# Comparing Rule-based Supply Temperature Control Strategies in Multi-Source District Heating Networks

Roberto Boghetti[1,2] and Jérôme H. Kämpf[1,2]

[1] Idiap Research Institute, Rue Marconi 19, Martigny, 1920, Switzerland
[2] EPFL, L'IDIAP, Station 14, Lausanne, 1015, Switzerland
roberto.boghetti@idiap.ch

This preprint has not undergone peer review or any post-submission improvements or corrections.
The Version of Record of this contribution is published in **Construction, Energy, Environment and Sustainability**,
and is available online at https://doi.org/10.1007/978-981-95-1826-5_38.

***Note:*** *The published Version of Record contains significant expansions compared to this preprint, including additional experiments and extended analysis.*

**Abstract.** District Heating Networks (DHNs) play a key role in sustainable energy transition scenarios, but their efficiency depends on effective supply temperature control. Traditional rule-based strategies, such as heating curves, are widely used due to their simplicity and stability. However, they are often tuned over extended periods, during which the network operates suboptimally. This paper proposes an alternative approach based on a differentiable dynamic simulation model, allowing direct optimization of control parameters through gradient descent. We apply this approach to optimize the parameters of various heating curve formulations and evaluate its performance on a real meshed network using monitoring data. Results show a reduction of heat losses even compared to manual regulation, demonstrating the potential of gradient-based methods for DHN control.

**Keywords:** district heating, control, thermal networks, heat distribution, optimization

## 1 Introduction

District Heating Networks (DHNs) play a critical role in sustainable energy transition scenarios [1] as they allow the efficient distribution of thermal energy from centralized or decentralized sources to consumers in urbanized settings [2]. Many existing networks follow the design principles of 3rd-generation DHNs [3], where there is a significant temperature difference between the working fluid and the surrounding soil. In such systems, supply temperature control is crucial: higher temperatures ensure that consumer demand is met but come at the cost of increased energy losses. Conventional control methods rely on predefined rules or

simple linear dependencies on outdoor temperature, known as heating curves [4]. These curves are typically tuned over extended periods of time, commonly a full year of operation, with parameters slowly adjusted based on observed performance. While this approach is straightforward and provides operational stability, it results in prolonged periods of suboptimal control, leading to unnecessary energy waste. Moreover, when major changes occur in the network, the tuning process must be repeated, making it difficult to adapt efficiently to evolving conditions.

This paper addresses these limitations by formulating an optimization-based approach to control strategy design. Using a differentiable dynamic thermal model, we optimize the control parameters of heating curves directly through gradient-based methods, allowing for systematic performance improvements without the need for extensive trial-and-error tuning. We study the performance of these optimized strategies on a real network and compare them against manual temperature setting, demonstrating their potential to enhance efficiency and reduce energy waste in DHN operations.

Given the constraints on data availability, we focus our analysis on a representative 5-day period. However, the proposed method can be extended to longer periods, provided that the gradient computation is adapted to manage computational costs and maintain numerical stability.

## 2 Methodology

### 2.1 Case Study Network

For this study, we used the DHN of Verbier, Switzerland, as described in [5]. At the time of data collection, the network comprised 165 active substations, 35 dedicated to ramp defrosting, and three heat plants (MAIN, SEC1, SEC2). These have no accumulators, limiting their potential for temperature control and increasing the importance of an accurate control strategy. The distribution network had a meshed topology with six internal loops, a total pipe length of 28 km, and covered an altitude difference of 126 m. Most pipes were buried at a depth of 0.8 m, though some ran through basements or parking lots.

Monitoring data, including flow, temperature, and pressure, was collected between January 10 and 22, 2022, at variable intervals and resampled to a 15-minute resolution. During the monitoring period SEC2 operated sporadically with insufficient data and was therefore excluded from the analysis. Outdoor air temperature data was collected from the nearby MeteoSwiss weather station in Montagnier, Val de Bagnes (839 m asl), and the soil temperature at network depth estimated as -2.23°C for the whole period. For further details, we refer the reader to the original work (see [5]).

