EUROSUN2026

MODULAR 1D STRATIFIED THERMAL ENERGY STORAGE MODEL WITH MULTI-PORT INTERFACE FOR DISTRICT HEATING CO-SIMULATION: DESIGN AND VALIDATION

Jonas Pfeiffer, Matthias Kunick
Zittau/Görlitz University of Applied Sciences, Zittau (Germany)
jonas.pfeiffer@hszg.de

Summary

Pit thermal energy storages allow solar district heating systems to cover winter demand with heat collected in summer. Planning such a system means simulating the storage together with the pipe network, and that calls for a storage model which predicts the temperature delivered to the network accurately, which is inexpensive enough to be evaluated in every network time step, and which can represent how real storages are connected, namely not through one pair of pipes but through several connections at different heights.

This paper presents a one-dimensional finite-volume model in which each connection, here called a port, is described by no more than a height, a mass flow and an inlet temperature, so that any number of ports can be placed at any height. On a storage with three connections a conventional two-port interface proves to be a strong baseline, but it cannot represent the middle one, at which the multi-port formulation lowers the mean absolute error (MAE) from 6.6 to 4.8 K; without charge/discharge switching logic the same interface is wrong by 13.6 K. Against a detailed reference model the agreement is 0.5–1.4 K at a computing cost three orders of magnitude lower, and against two Danish pit storages, with parameters taken from the plant documentation (Sifnaios et al., 2023, 2025) rather than fitted, the MAE is 2.1 K over a full year at Høje Taastrup and 3.4 K over eight months of the first and hardest operating year at Dronninglund. Coupled to the open-source solver pandapipes, it adds at most 1.5 % to the cost of solving the network.

Keywords: pit thermal energy storage, district heating, finite volume method, multi-port interface, co-simulation

Field of application: Solar district heating, seasonal thermal energy storage, building and district energy system simulation

Addressed audience: Simulation engineers, district heating planners, solar thermal researchers

1. Introduction

Large-scale seasonal thermal energy storage (TES) is a central component in solar district heating systems, decoupling summer solar heat collection from winter heat demand. Pit thermal energy storage (PTES) systems in particular, ground-embedded and lined water basins closed by a floating insulated lid, are widely deployed across northern Europe (Xiang et al., 2022); volumes between 60 000 and 200 000 m^3 are in operation and larger basins are under construction. Accurate simulation of thermal stratification is essential for predicting charging efficiency, supply temperatures, and seasonal energy yields (Dahash et al., 2020; Ochs et al., 2021).

The physics of a thermal energy storage are commonly described in a one-dimensional finite-volume (FV) model, and two spatial discretisations are in use. A *Lagrangian* discretisation follows individual volumes of water as they move through the storage and so reproduces sharp thermoclines with little numerical diffusion (Narusavicius et al., 2024), but its cost grows as the segment count rises during operation, which makes it expensive for co-simulation, where the storage is evaluated repeatedly inside every network time step. An *Eulerian* discretisation solves energy balances on a fixed grid at $O(K)$ cost that does not grow as the storage is cycled (Lago et al., 2019; Untrau et al., 2023; Rendall et al., 2021), but it blurs the thermocline if the control volumes are too large. Detailed stratified-storage models have long supported connections at freely chosen heights (Sect. 2.1), although as self-contained components rather than as lightweight elements inside a network solver, whereas reduced-order FV models built for system-level integration expose a fixed charge/discharge pair.

This paper presents a finite-volume formulation that is *independent of the number of ports*, namely a cumulative-flux advection scheme (Sect. 2.2) that handles any port configuration without code changes. It is packaged

as a modular, open-source model for co-simulation, validated against two PTES in operation with published geometry and heat transfer data, and coupled to the open-source pipe-flow solver pandapipes in Sect. 4.4.

2. Model Description

2.1 Requirements on the hydraulic connections

In a district heating context the storage often acts as a hydraulic separator (Fig. 1), in which generator circuit and network circuit attach at different heights and are otherwise hydraulically independent. The network results of Sect. 4.4 use this arrangement, and the connection model has to support it.

An operating storage shows what this requires. The Dronninglund PTES is charged and discharged through three diffusors at 0.4, 10.9 and 15.3 m above the floor, and the middle one changes direction over the year, acting as an inlet in summer and as an outlet in winter (Sifnaios et al., 2023). A model that offers only one fixed inlet and one fixed outlet cannot represent this storage without case-specific logic that lumps the flows together and switches the heights according to the operating mode.

Describing a storage through connections at freely chosen heights is not new: TRNSYS Type 340 has offered up to ten double ports at any height for decades (Drück, 2006), Modelica storage models let the user place the charging and discharging connectors (De la Cruz-Loredo et al., 2023), and a recent reduced-order PTES model allows the number and vertical position of ports to be configured freely (Bayer et al., 2025). The contribution here is not the idea of multiple ports but the formulation behind it (Sect. 2.2), which handles any port configuration through a single cumulative mass flow instead of a switching logic and is inexpensive enough to run inside every time step of a network solver.

The model separates the description of the storage from its operation. Before the simulation, the storage is defined by a list of ports, each fixed at a height z_p above the storage floor and assigned a diffusor sub-model. During the simulation, only a mass flow $\dot{m}_p$ (positive into the storage, negative out of it) and, for an inflow, an inlet temperature $T_{p,\mathrm{in}}$ are prescribed for each port; the model returns the outlet temperature at every port. Port heights never change; only the mass flows do. A port that carries no flow needs no special treatment, so the same formulation covers any combination of simultaneous charging, discharging and recirculation.

One limit follows from the discretisation. Two ports closer together than H/K act on the same control volume and become indistinguishable, so the grid has to be fine enough for the port spacing of the storage, which is not a restriction for the cases shown here.

2.2 Model assumptions, governing equations and numerics

The starting point is the one-dimensional form of Euler's equations for the conservation of mass, momentum and energy. The water is considered as incompressible and the pressure does not change across the storage, so the momentum equation is not solved and the mass balance and the thermal energy balance remain. The storage is modelled as a one-dimensional column of K control volumes, each of which is fully mixed and has a fixed volume V_k set by the geometry. Density and specific heat capacity follow temperature-dependent polynomial fits for water.

The discretisation itself is the standard 1D FV one (Lago et al., 2019; Untrau et al., 2023); what is added around it is that the physical sub-models are independently exchangeable (Table 1). The geometry, for instance, enters only through the cross-sectional area $A(z)$, so that one solver core covers cylindrical hot-water tanks and truncated-cone or truncated-pyramid pit storages without code modifications.

Table 1: Sub-models that can be exchanged without touching the solver core, and the variants used in this paper.

