Steady-state analysis of a small district heating loop: thermo-hydraulic state, energy accounting, and consumer duties inferred from return temperatures

Ricardo Marquez, Ph.D.
rmarquez123@gmail.com

September 2026

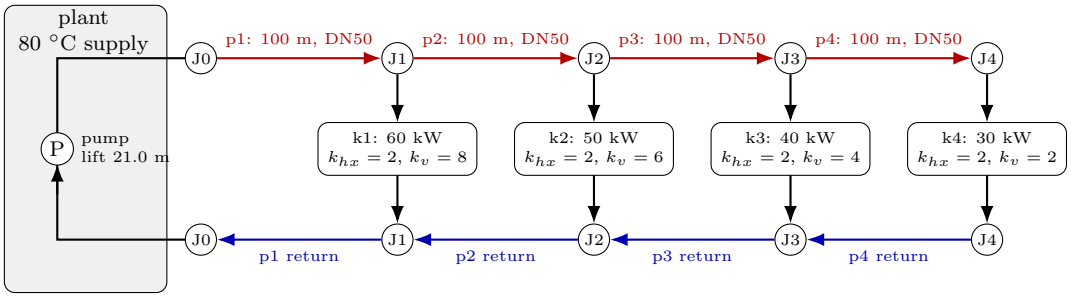
Abstract

A district heating network carries hot water from a plant through buried pipes to the buildings it serves and back. Its steady state follows from mass and energy conservation at the pipe junctions and from three laws on the pipes: a friction law for the pressure drop, an advection law for the heat the water carries, and a loss law for the heat that escapes to the ground. This study writes those laws for the smallest network with a return line and a pump, 4 substations of fixed duty on a 400 m supply line, and asks three questions of one set of equations. The state (heads, flows, temperatures) agrees with the open-source network simulator pandapipes, given the same constant fluid, to below 10^{-9} K and below 10^{-9} m; with pandapipes' temperature-dependent properties the return line differs by up to 0.93 K, a difference of input rather than of formulation. Two ledgers account for the supplied energy: of the plant's 183.8 kW, 97.9 % reaches the consumers and 2.1 % is lost to the ground; of the pump's 724 W, 88 % is spent on pipe friction and the rest in the substations' exchangers and balancing valves. Both ledgers close to round-off, the ground losses and the pump power agree with pandapipes to below 10^{-9} relative, and the friction heating both toolchains omit is 0.049 K. Finally, with one return-temperature sensor per substation and the plant's flow and return-temperature readings, a maximum-a-posteriori fit of the same equations recovers every consumer's duty to 1.0–1.4 % without a heat meter; over 20 noise draws 79 of 80 95 % intervals cover the truth, against 76 expected. Every number costs 21 engine solves.

1 Introduction

District heating supplies a large share of the heat used in the cities of northern and eastern Europe and is growing elsewhere; Werner (2017) reviews its extent and its technology. A network of this kind is a closed water loop: a plant heats water, a pump drives it through buried supply pipes to substations in the buildings, where heat exchangers transfer the heat to the buildings' own circuits, and the cooled water returns to the plant. Two questions follow an operator through every hour of such a network's life. What is the state of the loop, that is, what are the temperatures, pressures and flows at the points that are not measured, and where does the energy the plant supplies go, to the consumers, to the ground through the pipe insulation, or into pumping. The second question has gained weight with the move to lower supply temperatures that Lund et al. (2014) describe, which cuts the ground losses but raises the flows and hence the pumping power, and with the cost of heat meters, which not every substation carries.

The equations of the steady state are settled physics, and network simulators solve them routinely; pandapipes (Lohmeier et al., 2020) is an open-source one and serves here as the independent



ground at 8 °C, $U = 0.5$ W/(m^2 K) on every pipe; supply line above, return line below; k1–k4 are the substations

Figure 1: The network of the study: a plant holds the supply header at 80 °C and a circulation pump at a constant lift drives the water along a buried supply line to 4 substations and back along the return line. Each substation draws a fixed duty through an exchanger branch (loss coefficient k_{hx}) and a balancing valve (k_v, closing toward the plant). Every pipe is 100 m of DN50 steel with 0.1 mm roughness.

check. Writing the same conservation laws as a factor graph, a set of weighted residuals over the network's unknowns, and inferring the unmeasured state from sparse readings is the construction of Han, Eguchi, Mehrotra and Venkatasubramanian (2018) for water networks after an earthquake; it rests on the sparse least-squares machinery that Dellaert and Kaess (2017) describe for robotics. Applied to a heating network, the same construction turns a temperature sensor on a substation's return pipe into a heat meter, because the network's equations supply the flow the meter would otherwise have to measure. Nothing in this report is new; its purpose is to state the equations, check them against an independent solver on a network small enough to read, and fix the form in which larger heating studies report.

TERRA is a general-purpose engine for probabilistic inference on physical networks. It treats a network as a graph on which the physics is written locally: a conservation law at every node, a constitutive law on every edge, and, attached to the same graph, whatever the sensors report and whatever is known in advance. Each of these is a relation that holds only to within a stated tolerance, and the engine weighs them together to find the most probable state of the system, the uncertainty of every quantity in it, and the relative plausibility of competing explanations of the same readings. Domains enter through plugins that supply their laws, units and file formats; the inference and the bookkeeping are shared. This study uses its hydraulics plugin to write a heating loop's mass and energy balances, pipe laws, heat-loss laws and substation laws as one set of residuals, solves the loop's state and checks it against an independent network simulator, accounts for the heat and the pumping power in two ledgers, and infers the consumers' duties from return-temperature readings. Section 3 states the framework and each of these constructs precisely.

