FTI Initiative Energy Model Region

Publishable final report

Programme control:
Climate and Energy Fund

Programme management:
The Austrian Research Promotion Agency (FFG)

Final Report
created on
31.10.2025

Project title:

DOPPLER
Digital **OP**timisation **P**latform for DH systems with supp**L**ier and **E**nd user **R**esponse

Project number:
894580

Call	4. Call FTI Initiative Energy Model Region
Project Starting Date	01.11.2022
Project Ending Date	31.10.2025
Total duration of project (in months)	36 Months
Project holder (Institution)	Güssing Energy Technologies GmbH
Contact person	DI(FH) Klaus Paar
Postal address	7540 Güssing, Wiener Straße 49
Telephone	+43 3322 42606 322
Fax	Klicken oder tippen Sie hier, um Text einzugeben.
E-mail	k.paar@get.ac.at
Website	https://www.get.ac.at/doppler/

DOPPLER

Digital **OP**timisation **P**latform for DH systems with supp**L**ier and **E**nd user **R**esponse

Authors:
DI (FH) Klaus Paar, Güssing Energy Technologies GmbH
DI(FH) DI Christian Doczekal, Güssing Energy Technologies GmbH
Harald Trinkl, BSc., Güssing Energy Technologies GmbH
Konstantin Plank, BSc., Güssing Energy Technologies GmbH
Lucas Scheiber, Scheiber solutions GmbH
Wolfgang Trischnigg, Optimierma GmbH
Dr. Alexander Kaufmann, Industriewissenschaftliches Institut
DI Dr. Stefano Coss, Arteria Technologies GmbH
Ing. Otmar Schlager MSc, eKUT GmbH
DI Dr. Daniel Muschick, BEST – Bioenergy and Sustainable Technologies GmbH

1 Table of content

1 Table of content 4
2 Introduction 5
3 Content presentation 5
3.1 Digital Twin 5
3.1.1 Arteria Platform 5
3.1.2 GIS Platform 11
3.1.3 Connection and Data Interfaces 16
3.1.4 Connection to Smart Home System 21
3.1.5 Optimization and Demand Response 25
3.1.6 End-User Integration 39
3.2 Business Models 50
3.2.1 Core Concept: Common-Good Economy Foundation 50
3.2.2 End-User Perspectives and Flexibility Readiness 51
3.2.3 Integrated Platform Architecture and Technical Integration 51
3.2.4 Flexibility Mechanisms and User Engagement 52
3.2.5 Gamification and Community Engagement Strategy 52
3.2.6 Impact Visualization: Making Environmental and Financial Benefits Tangible 52
3.2.7 Stakeholder Value Creation and Revenue Model 53
3.2.8 Implementation Strategy and Risk Management 54
3.2.9 Validation and End-User Feedback 54
4 Results and conclusions 54
4.1 Integrated Digital Twin as a Technical Backbone 54
4.2 Demand-Response Optimization and Supply Temperature Control 55
4.3 Optimization of Heat Production via Predictive Supervisory Control 56
4.4 Data Interfaces and Cross-Site Integration Lessons 56
4.5 Customer-Side Integration with the meo PROPILOT Controller 57
4.6 Dashboards as a Shared Operational and Engagement Layer 57
4.7 End-User Motivation and Acceptance: Households vs. Manufacturing 58
4.8 Overall Conclusions: What DOPPLER Demonstrated 59
5 Outlook and recommendations 59
6 Bibliography 61
7 Appendix 61
8 Contact details 61

2 Introduction

The DOPPLER research project addresses one of the key challenges facing modern district heating (DH) systems: how to increase operational efficiency and flexibility in increasingly complex energy systems while maintaining user acceptance and economic viability. Against the backdrop of the energy transition, the project focuses on the integration of digital technologies, decentralized control, and user-centric approaches to unlock demand response and optimization potentials in existing DH networks. By combining technical innovation with socio-economic analysis, DOPPLER aims to bridge the gap between advanced system operation and real-world applicability in heterogeneous DH environments

The focal points of the project lie in three closely interlinked areas. First, the development of a **digital twin** for district heating systems enables real-time simulation, monitoring, and analysis of network behavior. Second, **decentralized optimization and control strategies**, including the integration of the MEO controller on the secondary side, are used to actively influence heat demand and system performance. Third, the project places strong emphasis on **end-user and stakeholder integration**, investigating behavioural aspects, acceptance, and motivation through surveys, workshops, dashboards, and gamification elements. Together, these focal points ensure that technical optimization is aligned with human and organizational factors.

To achieve these objectives, a diverse set of methods was applied. These include thermo-hydraulic and exergy-based modelling, real-time data integration via cloud platforms, predictive control algorithms, and field demonstrations in multiple DH networks. In parallel, socio-economic methods such as online surveys, stakeholder workshops, expert interviews, and usability evaluations were employed to capture end-user preferences and assess acceptance. This mixed-methods approach allows technical results to be continuously validated and refined through practical feedback.

The structure of this report reflects the project's work package logic. After outlining the methodological and conceptual foundations, the report presents the development of the digital twin and control strategies, followed by the characterization of use cases and end-users. Subsequent chapters document the demonstration results, business model considerations, and user integration outcomes. The report concludes with an outlook on scalability, long-term impacts, and recommendations for future implementation of digital optimization solutions in district heating systems.

3 Content presentation

3.1 Digital Twin

3.1.1 Arteria Platform

This part presents the technical work completed by Arteria Technologies within the framework of the DOPPLER project, focusing mostly on the contributions to Work Package 2 (Digital Twin Development)

and Work Package 4 (Demonstration of Optimization Potential). Over the duration of the project phase, Arteria implemented a complete set of digital models for the district heating networks of Mischendorf, Rohrbach and Güssing, integrated real-time data-streaming to the Arteria platform, developed and validated a demand-response driven optimization algorithm, and contributed to the operational integration between the digital twin and the MEO smart-home controllers. The overall goal was to enable decentralized optimisation of district heating systems through the combination of real-time measurements, high-precision simulation models, and advanced control strategies.

3.1.1.1 Digital Twin development

The initial stage of the work concentrated on building digital representations of the individual district heating networks. For each location

- Mischendorf,
- Rohrbach
- and Güssing,

the complete hydraulic network was imported and processed using Arteria's automated GIS-based parser. This resulted in operational digital twin environments capable of simulating energy flows, pipe temperatures, pressure conditions and end-consumer behaviour with high fidelity. The digital twins form the backbone of the entire DOPPLER concept, as they provide the required accuracy to test and validate control interventions before applying them to the real plants.

The digitalized representation of the Mischendorf network is shown in the following figure.

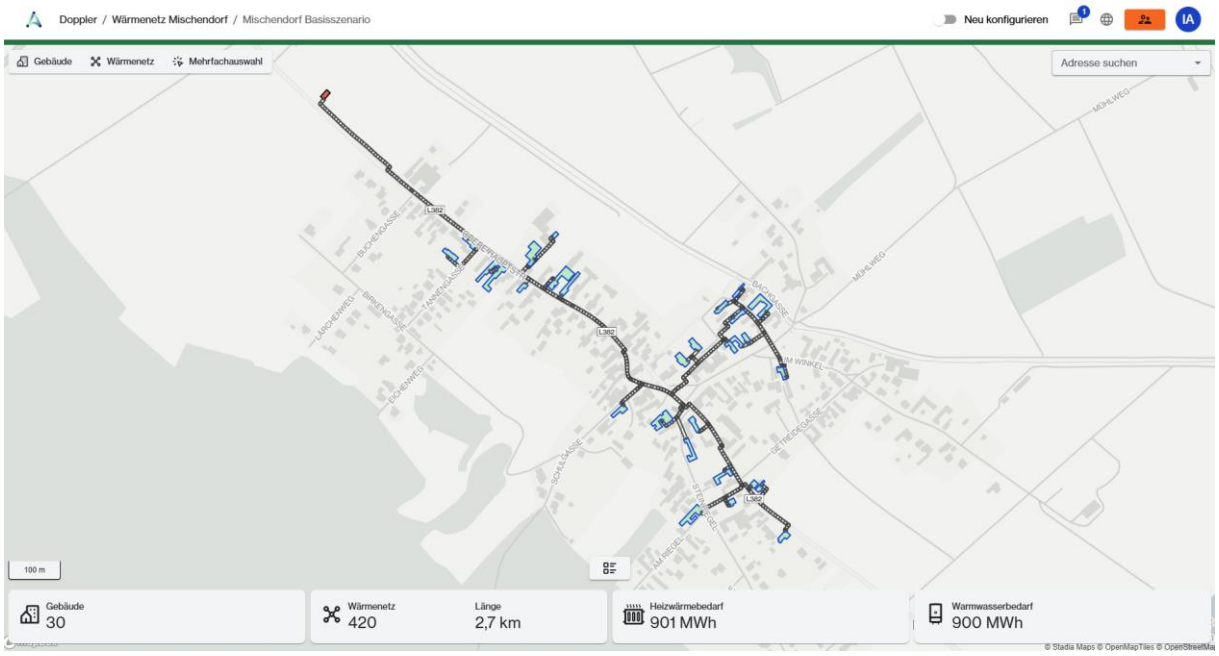

Figure 1: Mischendorf digital twin

The visualization presents the hydraulic layout, pipe connections, substations, and node distribution. Each node is georeferenced and linked to metadata and real-world attributes obtained from the network

operator. This data has been generated in collaboration with GET partner. The underlying model includes pipe lengths, diameters, insulation types, thermal losses, and dynamic behaviour. This level of detail is essential for thermal transient simulations and for estimating supply–return interactions.

The same methodology was applied to the Rohrbach district heating system. The following figure illustrates the Rohrbach network in its digitalized form, embedded in the Arteria GIS environment.

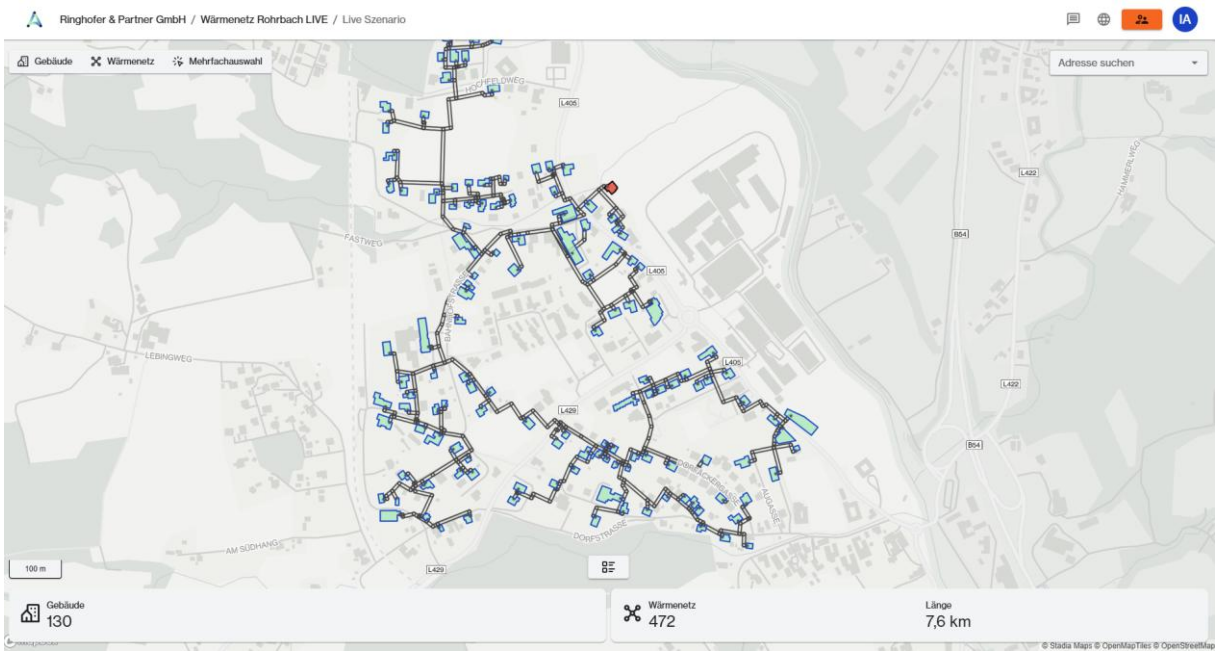

Figure 2: Rohrbach digital twin

Arteria verified node topology and ensured consistency of hydraulic flows. The Rohrbach network was of particular interest because it became the primary site for real-time optimization and algorithmic control integration.
Additionally, Figure 3 shows the digital representation of the district heating network in Güssing.

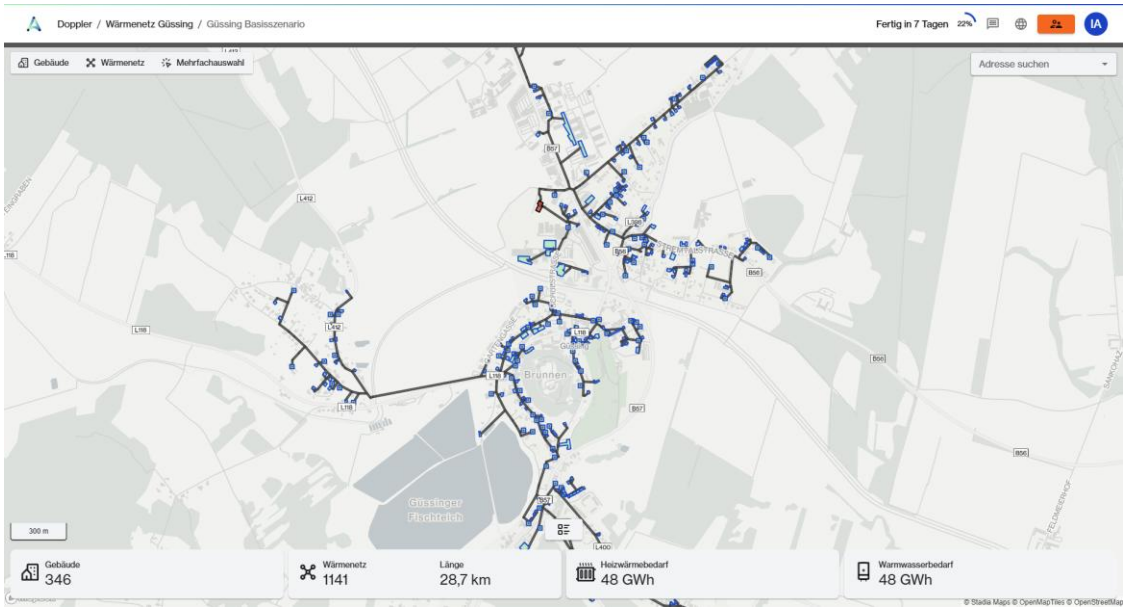

Figure 3: Güssing digital twin

As with the other locations, the digital twin combines static network geometry with time-series data interfaces to allow simulations of dynamic behaviour.

After model development, the next major task involved the connection of real-time data streams from the Schneid system of Rohrbach to the Arteria platform. A cloud-based data-handling architecture was set up, replacing the earlier file-based architecture with a unified geodata server and time-series database. This step created a continuous data pipeline in MS Azure enabling live monitoring of supply temperatures, return flows, pump states, pressure levels and substation behaviour. For Rohrbach in particular, a real-time link to the Schneid control system was implemented. Data were ingested through a dedicated interface, allowing the Arteria platform to receive updated conditions in operational time intervals and enabling detailed visibility of the system's operational behaviour. A screenshot of the integrated environment in MS Azure is given in the next figure.

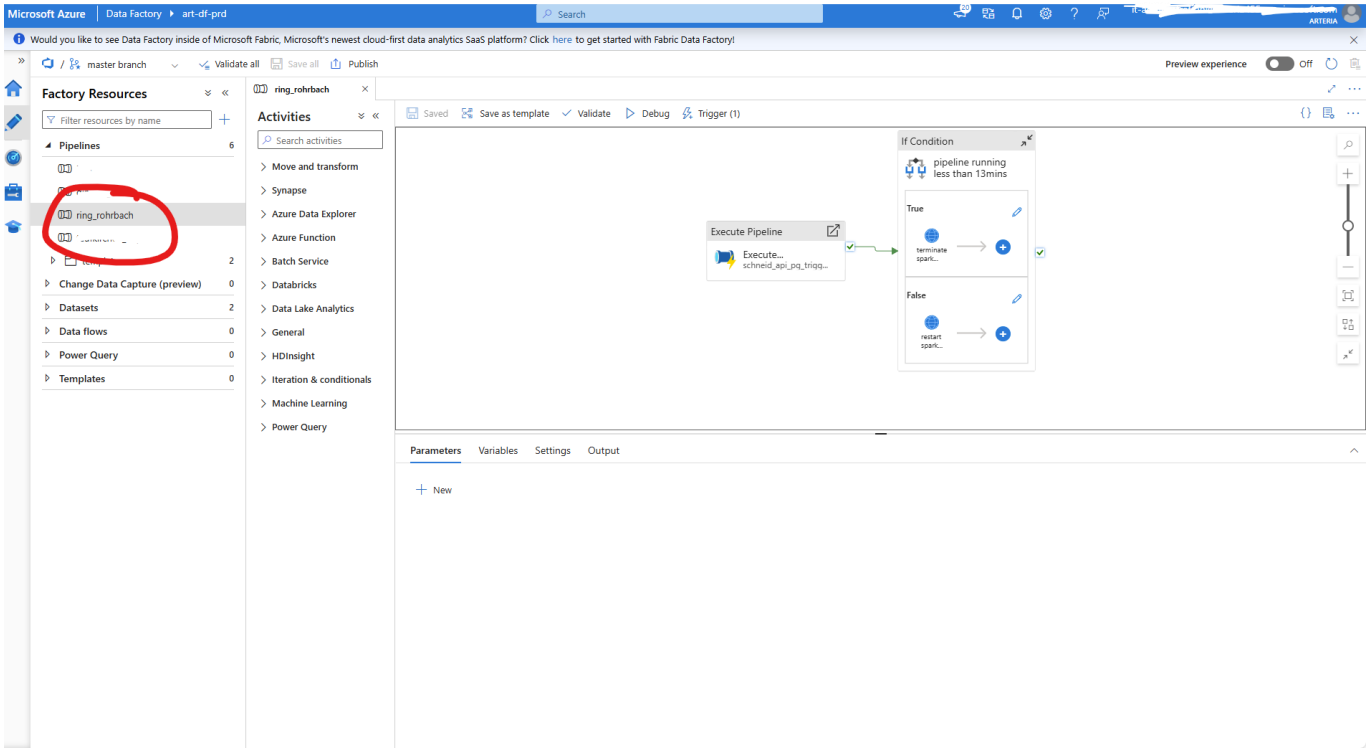

Figure 4: Data ingestion pipeline in MS Azure

In addition to the digital modelling tasks, Arteria redesigned the platform UI/UX to support the visualization of demand-response signals, optimization targets and network status. The updated user interface integrates thermal maps, plant performance indicators, live data trends and direct access to simulations. Live examples from the Rohrbach system are shown in Figure 5 and Figure 6.

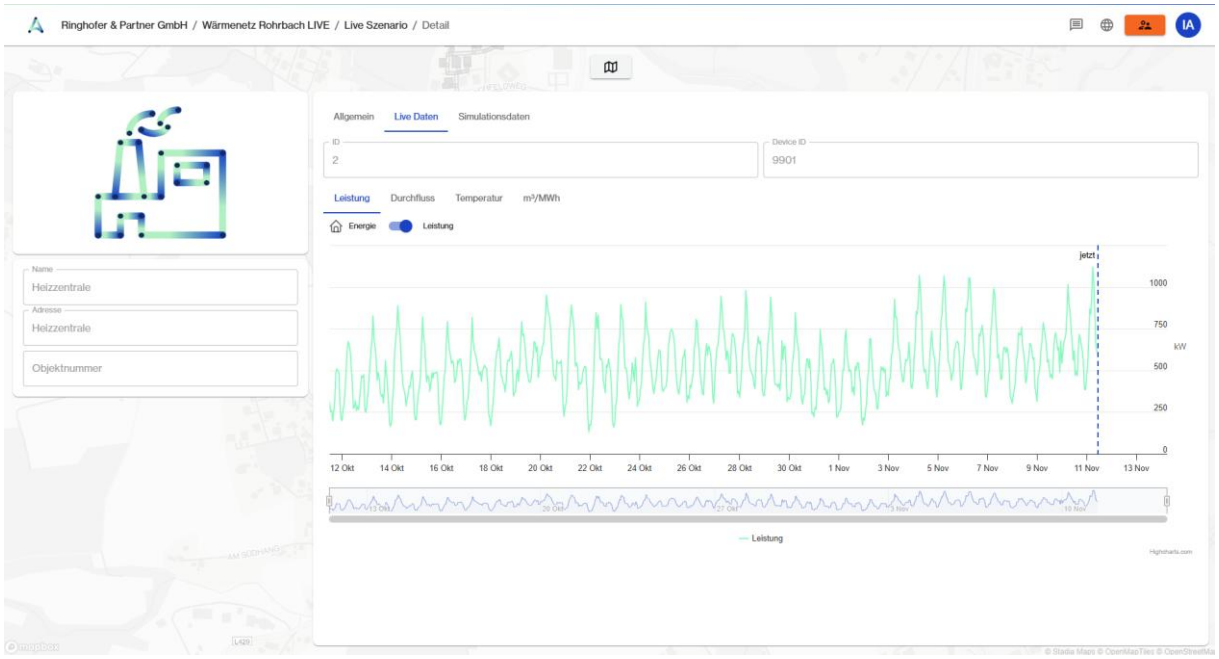

Figure 5: Live data integration of Rohrbach network

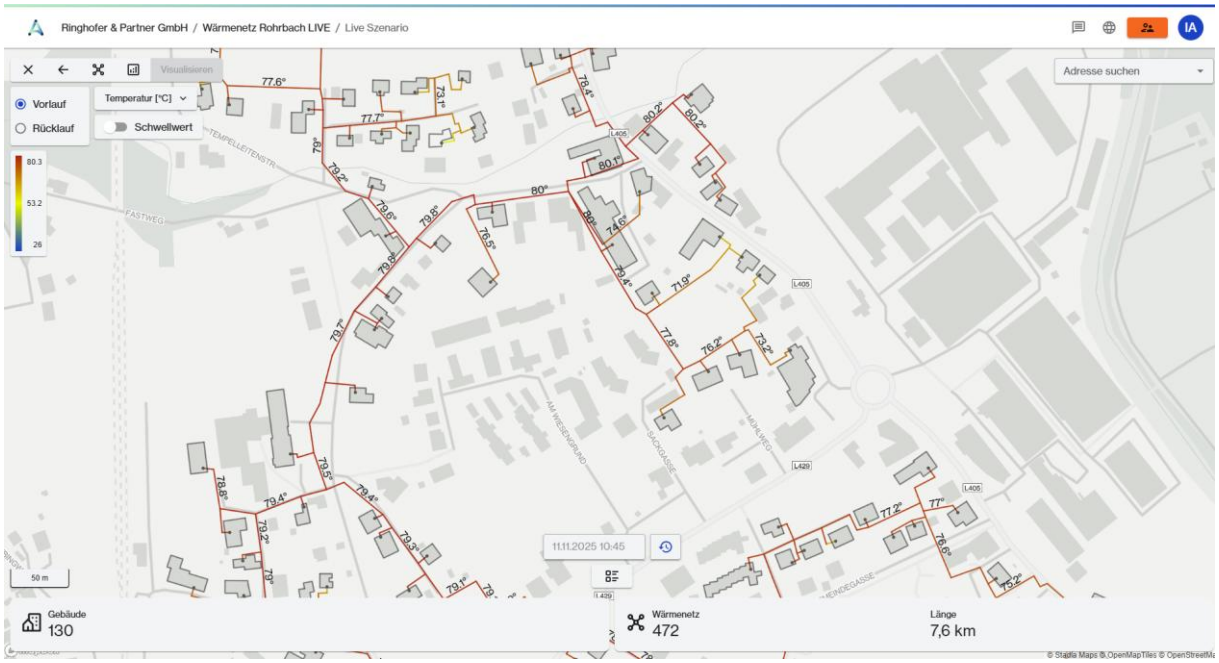

Figure 6: Live view of supply temperature at 11.11.25 10:45

One of the key developments within WP2 and WP4 was the creation of a real-time demand response algorithm specifically tailored to the Rohrbach district heating network. This algorithm uses live measurements from the heating plant, temperature sensors and substation data as inputs. Its objective is to continuously compute the lowest possible supply temperature that still satisfies all consumer heating requirements while minimizing distribution losses. By lowering the plant supply temperature, the thermal losses across the network decrease, particularly during periods with lower heat demand. Lower supply temperature also reduces stress on system components and improves operating efficiency.

