NUMERICAL ANALYSIS OF AN INDUSTRIAL SCALE CONTINUOUS CONTAINER GLASS ANNEALING FURNACE

HAZAR ŞIŞIK

ÖZYEĞİN UNIVERSITY

SEPTEMBER, 2025

NUMERICAL ANALYSIS OF AN INDUSTRIAL SCALE CONTINUOUS CONTAINER GLASS ANNEALING FURNACE

by
Hazar Şişik

Thesis
Submitted in Partial Fulfillment of the
Requirements for the Degree of

Master of Science

in
Mechanical Engineering

Advisor: Asst. Prof. Altuğ Melik Başol

Graduate School of Science and Engineering
Özyeğin University
İstanbul

September, 2025

NUMERICAL ANALYSIS OF AN INDUSTRIAL SCALE CONTINUOUS CONTAINER GLASS ANNEALING FURNACE

Approved by:

Assistant Prof. Altuğ Melik Başol, Advisor
Department of Mechanical Engineering
Özyeğin University

Professor M. Pınar Mengüç
Department of Mechanical Engineering
Özyeğin University

Professor Fırat Oğuz Edis
Department of Aeronautical Engineering
Istanbul Technical University

Approval Date: Aug. 4, 2025

DECLARATION OF ORIGINALITY

I hereby declare that I am the sole author of this thesis and that this is the true copy of my thesis, including the final revisions, approved by my thesis committee. All data and information have been obtained, produced, and presented in accordance with the rules of research ethics and principles of academic honesty. As required by these rules, to the best of my knowledge I have acknowledged ideas, thoughts, and any copyrighted material in accordance with the standard referencing rules. I certify that any part of this thesis has not been submitted for a degree or diploma in another educational institution.

Hazar Şişik

ABSTRACT

Annealing is a critical heat treatment process in container glass manufacturing, used to relieve residual stresses caused by rapid cooling during forming. This process takes place in continuous annealing furnaces, where glass is gradually heated to the annealing temperature and then cooled at a carefully controlled rate to avoid stress-induced defects. Given the transient nature of heat transfer in the furnace, precise adjustment of heating and cooling rates are essential. Numerical simulations offer a powerful tool to determine these parameters based on the geometry and thermal properties of the glass. This study presents a numerical investigation of the annealing process applied to a real container glass geometry inside a continuous annealing furnace. Furnace operating conditions were derived from actual measurements and applied as boundary conditions in the model. The simulation employs two coupled solvers: one resolves heat transfer in the furnace domain, while the other tracks the thermal behavior of moving glass objects on the conveyor belt. Radiative heat transfer is modeled using a Monte Carlo ray tracing-based surface-to-surface approach, capturing detailed radiation effects among 10,000 glass objects. Convective heat transfer to glass is modeled using a porosity-based method. Simulation results were validated against real furnace zone temperatures and glass exit temperatures. Finally, the effect of two different conveyor speeds on the temperature distribution of glass was analyzed, and the implications on process quality were discussed.

Keywords: Continuous Annealing furnace, Heat treatment, Numerical simulation, Conjugate Heat Transfer

ÖZET

Tavlama, konteyner camı üretiminde kritik bir ısıl işlem sürecidir ve şekillendirme sırasında hızlı soğumadan kaynaklanan artık gerilimleri gidermek için kullanılır. Bu süreç, camın tavlama sıcaklığına kadar kademeli olarak ısıtıldığı ve ardından gerilim kaynaklı kusurların oluşmasını önlemek amacıyla dikkatlice kontrol edilen bir hızda soğutulduğu sürekli tavlama fırınlarında gerçekleşir. Fırındaki ısı transferinin zamana bağlı doğası göz önüne alındığında, ısıtma ve soğutma oranlarının hassas bir şekilde ayarlanması gereklidir. Numerik Analizler (simülasyonlar), camın geometrisi ve termal özelliklerine dayalı olarak bu parametrelerin belirlenmesinde güçlü bir araç sunar.

Bu çalışma, bir sürekli tavlama fırını içerisinde, gerçek bir konteyner camı geometrisine uygulanan tavlama sürecinin sayısal bir incelemesini sunmaktadır. Fırın çalışma koşulları gerçek ölçümlerden elde edilerek modele sınır koşulu olarak uygulanmıştır. Simülasyon, iki bağlı çözücü (solver) kullanmaktadır: bunlardan biri fırın bölgesindeki ısı transferini çözerken, diğeri konveyör bandı üzerindeki hareketli cam nesnelerin termal davranışını izlemektedir. Işınımla Isı Transferi, 10.000 cam nesne arasındaki detaylı ışınım etkilerini yakalayabilen Monte Carlo ışın izleme tabanlı yüzeyden yüzeye bir yaklaşımla modellenmiştir. Cama olan konvektif ısı transferi ise gözeneklilik (porozite) tabanlı bir yöntemle modellenmiştir. Simülasyon sonuçları, gerçek fırın bölge sıcaklıkları ve cam çıkış sıcaklıkları ile doğrulanmıştır. Son olarak, iki farklı konveyör hızı altındaki sıcaklık dağılımı üzerindeki etkiler analiz edilmiş ve bu etkilerin süreç kalitesi üzerindeki yansımaları tartışılmıştır.

Anahtar Kelimeler: Sürekli Tip Tavlama Fırını, Isıl İşlem, Numerik Simulasyon, Birleşik Isı Transferi

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ACRONYMS AND ABBREVIATIONS

AI	Artificial Intelligence
CFD	Computational Fluid Dynamics
CHT	Conjugate Heat Transfer
DOM	Discrete Ordinates Method
GPU	Graphics Processing Unit
RANS	Reynolds Averaged Navier Stokes
STL	Stereolithography
S2S-MC	Surface to Surface Monte Carlo

1. INTRODUCTION

Annealing is a vital step in the manufacturing of container glass, serving to relieve internal stresses that develop during the rapid cooling phase of forming. If not properly managed, these residual stresses can lead to warping, fractures, and optical distortions, ultimately compromising product reliability and quality [2, 3, 4]. As emphasized in foundational studies, effective stress relief depends on a precisely controlled thermal cycle, particularly during the transformation range of the cooling stage where structural relaxation is most active. Comprehensive treatments of annealing fundamentals can be found in [5, 6], which outline the role of annealing in ensuring structural stability through viscosity–temperature relationships and relaxation kinetics.

Modern continuous annealing furnaces are engineered to ensure uniform and gradual heating and cooling of glassware. These systems typically consist of multiple temperature-controlled zones, including heating, soaking, and cooling sections. Radiative heat transfer from burner tubes and furnace walls plays a dominant role, though convective and conductive mechanisms also contribute to the overall thermal field. The control of such furnaces remains challenging. Despite their importance, industrial annealing processes are often tuned via trial-and-error strategies, adjusting zone setpoints, conveyor speeds, or fan settings. This empirical approach risks reducing thermal uniformity, especially under high-throughput conditions. Classical handbooks [7, 8] emphasize that even with decades of operational experience, annealing remains highly sensitive to zone balancing, equipment configuration, and air leakage.

Numerical modeling presents a powerful, cost-effective alternative for analyzing and optimizing furnace performance without disrupting production. By providing detailed spatial and temporal insight into the thermal field, simulation-based tools allow process engineers to evaluate design choices and operational regimes under controlled conditions. However, modeling annealing furnaces is inherently complex due to the interacting effects of geometry, material properties, flow characteristics, and radiative energy exchange. Modern heat transfer references [9, 10] provide frameworks for incorporating surface-to-surface radiation, spectral dependence, and Monte Carlo methods into glass furnace simulations, enabling more accurate descriptions of the coupled physics.

1.1 Annealing Process Characteristics

The annealing process involves heating glass to a temperature where internal stresses can relax, followed by controlled cooling to prevent the reintroduction of stress. For soda-lime glass used in container production, this typically occurs between 500 and 600, with precise dwell times determined by wall thickness and product geometry [2]. Early work [3] identified the importance of uniform cooling in preventing surface cracking and uneven shrinkage. Later studies [4] demonstrated that stress and volume relaxation are governed by both structural and thermal history, reinforcing the need for temperature homogeneity throughout the annealing cycle. Further clarification of the annealing and strain points is given in [5], which links them to the temperature–viscosity curve dictating cooling rates for different glass compositions.

In practice, annealing furnaces achieve this through a series of temperature zones. The glass articles enter the hot end of the lehr and are slowly cooled as they pass through successively cooler sections. This zone-based structure is common to both horizontal and vertical lehrs, as well as to more specialized continuous furnaces used in the flat and container glass industries. The practical design and operation of these lehrs are described extensively in [7, 8], which highlight both traditional and modern configurations.

1.2 Glass Annealing Furnaces

Glass annealing furnaces are tunnel kilns employed in container and flat glass production to relieve residual stresses through controlled cooling. In industrial-scale setups, continuous annealing furnaces transport glass via rollers or conveyors through three thermal zones: *heating*, *soaking*, and *cooling*.

The **heating zone** uses tubular burners or electric heaters and is dominated by radiative heat transfer, especially from hot surfaces to glass. In tempering furnaces studied in [11], thermal radiation from resistive heaters interacts with forced convection from air jets, demonstrating that radiation controls heating rate while convection supports uniformity—principles analogous to annealing furnaces. A similar view is presented in [9, 10], which provide the theoretical framework for modeling such combined transfer modes.

In the soaking zone, temperature is held nearly constant to facilitate stress relaxation. The cooling zone then applies forced or natural airflow to gradually cool glass, preventing thermal shock and residual stress development.

The key process parameters, zone length, burner configuration, glass spacing, conveyor speed, and airflow, critically affect temperature uniformity and energy efficiency. For instance, [11] revealed that proper airflow coupled with radiative control is essential for symmetric heating across glass surfaces. These observations are consistent with the guidance in [7, 8], which emphasize practical adjustments to airflow and burner placement to maintain uniform cooling.

Continuous annealing furnaces deliver high throughput, yet engineers are challenged to maintain uniform temperatures across wide glass widths. Advanced modeling and control strategies, especially Computational Fluid Dynamics (CFD) based simulations that combine radiation and convection, have shown promise in improving thermal uniformity and reducing energy consumption.

Figure 1.1: *Schematic of a Glass Annealing Furnace*

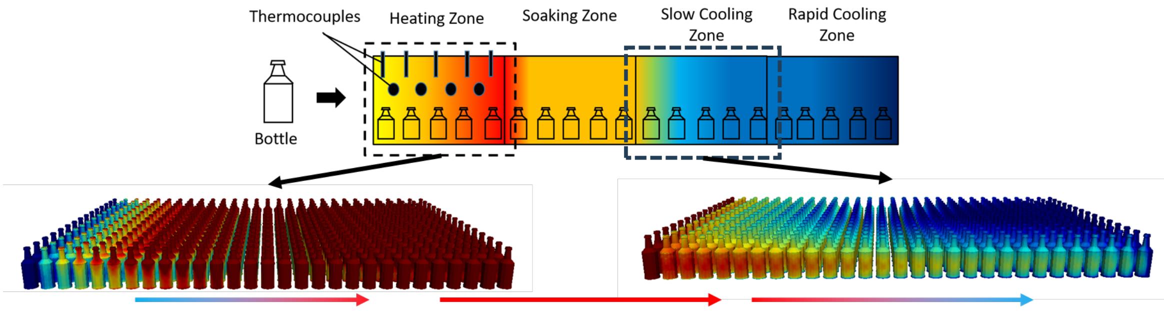

1.3 Challenges and Operational Considerations

Continuous-annealing furnaces are composed of multiple temperature-controlled zones. The tubular burners inside the heating section of the furnace are switched on or off to regulate the air temperature according to the desired set value. Often, the zone air temperatures are considered the primary control variable. However, regulating the heat transfer process based solely on zone temperature data can lead to suboptimal process conditions. The air temperature is only one factor in the thermal balance that determines the glass temperature: Radiative exchange, furnace geometry, and burner configuration are equally critical.

Despite advances in furnace design, several challenges continue to limit the robustness and efficiency of annealing operations. First, industrial practice often relies on empirical, trial-and-error adjustments of zone setpoints, conveyor speeds, or fan configurations. As highlighted in [7, 8], such empirical tuning can achieve acceptable throughput but does not guarantee thermal uniformity, especially under high production rates. This reliance on operator experience underlines the need for more systematic and model-based approaches.

Another difficulty lies in the measurement and control of local heat transfer mechanisms. Radiative exchange remains the dominant mode in the heating section, yet its accurate characterization is nontrivial. Surface emissivities, view factors, and spectral dependencies must be considered to avoid systematic errors [9, 10].

The representation of the glass load itself also poses modeling challenges. While radiation solvers require detailed surface resolution of individual containers to capture inter-object shading and heat exchange, flow solvers become prohibitively expensive if all bottles are meshed explicitly. Consequently, porous-medium approximations are often employed to represent pressure drop and momentum resistance [5]. However, accurate calibration of Darcy–Forchheimer coefficients remains essential, as errors in flow resistance translate directly into inaccuracies in convective heat transfer predictions.

Finally, multiphysics coupling significantly increases computational demands. Iterative exchange between radiative and convective solvers requires careful convergence monitoring, and Monte Carlo methods demand large numbers of rays to achieve statistical accuracy. This results in high computational cost, which can limit the resolution or time horizon of parametric studies. Moreover, common simplifying assumptions, such as treating glass and refractory surfaces as gray or black bodies, can introduce systematic bias in cases where spectral selectivity is important.

In summary, the challenges of annealing furnace operation stem from a combination of empirical control traditions and the need for complex and computationally intensive multiphysics simulations. These factors underscore the complexity of developing predictive models that can simultaneously capture industrial realism and remain computationally feasible.