This report asks three questions of a 4-consumer loop. What is its steady state, and do the equations reproduce an independent solver? Where do the plant's heat and the pump's power go, and do the two accounts close? And can each consumer's duty be read from a return temperature alone, with an honest uncertainty? Section 2 states the network, the data and the model, Section 3 the methods, Section 4 the results, Section 5 their limits, Section 6 the answers, and the appendices the reproduction recipe, the data and the references.

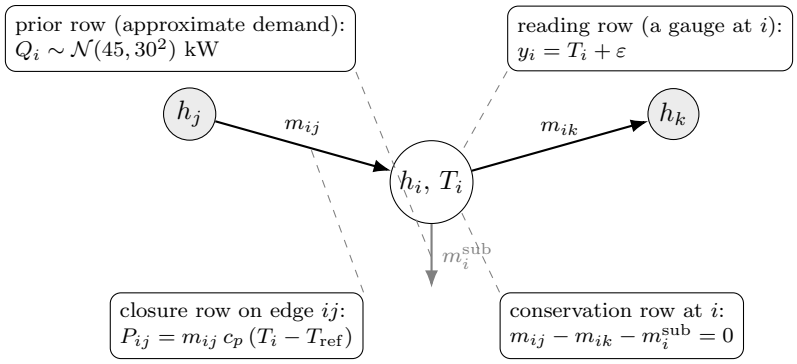


Figure 2: The model around one header of the loop. Unknowns are the head and temperature at the header, the mass flows on its pipes and into its substation and, at the substation, the duty; the equations tying them are a mass balance and an energy balance at the header, closure laws on each pipe (a friction law for the head drop, an advection law for the heat carried, a loss law for the heat that reaches the ground), a reading row per sensor and a prior row per approximately known quantity. Every row carries a weight; the fit minimizes the weighted squared misfits of all rows at once, and the uncertainty of each unknown follows from the curvature of that misfit. The same picture holds at every header; only the laws change.

2 System and data

The network. Figure 1 draws the loop. The plant holds its supply header J0 at 80 °C; a circulation pump lifts the returning water by 20.99 m of head (the pressure difference of 2 bar in meters of water column) into the supply header, and the return header is held at the head of 31.48 m that an expansion vessel would fix. From J0 a supply line of 4 pipes, each 100 m long with an inner diameter of 50 mm and a roughness of 0.1 mm, leads to headers J1 to J4; a return line of the same pipes leads back. At each header a substation takes water from the supply header, passes it through a heat exchanger of loss coefficient 2 and a balancing valve of loss coefficient 8, 6, 4, 2 (k1 to k4), both of 30 mm bore, and returns it to the return header. The substations draw 60, 50, 40, 30 kW, 180 kW in all. The pipes are buried in ground at 8 °C and lose heat through their insulation with an overall coefficient of 0.5 W/(m^2 K) on the inner surface, 7.85 W/K per pipe.

The data. Everything is synthetic and stated. The water has density 971.7 kg/m^3, viscosity 3.534×10^{-4} Pa s and specific heat 4,197 J/(kg K), the values of pandapipes' water tables at 80 °C, held constant so that the two toolchains solve the same equations. The reference state is pandapipes 0.14.0 solving the same loop, once with that constant fluid and once with its temperature-dependent tables. For the inference question the truth is the loop's own solution, and the readings are that truth plus Gaussian noise of 100 mK on temperatures and 1 % on the plant's flow, drawn afresh for each of 20 seeds. No field data are used.

The model. Figure 2 shows the equations around one header. The whole loop has 70 unknowns (heads and temperatures at 10 headers and 4 substation outlets, mass flows on 8 pipes, 8 substation branches and the pump, the heat carried and lost on every pipe, the plant's duty and the pump's power) and as many physical relations, stated in Section 3; the readings add 6 reading relations and the free duties 4 unknowns with priors.

The questions. (1) What are the heads, flows and temperatures of the loop, and do the equations reproduce pandapipes? (2) Of the heat the plant supplies and the power the pump supplies, how much goes where, and do the two accounts close? (3) Can each consumer's duty be inferred from one return-temperature sensor and the plant's readings, with an uncertainty that is honest?

