line shows the control actions.

It should be noted that although the STORM controller was active throughout the period shown in Fig. 12, the settings of the safety trigger on the cluster return temperature where different. During the first part of the period the setting was 12 °C and this was later lowered to 8 °C, during February 21st. This lowering gave the STORM controller the ability to impact the operational behaviour in practice, which was not possible during the period of the higher set-point. The higher setpoint effectively cancelled out all impact of the STORM controller, which makes it possible to use such periods as references for comparison. During the particular period shown above the weather conditions were also quite similar throughout the period, with an average outdoor temperature of 8.7 °C during the 12 °C period and 8.5 °C during the 8 °C period.

The primary performance indicators available in the Heerlen case are the temperature levels and the flow levels, and how they are affected by the STORM controller. The general idea in the Heerlen case is to lower the supply temperature in the cluster networks, which results in lower return temperature in the backbone. In this way, the temperature difference between the supply and return in the backbone increases, resulting in a lower flow rate

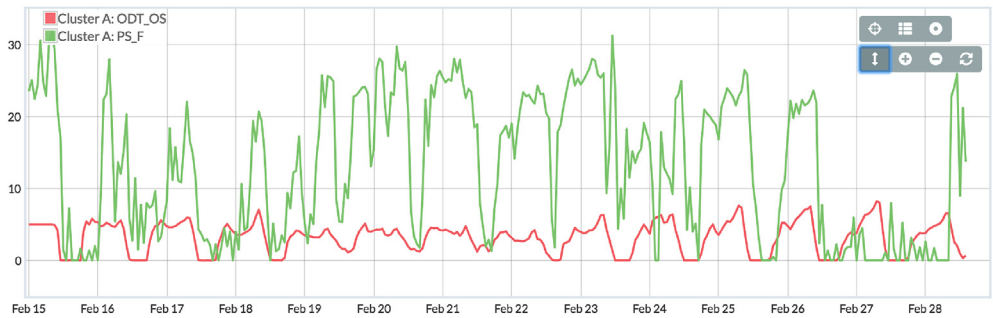

Fig. 12. Behaviour of the STORM controller in cluster A in the Heerlen demonstration site. Before February 21st, the cluster return temperature limit is set on 12 °C (reference); after February 21st, this setting is at 8 °C (test). ODT_OS: controller signal; PS_F: backbone flow rate.

Table 4
Shows the results of the testing with the active STORM controller compared to the reference period. The temperature and flow values are mean values over each period, measured on the backbone side of cluster A.

Variable	Reference	Test	Δ/%
Supply temperature, °C	20.2	21.3	+1.0
Return temperature, °C	13.9	11.8	−2.1
Δ*T*, °C	6.3	9.5	+3.1
Flow rate, m^3/h	15.0	13.9	−7.5%[a]
Capacity, MJ/m^3	26.4	39.4	+49%[b]

[a] Peak shaving.
[b] Cell balancing.

in the backbone. In this way, the backbone can serve more clusters. However, obviously the operational characteristics of the backbone influence the capacity within the individual clusters.

Table 4: Cell balancing results of testing the STORM controller in the Heerlen demonstration site (cluster A). All variables relate to the backbone side of the cluster station.

The results indicate an increase of the Δ*T* on the backbone side of 3.1 °C for cluster A — 2.6 °C for cluster B. This increase in temperature difference results in a potential backbone capacity improvement ranging from 37% (cluster B) to 49% (cluster A). The influence of the two clusters combined leads to a capacity improvement of 42% on system level.

At the same time the mean value of the flows from the backbone to the cluster installations decreased with 7.5% for cluster A — and 34% for cluster B. For the Mijnwater system this peak shaving potential is extremely important as system capacity depends to a large extent on the pump capacity of the wells. Combining the tested clusters, the peak shaving potential amounts to 17%.

