D5.9 - Lessons learned on advanced controls in DHC networks

Renewable and Waste Heat Recovery for Competitive District Heating and Cooling Networks

REWARDHeat

Project Title: Renewable and Waste Heat Recovery for Competitive District Heating and Cooling Networks

Project Acronym: REWARDHeat

Deliverable Title: Lessons learned on advanced controls in DHC networks

Lead beneficiary: EURAC

Roberto Fedrizzi, EURAC

Mohammad Hossein Fouladfar, EURAC

Tim Diller, EURAC

Francesco Turrin, EURAC

Amir Jodeiri, EURAC

Christian KEIM, EDF,

Karine PARPILLON, DALKIA

Clement FLINOIS, EDF

Come BISSUEL, EDF

Maxime Dufour, Artelys

Manfred Le Callonnec, Artelys

Nicolas Lair, Artelys

Daniel Stenberg, E.ON Sweden

Konrad Sikora, SEC

Dirk Pietruschka, ENISYST

Mathias Obermüller, ENISYST

Marcus Brennenstuhl, ENISYST

Adrià Serra Oliver, SAMPOL

Pau Joan Cortes Forteza, SAMPOL

Due date: 30.03.2024

This document has been produced in the context of the REWARDHeat Project.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 857811. The European Commission has no liability for any use that may be made of the information it contains

Table of Contents

1 Introduction

As global energy demands grow, the need for sustainable and efficient heating and cooling systems has become more urgent, particularly within urban and industrial sectors where energy-intensive District Heating and Cooling (DHC) networks are prevalent. The REWARDHeat project, funded under the European Union's Horizon 2020 program, aims to transform DHC systems by integrating renewable and waste heat sources, making these networks more sustainable, efficient, and cost-effective.

The primary focus of the activity reported in this report is to incorporate advanced control systems and leverage digital technologies to optimize the operational efficiency of DHC networks. Through strategic partnerships with several European demonstration sites, we tested various control methodologies and data-driven strategies that aim to improve energy recovery and distribution, lower emissions, and minimize operational costs. Each demonstration site brings unique characteristics and challenges, enabling the exploration of diverse approaches to network control and optimisation, which are documented throughout this report.

Advanced controls in DHC systems include innovations in predictive control, data mining, and real-time optimisation. These systems are designed to respond dynamically to fluctuating energy demands and integrate renewable and waste energy sources efficiently. In this context, the REWARDHeat project utilizes machine learning algorithms, SCADA systems, and advanced data analytics to create data-informed solutions. By harnessing digital tools and data-driven models, the project aspires to transition DHC networks from conventional, static control approaches to adaptive, optimized systems that can better accommodate variability in renewable energy availability and demand patterns.

Each section of the report delves into the technical and operational challenges faced at specific demonstration sites and highlights key findings from data-driven and machine-learning approaches to control optimisation. The insights gathered here serve as valuable benchmarks and inspiration for the future development of smart DHC systems.

2 Albertslund (Artelys)

2.1 Description of the demonstration site

The case study of Albertslund is located about fifteen kilometres west of central Copenhagen. Currently, the existing DHN supplies heat to a major portion of the municipality. The DHN was initially built in 1964 and covers around 90% of the municipality's thermal demand.

This network is connected to the *Greater Copenhagen* DH transmission network, which integrates heat produced by waste incineration, CHP plants and peak-load boilers, and provides most of the heat uses throughout the year, i.e., 100 MW of heat capacity from the transmission company *VEKS,* supplied at 100-110°C.

Additionally to withdrawing energy from the backbone, natural gas and oil boilers are also set up as local reserve sources, accounting for 145 MW of capacity installed, while waste heat from a data centre (0.35 MW approximately, recovered at ~20°C) is also supplied nearly constantly to the DHN by means of a heat pump.

Around 270 GWh are distributed along the DHN and 220 GWh are finally supplied to end users yearly.

Figure 1 – Albertslund municipality; Porsager district in orange

2.1.1 Objectives of the demonstration activity

The project involves two partners: Albertslund Forsyning's, the municipal utility company who manages the DHN and Artelys who have developed the datamining and optimisation tool.

Albertslund Forsyning's overall energy efficiency strategy focuses on moving from a 3rd generation, high-temperature (i.e. 85°C) and centralised DHN, to a 4th generation system with lower supply temperature (i.e. 60°C) and a more distributed heat generation. To implement their strategy, a range of activities is foreseen addressing both heat distribution and distributed generation.

In the framework of the REWARDHeat project, effort has been placed on lowering the supply temperature to a subnetwork, consisting of around 110 residential houses in Porsager, an area in the South-East of Albertslund. This has been pursued by:

- Installing a shunt valve to lower the district heating supply temperature from 85°C to 60°C across the entire area. This measure follows an extensive retrofit of most of the district's homes, completed in recent years in collaboration with Albertslund Forsyning. The change occurred on the fourth of January 2021.
- Exploitation of local waste heat from two supermarkets using heat pumps.
- Developing a datamining software to gather and structure monitoring data from two separated SCADA solutions in place and enable implementing performance optimisation through continuous supply temperature modulation (led by Artelys within WP5 scope).
- Exploration of a data-based approach to minimize DHN operation costs. In particular, it aims at adapting the operation of the shunt in order to control the return temperature under a certain threshold to prevent maluses and foster bonuses. (Incidentally reducing heat losses too) (led by Artelys within WP5 scope)

The rest of this chapter will focus on describing the implementation of the third and fourth measure.

2.2 Control problem formulation

As previously mentioned, the implemented controls centred around the widely available data. Key to this is the data mining software, which enables the creation of a clear picture of the Albertslund DHN, providing valuable insights for the daily operation of the DHN. Furthermore, a data-driven approach to DHN return temperature control through shunt operation optimisation has been explored in the specific case of the Porsager district.

Data mining tool

The boom of decentralized networks linked to smart meters with remote communication, which is the case for Albertslund, makes it possible to gain increased knowledge of real-time DH network's operation. Due to large amounts of data, it is at the same time increasingly difficult via ordinary control tools to perform assessments of whole networks. Moreover, as traditional control strategies do not necessarily suffice, the need arises for new approaches that can handle and summarize the vast amount of data into data patterns that indicate to the operators when the system does not operate optimally. Consequently, the following goals have been pursued:

As traditional operating strategies do not necessarily suffice, the need arises for new approaches and methods that can handle and summarize the amounts of data and the data patterns that, for

example, can provide optimal operation or give indications when the system does not function optimally. Consequently, the following goals have been pursued:

- importing data from several data sources: The municipality employs a variety of equipment and sensors from different companies but lacks software to unify all the network data from heat supply to heat consumption. This situation is common, as equipment providers typically offer multiple software services dedicated to their own hardware;
- visualising all data in a unified interface: this aspect is particularly important as it provides the operator with a global overview of the network's operation;
- cross-correlating information from the different data sources, e.g., visualising customer data on a geographical basis and connecting this to information on network operation;
- aggregating data at the district level, e.g., aggregating customers uses and providing information on the average forward temperature at the customer side.

Shunt optimisation

Additionally, the datamining software leverages these aggregating and structuring capabilities to forecast future thermal loads insisting on the network, based on weather forecasts already available, and to perform optimisation tasks.

Particularly, the optimisation task implemented is the continuous identification of the minimum supply temperature allowing to guarantee thermal comfort at user side. One considered option is to increase the network temperature punctually during the peak hours. However, there are constraints on thermal inertia inside the pipes, requiring anticipating the temperature rise. Hence the requirement for a reliable load forecast of the network which is being specified.

As part of the optimisation work, the goal was to explore operation solutions for the shunt valve set up, allowing to minimize energy losses and thus reducing operating costs. One way to decrease the energy losses is to reduce the return temperature. To tackle this measure, a data-driven approached has been chosen, relying on modelling the return temperature as a function of shunt valve control parameters (supply pressure and temperature).

2.3 Description of the control implemented

2.3.1 SCADA system

As mentioned, Albertslund Forsyning handles several data sources. The datamining software connects three different data sources:

- SCADA system delivering data from several network locations, including weather station providing weather data used for thermal loads forecasting, and supply data from the heat exchangers connecting the local DHN to Copenhagen's transmission network;
- Consumption data: Albertslund installed heat meters at the customers substations. These meters are referred to as the 'smart meters' in this document. They are not directly connected to the SCADA system and data is gathered over a different stream;
- Network data: is static GIS information containing a detailed geographical representation of the DHN.

SCADA system

Being part of the Greater Copenhagen system, the utility relies almost entirely on the transmission network. In particular, the local DHN is supplied by means of six heat exchangers. In addition, a small portion of the heat is gathered from a data centre and one supermarket.

The SCADA system collects all the data produced at the heat exchangers, including cumulative water and energy exchanged, water and energy flows, supply and return temperature, supply and return pressure. The data is retrieved as a csv file every 5 minutes.

In addition, the SCADA system integrates the weather data produced by the local weather station including the following information:

- wind speed and direction;
- outdoor temperature and humidity;
- solar irradiation.

Smart meters

Albertslund Forsyning can count on around 6300 heat meters which data is acquired on a hourly basis. The data is actually stored on Artelys servers. The metered data includes activation time, cumulative water and energy exchanged, water and energy flows, supply and return temperature. The data is retrieved as a csv file of around 1MB every hour.

Network data

The network data is considered static information since network updates are not so frequent. This data is shared as a GIS file containing the network layout. Albertslund's heating system is composed of around 25000 pipes. Each connection and the connections to customers are geographically identified and reported in the GIS file.

2.3.2 Data mining software

A data mining software is a platform that is not used to directly control a system, rather it helps initially the operators identifying eventual margins for improving the management of the system and secondly implementing a better management by computing optimal operation trajectories. With this respect, the optimal control options are communicated to the operator who decides for the actual utilisation, with eventual modifications.

Software architecture

The data mining software is a web application, the choice of a web version has several advantages:

- Granting access to the software to a new user is simple.
- The delivery of the software is simplified, in particular in case of a version upgrade.
- It offers a server able to handle way more efficiently intensive computing.
- Data protection on the servers is ensured.

The complete architecture for the Albertslund data mining software is described hereafter from data collection to the consumption of KPIs by the DHC network operators.

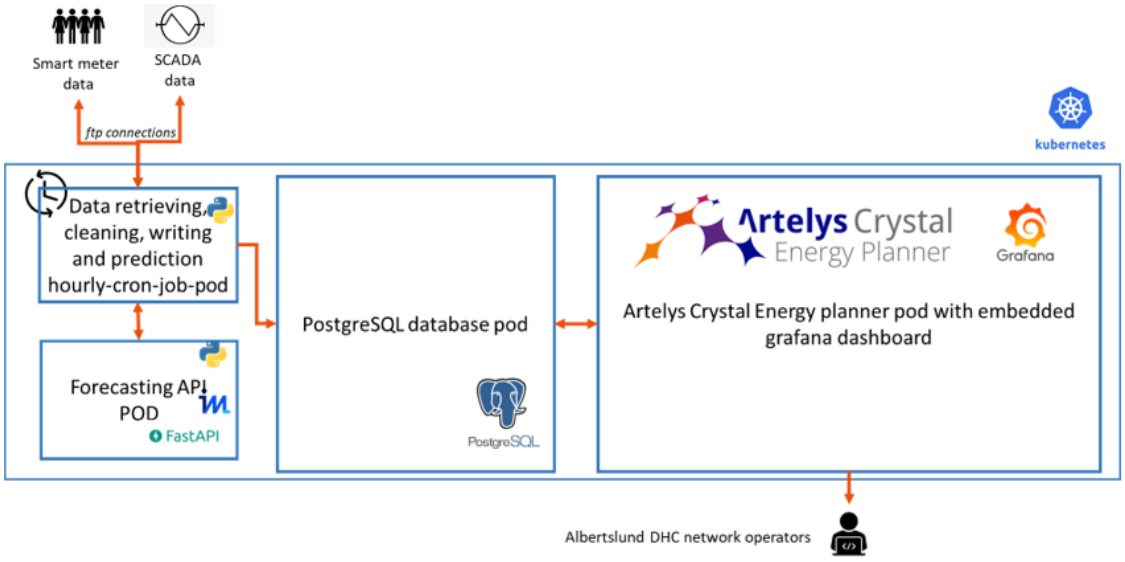

Figure 2 - Albertslund data mining software's architecture

User interfaces

At first, a cartographic view of the municipality, including all DHN users and network pipelines has been elaborated, as shown in Figure 3.

The pipe layout is displayed along with the other 4 asset types (namely district, weather station, heat exchangers and smart meters). The districts are defined as polygons on the map whereas the other asset types (weather station, heat exchangers and smart meters) are represented with points. The map is interactive as it is possible to zoom-in to investigate specific areas.

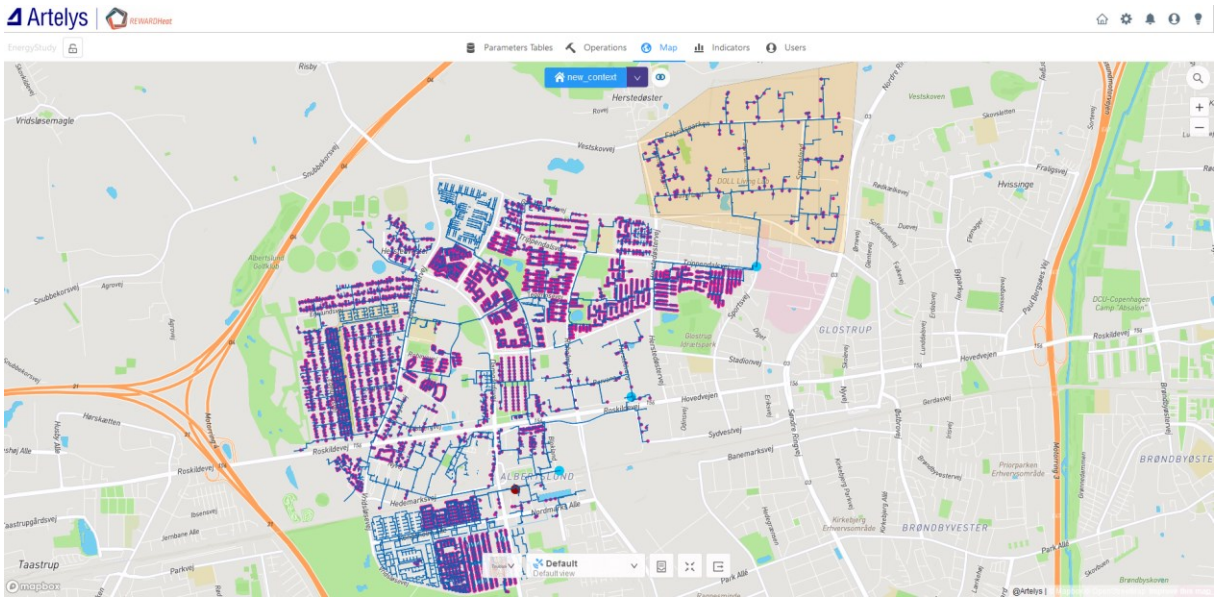

Figure 3 – Cartographic view with Albertslund DHC network.

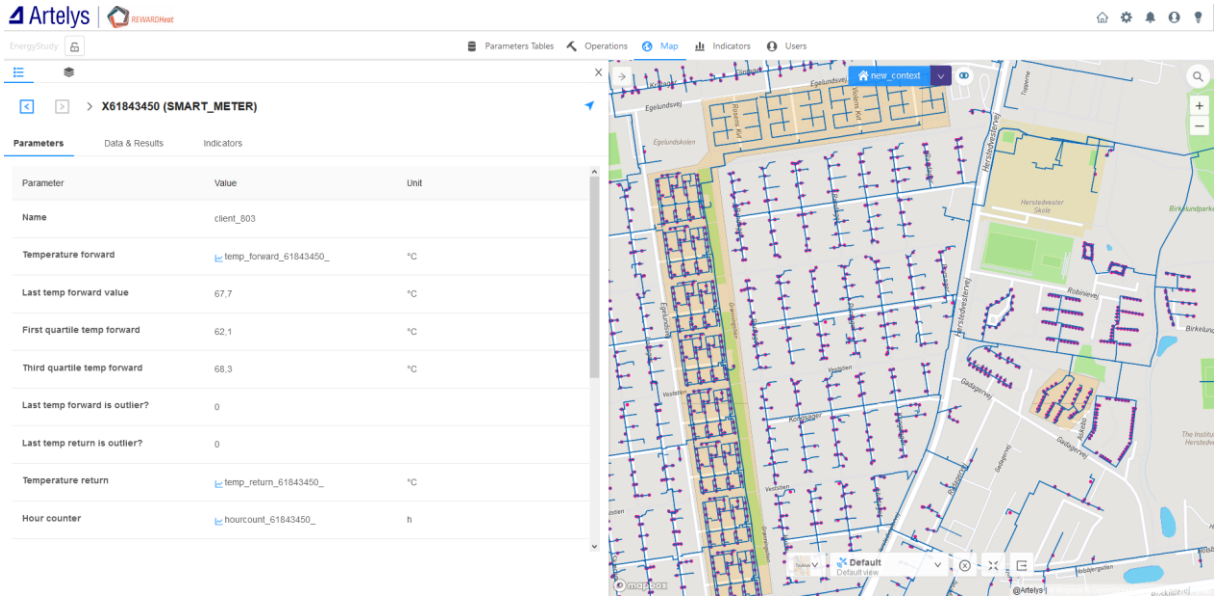

Figure 4 – Asset view from the cartographic view.

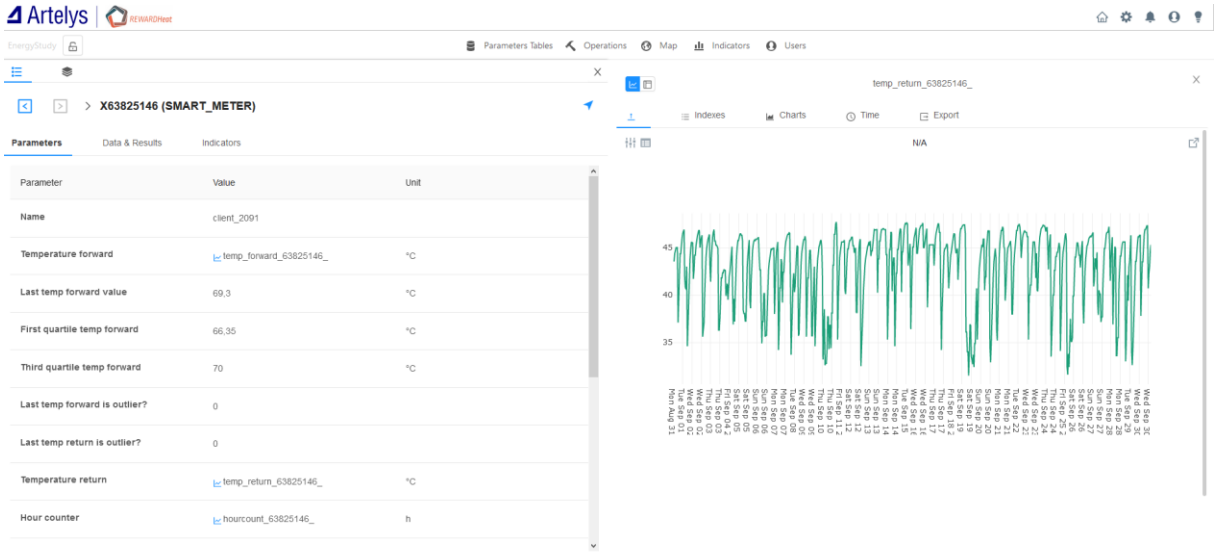

Figure 5 – Visualisation of timeseries in the asset view.

Apart from the zoom, the user can also open the asset view directly from the map by clicking on one of the assets. When doing so, the asset view opens alongside the map, providing an overview of the asset attributes. This includes an overview of each smart meter asset:

- Supply temperature (°C) timeseries
- Return temperature (°C) timeseries
- Volume (m3), accounting for the water volume delivered since the meter reset.
- Hour counter (h), i.e., elapsed time since the meter reset.

In addition, several attributes are computed to be used as indicators either on the map or the KPI view. The timeseries can also be visualized by the user simply clicking in the asset view as presented in Figure 5.

KPIs calculated

This section offers a first overview of the KPIs implemented in the datamining software. Using the geographical coordinates of the smart meters and districts, it is possible to assess the district of each customer. Several indicators have been computed at the district level in the form of aggregated KPIs as below where the average forward and return temperatures for different districts can be compared.

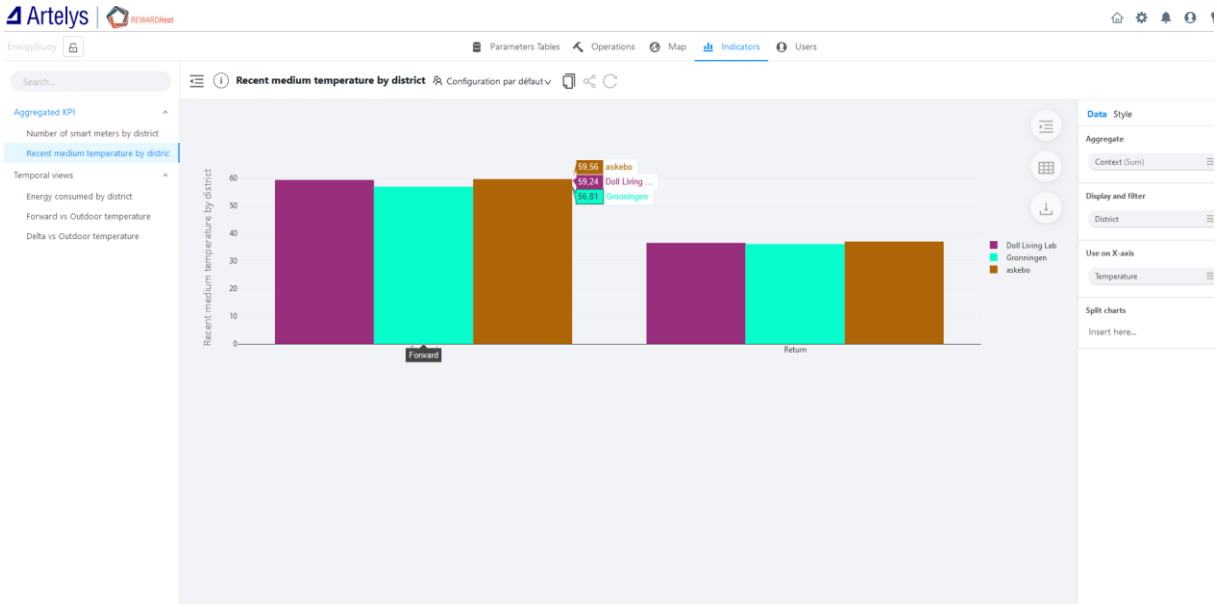

Figure 6 – Average supply and return temperature by district.

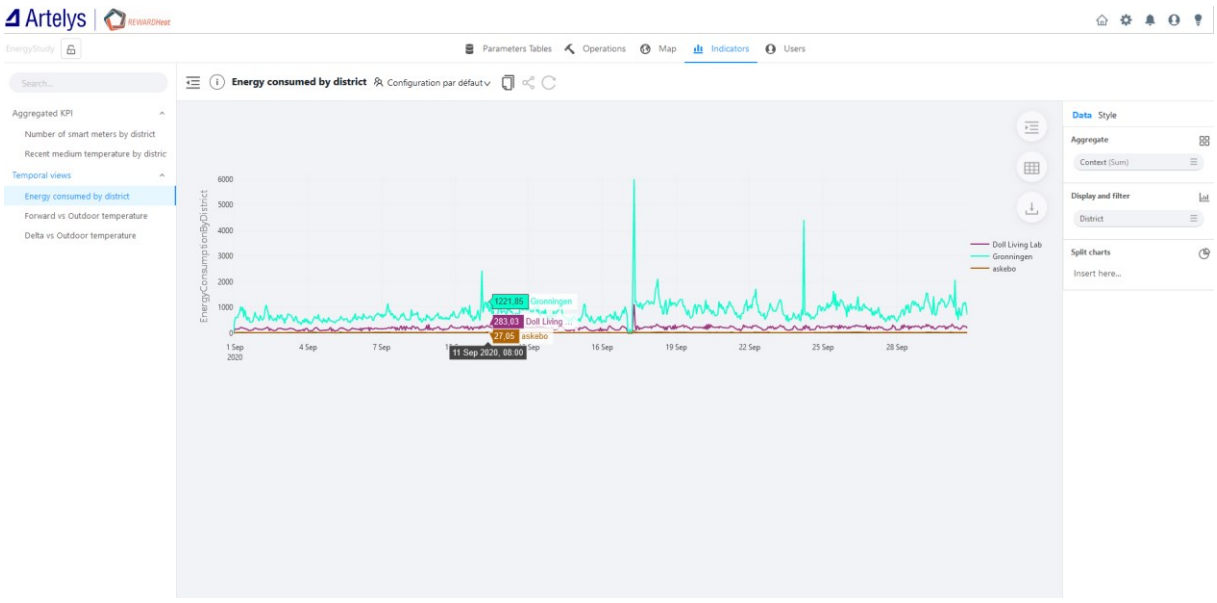

Figure 7 – Energy consumption by district.

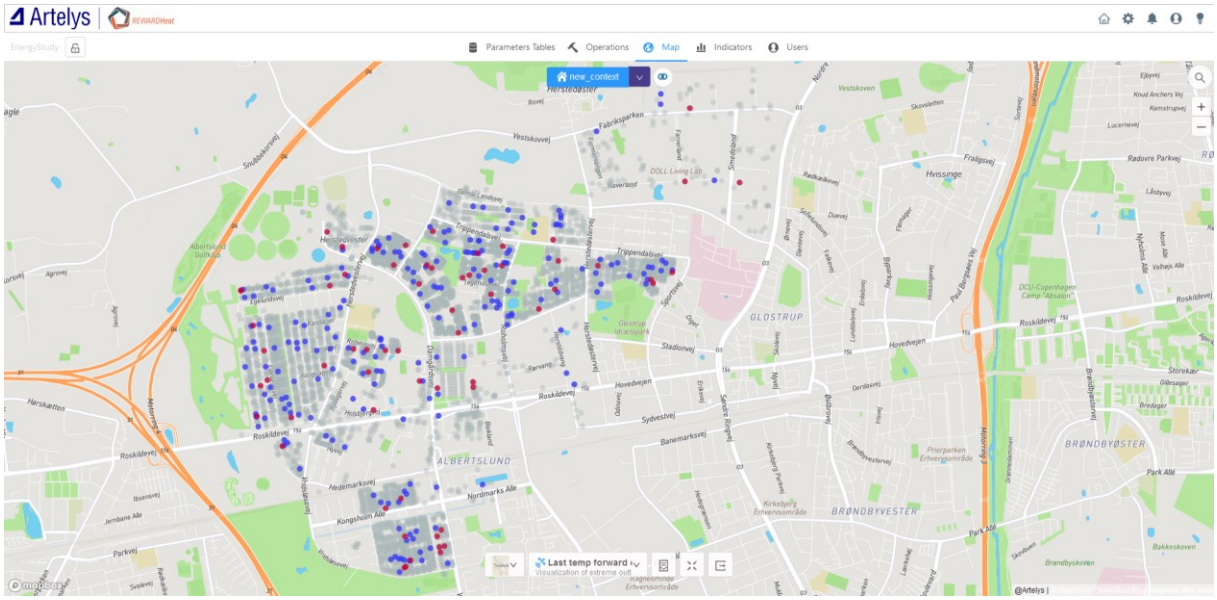

Figure 8 – Detection of outliers in the cartographic view.

In addition, timeseries indicators can be computed for each district, relying on the smart meter data as presented below.

Again, the interface of the datamining software is completely interactive to facilitate as much as possible the user workflow. To that end, it is possible to navigate in the timeseries values and zoom in for further investigation of potentially crucial time periods.

The last indicator is available directly in the cartographic view. It highlights all the smart meter outliers with regards to their current forward temperature as presented in Figure 8. The customers represented in blue (resp. red) are meters for which their current forward temperature is particularly low (resp. high) with regards to their historical data. The outlier analysis was tested to demonstrate the feature capabilities and integration in the cartographic view.

An asset-based model of the network for better scalability

One of the key concepts of the software is the ability to model any production and/or consumption unit or cluster (representing an aggregation of generation units) by a so-called "asset" in the energy system model. For instance, heat exchangers, gas boilers and network pipes are all "physical assets" that are used to model heat production and transmission.

Each asset type is defined by an asset model and a set of parameters. Several assets are available in the asset library; new models can be defined as needed, e.g., with the development of low-temperature district heating networks and the integration of distributed renewable and waste heat sources, new models with specific parameters and constraints might be needed.

The operational data is retrieved from ftp severs through a regular operation to collect, clean, store the data. The same script then calls the python forecasting api to request the 6 hour-rolling-forecast data and write it into a database. The data is then displayed through multiple UIs from both Artelys Crystal Energy Planner (GIS view, asset-based data) and its integrated Grafana dashboard (DHC network wide KPIs).

2.3.3 Shunt optimisation

In this section, we report on how the shut operation optimisation has been performed. Figure 9 provides a simplified representation of the Porsager heating network. There are two ways to control the operation of the Porsager grid through the shunt aimed to reduce the return temperature. One can change the supply temperature after the shunt (T_f2 in the figure), thanks to a mix between the supply and the return water inside the shunt. The second variable that can be changed is the supply flow.

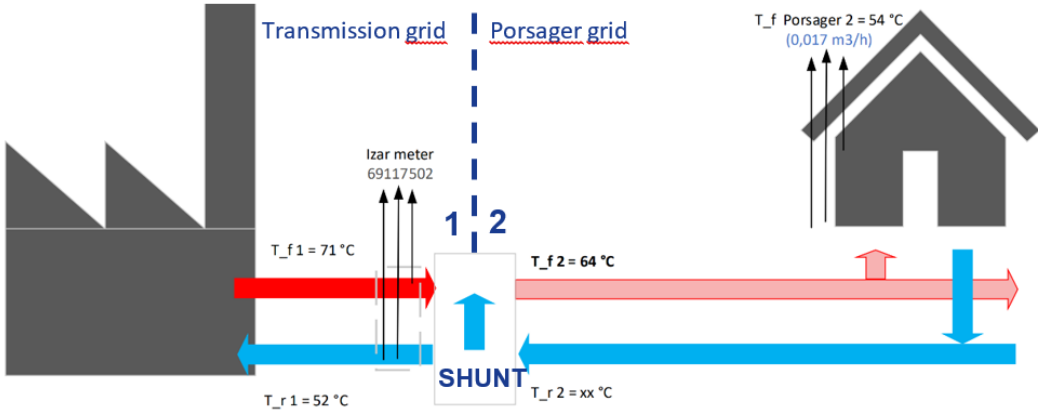

Figure 9 - Simplified representation of the Porsager heating network. Upstream of the shunt (left) is the transmission grid, downstream (right) is the Porsager distribution grid. Forward pipes are shown in red and return pipes in blue.

Several methodological choices could be made to tackle the problem stated above. For instance, one could choose to define an objective function, model the network with thermodynamic equations linking our decision variables and exogenous variables to our target variables, and eventually use an optimisation technique to solve it. Instead, we have experimented with a data-driven approach to benefit from the data environment implemented with the datamining tool.

This approach relies on statistically learning the relationships between the decision variables (supply temperature and supply pressure) and the target variables (return temperature). Once this relationship has been learnt from the data, one can knowingly select the right decision variable values given a certain exogenous situation to control the target variables. To learn such a relationship, we use machine learning techniques and exogenous variables, such as house meter consumption and some weather data, that are responsible for this relationship.

This process is summarised in the right part of Figure 10 for the case of return temperature optimisation. Once the learning process is complete (left part of Figure 10), the prediction model can be used to optimise shunt operation (right part of Figure 10). Here is an example of how this process could run in the case of return temperature optimisation:

1. Feed the model with current decision variable values and exogenous variables (or forecasts) and compute a prediction of the return temperature.

2. While the predicted return temperature is out of the acceptable range:

   a) Change the value of one of the decision variables.

   b) Compute the new predicted return temperature.

   c) Return to step 2.

3. Save the final values of the decision variables as the ones to use.

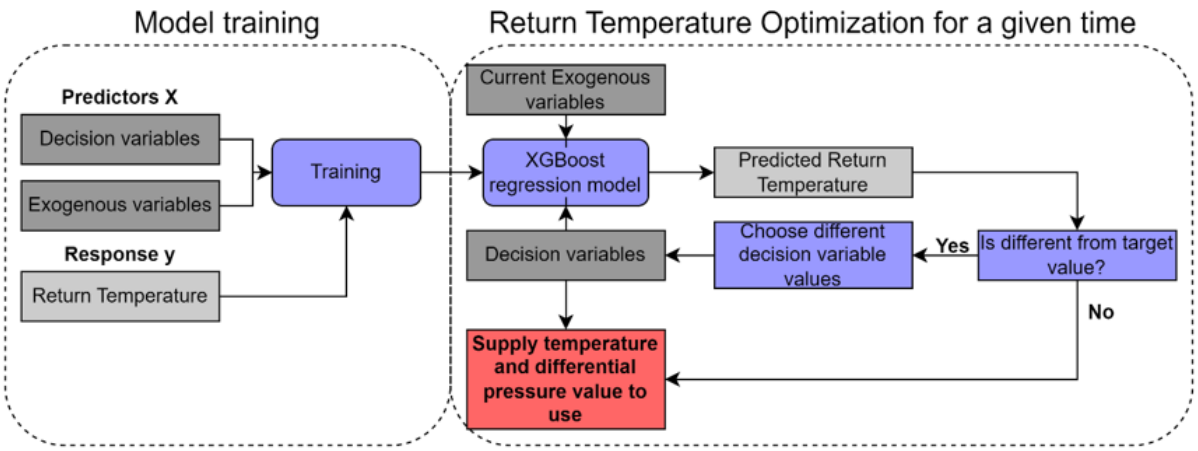

Figure 10 - Diagram representing the data-based shunt optimisation approach. The left part corresponds to learning the relationship between the decision variables and the target variables. The right part shows how this learnt relationship could be used to optimally operate the grid on a daily basis

2.4 Assessment of the data mining tool performance

2.4.1 Key Performance Indicators

First, the geographical components of the KPIs are a major factor to efficiently identify network disruptions by cross-checking different data sources. The objective is to provide complementary information between the cartographic view and the KPI view. Several KPIs can also be visualized on the cartographic view as data layers to facilitate the network operator analysis.

To identify network weaknesses or inefficiencies, highlighting the smart meters with extreme measurements (among all the smart meters) is particularly interesting. To that end, it is possible to visualize the list of smart meters with extreme values. Depending on the operator interest, several values can be considered e.g., the highest return temperature, the lowest forward temperature or the lowest temperature difference (between forward and return temperature).

These indicators can be visualized on the map with a dedicated layer highlighting the customers with such characteristics giving the operator a global overview of the network state.

In addition to extreme measurements among the smart meters, inconsistent measurements for a given smart meter can be identified, relying on its historical data. One efficient approach suggested by the literature is to use the interquartile range (IQR), which measures statistical dispersion as the difference between the 75th and 25th percentiles. For a distribution of smart meter energy consumption values over time, the IQR can be calculated and visualised as shown in the figure below. The IQR is often employed to identify outliers in data. Outliers are defined as observations that fall below Q1 − 1.5 IQR or above Q3 + 1.5 IQR. In a boxplot, the highest and lowest values within this range are indicated by the whiskers (often with an additional bar at the end), while any outliers are shown as individual points.

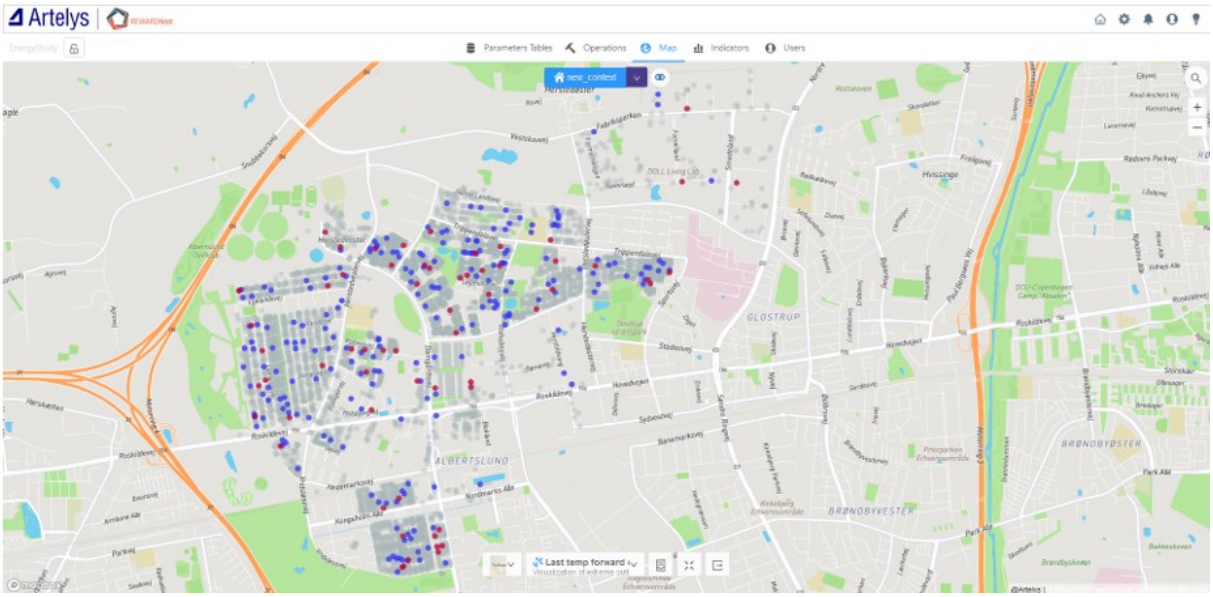

Figure 11 Detection of value outliers (in red) in the cartographic view.

Albertslund Forsyning particularly emphasised the importance of integrating information from multiple data sources. Consequently, efforts were made to combine data from both the smart meters and the SCADA system. The temperature readings from the smart meters are merged with the meteorological data from the weather station integrated into the SCADA system. The currently computed indicators are:

- the ratio of the smart meter supply temperature to the outdoor temperature provided by the weather station;
- the ratio of the smart meter temperature difference to the outdoor temperature.

Although these indicators are simple, they are useful for the network operator, who relies on their knowledge of the network and typical behaviours.

Asset model developed for Albertslund

A first asset model was implemented for Albertslund demo site considering the network specificities and available data. This model is composed of 6 different asset types: Heat exchanger, Shunt, Weather station, Smart meter, District, Pipe. Each asset contains several attributes either retrieved directly from various data sources or computed for visualisation and KPIs analysis.

2.4.2 Shunt optimisation

The load forecasting model uses 2 different algorithms in parallel and compares them at each simulation step, to reduce prediction uncertainty. The Explainable Boosting Machine (EBM) is a tree-based, cyclic gradient boosting Generalized Additive Model with automatic interaction detection. EBMs are often as accurate as state-of-the-art blackbox models while remaining completely interpretable. Additionally, a Generalized Additive Model (GAM) algorithm has been implemented. GAMs combine parametric and non-parametric techniques, making them suitable for a wide range of modelling problems, from standard linear regression to more complex tasks.

Both models perform well during training, as shown in Table 1, resulting in excellent performance in return temperature forecasting. The Mean Absolute Percentage Error (MAPE) is 3% for unseen data. Additionally, Figure 10 demonstrates that the predictions closely match the observations.

Table 1 – Prediction performance of the models employed

MAPE (%)	EBM	GAM	Return temp
TRAIN	7.0	7.3	2.1
TEST	7.0	5.9	3.1

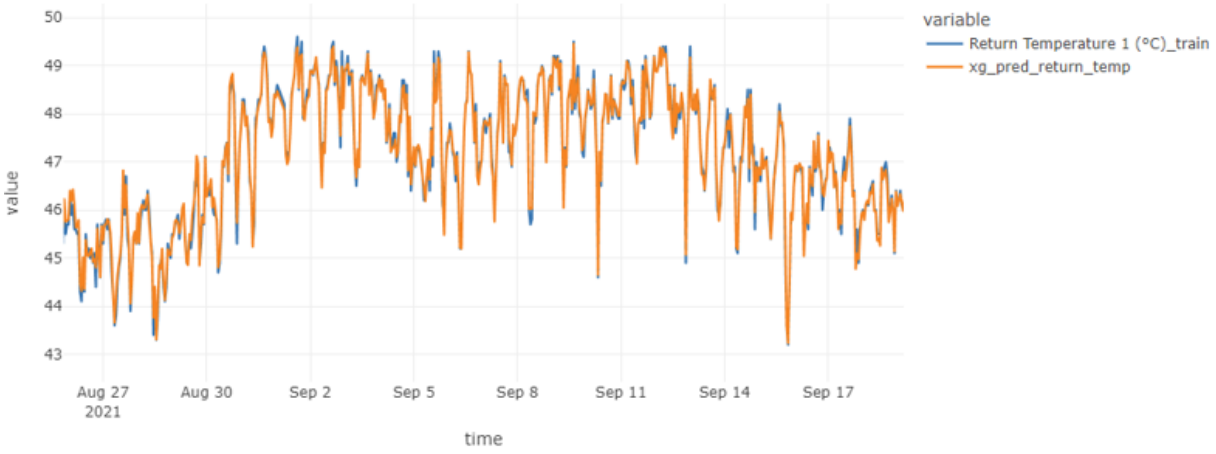

Figure 11 - Time series of the observed (blue) and predicted (orange) return temperature from 27th August 2023 to 17th September 2023

Despite the promising initial results, the implementation of the optimisation methodology has proven difficult: it turned out that the data made available are limited in terms of depth and quality. For an accurate optimisation, the model of the return temperature should have been trained with various sets of decision variables, but the historic data does not exhibit a large distribution over the decision variables as the operator always use the same set of values. The model is thus not allowed to learn an accurate estimate of the return temperature when changing the common decision variables. Furthermore, there is very little available data that can reflect on the network behaviour before the lowering of the supply temperature.

2.5 Lessons learned

The main learnings out of operating the data mining tool are:

- Barriers on SCADA and Data Quality: Although the Albertslund DHN is equipped with data retrieval systems, they sometimes fail to provide good quality data or any data at all, causing issues with the implemented tools.
- Shunt Optimisation: The limited quantity and quality of data made it difficult to accurately capture the relationship between decision variables and the return temperature, thus preventing us from fully achieving our initial objective.
- Continuous commissioning: Despite these limitations, the nature of the implemented tools fosters continuous improvement in data quality.

- Early objectives definition: The project has highlighted the importance and necessity of involving DHN network experts and operators (the software end-users) throughout the development of the data mining software. It will be crucial for future projects to emphasise this point and organise methodologies that facilitate smooth communication between both parties.
- Compliance with GDPR: led to some delays when handling sensitive data related to customers' addresses and consumption, as anonymisation work had to be carried out.
- Integrating different sources of information proved to be very insightful: The cartographic view provided by the data mining tool has been particularly appreciated, especially for the operational perspective it offers. Additionally, the ability for users of the data mining tool to directly add new components to the network (such as pipes) as it evolves would be beneficial.

3 Helsingborg (EURAC)

3.1 Description of the demonstration site

This demonstration site is located in the district of Drottninghög in Helsingborg (Sweden), and consists of a newly built, small-size, heating and cooling network supplying energy to four new apartment blocks (5 to 7 floors and 110 apartments, with a total living area of 7,795 m^2, see Figure 12. The construction has been implemented by Tornet, who are specialized in construction and management of affordable rental properties, with a commitment to energy efficiency, responsible material utilisation, and sustainable product choices.