3 Methods

3.1 The district heating factor graph

Each header i carries a head h_i (m) and a temperature T_i (K); each pipe e from header i to header j carries a mass flow m_e (kg/s), the heat P_e (W) it delivers to j and the heat L_e (W) it loses to the ground on the way. Each physical law is written as a residual, an expression that vanishes when the law holds and carries a weight, the inverse variance of how far it may be off; Section 3.2 makes this the general rule. Mass and energy are conserved at every header,

$$r_i^m = \sum_{e\to i} m_e - \sum_{e\leftarrow i} m_e - m_i^{\mathrm{sub}}, \tag{1}$$

$$r_i^E = \sum_{e\to i} P_e - \sum_{e\leftarrow i} P_e - \sum_{e\to i} L_e - Q_i, \tag{2}$$

where m_i^{sub} is the flow a substation draws (zero at a header without one, and the flow it returns enters the return header's balance with the opposite sign) and Q_i is a duty leaving the energy balance, present only at a substation's outlet node. On every pipe three laws hold: the head drop follows the Darcy–Weisbach law with the friction factor of Swamee and Jain, the heat carried is the enthalpy of the flow relative to a reference temperature T_{ref}, and the heat lost is the exact loss of plug flow through a pipe buried in ground at T_g with overall conductance UA,

$$r_e^h = h_i - h_j - \frac{8 f_e L_e}{\pi^2 g D_e^5 \rho^2}\, m_e|m_e|, \tag{3}$$

$$f_e = \frac{0.25}{\left[\log_{10}\!\left(\frac{\epsilon}{3.7 D_e} + \frac{5.74}{\mathrm{Re}_e^{0.9}}\right)\right]^2}, \tag{4}$$

$$r_e^P = P_e - m_e c_p (T_i - T_{\mathrm{ref}}), \tag{5}$$

$$r_e^L = L_e - m_e c_p (T_i - T_g)\left(1 - e^{-UA/(m_e c_p)}\right), \tag{6}$$

with $\mathrm{Re}_e = 4m_e/(\pi D_e \mu)$; the friction factor is one smooth function of the Reynolds number that reduces to $64/\mathrm{Re}$ in laminar flow and to (4) above $\mathrm{Re} = 4000$, which is where every pipe of this loop runs (Table 2). The loss law (6) is exact for a uniform ground temperature: the water leaving the pipe is at $T_g + (T_i - T_g)\, e^{-UA/(m_e c_p)}$, and L_e is the enthalpy it has given up. The enthalpy is carried relative to $T_{\mathrm{ref}} = 45$ °C, the design return temperature, so that the energy variables stay small; the reference cancels in every balance. A substation is a node X_k between two branches, each a minor loss of coefficient k and bore A with the advection law (5) on it,

$$r_b^h = h_i - h_j - \frac{k_b}{2 g A^2 \rho^2}\, m_b|m_b|, \tag{7}$$

$$r_{X_k}^E = m_k c_p (T_i - T_{\mathrm{ref}}) - m_k c_p (T_{X_k} - T_{\mathrm{ref}}) - Q_k, \tag{8}$$

so that the duty Q_k is what the primary water gives up between the supply header and the substation's outlet temperature T_{X_k}, which is also the temperature of the water it returns. The

pump holds a constant lift H between the return and the supply header and its hydraulic power is the product of flow and lift,

$$r^{\mathrm{pump}} = h_{J0,s} - h_{J0,r} - H, \tag{9}$$

$$r^{W} = W_{\mathrm{pump}} - g\, m_{\mathrm{pump}} H. \tag{10}$$

The plant's duty Q_{plant} enters the supply header's energy balance as a source, and the supply temperature there is held at 80 °C. The return header holds the loop's reference head and carries no mass balance, because the mass balances of a closed loop are dependent by one; its energy balance fixes the mixed return temperature. The physics residuals are weighted at $(1\ \mathrm{g/s})^{-2}$ on mass rows, $(1\ \mathrm{mm})^{-2}$ on head rows and σ_E^{-2} on energy rows,

$$\sigma_E = \max\!\big(m_{\mathrm{ref}}\, c_p \cdot 10^{-3}\ \mathrm{K},\ \sigma_m\, c_p\, (T_s - T_{\mathrm{ref}})\big)\,, \tag{11}$$

here 147 W: the energy slack that a mass slack of $\sigma_m = 1$ g/s implies at the largest temperature difference the loop carries, so that the two kinds of rows are trusted consistently. The readings and the priors attach to the same unknowns as residuals of the same form,

$$r = y - T_{X_k}, \tag{12}$$

$$r = Q_k - \bar{Q}, \tag{13}$$

weighted by σ_T^{-2} with $\sigma_T = 100$ mK (and σ_m^{-2} on the plant's flow reading, 1 % of the flow) and by τ^{-2} with $\bar{Q} = 45$ kW and $\tau = 30$ kW, a prior wide enough to say almost nothing about a duty of tens of kilowatts. In the inference the energy rows are written in kilowatts, a change of unit that leaves the model unchanged and keeps the information carried by a duty within the range the solver's rank test resolves.

3.2 The probabilistic model

The residuals of the previous subsection are all of one kind, and the engine treats them as such; this subsection states the general model they belong to. A physical network is a graph, and the engine writes its physics on that graph. The state is a vector z of unknowns: one potential per node, one flow per edge, and whatever else is uncertain, such as demands or parameters. Every relation the state must satisfy is a factor, written as a residual $r_k(z)$ that vanishes when the relation holds and carrying a weight w_k, the inverse variance of how far it may be off. A conservation law at a node, a constitutive law on an edge, a sensor reading and a prior on an approximately known quantity are all factors of this one kind; they differ only in their residual and their weight, a physics relation being trusted to a small tolerance, a reading to its sensor's noise, a prior to its stated spread. Treating each factor as an independent Gaussian belief about its residual gives the posterior density of the state given the readings y,

$$p(z \mid y) \propto \prod_k \exp\!\big(-\tfrac{1}{2}\, w_k\, r_k(z)^2\big). \tag{14}$$