Sub-model	Variants available	Used here
Geometry	cylinder, truncated cone, truncated pyramid	cylinder; truncated pyramid
Heat loss	uniform, split lid/wall, ground temperature	split lid/wall
Fluid properties	constant, polynomial fits for water	polynomial fits
Diffusor	point, uniform zone	uniform zone, 1 m

Fig. 1 fixes the notation. The storage is divided into K control volumes of height $\Delta z = H/K$, numbered from $k = 0$ at the top to $k = K-1$ at the bottom; the interfaces between them carry the half-indices $k \pm 1/2$. Each

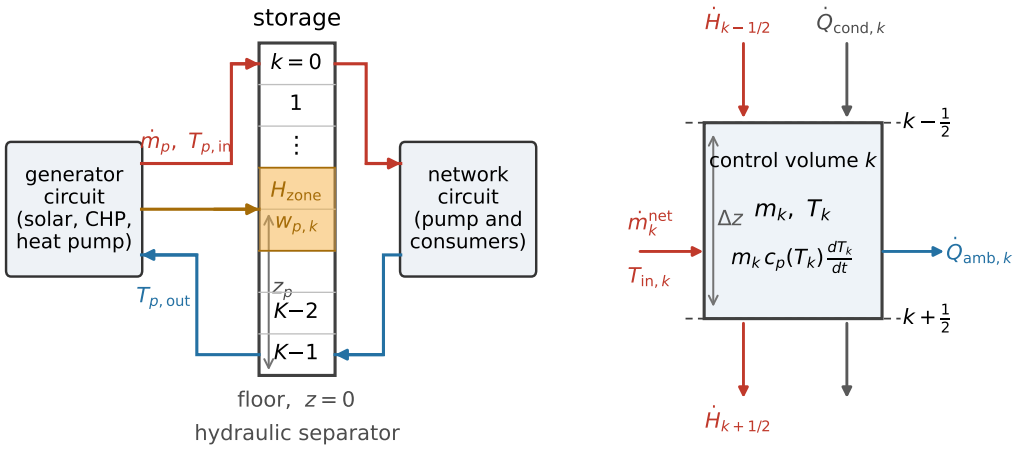


Fig. 1: The storage as a hydraulic separator, and its discretisation. Left: generator circuit and network circuit attach at their own heights and are otherwise hydraulically independent, exchanging heat only through the water in the storage. Each connection is a port fixed at a height z_p above the floor, and a diffusor of height H_{zone} distributes its flow over several control volumes with the weights $w_{p,k}$. Right: one control volume with the terms of its energy balance. Enthalpy flows $\dot{H}$ cross the interfaces at $k \pm 1/2$, the port contributes the net mass flow $\dot{m}_k^{\text{net}}$ at the inlet temperature $T_{\text{in},k}$, and conduction and the exchange with the surroundings act in addition.

control volume satisfies the energy balance (Dahash et al., 2020; Untrau et al., 2023)

$$m_k \, c_p(T_k) \, \frac{dT_k}{dt} = \dot{H}_{\text{adv},k} + \dot{Q}_{\text{cond},k} + \dot{Q}_{\text{amb},k} \tag{1}$$

with the control-volume mass $m_k = \rho(T_k)\, V_k$, evaluated at the current temperature. The volume V_k is fixed by the geometry, and the mass change caused by thermal expansion is neglected, so the internal mass balance below closes. The advective term is written as an enthalpy flow $\dot{H}$ rather than a heat flow, because it is energy carried by mass across the control-volume boundary.

Each control volume k carries a net port mass flow, positive inward,

$$\dot{m}_k^{\text{net}} = \sum_p w_{p,k}\, \dot{m}_p \qquad [\text{kg s}^{-1}], \tag{2}$$

which is a source term in the balance and not a change of the control-volume mass. The weights $w_{p,k}$ follow the diffusor sub-model. A point diffusor assigns $w_{p,k} = 1$ to the control volume nearest z_p. A uniform diffusor of height H_{zone} spreads the flow equally over the set $\mathcal{Z}_p$ of control volumes whose centres lie inside the zone,

$$w_{p,k} = \begin{cases} 1/|\mathcal{Z}_p| & k \in \mathcal{Z}_p \\ 0 & \text{otherwise} \end{cases} \qquad \mathcal{Z}_p = \left\{ k : |z_k - z_p| \le H_{\text{zone}}/2 \right\}, \tag{3}$$

so that $\sum_k w_{p,k} = 1$ in either case. The mass balance of a control volume of constant volume V_k with an incompressible fluid simplifies to $\dot{m}_{k-1/2} + \dot{m}_k^{\text{net}} = \dot{m}_{k+1/2}$, with the interface flows counted positive downwards. The mass flow crossing the interface between control volumes k and $k{+}1$ is therefore the cumulative sum

$$\dot{m}_{k+1/2} = \sum_{i=0}^{k} \dot{m}_i^{\text{net}} \qquad [\text{kg s}^{-1}], \qquad \dot{m}_{-1/2} = \dot{m}_{K-1/2} = 0, \tag{4}$$

since there is no mass flow through the lid, and none through the floor because the port flows sum to zero. The internal mass balance is a consequence, not an imposed condition. The enthalpy flow across each interface uses first-order upwinding with a van Leer TVD anti-diffusion correction Φ scaled by the local Courant number,

$$\dot{H}_{k+1/2} = \dot{m}_{k+1/2}\, c_p(T)\, T^{\text{up}}_{k+1/2} + \left(1 - \text{CFL}_{k+1/2}\right) \Phi_{k+1/2}, \tag{5}$$

so that the advective term of control volume k becomes $\dot{H}_{\text{adv},k} = \dot{H}_{k-1/2} - \dot{H}_{k+1/2} + c_p(T) \left[\max(\dot{m}_k^{\text{net}}, 0)\, T_{\text{in},k} + \min(\dot{m}_k^{\text{net}}, 0)\, T_k \right]$, with $T_{\text{in},k}$ the mass-flow-weighted inlet temperature of the ports feeding control volume k, and

where T^{up} is the upwind temperature and the anti-diffusive flow is $\Phi_{k+1/2} = \frac{1}{2}\dot{m}_{k+1/2}\, c_p(T)\, \psi(r)\, (T_{k+1} - T_k)$ with the van Leer limiter $\psi(r) = (r + |r|)/(1 + |r|)$ and r the ratio of the upwind temperature gradient to the local one. The Courant number is clipped at unity, so the correction vanishes once a face is swept in a single step. Whether that happens depends on grid and flow rate rather than on the time step alone: for the benchmark storage of Sect. 3 with $K = 50$ the Courant number stays at 0.30 even at $\Delta t = 3600\,\mathrm{s}$, so the limiter is active throughout, whereas the same case at $K = 200$ reaches 1.2 and is clipped. For a single charge/discharge pair the formulation reduces exactly to the conventional two-port scheme.

