POLITECNICO DI TORINO

Master of Science Course
in Energy Engineering - Renewable energy systems

Computational thermo-fluid dynamics study on the optimization of a thermocline energy storage system prototype

Tutor
Prof. Roberto Zanino

Candidate
Simone Donna S331481

Co-Tutor
Prof. Mattia Cagnoli

A.Y. 2024-2025

Abstract

Single-tank thermocline thermal energy storage (TES) systems using molten salt offer significant cost advantages for concentrating solar power plants, but maintaining sharp thermal stratification remains challenging due to natural convection-induced mixing. This thesis investigates a novel thermocline TES configuration comprising a cylindrical tank filled with HITEC XL molten salt and an internal recirculation channel equipped with electrical heating resistors, representing a hybrid solar-electrical storage approach.

The study employs computational fluid-dynamics (CFD) analysis using STAR-CCM+ to evaluate thermocline formation under natural and forced convection conditions. A multi-stage methodology is developed: first, a three-dimensional model characterizes the heating rod bundle across various operating conditions with pressure drop predictions and outlet temperatures values. A simplified two-dimensional axisymmetric model then simulates the complete tank charging transient with the computed Darcy-Forchheimer law parameters, comparing the two models result, the porous media simplification is validated and the full TES simulation can be run to analyze the stratification behavior and the flow characteristics.

The results are consequently compared to analyze if the forced convection improves the stratification of the fluid and the total efficiency of the system in comparison with the natural convection driven by the buoyancy forces. This research validates computationally efficient CFD modeling strategies combining detailed geometry with porous media approximations and demonstrates that controlled forced convection represents an effective approach for optimizing thermocline TES systems, improving both stratification and operational flexibility for renewable energy applications.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The global energy landscape is undergoing a profound transformation, driven by the dual imperatives of meeting increasing energy demands and mitigating climate change. This transition places renewable energy sources, particularly solar energy, at the forefront of sustainable energy production. Although solar energy technologies can generate electricity on a large scale, their intrinsic intermittency, along with the availability of sunlight, necessitates robust energy storage solutions. Among these, Thermal Energy Storage (TES) stands out as a highly effective and cost-efficient method for storing excess thermal energy, ensuring a stable and dispatchable power supply even after sunset.

According to a recent analysis [1], integrating TES systems into solar plants can significantly increase their overall efficiency, enabling them to operate for up to 24 hours a day. This not only maximizes power plant utilization, but also provides crucial grid stability by reducing reliance on fossil fuel backups. The ongoing research and development in this area is focused on overcoming key challenges such as lowering material costs and improving thermal efficiency.

This thesis contributes to this field by focusing on the analysis of a thermocline single tank TES system. This technology presents an economically attractive alternative to conventional two-tank systems using a single tank to store both hot and cold fluids. During the charging phase, hot fluid is introduced at the top, while the colder fluid remains at the bottom. A mobile thermal transition layer, known as the thermocline, forms between the two regions and acts as a natural thermal barrier. This stratification minimizes fluid mixing and heat losses, allowing for a simplified and cost-effective design with the main challenge being to maintain a stable and distinct thermocline. As noted in [2], natural convection can induce un-

desirable circulation, leading to fluid mixing and a degraded temperature gradient.

The goal is to study this phenomenon in a specific setup: a tank filled with HITEC XL molten salt and an internal heating channel equipped with immersed electrical resistors. This hybrid configuration, combining solar thermal energy with electrical heating for power-to-heat applications, is of particular interest.

The core of this work involves building and refining a numerical model to evaluate the formation of the thermocline. A key aspect of the study is to compare the effects of natural convection with those of forced convection, which will be simulated by adding an impeller. The objective is to determine whether forced circulation can improve thermal stratification and overall system efficiency. The computational thermo-fluid dynamic (CFD) model, built on an optimized geometry from previous research [3], will provide a detailed analysis of the fluid behavior under both scenarios. Ultimately, this research seeks to establish whether a stable and functional thermocline can be achieved in this configuration and to offer concrete recommendations for future design improvements.

Chapter 2

Storage system setup

2.1 TES working principles

The system under study is classified as a high temperature thermocline sensible thermal energy storage system. Energy (E) is stored as heat through to the temperature difference between the top and lower layer, both higher than 120°C. In fact, the energy needed to heat up the storage fluid is directly proportional to both the temperature difference and material thermal properties [4]:

$$E = \int_V \rho(T) \cdot c_p(T) \cdot (T(x, y, z) - T_{ref})\, dV \tag{2.1}$$

where $\rho(T)$ and $c_p(T)$ are respectively the material's density and specific heat according to the temperature while $T(x, y, z)$ is the fluid temperature and T_{ref} is the system's initial temperature.

To have an efficient sensible TES is evident that three main requirements are needed. The first, the most important one, is the guarantee of an optimal thermal stratification. The tank must be able to separate volumes at different temperatures so that the mixing layer is minimize also during the charging and discharging phase. Moreover, the tank's design must avoid any dead zone, an area where the fluid remains stagnant and does not effectively participate in the charging or discharging processes, in fact in a well-designed stratified tank, the effective portion is maximized, allowing most of the fluid to contribute to energy storage. Finally, the tank must minimize the heat losses to the environment thanks to a good thermal insulation.

Beyond minimizing dead zones and heat losses, another crucial aspect for the efficiency of a high-temperature thermocline TES is its capacity to store exergy (Ξ). While thermal energy is a quantitative measure of storage capacity, exergy evaluates the quality of the stored energy, or its potential to perform useful work. Since exergy is directly related to the temperature difference with a reference environment, operating at higher temperatures not only increases the stored energy but also significantly enhances the system's exergy. The 'three-zone model' [5] (hot temperature, thermocline, and cold temperature) is a common approach to analyze how the temperature distribution within the tank affects both exergy and the overall thermodynamic performance of the system. A key figure of merit to provide the level of stratification is the exergetic ratio:

$$\frac{\Xi}{\Xi_m} = \frac{(T_m - T_0) - T_0 \cdot ln(T_e/T_0)}{(T_m - T_0) - T_0 \cdot ln(T_m/T_0)} \tag{2.2}$$

in which Ξ_m is the exergy in a fully mixed tank with no stratification, T_m is the average fluid temperature, T_0 is the dead state temperature set as reference and T_e is a convenience quantity that takes into account the 3-zones model approximation, since it is dependent to the temperature delta between the lower and upper thermocline region layers.

Moreover, to evaluate the stability of thermal stratification and quantify the balance between buoyancy-driven and inertially-driven flows, the Richardson number (Ri) is introduced as an additional figure of merit [6]. The Richardson number is defined as the dimensionless ratio between buoyancy forces and inertial forces:

$$Ri = \frac{Gr}{Re^2} = \frac{g\beta(\Delta T)H}{u^2} \tag{2.3}$$

where β is the volumetric thermal expansion coefficient, ΔT represents the characteristic temperature difference in the thermocline, H is the tank height, and u is the characteristic flow velocity. Values of $Ri > 10$ indicate that buoyancy forces dominate, favoring stable stratification and minimal mixing. Conversely, $Ri < 0.1$ suggests that inertial effects overwhelm buoyancy, leading to significant mixing and thermocline degradation. Finally, $0.1 < Ri < 10$ means that the forces are balanced and neither of the two is predominant on the other [7]. This dimensionless number provides a quantitative assessment of whether natural convection alone can sustain the required thermal stratification in the storage system.

2.2 Geometry and material

The geometry of this numerical model is a direct representation of a prototype developed at the ENEA research center, specifically designed for a single-medium (molten salt) thermal energy storage system. The geometry consists of a tank with a 1.3 meters diameter, a total height of 2.65 m and a inner recirculating channel that aims at facilitating the buoyancy-driven motion of the salt and heat it trough a bundle of rods with a total power of 110 kW [3]. The internal recirculation channel is a critical structural element engineered to manage the fluid dynamics, essential for optimal thermal stratification. Its primary physical function is to control buoyancy-driven flow during the charging phase. When the internal resistors heat the molten salt, the resulting decrease in density generates a pressure head that drives the hot fluid upward according to the natural convection phenomenon. The channel acts as a low-impedance, confined pathway that guides this flow from the heating section to the tank's upper layer. By channeling the ascending flow, the system effectively bypasses the formation of large, chaotic recirculation cells within the main storage volume that would typically cause detrimental mixing. This crucial management of the convection currents is the mechanism by which the channel maintains a stable and sharp thermocline region. To investigate the impact of active mixing on thermal stratification, an impeller is introduced into the tank geometry. The impeller is modeled as a constant pressure-drop device applied over the internal channel region. This configuration is intended to study how mechanical circulation interacts with the naturally buoyancy-driven flow. Finally, the outer walls of the tank are provided with suitable insulation to minimize heat losses to the environment. The computational fluid dynamics (CFD) model is used to simulate only the charging phase of the tank.

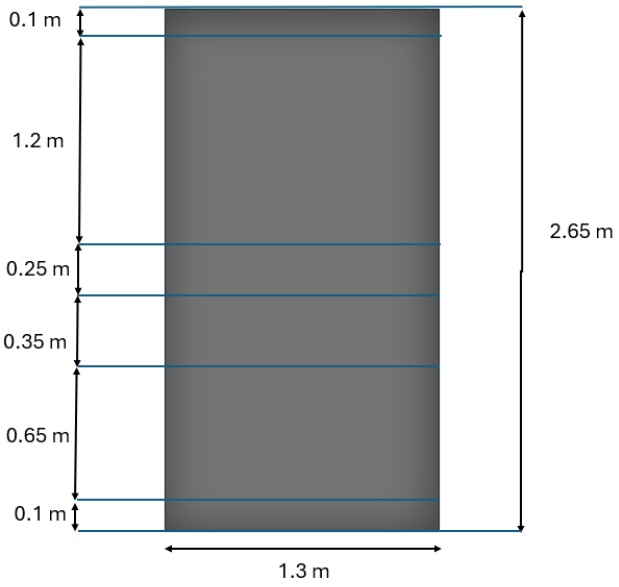

Figure 2.1: Tank dimensions

The internal recirculating channel is comprised of three main sections. The lower section serves as the inlet, where molten salt can be propelled by an impeller, acting as the primary driver for forced convection. The flow then splits into three separate channels to facilitate efficient heating with an optimized geometry of 75 heating rods for each separate channel.

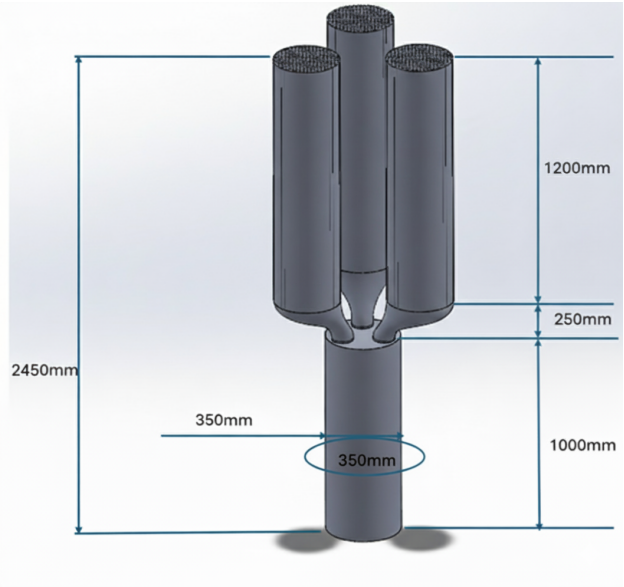

Figure 2.2: Recirculating channel

The optimal layout for the internal recirculating heating device is selected based on a previous parametric study [3], focused exclusively on the thermal-hydraulic performance of the heat exchanger bundle. The objective is to optimize the number and distribution of the heating rods within a single channel to efficiently transfer the required 36.67 kW of power, one third of the total 110 kW of power injected.

The study characterizes the most effective rods number and distribution in terms of three parameters.

The first one is the thermal uniformity and material integrity, measured in terms of wall temperature distribution across the rods. The chosen 75-rod layout is selected because it offers a sufficiently large heat transfer area to meet the thermal requirement while keeping the local surface temperature of the rods below critical thresholds. This mitigates the risk of molten salt decomposition or accelerated material corrosion, ensuring the long-term integrity of the metallic structure.

The second requirement is the minimal stagnation that is achieved via a velocity field analysis, in fact this configuration is found to minimize stagnation zones (dead zones) within the channel. Maximizing the active flow around the rods ensures that all of the fluid contributes to heat transfer, preventing localized overheating and improving the overall convective heat transfer coefficient.

Lastly, the alignment with theoretical performance is taken into consideration. The simulated outlet salt temperature of the chosen layout is close enough to the theoretical outlet temperature computed via the calculation of an energy balance based on the First Law of Thermodynamics.

The materials involved in this study are the molten salt used as the energy carrier and the metal used for the recirculating channel metallic structure. The molten salt mixture considered for this study is the HITEC XL[8], a tertiary salt composed of 45% KNO_3, 7% $NaNO_3$, 48% $Ca(NO_3)_2$, which is commonly used in this field thanks to its lower melting and crystallization point compared to binary solar salts. It prevents solidification at lower temperatures and reduces the need for freeze protection systems and associated cost while maintaining stability and functionality over a broad range of temperatures, allowing for efficient collection, transfer, and storage of thermal energy. The thermal properties, respectively specific heat (c_p), density (ρ), thermal conductivity (λ) and dynamic viscosity (μ), are computed according to the following formulas [9]:

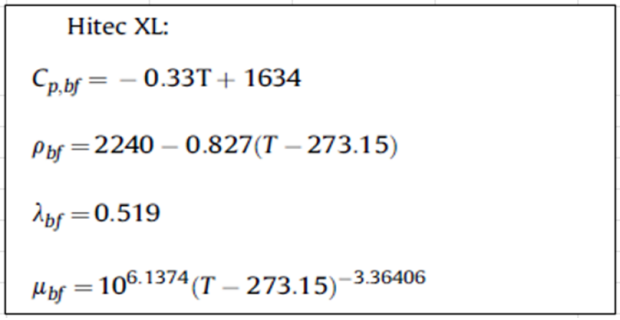

Figure 2.3: Thermal properties of HITEC XL

Hence, the thermal properties considered for the CFD simulation, introduced as tabular data, are the following:

T **[K]**	ρ **[kg/m^3]**	c_p **[J/kg·K]**	μ **[Pa·s]**	λ **[W/m·K]**
403.15	2132.49	1500.96	0.10616	0.519
424.15	2115.12	1494.03	0.06415	0.519
445.15	2097.76	1487.10	0.04140	0.519
466.15	2080.39	1480.17	0.02810	0.519
487.15	2063.02	1473.24	0.01985	0.519
508.15	2045.66	1466.31	0.01449	0.519
529.15	2028.29	1459.38	0.01086	0.519
550.15	2010.92	1452.45	0.00833	0.519
560.15	2002.65	1449.15	0.00756	0.519

Table 2.1: Thermal properties of the HITEC XL molten salt mixture.

