Ngoc Thuy Linh Nguyen

Development of a Graph-Based Tool for Planning of District Heating Network

School of Sustainability Energy
Master of Science, Smart Cities and Communities
Erasmus Mundus Joint Master

Vaasa 2026

UNIVERSITY OF VAASA
School of Computing and Electrical Engineering

Author: Ngoc Thuy Linh Nguyen

Title of the thesis: Development of a Graph-Based Tool for Planning of District Heating Network

Degree: Master of Science

Degree Programme: Smart Cities and Communities

Supervisor: Olivier Kaufmann

Year: 2026 **Pages:** 66

ABSTRACT:

As cities look for practical ways to reduce fossil-fuel use in heating, district heating offers a shared infrastructure through which buildings can be connected to more efficient and renewable heat sources. The planning of such networks is especially challenging when they are expected to grow over time, because each new consumer or heat source can change the pressure, flow, temperature, and cost profile of the whole system. This thesis develops HeatNet, a transparent graph-based pre design tool for testing radial district-heating networks before detailed engineering design is carried out. The tool represents the supply and return network as a radial tree and combines topology validation, demand-based flow estimation, pipe sizing, hydraulic balancing, thermal propagation, final-state checking, and distribution-side cost comparison in one reproducible workflow. It is applied to reference expansion cases and to a campus-scale case study at the University of Mons, using network layouts derived from geographic data and heat-demand information from building consumption records. The results show that HeatNet can compare alternative radial layouts, reveal when a passive single-pump network becomes hydraulically weak, and indicate balancing valves as a practical correction when short loops are overdriven. The thesis concludes that a topology-aware graph formulation is a useful basis for early-stage district-heating planning, particularly when several expansion or supply scenarios must be compared before committing to a final design.

KEYWORDS: District Heating Network, Pre-design Techno-economic Evaluation, Radial Tree, Oriented Objective Programming, Graph Traversal, Nonlinear Least Square

Acknowledgements

I would like to express my sincere gratitude to the SMACCs Erasmus Mundus Joint Master's Programme for giving me the opportunity to join the programme and for believing in my potential, even before I had completed my undergraduate studies. Being admitted to this programme opened doors that I had only dreamed of before, and I remain deeply thankful for this opportunity. My heartfelt thanks go to Professor Zacharie De Grève and Professor Hannu Laaksonen for his constant support throughout my studies. Whether it concerned administrative procedures, paperwork, or academic matters, he was always available to provide guidance and assistance. His dedication and kindness made many challenges much easier to overcome.

I am deeply grateful to Professor Olivier Kaufmann and Professor Eric Dumont for supervising this thesis. Their patience, guidance, and the many discussions we shared throughout this work have been invaluable. Beyond academic supervision, they introduced me to practical knowledge and real-world perspectives that I had never expected to encounter during my studies. I would also like to thank them for the opportunity to undertake a summer internship, which provided me with valuable insight into the industry and significantly contributed to the development of this thesis.

Last but certainly not least, I would like to thank my parents. Their unconditional love, trust, and support have been the foundation of every achievement in my life. They have always done everything within their means to provide me with opportunities and to help me pursue my ambitions. Completing this master's degree has been a dream of mine, and it would not have been possible without their sacrifices and encouragement. This accomplishment belongs as much to them as it does to me.

Thank you to everyone who has been part of this journey.

Use of Artificial Intelligence

Artificial intelligence tools, including ChatGPT, Claude, Grammarly, and QuillBot, were used during the preparation of this thesis to assist with language editing, grammar checking, text refinement, paraphrasing, and code debugging.

These tools served as supporting resources to improve the quality and readability of the manuscript and to facilitate the development process. All technical decisions, research design, analyses, interpretation of results, and conclusions remain the sole responsibility of the author.

The author has reviewed and validated all content produced with the assistance of these tools and takes full responsibility for the final version of the thesis.

Contents

Pictures

Tables

Abbreviations

Table 1. Abbreviations.

Abbreviation	Meaning
CAPEX	Capital expenditure
DH	District heating
DHN	District heating network
GIS	Geographic information system
HX	Heat exchanger
LCOH	Levelized cost of heat distribution
NPV	Net present value
OPEX	Operating expenditure
PUR	Polyurethane (insulation material)
SMACCs	Smart Cities and Communities (Erasmus Mundus programme)
UMONS	University of Mons

1 Context

1.1 Motivation and Background

Since 2021, the European Union has experienced one of the most severe energy disruptions in its recent history. Household natural gas prices reached 12.33 EUR per 100 kWh in the second half of 2024, the highest level recorded since 2008, even after a partial easing that followed the 2022 peak [1]. Electricity prices for households remained close to their post-crisis maximum over the same period [2] (Fig. 1.1). These price shocks reflect a deeper structural shift triggered by the Russian invasion of Ukraine in February 2022, which forced the EU to reduce its dependence on Russian fossil-fuel imports as a matter of strategic priority.

According to Eurostat data, the Russian share of EU oil imports fell from 29.2 % at the end of June 2021 to 2.3 % two years later, the share of natural gas from 38.5 % to 12.9 %, and the share of coal from 45 % to 0 % [3]. Although Russian gas still represented roughly 18 % of extra-EU natural-gas import value in early 2024 (Fig. 1.2), the absolute volume dropped from about 140 billion m^3 in 2021 to approximately 42 billion m^3 in 2023 [4]. The transformation is therefore both structural and ongoing, and energy security has become an explicit pillar of European energy policy rather than a background concern.

Compounding this transition, the European Climate Law commits the Union to reducing net greenhouse-gas emissions by at least 55 % by 2030 relative to 1990 levels, on a trajectory toward climate neutrality in 2050 [5]. The REPowerEU plan, launched in May 2022, reinforces these objectives by simultaneously phasing out Russian fossil fuels, accelerating energy savings, and expanding renewable deployment [6]. The combination of price shock, security concerns, and climate commitments therefore creates a coherent and unusually strong set of incentives to decarbonise the European energy system on every front where significant final-energy consumption occurs.

Among end uses, heating and cooling stand out as the largest single share of European final energy consumption. According to Eurostat, energy used for heating and cooling represents almost 50 % of the EU's total gross final energy consumption [7]. The dominance is even more pronounced at the household level: space and water heating together accounted for 77.6 % of the final energy consumed by EU households in 2023, while the share of renewables in heating and cooling reached only 26.7 % in 2024 [8]. Decarbonising heat is therefore not a peripheral question of the energy transition; it is the lever with the largest absolute leverage on final-energy emissions.

The picture varies considerably between Member States. Eurostat reports that space heating alone ranges from 22.0 % of household energy consumption in Malta to 79.3 % in Luxembourg, with an EU average of 62.5 %. Belgium sits near the upper end of this distribution at 70.8 % [9], which reflects both the country's climate and the predominance of individual fossil-fuel boilers in its building stock.

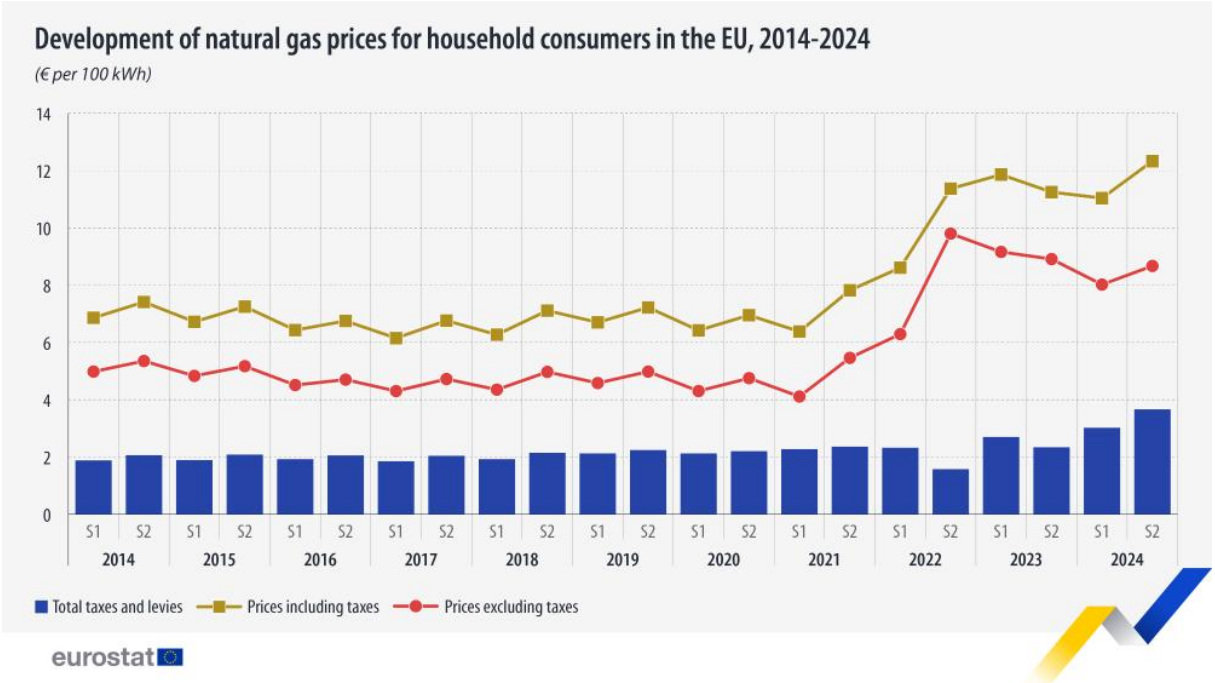

Figure 1.1. Evolution of natural gas prices in the EU, 2021–2024 [1].

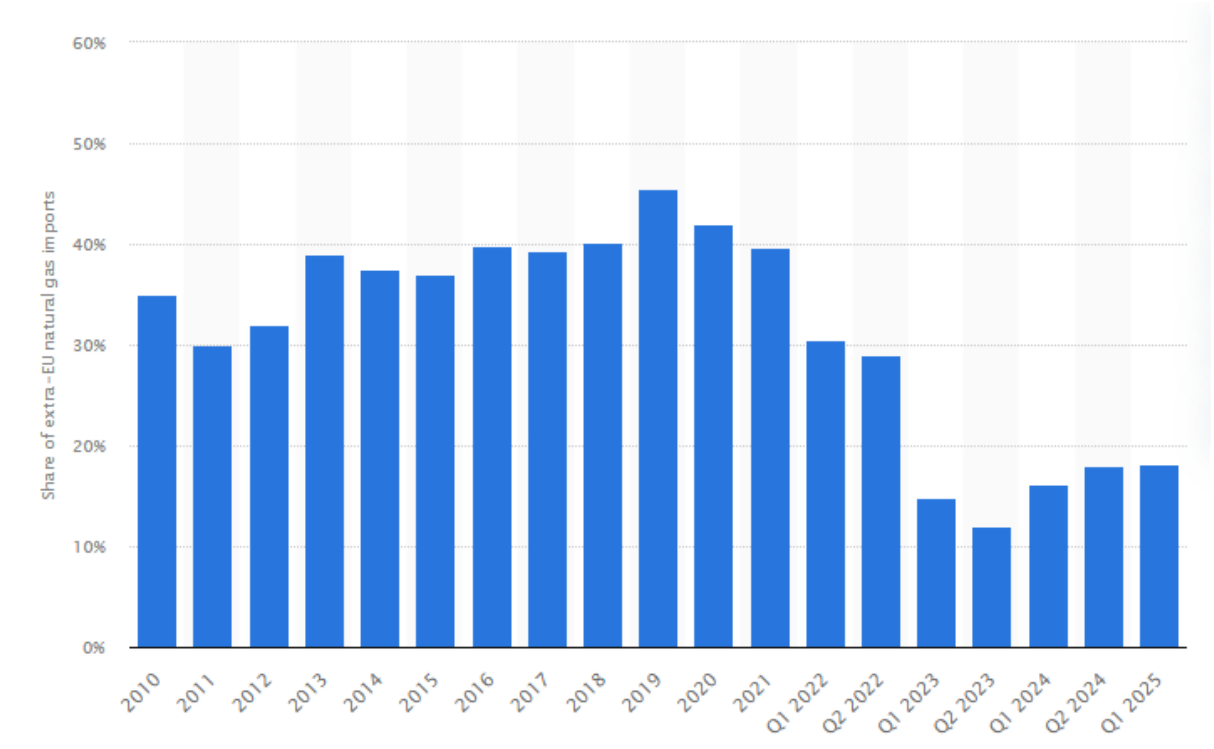

Figure 1.2. Share of extra-EU natural-gas import value from Russia, 2010–Q1 2025 [4].

1.2 Topologies of District Heating Networks

Understanding how heat is distributed in a system requires knowledge of the structural layout, or *topology*, of the network. District heating and cooling (DHC) networks are generally classified into the following basic configurations [10, 11, 12].

Radial (tree) network. The radial or tree topology is the simplest type of network, consisting of a central energy supply system (DES) that acts as the root, from which pipe branches extend to deliver hot water to buildings [10]. The characteristic of this network is that flow moves in only one direction from the source to the loads, and similarly in the return pipes. The flow directions remain fixed during operation, meaning there are no closed loops, and the pipe investment cost is the lowest among all network structures [12, 11]. A simple radial layout is shown in Fig. 1.3.

- **Advantages.** The radial layout requires the shortest total pipe length among all topologies, which translates directly into the lowest capital investment for trenching and pipe material [11].
- **Disadvantages.** Reliability is limited because a single pipe failure can disconnect all downstream consumers. Expandability is also constrained: adding a new consumer or integrating a distributed heat source alters the pressure profile in a complex way, potentially requiring repumping of the entire branch [10, 11].

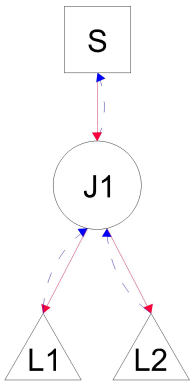

Figure 1.3. Example of a radial (tree) network with one source and two loads.

Ring and meshed networks. In contrast to the radial topology, a ring network connects pipes into a single closed loop, allowing the integration of multiple heat sources and significantly improving supply reliability [10]. Over time, as the heating network continues to expand with multiple interlocking loops, it evolves into a meshed network [11]. The meshed topology provides extremely high supply reliability, facilitates the integration of multiple distributed heat sources, and enables flow to move flexibly in both directions, but it requires the highest initial investment due to the large total pipe length and complex system [10, 12, 11]. The progression from radial to ring and meshed layouts is illustrated in Fig. 1.4.

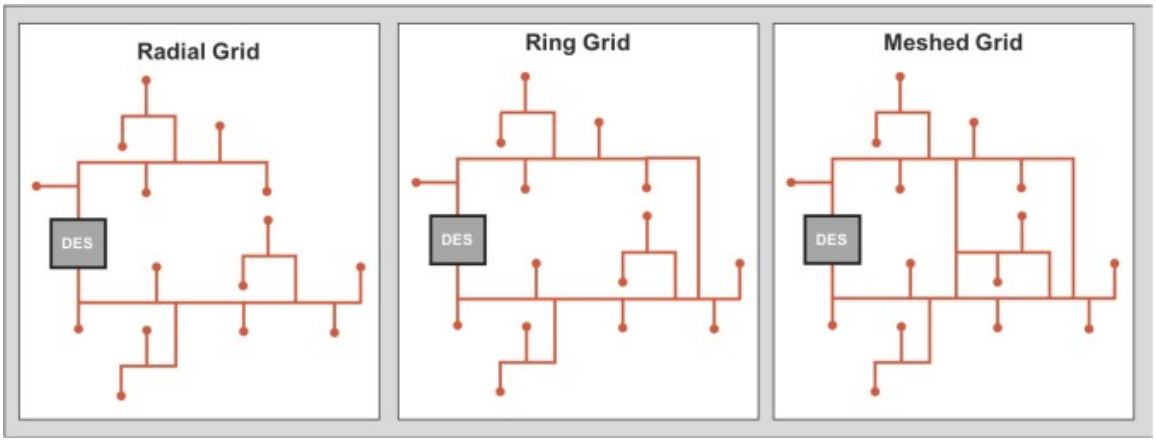

Figure 1.4. Comparison of the evolution from radial tree to single closed loop (Ring) to multiple interlocking loops (Meshed) [10]. DES: distributed energy system.

Bidirectional network. A bidirectional network is an advanced operating mechanism in which smart pumping stations route thermal energy precisely to the locations where it is needed, rather than forcing flow in a fixed direction [10]. Thanks to the flexibility of this bidirectional flow model, the concept of dedicated supply and return pipes is eliminated; the system instead consists only of "hot" or "cold" pipes [10].

Prosumer network. In a prosumer network, customers are not merely consumers but can also act as producers thanks to distributed thermal generation devices (such as solar thermal collectors, heat pumps) and thermal storage tanks located at the buildings [13, 14, 15]. Through *bidirectional substations*, these prosumers can self-consume the thermal energy they produce and flexibly pump any excess heat back into the shared district heating network for other buildings [14, 13]. The combination of the prosumer model and bidirectional substations is considered a key enabling technology for integrating renewable energy and decarbonising urban heating [15, 13, 14]. A representative prosumer configuration is shown in Fig. 1.5.

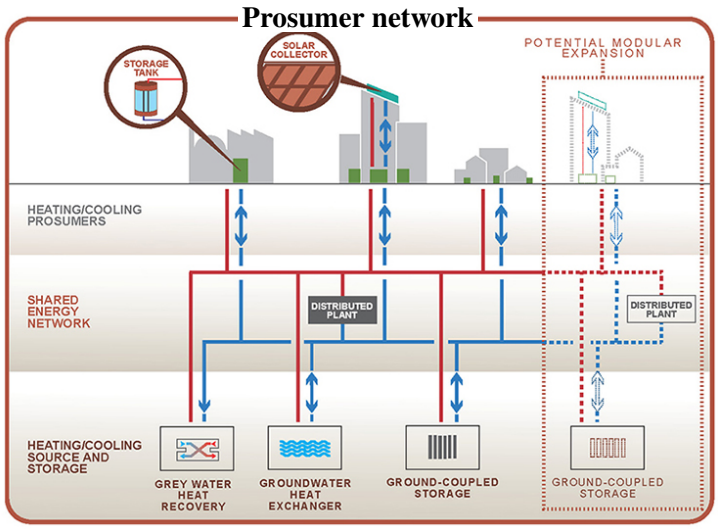

Figure 1.5. A building drawing heat from the network while collecting solar heat from rooftop panels for self-use and pumping excess heat back into the network.

1.3 Problem Statement

District heating (DH) is one of the few technologies that addresses heat decarbonisation at neighbourhood and city scale rather than building by building. By distributing hot water from one or several centralised generation units through an insulated pipe network, DH allows the use of heat sources that would otherwise be inaccessible to individual consumers: geothermal heat, industrial waste heat, large heat pumps, biomass, or solar thermal collectors. DH currently supplies approximately 12 % of total EU heating demand, with national penetration rates as high as 65 % of urban connections in Denmark and renewable shares above 70 % in Swedish DH systems [16] (Fig. 1.6). These figures demonstrate that high-share, low-carbon district heating is technically and economically feasible at country scale. Looking forward, the Heat Roadmap Europe study recommends expanding DH from roughly 13 % to 55 % of European heat supply by 2050, saving an estimated 280 TWh of final energy and requiring an investment of about €1.16 trillion across the network and generation infrastructure [17]. From a regulatory standpoint, the revised Renewable Energy Directive (RED III) sets an indicative target of an average increase of 2.2 percentage points per year in the share of renewables in district heating and cooling from 2021 to 2030, while the IEA Net Zero scenario projects renewables to exceed 8 % of DH supply by 2028 [18]. DH is thus simultaneously proven in practice, recommended by independent studies, and supported by binding European regulation.

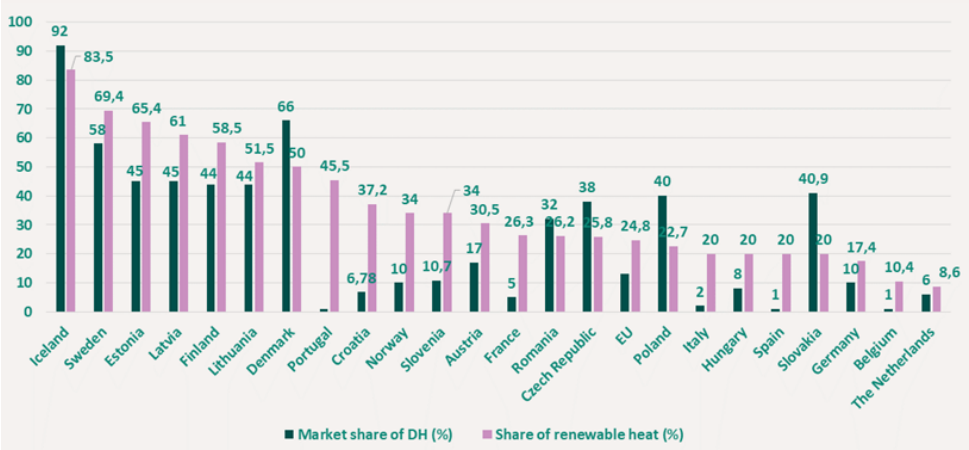

Figure 1.6. Shares of district heating and renewable energy in gross final consumption for heating and cooling (year 2022), source: Euroheat & Power and Eurostat [19].