Based on detailed simulations using the Arteria platform, an average reduction of approximately 5 °C in the supply temperature was achieved during pilot tests. The algorithm first evaluates the current state of the network and estimates the maximum substation temperature required. A predictive simulation (MPC) is then run using the digital twin to test whether the plant can operate safely at a lower temperature. If safe operation is confirmed, the algorithm transmits the new setpoint to the plant controller via POST request. This process repeats continuously, allowing dynamic adjustments.

To ensure robust performance in varying conditions, the algorithm includes anti-aliasing filtering to suppress noise in real-time measurements, limit detection to ensure safe operation bounds and fallback behaviour in case of network instability or missing data. The full technical description of the algorithm, including parameterization and structure, is provided in the WP3 documentation, and was developed in MATLAB and Python.

During the field tests, Arteria's optimization procedure demonstrated stability and predictability. Transitions between supply temperature setpoints were smooth, avoiding rapid oscillations in thermal behaviour. The system successfully maintained consumer comfort levels at all times. Moreover, by comparing historical uncontrolled data with the optimized mode, the benefit of the approach was clearly visible. The optimized mode exhibited fewer temperature peaks and less variability, indicating lower energy consumption and improved handling of heat demand fluctuations.

The same methodology used for Rohrbach was also evaluated for larger networks. Although outside the scope of this project for validation reasons, such as Taufkirchen, where it showed that the optimization strategy scales effectively across more complex hydraulics.

By completing the tasks of model development, UI redesign, real-time integration and algorithm deployment, Arteria successfully delivered its expected contributions to WP2 and WP4. The digital twins are fully operational and serve as the central technical layer connecting plant operators, data streams and optimization routines. The resulting platform now provides a continuous environment for full operational monitoring and optimization of district heating systems.

In addition to these central developments, a significant amount of work within WP2 focused on the internal consistency checks of the hydraulic and thermal models. Every digital twin created for the project underwent numerous validation stages. These stages ensured that measured real-world data aligned with simulated values under various operating conditions. Arteria implemented an iterative validation loop that compared supply and return temperatures, pressure differences and mass flow rates between the simulated environment and historic operational datasets. When discrepancies were identified, Arteria recalibrated pipe insulation values, adjusted thermal loss coefficients and refined spatial pipe-routing approximations. This calibration process was critical as even small deviations in the model could lead to inaccurate predictions of temperature propagation or pressure stability within the network.

The calibration of the Rohrbach network was particularly complex due to its operational characteristics. The network contains several branches with differing consumer types, ranging from residential buildings

to small commercial users. Each consumer group exhibits distinct heating profiles, resulting in diverse and often unpredictable demand patterns. To address this complexity, Arteria generated synthetic demand curves based on historical consumption profiles and tested them against the real-time measurements. This approach ensured that the digital twin could reproduce both steady-state and transient heating conditions under varying external temperatures and consumer behaviours.

Furthermore, Arteria used exergy costing technology within the digital twin infrastructure to evaluate how efficiently energy was being transported and consumed within the network. Exergy analysis allowed the team to identify areas within the network where the energy quality degraded disproportionately due to heat losses or inefficient heat transfer. These findings were particularly valuable in pinpointing specific sections of the network where optimization potential existed. For example, insulation deficiencies in certain pipe segments could be inferred by comparing expected thermal gradients with real-world data trends. In cases where the discrepancy was too large, Arteria recommended targeted improvements to the network operators to increase efficiency.

Real-time integration required a robust communication architecture to ensure uninterrupted data flow. Arteria implemented a dedicated time-series database optimized for high-frequency updates and large-scale streaming. The interface allowed the platform to maintain a rolling buffer of measured values such as supply temperatures, return temperatures, mass flow rates and pressure readings at multiple network nodes. The Schneid system used by the operator proved to be a reliable data source, although Arteria had to implement fallback data-handling mechanisms to mitigate risks associated with data latency or temporary communication outages. These fallback mechanisms ensured that in the event of missing or delayed data, the digital twin would continue operating based on the most recent validated values, thereby avoiding destabilization of the optimization routines.

3.1.2 GIS Platform

Development of a Unified Smart Data Model for Small District Heating Networks.

3.1.2.1 Objectives and Motivation

Small district heating (DH) networks are an important component of local energy supply systems and play a growing role in the transition toward renewable, decentralized energy infrastructures. However, these networks often suffer from heterogeneous data structures, proprietary formats, and isolated IT systems that make it difficult to exchange, maintain, and analyse information efficiently.
Operators typically deal with separate software tools for metering, billing, monitoring, and geospatial management. The lack of a shared data model hinders interoperability, increases maintenance costs, and prevents the use of modern digital technologies such as predictive analytics, smart metering, or digital twins.
The goal of this project was therefore the development of a standardized, open, and modular Smart Data Model specifically designed for small and medium-sized district heating networks. The model builds upon FIWARE Smart Data Models, which are open standards for structuring and exchanging context information in smart infrastructure applications.

All components were implemented using open-source technologies, ensuring flexibility, transparency, and independence from commercial vendors. The solution can be operated:

- Locally, e.g., in Docker containers or on-premise servers, or
- In the cloud, e.g., using Microsoft Azure Database Services.

The system integrates multiple data management and visualization tools and fulfils the following core functions:

- Centralized storage of both geospatial and attribute data within a PostgreSQL/PostGIS database.
- Preparation and maintenance of GIS data using QGIS Desktop.
- Web-based map visualization with QGIS Web Server.
- Interactive dashboards for time-series and performance data using Grafana.
- Optional gamification features to motivate users by comparing performance and efficiency indicators.

The overarching goal is to establish a sustainable, extendable, and transparent system architecture that provides district heating operators and their customers with valuable insights, key performance metrics, and actionable recommendations.

3.1.2.2 System Architecture and Components

The system architecture follows a modular approach, combining a geospatial database, GIS visualization tools, and data analytics dashboards under a unified data model.

3.1.2.2.1 Central Database: PostgreSQL with PostGIS

All project data is managed in a PostgreSQL database, an open-source relational system known for its stability and flexibility. The PostGIS extension adds support for geospatial objects and functions, enabling the storage, querying, and analysis of features such as pipelines, substations, buildings, and sensors.

Typical deployment options include:

- Cloud-based operation, hosted on Microsoft Azure Database Service, allowing scalable and secure access.
- Local operation, using Docker containers or dedicated servers, suitable for smaller networks or restricted environments.

PostGIS forms the technical foundation for spatial analysis and acts as the common data layer for GIS and analytical applications.

3.1.2.2.2 GIS Management and Editing: QGIS Desktop

QGIS serves as the main tool for creating, editing, and maintaining spatial data. It enables:

- The capture and correction of geometries (pipelines, stations, building connections).
- Data validation and quality assurance.
- Planning and simulation of network extensions.

QGIS provides an open environment where geospatial datasets can be linked directly to the PostgreSQL/PostGIS database. It thus allows operators to work with live data while maintaining complete control over accuracy and consistency.

3.1.2.2.3 Web GIS Visualization: QGIS Web Server

For non-technical users, spatial data can be published via QGIS Web Server, enabling browser-based access to maps and attribute information.
Key features include:

- Interactive navigation through the network layout.
- Object selection and display of detailed metadata.
- Optional hyperlinks connecting static features to live measurement data (e.g., real-time flow or temperature).

Questions of data protection and confidentiality are addressed by defining clear rules for which datasets can be made publicly available and which remain restricted to authorized staff.

3.1.2.2.4 Dashboard and Analytics: Grafana

Grafana is an open-source platform for visualizing time-based and operational data. Within this project, it provides intuitive dashboards for monitoring key performance indicators (KPIs) of the district heating system.
Grafana supports two primary user perspectives:

- Operator dashboard: An overview of the entire network, including total energy consumption, flow rates, and aggregated statistics.
- User dashboard: Personalized views focusing on individual buildings or substations, including comparisons with averages or rankings based on efficiency.

Extensions such as gamified leaderboards (e.g., ranking users by energy savings) can enhance motivation and promote more sustainable consumption behaviour.

3.1.2.3 Smart Data Model Design

3.1.2.3.1 FIWARE and the Smart Data Model Concept

The Smart Data Model framework is based on FIWARE, an open platform initiative developed to support Smart City and IoT applications. FIWARE defines a common method for managing context information—that is, data describing entities (devices, assets, processes) and their current state.
Using NGSI (Next Generation Service Interface) as the standard communication protocol, FIWARE ensures that systems can share and interpret data consistently, regardless of the vendor or platform. This interoperability is key for building modular and future-proof systems.
The project's data model adheres to FIWARE principles and reuses existing Smart Data Models wherever possible, extending them where necessary for district heating applications.
Main conceptual elements include:

- Device Type: Classification of measurement or control equipment (e.g., heat meter, flow sensor, temperature sensor).
- Device: Specific device instance, with unique ID, attributes, and location.
- Time Series: Historical measurements stored in a structured form and linked to the corresponding devices.

3.1.2.3.2 Entity-Relationship Model (ERD)

The logical relationships between entities were designed to represent both the physical infrastructure and its operational data:

- FW_Network: Represents the overall network structure.
- FW_Station: Interface points to buildings (heat transfer stations).
- FW_Pipeline: Connective infrastructure between stations.
- Device: Sensors or control units linked to specific components.
- TimeSeries: Associated measurement data over time.
- Customer: Links user accounts to physical assets and devices.

This ER structure provides a coherent basis for integrating spatial, technical, and operational information within one database.

3.1.2.4 Data Integration and Processing

3.1.2.4.1 ETL Workflow

Initial data import and harmonization were carried out via an ETL (Extract–Transform–Load) process:

1. Extract: Import of existing datasets from CAD, Excel, CSV, or shapefile sources.
2. Transform: Conversion and harmonization according to the Smart Data Model structure.
3. Load: Insertion of standardized data into the PostgreSQL/PostGIS database.

This process was executed once for the initial setup. Future updates can be performed directly through QGIS or automated import routines.

3.1.2.4.2 Real-Time Data Interfaces

To enable dynamic monitoring, real-time measurement data is integrated via defined interfaces. Currently supported systems include:

- Schneid (for heating and metering data)
- MEO (measurement and energy optimization systems)

Data streams are processed and continuously written into the database, from where Grafana retrieves them for visualization and trend analysis.

3.1.2.4.3 4.3 Integration with the Arteria Platform

The developed system is designed to interoperate with the Arteria platform, enabling data exchange and joint visualization. Two database user roles are defined:

- Read user: For viewing live and geospatial data (e.g., for dashboards or web maps).
- Read/Write user: For editing data, such as network planning or maintenance records.

This interface supports collaborative workflows and reduces redundancy between systems.

3.1.2.5 Visualization and Dashboards

3.1.2.5.1 GIS Visualization

In QGIS Desktop, the full network is visualized in detail, with symbols and colors representing attributes such as pipe material, construction year, temperature level, or operational status.
Through QGIS Web, selected data layers can be published and enriched with hyperlinks to live data, enabling intuitive access for users and stakeholders.

3.1.2.5.2 Grafana Dashboards

Grafana transforms raw sensor data into clear and interactive visual insights.
Examples of available dashboards:

- Time-series plots showing temperature and energy consumption trends.
- Bar charts comparing users or buildings.
- KPI panels summarizing network efficiency, peak loads, and energy balances in real time.

By offering both high-level overviews and detailed drill-downs, Grafana supports data-driven decision-making for both operators and end-users.

3.1.2.6 Applications and Service Potential

The developed system opens various possibilities for operational improvement and commercial services.

3.1.2.6.1 Data and Software Services

- Provision of data in standardized formats such as JSON, GeoJSON, CSV, or FIWARE-compatible NGSI.
- Setup and configuration of the complete open-source infrastructure (database + GIS + dashboard).
- Customized dashboard solutions tailored to the needs of specific operators or clients.

3.1.2.6.2 Analytical and Consulting Services

- Automatic generation of recommendations for efficiency improvements.
- Calculation of potential savings in kWh or euros through specific measures.
- Visualization of environmental benefits in relatable equivalents (e.g., trees saved, kilometres driven by an electric car).

3.1.2.6.3 6.3 Integration with Arteria and Future Extensions

The ongoing connection to the Arteria platform allows:

- Synchronization of maintenance and service data.
- Direct integration of expansion plans.
- Future extensions toward energy communities, sector coupling, or cross-network optimization.

3.1.3 Connection and Data Interfaces

Because the participating networks differ significantly in their existing SCADA and visualization environments and in site-specific data governance constraints, multiple integration paths were required, including OPC, HTTPS APIs, file exports, and periodic manual meter reads. The common objective was to provide reliable, time-stamped operational data streams that can be consumed by the project databases and by the Arteria platform for monitoring, analytics, and digital-twin parameterization.
From a technical perspective, the interfaces were designed as layered pipelines: (i) data acquisition at the network or customer interface, (ii) secure transport (where permitted) using VPN or on-premise buffering, (iii) intermediate storage in time-series databases (InfluxDB 2) to decouple acquisition from downstream services, and (iv) data harmonization to achieve comparable signals across customers and networks. This architecture reduced dependency on vendor systems and improved resilience against connectivity interruptions.

Cross-cutting data handling and quality assurance

Across all demonstrators, particular emphasis was placed on operational robustness and reproducibility of the data pipelines. To ensure that the interfaces deliver data suitable for modelling and platform services, the following cross-cutting measures were applied where technically and organizationally feasible:

- Timestamp normalization and consistent time zones to avoid offsets between vendor systems and databases.
- Buffering at the acquisition side (database or local InfluxDB 2) to bridge short-term connectivity interruptions.
- Plausibility checks (range checks, missing-value detection, and basic consistency checks such as supply ≥ return temperature).
- Signal harmonization using synonym mappings to support cross-site KPI definitions and digital-twin parameterization.
- Access control aligned with site-specific data protection requirements, VPN connectivity was used only where approved by the operators.

Rohrbach – OPC-based link to Arteria

In Rohrbach, current data from the substations are accessed via the built-in API of the visualization software. The system provides an OPC interface which is read by a service running on the BEST server. The retrieved values are buffered in a database and then forwarded to the Arteria server, allowing access to both live and historical data in a consistent format.
The pipeline was implemented with a focus on stable, low-latency transfer of the most relevant operational variables (e.g., temperatures, mass flows, and power). Buffering on the BEST server ensures continuity in case of short outages and provides a reproducible basis for model calibration and KPI generation.

Mischendorf – dual interface approach (OPC/API and CSV export)

In Mischendorf, two interfaces were established. First, the current network state was read via the OPC server of the network visualization. In addition, the vendor enabled a token-based HTTPS API that provides the current status of all meters in JSON format. The intended approach was to ingest these data continuously (analogous to Rohrbach) by connecting the BEST server to the visualization system via VPN and writing the values directly into the database. During the project period, the existing OpenVPN client failed and could not be reactivated. As a consequence, the API-based interface was limited to local data collection on site. Figure 7 shows the excerpt from the API response for a customer.


```
{
    "device" : "2",
    "items" : {
        "mError" : "0.00",
        "mError_upd" : "2026-01-21 08:22:04",
        "mFlow" : "2060.00",
        "mFlow_upd" : "2026-01-21 08:22:04",
        "mPower" : "63.10",
        "mPower_upd" : "2026-01-21 08:22:04",
        "mPrimFlowTemp" : "79.00",
        "mPrimFlowTemp_upd" : "2026-01-21 08:22:04",
        "mPrimReturnTemp" : "52.00",
        "mPrimReturnTemp_upd" : "2026-01-21 08:22:04",
        "rPrimValve" : "19.00",
        "rPrimValve_upd" : "2026-01-21 08:22:04",
        "rSecFlowTemp" : "67.90",
        "rSecFlowTemp_upd" : "2026-01-21 08:22:04",
        "rSecReturnTemp" : "51.80",
        "rSecReturnTemp_upd" : "2026-01-21 08:22:04",
        "sSecFlowTemp" : "74.90",
        "sSecFlowTemp_upd" : "2026-01-21 08:22:04",
        "sysCommLostRetard" : "0.00",
        "sysCommLostRetard_upd" : "2026-01-21 08:22:04",
        "xTime" : "08:20"
    }
},
```


Figure 7: Sample excerpt from the API response for a customer

Second, the visualization system offers an integrated export function that writes measurements to a CSV file every five minutes. The file is continuously appended, enabling access to historical values. An operational constraint is that these files are deleted twice per year. To prevent data loss, an R-based ingestion routine was implemented to parse the CSV file every eight hours and write newly available records into a local InfluxDB 2 time-series database (Figure 8).

Figure 8: Local InfluxDB2 database

In parallel, the processed Mischendorf data are mirrored to a third database operated by the project partner Scheiber Solutions GmbH. This database is directly connected to the Arteria platform, supporting unified access for analyses and visualization across partners.
The CSV export provides a substantially larger set of signals than the API. While the API exposes only a small number of core values (on the order of single digits), the CSV files contain several hundred data points. The richer CSV dataset includes, for example, cumulative energy (kWh), volume (m³), thermal power (kW), volume flow, supply and return temperatures, temperature spread, storage tank temperatures, outdoor temperature, controller setpoints, valve positions, and pressure readings.
A key challenge is that signal ordering and naming are configured per customer, resulting in multiple labels for the same variable. To harmonize the dataset across customers, a synonym mapping was created and non-essential fields were filtered out to obtain a consistent dataset for analytics and digital-twin calibration. Table 1 shows the applied mapping logic.

Table 1: Applied mapping logic

Name (canonical)	Synonyms / alternative labels
Timestamp	Date-Time
Heat amount	WZ_Waermemenge_kWh_
Volume	WZ_Volumen_m3_
Thermal power	WZ_Leistung_kW_
Volume flow	WZ_Durchfluss_l_h_
Return temperature	WZ_Ruecklauftemperatur__C_; WZ_Ruecklauftemp___C_
Supply temperature	WZ_Vorlauftemperatur__C_; WZ_Vorlauftemp___C_
Temperature spread	WZ_Spreizung__C_
Storage tank temp (bottom)	Boilertemp__unten__C_; Temp_Speicher_1_unten__C_; C-T10_Speicher_unten__C_
Storage tank temp (top)	Boilertemp__oben__C_; Temp_Speicher_1_oben__C_; C-T9_Speicher_oben__C_
Outdoor temperature	Aussentemp___C_; T6_Aussentemp___C_
Primary return temperature	Ruecklauftemp__primaer__C_; T7_RL_Primaer__C_
Secondary supply temperature	Vorlauftemp__sekundaer__C_; T8_VL_Sekundaer__C_
Heating circuit 1 supply temperature	Vorlauftemp__Kr_1__C_; 1-Vorlauftemp___C_; A-Vorlauftemp___C_
Heating circuit 2 supply temperature	Vorlauftemp__Kr_2__C_; 2-Vorlauftemp___C_; B-Vorlauftemp___C_
Heating circuit 3 supply temperature	Vorlauftemp__Kr_3__C_; 3-Vorlauftemp___C_; C-Vorlauftemp___C_
Secondary return temperature	Ruecklauftemp__Sekundaer__C_; T11_RL_Sekundaer__C_
HC1 setpoint	VL_Soll_Kr_1__C_; 1-Solltemp___C_
HC2 setpoint	VL_Soll_Kr_2__C_; 2-Solltemp___C_; B-Solltemp___C_

HC3 setpoint	VL_Soll_Kr_3__C_; 3-Solltemp___C_; C-Solltemp___C_
Secondary supply setpoint	Soll_VL_sekundaer__C_; Solltemp__VL_Sekundaer__C_
Primary valve position	Primaerventilstellung_proz_; Primaerventilstellung_%_; Ventilstellung_%_
Pressure sensor 1	Ain1_bar_; ZAin1_bar_
Pressure sensor 2	Ain2_bar_; ZAin2_bar_
Customer ID	KundenID

Güssing – periodic readout of Kamstrup heat meters

In the Güssing district heating network, no continuous data transmission from customer substations to a central system was available. Six large consumers, relevant for overall demand, were selected for detailed monitoring. Different Kamstrup heat meters are installed at these sites. The meters store sensor values (supply and return temperatures, volume flow, and power) in internal archives with different horizons: an annual archive (one snapshot per year), a daily archive (one value per day for approximately one year), and an hourly archive (approximately 36 days; older data are overwritten). Because the annual and daily archives are too coarse for dynamic modelling, longer time series were built by reading out the hourly archive periodically. Over more than one year, meter data were collected monthly using an optical readout head, thereby reconstructing higher-resolution histories despite the limited on-device retention.