For the metallic structure, the choice made it is the SS304 [3], also known as AISI 304. It is an austenitic stainless steel which is of common use globally due to its excellent corrosion resistance due to high chromium and nickel content, good mechanical properties and ductility, and good heat resistance up to high temperatures.

Property	Value
Thermal Conductivity (λ)	16.2 W/m $\cdot$ K
Specific Heat Capacity (c_P)	$500 - 550$ J/kg $\cdot$ K
Density (ρ)	7980 kg/m^3

Table 2.2: Thermal properties of AISI 304

2.3 Model description

The model is split into two primary domains: the first one is the tank filled with molten salt, which serves as the effective thermal energy storage, and the second one is the internal recirculating channel. As mentioned in the previous section, this channel is of key importance and, for a better understanding of its behavior, it is further segmented into three distinct regions to accurately reflect the prototype's geometry. Starting from the bottom, it is present the lower channel which is needed as the gatherer of the salt coming upward driven by the buoyancy forces and it will also host the impeller for the second step of the performance evaluation. Then, the duct is designed in order to split up the flow in three different channels, with the same dimensions of the lower one. The rationale behind it is to have an increment of the heating uniformity and a convection enhancement in the top and last section which is the heating rods bundle region that has been widely explained previously.

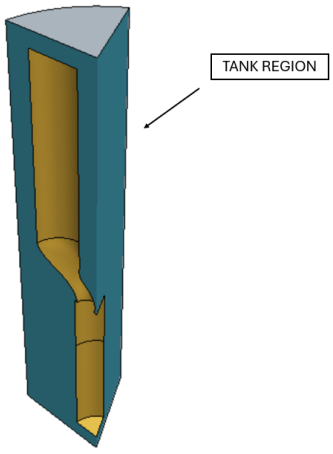

Figure 2.4: Tank computational domain

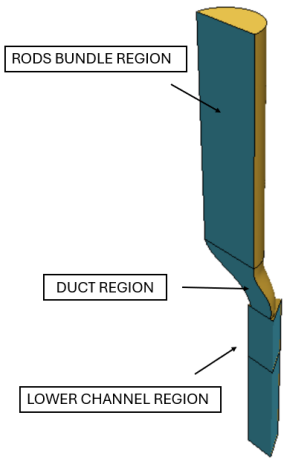

Figure 2.5: Recirculating channel computational domain

The figures above represent the minimal computational domain that will be taken into account in the simulation run. The detailed explanation on how the domain is

built in the STARCCM+ environment is provided in the following chapter. Nevertheless, it is important already to make the several strategic simplifications adopted clear to ensure fluent understanding of the model's structure.

The adoption of a minimal computational domain is the first strategic simplification, fundamentally driven by the need to significantly reduce computational costs while preserving the accuracy of the transient solution. This compromise is made possible by the inherent rotational symmetry of the physical setup, which applies to both the cylindrical main tank and the internal recirculating channel structure. By leveraging this symmetry, the model only represents the minimal non-divisible unit of the geometry. Specifically, the domain is divided into six identical sectors, allowing the simulation to utilize just one-sixth (60° sector) of the total physical volume. The radial boundaries of this sector are then defined as symmetry planes within the simulation environment, enabling the iterative solver to accurately extrapolate the computed fluid flow and temperature fields to the full geometry.

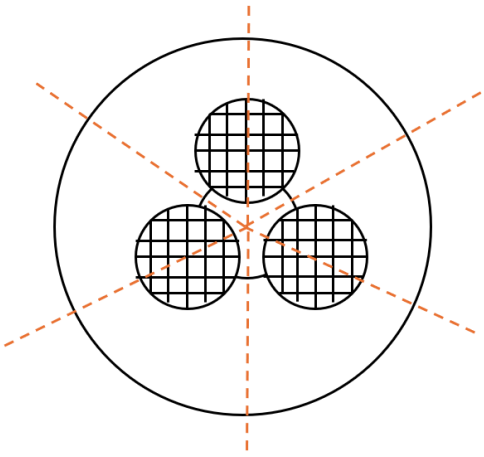

Figure 2.6: Tank view from above

Moreover, the rods bundle region is the core of the thermophysical heat exchange. A detailed representation of this section is therefore key to a high-fidelity model. On the other hand, explicitly resolving the complex geometry, including all 75 heating rods and the narrow fluid gaps between them, would require an extremely dense computational mesh. Such mesh would dramatically increase the number of computational cells, leading to prohibitively long simulation times and excessive computational resource consumption for the required transient analysis of the entire tank. Furthermore, the focus of the analysis is on the macro-scale thermocline region formation phenomena, hence capturing the cell-to-cell variations in flow and energy is not essential. Therefore, a porous media model approach is used. This approach replaces the complex rod geometry with an equivalent con-

tinuous fluid domain characterized by effective transport properties. By utilizing the empirically derived Darcy–Forchheimer coefficients permeability (K) and inertial coefficient (J), the model accurately replicates the essential pressure drop (ΔP) and distributed heat source effects of the rod bundle on the molten salt flow, preserving the crucial macro-scale thermo-fluid dynamic behavior while ensuring computational efficiency. This concept will be expanded and deeply explained in the section 3.3.

Lastly, for the purpose of simplification, all external walls of the studied domains are considered adiabatic unless otherwise specified. This is a reasonable assumptions since the prototype's walls are well insulated.

Chapter 3

Porous Media Characterization

3.1 Three-Stage Calibration Procedure

As discussed in the previous chapter, representing the bundle of heating rods through a porous media approach is fundamental to obtain a model that is both computationally efficient and physically accurate. The main advantage of this approximation is that it reproduces the pressure losses caused by the rod array without explicitly resolving its complex geometry. In this way, the mesh complexity and the simulation time are significantly reduced, while the essential thermo-fluid dynamic behavior of the molten salt in the heated section is preserved.

The characterization procedure is carried out in three stages:

In the first stage, a high-fidelity three-dimensional model of the rods bundle region is developed. This model explicitly includes the rods and makes it possible to compute the pressure drop as a function of both mass flow rate and temperature. The resulting dataset, which reflects the reference behavior of HITEC XL molten salt under realistic operating conditions, provides the basis for the porous media calibration.

In the second stage, the pressure drop data obtained from the resolved geometry are correlated using the dimensionless formulation of the Darcy–Forchheimer equation, which will be deeply analyzed in Section 3.3. A nonlinear least-squares regression is then performed to identify the permeability K and the inertial coefficient J of the equivalent porous medium. These parameters are finally used to construct the porous media model that replaces the explicit rod bundle in the subsequent CFD simulations. The MATLAB implementation of the regression procedure is reported in Appendix A.

Finally, the two-dimensional model is once again compared with a second iterating run for each mass flow rate and temperature dataset, activating the total heat source option. The values of outlet temperature are then compared with the theoretical temperatures manually calculated to check for the validity of the model built.

3.2 Rods bundle region model

According to what is stated above, the rods bundle region it is itself a rotational solid, hence the computational domain can be reduced to a representative portion of the whole volume. In this case, a sector corresponding to one-fifth (72° sector) of the total geometry is selected. This choice differs from the minimal computational domain of the whole tank because in this case the priority is to investigate the behavior of the flow through the narrow spaces of the rods and, thanks to this domain, exactly 15 rods out of the 75 total rods of the bundle can be designed inside the domain. This reduction ensures an accurate representation of the system while significantly improving the efficiency of the simulations.

The model of the rods bundle region is performed using a steady state simulation approach, since the primary goal is geometry characterization, not transient analysis. In particular, the aim is to find out the pressure drop along the region and the outlet temperature given a set of inlets of mass flow rates and temperatures. Moreover, modeling the full transient evolution would have been computationally expensive and not necessary at this stage. A segregated approach is selected for computational efficiency. This scheme solves the governing equations (continuity, momentum, and energy) sequentially rather than simultaneously (as in a Coupled Flow model). For the laminar flow regime encountered in the narrow channels of the rod bundle, the segregated solver often converges quickly and reliably, leading to significant memory and CPU time savings compared to the more resource-intensive Coupled solver. Finally, a laminar flow regime model is adopted based on a preliminary assessment of the Reynolds number (Re) within the narrow gaps between the heating rods.

$$Re = \frac{\rho \, u \, D_h}{\mu} \tag{3.1}$$

Although the molten salt flow is being circulated, the high viscosity of the HITEC XL salt combined with the small characteristic hydraulic diameter of the channels ensures that the flow velocity remains sufficiently low. The resulting Re values are not greater than 100, according to the data in 2.1, well within the range where the flow is predominantly viscosity-dominated, justifying the use of the laminar model. This selection is crucial for minimizing computational cost, as using a

Reynolds-Averaged Navier-Stokes (RANS) turbulence model for a laminar flow would introduce unnecessary complexity and potentially inaccurate results.[10]

The region domain boundaries are set as follows:

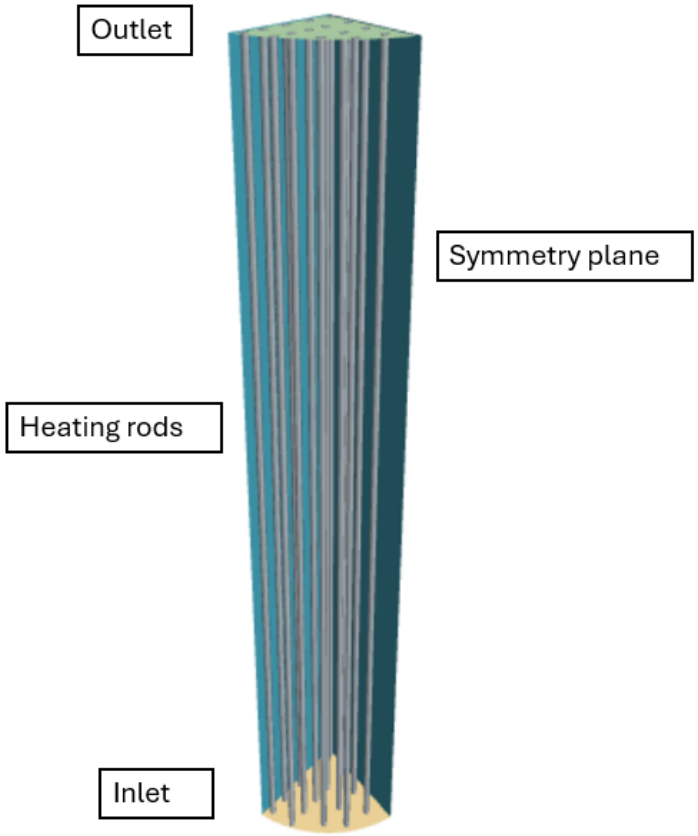

Figure 3.1: Rods bundle region boundaries

As we can see from the figure above, the boundaries are the inlet one in which the input dataset are inserted, the external wall and the rods walls, for which the no-slip condition is activated (Dirichlet condition in which the velocity is null) and set as adiabatic. In addition, the symmetry boundary condition is present representing an ideal boundary that divides the domain in half without altering the flow physics, imposing no flux trough the plane. To complete the setup the outlet boundary is also present.

The dataset used at the inlet boundary of the domain are the following:

$\dot{m}$ **[kg/s]**	T_{in} **[K]**
0.10	460.15
0.20	480.15
0.50	500.15
0.75	520.15
1.00	540.15
1.20	560.15

Table 3.1: Mass flow rate and inlet temperature values used for porous media characterization.

The selected values follow the considerations reported in [3], with the rationale being to cover the expected range in which the system will operate to capture both the viscous and inertial effects on pressure drop.

3.2.1 Mesh grid independence study

The grid independence study is a fundamental step to ensure the simulation tool provides accurate results without incurring excessive computational cost. This section details the necessary mesh accuracy. For this analysis, the pressure drop along the heating region at a mass flow rate of 0.1 kg/s and 460.15°C is set as the primary comparison parameter, as it is the most significant metric for the subsequent porous media characterization. Five meshes labeled as finest, finer, fine, coarse and coarser, will be evaluated to determine the optimal one. The setups for these five meshes, along with their relative pressure drop values and required computational time, are listed below:

	Finest	**Finer**	**Fine**	**Coarse**	**Coarser**
Cell size [m]	0.01	0.01	0.01	0.03	0.05
Volume growth rate	1.5	1.7	1.7	1.8	1.8
Surface growth rate	Slow	Slow	Default	Default	Fast
Number of cells	936 930	842 538	769 964	183 924	78 654
Computational time [s]	13 358	5 225	2 964	792	530
ΔP [Pa]	0.144	0.146	0.142	0.113	0.101
Relative error [%]	–	1.55	1.46	27.39	42.04

Table 3.2: Mesh refinement study results.

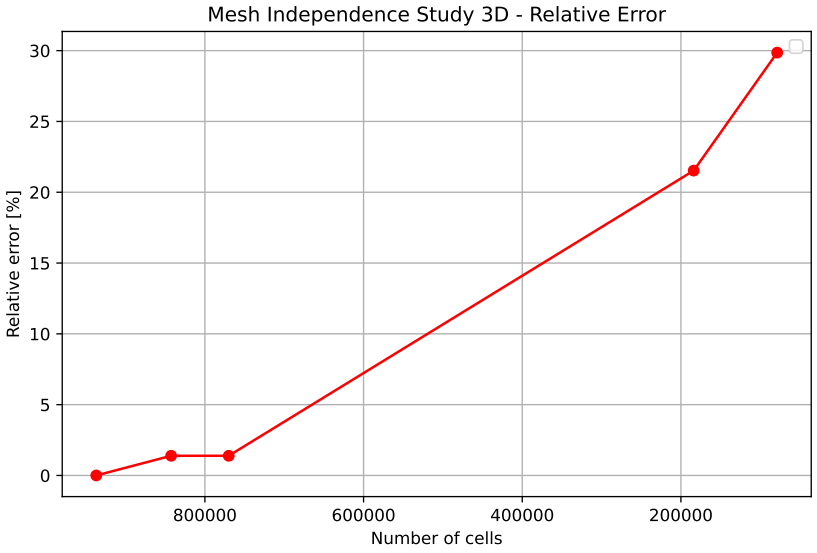

Figure 3.2: 3D Mesh independence study - Relative error plot

Thanks to the above dataset, the fine mesh is selected as the best option. It is the best trade-off between accuracy, along with the finer mesh, with a relative error around 1%, and computational cost. In fact it is considerably faster, almost the double, compared to the finer mesh. For these reasons, the fine mesh, shown below, will be the mesh used for the rest of the analysis.