2. LITERATURE REVIEW

In recent decades, numerical modeling of industrial furnaces—particularly for annealing applications—has advanced significantly. Earlier thermal analyses often relied on empirical correlations or simplified steady-state assumptions that overlooked the strong coupling between conduction, radiation, and convection within enclosed furnace geometries. However, with the increased availability of computational power and advanced numerical solvers, recent studies have adopted more rigorous Conjugate Heat Transfer (CHT) approaches that capture the multi-physics nature of industrial heat treatment processes. These developments are particularly crucial in continuous annealing furnaces, where thermal uniformity and energy efficiency are tightly linked to both product quality and operational cost. Classical references in glass science and furnace engineering [5, 6, 7, 8] emphasize that achieving stress-free glass depends on carefully controlled thermal cycles and that empirical adjustments often dominate in industrial practice.

Yıldız et al. [12] proposed a novel segregated modeling framework for the thermal simulation of continuous heat treatment furnaces. Their approach decouples solid conduction from surface-to-surface radiation, solving them in separate domains. This method enhances computational efficiency while maintaining sufficient accuracy, particularly in radiation-dominated heating zones typical in annealing operations. The model uses a finite-volume method for the solid domain and a radiosity-based radiative solver for wall-to-wall radiation transfer. Boundary conditions are carefully managed to ensure energy consistency at the interfaces. One of the key contributions was that the glass bottles behaved like a black surface absorbing 80 percent of the radiation emitted. Another key contribution is the ability to simulate furnace configurations with moving glass loads under steady-state conditions, mimicking industrial annealing conveyors. Their simulations showed that decoupling allows for flexible integration into existing solvers and significant reduction in simulation time without major loss in solution fidelity.

Altun and Başol [13] conducted a detailed numerical investigation into the radiative heating effectiveness of a continuous glass annealing furnace. Their study focused on how geometrical factors—particularly container spacing and arrangement—affect temperature uniformity within the furnace. A three-dimensional thermal model was developed, incorporating surface-to-surface radiation via a view-factor approach and conduction within the glass. By simulating different container spacing scenarios, the authors demonstrated that shadowing effects between containers could lead to significant thermal gradients, especially near furnace walls and in the centerline regions of the conveyor. The results showed that optimized spacing configurations improved temperature uniformity across the glass surface, minimizing thermal stress and energy losses. Their findings emphasize the critical role of furnace geometry in heat transfer behavior, suggesting that design-level modifications can enhance both thermal performance and product quality.

Dillon [14] proposed a dual-solver computational framework to model conjugate heat transfer in continuous thermal processing systems. In this approach, the solid and fluid domains are handled by separate solvers, which are then coupled through interfacial boundary conditions. This partitioned strategy allows for scalable and modular implementation of complex thermal models, particularly beneficial in large-scale industrial applications where computational cost is a limiting factor. The study demonstrated the model's capability to handle transient heat transfer problems with realistic furnace geometries and boundary conditions. Validation against experimental data confirmed the accuracy of the approach, showing good agreement in both transient temperature evolution and spatial thermal distribution. This work illustrates the effectiveness of dual-solver strategies in simulating thermal processes where both conduction and convection/radiation are tightly coupled.

Hajaliakbari and Hassanpour [15] performed a comprehensive thermal performance analysis of a continuous annealing furnace used in steel strip processing. Their study focused on improving energy efficiency through a combined assessment of furnace operating parameters, including strip velocity, radiant heating power, and thermal losses through furnace walls. The authors developed a numerical model based on the finite volume method, solving transient heat conduction in the steel strip and incorporating radiative and convective heat exchange within the furnace. Radiative heat transfer was handled using the Discrete Ordinates Method (DOM), which allowed directional dependence of thermal radiation to be resolved without the computational cost of stochastic techniques. The model was validated against real operational data, showing good agreement in temperature distribution and energy usage. A key finding was that simultaneous tuning of strip velocity and radiant heating power led to improved thermal uniformity and reduced energy consumption. Their methodology aligns with the growing trend toward conjugate heat transfer modeling.

In addition, various studies have explored different types of radiative solvers as components of conjugate heat transfer frameworks in furnace simulations. These include deterministic solvers such as view-factor-based models and the Discrete Ordinates Method (DOM), as well as stochastic Monte Carlo ray-tracing methods [9, 10, 16] that do not require explicit view factor calculation and offer greater robustness for complex geometries. These techniques have been applied successfully to industrial furnaces with diverse configurations and scales. For example, Depree et al. [17], Jurić et al. [18], and He et al. [19] presented studies where radiation–conduction coupling was achieved via these alternative solvers, further demonstrating the flexibility of modern thermal modeling strategies.

Filipponi et al. [20] conducted a thermal analysis of an industrial forging furnace with a focus on quantifying energy losses during door opening events. Using a validated CFD model, the study simulated the thermal behavior of the furnace during steel insertion and extraction, which are known to cause sharp temperature drops due to intense convection with ambient air. The authors found that a 10-minute door opening results in a total energy loss of 5606 MJ, with convection accounting for the majority (5020 MJ), followed by radiation and conduction. These findings underscore the importance of operational scheduling and door management strategies in minimizing energy waste and pollutant emissions.

Fei et al. [21] investigated the thermal behavior of steel strips during the heating stage of a vertical continuous annealing furnace. A numerical model was developed to simulate the strip heating process, incorporating both radiative and convective heat transfer. Their findings highlighted that uniform temperature profiles are highly sensitive to strip velocity and furnace zone temperature distribution. Optimizing these parameters is critical to ensure product consistency and avoid excessive thermal gradients that can degrade mechanical performance.

Strommer et al. [22] presented a first-principles mathematical model for a direct-fired continuous strip annealing furnace. The model integrated transient conduction in strips and rolls, flue gas mass–enthalpy balances, and radiative-convective coupling using the zone method. Validation against plant measurements confirmed high fidelity: flue gas and strip temperature profiles matched pyrometer readings within ± 5 K, and radiant tube/wall temperatures showed agreement within 10–20 K. The computational performance of the model was also notable, enabling a full 5-hour furnace cycle simulation in just 9 minutes on a standard desktop setup. These results demonstrate the industrial viability of high-fidelity conjugate models for both process optimization and real-time control.

Taken together, these studies illustrate the evolution of furnace modeling from simplified empirical correlations to high-resolution multiphysics simulations. At the same time, they demonstrate that industrial realism brings unresolved challenges: uncertain boundary conditions, reliance on empirical control traditions, and the trade-off between accuracy and computational feasibility. These themes directly motivate the present work, which aims to extend state-of-the-art modeling techniques to the specific context of container glass annealing, addressing both thermal uniformity and energy efficiency.

As emphasized in the Introduction, continuous-annealing furnaces are composed of multiple temperature-controlled zones. The tubular burners inside the heating section of the furnace are switched on or off to regulate the air temperature according to the desired set value. Often, the zone air temperatures are considered the primary control variable. However, regulating the heat transfer process based solely on zone temperature data can lead to suboptimal process conditions. The air temperature is only one factor in the thermal balance that determines the glass temperature: radiative exchange, furnace geometry, and burner configuration are equally critical.

Despite advances in furnace design, several challenges continue to limit the robustness and efficiency of annealing operations. First, industrial practice often relies on empirical, trial-and-error adjustments of zone setpoints, conveyor speeds, or fan configurations. As highlighted in [7, 8], such empirical tuning can achieve acceptable throughput but does not guarantee thermal uniformity, especially under high production rates. This reliance on operator experience underlines the need for more systematic and model-based approaches.

Another difficulty lies in the measurement and control of local heat transfer mechanisms. Radiative exchange remains the dominant mode in the heating section, yet its accurate characterization is nontrivial. Surface emissivities, view factors, and spectral dependencies must be considered to avoid systematic errors [9, 10]. More advanced stochastic approaches such as the Monte Carlo method [16] provide higher fidelity for complex geometries, but their computational expense often limits industrial applicability.

The representation of the glass load itself also poses modeling challenges. While radiation solvers require detailed surface resolution of individual containers to capture inter-object shading and heat exchange, flow solvers become prohibitively expensive if all bottles are meshed explicitly. Consequently, porous-medium approximations are often employed to represent pressure drop and momentum resistance [5]. as an example classical correlations such as the Ergun equation [23] or porous convection frameworks [24]

Leakage and heat losses add further uncertainty. Experimental and analytical studies of furnace thermal performance, such as the work of Rivas Evans [25], demonstrate that small inlet or outlet leakages can alter furnace flow balance, disturb heat transfer coefficients, and compromise thermal uniformity. Such effects are frequently ignored in numerical models, yet they can dominate discrepancies between simulations and industrial measurements.

Finally, multiphysics coupling significantly increases computational demands. Iterative exchange between radiative and convective solvers requires careful convergence monitoring, and Monte Carlo methods demand large numbers of rays to achieve statistical accuracy. This results in high computational cost, which can limit the resolution or time horizon of parametric studies. Moreover, common simplifying assumptions, such as treating glass and refractory surfaces as gray or black bodies, can introduce systematic bias in cases where spectral selectivity is important.

In summary, the challenges of annealing furnace operation stem from a combination of empirical control traditions, uncertain boundary conditions, and the need for complex and computationally intensive multiphysics simulations. These factors underscore the complexity of developing predictive models that can simultaneously capture industrial realism and remain computationally feasible.

3. NUMERICAL METHODOLOGY

3.1 Furnace Domain and Boundary conditions

The annealing furnace examined in this study is a continuous, industrial-scale unit designed for the treatment of container glass. This system utilizes a linear conveyor to move glass items through several zones, each with precisely controlled thermal conditions. A three-dimensional representation of the furnace is shown in Figure 3.1. The furnace features a double-wall construction, to regulate the air inside the furnace. Air drawn in by circulation fans installed at the top of the furnace is directed into this intermediate space and then reintroduced at the bottom of the conveyor system to the glassware. This recirculation mechanism promotes thermal uniformity throughout the zones and improves heat exchange between the furnace atmosphere and the glass products.

Figure 3.1: *Isometric view of the Furnace*

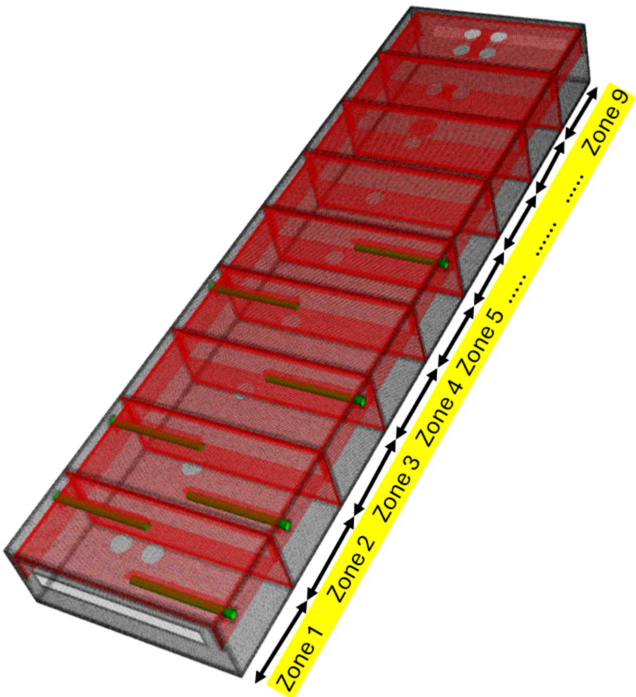

The furnace is composed of nine distinct thermal zones, which are grouped into five heating and four cooling sections. The heating areas incorporate tubular radiant burners located along the upper section of the furnace, marked in green in Figure 1. During standard operation, the burners in the initial two zones function continuously to preheat incoming glassware. The burners in subsequent zones are typically used intermittently and consume significantly less gas compared to the upstream burners. Therefore, for the purpose of validating the computational model, only the burners in the first zone were considered active. These active burners were assigned a tube surface temperature of 677°C, whereas the inactive ones were set to 527°C. The temperature data for the active tubes was acquired via visual inspection near the furnace inlet and confirmed through infrared measurements using a handheld thermometer.

The final two zones serve as the cooling stage and are equipped with air inlet vents that draw in ambient air to regulate the desired cooling profiles. Additionally, an exhaust vent is located just upstream of these cooling zones to extract hot gases from the furnace interior. The layout of both the air inlets and the exhaust outlet is presented in Figure 3.2. The furnace dimensions and the operating conditions can be found on Table 3.1

Figure 3.2: *Inner Furnace Mesh with Furnace Features*

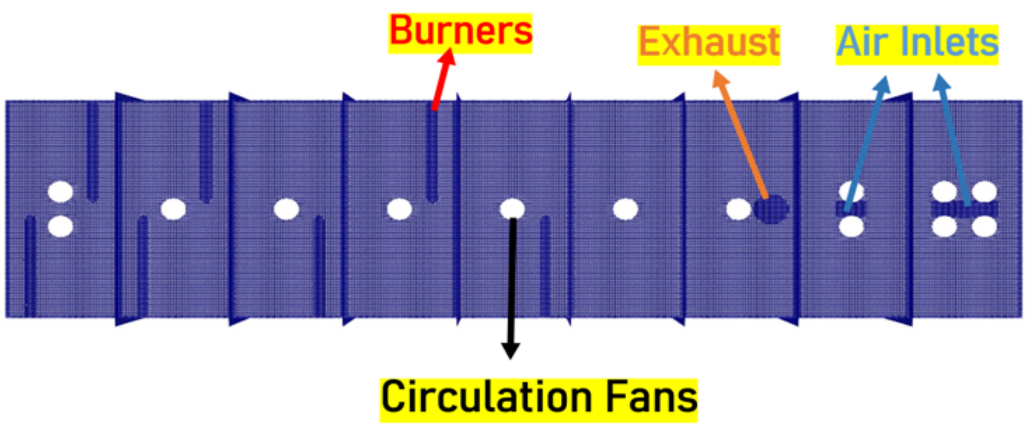

Table 3.1: *Furnace Geometric Parameters and Boundary Conditions used for the study.*

Parameter	**Value**
Total Furnace Length	20.25 m
Furnace Width	5.2 m
Furnace Height	2.0 m
Number of Zones	9 (5 Heating, 4 Cooling)
Number of Glassware Rows	203
Number of Glassware per Row	40
Total Number of Glassware	8,120
Glassware Temperature at Inlet	477 °C
Burner Tube Wall Temperature (Zone 1 – active)	677 °C
Burner Tube Wall Temperature (Zone 2-9 – inactive)	527 °C
Conveyor Speed	60 cm/min
Total Cooling Air Inlet Mass Flow Rate	2.6 kg/s
Cooling Air Temperatures	27 °C

3.2 Numerical Solvers and Models

The thermal processing within the continuous annealing furnace is modeled as a conjugate heat transfer problem, encompassing both solid and fluid domains with distinct governing equations and boundary conditions. These domains are resolved using separate solvers operating in an iterative fashion.