Güssing – Belimo energy valves as supplementary data source

A second source of customer-side measurements in Güssing are Belimo energy valves installed at selected interfaces. These devices provide local data storage, a web server, and an API. The API exposes only current sensor values (temperatures, volume flow, power); therefore, continuous recording would require a permanent connection to a readout computer. Permanent internet connectivity was generally not accepted by the participating companies.

Where feasible, data were collected via temporary VPN access to the valve web interface and its internal data stores. Data are available either in a monthly memory (values stored every two hours and retained for more than one year) or in a daily memory (values stored every 15 minutes; retention increased from 14 to 31 days after a firmware update).

Manual downloads via the web interface require extensive interaction. With support from the manufacturer, the data access was reverse engineered to enable automated downloads. A batch script generates the required file names and downloads the data via wget. The script is date-driven, starting from the current day and attempting to fetch the previous days. If a day does not exist or no data are available, an empty file is returned. Downloaded files are automatically checked and imported into an InfluxDB database using an R script. See details at Figure 9.

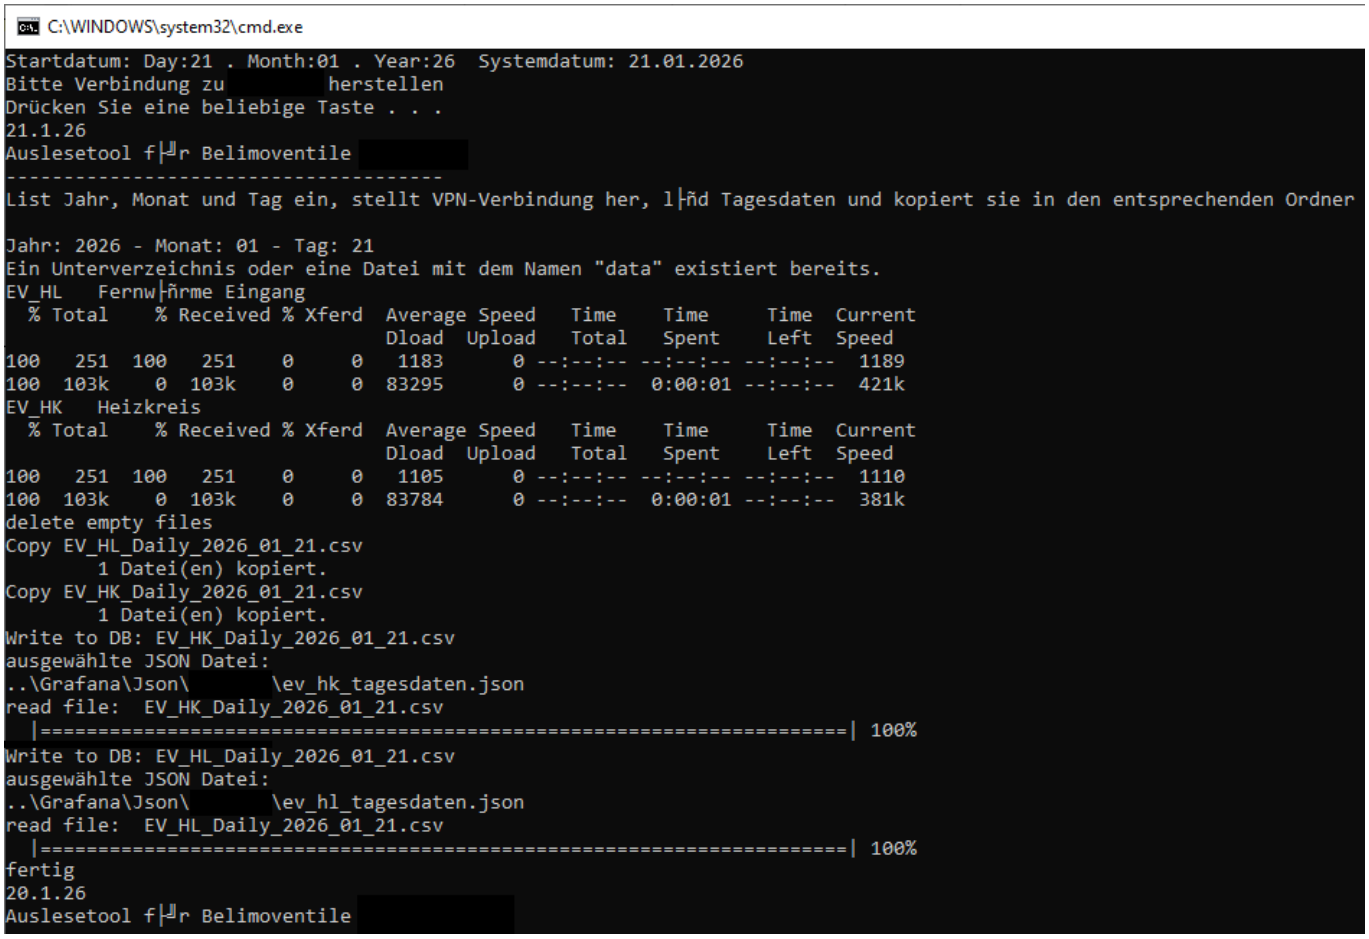

Figure 9: Readout script for Belimo valve - Download, check, and write to database

Japons – Danfoss interface via a dedicated Python client

In Japons, a Danfoss visualization system is used. Measurements are stored on Danfoss data servers and made accessible via an API. Since only an interface specification was available, and no ready-to-use client packages existed for integration into project software, a dedicated Python package was developed during the project. Separate functions were implemented for device metadata (customers), measurement subjects per customer (signals), last-value retrieval, and historical data queries. Although the API also offers write functions, these were intentionally not used within the project.

For bulk data acquisition, the workflow is as follows:

1) Retrieve the list of customers (DanfossAPI.get_all_devs()).
2) Retrieve the list of measurement points for each customer (DanfossAPI.set_subslist(devlist)).
3) Retrieve last values (DanfossAPI.get_all_last_data(devlist)).
4) Store the results in an InfluxDB time-series database using a dedicated ingestion module.

The Python package was made available to all project partners to enable consistent access to the data for analyses and model integration.

MEO Controller Interfaces and External Data Sources

In addition to the district heating interfaces described above, an existing Python package for the installed MEO controller was used to access controller data and, where required, implement parameter changes. These data streams were integrated into the digital twin.

Overall, the interface landscape combined standardized protocols, vendor-specific APIs, and pragmatic file-based exports. The main technical challenges were heterogeneous signal sets, inconsistent naming across installations, limited retention policies on local devices, and operational constraints such as VPN availability and site-specific data protection requirements. The buffering, harmonization, and ingestion routines implemented by GET established a robust basis for downstream project components, including the digital twin, dashboards, and optimization workflows.

3.1.4 Connection to Smart Home System

The DOPPLER project aims to increase the efficiency and flexibility of decentralized energy systems. A key element is the integration of intelligent control units and data-driven optimization methods. In this context, two meo PROPILOT controllers were installed in Rohrbach to evaluate their potential in terms of energy efficiency, user comfort, and district heating network optimization.

The installations serve as a practical test environment. By combining local control, central data analysis, and an open interface architecture, an important step towards a connected, adaptive energy system is being implemented.

3.1.4.1 Objectives of the integration

The integration of the meo PROPILOT controllers pursues several technical and strategic goals:

- Increasing energy efficiency:
  Adaptive control of supply temperatures reduces energy losses without compromising comfort.
- Improving user interaction:
  Homeowners gain insight into their building's heat flows via an intuitive user interface and can make individual adjustments.
- Optimizing system integration:
  The controllers communicate via a standardized 0–10 V interface with the Schneid transfer station to automatically and demand-oriented request the required supply temperature.
- Data-driven network optimization:
  Through the planned API interface to Arteria, operational data is provided for optimizing the district heating network, especially during peak loads or additional heating requirements.

These objectives form the basis for an intelligent, connected heat management system that combines local efficiency with system-wide stability.

3.1.4.2 Technical implementation

In two Rohrbach households, a meo PROPILOT controller was installed and integrated into the existing heating infrastructure. It handles central control of the heating circuits, monitoring of temperature values,

and communication with higher-level systems. The dashboard overviews are shown in Figure 10 and Figure 11.

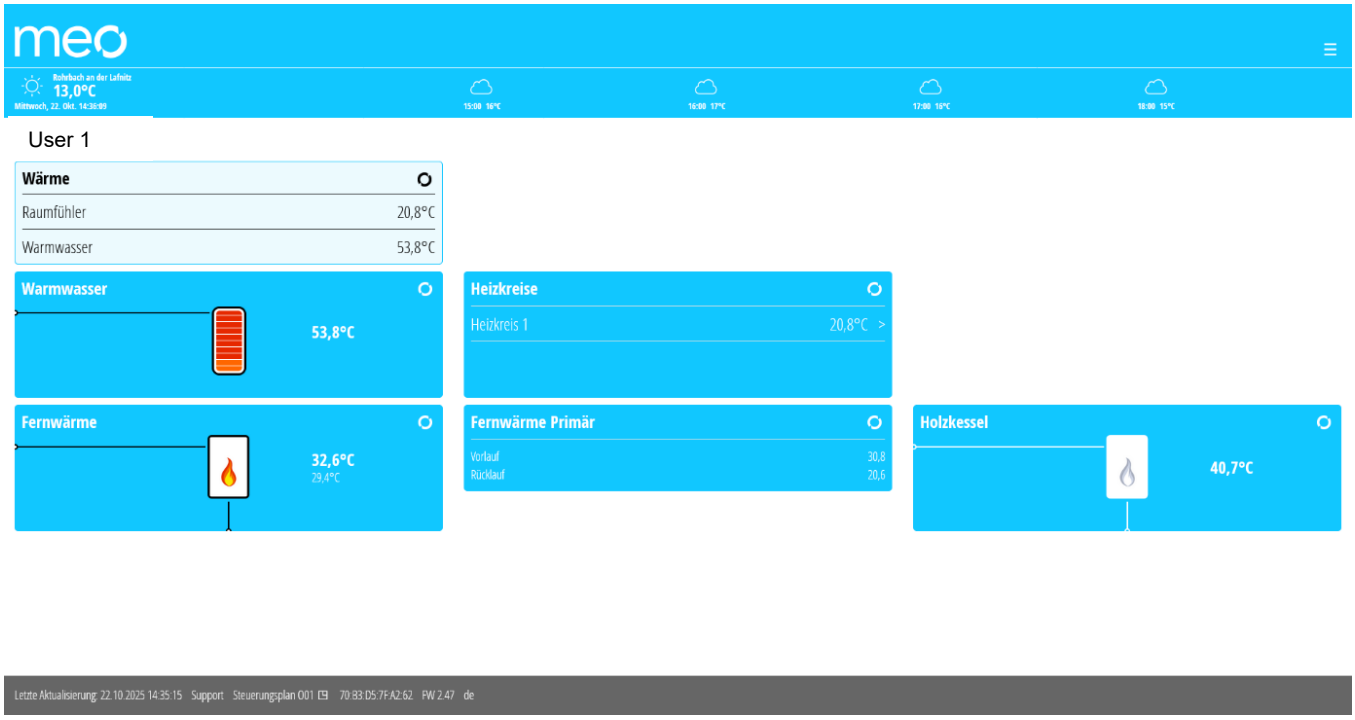

Figure 10: Dashboard overview for end user 1

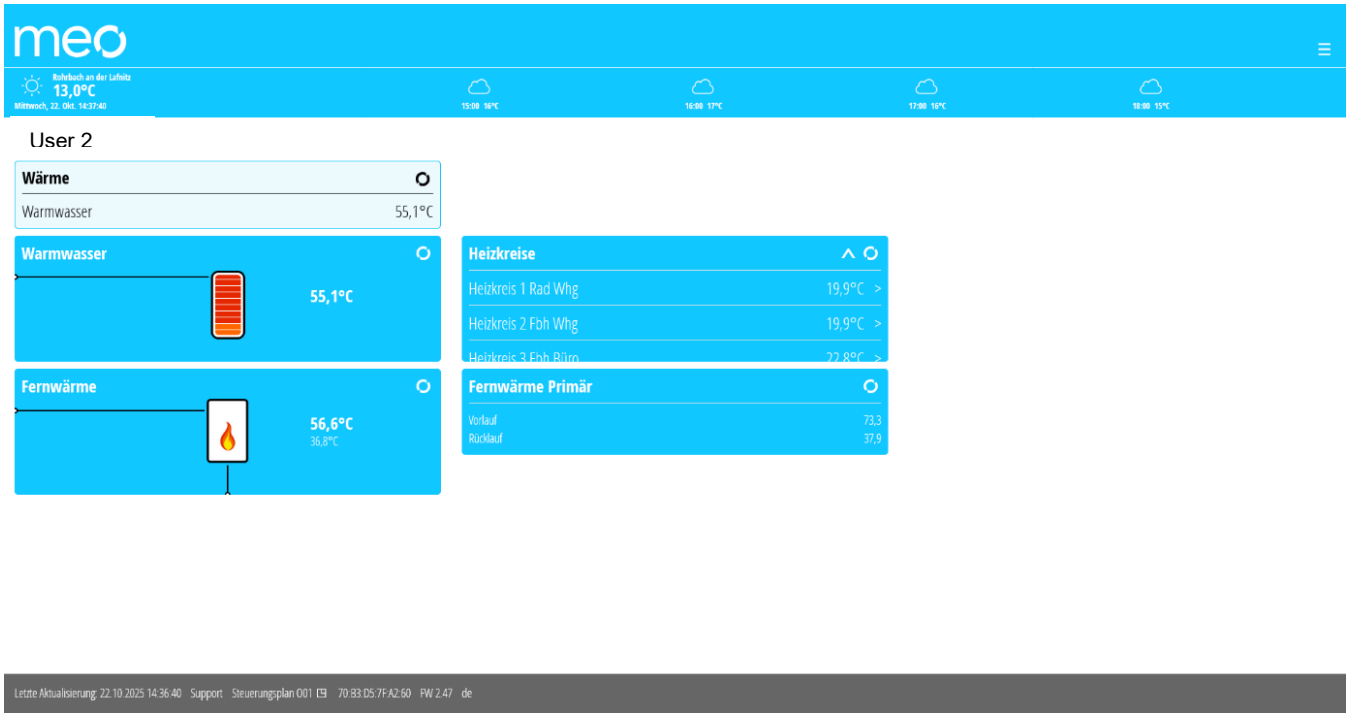

Figure 11: Dashboard overview for end user 2

3.1.4.3 Core functions

Following core functions were provided:

1. Personalized user experience:
   The interface allows homeowners to view supply and return temperatures, room temperatures, and setpoints, and adjust them if necessary.

2. Energy efficiency through weather forecast integration:
   The control strategy uses outdoor temperature forecasts (e.g., 4 hours ahead) to minimize heat losses and increase the efficiency of heat generation.

3. Interoperability and communication:
   The controllers request the required supply temperature from the Schneid transfer station via the 0–10 V interface.

The end user adjustment options are shown in Figure 12.

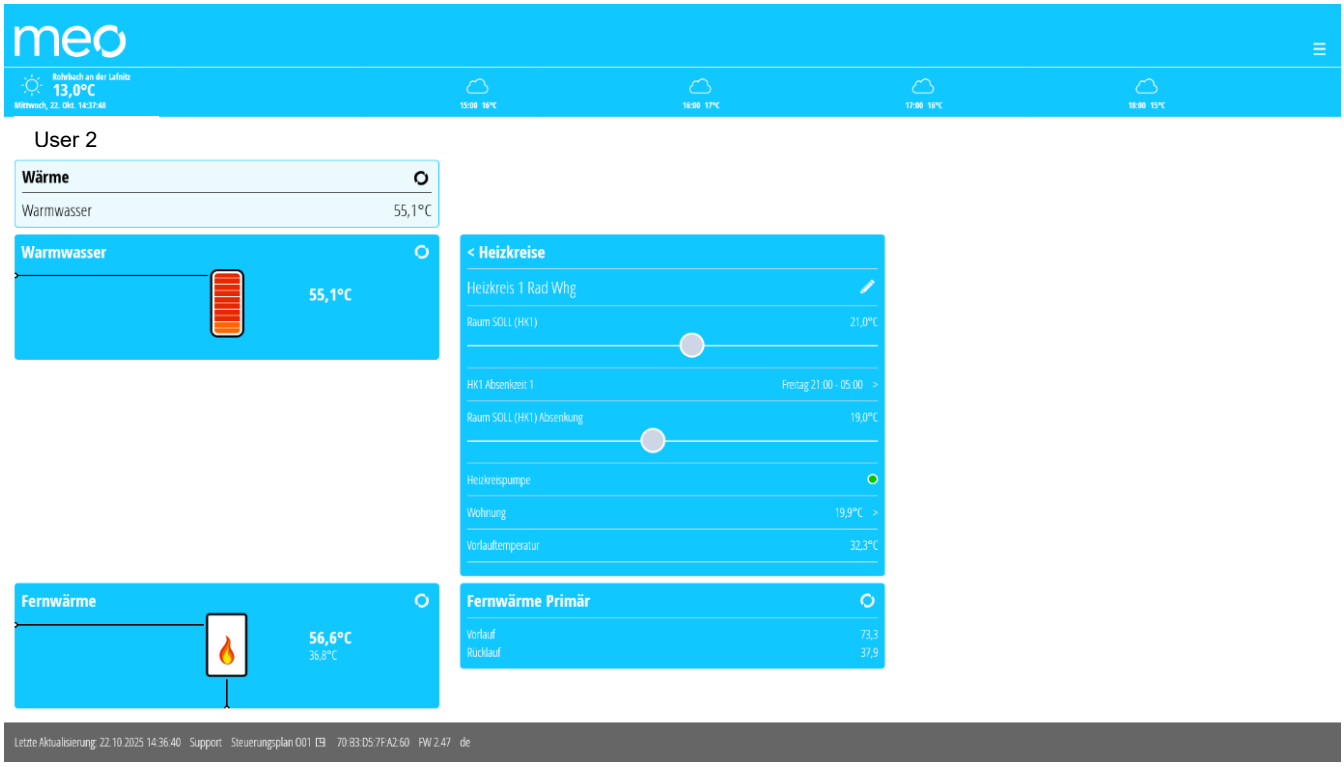

Figure 12: User adjustment options for end user

3.1.4.4 Extension via System_PROFILER

In addition to the standard functions of the PROPILOT controller, the System_PROFILER (Figure 13) is used. This module enables detailed analysis of the entire heating system and provides valuable insights for optimization:

- System analysis: Captures and evaluates all relevant thermal and electrical parameters of the heating circuits.
- Long-term monitoring: Stores and analyses historical data to support forecasts and control strategies.

An integrated alarm system notifies users of critical conditions or deviations via email:

- Monitoring critical parameters: Room and supply temperatures, domestic hot water, and heating circuits are continuously monitored.
- Automatic alerts: Notifications are triggered immediately when predefined thresholds are exceeded or system faults occur.
- Operational safety and comfort: This reduces downtime and ensures occupant safety and comfort.

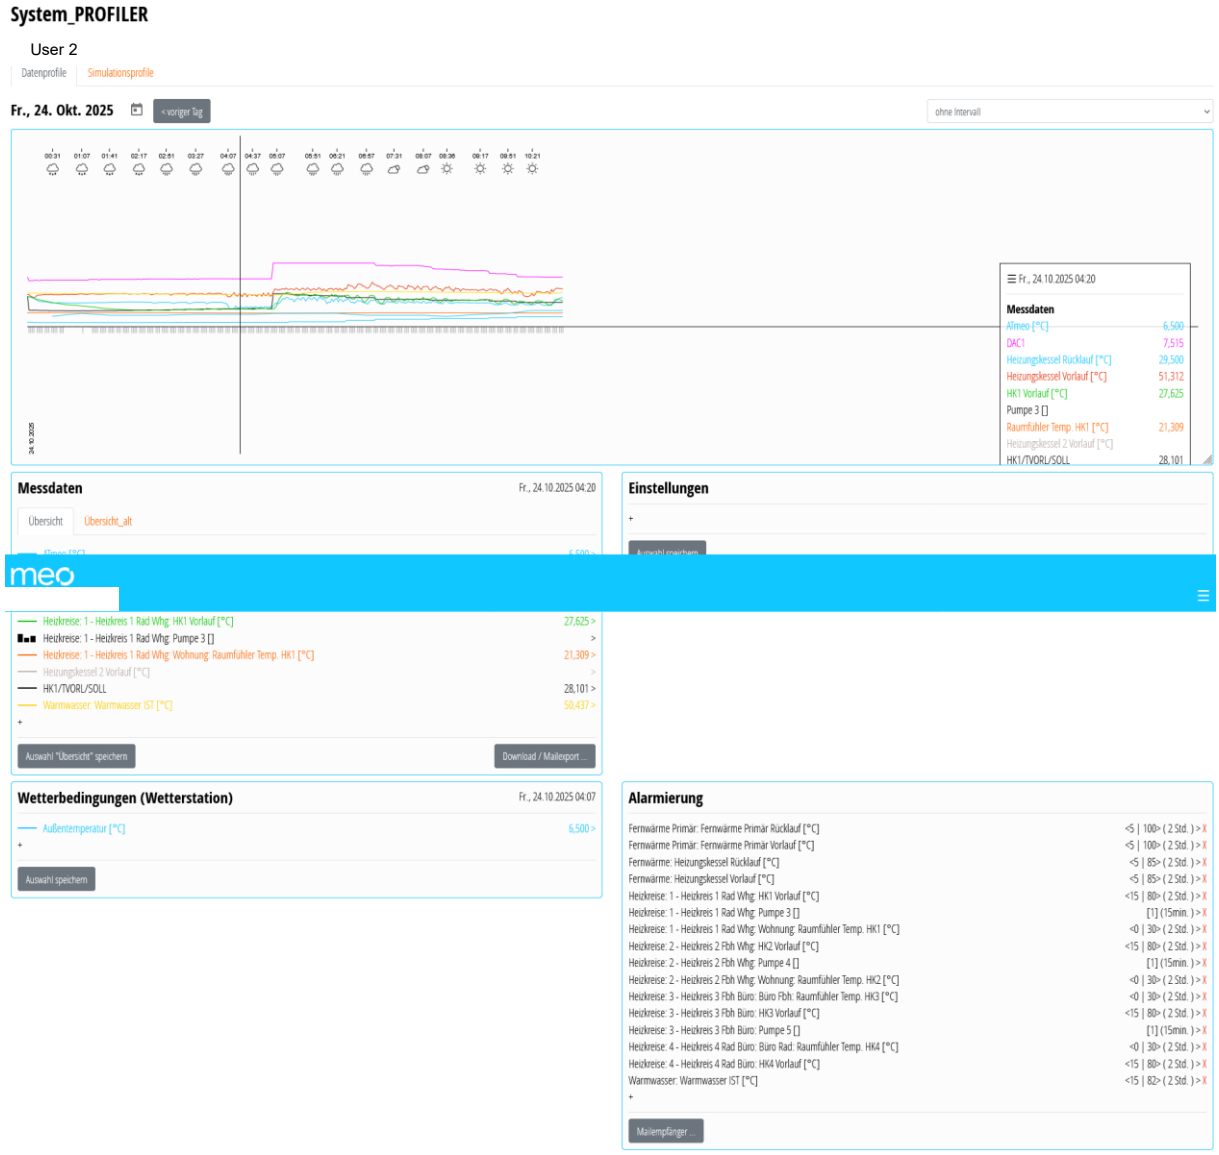

Figure 13: System_PROFILER

3.1.4.5 Collection of relevant data points

All relevant measurement and control data from the demonstration buildings are systematically recorded via the meo PROPILOT controllers. These include:

1. Supply and return temperatures (primary and secondary circuits)

2. Setpoints of the heating circuits (HK1–HK4)
3. Room temperatures
4. Domestic hot water temperatures
5. Outdoor temperature forecasts

These data are made available to project partners Arteria and Scheiber Solutions GmbH via an API interface.