Figure 3.3: Fine mesh of the rods region

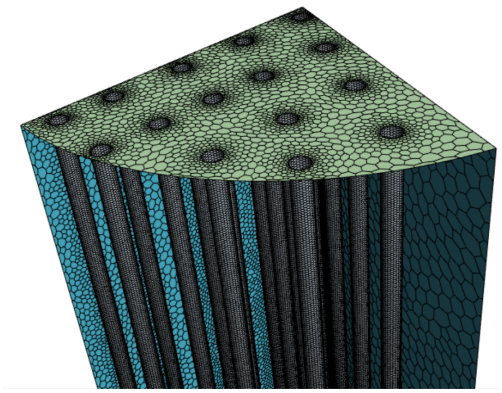

Figure 3.4: Fine mesh detail

3.2.2 Preliminary results

The simulation is run for each dataset up to 450 iterations, as it is sufficient for the pressure drop computed to stabilize. This indicates that the numerical model has accurately solved the fluid dynamics equations.

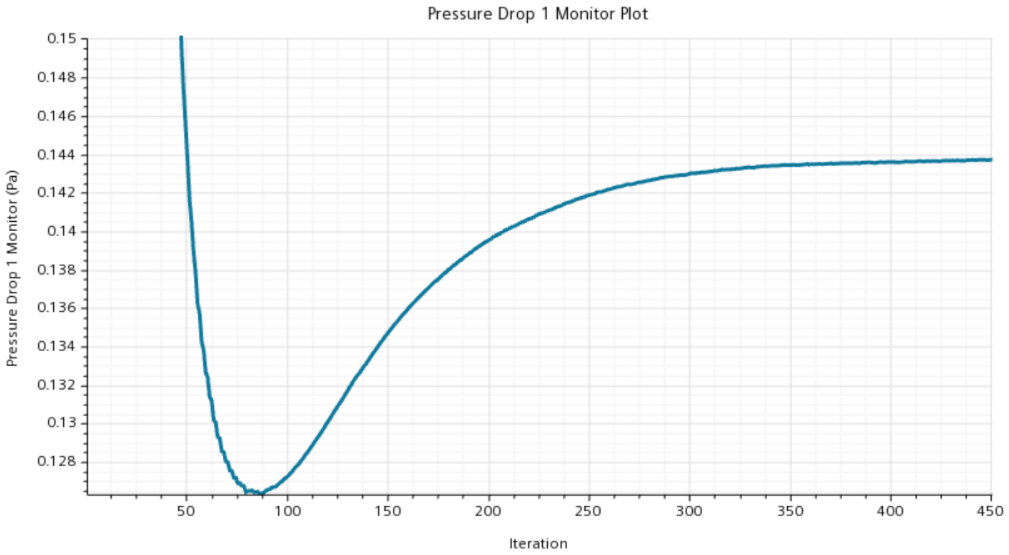

Figure 3.5: Pressure drop monitor along the rods bundle

As mentioned above, thermal effects are neglected to isolate the hydraulic behavior of the molten salt. The results of pressure drops are listed below:

$\dot{m}$ **[kg/s]**	T_{in} **[K]**	ΔP **[Pa]**
0.10	460.15	0.142
0.20	480.15	0.181
0.50	500.15	0.327
0.75	520.15	0.379
1.00	540.15	0.404
1.20	560.15	0.397

Table 3.3: Pressure drop values for the rods bundle region.

3.3 Darcy-Forchheimer equation approximation

The Darcy-Forchheimer law can be written in the following form:

$$\frac{\Delta P}{d} = \frac{\mu}{K}\phi\, u + \frac{J}{\sqrt{K}}\rho(\phi\, u)^2 \tag{3.2}$$

where d is the length of the bundle cylindrical channel, ϕ is the porosity of the bundle and u is the flow velocity. $\sqrt{K}$ is considered as the characteristic length, representing the typical dimension of the flow passage within the bundle. This law is of fundamental importance because accounts for both viscous and inertial effects. The first term, the proper Darcy's Law, represents the viscous contribution, dominant at low velocities, where pressure losses arise from friction between the fluid and the porous matrix. The second term, the Forchheimer contribution, captures the inertial component, which becomes relevant at higher velocities and represents additional pressure losses due to flow acceleration and turbulence within the pores. By fitting CFD data to this relation, K and J are determined (appendix: A).

Once the pressure drops for the 3D simulation have been evaluated, the porous media approximation thanks to the Darcy-Forchheimer law can be performed. The theoretical basis of this method can be found in [11], however an overview is important in order to understand the next steps.

To reproduce the hydraulic behavior of the rods bundle with an equivalent porous medium, the parameters of the Darcy–Forchheimer model are identified by fitting CFD results. The CFD simulations provides pressure drops corresponding to a range of mass flow rates and inlet temperatures as reported in Table 3.3. The geometrical characteristics of the porous region (flow area (A_f) , d and ϕ) are also considered, while the thermo-physical properties of the molten salt are computed as temperature-dependent functions for viscosity ($\mu(T)$) and density ($\rho(T)$), according to Figure 2.3.

The Darcy–Forchheimer correlation is expressed in its adimensional form:

$$f_K = \frac{1}{Re_K} + J \tag{3.3}$$

where:

$$Re_K = \frac{\dot{m}\,\phi\,\sqrt{K}}{\mu(T)\,A_f} \tag{3.4}$$

$$f_K = \frac{\Delta P\,\rho(T)\,\sqrt{K}\,A_f^2}{d\,\phi^2\,\dot{m}^2} \tag{3.5}$$

with K being the permeability of the medium and J the inertial coefficient.

The identification problem is solved by minimizing the residuals between the CFD-derived values of f_K and the model prediction. The nonlinear least-squares optimization was carried out in MATLAB using the function *lsqnonlin*, with the optimization variables being K and J. The adopted formulation ensures numerical stability in the estimation of permeability, avoiding negative or non-physical values. Once the best-fit parameters are obtained, the model predictions are compared against the CFD dataset.

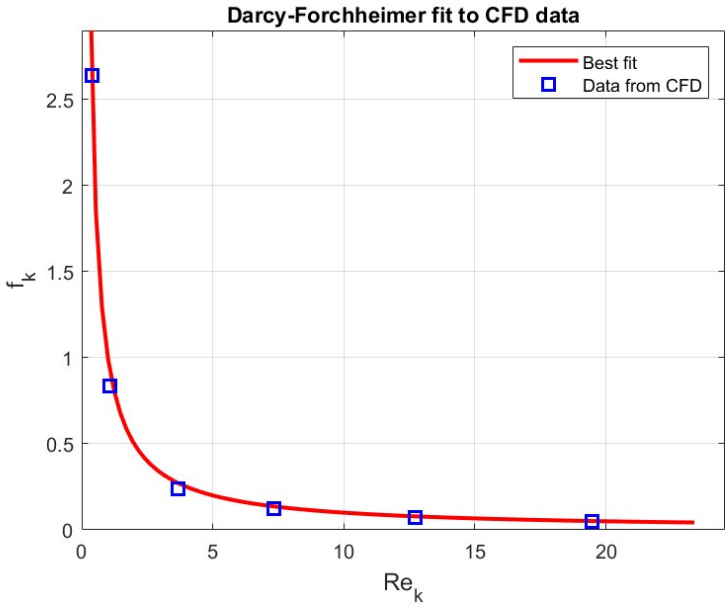

Figure 3.6: Darcy Forchheimer fit to CFD data

Figure 3.6 shows the correlation between f_K and Re_K, including both the simulation data and the Darcy–Forchheimer curve resulting from the regression. The quality of the fit is quantified through the coefficient of determination R^2, which confirms a satisfactory agreement between the model and the CFD results. The final parameters identified for the porous medium are:

Parameter	**Value**
Permeability K	$1.334 \cdot 10^{-4}$ m^2
Inertial coefficient J	$1.06 \cdot 10^{-13} \approx 0$
Coefficient of determination R^2	0.9979

Table 3.4: Final results of the nonlinear fitting for the Darcy–Forchheimer model.

These values are later implemented into the porous medium model for the remaining CFD analysis, allowing for a significant reduction in computational cost while preserving the correct thermo-hydraulic behavior of the system.

3.4 Two-dimensional model

Once the key parameters of the porous medium have been determined, a simplified two-dimensional model can be constructed. In this case, the geometry is even simpler than the full rod bundle configuration, so a rectangular computational domain is adopted. To further reduce computational cost without compromising physical accuracy, the axisymmetric boundary condition is applied. This option imposes a symmetry constraint around the central axis, assuming that physical quantities remain invariant with respect to the longitudinal coordinate. As a result, the original three-dimensional problem is reduced to a two-dimensional domain while still capturing both radial and axial variations of the flow. The boundary conditions are defined in close analogy with those used for the 3D rod bundle model. At the inlet, the dataset reported in Table 3.1 is prescribed. The external wall is set as adiabatic with no-slip condition, while the remaining boundaries correspond to the axisymmetric constraint and the outlet, respectively.

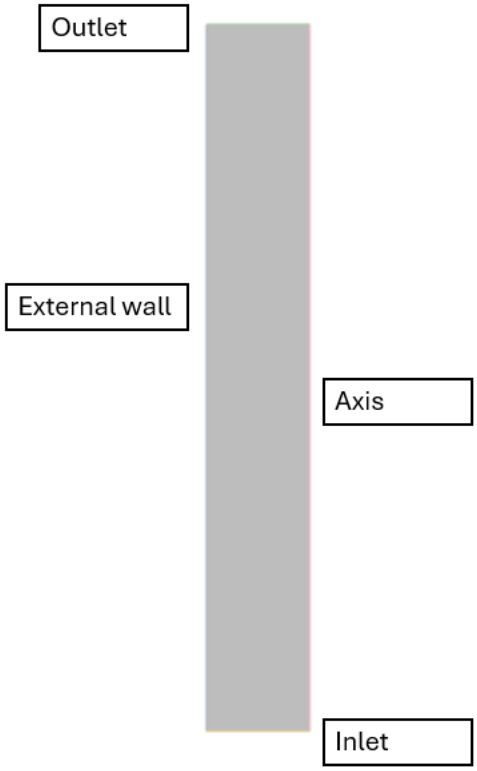

Figure 3.7: Two-dimensional model domain

In the rod bundle model, the region is defined as a fluid domain, with the molten salt thermal properties listed in Table 2.1. In the porous medium model, however, the region is instead defined as a porous domain, which makes it possible to directly include the properties computed in the previous section. STAR-CCM+ provides dedicated options for porous media modeling, allowing the user to specify the following parameters: ϕ, inertial resistance $[kg/m^4]$, viscous resistance $[kg/m^3\,s]$, tortuosity (τ), as well as the solid density, specific heat capacity, and thermal conductivity.

Regarding porosity, the value is calculated based on the geometry data reported in [9]. It is defined as the ratio between the free flow volume and the total heating channel volume. The resulting value is $\phi = 0.944$. The porous resistance coefficients are represented as tensors, accounting for the contribution of viscous and inertial forces to the overall pressure drop. In this study, the isotropic tensor formulation is adopted, and the coefficients are implemented according to the governing equations described in [10], supported by dimensional analysis to ensure consistency with the expected units. Moreover, tortuosity is set to its default value of 1, since the flow path between the inlet and outlet is straight and does not

involve deviations that would increase the distance traveled by the fluid particles. In addition, for the solid porous component, standard material properties at the operating temperatures of the molten salt, as reported in Table 2.2, are considered. The physics models used in this simplified domain are the same used for the rods bundle region and explained in section 3.2, with the only exception being the larger time-step considered thanks to the lower computational cost, accounting for a higher degree of accuracy.

Finally, it is important to state that, since the domain is considered to be axisymmetric, all extensive quantities applied at the boundaries, such as the mass flow rate and the heat input, must be rescaled with respect to the azimuthal angle. Specifically, since the complete three-dimensional geometry corresponds to a rotation of 2π radians, the boundary conditions in the axisymmetric domain are imposed by dividing the original values of these quantities by 2π. This ensures that the total flow rate and energy input are conserved once the domain is revolved into the full three-dimensional space.

3.4.1 Mesh grid independence study

The mesh grid independence study for this domain is carried out with the same procedure as in Subsection 3.2.1, with an advanced customized setting control of the prism layer mesher at the external wall boundary. The discretization of the boundary layer along the external wall of the channel, where both hydrodynamic and thermal gradients are most pronounced is needed to properly capture these effects. The prism layer enables accurate resolution of the viscous sub-layer and the near-wall thermal behavior, which is critical for predicting the heat transfer between the steel wall and the molten salt.

The prism layer specifications, based on empirical best practice in CFD meshing rather than on a strictly defined physical requirement, are summarized in the table below:

Parameter	**Value**
Number of prism layers	5
Prism layer stretching	1.5
Total thickness (relative to base)	33.33%

Table 3.5: Prism layer mesh settings applied at the external wall boundary.

The finest, finer, fine, coarse and coarser meshes settings are compared in the table below, considering the pressure drop and the computational time with a mass flow rate of 0.1 kg/s and 460.15°C as the main evaluation parameters.