Once influencing the flows to the thermal energy stations of individual buildings is up and running, the same benefits on capacity improvement and flow reduction may be expected, although on a smaller scale. Since the number of buildings in a cluster is greater than the number of clusters in the Mijnwater system, cell balancing will play a major role in the smart exchange of thermal energy flows within the cluster. Naturally, the influence of cell balancing and peak shaving of individual connections within a cluster will only have a marginal impact on the total system, although a larger Δ*T* will always be beneficial for the entire network.

5. Discussion

The presented results show a promising potential for the performance of the STORM controller in the demonstration sites that were part of the project. Nevertheless, some challenges remain regarding the results as well as the evaluation approach.

Limited heat load controllability. The results obtained in Rottne are due to the activation of nine buildings, representing around 34% of the heat load on average. Consequently, the impact of the STORM controller would be higher if the heat demand of all buildings could be controlled.

In Heerlen, the STORM controller could not control the buildings during the project because of lack of authorization from building owners. Therefore, the demand on the thermal grid is managed by controlling the network flow through the cluster and building thermal energy stations. Also in this case, the expected STORM impact would be higher if the building flexibility could be additionally activated.

Unaccounted heat load variability. Testing and evaluating a technology like the STORM controller in a realistic environment is challenging. It is not possible to make a test setup in a lab, because of the size and cost of DHC systems. Furthermore, the influences from weather and consumer behaviour would be difficult to simulate realistically. The only possible way is to test the controller live in an operational environment. With that comes the responsibility to maintain delivering Quality of Service. This requirement has two drawbacks.

On the one hand, the margins for testing are tight, meaning that the controller has to limit itself and behave very conservatively to account for uncertainties. This means that the actual potential for influencing the building heat demands could have been higher.

On the other hand, the controller performance can't be evaluated using a proper scenario analysis, i.e. testing with and without STORM controller under identical circumstances. This would require full control over all parameters influencing the DHC system performance, which is not feasible because it involves among others influences due to the weather and consumer behaviour. As a result of this challenge, test results can only be compared against assumed reference behaviour. In the presented results, the STORM controller was evaluated with respect to the best possible estimate for the reference performance. But still, unexpected behaviour in the test period itself can't be ruled out, or even verified. An example of this is the total heat demand increase with respect to the reference period in Rottne during peak shaving tests, especially noticeable in January 2019. Because the heat load of the controlled buildings was reduced, this is most likely due the large group of uncontrolled buildings. Furthermore, according to the district heating system operation VEAB, a large new building was connected to the DH grid in this period. Such events are hard to account for in an operational system, but the impact on the evaluated performance is high. If January 2019 is excluded from the analysis, then the peak heat production would be reduced by 19.8 MWh (12.7%), in contrast to the currently reported result of 7.4 MWh (3.1%). And even in the other months, the peak shaving performance was negatively affected by the behaviour of the uncontrolled buildings.

Limited testing period. One, if not the only, solution for the previous problem is to acquire more test data by testing for a longer time. This would allow better insights in the natural variations and trends in the heat demand pattern. Outlier behaviour could also be more easily detected and filtered out. Unfortunately, more data only becomes available at a pace dictated by the monitored DHC systems, which has important seasonal and yearly variations. Patience is needed to obtain more results to analyze.

Continued testing of the STORM controller. In order to address the aforementioned issues, further testing of the STORM controller will be done. Therefore, testing is continuing in the STORM-project demonstration sites. Furthermore, the STORM controller has in the meantime been installed in several more DHC networks in demonstration and commercial projects. Results from this continued testing will allow to get clearer and more general results, as well as provide input for further improving the STORM controller technology.

6. Conclusions

During the STORM project, an intelligent controller for DHC systems was developed. The controller is capable of shifting heat and cold demands in a DHC network in time. At present, this demand-side management technology can be utilised for three different control strategies: peak shaving, market interaction and cell balancing. They have been tested in two demonstration sites during the project: in Heerlen (The Netherlands) and Rottne (Sweden).