Where two streams meet, the formulation transports the merged flow and conserves its energy but does not resolve the small-scale entrainment and turbulent mixing. That mixing enters only implicitly, through the finite control-volume size, the diffusor mixing zone (Table 1) and f_λ, and is therefore lumped rather than predicted.

Axial conduction uses central differences with an effective conductivity $\lambda_{\mathrm{eff}} = f_\lambda\, \lambda_{\mathrm{fluid}}$ ($f_\lambda \geq 1$; Falkner (2014) reports $f_\lambda \approx 2$–5 for small storages and $f_\lambda \approx 1$ for large storages without internals, the value used here at PTES scale) that absorbs unresolved destratification, and the heat exchange with the surroundings is $\dot{Q}_{\mathrm{amb},k} = -U_k A_k (T_k - T_{\mathrm{amb},k})$ with separate coefficients, ambient temperatures and surface areas for the lid and for the rest of the envelope; the same term carries heat gains whenever the surroundings are warmer than the storage. Time integration uses implicit Euler with a tridiagonal (TDMA) solve, which is unconditionally stable and therefore suited to the coupling time steps of 900–3600 s used here. Inverted stratification is removed after each step by mass-weighted (caloric) mixing of unstable control-volume pairs (Lago et al., 2019; Xiang et al., 2022). This is a numerical correction, not a physical model: the buoyant plume that forms when cold water is injected above warmer water is not resolved, so results for strongly inverted charging should be read with that in mind.

3. Numerical Benchmark

3.1 Setup

The benchmark uses the reference tank configuration distributed with FreeTTES ($V = 50\,140\,\mathrm{m}^3$, $R = 20\,\mathrm{m}$, $H = 39.9\,\mathrm{m}$, $U = 0.2\,\mathrm{W\,m^{-2}K^{-1}}$, diffusor zones 1 m high at 1.0 and 39.65 m). Geometry, diffusor arrangement and loss coefficient are taken from that configuration, so the test case does not favour the model presented here. Three scenarios are imposed:

1. **Weekly cycle:** 168 h, daily 8 h charge / 4 h standby / 8 h discharge / 4 h standby, at 300 m^3/h, charging at 90 °C and discharging with a 60 °C return.
2. **Fast switching:** 48 h, 6 h cycle of 2 h / 1 h / 2 h / 1 h, same flow and the same temperatures.
3. **Annual seasonal:** 8 760 h at the same flow and temperatures; charging May–September, discharging November–March.

The agreement between two models is reported as the mean absolute deviation of the outlet temperature, averaged over those time steps in which flow actually crosses a port. Both values are model results, so this is a deviation and not an error; scored against a measurement below, the same quantity is called the mean absolute error (MAE). It is zero as long as the two models deliver the same outlet temperature, which during charging holds until the thermocline reaches the outlet.

FreeTTES is used as the reference for three reasons: it is a Lagrangian-type (segment) model and so resolves the thermocline by design rather than by grid refinement, which is precisely what the Eulerian scheme has to reproduce; it is openly available and validated against distributed-temperature-sensing measurements in operating atmospheric tanks (Narusavicius et al., 2024); and it provides the tank configuration used here. The deviation measures model-to-model agreement, not absolute accuracy: FreeTTES was developed for atmospheric tanks and its inversion sub-model uses manually chosen parameters whose general validity its authors state is not yet established. A fully mixed 0D model is carried through for information only; without stratification it cannot predict an outlet temperature at all. The absolute per-step cost is anchored by the segment reference, against which the present model at $K = 50$ is 1 800 to 2 800 times faster across the three scenarios (Table 2).

3.2 Results

For the weekly cycle (Table 2), $K = 50$ achieves a deviation of 0.47 K against the segment reference at 0.19 ms per time step, well within the budget of the 900–3600 s coupling steps used in Sect. 4.4. For rapid load switching, increasing from $K = 50$ to $K = 200$ reduces the deviation from 1.44 K to 0.79 K at twice the computation time, consistent with the thermocline-sharpening effect of finer spatial resolution reported by Dahash et al. (2020) and Ochs et al. (2021). Over the annual cycle the two resolutions differ by only 0.04 K and are indistinguishable at this timescale, especially since the reference is itself a model. The fully mixed 0D storage deviates by 16.3 K in the weekly cycle and 18.6 K under fast switching, but only by 1.94 K over the year, because the annual signal is dominated by the slow charge and discharge envelope rather than by the thermocline.

The stored useful energy at year end agrees within 3.1 MWh (0.11 %) across the 1D configurations, and lies 15 MWh (0.54 %) above the segment reference. An independent closure audit accumulates the port enthalpy flows and the losses to the surroundings separately and compares their sum with the change in stored energy. Run on the two measurement cases of Sect. 4, the residual is $1.7 \cdot 10^{-4}$ (Dronninglund) and $8.1 \cdot 10^{-5}$ (Høje Taastrup) of the gross enthalpy passed through the ports. Measured against the much smaller net change in stored energy, the same residuals read 0.8 and 3.2 %. Throughput is the relevant measure for an advection scheme, and on that measure the enthalpy-flow form $\dot{m}\, c_p(T)\, T$ conserves energy to better than 0.1 %.

Table 2: Benchmark results: mean absolute deviation of the outlet temperature and mean wall-clock time per time step over the scenario ($\Delta t = 3600$ s; 48, 168 and 8 760 steps, so sub-millisecond values carry ≈ 20 % run-to-run scatter). Single-threaded Python, Intel Core i7-13620H; reference cost is the annual value.

Scenario	Model configuration	Deviation [K]	Time [ms/step]
Weekly	K=50	0.47	0.19
	K=200	0.43	0.47
Fast switching	K=50	1.44	0.25
	K=200	0.79	0.50
Annual	K=50	1.20	0.16
	K=200	1.16	0.41
Reference	Segment model (FreeTTES)	(ref.)	≈447

3.3 Time-step convergence

The results above use one time step. To show that they are a property of the model and not of that choice, the weekly scenario was repeated with the time step halved successively from 3600 s down to 56.25 s, and each run compared with the finest one on a common hourly grid (Table 3). The deviation falls by roughly a factor of two per halving, which is the first-order behaviour expected of implicit Euler, and no error floor appears that would indicate a residual inconsistency in the spatial scheme. The Courant number stays well below unity throughout, so the anti-diffusion limiter is active in every one of these runs.

Table 3: Time step convergence of the weekly scenario ($K = 50$, implicit solver, TVD advection). Mean absolute deviation of the outlet temperature against the finest run ($\Delta t = 56.25$ s); the Courant number is the maximum reached during the run.

Δt [s]	max. Courant number	Deviation vs. finest run [K]
3600	0.302	0.145
1800	0.151	0.061
900	0.075	0.030
450	0.038	0.014
225	0.019	0.006
112.5	0.009	0.002

4. Validation

The model is validated on two levels: its physics, against operational measurement data from two pit thermal energy storages; and its suitability for its intended purpose as an inline co-simulation component, by coupling

it to a district heating network solver (Sect. 4.4).