	Finest	**Finer**	**Fine**	**Coarse**	**Coarser**
Cell size [m]	0,001	0,001	0,001	0,002	0,01
Surface growth rate	1.01	1.05	Slow	Slow	Slow
Number of cells	230 346	31 359	25 698	9 993	1 558
Computational time [s]	20 664	2 980	1 762	1 092	169
ΔP computed [Pa]	0,160	0,158	0,158	0,157	0,156
Relative error [%]	-	1,26%	1,40%	1,59%	2,51%

Table 3.6: Results of Mesh Refinement Study

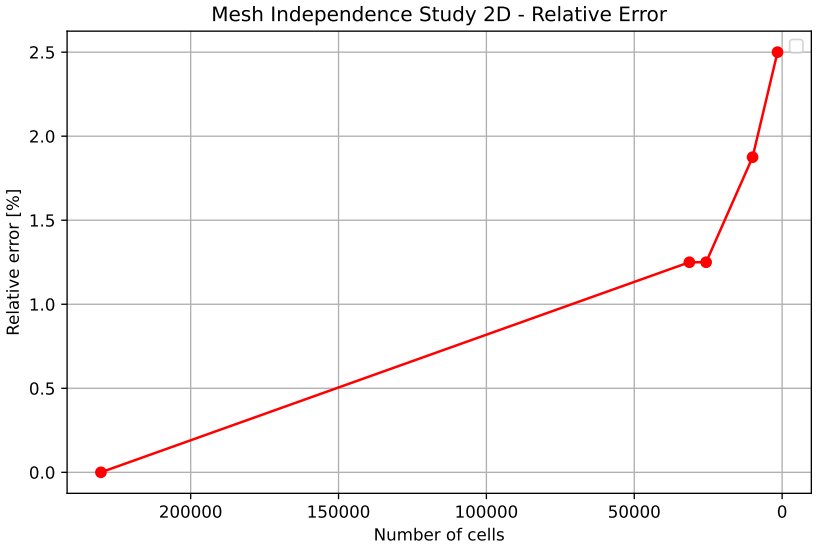

Figure 3.8: 2D Mesh independence study - Relative error plot

As it can be appreciated from the Table 3.6, the fine mesh is the right compromise between accuracy and computational time, according to what it is been stated in Subsection 3.2.1. The fine mesh is the considered one in the following section.

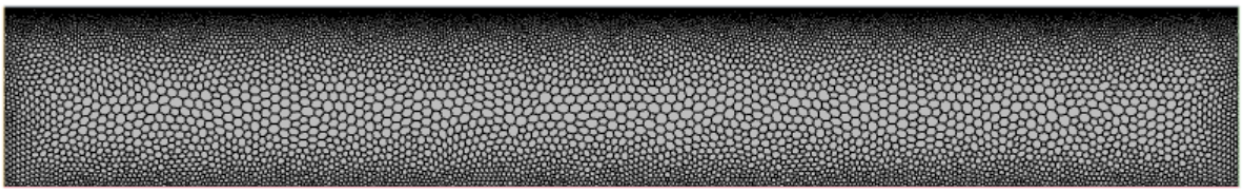

Figure 3.9: Porous media fine mesh

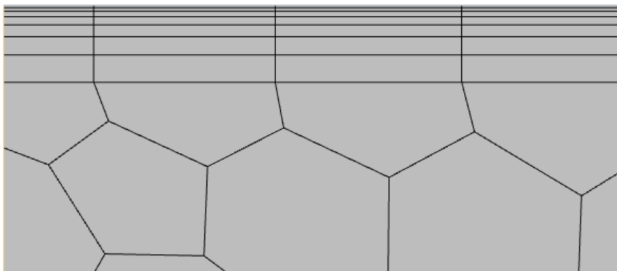

Figure 3.10: Detail of the porous media mesh

3.5 Validation of the porous media simplification

The validation of the porous media model confirms that the adopted geometric simplification does not compromise the overall thermo-hydraulic accuracy of the system. The crucial comparison is established between the pressure drop calculated by the high-fidelity reference model (explicitly resolving the rod bundle geometry) and the prediction from the equivalent porous media model using the derived Darcy–Forchheimer coefficients (K and J).

		RODS 3D	**POROUS MEDIA 2D**
m_{in} [kg/s]	T_{in} [K]	Pressure drop [Pa]	Pressure drop [Pa]
0.10	460.15	0.142	0.158
0.20	480.15	0.181	0.225
0.50	500.15	0.327	0.414
0.75	520.15	0.379	0.471
1.00	540.15	0.404	0.487
1.20	560.15	0.397	0.462

Table 3.7: Rods bundle VS porous media approximation pressure drop

The results, as presented in the table above, demonstrate a strong agreement between the two modeling approaches. This analysis shows that both the high-fidelity model and the porous media approximation exhibit a similar, non-linear dependency of the pressure drop on the mass flow rate and temperature inlet. This consistent trend alignment verifies that the derived K and J coefficients successfully capture the interplay between the viscous resistance, the Darcy term, dominant at lower velocities, and the inertial resistance, the Forchheimer term, significant at higher velocities. Furthermore, the magnitude of the pressure drop is accurately

replicated across the entire range of operational flow rates. The calculated ΔP is limited, as is expected for this internal recirculation system, but the near-perfect match in value between the complex geometry and the simplified model ensures that the overall hydraulic impedance of the heating section is preserved.

For intuitive visual interpretation, the data is plotted in the figure below. The graphical superposition clearly illustrates that the porous media model's curve closely tracks the data extracted from the 3D geometry. This high degree of correlation confirms the robustness of the simplified model and validates that the determined Permeability and Inertial Coefficient accurately model the hydraulic losses for the large-scale transient simulations of the thermocline storage tank.

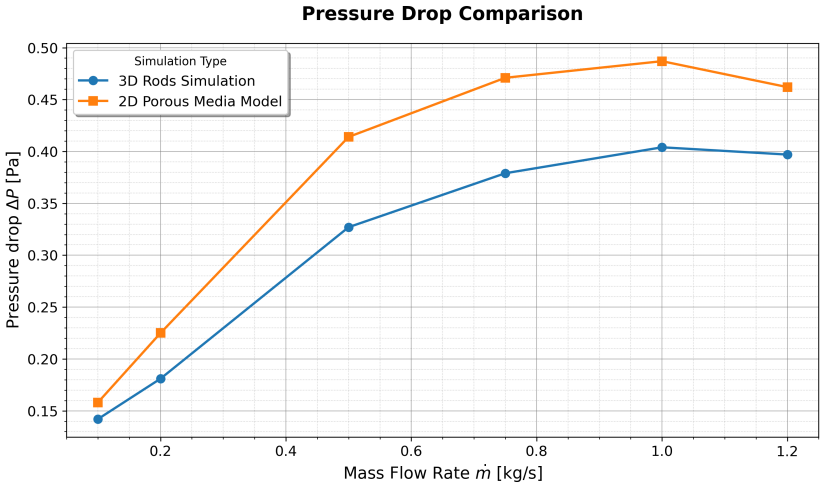

Figure 3.11: Pressure drop comparison between porous media and 3D rods bundle region

Moreover, as anticipated in the previous chapter, an additional validation is carried out. In this case, the comparison involves the theoretical outlet temperature and the numerically computed temperature obtained from the axisymmetric approximation. The objective is to verify whether the solution is consistent not only from a fluid-dynamic standpoint but also from a thermal perspective. To enable this comparison, the theoretical outlet temperature must first be determined for the given inlet heat flux and temperature.

The inlet enthalpy is evaluated using the correlation reported in [9]:

$$h_{in} = 1634(T_{in} - 273.15) - \frac{0.33}{2}(T_{in} - 273.15)^2 \tag{3.6}$$

Subsequently, the outlet enthalpy is obtained from the energy balance:

$$h_{out} = \frac{\dot{Q}}{\dot{m}} + h_{in} \tag{3.7}$$

Finally, the theoretical outlet temperature is derived by inverting Equation 3.6:

$$T_{out} = 273.15 + \frac{1634 \pm \sqrt{1634^2 - 0.66 \cdot h_{out}}}{0.33} \tag{3.8}$$

Since two mathematical solutions are possible, the positive branch is selected, as it represents the physically meaningful result. The computed outlet temperature is obtained from STAR-CCM+ as the surface-averaged value of the temperature scalar field at the outlet, while the heat is introduced as a total heat source of 36.7 kW, as the computational domain represents a single rods bundle channel, divided by 2π, due to the axisymmetric constrain. The results of the comparison are presented and discussed below.

$\dot{m}_{\text{in}}$ [kg/s]	T_{in} [K]	$T_{\text{out, theoretical}}$ [K]	$T_{\text{out, computed}}$ [K]	Error %
0,10	460,15	699,38	714,59	2,13%
0,20	480,15	598,74	606,16	1,23%
0,50	500,15	547,43	550,36	0,53%
0,75	520,15	551,75	553,71	0,35%
1,00	540,15	563,93	565,41	0,26%
1,20	560,15	580,04	581,29	0,21%

Table 3.8: Comparison of Outlet Temperature and Error Analysis

The results presented in Table 3.8 reveal a small relative error between the theoretical and the computed outlet temperatures. The deviation consistently remains around or below 2%, even under the most demanding flow conditions. Such a remarkable agreement attests to the robustness and predictive accuracy of the porous-media simplification.

This outcome confirms that the model, not only reproduces the hydraulic behavior with high fidelity, but also captures the thermal response of the system with impressive precision. Consequently, the porous-media approximation proves to be a reliable, efficient, and physically consistent tool capable of replacing the fully resolved geometry without compromising the overall thermo-hydraulic integrity of the analysis.

Its ability to replicate both pressure-drop and temperature-field trends demonstrates that the simplification preserves the essential physical phenomena while

drastically reducing computational complexity. A result that validates the approach as an excellent and scalable modeling strategy for large-scale transient simulations. The visual results can be appreciated from the figure below.

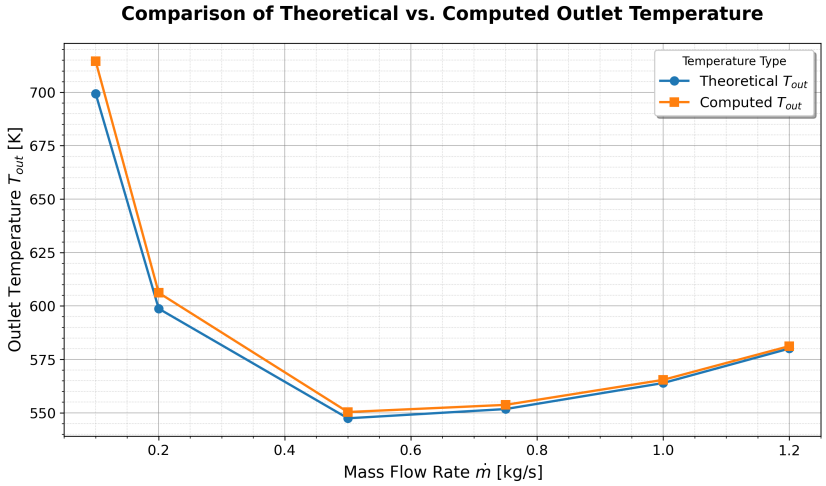

Figure 3.12: Temperature outlet comparison between theoretical and experimental data

Chapter 4

Thermal Storage Simulation

4.1 Computational domain

At this point, the full thermal storage can be simulated. The computational domain, as already briefly mentioned in the Section 2.3, it is computed of a minimal one, aimed at drastically reducing computational costs while maintaining the fidelity of the transient solution. This is made feasible by the strong rotational symmetry of the physical configuration, which characterizes both the cylindrical main tank and the internal recirculating channel. By exploiting this symmetry, the model is reduced to the smallest non-divisible segment of the geometry, consisting of a 60° sector (one-sixth) of the entire physical volume. The radial faces of this sector are prescribed as symmetry planes, allowing the solver to reconstruct the complete flow and thermal fields across the full geometry based on the solution in the reduced domain. It is important to note, however, that the boundary conditions are considerably more complex than this simple geometric reduction might suggest. In particular, they can be clearly distinguished between those applied to the tank region, which govern the external containment and free surface behavior, and those associated with the channel section, where internal recirculation, inflow and outflow conditions are imposed. This layered approach to boundary definition ensures that, despite the reduced domain, all the relevant physical phenomena are correctly represented in the simulation.

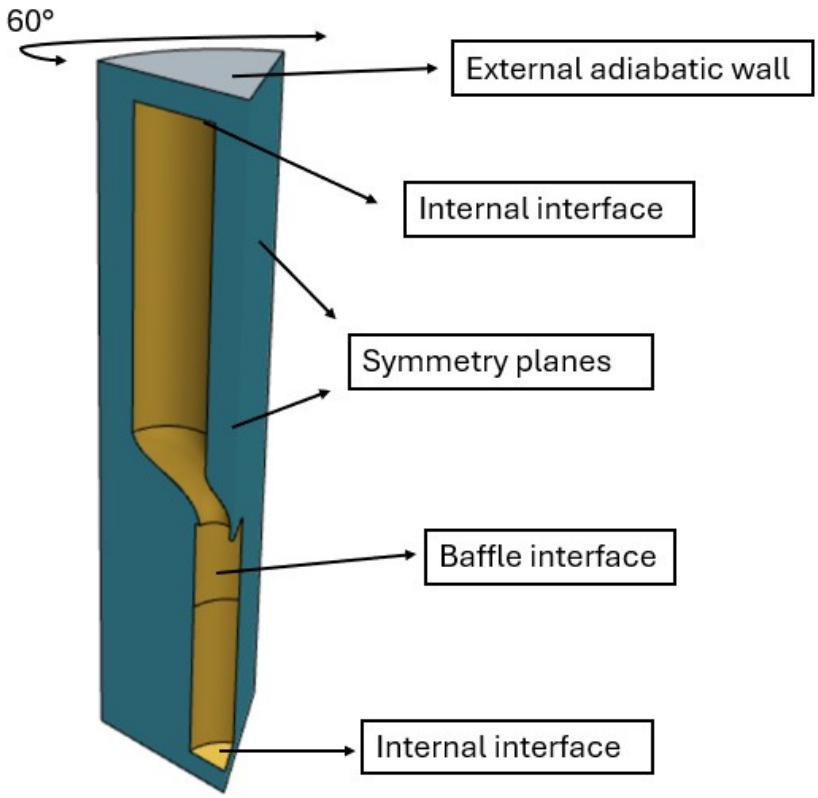

Figure 4.1: Tank boundary conditions

The figure above illustrates the modeled tank region. The grey surfaces correspond to the external adiabatic walls. On the upper boundary, a slip wall condition is imposed to represent the fluid free surface, allowing the tangential velocity to be non-zero while enforcing zero normal velocity and shear stress. The yellow surfaces define interface boundary conditions, which correspond to the areas where two distinct regions are in contact. Two types of interfaces are employed. The internal interfaces, located at the top and bottom of the tank, connect adjacent regions belonging to the same continuum. This ensures that the flow remains unaffected by the change in region, enabling smooth mass, momentum, and energy transfer across the interface. The second type, the baffle interfaces, are used to model the presence of thin, impermeable, and non-thermally conductive plates within the fluid domain. Baffles have infinitesimal thickness and are treated by the solver as wall-like surfaces. A no-slip boundary condition is applied in viscous flow, while the normal velocity component is set to zero, effectively preventing mass transfer across the baffle. Finally, the blue surfaces correspond to symmetry plane boundary conditions. These planes represent imaginary planes of symmetry within the simulation domain, effectively reducing the computational domain while preserving the physical accuracy of the flow field.