Peak shaving tests were performed in Rottne, with the objective to minimize the consumption of the expensive fuel of the peak heat

production unit. During the testing period, the controllable heat demand was reduced by 12.7 MWh. As a result, the peak heat production was reduced by 7.4 MWh, representing 3.1% of the reference peak heat production in that period. Due to inexplicable behaviour of the uncontrollable heat demand – especially in January 2019 – the results have been masked by external adverse influences. It is therefore concluded that the peak shaving potential of the STORM controller could be a lot higher than the present results show. More testing over a longer period is needed to reveal the real potential.

In Rottne, also the market interaction control strategy was tested. Its objective was to shift the heat demand to align with high electricity spot prices, such that electricity revenues from CHP heat production could be maximized. The tests demonstrated a proof of principle, because the district heating grid in Rottne doesn't have a CHP heat production unit. The tests focused on two important aspects: charging and consumption balance. Charging allows to concentrate the heat demand during hours with high electricity prices. The results show that the heat load could be increased by up to 559 kW (96%) during charging. The heat consumption balance is important to prevent increasing customers net heat demand, which would increase their heating costs. The tests have demonstrated that the net heat demand increase during a charge-discharge cycle is at most 5.8%, with typically lower values encountered in related commercial projects.

Tests in the Heerlen demonstration site, which features hydraulically disconnected cluster networks together with a backbone network, primarily focused on cell balancing. This control strategy aims at balancing heat and cold demands to promote thermal energy exchange within clusters. This reduces the dependence on external heat/cold sources (from the backbone network), representing a system capacity increase. The tests were performed in two clusters: the capacity increase was 49% in cluster A and 37% in cluster B. At the same time, the flow rate decreased by 7.5% (cluster A) and 34% (cluster B), representing the peak shaving potential.

The impact of the STORM controller in DHC networks is case-specific and depends on the chosen control strategy. Overall, the obtained results are very satisfactory. They present the following benefits to DHC system operators:

- Reduction in operational costs: e.g. by replacing consumption of expensive fuels by cheaper fuels, or maximizing CHP revenues
- Reduction in CO_2 emissions: e.g. by replacing CO_2-intensive heat production by more sustainable heat sources
- Increase in system capacity: e.g. by shaving peaks and maximizing thermal energy exchange

Several challenges with respect to the evaluation approach and results have been discussed. These are related to the limited controllability of the heat load in the demonstration sites, the variability of the heat demand, and the test period duration. These challenges will be tackled in the continued testing and monitoring of the STORM controller in existing and new projects.

Declaration of competing interests

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.

Acknowledgments

This work is supported by the European Union's Horizon 2020 research and innovation programme under grant agreements No 649743 (STORM) and No 768936 (TEMPO).

Appendix A. Calculation of expected load duration

A short derivation for Equation (1) is presented, which is used in Section 4.1 to calculate the expected duration corresponding to a specific heat load level power l^* in the case that the heat load profile is a random function $L = \{L(t)|t \in T\}$ of time over the time period T. For generality and ease of notation in the derivation, continuous-time notation is used, although practical calculations require time discretisation.

In the deterministic case, the duration d_l of a certain heat load level l^* given the load profile $l(t)$, i.e. the amount of time that the heat load profile $l(t)$ is higher than or equal to l^*, can be calculated using following time integral over the period T:

$$d_l(l^*) = \int_{t\in T} \mathbf{1}_{l(t)\geq l^*} \mathrm{d}t, \tag{A.2}$$

where $\mathbf{1}_{l(t)\geq l^*}$ is a function that indicates whether $l(t)$ is higher than or equal to l^*, i.e. $\mathbf{1}_{l(t)\geq l^*}$ evaluates to 1 when $l(t) \geq l^*$ is true and to 0 otherwise.