4.1 Dronninglund PTES (Denmark, 2014)

The first measurement case is the quality-controlled operational dataset from the Dronninglund pit thermal energy storage (Sifnaios et al., 2023), a pit storage in the shape of a truncated pyramid ($V \approx 60\,030\,\mathrm{m}^3$, bottom $26.4 \times 26.4\,\mathrm{m}$, top $90.4 \times 90.4\,\mathrm{m}$, $H = 16\,\mathrm{m}$) operated as part of a solar district heating system in Denmark. The truncated-pyramid geometry sub-model of Table 1 is used without modification. Three diffusor ports at their measured heights ($z_{\mathrm{top}} = 15.3\,\mathrm{m}$, $z_{\mathrm{mid}} = 10.9\,\mathrm{m}$, $z_{\mathrm{bot}} = 0.4\,\mathrm{m}$) receive hourly mass-flow boundary conditions and inlet temperatures derived directly from the dataset. They use the uniform-diffusor sub-model with a 1 m mixing zone, because a point diffusor would leave the water below the lowest port stagnant. Of the two temperature references the dataset offers, the *sensor strings inside the storage* (two strings, 1 m spacing, mutually offset by 0.5 m) are used throughout and not the *diffusor pipe sensors*, which sit in the hydraulic lines and relax towards ambient whenever their diffusor is inactive. The simulation starts from the measured profile on 1 May 2014, the first day for which flow-direction data are available, and runs through 31 December 2014 (5 872 hourly time steps). The heat transfer coefficients are not fitted: $U_{\mathrm{lid}} = 0.167\,\mathrm{W\,m^{-2}K^{-1}}$ comes from the in-situ lid heat-flux measurement supplied with the dataset, and $U_{\mathrm{wall}} = 0.35\,\mathrm{W\,m^{-2}K^{-1}}$ from soil conductivity $\lambda \approx 0.4\,\mathrm{W\,m^{-1}K^{-1}}$ and an effective depth of $\approx 1\,\mathrm{m}$. No ambient record exists for 2014, and the dataset's own sensor is documented as unreliable (Sifnaios et al., 2023). The lid boundary condition is therefore a sinusoidal air temperature (mean 8 °C, amplitude 8.5 K, minimum mid-January), and that for the walls and floor a constant 8 °C, the Danish annual mean, which lies 2–4 K below the ground temperatures measured beside the storage, so wall losses are biased high.

2014 is the plant's first operating year, the only one publicly released, and the hardest year of the record: the measured storage efficiency was ≈ 78 % against 90 and 91 % in 2015 and 2016, attributed to the start-up phase preheating the surrounding ground (Schmidt and Sørensen, 2018). A steady $UA\Delta T$ loss model cannot represent that transient, and the Modelica study of Reisenbichler-S. et al. (2025) discussed below excludes 2014 for this reason, using it only to preheat its ground domain before scoring 2015–2019. This site is therefore treated here as the stress case, and the quantitative cross-validation rests on Høje Taastrup (Sect. 4.3).

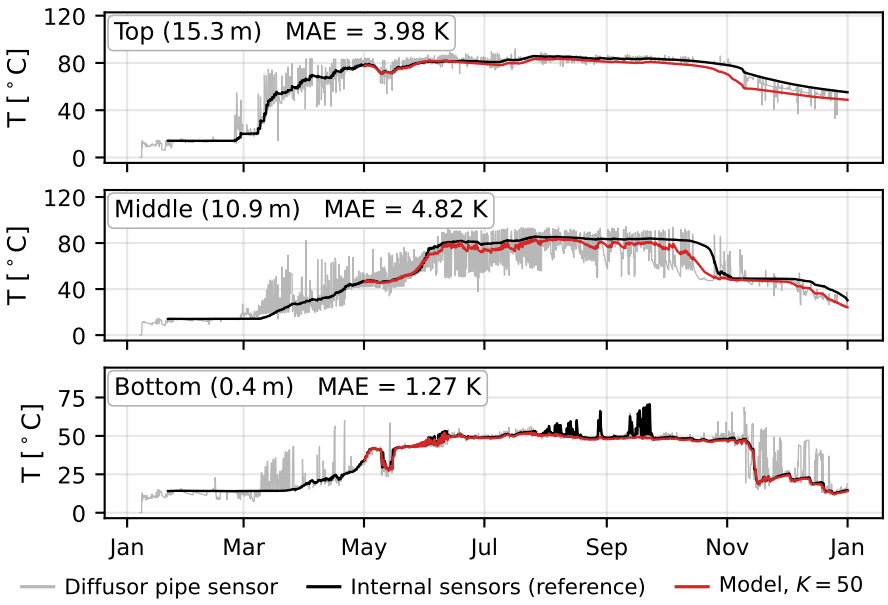


Fig. 2: Dronninglund PTES 2014: simulated (red) vs. measured (black) temperatures at the three diffusor positions (top 15.3 m, mid 10.9 m, bottom 0.4 m). Reference: the internal sensor strings. The diffusor pipe sensors (grey) relax towards ambient whenever their diffusor is inactive (dense downward spikes) and are shown only to illustrate why they are unsuitable. Initialised from the measured profile on 1 May; MAE per position is annotated in each panel.

The dominant error occurs at the middle diffusor (10.9 m, MAE 4.82 K; Table 4, Fig. 2), where the model positions the thermocline approximately 2–3 m too deep in October–November; the upper and lower diffusors give 3.98 and 1.27 K, and the mean over the three diffusor heights is 3.36 K. The strings also allow scoring over the *complete* measured profile. Across all 32 sensor heights (187 904 sensor-hours) the MAE is 5.40 K, higher

than the three-point figure because the error depends strongly on height: 1.21 K at the floor (0.5 m) against 9.27 K at 8.5 m, in the thermocline zone that the three diffusors straddle rather than sample (Fig. 4).

Scoring the energy content rather than individual heights separates two kinds of error. The energetic mean temperature, the temperature a fully mixed storage of the same energy content would have, is reproduced with a mean absolute error of 5.20 K, almost all of it a constant bias of −5.17 K: the simulated storage is too cold over essentially the whole period. That is an energy-balance and not a stratification error, but it cannot be attributed to the boundary parameters alone. Lowering U_{wall} from 0.35 to 0.20 $\mathrm{W\,m^{-2}K^{-1}}$ removes 0.8 K of it, and switching the ambient exchange off altogether (which removes heat gains as well as losses, since the same term carries both) removes 2.6 K, so at least half of the offset survives any parameterisation of it. Nor does the measured record close over this period. The plant energy meters and the measured temperature profiles differ by roughly 320 MWh, which corresponds to 4.6 K of storage temperature and is thus of the same order as the bias itself. The residual is therefore reported here rather than explained, and the absolute energy balance is assessed at Høje Taastrup instead.