The state that maximizes it, the *maximum-a-posteriori* (MAP) state z^*, minimizes the weighted sum of squared residuals $C(z)$ and is found by Gauss–Newton iteration on the sparse Jacobian J of the residuals, $J_{ki} = \partial r_k / \partial z_i$, with W the diagonal matrix of weights, from a flat start until the

gradient norm falls below a tolerance ϵ:

$$C(z) = \tfrac{1}{2} \sum_k w_k \, r_k(z)^2, \tag{15}$$

$$z^* = \arg\min_z C(z), \tag{16}$$

$$z \leftarrow z - (J^\top W J)^{-1} J^\top W \, r(z) \quad \text{until } \|J^\top W \, r(z)\| < \epsilon. \tag{17}$$

Near the minimum the cost (15) is approximately quadratic, so the posterior is approximately Gaussian about z^* with information matrix Λ. This is the *Laplace approximation*, and the *posterior standard deviation* σ_i of any unknown is read from the diagonal of Λ^{-1}, from the same sparse factorization the last Gauss–Newton step used:

$$p(z \mid y) \approx \mathcal{N}\big(z^*, \Lambda^{-1}\big), \tag{18}$$

$$\Lambda = J^\top W J, \tag{19}$$

$$\sigma_i = \sqrt{(\Lambda^{-1})_{ii}}. \tag{20}$$

When several factor graphs compete to explain the same readings — a leak here, a leak there, no leak — each is compared by its *evidence* $Z = \int \prod_k \exp(-\frac{1}{2} w_k r_k(z)^2)\, dz$, the probability of the readings under the model once the extra freedom of its unknowns has been paid for. In the Laplace approximation the evidence has a closed form, and normalizing every factor by its own weight lets graphs with different numbers of unknowns d and factors compare as densities,

$$\log Z = -C(z^*) + \tfrac{d}{2} \log 2\pi - \tfrac{1}{2} \log\det \Lambda + \sum_k \tfrac{1}{2} \log \frac{w_k}{2\pi}. \tag{21}$$

Differences of $\log Z$ between hypotheses, from (21), are in *nats*, natural-log units of odds; five nats is odds of about 150 to 1. The posterior probability of a hypothesis is its evidence normalized over all hypotheses under equal prior odds. Everything that follows is a choice of variables and factors on this graph, followed by these three operations.

3.3 The reference solution

pandapipes (Lohmeier et al., 2020) solves the same loop: the pipes with the same length, diameter, roughness, heat-transfer coefficient and ground temperature, a circulation pump holding 5 bar at the supply header and a lift of 2 bar, and each substation as a heat exchanger of fixed duty between two valves carrying the exchanger's and the balancing valve's loss coefficients (its heat-exchanger element applies a loss coefficient with the bore of its default diameter rather than the one given, so the exchanger's own loss is carried by a valve in series). It runs with the Swamee–Jain friction factor and tolerances of 10^{-9}. Two fluids are used. With a constant fluid at the properties above the two toolchains differ only in formulation, and the comparison is a parity target: temperatures within 10^{-3} K, heads within 10^{-3} m, flows, pump power and substation head drops within 10^{-4} relative, and every pipe's ground loss within 0.1 W or 10^{-4} relative. With pandapipes' water tables the density, viscosity and specific heat change along the return line, which is 12 K colder than the supply; that comparison is a reported difference of input, not a target. Wherever two toolchains disagree, the source is one of five: input, formulation, scenario, solver or aggregation; the tables name it.

3.4 The two ledgers

Summing the energy balances (2) over every header of the solved loop gives the thermal ledger, and summing the head drops (3) and (7) weighted by their flows around the loop gives the mechanical one,

$$Q_{\text{plant}} = \sum_k Q_k + \sum_e L_e, \tag{22}$$

$$g\, m_{\text{pump}} H = \sum_{e \in \text{pipes}} g\, m_e (h_i - h_j) + \sum_k g\, m_k (h_{k,s} - h_{k,r}), \tag{23}$$

where the second sum of (23) is each substation's dissipation, split between its exchanger branch and its balancing valve. Both are identities of the solved system, so their closing tests the solve, not the physics; the physics is tested by comparing the terms with pandapipes. The plant's duty is also $m_{\text{pump}} c_p (T_s - T_r)$ from the plant's own flow and temperatures, which the ledger reports as a check of the same identity from the other side. Neither toolchain returns the pump's work to the water as heat; the temperature rise that would add is

$$\Delta T_{\text{fric}} = \frac{gH}{c_p}, \tag{24}$$

and it is stated as a limit of both.

3.5 Virtual heat meters

The third question makes the 4 duties unknown, each with the prior (13), and adds 6 readings: the plant's flow m_{pump} and return temperature $T_{J0,r}$, and each substation's outlet temperature T_{X_k}, the sensor a substation's return pipe usually carries. The model is solved to its MAP state (16) and Laplace posterior (18); each duty's 95 % interval is ± 1.96 posterior standard deviations (20). Whether the readings fit the model at all is tested by the chi-square misfit of the reading residuals at the optimum,

$$\chi^2 = \sum_{k \in \text{readings}} w_k\, r_k(z^*)^2, \tag{25}$$

on as many degrees of freedom as readings minus fitted duties, rejected at $\alpha = 0.01$. The experiment is repeated for 20 noise draws so that the posterior can be checked against the spread of the estimates.