An overview of the monitored data is given in Figure 1. The left plot illustrates the range of air temperatures and the corresponding network demand, while the right plot shows the operational settings of MAIN. Unlike conventional setups where the supply temperature is regulated based on air temperature, here

it is manually set. The demand peaks at around 9 AM, which is met by increasing the mass flow rate and, consequently, lowering the supply temperature. Around 10 PM, as demand decreases, the mass flow rate is reduced, and the supply temperature is increased to compensate for the longer travel time and prevent excessively low temperatures at consumer substations. A different control strategy is used for SEC1, where the temperature and mass flow are kept almost constant at 82°C and 5 kg/s respectively.

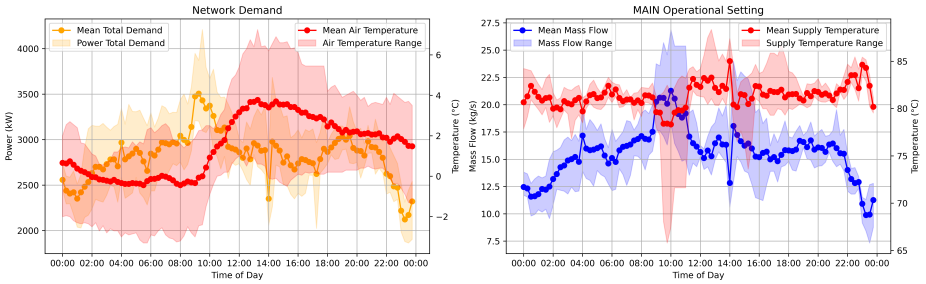


**Fig. 1.** Daily variations in monitoring data. (Left) The total power demand (kW) and air temperature (°C) over a day, showing mean values and ranges. (Right) The mass flow (kg/s) and supply temperature (°C) at the MAIN plant, illustrating operational dynamics. Shaded regions indicate the variability within the investigated period.

## 2.2 Dynamic Model of the DHN

In this section, we present the dynamic thermal model used in this work, that is adapted from [5] and has been validated on the considered case study. We focus solely on the thermal dynamics of the network, assuming a predefined mass flow based on the results of the above study. In the first subsection, we discuss the modeling assumptions. We then describe the considered components and the overall convergence strategy.

**Modeling assumptions.** We assume that the DHN is a closed system, with no leaks or injections of water, represented as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the edges $\mathcal{E}$ correspond to components - pipes, consumers, or producers - while the vertices $\mathcal{V}$ represent junctions. Producers and consumers have fixed mass flow directions, and their outlet vertices have no other incoming edges.

**Component functions.** For each component, we model the outlet temperature, $\theta_{\text{out}}$ (°C), as a function of the inlet temperature, $\theta_{\text{in}}$ (°C):

$$\theta_{\text{out}} = \psi(\theta_{\text{in}}). \tag{1}$$

In pipes, this relationship was originally adapted from the two-capacity Lagrangian model in [6]. Pipes are treated as hollow cylinders with three layers: internal wall, insulation, and casing. The water is discretized into volumes of varying size, each with a constant temperature. The internal wall is segmented using the same spatial discretization as the water. Unlike [6], at each time step $t$, we first solve the source problem to update the temperatures of water volumes and pipe wall segments. The heat loss of a water segment $v$ with volume $V_v$ and temperature $\theta_v$ over a time step $\Delta t$ (s) is given by:

$$Q_v^{(t)} = \left(\theta_v^{(t)} - \theta_v^{(t-1)}\right) \rho^{(t-1)} V_v c_p^{(t-1)} \frac{\Delta t}{h} \tag{2}$$

where $h = 3600$ s, and the water density $\rho$ and specific heat capacity $c_p$ are evaluated at the initial temperature of the time step, following the relationships in [7]. Next, volumes are advected, and the outlet temperature of the pipe is obtained as a volume-weighted sum of the exiting water temperatures. If the incoming volume $V_{\text{in}}$ exceeds the pipe's capacity $V_{\max}$, the excess must be added to the outflow. To account for this, we define a blending parameter $\alpha$ as:

$$\alpha = \max\left(0, \frac{V_{\text{in}} - V_{\max}}{V_{\text{in}}}\right) \tag{3}$$

which represents the fraction of inlet water that surpasses the pipe's capacity. The outlet temperature is then given by:

$$\theta_{\text{out}} = \alpha \cdot \theta_{\text{in}} + (1 - \alpha) \cdot \bar{\theta}_{\text{out}} \tag{4}$$

where $\bar{\theta}_{\text{out}}$ represents the volume-weighted average temperature of the portion of outflow originating from water already in the pipe at the start of the time step. The process is illustrated in Figure 2. Consumers are modeled as fixed