For scale, a detailed Modelica PTES model validated on this same site under a full ASME V&V uncertainty analysis reports ±2 K over five years (Reisenbichler-S. et al., 2025), although that is a validation uncertainty interval rather than an MAE. It uses a transiently discretised ground domain, where the present model uses a constant $UA\Delta T$ loss. The 8-month simulation requires 9.7 s wall-clock (1.65 ms per time step, $K = 50$): an order of magnitude more than the cylindrical benchmark of Table 2, because this configuration combines a non-prismatic cross-section, temperature-dependent fluid properties and three active ports.

Table 4: Dronninglund PTES validation against the internal sensor strings (May–December 2014, $K = 50$, $f_\lambda = 1$). Nominal sensor accuracy ±0.15 K (PT100, Sifnaios et al., 2023); the two independent strings differ from each other by 0.43 K on average, the effective measurement floor.

Diffusor	Height above floor	MAE [K]
Top	15.3 m	3.98
Middle	10.9 m	4.82
Bottom	0.4 m	1.27
Overall (3 diffusor heights)	—	**3.36**
Profile-wide (32 heights)	0.5–16.0 m	**5.40**

4.2 Effect of the multi-port interface

Because this storage has three diffusors at measured heights, the multi-port capability can be tested against measurement rather than asserted. Dronninglund carries this comparison even though its record cannot settle an absolute temperature, because it is the only one of the two sites at which the test is possible: its middle diffusor reverses over the year (Sect. 2.1), which is precisely what a two-port interface cannot represent, whereas at Høje Taastrup that diffusor is an outlet throughout (Sect. 4.3). The energy-balance offset described above acts on all three variants alike and cancels in their difference. One and the same model, with identical physics and identical measured boundary conditions, was coupled through three different hydraulic interfaces and scored against the same internal sensor strings (Table 5): the three diffusor ports at their measured heights; a conventional fixed charge/discharge pair, in which all simultaneous inflows are lumped into one inlet at their mass-weighted mixing temperature and the inlet/outlet heights follow the flow direction; and the same pair with the lumped inflow always injected at the top.

Implemented correctly, a two-port interface is a strong baseline, since it gives the better value in three of the five columns. It is marginally more accurate at the top and bottom diffusors and 0.89 K better over the whole profile, because lumping the inflow conserves the injected energy and merely places it at the wrong height. The multi-port configuration is better exactly where the geometry demands it. At the middle diffusor, the connection a two-port interface structurally cannot represent, the error falls from 6.59 to 4.82 K (−27 %). The two interfaces also differ at the network-relevant top port by up to 7.11 K at individual time steps, so even at the one diffusor where the two-port interface has the lower MAE the choice is not neutral for supply-temperature prediction. Without the direction switching, with the lumped inflow injected at the top, the error grows to 13.55 K.

The formulation therefore does not make the model more accurate everywhere; what it does is reproduce a storage as it is built, for any port count and flow pattern, without case-specific lumping and switching logic.

Table 5: Hydraulic-interface comparison on Dronninglund: identical measured boundary conditions, identical physics, three hydraulic interfaces. MAE against the internal sensor strings (May–December 2014, $K = 50$); bold marks the best value in each column.

	At the diffusor heights [K]				
Hydraulic interface	Top	Middle	Bottom	Mean [K]	Profile-wide [K]
Three diffusor ports (multi-port)	3.98	**4.82**	1.27	**3.36**	5.40
Two-port, direction-switching	**3.61**	6.59	**1.07**	3.76	**4.51**
Two-port, lumped inflow at top	13.31	10.13	17.22	13.55	11.98

4.3 Høje Taastrup PTES (Denmark, 2024)

That cross-validation uses the monitoring dataset of the Høje Taastrup PTES (Sifnaios et al., 2025): a full year of normal operation with a measured ambient record and a documented lid construction. It is a larger pit storage ($V \approx 70\,880\,\mathrm{m}^3$, bottom 29.4×29.4 m, top 105.4×105.4 m, $H = 14.05$ m) in operation since February 2023 and the first large PTES to be operated as a *short-term* storage, on cycles of one to two weeks rather than seasonally. The full calendar year 2024 is simulated (52 704 time steps at $\Delta t = 600$ s).

The real storage has a slightly slanted floor (14.05 m deep at the east end, 12 m at the west); it is modelled as a truncated pyramid with a horizontal floor at the greater depth. Three diffusor ports are placed at the heights given in Sifnaios et al. (2025): $z_{\mathrm{top}} = 13.80$ m, $z_{\mathrm{mid}} = 10.50$ m (volumetric midpoint), $z_{\mathrm{bot}} = 0.45$ m; here the middle diffusor is an outlet throughout, so unlike at Dronninglund it never carries an inflow. Mass-flow boundary conditions and inlet temperatures are taken directly from the dataset. The sensor strings inside the storage serve as the measurement reference, the diffusor pipe sensors being unsuitable for the same reason as at Dronninglund. The lid heat transfer coefficient is $U_{\mathrm{lid}} = 0.151\,\mathrm{W\,m^{-2}K^{-1}}$ (calculated from the XPS + NOMATEC lid construction, Sifnaios et al., 2025); the wall and floor coefficient is estimated at $U_{\mathrm{wall}} = 0.30\,\mathrm{W\,m^{-2}K^{-1}}$ (soil conductivity $\lambda = 2.23\,\mathrm{W\,m^{-1}K^{-1}}$, Table 2 in Sifnaios et al., 2025; effective depth ≈ 7 m). Ambient temperature is updated each time step from the measured outdoor sensor.

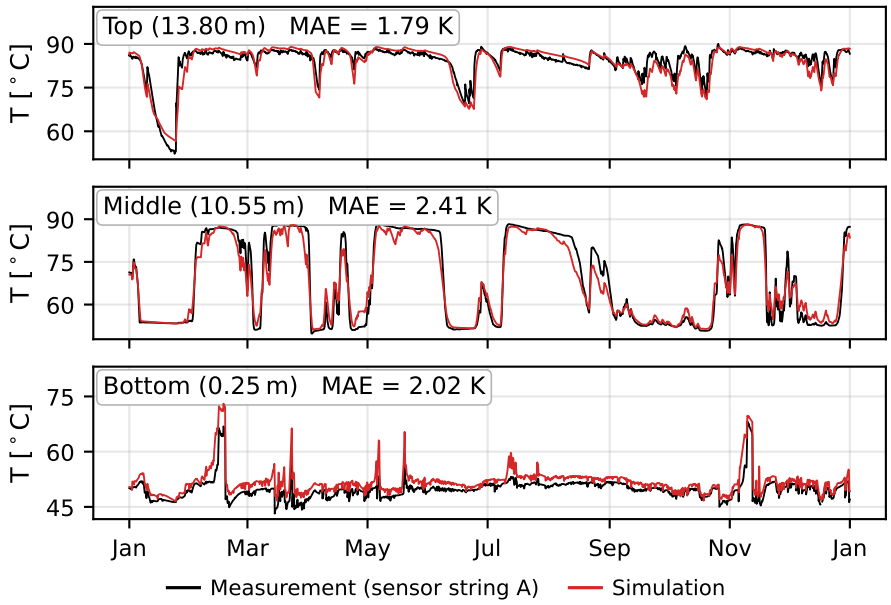


Fig. 3: Høje Taastrup PTES 2024: simulated (red) vs. measured (black) temperatures at the three sensor heights nearest the diffusors (13.80 m, 10.55 m, 0.25 m above the floor). Ambient temperature used as time-varying boundary condition; MAE per position is annotated in each panel.