3.1.4.6 API integration and data processing

A Python-based API is used for automated acquisition of historical and current data. The core function allows the retrieval of measurement data, including energy flows, temperature values, battery data, weather information, and energy price data.
As the data are not always available in uniform temporal resolution, the processed datasets are resampled and cleaned before being provided to the partners. These data form the basis for optimization algorithms, analyses and modelling.

3.1.4.7 Role as a data platform

For Arteria and Scheiber, meo acts as a central data hub:

- On-site: Local control and real-time regulation
- In the cloud: Aggregation and provision of data to partners

This architecture enables both local optimizations and network-wide analyses based on a consistent data foundation.

3.1.5 Optimization and Demand Response

An optimized operation of a heating network can be achieved by three separate measures: optimizing the heat production side, the network control side, and/or the substation and consumer side. In this project, the heat production side was optimized using a predictive, optimization-based controller by BEST. Existing models were extended to integrate flue gas heat recovery and the optimization problem adapted to the demonstrator Rohrbach. For the network control side, a predictive algorithm determining the minimally feasible feed temperature making use of the Arteria Digital Twin was developed and tested in simulations. Finally, the integration of the meo platform to provide customer side optimization potential was investigated and methods to integrate this flexibility potential into the supervisory optimization problem were developed.

3.1.5.1 Predictive supervisory control of heat production

Optimizing the heat production side focuses on unit commitment, i.e., deciding which producer should be turned on or off and at which part load they should be operated. Optimization potential stems from maximizing the yield from renewable sources, prioritizing production from cheaper heat sources, operating producers at points with higher efficiency (boilers) or COPs (heat pumps) or, in the case of

varying tariffs or market participation, when electrical power is cheaper (heat pumps) or more expensive (CHPs). This kind of high-level control is typically done using rules, i.e., switching boilers on when a temperature sensor reading in a thermal buffer drops below a threshold, and switching it off when a maximum temperature is exceeded at a lower sensor. To avoid simple on/off switching behaviour, additional low-level controllers such as PID controllers can be used to track, e.g., the desired mean temperature of the buffer storage, thus leading to modulated operation and avoiding unnecessary on/off switches which typically result in higher maintenance costs.
While simple and robust, these rules tend to be too conservative or need to be re-adjusted depending on the season or expected consumer behaviour. As an alternative, optimization-based methods can benefit from predictions of the yield from renewables or of the future heat demand to make more informed decisions, avoid unnecessary start/stop cycles and modulation of producers by making better use of available flexibility options such as thermal storage or demand shift capabilities of consumers.
Optimization-based methods formulate an optimization problem whose optimization variables contain the decision variables, or setpoints, that must be sent to the low-level controllers: On/off switching commands, modulation percentages, temperature setpoints etc. The goal of the optimization is formulated as an objective function, typically a cost function that needs to be minimized and depends on the fuel consumption and possibly other considerations. The schematic overview is shown below.

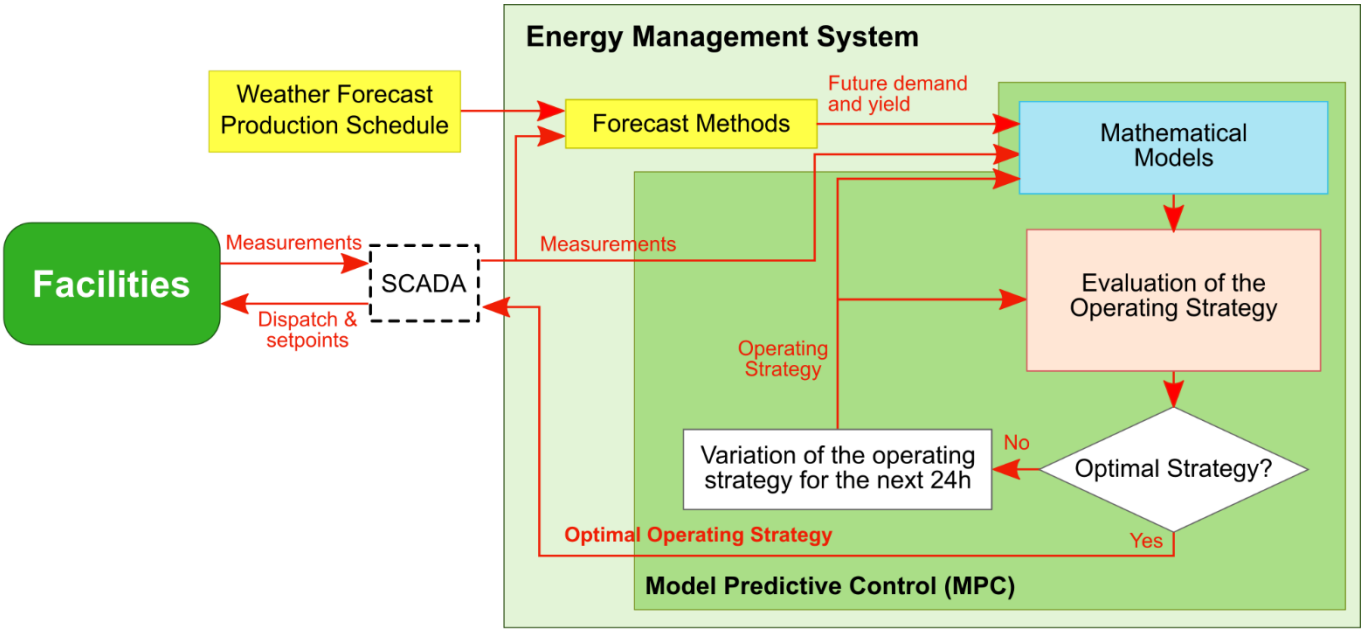

Figure 14: Schematic overview of the optimization-based supervisory controller (here called Energy Management System) using Model Predictive Control (MPC)

A common approach to using optimization to control energy systems is so-called Model Predictive Control, where these decision variables are linked together in a mathematical model that represents the energy system. In it, the effects of choosing specific setpoints are modelled and the dynamics of, e.g., the thermal storage or biomass boilers, are represented to predict the behaviour of the energy system in the near future, e.g., the next 48 hours. Since these decisions are based on uncertain forecasts of future yield or demand, and since the mathematical models also are only an approximation of reality, the resulting schedule is not followed through until the end, but instead only applied for a short duration, e.g.,

15 minutes. Then, new information is used to update the current state estimate and all forecasts, and the optimization is repeated. This is called the Moving Horizon principle.
The challenge with an optimization-based approach is that higher numbers of optimization variables, and especially of the binary decisions that need to be made (on/off, yes/no), increase the complexity of the optimization problem, in the case of binary variables even exponentially. Therefore, and because the optimization needs to be repeated in real time, detailed mathematical models such as used in the Digital Twin cannot realistically be used. Instead, simplified linear models are often used, resulting in Mixed Integer Linear Programs that need to be solved. The challenge thus is to define sufficiently simple models that still represent the real behaviour sufficiently well so that he resulting operating schedule is acceptable to the operator, and ideally more efficient than the operation based on simple rules and more conventional controllers.

For DOPPLER, the heating central of Rohrbach was investigated in more detail and modelled in the modular Energy System Management framework of BEST. Its schematics are depicted in Figure 15. The setup is sufficiently simple: A biomass boiler of 1,5 MW_{th} provides heat to the network, either directly or via a sensible storage of 52.000 l. The return flow can be deviated to a flue gas condensation unit. There is also a connection to a sawmill which previously required heat, but is now supplying the network with excess heat in summer. The representation of this energy system in the optimization framework is shown in Figure 15, where lines between components indicate balancing equations, i.e., ensure that all heat produced is also consumed and all wood chips burned are provided by a source and paid for.

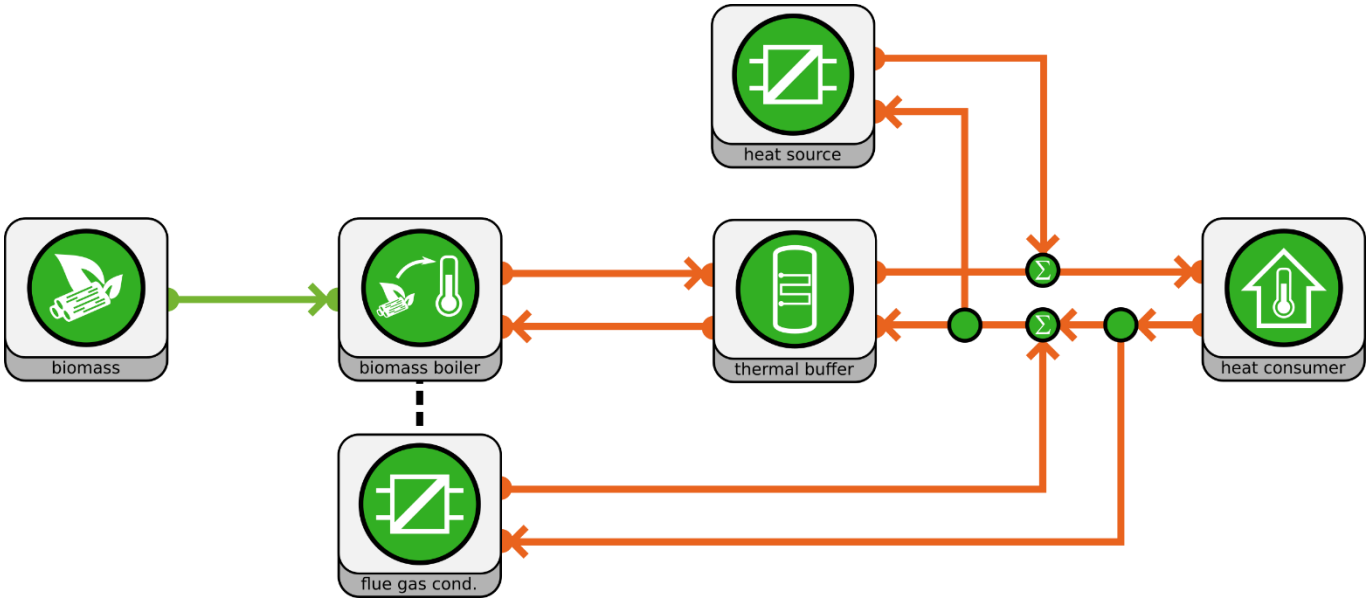

Figure 15: Configuration of the energy system representing Rohrbach inside of BEST's Energy Management System.

The heat source in the diagram corresponds to the sawmill, and the flue gas condensation is directly linked to the biomass boiler operation: It can only be activated if the biomass boiler is active itself.
The flue gas condensation was modelled as an optional heat source with a maximum output proportional to the current heat output of the boiler. The efficiency of heat recovery was determined from an affine relation between the respective heat outputs of the boiler and flue gas heat recovery, see Figure 16. For a time range from April 2023 to March 2024, this resulted in an offset of 150 kW and a slope of 0.239, i.e. an efficiency of about 23,9%. Dependencies on temperatures were currently ignored but could be

added to incorporate, e.g. predictions from the Arteria simulation, with higher network return temperatures resulting in lower efficiency values.

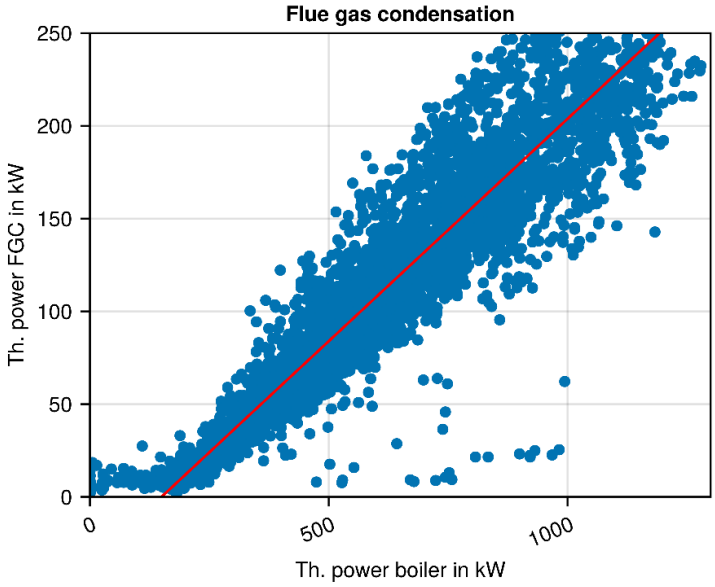

Figure 16: Identifying the efficiency of the flue gas condenser from measurement data in November 2023 using linear regression

The flexibility options of the setup are rather limited since there are no qualitatively different sources of heat in the system: The boiler is the only controllable source, and the utilization of flue gas heat recovery, whenever possible or necessary, is self-evident.
The predictive controller thus is limited to improving the following aspects of operation: reduce the on/off switches of the boiler and operate the boiler at its most efficient operating point. A typical efficiency curve of a biomass boiler is visualized in Figure 17. This was approximated in the optimization as a piecewise affine function (PWA) to stay within the MILP context. Depending on the current demand, operating at high efficiency can be difficult and might lead to more on/off switching cycles. This is typically difficult to tune and involves selecting appropriate costs for on/off switching. As a benefit, the operation will automatically adjust to the current demand profiles and needs no further tuning by the operator during the year.

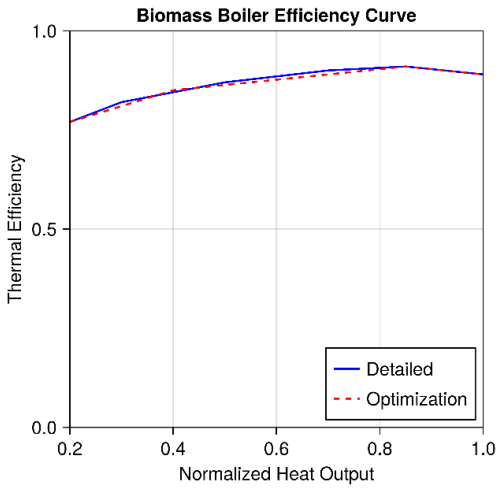

Figure 17: Biomass boiler efficiency curve: A detailed curve (blue) is approximated by a piecewise affine function (red, dashed).

Real operation is using a controller that tracks the mean storage temperature. This leads to the boiler output closely matching the current demand, with no on/off switching as long as the demand is within the operating range of the boiler. During summer, heat is provided by the sawmill, so the typical problems of summer operation, i.e., frequent on/off switching or operation at inefficient operating points, are not a real problem of the plant.

Optimization was performed in a simulation with a time resolution of 15 min for the first 4 ½ h, then 1 hour until 12 hours into the future, 2 up to 24h, 3 up to 36h, and finally 4h up to 48h. As perfect foresight was assumed, the optimization was repeated every 4h, with all the 15 min values in between stored. The boiler was restricted to ramping up and down by 100% in one hour, and virtual costs were introduced to avoid fast ramping as well as emptying the storage too much (< 30% state of charge) or letting it overheat too much (> 80% SoC).

During winter, heat demand is sufficiently high for the boiler to operate at better efficiency levels. Figure 18 shows a comparison of historic operation in 2024 and the corresponding duration when operated by the predictive controller. The very high buffer temperatures in historic operation offer a lot of safety but result in higher heat losses. The optimized operation makes more use of the stratified storage; if more energy should be maintained in the storage, this would be a simple parameter in the optimization indicating the required output power and provision time.

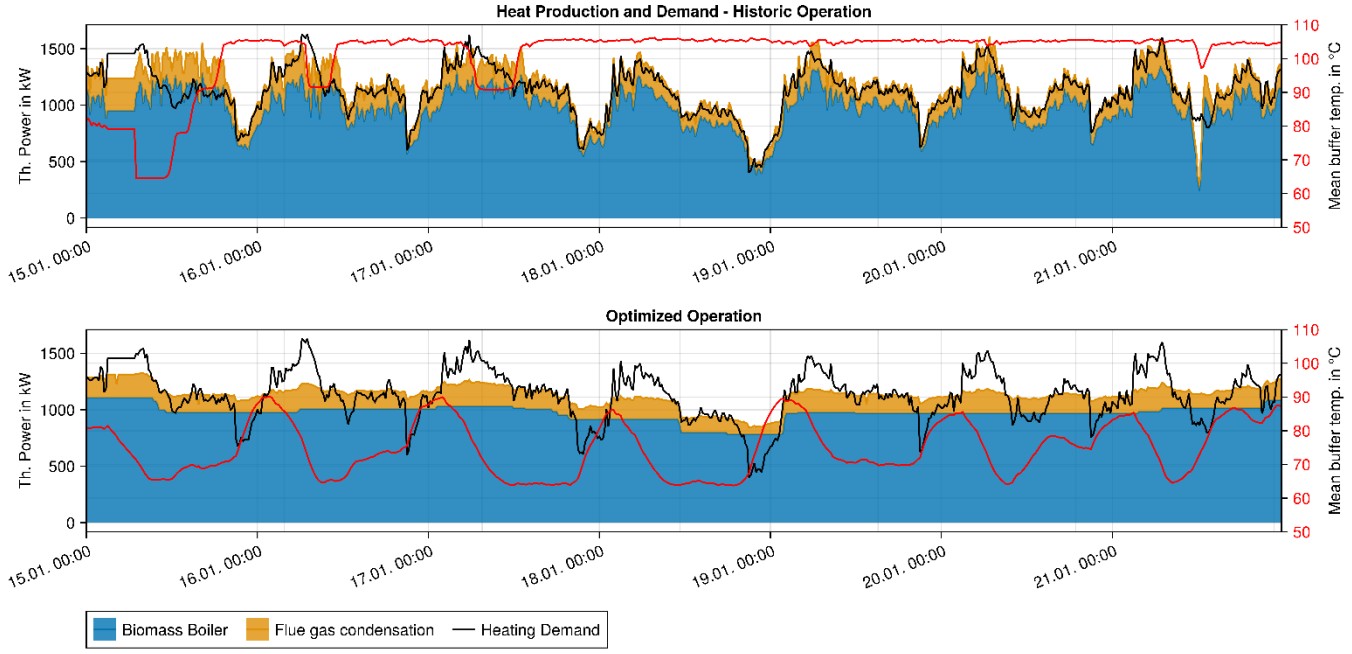

Figure 18: Comparison between historic operation and result from the predictive controller for a week in January 2024. The mean buffer temperature is shown in red.

The relative times the boiler is operated at which output power is shown in Figure 19. While historic operation followed the heat demand, the optimized operation maintains a more constant operating point close to the efficiency optimum.

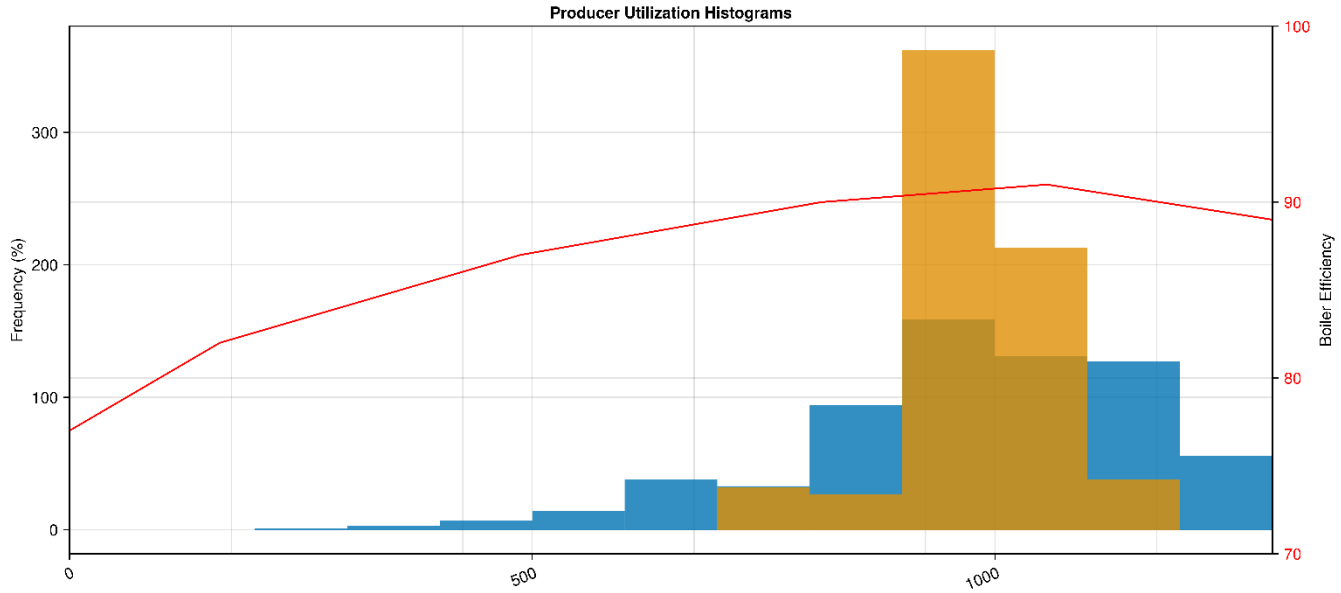

Figure 19: Histogram of output power frequency compared to efficiency curve for the week in January 2024. Historic operation in blue, result from the predictive controller in orange. The optimized operation has a slightly higher share of good operating conditions.