Belgium, however, lies far behind this trajectory. A recent feasibility study by Sensil et al. [20] on fifth-generation district heating and cooling (5GDHC) in Belgium concludes that DHC in the country remains at an "infant stage", with no operational 5GDHC system in Wallonia and only a handful of pilot initiatives in Flanders. In Wallonia, roughly 4 % of households are classified as energy-poor, and residential heating still relies heavily on individual fossil-fuel installations. The contrast between the EU-level momentum behind district heating and its near-absence on the Walloon territory defines the practical gap addressed by the present work.

Deploying a district-heating network in practice is rarely a one-shot operation. A first phase typically connects a small number of anchor consumers, and the network grows incrementally as additional buildings, neighbourhoods, or industrial sites request connection. Each such extension

raises a coupled hydraulic–thermal question that cannot be reduced to a simple load addition: when a new consumer is connected, the upstream branches must carry redistributed flow, the pump operating point may need to be raised, commercial pipe diameters may become locally undersized, and the supply temperature delivered at the farthest load may drop below its required minimum. The central engineering question of incremental network design is therefore: *can the existing network still deliver sufficient mass flow and supply temperature to every consumer, including the new ones, without violating its operational limits?*

Answering this question quickly, transparently, and reproducibly is essential if district heating is to scale at the pace required by EU climate targets. The challenge is particularly acute in regions such as Wallonia, where only a few operational DH networks exist to learn from (Fig. 1.7) [21] and where each new project must be justified on its own technical and economic merits, often before its long-term expansion path is fixed.

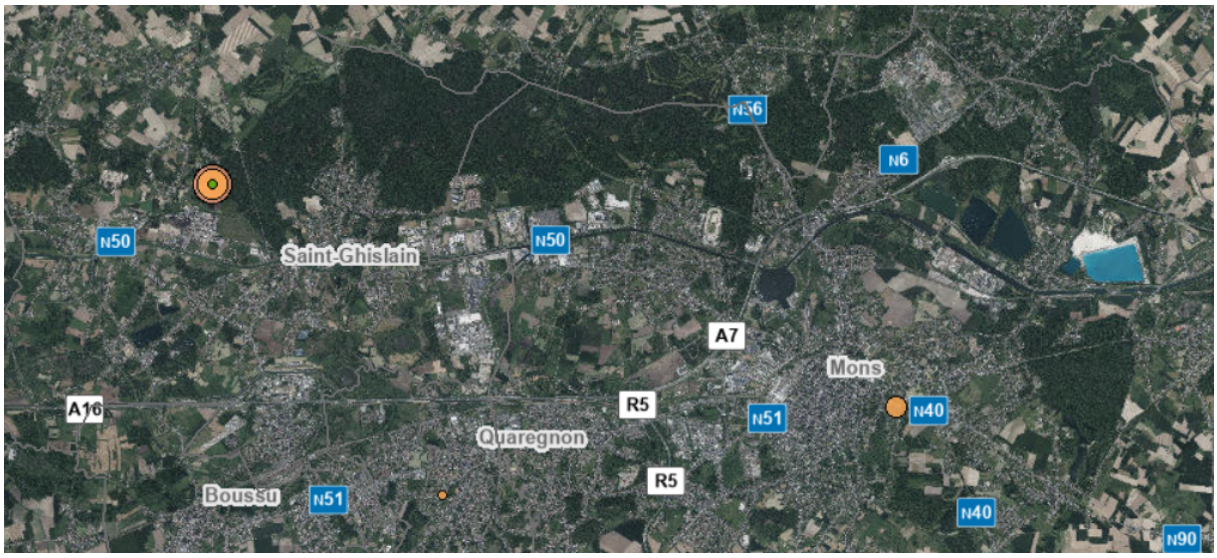

Figure 1.7. Geothermal well in Wallonia. [21] (accessed 05/06/2026).

Commercial software packages such as EcoStruxure District Energy (formerly TERMIS) [22] and Apros DH [23] can handle these calculations, but they are typically expensive, closed, and difficult to integrate into a fast iterative workflow when an engineer wishes to compare several expansion scenarios before committing to a layout. Open-source alternatives exist, but their internal logic is rarely documented at the level of granularity required for a master's thesis, and most of them are not specifically structured around the scenario-comparison question that motivates this work. There is therefore a clear methodological gap for a transparent, fully documented, topology-driven tool, compatible with network layouts derived from geographic (GIS) data, able to test network extensions against coupled hydraulic and thermal feasibility criteria in a reproducible way.

This thesis develops such a tool, referred to throughout the document as HeatNet, and demonstrates its use as a pre-design decision-support instrument for the analysis of radial district-heating networks under incremental expansion, providing rapid feasibility assessments before detailed engineering design is committed.

1.4 Scope and Limitations

The thesis pursues four interconnected objectives:

1. To develop a graph-based computational workflow capable of reading a network description from tabular CSV input files, validating its radial supply–return topology, and constructing a directed graph that supports every subsequent hydraulic and thermal calculation.

2. To implement a staged hydraulic–thermal solver that

   (i) constructs a design-state mass-flow field from the prescribed load demands (Phase 0),

   (ii) selects commercial pipe diameters by upward snapping from a catalogue,

   (iii) replays the design state to evaluate pipe pressure drops and identifies the critical loop that defines the initial pump head (Phase 1),

   (iv) solves the fixed-head pressure-balancing problem (Phase 2),

   (v) propagates supply and return temperatures through the network (Phase 3),

   (vi) applies a hydraulic correction on the pump head until every consumer receives at least its design mass flow (Phase 4), and

   (vii) applies a thermal correction of the source supply temperature until every consumer receives at least the required minimum supply temperature (Phase 5).

3. To guarantee traceability by exporting at every stage the intermediate quantities, including design and actual mass flows, theoretical and commercial diameters, Phase 1 and Phase 2 pressure drops, residual balancing quality, pipe heat losses, and a global first-law energy balance, so that every calculation step can be audited and reproduced.

4. To demonstrate the tool as a decision-support instrument by applying it to a reference radial network and to three expansion scenarios of increasing topological complexity, including a case in which a new internal junction is inserted within the existing tree.

A secondary objective is to couple the technical model with a first-order economic layer that estimates network CAPEX, OPEX, pump electricity, and net present cost, so that alternative expansion strategies can be compared not only in engineering terms but also in financial terms.

Scope. The present work is restricted to steady-state simulation of radial supply–return networks at the distribution boundary. Heat-generation plant sizing (geothermal wells, boilers, large heat pumps), transient and start-up behaviour, meshed or looped topologies, and detailed substation modelling are deliberately out of scope. Fluid properties are treated as fixed representative values within the pressure-balancing stage, and a single mean ground temperature is used in the thermal model.

Main contribution. The main contribution of this work is a transparent, fully auditable, topology-driven workflow in which every intermediate quantity used by the solver is exported and verifiable, and in which the same workflow is shown to handle structural topology changes, including the insertion of a new internal junction, without manual reformulation of the hydraulic residual system. To the author's knowledge, no openly documented thesis-scale tool currently offers this combination of staged hydraulic–thermal correction, exported energy-balance verification, and built-in expansion-scenario comparison for radial district-heating networks in the Belgian context.

2 State of the Art

2.1 Modeling paradigms for steady-state DHN simulation

District heating network (DHN) models are commonly classified according to two dimensions: temporal resolution and spatial granularity. In terms of temporal resolution, models may be steady-state, quasi-steady or quasi-dynamic, and fully dynamic. In terms of spatial representation, they may use lumped, nodal, or distributed pipe formulations. Steady-state and quasi-steady formulations remain widely used for design-level DHN studies, including network layout evaluation, pipe sizing, insulation sizing, sensitivity analysis, and static techno-economic assessment [24, 25, 26]. Their computational efficiency makes them attractive when the objective is to evaluate many design alternatives rather than to reproduce short-term operational transients. However, when the focus shifts to operational optimization, peak-load estimation, demand-side management, or control actions involving strong thermal transients, dynamic or quasi-dynamic formulations become more appropriate [24, 25, 27, 28].

Quasi-dynamic models usually rely on the assumption that the hydraulic state can be treated as quasi-steady while the thermal field evolves over time. This separation is commonly justified by the much faster propagation of pressure disturbances compared with thermal fronts in district-heating pipelines. Nevertheless, fully dynamic studies have shown that neglecting hydraulic transients may become inaccurate in long-distance DH systems or during fast operational changes [27].

Within steady-state DHN simulation, three main modeling paradigms can be identified. The first is the equation-based nodal formulation, where mass conservation at nodes and momentum conservation along pipes are assembled into a nonlinear algebraic system. Pipe pressure losses are usually represented using Darcy–Weisbach-type equations with suitable friction-factor correlations, while the thermal part is described through energy balances and steady heat-loss expressions [26, 29, 30]. The second paradigm is the equivalent-circuit or thermal-resistance analogy. In hydraulic applications, pressures and mass-flow rates may be treated analogously to voltages and currents, whereas in heat-loss calculations temperature differences and heat fluxes play the analogous roles. This approach is particularly useful for compact steady-state representations of buried-pipe heat losses. Wieland et al. [31], for instance, compared an equivalent-circuit-based method with standardized procedures such as EN 13941. The third paradigm is design-oriented sizing and techno-economic assessment, where steady-state or quasi-steady thermo-hydraulic calculations are used to compare pipe dimensions, insulation levels, pumping requirements, heat losses, and distribution costs [26, 32]. In some studies, these calculations are embedded in formal optimization frameworks [26], whereas in others they are used as calculation tools for scenario-based comparison of design guidelines or sizing alternatives [32].

The present work belongs to the equation-based family, but with an important specialization. Rather than assembling a global nodal nonlinear system for arbitrary meshed networks, the present work exploits the radial topology of tree-structured DHNs to impose an ordered solution procedure. In this setting, mass-flow accumulation, pressure propagation, and temperature

propagation can be performed directly through graph traversal along the unique supply and return paths.

2.2 Sequential and fully coupled thermo-hydraulic solvers

A central methodological issue in DHN simulation is how to treat the coupling between the hydraulic variables, namely pressures and mass-flow rates, and the thermal variables, namely supply temperatures, return temperatures, and heat losses. Two broad solver strategies are used in the literature: sequential and fully coupled formulations.

The sequential, or decoupled, approach solves the hydraulic and thermal subproblems separately. Typically, consumer mass-flow requirements are specified first, the hydraulic problem is solved to obtain pipe flows and pressure drops, and the thermal problem is then propagated on the resulting flow field. This approach is attractive because each subproblem remains smaller and better conditioned than a fully coupled thermo-hydraulic system [28, 30, 33].

In radial or directed acyclic networks, the thermal pass can exploit a fixed upstream-to-downstream ordering. Xiang et al. [33], for example, improved the efficiency of a decoupled steady-state solver by using a recursive search to determine the thermal propagation order on a directed acyclic graph. In contrast, Wang et al. [28] addressed meshed and bidirectional networks, where the thermal ordering may change because of flow reversals; they used a breadth-first search strategy to update the computational order during the sequential solution. Gumpert et al. [30] also adopted a sequential steady-state framework for a real DHN case study, but their contribution is better interpreted as a matrix-based hydraulic and thermal simulation workflow rather than as a radial traversal algorithm.

Fully coupled approaches instead assemble the hydraulic and thermal equations into one nonlinear system, usually solved by Newton–Raphson or related methods. This strategy is particularly relevant for low-temperature, meshed, bidirectional, and prosumer-based networks, where mass-flow directions may change and where the simple upstream/downstream ordering assumed by sequential methods may no longer be valid [34, 35]. Tang and Huang [34] introduced outflow indication variables to handle changing flow directions during Newton iterations. Dancker and Wolter [35] formulated a coupled Newton–Raphson framework in which hydraulic and thermal interdependencies appear explicitly in the Jacobian, facilitating integration with multi-energy system analyses.

It is important, however, to distinguish fully coupled mathematical formulations from sequential algorithms that are described as "coupled" because they exchange information between hydraulic and thermal modules. Zheng et al. [36], for example, solve the hydraulic and thermal problems alternately and use graph traversal to update the thermal calculation order. Therefore, their method is better interpreted as an advanced sequential or iterative coupling scheme rather than as a fully assembled Newton system.

The choice between sequential and fully coupled strategies therefore depends less on a universal hierarchy of accuracy than on network topology, operating regime, and the strength of thermo-hydraulic feedback. Sequential methods can also be extended to meshed networks through iterative ordering updates or matrix-based hydraulic formulations [28, 30, 33], but they are

especially simple and robust for radial, unidirectional networks where the upstream/downstream order is fixed. Fully coupled Newton-type solvers are most justified when bidirectional flows, prosumers, or strong thermo-hydraulic interactions make this ordering uncertain [34, 35].

2.3 Numerical methods for hydraulic network solving

The hydraulic subproblem in DHN simulation inherits many techniques from water-distribution and gas-distribution network analysis. The classical Hardy Cross method iteratively corrects guessed flow rates around network loops until mass and energy balance conditions are satisfied [37, 38]. Its appeal lies in its conceptual simplicity and ease of implementation. However, Hardy Cross-type methods may converge slowly for stiff or large networks and can become cumbersome when many loops and sign changes are present. This has motivated node-loop and simultaneous-loop variants, which solve more directly for the unknown flows or loop corrections [37].

In modern DHN research, Newton–Raphson methods are widely used for solving nonlinear hydraulic equations, especially in meshed networks. In these formulations, mass conservation and pressure-loss equations are assembled from the network graph, and the resulting nonlinear residual, which is the difference between the target and the real calculated value, is solved iteratively [29, 34, 35]. Newton-type methods offer fast local convergence, but they can be sensitive to initialization and ill-conditioning, particularly when pipe diameters, roughness values, or operating regimes vary strongly across the network. For this reason, robust variants such as damped Newton schemes, line-search strategies, or residual-minimization formulations are often introduced to improve convergence behavior [29].

For radial DHNs, however, the hydraulic problem is structurally simpler. Because each node has a unique path from the source, mass conservation can be enforced by a single graph traversal rather than by solving a global hydraulic system. Pipe-wise pressure losses can then be evaluated locally using Darcy–Weisbach and friction-factor correlations. When pressure balancing between terminal branches is required, the nonlinear solve does not need to be performed over the entire network. Instead, it can be reduced to a branch-level least-squares problem, where the unknowns are terminal branch velocities and each residual enforces consistency between the branch pressure drop and the prescribed pump-head target. In the present work, since the network is radial, the solver can use the fixed branch paths that are computed in advance, instead of using a more complex meshed-network method that must handle loop-cutting.

2.4 Graph-based and topology-aware approaches

Recent DHN literature increasingly emphasizes the role of graph theory in both model formulation and solver design. Directed graphs provide a natural representation of DHNs, where nodes correspond to sources, junctions, or consumers, and edges correspond to supply and return pipes. Graph traversal algorithms have already appeared in DHN studies: Xiang et al. [33] used a recursive search to determine thermal propagation order on a directed acyclic graph, Wang et al. [28] employed breadth-first search to update computational ordering during flow reversals,

and Zheng et al. [36] applied graph traversal to refresh the thermal calculation sequence within an iterative coupling scheme. However, in each of these studies traversal serves as an auxiliary ordering step inside a broader matrix-based or iterative-coupling framework, rather than as the primary computational mechanism of the solver itself.

For radial networks in particular, the tree topology admits a stronger exploitation. Post-order and pre-order traversal, standard algorithms in computer-science graph theory [39, 40], map directly onto the physical propagation directions of a tree-structured DHN: post-order traversal accumulates mass flow from terminal loads toward the source, pre-order traversal propagates supply pressure and supply temperature from the source toward the loads, and a second post-order pass with flow-weighted mixing computes return temperatures at junctions. Based on the reviewed literature, no prior work was found that formulates graph traversal as a complete staged workflow specifically for radial DHN modelling and analysis, in which mass-flow aggregation, pipe sizing, pressure propagation, and thermal correction are all resolved through ordered traversals on the same radial structure without assembling a global nodal nonlinear system.

For the present work, this motivates a topology-aware solver specialized for radial networks. The supply graph is treated as a tree, mass flows are accumulated from terminal loads toward the source, pressure and temperature fields are propagated along unique supply and return paths, and return temperatures are computed through flow-weighted mixing at junctions. This avoids the unnecessary complexity of a full meshed-network Newton formulation while retaining the physically relevant steady-state thermo-hydraulic balances for radial DHNs.

3 Methodology

This chapter follows the HeatNet calculation sequence from input validation to final verification. Before Phase 0, the CSV input files are loaded and checked to ensure that the supply network is a valid radial tree. The methodology then moves through design-flow construction, pipe sizing, hydraulic calculation, pressure balancing, thermal propagation, hydraulic and thermal corrections, and a final energy-balance check. These phases are summarized in Table 3.1.

Table 3.1. Methodology phase overview.

Phase	Main purpose	Main output
Phase 0: design flow and sizing	Convert load demands into design mass flows and select commercial pipe diameters	Node flows, pipe mass flows, diameters, R_{th}
Phase 1: hydraulic calculation	Replay the design state and identify the critical pressure path	Pressure drops, H_{initial}
Phase 2: pressure balancing	Balance all load loops for a prescribed pump head	Actual balanced load flows, balanced branch flows
Phase 3: thermal propagation	Propagate supply and return temperatures through the radial network	Supply and return temperature fields
Phase 4: hydraulic correction	Raise the pump head only if a load remains undersupplied	Corrected H^*
Phase 5: thermal correction	Raise the source temperature only if the delivered supply temperature is too low	Corrected $T^*_{\text{source,supply}}$
Phase 6: final verification	Check the energy balance of the selected operating point	Verified final operating state

3.1 Graph-Based Network Construction and Validation

This chapter uses the same four-node radial network throughout the methodology. The base topology is shown in Fig. 3.1: one source S, one junction $J1$, and two terminal loads $L1$ and $L2$. Subsequent figures reuse this layout with different information layers rather than introducing a new graph each time.

Throughout the methodology, each physical link is treated as one pipe pair $e = (u, v)$. The red supply pipe follows the tree direction from the source to the loads ($u \rightarrow v$); the blue return pipe is the reverse direction ($v \rightarrow u$). The two pipes share the same length, diameter, material, and thermal resistance, while their flow directions are opposite. This convention is used in all hydraulic, thermal, and economic calculations below.

The network is loaded from CSV tables containing source data, load data, and pipe-pair connectivity. As illustrated in Fig. 3.2, the source node provides the source temperature and pressure, while load nodes provide heat demand and design temperature drop. Default values are used only when optional fields are missing: the specific heat capacity $c_p = 4180\ \text{J}\,\text{kg}^{-1}\,\text{K}^{-1}$

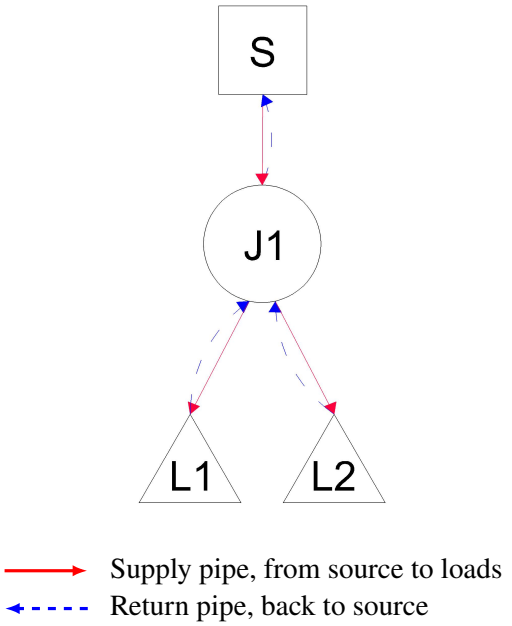

Figure 3.1. Running example topology and pipe-direction convention.

is the standard value for liquid water from the IAPWS-IF97 formulation [41], while the heat-exchanger pressure drop $\Delta p_{\mathrm{HX}} = 20000$ Pa is an assumed representative substation value, treated as a pre-design default in the same spirit as the economic parameters of Section 3.9.