Figure 12: The four multifamily homes during their construction

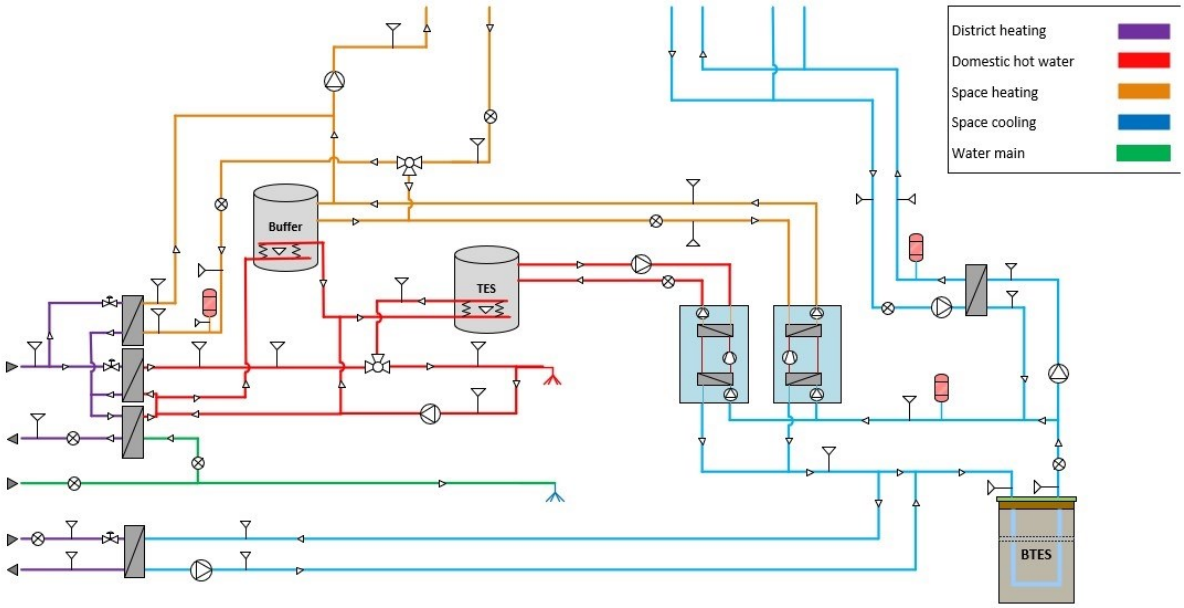

Figure 13. Layout of the Tornet Energy Centre in Helsingborg

The Energy Centre at the core of the small network consists of a thermal substation and a 6-pipe distribution network designed to cover the demands for DHW, SH, and SC. Figure 13 shows a simplified P&ID of the system: the energy system interfaces with the Helsingborg DHN network (purple lines) via two counterflow HXs—one dual-stage HX to meet DHW demand (red lines) and

one for SH demand (orange lines). The system uses two identical hot water tanks, each with a capacity of 750 litres: a high-temperature TES dedicated to meeting DHW peak loads and a low-temperature buffer that interfaces with the SH system and preheats DHW.

The Energy Centre also integrates a 4-pipe geothermal HP, equipped with a variable-speed compressor and a de-superheater that recovers heat from DHW preparation for space heating purposes. For sake of simplifying the system description, the latter is represented as two separate HPs. The HP is used solely for heating, while SC is achieved through free cooling (blue lines): heat is extracted from the building via the Air Handling Units (AHU) and directed to the boreholes.

Free cooling alone is insufficient to balance heat extraction during winter. To mitigate this, a recharging strategy for the geothermal installation is devised: PVT panels are employed to capture solar thermal energy; additionally, waste heat is recovered during space heating season from Air Handling Units' exhaust air duct, gathering an almost constant 15 kW thermal power. PVT and waste heat recovered are primarily used at HP evaporator for DHW and SH delivery to the building; excess heat, which cannot be directly exploited by the HP is fed to the BTES aimed to balance extraction.

The connection between the DH network and BTES, which could potentially supply thermal energy from the DH to the ground, is currently non-operational due to regulatory restrictions. Heat from the Helsingborg DHN is mainly drawn during summer when boreholes are being charged, while in winter it covers peak loads.

3.1.1 Objectives of the demonstration activity

The project involves three partners: INDEPRO, tasked with strategic planning, ARVALLA, overseeing construction management, and EURAC who developed an optimised RBC for the management of the boreholes field. The project aims to achieve the following objectives:

- Installation of a standardized substation, streamlining construction processes through the deployment of standardized design.
- Integration of a geothermal HP in the substation aimed to minimise the import of energy from the district heating network.
- Integration of PVT panels, reducing electricity consumption from the grid and charging the boreholes field in summer.
- Integration of a state-of-the-art smart monitoring and control hardware and software.
- Selection of the most effective way of sourcing heat for the buildings, between the DHN and the borehole field.

3.2 Control problem formulation

The latter bullet point is concerned in this report, as this has been the object of a parametric analysis leading to an optimised RBC managing sourcing thermal energy from the DHN.

As reported previously, heat from the Helsingborg DHN is only used to a limited extent in summertime, to cover the DHW demand. Indeed, this is to some extent justified by a very low price encountered along this season; on the contrary though, this management strategy requires using electricity to run the HP system during winter, when this vector is affected by high prices, lower fraction of RES compared to summer and eventual shortages are encountered during peak hours. This is particularly relevant nowadays in Scania, as most of the hydropower plants in Sweden are

located in the north of the country, and the transmission lines struggle transferring electricity to the south when subject to heavy loads.

In fact, the main driver for implementing such a management strategy is regulatory, as companies delivering affordable housing and profiting of national subsidies to this end are subject to limitations in terms of the energy imported from grids to their properties. While this is clearly driven by the aim to achieve high-quality, energy-efficient constructions, it can also result in suboptimal management of energy use both on-site and across the overall energy system.

As an example for a different exploitation of energy from the grids, Thermal energy from the DHN could be employed in summer to charge the boreholes thermal energy storage, achieving higher ground temperatures at the end of summer compared to what happens today, thus leading to higher seasonal COP values and lower electricity consumption associated to the HP system when operated in winter.

Having this in mind, in this project we have analysed the energy use at the demonstration site not only from the final energy and energy bills perspective, rather we extended the view to the yearly primary energy use and overall CO_2 emissions.

Specifically, we have modelled the buildings + Energy Centre system and run a parametric analysis aimed to infer simple and replicable management rules. The parametric simulation campaign has span over a range of varying boundary conditions.

3.3 Description of the control implemented

3.3.1 SCADA system

Communication devices, both wired and wireless, enable data exchange between the substation and central monitoring stations. A SCADA system facilitates data acquisition, monitoring, and analysis.

The data collection is primarily performed with heat meters set up on the system branches, allowing to monitor energy and mass flows, as well as supply and return temperatures. Additionally, temperature sensors are located in water storages and around heat exchangers, providing data to PLCs, which implement RBC strategies in feedback loop.

While this equipment is generally sufficient for fault detection and diagnostics, there are noticeable gaps in the data required for a comprehensive analysis of the system energy performance. Specifically, the electricity consumption of individual circulation pumps and HP compressors is not recorded. Consequently, we have calculate the electricity uses based on the mass flows distributed and on the COP values declared by the manufacturer, for what concerns the HP. Furthermore, the DHW circuit lacks both a temperature sensor to measure the main water temperature and a flow meter to monitor the recirculation flow rate.

The data is saved on a cloud server, with an acquisition period of 5 minutes; FastAPI is employed to develop web applications that interact seamlessly with the control system. This integration ensures precise control and monitoring of the Energy Centre, as data is available to operators who can enhance system efficiency and reliability.

3.3.2 Advanced controls

A rule-based control system is implemented on top of the energy system to manage the operation of various generation, distribution, and emission units. This control system follows a static strategy

based on pre-established rules. Consequently, its efficiency depends on the expertise of the person defining the rules. Despite this, rule-based control generally exhibits robustness and strong real-time performance by applying practical control rules based on the system's status. However, due to its reliance on experience, it becomes difficult to identify optimal control points for complex, dynamic systems. The rule-based control concerned is structured into blocks that interact with one another to generate the control signals for the energy system. The blocks constituting the control are:

- Feedback signals: information acquired from the sensors
- Hystereses: elaboration of the acquisition signals in Boolean format. The hysteresis, in thermal systems, is useful to avoid continuous oscillation between two states due to the nature of the system
- Schemes: represent the working modes used by the HP system. The schemes are defined as logical phrase of hysteresis
- Modulations: refers to pumps and valves and it is used to scale the control signal of the component. The modulation can be either a fixed value or a function of other independent variables (temperature or mass flow rate)
- Control signals: commands given to the devices to be controlled; it is the combination of schemes and modulations

A schematic of the rule-based control is shown in Figure 14. More details on the system control are reported in D4.5.

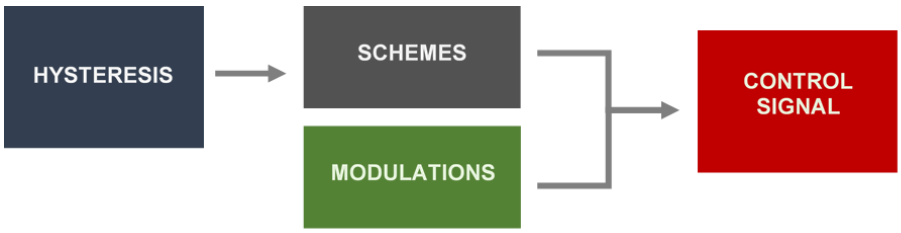

Figure 14. Block diagram for the structure of the rule-based control

During the analysis of the monitoring data, it was observed that the energy extracted from the BTES during winter is higher than the energy stored during summer. Therefore, its temperature decreases over time, reducing the COP of the heat pump. On the one hand, this increases the electricity consumption; on the other, the heat drawn from the DH also rises, because the heat pump will generate less heat during winter.

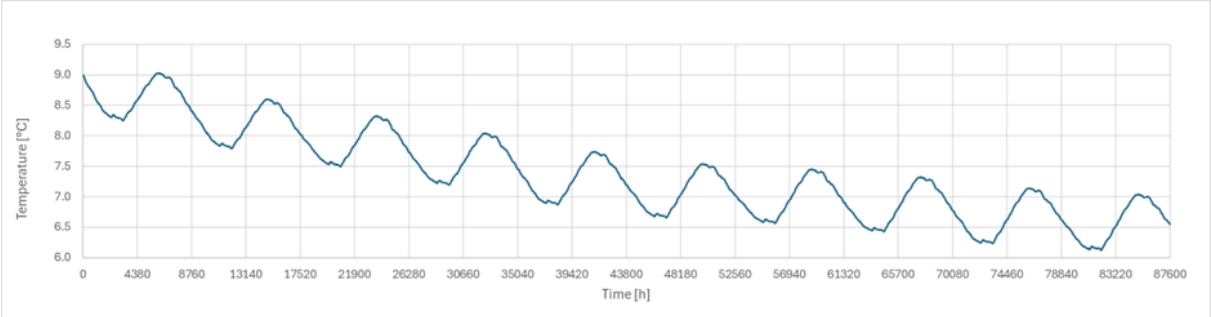

Figure 15. Average BTES temperature over 10 years of simulation

A 10-year simulation was conducted to better understand this phenomenon. Figure 15 shows the hourly average ground temperature in the BTES over the 10-year period. A seasonal behaviour of the temperature is noticeable. The BTES temperature decreases during the heating season because heat is extracted from the ground to run the heat pump. During the cooling season, the BTES temperature increases because heat is transferred from the air handling units used for space cooling to the BTES. During this period, there is also high production from the PVT system, allowing more thermal energy to be stored in the BTES.

However, the balance is negative, meaning the temperature at the end of the year is lower than at the beginning of the year. This occurs annually, causing the temperature to decrease over time. After 10 years, the average temperature is almost 3°C lower than the initial temperature. The decrease of the inlet temperature implies a decrease of the COP, especially when the heat pump is working in part-load conditions.

Figure 16 reports on the annual energy cost to cover the thermal loads of the building, divided in electrical and DH cost (see Table 3 for specific monthly costs). After 10 years the total energy cost increases by almost € 1'100, which corresponds to 9% of the initial value. The electricity cost increased by 2.9% (€ 210) and DH cost increased by 20.7% (€ 860).

To mitigate the effect of the temperature decrease, different strategies to charge the BTES during summer were defined, making use of DH energy during summer months, as represented in Figure 17.

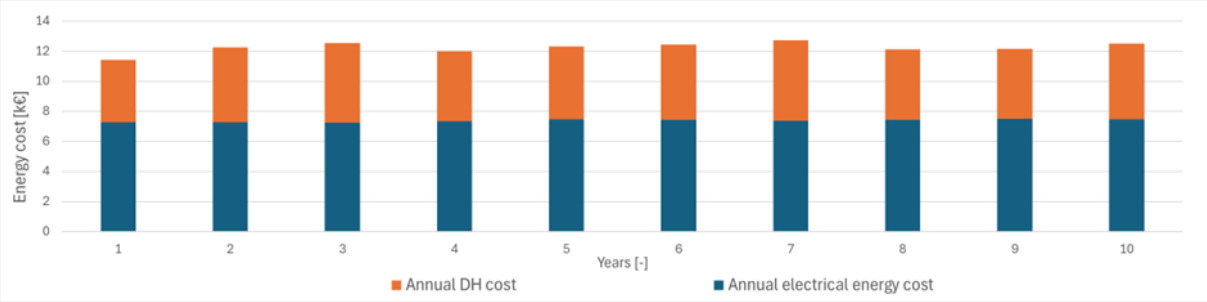

Figure 16. Variation of energy cost of the system over 10 years

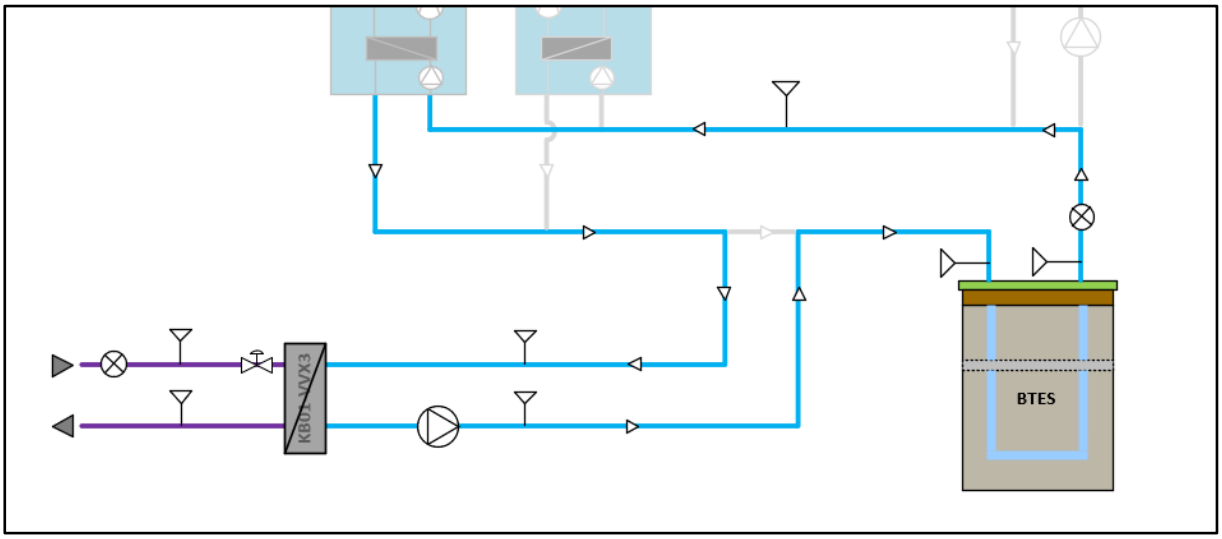

Figure 17. Scheme S7 to charge the BTES with DH in Summer

The HX between the DH and the BTES circuit has a nominal capacity of 50kW. In order to still permit space cooling, the temperature at the outlet of the BTES cannot exceed 15°C. Control logics elaborated to charge the BTES with DH account for these boundaries. The target energy to be stored in the BTES is 60MWh per year, corresponding to the annual imbalance assessed through monitoring data. Considering a maximum thermal capacity of 50kW, the drawing energy from DH for BTES charging should be operated for about 1200 hours, which correspond to roughly 2 full operation months. To limit the thermal losses of the BTES it is convenient to operate it when DH cost is low and right before winter season commences, i.e., August and September. Since feeding 50kW into the BTES circuit can raise the BTES outlet temperature above the 15°C limit, alternative charging powers and different periods might be considered. The selected parameters are reported in Table 2.

Table 2. Advanced strategies to charge the BTES

	Peak charging power [kW]	Start of charging period	End of charging period
1	0	-	-
2	50	1st August	15th September
3	25	15th April	30th September
4	20	1st April	30th September
5	16	1st April	30th September

Table 3. Monthly price for electricity and district heating

Month	Electricity price [€/MWh]	DH price [€/MWh]
January	87.2	70
February	69.8	70
March	77.0	70
April	65.6	30
May	43.1	30
June	51.9	30
July	39.9	30
August	41.4	30
September	30.6	30
October	40.7	70
November	84.7	70
December	80.9	70

3.4 Assessment of the control performance

Using different strategies has a different impact on the BTES. Figure 18 shows the internal energy variation of the BTES in one year. The initial internal energy is set to zero. Whenever energy is extracted from the BTES the internal energy drops and when energy is injected the BTES internal energy increases.

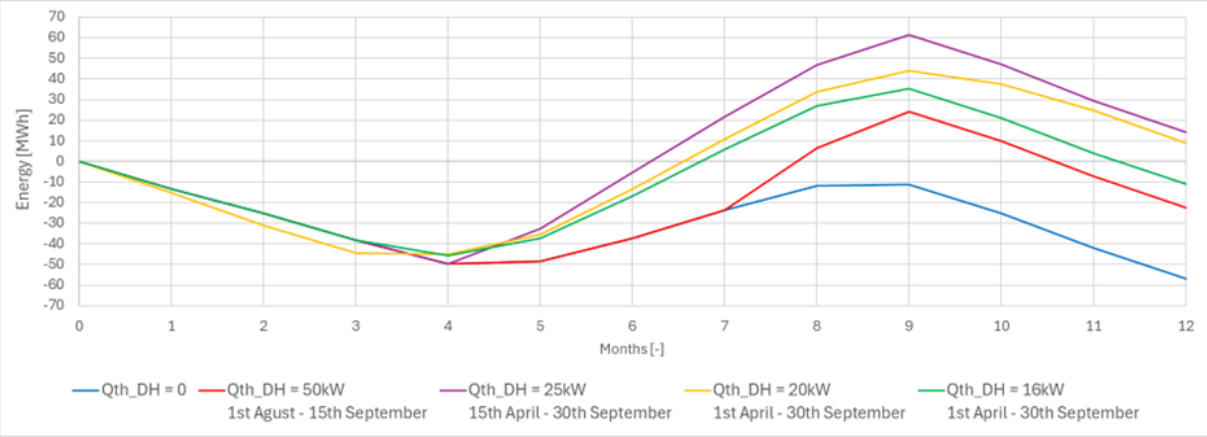

Figure 18. Internal energy variation of BTES with different charging strategies

For all these simulations the first part of the graph is the same: in this period the energy extracted from the BTES to run the heat pump is higher than the energy gain from the solar thermal circuit. Starting from April, the internal energy of the BTES increases.

The case feeding of 50kW from DH from the 1st of August to the 15th of September shows a negative energy balance. Theoretically, this combination should be sufficient to balance the injected and extracted energy. However, because the thermal power is high, the outlet temperature of the BTES often exceeds the limit temperature of 15°C. Therefore, this power cannot be continuously exploited and the actual stored energy results lower than the theoretically available.

By halving the thermal power and expanding the charging period, the energy fed to the BTES is higher than the extracted. Consequently, the BTES internal energy balance is positive and the average BTES temperature increases over the year. This benefits the performance of the system, but also results in higher operative costs. The analysis highlights that the charging strategy that keeps the internal energy balance unaltered consists in a charging power of 20kW between the 1st April to the 30th September. Figure 19 shows the average BTES temperature profile over 10 years of operation, considering the different strategies.

Figure 20 and Figure 21 show the monthly electrical energy and DH consumption of the two different controls: actual vs. latest described. Considering the system operation after 10 years, using about 60MWh of energy DH to load the BTES in summer allows to reduce the electricity use in winter by about 4.2MWh and DH energy consumption by 3.5MWh.

Considering the energy prices reported in Table 3, after 10 years the annual electrical energy cost is € 7'157.36 when using DH during summer, and € 7'487.01 for the base case scenario; the DH costs are € 6'367.30 and € 5'021.34 respectively. This means that charging the BTES allow to save € 329.65 for the electrical energy consumption but the total energy cost is € 1'016.31 higher.

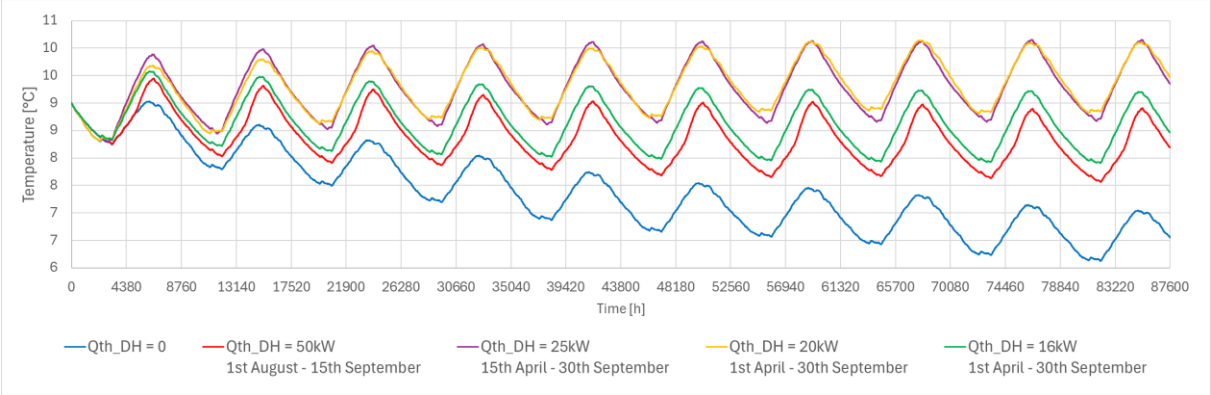

Figure 19. 10-year average BTES temperature profile according to different charging strategies

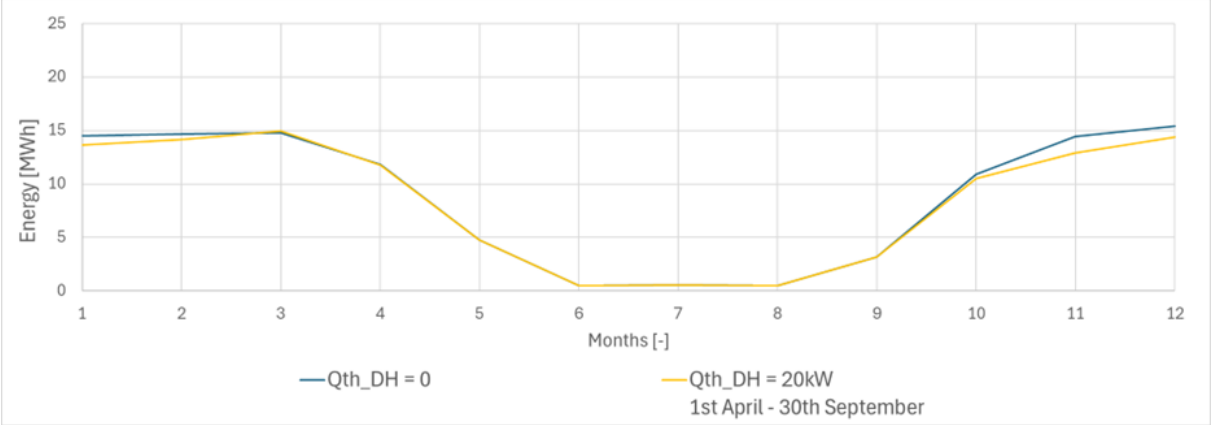

Figure 20. Electrical energy consumption of the 10th year

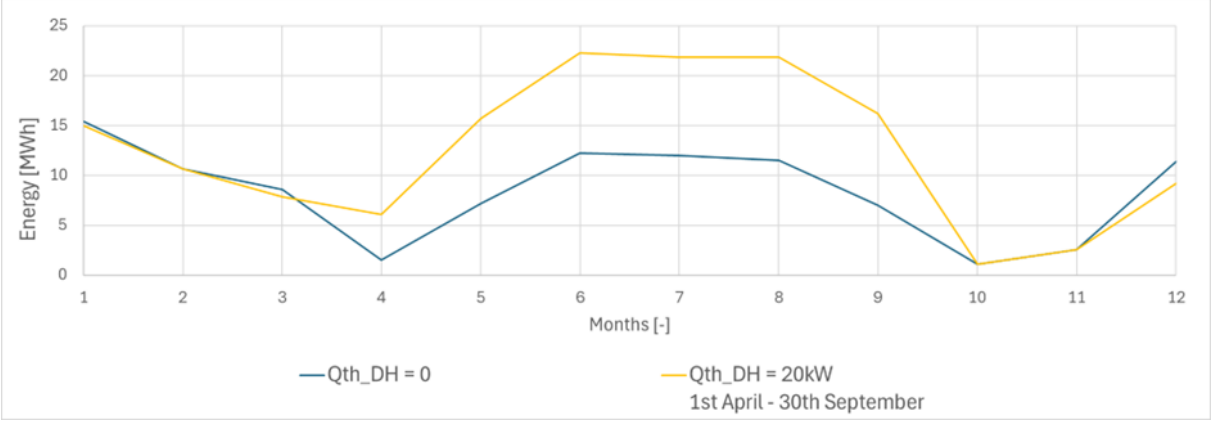

Figure 21. DH consumption of the 10th year

Figure 22, Figure 23 and Figure 24 show the cumulative energy cost (electricity, DH and total respectively) over the years of system operation considering the prices indicated in Table 3. Over 10 years, additional costs for about € 11'784.90 are accumulated. In order to have a breakeven, the price of DH from April to September should be less than 8€/MWh.

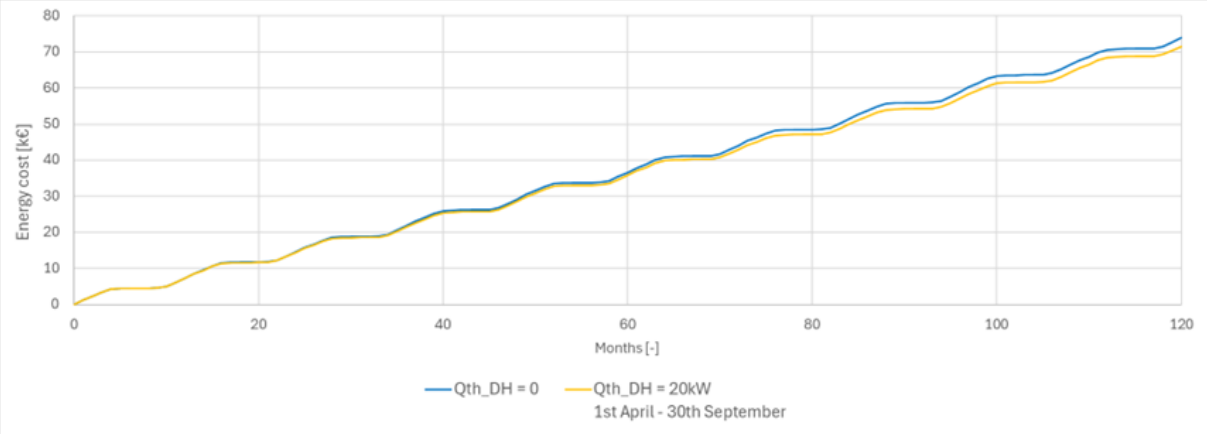

Figure 22. Cumulative cost of electrical energy (10 years)

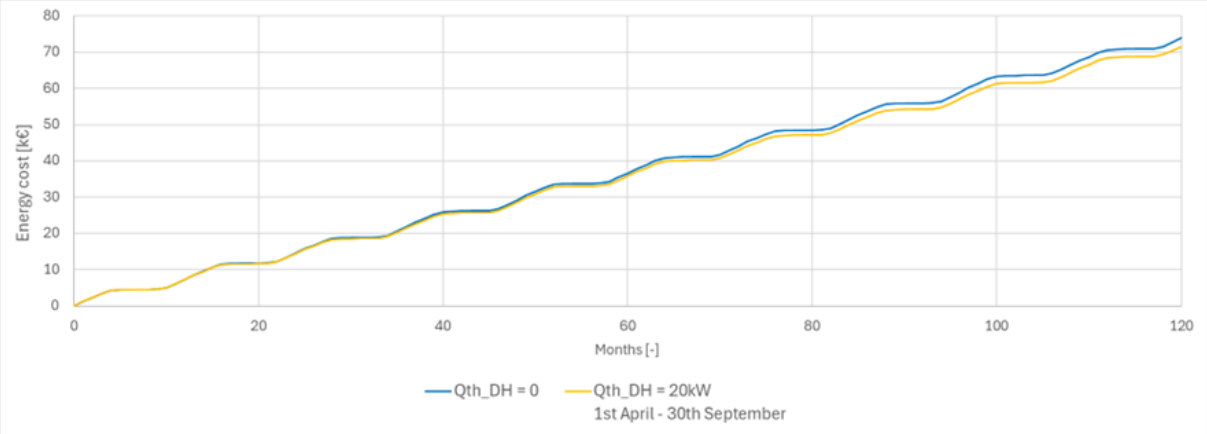

Figure 23. Cumulative cost of DH (10 years)

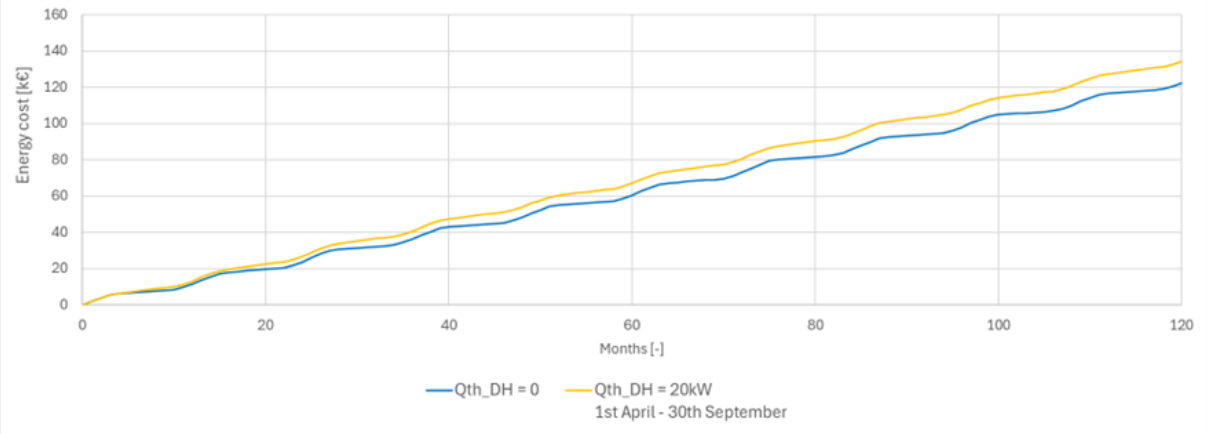

Figure 24. Cumulative cost of total energy (10 years)

Beside the economic analysis it is also relevant to consider environmental aspects, such as primary energy use CO_2 emissions. To start the analysis, it is crucial to define the primary energy and the CO_2 emission factors for district heating and for electricity:

- District heating primary energy factor → 0.07 [1]
- District heating CO_2 emission factor → 4.80 g/kWh [1]
- Electrical energy primary energy factor → 1.98 [2]
- Electrical energy CO_2 emission factor → 28 g/kWh [3]

Table 6 reports the primary energy consumption of the system without considering the BTES charge with district heating and considering the BTES charge with 20kW from district heating from the 1st of April to the 30th of September. In the primary energy calculation, the renewable primary energy taken from the BTES is not considered, only the primary energy and CO_2 emissions of electricity and district heating are considered.

Because the primary energy factor of district heating is much lower than the primary energy factor of electricity, the scenario in which district heating is used to charge the BTES presents a lower primary energy factor. As the decrease in electricity use is about 4.2 MWh/y, whilst the energy use from DH increases by 49.3 MWh/y, the primary energy balance is positive if the ration between primary energy factors is higher than 11.7, like in this case). On the opposite, since the ratio between CO_2 emission is about 6, CO_2 emissions increase in the advanced scenario.

Table 4. Primary energy consumption of the system without BTES charging and with BTES charging (20kW from April 1st to September 30th)

	Qth_DH = 0kW			Qth_DH = 20kW (1st April – 30th September)		
	Thermal energy from DH [MWh]	Electrical energy consumption [MWh]	Primary energy [MWh]	Thermal energy from DH [MWh]	Electrical energy consumption [MWh]	Primary energy [MWh]
January	15.4	14.5	29.8	15.0	13.6	28.0
February	10.6	14.7	29.8	10.6	14.2	28.8
March	8.6	14.8	29.9	7.9	15.0	30.2
April	1.5	11.8	23.5	6.1	11.8	23.8
May	7.2	4.8	10.0	15.7	4.8	10.5
June	12.2	0.5	1.8	22.3	0.5	2.6
July	12.0	0.6	1.9	21.9	0.6	2.6
August	11.5	0.5	1.8	21.8	0.5	2.5
Septemb	7.0	3.2	6.8	16.2	3.1	7.4
October	1.1	10.9	21.7	1.1	10.5	20.9
Novembe r	2.6	14.4	28.7	2.6	12.9	25.8
Decembe r	11.4	15.4	31.3	9.2	14.4	29.1
	101.1	106.0	217.0	150.4	101.8	212.1

Table 5. CO_2 emissions of the system without BTES charging and with BTES charging (20kW from April 1st to September 30th)

	Qth_DH = 0kW			Qth_DH = 20kW (1st April – 30th September)		
	Thermal energy from DH [MWh]	Electrical energy consump [MWh]	CO_2 emissions [kg]	Thermal energy from DH [MWh]	Electrical energy consumption [MWh]	CO_2 emissions [kg]
January	15.4	14.5	480.6	15.0	13.6	453.5
February	10.6	14.7	462.0	10.6	14.2	447.8
March	8.6	14.8	455.2	7.9	15.0	457.0
April	1.5	11.8	338.1	6.1	11.8	359.4
May	7.2	4.8	168.1	15.7	4.8	208.9
June	12.2	0.5	72.8	22.3	0.5	121.0
July	12.0	0.6	73.0	21.9	0.6	120.3
August	11.5	0.5	68.9	21.8	0.5	118.5
Septemb	7.0	3.2	122.5	16.2	3.1	165.9
October	1.1	10.9	311.5	1.1	10.5	299.3
Novembe r	2.6	14.4	416.1	2.6	12.9	374.0
Decembe r	11.4	15.4	485.5	9.2	14.4	446.3
	101.1	106.0	3454.1	150.4	101.8	3572.1

Table 6. Variation of energy mix, primary energy factor and CO_2 emission factor for electricity (Sweden, 2022 [3])

	Biomass [-]	Hydro [-]	Nuclear [-]	Wind [-]	Solar [-]	Wel PEF [-]	Wel CO_2 emission [g/kWh]
January	5.9%	44.2%	27.0%	22.9%	0.0%	2.12	30
February	5.9%	40.8%	29.3%	23.9%	0.1%	2.12	30
March	6.6%	39.1%	33.3%	20.6%	0.4%	2.19	31
April	6.3%	42.9%	33.6%	16.5%	0.7%	2.19	31
May	5.1%	41.8%	30.2%	21.9%	1.0%	1.98	28
June	3.6%	47.1%	32.1%	16.0%	1.2%	1.84	26
July	2.7%	45.3%	32.9%	18.0%	1.1%	1.70	24
August	3.0%	46.6%	34.1%	15.3%	1.0%	1.70	24
Septemb	4.2%	50.8%	29.0%	15.4%	0.6%	1.91	27
October	4.1%	38.7%	30.6%	26.4%	0.2%	1.77	25
Novemb	5.6%	46.8%	30.0%	17.6%	0.0%	2.12	30
Decemb	8.1%	44.2%	26.9%	20.8%	0.0%	2.48	35

Table 7. Hypothetical variation primary energy factor and CO_2 emission factor for Helsingborg's district heating

	DH PEF [-]	DH CO_2 emission [g/kWh]
January	0.11	7.20
February	0.11	7.20
March	0.11	7.20
April	0.04	2.40
May	0.04	2.40
June	0.04	2.40
July	0.04	2.40
August	0.04	2.40
September	0.04	2.40
October	0.11	7.20
November	0.11	7.20
December	0.11	7.20

The above analysis considers annual average values for the primary energy and CO_2 emission factors. In reality, these change over the year because the energy mix varies. Table 6 shows the variation of the energy mix, the primary energy factor and the CO_2 emission factor for electricity in Sweden for the year 2022.

The monthly values of primary energy and CO_2 emission factors are not available for Helsingborg's, DHN, however, it is reasonable to assume that they would be higher in cooler months and lower in warmer months. This is due to the excess of renewable energy sources (RES) and lower loads in warmer months, which reduce the need for fossil-fuelled generators. To consider this scenario the primary energy and CO_2 emission factors were considered three times higher in the cooler months if compared to the warmer months, keeping the same annual average value (see Table 7).

The primary energy and CO_2 emissions were recalculated for the Tornet Energy Centre considering these new monthly values, and the updated results reported in Table 8 and 9.

Since both the thermal energy from district heating and the electrical energy consumption are mostly concentrated in the cooler months, the updated values for primary energy and CO_2 emission factors result in higher annual primary energy consumption and CO_2 emissions compared to the previous scenario. However, it is noticeable that with these updated values, the primary energy consumption and CO_2 emissions are lower when the BTES is charged with district heating during summer, compared to when the BTES is not charged.

Table 8. Primary energy consumption of the system without BTES charging and with BTES charging (20kW from April 1st to September 30th) with the new monthly data for primary energy factor

	Qth_DH = 0kW			Qth_DH = 20kW (1st April – 30th September)		
	Thermal energy from DH [MWh]	Electrical energy consumption [MWh]	Primary energy [MWh]	Thermal energy from DH [MWh]	Electrical energy consumption [MWh]	Primary energy [MWh]
January	15.4	14.5	32.5	15.0	13.6	30.6
February	10.6	14.7	32.3	10.6	14.2	31.2
March	8.6	14.8	33.4	7.9	15.0	33.7
April	1.5	11.8	26.0	6.1	11.8	26.1
May	7.2	4.8	10.1	15.7	4.8	10.9
June	12.2	0.5	2.0	22.3	0.5	2.9
July	12.0	0.6	2.0	21.9	0.6	2.9
August	11.5	0.5	1.9	21.8	0.5	2.8
Septemb	7.0	3.2	6.7	16.2	3.1	7.5
October	1.1	10.9	19.5	1.1	10.5	18.7
Novemb	2.6	14.4	30.9	2.6	12.9	27.7
Decemb	11.4	15.4	39.3	9.2	14.4	36.6
	101.1	106.0	236.5	150.4	101.8	231.4

Table 9. CO_2 emissions of the system without BTES charging and with BTES charging (20kW from April 1st to September 30th) with the new monthly data for CO_2 emission factor

	Qth_DH = 0kW			Qth_DH = 20kW (1st April – 30th September)		
	Thermal energy from DH [MWh]	Electrical energy consumption [MWh]	CO_2 emissions [kg]	Thermal energy from DH [MWh]	Electrical energy consumption [MWh]	CO_2 emissions [kg]
January	15.4	14.5	546.6	15.0	13.6	516.7
February	10.6	14.7	516.9	10.6	14.2	501.7
March	8.6	14.8	520.2	7.9	15.0	520.8
April	1.5	11.8	369.8	6.1	11.8	380.1
May	7.2	4.8	150.9	15.7	4.8	171.2
June	12.2	0.5	42.4	22.3	0.5	66.5
July	12.0	0.6	42.0	21.9	0.6	65.7
August	11.5	0.5	39.3	21.8	0.5	64.1
Septemb	7.0	3.2	102.5	16.2	3.1	123.9
October	1.1	10.9	281.3	1.1	10.5	270.5
Novemb	2.6	14.4	451.1	2.6	12.9	406.0
Decemb	11.4	15.4	620.6	9.2	14.4	569.0
	101.1	106.0	3683.5	150.4	101.8	3656.1

3.5 Lessons learned

The above assessment brought to the following key takeaways:

- Thermal energy storage tank and borehole thermal energy storage: The incorporation of these components provides substantial benefits in balancing supply and demand by storing excess thermal energy during low-demand periods and utilizing it during peak demand.
- Advanced monitoring and control system: The deployment of a sophisticated monitoring and control system, incorporating sensors, programmable logic controllers, and a SCADA system, ensures optimal performance and reliability. This system improves energy efficiency and user comfort, while also supporting fault detection and diagnostics, enabling proactive maintenance and system optimisation.
- BTES loading strategy: Different control strategies could be used to charge the BTES during summer. With high thermal power and short charging periods, the stored energy is not sufficient to charge the BTES completely because the temperature at the outlet of the BTES exceeds the imposed limit very often. Exchanging 20kW between the DH and the BTES from April to September, limiting the BTES outlet temperature at 15°C, it is possible to store enough energy in the BTES to cover the energy extraction during the heating season.
- Cost savings: This charging strategy allows to reduce the electricity consumption over the years. At the same time more energy should be imported from district heating to charge the BTES. After 10 years the electricity consumption is 4% lower but the district heating consumption is 33% higher. Considering a 10-years span, to reach the financial breakeven the district heating price between April and September cannot exceed 8 €/MWh.
- Primary Energy savings: From an energetic point of view, charging of the BTES has a positive impact if the ratio between the primary energy factor of electricity and the primary energy factor of district heating is higher than the ratio between the energy taken from district heating to charge the BTES and the electrical energy savings.

$$\frac{PEF_{Wel}}{PEF_{DH}} \geq \frac{\Delta Qth\_DH}{\Delta Wel}$$

- CO_2 emissions savings: The same reasoning is valid for the environmental point of view. To have a positive impact, the ratio between the CO_2 emission factors of electricity and district heating should higher than the ratio between the energy taken from district heating to charge the BTES and the electrical energy savings.

$$\frac{CO2_{Wel}}{CO2_{DH}} \geq \frac{\Delta Qth\_DH}{\Delta Wel}$$

- Boreholes as a storage: Overall, boreholes can be used as a storage under prescribed conditions. However, system control must be elaborated specifically and in collaboration with DH and electric grid authorities, based on variable primary energy use and CO_2 emissions associated to the energy vectors. Additionally, specific vectors' prices must be set at energy system level, reflecting the renewable energy content in the specific vector and period of the year.

We conducted a breakeven assessment using informed assumptions for district heating energy prices and primary energy factors, which yielded only marginal benefits for the energy system overall. To enhance performance, a more detailed analysis at the hourly level is necessary, along with implementing active energy centre control. This approach would better leverage the system's flexibility and maximise benefits at the energy system level.