For the solid domain, which comprises the moving glassware, a transient heat conduction solver is employed. Specifically, the moving solid region includes all 40 glass items in a single row, with each object discretized into approximately 500 hexahedral volume cells and 1,200 surface triangles. The interaction between the glass surfaces and the furnace walls, including the tubular burners, is modeled using an in-house developed surface-to-surface radiation solver based on the Monte Carlo ray tracing method. In this approach, all surfaces, including the glassware, are assumed to behave as black bodies. Given the strong infrared absorptivity of glass, this assumption remains physically reasonable.

In addition to radiation, convective heat exchange between the glass objects and the surrounding furnace atmosphere is captured by the same transient conduction solver. Temperature distributions for the furnace walls and the surrounding air, necessary for these convective calculations, are extracted from the solution of the fluid domain.

The fluid flow and thermal behavior of the furnace atmosphere are resolved using a steady-state RANS solver, specifically the buoyantSimpleFoam solver from OpenFOAM. This solver applies the SIMPLE algorithm for solving the Reynolds Averaged Navier Stokes (RANS) equations, with turbulence represented via the k–omega model. The fluid domain, shown in Figures 3.1 and 3.2, consists of around 4 million hexahedral cells. Notably, the geometry of individual glass items is not explicitly resolved in this mesh; instead, a porous zone modeling strategy is used, significantly reducing computational cost.

Also shown in Figure 3.2 are the burners, circulation fans (indicated as white circles), exhaust outlet, and the air inlet vents. The air vents in the last two zones are assigned a total pressure equal to atmospheric pressure, and the inflow rate is regulated by adjusting the porosity values in those regions. At the exhaust outlet, a negative gauge static pressure is imposed to replicate the suction effect of the exhaust fan.

Radiative heat transfer between the furnace walls is again computed using the surface-to-surface radiation model, this time incorporating temperature inputs from the solid domain solver to determine the heat exchange between glassware and walls. To maintain high geometric fidelity in radiation modeling, the full 3D surface of all 8,120 glass items is included, resulting in over 10 million surface elements. Convective heat transfer between the glassware and furnace air is approximated using a constant heat transfer coefficient applied uniformly to the glassware surfaces.

3.2.1 Modeling of Furnace Walls

The furnace geometry consists of two main layers: the *inner* and *outer* furnace regions. Both regions are implemented with no-slip boundary conditions. However, the inner furnace is defined as a baffle interface, assigned as *master* and *slave* patches.

Figure 3.3: *Double Layer Structure View from the Inlet*

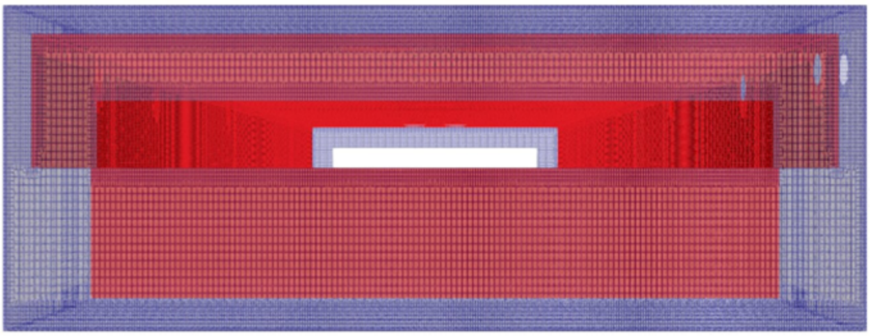

In the adopted modeling framework, the baffle geometry is represented with two distinct surface patches, allowing radiative heat fluxes to be mapped separately on each side in accordance with the solver's formulation.

Another practical benefit of modeling these surfaces as baffles is that it eliminates the need to explicitly prescribe an interface thickness. Because baffles are treated as infinitesimally thin internal surfaces in OpenFOAM, they simplify the mesh without altering cell topology or requiring additional resolution. In addition, baffle interfaces facilitate advanced numerical handling of sharp discontinuities in thermal or radiative properties, which is especially advantageous in furnace simulations involving steep internal gradients.

All furnace walls, including both the inner and outer regions, are assigned flux-based boundary conditions in OpenFOAM, specifically, the **externalWallHeatFluxTemperature** and **thermalBaffle1D** types. These boundary conditions are configured to calculate wall temperature implicitly as a function of the net heat flux, rather than prescribing temperature directly. Within the coupled solver framework, radiative heat fluxes are provided by the radiation solver, while convective heat fluxes are computed by OpenFOAM. These two contributions are mapped onto the primitives of the wall surfaces at each coupling step. This design enables a consistent and modular interface between the solvers, as all thermal exchange is handled via flux terms. The details of the coupling algorithm will be discussed in the relevant section (3.4,3.5).

3.2.2 Modeling of the Burners

In this simulation framework, the burners are not modeled through detailed combustion processes or participating media radiation. Instead, a simplified representation is employed by prescribing fixed wall temperatures and, in active cases, adding hot air injection as a source of mass, momentum, and thermal energy.

A total of seven burners are explicitly meshed and defined as distinct regions within both the radiative and convective solvers. Burner wall temperatures are assigned individually for each unit using a Python-based pre-processing script. This script ensures temperature mapping across both the radiative and OpenFOAM solvers. The temperature values are prescribed as fixed boundary conditions and are not solved for during the simulation. These values were approximated based on observations of surface colors. Check Figure 3.4 and Figure 3.5 for the table and the real view from the inlet.

Figure 3.4: *Temperature vs Glow Colour Table [1]*

color	approximate temperature		
	°F	°C	K
faint red	930	500	770
blood red	1075	580	855
dark cherry	1175	635	910
medium cherry	1275	0690	0965
cherry	1375	0745	1020
bright cherry	1450	0790	1060
salmon	1550	0845	1115
dark orange	1630	0890	1160
orange	1725	0940	1215
lemon	1830	1000	1270
light yellow	1975	1080	1355
white	2200	1205	1480

For burners in an *active* state, an additional air inlet is modeled to represent hot air injection into the furnace domain with $\vec{U}_{\mathrm{y}} = 4$ m/s. The velocities are deduced from the natural gas usage of the furnace.The incoming flow is directly introduced into the domain with an inlet boundary where the velocity and temperature boundary conditions are given. In contrast, *passive* burners are only modeled through their static wall temperature values, with no internal flow or energy input.

This simplified burner modeling strategy significantly reduces computational complexity while preserving the thermal effect of burner operation. The impact of burner activity is captured through two main mechanisms: the radiative heat flux emitted by the hot burner surfaces and the convective effect of the hot air introduced into the domain. These combined influences are sufficient to evaluate the global furnace temperature distribution and wall heat flux behavior.

Figure 3.5: *A typical annealing furnace viewed from inlet*

3.2.3 Modeling of the Circulation Fans

The furnace includes nine internal circulation fans, which are modeled as volumetric momentum sources using the **meanVelocityForce** feature of OpenFOAM's **fvOptions** framework. Due to the lack of direct information on the actual operating speeds, their inlet velocity was set to a constant value of 20 m/s.

Each fan is geometrically represented as a vertical cylinder defined via the **cylinderToCell** source in the **topoSetDict** utility. In certain cases, multiple cylindrical regions are combined using the **add** action to capture more complex fan geometries. The resulting **cellSet**s are then converted into **cellZone**s using the **setToCellZone** method. These zones serve as the regions over which the momentum sources are applied.

The applied mean velocity vectors, $\vec{U}_z = 20$ m/s, are oriented along the axial direction of each fan zone and uniformly distributed within the designated regions.

3.2.4 Modeling of the Air Intakes and Exhaust

The air intake and exhaust system is modeled using a combination of porous media zones and pressure boundary conditions to realistically simulate the cooling airflow within the furnace. The goal of this setup is to represent cold ambient air entering through dedicated intake vents and exiting through a passive outlet region, forming a natural cooling loop.

Each intake is modeled using the **pressureDirectedInletVelocity** boundary condition, allowing flow to enter along a prescribed direction in response to local pressure gradients.The pressure is set to 100 kPa. The inlet velocity is initialized as zero and evolves based on the overall pressure distribution. The inflow air temperature is fixed at 300 K, representing ambient cooling air entering the furnace domain.

To reflect the physical structure of intake vents (e.g., perforated panels or louvered ducts), each intake region is treated as a porous medium using OpenFOAM's **explicitPorositySource** model. The Darcy-Forchheimer formulation is applied, with user-defined resistance coefficients to control the flow rate through each zone. These porous regions are defined via **topoSetDict** using **boxToCell** selection, and then converted into **cellZone**s for use in the simulation.

The Exhaust as represented in Figure 3.2 is defined using an **inletOutlet** condition for velocity and a relatively low fixed pressure value (97 kPa) in the p_rgh field. This creates a passive suction effect, drawing air through the domain. Importantly, the porous treatment of the intake regions not only enables physically meaningful airflow restriction, but also ensures that the total volumetric inflow and outflow remain balanced, even under varying furnace conditions. As a result, the model maintains both numerical stability and thermodynamic consistency across the cooling system.

An important motivation for modeling the intake vents as porous media is to maintain mass flow rate through the domain without prescribing fixed inflow conditions. Using a **pressureDirectedInletVelocity** condition in combination with Darcy-Forchheimer resistance allows for passive control of the inflow, whereby the actual velocity adjusts dynamically in response to pressure gradients.

In contrast, the porous media formulation introduces tunable hydraulic resistance via the Forchheimer coefficients. By adjusting these coefficients, the simulation can achieve an approximate equilibrium between inflow and outflow ($\dot{m}_{\text{in}} \approx \dot{m}_{\text{out}}$), preserving mass conservation and enabling more flexible control over cooling airflow behavior under various thermal conditions.

3.2.5 Radiation Modeling

Radiative heat transfer within the furnace is simulated using a Graphics Processing Unit (GPU) accelerated Surface to Surface Monte Carlo (S2S-MC), implemented with **CUDA** and **NVIDIA OptiX**. This custom radiation solver directly handles ray emission, intersection, and energy exchange processes over the triangulated furnace geometry, which is exported in Stereolithography (STL) format and parsed by a dedicated mesh reader [26]. The solver architecture separates geometry processing from radiation transport logic, enabling modular integration with OpenFOAM for conjugate heat transfer coupling.

The method relies on stochastic emission and transport of energy packets (rays) from discretized surface elements to capture view factors and directional radiative interactions. Each triangle emits a predefined number of rays (here, 50 rays per triangle), each carrying a fraction of the total emitted energy. The emitted energy for a triangle is computed using the Stefan–Boltzmann law:

$$Q_i^{\text{emit}} = \varepsilon \sigma A_i T_i^4, \tag{3.1}$$

where A_i is the surface area, T_i is the surface temperature, ε is the emissivity (assumed to be 1.0), and σ is the Stefan–Boltzmann constant. This energy is equally divided across all rays as:

$$E = \frac{Q_i^{\text{emit}}}{N_{\text{ray}}}. \tag{3.2}$$

Each ray thus becomes an energy packet, tracked through emission, transport, and potential absorption events. The energy is encoded into a 32-bit floating point representation and passed to the ray tracing pipeline as CUDA payload.

The origin of each emitted ray is sampled from within the emitting triangle using an **edge-length-weighted centroid** method, often referred to as a geocentric scheme. Specifically, for a triangle with vertices $\vec{v}_0$, $\vec{v}_1$, and $\vec{v}_2$, and edge lengths a, b, and c, the ray origin $\vec{x}_{\text{origin}}$ is given by:

$$\vec{x}_{\text{origin}} = \frac{a\vec{v}_0 + b\vec{v}_1 + c\vec{v}_2}{a + b + c} + \delta \vec{d}, \tag{3.3}$$

where $\vec{d}$ is the sampled ray direction (described below), and δ is a small offset used to avoid self-intersections. This origin sampling ensures better spatial coverage and uniformity over highly irregular or anisotropic mesh elements as compared to barycentric sampling.

The ray direction is sampled from a hemisphere oriented about the surface normal, following a cosine-weighted Lambertian distribution. To achieve this, a local orthonormal coordinate frame $(\vec{U}, \vec{V}, \vec{W})$ is constructed:

- $\vec{W}$ is the normalized surface normal at the triangle (i.e., the emission direction)
- $\vec{U} = \vec{W} \times$ [1,1,1], and $\vec{V} = \vec{U} \times \vec{W}$.

Each of these vectors is then scaled by the norm of $\vec{W}$ to ensure consistent magnitude in the local frame. This defines a transformation basis for projecting randomly sampled directions.

Random numbers r_1 and r_2 drawn from a uniform distribution $[0, 1)$ are used to generate spherical coordinates via:

$$\phi = 2\pi r_1, \quad \theta = \cos^{-1}(\sqrt{1 - r_2}), \tag{3.4}$$

which correspond to cosine-weighted sampling over the hemisphere. The final ray direction in world space is given by:

$$\vec{d} = \sin\theta\cos\phi\,\vec{U} + \sin\theta\sin\phi\,\vec{V} + \cos\theta\,\vec{W}. \tag{3.5}$$

This Monte Carlo sampling strategy ensures denser sampling in the directions close to the surface normal, enhancing simulation accuracy in radiative heat transfer. The stochastic nature of the method enables each ray to explore a distinct path, allowing statistical convergence of flux distributions through repeated sampling.