Table 3.2. Example source input data.

Source ID	Connected node	Supply temperature (°C)	Source pressure (bar)	Supply mass flow (kg/s)
S	$J1$	90	5	1.5

Table 3.3. Example load input data.

Load ID	Connected node	Heat load Q (kW)	ΔT (K)	Design mass flow (kg/s)
$L1$	$J1$	40	20	0.5
$L2$	$J1$	80	20	1.0

Tables 3.2–3.4 show an example of the information contained in the CSV input files: source parameters, load demand data, and branch connectivity. For the radial branch S–$J1$, the mass flow is obtained by aggregating the downstream load flows, i.e., $\dot{m}_{S-J1} = \dot{m}_{L1} + \dot{m}_{L2} = 1.5\,\mathrm{kg\,s^{-1}}$.

Before any hydraulic or thermal calculation is performed, HeatNet checks the red supply arrow as a directed tree rooted at the source. The source has no parent link, every other node has exactly one incoming link, and a valid radial layer with N nodes has $N-1$ links, as expected for a spanning-tree network [42]. The validator also checks that every node is reachable from the source, so that no disconnected branch remains in the input data.

The blue return arrow is then generated automatically as the exact reverse of the supply path, rather than entered as an independent topology. This single pipe-pair convention ensures that

Table 3.4. Example branch/connectivity input data.

Branch ID	From node	To node	Branch type	Mass flow (kg/s)
S–$J1$	S	$J1$	Aggregated supply branch	1.5
$J1$–$L1$	$J1$	$L1$	Load branch	0.5
$J1$–$L2$	$J1$	$L2$	Load branch	1.0

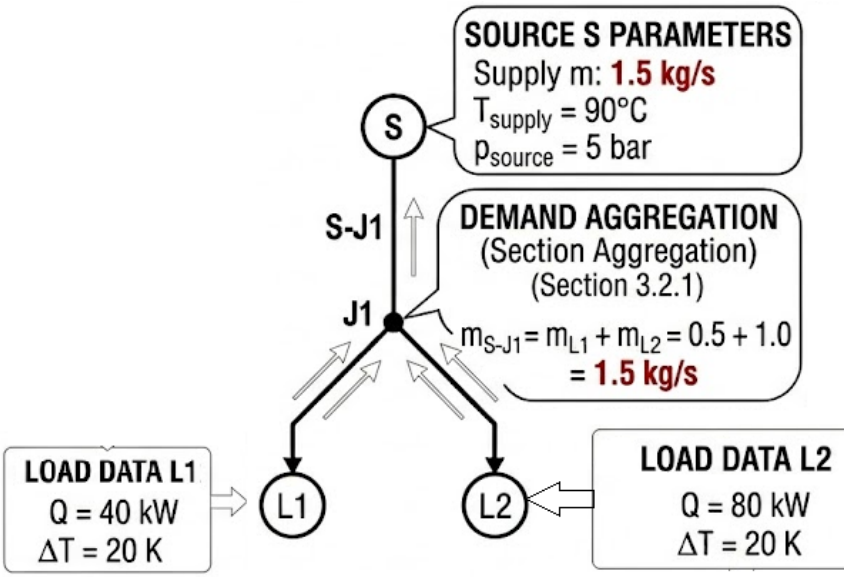

Figure 3.2. Input data assigned to the running example.

mass-flow aggregation, pressure-path tracing, and return-temperature mixing all operate on one unique path between each load and the source.

Terminal-load assumption. HeatNet assumes that consumers are terminal nodes, or leaves, of the red supply tree. Load nodes receive flow but do not carry flow onward to another consumer; through-flow is carried only by the source and junctions.

This assumption keeps the pressure-balancing problem well posed. When consumers are terminal nodes, each load loop has its own terminal branch, so a single pump head can be balanced by adjusting the terminal branch flows. If a consumer were placed on a shared path upstream of another consumer, the upstream loop would be contained inside the downstream loop. Equal loop pressure drops would then require the additional downstream branch to have zero pressure loss, which is impossible for a pipe with non-zero flow and length. For this reason, buildings that lie on a geographic routing path are converted into leaf connections before simulation.

3.2 Phase 0: Design-State Construction and Commercial Pipe Sizing

After topology validation, Phase 0 converts the thermal load data into a branch-wise design mass-flow field and uses this field to select commercial pipe diameters. This preprocessing step links three quantities directly: load heat demand, aggregated branch mass flow, and selected

pipe size.

The Phase 0 state is not the final operating point. It is the nominal flow field required if each load receives its prescribed heat demand at its design temperature drop. The resulting branch flows are used immediately for catalog-based pipe sizing and later for the Phase 1 hydraulic replay.

3.2.1 Demand Aggregation on the Radial Tree

Each load converts its heat demand into a design mass flow using the basic heat balance $\dot{m} = Q/(c_p \Delta T_{\text{design}})$. The model rejects non-physical inputs with $\Delta T_{\text{design}} \leq 0$.

The local load flows are then accumulated on the supply tree from the leaves back toward the source. In Fig. 3.3, $L1$ and $L2$ first receive their own design flows; junction $J1$ carries their sum; and the main branch $S \rightarrow J1$ carries the total network flow. By the pipe-pair convention, the same mass-flow magnitude is assigned to the paired return direction.

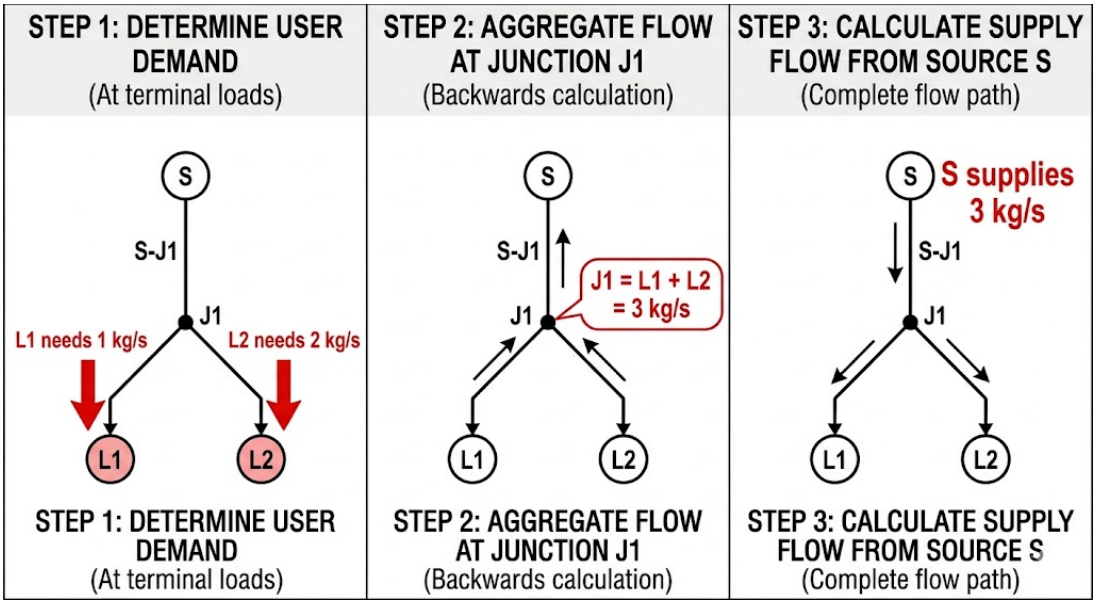

Figure 3.3. Bottom-up design-flow aggregation.

This bottom-up pass constructs the design mass-flow field used for sizing every pipe pair.

3.2.2 Discrete Sizing and Thermal Resistance

The Phase 0 flow field specifies how much water each pipe pair must carry, but it still gives a continuous hydraulic requirement rather than a buildable pipe size. The sizing step therefore converts each design mass flow into a commercial inner diameter. This diameter set is important because it is used by every later part of the model: cross-sectional area, velocity, pressure loss, thermal loss, and cost are all evaluated with the selected commercial diameter, not with the continuous theoretical value.

In the standard CSV workflow, pipe diameters are not prescribed in the pipe-pair input table. They are selected during this step and then carried through the hydraulic, thermal, and economic calculations. When an existing network is reused in a scenario analysis, the previously selected

diameters can instead be kept fixed so that the comparison reflects only the change in operating conditions or heat demand.

Velocity-based sizing. Sizing from a velocity criterion is a common design rule because it limits both pressure gradients and pumping costs [43]. For each pipe pair e, the design mass flow obtained from the upstream aggregation is converted into a continuous theoretical diameter,

$$D_{\text{theoretical},e} = \sqrt{\frac{4\dot{m}_e}{\rho \pi v_{\text{design}}}}, \tag{3.1}$$

Here, $D_{\text{theoretical},e}$ is the continuous inner diameter that would carry the design mass flow $\dot{m}_e$ at the target velocity v_{design}, and ρ is the water density used for the sizing calculation. The default target design velocity used in this work is

$$v_{\text{design}} = 1.0 \text{ m s}^{-1}. \tag{3.2}$$

If a pipe carries no design mass flow, Eq. (3.1) is not evaluated and the smallest available diameter is assigned. Otherwise, the theoretical diameter is passed to the discrete selection rule.

Commercial diameter and thermal resistance. Each theoretical diameter is snapped upward to the nearest commercial inner diameter:

$$D_e = \min \left\{ D \in \mathcal{D}_{\text{catalog}} \;\middle|\; D \geq D_{\text{theoretical},e} \right\}. \tag{3.3}$$

Here, D_e is the selected buildable inner diameter and $\mathcal{D}_{\text{catalog}}$ is the finite list of available commercial diameters. The upward snapping rule ensures that the installed pipe is never smaller than the continuous diameter required by the velocity criterion.

The adopted catalogue and thermal-resistance values are given in Table 3.5: DN20–DN110 are flexible PE-Xa branch pipes from Uponor Ecoflex Thermo Single [44], while DN125–DN300 are pre-insulated steel trunk pipes from Logstor Series 2 geometry [45]. The selected diameter and thermal resistance are stored once for the pipe pair and used in both directions:

$$R_{\text{th},e} = \mathcal{R}_{\text{th}}(D_e), \tag{3.4}$$

Here, $R_{\text{th},e}$ is the linear thermal resistance assigned to pipe pair e in K·m W^{-1}, and $\mathcal{R}_{\text{th}}(D_e)$ is the catalogue lookup associated with the selected diameter. A larger $R_{\text{th},e}$ means better insulation and therefore a smaller heat loss for the same temperature difference.

For PE-Xa branches, the catalogue U-value is converted as

$$R_{\text{th}} = \frac{1}{U}, \tag{3.5}$$

In this convention, U is the linear heat-transfer coefficient per metre of pipe. The reciprocal form is used so that the thermal model can treat all pipe families through the same resistance variable R_{th}.

For steel trunk pipes, the concentric PUR insulation model gives

$$R_{\text{th}} = \frac{\ln\left(D_{\text{casing}} / d_{\text{steel}}\right)}{2\pi\,\lambda_{\text{PUR}}}, \tag{3.6}$$

Here, D_{casing} is the outer casing diameter, d_{steel} is the steel carrier-pipe outer diameter, and λ_{PUR} is the thermal conductivity of the PUR insulation. The logarithmic term represents radial heat conduction through the cylindrical insulation layer.

Soil thermal resistance is intentionally neglected. The two pipe families are illustrated in Figure 3.4.

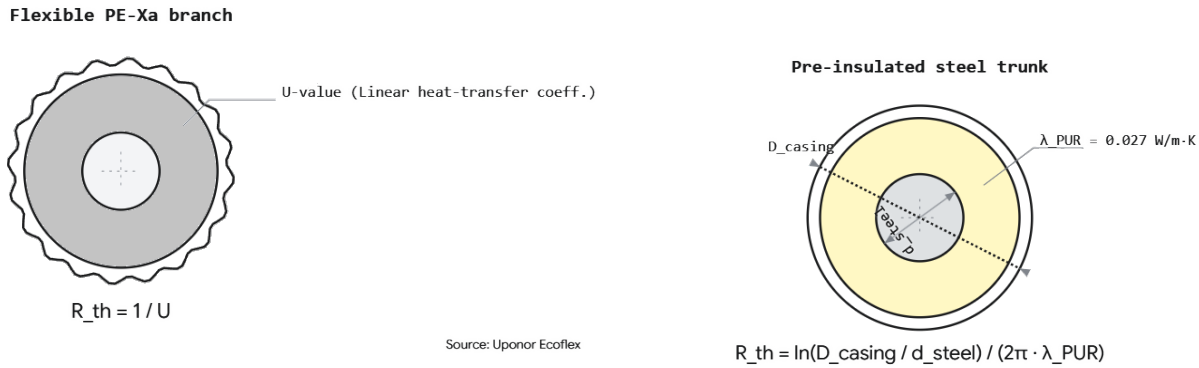

(a) Flexible PE-Xa branch pipe (b) Pre-insulated steel trunk pipe

Figure 3.4. Commercial pipe cross-sections used for thermal-resistance assignment.

Table 3.5. Commercial diameters and thermal resistances.

DN	Inner D (m)	Thermal basis	R_{th} (K·m W^{-1})
DN20	0.020	$U = 0.141$	7.092
DN25	0.025	$U = 0.141$	7.092
DN32	0.032	$U = 0.162$	6.173
DN40	0.040	$U = 0.162$	6.173
DN50	0.050	$U = 0.188$	5.319
DN63	0.065	$U = 0.226$	4.425
DN75	0.080	$U = 0.233$	4.292
DN110	0.100	$U = 0.356$	2.809
DN125	0.125	139.7/250	3.430
DN150	0.150	168.3/280	3.001
DN200	0.200	219.1/355	2.845
DN250	0.250	273.0/450	2.946
DN300	0.300	323.9/500	2.559

3.3 Phase 1: Hydraulic Calculation

Phase 1 is the design-state pressure replay performed after commercial diameters have been selected. The branch mass flows are still the Phase 0 design flows, so this stage is not the final hydraulically balanced operating state. Its role is to turn the sized network into a pressure-drop layer and to identify the path that sets the first pump-head estimate.

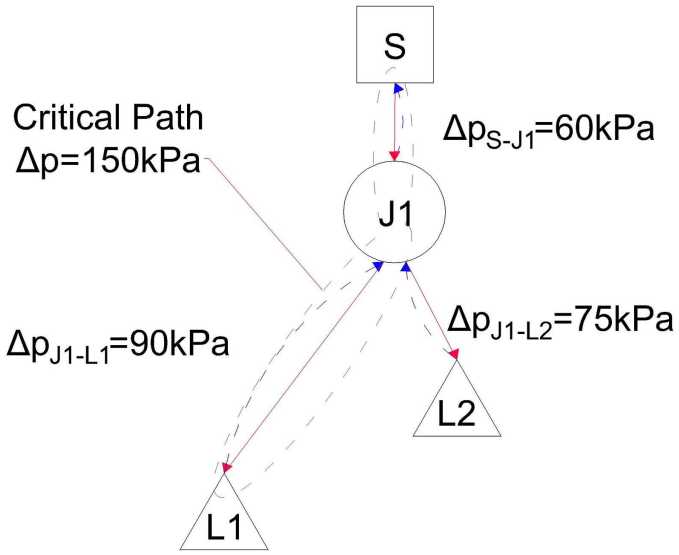

Figure 3.5. Illustrative Phase 1 pressure-drop layer and critical-path identification on the running example.

For each pipe, the pressure drop is evaluated from the current design flow, the selected commercial diameter, the fluid properties, and the minor-loss coefficient. The detailed Darcy–Weisbach relation and friction-factor model are documented in Appendix 2, Section 2.2. Figure 3.5 uses illustrative pressure-drop values to show the selection logic; the actual values used in a simulation are those exported by the Phase 1 replay.

The critical path is the source-to-load route with the largest accumulated pressure drop. In the illustrative example, the route S–$J1$–$L1$ has the largest cumulative pressure drop and is therefore highlighted as the critical path. Every other load is checked in the same way by traversing its unique path through the radial graph and comparing the accumulated pressure drop.

The initial pump head used to start Phase 2, denoted H_{initial}, is then derived from the real critical loop in the exported calculation: the red source-to-load path, the blue reverse path, and the heat-exchanger pressure drop at the terminal load are added. A Phase 1 snapshot of pipe mass flow, velocity, and pressure drop is kept so that this pump-head choice remains traceable when it is compared with the final balanced operating point.

3.4 Phase 2: Pressure Balancing

After commercial pipe sizing, the actual hydraulic operating point is recomputed because the design-flow field is generally no longer compatible with the discrete branch resistances. For a prescribed pump head, Phase 2 asks whether all load loops can be balanced under the same pump pressure. The adjustable quantities are the terminal load-branch velocities on the supply tree, while the full pipe-pair flow field is recovered automatically by upstream aggregation.

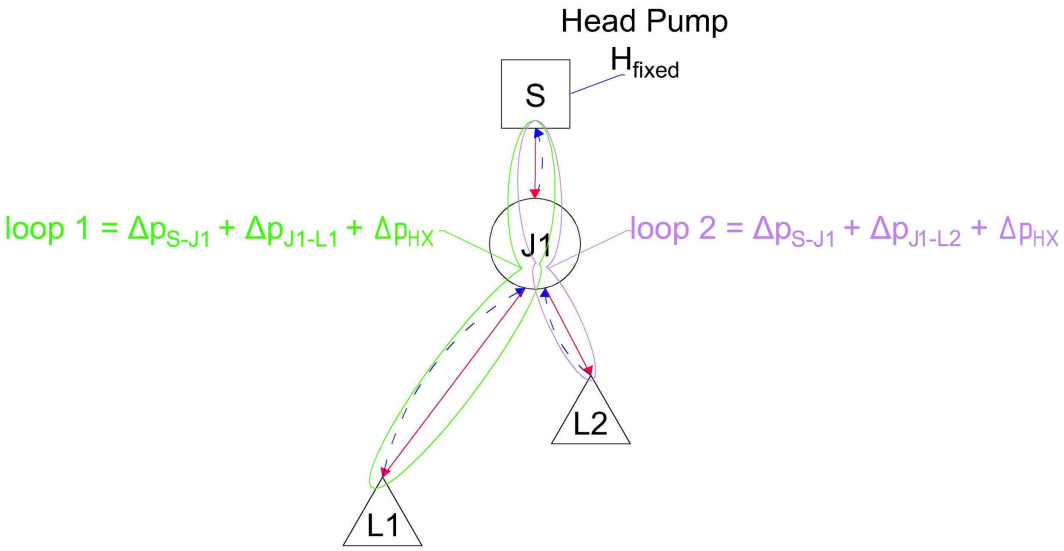

Figure 3.6. Illustrative Phase 2 pressure-balancing loops on the running example. The terminal velocities are adjusted until every load loop has the same pressure drop as the prescribed pump head H_{fixed}.

3.4.1 Unknown Vector and Loop Residuals

For a network with N_L load nodes, the unknown vector is

$$\mathbf{v} = [v_1, \ldots, v_{N_L}]^{\mathsf{T}} \tag{3.7}$$

For each trial vector, the corresponding load mass flows are

$$\dot{m}_i(\mathbf{v}) = \rho_{\text{sup}} A_i v_i, \tag{3.8}$$

and these are propagated upstream through the radial graph to reconstruct the complete pipe-pair flow field. This reconstruction uses the same bottom-up aggregation as Section 3.2.1, but replaces the design load flows by the trial terminal flows $\dot{m}_i(\mathbf{v})$.

In the running example of Fig. 3.6, the unknown vector is simply $\mathbf{v} = [v_{L1}, v_{L2}]^{\mathsf{T}}$. These two terminal velocities set the flows in $J1 \to L1$ and $J1 \to L2$, while the shared upstream branch $S \to J1$ carries their sum. The two pressure loops therefore share one upstream branch but have different terminal branches.

Using this reconstructed flow field, the full-loop pressure drop associated with load i is

$$\Delta p_{\text{loop},i}(\mathbf{v}) = \Delta p_{\text{HX},i} + \sum_{e \in \mathcal{P}_i^s} \Delta p_e(\mathbf{v}) + \sum_{e \in \mathcal{P}_i^r} \Delta p_e(\mathbf{v}). \tag{3.9}$$

Here, $\Delta p_{\text{loop},i}$ is the total pressure drop of the closed supply–return path serving load i, including the substation pressure drop $\Delta p_{\text{HX},i}$. The sets $\mathcal{P}_i^s$ and $\mathcal{P}_i^r$ are respectively the supply and return pipe paths connected to that load, and each Δp_e is the hydraulic loss of one pipe segment.