4 La Seyne-sur-Mer (DALKIA-EDF)

4.1 Description of the demonstration site

The demonstration site located in La Seyne-sur-Mer is a DHC network operated at different temperature levels to accommodate different needs encountered along the expansion. The DHC network is operational since 2008, while DALKIA took over ownership and management starting 2019. The network was extended along the project elaboration, as shown in the map of Figure 25, and the number of customers raised from 4 to 14 in 2024.

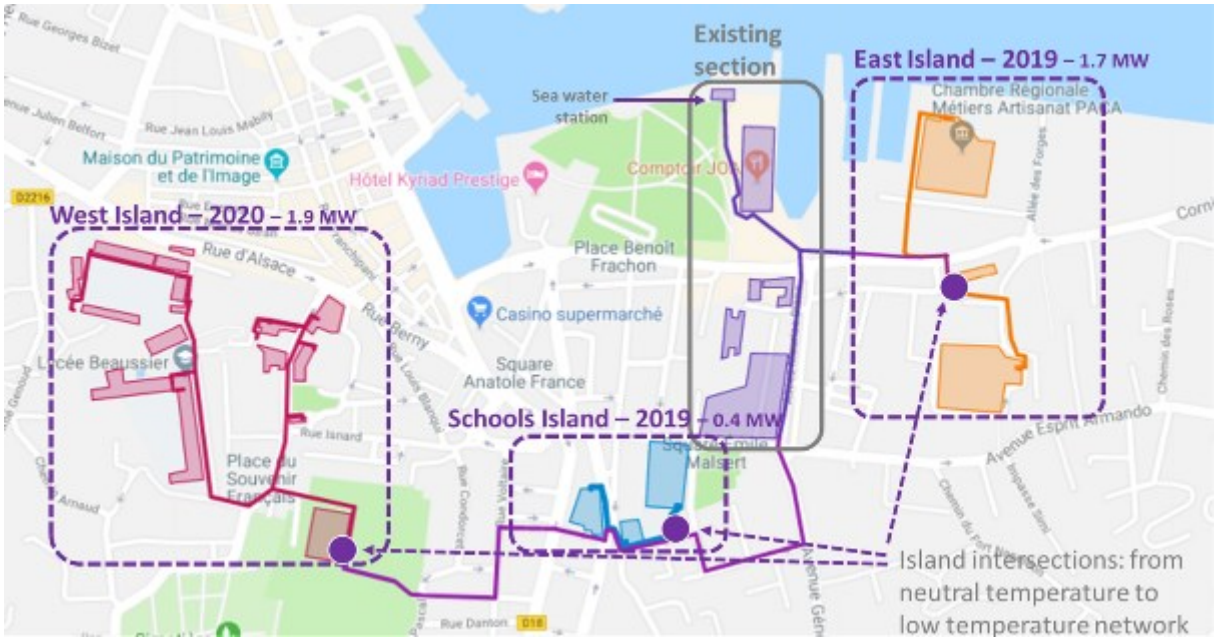

Figure 25: Expansion at the demonstration site in La Seyne-sur-Mer.

The neutral-temperature DHCN initially set up uses seawater as energy source and sink, and allows to cover both heating and cooling needs, as it is operated between 13-24°C, the temperature varying over the year according to the seawater temperature and the extent of heating and cooling, partially loads balancing out over the network. The buildings connected to this network are set up with substations exploiting water-to-water heat pumps to draw or reject thermal energy from/into the network.

The networks stemming from the core one are conceived as semi-decentralised ones, as substations integrating large HPs connect the extensions, as represented in Figure 26. Depending on the energy uses of the buildings integrated in the subnetworks, those can be operated in heating mode only or provide both heating and cooling.

While SH and, eventually, SC loads are consistently covered by the DHCN, DHW preparation is addressed occasionally. In most cases, the secondary network operator ensures the DHW preparation via existing or refurbished means. In 3 substations, gas boilers are still present and falling under the public delegation of service contract of DALKIA and are used basically as back-up: in HLM PRESENTATION (residential building) gas in used as back-up or when the gas is cheaper than electricity, taking COP of the systems into consideration; in School Malsert gas is only used as back-up; in School Jaurès gas is used for heating and DHW as HPs are sub-dimensioned and do not reach the needed set temperature, and are thus used for preheating the return flow of the boiler.

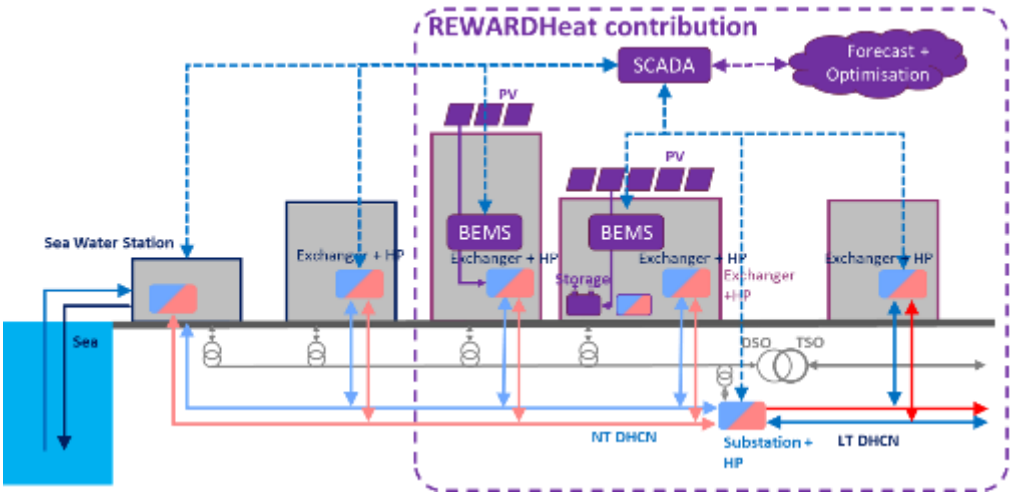

Figure 26: Overview of the technologies that will be implemented at the demo case of Toulon.

4.1.1 Objectives of the demonstration activity

The project involves two partners: DALKIA, who owns and manages the DHCN, and EDF who developed the advanced control tools and implemented them in DALKIA's platforms. Overall, the project aimed to improve the energy efficiency of the DHCN through new equipment installation and advanced control strategies integration. Particularly, the following objectives have been pursued:

- Replacement of the seawater filters and installation of variable speed pumps at the central pumping station in order to reduce flow, raise the temperature difference between supply and return temperature, hence lower electricity consumptions overall;
- Upgrade of the PLCs and the functional analysis implemented at the central pumping station and all customer substations. These have been improved continuously along the project elaboration, thanks to the feedback received while modelling or simulating the network substations for forecast, optimisation or simulation purposes;
- Establishment of a reliable 4G communication link between substations and the SCADA system, along with an upgrade of the entire DALKIA's ICT infrastructure to enhance the reliability and stability of data flows;
- Development and deployment of a centralised, advanced supervision and control system based on model predictive based optimisation techniques.

4.2 Formulation of the control problem

Neutral-temperature DHCNs distributing thermal energy with a low temperature difference between supply and return pipe (i.e., ΔT=10-15°C), require particular attention to minimising pumping costs both for sourcing and distributing energy, as these can considerably affect the overall operation costs. Similarly, management strategies are needed to reduce the electricity costs associated with operating HPs that interface the DHCN with customers and connect different branches of the network.

These matters can be tackled with supply temperatures and thermal storages management strategies that adapt to the specific context, climate and seasonal operating conditions, benefitting of the "best" combinations of RE (both thermal and electric) and Waste Heat encountered: network

temperatures can be adapted based on the ratio between heating and cooling loads insisting on the system; supply temperature can slide down when an excess of renewable electricity is available to run heat pumps covering thermal loads and/or charging thermal storages. Conversely, heat stored at higher temperatures is more effectively utilised to cover loads during peak hours on the electric grid.

If such an active management would have been unbearable 5 to 10 years ago, digital hardware incrementally more affordable and control platforms based on AI that start proving reliable, started to be game changers.

Due to the system's thermal and hydraulic inertia, active management cannot be performed using conventional feedback loop-based RBCs. Instead, it requires awareness of weather conditions and the ability to predict future loads on the network over an appropriate time horizon. To achieve this, EDF and Dalkia decided to implement an optimisation solution based on the MPC approach.

The overall scope of the MPC development has been to reduce the OPEX of the DHCN, by reducing electricity expenditures. Specifically, the objectives pursued by developing and implementing the MPC solution in La Seyne-sur-Mer have been:

- Validate the developed MPC in terms of forecasting and optimisation;
- Balancing heating and cooling loads insisting on the DHCN as much as possible to reduce seawater pumping costs;
- Delay the sea-water pumps start, letting the network deviate slightly in terms of temperature, when forecast/optimisation results show that it will re-equilibrate in the short term;
- Using the network as a thermal energy storage itself, profiting of the thermal mass “freely” available, enabling to balance loads among clients.

4.3 Description of the control implemented

4.3.1 SCADA system

The main metering and automation backbone of the network level has been realized based on a commercial system called “WIT”. The advantage of this solution compared to other market available systems is that it can be operated via ADSL 4G technology. The WIT solution chosen is called “READY Process” and being modular, enables the integration of different controllers in a hierarchical way, based on the number of connections to be handled.

The complete system architecture is represented in Figure 27. It shows the base of the system - the installed WIT system - responsible for assuring the correct operation of the DHC network. This is interfaced by the Regional Control Centre, using the generated data for monitoring and other O&M services for DALKIA. On top of these, the DEMix platform, responsible for the advanced control implementation, communicates bidirectionally with the WIT system through the Regional Control Centre.

The main supervision centre is located in the Armada-Santa Maria substation and is equipped with a “READY XL” controller (capable of handling up to 2500 connections), the pumping-station which has a “READY M” controller, while all other substations are equipped with “READY S” controllers.

As can be evinced from Figure 27, between the upper smart control level and the automation system on-site, the data and order flow pass via the Regional Control Centre (RCC) of DALKIA. Key figures concerning the installed metering and automation equipment are:

- 11 WIT READY modules;
- 330 data points measured every 10 min;
- 10 thermal energy meters with remote reading;
- 25 electric meters with remote reading: 16 installed for the project (9 smart meters and 7 Socomec metering stations) and 9 smart meters;
- 88 temperature sensors;
- 1 “delta P” meter;
- All heat pumps are controlled (i.e. temperature, pressure or alarms) via Modbus RTU and Bacnet MS/TP;
- 10 V2V on the low-temperature loop– Modbus compatible valves, enabling to read the position and flow rate.

The available historic data on heating/cooling power and energy values have been harmonized on by DALKIA and have been shared with EDF. In parallel, the raw data streaming from DALKIA’s Regional Control Centre towards EDF.

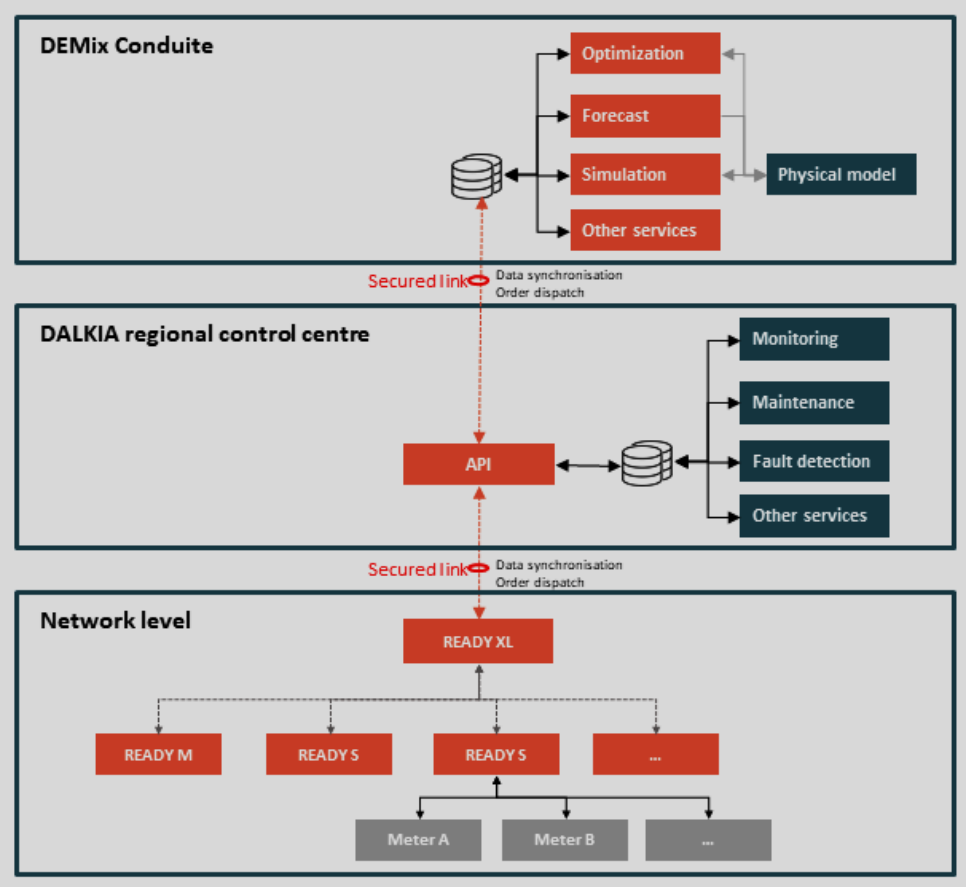

Figure 27 – Simplified overall SCADA architecture for the smart control of the LSSM DHCN. Source: EDF.

4.3.2 Advanced Controls

The smart control system developed is based on the platform under development in EDF, called DEMIX®, and uses a specific module called "DEMIX Conduite". This is aimed to become a centralised platform for monitoring data acquisition, load forecasting, and optimising and dispatching production and storage units (24h ahead) at any DHCN operated by the EDF Group.

The current development of DEMIX is based on the operation of traditional centralized networks (one centralized energy station), hence not fully adapted to handle a decentralized network as the one concerned in this project.

Thermal loads forecast

The first step towards a smart control system for a DHCN exploitation system is the capability to foresee load conditions. The module used for this in the DEMIX platform is the ForecastHeat® software developed by EDF. This software is based on AI algorithms, learning from past aggregated heating loads insisting on the DHN and weather information to forecast the future aggregated loads. It allows to identify the loads on different time horizons (day, week) and use the following input data:

- weather data as can be retrieved from the services of Meteo France (forecasts or historic data);
- site specific data as past consumptions or local outside temperature;
- calendar data as time, day of the week or events related to specific periods of the year.

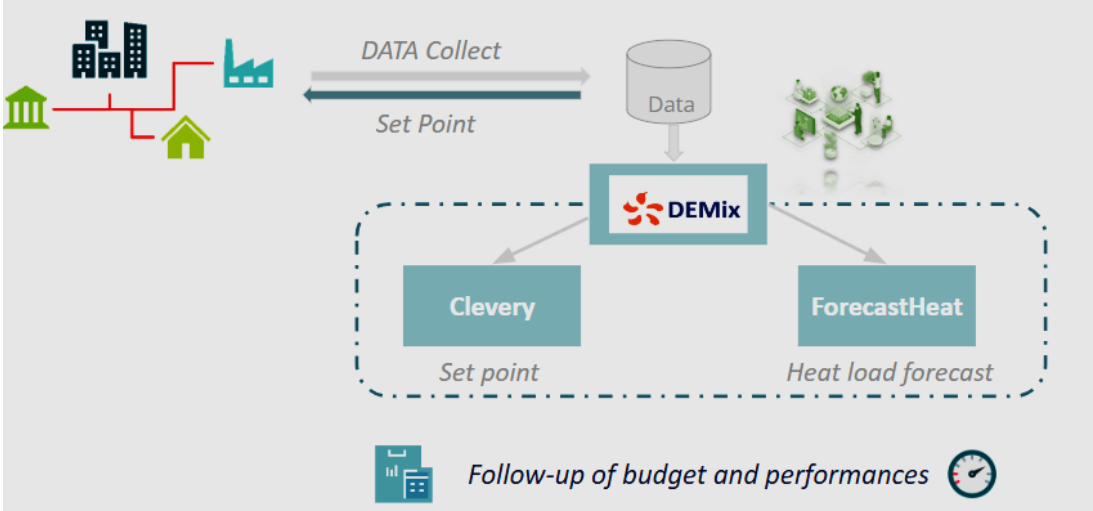

Figure 28: schematic working principle of the DEMix platform: real-time data is pushed to the platform, using these data for load forecast and subsequent optimisation, sending back to the site the calculated set points. Source: DALKIA.

Along the project elaboration this module has been upgraded to gather data from distributed substations and to forecast heating loads at the single customers. DALKIA has generated a 1-year dataset of heating and cooling loads, with 10 min time step for each thermal load at the substation level, and 14 forecast models have been set up and coupled with the DEMIX Conduite platform. These have been trained accordingly with the 1-year datasets (Figure 30).

The work has been focused on setting up the system and ensuring data consistency and quality, resulting in reliable and accurate enough predictions.

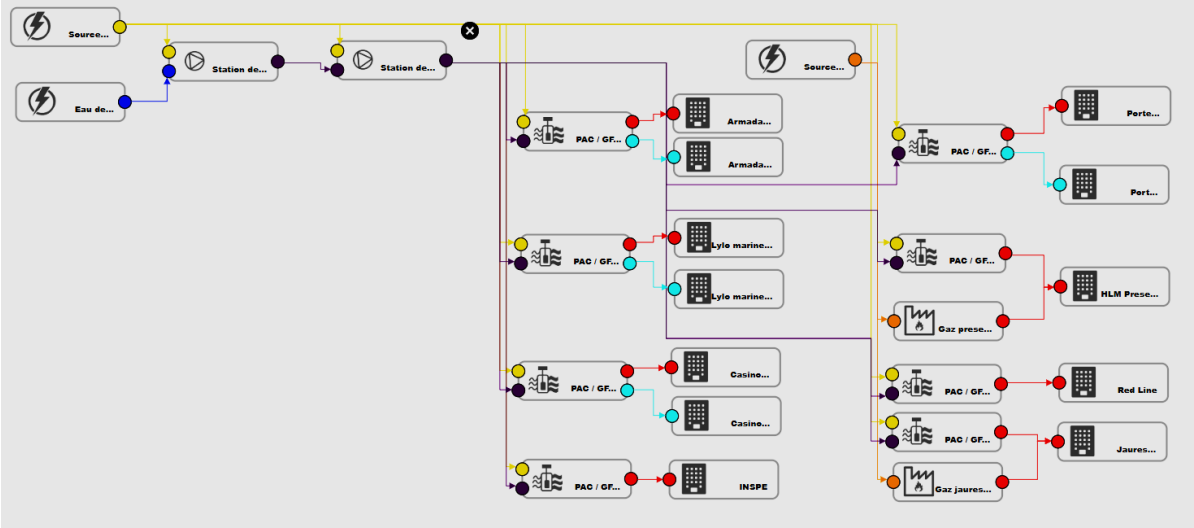

Figure 29 : screen shot of the forecast modules modelled in DEMIX Conduite as by the new industrialized temp/flow configuration – red indicated heating forecast and blue the cooling forecast. Source: DALKIA.

Additionally, cooling loads forecast capabilities have been implemented: cooling loads forecast needs additional information (predictors) compared to heating, e.g., calendar-related information and occupancy rates.

Cooling loads are produced mostly at the biggest load integrated in the network, the Casino. A dialogue with the Casino operator has been launched and daily historic occupancy data of its main spaces (casino and restaurants) have been retrieved for several years. These have been used to test the forecast accuracy of the model and it has been shown that the model can be calibrated to the data and provide accurate forecast. Hence, expected occupancy rates have been entailed among the predictors time series used for forecast model training.

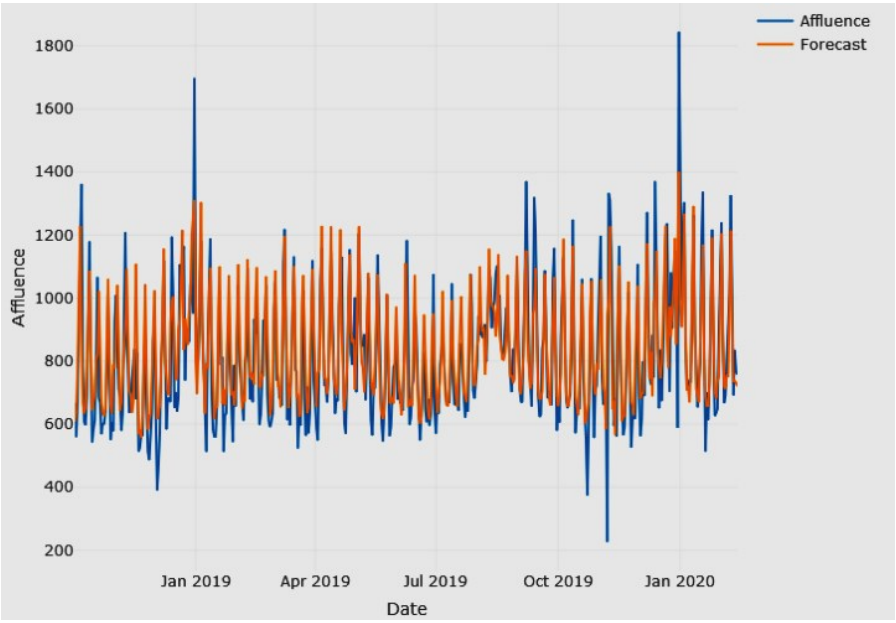

Figure 30 : comparison of forecasted (orange) and real (blue) occupancy data of the Casino. As can be evinced, except singular outlier events, the forecast matches well the real data. Source: EDF.

During the development process, additional predictors, e.g., wind speed, humidity, irradiation and other weather data sources have been explored. The model used has proven to be robust, with no

noticeable improvements observed. One potential area for enhancement identified is the use of different models or methodologies for the same prediction task. However, this approach cannot be pursued further, as it would significantly affect the requirements and usability for the end-user of the industrialized version of DEMIX.

Optimisation module

Forecasted data generated feed an optimisation tool aimed to retrieve the needed set points and dispatching plan for each system unit, over the chosen receding time horizon. EDF has developed a tool called Clevery®, designed to model and optimize the daily operations of multi-energy systems. This has been conceived to be a decision-making support tool during exploitation, exploring techno-economic criteria (as costs of assets or their coefficient of performance) for systems operation optimisation.

Before the project began, the tool integrated a library of system components based on the MILP approach, allowing for the modelling of energy flows across the system. However, this approach has significant limitations in calculating temperatures, flow rates, and pressures, which are crucial in neutral-temperature DHCNs where the performance of distributed heat pumps and circulation pumps is strongly affected by these parameters. An optimisation solver operating under constrained conditions was used on the overall system model to minimise an objective function.

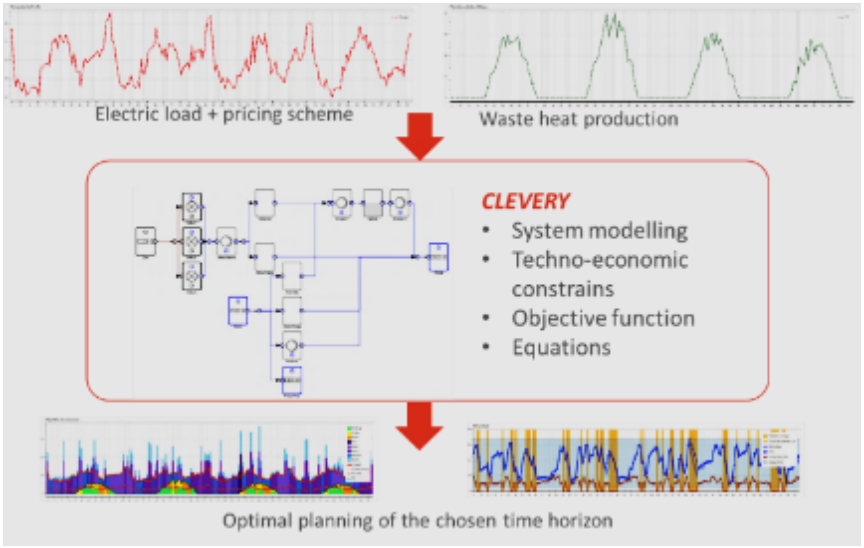

Figure 31: Schematisation of the functioning principle of Clevery and visualisation of its physical modelling interface. Source: EDF.

To overcome this limitation, upgrading the individual system component models was addressed as part of the project activities. Initial studies on the nonlinear optimisation of DHNs yielded inconclusive results due to convergence issues and very long computation times, making it impractical to study large-scale systems. This is because some fundamental physical equations, such as those for heat losses, are highly nonlinear and cannot be incorporated into a MILP model.

To address this problem, an iterative process combining detailed simulation and optimisation was tested, as simplified in Figure 32. This process involves the following steps:

- Initialisation: Boundary conditions are set.
- Detailed Simulation: Performed using physical models of the system components, developed in Modelica and MixSysPro, to determine the current system dynamics.

- Optimisation: The output from the simulation serves as the input for the optimisation process, resulting in a scheduling strategy for managing the DHCN.

If a convergence criterion is met (meaning the values of variables in the optimisation and simulation are very similar), the solution is accepted, even if it is likely suboptimal. If not, the optimised set points are used to initialise a new detailed simulation. This iterative process is repeated until convergence occurs for the agreed control variables, as seen in Figure 33.

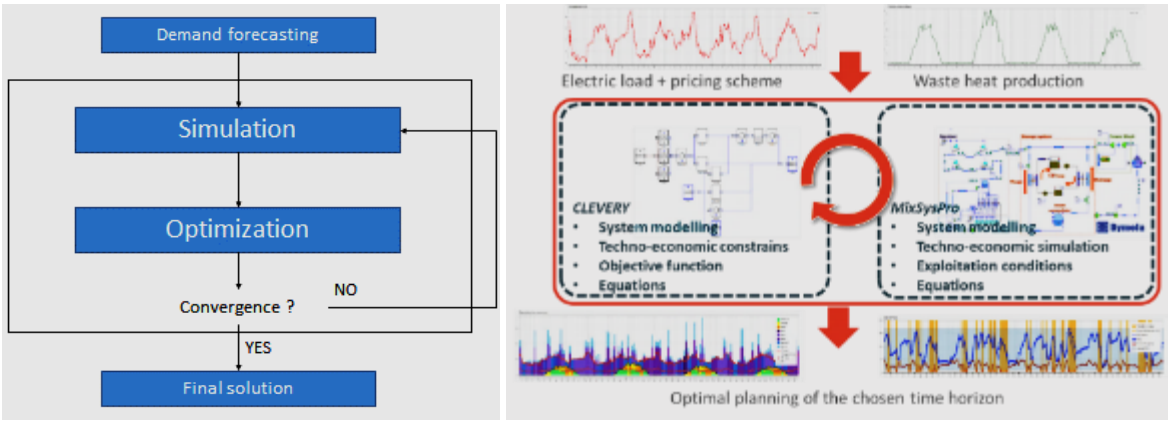

Figure 32: Schematisation of the iterative method proposed. Source: EDF.

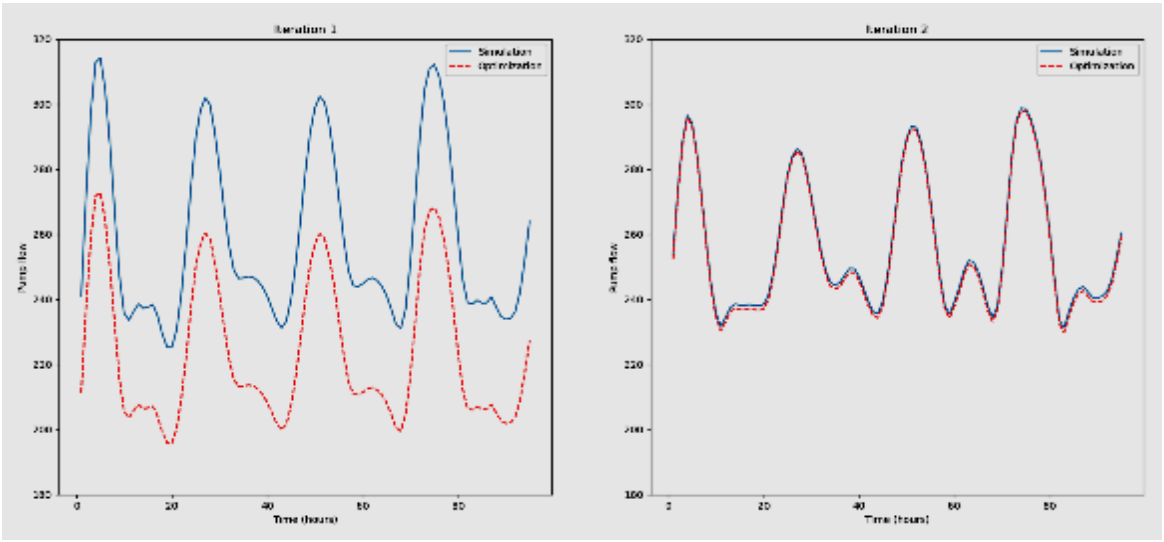

Figure 33: Example of convergence of the iterative method. Source: EDF.

The process described is complex and requires the implementation and calibration of a detailed physical model of the system. The objective of the project was to develop a generic framework to model and optimise the operation of the DHCN without requiring extensive preparation work.

The modelling of the DHCN's components, particularly heat pumps, has been adapted from a "generic behaviour" to the specific hardware installed on site, resulting in more realistic modelling. The interface between the DEMIX platform and the new physical module is developed in a Python framework. This interface was initially tested on a simplified generic network model and then applied to the complete DHC network of LSSM under operational conditions, using the model developed in Dymola-Modelica.

The advantage of Dymola-Modelica-based modelling is that the assets' behaviours are described as systems of physical equations, whether linear or not, and the system simulation is solved depending on the given simulation environment and its temporal or physical constraints. This approach maintains a white-box methodology, preserving the cause-and-effect relationships and enabling the explanation of achieved results, which cannot be ensured in grey or black-box approaches. This method has indeed enabled the identification of several faults or malfunctions in terms of operation, design, or control in the DHCN, through comparisons of real data with forecast and optimisation results.

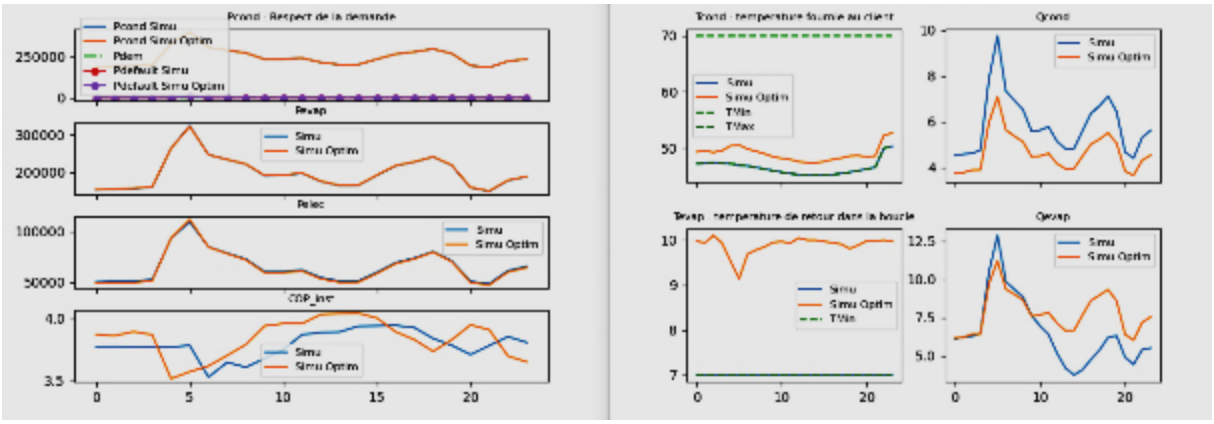

Figure 34 – Example of data output from the optimisation comparison – in this screenshot, the model outputs are compared to actual operation conditions considering the network temperatures at the customer side (delivery/return). As in this case, whilst the COP of individual HP might suffer for some degradation, the overall DHCN's electricity demand is lowered, as less pumping on the network side is needed. Source: EDF.

For what concerns optimisation techniques, the project, along with the dedicated organised exchanges and workshops among partners, has enabled a shift from a linear approach to a non-linear one. It is important to consider that previous solvers were linear, necessitating the linearisation of the entire DHCN, a very time-consuming process that reduces model accuracy. To address this, "mixed integer linear" solvers (MILP) were compared to more computationally intensive "mixed integer non-linear" solvers (MINLP), which account for the non-linear behaviours in the system.

Table 10 : Summary of generic results from the comparison between the linear and non-linear solvers. Source: EDF

	Nonlinear (MINLP) – Knitro solver	Linear (MILP) – Xpress Solver
Modelling	+	- - - All equations must be linearized
Readability	+ Equations can be directly transposed from the physical models	-
Calculation speed	+ + Less than 1 min in average for 24h	- - The optimisation has difficulties to prove that the « best solution » has been achieved
Reliability of results	- Problems with local minima, partially solved by *multi-starts*	+ Gaps are sometimes important, not enabling to guarantee a solution
EDF knowledge on the software	- Work in progress!	+ + Much experience on the subject

The comparison accounted for output quality, programming complexity, calculation duration and associated costs. Solvers vary widely in price and a cost-benefit evaluation (considering software license costs and benefits in terms of system efficiency and performance) is crucial for the replication and industrialisation plan (see Table 10).

The results of this work have been shared with EDF and DALKIA, who decided to adopt the non-linear/MINLP optimisation due to its simplicity of implementation and efficiency. The Modelica/MixSysPro/DEMix Conception models include two different system models: one in Dymola-Modelica and one in a MINLP formulation. The latter is challenging to extend to other networks and industrialise in a generic framework. One of the main advantages of the chosen approach is that using a non-linear solver eliminates the need for linearising physical models, while providing robust results in negligible time.

This new MINLP formulation led to a complete overhaul of the functional and technical requirements for the DEMix Conduite platform and the optimisation module. Once verified, these enhancements from the temperature/flow optimisation are transferred to the industrialised platform of DEMix Conduite, along with new requirements for human-machine interfaces (HMI) needed for additional monitoring of the flow/temperature approach (as shown in Figure 40, a screenshot of the newly developed GUI). Meanwhile, the optimisation core is still undergoing verification and continuous improvement as the team gains knowledge in running a non-linear optimisation environment. The solver itself (Knitro-Solver) is also being updated to become more efficient and reliable.

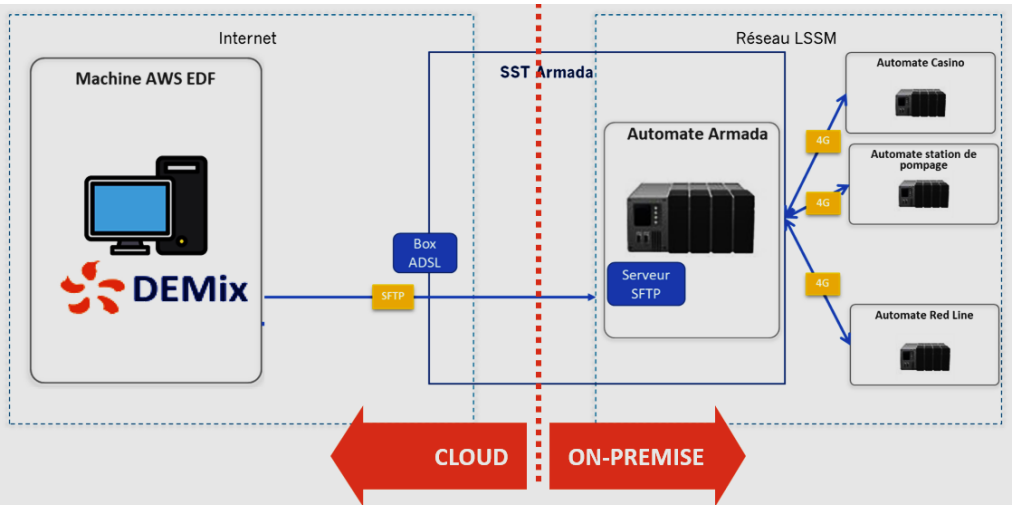

Figure 35 : architecture of the DEMix platform, as commissioned for the experimentation. Source: Dalkia.

4.4 Assessment of the control performance

Thermal loads forecast

Acceptable results have been achieved for the overall load forecast error, maintaining an average below 12% on an annual basis. However, results vary significantly among substations, with some exhibiting unpredictable or chaotic behaviour. Thanks to the commissioning work of DEMix, some of these erratic behaviours could be corrected, as they were related to metering or data aggregation issues, or specificities of the secondary network regulation or system design. ForecastHeat has been installed on the EDF server since November 2022, as shown in Figure 35. Some examples of good prediction are shown in Figure 36 and Figure 37; and some examples of bad forecasting are shown in Figure 38 and Figure 39.

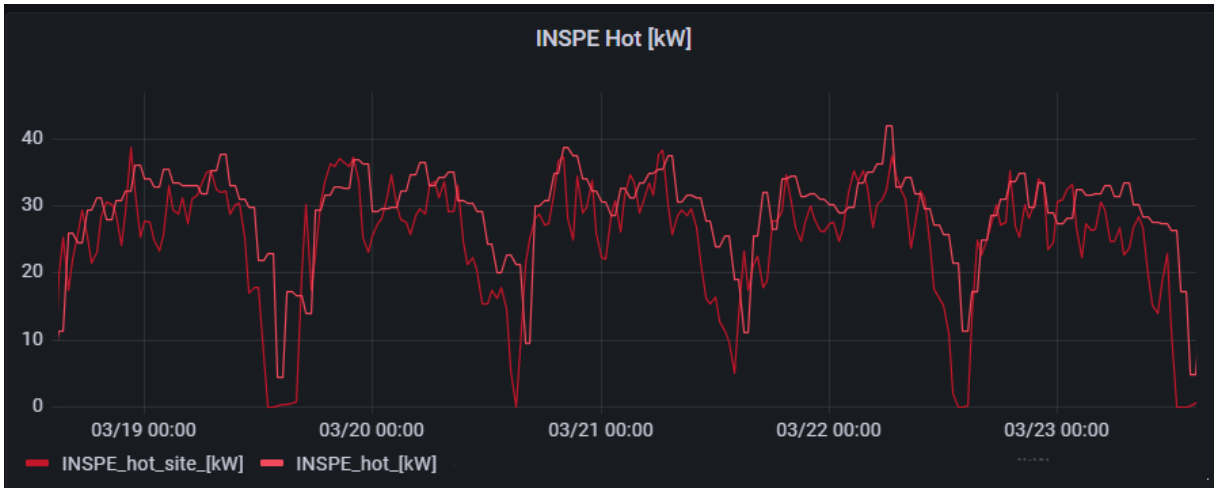

Figure 36 : Comparison of forecast (pink) and real (red) hot consumption of the INSPE substation, which is a low-demand substation. Source: EDF.

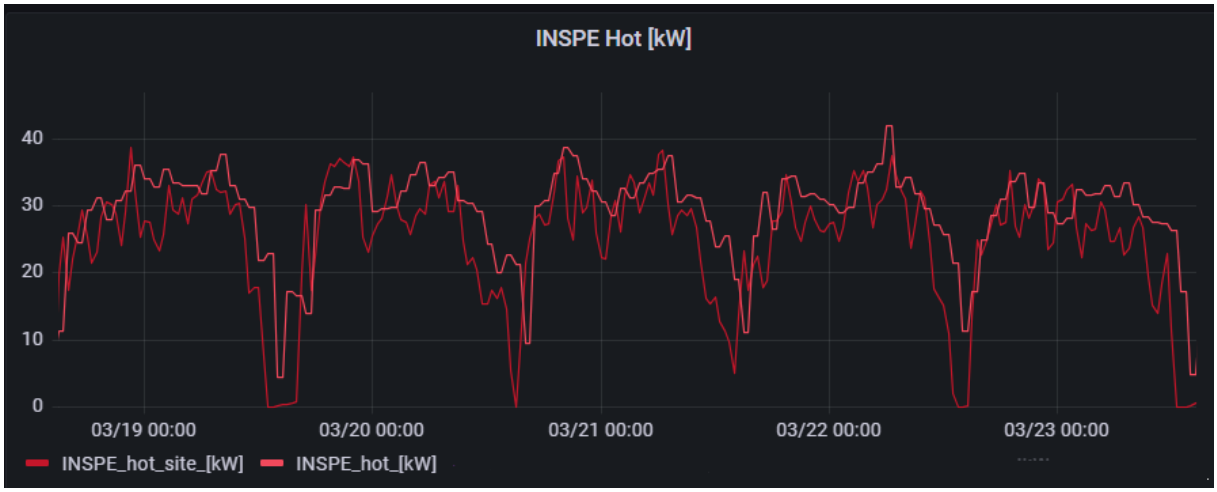

Figure 37 : Comparison of forecast (pink) and real (red) cold consumption of the Casino, which is a high-demand substation. Source: EDF.

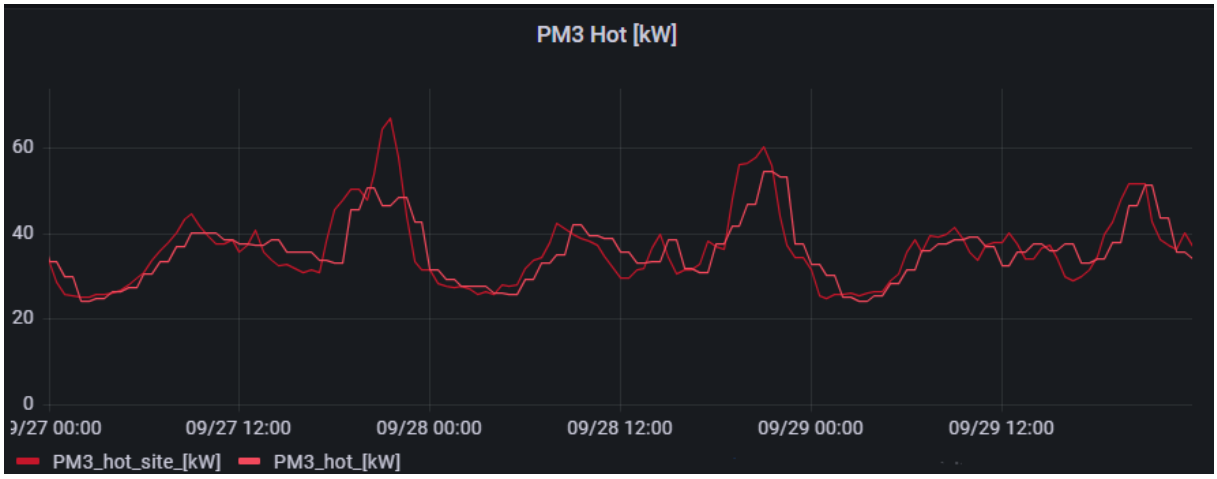

Figure 38 : Comparison of forecast (pink) and real (red) hot consumption of the PM3 substation. Example of bad forecasting: unpredictable consumption peak. Source: EDF.