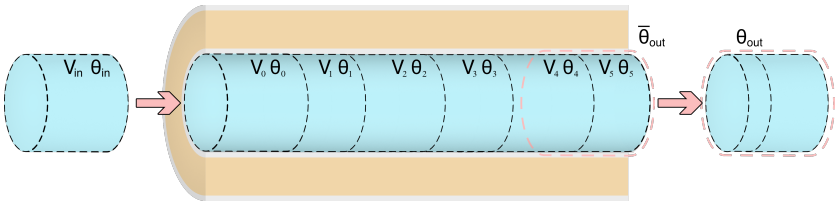


**Fig. 2.** Illustration of the pipe model. In this case, the incoming volume is smaller than the total pipe volume ($\alpha = 0$), and the outlet temperature is computed as the volume-weighted average temperature of the portion of water highlighted by the red dotted border.

temperature difference components, thus we enforce:

$$\theta_{\text{out}} = \theta_{\text{in}} - \Delta\theta_{set} \tag{5}$$

Finally, the outlet temperature of producers is imposed as boundary condition independently from the inlet temperature.

**Overall convergence.** The convergence criterion is based on nodal energy conservation. At each time step, we assume that the orientation of the $j$-th edge aligns with the positive direction of the flow. As a first condition to determine the temperature $\theta$ at each node in the network, we impose that the inlet temperature of an edge matches the temperature of its starting vertex:

$$\theta_{\mathrm{in},j} = \theta_{t(j)}, \quad \forall j \in \mathcal{E}. \tag{6}$$

Where the notation $t(j)$ indicates the starting vertex (tail) of edge $j$.

We then set the ending vertices of producers as slack nodes, indicated with the subscript $s$, and impose their temperatures as that of the corresponding unique incoming edge $\theta_{\mathrm{imp},i}$ (°C):

$$\theta_i = \theta_{\mathrm{imp},i}, \quad \forall i \in \mathcal{V}_{\mathrm{s}}. \tag{7}$$

Finally, we impose the conservation of energy at the remaining vertices. Following (1) and (6), this can be written as:

$$\sum_{j \in \mathcal{I}_i} \dot{m}_j \cdot c_p \cdot \psi_j(\theta_{t(j)}) - \left( \sum_{k \in \mathcal{O}_i} \dot{m}_k \right) \cdot c_p \cdot \theta_i = 0, \quad \forall i \in \mathcal{V} \setminus \mathcal{V}_{\mathrm{s}}. \tag{8}$$

where $\mathcal{I}_i$ is the set of edges entering node $i$, $\mathcal{O}_i$ is the set of edges leaving node $i$, $\dot{m}_j$ is the mass flow rate (kg/s) through edge $j$, $\theta_{h(j)}$ is the temperature (°C) of the upstream node of edge $j$ and $\theta_i$ is the temperature (°C) at node $i$. $c_p$ is the specific heat capacity that we assume to be constant at 4180 $\mathrm{J}/(\mathrm{kg} \cdot \mathrm{K})$ within the energy balance calculation. On the other hand, it ensures linearity in the system (8), provided the edge-wise temperature transfer functions $\psi_j(\cdot)$ are themselves linear. As this condition is met using the components model previously described, combining (7) and (8) we obtain a square system of linear equations in the form:

$$\mathbf{A}\vec{\theta} = \vec{b} \tag{9}$$

that can be solved using LU factorization. Note that if non-slack nodes with zero mass flow are present, $\mathbf{A}$ is singular. Therefore, the equations of such nodes are removed from system (9) and solved separately. In particular, these nodes are assigned with the average value of the temperatures of adjacent water volumes in pipes.

### 2.3 Control Strategies

We consider the following heating curves: a constant model and a conventional piece-wise linear model. In both cases, we test variants with pre-heating, night setback or both. The piece-wise linear model is controlled by 4 (learnable) parameters, $\theta_a, \theta_b, \theta_c, \theta_d$:

$$\theta_{out} = f_l(\theta_{\mathrm{air}}) = \begin{cases} \theta_c, & \theta_{\mathrm{air}} < \theta_a \\ \theta_c + \frac{(\theta_{\mathrm{air}} - \theta_a)(\theta_d - \theta_c)}{\theta_b - \theta_a}, & \theta_a \leq \theta_{\mathrm{air}} \leq \theta_b \\ \theta_d, & \theta_{\mathrm{air}} > \theta_b \end{cases} \tag{10}$$