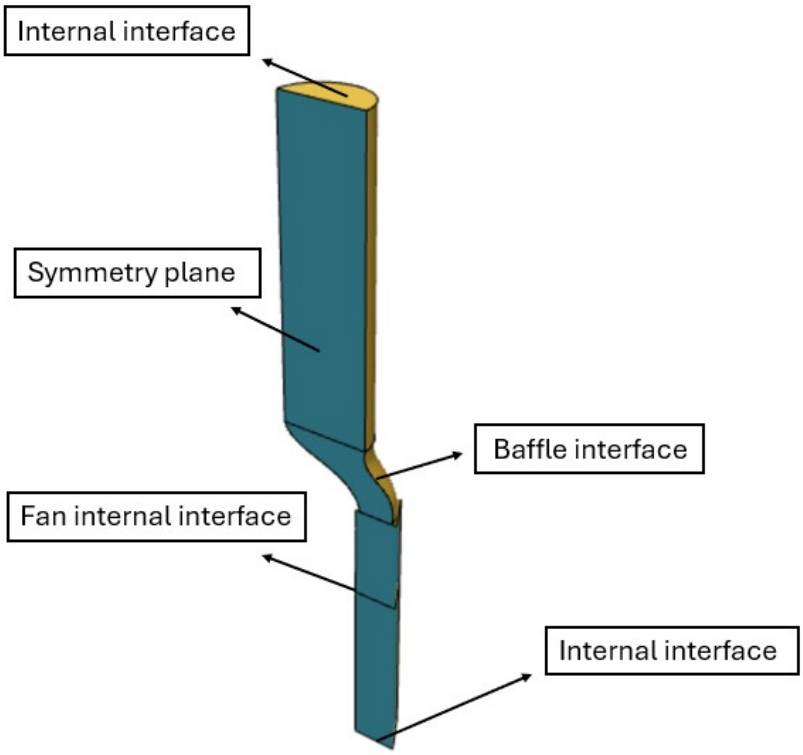

Figure 4.2: Internal channel boundary conditions

Complementarily, Figure 4.2 illustrates the boundary condition configuration of the inner channel. In this case, the computational domain is subdivided into three distinct regions, allowing for a more detailed representation of the complex geometry and improving control over the local boundary specifications.

Starting from the upper section, the porous media region is defined. This region includes an internal interface that allows the fluid to exit the channel and enter the tank domain, ensuring continuity of the flow field. A baffle interface is employed to represent thin impermeable plates embedded within the flow path, while another internal interface links this region to the one immediately below. Additionally, a symmetry plane is imposed to account for the azimuthal periodicity and enable the solver to reconstruct the fully revolved geometry while reducing computational cost.

The duct region follows, characterized by a similar structure: a baffle interface and symmetry plane are applied consistently with the adjacent regions, while internal interfaces ensure proper coupling with the neighboring domains. This configuration guarantees fluid dynamic and thermal continuity across the channel structure.

Finally, the lower channel region is defined in an analogous manner, with one notable difference. Here, an internal interface, referred to in Figure 4.2 as the fan internal interface, is introduced. In the baseline configuration, this interface behaves as a standard internal interface. However, in the second phase of the study, it is redefined as a fan interface, enabling the implementation of a simplified fan model. This approach treats the fan as a zero-thickness interface that imposes a constant pressure jump across the boundary, effectively replicating the fan's impact on the flow field without explicitly modeling its geometry. The fan interface models available in STAR-CCM+ are typically employed in electronic cooling simulations involving incompressible flows, but they are equally well-suited for low-velocity systems, such as those analyzed in the present work. This modeling choice offers a computationally efficient and physically consistent way to represent forced convection within the system.

4.2 Physics models

The simulation employs an implicit unsteady solver with a time-step of 5 seconds, ensuring numerical stability while accurately capturing the transient thermal evolution of the storage tank. The total simulated physical time corresponds to 3600 seconds (1 hour), which is sufficient to analyze the dynamic response of the system under the imposed thermal and hydraulic boundary conditions. A total of 50 inner iterations per time-step are performed, ensuring the convergence of both momentum and energy equations within each temporal increment. From a fluid-dynamic perspective, a laminar flow model is selected. This choice is justified by the characteristic flow regime of molten salts in thermal energy storage systems, where the Reynolds number typically remains below the critical threshold for turbulence due to the relatively high fluid viscosity and moderate flow velocities, as already explained in Section 3.2. According to documentation in [10], the laminar model provides an accurate and computationally efficient description of such flows, avoiding unnecessary turbulence modeling complexity while preserving the fidelity of viscous and conductive effects. For the energy equation, a segregated fluid temperature model is adopted. This formulation decouples the solution of the energy equation from the momentum equations, which is particularly advantageous for incompressible and thermally dominated flows such as those involving molten salts. It allows the solver to handle large temperature gradients with improved numerical stability and reduced computational cost compared to the coupled approach, while maintaining an accurate prediction of heat transfer and buoyancy-driven effects. Additionally, the gravity model is activated to capture natural convection phenomena that develop within the tank. Given the strong dependence of density on temperature in molten salts, buoyancy plays a significant role in the establishment of thermal stratification and in the overall heat transfer dynamics of the system. Enabling gravity ensures that these effects are correctly resolved, allow-

ing the simulation to reproduce the physical behavior observed in realistic storage conditions.

In summary, the selected physical models of laminar flow, segregated temperature, implicit unsteady formulation, and gravity inclusion constitute a coherent and physically consistent modeling framework. Together, they provide a balance between numerical efficiency and physical accuracy, ensuring that the simulation captures the essential thermo-hydraulic phenomena governing the molten-salt thermal energy storage system.

Moreover, as previously discussed, all the computational domains are defined as fluid regions characterized by the thermo-physical properties of HITEC XL molten salt, as specified in Section 2.2. The only exception is the porous media region, which, according to the considerations developed in Chapter 3, is implemented using the porous media model with the Darcy–Forchheimer parameters and specifications detailed in Section 3.4 and a volumetric heat source is imposed within the channel region, representing the primary driving mechanism of the flow circulation inside the tank. This energy input induces a temperature rise in the fluid, which generates buoyancy forces that promote upward motion and enable the establishment of thermal stratification within the storage volume. Specifically, considering a total heat input of 110 kW for the entire prototype system, a volumetric heat source of 18.33 kW, corresponding to one-sixth of the total thermal power, is applied in the simulated domain.

4.3 Mesh grid independence study

To conduct the final simulation, a mesh grid independence study is essential, as was performed in the previous chapters. The specifications for the prism layers are matched to those outlined in Section 3.4.1 and applied consistently at all baffle interfaces separating the tank region from the inner recirculation channel. As with prior simulations, the rationale for undertaking a mesh independence study is to verify that the mesh design is sufficiently refined to accurately reproduce the fluid's physical behavior. This is achieved by running simulations with varying levels of mesh resolution and cell counts, enabling the identification of the optimal mesh that serves as the baseline for subsequent analysis. The five mesh configurations considered for this assessment are presented below.

	Finest	**Finer**	**Fine**	**Coarse**	**Coarser**
Cell size [m]	0.005	0.006	0.007	0.01	0.03
Surface growth rate	Slow	Slow	Slow	Slow	Default
Volume growth rate	1.2	1.2	1.2	1.5	1.7
Number of cells	1 364 037	967 890	712 018	300 619	45 682
Computational time [s]	279 261	199 726	139 108	79 902	12 217
Number of prism layers	5	5	5	5	3
Prism layer stretching ratio	1.5	1.5	1.5	1.5	1.5
$\dot{m}_{out}$ [kg/s]	0.172	0.181	0.241	0.803	2.386
T_{bottom} [K]	487.41	488.10	488.33	496.00	498.29
T_{top} [K]	523.48	523.33	519.44	508.45	502.08
ΔT	36.07	35.23	31.11	12.45	3.79
Relative error [%]	-	2%	14%	65%	89%

Table 4.1: Results of mesh refinement study

The mesh independence study, while essential for ensuring solution reliability, presents considerable computational challenges when conducted on standard personal computing resources. The core difficulty lies in the trade-off between achieving high numerical accuracy through fine mesh resolution and maintaining practical feasibility in terms of computation time. This investigation demonstrates that a thoughtful compromise between these two competing demands can be reached through systematic analysis of the cost-benefit relationship.

The finest mesh configuration, comprising 1.3 millions cells, demands a computational time of 279261 seconds, which translates to approximately 78 hours of continuous processing. This extended duration makes it impractical for conducting systematic parametric studies or sensitivity analyses. However, the finer mesh configuration, containing 967890 cells, achieves a relative error with respect to T_{top} of only 2% compared to the finest mesh while simultaneously reducing the computational time by nearly 45%. This substantial time savings corresponds to approximately 41.9 hours per simulation, making it feasible to conduct multiple simulation runs within reasonable project timelines.

The selection of the finer mesh as the baseline configuration represents a correct engineering decision that balances computational practicality with solution accuracy. The numerical evidence from Table 4.1 supports this choice comprehensively. The temperature predictions at the tank bottom show a difference of only 0.69 K between the finer and finest meshes and similarly, the porous region temperature exhibits a difference of 0.74 K, representing negligible divergence from the finest mesh prediction. These minimal differences indicate that the solution has achieved

quasi-convergence, where further mesh refinement yields diminishing returns in terms of improved accuracy.

4.4 Results in natural convection

This chapter presents and discusses the main outcomes obtained from the numerical simulations carried out in this work using the STARCCM+ software. The primary objective is to evaluate the thermo-fluid dynamic behavior of the thermocline TES system under various operating conditions, considering firstly the natural convection scenario and later the forced one.

According to the previous chapter, the mesh used in the simulation is composed of 967890 cells and some fundamental details are shown below.

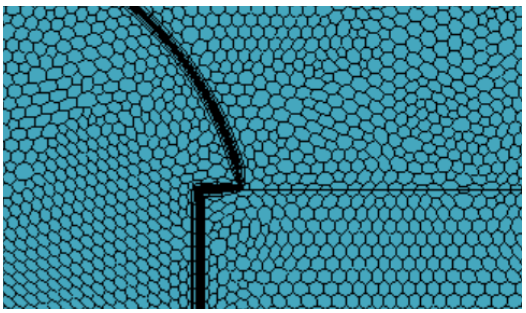

Figure 4.3: Mesh detail

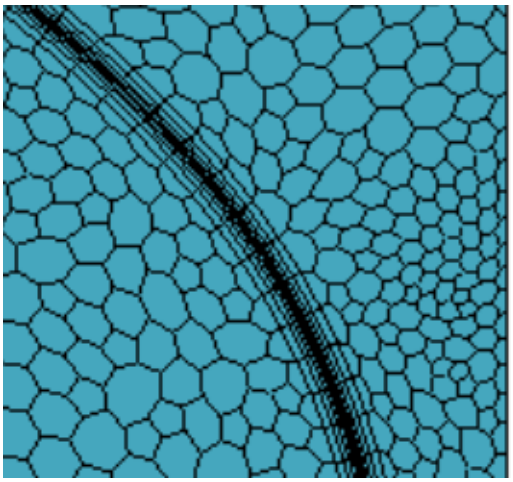

Figure 4.4: Mesh detail

The natural convection case serves as the reference scenario against which all subsequent modifications will be evaluated. This baseline simulation begins with the tank entirely filled with molten salt at a uniform temperature of 463 K, with heat input supplied solely through the internal electrical heating channel. Under these conditions, the system dynamics are driven exclusively by buoyancy forces arising from density variations induced by thermal gradients. The objective of this section is to characterize the fundamental mechanisms of thermocline formation and thermal stratification development in the absence of mechanical forcing, thereby establishing the baseline performance metrics against which the impeller-assisted configuration will be compared. Specifically, we aim to quantify three critical parameters: the thermal energy stored over time, the exergetic ratio of the stratification and the Richardson number describing the stability of the thermal layers, all of which will be evaluated throughout this section.

4.4.1 Temperature evolution

Firstly it is important to analyze some key parameters of the system, starting with the temperature at different time-steps and thermal zones.