For a fixed pump head H_{fixed}, the balancing residual is defined as

$$F_i(\mathbf{v}) = \Delta p_{\text{loop},i}(\mathbf{v}) - H_{\text{fixed}}, \qquad i = 1, \ldots, N_L. \tag{3.10}$$

Thus, F_i is the remaining pressure mismatch for load loop i: positive values mean that the loop requires more head than the pump provides, while negative values mean that the same pump head would over-drive that loop.

Hence, the hydraulic balance problem is to find $\mathbf{v}$ such that all loop residuals vanish simultaneously. In Fig. 3.6, this means that the algorithm adjusts v_{L1} and v_{L2} until the pressure drop of Loop 1 and Loop 2 both match the pump head imposed at the source. The residual evaluation uses the water-property convention documented in Appendix 2: density is held fixed within a hydraulic solve, while viscosity is evaluated from the available pipe temperature field and falls back to reference values before the thermal replay is available.

3.4.2 Nonlinear Least-Squares Solution

The residual system is solved as a bounded nonlinear least-squares problem,

$$\min_{\mathbf{v}} \ \Phi(\mathbf{v}) = \frac{1}{2} \sum_{i=1}^{N_L} F_i(\mathbf{v})^2, \qquad v_{\min} \le v_i \le v_{\max}, \tag{3.11}$$

using the trust-region reflective algorithm available in the SciPy scientific computing ecosystem [46, 47]. The SciPy stopping tolerances for the least-squares iteration are kept at their documented default values, namely `xtol`$= 10^{-8}$, `ftol`$= 10^{-8}$, and `gtol`$= 10^{-8}$. These defaults are sufficiently small for the present pressure-balancing problem; tightening them did not change the converged network results. In physical terms, the solver varies the terminal velocities shown in Fig. 3.6 until the loop-pressure mismatches are as small as possible within the velocity bounds. The initial guess is obtained from the design mass flows and the terminal pipe areas:

$$v_{i,0} = \frac{\dot{m}_{\text{design},i}}{\rho_{\text{sup}} A_i}. \tag{3.12}$$

The unknown velocities are constrained by a lower and an upper bound,

$$v_{\min} \le v_i \le v_{\max}, \qquad v_{\min} = 0.05 \ \text{m s}^{-1}, \qquad v_{\max} = 2.5 \ \text{m s}^{-1}. \tag{3.13}$$

The lower bound prevents numerically degenerate near-zero velocities, while the upper bound encodes the maximum admissible flow velocity in distribution pipes. Both bounds are configurable solver parameters. The upper bound is set to 2.5 m s^{-1}, the recommended maximum continuous water speed for flexible district-heating pipes given in the manufacturer catalogue [44]; a more conservative ceiling such as 1.5 m s^{-1} may be used when lower pumping noise and larger erosion margins are required, at the cost of a smaller feasible operating envelope.

The maximum number of function evaluations is chosen as

$$N_{\text{fev,max}} = \max\left(N_{\text{iter}} N_L, \ 200\right), \tag{3.14}$$

Here N_{iter} is the configured iteration allowance per load. This scaling gives larger networks

proportionally more residual evaluations while preserving a minimum allowance for small cases.

After the numerical solve returns, the maximum absolute loop residual is evaluated as

$$r_{\max} = \max_i |F_i(\mathbf{v}^*)|, \tag{3.15}$$

and the solution is declared converged only if it falls below a convergence tolerance ε_p. In this pressure-balancing context, convergence means that the remaining loop-pressure mismatch is smaller than the prescribed tolerance, so the residual is negligible at the scale of the hydraulic model. Because the loop pressure drops $\Delta p_{\text{loop},i}$ are themselves of the order of the pump head H, a fixed absolute tolerance is not scale-invariant: a value such as 0.01 Pa corresponds to a relative accuracy of about 10^{-7} for a loop head of order 10^5 Pa, far below both the physical resolution of the problem and the numerical conditioning of the coupled branch system. An over-strict tolerance causes the least-squares solve to be reported as non-converged even when the recovered flow field is physically balanced. The tolerance is therefore defined relative to the prescribed pump head,

$$\varepsilon_p = \max(\varepsilon_{\text{abs}},\ \eta_{\text{rel}}\, H)\,, \qquad \varepsilon_{\text{abs}} = 0.01 \text{ Pa}, \qquad \eta_{\text{rel}} = 10^{-3}, \tag{3.16}$$

so that the convergence threshold scales with the head (here 0.1 % of H) while retaining the small absolute value only as a floor for degenerate low-head cases. The balanced solution is converged when $r_{\max} \le \varepsilon_p$. Conversely, if $r_{\max} > \varepsilon_p$, the candidate operating point is rejected rather than forced by relaxing the numerical tolerance. In engineering terms, this means that the current combination of pump head, pipe sizes, velocity bounds, and design-flow assumptions does not provide a converged pressure-balanced solution; the network design or control strategy must then be revised.

Once convergence is achieved, the balanced terminal velocities are replayed through the same graph-based reconstruction to obtain the final branch-wise flow field. At this stage, the network is pressure-balanced for the prescribed pump head.

3.5 Phase 3: Thermal Propagation Model

The thermal model takes the balanced pipe-pair mass-flow field as a fixed input and propagates temperatures using the convention of Fig. 3.1. The red pass carries hot water from the source to the loads; the blue pass carries cooled water back to the source and mixes streams at junctions. No global nodal temperature system is assembled. The underlying water properties are collected in Appendix 2.

3.5.1 Pipe Heat Loss and Outlet Temperature

Figure 3.7 is the pipe-level thermal model. The water enters pipe e at $T_{\text{in},e}$, exchanges heat with the ground at T_g through the thermal resistance $R_{\text{th},e}$, and leaves at $T_{\text{out},e}$. For a pipe of length L_e, this gives

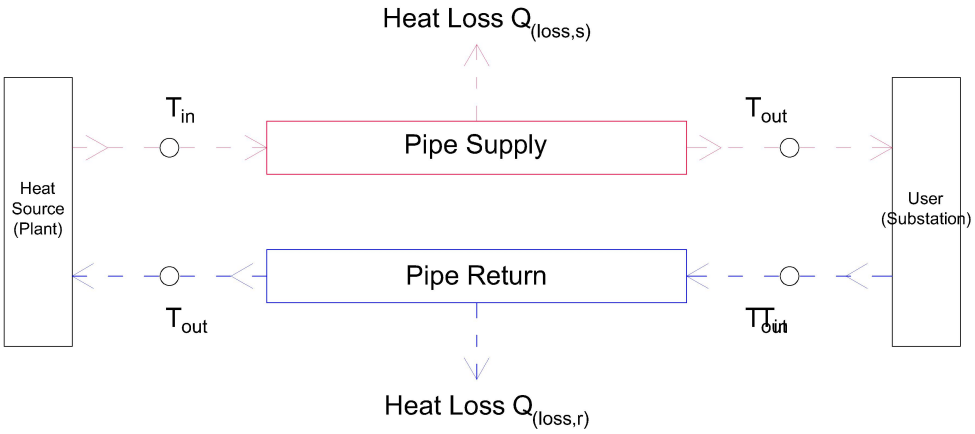

Figure 3.7. Pipe heat-loss schematic used in the thermal propagation model. The local flow direction defines $T_{\text{in},e}$ and $T_{\text{out},e}$.

$$T_{\text{out},e} = T_g + \left(T_{\text{in},e} - T_g\right) \exp\left(\frac{-L_e}{R_{\text{th},e}\, \dot{m}_e\, c_p}\right), \tag{3.17}$$

Here, $T_{\text{out},e}$ is the outlet temperature after pipe segment e, $T_{\text{in},e}$ is the inlet temperature, T_g is the ground temperature, L_e is the pipe length, $\dot{m}_e$ is the mass flow rate, and c_p is the specific heat capacity of water. The denominator $R_{\text{th},e}\dot{m}_e c_p$ shows the two mechanisms that reduce cooling: better insulation through larger thermal resistance and larger thermal inertia through higher mass flow.

The heat loss shown in the schematic is the enthalpy drop between inlet and outlet:

$$Q_{\text{loss},e} = \dot{m}_e\, c_p \left(T_{\text{in},e} - T_{\text{out},e}\right). \tag{3.18}$$

Here, $Q_{\text{loss},e}$ is the thermal power lost by pipe segment e during the considered pass. It is computed from the water enthalpy decrease, so it is directly linked to the temperature drop predicted by Eq. (3.17).

If $\dot{m}_e = 0$, $L_e = 0$, or $T_{\text{in},e} \leq T_g$, the pipe is treated as lossless for that pass: $Q_{\text{loss},e} = 0$ and $T_{\text{out},e} = T_{\text{in},e}$.

3.5.2 Network-Wide Thermal Propagation

After the pipe-level outlet temperature has been computed with Eq. (3.17), the network temperature field is obtained by following the two real flow directions shown in Fig. 3.8: red for supply and blue for return.

Supply propagation. Hot water starts from the source. At a junction, the same node temperature is used as the inlet temperature of every outgoing child branch. Along each pipe, for example $S \rightarrow J1$ or $J1 \rightarrow L1$, the temperature decreases according to the pipe heat-loss relation

3.5.2 Network-Wide Thermal Propagation

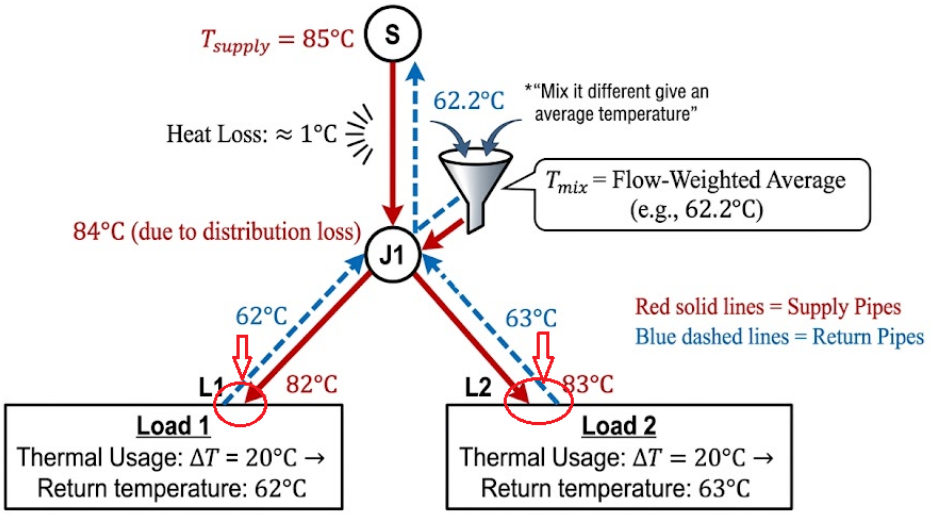

Figure 3.8. Network-wide thermal propagation on the running example. Red arrows show supply-temperature propagation, while blue arrows show return-temperature mixing and propagation back to the source.

of Eq. (3.17).

Return mixing. At each load, the return temperature is obtained by subtracting the user's design temperature drop, $T_{\text{return},i} = T_{\text{supply},i} - \Delta T_{\text{design},i}$. The cooled streams then flow back toward the source. When several return streams meet at a junction such as $J1$, the mixed temperature is the flow-weighted average

$$T_{\text{return},n} = \frac{\sum_{e \in \mathcal{R}_n} \dot{m}_e \, T_{\text{out},e}}{\sum_{e \in \mathcal{R}_n} \dot{m}_e}. \tag{3.19}$$

which conserves enthalpy at the junction. The mixed stream then continues toward the source and loses heat on each upstream return pipe using the same pipe relation, Eq. (3.17). The resulting source return temperature is used later in the final energy balance.

3.6 Phase 4: Hydraulic Correction

The pressure-balancing solver equalizes the loop pressure drops for a given pump head, but it does not automatically mean that every consumer receives enough water. The hydraulic correction therefore checks the actual mass flow at the load nodes, and raises the pump head only if at least one consumer is undersupplied.

The hydraulic service requirement is checked at the load-node level:

$$\dot{m}_{\text{actual,bal},i} \geq \dot{m}_{\text{design},i} \qquad \forall i = 1, \ldots, N_L. \tag{3.20}$$

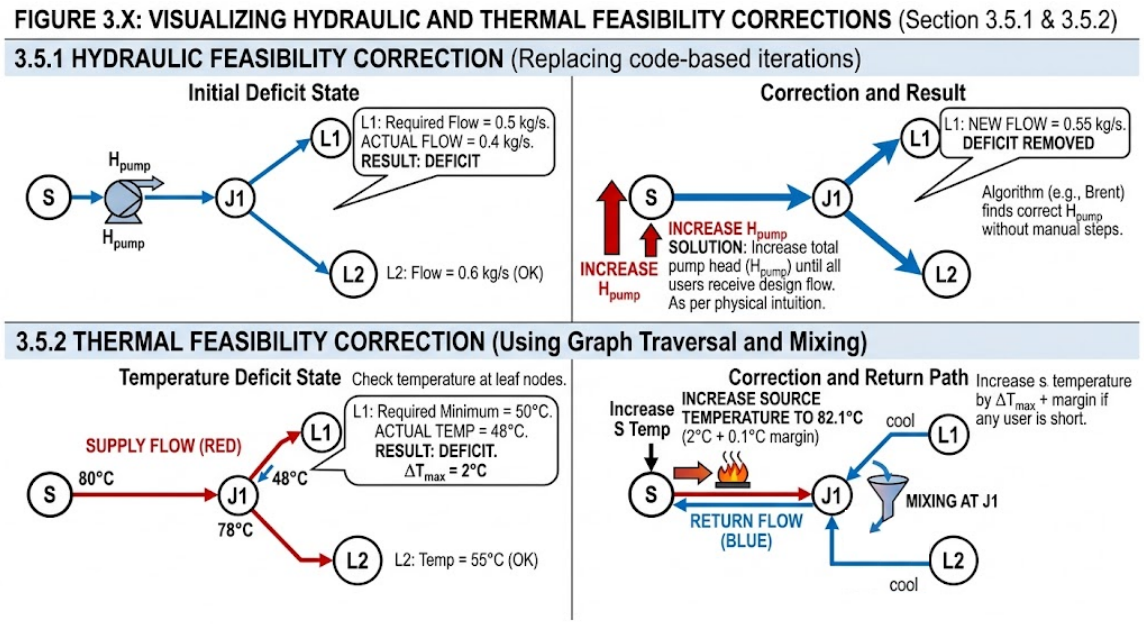

Figure 3.9. Phase 4 and Phase 5 corrections on the running example. The upper panels show the pump-head correction used to remove a mass-flow deficit; the lower panels show the source-temperature correction used to remove a delivered-temperature deficit.

Equivalently, the worst hydraulic deficit is defined as

$$\psi(H) = \max_i \left(\dot{m}_{\text{design},i} - \dot{m}_{\text{actual,bal},i}(H)\right). \tag{3.21}$$

The pump head satisfies the hydraulic requirement when

$$\psi(H) \leq 0. \tag{3.22}$$

Phase 4 starts at the Phase 1 initial pump head H_{initial}. If no deficient load is detected, this pump head is kept. Otherwise, as illustrated in Fig. 3.9, the pump head is increased with Brent's one-dimensional root-finding method until the worst deficit disappears [48]. The selected value is denoted H^*, and the hydraulic state is replayed once more at this pump head before the workflow proceeds to Phase 5.

3.7 Phase 5: Thermal Correction

Phase 5 starts once the flow field satisfies the load-flow requirement. The thermal model of Section 3.5.2 is replayed. The correction only checks whether the red supply temperature at every load remains above the required minimum; the blue return field is updated by the same replay for the final energy balance.

The delivered-temperature requirement is checked at the load nodes:

$$T_{\text{supply},i} \geq T_{\text{supply,min}} \qquad \forall i = 1, \ldots, N_L. \tag{3.23}$$

If every load satisfies this requirement, the current source temperature is kept. Otherwise,

the largest temperature deficit is computed:

$$\Delta T_{\max} = \max_i \left(T_{\text{supply,min}} - T_{\text{supply},i} \right). \tag{3.24}$$

The source supply temperature is then increased by this deficit plus a small numerical margin:

$$T^{k+1}_{\text{source,supply}} = T^{k}_{\text{source,supply}} + \Delta T_{\max} + 0.1^{\circ}\text{C}. \tag{3.25}$$

After each increase, the thermal replay is repeated using the branch flows obtained after Phase 4. The correction changes only the source temperature; it does not reopen the pressure-balancing problem. The final operating point is retained only after it satisfies both the load-flow requirement and the minimum delivered-temperature requirement.

3.8 Phase 6: Final State Verification and Energy Balance

After H^* and $T^*_{\text{source,supply}}$ have been selected, an explicit first-law consistency check is applied to the final operating point.

The source heat injection is computed from the final source mass flow, the specific heat capacity of water c_p, and the final source supply–return temperature difference:

$$Q_{\text{source}} = \dot{m}_{\text{source}} c_p \left(T_{\text{source,supply}} - T_{\text{source,return}} \right). \tag{3.26}$$

The useful heat delivered to consumers is evaluated by summing, over all load nodes, the product of the final actual balanced load mass flow, the specific heat capacity of water c_p, and the local supply–return temperature difference:

$$Q_{\text{delivered}} = \sum_{i=1}^{N_L} \dot{m}_{\text{local},i} c_p \left(T_{\text{supply},i} - T_{\text{return},i} \right). \tag{3.27}$$

The total pipe heat loss is obtained by summing the heat-loss term over both directions of every pipe pair at the final operating point:

$$Q_{\text{loss}} = \sum_{e \in \mathcal{P}^s \cup \mathcal{P}^r} Q_{\text{loss},e}. \tag{3.28}$$

The final energy imbalance is then evaluated as an implementation sanity check, verifying that the sum of heat delivered and thermal distribution losses equals the source heat injection [49]:

$$\Delta Q = \left| Q_{\text{source}} - Q_{\text{delivered}} - Q_{\text{loss}} \right|. \tag{3.29}$$

The source contribution must be strictly positive. A relative imbalance tolerance of 1% is adopted as the diagnostic threshold for this consistency check. This value is chosen as a practical intermediate threshold: a much smaller tolerance would be unnecessarily strict for a pre-design model, whereas a much larger tolerance would allow an energy mismatch too large to indicate a reliable implementation.

$$\frac{\Delta Q}{Q_{\text{source}}} \leq 0.01. \tag{3.30}$$

If this diagnostic check is violated, the result is treated as an implementation or bookkeeping error rather than as a physically meaningful operating point. In that case, the exported values must be inspected before the final state is reported.

3.9 Economic and Financial Model

The economic model estimates the network-side capital expenditure (CAPEX), operating expenditure (OPEX), and net present cost (NPC) of the district-heating network. The boundary of the calculation is restricted to the distribution network. Heat-generation plant investment, such as geothermal production equipment or boiler plant CAPEX, is excluded. Consumer-side equipment, including building substations and heat exchangers, is also excluded from the CAPEX boundary. These devices are assumed to be owned and financed by the consumers, since comparable building-side heat-transfer equipment would also be required in an individual local-heating alternative. The reported CAPEX should therefore be read as the additional investment required to build the heat-distribution network relative to conventional local heating, not as a full building-conversion cost. The net present cost aggregates the initial investment and all discounted future operating costs into a single indicator:

$$C_{\text{NPV}} = C_{\text{CAPEX}} + \sum_{y=1}^{N_{\text{proj}}} \frac{C_{\text{OPEX},y}}{(1+r)^y}. \tag{3.31}$$

From the net present cost, the levelized cost of heat distribution is

$$\text{LCOH} = \frac{C_{\text{NPV}}}{Q_{\text{design}}\, h_{\text{op}}\, N_{\text{proj}}}, \tag{3.32}$$

expressed in EUR/MWh: the net present cost spread over the heat delivered during the project horizon. It excludes heat-generation costs. The capital expenditure (CAPEX) and the LCOH are the two indicators used to compare scenarios throughout this work.

The following subsections detail the CAPEX and OPEX components that enter these expressions.