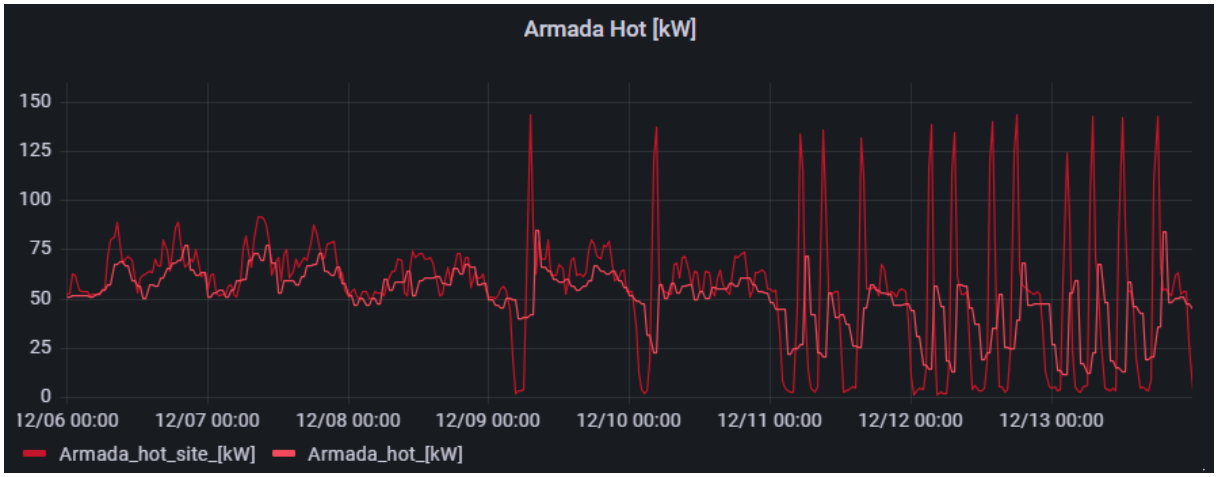

Figure 39 : Comparison of forecast (pink) and real (red) hot consumption of the Armada substation. Example of bad forecasting: abrupt changes in demand behaviour. Source: EDF.

In terms of both heating and cooling forecasts, the overall forecasting capabilities are appropriate for the optimisation problem formulation. Indeed, one of the valuable, yet unexpected, repercussions of the forecast-related control is that forecast verification has prompted a re-evaluation of the observed behaviours of assets. This has led to significant improvements in system operation (SCADA) and related automation. More details on this are provided in the following sections.

Optimisation module

RewardHeat has enabled EDF and DALKIA to gain a much deeper understanding of the optimisation of a low-temperature (LT) network and to develop the DEMix platform in its various components (hypervision, forecast, optimisation). For this experimentation, the optimisation part of the software was initially developed in Python due to its simplicity of implementation. This is currently being transferred to a more industrialised version in C++, as shown in the following figures. This work is still in progress and will continue, among other initiatives, in the European project Senergy-Nets. For larger replication, different specific assets and related modules/models, such as storage, still need development.

The original DEMix platform has been expanded in terms of technical and functional requirements and industrialised for the hypervision of a targeted asset, as shown below. The new temperature/flow optimisation interface enables navigation of individual assets and provides much richer monitoring of the assets, related to the decentralised approach implemented.

Specifically for the project, besides the Grafana interface developed to monitor all individual variables of the optimisation module, a dedicated section has been developed in DEMix to follow the orders sent by DEMix to the local SCADA system of LSSM (loop temperature setpoint). This led to the development of specific visuals to compare the sent setpoints with the values applied on site by the SCADA for the operator of the DHCN (see Figure 41).

Figure 40 : scrren shot of the main hypervision view of the industrialized DEMix platform and adapted the needs of the temp/flow optimisation developments as implemented first in LSSM. Source: DALKIA.

Figure 41: as by the new industrialized temp/flow configuration, in order to follow the loop temperature setpoint sent by Clevery, specific visuals have been integrated to compare sent setpoints, with the values applied on site by the SCADA. Source: DALKIA.

Independently of the optimisation environment, the objective function is the operational optimisation of the network. Preliminary commissioning campaigns have been performed to qualify and quantify the optimisation tool. To this end, different approaches have been simulated and compared, i.e.:

- Minimisation of total electricity consumption (in kWh) compared to minimisation of the total electricity costs (in €): a result of this comparison is shown for the Lylo-Marine substation in the figure below, where prices at the pumping station are higher than prices at the substation. When minimising the total electricity cost in €, the solution proposed by Clevery consists in increasing the electrical power of the HP (even if this means reducing the COP slightly) and reducing the flow rate between the substation and the main loop. However, this conclusion cannot be generalised, as it depends on prices and consumption at each time step.
- Maximisation of the overall HPs' efficiency (COP) compared to minimisation of the electricity consumption of head-pumps: the figure below shows an example of this comparison where maximising the overall HP's efficiency results in a higher flow rate in the network loop;
- Minimisation of electrical power demand to identify power subscription reduction potential (peak shaving capability): as shown in the figure below, this optimisation potential is limited under the given technical and operational constrains, hence is not been pursued.

These different objective functions have been tested offline, and it has been finally decided to adopt minimizing the total electricity costs in €.

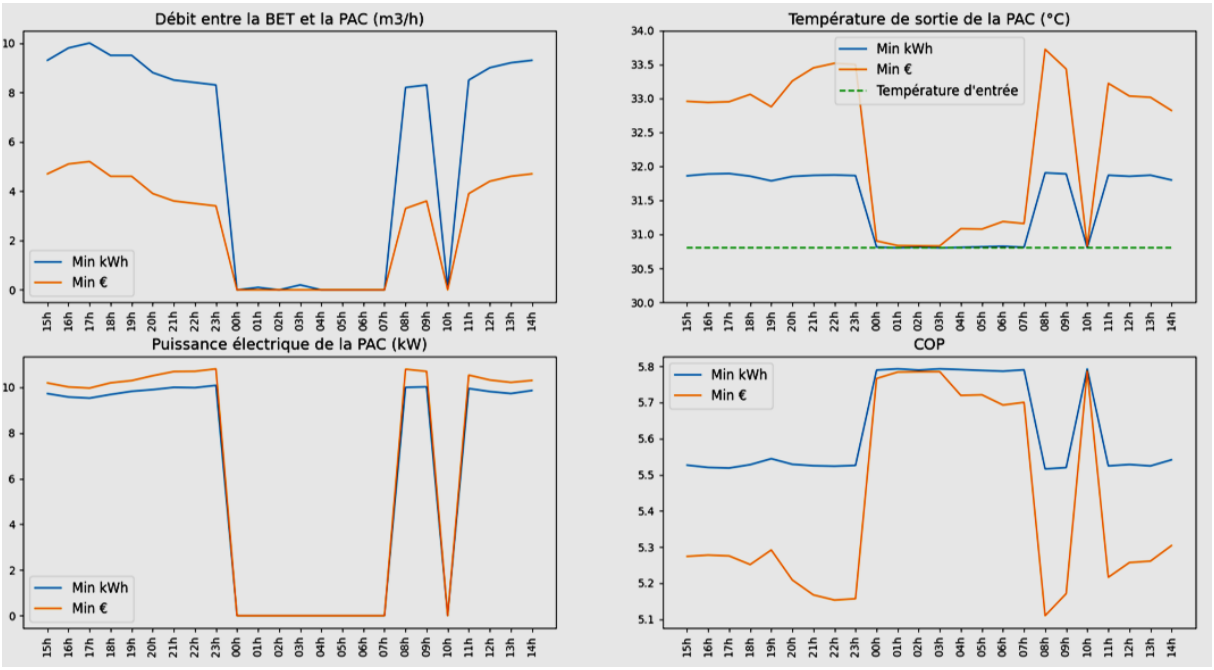

Figure 42 : Minimisation of total electricity consumption in kWh (blue) compared to minimisation of the total electricity costs in € (orange) for the Lylo-Marine substation. Source: EDF.

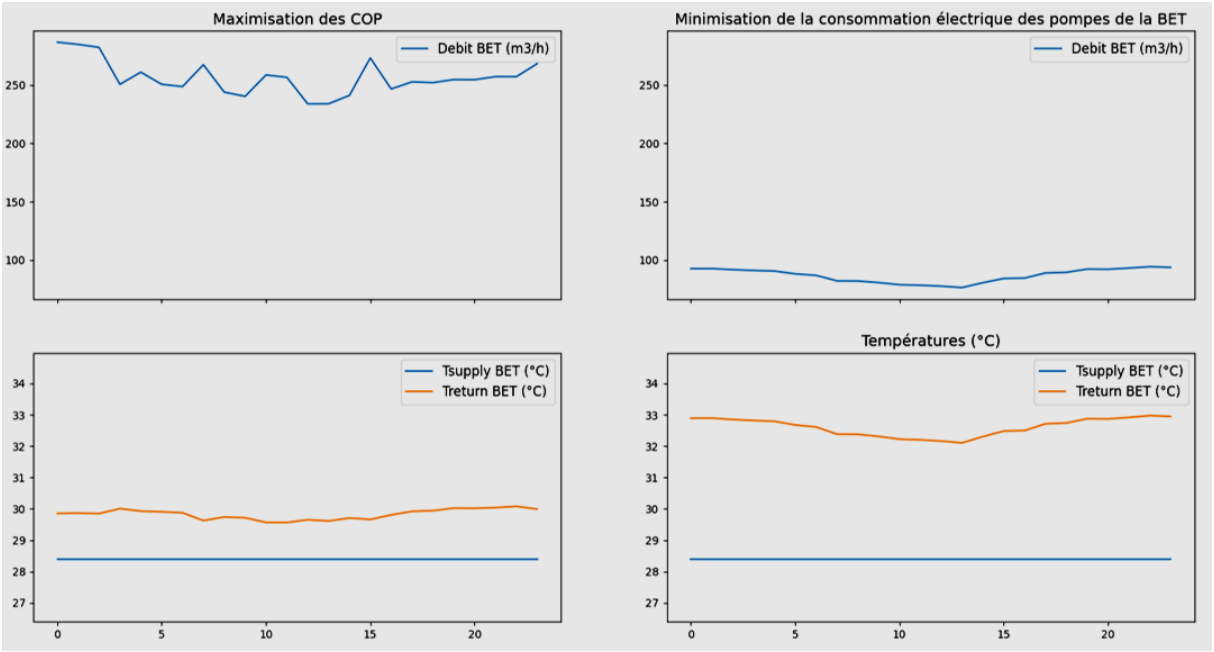

Figure 43 Maximisation of the overall HPs' efficiency (left) compared to minimisation of the electricity consumption of the pumping station (right). Source: EDF.

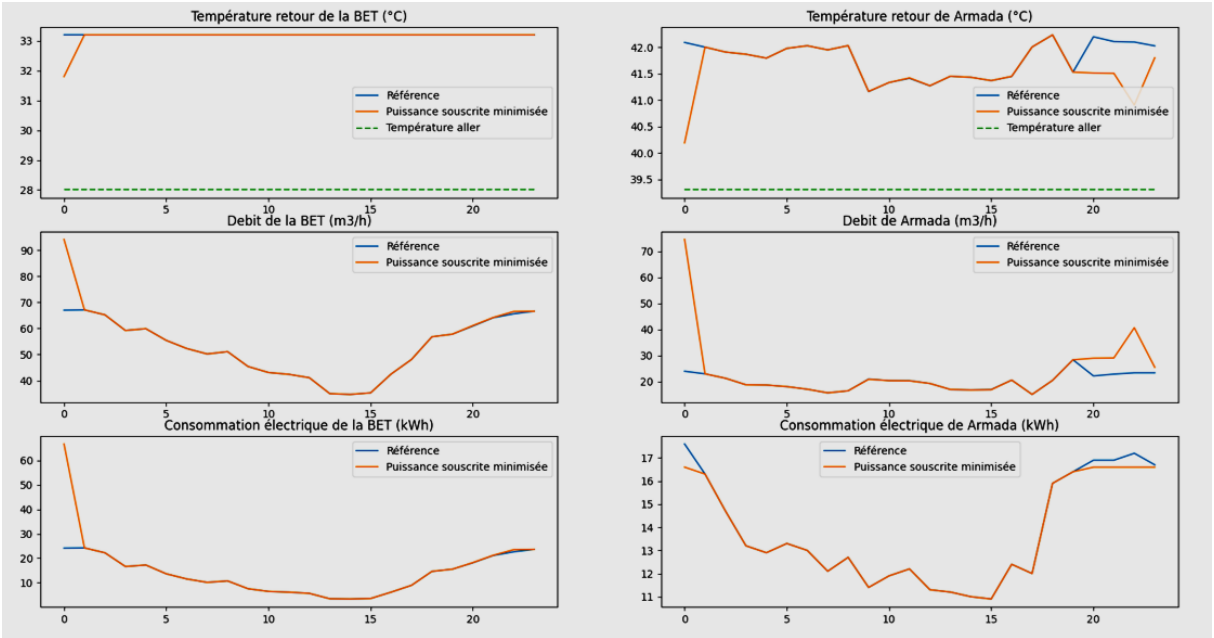

Figure 44 : Minimisation of electrical power demand (orange) compared to power subscribed for LSSM network (blue). Source: EDF.

With respect to the DEMix platform, and specifically the optimisation module (Clevery), the main objective achieved is that it has been tested in an operational real environment (TRL 6). The commissioning has enabled DALKIA to integrate the tool into daily operations, and the system has proven to be reliable, with the operator not needing to intervene in its operation, thus fully trusting the implemented solution.

As described, the system has been consistently enhanced in its optimisation functionalities, moving from a MILP to a MINLP approach, enabling the temperature/flow optimisation of the LSSM DHCN. The system is continuously being made more reliable, allowing for the expansion of modelling towards more accurate models/modules.

As of today, the optimisation module has enabled efficient operation of the DHCN of LSSM under real conditions. The DEMix platform has thus operated the main electricity-intensive asset, the pumping station, by controlling temperature setpoints and flow of the primary LT loop. Control of setpoints for individual substations has not been possible due to contractual and technical limitations. However, the optimisation problem considers these aspects and could, under different conditions, be tested to improve operations by leveraging control of both the LT loop and substations (e.g., HPs) relevant setpoints. This operation mode has the advantage of incurring relatively lower risks of impacting substations' operations while maximising the control's impact on the overall DHCN.

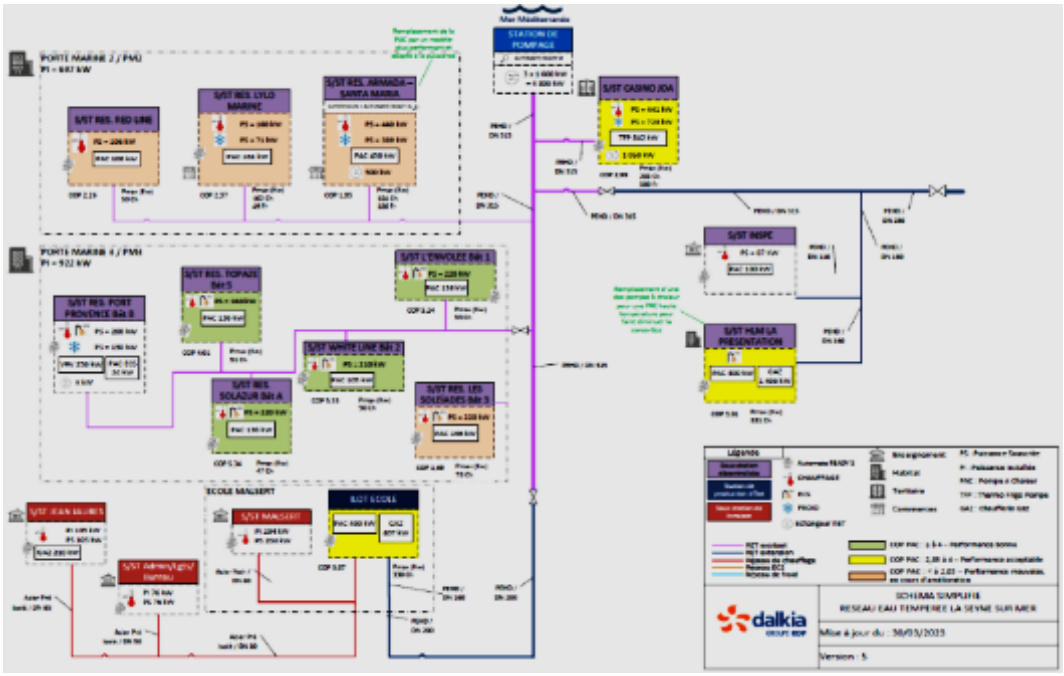

Figure 45 : scheme resuming the COP assessment per substation: green ones have COP above 4, yellow between 2,85<>4 and light red when below 2,85. The latter have been prioritized. Source: DALKIA.

DALKIA has been following the setpoints since February 2024 without any events or problems. Nevertheless, it is not possible to have a clear comparison between two similar periods (with and without optimisation) based solely on metering data, due to two main reasons:

- The DHCN has undergone many significant modifications of its assets and related controls, such as closing a main valve, removing some heat exchangers between the main loop and the HPs or coolers, adding an exchanger in other cases, and equipping head pumps with variable speed drives. These modifications have had a positive impact in terms of energy efficiency and COP but have made a proper comparison based on metered data impossible, allowing only a simulation-based approach for defining a reference scenario.
- During the project lifetime, individual sensors and metering equipment have shown to not work properly, with some parts just recently undergoing replacement or repair. More particularly, energy sensors, especially for heat, are very sensitive and prone to errors (e.g., a

temperature sensor not placed at its optimal location), showing consistent deviations between observed and expected behaviour. Thus, these measurements are not reliable and cannot be used for a fair comparison.

4.5 Lessons learned

The main learning out of the described work are:

- Predictive models // ergonomics: A trade-off between model diversity/complexity and forecast accuracy had to be integrated in the objectives of the forecast: considering the industrialisation of the DEMIX platform, a general end-user is assumed as being a non-specialist of forecast methods, forcing to reduce the variety/choice of available forecast models/methodologies, sometimes against possible accuracy loss of the forecast for specific substations. Some forecast might be improved, but the choice can't be automatized (with the current given constrains).
- Predictive models // Post-COVID situation: The COVID situation has made historical data almost unusable and delayed the validation work of the forecast. The implemented self-learning algorithms, need yearly timeframe for the training and are considerably impacted by such outlier events. Forecast methods, needing shorter timeframes (few month) could be better adapted to the problem at hand (see following point). The accuracy of the forecast under such time constrain should be investigated.
- Predictive models // new assets : New substations do not have data to learn from and thus, new substations are barely predictable in the first year of exploitation with the implemented method. Shorter time-frame applications could enhance the commissioning time a full forecast and optimisation platform as that developed for LSSM.
- Predictive models // predictors choice: More predictors/data does not equal to greater accuracy and depending on the methods implemented, such predictors are already, implicitly, entailed in the load data and well considered by the implemented forecast method. This is the case observed in the methodology deployed for LSSM, where no enhancement has been achieved.
- Predictive models // data reliability: The model must consider the fact that the data is not always complete as data streams might have outages and be able to work despite data-gaps (usually such accidents could be solved in a matter of hours). This has been taken into consideration for the developed forecast module, able to bridge several hours of data outage. Nevertheless, no solutions are at hand if the outage concerns one or several days of data stream outage and thus, a DHCN running MPC platform has to have a back-up solution as in the LSSM case, where the SCADA system takes over the operation in case of missing or impossible outputs from the MPC platform.
- Predictive models // local conditions: Local weather correction methods must be integrated; it has been shown that the implemented forecast module is able to correct recurrent weather forecast deviations compared to local conditions in a correct way. This has been tested against different temperature sources, showing that the model did correct such deviations in an acceptable manner.
- Predictive models // results accuracy: Acceptable results have been achieved for the overall load forecast error, staying below <12% in average on an annual basis, but results vary greatly among substations, even showing substations that may behave unpredictably or with chaotic

behaviour. Thanks to the commissioning work of DEMIX, part of such erratic behaviours could be corrected as related to metering or data aggregation issues or related to specificities of the secondary network regulation or system design.

- Optimisation // white box approach: Thanks to the fact that the model does not work as black-box, but is based on physical models, cause-effects can be identified and explained, and corrective measures more easily identified than on a grey/black-box approach. This has been shown to be very helpful for the commissioning process, enabling indeed to have a two-way improvement on the site. With the latter, is meant that not only the model could be verified against real data, but real data be challenged against consistency and on-site/implemented operational modes questioned.
- Optimisation // technical control constrains: the already implemented control and SCADA, is constraining the optimisation freedom from the forecast/optimisation platform. Indeed, a new mode of operating such LT DHCN should be revised, so to be able to implement regulation orders or specific set points for different assets, which are otherwise not possible to be implemented today.
- Optimisation // Regulatory/contractual constrains: another lesson learned is the close link between optimisation freedom and contractual or regulatory constraints. Indeed, regulation imposes a max limitation on the sea water injection, major determinant for the exploitation efficiency of the DHCN during summertime. Equally, contractual constrain with secondary network operators, at the building interface, make certain operation strategies not feasible (i.e. fixed set point for delivered temperature).
- Optimisation // System design constrains: the current "classic" design solution of substations, where a heat exchanger hydraulically isolates HP and other generation/storage assets from the primary network, is not well adapted for the system at hand. It is suggested, to replace the heat exchanger by a buffer tank, which would enable to have peak shifting/shading capacities, thanks to the intermediate storage. This specific design solution has however not been tested yet.
- Optimisation // Non-linear solvers: despite the fact that the integration of non-linear solvers has enabled to rethink the whole optimisation approach, these are relatively new and still under improvement as technology/service. If on the one hand, the linearisation of the system is yet not needed anymore, solvers can't handle the whole physical complexity at once on the other. Trade-offs must be done, between the level of detail of the physical description of assets or processes and the convergence time or number of unfeasible solutions resulting from the optimisation.
- Optimisation // industrialisation: the solution tested on site was developed in Python, for its simplicity of implementation. An industrialisation in C++ is currently under way, with many difficulties arising from the adaptability of the equations, i.e. equations were fitted for the LSSM test-case, but may not be perfectly adapted for other networks (i.e. difficulties to converge). This replicability problem still needs to be addressed, and EDF and DALKIA will continue to work on it further on.
- Design // ICT endowment and data reliability: for what concerns real-time operation, a fibre-optic wiring among substations and a central automation/control should be ensured. As in the case of LSSM, a 4G communication had to be implemented as only possible solution (against huge amount of CAPEX and related civil-works) as being the DHCN mostly existing and not equipped with a centralized automation. Such solution has proven not to be reliable

enough for a real-time platform, causing too long outages of communication of single assets, putting at stake the whole forecast/optimisation process. Despite many interventions and improvements implemented during the project, aiming at signal quality and stability improvements, the overall data stream reliability is not completely satisfying. A fibre-optic solution seems being yet the most reliable one. Solutions have to be found to fill data gaps. A fibre-optic solution appears to be the most reliable option. Additionally, solutions must be found to address data gaps.

5 Lund (E.ON)

5.1 Description of the demonstration site

The project area Medicon Village is a campus in Lund focused on medical research with more than 2500 employees working for different companies. E.ON uses an ectogrid™ Neutral-Temperature Heating and Cooling Network to supply Medicon Village and three large apartment blocks.

An ectogrid™ consists of conventional equipment, such as heat pumps, cooling machines, heat exchangers and piping. The innovation of the technology lies in the way these blocks are combined, enabling new hydraulic coupling allowing to achieve an automatic bidirectional heat exchange. Generation is mainly distributed, and excess energy from each building is used internally by heat recovery from the heat pumps, thus lowering the overall energy consumption in each building. The decentralized heat pumps and chillers in each building raise or lower the supplied temperatures to the levels that are specific to the building requirements, avoiding unnecessary energy consumption.

Figure 46 – Overview of the Medicon Village ectogrid™ with the new substation highlighted in red.

5.1.1 Objectives of the demonstration activity

Within the REWARDHeat project, a new kind of decentralized generation unit has been developed and tested. The heat pumps are connected in a novel "transversal" way to simultaneously act as both heat pumps and cooling machines. Both the inlet and outlet temperatures are controlled and steered to optimize the functionality of the grid.

This novel substation setup will have a significant impact on the potential for the overall system performance optimization. i.e. how the control of several buildings and substations in an ectogrid™ can be coordinated. To develop this knowledge and enable the dissemination of the results, numerous operational conditions and control strategies have been tested and evaluated, and control algorithms developed and implemented, leading to high replicability of the project results.

Figure 47 shows the main energy flow principles of the novel substation whereas *Figure 48* shows the P&ID, outlining the technical components.

5.2 Control problem formulation

Considering that heating and cooling demands are most often not balanced, there are several control challenges related to how the substation is to be run:

Meeting the demands for heating, domestic hot water and cooling

Handling the rejection of excess heating or cooling to the ectogrid™ warm and cold pipes at desired temperatures

Operating the heat pumps and adjacent equipment smoothly to avoid unnecessary wear and tear

Thus, there are multiple setpoints that must be simultaneously tracked by the substation PLC/BMS controller to distribute water via different circulation pumps and valves. Going forward, it will be important to standardize the substation control logic to enable scaling of the solution.

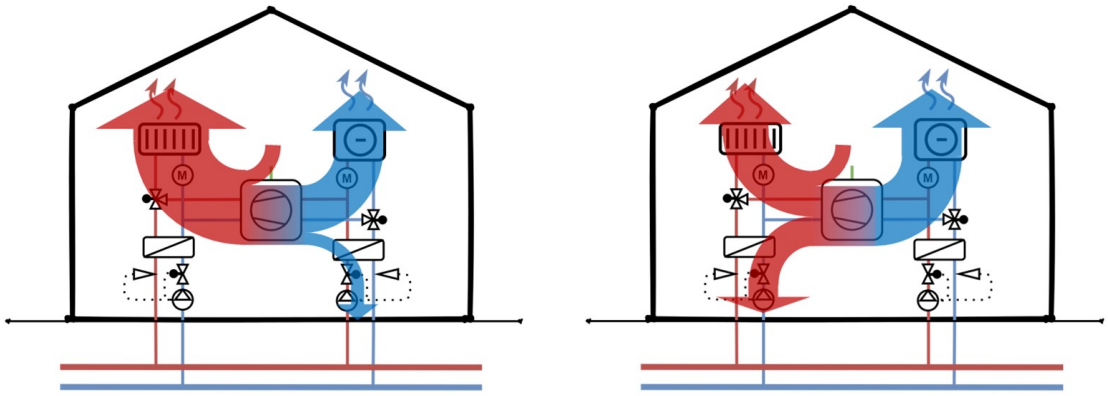

Figure 47 - Heat pump-based novel substation operation principle

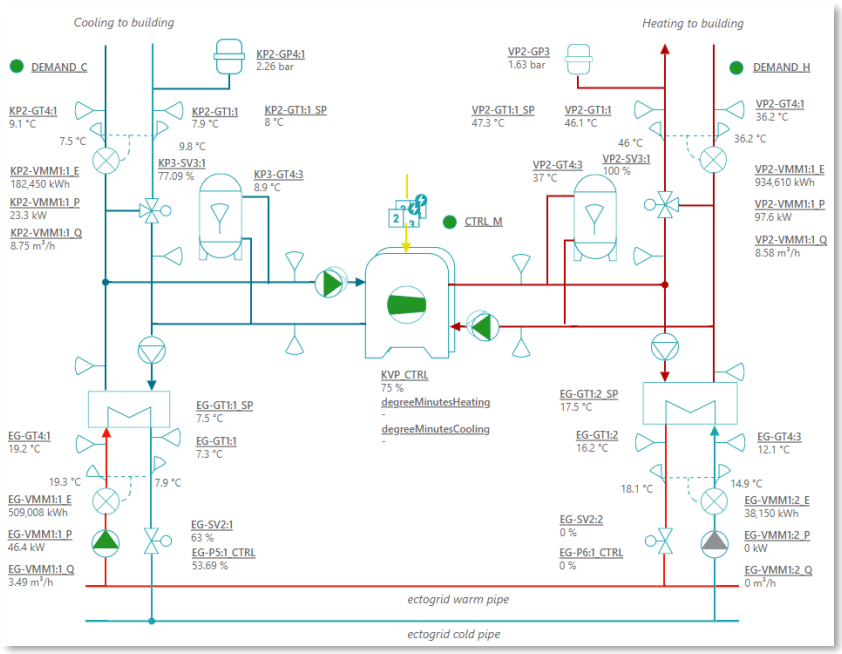

Figure 48 – Detailed P&ID of the substation developed

Additionally, different digital capabilities have been developed to unlock further values based on the cloud-based platform ectocloud™. These include:

Monitoring the solution in detail and in near real-time

Enable proper analysis through data aggregation and visualization such that evaluation of the substation energy performance is possible in both local and system-wide contexts

Intelligent optimization of thermal demand, e.g., through demand-side management of building thermal loads based on price signals

Inclusion of the novel substation into the overall steering of the ectogrid™ system, by adapting energy load forecasting to the novel substation configuration

5.3 Description of the control implemented

5.3.1 SCADA and Energy Analytics System

The digital platform ectocloud™ is a control, optimization, and data acquisition platform for decentralized, sector-coupled energy systems. It utilizes modern technologies such as cloud and edge computing, machine learning and IoT communication. It is designed to efficiently connect and digitally structure distributed assets and provides tools for performance monitoring, asset operations and for data-driven energy optimization and control purposes.

The connection of the building and substation to the ectocloud™ is facilitated through IoT technology with high resolution data obtained from meters and sensor. This includes an IoT gateway able to communicate with local PLC and BMS systems, which in turn gather data from hard-wired sensors, metering gateways and other underlying control systems such as ones built into the heat pumps. Additional data relating to weather forecasts and energy market day-ahead prices are integrated by cloud-to-cloud APIs.

In terms of ectocloud™'s architecture, illustrated by *Figure 49*, the core backend acts as the main integration and data storage point, fetching data from different external sources like weather forecast data or energy market data. As mentioned, IoT gateways ("Energy Manager") are used to interact with the physical energy systems via local PLCs or BMSs in near real-time, while machine learning services provide the computational tools necessary for advanced energy optimization use cases.

The main purpose of the energy system visualization, exemplified in *Figure 50*, is the operational supervision and energy analytics.

Special formulation of the individual COP of heating and cooling had to be made such that the common electricity use was separated in a fair and representative way depending on the demand balance of the building and the amount of excess heating or cooling. An important consideration has been to make substations, regardless of exact design, comparable in how they are benchmarked.

The system COP, which is an important metric, can thus be calculated as below for different time frames such a daily, monthly or yearly:

$$COP_{ecto} = \frac{Q_{H,delivered} + Q_{C,delivered}}{E_{input,total}}$$

Other energy data-based KPIs aggregations are also provided to give an overview and drill-down user experience.

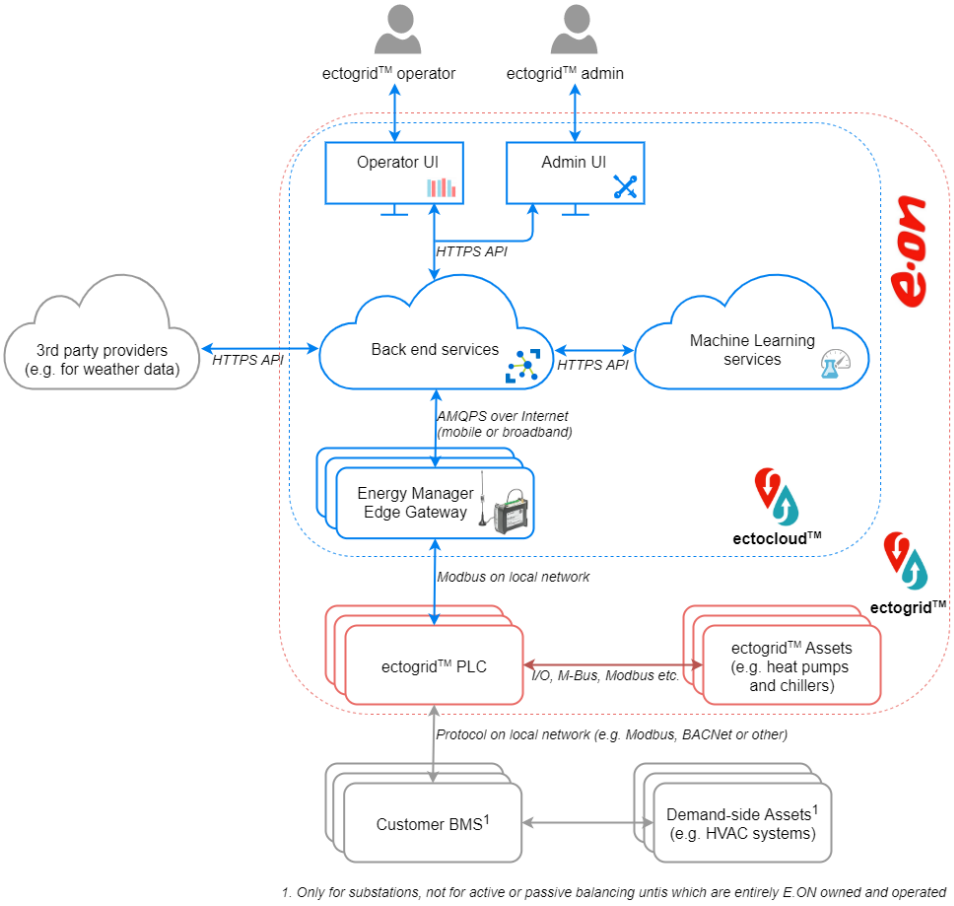

Figure 49 – ectocloud™ core architecture and its relationship with external services, local systems and hardware

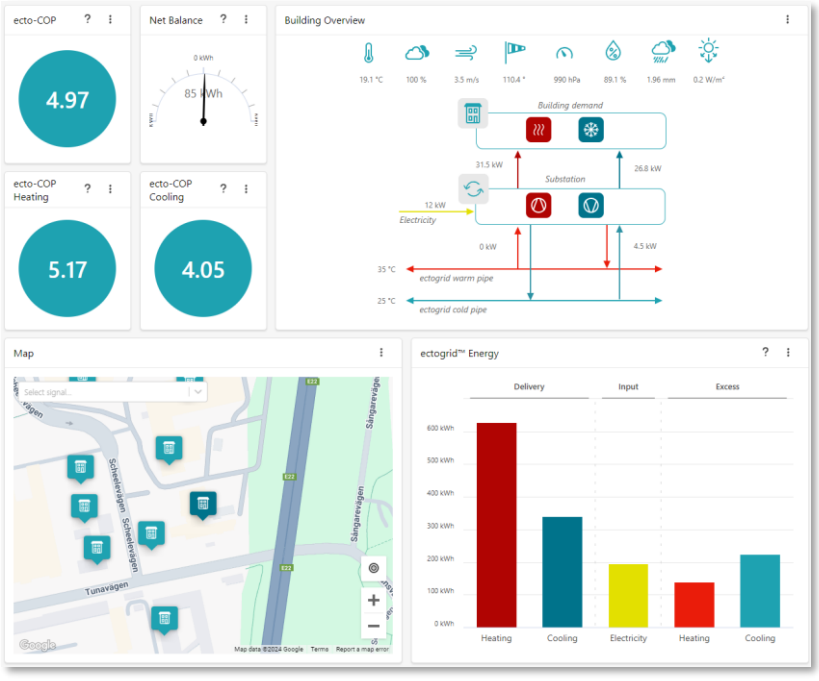

Figure 50 – ectocloud™ analytics and user interface example

5.3.2 Substation

As described, the substation real-time control is implemented in the local control system to track multiple setpoints. The setpoints for the delivery of heating and cooling are provided from the building HVAC control system (BMS), whereas setpoints for the excess heating and cooling are set by ectocloud™, based on system-wide grid temperature considerations.

Different controllers with specific tasks interact such that the combined heat pump run to meet the dominant demand and to distribute the excess energy back into the grid. These controllers are:

Distribution controller (heating and cooling): Controls the mixing of flow such that the building demand is met, and that potential excess energy gets redistributed

Production controller: Controls the heat pump compressor utilization such that the necessary forward temperatures on the condenser and evaporator sides are met (given by the building demand or excess energy rejection)

Grid temperature controller: Controls pump and valve to make sure excess energy is injected to the grid at the desired temperature

An illustrative diagram of the controllers interaction with the substation is shown in *Figure 51*.

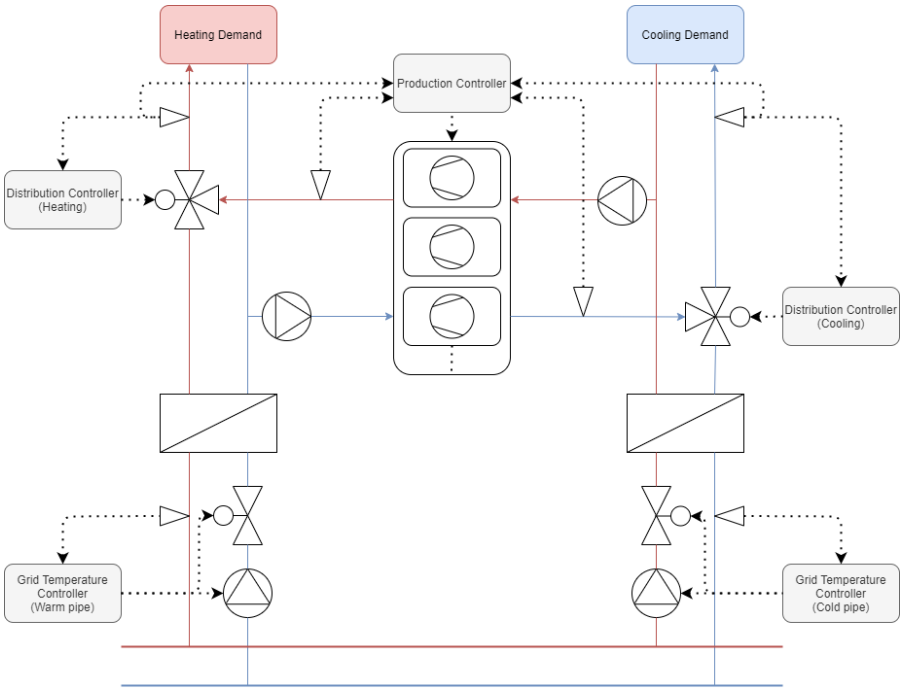

Figure 51 – Feedback loops acting collaboratively on the novel substation

Controller Tuning

To complement experience-based control tuning, more structured tuning strategies have been tested to create guidelines for how get satisfactory control performance. In addition, certain other factors influencing the control performance have also been documented, such as the sensor placement.

To improve the overall tracking of temperature setpoints, strategies for tuning the controllers have been investigated and designed according to methods like Ziegler-Nichols and Lambda to calculate

appropriate PID control parameters. A step response experiment is conducted by running a manually defined control signal to identify process dead-time, time-constant and amplification of control signal to measurement response.

Feed-forward Control

Rather than just having the conventional feedback control architecture, a feed-forward component has been tested for substations, where disturbances originating from compressors starting or stopping in discrete steps caused short-term but high temperature deviations. The control architecture of feed-forward and feedback control is shown in *Figure 52*.

As an indicator of cooling energy increase on the evaporator-side of the heat pumps, a compressor start signal was used to pre-emptively increase the flow, limiting the temperature impact of the sudden event of compressor start.

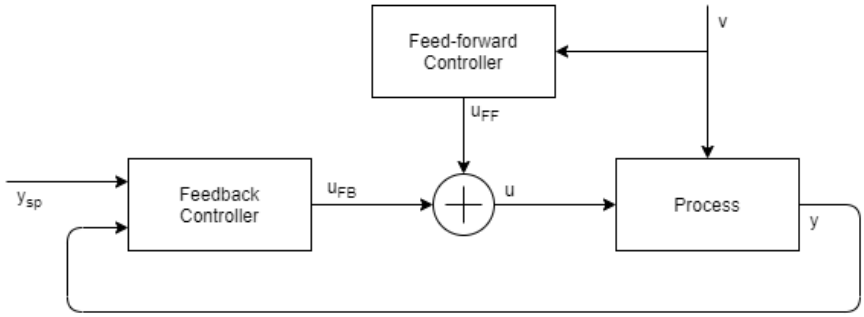

Figure 52 – Feed-forward control loop concept where the disturbance indicator v is measured as a compressor start signal

5.3.3 Advanced controls

In addition to the substation basic control, different more advanced controls have been developed and tested including:

Demand-side management of thermal loads based on day-ahead price signal

Grid temperature control for enhanced ambient energy harvesting

Control involving the main thermal storage in the grid with energy demand forecasting

These advanced and system-wide control strategies act on top of the low-level control of each substation and asset and are provided by ectocloud™.

Demand-side Management

For the demand-side management, comfort-related heating loads have been connected and controlled, allowing ectocloud™ to influence the building space heating consumption. By considering day-ahead price signals and the thermal inertia of the building, the developed algorithm moves heat consumption from expensive hours to cheaper hours of the day. As the heat is generated by heat pumps, savings on electricity costs can be achieved without noticeable indoor climate impact.

In *Figure 53* the setpoint control and response of the main heating circuit of the building is shown together with the day-ahead market electricity price. As can be seen, the supply temperature is shifted up and down depending on the price that during these five days varies dramatically.

Energy Load Forecasting

From the perspective of controlling and optimally balancing the grid, the novel substation has introduced a new method for deriving the net demand for heating and cooling. This change has impacted energy demand forecasting, which is a critical component for optimising all parts of the system, alongside models that describe the dynamics of energy source efficiency, thermal storage, and the distribution grid. The energy forecasting model uses machine learning trained on historical data to predict heating, cooling, and electricity consumption based on forecasted weather, time of day, and other factors.

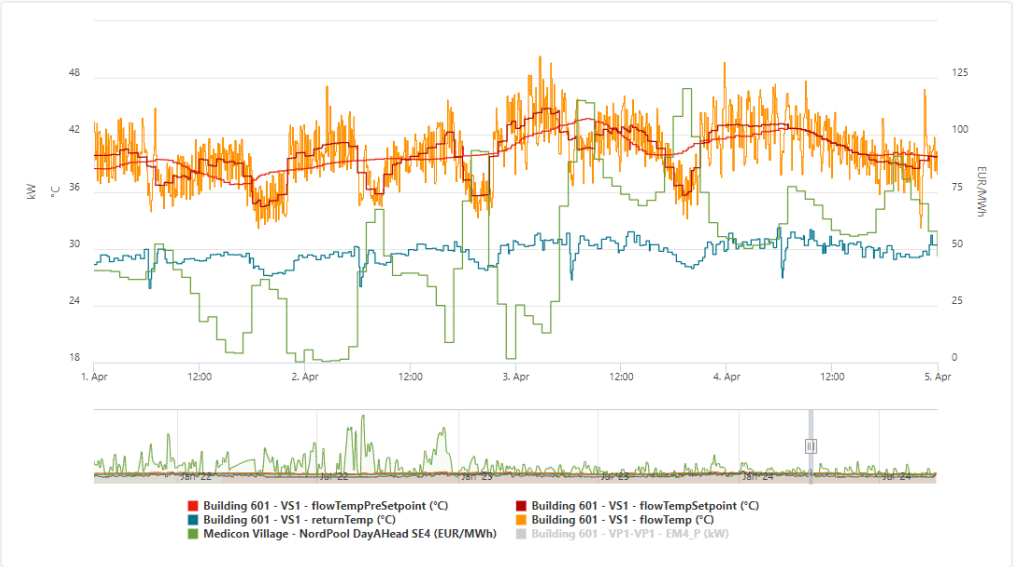

Figure 53 – Timeseries of supply and return temperatures from/to the substation during demand side management operation together with day-ahead price data (green line). The lighter red line is the calculated baseline temperature whereas the red is the actually controlled and optimized one.

