

Optimizing Supply Temperature Control in District Heating Networks via Differentiable Dynamic Simulation and Gradient Descent

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

Abstract. District Heating Networks (DHNs) are central to sustainable energy transition, but their efficiency depends on effective supply temperature control. Rule-based strategies, such as heating curves, remain common for their simplicity, but tuning them often requires long periods of suboptimal operation. In this paper, we propose an alternative approach that directly optimizes control parameters via gradient descent, based on a differentiable dynamic simulation model. 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 compared to both manual regulation and gradient-free optimization, 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 simple rules or linear dependencies on outdoor temperature, known as heating curves [4]. These curves are typically tuned over long periods, often a full year, with parameters slowly adjusted based on observed performance. While straightforward and stable, this approach leads to long periods of suboptimal control and increased 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 both manual settings and gradient-free optimization, demonstrating their potential to increase 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 0.8 m, with some passing through basements or parking lots.

Flow, temperature, and pressure data were collected (Jan 10–22, 2022) and resampled at 15-minute intervals. 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. 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.

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. First, we outline the modeling assumptions, then we describe the components and convergence strategy.




Fig. 1. Daily variations in monitoring data, showing mean daily values and ranges. Left: total power demand (kW) and air temperature (°C). Right: mass flow (kg/s) and supply temperature (°C) at the MAIN plant. Shaded regions indicate the variability.

Modeling assumptions. We assume that the DHN is a closed system, with no leaks or injections of water, represented as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with edges as components (pipes, consumers, producers) and vertices as 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, θ_{out} (°C), as a function of the inlet temperature, θ_{in} (°C):

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

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. Water is discretized into variable-sized volumes with 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 θ_v from time step $t - 1$ to t is given by:

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

where $h = 3600$ J/Wh, and the water density ρ 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 pipe outlet temperature is obtained as the volume-weighted sum of exiting water temperatures. If the incoming volume V_{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 α as:

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

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

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

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. Here, the incoming volume is less 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 dashed border.

temperature difference components, thus we enforce:

$$\theta_{\text{out}} = \theta_{\text{in}} - \Delta\theta_{\text{set}}. \quad (5)$$

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

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 θ at each node in the network, we impose that the inlet temperature of an edge matches the temperature of its starting vertex:

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

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 $\theta_{\text{imp},i}$ ($^{\circ}\text{C}$):

$$\theta_i = \theta_{\text{imp},i}, \quad \forall i \in \mathcal{V}_s. \quad (7)$$

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

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

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 , and θ_i is the temperature ($^{\circ}\text{C}$) at node i . c_p is the specific heat capacity that we assume to be constant at $4180 \text{ J}/(\text{kg} \cdot \text{K})$ within the energy balance calculation. This assumption neglects the temperature dependence of c_p , which has limited impact over typical DHN operating ranges ($0\text{--}120^{\circ}\text{C}$). However, it ensures linearity of Equation System (8), provided the edge-wise temperature transfer functions $\psi_j(\cdot)$ are themselves linear. As this condition is met using the component models previously described, combining (7) and (8) we obtain a square system of linear equations in the form:

$$\mathbf{A}\vec{\theta} = \vec{b}, \quad (9)$$

that is 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 assigned with the average value of the temperatures of adjacent water volumes in pipes.

2.3 Control Strategies

We consider two base heating curves: a constant model and a piecewise linear model. For both, we test variants with preheating, night setback, or both. The piecewise linear model is controlled by 4 (learnable) parameters, $\theta_a, \theta_b, \theta_c, \theta_d$:

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

with the requirements that $\theta_a < \theta_b$ and $\theta_d < \theta_c$. The models with preheating have 3 additional parameters: the starting time t_0 , the duration $\mathcal{T}$, and the temperature difference Δ . Given $t_1 = t_0 + \mathcal{T}$ and assuming $t_0 < t_1$:

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

$$\Delta^* = \Delta [\sigma(k(t - t_0)) - \sigma(k(t - t_1))] . \quad (12)$$

Here σ is the sigmoid function and $k = 100$. Δ^* is used to avoid discontinuity induced by a sudden change of temperature. For daily transitions, t is interpreted over a continuous daily cycle. 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, for MAIN and SEC1, and compare the performance of the optimized strategy against both the manually controlled supply temperatures used during the monitoring period and a gradient-free baseline. We frame this optimization problem as follows:

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

$$Q_j^{(t)} = \begin{cases} c_p \dot{m}_j^{(t)} \frac{\Delta t}{h} \delta_j^{(t)}, & \delta_j^{(t)} > 0 \\ 0, & \text{otherwise} \end{cases} \quad (14)$$

$$\delta_j^{(t)} = \theta_{\text{ret},j}^{\min} - \theta_{\text{ret},j}^{(t)} \quad (15)$$

Here, Δt (s) is the time-step size, $\mathcal{E}_p$ 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 $Q_j^{(t)}(P)$ accounts for violations of the minimum return temperature. The deviation $\delta_j^{(t)}$ (K) measures how much the simulated return temperature $\theta_{\text{ret},j}^{(t)}$ (°C) falls below the reference value $\theta_{\text{ret},j}^{\min}$ (°C), estimated as the minimum return temperature observed during the monitoring period.