3.9.1 CAPEX Estimation Model

Pipe and installation costs follow the nPro technology-cost model for low-temperature networks with insulated pipes [50]. These cost correlations are not arbitrary placeholders: nPro derives them from the technology catalogue of the KEA Baden-Württemberg energy agency [51], and they distinguish pipe material from civil works and scale with nominal diameter and installation environment. nPro provides, per metre of trench and for the supply–return pipe pair, a pipe (material) cost that is independent of the environment,

$$c_{\text{pipe}}(D_e) = 0.0033\,\text{DN}_e^2 + 0.6843\,\text{DN}_e + 54.033 \quad \text{[EUR/m]}, \tag{3.33}$$

and an environment-dependent total (pipe + installation) cost that is linear in the nominal diameter $\mathrm{DN}_e = 1000\, D_e$ (in mm),

$$c_{\text{total}}(D_e, a) = \alpha_a\, \mathrm{DN}_e + \beta_a, \qquad (\alpha_a, \beta_a) = \begin{cases} (4.122,\ 446.29) & a = \text{rural/unpaved}, \\ (5.854,\ 633.43) & a = \text{suburban}, \\ (7.956,\ 860.92) & a = \text{urban/paved}. \end{cases} \tag{3.34}$$

The installation (civil-works) cost per pair is the difference $c_{\text{total}}(D_e, a) - c_{\text{pipe}}(D_e)$, which is strongly location-dependent: paved urban streets are far more expensive than open rural corridors. The environment is assigned per segment from the network geometry (paved streets for the on-campus distribution, unpaved/rural for the source feeder crossing open ground). Valve cost is represented by an average number of valves per pipe pair, a network-level simplification of the assumption of approximately ten valves per loop, adjustable for a more detailed layout.

The CAPEX model may be written as

$$C_{\text{CAPEX}} = C_{\text{trench}} + C_{\text{pipe}} + C_{\text{valve}} + C_{\text{pump}}, \tag{3.35}$$

with

$$C_{\text{trench}} = \sum_e L_e \left[c_{\text{total}}(D_e, a_e) - c_{\text{pipe}}(D_e) \right], \tag{3.36}$$

$$C_{\text{pipe}} = \sum_e L_e\, c_{\text{pipe}}(D_e), \tag{3.37}$$

$$C_{\text{valve}} = \sum_e N_{\text{valve},e}\, c_{\text{valve}}, \tag{3.38}$$

and

$$C_{\text{pump}} = P_{\text{pump}}\, c_{\text{pump}}, \tag{3.39}$$

Here, C_{trench} is the civil-works cost, C_{pipe} is the pipe-material cost, C_{valve} is the balancing-valve cost, and C_{pump} is the pump capital cost. The selected commercial diameter D_e is reused from the pipe-sizing stage so that the economic model prices the same buildable network that is checked hydraulically and thermally.

The valve-count parameter is a modeling simplification. The present formulation uses an average number of valves per pipe pair, whereas the engineering interpretation may be linked back to an assumed average number of valves per hydraulic loop.

3.9.2 OPEX Estimation Model

The OPEX model includes electricity consumption for pumping and network-side maintenance. Pump electricity is calculated from the required pump power, annual operating hours, and electricity price. Maintenance is represented either as an annualized fraction of CAPEX or as scheduled fixed costs occurring every year, every five years, and every ten years. The case-study

calculations use the annualized fraction method; the scheduled-cost fields are kept only as an inactive alternative parameterization. The annualized maintenance term is interpreted as covering pipe inspection, pipe repair, pump maintenance, and other routine network-side interventions; consumer-owned substations and heat exchangers remain outside the cost boundary.

The annual pumping-energy demand is estimated as

$$E_{\mathrm{pump},y} = P_{\mathrm{pump}}\, h_{\mathrm{op}}, \tag{3.40}$$

and the corresponding electricity cost is

$$C_{\mathrm{elec},y} = E_{\mathrm{pump},y}\, c_{\mathrm{elec}}. \tag{3.41}$$

The maintenance cost in year y is evaluated as

$$C_{\mathrm{maint},y} = \begin{cases} C_{\mathrm{maint,annual}} + \mathbb{I}_{5|y} C_{\mathrm{maint,5y}} + \mathbb{I}_{10|y} C_{\mathrm{maint,10y}}, & \text{if the scheduled-cost option is used,} \\ f_{\mathrm{maint}}\, C_{\mathrm{CAPEX}}, & \text{otherwise.} \end{cases} \tag{3.42}$$

Accordingly, the annual operating expenditure is

$$C_{\mathrm{OPEX},y} = C_{\mathrm{elec},y} + C_{\mathrm{maint},y}. \tag{3.43}$$

Here, $\mathbb{I}_{5|y}$ and $\mathbb{I}_{10|y}$ are indicator functions for the five-year and ten-year maintenance events.

This parameterization is intentionally simple. Maintenance may be represented either as an annualized fraction of total CAPEX or as explicit fixed maintenance events at annual, 5-year, and 10-year intervals. In the present case study, the first option is active: maintenance is charged as $f_{\mathrm{maint}} C_{\mathrm{CAPEX}}$ each year, with $f_{\mathrm{maint}} = 1\,\%$ following the nPro economic-calculation convention for district-heating networks [52]. This keeps the model consistent with the supervisor's recommendation while preserving a compact computational structure.

Note on maintenance accounting. While Eq. (3.42) allows for explicit scheduled maintenance events ($C_{\mathrm{maint,annual}}$, $C_{\mathrm{maint,5y}}$, $C_{\mathrm{maint,10y}}$), this option is not used in the case-study parameter table (Table 3.6). The underlying scheduled inputs remain zero only to disable the alternative event-based representation and prevent double-counting. The economic model instead uses the continuous annualized fraction method ($f_{\mathrm{maint}} = 1\,\%$ of CAPEX [52]) to cover routine operational repairs, pipe inspections, and component replacements consistently across all evaluated scenarios.

Table 3.6 lists only the default numerical values used in the case study. The corresponding symbols, units, and definitions are collected in the nomenclature (Table 1.1).

Thus, the zero scheduled-maintenance inputs are not interpreted as zero maintenance cost. They only disable the alternative event-based representation. The maintenance cost reported in the economic results is $f_{\mathrm{maint}} \times C_{\mathrm{CAPEX}}$ for every year of the case study; if explicit scheduled values are activated in a sensitivity run, the fraction-based term is not applied.

Finally, a subset of the economic inputs is not drawn from a specific catalogue or data source but assumed as engineering and financial defaults, to be replaced by project-specific quotations

Table 3.6. Default financial-model inputs.

Input	Default value	Basis
$c_{\text{total}}(D_e, a_e)$	nPro formulas	Cost correlations given above [50]
c_{pump}	500 EUR/kW	Assumed pre-design unit cost
c_{valve}	250 EUR/valve	Assumed fixed unit cost
$N_{\text{valve,pair}}$	10 valves/pair	Bookkeeping proxy from the valves-per-loop rule
η_{pump}	0.70	Representative pump efficiency [53]
c_{elec}	0.25 EUR/kWh	Consistent with EU household electricity prices [2]
h_{op}	4000 h/year	Assumed equivalent full-load operation
Maintenance accounting	annualized	Uses $f_{\text{maint}} C_{\text{CAPEX}}$ rather than scheduled events
f_{maint}	1 %/year	Active maintenance method [52]
$C_{\text{maint,annual}}$, $C_{\text{maint,5y}}$, $C_{\text{maint,10y}}$	0 EUR	Disabled scheduled-cost option
r	4 %/year	Assumed discount rate
N_{proj}	30 years	Assumed evaluation horizon

and investor policies in a detailed design phase: the pump unit cost (500 EUR/kW), the valve unit cost (250 EUR/unit), the valve-count proxy (10 valves per pipe pair), the equivalent full-load operating hours (4000 h/yr), and the financial evaluation parameters, namely a 4 % discount rate and a 30-year project horizon. These are deliberately treated as placeholders; because every scenario shares the same cost basis, the *relative* comparison between scenarios, which is the outcome of interest here, is insensitive to their absolute values.

The pump efficiency ($\eta_{\text{pump}} = 0.70$) is used in the same spirit as a representative pre-design conversion efficiency, not as a selected commercial pump guarantee.

4 Case Study and Results

4.1 Impact of Adding New Consumers

This sensitivity analysis tests HeatNet on incremental network extensions rather than on a single fixed layout. The practical question is whether an existing radial network can serve new consumers while still satisfying hydraulic and thermal service constraints.

Each case modifies the input graph before simulation. The same workflow is then reapplied to rebuild the topology, recompute design flows and pipe sizes, balance the pressure field, correct the operating point, and verify the final state. This confirms that the implementation is topology-aware rather than hard-coded for the reference network.

4.1.1 Reference Data and Expansion-Scenario Evaluation

The scenario analysis is performed on a reference radial district-heating network defined directly from the original-topology HeatNet CSV inputs. The base case contains one source node, four junction nodes ($J1$, $J2$, $J4$, $J5$), and nine load nodes ($L1$, $L2$, $L3$, $L4$, $L7$, $L8$, $L9$, $L10$, $L11$). The source input conditions are a supply temperature of 50°C and an initial source pressure of 300,000 Pa. The reference topology is shown in Fig. 4.1.

Tables 4.1–4.2 give the corresponding load, and pipe-pair data in a reader-oriented form. Heat demands are shown in kW for compactness; the CSV files store the same values in W. The design mass-flow column is computed from $Q/(c_p \Delta T_{\text{design}})$ and is included to make the physical meaning of the load inputs explicit.

Table 4.1. Base-case load input data.

Load node	Heat demand Q (kW)	ΔT_{design} (K)	c_p (J kg^{-1} K^{-1})	$\dot{m}_{\text{design}}$ (kg s^{-1})
L1	45.0	20	4180	0.538
L2	50.0	20	4180	0.598
L3	36.0	23	4180	0.374
L4	55.0	22	4180	0.598
L7	45.0	20	4180	0.538
L8	50.0	20	4180	0.598
L9	47.0	21	4180	0.535
L10	48.0	22	4180	0.522
L11	52.0	21	4180	0.592

Three modified input datasets are compared with the reference case. Case 1 adds load $L5$ at junction $J2$; Case 2 adds $L5$ at $J2$ and $L6$ at $J4$; Case 3 also inserts a new junction $J6$ downstream of $J2$ and supplies $L4$ and $L5$ through this junction. The three topology variants are compared in Fig. 4.2.

For each scenario, the complete HeatNet workflow is executed again: the modified graph is reconstructed, the radial topology is validated, design mass flows and commercial pipe diameters

Table 4.2. Base-case pipe-pair connectivity input data (columns read by the solver).

start_node	end_node	length (m)
S	J1	100
S	J2	60
J1	J4	80
J2	J5	100
J1	L1	150
J1	L2	130
J2	L3	100
J2	L4	150
J4	L7	75
J4	L8	90
J5	L9	120
J5	L10	100
J5	L11	80

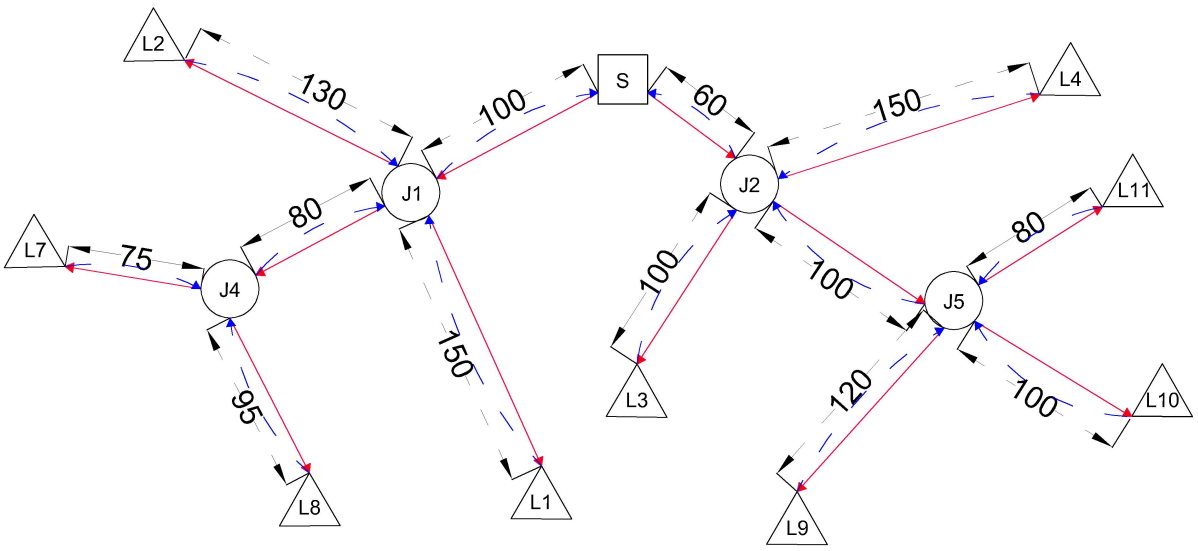

Figure 4.1. Radial supply–return topology of the reference district-heating network. Solid red arrows represent supply pipes (source to loads), dashed blue lines represent return pipes. Square: source node; circles: junction nodes; triangles: load nodes.

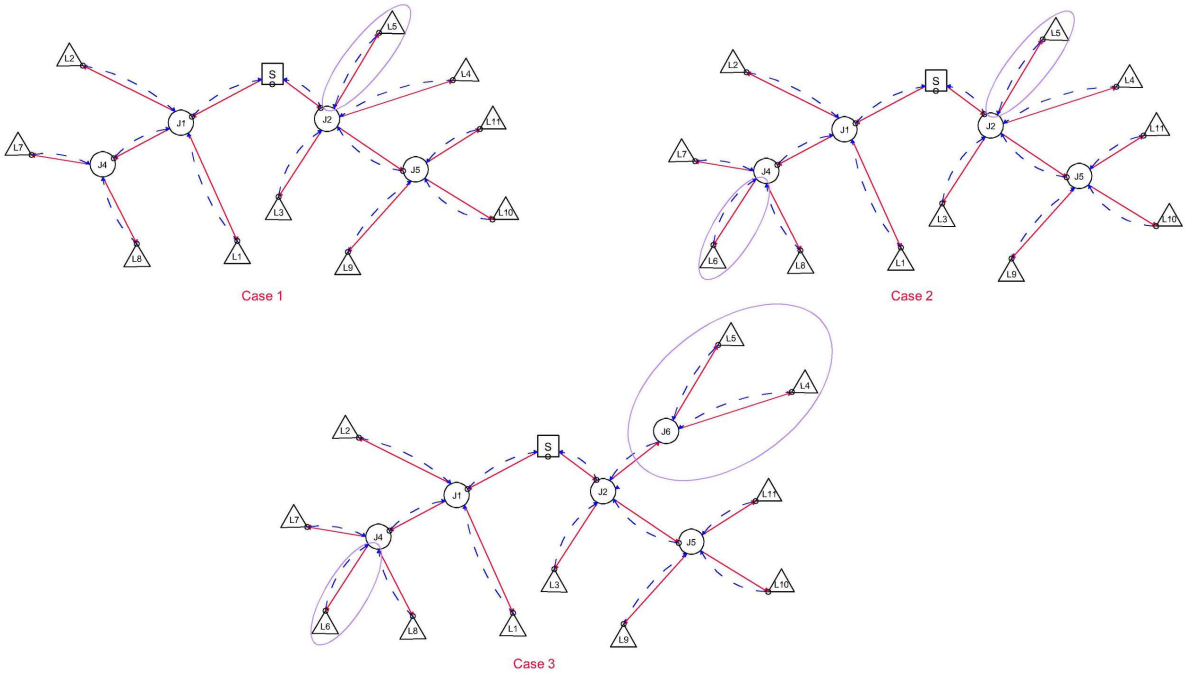

Figure 4.2. Topology comparison for the three expansion scenarios. The highlighted areas indicate the added consumer branches and, in Case 3, the inserted junction $J6$.

are recomputed, and Phases 4–5 are applied. The compact network-level comparison is reported in Table 4.3.

Table 4.3. Reference and expansion-scenario summary.

Scenario	Nodes	Pairs	Q (kW)	Length (m)	P_{pump} (kW)	$v_{\max}$ (m/s)	$T_{\min}$ (°C)	Feasible
Base case	14	13	428.0	1335	1.2	1.01	50.08	Yes
Case 1: add $L5$ at $J2$	15	14	480.0	1435	1.4	1.15	50.08	Yes
Case 2: add $L5$ and $L6$	16	15	534.5	1530	1.6	1.18	50.08	Yes
Case 3: add $J6$, $L5$, and $L6$	17	16	534.5	1555	1.5	1.07	50.08	Yes

All expansion cases remain hydraulically and thermally feasible. The added consumers increase the total heat demand from 428.0 kW to 534.5 kW and the pipe-pair length from 1335 m to 1555 m, while the maximum final velocity stays below 1.2 m/s. The minimum delivered supply temperature remains 50.08°C because Phase 5 raises the source supply temperature to approximately 50.6°C before the final state is exported.

Table 4.4 reports the reference-case terminal load service results only. The corresponding supply and return temperature map is shown in Fig. 4.3.

These results show that HeatNet is not limited to evaluating a fixed network. It supports comparative expansion studies, automatically rebuilds the hydraulic residual system after topology changes, and verifies whether additional consumers remain serviceable.

Table 4.4. Reference-case terminal-load service results.

Load	Q_i (kW)	$\dot{m}_{\mathrm{design}}$ (kg/s)	$\dot{m}_{\mathrm{actual}}$ (kg/s)	Flow margin (%)	T_{supply} (°C)	Feasible
L1	45	0.538	0.688	27.9	≥ 50.08	Yes
L2	50	0.598	0.742	24.0	≥ 50.08	Yes
L3	36	0.374	0.451	20.4	≥ 50.08	Yes
L4	55	0.598	0.707	18.1	≥ 50.08	Yes
L7	45	0.538	0.658	22.3	≥ 50.08	Yes
L8	**50**	**0.598**	**0.598**	**0.1**	≥ 50.08	**Yes**
L9	47	0.535	0.575	7.3	≥ 50.08	Yes
L10	48	0.522	0.632	21.1	≥ 50.08	Yes
L11	52	0.592	0.711	19.9	≥ 50.08	Yes

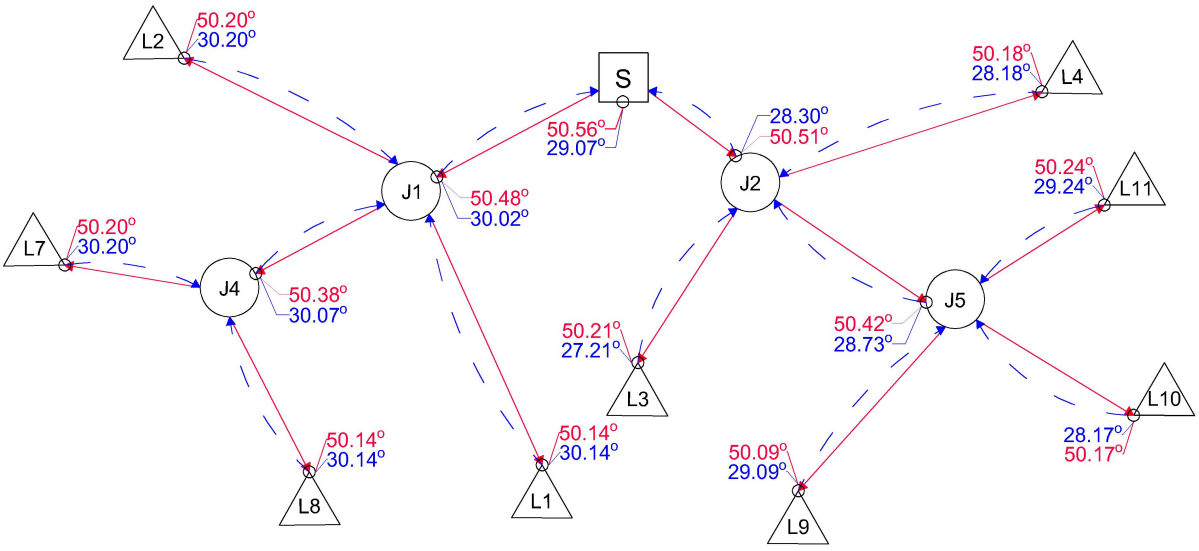

Figure 4.3. Reference-case temperature map after hydraulic and thermal correction. Red labels indicate supply temperatures and blue labels indicate return temperatures at the corresponding nodes.