In the probabilistic case, the expected duration $E[d_L]$ (in a statistical sense) of a certain heat load l^* given the random load profile L can then be calculated as follows:

$$E[d_L](l^*) = E[d_L(l^*)] \tag{A.3}$$

$$= E\left[\int_{t\in T} \mathbf{1}_{L(t)\geq l^*} \mathrm{d}t\right] \tag{A.4}$$

$$= \int_{t\in T} E\left[\mathbf{1}_{L(t)\geq l^*}\right] \mathrm{d}t \tag{A.5}$$

$$= \int_{t\in T} P[L(t) \geq l^*] \mathrm{d}t \tag{A.6}$$

$$= \int_{t\in T} (1 - F_{L(t)}(l^*)) \mathrm{d}t, \tag{A.7}$$

where $E[\,\cdot\,]$ and $P[\,\cdot\,]$ denote the expectation and likelihood operators respectively. $F_{L(t)}$ is the cumulative distribution function of the random variable $L(t)$, i.e. representing the likelihood that $L(t)$ is smaller than l^*:

$$F_{L(t)}(l^*) = P[L(t) < l^*]. \tag{A.8}$$

References

[1] Lund H, Werner S, Wiltshire R, Svendsen S, Thorsen JE, Hvelplund F, Mathiesen BV. 4th generation district heating (4GDH). Energy 2014;68:1–11. https://doi.org/10.1016/j.energy.2014.02.089.
[2] Lund H, Østergaard PA, Connolly D, Mathiesen BV. Smart energy and smart energy systems. Energy 2017;137:556–65. https://doi.org/10.1016/j.energy.2017.05.123.
[3] Lund H, Østergaard PA, Chang M, Werner S, Svendsen S, Sorknæs P, Thorsen JE, Hvelplund F, Mortensen BOG, Mathiesen BV, Bojesen C, Duic N, Zhang X, Möller B. The status of 4th generation district heating: research and results. Energy 2018;164:147–59. https://doi.org/10.1016/j.energy.2018.08.206.
[4] Vanhoudt D, Claessens B, Desmedt J, Johansson C. H2020 STORM project: self-

organising thermal operational resource management presentation of the project and progress update. In: 2015 international conference on renewable energy research and applications (ICRERA). IEEE; 2015. https://doi.org/10.1109/icrera.2015.7418502.
[5] Vanhoudt D, Claessens B, Desmedt J, Johansson C. Status of the Horizon 2020 storm project. Energy Procedia 2017;116:170–9. https://doi.org/10.1016/j.egypro.2017.05.065.
[6] Johansson C, Vanhoudt D, Brage J, Geysen D. Real-time grid optimisation through digitalisation – results of the STORM project. Energy Procedia 2018;149:246–55. https://doi.org/10.1016/j.egypro.2018.08.189.
[7] Koch S, Mathieu JL, Callaway DS. Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In: 17th power systems computation conference (PSCC); 2011.
[8] Vanthournout K, D'hulst R, Geysen D, Jacobs G. A smart domestic hot water buffer. IEEE Trans Smart Grid 2012;3:2121–7. https://doi.org/10.1109/tsg.2012.2205591.
[9] Vandael S, Claessens B, Hommelberg M, Holvoet T, Deconinck G. A scalable three-step approach for demand side management of plug-in hybrid vehicles. IEEE Trans Smart Grids 2012. https://doi.org/10.1109/TSG.2012.2213847.
[10] Benonysson A, Bøhm B, Ravn HF. Operational optimization in a district heating system. Energy Convers Manag 1995;36:297–314.
[11] Ernst D, Geurts P, Wehenkel L. Tree-based batch mode reinforcement learning. J Mach Learn Res 2005;6:503–56.
[12] Ruelens F, Claessens B, Vandael S, Iacovella S, Vingerhoets P, Belmans R. Demand response of a heterogeneous cluster of electric water heaters using batch reinforcement learning. In: Power systems computation conference; 2014.
[13] Johansson C, Bergkvist M, Geysen D, Somer OD, Lavesson N, Vanhoudt D. Operational demand forecasting in district heating systems using ensembles of online machine learning algorithms. Energy Procedia 2017;116:208–16. https://doi.org/10.1016/j.egypro.2017.05.068.
[14] Geysen D, Somer OD, Johansson C, Brage J, Vanhoudt D. Operational thermal load forecasting in district heating networks using machine learning and expert advice. Energy Build 2018;162:144–53. https://doi.org/10.1016/j.enbuild.2017.12.042.
[15] 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. https://doi.org/10.1016/j.energy.2018.05.111.