3.6 Metrics

The questions of Section 2 are answered with measures chosen so that each can be read against a reference. Agreement of the state is the largest absolute difference in temperature and in head over the headers and substation outlets, and the largest relative difference in flow, pump power, substation head drop and pipe loss, each against the parity target stated above, and separately against pandapipes' water tables, whose reference is the size of the property change along the return line. The ledgers are read as shares of the plant's duty and of the pump's power, and their closing as a relative residual whose reference is round-off; the ground losses and the pump power are read against pandapipes. The virtual meters are read by the posterior standard deviation of each duty relative to the duty, against the precision of a commercial heat meter (a few percent); by the root-mean-square error of the estimate across draws against the mean posterior standard

Table 1: Temperatures and heads at the ten headers: this study, pandapipes with the same constant fluid, and pandapipes with its temperature-dependent water tables. The constant-fluid columns agree with this study to below 10^{-9} K and below 10^{-9} m; the water-table columns differ by up to 0.93 K and 0.016 m, the effect of the properties' change along the return line.

node	T [°C]	T, pp const.	T, pp table	h [m]	h, pp const.	h, pp table
J0 (supply)	80.0000	80.0000	80.000	52.4710	52.4710	52.471
J0 (return)	67.5501	67.5501	67.546	31.4826	31.4826	31.483
J1 (supply)	79.9617	79.9617	79.962	44.0520	44.0520	44.042
J1 (return)	67.5818	67.5818	67.577	39.9016	39.9016	39.907
J2 (supply)	79.8753	79.8753	79.875	42.3415	42.3415	42.346
J2 (return)	61.2496	61.2496	61.139	41.6121	41.6121	41.624
J3 (supply)	79.6653	79.6653	79.662	42.0323	42.0323	42.045
J3 (return)	53.2592	53.2592	52.863	41.9213	41.9213	41.936
J4 (supply)	79.0736	79.0736	79.055	41.9878	41.9878	42.002
J4 (return)	47.4081	47.4081	46.480	41.9658	41.9658	41.982

deviation, which a calibrated posterior makes equal or smaller; by the number of 95 % intervals that cover the truth, about 76 of 80; and by the number of draws the misfit test (25) accepts, about 20 of 20. Cost is counted in engine solves rather than timed.

4 Results

The state. Table 1 lists the temperatures and heads at the ten headers and Figure 3 draws the temperatures along the loop. The supply line loses 0.93 K over 400 m; the return line, fed at its far end by k4's outflow at 47.4 °C, warms toward the plant as the nearer substations' warmer outflows mix in and reaches the plant at 67.55 °C. The pump moves 3.518 kg/s and the plant supplies 183.80 kW. With the same constant fluid pandapipes agrees to below 10^{-9} K in temperature, below 10^{-9} m in head and below 10^{-9} relative in flow, within the parity target by three orders of magnitude, so the two formulations are the same to the solvers' tolerances. With pandapipes' water tables the supply line is unchanged and the return line differs by up to 0.93 K and 0.016 m, with flows 2.7 % apart: a difference of input, the properties of 68 °C water against those of 80 °C water, and its size is what a study on a real network must carry into its uncertainty or remove by evaluating the properties per pipe.

Pipes and substations. Figure 4 draws the heads: the first pipe carries the whole flow and takes 8.42 m of the pump's 20.99 m on the way out and about as much on the way back, so the head available to the substations falls from 4.15 m at k1 to 0.022 m at k4. Table 2 lists the pipes: every one is turbulent, the four supply pipes each lose about 565 W whatever their flow, because the loss is set by the temperature difference to the ground and the conductance UA and only weakly by the flow, and the return pipes, 12 K colder or more, lose less. All eight losses agree with pandapipes within below 10^{-9} W. Table 3 lists the substations: their flows agree with pandapipes to below 10^{-9} relative and their head drops to below 10^{-9}; the far substations, with little head, run at large temperature drops, k4 cooling its water by 31.7 K.

The two ledgers. Table 4 and Figure 5 account for the supplied energy. Of the plant's 183.80 kW, 180 kW (97.9 %) reach the 4 consumers and 3.80 kW (2.1 %) go to the ground, 2.25 kW from the supply line and 1.55 kW from the return line; the ledger (22) closes below 10^{-9} relative, and

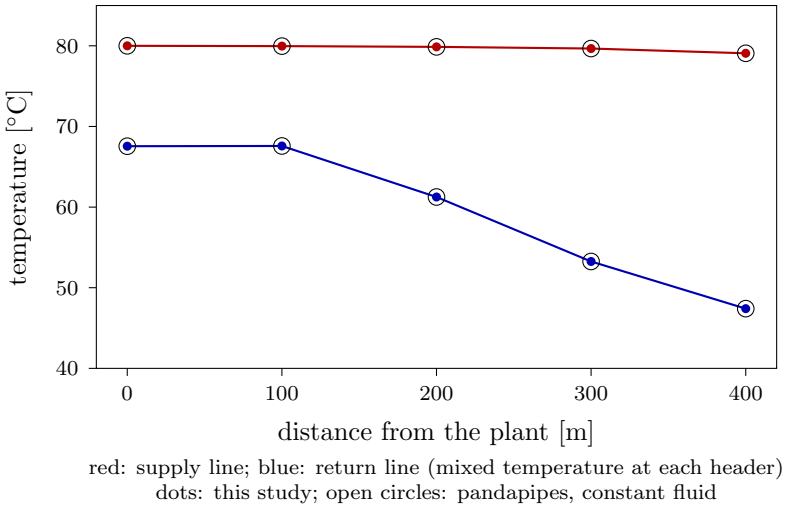


Figure 3: Temperature along the loop. The supply line loses 0.93 K over 400 m; the return line warms toward the plant as each substation's outflow mixes into it, from 47.4 °C at the far end to 67.6 °C at the plant. The pandapipes solution with the same constant fluid lies on the same points.