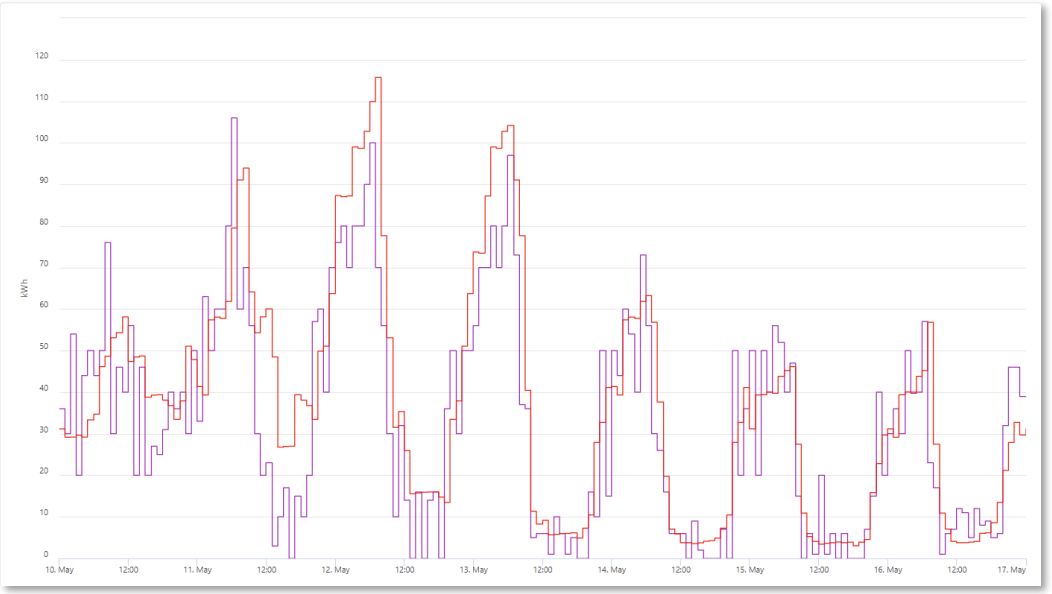

Figure 54 – Thermal load forecasting example timeseries with forecasted heat demand in red and actual value in purple.

As part of the system-wide optimisation, grid temperatures are determined by considering the temperature-dependent COP of the substations and energy sources, along with distribution losses or gains. Since all consumers are also producers in ectogrid™, the setpoints are sent to all connected systems to be adhered to by each of the local controllers.

5.4 Assessment of the control performance

Thanks to a continuous improvement and control tuning process, the performance of the novel substation control has reached a satisfactory state. Setpoints are tracked accurately without running machinery and equipment too aggressively, see *Figure 55*.

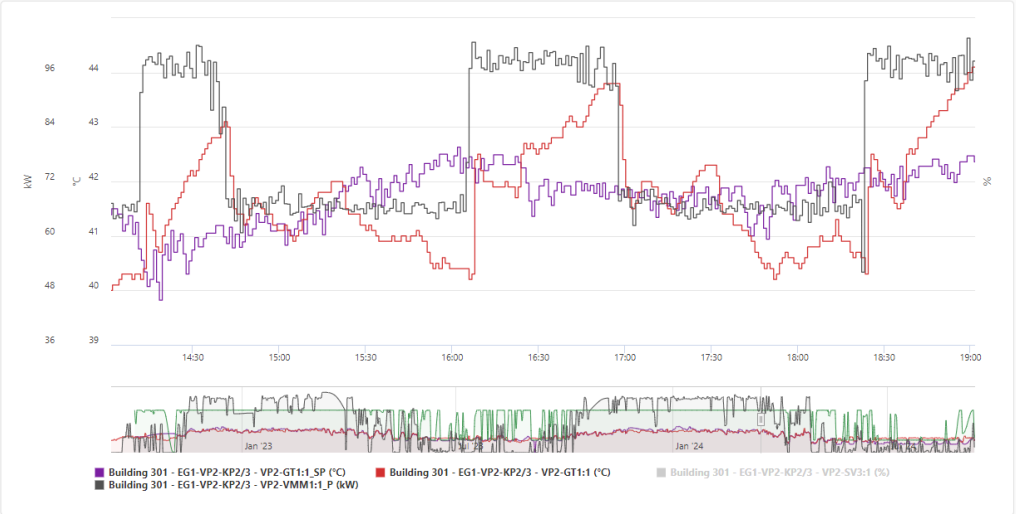

Figure 55 – Substation supply temperature control set point shown in purple with measured value in red. Note that over a time-period average the setpoint is kept but the step-controlled compressor is not triggered to often as indicated by the heating power in grey.

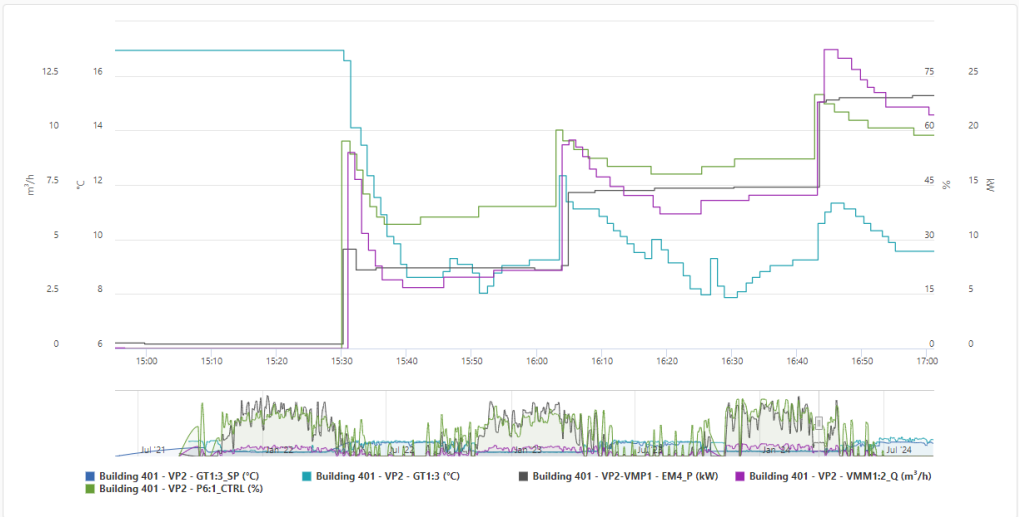

Figure 56 – The eed forward control in action. As a compressor steps in (indicated by electrical power in grey), the control signal is immediately increased resulting in an increase in flow on the condenser side (purple). Thus, the temperature (blue) settles at around 8 °C without undershooting.

In substations where excessively low evaporator temperatures could occur due to compressor start-ups, decoupling the control logic into feedback and feed-forward control blocks has proven useful. This approach allows low temperatures to be sent back to the grid without triggering freeze protection stops, as the feed-forward controller pre-emptively increases the flow, see *Figure 56*.

Overall, standard approaches have been developed and documented to ensure they can be reused in future substation installations, both in terms of control logic and efficient installation and connection to the cloud for data acquisition.

Regarding demand-side management control, significant load shifting has been performed over an extended period. The customer has been involved in the process and has not experienced notable changes in indoor comfort, demonstrating the effectiveness of building thermal inertia. At the same time, the potential for savings has been shown to be significant, assuming optimisation towards dynamic tariffs (day-ahead spot prices).

Given the electricity price volatility, the cost-saving potential for the tested buildings is estimated to be up to around 20%, with only a limited impact on indoor temperature. The general potential of such optimisation depends on future price volatility, customer sensitivity, and the properties of the building (insulation and thermal mass).

Example operation output of the demand-side management control is shown in *Figure 57*.

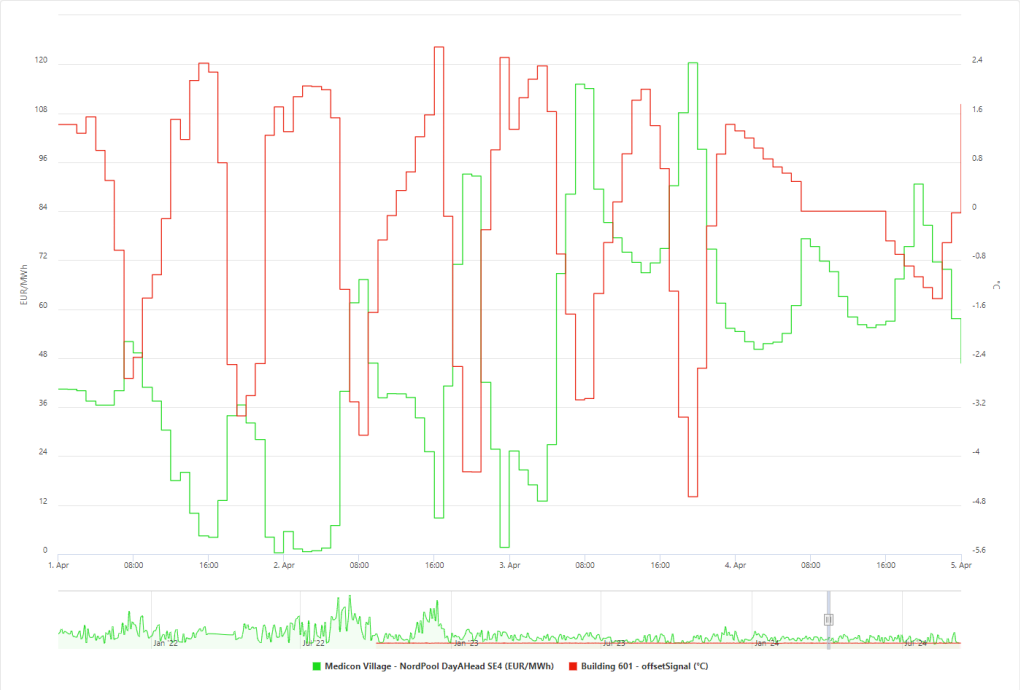

Figure 57 – Example of space heating setpoint offset control in relation to Day-ahead electricity price

The developed load forecasts for the novel substation and associated building have been generalised from previously used machine learning-based forecasting approaches, utilising historical data, weather forecast data, and temporal data as predictors. The model training shows promising results, consistent with the performance of models for other buildings at the site.

The dominant predictor for heating demand was the 1-hour average outdoor temperature, with some weaker dependencies on the time of day also present. For cooling, multiple relevant

predictors were identified, mainly temporal ones such as the time of day and whether it is a weekday or weekend.

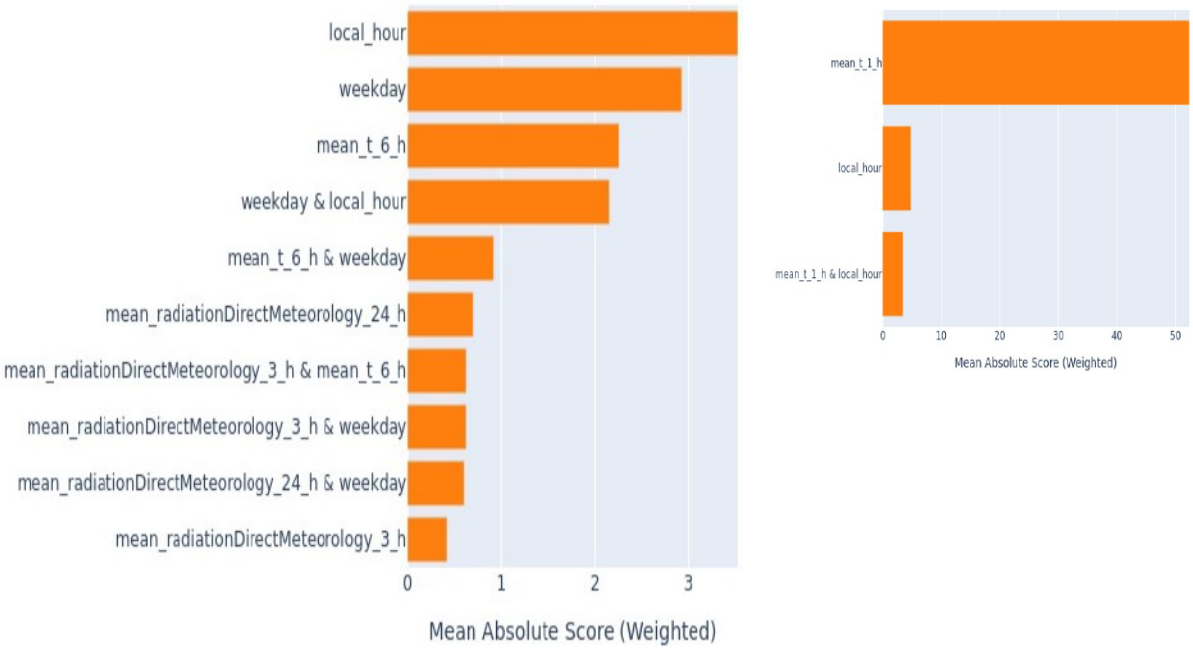

Figure 58 – Predictors influencing heating (right) and cooling (left) load forecast

The heating demand forecast model has reached the following performance:

R2: 0.95

RMSE: 11 kWh

The cooling demand forecast model has reached the following performance:

R2: 0.75

RMSE: 4 kWh

5.5 Lessons learned

There are multiple learnings and take-aways from the experience of introducing the novel substation into the ectogrid™ solution, including:

- It is important to consider all levels of control, taking a bottom-up approach starting from the control systems working closest to the hardware and considering how one can build more sophisticated control logic and optimisation on top of that
- Trying to apply different control tuning and other real-time control mechanisms such as feed-forward requires close collaboration with building automation contractors, especially for such sophisticated substations
- Harmonizing and assimilating data from different types of substations so that variations can be handled seamlessly can be challenging, but important to enable digital standardisation
- The potential of demand-side management can be significant in heat pump-based heating and cooling solutions, in the current times of price volatility
- The approach for creating machine learning-based energy forecasts is not straight-forward when it comes to the extent of modelling the interdependency of heating, cooling and electricity that exists in heat pump-based systems, which was exacerbated by the novel substation and the utilisation and effects of demand-side management
- In cases with less representative data for the substation or building behaviour, such as when it is new and undergoing its early phases of commissioning, it could be useful to initially use forecasting methods that are less data-intensive and more adaptable to real-time conditions and events.

6 Palma de Mallorca (SAMPOL)

6.1 Description of the demonstration site

The Parc Bit power plant is a Combined Cooling, Heating, and Power (CCHP) facility that supplies electricity, heating, and chilled water to the Parc Bit Innovation Centre and the University of the Balearic Islands (UIB), through a district heating and cooling network. The plant generates hot water using a combination of CHP engines, a biomass boiler, solar thermal panels, and a fuel boiler. To meet cooling demands, the facility employs three conventional electric chillers and two absorption chillers. The absorption chiller and one of the electric chillers use water for cooling and share the same cooling towers, while the remaining electric chiller utilizes ambient air for cooling purposes. Heating and cooling can be stored in TES water tanks.

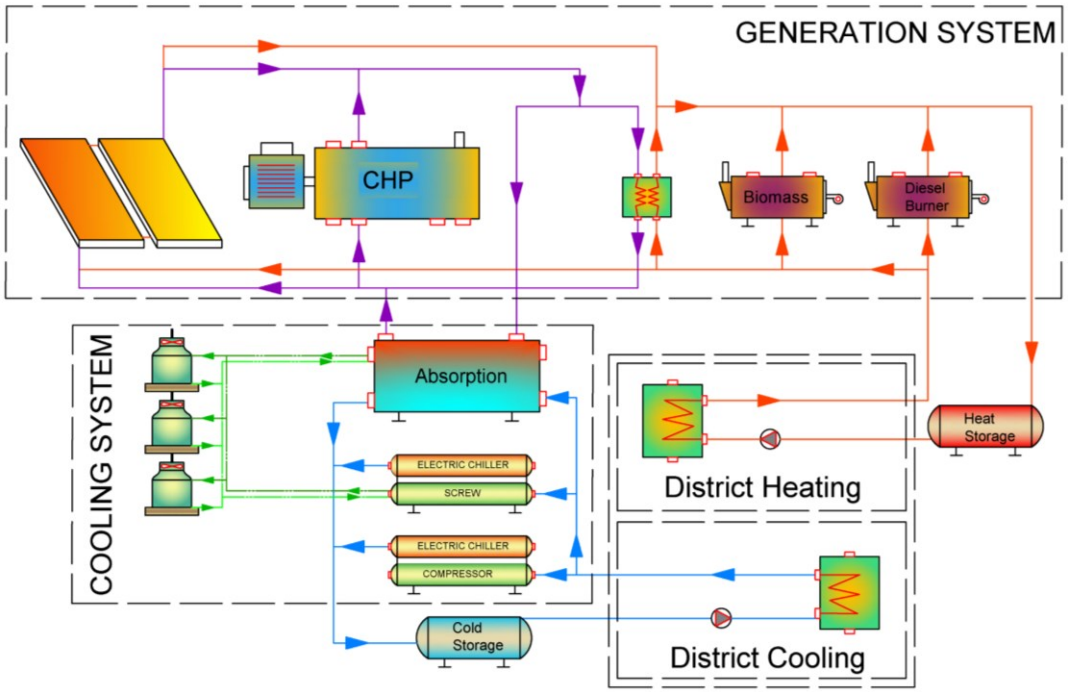

Figure 59 Parc Bit Power Plant scheme

The DHC was initially constructed in 2000, linking the CCHP plant to the office buildings within the Parc Bit complex via three branches. In 2002, the network underwent an expansion, adding another branch to connect with the university facilities, including student housing and the sports centre. Presently, the network services 25 different customers, providing either heating, cooling, or both. The customers on this network fall into three primary categories: office buildings, educational facilities, and specific usage buildings such as residential spaces, a swimming pool, and an IT room. The profiles vary considerably between workdays and weekends, adding complexity to the network management.

The network consists of four branches of pre-insulated steel pipes. Each branch contains two pairs of pipes, i.e., supply and return. The total length of the district DHC network is approximately 4.6 kilometres.

Figure 60 Parc Bit and UIB District Heating and Cooling Network

6.1.1 Objectives of the demonstration activity

The demonstration aimed to enhance the operational efficiency and energy performance of the DHC network through the application of advanced control strategies and optimisation techniques in both the CCHP and the DHC network. The key objectives of this demonstration included mainly developing an advanced control solution for both the CCHP and the substations distributed along the network, aimed to maximise economic performance, minimise thermal losses, and optimise CHP-DHC system operations under varying demand conditions.

6.2 Control problem formulation

To effectively address the control problem cantered on the innovative substation, the following objectives have been outlined:

- Enhancement of Control Systems: Upgrading PLCs and refining functional analysis at the central pumping station and customer substations. These systems were continuously improved based on feedback from network modelling and simulation, optimizing performance and energy use.
- Optimisation of Substation Control: Installation of three-way valves with controllers at each substation aimed to precisely regulate the flow of heating and cooling water, optimizing the temperature differential between supply and return lines. These controllers dynamically adjust the valves based on real-time demand, improving the efficiency of heat transfer and reducing energy consumption. This setup allowed for better control over the distribution network, ensuring that the right amount of energy is delivered where and when it is needed, thereby minimizing waste and lowering overall electricity use.

- DHC Demand Forecasting: The goal was to forecast demand for the next six hours in 15-minute intervals using an LSTM (Long Short-Term Memory) model. This model utilises inputs such as historical ambient and wet bulb temperatures, along with sinusoidal variables to capture seasonal information, enabling accurate energy demand forecasting and efficient management of the DHC and TES systems. Based on this, a better integration between TES systems at the power plant and CCHP was achieved, aimed to improve the flexible management of the plant, and allowing for strategic scheduling and operation to optimise economic and energy efficiency while reducing emissions.
- Deployment of Advanced Control Systems: A centralised control system using reinforcement learning (RL) techniques was implemented to optimise the next six hours of operation by adjusting the output temperature of TES to the DHC. The RL agent employs a combination of model-based algorithms and deep learning to make decisions that maximise the system's economic performance, minimise thermal losses, and optimise the operation of CHP-DHC systems under varying demand conditions. The following instruments were implemented:
- Platform Development: Develop a comprehensive platform that integrates all components required for the Digital Twin's operation. This includes processes from data acquisition and processing (data munging) to simulation optimisation and live evaluation of the proposed control strategies.
- Data Acquisition: Implement a data acquisition API that incorporates connectors to PLCs, weather APIs, PID controllers, and other essential data sources. This API ensures that all necessary data is collected and made available to the Digital Twin for seamless operation.
- Control Operations: Establish control operations that define the scheduling of simulations based on the inputs allowed by the control strategy. The control strategy in this case involves optimising the power plant's operation based on the startup schedule of machinery and the accumulation in district heating tanks.
- Model Retraining: Integrate a continuous model retraining process to ensure that the machine learning models remain accurate and effective. This involves regularly updating the models with new data to adapt to changes in system behaviour, ensuring that the control strategies are always based on the most current and relevant information.
- Live Operation Analysis: Design the evaluation process to run in real-time, with results visualised through an interface. This allows for the detection of any malfunctions in the simulation, enabling a quick switch to stable operation while the issue is analysed and resolved.

6.3 Description of the control implemented

6.3.1 Digital Twin Platform

The Digital Twin platform has been developed with a strong emphasis on replicability and scalability, ensuring that all components are containerized using Docker to facilitate seamless deployment across various servers. This modular architecture is bifurcated into edge and cloud components, which are delineated in the system's schematic by dashed line squares. This design enables continuous local operations, even in scenarios where network connectivity is compromised, a critical requirement for installations in remote or geographically dispersed

locations. While the Parc Bit site benefits from a reliable connection to Sampol's central office, this capability ensures robust performance across diverse operational environments.

At the edge, the platform is tasked with data acquisition from the Combined Heat and Power (CHP) plant and the District Heating and Cooling (DHC) network. This is accomplished through an API that interfaces with Programmable Logic Controllers (PLCs) and Supervisory Control and Data Acquisition (SCADA) systems, thus ensuring comprehensive data capture. The extracted data is subsequently channelled into a Kafka broker, a widely utilized messaging platform in industrial applications, which efficiently manages data streams. These data streams are then directed to two distinct databases: a cloud-based repository that archives the complete historical dataset and an edge-based database with a two-week retention policy, designed to support real-time operational needs while minimizing storage demands on edge devices.

In the cloud, the platform focuses on the continuous training and updating of machine learning models that simulate the operational dynamics of the CHP and DHC systems. This process, known as Machine Learning Operations (MLOps), is orchestrated by an Extract, Transform, Load (ETL) platform, specifically Apache Airflow. Airflow is employed to schedule tasks, manage dependencies across services, and perform data backfills to mitigate any data gaps resulting from service disruptions. It also oversees the iterative training of machine learning models, systematically comparing new models against existing ones to ensure optimal performance. Model lifecycle management is conducted using MLflow, with model storage and associated metrics maintained in MinIO, a NoSQL database. This infrastructure allows the platform to dynamically adapt to changes in system behavior, performance degradation, and operational adjustments due to maintenance activities.

The integration between the model lifecycle platform and the edge system is facilitated by an API that manages a MongoDB collection, which stores the top three versions of each model. This API allows for the selection of the best-performing model based on real-time data, ensuring that the most effective models are deployed. Additionally, the API serves models from the edge MongoDB database to the live simulation, scheduling, and evaluation systems. A live evaluation system continuously monitors model performance at 30-second intervals, with results stored in a TimescaleDB and visualized through Grafana dashboards. This setup allows for real-time performance tracking and automatic model adjustments if deviations are detected.

The platform's optimisation function utilizes these live models to execute a trained Deep Deterministic Policy Gradient (DDPG) reinforcement learning algorithm. This algorithm optimizes the operational strategy for the CHP and DHC systems over a six-hour horizon, with scheduling updates occurring every 15 minutes. The optimizer directly interfaces with the CHP and DHC systems via PLC connections, allowing for the execution of the optimized operational plan. This architecture ensures that the Digital Twin platform not only supports efficient edge operations but also continuously evolves through cloud-based model management, offering a scalable and robust solution for optimizing energy operations in a variety of settings.

The pilot project aimed at optimizing the live operation of both the CCHP plant and the district heating supply has successfully incorporated the use of a 1000m^3 heat accumulation tank as a critical flexibility medium. This thermal storage capability allows for the strategic shifting of heat generation away from peak consumption periods, enabling the plant to adapt production schedules in a way that maximizes revenue from electricity sales to the market. Before implementing these optimisation measures, the operational planning of the DHC system was primarily based on historical consumption data and the empirical knowledge of plant operators. Operational decisions, including the scheduling of heating and cooling, were guided by anticipated

demand fluctuations and expected changes in gas and electricity prices, drawing heavily on past experiences. To better utilize the system's thermal storage, operators adjusted the push temperature of water from the tanks on an hourly basis. However, this approach had its drawbacks, including significant thermal losses during storage and distribution, which reduced overall system efficiency. Additionally, regulatory requirements, such as maintaining a minimum temperature of 65 degrees Celsius in the DH4 pipeline to ensure the provision of sanitary hot water for the swimming pool and showers, imposed further constraints. These challenges underscored the complexity of managing a regulatory-compliant, efficient DH&C system and highlighted the necessity for an optimized operational strategy.

6.4 Advanced controls

6.4.1 District heating and cooling demand forecasting

To leverage the 1000m^3 TES capacity, it is essential to infer near-future consumption and determine the energy that must be produced during a given period. This forecast provides information on how much energy needs to be produced for the next six hours in 15-minute intervals, allowing the optimiser to decide whether to produce energy before, after, during consumption, or a combination of these.

This tool was built using deep learning, specifically a stack of LSTM layers. The input variables are listed in Table 11. In Figure 61, the resulting forecast is shown in red, compared to the actual values in blue. Most of the error is attributed to the weekend factor, and future work will include information on weekends and bank holidays to address this issue.

Table 11 Input variables of the DHC demand forecasting

Name	Description	Units
Heat Energy	Consumed energy of the district heating measured at power plant heat exchangers (Last 24 hours of data).	GWh
Cold Energy	Total consumed energy of the district cooling measured at power plant heat exchangers (Last 24 hours of data).	GWh
Ambient temperature	Temperature measured at the CHP power plant (Last 24 hours of data).	°C
Wet bulb temperature	Temperature measured at the CHP power plant at 100% humidity level (Last 24 hours of data).	°C
sin(15min)	Artificial variable, sinus wave with period of 1 day (Last 24 hours of data).	-
sin(dayofyear)	Artificial variable, sinus wave with period of 1 year (Last 24 hours of data).	-

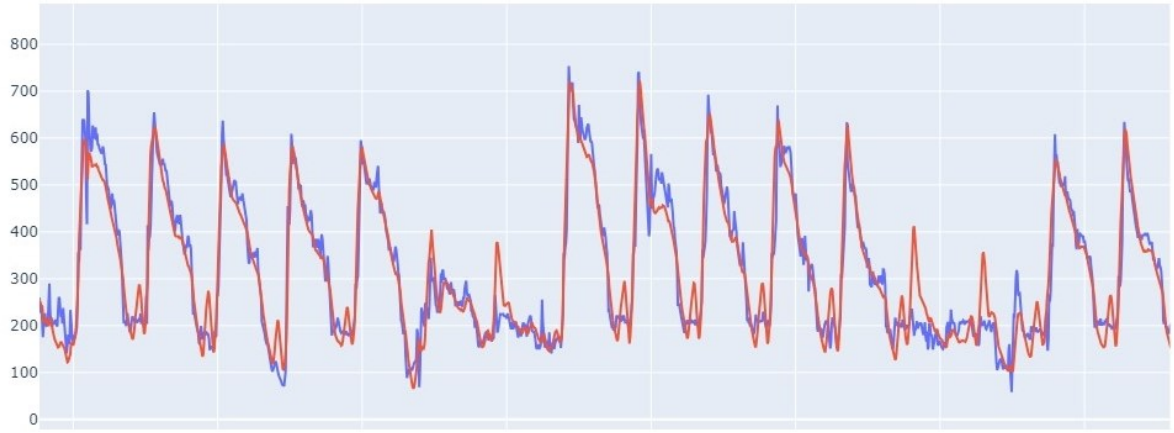

Figure 61 District Heating demand forecasting results on test data.

6.4.2 Electricity price forecasting

One of the main sources of income for the CCHP is selling the excess electricity generated after producing both heat and cold. The Spanish electricity market is regulated by OMIE, which facilitates the matching of electricity supply and demand in a competitive market environment. Prices in the Spanish electricity market are determined through a marginalist auction system, where bids and offers from generators and consumers are matched to set the market price for each hour of the day. This system is designed to reflect the true cost of electricity generation and consumption, allowing for efficient price signals and promoting the optimal dispatch of generation assets.

As shown in Figure 62, the OMIE price fluctuates significantly between days, following a consistent daily pattern. Most of this fluctuation is due to the generation of renewable energy sources, which enter the market at near-zero costs. On days when both solar and wind power have optimal production, they generate most of the necessary energy, allowing only the next most cost-efficient energy sources, such as hydroelectric power, to enter the market. Our model appears to capture these oscillations effectively.

Table 12 Input variables of the Omie Price forecast model

Name	Description	Units
Electricity Price	Price at which the power plant is selling electricity (Last 24 hours of data).	€
sin(15min)	Artificial variable, sinus wave with period of 1 day (Last 24 hours of data).	-
sin(dayofyear)	Artificial variable, sinus wave with period of 1 year (Last 24 hours of data).	-

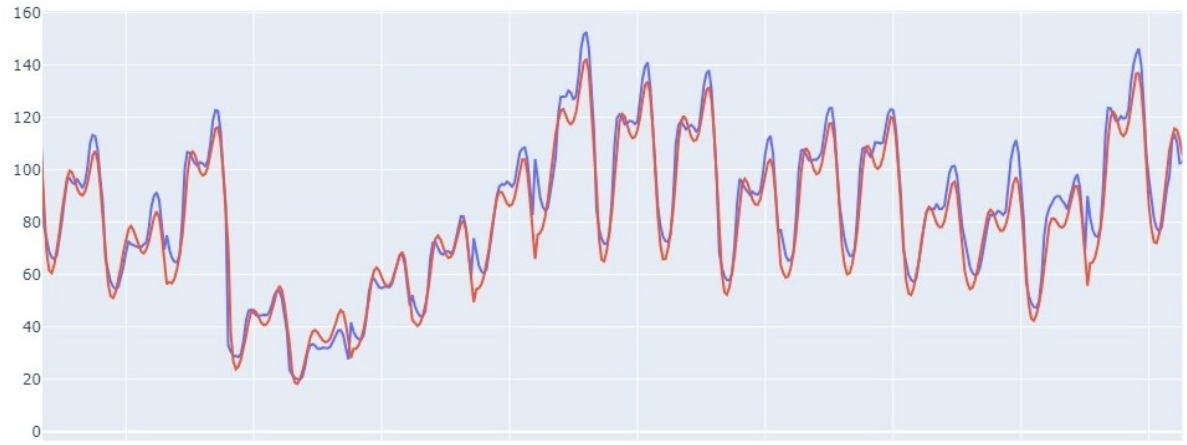

Figure 62 OMIE price forecasting results on test data.

6.4.3 Combined heat and power recurrent simulator

With the DHC demand and OMIE forecast available, we need to determine how to meet the expected heat and cold demands. As we plan to simulate the next six hours recursively, changing only the OMIE, gas prices, and demand forecast inputs, the power plant simulator must be recursive.

This simulator is built using machine learning models, where each machine simulation uses inputs from other machine simulation outputs, forecasting models, or action variables, such as the supply temperature of the DHC. This configuration allows us to recurrently calculate the state of the CHP system using only the action variables.

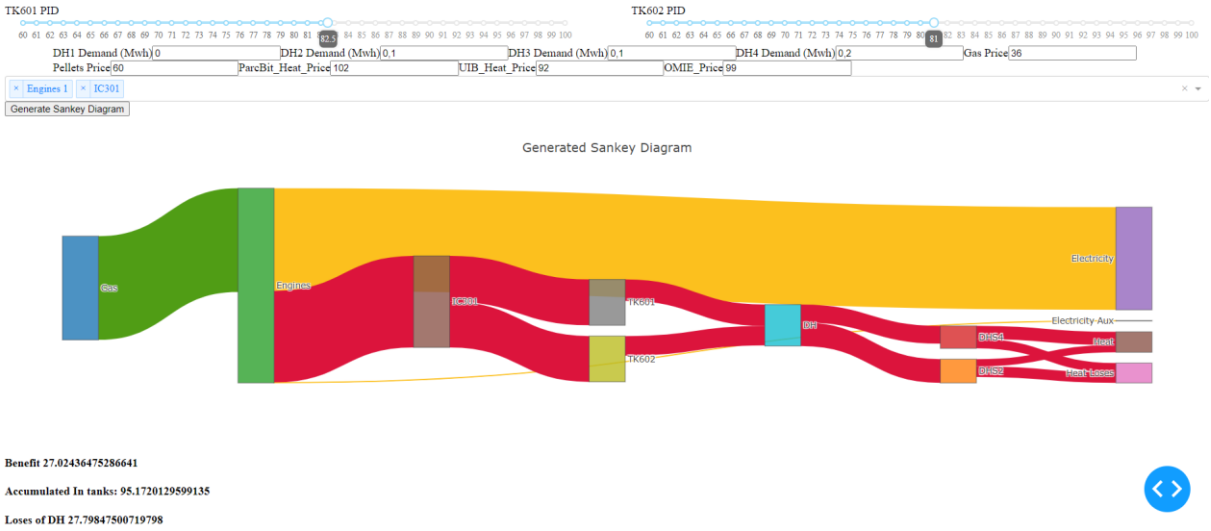

Figure 63 CHP operation simulation, economical sankey diagram

Our reinforcement learning agent optimizes the next 6 hours of operation with recommendations every 15 minutes using the CHP-DHC environment. The calculation that the agent uses as reward is the gross accumulated margin of the whole CHP-DHC system:

$$R = \sum_{i=0}^{24} \gamma^i R_i$$

Where Ri is margin calculated on the forecasted quantities:

$$R_i = HeatSold_i \cdot HeatPrice_i + ElectricitySold_i \cdot El.Price_i - PelletsCons_i \cdot PelletsPrice_i - GasCons_i \cdot GasPrice_i$$

The reward function in a reinforcement learning algorithm, especially one with a decay value, is crucial for guiding the agent's learning process. It provides essential feedback on actions, helps optimise policies, balances short-term and long-term goals, and enhances the agent's ability to navigate and learn effectively in complex, uncertain environments. In our case, adding the decay value allowed the agent to optimise the entire period while focusing more on the early stages of the forecast. This approach enables the agent to learn to optimise the whole period, decide to accumulate energy in the tanks, and simultaneously optimise short-term operations.

Figure 64 illustrates the results from the initial machine learning strategy for optimisation. The first graph reveals the total gains (profit in euros), the middle graph displays the control points for the tanks' outputs, and the last graph presents the predicted needs for each district's heating system. During this process, the algorithm opted to increase the temperature for the tank accumulation. This decision, influenced by the relationship between the cost of gas and the selling price of electricity, led to reducing CHPs activation, which resulted paradoxically in financial losses rather than gains.

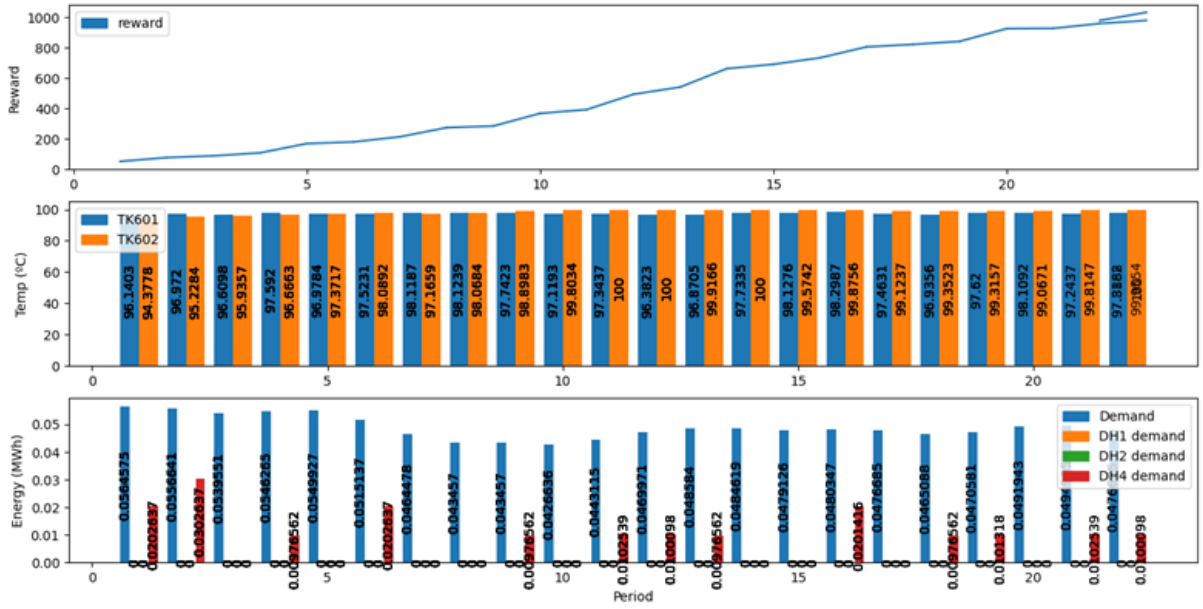

Figure 64 Optimized operation.

6.5 Assessment of the control performance

The average training MSE across all models is approximately 9.76%, while the average testing MSE is slightly higher at 11.44%, indicating that the models perform consistently, although they encounter some additional error during real-world application. The average training MAE is 8.50, and the testing MAE is 9.62, further supporting the observation that model accuracy slightly decreases when applied to unseen data. In terms of percentage error, the average MAPE for training is 5.18%, whereas for testing, it increases to 6.62%. These metrics suggest that while the models generally perform well, there is room for improvement, particularly in reducing the error during testing to ensure more robust predictions in live operations. The increase in error from training to testing phases highlights the importance of continuous model refinement and

retraining to better capture the complexities of the CHP and DHC systems under varying conditions.

The smart control was tested on historical data to benchmark its performance against traditional operational strategies. In Figure 65, the left y-axis displays the rewards calculated for the baseline operation (blue), the optimized operation (orange), and the accumulated difference (green). The right y-axis shows the aggregated heat demand of the district heating system (red), ambient temperature (blue), and the differential between the baseline and optimized operations.

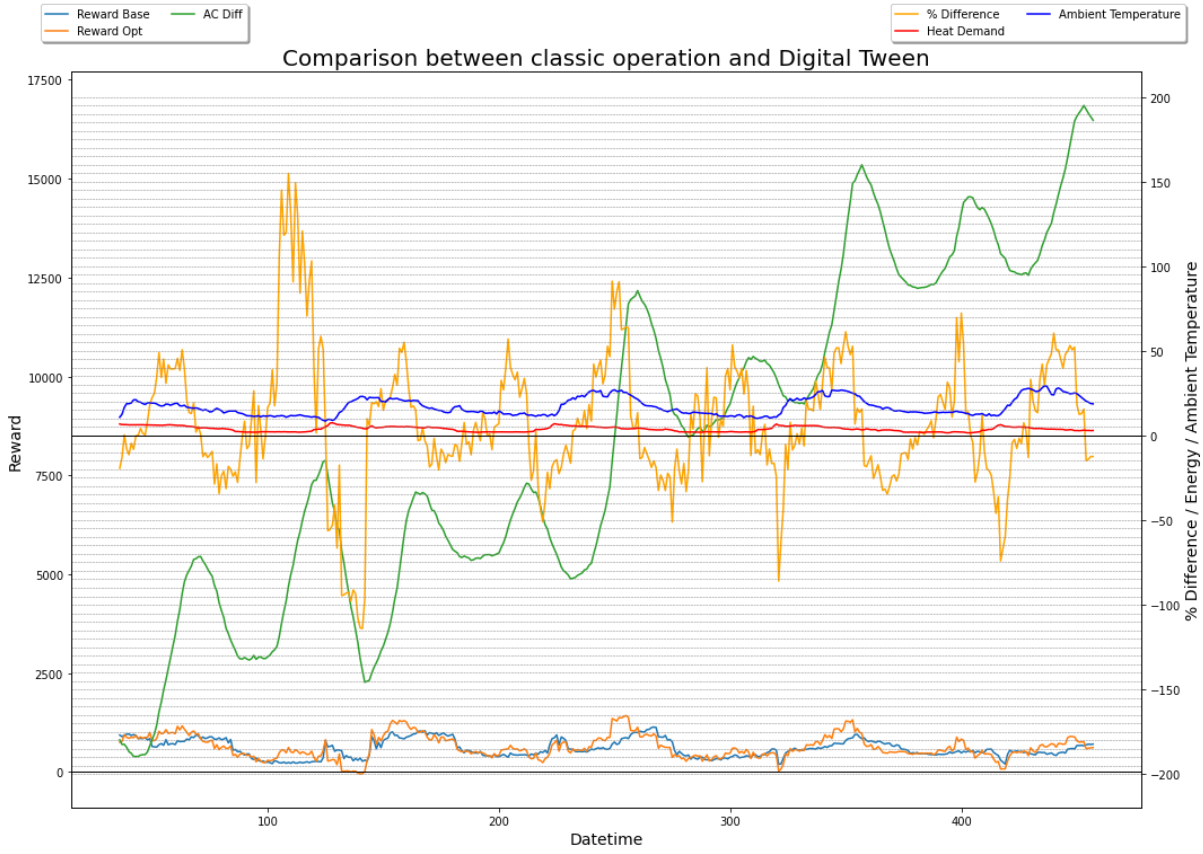

Figure 65 Operation Comparison between Digital Tween and Classic Operation strategy

The Digital Twin demonstrates its capability to optimize the scheduling of active machines in the power plant, manage heat accumulation in the tank, and regulate the supply temperature for district heating. Several conclusions can be derived:

- The Digital Twin does not consistently outperform the baseline in every 15-minute interval. This variability is due to the flexibility gained by adjusting the accumulation tank temperatures, which allows the Digital Twin to align heat generation with periods of high electricity prices.
- By leveraging this flexibility, the Digital Twin effectively coordinates the operation of engines, heat exchangers, and absorption chillers during peak electricity price periods, thereby enhancing the overall efficiency and profitability of heat generation. This increase in efficiency and benefits, however, must account for the losses incurred during heat accumulation.
- During the observed week, the Digital Twin capitalized on higher electricity prices in the mornings, primarily due to the reduced availability of solar power in the market. The system generates electricity during these periods, utilizing engine refrigeration and exhaust gases to increase the temperature in the accumulation tanks, while supplying the district heating

system at the lowest available temperature (60°C). Although this strategy leads to higher morning profits compared to the baseline, the cessation of engine operations at midday results in lower afternoon benefits due to the shift from electricity selling to consumption for pump operations, which favours the baseline operation instead.

- Overall, the Digital Twin operation yields an average daily increase in power plant revenues of approximately 7.95% compared to the baseline.
- As we reduced the operation temperature of the district heating from an average of 80 ºC to 65ºC, the loses of both the accumulation tanks and the distribution network deceased about a 22.12%.