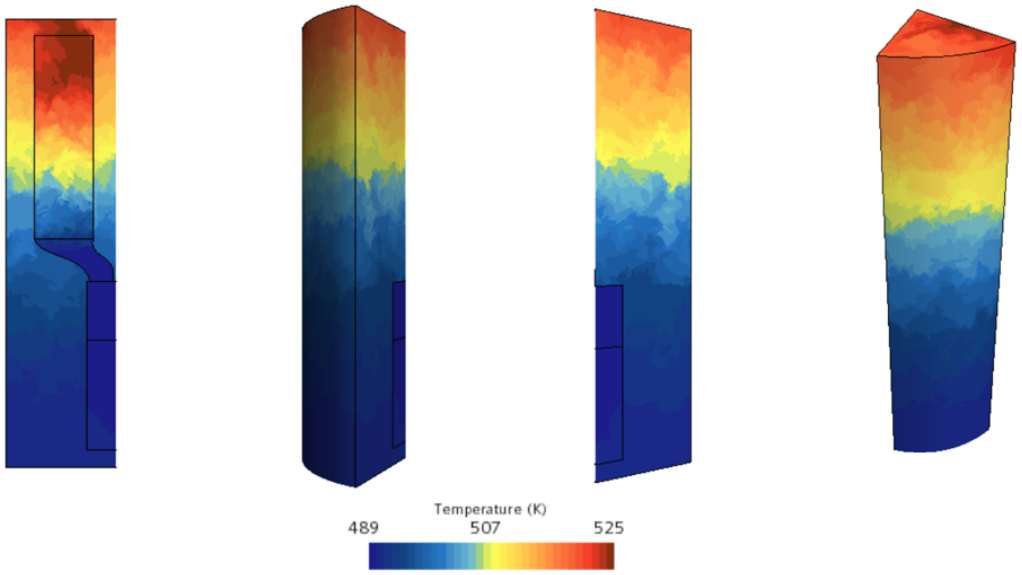

Figure 4.5: Natural convection temperature scene

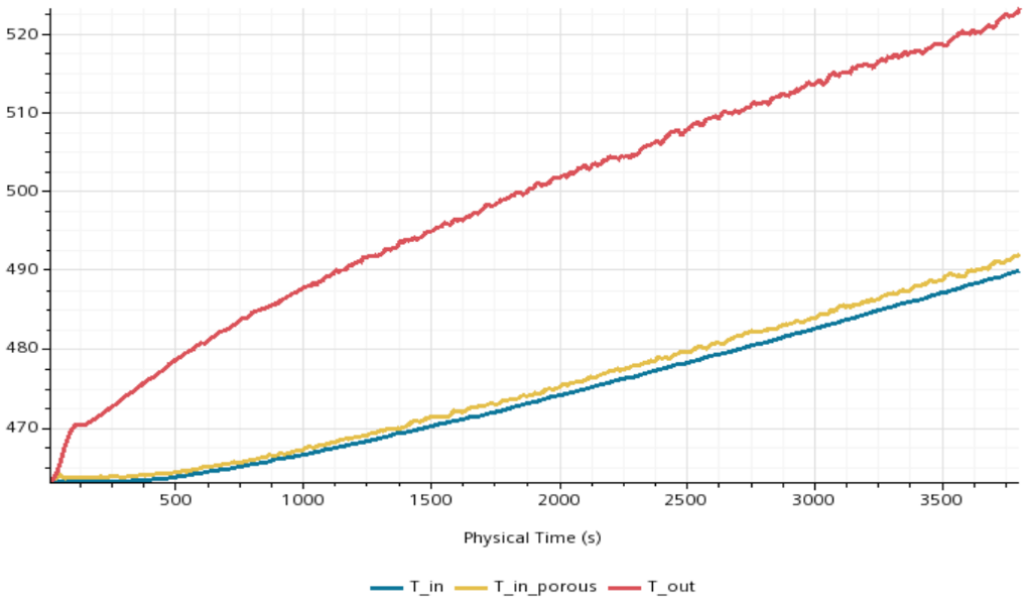

Figure 4.6: Temperature VS time

The thermal contour visualization in Figure 4.5 and temporal temperature evolution plot in Figure 4.6 presented above clearly illustrate the behavior and development of thermal stratification in the system. The temperature contour map reveals a monotonic increase in temperature with increasing height throughout the domain, which directly reflects the action of buoyancy forces. The heated molten salt, exhibiting lower density than its cooler surroundings, undergoes upward transport from the internal heating channel, establishing a vertical temperature gradient. This density-driven separation creates a distinct layering pattern in which hot fluid occupies the upper regions while cold fluid remains confined to the lower portions of the tank, making the stratification visible in temperature scene figure.

The temporal behavior captured in the Fig. 4.6 demonstrates more insights into the stratification development.

For the first approximately 500 seconds of physical time, the temperature at the tank bottom exhibits minimal change, remaining near the initial condition of 463 K. This stasis indicates that the cold fluid in the lower region has not yet been subjected to sufficient thermal forcing to drive convective transport upward. Conversely, the temperature at the outlet of the porous heating medium rises sharply during this same period, increasing from 463 K to approximately 480 K. This rapid increase reflects the direct influence of the internal heating elements within the porous region and demonstrates that heat is being generated faster than it is being transported into the main tank volume.

Approximately 250 seconds into the simulation, a distinct change in the upper temperature evolution trend becomes evident. From this inflection point forward, the slope of the temperature rise at the tank outlet transitions to a more gradual, almost linear profile. This evolution reflects the assumption of adiabatic boundary conditions applied to the tank walls mentioned in 3.3, in fact in practical applications, heat losses to the environment would introduce a thermal sink term, leading to a non-linear, asymptotic approach to an equilibrium temperature. However, the adiabatic assumption is conservative for this investigation, as it allows the characterization of the storage system's intrinsic thermal stratification capability without confounding factors from external heat transfer. This transition marks the establishment of a controlled convective regime in which heated fluid is transported upward at a relatively steady rate. This behavior persists through the remainder of the simulation, with the outlet temperature continuing to rise uniformly. Similarly, the temperatures at the tank and porous media bottom exhibit parallel trends. By the end of the one-hour charging period, this linear trend continues with only minor oscillations of the upper temperature line, with oscillation amplitudes remaining below approximately 1 K.

The physical interpretation of these observations is clear: the hot fluid produced by the heating elements remains concentrated in the upper portion of the tank, ensuring effective thermal stratification, while the cold fluid in the lower regions

does not possess sufficient buoyancy to rise and mix with the hot layer. This separation between hot and cold zones is maintained throughout the simulation, confirming that natural convection provides an adequate driving mechanism for sustained stratification over extended charging periods.

4.4.2 Velocity field characteristics

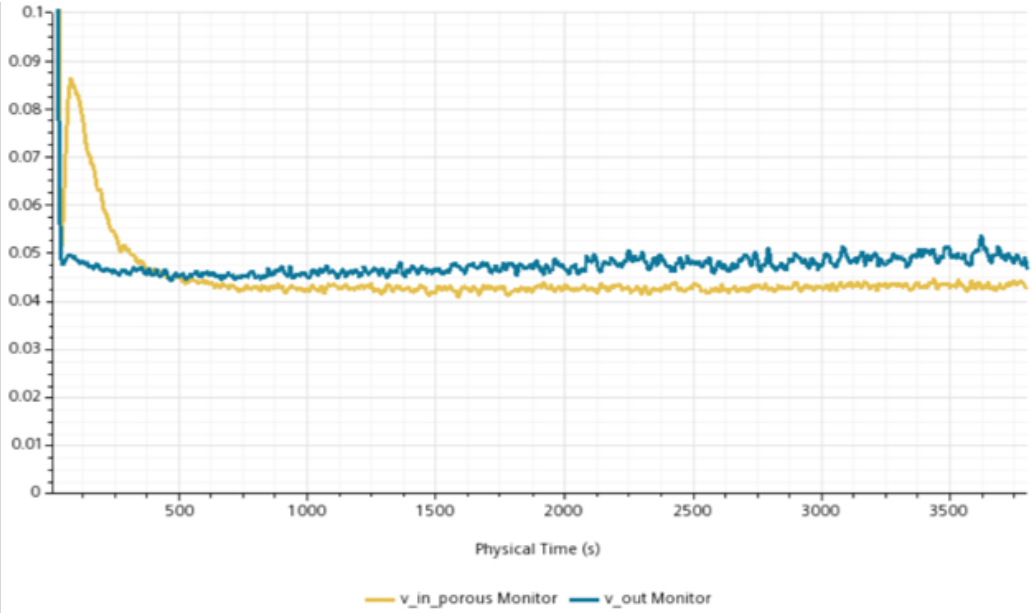

Figure 4.7: Velocity VS time

The velocity profile observed during natural convection operation reveals a significant physical phenomenon: the fluid velocity at the outlet of the porous media region exceeds the inlet velocity. This counterintuitive observation is not simply a result of thermal energy gain, but rather reflects a fundamental thermo fluid-dynamic coupling to natural convection-driven flows in thermally stratified systems.

The apparent velocity increase originates from the interaction between thermal expansion, buoyancy-driven acceleration, and mass continuity constraints within the confined geometry. When the molten salt flows through the heating rod bundle, it undergoes significant temperature increase. However, the mechanism driving velocity increase is not a direct transfer of thermal energy to kinetic energy. Instead, it is the consequence of density variation along the flow path. As the fluid is heated, its density decreases. This density reduction accelerates the fluid upward through the porous region, working with the buoyancy forces already present.

From a mass conservation perspective, the density change must be balanced by velocity variation:

$$\dot{m} = \rho A v = constant \tag{4.1}$$

When density decreases within the porous medium due to heating, the velocity must increase to maintain constant mass flow rate. This is the primary driver of the outlet velocity exceeding the inlet velocity observed in the experimental data.

The plot 4.7 clearly demonstrates this phenomenon. At time zero, both inlet and outlet velocities exhibit transient behavior due to the initiation of convection and thermocline formation. After approximately 500 seconds, the system reaches a quasi-steady state. The outlet velocity stabilizes at approximately 0.050 m/s, consistently exceeding the inlet velocity at roughly 0.045 m/s. This separation is maintained throughout the simulation despite minor oscillations due to numerical noise and thermal stratification dynamics. The steady-state offset between inlet and outlet velocities confirms that the velocity increase is a structural feature of natural convection in this geometry, not a transient phenomenon.

4.4.3 Mass flow rate stability

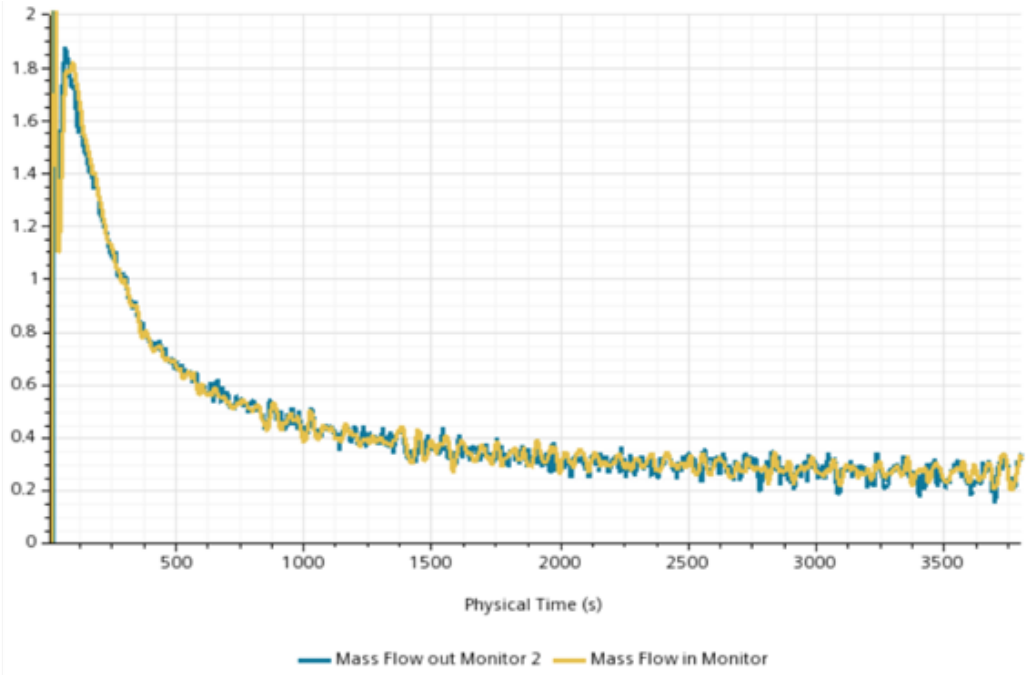

Figure 4.8: Mass flow VS time

The mass flow rate profile during the natural convection charging process reveals two distinct temporal regimes. Unlike the relatively stable velocity profiles observed previously, the mass flow rate exhibits more pronounced oscillatory behavior and a strong decaying trend, providing insight into the global energy balance and the system's response to thermocline formation and vertical stratification development.

The initial transient phase, spanning approximately 0 to 1000 seconds, is characterized by a rapid decrease in mass flow rate from roughly 2.0 kg/s at system initiation to approximately 0.4 kg/s. This sharp decay is not due to obstruction or increased flow resistance, but rather reflects the fundamental thermodynamic adjustment occurring within the storage tank. During this phase, the heating rods are continuously supplying thermal energy to the salt, causing widespread temperature increase throughout the domain. The internal energy of the fluid rises substantially, and the temperature stratification begins to establish itself as the initially isothermal fluid becomes thermally stratified into hotter upper regions and cooler lower regions. The role of energy absorption is critical to understanding this transient behavior. As thermal energy from the heating rods is absorbed by the molten salt, the fluid near the heating elements experiences rapid temperature rise. According to the first law of thermodynamics, this energy absorption increases the internal energy of the fluid, manifesting as temperature increase.

After approximately 1000 seconds, the system transitions into a quasi-steady state where the mass flow rate oscillates around a nearly constant mean value of approximately 0.25-0.35 kg/s. The oscillations evident in the data, are believed to originate from multiple coupled mechanisms: thermal plumes formed at the interface between hot and cold layers which periodically break away and alter the local flow structure; interactions between the internal recirculating channel and the main tank cavity and numerical effects related to the temporal discretizations of the transient equations.

The physical mechanism governing the quasi-steady mass flow rate differs from that of the transient decay. Once the thermocline has formed and the internal energy distribution has stabilized, the mass flow rate is controlled by a balance between the buoyancy-driven acceleration and the viscous dissipation. The energy input from the heating rods continues to warm the salt, but this energy is now primarily stored as internal energy rather than driving bulk acceleration.

The comparison between inlet and outlet mass flow rate measurements shows excellent agreement in both transient and steady-state phases, with negligible divergence between the two signals. This near-perfect overlap validates the mass continuity assumption and indicates that the flow undergoes neither significant accumulation nor depletion within the control volume. The oscillatory component visible in the mass flow rate but less apparent in the velocity measurements suggests that the oscillations result from temporal variations in the effective cross-sectional area

through which flow occurs, possibly due to dynamic changes in the thermocline position or the formation of recirculation eddies.

4.5 Results with impeller driven fluid

Now, the fan interface is activated. The purpose of it is to investigate whether an impeller at the lower channel region forcing the fluid at a constant pressure drop opposed to the natural rising flow direction would improve the performance of the storage system. In order to achieve that, the fan interface is treated as a constant pressure jump source with four different configurations: 10, 50, 60 and 100 Pa. This configuration should allow the stratification to be more effective because of the fundamental interaction between forced convection and natural convection in mixed convection regimes.

Although introducing forced convection against the natural buoyant rise to enhance stratification might appear paradoxical at first sight, it works because of three fundamental thermophysical mechanisms grounded in mixed convection theory.

First, the Richardson Number framework demonstrates that neither pure natural convection nor pure forced convection alone guarantees optimal thermal stratification.

Second, this balance suppresses the formation of large-scale vortices and secondary flow instabilities that characterize uncontrolled natural convection.

Third, by maintaining low flow velocities, the residence time of fluid in the heating zone is maximized without allowing diffusive mixing to degrade the thermocline, a principle consistent with the exergy-based analysis of stratified thermal storage systems.