the plant's duty read from its own flow and temperatures agrees to the same. Of the pump's 724.0 W, 637.2 W (88.0 %) are dissipated by pipe friction, 17.8 W by the exchanger branches and 69.1 W (9.5 %) by the balancing valves, most of it at k1, which must throttle the largest surplus of head; the ledger (23) closes below 10^{-8} relative. The pump's power is 0.39 % of the plant's duty. The ground losses and the pump power agree with pandapipes to below 10^{-9} and below 10^{-9} relative, which is the independent check of both totals. The friction heating (24) that neither toolchain returns to the water is 0.049 K at this lift, below the sensor noise of the third question.

Virtual heat meters. Table 5 and Figure 6 report the duties inferred from 6 readings. Every duty is recovered to a posterior standard deviation of 1.0–1.4 % of its value, 489 W at k2 and 432 W at k4, from a 100 mK temperature sensor and no flow meter, because the loop's equations supply the flow: a duty is $m_k c_p$ times a temperature drop, and m_k follows from the pump's lift, the valves and the pipes. The root-mean-square error across draws is at or below the posterior standard deviation for every consumer, 79 of 80 intervals cover the truth against 76 expected, 79 of the 80 standardized errors lie within ± 1.96 and the largest is 2.20, so the posterior is honest and, at the far consumers, conservative. There the reading alone would fix the duty to $m_k c_p \sigma_T$, about 95 W at k4, and the posterior width is set instead by the energy rows' weight (11), the 147 W of slack the physics is allowed, which the synthetic truth never uses; a smaller mass-row weight would tighten it. The misfit test accepts 19 of 20 draws at $\alpha = 0.01$, about the 20 expected.

Cost. 1 Newton solve and 20 MAP solves (21 in all) produce every number in this report.

5 Discussion

Everything here is synthetic, steady and small by design, and those choices set the limits of the result. The fluid's properties are held at their 80 °C values so that the comparison with pandapipes tests the formulation alone; the water-table columns show the size of that choice, up to 0.93 K on a return line 12 K colder than the supply, and a study on a real network should evaluate the

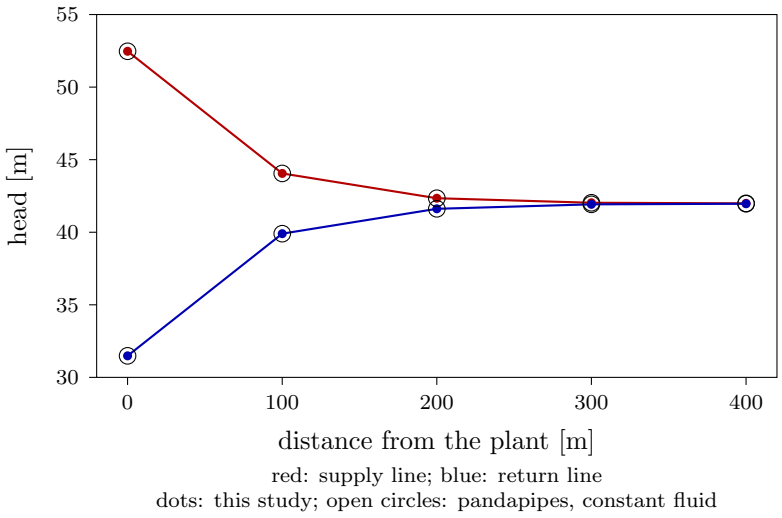


Figure 4: Head along the loop. The pump lifts the water by 20.99 m at the plant; the first pipe, carrying the whole flow, takes 8.42 m of it, and the head available to each substation shrinks from 4.15 m at k1 to 0.022 m at k4, which is why the balancing valves close toward the plant.

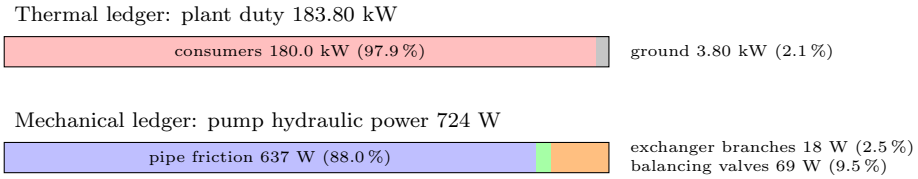


Figure 5: Where the supplied energy goes. Top: the plant's duty is the consumers' duties plus the heat lost to the ground by the eight buried pipes. Bottom: the pump's hydraulic power is dissipated by pipe friction and by the substations' exchanger branches and balancing valves; the segments too narrow to label are named beside the bar, in order. Each bar closes to 1.0×10^{-9} relative (thermal) and 1.0×10^{-8} relative (mechanical) by construction of the solved system; the ground losses and the pump power are also checked against pandapipes in Table 2.