During a week in May 2025, demand is rather low, but the sawmill did not yet provide any waste heat, leading to operation at very bad efficiencies as depicted in Figure 20. Depending on the setup, the optimization would like to switch off the boiler for some time and operate it at a better efficiency later. This is not possible if the boiler must be manually ignited. If forced not to switch off, optimized operation has no other choice but to also remain at low output levels. As soon as possible, however, it tries to operate at higher output levels.

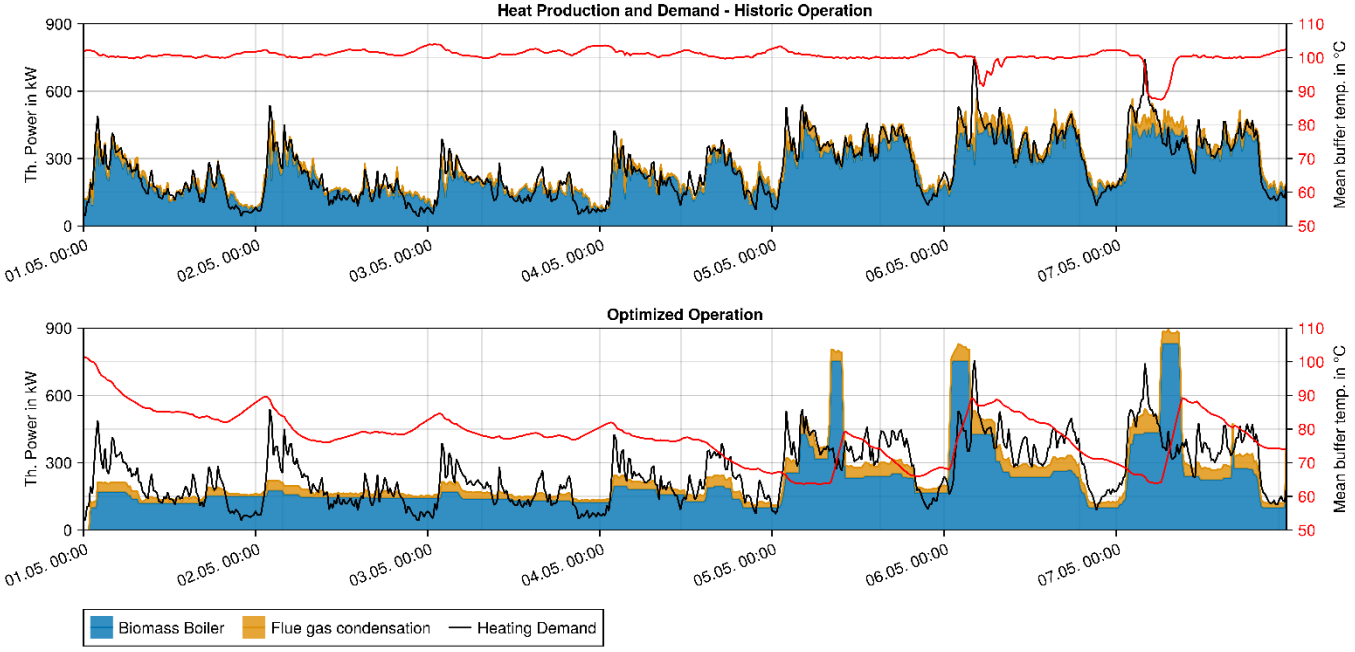

Figure 20: Comparison of historic and optimized operation during early May in 2025. Heat demand is low and the boiler needs to operate at a low power output in both cases.

Predictive operation of the energy production side based on optimization is most useful when dealing with multiple heat sources, fluctuating renewables or other circumstances that benefit from predictive

behaviour such as variable energy prices. For the simple setup found in Rohrbach, its possibilities are quite limited, but can lead to more constant operation, reducing the thermal stress of continuous modulation. Operation can be simplified by specifying goals, such as the energy to be stored in the storage, instead of having to tune PI controller parameters or selecting temperature sensors for 2-point controllers. This has the benefit of automatically adapting to different boundary conditions, e.g., lower or higher heat demand over the year.
After optimizing the heat production side, the next step is to lower the feed temperature to reduce losses in the network itself.

3.1.5.2 Optimal feed temperature selection using the Digital Twin

In this work, a control algorithm for power plant supply temperature was developed, utilising the Arteria district heating network model parameterised using measured data. The WP was completed in three stages, each designed to improve the control algorithm, increase computational efficiency and integrate the control methodology into the Arteria platform.

In the first stage, a model-based control algorithm was developed using an optimisation approach to determine the optimal supply temperature for the power plant. The optimisation procedure followed these steps:

1. Initially, the maximum allowable supply temperature was applied.
2. The network model was then used to calculate individual supply temperatures at each substation.
3. The substation with the lowest supply temperature was identified and compared to the minimum required substation supply temperature.
4. The power plant supply temperature was iteratively adjusted until the lowest substation temperature was as close as possible to the minimum required temperature.

The optimal power plant supply temperature was then defined as the temperature that ensured at least one substation met the minimum supply temperature requirement. This algorithm was extended to scenarios where two power plants operated in parallel, with the mass flow rate split between them in a fixed ratio. Figure 21 displays the optimisation results for a single power plant, while Figure 22 shows the optimisation results for dual power plant operation in the network. Although the algorithm successfully identified optimal supply temperatures, its convergence rate was too slow for practical use.

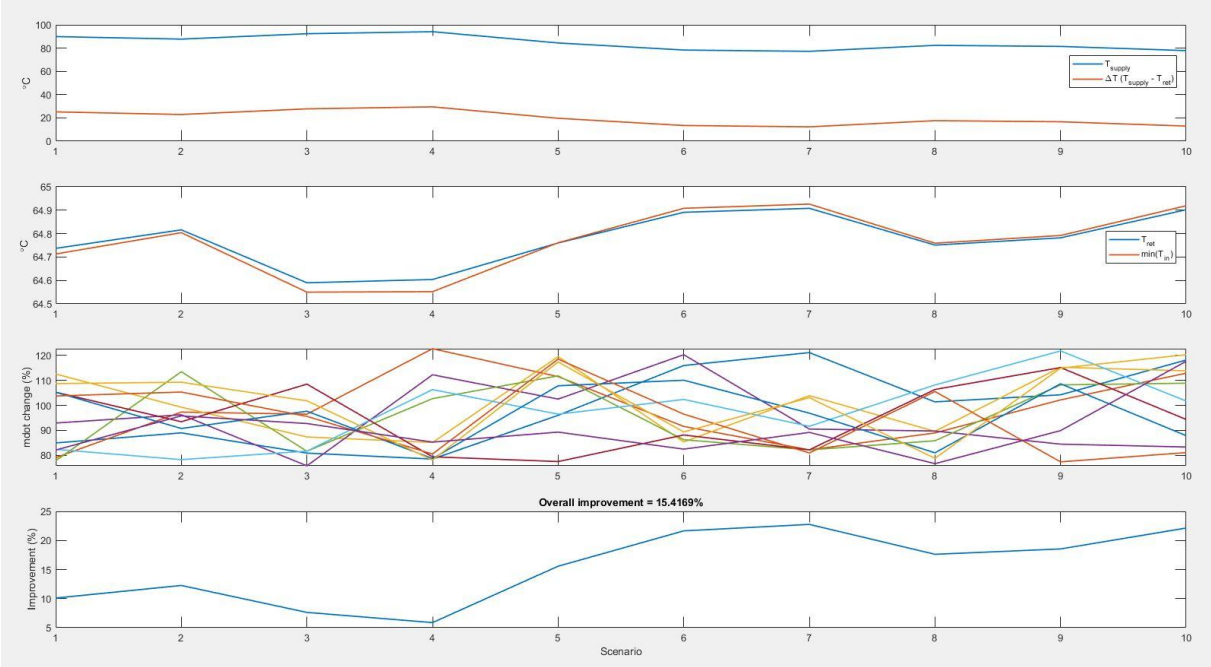

Figure 21: Optimization results for a single power plant.

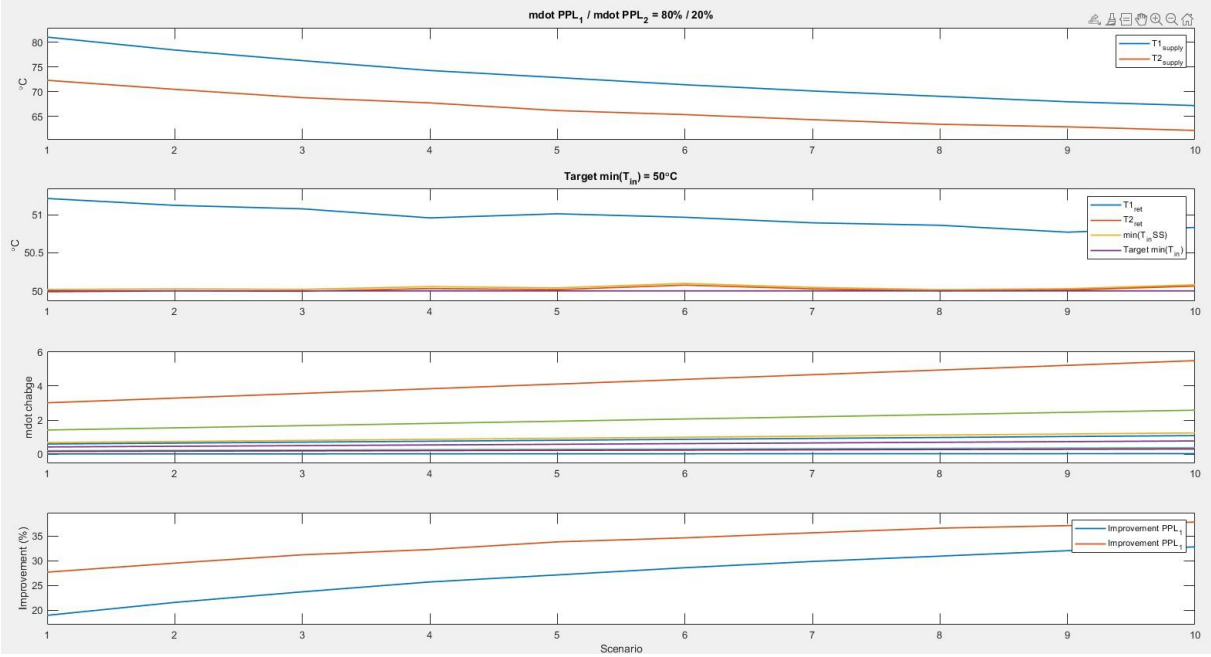

Figure 22: Optimization results with dual power plant operation with 80%/20% mass flow rate split.

To address the convergence issue, the control algorithm was parallelised and executed on a virtual machine using MATLAB's Parallel Toolbox. In this parallel optimisation approach:

1. Six instances of the network model were run concurrently for each time step, with each model set to a different power plant supply temperature.
2. These six supply temperatures were spaced equidistantly between the maximum and minimum allowable values.
3. Each model iteration calculated individual supply temperatures at substations, identifying those below and above the minimum required threshold.

4. Interpolation was then used to determine the exact temperature that met the minimum required substation supply temperature, allowing for the identification of the optimal power plant supply temperature.

With 64 cores available on the virtual machine, optimal power plant supply temperatures for 10-time steps were calculated simultaneously. Following figure illustrates the utilisation of 60 cores running in parallel. This parallelisation accelerated the control algorithm, reducing the time required to complete a 4-week optimisation with a 5-minute sampling interval to 3.36 days.

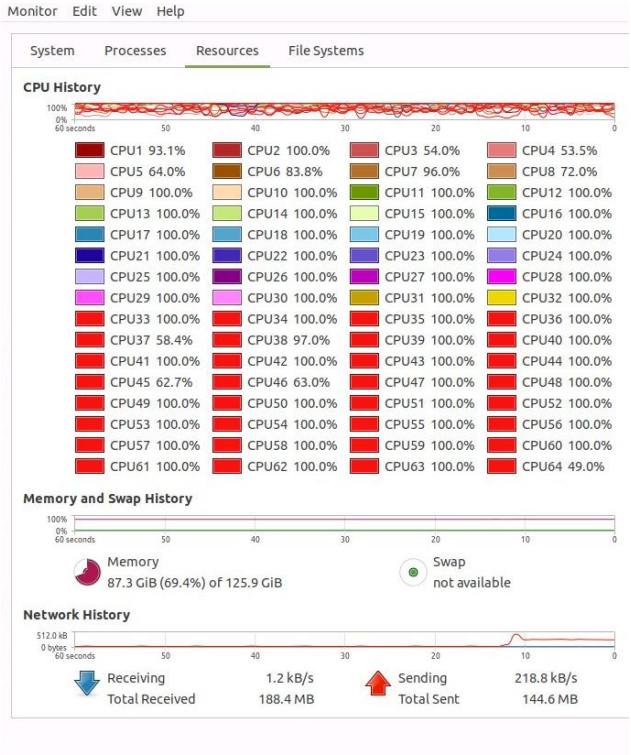

Figure 23: Core utilization in parallelized algorithm execution.

In the final stage of the WP, the MATLAB code was translated into Python and integrated into the Arteria platform. Next figure compares the results generated by the MATLAB code and the Arteria-integrated Python code, demonstrating a minor difference attributed to rounding at the third decimal place.

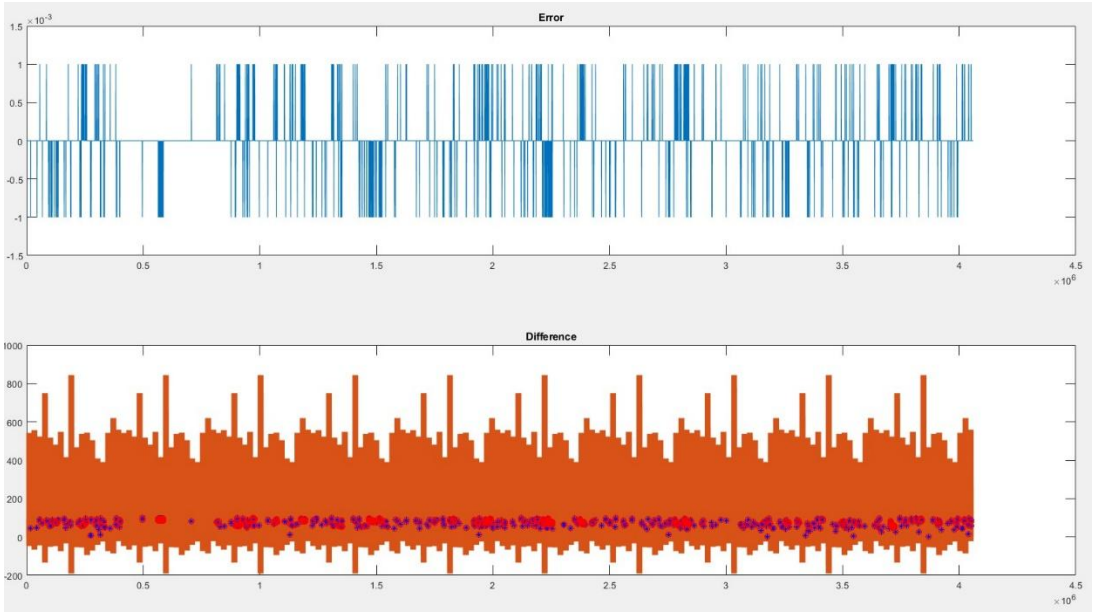

Figure 24: Comparison of results from MATLAB and Arteria platform.

The initial prototype of the control algorithm was developed using the Arteria model implemented in MATLAB. This prototype was extended with anti-aliasing filters and adapted for online, live operation. Data exchange with the SIOCS web server was handled via the curl utility. Data retrieval was performed by downloading JSON responses from the server. Data transmission was carried out through curl POST requests, sending JSON generated by the control algorithm. Two successful live tests were conducted and documented. Full parameterization of the control algorithm, including the design of anti-aliasing filters and its Python implementation, is available.

Following the successful control of a pilot network, the algorithm was further validated using an enhanced Arteria model representing the larger Rohrbach network

The results of the pilot network control simulation are presented in the following figure.

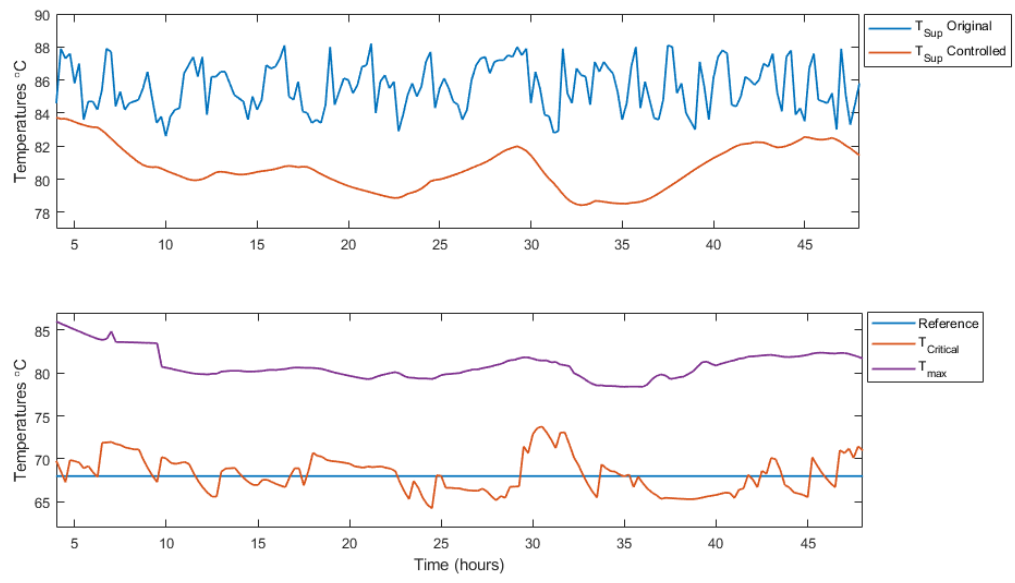

Figure 25: Development of optimum MPC controller

The lower plot in Figure 25 illustrates the critical setpoint temperature, defined at 68 °C (blue line). The critical temperature response (red curve) remains close to this setpoint throughout the simulation. Although the deviation is marginally higher than that observed in the pilot network, the control loop still demonstrates very robust and well-performing behavior, maintaining stability and effectively minimizing deviations from the setpoint despite the larger and more complex dynamics of the Rohrbach network.

The upper plot compares the power plant supply temperature achieved under the control strategy (red curve) with the historical operation without control (blue curve). Under control, the supply temperature is significantly reduced while exhibiting a smoother and less oscillatory profile. This smoothing directly translates to lower thermal stress on equipment, improved operational efficiency and reduced energy consumption, as the plant avoids unnecessary temperature peaks and rapid fluctuations that were present in the previous uncontrolled mode.

Also visible in the lower plot is the maximum temperature among all substations (purple curve). This value evolves more gradually than the critical temperature and stabilizes at approximately 80 °C. The observed behavior suggests that controlling either the maximum substation temperature directly, may yield even better overall network performance than relying on the critical temperature metric. Analysis of recorded data from the Rohrbach network supports this hypothesis and highlights opportunities for further efficiency optimization in future control strategies.

Efficient operation of the district heating network requires precise regulation of the individual mass flow rate at each substation. Currently, this mass flow rate is adjusted by manipulating the primary-side control valve, which is governed by the local substation controller. However, a detailed recorded data analysis indicated that the performance of these local controllers is frequently suboptimal due to poorly tuned mass flow control loops, which significantly reduce the efficiency of the overall network. To improve the overall efficiency, two main strategies are proposed

In this approach, the Arteria platform gains access to the existing substation controllers, enabling centralised tuning of their control parameters. The objective is to ensure consistent and optimised performance across all substations by systematically tuning controller parameters based on the measured system dynamics.

Before tuning can be applied, a detailed understanding of the existing controller architecture and actuation mechanisms is necessary. This includes:

- Determination whether the valve exhibits linear, equal-percentage, or other nonlinear flow characteristics.
- Investigation whether the valve is controlled via direct position commands (0-10 V, 4-20 mA) or motorised actuation timing (incremental open/close control).
- Determination of minimum and maximum valve openings.
- Estimation of time delays, actuator response times and hysteresis.

The tuning process begins with system identification, producing a simplified dynamic model of the substation[1], where the input (manipulated variable) is valve opening command and output (controlled variable) is secondary outlet temperature Tout_S. Identification can use historical operational data and/or open-loop excitation tests (e.g., low-amplitude PRBS signals) to derive a transfer function or state-space representation of the substation thermal dynamics.

A number of autotuning techniques are applicable once the system dynamics are known. For existing substation controllers tuning the most suitable approaches are:

[1] The substation model presented here differs from the substation model developed in WP1.

- Ziegler–Nichols-based tuning: Based on the identified model, the controller parameters can be calculated using Ziegler–Nichols recommendations.
- Optimisation-based tuning: Formulate an objective function (e.g., minimise mean squared error or weighted energy cost) and solve using numerical optimisation.
- Adaptive/gain-scheduled PI control: Dynamically adjusts parameters based on operating conditions.

The autotuning workflow can be summarised as:

1. Data acquisition:
    a. Collect input/output data (valve command vs. Tout_S).
    b. Perform optional open-loop PRBS tests for richer excitation.
    c. Apply anti-aliasing filtering, detrending, removing offsets.
2. System identification:
    a. Select a transfer function or state-space model.
    b. Estimate time constant, gain, and delay if present.
3. Controller parameter computation:
    a. Apply tuning method (e.g., Ziegler–Nichols, optimization-based, Adaptive/gain-scheduled PI…).
    b. Incorporate actuator limits and anti-windup logic.
4. Validation and deployment:
    a. Offline simulate closed-loop response.
    b. Deploy tuned gains to substation controller.