At very low pressures, the forced convection is insufficiently strong to suppress natural convection instabilities; the buoyancy forces remain dominant, as it is happening in the natural convection case. Conversely, at very high pressures, the forced flow becomes so powerful that it inverts the natural stratification entirely, driving the fluid downward and destroying the thermocline structure that is essential for efficient storage performance.

The critical insight is that an optimal balance exists at intermediate pressures. At this equilibrium point, buoyancy and forced convection are nearly balanced rather than dominated by either mechanism alone. This balance suppresses the chaotic vortex formation characteristic of pure natural convection while simultaneously avoiding the destructive flow inversion induced by excessive forcing. The result is a quasi-static, orderly flow pattern, allowing the thermal gradient to establish and maintain a sharp, stable thermocline.

In the table below some forced convection parameters are compared to natural convection ones in order to state if the forced convection solution does improve the stratification.

	Natural	**Forced 10 Pa**	**Forced 50 Pa**	**Forced 60 Pa**	**Forced 100 Pa**
Mass Flow [kg/s]	0.181	0.267	0.137	0.43	5.77
T_{bottom} [K]	489.9	485.6	484.52	488.67	501.65
T_{top} [K]	523.3	527.1	529.26	525.74	500.37
ΔT [K]	33.4	41.5	44.7	37.1	-1.3
v_{out} [m/s]	0.047	0.048	0.052	0.052	0.119

Table 4.2: Thermal stratification and hydrodynamic characteristics.

The thermal stratification behavior across different operating conditions reveals distinct convective regimes with markedly different characteristics. The natural convection baseline establishes a reference stratification of 33.4 K between the bottom and upper regions, achieved through uncontrolled buoyancy-driven flow at a mass flow rate of 0.181 kg/s. When a low-pressure forced convection is introduced at 10 Pa, the system exhibits an intermediate behavior: the mass flow increases to $0.267 kg/s$ while the thermal stratification improves modestly to 41.5 K. This suggests that the weak forced convection begins to suppress buoyancy-driven mixing but remains insufficient to achieve full control of the flow regime.

The 50 Pa configuration, shown in Fig.4.9 represents a critical transition point where the system attains optimal thermal performance. Despite operating at the lowest mass flow rate among all forced configurations (0.137 kg/s), even lower than the natural convection case, the thermal stratification reaches its maximum value of 44.7 K. This pressure creates an equilibrium in which buoyancy and pressure-driven forces achieve balance, resulting in an orderly, low-velocity flow pattern. The reduced mass flow, combined with enhanced stratification, suggests that the system operates with maximum thermodynamic efficiency at this point: thermal energy is retained within the storage domain longer, allowing sharper temperature gradients to develop and persist.

At 60 Pa, a significant discontinuity occurs. The mass flow significantly increases to 0.43 kg/s, more than tripling from the 50 Pa value, while the thermal stratification simultaneously collapses to 37.1 K, representing a 17% degradation relative to the optimal point. This ehavior indicates a transition from the equilibrium regime toward forced-convection dominance. The change in flow characteristics suggests the system undergoes a regime transition between these two pressure levels, where the balance between buoyancy and forced convection is destabilized.

Finally, at 100 Pa, shown in Fig.4.10, the forced convection becomes overwhelmingly dominant. The mass flow accelerates to 5.77 kg/s while the temperature difference across the porous region inverts to –1.3 K, signifying that the cooler outlet fluid is pushed downward by the fan, completely destroying the natural stratification and reversing the intended flow direction.

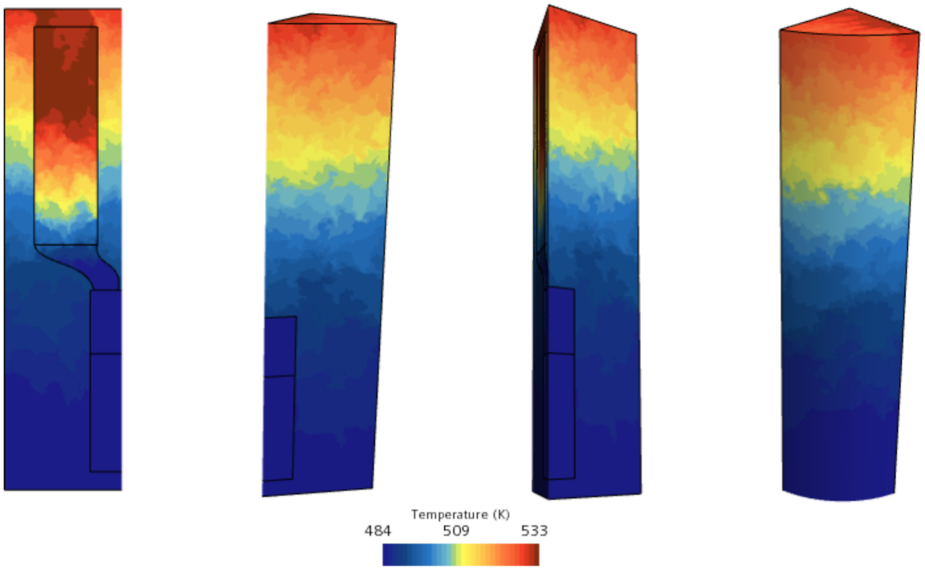

Figure 4.9: Forced convection at 50 Pa Temperature scene

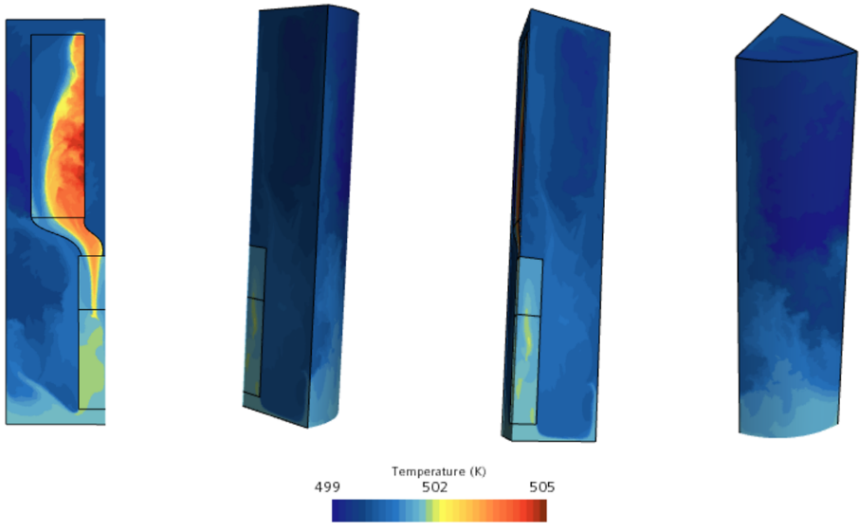

Figure 4.10: Forced convection at 100 Pa Temperature scene

In synthesis, the forced convection at 50 Pa emerges as the superior operating condition. It simultaneously minimizes parasitic mass flow, maximizes thermal stratification through controlled equilibration of competing forces, and establishes a stable, quasi-static flow regime where gravitational potential energy and thermal gradients are preserved rather than dissipated through turbulent mixing. This configuration represents the optimal balance between imposing sufficient convective control to suppress buoyancy-driven chaos while remaining gentle enough to preserve the thermocline structure that is fundamental to efficient thermal storage operation.

4.5.1 Quantitative performance metrics

The statement in the previous section is confirmed also with an insight on the key parameters indicators that can be considered to analyze a TES system: the energy stored; the exergetic ratio and the Richardson number.

	Natural	**Forced 10 Pa**	**Forced 50 Pa**	**Forced 60 Pa**	**Forced 100 Pa**
Energy Stored [MJ]	63.59	64.66	64.96	65.22	66.18
Exergetic Ratio Ξ/Ξ_0 [%]	12.47	18.58	22.23	14.70	0.0083
Richardson Number	162.58	14.69	7.92	17.64	-0.19

Table 4.3: Key performance indicators for thermal stratification across operating conditions

The natural convection configuration establishes the performance baseline for thermocline stratification in the absence of mechanical circulation. Operating at Ri = 162.58, the system is deeply buoyancy-dominated, achieving the poorest thermal performance: an energy storage of 63.59 MJ and exergetic ratio of only 12.47%. Rather than organizing the thermal field into sharp layers, natural convection creates chaotic mixing within the porous heating region, rapidly equilibrating temperature across the tank height and destroying the intended thermocline interface. The low exergetic ratio reflects the physical reality that natural convection, despite requiring zero auxiliary fan power, sacrifices the usable thermal content of the storage medium. The thermocline becomes diffuse, rendering the stored energy less valuable for discharge-phase power generation. This baseline analysis justifies the introduction of an internal circulation impeller: the parasitic fan power is economically justified by the demonstrated recovery of stratification quality.

Introduction of forced convection produces immediate and dramatic improvements across all pressure-drop configurations. The exergetic ratio increases to 18.58% at 10 Pa (Ri = 14.69), rises further to 22.23% at 50 Pa (Ri = 7.92), before gradually declining through 60 Pa (14.70%, Ri = 17.64) and collapsing at 100 Pa (0.0083%,

Ri = −0.19). This non-monotonic trend, culminating at 50 Pa, reveals a hidden physics: the optimal configuration is not simply "more forcing" but rather precise equilibration of competing convective forces.

At 50 Pa, the Richardson Number of 7.92 places the system within the critical mixed-convection window where buoyancy and forced-convection inertia are mechanically balanced rather than hierarchically dominated. This balance yields the maximum exergetic ratio of 22.23%, a 78% improvement relative to natural convection. The system achieves stratification with minimal energy dissipation in the flow field, preserving the temperature gradient across the thermocline interface throughout the 3600 s charging transient.

Increasing pressure to 60 Pa crosses a critical threshold. Although the Richardson Number increases to 17.64, the exergetic ratio paradoxically drops to 14.70%, and the stratification metric ΔT collapses from 44.7 K to 37.1 K. This apparent contradiction reveals that Richardson Number alone does not capture the full physics at its extremes.

The 100 Pa configuration represents a catastrophic departure from thermocline physics. The Richardson Number turns negative, signifying that forced convection inertia has completely overwhelmed any residual buoyancy influence. Temperature inversion develops: the outlet temperature measurement reveals $\Delta T = -1.3K$, meaning cold fluid is forced downward and hot fluid displaced upward, the opposite of the intended stratification. Despite this inversion, the energy storage paradoxically increases to 66.18 MJ, a 4.07% gain relative to natural convection. However, this apparent advantage is illusory: the higher energy reflects only the elevation of global mean tank temperature due to sustained auxiliary heating through the internal resistors, while the exergetic ratio collapses to 0.0083%, indicating that the stratification required for practical discharge-phase energy recovery has been completely annihilated.

The data establish that optimal thermocline storage is neither a natural-convection nor a pure forced-convection problem, but rather a mixed-convection optimization challenge. The 50 Pa configuration uniquely balances three competing objectives: maximum stratification sharpness, highest exergetic retention capacity and minimal parasitic circulation.

4.5.2 Optimal forced convection temporal behavior

The optimal configuration of 50 Pa pressure jump can be analyzed more in depth looking at the temporal evolution of the main characteristics such as temperature, velocity and mass flow rate.

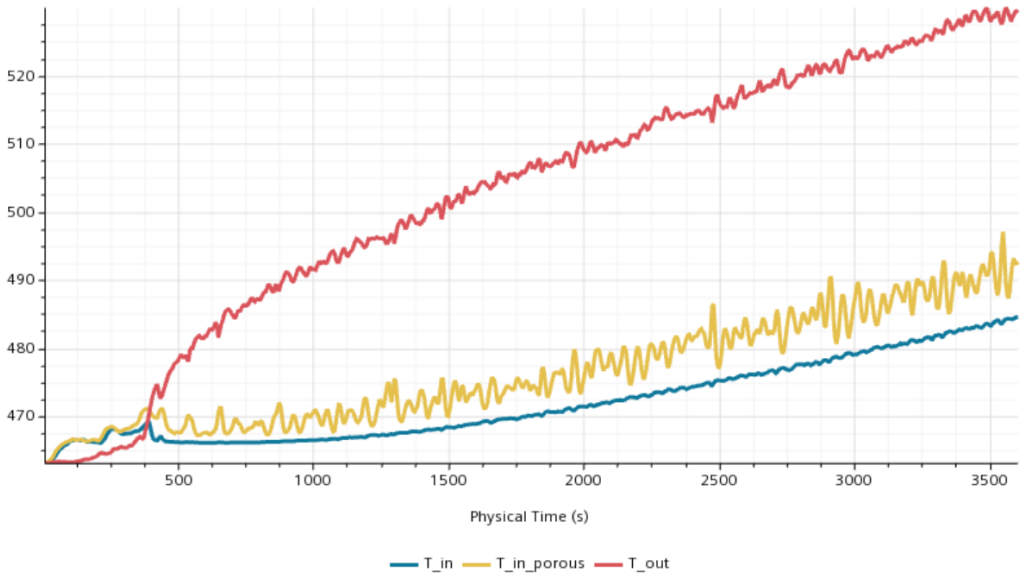

Figure 4.11: Forced convection temperature evolution plot

The transient temperature evolution over the 3600 second charging cycle at 50 Pa configuration reveals the fundamental advantage of operating within the Richardson equilibrium regime.

During the initial phase, between 0 and 500 seconds, the system exhibits rapid thermal transients. The outlet temperature climbs steeply from approximately 470 K to 490 K, driven by the incoming hot fluid and minimal thermal inertia of the storage medium. During this phase, the porous medium temperature lags behind the outlet, suggesting the formation of an emerging thermocline. The inlet temperature remains relatively constant near 470 K (representing the ambient reference), confirming steady thermal boundary conditions upstream. This early overshoot in outlet temperature is characteristic of thermocline systems where the hot fluid initially travels through the tank without significant thermal dispersion.

Between 500 and 2500 seconds, the system transitions into a quasi-linear charging regime where all three temperature profiles evolve more uniformly and predictably. Notably, the porous medium temperature rises steadily without significant oscillations, climbing from 470 K to 490 K over this extended period. This smooth, monotonic behavior is significant because it indicates that the low mass flow rate at 50 Pa allows sufficient residence time for thermal energy absorption without generating the chaotic mixing and secondary flows that would produce oscillatory temperature fluctuations. The separation between the outlet and porous temperature narrows gradually, suggesting the thermocline is progressively compressing as more cold fluid is displaced upward by incoming hot fluid. Critically, the outlet temperature continues to rise steadily without plateauing, evidence that the strat-

ification remains sharp and the thermocline has not yet reached the tank's upper boundary.