Once a ray origin and direction are defined, emission is executed using the **optixTrace** function. For surfaces defined as internal baffles (see Section 3.2.1), a second ray is launched in the opposite direction to simulate bidirectional exchange.

Each ray carries its energy E and is tested for intersection with the triangulated domain. If a hit is detected, the **closest-hit** shader computes the corrected absorbed energy as:

$$E_{\text{abs}} = E\left(\frac{T_{\text{hit}}}{T_{\text{emit}}}\right)^4, \tag{3.6}$$

and updates the corresponding surface accumulation array using atomic operations. The absorbed radiative energy is mapped back to the emitting surface element, as is characteristic of the backward Monte Carlo ray tracing approach. Rays that do not intersect any geometry are recorded in a **miss** shader, which contributes to diagnostics and energy closure checks.The CUDA implementation is fully parallelized over surface elements and rays. Random numbers are generated using seeded **curand** and **tea** hash functions to ensure statistically independent ray paths. Directional vectors are normalized on-device and energy values are tracked in global memory buffers. Absorbed energy is atomically accumulated per primitive to avoid race conditions.

Following each Monte Carlo sweep, the net heat flux on each surface element is computed as the difference between emitted and absorbed energy. This net radiative heat flux is then exported to disk and mapped onto the OpenFOAM mesh using a python code as a tool where the data is interpolated between Surface-to-Surface Radiative Solver and OpenFOAM. This mapped field is consumed by the **externalWallHeatFluxTemperature** and **thermalBaffle1D** boundary conditions within OpenFOAM, which combine it with convective heat fluxes to compute the surface temperature:

$$\dot{q}_{\text{net}} = \dot{q}_{\text{rad}} + \dot{q}_{\text{conv}} \tag{3.7}$$

3.2.6 Mesh of the Objects and Monte Carlo Ray Sizes

Two separate solvers utilize two different meshes. The steady-state OpenFOAM based convective solver only requires a volume mesh, the volume mesh is hexahedral. The steady-state surface-to-surface radiative solver utilizes the triangular surface meshes of both the furnace elements and the glassware, the real values of the surface mesh can be found in Table 3.2. The transient solver only requires 1 row of glassware mesh, its surface mesh for radiation, and volume mesh for conduction. The transient radiation model also utilizes the surface mesh of the furnace elements.

From the furnace elements surface meshes for each primitive 50 rays are emitted

Table 3.2: *Steady State Solver Mesh Sizes.*

Furnace Element	Mesh Size
Volume Mesh of Fluid Domain	4,045,543
Outer Furnace Surface Mesh	619,894
Inner Furnace Surface Mesh	408,995
Total Surface Mesh of Tubular Heaters	57,683
1 Glass Object Surface Mesh	1,260
1 Glass Object Volume Mesh	505
Total Surface Mesh of the Glassware	10,231,200
Total Surface Mesh Size	11,317,772

for 50 times, making a total of 54,328,600 rays for a single radiative iteration and the radiative solver operation continues until 50 iterations making a grand total of 2,716,430,000 rays emitted. For the results, however, average of the sum of each iteration is taken into consideration, meaning total heat flux at the end of radiative solver operation is divided by 50.

The modeling of the glass objects was performed using ANSYS ICEM-CFD .As illustrated in Figure 3.6, both the surface and volumetric mesh structures are presented. Given the presence of over 8000 individual glass items within the computational domain, a relatively coarse mesh resolution was utilized.

Figure 3.6: *Mesh of a Single Glass Object*

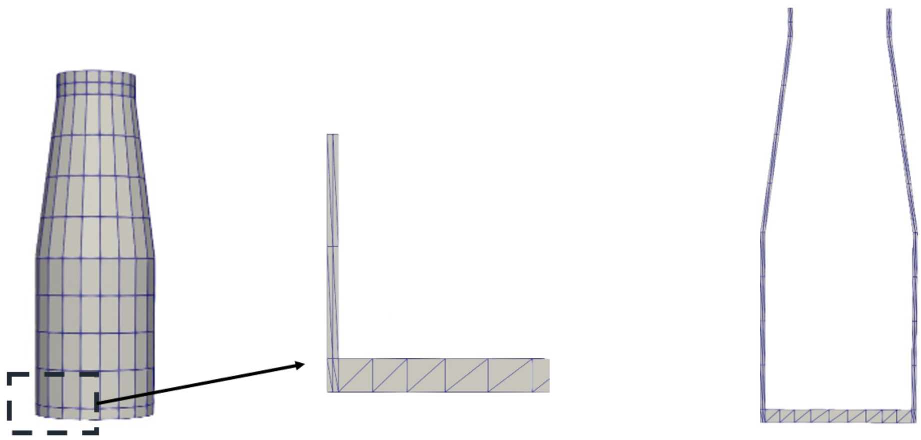

3.3 Porous Zone Modeling of the Glassware

To accurately model the airflow resistance introduced by the glassware inside the furnace, these regions are treated as porous zones governed by the Darcy–Forchheimer formulation. Proper estimation of the Darcy and Forchheimer coefficients is essential for capturing the pressure drop and velocity distribution in these regions, especially since the flow must interact realistically with the densely packed bottles. Therefore, a dedicated set of OpenFOAM simulations was conducted to calibrate these coefficients using simplified domains containing a single-bottle configuration and a three-bottle arrangement. The inclusion of the latter allows for the capture the flow interactions between adjacent bottles, which are critical for accurately representing the overall pressure resistance of the porous media. The tube walls are positioned to touch the bottle's top and bottom surfaces. Slip boundary conditions are applied to the tube walls to minimize shear effects, while the bottle surfaces are assigned no-slip conditions to represent realistic fluid–solid interaction. A fixed outlet pressure of 0 Pa is imposed, and the inlet pressure is determined implicitly via a zero-gradient boundary condition.

The pressure drop results obtained for varying inlet velocities are fitted to a polynomial function to extract the Darcy–Forchheimer coefficients. The fitting formulation is adopted from the OpenFOAM Wiki documentation on Darcy–Forchheimer modeling [27].

According to the Wiki, when velocity-dependent pressure drop data are available, the following empirical relation can be used:

$$\frac{\Delta p}{\Delta x} = Au + Bu^2 \tag{3.8}$$

Here, the coefficients A and B are defined as:

$$A = \mu \mathbf{D} \tag{3.9}$$

$$B = \frac{1}{2} \rho \mathbf{F} \tag{3.10}$$

By evaluating the linear and quadratic components of the velocity–pressure relationship, the Darcy ($\mathbf{D}$) and Forchheimer ($\mathbf{F}$) coefficients can be inferred. It is important to note that both $\mathbf{D}$ and $\mathbf{F}$ may exhibit directional dependence due to anisotropic resistance in the porous domain.

The computational domain includes an inlet, outlet, surrounding walls, and a centrally positioned bottle as previously described. The volumetric mesh comprises 529,264 cells, with 704 cells at both the inlet and outlet, 117,696 surface cells on the bottle, and 10,962 cells on the wall boundaries. Local mesh refinement is applied in the vicinity of the bottle to resolve flow features more accurately, particularly in the wake region, where turbulence effects are anticipated due to the obstruction created by the bottle. Figure 3.7 shows the refined mesh around the bottle geometry.

The isometric view of the case and the bottle position can be found on figure 3.8

Figure 3.7: *Mesh of the Bottle Zone*

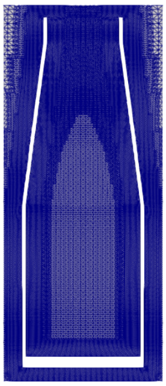

Figure 3.8: *Isometric View of the Setup*

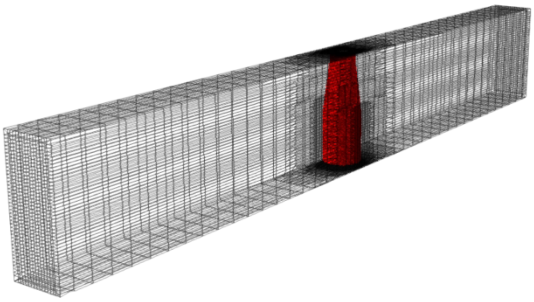

After the simulations the pressure drops relative to the velocities were collected. In the Table 3.3 below the results can be seen.

The corresponding fit is made with matlab polyfit function,the fitted function is

$$y = 34.35U^2 + 12.58U \tag{3.11}$$

the corresponding coefficients **A** and **B** is found to be then 12.58 and 34.35 in Figure 3.9 below.

Table 3.3: *Velocty vs Pressure Drop - 1 Bottle*

X-Velocity	Pressure Drop
1 m/s	6.86
2 m/s	24.919
3 m/s	55.084
4 m/s	95.131

Figure 3.9: *Polynomial fit*

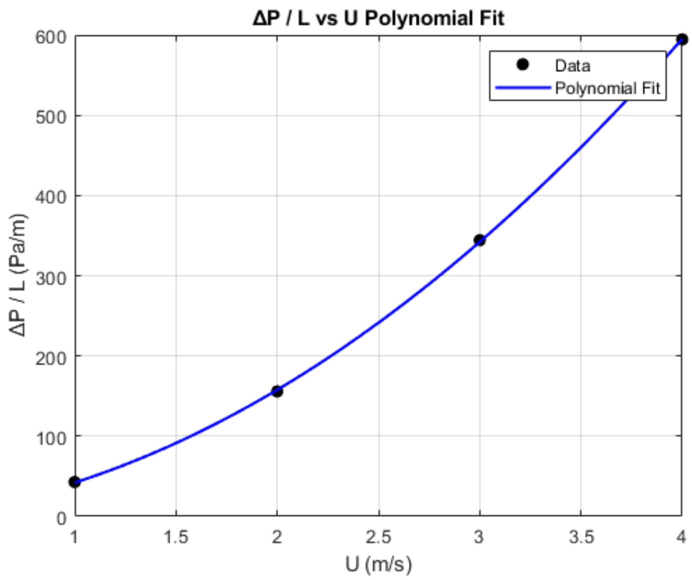

from Equations (2.2) and (2.1) the Darcy-Forchheimer Coefficients are found to be **Darcy = 0** and **Forchhemier = 116** for the single bottle case and then a validation case has been made, with larger tube domain and with 8 bottles.

Figure 3.10: *Top and Front View of the Larger Domain (Red Rectangle Represents Porous Zone)*

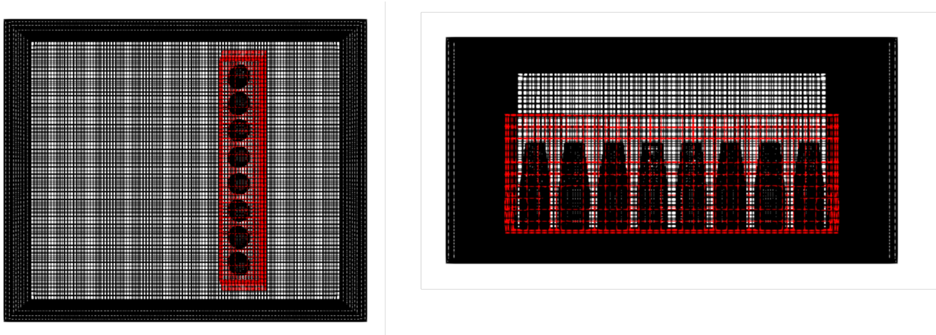

Figure 3.11: *Mesh of Porous Zone vs Bottle Zone*

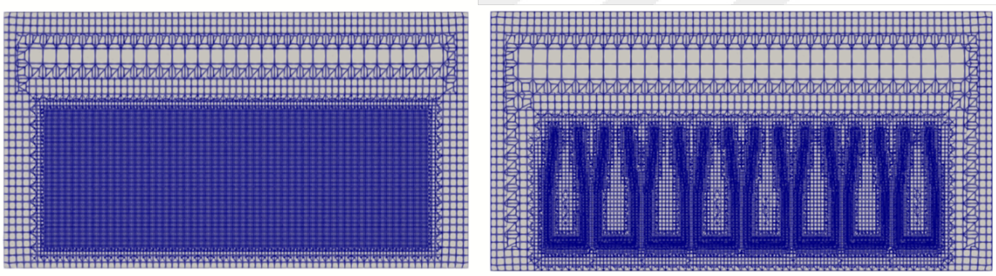

The results are for the larger domain with glass objects the pressure drop found out to be **3.72 Pa**, with the Porous zone with the found coefficients, the pressure drop found out to be **3.84 Pa**.

Figure 3.12: *Isometric view of the 3 Bottle Case*

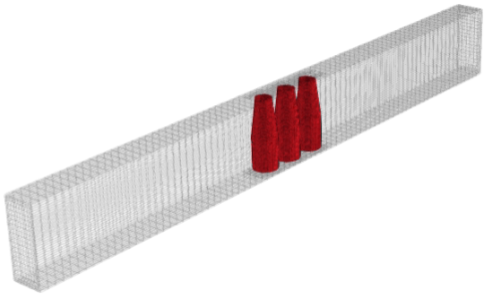

For the Case Study with 3 Adjacent Bottles (3.12) the results indicated a smaller Forchheimer Coefficient with the following fit 3.13.