Since the simulation model (Section 2.2) and the heating curves (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 is evaluated with a fixed set of parameters $P^{(k)}$, which is then updated as:

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

where $\mathcal{L}$ is the objective function to minimize in Eq. (13) and η 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 η of 0.01. If the loss does not improve after 5 iterations, η is reduced by half, up to a minimum of 0.0001. To prevent instabilities from large parameter updates in a single step, the update of temperature-related parameters is clipped to the range $[-1, 1]$. The λ 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 are shown in Table 1, alongside those obtained using a gradient-free method, Differential Evolution (DE). Under manual control, total heat losses amounted to 42,921 kWh, with average supply temperatures of 81.4°C (MAIN) and 81.7°C (SEC1), and standard deviations of 2.2°C and 1.3°C, respectively. For DE, we used SciPy's [8] implementation with default settings and ran each optimization for 40 minutes, the maximum time our method took across all heating curves. As DE does not require gradient computations, function evaluations are faster, reaching around 210 evaluations per run.

Our gradient-based approach outperformed DE in most cases, although a more rigorous study is needed to assess whether this holds across different initializations and DE settings. The performance gap widens as the number of parameters increases, suggesting that DE struggles with higher-dimensional parameter spaces within the fixed evaluation budget. The lowest losses were achieved using a simple constant curve with preheating, a strategy that closely mirrors the

Table 1. Performance of each curve optimized with our method and DE. Total heat loss differences relative to manual control (in kWh and %) are reported, along with the mean (μ) and standard deviation (σ) of supply temperatures, shown as (MAIN, SEC1). Bold values indicate the lowest heat loss for each variant; no bold means manual control performed best. Curve types: *C* (constant), *L* (linear), *p* (preheating), *s* (setback).

Curve	Loss Diff. (kWh)		Temp. μ ($^{\circ}\text{C}$)		Temp. σ ($^{\circ}\text{C}$)	
	Ours	DE	Ours	DE	Ours	DE
<i>C</i>	-90 (-0.2%)	16 (0.0%)	(81.0, 82.0)	(81.5, 81.1)	(0.0, 0.0)	(0.0, 0.0)
<i>C+p</i>	-283 (-0.7%)	284 (0.7%)	(80.8, 82.5)	(81.7, 82.1)	(2.6, 2.8)	(1.1, 1.1)
<i>C+s</i>	427 (1.0%)	278 (0.6%)	(82.0, 82.6)	(80.6, 86.3)	(2.8, 2.2)	(0.9, 1.5)
<i>C+p+s</i>	438 (1.0%)	1025 (2.4%)	(80.5, 88.0)	(84.0, 79.0)	(2.4, 1.5)	(2.3, 1.4)
<i>L</i>	526 (1.2%)	343 (0.8%)	(82.3, 82.3)	(82.5, 79.6)	(1.1, 1.1)	(0.0, 0.0)
<i>L+p</i>	69 (0.2%)	-39 (-0.1%)	(81.6, 82.1)	(81.4, 80.9)	(2.3, 2.5)	(0.4, 0.7)
<i>L+s</i>	105 (0.2%)	1330 (3.1%)	(81.0, 83.7)	(84.3, 80.8)	(1.0, 0.7)	(1.0, 3.2)
<i>L+p+s</i>	-82 (-0.2%)	1006 (2.3%)	(81.2, 82.9)	(81.8, 87.0)	(3.1, 3.0)	(2.4, 4.1)

manual operations observed during the monitoring period and used as boundary conditions. In this configuration, the supply temperature is held at 80.0 $^{\circ}\text{C}$ for MAIN and 81.7 $^{\circ}\text{C}$ for SEC1. preheating is activated at 10 PM for both plants, increasing the temperature by 9-10 $^{\circ}\text{C}$ for about 2 hours. This compensates for the decrease in mass flow and replicates the supply temperature rise observed under manual control.

In general, optimized curves consistently initiate preheating around sharp drops in MAIN’s mass flow, at 11 AM, 2 PM, or 10 PM, rather than during morning peaks. Setback variants performed poorly, with the optimizer minimizing their effect. These trends suggest that gradient descent effectively navigates the model’s solution space, confirming its viability.

Regarding scalability, the memory footprint of the simulation grows linearly with the number of water volumes, which is determined by the number and length of pipes, their flow rates, and the chosen time step size. In practice, the model can handle larger networks efficiently, and if memory becomes a constraint, adjacent volumes can be dynamically merged to reduce the footprint with minimal impact on accuracy. To compute gradients over longer horizons, the simulation trajectory must also be truncated, both to mitigate vanishing gradients and to reduce memory usage. At the same time, the physical influence of a control decision is inherently limited: it persists only while the affected water volume remains in the network or continues to influence pipe wall temperatures in a non-negligible way. This reinforces the case for truncation, typically implemented using (overlapping) Truncated Backpropagation Through Time (BPTT).

4 Conclusion and Future Work

In this work, we proposed a gradient-based approach to DHN control optimization using a differentiable dynamic thermal model. Specifically, we used it to tune

heating curves for a real network based on monitoring data over a 5-day period. Despite using boundary conditions from a monitoring period where mass flow was adjusted to demand, and supply temperature was not set based on outdoor air temperature, results show a reduction in heat losses compared to manual operation. Our method also outperformed a gradient-free baseline based on Differential Evolution in most cases, particularly for curves with more parameters.

Future work should expand this approach to longer time horizons, which will require addressing computational challenges in gradient calculation over extended periods of time. Additionally, extending the model to include the hydraulic aspects of the network would provide a more complete optimization framework, which could allow for a higher energy saving potential.

Acknowledgments

We thank ALTIS Groupe SA for providing the data used in this study.



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. *Int. J. Energy Res.*, 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. Boggetti, 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 transf. eng.*, 19(3), 87-101 (1998). <https://doi.org/10.1080/01457639808939929>
8. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., et al.: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods*, 17(3), 261-272 (2020). <https://doi.org/10.1038/s41592-019-0686-2>