Finally a late phase between 2500 and 3600 seconds is present. As the system approaches the end of the charging window, the outlet temperature slows its rise. The porous medium temperature concurrently reaches approximately 490–495 K. This deceleration phase reflects the final stage of thermocline compression: the hot-fluid arrival front is imminent at the outlet plane and the system approaches thermal saturation. Crucially, the lack of temperature oscillations, despite the system operating near its dynamic limit, demonstrates the superior stability granted by the low Richardson number. Were the system operating at higher pressures, where forced convection dominates, significant oscillations would be expected due to flow instability and vortex formation. The smooth trajectories here confirm that the balanced force regime maintains laminar, ordered flow even as thermal stratification intensifies.

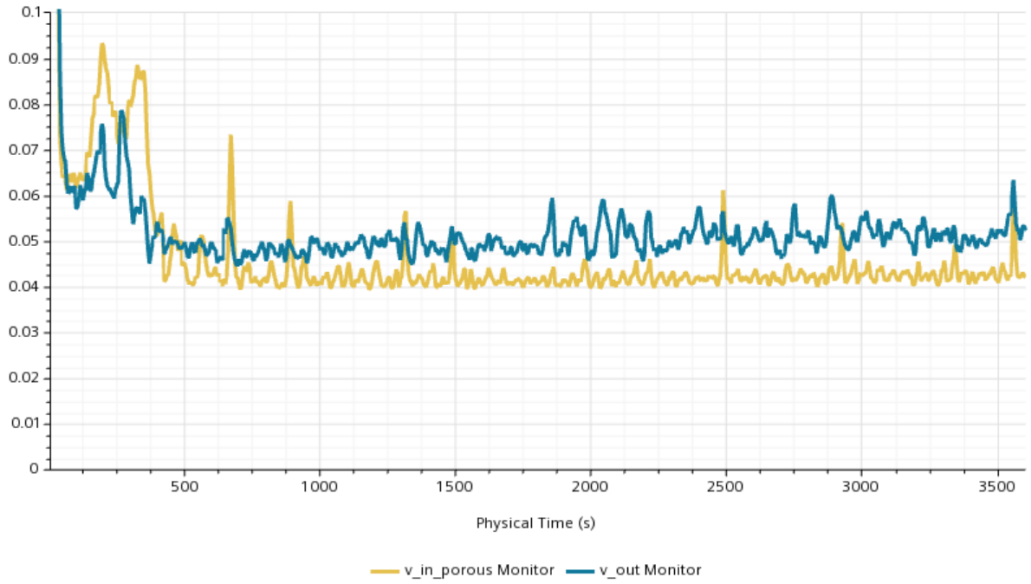

Figure 4.12: Forced convection velocity plot

The figure above shows the velocity evolution trough time. The velocity evolution at 50 Pa exemplifies the hydrodynamic signature of optimal thermocline operation. Initial transients are rapidly damped, the system achieves a quasi-steady, low-velocity state within 15 minutes and sustained stability is maintained thereafter. The inlet and outlet velocities remain closely coupled, indicating spatially uniform flow. Most significantly, the complete absence of velocity spikes, reversals, or long-period oscillations, features that would compromise thermal stratification through mixing, confirms that the Richardson-balanced regime successfully suppresses convective instabilities. Moreover, it is important to notice that, according

to what it has already been explained in the Section 4.4.2, the outlet velocity is constantly greater than the inlet one.

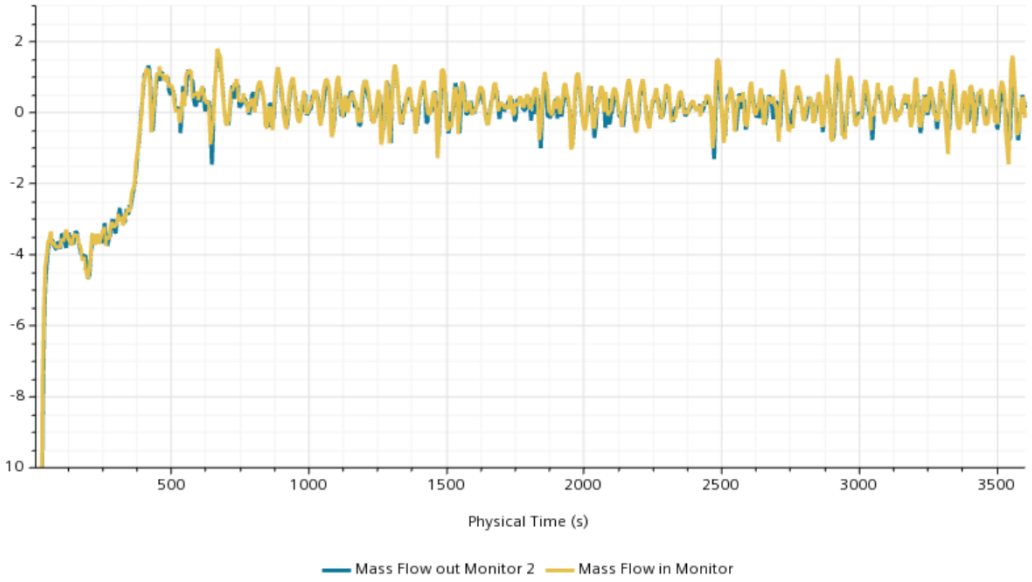

Figure 4.13: Forced convection mass flow rate plot

Finally, a look at the mass flow rate plot reveals a fundamental hydrodynamic phenomenon: the systematic oscillation of instantaneous mass flow between positive and negative directions around a net positive mean value. This behavior reflects the continuous dynamic equilibration between the competing buoyancy and forced convection forces that define the system's equilibrium state. At 50 Pa, the forced convection and buoyancy forces are nearly equal in magnitude but opposite in direction. The imposed low-pressure fan drives fluid upward into the porous region with controlled momentum, while simultaneously, the evolving thermal stratification generates local density inversions that produce counter-directed buoyancy driven flows. The result is not a true flow reversal at the global scale, but rather a continuous micro reversals at the local thermocline front where buoyancy and forcing counterbalance each other. Moreover, these localized flow reversals at the thermocline interface are not catastrophic for stratification quality. Rather, they represent the system's optimal mechanism for dissipating the tension between competing forces without generating large scale mixing or vortex formation. The oscillations occur within the thermal gradient region rather than spreading throughout the bulk fluid.

Chapter 5

Final Considerations

This project has fully achieved its objectives and can be considered a success. The first major contribution is the development of a reliable porous media model, rigorously validated through a three-stage methodology. By starting from a high-fidelity explicit rods simulation and systematically simplifying to a porous media approach, the thesis demonstrates that it is possible to drastically reduce computational requirements without sacrificing accuracy. This makes full-scale system simulations feasible even on limited computational resources.

Secondly, the original concept of modifying natural convection storage with an internal, inverted fan to introduce controlled forced convection has been demonstrated as both practical and effective. The study clearly shows that forced convection, when applied in a carefully tuned manner, can enhance the thermal stratification of the system by 25% in the best case scenario and maximize its usable energy storage capacity. The parametric analysis highlights the fine operational window in which such a system must be operated in order to retain its advantages, providing an important guideline for future system design and control strategies.

Nevertheless, several potential directions for future enhancement of the model and simulation process emerge. First, moving beyond the adiabatic wall assumption by incorporating thermal loss analysis would increase the realism of the predicted temperature profiles, allowing them to reach true asymptotic behavior reflective of real-world systems. This would allow system designers to more accurately forecast operational efficiency and storage duration in practical scenarios.

Furthermore, the long-term behavior of structural components under thermal cycling, particularly steel fatigue, warrants attention. Future work should investigate the effects of repeated charge and discharge cycles on the mechanical integrity of the vessel and internals, potentially coupling thermal and structural simulations.

This expanded analysis would also offer valuable insight into system reliability, maintenance scheduling, and safety.

Additionally, extending the simulation to include the discharge phase would deliver a more comprehensive understanding of the full system dynamics, and allow for the evaluation of strategies to mitigate mixing and stratification losses during energy extraction.

Bibliography

[1] Wan Afin Fadzlin et al. "Solar thermal energy storage: global challenges, innovations, and future directions for renewable energy systems". In: *Applied Thermal Engineering* 2025 (2025), p. 128346.

[2] Roberto Zanino et al. *Ibridizzazione del CSP con altre tecnologie energetiche rinnovabili - Sviluppo di un modello numerico di un sistema di accumulo termoclino ibridizzato con due serpentine e resistenze elettriche immerse nei sali: confronto con il sistema termoclino ibridizzato con tre serpentine e caldaia elettrica*. Tech. rep. Presentazione nell'ambito dell'Accordo di Programma Ministero dell'Ambiente e della Sicurezza Energetica – ENEA, Piano Triennale di Realizzazione 2022-2024. Dec. 2024.

[3] Rafaa Mohammad. *Preliminary Study on Hybrid Thermocline Thermal Energy Storage*. Final Report. Supervisor: Dr. Mattia Cagnoli. ENEA, June 2025.

[4] E. Chiavazzo. *Energy Storage*. Lecture notes, Politecnico di Torino. Corso di laurea magistrale. 2024.

[5] I. Dincer and M.A. Rosen. *Thermal Energy Storage Systems and Applications*. 2nd ed. John Wiley & Sons, 2011.

[6] Zilong Wang et al. "The thermal stratification characteristics affected by a novel equalizer in a dynamic hot water storage tank". In: *Applied Thermal Engineering* 126 (2017), pp. 1006–1016. ISSN: 1359-4311.

[7] Md. Shahneoug Shuvo, Mahmudul Hasan Hasib, and Sumon Saha. "Entropy generation and characteristics of mixed convection in lid-driven trapezoidal tilted enclosure filled with nanofluid". In: *Heliyon* 8.12 (2022), e12079.

[8] Mehdi Shokrnia et al. "Geometrical and PCM optimization of a thermocline energy storage system". In: *Journal of Energy Storage* 98 (2024), p. 113070.

[9] A. Kaood et al. "Thermal, rheological, and corrosive properties of a novel ternary molten salt mixture for concentrated solar power plants". In: *Renewable Energy* 177 (2021), pp. 1045–1062.

[10] Siemens Digital Industries Software. *Simcenter STAR-CCM+ User Guide*. Accessed: 2025-09-30. 2024.

[11] Mattia Cagnoli et al. "Multi-scale modular analysis of open volumetric receivers for central tower CSP systems". In: *Solar Energy* 190 (2019), pp. 195–211.

Appendix A

MATLAB code for the parametric evaluation

The following MATLAB script has been used for the K e J parameters identification according to the Darcy–Forchheimer law.

```
clear; clc;
% INPUT DATI REALI
m_dot = [0.1, 0.2, 0.5, 0.75, 1, 1.2]; %[kg/s]
T = [460.15, 480.15, 500.15, 520.15, 540.15, 560.15]; %[K]
delta_p = [0.142,0.181,0.327,0.379,0.404,0.397]; %[Pa]
Af = 9.08670946e-02;      %[m^2]
d = 1.2;                  % [m]
phi = 0.944453699;
mu = @(T) (10^6.1374*(T-273.15).^-3.36406);  %[Pa s]
rho = @(T) (2240 - 0.827*(T-273.15));   %[kg/m^3]
% FIT NON LINEARE PER K E J
cost_function = @(x) compute_residuals(x, m_dot, T, delta_p,
    Af, d, phi, mu, rho);
x0 = [1e-4, 0.01];  %[sqrt(K), J]
opt = optimoptions('lsqnonlin','Display','iter');
[x_opt, resnorm] = lsqnonlin(cost_function, x0, [1e-8, 0],
    [1, 10], opt);
sqrtK = x_opt(1);
K = sqrtK^2;
J = x_opt(2);
fprintf('\nFinal results:\nK = %.3e m^2\nJ = %.4f\n', K, J);
% CALCOLO f_k E Re_k
n = length(m_dot);
f_k = zeros(n,1);
Re_k = zeros(n,1);
for i = 1:n
    mu_i = mu(T(i));
```

```
    rho_i = rho(T(i));
    Re_k(i) = (m_dot(i) * phi * sqrtK) / (mu_i * Af);
    f_k(i) = (delta_p(i) * rho_i * sqrtK * Af^2) / (d * phi^2
        * m_dot(i)^2);
end
% PLOT CURVA FIT + DATI CFD
Re_fit = linspace(min(Re_k)*0.8, max(Re_k)*1.2, 100);
f_fit = 1 ./ Re_fit + J;
figure;
plot(Re_fit, f_fit, 'r-', 'LineWidth', 2); hold on;
plot(Re_k, f_k, 'bs', 'MarkerSize', 8, 'LineWidth', 1.5);
xlabel('Re_k', 'FontSize', 12);
ylabel('f_k', 'FontSize', 12);
legend('Best fit', 'Data from CFD', 'Location', 'northeast');
grid on;
xlim([0, max(Re_fit)*1.05]);
ylim([0, max(f_k)*1.1]);
title('Darcy-Forchheimer fit to CFD data');
% CALCOLO R^2
f_k_model = 1 ./ Re_k + J;
f_k_mean = mean(f_k);
SS_tot = sum((f_k - f_k_mean).^2);
SS_res = sum((f_k - f_k_model).^2);
R2 = 1 - (SS_res / SS_tot);
fprintf('Coefficiente di determinazione R^2 = %.4f\n', R2);
% FUNZIONE DI ERRORE
function residuals = compute_residuals(x, m_dot, T, delta_p,
   Af, d, phi, mu_fun, rho_fun)
    sqrtK = x(1);
    J = x(2);
    n = length(m_dot);
    residuals = zeros(n,1);
    for i = 1:n
        mu_i = mu_fun(T(i));
        rho_i = rho_fun(T(i));
        Re_k = (m_dot(i) * phi * sqrtK) / (mu_i * Af);
        f_k = (delta_p(i) * rho_i * sqrtK * Af^2) / (d * phi
            ^2 * m_dot(i)^2);
        f_k_model = 1 ./ Re_k + J;
        residuals(i) = f_k - f_k_model;
    end
end
```

Listing A.1: MATLAB script for the fitting Darcy–Forchheimer