6.6 Lessons learned

The main learning out of operating the data mining tool and the advanced controls embedded are:

- Advanced forecasting tools are of utmost importance: without the forecast tools developed for OMIE (electricity price), MIBGAS (natural gas price) and DHC demands, scheduling and optimizing the operation of the CHP-DHDC have shown to be not feasible.
- Flexibility between generation and demand is necessary to decrease costs and emissions, moving the generation into the high OMIE hours of the day allowed us to use the engines and absorption machines that are much more efficient than the pellets or gas burners.
- Lowering the temperature of a district heating is not always an easy task and highly depends on contracts with the prosumers, going beyond certain temperate needs the installation of smart substations with heat pumps that is not always feasible.
- The original data stored in the SCADA, used to operate a system using classical operation methods, may not be gathering all the necessary data values to be able to create a digital simulation of the physical power plant.
- Defining the actuators and the simulation loop in a CHP power plant can significantly increase the complexity of the simulation. The addition of an initial scheduler allows the live optimizer to use accumulation as a resource and increases the efficiency of the CHP power plant.
- Using Reinforcement Learning to simulate such a complex system as a CHP power plant is not straightforward, the actuators and the environment play a vital role on the accuracy and speed of the agent training process. We had to simplify the system variables to obtain the expected results.

7 Szczecin (SEC-E.ON)

7.1 Description of the demonstration site

Szczecin's demonstration site is located on the Łasztownia river island, and constitutes the first case of low temperature, hybrid DHC network in Poland.

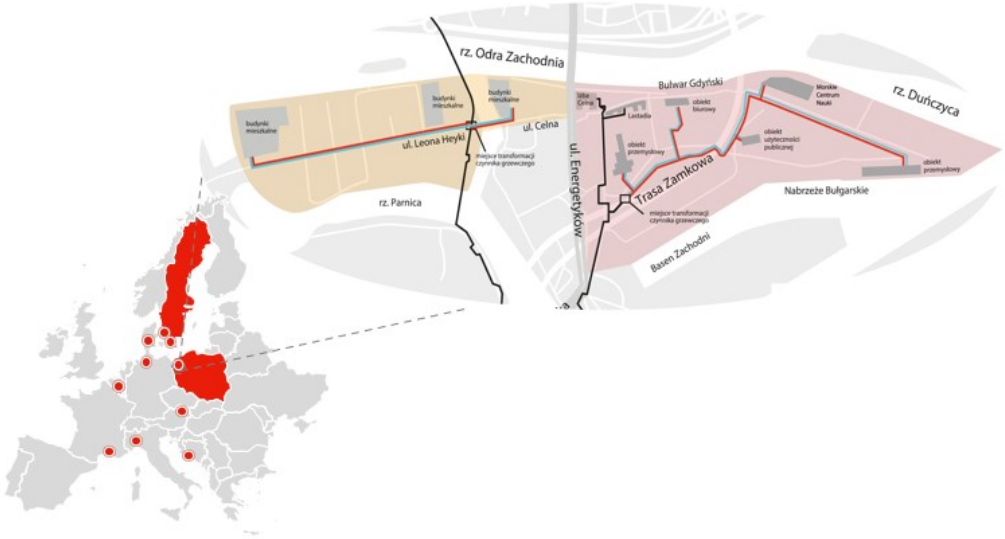

Figure 66: Łasztownia Island. On the Right, red highlighted, the area interested by the REWARDHeat Project

The newly developed network follows the E.ON ectogrid™ concept. Typical operating temperatures range from 30-50 °C at the supply pipe in summer and winter, respectively, and 25-35 °C at the return. A Heat Balancing Station (HBS) connects the new network to the existing high-temperature district heating (DH) system and adiabatic coolers, enabling the distribution of both heating and cooling to users. Currently, the HBS supplies the Maritime Science Centre (MSC) with approximately 400 kW for heating and 600 kW for cooling. It is designed to support future expansion, covering 3450 kW for heating and 2250 kW for cooling.

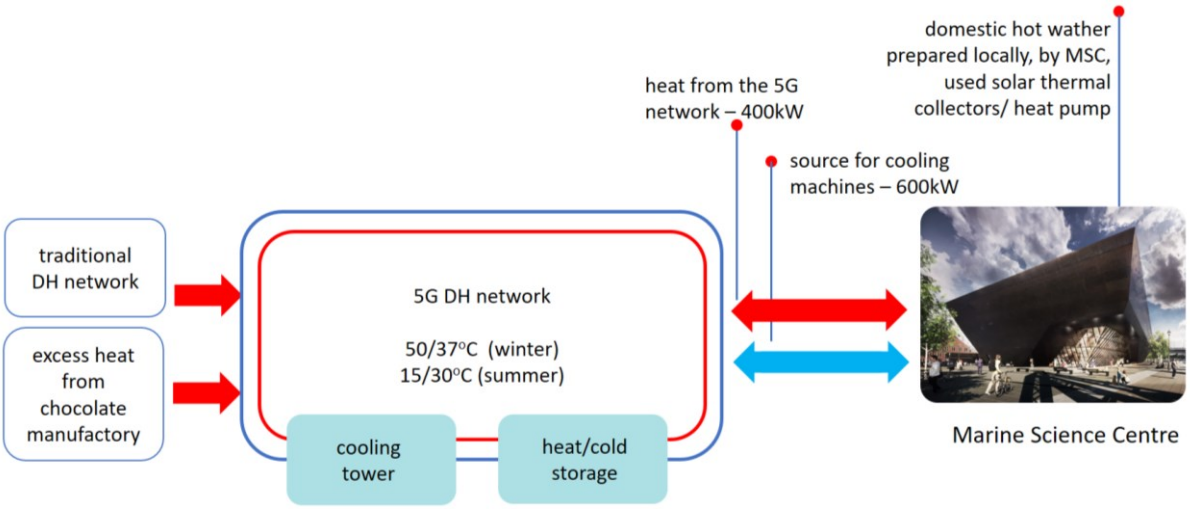

Figure 67: Łasztownia Island – 5G DH network

The MSC's space heating and DHW loads are covered with direct heating in winter. Polyvalent, 4-pipes, water-water chillers cover DHW and space cooling needs in mid- and summer seasons and reject waste heat into the network. The chillers draw water from the network's return pipe, ensuring the condenser operates within its permitted temperature ranges throughout the year.

All assets on the secondary side at the MSC are supervised by a local BMS and managed by the building energy manager. Both, HBS and MSC transfer data in two directions to and from the DHC network SCADA system.

7.1.1 Objectives of the demonstration activity

The scope of the demonstration activity is to supply heating and cooling to the MSC. The MSC needed thermal capacity was calculated in about 400 kW @ 45/35°C for what concerns heating, and 600 kW @ 7/12°C for cooling. Specifically, the objectives pursued during the project are:

Development and set up of the initial low temperature DHC network backbone;

Construction of the Heat Balancing Station (HBS) and set up of the first customer substation;

Implementation of the HBS management system, aimed to optimize mass flow and temperature in the network based on demand side loads assessment.

Assessment and improvement of the solutions implemented and development of recommendations for replication

Analysis of the available waste heat on the island from identified potential sources: a chocolate factory and a food cold storage facility.

7.2 Control problem formulation

As in the other cases, which we reported about previously (see section 4 and 5), also in this case, a network operated with low-temperature and low-DT between supply and return, requires special attention in terms of management in order to maximise thermal capacity distributed and efficient share of energy among prosumers. Specifically with respect to the smart control development, the optimization goals are:

to minimize energy drawn from municipal high-temperature DH network, by reducing heat distribution losses and maximising waste heat recovery from cooling the MSC;

to predict heating and cooling demands and cover them with stored energy as much as possible;

in the long run, balance out heating and cooling loads along the network and among the users connected.

7.3 Description of the control implemented

7.3.1 Heat Balancing Station

The Heat Balancing Station (HBS) is designed to maintain the desired temperatures in the warm and cold pipes using the ectocloud™ system. This system determines the average temperature of the thermal buffer tank to ensure that either heating or cooling is produced, while keeping the storage temperature within the specified range. The control system ensures that the storage temperature never exceeds the boundaries of the desired temperatures, thus optimising the

system's performance for both heating and cooling, while minimising the energy drawn from the conventional district heating backbone during winters.

Automation programming has been implemented to ensure smooth and robust operation of the Heat Balancing Station (HBS). This includes the integration of necessary interfaces and supporting functionality to enable intelligent control via the ectocloud™ system. The automation ensures that the system can dynamically adjust to operational demands, optimising both heating and cooling processes while maintaining the required temperature range for efficient energy management.

The main requirements on the local control at HBS is to in real-time:

Produce the right amount of heating and cooling loads

Shift between heating and cooling as indicated by the grid demand balance

Supply the right temperature to the network cold or warm pipe as requested by ectocloud™

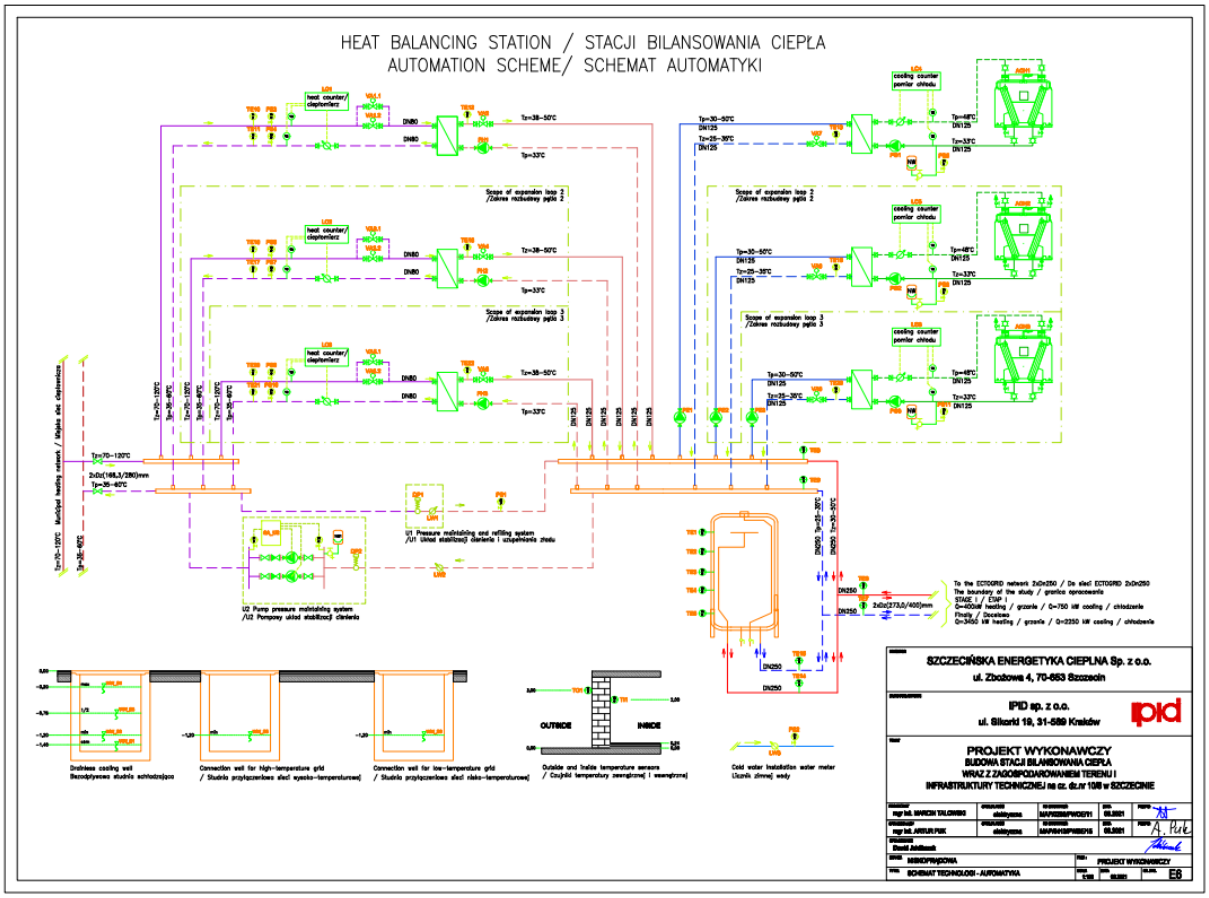

Figure 68 – Heat Balancing Station P&I diagram

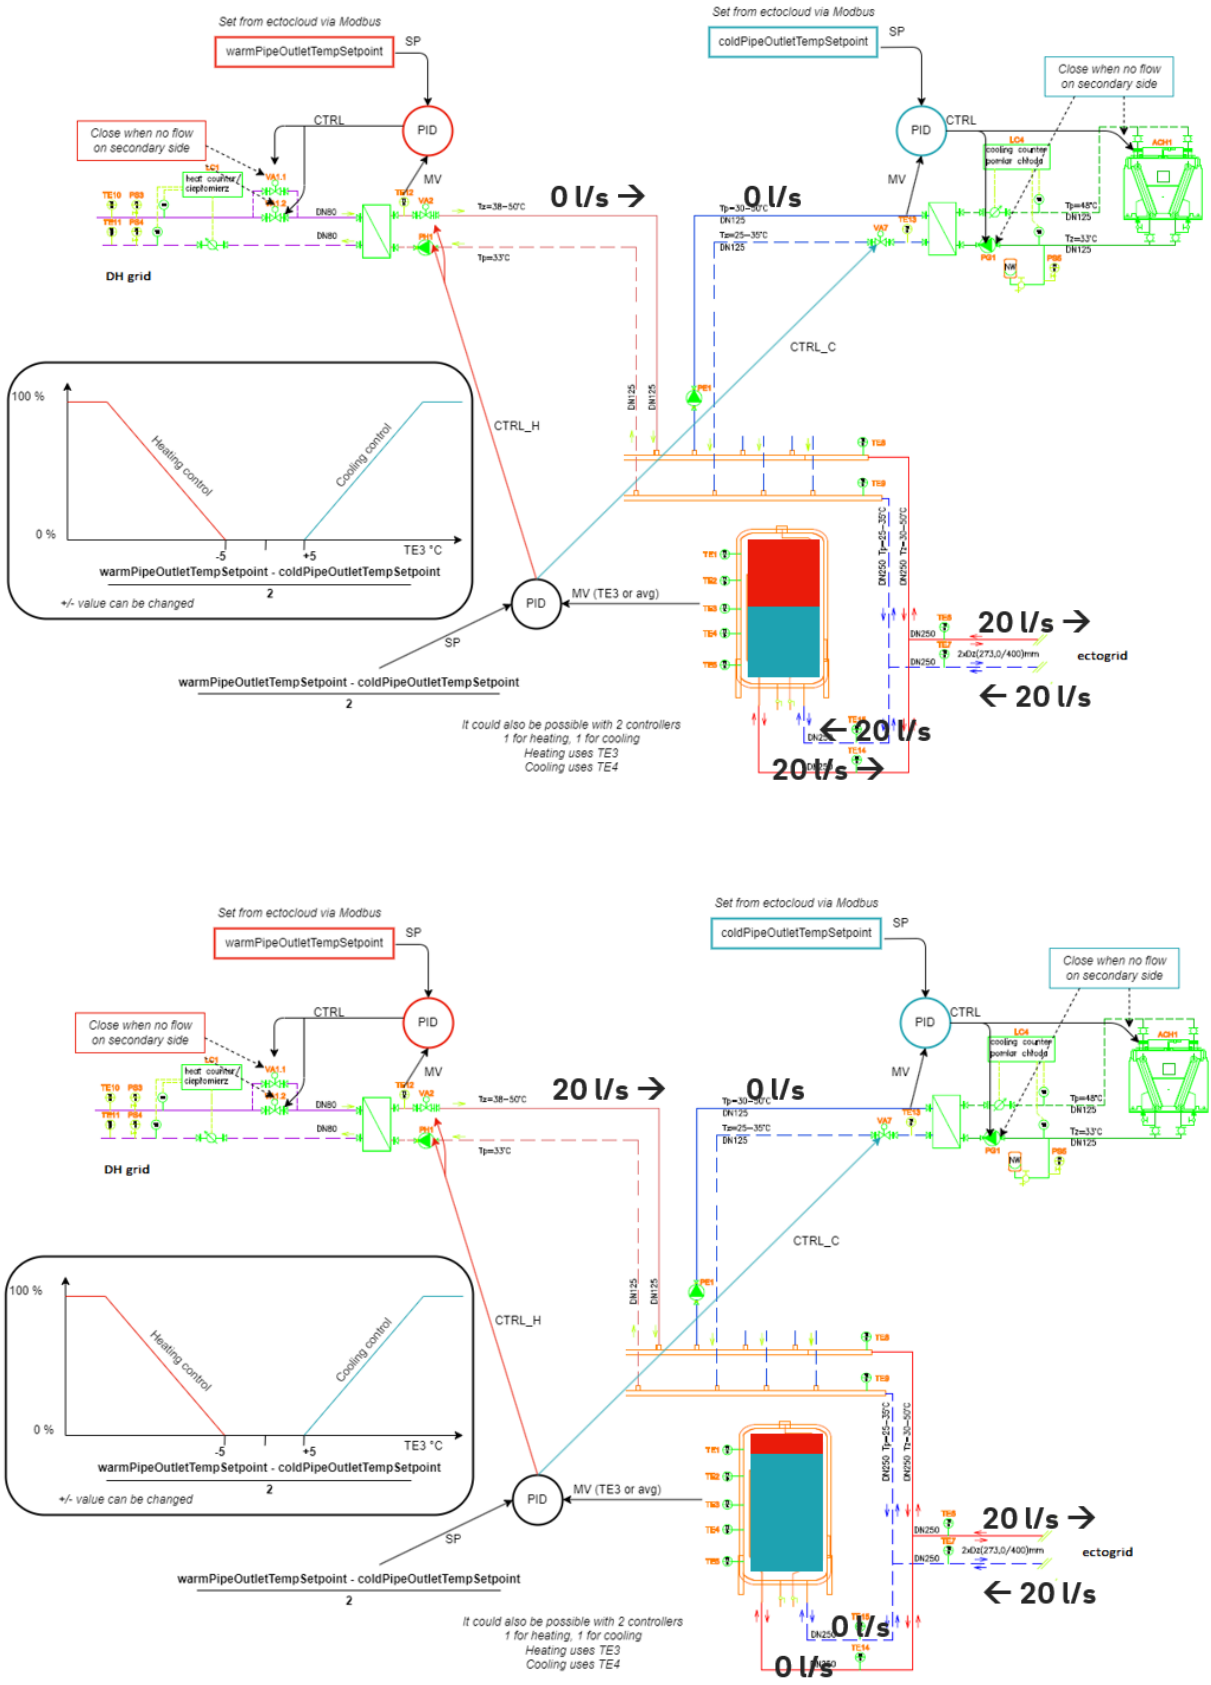

Figure 69 – Control concept sketch of thermal buffer tank control aimed to maintain supply temperatures to the Łasztownia low-temperature district heating and cooling grid within the desired boundaries

7.3.2 Maritime Science Center substation

At the Maritime Science Centre (MSC), the local control system works to meet the building's heating and cooling demands, while simultaneously aiming to return energy to the network at the appropriate temperature. This is mainly achievable for excess heat generated by the building's internal chillers. By precisely controlling this temperature, the system allows for internal heat recovery at the optimal temperature, while also aligning with broader system considerations such as energy balancing and waste heat integration, which will become increasingly important in the future.

To achieve this, the substation shown within the dashed contour in Figure 70 supplies hot water to meet heating demands through the hot pipe (in red) and supports the chillers through the cold pipe (in blue). Although the building generally requires either heating or cooling, this configuration enables the system to deliver both when needed. Sufficiently high temperatures can be supplied via the hot pipe, while the cold pipe maintains water temperatures compatible with chiller operations.

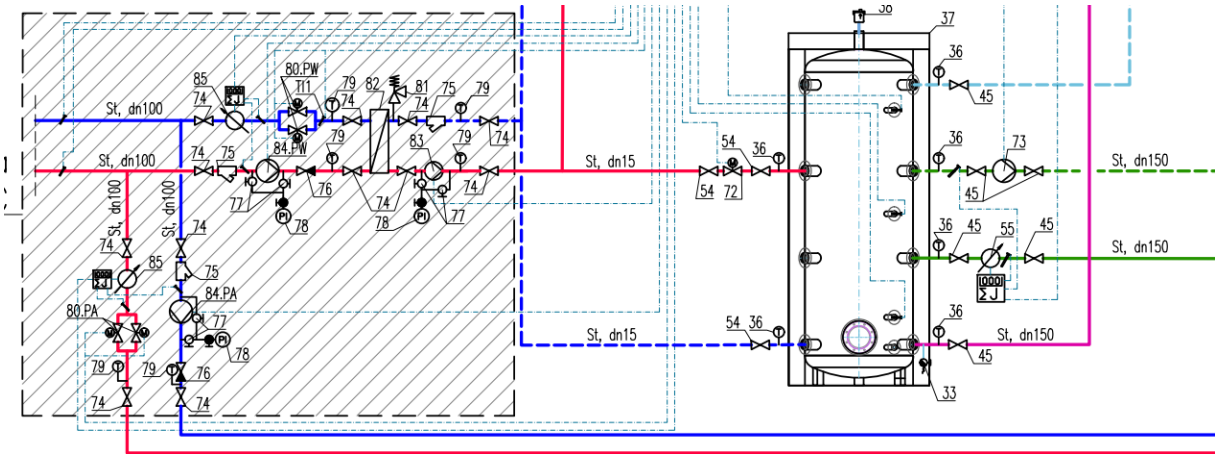

Figure 70 – DHC substation at prosumer side

7.3.3 SCADA and Energy Analytics System

Both the Heat Balancing Station and the Maritime Science Centre are integrated with the existing SCADA system at the Municipal Grid Control Room, allowing real-time monitoring of values and manual control of actuators if necessary. Data for the SCADA is transferred via a router with redundant communication channels.

For the commissioning of the digital energy management platform, ectocloud™, the installation of an IoT Gateway is required. This device facilitates bidirectional communication between ectocloud™ and the local control system, with communication conducted over secure protocols. The Energy Manager is a mandatory component for the supply substation and for all end-users connected to the low-temperature grid. Typically, it is installed on a DIN rail inside control cabinets near the PLC or the switch/router. It is powered by a 24VDC supply and requires 60W. It is equipped with an Ethernet port and a SIM card slot for GSM communication. It connects to the BMS to:

collect heat/cooling demand data;

transfer heat/cooling meter data;

transmit control signals.

To communicate with the PLC, the system uses Modbus TCP/Modbus RTU in master or slave mode. It employs outgoing WAN protocols with data encryption such as MQTT SSL/TLS, AMQP SSL/TLS, HTTPS, and SSH.

Before transmission, data is briefly stored. All communication links are continuously monitored, and each has a fallback strategy in case of failure. A "watchdog" feature monitors communication with the PLC, triggering alarms and activating the fallback strategy in case of faults.

Once the digital infrastructure is in place, the platform provides users with interfaces and data analytics tools to monitor the energy system in near real-time, track aggregated energy data, and follow key system performance indicators (KPIs). The platform enables analysis of the internal exchange of heating and cooling within the Maritime Science Centre, as well as the net demand supplied by the low-temperature grid.

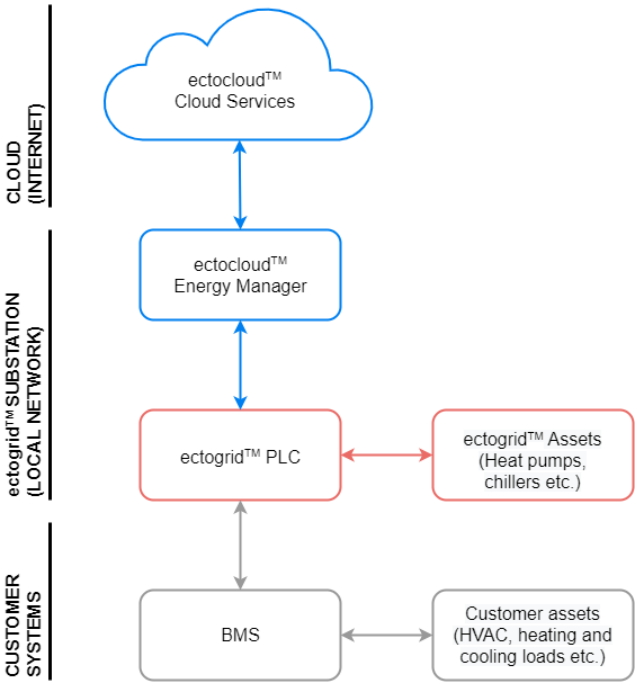

Figure 71 – Communication architecture implemented

Figure 72 – Typical control cabinet (left) and typical E.ON Energy Manager mounting (right)

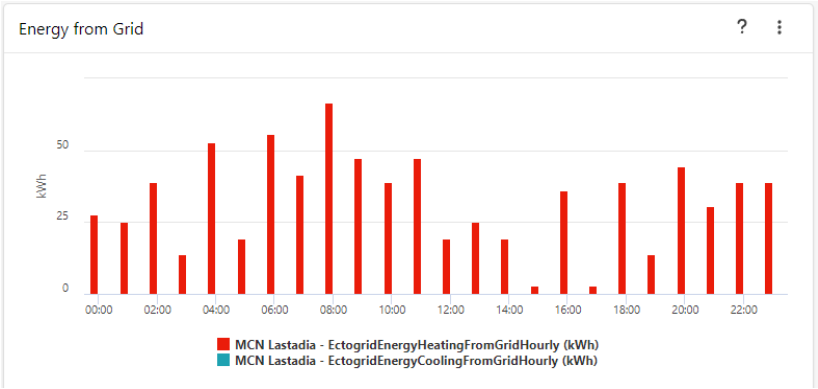

Figure 73 - Example of net load supplied by the low-temperature grid

Figure 74 – Example of heat balancing station assessment

7.3.4 Advanced controls

The advanced control of the ectogrid™ on Łasztownia Island focuses on steering network temperature to achieve system-wide benefits. After establishing basic control at the Heat Balancing Station (HBS) and Maritime Science Centre (MSC) to allow adjustments in emitted temperatures, an overarching control was implemented to coordinate grid temperatures and demonstrate the advantages of dynamically shifting temperatures based on the heating and cooling demand balance.

Managing thermal energy supply and distribution relies on historical data collected in the cloud, which is processed using a machine learning-based tool. This data reveals trends, customer behaviour, and supports the creation of consumption profiles. Considering weather forecast inputs, the tool provides predictions of expected demand.

Energy forecasting has been developed to support intelligent control strategies, forming the basis for anticipating demand for current and future advanced control scenarios. Predicted demand data is valuable for preloading buffer tanks with heat or cool energy in anticipation of periodic peak demands. Heat/cool storage, along with grid volume capacity, increases system inertia and helps minimise temporary reliance on the high-temperature district heating backbone.

Through demand measurement and machine learning-based forecasting, alongside temperature impact assessments on various system components, temperature setpoints for the warm and cold pipes were distributed hourly to the HBS and MSC.

7.4 Assessment of the control performance

The energy load forecasting of heating demand from the grid resulted in the following model training performance:

R2: 0.94

RMSE: 8 kWh

With the most dominant predictor being the 6-hour average outdoor temperature. As for the cooling demand forecast, the following training performance was achieved:

R2: 0.86

RMSE: 16 kWh.

Mostly the outdoor temperature on a 1-hour and 24-hour average basis plays a significant role in determining the forecast output, but also the time of day which indicated that on average more cooling is needed in afternoons.

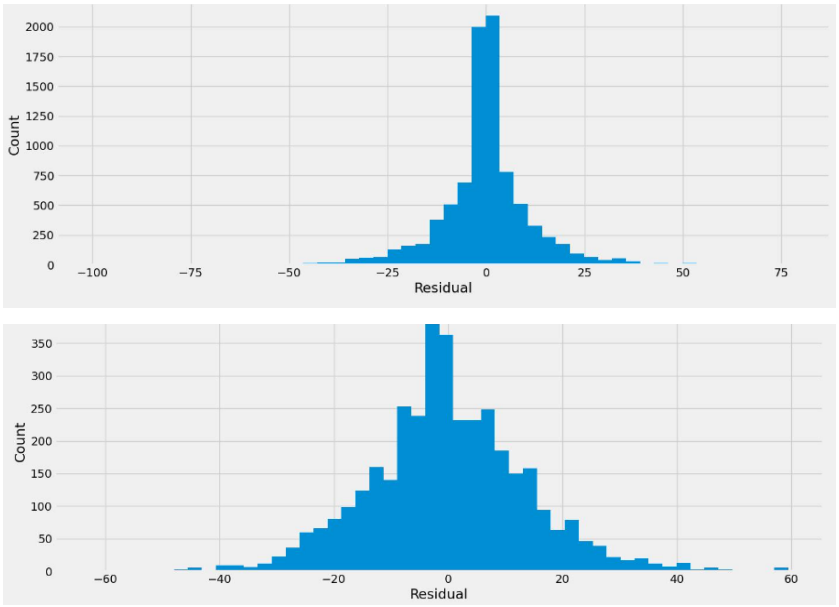

Figure 75 – Distribution of error achieved when training the heating (top plot) and cooling demands (bottom plot)

Although not all necessary data is currently available for a complete energy impact analysis of the control system, thermodynamic principles suggest that certain benefits are likely realised, specifically:

Lowered distribution losses due to dynamic control of the grid temperature

Increased efficiency of waste heat recovery at MSC (excess heat from chillers)

Efficient use of adiabatic cooling at HBS in relation to distribution

From the periods when network temperature control was tested, it was evident that the full temperature range was utilised, generally lowering grid temperatures as often as possible. During an analysed period in spring 2024, the mean temperature was controlled to 41°C in the warm pipe and 33°C in the cold pipe. Heat losses in transmission, compared to a static temperature regime, were found to be significantly reduced—up to around 15% less, depending on soil temperatures—demonstrating potential savings achievable with further analysis and refinement.

Additionally, as grid temperatures were lowered in response to reduced heating demand and increased cooling demand, the efficiency of the MSC chillers improved. Assuming Carnot cycle efficiency, a 10°C reduction in network temperature, as achieved in summer, would yield a chiller efficiency gain of over 30%, while adiabatic coolers at the HBS efficiently managed the release of unused excess heat.

Regarding adherence to overall temperature control, the heat supplied from the HBS and MSC chillers generally aligned with setpoints. However, temperatures at the HBS were often slightly below target, suggesting room for improved local control. Safety margins in setpoints were found to compensate adequately for these variations, with future improvements noted for enhanced control precision. Although local control responds to requested temperature changes, further refinements are likely to allow even better adherence, which has been earmarked for future development.

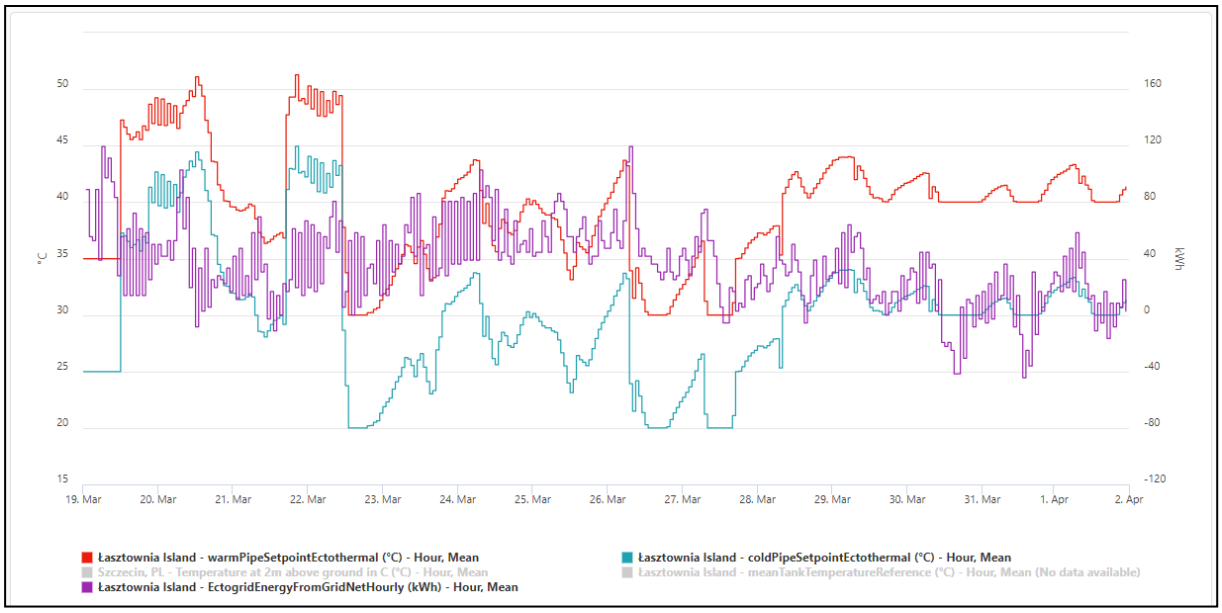

Figure 76 – Warm and cold pipe distribution temperatures and thermal energy distributed

Figure 77 – Heat Balancing Station set points and distribution temperatures achieved

7.5 Lessons learned

While implementing smart analytics and control at Łasztownia, there were some practical lessons learned:

- Implementation of PLC/BMS control logic: Integrating the PLC with the BMS control logic proved challenging, as automation contractors were not thoroughly acquainted with the innovative energy concepts used in this project. This unfamiliarity required additional time and support for accurate implementation, which could influence system setup and optimisation.
- Challenges in data collection: Acquiring all relevant data necessary to comprehensively analyse system effects and energy performance was hindered by network limitations, data access issues, and security concerns. These challenges impacted the ability to obtain a full, real-time picture of system performance, restricting data-driven adjustments.
- Importance of multilevel control attention: Focusing on all control levels, from real-time equipment regulation within substations to system-wide optimisation, is critical to fully unlock the potential of the energy system. Close attention to these layers ensures that performance gains are achieved at both local and network scales.
- Network temperature influence: The temperature within the network can be dynamically adjusted through intelligent control algorithms, resulting in tangible system benefits. This capability enhances the adaptability of the network to changing demands and operational conditions.
- Demand-Based temperature modulation: The flexibility to adjust temperatures based on varying demand presents advantages in energy production, recovery, and distribution. By tailoring temperature settings to demand, the system can achieve improved efficiency and heat recovery.

8 Topusko (ENISYST)

8.1 Description of the demonstration site

The area of Topusko is rich in thermal springs. The concessionaire for the extraction of geothermal hot water is Health Spa Topusko and TOP-TERME LCC, who manage a large healthcare structure including hotels, mud baths and swimming pools.

Heating and DHW preparation for all buildings and facilities takes place inside the central thermal station (CTS, Figure 78): geothermal water is collected from the TEB-4 well @ 62 °C and used to condition technical water to the different temperature levels needed; technical water is then distributed from here to the uses.

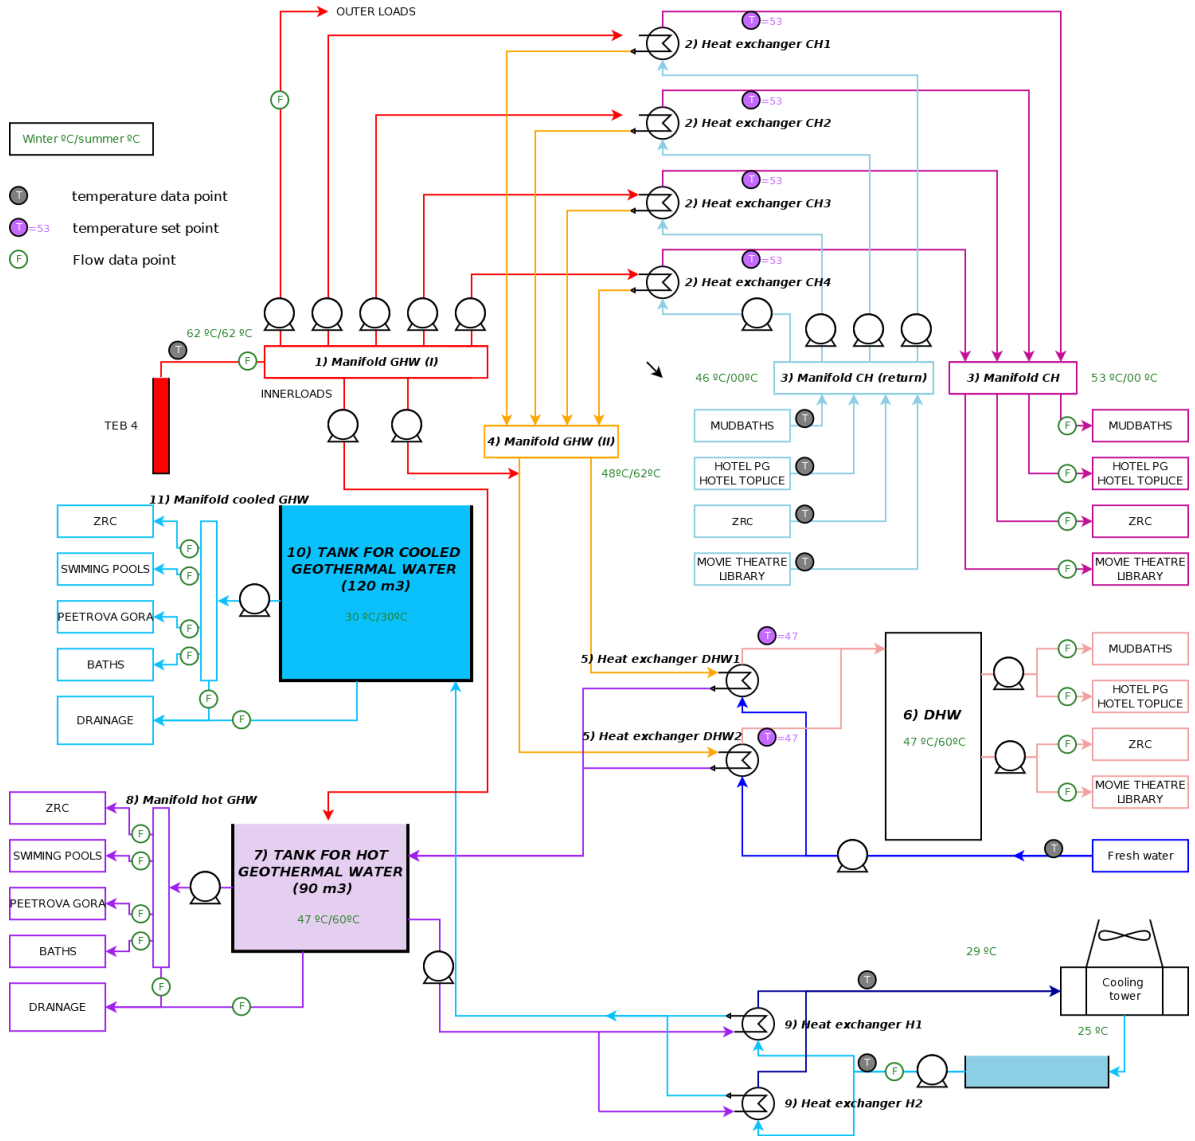

Figure 78: Scheme of energy facilities in Topusko.

Four heat exchangers, 712 kW_{th} capacity each, generate heating for all buildings @ 50°C. From here, geothermal water is cascaded to the heat exchangers devoted to the preparation of DHW stored @47 °C in one 10m^3 TES, and further distributed to the water storages for balneology and

swimming pools. Two concrete tanks store water at high temperature (90 m^3 @ 47°C) and low temperature (120 m^3 @ 30°C). In the latter case, water is preliminary cooled by means of a wet cooling tower. Water from the two TES is distributed to baths or swimming pools where it is mixed to obtain the needed temperature level.

In addition to distributing geothermal and technical water throughout the healthcare facility, high electricity consumption is also required to operate the wet cooling tower. This demand is particularly pronounced during the summer when there is no need for space heating, hence higher temperature water is available after DHW preparation.

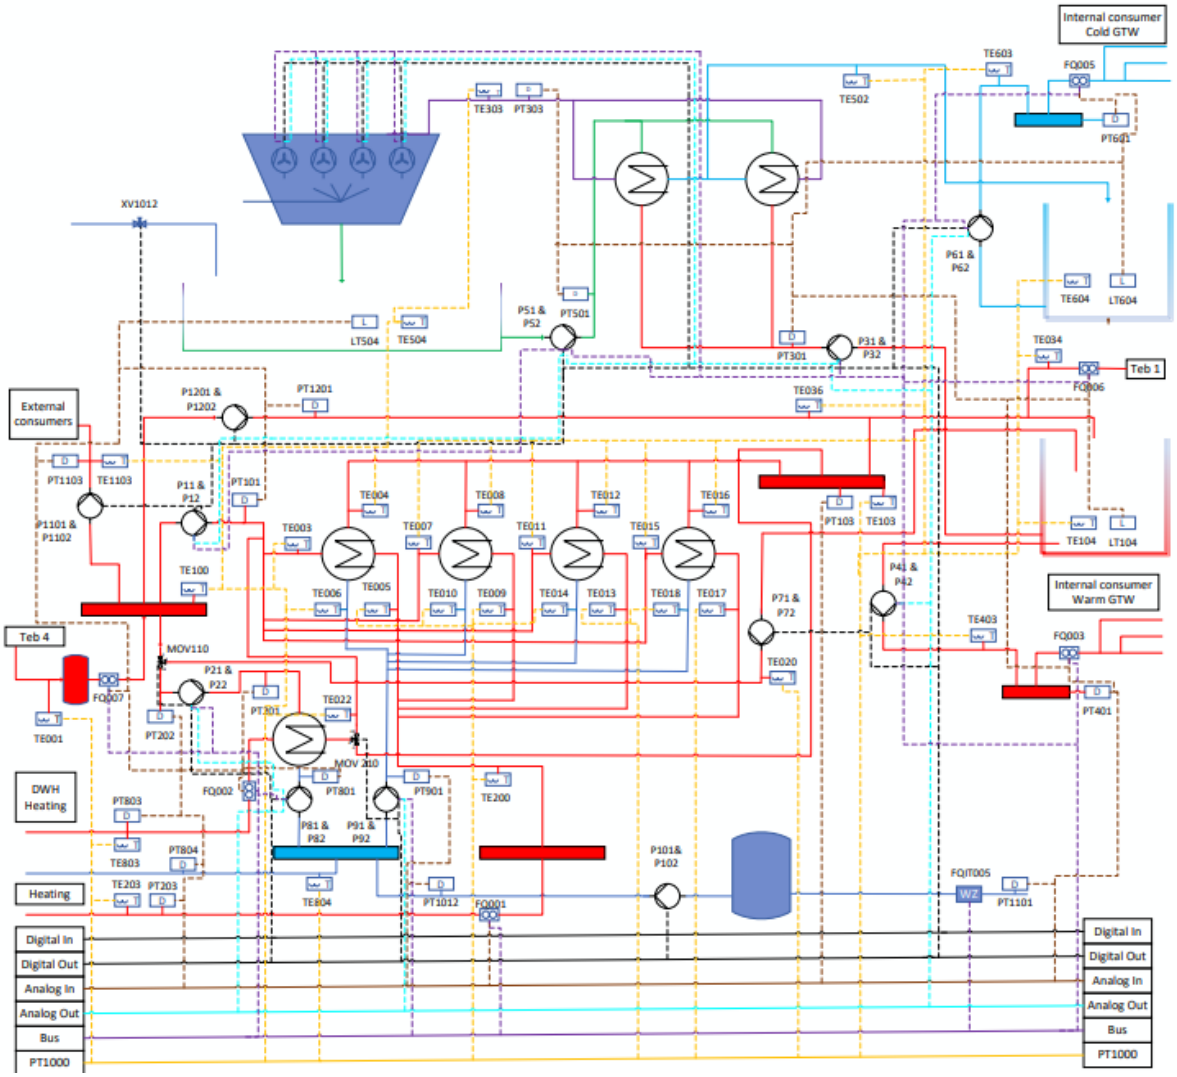

Figure 79 – Scheme of the central heat station (CTS) with flow chart of the control signals.