The overall MAE (Table 6, Fig. 3) is 2.07 K at the three diffusor heights and 2.05 K over the complete measured profile (26 heights). The middle position again shows the largest error (2.41 K), here from thermocline position uncertainty in the transition periods alone.

At the diffusor heights the two sites differ by 1.3 K (3.36 vs. 2.07 K), less than a surface-to-volume argument would suggest for the larger volume and better lid of Høje Taastrup. The profile-wide metric locates the real

difference. Høje Taastrup stays within 1.46–2.67 K at every measured depth, whereas Dronninglund ranges from 1.21 K at the floor to 9.27 K in the thermocline, so what separates the two sites is thermocline sharpness and not the accuracy of the loss model. The energetic mean temperature supports that reading: it is reproduced with a mean absolute error of 0.58 K and a bias of only +0.28 K (67.7 against 67.4 °C on average), so the stored energy is essentially correct and the profile-wide residual is a question of thermocline position alone, in contrast to the −5.17 K bias at Dronninglund. The full-year run completes in 68.7 s wall-clock (1.30 ms per time step, $K = 50$).

Table 6: Høje Taastrup PTES validation, full year 2024 ($K = 50$, $f_\lambda = 1$): MAE at the sensor nearest each diffusor, and over all 26 measured heights (0.25–13.80 m). Pt100 uncertainty ≈ 0.3 K at operating temperature (Sifnaios et al., 2025).

Position	Height [m]	MAE [K]
Top	13.80	1.79
Middle	10.55	2.41
Bottom	0.25	2.02
3 diffusors		**2.07**
Profile-wide		**2.05**

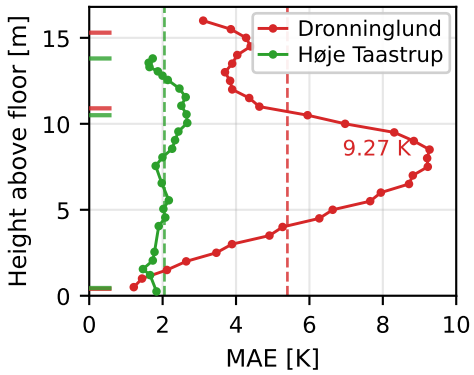


Fig. 4: Error structure over the full measured profile, both sites on a common error axis; short bars mark the diffusor heights, dashed lines the profile-wide means. At Dronninglund the error peaks in the thermocline zone, which the diffusors straddle rather than sample, so the three-point figure understates the profile-wide one (3.36 vs. 5.40 K); at Høje Taastrup the profile is flat and the two agree (2.07 vs. 2.05 K).

4.4 Network integration

The storage attaches to the network solver through the multi-port interface of Sect. 2.1, as an external component exchanging port data once per step in a sequential (Gauss–Seidel) scheme. The coupling target is the open-source pipe-flow solver pandapipes (Lohmeier et al., 2020), chosen because it makes the set-up openly reproducible and isolates the coupling cost. The storage is operated as a hydraulic separator (Fig. 5): on the network side pandapipes' `circ_pump_const_pressure` component supplies a set of heat consumers, its supply temperature set each step to the storage top temperature; on the generator side a simple state-of-charge hysteresis charges the storage through a dedicated port, a proof of concept that exercises the interface rather than a dispatch recommendation. The network is solved quasi-statically at each coupling step (Δt = 900–3600 s). pandapipes' dynamic mode and the designated heat-generator component currently under development (Zoschke, 2025) are the natural next steps.

Three findings result: First, driven by a real demand profile over a two-week shoulder-season window (Fig. 5), the storage-controlled 250 kW generator runs at a duty cycle of 50 %. The stratified model holds the network supply temperature at 75–90 °C throughout, whereas a fully mixed 0D storage of identical capacity swings between 65 and 90 °C (mean deviation 7 K), reproducing at network level the conclusion of Sect. 3. Second, the coupling overhead is small and shrinks with network size. Median per-step costs over 200 timed steps: a storage step ($K = 20$) 0.85 ms, against 56 ms for a 10-consumer pandapipes solve and 141 ms for a 400-consumer one, so the overhead falls from 1.5 to 0.6 % across a 40-fold range of network size; the storage cost follows K and not the network ($K = 50$: 1.45 ms; $K = 100$: 2.57 ms). Medians are quoted because the pandapipes step time is strongly bimodal in the solver iteration count. These coupled timings are several times higher than the standalone ones of Table 2 (0.16–0.25 ms at $K = 50$), because small-array operations dispatch differently once pandapipes is imported; storage and network are measured inside the same process, so the ratio between them is unaffected. Both sides are single-threaded Python; against a compiled network solver the storage share would rise rather than fall, though the storage step stays at or below the millisecond scale either way.

Third, the sequential exchange lets the network see a storage state one step old. Scoring the identical scenario from Δt = 3600 s down to 225 s against the finest run gives a supply-temperature deviation of 0.105 K at 3600 s, 0.022 K at 900 s and 0.007 K at 450 s (maxima 0.70, 0.12 and 0.04 K), which is first-order convergence. The lag therefore costs at most a few tenths of a Kelvin over the recommended range, an order of magnitude below the model's own measurement error.