Extensive prior experience in thermal control will be applied here, ensuring robust algorithms that handle nonlinearities, noise, and actuator constraints. Proven approaches such as anti-windup mechanisms, setpoint weighting and anti-aliasing filter integration will be adapted for substation control.

A more integrated and preferred strategy for optimising network efficiency is to migrate the entire control logic from local substation controllers to the Arteria platform. In this configuration, control algorithms are executed centrally, providing a unified control layer that governs all substations in conjunction with the power plant. Centralised control provides benefits such are:

- Dynamic coordination with power plant operations: Substation control dynamics can be matched to the central plants response characteristics, enabling overall optimisation of production.
- Return temperature minimisation: By directly coordinating substation valves with system-wide objectives, return temperatures can be minimised, improving the efficiency of the network.
- Simplified tuning and maintenance: Controller parameters can be updated centrally. This also enables faster deployment of improvements and eliminates inconsistencies between substations.
- Scalability and uniform structure: A single control architecture can manage hundreds of substations with uniform strategies which significantly reduce development time.

The proposed control law for this centralised approach is:

$u(k) = u(k-1) + w_1*(diff(k) - diff(k-1)) + w_2*diff(k)$

where: u(k) is the manipulated variable, in this case primary valve position. diff (k) is an error signal combining two control objectives:

- to drive the secondary side supply temperature Tout_S toward its setpoint.
- to minimise the primary-side return temperature Tout_P to its lowest feasible value.

The weights w_1 and w_2 are determined via data-driven optimization, following methodologies similar to those successfully applied in power plant control design. In this application, two approaches are available for obtaining the dynamic substation models required for tuning.

- Model identification from historical data: Utilising recorded data, a suitable simplified dynamic models of substations can be derived.
- System identification with external excitation: When feasible, low-amplitude PRBS (Pseudo-Random Binary Sequence) signals or step tests can be applied to the valve input to enrich the excitation spectrum and obtain more accurate models. This method is especially valuable for substations with insufficient natural variation in historical data.

Once models are available, model-based tuning can be performed to determine the optimal weights w_1 and w_2. The flexibility of this approach is in the ability to incorporate control objectives directly into the cost function. This enables a balance between secondary supply temperature tracking and minimizing primary return temperature.
Because the controller structure is highly flexible, this framework can achieve performance levels significantly higher than those of traditional substation controllers. The optimization strategy naturally supports adaptive re-tuning, where the controllers parameters can be updated as soon as new operational condition become available, therefore maintaining high performance over time.
To ensure robust and high quality control performance the anti-integral windup mechanism, anti-aliasing filters, valve nonlinearity compensation and actuator constraints integration will also be incorporated in the control algorithm.
During the development of the optimization algorithm, Arteria implemented a detailed simulation environment to test various control strategies before applying them to the live system. These simulations involved modelling thermal transients, heat propagation and the effects of changing plant supply temperatures on downstream consumers. The algorithm underwent extensive iterative testing, incorporating anti-aliasing filters to prevent measurement noise from destabilizing the control system. This approach was necessary because real-world sensors often report slightly fluctuating values due to environmental factors, sensor precision or communication jitter. Without filtering, these fluctuations could cause the algorithm to issue unnecessary or even harmful adjustments to the plant.

In addition to real-time optimization, Arteria’s digital twin environment allowed for post-processing and retrospective analysis of specific optimization periods. The platform provides tools for replaying historic operational data within the digital model, enabling the team to analyze how the system behaved during specific events. For example, on days with sudden weather changes or unexpected consumer demand spikes, the algorithm demonstrated resilience by adjusting supply temperatures while maintaining system stability. These analyses further validated the robustness of the entire optimization framework.

The network operators expressed positive feedback on the clarity and usability of the redesigned dashboards, highlighting the improved visibility of network performance indicators and the intuitive layout enabling efficient decision-making. The platform’s ability to show real-time sensor readings alongside simulation-based predictions helped operators understand why the algorithm recommended specific control interventions.

Overall, the coordinated efforts across WP2 and WP4 resulted in a comprehensive, integrated technical solution combining detailed network modelling, real-time data streaming, secondary-side data integration and dynamic optimization algorithms. The successful development and implementation of these technologies mark a significant step toward a fully automated and decentralized district heating optimization platform. As the DOPPLER project continues, these foundations will enable further enhancements such as more advanced predictive control, automated substation controller tuning and expanded integration with smart home devices, providing even higher efficiency and flexibility for district heating operators.

3.1.5.3 Considering flexibility provided by consumers in supervisory control

Usually when optimizing the supply side of a heating network, the heating plant is modelled while the supply network is omitted. The consumption of the individual consumers manifests itself as an aggregated power demand at the feed-in point. This simplification, however, does not allow to take the full flexibility of the network into account. Actively influencing end-user behaviour (demand-side management, DSM) is therefore gaining importance since it can reduce peak loads. In the context of fluctuating renewable generation and variable (electricity) prices, DSM can also be used more generally as an additional flexibility measure to deliberately shift consumption in time. However, this requires an appropriate representation of consumers within the optimization as well as a realistic way to influence substation behaviour without negatively impacting consumer comfort. This can, e.g., be achieved by connecting the optimization directly with the smart home systems.

During the project duration, the MILP-based energy management system for controlling the heat supply central presented in 3.1.5.1 was extended to be able to consider and recommend changes in the operating conditions of the connected consumers. The proposed workflow is as follows:
Step 1: Segmentation of consumers into actively controllable and passive consumers. Passive consumers represent critical infrastructure or customers that do not wish to participate in DSM measures. From a technical perspective, they can also encompass those whose demand profiles are difficult to model.
Step 2: For the actively controllable consumers, individual, simple dynamic models need to be formulated and parameterized using measurement data, enabling prediction of local demand as a function of weather forecasts and power limits. Based on simulations by, e.g., the Arteria framework, network losses are estimated and added to the individual demand data in order to estimate the expected power at the feed-in point. The demand of passive participants is forecast using standard multilinear regression (MLR) on various features (outside temperature, time of day, working day or public holiday), based on historical feed-in power and corrected for the influence of controllable consumers. This approach makes it possible to approximately factor out the effect of active load shifting from the load profile and to estimate the future total demand from the heating plant's perspective *without* DSM measures.
Step 3: Modelling DSM measures as a *virtual storage* unit. Once per day, a "coordinator" queries the load-shifting potential of individual consumers by sending a reference list of load-shifting curves to "agents" representing the customers. Using the dynamic mathematical models of step 2, these agents

compute to what extent a shift according to the request can be achieved by varying input variables (e.g., the setpoint temperature or a fictive ambient temperature that is communicated to the local heating controller) without, for example, reducing comfort or increasing overall energy consumption. The achieved shifting curves are reported back, aggregated across all consumers, and used to derive time-varying charging-power and energy constraints for a virtual storage unit. This virtual storage is then available to the optimization analogously to a conventional thermal storage in a day-ahead optimization, with the state of charge required to be zero at the beginning and end of the day (no increase or reduction of the daily energy demand) and allowed to become negative during the day (a temporary reduction in consumption corresponds to discharging the storage).
Step 4: Coordination of the desired load shift. After the day-ahead optimization, the charging and discharging power of the virtual storage is interpreted as a desired load-shifting curve and is distributed by the “coordinator” back to the agents. The consumers must then follow the committed trajectories as closely as possible.
This developed agent-based algorithm is more scalable than actively considering each consumer in the overall optimization problem; however, it depends on the consumers closely following the target trajectories supplied by the centralized controller, and future work is required to better handle uncertainties and model errors that lead to a different behaviour than expected.

3.1.6 End-User Integration

3.1.6.1 Dashboards

Purpose and Objectives of the Dashboards
The dashboards developed in the DOPPLER project are designed to visualize and evaluate the efficiency of district heating networks from both the operator’s and the end customer’s perspective. The goal is to present complex energy relationships in a clear and transparent way to promote data-driven decisions and behavioral adjustments. Two core aspects take center stage: operational optimization and user engagement through gamification.
For operators, the dashboard enables precise analysis of network operations. Continuously recorded measurements from substations are used to evaluate thermal outputs, volume flow rates, and temperature spreads. Using the metric m³/MWh, the hydraulic efficiency of individual transfer stations can be compared. This key figure represents the ratio of the consumed volume flow to the delivered energy and allows assessment of network and consumption efficiency independent of absolute power levels. High efficiency (a low m³/MWh value) indicates a large temperature spread and low return temperatures, whereas high values point to optimization potential. This makes it possible to identify inefficient consumers or network sections and initiate technical adjustments (e.g., hydraulic balancing, control optimization). For the operator, the result is a tool for data-driven efficiency improvement and monitoring of network quality. For end customers, the focus is on participation and motivation. The principle of gamification is used to encourage energy-efficient behavior in a playful way: color coding, leaderboards, and efficiency scores visualize one’s status in comparison with other consumers. As a result, end customers become active participants in the district heating system, which in the long term

leads to lower return temperatures, better utilization of generation assets, and thus a reduction in overall losses. Combining both perspectives, operator and end customer, creates a shared information space. While operators derive operational decisions and optimization measures from the dashboards, end customers receive feedback on how their individual behavior affects the overall system. The dashboards therefore fulfill a dual function. They serve as a planning and control tool for operators and as a motivation and learning platform for end users. The objective of the dashboard component within the project is therefore to foster a shared digital platform that unites transparency, efficiency awareness, and participation. By visually preparing the key metrics, technical data are translated into understandable information an essential step to ensure acceptance and the long-term success of digital optimization strategies in the district heating sector.

Operator Dashboard

- Purpose. Operational monitoring and efficiency assessment of the network at the level of individual transfer stations (KIDs) and the overall network. Focus on consumption efficiency (m³/MWh), leakage indicators, loss shares, and temporal patterns of volume flow and thermal power.
- Global elements. Time-range selection (rolling and fixed windows), KID filter (including exclusion of individual KIDs), consistent colour scales (green = efficient/low m³/MWh, red = inefficient/high m³/MWh).
- Panels (briefly described):

Heat quantities and losses (network level)

This panel monitors district heating performance at network level by comparing sold heat and distribution losses. Daily values are stacked to show delivered energy versus loss energy, while a secondary axis displays the loss share (%) as a line. This enables seasonal interpretation (summer/winter) and makes optimization effects visible over time (Figure 26).

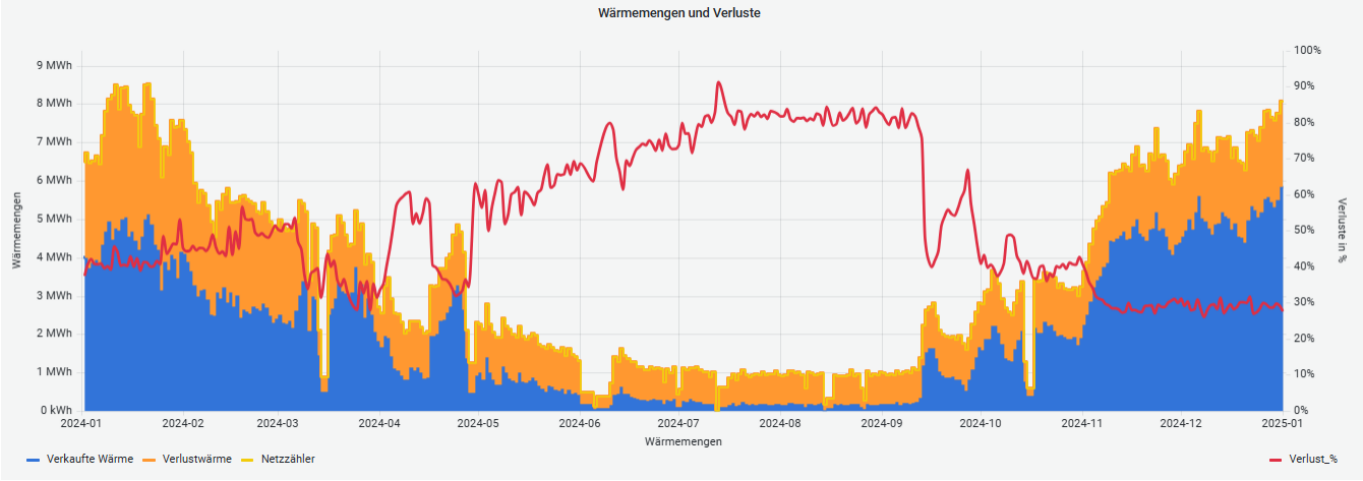

Figure 26: Line graph displaying heat produced and heat sold with losses calculated as a percentage.

Heatmap m³/MWh over the year

This heatmap visualizes volumetric heat offtake intensity (m³/MWh) per customer interface (KID) over time (Figure 27). Each row represents one delivery point, with color thresholds highlighting efficient versus inefficient operating states. It supports pattern detection across the year (e.g., commissioning, shutdowns, high return temperatures). Grey indicates no data/no load. Dark grey flags volume flow without thermal power.

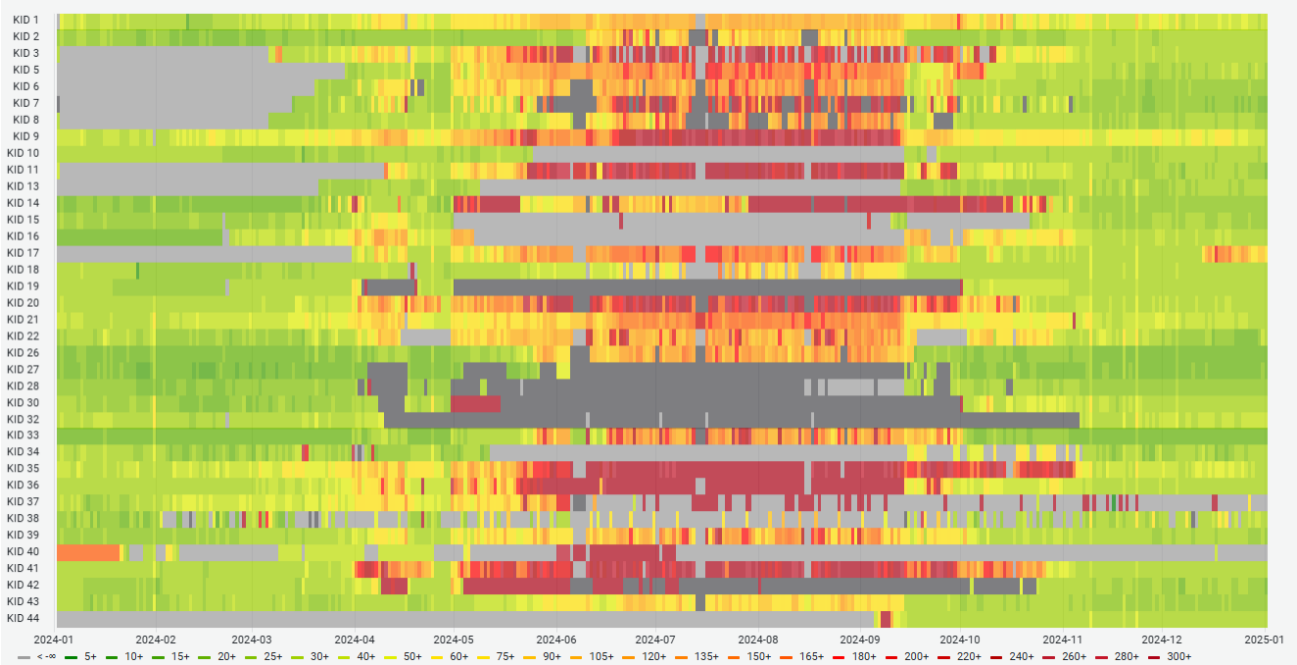

Figure 27: Heatmap displaying the colour-coded specific volume per energy

Heat generators

These two panels provide an operational view of heat generation. The time series tracks hourly thermal power from the biomass boiler, oil boiler, and the network meter to identify load coverage and peak events. The accompanying pie chart summarizes the heat contribution of each generator over the last 14 days (Figure 28).

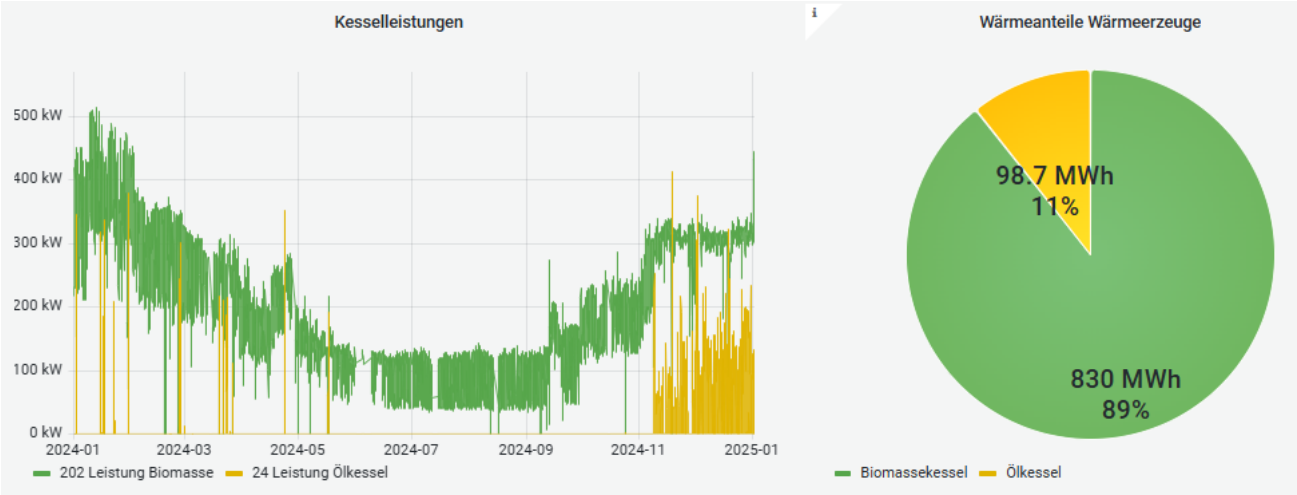

Figure 28: Visualization of boiler outputs and energy quantities

Treemap

This treemap (Figure 29) provides a compact portfolio view of all district heating delivery points (KIDs). Each tile represents one KID, tile area scales with total energy in the selected period, while tile color encodes volumetric intensity (m³/MWh). This highlights both high-impact consumers and inefficient operating patterns at a glance.

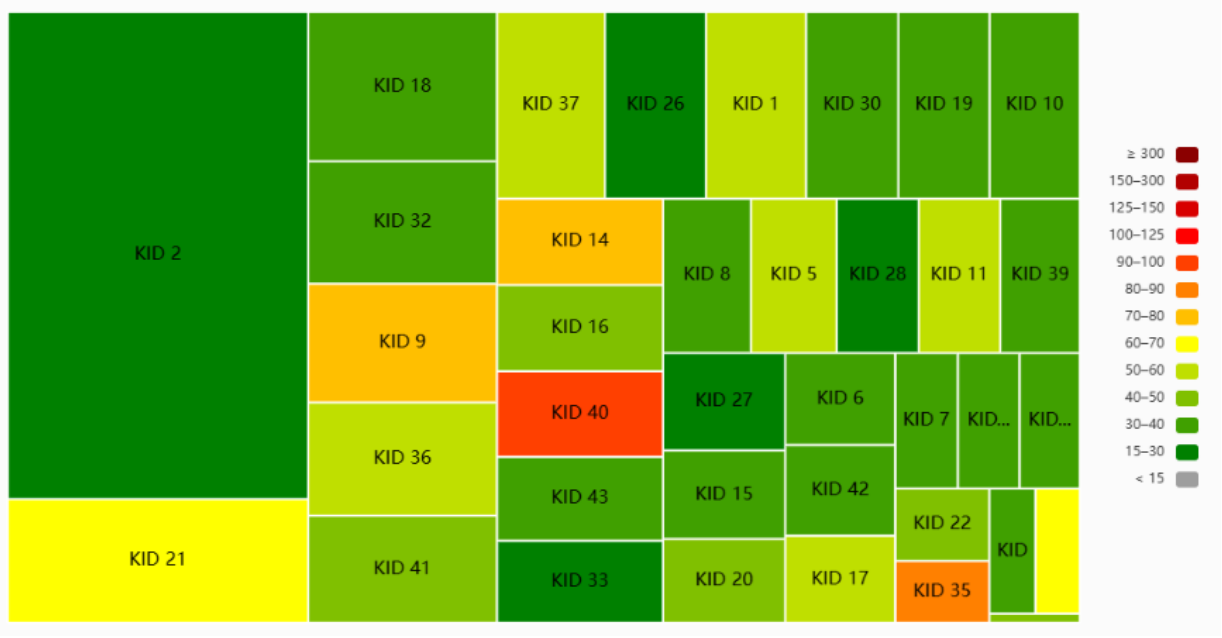

Figure 29: Treemap for a compact overview

Leak volume for the selected time range

This bar-gauge panel quantifies “leak volume” for the selected time step, aggregating volume flow during operating states where thermal power is approximately zero while flow remains positive. Values are computed per KID and summarized as total m³, enabling rapid identification of abnormal circulation (e.g., leaking valves) flagged as dark-grey states in the heatmap (Figure 30).