With the requirement that $\theta_a < \theta_b$. The model with pre-heating has 3 additional parameters: the starting time $t_0$, the duration $\mathcal{T}$, and the temperature difference $\Delta$. Given $t_1 = t_0 + \mathcal{T}$ and assuming $t_0 < t_1$:

$$\theta_{out} = f_l(\theta_{\text{air}}) + \Delta^* \tag{11}$$

$$\Delta^* = \Delta\left[\sigma(k(t - t_0)) - \sigma(k(x - t_1)\right] \tag{12}$$

Where $\sigma$ is the sigmoid function and $k = 100$. Here, $\Delta^*$ is used to avoid discontinuity induced by a sudden change of temperature. The same formulation is used for setback, with the temperature delta constrained to be negative.

### 2.4 Optimization Approach

Given a control strategy, we seek to find its optimal sets of parameters - one for MAIN and one for SEC1 - and compare the performance of the optimized strategy with the manually controlled supply temperatures used during the monitoring period. We frame this optimization problem as follows:

$$\min_P \left( \sum_{t=1}^{T} \sum_{i=1}^{\mathcal{E}_p} Q_i^{(t)}(P) + \lambda \sum_{t=1}^{T} \sum_{j=1}^{\mathcal{E}_{c\setminus r}} \mathcal{Q}_j^{(t)}(P) \right) \tag{13}$$

$$\mathcal{Q}_j^{(t)} = \begin{cases} c_p \dot{m}_j^{(t)} \Delta t \Delta\theta_j^{(t)}, & \Delta\theta_j^{(t)} < 0 \\ 0, & \text{otherwise} \end{cases} \tag{14}$$

$$\Delta\theta_j^t = \theta_{\text{ret,min},j} - \theta_{\text{ret},j}^{(t)} \tag{15}$$

Here, $\mathcal{E}_p$ represents the set of pipes where energy losses $Q_i^{(t)}(P)$ are evaluated, while $\mathcal{E}_{c\setminus r}$ denotes the subset of consumers excluding ramps. The penalty term $\mathcal{Q}_j^{(t)}(P)$ accounts for violations of the theoretical minimum return temperature $\theta_{\text{ret,min},j}$, with $\Delta\theta_j^{(t)}$ measuring the deficit between this minimum and the simulated return temperature $\theta_{\text{ret},j}^{(t)}$. The reference value $\theta_{\text{ret,min},j}$ is estimated as the minimum return temperature observed during the monitoring period.

Since the dynamic simulation model described in Section 2.2 and the control strategies presented in Section 2.3 are differentiable, we can adjust the parameters $P$ using gradient descent. At each $k$-th iteration of gradient descent, the full trajectory of $T$ time steps period is evaluated with a fixed set of parameters $P^k$, which is then updated as:

$$P^{(k+1)} = P^{(k)} - \eta \nabla \mathcal{L}(P^{(k)}), \tag{16}$$

Where $\mathcal{L}$ is the objective function to minimize in Eq. (13) and $\eta$ is the learning rate. In addition to the $T$ time steps, an initial period of 24 hours is also simulated at each iteration using the conditions of the first time step to initialize the temperatures of water volumes and pipe walls.

For each heating curve variant we perform 50 steps of gradient descent with an initial learning rate of 0.01. If there are no improvements on the loss for more than 5 iterations, the learning rate is reduced by half, up to a minimum of 0.0001. To prevent large parameter updates in a single step that could lead to instabilities, the update value for temperature-related parameters is clipped to the range [-1, 1]. The $\lambda$ factor in Eq. (13) is chosen as 1000. Further analysis is conducted using the optimal parameter set found for each heating curve.