8.1.1 Objectives of the demonstration activity

Within REWARDHeat a complete refurbishment of the CTS and of the structure's DH network has been performed, aimed to significantly increase the overall system energy efficiency. The following measures have been implemented:

- All pipelines from TEB-4 to CTS and from CTS to all uses have been replaced with insulated ones;

- CTS pipes have been substituted, constant speed pumps and wet cooling tower's fans have been replaced with inverter-controlled, variable-speed ones, and manual valves have been replaced with remotely controlled ones. This allowed to implement an active control of both mass flows and temperature set points at the CTS.

The control hardware of the whole system at the CTS has been renovated by ENISYST, as at project start the system was operated manually, including the following measures (see Figure 79):

- Temperature sensors have been installed along the pipelines, DHW and geothermal water TES.
- Flow meters have been installed along the pipelines to get a complete view of the water flows and the thermal energy distributed within and from the CTS.
- Electric meters have been used to measure the electricity consumption of the whole CTS and of the wet cooling tower.
- Although building substations could not be substituted nor can be controlled actively as they are owned by the building owners, heat meters have been installed also and connected to the CTS via LoraWAN technology, aimed to monitor space heating delivery and DHW preparation, hence allowing to forecast future energy uses.

8.2 Control problem formulation

To increase the overall efficiency of the system apart from the CTS hardware and thermal network refurbishment, an active control strategy has been implemented, including:

- Control of the mass flow, supply and return temperatures in the three-tube district heating network distributing energy to the buildings;
- Control of the geothermal water supply temperature and mass flow to the baths and swimming pools and its temperature in the dedicated TES;
- Control of the cooling tower's fans speed.

The three control problems of course strongly influence one another, hence requiring a high-level control optimisation to guarantee the optimum use of geothermal water and to minimize the electricity consumption. In the following the three control problems and their specific challenges are described. Afterwards, the overall optimised control is described.

8.3 Description of the control implemented

8.3.1 SCADA system

ENISYST delivered control cabinets which include the controller, internet router, IO-modules, low voltage power supply, ethernet switches, electrical fuses, emergency control level for manual operation, and a tablet for the access to the web-based control system. For the new controller generation, which is used as edge device / IoT gateway, ENISYST has evaluated different market available controllers and carried out software and operational performance tests. The Avnet AVTSE-RPI-IIOTG controller, has been selected for the system in Topusko. It includes Wifi and BLE on board, Internal Raspberry Pi HAT, mPCIe and USB ports.

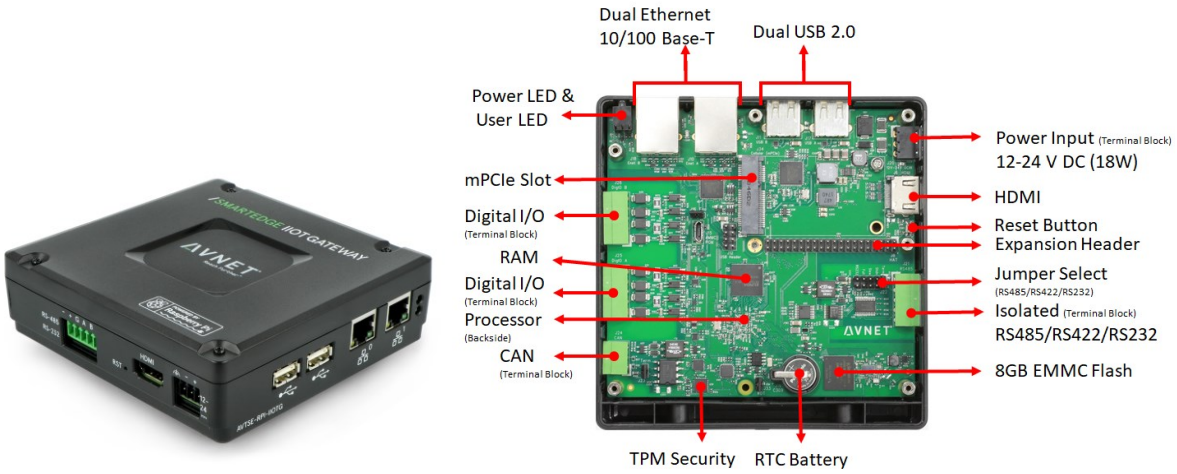

Figure 80 – Housing (left) and mainboard (right) of the Smart Edge device AVNET.

ENISYST software integrates the following features:

- eni.serv, which is a cloud-based high-level district/property management and control system running on a server in the cloud;
- eni.web, which is a low-level WEB based monitoring and control system running on the controller on site.

A secure connection via the internet to the district/property management system is implemented which includes a two-factor identification for the access to eni.serv; this is connected via VPN tunnels to the different controllers on site. For security reasons, the VPN tunnels and the https:\\ connections are established by the local controllers. Here the users get an overview of their connected systems, can compare system performances of different sites and can access the local WEB based BMS (eni.web), running on the controllers on site.

The plant operators can access the local BMS of the different sites via eni.serv. From the starting page, different levels and views can be opened like the hydraulic scheme with on-time visualisation operational parameters and system performance; the latter can be assessed by means of line graphs, scatter plots and carpet plots. From here it is also possible to operate single units in manual mode. For smart district or property management eni.serv integrates a central database and load forecast algorithms enabling AI-based controls.

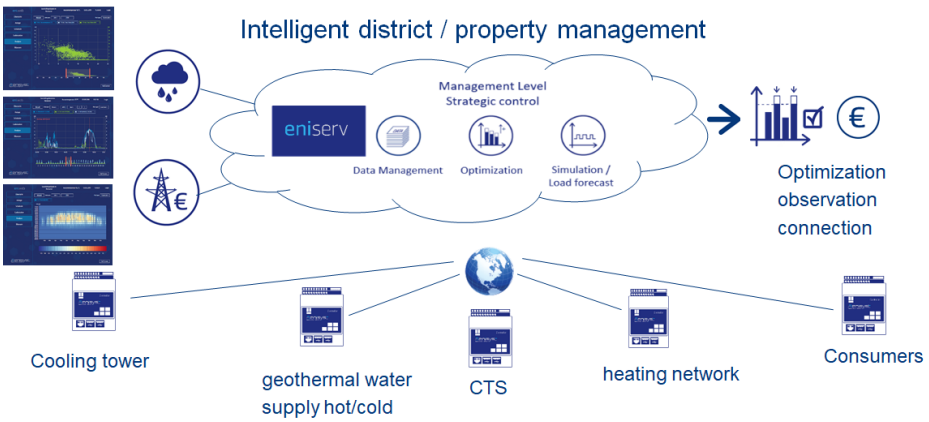

Figure 81 – eni.serv as data management and predictive optimisation platform.

8.3.2 Advanced controls

The three pipe district heating network is built with one pipe devoted to space heating supply, one pipe for domestic hot water preparation supply and one common return. The control of the network is performed with automated two-way valves that open when users require thermal load and are modulated to meet temperature set points. The newly installed, inverter driven pumps maintain pressure drop across the network constant, minimising mass flow distributed and maximising the temperature drop between supply and return. This shows several advantages:

- Electricity consumption for distribution and thermal losses are minimized
- By lowering the return temperature, the cooling load needed to keep the geothermal water at 30°C in the low-temperature concrete tank is also reduced. As a result, the electricity consumption of the wet cooling tower needed for this purpose is minimized.

In fact, geothermal water is pumped into the concrete basins in order to maintain a constant level and a minimum water temperature. The distribution to the baths and swimming pools uses variable speed pumps to maintain a constant pressure at the supply manifold. The wet cooling tower's fans are controlled to maintain a constant return temperature of the geothermal water to the cold basin. This significantly reduces both electricity and freshwater use.

AI-based high-level supervision

The principle of AI-based predictive control is shown in Figure 82. From the measured performance data, a demand model is generated using machine learning and statistical methods.

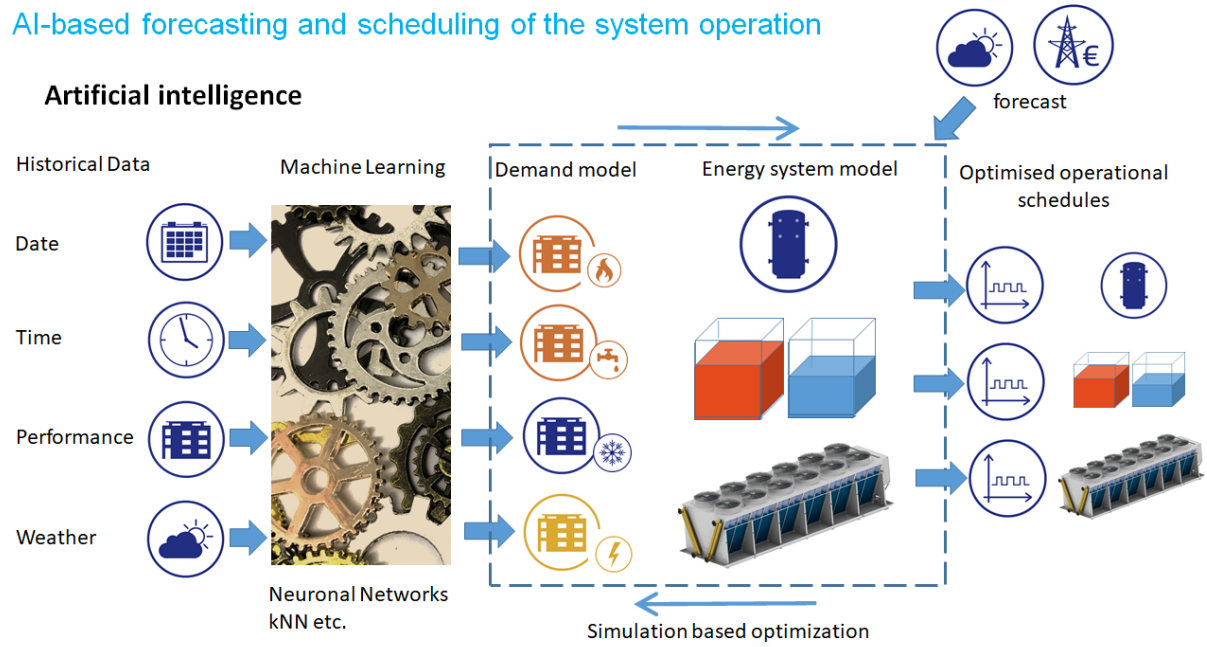

Figure 82 – The principles of the AI based control.

The aim is to forecast demand for the cold and hot water storage tanks in order to minimise the charge levels and reduce energy consumption. Up to now, storages have been charged with thermal water when the level fell below 60%. A predictive control system is now implemented which uses prediction methods.

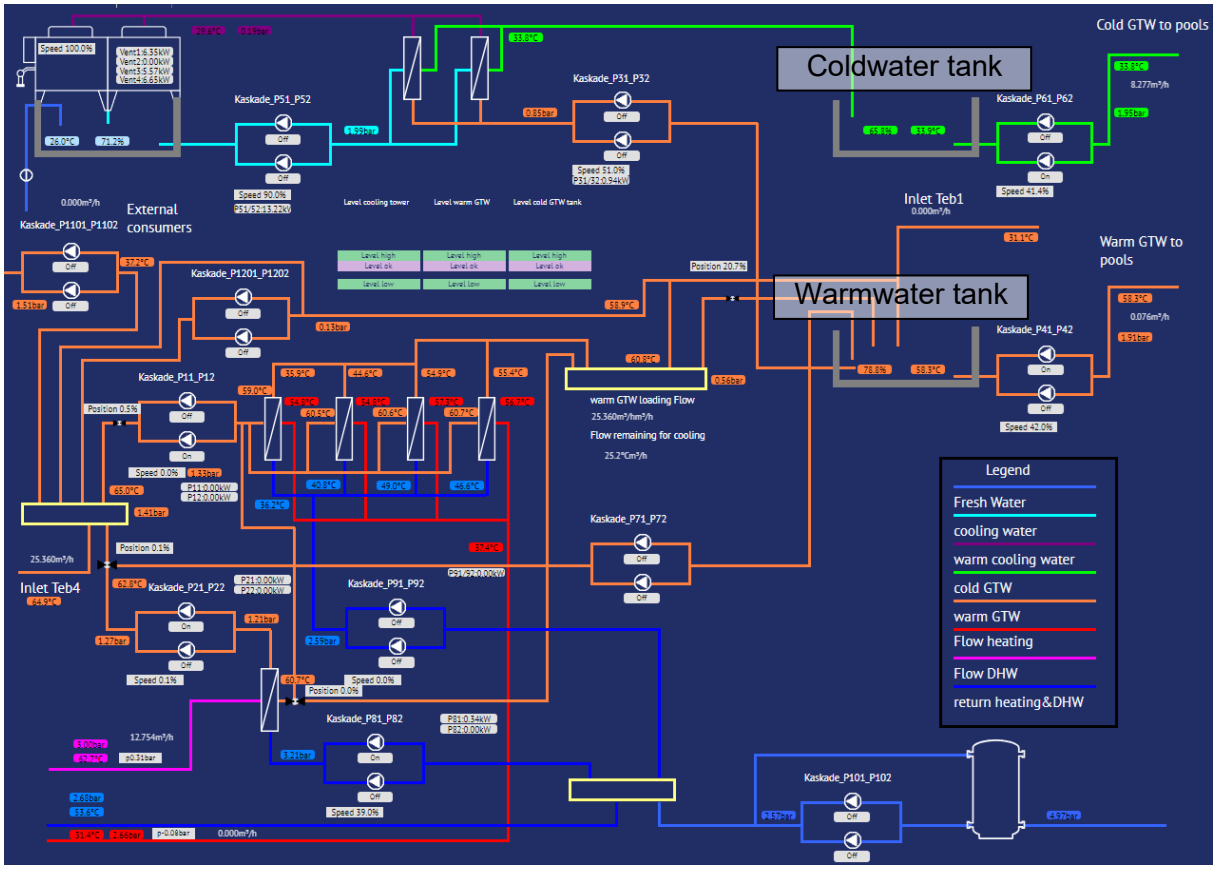

Figure 83 – Topusko system scheme

Prediction method

Various input data was used for model training, including outside temperature, month, day of the week and hour of the day. Weather forecast data from a database and live measurements from the pilot site, such as fill level data and flow rates, are used for the real-time forecast.

In particular, a random forest regressor algorithm, a meta-model based on the principle of the decision tree, was utilized. A random forest trains several decision trees on different subsets of a data set and then averages the results. This serves to increase the accuracy of the predictions and avoids overfitting the model. The decision trees within the random forest use optimal splitting strategies. A 20:80 split was made between the training and test data. To validate the prediction the coefficient of determination and the normalized root mean square error were used.

Coldwater demand

Figure 84 shows the prediction of cold-water demand, split into random train and test subsets. Figure 85 shows the corresponding prediction of cumulated cold-water demand based on the previously trained model. With training and test data obtained in the period from 2024-01-01 till 2024-07-14 results with test data of R2 = 0.82 and a NRMSE = 0.13 m^3 (120 m^3 cold water storage) could be archived. As such, the forecast quality was sufficient to be implemented in the control system.

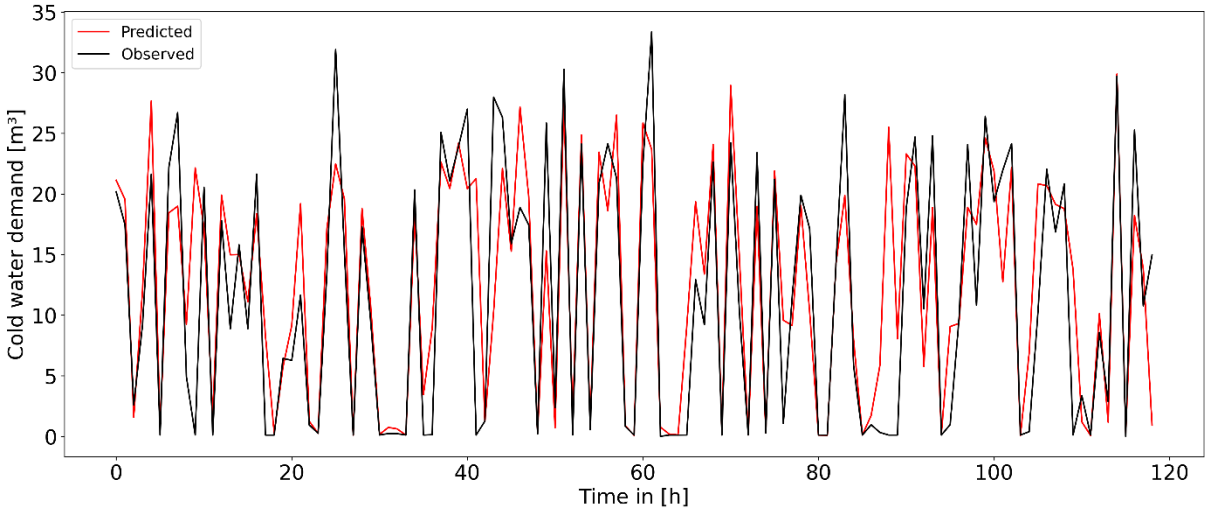

Figure 84 – Topusko prediction of coldwater demand; split into random training and test subsets

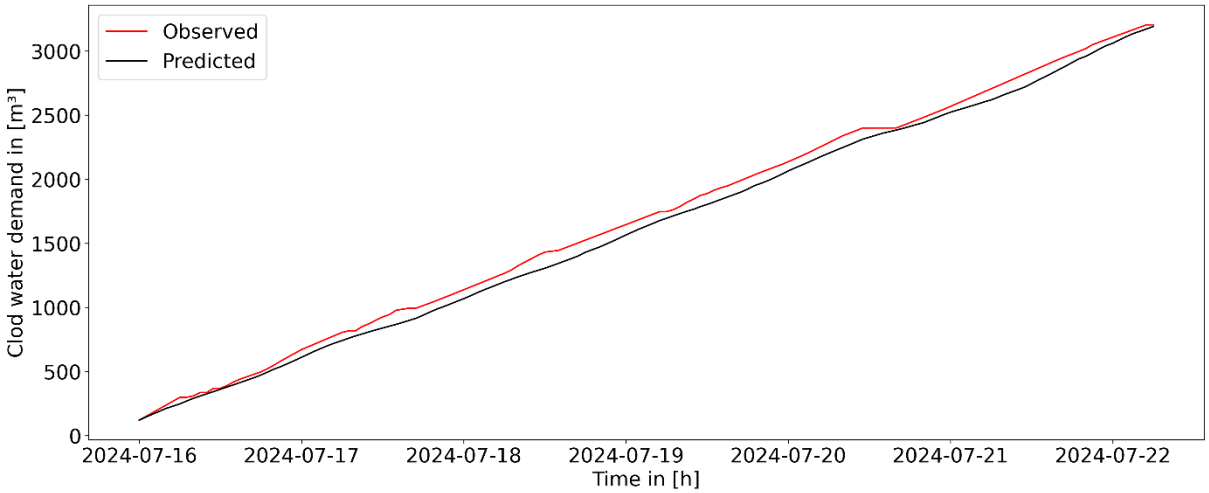

Figure 85 – Topusko prediction of coldwater demand; cumulated demand

Warmwater demand

Figure 86 shows the prediction of warmwater demand, split into random train and test subsets. Figure 87 shows the corresponding prediction of cumulated warmwater demand based on the previously trained model. With training and test data obtained in the period from 2024-01-01 till 2024-05-14 results with test data of R2 = 0.83 and a NRMSE (MinMax) = 0.12 m^3 (90 m^3 warmwater storage) could be archived. Again, the forecast quality was deemed sufficient to be implemented in the control system.

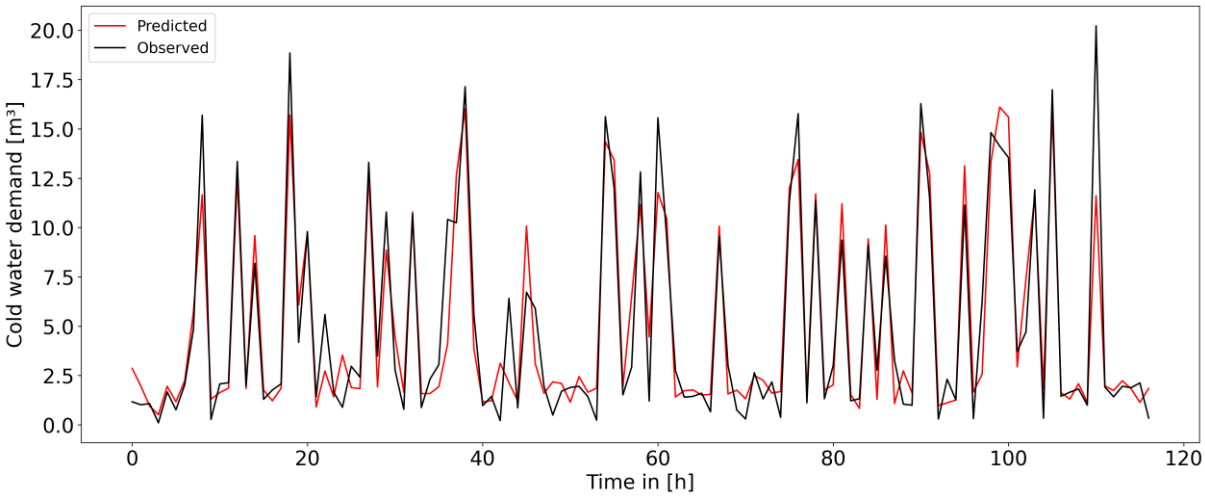

Figure 86 – Topusko prediction of warmwater demand; split into random training and test subsets

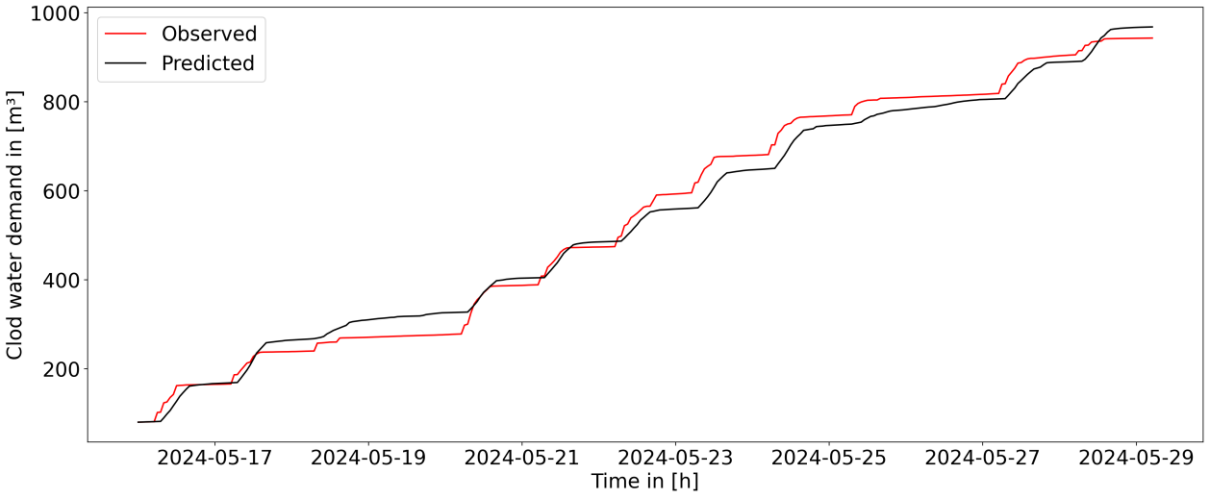

Figure 87 – Topusko prediction of warmwater demand; cumulated demand

The objective of the ML-based control system is to improve the efficiency of cooling process and minimize energy consumption by leveraging real-time environmental data and the operational needs of the system. For the geothermal water supply three different strategies were tested various environmental conditions:

- Night-time cooling through natural convection. This implies loading the thermal storage tank completely in the evening. This approach aimed to capitalize on the natural cooling process that occurs at night due to lower ambient temperatures. The hypothesis is that cooling the stored water through passive night-time convection will reduce the need for active cooling the next day. The advantages thereby are that it reduces the reliance on energy-intensive cooling systems during peak hours and that it uses naturally cooler night-time temperatures to dissipate heat, improving energy efficiency. The challenges are that ambient temperature at night might still be relatively high during summer months, limiting the effectiveness of natural cooling and that this strategy could result in inefficiency if storage capacity is predicted inaccurately for the following day.

- Loading the storage during the coldest hours of the night, leveraging lower ambient temperatures. The second strategy focusses on loading the tank during the period when outdoor temperatures are at their lowest, typically in the early morning hours. This technique attempts to maximize the active cooling effect by exploiting the coldest possible conditions to store cooler water for the upcoming day. The advantages thereby are a more effective cooling due to lower ambient temperatures. The challenges here are that it requires accurate prediction of when the coldest period will occur, which can vary depending on local weather patterns, especially when accurate weather data predictions are not available for the specific site that might also be affected by local weather phenomena and shading from hills, forests, etc.
- The third method involves cooling the water first before loading it into the tank. This approach uses cold geothermal water as an auxiliary cooling source, especially during warmer periods. The use of geothermal water reduces the energy required to operate the cooling tower. However, this method is unnecessary in the winter when ambient temperatures are naturally lower. The advantages thereby are that it reduces the overall demand on mechanical cooling systems, leading to energy savings. It is not applicable during colder seasons, necessitating a dynamic control that adapts to seasonal variations.

8.4 Assessment of the control performance

The implementation of a new Control and Monitoring System has significantly enhanced the efficiency of thermal energy management. Key components of this upgrade include the installation of controllable valves, new variable frequency drive pumps, and a state-of-the-art SCADA system with web access. In addition to these main components, numerous temperature probes, pressure sensors, and flow meters have been integrated into the system to facilitate improved operational control.

One of the most notable achievements of implementing advanced controls was the significant decrease in the utilisation of geothermal water. By optimizing system parameters and enhancing control mechanisms, the overall demand for geothermal water was reduced, thereby contributing to the sustainability of geothermal resources.

The integration of controllable valves, variable frequency drives (VFDs), and a robust SCADA system has led to an approximate 25% improvement in overall system efficiency. This increase not only demonstrates the effectiveness of the upgrades made to the thermal energy system but also underscores the importance of utilizing real-time data to fine-tune operations.

Through enhanced monitoring and control capabilities, the project successfully achieved a reduction in annual thermal energy consumption by approximately 27%. This accomplishment reflects the effectiveness of the adjustments made to the system's operational parameters, allowing for a more efficient utilisation of resources.

The implementation of energy-efficient technologies has resulted in a substantial reduction in electricity consumption, alongside a corresponding decrease in CO2 emissions, estimated at around 33% annually. This reduction highlights the positive environmental impact of integrating advanced controls and energy-efficient systems.

The achievements from operating advanced control systems underscore the potential for significant enhancements and better sustainability. These outcomes not only provide a roadmap for future projects but also emphasize the critical importance of continuous dialogue and collaboration between technicians, project engineers and scientists. Engaging stakeholders more

thoroughly during the design and implementation phases could facilitate smoother operations and further optimize system performance.

8.5 Lessons learned

The main learnings out of operating the data mining tool and the advanced controls embedded are:

Implementation and operation

- Initial assumptions during commissioning indicated that substations were adequately regulated for consistent return temperature. But it was discovered that substations were not adequately regulated, leading to the need for additional management mechanisms within the Centralized Thermal System (CTS).
- The requirement for more valves than originally planned was a direct result of this oversight. The experience emphasizes the necessity of thorough discussions and collaboration with onsite technicians, rather than solely relying on engineering designs. Improved communication during the planning phase could have helped mitigate the need for additional installations.
- The advancements made at the Topusko pilot site demonstrate the positive impact of modern intelligent control systems and enhanced monitoring on thermal energy management efficiency.
- Implemented changes have led to reduced energy consumption and emissions, contributing to a more sustainable operational model for future energy systems.

ML-based Control

- Decision tree regression was effective due to its interpretability, but careful attention must be paid to its limitations.
- Proper data preparation, including ensuring data quality and choosing the appropriate dataset length, is vital for building reliable regression models, especially for time-series forecasting. Proper preprocessing, including dealing with missing values and categorical variables, is essential to prevent biases in model predictions.
- Automated procedures for error correction and continuous monitoring of sensor data are necessary to maintain data integrity. Alerts and notifications are crucial for real-time data issue handling. Outliers can significantly distort results, particularly in decision tree models, leading to overfitting. Identifying and addressing outliers through techniques like transformation or trimming is necessary for robust predictions.
- Using a dataset that is too short or too long can negatively impact model accuracy. Balancing historical data with relevant, recent trends is key for effective forecasting. In dynamic environments, models must be retrained regularly to adapt to evolving patterns in the data, ensuring continuous accuracy over time. Regular cross-validation during retraining cycles helps avoid overfitting, ensuring that the model generalizes well to new data.
- Capturing seasonality and long-term trends in the data is crucial for accurate predictions, especially when dealing with time series forecasting.

9 Low-Temperature Substations (EURAC)

9.1 Description of the substation

In neutral-temperature district heating networks (NTDHNs), substations integrating heat pumps are essential components for the distribution of thermal energy, as they adjust the temperature of the energy distributed in the network to meet the specific heating or cooling needs of each building, acting as a prosumer.

Adopting decentralized, active substations in DHC networks offers multiple benefits. First, a more efficient distribution of thermal energy is achieved, reducing energy losses during transmission. Additionally, these substations provide greater control flexibility to individual prosumers.

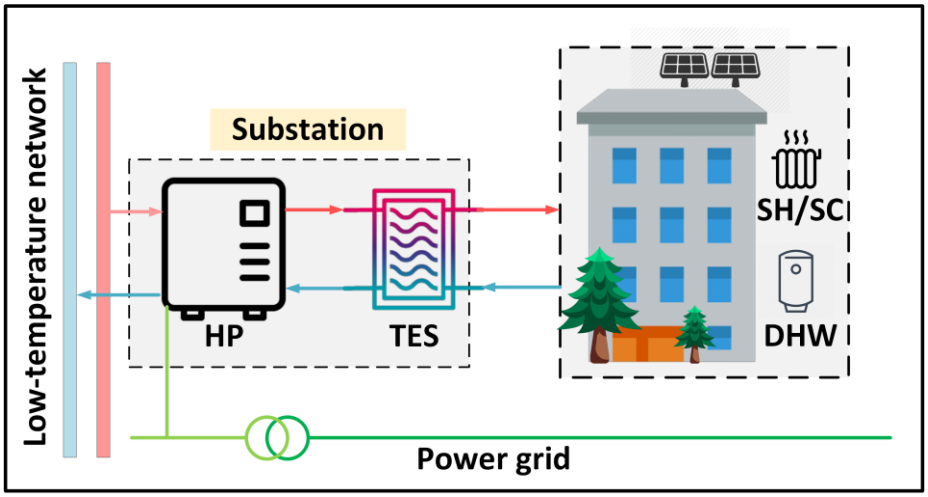

Figure 88. Overview of the substation's operation within the context of NTDHNs.

Figure 88 shows how the neutral-temperature network, operating at temperatures unsuitable for direct heating, integrates with a residential building. At the centre of the diagram, the substation is highlighted, equipped with a HP and TES. This substation supplies the necessary thermal energy for space heating and domestic hot water to the residential building shown on the right. This figure represents a straightforward configuration commonly used in substations.

In REWARDHeat, DANFOSS and EURAC developed a prefabricated substation for residential and tertiary applications, in which the HP is hydraulically connected to the DH and the building via pipes and valves in a way to enable optimal utilisation of the HP, hence the system's COP, for given DH supply and building load temperatures.

Figure 89 illustrates the different operational scenarios. Four different operation schemes can be made possible to enhance the effectiveness and efficiency of substation, tailored to specific temperature conditions:

- If the DH supply temperature is lower than the return from the building, scheme 1 is activated.
- If the DH supply temperature is higher than the supply set temperature to the building, scheme 4 is triggered.
- If the DH supply temperature is higher than the return from the building and lower than the supply set temperature to the building, scheme 2 or 3 are enabled; scheme 2 maximises system COP, while scheme 3 minimises return temperature to the DH, while still enabling high system COP.

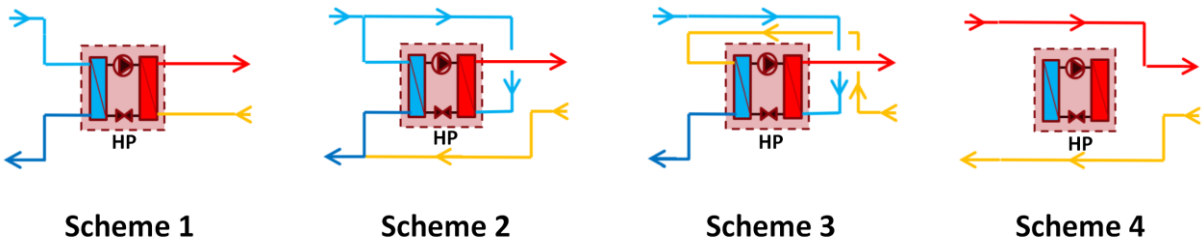

Figure 89. Overview of the four different configuration schemes considered in this study. The DH network is on the left and building on the right side of each scheme

The primary focus of this study is the optimal control of substation within a low-temperature network, specifically using a high-level controller based on MPC. Special attention is given to how MPC can manage the operation of HP + TES systems to meet the varying energy demands of buildings in a flexible and cost-effective manner. Figure 90 illustrates the concept of MPC approach: a rule-based control takes care of the substation operation, while the MPC overrides the set temperatures triggering TES charging, hence forcing or delaying the activation of the HP. To this end, a model of the substation is used to simulate the operation of the system in quasi- real-time over a 24-hour receding horizon, and an optimisation algorithm calculates optimal set point settings over the same timeframe (see Figure 90).

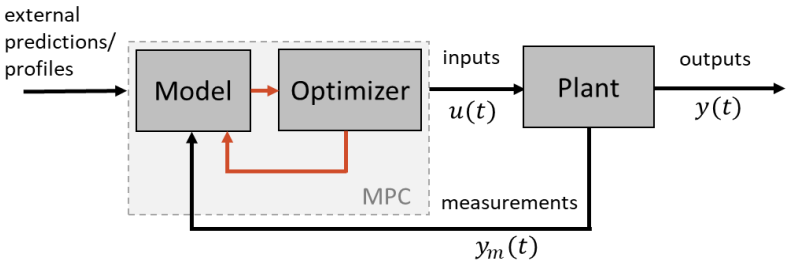

Figure 90. Schematic of MPC framework

9.2 Control problem formulation

The implementation of the MPC requires the development of several models. First, a model calculating the building's thermal demand based on outdoor conditions and user behaviour is essential, as MPC relies on accurate forecasts to anticipate future energy needs. Additionally, the MPC must account for the TES state of charge (SOC) and the operational status of the HP, as their interaction significantly impacts the overall system performance. Reduced-order models (ROM) for both the HP and TES are necessary to represent their behaviour efficiently while maintaining computational feasibility. Lastly, the optimisation process plays a key role in defining the best operational trajectory over a specific timeframe. To address these requirements, this study encompassed a range of activities:

- Building Demand Model: We built a ROM that not only captures the core dynamics of thermal load variations with computational simplicity but also remains robust and adaptable across various building types, sizes, and locations. The model can be trained during operation with limited monitoring data and is suitable for online execution, even on hardware with constrained computational resources.
- TES Model: We developed a mathematical model that achieves three key objectives: Simplicity, allowing it to capture the essential dynamics of TES behaviour while minimising computational

demands; maintaining a consistent energy balance that accounts for all relevant energy inputs and outputs during charging and discharging processes; ensuring an accurate representation of the thermocline, which influences performance and the operation of the connected HP.

- Heat Pump Model: The development of a ROM for heat pumps involves several key aspects to ensure accurate and efficient operation. The primary goal is to model compressor speed modulation to consistently track the setpoint temperature at the inlet of the load side. The model computes the thermal power and electrical consumption as functions of the compressor speed and other operational parameters. This requires a comprehensive understanding of the heat pump's performance characteristics and the ability to simulate these under various conditions. To ensure the model remains accurate and adaptive, continuous training procedures are integrated. These procedures update the performance map based on real-time data, allowing the model to adapt to different heat pump configurations and operational scenarios.
- Quasi-Real-Time Optimisation: We focused on the integrated operation of the substation, considering the dynamic interactions between the electricity grid, thermal grid, and buildings. This holistic approach ensures that all components work together seamlessly to achieve optimal performance.

9.3 Description of the advanced control implemented

9.3.1 Building thermal loads forecast

An Artificial Neural Network (ANN) was employed to develop a ROM capable of accurately predicting thermal loads. This data-driven approach effectively captures complex, multidimensional relationships between input variables. Unlike models dependent on specific building physics, the ANN is adaptable to diverse building types and sizes, requiring minimal monitoring data. Continuous training on new data enables the model to improve over time, adapting to changes in building characteristics and user behaviour.

The ANN architecture includes an input layer for relevant conditions, hidden layers for capturing intricate patterns, and an output layer for predicting thermal loads and indoor temperatures. Validated against both simulated and monitored data, this approach demonstrates superior performance.

Figure 91 illustrates the model's accuracy in predicting space heating load and indoor temperature for a multifamily apartment. The model aligns with the receding horizon and prediction horizon of the MPC. As shown, every hour (receding horizon), the model forecasts the next 24 hours (prediction horizon), ensuring continuous, up-to-date predictions that enhance MPC decision-making.

The zone-temperature ANN model and thermal load model's accuracy are evaluated through RMSE and MBE across different prediction intervals, including hourly, 3-hour, 6-hour, 12-hour, and 24-hour horizons. For zone temperature, the RMSE averages approximately 0.3 °C, while the mean MBE is 0.18 °C, on a 24-hour time horizon; for thermal load model, the average RMSE is 12 W/m^2, and the MBE is -2.4 W/m^2.

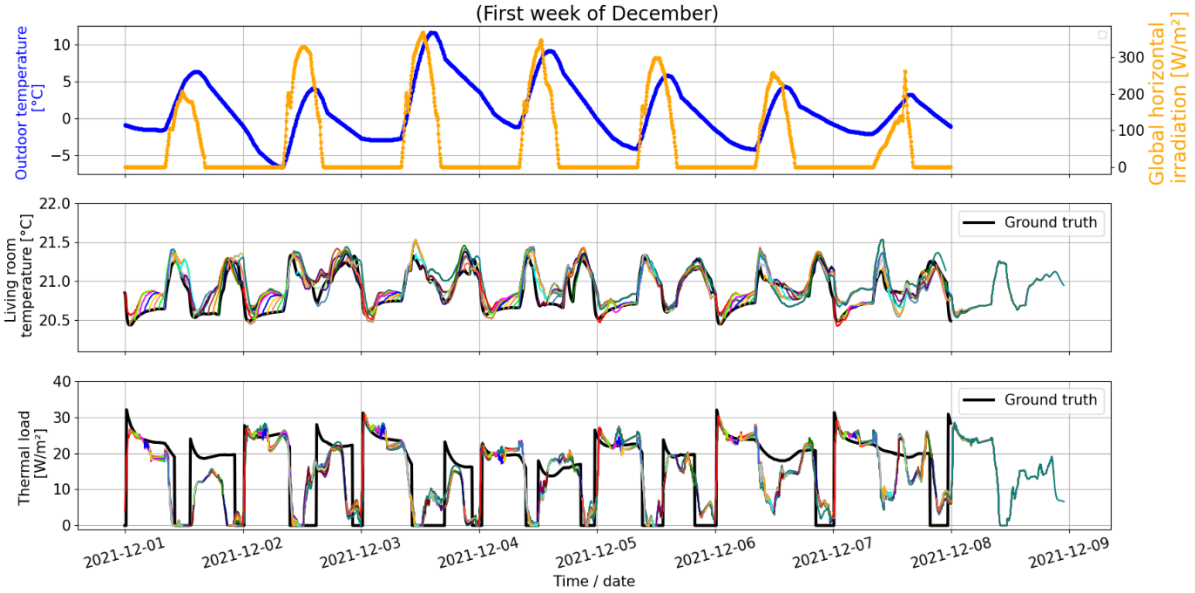

Figure 91. Living room's temperature and related thermal load prediction

9.3.2 Thermal Energy Storage modelling

The challenge with modelling a TES is to find a compromise between model accuracy and computational time required. TES models based on CFD simulations are widely available, but their execution time is usually prohibitively slow to use them in rigorous optimisation. TES models based on a simple energy balance, but they lack the accuracy in modelling the thermocline behaviour, making them unsuitable in a context where accurate temperature forecasts are required.

To achieve an optimal balance between runtime and model accuracy, a piston-based model was developed. This model greatly improved the ability to capture the formation, evolution, and collapse of thermoclines throughout the height of the TES. This model was designed to simulate key TES processes, focusing on energy transfer behaviour during charging and discharging phases. By concentrating on primary mechanisms, it reduces computational resource requirements. A critical feature of this model is its consistent energy balance, ensuring accurate tracking of all energy inputs and outputs. Additionally, it provides a detailed temperature gradient profile within the TES, which influences HP performance via the return temperature at its condenser.

To validate and refine the model, laboratory data from various charging and discharging cycles under different conditions was used, enabling fine-tuning to reflect real-life operation coupled with HPs. The model's outputs were compared against monitored data using key performance indicators, assessing uncertainty and consistency across various operating conditions. Figure 92 shows the predicted and observed tank temperatures during a preliminary dynamic test. This pattern was chosen, because it represents the standard dynamics when heating a TES with fixed ΔT while both mass flow and return temperature from the TES vary: when heating begins, a thermocline slowly starts to form as the heated water stays in the top of the tank. Once the thermocline is fully formed (in this case after around 50 minutes), the thermocline starts moving downwards, as the hot water is added at the top. Once the bottom of the thermocline starts touching the bottom port of the TES (around 80 minutes in this case), the outlet temperature starts rising. As this warmer temperature is fed back to the HP, the temperature of the inlet water to the TEs also rises, and a second thermocline starts to form. When the discharging starts (after around

100 minutes), all three temperatures start to drop, though this does not happen uniformly due to the development of the thermocline.

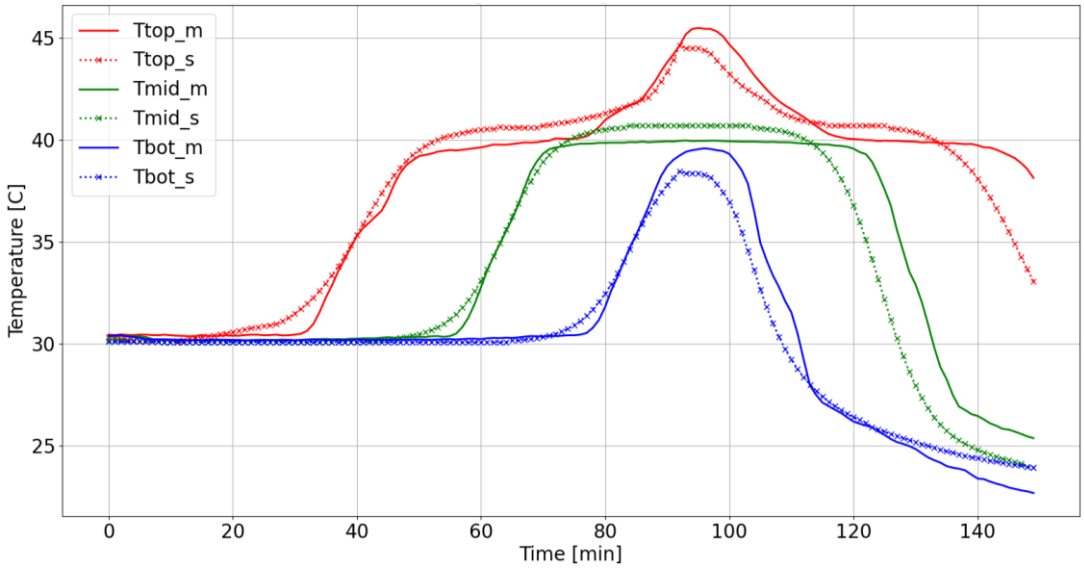

Figure 92. Temperature distribution along the TES height, monitored vs. ROM simulated

of a 1.2 m^3 TES tested in the laboratory under dynamic conditions over several days. The TES was equipped with four temperature sensors, located at the heights of the inlet and outlet ports (top and bottom, respectively) and at two intermediate levels (mtop and mbot). In addition, the inlet and outlet flows were monitored using dedicated heat meters.