Figure 3.13: *PolyFit of the 3 Bottle Case*

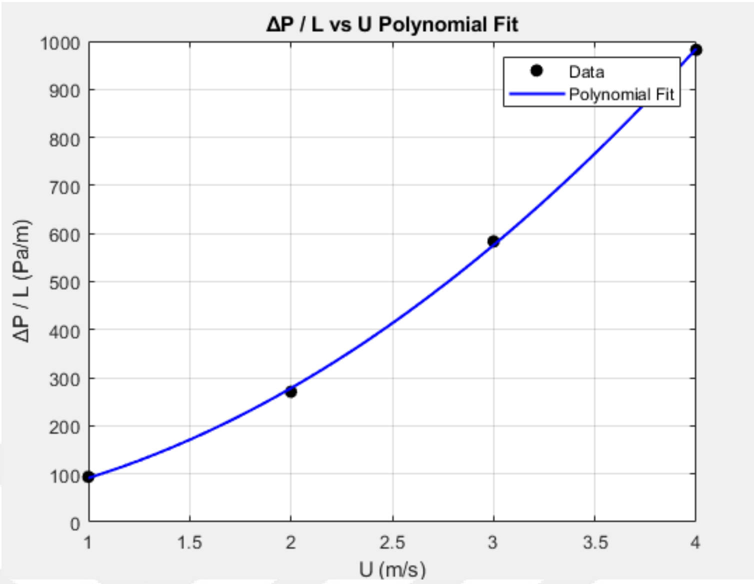

Table 3.4: *Velocity vs Pressure Drop - 3 Bottles*

X-Velocity	Pressure Drop
1 m/s	25.3742
2 m/s	73.11413
3 m/s	157.5830
4 m/s	265.1851

An interesting observation emerges when comparing the single-bottle and three-bottle configurations. The flow domain length increases from 0.08 m to 0.27 m, a factor of 3.375. Meanwhile, the corresponding pressure drop increases from 6.86 Pa to 25.37 Pa, approximately 3.69 times greater. When normalized by the domain length ($\Delta p/\Delta x$), the values remain nearly proportional, indicating that the presence of additional bottles does not introduce a significant deviation in pressure loss per unit length. In other words, despite the added complexity due to wake interactions and local disturbances, the three-bottle setup behaves similarly to the single-bottle case in terms of flow resistance. At higher inlet velocities, however, this proportionality begins to decline. In some cases, the pressure drop increase becomes less than the geometric scaling factor of 3.375.

Figures 3.14 and 3.15 illustrate the x-velocity field in the three-bottle configuration, showing that the airflow is predominantly channeled through the narrow neck regions between bottles, while the surrounding areas exhibit pronounced flow stagnation and separation.

Figure 3.14: *X-Velocity Profile of 3 Bottle Case*

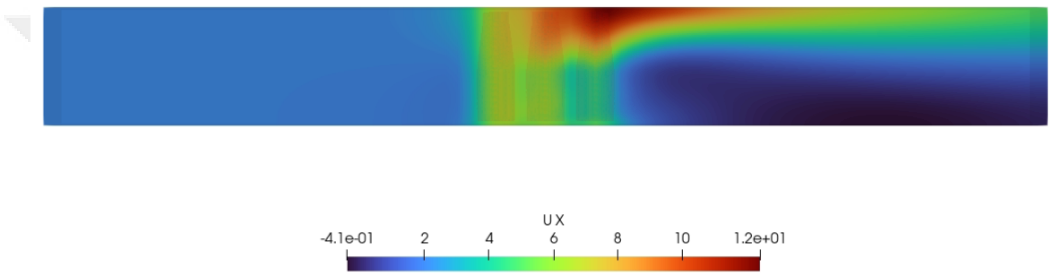

Figure 3.15: *X-Velocity Profile of 3 Bottle Case - Top View*

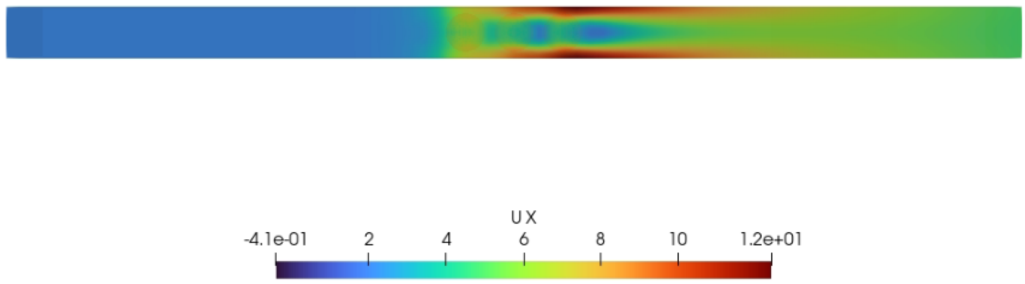

The initial simulation yielded a Forchheimer coefficient of 93.6. However, when this value was applied to the porous zone model in the validation setup (see Figure 3.16), the resulting pressure drop was calculated as 56 Pa, which significantly underestimated the pressure loss observed in the three-bottle configuration (73.1 Pa). To improve the agreement, the Forchheimer coefficient was adjusted further up increasing the Forchheimer coefficient until 123.0, which produced a pressure drop of 73.55 Pa, closely matching the reference value from the three-bottle case. Then the findings are tested with the furnace simulation check 3.22 for the final findings

Figure 3.16: *Porous Zone of 3 Bottle Case - Validation*

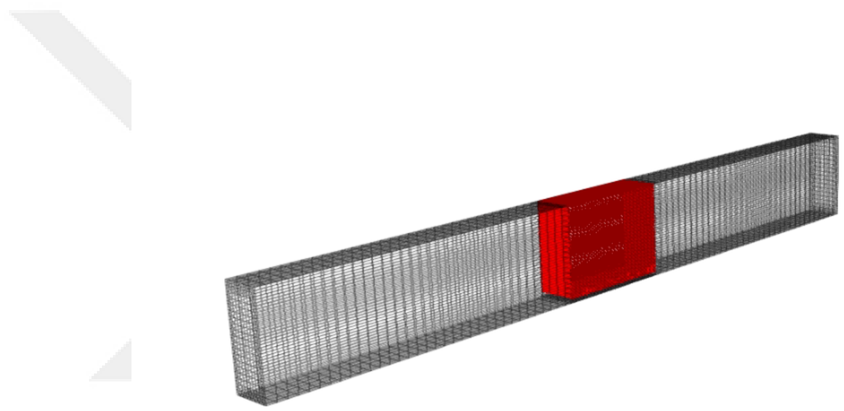

3.4 Solver Coupling Strategy

The simulation framework is based on two main solvers that are coupled through a hybrid iterative structure:

- A **steady-state fluid solver**, which simultaneously resolves radiative and convective heat transfer within the furnace atmosphere,
- A **transient solid conduction solver**, which models the time-dependent temperature field inside the moving glass containers.

Within the steady-state solver, radiative and convective heat transfer are solved using distinct sub-solvers. Radiative heat transfer is computed using a surface-to-surface Monte Carlo method while convective effects are resolved using OpenFOAM's buoyantSimpleFoam solver. The coupling between these two sub-solvers is managed through the exchange of surface heat fluxes and temperatures.

The steady-state solution loop begins with an initial guess for the wall and glass surface temperatures. The Monte Carlo radiative solver then computes directional radiative heat fluxes ($\dot{q}_{\text{rad}}$) for each surface patch using a ray tracing model. These fluxes are mapped onto OpenFOAM boundary patches and used as input in the **externalWallHeatFluxTemperature** and **thermalBaffle1D** boundary conditions.

OpenFOAM subsequently calculates convective heat fluxes ($\dot{q}_{\text{conv}}$) based on local airflow conditions. The total surface heat flux is obtained as the sum:

$$\dot{q}_{\text{net}} = \dot{q}_{\text{rad}} + \dot{q}_{\text{conv}} \tag{3.12}$$

This net flux is used to update surface temperatures within OpenFOAM. To balance computational cost and accuracy, the radiative solver is executed every 10 iterations of the OpenFOAM solver.

Once the steady-state solution converges, the updated wall temperatures are transferred to the transient conduction solver. This solver computes the internal temperature distribution of the glass bottles, accounting for their continuous motion along the furnace. The surface temperatures of the glassware, computed in the transient domain, are then mapped back to the steady-state domain as boundary conditions for the next radiative calculation.

This coupling loop is repeated until convergence is reached in both solvers. The interaction proceeds as follows:

1. The transient solver initializes glass temperatures and passes them to the steady

solver.

2. The steady solver resolves $\dot{q}_{\text{rad}}$ and $\dot{q}_{\text{conv}}$, and updates wall temperatures.

3. These updated wall temperatures are passed back to the transient domain to solve for the next time step.

Figure 3.17: *Solver Coupling Flowchart*

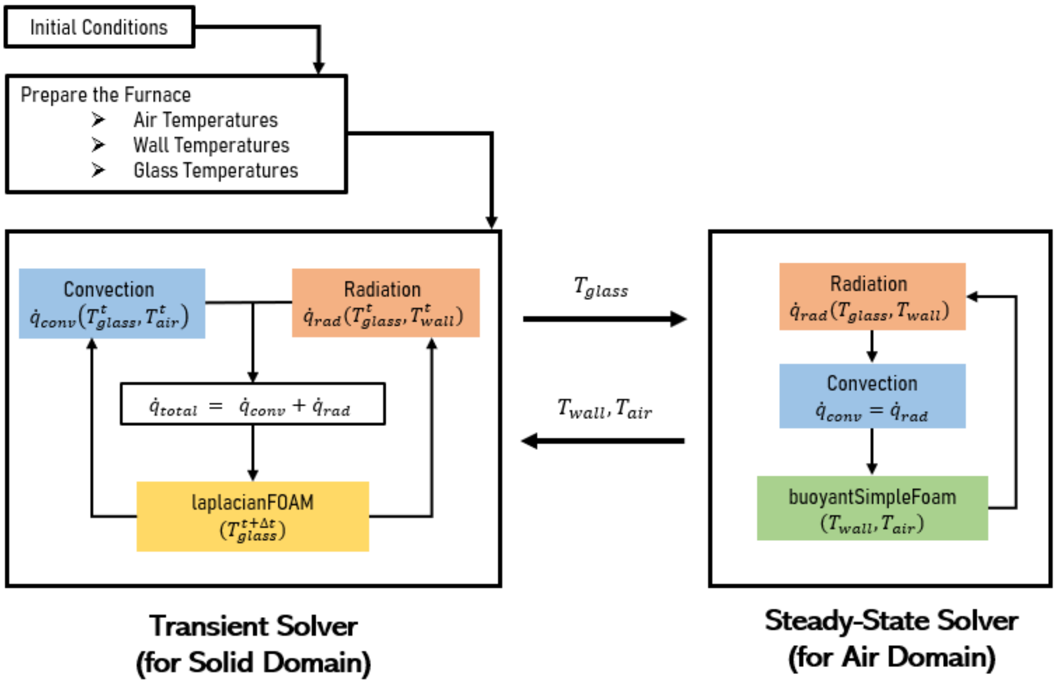

All scalar fields (e.g., temperature, radiative flux) are transferred between solvers using a python code as pipeline where the mapping is done based on writing on files ensuring the mesh exchange between S2S to OpenFOAM using interpolations. Special attention is given to preserve energy consistency across interfaces by ensuring that the energy emitted by one domain is accurately absorbed by the other.

The solver proceeds in a sequential loop: the steady-state solver executes until it reaches thermal convergence (typically within 3000 iterations), after which the transient solver is advanced. This alternating process continues until both solvers converge according to predefined tolerances on temperature change.

Figure 3.18: *Convergence Graph*

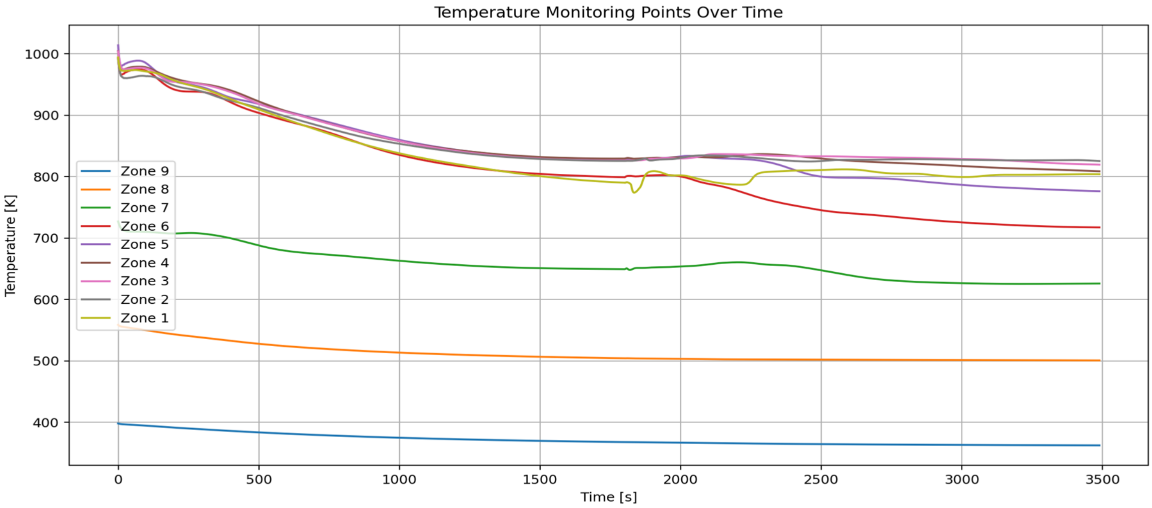

3.5 Interpolation of Data and Mesh

Since the objects in the simulation are represented by different surface discretizations, a direct field transfer between the Surface-to-Surface Radiative Solver and OpenFOAM is not possible without interpolation. Specifically, the radiative solver employs STL-based triangulated surface meshes, while OpenFOAM uses a predominantly structured quadrilateral mesh. In Figure 3.19, a side-by-side comparison of these meshes is presented. For instance, the surface mesh corresponding to the heater walls contains 53,683 triangular elements in the radiative solver, whereas the same region is represented by only 28,666 quadrilateral cells in the OpenFOAM domain.

Figure 3.19: *Interpolated Mesh (Left: Surface-to-Surface Radiative Solver, Right: OpenFOAM)*

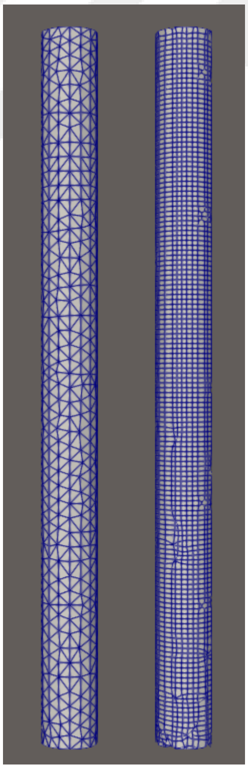

This difference in mesh topology necessitates a point-wise interpolation strategy for exchanging radiative and thermal fields. In this work, the interpolation is carried out using a simple nearest-neighbor approach based on spatial proximity between mesh nodes. Specifically, for every target node in one mesh, the nearest source node from the corresponding dataset is located and its scalar field value is assigned to the target.