properties per pipe or carry the difference as uncertainty. The ground loss law assumes a uniform ground temperature and a constant conductance, which buried pipes in a city do not offer exactly, and it omits the coupling between a supply pipe and the return pipe in the same trench. The friction heating is omitted by both toolchains and is small here. The virtual heat meters rest on the hydraulics being known: the valve positions, the pump's lift and the pipes' roughness fix each substation's flow, and a valve moved without the model knowing would be read as a change of duty; on a real network the valve settings would have to be read or the flows metered at a few points, and the pressure readings of the substations would then join the fit. A single steady snapshot is also blind to the transport delay of the pipes, which on a longer network makes the return temperature at the plant answer to the substations' duties of some minutes before. The comparison with pandapipes is one of formulation on one small loop, not of the two toolchains' merits.

Table 2: The eight buried pipes: mass flow, Reynolds number, head drop, heat lost to the ground (this study and pandapipes with the constant fluid) and the temperature drop along the pipe. Every pipe is turbulent (smallest Re 16.3×10^3) and every loss agrees within below 10^{-9} W (8 of 8 within 0.1 W or 10^{-4} relative). A supply pipe loses about the same heat whatever its flow, because the loss is set by the temperature difference to the ground and by UA, not by the flow; the return pipes, colder, lose less.

pipe	m [kg/s]	Re [10^3]	Δh [m]	loss [W]	loss, pandapipes [W]	$T_{\mathrm{in}} - T_{\mathrm{out}}$ [K]
p1 s	3.51759	253.5	8.4190	565.34	565.34	0.0383
p1 r	3.51759	253.5	8.4190	467.83	467.83	0.0317
p2 s	1.55793	112.3	1.7105	564.85	564.85	0.0864
p2 r	1.55793	112.3	1.7105	417.97	417.97	0.0639
p3 s	0.63945	46.1	0.3092	563.68	563.68	0.2100
p3 r	0.63945	46.1	0.3092	354.94	354.94	0.1322
p4 s	0.22573	16.3	0.0445	560.53	560.53	0.5917
p4 r	0.22573	16.3	0.0445	308.23	308.23	0.3253

Table 3: The four substations: duty, primary flow (this study and pandapipes), inlet and outlet temperatures, the head each takes and the power its two branches dissipate. Flows agree to below 10^{-9} relative and the substations' head drops to below 10^{-9} relative. The far substations, with little head available, run at large temperature drops.

substation	Q_k [kW]	m_k [kg/s]	m_k, pp	T_{in} [°C]	T_{out}	T_{out}, pp	Δh [m]	exch. [W]	valve [W]
k1	60	1.95966	1.95966	79.962	72.667	72.667	4.1504	15.95	63.81
k2	50	0.91848	0.91848	79.875	66.905	66.905	0.7294	1.64	4.93
k3	40	0.41372	0.41372	79.665	56.629	56.629	0.1110	0.15	0.30
k4	30	0.22573	0.22573	79.074	47.408	47.408	0.0220	0.02	0.02

6 Conclusion

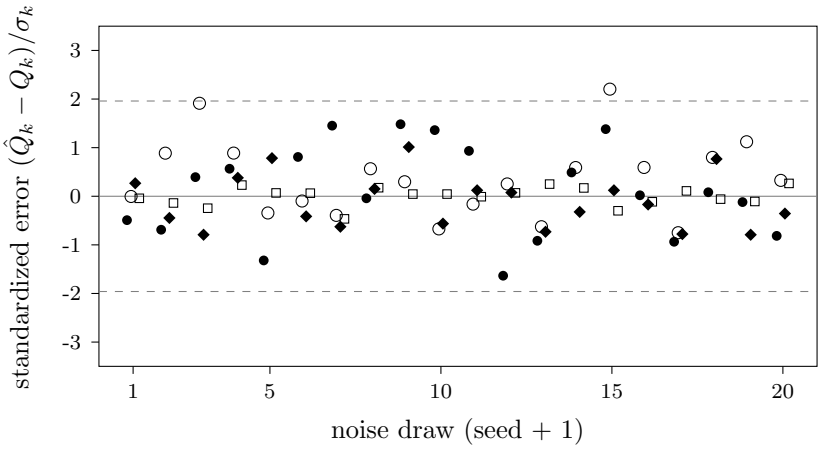
This report asked three questions of one set of conservation, friction, advection and loss residuals on a 4-consumer district heating loop with a return line and a constant-lift pump. What its steady state is and whether the equations reproduce an independent solver: the plant supplies 183.8 kW at 3.52 kg/s and the water returns at 67.6 °C, and pandapipes, given the same constant fluid, agrees to below 10^{-9} K and below 10^{-9} m, while its temperature-dependent properties move the return line by up to 0.93 K, a difference of input the report sizes rather than removes. Where the heat and the pumping power go: 97.9 % of the heat reaches the consumers and 2.1 % is lost to the ground; 88 % of the pump's 724 W is spent on pipe friction and 10 % in the balancing valves, most of it at the nearest substation; both ledgers close to round-off and their totals agree with pandapipes to below 10^{-9} relative, with a friction heating of 0.049 K omitted by both. Whether a consumer's duty can be read without a heat meter: from one 100 mK return-temperature sensor per substation and the plant's flow and return readings, every duty is recovered to 1.0–1.4 % with intervals that cover the truth in 79 of 80 draws, an honest posterior, conservative at the far consumers where the physics weight rather than the reading sets its width, provided the valve settings that fix the flows are known. None of the constructions is new; the study fixes the form in which the heating studies that follow report and leaves four exercise-sized problems in the corpus. The same equations are next applied to a measured network, with the pipes' resistances and the valves' coefficients as unknowns to be inferred from excitation data.