A qualitative assessment (Figure 93) indicates that our simplified model aligns well with the laboratory benchmark, although some discrepancies remain, particularly regarding the height and shape of the thermocline. For instance, at 215,000 seconds, the actual data shows a single uniform thermocline passing through the middle-top sensor (mtop, shown in blue), while our model represents the thermocline in two slightly distinct steps.

Quantitatively, the model's accuracy at specific sensor positions is illustrated in Figures 94 and 95, where temperature predictions for each time step are plotted against measured values. Overall, the model demonstrates strong agreement with the ground truth, though minor time lags between predicted and measured series introduce some scatter around the bisector.

The R^2 values of 0.925 and 0.941 for the top and mtop sensors, respectively, indicate that the simplified TES model reasonably captures real-life behaviour. Specifically, it accurately predicts the timing and magnitude of the thermocline as it passes through the sensors governing heat pump hysteresis.

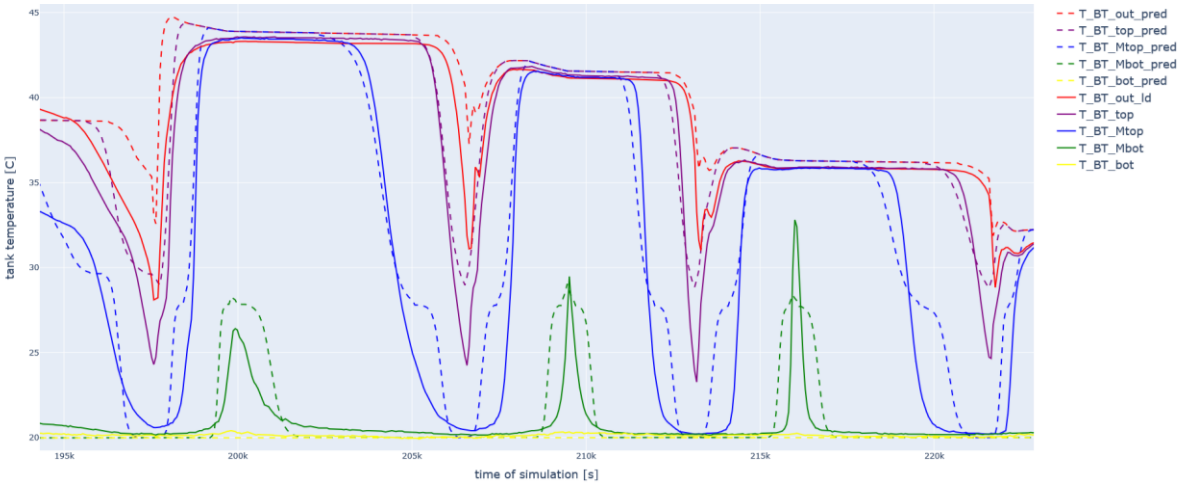

Figure 93: Predicted TES temperatures with dashed lines vs. measured ground truth.

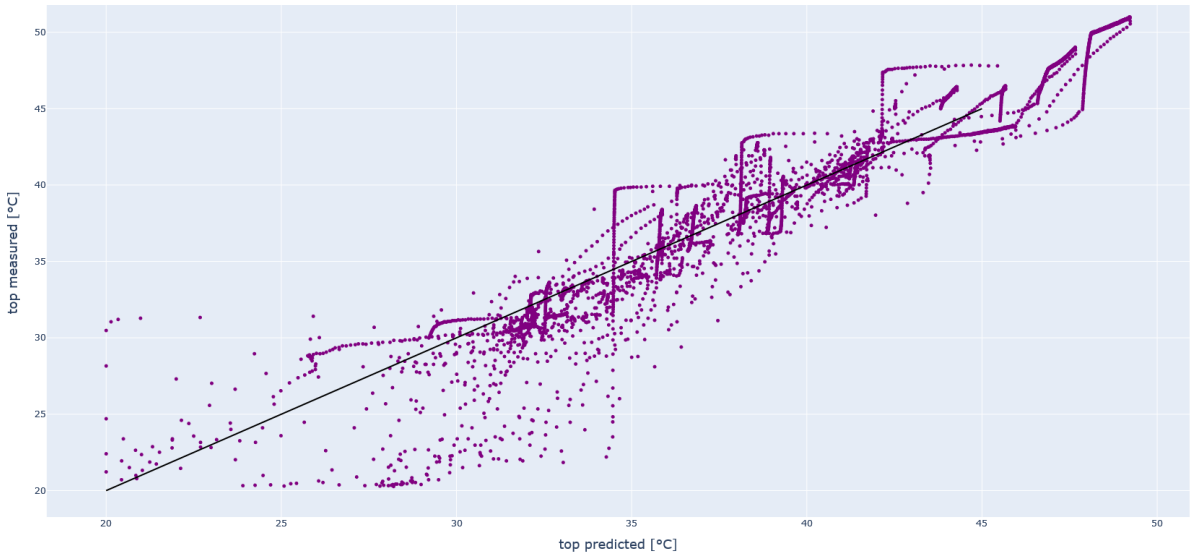

Figure 94: Scatter plot comparing the monitored TES top temperature sensors reading vs. ROM

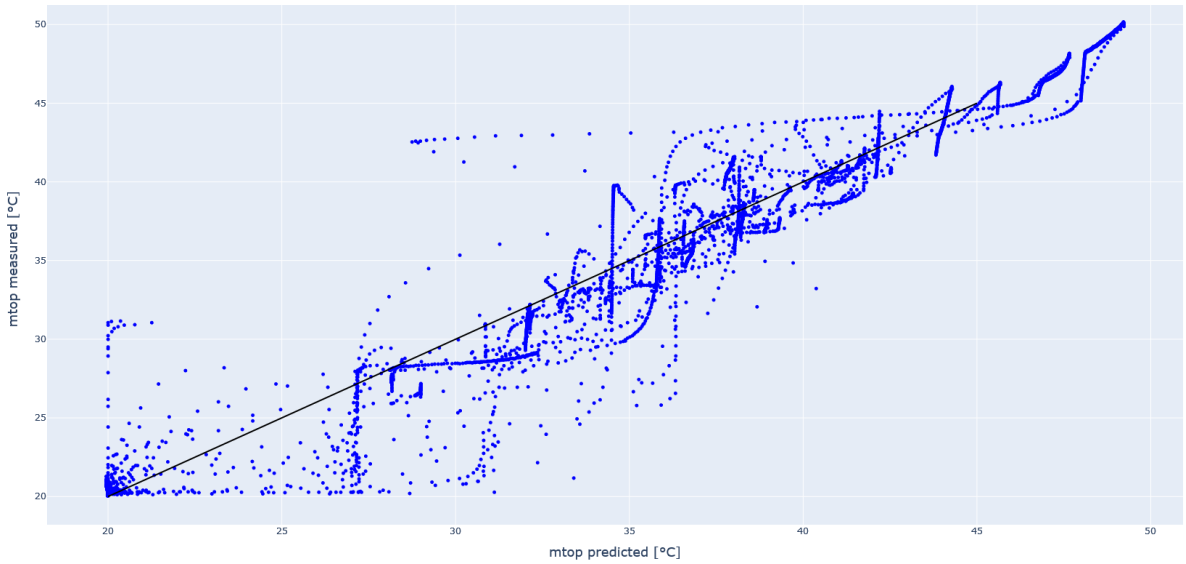

Figure 95: Scatter plot comparing the monitored TES mid-top temperature sensors reading vs. ROM

9.3.3 Heat pump modelling

The model effectively captures HP behaviour by combining discrete and continuous states. A key element of the ROM is the performance map, constructed as a multidimensional table with inputs such as compressor speed, inlet temperatures at both source and load sides, and load-side mass flow rate, while providing thermal power and electrical consumption as outputs. This performance map acts as a reference for predicting HP performance under various conditions.

Three distinct algorithms operate atop the performance map to modulate compressor speed, each tailored to specific operational goals. The first algorithm uses a standard state machine with ramps, gradually adjusting compressor speed to achieve the setpoint temperature. The second algorithm employs Proportional-Integral (PI) modulation, leveraging control theory for precise responses to temperature variations. The third algorithm focuses directly on the thermal load instead of the setpoint temperature, modulating compressor speed to match the heat demand.

To ensure ongoing accuracy and adaptability, a continuous training process updates the performance map and regression functions with new data, allowing the model to learn from real-time operations and refine predictions. When specific training data is unavailable, the model relies on default performance maps, scaled to the HP's nominal thermal power. Although these default maps may lack the precision of custom maps, they provide a reliable operational baseline, ensuring system functionality and reliability even without detailed input data. Figure 96 compares the performance of proposed HP model and monitored data.

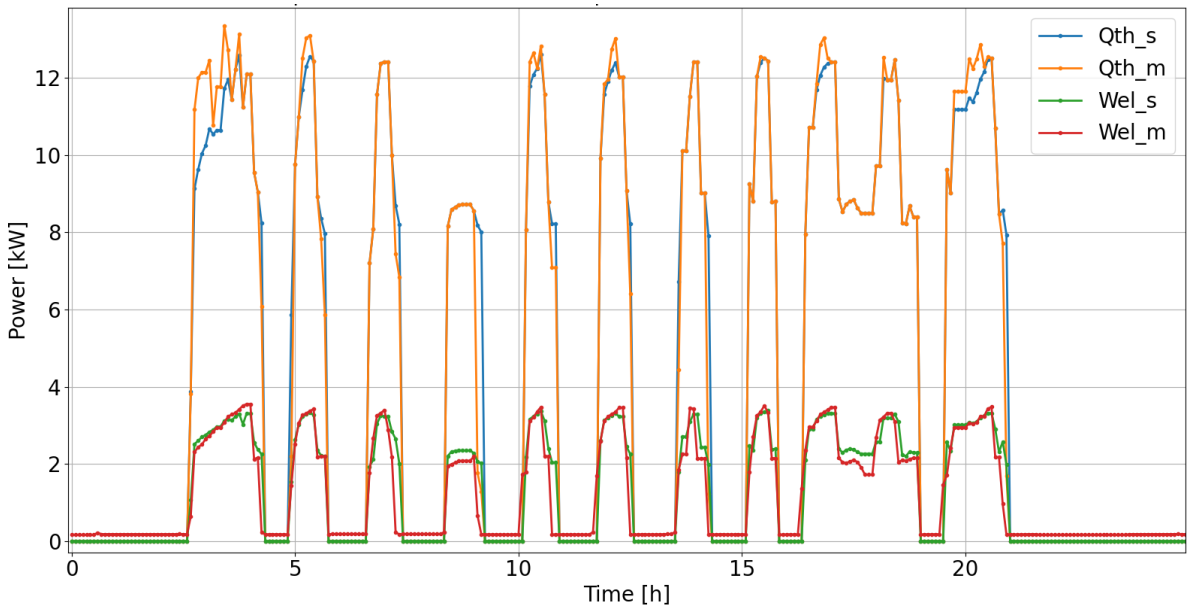

Figure 96. Thermal power and electrical consumption of HP, monitored (m) vs. ROM simulated (s)

The HP model's accuracy over 5-day simulations was assessed using RMSE and MBE for both electricity consumption (Wel) and output power (Qth). For Qth, the average RMSE was 0.05 kWh with an MBE of -2.3 Wh, while Wel had an RMSE of 0.03 kWh and MBE of -3.1 Wh.

9.3.4 Optimisation algorithm

Dynamic programming (DP) has been implemented as the optimisation method in the development of MPC for this project. This approach has been chosen for its potential in smart control applications, where the ability to handle complex, non-linear models is essential. The forward implementation of DP allows for efficient computation of system dynamics over time,

while its capability to conduct an exhaustive search of the solution space guarantees that the globally optimal control strategy is identified. This ensures that even in the presence of complex system behaviours, such as varying HP performance and thermal energy storage dynamics, the MPC framework operates with maximum efficiency and accuracy.

9.4 Assessment of the control performance

A simulation campaign has been carried out to test the MPC versus a Rule-Based Control (RBC), both used to steer the substation concerned under identical boundary conditions, when set up in a multifamily home.

A performance comparison has been conducted aimed to identify reliable and consistent operation of the MPC first, and to assess performance in terms of system COP and energy/costs savings overall.

The reference building, a renovated and energy-efficient multi-family house in Stuttgart, features a total area of 500 m^2 with a heating demand of 70 kWh/m^2 per year. The HP's thermal capacity is 35 kW at 80°C, whilst the TES has a volume of 2 m^3 and a height of 1.77 m. Weather conditions represent those of a typical December week. Along the work elaboration we considered several combinations of boundary conditions, including variable DH supply temperature and electricity price.

The MPC's optimisation algorithm has been allowed to change the TES set point, with a cost function aimed to minimise the overall operating (electricity) costs.

Test 1: Performance Assessment with Constant Electricity Price

In this test, the performance of MPC and RBC is compared with a constant electricity cost of 0.3€/kWh, while the DH temperature varies between 40°C and 55°C; the simulation is conducted over the first week of December.

The first plot in Figure 97 shows the thermal load, with the black line representing actual values and the colourful crosses indicating hourly predictions for the next 24 hours. The second plot displays setpoints: solid blue lines for RBC (55°C) and solid red lines for MPC. Dotted lines represent the user-side temperature delivered by RBC (blue) and MPC (red), showing that while the MPC sometimes violates the comfort condition, this effect is negligible. The third plot shows the stored energy in the tank, with the blue line for RBC and the red line for MPC, showing how the MPC decides to overcharge the TES when DH supply temperature is high and building load is low, e.g., at start of 2nd, 4th and 6th of December. The same happens along the day with less clear patterns, as the MPC also decides to delay the TES charge before low thermal load concretises. The fourth plot presents the cumulative electricity cost, with the blue line for RBC and the red line for MPC, demonstrating that MPC generally experiences lower energy costs, with a reduction quantified at 7.44%. Finally, the fifth plot depicts the DH temperature (black line) and the system COP, as scatter points. The colour of these points indicates the active control scheme: as can be qualitatively seen, the MPC operates the substation more likely with scheme 2 (higher COP), during periods of high DH supply temperature.

Figure 98 illustrates the activation time of the HP under different control schemes, enabling a quantitative comparison. The analysis evaluated both the total HP_{ON} time and the number of activation cycles.

Under RBC, the HP was activated for a total of 2,495 minutes, with scheme 1 being the predominant mode (1,753 minutes, 70.3% of the total HP_{ON} time). Scheme 2 was active for 742 minutes, making

up 29.7% of the total activation time. The RBC exhibited a relatively low number of activation cycles: 38 cycles observed during the test period.

In contrast, the MPC showed a slightly lower total HP_{ON} time of 2,341 minutes. Scheme 1 was used for 1,517 minutes, representing 64.8% of the total HP_{ON} time, while scheme 2 contributed 824 minutes, or 35.2% of the total. Notably, the MPC resulted in a higher frequency of activation cycles, as activation cycles have not been constrained in the optimiser's cost function.

Increasing Scheme 2 operation time results in this case in a COP rise from 7 to 7.5.

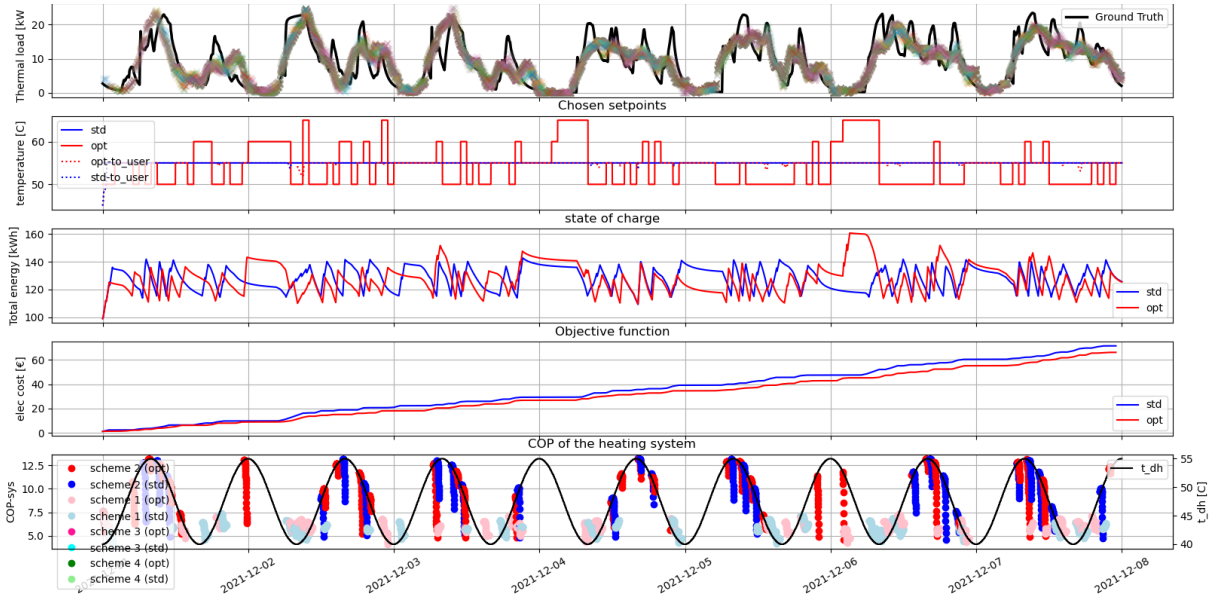

Figure 97. Result of comparison the performance of MPC vs. RBC with constant electricity price (test 1)

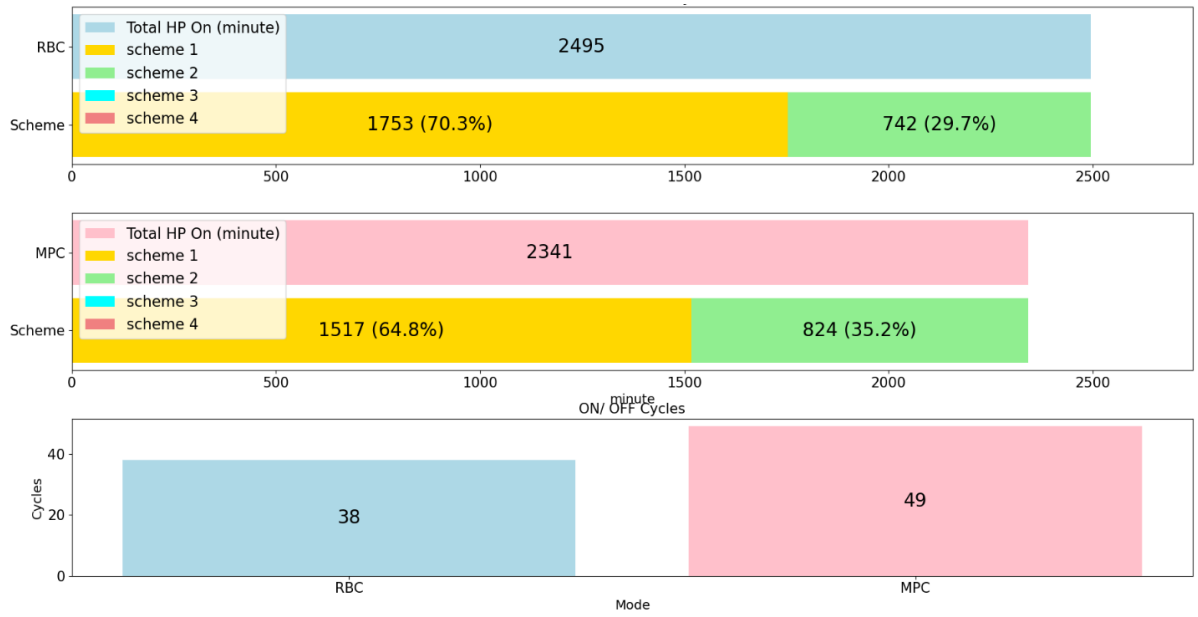

Figure 98. Activation time of the HP in test 1

Test 2: Performance assessment with time of use (TOU) electricity tariff

In Test 2, conditions are identical to those in Test 1, except that the electricity price from the grid follows a Time-of-Use (TOU) tariff: €0.15/kWh during off-peak hours (00:00 to 07:00) and €0.30/kWh during peak hours.

Figure 99 presents the results of Test 2, showing similar outcomes to the previous test, with the MPC reducing overall electricity costs by 6.7%. This is achieved by triggering substation operation under scheme 2 more frequently. In this scenario, the MPC selects even higher set temperatures during the night, when both high DH supply temperatures and low electricity prices coincide. The MPC strategy spending 40.4% of its operational time in scheme 2 results in a system COP increase to 7.5 also in this case.

Figure 100 depicts the activation time of the HP across different schemes. The MPC not only reduced the total HP_{ON} time but also increased its operation in scheme 2, contributing to improved system efficiency. Additionally, the HP's running time during the TOU period (00:00 to 07:00) increased by 9% in Test 2, denoting that the MPC tries to capitalise on TOU pricing incentives. This slight increase suggests that the MPC recognized the cost-saving opportunities presented by the TOU pricing, though the marginal change implies that the TOU scheme may not have been sufficiently attractive to significantly alter the MPC's overall strategy.

The MPC's responsiveness to dynamic pricing models demonstrates potential for further optimisation and efficiency improvements under more compelling pricing schemes.

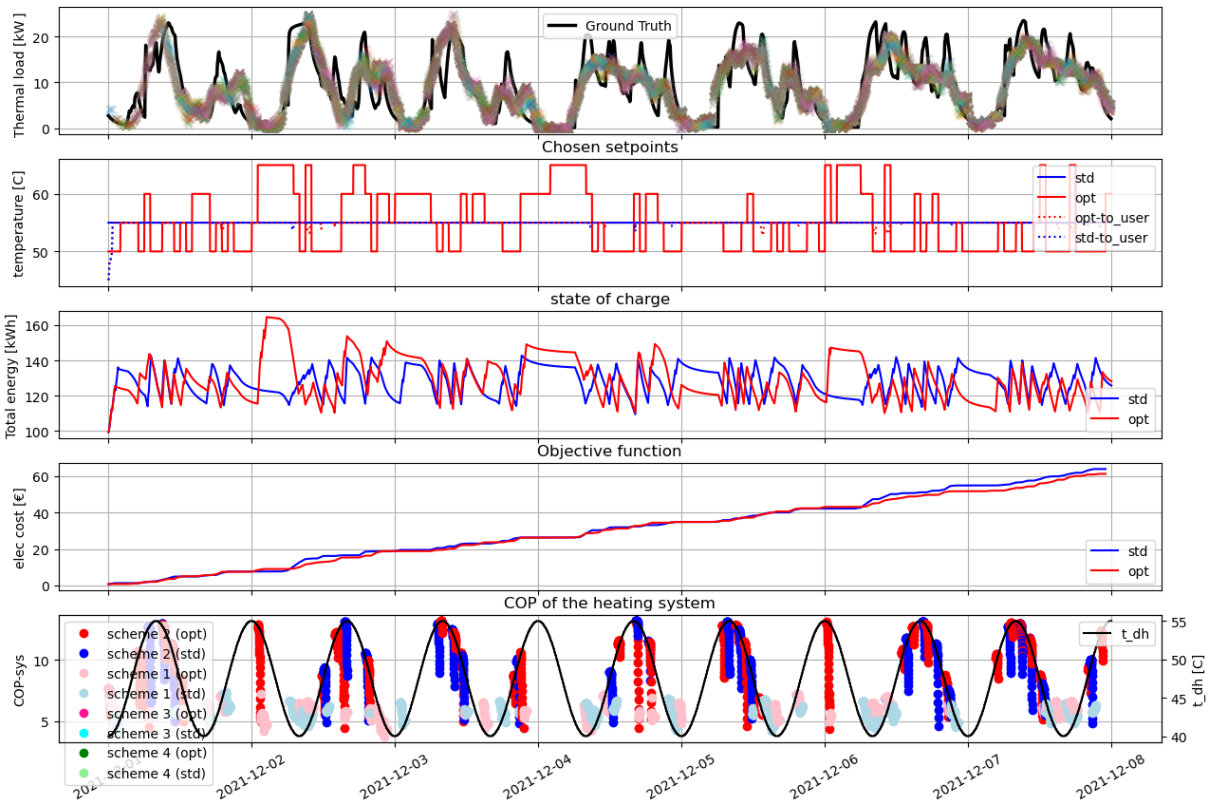

Figure 99. Result of comparison the performance of MPC vs. RBC with TOU (test 2)

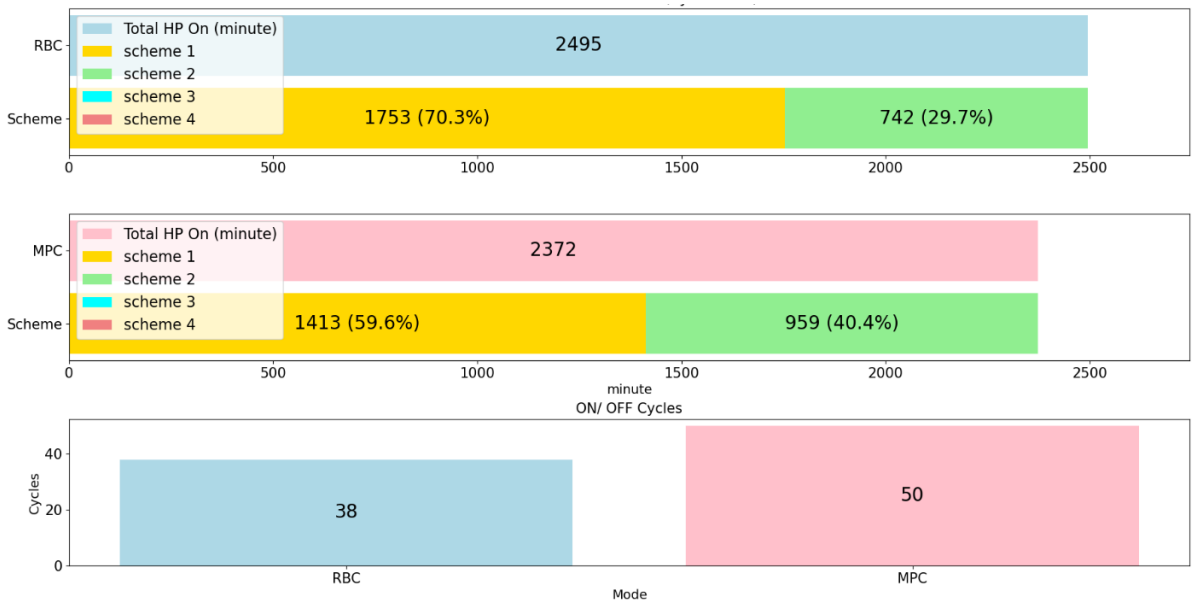

Figure 100. Activation time of the HP in test 2

Test 3: Performance assessment with dynamic electricity price

Test 3 assesses the performance of MPC and RBC under dynamic electricity pricing, building on the findings from the previous test. Figure 101 shows the identified hypothetical, variable electricity price profile, highlighting high values when DH supply temperatures are low.

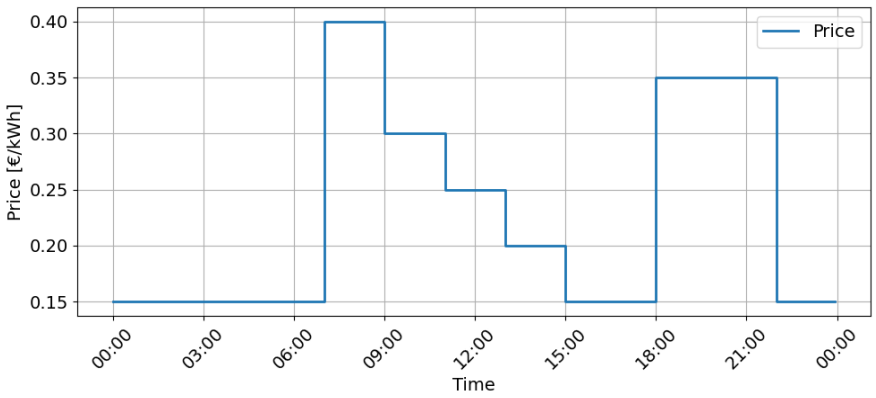

Figure 101. Electricity price in test 3

Figure 102 presents the test results, showing a 21.7% reduction in electricity costs achieved by the MPC compared to the RBC. The synchronisation of incentives enables the MPC to shift a larger share of HP use to cost-effective periods and operate in scheme 2 for 39% of the time, resulting in an overall system COP of 7.4.

Figure 103 displays the utilisation of different schemes, showing results overall similar to those of the previous test. The analysis shows that MPC not only reduced overall activation time by approximately 4.8% compared to RBC but also demonstrated higher adaptability to price variations.

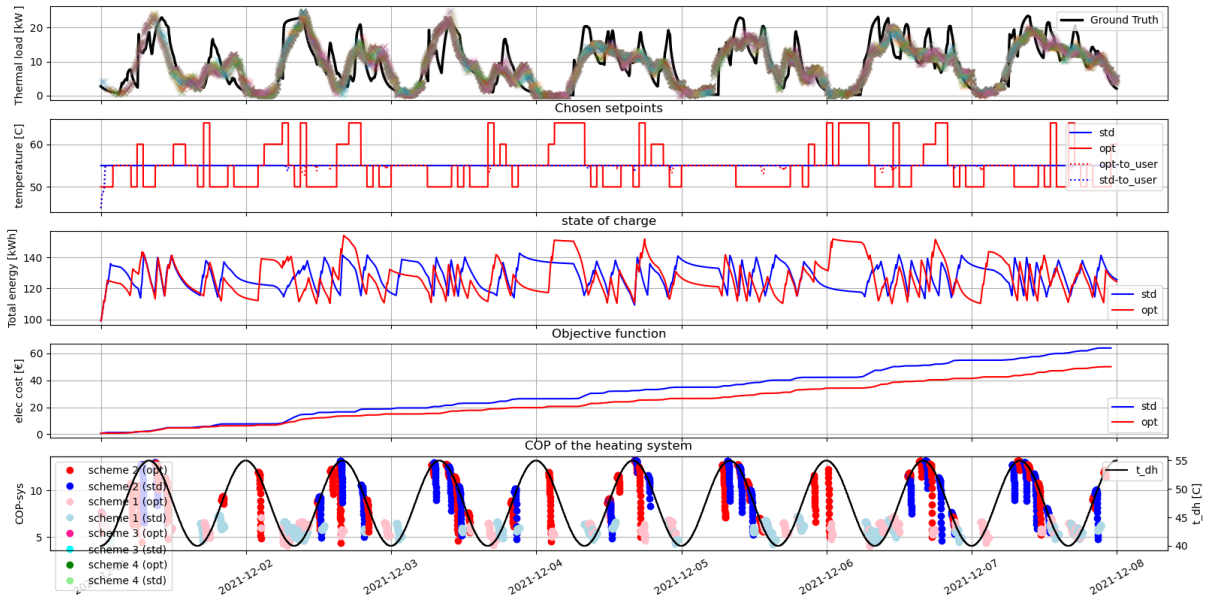

Figure 102. Result of comparison the performance of MPC vs. RBC with variable electricity price (test 3)

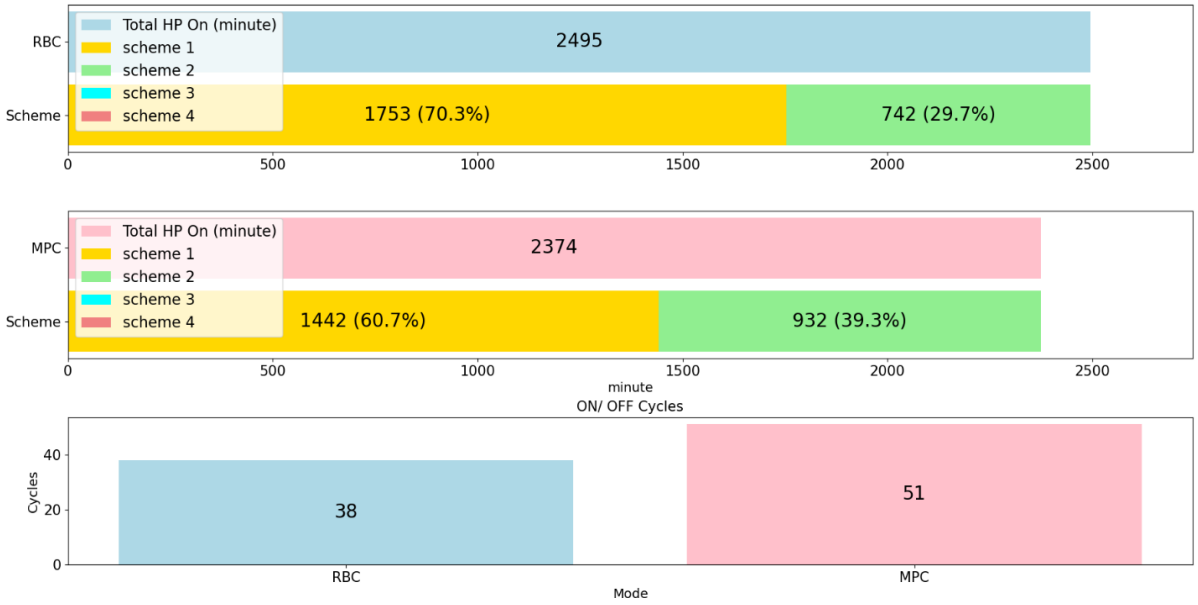

Figure 103. Activation time of the HP in test 3

Figure 104 illustrates the hourly distribution of HP operation over the simulation period. The black line indicates daily fluctuations in electricity prices, while the green line represents the average thermal load throughout the simulation week. The red and blue bars show the HP operation distribution by hour for MPC and RBC over the entire simulation.

The MPC strategy activates the HP more frequently during midnight hours and other low-cost periods, taking advantage of lower electricity prices. During these times, MPC extensively uses the HP to maximize TES charging, achieving its highest operation rates when electricity costs are lowest. In contrast, RBC shows the highest levels of HP operation during peak electricity price periods, such as between 07:00 and 08:00 in the morning.

MPC effectively manages HP operation during off-peak hours (00:00 to 06:00, 13:00 to 17:00, and 22:00 to 23:00), totalling 1,355 minutes, compared to RBC's 879 minutes -a 35% increase in off-peak operation time by MPC. Additionally, the RBC operation bars closely follow the load pattern, while MPC bars show less correlation. This difference is due to MPC's ability to predict future trends, including electricity prices and load peaks, allowing it to make more optimal decisions.

For example, at 17:00, MPC identifies an upcoming load peak and, recognizing the current low electricity price, opts to run the HP extensively. Conversely, at 18:00, when electricity prices are high, MPC strategically minimizes HP operation and discharges the TES instead.

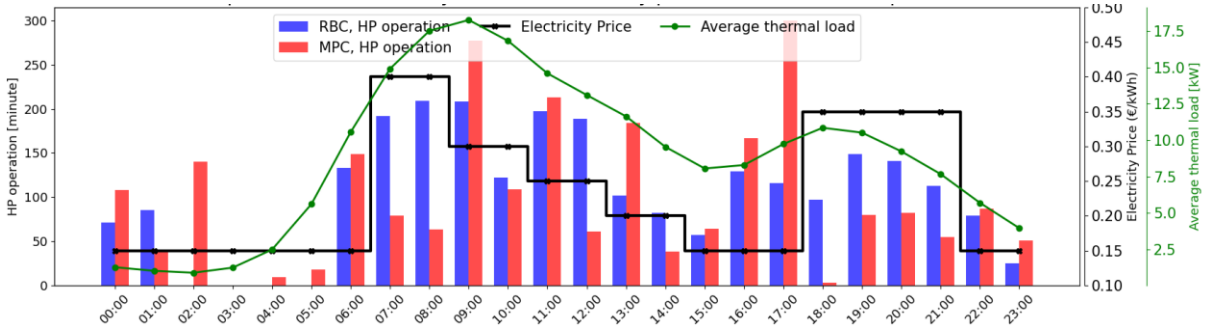

Figure 104. HP operation distribution by hour with electricity prices and average of thermal load (test 3)

Table 9-1 summarises the results from a costs perspective. In the first test, under the RBC strategy, the total electricity cost was 71.28 €, the thermal cost was 71.84 €. When MPC was applied, the total electricity cost dropped to 65.97 €, while the thermal cost slightly decreased to 71.41 €.

In the second test, with RBC, the total electricity cost was 65.64 €, the thermal cost was 71.84 €. Using MPC, the electricity cost was further reduced to 61.24 €, while the thermal cost increased slightly to 73.85 €. These results indicated a 6.7% reduction in electricity costs, a 2.7% increase in thermal costs. In the third test, RBC resulted in a total electricity cost of 63.88 €, with a thermal cost of 71.84 €. Under MPC, the electricity cost significantly decreased to 50.01 €, while the thermal cost increased to 73.13 €. This test demonstrated a substantial 21.7% reduction in electricity costs, and a 1.8% increase in thermal costs.

Table 9-1. Economic assessment of test 1, 2 and 3

	Cel,off-peak [€/kWh]	Cel,peak [€/kWh]	Cth,dh [€/kWh]	Cel,total [€]	Cth,total [€]	Total bill [€]
RBC	0.3	0.3	0.05	71.28	71.84	143.12
MPC-Test 1	0.3	0.3	0.05	65.97	71.41	137.38
Variation [%]				-7.44%	-0.5%	-4%
RBC	0.15	0.3	0.05	65.64	71.84	137.48
MPC-Test 2	0.15	0.3	0.05	61.24	73.85	135.09
Variation [%]				-6.7%	+2.7%	-1.73%
RBC	Variable		0.05	63.88	71.84	135.72
MPC-Test 3	Variable		0.05	50.01	73.13	123.14
Variation [%]				-21.7%	+1.76%	-9.26%

9.5 Lessons learned

Overall, MPC consistently outperformed RBC by reducing electricity costs and achieving notable savings on the total bill, despite slight increases in thermal costs in some scenarios.

The analysis indicates that MPC consistently resulted in lower total HP operation time, reflecting more optimized HP usage. However, the increase in activation cycles with MPC suggests a more dynamic response to changing conditions. To address this, it may be beneficial to include a term in the objective function aimed at reducing these cycles.

Although the amount of power that the MPC controller failed to cover was slightly higher, it remained minimal. Further investigation, including a penalty term in the objective function, could help reduce this uncovered power.

The main learnings out of operating the data mining tool and the advanced controls embedded are:

- Critical role of accurate thermal load forecasting: Precise prediction of thermal load demand is essential for effective MPC. Utilizing ANN to develop ROM provides accurate and adaptable forecasts, improving system responsiveness.
- Continuous adaptation through model training: Integrating continuous training procedures for models ensures they remain accurate over time, allowing the system to adapt to changing conditions and maintain optimal performance.
- Effective use of dynamic pricing and load shifting: MPC demonstrates strong capability in load shifting by scheduling HP operation during lower-cost periods and optimizing TES charging strategies. This not only weights dynamic electricity pricing to reduce operating costs but also effectively shifts the energy load to off-peak times, enhancing overall system efficiency and demand side management.
- Efficiency gains from multiple substation schemes: Considering different schemes for the substation managed to increase the efficiency of the system. Utilizing temperature-based transitions between these operational schemes and implementing hysteresis combinations improved substation efficiency by ensuring smooth and effective scheme changes.

While this study has shown significant benefits from using MPC in decentralized substations, there is room for further enhancement. Below are some suggestions for improving the system and directions for future research to build upon our findings:

- Incorporate equipment longevity constraints: While MPC increases efficiency, it also leads to more frequent HP switching rate, which may impact equipment wear in the long run.
- Expand dynamic pricing models: Further research could explore more complex and realistic dynamic pricing schemes, including real-time pricing and demand-response programs, to fully leverage MPC's capabilities in cost reduction and load shifting.
- Enhance model training techniques: Investigate advanced machine learning techniques for model training, such as reinforcement learning, to improve model adaptability and prediction accuracy under varying conditions.
- Optimize TES sizing and operation: Study the impact of different TES sizes and configurations on system performance, aiming to optimize TES design for specific applications and further improve energy storage utilisation.

- Integrate renewable energy sources: Explore the integration of renewable energy sources, such as solar thermal, PV panels or battery storages, into the MPC framework to enhance sustainability and reduce reliance on grid electricity.
- Extend the study to include cooling applications: This study focused mainly on heating applications. Future research could expand the scope to include cooling operations within the district heating network. Applying MPC to cooling systems can optimize performance and efficiency year-round, providing comprehensive solutions for both heating and cooling needs.

10 Key Takeaways

The REWARDHeat project revealed critical insights in several areas of DHC network optimisation, from monitoring and data management to advanced control and MPC. Here are the main key takeaways:

- Collaborative Development and Communication: Engaging DHC experts, operators, and local staff in the design and development phases was crucial for effective system integration and ensuring user-friendly software interfaces. This collaborative approach allowed for better system ergonomics and operational adaptability, making the software more effective in meeting practical requirements.
- Monitoring Data and System Integration: Reliable, real-time data is fundamental for effective DHC management. The project highlighted challenges related to data gaps and network inconsistencies, with SCADA systems sometimes failing to provide high-quality or continuous data. Advanced monitoring systems, with fibre-optic communication as the preferred infrastructure, improve performance but still require mechanisms for gap-filling and data integrity. Additionally, integrating diverse data sources, such as SCADA, and BMS at customer side, provides operators with enhanced system visibility and operational flexibility, although harmonizing data across varied sources remains complex.
- Layered Control Approach: Effective control systems require a bottom-up approach, starting with hardware-level control and building more complex control logic on top. This structured approach allows for robust system performance and aligns with varying automation needs across substations.
- Predictive Model Adjustments to Context: Balancing model complexity with forecast accuracy is essential. To improve forecasting, models must incorporate local weather correction to address recurrent deviations.
- Limitations of Non-Linear Solvers: Although non-linear solvers expand optimisation possibilities, they're still evolving. Trade-offs are required between the depth of physical modelling and the feasibility of solver convergence, highlighting the need for selective application in complex networks.
- Machine Learning Challenges for Forecasting: Modelling interdependencies between heating, cooling, and electricity in systems like heat pumps can be complex. For new substations with limited data, simpler forecasting methods adaptable to real-time changes are beneficial during early commissioning.
- Data Management in Predictive Models: Model accuracy is affected by data quality. A balance of historical and recent data, coupled with regular model retraining and cross-validation, ensures models stay adaptable to evolving conditions and prevent overfitting. Continuous training of models is essential for accuracy. Advanced machine learning techniques, like reinforcement learning, should be further explored to enhance adaptability and improve predictions under changing conditions.