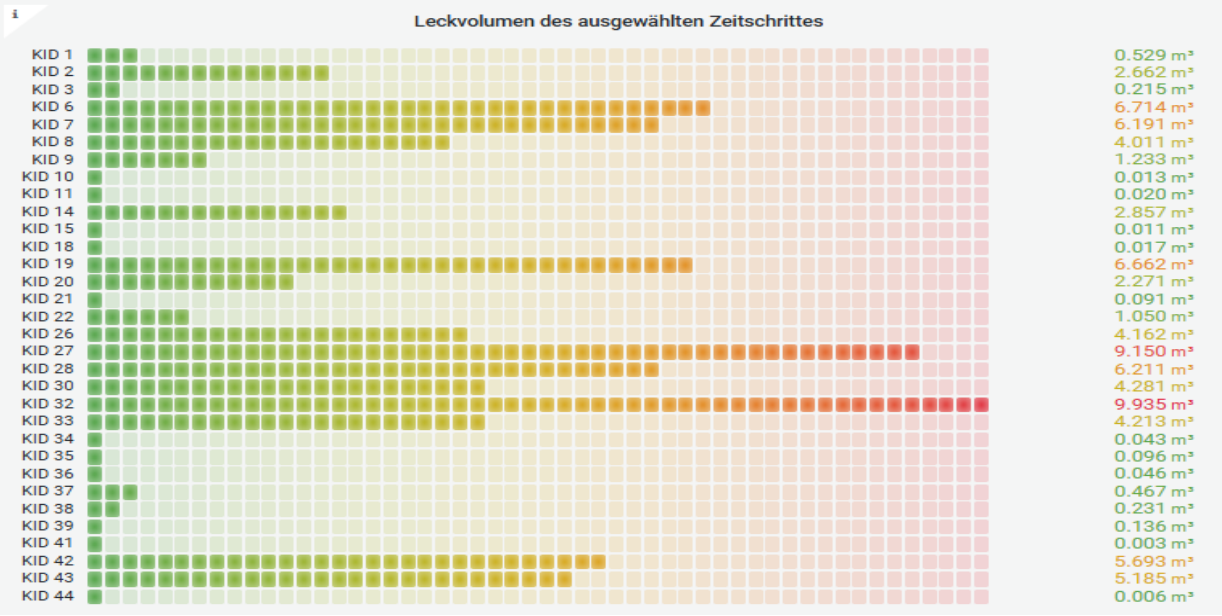

Figure 30: The leak-volume dashboard shows leaking valves

Detailed analysis of flow (l/h) and thermal power (kW)

These two detail panels (Figure 31) show hourly volume flow (l/h) and thermal power (kW) per delivery point (KID) for a selected period. By comparing both signals, they help identify load peaks, atypical baseloads, and behavior-driven patterns. This supports root-cause analysis and validation of efficiency indicators such as m³/MWh and leak-volume detection.

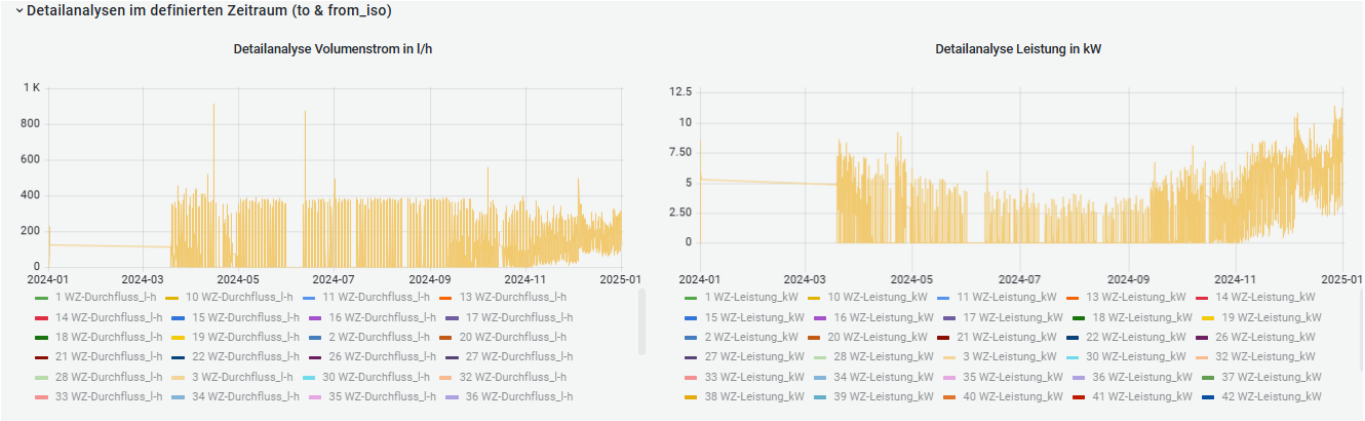

Figure 31: Analysis dashboard of power and flow

Operational benefits

The combination of matrix benchmarking, leak-volume mosaic, time series, and annual heatmap enables: targeted identification of inefficient consumers/sections, prioritization of measures (hydraulic balancing, control optimization), tracking of loss development, and effectiveness monitoring of interventions over defined reference periods.

End-Customer Dashboard

Purpose: Clear feedback on one's own consumption and efficiency behavior: promoting efficient use through transparent KPIs, reference values, and gamification elements.

Design principle: Reduction to a few meaningful KPIs and intuitive visualizations (traffic-light logic, gauge/score, ranking). Identical definitions to the operator dashboard, but with a simplified presentation.

Stat charts

This panel group (Figure 32) summarizes the selected customer station (KID) with key totals for energy (kWh), volume (m³), and temperature spread ΔT (K), including a derived ΔT based on energy and volume. Three supporting time series show daily energy use, daily volume, and hourly ΔT, enabling customers to interpret usage patterns (e.g., weekends, holidays, weather).

Figure 32: Overview panel for energy, volume and temperature difference.

Efficiency band (m³/MWh-classification)

This panel shows (Figure 33) the customer's offtake efficiency over time using the volumetric intensity metric m³/MWh (volume required to transfer 1 MWh of heat). Daily values are classified into color bands, where green indicates efficient operation and red indicates inefficient operation. Grey represents no load/no data, while black flags flow without heat transfer (power ≈ 0).

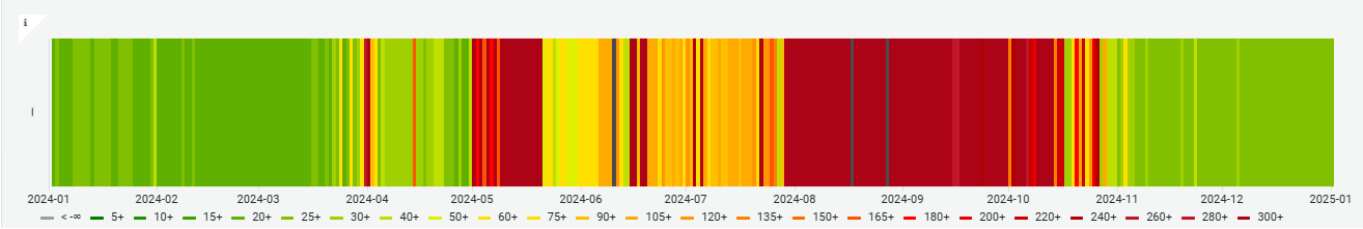

Figure 33: Efficiency band displaying current and past volume per energy

Comparison displays – last 24h, 30 days, and selected period

Three KPI gauges (Figure 34) benchmark the customer's offtake efficiency (m³/MWh). The first shows the customer's value for the selected period. The second displays the network-wide average as a reference. The third computes the customer's rank (1 = best) among all participants. Together they enable simple benchmarking and gamification. Replicated for last 24h, last 30d, and custom range.

Figure 34: Efficiency gauges for current volume per energy and ranking within the district heating

Detailed analyses of flow, power and temperatures

Linked time series of consumer station (Figure 35): left panel shows thermal power and, if available, volume flow, right panel displays primary and secondary supply/return temperatures. Together they help you interpret unusual periods (e.g. high return temperatures, low load, cycling) without needing raw data exports.

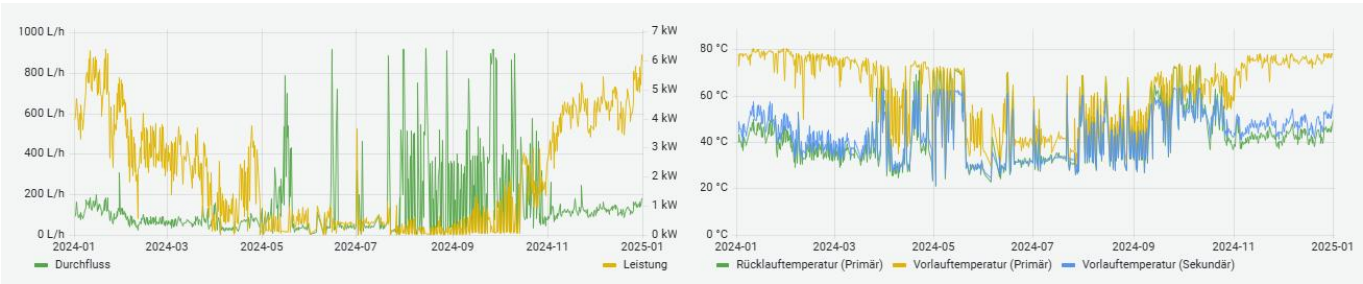

Figure 35: Detailed analysis panel for flow, power and temperatures

Quantification of improvements

The panel quantifies the year-over-year improvement by back-calculating the saved energy from the change in the m³/MWh ratio (Figure 36). The resulting savings are reported in kWh and additionally converted into an equivalent CO_2 reduction, which is displayed as a key metric in the panel.

Figure 36: Panel for quantification of improvements

Use of end-customer

The interface enables users without any specialist knowledge to see their own position relative to the network, understand how adjustments (heating curve, night setback, hydraulics) affect efficiency, and stay motivated by visible progress (score/ranking) over defined reference periods.

Methodology

The dashboards were developed in Grafana v8.3.3, an open-source visualization and monitoring platform for time-series data that allows analysts to interactively define arbitrary time ranges and adjust the temporal resolution (aggregation/granularity) of plots and KPIs depending on the analysis objective. An InfluxDB 2 backend was used to store and provide the measurements. Grafana retrieves the data via Flux queries and renders them into the corresponding visualizations. The underlying dataset comprises the 2024 measurement records of the Mischendorf district heating network and includes time series from the transfer stations (KIDs), specifically thermal power, volume flow, and supply/return temperatures. For preprocessing, fixed plausibility limits were applied in Flux (e.g. upper bounds for power and flow), and zero/invalid values were discarded. Energy totals were obtained via time-discrete integration of power values $E_{\mathrm{kWh}} = \sum P_{\mathrm{kW}} \cdot \Delta t_{\mathrm{h}}$(logging interval 5 min), Volumes were computed analogously as $V_{\mathrm{m^3}} = \sum \dot{V}_{\mathrm{l/h}} \cdot \Delta t_{\mathrm{h}}/1000$. The m³/MWh KPI was calculated as $V_{\mathrm{m^3}}/E_{\mathrm{MWh}}$. The temperature spread was derived from volume and thermal energy.

Challenges

Key challenges included crafting robust Flux queries, heterogeneous data quality (outliers, inactive periods, differing time coverage), and translating engineering metrics into easy-to-grasp displays for end users. The complexity of the subject required reducing to a few unambiguous KPIs and a visual language with clear traffic-light logic without sacrificing explanatory power.

Quantifying improvements

Effects of measures are evaluated relative to a defined reference period. For each station and for the overall network, changes in the core KPIs, $\Delta(\mathrm{m^3/MWh})$, are calculated. Similarly, the influence of transfer stations on the network return temperature can be assessed, allowing network losses to be recalculated.

3.1.6.2 Motivation

When dealing with the problem of how to motivate district heating (DH) customers to change their use patterns in a more flexible way enabling a more efficient operation of DH systems at least two groups of users have to be distinguished: households and manufacturing.
Additionally, further groups might be distinguished such as public organizations (e.g. schools and hospitals) and service companies (e.g. restaurants and shops). Basically their use of DH is very similar to households but differ in size, sometimes significantly. Furthermore, there are also very different decision processes to be considered. Nevertheless, the more important distinction is between households and production. Accordingly, these two groups of DH end-users have been analyzed in DOPPLER.

Households

A survey of private DH end-users (i.e., households) in Güssing (one of the demonstration DH-sites) showed customers' willingness to be more flexible in their DH use. The majority expressed a positive attitude towards flexibility: 58% said they would be ready to use DH more flexible concerning hot water preparation, 68% concerning room heating. In general, flexibility is accepted. But there are interesting details to be considered (see Figure 37).

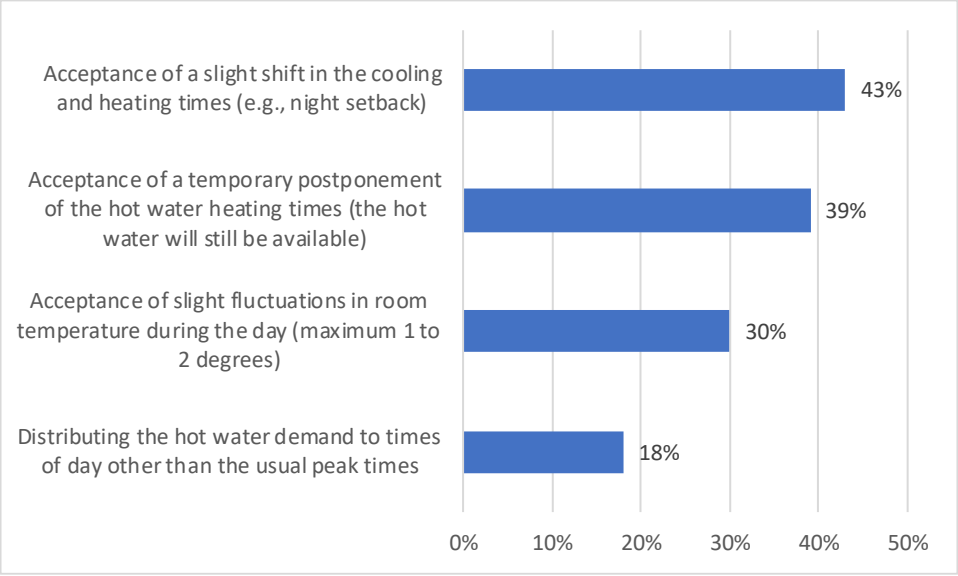

Figure 37: Acceptance of specific forms of DH-flexibility, source: DOPPLER survey of DH customers in Güssing, Austria

The willingness to shift hot water consumption is particularly low. It is evident that the associated need to plan or spontaneously postpone hot water use is rarely accepted. Shifting the charging times for hot water preparation fares significantly better. This is hardly surprising, as it does not, or only rarely, mean that hot water is unavailable for use. Regarding room heating, the option that does not involve any immediate restriction on the DH use was mentioned significantly more often. Shifting the setback and heating times is accepted much more frequently than a varying room temperature, even if the fluctuation is only slight. Whenever there is a marked restriction of convenience the willingness to be flexible is decreasing.

In addition to questions about their willingness to use district heating more flexibly, households were also asked whether they would be willing to take certain measures that would allow the operator to make the DH supply more efficient. The most widely accepted measure is an individualized automatic activation (i.e., setting according to your actual individual needs, deviating from the standard setting) with 59% approval, followed by an automatic limiting control (i.e., activation to avoid peak DH demand in the house) with 36%. Least accepted is the reduction in DH consumption following notification (i.e., voluntary reduction based on information from the district heating provider, e.g., via SMS) with only 21% approval. Efficiency measures that can be perceived as external interference are less often accepted.

For motivating households to use DH in a more flexible way elements of gaming have been included in the end-user dashboard. Gamification is the application of game-like elements and mechanics in non-game contexts to increase motivation and engagement and to bring about behavioral changes. The aim is to encourage real-world behaviors and changes through playful means. This includes elements such as points, badges, leaderboards, and progress indicators, which serve to make goals achievable and motivate participants to engage in a task or process. Gamification works by using typical game elements and methods. Elements commonly used in games include points, medals, leaderboards, progress indicators, or achieving levels that make increasing success visible.

Two gaming elements were integrated into the end-user dashboard: ‚Ranking' and ‚Status/Progress indicator'. The ranking is derived from a comparison of the consumption efficiencies of all DH customers. ‚Consumption efficiency' refers to the ratio of the amount of water required to the energy consumed. Green colors indicate a favorable (efficient) ratio, while red indicates an unfavorable one. A household's rank is determined by the ratio of the end-user's average consumption efficiency during the selected period to the average efficiency of the entire district heating network. The rank of each end-user is illustrated with corresponding colors, ranging from ‚green' (i.e., high efficiency) to ‚red' (i.e., low efficiency).

The first gaming element, ‚Ranking', reflects an end-user's status in comparison to other end-users. To make dynamic changes and their effects on the entire district heating network understandable, a second gaming element was designed that demonstrates the benefits for the entire network. The goal is to provide an incentive to increase the efficiency of DH use and to make the effects of such individual improvements comprehensible.

To this end, it is first necessary to improve consumption efficiency by implementing specific efficiency-enhancing measures (e.g., replacing a malfunctioning valve, gradually increasing the heating temperature). This can be achieved retrospectively by comparing the consumption efficiency after implementing the measure with that before. However, in order to incentivize the implementation of an efficiency-enhancing measure, it is more important to illustrate the prospective effects of such a measure. For this purpose, the consumption efficiency of a household that has already implemented the corresponding measure can be used as a reference.

However, the problem arises that the individual contribution of a single household to most district heating networks is extremely small and hardly noticeable or quantifiable. Exceptions are large commercial DH customers, who, individually, already make a significant contribution to optimizing the DH network. This is almost never the case in the residential sector. Therefore, an extrapolation method for the entire network must first be developed. A very simple approach is to assume that the efficiency measure in question, implemented by the end-user under consideration, is also adopted by all other end-users who exhibit the same efficiency deficit. The cumulative effect would then be clearly noticeable in the DH network.

Once such a projected overall network effect has been determined, the challenge is to present it in a way that is understandable to non-district heating experts. To clearly illustrate this progress, a memorable symbol was chosen – the tree. The tree symbol visualizes two positive effects for the district heating network:
1) Fuel savings: Increased efficiency means that the same required district heating output (energy delivered) can be provided with less fuel input (energy consumed). In other words, conversion and transmission losses are reduced.
2) Cost savings: The reduced amount of fuel required results in savings on procurement costs. These freed-up funds can be used elsewhere. In the case of the demonstration DH network in the small village of Japons in Austria, planting trees was considered, which would also help to sequester CO_2.

The more trees are displayed, the greater the efficiency gain achieved by the measure under consideration in the overall district heating network. Each tree represents the amount of CO_2 saved or sequestered as a result of the implemented efficiency-enhancing measures. This effect represents the collective benefit.

As part of the project, this gaming element was conceptually outlined but not yet implemented in the dashboard. The involvement of the respective DH operator is essential for implementing this gaming element. Only in this way can the collective benefits be reliably defined, the necessary information and assumptions specified, and a practical form of implementation found for the DH network and its end-users. Consequently, such a gaming element must be designed not generically, but specifically for each operator and network.

The gaming elements in the dashboard were assessed by the end-users at the final DOPPLER-workshop in Japons in a differentiated way, ‚Ranking' received a positive valuation, ‚Status/Progress indicator' a more sceptical one.

The information about their own ranking in terms of consumption efficiency was considered useful by some workshop participants and was rejected by nobody. The ranking information can therefore serve as an incentive for at least some end-users to improve their own position through efficiency measures (in addition to the aspect of cost savings). This information in the dashboard is considered generally useful and an incentive to improve one's own efficiency in using district heating. However, it has to be considered that the level of expertise of an end-user matters a lot. The more advanced end-users are,

the lower is their interest in gaming elements. That became clear at the first workshop in Güssing, where the ranking element had first been presented to more experienced DH customers. Most of them considered the ranking information irrelevant to them. They are more directly looking at technical data. Less advanced end-users tend to appreciate gaming elements more.

The collective benefit of implementing individual measures to improve the efficiency of DH use was not sufficiently recognized. The tree symbolism of the second gaming element is inadequate for understanding the causal relationship. The fact that the collective benefit for the DH network also lies in the individual interest of each end user was indeed acknowledged. However, it was criticized that the connections were not adequately explained. The „story“ is missing. Without a deeper understanding of the collective benefit, the display of corresponding symbols (regardless of their type) will not be seen as an incentive for behavioral change.

Manufacturing
Based on expert interviews in four manufacturing companies in Güssing – two cases in the wood industry, one case each in the food and metal industries, it became clear that the production factor ‚heat energy‘ is not important enough to motivate the companies to engage in substantial changes. The share of district heating (DH) costs in total production costs is too low to justify investments in equipment or the adaptation of processes that would be necessary to make district heating demand more flexible. To enable a significant increase in the flexibility of DH demand, production processes must first be made correspondingly more flexible. This therefore involves the flexible automation and digital control of manufacturing. The more rigidly automated the production processes are, the greater the investment required. And the greater the investment required, the less significant the cost savings in the area of DH costs become. In the specific case of the industrial companies examined in Güssing, the low to non-existent degree of flexible automation therefore represents a significant barrier to increasing the flexibility of DH demand. While there are some opportunities to make the use of district heating for production processes more flexible in the current operations, these are relatively insignificant and some are already being implemented anyway. A significant increase in flexibility is therefore only possible with substantial investments in new manufacturing processes. However, such a willingness to invest is currently not foreseeable in the case of the companies investigated in Güssing.

3.2 Business Models

3.2.1 Core Concept: Common-Good Economy Foundation

The development of a sustainable and scalable business model for district heating systems requires a fundamental shift in economic thinking. The DOPPLER project has developed an innovative common-good economy framework that extends beyond traditional profit-maximization approaches to incorporate shared value creation for all stakeholders.
The common-good economy model is grounded in constitutional and foundational values that enable successful relationships: trust-building, appreciation, cooperation, solidarity, and sharing. Rather than measuring economic success through financial metrics alone, this approach evaluates achievement

through common-good objectives: need satisfaction, life quality, and community well-being. This represents a paradigm shift from competition-based to cooperation-based economic incentives, where economic success is measured not by the means of economic activity (money, capital, financial profit) but by the goals achieved (need fulfillment, quality of life, common welfare).
This framework directly addresses the challenge of optimizing district heating utilization through flexible usage mechanisms while simultaneously reducing CO_2 emissions and achieving cost efficiency. The integration of incentives with community-oriented thinking enables intelligent control of heat consumption, reduction of peak loads, cost reduction, and improved environmental outcomes.

3.2.2 End-User Perspectives and Flexibility Readiness

To inform the development of this common-good business model, the DOPPLER project conducted comprehensive end-user research involving 106 respondents from the Güssing and Mischendorf district heating networks. Survey results demonstrated significant potential for demand-side management: Space heating flexibility was indicated by 68.1% of respondents, while 57.7% expressed willingness to adopt flexible hot water practices. Specifically, 42.6% of respondents accepted slight shifts in heating and cooling times through night setback adjustments, and 29.8% tolerated minor room temperature variations of ±1–2°C. Temperature variation acceptance was notably higher in detached houses compared to apartment buildings.
Regarding financial incentives, variable tariffs based on system load were preferred by 59.8% of respondents as mechanisms to encourage flexible consumption. Free consulting services were preferred by 41.3% of respondents, indicating strong support for technical guidance to optimize system performance.
These findings demonstrate that end-users constitute a willing and capable stakeholder group for participation in district heating optimization programs. The research revealed that flexibility without visible comfort impact is substantially more acceptable to users, providing critical guidance for implementation strategy.