This mapping is conducted bidirectionally throughout the coupled simulation loop:

- **From Radiative Solver to OpenFOAM:** The radiative heat flux values computed on the triangulated mesh are interpolated to the OpenFOAM quadrilateral mesh and applied as boundary conditions on the furnace walls.
- **From OpenFOAM to Radiative Solver:** The temperature field obtained from OpenFOAM is mapped back to the original triangulated surface to be used in recalculating radiative exchange factors in the next iteration.

The interpolation is implemented using the *SciPy* Python library's *griddata* function, with the interpolation method set to *nearest*. This approach assigns to each target node the value of its closest source node in Euclidean space:

$$\tilde{q}(\boldsymbol{x}_f) = q\left(\arg\min_{\boldsymbol{x}_c} \|\boldsymbol{x}_f - \boldsymbol{x}_c\|\right) \tag{3.13}$$

where $\boldsymbol{x}_f$ is a target point and $\boldsymbol{x}_c$ is a point from the source mesh. While this approach is computationally efficient, it does not perform any averaging or weighting between multiple neighboring values, which can lead to local discontinuities in mapped fields.

To facilitate this exchange, a file-based communication structure is adopted. The radiative solver and OpenFOAM share data via structured plain-text files:

- *furnaceCenters.txt* – coordinates of the surface nodes in the radiative solver,
- *innerFurnace_*.txt*, *outerFurnace_*.txt* – OpenFOAM patch face centers,

- *Furnace_heatFlux.txt* – surface heat fluxes from the radiative solver,
- *T* – OpenFOAM temperature field exported for radiative feedback.

In contrast, when radiative data are transferred to the quadrilateral OpenFOAM mesh, this granularity is partially lost. As shown in Figure 3.20, which compares temperature distributions on the same geometry, the left panel shows the radiative solver output, while the right panel shows the mapped field in OpenFOAM. The offset between mapped values and the original is a direct consequence of mesh resolution mismatch and the nearest-neighbor interpolation approach. This discrepancy is further influenced by how quadrilateral faces in OpenFOAM do not always align with the triangulated faces used in the radiative solver, potentially causing mapping artifacts.

It is also worth noting that the glassware, modeled as a porous zone in OpenFOAM, is excluded from surface-based radiative interactions. Therefore, no mapping is required for internal volumetric zones in this context.

Figure 3.20: *Temperature Mapping (Left: Surface-to-Surface Radiative Solver, Right: OpenFOAM)*

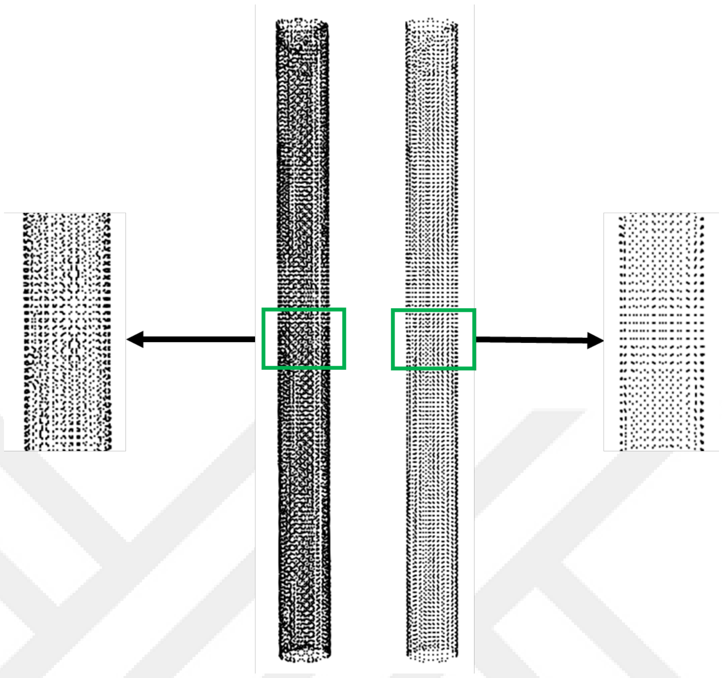

3.6 Validation of the Model

The numerical model was validated using measured temperature data from the actual furnace operation. As part of the convergence process, several temperature monitoring points were placed at the midpoints of each thermal zone. Initially, the furnace was modeled under adiabatic boundary assumptions, where both the inlet and outlet were defined as *no-slip walls*. This configuration was intended to prevent any convective heat and mass loss to the external environment, ensuring that all thermal energy remained confined within the furnace domain.

However, as shown in Figure 3.21, the resulting temperature predictions under this sealed configuration exhibited notable discrepancies. While the model successfully captured the trends in the first two zones, it significantly deviated from the measured values in the downstream heating and cooling zones. A partial recovery is observed only in the final cooling zones, but overall, the mismatch highlighted the inadequacy of fully closed boundary conditions in replicating the real furnace dynamics.

To capture the natural flow dynamics within the furnace and to ensure realistic air supply to the burner regions, a low-pressure condition was imposed at the inlet. This setup allowed the development of a flow from the outlet toward the inlet, driven by internal pressure gradients, and better represented the operational behavior observed in industrial furnaces.

Following this adjustment, the simulation results showed markedly improved agreement with the operational temperature data, as illustrated in Figure 3.22. The updated model reproduced the zone-wise temperature distribution with high fidelity, achieving agreement within approximately ± 5 °C across all zones, both in the heating and cooling regions. This confirms that allowing limited leakage at the inlet is essential for capturing the realistic pressure and thermal behavior of industrial-scale annealing furnaces.

Figure 3.21: *Validation of the Simulation – Without Inlet Leakage*

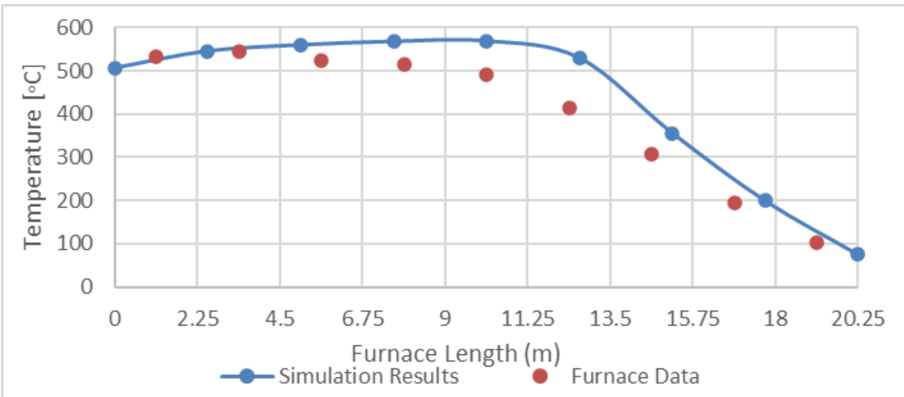

Figure 3.22: *Validation of the Simulation – With Inlet Leakage (Improved Agreement)*

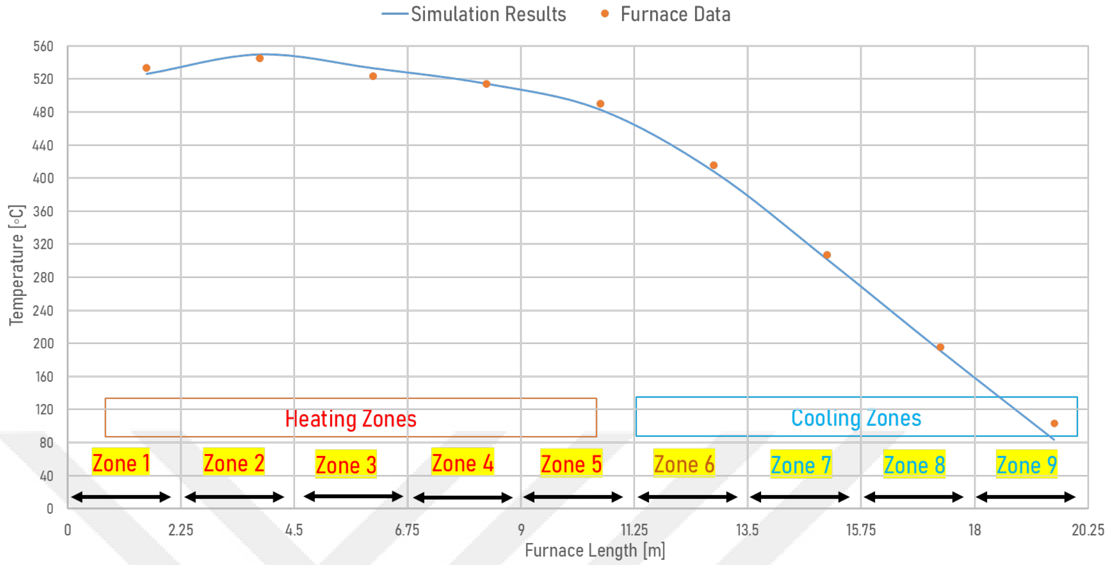

4. RESULTS

In this chapter, the results of the two case studies conducted using the proposed methodology and furnace model are presented and analyzed. The objective of these comparative simulations is to assess how different burner activation strategies influence the thermal performance and energy efficiency of the annealing furnace, as well as their effect on glass product quality.

In *Case 1*, only the burners in the first heating zone are activated, while the remaining zones remain passive. Conversely, *Case 2* represents a configuration where only the burners in the second heating zone are active. These two configurations were selected to examine how shifting the heating load affects the internal temperature distribution, energy consumption, and the thermal uniformity of the glassware.

Glass quality is particularly sensitive to temperature gradients during annealing; therefore, the standard deviation of glass temperatures is used for evaluating product consistency. Lower standard deviation values indicate more uniform heating and, by extension, equal stress relief, better product quality. The simulations aim to determine which configuration yields more favorable conditions for both operational efficiency and product quality.

Detailed information regarding the furnace setup and operational inputs used in the simulations is provided in Table 3.1. The same burner temperatures were utilized for active and passive zones across both cases.

4.1 Comparison of Zone Temperatures Across Case Studies

One of the primary evaluation metrics in the conducted simulations is the evolution of thermal zones along the furnace, particularly focusing on the steady-state temperature distributions achieved under different burner activation strategies. As illustrated in Figure 4.1, a notable difference is observed between the two cases: Case 2 results in significantly higher temperatures in the second heating zone compared to Case 1.

Figure 4.1: *Case 1 vs Case 2 Zone Temperatures*

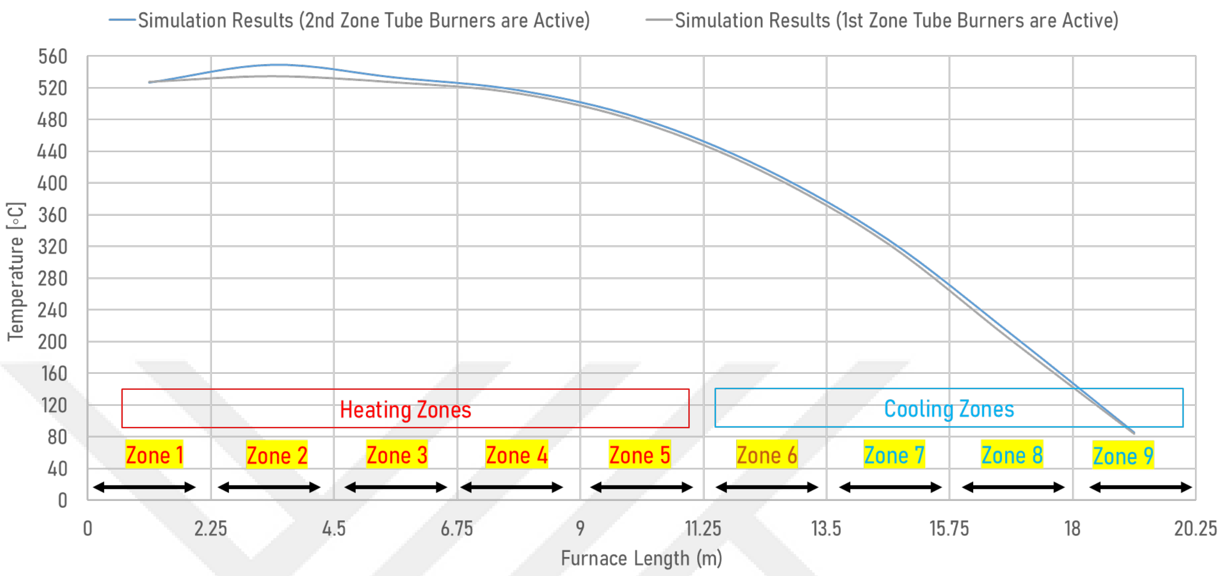

In Case 1, where only the first heating zone is active, the hot air generated by the burners does not propagate effectively to the downstream regions of the furnace. Figure 4.2 illustrates the temperature distribution across and within the furnace zones for Case 2. In this configuration, hot air flows centrally in a conical pattern, while colder air predominantly recirculates along the furnace walls. This behavior results from the natural flow tendency toward the furnace inlet. Since the burners are located adjacent to the side walls, this natural circulation facilitates the transport of oxygen-rich air toward the burners, which enter the system from the final two cooling zones, thereby supporting effective combustion near the wall regions.