4.2 Real-World Application: the UMONS Mons Campus

To demonstrate the workflow on a genuine engineering problem rather than a synthetic network, this section applies HeatNet to the heat supply of several buildings of the University of Mons (UMONS) campus. The building footprints and the road network are extracted from OpenStreetMap open data [54], the detailed geographic extraction method is documented in the author's internal internship report [**internship˙report**], and the consumers are taken as the centroids of the aggregated building clusters. Two heat-supply cases are defined once here and then referred to simply as *gas boiler* and *geothermal*. The gas boiler case uses an in-campus boiler, the full peak demand of Table 4.5, and the high-temperature regime 85/65°C. The geothermal case uses a remote well north of the campus, the additional source feeder, a 30 % peak-demand reduction, and the lower temperature regime 70/50°C. The comparison is restricted to the distribution network: it excludes gas purchase, boiler fuel cost, geothermal production-plant investment, carbon costs, and avoided emissions. The corresponding building and candidate heat-source locations are shown in Fig. 4.4.

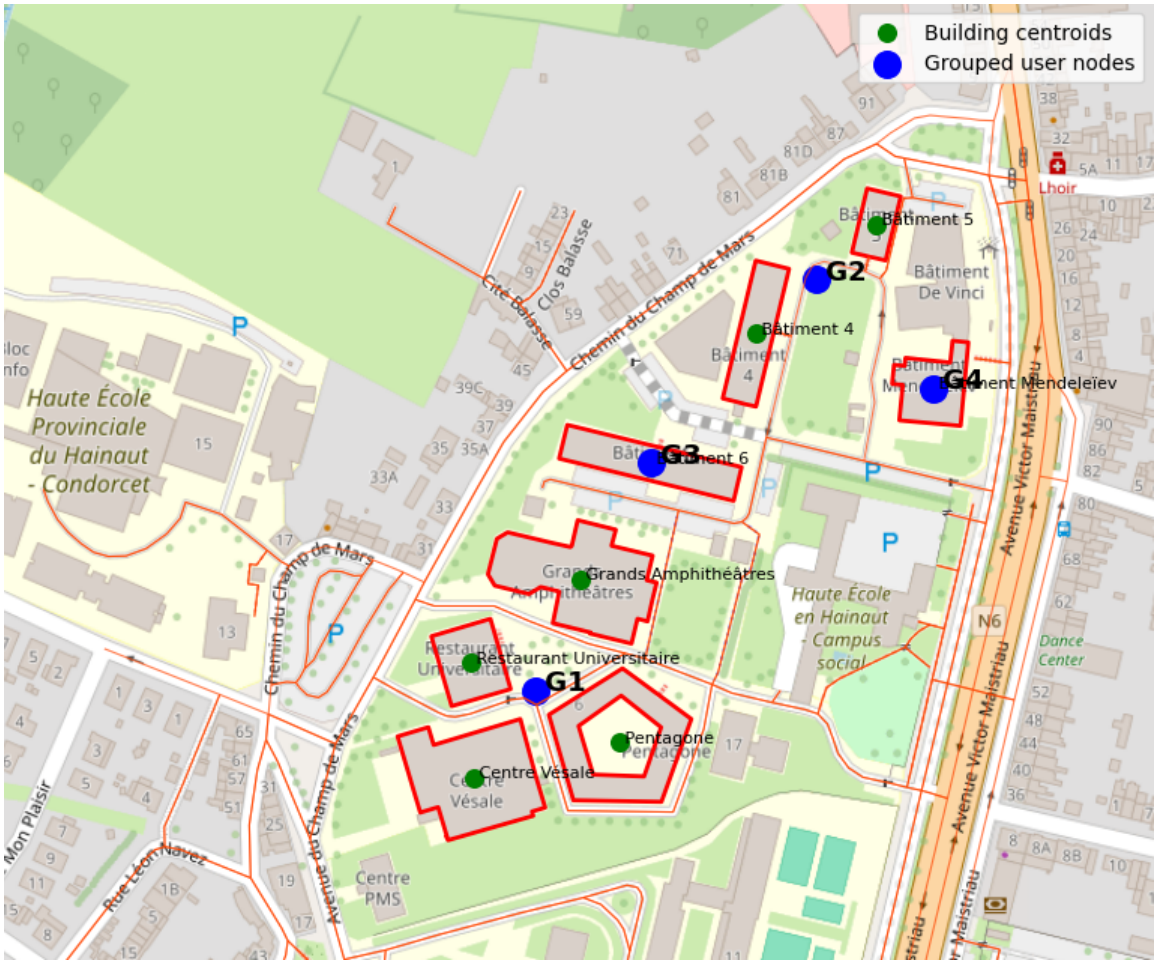

Figure 4.4. Position of the UMONS buildings and heat source used in the real-world case study. The mapped building locations define the consumer clusters represented in the HeatNet model.

4.2.1 Input Data and Demand Sizing

The design heat demand of each consumer is taken directly as the maximum hourly value of its 2022 natural-gas consumption record. The year 2022 is the most recent complete post-Covid record, and sizing for the single worst hour of the year guarantees that the network meets the total demand at the moment all buildings need the most energy. One kilowatt-hour of gas consumed within one hour equals one kilowatt of average power, so the maximum hourly consumption is numerically the peak power; it is treated directly as a heat demand without a boiler-efficiency conversion (a deliberately conservative hypothesis).

The metered consumption peaks are internal UMONS/local CSV input data available to the project rather than a public statistical dataset.

The 2022 natural-gas consumption is metered per cluster. These metered peaks are assigned directly to the four consumer clusters $G1$–$G4$ that correspond to the network connection points, as listed in Table 4.5. Because gas consumption is metered at the cluster level rather than per individual building, the internal pipe lengths connecting each building to its cluster connection point are not known and cannot be verified from the available data. The table also reports the design temperatures and design mass flows ($c_p = 4180\ \mathrm{J\,kg^{-1}K^{-1}}$). The buildings operate at two measured high-temperature regimes (80/60 and 85/65); since a single source must satisfy the strictest one, the source supply temperature is set to 85°C. All groups share a 20 K design temperature difference.

Table 4.5. UMONS demand and design mass flow.

Group	Aggregated buildings	Peak demand (kW)	Supply/return (°C)	$\dot{m}_{\mathrm{design}}$ (kg/s)
$G1$	Restaurant U. + Pentagone + Grands Amphis + Vésale	1635.7	80/60	19.57
$G2$	Bâtiment 4 + Bâtiment 5	671.2	85/65	8.03
$G3$	Bâtiment 6	493.7	85/65	5.91
$G4$	Bâtiment Mendeleïev	448.0	80/60	5.36
Total		3248.7		38.87

Pipe routing. The campus distribution network follows the OpenStreetMap road graph [54] and connects the four consumer clusters through road junctions and short building-connection stubs. This routing layer is the common campus network used by both supply cases. The gas-boiler base case uses only this campus network; the additional remote feeder used by the geothermal what-if is introduced in Section 4.2.3.

4.2.2 Base Scenario: In-Campus Gas Boiler

The gas-boiler case is used as the base scenario because it represents the existing high-temperature heating logic with an in-campus heat source. It uses the full peak demand of Table 4.5, the strictest measured temperature regime 85/65°C, and only the campus distribution network, without the remote geothermal feeder.

The detailed terminal-load results of this base case are reported in Table 4.6. The diameter column gives the commercial diameter of the terminal pipe pair connected to each consumer

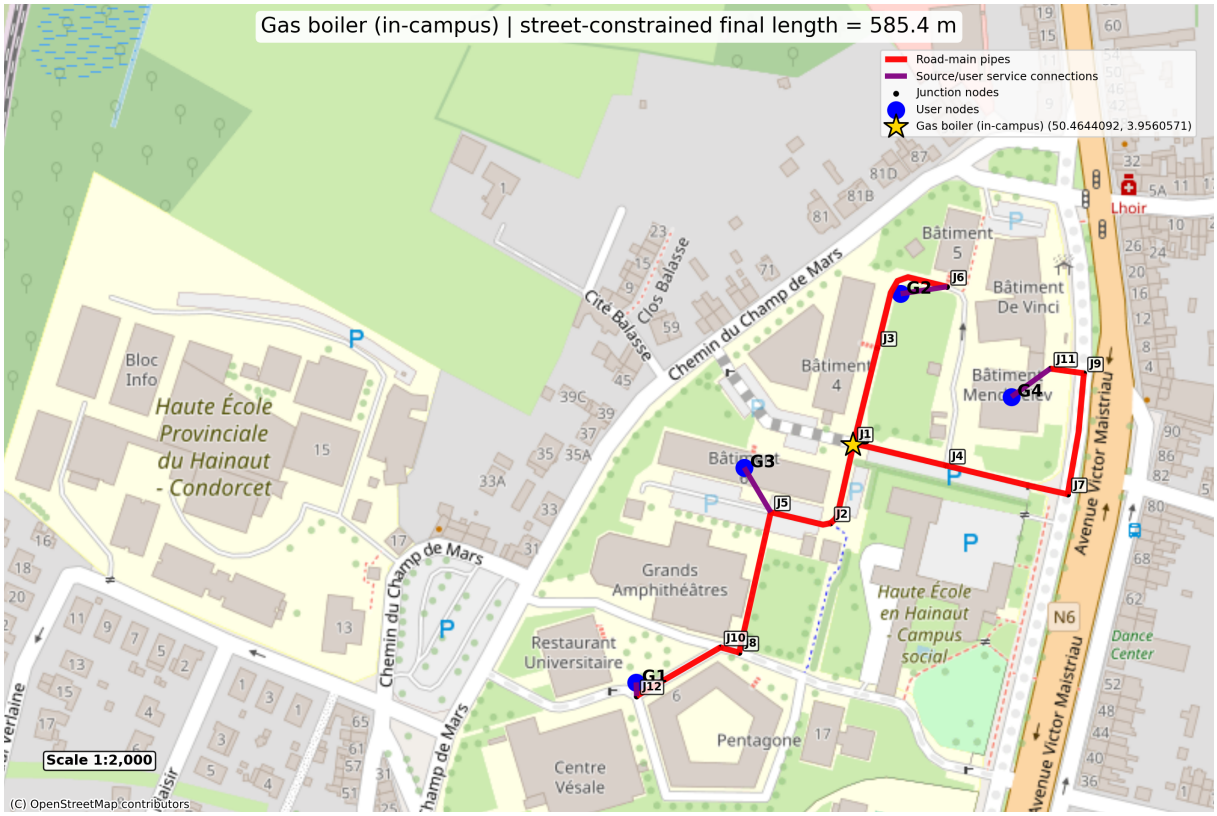

Figure 4.5. OpenStreetMap routing of the in-campus gas-boiler base scenario, showing pipe lengths along campus streets.

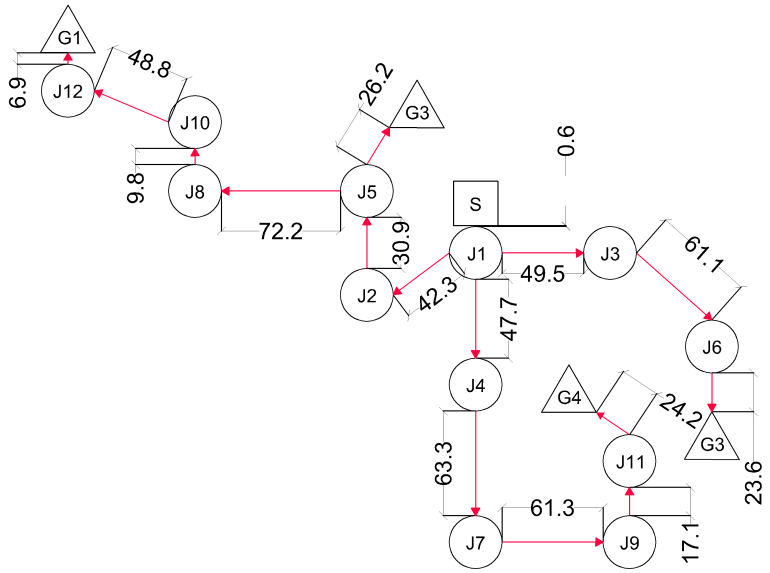

Figure 4.6. Simplified HeatNet topology of the in-campus gas-boiler base scenario, showing junctions and consumer groups.

group, the velocity column gives the corresponding terminal-branch velocity, and $\delta\dot{m}$ is the deviation of the balanced mass flow from the design mass flow.

Table 4.6. Detailed terminal-load results for the gas-boiler base scenario.

Group	Q_i (kW)	$\dot{m}_{\text{des}}$ (kg/s)	D_{term} (mm)	$\dot{m}_{\text{act}}$ (kg/s)	$\delta\dot{m}$ (%)	v_{term} (m/s)	T_s/T_r (°C)
$G1$	1635.7	19.57	200	29.24	+49.5	0.96	85.3/65.3
$G2$	671.2	8.03	125	12.25	+52.6	1.03	85.3/65.3
$G3$	493.7	5.91	100	11.03	+86.7	1.45	85.3/65.3
$G4$	448.0	5.36	100	5.36	0.0	0.70	85.1/65.1

The base scenario is feasible with the adopted velocity ceiling $v_{\max} = 2.5\,\text{m}\,\text{s}^{-1}$ and minimum delivered supply temperature $T_{\min} = 85°\text{C}$. The thermal correction raises the source temperature from 85.0 to about 85.4°C, while the accepted pump head is $H^* = 41.7$ kPa. The load $G4$ acts as the critical terminal consumer: it receives its design mass flow within rounding, while the shorter lower-resistance loops receive a positive flow margin in the passive single-pump configuration.

4.2.3 What-If: Geothermal Supply on the Existing Campus Network

The first what-if replaces the in-campus gas boiler by a remote geothermal heat source while keeping the comparison controlled. The consumer groups, the road-derived campus distribution topology, the common campus pipe diameters, and the distribution-only economic boundary are kept from the gas-boiler base scenario. The geothermal case differs by adding a 377 m source feeder from the remote well to the campus entry junction, applying a 30 % peak-demand reduction, and operating at the lower 70/50°C temperature regime with a minimum delivered supply temperature of 65°C.

Table 4.7 compares this geothermal what-if with the gas-boiler base scenario using the same network-side indicators.

Table 4.7. Gas boiler and geothermal network-side comparison.

Indicator	Gas boiler 85/65	Geothermal 70/50
Total design demand (kW)	3248.7	2274.1
Source supply / final (°C)	85 / 85.4	70 / 70.0
Min. delivered supply (°C)	85.1	69.6
Accepted pump head H^* (kPa)	41.7	75.5
Pump power (kW)	3.6	6.7
Max. branch velocity (m/s)	1.45	1.74
Network heat loss (kW)	24.2	38.8
as % of demand	0.75	1.70
Network CAPEX (€)	1 205 840	1 733 983
Annual maintenance (1 % CAPEX, €/yr)	12 058	17 340
Levelized distribution cost (€/MWh)	3.79	7.88

At this distribution-network boundary, the geothermal what-if is therefore not cheaper than the gas-boiler base case. The additional feeder increases pipe length, heat loss, pump head, CAPEX, maintenance, and levelized distribution cost. The geothermal case is included for

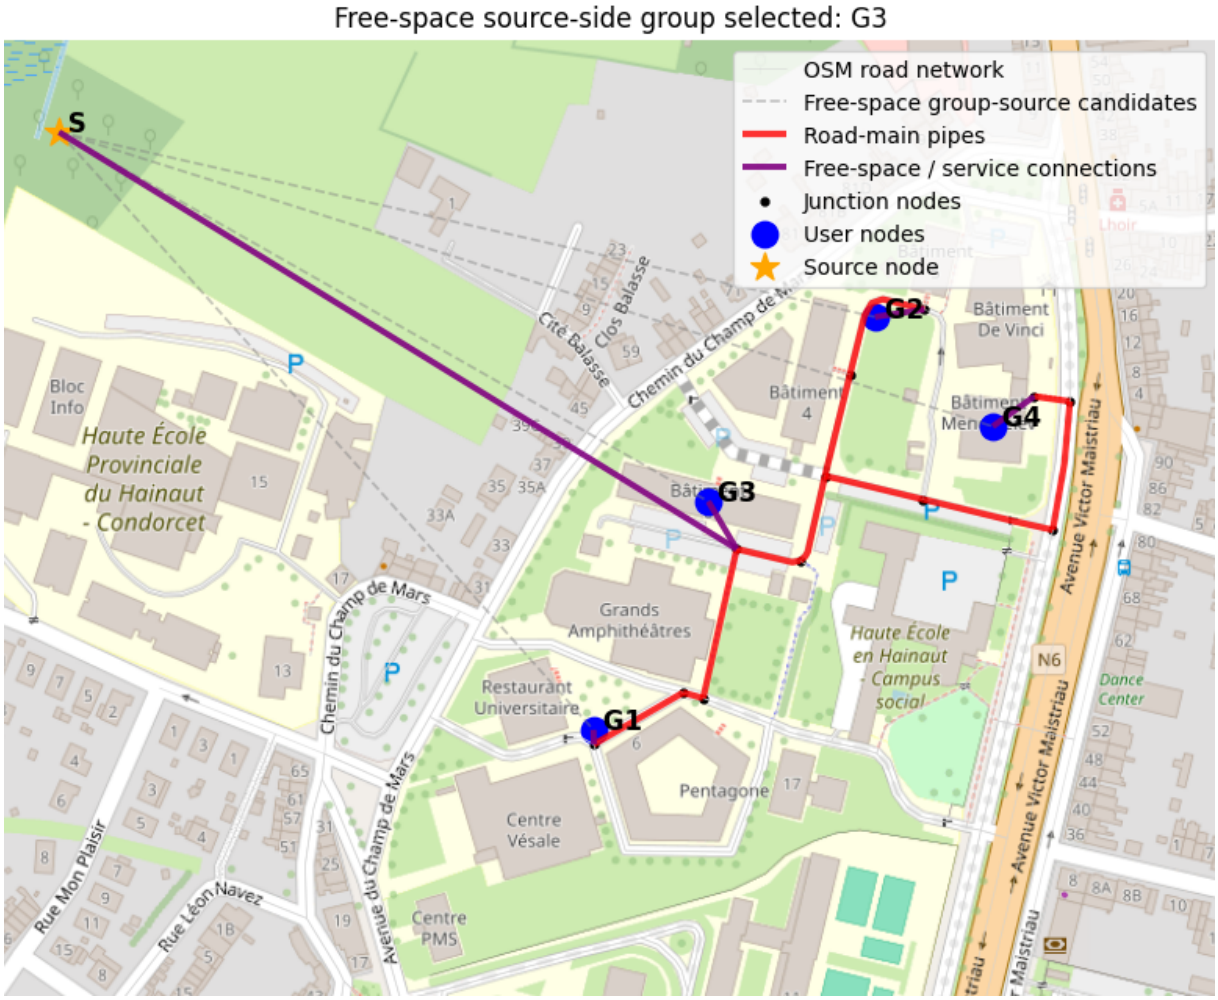

Figure 4.7. Simplified heat-network topology of the UMONS Mons campus derived from OpenStreetMap. The four consumer clusters $G1$–$G4$ are connected as leaves of the campus network. The campus distribution branches are shared with the gas-boiler base case, while the free-space source pipe is the additional feeder used only for the geothermal what-if.

a different reason: it represents the renewable heat-supply alternative that could reduce gas consumption, carbon emissions, and exposure to fuel-price volatility. Those benefits are outside the present network-only economic model. Consequently, Table 4.7 should be read as the distribution-cost penalty that geothermal must offset through generation-side and environmental benefits, not as a proof that geothermal is economically superior on the distribution network alone.

4.2.4 What-If: Re-Sizing the Network for the Geothermal Demand

For the geothermal case, a natural design question is whether the network can be resized for the reduced low-temperature demand rather than reusing the gas boiler pipe sizes. Re-sizing every pipe for the reduced design flow yields smaller diameters ($G1$: 150, $G2$: 100, $G3$: 80, $G4$: 80 mm) and, under the nPro cost model, a lower pipe-and-installation cost, from 1.69 to 1.48 M€, a 12.5 % saving. However, smaller pipes raise the resistance of the critical loop, so the pump head required to feed the farthest consumer $G4$ rises to $H^* = 163.9$ kPa (about 6.5 kW of pump power at design flow). In a passive single-pump network every loop sees that same head. The word "before" in the following comparison therefore means *before balancing valves are added* to the re-sized geothermal network. In that no-valve state, the nearer loops would have to run at the velocities reported in Table 4.8. The closest cluster $G3$ would require 2.60 m s^{-1}, exceeding the physical limit. This identifies the hydraulic condition that must be monitored: cost-driven downsizing can increase loop-resistance differences enough that a passive single-pump layout becomes infeasible. This is the case-study counterpart of the Phase 2 tolerance logic: a failed hydraulic condition is treated as a design infeasibility rather than as a reason to loosen the numerical tolerance. It therefore motivates the introduction of local flow control.