Table 4: The two ledgers. Above the rule, heat in kW: the plant's duty against the consumers' duties and the ground losses, and the same duty read from the plant's flow and temperatures. Below, mechanical power in W: the pump's hydraulic power against the three places it is dissipated. Both ledgers close to round-off; the pandapipes column is the independent check of the totals.

quantity	this study	pandapipes	rel. difference
plant duty Q_{plant}	183.803	183.803	below 10^{-9}
consumers $\sum_k Q_k$	180.000	180.000	—
ground, supply pipes	2.254		
ground, return pipes	1.549		
ground, total	3.803	3.803	below 10^{-9}
$m_{\text{pump}} c_p (T_s - T_r)$	183.803		
pump hydraulic power	724.01 W	724.01 W	below 10^{-9}
pipe friction	637.18 W		
exchanger branches	17.77 W		
balancing valves	69.06 W		

Table 5: Virtual heat meters over 20 noise draws: the true duty, the primary flow, the mean MAP estimate, its root-mean-square error across draws, the mean posterior standard deviation, the same relative to the duty, and how many 95 % intervals cover the truth. An honest posterior has RMSE at or below the posterior standard deviation and about 19 of 20 intervals covering.

consumer	true Q_k [kW]	m_k [kg/s]	mean $\hat{Q}_k$ [kW]	RMSE [W]	mean σ_k [W]	σ_k/Q_k [%]	cover
k1	60.0	1.960	60.08	746	790	1.32	20/20
k2	50.0	0.918	50.18	423	489	0.98	19/20
k3	40.0	0.414	39.95	221	395	0.99	20/20
k4	30.0	0.226	30.00	80	432	1.44	20/20



filled circle: k1; open circle: k2; diamond: k3; square: k4; dashed lines: ±1.96, the 95 % interval

Figure 6: Standardized error of the inferred duty of every consumer in every noise draw (80 values). A calibrated posterior puts about 95 % of them between the dashed lines; here 79 of 80 are, and the largest magnitude is 2.20.

A Reproduction

```
cd docs/terra_graph_engine/case_studies/district_heating/line_network_four_consumers
PYTHONPATH=../../../../../src python run.py
PYTHONPATH=../../../../../src python make_report.py --pdf
```

`run.py` writes `results.json` (every number at six significant digits) and `timings.json`; `make_report.py` reads only `results.json`. Seeds 0–19 drive the reading noise. The pandapipes solutions are read from the committed snapshots under `references/pandapipes/snapshots/`, regenerated by `snapshot.py` (the `oracles` extra).

B Data and problems

The loop is the study's own construction (no external data); its constants are those of the pandapipes snapshot, so geometry lives once. Under `problems/`, four directories: *buried pipe loss* (the outlet temperature and loss of one buried pipe, closed form against the solved closure, at two flows); *substation flow split* (how the head available at a header divides between an exchanger branch and a valve, and between two substations in parallel, closed form against Newton); *two ledgers* (the thermal and mechanical ledgers of the solved loop and the share each term takes); *duty from return temperature* (one consumer's duty and interval from one temperature reading, and how the interval scales with the flow). Each is a statement, a script that solves it with the plugin, and its results file.

References

[Werner 2017] Werner, S. (2017). International review of district heating and cooling. Energy 137, 617–631. doi:10.1016/j.energy.2017.04.045.

[Lund et al. 2014] Lund, H., Werner, S., Wiltshire, R., Svendsen, S., Thorsen, J. E., Hvelplund, F., Mathiesen, B. V. (2014). 4th Generation District Heating (4GDH): Integrating smart thermal grids into future sustainable energy systems. Energy 68, 1–11. doi:10.1016/j.energy.2014.02.089.

[Lohmeier et al. 2020] Lohmeier, D., Cronbach, D., Drauz, S. R., Braun, M., Kneiske, T. M. (2020). Pandapipes: An open-source piping grid calculation package for multi-energy grid simulations. Sustainability 12(23), 9899. doi:10.3390/su12239899.

[Han et al. 2018] Han, Q., Eguchi, R. T., Mehrotra, S., Venkatasubramanian, N. (2018). Enabling state estimation for fault identification in water distribution systems under large disasters. In: 2018 IEEE 37th Symposium on Reliable Distributed Systems (SRDS), pp. 161–170. doi:10.1109/SRDS.2018.00027.

[Dellaert and Kaess 2017] Dellaert, F., Kaess, M. (2017). Factor graphs for robot perception. Foundations and Trends in Robotics 6(1–2), 1–139. doi:10.1561/2300000043.