Figure 4.2: *Zone Temperature Distribution for Case 2*

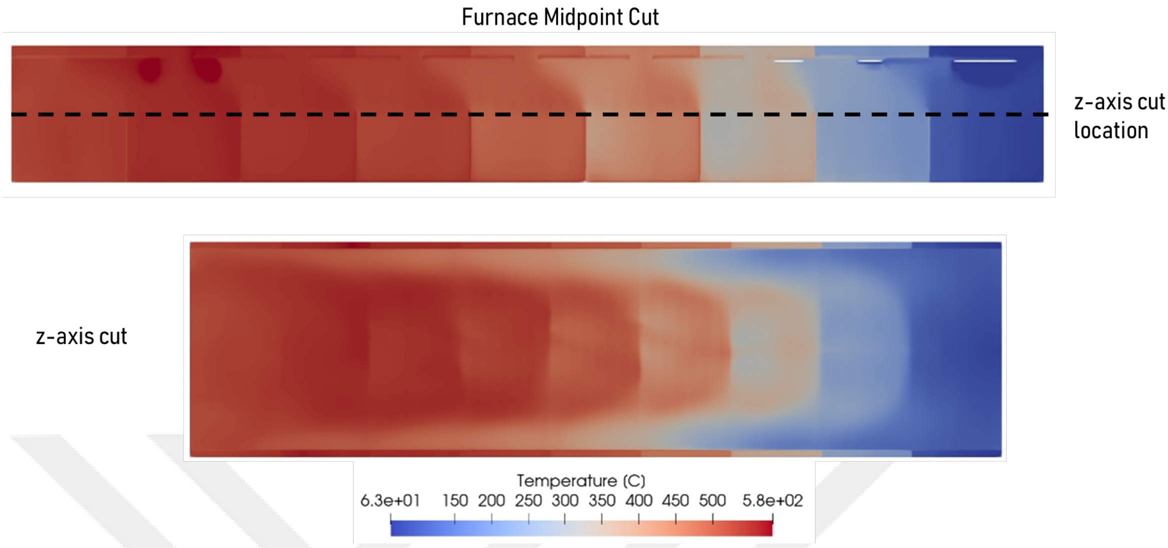

It is also important to note that the temperature differences between the two cases remain relatively minor across most furnace zones. Apart from Zone 2, where Case 2 exhibits a clearly elevated thermal profile, deviations in other zones are generally within a margin of 3–4 °C.

4.2 Comparison of Glass Temperature Profiles

In addition to the furnace zone temperatures, the temperatures of glassware, modeled through multiple rows of moving glass containers,serves as a key indicator to evaluate the effectiveness of each heating configuration. In this context, the average temperatures of each row, computed over the entire annealing trajectory, are compared between the two case studies.

Figure 4.3: *Case 1 vs Case 2 Average Bottle Temperatures*

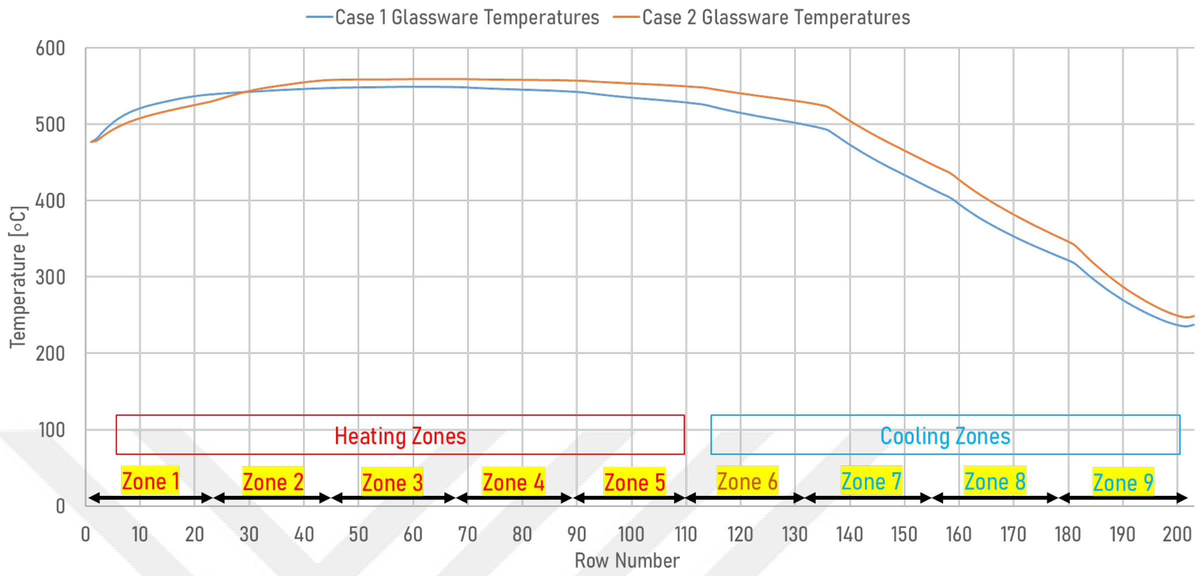

As shown in Figure 4.3, Case 1 exhibits a rapid initial rise in glass temperature, particularly in the early stages of the furnace. This is expected due to the direct exposure of the glass containers to the active burners located at the very beginning of the furnace Zone 1. Consequently, the glass in Case 1 initially surpasses Case 2 in terms of temperature within the first zone, 533 °C to 522 °C for Case 1 and Case 2 respectively.

However, as the containers progress along the furnace, this trend reverses. From the midpoint of Zone 2 onwards, the glass rows in Case 2 begin to achieve higher peak temperatures, 559 °C to 548 °C ,Case 2 and Case 1 respectively. This is a direct result of the enhanced heating in Zone 2 of Case 2, which creates more sustained and spatially extended thermal exposure downstream.

Furthermore, not only do the peak temperatures in Case 2 exceed those in Case 1, but the average temperatures of the glass containers also remain consistently higher throughout the remainder of the furnace length.

These findings suggest that although Case 1 initiates heating more aggressively the temperatures doesn't increase further than initially. In contrast, Case 2 maintains better thermal coverage and achieves superior glass heating performance over the entire furnace path, resulting in potentially improved product quality and more uniform stress relaxation across the glass articles.

To further evaluate the thermal uniformity within the furnace, standard deviation values of the glass temperatures were calculated for each furnace zone.

Figure 4.4: *Case 1 vs Case 2 Standard Deviations (Each Small Square Represents a Glass Object)*

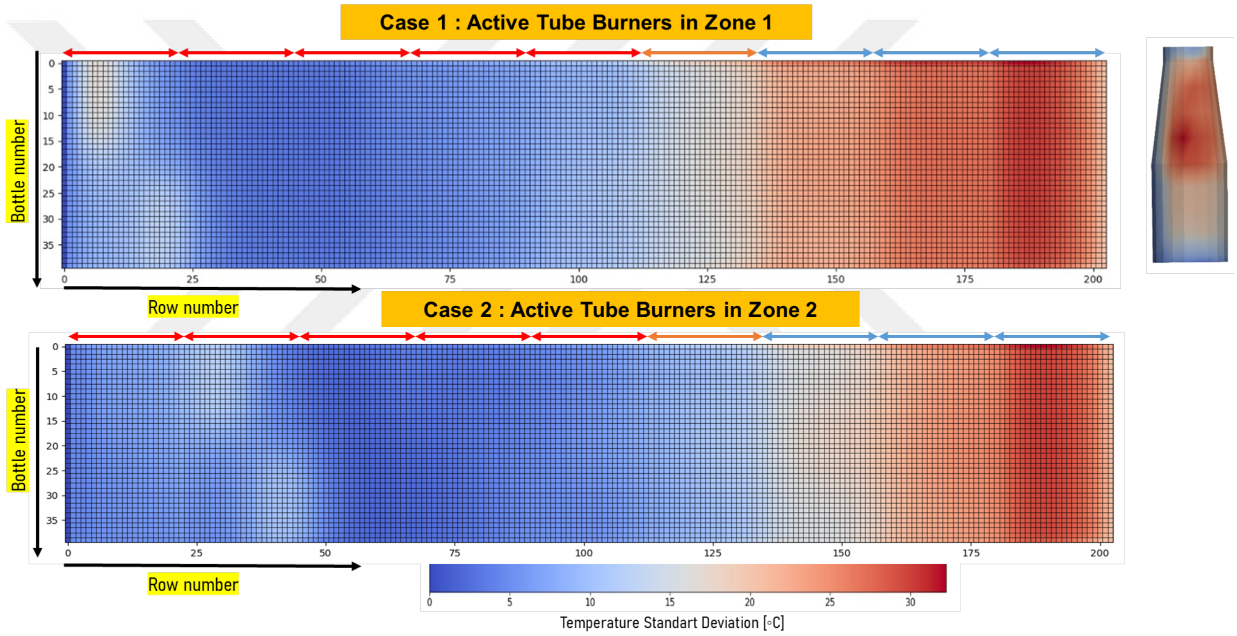

Figure 4.4 indicate that Case 2 exhibits lower standard deviations across the first five zones, which correspond to the heating region of the furnace. This suggests that Case 2 achieves more spatially uniform heating in the upstream segment, However, when considering the furnace as a whole, the advantage of Case 2 in terms of thermal uniformity becomes less definitive. A shift in the cooling region is observed, glass cooling in Case 1 begins slightly earlier compared to Case 2, as indicated by the standard deviations 4.4 and average glass temperatures 4.3. This shift suggests that the thermal response in Case 2 is

slightly delayed, with the temperature field appearing offset by approximately one zone.

Despite this, both configurations ultimately bring the glass containers to the desired exit temperature at approximately the same point along the furnace. Therefore, while Case 2 achieves better uniformity in the heating zones, its overall superiority in terms of product quality cannot be definitively claimed without further on-hand investigation.

4.3 Radiative Flux Distribution – Case Comparison

In this section, the radiative heat flux distributions obtained for Case 1 and Case 2 are compared to assess the impact of passive and active burner configurations on the heat transfer behavior. As stated before, all burners in both configurations are modeled with prescribed wall temperatures: *active burners* are maintained at 950 K (677 °C) and *passive burners* at 800 K (527 °C). This fixed-temperature approach implies that the burner surfaces do not respond thermally (i.e., do not heat up dynamically), but rather act as constant-temperature radiation sources.

In the resulting radiative flux maps, *negative values* correspond to net radiation emitted from the burner surfaces into the furnace (i.e., net heat loss), while *positive values* indicate incoming radiative heat absorbed by the burner walls. Although positive net heat flux might seem counterintuitive, it arises due to the fixed-temperature boundary condition; in reality, such heat input would imply that the burner walls would heat up due to incoming radiation from other surfaces.

A detailed visual comparison of top and bottom views of the furnace in Figure 4.5 reveals several insights. In Case 1, as the active Burners are on Zone 1 the majority of the net radiation originates from there, The heaters radiate strongly especially toward the colder incoming glass containers. In contrast, the downstream zones exhibit minimal radiative activity. The top sides of the burners show significantly reduced net radiative transfer compared to their bottom sides. This is attributed to thermal stratification in the furnace: during the operation, the *inner furnace walls above the burners absorb radiation and heat up*, reducing the temperature gradient and hence the net heat flux. Meanwhile, the colder glass containers incoming below the burners maintain a strong temperature contrast, leading to more pronounced radiative heat exchange in the lower sections.

Figure 4.5: *Comparison of radiative flux distribution for Case 1 (top row) and Case 2 (bottom row).*

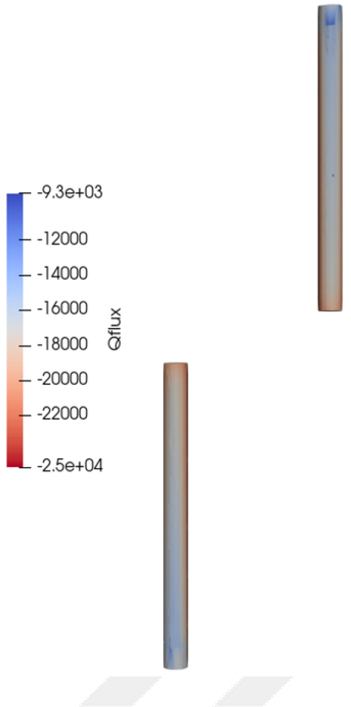

(a) Case 1 – Top View

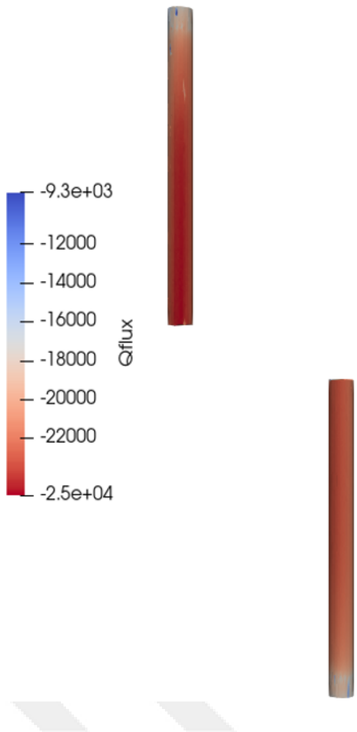

(b) Case 1 – Bottom View

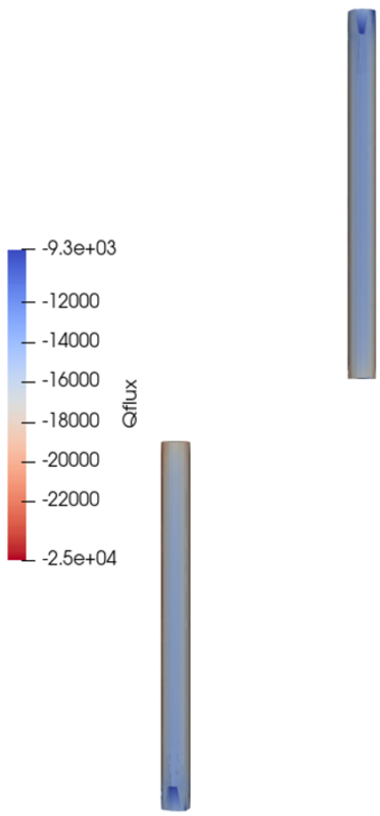

(c) Case 2 – Top View

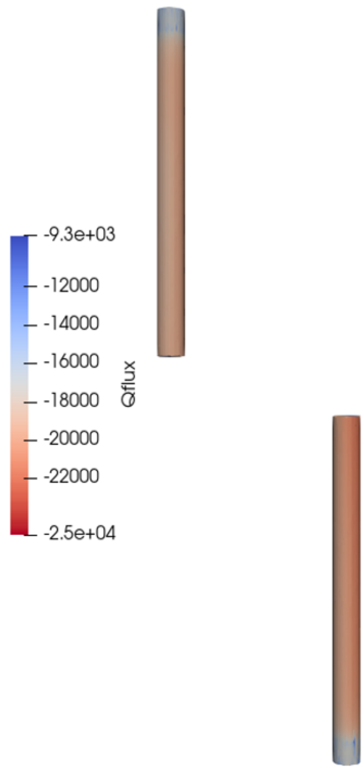

(d) Case 2 – Bottom View

When comparing all of the heaters active or passive, as shown in Figure 4.6, we observe that the radiative heat transfer from the burners is lower in Case 2, even though those burners are active in both configurations. This can be explained by the presence of residual heat in the upstream zone of Case 2. In realistic furnace operation, zone control mechanisms deactivate burners once the target temperature is achieved, but the heater walls remain hot and continue to radiate. Consequently, in Case 2, the glass containers entering the second zone are already partially preheated, reducing the demand for additional heat from the active burners. This results in lower radiative output in the second zone. The first zone, being passive in Case 2, still emits radiation due to having higher temperature than the glassware entering at 477 °C.