4.2.5 Balancing Valves: a Quantified Resolution

The residual over-supply that no passive layout can remove is resolved by a balancing valve at each consumer substation, introducing an adjustable local pressure drop $\Delta p_{\mathrm{valve},i}$. The valves are sized so that every loop balances at *design* flow under one common pump head. The passive no-valve probe first exposes the infeasible operating point ($H^* = 163.9$ kPa). The valve-controlled design then lowers the pump head to the critical loop's design-flow drop,

$$H_{\mathrm{pump}} \;=\; \max_i \, \Delta p^{\mathrm{design}}_{\mathrm{loop},i}, \qquad \Delta p_{\mathrm{valve},i} \;=\; H_{\mathrm{pump}} - \Delta p^{\mathrm{design}}_{\mathrm{loop},i} \;\geq\; 0, \tag{4.1}$$

with $\Delta p^{\mathrm{design}}_{\mathrm{loop},i}$ equal to the supply+return+HX pressure drop of loop i at design flow. The critical loop receives a zero valve, while every lower-resistance loop receives an added pressure drop so that all loops settle at the same new design-flow head. The combined before/after comparison is given in Table 4.8. The table is read from left to right: $\Delta p_{\mathrm{design}}$ is the loop pressure drop at design flow before any valve is added; H_{passive} is the common head required by the passive no-valve solution; and $\Delta p_{\mathrm{new}} = \Delta p_{\mathrm{design}} + \Delta p_{\mathrm{valve}}$ is the common loop pressure drop after balancing valves are installed. The excess head is the diagnostic gap $H_{\mathrm{passive}} - \Delta p_{\mathrm{design}}$ in the no-valve probe; it is not the same quantity as $\Delta p_{\mathrm{valve}}$, because the valve-controlled case also lowers the pump head from 163.9 to 87.0 kPa.

Table 4.8. Re-sized geothermal network: passive no-valve imbalance and balancing-valve correction.

Group	D (mm)	Δp_{design} (kPa)	H_{passive} (kPa)	v_{passive} (m/s)	Excess head (kPa)	Δp_{valve} (kPa)	Δp_{new} (kPa)	v_{new} (m/s)	Feasible
$G1$	150	56.7	163.9	1.66	107.3	30.3	87.0	0.79	yes
$G2$	100	67.8	163.9	1.12	96.1	19.1	87.0	0.73	yes
$G3$	80	50.9	163.9	**2.60**	**113.1**	**36.1**	87.0	0.84	yes
$G4$	80	87.0	163.9	0.76	77.0	0.0	87.0	0.76	yes (critical)

In the valve calculation, the valves throttle every loop back to its design flow, so all branch velocities return to the design range ($\leq 0.84\ \text{m s}^{-1}$) and the velocity violation disappears. The clearest case is again $G3$: after geothermal re-sizing, the passive no-valve state would require $2.60\ \text{m s}^{-1}$ under the common head of 163.9 kPa. Adding a 36.1 kPa balancing valve brings the loop to the new common design-flow pressure drop of 87.0 kPa and restores the branch velocity to $0.84\ \text{m s}^{-1}$. A further benefit is that the pump head drops from the infeasible passive value of 163.9 kPa (6.5 kW) to 87.0 kPa (3.5 kW). Balancing valves are therefore the local control mechanism to test if passive operation overdrives the short loops.

4.2.6 Economic Sensitivity Analysis

The net present cost and levelised cost of heat distribution reported in Table 4.7 depend on two financial parameters that were set as engineering defaults in Section 3.9: the discount rate r and the project evaluation horizon N_{proj}. Because the distribution-network CAPEX is incurred at year 0 while OPEX is spread over the full horizon, both indicators are sensitive to these inputs. This subsection quantifies that sensitivity by re-running the gas-boiler and geothermal economic calculations with the parameter matrix given in Table 4.9.

Table 4.9. Discount rate and project horizon combinations evaluated in the sensitivity analysis.

Parameter	Low	Default	High
Discount rate r (%/yr)	1	4	—
Project horizon N_{proj} (yr)	20	30	40

Because the geothermal scenario relies on frozen campus pipe diameters sized for the gas-boiler demand, its passive hydraulic solver does not converge at $v_{\max} = 2.5\ \text{m s}^{-1}$. The geothermal sensitivity analysis is therefore performed on the valve-controlled operating point, in which balancing valves throttle every loop to its design flow under a common pump head $H_{\text{pump}} = \max_i \Delta p_{\text{loop},i}^{\text{design}} = 86.4$ kPa, as described in Section 4.2.5.

Gas-boiler results

Table 4.10 reports the net present cost and levelised distribution cost for all nine (r, N) combinations. The default operating point ($r = 4\,\%$, $N = 30$ yr) is marked with an asterisk.

Two trends are visible. First, extending the project horizon from 20 to 40 years reduces the LCOH by roughly 48 % at every discount rate, because the fixed CAPEX is amortised over a larger cumulative heat delivery. Second, lowering the discount rate from 4 % to 1 % increases

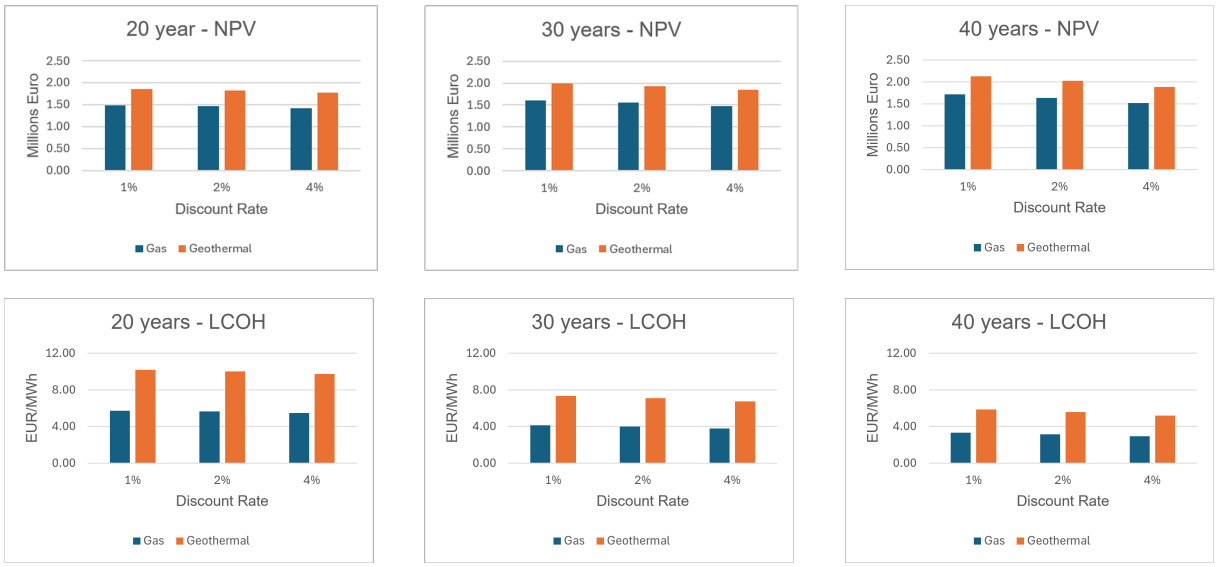

Figure 4.8. NPV and LCOH sensitivity to discount rate (r) and project horizon (N) for the gas-boiler and geothermal valve-balanced scenarios. Values cover the distribution network only; heat-generation cost is excluded.

Table 4.10. Gas-boiler scenario: NPV and LCOH (distribution network only) sensitivity to discount rate and project horizon.

	N_{proj} (years)		
r (%/yr)	20	30	40
1	1 487 632 / 5.72	1 608 843 / 4.13	1 718 573 / 3.31
2	1 461 177 / 5.62	1 555 574 / 3.99	1 633 012 / 3.14
4	1 418 061 / 5.46	1 475 865 / 3.79*	1 514 916 / 2.91

Values are NPV (EUR) / LCOH (EUR / MWh). Distribution network only; heat-generation cost is excluded.

* Default parameter set.

the NPV by approximately 9 %–13 % across horizons: future OPEX terms, which include pump electricity and 1 % annual maintenance on CAPEX, are discounted less aggressively and therefore contribute more to the present cost. The combined effect is that the cheapest distribution cost in the matrix (2.91 EUR / MWh) occurs at the lowest discount rate and longest horizon, while the most expensive (5.72 EUR / MWh) occurs at the shortest horizon with a low discount rate.

These results confirm that the pre-design LCOH figures are sensitive to the financial assumptions. In particular, a project financed at 1 % over 40 years would report a 26 % lower levelised distribution cost than the default case (3.31 vs. 3.79 EUR / MWh), which may be relevant when comparing district heating against individual heating alternatives in a full socio-economic appraisal. These LCOH values cover the distribution network only and exclude heat-generation costs (gas purchase, boiler efficiency, geothermal production plant); they should therefore be read as the distribution-cost component of the total levelised cost of heat.

Geothermal valve-balanced results

Table 4.11 reports the same indicators for the geothermal what-if evaluated at the valve-controlled operating point. The pump head is 86.4 kPa, the pipe network is re-sized for the reduced geothermal demand, and the source supply temperature is 70°C.

Table 4.11. Geothermal valve-balanced scenario: NPV and LCOH (distribution network only) sensitivity to discount rate and project horizon.

	N_{proj} (years)		
r (%/yr)	20	30	40
1	1 854 808 / 10.20	1 999 344 / 7.33	2 130 190 / 5.85
2	1 823 263 / 10.02	1 935 824 / 7.09	2 028 164 / 5.57
4	1 771 849 / 9.74	1 840 777 / 6.75*	1 887 342 / 5.19

Values are NPV (EUR) / LCOH (EUR / MWh). Distribution network only; heat-generation cost is excluded.
* Default parameter set.

The geothermal LCOH is consistently higher than the gas-boiler counterpart at every (r, N) combination: at the default parameters (r = 4 %, N = 30 yr), the geothermal valve-balanced levelised cost is 6.75 EUR / MWh compared with 3.79 EUR / MWh for the gas boiler, a 78 % premium. This difference is dominated by the additional feeder pipe length and the larger pump head required by the re-sized geothermal network. The same two financial trends remain: extending the project horizon from 20 to 40 years reduces the LCOH by roughly 46 % at every discount rate, and lowering r from 4 % to 1 % increases the NPV by approximately 5 %–8 %. Because the distribution-network CAPEX is identical in both scenarios (the same campus pipe diameters are used), the cost premium is entirely attributable to the additional geothermal feeder and the higher pump head, reinforcing the interpretation of Table 4.7: the geothermal distribution cost must be offset by generation-side and environmental benefits to be economically competitive.

4.2.7 Case-Study Interpretation, Assumptions, and Practical Limits

The UMONS case study should be read as a comparative pre-design analysis rather than as a final investment design. The comparison keeps the following assumptions explicit: (i) gas consumption is converted to heat with unit boiler efficiency; (ii) refurbishment in the geothermal case is represented by a 30 % peak-demand reduction; (iii) the minimum delivered supply temperature for the geothermal case is set at 65°C; and (iv) all economic indicators cover the distribution network only. A sensitivity study on the refurbishment fraction (20–40 %) would be a natural extension.

The pipe catalogue also frames the practical reading of the results. The velocity ceiling of 2.5 $\mathrm{m\,s^{-1}}$ is the manufacturer-recommended maximum continuous water speed for flexible district-heating pipes, so any violation in the re-sized geothermal case would be a real product-limit breach, not an arbitrary numerical warning [44]. In addition, flexible PE-Xa pipes are size-limited, while the remote geothermal feeder may require a larger diameter; a realistic design would therefore combine a pre-insulated steel trunk with flexible PE-Xa building connections. Finally, operating temperature affects service life: the gas boiler case runs near the upper end of the PE-Xa temperature range, whereas the 70°C geothermal case is consistent with a longer expected pipe lifetime [44]. Within the present distribution-only boundary, however, this practical benefit remains conditional: the added feeder and any hydraulic imbalance must be handled, and a full economic justification would require adding the generation-side fuel, carbon, and investment terms that are excluded here.

5 Conclusion and Perspectives

This thesis set out to address a concrete methodological gap: the absence of a transparent, fully documented, topology-driven tool able to test the incremental expansion of radial district-heating networks against coupled hydraulic and thermal feasibility criteria. The motivation is both timely and regional. District heating is one of the few technologies that can decarbonise heat at neighbourhood scale, European policy now actively pushes its expansion, and yet Wallonia remains almost entirely without operational networks to learn from. In that context, a fast and auditable pre-design instrument has practical value: it allows an engineer to compare several expansion or supply scenarios before committing to a detailed design, and to do so with every intermediate quantity exposed for verification.

5.1 Summary of Contributions

The central outcome of this work is HeatNet: an open, reproducible pre-design tool for radial district-heating networks. Its contribution is not only the implementation of a solver, but the ability to turn network topology, consumer demand, pipe sizing, thermal feasibility, and distribution cost into a single scenario-comparison workflow. Four contributions can be highlighted.

First, the thesis provides a complete decision-support workflow for radial district-heating extensions. Starting from tabular network and demand inputs, HeatNet produces the engineering quantities needed for early design screening: selected pipe diameters, serviceable load flows, delivered temperatures, pump requirements, heat losses, energy-balance diagnostics, and network-side cost indicators. This makes the tool useful before the project reaches the level of detailed component selection or commercial software modelling.

Second, the work demonstrates that topology changes can be evaluated as part of the normal input workflow rather than as a separate manual reformulation. The reference network and the three expansion scenarios in Chapter 4 are processed with the same calculation structure, including the case where a new internal junction is inserted. This is the main scenario-analysis contribution: the tool can compare alternative radial extensions and return a consistent feasibility diagnosis for each one.

Third, the thesis contributes a transparent reporting structure for thermo-hydraulic pre-design. Instead of reporting only a final feasible or infeasible status, HeatNet exports the intermediate evidence needed to explain the result: branch flows, commercial diameters, pressure drops, temperature levels, heat losses, pump power, and economic indicators. The calculation is therefore traceable enough for academic review and for early engineering discussion.

Fourth, the tool is demonstrated on a real campus-scale case study. Applied to the UMONS Mons campus (Section 4.2), HeatNet connects GIS-derived routing, metered building demand, high- and low-temperature supply scenarios, and distribution-cost comparison. The case study shows that the tool can identify when a passive single-pump radial network becomes hydraulically weak and can point toward a concrete design intervention, namely local balancing valves. This

demonstrates HeatNet's value as a diagnostic pre-design tool, not only as a numerical feasibility checker.

5.2 Limitations

The contributions above are bounded by the modelling choices made explicit throughout the report. The solver is restricted to steady-state simulation of radial supply–return networks at the distribution boundary; meshed and looped topologies, bidirectional and prosumer operation, transient and start-up behaviour, and detailed substation modelling are all out of scope. Consumers are assumed to be terminal leaves of the supply tree, which keeps the pressure-balancing problem well posed but requires through-path buildings to be converted to leaf connections beforehand. Fluid density is held fixed within a hydraulic solve, soil thermal resistance is neglected, and a single mean ground temperature is used in the thermal model. On the economic side, the figures cover the distribution network only and rest on predefined unit costs. Heat-generation plant sizing (geothermal wells, boilers, large heat pumps) is deliberately excluded.

Three further limitations should be made explicit. First, the results are verified only for internal consistency, through the first-law energy balance and the loop-residual tolerance, and have not been validated against an established commercial solver (e.g. EcoStruxure/TERMIS, Apros DH) or measured operating data; such a cross-validation is left for future work. Second, the network is sized for a single worst-hour design point rather than a time-resolved load profile, so the annual energy and economic figures rest on an assumed equivalent full-load duration rather than on simulated seasonal operation. Third, the balancing valves of Section 4.2.5 are sized analytically as a post-processing step and are not part of the automated pressure-balancing solver, which models only passive single-pump operation. Fourth, the geographic path-finding algorithm occasionally routes pipes through building footprints rather than following the street graph; constraining the routing to stay on the road network is identified as a priority improvement for future work.

5.3 Perspectives

Several extensions follow naturally from these limitations. The most immediate is a sensitivity study of the case-study conclusions to the key assumptions, in particular the refurbishment fraction of the geothermal case, which was fixed at 30 % but could reasonably be swept over 20–40 % to test the robustness of the supply-temperature and valve-sizing results. A second extension is to move from a single design-hour analysis toward a time-resolved or load-duration evaluation, which would let the economic layer reflect seasonal operation rather than a single worst hour. A third direction is to relax the fixed-fluid-property and mean-ground-temperature assumptions and to add a soil-resistance term, improving the accuracy of the thermal balance for buried trunks. Architecturally, the most significant generalisation would be to extend the traversal-based solver toward meshed and bidirectional topologies, which would require replacing the unique-path assumption with a loop-aware or Newton-type hydraulic formulation while retaining the same auditable, export-driven philosophy. Finally, coupling HeatNet with a generation-side model

and with automated extraction of consumer demand from GIS and metered data would turn the present pre-design instrument into a more complete planning workflow, a capability that is directly relevant to the early-stage projects that Wallonia will need if district heating is to scale at the pace European climate targets require.

Bibliography

[1] Eurostat. *EU household gas prices rise in the second half of 2024*. 2025. URL: https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20250506-3 (visited on 03/10/2026).

[2] Eurostat. *Energy prices stabilised at a high level in 2024*. 2025. URL: https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20250220-1 (visited on 03/10/2026).

[3] Eurostat/Interfax. *EU dependence on Russian imports of oil falls to 2.3% in two years, gas to 12.9%*. URL: https://interfax.com/newsroom/top-stories/94236/ (visited on 03/12/2026).

[4] Statista. *Share of extra-EU natural gas import value from Russia from 2010 to 1st quarter 2025*. URL: https://www.statista.com/statistics/1021735/share-russian-gas-imports-eu (visited on 03/12/2026).

[5] European Commission. *2030 Climate Targets*. URL: https://climate.ec.europa.eu/eu-action/climate-strategies-targets/2030-climate-targets_en (visited on 03/10/2026).

[6] European Commission. *REPowerEU*. URL: https://commission.europa.eu/topics/energy/repowereu_en (visited on 03/10/2026).

[7] Eurostat. *Heating and cooling from renewables gradually increasing*. 2023. URL: https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20230203-1 (visited on 03/10/2026).

[8] European Commission. *Heating and Cooling*. URL: https://energy.ec.europa.eu/topics/energy-efficiency/heating-and-cooling_en (visited on 03/10/2026).

[9] Eurostat. *Energy consumption in households (statistics explained)*. URL: https://ec.europa.eu/eurostat/statistics-explained/SEPDF/cache/58200.pdf (visited on 03/10/2026).

[10] J. von Rhein et al. "Development of a topology analysis tool for fifth-generation district heating and cooling networks". In: *Energy Conversion and Management* 196 (Sept. 2019), pp. 705–716. DOI: 10.1016/j.enconman.2019.06.057.

[11] M. Vesterlund and A. Toffolo. "Design Optimization of a District Heating Network Expansion, a Case Study for the Town of Kiruna". In: *Applied Sciences* 7.5 (May 2017), p. 488. DOI: 10.3390/app7050488.

[12] M. Tunzi et al. "Double Loop Network for Combined Heating and Cooling in Low Heat Density Areas". In: *Energies* 13.22 (Nov. 2020), p. 6091. DOI: 10.3390/en13226091.

[13] L. Frison et al. "Model predictive control of bidirectional heat transfer in prosumer-based solar district heating networks". In: *Applied Energy* 358 (Mar. 2024), p. 122617. DOI: 10.1016/j.apenergy.2023.122617.

[14] G. E. Dino et al. “Modeling of a bidirectional substation in a district heating network: Validation, dynamic analysis, and application to a solar prosumer”. In: *Energy* 284 (Nov. 2023), p. 128621. DOI: 10.1016/j.energy.2023.128621.

[15] P. Sdringola et al. “Prosumers and district heating: Experimental validation of strategies to improve thermal energy production and consumption”. In: *Energy and Buildings* 338 (May 2025), p. 115713. DOI: 10.1016/j.enbuild.2025.115713.

[16] Persistence Market Research. *District Heating Market Size, Share, and Growth Forecast 2026–2033*. URL: https://www.persistencemarketresearch.com/market-research/district-heating-market.asp (visited on 03/10/2026).