3.2.3 Integrated Platform Architecture and Technical Integration

The business model integrates three core technical systems to enable seamless optimization:
The Digital Twin platform provides system-wide simulation and scenario analysis capabilities. The Arteria simulation platform enables detailed forecasting and flexibility potential identification. Integration with Smart-Home Energy systems and the MEO Energy Controller ensures real-time monitoring and control coordination. These technical components form the backbone enabling transparency and performance tracking for all stakeholder groups.
Three distinct dashboard implementations address different user requirements: Household dashboards provide simplified interfaces with consumption visualizations and gamification elements; building manager dashboards enable tracking of multiple units with tenant analytics and anomaly detection; operator dashboards provide comprehensive system-wide analyses with advanced forecasting capabilities.

3.2.4 Flexibility Mechanisms and User Engagement

The common-good business model supports three primary flexibility mechanisms:
Load shifting through targeted control of heat requirements enables users to adjust consumption patterns in response to system signals and pricing incentives. Room temperature reduction during non-use periods achieves efficiency gains without requiring behavior change during occupied times.
Individual time-window selection for better consumption adaptation provides users with autonomy over when and how they participate in optimization.
Users can choose their level of participation based on personal comfort preferences, ensuring that optimization objectives do not compromise quality of life. This flexibility is essential for sustainable user engagement and long-term behavior change.
Beyond technical flexibility, the model incorporates personalized energy consulting services as standard offering to participants. These services support process optimization, equipment improvements, and behavioral adaptation, enabling long-term efficiency gains beyond immediate technical interventions.

3.2.5 Gamification and Community Engagement Strategy

The DOPPLER project developed a comprehensive gamification system specifically designed to motivate end-user participation without compromising autonomy or comfort. The system recognizes that different user segments possess different psychological motivations and respond to different engagement mechanisms:
The environmentally conscious segment is engaged through gamification elements highlighting ecological impacts of energy-saving efforts. This includes carbon badges in various tiers, environmental protection goal tracking, and opportunity to convert saved CO_2 emissions into concrete environmental impacts. Environmentally conscious users can optionally donate portions of energy cost savings to environmental protection projects, creating deep emotional connection to conservation actions.
The cost-conscious segment is addressed through gamification elements placing financial savings front and center. This includes visual savings trackers showing cumulative financial savings, financial milestones with tangible rewards such as vouchers for local businesses or energy credits, and special challenges offering monetary incentives for participation.
The tech-enthusiast segment is engaged through advanced functionality including granular data access, JSON/CSV export capabilities, API integration options, and access to advanced analysis features enabling deeper energy consumption insights.
A comprehensive points system allows users to earn one point per kWh saved below individualized baseline values. Points can be redeemed for segment-appropriate rewards. Social comparison and community elements are implemented with complete user control and privacy protection, with rankings and competitions entirely optional and restricted to geographically or community-limited groups.

3.2.6 Impact Visualization: Making Environmental and Financial Benefits Tangible

Effective behavior change requires transparent visualization of individual and collective impacts. The DOPPLER project developed multiple complementary visualization perspectives:

Consumption trend tracking enables users to understand daily, weekly, monthly, and annual consumption patterns with 50% historical transparency for long-term comparison. Cost breakdown visualization displays heating costs, hot water costs, and various charges separately, enabling optimization of specific cost components.
The CO_2 tracker converts energy usage into environmental impacts by showing cumulative CO_2 emissions and comparing these with baselines of similar households. Importantly, the system converts abstract metrics into concrete user-understandable comparisons for example, 100 kg CO_2 equals approximately 400 km of automobile driving.
The tree metaphor provides intuitive visualization of collective impact. One tree stores approximately 24.62 kg CO_2 annually. When a household saves 100 kg CO_2, this is displayed as 4 planted trees. Dual tree visualization shows both forest preservation through reduced fuel consumption and new afforestation through cost savings allocation.
Scenario analysis enables users to understand how different consumption pattern changes would affect energy bills and environmental impacts. Users can create and compare scenarios: "What would my monthly bill cost if I lowered room temperature by 1 degree?" or "How much could my emissions be reduced with a new high-efficiency heating system?"

3.2.7 Stakeholder Value Creation and Revenue Model

The common-good economy model creates distinct value propositions for each stakeholder group: District Heating Operators achieve more efficient use of existing infrastructure and substantial reduction of production costs through optimized load management and improved system reliability. The Digital Twin enables data-driven operational decision-making and proactive maintenance planning. Operators benefit from reduced peak load pressures and improved network stability.
End-users reduce energy costs through flexible participation in price-incentive programs while maintaining comfort and autonomy. Gamification elements and environmental impact visualization provide additional motivation for participation beyond financial savings. Users experience improved system transparency and understand concrete impacts of their consumption patterns.
Society achieves lower CO_2 emissions and more sustainable energy supply through collective conservation efforts. The business model can support energy poverty programs by allocating portions of collective savings to vulnerable populations, creating social value alongside environmental benefits. This alignment of individual and collective interests is fundamental to the common-good approach.
The business model generates sustainable revenue through multiple streams: platform licensing fees for district heating operators, service and consulting fees for implementation and optimization support, and data analytics services providing benchmarking and predictive modelling for energy service providers. This diversified revenue model ensures financial sustainability while maintaining alignment with common-good principles.

3.2.8 Implementation Strategy and Risk Management

Pilot implementation across three Austrian district heating systems Güssing, Mischendorf, and Japons demonstrates system effectiveness and validates business case assumptions. These demonstration sites provide evidence for scaling strategy and market entry approaches.
A collaborative stakeholder network including district heating operators, end-users, and energy experts supports platform development and identifies new application opportunities. This partnership-based approach ensures market relevance and user acceptance throughout implementation.
Risk management addresses several key implementation challenges: Aging infrastructure with low-efficiency boilers is addressed through targeted consulting programs and technical support. Market volatility in energy costs is mitigated through long-term energy supply contracts and diversification of energy sources. Customer flexibility constraints, particularly for large industrial customers, are addressed through customized solutions including flexible contract models and heat storage solutions. Infrastructure evolution is ensured through evaluation and development of new district heating systems at optimal locations, integrating renewable energy sources including geothermal, solar thermal, and industrial waste heat recovery.

3.2.9 Validation and End-User Feedback

End-user feedback from workshop presentations at Japons (October 2025) involving 22 district heating customers provided critical validation of the business model approach:
Ranking information indicating individual efficiency position relative to network average was perceived as useful by workshop participants and was viewed as potential incentive for behavior improvement. This element was recognized as fundamentally sound and motivating.
Visualization of collective network benefits through tree symbols was recognized as requiring stronger contextual explanation. While users acknowledged that collective benefits align with individual interests, they emphasized that relationships and cause-effect connections require clearer narrative explanation to function as effective incentives. This feedback indicates that the business model's theoretical foundations are sound, but implementation requires substantial attention to narrative explanation and user education regarding how individual actions create network-level benefits.

4 Results and conclusions

4.1 Integrated Digital Twin as a Technical Backbone

A central result of DOPPLER is the successful implementation of operational digital twins for multiple district heating networks and their positioning as the technical backbone of an integrated optimization approach. Digital network representations were created for Mischendorf, Rohrbach, Japons and Güssing, combining georeferenced hydraulic structures with metadata and time-series data interfaces. This provided a consistent environment to simulate thermo-hydraulic behavior, assess network states in

near real time, and evaluate control interventions before applying them to real operation. Beyond model creation, the platform was extended with improved UI/UX elements to support practical usage, including the visualization of network status, demand-response signals, and operational targets.
The digital twin development also required iterative validation and calibration loops. Measured real-world data were compared against simulated values under different operating conditions, and model parameters were refined accordingly. This systematic alignment of model and reality is an important prerequisite for meaningful optimization, since even small deviations may distort predictions of temperature propagation, pressure stability, and loss behavior.

Key findings

- Digital twins for DH grids were implemented and calibrated to enable high-fidelity simulation and operational analysis.
- Iterative validation against measured data proved essential to ensure that control decisions derived from simulations remain operationally trustworthy.
- The digital twin environment established a shared technical basis connecting live monitoring, simulation and optimization workflows.

4.2 Demand-Response Optimization and Supply Temperature Control

A key technical outcome is the development and validation of a demand-response driven control strategy for minimizing supply temperature while maintaining end-user comfort. In Rohrbach, a real-time optimization algorithm was created that continuously computes the lowest feasible supply temperature that still satisfies heating requirements at all substations. Lowering the supply temperature reduces distribution losses, smooths operational behavior, and can reduce thermal stress on components. During pilot tests, simulations and operational comparisons indicated an average reduction of approximately 5 K in supply temperature, achieved without compromising comfort. The control loop was designed for robust real-time behavior and includes mechanisms such as filtering to suppress sensor noise, safety bounds, and fallback strategies for data gaps or unstable conditions.
In addition, the control methodology evolved from an initial optimization procedure to a more computationally efficient and platform-integrated implementation. Early convergence limitations were addressed through parallelization, and later the algorithm was translated into Python and integrated into the Arteria platform, showing only minor numerical differences attributable to rounding. The project results demonstrate that such model-based control is technically feasible and can be embedded into an operational platform, provided data interfaces and safety logic are robust.

Key findings

- A real-time algorithm was validated in Rohrbach to reduce supply temperature while maintaining comfort. Pilot tests achieved ~5 K average reduction.
- Robustness measures (filtering, bounds, fallback behavior) were necessary to avoid instability from measurement noise or missing data.
- Translating and integrating the control approach into the platform (MATLAB → Python) enabled practical deployment within the operational toolchain.

4.3 Optimization of Heat Production via Predictive Supervisory Control

Beyond network-side temperature optimization, DOPPLER also assessed optimization potential on the heat production side using predictive supervisory control. For Rohrbach, the heating central was modelled within BEST's modular energy management framework and controlled via model predictive control principles. The project demonstrates how optimization-based scheduling can replace conservative rule-based operation by using forecasts and system models to reduce unnecessary switching, stabilize operation near efficient operating points, and define operational goals more directly (e.g. desired storage energy) rather than tuning conventional controller parameters.
At the same time the results highlight that the magnitude of achievable benefit depends strongly on system complexity. In Rohrbach, flexibility options are limited because there are no multiple competing heat sources. The biomass boiler remains the only controllable source and flue gas heat recovery is tied to boiler operation. Consequently, predictive optimization primarily supports smoother operation and improved setpoint strategies, while major cost-optimal source switching is structurally constrained in such a setup. The project therefore clarifies both the applicability and the boundaries of predictive supervisory control in real DH plants.

Key findings

- Predictive supervisory control is technically feasible and supports more stable, goal-driven operation than purely rule-based approaches.
- In simple production setups with limited flexibility, optimization benefits exist but remain constrained compared to multi-source or price-driven systems.
- Explicit modelling of components (e.g. flue gas condensation as proportional optional source) enables structured evaluation of operational strategies.

4.4 Data Interfaces and Cross-Site Integration Lessons

A major practical result of the project lies in the implementation of multiple data integration pathways across heterogeneous operator environments. Because demonstrator networks differed in SCADA systems, governance constraints and available interfaces. DOPPLER implemented a layered pipeline approach combining acquisition, secure transport, intermediate buffering, time-series storage and harmonization. Across sites, data were obtained via OPC links, HTTPS APIs, file exports (CSV) and periodic manual reads. Time alignment, buffering, plausibility checks and signal harmonization were applied to improve reproducibility and suitability for modelling and analytics.
The project also demonstrates that technical feasibility is often shaped by operational constraints. In Mischendorf, for example, a token-based API existed, but an OpenVPN client failure prevented the intended continuous ingestion approach, leading to reliance on local collection and CSV exports with periodic deletion cycles. In Güssing, continuous substation data streams were not available. Instead, longer time series had to be reconstructed through periodic readouts of Kamstrup meter archives and supplementary sources such as Belimo energy valves, with additional automation effort required for data extraction.

Key findings

- Cross-site integration required multiple interface types (OPC, APIs, CSV exports, manual reads), making harmonization and buffering essential.
- Data availability and governance constraints (VPN failures, limited device retention, restricted connectivity) significantly influenced what could be validated in practice.
- Creating consistent datasets required mapping of heterogeneous signal labels to canonical variables and filtering to a stable KPI core.

4.5 Customer-Side Integration with the meo PROPILOT Controller

On the customer side, DOPPLER implemented a practical test setup with two meo PROPILOT controllers installed in Rohrbach households. The integration targeted both efficiency and usability: local control of heating circuits and setpoints was combined with monitoring, visualization in end-user dashboards and a standardized 0–10 V interface to the transfer station for requesting required supply temperatures. The controllers also provided forecast-based control logic and served as structured data sources for project partners through an API interface, with processing steps such as resampling and cleaning to handle non-uniform temporal resolution.
A significant extension was the use of the System_PROFILER module, enabling long-term monitoring and alarms (e-mail notifications) for critical parameters such as room and supply temperatures and domestic hot water. As described in the report, the meo platform thereby acted as both a local optimization layer and a cloud-connected data hub supporting network-level analyses and modelling.

Key findings

- Two meo PROPILOT controllers were successfully installed and integrated, combining local control, user-facing dashboards, and standardized interfacing to the transfer station.
- API-based data provision (with cleaning/resampling) enabled structured sharing of customer-side measurements for modelling and optimization tasks.
- System_PROFILER added operational value via long-term monitoring and alarms, supporting comfort and reliability.

4.6 Dashboards as a Shared Operational and Engagement Layer

Another core outcome of DOPPLER is the development of dashboards that translate engineering metrics into actionable insights for both operators and end customers. For operators, the dashboards enable efficiency assessment and prioritization of measures using KPIs such as m³/MWh, temperature spread and aggregated energy and volume values, complemented by visualization patterns such as annual heatmaps, treemaps and leak-volume detection based on flow without thermal power. This supports identifying inefficient transfer stations, tracking interventions across reference periods and understanding seasonal patterns.
For end customers, the dashboard approach reduces complexity into a small set of intuitive indicators (efficiency score, ranking, spread indicator) and uses clear visual logic to provide feedback and

motivation. Importantly, the results show pronounced heterogeneity between installations and seasonal patterns, while also emphasizing practical limitations like data quality issues, outliers and inactive periods can affect interpretability and require robust preprocessing and KPI definitions.

Key findings

- Operator dashboards enabled targeted identification of inefficiencies (e.g. high m³/MWh, leak-volume patterns) and tracking of measures over time.
- End-customer dashboards translated thermo-hydraulic relationships into intuitive feedback (scores/ranking), supporting behavior-driven efficiency improvements.
- Data quality and inactive periods remain a key limiting factor, making robust preprocessing and careful KPI design essential.

4.7 End-User Motivation and Acceptance: Households vs. Manufacturing

DOPPLER produced clear evidence that motivational mechanisms must be adapted to fundamentally different end-user groups. For households, survey results from Güssing show generally positive attitudes towards flexibility: 58% expressed willingness for more flexibility in hot water preparation and 68% for room heating. However, acceptance drops when flexibility implies immediate convenience restrictions, such as postponing hot water use. Measures perceived as individualized and automated are more accepted than those perceived as external interference.
To address motivation, the project integrated gamification elements into dashboards, notably a ranking mechanism based on consumption efficiency. Workshop feedback indicates that ranking was perceived positively and not rejected, though it is more appealing to less advanced users. More experienced users tend to prefer direct technical data. By contrast, the intended collective-benefit visualization (tree symbolism) was not sufficiently understood. Participants indicated that the “story” explaining cause–effect relationships is missing, reducing its motivational power.
For manufacturing, expert interviews in four companies in Güssing show that DH costs are too small a share of production costs to motivate substantial investments or process adaptations. Increasing flexibility would require more flexible automation and digital control of production processes, investments that are currently not foreseeable in the investigated cases. This result sets a realistic boundary for what demand-side flexibility can be expected from industrial users under current conditions.

Key findings

- Household flexibility is broadly accepted in principle (58% hot water, 68% room heating), but drops with perceived comfort loss or “external interference.”
- Ranking-based gamification was positively received, while collective-benefit visualization lacked an explanatory narrative and therefore failed to motivate effectively.
- Manufacturing flexibility is strongly constrained by low DH cost relevance and rigid processes; meaningful shifts would require major investments that are not currently expected.

4.8 Overall Conclusions: What DOPPLER Demonstrated

Taken together, DOPPLER demonstrates that meaningful efficiency and flexibility gains in district heating require an integrated stack: high-quality data pipelines, interoperable information structures, validated digital twins and practical control strategies, combined with user-facing tools that translate system logic into understandable feedback. The project delivered this stack across several demonstrators, showing that (i) real-time, model-based optimization is feasible, (ii) customer-side systems can be integrated as both control and data layers, and (iii) dashboards can form a shared "interface" between operators and end users.
At the same time, the project clearly identifies the limiting factors that must be addressed for broader scaling. These include heterogeneous and sometimes fragile data interfaces, retention and access restrictions on devices and the need for stronger narrative framing of collective benefit if behavioral change is expected from households. Finally, the industrial results underline that flexibility potential is not just a technical question, it is also an economic one, and in some segments the business case for deep flexibility does not yet exist.

5 Outlook and recommendations

Digital Twin Arteria Platform
Following the completion of the project, several further developments are envisaged for the digital twin based on the Arteria platform. A primary next step is the digitization of additional district heating networks and their implementation as digital twins. This expansion will strengthen the applicability of the platform for both planning and operational support and will allow its benefits to be demonstrated across a wider range of network configurations. In parallel, the demand response algorithm is planned to be further enhanced. Beyond the consideration of temperatures at the heat transfer stations, mass flows will be incorporated in order to determine optimal flow rates. This optimization can be realized through the manipulation of the outside temperature value, with the next concrete implementation step being the establishment of a write-data connection via the splitter interface.

GIS Platform
The developed GIS platform represents a future-proof, modular, and open approach tailored to the requirements of small district heating systems. By combining standardized data structures with open-source software components and flexible deployment options, the platform ensures scalability and long-term sustainability. The integration of FIWARE Smart Data Models, PostGIS, QGIS, and Grafana creates a transparent and interoperable environment that unifies technical, spatial, and operational information within a single system. Potential future developments include automated analytics, AI-based anomaly detection, integration into broader Smart City data platforms, and enhanced user engagement through interactive or gamified interfaces. Overall, the project demonstrates how open standards and modern digital tools can significantly improve the efficiency, transparency, and sustainability of decentralized heating networks and support intelligent, data-driven energy management at the local level.

meo PROPILOT Controller

For the meo PROPILOT controller, several key directions for further development and implementation are recommended. An initial focus lies on the expansion to additional sites, enabling systematic testing of scalability and robustness under diverse technical and operational conditions. This step is essential for validating the adaptability and performance of the controller across different infrastructure environments. At the same time, the continued development of control algorithms is of central importance. The integration of machine-learning approaches offers the potential for individualized optimization of heating profiles and energy flows, thereby improving system efficiency and responsiveness to user needs and external influences. Further improvements are recommended with regard to data synchronization, as standardized measurement data structures would enhance interoperability between systems and facilitate seamless integration into broader energy management frameworks. Finally, long-term system evaluations across multiple heating periods are necessary to reliably quantify energy savings and assess the impacts of the controller on district heating networks, providing a robust basis for future deployment decisions.

Motivation of End-Users

Future work on motivating district heating customers should focus on a deeper understanding of behavioral patterns and the development of more targeted incentive mechanisms. For households, further refinement of gamification elements is recommended to strengthen their motivational effect. In particular, the integration of storytelling approaches and clearer visualizations of collective benefits can improve users' understanding of how individual actions contribute to overall network efficiency. Adaptive feedback mechanisms may further enhance effectiveness by dynamically responding to user behavior and personalizing motivational strategies. In contrast, manufacturing companies require a different approach, as their flexibility is constrained by rigid production processes. Here, technical and economic incentives need to be more closely aligned with operational realities. Collaborative pilot projects could explore how digital control systems or energy management tools might gradually increase flexibility without requiring substantial upfront investments, while financial support schemes or tax incentives could further encourage adoption of energy-efficient technologies.

Business Model Perspective

The name DOPPLER and the underlying concept of digital twins for district heating systems offer an attractive and engaging entry point for potential users. Attractiveness and entertainment value, alongside robust functionality and meaningful content, are central prerequisites for practical success. From the documented challenges of individual district heating operators, concrete business cases were derived and corresponding business models developed or combined. Within DOPPLER, the focus is deliberately placed on a business model that emphasizes shared benefits and common-good-oriented value creation rather than the optimization of individual stakeholder profits. In addition, a range of further application scenarios for the digital district heating twin emerged, addressing the needs of different user groups.

Dashboards

The dashboards developed within the project provide a consistent information and control framework for both district heating operators and end customers. On the operator side, key performance indicators

such as m³/MWh, temperature spread, and aggregated energy and volume values enable the targeted identification of inefficient transfer stations, the prioritization of technical measures, and continuous performance tracking across reference periods. On the end-customer side, simplified and visually clear indicators, including efficiency scores, rankings, and spread indicators, translate complex thermo-hydraulic relationships into intuitive feedback and support behavior-driven efficiency improvements. The results reveal substantial heterogeneity between customer installations, pronounced seasonal effects, and identifiable drivers of losses. Current limitations are primarily related to data quality and inactive periods, while future enhancements may include automated alerts, forecasting functions, cost and CO_2 accounting, and optimized mobile layouts.

6 Bibliography

None.

7 Appendix

None.

8 Contact details

Project manager: DI (FH) Klaus Paar
Institute/Company: Güssing Energy Technologies GmbH
Contact address: 7540 Güssing, Wiener Straße 49, +43 3322 42606 322, k.paar@get.ac.at ,
https://www.get.ac.at/doppler/

List of other project and cooperation partners company:

Arteria Technologies e.U.
BEST – Bioenergy and Sustainable Technologies GmbH
eKUT GmbH
Energieagentur Obersteiermark GmbH
Güssinger Fernwärme GmbH
Industriewissenschaftliches Institut (IWI)
KD Nahwärme GmbH
MEO Energy GmbH
Ringhofer & Partner GmbH
Scheiber Solutions GmbH