## 3 Results and Discussion

Results of the experiments are given in Table 1, along with the losses resulting from the manual control as a reference. On average, the piecewise-linear

**Table 1.** Results: total heat losses over the 5 days considered for each heating curve.

| **Model** | **Heat loss** | |
|---|---|---|
| | **Value (kWh)** | **Diff From Best (%)** |
| *Manual* | 42921 | +0.66 |
| *Constant* | 42831 | +0.45 |
| *Constant, pre-heating* | 42638 | - |
| *Constant, setback* | 43286 | +1.51 |
| *Constant, pre-heating, setback* | 43359 | +1.69 |
| *Linear* | 43446 | +1.90 |
| *Linear, pre-heating* | 43237 | +1.40 |
| *Linear, setback* | 43361 | +1.70 |
| *Linear, pre-heating, setback* | 43895 | +2.95 |

model shows worse performance than a constant supply temperature value. This counterintuitive result likely stems from the operational strategy used during the monitoring period, which was used as our boundary condition. The best supply temperatures found for the base constant model are 81.0°C for MAIN and 82.0°C for SEC1, slightly lower than the setpoints enforced by the manual control. This allowed to reduce heat losses to the ground while satisfying the constraints for the temperature at consumer substations. Interestingly, models with setback performed worse, and in most cases either the setback duration or magnitude was reduced to almost zero during the optimization. Pre-heating was beneficial for reducing losses, but contrarily to its usual usage it was pushed to the afternoon and not used for shaving the peak demand in the morning. Similarly to the manual control used, its starting times converged either around 11 AM, 14 PM or 22 PM, which correspond to sharp decreases in mass flow rate at MAIN.

These adaptations suggest that the gradient descent method is successfully navigating the complex solution space of the thermal model, confirming the viability of the approach. While our analysis was constrained to a 5-day period, this

methodology has the potential for extension to longer time-frames, provided that the computation of gradients is adapted accordingly. On this regard, techniques such as Backpropagation Through Time (BPTT) with truncated horizons might be needed for maintaining numerical stability and managing computational costs for longer trajectories. The approach also remains viable for larger networks, provided that constraints are placed on the number of water volumes modeled simultaneously across all pipes, for example by periodically joining adjacent volumes.

## 4 Conclusion and Future Work

In this work, we proposed a gradient-based approach to DHN control optimization using a differentiable dynamic thermal model. In particular, we used it as a mean to tune heating curves for a real network based on monitoring data over a period of 5 days. Despite the limitations imposed by using boundary conditions on mass flow that are not correlated with the outdoor air temperature, results show a reduction in heat losses compared to manual operations. This approach eliminates the need for extensive trial-and-error tuning, and consequently, periods of suboptimal operations. Future work should expand this approach to longer time horizons, which will require addressing computational challenges in gradient calculation over extended periods. Additionally, extending the model to include the hydraulic aspects of the network would provide a more complete optimization framework, that could allow for a higher energy saving potential.

## Acknowledgments

R.B. acknowledges additional support from the EPFLGlobalLeaders program, which has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 945363.

## References

1. World Energy Outlook 2023, IEA, Paris (2023) https://www.iea.org/reports/world-energy-outlook-2023
2. Zuberi, M.J.S., Narula, K., Klinke, S., Chambers, J., Streicher, K.N., Patel, M.K.: Potential and costs of decentralized heat pumps and thermal networks in Swiss residential areas. International Journal of Energy Research, 45(10), 15245-15264 (2021). https://doi.org/10.1002/er.6801
3. Lund, H., Østergaard, P.A., Nielsen, T.B., Werner, S., Thorsen, J.E., Gudmundsson, O., Arabkoohsar, A., Mathiesen, B.V.: Perspectives on fourth and fifth generation district heating. Energy, 227, 120520 (2021). https://doi.org/10.1016/j.energy.2021.120520

4. Ghane, S., Jacobs, S., Casteels, W., Brembilla, C., Mercelis, S., Latré, S., Verhaert, I., Hellinckx, P.: Supply temperature control of a heating network with reinforcement learning. In: 2021 IEEE International Smart Cities Conference (ISC2) (pp. 1-7). IEEE (2021). https://doi.org/10.1109/ISC253183.2021.9562966
5. Boghetti, R., Kämpf, J.H.: Verification of an open-source Python library for the simulation of district heating networks with complex topologies. Energy, 290, 130169 (2024). https://doi.org/10.1016/j.energy.2023.130169
6. Denarie, A., Aprile, M., Motta, M.: Heat transmission over long pipes: New model for fast and accurate district heating simulations. Energy, 166, 267-276 (2019). https://doi.org/10.1016/j.energy.2018.09.186
7. Popiel, C. O., Wojtkowiak, J.: Simple formulas for thermophysical properties of liquid water for heat transfer calculations (from 0 C to 150 C). Heat transfer engineering, 19(3), 87-101 (1998). https://doi.org/10.1080/01457639808939929