[17] Heat Roadmap Europe. *Towards a sustainable, resilient and competitive heating sector by 2050 – Executive Summary*. 2025. URL: https://coolheatingcoalition.eu/wp-content/uploads/2025/11/Heat-Roadmap-Europe-Towards-a-sustainable-resilient-and-competitive-heating-sector-by-2050-Executive-Summery.pdf (visited on 03/10/2026).

[18] International Energy Agency. *Renewables 2023*. Tech. rep. Accessed: 2026-03-10. IEA, 2023. URL: https://iea.blob.core.windows.net/assets/3f7f2c25-5b6f-4f3c-a1c0-71085bac5383/Renewables_2023.pdf.

[19] Euroheat & Power. *DHC Market Outlook 2024*. URL: https://www.euroheat.org/data-insights/outlooks/dhc-market-outlook-2024 (visited on 05/12/2026).

[20] S. Koutra et al. “Feasibility of fifth-generation district heating and cooling using mine water in Belgium: A multi-site techno-economic assessment”. In: *Sustainable Energy Technologies and Assessments* 86 (2026), p. 104853. DOI: 10.1016/j.seta.2026.104853.

[21] Service Public de Wallonie (SPW). *WalOnMap — Géoportail de la Wallonie*. 2026. URL: https://geoportail.wallonie.be/walonmap (visited on 06/05/2026).

[22] Schneider Electric. *EcoStruxure District Energy Brochure*. URL: https://www.se.com/es/es/download/document/EcoStruxureDistricEnergy/ (visited on 05/12/2026).

[23] VTT / Fortum. *Apros® District*. 2021. URL: https://www.apros.fi/wp-content/uploads/sites/27/2021/09/Apros-District.pdf (visited on 05/10/2026).

[24] R. Lund et al. *A review of modelling and optimisation techniques applied to district heating systems*. 2016. URL: https://www.sciencedirect.com/science/article/pii/S1364032115007704 (visited on 03/10/2026).

[25] Various. *A comprehensive review of district heating and cooling systems*. 2020. URL: https://www.sciencedirect.com/science/article/pii/S0360544220315000 (visited on 03/10/2026).

[26] Various. *Optimization of District Heating Network*. 2025. URL: https://asmedigitalcollection.asme.org/thermalscienceapplication/article-abstract/17/1/011007/1207341/Optimization-of-District-Heating-Network (visited on 03/10/2026).

[27] Various. *Dynamic modelling of district heating networks*. 2022. URL: https://www.sciencedirect.com/science/article/pii/S1359431122018427 (visited on 03/10/2026).

[28] H. Wang et al. *Modelling and simulation of meshed district heating networks*. 2020. URL: https://www.sciencedirect.com/science/article/pii/S0360544220310574 (visited on 03/10/2026).

[29] Various. *Graph-based modelling of district heating networks*. 2021. URL: https://www.sciencedirect.com/science/article/pii/S0378778821007969 (visited on 03/10/2026).

[30] T. Gumpert et al. *Matrix-based hydraulic and thermal simulation of a real district heating network*. 2018. URL: https://www.sciencedirect.com/science/article/pii/S0375650518303195 (visited on 03/10/2026).

[31] C. Wieland et al. *Comparison of equivalent-circuit and standardised heat-loss methods for buried district heating pipes*. 2024. DOI: 10.2478/rtuect-2024-0070. URL: https://reference-global.com/article/10.2478/rtuect-2024-0070 (visited on 03/10/2026).

[32] Various. *Techno-economic assessment of district heating pipe alternatives*. 2018. URL: https://www.sciencedirect.com/science/article/pii/S1876610218304752 (visited on 03/10/2026).

[33] Y. Xiang et al. *Efficient steady-state solver for radial district heating networks*. 2019. URL: https://ieeexplore.ieee.org/document/8865856 (visited on 03/10/2026).

[34] L. Tang and S. Huang. *Newton–Raphson power flow for district heating networks with changing flow directions*. 2020. URL: https://ieeexplore.ieee.org/document/9061996 (visited on 03/10/2026).

[35] J. Dancker and M. Wolter. "Improved quasi-steady-state power flow calculation for district heating systems: A coupled Newton-Raphson approach". In: *Applied Energy* 295 (2021), p. 116930. DOI: 10.1016/j.apenergy.2021.116930. URL: https://www.sciencedirect.com/science/article/pii/S0306261921004104.

[36] J. Zheng et al. *Graph-traversal-based thermal calculation order for district heating networks*. 2021. URL: https://www.sciencedirect.com/science/article/pii/S0360544221016376 (visited on 05/12/2026).

[37] Various. *The Hardy Cross method for pipe network analysis*. 2019. URL: https://www.mdpi.com/2311-5521/4/2/73 (visited on 03/10/2026).

[38] Various. *Computational methods for hydraulic network solving*. 2024. URL: https://www.mdpi.com/2079-3197/12/2/25 (visited on 03/10/2026).

[39] W3Schools. *DSA Binary Trees Post-order Traversal*. 2025. URL: https://www.w3schools.com/dsa/dsa_algo_binarytrees_postorder.php (visited on 03/15/2026).

[40] W3Schools. *DSA Binary Trees Pre-order Traversal*. 2025. URL: https://www.w3schools.com/dsa/dsa_algo_binarytrees_preorder.php (visited on 03/15/2026).

[41] Hans-Joachim Kretzschmar and Wolfgang Wagner. "IAPWS Industrial Formulation 1997 for the Thermodynamic Properties of Water and Steam". In: *International Steam Tables*. Berlin, Heidelberg: Springer, 2019, pp. 7–150. DOI: 10.1007/978-3-662-53219-5_3.

[42] Massachusetts Institute of Technology. *Network Models 8*. No date. URL: https://web.mit.edu/15.053/www/AMP-Chapter-08.pdf (visited on 05/29/2026).

[43] J. Vaillant Rebollar et al. "Urban heat distribution networks comparison using a GIS-based district heating design tool". In: *Proc. XX Scientific Conference on Engineering and Architecture (MACDES 2022)*. Havana, Cuba, 2022, pp. 3395–3407.

[44] Uponor GmbH. *Uponor Ecoflex piping systems — Technical information*. Tech. rep. 1120217 v4_01_2025_EN. Accessed: 2026-03-10. Uponor GmbH, 2025. URL: https://www.uponor.com/services/download-centre.

[45] LOGSTOR A/S. *Product Catalogue — Pre-insulated Solutions, Industry*. Tech. rep. Second issue, July 2022. LOGSTOR A/S, 2022. URL: https://www.logstor.com.

[46] P. Virtanen et al. "SciPy 1.0: fundamental algorithms for scientific computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

[47] SciPy Community. *least_squares — SciPy v1.17.0 Manual*. 2024. URL: https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.least_squares.html.

[48] SciPy Community. *brentq — SciPy v1.17.0 Manual*. 2026. URL: https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.brentq.html (visited on 05/29/2026).

[49] J. Rojer et al. "Integral techno-economic design and operational optimization for district heating networks with a mixed integer linear programming strategy". In: *Energy* 308 (2024), p. 132710. DOI: 10.1016/j.energy.2024.132710.

[50] nPro Energy GmbH. *Technology Costs — District Heating Network Pipes and Installation*. 2025. URL: https://www.npro.energy/main/en/help/technology-costs (visited on 05/22/2026).

[51] KEA Klimaschutz- und Energieagentur Baden-Württemberg GmbH (KEA-BW). *Einführung in den Technikkatalog zur kommunalen Wärmeplanung in Baden-Württemberg*. Version 1.1, im Auftrag des Ministeriums für Umwelt, Klima und Energiewirtschaft Baden-Württemberg. Section 4.2 covers low-temperature heating-network investment costs; underlying cost basis adopted by nPro Energy GmbH. Stuttgart, 2023. URL: https://um.baden-wuerttemberg.de/fileadmin/redaktion/m-um/intern/Dateien/Dokumente/2_Presse_und_Service/Publikationen/Energie/Kommunale-Waermeplanung-Einfuehrung-in-den-Technikkatalog.pdf (visited on 05/29/2026).

[52] nPro Energy GmbH. *Economic Calculation — Annuity Method and Operating Costs*. 2025. URL: https://www.npro.energy/main/en/help/economic-calculation (visited on 05/29/2026).

[53] Armstrong Fluid Technology. *Minimum Efficiency Index (MEI) Table*. Data sheet 43-17UK. No date. URL: https://armstrongfluidtechnology.com/~/media/documents/installation_maintenance_parts/data-sheets/43_17uk_mei_table_for_armstrong_literature.pdf (visited on 05/29/2026).

[54] OpenStreetMap contributors. *OpenStreetMap Copyright and License*. no date. URL: https://www.openstreetmap.org/copyright (visited on 05/29/2026).

[55] Pipe Flow Software. *Colebrook–White Formula — Pipe Flow Expert Software User Guide*. no date. URL: https://www.pipeflow.com/public/PipeFlowExpertSoftwareHelp/html/Colebrook-WhiteFormula.html (visited on 05/29/2026).

[56] Wolfram Research. *Colebrook–White Equation for Darcy–Weisbach Friction Factor*. Wolfram Formula Repository. no date. URL: https://resources.wolframcloud.com/FormulaRepository/resources/ColebrookWhite-Equation-for-DarcyWeisbach-Friction-Factor (visited on 05/29/2026).

[57] S. El Mrabet et al. "A brief overview of district heating pipe network progress". In: *Energy Conversion and Management: X* 23 (2024), p. 100641. DOI: 10.1016/j.ecmx.2024.100641.

Appendix 1 Nomenclature

The notation used throughout this report is summarized in Table 1.1.

Table 1.1. Nomenclature.

Symbol	Unit	Definition
Hydraulic variables		
$\dot{m}$	$\mathrm{kg\,s^{-1}}$	Mass flow rate
$\dot{m}_{\mathrm{design},i}$	$\mathrm{kg\,s^{-1}}$	Design mass flow of load i
$\dot{m}_{\mathrm{cum},n}$	$\mathrm{kg\,s^{-1}}$	Cumulative mass flow at node n
v	$\mathrm{m\,s^{-1}}$	Mean flow velocity
v_{design}	$\mathrm{m\,s^{-1}}$	Target design velocity for pipe sizing
v_{min}, v_{max}	$\mathrm{m\,s^{-1}}$	Lower and upper velocity bounds
ρ	$\mathrm{kg\,m^{-3}}$	Fluid density
μ	Pa·s	Dynamic viscosity
Re	—	Reynolds number
f	—	Darcy friction factor
ε	m	Absolute pipe-wall roughness
Δp_e	Pa	Pressure drop along pipe e
$\Delta p_{\mathrm{HX},i}$	Pa	Heat-exchanger pressure drop at load i
K_e	—	Minor-loss coefficient of pipe e
H	Pa	Pump head
H^*	Pa	Corrected pump head selected in Phase 4
D_e	m	Commercial inner diameter of pipe e
D_{theo}	m	Theoretical pipe diameter from design rule
L_e	m	Length of pipe e
p_{source}	Pa	Prescribed source-node pressure
N_L	—	Number of load (consumer) nodes
N	—	Total number of nodes
Thermal variables		
$T_{\mathrm{source,supply}}$	°C	Source supply temperature
$T_{\mathrm{supply},n}$	°C	Supply temperature at node n
$T_{\mathrm{return},n}$	°C	Return temperature at node n
$T_{\mathrm{in},e}$, $T_{\mathrm{out},e}$	°C	Inlet and outlet temperature of pipe e
T_g	°C	Ground temperature
$T_{\mathrm{supply,min}}$	°C	Required minimum supply temperature at loads
$\Delta T_{\mathrm{design},i}$	K	Design supply–return temperature drop at load i
c_p	$\mathrm{J\,kg^{-1}\,K^{-1}}$	Specific heat capacity of water
Q_i	W	Heat demand of load i
$Q_{\mathrm{loss},e}$	W	Heat loss of pipe e
Q_{source}	W	Total heat injection at the source
$Q_{\mathrm{delivered}}$	W	Total useful heat delivered to consumers
R_{th}	$\mathrm{K{\cdot}m\,W^{-1}}$	Linear thermal resistance per unit pipe length (= $1/U$)

Table 1.1. (continued).

Symbol	Unit	Definition
U	$\mathrm{W\,m^{-1}\,K^{-1}}$	Linear heat-transfer coefficient (per unit length)
λ_{PUR}	$\mathrm{W\,m^{-1}\,K^{-1}}$	Thermal conductivity of PUR insulation foam
D_{casing}	m	Outer diameter of the insulation casing (jacket pipe)
d_{steel}	m	Outer diameter of the steel carrier pipe
Economic variables		
C_{CAPEX}	EUR	Total capital expenditure (network only)
C_{NPV}	EUR	Net present cost
$C_{\mathrm{OPEX},y}$	EUR	Operating expenditure in year y
C_{trench}	EUR	Civil-works and trench installation cost
C_{pipe}	EUR	Total pipe-material cost
C_{valve}	EUR	Total valve capital cost
C_{pump}	EUR	Total pump capital cost
$C_{\mathrm{elec},y}$	EUR	Pump-electricity cost in year y
$C_{\mathrm{maint},y}$	EUR	Maintenance cost in year y
$C_{\mathrm{maint,annual}}$	EUR	Optional annual scheduled maintenance cost
$C_{\mathrm{maint,5y}}$	EUR	Optional five-year scheduled maintenance cost
$C_{\mathrm{maint,10y}}$	EUR	Optional ten-year scheduled maintenance cost
$c_{\mathrm{total}}(D_e, a_e)$	$\mathrm{EUR\,m^{-1}}$	Total pipe-plus-installation cost per metre of pipe pair in environment a_e
$c_{\mathrm{pipe}}(D_e)$	$\mathrm{EUR\,m^{-1}}$	Pipe-only material cost per metre of pipe pair
c_{pump}	$\mathrm{EUR\,kW^{-1}}$	Pump unit capital cost
c_{valve}	$\mathrm{EUR\,valve^{-1}}$	Unit cost per valve
c_{elec}	$\mathrm{EUR\,kWh^{-1}}$	Electricity price for pumping energy
$E_{\mathrm{pump},y}$	kWh	Pumping-electricity demand in year y
P_{pump}	kW	Pump electrical power requirement
η_{pump}	—	Pump efficiency
f_{maint}	$\mathrm{yr^{-1}}$	Annualized maintenance fraction of CAPEX
$N_{\mathrm{valve},e}$	—	Number of valves assigned to pipe pair e
$N_{\mathrm{valve,pair}}$	$\mathrm{valves\,pair^{-1}}$	Average number of valves per pipe pair
Q_{design}	W	Total design heat demand used in the LCOH calculation
LCOH	$\mathrm{EUR\,MWh^{-1}}$	Levelized cost of heat distribution
r	—	Discount rate
h_{op}	$\mathrm{h\,yr^{-1}}$	Equivalent full-load operating hours per year
N_{proj}	yr	Financial evaluation horizon

Appendix 2 Fluid Properties and Physical Models

This appendix gathers the supporting physical relations used by the hydraulic and thermal solvers so that the main methodology chapter can remain focused on topology-driven equation assembly and solver logic. The purpose of this section is not to introduce a second methodology, but to document the physical assumptions behind the residual evaluations.

2.1 Water Properties

The working fluid is liquid water. In the hydraulic model, the mass-flow conversion and pipe pressure-drop calculations require the fluid density ρ and the dynamic viscosity μ. The Reynolds number is evaluated from

$$\mathrm{Re} = \frac{\rho v D}{\mu}. \tag{2.1}$$

Here, Re is the Reynolds number, ρ is the water density, v is the mean flow velocity, D is the pipe inner diameter, and μ is the dynamic viscosity. This dimensionless number indicates whether the flow is laminar or turbulent.

For thermal calculations, the specific heat capacity c_p is also required in order to convert pipe heat loss into a temperature drop along the flow direction.

The density, dynamic viscosity, and specific heat capacity values used in the model are taken from standard thermodynamic data for liquid water based on the IAPWS-IF97 formulation [41]. The lookup table used for linear interpolation of these water properties is reported in Table 2.1.

Table 2.1. Water-property lookup table.

T (°C)	ρ (kg/m^3)	μ (Pa·s)	c_p (J/kg·K)
0	999.8	0.001792	4218
10	999.7	0.001307	4192
20	998.2	0.001002	4182
30	995.7	0.000798	4179
40	992.2	0.000653	4179
50	988.1	0.000547	4181
60	983.2	0.000467	4185
70	977.8	0.000404	4190
80	971.8	0.000354	4196
90	965.3	0.000315	4205
100	958.4	0.000282	4216
110	950.6	0.000255	4229
120	943.4	0.000232	4245

In HeatNet, the density is held fixed during each hydraulic residual solve. The supply-side density is evaluated at the reference source temperature T_{ref}, while the return-side density is evaluated at a network-average return temperature,

$$\bar{T}_{\mathrm{ret}} = \max\left(T_{\mathrm{ref}} - \frac{1}{N_L}\sum_{i=1}^{N_L}\Delta T_{\mathrm{design},i},\, 10^{\circ}\mathrm{C}\right), \tag{2.2}$$

so that $\rho_{\mathrm{ret}} = \rho(\bar{T}_{\mathrm{ret}})$. This avoids changing the density field inside one pressure-balancing residual evaluation.

The dynamic viscosity is evaluated pipe by pipe when a temperature field is available:

$$\mu_e = \mu\left(\tfrac{1}{2}\left(T_{\mathrm{in},e} + T_{\mathrm{out},e}\right)\right). \tag{2.3}$$

Before any thermal pass is available, the viscosity falls back to $\mu(T_{\mathrm{ref}})$ on the supply side and $\mu(\bar{T}_{\mathrm{ret}})$ on the return side. The final exported hydraulic state is replayed with the same temperature-dependent viscosity convention used inside the converged residual evaluation.

2.2 Friction Factor and Flow Regime

To determine the regular pressure drop along each pipe, the Darcy friction factor f is evaluated from the flow regime and the pipe roughness. The Reynolds number distinguishes laminar and turbulent behavior.

For laminar flow, the classical relation is

$$f = \frac{64}{\mathrm{Re}}. \tag{2.4}$$

Here, f is the Darcy friction factor, a dimensionless coefficient that converts the flow regime and pipe roughness into distributed pressure loss. In laminar flow it depends only on Re.

For turbulent flow, the friction factor depends on both Reynolds number and relative roughness ε/D. In general form, this dependence is represented by the Colebrook–White relation [55, 56],

$$\frac{1}{\sqrt{f}} = -2\log_{10}\left(\frac{\varepsilon}{3.7D} + \frac{2.51}{\mathrm{Re}\sqrt{f}}\right), \tag{2.5}$$

Here ε is the absolute roughness of the inner pipe wall. In practice, either this implicit relation or a standard explicit approximation such as Swamee–Jain may be used, provided that the same convention is applied consistently throughout the network model.

The role of roughness is particularly important when comparing different commercial pipe materials. Smooth polymer pipes such as PEHD typically have much smaller roughness than steel pipes, which reduces the friction factor for the same diameter and Reynolds number. Therefore, the selected roughness values should be interpreted as part of the technical input data rather than as arbitrary tuning parameters.

2.3 Minor Losses and Singular Pressure Drops

In addition to distributed friction losses, each branch may include local singular losses associated with fittings and equipment. These are represented through a lumped minor-loss coefficient K in the Darcy–Weisbach expression [57],

$$\Delta p_{\text{minor}} = K \frac{\rho v^2}{2}. \tag{2.6}$$

Here, Δp_{minor} is the pressure drop caused by local fittings and equipment, and K is a dimensionless minor-loss coefficient. The term $\rho v^2/2$ is the dynamic pressure, so larger velocities make these local losses grow quadratically.

Typical contributors include elbows, tees, valves, contractions, and expansions. In a compact network model, these losses may be aggregated at branch level rather than modeled one fitting at a time. Even when simplified default values are used, the modeling assumption should be stated explicitly because singular losses contribute directly to the loop pressure-drop residuals solved in the hydraulic balancing stage.

The total local hydraulic resistance of a pipe branch can therefore be interpreted as the sum of distributed friction and minor losses,

$$\Delta p_e = \left(f_e \frac{L_e}{D_e} + K_e \right) \frac{\rho_e v_e^2}{2}. \tag{2.7}$$

Here, Δp_e is the total pressure drop of branch e, $f_e L_e / D_e$ is the distributed friction contribution, and K_e collects the minor losses. The bracketed term acts as the branch hydraulic resistance coefficient: for a given velocity, a longer pipe, a smaller diameter, a rougher flow regime, or more fittings all increase the pressure head required to carry the flow.