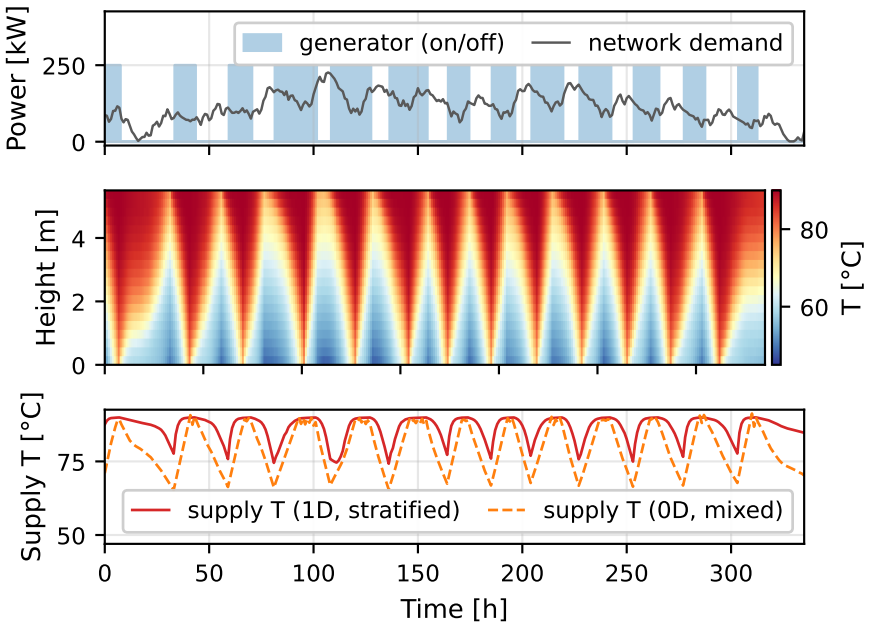


Fig. 5: Realistic-load pandapipes co-simulation ($V = 50\,\mathrm{m}^3$ buffer, $K = 20$, two-week window). Top: the real network demand and the on/off dispatch of the storage-controlled 250 kW generator. Middle: the storage stratification $T(z, t)$ as a carpet plot; the thermocline descends on charging and rises on discharge with each cycle. Bottom: the network supply temperature of the stratified 1D versus a fully mixed 0D storage of equal capacity.

A tighter integration was also implemented: In parallel with this work the same model was integrated *directly* into the district heating simulation environment VICUS Districts (solver core NANDRAD) within the first author's master's thesis (Pfeiffer, 2026), where the storage is not an external component but a flow element, so that its control-volume energies enter the solver's own implicit equation system and are advanced by its adaptive backward-differentiation integrator. The multi-port advection of Sect. 2.1 carries over unchanged. One adaptation is a genuine consequence of tight implicit coupling and transfers to any such integration: the standalone model removes inverted stratification by discontinuous caloric mixing applied *after* each step (Sect. 2.2), and a correction applied outside the right-hand side in this way is invisible to the error controller and makes the state trajectory discontinuous, which destabilises an adaptive BDF integrator. It has to become a continuous buoyancy source evaluated inside the right-hand side,

$$\dot{Q}_{\mathrm{buoy},k} = -\dot{Q}_{\mathrm{buoy},k+1} = \frac{C}{\tau_{\mathrm{buoy}}} \max(0,\ T_{k+1} - T_k)\,, \qquad C = \rho\, c_p V_k, \tag{6}$$

applied antisymmetrically to each unstable pair of control volumes, so that energy is conserved exactly when their volumes are equal, which is one reason that variant is so far restricted to prismatic geometry. The single parameter τ_{buoy} sets how fast the mixing acts: $\tau_{\mathrm{buoy}} \to 0$ recovers the instantaneous caloric mixing, larger values spread it over time. $\tau_{\mathrm{buoy}} = 60\,\mathrm{s}$ is a practical default, since the effective mixing time scales as $\tau_{\mathrm{eff}} \propto \tau_{\mathrm{buoy}}\, K^2$ while smaller values stiffen the ODE. Relaxation-type corrections of this form are also used in Modelica stratified-storage components (Wetter et al., 2014). That integration and its cross-validation against five reference models are reported in the thesis; they rest on a separate C++ implementation, so no results from them are claimed here.

5. Discussion

At both sites the dominant error occurs at mid height, where the thermocline sits. Turbulent mixing at a bidirectional diffusor and convection in the surrounding ground, which neither the one-dimensional advection nor a constant heat transfer coefficient can represent, are known limitations of this type of model (Ochs et al., 2021; Xiang et al., 2022), but the present results do not allow the residual to be attributed to any one of them. What the results do show (Fig. 4) is that a three-point validation can understate the error of a stratification model by a wide margin where the thermocline sits between the ports. Two mechanisms are lumped rather than predicted, and both are natural targets for improvement, namely the mixing where streams merge (Sect. 2.2)

and the choice of f_λ. The two-port interface is in fact marginally *better* over the whole profile (Table 5), which points to the injection at the middle port rather than to the merged transport. The network-relevant top and bottom positions are reproduced far better (1.8–4.0 and 1.3–2.0 K across the two sites), so this has limited bearing on supply-temperature prediction. Spatial resolution controls the sharpness of the front. Refining the grid from $K = 50$ to $K = 200$ changes the deviation by less than 0.05 K for the weekly and annual cycles, but still halves it under fast switching (Table 2), so K should be chosen for the fastest dynamics of interest.

Three quantities in the model are not measured: the effective-conductivity factor f_λ; U_{wall}, derived as $\lambda_{\text{soil}}/d_{\text{eff}}$ with an effective depth that is not directly observable; and the diffusor mixing-zone height. Five-point sweeps of the first two, scored as the MAE at the three diffusor heights, give an asymmetric picture. Over $f_\lambda = 1$–15 the Dronninglund error is flat (3.36–3.62 K) whereas Høje Taastrup degrades by 1.5 K (2.07–3.60 K); $f_\lambda = 1$ is the value Falkner (2014) gives for large storages without internals and is used at both sites. It is also the best-performing value here, the one parameter for which the literature choice and the data-optimal one coincide. Varying U_{wall} three-fold about the value used at each site leaves Høje Taastrup essentially unchanged (2.04–2.16 K, the well-insulated lid dominating) but moves Dronninglund by 1.5 K (2.78–4.31 K). That the *derived* 0.35 $\text{W m}^{-2}\text{K}^{-1}$ is *not* the best-performing value (0.20 would give 2.78 K instead of 3.36 K) is direct evidence that it came from construction and soil data rather than tuning. The mixing-zone height is fixed at the 1 m documented default; doubling it to 2 m changes the profile-wide MAE by less than 0.4 K at both sites.

The two coupling schemes differ in what they buy and what they cost. A sequential co-simulation keeps the storage a black box behind a stable port interface and adds only a few percent to the network solve; solving the storage inside the network solver's own equation system removes the one-step lag but constrains the model to a smooth right-hand side. Since the coupling deviation stays below a few tenths of a Kelvin at district-heating time steps of 900–3600 s (Sect. 4.4), the sequential scheme is the recommended default; the in-solver route is warranted where the storage shares one implicit system with a stiff network.

Two limitations remain: The measurement validation uses the standalone model with truncated-pyramid PTES geometry and a split lid/wall loss model, whereas the in-solver variant has so far been exercised only on cylindrical geometry with a uniform loss coefficient; extending the production solver and repeating the two PTES runs inside it is the natural next validation step. The second limitation concerns the ground: the constant-coefficient wall loss used here could be replaced by a discretised ground domain with a far-field boundary condition, which would represent the start-up transient of a first operating year that a constant coefficient cannot. On the evidence of Sect. 4 that would address part, but not all, of the Dronninglund offset.