Figure 4.6: *Comparison of radiative flux distribution for all burner surfaces in Case 1 (top) and Case 2 (bottom).*

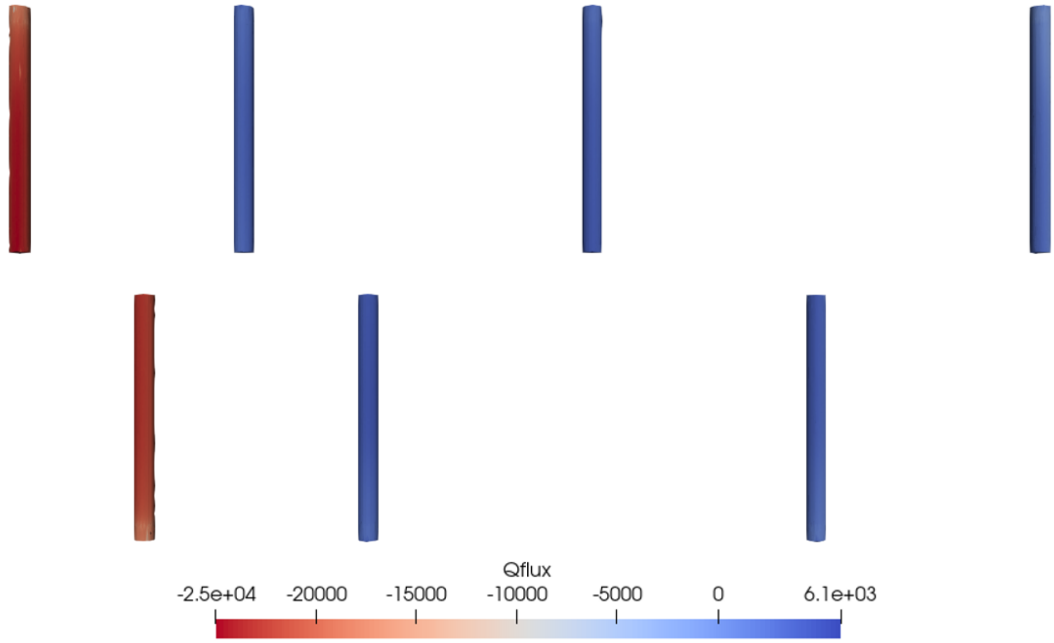

(a) Case 1 – All Heaters

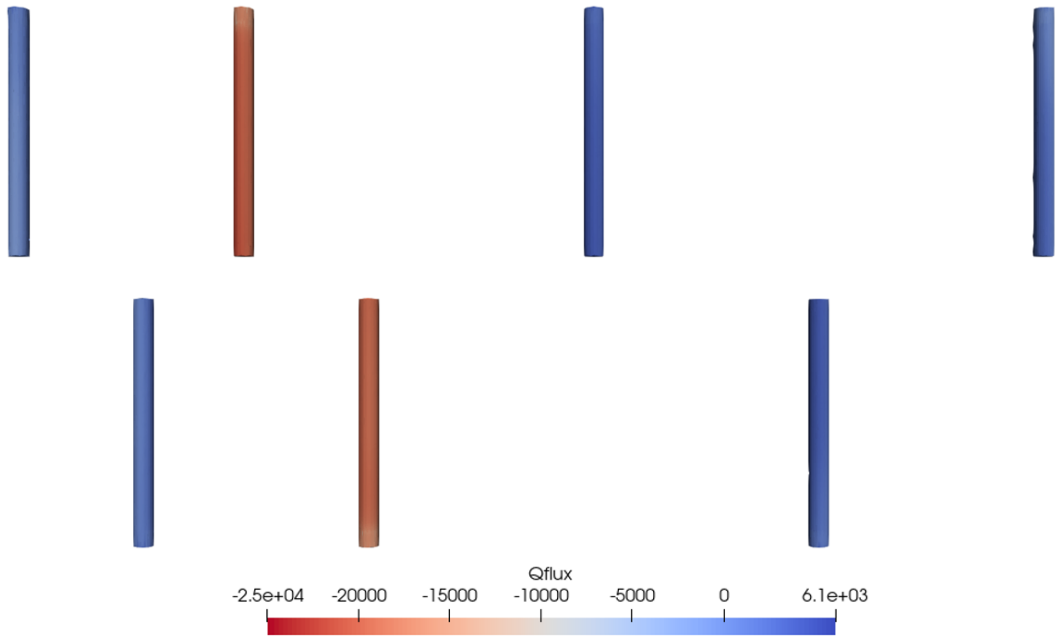

(b) Case 2 – All Heaters

When looking at the overall radiative performance, the trend is further clarified by the zone-wise net radiative heat transfer values presented in Table 4.1. In Case 1, the total net radiative output reaches approximately 61 kW for the active heaters, whereas in Case 2 this value is reduced to about 55 kW. This indicates that less energy is required in Case 2 to sustain the same thermal conditions, making it a more energy-efficient configuration. In practice, this implies that the burner system would consume less fuel, offering improved energy economy of the industrial operation.

Table 4.1: *Net Radiative Heat Transfer of the Heaters in Each Heating Zone.*

Zone	**Case 1**	**Case 2**
1	-61424.18 W	610.79 W
2	7357.15 W	-55031.05 W
3	3488.98 W	3817.73 W
4	2367.71 W	2568.63 W
5	-540.66 W	-475.56 W

5. CONCLUSION

This study developed a detailed thermal model of a continuous container glass annealing furnace and investigated the role of burner configuration, porous resistance, and flow–radiation interactions in determining process outcomes. The framework resolved the coupled effects of conduction, convection, and radiation and enabled quantitative assessment of both product quality and energy efficiency.

The bottle zone was represented as a porous medium using the Darcy–Forchheimer formalism. To calibrate the resistance coefficients, simulations were performed at different inlet velocities and bottle packing densities. A notable outcome was that the porous zone thickness Δx increased by a factor of 3.375 between the one-bottle and three-bottle cases (from 0.08 m to 0.27 m), while the corresponding pressure drop Δp increased by a factor of 3.69 at 2 m/s. The resulting 9% difference between the scaling of Δp and Δx indicates that additional glass packing introduces a measurable but limited increase in flow resistance, with moderate overall impact on slow-speed flows.

Two burner activation strategies were compared, revealing distinct thermal and energy outcomes. Although the average air temperatures across most zones were similar in both cases (within ±5 °C), significant differences emerged in glass temperature profiles. In the downstream heating configuration (Case 2), peak glass temperatures were approximately 14 °C higher than in the upstream case (Case 1) and have better temperature uniformity at the heating zones, despite comparable furnace setpoints. Moreover, Case 2 produced a delayed onset of cooling, shifting the cooling process one zone downstream. Energy analysis further showed that Case 2 operated with a lower net radiative output—approximately 55 kW compared to 61 kW in Case 1—corresponding to a 9% less energy input required to maintain the furnace temperatures. Another factor to mention is that the burners in Case 2 are located farther from the inlet where leakage persists. As a result, air temperatures in the first zone remain lower, reducing heat loss through the inlet and thereby enhancing overall efficiency.

Overall, the study demonstrates that burner configuration is a critical control parameter, with direct influence on glass temperature uniformity, thermal stress development, and furnace energy consumption. Radiative transfer emerged as the dominant mechanism governing glass temperature evolution, reinforcing that air temperature alone is an insufficient control metric for annealing quality.

Beyond providing industrial insights, the developed methodology also offers a foundation for data-driven furnace optimization. With the continued advancement of artificial intelligence, high-fidelity simulation data generated by such models can be used to train predictive algorithms. In the future, Artificial Intelligence (AI) based surrogates could deliver real-time predictions of furnace behavior, enabling burner configurations and operational strategies to be optimized without running full CFD simulations. This would transform the developed framework from a design tool into a basis for next-generation, AI-driven control mechanisms in glass annealing and other thermal processing systems.

REFERENCES

[1] H. Wiki, "Know temperature when metal glows red," Apr. 2013. [Accessed July-2025].

[2] A. Q. Tool, "Relaxation of stresses in annealing glass," *J. Res. Natl. Bur. Stand.(US)*, vol. 34, no. 2, pp. 199–211, 1945.

[3] L. H. Adams and E. D. Williamson, "The annealing of glass," *Journal of the Franklin Institute*, vol. 190, no. 5, pp. 597–631, 1920.

[4] R. Gardon and O. Narayanaswamy, "Stress and volume relaxation in annealing flat glass," *Journal of the American Ceramic Society*, vol. 53, no. 7, pp. 380–385, 1970.

[5] J. E. Shelby, *Introduction to Glass Science and Technology*. Royal Society of Chemistry, 2nd ed., 2005.

[6] A. K. Varshneya and J. C. Mauro, *Fundamentals of Inorganic Glasses*. Elsevier, 3rd ed., 2019.

[7] F. V. Tooley, *The Handbook of Glass Manufacture*. Books for Industry, 3rd ed., 1984.

[8] W. Trier, *Glass Furnaces: Design, Construction and Operation*. The Society of Glass Technology, 1987.

[9] J. R. Howell, M. P. Mengüç, R. Siegel, and K. Daun, *Thermal Radiation Heat Transfer*. CRC Press, 7th ed., 2020.

[10] M. F. Modest, *Radiative Heat Transfer*. Academic Press, 4th ed., 2021.

[11] M. Rantala, "Thermal radiation and forced convection in flat glass tempering furnaces," *Glass Processing Days 2019 Conference*, 2020. Presented at GPD 2019; published 2020.

[12] E. Yıldız, A. M. Başol, and M. P. Mengüç, "Segregated modeling of continuous heat treatment furnaces," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 249, p. 106993, 2020.

[13] G. C. Altun and A. M. Başol, "Numerical analysis of the radiant heating effectiveness of a continuous glass annealing furnace," *Applied Thermal Engineering*, vol. 204, p. 117943, 2022.

[14] P. Dillon, "A dual-solver cfd model for conjugate heat transfer in continuous thermal processing," *Case Studies in Thermal Engineering*, vol. 49, p. 103337, 2023.

[15] N. Hajaliakbari and S. Hassanpour, “Analysis of thermal energy performance in continuous annealing furnace,” *Applied Energy*, vol. 206, pp. 829–842, 2017.

[16] J. R. Howell, “The monte carlo method in radiative heat transfer,” 1998.

[17] N. Depree, J. Sneyd, S. Taylor, M. P. Taylor, J. J. Chen, S. Wang, and M. O’Connor, “Development and validation of models for annealing furnace control from heat transfer fundamentals,” *Computers & Chemical Engineering*, vol. 34, no. 11, pp. 1849–1853, 2010.

[18] F. Jurić, M. Vujanović, M. Živić, M. Holik, X. Wang, and N. Duić, “Assessment of radiative heat transfer impact on a temperature distribution inside a real industrial swirled furnace,” *Thermal science*, vol. 26, no. 6, pp. 3663–3672, 2020.

[19] F. He, J. Shi, L. Zhou, W. Li, and X. Li, “Monte carlo calculation of view factors between some complex surfaces: Rectangular plane and parallel cylinder, rectangular plane and torus, especially cold-rolled strip and w-shaped radiant tube in continuous annealing furnace,” *International Journal of Thermal Sciences*, vol. 134, pp. 465–474, 2018.

[20] M. Filipponi, F. Rossi, A. Presciutti, S. De Ciantis, B. Castellani, and A. Carpinelli, “Thermal analysis of an industrial furnace,” *Energies*, vol. 9, no. 10, 2016.

[21] W. Fei *et al.*, “Modeling of strip heating process in vertical continuous annealing furnace,” *Journal of iron and steel research, International*, vol. 19, no. 5, pp. 29–36, 2012.

[22] S. Strommer, M. Niederer, A. Steinboeck, and A. Kugi, “A mathematical model of a direct-fired continuous strip annealing furnace,” *International Journal of Heat and Mass Transfer*, vol. 69, pp. 375–389, 2014.

[23] S. Ergun and A. A. Orning, “Fluid flow through randomly packed columns and fluidized beds,” *Industrial & Engineering Chemistry*, vol. 41, no. 6, pp. 1179–1184, 1949.

[24] D. A. Nield and A. Bejan, *Convection in porous media*. Springer, 2006.

[25] C. Rivas Evans, *A thermal model for a float glass annealing lehr*. PhD thesis, Massachusetts Institute of Technology, 1982.

[26] S. Reiter, “Stl reader code,” Copyright (c) 2018-2023. Accessed July-2025.

[27] OpenFOAM, “Darcy forchheimer coefficients and explisit,” Aug. 2011. [Accessed July-2025].