6. Conclusions

A modular, open-source 1D FV model for stratified thermal energy storage with a multi-port hydraulic interface has been presented. Exchangeable sub-models adapt it from steel cylinders to large PTES without changes to the solver core, and the cumulative-flux formulation handles any number of connections uniformly. Exposing the same model through three hydraulic interfaces (Sect. 4.2) shows what this achieves, which is not accuracy everywhere but the ability to reproduce a storage as it is built. A correctly implemented two-port interface stays competitive, yet it cannot represent the middle diffusor, and without its switching logic it deviates by 13.6 K.

Benchmarked against a segment-model reference across three scenarios, the implicit solver ($K = 50$) reaches a deviation of 0.47–1.44 K at 0.16–0.25 ms per time step, well within the coupling-step budget used here; the fully mixed 0D model deviates by 16–19 K in dynamic cycling, confirming that stratification modelling is essential for network supply-temperature prediction.

Against the internal sensor strings of two Danish PTES sites, the overall MAE at the three diffusor heights is 2.07 K at Høje Taastrup ($V = 70\,880\ \text{m}^3$, full year 2024) and 3.36 K at Dronninglund ($V = 60\,030\ \text{m}^3$, 8 months of 2014); over the complete measured profile the corresponding values are 2.05 and 5.40 K. At Høje Taastrup the model also reproduces the stored energy to 0.6 K, so the residual there is one of thermocline position alone; the quantitative claim rests on that site. Dronninglund additionally carries a systematic −5.2 K energy-balance offset from the plant's first operating year, reported rather than resolved here.

The coupling path was demonstrated by co-simulation with the open-source solver pandapipes, with the storage acting as a hydraulic separator; the coupling adds 1.5 % or less to the network solve cost, and the sequential exchange deviation stays well below the model's own measurement error at district-heating time steps.

Benchmark, measurement validation and co-simulation together show that one lightweight element spans the chain from numerics to network operation, and they locate the gaps that remain: the thermocline zone at mid

height, and the mixing at merging ports that the formulation lumps into f_λ and the diffusor zone.

Acknowledgements

This work was carried out within the ESF-funded junior research group "Energy Storage" (ESF-Nachwuchsforschergruppe) at Zittau/Görlitz University of Applied Sciences (HSZG). It is co-funded by the European Union and co-financed from tax revenues on the basis of the budget adopted by the Saxon State Parliament.

Code and data availability

The model and both measurement validations are open source at `https://github.com/JonasPfeiffer123/thermal-energy-storage-1d`; the benchmark and co-simulation scripts will follow there. Both PTES datasets are openly available (Sifnaios et al., 2023, 2025).

References

Bayer, M., Meister, C., Schuetz, P., Villasmil, W., Walter, H., Dahash, A., 2025. Development of a reduced-order dynamic model for large-scale seasonal thermal energy storage applications. Energy 333, 137379.

Dahash, A., Ochs, F., Tosatto, A., Streicher, W., 2020. Toward efficient numerical modeling and analysis of large-scale thermal energy storage for renewable district heating. Appl. Energy 279, 115840.

De la Cruz-Loredo, I., Zinsmeister, D., Licklederer, T., Ugalde-Loo, C.E., Morales, D.A., Bastida, H., Perić, V.S., Saleem, A., 2023. Experimental validation of a hybrid 1-d multi-node model of a hot water thermal energy storage tank. Appl. Energy 332, 120556.

Drück, H., 2006. MULTIPORT Store-Model for TRNSYS, Type 340. Technical Report. Institut für Thermodynamik und Wärmetechnik (ITW), Universität Stuttgart.

Falkner, S., 2014. Modellierung und Simulation von thermischen Speichern. Diplomarbeit. Institut für Energietechnik und Thermodynamik (E302), Technische Universität Wien.

Lago, J., De Ridder, F., Mazairac, W., De Schutter, B., 2019. A 1-dimensional continuous and smooth model for thermally stratified storage tanks including mixing and buoyancy. Appl. Energy 248, 640–655.

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, 9899.

Narusavicius, B., Koch, F., Matthees, A., Stange, P., Rühling, K., 2024. Prediction of the temperature field development in large heat storage tanks for integration in network simulation, in: 34. Symposium Solarthermie und Innovative Wärmesysteme.

Ochs, F., Dahash, A., Tosatto, A., Reisenbichler, M., O'Donovan, K., Gauthier, G., Kok Skov, C., Schmidt, T., 2021. Comprehensive comparison of different models for large-scale thermal energy storage, in: Proceedings of the International Renewable Energy Storage Conference (IRES 2021).

Pfeiffer, J., 2026. Modellierung und softwaregestützte Planung zentraler Wärmespeicher in Wärmenetzen mit mehreren Erzeugern: Integration in VICUS Districts. Master's thesis. Hochschule Zittau/Görlitz.

Reisenbichler-S., M., Sifnaios, I., Wotawa, F., 2025. Validation of a pit thermal energy storage model: Demonstration of a comprehensive approach. J. Energy Storage 116, 115680.

Rendall, J., Karg Bulnes, F., Gluesenkamp, K., Abu-Heiba, A., Worek, W., Nawaz, K., 2021. A flow rate dependent 1D model for thermally stratified hot-water energy storage. Energies 14, 2611.

Schmidt, T., Sørensen, P.A., 2018. Monitoring results from large scale heat storages for district heating in Denmark, in: EnerSTOCK2018, 14th International Conference on Energy Storage, Adana, Turkey.

Sifnaios, I., Furbo, S., Jensen, A.R., 2025. Monitoring data of the Høje Taastrup water pit thermal energy storage. Data Brief 58, 111305.

Sifnaios, I., Gauthier, G., Trier, D., Fan, J., Jensen, A.R., 2023. Dronninglund water pit thermal energy storage dataset. Sol. Energy 251, 68–76.

Untrau, A., Sochard, S., Marias, F., Reneaume, J.-M., Le Roux, G.A.C., Serra, S., 2023. A fast and accurate 1-dimensional model for dynamic simulation and optimization of a stratified thermal energy storage. Appl. Energy 333, 120614.

Wetter, M., Zuo, W., Nouidui, T.S., Pang, X., 2014. Modelica Buildings library. J. Build. Perform. Simul. 7, 253–270.

Xiang, Y., Xie, Z., Furbo, S., Wang, D., Gao, M., Fan, J., 2022. A comprehensive review on pit thermal energy storage: Technical elements, numerical approaches and recent applications. J. Energy Storage 55, 105716.

Zoschke, T., 2025. pandapipes pull request #773: new component heat_generator.